Neutrino oscillations and earth tomography with KM3NeT-ORCA

Simon Bourret

▶ **To cite this version:**

Simon Bourret. Neutrino oscillations and earth tomography with KM3NeT-ORCA. Physics [physics]. Université Sorbonne Paris Cité, 2018. English. NNT : 2018USPCC247 . tel-02491394

HAL Id: tel-02491394
https://theses.hal.science/tel-02491394

Submitted on 26 Feb 2020

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Thèse de doctorat de l'Université Sorbonne Paris Cité
préparée à l'Université Paris Diderot

École doctorale STEP'UP – ED N°560
Laboratoire AstroParticule et Cosmologie

Neutrino oscillations and Earth tomography with KM3NeT-ORCA

Simon Bourret

sous la direction de

Véronique Van Elewyck et Édouard Kaminski

présentée et soutenue publiquement le 30 novembre 2018
devant le jury composé de

Pr. Alessandra TONAZZO	Présidente du jury
Dr. Véronique VAN ELEWYCK	Directrice de thèse
Pr. Édouard KAMINSKI	Co-directeur de thèse
Dr. Fabrice PIQUEMAL	Rapporteur
Dr. Walter WINTER	Rapporteur
Dr. Justin EVANS	Examinateur
Dr. Aart HEIJBOER	Examinateur

Contents

List of figures

List of tables

Introduction

Neutrinos have played a special part in the development of particle physics throughout the past century. Beyond the constituents of the atom, the neutrino was the first elementary particle to be introduced in the theoretical framework. Its existence was postulated in 1931 to reconcile the principle of energy conservation with the observed spectrum of β-decays. Predicted as a possibly undetectable 'ghost' particle, its experimental detection turned out to be possible but would not be achieved until 1956. Because they are neutral both electrically and in terms of the colour charge of quarks, neutrinos only interact with other particles via the weak force and thus travel essentially unimpeded through matter. Three flavours of neutrinos have been discovered: electron, muon, and tau neutrinos (ν_e, ν_μ, ν_τ), fitting in the three-generation landscape of the Standard Model of particle physics.

Neutrinos had been understood to be massless particles until 1998, when it was discovered that they are able to change from one flavour to another when propagating over macroscopic distances – a quantum mechanical phenomenon called *neutrino oscillations*. This was established unambiguously by two experiments (SNO and Super-Kamiokande) observing neutrinos produced respectively by nuclear reactions inside the Sun (solar neutrinos) and by the impact of cosmic rays in Earth's atmosphere (atmospheric neutrinos). The importance of this experimental discovery was acknowledged by the Nobel Prize in Physics 2015, awarded jointly to A. McDonald and T. Kajita. The existence of neutrino flavour oscillations imply that neutrinos have mass, and that the masses of the three neutrino states (distinct from the flavour states) must be different. To this day, neutrino oscillations are the only established experimental fact which indicates the need for physics beyond the Standard Model.

Most of the parameters governing neutrino oscillations have been measured over the past two decades by dedicated experiments using a variety of neutrino sources and set-ups. One of the remaining unknowns concerns the *neutrino mass hierarchy* (NMH), related to the ordering of neutrino masses. Due to its many theoretical and experimental implications, the determination of the NMH is of utmost importance for the field. The most promising avenue for its measurement makes use of the enhancement of the oscillations of neutrinos when propagating in matter. This resonance results from the coherent forward scattering of neutrinos on electrons (an analogue of the refraction of light in a transparent medium). One possible experimental approach for the NMH determination then consists in the observation of atmospheric neutrinos traversing the Earth, using very large volume detectors: *neutrino telescopes*, which rely on the Cherenkov effect to detect particles. To pursue this goal, the KM3NeT collaboration is building a megaton-scale Cherenkov detector in the depths of the Mediterranean Sea: ORCA (*Oscillation Research with Cosmics in the Abyss*). ORCA is expected to determine the NMH with a significance better than 3σ after a few years of operation, and to provide improved measurements of the atmospheric oscillation parameters. The relatively short timescale considered for ORCA contruction and subsequent measurement of the NMH makes it one of the most competitive experiments in the field.

The groundbreaking measurement of atmospheric neutrino oscillations by Super-Kamiokande was only sensitive to a first order phenomenon, well described by a two-flavour approxima-

tion ($\nu_\mu \to \nu_\tau$ transitions) and fairly independent of Earth matter effects. ORCA and other next-generation atmospheric detectors will investigate subleading, three-flavour effects such as matter-enhanced $\nu_\mu \leftrightarrow \nu_e$ transitions. Beyond the determination of neutrino parameters, the dependence of these effects on the electron density in the traversed medium might ultimately offer an unprecedented window on the interior of our planet: *neutrino oscillation tomography*, with the potential to bring new information on the chemical composition of its innermost layers. However, this measurement will be experimentally challenging, as the impact of small electron density variations is a smaller effect than that of the mass hierarchy.

The work presented in this thesis, carried on within the KM3NeT collaboration, is aimed at estimating and optimising ORCA's sensitivity to the NMH determination, the measurement of atmospheric oscillation parameters, and the study of the electron density in the deep Earth. The focus is given to detector response modeling, statistical methods, and the impact of systematic uncertainties. A new analysis method has been developed, based on a 'full Monte Carlo' modeling of the detector response. The approach is implemented in a general purpose analysis software framework, designed for the production of reference sensitivity results and the systematic exploration of new and improved strategies optimising the performance of the oscillation analyses. Based on an updated Monte Carlo (MC) simulation of the detector, I have produced the latest sensitivity estimations for ORCA, that I presented at the Neutrino 2018 Conference. Earth tomography studies with preliminary versions of the analysis were also presented in Neutrino 2016 and ICRC 2017. The framework has further been used in other physics studies within the collaboration, *e.g.* focusing on light sterile neutrino and non-standard interactions.

The scientific and technical background is presented in the first part of this thesis. Neutrinos and their flavour oscillations are introduced in Chap. 1. After a historical introduction and an overview of neutrino sources and detection channels, the discovery of neutrino oscillations is discussed. Neutrino masses and mixing, neutrino flavour oscillations in vacuum and in matter are then introduced from a theoretical point of view, before summarising the current state of experimental neutrino oscillation physics. Chap. 2 focuses on the oscillations of atmospheric neutrinos crossing the Earth and their experimental study with neutrino telescopes at the GeV scale. The NMH determination and the prospects for Earth tomography are considered in some depth. Chap. 3 then presents the ORCA experiment: detector design, data processing, simulation tools, and strategies for the reconstruction and classification of neutrino events.

The second part of this thesis presents the general methodology and analysis tools that I have developed for ORCA sensitivity studies. Chap. 4 details the calculation of expected rates of neutrino interactions, which are at the basis of the analysis. The models of flux, cross-section and Earth density profile as well as the oscillation calculator methodology are described in some detail. The approach to the modeling of the detector response is discussed in Chap. 5. The method is based on a fully correlated response matrix built from MC events, and the correction of statistical artifacts due to the finite size of the MC sample is adressed in detail. Chap. 6 then presents the general statistical approach to the mass hierarchy determination and discusses two methods employed in practice for the sensitivity study. Systematic uncertainties are adressed in Chap. 7, focusing on their modeling and implementation in the framework.

Results of the sensitivity studies are presented in the last part. Chap. 8 focuses on the neutrino mass hierarchy. The impact of detector resolutions, finite MC effects, oscillation parameter degeneracies and other sources of systematic uncertainties is detailed, and some preliminary investigations of analysis strategy optimisations are reported. The official sensitivity results obtained with a full frequentist statistical approach are presented. Finally, Chap. 9 investigates the potential of ORCA for the measurement of atmospheric oscillation parameters and Earth tomography studies.

Part I

Scientific and experimental context

Chapter 1

Overview of neutrinos and their flavour oscillations

Contents

This chapter first relates the birth of the idea of the neutrino, whose existence was postulated in the early 1930's. An overview is then given of the great variety of natural and man-made neutrino sources, along with some of their detection techniques. The existence of neutrino flavour oscillations, established at the turn of the 20th century, implies that neutrinos have mass. The experimental discovery of neutrino oscillations is discussed, before introducing the formalism of this quantum mechanical phenomenon. Some aspects of the description of neutrinos in modern particle physics theory are then presented, with an emphasis on neutrino masses and mixing as well as three-flavour oscillations in vacuum and in matter. Finally, the current state and upcoming challenges of experimental neutrino physics are described, focusing on oscillation research and in particular the determination of the neutrino mass hierarchy.

1.1 Introduction to neutrinos and historical overview

This section gives a general introduction to the neutrino and to neutrino oscillations. The focus is given to historical developments and simple experimental aspects, with the minimal amount of theoretical description necessary to understand the context of the thesis. Appendix A.1 and the glossary provide complements intended for readers with a limited background in the field, and are referenced when needed. Some theoretical elements are further developed in Appendix A.2. The main sources for this section and Appendix A are Refs. [1–4].

1.1.1 Early history of the neutrino

The history of the neutrino has long been interlinked with that of radioactivity and nuclear physics. The birth of nuclear physics dates back to the end of the 19th century, when Henri Becquerel discovered that uranium naturally emits an invisible radiation able to blacken a photographic plate through an opaque layer of paper. Further research by Ernest Rutherford, Marie and Pierre Curie and others showed that several elements other than uranium were *radioactive* (a term later coined by Marie Curie), and that the radioactive emissions could be classified, based on their different ability to penetrate dense materials, into three types of radiations designated as α, β and γ rays. It was later found that α and β rays were in fact charged particles: respectively helium nuclei (α particles, of positive charge) and highly accelerated electrons (β particles, of negative charge[1]).

In 1914, James Chadwick showed that the values of the energy of β particles emitted by a single type of radioactive nucleus are measured with a *continuous* distribution. This seemed to be fundamentally different from α and γ rays which, for a given nucleus and decay type, are always measured with approximately *the same* energy (or at least with a *discrete* set of values of energy). Fig. 1.1 shows such a continuous β spectrum, as measured in 1927 by Ellis and Wooster [5].

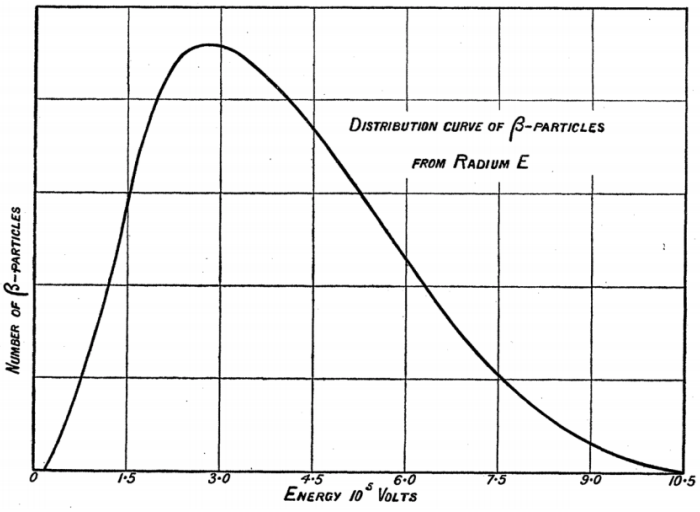

Figure 1.1: Energy distribution of the β particles (electrons) emitted by radioactive decays of *Radium E* (^{210}Bi) measured by Ellis and Wooster in 1927. This version of the figure was taken from [6] and the original paper is accessible at Ref [5].

The experiment of Ellis and Wooster included a calorimetric measurement of the total emitted energy, which showed that the continuous distribution of energy could not be explained by secondary effects, *e.g.* energy loss of the emitted electrons in the medium before their detection.

[1]The original β-decay, as categorised by Rutherford, is now sometimes referred to as β^-, in contrast with β^+ decays undergone by some nuclei – emitting a positron (positively charged electron) instead of an electron.

It was later demonstrated that this missing energy could not be associated with undetected neutral γ-rays either.

This result was extremely surprising, because it seemed to contradict one of the most important, foundational laws of physics: conservation of energy. Indeed, considering the β-decay of a nucleus:

$$ {}^{\mathrm{A}}_{\mathrm{Z}}\mathrm{X} \rightarrow {}^{\mathrm{A}}_{\mathrm{Z+1}}\mathrm{Y} + e^{-}, $$

the conservation of energy implies that the kinetic energy of the emitted electron should be equal to the difference of mass energy between the parent nucleus and the disintegration products:

$$ E_c(e^-) = \Big(M_X - (M_Y + m_{e^-}) \Big) \times c^2 $$

(noting that the recoil kinetic energy of the nucleus is negligible because its mass is much greater than the electron's). Therefore, the electrons emitted in β-decays should be monoenergetic.

Various ideas were put forward to remedy to this situation. A notorious one was Niels Bohr's radical proposal to abandon the principle of energy conservation for individual β-decays (energy would then have been conserved only in a statistical sense). This new principle however seemed to be in contradiction with the observation that the β spectrum had an upper bound in energy.

Wolfgang Pauli formulated the idea of the neutrino in a letter now become famous, which he addressed to the participants of a nuclear physics conference in December 1930. His proposal was that there exists a yet unknown neutral elementary particle which is emitted in the β-decay together with the electron. This particle must be assumed to escape the experimental apparatus undetected. The β-decay then becomes a three-body process:

$$ {}^{\mathrm{A}}_{\mathrm{Z}}\mathrm{X} \rightarrow {}^{\mathrm{A}}_{\mathrm{Z+1}}\mathrm{Y} + e^{-} + \nu, \tag{1.1} $$

where ν denotes the new particle – which Pauli called the "neutron" (the particle known today as neutron would only be discovered a few years later). In such a three-body process, the kinetic energy of the decay can be shared between the two outgoing particles e^- and ν. This explains that the observed energy spectrum for the electron is continuous (as the fraction of energy carried by the electron can vary) and has a finite endpoint corresponding to the electron carrying away essentially all of the decay energy.

Pauli's "neutron", which he assumed to be a fermion of spin 1/2, and a *constituent* of the nucleus (in his inital 1930 proposal), not only solved the continuous β spectrum problem, but also other paradoxical observations concerning the spin and statistics of certain nuclei[2]. Nevertheless the proposal of a new, undetected, and possibly undetectable particle was seen as a "desperate remedy" both by Pauli and the rest of the nuclear physics community. For this reason, and although he presented the idea publicly and discussed it with *e.g.* Fermi and Bohr at several occasions in-between 1930 and 1933, there was no publication of Pauli's idea until the transcription of a discussion in the proceedings of the Solvay conference of 1933 (Ref. [7], in French).

Following the discovery of the "heavy neutron" (the particle known today as the neutron) by Chadwick in 1932, Enrico Fermi renamed Pauli's particle as the *neutrino* – litterally "little neutron" in Italian. Shortly after the 1933 Solvay conference – where the hypothesis that neutrinos could be massless was acknowledged by Pauli and others – Fermi formulated a theory of β-decay, which was built in analogy with the newly created quantum electrodynamics. Fermi's theory explains the β-decay by a direct four-component interaction, where a pair formed by an electron and a neutrino is *produced* in the process of the quantum transition of a neutron to a

[2]The statistical behaviour followed by nuclei such as ^{6}Li and ^{14}N seemed to reveal an incompatibility in the quantum theory of the proton-electron atomic structure and the "exchange theorem" (now known as the spin-statistics theorem).

proton:

$$n \to p + e^- + \nu, \tag{1.2}$$

Assuming that the nucleus is a bound state of protons and neutrons (a hypothesis which was put forward shortly after Chadwick's discovery of the neutron), the nuclear β-decay of Eq. 1.1 can be reinterpreted as the decay of a neutron inside the nucleus according to Eq. 1.2.

At the approximate same time, the positron (antiparticle of the electron) was discovered in cloud chamber experiments observing cosmic rays. Radioactive β^+-decay, emitting a positron, were predicted by Fermi's theory and were first observed shortly after its formulation. This model of weak interactions is often considered as the first example of modern field theory, and as of today it is still in many respects relevant to describe weak decay processes. An excellent discussion of the evolution of the idea of the neutrino in these pivotal years of physics history can be found in Ref. [8], with an English translation of both Pauli's 1930 letter and an excerpt of the discussion at the 1933 conference.

It is interesting to note that Pauli had originally proposed his "neutron" as a *constituent* of the nucleus, with a small but non-zero mass which he estimated in his 1930 letter to be of the same order of magnitude as the electron mass. Hence the particle he proposed shared characteristics of Chadwick's neutron and of Fermi's neutrino. After the proton-neutron model of the nucleus was established, and although it had not been observed, the existence of the neutrino raised little doubt given the success of Fermi's theory in the description of radioactive decays. It was rapidly acknowledged that the neutrino mass had to be extremely small, and that neutrinos were possibly massless – implying that they travel at the speed of light just like photons.

Today, the particle denoted as ν in Eq. 1.2 is called an electron antineutrino (denoted as $\overline{\nu}_e$), and is believed to be one of the six "active" (detectable) types of neutrinos. Fermi's theory of β-decay is now seen, in modern terms, as an *effective formulation* of the more general theory of the weak interaction, which is itself encompassed by the Standard Model of particle physics along with the electromagnetic and strong interactions. As we will see, the mass of neutrinos is still one of the most important outstanding questions of particle physics. Indeed, in the Standard Model neutrinos are described as massless, while there is strong experimental evidence that neutrinos have mass. Non-zero neutrino masses are required by the existence of neutrino flavour oscillations, which has been demonstrated at the turn of the 21st century. This is, as of today, the only firmly established experimental fact which indicates the need for physics beyond the Standard Model.

In Appendix A, Sec. A.1 provides an elementary introduction to the view of neutrinos in modern particle physics. The intent is to provide to non-specialists the key elements needed for understanding the remaining part of this section and the following chapters. The glossary may also be useful. The Standard Model is introduced in Sec. A.2. The theoretical building blocks needed to discuss neutrino interactions and neutrino masses are presented, and the electroweak model (unifying electromagnetic and weak interactions) is briefly sketched.

1.1.2 Sources of neutrinos and detection channels

Even though the neutrino was, in some sense, the first "new" elementary particle to be considered in the 1930's, its first experimental observation was only realised in 1956 [9, 10]. This difficulty of detection was foreseen very early, and it was even hypothesized that the neutrino may be undetectable. At the time, ionising radiations (X-rays, γ-rays) and charged particles (β-rays, protons, α-rays) were typically observed via the charge deposited in ionisation chambers or the ionisation tracks produced in cloud chambers. Because they have no electric charge, neutrinos do not ionise atoms on their passage. This is also the case of the neutron, which was first detected indirectly through the observation of protons ejected by the impact of

neutrons on nuclei at rest. The ejection of protons in such collisions is however the result of the strong (nuclear) interaction, in which neutrinos do not take part.

Considering the possibility of such indirect detection of neutrinos, the cross-section for the interaction of neutrinos with nuclei in usual matter was first calculated in 1934 based on Fermi's theory. For neutrinos with an energy of about 1 MeV (the typical energy of ν's emitted by β-decays) the cross-section was estimated to be extremely small [3]:

$$\sigma < 10^{-44}\,\mathrm{cm}^2.$$

This corresponds to an absorption length (or equivalently, a mean free path):

$$L_{\mathrm{abs}} > 10^{14}\,\mathrm{km}.$$

This means that a single neutrino of energy $\sim$ 1MeV, in ordinary dense matter, travels undisturbed on an average distance of several light-years before interacting with the medium. This is the same order of magnitude as the distance between the Sun and the closest star Alpha Centauri. Present day calculations give a similar result. Even though this cross-section does increase about linearly with energy, it becomes "large", *i.e.* so that the absorption length is of the order of the Earth diameter, only for extremely high energy (TeV) astrophysical neutrinos.

The only possibility to detect particles having such a small cross-section is either to dispose of a very intense source, to build a very large detector, or both at the same time. As shown on Fig. 1.2, a large number of physical processes produce neutrinos at various scales of energy, and with extremely intense fluxes for some of them. The following natural phenomena are known to produce neutrinos in the Universe:

- nuclear processes in the Sun (*solar neutrinos*),
- impact of cosmic rays onto gas molecules in the atmosphere (*atmospheric neutrinos*),
- natural radioactivity in Earth's crust and mantle (*geoneutrinos*),
- supernovae,
- certain other astrophysical sources, among the so-called cosmic accelerators (*astrophysical neutrinos*),
- the Big Bang (*relic* or *cosmological neutrinos*),

In addition to the above natural sources, some man-made engines constitute intense sources and are very important for neutrino research:

- nuclear fission reactors produce a large flux of electron antineutrinos (*reactor neutrinos*) through the β-decays undergone by fission products,
- intense artifical beams of higher-energetic neutrinos (*accelerator neutrinos* or *beam neutrinos*) can also be obtained at particle accelerators.

All these sources fundamentally differ by the range of energies of the neutrinos they produce, overall spanning twelve decades in energy (see Fig. 1.2). Some of them also produce specific neutrino flavours. In the following, I briefly describe for each type of source the associated mechanisms of neutrino production and the typical means of neutrino detection in the corresponding energy range and flavour.

Terrestrial sources and the Sun are first discussed, following a historical order. The first detection ever realised was that of reactor neutrinos (of energies of the order $\sim$ 1 MeV) in 1956. This historical example is used to introduce experimental concepts common to most detection experiments. Reactor neutrinos were followed in the 1960's by accelerator and atmospheric neutrinos, which cover a wider range of energies from $\sim$ 0.1 GeV to $\sim$ 10 GeV (at accelerators) and up to several TeV (atmospheric). Solar neutrinos (keV to MeV) were discovered in 1970.

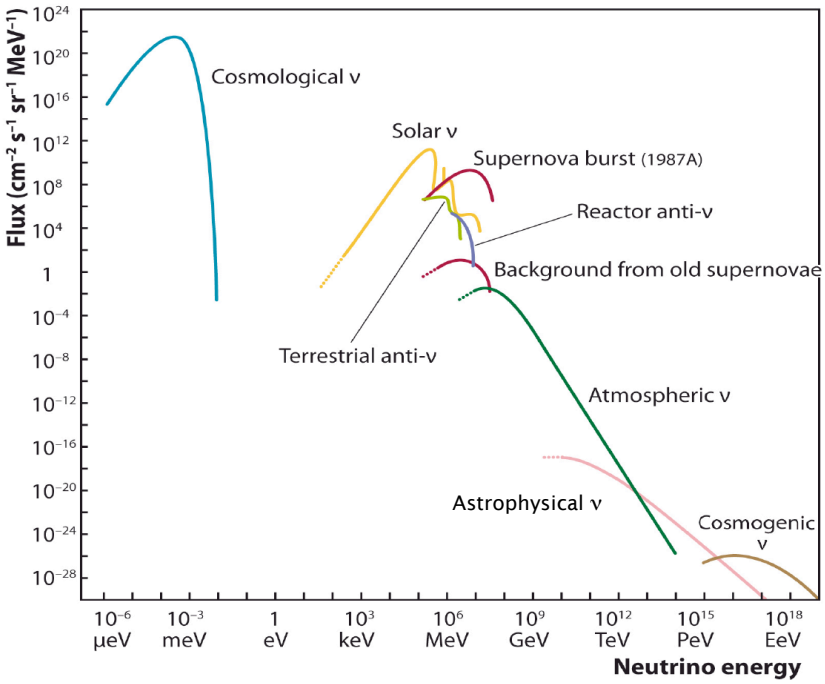

Figure 1.2: Typical energy spectrum for the fluxes of various natural and man-made neutrino sources. The flux of accelerator neutrinos (not shown), ranging around $\sim 0.1 - 10\,\mathrm{GeV}$ energies, is dependent on the production method and accelerator power. The label 'terrestrial anti-ν' refers to geoneutrinos (which are electron antineutrinos), while 'cosmogenic ν' are the yet hypothetical neutrinos produced by the interaction of ultra-high energy comsic rays with photons of the cosmic microwave background.

The first detection of geoneutrinos from Earth's natural radioactivity (MeV) was reported in 2005.

As to sources outside of the Solar System, low-energy (MeV) supernova neutrinos were detected first, in 1987. Much more recently the first evidence was reported for the detection of high-energy astrophysical neutrinos (up to $\sim 100\,\mathrm{PeV}$), both as a diffuse flux (2013) and as a point-like emission (2018). Relic neutrinos from the Big Bang (as known as *cosmic neutrino background*) have never been detected directly.

Reactor neutrinos and the Reines and Cowan experiment

In nuclear fission reactors, energy is produced through the fission of heavy nuclei (^{235}U, ^{238}U, ^{239}Pu, ^{241}Pu) into lighter nuclei. The neutron-rich fission fragments are generally unstable and undergo a sequence of β-decays (producing antielectron neutrinos) until stable isotopes are reached. On average, each fission produces about 200 MeV of energy and releases six $\overline{\nu}_e$'s. As a result, such reactors are a very intense source of electron antineutrinos, producing approximately 2×10^{20} $\overline{\nu}_e$ per second and per giga-watt of thermal power ($\mathrm{GW_{th}}$). A typical modern power plant has several reactor cores, each with a thermal power of about $3\,\mathrm{GW_{th}}$ [1]. The produced flux of neutrinos is essentially isotropic, and therefore decreases rapidly with the distance L from the reactor (proportionally to $1/L^2$).

In reactor neutrino experiments, neutrinos are detected through the inverse beta-decay (IBD) process:

$$\overline{\nu}_e + p \rightarrow n + e^+$$

Since a neutron is heavier than a proton, this process can only occur if the neutrino energy is sufficient to compensate the difference of mass energy between final and inital state. The corresponding energy threshold is:

$$E_{\mathrm{th}} \simeq 1.806\,\mathrm{MeV}$$

Neutrinos with $E_\nu < E_{\text{th}}$ are not detectable via the IBD process.

As a prototype of IBD-based detection techniques, let us quickly describe the Reines and Cowan experiment of 1956 which used as a source the Savannah river reactor in South Carolina, USA. This will also be the occasion to introduce some general concepts of experimental neutrino physics.

In a dense medium such as a liquid water target, the positron e^+ created in an IBD process travels on a short distance (losing kinetic energy) until it annihilates with an electron and produces two 0.51 MeV photons (γ-rays) through a process called *pair annihilation*:

$$e^+ + e^- \to \gamma + \gamma$$

This occurs shortly after the neutrino interaction, and it is a distinctive event since the produced photons have anticolinear momenta (*i.e.* they are emitted in opposite directions) and can be detected close to simultaneously. In the 1956 experiment two tanks filled with 200 L of water were used as the target for the interaction of neutrinos with protons, and the coincident photons from positron annihilation were detected using tanks of liquid scintillator placed on each side of the targets. In response to a γ-ray, the scintillator produces a flash of visible light which is detected by photomultiplier tubes (PMTs) placed on the external sides of the scintillator tanks.

In order to uniquely identify the IBD events and thus better distinguish signal from background (see next paragraph), a scheme was designed for detecting the neutron as well as the positron-induced photon coincidence. To this end, the detection tanks were filled with cadmium chloride to increase the probability of neutron absorption. Cadmium is an excellent neutron absorber, and its de-excitation after absorbing a neutron emits additional γ-rays:

$$^{108}\text{Cd} + n \to {}^{109}\text{Cd}^* \to {}^{109}\text{Cd} + \gamma,$$

These γ-rays are detected by the same means as the ones from positron annihilation. An essential difference is that before its energy is low enough to be absorbed, the IBD neutron diffuses in the target medium (*i.e.* it randomly scatters on nuclei in water, gradually losing its kinetic energy) for a significantly larger time than the positron. The simultaneous emission of a positron and a neutron therefore yields a first coincident pair of photons from positron annihilation (*prompt coincidence signal*), followed by another emission of γ-rays from neutron capture delayed which is delayed by a few μs (*delayed coincidence signal*). This provides a distinctive signature, characteristic of inverse beta decay and thus signaling a neutrino detection.

Note on signal and background in neutrino experiments In neutrino physics and more generally in particle physics, the main experimental challenge often consists in distinguishing *signal* from *background*. For instance, in an IBD-based experiment a *signal event* consists in the neutrino-induced simultaneous emission of a neutron and a positron somewhere in the detector. A background source (or simply "background") is any physical process emitting a neutron and/or a positron of similar characteristics, or anything that can be confused as such by the detector apparatus. In particular, the delayed coincidence signal can easily be faked by a so-called "stray neutron" originating from the interaction of a cosmic ray secondary in the detector. Cosmic rays and their secondaries are a typical background shared by most neutrino experiments. They are generically referred to as *cosmic background*. Another important source of background for the detection of low-energy neutrinos is natural radioactivity. In addition, in reactor experiments neutrons and γ radiation are emitted by the reactor itself.

Three years before their successful detection, Reines and Cowan had performed a first experiment where the rate of measured events was about five per minute, while the expected rate of signal events was less than one per minute ($0.1 - 0.3$/min). The experiment was deemed inconclusive, because the amount of background events largely exceeded the signal[3], mostly due to the cosmic background (the reactor-induced backgrounds were efficiently shielded by thick walls

of paraffin, borax and lead). In the 1956 experiment, the detector was placed 12 m underground, which further reduced the flux of atmospheric muons. In addition, two identical detection tanks were disposed in-between three scintillator tanks in a so-called "club sandwich" arrangement. This allowed to reject events from cosmic background using an *anticoincidence veto*: a signal event in the uppermost tank would not produce γ's able to reach the lowermost detector (and conversely), while an event from a cosmic ray shower would most likely trigger all three detectors in a random combination.

In the 1956 experiment, an event count consistent with the predicted IBD cross-section was measured, with a ratio of signal to overall background of $3:1$. The observed event count was dependent on the reactor power (five times larger when the reactor was running than when turned off), and several cross-checks were performed to rule out other reactor-induced backgrounds as the source of this signal. In June 1956, a telegram was sent to Pauli to announce that the neutrino had been detected, and the discovery was published in Nature [9] and Science [10] in the same year.

Reines was awarded the Nobel prize in 1995 for the discovery of the neutrino (Cowan had already passed away). As a complement of their original 1956 papers [9, 10], a detailed historical discussion of the successive Reines-Cowan experiments (including their initial plan to use a nuclear bomb as source of neutrinos) can be found in Ref. [11].

As will be seen in Sec. 1.3.1 and 1.3.2, reactor neutrino experiments were of great importance to the understanding of neutrino oscillations in the past two decades, and a new generation of experiments is being designed to tackle the remaining unknowns.

Accelerator neutrinos

Higher-energy neutrinos ($E \sim$ GeV) can be produced artificially by bombarding a fixed target (for instance made of solid carbon) with a beam of high-energy protons obtained in a particle accelerator. The collisions between protons and nuclei in the target produce secondary hadrons (via the strong interaction). The proton beam energy and target are chosen so that most of them are pions $\pi^{\pm}$ and kaons $K^{\pm}$, whose decays produce mainly muons $\mu^{\pm}$ and muon neutrinos ν_μ, $\overline{\nu}_\mu$:

$$\pi^+ / K^+ \to \mu^+ + \nu_\mu \tag{1.3}$$

$$\pi^- / K^- \to \mu^- + \overline{\nu}_\mu \tag{1.4}$$

The majority of beam neutrino experiments employ the *Decay In Flight (DIF)* design to generate the neutrino beam. As a modern day example, the NuMI beamline at Fermilab (currently the most powerful neutrino beam worlwide) is shown on Fig. 1.3.

In this type of beam, the pions and kaons are first focused in a narrow beam using a high-current magnetic *focusing horn*. This step also selects them based on their charge: for instance π^- and K^- are focused while π^+ and K^+ are directed towards a beam dump. The selected pions and kaons then propagate along a *decay tunnel* which length is adjusted so that most of them have time to decay freely, but the majority of the produced muons do not have time to decay. Finally, a dense material at the end of the decay pipe absorbs the remaining secondaries apart from muons.

The ability to switch the polarity of the horn current (and therefore the direction of the magnetic field) allows to select either positively or negatively charged pions and kaons, which

[3] More precisely when the signal-to-background ratio is small, the signal search can only be conclusive if the expectations for the measured background rate and measured signal rate are both calculated very precisely (small *systematic uncertainty*), and the total number of accumulated events is large enough to reduce the relative impact of statistical fluctuations (small *statistical uncertainty*).

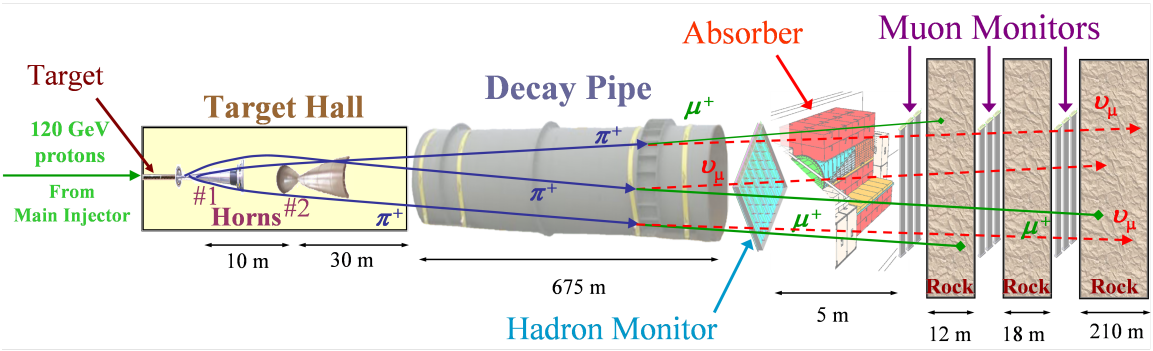

Figure 1.3: Schematic diagram of the NuMI beamline at Fermilab. Image taken from [12] (see also [13]).

respectively selects neutrinos or antineutrinos at the end of the beamline. Most modern beam neutrino experiments can thus function either in "ν focusing mode" or "$\overline{\nu}$ focusing mode".

A small contamination of the beam by wrong-flavour and wrong-sign neutrinos (ν_e and $\overline{\nu}_\mu$ in ν mode, $\overline{\nu}_e$ and ν_μ in $\overline{\nu}$ mode) comes mainly from $\mu^\pm$ which decay before being absorbed, for instance in ν mode:

$$\mu^- \rightarrow e^- + \overline{\nu}_e + \nu_\mu$$

To estimate precisely this contamination as well as the other properties of the beam, various monitors can be installed at different steps of the beamline (see Fig. 1.3).

Typically, the energy of neutrinos from such DIF beams is distributed around a few GeV. The energy range and the spread of the distribution can be tuned by adjusting:

- the energy of the primary proton beam,
- the focusing properties of the magnetic horn,
- and the position of the detector with respect to the beam axis.

The first notable accelerator neutrino experiment was the 1962 BNL-Columbia experiment at the Stanford Linear Accelerator Center (SLAC), which discovered the muon neutrino [14]. It made use of a spark chamber, a type of segmented tracker detector. Later, the Gargamelle detector at CERN, a large size heavy liquid bubble chamber, observed neutral current neutrino interactions (a few events) for the first time [15]. The tau neutrino ν_τ was also first detected directly in an accelerator experiment at Fermilab in 1997 [16]. In the past two decades, numerous accelerator experiments have contributed to the study of neutrino oscillations with a variety of detection techniques (see Sec. 1.3.1).

Atmospheric neutrinos

Atmospheric neutrinos are created by the interaction of cosmic rays (CR) (mostly protons and light nuclei such as helium) with nuclei in the upper part of the atmosphere. In a manner similar to the production of an artificial neutrino beam at a particle accelerator, the incoming cosmic rays (referred to as the *primary* CR) are highly-energetic (up to 10^{20} eV), while in comparison the target nuclei in gas molecules are essentially at rest. The collisions produce the so-called *secondary*[4]CR, which include a variety of hadrons – in particular, pions. These pions

decay to (anti)muons and muon (anti)neutrinos:

$$\pi^- \to \mu^- + \overline{\nu}_\mu \tag{1.5}$$
$$\pi^+ \to \mu^+ + \nu_\mu \tag{1.6}$$

Depending on the energy, muons are the main component of the flux of charged secondary CR seen at the Earth surface. They are referred to as *atmospheric muons* or sometimes as *cosmic muons* (though the muon component in the primary cosmic ray flux is very small). Nevertheless, a significant number of them decay before reaching us, producing two neutrinos per decay:

$$\mu^- \to e^- + \overline{\nu}_e + \nu_\mu \tag{1.7}$$
$$\mu^+ \to e^+ + \nu_e + \overline{\nu}_\mu \tag{1.8}$$

The flux of atmospheric neutrinos therefore contains ν_e, $\overline{\nu}_e$, ν_μ and $\overline{\nu}_\mu$ (τ leptons are not produced in significant quantities in the secondary CR, thus the $\nu_\tau/\overline{\nu}_\tau$ component is negligible). Their energy ranges from $\sim 0.1\,\text{GeV}$ up to very high energies (since the maximal energy of the primary CR is extremely high), with a steeply decreasing energy spectrum. The flavour composition, energy and zenith dependence of the atmospheric flux will be detailed further in Sec. 2.1.1, with a discussion of the associated uncertainties.

The detection of atmospheric neutrinos has been reported for the first time in 1965, by two experiments operating scintillator detectors in gold mines in India and South Africa. Note that the necessity to shield the detector from atmospheric muons is more stringent for atmospheric neutrino detectors because, in contrast with reactor or accelerator experiments, the neutrinos we aim at detecting do not come from a preferred direction. For that reason, atmospheric experiments operate either deep underground, in tunnels below mountains, under the sea, or in Antarctic ice. The effective reduction of the atmospheric muon flux depends on the integrated matter density over the detector. This shielding (called *overburden*) is often characterized in units of *meters water equivalent (mwe)*, corresponding to the equivalent shielding achieved with a water overburden. For instance, the 1965 detection was achieved with detectors operating with about 8000 mwe.

These early experiments detected only the muon neutrinos emerging from the interactions of ν_μ and $\overline{\nu}_\mu$:

$$\nu_\mu + N \to \mu^- + X$$

where N refers to a nucleus and X to an unspecified hadronic final state. Modern experiments are able to detect electron neutrino events as well.

The first large underground water-Cherenkov detectors IMB and Kamiokande started operating in the early 1980's. They were designed for the observation of nucleon decay, as predicted by theories beyond the Standard Model. Although they were not built with the primary goal of detecting atmospheric neutrinos, their observations provided the first hints pointing towards the existence of neutrino flavour oscillations.

In a water-Cherenkov detector, neutrinos are detected through the Cherenkov light (see Sec. 2.1.3) emitted by the charged secondary particles emerging from their interactions. For instance, an electron neutrino with about 1 GeV energy can undergo charged-current quasi-elastic (CCQE) scattering on a neutron:

$$\nu_e + n \to e^- + p$$

[4] In this thesis, the distinction between primary and secondary CR is understood from the point of view of the atmosphere. In astroparticle physics they are often called *atmospheric* primaries and secondaries, while the terms of primary and secondary CR are reserved for the distinction between the particles accelerated at astrophysical sources and those created by spallation, *i.e.* interaction of the primaries with interstellar gas.

A muon neutrino of equivalent energy produces a muon neutrino instead of an electron:

$$\nu_\mu + n \rightarrow \mu^- + p$$

The Cherenkov light produced by the outgoing charged lepton is usually detected by PMTs surrounding the water tank. An electron or a muon produce different patterns of Cherenkov light, allowing to identify the flavour of the incoming neutrino. The energy and incoming direction can also be estimated from the patterns of Cherenkov light. This will be described in details in the next chapter (Sec. 2.1.3).

IMB (the *Irvine-Michigan-Brookhaven* detector) was located in a salt mine in Ohio, USA. It consisted in a tank filled with 10 kton ultrapure water and surrounded by about 2000 photomultiplier tubes. Kamiokande (*Kamioka Nucleon Decay Experiment*) operated on the same principle in the Kamioka mine (Japan) with about 3 kton of target water. The first phase of Kamiokande started in 1983, and the upgraded Kamiokande-II took data from 1987 to 1995. Starting in 1996, the larger detector Super-Kamiokande (SK) followed on the same site, with 50 kton target water. Interestingly in Super-Kamiokande, the signification of the acronym had been changed to *Kamioka Neutrino Detection Experiment*.

Finally, very large-volume (up to 10^9 ton target mass) neutrino telescopes were built in the 2000's. Their purpose is to detect very high-energy neutrinos emitted by astrophysical sources, with a Cherenkov-based detection technique similar to the one of SK and predecessors. Neutrino telescopes will be described in the last paragraph of this section. For the purpose of neutrino astronomy, atmospheric neutrinos are an unwanted background. However, it was realised that the technology of neutrino telescopes could be adapted for the detection of atmospheric neutrinos in the $1 - 100$ GeV energy range. This is the motivation for the KM3NeT-ORCA experiment, which is the main subject of this thesis.

Solar neutrinos

Similarly to nuclear power plants, the Sun is a powerful source of low-energy (~ 1 MeV) neutrinos, generated as byproducts of the nuclear fusion processes in its core. In contrast with fission reactors (producing $\overline{\nu}_e$'s), only ν_e are produced in the Sun.

Thermonuclear energy is produced in the Sun by two chains of nuclear reactions, called the *pp chain* and the *CNO cycle*. The overall result of both chains is the conversion of four protons and two electrons into a helium nuclei and two electron neutrinos:

$$4p + 2e^- \rightarrow {}^4\mathrm{He} + 2\nu_e$$

with an energy release of about 26.7 MeV per process, converted into photons and kinetic energy of the neutrinos. In the *pp* chain, the main underlying process is the proton-proton fusion into a deuterium nucleus:

$$p + p \rightarrow {}^2\mathrm{H} + e^+ + \nu_e$$

In the CNO cycle, neutrinos are produced by the β^+ decays of nuclei created in fusion processes ($^{15}\mathrm{O}$, $^{13}\mathrm{N}$, ^{17}F).

The energy spectrum of solar neutrinos produced by the *pp* chain, CNO cycle and other decays is shown on Fig. 1.4. The fluxes are predicted rather precisely by the models of thermonuclear reactions in the Sun (from the early 1960's). While the *pp* chain produces the largest flux of neutrinos by far, their detection is more difficult due to their very low energies. Davis's chlorine experiment, for instance, was only sensitive to $^7\mathrm{Be}$ and $^8\mathrm{Be}$ neutrinos. With an upper limit in energy at about 20 MeV, there is no real overlap between the solar ν flux and atmospheric neutrinos ($E_\nu > 0.1$ GeV).

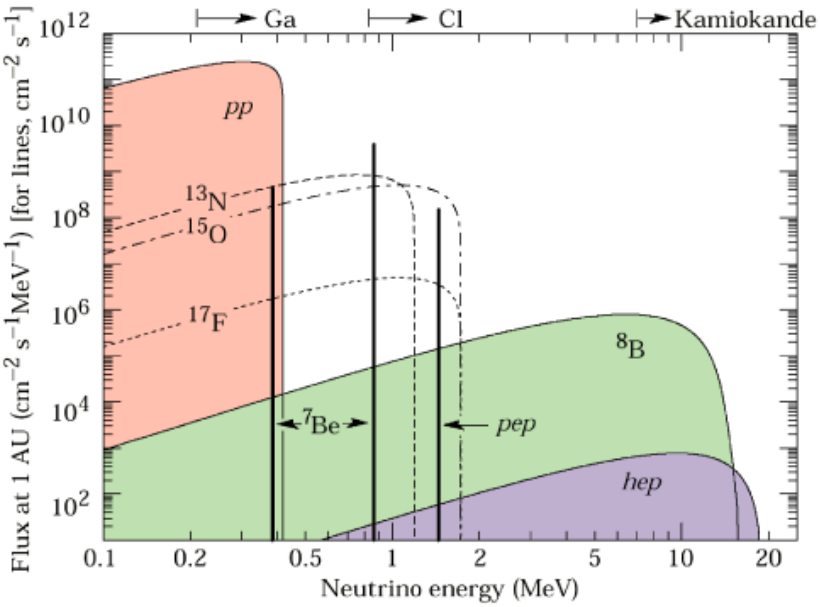

Figure 1.4: Flux of solar neutrinos. Taken from Ref. [17].

The solar neutrinos were first detected in the late 1960's by the Homestake experiment, also known as the Ray Davis chlorine experiment (it was headed by Raymond Davis, Jr. and John Bacall). Operating in a gold mine in South Dakota (USA), it employed a *radiochemical* detection method. Neutrinos from the Sun induced a nuclear transition $^{37}\mathrm{Cl} \to {}^{37}\mathrm{Ar}$ through the inverse β-decay process:

$$\nu_e + {}^{37}\mathrm{Cl} \to {}^{37}\mathrm{Ar} + e^-$$

which has an energy threshold of about 0.814 MeV. The target was a tank of 600 m^3 of liquid $C_2\,Cl_4$, amounting to 133 ton of ^{37}Cl. Neutrino IBD events were not detected on an individual basis as in the Reines-Cowan experiment. Instead, the argon produced in the tank was extracted chemically after a long period of exposure, and the amount of radioactive ^{37}Ar atoms was measured using classical counting devices for the measurement of radioactivity (proportional counters). In total, the Homestake experiment monitored the solar neutrino flux for about 25 years, until 1994.

In the 1990's, a second generation of radiochemical experiments was operated: GALLEX (*The Gallium Experiment*) and GNO (*Gallium Neutrino Observatory*) at the Gran Sasso underground laboratory in Italy and SAGE (*Soviet-American Gallium Experiment*) at the Baksan Neutrino Observatory in Caucasus (Russia). This second generation of experiments used the gallium reaction

$$\nu_e + {}^{71}\mathrm{Ga} \to {}^{71}\mathrm{Ge} + e^-$$

which has a lower energy threshold of about 0.233 MeV and thus allows to study the solar neutrino spectrum more broadly.

Starting in the 1980's, solar neutrinos were also observed in real-time by the large water-Cherenkov detectors Kamiokande and its successor Super-Kamiokande, which were already mentioned as atmospheric neutrino detectors. In addition, the Sudbury Neutrino Observatory (SNO) experiment operated from 1999 to 2006 in Ontario (Canada) with a 1 kton heavy-water vessel.

In Kamiokande and Super-Kamiokande, solar neutrinos are detected through their elastic scattering (ES) on electrons:

$$\nu_\alpha + e^- \to \nu_\alpha + e^- \quad \text{(ES)}$$

where the electron in the initial state (an atomic electron) is ejected from the atom, towards the direction opposite as the direction of the incoming neutrino. One can therefore measure the solar neutrino flux by considering the directionality of the recoil electron (the distribution of background events is mostly isotropic). This process is about six times more sensitive to ν_e than ν_μ and ν_τ neutrinos.

The SNO experiment, in addition to the ES channel, was also sensitive to the charged current (CC) and neutral current (NC) scattering processes on deuterium ^{2}H:

$$\nu_e + {}^2\mathrm{H} \rightarrow 2p + e^- \quad \text{(CC)} \tag{1.9}$$

$$\nu_\alpha + {}^2\mathrm{H} \rightarrow p + n + \nu_e \quad \text{(NC)} \tag{1.10}$$

At the energies of solar neutrinos, the CC channel is sensitive to ν_e only, while the NC channel has the same cross-section for all flavours of neutrinos.

We will see in Sec. 1.1.3, that the different neutrino detection channels in SK and SNO were of prime importance in solving the solar neutrino problem and establishing the existence of neutrino oscillations. Ray Davis (who designed the Homestake solar neutrino experiment) and Masatoshi Koshiba (who led the Kamiokande collaboration) were awarded the Nobel prize in 2002 "for pioneering contributions to astrophysics, in particular for the detection of cosmic neutrinos". Indeed, the observation and subsequent study of solar neutrinos marked the birth of *neutrino astronomy* – understood as the detection of neutrinos of extraterrestrial origin. Homestake was the ground-breaking experiment in this field, and Kamiokande confirmed their results. In addition, Kamiokande was, in 1987, one of the three experiments to first detect neutrinos from outside the solar system (generally referred to as *astrophysical neutrinos*).

Geoneutrinos

Geoneutrinos are emitted in decays of radioactive nuclei occurring naturally in the Earth. Most are $\overline{\nu}_e$ originating in β-decay branches of ^{40}K, ^{232}Th and ^{238}U. Only the latter two are detectable via the conventional IBD mechanism, having energies above the 1.8 MeV threshold. The first detection of geoneutrinos was reported by the KamLAND experiment [18]. KamLAND, primarily designed for observing $\overline{\nu}_e$ produced by surrounding nuclear reactors, was a scintillator detector in the cavity formerly occupied by Kamiokande. Collectively, geoneutrinos carry information about the abundances and spatial distribution of their radioactive sources inside the Earth. which is of importance for geophysics[5].

Astrophysical neutrinos

For the same reason which makes them difficult to detect (they interact only via the weak interaction), neutrinos have a great potential as messengers of the highest-energetic astrophysical phenomena. Neutrinos can escape even very dense environments without being absorbed, and since the interstellar medium is mostly formed by very low density gas they can travel undisturbed on cosmological distances. Moreover, the fact that they do not interact electromagnetically means that they are not deflected by magnetic fields, and do not interact with the cosmic microwave background (nor other photon backgrounds). This contrasts with charged cosmic rays, which generally cannot be traced back to their source because of magnetic deflection, and with γ-rays (and generally electromagnetic radiation) which are much more easily absorbed. Neutrinos from extraterrestrial sources have been detected in very distinct energy

[5] This relates to the long-standing question of Earth's internal 'heat budget' (contribution of the radiogenic heat to the total surface heat flux) which is an important test for many geophysical and geochemical models of the Earth. Ultimately, neutrino geology may help studying the mantle homogeneity and stratification and give insights into the processes of Earth's formation.

ranges where their flux can exceed the background of atmospheric neutrinos: at the MeV scale and below (solar and supernova neutrinos), and at the TeV scale and above (cosmic neutrinos).

On 24 February 1987, a very bright type-II supernova (SN) was discovered in the Large Magellanic Cloud, a neighbouring galaxy, at a distance of about 50 kpc from the Solar System. This event is known as SN1987A. It was the first supernova to be visible with the naked eye since the Kepler supernova in 1604 and it remains, as of today, the best studied SN. A few hours before the visible light from SN1987A reached the Earth, three different large underground detectors of neutrinos observed an unusual *burst* of neutrinos, *i.e.* an excess of detection events in a short duration. In a time window of about 10 s, Kamiokande-II (Japan), IMB (USA) and Baksan (Russia) detected respectively 12, 8 and 5 events, which all three cases was way above the usual background rate.

Very large-volume neutrino telescopes were built in the 2000's, after two decades of peparatory projects, for the detection of neutrinos from cosmic sources. These detectors instrument vast volumes (up to $1\,\mathrm{km}^3$) of deep water or ice in natural environments (the Mediterranean Sea, Antarctic ice and Lake Baikal) with photomultipliers arranged in three-dimensional arrays.

The first project for an underwater neutrino telescope, dubbed DUMAND, started in the late 1970s with the aim to deploy a detector off the coast of Hawaii at 5 km depth. The project was later cancelled due to technical difficulties. In Russia the Baikal collaboration began work at the approximate same time to construct a similar detector under the surface of the Lake Baikal, which was completed in 1998. The use of deep Antarctic ice as a detection medium was pioneered by AMANDA, starting in 1993[6]. Upon completion of AMANDA a much larger detector called IceCube, with a total instrumented volume of about $1\,\mathrm{km}^3$, was built from 2000 to 2010 [19]. In the Mediterranean, the ANTARES, NEMO and NESTOR collaborations demonstrated the viability of the sea-based detection technique, and ANTARES has been operating since 2008 with a total instrumented volume of about $0.015\,\mathrm{km}^3$ [20]. The construction of the next-generation Mediterranean underwater neutrino telescope KM3NeT was started in 2012 at two sites offshore France and Italy [21]. The third currently operating neutrino telescope is Baikal, and its successor Baikal-GVD is under construction as well [22].

The discovery of a diffuse astrophysical flux of high-energy neutrinos was announced by the IceCube collaboration in 2013 [25]. Much more recently, IceCube detected a high-energy neutrino event coincident in direction and time with a gamma-ray flare from the known blazar TXS 0506+056 [26]. The reinvestigation of 9.5 years of data prior to this event showed a 3.5σ excess of high-energy events at that position, which constitutes the first evidence of point-like astrophysical neutrino emission [27].

The water and ice-based Cherenkov detection technique will be discussed in details in the next chapter (Sec. 2.1.3), focusing in particular on the detection of few-GeV atmospheric neutrinos with undersea neutrino telescopes like KM3NeT-ORCA.

1.1.3 Neutrino flavour oscillations and their discovery

The idea of a mixing between ν_e and ν_μ was first proposed by Maki, Nakagawa and Sakata in 1962, and the first phenomenological model of $\nu_e \rightleftarrows \nu_\mu$ oscillations was proposed by Pontecorvo in 1967 [28]. Shortly after, Davis and his collaborators published the first measurements of the flux of ν_e from the Sun with the Homestake detector, which was about 30% of the flux predicted by the standard solar models (SSM) [29]. This inconsistency became known as the *solar neutrino problem*.

[6] More information about these early projects can be found *e.g.* in Refs. [23, 24] and references therein.

The solar neutrino problem

Homestake, started in 1968, continued taking data for more than 20 years, strenghtening its measurement of the deficit of solar ν_e. In the late 1980's this deficit was confirmed by the water-Cherenkov detector Kamiokande, measuring about 40% of the expected flux [30]. Due to their energy threshold , Homestake and Kamiokande were only sensitive to a fraction of the neutrinos produced in the Sun, mainly those produced by ^{8}B decays. In the 1990's, the Gallium-based radiochemical experiments (SAGE and Gallex/GNO) measured the ν_e flux on a wider energy range including neutrinos from the *pp* chain, the main neutrino production mechanism in the Sun. They consistently measured about one-half of the predicted flux. Finally, in the late 1990's Super-Kamiokande reiterated the measurement of the flux of ^{8}B neutrinos and found a consistent deficit of ν_e [31].

As explained in Sec. 1.1.2, the radiochemical experiments are only sensitive to ν_e and the water-Cherenkov detectors, via the elastic scattering of neutrinos on atomic electrons, are about six times more sensitive to ν_e than ν_μ and ν_τ. The decisive experimental result was achieved by SNO with the measurement of the all-flavour flux of neutrinos via the NC scattering (1.10). This process is equally sensitive to ν_e, ν_μ and ν_τ, with a threshold of 2.2 MeV. SNO also measured the rate of ES interactions (1.10) and of CC interactions (1.10), specifically sensitive to ν_e. The ratio of the CC to NC rates deviated from unity with a very high significance[7]:

$$\frac{\Phi_{\mathrm{CC}}}{\Phi_{\mathrm{NC}}} = 0.340 \pm 0.023\,(stat.)\,^{+0.029}_{-0.031}\,(syst.)$$

On the other hand, the all-flavour NC rate was found to be consistent with the SSM predictions. This finally proved that the deficit of ν_e was very likely to be explained by flavour transitions $\nu_e \to \nu_\mu$ or $\nu_e \to \nu_\tau$.

The atmospheric neutrino anomaly

As seen in Sec. 1.1.2, atmospheric neutrinos are produced via the series of decays

$$\pi^- \to \mu^- + \overline{\nu}_\mu, \quad \mu^- \to e^- + \overline{\nu}_e + \nu_\mu$$
$$\pi^+ \to \mu^+ + \nu_\mu, \quad \mu^+ \to e^+ + \nu_e + \overline{\nu}_\mu$$

At energies below 1 GeV, most of the (anti)muons decay before reaching the surface of the Earth. As a result, an underground atmospheric neutrino detector is expected to see approximately two muon (anti)neutrinos for each electron (anti)neutrino:

$$\frac{N_\mu}{N_e} \equiv \frac{N_{\nu_\mu+\overline{\nu}_\mu}}{N_{\nu_e+\overline{\nu}_e}} \sim 2 \tag{1.11}$$

Although the uncertainties on the overall normalisation of the atmospheric flux are rather large (up to 30%), these uncertainties essentially cancel when examining flux ratios.

The first deviations from the 2 : 1 expectation were reported in the 1980's when water-Cherenkov experiments like IMB and Kamiokande attempted to characterise the flux of atmospheric neutrinos, originally seen as an unwanted background for the observation of nucleon decay. The results were reported as the ratio of the measured flavour ratio to the expected flavour ratio:

$$R = \frac{(N_\mu/N_e)_{\mathrm{obs}}}{(N_\mu/N_e)_{\mathrm{th}}}$$

[7] Conventionally, the first error indicated in Eq. 1.1.3 corresponds to the statistical error, while the second one is the systematic error, which is asymmetric here.

IMB and Kamiokande reported respectively $R \simeq 0.54$ and $R \simeq 0.60$, indicating a deficit of muon (anti)neutrinos with respect to electron (anti)neutrinos. In addition, Kamiokande found a dependence of R on the measured incoming direction (zenith angle) of the neutrinos.

In 1998 the Super-Kamiokande collaboration reported compelling evidence for flavour oscillations of atmospheric neutrinos [32]. The evidence came from the asymmetry between the total measured fluxes of upward-going ("upgoing") and downward-going ("downgoing") neutrinos, defined for a given flavour $\alpha = e, \mu$ as:

$$A_\alpha^{\text{up-down}} = \frac{N_\alpha^{\text{up}} - N_\alpha^{\text{down}}}{N_\alpha^{\text{up}} + N_\alpha^{\text{down}}}$$

For multi-GeV neutrinos the influence of the geomagnetic field on the primary CR can be neglected, so that the rate and angular distribution of neutrino production are approximately the same everywhere in the atmosphere. At such energies neutrino absorption in the Earth is also negligible. Therefore, in the absence of flavour transitions the total upgoing and downgoing fluxes of neutrinos received at the location of a detector are expected to be equal, *i.e.* $A_\alpha^{\text{up-down}} = 0$ for both flavours $\alpha = e,\ \mu$.

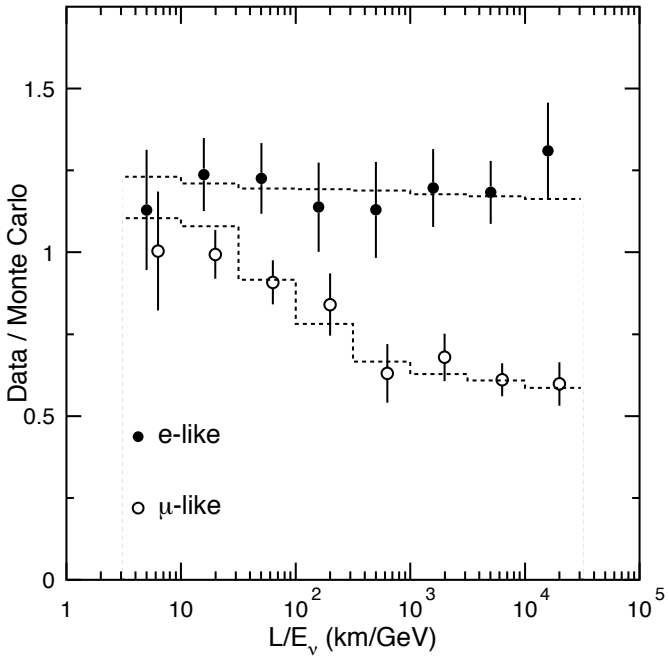

Figure 1.5: First evidence for atmospheric neutrino oscillations as presented by Super-Kamiokande in 1998 [32]. The ratio of the number of data events to Monte Carlo (MC) simulated events is shown as a function of reconstructed L/E_ν. Only events fully contained (FC) in the detector are selected. The points show the ratio of observed data to MC expectation in the absence of oscillations. The dashed lines show the expected shape for $\nu_\mu \leftrightarrow \nu_\tau$ oscillation assuming $\Delta m^2 = 2.2 \cdot 10^{-3}\text{eV}^2$ and $\sin^2 2\theta = 1$. The slight L/E_ν dependence for e-like events is due to the contamination (2-7%) of misidentified ν_μ CC interactions.

In Ref. [32], SK measured the up-down asymmetry for μ-like multi-GeV[8] events as

$$A_\mu^{\text{up-down}} = -0.296 \pm 0.048\,(stat.) \pm 0.01\,(syst.)$$

[8] More precisely the SK multi-GeV sample included neutrinos with "visible energy" $E_{\text{vis}} > 1.33\,\text{GeV}$, where E_{vis} is defined as the energy of the charged lepton which would produce on average the amount of Cherenkov light observed in the detector.

In contrast, the measured asymmetry for e-like events was consistent with zero. The discrepancy between the measured rate of μ-like events and the prediction in the absence of oscillations exhibited a dependance with zenith angle. In fact, as shown on Fig. 1.5 this rate was found to depend even more clearly on the ratio L/E, with

- L the reconstructed propagation length (distance between the production point in the atmosphere and the detector, directly inferred from the reconstructed zenith angle),
- and E the reconstructed neutrino energy.

As seen in the next section, the L/E dependance is the signature of neutrino oscillations. The SK data strongly suggested that the main oscillation channel was $\nu_\mu \to \nu_\tau$. In other words, muon (anti)neutrinos "disappeared" by changing into tau (anti)neutrinos which were mostly invisible in this energy range due to the large mass of the τ lepton[9].

The SK results were confirmed in the following years by two experiments based on completely different detection methods: *Soudan-2*, an iron tracking calorimeter located in Minnesota (USA), and *MACRO*, a scintillator detector in LNGS.

Two-flavour neutrino oscillations

Neutrino flavour oscillations are described as a consequence of the following features of massive neutrinos:

- the neutrino states with definite mass (ν_1, ν_2, ν_3) are *mixed* with the neutrino states with definite flavour (ν_e, ν_μ, ν_τ) which take part in the weak interaction,
- the masses of neutrinos are not all identical (implying that one, at least, is non-zero)

In the following the mechanism will be illustrated on the example of two-neutrino oscillations in vacuum. The two-neutrino case is useful to understand the principle of neutrino oscillations, and in many cases a two-neutrino approximation is sufficient to describe experimental data.

The mass eigenstates are denoted generically $|\nu_1\rangle$ and $|\nu_2\rangle$. By definition, they have a definite energy E_i, meaning that they are eigenstates of the free hamiltonian (describing the propagation of neutrinos in vacuum):

$$\widehat{H}\,|\nu_i(t)\rangle = E_i\,|\nu_i(t)\rangle \qquad (i=1,2)$$

with the energy eigenvalues E_i given by the relativistic relation

$$E_i = \sqrt{|\vec{p}|^2 + m_i^2}.$$

$\vec{p}$ is the momentum of the propagating neutrino state. In their own rest frame, the eigenstates have energy $E_i = m_i$ (hence their denomination as mass states).

The flavour eigenstates will be denoted generically $|\nu_\alpha\rangle$ and $|\nu_\beta\rangle$. By definition, the production or the detection of a neutrino always occurs via the weak interaction and therefore either in the state $|\nu_\alpha\rangle$ or in the state $|\nu_\beta\rangle$. The *mixing* of the flavour and mass eigenstates means that $|\nu_\alpha\rangle$ and $|\nu_\beta\rangle$ are linear superpositions of the massive states, with some complex coefficients $U_{\alpha1}$, $U_{\alpha2}$, $U_{\beta1}$, $U_{\beta2}$:

$$|\nu_\alpha\rangle = U_{\alpha1}\,|\nu_1\rangle + U_{\alpha2}\,|\nu_2\rangle \tag{1.12}$$

$$|\nu_\beta\rangle = U_{\beta1}\,|\nu_1\rangle + U_{\beta2}\,|\nu_2\rangle \tag{1.13}$$

[9] The SK multi-GeV data sample included neutrinos with energies up to approximately 8 GeV [33]. As will be detailed in Sec. 2.1.2, at these energies the Cherenkov signal related to ν_τ appearance is very faint, both because of the small ν_τ cross-section and the immediate decay of the τ lepton. At higher energies, the signal of ν_τ appearance would likely have been seen mainly in the e-like channel.

As a consequence, the states $|\nu_\alpha\rangle$ and $|\nu_\beta\rangle$ do not have definite masses. Due to the requirements of quantum mechanics, the coefficients U_{ij} are not completely arbitrary[10].

In the two-neutrino case the mixing can be parametrised with a single angle $0 \leq \theta \leq \pi/4$:

$$\begin{pmatrix} |\nu_\alpha\rangle \\ |\nu_\beta\rangle \end{pmatrix} = \begin{pmatrix} \cos\theta & \sin\theta \\ -\sin\theta & \cos\theta \end{pmatrix} \begin{pmatrix} |\nu_1\rangle \\ |\nu_2\rangle \end{pmatrix} \tag{1.14}$$

The corresponding transformation of the states is simply a rotation (see the adjacent figure 1.6). The inverse mixing relation is then given by the opposite rotation:

$$\begin{pmatrix} |\nu_1\rangle \\ |\nu_2\rangle \end{pmatrix} = \begin{pmatrix} \cos\theta & -\sin\theta \\ \sin\theta & \cos\theta \end{pmatrix} \begin{pmatrix} |\nu_\alpha\rangle \\ |\nu_\beta\rangle \end{pmatrix} \tag{1.15}$$

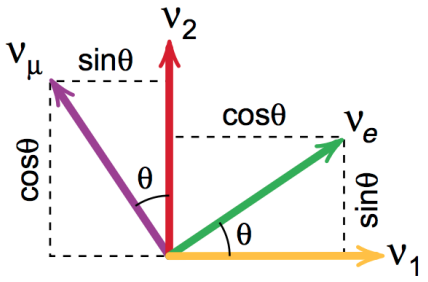

Figure 1.6: The flavour states ($|\nu_\alpha\rangle$, $|\nu_\beta\rangle$) and mass states ($|\nu_1\rangle$, $|\nu_2\rangle$) form two distinct bases of the vector space of neutrino states, related by a rotation of angle θ. In the figure, $\alpha = e$ and $\beta = \mu$. Taken from Ref. [34].

The wave functions of $|\nu_1(t)\rangle$ and $|\nu_2(t)\rangle$ verify the Schrödinger equation and are eigenstates of the hamiltonian, therefore they evolve in time as plane waves:

$$i\frac{\mathrm{d}}{\mathrm{d}t}|\nu_i(t)\rangle = \widehat{H}\,|\nu_i(t)\rangle = E_i\,|\nu_i(t)\rangle \Rightarrow |\nu_i(t)\rangle = e^{-iE_it}\,|\nu_i(t=0)\rangle\,. \tag{1.16}$$

Let us consider a neutrino created at $t = 0$ in the flavour state $|\nu_\alpha\rangle$:

$$|\nu_\alpha(t=0)\rangle = |\nu_\alpha\rangle = \cos\theta\,|\nu_1\rangle + \sin\theta\,|\nu_2\rangle\,.$$

The state at a later time t is denoted as $|\nu_\alpha(t)\rangle$ to keep track of its initial flavour. However, at $t \neq 0$ it is in general a superposition of flavour states. Indeed its time evolution is given by:

$$\begin{aligned} |\nu_\alpha(t)\rangle &= \cos\theta\, e^{-iE_1t}\,|\nu_1\rangle + \sin\theta\, e^{-iE_2t}\,|\nu_2\rangle \\ &= \cos\theta\, e^{-iE_1t}\big(\cos\theta\,|\nu_\alpha\rangle + \sin\theta\,|\nu_\beta\rangle\big) + \sin\theta\, e^{-iE_2t}\big(\sin\theta\,|\nu_\alpha\rangle + \cos\theta\,|\nu_\beta\rangle\big) \\ &= \big(\cos^2\theta\, e^{-iE_1t} + \sin^2\theta\, e^{-iE_2t}\big)\,|\nu_\alpha\rangle + \big(\cos\theta\sin\theta\, e^{-iE_1t} + \sin\theta\cos\theta e^{-iE_2t}\big)\,|\nu_\beta\rangle\,. \end{aligned}$$

Thus, in general $|\nu_\alpha(t)\rangle$ is no longer a flavour eigenstate. At any instant t, if the neutrino described by the state $|\nu_\alpha(t)\rangle$ interacts via a charged current ($W^\pm$ boson exchange), the probability that it interacts as a neutrino of the *other* flavour is given by the projection of $|\nu_\alpha(t)\rangle$ on the eigenstate $|\nu_\beta\rangle$:

$$\begin{aligned} P_{\alpha\to\beta}(t) &= |\langle\nu_\beta|\nu_\alpha(t)\rangle|^2 \\ &= \left|\cos\theta\sin\theta\, e^{-iE_1t} + \sin\theta\cos\theta e^{-iE_2t}\right|^2 \\ &= \sin^2(2\theta)\sin^2\left[(E_2 - E_1)t\right]. \end{aligned} \tag{1.17}$$

From (1.17) it is clear that the probability of flavour transition can be non-zero if

- $\theta \neq 0$ (the flavour and mass eigenstates are not identical),
- and the mass eigenstates have different energies $E_2 \neq E_1$.

[10] The orthonormality conditions $\langle\nu_\alpha|\nu_\alpha\rangle = \langle\nu_\beta|\nu_\beta\rangle = 1$, $\langle\nu_\alpha|\nu_\beta\rangle = 0$ impose that the matrix U is unitary. The most general parametrisation of a 2×2 unitary matrix requires one angle θ and three complex phases. In this case, the complex phases are all unphysical: they can be eliminated by rephasing three of the four neutrino states without physical effect. The mixing angle θ can be chosen as $0 \leq \theta \leq \pi/4$ for the same reason.

In that case, $P_{\alpha\to\beta}(t)$ has an oscillatory time dependence.

Physically, neutrino oscillations are a phenomenon of wave interference, as illustrated on Fig. 1.7. The relative phase of the waves[11] $|\nu_1(t)\rangle$ and $|\nu_2(t)\rangle$ shifts in time, because they do not have the same phase velocity $v_i = E_i/p_i$. The description of neutrino states as plane waves is an abstraction, but is sufficient to understand the most important features of flavour oscillations. Describing them as wave packets localised in space-time is more physical and allows to account for loss of coherence. As a consequence, $|\nu_1(t)\rangle$ and $|\nu_2(t)\rangle$ interfere. When this interference in constructive with respect to $|\nu_\beta\rangle$, it is destructive with respect to $|\nu_\alpha\rangle$. The $|\nu_\beta\rangle$ component of $|\nu_\alpha(t)\rangle$, which is initially zero, can thus become arbitrary large.

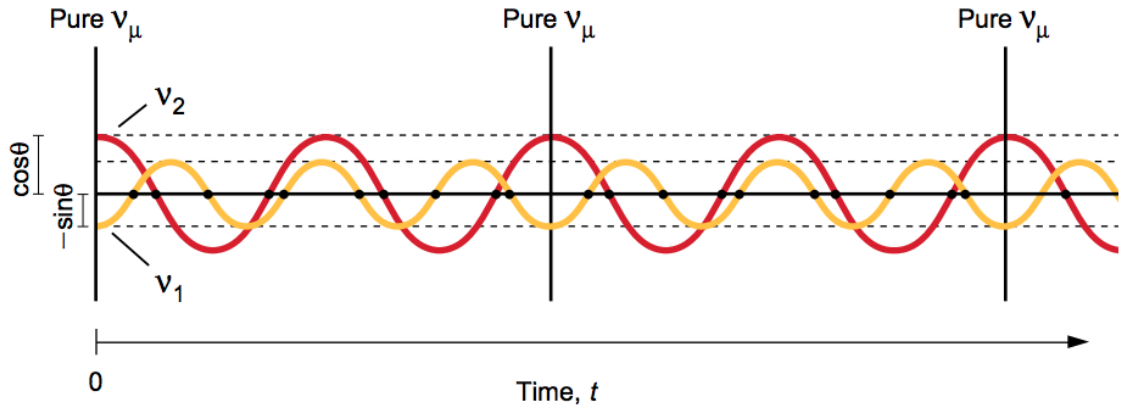

Figure 1.7: Two-neutrino oscillations as an interference phenomenon. The phases associated to the $|\nu_1\rangle$ and $|\nu_2\rangle$ components of the state $|\nu_\mu(t)\rangle$ are represented in yellow and red, following the colour code of Fig. 1.6. It is clear from the figure that interference arises from the fact that both these amplitudes are non-zero (neutrino mixing) and that the phases evolve with different frequencies ($E_2 \neq E_1$).

The neutrinos considered in most oscillation experiments propagate in an ultrarelativistic regime, since their masses are known to be at most $m_i \sim$ eV and conventional detection techniques are only sensitive to neutrinos with energies above a fraction of MeV. In the ultrarelativistic regime, $|\vec{p}| \ll m$ so that the energy can be approximated as

$$E_i = \sqrt{|\vec{p}|^2 + m_i^2} \simeq p + \frac{m_i^2}{2p}.$$

Thus, at first order in $m_i/|\vec{p}|$:

$$E_2 - E_1 \simeq \frac{m_2^2 - m_1^2}{2|\vec{p}|} \simeq \frac{\Delta m^2}{2E}, \tag{1.18}$$

where the squared mass splitting is defined as

$$\Delta m^2 = \Delta m^2_{21} = m_2^2 - m_1^2.$$

Eq. 1.18 shows that neutrino oscillations only occur if $m_1 \neq m_2$. The observation of neutrino flavour oscillations thus demonstrates that at least one of the states $|\nu_1\rangle$ or $|\nu_2\rangle$ has a non-zero mass.

The proper time t of the neutrino state in Eq. 1.17 cannot be easily measured. However, in the ultrarelativistic approximation neutrinos propagate at the speed of light and t (in natural units) is then equal to the distance L between the production and detection of the neutrino. L is usually called the *baseline*. The flavour transition probability in the ultrarelativistic approximation is then

$$P_{\alpha\to\beta}(L,E) = \sin^2(2\theta)\sin^2\left[\frac{\Delta m^2 L}{4E}\right] \tag{1.19}$$

The flavour transition $P_{\alpha\to\beta}$ is also referred to as an *appearance probability*. The so-called *survival probability* is the probability of no flavour transition:

$$P_{\alpha\to\alpha}(L,E) = 1 - P_{\alpha\to\beta}(L,E).$$

Both probabilities are oscillatory functions of the variable L/E. An illustration for a fixed neutrino energy can be seen on Fig. 1.8. The amplitude of the oscillations is determined by the mixing angle, with the maximum appearance probability $P_{\max} = \sin^2(2\theta)$. The so-called *maximal mixing* situation arises when $\theta = \pi/4$, so that $P_{\max} = 1$ and the flavour conversion is total for certain values of L/E.

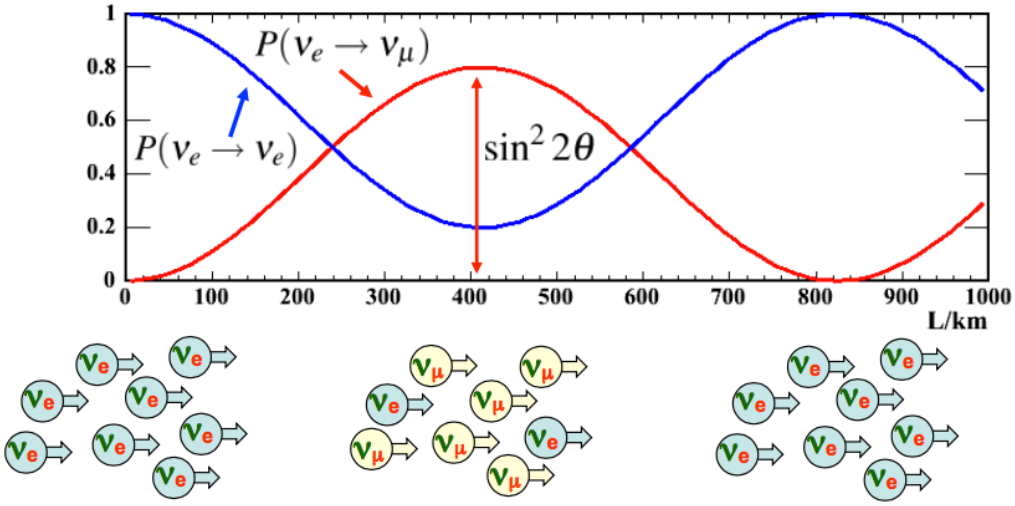

Figure 1.8: Two-neutrino oscillation probabilities as a function of the baseline L for a fixed energy $E_\nu = 1\,\text{GeV}$. A single neutrino has no defined flavour (quantum superposition) as long as it propagates; when interacting at a given L its flavour is randomly determined according to the probabilities P. The sketch of the change of flavour of a fraction of an ensemble of neutrinos thus illustrates the fact that measuring P at a given value of L/E_ν requires to observe the flavour of a large sample of neutrinos. Oscillation parameter values $\sin^2 2\theta = 0.8$ and $\Delta m^2 = 3\cdot 10^{-3}\,\text{eV}^2$ are assumed. The corresponding oscillation length is $L_{\text{osc}}(E_\nu) = 823$. Figure taken from Ref. [2].

The phase in Eq. 1.19 can be written in usual units

$$\Phi = \frac{\Delta m^2 L}{4E} = 1.27\frac{(\Delta m^2/\text{eV}^2)\times(L/\text{km})}{(E/\text{GeV})}$$

For a fixed neutrino energy E, the *oscillation length* $L_{\text{osc}}(E)$ is defined as the spatial period of the oscillations:

$$\Phi = 2\pi\frac{L}{L_{\text{osc}}}.$$

The oscillation length is inversely proportional to the absolute value of the squared-mass splitting $|\Delta m^2|$:

$$L_{\text{osc}}(E) = \frac{4\pi E}{|\Delta m^2|} \qquad \text{(in natural units)} \tag{1.20}$$

$$\simeq 2.47\times\frac{E}{|\Delta m^2|} \qquad \text{(in km)} \tag{1.21}$$

where in the second equality E is expressed in GeV and Δm^2 is expressed in eV^2. At a fixed neutrino energy the spatial oscillations are thus fast for large $|\Delta m^2|$, and slow for small $|\Delta m^2|$.

In general, in order to be sensitive to the flavour oscillation $\nu_\alpha \to \nu_\beta$, an experiment must probe a range of L/E so that $\Phi \sim 1$. If $\Phi \ll 1$, the flavour transition is not observable. If $\Phi \gg 1$, in general the oscillations are too fast and cannot be resolved: the experiment can only measure the average transition probability

$$< P_{\alpha\to\beta} >_{L/E} = \frac{1}{2}\sin^2(2\theta)$$

It can be shown easily that the appearance and survival probabilities for the other flavour β are symmetrical:

$$P_{\beta\to\beta}(L,E) = P_{\alpha\to\alpha}(L,E) \tag{1.22}$$
$$P_{\beta\to\alpha}(L,E) = P_{\alpha\to\beta}(L,E) \tag{1.23}$$

These properties are due to the fact that the mixing matrix is real. As a related fact, the mixing of antineutrino states is identical to the mixing of neutrino states, and thus the replacements of ν_α, ν_β with the antineutrino flavours $\overline{\nu}_\alpha$, $\overline{\nu}_\beta$ leaves the oscillation probabilities unchanged.

Finally, note that no assumptions were made on the ordering of the neutrino masses, *i.e.* whether $m_1 < m_2$ or $m_1 > m_2$. When neutrinos propagate in vacuum, the two-neutrino flavour transition probabilities are invariant under a change of sign of Δm^2. This sign is however physical, and can be revealed once matter effects come into play[12].

Conclusion on the atmospheric and solar flavour transitions

The SK and SNO collaborations were led respectively by T. Kajita and A. McDonald, who were jointly awarded the Nobel prize in 2015 *"for the discovery of neutrino oscillations, which shows that neutrinos have mass."*.

Although the precise description of the oscillations of Earth-crossing atmospheric neutrinos requires to take into account the effects of three-neutrino mixing and the influence of electrons in the Earth (matter effects), the SK atmospheric data was, in 1998, sufficiently well fit by a two-neutrino $\nu_\mu \to \nu_\tau$ oscillation model (for which the Earth matter has no influence). The two-neutrino atmospheric oscillation parameters $\theta_{\rm atm}$ and $|\Delta m^2_{\rm atm}|$ were estimated as:

$$|\Delta m^2_{\rm atm}| \simeq 2.3 \times 10^{-3}\,{\rm eV}^2, \quad \sin^2 2\theta_{\rm atm} \simeq 1.0 \tag{1.24}$$

The $\nu_\mu \to \nu_\tau$ channel was favoured over the $\nu_\mu \leftrightarrow \nu_e$ channel (for which matter effects needed to be accounted for). $\nu_\mu \leftrightarrow \nu_e$ oscillations were also disfavoured in the same period by the results of the CHOOZ experiment in Ardennes (France). CHOOZ studied a similar L/E range as SK, measuring the disappearance of reactor $\overline{\nu}_e$ ($E \sim$ MeV), at a baseline $L \simeq 1\,$km. Their results were consistent with the hypothesis of no $\overline{\nu}_e \to \overline{\nu}_\mu$ oscillations [35].

The theoretical solution to the solar neutrino problem is rather different. In this process, matter effects play a crucial role. ν_e's are produced in the core of the Sun where the electron density is extremely high. While they propagate away from the core, the electron density of the medium slowly decreases. As a result, solar ν_e's undergo *adiabatic* resonant flavour transitions, a phenomenon called the Mikheev-Smirnov-Wolfenstein (MSW) effect. The MSW effect is not similar to standard neutrino oscillations, although it also arises from the mixing of mass and flavour eigenstates. After propagating along a distance of the order of the solar radius, flavour oscillations are no longer relevant because of the loss of coherence of the neutrino wave packets (see for instance Refs. [36, 37]).

[12] The simultaneous flip of the ordering and octant ($\Delta m^2 \to -\Delta m^2$, $\theta \to \pi/2 - \theta$) is however unphysical. Here, the choice of assignment of the label 1 or 2 to the mass states is determined by the requirement on the octant of the mixing angle $0 \leq \theta \leq \pi/4$. Alternatively, one can define ν_1 as the lightest mass state, and allow $0 \leq \theta \leq \pi/2$.

The mechanism of the flavour transitions of solar neutrinos, although mostly non-oscillatory, is controlled in a large part by the mixing of the mass and flavour eigenstates. Again, it is well described by a two-neutrino mixing approximation. The MSW effect with large mixing angle (LMA solution) was one of the candidate mechanisms to solve the solar neutrino problem.

The same mixing regime was explored independently in vacuum oscillations of reactor neutrinos with the *Kamioka Liquid Scintillator Neutrino Detector* (KamLAND). The scintillator detector was located in the cavity formerly occupied by the Kamiokande precursor. It observed the $\overline{\nu}_e$'s produced by the surrounding 53 nuclear reactors in Japan (with baselines ranging from 80 km to 800 km), with an additional contribution from the reactors in the rest of the world estimated to $\sim 3\%$. In this case, matter effects were negligible. By observing $\overline{\nu}_e$ disappearance in the L/E range $\sim (2-6)\cdot 10^4\,\mathrm{km/GeV}$, KamLAND successfully narrowed down the solution to the solar neutrino problem as the MSW LMA mechanism. The solar mixing parameters were estimated from the combination of solar and KamLAND data as [38]

$$\Delta m^2_{\mathrm{sol}} \simeq 7.9^{+0.6}_{-0.5}\times 10^{-5}\,\mathrm{eV}^2, \quad \tan^2\theta_{\mathrm{sol}} \simeq 0.40^{+0.10}_{-0.7} \tag{1.25}$$

This picture of solar neutrino physics was completed and the LMA MSW scenario reinforced by the measurement of sub-MeV solar neutrinos at Borexino, an ultrapure liquid scintillator detector operating at LNGS since 2007 [39].

The solar and atmospheric neutrino oscillations thus appear to be controlled by two distinct mixing regimes, corresponding to separate ranges of L/E:

$$\left(\frac{L}{E}\right)_{\mathrm{sol}} \sim 10^4\,\mathrm{km/GeV}, \quad \left(\frac{L}{E}\right)_{\mathrm{atm}} \sim 10^2 - 10^3\,\mathrm{km/GeV} \tag{1.26}$$

As will be adressed in detail in Sec. 1.2.2, three rotation angles (plus one complex phase) are needed to parametrise three-neutrino mixing. Two of them can be identified in first approximation with θ_{sol} and θ_{atm}. There are two independent squared-mass splittings, which can be associated with $\Delta m^2_{\mathrm{sol}}$ and $\Delta m^2_{\mathrm{atm}}$. The third mixing can be seen as coupling the solar and atmospheric oscillation sectors. The two-neutrino model remains a good approximation in many cases, due to the hierarchy of neutrino masses ($\Delta m^2_{\mathrm{atm}} \gg \Delta m^2_{\mathrm{sol}}$) and the smallness of the third mixing angle.

1.2 Theoretical aspects

The main aspects of the theoretical description of massive neutrinos are presented in this section, with a discussion of the outstanding issues. The conventions of notations are listed at the beginning of Sec. A.2 (Appendix A), with an introduction to the description of spin-1/2 fermions in the context of relativistic quantum field theory. Elements of the electroweak Standard Model are also presented in the appendix, with the intent to adopt a didactic approach rather than an exhaustive one. The glossary can also be of use for non-specialists.

1.2.1 Neutrinos in the Standard Model and neutrino masses

The description of electromagnetic and weak interactions was unified in the 1960's by the Glashow theory, which then became the Glashow-Weinberg-Salam model with the addition of the Brout-Englert-Higgs mechanism. The model was finally proved to be mathematically self-consistent (*renormalisable*) in 1971. The Glashow-Weinberg-Salam model, together with the theory of strong interactions which is formulated in the same mathematical framework and acquired its final form in the mid-1970's, are currently known as the Standard Model (SM) of particle physics. The SM knew its first major success with the discovery of neutral-current neutrino interactions, and has been extremely successful ever since, especially in the

precise description of electroweak interactions. Several particles predicted by the model (most notably the quarks b and t and the $W^{\pm}$ and Z^0 bosons) were successfully detected, following huge experimental efforts throughout five decades. The SM reached its ultimate experimental confirmation with the discovery in 2012 and subsequent study at the LHC of the Higgs boson, a cornerstone of the model.

To this day, the only *experimentally proven* deviation from the Standard Model is the existence of neutrino flavour oscillations, established in 1998. Neutrino flavour oscillations are understood to unequivocally demonstrate that neutrinos are massive, while neutrinos in the Standard Model are described as massless.

The theoretical features useful for the discussion of SM neutrinos and neutrino masses are mostly introduced in Sec. A.2.2. They are briefly summarised here:

Massless fermion A two-component spinor field (Weyl spinor) is sufficient to describe a massless spin-1/2 fermion. In this case, one has to pick either the chiral left-handed (χ_L) or the chiral right-handed (χ_R) representation.

Massive Dirac fermion To describe a massive spin-1/2 fermion, a four-component Dirac spinor is needed. Such an object is described by the Dirac equation (or equivalently by the free Dirac Lagrangian). A Dirac spinor field ψ can be decomposed as the sum of two chiral fields ψ_L and ψ_R.

Helicity and chirality The helicity (orientation of the spin with respect to the direction of motion) coincides with the chirality for a massless field, or in the ultrarelativistic limit for a massive field.

Parity P Parity corresponds to the inversion of space coordinates:

$$P: \ (t, \vec{x}) \to (t, -\vec{x})$$

Parity reverses the helicity and exchanges right-handed and left-handed fields.

Symmetries C, P and T The charge conjugation operator C exchanges particles and antiparticles, while the time reversal operator T reverses the time (direction of motion). The free Dirac Lagrangian, describing non-interacting particles, is invariant under the action of all three individual discrete symmetries C, P and T. Invariance by the product CPT is a feature of all local quantum field theories; in a sense this invariance requires the existence of antiparticles.

The description of physical particles by a two-component Weyl spinor was rejected by Pauli in 1933, because it violated parity symmetry. More generally parity was largely considered to be a fundamental symmetry of the Universe. It did (and still does) seem counter-intuitive to most physicists that the laws of physics "prefer" one orientation of space over the other – or equivalently, one geometrical configuration over its image in a mirror. It was however discovered as soon as the 1950's that the weak interaction does violate parity symmetry. Parity violation appeared as the only way to solve a puzzling experimental observation in which two particles which were seemingly identical (having the exact same mass and lifetime) decayed into states with opposite parities. The possibility of parity violation was suggested in 1956 by Lee and Yang, along with propositions of experimental tests. Shortly after, parity violation in weak decays was observed in two different dedicated experiments by Wu *et al.* and Lederman *et al.*

Following this discovery the *two-component theory* of massless neutrinos was proposed in 1957. It assumes that neutrinos are *left-handed* massless fermions, and describes them with left-handed two-component Weyl spinors. One year later, the neutrino helicity was measured

in agreement with this description, via the measurement of the polarisation of photons emitted following an electron capture radioactive process:

$$^{152}\mathrm{Eu} + e^- \to {}^{152}\mathrm{Sm}^* + \nu_e \tag{1.27}$$

$$^{152}\mathrm{Sm}^* \to {}^{152}\mathrm{Sm} + \gamma \tag{1.28}$$

As of today, experimental observations unequivocally indicate that neutrinos are exclusively left-handed particles, while antineutrinos are exclusively right-handed. This is the way neutrinos are described in the SM. Therefore, neutrinos maximally violate both the parity (P) and the charge conjugation (C) symmetries. The combination CP (parity followed by charge conjugation), however, do relate left-handed neutrinos to right-handed antineutrinos:

$$\nu_L \overset{CP}{\rightleftharpoons} \overline{\nu}_R$$

After the discovery of parity violation in the weak interaction it was proposed that CP be the fundamental symmetry between particles and antiparticles. CP however turned out to be violated via the mixing of quarks (accounted for in the SM) and, as we will see, it may well be violated in the mixing of neutrinos (beyond the SM).

Massive Dirac neutrinos

As detailed in Appendix A.2, Sec. A.2, in the SM left-handed fields are introduced as *doublets* of the gauge symmetry group $\mathbf{SU}(2)$. For leptons there are thus three doublets:

$$L'_{eL} \equiv \begin{pmatrix} \nu'_{eL} \\ e'_L \end{pmatrix},\ L'_{\mu L} \equiv \begin{pmatrix} \nu'_{\mu L} \\ \mu'_L \end{pmatrix},\ L'_{\tau L} \equiv \begin{pmatrix} \nu'_{\tau L} \\ \tau'_L \end{pmatrix}$$

On the other-hand, right-handed fields are singlets of $\mathbf{SU}(2)$, *i.e.* they are unaffected by the corresponding gauge transformations. Moreover, there are no right-handed neutrino fields. Thus for leptons we have:

$$e'_R,\ \mu'_R,\ \tau'_R$$

The prime notation ($e_L \to e'_L$, *etc.*) is used to signify that the fields do not necessarily have a definite mass. The asymmetric treatment of the left-handed (doublets) and right-handed (singlets) chiral fields gives rise to the maximal parity-violating structure of the weak CC couplings, and the absence of right-handed neutrino fields forbids that neutrinos acquire mass through the Higgs mechanism.

To incorporate the fact that neutrinos have mass into the SM, the simplest approach is to introduce right-handed neutrino fields as $\mathbf{SU(2)_L}$ singlets:

$$\nu'_{eR},\ \nu'_{\mu R},\ \nu'_{\tau R}$$

The SM Lagrangian then naturally includes additional Higgs-neutrino Yukawa couplings, which give rise after symmetry breaking to the Dirac mass terms for leptons

$$\mathcal{L}_{\mathrm{mass},L} = \frac{v}{\sqrt{2}}\left(\overline{\boldsymbol{l}'_L} Y'^{l} \boldsymbol{l}'_R + \boldsymbol{\nu}'_L Y'^{\nu} \boldsymbol{\nu}'_R\right) + \mathrm{h.c.}$$

where v is the Higgs vacuum expectation value (see A.2) and Y'^{l} and Y'^{ν} are any two independent complex matrices. Similarly to Eq. A.86, each one can be diagonalised via biunitary transformations. This requires four different unitary matrices: V_L^l, V_R^l diagonalising Y'^{l}, and

$V_L^\nu,\ V_R^\nu$ diagonalising Y'^ν:

$$V_L^{l\dagger} Y'^l V_R^l = Y^l = \mathrm{diag}(y_e^l,\ y_\mu^l,\ y_\tau^l) \tag{1.29}$$

$$V_L^{\nu\dagger} Y'^\nu V_R^l = Y^\nu = \mathrm{diag}(y_e^\nu,\ y_\mu^\nu,\ y_\tau^\nu) \tag{1.30}$$

The charged leptons with definite mass are defined as in Eq. A.87, and the neutrino fields with definite mass, denoted as

$$\nu_i = \nu_{iL} + \nu_{iR}, \quad i = 1,\ 2,\ 3$$

are found via the transformations

$$\begin{pmatrix}\nu_{1L}\\ \nu_{2L}\\ \nu_{3L}\end{pmatrix} \equiv V_L^{\nu\dagger} \begin{pmatrix}\nu'_{eL}\\ \nu'_{\mu L}\\ \nu'_{\tau L}\end{pmatrix}, \quad \begin{pmatrix}\nu_{1R}\\ \nu_{2R}\\ \nu_{3R}\end{pmatrix} \equiv V_R^{\nu\dagger} \begin{pmatrix}\nu'_{eR}\\ \nu'_{\mu R}\\ \nu'_{\tau R}\end{pmatrix} \tag{1.31}$$

Neutrino mixing then arises in a similar way as quark mixing in the Standard Model. This will be detailed in the next section.

The introduction of right-handed neutrino fields to explain the observation of neutrino masses is somewhat similar to the 1930 postulate of the neutrino as a possibly undetectable particle. Indeed, as $\mathbf{SU(2)_L}$ singlets, the right-handed neutrino fields would not take part in the weak interaction, and therefore they would not be detectable via direct interactions with matter described by the SM. Such neutrino fields are referred to as *sterile*, while the ones coupling to other particles are called *active*. About his "desperate remedy", Pauli is said to have commented [40]:

> *"I have done something very bad today by proposing a particle that cannot be detected. It is something no theorist should ever do."*

Another issue with this mechanism is the smallness of neutrino masses, which are known to be of the order of a fraction of eV at most. Thus neutrinos are a least six order of magnitudes lighter than the electron which is the lightest fermion. In this minimal extension of the SM, the smallness of neutrino masses requires to assign Yukawa couplings for neutrinos at a radically different magnitude than those of the other fermions. This is generally considered *unnatural* and, as such, unsatisfactory[13].

Majorana neutrinos and seesaw mechanisms

In principle, another way of assigning mass to left-handed neutrinos arises from the fact that neutrinos can be described as *Majorana fermions*. Majorana fermions are different from Dirac fermions, in that they are identical to their antiparticle. Thus, they are described by spinors which are equal to their charge conjugate:

$$\psi = \psi^C = \mathcal{C}\overline{\psi}^T$$

Obviously, only neutral fermions can be described as Majorana fermions.

To the most general massive Dirac field are associated four particle modes: fermion with positive helicity, fermion with negative helicity, antifermion with positive helicity, antifermion with negative helicity. In the case of a Majorana field the number of degrees of freedom is reduced to two: positive and negative helicity. If neutrinos are massive Majorana fermions, the experimental distinction between ν and $\overline{\nu}$ remains, but is only interpreted as the fact that the negative helicity Majorana neutrino fields couple to charged leptons ($e^-,\ \mu^-,\ \tau^-$), while the

[13] The discussion of the role of *naturalness* in the construction of physical theory is deeply interesting, but is outside the scope of this thesis.

fields with positive helicity couple to charged antileptons (e^+, μ^+, τ^+), as is already encoded in the structure of the weak CC couplings.

For a Majorana field $\nu = \nu_L + (\nu_L)^{\mathrm{C}}$, Lorentz scalar products of the form

$$\overline{\nu_L}\,(\nu_L)^C + \text{h.c.} = \overline{\nu_L}\,\nu_R^{\mathrm{C}}\,\nu_L + \text{h.c.} \tag{1.32}$$

correspond to mass terms describing a free Majorana fermion. Therefore, unlike in the Dirac case, describing a massive Majorana neutrino does not require to introduce a right-handed field. However, such terms are incompatible with the electroweak symmetry and the Higgs mechanism[14], thus the assumption of a Majorana neutrino alone is not sufficient to describe massive left-handed neutrinos in the SM.

With the addition of a sterile right-handed Majorana field ν_R, the most general mass term involves both left-handed and right-handed Majorana mass terms as well as a usual Dirac mass term. This so-called *Dirac-Majorana* mass term can be expressed in matrix form as

$$\begin{aligned}\mathcal{L}_{\text{mass},D-M} &= \mathcal{L}_{\text{mass},L} + \mathcal{L}_{\text{mass},R} + \mathcal{L}_{\text{mass},D} \\ &= \frac{1}{2}\left(\overline{\nu_L},\ \overline{\nu_L^{\mathrm{C}}}\right)\begin{pmatrix} m_L & m_D \\ m_D & m_R \end{pmatrix}\begin{pmatrix} \nu_R^{\mathrm{C}} \\ \nu_R \end{pmatrix} + \text{h.c.}\end{aligned} \tag{1.33}$$

where m_L, m_R, m_D are all additional free parameters of the model ($m_L \neq 0$ being forbidden within the SM).

The ν_L and ν_R chiral fields do not have definite mass, however the mass matrix is symmetric and thus diagonalisable, with eigenvalues:

$$m_{1,2} = \frac{1}{2}\left(m_L + m_R \pm \sqrt{(m_L - m_R)^2 + 4m_D^2}\right) \tag{1.34}$$

which are the masses of the neutrino fields (mixtures of ν_R and ν_L) having definite mass.

When assuming no left-handed mass term ($m_L = 0$) and that the Dirac and Majorana masses are of very different scales $m_D \ll m_R$, we obtain one very heavy and one very light physical neutrinos (note that the minus sign in m_{light} can be reverted via a rephasing of the fields):

$$m_{\text{light}} \simeq (-)\frac{m_D^2}{m_R} \tag{1.35}$$

$$m_{\text{heavy}} \simeq m_R\left(1 + \frac{m_D^2}{m_R^2}\right) \simeq m_R \tag{1.36}$$

This type of mechanism is called a *seesaw* mechanism, since m_{light} becomes smaller as m_{heavy} becomes larger. For instance, if the Dirac mass is of the order of the Higgs mass $m_D \simeq 10^2\,\text{GeV}$, and the Majorana mass is of the order of "grand unification"[15] energies $m_R \simeq 10^{16}\,\text{GeV}$, the mass of the light neutrino state is then $m_{\text{light}} \simeq 10^{-3}\,\text{eV}$. Note that $m_L \neq 0$, typically obtained in larger extensions of the SM, may also lead to variants of seesaw mechanisms.

[14] Because ν_L belongs to a weak isospin doublet (with $I_3 = -\frac{1}{2}$ and $Y = -1$), Higgs-lepton couplings generating after SSB a mass term of the form (1.32) would break either the gauge invariance of the Lagrangian or the renormalisability of the theory (see for instance Ref. [1, 41] for details).

[15] *Grand Unified Theories* (GUT) is the designation of the class of models attempting to unify the electroweak and strong interaction into a single gauge interaction. Just like the specific effects of the electroweak unification are only visible above energies of the order of the W and Z masses, the manifestations of grand unification are expected to become apparent only at much higher energies, the so-called *GUT scale*: $\Lambda_{\text{GUT}} \approx 10^{16}\,\text{GeV}$. Some of the Grand Unified Theories include heavy ν_R singlets, for instance the models based on the **SU**(5) or **SO**(10) gauge groups.

Neutrinos are the only particles in the SM that can possibly be described as Majorana fermions and have a Dirac-Majorana mass term. Since a Majorana spinor requires less degrees of freedom than a Dirac spinor, it can be argued that the Majorana description is the most natural assumption. The experimental manifestations of the Dirac or Majorana nature of the neutrinos are however challenging to observe. As of today, it is thus not yet known whether neutrinos are Dirac or Majorana fermions. Experimental tests will be discussed briefly in Sec. 1.3.2.

1.2.2 Three-neutrino mixing and neutrino mass hierarchy

Two-neutrino mixing and flavour oscillations in vacuum were introduced in Sec. 1.1.3. In this section, we present the standard parametrisation of three-neutrino mixing, extend the derivation of vacuum oscillation probabilities for three flavours, and discuss the important differences with the two-neutrino case.

In terms of the lepton fields with definite mass (Eq. A.87, 1.31), the leptonic charged current writes as

$$j^{\mu}_{\mathrm{CC,lep}} = 2 \sum_{\alpha=e,\ \mu,\ \tau} \nu_{\alpha L}\, \gamma^{\mu}\, l_{\alpha,L} = \overline{\boldsymbol{\nu}_L}\, U^{\dagger}\, \gamma^{\mu}\, \boldsymbol{l}_L \tag{1.37}$$

with the mixing matrix U:

$$U \equiv V_L^{l\,\dagger}\, V_L^{\nu}.$$

The left-handed flavour neutrino fields ν_{eL}, $\nu_{\mu L}$, $\nu_{\tau L}$ which couple directly with the charged leptons of the corresponding generation in the weak CC interaction are therefore:

$$\begin{pmatrix} \nu_{eL} \\ \nu_{\mu L} \\ \nu_{\tau L} \end{pmatrix} \equiv U \begin{pmatrix} \nu_{1L} \\ \nu_{2L} \\ \nu_{3L} \end{pmatrix} = \begin{pmatrix} U_{e1} & U_{e2} & U_{e3} \\ U_{\mu 1} & U_{\mu 2} & U_{\mu 3} \\ U_{\tau 1} & U_{\tau 2} & U_{\tau 3} \end{pmatrix} \begin{pmatrix} \nu_{1L} \\ \nu_{2L} \\ \nu_{3L} \end{pmatrix} \tag{1.38}$$

On the other hand, there is no point in defining flavour neutrino fields for the right-handed components ν_{1R}, ν_{2R}, ν_{3R} since those do not couple with the charged leptons. The above mixing can be generalised very similarly to the case of Majorana neutrinos with a Dirac-Majorana mass term.

Note that the charged lepton fields with definite flavour, *by definition*, also have definite mass. Indeed, the mass is what distinguishes physically one charged lepton from another. For instance, in a magnetised tracking detector, muon and electrons with similar kinetic energies differ by the curvature of their trajectories, their energy deposition (via ionisation and other electromagnetic processes), and their potential absorption or decay into lighter particles. All these properties are determined by the difference between the electron and the muon mass. On the contrary, since the flavour of neutrinos is always observed indirectly via CC interactions, the only natural definition of the flavour neutrino fields is through their weak CC coupling with physical charged leptons. There is no way of directly observing the massive neutrino states.

PMNS matrix

The mixing matrix U is called the Pontecorvo-Maki-Nakagawa-Sakata (PMNS) matrix. It can be decomposed into a product of three real rotation matrices, parametrized by mixing angles θ_{12}, θ_{13} and θ_{23}, and a diagonal matrix involving a complex phase $e^{i\delta_{\mathrm{CP}}}$. The standard parametrisation of the PMNS matrix is:

$$U = \begin{pmatrix} 1 & 0 & 0 \\ 0 & c_{23} & s_{23} \\ 0 & -s_{23} & c_{23} \end{pmatrix} \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & e^{i\delta_{\mathrm{CP}}} \end{pmatrix} \begin{pmatrix} c_{13} & 0 & s_{13} \\ 0 & 1 & 0 \\ -s_{13} & 0 & c_{13} \end{pmatrix} \begin{pmatrix} c_{12} & s_{12} & 0 \\ -s_{12} & c_{12} & 0 \\ 0 & 0 & 1 \end{pmatrix} \tag{1.39}$$

with the short notations

$$c_{ij} \equiv \cos\theta_{ij}, \quad s_{ij} \equiv \sin\theta_{ij}.$$

All the mixing angles can be chosen in the interval $0 \leq \theta_{ij} \leq \pi/2$. In contrast with the two-neutrino case where all complex phases can be eliminated, here one physical degree of freedom remains for the complex phase $\delta_{\rm CP}$[16].

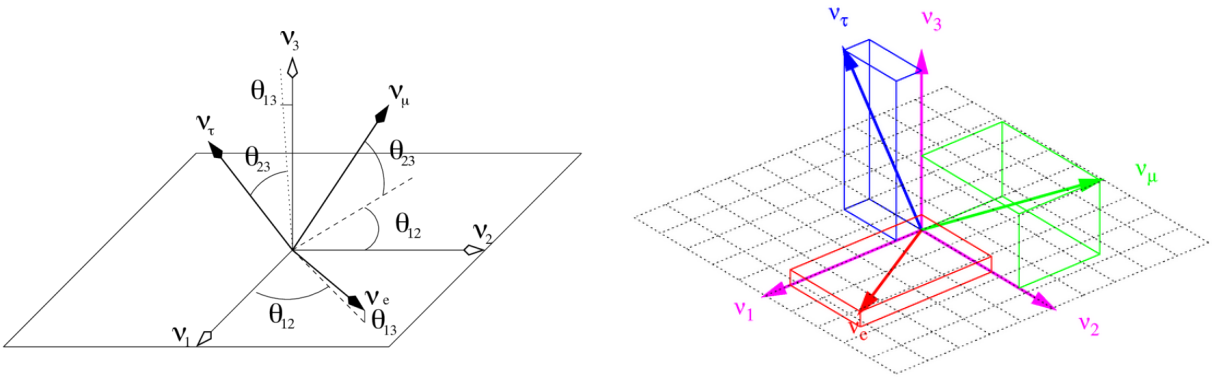

Figure 1.9: Neutrino mixing angles (assuming zero CP violation) may be represented as Euler angles relating the flavour eigenstates ν_e, ν_μ, ν_τ to the mass eigenstates ν_1, ν_2, ν_3. Figures taken from Refs. [42] (left) and [43] (right).

Neutrino oscillations in vacuum are then controlled by the following six independent parameters

$$\theta_{12}, \quad \theta_{23}, \quad \theta_{13}, \quad \delta_{\rm CP}, \quad \Delta m^2_{31}, \quad \Delta m^2_{21}.$$

The mixing angles θ_{ij} and the phase $\delta_{\rm CP}$ determine the mixing matrix U as in Eq. 1.39. Δm^2_{31} and Δm^2_{21} are the two independent squared-mass differences:

$$\Delta m^2_{31} = (m_3)^2 - (m_1)^2 \tag{1.40}$$

$$\Delta m^2_{21} = (m_2)^2 - (m_1)^2 \tag{1.41}$$

The $1-2$ mixing can be associated with the solar oscillation sector:

$$(\Delta m^2_{21},\ \theta_{12}) \longrightarrow \Delta m^2_{\rm sol},\ \theta_{\rm sol} \tag{1.42}$$

while the $2-3$ mixing corresponds to the atmospheric sector:

$$(\Delta m^2_{32} \simeq \Delta m^2_{31},\ \theta_{23}) \longrightarrow \Delta m^2_{\rm atm},\ \theta_{\rm atm}. \tag{1.43}$$

The atmospheric and solar oscillations are coupled by the $1-3$ mixing, parametrised by θ_{13} which is often referred to as the *reactor* mixing angle. As will be seen in Sec. 1.3.1, the $1-3$ mixing is comparatively very small. This contributes to explain the good approximation of solar and atmospheric oscillations by two-flavour models.

[16] The most general unitary U contains six complex phases. If neutrinos are assumed to be Dirac fermions, five of these phases can be eliminated by rephasing the fields with no physical effect. If neutrinos are of Majorana nature, only three phases can be eliminated. The two additional Majorana phases can however be ignored here since they have no effect on neutrino oscillations – more generally the Dirac or Majorana nature of neutrinos does not affect neutrino oscillations.

Hierarchy of neutrino masses

Experiments have shown that the neutrino masses follow a hierarchical pattern:

$$|\Delta m^2_{31}| \simeq 2.5 \cdot 10^{-3}\,\mathrm{eV}^2 \simeq |\Delta m^2_{32}|$$
$$\Delta m^2_{21} \simeq 7.5 \cdot 10^{-5}\,\mathrm{eV}^2 \ll |\Delta m^2_{31}|,\ |\Delta m^2_{32}|$$

It is conventional to label the lightest massive neutrino as ν_1 so that $\Delta m^2_{21} > 0$, which in turn imposes a constraint on θ_{12} from solar data[17]. In contrast, it is still not known whether m_3 is heavier or lighter than m_1. This is referred to as the question of the neutrino mass ordering (NMO). The two possible ordering scenarios are represented on Fig. 1.10.

Assuming that the lightest mass m_1 is much smaller than $\sqrt{\Delta m^2_{\mathrm{sol}}}$, each ordering implies a very different hierarchy of the masses. The case $m_3 > m_{1,2}$ implies

$$m_1 \ll m_2 \ll m_3$$

which is conventionnally referred to as the *normal* hierarchy (NH). The case $m_3 < m_{1,2}$ implies that two masses are quasi-degenerate

$$m_1 \ll m_2 \simeq m_3$$

This is called the *inverted* hierarchy (IH). Finally, the case where all three masses are quasi-degenerate cannot be excluded: if $m_1 \gg \sqrt{\Delta m^2_{\mathrm{sol}}}$, then

$$m_1 \simeq m_2 \simeq m_3 \tag{1.44}$$

Although in the latter case the hierarchy is quasi-degenerate independently of the ordering, the term *neutrino mass hierarchy* (NMH) is frequently employed in the neutrino physics community to refer to the sign of Δm^2_{31}. This will be the case in this thesis.

Mixing of the neutrino and antineutrino states

The particle *states* with definite mass created by the fields ν_i are denoted as $|\nu_i\rangle$ $(i = 1,\ 2,\ 3)$. These are the eigenstates relevant to describe the free propagation of neutrinos. The states created by the fields ν_α are denoted as $|\nu_\alpha\rangle$ $(\alpha = e,\ \mu,\ \tau)$. These are states with definite flavour, relevant to describe the interactions of neutrinos and in particular their production and detection.

The mixing relation between the neutrino states is anologous to the mixing of the fields (1.38), except that it involves the complex conjugates $U^*_{\alpha k}$ of the PMNS matrix entries. Using the fact that U is unitary we can write the inverse relation as well:

$$|\nu_\alpha\rangle = \sum_k U^*_{\alpha k}\,|\nu_k\rangle \tag{1.45}$$
$$|\nu_k\rangle = \sum_\alpha U_{\alpha k}\,|\nu_\alpha\rangle \tag{1.46}$$

[17] Following *e.g.* the review of particle physics [44], Chap. 14. Solar neutrino data constrain the product $\Delta m^2_{\mathrm{sol}} \cos 2\theta_{\mathrm{sol}}$ to be positive (see Sec. 1.2.4). With the labeling convention $m_1 < m_2$, this imposes $\theta_{12} \leq \pi/4$ (*lower octant*). One could equivalently use the opposite convention.

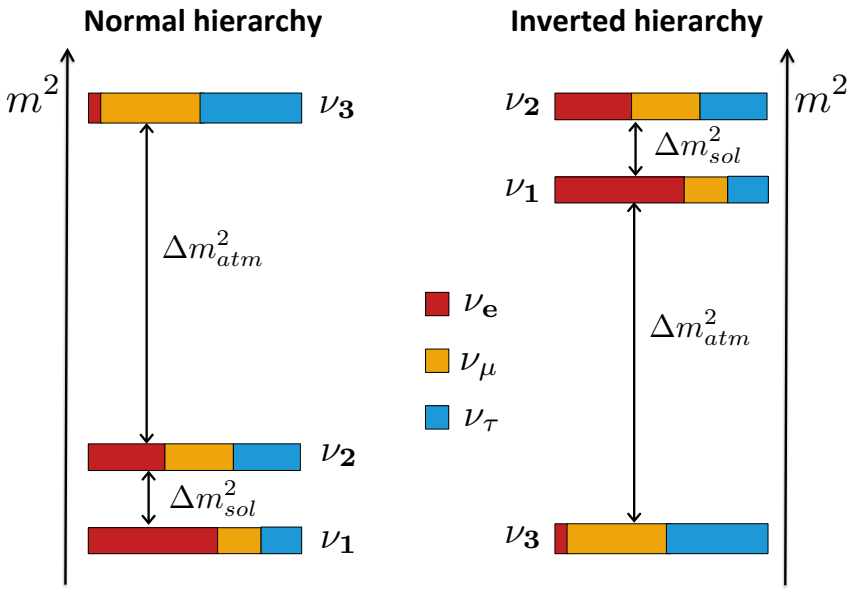

Figure 1.10: The normal and inverted scenarios for the neutrino mass hierarchy (NMH). The colour code indicates the fraction of each flavour (e, μ, τ) present in each of the mass eigenstates (1, 2, 3). Note that the spacing between mass eigenstates on the vertical axis does not scale exactly with the measured Δm^2 values. Figure reused from Ref. [21].

The mixing of antineutrino states[18] (annihilated by the fields ν_α or ν_i and created by their charge conjugated fields) is given by

$$|\overline{\nu}_\alpha\rangle = \sum_k U_{\alpha k} |\overline{\nu}_k\rangle \tag{1.47}$$

$$|\overline{\nu}_k\rangle = \sum_\alpha U^*_{\alpha k} |\overline{\nu}_\alpha\rangle \tag{1.48}$$

If all the entries of $U_{\alpha k}$ are real, the mixing of neutrinos and antineutrinos are identical. A sufficient condition is that $e^{i\delta_{\mathrm{CP}}} = \pm 1$, *i.e.* $\delta_{\mathrm{CP}} = k\pi$. The phase δ_{CP} is thus related to a potential violation of the CP-symmetry in neutrino mixing.

1.2.3 Three-flavour oscillations in vacuum

Flavour transition probabilities can be derived in the three-neutrino mixing case by the same method as in the two-flavour case of Sec. 1.1.3. A neutrino created at $t = 0$ in a state of given flavour α ($\alpha = e$, μ or τ) writes as:

$$|\nu_\alpha(t=0)\rangle = |\nu_\alpha\rangle = \sum_k U^*_{\alpha k} |\nu_k\rangle$$

Its time evolution is

$$|\nu_\alpha(t)\rangle = \sum_k U^*_{\alpha k} e^{-iE_k t} |\nu_k\rangle$$

[18] If neutrinos are Majorana fermions, the term *neutrino* is used conventionnally to refer to the states with negative helicity, and *antineutrino* is used to refer to the states with positive helicity.

By injecting the expression (1.48) for $|\nu_k\rangle$ we get:

$$|\nu_\alpha(t)\rangle = \sum_{\beta=e,\mu,\tau} \left(\sum_{k=1}^{3} U^*_{\alpha k} U_{\beta k} e^{-iE_k t} \right) |\nu_\beta\rangle \,. \tag{1.49}$$

At any instant t, if the particle state $|\nu_\alpha(t)\rangle$ interacts via a charged current (W boson exchange), the probability that it interacts as the flavour state $|\nu_\beta\rangle$ is given by the projection

$$\begin{aligned} P_{\alpha\to\beta}(t) &= |\langle\nu_\beta|\nu_\alpha(t)\rangle|^2 \\ &= \left| \sum_k U_{\alpha k} U_{\beta_k} e^{-iE_k t} \right|^2 \\ &= \sum_k \sum_j U^*_{\alpha k} U_{\beta k} U_{\alpha j} U^*_{\beta j} \, e^{-i(E_k - E_j)t} \end{aligned} \tag{1.50}$$

This is the exact probability for the transition of neutrino states $\nu_\alpha \to \nu_\beta$. The corresponding probability for flavour transitions of antineutrino states is obtained with the interversion

$$U \rightleftharpoons U^* .$$

In the ultrareltivistic regime the phases Φ_{kj} of the exponential factors in Eq. 1.50 can be approximated as

$$\Phi_{kj} = -i(E_k - E_j)t \simeq -\frac{\Delta m^2_{kj} L}{2E} = -2\pi \frac{L}{L^{\text{osc}}_{kj}(E)}$$

Each Δm^2_{kj} has a specific oscillation length

$$L^{\text{osc}}_{ij}(E) = \frac{4\pi E}{\Delta m^2_{ij}} \tag{1.51}$$

One way to express the general form of the flavour transition probability (1.50) in the ultra-relativistic approximation is, in analogy with the two-flavour case:

$$\begin{aligned} P_{\alpha\to\beta}(L,E) = \delta_{\alpha\beta} &- 4 \sum_{j=1}^{3} \sum_{k=j+1}^{3} \text{Re}\left[U^*_{\alpha k} U_{\beta k} U_{\alpha j} U^*_{\beta j} \right] \sin^2(\Phi_{kj}) \\ &+ 2 \sum_{j=1}^{3} \sum_{k=j+1}^{3} \text{Im}\left[U^*_{\alpha k} U_{\beta k} U_{\alpha j} U^*_{\beta j} \right] \sin(2\Phi_{kj}) \end{aligned} \tag{1.52}$$

To go to the antineutrino probability, the change $U \rightleftharpoons U^*$ simply flips the sign of the last term. To illustrate this formula on a simple case, let us assume $\delta_{\text{CP}} = 0$ such that all the coefficients $U_{\alpha i}$ are real and the last term vanishes. Then we obtain for the muon neutrino survival probability:

$$P_{\nu_\mu\to\nu_\mu}(L,E) \simeq 1 - K_{21} \sin^2\left[\frac{\Delta m^2_{21} L}{4E}\right] - K_{31} \sin^2\left[\frac{\Delta m^2_{31} L}{4E}\right] - K_{32} \sin^2\left[\frac{\Delta m^2_{32} L}{4E}\right] \tag{1.53}$$

with

$$K_{21} = 4U^2_{\mu 1} U^2_{\mu 2}, \quad K_{31} = 4U^2_{\mu 1} U^2_{\mu 3}, \quad K_{32} = 4U^2_{\mu 2} U^2_{\mu 3}$$

The appearance probabilities $P(\nu_\mu \to \nu_e)$ and $P(\nu_\mu \to \nu_\tau)$ are obtained with the corresponding replacements of the indices $\mu \to e$ and $\mu \to \tau$.

It is interesting to compare this expression to the corresponding one in the two-flavour case. The main difference is that three distinct oscillation terms coexist, driven respectively by Δm^2_{21}, Δm^2_{31} and Δm^2_{32}.

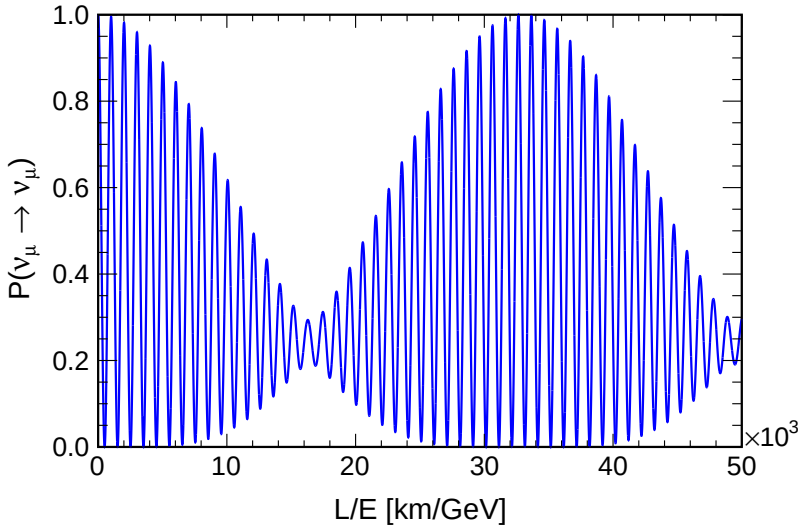

Figure 1.11: ν_μ disappearance probability in vacuum as a function of L/E. The oscillation parameters are set consistently with the ranges obtained by experimental measurements: $|\Delta m^2_{21}| \simeq 7.5 \cdot 10^{-5}\,\mathrm{eV}^2$, $|\Delta m^2_{31}| \simeq 2.5 \cdot 10^{-3}\,\mathrm{eV}^2$, $K_{21} \simeq 0.26$, $K_{32} \simeq 0.55$, $K_{31} \simeq 0.45$.

Fig. 1.11 shows the ν_μ survival probability of Eq. 1.53 as a function of L/E for a choice of oscillation parameters consistent with the values determined experimentally. Since $\Delta m^2_{31} \simeq \Delta m^2_{32}$, two main oscillation frequencies are visible. The faster frequency (short L_{osc}, large Δm^2) correspond to the atmospheric mixing; the slower ones (long L_{osc}, small Δm^2) correspond to the solar mixing.

1.2.4 Flavour oscillations in matter

Neutrinos propagating in ordinary matter undergo *coherent forward elastic scattering* on both electrons and nucleons. In contrast with the (incoherent) scattering processes mentioned previously, this type of interaction corresponds to a negligible transfer of energy and momentum from the point of view of the target particles. The effect on the propagation of neutrinos can be modeled by a potential which modifies the effective mass and flavour content of the propagation eigenstates. This mechanism is analogous to the way the propagation of photons is modified depending of the refractive index of a medium.

A slight reformulation of the vacuum evolution equation of neutrino states is convenient to study the impact of matter effects. In vacuum the hamiltonian $\widehat{H}_0$ is diagonal in the mass basis $|\nu_i\rangle$. In the ultrarelativistic approximation it can be written as:

$$\widehat{H}_{0,\mathrm{mass}} = \frac{1}{2E}\begin{pmatrix} 0 & 0 & 0 \\ 0 & \Delta m^2_{21} & 0 \\ 0 & 0 & \Delta m^2_{31} \end{pmatrix}$$

where the common phase $E_1 = (m_1)^2/2E$ has been subtracted from the diagonal (only relative phases between states are physical). In the flavour basis, $\widehat{H}_0$ is given by

$$\widehat{H}_{0,\mathrm{flv}} = U\,\widehat{H}_{0,\mathrm{mass}}\,U^\dagger.$$

With $t \simeq x$, the Schrödinger equation

$$i\frac{\mathrm{d}}{\mathrm{d}x}\left|\nu_\alpha(x)\right\rangle = \widehat{H}_0\left|\nu_\alpha(x)\right\rangle \tag{1.54}$$

is solved by

$$\left|\nu_\alpha(x)\right\rangle = \exp\left(-ix\widehat{H}_0\right)\left|\nu_\alpha(x=0)\right\rangle = \widehat{\mathcal{U}}(x)\left|\nu_\alpha\right\rangle, \tag{1.55}$$

where the evolution operator $\widehat{\mathcal{U}}(x) \equiv \exp\left(-ix\widehat{H}_0\right)$ has been introduced. Written explicitly in the flavour basis this corresponds to the evolution equation (1.49), factoring out the common phase $e^{-iE_1 t}$.

In the most general case the matter potential depends on the density of electrons and nucleons. The tree-level Feynman diagrams for the elastic scattering processes between (anti)neutrinos and matter are represented on Fig. 1.12. The electron flavour component of a propagating (anti)neutrino state can undergo both CC and NC interactions with electrons, while the μ and τ components undergo only NC interactions. The contributions from NC scattering on electrons and protons are opposite in sign, and cancel in neutral matter. The NC neutrino-nucleon scattering is flavour independent and thus affects the mass eigenstates identically, introducing a common phase which is irrelevant for neutrino oscillations.

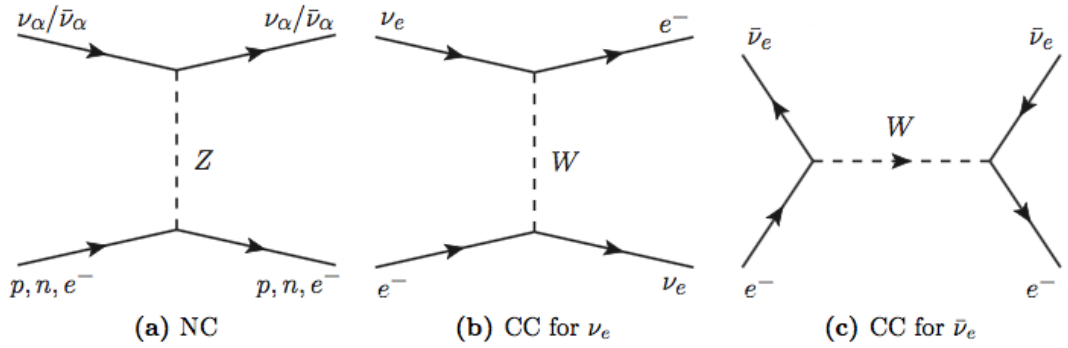

Figure 1.12: Elastic scattering processes between anti(neutrinos) and ordinary matter (n, p, e^-).

The net physical effect on neutrino oscillations then comes from contributions of the CC scattering on electrons. It is modeled by an effective potential V_{CC}, proportional to the electron number density of the medium at position x, denoted as $N_e(x)$:

$$V_{\mathrm{CC}}(x) = \pm\sqrt{2}G_F N_e(x), \tag{1.56}$$

where the positive and negative sign apply respectively to ν_e and $\overline{\nu}_e$, and G_F is the Fermi constant. As first discussed by Wolfenstein in 1978 [45], this potential has the effect of modifying the mixing between propagation and flavour neutrino states, since the propagation hamiltonian is no longer diagonal in the mass basis:

$$\widehat{H}_{M,\mathrm{mass}}(x) = \frac{1}{2E}\begin{pmatrix}0 & 0 & 0\\ 0 & \Delta m^2_{21} & 0\\ 0 & 0 & \Delta m^2_{31}\end{pmatrix} + U^\dagger \begin{pmatrix}V_{\mathrm{CC}}(x) & 0 & 0\\ 0 & 0 & 0\\ 0 & 0 & 0\end{pmatrix} U. \tag{1.57}$$

The effective matter propagation eigenstates at the position x are then found by diagonalising $\widehat{H}_M$. For non-zero electron density the mixing of these effective states with the flavour states, as well as the associated eigenvalues of $\widehat{H}_M$ (the effective masses), differ from the mixing and masses of the original vacuum propagation states.

In a medium with constant density N_e, the propagation hamiltonian $\widehat{H}_M(N_e)$ is position (time) independent and the use of the simple evolution operator (1.55) remains valid. For the purpose of accurate estimations of transition probabilities, any density profile along a baseline L can be approximated sufficiently well as a discrete sequence of N constant density steps $(x_0,\ x_1, ...x_N)$:

$$x_i \leq x \leq x_{i+1} \Rightarrow \begin{cases} N_e(x) = N_{e,i} \\ \widehat{\mathcal{U}}_i(x) = \exp\left(-i\,x\,\widehat{H}_M(N_{e,i})\right) \end{cases}$$

The evolution of the neutrino state over the whole density step i follows

$$|\nu_\alpha(x_{i+1})\rangle = \widehat{\mathcal{U}}_i(x_i+1)\,|\nu_\alpha(x_i)\rangle\,, \tag{1.58}$$

and the transition probability for the full path $L = [x_0,\ x_N]$ is given by the concatenation

$$P_{\alpha\to\beta}(x_N) = \left|\,\langle\nu_\beta|\nu_\alpha(x_N)\rangle\,\right|^2 = \left|\,\langle\nu_\beta|\widehat{\mathcal{U}}_{N-1}(x_N)\ldots\widehat{\mathcal{U}}_1(x_2)\widehat{\mathcal{U}}_0(x_1)\,|\nu_\alpha\rangle\,\right|^2 \tag{1.59}$$

where in each step $\widehat{\mathcal{U}}_i(x)$ is diagonalised by the specific effective mass states [46]. This discretisation approach is used in practice for numerical calculations (see Chap. 4). It is in particular applicable to density profiles containing physical discontinuities, for instance in the Earth.

In the following, the impact of matter effects is detailed in the case of two-neutrino mixing and constant electron density. The three-neutrino case for oscillations in the atmospheric L/E range will be discussed in Chap. 2, Sec. 2.2.1.

Matter resonance in the two-neutrino case

It is physically meaningful to consider matter effects in the two-flavour mixing case, as in Sec. 1.1.3[19]:

$$\begin{pmatrix} |\nu_e\rangle \\ |\nu_\mu\rangle \end{pmatrix} = \begin{pmatrix} \cos\theta & \sin\theta \\ -\sin\theta & \cos\theta \end{pmatrix} \begin{pmatrix} |\nu_1\rangle \\ |\nu_2\rangle \end{pmatrix} \tag{1.60}$$

The propagation hamiltonian in matter with density N_e can then be written explicitly in the flavour basis as

$$\widehat{H}_{M,\text{flv}} = \frac{1}{4E}\begin{pmatrix} \Delta m^2\cos 2\theta + A_{\text{CC}} & \Delta m^2 \sin 2\theta \\ \Delta m^2 \sin 2\theta & \Delta m^2\cos 2\theta - A_{\text{CC}} \end{pmatrix}$$

where A_{CC} is defined as:

$$A_{\text{CC}} \equiv 2EV_{\text{CC}} = \pm 2\sqrt{2}G_F N_e$$

$\widehat{H}_{M,\text{flv}}$ is diagonalised by a rotation of angle θ^M as:

$$U_M^\dagger\,\widehat{H}_{M,\text{flv}}\,U_M = \frac{1}{4E}\begin{pmatrix} -\Delta^M m^2 & 0 \\ 0 & \Delta^M m^2 \end{pmatrix} \text{ with } U_M = \begin{pmatrix} \cos\theta^M & \sin\theta^M \\ -\sin\theta^M & \cos\theta^M \end{pmatrix}$$

θ_M and $\Delta^M m^2$ are the effective mixing angle and the effective squared-mass splitting in matter. They can be expressed for instance as

$$\Delta^M m^2 = \xi\cdot\Delta m^2, \quad \sin^2(2\theta^M) = \frac{\sin^2(2\theta)}{\xi^2}, \tag{1.61}$$

with the mapping parameter

$$\xi = \sqrt{\sin^2 2\theta + \left(\cos 2\theta - \frac{A_{\text{CC}}}{\Delta m^2}\right)^2} \tag{1.62}$$

[19] Again we use the convention to fix $0 \leq \theta \leq \pi/4$ and let Δm^2 be either positive or negative.

bearing in mind that $A_{\rm CC}$ is positive for neutrinos and negative for antineutrinos. It is easy to check that $\sin 2\theta \leq \xi < +\infty$ and that $\xi = 1$ for zero electron density, recovering the vacuum mixing.

In a medium with constant density, the flavour transition probabilities then take the same form as in vacuum, with the replacements $\theta \to \theta^M$ and $\Delta m^2 \to \Delta^M m^2$:

$$P_{\alpha\to\beta}(L,E) = \sin^2(2\theta^M)\sin^2\left(\frac{\Delta^M m^2 L}{4E}\right)$$

Both the amplitude and the frequency of the oscillations are thus modified by matter effects. In certain conditions the oscillations can become resonant, as discovered by Smirnov and Mikheev in 1985 [47]. From Eq. 1.61, the amplitude $\sin^2 2\theta^M$ becomes maximal when $\xi^2 \to \sin^2 2\theta$, *i.e.*

$$A_{\rm CC} = \pm 2\sqrt{2} E G_F N_e \longrightarrow \Delta m^2 \cos 2\theta. \tag{1.63}$$

This condition corresponds to a match of the neutrino energy E and electron density N_e. In addition, the resonance occurs selectively either for neutrinos or for antineutrinos depending on the mass ordering. Indeed, with our convention ($\theta \leq \pi/4$) the sign of $\Delta m^2 \cos 2\theta$ is the sign of Δm^2. Thus, if $\Delta m^2 > 0$, the condition (1.63) is possible only for neutrinos ($A_{\rm CC} > 0$), and if $\Delta m^2 < 0$, only for antineutrinos ($A_{\rm CC} > 0$). The sign of Δm^2, which had no effect in vacuum, becomes observable with matter effects.

More generally, for $\Delta m^2 > 0$ a suppression of the oscillations for antineutrinos is associated to the enhancement of the oscillations for neutrinos, and conversely when $\Delta m^2 < 0$. Although the vacuum mixing (1.60) is CP-invariant, the oscillation probabilities are not, due to the fact that the medium does not contain positrons to compensate the effect of electrons. This is often referred to as a *matter-induced* or *extrinsic* CP violation. In contrast, the T-invariance is preserved here[20], and thus CPT is also violated (Earth matter being non CPT-symmetric).

Solar and atmospheric matter resonances

In the three-neutrino case, two of the three mixing angles (θ_{12} and θ_{13}), as well as the corresponding Δm^2's, acquire effective values similarly to (1.61). Two resonances can then occur in distinct energy regimes, enhancing either the $1-2$ mixing (*solar matter resonance*) or the $1-3$ mixing (*atmospheric matter resonance*). Either way, θ_{23} is unaffected because the matter potential equally affects ν_μ and ν_τ. As detailed in Chap. 2, the enhancement of the oscillation probabilities in the atmospheric regime thus affects subleading terms driven by $\sin^2\theta_{13}$. For neutrinos propagating in Earth matter, the atmospheric matter resonance occurs around $\sim 2-8\,$GeV, depending on their path across Earth's layers.

The observation of the MSW effect for solar ν_e imply that $\Delta m^2_{21}\cos 2\theta_{12} > 0$, since the matter resonance would occur solely for antineutrinos otherwise. This fixes the ordering of the masses m_1 and m_2 (or equivalently the octant of $\theta_{\rm sol}$, as discussed in Sec. 1.1.3). In a similar way, the atmospheric resonance in Earth matter provides a way to measure the NMH (sign of Δm^2_{31}) by determining whether it occurs for neutrinos or antineutrinos. The same resonance also enhances the sensitivity of three-neutrino oscillations to $\delta_{\rm CP}$ via the $\nu_\mu \to \nu_e$ channel. As discussed in Sec. 2.3, it can also be used as a probe of the electron density in the deep Earth, which is inaccessible via standard geophysical methods.

[20] Note that T-invariance is in general violated when neutrinos propagate in a varying density medium [48], as can be understood from Eq. 1.59 where in general $[\mathcal{U}_i, \mathcal{U}_j] \neq 0$ for $i \neq j$.

Resonance enhancement, MSW effect and parametric enhancement

The resonance (1.63), where the mixing can become maximal regardless of the magnitude of the mixing in vacuum, is often referred to as *MSW resonance* from the names of Mikheev, Smirnov and Wolfenstein who discovered it. The enhancement of the oscillation probabilities due to this resonance is known as *resonance enhancement*.

The term of *MSW effect* refers to[21] the description of the *adiabatic* flavour transitions of ν_e as they propagate in a medium of slowly varying density, for instance inside the Sun. In case of smooth density variation, the spatial dependence of $N_e(x)$ can be accounted for analytically in the evolution equation. The mixing is then position-dependent: $\theta_M(x)$. Flavour transitions $\nu_\alpha \to \nu_\beta$ can then occur without the need for interference between mass states with different phases, *i.e.* without neutrino oscillations. The regime of adiabatic flavour transitions corresponds to the case where the spatial dependence of $\theta^M(x)$ is slow enough so that the evolution of the amplitudes of the effective mass states are decoupled (adiabaticity condition). A depiction of this mechanism in the Sun (LMA MSW) can be found in Ref. [36, 49].

Finally, large flavour transitions can occur without the mixing angle becoming large, in the case of a periodic variation of the electron density profile along the neutrino path. This effect is known as *parametric enhancement* of the oscillation probabilities. As a typical example, the path of a neutrino traversing the Earth diametrically crosses the core-mantle boundary (CMB) twice, with symmetrical discontinuities of the density profile. The well-studied parametric enhancement of the oscillation probabilities for neutrinos crossing the Earth will be discussed in Sec. 2.2.1.

1.3 Recent experimental advances and future prospects

1.3.1 Latest oscillation experiments

Neutrino beams to probe the atmospheric sector

KamLAND had been designed to finally solve the solar neutrino problem with the help of artificial neutrinos, and similarly, the atmospheric oscillation regime was explored independently by accelerator experiments. Beam neutrinos produced at accelerator have typical energies $\sim$ 1 GeV, therefore oscillation baselines of about 10^{2-3} km are needed to probe Δm^2_{atm}. The first such long-baseline beam experiments were K2K [50] and MINOS [51]. Both used ν_μ beams with good purity (few % contamination at most). In addition to the main far detector(s), long-baseline beam experiments generally place one or several near detector(s) close to the beamline in order to monitor the neutrino flux before oscillations. K2K and MINOS specifically observed ν_μ disappearance by comparing the event rate at the far detector to the prediction obtained by extrapolating the near detector measurement. K2K (*KEK-to-Kamiokande*), operating from 1999 to 2004, used a neutrino beam produced at the KEK accelerator center in Tsukuba and pointed towards the Super-Kamiokande detector 250 km away. The MINOS beam was produced at the Fermilab NuMI beamline (introduced previously in Fig. 1.3). The far detector was located 700 m underground in the Soudan mine (Minnesota) for a baseline of 735 km. The near and far detectors were both segmented steel-scintillator tracking calorimeters. In both K2K and MINOS, near detectors were functionnally similar to the far detectors, allowing the cancellation of part of the detection-related systematic uncertainties.

The MINOS detector, being optimised for GeV neutrinos with a consequent overburden of 2070 mwe, was also sensitive to atmospheric neutrinos. By combining beam and atmospheric data, MINOS was able to provide the strongest constraints on the atmospheric parameters *circa*

[21] It is often used in a broader sense to refer the MSW resonance. About the terminology, see for instance Ref. [49].

2013 [52]:

$$\Delta m^2_{\rm atm} = 2.41^{+0.09}_{-0.10} \times 10^{-3}\,, \quad \sin^2\theta_{\rm atm} = 0.950^{+0.035}_{-0.036}.$$

The discovery of non-zero θ_{13}

The apparent decoupling of the solar and atmospheric oscillations suggested a small or zero value for the mixing parameter $|U_{e3}| = \theta_{13}$, in strong contrast with the apparent maximal atmospheric mixing ($\sin^2 2\theta_{23} \sim 1$) and large solar mixing ($\sin^2 2\theta_{12} \sim 0.8$). Since CP violation requires all three mixing angles to be non-zero, the investigation of the 1-3 mixing sector was a prerequisite for the study of the CP-violating nature of neutrino mixing. An intense experimental effort towards the measurement of θ_{13} was thus undertaken in the 2000s.

The most direct way to measure θ_{13} is through the observation of $\nu_\mu \leftrightarrow \nu_e$ oscillations. One option is to measure electron (anti)neutrino disappearance in the atmospheric L/E range, which is ruled at first order by the effective ν_e survival probability

$$P^{\rm eff}_{\nu_e \to \nu_e} \simeq 1 - \sin^2 2\theta_{13} \sin^2 \frac{\Delta m^2_{31} L}{4E}$$

With a small mixing angle, the deficit of $\overline{\nu}_e$ events to be measured is small and thus requires both very large statistics and a good control of the flux and detection-related systematic uncertainties.

In the late 1990's the CHOOZ and Palo Verde experiments were the first to attempt the observation of reactor $\overline{\nu}_e$ disappearance at baselines of $\sim 1\,$km. The data was consistent with the no-oscillations hypothesis, providing the limit [53]

$$\sin^2 2\theta_{13} < 0.15 \text{ at } 90\% \text{ C.L.}$$

To go further, another generation of reactor experiments was designed, with increased statistics and multiple-baselines detectors:

- *Double Chooz* in France, on the same two-reactor site as the CHOOZ precursor, operated first with a single far detector (1050 m), then with its near detector (400 m) since 2015,
- *RENO* in South Korea, with two identical detectors placed at 300 m and 1400 m of a six-reactor power plant,
- *Daya Bay* in China, with eight identical detectors within 1.9 km of two closeby power plants with a total of six reactors.

All these experiments operate detectors based on the same broad principle as the Reines-Cowan historical experiment: the delayed coincidence signal of $\overline{\nu}_e$-induced inverse beta decays are detected in large ($\sim 10\,$ton) liquid scintillator tanks. About 10^{2-3} candidate IBD events are detected per day.

The 1-3 mixing can also be probed at accelerator experiments, via the observation of $\nu_\mu \to \nu_e$ appearance with the effective probability

$$P^{\rm eff}_{\nu_\mu \to \nu_e} \simeq \sin^2\theta_{23} \sin^2 2\theta_{13} \sin^2 \frac{\Delta m^2_{31} L}{4E}$$

In that case the detection of a statistically significant excess of ν_e events in a ν_μ beam would be sufficient to achieve the discovery. The intrinsic ν_e contamination of the beam thus needed to be low and accurately monitored, and the possible misidentification of background events as ν_e well understood.

To this end, a follow-up of K2K using Super-Kamiokande as far detector with a baseline of 295 km started operating in 2010: the *Tokai-to-Kamiokande* experiment (T2K). The design of T2K was optimised for the observation of $\nu_\mu \to \nu_e$ appearance. A description of the experiment

can be found in Ref. [54]. One peculiarity is that the axis of the neutrino beam, produced at the J-PARC accelerator facility, is directed 2.5° off the direction of the far detector. This off-axis configuration takes advantage of pion decay kinematics so that neutrinos reaching the far detector have a narrow energy spectrum peaked at 0.6 GeV, maximum of the ν_e appearance probability. This suppresses very effectively the higher-energy part of the spectrum, which would otherwise induce a non-oscillating background. Two near detectors, one off-axis and one on-axis, are installed 280 m downstream of the J-PARC beamline.

The "race" for the last mixing angle found its conclusion in 2011 and 2012. θ_{13} turned out to be "relatively large", in the sense that it was measured slightly below the most stringent limit set by CHOOZ. The combined analysis of atmospheric, solar and long-baseline reactor data provided first hints for $\theta_{13} \neq 0$ in 2008 [55]. In July 2011, T2K reported an indication for beam $\nu_\mu \to \nu_e$ appearance, excluding $\theta_{13} = 0$ at 90% CL [56]. MINOS followed closely with a similar significance [57]. In the same year, Double Chooz reported the first indication for reactor $\overline{\nu}_e$ disappearance, also excluding $\theta_{13} = 0$ at 90% C.L. [58]. The 5σ discovery was finally achieved by Daya Bay in March 2012 [59], and shortly after confirmed by RENO [60].

In 2013, T2K provided 3σ evidence for $\nu_\mu \to \nu_e$ oscillations [61], observing 11 candidate ν_e events for an expected $3.3 \pm 0.4\,(\text{syst.})$ background events. It is remarkable that θ_{13} is now estimated with a relative precision of less than 2% in the combined analysis of global oscillation data [62, 63]:

$$\theta_{13} = (8.54 \pm 0.15)^\circ \quad \text{(assuming normal mass ordering)}$$

The precision of the measurement mainly comes from the reactor experiments, in particular Daya Bay.

Latest oscillation experiments

The $\nu_\mu \leftrightarrow \nu_\tau$ oscillation channel has been studied in 'appearance mode' by the long-baseline OPERA[22] experiment, operating from 2008 to 2012. Neutrinos produced at CERN by the CNGS[23] ν_μ beam were detected at the Gran Sasso underground laboratory, with a baseline of 730 km. The far detector, a segmented tracker based on the emulsion cloud chamber (ECC) technology, was optimised for the identification of τ leptons produced by ν_τ CC interactions. OPERA reported the discovery of $\nu_\mu \to \nu_\tau$ appearance in 2016, with the observation of 5 ν_τ candidate events (out of 19505 neutrino interactions in the target fiducial volume) for a background expectation of 0.25 ± 0.05 events [64].

Following MINOS at the Fermilab NuMI beam, the NoνA long-baseline experiment started operating in 2014. Similarly to T2K, NOνA is optimised for observing $\nu_\mu \to \nu_e$ appearance and operates off-axis (0.8°) with a narrow-band energy spectrum centered on 1.9 GeV. Its longer baseline of 810 km is an asset for measuring matter-induced oscillations, thus increasing the sensitivity to the mass hierarchy and to δ_{CP}. Its near and far detectors are functionnally identical tracking calorimeters, with a total target mass of 14 kton at the far detector.

As discussed in next chapter (Sec. 2.2), while the most stringent constraints on the atmospheric oscillaton parameters (θ_{23}, Δm^2_{atm}) are currently achieved by beam experiments (T2K and NOνA, MINOS), atmospheric Cherenkov detectors (Super-Kamiokande and very large volume neutrino telescopes) provide very competitive measurements.

1.3.2 Current knowledge and remaining unknowns

The current knowledge of oscillation parameters is summarised in Tab. 1.1, and as a way to recapitulate the overview of the experimental field given in the previous sections, the data

[22] Oscillation Project with Emulsion-tRacking Apparatus
[23] CERN neutrinos to Gran Sasso

contributing the most to the determination of each parameter are listed in Tab. 1.2. The parameters in Tab. 1.1 are determined by a combined analysis of most of the available oscillation data as of January 2018 [62, 63]. Other similar global fits can be found in Refs. [65, 66].

The discovery of $\theta_{13} \neq 0$, establishing that all mixing elements of the PMNS matrix are non-zero, completed a first phase of discoveries in neutrino oscillation physics. To go forward, the new generation of experiments must be sensitive to subleading effects in the full three neutrino-oscillation picture. $\theta_{13} \neq 0$ leaves open the possibility for physical CP violation via the δ_{CP} phase, and the mass hierarchy (ordering) is the other major remaining unknown accessible via neutrino oscillations. With the impressive precision on θ_{13} provided by recent reactor experiments, the least well constrained mixing angle is θ_{23}. In particular, whether $\theta_{23} = \pi/4$ (*maximal mixing*), $\theta_{23} < \pi/4$ (*lower octant*), or $\theta_{23} > \pi/4$ (*upper octant*) is of theoretical interest.

In the following, the theoretical and experimental implications of these three issues are introduced, before discussing the paths towards their determination in the next section.

Octant of θ_{23} The observed pattern of neutrino mixing is peculiar, with one small and two large mixing angles, and the possibility that the $2-3$ mixing be maximal. A number of theoretical ideas have been proposed to explain this pattern (see *e.g.* Ref. [68] for a review). Resolving the θ_{23} octant (if not maximal) would rule out some of these models, and more generally lead to a better understanding of neutrino masses and mixing.

Mass hierarchy The determination of the NMH is of importance both from the theoretical point of view and for its connections with other experimental efforts in neutrino physics. In terms of fundamental implications, the NMH generally relates to theories attempting to explain the pattern of fermion masses. It would be a clear discriminator for some neutrino mass models (for instance based on flavour symmetry) which predict a specific hierarchy. On the experimental side, the NMH has a strong impact on next-generation experiments aiming towards the measurement of δ_{CP}. It also has implications for both direct and indirect neutrino mass measurements, as well as tests of the Dirac or Majorana nature of the neutrino via neutrinoless double beta decay searches. Further discussion of the importance of the mass hierarchy can be found in Ref. [69, 70] and references therein.

CP violation (CPV) Similarly to θ_{23} and the NMH, the value of δ_{CP} has important implications for theoretical models of neutrino mixing and flavour. CPV in the weak interactions of quarks was discovered as early as 1964, and it is accounted for in the SM via the δ_{CP} phase of quark mixing (CKM matrix). It is natural to investigate the possibility of an analogous in the lepton sector, all the more so since it may help understand the striking discrepancy between the mixing patterns of quarks and leptons. In fact, the motivation for the hunt for leptonic CPV goes beyond particle physics: CP violation in general is related to the question of matter-antimatter asymmetry in the Universe via baryogenesis mechanisms[24] [71].

Finally, the absolute mass scale and Dirac or Majorana nature are crucial properties of neutrinos that cannot be determined via flavour oscillations. Direct measurements of the absolute neutrino massses, for instance at the *Karlsruhe Tritium Neutrino* (KATRIN) experiment [72], are based on the kinematics of β-decay: the imprint of neutrino masses can be observed as a distorsion of the β-spectrum close to its endpoint. Some cosmological observables are also indirectly sensitive to the sum of neutrino masses via their gravitational effect [73]. Combination of cosmological probes currently provide the most stringent upper limits on neutrino masses,

[24] The existence of CP violation is required by baryogenesis (Sakharov criteria). The CKM phase as the single source of CP violation in baryogenesis has been ruled out. In a model of baryogenesis called *thermal leptogenesis*, the CP violation generated in neutrino mixing might be sufficient to generate a net baryon asymmetry, although it is unclear whether this hypothesis can be tested experimentally. See Ref. [71] for a recent review of the subject.

	Normal Hierarchy (best fit)		Inverted Hierarchy ($\Delta\chi^2 = 4.14$)		Any Hierarchy
	bfp $\pm 1\sigma$	3σ range	bfp $\pm 1\sigma$	3σ range	3σ range
$\sin^2\theta_{12}$	$0.307^{+0.013}_{-0.012}$	$0.272 \to 0.346$	$0.307^{+0.013}_{-0.012}$	$0.272 \to 0.346$	$0.272 \to 0.346$
$\theta_{12}/^\circ$	$33.62^{+0.78}_{-0.76}$	$31.42 \to 36.05$	$33.62^{+0.78}_{-0.76}$	$31.43 \to 36.06$	$31.42 \to 36.05$
$\sin^2\theta_{23}$	$0.538^{+0.033}_{-0.069}$	$0.418 \to 0.613$	$0.554^{+0.023}_{-0.033}$	$0.435 \to 0.616$	$0.418 \to 0.613$
$\theta_{23}/^\circ$	$47.2^{+1.9}_{-3.9}$	$40.3 \to 51.5$	$48.1^{+1.4}_{-1.9}$	$41.3 \to 51.7$	$40.3 \to 51.5$
$\sin^2\theta_{13}$	$0.02206^{+0.00075}_{-0.00075}$	$0.01981 \to 0.02436$	$0.02227^{+0.00074}_{-0.00074}$	$0.02006 \to 0.02452$	$0.01981 \to 0.02436$
$\theta_{13}/^\circ$	$8.54^{+0.15}_{-0.15}$	$8.09 \to 8.98$	$8.58^{+0.14}_{-0.14}$	$8.14 \to 9.01$	$8.09 \to 8.98$
$\delta_{\rm CP}/^\circ$	234^{+43}_{-31}	$144 \to 374$	278^{+26}_{-29}	$192 \to 354$	$144 \to 374$
$\frac{\Delta m^2_{21}}{10^{-5}\,\mathrm{eV}^2}$	$7.40^{+0.21}_{-0.20}$	$6.80 \to 8.02$	$7.40^{+0.21}_{-0.20}$	$6.80 \to 8.02$	$6.80 \to 8.02$
$\frac{\Delta m^2_{3\ell}}{10^{-3}\,\mathrm{eV}^2}$	$+2.494^{+0.033}_{-0.031}$	$+2.399 \to +2.593$	$-2.465^{+0.032}_{-0.031}$	$-2.562 \to -2.369$	$\begin{bmatrix} +2.399 \to +2.593 \\ -2.536 \to -2.395 \end{bmatrix}$

Table 1.1: Three-flavor oscillation parameters as extracted by a global fit of the experimental data released as of January 2018 ('NuFit 3.2'), available at Ref. [63] (see also Ref. [62]). The best-fit parameters (bfp) and errors (at 1σ, 3σ) are reported. Numbers in the first (second) column are obtained assuming NH (IH), *i.e.* relative to the respective local $\Delta\chi^2$ minimum, whereas the third column corresponds to a global minimisation. Note that $\Delta m^2_{3\ell} \equiv \Delta m^2_{31} > 0$ for NH and $\Delta m^2_{3\ell} \equiv \Delta m^2_{32} < 0$ for IH. As indicated in the top row, the normal hierarchy hypothesis is favoured by the fit at the level of $\Delta\chi^2 = 4.14$ (the interpretation in terms of confidence level is discussed in Chap. 6). Table adapted from Refs. [62, 63].

Experiment	Dominant	Important
Solar Experiments	θ_{12}	Δm^2_{21}, θ_{13}
Reactor LBL (KamLAND)	Δm^2_{21}	θ_{12}, θ_{13}
Reactor MBL (Daya-Bay, Reno, D-Chooz)	θ_{13}	$\lvert\Delta m^2_{3\ell}\rvert$
Atmospheric Experiments	θ_{23}	$\lvert\Delta m^2_{3\ell}\rvert$, θ_{13}, $\delta_{\rm CP}$
Accelerator LBL $\nu_\mu \to \nu_\mu$ Disapp (Minos, NOνA, T2K)	$\lvert\Delta m^2_{3\ell}\rvert$, θ_{23}	
Accelerator LBL $\nu_\mu \to \nu_e$ (Minos, NOνA, T2K)	$\delta_{\rm CP}$	θ_{13}, θ_{23}, sign$(\Delta m^2_{3\ell})$

Table 1.2: Experiments contributing to the present determination of the oscillation parameters. The acronyms LBL and MBL refer to medium and long baseline experiments, respectively. Table adapted from Ref. [67].

though in a model-dependent manner. Assuming the standard cosmological model, the data from the Planck space observatory alone sets the limit [74, 75]:

$$\sum_{m_\nu} < 0.23\,\mathrm{eV} \quad (95\%\,\mathrm{C.L.}).$$

The Dirac or Majorana nature of the neutrino is tested by searching for lepton number violating processes (forbidden in the SM), in particular neutrinoless double β decay ($0\nu\beta\beta$). A possible diagram for the $0\nu\beta\beta$ process via the exchange of a Majorana neutrino is represented on the adjacent figure 1.13. Although the process could be mediated by other particles, its existence requires Majorana neutrino mass [44]. In case $0\nu\beta\beta$ is observed and is assumed to be mediated by Majorana neutrinos, its decay rate may be translated into an estimation of the absolute neutrino mass scale, and even provide indications on the mass hierarchy if combined with other neutrino mass measurements. Conversely, establishing unambiguously that Majorana neutrinos mediate $0\nu\beta\beta$ would require the external input of direct neutrino mass probes and the measurement of the NMH from oscillations [76]. A recent experimental review of $0\nu\beta\beta$ searches can be found in Ref. [77].

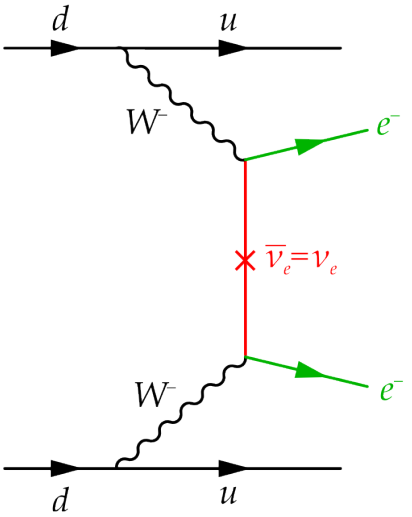

Figure 1.13: Feynman diagram for a neutrinoless double beta decay mediated by a Majorana neutrino. Taken from Ref. [78].

1.3.3 Towards the mass hierarchy and CP violation

In oscillation experiments, the main strategy for the determination of the NMH focuses on matter-induced oscillations in the atmospheric L/E range, probing both the ν_μ survival and $\nu_\mu \leftrightarrow \nu_e$ appearance channels. To maximise the impact of matter effects long oscillation baselines are required. On one hand, current accelerator experiments (NOνA, and to a lesser extent T2K) have some NMH sensitivity via the $\nu_\mu \rightarrow \nu_e$ appearance channel, accounting for the reactor constraints on θ_{13} (see for instance Ref. [79]). However, due to their moderate baselines they are affected by degeneracies between the NMH and δ_{CP}. Atmospheric neutrino detectors, on the other hand, benefit from much longer baselines and thus stronger matter effects, so that both the $\nu_\mu \rightarrow \nu_\mu$ and $\nu_\mu \leftrightarrow \nu_e$ channels become sensitive to the NMH. Some of the oscillation parameter degeneracies are reduced, thanks to the more intense matter resonance and the wide range of oscillation baselines available ($\sim 50 - 12800\,\text{km}$). The composition of the atmospheric flux as a mixture of ν_μ, ν_e, $\overline{\nu}_\mu$, $\overline{\nu}_e$ is a downside, since most detectors cannot distinguish ν from $\overline{\nu}$. This is however compensated by the difference in cross-sections ($\sigma_{\text{CC}}(\nu)/\sigma_{\text{CC}}(\overline{\nu}) \sim 2$) and the very large event statistics that can be accumulated in such detectors. While the latest global fits to neutrino oscillation data tend to find a mild preference for the normal hierarchy (see for instance Tab. 1.1), the definitive discovery will require a high confidence level determination of the NMH by a single experiment. To this end, high-statistics atmospheric detectors may be the best option in the next decade, before the next generation of beam experiments becomes operational.

The second possibility to measure the NMH is radically different and complementary, as it is independent of matter effects. It arises from the coexistence of oscillations with very close frequencies driven by Δm^2_{32} and Δm^2_{31}. Their combination in disappearance probabilities produces interference effects which are sensitive to the NMH when observed over many oscillation maxima. The JUNO experiment, currently under construction in China, implements this method by measuring reactor $\overline{\nu}_e$ disappearance at a baseline of about 50 km [80]. The main challenge of this measurement is the extreme requirement in energy resolution (about 3%).

Finally, a new generation of long-baseline experiments is being designed, dedicated to the discovery of CP violation and subsequent precise measurement of δ_{CP}. The main such projects are DUNE (*Deep Underground Neutrino Experiment*) planned to start operating in ~ 2026 with a baseline of $\sim 1300\,\text{km}$ in the USA, and T2HK in Japan (*Tokai-to-Hyper-Kamiokande*).

The atmospheric option for the NMH measurement is being implemented in the Mediterranean Sea by ORCA (*Oscillation Research with Cosmics in the Abyss*), the low-energy branch of the KM3NeT neutrino telescope. A similar project is underway as a low-energy upgrade of the IceCube neutrino telescope at the South Pole, called PINGU (*Precision IceCube Next Generation Upgrade*). Chap. 2 will discuss large volume atmospheric detectors and the phenomenology of Earth-crossing atmospheric neutrino oscillations; Chap. 3 will describe the ORCA experiment in details.

Chapter 2

Physics with atmospheric neutrinos

Contents

This chapter is an introduction to the experimental study of the flavour oscillations of Earth-crossing atmospheric neutrinos with a very large-volume Cherenkov detector, with the objectives of i) determining the neutrino mass hierarchy and measuring atmospheric oscillation parameters and ii) learning about the physical and chemical properties of the deep Earth by measuring the electron density in its inner layers. Sec. 2.1 first presents the flux of atmospheric neutrinos, generated as a byproduct of the interaction of cosmic rays in Earth's atmosphere. Some characteristics of the interactions of neutrinos in the $1-100\,\mathrm{GeV}$ energy range are then examined, before adressing Cherenkov detection in neutrino telescopes and the event topologies associated to each neutrino flavour and interaction channel. Sec. 2.2 adresses the phenomology of matter-enhanced flavour oscillations of neutrinos traversing the earth, with an emphasis on the measurement of the neutrino mass hierarchy. Finally, in Sec. 2.3 the bases of the current knowledge of Earth's structure are depicted, with a focus on the unknowns that could potentially be adressed with the help of neutrinos. A review of the general principles of neutrino tomography is then given, before discussing in more detail the prospects for oscillation tomography of the deep Earth using atmospheric neutrinos.

2.1 Atmospheric neutrinos and their detection

2.1.1 Flux of atmospheric neutrinos

The next generation of atmospheric neutrino oscillation studies must go beyond the simple description of $\nu_\mu \to \nu_\tau$ disappearance as the dominant process and study subleading effects, including matter-enhanced $\nu_\mu \leftrightarrow \nu_e$ or $\overline{\nu}_\mu \leftrightarrow \overline{\nu}_e$ depending on the mass hierarchy. To this end, a precise understanding of the atmospheric neutrino flux is crucial. In particular, the energy and zenith angle dependence of the flavour ratio $(\Phi_{\nu_e} + \Phi_{\overline{\nu}_e})/(\Phi_{\nu_\mu} + \Phi_{\overline{\nu}_\mu})$ and neutrino-to-antineutrino ratios $(\Phi_{\nu_e} + \Phi_{\nu_\mu})/(\Phi_{\overline{\nu}_e} + \Phi_{\overline{\nu}_\mu})$, $\Phi_{\nu_e}/\Phi_{\overline{\nu}_e}$, and $\Phi_{\nu_\mu}/\Phi_{\overline{\nu}_\mu}$ are of importance, since their mismodeling would bias the interpretation of the measured fluxes in terms of neutrino oscillation patterns.

This section describes the main production processes and the most important variabilities affecting these characteristics, focusing on the energy range of interest for ORCA ($\sim 1-100\,\mathrm{GeV}$). The numerical simulations and experimental measurements of the atmospheric neutrino flux are introduced, with a discussion of the associated uncertainties. The main references for this section are [1, 81–83]. Unless otherwise specified, the model of flux used in the figures is the one used as input for the analyses presented in the rest of the thesis, corresponding to the recent simulations of Honda *et al.* [84, 85].

As introduced in Sec. 1.1.2, atmospheric neutrinos are produced in the chains of decays of the charged mesons produced in air showers initiated by the impact of cosmic rays with nuclei in the atmosphere. Below 100 GeV, the main contribution comes from $\pi^\pm$ decays, but the contribution of $K^\pm$ decays increases progressively with energy. For the most part, neutrino production occurs in the lower stratosphere ($\sim 10-30\,\mathrm{km}$). In first approximation:

- The energy dependence of the flux of atmospheric neutrinos is similar to the one of the primary CR flux, *i.e.* it follows a power law

$$\Phi \sim \frac{\mathrm{d}N}{\mathrm{d}E} \sim E^{-\alpha}$$

  with a differential spectral index $\alpha \simeq 3$ (see Fig. 2.1).
- The electron-to-muon flavour ratio is about 1 : 2, due to the dominant chain of production processes formed by the charged meson decay

$$\pi/K^+ \to \mu^+ + \nu_\mu \tag{2.1}$$

  followed by the in-flight muon decay

$$\mu^+ \to e^+ + \nu_e + \overline{\nu}_\mu, \tag{2.2}$$

  the kinematics of $\pi^\pm$ and $\mu^\pm$ decay is such that each neutrino in the chain carries on average roughly the same energy.

Tau neutrinos are produced in negligible quantities, the tau lepton being heavier than $\pi/K^\pm$. The chain of processes starting with negatively charged mesons is obtained by replacing particles with their antiparticles in (2.1, 2.2). In terms of neutrino production, they can be summarised as

$$\pi/K^+ \rightrightarrows \nu_\mu + \overline{\nu}_\mu + \nu_e \tag{2.3}$$
$$\pi/K^- \rightrightarrows \overline{\nu}_\mu + \nu_\mu + \overline{\nu}_e \tag{2.4}$$

The neutrino-to-antineutrio ratios thus follow

$$\frac{\Phi_{\nu_\mu}}{\Phi_{\overline{\nu}_\mu}} \simeq 1, \quad \frac{\Phi_{\nu_e}}{\Phi_{\overline{\nu}_e}} \simeq \frac{\Phi_{\mu^+}}{\Phi_{\mu^-}}. \tag{2.5}$$

These $\nu/\overline{\nu}$ proportions as well as the $1:2$ flavour ratio are mostly valid at low energy ($E_\nu \lesssim 1\,\mathrm{GeV}$) when most muons decay before reaching the ground and kaon decays contribute negligibly to neutrino production.

The most important effect which modulates this first order picture comes from the variability of the contribution of in-flight muon decays (2.4) with energy and zenith angle. Muons are relatively long-lived, with a mean lifetime of about $2.2 \cdot 10^{-6}\,\mathrm{s}$. It is a well-known consequence of relativistic time dilation that a significant fraction of energetic muons reach the ground before decaying: a production height of 15 km corresponds to the decay length of 2.4 GeV muons. Besides, few-GeV muons quickly lose their energy in the atmosphere, thus producing neutrinos of lower energies once they decay. This effect also introduces a zenith dependence of the flux, since the muon path length in the atmosphere is larger for trajectories close to horizontal.

The energy and angular dependence of muons and neutrinos from the decays of pions and kaons (2.1) is determined by the competition between decay and reinteraction in the development of hadronic showers. At low energy the reinteraction is favoured for kaons, thus pion production dominates. This is no longer true in the higher-energy part of the spectrum, all the more so since kaon decays produce higher-energetic neutrinos. In addition, the approximately exponential decrease of the atmospheric density with altitude results in a slight hardening of the neutrino energy spectrum for showers developing close to horizontally. In contrast, above $\sim 1\,\mathrm{GeV}$ the effect of the transverse momentum of the secondaries in hadronic showers can be neglected in good approximation.

Flavour ratio

The overall flavour ratio in the energy range of interest is shown as a function of energy and for selected ranges of zenith angle in Fig. 2.1.

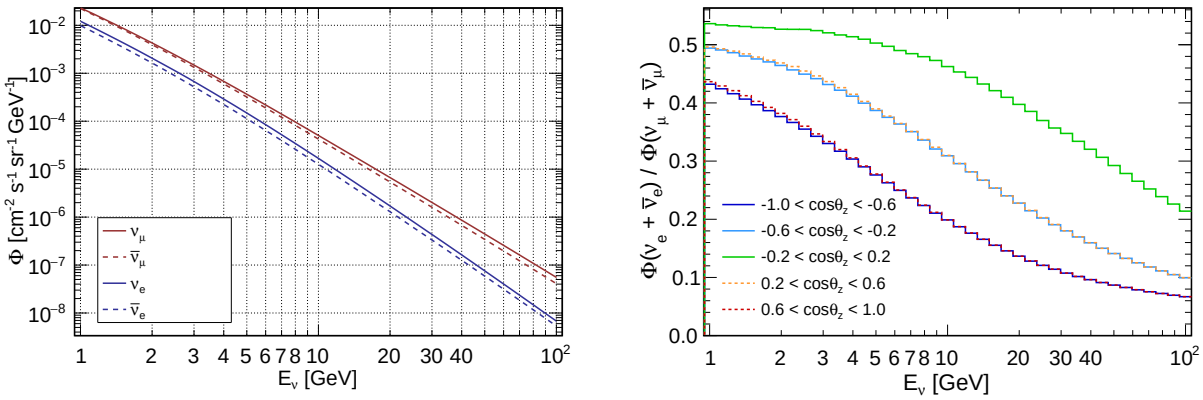

Figure 2.1: **Left:** Differential flux of atmospheric neutrinos per flavour (without neutrino oscillations). The flux is averaged over all upgoing zenith directions ($\theta_z = \pi/2 \to \pi$) and over the azimuth direction ($\varphi = 0 \to 2\pi$). **Right:** $(\nu_e + \overline{\nu}_e)/(\nu_\mu + \overline{\nu}_\mu)$ flux ratio averaged over five distinct bins of equal width of $\cos\theta_z$.

The effect mentioned above becomes more important at high energies, and affects the ν_e flux more strongly since ν_e are produced mostly via muon decays. As seen on Fig. 2.1, this results in a steepening of the energy spectrum of ν_e relatively to ν_μ, so that the e/μ flavour ratio decreases with neutrino energy, and is the lowest for close to horizontal directions.

Another effect comes into play in the higher energy part of the spectrum. While (2.1) is an almost exclusive decay channel for $\pi^{\pm}$ (99.9%), the branching ratio (BR) of this channel for charged kaons $K^{\pm}$ is only 63.6%. The other decays producing neutrinos are semileptonic:

$$K^+ \to \pi^0 + \mu^+ + \nu_\mu \quad (3.4\%) \tag{2.6}$$
$$K^+ \to \pi^0 + e^+ + \nu_e \quad (5.1\%) \tag{2.7}$$

In the case of (2.7), the chain ends, only producing one ν_e overall. The flavour ratio thus depends on the ratio of $\pi^{\pm}$ to $K^{\pm}$ production in the hadronic processes, which is poorly constrained in the energy range of interest.

Neutrino-to-antineutrino ratio

The $\nu_e/\overline{\nu}_e$ and $\nu_\mu/\overline{\nu}_\mu$ ratios are shown on Fig. 2.2. As a consequence of the overabundance of matter over antimatter in the cosmic ray primaries and secondaries, the $\nu_e/\overline{\nu}_e$ ratio is roughly equal to 1.3, with a moderate dependence on both energy and zenith angle. On the contrary, the $\nu_\mu/\overline{\nu}_\mu$ ratio shows a clear increase with energy and a zenith angle dependance, the predominance of ν_μ over $\overline{\nu}_\mu$ being stronger close to vertical. This is again related to the decrease of the contribution of muon decays, which are responsible for washing out the μ^-/μ^+ asymmetry.

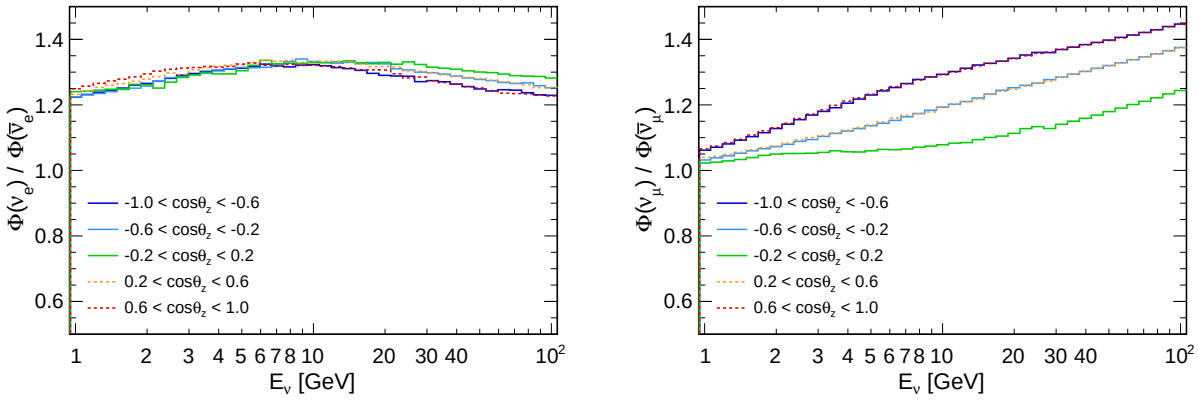

Figure 2.2: $\nu_e/\overline{\nu}_e$ and $\nu_\mu/\overline{\nu}_\mu$ flux ratios as a function of energy, averaged over five distinct bins of $\cos\theta_z$.

Effect of the geomagnetic field

The angular distribution of primary CR interacting in the atmosphere would be close to isotropic, were it not for the shielding effect of the geomagnetic field. Low-energetic primaries are filtered outside the atmosphere, most significantly for energies below $\sim 10\,\text{GeV}$ in the case of primary protons. More generally, the gyroradius of charged particles depends on their *rigidity*, *i.e.* the ratio of total momentum to total charge. Therefore precise calculations must account for the nuclear composition of the primary CR. The minimal rigidity for a primary CR to be able to reach Earth's surface (*geomagnetic cutoff*) strongly depends on the location: it is vanishingly small at the poles (almost no shielding) and maximal close to the equator. This results in a rather large variability of the flux among detector locations, affecting in particular its zenith dependence. A slight azimuthal dependence of the neutrino flux originates from the classical East-West geomagnetic effect (cutoffs are higher for particles traveling towards magnetic West than East).

The second effect of the geomagnetic field is the bending of the trajectories of low-energy charged secondaries in air showers. The effect is significant for muons, which are typically deviated by $\sim 3°$ from the direction of the parent meson before decaying. The deviation is

energy-independent, because the gyroradius and decay length have opposite dependence with energy [82]. However the charge-sign dependence of the muon bending affect neutrino and antineutrino fluxes differently [83, 86].

Monte Carlo simulations and uncertainties

Most of the essential features of the atmospheric neutrino flux were discussed in the seminal papers of the early 1960's [87, 88]. While analytical models are useful to understand the main mechanisms at play [1], these calculations quickly become untractable when the number of effects accounted for increases. Several Monte Carlo based numerical simulations have thus been developed and refined throughout the years [82, 84, 89]. Detailed models of the primary CR energy spectrum and composition are used as inputs, and the geomagnetic filtering is accounted for using particle back-tracing techniques and a precise model of Earth's magnetic field. A state-of-the-art location and seasonal-dependent model of the atmospheric density profile is also used in Ref. [84]. The development of air showers is followed step-by-step, including detailed particle energy loss and three-dimensional effects due to transverse momentum in hadron interactions and decays, and geomagnetic bending of charged secondaries. In this thesis, the most recent version provided by the Honda group [84, 85] has been used.

An extensive discussion of the uncertainties related to atmospheric flux calculations and numerical simulations can be found in Ref. [86]. The main source of uncertainty are the hadronic interaction models, which are challenging both theoretically and experimentally. In particular, constraints on nucleon-nucleus and nucleon-nucleon cross-sections from beam experiments are sparse in the range of parameters relevant for air showers. The second main source of uncertainty is the parametrisation of the primary CR flux. Even though this flux has been measured over a large energy range by numerous atmospheric balloon-borne as well as satellite experiments, large uncertainties remain, especially on the total normalisation. Atmospheric flux simulation and uncertainties have more recently been discussed in Refs. [90], making used of a direct numerical solving method of hadronic cascade equations [91, 92] which allows to propagate arbitrary hadronic model uncertainties to the predicted neutrino flux.

2.1.2 Neutrino interactions in the 1–100 GeV range

The cross-sections for neutrino interactions enter in the prediction of the observed event rates in all neutrino experiments. They are a crucial ingredient of neutrino oscillation analyses, and constitute a whole research subject[1]. A comprehensive review can be found in Ref. [93]. This section gives a rapid overview of neutrino-nucleon interactions, which are the most relevant in the $1-100\,\mathrm{GeV}$ energy range probed by ORCA. The overall structure and content of this section are partly inspired by other PhD theses [94, 95].

The scattering of neutrinos and fermions is described by the weak interaction, as seen in Chap. 1 (see also the elementary introduction in Appendix A.1). For the purpose of neutrino detection, neutrino scattering with electrons and nuclei must in general be considered. The scattering occurs either via the weak charged current (exchange of a $W^{\pm}$ boson) or the weak neutral current (exchange of a Z^0 boson). Fig. 2.3 represents a generic neutrino interaction and defines the notations for the four-momenta of incoming and outgoing particles, as well as the four-momentum exchange q.

[1] It is worth mentioning that neutrino scattering processes with fermions (electrons, nucleons) have been used in the past to probe various aspects of electroweak and nuclear physics, as well as the early confirmation of QCD and the quark parton model. See Ref. [96] for a review. The interest for precise cross-section measurements was increased due to the need for such information in the study of neutrino oscillations.

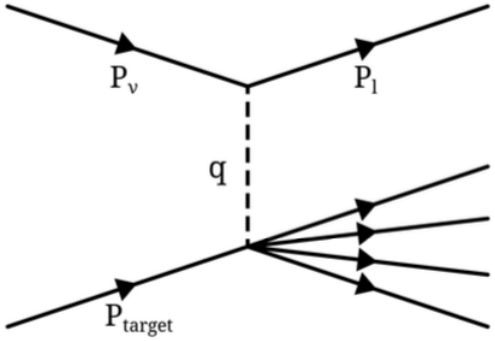

Figure 2.3: A generic neutrino interaction with a target, generating a (charged or neutral) lepton and an unspecified system of other outgoing particles. Taken from Ref. [94].

Neutrinos with energies in the range $\sim 1-100\,\mathrm{GeV}$ interact mostly via scattering on nuclei. Depending on the range of four-momentum transfer

$$Q^2 = -q^2 = -(p_\nu - p_l)^2\,, \tag{2.8}$$

the latter can be described considering as target the whole nucleus ($Q^2 \ll 1\,\mathrm{GeV}$), a nucleon ($Q^2 \sim 1\,\mathrm{GeV}$) or a parton ($Q^2 \gg 1\,\mathrm{GeV}$).

In charged current (CC) interactions, the outgoing lepton l is a charged lepton with the same flavour α as the incoming neutrino ν_α, while in neutral current (NC) interactions the outgoing lepton is a neutrino ν_α. Thus, CC interactions are sensitive to the neutrino flavour and can be used to probe neutrino flavour oscillations, while NC interactions form a background of flavour-insensitive events. The CC interaction is maximally parity-violating and involves only left-handed fermions and right-handed antifermions, as summarised by the diagrams of Fig. 2.4 – considering for the target fermion a quark or antiquark of the first generation.

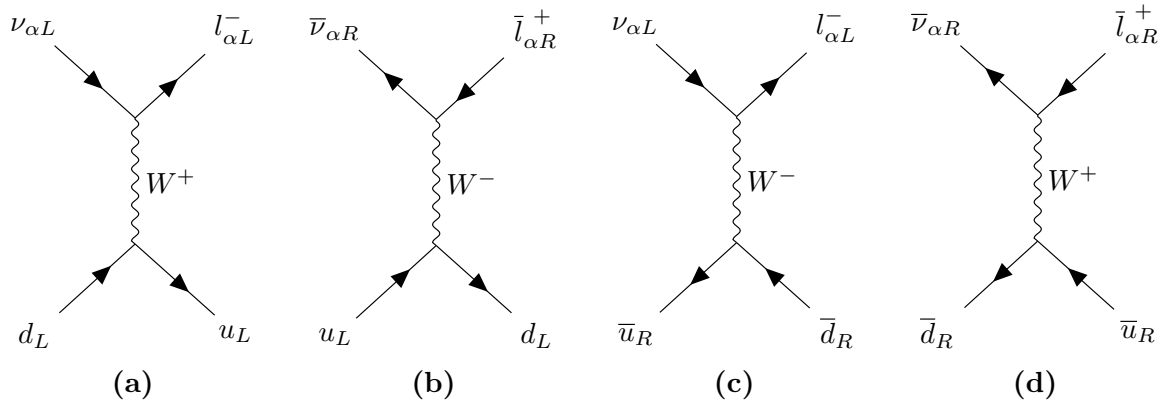

Figure 2.4: CC interactions of (anti)neutrinos and (anti)quarks. Note that interactions with antiquarks become relevant at high Q^2, where the neutrino can interact with non-valence quarks of the parton sea; likewise for quarks of the second generation.

In addition to the momentum transfer Q^2, other Lorentz variables are used to characterise neutrino interactions. Among them are the centre of mass energy of the system:

$$s = (p_\nu + p_{\mathrm{target}})^2 \tag{2.9}$$

and the inelasticity or *Bjorken-y* variable:

$$y = \frac{q \cdot p_{\text{target}}}{p_\nu \cdot p_{\text{target}}} \tag{2.10}$$

In the rest frame of the target fermion (laboratory frame), the inelasticity corresponds to the fraction of neutrino energy transferred to the target system:

$$y = \frac{E_\nu - E_l}{E_\nu}. \tag{2.11}$$

Very elastic CC interactions are those where most of the energy is transferred to the charged lepton. In such cases the energy, direction and flavour of the incoming neutrino are easier to infer, while inelastic interactions ($y \simeq 1$) are much less favourable.

The total neutrino-nucleus cross-sections is in general proportional to the number of nucleons, ignoring nuclear effects which become relevant at low Q^2. Fig. 2.5 shows the total cross-section per nucleon as a function of the neutrino energy, for the CC and NC interactions of neutrinos and antineutrinos of all flavours with a water target.

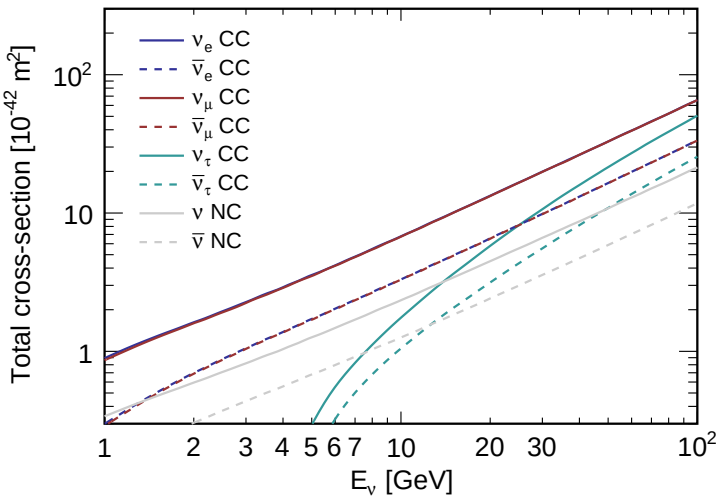

Figure 2.5: Inclusive cross-section per nucleon for the interaction of neutrinos with a water target, as predicted by the GENIE Monte Carlo generator [97, 98] (see Chap. 4 for details).

As a consequence of the universality of weak interactions (see Appendix A), the CC cross-sections are identical for all three neutrino flavours in the limit $E_\nu \to \infty$ where the effect of the outgoing lepton mass can be neglected. In this limit, cross-sections increase approximately linearly with energy. In practice, for $E_\nu > 1\,\text{GeV}$ the approximation $\sigma_{\nu_\mu} \simeq \sigma_{\nu_e}$ is good since $m_\mu \simeq 0.1\,\text{GeV}$. This does not hold for ν_τ due to the large mass of the τ lepton ($m_\tau = 1.78\,\text{GeV}$). The energy threshold for the quasi-elastic CC scattering $\nu_\tau + n \to \tau^- + p$ is

$$E_{\text{th}} = \frac{(m_p + m_\tau)^2}{2m_n} - \frac{m_n}{2} \simeq m_\tau \left(1 + \frac{m_\tau}{2m_n}\right) \simeq 3.45\,\text{GeV} \tag{2.12}$$

and the kinematic effect of the τ mass remains significant up to 100 GeV.

The total cross-sections are larger for neutrinos than for antineutrinos:

$$\frac{\sigma_{\text{CC}}(\nu)}{\sigma_{\text{CC}}(\overline{\nu})} \simeq 2, \quad \frac{\sigma_{\text{NC}}(\nu)}{\sigma_{\text{NC}}(\overline{\nu})} \simeq 3 \tag{2.13}$$

for an isoscalar target, *i.e.* a nucleon with equal number of protons and neutrons. This is due to the chiral-specific nature of the weak interaction (described previously and summarised on Fig. 2.4 for charged currents) and the fact that matter is predominantly[2] made of quarks over antiquarks.

The cross-sections shown on Fig. 2.5 are *inclusive*, *i.e.* they account for all neutrino interactions regardless of the final state and the exact mechanism involved. In the energy range of interest, three distinct mechanisms can be identified, and their contributions to the total cross-section are shown on Fig. 2.6:

Elastic and quasi-elastic[3] scattering (QE) This refers to the scattering of a neutrino off an entire nucleon, liberating one or multiple nucleons(s) from the target. Elastic and quasi-elastic scattering are dominant in the sub-GeV range, but remain important up to $\sim 20\,\text{GeV}$.

Resonant production (RES) This mechanism corresponds to the excitation of the nucleon to a baryonic resonance state, decaying to a variety of mesonic final states including nucleons and mesons (mainly single or multiple pions, and kaons). Resonant production starts to be significant for $E_\nu \sim 1\,\text{GeV}$ and dominates in the range 1 to 10 GeV.

Deep inelastic scattering (DIS) When neutrinos are sufficiently energetic to probe the nucleon at the quark level, the nucleon is effectively broken up and the outgoing quark and nucleon remnant *hadronise*, *i.e.* they produce a spray of hadrons, initiating a hadronic shower. DIS starts to dominate above 10 GeV and is the only relevant channel above $\sim 100\,\text{GeV}$.

From Fig. 2.6 it is clear that DIS is the most important channel for neutrino telescopes, while it is negligible when focusing on sub-GeV neutrinos. For the measurement of the NMH with atmospheric detectors the critical energy range is $\sim 2-8\,\text{GeV}$, where all three processes contribute significantly.

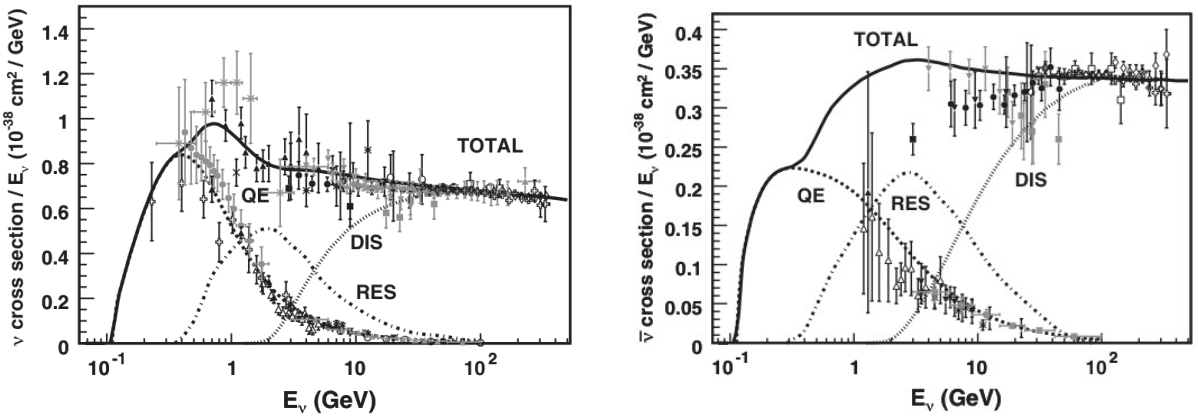

Figure 2.6: Measurements of the total CC cross-sections per nucleon for muon neutrinos (left) and antineutrinos (right) divided by neutrino energy E_ν and plotted as a function of E_ν. An isoscalar nuclear target ($Z = A - Z$) is assumed. Also shown are predictions of the NUANCE generator [99] for the various contributing processes: quasielastic scattering (dashed), resonance production (dot-dashed), and deep inelastic scattering (dotted). Taken from Ref. [93].

The larger total cross-section for neutrinos over antineutrinos is better understood as a consequence of the different behaviours of the angular differential cross-sections. In antineutrino-quark

[2] As a subdominant effect, neutrinos can interact in usual matter with antiquarks, which are present in nucleons as short-lived quark-antiquark pairs in the parton sea.

[3] The scattering is usually referred to as *quasi-elastic* for CC interactions (CCQE) due to the creation of the massive charged lepton, and *elastic* for NC interactions.

and neutrino-antiquark interactions (Fig. 2.4 (b), (c)), the conservation of angular momentum leads to a preference for forward scattering over backward scattering, while neutrino-quark and antineutrino-antiquark scattering (Fig. 2.4 (a), (d)) are isotropic. This can be translated as a dependence on the inelasticity y, which has the advantage to be observable from the final state energies. For the DIS cross-section [100]:

$$\frac{\mathrm{d}\sigma_{\nu}^{\mathrm{CC}}}{\mathrm{d}y} \simeq \left(-a_0 - a_1(1-y^2)\right) 10^{-38}\,\mathrm{cm}^2 \frac{E_\nu}{1\,\mathrm{GeV}} \tag{2.14}$$

$$\frac{\mathrm{d}\sigma_{\overline{\nu}}^{\mathrm{CC}}}{\mathrm{d}y} \simeq \left(-b_0 - b_1(1-y^2)\right) 10^{-38}\,\mathrm{cm}^2 \frac{E_\nu}{1\,\mathrm{GeV}} \tag{2.15}$$

with $a_0 = 0.72$, $a_1 = 0.06$, $b_0 = 0.09$ and $b_1 = 0.69$, the small coefficients a_1 and b_0 being due to the antiquark component in nucleons. Fig. 2.7 shows an early measurement of the Bjorken-y (inelasticity) differential cross-section of ν_μ-iron and $\overline{\nu}_\mu$-iron deep inelastic scattering at the CDHS experiment (*CERN, Dortmund, Heidelberg, Saclay*) [101].

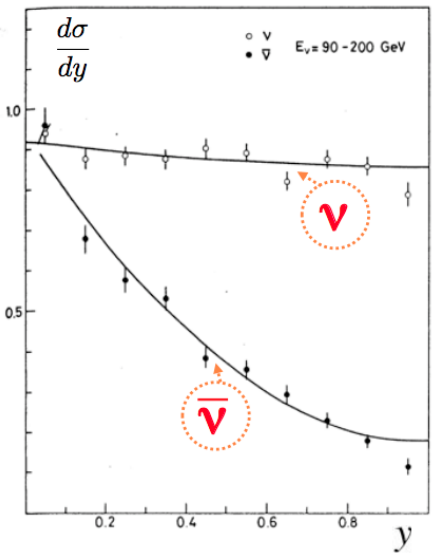

Figure 2.7: Measurement of the Bjorken-y (inelasticity) differential cross-section in ν_μ-iron DIS in the energy range $90 - 200\,\mathrm{GeV}$ (corresponding to kaon neutrinos) at the CDHS experiment [101]. The annotated figure was taken from [2].

Precision of cross-section predictions and measurements

As visible from Fig. 2.6, the experimental data is scattered, and at low energy statistical errors are large. The total cross-section is reasonably well described by predictions, but some discrepancies appear both in absolute normalisation and shape. The normalisation uncertainty at 10 GeV is estimated to be around 10%. Cross-sections are dominant sources of uncertainty in precision neutrino oscillation studies, and are especially crucial in next-generation long-baseline experiments. Significant research efforts, both theoretical and experimental, are being carried on. For recent dedicated experiments, see for instance Refs. [102–105]. Long-baseline oscillation experiments are also able to perform cross-section measurements at their near detectors [106]. On the theoretical side, precise models exist for each of the interaction channels (QE, RES, DIS), with input parameters that can be adjusted to fit the experimental data accurately; the description of the transition between different regimes is more problematic.

Both measurements and theoretical modeling are particularly challenging at low energies due to multiple entangled nuclear effects such as the impact of *Fermi motion* of the target nucleon inside the nucleus, *multinucleon* contributions to the neutrino interaction, or *final state interac-*

tions undergone by outgoing particles traversing the nucleus. Measurements and predictions in the DIS regime are less affected by nuclear effects because the contribution of the Fermi motion has a smaller relative impact, and because the precise number and kinematics of the outgoing particles are less important. Indeed most of the experiments in this energy regime, like neutrino telescopes, are not capable of reconstructing individual low energetic particles. The better agreement at high energy can be seen on Fig. 2.6.

2.1.3 Cherenkov detection with neutrino telescopes

The most efficient detection technique allowing to instrument very large volumes of target matter, as needed to observe the relatively small fluxes of atmospheric and extraterrestrial neutrinos, proceeds via the detection of Cherenkov light emitted in a transparent material by the ultrarelativistic secondary particles emerging from a neutrino interaction.

The theory of Cherenkov radiation was formulated in 1937 [107], and Cherenkov light emission in the atmosphere was first detected in 1953 [108]. The idea of installing photosensitive devices deep underwater to detect high-energy atmospheric neutrinos via the Cherenkov effect was proposed by Markov in 1960 [109, 110], only a few years after the detection of the neutrino. While Markov's proposal was motivated by the study of the fundamental properties of neutrinos, in the same year Greisen discussed the detection of neutrinos from cosmic sources like the Crab nebula [111]. He suggested to use a large ($\sim 3\,\text{kton}$) underground Cherenkov detector surrounded by photomultipliers, and enclosed in a shell of scintillating material to distinguish events induced by neutrinos and atmospheric muons.

As seen in Chap. 1, the pioneering large underground Cherenkov experiments, built in the 1980's with designs similar to Greisen's proposal, detected the first neutrino signal from an astrophysical source, the supernova SN1987A. The largest currently operating such detector is Super-Kamiokande (50 kton), which is only surpassed in instrumented volume by *neutrino telescopes*, built specifically to observe high-energy astrophysical neutrinos: ANTARES ($\sim 15\,\text{Mton}$), Baikal ($\sim 15\,\text{Mton}$[4]), and IceCube ($\sim 1\,\text{Gton}$) [19, 20, 22]. In underground water tanks, Cherenkov light is detected by a dense array of PMTs, arranged in a two-dimensional surface around the detection volume and facing inwards. In contrast, neutrino telescopes use a three-dimensional grid of PMTs, generally grouped in optical modules disposed in vertical structures spread out across the detector volume. Secondary charged particles from neutrino interactions emit Cherenkov light in characteristic patterns depending on their nature, energy, and trajectory. Neutrino events are thus reconstructed from the measured information on the number of emitted photons, their arrival times and their spatial distribution. Based on the technology developed for neutrino telescopes, more densely instrumented detectors like KM3NeT-ORCA in the Mediterranean will focus on neutrino oscillation physics. ORCA will lower the detection threshold down to $\sim 1\,\text{GeV}$, and thus accumulate samples of atmospheric neutrino events with unprecedented statistics in the $1-10$ GeV range, optimal for the study of matter-enhanced neutrino oscillations and the NMH measurement. In the following, the principle and most important characteristics of Cherenkov light emission are introduced, before discussing the specific event topologies in neutrino telescopes focusing on the GeV range such as ORCA.

Cherenkov radiation

Charged particles propagating in a dielectric medium with a velocity greater than the phase velocity of light in that medium emit Cherenkov radiation [107]. The mechanism of the Cherenkov effect is illustrated on Fig. 2.8. Charged particles polarise the surrounding medium, creating an electromagnetic perturbation which propagates towards all directions in a spherical wave with the phase velocity of light $V_P = c/n$ (n the refraction index of the medium). If the

[4] The Baikal-GVD upgrade is being constructed, the volume indicated refers to the construction status as of summer 2018 [].

particle's velocity $v = \beta c$ is greater than V_P, the spherical waves emitted along the particle's trajectory interfere constructively, forming a coherent wave front of conical shape. The condition $\beta > 1/n$ yields the total threshold energy for the production of Cherenkov radiation for a particle with rest mass m_0:

$$E^c_{\text{th}} = \frac{m_0}{\sqrt{1-1/n^2}} \simeq 1.49 \times m_0 \text{ in water} \tag{2.16}$$

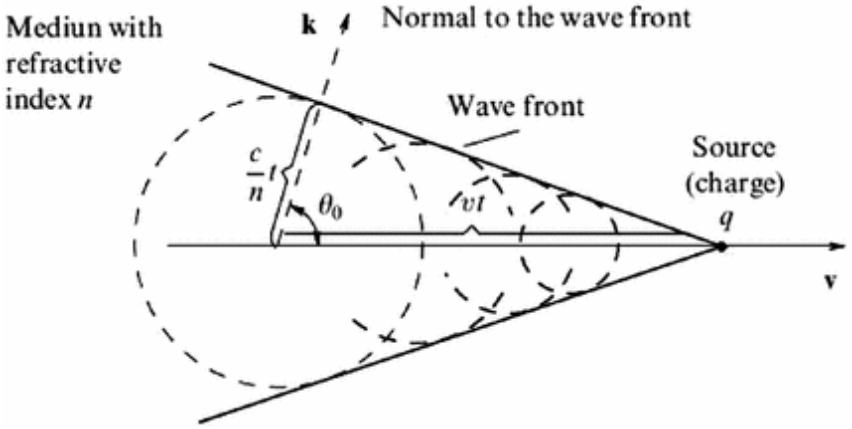

Figure 2.8: Illustration of the Cherenkov effect. Figure reused from [112].

The Cherenkov angle θ_c (θ_0 on Fig. 2.8) is defined as the angle between the particle trajectory and the emission of light; it is given by

$$\cos\theta_c = \frac{1}{\beta n}. \tag{2.17}$$

For ultrarelativistic particles $\beta \simeq 1$ and the Cherenkov angle is a characteristic of the medium: in seawater with $n = 1.35$, $\theta_c \simeq 42°$ independently of the particle energy. Water is transparent in a band of wavelengths in the visible and ultraviolet parts of the electromagnetic spectrum, and the number of Cherenkov photons emitted per unit length by a particle of charge $\pm e$ in an interval of wavelength $[\lambda, \lambda + \mathrm{d}\lambda]$ is roughly proportional to $1/\lambda^2$. Therefore, most photons are emitted in the ultraviolet and blue visible range, and Cherenkov light appears blue to the human eye. In the range of wavelengths relevant for Cherenkov detection ($\lambda \sim 300 - 600\,\text{nm}$), a particle with unit charge yields about 340 photons per cm of path length.

Secondary particles induced by neutrino events

In a generic neutrino-nucleus interaction $\nu + N \to l + X$, the lepton l can be a neutrino (in NC events) or a charged lepton e, μ or τ (in CC events). It is always accompanied by one or several hadrons, generically referred to as *hadronic system* X. Fig. 2.9 summarises the four main types of neutrino events observed at neutrino telescopes.

Due to their different masses, few-GeV muons and electrons have very different behaviours in a water-Cherenkov detector. Muons with energies up to $\sim 10\,\text{GeV}$ can be reasonably well approximated as *minimum ionising particles*: their energy deposited in matter is dominated by ionisation, with an approximately constant rate of energy loss of about 4 GeV per metre of track length independently of the muon energy. Radiative losses start to dominate over ionisation only above several hundred GeV. With a lifetime of about $\sim 10^{-6}\,\text{s}$, muons generally lose most of their energy before decaying and thus practically decay at rest. Thus muons leave clear track signatures in particle detectors, the track length being in good approximation proportional to the muon energy.

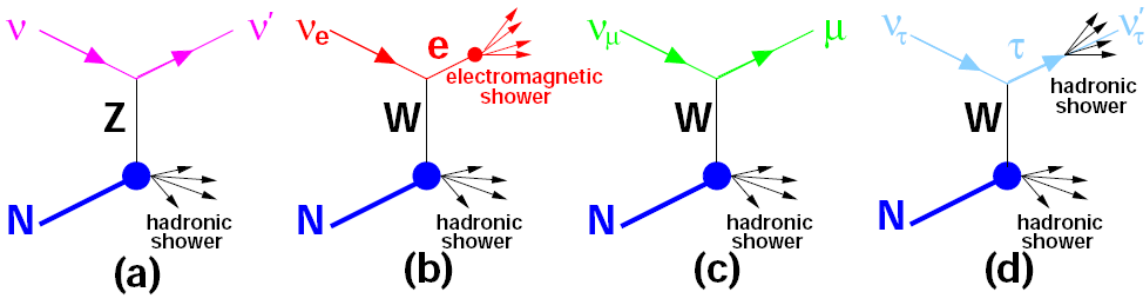

Figure 2.9: Summary of DIS neutrino event topologies in neutrino telescopes: flavour-insensitive NC (a), ν_e CC (b), ν_μ CC (c), ν_τ CC (d). In the case of the CCQE and RES interaction channels, the hadronic system may be formed simply by one or multiple nucleons or pions. Taken from Ref. [113].

Electrons, on the other hand, do not propagate far in water. They rapidly lose their energy via a cascade of repeated electromagnetic processes, most importantly Bremsstrahlung

$$e^- \to e^- + \gamma \tag{2.18}$$

and pair production

$$\gamma\ (+N) \to e^- + e^+, \tag{2.19}$$

thus initiating an *electromagnetic shower*. The relevant length scale is the *radiation length*, defined as the distance over which a high-energy electron loses 63% of its energy, which is also the order of magnitude of the mean free path for pair production by a high-energy γ. The radiation length in water is about 36 cm.

Although ν_τ CC events are suppressed at low energies due to the large τ mass, they become important at $\sim 20\,\text{GeV}$. The τ lepton has a lifetime of $\sim 10^{-13}\,\text{s}$ and thus decays very shortly after being created, producing either hadrons or leptons with almost equal branching ratios for the muonic and electronic decays:

$$\begin{aligned} \tau^- &\to \nu_\tau + \text{hadrons} \quad (64.79\%) \\ \tau^- &\to \nu_\tau + \mu^- + \overline{\nu}_\mu \quad (17.39\%) \\ \tau^- &\to \nu_\tau + e^- + \overline{\nu}_\mu \quad (17.82\%) \end{aligned}$$

At GeV energies the τ track between production and decay is very short, making it almost impossible to differentiate ν_τ from ν_e and ν_μ CC interactions on an event-by-event basis. In addition, a lesser fraction of the initial neutrino energy is visible due to the regeneration of a ν_τ which escapes the detector.

The composition of the hadronic system varies depending on the interaction channel. In water, unlike in light nuclear targets such as H_2, CCQE events typically liberate several nucleons and may even produce additional hadrons due to the higher probability of intranuclear reinteractions of the original products. In the RES channel, various combinations of nucleons and mesons (π, K) can be produced. Finally, in DIS events a spray of hadrons is produced via the hadronisation of the recoil quark. In all cases, a *hadronic shower* may develop through the cascade reinteractions of the original hadrons with the medium. The development of hadronic showers is driven by nuclear interactions, but electromagnetic processes are also involved. In particular, neutral pions are produced and decay into two photons, which initiate electromagnetic sub-cascades.

In both electromagnetic and hadronic showers, Cherenkov light is emitted by all charged particles above their Cherenkov threshold. The longitudinal elongation of the showers is similar, with a logarithmic dependence on the shower primary energy. Around a few GeV, the brightest

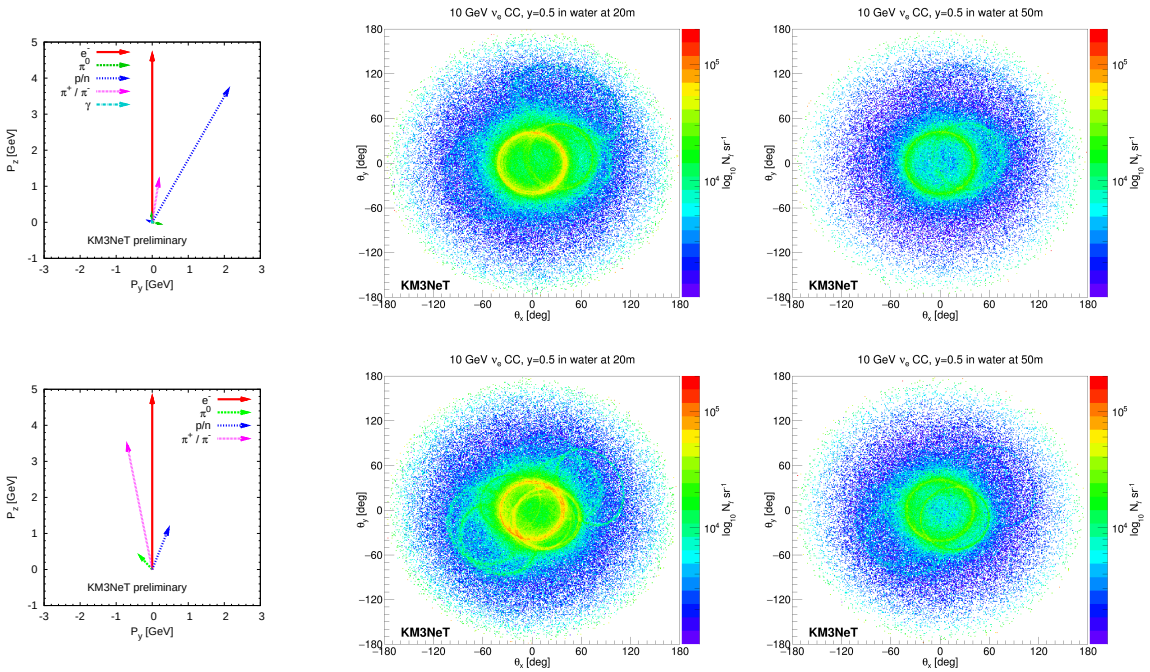

Figure 2.10: Two different simulated ν_eCC events with $E_\nu \approx 10\,\text{GeV}$ and $y \approx 0.5$ in the upper and lower row. Each event is rotated in such a way that the electron is in the z-direction. Left: Illustration of the particles produced in the two events. Each arrow represents one particle. The arrow direction and length correspond to the particle momentum in the p_y-p_z-plane, and the arrow colour indicates the particle type. Middle and right: Photon distributions in sea water recorded on shells at 20 m and 50 m around the neutrino interaction vertex. Each photon is weighted with the solid angle averaged effective area of a PMT for the photon wavelength. The Cherenkov ring from the electron is centred around $(0,0)$ with an opening angle of 42°, as the electron moves in the z-direction. Figure reused from [21].

point in the shower is situated in both cases a few metres from the initial starting position. The most important differences between electromagnetic and hadronic showers are the angular distribution of light production as well as the total light yield. For electromagnetic showers, the light yield is proportional to the initiating particle's energy with fluctuations smaller than 1%, and the lateral extension of energy deposition and light generation around the shower axis is about $\sim 0.1\,\text{m}$ (essentially energy-independent). In the energy range of interest hadronic showers emit less Cherenkov light, mostly due to the higher Cherenkov thresholds of the produced hadrons. Importantly, this light yield is subject to much larger stochastic fluctuations than for electromagnetic showers. This is due both to the variety of possible initial hadronic states (number and nature of produced hadrons, energy repartition), and to the variability of the interaction processes in the shower development. The lateral spread of hadronic showers is much larger and several distinct Cherenkov cones are usually produced, since the space angle between the initial hadrons can be large (up to $\sim 10°$). This is shown on Fig. 2.10, where the ring patterns are obtained as sections of the Cherenkov cones with a spherical shell centred on the interaction vertex.

Event topologies and reconstruction in neutrino telescopes

In CC interactions, the leading lepton (e, μ, or τ) can be used in principle to identify the neutrino flavour, bearing in mind that the τ decays almost immediately into hadrons or leptons. While a sufficiently energetic muon leaves a clear track signature, the individual particles present in electromagnetic or hadronic showers are most often indistinguishable in neutrino telescopes due to the moderate instrumentation density. Indeed, these particle showers develop on a

few metres, which is very small with respect to the typical distance between optical modules[5]. Thus, as compared to muon tracks, showers appear in first approximation as point-like emissions producing a spherical pattern of light in the detector.

All in all the distinction between hadronic and electromagnetic showers is challenging, so that flavour identification is generally limited to two classes of events:

- *Track-like* when the final state contains a clear muon track: ν_μ CC events and the minority of ν_τ CC events with muonic tau decay.
- *Shower-like* when only electromagnetic and hadronic showers are present: ν NC, ν_e CC and most ν_τ CC events.

In some cases, a ν_μ CC event may appear as shower-like, either when the neutrino energy is very low or when it is mostly transferred to the hadronic system, which is the case in very inelastic interactions (high Bjorken-y).

The main features of the reconstruction and classification of track-like and shower-like events specific to ORCA will be presented in Chap. 3. For a generic GeV-focused neutrino telescope they can be summarised as follows. In the case of track-like events, both the energy and incoming neutrino direction are typically inferred from the reconstructed muon track length, with a correction to the energy estimate from the total light yield (which includes the hadronic shower). Events with a fully contained muon track are thus more easily reconstructed. Reconstruction of hadronic showers is complicated by their large intrinsic fluctuations especially at low energy, while the clean light yield characteristics of electromagnetic showers generally benefit to the reconstruction of ν_e CC events.

As discussed in Chap. 3, the reconstruction accuracy for neutrino events is intrinsically limited by the kinematical smearing between the neutrino and the outgoing lepton: in the center-of-mass frame, the momentum of the outgoing charged lepton is not colinear with the momentum of the neutrino due to the transverse momentum going into the hadronic system, which can hardly be reconstructed. This intrinsic smearing becomes irrelevant at high energy (the momenta of outgoing particles in the detector frame being collimated forward by the effect of relativistic boost), and is also inelasticity-dependent. Other limiting effects are due to stochastic fluctuations in the development of hadronic showers, Cherenkov light generation and propagation (absorption, scattering), optical background as well as PMT and electronics response. The reconstruction accuracies are however close to the intrinsic limits arising from kinematic smearing [114].

Backgrounds

Although the shielding provided by the 2.5 km of seawater above the ORCA detector provides a reduction of the atmospheric muon flux, muons with TeV energies and above at the sea surface can still reach the detector volume before being absorbed or decay. This is frequently the case, both for single muons and muon 'bundles' (up to several hundred muons from a primary cosmic ray event). The downgoing atmospheric muon flux is thus larger than the neutrino flux by more than two orders of magnitude.

Since atmospheric muons are completely filtered by the earth, a neutrino oscillation analysis selecting only upward-going events is in principle not affected by this background, however the Cherenkov signal from atmospheric muons may produce a pattern that reconstruction algorithms misidentify as an upward-going event. Another efficient rejection method, with the downside of reducing the effective detection volume, is to require that the reconstructed interaction vertex

[5] For instance in ORCA, neighbouring optical modules are separated by 9 m vertically and about 20 m horizontally.

be strictly contained in the detector. The background rejection methods in ORCA are further described in Chap. 3.

Finally, photons produced by other sources than neutrinos and atmospheric muons constitute an additional background referred to as *optical noise*. As discussed in Chap. 3, the main sources of optical noise in seawater are the β-decays of the naturally abundant radioactive ^{40}K nuclei, as well as bioluminescence. In ice-based telescopes the optical background is more limited, which compensates for the less favourable optical properties of Antarctic ice.

2.2 Oscillation physics with neutrino telescopes at the GeV scale

2.2.1 Flavour oscillation of neutrinos crossing the earth

Three-neutrino matter effects in the atmospheric regime

As seen in Sec. 1.2.4, accounting for three-neutrino mixing the evolution equation for neutrino states in matter with constant electron number density N_e is:

$$i\frac{\mathrm{d}}{\mathrm{d}x}\left|\nu_\alpha(x)\right\rangle = \widehat{H}_M\left|\nu_\alpha(x)\right\rangle, \tag{2.20}$$

with the evolution hamiltonian given in the mass basis by

$$\widehat{H}_{M,\mathrm{mass}} = \frac{1}{2E}\left(\mathcal{D}_m + U^\dagger\,\mathcal{A}_{\mathrm{CC}}\,U\right), \tag{2.21}$$

$$\mathcal{D}_m = \begin{pmatrix} 0 & 0 & 0 \\ 0 & \Delta m^2_{21} & 0 \\ 0 & 0 & \Delta m^2_{31} \end{pmatrix}, \quad \mathcal{A}_{\mathrm{CC}} = \begin{pmatrix} A_{\mathrm{CC}} & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix}, \quad A_{\mathrm{CC}} = \pm 2\sqrt{2}G_F N_e \tag{2.22}$$

and the PMNS matrix U, following the standard parametrisation (1.39):

$$U = R_{23}\,D_{CP}\,R_{13}\,R_{12},$$

with R_{ij} the rotation of angle θ_{ij} in the plane (ij) and $D_{CP} = \mathrm{diag}\left(1,\ 1,\ e^{i\delta_{CP}}\right)$. Oscillations in matter are then solved by a set of effective mass states $(\nu_1^m,\ \nu_2^m,\ \nu_3^m)$ which diagonalise $\widehat{H}_M$.

Since R_{23} commutes with $\mathcal{A}_{\mathrm{CC}}$, the $2-3$ mixing can be factored out from the evolution equation. As a result, even in the most general case only the angles θ_{13} and θ_{12} acquire effective values in matter. When considering specifically oscillations in the atmospheric range of L/E, further simplifications can be achieved, for instance via a series expansion of the oscillation probabilities with respect to the ratio ε, known as the 'hierarchy parameter'

$$\varepsilon \equiv \frac{\Delta m^2_{21}}{\Delta m^2_{\mathrm{atm}}} \simeq 0.03 \ll 1,$$

To evaluate which terms are dominant, the smallness of the third mixing angle

$$\sin^2\theta_{13} \sim 0.02 \ll 1$$

is also taken into account. With this method the following approximate oscillation probabilities are obtained at the lowest order (adapted from Ref. [115]):

$$P(\nu_\mu \to \nu_e) \simeq \sin^2\theta_{23}\sin^2 2\theta_{13}^M \sin^2\left(\frac{\Delta^M m^2_{31} L}{4E_\nu}\right) \tag{2.23}$$

$$\begin{aligned} P(\nu_\mu \to \nu_\mu) \simeq 1 &- \sin^2 2\theta_{23}\cos^2\theta_{13}^M \sin^2\left(\frac{(\Delta m^2_{31} + \Delta^M m^2_{31})L}{8E_\nu} + \frac{V_{\rm CC} L}{4}\right) \\ &- \sin^2 2\theta_{23}\sin^2\theta_{13}^M \sin^2\left(\frac{(\Delta m^2_{31} - \Delta^M m^2_{31})L}{8E_\nu} + \frac{V_{\rm CC} L}{4}\right) \\ &- \sin^4\theta_{23}\sin^2 2\theta_{13}^M \sin^2\left(\frac{\Delta^M m^2_{31} L}{4E_\nu}\right) \end{aligned} \tag{2.24}$$

where $V_{\rm CC} = \pm 2\sqrt{2} E G_F N_e$, $A_{\rm CC} = 2EV_{\rm CC}$ and the effective $1-3$ mixing parameters are given by the same expressions as in the two-flavour case:

$$\Delta^M m^2_{31} = \xi \cdot \Delta m^2_{31}, \quad \sin^2(2\theta_{13}^M) = \frac{\sin^2(2\theta_{13})}{\xi^2}, \tag{2.25}$$

$$\xi = \sqrt{\sin^2 2\theta_{13} + \left(\cos 2\theta_{13} - \frac{A_{\rm CC}}{\Delta m^2_{31}}\right)^2} \tag{2.26}$$

Only θ_{13} and Δm^2_{31} acquire effective values in the matter potential. This can be understood by taking the lowest order in ε in Eq. 2.22, *i.e.* setting $\Delta m^2_{21} = 0$. Then R_{12} and D_{CP} can be factored out from the evolution equation, leading to an effective two-neutrino evolution where the presence of $V_{\rm CC}$ only affects the $1-3$ mixing.

The matter density thus modifies both the amplitude and phase of the appearance and survival probabilities via θ_{13}^M, $A_{\rm CC}$ and $\Delta^M m^2_{31}$. As in the two-flavour case, the interplay between the signs of Δm^2_{31} and $A_{\rm CC}$ in (2.26) imply that the matter resonance enhances the oscillations of neutrinos and suppresses the oscillations of antineutrinos if the hierarchy is normal, and conversely if the hierarchy is inverted – thus providing a handle to determine the NMH.

It is useful to examine the amplitude factors and dominant terms in Eqs. 2.23 and 2.24. $\nu_\mu \leftrightarrow \nu_e$ transitions are straighforward: they are controlled by $\sin^2\theta_{23} \sim 0.5$ (thus sensitive to the octant of θ_{23}) and $\sin^2\theta_{13}^M$ which equals ~ 0.02 far from the resonance and ~ 1 at the resonance. The case of $\nu_\mu \to \nu_\mu$ survival probability is more complex. Away from the MSW resonance, $\cos^2\theta_{13}^M \sim 1$ and the first oscillating term dominates. This term, being proportional to $\sin^2(2\theta_{23})$ which is invariant by $\theta_{23} \to \pi/2 - \theta_{23}$, is thus insensitive to the octant of θ_{23}. Closer to the resonance the other terms become dominant since $\cos^2\theta_{13}^M \to 0$ and $\sin^2\theta_{13}^M \to 1$. While the second and third terms have similar amplitudes in vacuum, the factor $\sin^2 2\theta_{13}^M$ vanishes at the resonance, so that the only term sensitive to the θ_{23} octant in the muon channel is essentially always subdominant.

Note that $\delta_{\rm CP}$ is absent from the expressions at lowest order (2.23, 2.24). Terms proportional to $\cos\delta_{\rm CP}$ and $\sin\delta_{\rm CP}$ appear at order ε in the series expansion, but these terms are further suppressed due to being proportional to $\sin 2\theta_{13}$ [115]. As a consequence, the probabilities for T-reversed transitions are identical at lowest order: $P(\nu_\mu \to \nu_e) \simeq P(\nu_e \to \nu_\mu)$.

Matter resonance in Earth densities

The radial profile $\rho(r)$ of matter density in the earth is shown on Fig. 2.11; the structure and composition of the earth as well as the reference model used here are discussed in Sec. 2.3.1. The definition of zenith angle θ_z used throughout the thesis is also reminded in the right panel. The variable $\cos\theta_z$ is generally used rather than θ_z, one of the reasons being that the baseline

L for a given trajectory is directly proportional to $\cos\theta_z$, with

$$L = -2R_\oplus \cos\theta_z \tag{2.27}$$

where $R_\oplus = 6371\,\text{km}$ is Earth's radius[6]. Note that for Earth-crossing neutrino trajectories, *i.e.* upward-going trajectories from the point of view of the detector, $\cos\theta_z < 0$. For instance the trajectory represented on Fig. 2.11 corresponds to $\cos\theta_z = -0.92$.

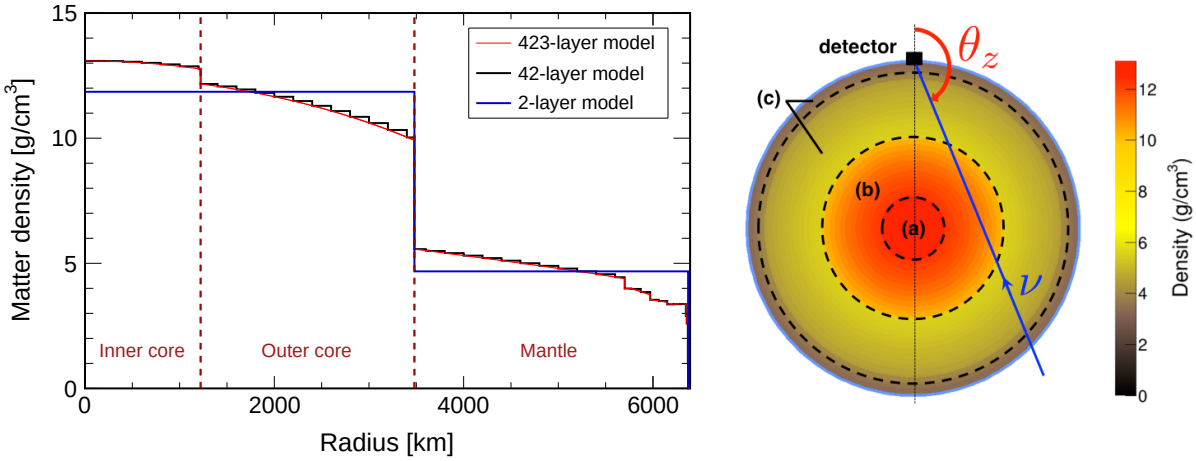

Figure 2.11: **Left:** radial models for Earth's matter density profile. Shown in red is a 423-layers model derived from the (piecewise-polynomial) Preliminary Reference Earth Model (PREM) [116]. The PREM is further discussed in the next section. Also shown are the 42-layer model used throughout the thesis and the 2-layer approximation discussed in this section. **Right:** color scale plot of the 42-layer model, with indication of the main compositional layers: (a) Inner core (IC); (b) Outer core (OC); (c) Mantle and crust (silicate Earth). Also shown is the definition of the zenith angle θ_z used throughout the text (defined with respect to the local vertical at the detector and incoming neutrino direction).

Given a spherical shell defined by an internal radius R_{min} and external radius R_{max}, the minimal and maximal zenith angles of neutrino trajectories crossing the shell are given by

$$(\cos\theta)_{\text{min/max}} = -\sqrt{1 - \left(\frac{R_{\text{min/max}}}{R_\oplus}\right)^2} \tag{2.28}$$

Tab. 2.1 shows the defining radius and corresponding zenith angle for the main layers of the reference Earth model used in the thesis (see Sec. 2.3.1 for details).

Layer	Inner core	Outer core	Lower mantle	Upper mantle	Crust
R_{max} (km)	1221.5	3480	5701	6346	6368
θ_z (°)	168.9	146.9	116.5	94.8	90.0
$180 - \theta_z$ (°)	11.1	33.1	63.5	85.2	90.0
$\cos\theta_z$	[-1.0, -0.981]	[-0.981, -0.837]	[-0.837, -0.446]	[-0.446, -0.083]	[-0.083, 0.0]

Table 2.1: Correspondance between external radius R_{max} and limiting zenith angle θ_z for Earth's layers. The indicated values correspond to the 42-layer model of Fig. 2.11. The oceanic crust model is used and the detector is assumed to be located 3 km below the sea level.

[6] In the case of ORCA which is located approximately at 3 km depth in the Mediterranean sea, the value $R'_\oplus = 6368\,\text{km}$ is used.

In first approximation the electron density N_e is proportional to the matter density ρ:

$$N_e = \frac{Z}{A}\frac{\rho}{m_p} \simeq 0.5 \times \frac{\rho}{m_p} \tag{2.29}$$

where m_p is the proton mass and the proton-to-nucleon ratio Z/A, depending on the chemical composition of the layer, is generally close to 0.5 (especially for the light elements found in the mantle).

The modified matter potential $A_{\rm CC}$ can then be expressed in SI units as

$$A_{\rm CC} \simeq \left(1.53 \cdot 10^{-4}\,{\rm eV}\right) \times \frac{Z}{A} \times \left(\frac{\rho}{\rm g.cm^3}\right) \times \left(\frac{E}{\rm GeV}\right) \tag{2.30}$$

As visible from Fig. 2.11, the essential feature of Earth's radial density profile is the large density jump at the core-mantle boundary (CMB). Except for this discontinuity, the density variations are slow with respect to the Δm^2_{31} oscillation length, so that the approximate two-layer model is qualitatively sufficient to describe the flavour transitions of core-crossing neutrinos.

For a layer with constant matter density ρ the MSW resonance $\sin^2\theta^M_{13} \to 1$ thus occurs for neutrinos (antineutrinos) in the normal (inverted) mass hierarchy at the resonance energy:

$$E_{\rm res} \equiv \frac{\Delta m^2_{31}\cos 2\theta_{13}}{2\sqrt{2}G_F N_e} \simeq 7\,{\rm GeV}\left(\frac{4.5\,{\rm g/cm^3}}{\rho}\right)\left(\frac{\Delta m^2_{31}}{2.4\times 10^{-3}\,{\rm eV^2}}\right)\cos 2\theta_{13} \tag{2.31}$$

where the average mantle density $\rho \simeq 4.5\,{\rm g/cm^3}$ and the corresponding resonance energy $E_{\rm res} \simeq 7\,{\rm GeV}$ have been introduced as reference values. For typical core densities $\rho \sim 12\,{\rm g/cm^3}$, the resonance energy is closer to $E_{\rm res} \sim 2.5\,{\rm GeV}$.

More generally, the dependence of the effective mixing parameters on the product $\rho \cdot E_\nu$ is shown on Fig. 2.12. Assuming a normal mass hierarchy, the MSW resonant regime corresponds to values of $\rho \cdot E_\nu$ approaching $x_0 \simeq 30\,{\rm GeV.g.cm^{-3}}$. The vacuum mixing, where θ_{13}-driven $\nu_\mu \leftrightarrow \nu_e$ transitions are subdominant, is recovered when $\rho E_\nu \ll x_0$. In contrast, when $\rho E_\nu \gg x_0$ a saturation effect occurs, and these transitions are strongly suppressed. In this regime the oscillations are effectively described by two-flavour $\nu_\mu \to \nu_\tau$ transitions, as obtained when setting $\sin^2 2\theta^M_{13} \to 0$ in Eqs. (2.23, 2.24).

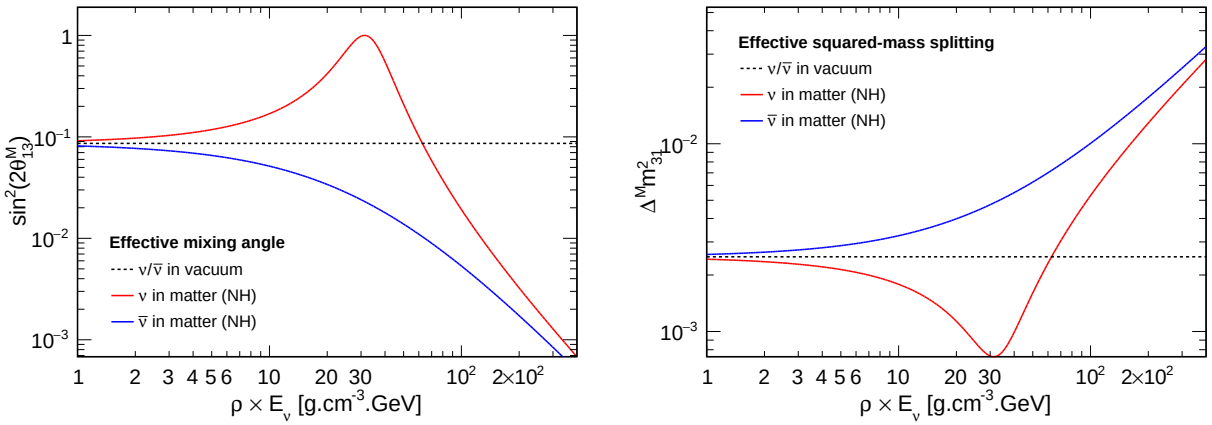

Figure 2.12: Effective mixing parameters in matter as a function of the product of matter density and neutrino energy. While $\sin^2 2\theta^M_{13}$ directly controls the amplitude of the oscillations in constant matter density, $\Delta^M m^2_{13}$ is closely related to the oscillation length – see Eqs. (2.23, 2.24). The resonant mixing $\sin^2 2\theta^M_{13} \to 1$ is associated with a minimum effective $\Delta^M m^2_{13}$, *i.e.* a maximum oscillation length.

Parametric enhancement

In the case of neutrinos crossing both Earth's mantle and core, new resonance-like effects become possible in the $\nu_\mu \to \nu_e$ and $\nu_e \to \nu_\mu, \nu_\tau$ channels. While the resonance enhancement in constant density maximises flavour transitions by amplifying neutrino mixing ($\theta \to \theta^m \simeq \pi/4$), these effects arise from a match of the variation of the matter density with the phase of neutrino oscillations. The transition probability can then become maximal even if the effective mixing angle in matter remains small (*i.e.* far from the MSW resonance). This occurs for instance when the enhancement accumulates as neutrinos traverse many repetitions of a periodic 'castle wall' density profile (square function). This effect was first studied in the late 1980's [117–119], and dubbed *parametric resonance* in analogy with the resonance in dynamical systems whose parameters vary periodically with time (for instance a swing).

Atmospheric neutrinos whose trajectories intersect Earth's core ($-1 \leq \cos\theta_z \leq -0.84$) go through a mantle-core-mantle density profile along their path. With the two-layer approximation shown on Fig. 2.11, the density sequence is analogous to a truncated periodic castle wall profile. Since less than two periods are described, there is no accumulation and the parametric-like enhancement occurs mostly for energies close to the MSW resonance where the effective mixing angle is already significant. Following the discovery of neutrino oscillations this effect[7] was extensively discussed both for solar and atmospheric neutrinos [120–124]. A review can be found in Ref. [125]. More generally, oscillations in an arbitrary inhomogeneous density profile can be studied by means of Fourier analysis, retrieving the parametric enhancement effect [126, 127].

Fig. 2.13 demonstrates the impact of the MSW resonance and its combination with the parametric-like effect. Several different Earth models are considered: empty Earth ($\rho = 0$), uniform density in the whole Earth with $\rho = \rho_{\rm M}$ or $\rho = \rho_{\rm C}$, two-layer and 42-layer models as in Fig. 2.11. To illustrate the oscillation patterns, the $\nu_\mu \to \nu_e$ oscillation probability is shown as a function of the path length inside the earth for a neutrino traversing the planet diametrically. However it should be kept in mind that only the probability at the detector (*i.e.* at $L \simeq 12600\,$km) is actually observed. The neutrino energy is fixed at 4 GeV, which is intermediate between the mantle and core resonance energies. Therefore the MSW resonance increases the amplitude and oscillation length with similar intensity in both one-layer models[8] with respect to the vacuum case. As a related fact, the parametric-like effect in the two-layer model is very strong. The close similarity of the oscillation probabilities obtained in the two-layer and 42-layer models reflects the fact that flavour oscillations are at first order sensitive to the matter density averaged over $\mathcal{O}(L_{\rm osc})$ distances, as further discussed in Sec. 2.3.3.

Early achievements with neutrino telescopes

The existing neutrino telescopes ANTARES and IceCube/DeepCore[9] have detection thresholds of 10 to 20 GeV and effective detection volumes of $\mathcal{O}(10\,$Mton). They are thus sensitive to the saturated regime dominated by $\nu_\mu \to \nu_\tau$ disappearance with unprecedented statistics, and can measure the atmospheric mixing parameters without being affected by the degeneracies introduced by matter-enhanced three-flavour effects (value of θ_{13}, mass hierarchy, $\nu/\overline{\nu}$ flux ratio, *etc*). More generally measurements from neutrino telescopes are complementary to those from long-baseline accelerator experiments (MINOS, T2K, NOνA) and Super-Kamiokande, being affected by different sources of systematic uncertainty. As compared to Super-Kamiokande, neutrino telescopes also benefit from the fact that neutrinos above 20 GeV almost exclusively in-

[7] The same mechanism is called *oscillation length resonance* in Ref. [122–124].

[8] Indeed at 4 GeV the mapping parameter ξ is almost equal for both models: $\xi(\rho_M) = 0.485$, $\xi(\rho_C) = 0.502$.

[9] DeepCore is the name given to the denser sub-array inside the IceCube detector, dedicated to low-energy studies [129]. The fiducial volume used in DeepCore analyses is approximately 2.5 times larger than ANTARES, with a higher instrumentation density in the dedicated sub-array.

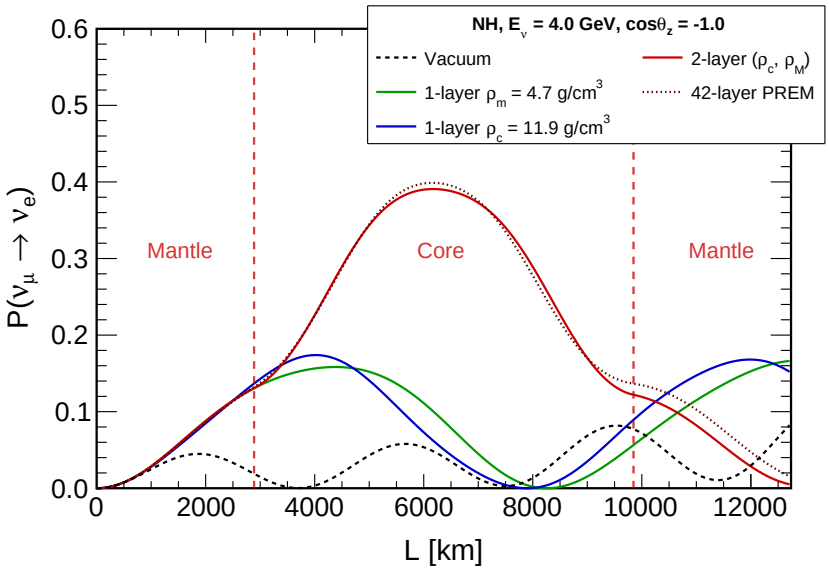

Figure 2.13: Effect of the mantle-core-mantle parametric resonance for $\nu_\mu \to \nu_e$ transitions of 4 GeV atmospheric neutrinos crossing Earth diametrically. Mind that the x-axis is the coordinate along the neutrino path inside Earth, with fixed energy and zenith angle. The oscillations probabilities have been computed numerically with the OscProb software [128], described in Chap. 4.

teract via the DIS channel (with a much better constrained cross-section). However, due to their sparser instrumentation they have worse energy-angular resolutions and particle identification capabilities.

ANTARES provided in 2012 the first measurement of atmospheric neutrino oscillation in the energy range 20-100 GeV [130]. The successive measurements reported by DeepCore since then have achieved better precision [131, 132]. As shown on Fig. 2.14, they are now competitive with those of Super-Kamiokande and accelerator experiments.

2.2.2 Measurement of the neutrino mass hierarchy

Measuring the NMH from the oscillations of atmospheric neutrinos requires to probe the MSW-resonant energy range, which cannot be done with sufficient precision by existing neutrino telescopes due to their relatively high detection thresholds. In the past two decades, various strategies have been discussed to go forward: magnetised detectors [138–148], and water-Cherenkov detectors [121, 122, 143, 149–152], including more recently the specific case of megaton-scale detectors deep under ice (PINGU, a proposed low-energy extension of IceCube) or under sea (ORCA, low-energy branch of KM3NeT) [69, 100, 153–162].

Oscillation probabilities $P(\nu_X \to \nu_\mu)$ and $P(\overline{\nu}_X \to \overline{\nu}_\mu)$, with $X = e$, μ[10], are shown in Fig. 2.15 as a function of the neutrino energy for various zenith angles θ_z (i.e various baselines and density profiles in the earth). In each case, both NMH hypotheses are represented. As expected from the approximate transition probabilities (2.23, 2.24), the NMH affects both the amplitude and phase of the oscillations, with the strongest impact in the resonance region $E_\nu \sim (4-8)$ GeV. The effect of the mantle-core-mantle resonant enhancement is visible in the region $E_\nu < 7$ GeV in the uppermost panel ($\cos\theta_z = -1$). Above 15 GeV the $\nu_e \to \nu_\mu$ transition probability becomes very small and differences from distinct hierarchies tend to disappear as well.

Fig. 2.15 shows that to first order, the effect for neutrinos in the NH case is the same as for antineutrinos in the IH case. Neutrino and antineutrino interactions in the GeV range are undistinguishable in a neutrino telescope, essentially blind to the charge of outgoing leptons. Nevertheless, a net asymmetry in the combined $(\nu + \overline{\nu})$ event rates between NH and IH for a given flavour can be observed. This mainly comes from the fact that in the GeV range the

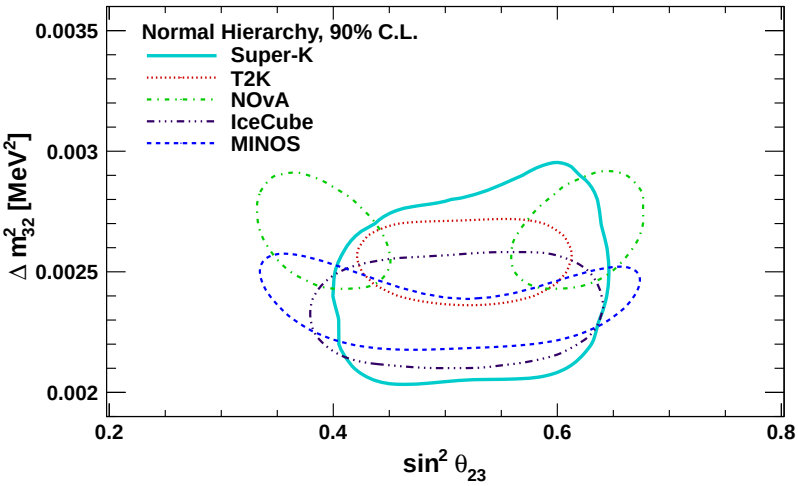

Figure 2.14: Measurement of the atmospheric oscillation parameters by accelerator experiments, SK and IceCube/DeepCore. Figure taken from Ref. [133]. Contours at 90% C.L. from analyses assuming the normal mass hierarchy are shown. The Super-K contour (cyan) is taken from the analysis with $\sin^2\theta_{13}$ assumed to be 0.0219 ± 0.0012 [133]. Contours from the T2K (violet) [134], NOvA (dashed green) [135], MINOS+ (dashed blue) [136], and IceCube (red) [137] experiments are also shown.

CC cross section is different (by about a factor of 2) for neutrinos and antineutrinos, as seen in Sec. 2.1.2 (see Fig. 2.5). The relative contribution of ν_e and ν_μ in the atmospheric flux also affects the number of events of each flavour that can be expected at the detector level. In particular, the θ_{13}-driven transitions between ν_μ and ν_e are observable even though (in first approximation) $P(\nu_\mu \to \nu_e) \simeq P(\nu_e \to \nu_\mu)$, because the initial flux is about twice larger for the muon flavour.

The expected event rates are thus calculated by accounting for the initial flux of atmospheric neutrinos, the full oscillation probabilities in the realistic Earth density profile, and the neutrino-nucleon cross-sections. This calculation is described in detail in Chap. 4. In Fig. 2.16, the rate of CC events per neutrino flavour (summing the ν and $\overline{\nu}$ components) is shown for a fixed zenith angle, as a function of neutrino energy and multiplied by E^2. Since the atmospheric neutrino flux has a spectral index ~ 3 and the cross-sections increase about linearly with energy, the resulting differential event rate $\mathrm{d}N/\mathrm{d}E$ roughly follows an E^{-2} power law, when considering the *unoscillated* spectrum (*i.e.* averaged with respect to energy). The deviation from this behaviour in the electron channel can be explained by the decrease with energy of the ν_e fraction in the atmospheric flux (discussed in Sec. 2.1.1).

In Fig. 2.16 and in the rest of the chapter, the normal and inverted hierarchy hypotheses are defined by $|\Delta m^2_{31}|_{\mathrm{NH/IH}}$ given by the respective best fit values from Ref. [63]. Unless otherwise specified, $\sin^2\theta_{23} = 0.5$, $\delta_{\mathrm{CP}} = 0$, and other parameters follow Ref. [63].

The signature of the NMH seen in Fig. 2.16 is very different in the electron and muon channels. In the electron channel, a clear excess of events is expected in NH around the resonance region, consistently for mantle-only and core-crossing trajectories. This essentially reflects the MSW and parametric enhancement of the $\nu_\mu \to \nu_e$ appearance probability. An accurate energy reconstruction is not required to observe this excess, as it is sustained over an extended range of energy. On the contrary, in the muon channel the imprint of the NMH affects fast oscillations and it can be expected that the discrepancy between the NH and IH expectations will be partly washed out by the averaging effect of a finite resolution detector.

For both channels, the change in amplitude of the oscillatory patterns due to the impact of the NMH via θ^M_{13} is more important than the effect due to the phase of the oscillations,

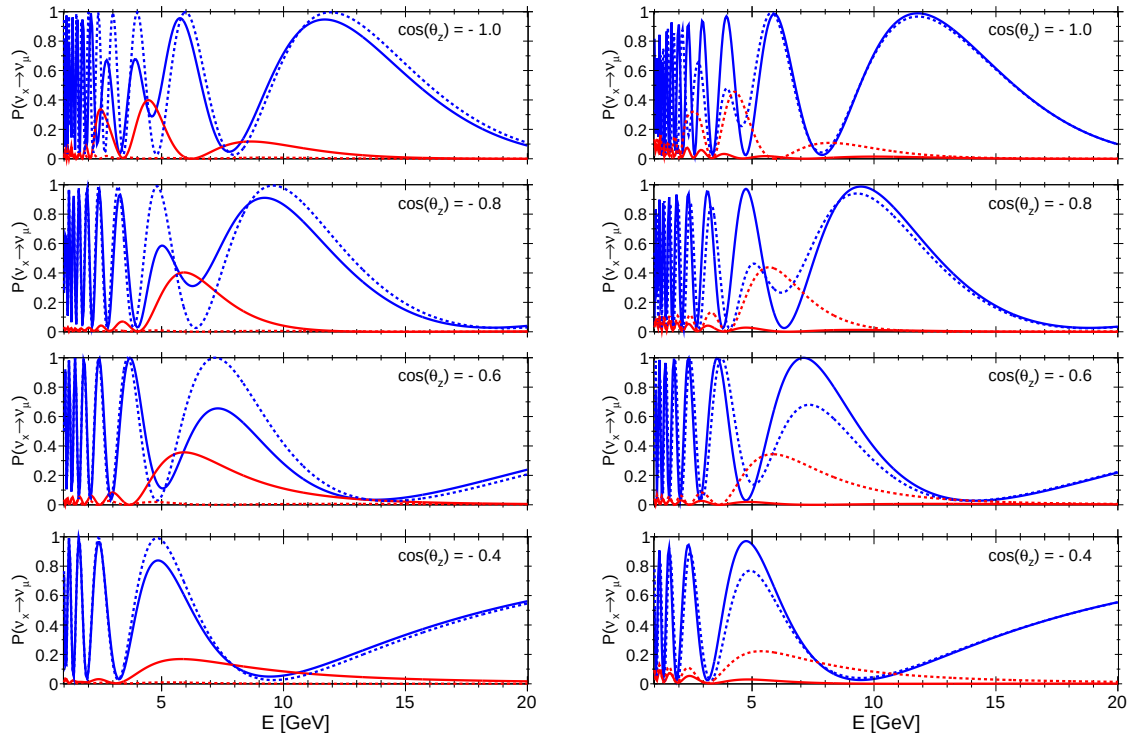

Figure 2.15: Oscillation probabilities $\nu_\mu \to \nu_\mu$ (blue lines) and $\nu_e \to \nu_\mu$ (red lines) as a function of the neutrino energy for several values of the zenith angle, corresponding to different baselines and density profiles. The top panel corresponds to a mantle-core-mantle diametrical trajectory, whereas the other three panels corresponds to mantle-only trajectories. The solid (dashed) lines are for NH (IH). Probabilities are shown both for neutrinos (left) and for antineutrinos (right). Figure reused from [21].

involving $\Delta^M m^2_{31}$. Nevertheless, this effect of energy scaling and shifting between the NH and IH expectations remains noticeable, mostly in the electron channel. It also explains the fine oscillating sub-structure, which comes from the addition of the phase-shifted ν_e and $\overline{\nu}_e$ oscillated rates. As expected, the region $E_\nu > 20\,\text{GeV}$ remains essentially unaffected by the NMH and can thus be used as a control sample to constrain unknown parameters that may be degenerate with the hierarchy in the signal region.

The $\nu_\tau + \overline{\nu}_\tau$ channel has not been discussed here. Although both the $\nu_\mu \to \nu_\tau$ and $\nu_e \to \nu_\tau$ oscillation channels are both affected by the MSW resonant enhancement and thus sensitive to the NMH, the observation of $\nu_\tau/\overline{\nu}_\tau$ CC events in the resonant region is doubly suppressed by the small ν_τ cross-sections and smaller deposited energy (see Sec. 2.1.2). Therefore, while these oscillation channels are important for ν_μ and ν_e disappearance, ν_τ appearance only becomes observable in the saturation region where it is fairly independent of the NMH. Moreover its observation is essentially indirect, due to the difficulty of tagging ν_τ events (most appear as shower-like).

The distributions of interacting event rates discussed here correspond to the idealised signal seen by a perfect detector. For a realistic, quantitative estimate of the NMH sensitivity of an atmospheric neutrino telescope, the detector response must be accounted for precisely so as to predict distributions of observed events, as function of their measured characteristics. The most important experimental features are the following:

- effective exposure, given by the efficiency of detection and reconstruction of neutrino events (multiplied by the total target mass and data-taking time),
- two-dimensional smearing of the detected event rate distributions due to the finite resolutions both in energy and zenith angle,

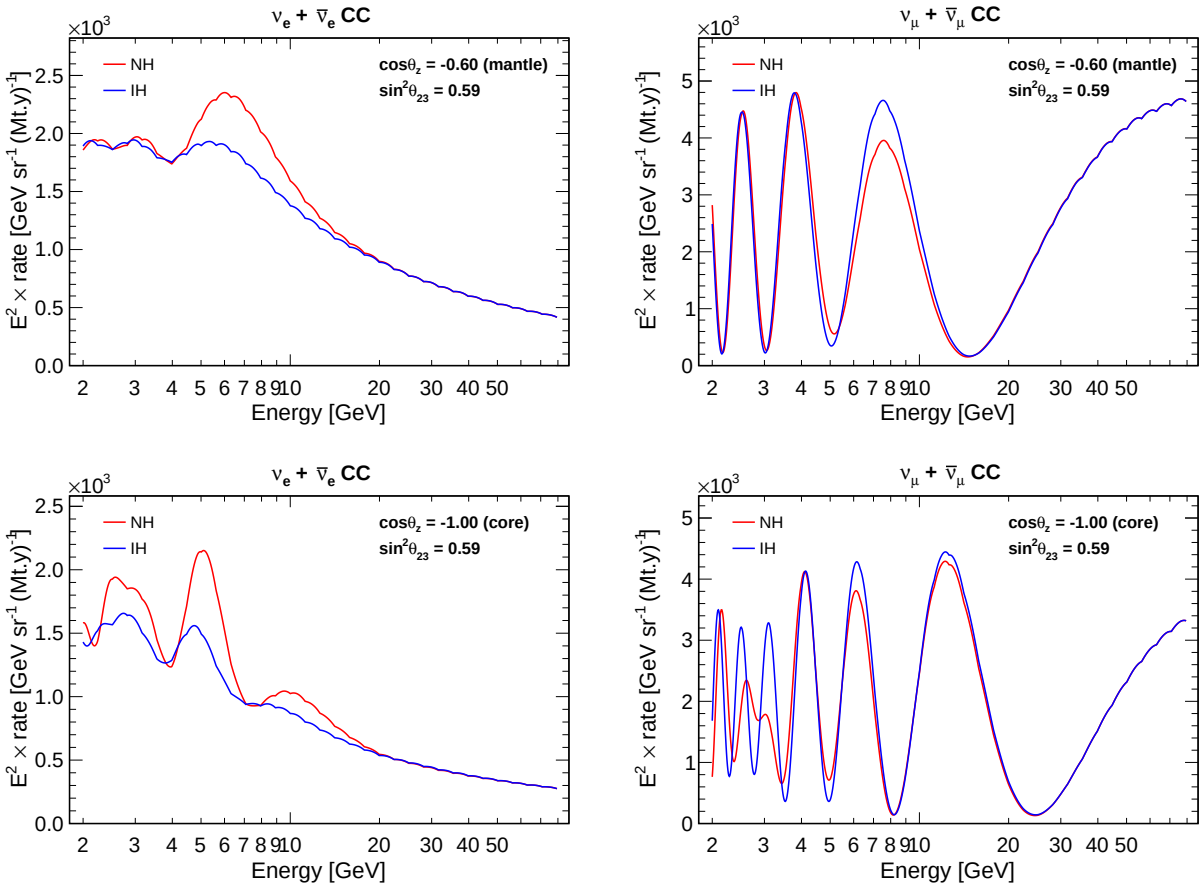

Figure 2.16: Rate of $\nu + \overline{\nu}$ CC events per neutrino flavour as a function of energy in NH and IH. Left: $\nu_e + \overline{\nu}_e$ CC events. Right: $\nu_\mu + \overline{\nu}_\mu$ CC events. The top panel corresponds to a mantle-only trajectory and the bottom panel to a diametrical mantle-core-mantle trajectory. The plotted quantity is a double differential rate of events $\mathrm{d}N/\mathrm{d}E\mathrm{d}\Omega$, per unit exposure (megaton-year, denoted as Mt.y), and multiplied by E^2 so that the oscillation patterns appear more clearly.

- limited flavour identification performance, partly blending together the electron and muon channels,
- impact of the background of flavour-insensitive NC events, misidentified atmospheric muons and optical noise events.

These response functions are in general dependent on the type of neutrino event (flavour, interaction channel), on the true neutrino energy and zenith angle as well as on the event kinematics which determines, for the most part, the observable topology. The dependence of the response functions on the true interaction inelasticity y is thus taken into account in this thesis. In addition, the y-distributions are significantly different for neutrinos and antineutrinos (see Fig. 2.7), so that a reasonably accurate inelasticity estimate may provide some statistical separation power between ν and $\overline{\nu}$, which would in turn yield a potentially large increase in NMH sensitivity [100].

The detailed calculations of two-dimensional (E,θ) oscillograms and (E,θ,y) event rate distributions are discussed in Chap. 4. The Monte Carlo based modeling of detector efficiencies, resolution and event classification functions for ORCA is then adressed in Chap. 5.

2.2.3 Oscillation parameters and degeneracies

Fig. 2.17 shows the values spanned by the oscillation probabilities for the normal mass hierarchy when oscillation parameters are allowed to vary within the range allowed by the global fit [163]. The uncertainties having a large impact here are those related to θ_{23}, Δm^2_{31} and δ_{CP}.

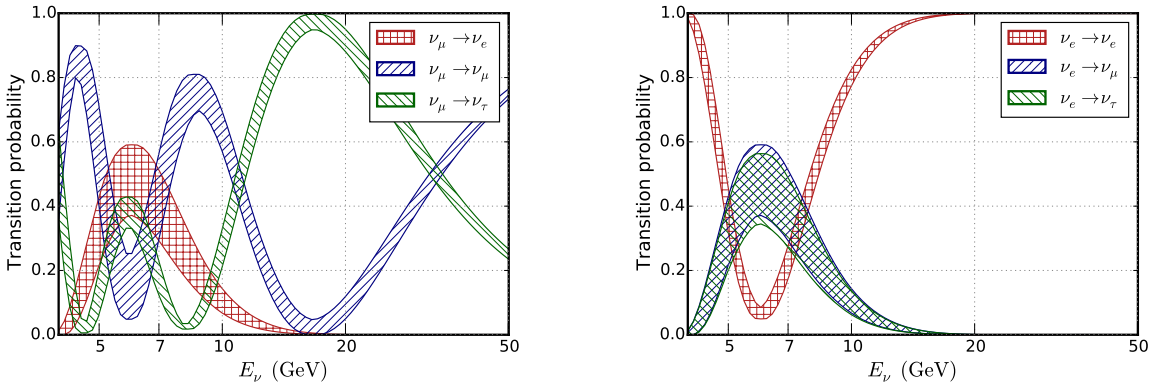

Figure 2.17: Transition probabilities $\nu_\mu \to \nu_X$ (left) and $\nu_e \to \nu_x$ for a neutrino trajectory in the earth defined by $\cos\theta_z = -0.7$ (mantle-crossing trajectory). Normal mass hierarchy is assumed. The bands encompass the results of the calculation when the oscillation parameters are varied within their uncertainties from Ref. [163]. Figure taken from Ref. [162].

The extent to which these uncertainties could spoil the NMH determination must be examined; another interesting issue is the measurement of these parameters, either assuming a fixed mass hierarchy or independently of the NMH.

Fig. 2.18 to 2.20 show how individual variations of these parameters affect the expected $\nu + \overline{\nu}$ event rates in NH and IH. For $\sin^2\theta_{23}$ and Δm^2_{31} the parameter ranges are chosen to approximately match the 3σ allowed range of the global fit [63], while for δ_{CP} four values $(0, \pi/2, \pi, 3\pi/2)$ spanning the whole physical range have been used without considering the global fit constraint.

Atmospheric mixing angle θ_{23}

As visible from Fig. 2.18, θ_{23} affects the amplitude of the oscillatory patterns in both channels. On the one hand, it is clear that in the electron channel a variation of $\sin^2\theta_{23}$ within its 3σ allowed range can mimick the effect of the hierarchy flip. For instance, the $(\nu_e + \overline{\nu}_e)$ oscillated rates obtained for the combinations $\{\text{NH}, \sin^2\theta_{23} = 0.4\}$ and $\{\text{IH}, \sin^2\theta_{23} = 0.6\}$ are closely similar in the mantle-crossing case. On the other hand, the degeneracy is not as pronounced in the muon channel. There is no significant overlap between the NH and IH sets of curves in the resonant energy ranges, while outside of the resonant regions the amplitudes are independent of the NMH and controlled by $\sin^2(2\theta_{23})$, as shown by the last visible oscillation minimum[11] (saturated regime). Therefore the combination of the electron and muon channels is expected to help in resolving the NMH-θ_{23} degeneracy in the ν_e appearance probability. A limitation to this statement is due to the *octant degeneracy*, resulting from the fact that the terms controlled by $\sin^2(2\theta_{23})$ are invariant by the reflection with respect to maximal mixing ($\theta_{23} = \pi/4$). This degeneracy is the most severe if θ_{23} is close to maximal:

$$\theta'_{23} = \frac{\pi}{2} - \theta_{23} \Rightarrow \begin{cases} \text{octant-degenerate terms:} & \sin^2(2\theta'_{23}) = \sin^2(2\theta_{23}) \\ \text{octant-sensitive terms:} & \sin^2\theta'_{23} - \sin^2\theta_{23} = \pm\sqrt{1 - \sin^2(2\theta_{23})} \end{cases} \quad (2.32)$$

Atmospheric mass splitting $|\Delta m^2_{31}|$ and CP-violating phase δ_{CP}

In Eqs. 2.23-2.24, the vacuum value of $|\Delta m^2_{31}|$ is a common factor in the phase of all oscillating terms, except for the energy-independent phase shift $V_{\text{CC}}L/4$. A change of $|\Delta m^2_{31}|$ thus results in a shift of the oscillatory pattern, as shown by Fig. 2.19. The best handle on this parameter

[11] This is also frequently called the first oscillation minimum, referring to oscillations with respect to L/E_ν.

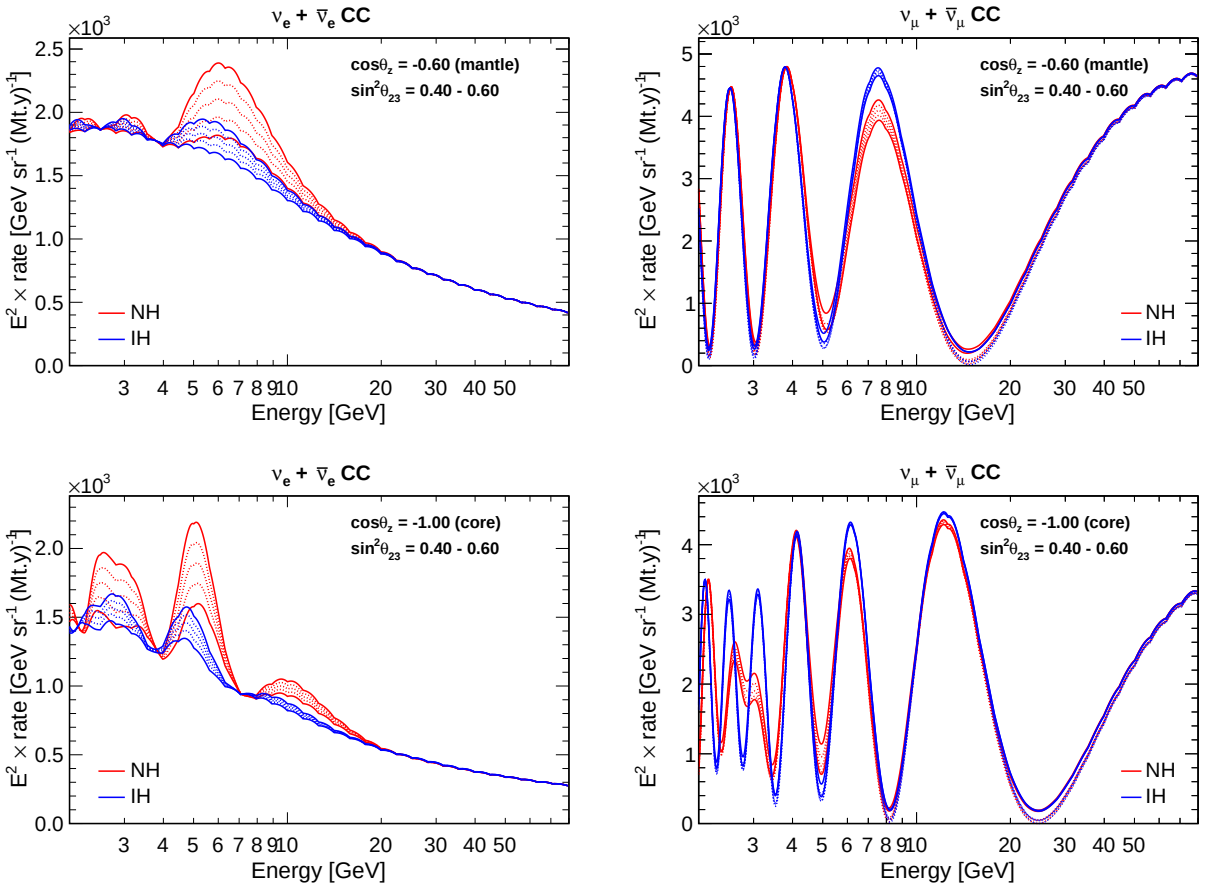

Figure 2.18: Effect of θ_{23} on the $\nu + \bar{\nu}$ event rate measured in neutrino telescopes. The plotted quantity and neutrino trajectories are the same as in Fig. 2.16. The range of values $0.40 \leq \sin^2\theta_{23} \leq 0.60$ corresponds approximately to the 3σ range of the global fit [63].

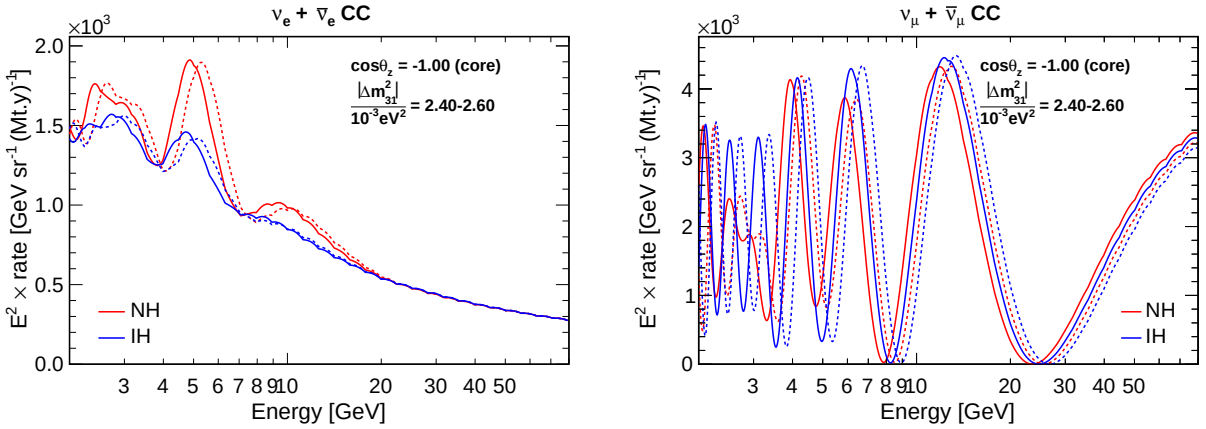

Figure 2.19: Effect of $|\Delta m^2_{31}|$ on the $\nu + \bar{\nu}$ event rate measured in neutrino telescopes. The plotted quantity is the same as in Fig. 2.16. Only the diametrical ($\cos\theta_z = -1$) trajectory is represented; the energy shift effect is identical for mantle-only trajectories. The two values of $|\Delta m^2_{31}|$ correspond approximately to the lower and upper bounds of the 3σ allowed range in the global fit [63].

is the measurement of the position of the first L/E oscillation minimum, corresponding to the minimum close to 25 GeV on Fig. 2.19. Clearly, this measurement requires a very good absolute energy calibration.

The effect of $\delta_{\rm CP}$ variations is shown on Fig. 2.20 for the diametrical core-crossing trajectory, where it has the most impact. The induced modifications in the electron channel are rather

complex, affecting both maximal amplitudes and shape of the oscillated rates. The degeneracy introduced for the mass hierarchy measurement in this channel is non-negligible, especially when combined with θ_{23} and $|\Delta m^2_{31}|$. However the muon channel remains essentially unaffected.

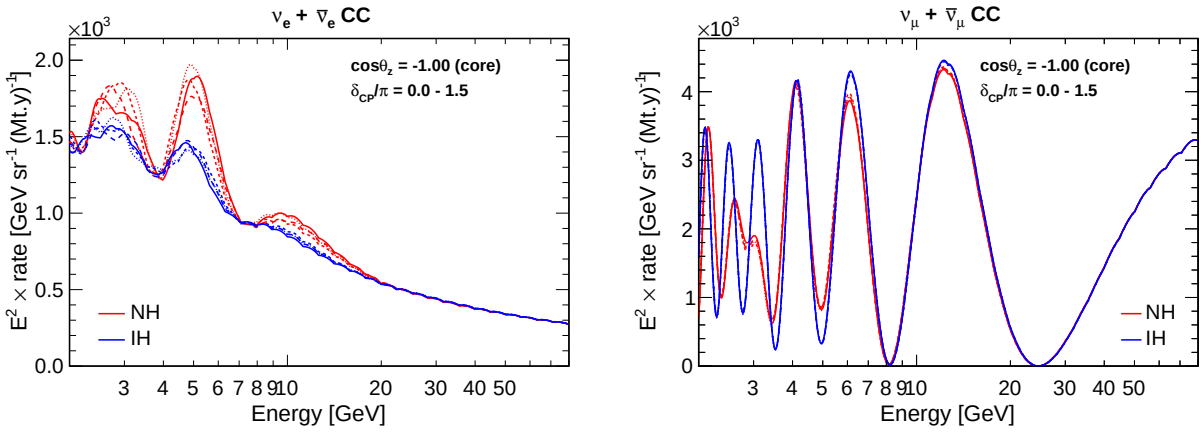

Figure 2.20: Effect of $\delta_{\rm CP}$ on the $\nu + \overline{\nu}$ event rate measured in neutrino telescopes. For each true hierarchy, four curves are represented, corresponding to $\delta_{CP} = 0, \pi/2,\ \pi,\ 3\pi/2$. Only the diametrical ($\cos\theta_z = -1$) trajectory is represented; the $\delta_{\rm CP}$ effect for mantle-only trajectories is qualitatively similar, though less intense.

The impact of the uncertainty on $\sin^2\theta_{13}$ is not shown here. Away from the resonance this parameter acts on the oscillated rates similarly as θ_{23}, simply controlling the oscillation amplitude. In the resonance region this gets more complicated, as the vacuum value of θ_{13} affects the oscillation phase via the mapping parameter ξ. Either way, due to the stringent constraints from reactor data the amplitude of this perturbation within the 3σ allowed range is relatively small as compared to the effects discussed above.

Generally speaking, the availability of many different oscillation baselines and density profiles is greatly beneficial to atmospheric neutrino experiments for resolving the degeneracies between the NMH, θ_{23} and $\delta_{\rm CP}$. The inclusion of external reactor constraints on θ_{13} is however necessary.

2.3 Probing Earth's interior with neutrinos

2.3.1 The composition of Earth and its unknowns

In the most simple description of its interior structure, Earth can be divided in three concentric, approximately spherical shells: the crust, mantle, and core. As shown on Fig. 2.11, the density steadily increases with depth, with a large discontinuity at the core-mantle boundary (CMB). The CMB corresponds to a radical change in composition: while the mantle and crust are essentially composed of silicate minerals (with Si and O as main chemical elements), the core is metallic and is predominantly made of iron.

The outermost layer is the crust, representing less than 1% of Earth's mass. Its composition and thickness (from 5 to 70 km) vary greatly depending on the location (continental and oceanic crust), and its dynamical evolution due to tectonic activity is important. It is the most easily accessible layer for experimental study. In particular, direct sampling methods are possible from drills (down to ~ 12 km) and old mountain ranges.

As to the deep Earth, the knowledge of its structure primarily comes from the study of seismic wave propagation. The mantle, representing by far the largest fraction of Earth's volume, is subdivided into the *uppermost mantle* (from the base of the crust to about 400 km depth), the *transition zone* ($\sim 400 - 660$ km)[12], and the *lower mantle* ($\sim 660 - 2900$ km). To first

[12] The *upper mantle* refers to the uppermost mantle together with the transition zone.

order, the mantle is globally uniform in chemical composition. The layers are distinguished by their different seismic properties, which vary due to cristalline phase transitions as the pressure increases with depth. Temperature gradients are small, due to convection.

While the uppermost part of the mantle can be studied by entrainment sampling, the state and composition of the deep mantle remain more uncertain. For the most part, the lower mantle is homogeneous from the point of view of seismic properties. However, both the transition zone and the $\sim 200\,\mathrm{km}$ thick layer directly above the CMB (known as the D'' region) are very complex and exhibit local seismic anomalies.

Earth's core is believed to be mostly formed of iron, and subdivided into a solid *inner core* (IC) with an approximate radius $R \sim 1200\,\mathrm{km}$, surrounded by the molten *outer core* (OC). The outer core is primarily a $\mathrm{Fe}-\mathrm{Ni}$ alloy, with an admixture of light elements of unknown nature and proportions. Convective motion in the fluid outer core, presumably driven by thermal and compositional buoyancy and shaped by the Coriolis force, is believed to be the mechanism generating Earth's magnetic field.

Seismic measurements and density profile

The picture of Earth as a metallic core surrounded by a silicate shell was first hypothesized in the late 18th century based on estimations of its average density, found to be significantly larger than the density of surface rocks. The existence of stony meteorites and iron meteorites suggested that the metallic core could be primarily composed of iron. This picture was reinforced by the measurement of the moment of inertia of the planet, indicating a concentration of the mass towards its center. Based on these observations, the hypothesis of a large and dense iron-nickel core was formulated by E. Wiechert in 1897, and the first accurate determination of the location of the core-mantle boundary at about 2900 km depth is due to B. Gutenberg in 1914. The conclusive evidence for the existence of a distinct core came from the observation of a sharp decrease with depth of the velocities of seismic waves of type P (primary) and S (secondary), which are the two types of body waves that propagate in the interior of the earth. While P-waves are compressional (longitudinal deformation), S-waves are shear waves (transverse deformation). Their velocity is determined by the density ρ and elasticity properties of the medium, the incompressibility (or bulk modulus) K and the shear modulus μ:

$$V_P = \sqrt{\frac{K + 4\mu/3}{\rho}}, \quad V_S = \sqrt{\frac{\mu}{\rho}} \tag{2.33}$$

S-waves generally travel slower (hence their denomination as secondary). Due to the variation of the above physical properties the velocities of both S- and P-waves are depth-dependent, causing a refractive bending of their trajectories as illustrated on Fig. 2.21. Refraction and reflection also occur at discontinuities. In addition, S-waves do not propagate in the absence of a shear modulus, *i.e.* in fluid materials. The existence of the liquid outer core is therefore characterised by a shadowing effect, where no S-waves are detected on the side of the earth opposite to the location of an earthquake. In fact, as seen on the left panel of Fig. 2.21 a more modest shadowing effect is also observed for P-waves, due to refraction (important decrease of V_P in the outer core). Finally, the inner core was established to be solid only relatively recently, via the study of whole-Earth free oscillation eigenmodes [164].

The Preliminary Reference Earth model and density uncertainties

Using the theoretical knowledge on wave propagation, seismic data can be inverted to obtain a map of the spatial variation of mechanical properties inside the earth as a function of planetary radius (or equivalently, depth): wave velocities, density, elasticity constants, anelastic attenuation factors, pressure, *etc.* Although Earth is not exactly spherically symmetric and

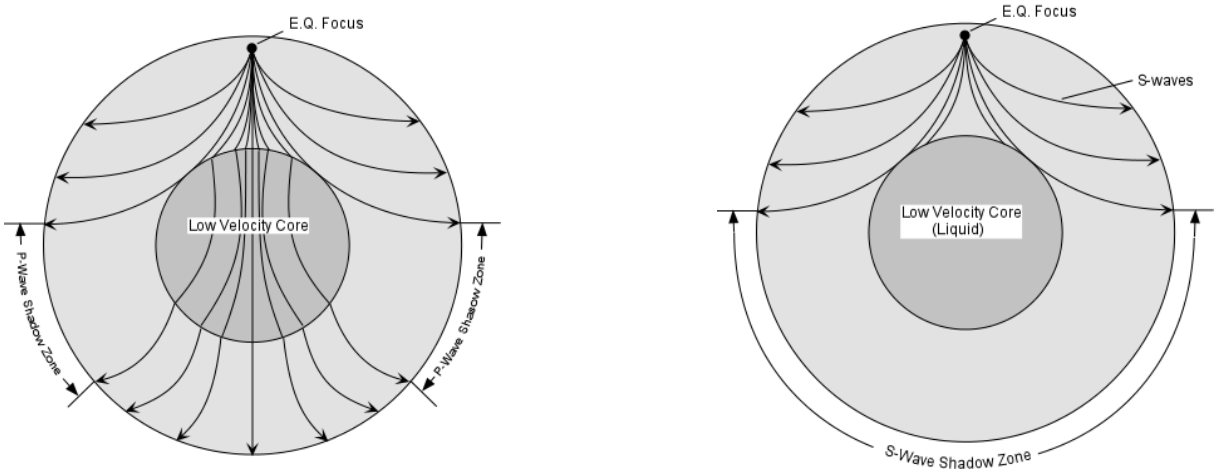

Figure 2.21: Illustration of the propagation of S-waves and P-waves in the deep Earth. Note that the inner core is not represented. Figure taken from Ref. [165].

local inhomogeneities are known to exist in its interior, such radial models are used as reference for various purposes. This the case of the Preliminary Reference Earth Model (PREM) [116], which was built following a global effort of the geophysics community in the late 1970s, and published in its final form in 1981. The PREM is based on the inversion of a large set of seismic data including eigenfrequencies for the whole-Earth free oscillation modes as well as dispersion and travel time data for both surface waves and body waves (the S- and P-waves mentioned previously). In addition, the radius, total mass and moment of inertia of Earth were included as integral constraints on the mass distribution. The original PREM is a parametric model, using low-order polynomials to describe radial functions in 9 predefined layers. Discretised versions are often used as well (as in Fig. 2.11 and the rest of this thesis). More recent one-dimensional velocity models are available. In particular, the model dubbed *AK135-F* [166, 167] includes a radial density profile, which differs from the PREM in the upper mantle by few % density discrepancies, and in the core by the precise location of the inner core boundary. A short comparison of PREM and AK135 from the methodological point of view (data sets, hypotheses, *etc*) can be found in Ref. [168]. The radial density profiles are also compared in Ref. [169].

The uncertainties associated to the determination of the velocities, matter densities or elasticity moduli were not evaluated in the original PREM [116]. The precision of density estimation has been discussed by several authors since then [170–172]. In general, densities are much less well constrained than the seismic velocities. Indeed, the only direct constraints on the mass distribution are the astronomic-geodetic parameters (mass and moment of inertia). The information on densities from body waves data is doubly indirect: velocities are obtained via nonlinear inversion of the travel times, and the densities and elasticity moduli can only be inferred from the velocities via model-dependent assumptions (Eq. 2.33 is used together with an additional relation, which can be empirical or follow from thermodynamical assumptions [170]). In currently available data, the most useful information comes from the lowest-order free oscillation modes, which are sensitive to density through self-gravitation effects induced in deformation [171]. The related inversion problem is however strongly non-linear and the use of perturbative methods is generally necessary. Significant discrepancies are observed between the density profile inferred from normal mode data and body waves data, especially in the outer core [173].

All in all, various estimates of the precision of the density profile $\rho(r)$ in the PREM and later models are put forward in the literature. In a 1991 review of the subject [170], B. Bolt concludes that density values averaged over 100 km are 'probably' known to a 5% accuracy at nearly all depths, while density gradients over the same length scale are not well constrained except in most of the outer core; furthermore he advocates the need for further investigation of the

covariances and resolution bounds on $\rho(r)$. In Ref. [171], the PREM density profile is reassessed in light of new free oscillation data including error estimates, using the nonlinear inversion approach and a random sampling method. Results indicate a good accuracy of the PREM model within the limits of its parametrisation choices. In Ref. [172] a resolution kernel analysis is carried out within the classical perturbative method: it is found that the average density in the lowermost 500 km of the mantle is constrained within 0.4 %. In both these studies, error estimates are however not easily interpreted from a probabilistic point of view. Very recently a new model of elastic parameters restricted to the outer core has been obtained from state-of-the-art free oscillation data [173]. Uncertainties are evaluated with a classical Bayesian inference approach. The outer core density is found to be approximately 1% larger than in the PREM, with an estimated uncertainty of about $\pm 0.5\%$ at 1σ confidence level. Another study involving uncertainty estimates for the whole-Earth model using a Bayesian inversion method was performed in Ref. [174].

Finally, it should be emphasized that numerous examples of local deviations from radially symmetric reference models are known to exist and have been actively studied since the advent of the PREM. In the deep Earth, the most prominent large-scale lateral inhomogeneities are the so-called *large low-shear-velocity provinces* (LLSVPs) [175]: regions of low S-wave velocity in the lowermost mantle, extending over thousands of kilometres laterally and up to ~ 1000 km above the CMB. There are two main such regions (also referred to as *superplumes*), below Africa and the Pacific ocean. The nature and origin of LLSVPs remain ambiguous. On a lower spatial scale (~ 10 km thick), zones of 'ultra-low' S-wave velocity and reduced P-wave velocity are found at the core-mantle boundary, mostly correlated with the edges of LLSVPs; these are referred to as *ultra-low velocity zones* (ULVZs). In the past decade an effort has been undertaken towards the development of a reference 3D seismological model [176].

Chemical composition

Global Earth models intend to describe both the chemical composition and the physical properties of all layers. They are generally closely related to a theory of the formation of Earth accounting for cosmochemical data and information from geodynamics and geomagnetism. These models can be tested with respect to the seismological measurements by relating composition to mechanical properties via an equation of state, obtained from a combination of theoretical 'ab initio' calculations and laboratory experiments simulating the very high pressure and temperature in the deep Earth. Both theoretical and experimental approaches are challenging.

The largest uncertainties in terms of global chemical composition concern the outer core. As first proposed by Birch [177], comparison between the PREM and high-pressure/high-temperature (HP/PT) properties of Fe-Ni alloys shows that the density jump at the inner core boundary (ICB) as well as the seismic velocities in the core requires the presence of a few percent of light elements. Si, O, S, C and H (and some of their combinations) are the most popular elements that have been considered so far [178, 179]. But the precise chemistry and amounts of the light element(s) involved cannot be fully determined based on seismology and HP/HT experimental petrology or ab initio calculations alone (see *e.g.* [180, 181]), and models of core composition must rely also on specific scenarios of Earth formation [182]. However, models favoring Si for example, can be based both on heterogeneous accretion [183] or on the at odds hypothesis of homogeneous accretion of materials parented to Enstatite chondrites [184]. Amongst the different light elements, H has received a renewed interest in the past years based on HP-HT experiments that confirmed the possibility to put a significant amount of H in the core [185]. H is the most abundant element in the proto-solar nebula, and it has been argued that the incorporation of a few percent of H in the core was made possible through the accretion of H_2O-bearing materials [186]. Furthermore, the effect of H on density and seismic velocities is such that about 1 wt% of H in the outer core could be enough to fit PREM [187], which would

in turn rule out the presence of other light elements. Up to now however there has been no method available to directly constrain the amount of H in the core.

2.3.2 Neutrino absorption tomography

The possibility that high-energetic neutrinos be absorbed by Earth's matter was already mentioned by Greisen as a concluding remark of the 1960 review where the water-tank Cherenkov detection of cosmic neutrinos was first discussed [111]. The neutrino interaction cross-section needs to be $\mathcal{O}(10^{-34}\text{cm}^2)$ per nucleon in order for neutrino absorption in the earth to be significant, *i.e.* for the neutrino mean free path to become comparable with Earth's radius. Based on an extrapolation of Fermi's theory, Greisen estimated the required neutrino energies as $E_\nu \sim 10^{13}\,\text{eV}$ and expressed the possibility, which he qualified as 'fanciful', that the cosmic neutrino flux may be sufficient to ultimately detect its absorption by the earth.

The first proposals to actually employ neutrinos to probe the structure of Earth considered as a source an artificial beam of TeV energy neutrinos [188–190]. Because neutrino absorption is not sensitive to the shape of the matter structure along the neutrino path but only to the integrated density, a single baseline is not sufficient to obtain structural information. In order to perform 'whole-Earth' tomography, it was thus proposed to move the detector around Earth's surface – with the additional challenge to design a steerable neutrino beam (see Fig. 2.22(a)).

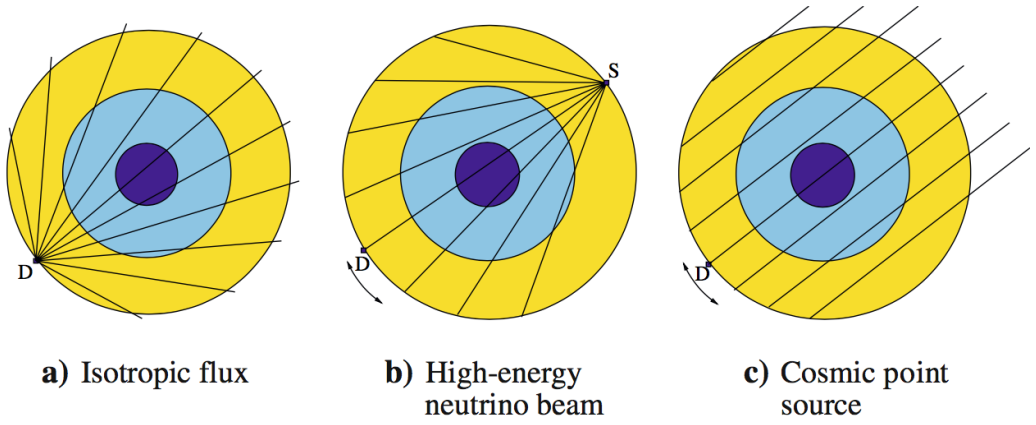

Figure 2.22: Possible source-detector configurations for whole-Earth absorption tomography using beam, atmospheric or cosmic neutrinos. The lines represent different baselines along which the attenuation of the neutrino flux is measured. A similar categorisation apply to sources for neutrino oscillation tomography. Taken from Ref. [191].

Later on, the 'neutrino sky' was again discussed as a potential source. As represented on Fig. 2.22, this one can in fact be subdivided into two cases. In the first case, the flux is isotropic and therefore moving the detector is not required. A good resolution on the incoming direction of detected neutrinos is necessary, but can be considered realistic at TeV energies. The source can be realised either as a diffuse astrophysical flux [192, 193] or as the high-energy tail of the flux of atmospheric neutrinos [193, 194]. In the second case, the flux from a single astrophysical point source of neutrinos reaches the detector after traversing Earth along a baseline which varies in time due to Earth's rotation [195, 196].

In addition to whole Earth tomography, the feasibility of using neutrino beams to search for geological deposits (oil, gas) has been evaluated, notably considering a surface-based acoustic detection technique [197]. A review of the potential of neutrino absorption tomography can be found in Ref. [191], where the various options mentioned here are further discussed and compared. Although it appears clear that the density measurements performed with this method will not reach a precision comparable to that of standard geophysical techniques in the near future, it has the advantage of accessing the mass (nucleon density) distribution in a more direct manner.

Moreover, the absorption measures the integrated density along the neutrino path and is thus equally sensitive to all Earth layers, unlike for instance body waves which are attenuated in the mantle and partly reflected at the CMB.

The first whole-Earth tomography analysis from real neutrino data has recently been performed by A. Donini *et al.* (Ref. [198]). They used a 1-year sample of atmospheric neutrino events collected at the IceCube neutrino telescope during 2011-2012 and publicly released by the experimental collaboration in 2016 [199]. This dataset (known as IC86) contains 20145 up-going muon events with reconstructed muon energies between 400 GeV and 20 TeV. Fig. 2.23 shows the zenith angle distribution the events, normalised to the expectation without Earth attenuation. The effect of Earth attenuation can be clearly seen in the data, although the statistical errors are very large. Considering events with energies above $\sim 5\,\text{TeV}$, up to 50% of the flux is expected to be absorbed, but the available statistics with a one-year sample is still low.

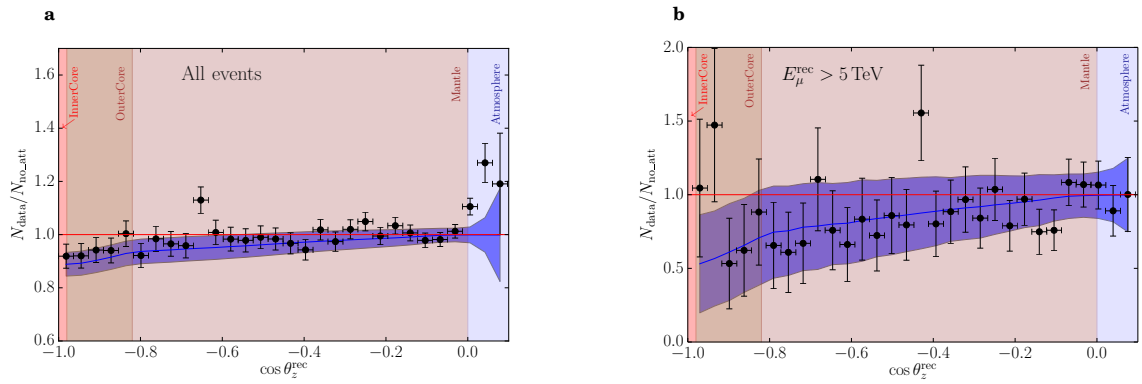

Figure 2.23: Ratio of the number of observed events in the IC86 sample to the number of expected events without including Earth attenuation. Left: zenith distribution of the ratio, including all events in the IC86 sample. Right: zenith distribution of the ratio, but only considering events with a minimum reconstructed muon energy of 5 TeV. In both panels, the solid blue line represents the expectation using the PREM for the density profile, with the expected statistical error represented by the blue band.

In the analysis the density profile is parametrised by a set of five concentric layers of constant density and of approximately equal thicknesses (both the outer core and mantle being subdivided in two). The five densities are fitted, without further geophysical constraint, to the observed event counts binned in terms of reconstructed muon energy and zenith angle. Bayesian credible intervals are obtained for the total mass and moment of inertia of the planet, as well as the density in each of the five layers. These results are compatible with the conventional values (PREM), although with large uncertainties. The difference between the average densities in the core and mantle is thus measured at

$$\overline{\rho}_C - \overline{\rho}_M = 13.1^{+5.8}_{6.3}\,\text{g/cm}^3 \tag{2.34}$$

The core is thus established as denser than the mantle with a p-value evaluated at 0.011.

With the increase of the IceCube dataset and the start of operation of KM3NeT, statistical errors could be reduced to a few % in less than a decade. The control of systematics may then become challenging, due to the limited knowledge of the atmospheric neutrino flux at these energies[13]. Obviously, uncertainties in neutrino-nucleon cross-sections are also of prime importance. While the Standard Model prediction for neutrino-nucleon cross-section beyond the TeV is considered accurate at the 5% level, the only available measurement has recently been performed by the IceCube collaboration [200] (using another dataset of comparable size from 2009-2010) and exhibits a large uncertainty ($\sim 40\%$). For the cross-section analysis, which is conceptually very similar to absorption tomography, the dominant systematics were related to

the atmospheric flux and the detection method (optical properties of ice, photosensor efficiencies and reconstruction performances).

2.3.3 Oscillation tomography with atmospheric neutrinos

Absorption tomography requires neutrinos in the TeV range and above, and one of the main limitations is the small flux of existing sources. Much larger fluxes are available at lower energies, where the propagation of neutrinos is also sensitive to the matter profile via a totally different mechanism, the MSW enhancement of flavour oscillations. From a practical point of view, two main differences with absorption tomography can be identified. First, flavour oscillations are sensitive to the electron density rather than directly to the mass (nucleon density), introducing a dependence on chemical composition via the electron-to-nucleon ratio. Second, while attenuation only measures the total integrated density along the neutrino path, oscillations are in principle sensitive to the shape of a density profile. Indeed, the total evolution operator for a sequence of constant density layers (Eq. 1.59) may be different depending of the ordering of the layers due to the non-commutation of the evolution operators of pairs of individual layers [46]. This effect however remains subdominant in the oscillation probabilities, which are at first order sensitive to the baseline-averaged density.

Matter effects both in the solar and atmospheric oscillation regimes are potential probes of Earth's interior, provided suitable neutrino sources and detection methods exist in the corresponding energy range. As in the case of absorption tomography, potential sources can be categorised as in Fig. 2.22. To probe the solar resonance, neutrino energies $E \sim 0.1 - 1\,\mathrm{GeV}$ would be optimal. Solar and supernova neutrinos (case of cosmic point sources, Fig. 2.22c) have been discussed by several authors [201–205]. Tomography of the crust with MeV neutrino beams ('β beams') was also considered more recently [206, 207].

For solar neutrinos, $E_\nu < 20\,\mathrm{MeV}$ so that matter effects occur away from the resonance. Although the hypothesis of small matter effects leads to useful analytical simplifications for the matter profile inversion problem, this regime induces small modifications of oscillation probabilities [204, 205]. While the overall influence of matter effects is clearly observed as a 'day-night' asymmetry of the solar neutrino event rates [36, 208], the effect is not strong enough to allow for precise measurements of Earth's electron density with current experiments. Moreover, the contribution of remote density structures (such as discontinuities in the core) to the measured oscillation probabilities is attenuated by the finite energy resolution of a realistic detector [209]. The potential of the high-precision detection technique developed for the DUNE experiment for a precise study of Earth's matter effect on solar neutrino oscillations has recently been evaluated [210]. The results show that the detectable effect of the density structure will at best resolve the two density jumps closest to Earth's surface, at a few tens of kilometres depth. The energy spectrum of supernova (SN) neutrinos extends to energies closer to the solar MSW resonance, and the neutrino burst from a closeby SN, if detected by two conveniently positioned detectors, may produce a sufficiently large signal to measure the density of Earth's core with good precision [203]. A special feature of solar and supernova neutrinos is that due to the loss of coherence of neutrino wave packets while propagating along astrophysical distances, upon reaching Earth these neutrinos can practically be considered as mass eigenstates [37], with interesting consequences for the purpose of neutrino tomography [204].

In the atmospheric regime, matter effects for Earth densities are the strongest in the interval $E_\nu \sim 2 - 8\,\mathrm{GeV}$, which can be probed using both artificial neutrino beams and atmospheric neutrinos. While the focus of this thesis is on atmospheric neutrinos, considering first the case of single baselines is instructive. Moreover in the past two decades the detailed analysis of

[13] The contribution of prompt charm decays to the total atmospheric neutrino flux becomes significant above TeV energies. In the PeV range and above this so-called 'prompt component' is expected to dominate the 'conventional component' described in Sec. 2.1.1. The prompt component is affected by large theoretical uncertainties.

degeneracies between oscillation parameters and the electron density profile for very long baseline accelerator experiments (in particular before the measurement of θ_{13}) produced theoretical insights relevant to neutrino oscillation tomography in general [46, 211–216].

Single baseline tomography in the GeV range

Mathematically, inferring an unknown density profile from an oscillation pattern $P_{\text{osc}}(E)$ at a given baseline is a difficult inversion problem. It becomes tractable for solar and SN neutrinos in Earth matter (away from the solar resonance), where a perturbative approach can be used [217]. However the approximation does not hold for GeV neutrinos. One possibility is to consider a simplified profile parametrised with a limited number of parameters: the inversion then reduces to a numerical optimisation problem. In a constrained case, for example when considering small variations from a reference profile, standard minimisation methods (as routinely used in particle physics analyses) can be sufficient. With this approach the possibility to detect and locate small density contrasts in the crust (water-filled cavity, oil and gas reservoir, *etc.*) using a high-luminosity beam is evaluated in Ref. [218], demonstrating the capability of the technique to obtain positional information. Exploring more general profiles with the same method amounts to solving a higher-dimensional minimisation problem. In the unconstrained case, global minimisation becomes more challenging but can be attempted for instance using non-deterministic algorithms. This method is employed in Ref. [219] to study the resolution of structure and edges in Earth's density profile using a single core-crossing baseline. It is shown that a PREM-like profile can be inferred from neutrino oscillations without any prior information except the hypothesis of radial symmetry. However, degeneracies between candidate profiles remain and the precision of the method is limited.

The above example illustrates an important limitation of single-baseline oscillation tomography: structures much smaller than the oscillation length in matter cannot be well resolved. For instance, considering an electron density profile of the form

$$n_e(x) = N_e + \eta(x), \quad \eta(x) = A\cos\frac{2\pi x}{\lambda}, \tag{2.35}$$

with $A \ll N_e$, it can be shown using perturbation theory that in the two-neutrino oscillation probabilities the term arising from the perturbation $\eta(x)$ vanishes in the limit [219]

$$\lambda \ll L_{\text{osc}}(N_e, E_\nu) = \frac{4\pi E_\nu}{\xi(N_e, E_\nu)\Delta m^2}\,. \tag{2.36}$$

Considering that a given density is optimally probed by neutrinos around the corresponding resonance energy, with a single baseline a resolution of a few hundred kilometres can be expected at best, since $L_{\text{osc}}(E_{\text{res}}) = \mathcal{O}(10^4\,\text{km})$ for both mantle and core densities. Although the oscillation length decreases above the resonance energy ($\xi \to \infty$), this regime also corresponds to a suppression of the matter-induced oscillations.

More generally, when considering the Fourier series decomposition of the electron density profile $n_e(x)$ along a baseline L with average density $\overline{n}_e$, truncating the series after the first few modes is sufficient to reproduce the oscillation probabilities with very good accuracy. In fact, the contribution of the n-th mode (corresponding to a length scale L/n) to the evolution operator is suppressed by a factor $1/n$ [220]. In oscillation probabilities $P_{\text{osc}}(E)$, this mode selectively impacts the n-th oscillation minima on either side of the resonance energy $E_{\text{res}}(\overline{n}_e)$ [127]. From the above considerations, the achievable spatial resolution should in principle be better for core densities. However, realistic sensitivities depend strongly on the characteristics of the detector, as illustrated by the attenuation effect discussed previously for solar neutrino tomography [209].

While the resolution of structures (*e.g.* LLSVPs or precise localisation of the CMB or inner core boundary) thus appears out of reach of single baseline tomography, the measurement of av-

erage densities can be rather precise. For instance it is shown in Ref. [221] that a vertical neutrino beam produced at a future 'neutrino factory' could measure the inner core average density $\overline{\rho}_{IC}$ at the percent level precision. Moreover, it is pointed out that the measurement of line-averaged densities is very interesting when targeting Earth's inner layers, since the weight of each Earth layer in this quantity is proportional to its radius, with for the inner core $R_{IC}/R_{\text{Earth}} \simeq 19\%$. In comparison, the constraints on $\overline{\rho}_{IC}$ from the total mass and moment of inertia are much weaker since the inner core represents only about 0.7% of Earth's volume and 3.7% of $\int r^2 \mathrm{d}V$.

Atmospheric neutrinos as a source

The limitations of single-baseline tomography for resolving Earth's structure would be overcome by the combination of a steerable neutrino beam and a movable detector able to scan the whole Earth – see Fig. 2.22(b). Unfortunately, from the technological point of view this appears unlikely to become feasible in the upcoming decades. A geometrically equivalent scan may however be performed with a single detector and a distributed neutrino source, as seen in Fig. 2.22(a). Such a source is provided 'for free' by the production of neutrinos in the atmosphere, which does cover the relevant energy range in a close to isotropic manner. While no existing experiment has a sufficient sensitivity to matter-enhanced oscillations to realistically access the electron density in the deep Earth, the next generation of megaton-scale atmospheric detectors will have such a potential [154, 169, 222].

To some extent the coexistence of muon and electron flavours and of neutrinos and antineutrinos in the flux is a limitation with respect to a pure ν_μ beam[14]. Nevertheless, a few % variation of the average density in a whole layer has a sizeable effect on the atmospheric rates, as is shown on Fig. 2.24 for the lower mantle and in Fig. 2.25 for the outer core. In those figures each event rate curve corresponds to a uniform variation of the electron density in a given layer with respect to the PREM, by a factor $f = \{0.90,\ 0.95,\ 1.0,\ 1.05,\ 1.10\}$. A normal mass hierarchy is assumed in both cases. Although the precision on f should reach $\sim 1\%$ to become competitive with the existing precision from geophysics, larger values are chosen here to illustrate the potential of oscillation tomography. Note that an inverted mass hierarchy (not shown here) would be less favourable since the MSW resonance would then occur for antineutrinos whose interaction cross-section is smaller by a factor ~ 2.

Based on single-baseline rates of atmospheric neutrino events with respect to energy, the sensitivity to the average electron density thus appears promising in both the deep mantle and the outer core. This must however be put in perspective with caveats related to the angular event distribution. First, unlike in the case of an artificial beam, the incoming direction of an atmospheric neutrino interacting in the detector cannot be assumed and has to be reconstructed from secondary particles – and so is the energy. In neutrino telescopes, the accuracy in angular reconstruction is essentially limited by the intrinsic kinematic smearing (misalignment of the momenta of the incoming neutrino and outgoing lepton), which increases with decreasing energy and becomes $\mathcal{O}(10°)$ for few-GeV neutrinos. Errors in zenith angle translate as errors on the baseline estimation, which can obviously spoil the measurement of a given layer's density. Fast oscillatory patterns in L/E may even be averaged out completely in case of large errors on E and/or L. The second caveat concerns the angular distribution of the flux of atmospheric neutrinos. Assuming for simplicity that neutrino production in the atmosphere is isotropic, the flux received at the detector is uniformly distributed in solid angle Ω, so that the angular

[14] It should not be overlooked that besides 'conventional' beams producing an (almost) pure ν_μ beam from pion decays in flight, two alternative methods have been investigated in great detail in the last two decades and discussed for Earth tomography: i) neutrino factories, producing ν_e and $\overline{\nu}_\mu$ in equal intensities from the decays of accelerated muons circulating in a storage ring; ii) beta-beams, producing purely $\overline{\nu}_e$ from the β-decays of stored ions.

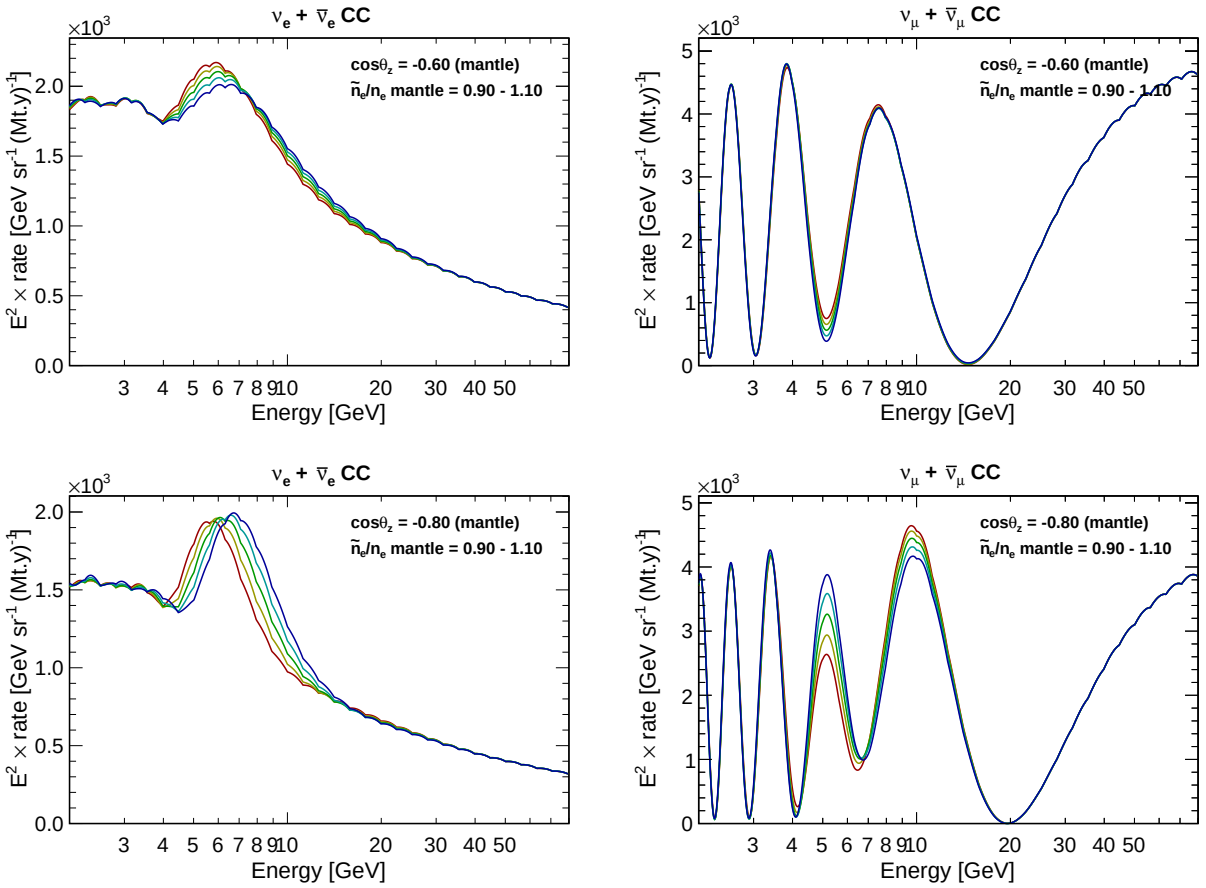

Figure 2.24: Effect of a variation of the average electron density in the lower mantle on the $\nu + \overline{\nu}$ event rate at an atmospheric detector (NH is assumed). Starting from the PREM profile, the electron density is multiplied by a uniform factor f in the whole lower mantle, while it remains constant in the other layers. The values $f = 0.90,\ 0.95,\ 1.0,\ 1.05,\ 1.10$ correspond respectively to the blue, cyan, green, yellow and red curves.

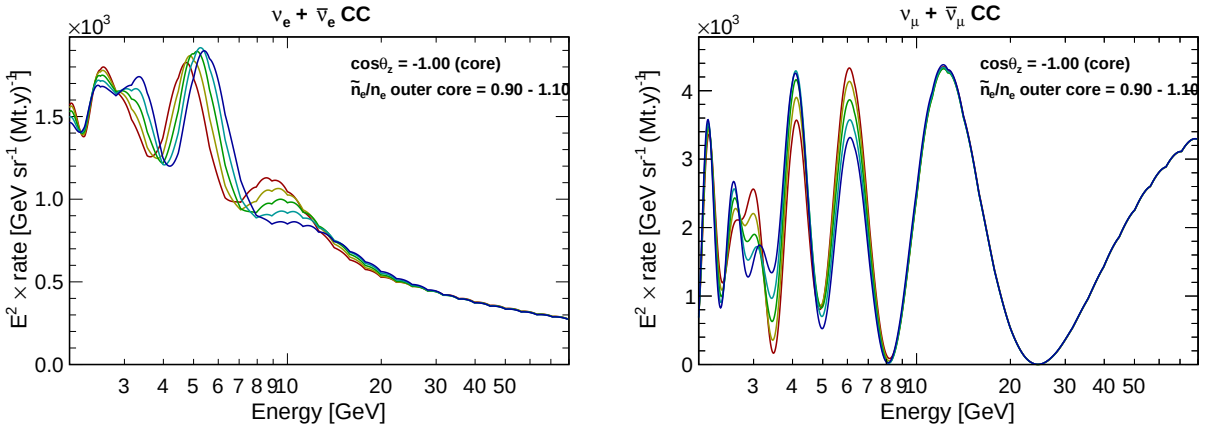

Figure 2.25: Effect of a variation of the average electron density in the outer core on the $\nu + \overline{\nu}$ event rate at an atmospheric detector (NH is assumed). Starting from the PREM profile, the electron density is multiplied by a uniform factor f in the outer core, while it remains constant in the other layers (inner core and mantle). The values $f = 0.90,\ 0.95,\ 1.0,\ 1.05,\ 1.10$ correspond respectively to the blue, cyan, green, yellow and red curves.

distribution is flat in terms of both $\cos\theta_z$ (cosine zenith) and φ (azimuth angle):

$$\frac{\mathrm{d}N}{\mathrm{d}E\,\mathrm{d}\Omega} = \frac{\mathrm{d}N}{\mathrm{d}E\,\mathrm{d}(\cos\theta_z)\,\mathrm{d}\varphi} = f(E). \tag{2.37}$$

The fraction of the upgoing neutrinos received at the detector whose trajectories cross a shell defined by ($R_{\min}$, $R_{\max}$) is then given by $(\cos\theta)_{\max} - (\cos\theta)_{\min}$ (see Eq. 2.28). This yields isotropic flux fractions of about 2% for the inner core, 14% for the outer core, 40% for the lower mantle and 44% for the upper mantle and crust. In addition, as seen in Sec. 2.1.1 the real atmospheric flux is stronger close to the horizontal direction, which further decreases the fraction of upgoing neutrinos crossing the deep Earth. Note however that this issue is purely statistics-related and can be overcome by increasing the total exposure (detector size × time), whereas the limitation related to angular resolution is a more fundamental one.

Other questions adressed in this thesis regard the impact of detector effects such as low-energy efficiencies, energy resolutions and particle identification performance; the effect of oscillation parameters and other systematic uncertainties (*e.g.* flux, cross-section) on a measurement of the electron density; and conversely, the importance of electron density uncertainties for the NMH measurement.

Chemical composition sensitivity

Although not as performant as in the outermost layers, the determination of $\rho(r)$ averaged over $\mathcal{O}(100\,\text{km})$ by geophysical methods remains accurate at the sub-percent level in the deep Earth, and reaching a similar performance with the upcoming generation of GeV-focused neutrino telescopes is not realistic.

Nevertheless, the direct sensitivity of neutrino oscillation tomography to the electron density profile $n_e(r)$ can provide complementary information inaccessible to geophysics, which has the potential to constrain compositional models when combined with the existing information on $\rho(r)$ [169]. Indeed, in a material composed of the chemical elements X_i with respective weight fractions w_i (usually expressed in 'wt%'), the electron number density (m^{-3}) can be expressed as

$$n_e = \left(\sum_i w_i \frac{Z_i}{A_{r,i}}\right) \times \frac{N_A}{10^{-3}\,\text{kg.mol}^{-1}} \times \rho \tag{2.38}$$

where Z_i and $A_{r,i}$ are the atomic number and standard atomic weight[15] of element X_i, N_A is the Avogadro number, and ρ is the mass density in kg.m^{-3}. The comparison of ρ and n_e thus gives access to the average 'proton-to-nucleon' ratio of the medium $\sum w_i Z_i/A_{r,i}$, generically denoted as Z/A in the following.

The values of Z/A for a selection of chemical elements and compounds are listed in Tab. 2.2. Most of the chemical elements found in the minerals entering in the composition of the silicate Earth are light elements with $Z/A \simeq 0.5$, whereas heavier elements tend to be more neutron-rich ($Z/A < 0.5$). This is the case of the primary components of Earth's core, iron and nickel. As introduced in Sec. 2.3.1, the nature and amount of light elements present in the outer core besides the Fe-Ni alloy remains uncertain. Si, O, S, C and H as well as some of their combinations have been considered. The presence of light elements implies an increased average Z/A ratio with respect to a pure iron or Fe-Ni composition. An overview of compositional models and corresponding ratio can be found in Ref. [169]. A measurement of the average Z/A, even infinitely accurate, would not identify a compositional model uniquely. Nevertheless oscillation

	${}_{1}\text{H}$	${}_{8}\text{O}$	${}_{12}\text{Mg}$	${}_{13}\text{Al}$	${}_{14}\text{Si}$	${}_{20}\text{Ca}$	${}_{26}\text{Fe}$	${}_{28}\text{Ni}$
Z/A	0.992	0.500	0.494	0.482	0.498	0.499	0.466	0.477

Table 2.2: Values of Z/A for some common elements in Earth's composition.

[15] Note that the standard atomic weight $A_{r,i}$ reflects the isotopic abundance of element X_i in Earth's crust and atmosphere, which may not necessarily be extrapolated safely to the deep Earth. However, variations of chemical composition affect the average Z/A more strongly than isotopic uncertainties.

tomography has a potential to help in excluding the most extreme options, such as hydrogen-rich models discussed in Sec. 2.3.1.

Having this application in mind, the factor f rescaling the electron density in each main layer, with respect to a reference proton-to-nucleon ratio of 0.5 and PREM-based $\rho(r)$ profile, is referred to as Z/A throughout the thesis. This convention corresponds to assigning zero uncertainty to $\rho(r)$ and assuming uniform chemical composition in each layer. In reality, both $\rho(r)$ and $[Z/A](r)$ are currently known with finite accuracy, and variations in chemical composition may affect both variables. Moreover local inhomogneities breaking radial symmetry are known to exist. Strictly speaking an analysis of atmospheric neutrino oscillations by itself only provides an estimate of $n_e(x)$ along Earth-crossing baselines in the detector 'field of view', and this information then has to be interpreted in light of geophysical knowledge.

Chapter 3

KM3NeT/ORCA: detector, simulation and reconstruction

Contents

After a short introduction to the KM3NeT project, Sec. 3.1 describes the design and principles of operation of the ORCA undersea detector. The main technological components (photomultipliers, digital optical modules, detection unit) and the layout of the detector are first depicted. The detector deployment status and plans are quickly discussed. The most important sources of background and the methods for processing and filtering the data are introduced. Simulated datasets are a central ingredient of this work: the simulation methods are detailed in Sec. 3.2, from the generation of neutrino interactions and atmospheric muon background events up to the simulation of the photomultiplier response. The specifics of the sample of Monte Carlo events used in this work are presented and motivated. Finally, Sec. 3.3 discusses the strategies for selecting signal events and discarding background, reconstructing neutrino interactions from the observed patterns of Cherenkov light, and classifying reconstructed events for the purpose of analysing neutrino flavour oscillations.

3.1 The ORCA detector

The ORCA detector configuration studied in this work consists in one building block of the KM3NeT infrastructure – the next generation neutrino telescope in the Mediterranean. The ORCA building block instruments a total mass of seawater of about 8 megatons at a depth of 2450 m.

In this section the KM3NeT project is first introduced with its components ARCA (focused on neutrino astronomy) and ORCA (oscillations). The technology and fundamental detector components (photomultiplier, optical module, detection string) common to both detector designs are briefly depicted. The planned detector layout and data processing systems for ORCA are then presented briefly.

3.1.1 The KM3NeT project and detector technology

The KM3NeT collaboration[1] has undergone the construction of a network of deep underwater Cherenkov neutrino detectors in the Mediterranean Sea [21]. The objective pursued by KM3NeT is twofold: discovering and characterising high-energy neutrino sources in the Universe, and measuring the oscillations of Earth-crossing atmospheric neutrinos with the primary goal of determining the neutrino mass hierarchy. Dedicated building blocks, using the same detection technology but different detector layouts, are being built in two different locations offshore Italy and France. The KM3NeT project builds upon the expertise acquired by the NESTOR and NEMO prototypes as well as the ANTARES neutrino telescope, in operation since 2008 [20]. Each building block of the KM3NeT infrastructure comprises 115 vertical lines anchored to the seabed and connected to the shore through an electro-optical cable. This network supports a three-dimensional array of 2070 spherical optical modules, each one containing 31 small photomultipliers (PMTs) facing towards all directions.

The astronomy-focused component of the project is called KM3NeT/ARCA, where ARCA stands for *Astroparticle Research with Cosmics in the Abyss*. The detector is being deployed off the East coast of Sicily, offshore Capo Passero, at a depth of 3500 m. Looking downwards from the Northern hermisphere with an instrumented mass of ~ 1, ARCA's field of view will encompass a large portion of the Galactic Plane, including the Galactic Centre, and will be complementary to that of the IceCube observatory located at the South Pole. ARCA shall provide an independent confirmation of the high energy flux of astrophysical neutrinos observed by IceCube [25] and will further seek to discover point sources of neutrinos. Thanks to the very good optical properties of seawater, ARCA will have excellent pointing accuracy, with an angular resolution reaching 0.2° for muon neutrinos with energies above 10 TeV [21]. A rich multi-messenger program is also foreseen.

The second detector, dubbed ORCA (*Oscillation Research with Cosmics in the Abyss*), is located at a depth of 2450 m about 40 km offshore Toulon (France). The exact coordinates are 42° 48′ N, 06° 02′ E which is about 10 km west of the ANTARES site. Both locations are shown on Fig. 3.1. The ORCA detector, begin optimised for the detection of $\sim$ GeV neutrino, is much more densely instrumented than ARCA. In addition to neutrino astronomy and oscillation physics, KM3NeT will also provide a platform for Earth and Sea science measurements in the deep-sea environment.

The detection technology, precise layout of the detector and data acquisition systems are described in Sections 3.1.1 and 3.1.3. As for the physics reach of ORCA, it is one of the main topics of this work. The sensitivity results for neutrino oscillation measurements are presented in Chapters 8 and 9.

[1] KM3NeT is short for 'Cubic Kilometre Neutrino Telescope'. The KM3NeT collaboration, as of February 2017, consists of more than 250 participants from 51 universities and research institutes in 15 different countries. Most member institutions are located in France, Italy, Germany, Spain, The Netherlands, and Greece [223].

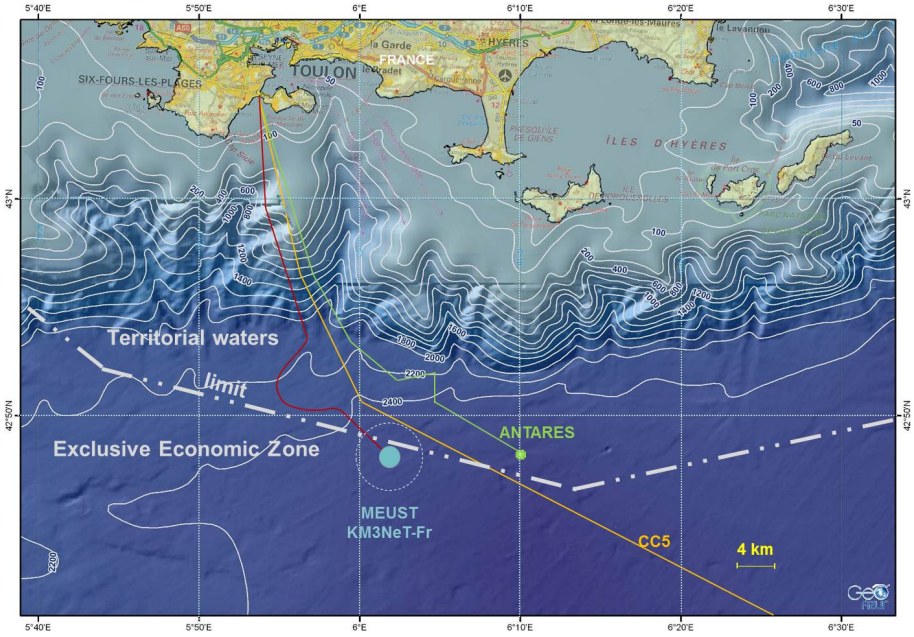

Figure 3.1: Map of the Mediterranean Sea south of Toulon, France. The respective locations of the KM3NeT-France (MEUST) and ANTARES infrastructures are indicated as well as the approximate paths of the electro-optical cables running from the detector sites to the shore station in La Seyne-sur-Mer. The acronym MEUST refers to the *Mediterranean Eurocentre for Underwater Sciences and Technologies* [224]. Figure taken from Ref. [21].

The Digital Optical Module

The detection technology is common to both ARCA and ORCA. The key technical component is the Digital Optical Module (DOM). The DOM consists of a pressure-resistant glass sphere of 43 cm diameter, housing 31 3-inch Hamamatsu photomultipliers (PMTs) together with their associated readout electronics and sensors. Pictures of the DOM and PMT are shown in Fig. 3.2. Photomultipliers are pointing towards all directions with close to uniform angular coverage. They are positioned in a 3D printed support, in 5 rings of 6 equally spaced PMTs plus a single PMT pointing vertically downwards, for a total of 12 PMTs in the upper hemisphere and 19 in the lower hemisphere. A reflector ring around the face of each PMT is added to increase the photon collection efficiency, and an optical gel fills the space between the PMT photocathode, reflector ring and glass sphere to ensure optical contact.

The multi-PMT design is a distinguishing feature of the KM3NeT optical module. While traditional designs typically use single larger PMTs, as done for instance in ANTARES or IceCube, the segmentation of the photocathode area in the multi-PMT design is very advantageous. The directional information of the photon arrival direction can be used for event reconstruction, and the ability to detect multiple photons with high efficiency greatly improves the optical background rejection capabilities. In addition the overall cost per unit of photocathode area is reduced.

Photomultipliers

In a photomultiplier tube, the photodetection is made possible by the application of a high voltage (HV). A photon reaching the photocathode (*photon hit*) from outside the DOM can issue the release of a photoelectron (*p.e.*) via the photoelectric effect. In the tube, containing a succession of dynodes, the high electric field accelerates electrons so that they free more electrons when hitting the next dynode. The initial photoelectron can thus be amplified into an electron cascade, whose charge is measured under the form of a voltage pulse at the final anode. The

Figure 3.2: Photographs of a KM3NeT DOM (left) and Hamamatsu R12199-02 photomultiplier tube (right). In the DOM, the visible part of the individual PMTs comprise the photocathode areas, circled by the silver-coloured reflector rings.

KM3NeT photomultipliers are operated at a gain (amount of electrons after the multiplication stage for a single initial p.e.) of $3 \cdot 10^6$, which is set by tuning the applied high voltage.

In the PMT base, the analog pulse is further amplified electronically and a so-called *PMT hit*. In the following the term 'hit' without precision may be used to refer to PMT hits. is recorded when the amplified voltage exceeds a certain threshold. For each hit, the time of arrival is the instant when the threshold was exceeded and the charge information is summarised by the duration for which the pulse remains above the threshold, as known as 'ToT' (Time over Threshold). This analog signal processing is done in the PMT base (preamplifier, comparator); the time digitisation is left to a time-to-digital converter implemented on an FPGA (*Field Programmable Gate Array*) on the central logic board of the DOM. The transmission of these data is discussed in Sec. 3.1.3.

To suppress electrical noise, the threshold is typically set to 0.3 p.e., *i.e.* at 30% of the height of the mean pulse height for a single initial photoelectron. Both the high voltage (700-1300 V) and the threshold (0.8-2.4 V) of individual PMTs can be set from the shore. Further details on the electronics readout of the KM3NeT DOM can be found in Ref. [225]. The application of the threshold and digitisation of the analog pulse constitute the first level of data filtering, applied at the hardware level, which is conventionally referred to as the 'level 0' filter (L0). Hence the PMT hits are also called 'L0 hits'. In the data acquisition system each L0 hit is characterised by its PMT address, hit time and time-over-threshold (both with a precision of 1 ns).

Understanding the PMT response properties is of key importance for simulating the detector and analysing data. Some of the most essential characteristics are the following:

- optical properties such as the wavelength-dependent quantum efficiency of the photocathode (fraction of the photon hits releasing a photoelectron), and the effective collection area of the PMT and reflector system;
- high-voltage tuning properties (relation between HV and gain in pulse mode);
- spread of the distribution of the transit time of electrons in the tube, which is known as the 'transit time spread' (TTS) and is the principal factor limiting the hit timing accuracy;
- spread of the gain of the electron multiplier stage.

Optical or radiative phenomena in the DOM glass and optical gel may be relevant as well. All these are addressed by the combination of extensive characterisation tests in the laboratory [226, 227], and high-precision simulations [228].

Detection Unit and detector layout

The DOMs are arranged along vertical flexible string structures called *Detection Units* (DUs). Each DU comprises 18 DOMS maintained by two parallel fiber ropes, which are themselves attached to a seafloor anchor. An electro-optical cable running vertically along the strings connects the DOMs to the seabed infrastructure for power supply and optical data transmission. Though not rigid, the ensemble remains close to vertical even in strong sea currents, due to the natural buoyancy of the DOMs and of an additional buoy at the top of the string.

The vertical spacing between DOMs in a DU and the horizontal spacing between DUs are different in ARCA and ORCA since the energy of targeted neutrinos are different. The vertical spacing for ORCA has been optimised for the target physics study (NMH): a 9 m spacing was found to be optimal [21]. In the latest simulated geometry which is studied in this work, realistic technical constraints are accounted for. The precise vertical spacing thus varies slightly between successive storeys (8.7, 9.4, 10.9 m) to accomodate the furling of the string in the deployment module. As to the horizontal spacing, a lower bound of ~ 20 m is imposed as a technical requirement for safe deployment and detector operation. The spacing between two neighbouring strings is distributed around an average of 23 m with a minimum of 19 m. For comparison, in the ARCA geometry the vertical and horizontal spacing are distributed around ~ 36 m and ~ 95 m respectively.

Fig. 3.3 shows the most recent version of the ORCA detector layout. Optical fibres and power supply are gathered in the main electro-optical cables (MEOC) running from the shore to the detector site. ORCA's newly installed MEOC contains 36 fibres and can accomodate for 3 nodes (72 DUs). The MEOC currently in use for ANTARES (24 fibres) will be redirected and used for the remaining DUs of ORCA after the decommissioning of ANTARES. In addition to the optical data stream of L0 hits, the acoustic data for position calibration is transmitted to the shore station by the MEOCs. They also convey in both directions the signals for the remote control and global monitoring of the detector. These concern the nodes, general DU control as well as individual DOM slow control: HV tuning for each PMT, LED beacon, firmware, *etc.*

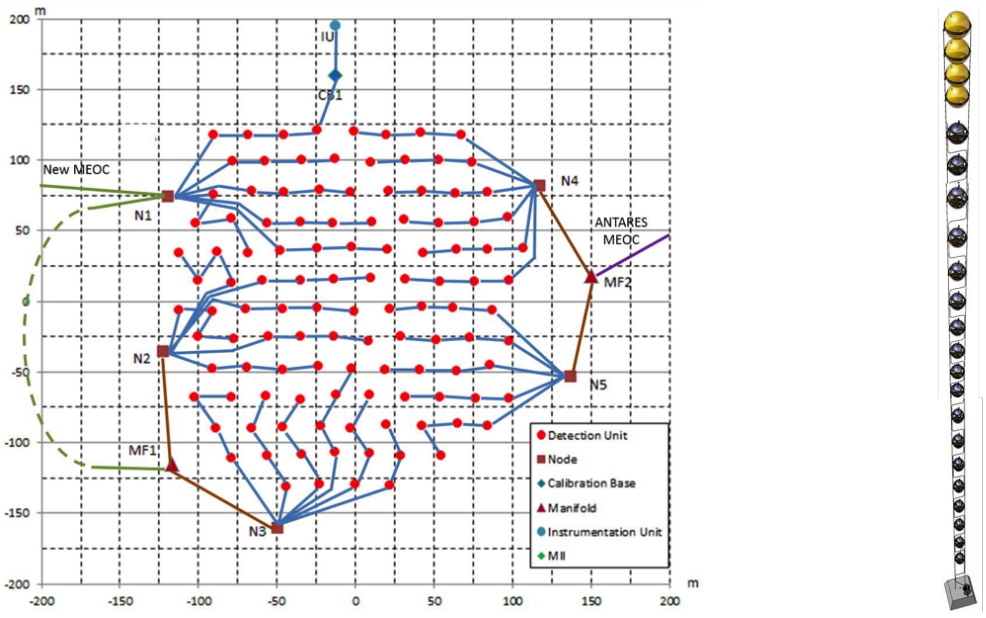

Figure 3.3: **Left:** Footprint of the planned layout of the ORCA detector (top view). The DUs are represented as red dots and the power supply and optical fibre connections between the DU bases as blue lines. Each chain of 4 DUs is connected to a node, where the optical signal is multiplexed. Each node controls in total 24 DUs and is connected either to the next node or to one of the two main electro-optical cables (MEOC). More technical information can be found in Ref. [224]. **Right:** Sketch of a DU. In the ORCA configuration the inter-DOM vertical distance is about 9 m.

Deployment and current status

A dedicated technology has been developed for deploying KM3NeT DUs. The DU is furled around a launcher vehicle (a spherical metal frame), which then descends towards the detector site. After the DU is anchored to the seabed and connected to the infrastructure by a remotely operated vehicle, the launcher unrolls and releases the DOMs one-by-one until it reaches the surface to be recovered and reused. A prototype DOM was installed on an ANTARES instrumentation line in April 2013 and took data for several months [229]. A prototype DU (three DOMs) was then deployed at the ARCA site in May 2014 and operated for about a year [230].

At the ORCA site, the MEOC was connected in December 2014 and the first ORCA junction box deployed in April 2015. After a fault in the MEOC, the junction box had to be redeployed in October 2016. The first ORCA DU was deployed and began taking data in September 2017. The first results of data-taking with a single DU are summarised in Ref. [231]. Data-taking was stopped in December 2017 due to a new fault in the MEOC. As of October 2018, a new cable is being installed and six DUs are ready for deployment, which will restart as soon as the MEOC is operational. The currently secured funding covers more than 50% of the two ARCA and one ORCA building blocks, each building block comprising 115 DUs.

3.1.2 Optical background and calibration

The main source of optical background are the radioactive decays of a potassium isotope, ^{40}K, which is naturally present in seawater. The seawater salinity is measured on-site to be about 3.8%, while the fraction of potassium in the Mediterranean salt is $\sim 1.11\%$ and the isotopic abundance of ^{40}K about $1.17 \cdot 10^{-4}$ [232]. The ^{40}K nuclei decay mostly into ^{40}Ca via β-decay or to ^{40}Ar via electron capture:

$$^{40}\mathrm{K} \rightarrow\ ^{40}\mathrm{Ca} + \overline{\nu}_e + e^- \qquad (89.3\%) \qquad (3.1)$$

$$^{40}\mathrm{K} + e^- \rightarrow\ ^{40}\mathrm{Ar} + \gamma + \nu_e \qquad (10.7\%) \qquad (3.2)$$

The electrons produced in the β-decays have an energy of about 1 MeV, which is sufficient to generate Cherenkov light. In the second decay channel the de-excitation of ^{40}Ar* emits a 1.460 MeV γ-ray (photon) which can release electrons in the medium via Compton-scattering. These electrons may be energetic enough to produce Cherenkov light as well [232].

Radioactive decays occur randomly everywhere in the detection volume. Although several hits from the same ^{40}K decay in the vicinity of a DOM can be detected in time coincidence (a so-called *genuine* coincidence), the vast majority of ^{40}K-induced hits are uncorrelated in space and time. The chance for a combination of such hits to be mistaken for the signature of a few-GeV neutrino event is thus rather low. Nevertheless, the steady background of ^{40}K hits does impact the requirements on data acquisition and triggering and complicate the reconstruction of neutrino events, especially at low energy.

Bioluminescence is another important source of optical background for deep-sea Cherenkov detectors. A variety of marine organisms are bioluminescent, including fish, crustaceans, molluscs, plankton, bacteria, *etc.* [233]. They mostly emit light in the blue and green visible wavelength range to which seawater is the most transparent. The effect of bioluminscence on the measured hit rate is twofold. First, a slowly varying overall rate distributed uniformly in the detector is believed to be due to continuously emitting organisms (*e.g.* bacteria). In addition, *bursts* of bioluminescent light are observed, illuminating local clusters of optical modules for up to several seconds. Such bursts can be extremely bright: while ^{40}K decays generate $\mathcal{O}(10-100)$ Cherenkov photons, bioluminescence bursts can emit up to 10^{13} photons. Bioluminescent activity has been observed to be correlated with the intensity of sea currents, which varies periodically. Seasonal variations are also observed [234, 235].

Some KM3NeT design choices were motivated partly by the objective of reducing the impact of bioluminescence. For instance, bioluminescent bursts are understood to be emitted by larger organisms in response to stimuli, *e.g.* collision with mechanical elements of the DUs [229, 230]; therefore the amount of mechanical parts surrounding each DOM has been reduced drastically in the KM3NeT design with respect to the ANTARES one. The multi-PMT design is also an asset for identifying bioluminescent activity. An extensive discussion of the effect of bioluminescence in ANTARES and its expected effect in ORCA can be found in Ref. [94].

Fig. 3.4 shows the distribution of the rate of single L0 hits ('singles rate') per timeslice for one PMT of a KM3NeT prototype DU. The contributions from the steady background rate and from bioluminescent bursts can be clearly identified as the gaussian peak centered on $\sim 6\,\mathrm{kHz}$ and the high-rate tail, respectively. The singles rate averaged over the whole detector and a time period of a few hours typically ranges between 6 kHz (low bioluminescence conditions) and 8-10 kHz (high bioluminescence conditions).

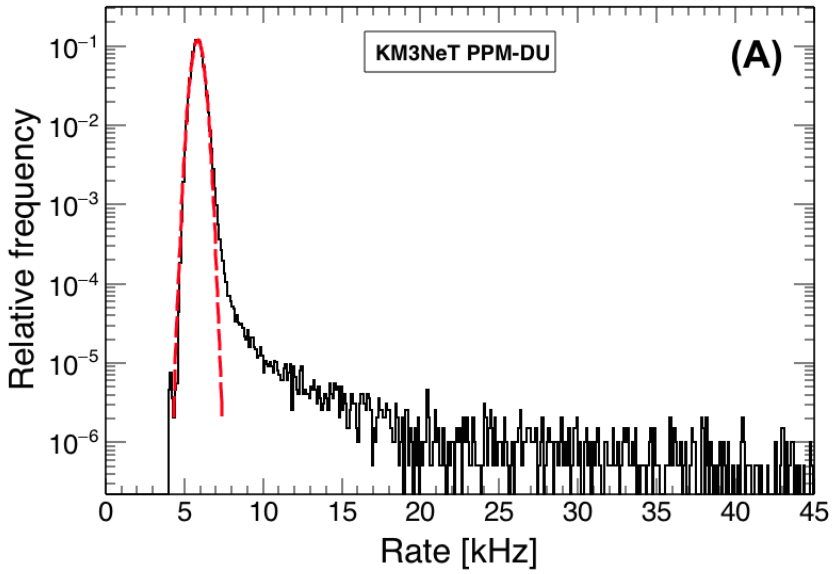

Figure 3.4: Distribution of the average rate of single hits per 134 ms timeslice, for one PMT of the prototype Detection Unit deployed in Capo Passero at a depth of 3500 m (ARCA site) and operated from May 2014 to July 2015. The red curve shows a gaussian fit, centered on 5.9 kHz. Taken from Ref. [230].

In addition to ^{40}K and bioluminescence, dark pulses (random pulses in the absence of an external source of light) constitute an 'internal' source of uncorrelated PMT hits. The most significant source of such random noise is the spontaneous thermionic emission of electrons by the photocathode. Radioactive decays inside the PMT structure (in particular in the glass enveloppe) also contribute, similarly to ^{40}K decays in water [227]. The rate of L0 hits due to dark pulse ('dark count rate') can fluctuate from one PMT to the other, and depends on environmental conditions. It is a constructor requirement of the KM3NeT PMTs that the dark count rate at 20° C be lower than 2000 counts per second (cps). Black box studies have measured a mean dark count rate of about 700 cps for the PMTs satisfying this requirement (while a few % are found above 2000 cps) [227]. This is therefore a small, but non-negligible contribution to the steady background rate.

Calibration and monitoring

The detection and accurate reconstruction of neutrino events in ORCA requires a precise knowledge of the optical properties of the seawater and DOM, of the PMT efficiencies and timing, and of the positioning of the optical modules. A timing accuracy at the ns level is

necessary, which translates into a position accuracy of $\sim 10\,\text{cm}$. Besides the PMTs, each DOM holds additional calibration and monitoring instruments:

- pressure, temperature and humidity sensors,
- a LED beacon, used to illuminate the above neighbouring DOM for timing and position calibration,
- a compass and tilt-meter for monitoring the DOM orientation,
- an acoustic sensor fixed to the inner surface of the glass sphere, used for acoustic positioning.

The DOMs positioning is monitored in real-time using acoustic triangulation with the network of acoustic sensors installed on the DOMs, DU base containers, and on dedicated Calibration Units (CU) positioned at the periphery of the detector (see Fig. 3.3). The timing calibration is done using the LED beacons mounted on the DOMs, as well as laser beacons installed on the CUs. These optical beacons also serve for measurements of the water transparency.

In addition, the steady ^{40}K background provides methods for *intrinsic* calibration. The rate of ^{40}K decays is hardly expected to vary, as it mostly depends on salinity which is rather stable and measured precisely with dedicated sensors. The expected rate of single hits and twofold hit coincidences are both largely dominated by the ^{40}K background. Their study can therefore be used in determining water optical properties, individual PMT efficiencies, relative time offsets of PMTs in a DOM, as well as their time spread properties. Fig. 3.5 shows the rate of two-fold hit coincidences as a function of the hit time difference for an arbitrarily chosen pair of adjacent PMTs in one DOM of the first two ARCA DUs. The distribution of the time difference between coincident hits is expected to be approximately normally distributed, the mean being the difference between the transit times of the involved PMTs, the width related to their respective transit time spread (TTS), and the integral scaling with the product of the PMT efficiencies [236]. , These parameters can be estimated by performing a gaussian fit to the coincident rates as in Fig. 3.5, simultaneously for all pairs of PMTs.

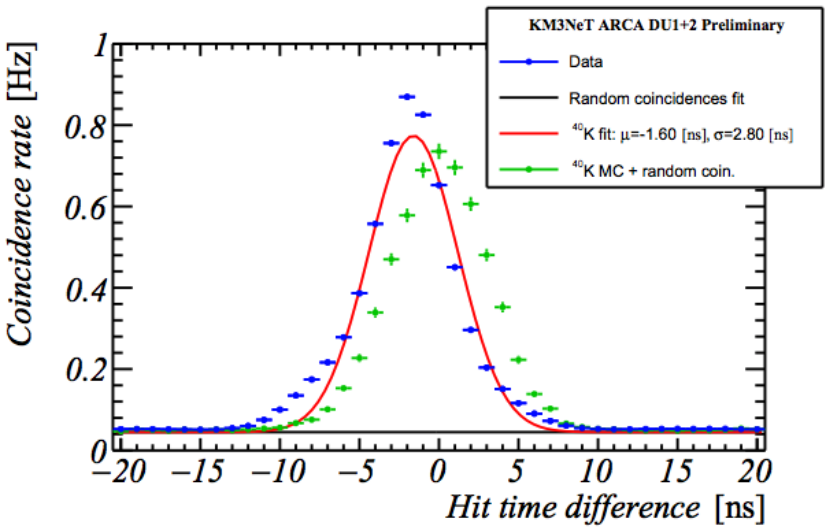

Figure 3.5: Rate of coincident pairs of hits as a function of the hit time difference, for an arbitrarily chosen pair of adjacent PMTs in one DOM of the first installed ARCA DUs. Shown are the rate observed in data without calibration (blue points), with gaussian fit (red line), and from Monte Carlo (MC) simulation (green points). The offset between the data and simulated distributions corresponds to the intrinsic time offset between the PMTs due to their different transit times. The black solid line shows the fitted rate of truly random ('accidental') coincidences, useful for absolute calibration. Taken from Ref. [236].

While ^{40}K decays provide for intrinsic methods of relative time calibration of PMTs within a DOM, an additional time offset between DOMs can be expected due to offsets between the clocks of their CLBs. As discussed in Ref. [236], intrinsic *inter-DOM* calibration can also be performed, using the physical background of atmospheric muons.

3.1.3 Data processing and triggering

The data processing follows the so-called 'all data to shore' approach, as successfully implemented by ANTARES: all data processing and filtering beyond hit digitisation is done in a shore station. Although this requires a high continuous data throughput between the detector and the shore station, it has the advantage of minimising the complexity of the underwater data acquisition systems. The L0 hit data ('optical data stream') is transmitted through a dedicated optical fibre in the cable running along the DU. In a group of 4 DUs, each DOM is attributed a dedicated wavelength so that the signal from all 72 DOMs can be multiplexed downstream and transmitted to shore over a single fibre.

The data rate of a single DOM in average bioluminescence conditions (assuming a 7 kHz overall singles rate) is ~ 10 Mb/s, which amounts to about 20 Gb/s for the whole detector. The data stream must therefore be filtered, with the objective of reducing the amount of data written to disk by a factor $\sim 10^3 - 10^5$. To this end, trigger algorithms are applied to select events of interest and discard background.

As seen in Sec. 3.1.2, most optical background hits are uncorrelated in space and time. In contrast, hits originating from neutrino secondaries or atmospheric muons are expected to be correlated at the level of a DOM, and to show a distinct space-time pattern among DOMs across the detector. The trigger algorithms generally search for such patterns aggregating local clusters of hits in time coincidence on the same DOM or on neighbouring DOMs.

The data processing is done continuously on a computing farm in the shore station. The 'frames' of data sent by all active DOMs in a time period of 100 ms, together forming a so-called *time slice*, are processed by a data filter instance running on a single CPU core. The data filter first calibrates the received raw data, typically correcting for time offsets between DOMs and PMTs, and runs the trigger algorithms on the calibrated data over the time slice. The computing farm processes as many time slices in parallel as there are CPU cores available.

When a trigger 'fires', a snapshot of all hits recorded in the detector in a window of $\pm 1.3\,\mu s$ around the hits that fired the trigger is written to disk: this constitutes an 'event'. The time window corresponds to the travel time of a photon across the whole detector. Considering that an ultra-relativistic particle also traverses the detector in $\mathcal{O}(1\,\mu s)$ and that the time slice duration is 10^5 times larger, the probability for a physics event to be split between two time slices is very small.

In addition to triggered events, the data acquisition status and singles rates over the time slice are recorded for each PMT, as a way to monitor the global status of the detector. This summary data can be used by simulations and reconstructions, *e.g.* to estimate the optical background conditions at the time of an event or over the course of an acquisition run.

Trigger patterns and optimisation

Three different hit coincidence levels are conventionnally defined. **L0 hits** refers to all hits having a PMT pulse charge above the 0.3 p.e. threshold, *i.e.* all hits sent to shore. **L1 hits** are pairs of hits taking part in a local coincidence, defined as two or more hits recorded in the same DOM within a certain time window Δt. The time window, typically set to 10 ns for ORCA standard physics runs, can be adjusted depending on data-taking conditions and run type (high bioluminescence conditions, calibration run, *etc*). **L2 hits** are hits taking part in one of the

higher-level coincidence patterns triggering data acquisition. There are three of such trigger algorithms.

- The **3D muon** trigger targets track-like events specifically. The algorithm scans over 200 uniformly distributed track directions and looks for clusters of at least $N_{3D,\mu} = 4$ causally connected L1 hits contained in a cylinder centred on the track direction with a predefined radius $R_{3D,\mu}$. Two hits occuring at (t_i, x_i, y_i, z_i) and (t_j, x_j, y_j, z_j) (where the z-axis is the cylinder axis and t is measured from the start of the recorder event) are considered as causally connected if their coordinates are compatible with an emission along the track hypothesis, *i.e.* if they satisfy the condition

$$c|t_i - t_j| < (z_i - z_j) + n \sin\theta_c \sqrt{(x_i^2 - x_j^2) + (y_i^2 - y_j^2)} + ct \tag{3.3}$$

  with n the refractive index, θ_c the Cherenkov angle, and δt a tunable additional time window.
- The **3D shower** algorithm is designed for shower-like topologies. It searches for $N_{3D,sh} = 3$ causally connected L1 hits on DOMs separated by less than a predefined maximum distance $D_{max,sh}$. For the shower trigger, hits are causally connected if

$$c|t_i - t_j| < n|\vec{r}_i - \vec{r}_j| + c\delta t, \tag{3.4}$$

  i.e. an isotropic light emission pattern is assumed.
- The **MX trigger** is similar to the 3D shower trigger, but is designed to lower the detection threshold energy. It searches for a cluster of L0 hits around a single L1, in order to include faint events which do not emit enough light to produce multiple causally connected L1.

The main technical limitation to the design of trigger algorithms is that they must be fast enough to be applicable in real-time. Otherwise the performance of a given trigger strategy is measured by its signal efficiency, which trades off with the purity of the triggered events. Efficiency is the proportion of neutrino events that are triggered, while purity corresponds to a low rate of pure noise events (where the trigger fires on random hit coincidences). In ORCA, optimising the trigger for the efficient detection of ~ 1 GeV neutrinos while keeping a manageable trigger rate is challenging, since such low-energy events yield few hits. A general requirement is that the rate remains dominated by atmospheric muon events (vastly more frequent than neutrino events) rather than by pure noise. The trigger parameters $R_{3D,\mu}$, D_{max}, δt are tunable in real time to adapt to bioluminescence condition.

Detailed trigger optimisation studies have been reported in the 2016 Letter of Intent [21]. The MX trigger strategy, introduced more recently, significantly increases the low-energy detection efficiency while keeping manageable noise rates. One of the initial goals of this thesis work was the evaluation of the subsequent performance improvement, with the sensitivity to the NMH determination as figure of merit.

3.2 Simulation

In this section the software tools used for ORCA simulations are introduced. The whole chain is summarised on Fig. 3.6, which also includes the subsequent steps of event reconstruction and classification. The chain starts with the neutrino and atmospheric muon event generators, introduced in Sec. 3.2.1. The simulation of the propagation of secondary particles as well as the production, propagation, and detection of Cherenkov light are then presented in Sec. 3.2.2. The final step is the addition of optical background, simulation of the PMT response, and the application of trigger algorithms (Sec. 3.2.3). The main simulation sample employed for the sensitivity studies presented in this thesis, labeled 'ORCA2016', was started in 2016 and

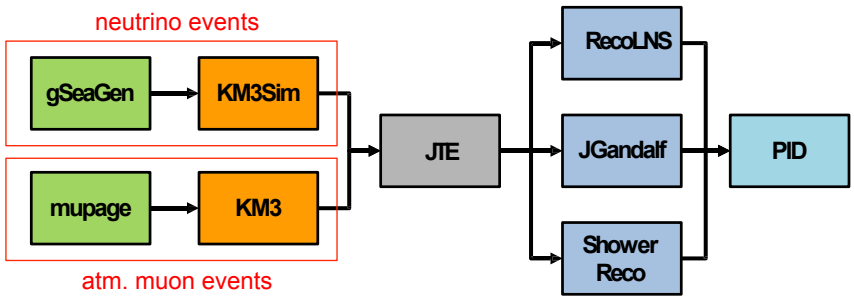

Figure 3.6: Flowchart of the simulation chain for ORCA Monte Carlo. The boxes indicate the software used for the successive steps: event generation (`gSeaGen`, `MUPAGE`), particle and light propagation (`KM3Sim`, `KM3`), PMT simulation and triggering (`JTE`), event reconstruction (`RecoLNS`, `JGandalf`, `ShowerReco`), and finally event classification, also referred to as PID (particle identification).

finalised early 2018 (including the subsequent reconstruction and classification steps). In this MC production, I was in charge of running the simulation of atmospheric muon events, which represented about 45 CPU-years of computation time. In Sec. 3.2.4 some of the specifics of the 'ORCA2016' sample are outlined, as they are useful for understanding some developments in Chap. 5.

3.2.1 Event generators

Neutrino interactions

The `gSeaGen` code [237] has been designed within the ANTARES and KM3NeT collaborations for simulating neutrino events. It is based on the widely used `GENIE` neutrino event generator [97, 98]. `gSeaGen` defines a volume surrounding the detector called the *can*, large enough so that Cherenkov light produced outside it may not be seen by the detector.

In the 'ORCA2016' production the can was set to extend the instrumented volume[2] by 60 m, which is about the light absorption length for a photon wavelength of 440 nm. It is a cylinder of radius $R = 205.4\,\text{m}$ and height $H = 256.7\,\text{m}$. The top and bottom of the can are located at $z_{\text{min}} = -117.2\,\text{m}$ and $z_{\text{max}} = 139.5\,\text{m}$, where $z = 0$ is at the detector center. The cylinder is not symmetrical with respect to the detector center because the can is bounded on its bottom side by the seabed rock, where no Cherenkov light is produced. The can volume is then

$$V_{\text{can}} = 3.402 \cdot 10^7\,\text{m}^3 \tag{3.5}$$

and the corresponding mass of seawater is:

$$M_{\text{can}} = 3.528 \cdot 10^{10}\,\text{kg} = 35.28\,\text{Mt}, \tag{3.6}$$

The value

$$\rho_{\text{water}} = 1.037 \cdot 10^3\,\text{kg.m}^{-3} \tag{3.7}$$

has been used for the density of seawater at this depth[3], and the mass and volume represented by the detector parts are neglected. These values are used all along in the calculations described in Chap. 4 and 5.

[2] In practice, `gSeaGen` reads the detector geometry from an input file, and defines the instrumented volume as a cylinder centered on the detector center of gravity and containing all PMTs.

[3] The density can be calculated based on the measured pressure, temperature and salinity using an equation of state, see for instance Refs. [238, 239]. The value $\rho \simeq 1.037 \cdot 10^3\,\text{kg.m}^{-3}$ was calculated using Ref. [238] with a pressure of 207.5 kPa, temperature of 13.2°C and a salinity of 37 g/kg. The `gSeaGen` code uses a slightly higher value ($1.039 \cdot 10^3\,\text{kg.m}^{-3}$). The difference has a negligible impact on the calculated rates relatively to other sources of uncertainty, as does the difference induced by salinity variations.

The volume in which `gSeaGen` simulates neutrino events depends on the event type. For NC events and CC $\nu_e/\overline{\nu}_e$ which produce no long particle tracks, the interaction volume coincides with the can and is thus composed only of seawater. In contrast, CC interactions of muon and tau (anti)neutrinos can procude muons propagating towards the detector which may eventually produce light susceptible to reach the instrumented volume. Therefore, for these interaction types

- the interaction volume is extended by an additional length corresponding to the energy-dependent maximum muon range (as a result the interaction volume can include a layer of bedrock beneath the can),
- the outgoing muons are propagated, and stored if they reach the can surface.

The neutrino-induced particles are written to the output only if they are generated inside the can or have propagated up to its surface. As a result, the density of neutrino events per unit volume found in the `gSeaGen` runs is uniform in the can (up to the effect of statistical fluctuations), while outside the can it depends on the event type and energy. This must be taken into account in the detector response modeling (see Chap. 5).

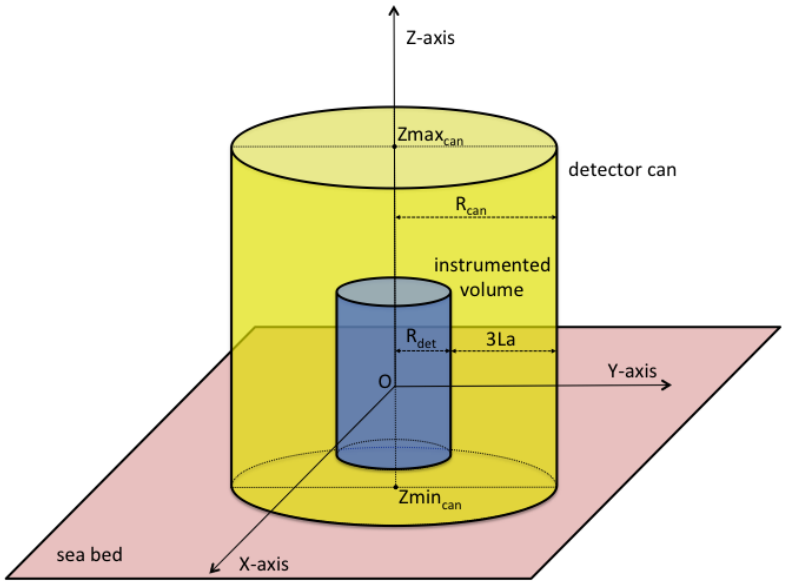

Figure 3.7: Definition of the detector can in `gSeaGen`. Figure identical to [237] (high resolution version provided by Carla Distefano in KM3NeT internal documentation).

Atmospheric muon bundles

The generation of atmospheric muons is done using the `MUPAGE` code [240] [241]. `MUPAGE` is a fast Monte Carlo generator developed specifically for underwater and ice-based neutrino telescopes. It generates single and multi-muon events based on parametrisations of the flux and the distributions of multiplicity and lateral spread of muon bundles, for vertical depths ranging from 1.5 to 5 km water equivalent of water or ice. The parametrisations were obtained by a full Monte Carlo simulation of the primary cosmic ray interactions and propagation of the induced showers in the atmosphere; this simulation used the `HEMAS-DPM` code [242] which was cross-checked against muon bundle measurements at the Gran Sasso underground laboratory [243].

In `MUPAGE`, an event is a muon bundle with a given multiplicity on the surface of a virtual cylinder surrounding the detector. As in `gSeaGen`, the `MUPAGE` can corresponds to a cylinder

enclosing the instrumented volume and extended by a certain amount, in this case *three* times the absorption length, in all directions. The base can cylinder is chosen larger for atmospheric muons than neutrinos because their light yield is generally more important. Furthermore, its radius is extended by an additional 100 m, because muons in high-multiplicity bundles can be hundreds of meters far from the shower axis. This surface is called the *extended can* in `MUPAGE`.

Muon bundles with multiplicities up to 200 and lateral spread up to 100 m were simulated, with a threshold at 1 GeV for the total bundle energy, and a maximal muon energy of 500 TeV.

3.2.2 Secondaries and Cherenkov light propagation

The generators described in the last section provide the kinematics of the interactions of neutrinos in and around the detector, and of atmospheric muon bundles reaching the detector surroundings. More precisely, `gSeaGen` provides (among other global event information) a list of particles emerging from the interaction with their individual four-momenta in the detector frame; MUPAGE provides a list of muon tracks with the space and time coordinates of their impact on the extended can surface. The next step in the simulation chain is the propagation of the secondary particles and the Cherenkov light simulation. For neutrino events this is done using a `Geant4`-based simulation software tracking individual particles, called `KM3Sim`. For atmospheric muon events which produce much brighter events, using the full simulation would be excessively time consuming. Therefore a faster tool (`KM3`) is employed, which uses pre-computed parametrised tables to propagate light. These codes are shortly described in the next two paragraphs.

Neutrino events: `KM3Sim`

The `KM3Sim` software was developed for KM3NeT, originally as part of the `HOURS` package [244] [245]. The propagation of the charged particles emerging from neutrino interactions is simulated using `Geant4` tools [246]. It simulates the production of Cherenkov light by primary and secondary particles, and tracks the individual photons in seawater, simulating absorption and scattering. Finally, it simulates the light detection, accounting for the following PMT and DOM characteristics: photocathode area, quantum efficiency and angular acceptance, as well as the light transmission in the DOM glass and optical gel.

Atmospheric muon events: `KM3`

The older Fortran-based package `KM3` is used for simulating the Cherenkov light from atmospheric muons. `KM3` was developed for and is used standardly in ANTARES for simulating both muons and electromagnetic showers. For the light propagation it uses parametrised tables (obtained from a full `GEANT3`-based simulation) describing the distribution of the numbers and arrival times of PMT hits at different distances, positions and orientations with respect to a given muon track or electromagnetic shower [247]. As a consequence `KM3` is much faster to run than `KM3Sim`. `KM3` also simulates the photon detection, implementing (in its modified KM3NeT version) the same PMT and DOM characteristics as `KM3Sim`. A comparison of the results of `KM3Sim` and `KM3` can be found in Ref. [21].

3.2.3 Optical noise, PMT simulation, and triggering

The optical background is simulated by the addition of random photon hits, uncorrelated and correlated, on all PMTs in the detector. Uncorrelated single photon hits are added with a rate of 10 kHz per PMT. In addition, time-correlated hit coincidences of multiplicity 2 to 5 are added on each DOM with the following rates: 500 Hz twofold, 50 Hz threefold, 5 Hz fourfold and 0.5 Hz fivefold. These are conservative assumptions with respect to the rates measured in the prototype in-situ tests and first Detection Units – typically around 8 kHz for singles and 340 Hz

for twofold coincidences [229, 230, 236, 248]. Thus, although localised bioluminescence bursts are not added to the simulation, the overall background rate corresponds to higher than average bioluminescence activity.

The final step of the detector simulation consists in the conversion of photon hits on each PMT's photocathode into L0 hits, then into time frames of raw data as sent by the CLBs to the shore station. The quantum efficiency of the photocathode is already included at the `KM3Sim/KM3` stages, so the PMT response is simulated starting at the photoelectron level. This simulation includes the TTS distribution (gaussian smearing of the hit arrival times), gain spread, and merging of coincident hits into single pulses with higher ToT.

The above step can be seen as a 'decalibration' of the MC truth information, so as to obtain pseudo-data affected by the same calibration uncertainties as real physics events. Moreover, the MC simulated data frames are formatted in the same way as the real ones, so that after this stage the exact same processing is applied to the Monte Carlo as to the real data stream. This data processing mostly consists in applying the trigger algorithms and retaining triggered events. The addition of optical background, PMT simulation, and triggering steps are handled together by a standard KM3NeT software tool called `JTriggerEfficiency` (`JTE`).

3.2.4 The 'ORCA2016' Monte Carlo production

For a fixed total cost (assumed to scale linearly with the total number of DOMs to produce), various geometrical configurations can be considered. A dense configuration allows to detect fainter events, thus lowering the energy threshold, and generally results in better reconstruction performance; this trades off with a smaller effective detection volume and less events collected overall.

An optimisation of the ORCA geometry was presented in the KM3NeT Letter of Intent [21]. Values of the inter-DOM vertical spacing of 6 m, 9 m, 12 m and 15 m spacing were considered, while the horizontal spacing was kept fixed to the minimum allowed by technical constraints. In order to keep the running time for simulations to a manageable level, a *masking* technique was employed. The full chain was ran only for a single dense configuration with 6 m spacing, and sparser configurations with 9, 12 and 15 m spacing were obtained by masking part of the DOMs in the simulated events (*i.e.* ignoring the hits on these DOMs in the triggering and reconstruction steps). The optimal spacing was found to be 9 m. In the rest of the thesis, the Monte Carlo sample corresponding to the 9 m-masked geometry will be referred to as the 'LoI' sample. The 'ORCA2016' Monte Carlo production was started in summer 2016 with the goal to confirm this result using a full simulation. It was the first full ORCA simulation based on a realistic detector geometry, accounting for technical requirements following deployment tests. This implied in particular a larger horizontal spacing than in the 'LoI' geometry.

In the production, the total size of the MC sample ('MC statistics') was increased, as detailed in Tab. 3.1. The comparatively lower overall efficiency (ratio of selected to generated events) in

	LoI		ORCA2016	
$\times 10^6$ events	Generated	Selected	Generated	Selected
$\nu_e + \overline{\nu}_e$ CC	3.4	0.63	14.6	1.4
$\nu_\mu + \overline{\nu}_\mu$ CC	6.4	0.60	15.9	1.5
$\nu_\tau + \overline{\nu}_\tau$ CC	2.1	0.35	23	0.86
$\nu + \overline{\nu}$ NC	3.5	0.31	103	1.1
Total	15.4	1.9	150	4.8

Table 3.1: Comparison of the available statistics of generated events (`gSeaGen` stage) and 'selected' events (*i.e.* events that are triggered, successfully reconstructed, and passing the background rejection cuts at the pre-PID stage).

the new production is due to an optimisation of the energy distribution favouring low-energy events (less efficiently selected). This optimisation was motivated by the development of full MC analysis methods. As discussed in Chap. 5 and Appendix B, using sparsely sampled MC in such applications leads to statistical artifacts whose effect, on average, is to overestimate sensitivities. This issue quickly became critical to my work, focused on developing a full MC analysis method for the NMH sensitivity study. I was therefore involved in optimising the event statistics for the new production. As shown on Fig. 3.8, the energy spectrum of the `gSeaGen` event generation was tuned in order to maximise the statistics of selected events available in the $\sim 2-10\,\mathrm{GeV}$ range. To estimate the impact of MC statistical uncertainties in a given bin, a common figure of merit is the ratio of the count of *unweighted* Monte Carlo events to the count of (expected) data events. While this ratio can easily be made uniform by aligning the event generation scheme on the expected interaction rate, here this option would be impractical. Indeed, in the neutrino simulation chain (Fig. 3.6) the computational bottleneck is light propagation (`KM3Sim`). At low energies ($\sim 1-2\,\mathrm{GeV}$), the event selection efficiency at the triggering stage (`JTE`) is very low, so that only a few % of the simulated events pass this stage. By using a realistic spectral index (E^{-3}) for the `gSeaGen` event generation, most of the computation time would then be wasted on simulating few-GeV events rejected by the trigger immediately after. Therefore the MC production was subdivided into two overlapping energy samples with different spectra[4]:

$$\frac{\mathrm{d}N}{\mathrm{d}E} \sim E^{-3}, \quad 1\,\mathrm{GeV} \leq E \leq 5\,\mathrm{GeV}, \tag{3.8}$$

$$\frac{\mathrm{d}N}{\mathrm{d}E} \sim E^{-1}, \quad 3\,\mathrm{GeV} \leq E \leq 100\,\mathrm{GeV}. \tag{3.9}$$

Note that the indicated spectral index corresponds to the flux generation index, the effect of cross-sections being folded in by `gSeaGen`. The resulting energy distribution is shown on Fig. 3.8, in comparison with the 'LoI' MC sample where a E^{-2} spectrum was used uniformly. As can be seen from the right panel, the uniformity of the MC/data ratio is greatly improved and fairly high MC statistics is achieved at low energy.

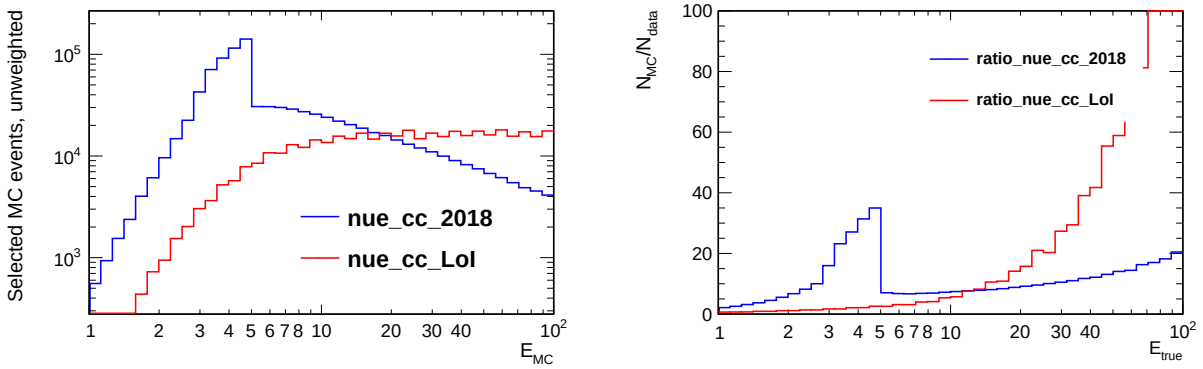

Figure 3.8: Comparison of the available statistics of unweighted events, as a function of true neutrino energy, in the 'LoI' and 'ORCA2016' (labeled '2018' in the figure) Monte Carlo samples. Distributions are restricted to truly upgoing ν_e CC events. **Left**: raw distributions of 'selected' events (triggered, successfully reconstructed, and passing the background rejection cuts at the pre-PID stage). **Right**: ratio of the raw distribution of selected events to the expected number of data events (calculated based on flux, cross-section, detection efficiency).

3.3 Event reconstruction and classification

3.3.1 Reconstruction of track-like and shower-like events

Dedicated algorithms have been developed to optimally reconstruct either $\nu_\mu/\overline{\nu}_\mu$ CC or $\nu_e/\overline{\nu}_e$ CC events. In both cases the reconstruction strategy is based on an assumption on the event topology, therefore they are generally referred to as 'track reconstruction' and 'shower reconstruction'. Track reconstruction strategies are focused on determining the length and direction of the main track, assumed to correspond to a $\mu^\pm$. The reconstructed track length is used as an estimate of the muon energy. The shower reconstruction method first locates the interaction vertex, initially assuming spherical light emission, and then refines the energy and direction estimates based on the expected angular light distribution (see Fig. 2.10). As shown on Fig. 3.6, in the simulation chain three different reconstruction algorithms are applied in parallel to all events: `recoLNS` (track), `JGandalf` (track) and `ShowerReco` (shower, also frequently referred to as `Dusj`). While both `recoLNS` and `JGandalf` have been applied for the 'ORCA2016' sample used in this work, `JGandalf` has proven more performant and has now superseded `recoLNS`. The common concepts and methods are first introduced, before discussing specific aspects and their respective performance.

Hit selection The first step of the reconstruction is hit selection, which purpose is to discard hits that are more likely to be due to optical background than to charged neutrino secondaries. Hit selection may be refined in successive steps of the reconstruction. As for event triggering, the selection is based on space-time coincidences and clustering. Since the reconstruction is performed offline, more sophisticated and CPU-intensive approaches can be used.

Time residual The time residual t_{res} of a hit with respect to a certain hypothesized 'source' in space and time $(\boldsymbol{x}, t)$ is the difference between the observed hit time and the predicted time assuming the light has travelled from $(\boldsymbol{x}, t)$ with velocity c_{water}. For instance, when fitting the vertex position of a shower event, the time residual δt with respect to the vertex hypothesis (x_{vtx}, t) is given by

$$\delta t = t_{\text{hit}} - t_{\text{vtx}} - d/c_{\text{water}} \tag{3.10}$$

where d is the distance between the vertex position x_{vtx} and the PMT. In the case of track reconstruction, the source hypothesis may be for instance a track direction and starting point.

Likelihood fit A example of approach to find a good source candidate (vertex, track) is then to minimise a function of the time residuals of all selected hits, $g(\{\delta t_i\})$, with respect to the source hypothesis. Hit characteristics beyond the simple time residuals, such as the arrival direction of the photon inferred from the PMT orientation, can be included as well. The functional form of the objective function g is typically related to a likelihood, which may account for various levels of physical complexity. It can also be adjusted for practical reasons (*e.g.* numerical stability).

Prefit scan If an arbitrary starting point is used, the full likelihood fit may be overly time consuming. In order to speed it up a preliminary set of starting points can be obtained by scanning over a grid spanning the whole space (or a large region) and selecting the best performing points. This also helps in avoiding that the maximisation misses the global maximum because of an ill-chosen starting point.

Track reconstruction: `JGandalf`

A first prefit step scans over a total of 844 equidistant track direction hypotheses. A clustering algorithm selects hits most compatible with each track hypothesis. Given an ensemble of preselected hits and assumed track direction, the most likely track starting point can then be solved for as a linear inversion problem, assuming light emission at the Cherenkov angle and ignoring photon scattering. The prefitted tracks are then ranked according to a custom discriminator function based on the number of DOMs participating and likelihood of the best-fit track. The best ranking prefitted tracks are then passed on to the 'PDF fit', a full likelihood fit using accurate probability density functions (PDFs) for the time residual distributions at individual PMTs. These PDFs are dependent on the PMT orientation with respect to the track hypothesis and account for the light emission profile of muons in water, single photon scattering probabilities, as well as the area, angular acceptance and quantum efficiency of the photocathode. Finally the best fit track's length is calculated based on the first and last DOMs hit along the track. In very good approximation, a muon track length scales linearly with energy below 100 GeV (minimum ionising particle regime), so the muon energy is estimated proportionally to track length.

The hadronic shower energy can be estimated based on the density of selected hits around the track relatively to the track length: a highly inelastic interaction (high-energetic hadronic shower) is expected to produce a short track with many hits, and conversely an elastic interaction will result in a long track with relatively few hits. An empirical correction depending on the number of selected hits is therefore applied to the track length based energy estimate, so as to yield the final neutrino energy estimate as well as a Bjorken-y estimate.

Shower reconstruction: `Dusj`

The shower reconstruction is performed in two steps. In the first step, the interaction vertex is reconstructed based on the recorded arrival time of the hits, and in the second step the direction, energy and inelasticity are reconstructed based on the number of hits and their distribution in the detector. This factorisation of the fitting procedure works well due to the homogeneity of seawater and its large scattering length, which allows for a precise vertex reconstruction independently of the shower direction [94]. The overall structure of the algorithm has been adapted from a shower reconstruction developed for ANTARES.

The vertex reconstruction proceeds in two steps. First, a maximum likelihood fit to the distribution of hit time residuals is performed, including penalty factors to account for PMT directionality information. This first estimate is then refined by reselecting hits on PMTs facing the vertex and using time residuals PDFs accounting for distance-dependent noise levels. The final fit (energy, Bjorken-y, direction) uses PDFs for the number of expected photons per DOM depending on the position and orientation of the DOM relative to the vertex, neutrino energy, Bjorken-y and emission angle. Dedicated hit selections are reiterated at each step. Moreover the algorithm performs intermediate checks to avoid attempting to reconstruct topologies which do not resemble point-like emission and very bright events (typically induced by atmospheric muons) which would be excessively computer time-consuming.

Reconstruction performance

The performance of the track and shower reconstructions are shown on Fig. 3.9 and 3.10. The directional estimate is generally more better for track-like events, while the energy estimate is better for shower-like events. The accuracy is limited at low energy by several factors. Shorter muon tracks and compact electromagnetic showers are difficult to resolve given the limited instrumentation density. Moreover, both the intrinsic kinematic smearing (opening angle between neutrino and leading lepton) and the relative stochastic fluctuations in light production become

larger. At high energy, the muon energy resolution is degraded by the larger fraction of muon tracks being only partially contained in the detector. Note that the degradation in shower energy resolution observed in 3.10 between the 'LoI' and 'ORCA2016' stages simulation samples is due to the enlarged horizontal spacing of the lines and the inclusion of fainter events (MX trigger); it is the counterpart of the large increase in effective volume acheved in parallel. Further discussion of the reconstruction strategies and their performance can be found in Refs. [21, 114].

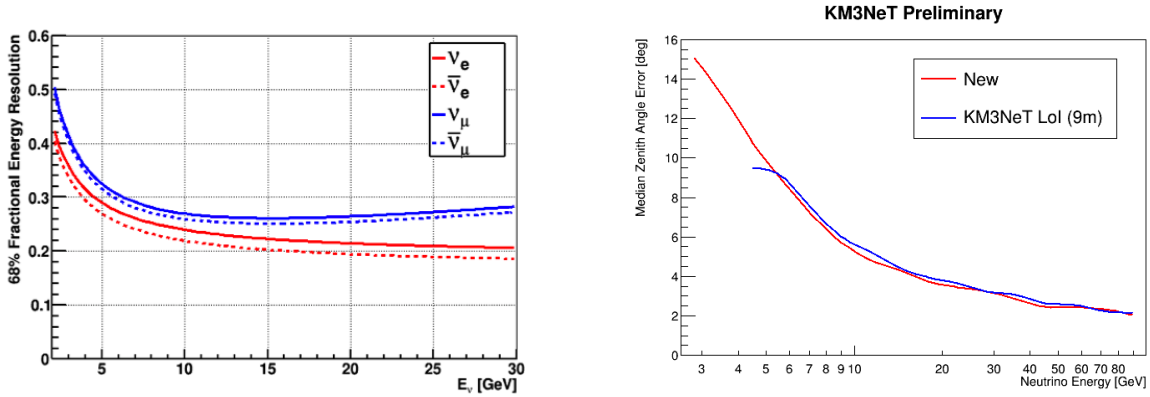

Figure 3.9: **Left:** Parametrisation of the 68% fractional energy resolution (*i.e.* σ_E/E if the energy response is fitted by a gaussian) for the track reconstruction on $\nu_\mu/\overline{\nu}_\mu$ CC events (blue) and for the shower reconstruction on $\nu_e/\overline{\nu}_e$ CC events (red). **Right:** median error on the zenith angle for the track reconstruction applied to $\nu_\mu/\overline{\nu}_\mu$ CC events.

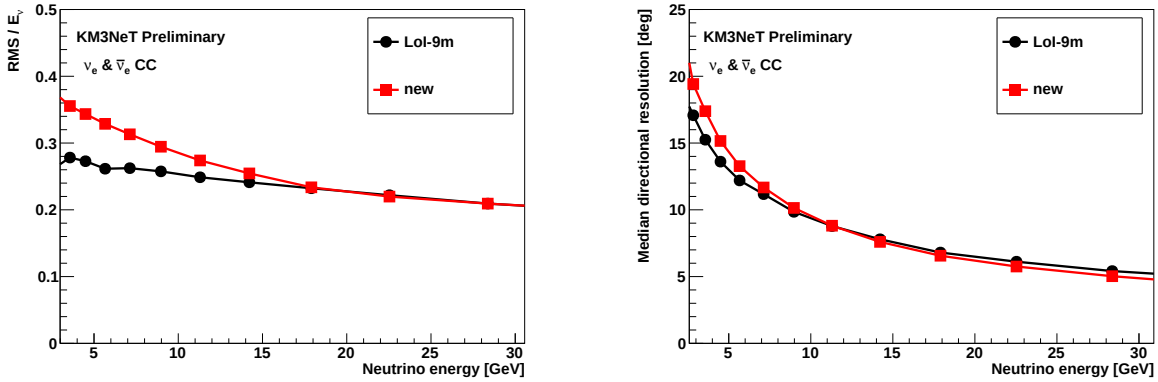

Figure 3.10: Performance of the shower reconstruction on ν_e and $\overline{\nu}_e$ CC events. The performance achieved with the 'ORCA2016' detector configuration, new trigger and improved reconstruction (red) is compared with the one published in Ref. [21] (black). **Left:** 68% fractional energy resolution (*i.e.* σ_E/E if the energy response is fitted by a gaussian). **Right:** median error on the space angle.

3.3.2 Background rejection and event classification

As discussed previously, the rate of triggered events is dominated by atmospheric muons. While this background can be reduced by several orders of magnitude by selecting events reconstructed as upgoing and with a starting point contained in the detector, the remaining contamination remains important due to the misreconstruction of some events as upgoing. Pure noise events (triggered on optical background) also dominate over neutrino events, although less frequent than atmospheric muons. The reconstructed starting point and direction of these events are distributed rather uniformly, however they generally involve fewer hits per DOM and are reconstructed with lower quality parameters than signal events. Further background rejection is therefore achieved by cutting on reconstruction quality variables, reconstructed vertex position, and applying dedicated classification algorithms: decision tree learning, as described hereafter.

Particle identification classifiers

The last step of the data processing chain is referred to as particle identification (PID). PID is generally aimed at rejecting background and determining the flavour of selected CC events. As explained in Sec. 2.1.3, for megaton-scale detectors the flavour tagging of neutrino events is essentially limited to classifying events as track-like or shower-like. The rejection of pure noise and atmospheric muons events is also important. The goal of a classification algorithm is then to predict to which predefined *class* an event belongs, based on a set of *features* of the event. Machine learning techniques are often used in experimental particle physics for such classification problems, since the availability of Monte Carlo simulated data allows to train the algorithm with extremely large samples of events whose true class membership is known.

The standard PID approach in ORCA uses a random decision forest (RDF), a machine learning technique based on the aggregation of classification trees [249]. A classification tree is a type of decision tree: an event goes through a path in the tree where each internal node performs a test on one feature and, depending on the outcome, pushes the event to one of its child nodes. The sequence goes on until the event reaches a leaf, associated with a class label. Learning algorithms are designed to build optimal decision trees based on the provided training data. A RDF is an ensemble of classification trees where each tree is trained on a randomly drawn fraction of the available features. RDFs were introduced to correct the tendency of single decision trees to *overfit* on their training sample, *i.e.* build classification models that fail to apply to new data, typically picking up non physical features (random noise) of the training sample.

The RDF-based classification tools for ORCA are referred to as 'ECAP PID' or 'ECAP classifier'. They were originally developed at ECAP (Erlangen) for the ANTARES neutrino telescope [250], and later adapted to the classification of ARCA and ORCA events [21]. A large spectrum of event features enter as classification inputs:

- low-level information, for instance the total number of hits or DOMs participating in the event;
- raw reconstruction outputs, including both final reconstruction estimates and intermediate parameters (*e.g.* fit quality indicators at successive steps);
- dedicated features specifically calculated for PID purposes, such as the tensor of inertia of the hit distribution, or variables combining estimates from the track and shower algorithms.

Three distinct binary classifiers are trained separately and run independently on all reconstructed events, attributing a score $0 \leq p \leq 1$ to each event:

- $p_{\text{bkg,noise}}$ for noise-like versus signal-like classification (aimed at rejecting pure noise events),
- $p_{\text{bkg},\mu}$ for atmospheric muon-like versus neutrino-like classification,
- p_{track} for track-like versus shower-like classification.

In each case the output score is simply the fraction of the classification trees that voted for the predicted class out of 101 randomised trees. Example distributions of p_{track} in a high an low neutrino energy range are shown on Fig. 3.11. Arbitrary combinations of these three scores (possibly together with additional cuts) can then be used to define event classes in the oscillation analysis. In the typical strategy three classes are defined (background, track-like, shower-like), and events are binned according to the reconstruction estimates of E_ν, $\cos\theta_z$ and y provided by the track (resp. shower) reconstruction for events classified as track-like (resp. shower-like).

To first order, the performance of the track/shower classifier is determined by its ability to tag a muon track emerging from the interaction. As can be seen from Fig. 3.11 and 3.12, this is far easier at high energy whereas at low energy the classifier shows an overall bias towards the shower-like topology. More specifically, the $\nu_\mu/\overline{\nu}_\mu$ tagging performance mainly depends on the (anti)muon energy; for $E_\mu > 10\,\text{GeV}$ more than 95% of the events are correctly classified

as track-like. The significantly better performance for $\overline{\nu}_\mu$ classification (as compared to ν_μ) observed on Fig. 3.12 is thus due to the fact that the interaction of antineutrinos tend to be more inelastic, *i.e.* produce higher energetic charged leptons (se Sec. 2.1.2, Fig. 2.7).

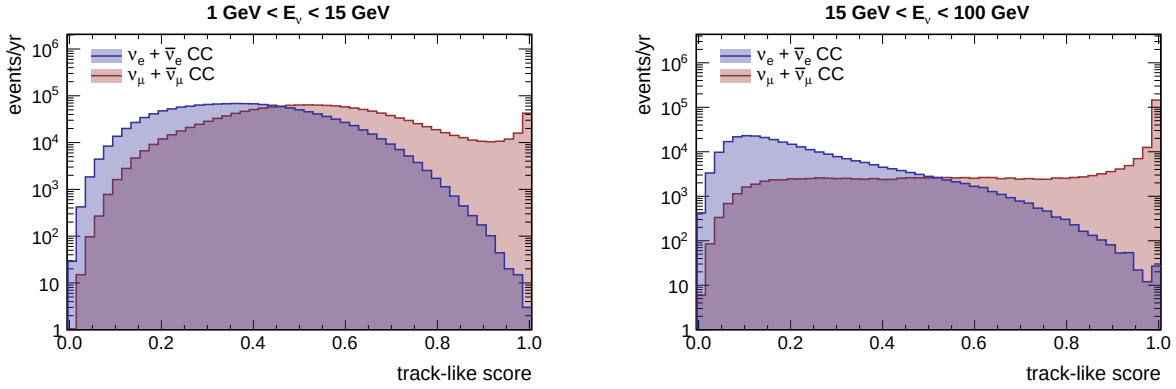

Figure 3.11: Distributions of the track/shower classification score for muon and electron (anti)neutrino CC events in the energy range 1-15 GeV (left) and 15-100 GeV (right). A track-like score close to 1 means 'highly track-like'. Events are weighted according to the unoscillated atmospheric flux.

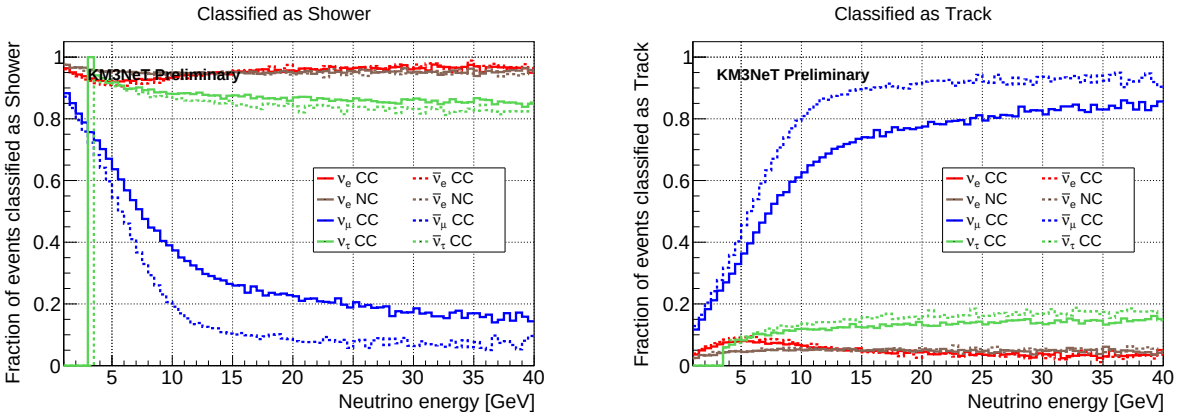

Figure 3.12: Event classification probabilities as a function of energy for each channel ν_x. Figure re-used from Ref. [251]. Events are classified as track if $p_{\text{track}} > 0.6$ and as shower otherwise.

As an alternative to the RDF-based approach, a classifier based on deep convolutional neural networks is being developed, relying on state-of-the-art open source software tools [252]. The technique, originally designed for image recognition, can be successfully applied to the space-time patterns of hits in ORCA events and has already achieved promising results [251, 253].

As discussed in Chap. 5, being able to systematically compare the performance (in terms of sensitivity to neutrino oscillations) of different options for reconstruction, classification, or analysis strategy was one of the motivations of the analysis methodology and software framework developed in this work.

Part II

Development of analysis methods

Summary of notations

A summary of the notations used in the following chapters is given here:

- neutrino interaction channels are indexed by the symbol ν_x:

$$\nu_x \in \{\nu_e\,\mathrm{CC}, \overline{\nu}_e\,\mathrm{CC}, \nu_\mu\,\mathrm{CC}, \overline{\nu}_\mu\,\mathrm{CC}, \nu_\tau\,\mathrm{CC}, \overline{\nu}_\tau\,\mathrm{CC}, \nu\,\mathrm{NC}, \overline{\nu}\,\mathrm{NC}\}$$

- event classes in term of particle identification (PID) are indexed by i

$$0 \leq i \leq N_{\mathrm{class}} - 1$$

  (note that the number N_{class} of PID classes is arbitrary)
- the bins in "true space" are indexed by the values (E, θ, y) of the true variables at the bin center[5],
- the bins in "reconstructed space" are indexed by the reconstructed values at bin center (E', θ', y').

It must be emphasized that (E, θ, y) and (E', θ', y') take **discrete values**; they can be identified with the corresponding bin indices. Note also that the notation (E, θ, y) does not imply that a uniform bin width is employed: instead the bin width is generally constant in $\log_{10} E$, $\log_{10} E'$, $\cos\theta$, $\cos\theta'$.

[5] In the actual implementation the variables E and $\cos\theta$ are employed regardless of the binning shape. The notation θ is used here to avoid overburdening the equations.

Chapter 4

Calculation of neutrino interaction rates

Contents

This chapter details the method for calculating the expected rate of neutrino interactions at the detector site, and discusses the NMH signal at this level, *i.e.* without realistic detector effects. After introducing the overall calculation, flux and cross-section models are presented from the technical point of view. The calculation of oscillation probabilities is then described, and examples of $(E_\nu,\ \cos\theta_z)$ oscillograms are discussed. Some practical issues relevant in the context of the full sensitivity analysis are adressed in both Sec. 4.1 and 4.2, while pure implementation aspects are discussed in Appendix C. Finally, Sec. 4.3 discusses the measurement of the NMH based on the obtained event rate distributions, complementing the phenomenological discussion initiated in Chap. 2.

4.1 Event rate calculation and input models

This section has evolved from an internal note documenting the analysis as a reference for the KM3NeT collaboration: technical aspects are thus covered in some detail. The analysis code I developed for the NMH sensitivity study has become an official KM3NeT software, which is described in Appendix C.

4.1.1 Overall calculation and implementation

The first step of the sensitivity analysis is to compute the expected rate of neutrino interactions per unit exposure[1] at the detector site. The rate is calculated as a function of neutrino energy E, zenith angle θ_z, and interaction inelasticity y. The following ingredients are involved in this calculation:

- incoming flux of atmospheric neutrinos,
- oscillations of neutrinos along their propagation path inside the Earth,
- interaction cross-sections, including the distribution of inelasticity y for the final state,
- volume density of target nuclei in the seawater at the detector site.

In practice neutrino interactions are separated into 8 channels:

- charged current interactions (CC): ν_e, $\overline{\nu}_e$, ν_μ, $\overline{\nu}_\mu$, ν_τ, $\overline{\nu}_\tau$
- neutral current interactions (NC), which are insensitive to flavour: ν, $\overline{\nu}$

The rate per unit of exposure in channel ν_x $\frac{\mathrm{d}n_{\mathrm{int}}^{[\nu_x]}}{\mathrm{d}M\mathrm{d}t}(E,\theta,y)$ is computed as:

$$\frac{\mathrm{d}n_{\mathrm{int}}^{[\nu_x]}}{\mathrm{d}M\mathrm{d}t}(E,\theta,y) = \Delta y \cdot \frac{\mathrm{d}\sigma_{\nu_x}}{\mathrm{d}y}(E,y) \cdot \sum_{\nu_f} 2\pi \cdot \Delta E \cdot \Delta(\cos\theta) P_{\mathrm{osc}}(\nu_f \to \nu_x) \cdot \frac{\mathrm{d}\Phi_{\nu_f}}{\mathrm{d}E\mathrm{d}(\cos\theta)\mathrm{d}\phi}(E,\theta) \quad (4.1)$$

where

- $\nu_f \in \{\nu_e, \nu_\mu\}$ (resp. $\nu_f \in \{\overline{\nu}_e, \overline{\nu}_\mu\}$) when ν_x is a neutrino (resp. antineutrino) channel
- $\frac{\mathrm{d}\Phi_{\nu_f}}{\mathrm{d}E\mathrm{d}(\cos\theta)\mathrm{d}\phi}(E,\theta)$ is the double differential flux of atmospheric neutrinos at the approximate detector location;
- $P_{\mathrm{osc}}(\nu_f \to \nu_x)$ is the probability of flavour transition from ν_f to ν_x (1 when $\nu_x = \nu\mathrm{NC}$ or $\overline{\nu}\mathrm{NC}$);
- $\frac{\mathrm{d}\sigma_{\nu_x}}{\mathrm{d}y}(E,y)$ is the cross-section for an incoming neutrino of energy E interacting in channel ν_x (differential cross-section with respect to the final state inelasticity y);
- Δy, ΔE and $\Delta(\cos\theta)$ are the inelasticity, energy and zenith angle widths of the bin (E,θ,y), and the 2π factor stems from the integration over azimuth.

The rate of NC events is insensitive to neutrino oscillations, since their flavour does not affect the final state of neutrino-nuclei interactions. This amounts to setting the oscillation probability as 1 in Eq. 4.1. For simplification the systematic effects have not been included in Eq. 4.1; they are detailed in Chap. 7.

Choice of binning scheme

The expected number of events per unit exposure is computed according to 4.1 for 3-dimensional bins of true neutrino energy E_ν, true cosine of the zenith angle $\cos(\theta_z)$, and true inelasticity of the interaction, denoted y_{Bjorken} or y. The implementation leaves the possibility

[1] Exposure is defined as the product of data-taking time and mass of the detection target. In this thesis exposure is generally expressed in units of megaton-year (Mt.y).

for $\cos(\theta_z)$ bins to be defined such that the bin width is constant in zenith angle θ_z. Typical binning schemes used for interacting event rates in the mass hierarchy analysis is the following:

- 40 to 100 bins in true $\log_{10} E_\nu$ ranging from 1 to 100 GeV. The actual variable used for histogramming is E and the bin edges are defined such as to have a constant width in $\log_{10}(E_\nu)$).
- 40 to 100 bins in true $\cos\theta_z$ from -1 to 0 (upgoing events) and 10 bins from 0 to 1 (downgoing events).
- 4 bins of constant width in true y_{Bjorken}

Energy binning shape

Since the atmospheric neutrino fluxes roughly follow power laws, it is natural to use energy bins of equal logarithmic width. The same choice for reconstructed event histograms is also consistent with the fact that ORCA's relative energy resolution is, in first approximation, constant above a few GeV:

$$\delta E/E = \delta(\log E) \approx 25\%. \tag{4.2}$$

Zenith angle binning shape

A binning of constant width in $\cos\theta_z$ is used for the NMH analysis. This is motivated first by the fact that the solid angle covered by an interval of zenith angle $\theta_{z1} \leq \theta \leq \theta_{z2}$ is proportional to $|\cos\theta_{z1} - \cos\theta_{z2}|$. Considering to first order the atmospheric neutrino flux as isotropic, a binning in $\cos\theta_z$ thus yields equally populated bins along the zenith angle axis. In addition, oscillation probabilities depend on L/E_ν with $L = 2R_\oplus \cos\theta_z$ ($R_\oplus$ denoting Earth's radius. It is then more natural to compute oscillation probabilities at points equally spaced in $\cos\theta_z$.

A disadvantage of this choice is a reduction of the number of bins for core-crossing trajectories (15% of the upgoing $\cos\theta_z$ range) as compared to a binning in θ_z (about 30%), due to $\mathrm{d}(\cos\theta)/\mathrm{d}\theta$ approaching zero close to vertical. A binning with constant width in θ_z is thus used for Earth tomography analysis. To ensure an equivalent precision of the neutrino oscillation probability calculations the number of true $\cos\theta$ bins is generally increased.

Binning granularity

The atmospheric flux and neutrino cross-sections are modelled by smooth functions, and their values at arbitrary points can be interpolated from the grid provided by the simulation inputs, the computational cost of this step being negligible. Hence the dominant constraints on the choice of binning granularity are the computation of oscillation probabilities and application of the detector response. It is desirable to compute oscillation probabilities for a sufficiently large number of $(E, \cos\theta)$ points in order to sample fast oscillation patterns occuring in the resonance region precisely enough. However, the computation of oscillation probabilities is a bottleneck in the analysis chain. In addition, the computation time needed to apply the detector response to the true event distributions also scales linearly with the total number of bins in (E, θ, y).

As a side note, for reconstructed event histograms the choice of binning granularity is dominated by the detector resolutions (the bin width for a variable X should be comparable with the typical error on reconstructed X), and the need to sample the detector response functions in a sufficiently smooth manner (Monte Carlo statistics issues). This is discussed in Chap. 5.

4.1.2 Atmospheric flux model

The model of atmospheric neutrino flux follows the simulations of M. Honda *et al.* labeled 'HAKKM2014' [84]. The tabulated simulation results can be accessed online (see Ref. [85]).

There is one table per relevant type of incoming neutrino: ν_e, $\overline{\nu_e}$, ν_μ, $\overline{\nu_\mu}$. The data tables used here are the ones corresponding to the Gran Sasso site without mountain over the detector, assuming minimal solar activity. The calculated fluxes are all-year averaged and averaged over azimuth angle. While other KM3NeT publications have used instead the simulation output corresponding to the Fréjus site, the results are very similar and well within the systematics uncertainties considered for the atmospheric neutrino flux, as discussed in Chap. 7.

The data tables are organised in the following way:

- the tabulated quantity is the differential flux in $\mathrm{m}^{-2} \cdot \mathrm{s}^{-1} \cdot \mathrm{GeV}^{-1} \cdot \mathrm{sr}^{-1}$;
- with respect to energy, the value of this differential flux is given for values E_i equally spaced in $\log_{10} E$, from $E = 0.1\,\mathrm{GeV}$ to $E = 10$, with 20 points per decade;
- with respect to zenith angle, the flux averaged over 20 bins of width $\Delta(\cos\theta_z) = 0.10$ is reported.

For the purpose of this analysis these data tables have been converted into two-dimensional histograms with $X = \log_{10}(E), Y = \cos\theta$. The bins in X-axis are centered on the values given in the data tables. All data points from 1 GeV to 100 GeV are included, for a total of 41 bins. For the Y-axis the same 20 bins as in the data tables are used. An example of resulting histogram is shown on Fig. 4.1.

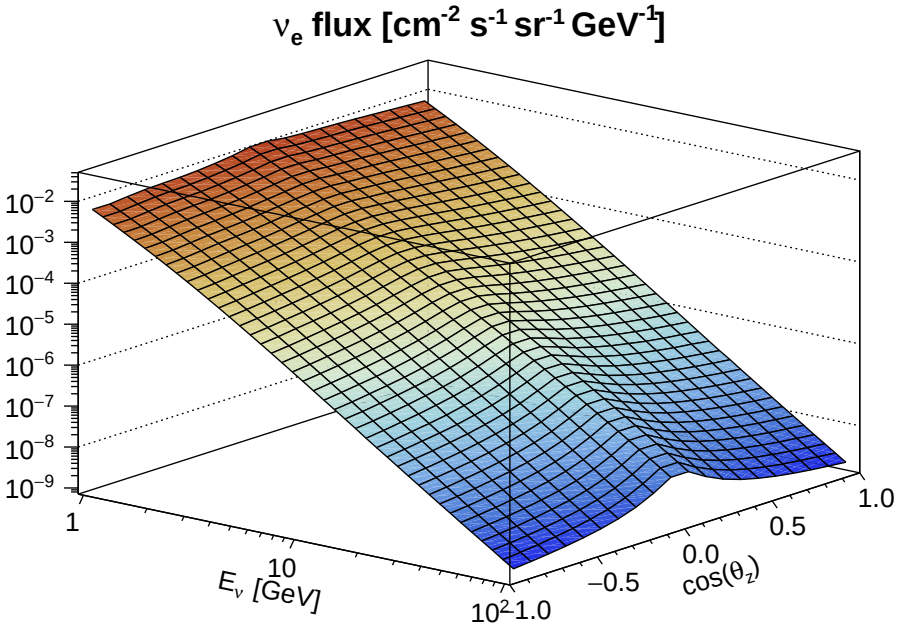

Figure 4.1: Example of two-dimensional flux histogram extracted from the HAKKM2014 tables [84, 85].

To compute interacting event distributions, the differential flux for a given value of $(E, \cos\theta_z)$ is interpolated from these histograms using bilinear interpolation. As described for instance in Refs. [21, 254], an alternative method consists in interpolating instead the cumulative flux distribution, which results in conserving the bin-wise flux integral. A spline interpolator implementing this approach has been developed by a bachelor student at APC under my supervision, but the method has not yet been incorporated in this analysis. The expected improvement in accuracy is however modest [254].

4.1.3 Cross-section model

The differential neutrino-nucleon cross-section per nucleon can be recast as:

$$\frac{\mathrm{d}\sigma_{\nu_x}}{\mathrm{d}y}(E,y) = \sigma(E_\nu)\cdot P_E(y) \tag{4.3}$$

where $\sigma(E)$ is the total cross-section and $P_E(y)$ is a normalised probability distribution for the interaction inelasticity y at energy E_ν. In practice $\sigma(E)$ is interpolated from graphs ($x = \log_{10}(E)$, $y = \frac{1}{E}\times\sigma(E)$). The graphs for each interaction channel are shown on Fig. 4.2. These are inclusive cross-sections, and do not include any final state dependence (e.g. Bjorken-y). The graphs (passed on from previous ORCA analyses) were obtained using the `GENIE` neutrino interaction generator [97], as a weighted sum of the cross-sections for target proton and target neutron:

$$\sigma = \frac{2\,\sigma(^1_1\mathrm{H}) + \sigma(^{16}_8\mathrm{O})}{18} \tag{4.4}$$

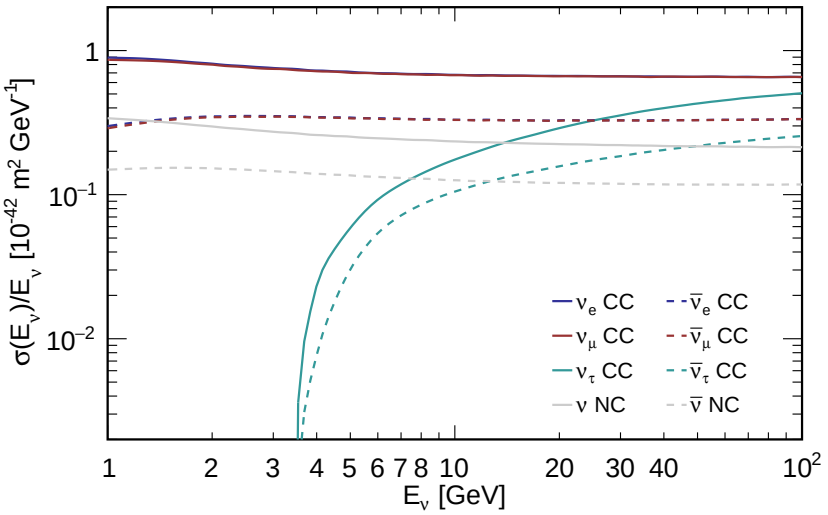

Figure 4.2: Total cross-section per nucleon divided by energy for each interaction channel ν_x.

To model the differential cross-sections $\frac{d\sigma}{dy}$, the distributions of Bjorken-y in neutrino interactions for each channel and each bin of true neutrino energy have been extracted from the ORCA Monte Carlo sample. In the MC simulation, neutrino interactions are generated with `gSeaGen`, based on `GENIE`. To extract these distributions, a 2-dimensional histogram (E_ν, y) is filled for each channel ν_x using the whole set of generated neutrino interactions, *i.e.* about 3×10^6 events in total (see Sec. 3). For each E_ν bin, the corresponding slice of the histogram is normalised to 1 with respect to the y variable, so that the bin content at coordinates (E_ν, y) is the probability $P_E(y)$ of Eq. 4.3.

For extracting these distributions, 40 logarithmic bins of $1\,\mathrm{GeV} \leq E_\nu \leq 100\,\mathrm{GeV}$, and 20 uniform bins of inelasticity $0 \leq y \leq 1$ are used. This histogram is rebinned depending on the binning in true y chosen for a specific analysis[2] and the values of $P_E(y)$ are taken from the rebinned histogram, with interpolation in the E_ν dimension.

[2]The true Bjorken-y binning is therefore required to be uniform with a compatible number of bins (1, 2, 4, 5, 10, 20).

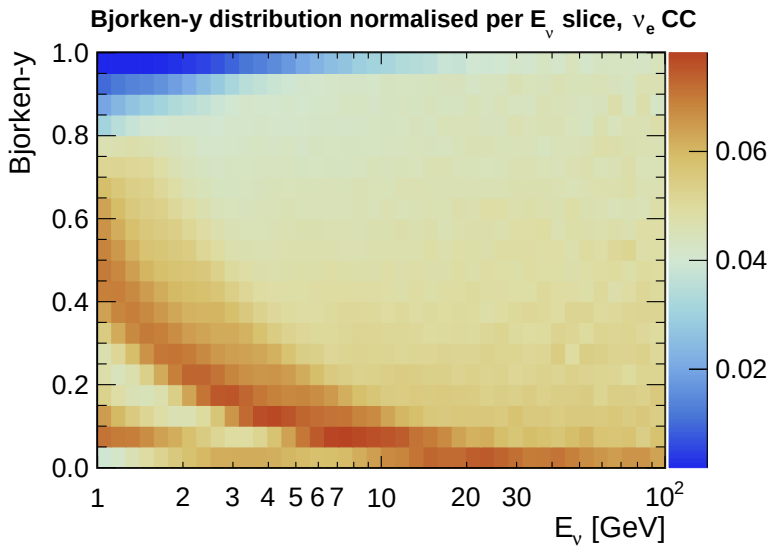

Figure 4.3: Example of Bjorken-y distribution $P_E(y)$ for 10 true Bjorken-y bins in the ν_e CC channel. Each energy slice is normalised to 1. Obtained from the full 'ORCA2016' sample of generated MC events (see Chap. 3).

4.2 Oscillation calculator and oscillograms

4.2.1 Earth model

The Earth model is based on the 42-steps model plotted on Fig. 2.11 and derived from the Preliminary Reference Earth Model (PREM) [116]. The concentric layers of constant matter density are called 'shells' hereafter, whereas the term 'layer' is reserved for compositional (chemical) layers which generally comprise several density shells. While the conventional Earth radius is $R_\oplus = 6371\,\text{km}$, the model has been modified as follows to account for ORCA's specificity:

- the last shell of Earth matter ends at $R = 6368\,\text{km}$,
- the shell from $R = 6368$ to $R = 6371\,\text{km}$ is composed of water, with an approximated constant density $\rho = 1 \cdot 10^3\,\text{kg.m}^{-3}$,
- an additional 15-km thick shell of atmosphere ($\rho = 1\,\text{kg.m}^{-3}$) is added, at the top of which neutrino production is assumed to occur.

The model thus has 44 shells in total, the detector being positioned at the bottom of the sea, *i.e.* at the inner boundary of the water shell. As explained in Sec. 2.3.3, an assumption on the average proton-to-nucleon ratio Z/A is needed to obtain the electron density in each shell according to Eq. 2.38. Seven distinct compositional layers are defined in the implementation. In each layer the value of the Z/A ratio is uniform. The full 44-layer model can be consulted online at Ref. [128] (see the `PremTables` directory). A more precise 425-layer (see Fig. 2.11) and simplified 15-layer models are available as well.

A summary of the compositional layers is given in Tab. 4.1. The assumed chemical composition for calculating the average Z/A in each layer is based on the *Geochemical Earth Reference Model* (GERM) database [255]. The most important elements entering in the composition of each layer are listed in Tab. 4.2.

[3] In the case of the inner core the thickness is ill-defined, the radius is indicated.

Layer	Id.	Thickness (km)	Shells	$[R_-,\ R_+]$ (km)	Approx. depth (km)	Z/A
Inner core	0	1221.5	4	0 - 1221.5	5100 - 6400	0.4691
Outer core	1	2258.5	20	1221.5 - 3480	2900 - 5100	0.4691
Lower mantle	2	2221	11	3480 - 5701	670 - 2900	0.4954
Upper mantle	3	645	5	5701 - 6346	22 - 670	0.4954
Crust	4	6346	2	6346 - 6368	0 - 22	0.4956
Water	5	3	1	6368 - 6371	-	0.5525
Atmosphere	6	15	1	6371 - 6386	-	0.4991

Table 4.1: Compositional layers in the Earth model used in the analysis. The model is defined in the `OscProb` package [128]. The columns indicate the layer index (with reference to the `OscProb` interface), the total thickness[3], the number of constant density shells, the exact innermost and outermost radius, the approximate range of depth, and the assumed Z/A value.

Core										
Element	Fe	Si	Ni	S	Cr	C	P	H	Mn	
wt%	85.5	6	5.2	1.9	0.9	0.2	0.2	0.06	0.03	
Z/A	0.4655	0.4984	0.4771	0.4989	0.4615	0.4996	0.4843	0.9901	0.4550	
Mantle										
Element	O	Mg	Si	Fe	Ca	Al	Na	Cr	Ni	Mn
wt%	44	22.8	21	6.26	2.53	2.35	0.27	0.26	0.2	0.1
Z/A	0.5000	0.4936	0.4984	0.4655	0.4990	0.4818	0.4785	0.4615	0.4771	0.4550

Table 4.2: Reference chemical composition used for the calculation of Z/A for the mantle and core in the `OscProb` Earth model, based on Ref. [255]. For each element the mass fraction (wt%) and Z/A based on the standard atomic weight are indicated. Differences in average chemical composition between the upper and lower mantle are very marginal and not accounted for here. In the second table (mantle) only the 10 dominant elements have been reported, following Ref. [255]. Additional elements (S, C, H, P) contribute for less than 0.1 wt% of the total.

4.2.2 Oscillation calculator

The `OscProb` neutrino oscillation calculator [128] has been used throughout the thesis. The context of the implementation in `OscProb` of the Earth model described in the previous section is discussed in Appendix C. Given a zenith angle θ_z the code calculates a neutrino trajectory (*baseline*) as represented on Fig. 4.4. The baseline starts at the assumed production height (15 km above the sea level) and ends at the detector position (3 km below the sea level). Denoting by N_s the number of distinct constant density shells intersected by the baseline, including the atmosphere and water shells, the baseline can be subdivided as a sequence of $N_{bl} = 2N_l - 3$ elementary steps (or *path segments*) defined by a length l_i, matter density ρ_i and proton-to-nucleon ratio $(Z/A)_i$:

$$S_i = (l_i,\ \rho_i,\ (Z/A)_i), \quad 0 \leq i < N_{bl} \tag{4.5}$$

Indeed all shells are traversed twice except the innermost layer as well as the atmosphere and water layers (traversed once before entering the Earth). Due to the radial symmetry of the Earth model the subdivision (S_i) is symmetrical after the first two steps:

$$S_{2+j} = S_{N_{bl}-j}, \quad 0 \leq j < N_{bl} - 2 \tag{4.6}$$

Given the position of the detector and production height, the expression for the total baseline length introduced in Sec. 2.2.1

$$L \simeq -2\cos\theta_z\, R_\oplus \tag{4.7}$$

is only approximate and valid for upgoing long-baseline trajectories. It becomes incorrect close to vertical, where the position of the detector with respect to the neutrino source at $H = 15$ km becomes dominant. For instance in the model used here, zenith angles $\cos\theta_z = 0.0,\ 0.2$ correspond respectively to baselines of 479 km and 87 km.

More generally, `OscProb` allows arbitrary definitions of a radial planetary model and position of the detector within this model, so that a detector located at the Earth surface or at any depth in the crust can be handled easily. To simulate a long-baseline accelerator neutrino experiment with neutrino production at the surface it is sufficient to remove the shell of atmosphere (this is provided as an alternative version of the 423-layer Earth model).

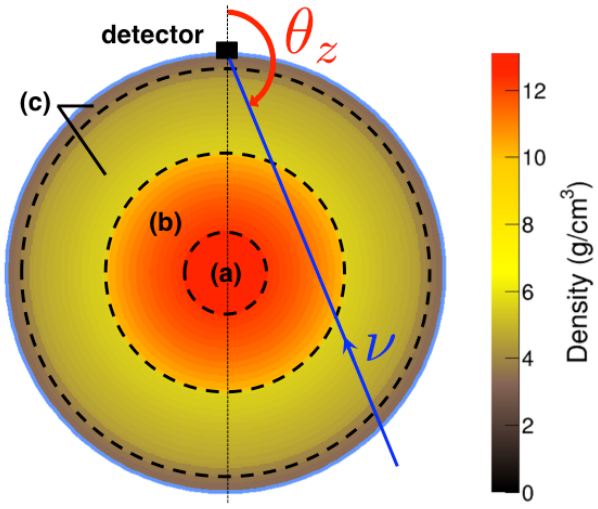

Figure 4.4: The 42-layers radial model of the Earth used in this analysis (PREM) [116]. The following major compositional layers are indicated: inner core (a), outer core (b), silicate Earth (c). With respect to Tab. 4.1 the silicate Earth includes the lower and upper mantle as well as the crust.

The method for calculating oscillation probabilities follows the approach described in Sec. 1.2.4: the quantum evolution operator of the neutrino state is diagonalised for each constant density step, and the overall evolution follows Eq. 1.59. For the three-flavour mixing studied in this thesis, dedicated algorithms for fast numerical diagonalisation of 3×3 hermitian matrices are employed [256]. These algorithms were developed for the oscillation calculator of the widely used *General Long Baseline Experiment Simulator* (`GLoBES`) package [257, 258]. I performed a detailed validation of `OscProb` against `GLoBES` at the earliest stage of developing this analysis.

`OscProb` also allows the use of extended mixing and propagation models, including $3+N$ models (N sterile neutrino states), non-standard interactions, and neutrino decoherence [259]. Calculations are significantly less efficient in these more general cases as the simplifications exploited by Ref. [256] are no longer valid (standard linear algebra libraries are used instead). Taking advantage of this flexibility, the analysis framework I developed is designed to allow the study of both the standard 3ν mixing models and more exotic ones, and is being used for both purposes within the KM3NeT collaboration. This is further discussed in Appendix C.

Reference oscillation parameters

Throughout the thesis the reference oscillation parameter values follow the 'NuFit3.2' [63] global fit values reported in Tab. 1.1, with some exceptions. Firstly, θ_{13} is set to the NH best-fit value ($\theta_{13} = 8.54°$) independently of the hierarchy (note that the IH best-fit from Tab. 1.1 is very close). Besides, the parameters θ_{23} and δ_{CP} still have large associated uncertainties and have a significant impact on NMH sensitivities, thus the best-fit values of Tab. 1.1 are not used: instead the dependance of the results on the true values of both these parameters is generally considered. If unspecified, our default values are $\delta_{\mathrm{CP}} = 0$ and $\sin^2\theta_{23} = 0.5$ (maximal mixing).

4.2.3 Discussion of oscillograms

Oscillation probabilities are calculated as the first step of the computation chain and stored in so-called 'oscillograms': two-dimensional histograms $(E_\nu, \cos\theta_z)$ with the same binning scheme as interacting event histograms (usually logarithmic in energy). One value of the oscillation probability is computed at each logarithmic bin center. During a fit (likelihood maximisation), oscillation probabilities are retrieved from these histograms, which are only recomputed if the trial oscillation values or Earth parameters have changed.

Densely sampled oscillograms, calculated with the reference oscillation probabilities, are shown for all relevant oscillation channels assuming a normal hierarchy in Fig. 4.6 and an inverted hierarchy in Fig. 4.7. Note that the interval $0 \leq \cos\theta_z \leq 0.2$, corresponding to a fraction of the downgoing neutrino flux, is shown on the oscillogram. Oscillation probabilities in this range of zenith are small but non-zero at low energy due to the shortening of the oscillation length.

As expected from the theoretical discussion of oscillation probabilities in Chap. 1 and 2, the effect of the MSW resonant enhancement can be seen for neutrinos in NH and antineutrinos in IH. In the resonant channels, the impact of the density discontinuity at the core-mantle boundary is clearly visible at $\cos\theta_z \simeq -0.84$. For trajectories crossing the mantle only ($\cos\theta_z > -0.84$), the enhancement is the most intense for the first oscillation extremum, around $E_\nu \sim 7\,\text{GeV}$ and for trajectories crossing the deep mantle ($\cos\theta_z < -0.4$). For core-crossing baselines several intense oscillation extrema appear in the $\sim 2-8\,\text{GeV}$ energy range, as an effect of the mantle-core-mantle parametric enhancement.

In oscillation channels where the MSW resonance is absent, the $\nu_\mu \leftrightarrow \nu_e$ and $\nu_e \rightarrow \nu_\tau$ channels remain essentially 'closed', as they are controlled by $\sin^2\theta_{13} \sim 0.02$. For instance, in the NH case this results in $P(\overline{\nu}_\mu \rightarrow \overline{\nu}_\mu) \simeq 1 - P(\overline{\nu}_\mu \rightarrow \overline{\nu}_\tau)$ and $P(\overline{\nu}_e \rightarrow \overline{\nu}_e) \simeq 1$ (two-neutrino approximation). Moreover, in those channels the L/E dependence of the oscillation probabilities is preserved: the lines of constant probability correspond to constant L/E. As can be seen more clearly from Fig. 4.5, this structure is also approximately preserved far from the resonance region in the resonant oscillation channels. The saturation regime induced by matter effects at high energy and high densities, is visible in both resonant and non-resonant oscillograms as a suppression of the $\nu_\mu \leftrightarrow \nu_e$ oscillations ($\sin^2\theta_{13}^M < \sin^2\theta_{13}$).

As discussed in Chap. 1 and Chap. 2, flipping simultaneously $\nu \rightleftharpoons \overline{\nu}$ (CP symmetry) and the mass hierarchy NH $\rightleftharpoons$ IH results in visually indistinguishable oscillograms. Similarly, to first order the time reversal (T) leaves the probabilities unchanged $P(\nu_a \rightarrow \nu_b) \simeq P(\nu_b \rightarrow \nu_a)$. The second-order effects breaking those symmetries are due to i) intrinsic CP violation, which is inexistant here since $\delta_{\text{CP}} = 0$ (though the effect would hardly be noticeable by eye even with $\delta_{\text{CP}} = \pi/2$); and ii) extrinsic violation due to the non-CP symmetric matter profile.

Finally, the similarity $P(\nu_e \rightarrow \nu_\mu) \simeq P(\nu_e \rightarrow \nu_\tau)$ is less fundamental; it is a consequence of assuming maximal mixing.

L/E_ν oscillograms and sampling

The computation of the oscillation probabilities is time-consuming and has to be performed from several hundreds to several thousand times in analysis where oscillation or Earth parameters are allowed to vary, *e.g.* as nuisance parameters in likelihood maximisations. In order to keep manageable CPU consumption the sampling density used in the analysis is generally lower (larger bin widths) than shown in the examples of Fig. 4.6 and 4.7.

Fig. 4.5(a) shows a $(E_\nu,\ \cos\theta_z)$ oscillogram computed with an example of practical binning density. In Fig. 4.5(b), the same oscillogram is shown in an $(L/E_\nu,\ \cos\theta_z)$ sampling space, with $L \simeq -2\cos\theta_z\, R_\oplus$ the total baseline. The correspondance between the two sampling schemes is also shown. This illustrates more clearly how the overall L/E_ν structure of the oscillatory pattern is perturbated by the matter resonance, implying that binning events in L/E_ν only

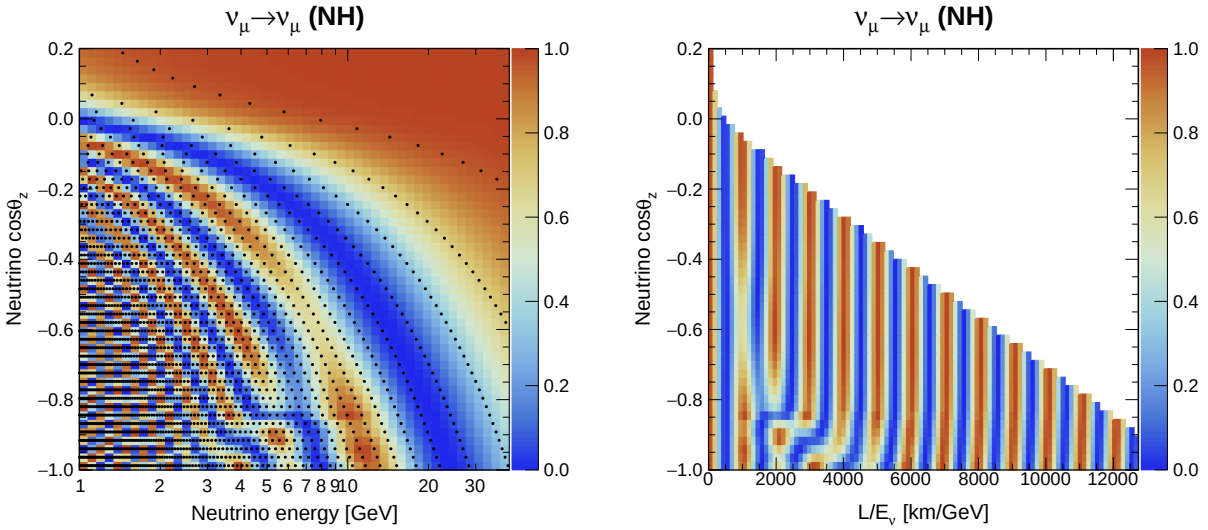

(a) (b)

Figure 4.5: Comparison of $(E_\nu,\ \cos\theta_z)$ and $(L/E_\nu,\ \cos\theta_z)$ oscillograms and sampling. The oscillation probabilities are calculated at the same number of points in both grids, and the black dots in Fig. (a) indicate the positions of the centers of the bins (sampling points) of histogram (b).

would not capture the whole available information, unlike in the case of vacuum oscillations. Nevertheless this structure can be retrieved, as expected, outside of the resonance region. An analysis strategy histogramming observed events according to the reconstructed $(L/E_\nu,\ \cos\theta_z)$ could be considered. As a benefit, this would yield more uniformly populated bins since the interaction rate roughly follows an E^{-2} power law. However, those two variables are not independent, and the errors on both the energy and angular measurement would be combined in the L/E_ν variable.

Even when histogramming observed events according to $(E_\nu,\ \cos\theta_z)$ (as done in this thesis), the comparison of the sampling schemes shown on Fig. 4.5 suggests that calculating oscillation probabilities in a $(L/E_\nu,\ \cos\theta_z)$ equidistant grid could be significantly more efficient in order to sample the fast oscillations seen at low energy.

4.3 Interaction rates and NMH signal

Fig. 4.8 shows the results of the calculations of the expected rates of CC events per unit exposure for each neutrino flavour, accounting for the absence of $\nu/\overline{\nu}$ discrimination ($\nu+\overline{\nu}$ rates). The plotted quantity is a differential rate $\mathrm{d}N/\mathrm{d}E$, multiplied by E^2 as in Chap. 2 (see Fig. 2.16 -2.25). From the left and middle panels of Fig. 4.8 it appears clearly that the $\nu_\mu \rightarrow \nu_\tau$ oscillation channel remains the dominant feature. Experimentally, in the few-GeV range this indeed amounts to ν_μ *disappearance* due to the effect of the τ threshold suppressing ν_τ interactions at low energy.

The NMH signal is shown qualitatively as a relative difference between the NH and IH expectation. While the presence or absence of the MSW resonance in ν or $\overline{\nu}$ oscillation channels is an unequivocal signature of the NMH at the level of oscillation probabilities, this signature is somewhat blurred in the observed event rates due to the coexistence of ν_e, $\overline{\nu}_e$,ν_μ, $\overline{\nu}_\mu$ in the flux, combined with the symmetries discussed in the last section. Nevertheless, as already discussed and observed in Sec. 2.2.2, relative differences between the NH and IH ($\nu+\overline{\nu}$) expected interaction rates reach the level of $\sim 20\%$ on an extended region of the $(E_\nu,\ \cos\theta_z)$ planes. In the muon and tau channels the relative difference reaches up to 100% and more, but the signal is localised at oscillation extrema. It should also be noted that the relative event difference has no interpretation in terms of statistical significance without an indication of the total exposure and scaling of the event count in each channel. As an example the very large relative difference at low energy in the tau channel will generally be less statistically significant experimentally (even assuming perfect detection performance) than a few % difference in a more populated region of the electron or muon channel. A more relevant sensitivity indicator will be used in Chap. 8.

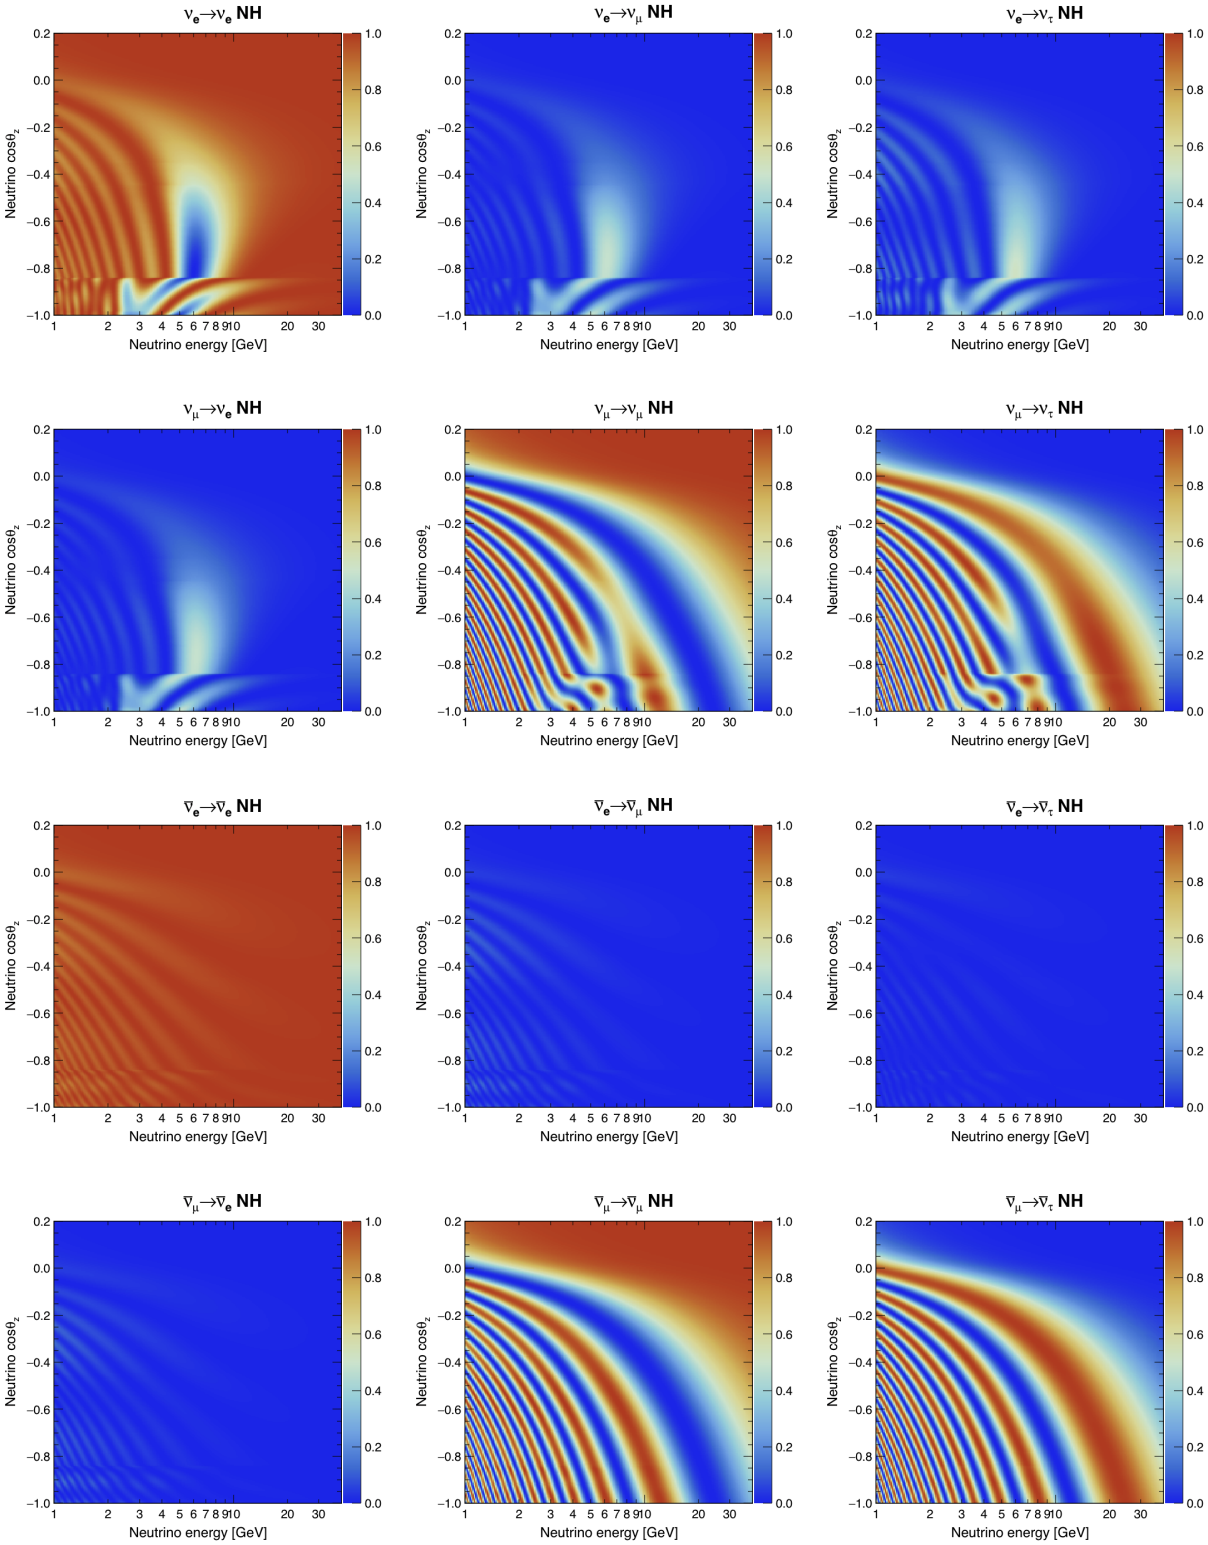

Figure 4.6: $(E_\nu, \cos\theta_z)$ oscillograms for all relevant oscillation channels, assuming a normal neutrino mass hierarchy (NH). The oscillation parameters are set to their reference values (maximal mixing, $\delta_{\mathrm{CP}} = 0$).

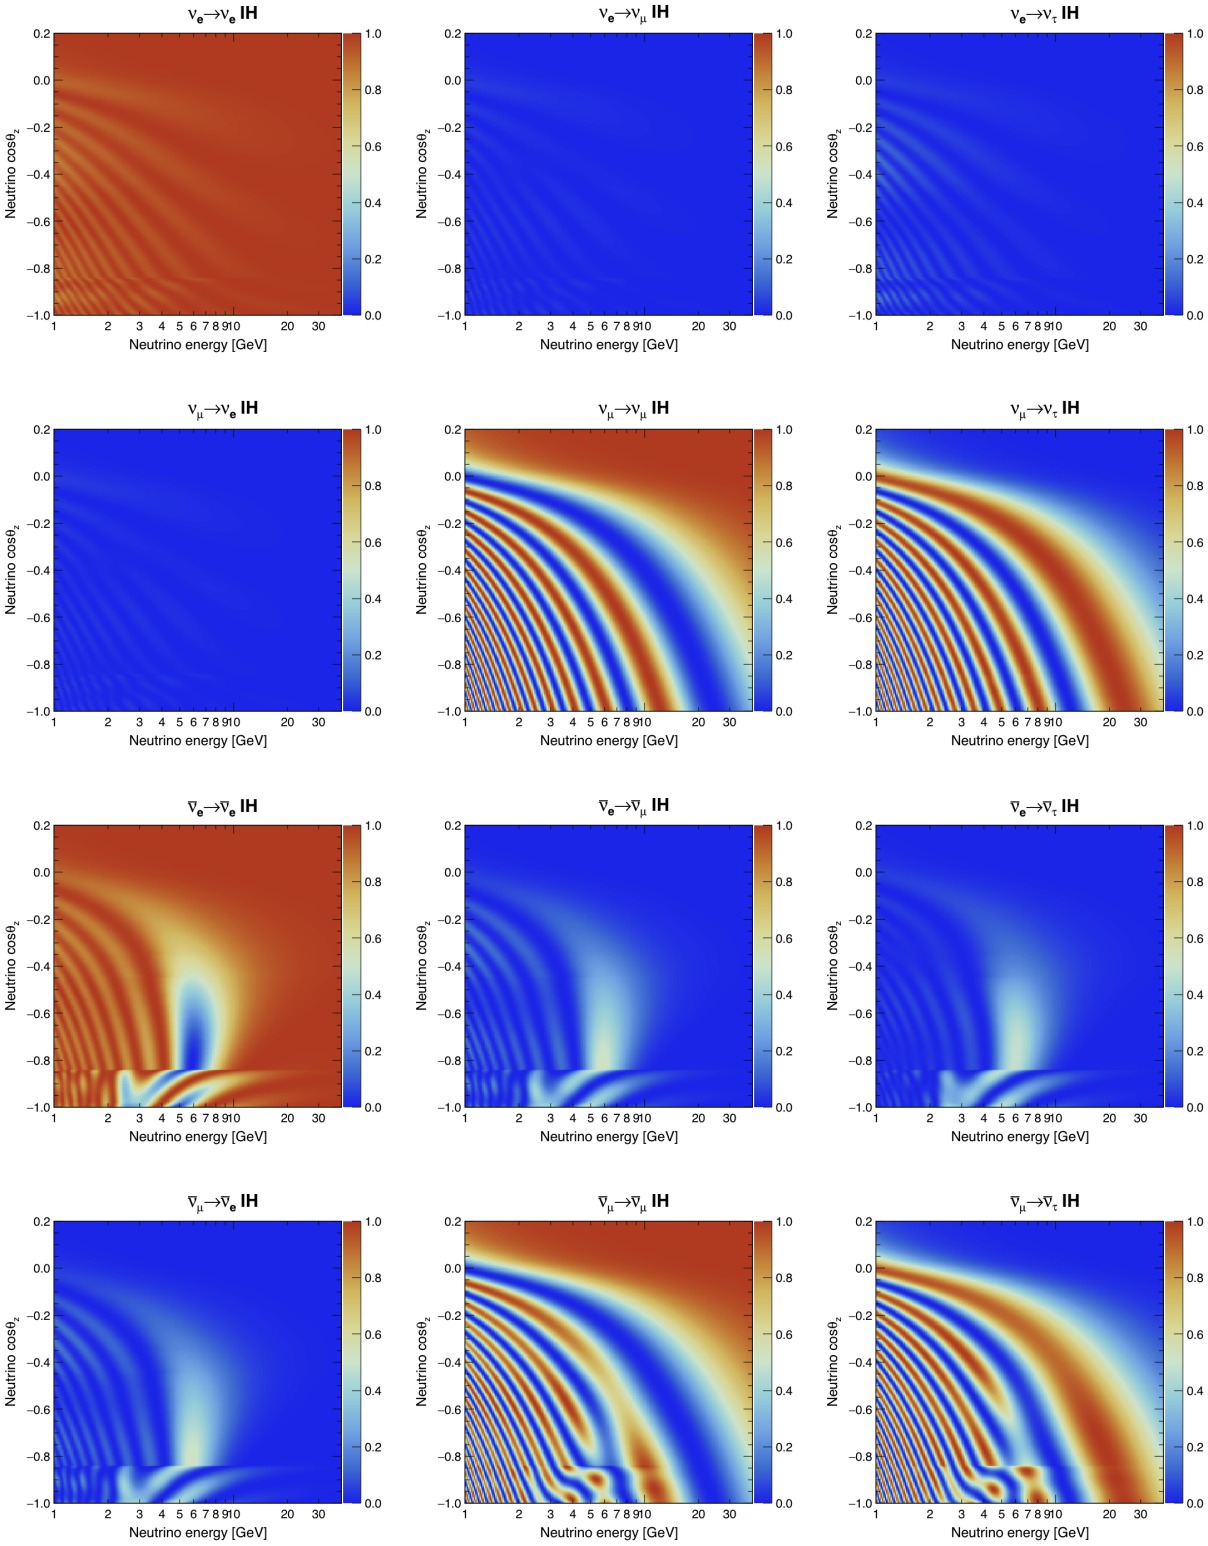

Figure 4.7: $(E_\nu,\ \cos\theta_z)$ oscillograms for all relevant oscillation channels, assuming an inverted neutrino mass hierarchy (IH). The oscillation parameters are set to their reference values (maximal mixing, $\delta_{\mathrm{CP}} = 0$).

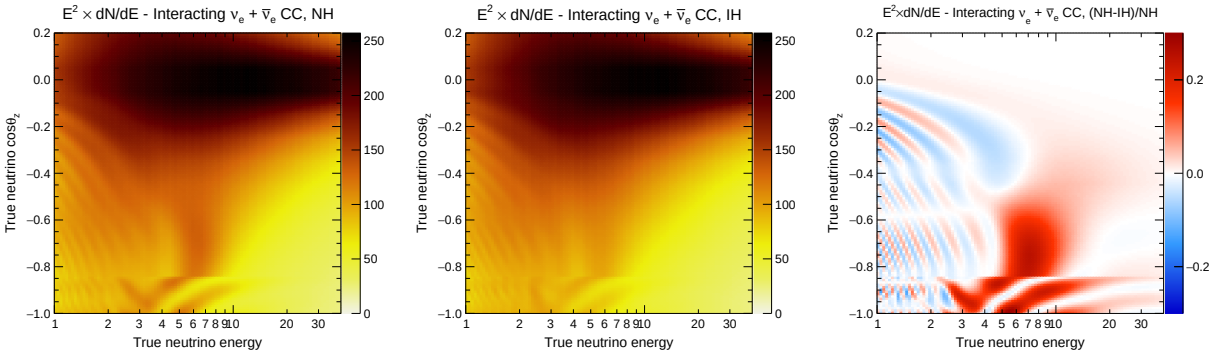

$\nu_e + \overline{\nu}_e$ CC events

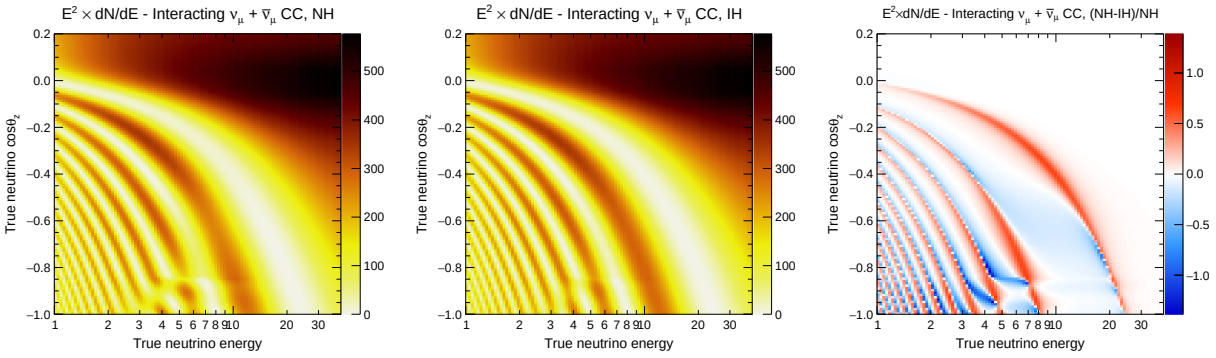

$\nu_\mu + \overline{\nu}_\mu$ CC events

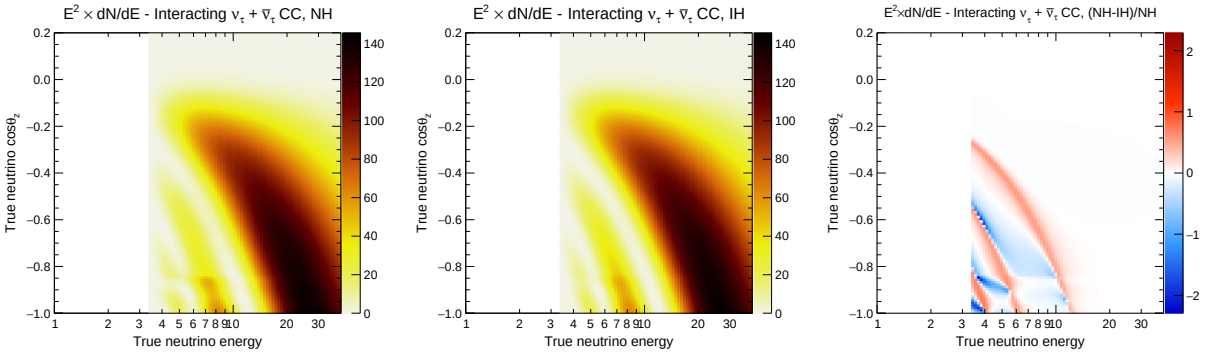

$\nu_\tau + \overline{\nu}_\tau$ CC events

Figure 4.8: Rate $\mathrm{d}N/\mathrm{d}E$ of $\nu + \overline{\nu}$ CC events, multiplied by E^2, shown as a function of $(E_\nu, \cos\theta_z)$. The left and middle panels show the expectation for NH and IH, while the rightmost panel shows the relative difference (NH-IH)/NH.

Chapter 5

Modeling the detector response

Contents

Once tools for the calculation of neutrino interaction rates are in place, the next ingredient for predicting observed event distributions is a model for the response of the detector to neutrino interactions. Denoting as $\boldsymbol{x}$ the ensemble of true characteristics for a neutrino event and $\boldsymbol{x'}$ the measured characteristics (reconstruction, classification) for an event selected in the analysis sample, the predicted event count in a bin of the variables $\boldsymbol{x'}$ is obtained schematically as

$$n_{\text{reco}}(\boldsymbol{x'}) = \sum_{\boldsymbol{x}} R(\boldsymbol{x}, \boldsymbol{x'}) \times n_{\text{int}}(\boldsymbol{x}) \tag{5.1}$$

where $R(\boldsymbol{x}, \boldsymbol{x'})$ represents the conditional probability for an event occuring in true bin $\boldsymbol{x}$ to be detected, successfully reconstructed, selected and classified in the bin of measured characteristics $\boldsymbol{x'}$. Written as such, the probability R incorporates the totality of the effects related both to the detector and to the experimenter's methodology. Modeling this response is a challenging task. This chapter goes through the methods I developed and implemented towards this aim for ORCA neutrino oscillation studies.

The initial method is based on the direct sampling of $R(\boldsymbol{x}, \boldsymbol{x'})$ as a fully correlated response matrix built from Monte Carlo events. The motivation for this approach is exposed in Sec. 5.1. Formal definitions are given in Sec. 5.2, along with some illustrations of the method and of the ORCA-specific response features. Finally, the impact of finite Monte Carlo statistics and the subsequent adaptations of the method are adressed in Sec. 5.3, including the comparison of two techniques to evaluate the impact of Monte Carlo statistical uncertainties on the final predictions $n_{\text{reco}}(\boldsymbol{x'})$.

5.1 Monte Carlo based detector model

5.1.1 Motivation

A specificity and founding motivation of my analysis was the intent to take into account as precisely as possible the detailed features of the detector response. Specific aspects which were identified as potentially important but ignored or incompletely accounted for in previous analyses included the interplay of the energy-angular reconstructions and event classification, the dependence of the response functions on the inelasticity of neutrino interactions, as well as the reconstruction capabilities for said inelasticity. The role of inelasticity with potential sensitivity improvement associated to its measurement, and more generally the correlations between reconstruction variables were not considered in previous sensitivity studies, and were targeted to be modeled here.

To this end, the natural approach is to base detector modeling on event-by-event Monte Carlo simulations[1]. Allowing such event-by-event studies with additional analysis binnings (*e.g.* inclusion of measured Bjorken-y) was one of the motivations for the increase in Monte Carlo statistics in the latest production, presented in Chap. 3. A technique frequently used to simulate experiments in a full Monte Carlo approach and obtain expected event distributions in binned form is event-by-event reweighting. Here, one would perform a loop on selected MC events and for each event fill the relevant bin $\boldsymbol{x'}$ with a weight obtained as a product of an intrinsic event weight, related to the Monte Carlo generation scheme, and a reweighting factor depending on $n_{\text{int}}(\boldsymbol{x})$ and thus on the test values of the model parameters.

An important technical requirement to take into consideration is computation speed. The study of systematic uncertainties typically involves repeating the whole computation chain many times, either in the process of maximising or marginalising a likelihood function with respect to nuisance parameters and/or when drawing pseudo-experiments for frequentist studies. The computation method was therefore designed with the objective that the whole chain be executed in $\mathcal{O}(1\,\text{s})$. When using a genuine event-by-event approach, this speed requirement somewhat conflicts with the requirement of maximising Monte Carlo statistics: even though the weights can be approximated or interpolated from a pre-computed grid, the time needed to loop on $\mathcal{O}(10^6)$ MC events stored in a ROOT tree is non-negligible, and grows linearly with the number of entries.

A solution to this conflict is to summarise the detector response properties in a binned manner by filling events, accounting for their intrinsic weights, into a multi-dimensional histogram. Such a multidimensional response histogram is often called *response matrix*, interpreting Eq. 5.1 as a matrix-vector multiplication. This way, the loop and the reweighting can be done on the filled bins of the matrix rather than on individual events, thereby limiting the growth of running time with MC statistics. The dimensionality of the response matrix depends on the number of variables included and on the binning of continuous variables, which can be set in accordance with the analysis requirements. To limit the computation time only the variables most relevant to the analysis are kept, at the cost of losing some individual event information.

5.1.2 Correlated response matrix

In this analysis, the true variables $\boldsymbol{x}$ are limited to the interaction channel ν_x, energy E, zenith angle θ and inelasticity y. The measured variables $\boldsymbol{x'}$ include the same continuous variables E', θ' and y' as well as a discrete event category indexed by i. In the standard analysis strategy ('two-class strategy'), i has three possible values corresponding to events classified as track-like, shower-like, or rejected. This classification is typically based on cuts applied on

[1] Mathematically, one estimates the probability density function represented by R by means of Monte Carlo integration.

containment, reconstruction quality parameters, and scores provided by the background rejection and track/shower classifiers. This standard strategy is detailed in Sec. 5.1.3, along with alternatives which define a larger number of classes.

The full response matrix R is thus meant to represent the probability density for the transformation:

$$\{\nu_x, E, \theta, y\} \rightarrow \{\text{rejected}\} \cup \{i, E', \theta', y'\} \tag{5.2}$$

In the method referred to in this thesis as "correlated smearing", this transformation is applied in a fully correlated manner. This is detailed in section 5.2: I explain how the entries of the response matrix are calculated based on ratios of MC event counts, and break down these entries into parts which have an easier physical interpretation: detection and selection efficiency, effective mass, (mis-)classification probabilities, and conditional PDFs describing the energy, angular, and inelasticity response for each discrete event category.

The available number of Monte Carlo events is large, but finite. To build the response matrix these events are binned in 6 continuous dimensions for each discrete category, some of the categories being sparsely populated. As a result, the statistical uncertainties associated to the calculated entries of R are often non-negligible. In Sec. 5.3, I show that the effect of these uncertainties on sensitivity estimations must be accounted for, and that in the context of sensitivity studies the sparseness of the response leads to overestimating the performance of the measurement. This phenomenon is referred to in the thesis as "Monte Carlo sparseness effect" and is studied from a mathematical point of view in Appendix B. Solutions to limit the impact of this effect are presented, including adaptations of the correlated method in Sec. 5.3.2 and an alternative uncorrelated smearing approach in Sec. 5.3.3. Making precise estimates of the aforementioned Monte Carlo statistical uncertainties is necessary in order to include them in the statistical data analysis or sensitivity study. Two estimation techniques are presented in Sec. 5.3.4, based on bootstrapping and on classical combination of binomial errors.

5.1.3 Event selection and classification

A succession of filters and selection cuts is applied to the data, starting with L0 filtering and triggering, as described in Chap. 3. Triggered events are then processed in parallel by the reconstruction algorithms (`recoLNS`, `JGandalf`, `Dusj`). Each reconstruction applies a set of preselection cuts before attempting to reconstruct an event. In simulated MC data, the sample of events that are processed at least by one reconstruction goes on to the 'pre-PID' filtering stage. A further filtering (as known as 'pre-PID cuts') is applied at this stage to select the sample of events to be processed by the classifiers. The main objectives of these cuts are i) to reduce the atmospheric background by several orders of magnitude and ii) to reject events for which none of the reconstructions has provided a 'successful' output (for instance, an event is considered as successfully reconstructed by the track reconstruction if the likelihood of the reconstructed track exceeds a predefined minimal value). An event is thus selected if it satisfies at least one of the *"good track"* or *"good shower"* conditions, where

- *"good track"* requires that the event satisfies the quality cuts, is reconstructed as upgoing and has a reconstructed vertex contained within the detector volume for both `JGandalf` and `recoLNS` reconstructions;
- *"good shower"* requires that the event satisfies the `Dusj` quality cuts and is reconstructed as upgoing by `Dusj`.

The containment and directional conditions are efficient at rejecting atmospheric muons, while quality cuts help rejecting both atmospheric muons and pure noise events.

The events passing the PID pre-cuts constitute the minimal selection level for events entering in the oscillation analysis. Two denominations will be used for event samples throughout this chapter:

- *generated events* refer to all events simulated at the `gSeaGen` level (physically representing all neutrino interactions);
- *selected events* refer either to the whole sample of events passing the PID pre-cut, or to a subset of this sample restricted by additional selection cuts (defining an event class i).

In practice, the full output of the PID classification is converted into a lightweight format containing all variables potentially useful for the analysis. Similarly, a lightweight tree is constructed from the output files of the `gSeaGen` simulation runs (care is taken to remove runs that are not present in the final PID output due to occasional failures of the simulation chain).

Analysis strategies

As seen in Sec. 3.3.2, PID classifiers provide two anti-background classification scores $p_{\text{bkg,noise}}$ and $p_{\text{bkg},\mu}$ as well as one track/shower classification score p_{track} for each event. In the standard analysis strategy, the definition of event classes proceeds by cutting first on $p_{\text{bkg,noise}}$ and $p_{\text{bkg},\mu}$ so as to reject background (thus defining a background class) and then forming two event classes out of the remaining events: 'track-like' and 'shower-like', defined by a single cut p_{cut} on the track/shower score:

- Events classified as track if $p_{\text{track}} > p_{\text{cut}}$,
- Events classified as shower if $p_{\text{track}} \leq p_{\text{cut}}$.

The oscillation analysis then performs a joint fit (as detailed in Chap. 6) on these two signal channels, ignoring the background class. For reconstructed values, the track channel uses the `JGandalf` output if available (`recoLNS` otherwise), while the shower channel uses the `Dusj` output.

As an alternative strategy to further improve the purity of each signal sample, a stricter selection on the track/shower score can be achieved by defining two cuts $0 < p_1 < p_2 < 1$ and

- events classified as "good tracks" if $p_2 < p_{\text{track}}$,
- events classified as "mixed sample" if $p_1 < p_{\text{track}} \leq p_2$,
- events classified as "good showers" if $p_{\text{track}} \leq p_1$.

Although the events classified in the mixed sample have a lower purity in terms of flavour, their addition as a third channel in the fit is potentially interesting: even under the pessimistic assumption that this event sample carry no information on flavour oscillations, it may prove useful in constraining nuisance parameters such as flux normalisation and shape.

Following the above reasoning it seems interesting to evaluate the potential of a strategy defining an increased number of PID classes, up to the case where events are effectively binned in terms of their track/shower PID score. Therefore the possibility to define an arbitrary number (denoted N_{class}) of such event classes has been implemented in the analysis framework. In practice, each class is defined by a combination of cuts on the variables available in the selected events sample. In addition, the definition of a class must provide the identifier of the reconstruction variables (*e.g.* `JGandalf`, `recoLNS`, `Dusj`) to be used for events falling in this class. The classes are defined with a priority rank such that overlapping classes definitions do not result in double counting of events.

5.2 Correlated smearing

5.2.1 Definition of the response matrix entries

The *correlated* approach uses the full set of MC events to compute the entries of the response matrix. For each interaction channel ν_x and each classification i, a 6-dimensional response matrix $R^{[\nu_x \to i]}(E, \theta, y, E', \theta', y')$ is defined. Each entry of this matrix summarises in a single dimensionless coefficient the efficiency of detection, classification and probability of reconstruction for a given true bin (E, θ, y). An entry is computed as the ratio of the number of "selected" (i.e. triggered, reconstructed and classified) MC events to the number of generated MC events in the corresponding true bin:

$$R^{[\nu_x \to i]}(E, \theta, y, E', \theta', y') = \frac{N_{\text{sel}}^{\text{MC}}[\nu_x \to i](E, \theta, y, E', \theta', y')}{N_{\text{gen}}^{\text{MC}}[\nu_x](E, \theta, y)} \tag{5.3}$$

Implementation

In practice, the set of response matrices is implemented as a single 7-dimensional sparse histogram, using the class `THnSparse` provided by `ROOT` [260]. 6 dimensions are used for continuous variables (3 true, 3 reconstructed), and an additional dimension is reserved for a discrete integer flag f indexing the interaction channel and reconstructed class, with a one-to-one correspondance $(\nu_x,\ i) \leftrightarrow f$. This full 7-dimensional response matrix R may be used in the following as an alias of the set of matrices $\{R^{[\nu_x \to i]}\}$.

Building the response matrix involves looping on both selected and generated events and filling a high-dimensional histogram, and is thus computationnally expensive. However, once the matrix has been built, looping on its entries is very fast (comparable to looping on entries of a `TTree`). Concretely, a response matrix defined by a typical binning in the seven dimensions $(E,\ \theta,\ y,\ E',\ \theta',\ y',\ f)$ of

$$20 \times 40 \times 4 \times 20 \times 40 \times 1 \times 32 \text{ bins} \tag{5.4}$$

amounts to a total of about $8.2 \cdot 10^7$ bins ($5.2 \cdot 10^8$ when counting in the underflow and overflow bins of each axis). Such a matrix, filled with the $\mathcal{N} \simeq 6.1 \cdot 10^6$ selected MC events of the 'ORCA2016' sample described in Sec. 3.2.4, has about $1.9 \cdot 10^6$ non-zero entries (*i.e.* filled bins). The fraction of filled bins is only 0.35%, justifying the use of a sparse histogram implementation, where empty bins are not stored in memory.

With respect to the event-by-event reweighting approach, the gain is thus somewhat moderate size-wise, because of the fairly high dimensionality of the matrix and of the moderate size of the selected MC sample. However, an additional gain of computation time arises from the facts that events are already regrouped into the true and reconstructed binning used in the analysis, and no additional weights need to be computed on-the-fly.

In practice, the event-by-event reweighting method currently developed in parallel within the collaboration takes $\sim 50\,\text{s}$ to simulate an 'experiment' (*i.e.* obtain a set of reconstructed event distributions), versus $\sim 1\,\text{s}$ for the method described here – precomputed grids of oscillation probabilities being used in both cases. The size of MC productions for comparable currently running experiments (*e.g.* ANTARES, SK, IceCube), and foreseen for the full ORCA detector, are larger in statistics by several order of magnitudes. In a response matrix approach the fraction of filled bins is also increased with increased MC statistics, as the response is better sampled (*e.g.* the tails of the response functions become more populated). Nevertheless, the relative gain in running time with respect to an event-by-event approach is expected to increase with MC statistics. Binning schemes for the response matrix are discussed in more detail in Sec. 5.3.2. Some of the drawbacks of the approach, as compared to event-by-event, are discussed in Chap. 7.

MC generation volume

In the real experiment, events originating from a neutrino interaction vertex situated outside of the instrumented volume can be triggered and successfully reconstructed. Therefore, it is necessary to simulate events in a reference volume larger than the instrumented volume itself. This is especially true for higher energy events, which produce long tracks and/or large amounts of Cherenkov light, as these are more likely to reach the detector from outside the instrumented volume. The actual generation volume used by the `gSeaGen` code hence depends on the true energy, as can be seen from Fig. 5.1. In the *can volume* (energy-independant), `gSeaGen` writes all generated events to the output. These events are generated with a uniform distribution of interaction vertex positions, as can be verified from the figure. However, the generation volume is extended beyond the can volume by an additional, energy-dependent volume. In this volume, events which will not produce light susceptible to reach the detector are immediately discarded and not written to the output.

Therefore, one must choose a reference volume V_{ref} where all generated events are written out by `gSeaGen`. Events having an interaction vertex outside this reference volume are not included in the count $N_{\text{gen}}^{\text{MC}}[\nu_x]$ of Eq. 5.3 since only a fraction of those is present in the `gSeaGen` files. On the contrary, all selected events are included in the count $N_{\text{sel}}^{\text{MC}}[\nu_x \to i]$ regardless of the position of their interaction vertex. Thus, one needs to multiply the response matrix entries with numbers of neutrino interactions expected to occur *in the reference volume*. This calculation is correct as long as the events are generated uniformly in space, with the same density both in the reference volume and in the extended generation volume. Indeed, denoting V' as the largest generation volume used by `gSeaGen` ($V' > V_{\text{ref}}$), the assumption of uniform generation writes:

$$\frac{N_{\text{gen}}^{\text{MC}}(V')}{V'} = \frac{N_{\text{gen}}^{\text{MC}}(V_{\text{ref}})}{V_{\text{ref}}}, \tag{5.5}$$

which implies:

$$V' \times \frac{N_{\text{sel}}^{\text{MC}}(\text{all})}{N_{\text{gen}}^{\text{MC}}(V')} = V_{\text{ref}} \times \frac{N_{\text{sel}}^{\text{MC}}(\text{all})}{N_{\text{gen}}^{\text{MC}}(V_{\text{ref}})}. \tag{5.6}$$

In practice, in order to maximise the MC statistics the typical choice of reference volume is either the `gSeaGen` can volume (largest volume where the distribution of MC vertices in `gSeaGen` files is known to be uniform independently of the energy), or a slightly smaller volume in order to avoid potential edge effects.

Finally, we assume that V' is sufficient, in the sense that no high-energy events having an interaction vertex outside of V' would ever be selected.

Application of the detector response

With the above definitions, the number $n_{\text{reco}}^{i}(E',\theta',y')$ of reconstructed events expected to be classified as PID class i writes:

$$n_{\text{reco}}^{[i]}(E',\theta',y') = \sum_{\nu_x} \sum_{E,\theta,y} R^{[\nu_x \to i]}(E,\theta,y,E',\theta',y') \times M_{\text{ref}} \times \frac{\mathrm{d}n_{\text{int}}^{[\nu_x]}}{\mathrm{d}M}(E,\theta,y) \tag{5.7}$$

where M_{ref} is the mass of water corresponding to the reference generation volume used for computing the response matrix entries:

$$M_{\text{ref}} = \rho_{\text{water}} \times V_{\text{ref}} \tag{5.8}$$

and $\frac{\mathrm{d}n_{\text{int}}^{[\nu_x]}}{\mathrm{d}M}(E,\theta,y)$ is the number of expected interacting events per unit of target mass in channel ν_x.

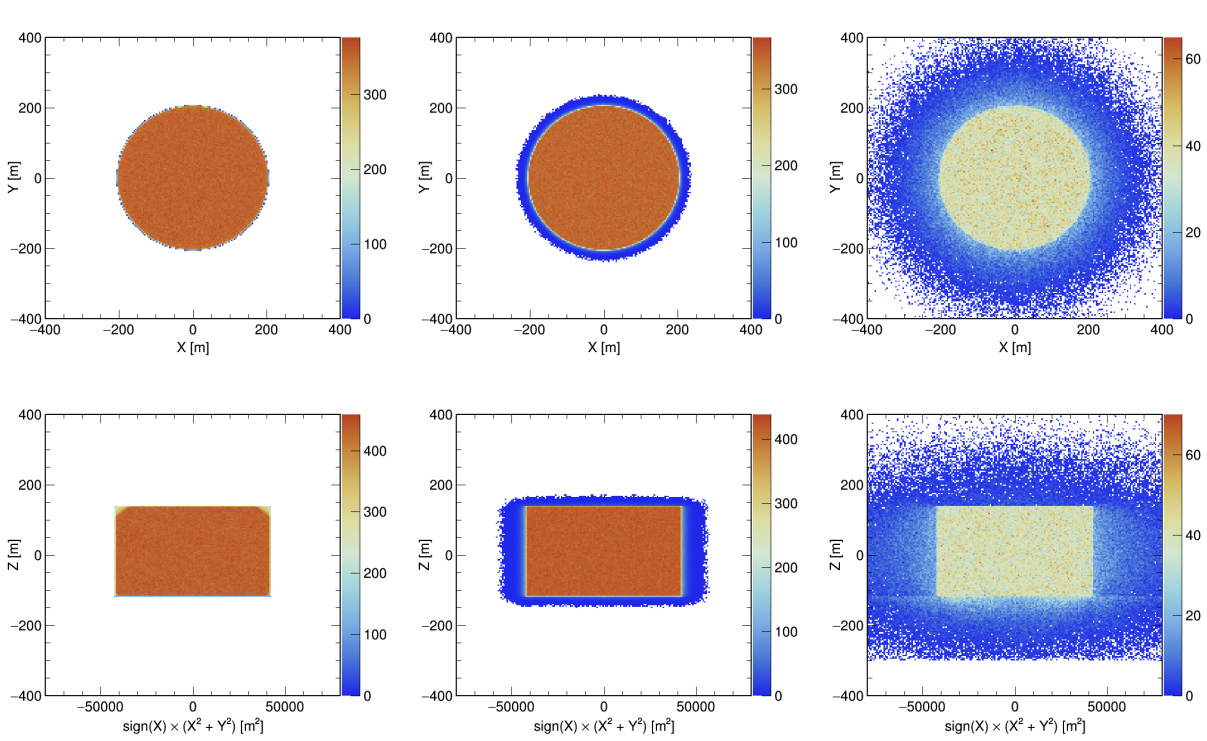

Generated $\nu_e + \overline{\nu}_e$ CC $1 < E_\nu/\text{GeV} < 100$

Generated $\nu_\mu + \overline{\nu}_\mu$ CC $1 < E_\nu/\text{GeV} < 10$

Generated $\nu_\mu + \overline{\nu}_\mu$ CC $50 < E_\nu/\text{GeV} < 100$

Figure 5.1: Spatial distribution of the interaction vertices of the neutrino events generated by `gSeaGen`. The top and bottom panels correspond to a top view and a side view of the volume surrounding the detector. The `gSeaGen` can is a cylinder of radius $R = 205.4\,\text{m}$ ($R^2 \simeq 4.2 \cdot 10^4 m^2$, see bottom panel) and height $H = 256.7\,\text{m}$. The top and bottom boundaries of the can are located at $z_{\text{min}} = -117.2\,\text{m}$ (corresponding to the bedrock) and $z_{\text{max}} = 139.5\,\text{m}$, where $z = 0$ is at the detector center. The event vertices inside the can are uniformly distributed in volume. As can be seen from the middle and right panels, the generation volume is extended for muon events, with dependence on the energy, as some of the higher-energetic ν_μ and $\overline{\nu}_\mu$ CC events, when pointing towards to detector, may produce light reaching the instrumented volume even from far away. Those events kept and passed on to the `KM3Sim` stage.

5.2.2 Correspondance with usual response models

In the following, we break down the correspondance between the response matrix entries and physical quantities that are easier to interpret: efficiency and effective mass, PID probability, and PDF for reconstructed values. Note however that in the implementation of the fully correlated model, these derived quantities are not needed to apply the detector response: instead Eq. 5.7 is used directly. Such quantities will be used in the uncorrelated model described in Sec. 5.3.3.

Pseudo-efficiency and reconstruction PDF

We define the pseudo-efficiency for detection and classification of events as:

$$\mathcal{E}_{\text{class}}^{[\nu_x \to i]}(E, \theta, y) = \sum_{E', \theta', y'} R^{[\nu_x \to i]}(E, \theta, y, E', \theta', y') \tag{5.9}$$

This quantity describes the probability for a ν_x event occuring in the reference generation volume with truth (E, θ, y) to be both selected and classified into PID class i. Note that the normalisation of this quantity is somewhat arbitrary as it scales with $1/V_{\text{ref}}$.

We can now define the conditional probability for a selected and classified event with truth (E,θ,y) to be reconstructed as (E',θ',y'):

$$p_{\text{reco}}^{[\nu_x\to i]}(E,\theta,y,E',\theta',y') = \frac{R^{[\nu_x\to i]}(E,\theta,y,E',\theta',y')}{\mathcal{E}_{\text{class}}[\nu_x\to i](E,\theta,y)} \tag{5.10}$$

With this definition, the interaction channel ν_x, PID class i and true bin (E,θ,y) all being fixed, the function

$$(E',\theta',y') \to p_{\text{reco}}^{[\nu_x\to i]}(E,\theta,y,E',\theta',y') \tag{5.11}$$

is a proper PDF for the reconstructed quantities (E',θ',y'), since

$$\sum_{(E',\theta',y')} p_{\text{reco}}^{[\nu_x\to i]}(E,\theta,y,E',\theta',y') = 1 \tag{5.12}$$

Examples of such reconstruction PDFs obtained as projection of a typical response matrix are shown on Fig. 5.2. While the reconstructed values of correctly classified events (ν_e CC $\to$ shower, ν_μ CC $\to$ track) are well sampled and yield $\sim$ smooth empirical distributions, the misclassified channels (ν_e CC $\to$ track, ν_μ CC $\to$ shower) suffer from limited statistics, especially at high energy where the PID capabilities are good.

Note that the underflow and overflow bins have to be included in the sum on (E',θ',y') for correct normalisation. In other words, the leakage of events outside the analysis range (e.g. events reconstructed with $E' < 1\,\text{GeV}$, or upgoing events misreconstructed as downgoing) is taken into account by design. Conversely, the contamination of the analysis range with truly outside events has to be taken into account, by defining e.g. a larger range for the true energy E than for the reconstructed energy E'.

Overall detection efficiency and effective mass

The overall detection pseudo-efficiency is:

$$\mathcal{E}_{\text{det}}^{[\nu_x]}(E,\theta,y) = \sum_i \mathcal{E}_{\text{class}}^{[\nu_x\to i]}(E,\theta,y) \tag{5.13}$$

It describes the probability for a ν_x event occuring in the reference generation volume with truth (E,θ,y) to be selected, regardless of its PID classification. Like $\mathcal{E}_{\text{class}}$, $\mathcal{E}_{\text{det}}$ scales with $1/V_{\text{ref}}$, therefore it is not a properly normalised efficiency in the sense that $\mathcal{E}_{\text{class}} < 1$ even with perfect detector capabilities. The relevant quantity is the effective detection volume:

$$\mathcal{V}_{\text{eff}}^{[\nu_x]}(E,\theta,y) = V_{\text{ref}} \times \mathcal{E}_{\text{det}}^{[\nu_x]}(E,\theta,y), \tag{5.14}$$

or equivalently the effective detection mass or "effective mass":

$$\mathcal{M}_{\text{eff}}^{[\nu_x]}(E,\theta,y) = \rho_{\text{water}} \times V_{\text{ref}} \times \mathcal{E}_{\text{det}}^{[\nu_x]}(E,\theta,y). \tag{5.15}$$

Note that the definition 5.15 is strictly equivalent to

$$\mathcal{M}_{\text{eff}}^{[\nu_x]}(E,\theta,y) = \rho_{\text{water}} \times V_{\text{ref}} \times \frac{N_{\text{sel}}^{\text{MC}}[\nu_x\to i](E,\theta,y)}{N_{\text{gen}}^{\text{MC}}[\nu_x](E,\theta,y)}. \tag{5.16}$$

For a given data-taking time, the total number of ν_x selected events will scale with $\mathcal{M}_{\text{eff}}^{[\nu_x]}$, whatever the choice of V_{ref}.

Fig. 5.3 shows example of two-dimensional representations of the effective mass $\mathcal{M}_{\text{eff}}^{[\nu_x]}(E,\theta)$ for two typical channels. In addition, Fig. 5.4 shows the one-dimensional $\mathcal{M}_{\text{eff}}(E)$ for all interac-

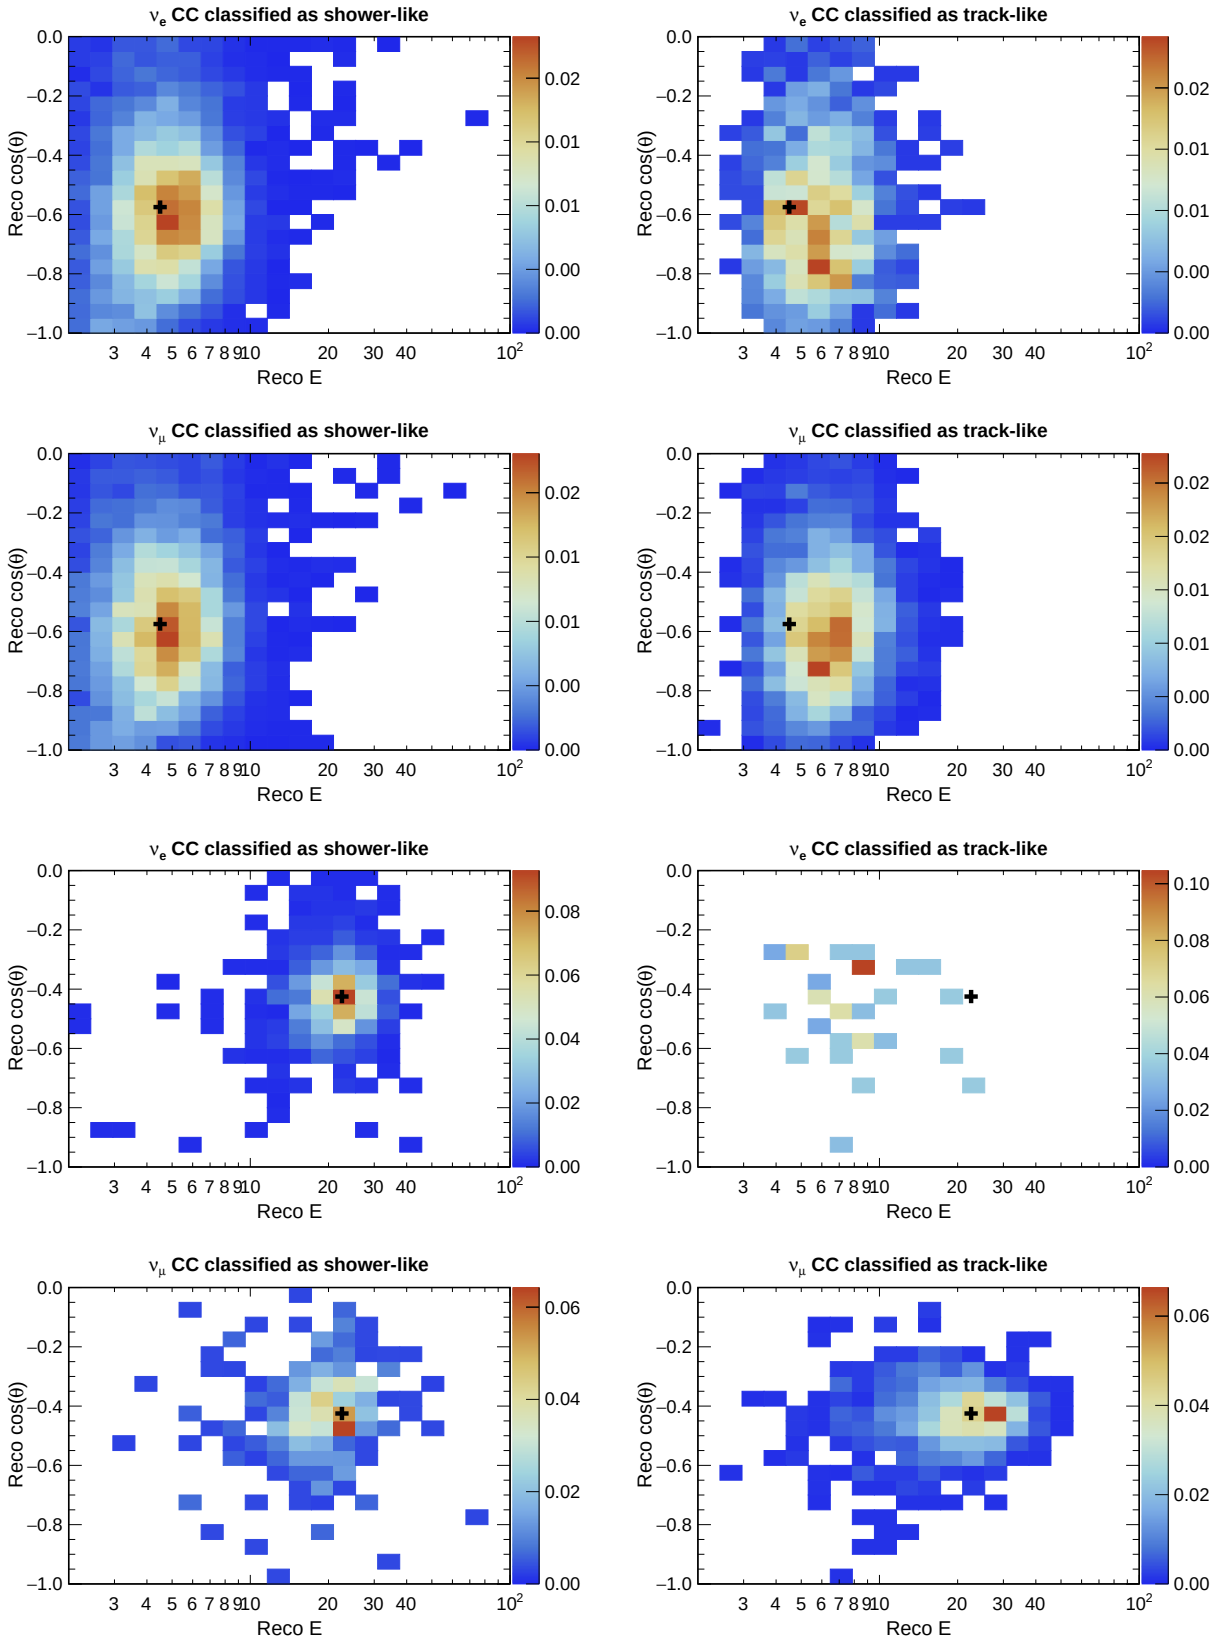

Figure 5.2: Examples of correlated response PDFs obtained from the response matrix. The standard PID approach (two-class) is used and the reconstruction PDFs are shown for both correctly and incorrectly classified events. The plotted quantity is $p_{\text{reco}}^{[\nu_x \to i]}(E, \theta, y, E', \theta', y')$ as defined in Eq. 5.10, with fixed selected values of (E, θ) indicated by the black crosses, averaged over y and marginalised over y'.

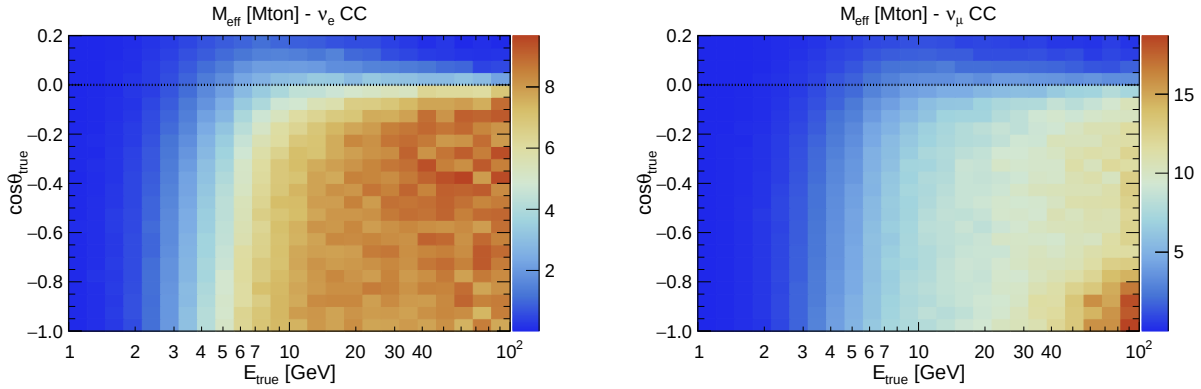

Figure 5.3: Effective mass of the detector as a function of true neutrino energy and zenith angle, for ν_e CC event (left) and ν_μ CC event (right). Mind the different color scales in left and right plots.

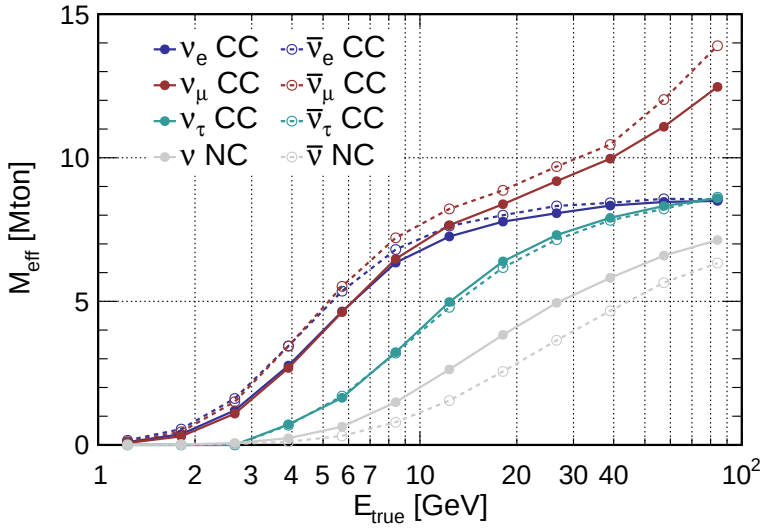

Figure 5.4: Effective mass of the detector as a function of true neutrino energy for all considered interaction channels.

tion channels. Note that these representations are not obtained via a direct unweighted average of $\mathcal{M}_{\text{eff}}^{[\nu_x]}(E,\theta,y)$ with respect to y, as this would fail to capture the effect of the distribution of the inelasticity y, which differs among interaction channels. Instead a ratio similar to Eq. 5.16 is used, where both $N_{\text{sel}}^{\text{MC}}$ and $N_{\text{gen}}^{\text{MC}}$ are summed over the Bjorken-y bins. The effect of the y distribution is better rendered this way, since the (`GENIE`-based) Monte Carlo generation does follow realistic y distributions.

The effective mass reaches a plateau around the intrumented mass in the electron CC, tau CC and NC channels; in the muon CC channel it keeps increasing up to values much larger than the instrumented mass, especially for events having an upgoing trajectory close to vertical ($\cos\theta_z \approx -1$). This is due to the containment and directional cuts described in Sec. 5.1.3. Such cuts are designed to reject the mostly downgoing atmospheric muons. The muons produced by the CC interactions of highly energetic, close to vertically-upgoing ν_μ and $\overline{\nu}_\mu$ can reach the detector and be reconstructed even if the interaction vertex is far below the instrumented volume – this is not the case for other interaction channels.

For the ν_τ and $\overline{\nu}_\tau$ CC channels, the shift of the effective mass 'turn-on' to higher energies is due to the $\tau^\pm$ leptons decaying into hadrons and thus producing less bright events most of the time (with a branching ratio of 65%); moreover when a $\tau^\pm$ decays into leptons part of its

energy is converted into neutrinos (see Sec. 2.1.2). A similar effect occurs for NC interactions, where on average only about half of the neutrino energy is converted into charged particles.

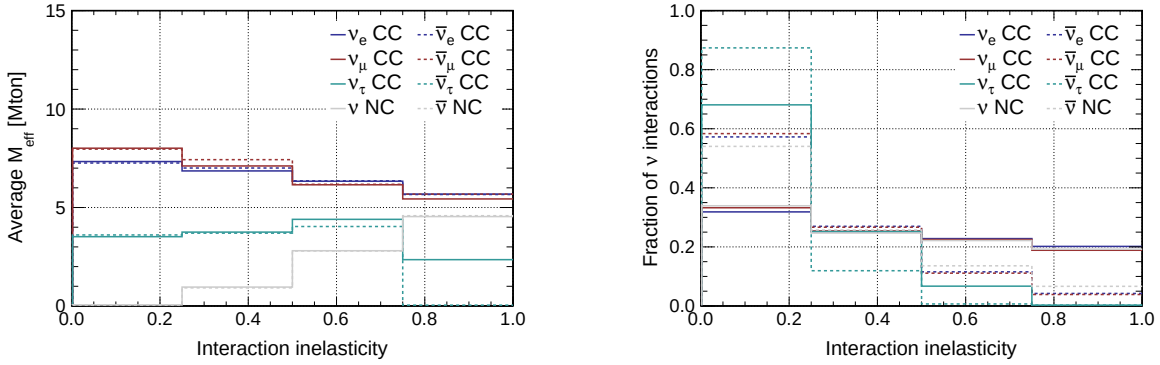

Figure 5.5: Left: detector effective mass $\mathcal{M}_{\text{eff}}(y)$ as a function of the inelasticity (Bjorken-y) of neutrino interactions, averaged over the whole range of upgoing trajectories and a single energy bin $8.25 < E < 10\,\text{GeV}$. Right: distributions of true inelasticity in the same energy bin.

Finally, the difference in detection efficiency between the neutrino and antineutrino channels of a given flavour arises from the different inelasticity distributions, as demonstrated in Fig. 5.5. The effective mass $\mathcal{M}_{\text{eff}}(y)$ (left plot) is essentially the same for neutrino and antineutrino channels, but its dependence on y is non-negligible. Highly inelastic (large y) CC interactions of $\nu_e/\overline{\nu}_e$ and $\nu_\mu/\overline{\nu}_\mu$ produce less visible light, since a larger part of the neutrino energy is converted into hadrons rather than an electromagnetic shower or a muon. Considering that the interactions of antineutrinos tend to be less inelastic (right plot), one retrieves that $\mathcal{M}_{\text{eff}}(\overline{\nu}_e\,\text{CC}) > \mathcal{M}_{\text{eff}}(\nu_e\,\text{CC})$, $\mathcal{M}_{\text{eff}}(\overline{\nu}_\mu\,\text{CC}) > \mathcal{M}_{\text{eff}}(\nu_\mu\,\text{CC})$ as observed on Fig. 5.4. For NC interactions the effect is opposite, the outgoing leptons being neutrinos: thus $\mathcal{M}_{\text{eff}}(\overline{\nu}) < \mathcal{M}_{\text{eff}}(\nu)$. For $\nu_\tau/\overline{\nu}_\tau$ CC interactions the dependence on y is less significant since most of the produced energy is hadronic regardless[2], and as a result $\mathcal{M}_{\text{eff}}(\overline{\nu}_\tau\,\text{CC}) \approx \mathcal{M}_{\text{eff}}(\nu_\tau\,\text{CC})$.

PID probability

Lastly we define the PID probability:

$$P_{\text{ID}}^{[\nu_x \to i]}(E,\theta,y) = \frac{\mathcal{E}_{\text{class}}^{[\nu_x \to i]}(E,\theta,y)}{\mathcal{E}_{\text{det}}^{[\nu_x]}(E,\theta,y)} \tag{5.17}$$

which is also properly normalised:

$$\sum_i P_{\text{ID}}^{[\nu_x \to i]}(E,\theta,y) = 1 \tag{5.18}$$

One-dimensional PID probabilities $P_{\text{ID}}(E)$ were shown and discussed in Chap. 3 (see Fig. 3.12). They are observed to be fairly independent of zenith angle. As discussed in Chap. 3, the Bjorken-y dependence is non-negligible and causes the PID to behave differently in neutrino and antineutrino channels.

[2] Note that the last y bin (0.75–1) is barely populated for the ν_τ and $\overline{\nu}_\tau$ channels, therefore the values of $\mathcal{M}_{\text{eff}}(y)$ corresponding to this bin on the left plot of Fig. 5.5 should not be considered accurate.

Summary

Finally, Eq. 5.7 can be rewritten equivalently as the successive application of the steps listed previously:

$$n_{\text{reco}}^{[i]}(E',\theta',y') = \sum_{\nu_x} \sum_{E,\theta,y} p_{\text{reco}}^{[\nu_x \to i]}(E,\theta,y,E',\theta',y') \times P_{\text{ID}}^{[\nu_x \to i]}(E,\theta,y) \times \mathcal{M}_{\text{eff}}^{[\nu_x]}(E,\theta,y) \times \frac{\mathrm{d}n_{\text{int}}^{[\nu_x]}}{\mathrm{d}M}(E,\theta,y) \tag{5.19}$$

5.2.3 Background contamination

As described in Chap. 3, two sources of non-neutrino background events are simulated: atmospheric muons and pure noise. Their impact on the analysis is twofold:

(a) background rejection cuts also filter out neutrino events, incidentally reducing the statistics of signal events;
(b) the background events remaining after the cuts contaminate the neutrino event samples and their impact must be accounted for in the oscillation analysis.

The two effects are fully interdependent, and the choice of selection cuts is driven by a trade-off between signal efficiency (a) and purity (b).

Regarding practical modeling and implementation, the first effect (a) is naturally accounted for in the application of selection cuts to the neutrino event samples. However, the background contamination (b) needs to be treated separately from neutrino events, because i) the calculation of the background event rates does not involve the same ingredients and model dependencies as for the neutrino rates, and ii) the available statistics of MC background events surviving the selection cuts is comparatively very small.

While the `gSeaGen` neutrino events follow a user-defined generation spectrum tuned to the requirements of the analysis (as described in Sec. 3.2.4), `MUPAGE` directly generates atmospheric muons, given a predefined detector livetime, following parametrisations of the expected flux. The obtained sample of events, after being processed by the reconstruction and selection chain, can thus be injected into the analysis without reweighting except for a livetime rescaling. Pure noise MC events are obtained by letting the trigger algorithms run on simulated optical background as in continuous data-taking conditions, and attempting to reconstruct the triggered events. In both cases the available MC statistics are quantified in equivalent detector livetime.

For atmospheric muons which are typically very bright, the limitation is the computational cost of the Cherenkov light simulation stage (`KM3`). During the production of the 'ORCA2016' MC sample I was in charge of running these simulations (`MUPAGE` + `KM3`) at the IN2P3 computing centre [261]. While this stage represented about 45 CPU-years of computation time and was the most time-consuming of the whole chain presented in Sec. 3.2, the total simulated livetime only amounts to $T_{\text{bkg},\mu} = 10.83\,\text{days}$, orders of magnitude below the equivalent neutrino livetime.

Pure noise events are also triggered with a fairly high rate ($\sim 10^3$ events/day) but very rarely pass the whole selection chain; nevertheless all triggered events need to be processed by the reconstructions, which limits the livetime that can be realistically simulated. Based on previous studies, it was anticipated that the selection chain would reduce the contamination from this background source to a close to negligible level. Therefore the simulated livetime was scaled down to about one day ($T_{\text{bkg,noise}} = 26.3\,\text{hours}$) so as to reduce the amount of computing resources spent.

The combination of pre-PID cuts and of the performant signal/background classifiers described in Sec. 3.3.2 is very successful in reducing the contamination to the percent level. When

applying, for instance, the following post-PID selection cuts:

$$p_{\text{bkg,noise}} < 0.10 \tag{5.20}$$

$$p_{\text{bkg},\mu} < 0.05 \tag{5.21}$$

a total of 8 and 11 simulated atmospheric muon events fail to be rejected and are tagged as track-like and shower-like, respectively. No pure noise events pass the rejection cuts. The loss of signal events due to these cuts is smaller than 5% over the whole analysis range and less than 1% above 5 GeV.

The contamination of atmospheric muon events into the signal samples, while not completely negligible, cannot be modeled from so few events. Therefore, their expected (E', θ', y') distribution is estimated by an extrapolation method. The PID rejection cuts are loosened up to a point where the statistics of surviving events becomes sufficient to infer a distribution (in a coarser binning). The obtained distribution is then rescaled down to the normalisation obtained with the tight cut: in this case, 8 and 11 events per 10.83 days of livetime.

5.3 Impact of Monte Carlo statistical uncertainties and solutions

5.3.1 The sparse Monte Carlo effect

Due to the limited MC statistics available in some channels, the response matrix entries can show bin-to-bin fluctuations which do not stem from any physical effect but are instead purely due to statistical fluctuations in the Monte Carlo sampling. This effect is for example visible on Fig. 5.2 in the tails of the distributions and in the misclassified channels (especially for ~ 20 GeV ν_e CC misclassified as tracks). As demonstrated in Chap. 8 (Sec. 8.1.3) and in Appendix B, using such sparsely sampled response matrices can result in drastic overestimations of sensitivity, an effect we will refer to as the 'sparse Monte Carlo effect'.

The motivation and conclusions of the mathematical study developed in Appendix B are summarised here. Considering as an estimator of the NMH sensitivity the Asimov $\Delta\chi^2$ function (introduced in details in Chap. 6), and denoting as $\mathcal{N}$ the size of the whole available sample of MC events, the dependence of the $\Delta\chi^2$ estimator on the amount $\mathcal{N}' < \mathcal{N}$ of MC events included in the construction of the response matrix is studied. The following observations are made:

- $\Delta\chi^2(\mathcal{N}')$ fluctuates when repeatedly drawing random subsets of size $\mathcal{N}'$ from the whole sample;
- the variance of $\Delta\chi^2(\mathcal{N}')$ increases as $\mathcal{N}'$ decreases;
- the average $\langle\Delta\chi^2\rangle(\mathcal{N}')$ is fairly well fitted by a power law:

$$\Delta\chi^2(\mathcal{N}') = p_0 + p_1 \times (\mathcal{N}'/\mathcal{N}_0)^{p_2} \tag{5.22}$$

  with p_0 the hypothetical 'infinite MC' limit and $p_2 \simeq -1$.

An asymptotic expansion of the $\Delta\chi^2$ function in the limit of small fluctuations of the response matrix entries indeed yields

$$\langle\Delta\chi^2\rangle(\mathcal{N}) = \Delta\chi^2_\infty + \frac{K}{\mathcal{N}} + \mathcal{O}\Big(\frac{1}{\mathcal{N}^{3/2}}\Big) \quad \text{with } K > 0. \tag{5.23}$$

As a heuristical interpretation of this effect, it can be pointed out that a very sparsely sampled response function such as the ones represented on Fig. 5.2 for the channel $\{\nu_e$ CC $\rightarrow$ track$\}$, amounts to an overestimation of the performance of the detector response. Indeed, this function

will send the oscillation signal contained in a bin of the 'true space' towards a few bins of the 'reconstructed space'. In the limit of small $\mathcal{N}$ where the response PDF is sampled with a single event, the signal is even imaged onto a single reconstructed bin instead of being smeared over a region of the reconstructed space. The response matrix method can be adapted in various ways in order to mitigate the sparse MC effect.

5.3.2 Adaptations of the correlated response matrix method

Adaptive response matrix binning and oversampling of the oscillated rates

Since we observe that the impact of the sparse MC effect increases with the binning density, an easy direction of improvement is to make sure the bin widths in both the true (E,θ,y) and reconstructed (E',θ',y') dimensions of the response matrix are the largest possible values, while still allowing to correctly capture the physical features of the detector response and avoiding the loss of real underlying information in the final event distributions.

In the early implementations of the analysis framework, the bins used for the true variables E and θ in the response matrix E were forced to be identical to the bins used for oscillated event rates n_{int}; moreover the oscillation probabilities were only calculated at the center of each of those bins. With such requirements, the number of bins for the true dimensions (E,θ) of the response matrix must be kept sufficient to sample the oscillation probabilities finely enough. On the other hand, projections of R onto the truth dimensions (*e.g.* efficiency functions $\mathcal{E}_{\text{class}}$ and $\mathcal{E}_{\text{det}}$) show rather extreme statistical fluctuations already at the typical binning used in oscillograms (50×50).

A solution which allows some averaging of R with respect to the true dimensions is to *oversample* the computation of the oscillated event rates. Instead of being computed at the bin center only, each bin b is subdivided into $N_{ov}\times N_{ov}$ sub-bins $b_1,...,b_{N_{ov}^2}$, where N_{ov} is the oversampling factor (per dimension), and the total oscillated event rate integrated over the bin b is computed as a sum over the sub-bins:

$$\frac{\mathrm{d}n_{\text{int}}^{[\nu_x]}}{\mathrm{d}M\mathrm{d}t}(E,\theta,y)=\sum_{i=1}^{N_{ov}}\sum_{j=1}^{N_{ov}}\frac{\mathrm{d}n_{\text{int}}^{[\nu_x]}}{\mathrm{d}M\mathrm{d}t}(E_i,\theta_j,y), \tag{5.24}$$

where $\frac{\mathrm{d}n_{\text{int}}^{[\nu_x]}}{\mathrm{d}M\mathrm{d}t}(E_i,\theta_j,y)$ is calculated following Eq. 4.1. Thus the oscillated event rates calculation can be made arbitrarily precise independently of the response matrix binnings. With a minimal oversampling factor of 2 in both energy and zenith angle, typical bin settings for the response matrix are as follows:

- 25 logarithmic bins in true E (1–100 GeV)
- 20 logarithmic bins in reconstructed energy E' (2–80 GeV)
- 20 bins in true zenith angle $\cos\theta_z$ for upgoing neutrinos (-1–0), plus 5 bins for downgoing neutrinos (0–1)
- 20 bins in reconstructed zenith angle $\cos\theta_z$ for upgoing neutrinos (-1–0)
- 4 bins in true Bjorken-y
- 1 to 4 bins in reconstructed Bjorken-y

The energy and zenith angle ranges in the true dimensions (E,θ) match those of the oscillated rate histograms. Likewise, the binnings and ranges in the reconstructed dimensions (E',θ',y') match those of the final event histograms which are used as input of the statistical analysis. The inclusion of a larger range of energy (1–2 GeV, 80–100 GeV) and zenith angle (downgoing neutrinos) for interacting events accounts for the contamination of events from these regions mis-reconstructed into the signal region. The use of identical ranges for true and reconstructed

dimensions would lead to an underestimation of the reconstructed event rate in the signal region, since the leakage of true events into the underflow and overflow bins would be accounted for but not the reverse migration.

At present, neutrinos reconstructed as downgoing are not used in the analysis, though the framework has been designed with the intention to include them later on. The potential benefits of measuring downgoing neutrinos for constraining systematics are discussed briefly in Chap. 7; currently the event selection (optimised to reject the atmospheric muon background) is set to discard events that appear as downgoing.

Parametrised detector response

To circumvent problems related to Monte Carlo statistics, an obvious possibility is to use a model of detector response based on parametrised analytical functions, fitted to the efficiency and response histograms. This has been developped in the KM3NeT collaboration in parallel to the method presented in this thesis. Independently, simpler models using published parametrisations have been employed in several works from outside the collaboration investigating ORCA's sensitivity to the mass hierarchy [161, 222].

Even in the advanced parametrised model used inside the collaboration, in order for the complexity of the model to remain manageable some simplifying assumptions are made: gaussian-like shape of the response functions, absence of bias or correlations between reconstructed variables, absence of dependance of the response functions on some true parameters, *etc.*

For ORCA's sensitivity study it has been observed that relevant reconstruction effects were missing from the parametrised model and that the simplifying assumptions led to overoptimistic sensitivities. Such observations confirm the relevance of the full MC approach even for a prospective sensitivity study.

Non-parametric smoothing

Another approach to remove Monte Carlo fluctuations is to apply smoothing algorithms to the sample response obtained from MC events. In contrast with the use of parametrised functions, these methods do not require any assumption on the functional form of the underlying response functions: as such they constitute *non-parametric* inference techniques.

One of the simplest smoothing techniques for a set of data points is to replace each point by an average (or weighted average) of the neighbouring points. If an unweighted average is used and the smoothing is applied to a histogram, this can be emulated rather closely by simply rebinning the histogram (*e.g.* going from $M \times N$ bins to $\frac{M}{2} \times \frac{N}{2}$), then optionnally going back to the original binning with interpolation. The smoothing can be also be applied directly to the data points without binning, and a variety of weight functions can be used to form a weighted average. This is referred to as *kernel smoothing*, which generally consists in a convolution of the empirical data distribution with smooth analytical kernels – *e.g.* gaussian. In the context of statistical inference, this is also called *kernel density estimation* (KDE).

A discussion of the application of KDE techniques for atmospheric neutrino oscillation analyses with large-volume Cherenkov detectors can be found in [254]. Apart from very simple methods based on rebinning and interpolation, no such approach was tried in this work; however they appear as very promising, especially given the low minimal number of MC events they require to obtain a reliable sensitivity estimate [254]. They are therefore identified as one of the main avenues for continuation of this work.

5.3.3 Uncorrelated smearing

The problems related to insufficient Monte Carlo statistics are inherent to the use of *event-by-event* smearing methods. As explained in Sec. 5.2, the correlated smearing matrix technique is equivalent to event-by-event reweighting. Another classical procedure is the use of sets of *smearing histograms*, which typically separate (decorrelate) the response of the detector for each measured variable. In particular, this was the approach adopted in the KM3NeT Letter of Intent [21]: the energy smearing and zenith angle smearing were performed separately, using a 3-dimensional smearing matrix (E, θ, θ') for the zenith angle smearing, and a 2-dimensional smearing matrix (E, E') for the energy smearing. The Bjorken-y variable was not included at all, and the choice was made to ignore the dependance of the energy response on θ after observing that it was a very marginal dependance.

In this work, as an intermediate step between the fully correlated response and the minimal smearing model used in [21] we opt for an uncorrelated smearing model where correlations between reconstructed variables are removed, but the dependance of the smearing functions on all true variables are kept. All the response functions used in the uncorrelated model are obtained directly as extractions (projections) of the correlated response R.

1. As a first step the efficiency functions $\mathcal{E}_{\text{class}}^{[\nu_x \to i]}(E, \theta, y)$ are obtained, following Eq. 5.9.
2. From there the marginalised versions of the full 6-dimensional reconstruction probability $p_{\text{reco}}^{[\nu_x \to i]}(E, \theta, y, E', \theta', y')$ (Eq. 5.10) are obtained: they are denoted
   (a) $\zeta_E^{[\nu_x \to i]}(E, \theta, y, E')$ for the energy smearing

$$\zeta_E^{[\nu_x \to i]}(E, \theta, y, E') = \sum_{\theta'} \sum_{y'} p_{\text{reco}}^{[\nu_x \to i]}(E, \theta, y, E', \theta', y'), \tag{5.25}$$

   (b) $\zeta_\theta^{[\nu_x \to i]}(E, \theta, y, \theta')$ for the zenith angle smearing

$$\zeta_\theta^{[\nu_x \to i]}(E, \theta, y, \theta') = \sum_{E'} \sum_{y'} p_{\text{reco}}^{[\nu_x \to i]}(E, \theta, y, E', \theta', y'). \tag{5.26}$$

   (c) $\zeta_y^{[\nu_x \to i]}(E, \theta, y, y')$ for the Bjorken-y smearing

$$\zeta_y^{[\nu_x \to i]}(E, \theta, y, y') = \sum_{E'} \sum_{\theta'} p_{\text{reco}}^{[\nu_x \to i]}(E, \theta, y, E', \theta', y'). \tag{5.27}$$

In the following we omit the superscript $[\nu_x \to i]$ for lightening the equations. The smearing functions are nevertheless still individual functions for each pair of interaction channel ν_x *and* event class i. The uncorrelated equivalent of the 6-dimensional p_{reco}, denoted Z, is obtained by recombining the decorrelated response density functions:

$$Z(E, \theta, y, E', \theta', y') = \zeta_E(E, \theta, y, E') \cdot \zeta_\theta(E, \theta, y, \theta') \cdot \zeta_y(E, \theta, y, y') \tag{5.28}$$

The overall application of the detector response, which was given by Eq. 5.7 or equivalently Eq. 5.19 for the correlated model, can be written as:

$$n_{\text{reco}}^{[i]}(E', \theta', y') = \sum_{\nu_x} \sum_{E, \theta, y} Z^{[\nu_x \to i]}(E, \theta, y, E', \theta', y') \times \mathcal{E}_{\text{class}}^{[\nu_x \to i]}(E, \theta, y) \times M_{\text{ref}} \times \frac{\mathrm{d}n_{\text{int}}^{[\nu_x]}}{\mathrm{d}M \mathrm{d}t}(E, \theta, y) \tag{5.29}$$

Although $\mathcal{E}_{\text{class}} \times M_{\text{ref}}$ is directly used in the implementation, we recall that it can also be expressed as a product of the PID probability and effective mass:

$$\mathcal{E}_{\text{class}}^{[\nu_x \to i]}(E, \theta, y) \times M_{\text{ref}} = P_{\text{ID}}^{[\nu_x \to i]}(E, \theta, y) \times \mathcal{M}_{\text{eff}}^{[\nu_x]}(E, \theta, y) \tag{5.30}$$

5.3.4 Estimating Monte Carlo statistical uncertainties

The use of a statistical technique to correct the sensitivity overestimation related to the MC sparseness, while keeping the original full MC approach has been evaluated. This is done by accounting for the MC statistical uncertainties in the likelihood. Two procedures for estimating the size of these uncertainties in the final event distributions are described in this section, and the method to include them in the statistical analysis is detailed in Sec. 6.4.

The uncertainty due to Monte Carlo statistical fluctuations on the final event count in a given bin $n_{\text{reco}}^{[i]}(E',\theta',y')$ corresponds to the standard deviation (or variance) of this quantity when reproducing the whole chain of simulation, reconstruction and classification described in Chap. 3 – for the same number of simulated events and with the same software, but providing a different set of input seeds to the many random generators involved in the simulation step. Clearly, if drawing a large number of such independent repetitions of the chain was within our reach computationnally, the sparse MC effect would no longer be a concern. Be that as it may, estimating the underlying variance of a sample statistic (a given function of the data) from a finite dataset is a common statistician's problem, and a variety of methods exist to adress it. One of them is *bootstrapping*, a particular case of *resampling* technique.

Bootstrapping

Given a dataset $\{x_i\}$ with N entries (events) drawn from the random variable X, the general idea of bootstrapping is to estimate the underlying statistical properties of a statistic $q(X)$ by randomly drawing resampled datasets of the same size N, with replacement, from the full original dataset $\{x_i\}$. Note that in the resampled dataset a significant number of the original events will appear multiple times, since it has the same size as the original dataset. In our case:

- the entries in $\{x_i\}$ correspond to neutrino generated MC events with their true characteristics ν_x and (E,θ,y) and, *if selected*, their reconstructed values (E',θ',y') and classification i;
- N is the total number N_{gen} of available generated events ($N_{\text{gen}} \sim 3.12\cdot 10^8$);
- q is our prediction for $n_{\text{reco}}^{[i]}(E',\theta',y')$ the number of events expected in a given class i and bin (E',θ',y');
- the statistical property we want to estimate is the variance $\sigma^2(q)$.

The prediction q is made, *e.g.* in the correlated case, through Eq. 5.7. In this equation none of the terms are random except the response matrix R itself. The randomness of X enters in the calculation of the entries of R (Eq. 5.3) both through the numerator $N_{\text{sel}}^{\text{MC}}[\nu_x \to i](E,\theta,y,E',\theta',y')$ and the denominator $N_{\text{gen}}^{\text{MC}}[\nu_x](E,\theta,y)$. Therefore, the set of MC events is resampled starting at the level of generated events. In practice, each event in the generated set is assigned an entry number $1 \leq k \leq N_{\text{gen}}$, and, if the event was detected and selected, a pointer to the corresponding event in the final set. Then, before building the response matrix a vector of N_{gen} independent random entry numbers is drawn, and these entries are used to fill the sets of histograms $N_{\text{gen}}^{\text{MC}}$ and $N_{\text{sel}}^{\text{MC}}$ entering in 5.3. The rest of the computation of the event distributions is done as usual, fixing the input parameters (oscillation, NMH, systematics *etc*) to a set of reference values.

Fig. 5.6 shows two examples of such bootstrap distributions of the event counts in individual bins of the reconstructed space, after three years data-taking and assuming a normal mass hierarchy. The two-class analysis strategy is used (track-like and shower-like classes). The left panel shows the distribution of the event count for all events classified as shower-like in a bin of reconstructed $E_\nu \sim 4\,\text{GeV}$ and reconstructed $\cos\theta_z \sim -0.7$. The event count is visibly normal distributed and exhibits a relative standard deviation of 1%. In the right panel, a low populated channel has been chosen: ν_e CC events misclassified as track-like, in a higher-energy bin (reconstructed $E_\nu \sim 15\,\text{GeV}$). The event count is much lower, with only a few events

predicted in three-years. Nevertheless the distribution is also very well fitted by a gaussian, with a larger relative standard deviation of about 18%.

In Fig. 5.7, the evolution of the relative standard deviation as a function of the reconstructed energy is shown for both event classes. The values obtained using the bootstrap technique are represented in solid lines (full markers), while dotted lines (hollow markers) correspond to the values obtained from a classical combination of binomial-like errors, as described hereafter. An excellent agreement between the two methods is observed. All in all, the relative fluctuations are rather closely related to the statistics of raw MC events available in each energy range (see Sec. 3.2.4).

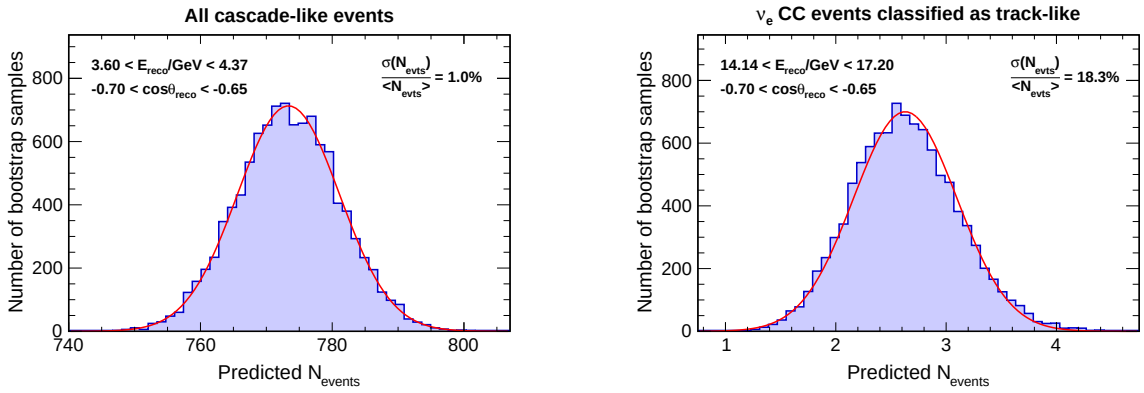

Figure 5.6: Examples of distributions of the predicted event count in a single bin, computed over an densemble of 10.397 bootstrap response matrices. The two-class analysis strategy is used (track-like and shower-like classes).

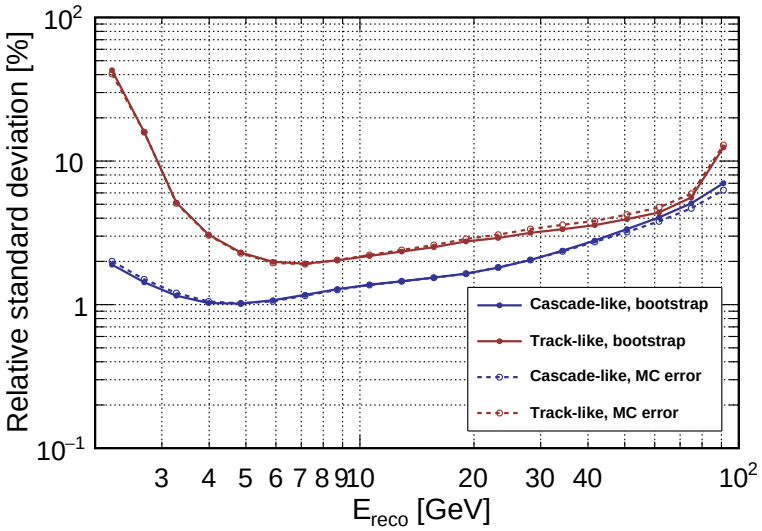

Figure 5.7: Relative standard deviation of the event count as a function of the reconstructed neutrino energy. In this example, the correlated smearing method was used. For the bootstrap technique (solid lines), this quantity is computed for each three-dimensional bin as the sample standard deviation divided by the sample mean of the bootstrap distribution of the event count, and is then averaged over the reconstructed zenith (a single bin of reconstructed inelasticity is used). The dependence on reconstructed zenith is observed to be negligble. For the curve labeled "MC error", the standard deviation is obtained as a weighted quadratic sum of binomial errors (see text), and divided by the expected event count in the bin. An excellent agreement between the two methods is observed.

Combination of binomial errors

The bootstrapping technique is computationnally expensive, as the whole procedure must be repeated at least $\mathcal{O}(10^4)$ times to obtain precise estimates of the fluctuations of the predicted event counts. The procedure thus cannot be applied on-the-fly during analysis.

In fact, the expected variance of the predicted event count can also be estimated via a probabilistic model applied to the calculation of the response matrix coefficients [262, 263]. To explain this method, simplified notations for the quantities entering in the efficiency ratio (5.3) are introduced:

$$R^{[\nu_x \to i]}(E,\theta,y,E',\theta',y') = \frac{N_{\text{sel}}^{\text{MC}}[\nu_x \to i](E,\theta,y,E',\theta',y')}{N_{\text{gen}}^{\text{MC}}[\nu_x](E,\theta,y)} \tag{5.31}$$

$$\hookrightarrow R = \frac{k}{N}, \tag{5.32}$$

Since $k \ll N$ in most practical cases, the calculation or R can be modeled to first order by a selection process applied to a fixed number N of events: each event may be either

- selected in the class i and bin (E',θ',y'),
- or not selected as such (either rejected or reconstructed in a different class or bin).

In this model k is then a binomially distributed random variable with N trials and binomial probability p equal to the (unknown) *true* underlying value of R, and its variance is given by

$$\sigma_k^2 = p\,(1-p)\,n. \tag{5.33}$$

As a way to estimate the variance of the ratio $R = k/N$, the unknown parameter p can be replaced by its estimate R [263], which ultimately yields

$$\frac{(\sigma_R)^2}{R^2} = \frac{1-R}{k}. \tag{5.34}$$

The error on $n_{\text{reco}}^i(E',\theta',y')$ is calculated as a weighted quadratic sum of the errors coming from all contributing bins $\{\nu_x,\ (E,\theta,y)\}$:

$$\sigma\left(n_{\text{reco}}^i(\boldsymbol{x}')\right)^2 = \sum_{\nu_x}\sum_{\boldsymbol{x}} \left(n_{\text{reco}}^{[\nu_x \to i]}(\boldsymbol{x},\boldsymbol{x}')\right)^2 \times \frac{\left(\sigma\left(R^{[\nu_x \to i]}(\boldsymbol{x},\boldsymbol{x}')\right)\right)^2}{\left(R^{[\nu_x \to i]}(\boldsymbol{x},\boldsymbol{x}')\right)^2} \tag{5.35}$$

$$= \sum_{\nu_x}\sum_{\boldsymbol{x}} \left(n_{\text{reco}}^{[\nu_x \to i]}(\boldsymbol{x},\boldsymbol{x}')\right)^2 \times \frac{1-R^{[\nu_x \to i]}(\boldsymbol{x},\boldsymbol{x}')}{N_{\text{sel}}^{\text{MC}}[\nu_x \to i](\boldsymbol{x},\boldsymbol{x}')} \tag{5.36}$$

with the notations $\boldsymbol{x} = (E,\theta,y)$ and $\boldsymbol{x}' = (E',\theta',y')$.

Modeling the calculation of R as a single binomial process, leading to Eq. 5.34, is only valid for the correlated reponse matrix approach. In the uncorrelated smearing case, R writes as

$$R^{[\nu_x \to i]}(\boldsymbol{x},\boldsymbol{x}') = \mathcal{E}_{\text{class}}^{[\nu_x \to i]}(\boldsymbol{x}) \times \zeta_E(\boldsymbol{x},E') \times \zeta_\theta(\boldsymbol{x},\theta') \times \zeta_y(\boldsymbol{x},y'). \tag{5.37}$$

A similar reasoning can be applied to each term of the product, leading to

$$\frac{(\sigma_R)^2}{R^2} = \frac{1}{N_{\text{sel}}^{\text{MC}}[\nu_x \to i](\boldsymbol{x},\boldsymbol{x}')} \times \left(1 - \mathcal{E}_{\text{class}}^{[\nu_x \to i]}(\boldsymbol{x}) + \frac{1-\zeta_E(\boldsymbol{x},E')}{\zeta_E(\boldsymbol{x},E')} + \frac{1-\zeta_\theta(\boldsymbol{x},\theta')}{\zeta_\theta(\boldsymbol{x},\theta')} + \frac{1-\zeta_y(\boldsymbol{x},y')}{\zeta_y(\boldsymbol{x},y')}\right). \tag{5.38}$$

The comparison of the relative errors obtained for the correlated and uncorrelated smearing methods is shown on Fig. 5.8. The relative fluctuations are about twice smaller for the uncor-

related smearing. This is a somewhat surprising result: a more important decrease would have been expected intuitively.

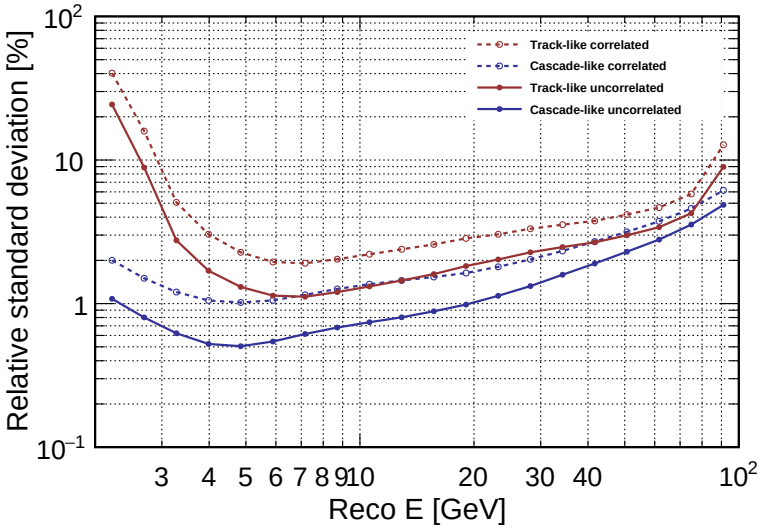

Figure 5.8: Comparison of the relative standard deviation of the event count per bin associated to MC fluctuations between the correlated and uncorrelated smearing methods. A fairly coarse reconstructed binning of $20 \times 20 \times 1$ bins has been used, as for the other figures of the section.

The method for including the estimated MC statistical fluctuations in the statistical analysis will be detailed in Chap. 6. The overall event distributions and sensitivity results obtained with the correlated and uncorrelated smearing methods will be further compared in Chap. 8.

Chapter 6

Statistical methods

Contents

Up to now we have described the modeling and calculation of the distributions of neutrino events expected to be observed once the experiment is performed. The sources of disagreement between real data and such predicted expectations may be separated into three categories:

(a) statistical fluctuations in data, inherent to the random nature of the measured processes;
(b) mismatch between the theoretical model and physical reality;
(c) and statistical fluctuations in the Monte Carlo simulations used to compute the predictions.

Estimating the statistical significance with which observed data can, for instance, reject a given neutrino mass hierarchy hypothesis requires understanding and modeling all three aspects. This chapter introduces the methods used in the estimation of the performance (statistical significance) expected to be achieved by the ORCA experiment for neutrino oscillation analyses, and in particular the NMH determination.

Sec. 6.1 introduces the basic tools of the statistical analysis, adressing items (a) and (b) of the above categorisation. The classical frequentist method for quantifying the sensitivity to the binary hypothesis testing problem of the NMH determination is then presented. Frequentist statistical inference considers the outcome of an experiment with respect to the ensembles of possible outcomes, and so does the NMH sensitivity study. The simplified but widely used 'Asimov dataset' approach uses instead the 'most representative' outcome to estimate the median sensitivity. This method and its comparison to the frequentist one are presented in Sec. 6.2. Methods for estimations of continuous parameters are discussed in Sec. 6.3. Finally, the methodology used to account for Monte Carlo statististical fluctuations (c) is adressed in Sec. 6.3.

6.1 Sensitivity to NMH hypothesis testing: general approach

6.1.1 Ratio of profiled likelihoods

In this work we treat the determination of the neutrino mass hierarchy as a classical binary hypothesis testing problem where both tested hypotheses are composite. We consider a number of *nuisance parameters*, which model sources of systematic uncertainty. We denote this ensemble of N_{syst} nuisance parameters as $\boldsymbol{\eta} = (\eta_0, \eta_1, ..., \eta_{N_{\text{syst}}-1})$.

The sensitivity estimation is based on the study of the distribution of a suitable test statistic over a large ensemble of simulated pseudo-experiments (PEs). This is in principle a purely frequentist approach. However, we do include prior knowledge of some nuisance parameters in the likelihood: our method is thus somewhat mixed between the frequentist and Bayesian paradigms.

We will denote the outcome of a PE as $\boldsymbol{d}$, with

$$\boldsymbol{d} = (\boldsymbol{n}^0, \boldsymbol{n}^1, ..., \boldsymbol{n}^{N_{\text{class}}-1}) \tag{6.1}$$

where $\boldsymbol{n}^i$ is the histogram of the reconstructed event distribution for PID class i. Each of these histograms has a total number of bins $N_{\text{bins}} = N_{\text{bins}\,E'} \times N_{\text{bins}\,\theta'} \times N_{\text{bins}\,y'}$ (the underflow and overflow bins being discarded), thus we write:

$$\boldsymbol{n}^i = (n^i_0, n^i_1, ..., n^i_{N_{\text{bins}}-1}) \in \mathbb{N}^{N_{\text{bins}}} \tag{6.2}$$

where n^i_b is the *measured* number of reconstructed events in bin b for PID class i.

Similarly we denote by $(\mu_{\text{NH}})^i_b = (\mu_{\text{NH}})^i_b(\boldsymbol{\eta})$ (resp. $(\mu_{\text{IH}})^i_b$) the *expected* number of reconstructed events in bin b for PID class i under the hypothesis that the normal (resp. inverted) hierarchy is true, the nuisance parameters being fixed to the set of values $\boldsymbol{\eta}$.

With these notations we define the test statistic T as:

$$T(\boldsymbol{d}) = \frac{\max_{\boldsymbol{\eta}} \left[\mathcal{L}_{\text{NH}}(\boldsymbol{d}\,|\,\boldsymbol{\eta})\right]}{\max_{\boldsymbol{\eta}} \left[\mathcal{L}_{\text{IH}}(\boldsymbol{d}\,|\,\boldsymbol{\eta})\right]} \tag{6.3}$$

where $\mathcal{L}_{\text{hyp}}(\boldsymbol{d}\,|\,\boldsymbol{\eta})$ (hyp = NH or hyp = IH) is the likelihood of measuring data $\boldsymbol{d}$, when the normal/inverted hierarchy is assumed to be true, and the nuisance parameters are given the values η. This test statistic (TS) is commonly referred to in the particle physics community as the *ratio of profiled likelihoods*[1] (RPL) [264].

We will first detail the basic statistical likelihood function $\mathcal{L}_{\text{stat}}$. We can optionnally add a term $\mathcal{L}_{\text{syst}}$ accounting for prior knowledge of the systematic parameters, in the fashion of a Bayesian prior likelihood. The total likelihood then writes:

$$\mathcal{L} = \mathcal{L}_{\text{stat}} \times \mathcal{L}_{\text{syst}} \tag{6.4}$$

[1] Note that this test statistic differs from the profile likelihood ratio TS which is more adapted to the test of nested parametric hypotheses where one hypothesis has a special status – typically, search for a signal of unknown strength compared to the background only hypothesis. Here, the two hierarchies are treated symmetrically in the definition of the TS.

Statistical likelihood

$\mathcal{L}_{\text{stat}}$ is a binned likelihood combining all PID classes i and assuming that in each bin the measured event number follows Poisson statistics:

$$\mathcal{L}_{\text{hyp}}(\boldsymbol{d}\,|\,\boldsymbol{\eta}) = \prod_{i=0}^{N_{\text{class}}-1} \prod_{b=0}^{N_{\text{bins}}-1} \exp\Big[-(\mu_{\text{hyp}})_b^i\Big] \frac{[(\mu_{\text{hyp}})_b^i]^{n_b^i}}{n_b^i!} \tag{6.5}$$

As an actual test statistic, instead of T a conventional choice is to use the so-called $\Delta\chi^2$ function:

$$\Delta\chi^2 = -2\ln T = \min_{\boldsymbol{\eta}} \Big[-2\ln\mathcal{L}_{\text{NH}}(\boldsymbol{d}\,|\,\boldsymbol{\eta})\Big] - \min_{\boldsymbol{\eta}} \Big[-2\ln\mathcal{L}_{\text{IH}}(\boldsymbol{d}\,|\,\boldsymbol{\eta})\Big] \tag{6.6}$$

From the point of view of numerical computation, the bin-by-bin likelihood is often an extremely small number and using a logarithm avoids floating point underflow problems. The quantity $\ln T$ is thus widely used and often referred to as 'log-likelihood ratio' (LLR). Multiplying by a factor -2 is a conventional choice which allows a straightforward conversion from $\Delta\chi^2$ values to confidence levels in cases such as parametric hypothesis testing and parameter estimation (where Wilk's theorem holds).

We can then write (with hyp = NH, IH):

$$-2\ln\mathcal{L}_{\text{hyp}} = \sum_{i=0}^{N_{\text{class}}-1} \sum_{b=0}^{N_{\text{bins}}-1} -2\Big[-(\mu_{\text{hyp}})_b^i + n_b^i \ln(\mu_{\text{hyp}})_b^i - \ln\Big(n_b^i!\Big)\Big] \tag{6.7}$$

and we further define[2]

$$\chi^2_{\text{hyp}}(\boldsymbol{d}\,|\,\boldsymbol{\eta}) = \sum_{i=0}^{N_{\text{class}}-1} \sum_{b=0}^{N_{\text{bins}}-1} 2\Big[(\mu_{\text{hyp}})_b^i - n_b^i + n_b^i \ln\frac{n_b^i}{(\mu_{\text{hyp}})_b^i}\Big]. \tag{6.8}$$

Note that χ^2_{hyp} and $-2\ln\mathcal{L}_{\text{hyp}}$ are equal up to a term that only depends on the data $\boldsymbol{d}$, and neither on the hypothesis (NH or IH) or on the nuisance parameters $\boldsymbol{\eta}$:

$$\chi^2_{\text{hyp}}(\boldsymbol{d}\,|\,\boldsymbol{\eta}) = -2\ln\mathcal{L}_{\text{hyp}} + 2\sum_{i,b}(n_b^i \ln n_b^i - 2\ln n_b^i!). \tag{6.9}$$

The term $2\sum_{i,b}(n_b^i \ln n_b^i - 2\ln n_b^i!)$ then vanishes in the likelihood ratio, such that our TS finally writes:

$$\Delta\chi^2 = \min_{\boldsymbol{\eta}} \Big[\chi^2_{\text{NH}}(\boldsymbol{d}\,|\,\boldsymbol{\eta})\Big] - \min_{\boldsymbol{\eta}} \Big[\chi^2_{\text{IH}}(\boldsymbol{d}\,|\,\boldsymbol{\eta})\Big] \tag{6.10}$$

Prior likelihood for nuisance parameters

Some of the nuisance parameters are better constrained by external experiments than by our own data. We take this prior knowledge into account by including such parameters as nuisance parameters in the analysis, along with constraint terms which are incorporated into the likelihood and called 'prior likelihood' terms, with

$$\mathcal{L}(\boldsymbol{d}\,|\,\boldsymbol{\eta}) = \mathcal{L}_{\text{stat}}(\boldsymbol{d}\,|\,\boldsymbol{\eta}) \times \mathcal{L}_{\text{syst}}(\boldsymbol{\eta}) \tag{6.11}$$

where $\mathcal{L}_{\text{stat}}(\boldsymbol{d}\,|\,\boldsymbol{\eta})$ is the likelihood defined in the previous paragraph.

The full set of nuisance parameters is denoted $\boldsymbol{\eta} = (\eta_p)_{p\in\mathcal{P}_{\text{syst}}}$ with $\mathcal{P}_{\text{syst}} = 1, ..., N_{\text{syst}}$. Let us denote by $\mathcal{P}_{\text{prior}} \subset \mathcal{P}_{\text{syst}}$ the set of indices for nuisance parameters constrained by a prior likelihood term. In the most simple approach, to each constrained parameter η_p ($p \in \mathcal{P}_{\text{prior}}$) is

[2] Using this definition of χ^2_{hyp} is again a conventional choice.

associated a gaussian PDF of mean H_p and standard deviation σ_p:

$$\mathcal{L}_{\text{syst}}(\boldsymbol{\eta}) = \prod_{p\in\mathcal{P}_{\text{prior}}} \frac{1}{\sqrt{2\pi\sigma_p^2}} e^{-\frac{-(\eta_p - H_p)^2}{2\sigma_p^2}} \tag{6.12}$$

Similarly to 6.8 we define

$$\chi^2_{\text{syst}}(\boldsymbol{\eta}) = \sum_{p\in\mathcal{P}_{\text{prior}}} \frac{(\eta_p - H_p)^2}{\sigma_p^2} = -2\ln\mathcal{L}_{\text{syst}}(\boldsymbol{\eta}) - \sum_{p\in\mathcal{P}_{\text{prior}}} \ln\left(2\pi\sigma^2\right) \tag{6.13}$$

With this definition, χ^2_{syst} and $-2\ln\mathcal{L}_{\text{syst}}$ again differ by a term that does not depend on $\boldsymbol{\eta}$ or on the hypothesis, and thus vanishes in the likelihood ratio. Therefore:

$$\Delta\chi^2 = \min_{\boldsymbol{\eta}} \left[\chi^2_{\text{NH,stat}}(\boldsymbol{d}\,|\,\boldsymbol{\eta}) + \chi^2_{\text{NH,syst}}(\boldsymbol{\eta})\right] - \min_{\boldsymbol{\eta}} \left[\chi^2_{\text{IH,stat}}(\boldsymbol{d}\,|\,\boldsymbol{\eta}) + \chi^2_{\text{IH,syst}}(\boldsymbol{\eta})\right] \tag{6.14}$$

The prior likelihood terms act as 'penalties' preventing the minimisation procedure to find minima with values of θ_p ($p \in \mathcal{P}_{\text{prior}}$) far from those provided by external experiments. As a side remark, note that in a purely Bayesian approach a choice of prior would be needed for *all* nuisance parameters for proper marginalisation. In our case no marginalisation is performed, therefore all priors are facultative.

Additional parameters (typically bin-by-bin normalisations) accounting for Monte Carlo statistical uncertainties can be added, formally acting in the same way as nuisance parameters constrained *e.g.* by gaussian priors. This is discussed in Sec. 6.4.

Compact notation

In the following we will use the index k as a compact index encompassing the indices i (PID classes) and the bins b, *i.e.*

$$\mu_k = \mu_b^i \tag{6.15}$$

$$n_k = n_b^i \tag{6.16}$$

with $0 \leq k \leq N_{\text{class}} \times N_{\text{bins}} - 1$. With this notation the total χ^2 function writes

$$\chi^2(\boldsymbol{d},\boldsymbol{\eta}) = \sum_k -2\cdot\left(n_k \ln(\mu_k(\boldsymbol{\eta})) - \mu_k(\boldsymbol{\eta})\right) + \sum_{p\in\mathcal{P}_{\text{prior}}} \frac{(\eta_p - H_p)^2}{\sigma_p^2} \tag{6.17}$$

6.1.2 Frequentist hypothesis testing with pseudo-experiments

In our main statistical approach, the evaluation of the sensitivity relies on the study of the probability density functions (PDFs) for the TS $t = \Delta\chi^2$ when NH is assumed to be true (PDF ϕ_{NH}), and when IH is assumed to be true (PDF ϕ_{IH}). The principle is illustrated on Fig. 6.1 with dummy PDFs based on analytical assumptions. In practice in the sensitivity study these PDFs are estimated by generating a large number of pseudo-experiments (PEs), which represent possible outcomes of the measurement.

Let us quickly describe the approach to the sensitivity evaluation using the standard hypothesis testing formulation. Recall that the test statistic t is defined as:

$$t = -2\ln T = \Delta\chi^2 = -2\cdot\ln\frac{\max\mathcal{L}_{\text{NH}}}{\max\mathcal{L}_{\text{IH}}} \tag{6.18}$$

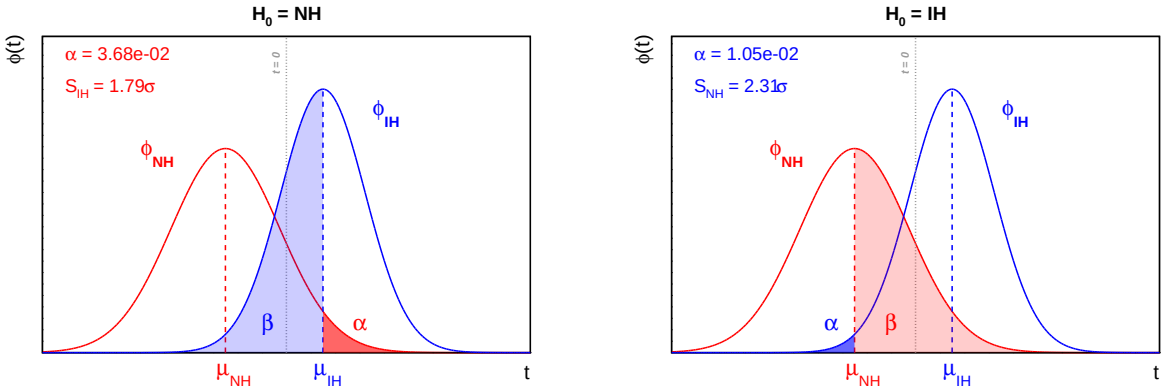

Figure 6.1: Example of probability density functions for the test statistic t under both hierarchy hypotheses, assuming normal distributions with $\mu_{\rm NH} = -5$, $\mu_{\rm IH} = 3$ and $\sigma = 2\sqrt{\mu}$. The areas corresponding to the probability α of erroneously rejecting H_0 (significance level) and probability β of erroneously accepting H_0 are represented for both $H_0 = \text{NH}$ (left) and $H_0 = \text{IH}$ (right). In each case the values of the significance level and corresponding sensitivity S in σ are reported.

Therefore, large (positive) values of t correspond to a higher likelihood for the IH hypothesis, while low (negative) values of t correspond to a higher likelihood for the NH hypothesis.

First one must decide which hypothesis is under test. The hypothesis under test is usually called the *null* hypothesis and denoted H_0, while the other hypothesis is the *alternative* hypothesis, denoted H_1. Here we will take the example of the NH hypothesis being under test (H_0). Assuming that the the PDFs of the TS $\phi_{\rm NH}(t)$ and $\phi_{\rm IH}(t)$ are known, we choose a cut value $t_{\rm cut}$ which defines our test. The decision to accept or reject the NH hypothesis depends on the TS value $t_{\rm obs}$ computed on the observed data:

- if $t_{\rm obs} > t_{\rm cut}$, the NH hypothesis is rejected
- if $t_{\rm obs} < t_{\rm cut}$, the NH hypothesis is accepted

Before doing the experiment, from the distributions $\phi_{\rm NH}$ and $\phi_{\rm IH}$ we can compute the probability of error of the first kind α:

$$\alpha = P(\text{reject NH} \,|\, \text{NH true}) = \int_{t_{\rm cut}}^{+\infty} \phi_{\rm NH} \tag{6.19}$$

and the probability of error of the second kind[3] β:

$$\beta = P(\text{accept NH} \,|\, \text{NH false}) = \int_{-\infty}^{t_{\rm cut}} \phi_{\rm IH} \tag{6.20}$$

The probability α is called the *significance level* for the test of NH. The probability to correctly accept the NH hypothesis is then $1-\alpha$, called the *confidence level* (CL), or *size* of the test. The probability $1-\beta$ to correctly reject NH is the *power* of the test. At this point let us emphasize that:

1. The definition of the confidence level $1-\alpha$ does not even require an alternative hypothesis to be formulated: to address the choice between two hypotheses the statistical power must be considered as well.
2. The two hypotheses are treated asymetrically: in most cases 'NH is accepted' and 'IH is rejected' are not equivalent statements.

[3] Errors of the first and second kind are sometimes referred to as 'false positive' and 'false negative', respectively – with the convention that the test is 'positive' when the null hypothesis is *rejected*.

Several approaches are possible to guide the choice of t_{cut}, with a trade-off between confidence level and statistical power. In the field of particle physics it is customary to report the *median sensitivity* of an experiment, which corresponds in our example to choosing t_{cut} as the median m_{IH} of the TS distribution under the alternative hypothesis IH. With the previous definitions this is the confidence level $\text{CL}_{H_0=\text{NH}}$ for accepting NH with a statistical power of 50%. We can rephrase this definition in two equivalent formulations:

- $\alpha_{H_0=\text{NH}} = 1 - \text{CL}_{H_0=\text{NH}}$ is the *p-value* for the rejection of NH when we obtain the 'average' experiment expected under the IH hypothesis,
- if NH were true, $\alpha_{H_0=\text{NH}}$ would be the probability to obtain a value t_{obs} equal or greater (*i.e.* 'worse') than its expectation under the IH hypothesis.

Fig. 6.1 shows a graphical illustration of this definition. Such a choice of t_{cut} has an obvious caveat: assuming that NH is false the test has an equal probability to accept or reject it; only the probability of *false rejection* is optimised. It is nevertheless a widely used convention for future experiments in neutrino physics to report median sensitivities for the NMH determination. Further discussion can be found in Sec. 6.1.3.

In the rest of the thesis, the term 'LLR sensitivity' is used to refer to the median sensitivity S to *reject the wrong hierarchy*. Importantly, the notation S_{NH} stands for 'true NH' *i.e.* $H_0 = \text{IH}$, and conversely S_{IH} corresponds to $H_0 = \text{NH}$. In addition the sensitivity S is expressed in number of σ rather than as a probability in $0 < p < 1$, following the conversion detailed below.

Sensitivity in number of σ

It is common practice in particle physics and other fields to convert a p-value or a level of significance expressed as a probability $0 < \alpha < 1$ into the equivalent number S of standard deviations σ. This conversion is defined such that α is the probability for a normally distributed random variable $X \sim \mathcal{N}(0, \sigma)$ to be drawn farther than $S\sigma$ from its mean. Specifically, two conventions exist:

- $\alpha = P(X > S\sigma)$ is the one-sided convention,
- $\alpha = P(|X| > S\sigma)$ is the two-sided convention.

With the one-sided convention, $S = 0\sigma$ corresponds to a level of significance of 50%. This choice seems logical in the case of binary hypothesis testing, as it amounts to an equal probability of accepting or rejecting H_0 when it is true – the test is meaningless. Therefore, in the following we adopt the one-sided convention[4], for which the explicit conversion reads:

$$\alpha = \frac{1}{\sqrt{2\pi\sigma^2}} \int_{S\sigma}^{+\infty} e^{-\frac{t^2}{2\sigma^2}} \mathrm{d}t = \frac{1}{\sqrt{\pi}} \int_{\frac{S}{\sqrt{2}}}^{\infty} e^{-t^2} \mathrm{d}t = \frac{1}{2} \operatorname{erfc}\left(\frac{S}{\sqrt{2}}\right), \tag{6.21}$$

i.e.

$$S = \sqrt{2} \operatorname{erfc}^{-1}(2\alpha), \tag{6.22}$$

where erfc is the complementary error function:

$$\operatorname{erfc}(x) = 1 - \operatorname{erf}(x) = \frac{2}{\sqrt{\pi}} \int_x^{\infty} e^{-t^2} \mathrm{d}t \tag{6.23}$$

and erfc^{-1} is its inverse. Tab. 6.1 lists a few values of S and the corresponding probabilites.

[4]Note that some authors (see for instance Ref. [159]) choose the two-sided convention also for hypothesis testing, for consistency with the conventions used in parameter estimation (standard error) and confidence intervals (coverage probability). In the two-sided convention 0, 1, 2, 3 and 5σ correspond to confidence levels (1-α) of 0, 68.3%, 95.5%, 99.7%, and $1 - 5.73 \cdot 10^{-7}$, respectively.

Sensitivity S	0σ	1σ	2σ	3σ	5σ
confidence level 1-α	50%	84.1%	97.7%	99.9%	99.99...%
significance α	50%	15.9%	2.28%	0.14%	$2.87 \cdot 10^{-7}$

Table 6.1: Correspondance between standard probabilities and number of σ with the one-sided convention.

Normally distributed test statistic

Consider the case, frequent in practice, where the test statistic t follows normal distributions both for true NH and true IH:

$$t_{\mathrm{NH}} \sim \mathcal{N}(\mu_{\mathrm{NH}}, \sigma_{\mathrm{NH}}) \tag{6.24}$$

and

$$t_{\mathrm{IH}} \sim \mathcal{N}(\mu_{\mathrm{IH}}, \sigma_{\mathrm{IH}}). \tag{6.25}$$

Let us assume that the true hierarchy is inverted (IH), so that we want to evaluate the sensitivity S_{IH} for the test of $H_0 = \mathrm{NH}$. The test is defined by

$$t_{\mathrm{cut}} = \mu_{\mathrm{IH}} \tag{6.26}$$

and therefore

$$\alpha = \int_{\mu_{\mathrm{IH}}}^{+\infty} \phi_{\mathrm{NH}}(t)\mathrm{d}t = \frac{1}{\sqrt{2\pi\sigma_{\mathrm{NH}}^2}} \int_{\mu_{\mathrm{IH}}}^{+\infty} e^{-\frac{(t-\mu_{\mathrm{NH}})^2}{2\sigma_{\mathrm{NH}}^2}} \mathrm{d}t = \frac{1}{\sqrt{\pi}} \int_{\frac{\mu_{\mathrm{IH}}-\mu_{\mathrm{NH}}}{\sqrt{2}\cdot\sigma_{\mathrm{NH}}}}^{+\infty} e^{-x^2} \mathrm{d}x \tag{6.27}$$

Comparing to Eq. 6.21, one can finally write S_{IH} under the simple form:

$$S_{\mathrm{IH}} = \frac{\mu_{\mathrm{IH}} - \mu_{\mathrm{NH}}}{\sigma_{\mathrm{NH}}} \tag{6.28}$$

and conversely for true NH:

$$S_{\mathrm{NH}} = \frac{\mu_{\mathrm{IH}} - \mu_{\mathrm{NH}}}{\sigma_{\mathrm{IH}}} \tag{6.29}$$

Provided that the assumption of gaussianity can be validated, the use of Eq. 6.28–6.28 for the sensitivity evaluation is very practical. Indeed, it only requires estimates of the mean and standard deviation of normally distributed variables, regardless of their actual relative positions (the scale of S). These estimates are very precise already with a few thousand PEs. In contrast, the direct sampling method (Eq. 6.27) consists in counting the fraction of PEs observing a value of t farther than the specified t_{cut} threshold, and therefore requires to sample the *tails* of the simulated distributions, all the more so when S is large. For high confidence levels ($S > 3\sigma$) several orders of magnitude more PEs would be needed, which would rapidly become infeasible in practice.

Expected range for the p-value of the test

Once the experiment is performed and a value t_{obs} of the TS is measured, one evaluates the p-value of the test, *i.e.* the probability of obtaining a more extreme result than the observed one assuming that H_0 is true. In the case $H_0 = \mathrm{NH}$ this is:

$$p = \int_{t_{\mathrm{obs}}}^{+\infty} \phi_{\mathrm{NH}}(t)\mathrm{d}t \tag{6.30}$$

Regardless of whether H_0 is rejected ($p < \alpha$) or not, the p-value stands as a result of the test. For the sensitivity study, evaluating an expected range for p thus gives more information on the

projected performance of the experiment than simply stating its median value (which equals α in our case). This expected range is simple to predict if the TS distributions are gaussian. For example, defining:

$$S^{-}_{\mathrm{IH}} = \frac{(\mu_{\mathrm{IH}} - \sigma_{\mathrm{IH}}) - \mu_{\mathrm{NH}}}{\sigma_{\mathrm{NH}}} = S_{\mathrm{IH}} - \frac{\sigma_{\mathrm{IH}}}{\sigma_{\mathrm{NH}}} \tag{6.31}$$

and

$$S^{+}_{\mathrm{IH}} = \frac{(\mu_{\mathrm{IH}} + \sigma_{\mathrm{IH}}) - \mu_{\mathrm{NH}}}{\sigma_{\mathrm{NH}}} = S_{\mathrm{IH}} + \frac{\sigma_{\mathrm{IH}}}{\sigma_{\mathrm{NH}}}, \tag{6.32}$$

we obtain an interval $I_S = [S^{-}_{\mathrm{IH}}, S^{+}_{\mathrm{IH}}]$ centered on the median sensitivity, such that the p-value (converted in number of σ) for exclusion of the normal hierarchy will be found in I_S for 68.3% of the possible outcomes of the experiment. As a side remark, other intervals with the same probabilistic property can be constructed, for example as a symmetrical interval in terms of raw p-values $I_p = [\alpha - \Delta p, \alpha + \Delta p]$. Such an interval would however be wider: I_S is actually the *shortest* possible interval with this property. This is an argument in favour of using the $n\sigma$ scale to report p-values or sensitivities.

Practical procedure for the sensitivity evaluation

1. A set of assumptions on the true parameter values ($\boldsymbol{\theta}_{\mathrm{true}}$) is fixed, including an assumption on the true NMH.
2. The expected events counts $\boldsymbol{\mu}(\boldsymbol{\theta}_{\mathrm{true}})$ are calculated following the methods described in Chap. 4 and 5.
3. A number N_{PEs} of PEs are generated based on $\boldsymbol{\mu}(\boldsymbol{\theta}_{\mathrm{true}})$. A single PE consists in a distribution $\boldsymbol{d}$ of randomly drawn numbers n_k of observed events per bin k:

   $$\text{for each bin } k,\ P(n_k = m) = \exp\{-\mu_k(\boldsymbol{\theta}_{\mathrm{true}})\}\frac{(\mu_k(\boldsymbol{\theta}_{\mathrm{true}}))^m}{m!} \tag{6.33}$$

   Distinct PEs differ from one another by the integer seed passed to the (pseudo-)random number generator. Seeds are used sequentially and the seed used for each PE is kept in its output, so that ensembles of pseudo-experiments are actually reproducible.
4. For each PE, the test statistic $t = \Delta\chi^2$ is computed by minimising the function $\chi^2(\boldsymbol{d}, \boldsymbol{\eta})$ with respect to the nuisance parameters $\boldsymbol{\eta}$, once under the NH hypothesis and once under the IH hypothesis.

After repeating the above steps for both true NMH hypotheses and fitting the obtained empirical TS distributions ϕ_{NH} and ϕ_{IH} with gaussian distributions, the median sensitivities S_{NH} and S_{IH} and the associated expected ranges of p-value can be calculated following Eqs. 6.28, 6.29, 6.31, and 6.32.

In general, the whole procedure is repeated to scan a range of assumed values $\boldsymbol{\theta}_{\mathrm{true}}$ for the parameters which unknown true values are expected to have a strong impact (typically θ_{23} and δ_{CP}).

6.1.3 Discussion

Caveats of the median sensitivity and alternatives

As stated previously, an important feature of the above definition of sensitivity is its asymmetric treatment of the null and alternative hypotheses: the null H_0 is given a special status. This is very natural in cases such as searches for an astrophysical signal or new physics, where H_0 is the background-only hypothesis and rejecting it is interpreted, depending on the significance, as an evidence or discovery of signal.

The definitions of the confidence level (Eq. 6.19) and power (Eq. 6.20) of the test are swapped if we change the reference hypothesis, yet we use the CL as a measure of the performance of the experiment while keeping the power fixed to 50%: therefore the test has to be applied separately to NH and IH. However, although the TS itself does not depend on the choice of H_0, applying the same test to both hypotheses will not result in accepting one and rejecting the other for all possible outcomes. In particular, for the median sensitivity approach where $t_{\text{cut}} = m_{\text{alt}}$ (median under the alternative hypothesis), all values t_{obs} such that $m_{\text{NH}} < t_{\text{obs}} < m_{\text{IH}}$ result in both hypotheses being accepted. As pointed out and extensively discussed in [159], an alternative, which ensures that the experiment will always accept one hierarchy and rule out the other, is to define t_{cut} such that

$$\int_{t_{\text{cut}}}^{+\infty} \phi_{\text{NH}}(t)\mathrm{d}t = \int_{-\infty}^{t_{\text{cut}}} \phi_{\text{IH}}(t)\mathrm{d}t = \alpha_{\text{cross}} \tag{6.34}$$

The corresponding $\text{CL} = 1 - \alpha_{\text{cross}}$ is called the crossing sensitivity.

Choice of alternative hypothesis

In principle, H_0 actually has to be compared with the ensemble of possible H_1. To be conservative, we need in principle to quote the minimal CLs obtained over the whole range of H_1 hypotheses (typically varying the true values of θ_{23} and δ_{CP}) [159]. Another possible approach is to compare H_0 with the most likely H_1. This was done for example in [21] with only the variation of true θ_{23} considered in H_1.

True parameter values in PEs

In the studies performed up to now, true values of all oscillation and systematic parameters have been kept fixed in the generation of a given ensemble of PEs, and never fluctuated. This corresponds to a purely frequentist approach: the true parameters are not assigned any stochastic character. This includes parameters such as the solar parameters Δm^2_{21} and θ_{12} to which the experiment is weakly sensitive and therefore are not fitted. In the Letter of Intent, the *model* (best-fit) values of these parameters are fluctuated among PEs. For small variations this approach is expected to have the same effect as fluctuating their true values. In the present analysis, these parameters are neither fluctuated nor fitted. It has been pointed out internally to the collaboration that not fluctuating true values can lead to incorrect sensitivity results, in particular for parameters to which the experiment's sensitivity is weak. Therefore it is planned to check what is the impact of fluctuating true values in our case.

Practical pseudo-experiments generation

Statistical analyses based on pseudo-experiments sometimes simulate neutrino events one-by-one, in particular when an unbinned likelihood is employed [155]. The question of the equivalence between event-by-event PEs and

The expected number of neutrino interactions $n_{\text{int}}^{[\nu_x]}(l)$ occuring in the detector in true bin $l = (E, \theta, y)$ in a given interval of time is Poisson distributed, as a consequence of the fact that the interaction rate is constant in time and space (neutrinos do not lose energy in the medium unless they interact). Assuming that the detector response for channel ν_x is modelled by a probability $R^{\nu_x}(l, b)$ for an interaction occurring in true bin l to be detected and reconstructed[5] in bin $b = (E', \theta', y')$, it follows that the histogram $(n_1^{\nu_x,l}, n_2^{\nu_x,l}, ...n_{N_{\text{bins}}}^{\nu_x,l})$ for reconstructed events originating from interaction channel ν_x and true bin l, *given that* $n_{\text{int}}^{\nu_x,j} = k$, follows a multinomial distribution with number of trials k and parameters $\left(R^{\nu_x}(l,1), R^{\nu_x}(l,2), ..., R^{\nu_x}(l, N_{\text{bins}})\right)$. Such a composition of a Poisson distribution and a multinomial distribution has the property that the $n_b^{\nu_x,l}$ are independent and Poisson-distributed with mean $\left(R^{\nu_x}(l,b) \times n_{\text{int}}^{[\nu_x]}(l)\right.$. The overall

reconstructed event count $n_b = \sum_{\nu_x,i} n_b^{\nu_x,l}$ is then a sum of Poisson distributed independent random variables, and therefore is Poisson distributed.

Hence, to simulate a PE for our binned measurement it is sufficient to compute the expectation for the number of events in each bin b of the signal histograms, and draw a random number of observed events following Poisson statistics. This approach is valid even when the event counts are small, for any number of contributing channels and for any shape of the response R.

6.2 The Asimov dataset approach

6.2.1 Technical limitations of the LLR approach

The method described in the last section has some caveats, most importantly that one needs a very large number of PEs to model the test statistic PDFs. Consider the following practical issues:

- the χ^2 minimisation has an order of 10 free parameters, requiring a few hundreds to a few thousands iterations to find the best-fit,
- for some parameters multiple likelihood maxima are present, meaning that the minimisation has to be repeated with various fit starting values to ascertain that the global maximum is found,
- while the simulation of one model experiment takes an order of 1 s in the worst case

All in all, the TS calculation for a single experiment can require several CPU-hours. This is especially a problem if high confidence levels are to be computed directly from the empirical TS distributions, since it requires the tails of these distributions to be sufficiently sampled. Generating an order of 10^5 to 10^8 PEs can be needed, depending on the achieved CL. In practice, with assumptions on the shape of the TS distributions this requirement can be mitigated by fitting these PDFs with parametrised functions, but several thousands PEs are still needed for a reliable fit.

6.2.2 Asimov sensitivity

A less CPU-intensive alternative is the use of the so-called Asimov dataset approach. This consists in evaluating the sensitivity from the 'most representative' outcome (the average experiment[6]), *i.e.* replace the fluctuated pseudo-data $\boldsymbol{d}$ in Eq. 6.10 by expectation values $\boldsymbol{\mu}$. The $\Delta\chi^2$ test statistic defined in Eq. 6.18, instead of a random variable, becomes a function of the models. It is customary to denote it $\overline{\Delta\chi^2}$. Assuming a given true hierarchy, the χ^2 difference simplifies to a single term:

$$\overline{\Delta\chi^2}_{\mathrm{NH}} = \min_{\boldsymbol{\eta}} \left[\chi^2_{\mathrm{NH}}(\boldsymbol{\mu}_{\mathrm{NH}}\,|\,\boldsymbol{\eta})\right] - \min_{\boldsymbol{\eta}} \left[\chi^2_{\mathrm{IH}}(\boldsymbol{\mu}_{\mathrm{NH}}\,|\,\boldsymbol{\eta})\right] = -\min_{\boldsymbol{\eta}} \left[\chi^2_{\mathrm{IH}}(\boldsymbol{\mu}_{\mathrm{NH}}\,|\,\boldsymbol{\eta})\right] \tag{6.35}$$

and conversely

$$\overline{\Delta\chi^2}_{\mathrm{NH}} = \min_{\boldsymbol{\eta}} \left[\chi^2_{\mathrm{NH}}(\boldsymbol{\mu}_{\mathrm{NH}}\,|\,\boldsymbol{\eta})\right] \tag{6.36}$$

The χ^2 terms corresponding to the fit under the correct hierarchy assumption vanish, since the fitted 'pseudo-data' (Asimov set) can be matched *exactly* with the best-fit parameters all set equal to their true values.

A widely used estimate of the significance to exclude the wrong hierarchy is given by

$$S_{\mathrm{Asimov}} = \sqrt{|\overline{\Delta\chi^2}|} \tag{6.37}$$

[5]For simplicity we consider a single PID class, or equivalently that b runs on both (E', θ', y') and PID classes.

Note that this is again a confidence level converted in number of sigmas with the one-sided convention. In the following we will refer to this estimate as the 'Asimov sensitivity'.

6.2.3 Agreement with the LLR sensitivity

Let us quickly examine some conditions for the Asimov sensitivity to equal the median sensitivity evaluated with the LLR method (Eq. 6.28–6.29). A simple sufficient condition is that distributions of the test statistic $t = \Delta\chi^2$ (Eq. 6.18) meet all the following criteria:

1. for each true hypothesis t follows a normal distribution with mean $t_0 = \overline{\Delta\chi^2}$,
2. the mean t_0 and standard deviation σ_t are related by $\sigma = 2\sqrt{|t_0|}$,
3. and the distributions under both hypotheses are symmetrical: $t_{0,\mathrm{NH}} = -t_{0,\mathrm{IH}}$.

Indeed, in such case we have

$$S_{\mathrm{NH}} = \frac{|\mu_{\mathrm{NH}} - \mu_{\mathrm{IH}}|}{\sigma_{\mathrm{IH}}} = \frac{2|t_0|}{2\sqrt{t_0}} = \sqrt{|\overline{\Delta\chi^2}|} = S_{\mathrm{IH}} \tag{6.38}$$

In addition we can relate the crossing sensitivity S_{cross} to the Asimov one: when the test statistic distribution are symmetrical, the condition 6.34 is met by choosing $t_{\mathrm{cut}} = 0$. In that case we simply test the sign of $\Delta\chi^2$, *i.e.* which hypothesis is more compatible with the data. It follows immediately that the crossing sensitivity in number of sigmas is

$$S_{\mathrm{cross}} = \frac{t_0}{2\sqrt{t_0}} = \frac{\sqrt{t_0}}{2} \tag{6.39}$$

With some simplifying assumptions on the dependance of the hypotheses with the model parameters, the three above conditions can be proven to hold: see *e.g.* [159] and [265]. In the concrete case of the mass hierarchy determination with ORCA, we will show in this thesis that they hold approximately for some, but not all considered values of the true parameters. Therefore the reference result remains the LLR sensitivity. Nevertheless the Asimov approximation can advantageously be used for quick estimations, comparisons of the impact of various systematics and optimisations of the analysis strategy.

Example of disagreement

Consider again the TS distributions of Fig. 6.1. These toy examples assume normally distributed t with

$$\sigma = 2\sqrt{\mu}, \tag{6.40}$$

but

$$|\mu_{\mathrm{IH}}| \neq |\mu_{\mathrm{NH}}| \tag{6.41}$$

The corresponding Asimov sensitivities are $\sqrt{\mu_{\mathrm{NH}}} = 1.73\sigma$ (true IH, left) and $\sqrt{\mu_{\mathrm{IH}}} = 2.24\sigma$ (true NH, right), to be compared to the LLR median sensitivies (1.79σ and 2.31σ, respectively). Even though the prescription 6.40 is respected, the fact that the LLR distributions are not symmetrical is sufficient to cause a discrepancy, the Asimov estimation being too pessimistic for both true cases. In this case the discrepancy is small. We will see in Chap. 8 that the hypothesis 6.40 does not hold exactly either: instead we find $\sigma < 2\sqrt{\mu}$, inducing an even larger discrepancy.

[6] Note that strictly speaking the 'average experiment' (the experiment yielding exactly the expectation values of the bin event counts) is not a possible outcome since in real data the event counts will be integers.

6.3 Parameter estimation and confidence intervals

6.3.1 Frequentist confidence intervals from pseudo-experiments

Besides the mass hierarchy determination, ORCA has an unprecedented potential for measuring the parameters Δm^2_{31} and θ_{23} from atmospheric oscillations. In addition we are interested in the measurement of the electron density in the Earth's core and mantle. To assess ORCA's sensitivity to these parameters we want to evaluate the typical width of (frequentist) confidence intervals we expect to obtain after performing the experiment. Let us generically denote θ the parameter which measurement is considered, and θ_{MLE} its maximum likelihood estimator (MLE). For the mass hierarchy study, following the method described in 6.1.2 we draw PEs and perform simultaneous maximum likelihood fits of all nuisance parameters, including θ, under both mass hierarchy assumptions. This provides the expected distributions of θ_{MLE}, taking into account the impact of the relevant nuisance parameters which are fitted simultaneously. Various cases can be considered regarding the mass hierarchy assumption:

- MH assumed to be known and correct: fits to the wrong hierarchy are discarded,
- MH assumed to be unknown: for each PE the fits under both MH hypotheses are compared, and θ_{MLE} is taken from the fit yielding the overall maximum likelihood
- wrong MH assumption: fits to the correct hierarchy are discarded.

Fig. 6.2 show examples of such distributions of the MLE for the parameter θ_{23}.

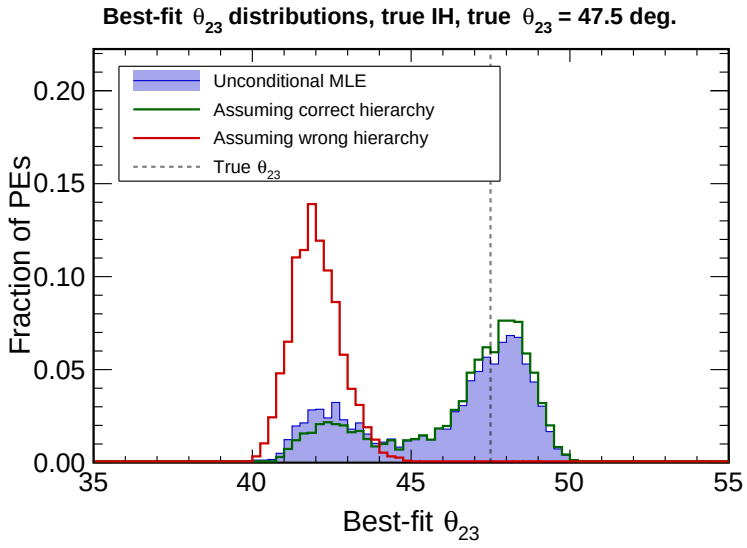

Figure 6.2: Distributions of the maximum likelihood estimator for θ_{23}, with various assumptions on the mass hierarchy (see text).

We are thus able to sample the PDF of the maximum likelihood estimator for θ. From there, the classical frequentist approach is to repeat this study for a series of assumed true values of θ and apply the method of Neyman's confidence belt construction, as described for example in the review of statistics (Chap. 40) of [44] or in Refs. [264, 266]. A choice of ordering rule is required to define the acceptance interval $[\theta_-, \theta_+]$ for each true value θ_{true} – further detail on this matter can be found *e.g.* in Refs. [264, 267].

6.3.2 Profile likelihood ratio

An alternative method is available to build approximate confidence intervals, which does not require simulating large numbers of PEs. Denoting again the parameter of interest as θ we define the profile likelihood ratio test statistic:

$$\lambda(\boldsymbol{d},\theta) = \frac{\max_{\boldsymbol{\eta}}\ \mathcal{L}(\boldsymbol{d}\,|\,\theta,\boldsymbol{\eta})}{\max_{\boldsymbol{\eta},\theta}\ \mathcal{L}(\boldsymbol{d}\,|\,\theta,\boldsymbol{\eta})}, \tag{6.42}$$

and the corresponding logarithmic ratio[7] τ:

$$\tau(\boldsymbol{d},\theta) = -2\cdot\ln\lambda(\boldsymbol{d},\theta) \tag{6.43}$$

This quantity is always positive and tests the incompatibility of the data (or pseudo-data) $\boldsymbol{d}$ with a test value of θ: large values of τ ($\lambda \ll 1$) indicate large incompatibility, while $\tau = 0$ ($\lambda = 1$) is found for the maximum likelihood estimate of θ. At the denominator, the likelihood is evaluated at the best-fit values of both θ and the nuisance parameters $\boldsymbol{\eta}$ (unconditional MLE), whereas at the numerator θ is fixed at the test value and $\mathcal{L}$ is evaluated at the corresponding conditional MLE for $\boldsymbol{\eta}$. Fig. 6.3 shows typical shapes of the τ test statistic, again applied to the measurement of θ_{23}. For comparison, the pseudo-data used in Fig. 6.3 is the Asimov set with the same true parameters as for Fig. 6.2. For a symetrically constrained parameter without degeneracies, the τ function is expected to display a symmetrical parabolic shape around its minimum, corresponding to a gaussian-like behaviour of the profile likelihood ratio [264, 268]. Fig. 6.3 shows deviations from this behaviour, with non-symmetrical shapes and multiple minima. The global minimum is found in the wrong octant when assuming the wrong hierarchy, which demonstrates the hierarchy-octant degeneracy (*cf.* Chap. 1). For the same pseudo-data, the loss of information on θ_{23} induced by the presence of nuisance parameters is reflected by the widening of the curves around the local minima, and by the comparatively smaller height difference between the two minima.

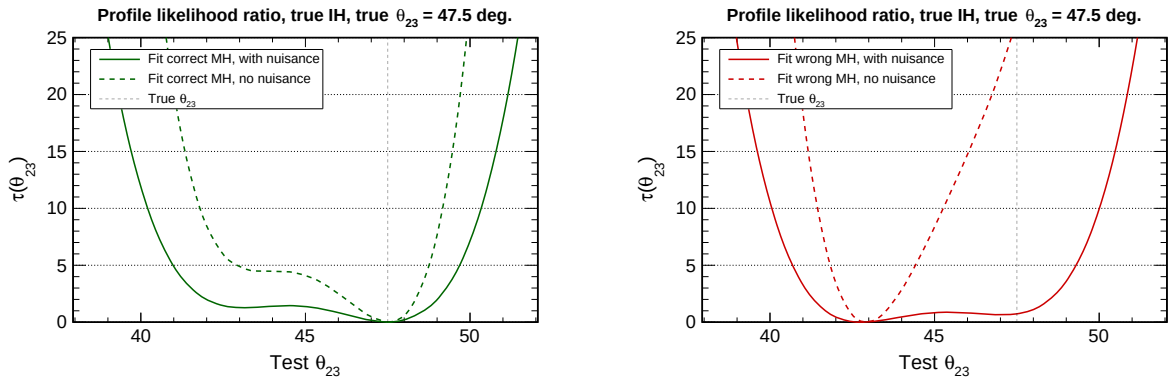

Figure 6.3: Test statistic τ as a function of the test value of θ_{23}. The pseudo-data is replaced by the Asimov data for true inverted hierarchy, with true $\theta_{23} = 47.5$ deg, $\delta_{CP} = 0$. The correct (wrong) hierarchy is assumed in the likelihood maximisation for the left (right) plot. Solid lines correspond to the maximisation including the same 9 nuisance parameters as in Fig. 6.2, while for dotted lines nuisance parameters are all kept fixed to their nominal values.

[7] Note that τ, like the test statistic t (also written $\Delta\chi^2$) defined previously for the mass hierarchy hypothesis test, can be written as a difference of χ^2, and is often denoted as $\Delta\chi^2$ as well. To avoid confusion with t we keep a separate notation here.

6.4 Accounting for Monte Carlo statistical uncertainties

As stated in introduction of this chapter, statistical fluctuations in Monte Carlo simulations, just like statistical fluctuations in data, can be taken into account in the overall likelihood function describing the agreement or disagreement between a model prediction and observed data. Following the notations and definitions used earlier, let us consider explicitly the observed response matrix R as a realisation of a random process, with the expectation values for each entry described by $\underline{R}$. This means that for all channels $[\nu_x \to i]$ and bins $(E,\theta,y,E',\theta',y')$:

$$R^{[\nu_x\to i]}(E,\theta,y,E',\theta',y') \xrightarrow[N^{\mathrm{MC}}\to\infty]{} \underline{R}^{[\nu_x\to i]}(E,\theta,y,E',\theta',y'), \tag{6.44}$$

and setting aside systematic uncertainties and model limitations, we can think of $\underline{R}$ as the underlying 'true' detector response that would be obtained with infinite Monte Carlo statistics. Likewise, the underlying true predictions for the event counts are denoted $\underline{\mu}(\boldsymbol{\eta},\underline{R})$, where the model parameters $\boldsymbol{\eta}$ comprise both the nuisance parameters and, if applicable, the parameters of interest. In contrast the computed predictions $\mu(\boldsymbol{\eta},R)$ depend on the particular realisation R obtained from a finite Monte Carlo sample and are therefore interpreted as random. We can rewrite and expand the likelihood of Sec. 6.1.1 to obtain $\mathcal{L}'$ as follows:

$$\mathcal{L}' = \mathcal{L}(\boldsymbol{d}\,|\,\underline{\mu}(\boldsymbol{\eta},\underline{R}))\cdot\mathcal{L}(R\,|\,\underline{R}) \tag{6.45}$$

Provided a probability model describing the fluctuations of R around $\underline{R}$ (expected to depend on the size of the MC sample), it is in principle possible to maximise $\mathcal{L}'$ with respect to the entries of R and the model parameters $\boldsymbol{\eta}$ simultaneously. Although this implies a very large number of free parameters (one per bin of the response matrix), the assumption that the entries of R are bin-to-bin uncorrelated and that their fluctuations are independent of the model parameters $\boldsymbol{\eta}$ allows to separate the maximisations and, in some cases, to treat part of the problem analytically. This procedure is known as the 'Beeston and Barlow method', originally described in [269]. In this work we have implemented an approximation of the above method known as 'Beeston and Barlow light method', widely used in particle physics analyses

$$\underline{\mu}_k = \beta_k \mu_k \tag{6.46}$$

and assume that β_k is normally distributed[8]:

$$\beta_k \sim \mathcal{N}(1,\sigma_k^2) \tag{6.47}$$

The statistical and systematic terms (*cf.* Eq. 6.8, 6.13), previously written

$$\chi^2 = \chi^2_{\mathrm{stat}}(\boldsymbol{d}\,|\,\mu_0(\boldsymbol{\eta}),\mu_1(\boldsymbol{\eta}),\ldots) + \chi^2_{\mathrm{syst}}(\boldsymbol{\eta}) \tag{6.48}$$

then become

$$\chi^2 = \chi^2_{\mathrm{stat}}(\boldsymbol{d}\,|\,\beta_0\mu_0(\boldsymbol{\eta}),\beta_1\mu_1(\boldsymbol{\eta}),\ldots) + \sum_k \frac{(\beta_k-1)^2}{\sigma_k^2} + \chi^2_{\mathrm{syst}}(\boldsymbol{\eta}) \tag{6.49}$$

Mathematically the new parameters β_k act exactly as standard nuisance parameters constrained by uncorrelated gaussian priors, modelling bin-by-bin uncorrelated uncertainties. To evaluate the σ_k prior uncertainties the bootstrapping method described in Sec. 5.3.4 is applied – noting also that the bootstrap distributions observed in Fig. 5.6 seem consistent with the gaussian model.

[8] From the probabilistic point of view it seems more natural to take $\frac{1}{\beta_k}$ as normally distributed. As long as σ_k is of the order of a few percent this is practically equivalent, and having a gaussian constraint on β_k is easier in the calculations.

An important feature of the method is that there is no need to perform a joint minimisation over the $\boldsymbol{\beta}$ and $\boldsymbol{\eta}$ parameters. Indeed, since $\boldsymbol{\mu}$ and $\boldsymbol{\eta}$ do not depend on $\boldsymbol{\beta}$, the analytical extremum condition for a single parameter β_k simplifies to

$$\frac{\partial \chi^2}{\partial \beta_k} = \frac{\partial}{\partial \beta_k}\Big[-2(d_k \ln(\beta_k \mu_k) - \beta_k \mu_k) + \frac{(\beta_k - 1)^2}{\sigma_k^2}\Big] = 0. \tag{6.50}$$

This leads to a quadratic equation which only positive solution is

$$\beta_k = \frac{1}{2}\Big[1 - \mu_k \sigma_k^2 + \sqrt{(1 - \mu_k \sigma_k^2)^2 + 4 n_k \sigma_k^2}\Big] \tag{6.51}$$

In practice, for each trial value of $\boldsymbol{\eta}$ in the numerical χ^2 minimisation, after computing $\mu_k(\boldsymbol{\eta})$ with the normal procedure the corresponding MLE of β_k and therefore of $\underline{\mu}_k$ are inferred at no cost, and can be used in the χ^2 computation. The introduction of the $\boldsymbol{\beta}$ uncertainties thus require no additional minimisation step. At the moment, the main computational limitation of this method is the application of the bootstrapping method described in Sec. 5.3.4, which should in principle be repeated for each analysis setup – *e.g.* when changing the definition of PID classes.

Chapter 7

Systematic uncertainties

Contents

In the context of particle physics experiments, *systematic* uncertainties, set apart from *statistical* uncertainties, are *"the uncertainties associated with the nature of the measurement apparatus, assumptions made by the experimenter, or the model used to make inferences based on the observed data"* [270]. From the mathematical point of view, the general treatment of systematic uncertainties as nuisance parameters in the statistical analysis has been exposed in Chap. 6. This chapter focuses on the modeling of the concrete sources of systematic uncertainties and their implementation in the analysis framework. For ORCA, the imperfect knowledge of oscillation parameters is a very important source of systematic uncertainty. The term 'systematics' will generally be used to refer to *other* sources, the most important for ORCA being

(a) the atmospheric neutrino flux,
(b) neutrino interactions,
(c) and detection effects.

The first category involves uncertainties on the normalisation, composition (flavour and neutrino-to-antineutrino fractions), and the dependence (shape) of the flux in energy and zenith angle. The modeling and implementation of flux uncertainties in the analysis are detailed in Sec. 7.1.1 and 7.1.2, while Sec. 7.1.3 discusses some of the existing constraints. The second category concerns the normalisation and shape of the overall interaction cross-section, as well as the relative contributions of the different interaction channels to which the detector responds differently. Cross-section systematics are discussed in Sec. 7.2.1. The detection systematics potentially include a wide range of effects from the uncertainties in the seawater optical properties, light propagation model, response of the photomultipliers and associated electronics, to systematic uncertainties arising from the Monte Carlo simulation methods and their interplay with the reconstruction and classification algorithms. They are discussed briefly in Sec. 7.2.2. Finally, the treatment of oscillation and Earth model uncertainties in the standard NMH analysis is shortly summarised in Sec. 7.3.

7.1 Flux systematics

7.1.1 Flux normalisation and composition

In the following we denote the flux component for $\nu_\alpha = \nu_e,\ \overline{\nu}_e,\ \nu_\mu,\ \overline{\nu}_\mu$ as $\phi_\alpha = \phi_{\nu_\alpha}(E,\theta)$ with $\alpha = e,\overline{e},\mu,\overline{\mu}$, and the corresponding flux integral as

$$I_\alpha = \int\int \phi_\alpha(E,\theta)\ \mathrm{d}E\ \mathrm{d}\theta, \tag{7.1}$$

computed over the range of true energy and zenith angle set for the analysis. The flux "perturbated" by the action of nuisance parameters is denoted $\widetilde{\phi}_\alpha$ and its integral $\widetilde{I}_\alpha$.

Overall normalisation

When considering normalisation uncertainties, the signal shape (distribution with respect to E,θ,y) is ignored. The contribution to the event channel ν_β (oscillated flavour) from the "source" ν_α (unoscillated flavour) can then be written schematically as

$$N(\beta) = R(\beta) \times \sigma(\beta) \times P(\alpha \to \beta) \times \phi(\alpha), \tag{7.2}$$

where R, σ, P and ϕ denote respectively the efficiency (response matrix), cross-section, oscillation probability and flux. Setting apart NC events, the total ν_β event count is obtained by summing on the unoscillated flavours ν_α.

As a first step to model flux uncertainties, an overall flux normalisation parameter f_{all} can be introduced:

$$\phi_\alpha \to \widetilde{\phi}_\alpha = f_{all} \times \phi_\alpha \tag{7.3}$$

f_{all} is then treated as a nuisance parameter, *i.e.* it is allowed to float – with or without constraint – in the likelihood maximisation.

If it is considered as independent of the flavour α, f_{all} factors out from the sum on α and therefore acts as a global normalisation effect, effectively absorbing any flavour-independent normalisation effect (*e.g.* overall cross-section normalisation, detection efficiency, *etc*). In case flavour-dependent of channel-dependent normalisations are considered, care should be taken to avoid introducing degeneracies with f_{all}.

Flux composition

To model flux composition uncertainties, individual normalisation parameters f_α are introduced:

$$\widetilde{\phi}_{\nu_\alpha} = f_{\nu_\alpha}\phi_{\nu_\alpha}. \tag{7.4}$$

To avoid degeneracies with the overall normalisation the conservation of the total flux is enforced:

$$\widetilde{I}_e + \widetilde{I}_{\overline{e}} + \widetilde{I}_\mu + \widetilde{I}_{\overline{\mu}} = I_e + I_{\overline{e}} + I_\mu + I_{\overline{\mu}} \tag{7.5}$$

A degree of freedom is thus removed and $f_{\nu_e},\ f_{\overline{\nu}_e},\ f_{\nu_\mu},\ f_{\overline{\nu}_\mu}$ can be reparametrised with only three parameters:

- the flavour skew $s_{e\mu}$
- the $\nu_e/\overline{\nu}_e$ skew $s_{e\overline{e}}$
- the $\nu_\mu/\overline{\nu}_\mu$ skew $s_{\mu\overline{\mu}}$

The parametrisation reads:

$$f_e = (1+s_{e\mu}) \times (1+s_{e\bar{e}}) \tag{7.6}$$

$$f_{\bar{e}} = (1+s_{e\mu}) \times (1-\frac{I_e}{I_{\bar{e}}} \cdot s_{e\bar{e}}) \tag{7.7}$$

$$f_\mu = (1-\frac{I_e+I_{\bar{e}}}{I_\mu+I_{\bar{\mu}}} \cdot s_{e\mu})(1+s_{\mu\bar{\mu}}) \tag{7.8}$$

$$f_{\bar{\mu}} = (1-\frac{I_e+I_{\bar{e}}}{I_\mu+I_{\bar{\mu}}} \cdot s_{e\mu})(1-\frac{I_\mu}{I_{\bar{\mu}}} \cdot s_{\mu\bar{\mu}}) \tag{7.9}$$

With this implementation the action of $s_{e\bar{e}}$ (resp. $s_{\mu\bar{\mu}}$) preserves the total flux of the electron (resp. muon) flavour, *i.e.* $\widetilde{I}_e + \widetilde{I}_{\bar{e}}$ and $\widetilde{I}_\mu + \widetilde{I}_{\bar{\mu}}$ are functions of $s_{e\mu}$ only.

7.1.2 Shape uncertainties

Flux shape uncertainties are introduced as perturbations

$$\phi_\alpha(E,\theta) \to f_{\text{shape}}(E,\theta) \times \phi_\alpha(E,\theta), \tag{7.10}$$

where the normalisation of the f_{shape} function is set so that it does not modify the total flux integral

$$\int\int f_{\text{shape}}(E,\theta) \times \phi_\alpha(E,\theta) = \int\int \phi_\alpha(E,\theta) \tag{7.11}$$

In the analysis framework, a generic module implementing f_{shape} as a polynomial function of $(X = \ln E, Y = \cos\theta)$ has been designed. The typical shape variations considered will be continuous functions varying reasonably slowly, so that in most cases they can be approximated to a good precision with a finite degree polynomial. In practice, the module allows the use of polynomials up to a total degree $D = 4$ in both variables: the most generic f_{shape} is therefore

$$\frac{f_{\text{shape}}(E,\theta)}{\text{Norm}(\{c_{ij}\})} = 1 + \sum_{1\le i+j\le 4} c_{ij} (\ln E)^i (\cos\theta)^j = 1 + \sum_{1\le i+j\le 4} c_{ij} X^i Y^j, \tag{7.12}$$

where $\text{Norm}(\{c_{ij}\})$ is a renormalising term ensuring the condition 7.11. Implementing a particular shape function is done by specifying a parametrisation of the coefficients c_{ij} with respect to a set of nuisance parameters. Two one-dimensional examples are detailed thereafter: spectral index uncertainty and zenith slope uncertainty.

This approach has the practical advantage that the correct normalisation of f_{shape} can be ensured in a completely automatic manner and with almost no numerical overhead. Indeed, it is sufficient to set

$$\text{Norm}(\{c_{ij}\}) = \frac{I_\alpha}{I_\alpha + \sum_{i,j} c_{ij} I_{\alpha,ij}} \tag{7.13}$$

where $I_{\alpha,00} = I_\alpha$ and

$$I_{\alpha,ij} = \int\int X^i Y^j \phi_\alpha(X,Y) \mathrm{d}X \mathrm{d}Y \tag{7.14}$$

and the integrals $I_{\alpha,ij}$ only need to be computed once, at initialisation of the module.

Spectral index uncertainty

Physically, the uncertainty on the atmospheric spectral index is modeled by introducing a 'tilt' parameter γ:

$$\phi_\alpha(E,\theta) \to E^{-\gamma} \times \phi_\alpha(E,\theta) \tag{7.15}$$

At small trial values $\gamma \ll 1$, $E^{-\gamma}$ is well approximated by its fourth order series expansion:

$$f_{\text{spectral}} = \text{Norm}(\gamma) \cdot \left(1 + \gamma \ln E + \frac{1}{2}\gamma^2 \ln^2 E + \frac{1}{6}\gamma^3 \ln^3 E + \frac{1}{24}\gamma^4 \ln^4 E\right), \tag{7.16}$$

which is the implementation used in the framework. The explicit form for $\text{Norm}(\gamma)$ need not be provided, as it is computed automatically following 7.13. Fig. 7.1 illustrates the effect of varying the tilt from $\gamma = -0.2$ to $\gamma = 0.2$.

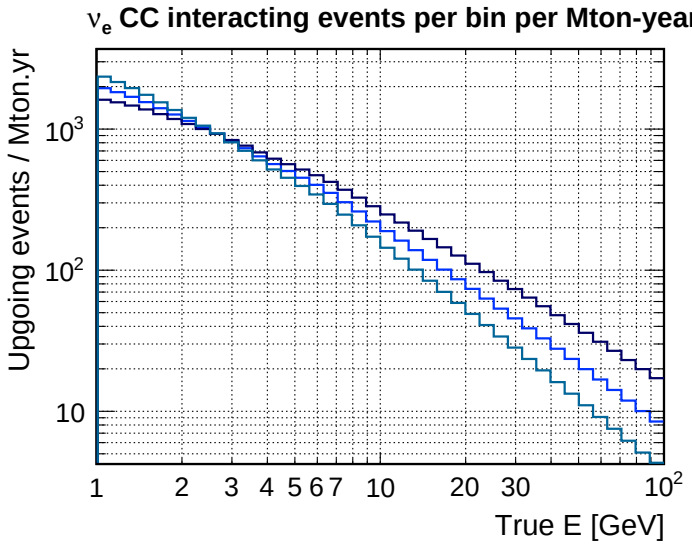

Figure 7.1: Illustration of the effect of the spectral index systematic on the ν_e CC interacting event rate: $\gamma = \{-0.2,\ 0,\ +0.2\}$.

Zenith slope uncertainty

A zenith slope tilt is implemented similarly to the energy slope tilt, except that the perturbation is taken to be linear in $\cos\theta$:

$$f_{\text{zenith}} = \text{Norm}(\epsilon_\theta) \cdot \left(1 + \epsilon_\theta \cos\theta\right) \tag{7.17}$$

Other parametrised shape effects

Up to now, dedicated shape systematics modules and their interface with the minimisation and experiment simulation modules have been implemented in the software framework to enable

- the energy slope systematic alone;
- the zenith slope systematic alone;
- energy and zenith slope systematics together[1];
- generic polynomial shape perturbations, with up to 14 free parameters.

The impact of the first three on the oscillation analysis has been characterised, as reported in the next chapter. Generic shape perturbations have not yet been characterised in detail.

[1] The introduction of a dedicated module for applying energy and zenith slope systematics together, instead of simply enabling both independent modules, is motivated by the fact that $\text{Norm}(\gamma, \epsilon_\theta) \neq \text{Norm}(\gamma) \times \text{Norm}(\epsilon_\theta)$ (with reference to Eqs. 7.16 and 7.17).

7.1.3 Constraints on flux uncertainties

A detailed study of the uncertainties associated to the calculations of neutrino production in the atmosphere was done in Ref. [86]. The study includes overall normalisation, flavour and neutrino/antineutrino ratios, as well as some shape effects (*e.g.* up/down and up/horizontal ratios). As discussed in Sec. 2.1.1, hadronic interaction models and the parametrisation of the primary CR flux are the main sources of uncertainty. Fig. 7.2, shows the uncertainties in the $\nu_\mu/\overline{\nu}_\mu$, $\nu_e/\overline{\nu}_e$, and $(\nu_\mu+\overline{\nu}_\mu)/(\nu_e+\overline{\nu}_e)$ ratios as evaluated in Ref. [86]. It can be seen that the uncertainties increase with energy. As the energy range of interest here is $1-100\,\text{GeV}$, a conservative set of constraints based on the results from Ref. [86] is the following:

- 10% prior on the $\nu_\mu/\overline{\nu}_\mu$ and $\nu_e/\overline{\nu}_e$ ratios,
- 3% prior on the flavour ratio.

More generally, the impact of various levels of constraints is studied in the next chapter (Sec. 8.3).

Implementation-wise, it is important to note that the gaussian prior constraints applied in the analysis act directly onto the skew parameters $s_{e\mu}$, $s_{e\overline{e}}$, and $s_{\mu\overline{\mu}}$. From the definitions (7.6-7.9) it can be deduced that the perturbation of the $\nu/\overline{\nu}$ flux ratios are related to the skew parameters as:

$$\frac{f_e}{f_{\overline{e}}} = 1 + \left(1 + \frac{I_e}{I_{\overline{e}}}\right) s_{e\overline{e}} + \mathcal{O}(s^2_{e\overline{e}}) \simeq 1 + 2.3 \times s_{e\overline{e}}, \tag{7.18}$$

$$\frac{f_\mu}{f_{\overline{\mu}}} = 1 + \left(1 + \frac{I_\mu}{I_{\overline{\mu}}}\right) s_{\mu\overline{\mu}} + \mathcal{O}(s^2_{\mu\overline{\mu}}) \simeq 1 + 2.1 \times s_{\mu\overline{\mu}}. \tag{7.19}$$

As to the flavour ratio, defining

$$r_{e\mu} \equiv \frac{I_e + I_{\overline{e}}}{I_\mu + I_{\overline{\mu}}}, \tag{7.20}$$

the perturbation to the flavour ratio can then be expressed as:

$$\begin{aligned} \frac{\widetilde{r}_{e\mu}}{r_{e\mu}} &= \Big(1 + (1 + r_{e\mu}) s_{e\mu}\Big) \\ &\simeq (1 + 1.5 \times s_{e\mu}). \end{aligned} \tag{7.21}$$

The correspondance between constraints on overall (integral) flux ratios and skew parameters can thus be summarised as:

$$\pm 0.01 \text{ in } s_{e\overline{e}} \Rightarrow \text{approx. } \pm 2.3\% \text{ in } \Phi_e/\Phi_{\overline{e}} \tag{7.22}$$

$$\pm 0.01 \text{ in } s_{\mu\overline{\mu}} \Rightarrow \text{approx. } \pm 2.1\% \text{ in } \Phi_\mu/\Phi_{\overline{\mu}} \tag{7.23}$$

$$\pm 0.01 \text{ in } s_{e\mu} \Rightarrow \text{approx. } \pm 1.5\% \text{ in } r_{e\mu} \tag{7.24}$$

Flux shape uncertainties have not been characterised as precisely as normalisation and composition uncertainties in the literature, though also discussed in Ref. [86]. Sensitivity studies focusing on PINGU and ORCA-like detectors have shown that the impact of shape uncertainties in general should not be neglected [160, 161]. As an illustration of potential mismodeling effects leading to variations of the overall shape of the event distribution, Fig. 7.3 shows the difference (signed χ^2) between the event distributions obtained when using the HAKKM2014 flux models [85] simulated for the Gran Sasso and Fréjus sites (both without mountain over the detector). Both these sites are geographically close to the ORCA site, Fréjus being the closest.

Fig. 7.4 shows the impact of modeling neutrino production as distributed uniformly between 8 and 30 km height instead of at a fixed height of 15 km. In this case, the dominant effect is the

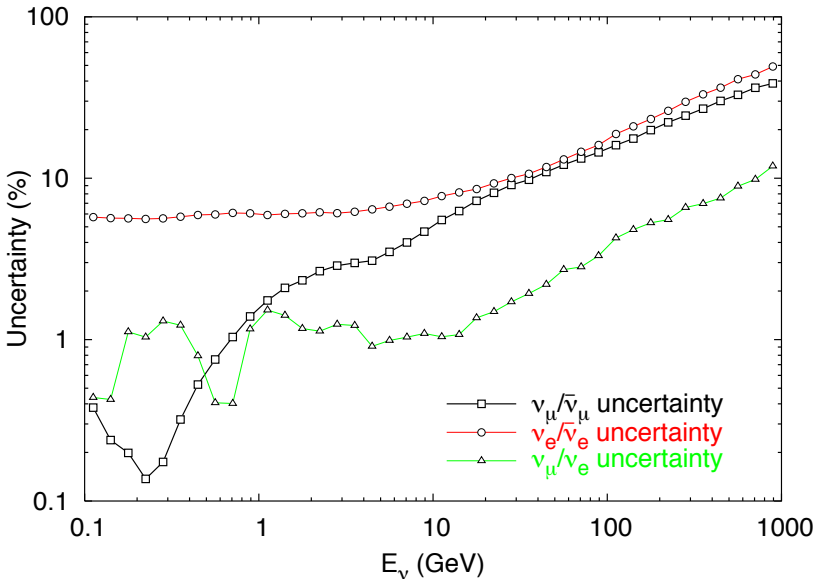

Figure 7.2: Uncertainties in neutrino-type ratios as a function of neutrino energy. $\nu_\mu/\overline{\nu}_\mu$ is shown with (black) lines with squares, $\nu_e/\overline{\nu}_e$ with (red) lines with circles and $(\nu_\mu+\overline{\nu}_\mu)/(\nu_e+\overline{\nu}_e)$ with (green) lines with triangles. Figure taken from Ref. [86].

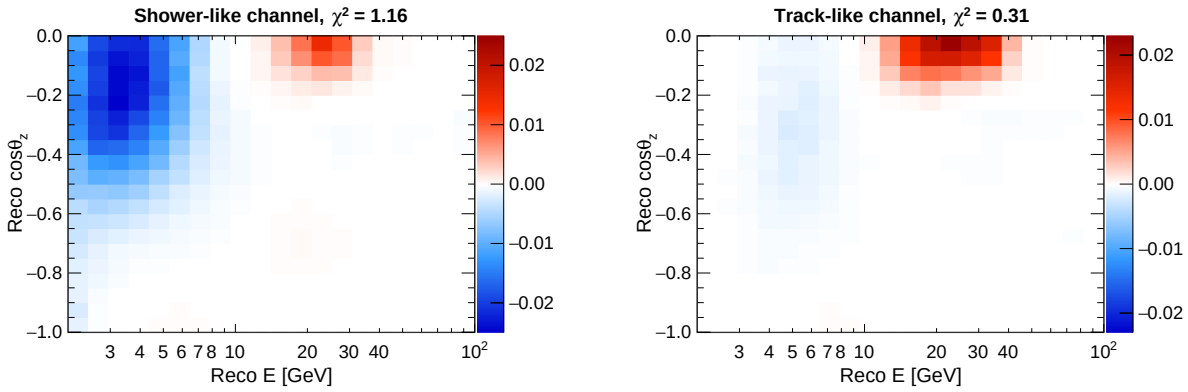

Figure 7.3: Impact of the choice of flux model (detector site) on the reconstructed event distributions. The plotted quantity is the $\chi^2(n_A,n_B)$ introduced in Chap. 6, multiplied with the sign of (n_A-n_B). A = Gran Sasso site without mountain over the detector. B = Fréjus site without mountain over the detector.

existence of a continuous distribution of oscillation baselines instead of a single one at a given zenith angle. As expected, this effect is mainly relevant for close to horizontal trajectories.

7.2 Cross-section systematics and detector effects

7.2.1 Cross-section

The dedicated modeling of cross-section uncertainties is somewhat limited in this work. In contrast with flux normalisation, cross-section normalisation or skew parameters should be applied according to the oscillated flavour β (referring to Eq. 7.2). In addition the CC or NC interaction channel has to be considered. A major simplification follows from lepton universality: electron and muon neutrino cross-sections are practically equal at the considered energies $E_\nu >> m_e, m_\mu$), and the cross-section ratios $\frac{\sigma(\nu_\tau)}{\nu_{\mu,e}}$ and $\frac{\sigma(\overline{\nu}_\tau)}{\overline{\nu}_{\mu,e}}$ are known to a very good precision. Furthermore, due to the absence of mixing between neutrinos and antineutrinos the $\nu/\overline{\nu}$ ratio

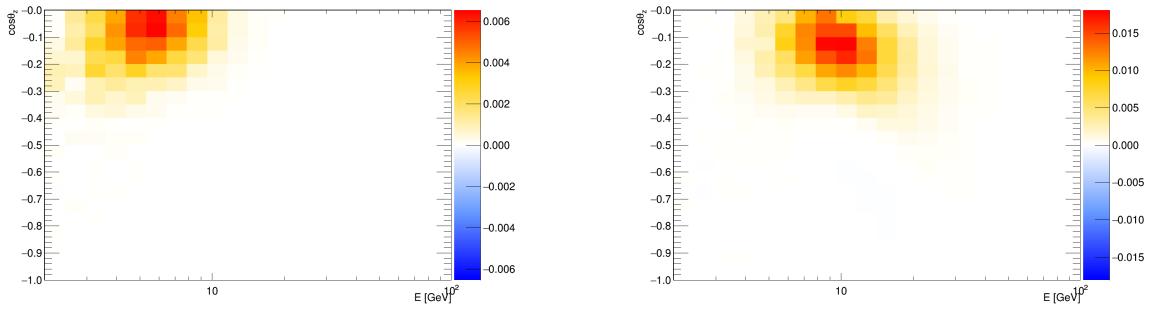

Shower-like channel
$\sum(n_A - n_B)/(n_A) = 0.06\%$
$\sum|\chi^2(n_A, n_B)| = 0.18$

Track-like channel
$\sum(n_A - n_B)/(n_A) = 0.2\%$
$\sum|\chi^2(n_A, n_B)| = 0.5$

Figure 7.4: Impact of the atmospheric neutrino production height on the shower-like (left) and track-like (right) reconstructed event distributions seen in ORCA. The plotted quantity is the $\chi^2(n_A, n_B)$ introduced in Chap. 6, multiplied with the sign of $(n_A - n_B)$. A = fixed neutrino production at 15 km altitude. B = uniformly distributed neutrino production between 8 and 30 km. The total relative difference in event count and total $|\chi^2(A, B)|$ corresponding to 3 years exposure to a full detector are indicated. These figures are preliminary results of a study by N. Chau, master intern at APC (under my supervision).

in terms of *rate* (flux × cross-section) is closely related to the corresponding flux ratio (they are proportional except for the effect of ν_τ appearance).

Therefore, additional skew parameters describing the $\nu/\overline{\nu}$ cross-section ratio would be mostly degenerate with $s_{e\overline{e}}$ and $s_{\mu\overline{\mu}}$ introduced in the previous section. The only dedicated normalisation parameter introduced standardly in the framework is then a scaling f_{NC} applied to the rate of neutral current events. A normalisation parameter for ν_τ CC events is also implemented, originally to study the sensitivity to ν_τ appearance. It can also be used to evaluate the impact of a systematic uncertainty on ν_τ CC cross-section modeling.

Discussion and outlook

As seen in Chap. 3, up to neutrino energies of approximately 20 GeV several interaction channels contribute significantly to the overall cross-section, and their relative contributions vary with energy. Different uncertainties are associated to each interaction channel, moreover the size of these uncertainties is also energy-dependent. As a consequence, energy-independent normalisation parameters only roughly model cross-section uncertainties, particularly in the low to intermediate energy range 1–20 GeV. To some extent, generic shape systematics may cover part of the energy-dependent effects induced in this range. However such generically parametrised uncertainties can easily be too conservative, or on the contrary overoptimistic. A detailed study of the impact of cross-section uncertainties has been reported for the PINGU experiment in Ref. [271], showing that such uncertainties are not dominant in this type of experiment given the energy range of interest (see also the discussion of cross-section uncertainties in Sec. 2.1.2).

As a first step for improvement which would require only a minimal adaptation of the framework, the uncertainty could be parametrised by calculating the total cross-section as a sum of contributions from the quasi-elastic, resonant and deep inelastic scattering processes, assigning a specific uncertainty (normalisation) to each process, and fitting these parameters separately. To go further and include genuine cross-section uncertainty models in the analysis, an effort to integrate a smearing technique based on event-by-event reweighting into this analysis framework is currently undergone within the KM3NeT collaboration. With such a detector response implementation, uncertainties on cross-section model parameters can be propagated to individual event weights, making use of existing features of neutrino event generators like `GENIE`.

7.2.2 Detector effects

A variety of detection-related effects can introduce additional systematic errors. As a prototypical example, incorrect assumptions on the light absorption or scattering properties of the detection medium would lead to misevaluating the light yield, and therefore the detection efficiency as well as the energy response properties of the detector. Similarly to the case of cross-section effects, it can be argued that part of the resulting uncertainties are effectively absorbed by generic normalisation or shape parameters, when such parameters are fitted unconstrained. To go further and study this systematic separately, a "migration" mechanism can be introduced.

Energy shift

Consider events simulated with a Monte Carlo neutrino energy in a bin E_{MC}. The light yield may be underevaluated such that in reality, neutrinos with an energy $E_{\mathrm{real}} < E_{\mathrm{MC}}$ produce a signal resembling the one expected for an energy E_{MC}. Assuming that the simulation of the detector is otherwise accurate, the data will therefore be fitted better by allowing the energy E_{detector} used as input to the detector model to shift with respect to the energy E_{rate} used in the calculation of the interacting event rates:

$$E_{\mathrm{detector}} = (1+\alpha_E)E_{\mathrm{rate}} \tag{7.25}$$

This is implemented in practice by migrating the events in the interacting event rate histograms. When the trial value of α_E is non-zero, for each bin $[E_i - E_{i+1}]$ in a rate histogram the event count $N(E_i \leq E_{\mathrm{rate}} \leq E_{i+1})$ is redistributed to the bins covered by the shifted interval $[(1+\alpha_E)E_i - (1+\alpha_E)E_{i+1}]$, proportionnally to the overlapping area.

Normalisations of the final event distributions

In Sec. 7.1.1 we discussed the introduction of an overall floating normalisation f_{all} common to all channels, initially thought as modelling a flux normalisation uncertainty. A more conservative approach is to introduce an independent normalisation f_β for each reconstructed event class β. In the standard 2-class approach, this introduces two parameters f_{track} and f_{cascade}. As pointed out in Sec. 7.1.1, applying an overall f_{all} together with the f_β parameters would introduce a degeneracy in the fit. For this reason it is avoided; instead only the f_β's – or f_{all} alone – are fitted. Separating an overall normalisation effect from asymmetry effects can still easily be done, *e.g.* with the convention:

$$f'_{\mathrm{all}} = \frac{1}{n_\alpha}\sum_{1}^{n_\alpha} f_\alpha \tag{7.26}$$

$$f'_\alpha = \frac{f_\alpha}{f'_{\mathrm{all}}} \tag{7.27}$$

For example in the 2-class case, either the normalisation ratio $f_{\mathrm{track}}/f_{\mathrm{cascade}}$ or relative difference $f'_{\mathrm{track}} - f'_{\mathrm{cascade}}$ can be used as proxies for asymmetry effects in the particle identification performance. This can be used to apply separate constraints to the overall normalisation and PID asmmetry effects, or retrospectively to evaluate their relative impacts from the best-fit normalisations ("pulls").

Outlook

The introduction of the separate f_β parameters accounts for detector-related uncertainties only to the extent of normalisation effects. Other physically motivated uncertainties concern the shape of detection or classification efficiency functions (see Fig. 3.12 and 5.4). A detailed

discussion of these effects is however outside the scope of this work. Currently, physically motivated shape variations are more easily studied using the alternative analysis framework based on parametrised functions discussed in Sec. 5.3.2. In the long term a shape systematics module similar to the one used for flux systematics could be implemented in the full MC framework as well. Likewise, the event-by-event reweighting technique will allow more advanced studies taking into account individual event characteristics such as particle multiplicity or decay channels – at the cost of vastly increased running time.

7.3 Oscillation and Earth parameters

The true values and ranges of all oscillation parameters, as well as their treatment as nuisance parameters, are summarised in Tab. 7.1 for the 'standard' NMH analysis presented in this thesis. The default true values mostly correspond to the global fit values reported in Tab. 1.1, taken from Refs. [63, 163]. Since θ_{23} and δ_{CP} are still weakly constrained, scanning over a large range of potential truth is favoured. As demonstrated in Sec. 2.2.3, the imperfect knowledge of θ_{23}, $|\Delta m^2_{31}|$ and δ_{CP} is expected to have the largest impact on the NMH determination performance, because their effect on the oscillation probabilities is partly degenerate with that of the NMH. Since they are also susceptible to be measured by ORCA (θ_{23}, $|\Delta m^2_{31}|$) or no stringent constraint exists yet (δ_{CP}), these parameters are thus fitted without external constraint. On the contrary, the solar parameters have very little effect on the measurable signal at ORCA and are thus kept fixed in the analysis. The reactor mixing angle θ_{13} has a large impact on oscillation probabilities, but cannot be constrained from atmospheric data nearly as precisely as from the current reactor data. As demonstrated in Sec. 8.2, applying a prior is thus necessary to exploit the full potential of the experiment for measuring the NMH and the atmospheric parameters.

Parameter	True value or range	Treatment	Prior width
θ_{12} (°)	33.62	fix	-
Δm^2_{21} (eV^2)	$7.40 \cdot 10^5$	fix	-
θ_{13} (°)	8.54	fitted	0.15
Δm^2_{31} NH (eV^2)	$+2.494 \cdot 10^{-3}$	fitted	free
Δm^2_{31} IH (eV^2)	$-2.391 \cdot 10^{-3}$	fitted	free
θ_{23} default (°)	45	fitted	free
$\sin^2\theta_{23}$ default	0.5	fitted	free
θ_{23} range (°)	$38 - 52$	-	-
$\sin^2\theta_{23}$ range	$0.38 - 0.62$	-	-
δ_{CP} default	0	fitted	free
δ_{CP} range	$0 - 2\pi$	-	-

Table 7.1: Oscillation parameters: default true values (or ranges) and treatment as nuisance parameters in the default NMH analysis.

In cases where the large mass splitting is not fitted as a nuisance parameter, the choice of definition of the MH flip has an importance, as shown in Sec. 8.2. Unless otherwise specified, the definition of Tab. 7.1 is used for true values and the model value of $|\Delta m^2_{31}|$ is fitted freely.

Finally, the Earth model parameters (Z/A scalings) are kept fixed to the values listed in Tab. 4.1 – in the standard NMH analysis. Nevertheless, their impact as nuisance parameters is studied in Sec. 8.2.

Minimisation

The χ^2 minimisation is implemented using the `Minuit2` library [272], developed specifically for high energy physics. As a default, the minimisation is performed by the `MIGRAD` method, a general purpose, widely used minimisation method based on a gradient descent algorithm with a variable metric builder.

Due to the non-linear internal transformation of parameters performed by `Minuit2` when limits are provided, its authors recommend avoiding the definition of parameter limits when possible [272]. For mixing angles and the CP phase no limits are needed since they enter as cyclic parameters in the physical model; their transformation back to the conventional range, when needed, is done after the fit values are returned (for instance $0 \leq \theta_{23} \leq \pi/2$, or $0 \leq \delta_{\text{CP}} < 2\pi$). Regarding the mass hierarchy, the sign of Δm^2_{31} is never effectively fixed in the fit. Instead, when a fixed NMH hypothesis is assumed, the value passed by the minimiser is flipped sign, if needed, before passing it to the oscillation calculator.

As will be seen in the next chapter, the χ^2 function frequently exhibits two local minima with respect to θ_{23}, located in the lower and upper octant. This is due to the octant degeneracy introduced in Sec. 8.2: most dominant oscillation terms are controlled by $\sin^2 2\theta_{23}$. Therefore, to ascertain that the global best-fit value is always found, the fit is repeated twice with one starting value per octant (usually 40° and 50°). No similar behaviour has been found for other parameters, *i.e.* the starting value is found to have no influence on the found minimum.

Part III

Sensitivity results

Chapter 8

Sensitivity of ORCA to the Mass Hierarchy determination

Contents

This chapter focuses on the results of my analysis aimed at estimating ORCA's sensitivity to the determination of the neutrino mass hierarchy. The NMH signal is studied in Sec. 8.1, including detector effects in successive stages. The flavour content and (E', θ') distributions of track-like and shower-like events are then discussed along with the associated statistical signal, with a qualitative comparison of the correlated and uncorrelated smearing methods. The importance of sparse MC effects is demonstrated, and it is shown that the correction methods presented in Chap. 5 and 6 allow to obtain unbiased sensitivity estimates. Oscillation parameter and systematic uncertainties are then investigated in Sec. 8.2 and 8.3, and some preliminary optimisations of the analysis strategy are presented in Sec. 8.4. Finally, the results of a full frequentist analysis, as presented officially at the Neutrino 2018 conference, are discussed in Sec. 8.5.

8.1 Event distributions and NMH statistical signal

In this section the NMH 'signal', *i.e.* the statistical separation between the event distributions expected assuming NH and IH, is studied. Starting at the level of detected events, the effect of the inability to distinguish neutrinos from antineutrinos, the energy and angular resolutions, and the realistic flavour identification capabilities all contribute to a decrease of the very large initial statistical separation.

Conventions to study the NMH statistical separation

The indicator chosen to visualise this statistical separation in the (E,θ) space is a χ^2-like quantity, the so-called 'signed χ^2'. In a given bin of the event distribution it is defined as:

$$\chi^2_s = \frac{n_{\mathrm{NH}} - n_{\mathrm{IH}}}{|n_{\mathrm{NH}} - n_{\mathrm{IH}}|} \cdot \frac{1}{2}\Big(\chi^2_{\mathrm{bin}}(\mathrm{NH}\,|\,\mathrm{IH}) + \chi^2_{\mathrm{bin}}(\mathrm{IH}\,|\,\mathrm{NH})\Big) \tag{8.1}$$

where n_{NH} and n_{IH} are the respective event counts in this bin, and (following the notation in Chap. 6):

$$\chi^2_{\mathrm{bin}}(A\,|\,B) = 2\cdot(n_A - n_B + n_A \ln\frac{n_A}{n_B}). \tag{8.2}$$

A positive (resp. negative) χ^2_s means that an excess (resp. deficit) of events in the bin is expected in the NH assumption with respect to the IH assumption. This definition allows to correctly render the relative contribution of each bin to the total statistical significance[1], with an average over the true NH and true IH cases.

When the "total χ^2" (or simply χ^2) is indicated, it corresponds to

$$\chi^2 = \sum_{\mathrm{bins}} |\chi^2_s|. \tag{8.3}$$

Importantly, for the simple comparisons performed in this section, event distributions are calculated using the respective central values of Δm^2_{31} found specifically under the NH and IH assumptions in the global fit [63]. The importance of this convention and the relevance of this choice with respect to the standard case where $|\Delta m^2_{31}|$ is fitted as a nuisance parameter are discussed in Sec. 8.2.1. The other oscillation parameters are set identically in both hierarchies, following the NH best-fit from Ref. [63] except for the atmospheric mixing angle and the CP violating phase, arbitrarily set to $\sin^2\theta_{23} = 0.5$ and $\delta_{\mathrm{CP}} = 0$.

8.1.1 Impact of the detector effects

In order to give an idea of the relative importance of the various detector effects in reducing the NMH statistical signal, the χ^2 sensitivity indicator has been computed channel-by-channel, including them in successive stages. The results[2] are reported in Tab. 8.1, assuming an exposure of 1 Mt.y throughout. In addition, the expected distributions of detected events in the NH case and the corresponding maps of the χ^2_s indicator are shown as functions of the true neutrino energy and zenith angle, for the electron channel in Fig. 8.1 and for the muon channel in Fig. 8.2. In the following the succesive stages of Tab. 8.1 are discussed.

[1] For instance the relative event difference $(n_{\mathrm{NH}} - n_{\mathrm{IH}})/n_{\mathrm{NH/IH}}$ fails in doing so and can be rather misleading when both low and high statistics bins are present; the statistical asymmetry $(n_{\mathrm{NH}} - n_{\mathrm{IH}})/\sqrt{n_{\mathrm{NH/IH}}}$ is better suited but tends to overevaluate the contribution of lower statistics bins as well. The choice to use a χ^2 rather than an asymmetry is somewhat important, while the use of a Poisson χ^2 (rather than gaussian) and the average of true NH and true IH have little impact.

[2] Note that the extremely high χ^2 values in the idealised cases are only meant as indicative of the order of magnitude. They should not be taken as precise estimates since at these extreme levels they depend heavily on the choice of binning, oscillation parameters, *etc.*

Channel	Interacting		Detected		Reconstructed		Class.	100% PID	Real PID
ν_e	89	33	50	14	27	4.9	Shower	2.89	0.97
$\overline{\nu}_e$	43		26		16				
ν_μ	$1.2{\cdot}10^3$	$3.3{\cdot}10^2$	$2.7{\cdot}10^2$	63	3.7	0.57	Track	0.57	0.23
$\overline{\nu}_\mu$	$4.5{\cdot}10^2$		$1.3{\cdot}10^2$		2.5				
ν_τ	36	8	9	2	0.29	0.06			
$\overline{\nu}_\tau$	15		4		0.14				

Table 8.1: $|\chi^2|$ indicator of the statistical separation between the event distributions at various stages of the application of the detector response (see text). All distributions are normalised to 1 Mt.yr exposure.

In Tab. 8.1, *interacting events* correspond to a folding of oscillated flux and cross-sections, with no detector effect applied. Distributions of $(\nu+\overline{\nu})$ interacting events have been discussed in Chap. 4. For *detected events*, which distributions are represented in Fig. 8.1 and 8.2, the energy-dependent detection efficiency (effective mass normalised by the instrumented volume) has been applied to the interacting event distributions. The smooth one-dimensional projections of the efficiency shown in Chap. 5, Fig. 5.4 were used. For both interacting and detected events, the case where neutrinos and antineutrinos can be distinguished has been considered separately in the first sub-column.

From the left panel plots in Fig. 8.1 and 8.2 it can be seen that the interaction rate peaks in energy around the region of interest for the NMH. The dependance of the flux in zenith angle is not optimal, though, as most neutrinos reach the detector with trajectories close to the horizon, undergoing no matter-enhanced oscillations. In both figures the normal hierarchy is assumed, thus the matter effect resonance is present in the ν channel and absent in the $\overline{\nu}$ channel. In the IH case (not shown) the opposite is true, therefore the χ^2_s maps show close to symmetrical excess and deficit patterns for ν and $\overline{\nu}$. As discussed in Chap. 2 and 4, a net effect remains after combining the ν and $\overline{\nu}$ channels, mainly due to the larger cross section for neutrinos. The inability to differentiate ν from $\overline{\nu}$ however results in a drastic reduction in overall statistical separation, as can be appreciated from Tab. 8.1.

For the column of *reconstructed events* of Tab. 8.1, the energy-angular smearing has been applied using the full response matrix method. In that case the events are binned into a coarse binning of 20 logarithmic bins in reconstructed E' and 20 bins in reconstructed $\cos\theta'$, using the results of the track reconstructions for $\nu_\mu/\overline{\nu}_\mu$ CC events and the shower reconstruction for all other channels. Perfect particle identification is still assumed at this stage, thus the backgrounds of ν_τ CC and NC events are ignored.

It is informative to compare the way the muon and electron channels are affected successively by the limited detection efficiency and energy-angular resolutions. In the muon channel (see Fig. 8.2), most of the visible NMH signal consists in fast oscillation patterns, *i.e.* small regions of deficit or excess of events which are precisely localised both in energy and zenith, with a rapid alternance of excess and deficit regions. Assuming for instance an energy resolution of $\sim 20\%$ and a zenith angle resolution of $\sim 10°$ (a few % in $\cos\theta_z$), it is easy to imagine how these oscillation patterns may be blurred for the most part and become effectively invisible to the detector. The associated signal is also most intense in the few GeV region where a large fraction of the events go undetected. As a result, the χ^2 reduces by an overall factor of several hundreds with the combined effect of the poor detection efficiency and of the limited $(E,\cos\theta)$ resolutions. In contrast the NMH signature in the electron channel (see Fig. 8.1) mostly consists in an excess of events (for the NH case), extending over a large region of the $(E,\cos\theta)$ plane. The total χ^2 is then reduced only by a factor ~ 2 when applying the detection efficiency, then by a factor ~ 3 by applying the energy and angular resolutions.

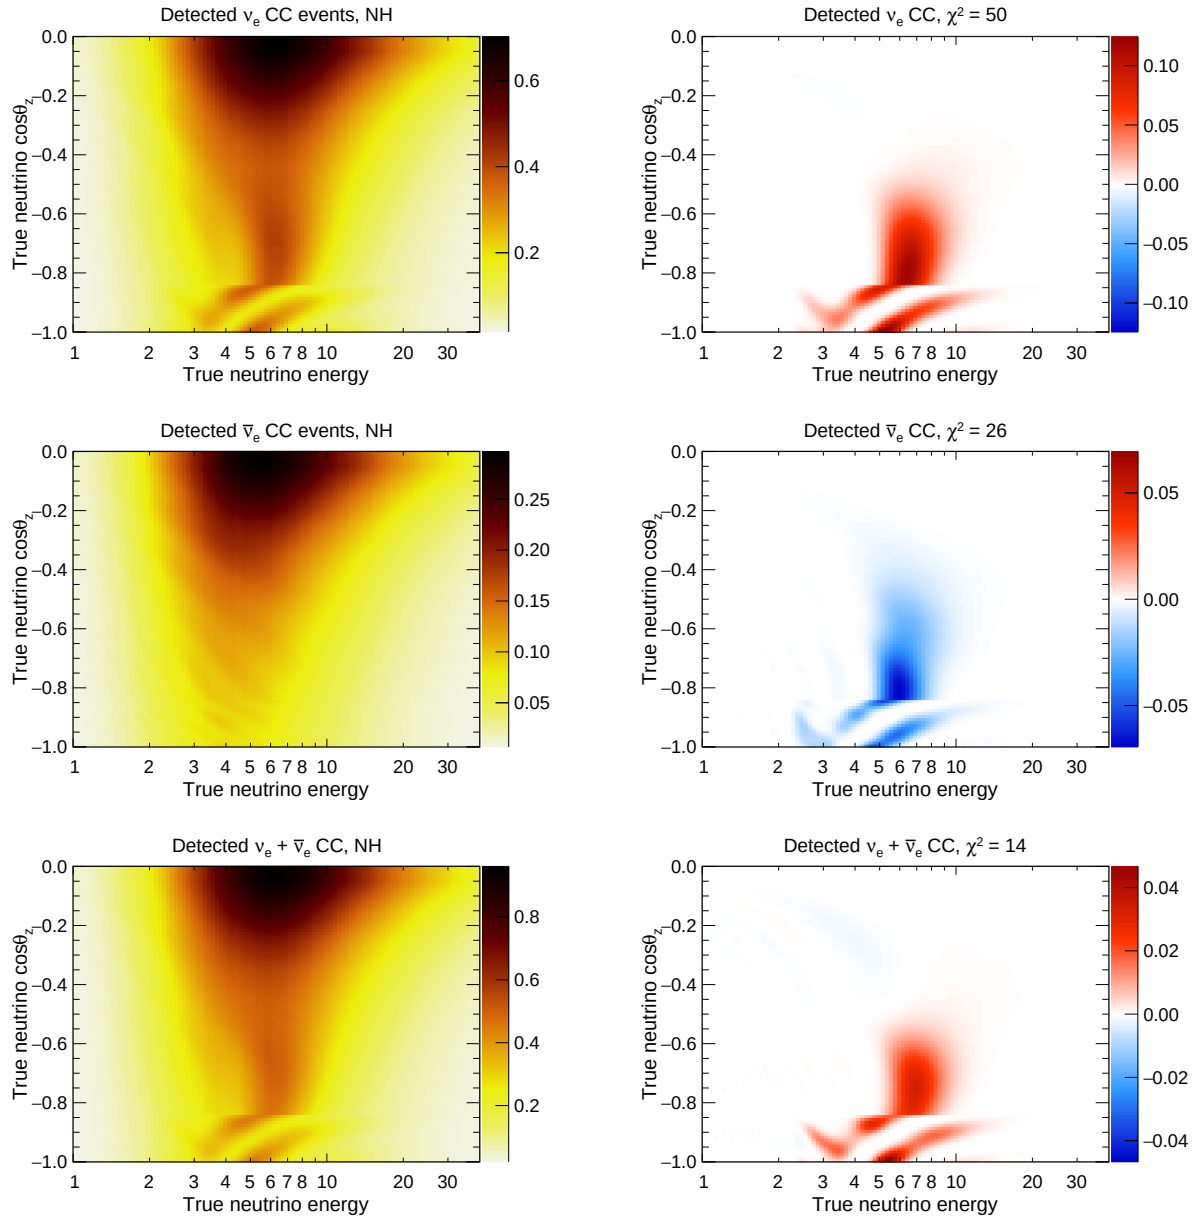

Figure 8.1: Left: expected numbers of **electron** neutrino and antineutrino CC events detected in ORCA as a function of true energy and cosine of the zenith angle. Right: χ^2_s indicator computed between the left distribution and its IH counterpart (not shown). Both the event distributions and χ^2_s distributions are normalised to 1Mt.y exposure. The top and middle panel show the separation of the signal into its neutrino and antineutrino components; the bottom panel shows the signal when neutrinos and antineutrinos are not distinguished, which is the case for ORCA.

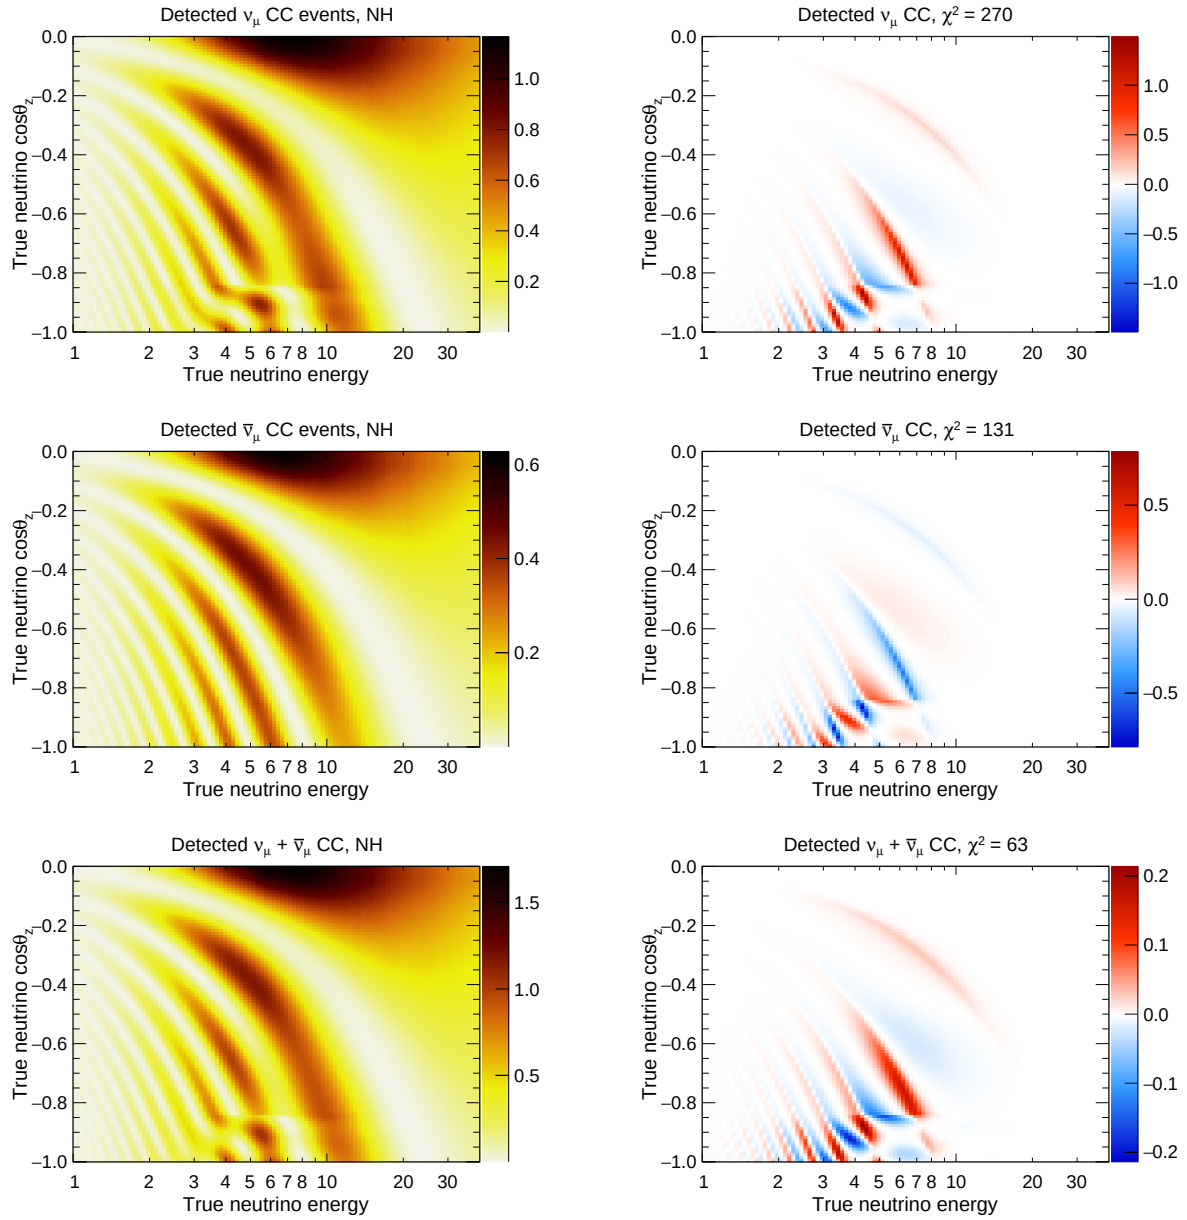

Figure 8.2: Left: expected numbers of **muon** neutrino and antineutrino CC events detected in ORCA as a function of true energy and cosine of the zenith angle. Right: χ^2_s indicator computed between the left distribution and its IH counterpart (not shown). Both the event distributions and χ^2_s distributions are normalised to 1Mt.y exposure. The top and middle panel show the separation of the signal into its neutrino and antineutrino components; the bottom panel shows the signal when neutrinos and antineutrinos are not distinguished, which is the case for ORCA.

The last step applied here is flavour identification. The standard two-class analysis strategy is applied. In the column labeled '100% PID', the events are assumed to be classified with perfect accuracy, as track-like for $\nu_\mu/\overline{\nu}_\mu$ CC events and as shower-like for the other channels including $\nu_e/\overline{\nu}_e$ CC, $\nu_\tau/\overline{\nu}_\tau$ CC (for simplification the $\sim 17\%$ of muonic tau decays are assumed to be classified as showers), as well as the background of $\nu/\overline{\nu}$ NC events. While the muon channel is unaffected by this step, the contamination of ν_τ and NC events significantly affects the electron channel. Finally, in the last column the realistic event classification performance is accounted for using the Monte Carlo based response. The cut value on the track/shower classification score is set to $P_{\text{cut}} = 0.6$ (an optimisation of this cut is discussed later). After this stage the final event distributions entering the statistical analysis are obtained.

8.1.2 Final event samples

The event distributions and sensitivity results shown in the rest of the chapter assume an exposure of three-years with a full detector (*i.e.* approximately 8 Mt effective mass). Fig. 8.5 shows the flavour content of the final shower-like and track-like samples as a function of energy. It can be seen that a large fraction of muon neutrino events are misidentified as shower-like (especially at low energy). On the contrary, the track-like class is a fairly pure sample of ν_μ and $\overline{\nu}_\mu$ CC events. This explains that the shower-like χ^2 is more degraded than the track-like one in the last step of Tab. 8.1.

The reconstructed energy-angular distributions for these two samples of events are shown on Fig. 8.3 (showers) and 8.4 (tracks). A good part of the oscillation features visible on the true energy-angular distributions are smeared out and are no longer visible by eye. The remaining separation between the NH and IH expectations is shown by the maps of χ^2_s indicator. As expected, the extended excess of $\nu_e + \overline{\nu}_e$ CC events in NH survives to the energy-angular smearing and contamination by other interaction channels, while most of the fast oscillation patterns in the muon channel are no longer visible. As a result, the overall statistical $|\chi^2|$ is largely dominated by the contribution from the shower channel.

The figures also illustrate the concrete difference between the correlated and uncorrelated smearing methods. Bin-to-bin statistical fluctuations in the event distributions, inherent to the use of a direct Monte Carlo sampling approach, are more clearly visible in the correlated case. Fluctuating signal 'spots' are also visible at low energy in the correlated signal maps, and are interpreted to be mostly due to the sparse Monte Carlo artifacts discussed in Sec. 5.3. These statistical artifacts appear to be fairly efficiently suppressed by the use of the uncorrelated smearing method.

The sparse MC artifacts are understood qualitatively as a consequence of using the *same* low-statistics MC samples to simulate both pseudo-data and model experiment. These effects motivate the use of adapted binning schemes (as coarse as possible without losing physical signal), implying in particular that the reconstructed Bjorken-y information is not included by default in the analysis.

8.1.3 Finite MC effects and correction

In order to evaluate the quantitative impact of sparse MC artifacts, we study the behaviour of the χ^2 when artificially reducing the amount $\mathcal{N}$ of MC events injected into the construction of the response matrix. The result of such a study is shown on Fig. 8.6, comparing the two smearing methods. The characteristic $1/\mathcal{N}$ behaviour, discussed in Sec. 5.3.4 and Appendix B, is observed for both the correlated and uncorrelated smearing when no correction scheme is applied (dotted lines). As expected, the uncorrelated smearing, with a flatter dependence on $\mathcal{N}$, appears to be closer to the effective infinite MC situation. For the solid line curves (labeled 'with MC error'), the Beeston-Barlow light method presented in Sec. 6.4 is applied, the MC error estimates being calculated with the binomial error model introduced in Sec. 5.3.4.

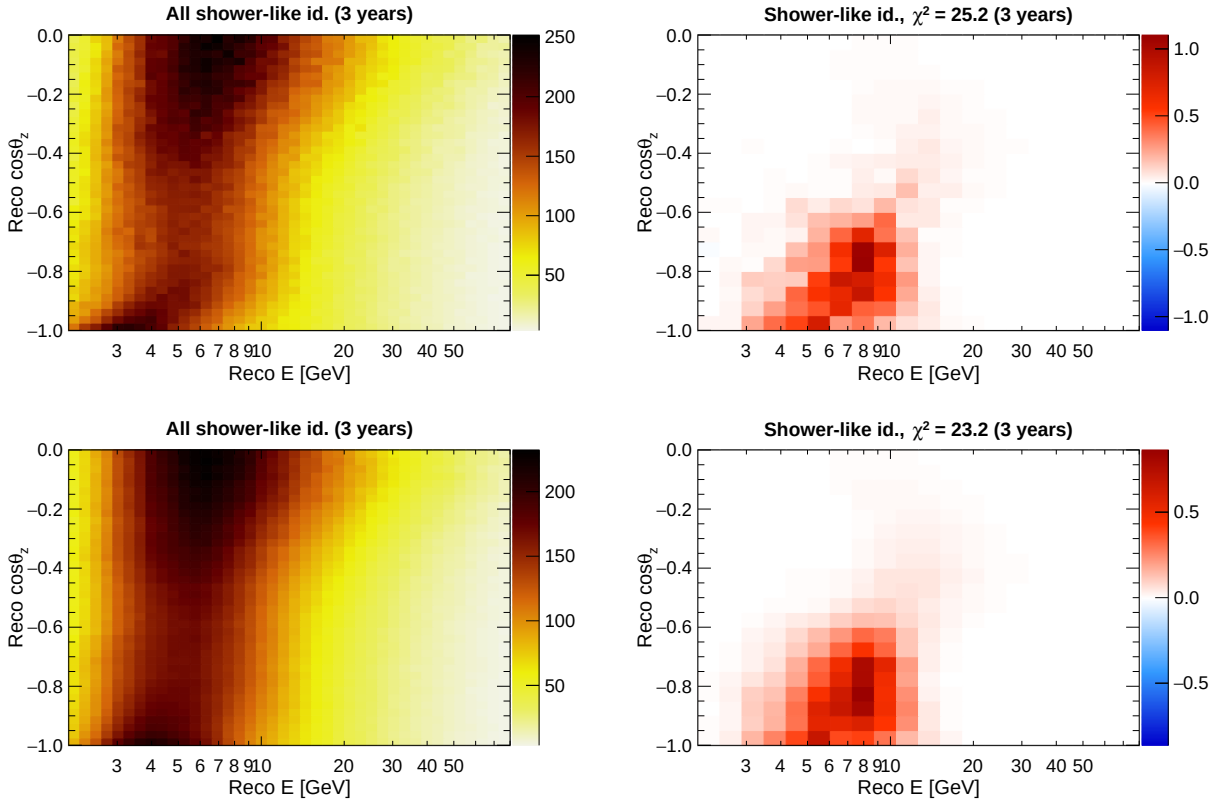

Figure 8.3: Final distributions of selected events classified as shower-like, shown as function of the reconstructed energy and zenith angle. Right: Corresponding map of χ^2_s indicator. The exposure corresponds to 3 years data-taking with a full detector. In the top (bottom) panel the correlated (uncorrelated) smearing method has been used. The event distributions (left) are shown with a finer binning than the one used in the analysis (right) in order for the difference between the two smearing methods to be more visible. Also note that the overall normalisation of both event distributions is identical by design.

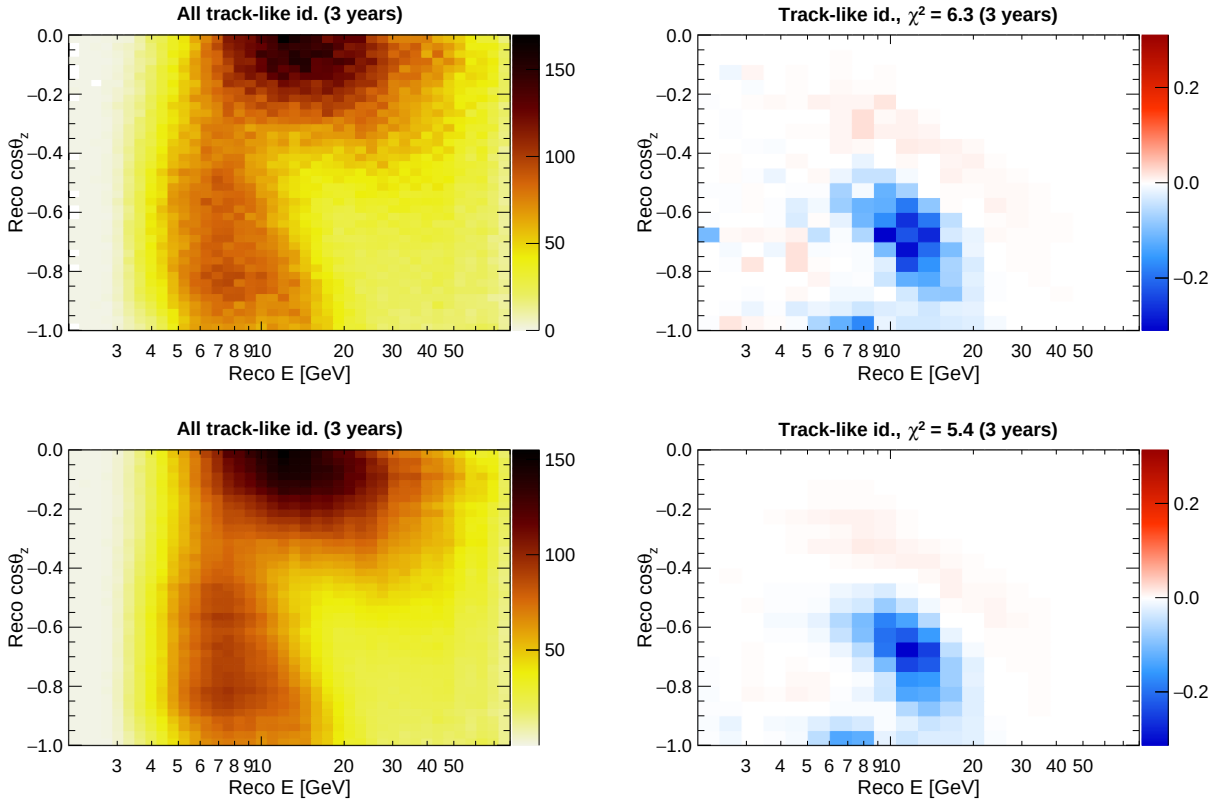

Figure 8.4: Left: final distributions of selected events classified as track-like, shown as function of the reconstructed energy and zenith angle. Right: Corresponding map of χ^2_s indicator. The exposure corresponds to 3 years data-taking with a full detector. In the top (bottom) panel the correlated (uncorrelated) smearing method has been used. The event distributions (left) are shown with a finer binning than the one used in the analysis (right) in order for the difference between the two smearing methods to be more visible. Also note that the overall normalisation of both event distributions is identical by design.

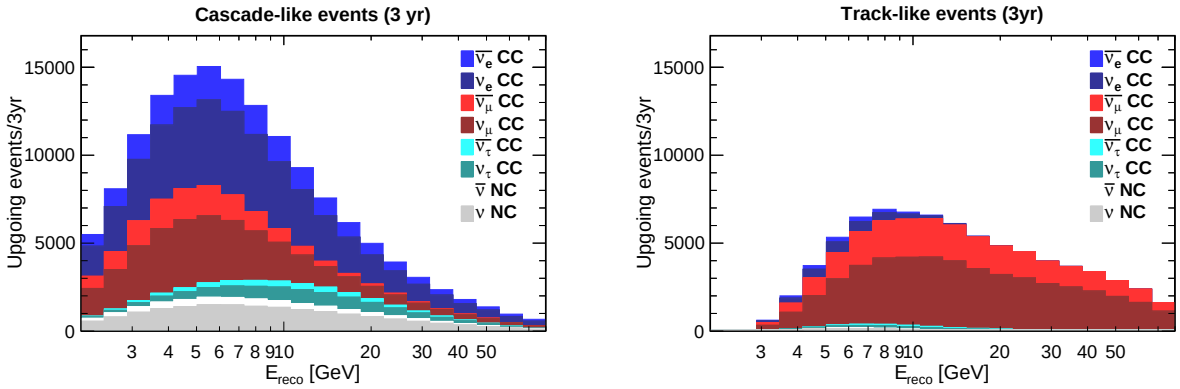

Figure 8.5: Final distributions of events reconstructed as upgoing and classified as shower-like (left) and track-like (right). The event count in the bin is represented by the visible coloured area rather than by the absolute height (stack histograms). The correlated smearing method is used, and the exposure corresponds to 3 years data-taking with a full detector.

Fig. 8.6 establishes the effectiveness of the Beeston-Barlow method as a correction scheme for the sparse MC artifacts: the corrected and non-corrected $\overline{\Delta\chi^2}$ can reasonably be expected to converge towards a common value in the infinite MC limit, and the corrected $\overline{\Delta\chi^2}$ depends much more weakly on MC statistics than the uncorrected one. Although the ideal correction scheme would yield strictly no dependence on MC statistics, this method has the advantage of being conservative in the low statistics limit.

The comparison of the relative impact of the Beeston-Barlow correction in the full statistics case (rightmost point in Fig. 8.6) confirms that statistical artifacts are reduced by the use of the uncorrelated smearing approach. However, extrapolating to the infinite MC limit also suggests a genuinely higher sensitivity when using the correlated smearing, *i.e.* that ignoring the correlation between reconstructed variables results in a non-negligible loss of information on the underlying oscillation phenomenon. The correlated smearing method with MC error correction would then be the most relevant choice to report sensitivity estimates.

While the developed correction scheme appears as very effective on this example, it should be noted that when the MC statistics per bin is lowered, such as when a finer binning or additional binning dimensions are used, the $\overline{\Delta\chi^2}$ may be observed to decrease with $\mathcal{N}$ even with the MC error correction applied. An example of this behaviour will be shown in Sec. 8.4.3. The conservative character of the correction method should therefore not be taken for granted and ideally checked for each analysis case.

The evolution with exposure of the sensitivity estimate with and without MC correction scheme can be seen in Fig. 8.7. Concretely, the increase of exposure in the Asimov analysis simply amounts to an artificial, uniform rescaling of the expected event counts (before minimising with respect to any nuisance parameters), while the MC error estimates stay at the same level. As expected, the correction scheme thus effectively acts as an uncorrelated bin-to-bin systematic uncertainty, whose impact becomes more and more dominant with time as statistical errors in data decrease (a similar effect is found in Ref. [161]). Of course, in a real experiment Monte Carlo keeps being produced in parallel of data-taking so that this situation is unlikely to occur, the dominant uncertainties generally being either statistical errors in data or physical systematics.

8.2 Impact of oscillation parameters and their uncertainties

This section examines how the NMH sensitivity is affected by the uncertainties on the neutrino oscillation parameters and Earth model. These studies are performed using the Asimov method described in Sec. 6.2. The atmospheric mass-squared splitting has a special status in a

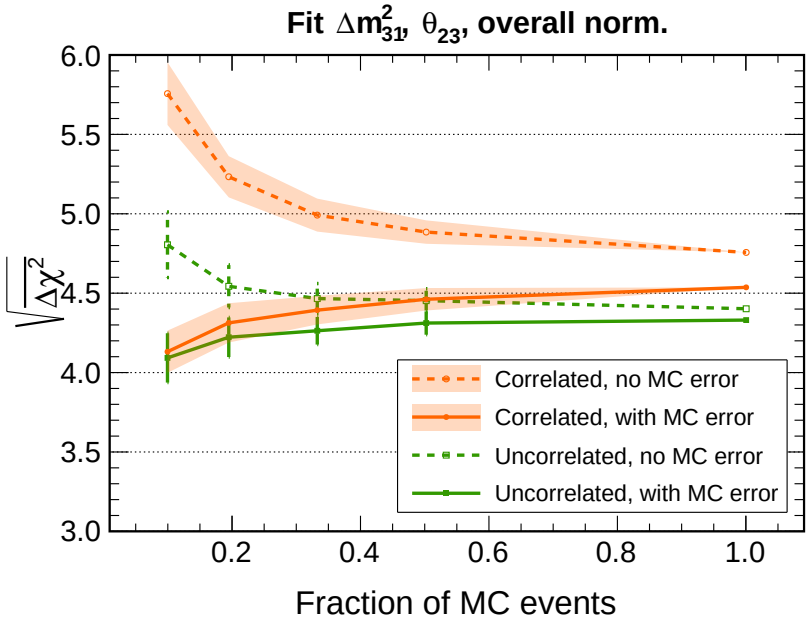

Figure 8.6: Dependence of $\sqrt{\Delta\chi^2}$ with the amount of MC events included to construct the response matrix. Each point corresponds to the average $\sqrt{\Delta\chi^2}$ obtained over 100 randomly drawn sub-samples with size 1/2, 1/3, 1/5, and 1/10 of the whole available MC sample. The errors bars (bands) indicate the sample standard deviation for each point. In addition to the purely statistical NH/IH separation, the sensitivity estimate shown here includes three of the most important sources of systematic uncertainty as nuisance parameters (thus anticipating on the next section).

mass hierarchy analysis, and the impact of its absolute value in the fit is discussed first. It is then shown how the degeneracy between the mass hierarchy and the atmospheric mixing angle θ_{23} strongly affects ORCA's sensitivity. The impact of the unknown CP phase δ_{CP} is also non-negligible, while the reactor mixing angle θ_{13} and solar parameters have little to no impact within their current uncertainties. Earth electron density parameters are also, to some extent, degenerate with the mass hierarchy. Several levels of prior constraints are considered, and it is shown that the assumption that the radial electron density profile is known at the sub-percent level makes the impact of these uncertainties close to negligible.

8.2.1 Atmospheric mass splitting

As a first consistency check for this analysis, the "statistics-only" mass hierarchy sensitivity (*i.e.* without including any systematic uncertainty) was calculated using the 'LoI' MC sample to model the detector response, to be compared with previous independent results. The sensitivity is evaluated using the Asimov sets as, assuming for instance a true normal hierarchy:

$$S_{\mathrm{NH}} = \sqrt{\chi^2(\mathrm{NH}|\mathrm{IH})} = \sqrt{\chi^2_{\mathrm{tr}}(\mathrm{NH}|\mathrm{IH}) + \chi^2_{\mathrm{sh}}(\mathrm{NH}|\mathrm{IH})}, \tag{8.4}$$

where χ^2_{tr} and χ^2_{sh} are the contributions from the track-like and shower-like channels. A large disagreement was initially found in the relative importance of these contributions. This turned out to be due to a different definition of the transformation of the atmospheric mass splittings Δm^2_{31}, Δm^2_{32} under the 'hierarchy flip' NH $\leftrightarrow$ IH.

Fig. 8.8 shows the statistics-only sensitivity as a function of the choice of Δm^2_{31} in the alternative hierarchy hypothesis. The pseudo-data is generated with a fixed true value Δm^2_{31} denoted as Δ_{true}, and the sign-flipped value $\Delta_{\mathrm{test}} = -\Delta_{\mathrm{true}} + \delta_{\mathrm{shift}}$ for the model hypothesis spans the 3σ allowed range from Ref. [63]. It is clear from the figure that the contribution from

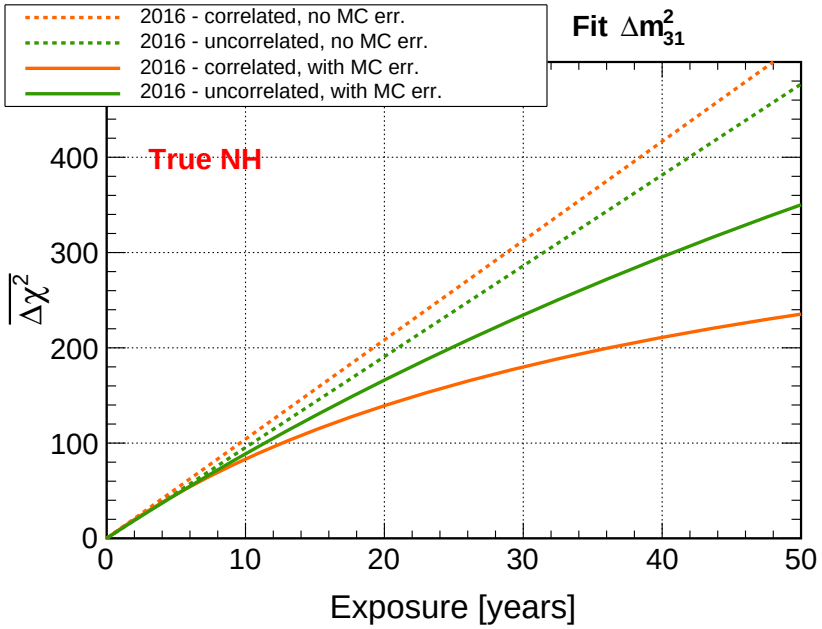

Figure 8.7: Increase with exposure of the Asimov $\overline{\Delta\chi^2}$ for NMH measurement, with and without MC uncertainties. The $\Delta\chi^2$ is computed with a single fitted parameter (Δm^2_{31}) to avoid known artifacts in the relative statistical contributions of the track and shower channels (see Sec. 8.2.1).

the track channel is very sensitive to the choice of Δm^2_{31} in the wrong hierarchy assumption, while the shower channel is essentially insensitive to it. This reflects the fact that in the 20–80 GeV energy range the track channel is a very pure sample of $\nu_\mu/\overline{\nu}_\mu$ CC events, and (anti)muon neutrino disappearance in that energy range is directly and unambiguously related to $|\Delta m^2_{31}|$ and $\sin^2 2\theta_{23}$. In contrast, in the region dominated by matter effects and $\nu_\mu \leftrightarrow \nu_e$ oscillations the measurement of $|\Delta m^2_{31}|$ is somewhat degenerate with the mass hierarchy. These observations match the expectations from the discussion of Chap. 2 (see Fig. 2.19).

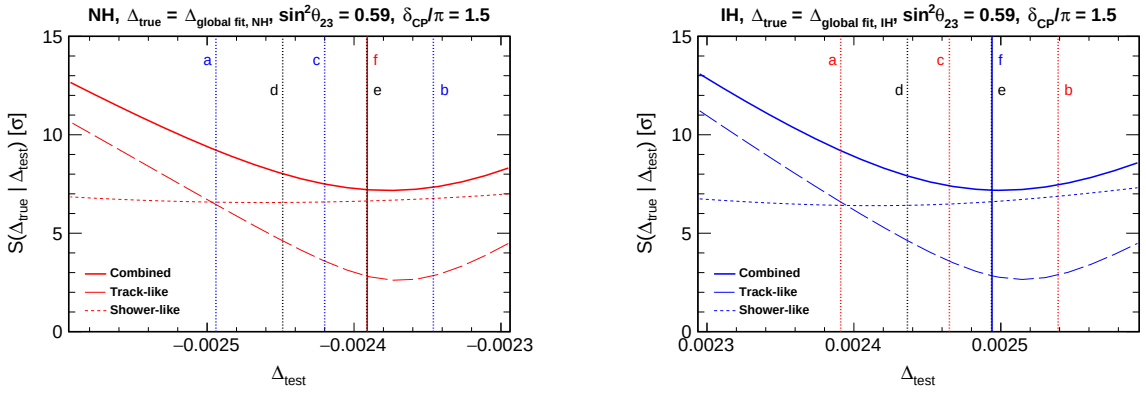

Figure 8.8: Effect of the choice of Δm^2_{31} (denoted as Δ_{test}) in the alternative MH hypothesis. In the left (right) plot, the truth (pseudo-data) is generated fixing Δ to the global best-fit value assuming a normal (inverted) hierarchy. The total sensitivity $S = \sqrt{\chi^2_{\text{tr}} + \chi^2_{\text{sh}}}$ is shown as a plain line and the contributions $\sqrt{\chi^2_{\text{tr}}}$ and $\sqrt{\chi^2_{\text{sh}}}$ of the track and shower channels are shown as thinner dashed lines. The vertical lines indicate special values corresponding to possible conventions for the transformation $\Delta \to \Delta'$ under the MH flip (see text).

The vertical lines in Fig. 8.8 indicate special values corresponding to possible conventions for the transformation $\Delta \to \Delta'$ under the MH flip. A large overestimation of the NMH sensitivity

in the track-like channel can be obtained when a naive definition of the hierarchy flip is used. For instance it can be defined arbitrarily as

- an inversion of the sign of Δm^2_{31}, yielding $\Delta' = -\Delta$ (choice (a)),
- an inversion of the sign of Δm^2_{32}, yielding $\Delta' = -\Delta + 2\delta_{\rm sol}$ (choice (b)).

There is no reason to prefer one over the other, and for a symmetrical treatment an alternative is to flip the sign of

$$\Delta M^2 = \frac{\Delta m^2_{31} + \Delta m^2_{32}}{2}, \tag{8.5}$$

corresponding to choice (c) in the figure with $\Delta' = -\Delta + \delta_{\rm sol}$, implying $|\Delta m^2_{32}({\rm IH})| = |\Delta m^2_{31}({\rm NH})|$ and conversely.

More physically motivated choices arise from phenomenological considerations. With three neutrino flavours, disappearance probabilities $P_{\alpha\alpha}$ *in vacuum* are in principle sensitive to the mass hierarchy since oscillations driven by Δm^2_{31} and Δm^2_{32} coexist. Since $\Delta m^2_{31} \approx \Delta m^2_{32}$, though, $P_{\alpha\alpha}$ is well-described by a two-flavour approximation with an effective mixing angle $\Delta_{{\rm eff},\alpha}$ which depends on the considered flavour α [273] [274]. At the L/E range relevant to ORCA, vacuum oscillations are then blind to the NMH in the disappearance channel, in the sense that $P_{\alpha\alpha}$ is unchanged under $\Delta_{{\rm eff},\alpha} \to -\Delta_{{\rm eff},\alpha}$. Using the effective mass splitting describing ν_e disappearance [273], one obtains[34]

$$\Delta' = -\Delta + 2\sin^2\theta_{12}\delta_{\rm sol} \approx -\Delta + 0.6 \cdot \delta_{\rm sol}, \tag{8.6}$$

corresponding to choice (d) in the figure. Using instead the description for ν_μ disappearance leads to choice (e):

$$\Delta' = -\Delta + 2(\cos^2\theta_{12} - \cos\delta_{\rm CP}\sin\theta_{13}\sin 2\theta_{12}\tan\theta_{23})\delta_{\rm sol} = -\Delta + x \cdot \delta_{\rm sol} \tag{8.7}$$

where x ranges approximately from 1 to 1.7 depending on θ_{23} and $\delta_{\rm CP}$. Finally, the global best-fit value[63] for the alternative MH is indicated with the line labeled as (f). Interestingly, the lines (e) and (f) are indistinguishable, meaning that the shift $\Delta m^2_{31}({\rm NH}) - |\Delta m^2_{31}({\rm IH})|$ in the global fit values is perfectly described[5] by the use of Eq. 8.7.

In this thesis, in any result including nuisance parameters Δm^2_{31} is fitted without constraint. In those cases, a choice of convention for the hierarchy flip is not needed. Instead the minimisation in the wrong hierarchy will typically converge to the minimum observed on the combined sensitivity curves of Fig. 8.8. When no nuisance parameters are included, the hierarchy-specific global fit values (f) are used, as this choice appears to be consistently close to the best-fit value observed by ORCA.

We found that Eq. 8.7 does not always describe ORCA's measurement of Δm^2_{31} as accurately as in the case shown in Fig. 8.8, depending on the choice of true values for θ_{23} and $\delta_{\rm CP}$. In fact, the wrong hierarchy best-fit Δm^2_{31} is relatively insensitive to these values, showing much smaller variations than expected from Eq. 8.7. Although most of the sensitivity to $|\Delta m^2_{31}|$ lies in the track channel where the rate of ν_μ events is measured, there is a non-negligible contribution from $\nu_e \to \nu_\mu$, which may contribute to explain this observation. This result is in agreement with studies found in the literature [155] [69].

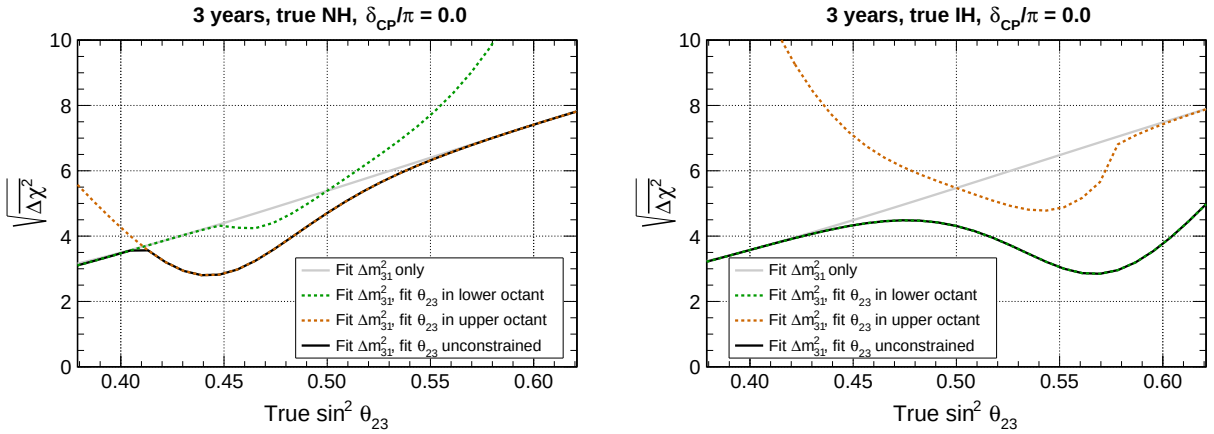

Figure 8.9: Asimov-based NMH sensitivity as a function of the true value of $\sin^2\theta_{23}$, for true Normal hierarchy in the left panel, true inverted hierarchy in the right panel.

8.2.2 θ_{23} and the octant degeneracy

The degeneracy in the current knowledge of neutrino mixing regarding the octant of θ_{23} (if non-maximal) has already been mentioned in Chap. 1, and the impact of this uncertainty on oscillation probabilities relevant to ORCA has been illustrated in Chap. 7. Fig. 8.9 shows the Asimov-based NMH sensitivity with respect to the true value of θ_{23}, with and without allowing θ_{23} to vary in the fit. Several important observations can be made. First, consider the case where θ_{23} is assumed to be known perfectly (grey curve). The NMH sensitivity then steadily increases with $\sin^2\theta_{23}$, ranging from $\sim 3.5\sigma$ to almost 8σ. This is explained by the fact that $\sin^2\theta_{23}$ is the amplitude of the leading term in the matter-enhanced $\nu_\mu \leftrightarrow \nu_e$ oscillations. The sensitivity is also found to be nearly identical for true NH and true IH at this stage.

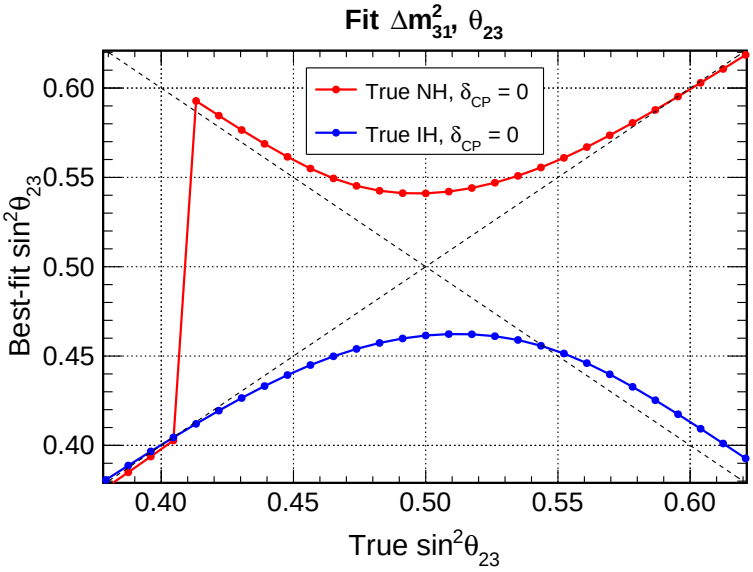

Figure 8.10: Best-fit value of $\sin^2\theta_{23}$ in the wrong hierarchy, as function of the true value of $\sin^2\theta_{23}$. The color refers to the true hierarchy. Diagonal lines correspond respectively to the identity and to a symmetrical octant switch $\theta'_{23} = \pi/2 - \theta_{23}$. The interval considered for true $\sin^2\theta_{23}$ is consistent with, but slightly larger than the 3σ range ($\sim 0.42 - 0.61$ in NH) currently allowed by the global fit [63].

[4] Equations 16 and 17 in [273] are used respectively for the ν_e disappearance and ν_μ disappearance descriptions.

[5] Note that this agreement occurs for the particular choices of θ_{23} and $\delta_{\rm CP}$ (close to the world best-fit values) used for 8.8, while the agreement is not perfect when using *e.g.* $\delta_{\rm CP} = 0$. The respective importance of accelerator long-baseline data (ν_μ disappearance) and reactor data ($\overline{\nu}_e$ disappearance) is discussed in [62].

The octant-hierarchy degeneracy (discussed in Chap. 2) comes into play once θ_{23} enters as a nuisance parameter. In the NH case, we observe in Fig. 8.9 that for $\sin^2\theta_{23}$ far from maximal mixing, fitting θ_{23} has no impact on the NMH sensitivity, implying that it can be very well constrained independently of the hierarchy. This is confirmed by examining the wrong hierarchy best-fit value of θ_{23}, as shown on Fig. 8.10 as a function of its true value (the red curve corresponds to true NH, fitted IH).

Around $\sin^2\theta_{23} \approx 0.41$, the NH data starts to be better fitted in the IH assumption by placing θ_{23} in the upper octant. This results in the slope discontinuity seen in the left panel of Fig. 8.9, and in a large sensitivity loss of up to almost 2σ with respect to the Δm^2_{31}-only fit. In contrast, when the inverted hierarchy is true and the normal hierarchy is fitted, a value of θ_{23} in the lower octant is preferred whatever its true value. As seen on the right panel of Fig. 8.9, this has a drastic impact on the sensitivity for a combination of true IH and second octant.

At this stage the knowledge of the octant alone would greatly reduce the impact of the θ_{23} uncertainty for the NMH determination, even more so if θ_{23} is far from maximal mixing. This is visible in Fig. 8.9 by the superimposed dotted lines for which θ_{23} is restrained to the upper or lower octant in the minimisation.

A dependence on the true value of δ_{CP} is also found. Its impact is actually more important than the one of the true hierarchy when θ_{23} is fixed, suggesting that the addition of δ_{CP} as a nuisance parameter may have a sizeable effect. This is studied in the next section. Regarding Δm^2_{31} we find that its best-fit value has a minor dependence (less than $\pm 10^{-6}\,\mathrm{eV}^2$) on both θ_{23} and δ_{CP}. This dependence does not match the one expected from Eq. 8.7, either in scale or functional form.

Interplay of the track and shower channels

With respect to the case where only Δm^2_{31} is fitted, it is interesting to consider the interplay of the track-like and shower-like channels in the presence of the uncertainty on θ_{23}. As discussed in Sec. 2.2.3, the oscillation probabilities in the $\nu_\mu \leftrightarrow \nu_e$ oscillation channel show a degeneracy between the NMH and θ_{23}, and are primarily sensitive to $\sin^2\theta_{23}$, while the $\nu_\mu \to \nu_\mu$ channel is less affected by this degeneracy, the dominant term being controlled by $\sin^2 2\theta_{23}$.

This expected behaviour can be verified on Fig. 8.11. The dependence of

$$\chi^2 = \chi^2_{\mathrm{tr}} + \chi^2_{\mathrm{sh}} \tag{8.8}$$

on the test value of θ_{23} in the wrong hierarchy fit is shown for both true hierarchies and two different true θ_{23} values: maximal mixing, and upper octant $\sin^2\theta_{23} = 0.55$. The minimum values of the χ^2_{tr} (red) and χ^2_{sh} (blue) graphs thus correspond respectively to the NMH sensitivities (with $S = \sqrt{\chi^2}$) that would be obtained with a track-only fit and a shower-only fit, while the result of the combined fit is shown in black.

It can be seen that the shower-only fit systematically favours a value of $\sin^2\theta_{23}$ off from the correct one by a large amount. As can be understood easily by re-examining Fig. 2.18, the best-fit $\sin^2\theta_{23}$ is higher than the truth for {true NH, fit IH} and lower for {true IH, fit NH}. In the track channel, the best-fit $\sin^2\theta_{23}$ is much closer to the true value, although affected by the octant degeneracy (most visibly in the combination {true IH, $\sin^2\theta_{23} = 0.55$}). Most importantly, the best-fit $\sin^2\theta_{23}$ found in the shower-only fit is generally strongly disfavoured by the track-only fit, resulting in a significant increase of sensitivity when the two channels are combined.

Therefore, although the quantitative contribution of the track channel χ^2_{tr} to the total χ^2 may appear as subdominant in the statistics-only sensitivity or when considering the result of the combined fit, the track channel plays a crucial part in reducing the impact of the NMH-θ_{23} degeneracy.

In spite of this partial removal of the degeneracy, the wrong-hierarchy best-fit $\sin^2\theta_{23}$ is still found in the wrong octant in the case {true IH, $\sin^2\theta_{23}=0.55$} (bottom right panel of Fig. 8.11), where the lower octant is favoured by both channels. The case {true NH, $\sin^2\theta_{23}=0.55$} is more favorable, as can also be deduced from Fig. 2.18.

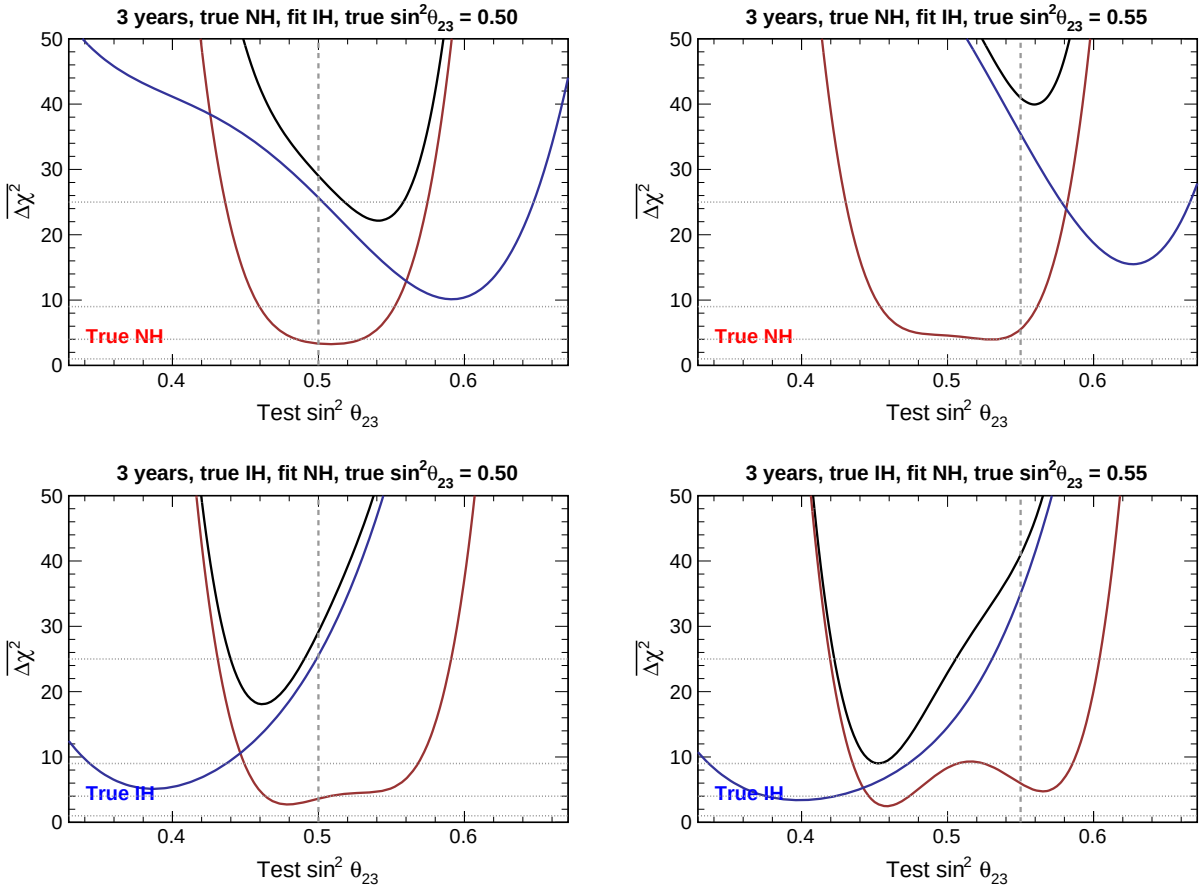

Figure 8.11: Interplay of the track and shower channels for the NMH measurement in the presence of θ_{23} as a free nuisance parameter. The red lines correspond to the track channel $\chi^2_{\rm tr}$, the blue lines correspond to the the shower channel $\chi^2_{\rm sh}$, and the black line is the combination $\chi^2_{\rm tot}=\chi^2_{\rm tr}+\chi^2_{\rm sh}$. The values of $\overline{\Delta\chi^2}$ corresponding to 1σ, 2σ, 3σ and 5σ mass hierarchy sensitivity are indicated as horizontal grey lines and the vertical grey line indicates the true θ_{23} value.

8.2.3 CP phase, reactor mixing angle and solar parameters

To further quantify the impact of oscillation uncertainties, $\delta_{\rm CP}$ is allowed to vary freely in the wrong hierarchy fit, in addition to the previous parameters. As visible from Fig. 8.12 the χ^2 minimisation tends to prefer values in the interval $[\pi/2-3\pi/2]$ for which $\cos\delta_{\rm CP}<0$. The best-fit pattern is mostly independent of the true value, while it has a clear dependence on the true hierarchy and in a lesser extent on θ_{23} – notably reflecting the octant jump in true NH around $\sin^2\theta_{23}\approx 0.41$. The fit generally prefers $\sin\delta_{\rm CP}<0$, except when the true value of θ_{23} is far from maximal mixing and the best fit lies in the wrong octant.

The loss of NMH sensitivity resulting from the freedom of $\delta_{\rm CP}$ in the fit can be as large as $\sim 1\sigma$, depending mostly on the true hierarchy and true value of $\delta_{\rm CP}$. This is shown on Fig. 8.13 for both true NH and IH in the maximal mixing case. As expected from Fig. 8.12, the effect is the largest when the true value is very different from π, which tends to reduce the initial dependence of the sensitivity on true $\delta_{\rm CP}$.

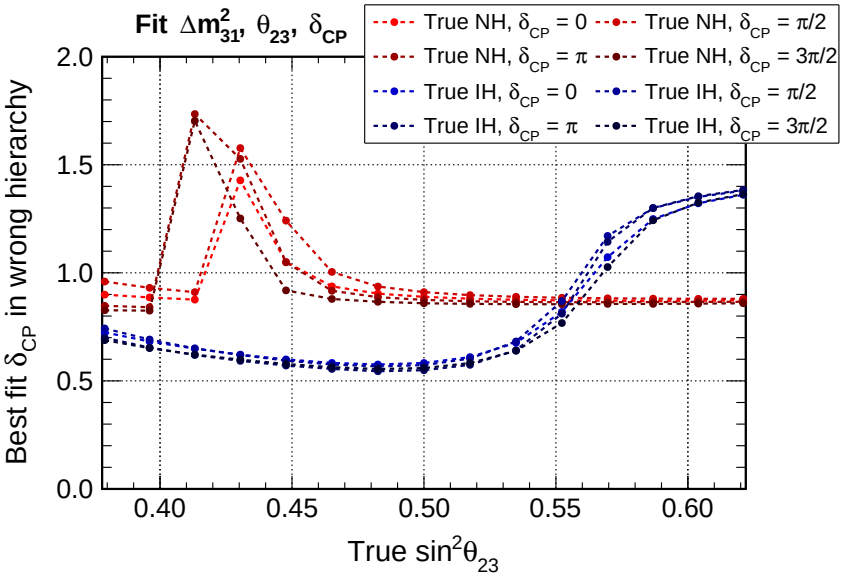

Figure 8.12: Best-fit value of δ_{CP} in the wrong hierarchy, as function of the true value of $\sin^2\theta_{23}$.

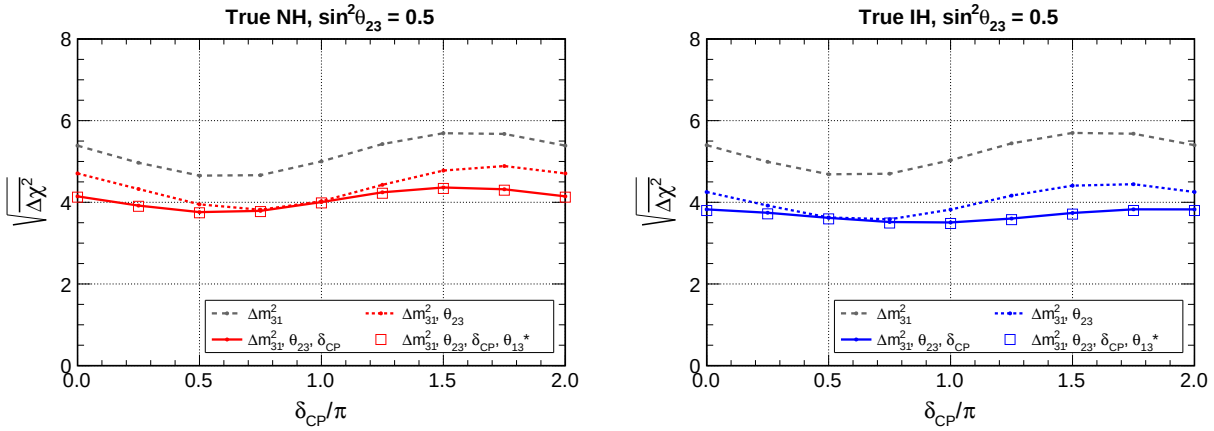

Figure 8.13: Mass hierarchy sensitivity as a function of the true value of δ_{CP}, assuming true maximal mixing and a normal hierarchy (left) or inverted hierarchy (right). The legend indicates the nuisance parameters included in the fit. For θ_{13} a gaussian prior is applied, following [63].

As discussed in Chap. 2, the precise knowledge of the reactor mixing angle θ_{13} is crucial for the NMH measurement with ORCA because this parameter is directly involved in the amplitude of the matter effect resonance. The measurement of θ_{13} by reactor and accelerator experiments is however much more precise than what can be achieved from ORCA's oscillation measurement alone, as illustrated on Fig. 8.14. Therefore, θ_{13} is included in the fit accounting for this external information with a gaussian prior constraint of width 0.15°, following the global fit [63].

As can be seen in Fig. 8.13, for the NMH sensitivity at this stage this is close to equivalent to keeping θ_{13} fixed, *i.e.* the prior constrains θ_{13} tightly enough to remove the potential degeneracies with the other fitted parameters – the patterns seen in Fig. 8.10 and Fig. 8.12 are also unchanged. Only a very small decrease in sensitivity is observed in the region where the inital significance is the highest (NH second octant). The fitted value of θ_{13} typically deviates from the prior mean by about 10^{-3} rad ($\sim 0.05°$) or less, which corresponds to a penalty of about 0.1 in units of $\Delta\chi^2$. The deviation is observed towards larger θ_{13} for true NH (fitting IH), which is consistent with Fig. 8.14. For true IH the opposite is observed.

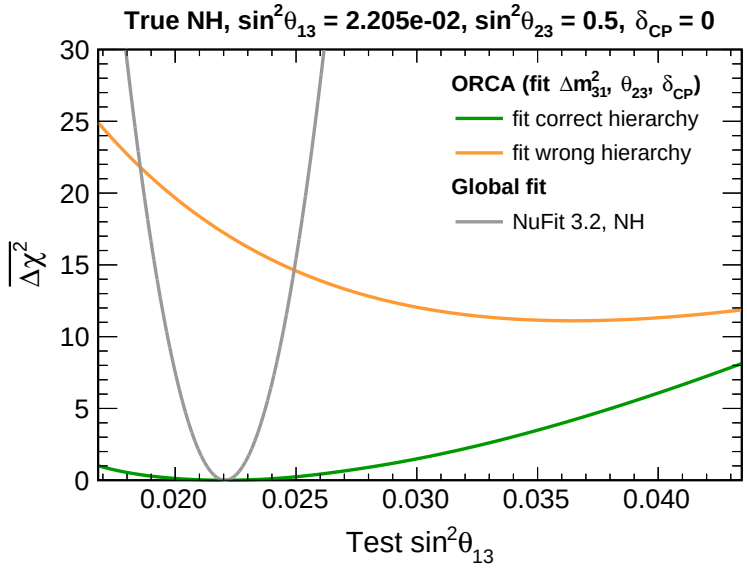

Figure 8.14: Impact of θ_{13} as a systematic for the NMH measurement with ORCA. The loss of NMH sensitivity in the absence of external constraint on θ_{13} can be appreciated by the difference in the wrong hierarchy fit between the value of $\Delta\chi^2$ close to true θ_{13} and its minimum value.

8.2.4 Earth model uncertainties

To assess the impact of uncertainties on the inner Earth density or chemical composition, we let the Z/A value vary in the fit, independently for each of the three innermost chemical layers defined in Chap. 4 (lower mantle, outer core and inner core), but uniformly in each layer. Although the electron density is generally assumed to be known from geophysics with a precision of a few percent or better, we are here interested in evaluating the possibility for ORCA to measure the NMH without such strict external constraints. Three different levels of constraint have therefore been evaluated by appling gaussian priors of width 0.15, 0.05, and 0.02 in Z/A. These correspond to relative constraints of about 30%, 10% and 4% respectively. In each case the priors are independent but the same width is applied to each of the three Z/A parameters. In addition the unconstrained case has been tried. The resulting sensitivities are shown on Fig. 8.15 for both true hierarchies.

The Z/A uncertainties are found to have a modest impact in the normal hierarchy case. The inspection of the Z/A values obtained for each layer in the unconstrained fit reveals that ORCA has a negligible sensitivity to the inner core density, as is expected considering that the corresponding solid angle is only about 2% (see Chap. 2, Tab. 2.1). In the case where NH is true and IH fitted, the unconstrained fit favours a high electron density in the outer core ($Z/A \sim 0.52$) and a lower density in the lower mantle ($Z/A \sim 0.43$). When NH is true and IH fitted, conversely the minimisation finds a low density core and high density lower mantle (with a large dependence on θ_{23}).

Note that here only uniform variations of Z/A in a layer were applied, while the overall radius and density profile of each chemical layer was kept fixed. As explained in Chap. 2, small-scale shape variations have little impact on oscillation probabilities. They were also demonstrated to be negligible in the context of NMH studies, for instance in Ref. [155].

8.3 Impact of systematics

In this section, the individual impact of each modeled systematic effect are assessed separately. This approach is motivated by the observation that when systematics and priors are

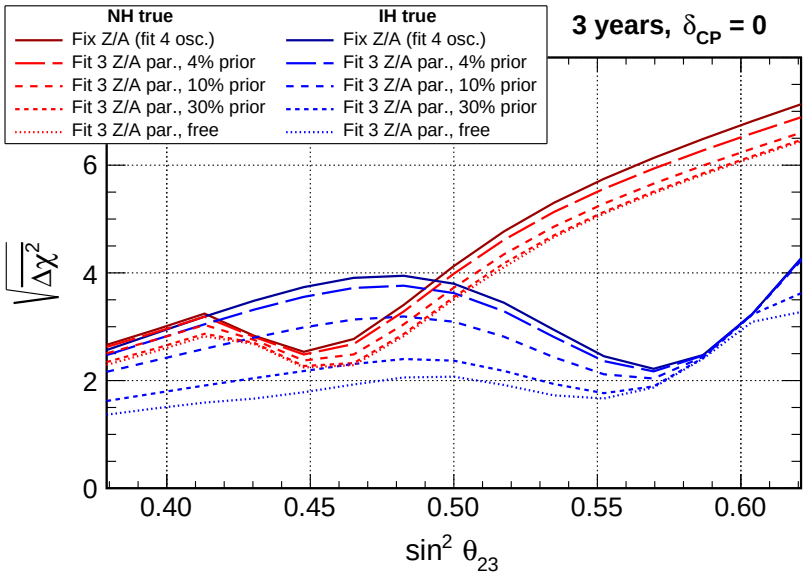

Figure 8.15: Asimov-based NMH sensitivity as a function of the true value of $\sin^2\theta_{23}$, for true NH in red, true IH in blue. For all curves the previous 4 oscillation parameters are fitted: Δm^2_{31}, θ_{23}, $\delta_{\rm CP}$, θ_{13} with prior.

added incrementally to the set of nuisance parameters, their apparent impact is very much dependent on the order of addition, so that the evaluation of their relative effects is difficult.

The starting point is the set-up where the oscillation parameters studied in the last section are all fitted, with the exception of Earth parameters. The normalisation parameters are studied in Sec. 8.3.1 and the flux composition and shape parameters in Sec. 8.3.2.

To some extent, the results of this section are a work in progress, and may be completed in the near future.

8.3.1 Normalisation parameters

The following normalisation systematics have been studied:

- Overall normalisation parameter $f_{\rm norm}$
- Separate normalisation parameters for the track-like and cascade-like channels (replacing $f_{\rm norm}$)
- NC events normalisation
- ν_τ CC events normalisation

There are intentionnally no plain normalisation parameters affecting specifically ν_μ and ν_e events, be it based on the oscillated or unoscillated flavour. As explained in Chap. 7, the flux composition uncertainties are taken into account by the combination of an overall normalisation parameter and skew parameters related to flavour and neutrino/antineutrino ratios (studied in the next section). Additional normalisations based on the oscillated flavour may be added in the near future to model uncertainties on cross-sections or flavour-dependent detection efficiencies.

The results are shown as sensitivity curves with dependence on the true value of θ_{23} on Fig. 8.16. The parameter set used as a basis, labeled as 'set A' in the caption, includes a fit to Δm^2_{31}, θ_{23} and $\delta_{\rm CP}$ with no external constraint, plus θ_{13} with its standard prior. Additional parameters are added individually on top of this set (without combination). In the case of the NC and ν_τ CC normalisation parameters, their impact with or without prior (10%) has been evaluated. The most important observations are the following:

- normalisation parameters have close to no impact for true θ_{23} values where the NMH-θ_{23} degeneracy already strongly degrades the sensitivity, but they significantly degrade the sensitivity for the best-performing true θ_{23} values;
- the overall normalisation has the least impact, as it can be well constrained in control sample regions (typically in the saturation regime of oscillations and for close to horizontal trajectories);
- on the contrary, distinct track and shower normalisations have the largest impact;
- the inclusion of 10% priors on the NC and ν_τ CC normalisations barely affects their impact.

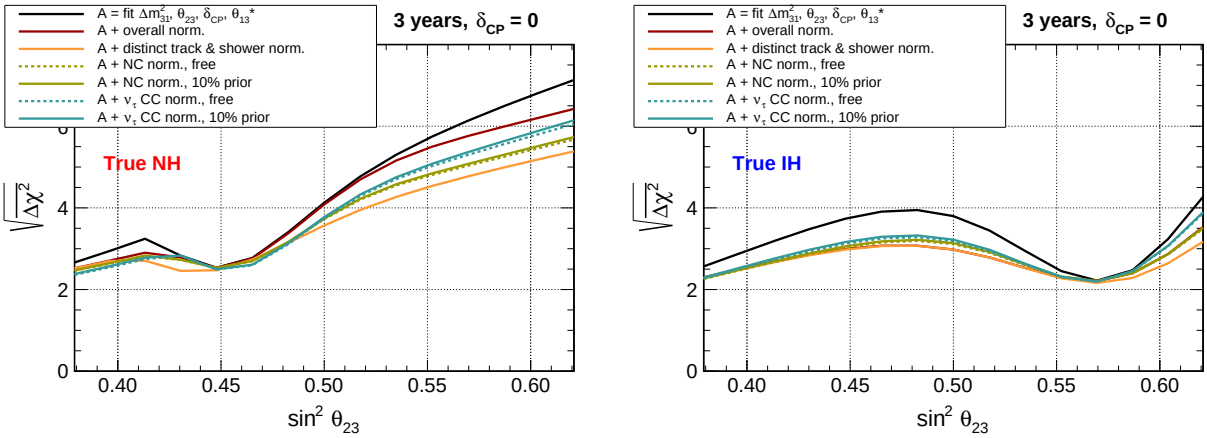

Figure 8.16: Impact of individual normalisation uncertainties on the Asimov-based sensitivity. Each curve corresponds to a single normalisation parameter being fitted, with the exception of the track and shower normalisations (2 free parameters).

8.3.2 Flux composition and shape

The comparison of the individual impact of flux composition parameters (skews) is shown on Fig. 8.17. The basis nuisance parameter set, labeled as 'set B', includes the overall normalisation (in addition to the oscillation parameters of set A), so that the flavour-specific normalisations are effectively free to take arbitrary values. The effect of including priors (loosely based on the constraints discussed in Chap. 7) is evaluated.

As in the previous section, a 'saturation' of the impact of systematics is observed when dominated by the θ_{23} degeneracy. While the flavour skew and $\nu_\mu/\overline{\nu}_\mu$ skews have comparable effects in both true hierarchy cases, the $\nu_e/\overline{\nu}_e$ skew has a very strong impact for true NH (and milder impact for true IH). The addition of reasonable priors on the $\nu/\overline{\nu}$ skews appears to be effective (especially for $\nu_e/\overline{\nu}_e$), but has essentially no impact for the flavour skew.

The strong impact of the $\nu/\overline{\nu}$ skews is not surprising, considering that the NMH sensitivity essentially comes from the determination of whether $\nu_\mu \to \nu_e$ appearance (NH) or $\overline{\nu}_\mu \to \overline{\nu}_e$ appearance (IH) is enhanced in the MSW resonant region, without being able to distinguish ν and $\overline{\nu}$. However, the reason for the observed discrepancy in the impact of the $\nu_e/\overline{\nu}_e$ skew for true NH and IH remains unclear and would deserve further investigation. As to the flavour ratio, the shower channel is strongly contaminated by $\nu_\mu/\overline{\nu}_\mu$ events and can hardly constrain the flavour ratio by itself, whereas the track channel, as a quite pure $\nu_\mu/\overline{\nu}_\mu$ sample, should be more helpful. It can thus be conjectured that the combined fit is again profitable here.

Finally, flux shape systematics as well as the energy scale uncertainty are studied in Fig. 8.18. The energy slope need not be constrained at all in the NMH fit, as confirmed by the investigation of best-fit values, which reveal that the tilt parameter is correctly fitted at the level of $\mathcal{O}(10^{-3})$ even when fitted in combination with other parameters. The zenith slope systematic is found to have a much larger impact. The combination of the flux and zenith slopes, in the inverted

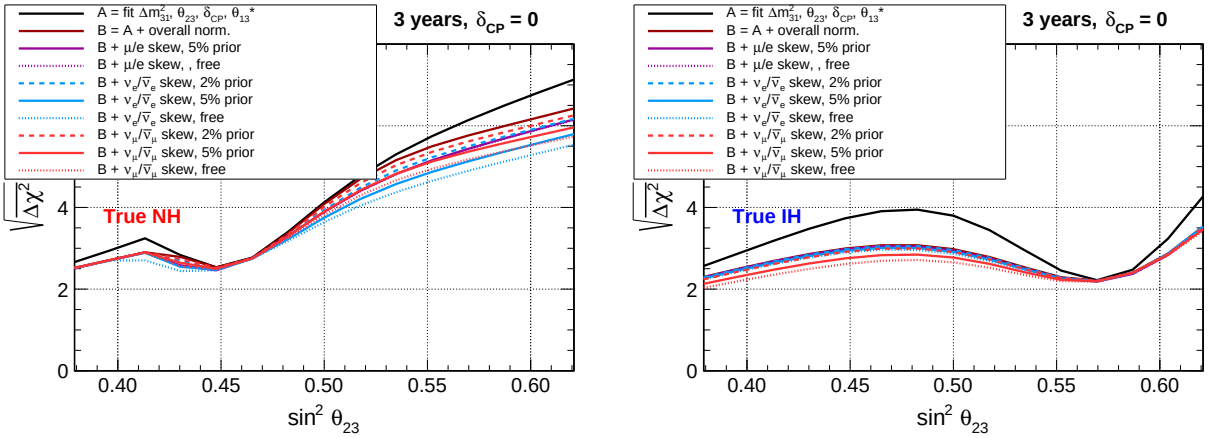

Figure 8.17: Impact of individual flux composition uncertainties (skew parameters).

hierarchy case, is found to have the largest impact amoung the 'individual' uncertainties tested up to now. Investigating the effect of a prior appears necessary here but has not been done yet.

As mentioned in Chap. 7, the investigation of generic flux shape uncertainties, although allowed by the (validated) implementation, has not been done yet as the required computing time is vastly increased, due to the large number of free parameters. This study may be performed in the near future.

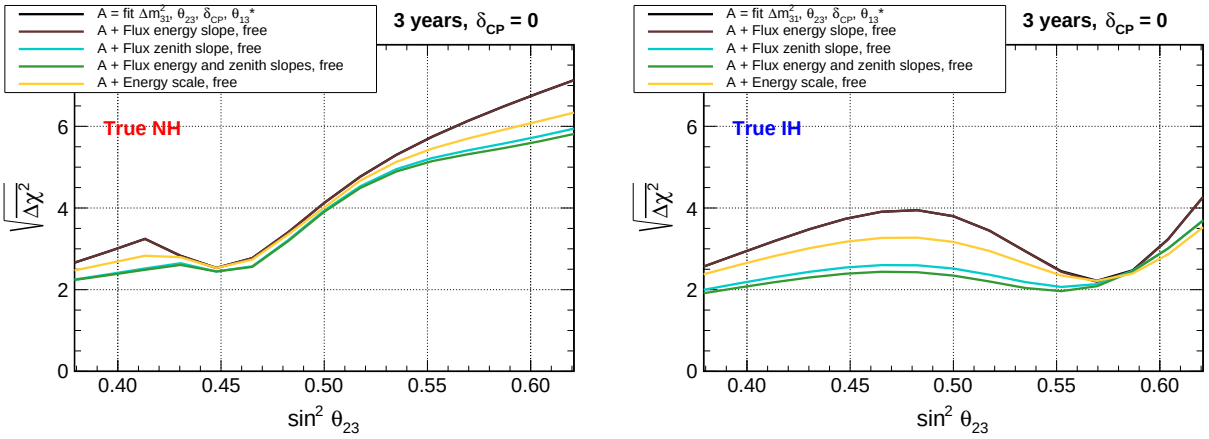

Figure 8.18: Impact of individual flux shape uncertainties (energy, zenith, plus their combination), as well as the energy scale uncertainty. Note that the curves for 'set A' and 'set A + Flux energy slope, free' are superimposed.

8.4 Improvements of the analysis strategy

This section presents the results of some preliminary investigations towards optimising the analysis strategy. As explained previously, the ability to perform such optimisations was a primary motivation for the development of a full MC method of detector response simulation. Unfortunately, since they generally imply a decrease of the MC statistics per bin of signal, these analysis refinements are subject to artifacts arising from the MC sparseness effect. Therefore, the investigation of this effect and development of correction schemes were a prerequisite to enable such studies. While the development step has taken up most of the time allowed for this thesis, the analysis framework now implements a satisfactory way of studying these optimisations and obtaining performance estimates free from MC sparseness biases. The preliminary optimisation

studies presented here should be continued in the near future, either by myself or within the collaboration.

8.4.1 PID cut optimisation in the two-class approach

As described in Chap. 5, the standard 'two-class' strategy classifies events as track-like if their RDF output score P_{track} is greater than a predefined value P_{cut}, and as shower-like otherwise. Optimising the choice of cut further than the naive $P_{\text{cut}} = 0.5$ seems especially relevant considering the moderate performance of the classifier at low energy and its tendency to lean towards the shower classification.

Fig. 8.19 shows the Asimov $\overline{\Delta\chi^2}$ for the NMH determination as a function of P_{cut} for various sets of nuisance parameter settings. The NMH fit is repeated, building a dedicated response matrix, for each value of the cut. The overall $\overline{\Delta\chi^2}$ and the respective contributions of the track-like and shower-like channels are shown.

In all panels of the figure, $P_{\text{cut}} \to 1$ corresponds to a high purity, low statistics track-like sample and a high statistics, low purity shower sample (and conversely for $P_{\text{cut}} \to 0$). Interestingly, the sensitivity is not zero in the extreme cases $P_{\text{cut}} = 0$ and $P_{\text{cut}} = 1$ even though there is no explicit flavour separation. This indirect sensitivity primarily reflects the fact that the detector response is flavour-dependent[6]. The difference in $\overline{\Delta\chi^2}$ between $P_{\text{cut}} = 0$ and $P_{\text{cut}} = 1$ is explained by the better efficiency of the `Dusj` selection cut at low energy, and the fact that `Dusj` performs better on ν_μ events than the track reconstructions on the shower topologies.

In Fig. 8.19(a) no nuisance parameters are included except Δm^2_{31}, which is known to be especially important for track/shower relative sensitivities (Sec. 8.2.1). As discussed previously, the contribution of the track channel then appears very small. In Fig. 8.19(b) and (c), once the θ_{23} uncertainty is included in the fit, the values of $\overline{\Delta\chi^2}_{\text{NMH}}$ obtained with a track-only and shower-only fit are also reported (grey curves). The strong increase for both contributions in the combined fit illustrates again their interplay in constraining oscillation parameters and systematics.

While the optimal choice of P_{cut} turns out to be similar for all parameter sets considered here (although the width of the close-to-optimal interval varies), the result may very well be different when the analysis purpose is not the NMH determination. The ability to easily perform this optimisation for each analysis case will likely be profitable.

8.4.2 N-class approach

The motivation for analysing event distributions in more than two PID classes has been introduced in Sec. 5.1.3. As a preliminary result, the sensitivity improvement achieved by defining five classes, as bins of width 0.2 in P_{track}, is shown on Fig. 8.20.

The uncorrelated smearing has been used in order to reduce the impact of sparse MC artifacts expected from the factor ~ 5 decrease of MC statistics per bin, and the reduced statistics study is performed to ascertain that the increase is genuine. A clear dependence on MC statistics is still observed without the MC error correction. Thus, the effect would have prevented to draw an unequivocal conclusion in the absence of a correction scheme. Applying the MC correction allows to validate an increase of $\sim 20\%$ in $\overline{\Delta\chi^2}$. This result is very encouraging considering that the choice of class definition was not optimised whatsoever.

To further validate the gain arising from the inclusion of detailed PID information in the fit, the issue of potentially increased systematic uncertainties should not be overlooked. Here, the

[6] In fact, due to the slight sensitivity of the ν_τ appearance channel to the NMH, the overall event distribution should have non-zero sensitivity even when assuming a detector response strictly independent of the flavour or interaction channel ν_x.

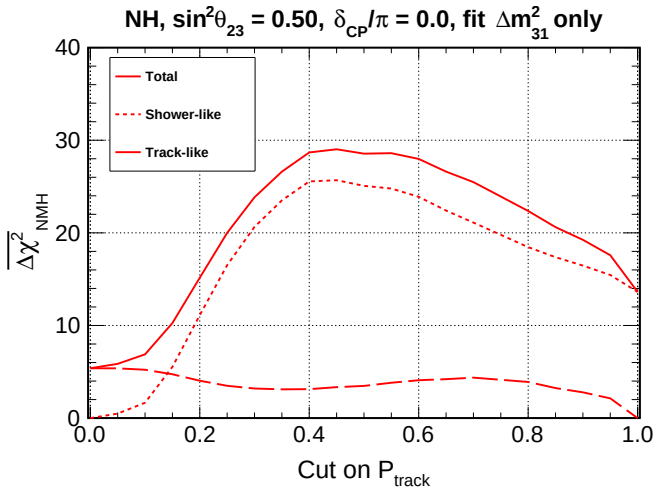

(a)

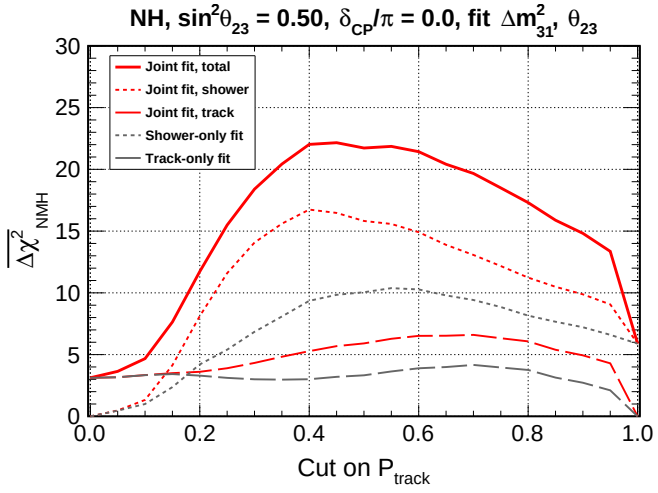

(b)

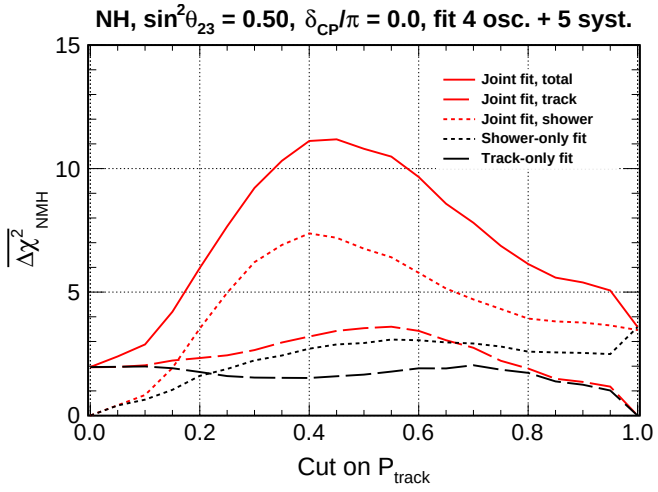

(c)

Figure 8.19: NMH sensitivity as function of the track/shower classification cut in the two-class approach.

impact of adding one free normalisation parameter per class should ideally be studied. Additional suggestions for potential improvements based on similar considerations will be discussed in the outlook section of this thesis.

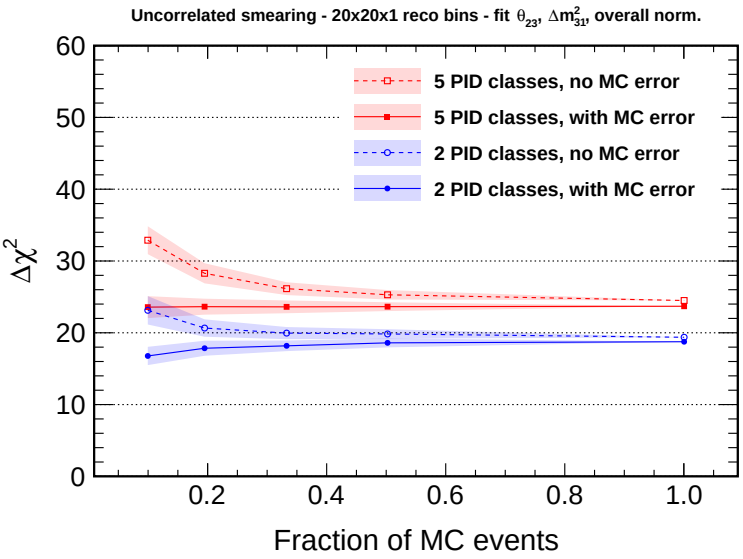

Figure 8.20: Improvement of NMH sensitivity from going to the standard two-class strategy to a fit in five PID classes. The classes are defined as bins of width 0.2 in P_{track}. The other cuts are identical to the two-class strategy, and the events are binned in (E', θ') using the shower reconstruction output for the first 3 classes ($0 < P_{\text{track}} < 0.6$) and the track reconstructions outputs for the last two classes ($0.6 < P_{\text{track}} < 1.0$).

8.4.3 Using reconstructed Bjorken-y for neutrino/antineutrino separation

The impact of including Bjorken-y information in the fit was initially evaluated using the correlated smearing method (and a fine binning), without realising the impact of the MC sparseness effect. A very large increase in sensitivity was found, but it turned out to be essentially due to a statistical artifact. As can be observed from Fig. 8.21, the dependence on MC statistics is indeed very important, and is even present when applying the MC error correction scheme (solid red line).

Drawing a reliable conclusion about a genuine sensitivity improvement (in the infinite MC limit) seems difficult from this preliminary result. In any case, the prospects do not appear as promising as in the case of the N-class analysis strategy.

8.5 Frequentist sensitivity

This section discusses the results of a full frequentist statistical treatment of the NMH sensitivity study, presented officially in the Neutrino 2018 conference (June 2018) [275, 276]. The analysis was performed by myself, using the `Swim` software framework which implements the methods described in this thesis. I was the main developer of both the analysis methodology and the software, with contributions from APC colleagues as detailed in Appendix C. In collaboration with L. Quinn (co-author of Ref. [275]), my analysis was extensively cross-checked with another framework developed within the KM3NeT collaboration, based on a parametrised model of detector response. A sensitivity result produced with the parametrised framework ($\theta_{23} - \Delta m^2_{31}$ measurement contour) was also presented in Ref. [275].

The detector model for these results is based on the 'ORCA2016' MC simulation sample, as in the rest of the thesis. As discussed in Chap. 3, in the production of this Monte Carlo simulation

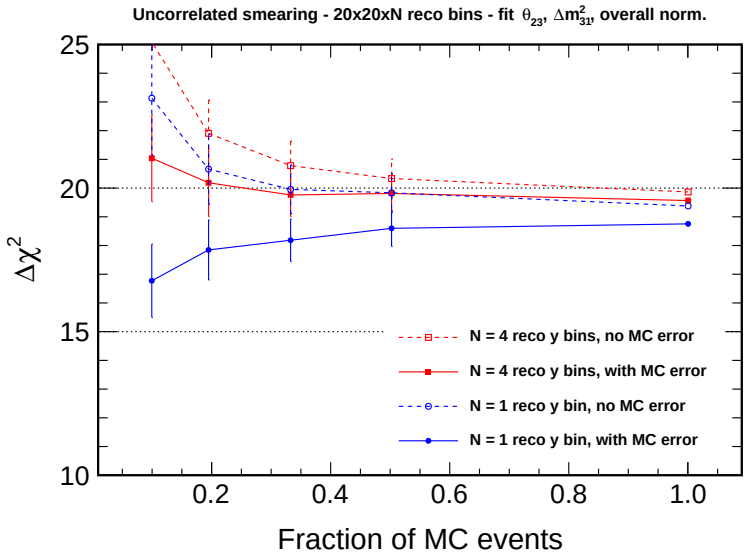

Figure 8.21: Improvement of NMH sensitivity from the inclusion of reconstructed Bjorken-y information.

sample I was only involved in a service task of running the generation and light propagation for atmospheric muon background events (`mupage`+`KM3`). Other steps of the simulation, as well as the reconstructions, event selection and PID steps were developed and run by other students within the collaboration.

Importantly, the results presented here do not include the latest version of the systematics implementation described in Chap. 7 (the practical difference is a minor one), nor the correction scheme accounting for Monte Carlo statistical error. The reason is that both were implemented and validated after Neutrino 2018 – though the MC statistics correction had been under development for longer. Nevertheless, the reported sensitivity is somewhat conservative in that it does not implement any of the potential improvements of the analysis strategy currently under study. The results are also more conservative than the ones obtained with the parametrised framework. An update of these results is planned to be carried out and published in the coming months, applying the latest version of the analysis described in this thesis. A new, recently simulated MC sample implementing an updated detector geometry will also be used.

8.5.1 Analysis settings

The treatment of oscillation parameters and systematics is summarised in Tab. 8.2. The oscillation parameter values were aligned on the settings used in Ref. [21]. With respect to the description of Chap. 7 and the results discussed in Sec. 8.3, the main difference is the implementation of the $\nu/\overline{\nu}$ skew systematic as a single parameter acting identically on both electron and muon flavours (a more optimistic assumption). Distinct track and shower normalisations were fitted freely, replacing the μ/e flavour skew.

Tab. 8.3 shows the reported number of selected events per interaction channel in the analysis range. The event counts are calculated for the NH expectation and assuming three years of data-taking.

8.5.2 The test statistic distributions and frequentist sensitivity

The analysis follows the procedure described in Chap. 6. A number $N_{\text{PEs}} = 3000$ pseudo-experiments is simulated and fitted for each true hierarchy hypothesis and for 7 different assumed true values of θ_{23} in the range $[38° - 42°]$. In total this amounts to $\mathcal{O}(10^5)$ repetitions of the minimisation process. This justifies the need to optimise the speed of the computations per-

parameter	*treatment*	*true value*	*prior*
$\lvert\Delta M^2\rvert\,(\mathrm{eV}^2)$	fitted	$2.48\cdot10^{-3}$	free
$\Delta m^2_{21}\,(\mathrm{eV}^2)$	fix	$7.53\cdot10^{-5}$	–
$\theta_{13}\,(^\circ)$	fitted	8.42	0.26
$\theta_{12}\,(^\circ)$	fix	33.4	–
$\theta_{23}\,(^\circ)$	fitted	$38-52$	free
δ_{CP}	fitted	$0-2\pi$	free
Flux spectral tilt	fitted	0	free
$\nu/\bar{\nu}$ skew	fitted	0	0.03
Tracks normalisation	fitted	1	free
Cascades normalisation	fitted	1	free
NC events normalisation	fitted	1	0.10

Table 8.2: Treatment of oscillation parameters and systematics in the LLR analysis.

Channel	Events/3y	Channel	Events/3y
ν_e CC	44100	ν_τ CC	8800
$\bar{\nu}_e$ CC	17100	$\bar{\nu}_\tau$ CC	3900
ν_μ CC	64000	ν NC	16000
$\bar{\nu}_\mu$ CC	29600	$\bar{\nu}$ NC	4500

Table 8.3: Number of selected events reconstructed as upgoing per channel in the 2-100 GeV reconstructed energy range.

formed at each call of the objective χ^2 function, and to include only the most relevant systematic parameters so as to reduce the number of calls per fit. Here, a few hundred calls were needed for each fit to converge, each call being executed in $\mathcal{O}(1\,\mathrm{s})$ on average (the exact execution time depends on whether the oscillation probabilities need to be recalculated).

As shown in Fig. 8.22, the empirical TS distributions are very well approximated as normally distributed, so that the median sensitivity and p-value expected range can be calculated easily according to Eqs. 6.28, 6.29, 6.31, and 6.32.

Results are shown on Fig. 8.24, comparing the LLR sensitivity to the simplified Asimov result $S=\sqrt{\overline{\Delta\chi^2}}$ for $\delta_{\mathrm{CP}}=0$. Asimov-only results, comparing both true hierarchy hypotheses and showing the impact of true δ_{CP}, are shown on Fig. 8.23. The Asimov results are very consistent with the results discussed in this chapter, showing the characteristic shape (and discontinuity in NH) mostly due to the θ_{23} degeneracy. The LLR-based sensitivity estimates are found to be more optimistic, and the origin of this discrepancy is discussed in the next section.

8.5.3 Comparison with Asimov sensitivity

The discussion done in Chap. 6 (Sec. 6.2.3) applies here to explain the difference between the LLR-based and Asimov sensitivity estimates observed in Fig. 8.24. While the TS distributions are indeed normally distributed, and their mean coincides with the Asimov $\overline{\Delta\chi^2}$ (see Fig. 8.25(a)), the condition $t_{0,\mathrm{NH}}=-t_{0,\mathrm{IH}}$ is not met, as clearly visible from the right panel of Fig. 8.24. This is mostly due to the θ_{23} degeneracy, as discussed extensively in this chapter. Moreover, the condition $\sigma_t=2\sqrt{|t_0|}$ is only approximately met, as shown in Fig. 8.25(b).

In order to obtain more accurate analytical estimates of the median sensitivity, at a lesser computational cost than with the LLR method and accounting for the fact that $\overline{\Delta\chi^2}_{\mathrm{NH}}\neq\overline{\Delta\chi^2}_{\mathrm{IH}}$, the use of the expressions derived in Ref. [159] could be an option. However, these expressions are still based on the assumption $\sigma_t=2\sqrt{|t_0|}$. While the ratio $\sigma_t/\sqrt{|t_0|}$ has been found to be

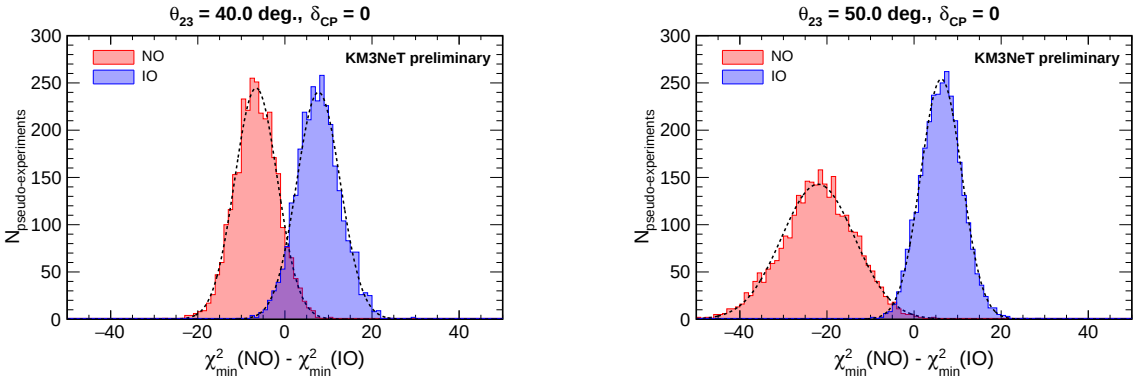

Figure 8.22: Example distributions of the test statistic $t = \Delta\chi^2$ for true normal ordering (NO) and inverted ordering (IO).

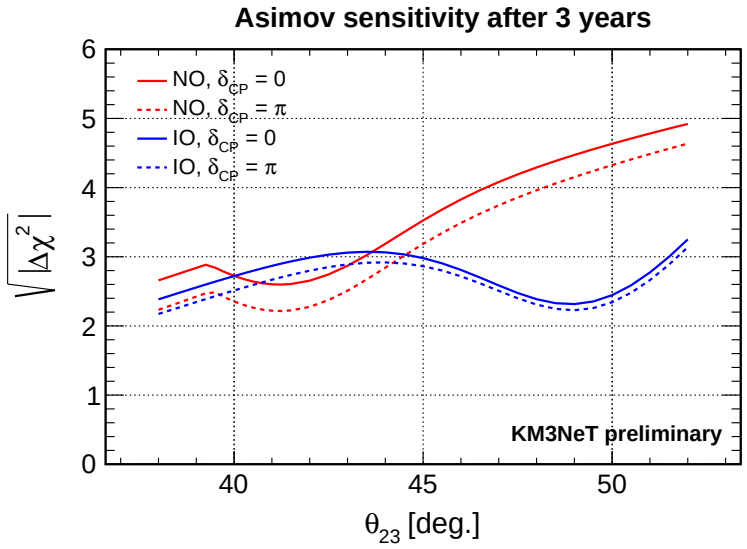

Figure 8.23: Asimov sensitivity result for the Neutrino 2018 analysis, showing both true hierarchy hypotheses and the impact of true δ_{CP}.

closer to 2 in dedicated cross-checks including few nuisance parameters and no priors, in general, as the analysis gets more complex the use of the full LLR method is likely to remain necessary.

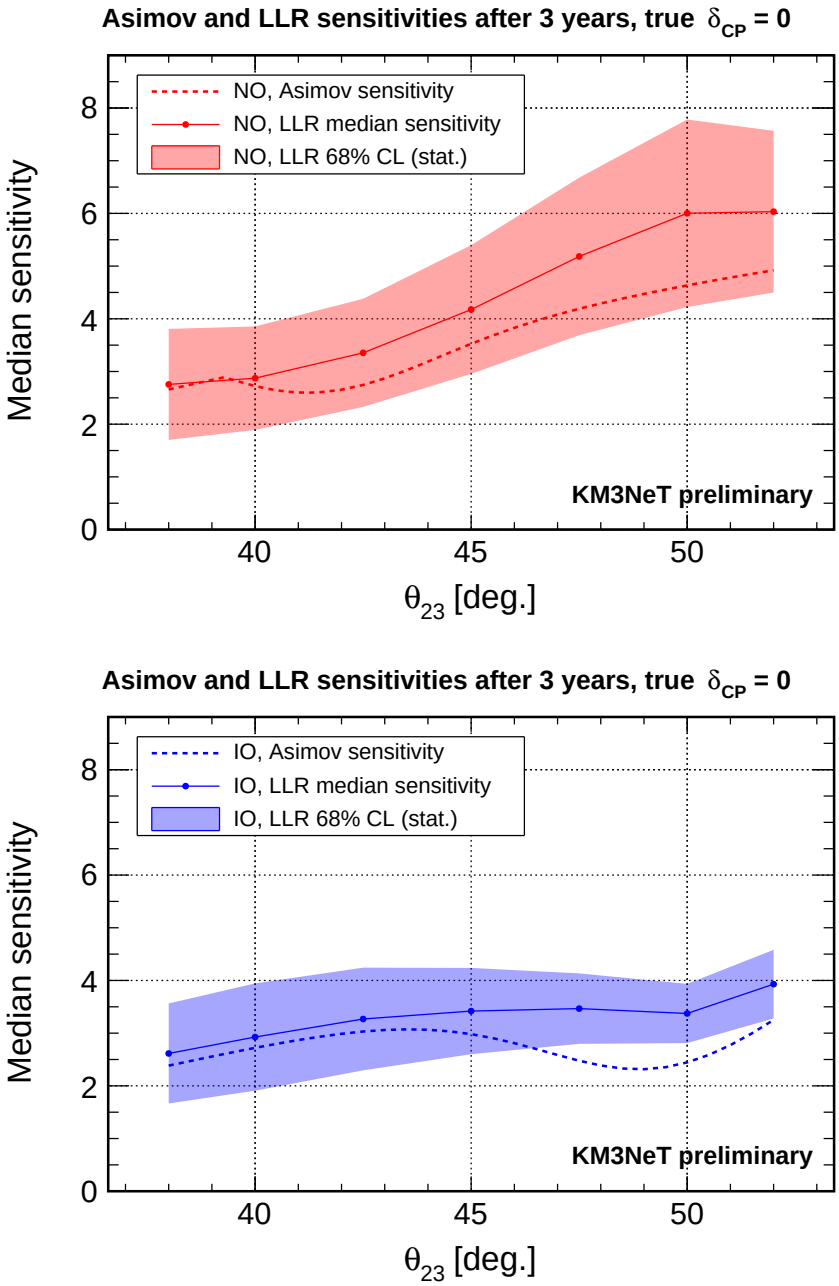

Figure 8.24: NMH sensitivity obtained with the full LLR method. The sensitivity is expressed in number of σ's (one-sided convention). The comparison with the simplified Asimov result $S = \sqrt{\overline{\Delta\chi^2}}$ is shown.

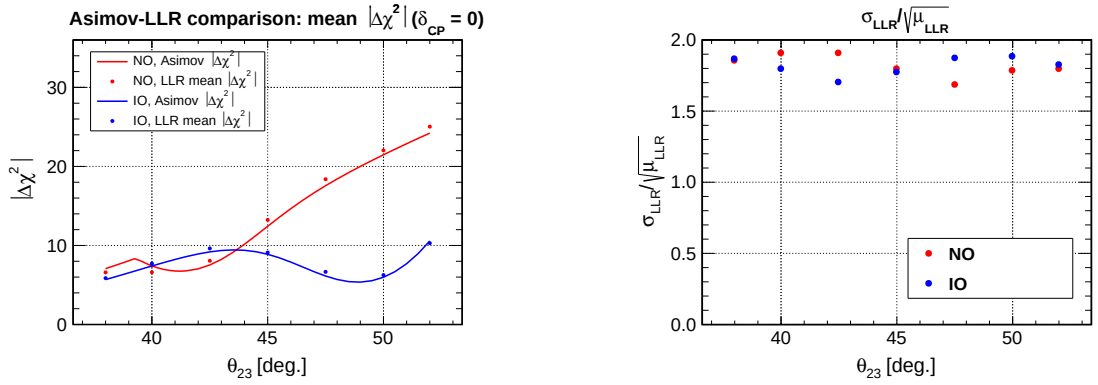

(a) **(b)**

Figure 8.25: Cross-checks for the comparison of LLR-based and Asimov sensitivity results. In the left panel, the mean of the LLR distributions (error bars are negligibly small) are compared with the Asimov $\overline{\Delta\chi^2}$, showing very good agreement. The right panel shows the ratio $\sigma_t/\sqrt{|t_0|}$, which equals exactly 2 under certain simplifying assumptions (see Sec. 6.2.3) [159].

Chapter 9

Further physics potential of ORCA

Contents

Beyond the measurement of the mass hierarchy, ORCA offers a rich potential for other oscillation (and non-oscillation) physics studies. This chapter focuses on the prospects for the precise measurement of atmospheric oscillation parameters and for oscillation tomography of the deep Earth. The expected performance for measuring θ_{23} and $|\Delta m^2_{31}|$, under a given mass hierarchy hypothesis, is investigated in Sec. 9.1. In order to compare the respective impact of oscillation parameters degeneracies and systematics, each measurement is first considered individually. The quantitative performance for the combined measurement is then discussed, considering a more complete panel of assumptions on the true hierarchy and true θ_{23}. The prospects for measuring the electron density in the deep Earth are then examined in Sec. 9.2. Following the approach introduced in Chap. 2, the sensitivity to the average proton-to-nucleon ratio (Z/A) in extended compositional layers is considered, assuming the mass density profile to be very well constrained by geophysics. The expected performance for the lower mantle and outer core are investigated more specifically.

9.1 Atmospheric oscillation parameters

In this chapter, sensitivity estimates are computed for a data-taking time of three years and under the assumption that the correct mass hierarchy has been determined beforehand. The measurement of oscillation parameters without assumption on the hierarchy (for instance, short-term determination of the octant of θ_{23} regardless of the NMH) may also offer interesting potential, but has not been studied here. Confidence intervals and contours are obtained from the profile likelihood ratio approach, based on Asimov datasets, as described in Chap. 6. The analysis strategy follows the standard two-class approach and, unless otherwise specified, the detector response uses the correlated smearing method with MC error correction enabled.

9.1.1 Measurement of θ_{23}

Fig. 9.1 shows the profile likelihood ratio for the measurement of θ_{23}, breaking down the impact of various sets of nuisance parameters. The combination of true NH and true $\sin^2\theta_{23} = 0.55$ (close to the current world best-fit) is assumed, as well as $\delta_{\mathrm{CP}} = 0$. Nuisance parameters have been included incrementally in the following sequence:

(a) Oscillation parameters: Δm^2_{31} (free), δ_{CP} (free), θ_{13} (with standard prior).
(b) Addition of free normalisation for the track-like and shower-like samples.
(c) Addition of the flux composition systematics: μ/e flavour skew (5% prior), $\nu_e/\overline{\nu}_e$ skew (5%), $\nu_\mu/\overline{\nu}_\mu$ skew (5%); plus NC events normalisation (10%), energy and zenith slope parameters (free).
(d) Addition of the energy scale systematic (free).

The partial octant degeneracy is reflected by the non-symmetricity of the profile with respect to $\sin^2\theta_{23} = 0.5$. The exclusion of high values of $\sin^2\theta_{23}$ is unaffected by oscillation parameter degeneracies, being dominated by $\nu_\mu \to \nu_\mu$ disappearance in the saturation regime. On the contrary, the octant sensitivity (coming from the electron channel) is significantly affected.

The addition of free channel-wise normalisations has a strong impact, as expected considering that $\sin^2\theta_{23}$ is an amplitude parameter. The remaining set of systematics have a moderate but non-negligible impact, whereas the addition of a free energy scale parameter does not affect the measurement further.

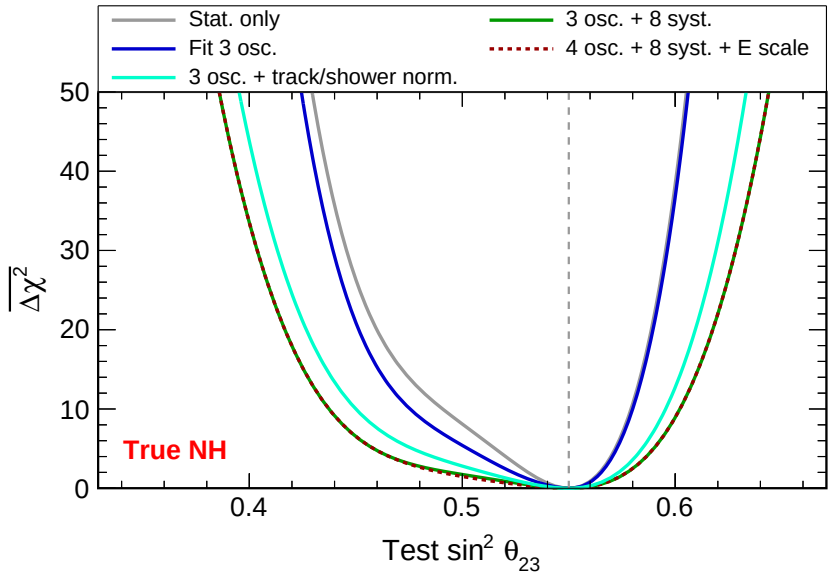

Figure 9.1: Impact of oscillation and systematic uncertainties for the measurement of θ_{23}. The vertical dotted line indicates the assumed true value of $\sin^2\theta_{23}$. The true value of the CP violating phase is set to $\delta_{\mathrm{CP}} = 0$.

9.1.2 Measurement of Δm^2_{31}

The expected performance for the measurement of Δm^2_{31} is considered in Fig. 9.2, with a similar sequence of addition of systematic parameters. Oscillation parameter degeneracies appear to have a close to negligible impact. Unlike in the case of the θ_{23} measurement, the energy scale uncertainty is by far the dominant systematic, as is expected considering that a shift of Δm^2_{31} essentially corresponds to shifting oscillation patterns with respect to energy (see Fig. 2.19). The strong anti-correlation between Δm^2_{31} and the energy scale systematic parameter in the χ^2 minimisation is further demonstrated in Fig. 9.3.

Various sources of uncertainty may enter in the absolute energy calibration error for the measurement of neutrino events in ORCA. While a detailed discussion and estimation of these effects is beyond the scope of this thesis, it is clear that controlling this parameter to a level better than a few percent will be challenging.

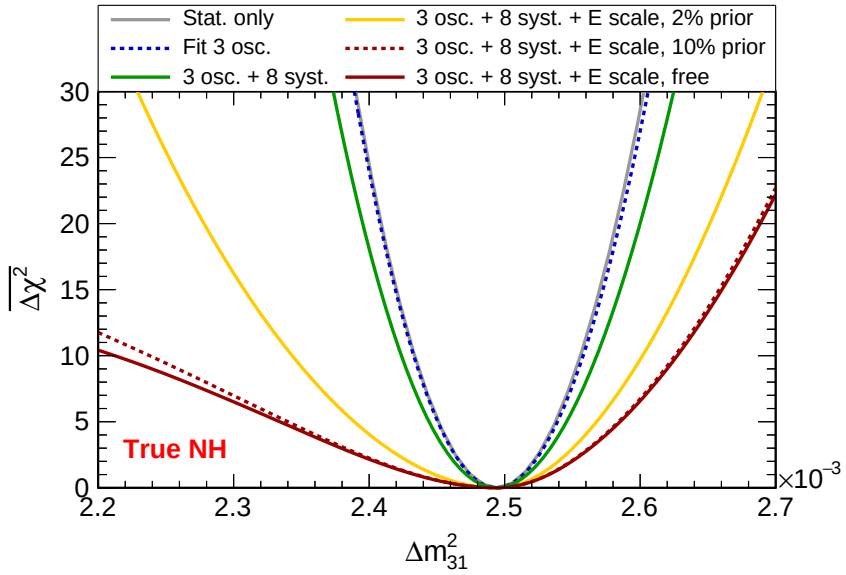

Figure 9.2: Impact of oscillation and systematic uncertainties for the measurement of Δm^2_{31}. The assumed true values are $\sin^2\theta_{23} = 0.5$ and $\delta_{\mathrm{CP}} = 0$. Unlike Fig. 9.1, the specific impact of normalisation parameters is not shown, as the energy scale uncertainty dominates the impact of systematics overall.

9.1.3 Combined measurement

Fig. 9.4 shows the expected confidence regions for a combined measurement of θ_{23} and Δm^2_{31} in six different combinations of true assumptions: normal and inverted mass hierarchy, and θ_{23} at maximal mixing, lower and upper octant ($\sin^2\theta_{23} = 0.4,\ 0.5,\ 0.6$). All nuisance parameters are included as in the last two sections, with the exception of the energy scale systematic.

It can be observed that

- there is little correlation between the two measurements;
- similar, close to symmetrical expected confidence regions are found when flipping the true octant;
- the true NH case is slightly more favourable for the θ_{23} measurement (for the non-maximal test-values the wrong octant is excluded at 3σ, while it is not the case for true IH);
- as seen in the last section, extreme high and low values of $\sin^2\theta_{23}$ are more easily excluded, while the maximal mixing case is unfavourable.

The expected performance of ORCA for the measurement of atmospheric oscillation parameters with only three years of data clearly improves on the current precision of existing

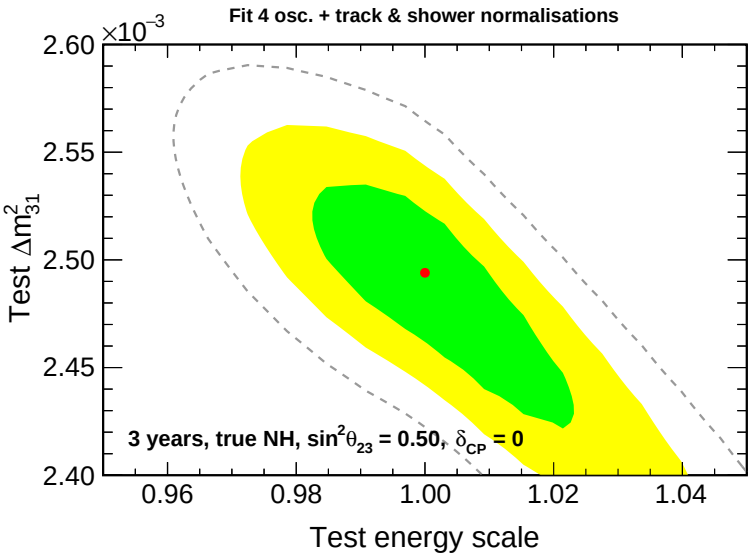

Figure 9.3: Correlation of Δm^2_{31} and the energy scale systematic parameter in the fit. The green, yellow, and grey contours correspond respectively to an expected coverage of 1σ (68.3%CL), 2σ (95.5% CL) and 3σ (99.7%CL). Only oscillation parameters and free normalisations have been included as nuisance here for practical reasons.

measurements, shown for instance on Fig. 2.14 (a direct comparison was also shown in Ref. [275] for different sets of contours). Although controlling the energy scale uncertainty at the $\sim 2\%$ level is definitely challenging for an experiment such as ORCA, this systematic does not affect the measurement of θ_{23}, which is the most important from the point of view of theoretical implications (see Chap. 1).

Most of the experiments leading the current determination of the atmospheric parameters (Fig. 2.14) will continue taking data and are expected to improve the precision of their results in the following years. A proper comparison would require an extrapolation of these results to the time period where ORCA will have accumulated three years of data with a full detector, which cannot be done easily here. In any case, it seems clear that the measurements provided by ORCA at least contribute very significantly to the global precision achieved by ~ 2025.

It should be emphasized that the sensitivity results of this section are only preliminary, and expected to be somewhat conservative. Although cross-section uncertainties have not been studied in detail and should be included in the near future, normalisations and flux uncertainties are here set to conservative assumptions: normalisations are fitted freely, and the priors applied to the flavour and $\nu/\overline{\nu}$ skew parameters correspond respectively to about 7-8% and 10% uncertainties on the flux ratios (see Chap. 6). Moreover, the standard two-class analysis strategy (with cuts optimised for the NMH determination) was used without dedicated optimisation.

9.2 Neutrino oscillation tomography of the Earth

In this work, the potential of ORCA for probing Earth's electron density has been studied by considering the measurement of uniform variations (rescaling) of the electron density in whole compositional layers. Assuming the mass density profile to be known exactly from external measurements (geophysics), the rescaling factor can be identified to the average electron-to-nucleon ratio of the medium. The motivation and limitations for this approach were discussed in Chap. 2 (Sec. 2.3.3). Details on the Earth model can be found in Chap. 4 (Sec. 4.2.1).

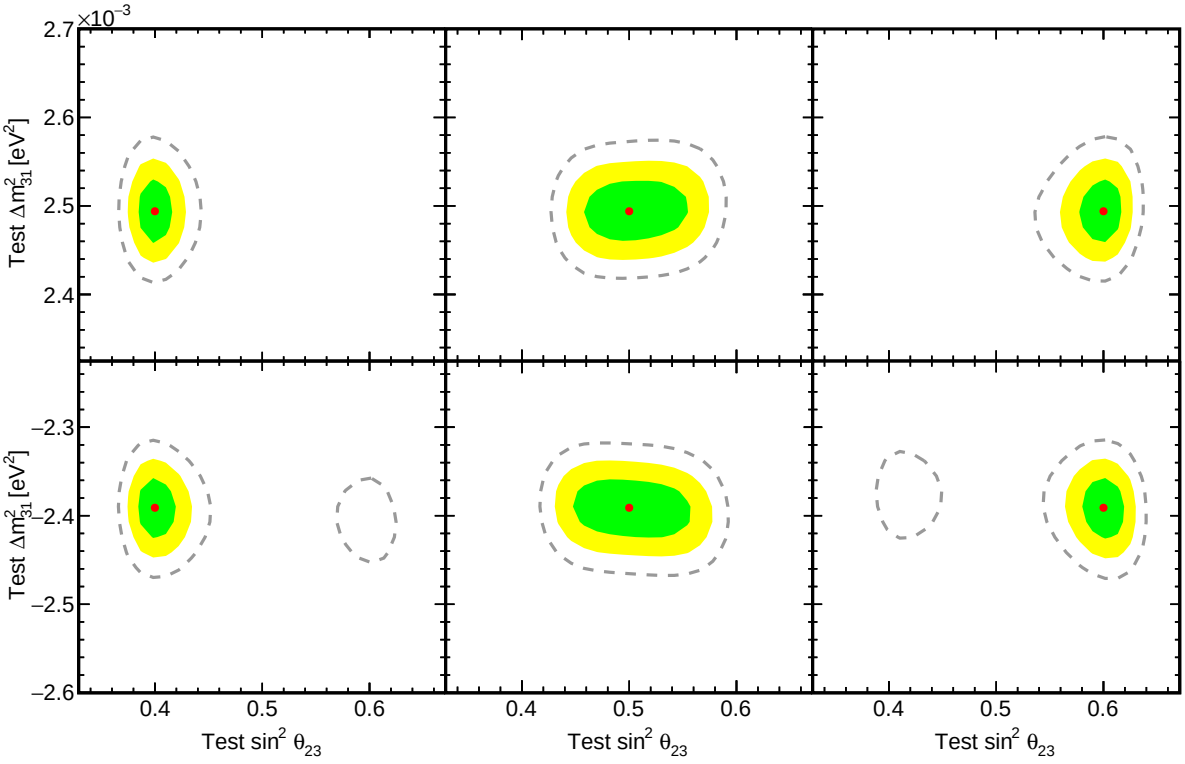

Figure 9.4: Expected confidence regions for the combined measurement of θ_{23} and Δm^2_{31} after 3 years of operation of ORCA. The energy scale systematic is not included. True NH (IH) is assumed in the top (bottom) panel. The assumed truth for the parameters of interest are indicated by the red dots, and the green, yellow, and grey contours correspond respectively to an expected coverage of 1σ (68.3%CL), 2σ (95.5% CL) and 3σ (99.7%CL).

9.2.1 Underlying signal and its measurement in ORCA

The statistical signal induced by a variation of Z/A in a given layer, at the level of $\nu + \overline{\nu}$ interaction rates, is shown in the $(E_\nu,\ \theta_z)$ plane on Fig. 9.5. Variations of Z/A of +5% with respect to the nominal value have been applied separately in

- the whole mantle together with the crust (top panel);
- the whole outer core (bottom panel).

Bins of constant width in zenith angle are used instead of the usual $\cos\theta$. While this has the advantage to allocate a larger fraction of the zenith range to core-crossing trajectories (see Tab. 2.1), the expected event statistics in these bins (already comparatively low due to the flux angular dependance) decreases even more.

For both the mantle and outer core, the electron channel appears to exhibit extended two-dimensional regions of statistical significance, whereas the muon channel shows faster but much more intense oscillatory patterns. The statistical separation at this level is clearly larger in the muon channel. While this comparison has some similarities with the NMH signal, there are qualitative differences that can be better appreciated by reconsidering Fig. 2.16 (NMH signal) and Fig. 2.24 and 2.25 (mantle and core) in Chap. 2.

As can be seen from Fig. 9.5, the signal arising from the Z/A variation in the upper mantle and crust (defined by $180 - \theta_z > 63.5°$) is negligibly small when compared to the signal in the lower mantle. At the same time, the potential measurements in the lower mantle and outer core are of greater interest for geophysics. The results presented in the next section thus focus on ORCA's sensitivity to these two layers.

The signal expected to be seen in ORCA, once all detector effects are accounted for, is shown on Fig. 9.6 for similar variations of +5% of Z/A in the lower mantle and outer core (mind that the shown energy range is larger than in Fig. 9.5). As expected, fine structure effects are mostly washed out by the energy-angular smearing. In addition, the low detection efficiency and limited overall detector performance at low energy tend to shift the observable signal to higher energies, to the point that some of the dominant patterns seen in reconstructed events are barely visible at the level of interaction rates (being dominated by the low energy signal). These effects impact the muon channel more strongly, so that the overall expected statistical separation becomes comparable in the cascade-like and track-like channels.

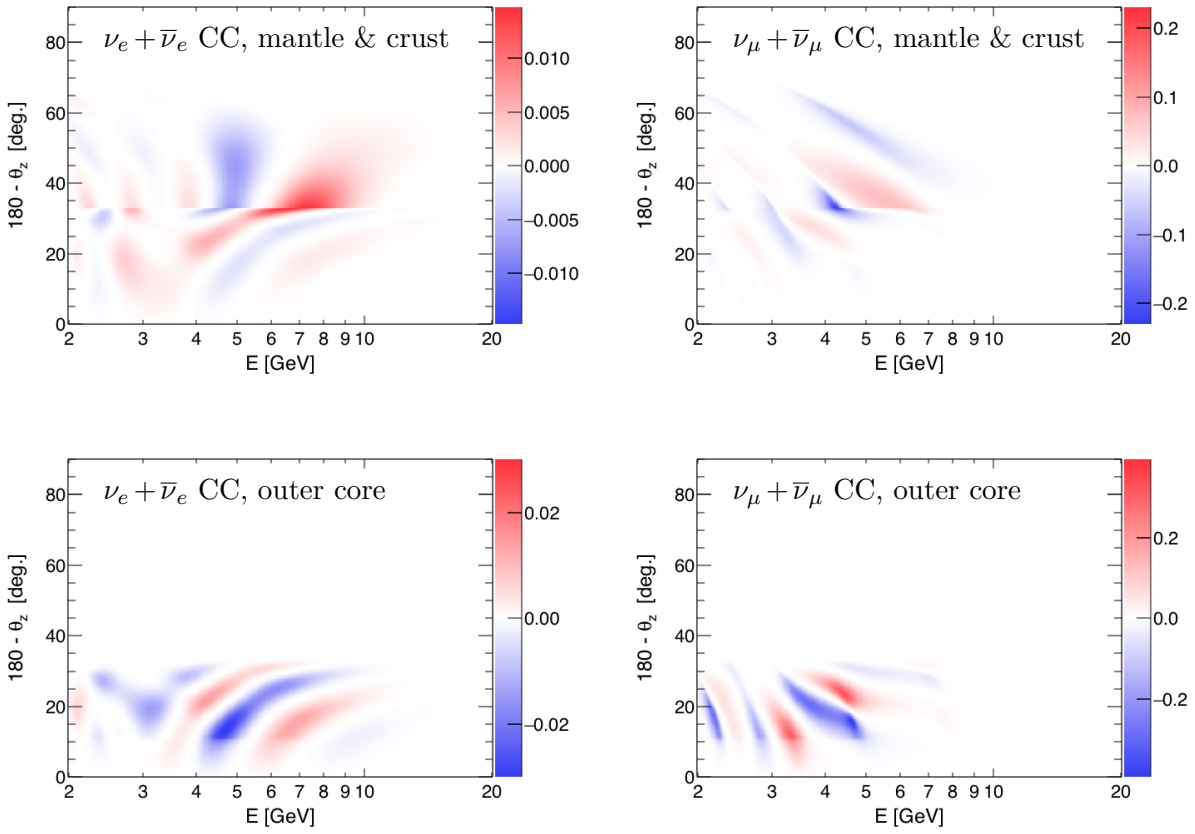

Figure 9.5: Expected signal at the level of $\nu + \overline{\nu}$ interaction rates, shown as maps of the χ^2_s indicator (mind the different color scales on the left and right plots). The plotted quantity is $\chi^2_s(A, B)$ where A is the nominal Earth model and B is the same model with a +5% uniform variation of Z/A in the relevant layer. No detector effects are included, except that ν and $\overline{\nu}$ are not distinguished. Figure reused from Ref. [277].

9.2.2 Timescale and consistency with oscillation analyses

As could be expected from the qualitative study of Chap. 2 (Fig. 2.16, 2.24, 2.25), the overall statistical signal associated to a 5% variation in Z/A in either layer is significantly smaller than the one associated to the mass hierarchy flip. This indicates that a larger exposure will be needed for ORCA to achieve a few percent Z/A measurement, even at the $\sim 1\sigma$ confidence level. In the following, an effective data-taking duration of 10 years is therefore considered.

In fact, another motivation for considering the tomography analysis on a longer timescale comes from its interplay with the mass hierarchy and oscillation parameter measurements. On the one hand, considering no external constraint on the Z/A scaling parameters affects ORCA's sensitivity to the NMH in a non-negligible manner. This was demonstrated quantitatively in

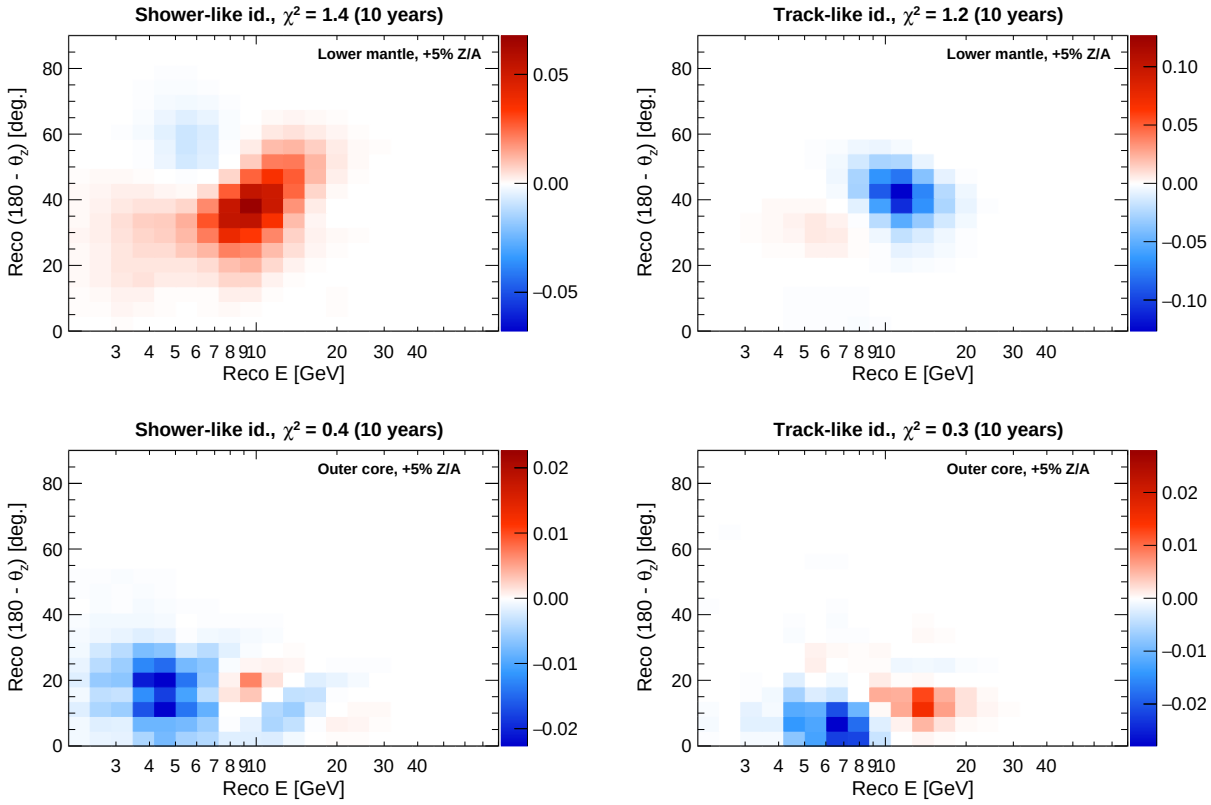

Figure 9.6: Expected signal in the final track-like and cascade-like samples, shown as maps of the χ^2_s indicator. Mind that the energy range shown here is extended with respect to Fig. 9.5. The Z/A factor is varied by +5% separately in each layer (lower mantle in top panel, outer core in bottom panel). 10 years of exposure with a full detector, true NH and reference oscillation parameters are assumed, and the detector response is applied using the uncorrelated smearing method.

Sec. 8.2.4, and the interplay can also be understood by comparing the NMH and tomography signal maps, which have a large overlap. It is then natural to use the knowledge of the Earth available from geology and geophysics to support neutrino oscillation studies, and keep the Z/A parameters (more generally the Earth model) either fixed or constrained by a tight prior in these analyses. A later reanalysis of the same data can be performed with the aim to measure Z/A parameters. Clearly, to preserve consistency the new analysis must be independent and not account for the results of the previous oscillation analysis. More generally, it should not include any constraint on oscillation parameters coming from measurements of matter effects where a model of Earth's electron density has been used as input.

Strong external constraints on θ_{13} and on the mass hierarchy are expected to be required to perform an oscillation tomography analysis with ORCA. While the current knowledge of θ_{13} is dominated by reactor experiments (independent of matter effects), a determination of the NMH independently of ORCA would be needed. On a 10 to 15 years timescale it is very likely that this will be achieved with high confidence level, either by JUNO or next-generation long-baseline experiments (DUNE, Hyper-Kamiokande) – JUNO having the advantage of being completely independent of matter effects. While DUNE and Hyper-Kamiokande do rely on matter effects, their neutrino beams only traverse Earth's crust, with far better known properties than the deep mantle and core of the Earth.

9.2.3 Sensitivity results

Fig. 9.7 shows one-dimensional scans of profile likelihood ratio for the Z/A measurement in both the lower mantle and outer core. The mass hierarchy is assumed to be known and kept fixed

in the fit, assuming a normal (inverted) hierarchy in the top (bottom) panel. Other parameters are fixed to their reference values. As a complement, the widths of the confidence intervals at 1σ inferred from each $\Delta\chi^2$ graph are reported in Tab. 9.1.

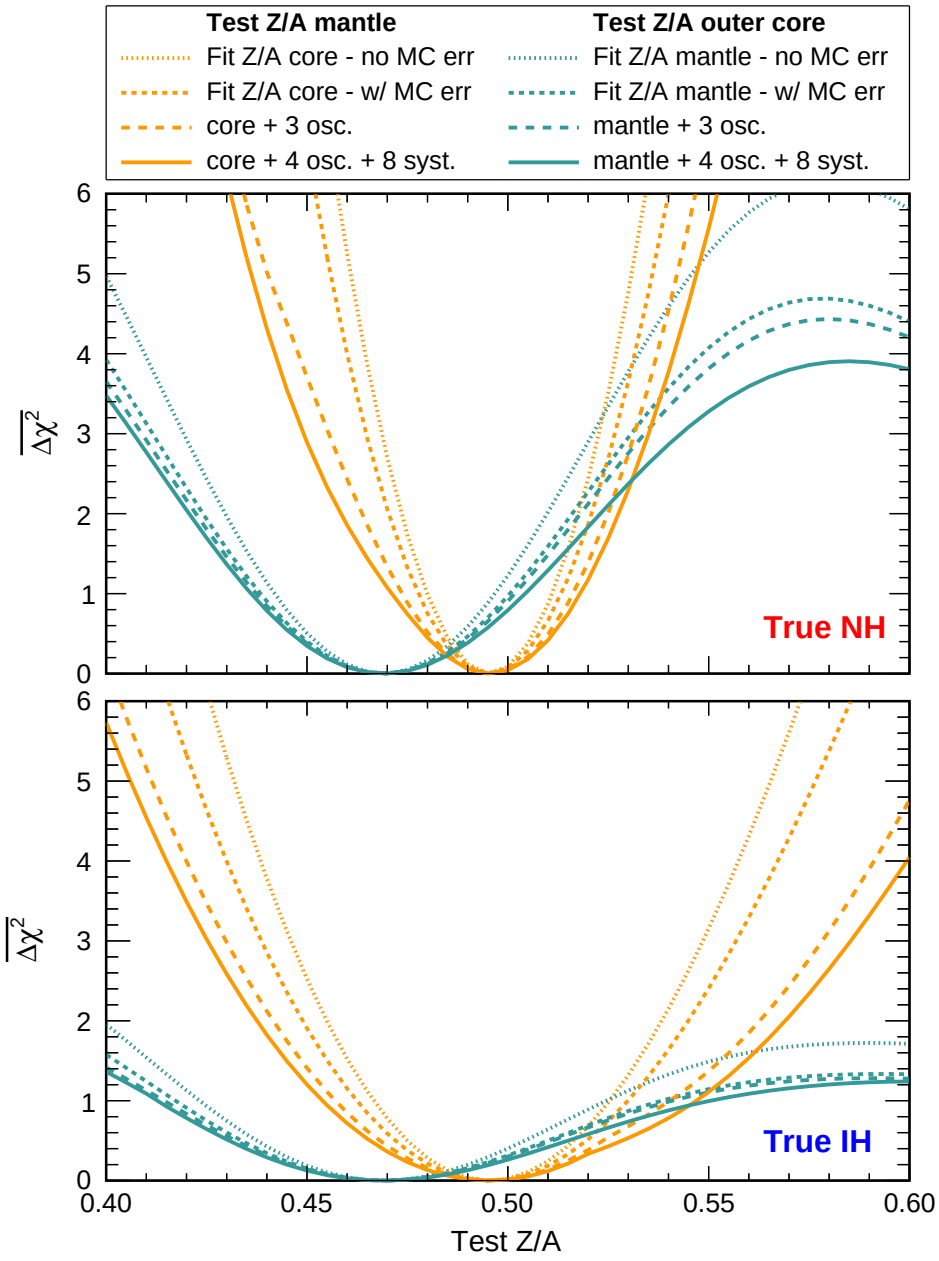

Figure 9.7: One-dimensional profile likelihood ratio for the measurement of the Z/A scaling in the lower mantle and outer core after 10 years of operation of ORCA. The top (bottom) panel correspond to a true normal (inverted) hierarchy. In both cases $\sin^2\theta_{23} = 0.5$, and $\delta_{\mathrm{CP}} = 0$. The data-taking time is set to 10 years.

The correlated smearing method has been used to produce these profiles (unlike in Fig. 9.6). For all curves the Z/A scaling in the other layer is profiled upon as a systematic, and additional nuisance parameters are added incrementally:

(a) without MC error correction scheme;
(b) enabling the MC error correction scheme;
(c) addition of θ_{23}, Δm^2_{31} and δ_{CP} as nuisance parameters (free);

	Lower mantle		Outer core	
	NH	IH	NH	IH
no MC err.	-3.1, 3.2	-5.7, 6.1	-6.0, 5.7	-9.9, 11.6
w/ MC err.	-3.5, 3.7	-6.6, 7.1	-6.7, 6.5	-11.2, 14.6
+ 3 osc.	-4.3, 4.3	-7.7, 9.2	-6.9, 6.8	-11.8, 15.1
+ 8 syst.	-4.8, 4.6	-8.3, 10.5	-7.2, 7.3	-12.2, 17.1

Table 9.1: Relative widths (in %) of the expected 1σ confidence intervals for the measurements of the Z/A scaling in the lower mantle and outer core. The labels and widths correspond to the $\Delta\chi^2$ graphs shown on Fig. 9.7.

(d) addition of θ_{13} (standard prior) and the full set of systematics, as in Sec. 9.1: μ/e flavour skew (5% prior), $\nu_e/\overline{\nu}_e$ skew (5%), $\nu_\mu/\overline{\nu}_\mu$ skew (5%); plus NC events normalisation (10%), energy and zenith slope parameters (free).

In addition, two-dimensional contours for the combined measurement of both Z/A parameters are shown in Fig. 9.8, considering larger ranges of Z/A.

As expected from the statistics-only results of the previous section, the measurement is more performant in the lower mantle than in the outer core. All systematics included and assuming true NH, a level of precision of approximately ±5% (at 1σ) can be reached in the lower mantle. In the outer core, this performance is reduced to about ±7%. The inverted hierarchy is also much less favourable, mostly due to the fact that the MSW resonance then occurs for antineutrinos instead of neutrinos, thus reducing the statistical signal since their CC cross-section is about twice smaller.

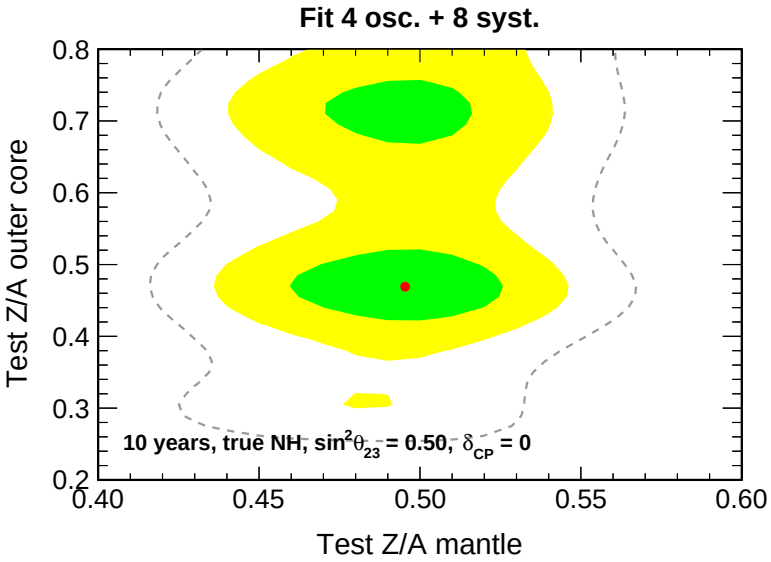

Figure 9.8: Expected confidence regions for the combined measurement of the Z/A scaling in the lower mantle and outer core. The green, yellow, and grey contours correspond respectively to an expected coverage of 1σ (68.3%CL), 2σ (95.5% CL) and 3σ (99.7%CL).

These results show that ORCA is very unlikely to ever achieve a level of $\sim 1\%$ precision, which would be a minimal requirement to reach the Z/A range spanned by realistic average chemical models (in either mantle or core). Nevertheless, the study demonstrates the concrete potential of neutrino oscillation tomography: ORCA's measurement, a completely new technique based solely on the weak interaction of neutrinos, could exclude the hypothesis of identical chemical compositions of the lower mantle and outer core at a confidence level better than 1σ. Very exotic

compositional models, such as the hydridic Earth model predicting an average $Z/A \sim 0.57$ in the core [278, 279], could also be excluded with good significance.

The fact that the measurement in the outer core is not strongly correlated with that in the mantle (as shown by the contours of Fig. 9.8) is also interesting, as it demonstrates the ability of the technique to access the inner Earth without being affected by the 'shielding' of the outermost layers. This arises from the possibility to measure large samples of neutrinos with both core-crossing and mantle-only trajectories. In that respect, using the flux of atmospheric neutrinos is advantageous for probing the deep Earth, with respect to using a single baseline beam-detector setup.

Discussion of the impact of degeneracies and systematics

In Fig. 9.7, the mantle measurement is found to be affected rather strongly by oscillation uncertainties, especially for true NH (IH) in the lower (upper) part of the range of test Z/A values. The examination of fit results in the case where only θ_{23}, Δm^2_{31} and δ_{CP} are included as nuisance reveals that the minimisation exploits the freedom of all three fitted parameters, notably including a rather large drift of Δm^2_{31} ($\pm 1\mathrm{eV}^2$) in the tested range of Z/A.

These results suggest to evaluate the impact of applying priors on these parameters. In consistency with the foreseen timescale of the measurement, the level of constraint expected to be reached by future precision experiments on the long term (~ 2030) could be considered. The impact of oscillation parameter degeneracies is then expected to be strongly reduced, so that flux-related and other systematics would likely dominate. On a side note, although the energy scale systematic has not been considered here, its impact would deserve to be evaluated. Indeed, based on the qualitative impact of density changes on oscillation patterns in the electron channel (see Fig. 2.24-2.25), can be expected to affect the measurement.

The impact of the sparse MC effect and correction is strong overall, and very strong in the case of the outer core measurement. This is somewhat expected, considering that the signal is concentrated in a rather small fraction of the phase space (recall that the MC is generated with a flat distribution in $\cos\theta$). With respect to the NMH analysis, the increase in exposure is also expected to amplify the discrepancy between corrected and non-corrected estimates. In fact, as suggested by Fig. 8.7, the MC correction may be overconservative in this case. Generally speaking, in an analysis where neutrino events are binned according to reconstructed zenith angle, estimating the sensitivity to the outer core electron density presents some challenges – be it based on a full MC approach as described here, or based on a parametrised method as briefly mentioned in the outlook section.

Finally, the impact of the true value of θ_{23} on the tomography performance is worth mentioning. As in the case of the NMH measurement, a large value of θ_{23} is more favourable. Considering for instance the mantle measurement in true NH, the width of the 1σ interval (all systematics included) decreases from about $\pm 6-7\%$ for $\sin^2\theta_{23} = 0.38$ to about $\pm 3\%$ for $\sin^2\theta_{23} = 0.62$.

It is interesting to note that high values of Z/A in the outer core cannot be excluded with high confidence level, as illustrated by the decrease of the one-dimensional $\Delta\chi^2$ profile above $Z/A \simeq 0.58$ in Fig. 9.7, and by the existence of two 'islands' at the 1σ confidence level in Fig. 9.8. As pointed out in Ref. [222], this can be explained by the existence of self-similarities of the oscillation probabilities for $E_\nu > 5\,\mathrm{GeV}$ as the electron density increases: the difference between the nominal model and $Z/A \simeq 0.7$ is faint and only observable at low energy, in a range where most of the signal is smeared out in ORCA. More generally, from the comparison of the signal at the level of interacting event rates (Fig. 9.5) with the signal effectively observed (Fig. 9.6), it appears that the most efficient solution to increase the sensitivity of ORCA to the outer core electron density would require both to lower the detection threshold and increase the reconstruction and classification performance at low energy.

Summary and Outlook

The KM3NeT/ORCA experiment has been presented within the landscape of today's experimental neutrino physics. At the turn of the century, ground-breaking experiments detecting solar and atmospheric neutrinos have discovered the existence of neutrino flavour oscillations, implying that neutrinos have mass and providing, as of today, the most direct window on physics beyond the Standard Model. In the past two decades, a first generation of experiments using solar, atmospheric, reactor and accelerator neutrinos have established a compelling picture of three-flavour neutrino oscillations.

A few parameters of the model remain to be determined, requiring a new generation of experiments. The neutrino mass hierarchy (NMH) stands out as one of the most crucial of these unknowns. Beyond its theoretical importance, relating to the long-standing question of the observed pattern of fermion masses, the knowledge of the NMH impacts the design of future precision oscillation experiments aiming for the discovery of leptonic CP violation, and the intepretation of neutrinoless double beta decay searches, seeking to determine the fundamental nature of the neutrino. Dedicated experimental efforts have therefore been undertaken, targeting an unequivocal determination of the NMH on a few years timescale. One of the most promising measurement strategies exploits the NMH-dependent enhancement of the oscillations of atmospheric neutrinos and antineutrinos crossing the Earth matter along baselines of thousands of kilometers.

The work summarised in this thesis has been carried out within the European collaboration KM3NeT, which has started the construction of ORCA, a megaton-scale underwater Cherenkov detector in the Mediterranean Sea off the coast of France. The estimation of the performance of the experimental design and data analysis strategies require the development of complex and computationally demanding Monte Carlo (MC) simulations, state-of-the-art statistical methods, and accurate modeling of systematic uncertainties.

The first results of such sensitivity studies were reported in the KM3NeT Letter of Intent [21], optimising the detector geometry and indicating a median sensitivity to the NMH determination of about 3σ after three years of operation. These results were based on preliminary simulations (masking technique) and the analysis did not account for the full complexity of detector response effects, neglecting for instance some correlations between reconstructed variables. Marginal optimisations of the detector geometry are still possible, and the trigger, reconstruction and particle identification (PID) algorithms are an ongoing work. The exploration of improved analysis strategies, such as exploiting reconstructed Bjorken-y information, was also a priority.

This motivated the development of a new analysis methodology and software framework for sensitivity studies, which has been the focus of my research work for the past three years. The goal was to enable the study of complementary physics cases in a consistent framework: neutrino mass hierarchy and atmospheric oscillation parameters, Earth tomography, sterile neutrino and non-standard interaction searches, *etc.* As design requirements, the analysis framework was intended to account for detector response effects as extensively as possible, and to offer full flexibility and robustness with respect to refinements of the analysis strategy and extensions of the set of considered systematic uncertainties.

To this end, the analysis methodology I developed is based on the simulation of the detector response via fully correlated response matrices built from MC events. This method retains a good part of the advantages of event-by-event reweighting techniques, while being far more efficient computationally. Statistical artifacts may arise from the sparse sampling of the response matrix, a result of the fairly high dimensionality of the problem and finite amount of MC simulated events available. The average effect of these 'sparse MC' artifacts is shown in this thesis to amount to an overestimation of sensitivity. This was initially an important obstacle for the evaluation of the potential of refined analysis strategies, which generally imply an increase of dimensionality. Adaptations of the response matrix methodology have been developed, as well as schemes to account for the sparse MC effect in the statistical analysis. These corrections schemes are successful in i) demonstrating the interest of preserving correlations between reconstructed variables and ii) obtaining unbiased estimates of the potential of optimised analysis strategies. Preliminary results are reported, regarding the inclusion in the fit of reconstructed Bjorken-y and of more detailed flavour identification information.

The framework is also used in this thesis to describe the internal mechanics of the NMH measurement with ORCA, breaking down the impact of detector effects and oscillation parameter degeneracies. New sources of systematic uncertainties are studied, including Earth model assumptions as well as neutrino flux composition and shape uncertainties. While Earth electron density parameters have little impact when constrained to the current geophysical precision, some of the newly introduced flux systematics have proven quite critical, in particular the flavour-specific neutrino/antineutrino ratios.

Frequentist LLR-based estimates for the NMH sensitivity have been produced, based on the most recent MC simulation of the detector to date. These were reported as official KM3NeT results in international conferences. The median sensitivity to the NMH determination heavily depends on the true value of θ_{23}, the most favourable case being the combination of true NH and θ_{23} in the upper octant. The dependence of the median sensitivity on whether the true hierarchy is normal or inverted is shown to arise primarily from the NMH-θ_{23} degeneracy. As a related fact, Asimov-based simplified estimates are found to be significantly more pessimistic than median sensitivity estimates following the frequentist definition.

Outlook for the NMH determination and other physics studies with ORCA

After a week of sea operations, on October 27th, 2018, the main electro-optical cable (MEOC) of ORCA has been repaired and the first node of the infrastructure redeployed successfully. All tests indicate nominal operation. The deployment of ORCA detection units will start again shortly, and the full detector is planned to be operational within a few years ($\sim 2020-2021$). A determination of the neutrino mass hierarchy at 3σ (99.9% confidence level) should then be achievable by 2025 at the latest, under the most pessimistic assumptions on oscillation parameters. If the hierarchy is normal and $\sin^2\theta_{23}$ is large, as indicated by the latest results from accelerator experiments ($\sin^2\theta_{23} \sim 0.53-0.58$ [280, 281]), the results of this thesis show that a 5σ discovery may be in reach of ORCA within the same timescale.

Currently operating long-baseline experiments have already reported a preference for the normal mass hierarchy in the analysis of their data, respectively at the level of a 87% Bayesian posterior probability for T2K [280] and at 95% confidence level for NOνA [79] (updated to only 1.8σ with the first joint $\nu/\overline{\nu}$ fit [281]). Based on an extrapolation of these results, NOνA in particular may reach the 3σ level before 2024 [281]. However the results of such projections do depend largely on the assumed true parameter values and in particular on δ_{CP}, whereas ORCA is less affected by this dependence. The JUNO experiment, also targeting the NMH, is planned to start operating in 2021 and would likely reach $\overline{\Delta\chi^2} > 9$ after ~ 6 years of operation, depending on the level of external constraint included [282].

All in all, the NMH will almost certainly be determined at the 3σ level around the 2025 mark. Whether or not ORCA will be the first experiment to reach this sensitivity depends on the true values of oscillation parameters (NOνA is likely to be first mostly when assuming a normal hierarchy and $\delta_{\mathrm{CP}} = 3\pi/2$). The timeline of the completion of funding requirements and the technical capacity of deploying the full ORCA detector in compliance with the planned schedule will also be critical. In any case, a combined analysis of T2K, NOνA, ORCA and JUNO data by the end of the next decade might be the only possibility to reach a global 5σ determination of the NMH before the next generation of precision accelerator experiments (DUNE and Hyper-Kamiokande) start operating.

As a matter of fact, the physics reach of ORCA goes far beyond the NMH determination. As a first example, it is shown in this thesis that ORCA will provide very competitive measurements of the atmospheric oscillation parameters θ_{23} and Δm^2_{31}. As further oscillation-related topics, sensitivity analyses are underway within the KM3NeT collaboration regarding the measurement of tau neutrino appearance (testing the unitarity of the PMNS matrix), tests of the mixing of light sterile neutrino states with active neutrinos, and tests of the existence of non-standard neutrino interactions. For further details about the physics motivations and ORCA-specific sensitivities, see Refs. [251, 259] and references therein.

In all of the above studies, the software framework that I developed and presented in this thesis has been used either as main analysis framework or as secondary cross-check tool, thus demonstrating its flexibility and robustness.

Further studies, ongoing or under consideration, concern ORCA's sensitivity to quantum decoherence and Lorentz invariance violation effects, as well as astrophysics topics, such as supernovae and dark matter detection.

Earth tomography with ORCA and beyond

Being directly sensitive to the electron density, the measurement of matter-enhanced neutrino oscillations offers an interesting alternative to the standard geophysical techniques for probing the deep Earth. Assuming the radial mass density profile to be well constrained by geophysical methods, neutrino oscillation tomography can be used in principle to constrain the chemical composition of the deep mantle or outer core of the Earth, via the average proton-to-nucleon ratio (Z/A).

While the effect of the mass hierarchy can already be seen as subdominant with respect to the global, first-order picture of atmospheric neutrino oscillations, the effect of small variations of chemical composition is yet a higher-order small perturbation and is therefore challenging to observe. In fact, it will most likely not be possible to study these effects before the NMH is measured independently of ORCA. The sensitivity studies carried out in this thesis have shown that after ten years of operation, ORCA can achieve a measurement of the Z/A ratio with a precision of about $\pm 5\%$ in the lower mantle, and about $\pm 7\%$ in the outer core, assuming a normal mass hierarchy and maximal $2-3$ mixing – the inverted hierarchy case being significantly less favourable. While this level of precision does allow to demonstrate the potential of the technique and to exclude very exotic models, it will not be sufficient to distinguish between the more restricted range of compositional models considered as most realistic by geophysics.

In the outer core, the potential of neutrino oscillation tomography is of great interest in principle, due to the presence of light elements (which exact nature and amount remain undetermined) along with the base composition of iron-nickel alloy. Regarding this measurement, the ORCA sensitivity study presented here suggests clear paths for improvement. While ORCA has the advantage of maximising event statistics, a smaller detector with lower detection threshold and better energy-angular reconstruction capabilities may have a better potential to capture fast oscillatory patterns, which are the most intense in the few-GeV region.

An analysis of the sensitivity of next-to-next generation detectors for this measurement has thus been undertaken, with the aim to identify the most important characteristics to be improved. While the first results of this study have been presented at a major international geophysics conference [283], it was decided to not include them in the thesis, primarily focused on ORCA-related development and results. The study adapts the analysis framework presented here, implementing a generic but realistic parametrised model of detector response in place of the Monte Carlo based model. The model accounts for the energy-dependent detection efficiency and flavour identification capabilities (modeled by two-parameter activation functions), and energy-dependent resolution functions for neutrino energy and zenith angle. Three approximated models of future detectors are considered, loosely based on ORCA (see also the 'super-ORCA' proposal [284]), Hyper-Kamiokande, and DUNE. It is shown that the three options have comparable statistical sensitivity to the outer core Z/A. Hypothetical next-to-next generation detectors are considered, taking as figure of merit the ability to distinguish a compositional model containing 1 wt% hydrogen from a pure Fe-Ni composition. The impact of systematics, potentially more important at this level of statistical sensitivity than *e.g.* for ORCA, has not been evaluated yet.

Areas of improvement and continuation of this work

First of all, the sensitivity to the mass hierarchy and (θ_{23}, Δm^2_{31}) measurement presented at Neutrino 2018 are planned to be updated for publication in the coming months, accounting for a recent full simulation with adapted detector geometry (reduced horizontal spacing). These results will make use of the latest version of the analysis tools including the MC errors correction schemes, and would ideally exploit improved analysis strategies (*e.g.* N-class).

For continuation of systematics studies, an abvious area of improvement is the study of cross-section uncertainties, which has already started. In addition, the impact of generic flux shape systematics should be evaluated. The results of this thesis have shown that flux composition systematics are of prime importance. However, these systematics are modeled based on potentially overconservative assumptions. A more detailed treatment is planned to be implemented, based on recently developed simulation tools [90, 92]. In principle, the impact of flux systematics could be canceled for a large part by measuring the downgoing flux of neutrinos. While this is in principle challenging because of the very high downgoing atmospheric muon background, preliminary studies within the collaboration have produced promising results.

As to the analysis methodology, the practical MC error correction scheme deserves further investigation. As discussed in Chap. 8, some aspects of the current behaviour, for instance overcorrection effects, remain to be completely understood. The mathematical model of the sparse MC effect and extrapolation scheme presented in Appendix B has the potential to be extremely useful for a refined understanding of these subtle statistical aspects. Nevertheless, the theoretical results of Appendix B somewhat lack proper mathematical foundation and should be validated by toy Monte Carlo studies, whose implementation in the analysis framework has been started.

Other planned studies concern statistical methods and the comparison of Asimov-based and frequentist estimates, both for the NMH and the measurement of continuous parameters. The evaluation of the use of the more sophisticated analytical estimates derived in Ref. [156], and the comparison of frequentist and Asimov $\Delta\chi^2$-based (asymptotic) confidence intervals were planned to be included in this thesis, but have not been completed yet.

As mentioned already, the potential of N-class strategies will be further investigated, without neglecting associated systematics. The inclusion of Bjorken-y reconstructed information in the statistical analysis does not seem to bring a sizeable sensitivity improvement, contrarily to expectations. The reasons for this lack of positive result merit practical investigation. As a way of indicating the necessary level of reconstruction accuracy for developers of reconstruction strategies, the impact of an arbitrary Bjorken-y measurement performance could be evaluated

relatively easily in the framework implemented by the correlated smearing machinery. More generally, many other promising ideas can be investigated thanks to the analysis methodology developed in this thesis. For instance, a refinement of the N-class analysis strategy where events are also separated into quality-based samples depending on reconstruction quality parameters seems very promising.

Conventions and notations

Zenith angle

The incoming direction of a particle reaching a point M at the surface of the Earth is associated two angular coordinates θ_z (*zenith*) and φ (*azimuth*). θ_z (sometimes simply θ) is defined as the oriented angle between the local vertical $\overrightarrow{OM}$ and the particles's momentum $\vec{p}$:

$$\theta_z = (\widehat{\overrightarrow{OM},\ \vec{p}}) \tag{9.1}$$

where O is the center of the Earth (assimilated to a sphere). M typically represents the location of a detector. The cosine $\cos\theta_z$ is often used instead of θ_z.

Direction	θ_z	$\cos\theta_z$
vertical upgoing	π (180°)	-1
horizontal upgoing	$\pi/2$ (90°)	0
vertical downgoing	0 (0°)	1

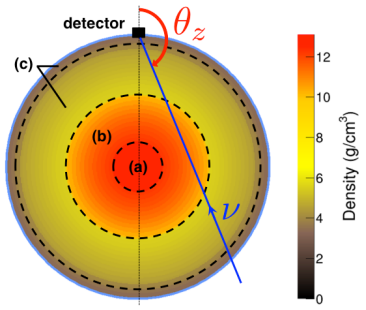

Units

A system of natural units is used so that $c = 1,\ \hbar = 1$.

Algebra and four-vectors

The customary notations of quantum mechanics are used for complex conjugates and related operations:

- z^* is the complex conjugate of the complex number z.
- $M^\dagger$ is the hermitian conjugate (h.c.) of a complex matrix M, *i.e.* $(M^\dagger)_{ij} = (M_{ji})^*$.
- M is hermitian if $M^\dagger = M$, this is the case of operators representing observables in quantum mechanics.
- M is unitary if $M^\dagger = M^{-1}$, this the case of matrices representing a change between orthonormal bases.
- $\overline{\psi}$ designates the *adjoint spinor* of the complex spinor ψ (see next section).

Regarding vectors and four-vector objects, following the usual notations of special relativity:

- Latin letters $(i, j...)$ are generally used for the indices of space coordinates $(1, 2, 3)$ while greek letters $(\mu, \nu...)$ are used for the indices of relativistic four-vectors $(0, 1, 2, 3$ with x^0 the time coordinate).

- A four-vector object may be referred to as a letter with a greek superscript, as in

$$x^\mu = (x^0, x^1, x^2, x^3) = (t, \vec{x}), \quad p^\mu = (E, \vec{p}).$$

- The Minkowski spacetime metric has signature $(-+++)$, *i.e.*

$$\eta_{\mu\nu} = \eta^{\mu\nu} = \mathrm{diag}(-1, 1, 1, 1).$$

  Lower indices are covariant coordinates:

$$x_\mu = \eta_{\mu\nu} x^\nu = (t, -\vec{x}),$$

  and repeated indices correspond to an implicit summation, as in

$$p^\mu x_\mu = \sum_{\mu=0}^{3} p^\mu x_\mu = \sum_{\mu=0}^{3} p^\mu \Big(\sum_{\nu=0}^{3} \eta_{\mu\nu} x^\nu \Big) = Et - \vec{p}.\vec{x}.$$

List of abbreviations

AMANDA Antarctic Muon And Neutrino Detector Array. p. 18
ANTARES Astronomy with a Neutrino Telescope and Abyss environmental RESearch. p. 18
ARCA Astroparticle Research with Cosmics in the Abyss. p. 86
BNL Brookhaven National Laboratory. pp. 13, 236
CC Charged current. pp. 17, 231
CCQE Charged current quasi elastic. p. 54
CERN European Organization for Nuclear Research (the acronym derives from French *Conseil Européen pour la Recherche Nucléaire*). p. 236
CMB Core-mantle boundary (in this thesis). p. 64
DIS Deep inelastic scattering. p. 54
DOM Digital Optical Module. p. 87
DU Detection Unit. p. 89
DUMAND Deep Underwater Muon And Neutrino Detector Project. p. 18
DUNE Deep Underground Neutrino Experiment. p. 45
eV electron-Volt. p. 9
Fermilab Fermi National Accelerator Laboratory (also referred to as FNAL). p. 236
IC Inner core. p. 73
IMB the Irvine-Michigan-Brookhaven detector. p. 14
Kamiokande Kamioka Nucleon Decay Experiment. p. 14
KM3NeT Kilometre Cube Neutrino Telescope. pp. 18, 46
LLR Log-likelihood ratio. p. 145
MC Monte Carlo. pp. 20, 92
MEOC Main electro-optical cable. p. 89
mwe meters water equivalent. pp. 14, 40
NC Neutral current. pp. 17, 232
NEMO Neutrino Ettore Majorana Observatory. p. 18
NESTOR Neutrino Extended Submarine Telescope with Oceanographic Research. p. 18
NMH Neutrino Mass Hierarchy. p. 33
OC Outer core. p. 73
ORCA Oscillation Research with Cosmics in the Abyss. pp. 46, 66
PDF Probability density function. p. 101
PE Pseudo-experiment. pp. 144, 146
PID Particle identification. p. 95
PINGU Precision IceCube Next Generation Upgrade. pp. 46, 66
PMT Photomultiplier tube. pp. 15, 86
PREM Preliminary Reference Earth Model. pp. 63, 74
QE Quasi elastic. p. 54
RES Resonant pion production. p. 54
SK Super-Kamiokande. p. 15
SM Standard Model. p. 26
ToT Time over threshold. pp. 88, 98
TS Test statistic. p. 144
TTS Transit time spread. p. 88

Glossary

Antiparticle In particle physics, every type of particle is associated an antiparticle which has the same mass and opposite physical charges, including the electric charge and other additive quantum numbers. The existence of antiparticles was predicted theoretically in 1931 (see Sec. A.2.2) and the the first antiparticle (the positron, associated to the electron) was discovered experimentally in 1935 in the study of cosmic rays. Some neutral particles, like the photon γ or the neutral pion π^0, are their own antiparticles. It is yet to be determined whether neutrinos are their own antiparticles, in which case they would be *Majorana* particles (see Sec. 1.2.1).

The CP symmetry (composition of the space parity P and charge conjugation C) fundamentally relates particles to antiparticles. As explained in Chap. 1, the CP symmetry could be violated in neutrino mixing. This is yet to be discovered. If confirmed, CP violation in neutrino mixing could have profound implications on our understanding of the Universe and its content (see baryogenesis).

Referenced on pp. 219, 230, 232

Baryogenesis Generic mechanism believed to be responsible for the observed predominance of matter over antimatter in the Universe. Baryogenesis should take place in the early Universe, as a dynamical mechanism generating a stable matter-antimatter asymmetry from the initially symmetric state predicted by the Big Bang theory. One of the necessary conditions for baryogenesis (*Sakharov criteria*) is the existence of C and CP violation in the laws of particle physics.

Regarding the C and CP symmetries, see antiparticle and parity. Regarding their violation by the weak interaction and neutrinos, see Sec. 1.2.1.

Referenced on pp. 43, 219

Boson Particle with integer spin ($S = 0,\ 1,\ 2...$), as opposed to fermions which have half-integer spin ($S = 1/2,\ 3/2...$). In the Standard Model, the only elementary bosons are the Higgs boson ($S = 0$) and the "force carriers" ($\gamma,\ Z^0,\ W^{\pm}$, gluons) which are vector bosons ($S = 1$).

Referenced on pp. 221, 224, 225, 232

Bremsstrahlung Electromagnetic radiation emitted by decelerating charged particles when deflected by other charged particles. Bremsstrahlung is the main cause of energy loss of relativistc electrons in matter. German word formed from *bremsen* (to brake) and Strahlung (radiation).

Referenced on p. 58

Bubble chamber Type of particle detector based on the same principle as a cloud chamber, except that it uses a superheated transparent liquid (generally liquid hydrogen or heavy liquids such as Freon) in which bubble of gas form around ionised atoms. Bubble chambers were invented 1952 and replaced cloud chambers in the second half of the century. Some of the early neutrino detectors were bubble chambers, like the Gargamelle experiment at CERN (1970-1979), famous for its dicovery of weak neutral currents. Today they are mostly replaced by wire chambers and spark chambers.

Referenced on p. 13

Chirality Geometrical property which is central to the description of neutrinos. A chiral object is an object that is not identical to its image in a mirror. A typical example are the human left and right hands: they are mirror images of each other, but are not identical objects (cannot be superposed onto one another). By analogy, the two possible versions of a chiral object are often called *left-handed* and *right-handed*. For instance, a right-handed screw is the usual version of a screw; a left-handed version would need one to invert the rotation of a screwdriver.

The chirality of molecules plays a very important role in biological processes. As to fundamental physics, the current theory of particle physics (see Standard Model) suggests that the laws of the

Universe are themselves chiral, *i.e.* the theory does not treat right-handedness and left-handedness identically (see parity violation). In particular, only left-handed neutrinos have ever been observed. Left-handedness corresponds, for massless neutrinos, to the fact that their spin is oriented in the opposite direction as their direction of motion.

Referenced on pp. 27, 223, 224, 240

Cloud chamber Particle detector used, mostly in the 1920's to 1950's, for visualizing the passage of ionizing radiation (*e.g.* charged particles). A cloud chamber is filled with water or alcohol in a gaseous *supersaturated* state, which means that it can condensate around ionised atoms resulting from the passage of radiation, forming a visible track of small droplets of condensated liquid. These tracks have characteristic shapes depending on the type of radiation (ionizing properties, charge, mass, *etc*).

Referenced on pp. 8, 219

Cosmic accelerator Hypothetical sources of the highly-energetic cosmic rays (CR) received on Earth. Because cosmic rays are charged particles, their trajectories are deflected by magnetic fields in space, thus the CR received on Earth cannot be traced back to their source. Examples of hypothesized CR sources are sources are supernovae, active galactic nuclei, gamma-ray bursts, *etc*. Some of the classes of CR sources are likely to produce high-energy neutrinos, which, contrarily to charged CR and electromagnetic radiation, can travel undisturbed on cosmological distances (see Sec. 1.1.2, Ref. [26, 27]).

Referenced on pp. 9, 220

Cosmic Neutrino Background Background particle radiation in the Universe formed by very low-energetic neutrinos (also known as *relic neutrinos*), as predicted by the *Big Bang* cosmological model (prevailing model describing the early evolution of the Universe as a rapid expansion from a very high-density and high-temperature initial state). Because of their very low energy, relic neutrino may never be detectable directly, but there is very strong indirect evidence of their existence.

Referenced on p. 9

Cosmic rays (CR) High-energy charged particles reaching the Earth, mainly originating outside the Solar System and even from distant galaxies. Part of the CR have extremely high energies, exceeding by far the highest energy reached on Earth at particle accelerators. The *primary* cosmic rays (those reaching the Earth's atmosphere) are mostly protons (∼99%), other nuclei (∼9%, dominantly helium), and solitary electrons (∼1%). When colliding with the Earth's atmosphere, cosmic rays can produce showers of secondary particles, some of them eventually able to reach the Earth surface (the *secondary* CR): mostly muons, and importantly to the subject of this thesis, *atmospheric neutrinos* (see Sec. 1.1.2). The discovery of cosmic rays dates back to 1912, marking the birth of particle physics. However, their origin is still largely unelucidated (see cosmic accelerator).

Referenced on pp. 9, 11, 13, 219, 220

Cross-section Physical quantity measuring the probability for two particles moving towards one another to interact. Cross-sections are expressed in units of area (*e.g.* cm^2). This can be understood by analogy to the collisions of hard spheres which interact only upon contact: their scattering cross section is related to their geometric size.

Referenced on pp. 9, 51

Electronvolt (eV) Unit of energy widely used in atomic and subatomic physics. It corresponds to the kinetic energy acquired by an electron accelerated by an electric difference of potential of one volt (V). In SI units of energy (joules), $1\,\mathrm{eV} = 1.6 \cdot 10^{-19}\,\mathrm{J}$.

In particle physics, the system of natural units is generally used, so that the mass of a particle is identified to its mass energy (see mass-energy equivalence). Therefore, in this thesis the electronvolt and its multiples are used to refer to the mass of particles. For instance, the mass energy of an electron is about half a MeV ($10^6\,\mathrm{eV}$); of a proton, $\sim 1\,\mathrm{GeV}$ ($10^9\,\mathrm{eV}$); of a neutrino, less than 1 eV. The highest known single-particle kinetic energy are reached on Earth at particle accelerators ($\sim \mathrm{TeV} = 10^{12}\,\mathrm{eV}$) and in space in the cosmic rays ($\sim 10^{21}\,\mathrm{eV}$). For comparison, photons of visible light have a few eV energy; the binding energy of electrons in atoms is of the order of 10 eV; the required energy (heat) to vaporise 1 kg of liquid water into steam is about $10^{22}\,\mathrm{eV}$ (2.3 kJ).

Referenced on pp. 217, 222, 223

Fermion Particle with half-integer spin ($S = 1/2,\ 3/2...$), as opposed to bosons which have integer spin ($S = 0,\ 1,\ 2...$). All elementary matter particles in the Standard Model are fermions with spin 1/2. This includes quarks, charged leptons ($e^{\pm}$, $\mu^{\pm}$, $\tau^{\pm}$), and neutrinos (ν_e, ν_μ, ν_τ and antineutrinos) – see Sec. 1.1.1.

Referenced on pp. 7, 219, 221, 224, 230

Flavour Type of elemtary particle, as "seen" by the weak interaction. Strictly speaking, in the Standard Model leptons thus come in six flavours (three charged leptons and three neutrinos), and quarks come in six flavours as well. However, the notion of flavour for leptons is often understood in neutrino physics as englobing the charged lepton and corresponding neutrino of a family, *i.e.* ν_e is said to be of "electron flavour".

Referenced on p. 234

Generation Family of elementary fermions formed by two quarks, one charged lepton and one neutrino – plus their antiparticles. Currently exactly *three* generations have been found to exist in Nature. Ordinary matter is made of particles of the first generation only (quarks u and d, electron e^-). Schematically, the second and third generation are "copies" of the first one, except that the corresponding fermions are heavier and, as a consequence, are unstable particles. See Sec. 1.1.1 (in particular Tab. A.2) for details and examples. The number of fermion generations is, as of today, a fully empirical ingredient of the Standard Model.

Referenced on pp. 224, 233

Hadron Composite particle made of an aggregate of quarks. Examples: protons (uud), neutrons udd, neutral pion π^0 ($u\overline{u}$), charged pions π^+ ($u\overline{d}$) and π^- ($\overline{u}d$). The only genuinely stable hadrons are protons. Neutrons are stable only when bound with protons in nuclei, otherwise a free neutrino decays to a proton via β-decay, with a mean lifetime of about 15 min.

Other hadrons are unstable. As an example, the positively charged pion π^+ is one of the longest-lived hadrons with a mean lifetime $\tau \simeq 2.6 \cdot 10^{-8}$ s; it decays mostly into an antimuon and a muon neutrino:

$$\pi^+ \to \mu^+ + \nu_\mu.$$

The neutral pion π^0 can decay electromagnetically into two photons (99.8%) and therefore has a much shorter mean life of $\sim 10^{-17}$ s. Hadrons containing quarks of the second and third generations are extremely unstable and only observed at particle accelerators.

Among hadrons, those formed by three quarks (like protons and neutrons) are called *baryons*. Those formed by one quark and one antiquark are called *mesons*.

Referenced on pp. 223–225, 235, 236

Ionisation chamber Gas-filled radiation detector in which an electric field (voltage) is applied in order to measure the number of ion pairs created by incident ionising radiation.

Referenced on p. 8

Lepton Elementary particle (fermion) which does not take part in the strong interaction. Leptons include the *charged leptons* (e^-, μ^-, τ^- and antiparticles) and the neutrinos (ν_e, ν_μ, ν_τ and antiparticles). The other category of elementary fermions are the quarks.

Referenced on pp. 221, 224, 230

Lie algebra Mathematical object used in quantum field theory, in relation with a Lie group, to translate physical symmetries into properties of the quantum world.

For instance, consider the assumption that physical laws are the same everywhere in the Universe and do not evolve with time, *i.e.* they are invariant by space-time translations. In the quantum world this assumption becomes a set of commutation properties of the *generators* of these transformations: the hamiltonian $\hat{H}$ and the momentum operators $\hat{P}_x$, $\hat{P}_y$, $\hat{P}_z$. In this correspondance, space-time translations form a *Lie group*, and the commutation properties are determined by the structure of the corresponding *Lie algebra* (which is, in this case, very simple).

In mathematics, an *algebra* is a set in which elements can be added together and multiplied by constants (as in a *vector space*), but also multiplied between themselves. Formally, the vector space tangent to a Lie group G at its unit element (identity) I has a natural structure of Lie algebra $\mathcal{A}$, inherited from the (infinitesimal) multiplicative structure of G. In addition, the basis vectors

(ω_k) of the algebra are (infinitesimal) *generators* of the group, in the sense that (if G is simply connected) any $g \in G$ can be decomposed as $g = \exp(I + i \cdot \sum_k x_k \omega_k)$ (see *e.g.* Ref. [41], Chap. 2).

In the previous example, the action of time translations onto the quantum fields is obtained by exponentiating the action of the (infinitesimal) generator $\hat{H}$:

$$\boldsymbol{\psi}(t, \vec{x}) = \exp\big[-i(t - t_0)\hat{H}\big] \cdot \boldsymbol{\psi}(t_0, \vec{x}).$$

and similarly for the space translations with generators $\hat{P}_x$, $\hat{P}_y$, $\hat{P}_z$.

Referenced on pp. 222, 240

Lie group Mathematical object used in physics to represent continuous symmetries.

As an example, the set of all possible rotations of a three-dimensional object is

- a *group* because the composition of two rotations is a rotation and each rotation has an inverse rotation;
- a *Lie group* because each rotation can be parametrised (by specifying two angles) and this parametrisation is "smooth", *i.e.* differentiable at will.

Lie groups are generally considered in physics through their *action* on vector spaces, *i.e.* the elements of the group are seen as linear transformations, for instance of the quantum states. Specifying a certain vector space and an action of a group on this vector space is called a *group representation*.

In quantum field theory, Lie groups and the mathematically related Lie algebras are used to translate physical symmetries into structural properties in the "quantum world" – see an example in the Lie algebra entry.

Formally, Lie groups are groups equipped with a structure of smooth manifold, *i.e.* they can locally be $\mathcal{C}^\infty$-mapped to subsets of $\mathbb{R}^N$, with compatibility of the differential structure (maps) and group structure (multiplication and inversion). Simple examples of Lie groups which often appear in physics are matrix subgroups of $\mathbf{GL}_n(\mathbb{K})$ ($\mathbb{K} = \mathbb{R}$ or $\mathbb{C}$) such as $\mathbf{SO}(n)$, $\mathbf{U}(n)$, $\mathbf{SU}(n)$. The Lie algebra $\mathcal{A}$ associated to a Lie group G is the vector space tangent at the unit element (identity) of G, which inherits the (infinitesimal) multiplicative structure of G, and essentially contains "all" the information on G.

Referenced on pp. 221, 239, 240, 243

Mass-energy equivalence In nuclear processes and interactions of elementary particles at high-energy, mass can be converted to energy and conversely.

This was expressed theoretically in Einstein's theory of relativity (1905), with the notorious formula $E = mc^2$ where E is the energy, m the mass, and c is the velocity of light.

For instance, in a nuclear β-decay

$$^{\mathrm{A}}_{\mathrm{Z}}\mathrm{X} \to {}_{\mathrm{Z}+1}^{\mathrm{A}}\mathrm{Y} + e^- + \overline{\nu}_e,$$

the difference of mass energy between the inital and final state is positive, it is converted into kinetic energy of the electron e^- and neutrino $\overline{\nu}_e$.

As an example of the opposite conversion, in the vicinity of a nucleus a high-energy γ photon (massless) can be converted to massive particles, *e.g.* a pair of electron e^- and positron (antielectron) e^+:

$$\gamma \to e^- + e^+$$

This process is called *pair production*.

See also electron-Volt (eV).

Referenced on pp. 58, 220

Natural units System of units used in theoretical and high-energy physics to simplify calculations by setting the omnipresent physical constants c, $\hbar$, *etc* to 1. For instance, the expression of the relativistic energy of a free particle:

$$E^2 = |\vec{p}|^2 c^2 + m^2 c^4$$

becomes:

$$E^2 = |\vec{p}|^2 + m^2$$

i.e. the energy, momentum and mass all have the same dimension. They are generally expressed in multiples of the electronvolt (eV).

Referenced on pp. 215, 220

Parity Mirror symmetry of space, which exchanges right-handed and left-handed objects (see chirality). Mathematically, parity P is defined as the inversion of space coordinates:

$$P: \ (t, \vec{x}) \to (t, -\vec{x}) \tag{9.2}$$

In three-dimensional space and from the point of view of physics, parity can be seen equivalently as the symmetry transformation of a usual plane mirror (inversion of a single space coordinate), because the two transformations are related by a rotation.

Referenced on pp. 27, 219

Parity violation The laws of physics were believed to be invariant under the parity symmetry, *i.e.* treat chiral left-handed or right-handed objects symmetrically, until it was discovered (1957) that the weak interaction is not. As a related fact, only left-handed neutrinos (and right-handed antineutrinos) have ever been observed.

Referenced on p. 220

Parton model A model of hadronic structure as a composition of point-like constituents called *partons*, used to describe high-energy hadron collisions. The parton model is a good approximation to describe *deep inelastic scattering*, the dominant process of multi-GeV neutrino-nucleon interactions (Sec. 2.1.2).

In quantum chromodynamics, the correspondence of a parton with particle states of quarks and gluons depends on the length scale under consideration. In neutrino-nucleon interactions, the relevant length scale is related to the inverse of the momentum transfer. For large momentum transfer, the *parton sea* description becomes relevant, where for instance a proton contains short-lived quark-antiquark pairs and gluons in addition to its valence quarks *uud*.

See also quark, strong interaction, hadron.

Referenced on pp. 52, 54

Photomultiplier tube (PMT) Extremely sensitive light detectors, able to detect single photons via the photoelectric effect and the avalanche multiplication of the photoelectric current produced by incident light. A description of the broad principle of photomultiplier tubes is given in Sec. 3.1.1.

Referenced on pp. 11, 217

Quantum eigenstate In quantum mechanics, the *state* of a system is represented by an element $|\psi\rangle$ of a vector space V. A physical observable is represented by an operator $\widehat{O}$ which acts on the states $|\psi\rangle$, and the states that can be observed are the *eigenvectors* $|\psi_i\rangle$ of $\widehat{O}$, *i.e.* they are such that:

$$\widehat{O}\,|\psi_i\rangle = \lambda_i \psi_i$$

The values λ_i (called in mathematics the *eigenvalues* of $\widehat{O}$ associated with each eigenvector $|\psi_i\rangle$) form the set of values that can be obtained when measuring the physical quantity O. This set is generally discrete, as in the example of the *quantised* energy levels of an atom.

In quantum mechanics, a physical state $|\psi'\rangle$ can be any linear superposition of eigenstates, for instance (with α_1, α_2 some numbers)

$$|\psi'\rangle = \alpha_1\,|\psi_1\rangle + \alpha_2\,|\psi_2\rangle\,,$$

but when measuring O, the experimenter will observe either α_1 (in which case the state of the system will become $|\psi_1\rangle$) or α_2 (in which case the state of the system will become $|\psi_2\rangle$). This rule, called *wave function collapse*, is perhaps both the most "defining" and the most intriguing feature of quantum mechanics.

Referenced on pp. 222, 236

Quantum electrodynamics (QED) Relativistic quantum field theory of electrodynamics, describing the interactions between charged particles (matter) and electromagnetic radiation (light) via the exchange of photons. QED was the first theory to achieve a full agreement between quantum mechanics and special relativity, and successfully predicts physical quantities in atomic and particle physics with extremely high accuracy. Mathematically, it is formulated as a *gauge theory* (see Appendix A, Sec. A.2.3), and largely inspired the formulation of modern particle physics theory.

Referenced on pp. 7, 8, 224, 234, 244

Quantum field theory (QFT) The mathematical framework for the theoretical description of particle physics. One way to introduce QFT is as an extension of the original formulation of *quantum mechanics* so that it accomodates the concept of causality according to *special relativity*.

The fundamental objects of QFT are *fields*, which are simply functions of space and time. Fields are generally not restricted to a finite volume. In fancy words, they are said to "pervade space and time". They can be scalar (*i.e.* with one component), like the simplest wave function of a particle in quantum mechanics. They can be four-vectors (with one time-like and three space-like coordinates) like the electromagnetic potential in the relativistic formulation of Maxwell's theory. Or, as introduced in Sec. A.2.2, they can be of other kinds.

In QFT, fields are promoted to *operators* which act on *particle states*. The general idea is that fields can annihilate and create particles, based on a fundamental *vacuum state* $|0\rangle$. Particles and antiparticles can then be described as *quanta* of the fields, *i.e.* excitations "over" the vacuum state. Quantised fields are the basis for the description of the subatomic world, not limited to the relativistic models used in high-energy particle physics. They are also extensively used in condensed matter physics.

Referenced on p. 239

Quark Elementary particle. Spin-1/2 fermion, characterised by the fact that it takes part in the strong interaction. In everyday matter, quarks are the primary constituents of the protons and neutrons which form the nuclei of atoms. More generally quarks do not exist as free particles and are always encountered under the form of aggregates of two, three, or more quarks. Such composite particles are called hadrons. Six types of quarks (plus antiparticles) are accounted for in the Standard Model and are experimentally observed:

- up (u) and down (d) in the first generation
- charm (u) and strange (s) in the second generation
- top (t) and bottom (b) in the third generation

Up-type quarks (u, c, t) and antiquarks ($\overline{u}$, $\overline{c}$, $\overline{t}$) have an electrical charge $Q = \pm 2/3$ in units of the elementary charge e, while down-type quarks (d, s, b) and antiquarks ($\overline{d}$, $\overline{s}$, $\overline{b}$) have $Q = \mp 1/3$. The other category of elementary fermions are the leptons.

Referenced on pp. 221, 223, 225

Scintillator Material that can absorb the energy of ionizing radiation (for instance, a high-energy γ photon) and re-emit this energy in the form of light.

Referenced on p. 11

Spin Fundamental, characteristic property of elementary particles. Spin schematically relates to a form of intrinsic angular momentum (self-rotation). The concept of spin is central in the mathematical formulation of particle physics theory. Moreover the relation between the spin and the collective behaviour of particles (in particular for atomic and molecular electrons) has profound consequences on many areas of physics and on the foundations of chemistry.

The notion of *helicity* is related to spin: it corresponds to the relative orientation of the spin of a particle with respect to its direction of motion. A classical analogy is a screw turning clockwise (usual screw) or counterclockwise (mirror screw).

See also Fermion, Boson, Vector boson, Chirality.

Referenced on pp. 7, 219–221, 225

Standard Model (SM) of particle physics Model which describes the known elementary particles and their interactions, encompassing three of the four fundamental forces into a common theoretical description: the electromagnetic, weak, and strong interactions. In particular, the electromagnetic

and weak interactions are unified in a single *electroweak* theory which is broadly described in Sec. A.2. The term "Standard Model" is sometimes used to refer specifically to the electroweak theory.

Referenced on pp. 8, 26, 217, 219, 221, 224

Strong interaction Short-range interaction which applies only to quarks and, indirectly, to hadrons (composite particles, such as protons and neutrons, formed by aggregates of quarks). The *strong nuclear force*, which tightly binds together protons and neutrons in nuclei, is a residual effect of the strong interaction undergone by their constituents, the quarks. On nuclear distances, the strong interaction is much more intense than the electromagnetic interaction, while it becomes negligible on larger distances in ordinary matter. The intermediate vector bosons mediating the strong interaction are called *gluons* – as described by *quantum chromodynamics* (see A.2).

Referenced on pp. 8, 9, 221, 223, 224, 251

Vector boson Boson with spin $S = 1$. Vector bosons are designated as such because their mathematical representations transform as vectors under a change of inertial reference frame. In the Standard Model the "force carriers" are all vector bosons: photon γ, weak force carriers Z^0 and $W^{\pm}$, and gluons. Such mediating vector bosons are also called *gauge bosons*, since they appear mathematically as a result of requiring local *gauge invariance* of the interactions.

Referenced on pp. 219, 224, 232

Weak interaction Also called the *weak nuclear force*, this is the interaction responsible for nuclear decays. The weak interaction affects all elementary particles, and is the only interaction affecting neutrinos (apart from gravitation). It is both very short-ranged and much less intense that electromagnetic and strong interactions. Three intermediate vector bosons can be involved in mediating the weak interaction: the charged W^+ and W^- bosons, and the neutral Z^0 boson. All of them are massive, as a result of the *Higgs mechanism* (see A.2).

Referenced on pp. 8, 221, 223, 224

Appendices

Appendix A

Introduction to neutrinos and the weak interaction

Contents

This appendix provides in Sec. A.1 an introduction to neutrinos and other flavours of elementary particles intended for the non-specialists. The reader should refer as often as needed to the glossary, where terms are introduced without assuming an advanced knowledge of physics and mathematics.

In Sec. A.2, some elements of the Standard Model of particle physics are introduced, with an emphasis on the concepts needed to discuss neutrino masses and the bases of the description of the electroweak interaction. The intent is to adopt a didactic approach rather than an exhaustive one. Some (even remote) previous kowledge of quantum mechanics and special relativity is however preferrable to read Sec. A.2.

A.1 Flavours of particles: an elementary introduction

The elementary fermions described in the Standard Model are presented, starting with the first generation which constitutes ordinary matter. A phenomenological description of the weak interactions of neutrinos is then given. The second and third generation of leptons and quarks are then introduced. Finally the mixing of quarks and leptons are discussed.

A.1.1 First generation of quarks and leptons

Tab. A.1 lists the first *generation* (or *family*) of elementary matter particles. The first generation includes the electron and electron neutrino (the leptons, as well as the quarks *up* (u) and *down* (d) which make up the proton and neutron. All these particles are fermions with spin 1/2.

In a simplified approach, the proton can be described as an assembly of two up quarks and one down quark (uud), and the neutron as an assembly of one up quark and two down quarks (udd). Besides the electrical charge Q (expressed in units of the fundamental charge $e = 1.6 \cdot 10^{-19}\,\mathrm{C}$), another additive quantum number (which sum is conserved in any interaction) is indicated in Tab. A.1: the lepton number L (also known as "leptonic charge"). Leptons of the first generation have a lepton number $L = 1$, quarks have $L = 0$.

In addition, to each elementary particle is associated a "mirror" antiparticle which has the same mass and opposite electrical charge and lepton number: the positron e^+ is thus the antiparticle of the electron, and the electron antineutrino $\overline{\nu}_e$ is, rather obviously, the antiparticle of the electron neutrino. Likewise the antiquark up $\overline{u}$ and antiquark down $\overline{d}$ mirror the quarks u and d.

A.1.2 Weak interaction processes of neutrinos

As mentioned in the historical introduction (Sec. 1.1.1), the electron antineutrino ν_e is produced in association with electrons in β^--decays (the usual β-decays):

$$ {}^{\mathrm{A}}_{\mathrm{Z}}\mathrm{X} \rightarrow {}^{\mathrm{A}}_{\mathrm{Z+1}}\mathrm{Y} + e^- + \overline{\nu}_e, \tag{A.1} $$

which is seen as the decay of a neutron inside the nucleus:

$$ n \rightarrow p + e^- + \overline{\nu}_e. \tag{A.2} $$

In modern terms it corresponds to the decay of a down quark d with charge $-1/3$ into an up quark u with charge $+2/3$:

$$ d \rightarrow u + e^- + \overline{\nu}_e \tag{A.3} $$

Notice how both the total electrical charge and total lepton number are conserved in these decays.

	Particle	Q	L	Antiparticle	Q	L
quarks	u	+2/3	0	$\overline{u}$	-2/3	0
	d	-1/3	0	$\overline{d}$	+1/3	0
leptons	e^-	-1	+1	e^+	+1	-1
	ν_e	0	+1	$\overline{\nu}_e$	0	-1

Table A.1: List of elementary quarks and leptons of the first generation with their antiparticles. Q designates the electrical charge quantum number ($Q = +1$ corresponds, in SI units, to the fundamental electrical charge of $e = 1.6 \cdot 10^{-19}\,\mathrm{C}$). L designates the lepton number (dimensionless).

Similarly, the electron neutrino ν_e is associated with positrons in β^+-decays (more rarely observed):

$$ {}^{\mathrm{A}}_{\mathrm{Z}}\mathrm{X} \rightarrow {}_{\mathrm{Z}-1}{}^{\mathrm{A}}\mathrm{Y} + e^+ + \nu_e. \tag{A.4} $$

Likewise, this is reinterpreted as

$$ p \rightarrow n + e^+ + \nu_e, \tag{A.5} $$

or at the level of elementary particles as

$$ u \rightarrow d + e^+ + \nu_e. \tag{A.6} $$

A few additional comments are needed on the decays of Eqs. (A.4) to (A.6):

- the decay of neutrons into protons is observed for free neutrons (with a half-life of approximately 15 minutes), while the decay A.5 does not happen spontaneously to a free proton (which is lighter than a neutron, and believed to be a stable particle);
- β^+-decays can occur spontaneously in heavy proton-rich nuclei, or it can occur in a stimulated manner in lighter nuclei, *e.g.* when ${}^4_2\mathrm{He}$ is bombarded with α particles (this is the way it was discovered by Frédéric and Irène Joliot-Curie in 1934);
- the quark decay process A.6 is also involved for instance in proton-proton fusion (*e.g.* inside the Sun).

In the electroweak theory, the decays of Eq. A.3 and A.6 are not described, as in Fermi's theory, by a direct, four-component interaction. Instead they are described as mediated by a (charged) *vector current*. The interaction can be seen as the propagation of an additional elementary particle, interacting on one hand with the quarks u and d, and on the other hand with the leptons e^- and $\overline{\nu}_e$ (resp. e^+ and ν_e). This is represented under the form of Feynman diagrams in Fig. A.1 – the next paragraph provides a short explanation of the interpretation of such diagrams.

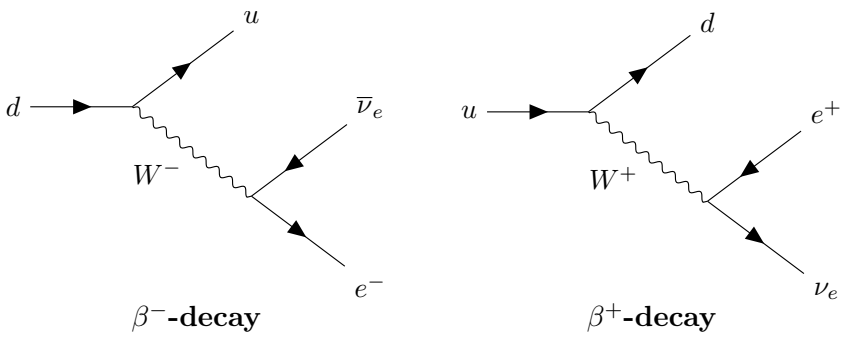

Figure A.1: Feynman diagrams[1] for the weak charged current (CC) processes corresponding to the β^- and β^+ decays.

Note on Feynman diagrams Feynman diagrams are used in particle physics to represent mathematical calculations of the *transition amplitudes* of particle interactions, which essentially describe how likely it is that a certain interaction (involving a set of incoming particles and a set of outgoing particles) occurs with a certain exchange of energy and momentum between them. Every element in a Feynman diagram (lines and vertices) has a precise mathematical meaning, which will not be described in detail here – see for example Ref. [41, 286, 287]. For the

[1] All Feynman diagrams in this document have been produced using the `TikZ-Feynman` package [285].

purpose of this thesis, it is sufficient to understand the following rules giving the interpretation of a diagram:

- the time evolution is read from left to right, *i.e.* the incoming particles (*initial state*) are represented on the left and outgoing particles (*final state*) on the right;
- fermions are represented by straight lines with arrows;
- in fermion lines, backard arrows are used to represent anti-fermions (the reason for this choice is briefly explained in Sec. A.2.2 and in the glossary.)
- wavy arrows represent intermediate vector bosons
- at each vertex (point connecting fermion and/or boson lines), the conservation of quantum numbers (electrical charge Q, lepton number L) must be enforced

More precisely, each diagram corresponds to a term in a perturbative expansion of an abstract object, the *S-matrix*, representing the "whole" interaction by connecting a given *initial state* and *final state*. In this thesis, the Feynman diagram we associate to a given interaction is only the first order term in the perturbative expansion (the so-called *tree level* diagram). Feynman diagrams, since they connect the canonical formulation of quantum field theory to measurable properties of particle scattering processes, are considered by many authors to be as good a representation of physical reality – see for example the introduction of Ref. [288] (available online).

As seen in Fig. A.1, in the case of charged current weak interactions of Eq. A.3 and A.6, the mediating particle is respectively a W^+ or a W^- boson. Other vector bosons in the electroweak theory are the neutral Z^0 boson, which mediates *neutral current* (NC) weak interactions, and the photon γ which mediates electromagnetic interactions. The $W^\pm$ and Z^0 bosons are massive, while the photon is massless: this has consequences onto the relative range of the electromagnetic (long-ranged) and weak (short-ranged) interactions.

To complete the picture of intermediate vector bosons, those mediating the strong interaction are called the *gluons*. Generally, intermediate vector bosons are also called *gauge bosons*, which is related to the fact that the SM is formulated as a *gauge theory* (see Sec. A.2).

Finally, let us illustrate further the weak interaction with processes which are involved in experimental neutrino physics and mentioned in the main body of this thesis. The Feynman diagram of the β-decay can be "rearranged" into several other conversion and particle scattering processes. For instance, *inverse beta decay (IBD)* corresponds to the conversion of a quark up to a quark down (converting a proton to a neutron) through the absorption of an electron antineutrino and emission of a positron:

$$u + \overline{\nu}_e \rightarrow d + e^+ \tag{A.7}$$

As seen in Sec. 1.1.2, this was the channel used for the very first detection of the neutrino in the Reines and Cowan experiment. Neutrino scattering processes include the charged current elastic scattering (ES) on an atomic electron:

$$\nu_e + e^- \rightarrow e^- + \nu_e, \tag{A.8}$$

Neutrinos and antineutrinos also scatter on nuclei via the charged current quasi-elastic (CCQE) scattering:

$$\nu_e + n \rightarrow e^- + p \tag{A.9}$$

$$\overline{\nu}_e + p \rightarrow e^+ + n \tag{A.10}$$

The above CCQE-like process is the dominant interaction channel for ν's and $\overline{\nu}$'s in the energy range $0.1 - 1$ GeV. As detailed in Sec. 2.1.2, neutrino-nucleus scattering processes at higher energy can become more complex and generally produces more particles in the final state.

Fig. A.2 gathers the tree-level Feynman diagrams corresponding to all charged current processes mentioned up to now (β-decay, IBD, ES, CCQE).

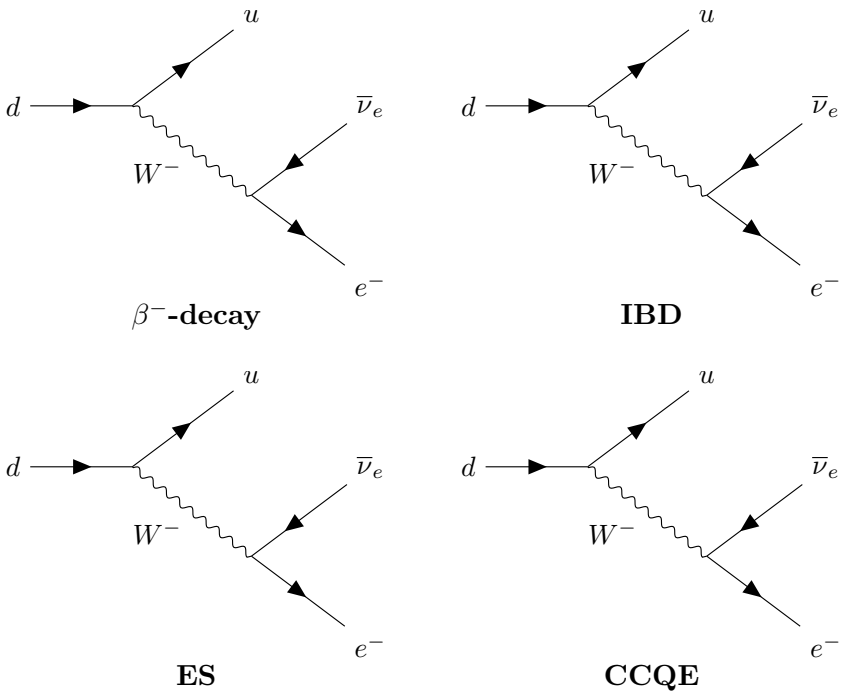

Figure A.2: Examples of weak charged current (CC) processes.

In addition, neutrinos and antineutrinos can interact through neutral current processes, mediated by the neutral Z^0 boson.

A.1.3 Second and third generation

Until now we have described only the first generation of elementary fermions. Two additional generations (or "families") of fermions are known to exist in Nature. As listed on Tab. A.2, each additional generation is formed by two quarks, one charged lepton and one neutrino – plus their antiparticles.

	First generation	Second generation	Third generation
quarks	u (up) d (down)	c (charm) s (strange)	t (top) b (bottom)
leptons	e^- ν_e (electron neutrino)	μ^- (muon) ν_μ (muon neutrino)	τ^- (tau lepton) ν_τ (tau neutrino)

Table A.2: Elementary fermions of the first, second, and third generations. To each one is associated an antiparticle. Antiparticles are not indicated in the table.

These particles have very similar properties to their counterparts in the first generation. In particular, they have the same electrical charge and are subject to the same fundamental interactions. The essential difference between successive generations is that particles in generation $N+1$ are generally significantly heavier than their counterparts in generation N. As an example

the mass of the charged leptons are:

$$m_{e^-} \simeq 0.511\,\text{MeV} \tag{A.11}$$
$$m_{\mu^-} \simeq 106\,\text{MeV} \tag{A.12}$$
$$m_{\tau} \simeq 1.78\,\text{GeV} \tag{A.13}$$

i.e. the mass of charged leptons increases by at least one order of magnitude between generations.

A similar pattern is found in the quark sector: the lightest one (up quark) has a mass of about $1.5-4$ MeV and the heaviest (top quark) a mass of about $170-180$ GeV (five orders of magnitude larger). Note that the masses of quarks are measured less precisely, and do not even have a univocal definition since free quarks do not exist. The case of neutrino masses is complicated as well (though for different reasons) and is detailed in Sec. 1.2. Let us simply note for now that their masses are known to be extremely small ($< 10\,\text{eV}$).

As a consequence of being more massive, fermions of the second and third generation are unstable, and tend to decay, shortly after being created, into fermions of a "lower" generation (and ultimately into fermions of the first generation). For instance, the muon μ^- (resp. antimuon μ^+) has a mean lifetime of about $2.2 \cdot 10^{-6}\,s$ and decays to an electron (resp. positron), one neutrino and one antineutrino:

$$\mu^- \to e^- + \overline{\nu}_e + \nu_\mu \tag{A.14}$$
$$\mu^+ \to e^+ + \nu_e + \overline{\nu}_\mu \tag{A.15}$$

These naturally occuring decays are formally very similar to the β-decays seen previously. They also correspond to weak charged current processes, as represented by the Feynman diagrams of Fig. A.3.

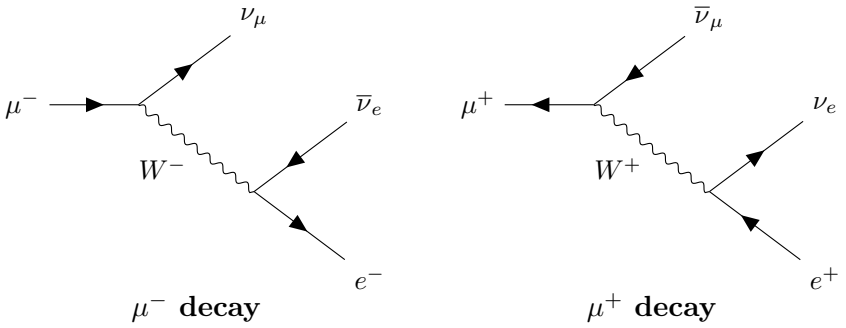

Figure A.3: Feynman diagrams for the weak charged current processes corresponding to the decay of a muon (left) and an antimuon (right).

While other decays (for example $\mu^\pm \to e^\pm + \gamma$) are technically allowed by kinematics and by the interaction rules of the electromagnetic interaction (quantum electrodynamics), the decays of Eqs. (A.14) and (A.15) are the only ones to be observed experimentally. This empirical fact is incorporated in the SM by introducing *flavour*[2] *lepton numbers* L_e ("electron number"), L_μ ("muon number"), and L_τ ("tau number") for leptons of the first, second and third generation respectively. As is summarised in Tab. A.3, in the family α ($\alpha = e$, μ or τ), leptons have $L_\alpha = +1$ and antileptons have $L_\alpha = -1$; leptons of the other families have $L_\alpha = 0$. With these

new definitions, the global lepton number seen previously is simply

$$L = L_e + L_\mu + L_\tau \tag{A.16}$$

	L_e	L_μ	L_τ		L_e	L_μ	L_τ
(ν_e, e^-)	+1	0	0	$(\overline{\nu}_e, e^+)$	-1	0	0
(ν_μ, μ^-)	0	+1	0	$(\overline{\nu}_\mu, \mu^+)$	0	-1	0
(ν_τ, τ^-)	0	0	+1	$(\overline{\nu}_\tau, \tau^+)$	0	0	-1

Table A.3: Assignment of the lepton flavour number for elementary fermions and antifermions of the first, second, and third generations.

Each individual flavour number is conserved in lepton decays. For instance, note how the total electron number is $L_e = 0$ in both decays A.14 and A.15, while the total muon number is $L_\mu = +1$ in the first decay and $L_\mu = -1$ in the second decay. An electromagnetic decay like $\mu^- \to e^- + \gamma$ would violate the conservation of L_e and L_μ.

The conservation of the individual lepton flavour number in the SM is a consequence of the mathematical representations chosen for the fermion fields, which are themselves based on the empirical observation that flavours don't "mix" in lepton decays. As explained in Sec. A.1.4, neutrino oscillations do violate the conservation of the flavour lepton number, thus requiring an extension of the Standard Model.

Experimental discovery of the second and third fermion generations

Historically, the muon was the first fermion beyond the first generation to be discovered: it was observed in the study of cosmic rays in 1936. It was however understood only in the late 1940's, when fundamental investigations were resumed after the Second World War, that muons were no more than "heavier electrons" and thus took no part in the strong interaction. This was clarified by the separate observation of the decays of *pions* into muons, and of muons into electrons [289].

This led Bruno Pontecorvo to propose the universality of the Fermi interaction in 1947, involving pairs $(e,\ \nu)$ and $(\mu,\ \nu)$ [290]. The hypothesis of an intermediate boson $W^\pm$ was formulated by Lee *et al.* in 1949 [291]. As mentioned before, the neutrino was first observed experimentally in 1953. The conservation of the global lepton number L was proposed in the same year and that of an individual flavour number L_α was introduced somewhat later to explain that $\mu^\pm \to e^\pm + \gamma$ was not observed.

On the basis of lepton flavour conservation and assuming that ν_μ and ν_e are different particles, the process

$$\nu_\mu + n \to p + e^- \tag{A.17}$$

is forbidden. Instead the charged-current interaction of ν_μ can only produce muons:

$$\nu_\mu + n \to p + \mu^- \tag{A.18}$$

This was demonstrated in the BNL-Columbia accelerator experiment in 1962 [14]. It used ν_μ produced by the decays of pions obtained at a particle accelerator (see Sec. 1.1.2):

$$\pi^+ \to \mu^+ + \nu_\mu \tag{A.19}$$

[2] Strictly speaking, the flavour designs a type of particle (as "seen" by the weak interaction), thus in the SM leptons come in six flavours (three charged leptons and six neutrinos), and quarks come in six flavours as well. However, the notion of flavour for leptons is sometimes interpreted as englobing the charged lepton and corresponding neutrino of a family (*i.e.* ν_e is said to be of "electron flavour").

[2]The pion or "π *meson*" is presented in the glossary entry for hadrons.

Thus the existence of the muon neutrino was established.

The leptons of the third generation were discovered much later: the τ lepton was observed in 1975 in an electron-positron collider experiment at the Stanford Linear Accelerator Center (SLAC), and the τ neutrino was discovered in 2000 at Fermilab (the Fermi National Accelerator Laboratory, also referred to as FNAL). Note, however, that the number of *active* neutrinos (interacting with the Z^0 boson) had been indirectly measured to be three via the measurement of the Z^0 decay width at the Large Electron Proton collider (LEP) at CERN in 1989.

Other milestones of the electroweak interaction were the discovery of weak neutral currents in neutrino interactions by the Gargamelle experiment in 1973, as well as the discovery of the $W^{\pm}$ and Z^0 bosons in proton-antiproton collisions at the Super Proton Synchrotron (SPS) in 1983. Finally, the Higgs boson was discovered in 2012 in proton-proton collisions at the Large Hadron Collider (LHC). All these experiments were conducted at CERN.

As to quarks, their existence was postulated in 1964 by Murray Gell-Mann as a tentative explanation of the symmetrical classification pattern of the known hadrons. "Strange" hadrons (containing the strange quark s) had been known since the discovery of kaons in 1947. The remaining quarks (charm c, top t and bottom b) were discovered at collider experiments:

- the charm quark via the existence of "charmed hadrons" like the J/ψ meson, first observed in 1974 both in the SLAC linear $e^+ - e^-$ collider and in proton-nucleus collisons (a proton beam onto a fixed target) at the Brookhaven National Laboratory (BNL);
- the bottom quark (as a $b\bar{b}$ meson) in fixed target proton collisions in 1977 at Fermilab;
- the top quark in 1995 at the Tevatron proton-antiproton collider at Fermilab;

A.1.4 Mixing the generations

While processes violating the conservation of the *total* lepton number L have never been observed, certain processes do not conserve the individual flavour number L_α. For instance it has been experimentally established that neutrinos can undergo flavour transitions:

$$\nu_\alpha \to \nu_\beta, \quad \alpha \neq \beta \tag{A.20}$$

where a neutrino is produced in a given flavour α and, after propagating in either vacuum of matter for a while, is detected in a different flavour β. Obviously, such a process violates the conservation of L_α and L_β. This phenomenon is generally referred to as *neutrino flavour oscillations* or simply *neutrino oscillations*, because the probability for the flavour transition to occur is an oscillating function of the variable L/E, where L is the distance between the production and detection of the neutrino, and E is the energy of the neutrino.

These oscillations are very well described theoretically as a consequence of *neutrino mixing*:

- the quantum eigenstates in which neutrinos are produced and detected are states of definite flavour,
- while the eigenstates relevant to describe the propagation of neutrinos are states of definite mass.

Neutrino mixing describes the fact that the flavour and mass eigenstates are not identical, but are related through a certain "mixing": a given flavour eigenstate is a mixture of the three mass eigenstates, and conversely.

Neutrino flavour oscillations are the subject of this thesis, therefore the current knowledge of this process is adressed in detail in the main body of this thesis. In Sec. 1.1.3 the context of their

discovery is presented, and their mechanism is illustrated in the two-flavour case. In Sec. A.2, three-flavour oscillations both in vacuum and matter are discussed in more ample detail.

To complete this introduction to flavour, let us consider the quark sector. As was summarised on Tab. A.2, the quarks come in 6 flavours (u, d, c, s, t, b) which are divided into three generations. Each generation includes one *up-type quark* (u, c, t) and one *down-type quark* (d, s, b). As seen on the examples of β-decays and IBD processes, the weak charged current always couples an up-type quark with a down-type quark. Therefore weak charged currents are sometimes said to *maximally violate* the conservation of flavour in the quark sector. This is in contrast with the electromagnetic and strong interactions as well as weak neutral currents, all conserving the flavour of quarks.

In fact, quarks are even allowed to change from one generation to another via weak charged currents. As an example a strange quark s (second generation) can decay into an up quark u (first generation) and a pair (e^-, $\overline{\nu}_e$):

$$s \to u + e^- + \overline{\nu}_e \tag{A.21}$$

This is formally similar to a lepton flavour violating process for quarks, since the total number of quarks of the first and second generation is not identical in the initial and final states. The corresponding Feynman diagram can be seen on Fig. A.4(a). Such processes explain for example that strange hadrons such as kaons (*e.g.* $K^0 = u$) can decay into pions, as known since the 1960's.

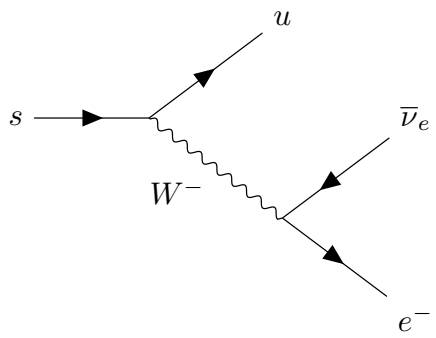

(a) Strange quark decay into an up quark

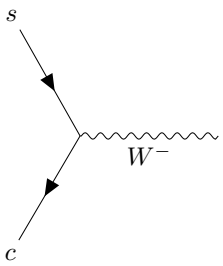

(b) Normal CC coupling

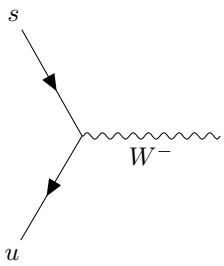

(c) Effect of quark mixing

Figure A.4: Flavor changing weak charged current decay of a strange quark.

Such processes are interpreted, as in the case of neutrinos, by a *mixing* phenomenon for quarks. More precisely, the charged current coupling of the quark fields in the absence of mixing would be the one of Fig. A.4(b), and the effect of the mixing is that instead of c only, the up-type quarks of other generations can couple to s (with different amplitudes, which are determined by the quark mixing parameters) as in Fig. A.4(c).

As to charged leptons, there are theoretical reasons within the SM which forbid their mixing. Searches for processes beyond the SM under the form of charged lepton flavour-violating (CLFV) decays however constitute an active area of experimental research.

A.2 Elements of the Standard Model

A.2.1 Fields and particles

The mathematical framework for the theoretical description of particle physics is quantum field theory, as introduced in the glossary. In this section we will not discuss the quantisation of fields and the particle states, except when necessary. Fields are the fundamental objects to construct the SM.

A.2.2 Free Dirac fermions

Searching for a relativistic equivalent of the Schrödinger equation for a free massive spin-1/2 particle, with the requirement that the obtained equation be of first order in time, leads to the Dirac equation (see for instance Ref. [41], Chap. 1):

$$(i\gamma^\mu\partial_\mu - m)\psi = 0. \tag{A.22}$$

where m is the rest mass of the particle and the γ^μ are matrices which must satisfy the anti-commutation relation

$$\{\gamma^\mu, \gamma^\nu\} = 2\eta^{\mu\nu}. \tag{A.23}$$

and

$$\gamma^0\gamma^{\mu\dagger}\gamma^0 = \gamma^\mu \tag{A.24}$$

The minimal dimension to realise the algebra A.23 is four. This implies that the solutions of the Dirac equation are four-component objects called *Dirac spinors*. Note that these objects are very different from the usual four-vectors of special relativity. To each particular choice (called a *representation*) of the γ matrices satisfying Eq. A.23 is associated a particular basis of the Dirac spinors.

Under a Lorentz transformation represented by the matrix (written in tensor form) $\Lambda^\mu{}_\nu$:

$$x^\mu \to x'^\mu = \Lambda^\mu{}_\nu x^\nu, \tag{A.25}$$

a Dirac spinor field transforms as:

$$\psi(x^\mu) \to \psi'(x'^\mu) = S(\Lambda)\psi(x) \tag{A.26}$$

where the 4×4 matrix $S(\Lambda)$ is determined as:

$$S(\Lambda)^{-1}\gamma^\mu S(\Lambda) = \Lambda^\mu{}_\nu\gamma^\nu \tag{A.27}$$

This transformation is explicitly required in order for the Dirac equation to be *Lorentz covariant*, meaning that its form remains the same under the action of Lorentz transformations. This is fundamentally different from the way a scalar field or a vector field would transform. A scalar field s (one component) would transform with $S(\Lambda) = 1$. A vector field V has the same number of components as a spinor field ($V(x^\mu)$ is a four-vector) but would transform with $S(\Lambda) = \Lambda$. Mathematically: scalars (spin 0), spinors (spin 1/2) and vectors (spin 1) belong to different *representations* of the Lie group $\mathcal{O}(1,3)$ formed by Lorentz transformations.

In the Lagrangian formalism, one possibility to write the free Dirac Lagrangian is

$$\mathcal{L} = \overline{\psi}(i\gamma^\mu\partial_\mu - m)\psi \tag{A.28}$$

where $\overline{\psi}(x)$ is the adjoint field:

$$\overline{\psi} \equiv \psi^\dagger\gamma^0 \tag{A.29}$$

Several ways to obtain the Dirac equation are discussed in the literature. Let us quickly refer to two rather different approaches, which are both instructive. The first one proceeds via the application of the correspondance principle of quantum mechanics to the relativistic expression

$$E^2 = p^2 + m^2. \tag{A.30}$$

Naively, one obtains the second order *Klein-Gordon* equation which is simply an extension of Maxwell's wave equation to a massive wave function. When searching for a first order equation which implies the Klein-Gordon equation (this was motivated historically by quantisation and causality considerations), one is led to write the Dirac equation and its algebra. This approach is described for instance in Chap. 1 of Ref. [286] (available online).

The second classical derivation proceeds in a more abstract manner by considerations of Lorentz covariance, employing tools of the theory of representations of Lie groups and Lie algebras. From that point of view, Dirac spinors arise as the minimal relativistic extension of the treatment of spin-1/2 particles in non-relativistic quantum mechanics. This is discussed in details in Ref. [41] (Chap. 5) and more briefly in Ref. [286] (Chap. 33).

The so-called *standard representation* of the γ matrices can be found in any of Ref. [1, 41, 286, 287]. In the following we will mostly be interested in another one, the *chiral representation*:

$$\gamma^0 = \begin{pmatrix} 0 & I_2 \\ I_2 & 0 \end{pmatrix}, \ \gamma^i = \begin{pmatrix} 0 & \sigma^i \\ -\sigma^i & 0 \end{pmatrix} \tag{A.31}$$

where the σ^i are the two-dimensional Pauli matrices

$$\sigma^1 \equiv \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, \ \sigma^2 \equiv \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix}, \ \sigma^3 \equiv \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix} \tag{A.32}$$

which are known for generating the spin-1/2 algebra in non-relativistic quantum mechanics, defined by the commutation relations

$$[\sigma_i, \sigma_j] = \varepsilon_{ijk}\sigma_k. \tag{A.33}$$

It is also convenient to introduce the *chirality matrix* γ^5:

$$\gamma^5 \equiv i\gamma^0\gamma^1\gamma^2\gamma^3 \tag{A.34}$$

which is diagonal in the chiral representation:

$$\gamma^5 = \begin{pmatrix} I_2 & 0 \\ 0 & -I_2 \end{pmatrix} \tag{A.35}$$

The chiral representation of the γ matrices (also called the Weyl representation) is the one arising naturally in the derivation of the Dirac equation based on spin and Lorentz covariance.

Chirality, parity and mass

Any Dirac spinor ψ can be decomposed as

$$\psi = \psi_R + \psi_L \tag{A.36}$$

where ψ_R and ψ_L are eigenvectors of the chirality matrix γ^5:

$$\gamma^5 \psi_R = +\psi_R \tag{A.37}$$
$$\gamma^5 \psi_L = -\psi_L \tag{A.38}$$

In the chiral basis, γ^5 is diagonal and this decomposition simply reads:

$$\psi_R = \begin{pmatrix} \chi_R \\ 0 \end{pmatrix}, \ \psi_L = \begin{pmatrix} 0 \\ \chi_L \end{pmatrix}, \ \psi = \begin{pmatrix} \chi_R \\ \chi_L \end{pmatrix}. \tag{A.39}$$

where χ_R and χ_L are two-components spinors, called respectively *right-handed* and *left-handed* Weyl spinors.

χ_R and χ_L are similar to the spinors describing spin-1/2 particles in non-relativistic quantum mechanics. In fact, they belong to equivalent representations of the *restricted Lorentz group* $\mathbf{SO}^{\uparrow}(1,3)$, which is the subset of Lorentz transformations conserving both the direction of time and the orientation of space (see *e.g.* Ref. [1], Appendix B).

However, the discrete symmetries of *parity* (*i.e.* space inversion)

$$P: \ (t,\vec{x}) \to (t,-\vec{x}) \tag{A.40}$$

and *time reversal*

$$T: \ (t,\vec{x}) \to (-t,\vec{x}) \tag{A.41}$$

are not elements of $\mathbf{SO}^{\uparrow}(1,3)$. As a consequence, the left-handed and right-handed representations are not equivalent representations of the full Lorentz group $\mathbf{O}(1,3)$. For instance, spatial rotations act identically on χ_R and χ_L, but Lorentz boosts act in a reversed manner (see Ref. [1], Chap. 1 or any of Ref. [41, 286, 287]).

In fact, left-handed and right-handed Weyl spinors are exchanged by the parity operator:

$$P\chi_L = \chi_R \tag{A.42}$$
$$P\chi_R = \chi_L \tag{A.43}$$

Thus four-component spinors are necessary to describe massive Dirac fermions in a fully covariant manner. This is also visible from the fact that the space-time evolution of the left-handed and right-handed components of the free Dirac field are coupled:

$$i\gamma^\mu \partial_\mu \psi_R = m\psi_L \tag{A.44}$$
$$i\gamma^\mu \partial_\mu \psi_L = m\psi_R \tag{A.45}$$

In a related manner, by inserting the explicit decomposition A.36 into the Dirac Lagrangian A.28 we see that only the contraction of left-handed components with right-handed components contribute to the mass term:

$$-m\overline{\psi}\psi = -m\Big(\overline{\psi_R} + \overline{\psi_L}\Big)\Big(\psi_R + \psi_L\Big) = -m\Big(\overline{\psi_R}\psi_L + \overline{\psi_L}\psi_R\Big) \tag{A.46}$$

Again, this shows that a Dirac spinor field needs both left-handed and right-handed chiral components to be described as massive. We will see that this has important consequences regarding the generation of mass in the Standard Model and beyond.

The situation is simplified in the case of a massless fermion. Indeed, the Dirac equation then becomes the Weyl equation

$$i\gamma^\mu \partial_\mu \psi = 0 \tag{A.47}$$

which is solved by two-component Weyl spinors. For a massless four-component Dirac spinor in the chiral basis, the space-time evolutions of ψ_R and ψ_L become decoupled (this is clear when setting $m = 0$ in Eqs. (A.44) and (A.45)). Furthermore it can be shown explicitly that each one has only two independent components. The chiral (2-components) representation is thus sufficient to describe massless fermions. In such a description, the field is either left-handed or right-handed, and its chirality is reversed by the parity operator.

Helicity

The *helicity* h of a particle with spin s is defined as the projection of its spin $\vec{S}$ on its three-momentum $\vec{P}$:

$$h = \frac{\vec{S}.\vec{P}}{s|\vec{P}|} \tag{A.48}$$

For a spin-1/2 fermion, the helicity operator has two eigenvalues $h = \pm 1$. In the massless case helicity coincides with chirality. A massless right-handed chiral field ψ_R (with definite four-momentum) has positive helicity; its left-handed counterpart ψ_L has negative helicity.

For a massive fermion, however, this correspondance does not hold. A helicity eigenstate is a mixture of left-handed and right-handed chirality eigenstates, with dependence on the frame of reference. This can be understood physically by the fact that the direction of the momentum of a massive particle can always be reverted by considering an inertial frame moving in the same direction as the particle at a sufficient velocity. The same cannot be done for massless particles.

For instance the chiral eigenstate $\psi_R(\vec{p})$ can be expressed schematically as:

$$\psi_R(\vec{p}) \sim \frac{1}{2}\left(1 + \frac{|\vec{p}|}{E+m}\right)\psi_+(\vec{p}) + \frac{1}{2}\left(1 - \frac{|\vec{p}|}{E+m}\right)\psi_-(\vec{p}) \tag{A.49}$$

where ψ_+ and ψ_- are helicity eigenstates with respective eigenvalues ± 1. In the ultrarelativistic limit where $\frac{|\vec{p}|}{E+m} \to 1$ the helicity-chirality correspondance is recovered.

Antiparticles and charge conjugation

When examining the solutions of the Dirac equation via a Fourier transform approach, *i.e.* as plane waves with the form

$$\psi(t, \vec{x}) = \omega(p^\mu) e^{-ip_\mu x^\mu} = \omega(p^\mu) e^{-i(E.t - \vec{p}.\vec{x})} \tag{A.50}$$

one encounters what appears to be independent solutions with both positive energy

$$E = +\sqrt{|\vec{p}| + m^2} \tag{A.51}$$

and negative energy

$$E = -\sqrt{|\vec{p}| + m^2}. \tag{A.52}$$

Alternatively (via a rephasing with a factor $e^{i\pi}$) the solutions with negative energy can be interpreted as positive energy solutions propagating backwards in time, which is just as much intriguing. This led Dirac to predict the existence of antimatter in 1931, one of the most celebrated examples of the predictive power of theoretical physics.

Dirac formulated his equation in 1928 as a relativistic description of the electron, which, along with the proton, was one of the two known particles at the time. In 1930 he proposed an idea based on the Pauli exclusion principle to solve the negative energy paradox. The idea was to postulate that the vacuum is in fact a quantum state where all energy levels with $E < 0$ are occupied (this became known as the *Dirac sea*), therefore preventing the electrons to fall into

these unphysical states. He predicted that in case some negative energy states become available, these "holes" should be detectable as positively charged particles. After some adjustments of the initial idea, he was led to predict in 1931 the existence of a new particle, positively charged and with the same mass as the electron. One year later, the positron e^+ was discovered in the study of cosmic rays. In spite of this undeniable success, the Dirac theory of holes was still unsatisfactory: what about the infinite charge density created by the sea of electrons? How to explain that the bosons, to whom the Pauli exclusion principle does not apply, also have their antiparticles? The formalism of QFT allows to solve this problem and to achieve a symmetrical description of particle and antiparticle states.

In the modern formulation of QFT, a quantised spinor field solution of the Dirac equation is used to form operators creating particle states and annihilating antiparticle states. The spinor fields which conversely create antiparticles and annihilate particles are obtained via the charge conjugation operator C:

$$\psi(x^\mu) \xrightarrow{C} \psi^C(x) = \eta_c \mathcal{C} \overline{\psi}^T(x^\mu) = \eta_c \gamma^0 \mathcal{C} \psi^*(x^\mu) \tag{A.53}$$

where $\mathcal{C}$ is the charge conjugation matrix, given in the chiral basis by

$$\mathcal{C} = -i \begin{pmatrix} \sigma^2 & 0 \\ 0 & \sigma^2 \end{pmatrix}, \tag{A.54}$$

bearing in mind that σ^2 denotes the second Pauli matrix of Eq. A.32. η_c is a phase factor ($|\eta_c|^2 = 1$) which has no physical effect and is thus arbitrary for free fields (we can take it equal to 1).

If ψ is a solution of the Dirac equation (A.22), then ψ^C is also a solution. Equivalently, the Lagrangian (A.28) is invariant by the transformation (A.53). The charge conjugation, exchanging particle states and antiparticle states, reverses the sign of all the additive quantum numbers (electrical charge, lepton number, *etc*). However it does not reverse the helicity of a given state.

Finally, let us mention that the invariance by the product CPT (time reversal, parity, and charge conjugation) is a necessary feature of all local quantum field theories. In a sense this invariance requires the existence of antiparticles. It is also related to the interpretation of antiparticles as the equivalent of particles moving backwards in time (as mentioned hereabove) and with mirrored helicities.

A.2.3 The electroweak interaction as a chiral gauge theory

Up to now we have only described the theory of free fields, *i.e.* fields that do not interact with each other. This section provides an introduction to the electroweak Standard Model, which unifies the electromagnetic and weak interactions in a common description.

Interactions from local gauge symmetry

In the Standard Model, interactions are introduced via the promotion of global gauge invariance to local gauge invariance. As an example of gauge invariance, the Dirac Lagrangian is invariant by the transformation

$$\psi(x) \to \psi'(x) = e^{i\theta} \psi(x) \tag{A.55}$$

This is referred to as a *global* $\mathbf{U}(1)$ invariance:

- it corresponds to the action of the one-parameter Lie group $\mathbf{U}(1)$ onto the vector space of spinors
- this action is *global*, in the sense that the parameter θ does not depend on x

Via Noether's theorem, this invariance is related to the conservation of the electromagnetic four-current. The *global* symmetry is promoted to a *local* symmetry by requiring that the lagrangian be invariant under any transformation, parametrised by a function $\theta(x)$:

$$\psi(x) \to \psi'(x) = e^{i\theta(x)}\psi(x) \tag{A.56}$$

This leads to replace in the Dirac Lagrangian the derivative ∂_μ with the *covariant derivative*

$$D_\mu = \partial_\mu - iqA_\mu \tag{A.57}$$

where q is the electric charge of the particle and A_μ is the electromagnetic potential. A field equation for A_μ can be derived from considerations of Lorentz covariance. This leads altogether to retrieve the relativistic formulation of Maxwell's theory of electromagnetism. From there, the theory of *quantum electrodynamics* is obtained via the application of the canonical quantisation procedure to both the spinor field ψ (electrons and positrons) and the vector field A_μ (photons).

Chiral gauge symmetry

In the electroweak part of the Standard Model, the local gauge group is $\mathbf{SU}(2)_L \times \mathbf{U}(1)_Y$. The symmetry group $\mathbf{SU}(2)_L$ is called *weak isospin*, and the subscript L refers to left-handed chirality. Indeed, a crucial feature of the SM is that the left-handed and right-handed chiral components of the fermion fields, with respect to weak isospin transformations, do not live in the same representation of $\mathbf{SU}(2)$. This asymmetrical treatment accounts for the parity violating nature of the weak interaction. Because of this structure, the SM is sometimes referred to as a *chiral gauge theory*.

The left-handed components of the leptons and quark fields are grouped into *weak isospin doublets* (L for leptons and Q for quarks):

$$L_{eL} \equiv \begin{pmatrix} \nu_{eL} \\ e_L \end{pmatrix},\ L_{\mu L} \equiv \begin{pmatrix} \nu_{\mu L} \\ \mu_L \end{pmatrix},\ L_{\tau L} \equiv \begin{pmatrix} \nu_{\tau L} \\ \tau_L \end{pmatrix} \tag{A.58}$$

$$Q_{1L} \equiv \begin{pmatrix} u_L \\ d_L \end{pmatrix},\ Q_{2L} \equiv \begin{pmatrix} c_L \\ s_L \end{pmatrix},\ Q_{3L} \equiv \begin{pmatrix} t_L \\ b_L \end{pmatrix} \tag{A.59}$$

$\mathbf{SU}(2)$ has three generators denoted as $\widehat{I}_a$ ($a = 1,\ 2,\ 3$), and the corresponding Lie algebra is defined by

$$[\widehat{I}_a,\ \widehat{I}_b] = i\epsilon_{abc}\widehat{I}_c \tag{A.60}$$

This algebra is similar to the one of $\mathbf{SO}(3)$ (non-relativistic spinor algebra). The action of an infinitesimal $\mathbf{SU}(2)_L$ symmetry transformation on the weak isospin doublets is represented by the Pauli matrices (A.32), up to a factor 1/2:

$$\widehat{I}_a L_L = \frac{\sigma_i}{2} L_L, \tag{A.61}$$

$$\widehat{I}_a Q_L = \frac{\sigma_i}{2} Q_L \tag{A.62}$$

This is the simplest non-trivial representation of $\mathbf{SU}(2)$.

On the other hand, the right-handed components of the fermion fields are considered as *singlets* under the weak isospin symmetry, *i.e.* they belong to the trivial representation in which weak isospin transformations have no effect. Moreover, there is no right-handed neutrino singlet. Thus there are only three weak isospin singlets for leptons:

$$e_R,\ \mu_R,\ \tau_R \tag{A.63}$$

while there is one singlet per quark flavour:

$$u_R,\ d_R,\ c_R,\ s_R,\ t_R,\ b_R. \tag{A.64}$$

For any fermion singlet f_R, the action of the infinitesimal generators $\widehat{I_a}$ is:

$$\widehat{I_a} f_R = 0 \tag{A.65}$$

In terms of symmetry, this means that the corresponding field is unaffected by $\mathbf{SU}(2)_L$ transformations. Physically, the effect is that the right-handed fermion fields do not take part in the weak interaction, which in turn relates to the experimental fact that the weak interaction violates parity conservation.

The $\mathbf{U}(1)_Y$ part of the gauge group is called the *weak hypercharge* symmetry. Note that it is not the same symmetry as the $U(1)$ symmetry group of QED. $\mathbf{U}(1)$ is the most simple continuous symmetry, and its action on any fermion field f can only be represented by a homothetic transformation:

$$\widehat{Y} f = Y \times f \tag{A.66}$$

In the Standard Model, the value of the constant Y (the weak hypercharge) for each fermion field f is determined by the Gell-Mann-Nishijima relation

$$Q = I_3 + \widehat{Y} 2 \tag{A.67}$$

where Q is the electric charge of the field f.

One can then deduce the covariant derivative D_μ to use for the kinetic term of each fermion field f

$$\gamma^\mu \partial_\mu f \to \gamma^\mu D_\mu f \tag{A.68}$$

so that the overall Lagrangian is invariant under any *local* $\mathbf{SU}(2)_L \times \mathbf{U}(1)_Y$ transformation. This requires to introduce *gauge boson fields* in the theory, similarly to the photon field in QED. The replacement (A.68) introduces coupling terms between the gauge boson fields and the fermion fields. The covariant derivative obtained for $\mathbf{SU}(2)_L \times \mathbf{U}(1)_Y$ is the following:

$$D_\mu \equiv \partial_\mu + ig\vec{W}_\mu.\vec{I} + ig' B_\mu \frac{\widehat{Y}}{2} \tag{A.69}$$

where $\vec{W}_\mu$ is a short notation for the three gauge boson fields $W_\mu^1,\ W_\mu^2,\ W_\mu^3,$ associated with the weak isospin symmetry and $\vec{I} = (\widehat{I_1},\ \widehat{I_2},\ \widehat{I_3})$. In turn, B_μ is the gauge boson associated with the weak hypercharge symmetry. At this point, the coupling constants g and g' are the only free parameters of the model.

The requirement of gauge invariance also determines how the gauge boson fields must transform under the action of the symmetry group, and the kinetic terms and self-interaction terms to include in the Lagrangian for these fields. This procedure is described for example in Ref. [1] (in Appendix D for a general gauge theory and in Chap. 3 for the particular case of the SM).

Restricting ourselves to the first generation of quarks and leptons, the corresponding Lagrangian is:

$$\mathcal{L} \equiv i\overline{L}_L \gamma^\mu D_\mu L_L + i\overline{Q}_L \gamma^\mu D_\mu Q_L + i \sum_{f=e,u,d} f_R \gamma^\mu D_\mu f_R + \mathcal{L}_{\text{gauge}} \tag{A.70}$$

It describes massless Dirac fermions, massless gauge boson fields, and their interactions. $\mathcal{L}_{\text{gauge}}$ contains the kinetic terms and self-interaction terms for the gauge bosons $W_\mu^1,\ W_\mu^2,\ W_\mu^3,$ and B_μ. These terms will not be discussed here.

Physical vector bosons and couplings

When expanding the covariant derivative (A.69) explicitly in the Lagrangian (A.70), the fermion-boson interaction terms corresponding to charged currents can be isolated by introducing a charged vector boson field:

$$W_\mu^+ = \frac{1}{\sqrt{2}}(W_1 + iW_2). \tag{A.71}$$

The remaining terms are neutral currents. In order to retrieve the QED Lagrangian as a part of the electroweak neutral current Lagrangian, a mixing between the gauge boson fields W_3^μ and B^μ is introduced:

$$A^\mu = \quad \cos\theta_W B^\mu + \sin\theta_W W_3^\mu \tag{A.72}$$
$$Z^\mu = -\sin\theta_W B^\mu + \cos\theta_W W_3^\mu \tag{A.73}$$

where θ_W is a parameter called the *Weinberg mixing angle*. It is chosen so that the coupling between the electromagnetic field A^μ and the fermion fields coincides with QED, which imposes:

$$g\sin\theta_W = g'\cos\theta_W = e \tag{A.74}$$

where e is the elementary electric charge.

The W_μ^+ field, in terms of particle states, creates W^+ bosons and annihilates W^- bosons (antiparticle of W^+), associated in Sec. A.1 to CC neutrino interactions. The A^μ and Z^μ fields correspond respectively to the photon γ and to the Z^0 boson associated in Sec. A.1 to NC neutrino interactions (both γ and Z^0 are their own antiparticle). At this stage, all four of them are massless.

Finally, we can rewrite the Lagrangian (A.70) as

$$\mathcal{L} = \mathcal{L}_{\text{kin}} + \mathcal{L}_{\text{int}} + \mathcal{L}_{\text{gauge}} \tag{A.75}$$

where $\mathcal{L}_{\text{kin}}$ contains the purely kinetic terms for the fermion fields, and $\mathcal{L}_{\text{int}}$ writes:

$$\mathcal{L}_{\text{int}} = \mathcal{L}_{\text{weak,CC}} + \mathcal{L}_{\text{weak,NC}} + \mathcal{L}_{\text{elm}} \tag{A.76}$$

where $\mathcal{L}_{\text{weak,CC}}$, $\mathcal{L}_{\text{weak,NC}}$ and $\mathcal{L}_{\text{elm}}$ contain terms corresponding to the trilinear couplings represented respectively by the diagrams of Fig. A.5, A.6, and A.7.

For instance, the weak CC coupling for leptons writes as:

$$\mathcal{L}_{\text{weak,CC}} = -\frac{g}{2\sqrt{2}} j^\mu_{\text{CC,lep}} W_\mu^+ + \text{h.c.} \tag{A.77}$$

where $j^\mu_{\text{CC,lep}}$ is the leptonic charged current

$$j^\mu_{\text{CC,lep}} = 2\sum_{\alpha=e,\ \mu,\ \tau} \nu_{\alpha L}\,\gamma^\mu\, l_{\alpha,L} = \overline{\boldsymbol{\nu}_L}\gamma^\mu \boldsymbol{l}_L + \text{h.c.} \quad \text{with} \quad \boldsymbol{l}_L = \begin{pmatrix} e_L \\ \mu_L \\ \tau_L \end{pmatrix},\ \boldsymbol{\nu}_L = \begin{pmatrix} \nu_{1L} \\ \nu_{2L} \\ \nu_{3L} \end{pmatrix} \tag{A.78}$$

Only the left-handed fields are coupled via the weak charged current. In terms of the full lepton fields $l_\alpha = l_{\alpha,L} + l_{\alpha,R}$ and $\nu_\alpha = \nu_{\alpha L}$, Eq. A.78 can be rewritten as

$$j^\mu_{\text{CC,lep}} = \sum_{\alpha=e,\ \mu,\ \tau} \nu_\alpha\,\gamma^\mu(1-\gamma^5)\,l_\alpha \tag{A.79}$$

where $P_L = \frac{1-\gamma^5}{2}$ is the chiral left-handed projector.

Since W_μ^+ is a combination of the gauge fields of $\mathbf{SU}(2)_L$ which acts selectively on left-handed fields (A.71), charged currents maximally violate parity. On the other hand, weak neutral currents involve both the left-handed and right-handed components of the charged lepton fields, because of the mixing (A.72, A.73) of the gauge fields W_3^μ associated to $\mathbf{SU}(2)_L$ and B^μ associated to $\mathbf{U}(1)_Y$. Finally, the electromagnetic current involves only the charged leptons, and it writes in terms of the full fields l_α:

$$\mathcal{L}_{\text{elm,lep}} = \sum_{\alpha=e,\mu,\tau} -e\,\overline{l_\alpha}\,\gamma^\mu\, l_\alpha\, A_\mu \tag{A.80}$$

where e is the elementary electric charge. The electromagnetic coupling is thus independent of chirality.

Generation of mass and electroweak symmetry breaking

Because the left-handed and right-handed components of the fermion fields f do not belong to the same representations of the gauge group, introducing simple Dirac mass terms of the form

$$-m\overline{f}f = -m\Big(\overline{f_R}+\overline{f_L}\Big)\Big(f_R+f_L\Big) = -m\Big(\overline{f_R}f_L+\overline{f_L}f_R\Big) \tag{A.81}$$

would break the gauge invariance of the Lagrangian, which would in turn cause the theory to be non-renormalisable[3]. Direct mass terms for the boson fields would have the same effect. The mass terms must then be introduced via a different mechanism.

In the Standard Model, this is realised via the *Higgs mechanism*, which will only be summarised qualitatively here. Two complex scalar fields ϕ^+ (positively charged) and ϕ^0 (neutral) are added to the theory. Similarly to the left-handed components of fermion fields, they transform as a doublet, called the *Higgs doublet*, under the weak isospin symmetry:

$$\Phi = \begin{pmatrix}\phi^+\\ \phi^-\end{pmatrix} \tag{A.82}$$

The Lagrangian of the electroweak Standard Model is then defined as the most general form of Lagrangian meeting the following requirements:

- contains the above fermion fields, gauge boson fields, and the Higgs doublet;
- is invariant under local $\mathbf{SU}(2)_L \times \mathbf{U}(1)_Y$ transformations;
- yields a *renormalisable* theory.

In addition to the terms previously seen, this Lagrangian includes kinetic terms and self coupling terms for the Higgs doublet, coupling terms between the Higgs doublet and the gauge boson fields W_μ^i and B_μ, and coupling terms (called *Yukawa couplings*) between the Higgs doublet and the fermion fields.

Furthermore, the *vacuum state* of the theory is assumed to acquire a non-zero expectation value for the neutral part of the Higgs doublet, ϕ^0, through a process called *spontaneous symmetry breaking* (SSB). As a result:

- the $\mathbf{SU}(2)_L \times \mathbf{U}(1)_Y$ symmetry is broken for the vacuum state (and consequently for the particle states built upon it);
- the gauge symmetry $\mathbf{U}(1)_Q$ corresponding to the conservation of the electric charge ($Q = I_3 + \frac{Y}{2}$) is however preserved;

[3] When calculating the probabilities of particle interactions (scattering amplitudes) as perturbative expansions, some quantities tend to diverge. A renormalisable theory is a theory in which a scheme can be defined to "absorb" these infinite quantities by redefining the coupling constants, so that scattering amplitudes can be calculated.

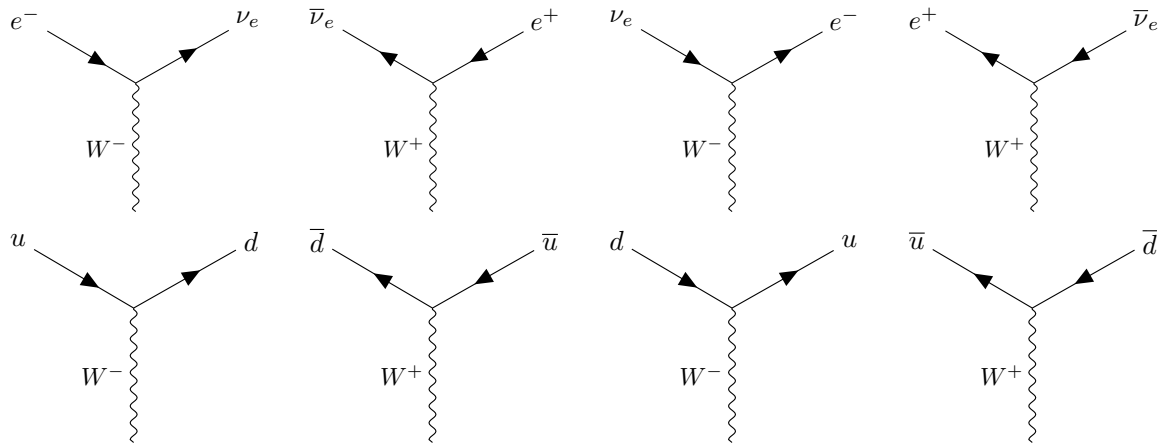

Figure A.5: Weak charged current trilinear couplings ($\mathcal{L}_{\text{weak,CC}}$ in Eq. A.76) for fermions of the first generation.

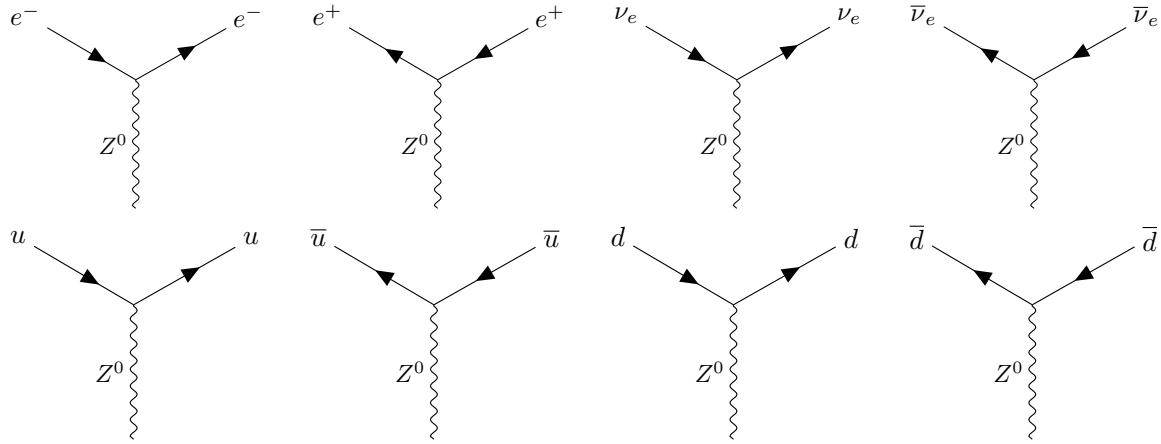

Figure A.6: Weak neutral current trilinear couplings ($\mathcal{L}_{\text{weak,NC}}$ in Eq. A.76) for fermions of the first generation.

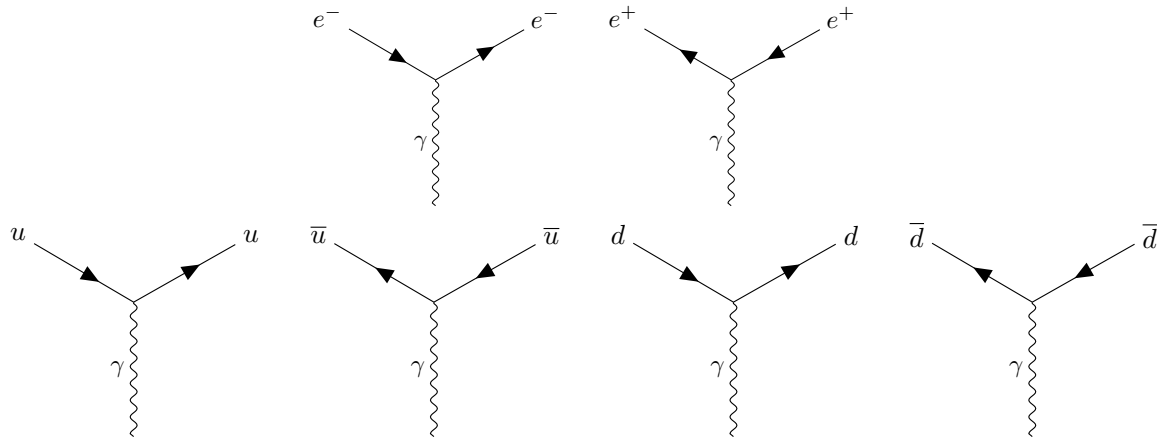

Figure A.7: Electromagnetic trilinear couplings ($\mathcal{L}_{\text{elm}}$ in Eq. A.76) for fermions of the first generation.

- a mass term for the Higgs field appears;
- mass terms appear for the $W^{\pm}$ and Z^0 bosons, but not for the photon;
- trilinear and quadrilinear coupling terms appear, coupling the Higgs boson with itself and the Higgs boson with the $W^{\pm}$ and Z^0 bosons.

The $\mathbf{SU}(2)_L \times \mathbf{U}(1)_Y \to \mathbf{U}(1)_Q$ symmetry breaking and the related fact that the $W^{\pm}$ and Z^0 acquire mass while the photon stays massless are very important physically. The large mass of the $W^{\pm}$ and Z^0 bosons explains that the electroweak theory reduces in the low-energy limit to the four-fermion point-like interaction described by the Fermi theory, while the range of the electromagnetic interaction remains infinite whatever the energy scale. Note that the scale of the masses of the Higgs boson and of the $W^{\pm}$ and Z^0 bosons ($\sim 10^2\,\mathrm{GeV}$) is not predicted by the model but empirical.

A.2.4 Fermion masses and mixing

In the above description we have only considered the first generation of fermions. When considering three generations of massive fermions with the most general form of Lagrangian, the Yukawa couplings between the Higgs field and the fermion fields can couple the generations between themselves. As a result, the fermion fields introduced as in (A.58), (A.59) and (A.63) do not necessarily have a definite mass. Instead, the fields with definite masses can be a mixture of these fields.

The fields with definite mass are important physically, because they correspond to the particle states relevant to describe the propagation of particles, and thus the initial and final state in a scattering process. Thus it is customary to introduce the initial fermion fields (without definite mass) as ν'_{eL}, e'_L, u'_L, d'_L, *etc.*

For leptons, the Yukawa couplings then give rise to mass terms of the form (omitting the Higgs-lepton trilinear couplings):

$$\mathcal{L}_{\mathrm{mass,L}} = -\frac{v}{\sqrt{2}}\left(\sum_{\alpha=e,\mu,\tau}\sum_{\beta=e,\mu,\tau} Y'^{l}_{\alpha\beta}\overline{l'_{\alpha L}}\,l'_{\beta L}\right) + \mathrm{h.c.} \tag{A.83}$$

$$= -\frac{v}{\sqrt{2}}\left(\overline{\boldsymbol{l}'_L}\,Y'^{l}\,\boldsymbol{l}'_R\right) + \mathrm{h.c.} \tag{A.84}$$

where v is the *vacuum expectation value* (vev) of the Higgs field, $\boldsymbol{l}'_L$ and $\boldsymbol{l}'_R$ are vector notations for the charged lepton fields

$$\boldsymbol{l}'_L = \begin{pmatrix} e'_L \\ \mu'_L \\ \tau'_L \end{pmatrix}, \quad \boldsymbol{l}'_R = \begin{pmatrix} e'_R \\ \mu'_R \\ \tau'_R \end{pmatrix}, \tag{A.85}$$

Y'^{l} is the matrix of Yukawa coupling parameters, and h.c. means "hermitian conjugate". Since no right-handed field $\nu_{\alpha R}$ has been included from the start, there are no Yukawa couplings and thus no mass term for neutrinos.

The charged lepton fields with definite mass are found by reducing the Yukawa matrix to a diagonal form. Since the Yukawa couplings are arbitrary parameters of the model, Y'^{l} can be any complex matrix. Thus we can only write a "biunitary diagonalisation" with two distinct unitary matrices V_R^l and V_L^l :

$$V_L^{l\dagger}\, Y'^{l}\, V_R^l = Y^l = \mathrm{diag}(y_e^l,\ y_\mu^l,\ y_\tau^l) \tag{A.86}$$

where V_R^l and V_L^l can be chosen so that the diagonal values y_e^l, y_μ^l, y_τ^l are real and positive.

We can finally define the charged lepton fields with definite mass as:

$$\boldsymbol{l}_L = \begin{pmatrix} e_L \\ \mu_L \\ \tau_L \end{pmatrix} \equiv V_L^{l\dagger} \boldsymbol{l}'_L, \quad \boldsymbol{l}_R = \begin{pmatrix} e_R \\ \mu_R \\ \tau_R \end{pmatrix} \equiv V_R^{l\dagger} \boldsymbol{l}'_R \tag{A.87}$$

such that the Yukawa mass terms (A.83) can be rewritten as usual Dirac mass terms:

$$\mathcal{L}_{\text{mass,L}} = -m_e \overline{l_e} l_e - m_\mu \overline{l_\mu} l_\mu - m_\tau \overline{l_\tau} l_\tau \tag{A.88}$$

with the overall charged lepton fields

$$l_e \equiv e_R + e_L, \quad l_\mu \equiv \mu_R + \mu_L, \quad l_\tau \equiv \tau_R + \tau_L \tag{A.89}$$

and their masses

$$m_e \equiv \frac{y_e v}{\sqrt{2}}, \quad m_\mu \equiv \frac{y_\mu v}{\sqrt{2}}, \quad m_\tau \equiv \frac{y_\tau v}{\sqrt{2}}. \tag{A.90}$$

Importantly, the *flavour neutrino fields* are then defined as the fields which couple directly to the charged leptons fields l_e, l_μ, l_τ in the charged current weak interaction. This requires to rotate the initial neutrino fields as:

$$\boldsymbol{\nu}_L \equiv V_L^{l\dagger} \boldsymbol{\nu}'_L = \begin{pmatrix} \nu'_{eL} \\ \nu'_{\mu L} \\ \nu'_{\tau L} \end{pmatrix} \tag{A.91}$$

With this redefinition, the leptonic charged current connects each charged lepton field with the flavour neutrino field of the same generation: there is no lepton mixing, and the lepton flavour number is strictly conserved. The diagrams of the first line of Fig. A.5 can be extended to the second and third generation, with the doublets (μ^-, ν_μ), (μ^+, $\overline{\nu}_\mu$), (τ^-, ν_τ) and (τ^+, $\overline{\nu}_\tau$).

For quarks, the situation is different because there is a right-handed singlet for both the down-type and up-type quarks and the Lagrangian thus contains more free physical parameters. The Higgs-quark coupling terms give rise to mass terms of the form

$$\mathcal{L}_{\text{mass},Q} = \frac{v}{\sqrt{2}} \Big(\sum_{\alpha,\beta=d,s,b} Y'^{D}_{\alpha\beta} \, \overline{q'^{D}} \, q'^{D}_{\beta R} + \sum_{\alpha,\beta=u,c,t} Y'^{U}_{\alpha\beta} \, \overline{q'^{U}} \, q'^{U}_{\beta R} \Big) + \text{h.c.} \tag{A.92}$$

The same procedure can be used as for the charged leptons. However, there are now two different Yukawa matrices to diagonalise: Y'^{U} and Y'^{D}, applying separately to up-type and down-type quarks. As a consequence, it is impossible to redefine the down-type quark fields so that they couple in the weak charged current only with the up-type quark fields of the same generation. Therefore, there is a theoretical possibility for a *physical* quark mixing. As introduced in Sec. 1.1.1 (see Fig. A.4), this means for instance that a strange quark (second generation) can decay into an up quark (first generation) via the weak charged current coupling of the fields s and u, which are the fields with definite mass. In fact, many hadron decays observed experimentally, even before the introduction of the quark model, are explained by quark mixing. The quark mixing matrix is called the *Cabbibo-Kobayashi-Maskawa* (CKM) matrix. It gives the strength of each of the possible weak CC couplings between physical quarks.

A.2.5 The complete Standard Model of particle physics

As we have seen, the way the fermions interact is determined by the choice of gauge group of the theory (up to value of the coupling constants which are free parameters). The content of the model in terms of fermion fields is chosen heuristically to meet the experimental observations.

There remains an additional number of free parameters in the model, which must be determined experimentally instead of being predicted from first principles.

Most of the free parameters come from the "Higgs sector", *i.e.* the part of the Lagrangian which describes the Higgs field and its coupling with the other fields. The strengths of the Higgs self-coupling terms (2 parameters) determine the vacuum expectation value v of the Higgs field and the Higgs mass m_H. These two parameters, along with the gauge couplings g and g' of Eq. A.69, determine the Weinberg angle as well as the $W^{\pm}$ and Z^0 masses. Along with the Higgs vev, the Yukawa coupling parameters determine the masses of the quarks and charged leptons (9 parameters), as well as the quark mixing parameters (4 parameters). In total, this makes for 17 free physical parameters in the electroweak SM.

The so-called Standard Model of particle physics is completed with the description of the strong interaction by the theory of *quantum chromodynamics* (QCD), whose gauge group is $\mathbf{SU}(3)$. As a Lie group, $\mathbf{SU}(3)$ is eight-dimensional, and the corresponding eight gauge boson fields are called *gluons*. The SM gauge group is thus extended to $\mathbf{SU}(3)_C \times \mathbf{SU}(2)_L \times \mathbf{U}(1)_Y$. The subscript C ("color") denotes an additional degree of freedom which, unlike the weak isospin and weak hypercharge, only concerns quarks. Each quark field transforms under the fundamental representation of $\mathbf{SU}(3)_C$, *i.e.* as a *color triplet*.

QCD can be studied separately from the electroweak SM, because the spontaneous symmetry breaking of the electroweak sector does not affect the $\mathbf{SU}(3)_C$ part. As a consequence, gluons remain massless. QCD adds only two empirical parameters to the whole Standard Model, bringing the total to 19. QCD is a rich model and is challenging to study. One of the reasons is that the group $\mathbf{SU}(3)$ is non-abelian, like $\mathbf{SU}(2) \times \mathbf{U}(\mathbf{1})$ but unlike $\mathbf{U}(1)_{\mathrm{QED}}$. Therefore the gluons can interact with themselves instead of being solely "force carriers". Moreover the QCD coupling constant is not small, preventing the use of perturbative methods at low energy. The understanding of QCD is still a very active area of theoretical research.

Appendix B

Sensitivity overestimation with sparse Monte Carlo

Contents

In particle or astroparticle physics, when assessing the sensitivity of future experiments, Monte Carlo (MC) simulations are extensively used in order to simulate the "response" of detectors (e.g. detection efficiency and errors on measured parameters). We consider the case where a "full MC" method is used to simulate this response rather than using an approximation based on parametrized (smooth) response functions. It is often stated that using an insufficient number of MC events to sample the detector response induces a systematic bias in the sensitivity estimation, yielding (in average) an *overestimation* of the sensitivity.

In Sec. B.1, we introduce the principle of the study along with some formal definitions. An example of empirical observation which motivated this study is also given. In Sec. B.2, we perform an asymptotic expansion for the $\Delta\chi^2$ function with respect to the fluctuations of the response matrix. Under reasonable assumptions, which are detailed in B.2.2, the response matrix fluctuations follow simple Poisson statistics, which allows us to conclude on the (second order) asymptotic behaviour of the $\Delta\chi^2$ function with respect to the number of MC events. Some elements of interpretation are given in Sec. B.2.3, along with prospects for verifying and using this formula in practice. We also discuss in Sec. B.3.2 the validity of our mathematical approach and

its limits, and give additional details about the motivation for our choice of sensitivity estimator and how it relates to a usual hypothesis testing procedure based on a binned likelihood ratio test statistic.

B.1 Motivation and principle of the study

B.1.1 Formal definitions

Our goal is to estimate the statistical sensitivity of a *future* experiment trying to distinguish between 2 physical models (hypotheses) A and B. In this study we consider that A and B correspond to definite theoretical predictions that we can compute with perfect accuracy. The probability of error on the decision to accept or reject hypotheses A or B then fully comes from the errors on reconstructed variables, and statistical fluctuations in the data. Thus we do not consider any source of systematic uncertainty. The predictions A and B consist in the distribution of *events* in terms of the true values of D physical observables. We consider these distributions in a discretized form: histograms with M bins, and we denote α_i^A (resp. α_i^B) $\in]0;+\infty[$ the predicted average number of events in bin i ($i = 1..M$) under hypothesis A (resp. B). These are referred to as the distributions of *true* events.

The experiment consists in the detection and selection of these events and measurement of the associated variables. Each event can be either selected or not selected, and we associate to each selected event a measured (or *reconstructed*) value for each of its D variables – with some error with respect to the true values. Similarly, we can build the D-dimensional histogram of *reconstructed* events, with a total of M' bins: we denote N_j^A (resp. N_j^B) $\in]0;+\infty[$ the predicted average number of events in bin j ($j = 1..M'$) under hypothesis A (resp. B).

The experiment itself is modelled by the *detector response matrix*, a $M' \times M$ matrix (Ω_{ji}). The entry Ω_{ji} is defined as the discrete conditional probability for an event in true bin i to be reconstructed in bin j:

$$\Omega_{ji} = P(\text{Reco } j \,|\, \text{True } i) = \frac{\#(\text{Reco } j \cap \text{True } i)}{\#(\text{True } i)} \tag{B.1}$$

such that:

$$N_j = \sum_{i=1}^{M} \Omega_{ji} \alpha_i \tag{B.2}$$

For a fixed true bin i, it is illustrative to separate the action of the response matrix in two parts. First, it models the detection efficiency, accounting for the fact that a fraction of these events will not be detected, or discarded by the analysis due to poor reconstruction quality. The detection efficiency in true bin i is defined as:

$$e_i = \sum_{j=1}^{M'} \Omega_{ji} < 1 \tag{B.3}$$

Then the normalized distribution in reconstructed space $(\frac{\Omega_{ji}}{e_i})_{j=1..M'}$ is analogous to a probability density for the response to the true bin i. With this point of view the true events distribution is first multiplied by a simple efficiency function, then convoluted with a normalized response function, which depends on i (truth) as well.

An example of physical motivation is given in Sec. B.1.3 along with graphical illustrations. At this point we can make the following additional remarks:

- We make no assumption on the number of dimensions D of the histogram, which corresponds to the number of distinct physical variables considered in our analysis.
- For simplification one can assume $M' = M$, although in practice we will often chose $M' > M$ since the physical model can be evaluated with good accuracy, whereas we will tend to chose a larger bin size in the reconstructed space due to limited resolutions (i.e. large reconstruction errors) and in order to reduce the amplitude of statistical fluctuations.

See Fig B.1 for a summary of the notations.

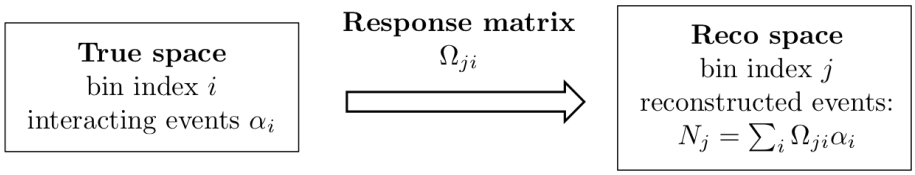

Figure B.1: Summary of the notations.

B.1.2 Statistical sensitivity

Typical experiments may detect a relatively small amount of events: in that case the number of *true* (resp. *reconstructed*) events in each bin will follow Poisson statistics with mean value α_i (resp. n_j). Assuming that the true event distributions ($\alpha_i^{A/B}$) as well as the response matrix (Ω_{ji}) are perfectly known, the error in intepreting an experimental result is then fully due to these statistical fluctuations. We then use the following classical $\Delta\chi^2$ function:

$$\Delta\chi^2 = 2\sum_{j=1}^{M'}\left[N_j^B - N_j^A + N_j^A \ln\frac{N_j^A}{N_j^B}\right] \tag{B.4}$$

This formula derives from a likelihood ratio analysis, assuming that the observed number of events in each bin follows Poisson statistics:

$$\Delta\chi^2{}_H = -2\ln\left[\frac{L(\text{data}(H)|B)}{L(\text{data}(H)|A)}\right]$$

Where we take the average expectation under hypothesis H as fictitious "data". This is also known as evaluating the sensitivity on the *Asimov dataset*. In what follows we will arbitrarily take $H = A$). See Chap. 6 for further details regarding the choice of $\Delta\chi^2$ and its translation in terms of confidence level for hypothesis testing and parameter estimation.

B.1.3 Sampled response matrix and empirical observation

We consider the case where the response matrix entries are estimated thanks to extensive MC simulations: individual events are generated, propagated to the detector, reconstructed and selected, and the entries Ω_{ji} are built according to eq. (B.1). In such simulations, the number of events $\mathcal{N}$ used to sample the response can be limited by the available computing power. If the number of MC events used is insufficient with respect to the size of the parameter space, the estimated matrix is said to be *sparse*: its entries will show statistical fluctuations with respect to the true physical response. It is generally stated that estimating the sensitivity of an experiment using insufficient MC statistics can only lead to an overestimation.

In order to check the dependency of the sensitivity estimation on $\mathcal{N}$ in a practical case, one can evaluate the $\Delta\chi^2$ using subsamples of MC events of size $\mathcal{N}' < \mathcal{N}_0$, where $\mathcal{N}_0$ is the total available number of simulated MC events. Such subsamples can be obtained by randomly drawing subsamples of the full available set of MC events.

Fig. B.3 shows an example of such a study in the case of the sensitivity of the ORCA experiment to the determination of the neutrino mass ordering. Hypothesis A corresponds to normal ordering (NO), while hypothesis B corresponds to inverted ordering (IO). In the case $D = 2$ dimensions (blue curve), the variables considered are neutrino energy (40 bins) and neutrino zenith angle (40 bins). In the case $D = 3$ dimensions (red curve), an additional variable is considered: the neutrino interaction inelasticity ($0 < Y_{\text{Bjorken}} < 1$), measured in 4 uniform bins. In addition, in both cases the events are split into discrete channels corresponding to neutrino flavour and interaction channel (for the true space) and track/cascade classification (for the reconstructed space). This yields a total of $M' = 3\,200$ bins for the 2D analysis and $M' = 12\,800$ bins for the 3D analysis.

It can be observed in Fig. B.2 that for a given value of $\mathcal{N}'/\mathcal{N}_0 < 1$, the $\Delta\chi^2$ value shows fluctuations depending on the subsample used. As a general trend the average $\Delta\chi^2$ value is increasing when reducing the size of the MC sample. This means that the sensitivity will on average be overestimated in case the response matrix is estimated using an insufficient number of MC events.

It is known that the Bjorken-Y variable has a potential for improving the sensitivity since it will be distributed differently under hypotheses A (normal ordering) and B (inverted ordering). The motivation of our study is rather clear from Fig. B.3: it appears that for $D = 2$ the end point corresponding to the full MC set approximately lies in a *plateau* region, hence the sensitivity seems to be correctly estimated, while in the $D = 3$ case the sensitivity might be overestimated with a rather large bias.

It is tempting to extrapolate this behaviour in order to infer the realistic gain in adding the information on reconstructed Y_{Bjorken}. Hence we tried to fit both curves to a power law of the following form:

$$\Delta\chi^2(\mathcal{N}') = p_0 + p_1 \times (\mathcal{N}'/\mathcal{N}_0)^{p_2} \tag{B.5}$$

When using only data points relatively close to the "full set" ($0.1 \leq \mathcal{N}'/\mathcal{N}_0 \leq 1$) for the fit and keeping all 3 parameters free, the agreement with a power law is very good with a value of $p_2 \simeq -1$. The resulting fit function is shown in Fig. B.3 along with its extrapolation to a 100 times larger MC sample. This motivates a theoretical study of the problem. Indeed, if this behaviour is expected from statistical theory, one could imagine using such a method to extrapolate to "infinite Monte Carlo" and take:

$$\Delta\chi^2(\mathcal{N}' \to \infty) = p_0$$

In particular if we expect an index of exactly $p_2 = -1$, the fit could be done with only p_0 and p_1 as free parameters, yielding better precision on the estimate of $\Delta\chi^2(\mathcal{N}' \to \infty)$.

B.1.4 Position of the problem

In order to explain the $\Delta\chi^2$ dependency on the number of MC events, we use a probabilistic model for the response matrix sampling: its entries are considered as random variables, denoted ω_{ji}, verifying:

$$\langle \omega_{ji} \rangle = \Omega_{ji}$$

which can also be put as:

$$\lim_{\mathcal{N} \to +\infty} \omega_{ji} = \Omega_{ji}$$

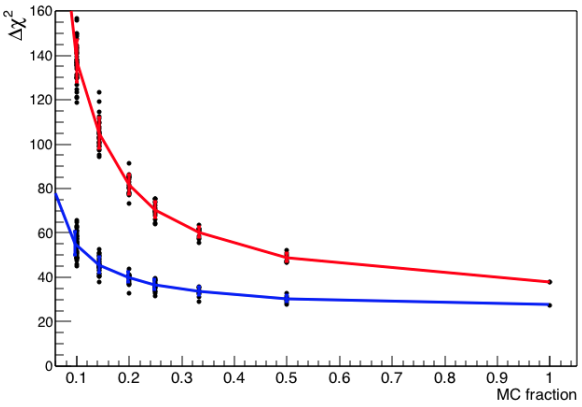

Figure B.2: $\Delta\chi^2$ vs. $\mathcal{N}'/\mathcal{N}_0$. The blue curve corresponds to $D=2$ ($M'=2\times40\times40$) and the red curve to $D=3$ ($M'=2\times40\times40\times4$, Y_{Bjorken} added). Individual trials in transparent grey, sample average and standard deviation in colored error bars.

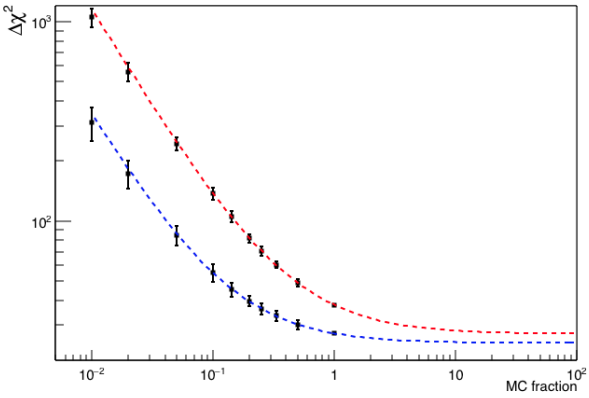

Figure B.3: 3-parameter fits of the sample average $\langle\Delta\chi^2\rangle(\mathcal{N}')$ with functions of the form $p_0+p_1\times(\mathcal{N}'/\mathcal{N}_0)^{p_2}$. The fit range is restrained to $0.1\leq\mathcal{N}'/\mathcal{N}_0\leq1$ and p_0, p_1 and p_2 are fitted, accounting for error bars (sample standard deviation).

On the contrary the physical predictions $\alpha_i^{A/B}$ are deterministic average values and do not depend on $\mathcal{N}$. We denote the response fluctuation as $\delta\omega_{ji}$:

$$\delta\omega_{ji}=\omega_{ji}-\Omega_{ji}$$

With these notations we have

$$\langle\delta\omega_{ji}\rangle=0 \tag{B.6}$$

$$\langle{\delta\omega_{ji}}^2\rangle=\text{Var}(\omega_{ji}) \tag{B.7}$$

The dependence of $\langle{\delta\omega_{ji}}^2\rangle$ on $\mathcal{N}$ is examined in Sec. B.2.2. Similarly to eq. (B.2) we define:

$$n_j=\sum_{i=1}^{M}\omega_{ji}\alpha_i=\sum_{i=1}^{M}(\Omega_{ji}+\delta\omega_{ji})\alpha_i \tag{B.8}$$

and we write

$$\delta n_j=\sum_{i=1}^{M}\delta\omega_{ji}\alpha_i$$

such that

$$n_j^{A/B}=N_j^{A/B}+\delta n_j^{A/B}$$

The $\Delta\chi^2$ function, evaluated on the predictions obtained with the sampled response matrix, is itself a random variable:

$$\Delta\chi^2 = 2\sum_{j=1}^{M'}\left[n_j^B - n_j^A + n_j^A \ln\frac{n_j^A}{n_j^B}\right] = \sum_{j=1}^{M'}\Delta\chi^2{}_j \tag{B.9}$$

Our goal is to demonstrate that the expectation $\langle\Delta\chi^2\rangle$ depends on $\mathcal{N}$ and behaves in the following form when $\mathcal{N}\to\infty$:

$$\langle\Delta\chi^2\rangle(\mathcal{N}) = \Delta\chi^2_\infty + K\times\frac{1}{\mathcal{N}} + \mathrm{o}\Big(\frac{1}{\mathcal{N}}\Big) \quad \text{with } K>0 \tag{B.10}$$

B.2 Statistical model

B.2.1 Asymptotic expansion

We start by writing explicitly the expression of $(\Delta\chi^2)_j$ (as defined in eq. (B.9)) by injecting the expression of the reconstructed numbers of events $n_j^{A/B}$ according to qq. (B.8). We then make an asymptotic expansion of the obtained equation, assuming that in the limit $\mathcal{N}\to\infty$:

$$\boxed{|\delta\omega_{ji}| << \Omega_{ji} \quad \text{(i.e.)} \quad |\epsilon_{ji}| = \frac{|\delta\omega_{ji}|}{\Omega_{ji}} << 1} \tag{B.11}$$

The validity of this assumption in practical examples will be discussed in Sec. B.3.2. Note that this assumption implies that

$$\delta n_j << N_j$$

First order development

To go to first order we can directly use the Taylor development:

$$\ln\left(N_j+\delta n_j\right) \simeq \ln N_j + \frac{\delta n_j}{N_j}$$

For readability, the index j will be omitted from now on, since the development is identical for all $(\Delta\chi^2)_j$. The first order development then reads:

$$\Delta\chi^2{}_{(j)} = 2\left(N^B - N^A + N^A\ln\frac{N^A}{N^B} + \delta n^A - \delta n^B + \delta n^A\ln\frac{N^A}{N^B} + N^A\Big(\frac{\delta n^A}{N^A} - \frac{\delta n^B}{N^B}\Big)\right) + \mathcal{O}\Big((\frac{\delta n}{N})^2\Big) \tag{B.12}$$

Where we identify

$$(\Delta\chi^2_\infty)_{(j)} = 2\Big(N^B - N^A + N^A\ln\frac{N^A}{N^B}\Big)$$

Let us evaluate now the expectation of eq. (B.12). Since, by definition,

$$\langle\delta n^A\rangle = \langle\delta n^B\rangle = 0$$

the expression reduces to:

$$\langle \Delta\chi^2{}_{(j)} \rangle = (\Delta\chi^2_\infty)_{(j)} + \mathcal{O}\Big(\langle (\frac{\delta n}{N})^2 \rangle\Big)$$

So the desired effect has to appear when pushing the development at higher order in $\frac{\delta n}{N}$, or more generally in $\frac{\delta\omega_i}{\Omega_i}$. Using eq. (B.7) we can anticipate that going to second order will indeed suffice if the variance of ω_i is similar to that of a Poisson-distributed variable:

$$\big\langle (\delta\omega_i)^2 \big\rangle = \mathrm{Var}(\omega_i) \sim \frac{K}{\mathcal{N}_i} \quad \text{with } K > 0 \tag{B.13}$$

Second order development

To go to second order we choose to write the δn's in terms of the $\delta\omega_i$ since we expect crossed products of the form $\delta\omega_{i_1} \cdot \delta\omega_{i_2}$ to appear which may contribute to the effect. We will then use a multiple variable Taylor development.

$$n = N + \sum_i \delta\omega_i \alpha_i = N\big(1 + \sum_i \eta_i\big) \quad \text{where} \quad \eta_i^{(A/B)} \equiv \frac{\delta\omega_i \alpha_i^{(A/B)}}{N^{(A/B)}} = \frac{\delta\omega_i \alpha_i^{(A/B)}}{\sum_i \Omega_i \alpha_i^{(A/B)}}$$

We denote $\vec{\eta} = (\eta_1, \eta_2, ..., \eta_n)$ and $f(\vec{\eta}) = \ln\big(1 + \sum_i \eta_i\big)$. The condition $\delta\omega_i << \Omega_i$ implies $\eta_i << 1$. The quadratic Taylor development of f for $\|\vec{\eta}\| \to 0$ reads:

$$f(\vec{\eta}) = f(\vec{0}) + \sum_i \frac{\partial f}{\partial \eta_i}(\vec{0}) \cdot \eta_i + \frac{1}{2} \sum_{i,j} \frac{\partial^2 f}{\partial \eta_i \partial \eta_j}(\vec{0}) \cdot \eta_i^2 + \mathcal{O}(\|\vec{\eta}\|^3)$$

With here:

$$\frac{\partial f}{\partial \eta_i} = \frac{1}{1 + \sum_i \eta_i} \to \frac{\partial f}{\partial \eta_i}(0) = 1$$
$$\frac{\partial^2 f}{\partial \eta_i \partial \eta_j} = -\frac{1}{(1 + \sum_i \eta_i)^2} \to \frac{\partial^2 f}{\partial \eta_i \partial \eta_j}(0) = -1$$

Let us finish the development. We write

$$\Delta\chi^2{}_{(j)} = (\Delta\chi^2_\infty)_{(j)} + 2\Big(\delta n^A - \delta n^B + \delta n^A \ln\frac{N^A}{N^B} + N^A\Big[f(\vec{\eta}^A) - f(\vec{\eta}^B)\Big] + \delta_n^A\Big[f(\vec{\eta}^A) - f(\vec{\eta}^B)\Big]\Big)$$

where:

$$\begin{aligned}
&N^A\Big[f(\vec{\eta}^A) - f(\vec{\eta}^B)\Big] + \delta_n^A\Big[f(\vec{\eta}^A) - f(\vec{\eta}^B)\Big] = \\
&N^A\Big[\sum_i \delta\omega_i\big(\frac{\alpha_i^A}{N^A} - \frac{\alpha_i^B}{N^B}\big) - \frac{1}{2}\sum_i (\delta\omega_i)^2\Big(\big(\frac{\alpha_i^A}{N^A}\big)^2 - \big(\frac{\alpha_i^B}{N^B}\big)^2\Big) - \sum_{i_1 \neq i_2} \delta\omega_{i_1}\delta\omega_{i_2}\big(\frac{\alpha_{i_1}^A \alpha_{i_2}^A}{(N^A)^2} - \frac{\alpha_{i_1}^B \alpha_{i_2}^B}{(N^B)^2}\big)\Big] \\
&+ \Big(\sum_i \delta\omega_i \alpha_i^A\Big) \sum_i \delta\omega_i \big(\frac{\alpha_i^A}{N^A} - \frac{\alpha_i^B}{N^B}\big) + \mathcal{O}(\|\vec{\eta}\|^3)
\end{aligned} \tag{B.14}$$

Now we evaluate the expectation of the previous expression. We have $\langle \delta\omega_i \rangle = 0$ by definition, and moreover we will assume:

$$\boxed{\text{for } i_1 \neq i_2, \quad \langle \delta\omega_{i_1} \delta\omega_{i_2} \rangle = \langle \delta\omega_{(j)i_1} \delta\omega_{(j)i_2} \rangle = 0} \tag{B.15}$$

The fluctuations of the response matrix entries are in principle anticorrelated between different (neighbouring) ***reconstructed*** bins: if an event is reconstructed in bin j_1, it is not reconstructed in bin j_2. On the contrary, the hypothesis (B.15) means that these fluctuations are uncorrelated between different ***true*** bins, which is to be expected (see Sec. B.2.2 for a discussion of this hypothesis).

Thus all first order and crossed second order terms in (B.14) vanish, and the expectation reduces to:

$$\langle \Delta\chi^2{}_{(j)} \rangle = (\Delta\chi^2_\infty)_{(j)} - 2\sum_i \left[\frac{N^A}{2} \Big((\frac{\alpha_i^A}{N^A})^2 - (\frac{\alpha_i^B}{N^B})^2 \Big) - \alpha_i^A \Big(\frac{\alpha_i^A}{N^A} - \frac{\alpha_i^B}{N^B} \Big) \right] \langle (\delta\omega_i)^2 \rangle + \mathcal{O}\Big(\langle (\frac{\delta\omega}{\Omega})^3 \rangle \Big)$$

$$\langle \Delta\chi^2{}_{(j)} \rangle = (\Delta\chi^2_\infty)_{(j)} + \sum_i N^A \Big(\frac{\alpha_i^A}{N^A} - \frac{\alpha_i^B}{N^B} \Big)^2 \langle (\delta\omega_i)^2 \rangle + \mathcal{O}\Big(\langle (\frac{\delta\omega}{\Omega})^3 \rangle \Big)$$

Finally, after summing on the reconstructed bins (index j, no longer omitted):

$$\langle \Delta\chi^2 \rangle = \sum_j \langle \Delta\chi^2{}_j \rangle = \Delta\chi^2_\infty + \sum_j \sum_i N_j^A \Big(\frac{\alpha_i^A}{N_j^A} - \frac{\alpha_i^B}{N_j^B} \Big)^2 \langle (\delta\omega_{ji})^2 \rangle + \mathcal{O}\Big(\langle (\frac{\delta\omega}{\Omega})^3 \rangle \Big) \tag{B.16}$$

Assuming that (B.13) is realized (see Sec. B.2.2), and that $\mathcal{N}_{i_1} \simeq \mathcal{N}_{i_2}$ for all i_1, i_2 (which is usually obtained with a relevant choice of binning and MC generation), we obtain the behaviour of eq. (B.10).

In the next section we will detail a probabilistic model for the sampled response matrix, which will allow us to make (B.16) more explicit. We will then be able to come back on the interpretation of this equation.

B.2.2 Probabilistic model for the response matrix

By definition

$$\omega_{ji} = \frac{\widetilde{n}_{\text{reco}}(j,i)}{\widetilde{N}_{\text{gen}}(i)} = \frac{\widetilde{n}_{\text{reco}}(j,i)}{\mathcal{N}_i} \tag{B.17}$$

where

- $\widetilde{N}_{\text{gen}}(i) = \mathcal{N}_i$ is the number of MC events generated in true bin i
- $\widetilde{n}_{\text{reco}}(j,i)$ is the number of MC events generated in true bin i that are detected and reconstructed in bin j

We will consider here that $\widetilde{N}_{\text{gen}}(i) = \mathcal{N}_i$ is under our control and not random. In practice, MC events are usually generated randomly according to some predefined continuous distribution, but if $\mathcal{N}_i$ is large enough the fluctuations of ω_{ji} will be effectively dominated by the fluctuations in $\widetilde{n}_{\text{reco}}(j,i)$.

On the contrary $\widetilde{n}_{\text{reco}}(j,i)$ is a random variable, taking integer values. Its expectation is denoted $\widetilde{N}_{\text{reco}}(j,i)$. As explained in Sec. B.1.1, we can separate the detector response into efficiency and reconstruction error:

$$\widetilde{N}_{\text{reco}}(j,i) = P_{ji} \cdot e_i \cdot \mathcal{N}_i$$

where e_i is the total selection efficiency as defined in (B.3) and P_{ji} is the probability for a *selected* event to be reconstructed in bin j. It can be relevant to further express this probability as

$$P_{ji} = \int_{\text{bin } j} \Phi_i(X) \mathrm{d}X$$

where Φ_i would be the underlying D-dimensional probability density function modeling the detector resolution for the true bin i. In a parametrized model of detector response (not the approach chosen here), the most basic choice to approximate Φ_i would be a D-dimensional normal PDF. In general, Φ_i cannot be derived from first principles, and we thus have to rely on MC simulations to sample it. Note that the probability P_{ji} then depends on the choice of binning in reconstructed space. Indeed it can be approximated as:

$$P_{ji} \simeq \Phi_i(X_j) \cdot V_j$$

where X_j and V_j are resp. the center and the volume of the D-dimensional reconstructed bin j.

It is natural to model the sampling of the response matrix with MC simulation as a repetition of independent random drawings of events according to the underlying distribution that we just described. Then the number $\widetilde{n}_{\text{reco}}(j,i)$ of reconstructed events drawn in bin j will follow a Poisson distribution with mean $\widetilde{N}_{\text{reco}}(j,i)$. Hence, its variance will be $\widetilde{N}_{\text{reco}}(j,i)$ as well. Thus:

$$\left\langle (\delta\omega_{ji})^2 \right\rangle = \mathrm{Var}(\omega_{ji}) = \mathrm{Var}\left[\frac{\widetilde{n}_{\text{reco}}(j,i)}{\mathcal{N}_i} \right] = \frac{\mathrm{Var}(\widetilde{n}_{\text{reco}}(j,i))}{{\mathcal{N}_i}^2} \simeq \Phi_i(X_j) \cdot V_j \cdot e_i \cdot \frac{1}{\mathcal{N}_i} \tag{B.18}$$

B.2.3 Final result and interpretation

We can finally inject the expression (B.18) into eq. (B.16). We obtain:

$$\langle \Delta\chi^2 \rangle = \Delta\chi^2_\infty + \sum_j \sum_i \frac{K_{ji}}{\mathcal{N}_i} + \mathcal{O}\Big(\frac{1}{\mathcal{N}^{3/2}}\Big) \tag{B.19}$$

where

$$\boxed{K_{ji} = N_j^A \Big(\frac{\alpha_i^A}{N_j^A} - \frac{\alpha_i^B}{N_j^B} \Big)^2 \cdot \Phi_i(X_j) \cdot V_j \cdot e_i > 0} \tag{B.20}$$

Let us introduce the fraction of MC events generated in bin i:

$$f_i = \frac{\mathcal{N}_i}{\mathcal{N}} \tag{B.21}$$

The binning and MC generation are often chosen together such that the number of generated events per true bin is approximately uniform $f_i \simeq 1/M$. Here, it will be sufficient to assume that the fraction remains constant as the total number $\mathcal{N}$ varies.

We then get the final asymptotic expansion in the limit $\mathcal{N} \to \infty$:

$$\boxed{\langle \Delta\chi^2 \rangle(\mathcal{N}) = \Delta\chi^2_\infty + \frac{K}{\mathcal{N}} + \mathcal{O}\Big(\frac{1}{\mathcal{N}^{3/2}}\Big)} \tag{B.22}$$

with

$$\boxed{K = \sum_j \sum_i \frac{K_{ji}}{f_i} > 0} \tag{B.23}$$

The positiveness of the coefficient K shows that on average, the sensitivity will indeed be overestimated – provided that the contribution of terms of order 3/2 and further in $1/\mathcal{N}$ can be neglected. We also obtain the approximate $1/\mathcal{N}$ dependance that matches the empirical observation of Sec. B.1.3.

B.3 Discussion

B.3.1 Interpretation

In this section, we will first consider simplified, limiting cases in order to try and get a practical interpretation of the formula. Then we will discuss which parameters in the final result (B.19) are under our control and how the average overestimation depends on these parameters, so as to get a hint of how the MC generation and binning can be chosen to minimize this effect.

Case of a perfect response

For a hypothetical detector with extremely peaked resolution functions, one can consider the limit $P_{ji} \simeq \delta_{ji}$ i.e. $\Omega_{ji} \simeq \delta_{ji} e_i$. This can also be achieved by chosing bins that are very large with respect to the spread of the resolution functions. In that case: if $i \neq j$, $N_j^A = e_i \delta_{ji} \alpha_i^A = 0$ and $K_{ji} = 0$, and if $i = j$, then $\frac{\alpha_i^A}{N_j^A} = \frac{\alpha_i^B}{N_j^B} = \frac{1}{e_i}$ and $K_{ji} = 0$ as well. Thus, as expected, there is no overestimation effect: MC events can only fall in one bin, hence there cannot be any fluctuation. Of course, using Monte Carlo is anyway pointless in such a case.

True bins with no signal

Consider a true bin i having no difference between A and B i.e. $\alpha_i^A = \alpha_i^B$ (typically an event coming from a background channel or lying in a "control sample" region). It must be emphasized that this does not necessarily imply that $K_{ji} = 0$ for all reco bins j. Indeed, since the reconstructed event counts N_j^A and N_j^B potentially include contributions from any other true bin, we can have $N_j^A \neq N_j^B$.

Simple unbiased realistic response

Let us consider a simple case where true and reconstructed space are identical with the same binning ($M = M'$). Moreover we assume that the response functions are unbiased and reasonably peaked, verifying:

$$\text{if } i \neq j, \quad \Omega_{jj} >> \Omega_{ji}$$

Then, assuming all the α_k's are of the same order of magnitude:

$$N_j^A = \Omega_{jj} \alpha_j^A + \sum_{k \neq j} \Omega_{jk} \alpha_k^A \simeq \Omega_{jj} \alpha_j^A$$

This results again in $K_{jj} \simeq 0$, and for $j \neq i$, K_{ji} will approximately scale as:

$$K_{ji} \sim \left(\frac{\alpha_i^A}{\Omega_{jj}\alpha_j^A} - \frac{\alpha_i^B}{\Omega_{jj}\alpha_j^B} \right)^2 \cdot N_j^A \cdot \Phi_i(X_j) \cdot V_j \cdot e_i \tag{B.24}$$

In this approximation we can see that the K_{ji} term will contribute significantly to the overestimation effect when, for example, $\alpha_j^A > \alpha_i^A$ while $\alpha_j^B < \alpha_i^B$. This corresponds to the contribution of events misreconstructed from true bin i into reco bin j, when those two bins are relevant to the physical measurement (clear distinction between A and B in true space). Qualitatively, if the response is smooth we can expect the physical signal to be blurred, whereas with a sparse matrix it may artificially remain visible due to the non-smoothness of the response.

Choice of binning

Let us examine the impact of the choice of binning on the average overestimation of eq. (B.19). Among the terms appearing in K_{ji}, only e_i and $\Phi_i(X_j)$ are intrinsic to the underlying detector response. All the other terms scale linearly with the volume of the bins, which we can tune. Denoting V_i (resp. V_j) the volume of true bin i (resp. reco bin j) we have $\alpha_i \sim V_i$ and $N_j \sim V_j$, hence:

$$K_{ji} \sim s_{ji} \cdot {V_i}^2$$

where s_{ji} is an intrinsic (intensive) parameter which, in reasonable cases, would be related to the theoretical signal (separation between A and B) in true bin i. For example one can get an explicit approximation of s_{ji} starting from eq. (B.24). Denoting $s_j = \sum_j s_{ji}$ we further get:

$$K = \sum_j \sum_i K_{ji} \sim \sum_i \frac{V_i^2}{f_i} \cdot s_i \tag{B.25}$$

We can now use that $s_i = \sum_{j=1}^{M'} s_{ji}$ must, at first order, scale linearly with the number of reco bins M' since s_{ji} is by definition an intensive parameter. All in all we are left with

$$K \sim M' \cdot \sum_i \frac{{V_i}^2}{f_i} \tag{B.26}$$

Thus we see that the average overestimation scales linearly with the number of reco bins M'. This is relatively intuitive, and in agreement with the empirical observations. Moreover, in the simplifying assumption that the MC generation is uniform i.e. $f_i \sim V_i$, the binning of the true space has no influence since $\sum_i \frac{{V_i}^2}{f_i} \sim \sum_i V_i = V_{\text{tot}}$.

Choice of MC generation

Examining eq. (B.25) again, and assuming that for a given binning we are able to evaluate the s_{ji} parameters beforehand, it is then relevant to tune the MC generation such that the true bins where s_i is large are attributed an important fraction of the total MC statistics. This will allow to reduce the overestimation effect. In most cases, this will correspond to the bins where the theoretical signal is the strongest and the detector performance not excellent.

Extrapolating to infinite MC

The result of eq. (B.19) could in principle be used for predictions and to extrapolate from a finite MC sample to $\Delta\chi^2_\infty$. However, the formula involves the reconstructed number of events

obtained with the matrix Ω, which is by definition inaccessible if we do not have sufficient MC statistics. The extrapolation can be done without being able to predict the coefficient K: as long as the general asymptotic behaviour is reliable we can fit it from reduced statistics studies (see Sec. B.1.3). To test the actual predictive power of the formula, one possibility would be to use a MC sample sufficiently close to infinite. Such a sample could be obtained for example as a "dummy MC" generated from a parametrized model of detector response. Even with our current ORCA simulation MC sample, by adapting the binning one could create a situation of effectively infinite MC statistics. Such tests are in preparation and the results will be reported as soon as possible.

B.3.2 Discussion of mathematical hypotheses

Mathematical approach

We have done a Taylor development of the $\Delta\chi^2$ estimator with respect to the parameter $\epsilon_{ji} = \frac{\delta\omega_{ji}}{\Omega_{ji}}$ which was assumed to be small. The mathematical correctness of this approach can already be questioned: indeed $\delta\omega_{ji}$ is a random variable which values are not necessarily restrained to a small range. It would be more sound to operate on the variance of this random variable, which is a number. After that, we arrived at our conclusion by evaluating the expectation of the obtained Taylor development (getting back the variance of $\delta\omega_{ji}$). This interversion may pose mathematical problems as well.

Hypothesis of small fluctuations

The hypothesis $|\epsilon_{ji}| = \frac{|\delta\omega_{ji}|}{\Omega_{ji}} << 1$ cannot be realized in all bins i and j and for all MC random realisations. In some regions of the parameter space, the opposite may even be realised most of the time. Indeed, using the same notations as in Sec. B.2.2:

$$\frac{\delta\omega_{ji}}{\Omega_{ji}} = \frac{\delta\widetilde{n}_{\text{reco}}(j,i)}{\widetilde{N}_{\text{reco}}(j,i)} = \frac{\widetilde{n}_{\text{reco}}(j,i) - \widetilde{N}_{\text{reco}}(j,i)}{\widetilde{N}_{\text{reco}}(j,i)} \tag{B.27}$$

where $\widetilde{n}_{\text{reco}}(j,i)$ is a Poisson-distributed random variable with mean $\widetilde{N}_{\text{reco}}(j,i)$. The problems arise in the *tails* of the response distributions, where $\widetilde{N}_{\text{reco}}(j,i)$ can have an arbitrarily small value – consider for example $\Phi_i(X_j)$ a normal pdf. On the contrary, $\widetilde{n}_{\text{reco}}(j,i)$ only takes integer values. Hence in the case $\widetilde{N}_{\text{reco}}(j,i) << 1$, we will have most of the time $\widetilde{n}_{\text{reco}}(j,i) = 0$ i.e. $\epsilon_{ji} = -1$, and rarely $\widetilde{n}_{\text{reco}}(j,i) = 1$ yielding $\epsilon_{ji} \simeq \frac{1}{\widetilde{N}_{\text{reco}}(j,i)}$ which can be arbitrarily large !

The idea of the correct mathematical treatment would then be to argue that these "outlier" bins do not matter much to the final sensitivity estimation, since they are far from the central peak of the response. This would likely require more elaborate tools than what we have used here.

Appendix C

The Swim software package

Contents

An important achievement of the three years of work summarised in this thesis has been the production of a software framework called *Swim*, providing high-level tools for sensitivity studies (and later on, real data analyses) addressing the measurement of neutrino oscillations with ORCA or other similar detectors. It is suitable for 3-flavour oscillation analyses, as well as sterile neutrino and non-standard interaction studies. The main specificity of Swim is that the detector response model is entirely based on Monte Carlo simulations, as described in Chap. 5. Any sample of Monte Carlo events can be used, with some minimal requirements on the format of the analysis variables. The interface is intended to be flexible, allowing the user to input arbitrary definitions of selection cuts and event classes without modifying the core code. The software can therefore be used to compare the measurement performance associated to different detector configurations, triggering, reconstruction or selection strategies in a unified framework and "automatised" manner, by running the same analysis on several input MC samples.

The software is now an official tool in the KM3NeT collaboration. I have used it to produce official sensitivity estimates for the Mass Hierarchy determination which were publicly presented in the Neutrino 2018 conference [275] [276], as well as Earth tomography studies for which a first version was presented in Neutrino 2016 [277] and a more complete treatment was presented in ICRC 2017 [292]. In addition, it has been used at APC by Joao Coelho for investigating ORCA's potential for sterile neutrinos and non-standard interaction searches, for which results have been presented publicly at several occasions [293] [294]. PhD students in APC and other teams of the collaboration have recently taken over these studies and started new ones making use of the Swim package, as detailed in Sec. C.2.

This appendix is intended to make a general overview of the software, detailing aspects (design, interface, optimisations) that are not mentioned in the main text. I have tried to limit the amount of technical details in most sections, hopefully to a level suitable to readers that are not also users of the software. For current and future users and developers of the software I hope

that this appendix can be useful in providing connections between the mathematical descriptions of Chap. 4 to 7 and their implementation and corresponding interface, complementing the Swim manual and the Doxygen documentation.

C.1 Design and optimisations

The software is written in C++ with an object-oriented design. It relies on the ROOT data analysis framework [260] for data formats (`TTree`) and containers such as standard histograms (`THx`), sparse higher-dimensional histograms (`THnSparse`), graphs, interpolation, fitting, *etc.* For computation of oscillation probabilities the *OscProb* software is used [128]. OscProb was initially written by Joao Coelho as a tool for long-baseline oscillation experiments. In the context of its use in the KM3NeT collaboration and in particular in Swim, it has been extended for simulating oscillations of atmospheric neutrinos with a detailed Earth model implementation The Minuit2 library [272] is used for minimisation.

C.1.1 Structure and main components

The source code is subdivided into the following packages:

- `OrcaSim`: main classes for simulating experiments, gluing together all the lower-level objects from following packages
- `OscProb`: oscillation probability calculator (external)
- `InputModels`: flux, cross-section and Earth models
- `OrcaDet`: classes for simulating the detector response based on MC events
- `NuFitter`: parameter classes and interface with the Minuit2 minimiser
- `OrcaSyst`: systematics manager for `OrcaSim`
- `UtilsOrca`: common definitions and utilities
- `Data`: common location for data files (input ROOT files and cached response matrices) and data processing scripts
- `Applications`: tutorial examples, general examples, and testing suite

A flowchart symbolising the interaction of the packages can be seen in Fig. C.1.

The central Experiment object

The central class of the framework is the `Experiment`. A typical analysis has one `Experiment` object representing pseudo-data (be it Asimov or a pseudo-experiment), and one holding the model prediction used to fit a model to the data. A basic `Experiment` holds

- oscillograms (two-dimensional histograms (E,θ)) for each oscillation scenario
- distributions for interacting event rates per Mt.y (three-dimensional histograms $(E,\theta\ y)$)
- distributions for the reconstructed event numbers in each event class (following user-provided event class definitions), under the form of (E',θ',y') histograms normalised to a full detector exposure with set data-taking time (three years by default)

The bin-by-bin error due to Monte Carlo statistical uncertainties is computed along with reconstructed event numbers (defining a mirror set of histograms), and used in subsequent χ^2 calculations, if the corresponding option is explicitly enabled. Intermediate histograms, *e.g.* holding the distributions of detected events with respect to true variables or assuming perfect particle identification are created and filled optionnally as well.

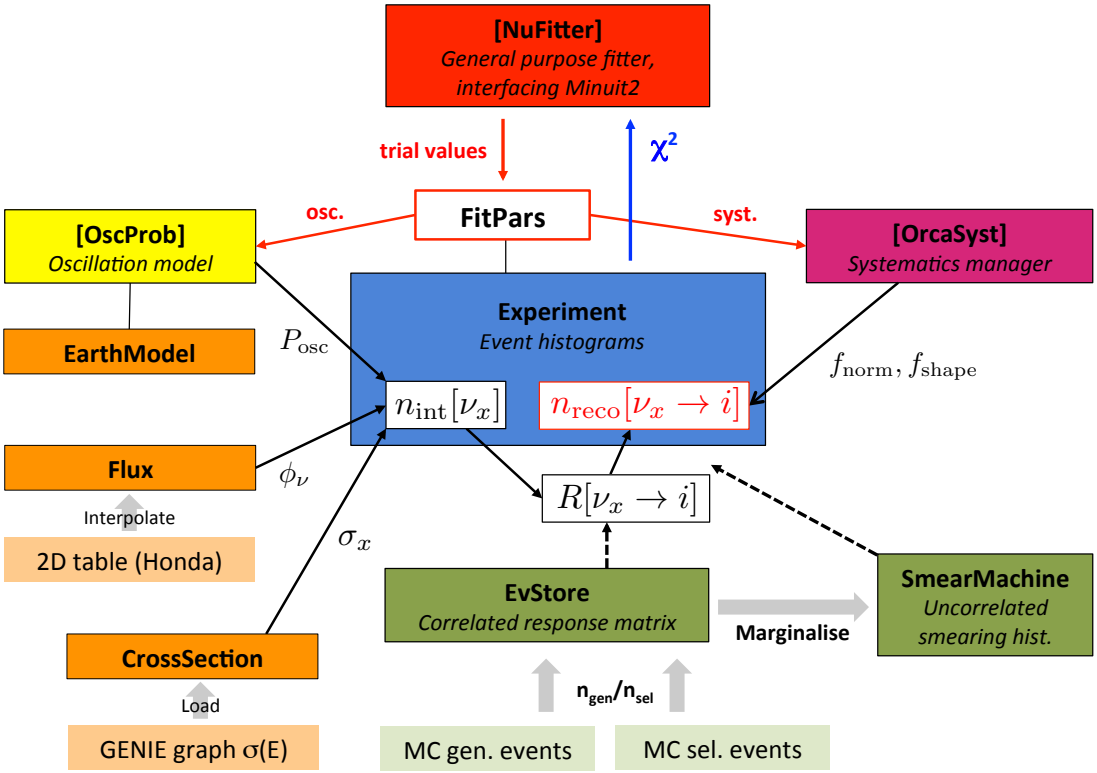

Figure C.1: Flowchart describing the interaction between the main classes in the Swim framework. Class names are in bold type, while common packages are represented by background colours, and package names are in bold type and in square brackets. Boxes in lighter background colour indicate user input.

Computation steps when fitting an `Experiment`

The steps of the computation chain internal to the `Experiment` class can be summarised in the following way:

1. Computation of the oscillation probabilities
2. Computation of the interacting event rates (essentially multiplying the oscillograms with flux and cross-sections, applying relevant systematics)
3. Application of the detector response, using either the `EvStore` correlated response matrix or the uncorrelated `SmearMachine`; after this step one reconstructed event histogram per interaction channel and event class (*e.g.* ν_e CC events classified as cascade-like) is obtained
4. Recombination of the channels, applying relevant normalisation systematics

Parameters and other members

All parameters of an `Experiment` susceptible to be fitted are held by an object of the `FitPars` class, which derives from the Minuit2 parameter class and is used to interact with the fitter.

In addition, an Experiment object holds the following members:

- an object of the base class `OscProb::PMNS_Base` which is instantiated to the corresponding derived type depending on the type of `Experiment` (standard, sterile, NSI),

- an `EarthModel` object which precalculates the baselines to be passed to the oscillation calculator, following the θ binning definition (both a custom implementation and a wrapper to `OscProb::PremModel` are available),
- an `EvBins` object (detailed below),
- an `IntRate` object which interfaces it with the `Flux` and `CrossSection` objects,
- an `EvStore` which loads MC events into the response matrix structure described in Chap. 5, depending on event classes definitions and selection cuts defined by the user through a classifier object (class `EvClasses`)
- optionnally when the uncorrelated smearing method is used, an object holding the "smearing machinery" is created (class `SmearMachine`) and builds the efficiency and smearing histograms as projections of the response matrix

Bins and oversampling

The energy variable used in the histogram axes and throughout the code is E (linear), though using bins with logarithmic shape is the default. For zenith angle the variable is actually $\cos\theta$. Though the notation θ is used in this thesis, bins of constant width in $\cos\theta$ are used by default in the code ; the option to use θ bins is provided. Three different levels of binnings are defined: "reconstructed", "true", and "oscillograms". Their definition is held by an instance of the class `EvBins`, which is shared by all the analysis objects which need to be aware of binning: `Experiment`, `EvStore`, `SmearMachine`, `IntRate`, `Flux`, `OrcaSystematics`. Bins of reconstructed neutrino energy E', zenith angle θ', and inelasticity y' are common to the axes for these variables in the response matrix (multi-dimensional histogram) or in the smearing machinery (multiple efficiency and smearing histograms) and to the event distribution histograms which enter in the χ^2 calculation. They are designated in the code with the suffix `reco`. Bins for the true axes E, θ, and y in the response histograms have the suffix `true`. Last, $(E_{\text{osc}},\theta_{\text{osc}})$ bins shared by oscillograms, flux and interaction rate histograms have the suffix `osc`. They are identical to the `true` bins if the oversampling factor is set to $N_{\text{ov}} = 1$, otherwise they correspond to a subdivision of the `true` bins of a factor N_{ov} on each axis (see Sec. 5.3.2). The y axis in the interaction rate histograms is always identical to the true y axis of the response histograms.

User input: Monte Carlo event samples and PID classes definition

When defining the response matrix to be used by the `Experiment` objects, the user specifies a label identifying the sample of Monte Carlo events to be used. Such a sample consists in a pair of trees (ROOT `TTree` objects):

- the "generated" sample contains all generated events with the true values ("MC truth") of the analysis variable (interaction channel ν_x, true variables E, θ, y and coordinates of the true interaction vertex)
- the "selected" sample contains detected MC events with their MC truth as well as the reconstruction and classification outputs relevant to the analysis and selection cuts

In addition, the user provides definition of event classes through the `EvClasses` interface. A class is defined by a name, a set of selection cuts, and reconstruction identifiers. The selection cuts can involve any variable appearing in the `TTree` of selected events. Typical quality cuts involve reconstruction quality parameters, position of reconstructed vertices, *etc*; classification cuts normally involve the scores provided by the background rejection and track/cascade classifiers. Cuts are passed as strings using the `TCut`-like syntax, well-known to ROOT users for its use in passing selection cuts to `TTree::Draw()` commands. In addition, to each class C_i is associated a set of labels that identify the particular reconstruction output to be used for events classified as C_i when building the response matrix.

An arbitrary number of PID classes can be defined. In case overlapping definitions are provided by the user, events are selected following a predefined order of priority (by default highest to lowest in the order of class definition), making use of the logical functionalities provided by the ROOT `TCut` and `TEntryList` classes. This way events cannot be double counted.

Sterile and NSI specificities

Classes `Experiment_Sterile` and `Experiment_NSI` derive from `Experiment` and implement the few necessary adaptations. Mostly, they own different types of oscillation calculator, making use of the polymorphism of the `PMNS` classes in Oscprob. The standard `Experiment` class is declared to hold a pointer to a `OscProb::PMNS_Base` which is instantiated at initialisation as a `OscProb::PMNS_Fast` object, a fast calculator which uses optimisations only valid for the three-neutrino standard case. For the `Experiment_Sterile` and `Experiment_NSI` classes it is instantiated as a `PMNS_Sterile` and `PMNS_NSI` respectively. There is no specific definition of the `FitPars` object for these types of analyses; the usual object encompasses all parameters, and the ones relating to sterile neutrino masses and mixing or to the NSI effective mixing matrix are simply ignored in the `Experiment` class and only passed to the oscillation calculator in the derived classes. Finally, the `Experiment_Sterile` class reimplements the calculation of the NC event rate, since oscillation probabilities do not necessarily sum to unity when sterile flavour states are introduced:

$$P(\nu_\alpha \to \nu_e) + P(\nu_\alpha \to \nu_\mu) + P(\nu_\alpha \to \nu_\tau) \neq 1 \tag{C.1}$$

C.1.2 Advanced functionalities

Advanced statistical functionalities which were presented and used in the main text are available in the maintained versions of the software: `EvStore` resampling and bootstrapping, use of a user-defined resampling split number in `Experiment` (including with `SmearMachine` response), use of external MC error evaluation (*e.g.* from bootstrapping) in `Experiment`. All of these can be enabled from the `Experiment` or `EvStore` interfaces.

C.1.3 Optimisations

In this section the optimisations that were left out of the main text are mentioned, regarding computational speed and stability, memory usage, or the flexibility and ease of use of the interface.

Response matrix caching

With $\mathcal{O}(10^8)$ generated and $\mathcal{O}(10^6)$ selected events, the process of building a response matrix can take a significant time, the longest part being the step of filling generated events (up to ten minutes). In addition, it requires to access the MC event samples which are large files ($\mathcal{O}(10\,\text{GB})$), even when skimmed down to the minimal necessary amount of information per event. Even though ROOT is optimised for such situations and handles reading data on-disk rather well, the disk access requirement can be problematic when running large amounts of parallel jobs on a computing cluster. In contrast, the response matrix (class `THnSparse`) retains only the binned information necessary to the application of the response, the number of bins is kept moderate to mitigate the impact of MC sparseness, and most importantly the empty bins are not stored in memory: the typical file size is then much smaller with $\mathcal{O}(10\,\text{Mb})$ to $\mathcal{O}(100\,\text{Mb})$ depending on the binning densities and number of classes.

Therefore, a simple caching system has been implemented. The whole set of characteristics of a response matrix is listed in a string which is then transformed into an integer hash code. When asked to build a response matrix with certain settings, the program first checks if the response

matrix can be loaded from the cache directory, where previously built response matrices are saved using the identifier hash code as filename. When a filename match is found an additional check for the identity of full string identifiers is done in case that a hash collision were to happen. The response matrix characteristics listed in the identifier string include the software version used to build it, the full list of bin edges for each dimension, the MC sample tag, the full definition of PID classes and background classes, as well as the resampling and bootstrapping options and if applicable, the RNG seed. In addition to the full response matrix, the generated events distributions (intermediate step) are also saved in a different file. Indeed, filling these generated distributions is the most-time consuming step, and most often when building a sequence of response matrix only the PID and background classes definitions are changed, allowing to reuse the generated part.

Avoiding unnecessary recomputation steps in a fit

The most time-consuming steps in the computation chain are the calculation of oscillation probabilities and the application of the detector response – this step being faster when a correlated response matrix is used, and considerably slower with the current implementation of the uncorrelated smearing. It is thus desirable to avoid performing these steps unnecessarily in the context of a fit. Such an optimisation is especially relevant considering that most often the minimisation algorithm will only modify the value of one single parameter between one trial and the next. The parameters in `FitPars` are therefore separated into three categories:

1. Parameters which affect the oscillation probabilities (neutrino parameters, Earth model)
2. Parameters leaving oscillations unaffected but acting with a dependence on the unoscillated flavour (flux systematics) or on the true variables E, θ or y (shape and migration systematics)
3. Parameters acting as overall normalisations depending on the oscillated flavour, interaction channel, or event classification

Changing the value of a parameter in the first category requires to repeat the whole chain of computation, while for the second and third categories the oscillograms can be kept, and for the third one the interaction rates can be kept as well, leaving only the very fast step of channels recombination to be performed. In the context of a fit, when the model `Experiment` object is called to compute a χ^2 value for a set of trial parameter values it internally compares the new trial values to the ones used in the previous computation and determines which steps must be performed again.

C.2 Beyond the code

C.2.1 Motivation and history

As soon as I started working on the sensitivity studies to both the Mass Hierarchy determination and Earth tomography with ORCA, it became evident that the two analyses should be undertaken in a similar framework. The decision was made to design the software tools with the long-term goal of making them available to the whole collaboration, eventually becoming an official analysis software. Therefore, major design concerns where robustness, modularity and usability of the interface. Since OscProb had built-in capabilities for computing oscillation probabilities including an arbitrary number of neutrino states (3 active + N steriles) and modified versions of the 3-neutrino mixing matrix (NSI), care was taken from the early stages to design the experiment simulation classes to allow for such studies. A few of the earliest elements of the code were taken from analysis tools by Salvatore Galatà – evolved from the "toy Monte Carlo" approach used in [155]. An important inspiration in terms of methodology came from the analysis developped by Martijn Jongen [295] [95] which results were published in the KM3NeT Letter of Intent [21]. The initial object-oriented approach was also inspired by software used for LBNO

sensitivity studies [296], that I had adapted to use with ORCA Monte Carlo in the context of my master thesis [297]. The earliest developments of the analysis method and code benefited considerably from the guidance of Salvatore Galatà, Martijn Jongen and Jannik Hofestädt and more generally from discussion within the ORCA working group. Later on, Swim really took off as a collaborative software project thanks to the kick-starting contribution and continued supervision of Joao Coelho.

C.2.2 Contributions, current and future use

I have been the main developper of the software throughout the three years of my thesis in APC. The main other contributors are Joao Coelho (developped OscProb, contributed statistical utilities and his general purpose fitter package `NuFitter`) and Christine Nielsen (initiated the implementation of systematics in the separate `OrcaSyst` module, implemented polynomial shape systematics). Recently, Nhan Chau has joined the ORCA team in APC to work on flux systematics, and has started contributing to the systematics module. The framework is also being used in research teams of the KM3NeT collaboration outside of APC. Sterile neutrino studies are currently taken over by Alba Domi (INFN Genova, CPPM Marseille) with contributions to the analysis examples; the non-standard interaction study is pursued by Nafis Rezwan Khan Chowdury (IFIC Valencia). In addition, the framework is used for estimating ORCA's sensitivity to tau neutrino appearance by Lukas Maderer (ECAP Erlangen). A detector response model based on an event-by-event reweighting method, developped by Tarak Thakore (IFIC Valencia), is currently being ported to Swim and will allow to carry out refined systematics studies.

C.2.3 Documentation, tutorials and support

Extensive code documentation is auto-generated based on the source code itself, completed with suitably formatted comments, using the Doxygen tool [298]. This was set from the beginning of the collaborative development at APC, although the Doxygen-specific commenting was initially only marginal. In addition I wrote a wiki page providing an overall description of the software and its installation procedure.

I undertook a more serious documentation effort in the last year of my thesis work. The primary objective was to facilitate the widespread use of the framework for physics analyses (sensitivity studies), including its use by newcomers to the collaboration and students with little programming experience. I first made a short presentation of the software in a "software sprint" session at a meeting of the KM3NeT collaboration (Bari, June 2017). A few tutorial examples were written with extensive documentation and option parsing scripts; I presented them along with OscProb tutorials in the KM3NeT software bootcamp (Valencia, November 2017). I later wrote an extensive manual for the Swim software which sources and PDF version are now included in the software repository. The secondary objective of the documentation is to provide the necessary reference to complement the Doxygen documentation and in-code comments in helping future developers and maintainers of the software. With respect to this objective the Swim manual is still being improved, but already enters into a reasonable level of detail. As an indication its table of contents is reported at the end of this appendix.

The prerequisites for installing and running Swim are the GSL libraries and a ROOT installation (version 5 or 6) including the Minuit2 library. A significant amount of work has been put in ensuring compatibility with ROOT versions 5 and 6, including usability of most classes in interpreted mode with `cint` for ROOT 5 and `cling` for ROOT 6. Likewise, compatibility has been ensured and is regularly checked for most operating systems and corresponding default compilers relevant to the collaboration: Scientific Linux 6 (`gcc-4.4`) and CentOS 7 (`gcc-4.8`) in use at the CC-IN2P3 computing center, as well as Ubuntu Debian (LTS 14.04 with `gcc-4.8`, LTS 16.04 with `gcc-5.4`) and macOS (10.12 with Apple LLVM `clang-800`) for use on personal machines. Providing support to new users for installation troubleshooting and first steps with

the framework was also a demanding task. Therefore I added a dedicated troubleshooting section to the software manual, intended to be detailed enough for newcomers to C++ and ROOT with limited knowledge of Unix systems.

The version control system used for the earliest developments was SVN. The project was later migrated to git for its daily development, making use of a common Gitlab server for KM3NeT developpers. Scripts are provided to automatically push tagged releases to the SVN repository centralising all KM3NeT official software.

C.3 Contents of the Swim manual

The cover page and the table of contents of the Swim manual are reported in the next page.

KM3NeT/ORCA

Swim, how to

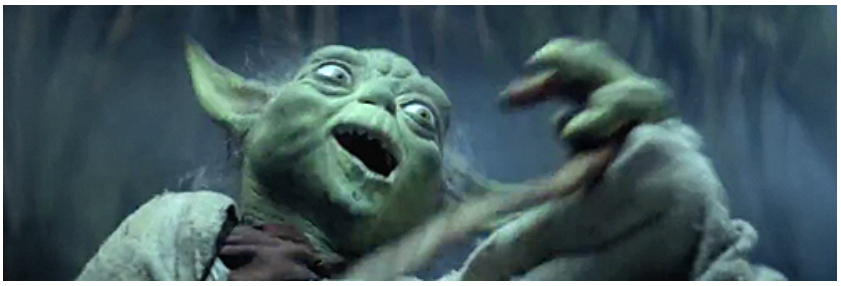

Manual for Swim v1.1

Authors
Simon Bourret `bourret@apc.in2p3.fr`
Christine Nielsen `cnielsen@apc.in2p3.fr`

June 7, 2018

Table of Contents

References

[1] Carlo Giunti and Kim Chung Wook. *Fundamentals of Neutrino Physics and Astrophysics*. Oxford: Oxford Univ., 2007 (cit. on pp. 6, 10, 30, 48, 51, 240, 241, 245).

[2] Andre Rubbia. *Lectures on Neutrino Physics at ETH Zürich*. 2014 (cit. on pp. 6, 24, 55).

[3] S. M. Bilenky. "Neutrino. History of a unique particle". In: *Eur. Phys. J.* H38, 2013, pp. 345–404. URL: https://arxiv.org/pdf/1210.3065.pdf (cit. on pp. 6, 9).

[4] Carlo Giunti. *Cours d'hiver 2017 du LAL (Lectures on Neutrino Physics at LAL, Orsay)*. http://personalpages.to.infn.it/~giunti/slides/2017/giunti-170103-LAL.pdf. 2017 (cit. on p. 6).

[5] C. D. Ellis and W. A. Wooster. "The average energy of disintegration of radium E". In: *Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences* 117.776, 1927, pp. 109–123. URL: http://rspa.royalsocietypublishing.org/content/117/776/109 (cit. on p. 6).

[6] Pertti O. Tikkanen. *Milestones in the history of beta decay*. Lectures on Nuclear Physics at the University of Helsinki. http://www.courses.physics.helsinki.fi/fys/nucphys/LN9-17.pdf. 2017 (cit. on p. 6).

[7] W. Pauli. "Discussion du rapport de M. Heisenberg". In: *Structure et propriétés des noyaux atomiques, rapports et discussions du septième Conseil de physique tenu à Bruxelles du 22 au 29 octobre 1933, sous les auspices de l'Institut international de physique Solvay*. 1934. URL: https://gallica.bnf.fr/ark:/12148/bpt6k5696894m (cit. on p. 7).

[8] Laurie M. Brown. "The idea of the neutrino". In: *Phys. Today* 31N9, 1978, pp. 23–28 (cit. on p. 8).

[9] Frederick Reines and Clyde L. Cowan. "The neutrino". In: *Nature* 178, 1956, pp. 446–449 (cit. on pp. 8, 12).

[10] C. L. Cowan et al. "Detection of the Free Neutrino: A Confirmation". In: *Science* 124.3212, 1956, pp. 103–104. URL: http://www.jstor.org/stable/1751492 (cit. on pp. 8, 12).

[11] N.G. Cooper et al. "The Reines-Cowan experiments". In: *Los Alamos Science* 25, 1997. URL: https://www.osti.gov/servlets/purl/569122-Zeu8O0/webviewable/ (cit. on p. 12).

[12] R. M. Zwaska. *Diagram of the NUMI beamline*. http://home.fnal.gov/~zwaska/beamline/index.html. 2005 (cit. on p. 13).

[13] R. M. Zwaska. "Accelerator Systems and Instrumentation for the NuMI Neutrino Beam". PhD thesis. Texas U., 2005. URL: http://lss.fnal.gov/cgi-bin/find_paper.pl?thesis-2005-73 (cit. on p. 13).

[14] G. Danby et al. "Observation of High-Energy Neutrino Reactions and the Existence of Two Kinds of Neutrinos". In: *Phys. Rev. Lett.* 9, 11962, pp. 36–44. URL: https://link.aps.org/doi/10.1103/PhysRevLett.9.36 (cit. on pp. 13, 235).

[15] F. J. Hasert et al. "Observation of Neutrino Like Interactions Without Muon Or Electron in the Gargamelle Neutrino Experiment". In: *Phys. Lett.* B46, 1973. [,5.15(1973)], pp. 138–140 (cit. on p. 13).

[16] K. Kodama et al. "Observation of tau neutrino interactions". In: *Phys. Lett.* B504, 2001, pp. 218–224 (cit. on p. 13).

[17] M. Cribier et al. *Solar neutrinos*. https://neutrino-history.in2p3.fr/solar-neutrinos/. 2018 (cit. on p. 16).

[18] T. Araki et al. “Experimental investigation of geologically produced antineutrinos with KamLAND”. In: *Nature* 436, 2005, pp. 499–503 (cit. on p. 17).

[19] R. Abbasi et al. “The IceCube Data Acquisition System: Signal Capture, Digitization, and Timestamping”. In: *Nucl. Instrum. Meth.* A601, 2009, pp. 294–316 (cit. on pp. 18, 56).

[20] M. Ageron et al. “ANTARES: the first undersea neutrino telescope”. In: *Nucl. Instrum. Meth.* A656, 2011, pp. 11–38 (cit. on pp. 18, 56, 86).

[21] KM3NeT Collaboration. “Letter of intent for KM3NeT 2.0”. In: *Journal of Physics G: Nuclear and Particle Physics* 43.8, 2016, p. 084001 (cit. on pp. 18, 34, 59, 68, 86, 87, 89, 94, 97, 98, 102, 103, 112, 138, 151, 194, 209, 270).

[22] A.D. Avrorin et al. “The prototyping/early construction phase of the BAIKAL-GVD project”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 742, 2014. 4th Roma International Conference on Astroparticle Physics, pp. 82 –88. URL: http://www.sciencedirect.com/science/article/pii/S0168900213014551 (cit. on pp. 18, 56).

[23] Francis Halzen. “Astroparticle physics with high energy neutrinos: from amanda to icecube”. In: *Eur. Phys. J.* C46, 2006, pp. 669–687 (cit. on p. 18).

[24] T. Chiarusi and M. Spurio. “High-Energy Astrophysics with Neutrino Telescopes”. In: *Eur. Phys. J.* C65, 2010, pp. 649–701 (cit. on p. 18).

[25] M. G. Aartsen et al. “Evidence for High-Energy Extraterrestrial Neutrinos at the IceCube Detector”. In: *Science* 342, 2013, p. 1242856 (cit. on pp. 18, 86).

[26] M. G. Aartsen et al. “Multimessenger observations of a flaring blazar coincident with high-energy neutrino IceCube-170922A”. In: *Science* 361.6398, 2018, eaat1378 (cit. on pp. 18, 220).

[27] M. G. Aartsen et al. “Neutrino emission from the direction of the blazar TXS 0506+056 prior to the IceCube-170922A alert”. In: *Science* 361.6398, 2018, pp. 147–151 (cit. on pp. 18, 220).

[28] B. Pontecorvo. “Neutrino Experiments and the Problem of Conservation of Leptonic Charge”. In: *Sov. Phys. JETP* 26, 1968. [Zh. Eksp. Teor. Fiz.53,1717(1967)], pp. 984–988 (cit. on p. 18).

[29] Raymond Davis Jr., Don S. Harmer, and Kenneth C. Hoffman. “Search for neutrinos from the sun”. In: *Phys. Rev. Lett.* 20, 1968, pp. 1205–1209 (cit. on p. 18).

[30] K. S. Hirata et al. “Observation of B-8 Solar Neutrinos in the Kamiokande-II Detector”. In: *Phys. Rev. Lett.* 63, 1989, p. 16 (cit. on p. 19).

[31] Y. Fukuda et al. “Measurements of the solar neutrino flux from Super-Kamiokande’s first 300 days”. In: *Phys. Rev. Lett.* 81, 1998. [Erratum: Phys. Rev. Lett.81,4279(1998)], pp. 1158–1162 (cit. on p. 19).

[32] Y. Fukuda et al. “Evidence for oscillation of atmospheric neutrinos”. In: *Phys. Rev. Lett.* 81, 1998, pp. 1562–1567 (cit. on p. 20).

[33] Y. Fukuda et al. “Study of the atmospheric neutrino flux in the multi-GeV energy range”. In: *Phys. Lett.* B436, 1998, pp. 33–41 (cit. on p. 21).

[34] N.G. Cooper et al. “The Oscillating Neutrino: an introduction to neutrino masses and mixings”. In: *Los Alamos Science* 25, 1997. URL: https://www.osti.gov/servlets/purl/569122-Zeu8O0/webviewable/ (cit. on p. 22).

[35] M. Apollonio et al. “Limits on neutrino oscillations from the CHOOZ experiment”. In: *Phys. Lett.* B466, 1999, pp. 415–430 (cit. on p. 25).

[36] Michele Maltoni and Alexei Yu. Smirnov. “Solar neutrinos and neutrino physics”. In: *Eur. Phys. J.* A52.4, 2016, p. 87 (cit. on pp. 25, 40, 78).

[37] Yasaman Farzan and Alexei Yu. Smirnov. “Coherence and oscillations of cosmic neutrinos”. In: *Nucl. Phys.* B805, 2008, pp. 356–376 (cit. on pp. 25, 78).

[38] T. Araki et al. “Measurement of neutrino oscillation with KamLAND: Evidence of spectral distortion”. In: *Phys. Rev. Lett.* 94, 2005, p. 081801. URL: https://journals.aps.org/prl/pdf/10.1103/PhysRevLett.63.16 (cit. on p. 26).

[39] M. Agostini et al. “First Simultaneous Precision Spectroscopy of *pp*, ^{7}Be, and *pep* Solar Neutrinos with Borexino Phase-II”. In: 2017 (cit. on p. 26).

[40] Frederick Reines. *quoting Pauli in his foreword to the book 'Spaceship Neutrino' (Christine Sutton, 1992)*. 1992 (cit. on p. 29).

[41] Steven Weinberg. *The Quantum theory of fields. Vol. 1: Foundations*. Cambridge University Press, 2005. ISBN: 9780521670531, 9780511252044 (cit. on pp. 30, 222, 231, 239–241).

[42] Stephen F King. "Discrete Symmetries and Models of Flavour Mixing". In: 631, 2015, p. 012005 (cit. on p. 32).

[43] Fanny Dufour and Dave Wark. "Experimental Status of Neutrino Physics". In: *J. Phys. Conf. Ser.* 408, 2013, p. 012001 (cit. on p. 32).

[44] C. Patrignani et al. "Review of Particle Physics". In: *Chin. Phys.* C40.10, 2016, p. 100001 (cit. on pp. 33, 45, 154).

[45] L. Wolfenstein. "Neutrino oscillations in matter". In: *Phys. Rev. D* 17.9, 1978, pp. 2369–2374. URL: http://link.aps.org/doi/10.1103/PhysRevD.17.2369 (cit. on p. 37).

[46] Tommy Ohlsson and Hakan Snellman. "Neutrino oscillations with three flavors in matter of varying density". In: *Eur. Phys. J.* C20, 2001, pp. 507–515 (cit. on pp. 38, 78, 79).

[47] S. P. Mikheev and A. Yu. Smirnov. "Resonance Amplification of Oscillations in Matter and Spectroscopy of Solar Neutrinos". In: *Sov. J. Nucl. Phys.* 42, 1985. [Yad. Fiz.42,1441(1985)], pp. 913–917 (cit. on p. 39).

[48] Evgeny K. Akhmedov et al. "T violation in neutrino oscillations in matter". In: *Nucl. Phys.* B608, 2001, pp. 394–422 (cit. on p. 39).

[49] A. Yu. Smirnov. "Solar neutrinos: Oscillations or No-oscillations?" In: 2016 (cit. on p. 40).

[50] M. H. Ahn et al. "Measurement of Neutrino Oscillation by the K2K Experiment". In: *Phys. Rev.* D74, 2006, p. 072003 (cit. on p. 40).

[51] D. G. Michael et al. "Observation of muon neutrino disappearance with the MINOS detectors and the NuMI neutrino beam". In: *Phys. Rev. Lett.* 97, 2006, p. 191801 (cit. on p. 40).

[52] P. Adamson et al. "Measurement of Neutrino and Antineutrino Oscillations Using Beam and Atmospheric Data in MINOS". In: *Phys. Rev. Lett.* 110.25, 2013, p. 251801 (cit. on p. 41).

[53] M. Apollonio et al. "Search for neutrino oscillations on a long baseline at the CHOOZ nuclear power station". In: *Eur. Phys. J.* C27, 2003, pp. 331–374 (cit. on p. 41).

[54] K. Abe et al. "The T2K Experiment". In: *Nucl. Instrum. Meth.* A659, 2011, pp. 106–135 (cit. on p. 42).

[55] G. L. Fogli et al. "Evidence of $\theta_{13} \neq 0$ from global neutrino data analysis". In: *Phys. Rev.* D84, 2011, p. 053007 (cit. on p. 42).

[56] K. Abe et al. "Indication of Electron Neutrino Appearance from an Accelerator-produced Off-axis Muon Neutrino Beam". In: *Phys. Rev. Lett.* 107, 2011 (cit. on p. 42).

[57] P. Adamson et al. "Improved search for muon-neutrino to electron-neutrino oscillations in MINOS". In: *Phys. Rev. Lett.* 107, 2011, p. 181802 (cit. on p. 42).

[58] Y. Abe et al. "Indication of Reactor $\bar{\nu}_e$ Disappearance in the Double Chooz Experiment". In: *Phys. Rev. Lett.* 108, 2012, p. 131801 (cit. on p. 42).

[59] F. P. An et al. "Observation of electron-antineutrino disappearance at Daya Bay". In: *Phys. Rev. Lett.* 108, 2012, p. 171803 (cit. on p. 42).

[60] J. K. Ahn et al. "Observation of Reactor Electron Antineutrino Disappearance in the RENO Experiment". In: *Phys. Rev. Lett.* 108, 2012, p. 191802 (cit. on p. 42).

[61] K. Abe et al. "Evidence of Electron Neutrino Appearance in a Muon Neutrino Beam". In: *Phys. Rev.* D88.3, 2013, p. 032002 (cit. on p. 42).

[62] Ivan Esteban et al. "Updated fit to three neutrino mixing: exploring the accelerator-reactor complementarity". In: *JHEP* 01, 2017, p. 087 (cit. on pp. 42–44, 183).

[63] Esteban, Ivan and Gonzalez-Garcia, M. C. and Maltoni, Michele and Martinez-Soler, Ivan and Schwetz, Thomas. *NuFiT 3.2*. http://www.nu-fit.org. See also arXiv:1611.01514v2 [hep-ph]. 2018 (cit. on pp. 42–44, 67, 70, 71, 116, 167, 172, 180, 182, 183, 186).

[64] N. Agafonova et al. “Discovery of τ Neutrino Appearance in the CNGS Neutrino Beam with the OPERA Experiment”. In: *Phys. Rev. Lett.* 115.12, 2015, p. 121802 (cit. on p. 42).

[65] P. F. de Salas et al. “Status of neutrino oscillations 2018: 3σ hint for normal mass ordering and improved CP sensitivity”. In: *Phys. Lett.* B782, 2018, pp. 633–640 (cit. on p. 43).

[66] F. Capozzi et al. “Status and prospects of global analyses of neutrino mass-mixing parameters”. In: *Journal of Physics: Conference Series* 888.1, 2017, p. 012037. URL: http://stacks.iop.org/1742-6596/888/i=1/a=012037 (cit. on p. 43).

[67] M. C. Gonzalez-Garcia, Michele Maltoni, and Thomas Schwetz. “Global Analyses of Neutrino Oscillation Experiments”. In: *Nucl. Phys.* B908, 2016, pp. 199–217 (cit. on p. 44).

[68] R. N. Mohapatra and A. Y. Smirnov. “Neutrino Mass and New Physics”. In: *Ann. Rev. Nucl. Part. Sci.* 56, 2006, pp. 569–628 (cit. on p. 43).

[69] Walter Winter. “Neutrino mass hierarchy determination with IceCube-PINGU”. In: *Phys. Rev.* D88.1, 2013, p. 013013 (cit. on pp. 43, 66, 182).

[70] A. de Gouvea et al. *Neutrinos : Report of the Community Summer Study 2013 (Snowmass) Intensity Frontier Neutrino Working Group*. arXiv:1310.4340v1. 2013 (cit. on p. 43).

[71] C. Hagedorn et al. “CP Violation in the Lepton Sector and Implications for Leptogenesis”. In: *Int. J. Mod. Phys.* A33.05n06, 2018, p. 1842006 (cit. on p. 43).

[72] Susanne Mertens. “Direct Neutrino Mass Experiments”. In: *J. Phys. Conf. Ser.* 718.2, 2016, p. 022013 (cit. on p. 43).

[73] Massimiliano Lattanzi and Martina Gerbino. “Status of neutrino properties and future prospects - Cosmological and astrophysical constraints”. In: *Front.in Phys.* 5, 2018, p. 70 (cit. on p. 43).

[74] P. A. R. Ade et al. “Planck 2015 results. XIII. Cosmological parameters”. In: *Astron. Astrophys.* 594, 2016, A13 (cit. on p. 44).

[75] Julien Lesgourgues and Sergio Pastor. “Neutrino cosmology and Planck”. In: *New J. Phys.* 16, 2014, p. 065002 (cit. on p. 44).

[76] Walter Winter. “Neutrino mass hierarchy: Theory and phenomenology”. In: *AIP Conf. Proc.* 1666, 2015, p. 120001 (cit. on p. 45).

[77] Igor Ostrovskiy and Kevin O’Sullivan. “Search for neutrinoless double beta decay”. In: *Mod. Phys. Lett.* A31.18, 2016. [Erratum: Mod. Phys. Lett.A31,no.23,1692004(2016)], p. 1630017 (cit. on p. 45).

[78] The CUORE collaboration. *CUORE - A search for neutrinoless double beta decay*. https://cuore.lngs.infn.it/en/about/physics. 2018 (cit. on p. 45).

[79] M. A. Acero et al. “New constraints on oscillation parameters from ν_e appearance and ν_μ disappearance in the NOvA experiment”. In: *Phys. Rev.* D98, 2018, p. 032012 (cit. on pp. 45, 210).

[80] Fengpeng An et al. “Neutrino Physics with JUNO”. In: *J. Phys.* G43.3, 2016, p. 030401 (cit. on p. 45).

[81] T. K. Gaisser and M. Honda. “Flux of atmospheric neutrinos”. In: *Ann. Rev. Nucl. Part. Sci.* 52, 2002, pp. 153–199 (cit. on p. 48).

[82] G. D. Barr et al. “A Three - dimensional calculation of atmospheric neutrinos”. In: *Phys. Rev.* D70, 2004, p. 023006 (cit. on pp. 48, 51).

[83] Paolo Lipari. “The Geometry of atmospheric neutrino production”. In: *Astropart. Phys.* 14, 2000, pp. 153–170 (cit. on pp. 48, 51).

[84] M. Honda et al. “Atmospheric neutrino flux calculation using the NRLMSISE-00 atmospheric model”. In: *Phys. Rev.* D92.2, 2015, p. 023004 (cit. on pp. 48, 51, 111, 112).

[85] M. Honda et al. *Atmospheric Neutrino Flux Tables for One-Year-Average (HAKKM, 2014)*. http://www.icrr.u-tokyo.ac.jp/~mhonda/nflx2014/index.html (cit. on pp. 48, 51, 111, 112, 163).

[86] G. D. Barr et al. “Uncertainties in Atmospheric Neutrino Fluxes”. In: *Phys. Rev.* D74, 2006, p. 094009 (cit. on pp. 51, 163, 164).

[87] GT Zatsepin and VA Kuz’min. “Neutrino production in the atmosphere”. In: *Sov. Phys. JETP* 14, 1962, p. 1294. URL: http://jetp.ac.ru/cgi-bin/dn/e_014_06_1294.pdf (cit. on p. 51).

[88] M. A. Markov and I. M. Zheleznykh. “On high energy neutrino physics in cosmic rays”. In: *Nucl. Phys.* 27, 1961, pp. 385–394 (cit. on p. 51).

[89] Morihiro Honda et al. “A New calculation of the atmospheric neutrino flux in a 3-dimensional scheme”. In: *Phys. Rev.* D70, 2004, p. 043008 (cit. on p. 51).

[90] A. Fedynitch et al. “A state-of-the-art calculation of atmospheric lepton fluxes”. In: *35th International Cosmic Ray Conference, July 10-20, 2017, Bexco, Busan, Korea*. PoS(ICRC2017)1019. 2017. URL: https://pos.sissa.it/301/1019/pdf (cit. on pp. 51, 212).

[91] Anatoli Fedynitch et al. “Calculation of conventional and prompt lepton fluxes at very high energy”. In: *EPJ Web Conf.* 99, 2015, p. 08001 (cit. on p. 51).

[92] Anatoli Fedynitch. *MCeq - Matrix cascade equation*. https://github.com/afedynitch/MCEq (cit. on pp. 51, 212).

[93] J. A. Formaggio and G. P. Zeller. “From eV to EeV: Neutrino Cross Sections Across Energy Scales”. In: *Rev. Mod. Phys.* 84, 2012, pp. 1307–1341 (cit. on pp. 51, 54).

[94] Jannik Hofestädt. “Measuring the neutrino mass hierarchy with the future KM3NeT/ORCA detector”. PhD thesis. Friedrich-Alexander-Universität Erlangen-Nürnberg, 2017 (cit. on pp. 51, 52, 91, 101).

[95] M.H.G. Jongen. “Neutrino physics in the Mediterranean Sea: Determining the neutrino mass hierarchy with KM3NeT/ORCA”. PhD thesis. University of Amsterdam, 2018 (cit. on pp. 51, 270).

[96] Janet M. Conrad, Michael H. Shaevitz, and Tim Bolton. “Precision measurements with high-energy neutrino beams”. In: *Rev. Mod. Phys.* 70, 1998, pp. 1341–1392 (cit. on p. 51).

[97] C. Andreopoulos et al. “The GENIE Neutrino Monte Carlo Generator”. In: *Nucl. Instrum. Meth.* A614, 2010, pp. 87–104 (cit. on pp. 53, 95, 113).

[98] Costas Andreopoulos et al. “The GENIE Neutrino Monte Carlo Generator: Physics and User Manual”. In: 2015 (cit. on pp. 53, 95).

[99] D. Casper. “The Nuance neutrino physics simulation, and the future”. In: *Nucl. Phys. Proc. Suppl.* 112, 2002. [,161(2002)], pp. 161–170 (cit. on p. 54).

[100] Mathieu Ribordy and Alexei Yu Smirnov. “Improving the neutrino mass hierarchy identification with inelasticity measurement in PINGU and ORCA”. In: *Phys. Rev.* D87.11, 2013, p. 113007 (cit. on pp. 55, 66, 69).

[101] J G H De Groot et al. “Inclusive interactions of high-energy neutrinos and antineutrinos in iron”. In: 1, 1979 (cit. on p. 55).

[102] J. Devan et al. “Measurements of the Inclusive Neutrino and Antineutrino Charged Current Cross Sections in MINERvA Using the Low-ν Flux Method”. In: *Phys. Rev.* D94.11, 2016, p. 112007 (cit. on p. 55).

[103] L. Ren et al. “Measurement of the antineutrino to neutrino charged-current interaction cross section ratio in MINERvA”. In: *Phys. Rev.* D95.7, 2017. [Addendum: Phys. Rev.D97,no.1,019902(2018)], p. 072009 (cit. on p. 55).

[104] A. A. Aguilar-Arevalo et al. “First Measurement of the Muon Neutrino Charged Current Quasielastic Double Differential Cross Section”. In: *Phys. Rev.* D81, 2010, p. 092005 (cit. on p. 55).

[105] R. Acciarri et al. “Design and Construction of the MicroBooNE Detector”. In: *JINST* 12.02, 2017, P02017 (cit. on p. 55).

[106] K. Abe et al. “Measurement of inclusive double-differential ν_μ charged-current cross section with improved acceptance in the T2K off-axis near detector”. In: *Phys. Rev.* D98, 2018, p. 012004 (cit. on p. 55).

[107] erenkov, familyi=v., given=P. A., giveni=P. A., „ “Visible Radiation Produced by Electrons Moving in a Medium with Velocities Exceeding that of Light”. In: *Phys. Rev.* 52, 41937, pp. 378–379. URL: https://link.aps.org/doi/10.1103/PhysRev.52.378 (cit. on p. 56).

[108] W. GALBRAITH and J. V. JELLEY. “Light Pulses from the Night Sky associated with Cosmic Rays”. In: *Nature* 171, 1953, 349 EP –. URL: http://dx.doi.org/10.1038/171349a0 (cit. on p. 56).

[109] M. A. Markov. "On high energy neutrino physics". In: *Proceedings, 10th International Conference on High-Energy Physics (ICHEP 60): Rochester, NY, USA, 25 Aug - 1 Sep 1960*. 1960, pp. 578–581 (cit. on p. 56).

[110] M.A. Markov and I.M. Zheleznykh. "On high energy neutrino physics in cosmic rays". In: *Nuclear Physics* 27.3, 1961, pp. 385 –394. URL: http://www.sciencedirect.com/science/article/pii/0029558261903315 (cit. on p. 56).

[111] K Greisen. "Cosmic Ray Showers". In: *Annual Review of Nuclear Science* 10.1, 1960, pp. 63–108. URL: https://doi.org/10.1146/annurev.ns.10.120160.000431 (cit. on pp. 56, 76).

[112] Aleksandr I. Volokitin and Bo N.J. Persson. "Radiation by Uniformly Moving Sources". In: *Electromagnetic Fluctuations at the Nanoscale: Theory and Applications*. Berlin, Heidelberg: Springer Berlin Heidelberg, 2017, pp. 227–242. ISBN: 978-3-662-53474-8. URL: https://doi.org/10.1007/978-3-662-53474-8_11 (cit. on p. 57).

[113] A. Trovato. "Development of reconstruction algorithms for large volume neutrino telescopes and their application to the KM3NeT detector". PhD thesis. Università degli Studi di Catania, Scuola Superiore di Catania, 2010 (cit. on p. 58).

[114] S. Adrián-Martínez et al. "Intrinsic limits on resolutions in muon- and electron-neutrino charged-current events in the KM3NeT/ORCA detector". In: *JHEP* 05, 2017, p. 008 (cit. on pp. 60, 102).

[115] Martin Freund. "Analytic approximations for three neutrino oscillation parameters and probabilities in matter". In: *Phys. Rev.* D64, 2001, p. 053003. URL: https://arxiv.org/pdf/hep-ph/0103300.pdf (cit. on p. 62).

[116] Adam M. Dziewonski and Don L. Anderson. "Preliminary reference Earth model". In: *Physics of the Earth and Planetary Interiors* 25.4, 1981, pp. 297–356 (cit. on pp. 63, 74, 114, 116).

[117] V. K. Ermilova, V. A. Tsarev, and V. A. Chechin. "Buildup of Neutrino Oscillations in the Earth". In: *JETP Lett.* 43, 1986. [Pisma Zh. Eksp. Teor. Fiz.43,353(1986)], pp. 453–456 (cit. on p. 65).

[118] Evgeny K. Akhmedov. "Neutrino oscillations in inhomogeneous matter. (In Russian)". In: *Sov. J. Nucl. Phys.* 47, 1988. [Yad. Fiz.47,475(1988)], pp. 301–302 (cit. on p. 65).

[119] P. I. Krastev and A. Yu. Smirnov. "Parametric Effects in Neutrino Oscillations". In: *Phys. Lett.* B226, 1989, pp. 341–346 (cit. on p. 65).

[120] Evgeny K. Akhmedov. "Parametric resonance of neutrino oscillations and passage of solar and atmospheric neutrinos through the earth". In: *Nucl. Phys.* B538, 1999, pp. 25–51 (cit. on p. 65).

[121] Evgeny K. Akhmedov et al. "Atmospheric neutrinos at Super-Kamiokande and parametric resonance in neutrino oscillations". In: *Nucl. Phys.* B542, 1999, pp. 3–30 (cit. on pp. 65, 66).

[122] M. Chizhov, M. Maris, and S. T. Petcov. "On the oscillation length resonance in the transitions of solar and atmospheric neutrinos crossing the earth core". In: 1998 (cit. on pp. 65, 66).

[123] S. T. Petcov. "Diffractive - like (or parametric resonance - like?) enhancement of the earth (day - night) effect for solar neutrinos crossing the earth core". In: *Phys. Lett.* B434, 1998, pp. 321–332 (cit. on p. 65).

[124] M. V. Chizhov and S. T. Petcov. "Enhancing mechanisms of neutrino transitions in a medium of nonperiodic constant density layers and in the Earth". In: *Phys. Rev.* D63, 2001, p. 073003 (cit. on p. 65).

[125] Evgeny K. Akhmedov. "Parametric resonance in neutrino oscillations in matter". In: *Pramana* 54, 2000, pp. 47–63 (cit. on p. 65).

[126] Masafumi Koike et al. "Fourier Analysis of the Parametric Resonance in Neutrino Oscillations". In: *Phys. Lett.* B675, 2009, pp. 69–72 (cit. on p. 65).

[127] Masafumi Koike et al. "Parametric Resonance in Neutrino Oscillation: A Guide to Control the Effects of Inhomogeneous Matter Density". In: *Phys. Lett.* B759, 2016, pp. 266–271 (cit. on pp. 65, 79).

[128] João A. B. Coelho. *OscProb neutrino oscillation calculator*. https://github.com/joaoabcoelho/OscProb (cit. on pp. 66, 114, 115, 266).

[129] R. Abbasi et al. “The design and performance of IceCube DeepCore”. In: *Astroparticle Physics* 35.10, 2012, pp. 615 –624. URL: http://www.sciencedirect.com/science/article/pii/S0927650512000254 (cit. on p. 65).

[130] S. Adrian-Martinez et al. “Measurement of Atmospheric Neutrino Oscillations with the ANTARES Neutrino Telescope”. In: *Phys. Lett.* B714, 2012, pp. 224–230 (cit. on p. 66).

[131] M. G. Aartsen et al. “Measurement of Atmospheric Neutrino Oscillations with IceCube”. In: *Phys. Rev. Lett.* 111.8, 2013, p. 081801 (cit. on p. 66).

[132] M. G. Aartsen et al. “Measurement of Atmospheric Neutrino Oscillations at 6–56 GeV with IceCube DeepCore”. In: *Phys. Rev. Lett.* 120.7, 2018, p. 071801 (cit. on p. 66).

[133] K. Abe et al. “Atmospheric neutrino oscillation analysis with external constraints in Super-Kamiokande I-IV”. In: *Phys. Rev.* D97.7, 2018, p. 072001 (cit. on p. 67).

[134] K. Abe et al. “Combined Analysis of Neutrino and Antineutrino Oscillations at T2K”. In: *Phys. Rev. Lett.* 118.15, 2017, p. 151801 (cit. on p. 67).

[135] P. Adamson et al. “Measurement of the neutrino mixing angle θ_{23} in NOvA”. In: *Phys. Rev. Lett.* 118.15, 2017, p. 151802 (cit. on p. 67).

[136] P. Adamson et al. “Combined analysis of ν_μ disappearance and $\nu_\mu \to \nu_e$ appearance in MINOS using accelerator and atmospheric neutrinos”. In: *Phys. Rev. Lett.* 112, 2014, p. 191801 (cit. on p. 67).

[137] M. G. Aartsen et al. “Neutrino oscillation studies with IceCube-DeepCore”. In: *Nucl. Phys.* B908, 2016, pp. 161–177 (cit. on p. 67).

[138] J. Bernabeu et al. “The Earth mantle core effect in matter induced asymmetries for atmospheric neutrino oscillations”. In: *Phys. Lett.* B531, 2002, pp. 90–98 (cit. on p. 66).

[139] Sergio Palomares-Ruiz and S. T. Petcov. “Three-neutrino oscillations of atmospheric neutrinos, theta(13), neutrino mass hierarchy and iron magnetized detectors”. In: *Nucl. Phys.* B712, 2005, pp. 392–410 (cit. on p. 66).

[140] D. Indumathi and M. V. N. Murthy. “A Question of hierarchy: Matter effects with atmospheric neutrinos and anti-neutrinos”. In: *Phys. Rev.* D71, 2005, p. 013001 (cit. on p. 66).

[141] Raj Gandhi et al. “Earth matter effects at very long baselines and the neutrino mass hierarchy”. In: *Phys. Rev.* D73, 2006, p. 053001 (cit. on p. 66).

[142] S. T. Petcov and T. Schwetz. “Determining the neutrino mass hierarchy with atmospheric neutrinos”. In: *Nucl. Phys.* B740, 2006, pp. 1–22 (cit. on p. 66).

[143] Raj Gandhi et al. “Mass Hierarchy Determination via future Atmospheric Neutrino Detectors”. In: *Phys. Rev.* D76, 2007, p. 073012 (cit. on p. 66).

[144] Abhijit Samanta. “Discrimination of mass hierarchy with atmospheric neutrinos at a magnetized muon detector”. In: *Phys. Rev.* D81, 2010, p. 037302 (cit. on p. 66).

[145] Abhijit Samanta and A. Yu. Smirnov. “The 2-3 mixing and mass split: atmospheric neutrinos and magnetized spectrometers”. In: *JHEP* 07, 2011, p. 048 (cit. on p. 66).

[146] Mattias Blennow and Thomas Schwetz. “Identifying the Neutrino mass Ordering with INO and NOvA”. In: *JHEP* 08, 2012. [Erratum: JHEP11,098(2012)], p. 058. URL: https://arxiv.org/pdf/1203.3388.pdf (cit. on p. 66).

[147] Anushree Ghosh, Tarak Thakore, and Sandhya Choubey. “Determining the Neutrino Mass Hierarchy with INO, T2K, NOvA and Reactor Experiments”. In: *JHEP* 04, 2013, p. 009 (cit. on p. 66).

[148] Anushree Ghosh and Sandhya Choubey. “Measuring the Mass Hierarchy with Muon and Hadron Events in Atmospheric Neutrino Experiments”. In: *JHEP* 10, 2013, p. 174 (cit. on p. 66).

[149] J. Bernabeu, Sergio Palomares Ruiz, and S. T. Petcov. “Atmospheric neutrino oscillations, theta(13) and neutrino mass hierarchy”. In: *Nucl. Phys.* B669, 2003, pp. 255–276 (cit. on p. 66).

[150] S. T. Petcov and Sergio Palomares-Ruiz. “On the atmospheric neutrino oscillations, theta(13) and neutrino mass hierarchy”. In: *Neutrino oscillations and their origin. Proceedings, 5th International Workshop, NOON2004, Tokyo, Japan, February 11-15, 2004*. 2004, pp. 220–227 (cit. on p. 66).

[151] Patrick Huber, Michele Maltoni, and Thomas Schwetz. “Resolving parameter degeneracies in long-baseline experiments by atmospheric neutrino data”. In: *Phys. Rev.* D71, 2005, p. 053006 (cit. on p. 66).

[152] Olga Mena, Irina Mocioiu, and Soebur Razzaque. “Neutrino mass hierarchy extraction using atmospheric neutrinos in ice”. In: *Phys. Rev.* D78, 2008, p. 093003 (cit. on p. 66).

[153] E. Kh. Akhmedov, Soebur Razzaque, and A. Yu. Smirnov. “Mass hierarchy, 2-3 mixing and CP-phase with Huge Atmospheric Neutrino Detectors”. In: *JHEP* 02, 2013. [Erratum: JHEP07,026(2013)], p. 082 (cit. on p. 66).

[154] Sanjib Kumar Agarwalla et al. “Exploring the Earth matter effect with atmospheric neutrinos in ice”. In: 2012 (cit. on pp. 66, 80).

[155] D. Franco et al. “Mass hierarchy discrimination with atmospheric neutrinos in large volume ice/water Cherenkov detectors”. In: *JHEP* 04, 2013, p. 008 (cit. on pp. 66, 151, 182, 187, 270).

[156] Mattias Blennow and Thomas Schwetz. “Determination of the neutrino mass ordering by combining PINGU and Daya Bay II”. In: *JHEP* 09, 2013, p. 089 (cit. on pp. 66, 212).

[157] Shao-Feng Ge, Kaoru Hagiwara, and Carsten Rott. “A Novel Approach to Study Atmospheric Neutrino Oscillation”. In: *JHEP* 06, 2014, p. 150 (cit. on p. 66).

[158] Sandhya Choubey and Anushree Ghosh. “Determining the Octant of θ_{23} with PINGU, T2K, NOvA and Reactor Data”. In: *JHEP* 11, 2013, p. 166 (cit. on p. 66).

[159] Mattias Blennow et al. “Quantifying the sensitivity of oscillation experiments to the neutrino mass ordering”. In: *JHEP* 03, 2014, p. 028 (cit. on pp. 66, 148, 151, 153, 195, 198).

[160] F. Capozzi, E. Lisi, and A. Marrone. “PINGU and the neutrino mass hierarchy: Statistical and systematic aspects”. In: *Phys. Rev.* D91, 2015, p. 073011 (cit. on pp. 66, 163).

[161] Francesco Capozzi, Eligio Lisi, and Antonio Marrone. “Probing the neutrino mass ordering with KM3NeT-ORCA: Analysis and perspectives”. In: 2017 (cit. on pp. 66, 137, 163, 179).

[162] J. P. Yañez and A. Kouchner. “Measurement of atmospheric neutrino oscillations with very large volume neutrino telescopes”. In: *Adv. High Energy Phys.* 2015, 2015, p. 271968 (cit. on pp. 66, 70).

[163] M. C. Gonzalez-Garcia, Michele Maltoni, and Thomas Schwetz. “Updated fit to three neutrino mixing: status of leptonic CP violation”. In: *Journal of High Energy Physics* 2014.11, 2014, pp. 1–28 (cit. on pp. 69, 70, 167).

[164] AM Dziewonski and F Gilbert. “Solidity of the inner core of the Earth inferred from normal mode observations”. In: *Nature* 234.5330, 1971, p. 465 (cit. on p. 73).

[165] Stephen A. Nelson. *Earthquakes and the Earth interior*. Lectures on Physical Geology at Tulane University. http://www.tulane.edu/~sanelson/eens1110/earthint.htm. 2015 (cit. on p. 74).

[166] B. L. N. Kennett, E. R. Engdahl, and R. Buland. “Constraints on seismic velocities in the Earth from traveltimes”. In: *Geophysical Journal International* 122.1, 1995, pp. 108–124. URL: http://dx.doi.org/10.1111/j.1365-246X.1995.tb03540.x (cit. on p. 74).

[167] J.-P. Montagner and B. L. N. Kennett. “How to reconcile body-wave and normal-mode reference earth models”. In: *Geophysical Journal International* 125.1, pp. 229–248. URL: https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1365-246X.1996.tb06548.x (cit. on p. 74).

[168] Cooperative Institute for Dynamic Earth Research. *Reference Models*. https://www.deep-earth.org/wiki_cider/REFERENCE_MODELS (cit. on p. 74).

[169] C. Rott, A. Taketa, and D. Bose. “Spectrometry of the Earth using Neutrino Oscillations”. In: *Scientific Reports* 5, 2015, p. 15225. URL: http://dx.doi.org/10.1038/srep15225 (cit. on pp. 74, 80, 82).

[170] B. A. Bolt. “The Precision of Density Estimation Deep in the Earth”. In: *Q. J. R. Astron. Soc.* 32, 1991, p. 367 (cit. on p. 74).

[171] B. L. N. Kennett. “On the density distribution within the Earth”. In: *Geophysical Journal International* 132.2, 1998, pp. 374–382. URL: http://dx.doi.org/10.1046/j.1365-246x.1998.00451.x (cit. on pp. 74, 75).

[172] Guy Masters and David Gubbins. “On the resolution of density within the Earth”. In: 140, 2003, pp. 159–167 (cit. on pp. 74, 75).

[173] Jessica C. E. Irving, Sanne Cottaar, and Vedran Lekić. “Seismically determined elastic parameters for Earth’s outer core”. In: *Science Advances* 4.6, 2018. URL: http://advances.sciencemag.org/content/4/6/eaar2538 (cit. on pp. 74, 75).

[174] R.W.L. de Wit et al. “Bayesian inversion of free oscillations for Earth’s radial (an)elastic structure”. In: *Physics of the Earth and Planetary Interiors* 237, 2014, pp. 1 –17. URL: http://www.sciencedirect.com/science/article/pii/S0031920114002039 (cit. on p. 75).

[175] Sanne Cottaar and Vedran Lekic. “Morphology of seismically slow lower-mantle structures”. In: *Geophysical Journal International* 207.2, 2016, pp. 1122–1136. URL: http://dx.doi.org/10.1093/gji/ggw324 (cit. on p. 75).

[176] Raj Moulik et al. *Three-dimensional Reference Earth Model webpage*. https://www.geol.umd.edu/facilities/seismology/rem-3d/ (cit. on p. 75).

[177] Francis Birch. “Elasticity and constitution of the Earth’s interior”. In: *Journal of Geophysical Research* 57.2, 1952, pp. 227–286 (cit. on p. 75).

[178] Jean-Paul Poirier. “Light elements in the Earth’s outer core: a critical review”. In: *Physics of the earth and planetary interiors* 85.3-4, 1994, pp. 319–337 (cit. on p. 75).

[179] Kei Hirose, Stéphane Labrosse, and John Hernlund. “Composition and state of the core”. In: *Annual Review of Earth and Planetary Sciences* 41, 2013, pp. 657–691 (cit. on p. 75).

[180] G Morard et al. “The Earth’s core composition from high pressure density measurements of liquid iron alloys”. In: *Earth and Planetary Science Letters* 373, 2013, pp. 169–178 (cit. on p. 75).

[181] Daniele Antonangeli and Eiji Ohtani. “Sound velocity of hcp-Fe at high pressure: experimental constraints, extrapolations and comparison with seismic models”. In: *Progress in Earth and Planetary Science* 2.1, 2015, p. 3 (cit. on p. 75).

[182] James Badro, Alexander S Côté, and John P Brodholt. “A seismologically consistent compositional model of Earth’s core”. In: *Proceedings of the National Academy of Sciences* 111.21, 2014, pp. 7542–7545 (cit. on p. 75).

[183] David C Rubie et al. “Heterogeneous accretion, composition and core–mantle differentiation of the Earth”. In: *Earth and Planetary Science Letters* 301.1-2, 2011, pp. 31–42 (cit. on p. 75).

[184] Edouard Kaminski and Marc Javoy. “A two-stage scenario for the formation of the Earth’s mantle and core”. In: *Earth and Planetary Science Letters* 365, 2013, pp. 97–107 (cit. on p. 75).

[185] Tatsuya Sakamaki et al. “Measurement of hydrous peridotite magma density at high pressure using the X-ray absorption method”. In: *Earth and Planetary Science Letters* 287.3-4, 2009, pp. 293–297 (cit. on p. 75).

[186] Eiji Ohtani et al. “Iron-water reaction at high pressure and temperature, and hydrogen transport into the core”. In: *Physics and chemistry of minerals* 32.1, 2005, pp. 77–82 (cit. on p. 75).

[187] Koichiro Umemoto and Kei Hirose. “Liquid iron-hydrogen alloys at outer core conditions by first-principles calculations”. In: *Geophysical Research Letters* 42.18, 2015. 2015GL065899, pp. 7513–7520. URL: http://dx.doi.org/10.1002/2015GL065899 (cit. on p. 75).

[188] Alfredo Placci and Emilio Zavattini. *On the possibility of using high-energy neutrinos to study the Earth’s interior*. Tech. rep. Geneva: CERN, 1973. URL: https://cds.cern.ch/record/2258764 (cit. on p. 76).

[189] L.V. Volkova and G.T. Zatsepin. “Passage of neutrinos through the Earth”. In: *Bull. Acad. Sci. USSR, Phys. Sci., v. 38, no. 5, pp. 151-154*, 1974 (cit. on p. 76).

[190] IP Nedyalkov. “Study of the Earth structure by means of neutrinos”. In: *mm I82*, 1982 (cit. on p. 76).

[191] Walter Winter. “Neutrino tomography: Learning about the earth’s interior using the propagation of neutrinos”. In: *Earth Moon Planets* 99, 2006, pp. 285–307 (cit. on p. 76).

[192] Pankaj Jain, John P. Ralston, and George M. Frichter. “Neutrino absorption tomography of the Earth’s interior using isotropic ultra-high energy flux”. In: *Astroparticle Physics* 12.3, 1999, pp. 193–198. URL: http://www.sciencedirect.com/science/article/pii/S0927650599000882 (cit. on p. 76).

[193] Matias M Reynoso and Oscar A Sampayo. “On neutrino absorption tomography of the Earth”. In: *Astroparticle Physics* 21.3, 2004, pp. 315 –324. URL: http://www.sciencedirect.com/science/article/pii/S0927650504000179 (cit. on p. 76).

[194] M. C. Gonzalez-Garcia et al. “Radiography of Earth’s Core and Mantle with Atmospheric Neutrinos”. In: *Phys. Rev. Lett.* 100.6, 2008, p. 061802. URL: http://link.aps.org/doi/10.1103/PhysRevLett.100.061802 (cit. on p. 76).

[195] Thomas L Wilson. “Neutrino tomography: Tevatron mapping versus the neutrino sky”. In: *Nature* 309.5963, 1984, p. 38 (cit. on p. 76).

[196] C. Kuo et al. “Extraterrestrial neutrinos and Earth structure”. In: *Earth and Planetary Science Letters* 133, 1995, pp. 95–103. URL: http://adsabs.harvard.edu/abs/1995E%26PSL.133...95K (cit. on p. 76).

[197] A. De Rújula et al. “Neutrino exploration of the Earth”. In: *Physics Reports* 99.6, 1983, pp. 341 –396. URL: http://www.sciencedirect.com/science/article/pii/0370157383901084 (cit. on p. 76).

[198] Andrea Donini, Sergio Palomares-Ruiz, and Jordi Salvado. “Neutrino tomography of the Earth”. In: 2018 (cit. on p. 77).

[199] M. G. Aartsen et al. “Searches for Sterile Neutrinos with the IceCube Detector”. In: *Phys. Rev. Lett.* 117.7, 2016, p. 071801 (cit. on p. 77).

[200] M. G. Aartsen et al. “Measurement of the multi-TeV neutrino cross section with IceCube using Earth absorption”. In: *Nature* 551, 2017, pp. 596–600 (cit. on p. 77).

[201] A. Nicolaidis. “Neutrinos for Geophysics”. In: *Phys. Lett.* B200, 1988, pp. 553–559 (cit. on p. 78).

[202] A. Nicolaidis, M. Jannane, and A. Tarantola. “Neutrino Tomography of the Earth”. In: *J. Geophys. Res. Solid Earth* 96, 1991, pp. 21811–21817 (cit. on p. 78).

[203] Manfred Lindner et al. “Tomography of the earth’s core using supernova neutrinos”. In: *Astropart. Phys.* 19, 2003, pp. 755–770 (cit. on p. 78).

[204] Evgeny K. Akhmedov, M. A. Tortola, and J. W. F. Valle. “Geotomography with solar and supernova neutrinos”. In: *JHEP* 06, 2005, p. 053 (cit. on p. 78).

[205] Ara N. Ioannisian and Alexei Yu. Smirnov. “Matter effects of thin layers: Detecting oil by oscillations of solar neutrinos”. In: 2002 (cit. on p. 78).

[206] C. A. Argüelles, M. Bustamante, and A. M. Gago. “Searching for cavities of various densities in the Earth’s crust with a low-energy $\bar{\nu}_e$ β-beam”. In: *Mod. Phys. Lett.* A30.29, 2015, p. 1550146 (cit. on p. 78).

[207] Takehiko Asaka et al. “Tomography by neutrino pair beam”. In: *Phys. Lett.* B785, 2018, pp. 536–542 (cit. on p. 78).

[208] K. Abe et al. “Solar Neutrino Measurements in Super-Kamiokande-IV”. In: *Phys. Rev.* D94.5, 2016, p. 052010 (cit. on p. 78).

[209] A. N. Ioannisian and A. Yu. Smirnov. “Attenuation effect and neutrino oscillation tomography”. In: 2017 (cit. on pp. 78, 79).

[210] Ara Ioannisian, Alexei Smirnov, and Daniel Wyler. “Scanning the Earth with solar neutrinos and DUNE”. In: *Phys. Rev.* D96.3, 2017, p. 036005 (cit. on p. 78).

[211] Tommy Ohlsson and Håkan Snellman. “Neutrino oscillations with three flavors in matter: Applications to neutrinos traversing the Earth”. In: *Physics Letters B* 474.1, 2000, pp. 153 –162. URL: http://www.sciencedirect.com/science/article/pii/S0370269300000083 (cit. on p. 79).

[212] Robert J. Geller and Tatsuhiko Hara. “Geophysical aspects of very long baseline neutrino experiments”. In: *Nucl. Instrum. Meth.* A503, 2003, pp. 187–191 (cit. on p. 79).

[213] Gian Luigi Fogli, G. Lettera, and E. Lisi. “Effects of matter density variations on dominant oscillations in long baseline neutrino experiments”. In: 2001 (cit. on p. 79).

[214] Toshihiko Ota and Joe Sato. “Yet another correlation in the analysis of CP violation using a neutrino oscillation experiment”. In: *Phys. Rev.* D67, 2003, p. 053003 (cit. on p. 79).

[215] Raj Gandhi and Walter Winter. “Physics with a very long neutrino factory baseline”. In: *Phys. Rev. D* 75.5, 2007, p. 053002. URL: http://link.aps.org/doi/10.1103/PhysRevD.75.053002 (cit. on p. 79).

[216] Hisakazu Minakata and Shoichi Uchinami. "In situ determination of Earth matter density in a neutrino factory". In: *Phys. Rev. D* 75, 72007, p. 073013. URL: https://link.aps.org/doi/10.1103/PhysRevD.75.073013 (cit. on p. 79).

[217] A. N. Ioannisian and A. Yu. Smirnov. "Neutrino oscillations in low density medium". In: *Phys. Rev. Lett.* 93, 2004, p. 241801 (cit. on p. 79).

[218] T. Ohlsson and W. Winter. "Could one find petroleum using neutrino oscillations in matter?" In: *EPL (Europhysics Letters)* 60.1, 2002, p. 34. URL: http://stacks.iop.org/0295-5075/60/i=1/a=034 (cit. on p. 79).

[219] Tommy Ohlsson and Walter Winter. "Reconstruction of the Earth's matter density profile using a single neutrino baseline". In: *Phys. Lett.* B512.3-4, 2001, pp. 357–364. URL: http://dx.doi.org/10.1016/S0370-2693(01)00731-6 (cit. on p. 79).

[220] Masafumi Koike and Joe Sato. "Effects of matter density fluctuation in long baseline neutrino oscillation experiments". In: *Mod. Phys. Lett.* A14, 1999, pp. 1297–1302 (cit. on p. 79).

[221] Walter Winter. "Probing the absolute density of the Earth's core using a vertical neutrino beam". In: *Phys. Rev. D* 72.3, 2005, p. 037302. URL: http://link.aps.org/doi/10.1103/PhysRevD.72.037302 (cit. on p. 80).

[222] Walter Winter. "Atmospheric Neutrino Oscillations for Earth Tomography". In: *Nucl. Phys.* B908, 2016, pp. 250–267 (cit. on pp. 80, 137, 208).

[223] The KM3NeT collaboration. *KM3NeT public website, accessed July 13th 2018*. http://www.km3net.org (cit. on p. 86).

[224] M. Ageron et al. *Technical Design Report of the MEUST infrastructure*. Tech. rep. 2013. URL: http://meust.cnrs.fr/TDR-MEUST_infrastructure.pdf (cit. on pp. 87, 89).

[225] Diego Real. "The electronics readout and data acquisition system of the KM3NeT neutrino telescope node". In: *AIP Conference Proceedings* 1630.1, 2014, pp. 102–105. URL: https://aip.scitation.org/doi/abs/10.1063/1.4902782 (cit. on p. 88).

[226] The KM3NeT collaboration. "Expansion cone for the 3-inch PMTs of the KM3NeT optical modules". In: *Journal of Instrumentation* 8.03, 2013, T03006. URL: http://stacks.iop.org/1748-0221/8/i=03/a=T03006 (cit. on p. 88).

[227] S. Aiello et al. "Characterisation of the Hamamatsu photomultipliers for the KM3NeT Neutrino Telescope". In: *JINST* 13.05, 2018, P05035 (cit. on pp. 88, 91).

[228] Christophe Hugon. "GEANT4 simulation of optical modules in neutrino telescopes". In: *Proceedings of the 34th International Cosmic Ray Conference, 29 Jul-6 Aug 2015, The Hague, The Netherlands*. PoS(ICRC2015)1106. URL: https://pos.sissa.it/236/1106/pdf (cit. on p. 88).

[229] S. Adrian-Martinez et al. "Deep sea tests of a prototype of the KM3NeT digital optical module". In: *The European Physical Journal C* 74.9, 2014, p. 3056. URL: https://doi.org/10.1140/epjc/s10052-014-3056-3 (cit. on pp. 90, 91, 98).

[230] S. Adrián-Martínez et al. "The prototype detection unit of the KM3NeT detector". In: *Eur. Phys. J.* C76.2, 2016, p. 54 (cit. on pp. 90, 91, 98).

[231] Irene DI PALMA. *Results from the first ARCA and ORCA detector units*. 2018. URL: https://doi.org/10.5281/zenodo.1300849 (cit. on p. 90).

[232] A. Albert et al. "Long-term monitoring of the ANTARES optical module efficiencies using 40K decays in sea water". Submitted to EPJC. URL: https://arxiv.org/abs/1805.08675 (cit. on p. 90).

[233] Steven H.D. Haddock, Mark A. Moline, and James F. Case. "Bioluminescence in the Sea". In: *Annual Review of Marine Science* 2.1, 2010. PMID: 21141672, pp. 443–493. URL: https://doi.org/10.1146/annurev-marine-120308-081028 (cit. on p. 90).

[234] J. A. Aguilar et al. "First results of the Instrumentation Line for the deep-sea ANTARES neutrino telescope". In: *Astropart. Phys.* 26, 2006, pp. 314–324 (cit. on p. 90).

[235] Stephanie Escoffier. "The ANTARES detector: background sources and effects on detector performance". In: *Proceedings, 30th International Cosmic Ray Conference (ICRC 2007): Merida, Yucatan, Mexico, July 3-11, 2007*. Vol. 5. 2007, pp. 1349–1352. URL: http://indico.nucleares.unam.mx/contributionDisplay.py?contribId=547&confId=4 (cit. on p. 90).

[236] Karel Melis. "In-Situ Calibration of KM3NeT". In: *35th International Cosmic Ray Conference, July 10-20, 2017, Bexco, Busan, Korea*. PoS(ICRC2017)1059. 2017. URL: https://pos.sissa.it/301/1059/pdf (cit. on pp. 92, 93, 98).

[237] Distefano, Carla. "gSeaGen: A GENIE-based code for neutrino telescopes". In: *EPJ Web of Conferences* 116, 2016, p. 08001. URL: https://doi.org/10.1051/epjconf/201611608001 (cit. on pp. 95, 96).

[238] Tomczak, M. *Seawater density calculator*. https://www.mt-oceanography.info/Utilities/density.html (cit. on p. 95).

[239] N P. Fofonoff and R Millard. "Algorithms for Computation of Fundamental Properties of Seawater". In: 44, 1983 (cit. on p. 95).

[240] Y. Becherini et al. "A Parameterisation of single and multiple muons in the deep water or ice". In: *Astropart. Phys.* 25, 2006, pp. 1–13 (cit. on p. 96).

[241] G Carminati A et al. "Atmospheric MUons from PArametric formulas: a fast GEnerator for neutrino telescopes (MUPAGE)". In: 179, 2008, pp. 915–923. URL: https://arxiv.org/abs/0802.0562 (cit. on p. 96).

[242] G. Battistoni et al. "Calculation of the TeV prompt muon component in very high energy cosmic ray showers". In: *Astroparticle Physics* 4.4, 1996, pp. 351 –363. URL: http://www.sciencedirect.com/science/article/pii/092765059500050X (cit. on p. 96).

[243] M. Ambrosio et al. "High statistics measurement of the underground muon pair separation at Gran Sasso". In: *Phys. Rev. D* 60, 31999, p. 032001. URL: https://arxiv.org/abs/hep-ex/9901027 (cit. on p. 96).

[244] A. G. Tsirigotis, A. Leisos, and S. E. Tzamarias. "HOU Reconstruction & Simulation (HOURS): A complete simulation and reconstruction package for very large volume underwater neutrino telescopes". In: *Nucl. Instrum. Meth.* A626-627, 2011, S185–S187 (cit. on p. 97).

[245] A. G. Tsirigotis et al. "Hellenic Open University Reconstruction & Simulation (HOURS) software package: User Guide & short reference of Event Generation, Cherenkov photon production and Optical Module simulation". In: 2017 (cit. on p. 97).

[246] S. Agostinelli et al. "Geant4—a simulation toolkit". In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 506.3, 2003, pp. 250 –303. URL: http://www.sciencedirect.com/science/article/pii/S0168900203013688 (cit. on p. 97).

[247] David Bailey. "Monte Carlo tools and analysis methods for understanding the ANTARES experiment and predicting its sensitivity to dark matter". PhD thesis. Univesrity of Oxford, 2002 (cit. on p. 97).

[248] S. Adrián-Martínez et al. "Long term monitoring of the optical background in the Capo Passero deep-sea site with the NEMO tower prototype". In: *Eur. Phys. J.* C76.2, 2016, p. 68 (cit. on p. 98).

[249] Leo Breiman. "Random Forests". In: *Machine Learning* 45.1, 2001, pp. 5–32. URL: https://doi.org/10.1023/A:1010933404324 (cit. on p. 103).

[250] Stefan Geisselsöder. "Model-independent search for neutrino sources with the ANTARES neutrino telescope". doctoralthesis. Friedrich-Alexander-Universität Erlangen-Nürnberg (FAU), 2016, p. 180 (cit. on p. 103).

[251] Thomas Eberl et al. *Tau neutrino appearance with KM3NeT/ORCA*. 2018. URL: https://doi.org/10.5281/zenodo.1292823 (cit. on pp. 104, 211).

[252] Martín Abadi et al. *TensorFlow: Large-Scale Machine Learning on Heterogeneous Systems*. Software available from tensorflow.org. 2015. URL: https://www.tensorflow.org/ (cit. on p. 104).

[253] Michael Moser. *OrcaNet CNN classifier for ORCA*. https://github.com/ViaFerrata/OrcaNet (cit. on p. 104).

[254] M. G. Aartsen et al. "Computational Techniques for the Analysis of Small Signals in High-Statistics Neutrino Oscillation Experiments". In: 2018 (cit. on pp. 112, 137).

[255] Database. *Geochemical Earth Model Reference*. https://earthref.org/GERMRD/datamodel/ (cit. on pp. 114, 115).

[256] Joachim Kopp. “Efficient numerical diagonalization of hermitian 3 x 3 matrices”. In: *Int. J. Mod. Phys.* C19, 2008, pp. 523–548 (cit. on p. 116).

[257] Patrick Huber, M. Lindner, and W. Winter. “Simulation of long-baseline neutrino oscillation experiments with GLoBES (General Long Baseline Experiment Simulator)”. In: *Comput. Phys. Commun.* 167, 2005, p. 195 (cit. on p. 116).

[258] Patrick Huber et al. “New features in the simulation of neutrino oscillation experiments with GLoBES 3.0: General Long Baseline Experiment Simulator”. In: *Comput. Phys. Commun.* 177, 2007, pp. 432–438 (cit. on p. 116).

[259] João A. B. Coelho. “Probing new physics with atmospheric neutrinos at KM3NeT-ORCA”. In: *J. Phys. Conf. Ser.* 888.1, 2017, p. 012115 (cit. on pp. 116, 211).

[260] Rene Brun and Fons Rademakers. “An Object Oriented Data Analysis Framework”. In: *AIHENP’96 Workshop, Lausanne, Sep. 1996, Nucl. Inst. & Meth. in Phys. Res. A 389 (1997) 81-86*. 1997. URL: http://root.cern.ch/ (cit. on pp. 127, 266).

[261] *IN2P3 Computing Centre*. https://cc.in2p3.fr/en/ (cit. on p. 134).

[262] Diego Casadei. “Estimating the selection efficiency”. In: *JINST* 7, 2012, P08021 (cit. on p. 141).

[263] Marc Paterno. “Calculating efficiencies and their uncertainties”. In: 2004 (cit. on p. 141).

[264] Kyle Cranmer. *Practical Statistics for Particle Physics*. Lectures at the CERN School HEP, Romania. 2011 (cit. on pp. 144, 154, 155).

[265] X. Qian et al. “Statistical Evaluation of Experimental Determinations of Neutrino Mass Hierarchy”. In: *Phys. Rev.* D86, 2012, p. 113011 (cit. on p. 153).

[266] W J Metzger. *Statistical methods in data analysis*. Nijmegen: Nijmegen Univ. Fys. Lab., 2002. URL: https://cds.cern.ch/record/482882 (cit. on p. 154).

[267] Gary J. Feldman and Robert D. Cousins. “A Unified approach to the classical statistical analysis of small signals”. In: *Phys. Rev.* D57, 1998, pp. 3873–3889 (cit. on p. 154).

[268] Frederick James. *Statistical Methods in Experimental Physics: 2nd Edition*. 2006 (cit. on p. 155).

[269] Roger J. Barlow and Christine Beeston. “Fitting using finite Monte Carlo samples”. In: *Comput. Phys. Commun.* 77, 1993, pp. 219–228 (cit. on p. 156).

[270] Pekka Sinervo. “Definition and Treatment of Systematic Uncertainties in High Energy Physics and Astrophysics”. In: *eConf* C030908, 2003. [,122(2003)], TUAT004 (cit. on p. 159).

[271] IceCube-Gen2. “PINGU: A Vision for Neutrino and Particle Physics at the South Pole”. In: 2016 (cit. on p. 165).

[272] James, Frederick and Winkler, Matthias. *Minuit2 user manual*. https://root.cern.ch/root/htmldoc/guides/minuit2/Minuit2.pdf (cit. on pp. 168, 266).

[273] Hiroshi Nunokawa, Stephen J. Parke, and Renata Zukanovich Funchal. “Another possible way to determine the neutrino mass hierarchy”. In: *Phys. Rev.* D72, 2005, p. 013009. URL: https://arxiv.org/pdf/hep-ph/0503283.pdf (cit. on pp. 182, 183).

[274] Andre de Gouvea, James Jenkins, and Boris Kayser. “Neutrino mass hierarchy, vacuum oscillations, and vanishing |U(e3)|”. In: *Phys. Rev.* D71, 2005, p. 113009 (cit. on p. 182).

[275] Simon Bourret and Liam Quinn. *Sensitivity of ORCA to the neutrino mass ordering and oscillation parameters*. Poster presentation in the 28th International Conference on Neutrino Physics and Astrophysics (Neutrino 2018) Heidelberg, Germany, June 4-9, 2018. 2018. URL: https://doi.org/10.5281/zenodo.1300771 (cit. on pp. 193, 202, 265).

[276] Uli Katz. *Future neutrino telescopes in water and ice*. Talk in the 28th International Conference on Neutrino Physics and Astrophysics (Neutrino 2018) Heidelberg, Germany, June 4-9, 2018. 2018. URL: https://doi.org/10.5281/zenodo.1287686 (cit. on pp. 193, 265).

[277] Simon Bourret, João A. B. Coelho, and Véronique Van Elewyck. “Neutrino oscillation tomography of the Earth with KM3NeT-ORCA”. In: *27th International Conference on Neutrino Physics and Astrophysics (Neutrino 2016) London, United Kingdom, July 4-9, 2016*. 2017. URL: http://inspirehep.net/record/1513140/files/arXiv:1702.03723.pdf (cit. on pp. 204, 265).

[278] L Bezrukov and V Sinev. “Atmospheric neutrinos for investigation of Earth interior”. In: *Physics of Particles and Nuclei* 47.6, 2016, pp. 915–917 (cit. on p. 208).

[279] Vladimir N. Larin. *Hydridic Earth: the New Geology of Our Primordially Hydrogen-rich Planet.* 1993 (cit. on p. 208).

[280] K. Abe et al. "Search for CP violation in Neutrino and Antineutrino Oscillations by the T2K experiment with 2.2×10^{21} protons on target". In: 2018 (cit. on p. 210).

[281] Mayly Sanchez. *NOvA Results and Prospects.* 2018. URL: https://doi.org/10.5281/zenodo.1286758 (cit. on p. 210).

[282] Bjoern Wonsak. *Status and Prospects of the JUNO Experiment.* 2018. URL: https://doi.org/10.5281/zenodo.1286850 (cit. on p. 210).

[283] Simon Bourret et al. *What Would It Take for an Atmospheric Neutrino Detector to Constrain the Hydrogen Content of the Earth's Core ?* Poster presentation at the American Geophysical Union (AGU) Fall Meeting, New Orleans, December 2017. URL: https://agu.confex.com/agu/fm17/meetingapp.cgi/Paper/231578 (cit. on p. 212).

[284] Jannik Hofestädt, Thomas Eberl, and Marc Bruchner. *KM3NeT/Super-ORCA: Measuring the leptonic CP-phase with atmospheric neutrinos — a feasibility study.* 2018. URL: https://doi.org/10.5281/zenodo.1292936 (cit. on p. 212).

[285] Joshua P. Ellis. "TikZ-Feynman: Feynman diagrams with TikZ". In: *Computer Physics Communications* 210, 2017, pp. 103 –123. URL: http://www.sciencedirect.com/science/article/pii/S0010465516302521 (cit. on p. 231).

[286] Mark Srednicki. *Quantum Field Theory.* Freely available online as a preprint. Cambridge: Cambridge Univ. Press, 2007. URL: http://web.physics.ucsb.edu/~mark/qft.html (cit. on pp. 231, 240, 241).

[287] Michael E. Peskin and Daniel V. Schroeder. *An Introduction to quantum field theory.* Reading, USA: Addison-Wesley, 1995. ISBN: 9780201503975, 0201503972. URL: http://www.slac.stanford.edu/~mpeskin/QFT.html (cit. on pp. 231, 240, 241).

[288] Gerardus 't Hooft and Martinus J G Veltman. *Diagrammar.* CERN Yellow Reports: Monographs. Geneva: CERN, 1973. URL: http://cds.cern.ch/record/186259 (cit. on p. 232).

[289] T. D. Lee. "A brief history of the muon". In: *Hyperfine Interactions* 86.1, 1994, pp. 439–453. URL: https://doi.org/10.1007/BF02068933 (cit. on p. 235).

[290] B. Pontecorvo. "Nuclear Capture of Mesons and the Meson Decay". In: *Phys. Rev.* 72, 31947, pp. 246–247. URL: https://link.aps.org/doi/10.1103/PhysRev.72.246 (cit. on p. 235).

[291] T. D. Lee, M. Rosenbluth, and C. N. Yang. "Interaction of Mesons with Nucleons and Light Particles". In: *Phys. Rev.* 75, 51949, pp. 905–905. URL: https://link.aps.org/doi/10.1103/PhysRev.75.905 (cit. on p. 235).

[292] Simon Bourret, João A. B. Coelho, and Véronique Van Elewyck. "Neutrino oscillation tomography of the Earth with KM3NeT-ORCA". In: *35th International Cosmic Ray Conference, July 10-20, 2017, Bexco, Busan, Korea.* PoS(ICRC2017)1020. 2017. URL: https://pos.sissa.it/301/1020/pdf (cit. on p. 265).

[293] Antoine Kouchner and Joao Coelho. "Measuring the neutrino mass ordering and other oscillation parameters with KM3NeT-ORCA". In: *35th International Cosmic Ray Conference, July 10-20, 2017, Bexco, Busan, Korea.* PoS(ICRC2017)1027. 2017. URL: https://pos.sissa.it/301/1027/pdf (cit. on p. 265).

[294] Simon Bourret. *KM3NeT/ORCA: Measuring neutrino oscillations and the mass hierarchy in the Mediterranean.* Talk in the 26th International Workshop on Weak Interactions and Neutrinos (WIN2017), UC Irvine, Irvine, CA, USA, June 18-24, 2017. 2017. URL: https://indico.fnal.gov/event/9942/session/13/contribution/64 (cit. on p. 265).

[295] Martijn Jongen. "Sensitivity to the Neutrino Mass Hierarchy of KM3NeT/ORCA". In: *Proceedings of the 34^{th} International Cosmic Ray Conference, 29 Jul-6 Aug 2015, The Hague, The Netherlands.* PoS(ICRC2015)1092 (cit. on p. 270).

[296] S. K. Agarwalla et al. "The mass-hierarchy and CP-violation discovery reach of the LBNO long-baseline neutrino experiment". In: *JHEP* 05, 2014. arXiv:1312.6520v3, p. 094 (cit. on p. 271).

[297] Simon Bourret. "Sensitivity to the neutrino mass hierarchy and discovery of CP violation with the Protvino-to-ORCA proposal". MA thesis. ETH Zürich (Switzerland) and École polytechnique (France), 2015 (cit. on p. 271).

[298] Dimitri van Heesch. *Doxygen: generate documentation from source code*. http://www.stack.nl/~dimitri/doxygen/index.html (cit. on p. 271).

Abstract:

Neutrino flavour oscillations, discovered at the turn of the 21st century, currently provide the most direct window on physics beyond the Standard Model. The KM3NeT collaboration has started the construction of a megaton-scale Cherenkov neutrino detector deep underwater in the Mediterranean: ORCA (*Oscillation Research with Cosmics in the Abyss*). ORCA's main goal will be the determination of the neutrino mass hierarchy, a major unknown in the current understanding of neutrino oscillations, via the detection of large samples of atmospheric neutrinos crossing the Earth and the analysis of their flavour oscillations. These oscillations are enhanced by matter effects, sensitive to the electron density along the neutrino paths. Beyond the neutrino mass hierarchy and the measurement of atmospheric oscillation parameters, this will allow ORCA to measure the electron density of the deep Earth, and possibly provide constraints on the chemical composition of its innermost layers.

This thesis presents a new study, undertaken within the KM3NeT collaboration, of the sensitivity of ORCA to the determination of the neutrino mass hierarchy, oscillation parameters, and the electron density in the outer core and lower mantle of the Earth. The focus is given to the development of a new analysis methodology intended to account for detector effects as precisely as possible by the use of a Monte Carlo based correlated model of detector response. Statistical methods and systematic uncertainties are also addressed in some detail. The methodology allows for a preliminary exploration of the potential of improved analysis strategies. Sensitivity studies show that ORCA is expected to achieve a median sensitivity to the NMH determination at the level of 3σ or better after a few years of operation, depending on the true hierarchy and true value of the oscillation parameters. Competitive performance for the measurement of atmospheric oscillation parameters is also expected. Assuming a normal mass hierarchy, the electron density can be measured with a precision of about 5% in the lower mantle, and 7% in the outer core.

Keywords: oscillations, flavour, neutrino, atmospheric, matter, Earth, mass, hierarchy.

Résumé :

Les oscillations de saveur des neutrinos, découvertes au tournant du 21ème siècle, constituent à ce jour le signe le plus direct de physique au-delà du Modèle Standard. La collaboration KM3NeT a commencé la construction d'un détecteur à effet Tcherenkov de neutrinos au fond de la mer Méditerranée : ORCA (*Oscillation Research with Cosmics in the Abyss*). Le but principal d'ORCA sera la détermination de la hiérarchie de masse des neutrinos, un paramètre fondamental, encore inconnu, du modèle théorique des oscillations de neutrinos. Le principe de la mesure est la détection d'un large échantillon de neutrinos atmosphériques ayant traversé la Terre, et l'analyse de leurs oscillations de saveur. Ces oscillations sont modifiées et augmentées par les effets de matière, dus à la présence d'une densité d'électrons le long de la trajectoire des neutrinos. Au-delà de la hiérarchie de masse et de la mesure des paramètres d'oscillation atmosphériques, ORCA sera alors capable de mesurer la densité en électrons et donc, en principe, de fournir des informations nouvelles sur la composition chimique des couches profondes de la Terre.

Cette thèse présente une nouvelle étude, menée au sein de la collaboration KM3NeT, de la sensibilité d'ORCA à la détermination de la hiérarchie de masse des neutrinos, aux paramètres d'oscillations atmosphériques, et à la densité en électrons dans le noyau externe et le manteau inférieur de la Terre. La thèse se concentre sur le développement d'une nouvelle méthode d'analyse dont le but est de rendre compte de la réponse du détecteur de la manière la plus précise possible. Pour cela, la méthode est basée sur l'utilisation de matrices de réponses corrélées et construites à partir de simulations Monte Carlo. Les méthodes statistiques et sources d'incertitude systématique sont aussi étudiées en détail. La méthode développée permet une première exploration de nouvelles stratégies d'analyse optimisées. Les études de sensibilité montrent qu'ORCA atteindra une sensibilité médiane à la hiérarchie de masse d'au moins 3σ après quelques années d'opération, le résultat dépendant en grande partie de la valeur réelle de la hiérarchie de masse et des autres paramètres d'oscillation encore inconnus. Une performance compétitive pour la mesure des paramètres d'oscillation atmosphériques est aussi attendue. Dans le cas où la hiérarchie de masse est normale et après 10 ans de prise de données, la densité en électrons pourra être mesurée avec une précision d'environ 5% dans le manteau inférieur, et 7% dans le noyau externe.

Mots-clés : oscillations, saveur, neutrino, atmosphérique, matière, Terre, masse, hiérarchie.