



HAL
open science

Quantifying uncertainties in Large Eddy Simulations of auto-ignition-stabilized H₂ flames

Guilhem Lavabre

► **To cite this version:**

Guilhem Lavabre. Quantifying uncertainties in Large Eddy Simulations of auto-ignition-stabilized H₂ flames. Probability [math.PR]. Université Paris-Saclay, 2023. English. NNT : 2023UPAST100 . tel-04457295

HAL Id: tel-04457295

<https://theses.hal.science/tel-04457295>

Submitted on 14 Feb 2024

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.

Quantifying uncertainties in Large Eddy Simulations of auto-ignition-stabilized H₂ flames

Thèse de doctorat de l'Université Paris-Saclay

École doctorale n° 579, S MEMAG
Spécialité de doctorat: Energétique

Graduate School: Science de l'ingénierie et des systèmes
Réfèrent: CentraleSupélec

Thèse préparée dans l'unité de recherche CNRS UPR 288 sous la
direction de **Ronan VICQUELIN**, Professeur à CentraleSupélec (EM2C)
et le co-encadrement d'**Olivier GICQUEL**, Professeur à
CentraleSupélec (EM2C)

Thèse présentée et soutenue à Gif s/Yvette, le 05/07/2023, par

Guilhem LAVABRE

Composition du jury

Membres avec voix délibérative

Arnaud MURA

Directeur de recherche, CNRS (PPRIME)

Alessandro PARENTE

Professeur, Université Libre de Bruxelles

Laurent SELLE

Directeur de recherche, CNRS (IMFT)

Pascale DOMINGO

Directrice de recherche, CNRS (CORIA)

Nathalie LAMOUREUX

Ingénieure de recherche, Université de Lille (PC2A)

Didier LUCOR

Directeur de recherche, CNRS (LISN)

Président

Rapporteur & Examineur

Rapporteur & Examineur

Examinatrice

Examinatrice

Examineur

Titre: Quantification d'incertitudes dans les simulations aux grandes échelles de flammes d'hydrogène stabilisées par auto-allumage

Mots clés: Quantification d'incertitudes, Combustion turbulente, Hydrogène, Jumeau numérique

Résumé: La combustion est une source d'énergie omniprésente dans notre société. Afin de minimiser ses émissions polluantes, dans un contexte où les conditions opératoires et les modèles de calculs sont incertains, nous devons être capables de quantifier les incertitudes de nos simulations prédictives.

Dans cette thèse, nous propageons les incertitudes provenant du mécanisme de cinétique chimique, de mesures expérimentales et du modèle de turbulence à travers la simulation aux grandes échelles (LES) d'une flamme turbulente. Nous choisissons de travailler sur la flamme de Cabra H_2 car son mécanisme de stabilisation principal est l'auto-allumage, ce qui la rend extrêmement sensible aux paramètres qui affectent les taux de réaction. Nous simulons la flamme de Cabra en utilisant une approche basée sur la tabulation de flammelettes instationnaires et sur des PDF présumées. La propagation d'incertitudes repose sur la réduction a priori de la dimension incertaine en utilisant une configuration simplifiée : un réacteur homogène isobare qui s'auto-allume. Une

étude Monte-Carlo sur ce réacteur dévoile un comportement très non-linéaire dans la plage d'incertitudes considérée. Ce comportement a été identifié à la transition entre les auto-allumages "faible" et "fort". Cette étude a aussi permis de réduire les incertitudes venant du mécanisme chimique et de la température du co-flow à un espace de dimension 2.

Dans cet espace réduit, une approche basée sur une surface de réponse est réalisée en utilisant seulement un faible nombre d'échantillons LES. Nous étudions également l'influence du coefficient de dissipation de la variance de fraction de mélange. Dans un premier temps, seule l'incertitude de la hauteur de stabilisation de la flamme est quantifiée. Dans un second temps, nous quantifions les incertitudes de toutes les quantités mesurées expérimentalement à toutes les positions dans le domaine de calcul. Cela est rendu possible par une analyse en composantes principales (PCA) et trois surfaces de réponses pour prédire les coefficients des plus importants modes de la PCA.

Title: Quantifying uncertainties in Large Eddy Simulations of auto-ignition-stabilized H₂ flames

Keywords: Uncertainty Quantification, Turbulent combustion, Hydrogen, Digital twin

Abstract: Combustion is a ubiquitous source of energy in our society. To minimize its pollutant emissions in a context where operating conditions and computational models are uncertain, we need to quantify the uncertainties of our predictive simulation results.

In this thesis, we propagate kinetic, experimental, and model uncertainties in the Large Eddy Simulation of a turbulent flame. We chose to work with the H₂ Cabra flame because its main stabilization mechanism is auto-ignition, which makes it extremely sensitive to parameters that impact reaction rates. We simulate the Cabra flame using an unsteady flamelet and presumed PDF approach. The uncertainty propagation relies on an a priori uncertain dimension reduction using a simplified configuration: an auto-igniting homogeneous reactor at constant pressure. An a priori Monte-Carlo study allows highlighting an extremely non-linear phenomenon in the uncertainty

range: the transition between weak and strong auto-ignition. It also allows reducing the uncertainty coming from the co-flow temperature and kinetic mechanism to a mere two dimensions.

In this reduced space, a surrogate-model-based approach is viable using a limited number of LES observations. We later investigate the sensitivity of the simulation to the dissipation coefficient of mixture fraction variance. We first quantify the uncertainty of the Cabra flame lift-off height. Then, we turn our gaze towards the quantification of uncertainties of every experimentally measured quantity at every location in the computational domain. This is done using a Principal Component Analysis on the training samples and three surrogate models to predict the coefficients of the most contributing modes brought forth by the PCA.

Remerciements

Je tiens à adresser mes remerciements à celles et ceux qui m'ont accompagné tout au long de ces années passées au laboratoire EM2C et qui m'ont aidé dans les divers aspects de la recherche, de la rédaction, et de la vie en général :

Tout d'abord, à Ronan Vicquelin et Olivier Gicquel, qui m'ont offert cette opportunité, qui ont fait preuve de pédagogie dans leurs explications, d'enthousiasme pour mes travaux, et de patience dans les (nombreux) moments difficiles,

A l'ensemble des contribuables français, dont les impôts ont, bien involontairement, financé mon travail,

A mon jury, pour sa bienveillance et ses intéressants retours,

A ma famille et à mes collègues et amis pour leur soutien,

En particulier, à Luc Lecointre, ami fidèle, colocataire exemplaire, collègue enjoué et véritable locomotive à vélo, à Corentin 'Helios' Grimaldi, camarade d'infortune et rayon de soleil des grises journées de rédaction, à Victorien Blanchard, pour sa bonne humeur communicative et son entrain sans faille, et à Junghwa Yi, pour son immense gentillesse et les folles soirées parisiennes,

Pour finir, à Lisa Bonatti, qui m'a ramassé à la petite cuillère et dont l'écoute et la sagesse m'ont peu à peu remis sur pieds, et à qui je dois d'être aujourd'hui à la fois vivant et diplômé.

Table of Contents

Introduction	1
Chapter 1: Handling uncertainties	11
1.1 Methods overview	11
1.2 Random sampling	13
1.3 Surrogate modeling	21
1.4 Sensitivity analysis	35
1.5 Uncertain dimension reduction	37
1.6 Error metrics	43
1.7 Takeaway points	46
Chapter 2: Uncertainty propagation in H ₂ /O ₂ auto-ignition simulations	47
2.1 Physical configurations and numerical setups	48
2.2 A qualitative comparison of deterministic H ₂ /O ₂ mechanisms	50
2.3 Uncertainties in hydrogen flame ignition	51
2.4 Numerical methods for uncertainty propagation	53
2.5 Analysis of uncertainty in the homogeneous reactor’s results and dimension reduction	54
2.6 Analysis of the auto-ignition delay time’s uncertainty in the complex configuration	57
2.7 Takeaway points	65
Chapter 3: The H ₂ Cabra flame	67
3.1 The experimental setup	67

3.2	Numerical simulation of the cabra flame	70
3.3	Results of the nominal simulation	86
3.4	Convergence study	87
3.5	Sensitivity of the simulation to the co-flow temperature	94
3.6	Sensitivity to the numerical setup	96
3.7	Takeaway points	99
Chapter 4:	Quantifying the uncertainty of the simulated Cabra flame’s lift-off height	101
4.1	The input uncertainties	101
4.2	Reduction of the input uncertain dimension	102
4.3	The computer experiment	108
4.4	Uncertainty propagation results for the Cabra flame’s lift-off height in Span(ξ_T, ξ_{12})	113
4.5	Investigating the uncertainty in the turbulence-combustion model	118
4.6	Takeaway points	125
Chapter 5:	Uncertainty Quantification in averaged LES fields	127
5.1	Reducing the input uncertain dimension	127
5.2	Reducing the output uncertain dimension	129
5.3	Surrogate modeling	133
5.4	Validation of the field prediction, as a whole	136
5.5	Analysis of the uncertain fields	139
5.6	Limitations of the study	144
5.7	Takeaway points	147
Conclusion	149
Bibliography	151

Chapter A: The Boiled-down Konnov (2019) mechanism	161
Chapter B: The UFPV table	163
Chapter C: Profiles of the simulated Cabra flame	173
Chapter D: Convergence study for the LES of the Cabra flame	183
Chapter E: Surrogate evaluations for the prediction of the Cabra flame's lift-off height	197
Chapter F: Averaged fields of the uncertain Cabra flame's training observations . . .	199
Chapter G: Modes retained for the prediction of the Cabra flame's averaged fields .	207
Chapter H: Prediction of the Cabra flame validation samples' averaged fields	215
Chapter I: Prediction of the Cabra flame validation samples' averaged profiles . . .	223

Introduction

The primary motivation: Uncertainty Quantification

Uncertainty Quantification (UQ) is a framework that allows taking into account unknown (epistemic uncertainties) and/or randomly fluctuating (aleatoric uncertainties) variables in a given model, in order to determine the uncertainty of a quantity of interest. These uncertainties can stem from many sources: measurement errors, inadequate computational models, unknown parameter values...

From this wide concept, diverse procedures can be derived. Among them, forward and inverse uncertainty propagations are becoming more and more commonplace:

- Forward uncertainty propagation is a procedure that uses the uncertain inputs of a model to evaluate the probability distribution of an output quantity of interest in the model. This can be useful, for example, to evaluate risks, compute safety margins, or estimate performance variability.
- Conversely, given a model and a set of uncertain outputs, inverse uncertainty propagation infers the probability distributions of the input variables of the model. This allows calibrating computational model parameters given uncertain measurement results.

One of the major interests of the uncertainty quantification framework, in particular its non-intrusive branch, is that "the model" can be practically anything that takes a set of variables as inputs and produces another set of output variables. "The model" can be a computer simulation, an analytical model, an experiment, etc. Therefore, uncertainty quantification can be used for extremely various problems. Hence, the methods presented and developed in this thesis are not limited to our particular use case. Three, very diverse, examples of fields that use uncertainty quantification are given below:

- Geosciences can use uncertainty quantification to predict the probability distributions of natural phenomena. One such example is the study by Li et al. [62] which evaluated the impact of initial conditions and wind-forcing uncertainties on the currents of the Gulf of Mexico. Another example is the seminal work of Krige [57] who introduced the Kriging method as a prospection tool for the Witwatersrand gold mines in South Africa.
- Modern Portfolio Theory, first introduced by Markowitz in 1952 [67], is the foundation of the modern financial sector. By applying the probability theory to the financial markets, he developed a new notion of risk and proposed to evaluate investment portfolios by evaluating both the expectation and variance of return. His work paved the way for previously-unseen profits by banks and investment funds. However, over-reliance on his risk-management method and its successors, coupled

with over-simplifications of the market mechanisms notoriously led to periodic financial crises in the 20th and 21st centuries (1987: Black Monday, 1990: LTCM collapse, 1994: Bond market crisis, 2007: Subprime mortgage crisis...) [49].

- Finally, engineering sciences are fond of uncertainty quantification. In particular, the combustion community, which initiated this thesis, is in demand for uncertainty quantification for purposes of robustifying efficient combustors, better predicting pollutant emissions, etc.

A more practical motivation: Combustion

Modern human civilization is addicted to combustion. From the myth of Prometheus to the industrial revolution and the wonders of the Saturn V moon rocket, humans have used fire as a means to exploit the chemical energy contained in their environment. A major tipping point was the 17th and 18th centuries when combustion moved from being a source of thermal energy (for cooking, metallurgy, or simply as a source of heat) to being a source of mechanical energy thanks to a flourishing of innovations in the steam engine. By the 19th century, combustion had become ubiquitous in human activities: Draught animals were replaced by engines; wood and peat were gradually replaced by coal, gas, and oil.

Mechanical energy, made cheaper and way more plentiful by this new technology, was consumed at an ever-increasing rate [44]. Among other examples, transportation habits have dramatically evolved since then. Where an intercontinental journey might have taken a few weeks aboard a sailing ship, the same journey is now done in a few hours aboard any commercial airplane. As a result, intercontinental flights have become commonplace for citizens of developed countries. This rebound effect probably reached its climax with the Concorde, a supersonic airliner which was routinely used for back-and-forth trips between Europe and America for business meetings.

For several decades now, the combustion of fossil fuels has been known to degrade our environment and its ecosystems by upsetting the global climate. Yet, despite a clear scientific consensus and collective awareness, mankind continues to prospect and exploit new fossil fuel deposits. At this point, it seems rather clear that our economic organization traps us in a prisoner's dilemma: the first institution, either company or country, trying to dispense with combustion would probably do without potential profits and spending power respectively. This is highlighted by several studies linking the per-capita Gross Domestic Product (GDP) with the per-capita greenhouse gas emissions [114, 38]. Therefore, it seems that mankind will likely continue to rely on combustion for as long as it finds fuel to burn. This tendency is suggested by Figure 1, which plots the evolution and shares, by source, of the world's total primary energy supply between 1800 and 2021 [88]. Since 1997 and the signature of the Kyoto Protocol, fossil fuel consumption has not shown any sign of decrease. On the contrary, mankind, as a whole, has burnt more and more fossil fuel since then, with only small drops in 2008 and 2020 caused by the subprime crisis, the Covid-19 pandemic, and the respective global recessions that followed.

The primary means of action to reduce greenhouse emissions is therefore to dramatically reduce the global energy supply and the share of combustion in this supply. Hence, it is mostly a drastic cultural and political change that is needed. Greenhouse emissions




Figure 1: Global primary energy consumption, in TWh, by source, between 1800 and 2021. Source: OWID [88]

and global warming are not problems that engineering sciences alone can solve.

For the activities that will be judged absolutely necessary for society, and for which combustion is definitely irreplaceable, the environmental cost must be minimized. In this prospect, some potential alternatives to fossil fuels are presently gaining traction. Among them, biofuels and dihydrogen are particularly in demand these days and draw in human resources and corporate and state funding.

Despite significant drawbacks at the production stage of these alternative fuels (soil depletion [104] and overall debatable carbon balance [121], and electric power requirements respectively), they are seen by decision-makers as valid substitutes that may reduce pollution without thwarting too much the energetic bulimia of our productivist society. This belief may lead to yet another rebound effect and should be considered with great caution.

Another area of work is the increase of fuel efficiency of (present and irreplaceable future) combustion devices and the reduction of their toxic byproducts like CO, NO_x, soot, etc. The motivation of this thesis stems from this context: Cautious uncertainty quantification is required to ensure that industrial systems stay in their safe and regulatory operating ranges, the alternative being the implementation of conservative safety margins which typically decrease performance.

As engineering has come to rely more and more on numerical simulation, Computational Fluid Dynamics (CFD) is now as commonplace in design offices as in research labs. Still, these simulations are subject to uncertainties, coming from the operating conditions (temperature, mixture composition in the case of biofuels, etc.), the combustion or turbulent models, or even from the chemical kinetic model itself.

Kempf et al. [50], for example, investigated the influence of the Smagorinsky constant, a key parameter of the turbulence model in Large Eddy Simulation (LES). Figure 2

presents the qualitative influence of its value on the dynamics of the Sydney flame, a non-premixed bluff-body configuration. This example exhibits that the choice of turbulence model constants is paramount when trying to correctly predict a turbulent flame in LES.




Figure 2: Influence of the value of the Smagorinsky constant on the dynamics of the LES of the Sydney bluff-body flame. Adapted from Kempf et al. [50]

Fiorina et al. [28] published another interesting study, where several research groups were asked to simulate the same flame. It corresponds to the unconfined Turbulent Stratified Flame from Darmstadt. Each group used its own numerical setup: each defined its own computational domain, mesh type and mesh size, kinetic mechanism and turbulent combustion model. Figure 3 presents the radial profiles of the Favre-averaged gas temperature at an axial location of 15 mm from the burner found by the different groups (plain and dashed lines denote different wall boundary conditions). Although all results are in global agreement with each other and with the experiment, some discrepancies remain between the prediction of the different groups, due to their modelization choices.




Figure 3: Radial profiles of the gas temperature in the Turbulent Stratified Flame, as simulated by each research group. Adapted from Fiorina et al. [28]

The two previous examples show that even high-fidelity simulations like Large Eddy Simulation are still plagued by epistemic uncertainties resulting from the modelization

choice. In other cases, experimental measurement uncertainties also make measurements and simulations seemingly disagree. This is the case of the H₂ Cabra flame [11]: a turbulent, non-premixed, auto-ignition-stabilized flame. Figure 4 shows a collection of simulations and experimental measurements of the Cabra flame lift-off height.

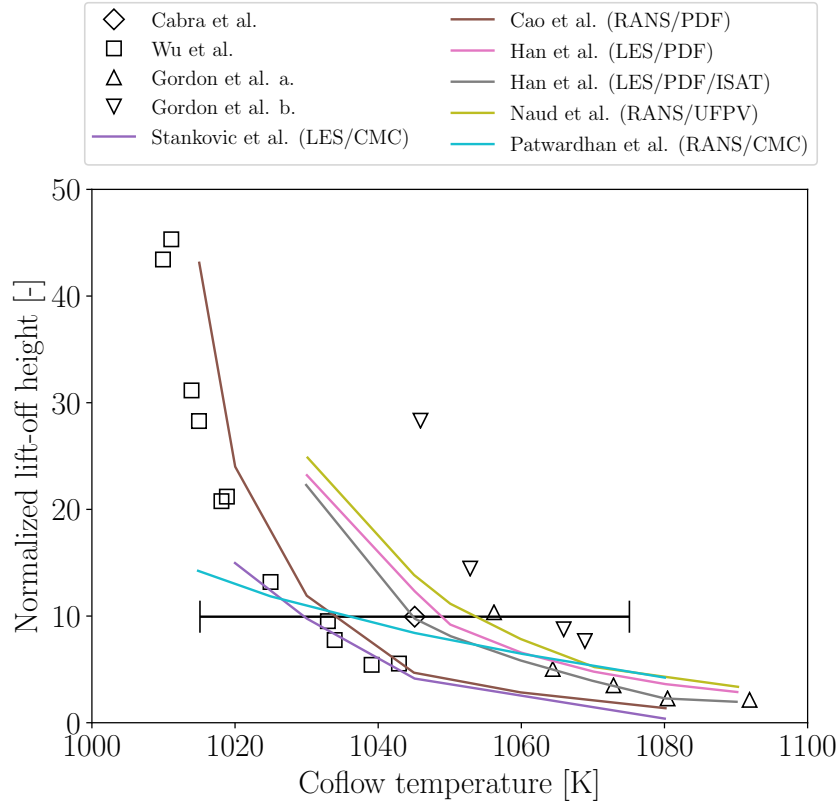


Figure 4: Lift-off height of the Cabra flame. Measurements [11, 34, 120] are the black symbols, simulations [105, 13, 35, 74, 78] are the colored lines. Note the measurement uncertainty of the co-flow temperature (black error bar)

Again, in this case, discrepancies appear between the simulation results of those authors. Perhaps more importantly, the experimental measurements show both a strong sensitivity of the lift-off height to the co-flow temperature and a significant uncertainty in the measurement of this temperature. In this configuration, proper uncertainty quantification is much-needed to evaluate the agreement between experiments and simulations and highlight the underlying phenomena.

Uncertainty quantification in inexpensive combustion simulations

In the literature, most of the intersection between combustion and uncertainty quantification revolves around the uncertainties in kinetic mechanisms. A kinetic mechanism is a set of reactions and equations to determine their rate constants as a function of the ambient conditions. In the vast majority of cases, rate constants are found using an Arrhenius law:

$$k_j = A_j T^{\beta_j} \exp\left(\frac{-E_{a_j}}{RT}\right)$$

Where k_j is the rate constant of reaction j , A_j is its pre-exponential factor, T is the gas temperature, β_j is a dimensionless constant, E_{a_j} is the energy activation of reaction j , and R is the universal gas constant. A_j , β_j , and E_{a_j} are either inferred from experiments [39] or with collision theory computations [48] which are currently affordable only for extremely simple molecules. In the former case, measurement uncertainties and dependencies between all the coefficients lead to uncertainties of the coefficients in the Arrhenius laws of the kinetic mechanism.

Theoretically, a joint probability distribution should be considered for all the coefficients of all the reactions in the mechanism. In practice, those coefficients are often assumed independent, and the pre-exponential constants are the only ones considered uncertain. Their uncertainty is typically represented by a log-normal distribution parametrized by an "Uncertainty Factor" (UF), such that:

$$\xi_j = \frac{\log(A_j/A_j^0)}{\frac{1}{3}\log(\text{UF}_j)} \sim \mathcal{N}(0, 1)$$

Forward uncertainty propagation methods are readily available to propagate the variability of some simulation inputs to a given quantity of interest. They have been applied to simple combustion problems such as homogeneous reactors and one-dimensional laminar flames [81, 87, 86, 126, 53]. For example, Phenix et al. [81] propagated the kinetic uncertainties from an 8-reaction reduced hydrogen combustion mechanism through the simulation of a plug-flow reactor in a supercritical state. Figure 5 reports the results of a Monte-Carlo study of this configuration. Significant uncertainties are observed. In particular, Phenix et al. show that they are much larger than the difference between simulations performed with and without real-gas corrections in the equation of state.

In most practical cases, non-intrusive methods are used, as they do not require any modification of the CFD solver. It must be noted, though, that intrusive methods have also been investigated [86]. In the non-intrusive studies listed above, random sampling was used to propagate uncertainties. This is not a realistic approach for more expensive simulations because the simplest laboratory-scale configurations can cost several hundred to several thousand hours of CPU time to simulate and several thousand runs may be needed to converge meaningful statistics. Surrogate-based approaches are more suited to this kind of problem because a response surface is built using relatively few observations, corresponding to equally few deterministic simulation runs. This response surface is then resampled to collect statistics. This approach has been implemented by Sheen and co-workers [99, 98, 119] or more recently Yousefian et al. [124] for simple combustion problems. However, surrogate models still require many training samples when the

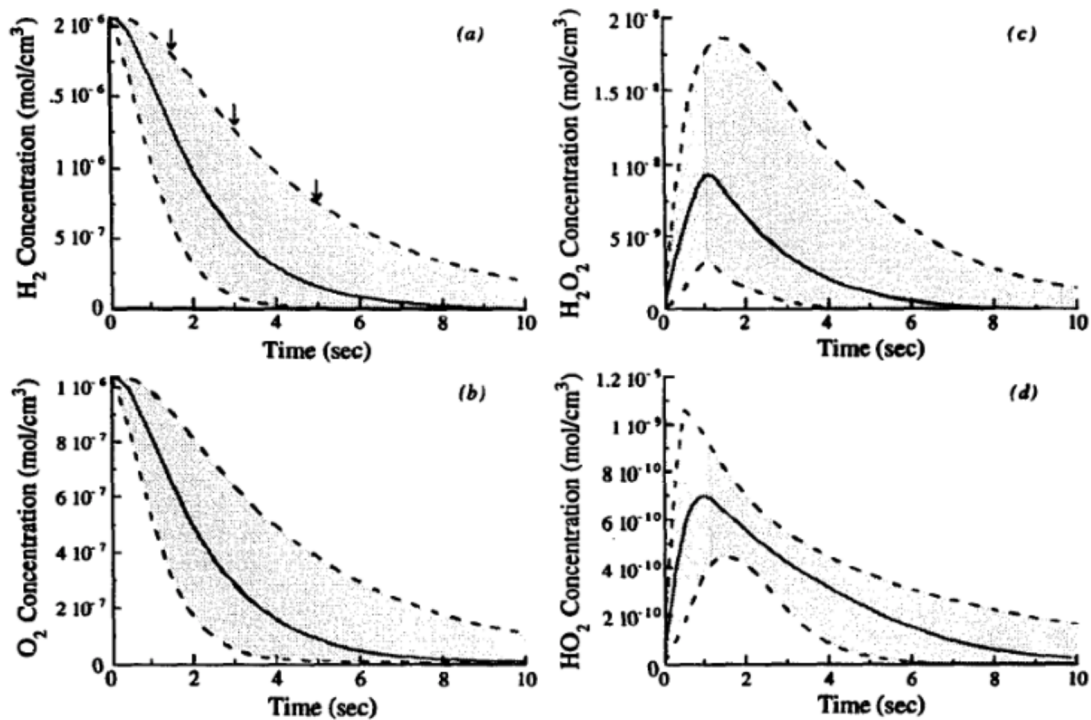


Figure 5: Temporal evolution of the probability distribution of species concentration in the plug-flow reactor, using 15000 Monte-Carlo samples. Solid lines represent median values and dashed values encompass the 95% confidence interval. Adapted from Phenix et al. [81]

uncertain dimension is high, which can be problematic for costlier numerical setups.

Uncertainty quantification in turbulent combustion simulations

More recent works have propagated uncertainties through turbulent combustion simulations. Notably Mueller and co-workers [72] used a two-step uncertainty propagation to build an uncertain tabulation for the chemistry. This allowed them to deal with only a reduced set of dependent random variables to propagate through the Large Eddy Simulation (LES) of the Sandia flame. A few years later, Mueller and Raman [73] propagated model form and parameter of the mixture fraction variance dissipation model in that same simulation. Figure 6 shows the distribution of time-averaged temperature and CO mass fraction, estimated using a collocation method, along several radial profiles. Again, notable uncertainties are found by these authors on the quantities of interest, especially on the most downstream profile.

Concomitantly, Khalil et al. [52] propagated the uncertainties of the turbulent model (Smagorinsky coefficient and turbulent Prandtl and Schmidt numbers) through the LES of a bluff-body-stabilized flame. They did not need to reduce the uncertain dimension as only three random variables were considered, and uncertain chemistry was disregarded. A similar study was performed by Lucor et al. [65] in the context of non-reactive LES.

All these uncertain turbulent flame studies cited above were performed with either

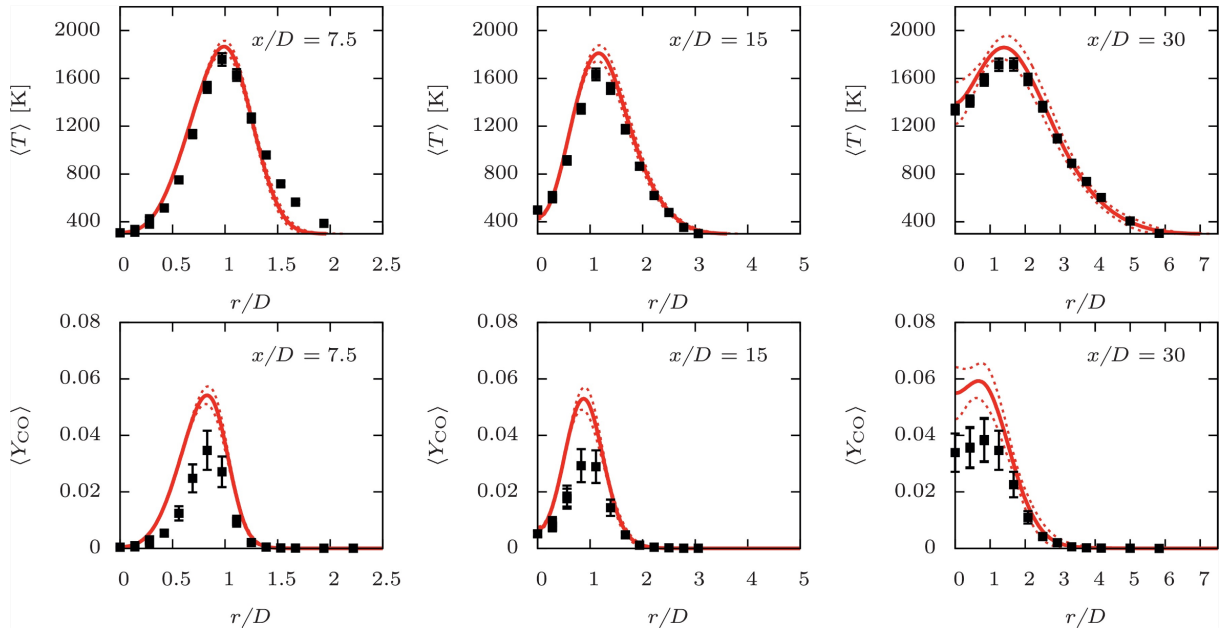


Figure 6: Radial profiles of time-averaged Temperature and CO mass fraction. Solid lines are the LES mean, dashed lines are $\pm 3\sigma$, and symbols are experimental measurements. Adapted from Mueller and Raman [73]

a limited number of uncertain variables [65, 52, 73] or an a priori dimension reduction [72]. In a radically different approach, other studies made an a posteriori dimension reduction of the uncertain space using the concept of "active subspace" [17]. In this framework, the physical model is first sampled randomly on the entire uncertain space, and the observations are used to infer a smaller uncertain space on which the surrogate is built. Constantine and co-workers [18] used it to quantify the uncertainty of the integrated exit pressure of a scramjet by propagating experimental and boundary conditions uncertainties. Ji et al. [45] used an active subspace to estimate the distribution of the H_2 Cabra flame lift-off height considering only the kinetic mechanism to be uncertain. This approach is only affordable for numerical setups that can be run several times the number of uncertain dimensions, i.e. it is limited to highly symmetric setups or RANS simulations and/or to cases where the uncertain dimension is low.

Objective of the thesis

The literature for uncertainty quantification in combustion contains few examples of CFD, and even fewer examples of Large Eddy simulations. LES is more precise than the Reynolds-Averaged Navier-Stokes approach, and it is also much more expensive due to its spatial and temporal resolution. However, LES is not immune to uncertainties, in particular when the physical configuration features high sensitivities to certain parameters.

This thesis aims to expand on the studies cited above and synthesize high-fidelity simulations and high-dimensional uncertainties. In other words, we aim to propagate high-dimensional uncertainties through expensive computer simulations, that cost several tens of thousand to several million of CPU hours for each run. In particular, we aim to propagate kinetic, turbulence model, and measurement uncertainties in the Large Eddy Simulation of a turbulent flame. We also choose to work on a flame stabilized by auto-

ignition because their lift-off height is first-order-determined by the chemical kinetics of the mixture. Hence, our preferred choice of configuration was the Cabra flame [11], which fulfills our criteria and has already been extensively studied, both numerically and experimentally. Furthermore, we choose to work on the H_2 variant of the Cabra flame because the contribution of auto-ignition is clearer in this flame than in the CH_4 variant [11, 22, 120, 33, 123] and because uncertain H_2/O_2 combustion mechanisms are readily available.

In short: **We aim to propagate several sources of uncertainty through an LES of the H_2 Cabra flame. We will take into account the co-flow temperature (uncertain experimental measurement), the detailed kinetic mechanism of H_2 oxidation, and a closure coefficient of the turbulent combustion model.**

Structure of the manuscript

- **Chapter 1** presents a set of uncertainty quantification methods that will be useful in the following: random sampling, surrogate modeling, sensitivity analysis, dimension reduction. We hope that this chapter may serve as a point of entry to uncertainty quantification for future researchers. This chapter also summarizes the methods implemented in the in-house code SURE during this thesis.
- **Chapter 2** is a preliminary study on auto-igniting homogeneous reactors and 1D laminar flames. Its purpose is to highlight the uncertainty of the auto-ignition delay in typical combustion configurations. It also serves as a first benchmark of the methods presented in Chapter 1.
- **Chapter 3** presents this thesis' main case of study: the Cabra flame and our LES of this flame. It details the key equations and the choice of turbulence and combustion models which rely on tabulated chemistry and presumed filtered-density function, which were implemented in the AVBP LES solver. Finally, it presents the deterministic simulation results and their validation.
- **Chapter 4** studies the uncertainty of the lift-off height of the simulated Cabra flame. Using a reduction of the initial uncertain space, a surrogate-based approach is leveraged to estimate the probability distribution of the lift-off height. 39 LES were run to train the surrogate and 20 more were run to serve as validation samples. This study allows the synthesis between seemingly disagreeing experimental observations and simulation results from several research groups.
- **Chapter 5** extends the methodology introduced in Chapter 4 in order to quantify the uncertainties of every averaged LES field at every location in the computational domain. These local uncertainties are obtained by applying Principal Component Analysis to a training set of averaged flames and building surrogate models to estimate the weights of the most important modes. From these local field predictions, the uncertainty of global quantities, like the lift-off height, can be rediscovered.

Chapter 1

Handling uncertainties

This chapter gives an overview of the methods of the uncertainty quantification (UQ) methods employed during this thesis. It provides the necessary elements to understand and implement these methods, as well as references to further investigate the points of detail. As such, it is written in the form of a beginner's guide towards uncertainty propagation, in the hope that it may provide a helpful entry point for future Ph.D. students.

The framework of this thesis is forward uncertainty propagation, which means that we have information on the uncertainty of uncertain inputs, and we seek to determine statistics (probability distribution, sensitivity, or at least expectation and variance) of a Quantity of Interest (QoI).

Unless specified otherwise, the methods presented in this chapter have been implemented in the in-house UQ code SURE.

This chapter begins with a quick overview of the different classes of methods that will be developed later. The second part focuses on how to sample an uncertain space. The third part is about surrogate modelling, or how to emulate a physical model given only a few observations. The fourth part deals with sensitivity analysis. Part 5 presents a few ways to reduce the uncertain dimension. Finally, part 6 introduces the error metrics used in this thesis to verify the quality of the methods presented in this chapter.

1.1 Methods overview

The methods employed in this thesis have been hand-picked to answer the specific needs of propagating kinetic and experimental uncertainties in a Large Eddy Simulation (LES):

- The simulation software are already coded with deterministic applications in mind, which means going for non-intrusive methods only.
- One observation in LES costs at least several tens of thousands of CPU-hours. Therefore, sample-efficient methods are wished for.
- The dimensionality of kinetic is high (cf Chapter 2), which means we need methods able to handle several tens of uncertain inputs, and/or ways to reduce the uncertain dimension

1.1.1 Deterministic methods

1.1.1.1 Linearization technique

A local gradient of the model is evaluated at the nominal input point. This allows performing a local sensitivity analysis by comparing the components of the gradient. Furthermore, if the input distributions are peaked enough, one can propagate uncertainties through a local linear surrogate of the model and retrieve an approximate output distribution, either by resampling or by analytical developments.

On the one hand, this method is quite cheap, because the minimum number of evaluations of the model is $d + 1$ in a non-intrusive framework. Even better, when an adjoint solver is used, gradients are already part of the simulation results. For "simple" analytical models, the gradient can even be analytically retrieved. On the other hand, this analysis is local. It is not valid if the input distributions are wide enough that the uncertain inputs frequently exit the neighborhood of the point of gradient evaluation.

1.1.1.2 Perturbation techniques and Neumann expansions

These techniques obey the same idea as the linearization technique but function with higher-order derivatives, making them more precise. Still, they remain local.

1.1.1.3 Moments methods

The idea behind these methods is not to solve the model for a given set of inputs, but rather to solve a system of equations on the moments of the inputs to retrieve moments of the output. These methods suffer from a problem of closure since the equations on the moments are often not closed and closure terms have to be modeled. Furthermore, they enter the domain of intrusive methods and do not allow the use of a legacy deterministic solver. Therefore, they are out of the scope of this thesis.

1.1.2 Simulation methods

Simulation methods rely on the evaluation of the stochastic process on a large number of samples, either directly or indirectly, to approximate the Probability Distribution Function (PDF) of the uncertain output.

1.1.2.1 Sampling techniques

Sampling techniques, as their name suggests, rely on sampling extensively the uncertain input space according to its joint probability distribution, evaluating the physical model, and storing the quantity of interest for each sample. The probability distribution and relevant statistics of the QoI can then be retrieved from the data set. These methods are very robust and do not require any particular property from the physical model (non-continuous outputs are fine, for example), and they allow the use of legacy/deterministic solvers. However, they suffer from slow convergence rates, meaning that a considerable number of observations may be necessary to retrieve significant statistics. This is not a problem for cheap models featuring simple or reduced physics, but it will become an issue for more complex and expensive models.

1.1.2.2 Surrogate modeling

Surrogate modeling alleviates the cost of sampling techniques by building a surrogate, also called "response surface" or "emulator", using a few observations of the physical

model. The surrogate is in turn extensively sampled instead of the physical model. As such, it is way cheaper to implement than a naive sampling technique, when the evaluation of the physical model is expensive. This cost may lay, for example, in the monetary costs of an experiment or in the computational power required to run a simulation. The surrogate, by construction, is very cheap to evaluate, which makes resampling affordable whatever the number of samples needed to obtain significant statistics. However, the precision of the resulting PDF relies entirely on the accuracy of the surrogate.

1.2 Random sampling

1.2.1 Monte-Carlo sampling




Figure 1.1: 800 points of a Monte-Carlo sampling of a 2D uniform random vector on the unit hypercube

Monte-Carlo simply consists in randomly sampling the input space and computing the corresponding outputs. An example of Monte-Carlo sampling on the unit hypercube, in 2D, is given in Figure 1.1. Let $F(\mathcal{S})$ be our quantity of interest, where \mathcal{S} is the output of the physical model. The function F is defined depending on what you need to compute. For example, if you want to evaluate:

- the expectation of the output: $F(\mathcal{S}) = \mathcal{S}$
- the second moment of the output: $F(\mathcal{S}) = \mathcal{S}^2$
- the probability of the value exceeding a given threshold:

$$F(\mathcal{S}) = \begin{cases} 1 & \text{if } \mathcal{S} > \text{Threshold} \\ 0 & \text{otherwise} \end{cases}$$

Convergence of a Monte-Carlo estimator

$F(\mathcal{S})$ is a function of the simulation \mathcal{S} which is itself a deterministic function of the uncertain inputs. Therefore, F is a function of the uncertain inputs ξ and we will denote

indifferently $F(\mathcal{S}) = F(\mathcal{S}(\boldsymbol{\xi})) = F(\boldsymbol{\xi})$ in the following.

The random vector $\boldsymbol{\xi}$ is defined on the set Ξ of joint probability $\pi_{\boldsymbol{\xi}}$.

Let \widehat{F}_N be the Monte-Carlo estimator of $\mathbb{E}[F]$ defined by:

$$\widehat{F}_N = \frac{1}{N} \sum_{i=1}^N F(\mathcal{S}(\boldsymbol{\xi}^{(i)})) \quad (1.1)$$

Where N is the number of input samples $\boldsymbol{\xi}^{(i)}$ drawn according to the probability distribution $\pi_{\boldsymbol{\xi}}$.

The central limit theorem gives, in the asymptotic regime, the following property:

$$\widehat{F}_N = \mathbb{E}[F] + \varepsilon_N, \quad \text{with} \quad \varepsilon_N \sim \mathcal{N}\left(0, \frac{\mathbb{V}[F]}{N}\right) \quad (1.2)$$

Where $\mathbb{V}[F]$ is the variance of F . Substituting the estimator \widehat{V}_N for $\mathbb{V}[F]$, we can show that once in the asymptotic regime, a Monte-Carlo estimator has a universal convergence rate of $1/2$, meaning that the slope of the upper bound for 50% confidence on the residual ε_N has a $-1/2$ slope in a log-log diagram, as shown in Figure 1.2. We can also say that it converges in $O(1/\sqrt{N})$.

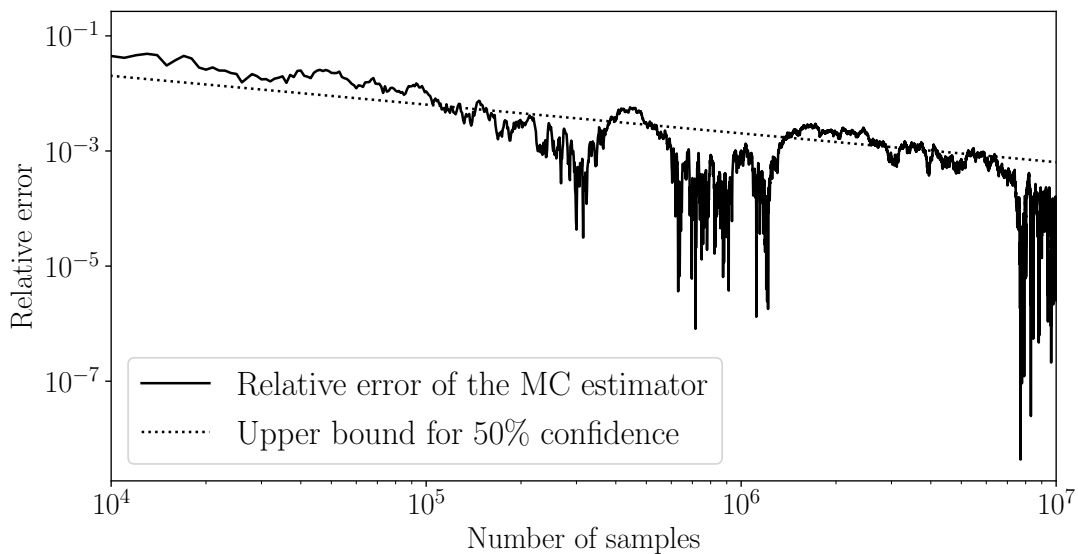


Figure 1.2: Illustration of the convergence of a Monte-Carlo estimator

As highlighted by eq. 1.2, the Monte-Carlo estimator's precision depends exclusively on the number of samples drawn and on the variance of the quantity of interest. In particular, it is independent of:

- The dimensionality of the sampling space
- The regularity of the model
- The oscillarity of the model (For example, $\mathbb{V}[\sin(kx)]$ does not depend on k)

When trying to converge a Monte-Carlo estimate, it is good practice to make sure you at least double the number of samples when the error estimator is too high. This is because \widehat{F}_N is random and fluctuates. Thus, adding just a few more samples will not

necessarily help with the convergence and may actually worsen the estimator (see figure 1.2). Finally, it is important to note that the error estimate ε_N is an absolute value. Only the relative error $\varepsilon_N/\widehat{F}_N$ should be used to characterize the estimator's level of convergence.

1.2.2 Sampling randomly from an arbitrary distribution

Random numbers generators (or rather pseudo-random generators [59]) are typically only capable of sampling random variables X that follow a uniform distribution over the segment $[0, 1]$ ($X \sim U([0, 1])$). However, uncertain variables that do not follow a $U([0, 1])$ distribution are plentiful. Thus, we need a way to sample from an arbitrary PDF.

Two cases can be identified:

- When we know (or can approximate) its cumulative distribution function (CDF) and we can inverse it
- When we only have access to the value of the probability density at any value taken by the variable

These two cases can be treated respectively by a measure change and Markov-chain Monte-Carlo sampling. These methods will be detailed in the following sections.

1.2.2.1 Measure change

The measure change is an analytical transformation that generates samples of a random variable using its CDF.

Let ξ be a one-dimensional random variable of cumulative distribution function F_ξ . If F_ξ is strictly monotonic, it is invertible and the random variable $F_\xi^{-1}(u)$, where u follows a uniform distribution on $[0, 1]$, follows the distribution of ξ . This property is illustrated in Figure 1.3. In short:

$$u \sim U([0, 1]) \iff F_\xi^{-1}(u) \sim \xi \quad (1.3)$$

F_ξ^{-1} is called the quantile function of ξ . Some examples are given here for usual distributions:

- Normal distribution:

$$F_\xi^{-1}(u) = \mu + \sigma\sqrt{2}\operatorname{erf}^{-1}(2u - 1) \quad (1.4)$$

Where μ is the expectation of the distribution and σ its standard deviation

- Log-normal distribution:

$$F_\xi^{-1}(u) = \exp\left(\mu + \sigma\sqrt{2}\operatorname{erf}^{-1}(2u - 1)\right) \quad (1.5)$$

Where μ and σ are defined by the associated normal distribution

It follows that, for a given quantity of interest Q ,

$$\mathbb{E}[Q] = \int_{\Xi} Q(\xi)\pi(\xi)d\xi = \int_0^1 Q(F_\xi^{-1}(u))\pi(u)du = \int_0^1 Q \circ F_\xi^{-1}(u)du \quad (1.6)$$

A Monte-Carlo estimator of Q will then be:

$$\mathbb{E}[Q] \approx \widehat{Q}_N = \frac{1}{N} \sum_{i=1}^N Q \circ F_{\xi}^{-1}(u^{(i)}) \quad (1.7)$$

Or more generally, for any weighted estimator:

$$\mathbb{E}[Q] \approx \widetilde{Q}_N = \sum_{i=1}^N w^{(i)} Q \circ F_{\xi}^{-1}(u^{(i)}) \quad (1.8)$$

where $u^{(i)}$ are samples of the random variable u and $w^{(i)}$ the weights associated with the integration point $u^{(i)}$.

This is an artificial way of sampling from the distribution of ξ because what is ultimately sampled is still $u \sim U([0, 1])$. However, it allows to efficiently evaluate an expectation estimator, and it is compatible with every Monte-Carlo convergence acceleration technique presented in the following.




Figure 1.3: Illustration of the measure change method for a reduced and centered normal distribution

In the case of a random vector made of a number of independent random variables verifying the conditions for the measure change, the method can be generalized and is then called the Rosenblatt inverse transform [89].

This procedure works only in the case where the probability distribution is simple/canon enough so that a quantile function, or at least a good approximation of it, is available. In the multi-dimensional case, it works only if the components of the random vector are independent.

1.2.2.2 Markov-Chain Monte-Carlo

Markov Chain Monte-Carlo (MCMC) [37] is a method that allows drawing samples directly from an arbitrary distribution, of which we only need to be able to evaluate the

probability distribution at any point. This means we do not need an analytical form for the Probability Distribution Function (PDF), Cumulative Distribution Function (CDF), or quantile function. MCMC is, therefore, the default choice for sampling non-standard distributions or random vectors whose components are not independent.

MCMC is based on the concept of Markovian processes, or Markov chains, which are memory-less random walks. In practice, the Metropolis-Hastings algorithm [37] and Gibbs sampling [31] are widely used, but other, more advanced algorithms of Markov-Chain Quasi Monte-Carlo [60] can be used to combine the ability of sampling from any distribution with the advantages of Quasi-Monte-Carlo (presented in Section 1.2.3.2).

1.2.3 Improving the convergence of the Monte-Carlo estimator

Monte-Carlo estimators do not converge quickly. For this reason, it is common to try and accelerate this convergence, which means minimizing the residual ε_N for a given computational effort. To achieve this you can either:

- Decrease the variance of the QoI (Linked to the intercept of the convergence curve)
- Increase the rate of convergence (The slope of the convergence curve)

1.2.3.1 Reducing the variance of the quantity of interest

Separation of the principal part

The idea here is to find a "principal part" H of F , that is, a function with known expectation $\mathbb{E}[H]$ that approximates the behavior of F . We then define the error Θ as: $\Theta(\xi) = F(\xi) - H(\xi)$. The expectation being a linear operator, we thus have: $\mathbb{E}[F] = \mathbb{E}[H] + \mathbb{E}[\Theta]$. We can then define a new estimator:

$$\widetilde{F}_N = \widehat{\Theta}_N + \mathbb{E}[H] \quad (1.9)$$

This new estimator converges as $\frac{\sigma(\Theta)}{\sqrt{N}} = \frac{\sigma(F - H)}{\sqrt{N}} \ll \frac{\sigma(F)}{\sqrt{N}}$

This idea is used in:

- Surrogate Variance Reduction / Reduced Model Variance Reduction (exactly the method presented)
- Multifidelity (degrade physics) and Multilevel (degrade discretization) Monte-Carlo methods.

The separation of the principal part is illustrated in Figure 1.4.

In practice, it can be simply implemented by spending a few samples to get the rough shape of F and construct a surrogate (for example a linear fit) to get a suitable H which is easily samplable/integrable. The better the surrogate, the faster the convergence.

Importance sampling

The idea of Importance sampling is to draw more samples in the areas where F has large values. That way, the integral computation is more efficient. This is done by using a variable change and an alternative distribution π^* to draw the samples.




Figure 1.4: Illustration of the separation of the principal part

$$\begin{aligned}\mathbb{E}_\pi[F] &= \int_{\Xi} F(\xi)\pi(\xi)d\xi = \int_{\Xi} F(\xi)\frac{\pi(\xi)}{\pi^*(\xi)}\pi^*(\xi)d\xi = \int_{\Xi} G(\xi)\pi^*(\xi)d\xi \\ &= \mathbb{E}_{\pi^*}[G]\end{aligned}\quad (1.10)$$

Where $G(\xi) = F(\xi)\frac{\pi(\xi)}{\pi^*(\xi)}$. Of course, π^* Must be non-zero on Ξ . The only difficulty resides in finding a probability distribution π^* such that $\mathbb{V}_{\pi^*}[G] < \mathbb{V}_\pi[F]$.

Analytically, the π^* that minimizes $\mathbb{V}_{\pi^*}[G]$ is defined as:

$$\pi^*(\xi) = \frac{|F(\xi)|\pi(\xi)}{\left|\int_{\Xi} F(\eta)\pi(\eta)d\eta\right|}\quad (1.11)$$

The denominator of this expression is not known. Indeed, it is the absolute value of $\mathbb{E}_\pi[F]$, the final result of the study. However, equation 1.11 states that any distribution proportionnal to $|F(\xi)|\pi(\xi)$ will be suitable. Fortunately, the Metropolis-Hastings algorithm, mentioned in Section 1.2.2.2, allows sampling from any arbitrary distribution known within a multiplication factor.

1.2.3.2 Increasing the convergence rate: Low discrepancy sampling

In this section, we will restrict ourselves to the sampling of uniform random variables on the hypercube $[0, 1]^d$, without loss of generality.

Latin Hypercube Sampling

The idea behind Latin Hypercube Sampling (LHS) [71] is to control the sampling so that it covers the input space more evenly than Monte-Carlo, while keeping it random. In 1D, the segment $[0, 1]$ is partitioned into N non-overlapping intervals having the same probability (equal lengths for uniform distributions). In d-D, the space is partitioned into N rows in each dimension, and the algorithm draws one and only one sample per row in each dimension. This is illustrated in Figure 1.5.




Figure 1.5: LHS sampling of a 2D uniform random vector on the unit hypercube

In practice, what can be done is an LHS is 1D, then one random permutation of this initial sampling is done for each remaining dimension. For example, to draw 4 LHS points in the uniform unit cube (3D), we first draw 4 LHS points in the segment $[0, 1]$: $\{0.2, 0.37, 0.51, 0.82\}$. We then perform two random permutations of this list (two remaining dimensions), which gives us: $\{0.37, 0.51, 0.2, 0.82\}$ and $\{0.82, 0.37, 0.2, 0.51\}$. Then, we read the three coordinates of the first 3D LHS point as the first elements of the three lists: $(0.2, 0.37, 0.82)$. The second point is $(0.37, 0.51, 0.37)$, and so on.

This ensures better coverage of the uncertain space than brute-force Monte-Carlo. In terms of even coverage, an improvement on this method is orthogonal sampling [108]. In orthogonal sampling, the input space is first divided into N equiprobable subspaces, and one sample is randomly drawn from each subspace in such a way that the total set of samples follows the rules of LHS.

Mc Kay et al. [71] showed that LHS could not have a worse convergence rate than brute-force Monte-Carlo, and that its convergence is $O(1/N)$ if $F(\boldsymbol{\xi})$ is monotonous in each of its arguments.

However, LHS presents one major issue. A fixed number of samples is decided a priori. In the case where the convergence of the quantity of interest's estimator is not satisfactory, increasing the number of samples while retaining the properties of LHS is not trivial. Tong [111] gives useful pointers in that regard.

The implementation of LHS used in this thesis is the one proposed by the OpenTURNS library [6].

Quasi Monte-Carlo

Quasi-Monte-Carlo sampling (QMC) makes use of deterministic space-filling (or "low-discrepancy") sequences. They are designed to fill the space as uniformly as possible but lose in the process the randomness of standard Monte-Carlo sampling.

Their advantage is that the sampling is sequential, and the sequences are not periodic, so samples can always easily be added to the sample set if the estimator is not converged enough. This is a very interesting property compared to LHS.

The easy refinement of the sampling comes at a cost, though, and this cost is the

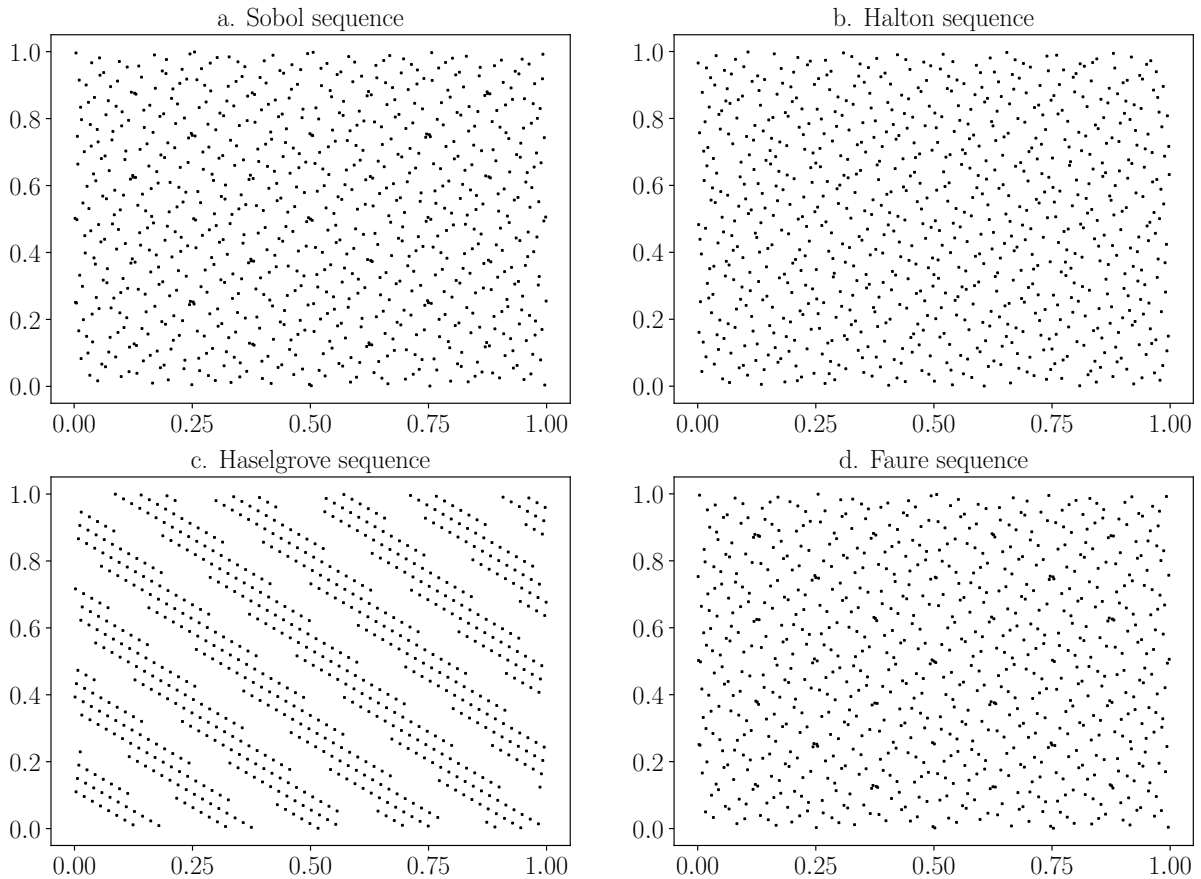


Figure 1.6: Examples of the 800 first points of low-discrepancy sequences of a 2D uniform random vector on the unit hypercube

convergence rate. The best possible convergence rate for a Quasi Monte-Carlo method is believed to be $O(\log(N)^d)/N$ (proven for $d = 2$, but still a conjecture for $d \geq 3$ [61]). This is asymptotically always better than a brute-force Monte-Carlo estimator, but it may not be as good as LHS depending on the case. However, they are impaired by some important drawbacks:

- The sampling sports patterns. If the model is oscillatory, some features may be difficult to retrieve
- The efficiency of QMC tends to decrease with the increase of dimensionality, while the convergence rate of Monte-Carlo and LHS is independent of the dimension. In large dimensions, some parts of the sampling space will not be covered at all if the number of samples is not sufficient, because the exploration of the space is stratified. An optimal number of samples, for example $N = 2^n$ for the Sobol' sequence, is required to maximize the uniformity of sampling.
- As QMC is purely deterministic, Monte-Carlo-style error estimates cannot be computed for this kind of sampling, because they rely on the intrinsic randomness of the sampling and the Central limit theorem. This drawback can be alleviated by the use of Randomized Quasi-Monte-Carlo sampling, which is presented in the next section.

The implementation of QMC used in this thesis is the one proposed by the OpenTURNS library [6].

Randomized Quasi-Monte-Carlo

The randomization of QMC allows the determination of the estimator's error in the same way as Monte-Carlo [113]. The randomization is not a trivial task if the space-filling properties of QMC are to be kept. Dumont's thesis [24] provides an interesting review of the different randomization methods.

1.3 Surrogate modeling

Surrogate modeling is the science of reducing the cost of sampling expensive functions by building "surrogates", also called "emulators" or "response surfaces". These surrogates are functions that should be cheap to evaluate and that give a good enough approximation of their underlying "true" physical model.

At this point, you might be wondering one of the following three things:

1. How do I recognize a "good enough" approximation of my function?
2. How cheap is "cheap to evaluate"?
3. Theory is all very well, but where do I start in practice?

First, a "good enough" approximation depends on the problem you want to tackle. It is an approximation that allows you to retrieve the quantities that you seek with sufficient accuracy. In practice, you might find that it relates to recovering the probability distribution's shape and moments of the QoI you want to emulate, but watch out, the level of accuracy is highly case-dependent. This is best illustrated with two - perhaps caricatural - scenarii:

Scenario 1: Global performance analysis

Say you are designing a brand new run-of-the-river hydroelectricity plant to power a city in a remote area. It happens that the only available river is capricious, meaning that its flow rate is far from constant. This flow rate does not appear to be predictable for now, but you can model it as a random variable of a given probability distribution based on previous observations. You then aim to propagate this randomness - this uncertainty - through a model of the plant to assess the uncertainty of your power output.

Alas, the decision has been made to design a field of immersed stream generators to reduce environmental impact. The precise computation of the power output of such a field requires advanced - and expensive - CFD simulations due to the interactions between the stream generators.

Thus, you need a surrogate model of the CFD simulation to obtain meaningful statistics for your power output at a reasonable cost.

Now, what are you really trying to figure out with this study? The expectation of the power output, to be sure, and probably its standard deviation to design adequate buffers or annex generators.

This means that while a very coarse surrogate only allows us to capture the first two moments of the power output distribution, it should be enough to suit your needs.

Scenario 2: Rare events prediction

Still in the energy business, you are now working in a nuclear power plant as a safety engineer. Your goal is now to assess the probability of the plant blowing up due to, for instance, excessive pressure in the primary coolant circuit.

Let's assume that the QoI is a function of the mechanical strength of the circuit and of the stress applied by the coolant's pressure. Uncertain inputs, on the other hand, refer to the composition, crystal structure and mechanical fatigue of the pipes, coolant temperature and pressure, and so on. Coupled simulations where heat transfer, fluid mechanics, and solid mechanics interact are very expensive. Therefore, you need a surrogate again to propagate uncertainties in this problem. Now, a nuclear plant blowing up is - hopefully - a rare event, so the events we seek to predict are not in the belly of a bell curve, but rather on the tails of the distribution - regions of low probability. Thus, the surrogate must be accurate enough so that the shape of the probability distributions, and in particular the tails, are well predicted.

Second, as a rule of thumb, a cheap function relies on an analytical formula, even a complicated one, ranging from simple piece-wise constant expressions to complex neural networks. The key here is that no complex iterative process should be involved in the evaluation of the surrogate, as these are generally expensive.

Last but not least, you can start building surrogates by using the two relatively common surrogate types used in this thesis. In the following, you will find a brief presentation of these surrogate models and how to build them.

1.3.1 Polynomial Chaos Expansion

Despite its daunting denomination, a Polynomial Chaos Expansion (PCE) is nothing more than a polynomial approximation of a stochastic process, which is itself a function of random variables. It must be noted that the random input variables must be independent to build a valid PCE.

Polynomials are very practical functions to approximate continuous physical phenomena. It is in fact possible to demonstrate that PCEs have a spectral convergence towards any continuous stochastic process [66]. That is, the more polynomials in the expansion, the better the approximation of the process, provided that the weights are exactly computed.

Formally, a PCE is defined as follows:

Let $(F : \boldsymbol{\xi} \in \Xi) \in \mathbb{R}^d \rightarrow \mathbb{R}$ be the physical model where $\boldsymbol{\xi} = (\xi_1, \dots, \xi_d)$ are independent, real-valued, random variables of joint-PDF $\pi_{\boldsymbol{\xi}}$.

Let $\alpha = \{\alpha_i\} \in \mathbb{N}^d$ be a multi-index and $\{\psi_{\alpha}\}$ be the set of d -variate orthogonal polynomials:

$$\psi_{\alpha}(\boldsymbol{\xi}) = \prod_{i=1}^d \psi_{\alpha_i}^{(i)}(\xi_i) \quad (1.12)$$

If $F \in \mathcal{L}_2(\Xi, \pi_{\boldsymbol{\xi}})$, its PCE exists and is written:

$$F(\boldsymbol{\xi}) = \lim_{N \rightarrow \infty} \sum_{\|\alpha\| \leq N} \psi_{\alpha}(\boldsymbol{\xi}) F_{\alpha}, \quad \|\alpha\| = \sum_{i=1}^d |\alpha_i| \quad (1.13)$$

In essence, a Polynomial Chaos Expansion is therefore the combination of a polynomial basis $\{\psi_\alpha\}$ and a set of weights $\{F_\alpha\}$. Here is how it shakes down:

1.3.1.1 The basis

Building the polynomial basis

The polynomial basis must be an orthogonal family of polynomials. Being orthogonal means that the dot product between two basis elements is zero. You may remember that the dot product between the polynomials ψ_1 and ψ_2 is written as follows:

$$\langle \psi_1, \psi_2 \rangle = \int_{\Omega} \psi_1(\boldsymbol{\xi}) \cdot \psi_2(\boldsymbol{\xi}) \cdot \pi(\boldsymbol{\xi}) d\boldsymbol{\xi} \quad (1.14)$$

The nullity of this expression obviously depends on $\pi(\boldsymbol{\xi})$, the probability distribution of the random vector $\boldsymbol{\xi}$. Consequently, a different polynomial basis must be built for each distribution.

Case 1: Canonical distributions

For canonical distributions, canonical polynomial bases are available (see Table 1.1). These bases can be tensorized for multivariate PCEs.

In the case where two random input variables are canonical but do not have the same probability distribution, one can (and must) compose the bases.

For example, if ξ_1 is normal and ξ_2 is uniform, the polynomial basis will be the tensorization of Hermite and Legendre polynomials.

	distribution ξ	Polynomials $\psi_k(\xi)$	Support
Continuous RV	Gaussian	Hermite	\mathbb{R}
	γ	Laguerre	\mathbb{R}^+
	β	Jacobi	$[a, b]$
	Uniform	Legendre	$[a, b]$
Discrete RV	Poisson	Charlier	\mathbb{N}
	Binomial	Krawtchouk	$\llbracket 0, n \rrbracket$
	Negative binomial	Meixner	\mathbb{N}
	Hypergeometric	Hahn	$\llbracket 0, n \rrbracket$

Table 1.1: Canonical distributions and associated polynomial bases

Case 2: Arbitrary distributions

In the case of arbitrary distributions, a custom polynomial basis can be built iteratively using the Gram-Schmidt orthonormalization process.

Truncating the polynomial basis

The polynomial bases constructed previously are infinite families of polynomials. The complete expansion - i.e. the infinite sum of correctly-weighted basis polynomials - is a perfect representation of the stochastic process.

However, in practice, you must truncate the basis. This means keeping a certain number of relevant polynomials while discarding all the other ones. The most common ways of truncating the basis are:

1. Total order truncation
2. Hyperbolic truncation
3. Adaptive basis construction

Total order truncation

This is the simplest of the truncation methods. All polynomials in the basis which are under a given total order are retained. This allows describing high-order interactions. However, this comes at the cost of rapidly increasing the cardinal of the basis. The number of basis polynomials M for a PCE of total order $p+1$ involving d uncertain variables is given by Equation 1.15.

$$M = \frac{(d+p)!}{d!p!} \tag{1.15}$$

As illustrated in Figure 1.7, the number of polynomials (and therefore the cost of finding the weights) reaches quickly unreasonable amounts for uncertain spaces of high dimensionality. This is called the "Curse of Dimensionality". Total-order truncation is therefore affordable only when the dimensionality of the problem is not too high.




Figure 1.7: Evolution of the number of basis polynomials M with the number of uncertain variables d and the total order of the basis p

Hyperbolic truncation

According to the "sparsity-of-effects principle", or "hierarchical ordering principle" [10], systems are usually governed by main effects (involving a single variable) and low-order interactions. Therefore, neglecting high-order interactions should not result in a significant loss of representativeness for the PCE, in the majority of cases. A systematic way to do so is to perform a so-called hyperbolic truncation, as introduced by Blatman and Sudret in [9].

A quasi-norm, so called q -quasi-norm $\|\cdot\|_q$ is defined as:

$$\forall 0 < q < 1, \|\alpha\|_q = \left(\sum_{i=1}^d \alpha_i \right)^{1/q} \quad \text{and} \quad \|\alpha\|_0 = \sum_{i=1}^d \mathbf{1}_{\{\alpha_i > 0\}} \quad (1.16)$$

The truncated PCE of total order p and q -norm is then defined as:

$$F(\boldsymbol{\xi}) \approx \tilde{F}_{p,q}(\boldsymbol{\xi}) = \sum_{\|\alpha\|_q \leq p} \psi_\alpha(\boldsymbol{\xi}) F_\alpha \quad (1.17)$$

The 1-quasi-norm is actually the absolute norm of \mathbb{N}^d and corresponds to the total order truncation. For $0 < q < 1$, the truncation follows a hyperbolic-like pattern illustrated in Figure 1.8, which cuts presumably-unneeded high-order interactions from the PCE. By weighting the different directions in the definition of the quasi-norm, you can also get an anisotropic hyperbolic truncation, which can help describe in more detail the evolution of the QoI in certain directions relative to others.

Adaptive basis

To further complexify the construction of the basis and retain only the useful terms of the polynomial basis, Blatman [8] developed adaptive methods to determine conjointly the optimal basis polynomial and weights of the PCE based on Least-angle Regression [25, 9].

1.3.1.2 The weights

There are several ways of computing the PCE weights, both intrusive and non-intrusive. In the following, we will discuss only the non-intrusive methods, which are projection and regression.

a. Projection

Because the $\{\psi_\alpha\}$ form an orthogonal basis of the Hilbertian space, the weights F_α can be found by simply projecting F on the basis:

$$F_\alpha = \frac{\langle F, \psi_\alpha \rangle}{\langle \psi_\alpha, \psi_\alpha \rangle} = \frac{1}{\langle \psi_\alpha, \psi_\alpha \rangle} \int_{\Xi} F(\mathbf{x}) \psi_\alpha(\mathbf{x}) \pi_{\boldsymbol{\xi}}(\mathbf{x}) d\mathbf{x} \quad (1.18)$$

The integral can then be approximated with a finite number of samples using a proper quadrature rule which is a proper numerical way to compute the integral:

$$F_\alpha \approx \frac{1}{\langle \psi_\alpha, \psi_\alpha \rangle} \sum_{i=1}^N w^{(i)} F(\boldsymbol{\xi}^{(i)}) \psi_\alpha(\boldsymbol{\xi}^{(i)}) \quad (1.19)$$

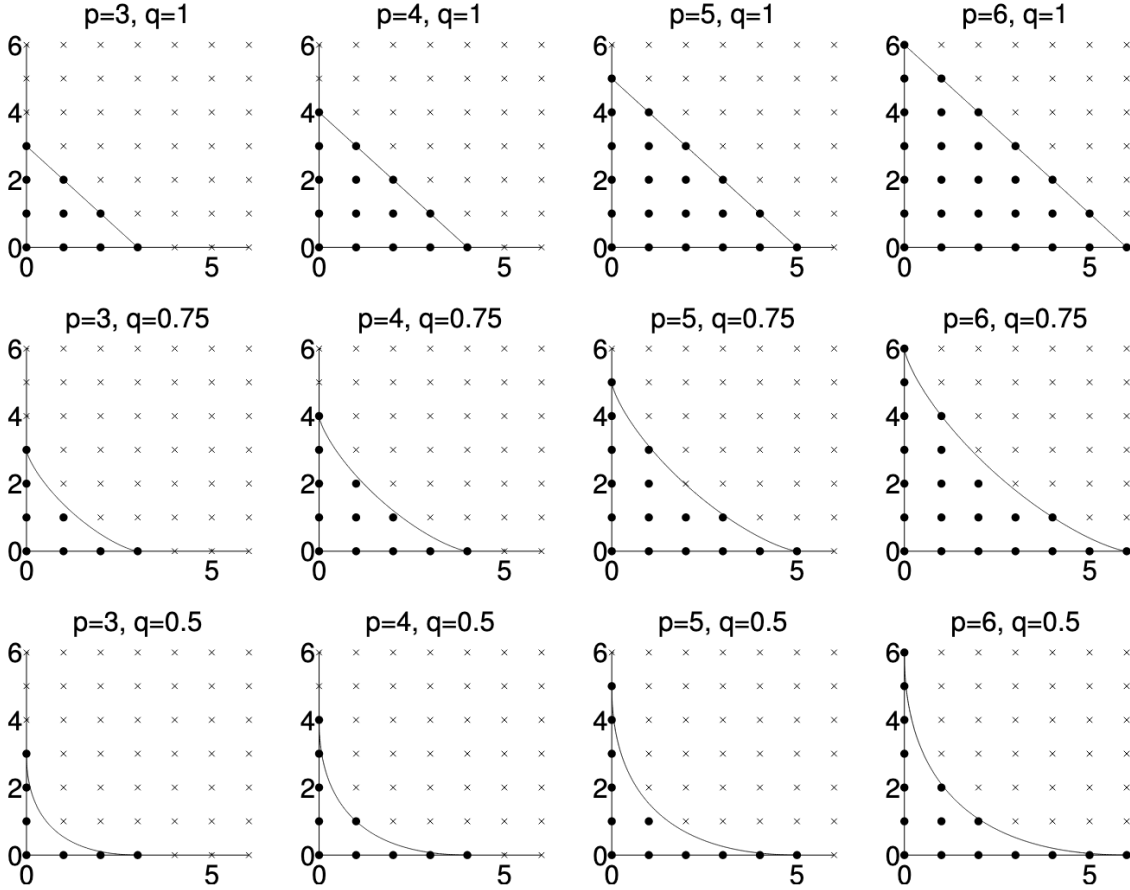


Figure 1.8: Illustration of the hyperbolic truncation [95]: Retained basis terms in the PCE when varying the parameter q and the total degree $p = 3, 4, 5, 6$ for $d = 2$ are highlighted by solid black points. The x and y axes correspond to the partial degree of the polynomial in ξ_1 and ξ_2 respectively.

The choice of the quadrature rule $(\xi^{(i)}, w^{(i)})$ is crucial as it conditions the number of samples needed to evaluate a given weight with a given precision. Most of these rules can be split into three main families:

- Brute-force integration
- Integration of polynomials that interpolate the integrand
- Subtler and more powerful change of variables-based approaches

An extensive discussion and benchmarks of numerous quadrature methods can be found in Dumont's thesis [24]. This section will only give a quick overview of these families, as well as the formal definition of the method primarily used during this thesis.

Brute force integration

This approach relies on Monte-Carlo estimates of the integral obtained using the methods described in Section 1.2. This approach is obviously not very efficient in this context because a large number of observations (i.e. a large training set) will be needed to compute precisely the integrals, which defeats the purpose of surrogate modeling.

Integration of polynomials that interpolate the integrand

This family of rules relies on the exact (analytical) integration of polynomials that interpolate the integrand. It comprises the Rectangle method (polynomials of order 0), the Trapezoidal rule (order 1), and the Newton-Cotes formulas (general case). Obviously, the closer the interpolating polynomial is to the integrand, the more precise the evaluation of the integral. However, high-order polynomials tend to oscillate out of hand and do not necessarily improve the precision of the estimation. Hence the popularity of the Composite Simpson rule, which is a special case of the Newton-Cotes rule, where the integrand is approached by a piecewise second-order polynomial on a regular grid.

Change of variable-based approaches

This family includes the well-known Gauss quadrature rules, which are designed to return the exact integral of a polynomial of order $2n + 1$ using only n quadrature points, i.e. n observations, which must be taken at very specific locations. The Gauss quadrature rules have the best convergence rate when the integrand is well approximated by polynomials. However, it is not possible to re-use observations when a more precise estimation of the integral is desired. For any refinement of the numerical integration, previous samples have to be put aside and entirely new observations have to be made. This is an important drawback because the number of quadrature points needed to get a "good enough" projection is not known a priori, in the general case. To tackle this problem, nested quadrature rules need to be used. "Nested" means that there are quadrature "levels" such that for a given level k , the quadrature level $k + 1$ reuses all the observations of level k , and a reduced number of additional samples are evaluated. This allows to refining the estimation of the integral by observing a reduced amount of new samples for each new quadrature level. This also allows sparse tensorization, which will be discussed in the following.

The Clenshaw-Curtis [14] rule and the Fejér quadrature of the second kind [27] are nested rules that feature a similar convergence rate as Gauss quadrature rules [112]. Clenshaw-Curtis is a closed quadrature rule, meaning that it includes the boundaries of the evaluation interval, while the second Fejér rule is an open rule, making it suitable to integrate a quantity over an open interval (typically \mathbb{R} for normal random variables).

The evaluation points and the associated weights are presented for the two methods, without loss of generality, for an integration over the uniform interval $[-1, 1]$. Clenshaw-Curtis quadrature points and weights are given respectively by Equations 1.20 and 1.21:

$$x_i = \cos\left(\frac{i\pi}{N}\right), \quad i \in \llbracket 0, N \rrbracket \quad (1.20)$$

$$\begin{cases} w_i = \frac{c_i}{N} \left(1 - \sum_{j=1}^{E[N/2]} \frac{b_j}{4j^2 - 1} \cos\left(\frac{2ji\pi}{N}\right) \right) \\ b_j = \begin{cases} 1, j = N/2 \\ 2, j < N/2 \end{cases}, c_i = \begin{cases} 1, i = 0 \text{ or } i = N \\ 2, \text{ else} \end{cases} \end{cases} \quad (1.21)$$

The Fejér quadrature of the second kind uses the same evaluation points, excluding the boundaries (-1 and 1), with different weights. They are given respectively in Equations 1.22 and 1.23.

$$x_i = \cos\left(\frac{i\pi}{N}\right), \quad i \in \llbracket 1, N-1 \rrbracket \quad (1.22)$$

$$w_i = \frac{4}{N} \sin\left(\frac{i\pi}{N}\right) \sum_{j=1}^{E[N/2]} \frac{\sin\left(\frac{(2j-1)i\pi}{N}\right)}{2j-1} \quad (1.23)$$

In the Clenshaw-Curtis rule, a quadrature of level k which uses $N_k = 2n + 1$ evaluations re-uses all the evaluations of the quadrature with $N_{k-1} = n + 1$ points. For the second Fejér rule of level k , the quadrature with $N_k = 2n - 1$ points re-uses the evaluation of the quadrature with $N_{k-1} = n - 1$ points.

For multi-dimensional integration, the quadrature rules need to be tensorized. Sometimes, the tensorized quadrature rule is called the "cubature rule". The full tensorization follows the following principles:

- For each quadrature point in a given dimension, the other dimensions must be sampled on each quadrature point. As a consequence using a quadrature rule which needs n evaluations in 1D, and tensorizing it in a space of dimension d requires making n^d evaluations. This exponential growth of the number of required evaluations is also a curse of dimensionality.
- The weights of a multi-dimensional quadrature rule are computed as the product of the 1D quadrature weights used in each dimension

To alleviate the curse of dimensionality in the integration, a sparse tensorization method may be used, such as Smolyak's tensorization [102]. More recently, Gernster and Griebel [32] have proposed an adaptive sparse tensorization method to improve on Smolyak's idea and prioritize the precision of the cubature on certain dimensions relative to others.

We will not dive further into the mathematical details of sparse tensorization. Instead, we refer to Dumont [24] for details on these methods and benchmarks of their respective performance. We showcase the dramatic decrease in evaluation cost, for a given level of quadrature, with Smolyak's method compared to the full tensorization in Figure 1.9.

Sparse cubature rules have some negative weights. Therefore, the positivity of the approximation of the dot product (defined in equation 1.14) is not guaranteed. While this may seem problematic from a mathematical point of view, we did not, in practice, encounter any unexpected behavior using Smolyak's tensorization. However, Constantine et al. [20] and Conrad et al. [16] explored this issue. They proposed alternative sparse grid algorithms for orthogonal projection to ensure, in particular, the absence of internal aliasing (ensure that $\langle \psi_i, \psi_j \rangle = \delta_{i,j}$).

b. Regression

The other family of methods is based on the minimization of the surrogate's error. We search the optimal PCE \tilde{F} of F such that:

$$\tilde{F} = \operatorname{argmin}_{F^* \in \Psi} \left(\mathbb{E} \left[\text{distance}(F^*(\boldsymbol{\xi}), F(\boldsymbol{\xi})) \right] \right), \quad \Psi = \operatorname{Span}(\{\psi_\alpha\}) \quad (1.24)$$




Figure 1.9: Example of tensorization of level 4 quadrature rules on the hypercube $[0, 1]^2$. On the left are the evaluation points needed for the full tensorization. On the right are the points needed by Smolyak's method. Up is Clenshaw-Curtis' rule, and bottom is Fejer's second rule.

The choice of the distance is critical to get a good surrogate \tilde{F} . The basic choice is the Ordinary Least-Square method, which transforms the problem into:

$$\tilde{F} = \operatorname{argmin}_{F^* \in \Psi} \frac{1}{N} \sum_{i=1}^N \|F^*(\boldsymbol{\xi}^{(i)}) - F(\boldsymbol{\xi}^{(i)})\|^2, \quad \Psi = \operatorname{Span}(\{\psi_\alpha\}) \quad (1.25)$$

Where N is the number of observations $F(\boldsymbol{\xi}^{(i)})$ used to calibrate the PCE.

The least-square method can lead to an overfit of the surrogate if the number of samples is not large enough compared to the number of polynomials in the truncated basis. This phenomenon is illustrated in Section 1.3.1.3. A good rule of thumb is to use about three to five times as many samples as coefficients to calibrate.

To improve the quality of the PCE, you can also improve the distance function:

- By adding $\varepsilon \cdot \sum_{k=1}^p |F_k|^2$, $\varepsilon > 0$:
This is called the "Tikhonov regularization" or " \mathcal{L}_2 regularization" [110]. It is a precision/variance trade-off, which means that it trades some discrepancy with the

observations for lower PCE weights. This reduces the variance of the PCE and its tendency to overfit.

- By adding $\gamma \cdot \sum_{k=1}^p |F_k|$, $\gamma > 0$:
This is called "Compressed sensing" or " \mathcal{L}_1 regularization" [109]. This promotes sparsity in the weights by pulling their values to zero.
- By adding both of the previous regularizations. This is called the "Elastic net" [125].

The hyperparameters ε and γ can be adjusted using Leave-One-Out or Cross-validation algorithms [2]. Finally, more advanced algorithms like Least Absolute Shrinkage and Selection Operator (LASSO) [109] or Least-Angle Regression (LARS) [25] can be used for the \mathcal{L}_1 regularization.

This family of methods is much more computation-intensive than the projection method because there is a minimization problem to solve instead of simple arithmetics to evaluate. This minimization problem can even become overwhelming as the cost increases exponentially with both the dimensionality and the highest polynomial order. Still, for a similar sampling cost, regression-based PCE can yield better precision than their projection-based counterparts if the pitfall of overfitting is avoided. This performance topic is addressed in Chapter 2.

1.3.1.3 About that coveted spectral convergence...

The spectral convergence of the PCE is no unicorn in real life. Whatever the evaluation process, one can not simply enrich the polynomial basis and hope the approximation will improve. The spectral convergence is achieved if and only if the coefficients given to each basis function are exactly computed.

In practice, for non-trivial stochastic processes, this comes at a cost:

- For regression-based PCE determination, increasing the number of polynomials - increasing the number of coefficients to determine - will lead to overfitting if new samples are not computed to train the surrogate. That is, the surrogate will eventually pass through every sample point but will oscillate between them. In this situation, the surrogate becomes too specialized to efficiently generalize the dataset and emulate the "true" physical model. This is illustrated in the left part of Figure 1.10.
- For projection-based PCE determination, the effect is a bit subtler. A quadrature method will always give a result for the coefficient of a given polynomial, so you can always compute as many weights as you want given a dataset. The trick is that the higher the degree of the polynomial, the more points are needed to accurately compute its coefficient. A PCE of high degree built with an insufficient number of points will sport ill-determined high-degree terms that will make the surrogate oscillate out of hand. This does not qualify as overfit, though, because the surrogate will not even pass through the sample points. This is illustrated in the right part of Figure 1.10.




Figure 1.10: Examples of PCE (in red) of too large polynomial degree ($p = 6$) compared to the number of training points ($N = 7$, represented by green crosses) for a non-trivial physical model (in black), with an uncertain variable following the standard normal distribution: $\xi \sim \mathcal{N}(0, 1)$. On the left-hand side, a PCE is trained with an Ordinary Least Square Algorithm on a set of $N = 7$ points drawn by a Monte-Carlo algorithm. On the right-hand side, a PCE is trained by projection using a level 3 Fejer quadrature of the second kind ($N = 7$).

1.3.2 Kriging

Kriging, also known as Gaussian process regression, is another surrogate modeling method first introduced by Krige [57].

Consider a set of training data

$$\mathcal{D} = \{(\boldsymbol{\xi}^{(i)}, y^{(i)}), i \in [1, N]\} \quad (1.26)$$

Such that $\forall i, y^{(i)} = F(\boldsymbol{\xi}^{(i)})$.

The Ordinary Kriging approximates F with the mean of the stochastic process Y defined as:

$$Y(\boldsymbol{\xi}) = \mu + Z(\boldsymbol{\xi}) \quad (1.27)$$

Where μ is the sum of the global mean contribution and local variations are represented by the stationary Gaussian process $Z(\boldsymbol{\xi})$. The mean contribution μ can be replaced by a deterministic function of the coordinates $\mu(\boldsymbol{\xi})$ when F has a principal part. This is the principle behind the Universal Kriging introduced by Matheron [69]. The following developments will be restricted to Ordinary Kriging. However, the Universal Kriging case is covered if F is replaced by $F'(\boldsymbol{\xi}) = F(\boldsymbol{\xi}) - \mu(\boldsymbol{\xi})$.

The Kriging Y is defined, at each point $\boldsymbol{\xi}$ by its predictor $\tilde{y}(\boldsymbol{\xi})$ and its variance $\tilde{\sigma}(\boldsymbol{\xi})$. $\tilde{y}(\boldsymbol{\xi})$ and $\tilde{\sigma}(\boldsymbol{\xi})$ are determined using the vector \mathbf{y} of the observations $\{y^{(i)}\}$ and the auto-correlation matrix \mathbf{R} :

$$R_{ij}(\boldsymbol{\theta}) = R(\boldsymbol{\xi}^{(i)}, \boldsymbol{\xi}^{(j)}, \boldsymbol{\theta}) \quad (1.28)$$

where R is a correlation function, or "kernel", and $\boldsymbol{\theta}$ is a set of correlation parameters, "hyperparameters", or "length-scales". Common one-dimensional kernel choices are given

in Table 1.2 and illustrated in Figure 1.11.

Table 1.2: Common choices of one-dimensional kernels. $d_{ij} = |\xi^{(i)} - \xi^{(j)}|$ is the distance between the samples i and j .

Name	$R(\xi^{(i)}, \xi^{(j)}, \theta)$
Linear	$\max(0, 1 - d_{ij}/\theta)$
Exponential	$\exp(-d_{ij}/\theta)$
Squared-Exponential	$\exp(-d_{ij}^2/\theta)$
Matérn 3/2	$(1 + k_{3/2})\exp(-k_{3/2})$ $k_{3/2} = \sqrt{3}d_{ij}/\theta$
Matérn 5/2	$(1 + k_{5/2} + k_{5/2}^2/3)\exp(-k_{5/2})$ $k_{5/2} = \sqrt{5}d_{ij}/\theta$

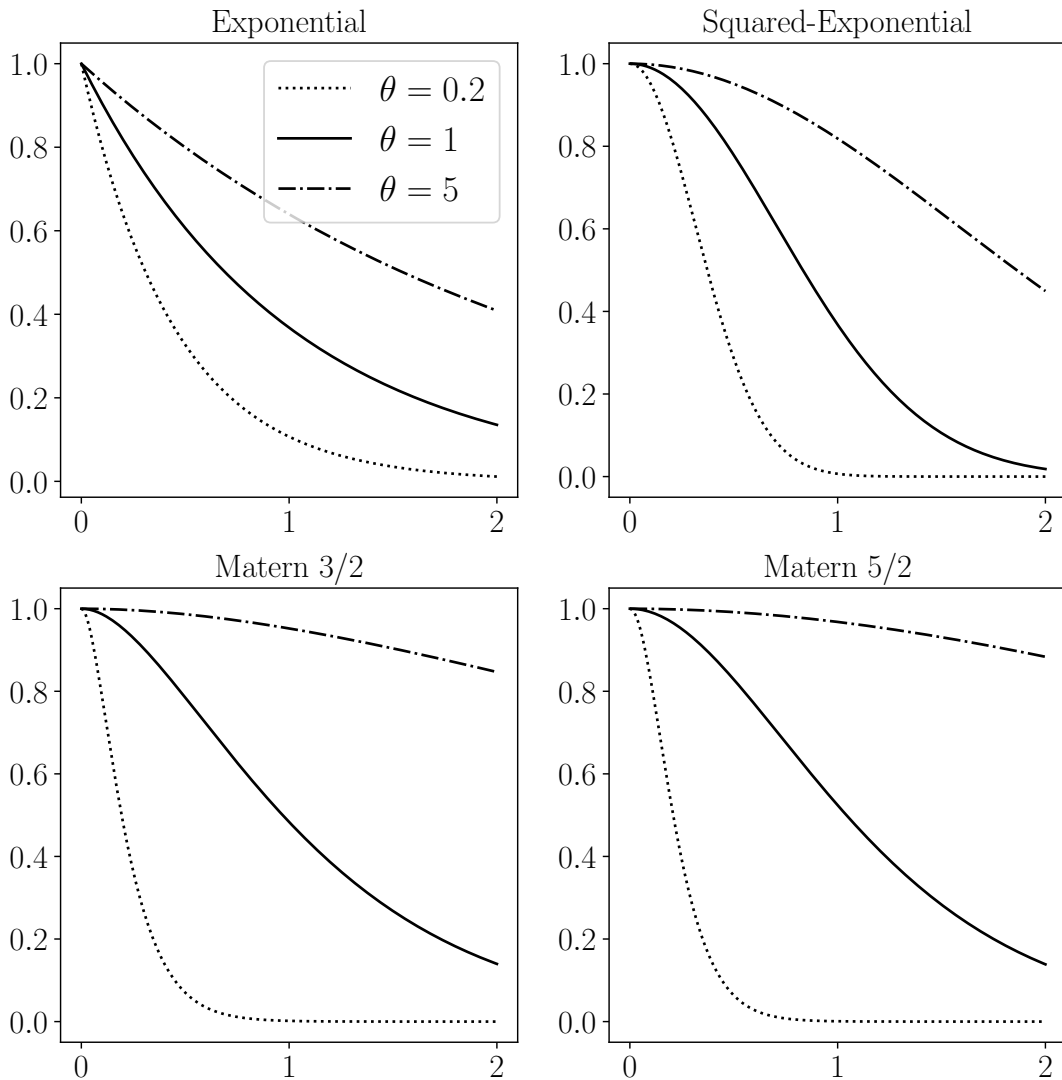


Figure 1.11: Examples of kernel functions. From left to right and from up to down: Exponential, Squared-Exponential, Matérn 3/2 and Matérn 5/2, for $d_{ij} \in [0, 2]$

Multi-dimensional kernels, in dimension d , can be obtained by tensorizing 1D kernels:

$$R_{ij} = \prod_{k=1}^d R(\xi_k^{(i)}, \xi_k^{(j)}, \theta_k) \quad (1.29)$$

The kriging predictor \tilde{y} at any point $\boldsymbol{\xi}$ is defined as the expectation of the realizations of the gaussian process presented in Equation 1.27:

$$F(\boldsymbol{\xi}) \approx \mu_{\tilde{Y}(\boldsymbol{\xi})} = \tilde{y} \quad (1.30)$$

And is computed with:

$$\tilde{y} = \hat{\boldsymbol{\mu}} + \mathbf{r}^T \mathbf{R}^{-1}(\mathbf{y} - \mathbf{1}\hat{\boldsymbol{\mu}}) \quad (1.31)$$

Where \mathbf{y} is the vector of the observations, and $\mathbf{1}$ is a vector of size N containing 1 for each coordinate. The prior estimation $\hat{\boldsymbol{\mu}}$ of the global mean $\boldsymbol{\mu}$ is found with:

$$\hat{\boldsymbol{\mu}} = (\mathbf{1}^T \mathbf{R}^{-1} \mathbf{1})^{-1} \mathbf{1}^T \mathbf{R}^{-1} \mathbf{y} \quad (1.32)$$

And the vector \mathbf{r} is defined by:

$$r_i = R(\boldsymbol{\xi}, \boldsymbol{\xi}^{(i)}, \boldsymbol{\theta}) \forall i \in \llbracket 1, N \rrbracket \quad (1.33)$$

The variance of the gaussian process $\tilde{\sigma}$ can also be computed for any given point $\boldsymbol{\xi}$, with:

$$\tilde{\sigma}^2 = \sigma_{\tilde{Y}(\boldsymbol{\xi})}^2 = \hat{\sigma}^2 (1 - \mathbf{r}^T \mathbf{R}^{-1} \mathbf{r} + u^2 (\mathbf{1}^T \mathbf{R}^{-1} \mathbf{1})^{-1}) \quad (1.34)$$

With:

$$u = \mathbf{1}^T \mathbf{R}^{-1} \mathbf{r} - 1 \quad \text{and} \quad \hat{\sigma}^2 = \frac{1}{N} (\mathbf{y} - \mathbf{1}\hat{\boldsymbol{\mu}})^T \mathbf{R}^{-1} (\mathbf{y} - \mathbf{1}\hat{\boldsymbol{\mu}}) \quad (1.35)$$

The fact that the variance of the gaussian process can be computed at any point means that credible intervals can be predicted along the expectation of the QoI.

The choice of an appropriate kernel and its proper calibration are extremely important, as illustrated in Figure 1.12. Ideally, in the case where a large number of observations is available, a variogram of the data can be plotted. A variogram, first defined in [70], is a plot of the squared difference between the values of each couple of observations as a function of the distance between these observations. A kernel that fits the data can then be chosen and calibrated on this variogram.

In the absence of a sufficient amount of data to plot a meaningful variogram, the hyperparameters can be determined using a Maximum-Likelihood approach [57, 4].

As for the choice of a kernel, Santner et al. [91] recommend using the Matérn 5/2 kernel as a default choice. Furthermore, Schobi et al. [93] state that the influence of the kernel is low in small dimensional problems.

Kriging can also be trained on noisy data, by adding a regularization constant to the diagonal of the correlation matrix \mathbf{R} . This regularization constant is called the 'nugget' variance, and translates a lack of correlation with distance between neighboring samples. This lack of correlation itself can result from uncertain realizations (noisy measurements, variance introduced by hidden variables...) or from a very rough underlying function. This nugget variance is the intercept of the variogram of the data, i.e. its value for a zero distance. Even in the case of non-noisy data, introducing a small regularization constant can also be used to facilitate the numerical inversion of \mathbf{R} .

Finally, Kriging estimates can be used to efficiently resample the uncertain space to sharpen the surrogate model where it is needed. A state-of-the-art review of such techniques can be found in [30].

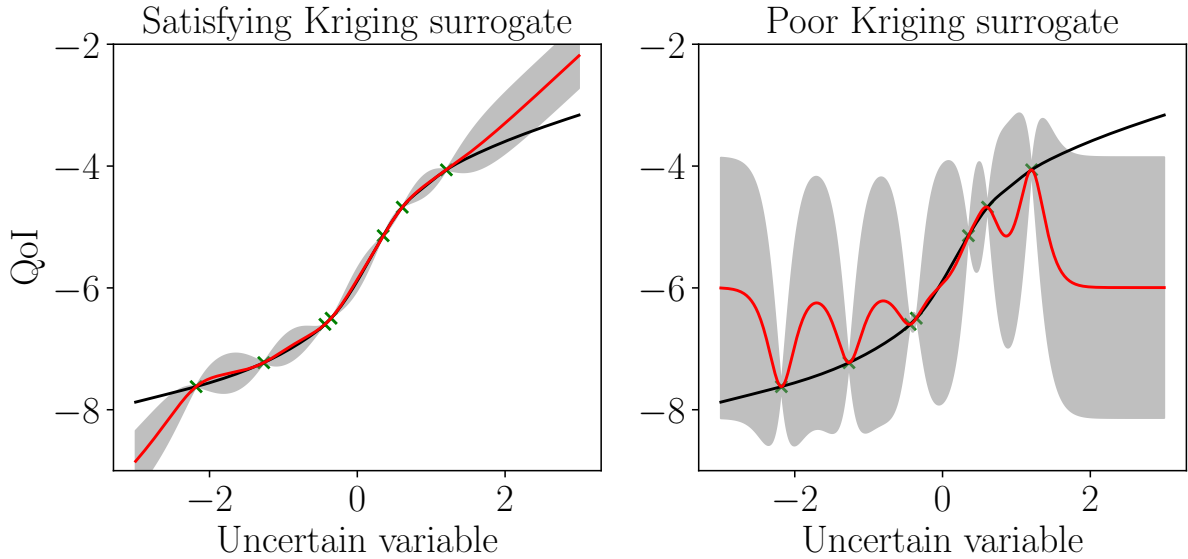


Figure 1.12: Examples of Kriging surrogates (in red) using $N = 7$ Monte-Carlo training points (represented by green crosses) for a non-trivial physical model (in black), with an uncertain variable following the standard normal distribution: $\xi \sim \mathcal{N}(0, 1)$. The 90% credible interval predicted by the surrogate is plotted in grey.

On the left-hand side, a Regression-Kriging with a linear trend and a Matérn 5/2 kernel of hyperparameter $\theta = 0.5$.

On the right-hand side, an Ordinary-Kriging (no trend) and a Matérn 5/2 kernel of hyperparameter $\theta = 0.2$.

1.3.3 Preconditioning

In certain cases, the physical model does not have an agreeable-enough behavior for the accurate calibration of a surrogate model. For example, the model may not be very smooth or comprise a plateau. It may have saturation behaviors. In some cases, the emulator may have positivity constraints. In these cases, the naive application of a surrogate modeling method may not lead to a satisfactory result. To alleviate some of these issues, it may be interesting to apply a pre-conditioning to the observations of the physical model.

A pre-conditioning Φ is an invertible transformation of the observations such that:

$$F'(\boldsymbol{\xi}) = \Phi \circ F(\boldsymbol{\xi}) \quad (1.36)$$

And F' , the pre-conditioned model, has more interesting properties than F .

The best pre-conditioning is case dependent but here are a few pointers:

- To ensure the positivity of a surrogate, it is common practice to use a logarithmic pre-conditioning: $\Phi = \log$
- If F' has a tighter spectrum than F , it will be more accurately represented by a low-order PCE
- The subtraction of the physical model's principal part (cf sec. 1.2.3) can help with the quality of a Kriging surrogate. This is the idea behind Regression Kriging [69] and PC-Kriging [95]

1.4 Sensitivity analysis

Sensitivity analysis is a step sideways from the previous sections. While Monte-Carlo algorithms and surrogate modeling are primarily aimed at estimating a Quantity of Interest or its probability distribution given a set uncertain input, sensitivity analysis aims at determining the relative impact of the uncertain parameters on the QoI.

1.4.1 Local sensitivity analysis

Local sensitivity analysis, as its name suggests, investigates the immediate vicinity of the nominal conditions. The simplest method of this family is gradient analysis, which is very close to the linearization method discussed in Section 1.1.1.1. The gradient of the model F with regard to the components of ξ is computed at a nominal point ξ^* . The components of the gradient may then be multiplied by the variability of their respective components of ξ and finally normalized to obtain local sensitivity coefficients.

The advantage of this method is that it is quite cheap because it requires a minimum of $d + 1$ observations of the model in the non-intrusive framework.

Its downsides are:

- The choice of the nominal point to compute the gradient is arbitrary: It may be more relevant, for a given case, to use either the expectation ($\bar{\xi} = \mathbb{E}[\xi]$), the median value (ξ), or the location of the most probable value of the uncertain input.
- Local derivatives may not fairly represent the variability of the model on the whole input space. This is the same issue as with the other local UQ methods discussed in Sections 1.1.1.1 and 1.1.1.2.
- Interaction effects cannot be retrieved.

Some of these issues may be mitigated by using higher-order derivatives, but the locality of the method remains an important limitation. For this reason, advanced UQ studies tend to use global sensitivity analysis.

1.4.2 Global sensitivity analysis

Global sensitivity analysis aims at filling the gaps of local sensitivity analysis. Ergo, its desirable properties are:

- Taking into account the behavior of the QoI on the whole range of possible values of the uncertain inputs
- Taking into account the probability distribution of the uncertain inputs
- If possible, representing the effect of interactions between random inputs

A powerful metric for sensitivity analysis is the set of Sobol' indices [103]. Each index represents the portion of the variance of the QoI explained by the associated variable or group of variables, which is precisely what we need to discriminate uncertainties of consequence.

In most practical cases, the Sobol' indices are not accessible analytically, and two main

methods co-exist to approximate them numerically.

Brute-force approach

The first approach, originally proposed by Sobol' [103] is a specialized Monte-Carlo algorithm that allows for a precise and controlled approximation of any Sobol' index. The downside is that a specific design of experiment is necessary to evaluate each Sobol' index. The number of evaluations needed to evaluate n_S Sobol' indices is $N(n_S + 1)$ using this method, where N is the number of evaluations needed to estimate the QoI at the desired level of convergence. This very large number of evaluations makes this approach impractical when several indices need to be computed and/or when observations are expensive. Saltelli et al. [90] proposed an alternate method called FAST to evaluate all the first-order Sobol' indices using only the N evaluations used to estimate the QoI. While this approach is much cheaper for the first-order Sobol' indices (the variance explained by one, and only one, variable at a time), it does not allow to compute higher-order indices, meaning the influence of interacting effects of two or more variables.

Indirect approach via the PCE

The other approach, proposed by Schöbi and Sudret [94], uses a PCE of the quantity of interest. Algebraic operations are performed on the weights of the PCE to find all the Sobol' indices at once. This is an obvious advantage over the other method. However, the precision of the estimation of the Sobol' indices depends on the fidelity of the PCE. This entails that the precision of this method is less controlled. Furthermore, when the uncertain dimension is high, the maximum order and the number of interactions considered must be limited so that the determination of the PCE is tractable. Still, this method gives access to sound approximations of the Sobol' indices and allows ranking the effects of all the input variables, which is the key feature of sensitivity analysis when trying to reduce the uncertain dimension. Ergo, this second approach is the one that we selected in this thesis.

To simplify notations, let us assume that the components $\xi^{(i)}$ of ξ are independent uniform random variables on the $[0, 1]$ interval. That way, no pdfs will show in the following formulas. F is then defined on the unit hypercube U^d . Let us also assume that F is square-integrable, that is:

$$F \in \mathcal{L}^2(U^d) \iff \int_{U^d} F(\xi)^2 d\xi < \infty \quad (1.37)$$

The Sobol' indices of the QoI appear in its Sobol'-Hoeffding (S-H) decomposition, which writes:

$$F(\xi) = F(\xi_1, \dots, \xi_d) = F_\emptyset + \sum_{i=1}^d F_i(\xi_i) + \sum_{i=1}^d \sum_{j=1}^d F_{i,j}(\xi_i, \xi_j) + \dots + F_{1,\dots,d}(\xi_1, \dots, \xi_d) \quad (1.38)$$

Or in a more compact form:

$$F(\xi) = \sum_{\iota \subseteq \mathcal{D}} F_\iota(\xi_\iota) \quad (1.39)$$

With F the QoI, ξ the vector of its uncertain inputs and \mathcal{D} the set of indices of the components of ξ , ι a subset of \mathcal{D} , ξ_ι the corresponding input vector, and F_ι the associated S-H functional.

The Sobol' indices are defined for each group of variables ι as:

$$\mathcal{S}_\iota = \frac{\mathbb{V}[F_\iota]}{\mathbb{V}[F]} \quad (1.40)$$

Where $\mathbb{V}[X]$ is the variance of the functional X .

To approximate the Sobol' indices, the S-H functionals can be approximated by their corresponding truncated Polynomial Chaos (PC), that is:

$$F_\iota(\boldsymbol{\xi}_\iota) \approx \sum_{\alpha \in \mathcal{A}_\iota} \psi_\alpha(\boldsymbol{\xi}_\iota) F_\alpha \quad (1.41)$$

Where ψ_α is the α -th basis polynomial, F_α its weight in the PC expansion of F , and \mathcal{A}_ι the set of indices of PC basis polynomials that involve the group of variables ι in the truncated PCE.

Replacing Equation 1.41 in Equation 1.40, we get:

$$\mathcal{S}_\iota(F) = \frac{\mathbb{V}[F_\iota]}{\mathbb{V}[F]} = \frac{\langle F_\iota, F_\iota \rangle}{\langle F, F \rangle} \approx \frac{\langle \sum_{\alpha \in \mathcal{A}_\iota} \psi_\alpha F_\alpha, \sum_{\alpha \in \mathcal{A}_\iota} \psi_\alpha F_\alpha \rangle}{\langle \sum_{\alpha \in \mathcal{A}} \psi_\alpha F_\alpha, \sum_{\alpha \in \mathcal{A}} \psi_\alpha F_\alpha \rangle} \quad (1.42)$$

Where \mathcal{A} is the set of indices of all PC basis polynomials in the truncated PCE.

Finally, the orthonormality of the PC basis gives:

$$\mathcal{S}_\iota(F) \approx \frac{\sum_{\alpha \in \mathcal{A}_\iota} F_\alpha^2}{\sum_{\alpha \in \mathcal{A}} F_\alpha^2} \quad (1.43)$$

This is a simple algebraic formula which uses only the coefficient of a PCE of the QoI, which makes the evaluation of any Sobol' index very cheap once a PCE has been computed.

1.5 Uncertain dimension reduction

When the dimensionality of an uncertain problem is too high, some propagation methods (other than Monte-Carlo whose convergence is dimension-independent) may not be tractable. In particular, as a rule of thumb, surrogate modeling with PCE or Kriging becomes too expensive past about ten dimensions. As a perspective, when dealing with chemical kinetic uncertainties, the number of uncertain parameters ranges from a few tens, for simple fuels, to several hundred when complex hydrocarbons are involved. In this context, reducing the uncertain dimension is crucial to enabling uncertainty propagation methods. Two categories of methods can be distinguished: supervised and unsupervised methods.

1.5.1 Supervised methods

Supervised methods are applied to problems where the dataset comprises both the coordinates $\{\boldsymbol{\xi}^{(i)}\}$ and values $\{y^i(i)\}$ of the observations. In the context of uncertain dimension reduction, it means determining which components or combination of components of $\boldsymbol{\xi}$ are important to explain the variations of y .

1.5.1.1 Additive problems

An additive problem is the simplest possible case for uncertain dimension reduction. It means that the uncertain inputs have an independent impact on the QoI. This translates into cross-sensitivity indices being null or negligible. In terms of Sobol' indices, this can be formalized by:

$$\mathcal{S}_i(F) \ll 1, \quad \forall \iota \mid \text{Card}(\iota) > 1 \quad (1.44)$$

If this condition is verified, the interactions between uncertain variables can be neglected and an uncertainty propagation study can be performed for each dimension independently, which greatly reduces the cost.

1.5.1.2 Dominant variables

An uncertainty propagation problem with dominant variables is another relatively simple case to deal with. As its name suggests, it is a problem where a few uncertain variables explain most of the variance of the QoI. The starting point of dealing with such a problem is to rank the effects of variables and their interactions according to a relevant metric, for example, their Sobol' indices. Then, it may be possible to retain only the most prominent uncertain variables and interactions in the study. Several courses of action can be considered for this segregation:

- Choosing an effect threshold (for example: "A variable or interaction should explain at least 1% of the variance to be kept in the study").
- Choosing a total portion of explained variance in the study: Sobol' indices are summed in decreasing order until this chosen portion of variance is reached. All other variables and interactions are discarded.
- Spotting a significant drop in the values of the Sobol indices sorted by decreasing order. Variables and interactions prior to the drop are kept. All the others are discarded.

Of course, this approach requires a priori knowledge of the behavior of the physical model. This can be achieved, depending on the case, by relying on previous knowledge of the model, by running an a priori study with degraded physics or resolution, by examining an approximate analytical model, etc.

1.5.1.3 Active subspace

An active subspace (AS) is a set of directions of the total uncertain space Ξ on which the QoI varies most [17]. This subspace may be aligned with the directions defined by the original uncertain variables. In this case, the "Dominant variables" approach of Section 1.5.1.2 applies. If the subspace is not aligned, a rotation can be applied to the uncertain space by performing a base change. The goal of this base change is to sort the directions in which the QoI F encounters the greatest gradients. This rotation of the basis is found with:

$$\mathbf{C}_F = \int_{\Xi} (\nabla_{\xi} F(\mathbf{x})) (\nabla_{\xi} F(\mathbf{x}))^T \pi_{\xi}(\mathbf{x}) d\mathbf{x} = \mathbf{W} \mathbf{\Lambda} \mathbf{W}^T \quad (1.45)$$

Where $\mathbf{W} = (\mathbf{w}_1, \dots, \mathbf{w}_d)$ is a unitary matrix defining the rotated basis and $\mathbf{\Lambda}$ the diagonal matrix of eigenvalues sorted in decreasing order:

$$\mathbf{\Lambda} = \mathbf{Diag}(\lambda_1, \dots, \lambda_d) \mid \forall i \in \llbracket 1, d-1 \rrbracket, \lambda_i > \lambda_{i-1} \quad (1.46)$$

The set $\{\mathbf{w}_1, \dots, \mathbf{w}_d\}$ defines a new basis of the uncertain space. This basis can be truncated at the first rank r such that $\lambda_r \gg \lambda_{r+1}$. $\text{Span}(\mathbf{w}_1, \dots, \mathbf{w}_r)$ is called the active subspace.

The expectations for an AS are that it should explain most of the variability of F and that a surrogate model built on this active subspace should be a good approximation of F .

In practice, the necessary number of observations M to accurately determine an active subspace is

$$M = \alpha \beta \log(d) \quad (1.47)$$

in the case of an adjoint simulation where the gradient is obtained "for free" with each sample, and

$$M = \alpha \beta \cdot d \cdot \log(d) \quad (1.48)$$

when the gradient information is obtained by finite differences, where d is the initial dimension of the problem, α an over-sampling factor, and β the largest dimension of active subspace acceptable.

In these situations, the cost of finding an active subspace is untenable when observations are expensive. Constantine et al. propose an alternate solution [18] if certain conditions are met [21]:

- F is monotonous in each direction ξ_i
- An active direction (i.e. an AS of dimension 1) can be discovered

Then, an active direction can be determined with a linear regression of F :

$$F(\boldsymbol{\xi}) = \mathbf{b}^T \boldsymbol{\xi} + b_0 + \varepsilon_F(\boldsymbol{\xi}) \quad (1.49)$$

The active direction \mathbf{w}_1 is expressed as:

$$\mathbf{w}_1 = \mathbf{b}^T / \|\mathbf{b}\|_2 \quad (1.50)$$

And the cost of finding such a direction is only:

$$M = \alpha \cdot d \quad (1.51)$$

Finally, if this method was actually appropriate, a single direction \mathbf{w}_1 on which most of the variability of the QoI F is found. A reasonable low-dimensional approximation can be written as:

$$F(\boldsymbol{\xi}) \approx G(\mathbf{w}_1^T \boldsymbol{\xi}) \quad (1.52)$$

A surrogate-based study is then very easy to conduct in this new 1D space.

Finally, Constantine and Diaz [19] have proposed to derive global sensitivity indices from the determination of the active subspace. When the AS is one-dimensional, they recommend using the components of the eigenvector \mathbf{w}_1 . These components provide insight into both the magnitude and sign of the effect of each original variable on the QoI. When the AS is multi-dimensional and the brute-force approach is used to determine the eigenvectors $\{\mathbf{w}_i\}$ and associated eigenvalues $\{\lambda_i\}$, they recommend using their "activity scores" α_i defined as:

$$\alpha_i = \sum_{j=1}^r \lambda_j w_{i,j}^2, \quad i \in \llbracket 1, r \rrbracket \quad (1.53)$$

1.5.2 Unsupervised methods

Unsupervised methods used datasets where only observations are available. They try to identify patterns in the dataset to build a more concise representation of the data.

1.5.2.1 Principal Component Analysis

Principal Component Analysis (PCA) [79], also called discrete Karhunen–Loève Transform (KLT), Proper Orthogonal Decomposition (POD), or Singular Value Decomposition (SVD) is a statistical tool to reduce, or compress, the dimensionality of a dataset while controlling the variance left on the table by the compression.

In concrete terms, PCA is the process of finding the principal components, or modes, of a dataset. The i -th principal component is the direction of the line that best fits the dataset, with the constraint of being in the subspace orthogonal to the first $i - 1$ principal components. This is illustrated in Figure 1.13.

In practice, we used the SVD formalism implemented in the `numpy.linalg` library [36]. For a dataset matrix \mathbf{M} of shape (n, d) , where n is the size - or the number of samples of the dataset - and d is the dimension - or the number of features - we can write:

$$\mathbf{M} = \mathbf{U} * \mathbf{\Sigma} * \mathbf{V}^T \quad (1.54)$$

Where $*$ denotes the matrix multiplication, $\mathbf{\Sigma}$ is the diagonal matrix of the singular values of \mathbf{M} , and \mathbf{U} and \mathbf{V} are the singular matrices, containing orthonormal vectors in their rows and columns respectively.

The principal components are the rows of \mathbf{P} , such that:

$$\mathbf{P} = \mathbf{\Sigma}^{1/2} * \mathbf{V}^T \quad (1.55)$$

The portion of variance V_i explained by the mode i is found with:

$$V_i = \frac{\Sigma_{i,i}^2}{n} \quad (1.56)$$

The new basis of modes can then be truncated to compress the data optimally - in the sense of retained variance. The possible truncature strategies are the same as those mentioned in Section 1.5.1.2.

This process is useful to compress correlated inputs. An excellent example of this is [62] where the authors reduce an uncertain field of data (hundreds or thousands of correlated uncertain initial conditions) into a few uncertain coefficients associated with PCA modes.

Following the same idea, instead of propagating uncertainties for a very large number of correlated QoIs (such as fields of data in a fluid flow, for example), it is possible to compress these uncertain outputs and perform the propagation of uncertainties on the few coefficients associated to the modes of the output.

Finally, PCA can be used to "clean" data by projecting noisy observations on the truncated basis of modes, which allows representing most of the variance of the data. This cleaning must be exercised with caution because it may erase useful data contained

in the original, "noisy" observations.

PCA is very sensitive to the amplitude of variations in each direction. Therefore, a proper normalization of the data is required to get the best of a PCA.

The formalism of PCA and Active Subspaces are strikingly similar, at least at first glance. In both cases, particularly interesting directions of the initial uncertain space are derived, ranked, and selected using the factorization of a certain matrix. However, the nature of the matrix that is factorized is extremely different between those cases. As a result, this allows both methods to pursue completely different objectives:

- In the case of the AS: The matrix \mathbf{C}_F , defined in Equation 1.45 is a square and symmetric matrix of size $d \times d$, where d is the dimension of the input space of the scalar function F . The components of \mathbf{C}_F represent the expected squared derivative of F in each direction in the original basis. Its diagonalization highlights new directions of the input space in which the scalar F varies most.
- In the case of the PCA: The matrix \mathbf{M} , defined in Equation 1.54, is rectangular of size $n \times d$, and is the juxtaposition of n realizations of a random vector of \mathbb{R}^d . Its factorization highlights a subspace of \mathbb{R}^d where the n realizations can be projected with a minimal loss of variance. Each basis vector of this subspace is also called a "mode" of the dataset.

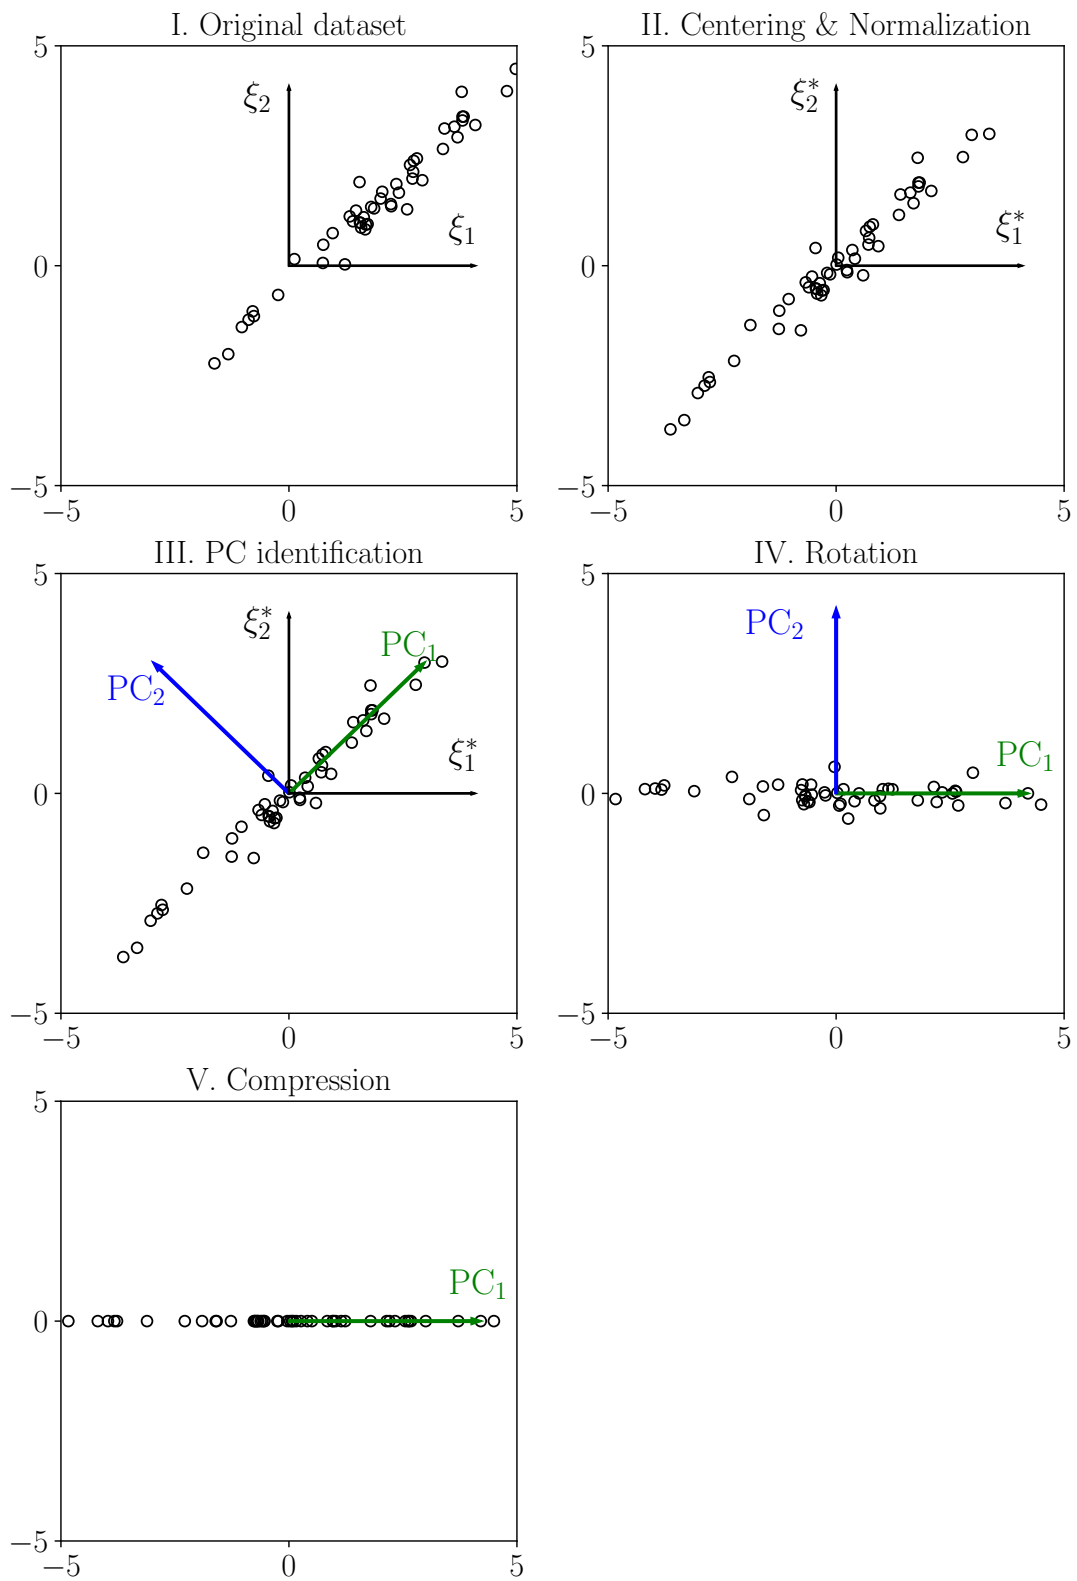


Figure 1.13: Illustration of a PCA in a 2D dataset, compressed to 1 dimension

1.6 Error metrics

The methods presented in Sections 1.3 and 1.5 have in common that they induce a loss of information compared to the reference uncertain problem:

- A surrogate model blurs the features of the physical model it emulates to avoid overfitting. It may also introduce artifacts, like the oscillations of high-order polynomials. Furthermore, any prediction from a surrogate model is only "the best guess" that we could make given the training observations and prior knowledge. A new observation of the physical model at a new coordinate could always yield a different, surprising value.
- A dimension reduction neglects the influence of some variables or combination of variables.

1.6.1 A visual verification: the summary plot

A summary plot of a method is a comparison of observations and predictions for a set of samples. It is mainly applicable to surrogate models, but it can also be used in the case of some dimension reduction methods.

Figure 1.14 shows two examples of summary plots. For each sample, the original model observation is plotted in x -axis, while the prediction is reported in y -axis. If the method which is investigated does not lose or modify any information, all the points collapse on the first bisector ($x = y$).

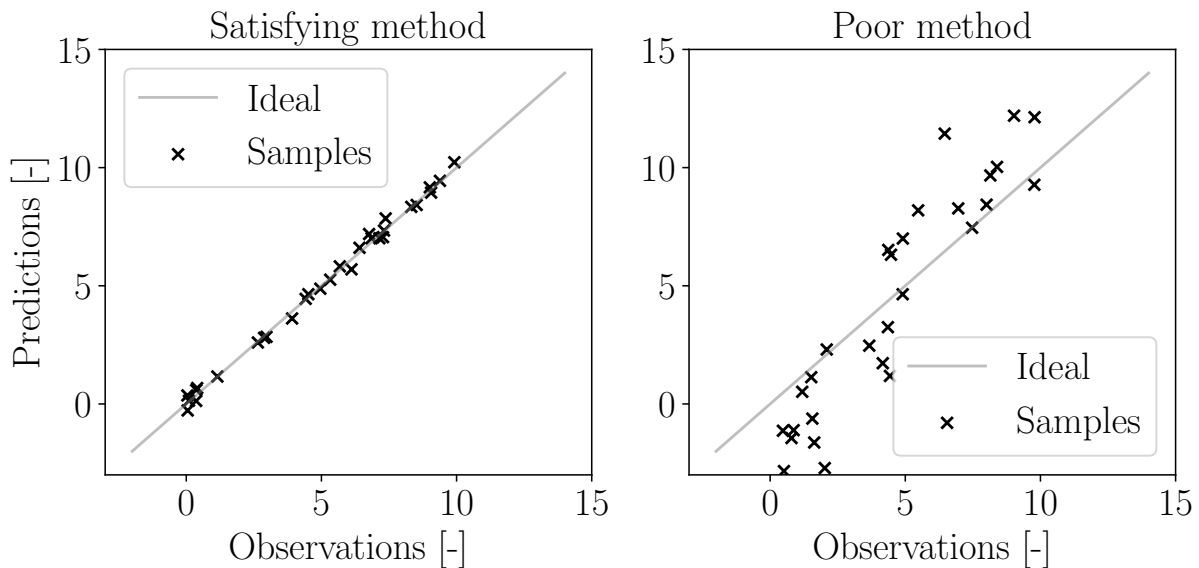


Figure 1.14: Illustration of the summary plot of a satisfying UQ method (left) and of a poor one (right)

In practice, most predictions are not perfect and the points are not exactly on the first bisector. This is the case of the plot on the left side of Figure 1.14, where a "good" surrogate is evaluated. On the right side, however, a very poor method (i.e. with a lot of information loss or modification) is evaluated. While the summary plot gives only visual, qualitative information on the agreement between predictions and observations, it is enough to rapidly rule out blatantly substandard surrogates or dimension reductions.

1.6.2 Distance of probability distribution functions

For more quantitative evaluations, a reference metric is the distance between the predicted PDF and the PDF of observations. Obviously, this metric is only appropriate when observations are sufficiently cheap to allow an extensive sampling of the uncertain space.

The metric used in this thesis is the Wasserstein distance between the reference distribution for a given case - i.e. the result of a brute-force Monte-Carlo study on the original model, accounting for all the uncertainties - and the distribution predicted by the method currently studied. The absolute value of the distance does not mean much in itself. However, comparing the distances obtained with different methods allow ranking them by precision: the lower the distance, the more precise the method.

The Wasserstein distance, also known as the "earth-mover's distance" or "Kantorovitch distance", was first introduced by Kantorovitch in [47]. It can be understood as the cost of changing the value of predictions to match the distribution of observations. The higher the number of observations to change and the higher the change amplitude, the higher the distance. Its precise mathematical definition can be found in [115]. Finally, this choice of metric is an arbitrary choice from our part, and other error metrics like the energy distance give similar results.

Probability distributions can also be compared by evaluating and comparing their successive moments: expectation, variance, skewness, kurtosis, etc.

1.6.3 Performance estimation with validation sets of low cardinality

When observations are more expensive, a reference distribution may not be available. Therefore, another metric needs to be used in order to check a method's quality. In this thesis, we used two such metrics: the Normalised Root Mean Squared Error (NRMSE) and the Normalised Maximum Absolute Error (NMAE) between a set of validation observation and the predictions made by the method at the same coordinates.

These metrics are defined as:

$$\text{NRMSE} = \frac{\|\mathbf{Y}_{\text{obs}} - \mathbf{Y}_{\text{pred}}\|_2}{\|\mathbf{Y}_{\text{obs}}\|_2} \quad (1.57)$$

$$\text{NMAE} = \frac{\|\mathbf{Y}_{\text{obs}} - \mathbf{Y}_{\text{pred}}\|_\infty}{\|\mathbf{Y}_{\text{obs}}\|_\infty} \quad (1.58)$$

Where \mathbf{Y}_{obs} is the vector of the validation observations and \mathbf{Y}_{pred} is the vector of the predictions at the same coordinates.

The NRMSE, whose definition is given in Equation 1.57, measures the global fit of the predictions with the observations. The NMAE, defined by Equation 1.58, measures the maximum discrepancy between the predictions and the observations. Ergo, it is more of a local indicator of the maximum error that may cripple the predictions.

On the one hand, the NRMSE is an efficient - though global - measure of the surrogate fit, even for validation sets of low cardinality, but it may overlook significant local discrepancies between the predictions and the observations. On the other hand, the NMAE will catch these local errors of prediction, provided that the validation set covers the relevant

areas. Still, the NMAE is blind to the fact that the prediction may be excellent over most of the uncertain space in the event that it catches one important discrepancy. This may be especially penalising for polynomial surrogates, which tend to oscillate on the fringes of the uncertain space while having a good fit in high-probability regions.

Both metrics have their pros and cons, but their combination gives a rather complete overview of the quality of a surrogate's predictions, both in terms of global fit and local errors.

1.6.4 Cross-validation

When a separate validation dataset is not available, or when trying to optimize a set of hyperparameters with regards to a certain error metric, cross-validation [107, 106] can be used. Cross-validation is a resampling method, where the available dataset is split randomly in a training and a validation dataset. The model is trained on the former and the error metric is evaluated on the latter. To limit the dependence of the error evaluation on the split of the dataset, it is good practice to repeat this process a number of times and average the results.

A commonly used variation of this method is k -fold cross validation, where the dataset is randomly split in k mutually exclusive subsets of equal size. The model is trained k times using $k - 1$ subsets and evaluated on the remaining subset. Finally, Leave-One-Out (LOO) is an extreme form of k -fold cross-validation, where $k = n$, the number of samples in the dataset.

1.7 Takeaway points

- A set of methods has been selected for the forward uncertainty propagation in the context of very expensive observations.
- Most of these methods have been implemented from scratch in the in-house UQ code SURE, which was entirely developed during this thesis. A few others relied on on-the-shelf libraries. SURE is designed as a modular software which allows for combining UQ methods and plugging in any deterministic model. For the needs of this thesis, interface modules for AGATH (an in-house 0D and 1D flame simulation software) and AVBP [96] (reactive, compressible LES software) were developed.
- This chapter is written in the form of an accessible point of entry for future PhD students discovering uncertainty quantification.

Chapter 2

Uncertainty propagation in H_2/O_2 auto-ignition simulations

The goal of this chapter is twofold. First, the importance of quantifying uncertainties in realistic hydrogen combustion simulations will be substantiated. Second, an efficient framework will be investigated to tackle kinetic uncertainty propagation in CPU-heavy combustion simulations. The case study will be a 1D strained and auto-igniting non-premixed flame. This configuration emulates the physical complexity of more expensive combustion simulations while still being affordable to sample with a Monte-Carlo method. This property is essential as the Monte-Carlo study will provide a reference distribution against which different uncertainty propagation methods can be measured, which allows for validating the developed framework. With the help of the framework developed in this chapter, uncertainty quantification in CPU-expensive simulations will be done in Chapters 4 and 5.

The crux of the framework is to build a surrogate model, or response surface, which is an analytical formula that approximates sufficiently well the behavior of the costly simulation for a negligible fraction of the computational effort. This surrogate, in turn, lends itself perfectly to extensive Monte-Carlo resampling from which any statistic of interest can be retrieved.

The challenge is to build a satisfying surrogate model using as few training samples as possible because each training sample needs a whole simulation to be run. While this may not be particularly difficult in low-dimensional uncertain spaces, it certainly becomes a hurdle for high-dimensional spaces. De facto, this is the case when considering an uncertain kinetic mechanism, even as "simple" a mechanism as one for hydrogen combustion. To tackle this problem, we perform an a priori uncertain dimension reduction using a simplified physical problem that features the same dominant physical phenomena as the original one.

The general framework is then a two-stage process. First, we define the simplified physical problem, which can undergo an extensive Monte-Carlo simulation, and identify the dominant uncertainties among all the input uncertainties. Only uncertainties leading to significant amounts of variance in the quantity of interest are retained. This truncation of this uncertain space results in a much-reduced set of uncertain variables, making surrogate building manageable with few training samples in the second stage. Four surrogate modeling techniques (projection-PCE, regression-PCE, Kriging, and PC-Kriging) are compared to propagate the uncertainty in the 1D igniting flame case using a minimal

number of simulation runs.

The following sections detail the corresponding framework and their results on the 1D strained auto-igniting non-premixed flame. Section 2.1 presents the strained auto-igniting non-premixed 1D flame configuration along with its simplified physical proxy, an auto-igniting homogeneous reactor. A qualitative comparison of the results from several kinetic mechanisms for H_2 combustion is given in Section 2.2. The considered high-dimensional uncertain space is presented in Section 2.3. The methods chosen for the study are very briefly presented in Section 2.4, which refers to Chapter 1 for all the mathematical details of the methods. An extensive Monte-Carlo study is performed for the simplified setup, and an analysis of the resulting uncertainties and underlying physical phenomena is presented in Section 2.5. In particular, quantitative sensitivity analysis is performed in Section 2.5.2 to reduce the uncertain dimension space. Finally, uncertain dimension reduction and surrogate modeling techniques are applied to the target configuration in Section 2.6, which presents and analyses the propagated uncertainties.

2.1 Physical configurations and numerical setups

This section presents the two configurations explored in this chapter and the numerical setups which simulate them. The complex configuration, a strained auto-igniting non-premixed 1D flame, is presented in Section 2.1.1. Quantifying uncertainties in this configuration using as few realizations of the simulation as possible is the goal of this chapter. Then, a simplification of this setup is presented in Section 2.1.2. This simplified configuration, an auto-igniting homogeneous reactor, will be helpful to analyse the physics and uncertainties that dominate hydrogen auto-ignition in a vitiated mixture. In the following, this analysis will allow sample-efficient uncertainty propagation methods to be used on the complex case.

2.1.1 The target configuration: a strained and auto-igniting non-premixed 1D flame

This section presents a setup that emulates complex combustors. The interest is double: (i) we will verify that a non-trivial uncertain behavior can be found in realistic setups; (ii) we will benchmark several UQ methods to handle this uncertain problem using only a few samples, while still being able to afford the complete reference distribution, using Monte-Carlo sampling.

Table 2.1: Nominal boundary conditions for the auto-igniting 1D flame: temperature and species molar fractions. The pressure is set to 101 325 Pa.

Fuel side			Oxidizer side		
T_{fuel}	[K]	305	T_{ox}	[K]	1020
X_{H_2}	[-]	0.25	X_{H_2O}	[-]	0.1
X_{N_2}	[-]	0.75	X_{O_2}	[-]	0.15
			X_{N_2}	[-]	0.75

The setup that will be studied in the following is a 1D auto-igniting non-premixed H_2/O_2 flame. Details of the boundary conditions are given in Table 2.1 and are similar to conditions found in the Cabra H_2 flame [11], only with a slightly colder oxidizer side.

These conditions mimic what is typically found in the later stages of staged industrial burners and combustion chambers. The initial conditions are given by the pure mixing solution from the boundary conditions, which corresponds, in the mixture fraction space, to a linear evolution of the mass fraction composition and enthalpy from one boundary to the other. A unity Lewis number is assumed so that this 1D flame is representative of turbulent combustion [80]. The corresponding evolution of species mass fractions Y_k (with unity Lewis number) in the mixture fraction space is then driven by Equation 2.1:

$$\rho \frac{\partial Y_k}{\partial t} = \frac{\rho \chi}{2} \frac{\partial^2 Y_k}{\partial z^2} + \rho \dot{\omega}_k \quad (2.1)$$

With:

$$\chi(z) = \frac{a}{\pi} \exp\left(-2 \left[\text{erf}^{-1}(2z - 1)\right]^2\right) \quad (2.2)$$

Where ρ is the density of the mixture, z is the mixture fraction, $\dot{\omega}_k$ is the reactive source term of species k , and a is the strain rate of the flame. A similar equation can be written for the evolution of the temperature. In the early stages of auto-ignition, if the transport term - proportional to the strain rate a - is too large compared to the source term $\rho \dot{\omega}_k$ for the intermediate species like O, H, and OH, the flame never ignites. A low strain rate is chosen to ensure the ignition of all sampled flames, even for low initial temperatures or slow realizations of the kinetic mechanism: $a = 10 \text{ s}^{-1}$. This allows a more straightforward comparison between surrogate modeling methods, but similar results can be obtained with higher strain rates by conditioning the surrogate models to the observations where the flame actually ignites.

The 1D flame equations are solved on a uniform 400-point mesh using the 2nd-order stabilized explicit Runge-Kutta method (ROCK2) with error control [1]. The auto-ignition delay is determined when the mass fraction of the hydroxyl radical reaches the threshold $Y_{\text{OH}} = 600 \text{ ppm}$ in the computational domain. The computation time per sample can rise up to a dozen minutes, making it just affordable in Monte-Carlo.

This setup, while featuring species and heat transport and being the backbone of flamelet-based tabulated chemistry for 3D simulations [82, 43], is driven by auto-ignition. Its temporal behavior is presented in Figure 2.1. Starting from the pure mixing line, the gas temperature first starts to rise at the most-reactive mixture $Z_{\text{MR}} = 0.04$ [68], which is a very lean and hot mixture. Then, heat and reactive species are transported towards the slower-reacting mixtures, progressively igniting the whole flame.

2.1.2 The simplified configuration: Auto-igniting homogeneous reactor

This section presents an auto-igniting constant-pressure homogeneous reactor, which will serve as a physics-reduced simplification of the setup presented in section 2.1.1. Its initial composition and temperature are chosen as the conditions initially reigning at the most reactive mixture Z_{MR} of the non-premixed flame. Thus, it represents the chemistry-dominated part of the evolution of the 1D flame, where ignition first occurs. As shown in Table 2.2, the initial temperature is already quite hot and the initial composition includes both N_2 and H_2O .

This case is much simpler, as it features a purely chemical behavior and omits the "transport" part of the equations that govern the evolution of the physical quantities (Equation 2.3). Therefore, it is a suitable reduced model of the costly setup when studying only chemical kinetics.




Figure 2.1: Instantaneous solutions of an example of a 1D auto-igniting flame with the Konnov 2019 mechanism, plotted in mixture fraction space. Solutions are plotted darker with increasing time from $t = 0$ ms up to 60 ms every 2 ms.

$$\rho \frac{\partial Y_k}{\partial t} = \rho \dot{\omega}_k \quad (2.3)$$

Table 2.2: Nominal initial conditions for the auto-igniting homogeneous reactor (Pressure, Temperature and species molar fractions)

P	[Pa]	101325
T	[K]	987
X_{O_2}	[-]	0.1424
X_{H_2}	[-]	0.0126
$X_{\text{H}_2\text{O}}$	[-]	0.0950
X_{N_2}	[-]	0.7500

Similarly to the 1D flame setup, auto-ignition is detected when $Y_{\text{OH}} = 600$ ppm. The temporal integration is performed using the stiffly-accurate implicit Runge-Kutta method of fifth-order accuracy (Radau IIA) with error control [26]. The computational cost per simulation is about 0.1 CPU-second. This makes this setup suitable for a priori studies with extensive Monte-Carlo sampling. As shown in section 2.5, this minimalistic setup (H_2/O_2 chemistry only) is already enough to showcase a non-trivial uncertain behavior.

2.2 A qualitative comparison of deterministic H_2/O_2 mechanisms

This section provides an (admittedly non-exhaustive) overview of the variety of ignition delay times (IDT) that can be found in the case of the auto-igniting homogeneous reactor presented in Section 2.1.2. Figure 2.2 presents the evolution of the IDT as a function of the temperature, for several deterministic combustion mechanisms available in the literature.




Figure 2.2: Comparison of the auto-ignition delay time of the reactor using GRI-Mech 3.0 [101], Hong et al.’s 2011 [40], Konnov’s 2008 [54], and Konnov’s 2019 [55] deterministic mechanisms, at different initial temperatures.

On the one hand, this Figure shows that the IDT varies widely with the initial temperature of the reactor, with two different regimes clearly emerging: "weak" or "mild" combustion at low initial temperatures and "strong" or "sharp" combustion for high initial temperatures [92, 122]. In industrial systems and, in particular, in the context of staged combustion, this initial temperature may not be exactly known, leading to potentially large uncertainties in the behavior of the system.

On the other hand, while the different kinetic mechanisms agree for the most part in both regimes, the transition zone between them is subject to strong disagreement between them. For example, two orders of magnitude separate GRI-mech 3.0 [101] and Konnov (2008) [54] at initial temperatures of about 1000K. Perhaps more interestingly, Figure 2.2, also compares Konnov’s 2008 [54] and Konnov’s 2019 [55] mechanisms, the latter being an update of the former. These two closely related mechanisms still present an order of magnitude of difference in the transition zone. This discrepancy entails that the IDT is very sensitive to the parameters of the kinetic mechanism.

Furthermore, Arrhenius coefficients that parametrize the kinetic mechanisms are uncertain because they are calibrated from sets of measurements of finite precision. Uncertainty factors are reported in some kinetic mechanisms, which means it is both possible and interesting to properly propagate kinetic uncertainties through an auto-ignition simulation, to assess the precision of numerical auto-ignition predictions.

2.3 Uncertainties in hydrogen flame ignition

While hydrogen combustion features a relatively simple kinetic mechanism, the simulation of ignition is still subject to uncertainties that can introduce a significant

variance in the results. In the configuration studied here, a 1D non-premixed auto-igniting flame emulating the staged combustion, only initial oxidizer temperature and kinetic constants are delved into.

2.3.1 Initial temperature

The scope of this study is on staged combustion, and more specifically on the second combustion stage of that process. In such circumstances, the initial temperature of the vitiated gases can be uncertain. This uncertainty can come from thermal losses in the combustor or mixing inhomogeneities, for example. Furthermore, precisely measuring the temperature inside the combustor can be challenging. Laser-based measurements have a typical uncertainty of $\pm 3\%$ [11], and thermocouples are difficult to calibrate in radiative environments such as heated pipes or combustors, also leading to measurement uncertainties. For these reasons, we set a representative uncertainty of ± 30 K for the vitiated gases. We also choose a uniform distribution for this temperature, as described in Equation 2.5, because it is the least-informative distribution for a given range of values. For the auto-igniting 1D flame, this corresponds to an initial oxidizer temperature ranging from 990K to 1050K. The fuel temperature, on the other hand, is considered to be precisely known. For the homogeneous reactor at Z_{MR} , this translates into an initial temperature ranging from 959K to 1016K.

$$T \sim \mathcal{U}(T_{\min}, T_{\max}) \quad (2.4)$$

To represent this uncertainty in a normalized way, a unit uniform random variable ξ_T is introduced:

$$\xi_T = \frac{T_{\text{cof}} - T_{\text{Min}}}{T_{\text{Max}} - T_{\text{Min}}} \sim \mathcal{U}(0, 1) \quad (2.5)$$

2.3.2 Kinetic mechanism

In 2008, Konnov first published a detailed combustion mechanism for H_2 containing uncertainty data for all reactions [54]. It models the Arrhenius pre-exponential factors as log-normal random variables, while the other kinetic parameters remain constant:

$$k_j = A_j T^{\beta_j} \exp\left(\frac{-E_{a_j}}{RT}\right) \quad (2.6)$$

$$\xi_j = \frac{\log(A_j/A_j^0)}{\frac{1}{3}\log(\text{UF}_j)} \sim \mathcal{N}(0, 1) \quad (2.7)$$

With k_j being the rate constant of reaction j , A_j the pre-exponential factor, A_j^0 its nominal value, and UF_j its uncertainty factor.

In this study, we choose to work on an updated version of this mechanism, published in 2019 [55], which brings state-of-the-art nominal Arrhenius coefficients and estimations of their uncertainty factor.

Konnov's 2019 mechanism contains excited species and reactions that involve them. While these are useful for predicting species observed in experimental diagnostics, this extra complexity only leads to more expensive simulations without changing the evolution of predominant species significantly [55]. Therefore, we boiled-down the kinetic mechanism to retain only non-excited species and the relevant reactions. In Figure 2.3, we plot the probability distribution of the auto-ignition delay time of the homogeneous reactor introduced in Section 2.1.2 for several temperatures in the operating range, using

the original mechanism from Konnov [55] (in red) and its reduction to non-excited species and relevant reactions (in turquoise color). This boiled-down mechanism is detailed in Annex A.




Figure 2.3: Violin plots of the logarithm of the homogeneous reactor’s auto-ignition delay time at different initial temperatures, using both the original and boiled-down versions of Konnov’s 2019 kinetic mechanism, with the initial composition from Sec. 2.1.2

Figure 2.3 shows that no significant discrepancy in the probability distribution of the auto-ignition delay time is found over the operating range of temperatures. Finally, the boiled-down Konnov’s 2019 mechanism contains 31 reactions, and therefore 31 uncertain pre-exponential factors.

In total, 32 uncertain variables are accounted for in the study: $\{\xi_1, \dots, \xi_{31}, \xi_T\}$. This relatively high dimensionality means that only Monte-Carlo-like methods will be tractable to tackle this uncertain problem directly.

2.4 Numerical methods for uncertainty propagation

To propagate these uncertainties through the costly simulation, specific uncertainty methods will be needed. The brute-force approach of uncertainty propagation is called Monte-Carlo sampling and is presented in Section 1.2. Although robust and suitable for high dimensional uncertain spaces, the Monte-Carlo approach has a slow convergence rate, meaning that many samples are needed to obtain meaningful statistics of the quantity of interest. This leads to tractability issues when the computation time per sample is significant.

A way to alleviate this problem is to build a surrogate of the costly simulation. A surrogate model, or response surface, emulates the costly simulation for a fraction of the cost, making an extensive sampling affordable. Two kinds of methods are explored in this study: Polynomial Chaos expansion, with its two variants, projection-PCE and regression-PCE, introduced in Section 1.3.1, and Kriging, presented in Section 1.3.2.

Alas, surrogate models need a rising number of training samples when the uncertain

dimension increases. This makes the construction of surrogate models impractical, or even intractable when the uncertain dimension is too high and the computation time per sample is significant.

To allay this problem, we propose to reduce the uncertain dimension a priori by analyzing a cheaper simulation representative of the physical phenomena at play. The mathematical tools that we used are presented in Section 1.5.

2.5 Analysis of uncertainty in the homogeneous reactor’s results and dimension reduction

This section determines and analyses the uncertainty of the simplified case’s auto-ignition delay time. This case is the homogeneous reactor defined in Section 2.1.2, which features hydrogen auto-ignition in a vitiated mixture. A brute-force uncertainty propagation is performed in Section 2.5.1. An analysis of the resulting uncertainty is then carried out, with a search of the dominant input uncertainties (Section 2.5.2) and an interpretation of the shape of the auto-ignition delay time’s probability distribution (Section 2.5.3).

2.5.1 Brute force uncertainty propagation

The homogeneous reactor’s behavior is piloted by the physical phenomenon of auto-ignition. Therefore, a relevant quantity to study is the auto-ignition delay time of the reactor. However, this delay can easily span several orders of magnitude for slightly different initial conditions, as shown in Figure 2.2. This makes the analysis and uncertainty quantification difficult. Therefore, a better Quantity of Interest (QoI) is the logarithm of the auto-ignition delay: $\text{QoI} = \log_{10}(\tau)$. This new QoI is contained in one order of magnitude only, and it has the added benefit of being defined over the entire real axis, whereas the auto-ignition delay must be strictly positive.

The simulations of auto-igniting homogeneous reactors are extremely cheap, which makes a brute-force Monte-Carlo approach affordable. The 32-dimension space $\text{Span}(\xi_1, \dots, \xi_{31}, \xi_T)$, defined in Section 2.3 is therefore sampled extensively. The temporal evolution of the temperature of a handful of these samples is presented in Figure 2.4 to illustrate the wide variety of thermal runaway starting points and slopes. ²¹⁹ QMC samples are drawn using the Sobol’ sequence, and the resulting probability distribution function of the QoI is shown in Figure 2.5. Table 2.3 gathers major statistics of this distribution.

Table 2.3: Major statistics of the auto-ignition delay time of the homogeneous reactor from Section 2.1.2

Auto-ignition delay time τ (ms)			
mean	median	nominal	95% confidence interval
45.4	15.3	14.4	[1.7 , 180.6]

First of all, it should be noted that the PDF of the auto-ignition delay time $\tau = 10^{\text{QoI}}$ spans two orders of magnitude. This uncertainty is far from negligible, and the uncertainty propagation framework will be helpful in this case.




Figure 2.4: Temporal evolution of the temperature for the first 20 QMC samples




Figure 2.5: PDF of the logarithm of the homogeneous reactor's auto-ignition delay time using 2^{19} QMC samples

Second, the PDF of the QoI is bimodal. This shape highlights strong non-linearities in the underlying physical phenomena. Therefore, this non-trivial statistical behavior must be represented by its PDF because usual metrics like mean and standard deviation are ill-suited to describe it.

Furthermore, such a behavior cannot be retrieved using local (linearized) sensitivity analysis. This justifies the use of a full-blown global uncertainty propagation approach for combustion simulations.

2.5.2 Quantitative sensitivity analysis and uncertain dimension reduction

From these QMC samples, Sobol’ sensitivity indices are obtained using the indirect method described in section 1.4.2. As a reminder, the Sobol’ indices are the portion of QoI variance explained by each variable or group of variables. The PCE used to retrieve the Sobol’ indices is determined using least square regression with 10-fold cross-validation using 2^{19} samples. The PC basis is truncated at a maximum order of 5 and first-order interactions ($p=5, q=0.5$) to contain the computational cost of determining the PCE weights.

This PCE is verified visually using the summary plot of Figure 2.6. The fit is not perfect because we constrained the highest order of polynomials and limited ourselves to first-order interactions due to tractability reasons. However, the behavior of the QoI is mostly retrieved, and we will proceed with confidence in the highest Sobol’ indices we can extract from it. These highest indices are plotted in decreasing order in Figure 2.7.




Figure 2.6: Summary plot of the PCE used to compute the Sobol’ indices

This sensitivity analysis highlights the dominance of two uncertain parameters in the problem:

- The initial temperature T : As already exposed in Section 2.2, the QoI is highly sensitive to the initial temperature. The large uncertainty range associated with this variable ($T \sim \mathcal{U}(959 K, 1016 K)$) makes it the most impacting uncertainty in this configuration. Its index stands alone in this study, dwarfing all effects from other uncertain parameters.
- The pre-exponential constant A_{12} : The 12th reaction, $H + O_2 = OH + O$, has been known for its very high sensitivity for a long time [7]. Because of this, it has since been extensively studied, and as a result, its uncertainty factor ($UF_{12} = 1.1$) is the lowest of Konnov’s kinetic mechanism. However, it remains the most impacting uncertain kinetic parameter as already underlined for laminar flame speed [119, 55].




Figure 2.7: Most significant Sobol' indices for the homogeneous reactor, plotted in decreasing order

2.5.3 Interpretation of the bimodal distribution

To understand the bimodal distribution observed for this uncertain problem in Figure 2.5, we propose Figure 2.8 which represents the evolution of the auto-ignition delay as a function of the initial temperature, while also taking into account the uncertainties of the kinetic mechanism. This allows for visually isolating the effect of the initial temperature, which dominates the uncertain problem, from all other uncertain effects.

This figure showcases a "step" in the auto-ignition delay between the initial temperatures of 970 K and 1000 K. This is the transition between the high- and the low- temperature regimes (or "strong" and "weak" regimes respectively). This transition can be observed at slightly different temperatures, depending on the initial composition. For numerical studies like this one, its position also obviously depends on the kinetic mechanism in use.

Considering the probability distribution of the initial temperature - a uniform distribution between 959 and 1016 K - the samples populate both regimes, in which the QoI is much more probable than in the transition zone between the two. Hence, the two bellies of the distribution. A similar conclusion can be reached using Figure 2.3, which shares the same axes but aims at determining the probability distribution of auto-ignition delay at several deterministic initial temperatures.

2.6 Analysis of the auto-ignition delay time's uncertainty in the complex configuration

Section 2.5 proved the necessity of proper uncertainty propagation for the simulation of the auto-ignition of the mixtures typically found in the second stage of staged




Figure 2.8: Logarithm of the auto-ignition delay time of the homogeneous reactor as a function of the temperature. The points are QMC samples from $\text{Span}(\xi_1, \dots, \xi_{31}, \xi_T)$. Note the consistence of the results with the per-temperature investigation shown in Figure 2.3

hydrogen combustion. It also provided insights into the dominant input uncertainties in such simulations and into the behavior of the QoI. Based on these insights, this section pushes forward to propagate uncertainties in the more complex and realistic configuration: the 1D non-premixed auto-igniting flame presented in Section 2.1.1. A Monte-Carlo uncertainty propagation is first conducted in Section 2.6.1 to provide a reference probability distribution of the auto-ignition delay time of the 1D transient flame. This study is allowed by the still manageable cost of such a simulation. Other uncertainty propagation methods are explored in the following because the Monte-Carlo approach is not affordable for typical 3D simulations of turbulent flames. To tackle this problem, the two-stage framework consists first of an uncertainty dimension reduction based on a simplified configuration. The reduction proposed in the previous section is assessed on the actual 1D problem in Section 2.6.2. This allows surrogate models to be built in a second stage using a limited number of training samples. A comparison of surrogate modeling methods is given in Section 2.6.3. Finally, error metrics for the surrogates, usable with a modest set of validation samples, are presented in Section 2.6.4.

2.6.1 Reference distribution from a Monte-Carlo study

The 1D auto-igniting non-premixed flame is still accessible in Monte-Carlo despite being several orders of magnitude more expensive to simulate than a homogeneous reactor. 2^{13} QMC samples are drawn in the 32-dimension uncertain space as a reference for the methods investigated in the following. The resulting probability distribution of the QoI is shown in Figure 2.9. Significant statistics are presented in Table 2.4.

A bimodal distribution is found, much like in the homogeneous reactor case, although this distribution is skewed by heat and species diffusion. This confirms that complex, uncertain kinetic behaviors can be observed in more physically complete setups. In turn, this entails that proper uncertainty quantification can be needed when such situations




Figure 2.9: PDF of the logarithm of the 1D flame’s auto-ignition delay time

Table 2.4: Major statistics of the auto-ignition delay time of the 1D flame

Auto-ignition delay time τ (ms)			
mean	median	nominal	95% confidence interval
52.1	12.0	12.8	[2.6 , 340.2]

arise.

In the case of the design of a complex burner, though, a Monte Carlo experiment is not tractable because the cost of a 3D turbulent combustion simulation typically ranges from a few hundred CPU hours (in RANS) to tens or hundreds of thousand CPU hours (in LES). Therefore, only a few observations can be obtained, on which surrogate models can be built. However, surrogate modeling is problematic in high-dimensional uncertain spaces. This problem will be addressed in the next section.

2.6.2 Impact of the uncertain dimension reduction

Building a surrogate in high dimensions requires many samples, preferably twice as many as uncertain dimensions, to capture at least first-order effects in all directions. In the case of simulations featuring H_2/O_2 combustion, this means at least 64 simulations of the physical model. For complex systems that may be expensive to simulate, especially in the context of design and optimization where several iterations of the system will be compared, such a high number of simulations to build a surrogate can be unaffordable. Therefore, we propose to perform a priori uncertain dimension reduction to drastically reduce the number of samples required to build the surrogate. The most forward way to do it is to conduct a global sensitivity analysis, i.e. sorting the variables whose uncertainty lead to the greatest uncertainty in the QoI. The researcher can then truncate the uncertain space at the level of uncertainty that is chosen to be retained.

The problem, however, is that a global sensitivity analysis requires either sufficiently abundant Monte-Carlo samples or a good surrogate of the phenomenon, which we cannot

afford in the general case. Then, the idea is to use a simplified setup that requires much less computational effort to simulate. It should feature approximately the same physical phenomena and be sensitive to the same uncertain variables. In these conditions, the sensitivity analysis results can be - at least qualitatively - conveyed onto the complex system.

In our case, we need a simple setup that is piloted by auto-ignition in the same conditions as the 1D flame. The homogeneous reactor presented in Section 2.1.2, being chosen at the most reactive mixture of the 1D flame, matches these requirements. In this study, we choose to retain only the uncertain parameters whose influence on the QoI's uncertainty is greater than 1%, i.e. those whose Sobol' index is greater than 10^{-2} . This leaves only the initial temperature T and the pre-exponential constant A_{12} . A new uncertain space is then assembled, spanned by these two variables alone.

The reduced uncertain space, $\text{Span}(\xi_{12}, \xi_T)$, is then resampled in a separate Monte-Carlo simulation, and the simplified physical model is re-evaluated to check the validity of the uncertain space reduction. As can be seen on the left of Figure 2.10, the two pdfs mostly collapse, indicating that most of the original uncertainty of the QoI is retained in the reduced uncertain space.




Figure 2.10: Comparison of Monte-Carlo experiments in the full and in the reduced uncertain space $\text{Span}(\xi_{12}, \xi_T)$. Left: On the homogeneous reactor. Right: On the 1D autoigniting flame.

The more complex physical model is also resampled on the reduced uncertain space as a means of step-by-step validation. Of course, in the general case, neither this resampling nor the original "reference" sampling can be afforded and one will have to rely on the error metrics described in section 1.6. The results of this resampling are shown on the right of Figure 2.10. Again, the pdfs are very similar, indicating that the sensitivity analysis

and the resulting uncertain dimension reduction could indeed be extrapolated to the more complex setup.

2.6.3 Comparison of the surrogate modeling methods

The performance of the four surrogate modeling methods presented in Section 2.4 (Regression PCE, Projection PCE, Kriging, and PC-Kriging) are now compared on this particular setup. Surrogate models will be trained on the reduced uncertain space $\text{Span}(\xi_{12}, \xi_T)$, using different sizes of training sets: The projection-PCE will use 17, 49, and 129 samples, corresponding to the number of points needed in the second, third, and fourth levels of the sparse Fejer cubature. The regression-PCE, Kriging, and PC-Kriging will be trained on a Quasi Monte-Carlo training set, using 10, 15, and 40 observations. To avoid overfitting, the PCEs are limited to a total order of 4. As shown in Figure 2.11 on the example of the regression PCE, the quality of the surrogate increases with the number of samples in the training set, which translates into a better fit between the PDF predicted by the surrogate and the reference PDF in the reduced uncertain space.




Figure 2.11: Convergence of the pdf obtained by resampling the regression-PCE, for several sizes of training sets

This better fit between PDFs can be quantitatively measured by their Wasserstein distance. The evolution of this distance for each type of surrogate and each size of training set is plotted in Figure 2.12.

This plot highlights that projection-PCE is not performing as well as Kriging and regression-PCE in this configuration. The projection-PCE has a spectral convergence towards the "real" evolution of the QoI, which means that given enough exactly computed polynomials, the real evolution of the QoI can be retrieved with as much accuracy as desired. However, in the case of expensive simulations, the weights of higher-order polynomials can not be adequately computed because their evaluation requires a large number of samples. Therefore, with few available samples, projection-PCE is stuck with low-order effects and interactions. This is not a problem per se, as highlighted by the performance of




Figure 2.12: Evolution of the Wasserstein distance between the reference pdf (in the 2D uncertain space) and the resampling for each surrogate modeling method as a function of the size of the training set

the regression-PCE which uses the same limited polynomial basis. However, the weights computed using projection do not optimize the global fit of the surrogate to the observations but are merely projections on basis polynomials.

On the other hand, both Kriging and regression-PCE make good use of the available samples to mimic the behavior of the QoI. They have a higher convergence rate, and the fit is already more than acceptable for low training set sizes (15 samples). However, with these methods, one must take care of the pitfall of overfitting. For the PCE, a rule of thumb is to avoid using a number of polynomials higher than half of the cardinality of the training set. However, this may not be enough to avoid overfitting, and we recommend training PCE surrogates with several sizes of polynomial basis and comparing them using the available quality metrics. For Kriging, overfitting is avoided by correctly choosing the Kernel and length scales. In our experience, Matérn 5/2 kernel and MLE determination of length scales work in most cases.

It is also important to note that Kriging and regression-PCE can use the same training samples, i.e. random or quasi-random samples, while projection-PCE needs to use its own set of training samples. Kriging and regression-PCE also yield similar performances while the projection-PCE surrogate lags behind.

Projection-PCE also suffers from a very specific issue. For some sample points, it may not be possible to properly define an observed QoI. In the case of a 1D auto-igniting flame, if the kinetics are too slow, heat and intermediate species are diffused too fast in comparison, and the flame never ignites. This can typically happen at the edges of the uncertain space, which are reached by the quadrature scheme. Discontinuities of that sort can be handled with regression methods by building a surrogate conditioned by auto-ignition and ignoring ill-defined observations in the training process. However, with projection methods, this is not possible because a value for the QoI must be specified at each quadra-

ture point. The chosen value impacts the shape of the response surface, and therefore the fidelity of the response surface. In this study, for the sake of comparing different surrogate modeling methods, this problem was avoided by setting a very low strain rate for the 1D flames, as specified in section 2.1.1, ensuring that all sampled flamelets actually ignite.

Finally, in this case, PC-Kriging using a 2nd order regression-PCE as trend outperforms the other surrogate model techniques by combining the smooth global trends of PCE with the possibility of local variations offered by Kriging.

Therefore, to maximize both the performance and the robustness of the surrogate modeling step, we recommend sampling the reduced uncertain space with a space-filling sampling method, like a Sobol' sequence, then building regression-PCE, Kriging, and PC-Kriging surrogates over these training points. The best surrogate model can then be selected a posteriori depending on which performs better in each specific situation.

2.6.4 Error measurement with a modest validation set

In the general case, when the reference distribution is not available, the alternative metrics presented in Section 1.6 can be used. Using 20 validation samples drawn with Latin Hypercube Sampling (LHS), we can compute the Normalized Root Mean Squared Error (NRMSE) and Normalized Maximum Absolute Error (NMAE) of each surrogate model. Figure 2.13 presents their evolution with the size of the training set. On the one hand, the evolution of the NRMSE is in qualitative agreement with the reference metric, making it suitable to compare the global quality of surrogates when the reference distribution can not be obtained. The NMAE, on the other hand, gives a measure of the maximum error of the surrogate. In particular, it can translate the increasing oscillations of the regression-PCE at the edges of the uncertain space when the maximum order of polynomial increases. The NRMSE and NMAE also have the advantage of being more quantitatively intuitive measures than the Wasserstein distance. Overall, their combination is suitable to evaluate the quality of a surrogate method, both in terms of global fit and potential localized error.




Figure 2.13: Evolution of the alternative metrics for each modeling method as a function of the size of the training set, using 20 validation LHS samples. Left: NRMSE. Right: NMAE

2.7 Takeaway points

On the physical interpretation side:

- This study has shown that typical uncertainties in the initial temperature of a mixture can significantly impact H_2/O_2 combustion when the auto-ignition phenomenon is involved.
- Simulations of such configurations are also impacted by the uncertainties in the kinetic mechanism itself.
- These effects were shown to happen not only in the relatively simple case of homogeneous reactor auto-ignition but also in the more complex case of auto-igniting non-premixed 1D flames.
- Significant non-linearities in the transition zone between high- and low-temperature chemistry can lead to heavily skewed or even bimodal probability distributions of the auto-ignition delay.
- A two-dimensional space, $\text{Span}(\xi_{12}, \xi_T)$, was shown to explain most of the variability of both auto-igniting homogeneous reactors and non-premixed 1D flames in the investigated conditions.

On the methodological side:

- In configurations where auto-ignition is involved, a proper uncertainty propagation framework must be used to tackle these uncertainties in combustion simulations to provide researchers and designers with reliable, uncertainty-enriched simulation results
- Propagating high dimensional uncertainties in combustion simulations is challenging because the Monte-Carlo approach is typically not affordable, and surrogate modeling methods require large training samples due to the curse of dimensionality.
- We put forward a framework based on a simplified configuration that shares the same underlying physical phenomena. Extensive sampling and sensitivity analysis can be performed on this less costly simulation, allowing a priori uncertain dimension reduction for the more complex configuration.
- We recommend sampling the reduced uncertain space with a space-filling method that allows building both Kriging surrogates and regression-PCE and picking whichever surrogate method works best. The comparison between the two methods can be done using the NRMSE and NMAE error metrics over a set of validation samples.
- With these tools, we are ready to tackle uncertainty propagation in expensive simulations such as 3D Large Eddy Simulations of reacting flows.

Chapter 3

The H₂ Cabra flame

This chapter presents the H₂ Cabra flame [11] and its Large Eddy Simulation. Section 1 presents the experimental setup and the flame's characteristics of interest. Section 2 presents the numerical setup used in this thesis to simulate the Cabra flame. Section 3 presents the results of the nominal simulation. Section 4 studies the convergence of the results with regards to the mesh size and length of the averaging time period. Sections 5 and 6 explore qualitatively the sensitivity of the Cabra flame to various parameters.

3.1 The experimental setup

The Cabra flame is a lifted non-premixed flame in a vitiated co-flow. As shown in Figure 3.1, a jet of H₂ diluted in N₂ exits from the central nozzle at ambient temperature. It exits in a co-flow of hot, burned gases composed of H₂O, O₂, and N₂. This hot co-flow is obtained by surrounding the base of the central nozzle with a myriad of small premixed H₂-air burners in a lean mixture.

3.1.1 Choice of operating conditions

Many combustor designs use the recirculation of hot combustion products. This can be achieved for example by using a swirled injector which creates an inner and an outer recirculation zone around the cone of injection. With lower local velocities and bountiful activation energy coming from the high temperature of the burnt gases, this recirculation helps the stabilization of flames in highly turbulent flows.

The modelization of such combustors is both difficult and computationally intensive due to the coupling between a complex, turbulent flow and the chemical kinetics. A corollary of this state of affairs is that simulations of such combustors are difficult to validate, because the sources of error are plenty, and modelization errors can sometimes compensate, leading to a "good-looking" solution for the wrong reasons.

For this reason, Cabra et al. [11] introduced the Cabra flame configuration, which partly decouples the complex flow configuration of industrial burners from the chemical kinetics by employing a jet in co-flow configuration. The fuel is injected in a cold, diluted central jet while the oxidizer is present in the hot, vitiated co-flow at atmospheric pressure. This allows studying the kinetics of a non-premixed flame in a vitiated environment in a rather simple flow configuration.

On the one hand, from an experimental point of view, this is practical because it allows easy optical access to the flame, which in turn allows many different diagnostics to be




Figure 3.1: The experimental Cabra flame setup [11]

performed. On the other hand, from the numerical point of view, this configuration also eliminates sources of error by using a simple and well-known flow configuration and reducing the sensitivity to wall boundary conditions.

Lastly, the flame is lifted in its nominal operating conditions, meaning that the flame front is not attached to the lips of the burner. It rather hovers a few jet diameters above them. This property is desirable in many practical applications because a lifted flame does not heat the injector as much as an attached flame.

Two versions of this configuration were proposed by Cabra et al.: A methane flame [12] and a dihydrogen flame [11, 12]. In this thesis, we studied only the H_2 flame. Quantitative details of the operating conditions are given in Table 3.1.

3.1.2 Flame structure and behavior

Cabra et al. noted in [11] that the flame spontaneously ignites downstream soon after the fuel inlet is opened. This event happens at $H/D_{\text{jet}} = 660$, which is far above the steady state flame front which hovers at about 10 diameters of the fuel inlet. This first observation indicates that the Cabra burner can - and does - auto-ignite. This initial auto-ignition is followed by "a sequence of loud auto-ignition events, each successively closer to the nozzle exit" [12], until the lifted flame stabilizes at its steady-state position.

Cabra et al. also asserted that the stabilization of the flame in its steady-state is mostly driven by auto-ignition, based on intermediate species budgets from a RANS-

Table 3.1: H₂ Cabra flame experimental conditions [11]

Central jet		Co flow	
Q _{H₂} (slm)	25	Q _{H₂} (slm)	225
Q _{N₂} (slm)	75	Q _{air} (slm)	2100
T _{jet} (K)	305	T _{coflow} (K)	1045
V _{jet} (m/s)	107	V _{coflow} (m/s)	3.5
Re _{jet}	23600	Re _{coflow}	18600
D _{jet} (mm)	4.57	d _{coflow} (mm)	210
X _{H₂}	0.25	X _{H₂O}	0.1
X _{N₂}	0.75	X _{O₂}	0.15
t _{lips} (mm)	0.89	X _{N₂}	0.75

Q: volumetric flow rate; X: mole fraction;

Re: Reynolds number; D: diameter; t:thickness

PDF simulation. This conclusion is shared by another experimental study by Wu et al. [120], another RANS-PDF simulation by Gordon et al. [33], and a Direct Numerical Simulation by Yoo et al. [123].

Finally, the structure of the steady-state flame can be described as follows:

- The fuel and oxidizer mix in the peripheral regions of the central jet.
- Small auto-igniting pockets form in places where a highly-reactive, hot, and lean mixture is found.
- These pockets swell and fuse while being convected downstream, ultimately forming a plume of burnt gases

3.1.3 The Quantity of Interest: the flame lift-off height

A relevant quantity of interest to describe the Cabra flame is its lift-off height. [11] defines this lift-off height as the axial location where the Favre-filtered hydroxyl mass fraction Y_{OH} first reaches 600 ppm.

3.1.4 Sensitivity to the co-flow temperature

The auto-igniting nature of the Cabra flame entails that the flame position and shape are affected by how reactive the initial mixture is.

The co-flow is formed by a hot, oxygenated gas (1045 K in the nominal conditions). It is the thermal energy brought by this gas that jumpstarts the combustion reaction with the cold fuel jet. The initial reactivity of an auto-igniting reactor varies exponentially with its initial temperature. Therefore, even small variations in the co-flow temperature should lead to significant modifications of the flame position and/or shape. Figure 3.2 presents the results of experiments by Cabra et al., Gordon et al., and Wu et al. [11, 34, 120]. It shows the normalized lift-off height plotted against the measured co-flow temperature. As expected, the sensitivity of the lift-off height to the co-flow temperature is significant in the vicinity of the nominal operating conditions. Also interesting is the fact that these three authors seem to be in severe disagreement unless the uncertainty on the measurement of the co-flow temperature is taken into account. This uncertainty, evaluated as $\pm 3\%$ in [11], substantiates the agreement between the four experiments.

However, it also highlights the difficulty to validate a deterministic simulation, or worse, to discriminate between two modelization choices, based on these experimental results.




Figure 3.2: Experimental measurements of the normalized lift-off height of the Cabra flame as a function of the measured co-flow temperature. [11, 120, 34]. The red error bar denotes the experimental uncertainty on the measurement of the co-flow temperature as reported in [11]

3.2 Numerical simulation of the cabra flame

Since the Cabra flame is mainly piloted by the local and transient phenomenon of auto-ignition, and since we wish to simulate it in the most precise possible way, the framework of choice is Large Eddy Simulation (LES).

LES is a framework of turbulent simulation that resolves the large scales of fluid motion, while smaller scales are filtered and their effect are modelled. Cost- and precision-wise, it falls between Direct Numerical Simulation (DNS) which resolves all scales of turbulence and Reynolds-Averaged Navier-Stokes simulations (RANS) which resolves only the mean fields. The large scales are geometry-dependent and contain most of the kinetic energy of the fluid. The smaller scales are more geometry-independent and contain less energy, which allows them to be modeled more accurately than the large scales.

To perform said LES of a reactive flow, a number of physical modelization choices need to be made, along with very practical discretization choices. To be precise, we need to choose:

1. A turbulence model
2. A chemical kinetic mechanism
3. A combustion model
4. A turbulence-combustion interaction model

And :

1. A computational domain
2. Boundary conditions
3. A spatial mesh
4. Numerical methods

3.2.1 Favre-filtered Navier Stokes equations

A Large Eddy Simulation solves a set of Favre-filtered Navier-Stokes equations. In the following, for any variable ϕ , Favre-filtered quantities are noted $\tilde{\phi}$, and Reynolds-filtered quantities are noted $\bar{\phi}$, such that $\bar{\rho}\tilde{\phi} = \overline{\rho\phi}$.

The equations that will need to be solved, regardless of the modelization choices, are the following:

Mass balance $\bar{\rho}$

$$\frac{\partial \bar{\rho}}{\partial t} + \nabla \cdot \bar{\rho} \tilde{\mathbf{u}} = 0 \quad (3.1)$$

Where ρ is the gas density and \mathbf{u} is its velocity vector.

Momentum balance $\bar{\rho} \tilde{\mathbf{u}}$

$$\frac{\partial \bar{\rho} \tilde{\mathbf{u}}}{\partial t} + \nabla \cdot (\bar{\rho} \tilde{\mathbf{u}} \tilde{\mathbf{u}}) = -\nabla \bar{P} + \nabla \cdot \bar{\boldsymbol{\tau}} + \nabla \cdot \boldsymbol{\tau}_{\mathbf{u}} \quad (3.2)$$

Where P is the gas pressure, $\bar{\boldsymbol{\tau}}$ is the laminar stress tensor and $\boldsymbol{\tau}_{\mathbf{u}}$ is the subgrid stress tensor.

The laminar and subgrid stress tensors are defined respectively as:

$$\bar{\tau}_{ij} = \bar{\mu} \left(\frac{\partial \tilde{u}_i}{\partial x_j} + \frac{\partial \tilde{u}_j}{\partial x_i} \right) - \frac{2}{3} \bar{\mu} \frac{\partial \tilde{u}_k}{\partial x_k} \delta_{ij} \quad ; \quad \tau_{\mathbf{u}ij} = -\bar{\rho} (\widetilde{u_i u_j} - \tilde{u}_i \tilde{u}_j) \quad (3.3)$$

With $\bar{\mu}$ the laminar dynamic viscosity of the fluid.

Total energy balance \tilde{e}_t

$$\frac{\partial \bar{\rho} \tilde{e}_t}{\partial t} + \nabla \cdot (\bar{\rho} \tilde{\mathbf{u}} \tilde{e}_t) = -\nabla \cdot (\bar{P} \tilde{\mathbf{u}}) + \nabla \cdot (\bar{\boldsymbol{\tau}} \tilde{\mathbf{u}}) - \nabla \cdot \bar{\mathbf{q}} - \nabla \cdot \mathbf{q}^t \quad (3.4)$$

e_t is the gas total energy, $\bar{\mathbf{q}}$ is the laminar heat flux and \mathbf{q}^t is the subgrid heat flux. The viscous flux $\bar{\boldsymbol{\tau}} \tilde{\mathbf{u}}$ can also be written $\bar{\boldsymbol{\tau}} \tilde{\mathbf{u}}$. The laminar and subgrid heat fluxes are respectively defined as:

$$\bar{\mathbf{q}}_i = \frac{\bar{\lambda}}{\bar{C}_p} \frac{\partial \tilde{h}}{\partial x_i} \quad ; \quad \mathbf{q}_i^t = \overline{(\rho e_t + P) u_i} - (\bar{\rho} \tilde{e}_t + \bar{P}) \tilde{u}_i \quad (3.5)$$

Where $\bar{\lambda}$ is thermal conductivity of the gas, \bar{C}_p is its laminar heat capacity at constant pressure and h is its enthalpy.

3.2.2 Physical choices

3.2.2.1 Turbulence model

The subgrid stress tensor $\tau_{\mathbf{u}}$ is modeled using the eddy-viscosity assumption:

$$\tau_{\mathbf{u}_{ij}} - \frac{1}{3}\delta_{ij}\tau_{\mathbf{u}_{ll}} = 2\bar{\rho}\nu_t \left(\tilde{S}_{ij} - \frac{\delta_{ij}}{3}\tilde{S}_{kk} \right) \quad (3.6)$$

Where:

$$\tilde{S}_{ij} = \frac{1}{2} \left(\frac{\partial \tilde{u}_i}{\partial x_j} + \frac{\partial \tilde{u}_j}{\partial x_i} \right) \quad (3.7)$$

The subgrid-scale viscosity ν_t is modeled in this study using the so-called Sigma model [76], which is written:

$$\nu_t = (C_\sigma \Delta)^2 \cdot \frac{\sigma_3(\sigma_1 - \sigma_2)(\sigma_2 - \sigma_3)}{\sigma_1^2} \quad (3.8)$$

where Δ is the filter characteristic length, $C_\sigma = 1.5$ is the model constant and $\sigma_1 \geq \sigma_2 \geq \sigma_3 \geq 0$ are the singular values of the resolved velocity gradient tensor.

This turbulence model is the recommended model for 3D simulation in the AVBP software [96]. It was developed to alleviate some drawbacks of the previous static turbulence models, notably by converging to zero subgrid-scale effects in laminar flows and having the correct asymptotic behavior near solid boundaries.

The subgrid heat flux \mathbf{q}^t is modeled using the Standard Gradient Diffusion Hypothesis (SGDH):

$$\mathbf{q}^t = -\frac{\lambda_t}{\tilde{C}_p} \nabla \tilde{h} \quad (3.9)$$

Where the turbulent thermal conductivity λ_t is given by:

$$\lambda_t = \frac{\bar{\rho}\nu_t\tilde{C}_p}{Pr_t} \quad (3.10)$$

Where Pr_t is the turbulent Prandtl number whose value is set at $Pr_t = 0.6$, the default in the AVBP software.

3.2.2.2 Kinetic mechanism

The ultimate goal of this thesis is to study the uncertainties in the Cabra flame, and in particular the uncertainties introduced by the kinetic mechanism. Therefore, for the LES of the Cabra flame, we chose to use the boiled-down Konnov 2019 [55], which was introduced in chapter 2, section 2.3.2.

3.2.2.3 Combustion and flame-turbulence interaction model

Since the flame is a diffusion flame and is auto-igniting [120], we leverage the Unsteady Flamelet Progress Variable (UFPV) approach with a presumed PDF closure model, proposed by Ihme and See [43]. The UFPV approach consists of the tabulation of non-premixed auto-igniting flamelets. It is important to note that this approach will predict accurate source terms solely for the auto-ignition of a mixture in the flamelet regime.

The flamelet model

The flamelet model, on which the UFPV tabulation is based, assumes that the flame front is thin compared to its wrinkling scales and other flow scales. This allows to assimilate the flame front to a locally-1D laminar flame, or flamelet, whose structure depends only on time and on the coordinate normal to the flame front. This structure can be pre-computed and stored in a table, from which thermophysical properties, composition and reaction source terms can be read at each time step and at every point of a 3D simulation.

In our case, a non-premixed jet flame, the flamelet of choice is a diffusion flamelet, which describes the structure of a counter-flow, non-premixed flame. The relevance of this choice is illustrated in Figure 3.3, which presents, in a simplified manner, the mixing in a shear layer behind a splitter plate, which can also be seen as a cross-section of one side of a non-premixed jet.




Figure 3.3: Mixing in a shear layer behind a splitter plate [80]

This figure shows that the shear layer can very locally be seen as a counter-flow flame, which calls for a non-premixed flamelet modelization.

Furthermore, it also illustrates that the properties of the flow, in this case the mixture fraction, can vary significantly across very short distances in a turbulent flow. These short distances may be smaller than the mesh size, which means that these variations must be modeled and taken into account when building the table. This is where the presumed PDF closure model comes into action. We also chose unsteady flamelets to correctly represent the auto-ignition of the mixture. Finally, a unity Lewis number is chosen for the simulation of the flamelets, as it is representative of turbulent combustion [80]. As a reminder, these flamelets are the ones studied in Chapter 2.

The UFPV table

The flamelets whose solutions are used to build the table should span the domain where they can auto-ignite. This domain can be determined by plotting the so-called S-curve.

The S-curve, shown in Fig. 3.4, displays the maximum temperature of the flamelet as a function of the strain rate it is subject to, all other parameters being equal. It contains two critical points:

- The first critical point corresponds to the quenching strain rate a_q beyond which heat and radical species diffusion become too strong for the flamelet to sustain combustion.

- The second critical point corresponds to the ignition strain rate a_i beyond which heat and radical species diffusion become too strong for the cold mixture to auto-ignite. This second critical point is the upper bound of our table.




Figure 3.4: S-Curve computed in the operating conditions of the Cabra H_2 flame. The first and second critical points are circled in red. The extinct branch is shown in dotted line.

The table should include flamelets with a strain rate going from zero up to the second critical point of the S-curve. Alas, the transient simulation of the auto-igniting flamelet is not stable for too small strain rates. Thus, an empirical value of $a = 5 \text{ s}^{-1}$ was chosen as a lower bound of the table, because it was the smallest strain rate to guarantee numerical stability.

A set of 10 flamelets are computed, each one with a different strain rate, ranging from this empirical lower bound up to the strain rate of the second critical point of the S-curve. The evolution of the gas temperature in such a flamelet is plotted in Figure 3.5.

From these flamelets, we could pre-compute the auto-ignition of non-premixed laminar flames, and tabulate, for each quantity ϕ , the function \mathcal{F}_ϕ :

$$\phi = \mathcal{F}_\phi(z, \Lambda, \chi_{z,st}) \quad (3.11)$$

Where z is the mixture fraction, $\chi_{z,st}$ is the stoichiometric scalar dissipation rate and Λ is a mixture fraction-independent reaction progress parameter [42]. In this work, Λ is defined as:

$$\Lambda = Y_C(z_\Lambda) \quad , \quad z_\Lambda = 0.1 \quad (3.12)$$

Where Y_C is the progress variable defined in Equation 3.14. The choice of z_Λ is arbitrary. In this work, we chose to place it near the most reactive mixture $z_{mr} = 0.04$. This definition of Λ allows a unique mapping of instantaneous flamelet solutions in the (t, a) -space (time and strain rate) to the $(\Lambda, \chi_{z,st})$ -space (progress parameter and stoichiometric scalar dissipation rate).




Figure 3.5: Evolution of the gas temperature in a flamelet of strain rate $a = 10 \text{ s}^{-1}$, plotted in mixture fraction space. Solutions are plotted every 10 ms, darker with increasing time.

The progress variable is typically defined as a linear combination of the mass fractions Y_k of the reactive species:

$$Y_c = \sum_{k=1}^{N_{\text{species}}} \alpha_k Y_k \quad (3.13)$$

The "correct" way to define the weights α_k is still an open problem for complex fuels and reactions [77], but it is much more manageable for simpler reactions such as $\text{H}_2\text{-O}_2$ combustion. The progress variable, in our case, must be defined such that:

1. It is an increasing function of time during ignition
2. Its derivative is non-zero at $t = 0$. This condition ensures that auto-ignition can occur in the simulation

To enforce the first condition, it is common to give a strictly positive weight to at least one product of the reaction: in this case, H_2O . However, this is not enough to enforce the second condition, because the final products of the reaction are not immediately produced. So-called intermediate species are first produced, and are in turn consumed to produce the final products. Therefore, an additional negative weight on at least one reactant ensures a strict, non-zero increase of the progress variable from the very start of the reaction. In our case, we chose to define Y_c as:

$$Y_c = Y_{\text{H}_2\text{O}} - 0.3778Y_{\text{O}_2} \quad (3.14)$$

The weight of O_2 could have any negative value, but we chose this one to set $Y_c = 0$ for $Z = 0$ and $Z = 1$. It also leads, that, for all Z values, the value $Y_{C,u}$ of the progress variable in the unburnt conditions is null: $\forall Z, Y_{C,u}(Z) = 0$. This is a purely esthetical consideration and this arbitrary choice has no impact on the transport of the progress variable or the construction of the table.

For turbulent flames in LES, a relation equivalent to Equation 3.11 must be defined for filtered quantities. A presumed joint PDF \tilde{P} for the mixture fraction, reaction progress

parameter, and stoichiometric scalar dissipation rate is defined, so that:

$$\tilde{\phi} = \int \int \int \mathcal{F}_\phi(z^*, \Lambda^*, \chi_{z,st}^*) \tilde{P}(z^*, \Lambda^*, \chi_{z,st}^*) dz^* d\Lambda^* d\chi_{z,st}^* \quad (3.15)$$

Λ is statistically independent of z and $\chi_{z,st}$ by definition, and we assume that z and $\chi_{z,st}$ are also independent. This gives:

$$\tilde{P}(z^*, \Lambda^*, \chi_{z,st}^*) = \tilde{P}(z^*) \tilde{P}(\Lambda^*) \tilde{P}(\chi_{z,st}^*) \quad (3.16)$$

Furthermore, we assumed a beta-PDF for z [46, 118], a Dirac distribution for $\chi_{z,st}$ [83], and we chose a Dirac distribution for Λ . This last choice is a significant simplification compared to Ihme and See [43], who chose a statistically most likely distribution for the progress parameter. This greatly simplified the implementation of the model in the AVBP solver, reduced the computer memory taken by the table, and reduced the computational time required to run a simulation. Most importantly, as shown in Section 3.3.2, this choice did not prevent the simulation from accurately reproducing the experimental data. With these assumptions, we get:

$$\tilde{P}(z^*, \Lambda^*, \chi_{z,st}^*) = \beta(z^*, \tilde{z}, z_v) \delta(\Lambda^* - \tilde{\Lambda}) \delta(\chi_{z,st}^* - \widetilde{\chi_{z,st}}) \quad (3.17)$$

Joining Equations 3.15 and 3.17 gives:

$$\tilde{\phi} = \mathcal{G}_\phi(\tilde{z}, z_v, \tilde{\Lambda}, \widetilde{\chi_{z,st}}) \quad (3.18)$$

However, the definition of $\tilde{\Lambda}$ makes it difficult to transport. The progress variable \widetilde{Y}_C is easier to transport, and its normalization \tilde{C} is used as a coordinate of the table, using the unique transformation \mathcal{C} such that:

$$\tilde{C} = \mathcal{C}(\tilde{\Lambda}, \tilde{z}, z_v, \chi_{z,st}) \quad (3.19)$$

With C defined as:

$$C = \frac{Y_C - Y_{C,u}}{Y_{C,eq} - Y_{C,u}} \quad (3.20)$$

Where $Y_{C,eq}$ and $Y_{C,u}$ are respectively the values taken by the progress variable at equilibrium and in unburnt conditions.

We can also replace the subgrid variance of mixture fraction z_v with the mixture fraction segregation S_z , defined in Eq. 3.25. Finally, this gives, for every quantity ϕ the relation:

$$\tilde{\phi} = \mathcal{T}_\phi(\tilde{z}, S_z, \tilde{C}, \widetilde{\chi_{z,st}}) \quad (3.21)$$

Where the function \mathcal{T}_ϕ is precomputed and tabulated on the mesh specified in Table 3.2.

Table 3.2: Discretization of the UFPV table

Variable	Discretization type	Num. of points	Bounds
\tilde{z}	Regular	400	[0, 1]
S_z	Geometric (increasing steps)	20	[0, 1]
\tilde{C}	Double-geometric (increasing then decreasing steps)	153	[0, 1]
$\widetilde{\chi_{z,st}}$	Geometric (increasing steps)	10	[1.6, 555]

Special attention is given to the discretization in \tilde{C} to ensure a correct prediction of the source term at the very start of the auto-ignition. The \tilde{C} -mesh is refined towards both

the lower and upper bounds. For \tilde{C} near 0, the mesh needs to be very fine, because the reaction is exponential in its first instants. Therefore, any error committed early on would have important repercussions on the predicted auto-ignition delay-time. Conversely, for \tilde{C} near 1, we must correctly refine the reaction termination, especially in the richer mixtures, where every quantity still varies while they are already nearly settled at z_Λ .

Figure 3.6 provides a verification that the \tilde{C} -mesh is adequate, by comparing the auto-ignition of a homogeneous reactor simulated with both detailed chemistry and tabulated chemistry using the same C -mesh as the UFPV table.




Figure 3.6: Comparison of detailed and tabulated chemistry on the auto-ignition of a homogeneous reactor at $z = z_{mr} = 0.04$. The evolution of temperature is plotted in plain line for the detailed chemistry and in dotted line for the tabulated chemistry.

A selection of slices of this table for the hydroxyl mass fraction Y_{OH} are plotted in Figure 3.7 as a function of the Favre-filtered mixture fraction \tilde{z} and of the Favre-filtered normalized progress variable \tilde{C} , for a few strain rate and mixture fraction segregation values. Note the significant shape change and overall amplitude decrease with increasing mixture fraction segregation. This species, which is a marker of the reaction, has a very similar behavior as the source term of progress variable. Slices of this source term, and other selected species mass fractions and thermophysical properties can be found in Annex B.

At each time step and at every point, a set of coordinates must be computed

As shown in Equation 3.21, the UFPV table used in this work involves four coordinates:

- The Favre-filtered mixture fraction balance \tilde{z}
- The mixture fraction segregation S_z
- The Favre-filtered scalar dissipation rate at stoichiometry $\widetilde{\chi_{z,st}}$
- The normalized Favre-filtered progress variable \tilde{C}

\tilde{z} , S_z , and \tilde{C} are found by solving transport equations, while $\widetilde{\chi_{z,st}}$ is computed using an algebraic equation. These equations are detailed in the following:




Figure 3.7: Slices of the filtered mass fraction of OH \widetilde{Y}_{OH} as a function of the filtered mixture fraction \tilde{z} and of the filtered normalized progress variable \tilde{C} , for different strain rates a and different mixture fraction segregation values S_z . Note the evolution of the vertical scale

Favre-filtered mixture fraction balance \tilde{z}

\tilde{z} is directly transported by the following equation:

$$\frac{\partial \bar{\rho} \tilde{z}}{\partial t} + \nabla \cdot (\bar{\rho} \tilde{\mathbf{u}} \tilde{z}) = \nabla \cdot \bar{\rho} D \nabla \tilde{z} - \nabla \cdot \tau_z \quad (3.22)$$

Where D is the mass diffusivity, determined using the unity Lewis number assumption: $Le = Sc/Pr = 1$. The subgrid flux of mixture fraction τ_z is defined as:

$$\tau_z = \overline{\rho \mathbf{u} \tilde{z}} - \bar{\rho} \tilde{\mathbf{u}} \tilde{z} \quad (3.23)$$

This subgrid flux is modeled using a turbulent diffusivity $D_t = Sc_t \nu_t$, where Sc_t is the turbulent Schmidt number. Although the determination of the optimal value for the turbulent Schmidt number remains an open question [64], we used the default value proposed by the AVBP software: $Sc_t = 0.6$, which is well within the range of acceptable values found in the literature.

Finally, τ_z is modeled using the SGDH:

$$\tau_z = -\bar{\rho} D_t \nabla \tilde{z} \quad (3.24)$$

Mixture fraction segregation S_z

$$S_z = \frac{z_v}{\tilde{z}(1 - \tilde{z})} \quad (3.25)$$

Where $z_v = \tilde{z}^2 - \tilde{z}^2$ is the subgrid variance of the mixture fraction and respects the following transport equation:

$$\frac{\partial \bar{\rho} z_v}{\partial t} + \nabla \cdot (\bar{\rho} \tilde{\mathbf{u}} z_v) = \nabla \cdot (\bar{\rho} D \nabla z_v) - \nabla \cdot \tau_{z_v} - 2\tau_z \cdot \nabla \tilde{z} - \overline{s_{\chi_z}} \quad (3.26)$$

The subgrid flux of mixture fraction variance τ_{z_v} , is defined as:

$$\tau_{z_v, i} = \bar{\rho} \left(\widetilde{u_i z^2} - \tilde{u}_i \tilde{z}^2 \right) - 2\tilde{z} \tau_{z_i} \quad (3.27)$$

And it is modeled using the SGDH:

$$\tau_{z_v} = -\bar{\rho} D_t \nabla z_v \quad (3.28)$$

Finally, the subgrid mixture fraction scalar dissipation rate $\overline{s_{\chi_z}}$ is modeled using a linear relaxation assumption:

$$\overline{s_{\chi_z}} = C_D \bar{\rho} \frac{z_v}{\Delta^2 / \nu_T} = \frac{C_{\chi_z} C_\varepsilon}{C_u} \bar{\rho} \frac{z_v}{\Delta^2 / \nu_T} \quad (3.29)$$

Caution must be exercised when dealing with this equation, for the definition of C_D is not standard in the literature. Other works, such as [23] introduce a factor 2 before $\overline{s_{\chi_z}}$ in Equation 3.26, and a factor 0.5 in the definition of C_D . For the sake of simplicity, we kept Ihme's definition [43] which removes these unnecessary factors. With this definition, we used a nominal value of $C_D = 20.0$ in the following. The value chosen for this coefficient is discussed in Section 3.6.2.

Favre-filtered stoichiometric scalar dissipation rate $\widetilde{\chi_{z,st}}$

The stoichiometric scalar dissipation rate $\chi_{z,st}$ is computed from the following algebraic equation:

$$\chi_{z,st} = \chi_z \frac{F(z_{st})}{F(z)} \quad (3.30)$$

With the local scalar dissipation rate χ_z defined as:

$$\chi_z = 2D |\nabla z|^2 \quad (3.31)$$

And:

$$F(z) = \frac{1}{\pi} \exp\left(-2 [\operatorname{erfc}^{-1}(2z)]^2\right) \quad (3.32)$$

This function can also be evaluated with the following equivalent expression, as seen in [117]:

$$F(z) = \frac{1}{\pi} \exp\left(-2 [\operatorname{erf}^{-1}(2z - 1)]^2\right) \quad (3.33)$$

The stoichiometric scalar dissipation rate is equivalent to the strain rate a , which follows the relation [80]:

$$\chi_{z,st} = aF(z_{st}) \quad (3.34)$$

In LES, χ_z is not accessible because it is not a filtered quantity. Instead, we can compute the Favre-filtered scalar dissipation rate $\widetilde{\chi_z}$ with:

$$\widetilde{\chi_z} = 2D |\nabla \tilde{z}|^2 + \frac{\overline{s\chi_z}}{\bar{\rho}} \quad (3.35)$$

The Favre-filtered stoichiometric scalar dissipation rate is then computed with:

$$\widetilde{\chi_{z,st}} = \frac{\widetilde{\chi_z}}{\mathcal{F}(\tilde{z}, S_z)} \quad (3.36)$$

Where \mathcal{F} is defined as:

$$\mathcal{F}(\tilde{z}, S_z) = \int_0^1 \frac{F(z^*)}{F(z_{st})} \tilde{P}(z^*) dz^* \quad (3.37)$$

Favre-filtered normalized progress variable \tilde{C}

$$\tilde{C} = \frac{\tilde{Y}_c - Y_{C,u}}{Y_{C,eq} - Y_{C,u}} \quad (3.38)$$

Where \tilde{Y}_c is the Favre-filtered progress variable and respects the following balance equation:

$$\frac{\partial \bar{\rho} \tilde{Y}_c}{\partial t} + \nabla \cdot (\bar{\rho} \tilde{\mathbf{u}} \tilde{Y}_c) = \nabla \cdot (\bar{\rho} D \nabla \tilde{Y}_c) - \nabla \cdot \tau_{Y_c} + \bar{\rho} \widetilde{\omega}_{Y_c} \quad (3.39)$$

Where $\widetilde{\omega}_{Y_c}$ is the production term of the progress variable, and is read from the table. The subgrid flux of progress variable, τ_{Y_c} , is defined as:

$$\tau_{Y_{ci}} = \bar{\rho} \left(\widetilde{u_i Y_c} - \tilde{u}_i \tilde{Y}_c \right) \quad (3.40)$$

It is then modeled with the SGDH:

$$\tau_{Y_c} = -\bar{\rho}D_t\nabla\tilde{Y}_c \quad (3.41)$$

Finally, $Y_{C,eq}$ and $Y_{C,u}$ are the values taken by the progress variable respectively at equilibrium and in unburnt conditions. They are functions of \tilde{z} , S_z , and $\widetilde{\chi_{z,st}}$ only. Hence, they can be tabulated along with other values of interest and retrieved for the computation of \tilde{C} .

3.2.3 Numerical choices

3.2.3.1 Computational domain and spatial mesh

The computational domain is a cylinder in which the fuel inlet tube protrudes. It is presented in Figure 3.8. It is large enough that the jet is unaffected by the domain boundaries in the region of interest. Figure 3.8 (a) presents an overview of the domain and the gas flows involved.

The mesh is composed of tetrahedral elements defining 5.5 million nodes. An axial slice of this mesh is presented in Figure 3.8 (b), with a zoom on the fuel inlet zone in Figure 3.8 (c). The domain is 100 diameters long and 65 diameters wide. The cone, where the jet develops and the flame is located, is finely discretized.

Stability issues were encountered when trying to simulate the inlet tube itself, as acoustic waves were growing uncontrollably in this tube. Therefore, it was decided to remove the tube from the computational domain. Instead, the fuel inlet boundary is placed on a surface leading directly to the main volume of the computational domain. It is highlighted by the red disk on Figure 3.8 (d). The lips of the inlet tube are highlighted in blue in this same figure.

3.2.3.2 Boundary conditions

Proper boundary conditions must be chosen for the different surfaces enclosing the computational domain, which are:

- Fuel inlet
- Co-flow inlet
- Inlet tube walls
- Side walls
- Outlet

For all fluid boundaries, i.e fuel inlet, co-flow inlet, and outlet, Navier-Stokes Characteristics Boundary Conditions (NSCBC) [84] are used. These conditions do not enforce any one quantity at the boundary, but rather control the waves crossing the boundaries, such that:

- The mean quantities are kept at their target values
- Acoustic waves can be generated or let through without reflection
- When appropriate, turbulence can be injected.




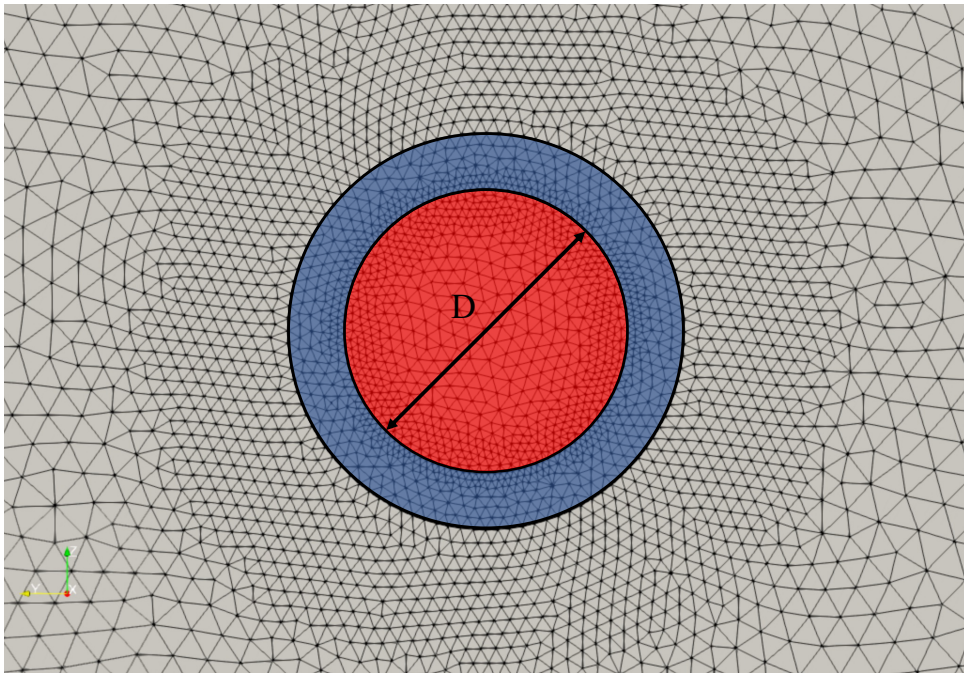


Figure 3.8: Mesh used in the LES of the Cabra flame. (a): 3D view and fluid boundary conditions. (b): Axial slice in the y-normal plane



(c)



(d)

Figure 3.8: Mesh used in the LES of the Cabra flame (cont.). (c) Zoom on the axial y -normal slice near the fuel inlet. (d) Zoom on the fuel inlet in the x -normal plane. The fuel inlet is highlighted in red and the pipe lips are highlighted in blue.

Fuel inlet

The fuel inlet is probably the most critical boundary condition, because it conditions the shape of the jet, especially regarding the length of the potential core. A non-reflective (NSCBC) inlet condition is used, with turbulence injection. The injection turbulence method applied is described in [100]. The quantities controlled are the velocity, temperature, and composition.

To emulate the effect of the inlet tube on the velocity and turbulence profile of the fuel inlet, a RANS simulation was performed on a tube of the same diameter, with a length $L = 200 D$ to ensure the convergence of the various profiles of the flow. It was run using the Ansys Fluent software, with a $k - \omega$ turbulence model and resolved flow at the wall ($y^+ = 1$).

The resulting profiles for the axial velocity U and root mean square velocity U_{rms} are presented in Figures 3.9 and 3.10. The velocity profile was kept unaltered and the velocity values at the mesh points of the boundary were interpolated from this profile.

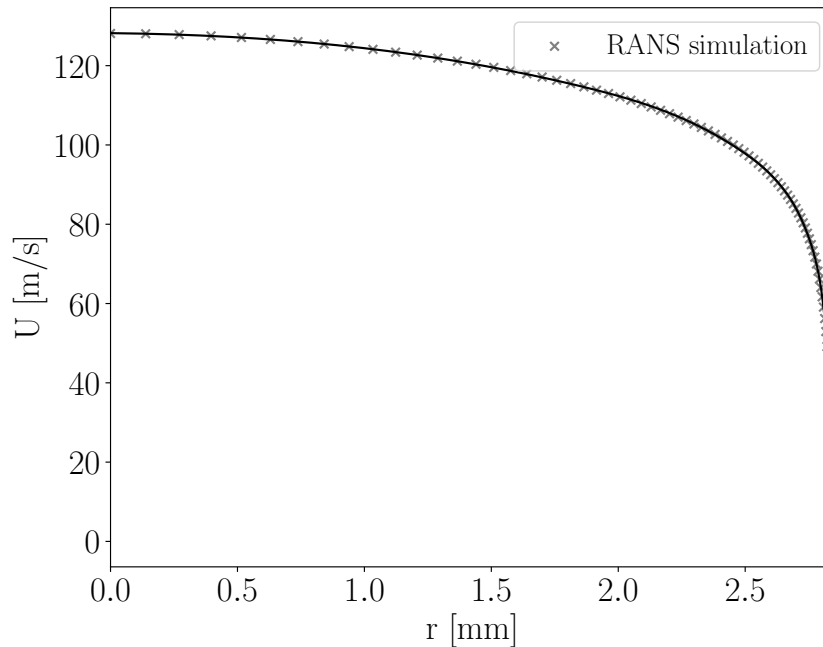


Figure 3.9: Velocity profile obtained from the RANS simulation

However, using the U_{rms} profile resulted in a much too short potential core. This may be explained by the fact that the profile is under-resolved in our LES, in particular near the wall. The turbulent intensity is thus interpolated only from the parts of the profile where it is high, ignoring the area close to the wall where the turbulent intensity plummets. This may lead to a premature destabilization of the jet and a shortening of the potential core. This is an important problem because it perturbs the mixing process, which will in turn modify the flame lift-off height. Because the most important feature to reproduce is the potential core length, and because the intensity profile only affects this critical length, an arbitrary profile can be used for U_{rms} as long as the potential core length fits the experiment. Therefore, an empirical, uniform value of $U_{rms} = 0.5$ m/s was used. This arbitrary choice of a uniform U_{rms} profile does not affect the flow beyond the potential core, because the flow becomes independent of the turbulence injection conditions once it reaches the developed jet zone.




Figure 3.10: Root mean square velocity profile obtained from the RANS simulation (in black), and empirical value found to fit the potential core length with the experiment (in red)

Co-flow inlet

For the co-flow inlet, a non-reflective condition was again used. This time, however, due to the low mean velocity (3.5 m/s) compared to the jet velocity (107 m/s), a flat velocity profile without turbulence injection was used. Furthermore, we did not simulate all the small laminar flames present in the experiment setup, but rather injected directly the equilibrium (post-combustion) mixture, as described in Section 3.1.

Inlet tube walls

The inlet tube walls are defined as adiabatic walls with a wall law. It sets the normal heat flux and velocity to zero and computes the wall shear stress according to the logarithmic law of the wall.

Side walls

The side walls of the domain are defined as adiabatic slip walls. They are placed far enough from the jet that they do not influence the central jet.

Outlet

The co-flow outlet is a partially reflective (relaxed) outlet condition with controlled pressure, which acts as a high-frequency filter for the acoustic waves crossing the boundary.

3.2.4 Numerical methods

The numerical scheme applied is the Two-step Taylor-Galerkin C (TTGC) [15], which is third-order in space and time. It emphasizes low-dissipation and dispersion. Although

about twice more expensive as a Lax-Wendroff scheme [75], it was chosen to minimize simulation error and to allow us to focus on other sources of uncertainties later on. We used the implementation of this numerical scheme provided by the AVBP solver [96]. We implemented in this solver the UFPV model described in Section 3.2.2.3, using the TTC formalism for compressible equations [3].

3.3 Results of the nominal simulation

3.3.1 Qualitative behavior of the simulated flame

The flame does ignite spontaneously, as described by Cabra et al. [12]. Furthermore, once the steady-state is reached, the correct behavior, explained in section 3.1.2 prevails: mixing happens in the periphery of the jet. Then some pockets of mixture auto-ignite, swell, fuse, and ultimately form a continuous flame front. *Nota bene*: reaching the steady-state does not mean that quantities become constant at every point of the flame. Rather, it means that, on average, the flame occupied the same position. A y-normal cut of an instantaneous solution is presented in Figure 3.11. It presents, on the upper part, the field of Favre-filtered mixture fraction \tilde{z} , and on the lower part, the field of Favre-filtered normalized progress variable \tilde{C} .




Figure 3.11: y-normal cut of the Favre-filtered instantaneous fields of mixture fraction and normalized progress variable

3.3.2 Comparison of the mean quantities with the experimental results

In this section, experimental measurements from Cabra et al. [11] are compared to the simulated average fields of the Cabra flame. Measurement uncertainties in this setup, using the laser Raman-Rayleigh-LIF method are estimated in [5].

For brevity's sake, not every available profile is shown and commented on. For a more exhaustive display of the comparison between experimental and simulation profiles,

please refer to Annex C. For clarity's sake, in this section, unless specified otherwise, the averaged Favre-filtered scalar $\{\tilde{\phi}\}$ will simply be denoted ϕ .

A synthetic quantity of the mixing of the jet is the mixture fraction of the gas z . The axial profile of this scalar is plotted in Figure 3.12, and several radial profiles are presented in Figure 3.13.




Figure 3.12: Axial profile of the mixture fraction Z in the LES, compared to the measurements from [11]

The axial profile of z shows the correct length of the potential core, as was expected from the work on the turbulent injection. However the initial slope of decrease (from $x/D \approx 6$ onwards) is too steep by about 30%. While this may be concerning, this issue has already been seen in works that computed similar jet flames in LES with AVBP, for example [97] and [117] for the methane Cabra flame. However, radial profiles of z , plotted in Figure 3.13, show a satisfactory agreement between experimental and simulated data. In particular, the evolution of the width of the jet is very well reproduced.

On the combustion side, the radial profiles of the reaction marker OH are plotted in Figure 3.14. A good agreement is found in the vicinity of the flame foot, i.e. for heights around $x = 10 D$. However, the mass fraction of OH is over-predicted higher up in the flame. This is visible in the profile taken at $x = 26 D$, and is best illustrated by Figure 3.15 which compares the entire fields of Y_{OH} in the simulation and experiment.

Finally, as shown in Annex C, Temperature profiles and major species' profiles show a correct flame temperature and consumption rate of reactants, meaning that most of the reaction has already occurred when the discrepancies in hydroxyl mass fraction appear. Overall, the agreement between experimental and simulation data is very satisfactory.

Finally, the simulation of the Cabra flame costs about 100k CPU hours on the Irene KNL supercomputer with Intel Xeon Phi 7250 1.4 GHz processors.

3.4 Convergence study

Convergence is paramount to putting any trust in the simulation's result. We can distinguish two kinds of convergence in this case:

1. Temporal mean convergence: The LES is transient, and we are interested in the mean quantities of the flame, such as its mean lift-off height. Therefore, we must




Figure 3.13: Radial profiles of the mixture fraction Z in the LES, compared to the measurements from [11]




Figure 3.14: Radial profiles of the mass fraction of OH Y_{OH} in the LES, compared to the measurements from [11]




Figure 3.15: Hydroxyl mass fraction mean field. Left: Experimental measurements from [11]. Right: This study. The white dots on the left-hand side denote the locations of multi-scalar measurements, reported in Figures 3.12, 3.13 and 3.14, as well as in Annex C for every other available scalar of interest

ensure that the fields have been averaged long enough for the mean quantities to stabilize.

2. Mesh convergence: The mesh must be fine enough so that its influence on the simulation is negligible. The repartition between resolved and subgrid quantities will of course evolve when the mesh is refined, but the total quantities must remain stable.

3.4.1 Temporal mean convergence

The quantities of interest are averaged over 7.5 ms of simulated time in the nominal simulation. This corresponds to roughly 7 flow-through times in the region covered by experimental measurements, i.e from the fuel inlet to 26 diameters in the axial direction.

The convergence of the averaged quantities can be assessed by comparing profiles of the simulated flame, averaged respectively on the nominal simulation time and double this time. The simulation will be considered as converged if no discrepancies are spotted. Figures 3.16 and 3.17 display such plots for Y_{OH} and \tilde{z} respectively. Other comparisons of profiles of species and quantities of interest are plotted in Annex D at several axial locations.

Additionally, since we are interested in the lift-off height of the Cabra flame, it is also interesting to quantify the convergence of this global property. To that end, a bootstrap estimate of the lift-off height is performed.

The lift-off height H is extracted from the averaged solution of the flame via the




Figure 3.16: Temporal convergence of the Hydroxyl mass fraction mean field.




Figure 3.17: Temporal convergence of the mixture fraction mean field.

function \mathcal{H} , as follows:

$$H = \mathcal{H}(\mathbb{E}[\psi]) \quad (3.42)$$

Where ψ represents the fields of the simulated flame. The ergodicity of the flow is leveraged to replace the statistical expectation of the fields by its temporal average. However, this temporal average can not be perfectly converged, leading in turn to an incomplete convergence of the predicted lift-off height. To characterize the convergence error on the lift-off height, a bootstrap procedure is proposed.

The successive instantaneous solutions of the flame (i.e. snapshots of the flame at each time step of the simulation) are dependent, making the straightforward bootstrap procedure inapplicable. Instead, a stationary block bootstrap scheme [85] is implemented: instantaneous solutions of the flame are averaged in N non-overlapping blocks of data.

Each of these blocks produces a partially converged Cabra flame ψ_i . This is illustrated in figure 3.18. The resulting partially converged flames are much less dependent than the original instantaneous solutions. Therefore, they are assumed independent in the following.

N bootstrap resamples are then drawn with replacement from this set of partially converged flames, and averaged to form a "simulated bootstrap sample" $\hat{\psi}_k$. Such a simulated bootstrap sample may be, for example (with $N = 10$):

$$\hat{\psi}_k = \frac{1}{N} \sum_{i \in I_k} \psi_i \quad , \quad I_k = \{1, 1, 3, 5, 6, 6, 6, 6, 8, 9\} \quad (3.43)$$

The lift-off height $\hat{H}_k = \mathcal{H}(\hat{\psi}_k)$ is computed for this simulated bootstrap sample. Finally, this "draw, average, and post-process" procedure is repeated a large number of times B , giving birth to the empirical bootstrap distribution of the simulated Cabra flame's lift-off height \hat{H} .




Figure 3.18: Illustration of the partial averaging of instantaneous solutions to produce independent blocks of data suitable for the bootstrap procedure

We used $N = 20$ blocks, each averaging 500 instantaneous solutions spanning 0.4 ms of simulated time, and a resampling size of $B = 1000$. This resulted in the bootstrap mean and standard deviation being respectively:

$$\hat{H} = 11.75D \quad ; \quad \sqrt{\sigma^2(\hat{H})} = 0.074D \approx 0.6\% \text{ of } \hat{H} \quad (3.44)$$

Therefore, with the time accumulation used for the nominal simulation, we are confident that the lift-off estimate will not be off by more than a percent, which is negligible before the variations we expect to witness in the following.

Finally, we can conclude that the average of the simulated flame is converged enough in our case.

3.4.2 Mesh convergence

The mesh convergence can be performed in two ways:

1. Verifying that most of the biggest turbulence scales are well resolved, meaning that the variance of any quantity of interest is mostly resolved rather than mostly modeled.
2. Verifying that the averaged solution of the flame does not change when the mesh is refined.

The first point can be addressed by comparing the portion of resolved (RV) and modeled (MV) mixture fraction variance. The total variance of mixture fraction Z_{var} is written:

$$Z_{var} = \underbrace{\{z_v\}}_{\text{Modeled variance}} + \underbrace{\{\tilde{z}^2\} - \{\tilde{z}\}^2}_{\text{Resolved variance}} \quad (3.45)$$




Figure 3.19: Comparison of resolved and modeled mixture fraction variance on the flame axis




Figure 3.20: Comparison of resolved and modeled mixture fraction variance on two radial profiles around the flame foot

Figures 3.19 and 3.20 show that the total variance of the mixture fraction is indeed mostly explained by its resolved component.

To address the second point, a new mesh was generated on the same computational domain, with a uniform factor of refinement over the mesh used for the nominal simulation. This new mesh featured 7.5 million nodes, i.e. about 1.3 times more nodes than the original mesh. Figure 3.21 features a selection of radial profiles comparing the simulated hydroxyl mass fractions on the original and finer mesh. More comparisons of profiles can

be found in Annex D. The profiles overlap for the most part, which indicates that the simulation is converged mesh-wise.

Finally, the lift-off height computed on the finer mesh is only 0.7% higher than the original, which is considered converged in our application.




Figure 3.21: Comparison of two hydroxyl mass fraction profiles for the original and finer mesh

3.5 Sensitivity of the simulation to the co-flow temperature

Figure 3.2 shows that the experimental flame is sensitive to the co-flow temperature, with the lift-off height exponentially increasing when the co-flow temperature decreases. To verify whether the LES of the Cabra flame reproduces this pattern, a set of four additional simulations were performed at varying temperatures, spanning the range of uncertainty reported by Cabra. One new UFPV table was generated for each of these conditions. Figure 3.22 shows the mean hydroxyl mass fraction Y_{OH} for each of these simulations. This figure does highlight a noticeable sensitivity of the flame to the co-flow temperature.

Figure 3.23 reprises the experimental data from Figure 3.2 to compare with the lift-off height predicted by the simulation. The nominal simulation, which features a co-flow temperature of 1045 K, is in good agreement with Cabra’s experiment. The trend given by [120] and [34] is also correctly reproduced. The exponential increase of the lift-off height with decreasing co-flow temperatures was unfortunately not observed in this study. However, we believe that simulations with co-flow temperatures of about 1000 K would show this increase, because the exponential increase of the auto-ignition delay-time with decreasing oxidizer temperature was observed with this kinetic mechanism during the uncertainty propagation studies presented in chapter 2, both for homogeneous auto-igniting reactors and non-premixed auto-igniting flamelets.

For reference, we present in Figure 3.24 a comparison of this LES with other numerical works from the literature. Again, our simulation fares well against other state-of-the-art simulations. Nota Bene: In this case, the results’ discrepancies are not associated




Figure 3.22: Hydroxyl mass fraction mean fields plotted for five simulations with varying co-flow temperature (1015, 1030, 1045, 1060 and 1075 K)




Figure 3.23: Lift-off height as a function of the co-flow temperature. Numerical simulations are compared with results from [11, 120, 34].

with co-flow temperature uncertainties but with different modeling choices, and modeling uncertainties (chemical kinetics, turbulence, turbulent combustion).

Finally, it must be noted that the measurement uncertainty on the co-flow temperature forbids any categorical statement about whether our simulation is better or worse than any other simulation that falls within the uncertainty range reported by Cabra.




Figure 3.24: Lift-off height as a function of the co-flow temperature. Numerical simulations from this thesis are compared with other numerical results from [105, 13, 35, 74, 78].

3.6 Sensitivity to the numerical setup

3.6.1 Sensitivity to the chemical kinetics

In chapter 2, we showed that the auto-ignition delay times of homogeneous reactors and non-premixed flamelets were very sensitive to the kinetic mechanism employed to model the chemical reactions. Different kinetic mechanisms lead to significantly different results. Kinetic mechanisms which reported uncertainties, such as Konnov (2019) [55], allowed the determination of wide probability distributions of the auto-ignition delay. Consequently, we can expect a noticeable sensitivity of the lift-off height - and therefore of the flame shape - to the chemical kinetics.

3.6.2 Sensitivity to the mixture fraction variance dissipation coefficient

The flame-turbulence interaction model used in this LES contains the coefficient C_D , also written [41]:

$$C_D = C_{\chi_z} \frac{C_\varepsilon}{C_u} \quad (3.46)$$

C_{χ_z} is the mechanical-to-scalar time scale ratio:

$$C_{\chi_z} = \frac{\tau_u^{\text{sgs}}}{\tau_z^{\text{sgs}}} = \frac{\{k^{\text{sgs}}\} \{\widetilde{\chi_z^{\text{sgs}}}\}}{\{\varepsilon^{\text{sgs}}\} \{\widetilde{z''^2}\}} \quad (3.47)$$

And C_ε and C_u are defined as:

$$C_\varepsilon = \{\varepsilon^{\text{sgs}}\} \frac{\Delta}{\{k^{\text{sgs}}\}^{3/2}}, \quad C_u = \frac{\{\nu_t\}}{\Delta \{k^{\text{sgs}}\}^{1/2}} \quad (3.48)$$

Where Δ is the LES filter size, k^{sgs} is the subgrid-scale turbulent kinetic energy and ε^{sgs} is the subgrid-scale turbulent energy dissipation rate.

Although there is a consensus on the proper value to choose for other model coefficients like the Sigma model constant C_σ , this mixture fraction variance dissipation coefficient takes significantly different values in previous works. Domingo et al. [23] or Vicquelin [117] set its value at $C_D = 2.0$, for example, while Ihme and See [43] use $C_D = 4.0$. Ihme's thesis [41] (pp. 33-37) gives more details on how to best determine this coefficient for a given simulation. Ihme shows that the value of C_D is dependent on the Taylor-scale Reynolds number, on the filter width, and on the Kolmogorov length scale. This makes the ideal C_D value dependent on the flow configuration and the local mesh size. Therefore, any attempt to define a global C_D value for a given simulation must be a compromise designed to minimize the discrepancy between the experiment and the simulation.

- A value of $C_{\chi_z} = 2.0$ is proposed by Ihme for typical LES filter size, which is in agreement with other propositions in the literature ([80, 29]).
- Ihme also shows that the ratio C_ε/C_u is in the range $\pi^2 \leq C_\varepsilon/C_u \leq 20$ for typical LES filter size, with an asymptotic value of π^2 for large Reynolds numbers.

In this study, a value of $C_D = 20$ ($C_{\chi_z} = 2.0$ and $C_\varepsilon/C_u = 10$) was found to lead to a satisfying agreement between the simulation and Cabra's experimental data. However, given the wide variety of values found in the literature, between theoretical arguments and actual values implemented, we deemed it important to check the sensitivity of the macroscopical behavior of the flame to this coefficient.

Figure 3.25 presents the mean field of OH mass fraction Y_{OH} computed with different mixture fraction variance dissipation coefficient values. We witnessed a significant sensitivity of the flame shape, and in particular of the flame lift-off, to the value of C_D .




Figure 3.25: Hydroxyl mass fraction mean fields plotted for three simulations of the nominal Cabra flame, for different values of C_D (2.0, 20.0 and 40.0)

This sensitivity was expected because the equilibrium value of z_v is proportional to C_D^{-1} . This is deduced from equations 3.26 and 3.29. Therefore, as a general trend, the higher C_D , the lower the subgrid mixture fraction variance. Additionally, less subgrid variance entails more reaction source term. This can be seen, indirectly, in Figure 3.7, or directly in Annex B. Finally, a higher reaction source term leads to a shorter lift-off height. This is illustrated in Figure 3.26, which displays the subgrid mixture fraction variance, progress variable source term, and hydroxyl mass fractions for the extreme values of C_D considered in the study.




Figure 3.26: From left to right: Subgrid mixture fraction variance, Progress variable source term and Hydroxyl mass fraction mean fields plotted for two different simulations of the nominal Cabra flame, for the extreme values of C_D (2.0 and 40.0)

3.7 Takeaway points

- A Large Eddy Simulation of the Cabra flame was performed.
- The UFPV framework was implemented in the AVBP solver to tackle the modelization of the turbulent combustion in the non-premixed auto-igniting flamelet regime.
- Although some discrepancies between the experimental measurements and the simulation were found, the LES of the Cabra flame was deemed satisfactory overall, in particular in its reproduction of the lift-off height.
- In light of the uncertainties on the co-flow temperature and the chemical kinetics, and the substantially different values of mixture fraction variance dissipation coefficient found in the literature, the deterministic approach to simulation is not sufficient to ascertain the agreement between the experiment and simulation of the Cabra flame.
- Additional work is therefore required, and Chapter 4 tackles this issue with the framework of uncertainty propagation.

Chapter 4

Quantifying the uncertainty of the simulated Cabra flame's lift-off height

This chapter investigates the uncertainties introduced in the LES of the Cabra flame and their influence on its lift-off height. It does not aim to be an exhaustive uncertainty propagation, meaning that we do not claim to have taken into account every input uncertainty. The first part of this chapter, (Sections 1-4) details the propagation of uncertainties directly linked to chemical kinetics. The second part introduces a new uncertain parameter from the turbulence-combustion model and performs again the uncertainty propagation with the addition of this new coefficient.

Section 1 introduces the kinetic and temperature propagated in the first study. Section 2 reduces the input uncertain dimension to a manageable number. Section 3 details the computer experiment and section 4 presents the results of the uncertainty propagation. Finally, section 5 repeats the same procedure with the addition of the uncertain dissipation coefficient of the mixture fraction variance.

4.1 The input uncertainties

4.1.1 Kinetic mechanism

The combustion model used in the LES of the Cabra flame relies on a kinetic mechanism, that is, a set of elementary chemical reactions and laws to determine their associated reaction rates. Chapter 2 outlined the uncertainties associated with these reaction rates and the effects that they have on the auto-ignition delay time of a reactive mixture. As explained in Chapter 3, the driving process of flame stabilization in the H₂ Cabra setup is auto-ignition. Therefore, uncertainties in the kinetic mechanism must have an impact on the flame lift-off height.

Similarly to the study presented in Chapter 2, we use the boiled-down Konnov (2019) mechanism [55] detailed in Annex A. This mechanism models the reaction rates of each reaction j with an Arrhenius law:

$$k_j = A_j T^{\beta_j} \exp\left(\frac{-Ea_j}{RT}\right) \quad (4.1)$$

Where T is the gas temperature, Ea_j is the activation energy of reaction j , β_j is a corrective dimensionless number and R is the universal gas constant. A_j is the pre-exponential factor of the Arrhenius law and is modeled as a log-normal random variable:

$$\xi_j = \frac{\log(A_j/A_j^0)}{\frac{1}{3}\log(UF_j)} \sim \mathcal{N}(0, 1) \quad (4.2)$$

Where A_j^0 is the nominal value of A_j and UF_j is its "uncertainty factor". ξ_j is a unit normal random variable.

The boiled-down Konnov mechanism has 31 reactions and therefore 31 uncertain pre-exponential factors are considered in the following.

The results of the following study are of course sensitive to the choice of kinetic mechanism and considered uncertainties. In particular, the fact that only the pre-exponential constants are assumed uncertain is a modelization choice from the authors of the mechanisms and their sources. We choose to use Konnov's (2019) mechanism [55] because it is state-of-the-art at the time of this thesis, but the methodology would be compatible with other mechanisms should better inputs become available in the future.

4.1.2 Co-flow temperature

As stated in the original publication by Cabra et al. [11], the co-flow temperature T_{cof} can be measured with an uncertainty of about $\pm 3\%$, which corresponds to $\pm 30\text{K}$. This uncertainty is dependent on the measurement method used (the laser Raman-Rayleigh-LIF method in the case of Cabra et al. [11] and Wu et al. [120]), but Gordon et al. [34] report the same uncertainty using a K-type thermocouple.

Figure 3.22, in Chapter 3, shows that such uncertainty on the co-flow temperature introduces significant variability in the lift-off height. Therefore, it must be taken into account to compare simulations with the experimental results.

As far as we are aware, no more information is available about this uncertainty. The least-informative probability distribution - the uniform distribution between $T_{\text{Min}} = 1015\text{ K}$ and $T_{\text{Max}} = 1075\text{ K}$ - is consequently chosen for the co-flow temperature:

$$T_{\text{cof}} \sim \mathcal{U}(T_{\text{Min}}, T_{\text{Max}}) \quad (4.3)$$

To represent this uncertainty in a normalized way, a unit uniform random variable ξ_T is introduced:

$$\xi_T = \frac{T_{\text{cof}} - T_{\text{Min}}}{T_{\text{Max}} - T_{\text{Min}}} \sim \mathcal{U}(0, 1) \quad (4.4)$$

In total, 32 uncertain variables are accounted for in the study: $\{\xi_1, \dots, \xi_{31}, \xi_T\}$.

4.2 Reduction of the input uncertain dimension

Large Eddy simulations of the Cabra flame are expensive (about 100k CPU hours on the Irene KNL supercomputer) which means that a brute-force, Monte-Carlo approach is unaffordable in this case. A surrogate-based approach will therefore be preferred. However, Chapter 1 pointed out that building a surrogate model in high dimensions requires a large number of observations, which, again, are not affordable.

Therefore, an a priori uncertain dimension reduction is mandatory. In this study, we chose to use the "dominant variables" approach (cf Section 1.5.1.2). To this end, a

global sensitivity analysis has to be performed to sort out important variables from less important ones. That is, to highlight the variables whose uncertainties lead to the greatest uncertainty in the quantity of interest.

Again, global sensitivity analysis cannot be implemented directly on the 3D flame lift-off height for tractability reasons. Instead, a representative physical problem of much reduced computational cost must be used.

4.2.1 A representative reduced physical problem

Auto-ignition being the dominant stabilization mechanism for the Cabra flame in most of the operating conditions [120], it can be assumed that the lift-off height of the flame is first-order dependent on the auto-ignition delay (IDT) of the most reactive mixture in the jet [45]. In Figure 4.1, we compare the normalized auto-ignition delay time of the most reactive mixture to the normalized lift-off height for different co-flow temperatures, with different nominal kinetic mechanisms. The auto-ignition is detected in the same way as lift-off is detected in the 3D flame, i.e. when the hydroxyl mass fraction first reaches $Y_{\text{OH}} = 600$ ppm. The nominal initial conditions of the reactor are given in Table 4.1. Figure 4.1 shows a very similar response of the auto-ignition delay of 0D reactors and the Cabra lift-off height in actual LES to both temperature and kinetic mechanism changes.

0D auto-ignition simulations are, therefore, a representative physical problem for most of the Cabra flame’s operating conditions. Furthermore, each 0D simulation costs about 0.1 CPU-second, making them suitable for a Monte-Carlo-based sensitivity analysis.

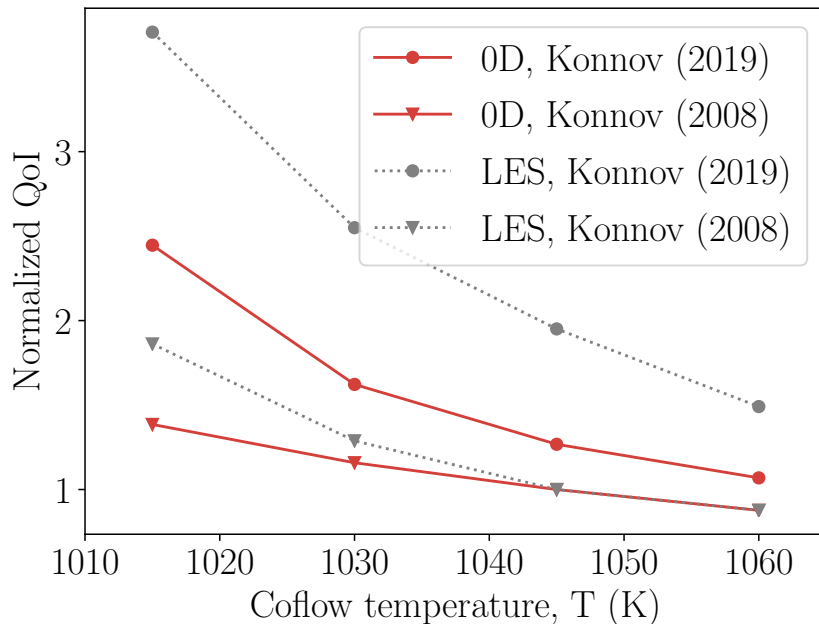


Figure 4.1: Comparison of the evolution of IDT (0D) to lift-off height (LES) for 2 kinetic mechanisms. All auto-ignition delays are log-preconditioned and normalized by the IDT at a co-flow of 1045K with the Konnov (2008) mechanism. All lift-off heights are normalized by the lift-off height at a co-flow of 1045K with the Konnov (2008) mechanism.

A Monte-Carlo study is conducted to account for the uncertainties of the co-flow temperature and the kinetic mechanism on the 0D auto-ignition delay. $2^{17} \approx 100\text{k}$ samples

Table 4.1: Initial conditions for the 0D reactor at the most reactive mixture $z_{mr} = 0.04$ [68] in the nominal case (corresponding to $T_{\text{cof}} = 1045$ K)

P (Pa)	101325	Y_{O_2}	0.1669	$Y_{\text{H}_2\text{O}}$	0.0626
T (K)	1011.3	Y_{H_2}	0.0009	Y_{N_2}	0.7696
T : Temperature; P : Pressure; Y : Mass fraction					

are drawn from the 32D Sobol' sequence and a homogeneous reactor autoignition is simulated for each sampled condition. Figure 4.2 shows the temporal evolution of the mixture temperature for the first 20 realizations of the homogeneous reactor. Table 4.2 records some statistics of this study. The uncertainty on the QoI is considerable, as shown by the 95% confidence interval, which spans an entire order of magnitude. Moreover, the important discrepancy between the mean of the distribution and the nominal case suggests substantial non-linearities in the underlying physical phenomenon.




Figure 4.2: Evolution of the temperature, as a function of time, of the first 20 realizations of the homogeneous reactor

Table 4.2: Statistics of the Monte-Carlo study on the auto-ignition delay τ of the 0D reactor

Auto-ignition delay time τ (ms)			
mean	median	nominal	95% confidence interval
5.10	2.07	2.04	[0.99 , 21.9]

4.2.2 Global sensitivity analysis

The global sensitivity analysis method we chose is Sobol' sensitivity analysis [103], approximated using the PCE approach proposed by Schöbi and Sudret [94] (cf Section 1.4.2). The target quantity of the surrogate is the logarithm of the IDT $\log_{10}(\tau)$, to

enforce the positivity of the emulated auto-ignition delay time.

The PCE is determined using least square regression with an elastic net and 10-fold cross-validation based on $2^{17} \approx 100\text{k}$ Quasi Monte-Carlo samples. To attenuate the computational cost of determining the PCE weights in high dimensions, the PC basis is truncated at a maximum order of 3 and first-order interactions ($p = 3$, $q = 0.75$, see Figure 1.8). A summary plot of the PCE is given in Figure 4.3 as a means of visual validation of the surrogate.




Figure 4.3: Summary plot of the PCE used to compute the Sobol indices

The Sobol indices, plotted in Figure 4.4, highlight the dominance of two uncertainties in the problem:

- The co-flow temperature: As already exposed in Section 3.1.4 and Figure 4.1, the lift-off height of the Cabra flame and the auto-ignition delay time of the corresponding homogeneous reactor are highly sensitive to the co-flow temperature. The large uncertainty range associated with this variable makes it the most impacting uncertainty in this configuration. Its index stands alone in this study, dwarfing all effects from other uncertain parameters.
- The pre-exponential constant A_{12} : The 12th reaction, $\text{H} + \text{O}_2 = \text{OH} + \text{O}$, has been known for its very high sensitivity for a long time [7]. Because of this, it has since been extensively studied and, as a result, its uncertainty factor (1.1) is the lowest of Konnov’s kinetic mechanism. However, it remains the most impacting uncertain kinetic parameter.

4.2.3 Dimension reduction

Even if the PCE is truncated, the resulting Sobol indices are sufficient to rank variables and get an adequate estimate of the portion of variance explained by each variable or group of variables.




Figure 4.4: First- and second-order Sobol indices of the auto-ignition delay

In this study, we choose to truncate the uncertain space to ensure that 95% of the variance of the QoI is retrieved, based on the prediction of the PCE. Only two uncertain variables are needed to explain that portion of the variance of the QoI: ξ_T and ξ_{12} , corresponding to the uncertainty on the co-flow temperature and 12th reaction rate respectively.

The reduced uncertain space spanned by these two variables is resampled to validate the truncation. The resulting probability distribution of the QoI is compared to its original distribution in Figure 4.5. The shape of the distribution is well captured. The peak has both the correct position and the correct amplitude and the kurtosis of the distribution is also well reproduced.

Therefore, the uncertainties of the co-flow temperature and of the 12th reaction pre-exponential factor are enough to explain most of the variability of the auto-ignition delay-time of the homogeneous reactor corresponding to the most reactive mixture in the Cabra flame configuration.

4.2.4 Verification of the dimension reduction on a slightly more complex problem

The dimension reduction proposed in the previous section is valid for a homogeneous reactor. In this section, we aim, to verify if it still holds in a configuration more representative of the real Cabra flame. Therefore, we test it on a strained and auto-igniting non-premixed laminar 1D flame. As shown in Chapter 3, this flamelet is the basis of the combustion model used in the LES. Therefore, if the dimension reduction holds well for this configuration, we trust that it will also hold for the uncertain behavior of the lift-off height in the Cabra flame's LES.

A pair of Monte-Carlo experiments are conducted in the original, 32D space and in the




Figure 4.5: Comparison of auto-ignition delay probability distribution of the homogeneous reactor, between the reference sampling of the 32D uncertain space and the sampling of the reduced uncertain space spanned by ξ_T and ξ_{12}

reduced, 2D space, for reference and evaluation of the reduced uncertain space respectively.

The conditions are very similar to the flamelets presented in chapter 2. The only difference is the boundary temperature of the oxidizer size, which is slightly higher in the case of the Cabra flame. Boundary conditions in the nominal case are summarized in Table 4.3.

Table 4.3: Nominal boundary conditions for the auto-igniting 1D flame: temperature and species molar fractions. The pressure is set to 101 325 Pa.

Fuel side			Oxidizer side		
T_{fuel}	[K]	305	T_{ox}	[K]	1045
X_{H_2}	[-]	0.25	$X_{\text{H}_2\text{O}}$	[-]	0.1
X_{N_2}	[-]	0.75	X_{O_2}	[-]	0.15
			X_{N_2}	[-]	0.75

For the same reasons as in the study presented in Chapter 2, the strain rate a is set to $a = 10 \text{ s}^{-1}$ to ensure the ignition of the flamelet even for slow realizations of the uncertain mechanism or low initial temperatures.

The flamelet equations (Equation 2.1) are solved in the mixture fraction space on a uniform 400-point mesh using the 2nd-order stabilized explicit Runge-Kutta method (ROCK2) with error control [1]. The auto-ignition delay is determined when the mass fraction of the hydroxyl radical first reaches the threshold $Y_{\text{OH}} = 600 \text{ ppm}$ in the computational domain. The quantity of interest, again, is the logarithm of the auto-ignition delay time. The computation time per sample can rise up to a dozen minutes. The cost of these simulations is still moderate but no longer negligible like it was for the homogeneous reactor. For this reason, the PDFs are approximated using only $2^{11} = 2048$ Quasi-Monte-Carlo samples drawn from the 32D Sobol' sequence for the reference PDF

and 2^{11} more Quasi-Monte-Carlo samples from the 2D Sobol' sequence for the PDF in the reduced uncertain space.

The resulting PDFs are presented in Figure 4.6. Again, the two PDFs collapse for the most part, which indicates that the dimension reduction holds for strained auto-igniting non-premixed flamelets. Therefore, we are confident that the uncertain behavior of the Cabra flame will also mostly take place in the same reduced uncertain space: $\text{Span}(\xi_T, \xi_{A_{12}})$.




Figure 4.6: Comparison of auto-ignition delay probability distribution of the strained flame, between the reference sampling of the 32D uncertain space and the reduced uncertain space spanned by ξ_T and ξ_{12}

4.3 The computer experiment

4.3.1 Sampling budget, strategy, and surrogate choice

A single LES of the Cabra flame costs about 100k CPU hours on the Irene KNL supercomputer that we used. Our CPU time allotment was enough for 30 simulations. We chose to reserve a third of these simulations for validation, while the two other thirds were assigned to the surrogate training.

We use Quasi Monte Carlo sampling (a Sobol' sequence), to draw the 20 training samples. This choice is motivated by:

- The superiority of regression surrogate methods over projection methods observed in the preliminary study in Chapter 2 which disqualifies cubature sampling in this case
- The low-discrepancy property of the sequence, although we recognize that a set of 20 samples does not maximize this property

- The convenience of increasing the cardinality of the training set, if deemed necessary and if more resources were available later on

The 10 validation samples were drawn using Latin Hypercube sampling, because:

- It has a good space coverage
- Its randomness protects against potential sampling bias

Each of these samples is a Large Eddy Simulation of the Cabra flame, with a setup identical as the one presented in Chapter 3. Instantaneous solutions of these simulations are shown in Figure 4.7 for the nominal case and for the samples where the highest and lowest lift-off heights are found.




Figure 4.7: y-normal slices of the instantaneous fields of water vapor mass fraction for the nominal (left), lowest (center), and highest (right) observed flames in the training set

The mean fields of hydroxyl mass fractions, from which the lift-off heights are extracted, are plotted for every training sample in Figure 4.8.

A wide variety of flame shapes and heights is observed. From these mean fields, we compute the lift-off height as the first streamwise location where the hydroxyl mass




Figure 4.8: Mean fields of hydroxyl mass fraction for every training sample

fraction first reaches 600 ppm. A logarithmic preconditioning is then applied to these observations of the lift-off height in order to guarantee its positivity in the following.

Finally, we chose to use a Regression Kriging as a surrogate model because it gave the best results among tested methods in this case. The chosen surrogate's properties are given in Table 4.4 and the response surface itself is plotted in Figure 4.9. The other surrogate types that were investigated are detailed, plotted, and evaluated in Annex E.

Table 4.4: Properties of the surrogate model used for the Lift-off height in $\text{Span}(\xi_{A_{12}}, \xi_T)$

Surrogate Type	Regression Kriging
Trend	Linear
Kernel	Matérn 5/2
Length-scales (θ) determination	Maximum likelihood approach
Nugget	0.05




Figure 4.9: Response surface of the Regression Kriging, plotted alongside the training and validation samples

Figure 4.9 suggests a good agreement between the data and the surrogate. This agreement is quantified in the following section.

4.3.2 Surrogate validation

The first step of validating the surrogate is to produce a summary plot. While the model plot of Figure 4.9 informs on the overall shape of the response surface, it does not really frame the agreement between the surrogate and the observations. The summary plot, shown in Figure 4.10 drops any notion of surrogate shape to focus solely on the comparison between the observations and their approximations by the surrogate.




Figure 4.10: Summary plot of the Regression Kriging

The summary plot in Figure 4.10 shows that the Kriging surrogate reproduces correctly its training set. More importantly, the validation samples which, by definition, were not included in the training of the surrogate, are also correctly reproduced. This indicates that there is little to no overfitting in this surrogate.

To quantify the error of the Kriging surrogate relative to the validation samples, we use the error metrics introduced in Section 1.6: The Normalized Root Mean Squared Error (NRMSE), and the Normalized Maximum Absolute Error (NMAE):

$$\text{NRMSE} = \frac{\|\mathbf{y}_{\text{validation}} - \mathbf{y}_{\text{predicted}}\|_2}{\|\mathbf{y}_{\text{validation}}\|_2} = 1.4\% \quad (4.5)$$

$$\text{NMAE} = \frac{\|\mathbf{y}_{\text{validation}} - \mathbf{y}_{\text{predicted}}\|_\infty}{\|\mathbf{y}_{\text{validation}}\|_\infty} = 2.1\% \quad (4.6)$$

These error results are extremely satisfying, especially compared to those of the best PCE (presented in Annex E) which were about 4 times higher for both metrics. Finally, the Regression Kriging whose response surface is plotted in Figure 4.9 will be used in the following.

4.4 Uncertainty propagation results for the Cabra flame's lift-off height in Span(ξ_T, ξ_{12})

4.4.1 Uncertain behavior analysis

As a whole, the lift-off height of the Cabra flame is mostly sensitive to the co-flow temperature, but the 12th reaction of the mechanism ($\text{H} + \text{O}_2 = \text{OH} + \text{O}$) also has a sensible influence, especially at low temperatures. This behavior is consistent with the sensitivity analysis performed a priori. The surrogate also shows a sudden increase in slope, both ξ_T - and ξ_{12} -wise when the temperature decreases. This is consistent with the results of Figure 2.3 (from Chapter 2, reproduced below), which plots the probability distribution of the auto-ignition delay time of a homogeneous reactor at z_{mr} for different initial temperatures.




Figure 2.3 (reproduced): Violin plots of the auto-ignition delay τ in the conditions of the Cabra flame for different co-flow temperatures for the original and boiled-down Konnov (2019) [55] mechanisms

This sudden increase in sensitivity to both variables at the low end of the temperature range translates the transition from two auto-ignition modes allowed by the kinetic mechanism:

- High-temperature auto-ignition, very fast, found in the "plateau" region which begins at about $\xi_T \approx 0.5$
- Low-temperature auto-ignition, much slower, which would be found for even lower co-flow temperatures

We did not investigate the Cabra flame in the conditions of weak auto-ignition [92, 122] in this study. However, this combustion mode is encountered in Chapter 2 for homogeneous reactors and non-premixed auto-igniting flamelets. The increase in sensitivity to ξ_{12} only reflects the higher uncertainty of the mechanism in the transition zone between weak and

strong auto-ignition.

The surrogate model, duly validated, can then be resampled extensively to obtain the probability distribution of the lift-off height in the uncertain space spanned by ξ_T and ξ_{12} . $2^{19} \approx 500k$ samples are drawn from the 2D Sobol' sequence and the surrogate is evaluated for each sample. The resulting PDF is presented in Figure 4.11.

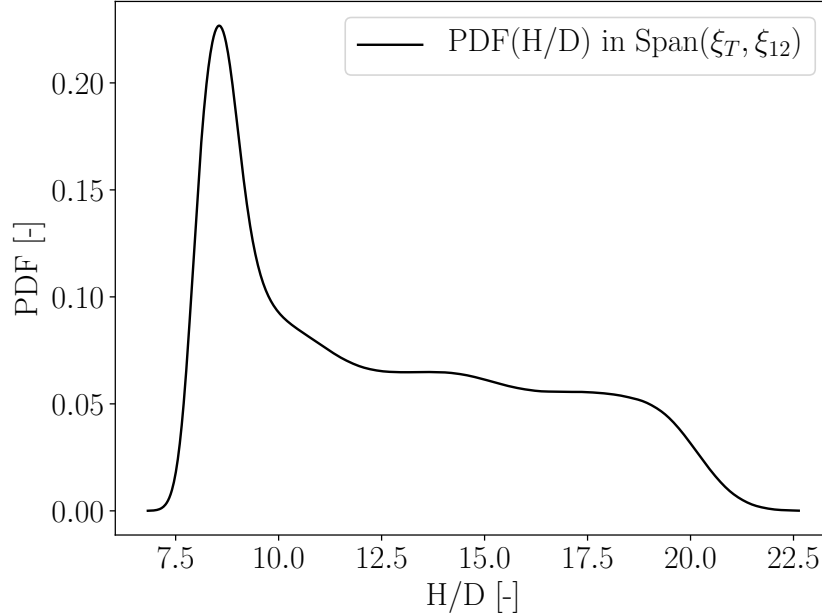


Figure 4.11: Probability distribution of the lift-off height in the uncertain space spanned by ξ_T and ξ_{12} using 2^{19} surrogate evaluations

The peak of the distribution, located around $H/D \approx 8.5$, corresponds to the plateau of high-temperature auto-ignition, while the near-uniform component between $H/D = 10$ and $H/D = 20$ corresponds to the transition zone, for $\xi_T < 0.5$, or $T_{\text{cof}} < 1045$ K.

4.4.2 Comparison with Ji et al. (2019)

Ji et al. [45] proposed an uncertainty propagation study for the lift-off height of the simulated H_2 Cabra flame in 2019. Their simulation framework is Reynolds Averaged Navier-Stokes simulations, with a $k-\varepsilon$ turbulence model, the Transported Probability Density Function approach to model combustion, and Interaction-by-Exchange-with-the-Mean approach for micro-mixing. They used the kinetic mechanism from Li et al. (2004) [63] with uncertainty factors borrowed from Konnov (2008)'s uncertain mechanism [54]. They did not propagate the uncertainties on the temperature measurement. However, they did qualitatively investigate the influence of co-flow temperature on the flame's lift-off height.

Using an active direction approach to reduce the kinetic uncertain dimension and a second-order PCE as a response surface, they obtained the PDF that we plotted in Figure 4.12 in the dotted line. It is visually different from the one plotted in their article [45] (Fig. 4b.) because they used a logarithmic scale in abscissa. Otherwise, the data points

are strictly identical.




Figure 4.12: Probability distribution of the lift-off height as estimated by Ji et al. [45] and in our work for the fixed, nominal co-flow temperature of $T_{\text{cof}} = 1045$ K and uncertain kinetic mechanism

Ji et al. found an expected lift-off height of $H/D_{\text{expected}} = 4.6$, which is in agreement with the experimental measurement of Wu et al. [120] (as seen in Figure 4.13, in the next Section) and a 95% confidence interval of [1.1, 28.3] which covers all the experimental measurements for a co-flow temperature of 1045 K.

We plotted on the same graph, in plain line, a PDF of the lift-off height obtained by resampling a slice of our surrogate model (trained with $UF_{12} = 1.1$ as given by the Konnov 2019 mechanism) at $\xi_T = 0.5 \Leftrightarrow T_{\text{cof}} = 1045$ K in order to place ourselves in the same conditions as Ji et al.. The PDF is much narrower than Ji’s proposed PDF and the expected value is also much higher: $H/D_{\text{expected}} = 11.8$.

The divergence in expected value between [45] and our study can be explained by the differences between Li (2004)’s and Konnov (2019)’s nominal kinetic mechanisms. In particular, for the most sensitive reaction of both mechanisms, $\text{H} + \text{O}_2 = \text{OH} + \text{O}$, Li et al. propose a rate constant of 3.55×10^{15} while Konnov proposes a rate constant of 1.04×10^{14} .

We have three leads to explain the difference in PDF width showcased in Figure 4.12:

- Konnov (2008)’s uncertainty factor for $\text{H} + \text{O}_2 = \text{OH} + \text{O}$ was 1.5, while its updated (2019) version only has an uncertainty factor of 1.1 for this reaction. This uncertainty reduction, stemming from a decade of further research, must have a significant influence on this problem. To evaluate this influence, we resampled our surrogate model for an uncertainty factor of 1.5 on A_{12} . The response surface was not trained for such a wide range of possible values, but the linear trend of the

Kriging guarantees an asymptotic behavior which is at least sound compared to a PCE. The result of this resampling is also plotted in Figure 4.12 in dashed line. This manipulation widens the PDF about three times and brings out the signature of the transition between low- and high-temperature combustion (or "weak" and "strong" combustion respectively): the onset of a bipolar distribution with the right-side bump in the PDF. However, this lead is not enough to explain the whole difference between Ji's and ours PDF.

- Our study reduced the uncertain dimension by discarding all variables that were not sensitive enough. Most of the uncertain reactions were discarded because of their insignificance before the uncertainty on the co-flow temperature. In particular, the next reaction in terms of sensitivity, $\text{H} + \text{HO}_2 = 2 \text{OH}$ has a first-order Sobol' index about 2 times lower than $\text{H} + \text{O}_2 = \text{OH} + \text{O}$ in 0D homogeneous reactor simulations in the conditions of the Cabra flame (cf Figure 4.4). Taking this uncertain parameter into account may help to further explain the difference in PDF width between [45] and our study.
- Finally, our most speculative lead is that Ji et al.'s choice of surrogate type may have artificially widened their proposed PDF. Fig. 4a. of their article [45] suggests that their second-order polynomial fit lacks curvature, which would result in an over-representation of high lift-off heights in the resampling.

As a conclusion, further work is needed to benchmark simulations and uncertainty propagation methods from different teams with modern uncertainty data.

4.4.3 Comparison with deterministic simulation results and experimental measurements

Section 4.4.2 showed the PDF of lift-off height conditioned to $T_{\text{cof}} = 1045 \text{ K}$. This procedure can be generalized, to several co-flow temperatures in the range of measurement uncertainties. The resulting violin plots are shown in Figures 4.13 and 4.14, compared respectively to experimental measurements and to other numerical works.

Figure 4.13 shows a good agreement between the experimental measurements and our simulation: the predicted lift-off heights are well within the range of measurements from [11, 120, 34]. Its evolution with regards to the co-flow temperature is also well reproduced. It also confirms that, in the H_2 Cabra flame configuration, taking into account the uncertain kinetic mechanism is only a refinement compared to the uncertainties introduced by uncertain measurements of the co-flow temperature.

Figure 4.14, on the other hand, shows that state-of-the-art kinetic uncertainties are not enough to explain the spread of numerical results. Section 4.4.2 already highlighted that older mechanisms were more uncertain. According to the PDF proposed by Ji et al., the relatively high uncertainties from [54] are representative of the spread of kinetic mechanisms from that era and suffice to explain the different experimental results from the literature.

The reduction of kinetic uncertainties over the last decade, compiled in [55], allows for much less uncertain simulations nowadays.




Figure 4.13: Violin plots of the lift-off height of the Cabra flame for several co-flow temperatures, compared to the measurements from [11, 120, 34]. For enhanced readability, measurement uncertainties of the co-flow temperature are only plotted for the experiment from Cabra et al. [11]. The error bars are the same for the other experiments.

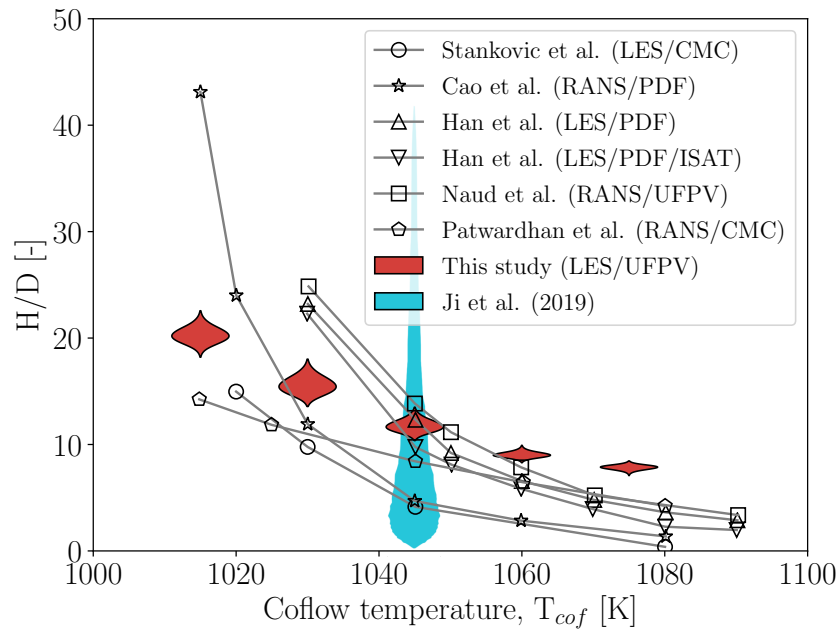


Figure 4.14: Violin plots of the lift-off height of the Cabra flame for several co-flow temperatures, compared to deterministic simulations from [105, 13, 35, 74, 78], and a violin plot of the uncertainty quantification study from [45]

However, there is another major axis that introduces uncertainties in numerical simulations of the H_2 Cabra flame: the turbulence and turbulence-combustion interaction models. The next section investigates this new aspect of the problem.

4.5 Investigating the uncertainty in the turbulence-combustion model

4.5.1 The uncertain dissipation coefficient of mixture fraction variance

Section 3.6.2 showed that the flame shape is extremely sensitive to the dissipation coefficient of the mixture fraction variance C_D . Moreover, while the theoretical value of C_D should lie between $2\pi^2$ and 40 for typical LES filter sizes, much lower values can be found in the literature [23, 117, 43]. This disagreement between the values proposed by different authors and values obtained by theoretical developments can be considered as a parametric uncertainty in the turbulence-combustion interaction model.

Without any more information on the probability distribution of C_D , we assume that it is equiprobable over the entire range of possible values. Therefore, we model C_D as a uniform random variable:

$$\xi_{C_D} = \frac{C_D - C_{D,\text{Min}}}{C_{D,\text{Max}} - C_{D,\text{Min}}} \sim \mathcal{U}(0, 1) \quad (4.7)$$

Where $C_{D,\text{Min}} = 2.0$ is the lowest value we found in the literature [23] and $C_{D,\text{Max}} = 40.0$ is the highest possible value predicted by analytical developments [41].

Therefore, the uncertainties must now be propagated from the 3D space: $\text{Span}(\xi_{A_{12}}, \xi_T, \xi_{C_D})$. To this end, 29 more LES are run for this study: 19 more training samples are drawn from the 3D Sobol' sequence and 10 more validation samples are drawn with LHS.

The 30 LES drawn from $\text{Span}(\xi_{A_{12}}, \xi_T)$ are also recycled to train and evaluate the surrogate model. In total, 39 training samples and 20 validation samples are used in the following.

As a side note, the first sample sample from the 2D and 3D Sobol' sequence are located at the same coordinates: they correspond to the nominal case. This is why only 19 samples were drawn from the 3D sequence, starting from the second element of the sequence.

4.5.2 Surrogate modeling

Several surrogate modeling methods were tested to emulate the behavior of the lift-off height in $\text{Span}(\xi_{A_{12}}, \xi_T, \xi_{C_D})$. They are presented and evaluated in Annex E. The best candidate among them was a PC-Kriging, whose parameters are given in Table 4.5. However, the performance of this surrogate was unsatisfying:

Table 4.5: Properties of the best surrogate model found in $\text{Span}(\xi_{A_{12}}, \xi_T, \xi_{C_D})$

Surrogate Type	PC-Kriging
Trend	Regression PCE, $p = 3, q = 0.75$
Kernel	Matérn 5/2
Length-scales (θ) determination	Maximum likelihood approach
Nugget	0.05

$$\text{NRMSE}_{\text{PCK}} = 7.5\% \quad (4.8)$$

$$\text{NMAE}_{\text{PCK}} = 18.6\% \quad (4.9)$$

To improve the quality of the surrogate model, we tried to reduce again the uncertain dimension in order to reduce the distance between observations. The discovery of an active direction was attempted but to no avail. Then, we noticed that the problem was mostly additive, which means that interactions between ξ_{C_D} and the other uncertain variables were small enough to be neglected. This allows writing:

$$F(\xi_{12}, \xi_T, \xi_{C_D}) \approx F_{\text{kinetics}}(\xi_{12}, \xi_T) + F_{\text{turbulence}}(\xi_{C_D}) \quad (4.10)$$

In this case, the 2D surrogate proposed in Section 4.3 was recycled and substracted from the observations in the 3D space. A 1D surrogate could then be constructed from the remaining data projected on the direction of ξ_{C_D} . The best surrogate model we found, in this case, was a regression Kriging with linear trend, presented in Table 4.6 and plotted in Figure 4.15.

Table 4.6: Properties of the surrogate used in $\text{Span}(\xi_{C_D})$

Surrogate Type	Regression Kriging
Trend	Linear
Kernel	Matérn 5/2
Length-scales (θ) determination	Maximum likelihood approach
Nugget	0.05




Figure 4.15: Response surface of the 1D surrogate

In this figure, the cluster of training and validation points located at $\xi_{C_D} \approx 0.47$ corresponds to the observations from Section 4.3.

Using the additivity of the problem, both the 2D and 1D surrogates are finally aggregated to form a 3D surrogate in $\text{Span}(\xi_{A_{12}}, \xi_T, \xi_{C_D})$, as shown in Equation 4.10. As a side note, because we submitted the observations to a logarithmic preconditioning

(Sec. 4.3), the response surfaces shown in Figures 4.9 and 4.15 are actually multiplied, not added, to obtain the aggregated surrogate of the flame lift-off. Its validation scores are the following :

$$\text{NRMSE}_{\text{agg}} = 6.8\% \quad (4.11)$$

$$\text{NMAE}_{\text{agg}} = 13.0\% \quad (4.12)$$

The resulting response surface is much more satisfactory. Its root mean square is slightly lower and its maximum error is significantly reduced. A slice of the surrogate in the hyperplane $\xi_{12} = 0$ is presented in Figure 4.16.




Figure 4.16: Response surface of the aggregated surrogate in the hyperplane $\xi_{12} = 0$

Training and validation observations are absent from this plot because most of them don't belong in the hyperplane $\xi_{12} = 0$ and it would make no sense to represent them.

As expected from the qualitative study of Section 3.6.2, low dissipation coefficients of mixture fraction variance lead to large simulated flame lift-off. Furthermore, no significant influence of this parameter is found for values of $\xi_{C_D} > 0.5$, which loosely corresponds to the range of values that should be used for typical LES sizes ($2\pi^2 < C_D < 40$ [41]). This is interesting for future deterministic LES of this type of flame because it suggests a "perfect" value of C_D does not need to be determined prior to the simulation. At least in this configuration, the lift-off height of the flame is only weakly dependent on this parameter, as long as it is in the range recommended by Ihme [41]. However, the flame shape is more and more sensitive to C_D with decreasing values. Increased caution should

therefore be exercised when analyzing simulations using C_D values in the lower range.

Finally, the minor bulge in the response surface, which can be observed in Figures 4.15 and 4.16 near $\xi_{C_D} = 1$, is probably an artifact of the surrogate model and probably has no physical signification.

Aside from the error metrics presented above, the aggregated surrogate can be visually verified against its training and validation observations using the summary plot presented in Figure 4.17, which compares the observations (on the x-axis) with the surrogate evaluations at the same coordinates in the uncertain space (on the y-axis).




Figure 4.17: Summary plot of the aggregated surrogate, in $\text{Span}(\xi_{A_{12}}, \xi_T, \xi_{C_D})$

The surrogate is not perfect. In particular, it fails to accurately predict lift-off heights in the higher range, i.e. when the flame is lifted above 15 D. However, the general agreement of the surrogate with the training and validation observations is satisfactory.

4.5.3 Probability distribution of the lift-off height in the 3D space

The 3D surrogate constructed in Section 4.5.2 is sampled extensively using 2^{19} evaluations at coordinates drawn from the 3D Sobol' sequence. The resulting PDF is presented in Figure 4.18.

The shape of the PDF in $\text{Span}(\xi_{C_D}, \xi_T, \xi_{12})$ is very close to the one obtained in $\text{Span}(\xi_T, \xi_{12})$, presented in Figure 4.11. However, the region of high probabilities around $H/D = 9$ is less peaked and the tail of the distribution is much heavier, due to the large lift-off heights permitted by the combination of low co-flow temperatures, slow-chemistry - which both slow down auto-ignition - and low dissipation coefficient of the mixture fraction variance - which enlarges the zone where auto-ignition does not even start due to a too high strain rate.




Figure 4.18: Probability distribution of the lift-off height in the uncertain space spanned by ξ_{CD} , ξ_T , and ξ_{12} using 2^{19} surrogate evaluations. Note that the x-scale is different from the PDF in the 2D uncertain space shown in Figure 4.11

We can also slice and resample the response surface for different temperatures and compare it to deterministic simulations (Figure 4.19) and experimental results (Figure 4.20) from the literature.

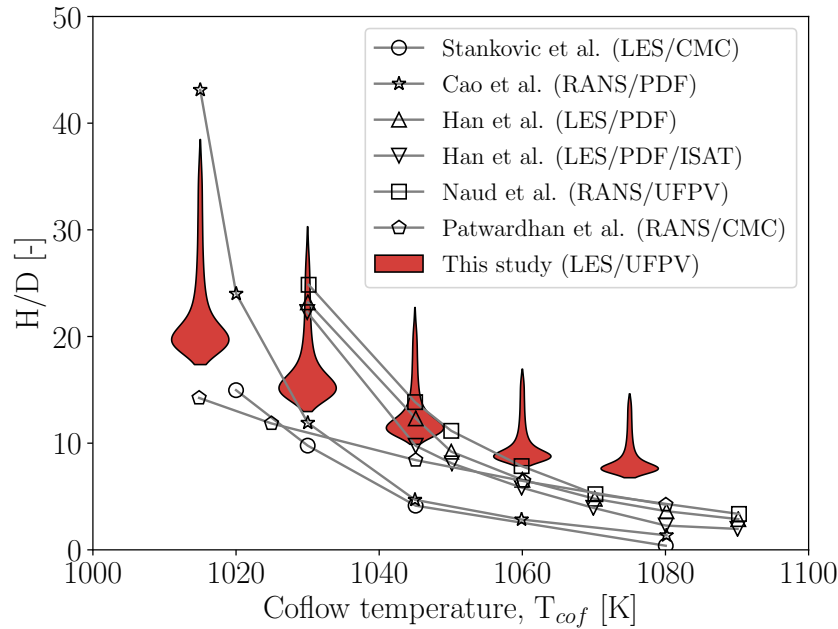


Figure 4.19: Violin plots of the lift-off height of the Cabra flame for several co-flow temperatures, compared to deterministic simulations from [105, 13, 35, 74, 78]

As expected, the resulting probability distributions are much wider than those obtained when considering only the chemical kinetics uncertain, as seen in Figure 4.14.




Figure 4.20: Violin plots of the lift-off height of the Cabra flame for several co-flow temperatures, compared to the measurements from [11, 120, 34]. For enhanced readability, measurement uncertainties of the co flow temperature are only plotted for the experiment from Cabra et al. [11]. The error bars are the same for the other experiments.

The abrupt bottom side of the violin plots is due to the same artifact of the surrogate model that caused the minor bulge in Figures 4.15 and 4.16 near $\xi_{C_D} = 1$.

Taking into account the uncertainty of the turbulence-combustion interaction model can explain some of the variety of simulation results found in the literature. Moreover, the state-of-the-art Konnov (2019) mechanism [55] is much less uncertain than other, more ancient mechanisms that were used when the other simulations were run and Section 4.4.3 already showed that the uncertainties in kinetic mechanisms of that era were already enough to explain most of the discrepancies between the numerical results of the literature. This explains why our study, as expected, using modern uncertainty factors, does not cover the whole range of previous simulation results. Finally, our uncertainty propagation study also explains most of the experimental measurements, considering that their uncertainty on the co-flow temperature is $\pm 3\%$.

Only the smallest of lift-off heights observed in the experiments and predicted by other simulations and, most importantly, observed by Gordon et al. are seemingly not explained by our study.

However, looking more closely at the foot of the flame for our observation with the shortest flame lift-off, we identified an early zone of progress variable production at the foot of the jet. This is illustrated in Figure 4.21.

This early ignition zone leads to a small - but visible - increase in hydroxyl mixture fraction near the jet exit. This suggests that smaller lift-off heights are accessible with our simulation setup, but that they would need slightly higher co-flow temperatures to appear. We are confident that future investigations at higher temperatures would confirm




Figure 4.21: Comparison of the progress variable and hydroxyl mass fraction fields between the lowest flame simulated (left) and the nominal case (right)

this hypothesis. Moreover, given the uncertainty of the temperature measurement for experimental results, our simulation is already within Gordon et al.'s measurement range.

4.6 Takeaway points

On the methodological side:

- A procedure to take into account highly dimensional kinetic and experimental uncertainties in LES is proposed in this chapter.
- Uncertainties not directly related to chemical kinetics can also be taken into account. We demonstrate this by taking into account an uncertain variable in the turbulent combustion model.

On the physical interpretation side:

- Taking into account state-of-the-art kinetic and temperature uncertainties shows that deterministic simulations of the Cabra flame are not suited to predict its behavior. It may not be a very useful configuration to compare quantitatively turbulence or combustion models in a deterministic framework.
- The reduction of kinetic uncertainties in the last decade has led to a dramatic decrease in lift-off height uncertainty in the simulated H₂ Cabra flame.
- The uncertainties on the turbulence-combustion interaction model are about as sensitive as the uncertainty on the co-flow temperature. Therefore, special care must be exercised when setting model parameters, in configurations similar to the H₂ Cabra flame, i.e. auto-igniting non-premixed lifted flames.

Chapter 5

Uncertainty Quantification in averaged LES fields

In this chapter, we aim to quantify the uncertainties of the averaged LES fields of the Cabra flame. That is, we aim to evaluate the PDF of every time-averaged quantity at every point in the computational domain. This will allow comparing the experimental and numerical ranges of uncertainty on local quantities. Furthermore, this work is a proof of concept for the prediction of uncertain averaged fields using LES for turbulent reacting flows.

We propagate the uncertainties coming from the kinetic mechanism, the co-flow temperature, and the dissipation coefficient of the mixture fraction variance. The fields considered in this study are the fields that were measured experimentally by Cabra et al. [11].

Sections 1 and 2 deal with uncertain dimension reduction of the input and output space respectively. The third section presents and validates the surrogate models used to predict uncertain fields. Section 4 validates the method as a whole and presents the results of the uncertainty propagation on the averaged LES fields of the Cabra flame.

5.1 Reducing the input uncertain dimension

In this study, we consider the uncertainties from the kinetic mechanism, the co-flow temperature, and the mixture fraction variance dissipation coefficient. We employ the same reduction of the input uncertain space as the one used in Section 4.5: The dominant variable approach on the auto-ignition delay time of a homogeneous reactor corresponding to the most reactive mixture of the Cabra flame eliminates a large number of weakly sensitive reaction rates. Finally, the reduced uncertain space is $\text{Span}(\xi_{A_{12}}, \xi_T, \xi_{C_D})$.

Therefore, we can re-use the simulations used in Chapter 4: the results of 39 LES are used as training observations, and 20 more are used as validation observations. The averaged fields of hydroxyl mass fraction for the training samples are shown in Figure 5.1. The fields of every considered averaged LES variable for the training samples are shown in Annex F.




Figure 5.1: Averaged Y_{OH} fields of the training samples

5.2 Reducing the output uncertain dimension

The computational mesh of the LES of the Cabra flame comprises 5.5 million nodes. Considering that Cabra [11] measured local values for Y_{OH} , Y_{O_2} , Y_{H_2} , Y_{H_2O} , Z , and the temperature, a total of 33 million local surrogates should be constructed and 33 million PDFs evaluated.

This brute-force approach is quite expensive, and clearly suboptimal:

- There is a spatial auto-correlation of each field and a cross-correlation between fields which means that there would be some redundant work in constructing every surrogate independently.
- With only 39 training observations, it is highly probable that the surrogates would not be perfectly determined. Therefore, inconsistencies may appear, such as spatial discontinuities or the sum of mass fractions locally not equal to 1.

These reasons call for a drastic dimension reduction of the uncertain outputs.

5.2.1 Preliminary work on the spatial mesh

The averaged fields are interpolated on a structured, axisymmetric, hexahedral mesh which discretizes the region covered by experimental measurements ($x \in [0, 26D]$ and $r \in [0, 3D]$). This 3D mesh has $315 \times 250 \times 20 \approx 1.5M$ nodes. An azimuthal mean is then applied to the fields to enforce the axisymmetric nature of the fields and further reduce the mesh size to $n_{\text{nodes}} = 315 \times 250 \approx 79k$.

This preliminary work brings the uncertain output dimension to about 4.8 million, which is still extremely large. Moreover, it does not address the two problems cited above. Further dimension reduction is therefore needed.

5.2.2 Uncertain dimension reduction using Principal Component Analysis

The ideal dimension reduction tool would have the following properties:

- Reduce the dimension of the problem to a tractable number ($d_{\text{out}} < 10$)
- Allow the prediction of the averaged fields from the reduced variables without too much information loss
- Preserve the auto-correlation and cross-correlations of the fields

Principal Component Analysis (PCA) fills these specifications and is used in this work.

For each observation, the averaged fields are centered and normalized using the "standard score" normalization [56]. Then, the normalized averaged fields are queued up in a unique "features' vector" of size $n_{\text{features}} = n_{\text{nodes}} \times n_{\text{fields}} \approx 4.7 \times 10^5$.

A PCA is applied to the set of observations. The modes obtained with this PCA represent the variations of the six different fields that were initially considered. This way,

their cross-correlation is conserved.

The first eigenvalues of the PCA are plotted in Figure 5.2. The basis of PC modes is truncated such that only modes that correspond to more than 1% of the variance of the observations are retained. This truncature is materialized by the red line in Figure 5.2. With only three modes left, 98.5% of the variance of the observation is explained. These three modes are plotted in Figure 5.3 for Y_{OH} and in Figure 5.4 for the gas temperature. The modes are plotted for the other variables in Annex G.




Figure 5.2: Variance explained by each mode of the PCA

These modes are enough to account for most of the possible variations of the Cabra flame in the uncertain space. A set of three weights - which can also be considered as coordinates in the basis formed by the three first modes of the PCA - are therefore enough to characterize a flame shape. The output uncertain dimension is reduced to three, compared to about 33 million initially. Equation 5.1 summarizes the prediction process of all the fields:

$$\mathbf{Y}_{\text{prediction}}(\boldsymbol{\xi}) = \bar{\mathbf{Y}} + \mathbf{N} \left[w_1(\boldsymbol{\xi}) \cdot \mathbf{M}_1 + w_2(\boldsymbol{\xi}) \cdot \mathbf{M}_2 + w_3(\boldsymbol{\xi}) \cdot \mathbf{M}_3 \right] \quad (5.1)$$

Where $\mathbf{Y}_{\text{prediction}}(\boldsymbol{\xi})$ is the vector of predicted field values, for each considered field and at each spatial location. It can be written as:

$$\mathbf{Y}_{\text{prediction}}(\boldsymbol{\xi}) = \begin{pmatrix} Y_{\text{OH},1}, \dots, Y_{\text{OH},n_{\text{nodes}}}, \\ Y_{\text{O}_2,1}, \dots, Y_{\text{O}_2,n_{\text{nodes}}}, \\ \vdots \\ Z_1, \dots, Z_{n_{\text{nodes}}} \end{pmatrix} \quad (5.2)$$

$\bar{\mathbf{Y}}$ is the vector of the mean of the training observations at each location for each field. Its shape is the same as $\mathbf{Y}_{\text{prediction}}$.




Figure 5.3: Mean and retained modes for Y_{OH}

$\mathbf{M}_1, \mathbf{M}_2, \mathbf{M}_3$ are the first three modes of the PCA. They are also vectors whose shape is the same as $\mathbf{Y}_{\text{prediction}}$: the first n_{nodes} components correspond to the mode intensity for the variable Y_{OH} at every node, and so on.

w_1, w_2, w_3 are the scalar weights associated with the first three modes of the PCA. They depend on the position $\boldsymbol{\xi}$ in the uncertain space and are predicted in Section 5.3.

Finally, \mathbf{N} is a $(n_{\text{features}}, n_{\text{features}})$ diagonal matrix containing the n_{fields} normalization coefficients applied to each field before computing the PCA. This allows rescaling the weighted modes and obtain dimensionalized quantities:

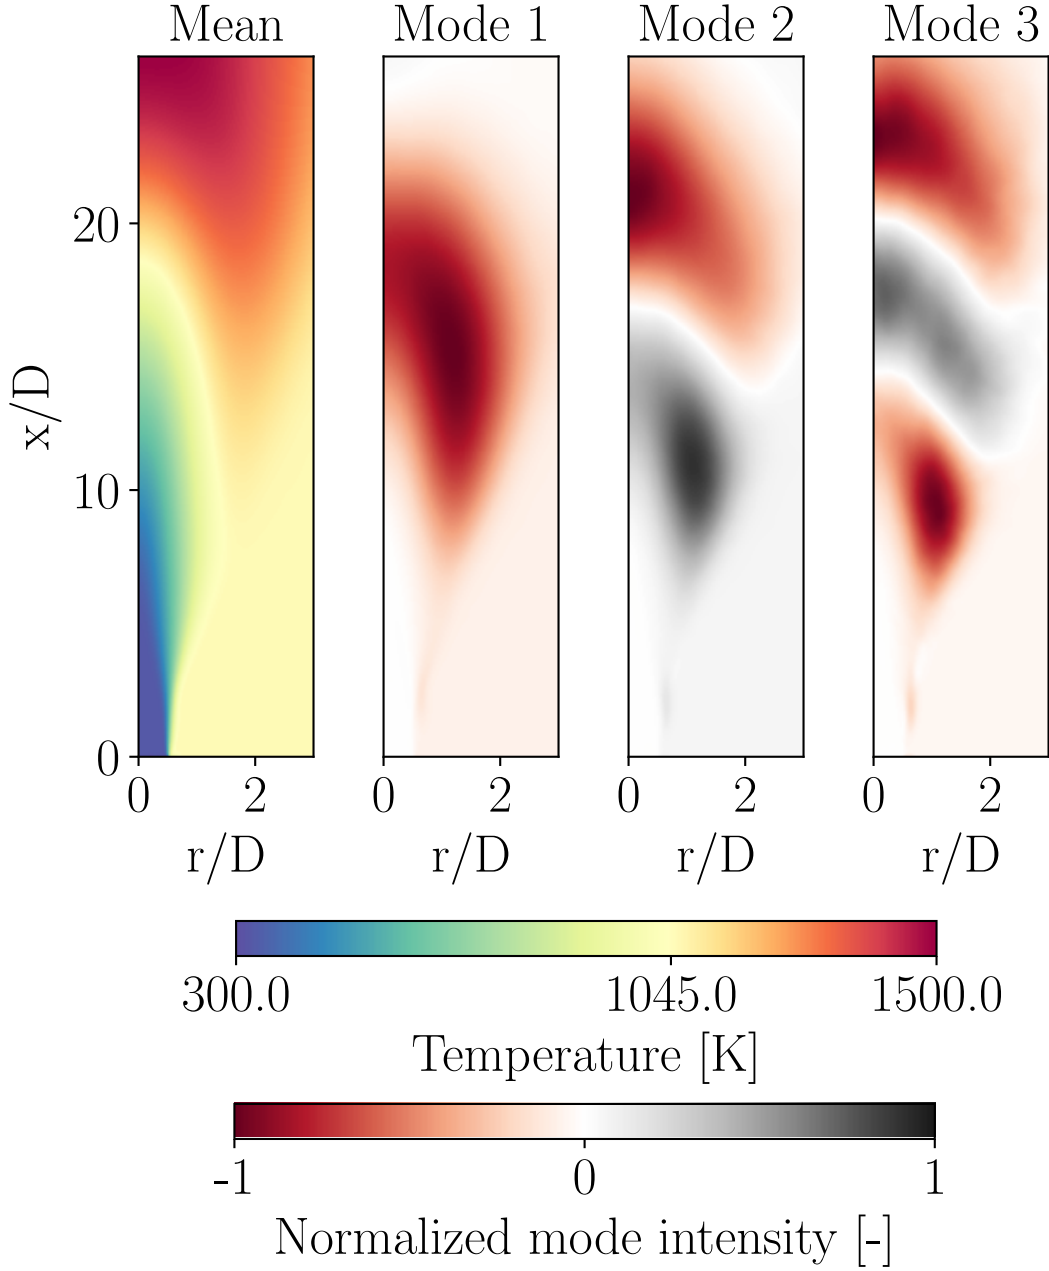


Figure 5.4: Mean and retained modes for the gas temperature

$$\mathbf{N} = \text{Diag} \left(\underbrace{\text{norm}_{Y_{\text{OH}}}, \dots, \text{norm}_{Y_{\text{OH}}}}_{n_{\text{nodes}} \text{ times}}, \right. \\
 \vdots \\
 \left. \text{norm}_Z, \dots, \text{norm}_Z \right) \quad (5.3)$$

As with any type of dimension reduction, we can compute error metrics to measure the loss of information induced by the compression. For each field, we concatenate the field value at each point and for each validation sample in a single vector. This vector is created for the LES averaged fields and their projections on the first three PCA modes.

The results are given in Table 5.1.

Table 5.1: Error scores of the projection of the validation observations on the first three PCA modes, per field

	Y_{H_2}	Y_{H_2O}	Y_{O_2}	Y_{OH}	T	Z
NRMSE	2.27%	1.53%	1.21%	7.31%	1.16%	1.29%
NMAE	5.92%	9.06%	5.19%	21.96%	6.49%	4.29%

These error scores are satisfying for the most part, in particular regarding the extremely low mean error (NRMSE). Only the Y_{OH} field’s projection is worse, with the validation samples VS12 and VS16 contributing the most to the NRMSE and VS12 being responsible for the disappointing NMAE. To better understand where these errors come from, we propose Figure 5.5.

Each subfigure of Figure 5.5 presents the fields of hydroxyl mass fraction Y_{OH} computed by the LES (left), compared with their projection on the first three modes of the PCA (right) for every validation sample. This visual representation of the projection suggests that no significant amount of information was lost in the compression. In fact, in most cases, the flame shapes are very well preserved by the compression. In particular, the flame lift-off height and axial evolution of the hydroxyl mass fraction are well reproduced. Furthermore, the explanation of the relatively high error scores observed with samples VS12 and VS16 appears clearly.

Sample VS12 has a higher lift-off flame than any of the training samples. As a reminder, the PCA only generalizes its training dataset and cannot accomplish miracles for samples outside its training range. The case of VS16 is a bit subtler: The lift-off height ($\approx 17D$) is well within the range of previously-seen values. However, the foot of the flame - i.e. the area where Y_{OH} values are non-negligible but lower than the auto-ignition threshold - is much longer than any of the training samples.

In both of these problematic cases, using more PCA modes did not improve significantly the agreement between the averaged LES fields and their projection. In terms of error metrics, using 5 modes, for example, the NMAE for the Y_{OH} field stayed around 22% and the NRMSE decreased slightly to reach 5.5%. To improve the quality of the projection, a more extensive sampling of the input uncertain space would be needed - meaning more LES should be run.

Overall, the compression satisfyingly preserves flame shapes given our training sample budget.

5.3 Surrogate modeling

5.3.1 Preliminary remarks

The dimension reduction of inputs and outputs defined in the previous sections results in only three input variables ($\xi_{A_{12}}, \xi_{Tcof}, \xi_{C_D}$) and three output variables (the weights of the three most important modes of the PCA, see Equation 5.1). The input variables are independent by nature. The three coefficients of the PCA modes may be statistically

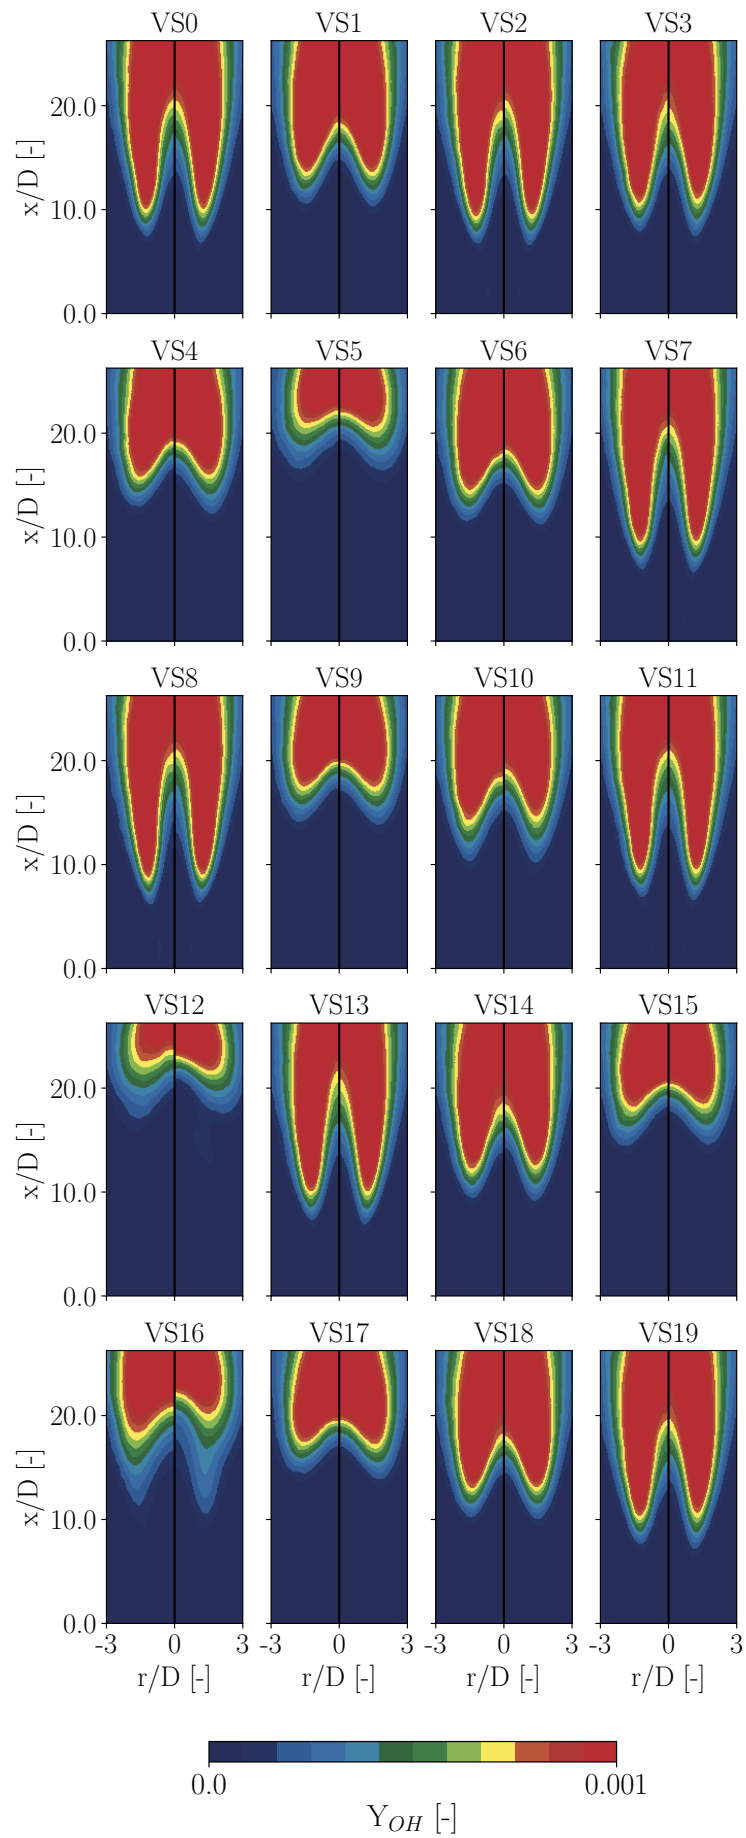


Figure 5.5: Y_{OH} fields of the validation samples **compressed** by the PCA (**right**) compared to the LES (left)

dependent. This is not a problem for us because we determine them as a function of the initial, independent uncertain parameters.

5.3.2 Choice of surrogates

Unlike in Section 4.5.2, we did not find any further simplification of the problem of emulating the coefficients of the PCA modes. Therefore, three surrogates must be fitted in the 3D input space $\text{Span}(\xi_{A_{12}}, \xi_T, \xi_{C_D})$. We select PC-Kriging (Kriging using a Polynomial Chaos Expansion as a trend) to emulate the PCA coefficients. PCEs of total order 2 are used for the first two coefficients, and a PCE of total order 1 (a linear trend) is used for the third coefficient. The hyperparameters of the Kriging models are fitted by minimizing the NRMSE in a Leave-One-Out cross-validation procedure (Section 1.6.4).

5.3.3 Verification of Surrogates

Unfortunately, 39 observations, 20 of which are in the 2D plane $\text{Span}(\xi_{A_{12}}, \xi_T)$ (see Section 4.5.2), do not fill the 3D input space very densely. Therefore, it is doubly important to check the quality of the surrogate models. Figure 5.6 and Table 5.2 give a summary of the quality of the surrogates used to emulate the three PCA modes.




Figure 5.6: Summary plots of the surrogates used to predict each mode's coefficient

Table 5.2: Error metrics for the surrogate models

Mode	Surrogate Type	PCE trend order	NRMSE	NMAE
1	PC-Kriging	2	27.36%	49.10%
2	PC-Kriging	2	32.32%	39.44%
3	PC-Kriging	1	55.45%	67.10%

Figure 5.6 suggests that the surrogate for the coefficient of mode 1 is relatively satisfying, although one of the validation samples is poorly predicted. This one poor prediction mostly drives the relatively high values of NRMSE and NMAE computed for this surrogate. The surrogates of the coefficients of modes 2 and 3 are a bit worse, both in terms of error metrics and visual agreement on the summary plot. Nevertheless, both of them correctly reproduce the trend of the coefficients they emulate, and they both avoid overfitting. Therefore, we consider these surrogate models satisfying despite their imperfections.

Moreover, Figure 5.2 showed a step decrease in the variance explained by the successive modes. This means that the greater the mode index, the less a slightly off coefficient prediction matters in the reproduction of the fields.

5.4 Validation of the field prediction, as a whole

Both the projection phase and the surrogate modeling of the coefficients have been verified. However, the combination of these two methods to predict a flame shape requires the validation of the method as a whole. Similarly to Section 5.2.2, we compute error scores for the prediction of the validation samples. The results are given, per field, in Table 5.3.

Table 5.3: Error scores for the prediction of the validation observations, per field

	Y_{H_2}	Y_{H_2O}	Y_{O_2}	Y_{OH}	T	Z
NRMSE	3.98%	3.76%	3.15%	11.15%	2.49%	1.79%
NMAE	7.97%	16.06%	14.37%	29.56%	11.69%	4.47%

The error scores are only slightly higher than those obtained in Section 5.2.2 to evaluate the projection step. This means that while the surrogate models that predict the coefficients of the modes are not perfect, most of the prediction error comes from a lack of richness of the PCA modes, in the sense that they restrict the possibilities of flame shape variation. Ultimately, as seen in Section 5.2.2, this comes from a lack of variety in the training observations. To enrich the training dataset, more LES should be run, each one costing 100k hours of CPU time. Due to our finite computational budget, we did not include any more observations in the study. Besides, the fidelity of the fields prediction is already more than satisfying in the probable regions of the uncertain space.

Finally, we provide a visual verification of the predictions with Figures 5.7 and 5.8 which compare the surrogate predictions and LES observations for the hydroxyl mass fraction and temperature fields respectively. The comparisons of the other fields can be found in Annex H.

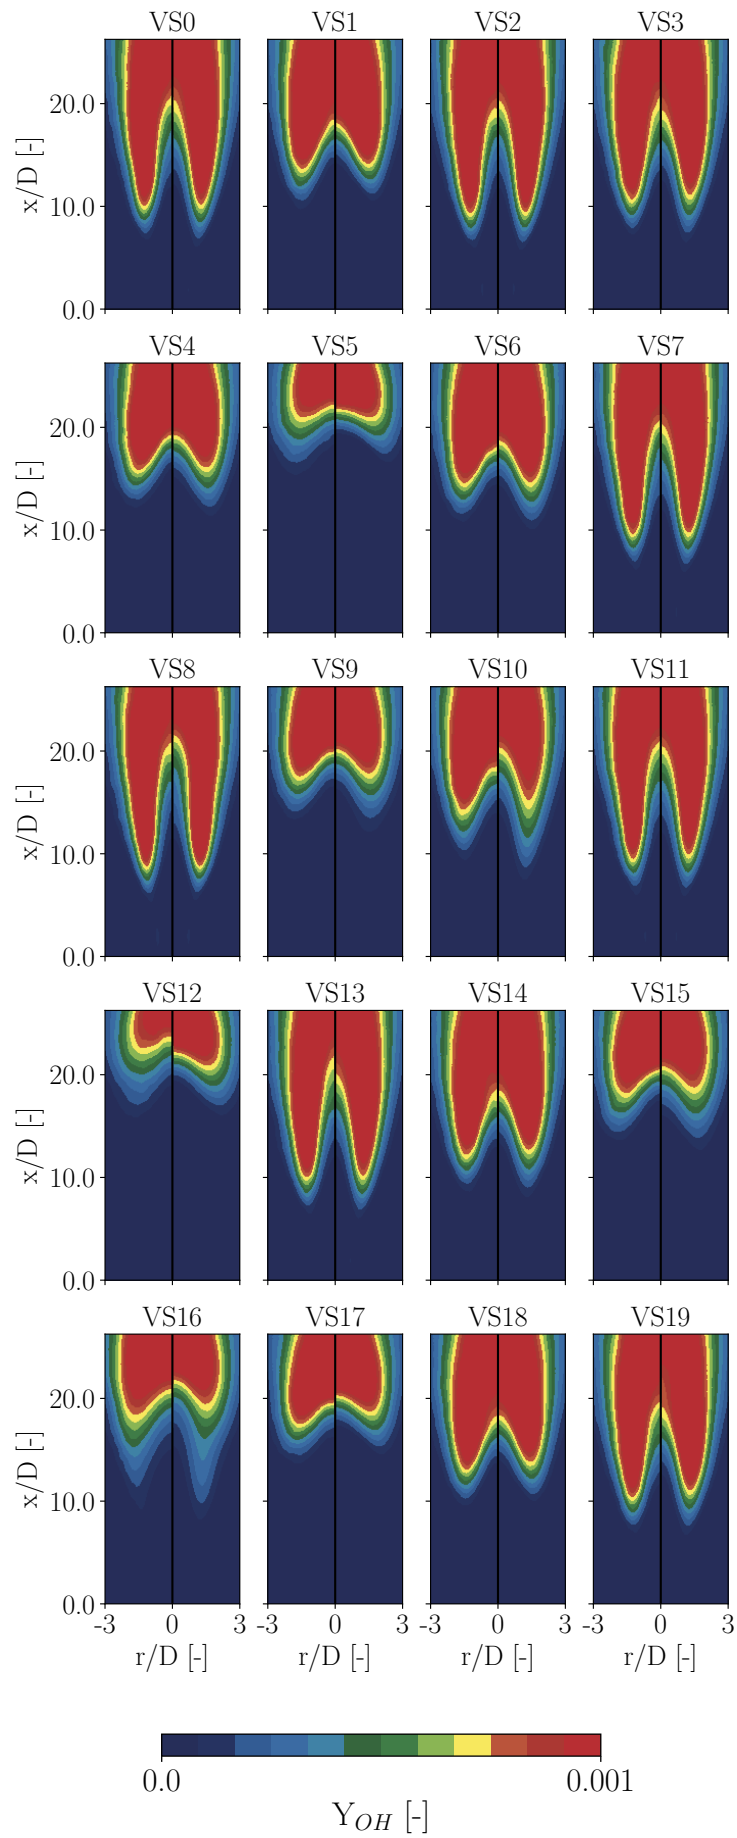


Figure 5.7: Y_{OH} fields of the validation samples **predicted** using the PCA modes and their coefficients predicted by their surrogates (**right**) compared with the LES (left)




Figure 5.8: Gas temperature fields of the validation samples **predicted** using the PCA modes and their coefficients predicted by their surrogates (**right**) compared with the LES (left)

5.5 Analysis of the uncertain fields

5.5.1 Comparison of the uncertain lift-off height with the results from Chapter 4

Chapters 3 and 4 highlighted the importance and the variability of the Cabra flame's lift-off height and its variations in the uncertain space. Therefore, a proper prediction of the time-averaged fields should also lead to a correct prediction of the lift-off height. Figure 5.9 shows a summary plot of the lift-off height, as observed in the LES (x-axis) and predicted using the method presented in this chapter (y-axis).




Figure 5.9: Summary plot of the lift-off height computed from the predicted fields, compared to the LES results

The agreement between the observations and the predictions is almost perfect. This further validates our field prediction method.

It is also possible to resample extensively the 3D uncertain space $\text{Span}(\xi_{A_{12}}, \xi_T, \xi_{C_D})$, predict the time-averaged fields, and determine a PDF of the lift-off height. This is shown in Figure 5.10.

Figure 5.10 shows a satisfying agreement between the PDF determined with the direct surrogate (in dotted line, from Chapter 4) and the PDF determined by post-processing the predicted fields. The general shape of the distribution is mostly respected. Perhaps more importantly, the key properties reported in Table 5.4 are extremely well reproduced.

Finally, Figure 5.10 shows a fundamental difference between both approaches. The lift-off height distribution determined from the predicted fields is bounded, as opposed to the non-bounded result of the direct approach from Chapter 4. This difference is related to an issue already reported in Section 5.2.2: The predicted fields are linear combinations of the PCA modes, which are themselves linear combinations of the fields observed in the training set. Therefore, while lift-off heights inside the training set range are extremely well predicted, as shown in Figure 5.9, our method cannot predict lift-off heights outside

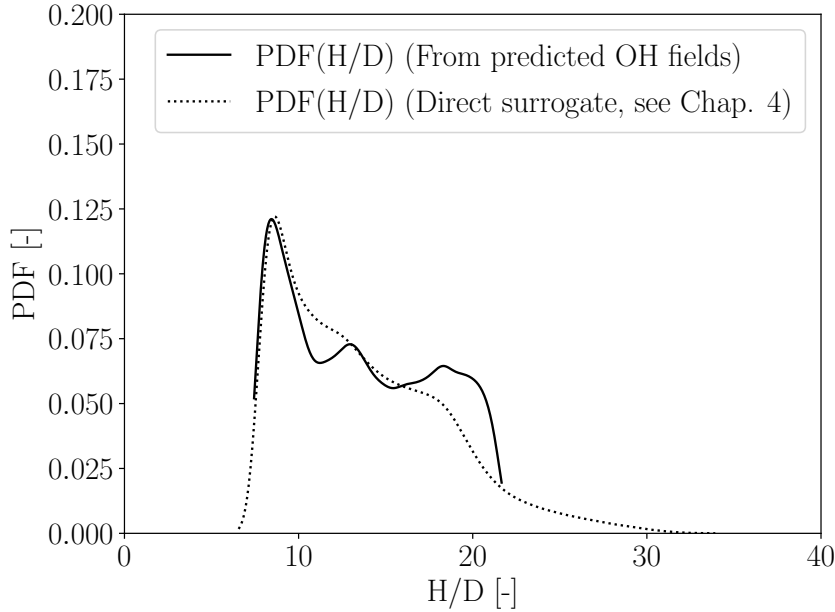


Figure 5.10: PDF of the Cabra flame lift-off height in $\text{Span}(\xi_{A_{12}}, \xi_T, \xi_{C_D})$, determined by post-processing the predicted field (this Chapter) in plain line, and determined by the direct surrogate model (cf. Chapter 4) in dotted line, using 2^{19} Quasi Monte Carlo samples

Table 5.4: Key features of the lift-off height distribution

	Direct surrogate	Predicted fields
Expectation	13.7	13.6
Standard deviation	4.7	4.2
Mode position	8.6	8.5

of this range. This is also the cause of the bulge in the PDF around $H/D = 20$: lift-off heights that should have been in the tail of the distribution are wrongly predicted at the upper end of the range of the training range, which, in turn, leads to an overestimation of the probability at this location.

To improve the prediction of the lift-off height, and of the general flame shape for the rare events above $H/D = 20$, observations of such flames should be included in the training set of the PCA. However, an over-representation of rare events may tamper with the optimal variance reproduction property of the PCA. Due to a lack of time and computational budget to run more LES, we did not investigate this aspect any further. Furthermore, despite this issue, the proposed method is already very satisfying in the probable regions.

5.5.2 Comparison of the uncertain fields with the multi-scalar measurements from Cabra et al. [11]

The ability to predict time-averaged fields given a set of coordinates in the input uncertain space allows the determination of a PDF for every time-averaged quantity at every location in the original LES computational domain. Among other possible use cases, it allows more relevant comparisons with experimental measurements than simply comparing the deterministic simulation with a set of experimental readings.

In this section, we complete the work of Section 3.3.2 by plotting the nominal, 50%, and 95% confidence intervals of the uncertain simulation for all the profiles investigated by Cabra et al. [11]. Figure 5.11 shows the axial profile of dioxygen mass fraction, Figure 5.12 shows several radial profiles of water vapor mass fraction, and Figure 5.13 shows several radial profiles of hydroxyl mass fraction. In these figures, we also represent the measurement uncertainties as reported by Cabra et al. All the other profiles are given in Annex I.




Figure 5.11: Nominal, 50%, and 95% confidence interval of the O_2 mass fraction’s axial profile in the LES, compared to the measurements from Cabra et al. [11]

The profiles shown in Section 3.3.2 showcased some minor disagreements between the experimental measurements from [11] and our deterministic LES. However, for the most part, these measurements fall well within our confidence interval when taking into account the uncertainties coming from the kinetic mechanism, the co-flow temperature, and the turbulence-combustion interaction model. Therefore, our modelization choices are suitable to explain most of the experimental measurements from [11].

Two discrepancies of note remain:

- The initial slope of the axial mixture fraction decrease, already outlined in Section 3.3.2, is still off (see Figure I.6). This issue is also noticeable in other works that used the same LES solver in a similar configuration (the CH_4 Cabra flame) [117, 97]. There may be an issue with the numerical setup that our numerous tests could not ferret out, or a problem with the experimental measurements.
- The hydroxyl mass fraction confidence interval is still 50 to 100% higher than the measurements from [11] in the far-downstream locations. These high values are directly linked to the mixture fraction values found in the UFPV table. In the nominal case, Figure B.4 shows maximum OH mass fraction values higher than 1000 ppm at $\tilde{C} = 1$ (end of reaction), increasing with the strain rate. The relatively high values of OH found downstream in the LES are therefore not surprising. However, we did not have the time to investigate further the discrepancy with experimental measurements in the Cabra flame.

The 95% confidence intervals displayed in these profiles are extremely large. Figure 5.13, in particular, is striking: At the flame foot, between $X = 9D$ and $X = 14D$, the hydroxyl mass fraction can range from 0 to over 1000 ppm, depending on whether

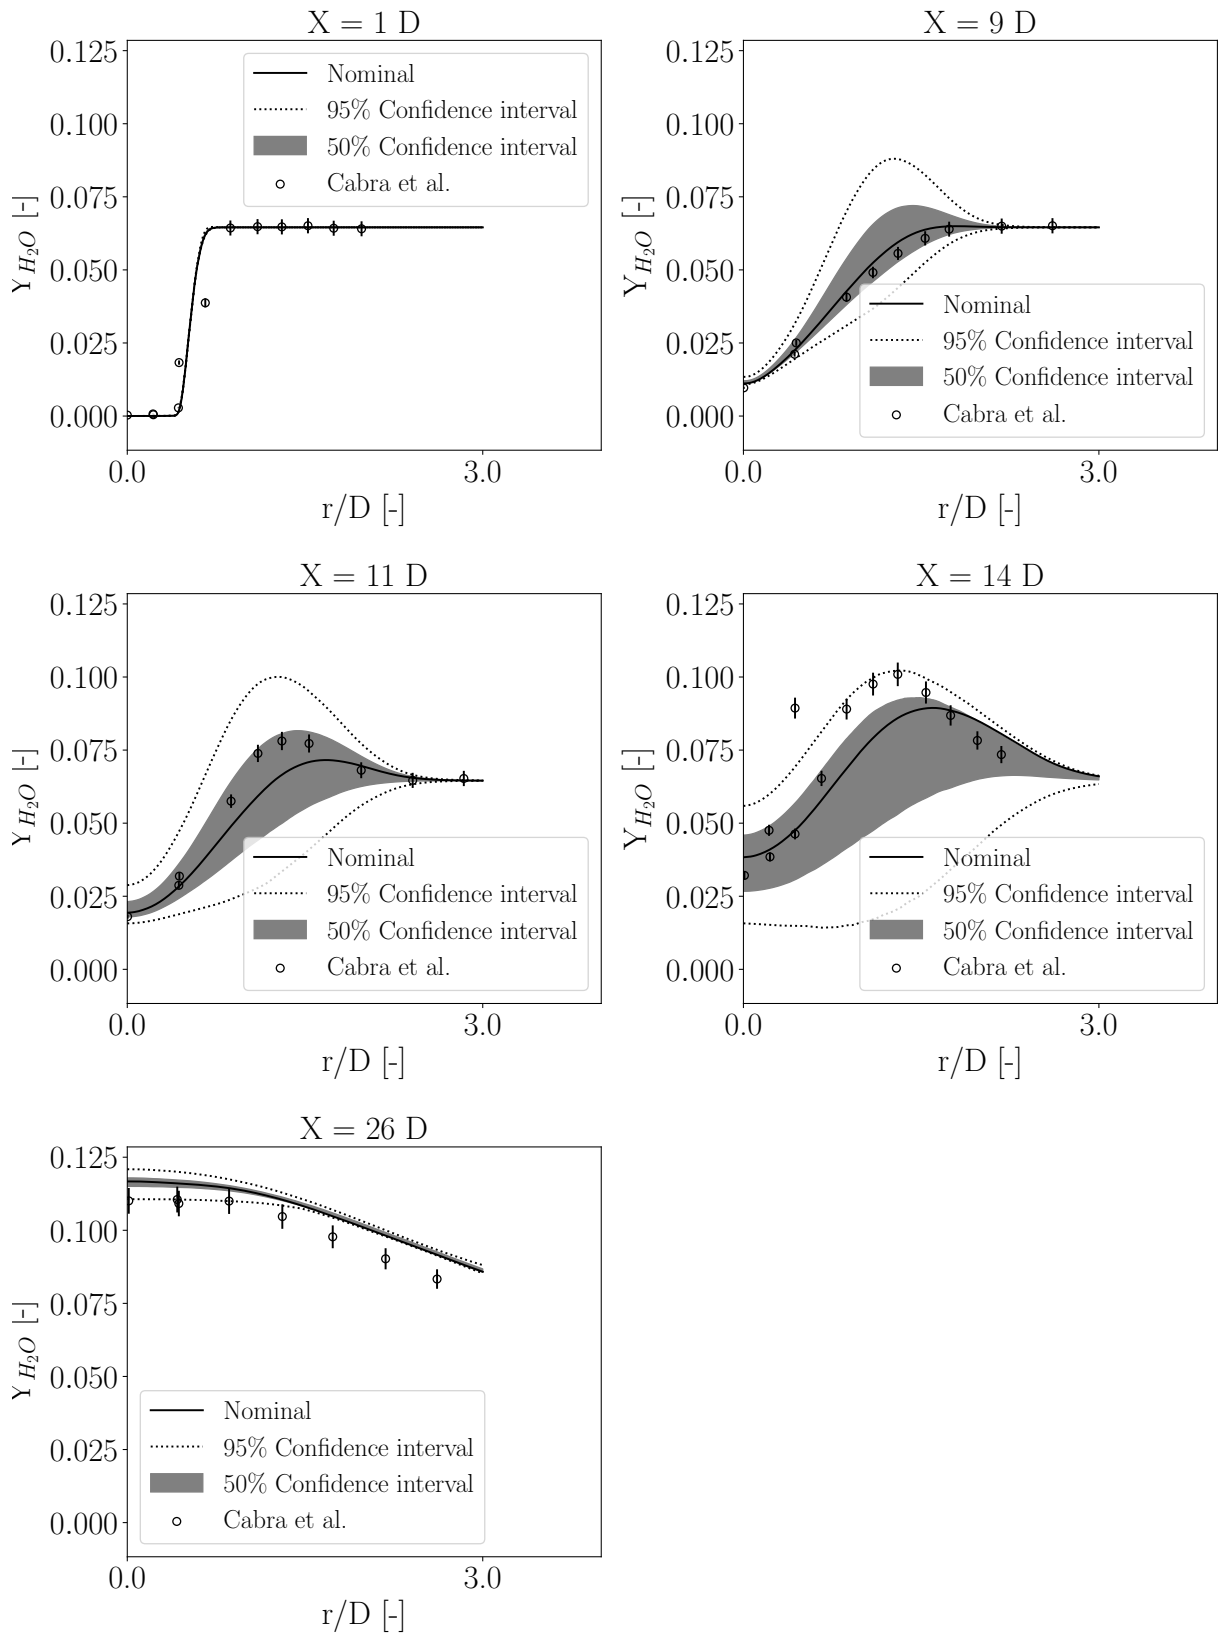


Figure 5.12: Nominal, 50%, and 95% confidence interval of the H_2O mass fraction's radial profiles in the LES, compared to the measurements from Cabra et al. [11]

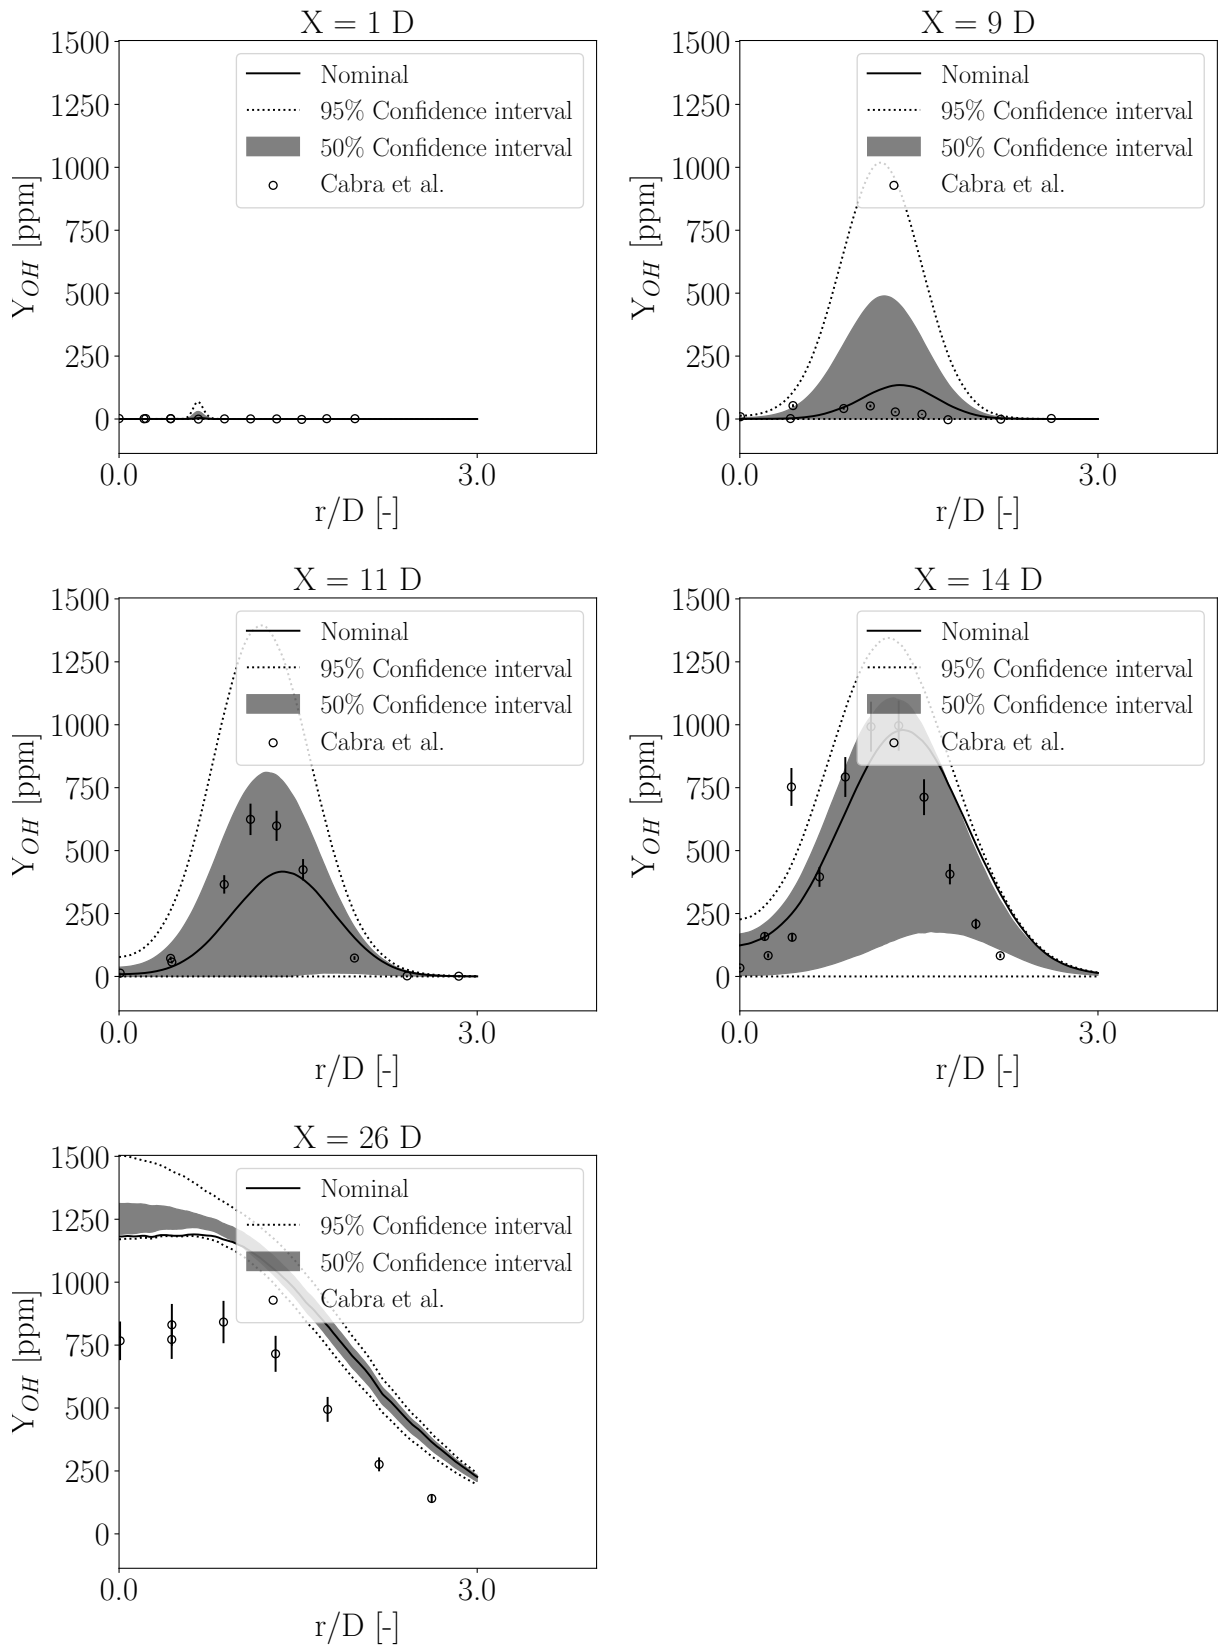


Figure 5.13: Nominal and 95% confidence interval of the OH mass fraction's radial profiles in the LES, compared to the measurements from Cabra et al. [11]

auto-ignition has already occurred. This puts into perspective the "high fidelity" characterization of Large Eddy Simulation if uncertainties are not taken into account.

In light of this study, we conclude that the Cabra flame is an excellent benchmark case for uncertainty quantification in reactive flow simulations, which would warrant further work from other groups, but a poor benchmark case for turbulence or combustion models in a deterministic framework.

5.6 Limitations of the study

This study presented in this chapter is subject to a few limitations, due to work hypotheses that would warrant further verifications and/or corrections. They are presented in the following in increasing order of importance.

Sampling uniformity

The sampling of the 3D uncertain space $\text{Span}(\xi_{A_{12}}, \xi_T, \xi_{C_D})$ was not performed according to its probability distribution: of the 39 training samples and 20 validation samples, 20 and 10 respectively were actually sampled from the 2D uncertain space $\text{Span}(\xi_{A_{12}}, \xi_T)$. This is due to our desire to recycle the LES results obtained for the study of Section 4.3. The idea was to enrich our knowledge of the 3D uncertain space with already-available observations in one of its subspaces.

To the author's knowledge, this inclusion of additional observations does not constitute an issue for the determination of the various surrogate models. On the contrary, it allows more accurate prediction on the plane $\text{Span}(\xi_{A_{12}}, \xi_T)$ and does not degrade the predictions elsewhere.

On the contrary, over-representing an area of the uncertain space (compared to its associated probability distribution) may tamper with the optimality of variance explained by the PCA. In this case, however, we found that the inclusion of the observations from the 2D study did not fundamentally alter the shape of the three principal components of the Cabra flame. Rather, it resulted in "smoother" modes by removing some minor artifacts.

Overall, in this case, the addition of more samples in the training set of the PCA improved the quality of the modes. Although, we would recommend, when adding new observations in such a training set, to sample according to the uncertain space's probability distribution when possible.

The dimension reduction of input uncertainties

In this chapter, we used the same dimension reduction of the kinetic uncertainties as in Chapter 4 without justifying further. We showed in Chapter 4 that using a dominant variable approach allowed neglecting most of the uncertain Arrhenius constants. However, this dimension reduction was shown valid only for the lift-off height determination. In the case of the whole flame shape, a more complete study of the sensitivity of each quantity (mass fractions, gas temperature, reaction source term...) in homogeneous reactors as a function of time should be performed to guarantee the relevance of the dimension reduction. Such a study can be found in Dumont's thesis [24].

However, given the predominance of the sensitivity of the lift-off height (or auto-ignition delay in the case of the homogeneous reactor) to the co-flow temperature (or initial temperature), compared to kinetic parameters, we are confident that the predicted PDF for each quantity at each location in the Cabra flame would not change much if more kinetic parameters had to be taken into account.

Choice of the probability distribution of C_D

Our choice of the probability distribution of C_D , both in terms of PDF shape and amplitude, is debatable. We chose what appeared to be the most sensible option at the time:

- A uniform distribution that reflects our lack of a priori knowledge on the "best" value to use
- A range that covered both the empirical range of values found in the literature and the theoretical range of values proposed by Ihme [41] for typical LES sizes. Confusingly, these two ranges were distinct, which lead us to consider a wider interval that included both empirical and theoretical ranges.

This wide distribution can be considered as a prior estimation of the parameter C_D that could be narrowed down in the future using bayesian optimization.

We show in Figure 5.14 the impact of choosing the most extreme cases of probability distributions for the Y_{H_2O} radial profile at $X = 11D$. On the left is the uncertain radial profile obtained with the uniform distribution of C_D considered in the whole chapter. On the right is the same uncertain profile obtained with a Dirac distribution of C_D at its nominal value ($C_D \sim \delta(20.0)$).




Figure 5.14: Uncertain radial profile of Y_{H_2O} at $X = 11D$. Left: with the uniform distribution of C_D . Right: With a Dirac of C_D , at its nominal value.

Figure 5.14 shows little differences between the uncertain profiles of Y_{H_2O} at $X = 11D$, aside from the location of the 95% confidence interval, although $X = 11D$ is a location

where the gas can be either ignited or not depending on the coordinates in the uncertain space. Similar conclusions can be reached by comparing other fields. This is consistent with Figure 4.18, in Chapter 4, which compared the PDF of the lift-off height of the Cabra flame in $\text{Span}(\xi_{A_{12}}, \xi_T, \xi_{C_D})$ and $\text{Span}(\xi_{A_{12}}, \xi_T)$: The probability distributions of every variable of interest are more heavy-tailed in the 3D uncertain space than they are in the 2D space, but the probable regions are similarly populated.

5.7 Takeaway points

On the methodological side:

- In this chapter, we propose a method to quantify the uncertainties of averaged LES fields.
- A combination of PCA, to reduce drastically the number of uncertain outputs, and surrogate modeling allows predicting every desired field at every location of the uncertain space.
- While our choices of uncertain inputs may be debatable, the method proposed in this chapter can be applied again, should further research reveal a need to change our uncertain inputs.

On the physical interpretation side:

- In this chapter, we predicted LES averaged fields using dimension reduction and surrogate modeling. This allowed the propagation of high-dimensional uncertainties and the determination of the probability distribution of every relevant field at any location in the computational domain.
- Taking into account state-of-the-art kinetic, temperature, and turbulence model uncertainties shows, again, that deterministic simulations of the Cabra flame are not suited to predict its behavior. This point is redundant with the previous chapter, but the confidence intervals of the averaged LES fields are even more eloquent.

Conclusion

Main achievements and results

- A set of uncertainty quantification methods, suitable for forward uncertainty quantification in the context of very expensive observations, has been selected and implemented in the in-house code SURE, which was entirely developed in this thesis. Its modularity and non-intrusive nature allow for combining different methods and interfacing with any "observations generator". For now, it has been interfaced with the in-house 0D-1D code AGATH, and with the LES code AVBP.
- The extreme uncertainty of H_2/O_2 auto-ignition simulations has been demonstrated in configurations where the transition between strong and weak combustion is involved. This uncertainty stems from the kinetic mechanism itself - which remains quite uncertain even when using state-of-the-art mechanisms - and from the uncertainty on the initial temperature of the mixture: uncertainties as small as $\pm 3\%$ in the vicinity of 1000K can lead to orders of magnitude of uncertainty in the auto-ignition delay. The resulting uncertainty of the auto-Ignition Delay Time is non-trivial: the probability distribution of the IDT can even be bimodal in the worst-case scenario.
- The deterministic Cabra flame was satisfyingly simulated using LES and UFPV tabulation.
- Thanks to an a priori study on a simplified configuration, which dramatically reduced the input uncertain dimension, we made affordable the surrogate-based uncertainty quantification of the Cabra flame lift-off height stemming from the kinetic mechanism and co-flow temperature. We also investigated the uncertainties introduced by the combustion-turbulence interaction model.
- Finally, we introduced a much richer uncertainty propagation study: by predicting entire averaged fields at any location in the uncertain space without resorting to LES other than for a few training observations, we were able to determine the probability distribution of every measured quantity and compare them with experimental measurements. This study also proved capable of rediscovering the results of the direct uncertainty quantification of the lift-off height.
- This work shows that, while recent works in chemistry greatly reduced the uncertainty of H_2/O_2 auto-ignition, the simulated Cabra flame remains very uncertain, mostly due to the experimental uncertainty on the co-flow temperature. This uncertainty affects first and foremost the lift-off height of the flame. In turn, every measured fields are eventually impacted.

A few words of nuance about the uncertainties uncovered in this thesis

The specificities of the Cabra flame

The large confidence intervals discovered in this study are likely not the norm in more conventional flame configurations. The Cabra flame is cursed - or gifted, depending on the point of view - with a very high sensitivity to parameters that happen to be slightly uncertain (the kinetic mechanism and the co-flow temperature) or that are lacking a scientific consensus, as far as the author is aware (the dissipation coefficient of mixture fraction variance).

The high sensitivity of the Cabra flame to these parameters is due to the driving mechanism of flame stabilization: auto-ignition delay. In more conventional configurations, the driving stabilization mechanism tends to be either mixing (for non-premixed flames) or the presence of low velocity and strain-rate areas behind a bluff body or in the recirculation of a swirler. In these configurations, the uncertain parameters investigated in this study have a much lesser sensitivity.

The choice of uncertain inputs

- We worked with a state-of-the-art kinetic mechanism, which assumed every parameter to be certain, except for the pre-exponential constants of the Arrhenius laws. This is a choice of the chemistry community to represent kinetic uncertainties that way. However, to the author's knowledge, there is no reason other than convention to assume the other parameters to be certain. Should further research from the chemistry community come up with an uncertain mechanism featuring uncertain activation energies as well as pre-exponential constants, the results of the uncertainty of the Cabra flame lift-off height may be different. In any case, the methodology proposed in this thesis can work with any parametric uncertainty in the kinetic mechanism.
- Section 5.6 already pointed out the arbitrariness of the choice of the C_D probability distribution. We chose to use the least informative distribution which covered both the empirical values found in the literature and those deduced from analytical developments. We are not claiming to hold any absolute truth about this probability distribution and another choice may have been wiser. Then again, our methodology can accommodate other distributions for C_D .
- Finally, other sources of uncertainty may have been overlooked.

Perspectives

Further improving the prediction of the averaged fields

This section proposes a few ideas to improve the study from Chapter 5.

- The statistical dependence of the mode coefficients may even have been an asset if we used a joint surrogate model of mode coefficients that takes into account the statistical dependence of its outputs, like co-kriging [51, 58]. Joint surrogate models were not implemented during this thesis due to a lack of time but would be an interesting improvement in future works.
- We saw that an important limitation of our results was our inability to predict rare events, in particular very high lift-off heights. We also saw that this inability was due to the lack of such examples in the training set, not due to an insufficient number of modes retained in the PCA. Training observations being expensive LES runs, adding more of them was not an option during this thesis. This problem is inherent to the PCA algorithm, and we leave finding a satisfying solution to future work.

Bayesian inference of the distributions of the co-flow temperature and dissipation coefficient of the mixture fraction variance: narrowing down our uncertain inputs

Our ability to predict the averaged fields of the Cabra flame nearly anywhere in the uncertain space indicates that we could also infer conjointly the probability distribution of the experimental co-flow temperature and of the "ideal" value for C_D using the experimental multi-scalar measurements from Cabra et al. [11]. We did not have the time to investigate this during this thesis. However, we believe that it would be very interesting for the combustion community:

- Reducing the uncertainty of the co-flow temperature in Cabra's experiment would allow numerical scientists to produce more relevant deterministic simulations, which would in turn allow more relevant combustion model comparisons on this flame.
- Fine-tuning the value of C_D would bring fresh arguments to the scientific debate over the choice of dissipation coefficient value.

Applying the methods to other use cases

The methods presented in this thesis could be applied to other combustor configurations. Given their non-intrusive nature, they could also be applied in other domains. In the chemistry sector, they could for example be applied to improve [116] the (currently) extremely polluting Haber-Bosch process that keeps afloat our intensive agriculture model. Finally, they may also be used in any scientific domain where expensive observations meet high-dimensional uncertainties and where cheaper sub-problems can be used to reduce the uncertain dimension a priori.

Bibliography

- [1] Assyr Abdulle and Alexei Medovikov. “Second order Chebyshev methods based on orthogonal polynomials”. In: *Numerische Mathematik* 90 (Nov. 2001), pp. 1–18.
- [2] David M. Allen. “The Relationship Between Variable Selection and Data Augmentation and a Method for Prediction”. In: *Technometrics* 16.1 (Feb. 1974), pp. 125–127.
- [3] Pierre Auzillon, Ronan Vicquelin, Olivier Gicquel, Nasser Darabiha, Denis Veynante, and Benoit Fiorina. *A Filtered Tabulated Chemistry Model for Large Eddy Simulation of Reactive Flows*. Orlando, United States, Jan. 4, 2010, AIAA–2010–205.
- [4] François Bachoc. “Cross Validation and Maximum Likelihood estimations of hyperparameters of Gaussian processes with model misspecification”. In: *Computational Statistics & Data Analysis* 66 (2013), pp. 55–69.
- [5] R. S. Barlow, A. N. Karpetsis, J. H. Frank, and J. Y. Chen. “Scalar profiles and NO formation in laminar opposed-flow partially premixed methane/air flames”. In: *Combustion and Flame* 127.3 (2001), pp. 2102–2118.
- [6] Michaël Baudin, Anne Dufloy, Bertrand Iooss, and Anne-Laure Popelin. “Open-TURNS: An industrial software for uncertainty quantification in simulation”. In: (2015).
- [7] D.L. Baulch, Drysdale D.D., Hrone D.G., and Lloyd A.C. “Evaluated kinetic data for high temperature reactions Volume I, Homogeneous gas phase reactions of the H₂-O₂ systems”. In: *Journal of Molecular Structure* 15.1 (1973), pp. 169–170.
- [8] Géraud Blatman. “Chaos polynomial creux et adaptatif pour la propagation d’incertitudes et l’analyse de sensibilité”. PhD thesis. Université Blaise Pascal - Clermont-Ferrand II, 2009.
- [9] Géraud Blatman and Bruno Sudret. “Adaptive sparse polynomial chaos expansion based on least angle regression”. In: *Journal of Computational Physics* 230.6 (2011), pp. 2345–2367.
- [10] Michael S. Hamada C. F. Jeff Wu. *Experiments: Planning, Analysis, and Optimization, 2nd Edition*. 2nd. Applied Probability & Statistics. Wiley Subscription Services, Inc., A Wiley Company, 2009.
- [11] R. Cabra, T. Myhrvold, J. Y. Chen, R. W. Dibble, A. N. Karpetsis, and R. S. Barlow. “Simultaneous laser raman-rayleigh-lif measurements and numerical modeling results of a lifted turbulent H₂/N₂ jet flame in a vitiated coflow”. In: *Proceedings of the Combustion Institute* 29.2 (2002), pp. 1881–1888.
- [12] Ricardo Cabra. *turbulent Jet Flames Into a Vitiated Coflow*. Tech. rep. University of California, Berkeley, 2004.
- [13] Renfeng Richard Cao, Stephen B. Pope, and Assaad R. Masri. “Turbulent lifted flames in a vitiated coflow investigated using joint PDF calculations”. In: *Combustion and Flame* 142.4 (2005), pp. 438–453.

- [14] C. W. Clenshaw and A. R. Curtis. “A method for numerical integration on an automatic computer”. In: *Numerische Mathematik* 2.1 (1960), pp. 197–205.
- [15] Olivier Colin and Michael Rudgyard. “Development of High-Order Taylor-Galerkin Schemes for LES”. In: *Journal of Computational Physics* 162 (2000), pp. 338–371.
- [16] Patrick R. Conrad and Youssef M. Marzouk. “Adaptive Smolyak Pseudospectral Approximations”. In: *SIAM Journal on Scientific Computing* 35.6 (2013), A2643–A2670.
- [17] P. Constantine, E. Dow, and Q. Wang. “Active Subspace Methods in Theory and Practice: Applications to Kriging Surfaces”. In: *SIAM Journal on Scientific Computing* 36.4 (2014), A1500–A1524. eprint: <https://doi.org/10.1137/130916138>.
- [18] P. G. Constantine, M. Emory, J. Larsson, and G. Iaccarino. “Exploiting active subspaces to quantify uncertainty in the numerical simulation of the HyShot II scramjet”. In: *Journal of Computational Physics* 302 (2015), pp. 1–20.
- [19] Paul G. Constantine and Paul Diaz. “Global sensitivity metrics from active subspaces”. In: *Reliability Engineering & System Safety* 162 (2017), pp. 1–13.
- [20] Paul G. Constantine, Michael S. Eldred, and Eric T. Phipps. “Sparse pseudospectral approximation method”. In: *Computer Methods in Applied Mechanics and Engineering* 229-232 (2012), pp. 1–12.
- [21] R Dennis Cook. *Regression graphics: Ideas for studying regressions through graphics*. Vol. 482. John Wiley & Sons, Ltd, 2009.
- [22] P Domingo, L. Vervisch, and D. Veynante. “Auto-ignition and flame propagation effects in LES of burnt gases diluted turbulent combustion”. In: *Proceedings of the Summer Program 2006, Center of turbulence Research*. 2006.
- [23] P. Domingo, L. Vervisch, and D. Veynante. “Large-eddy simulation of a lifted methane jet flame in a vitiated coflow”. In: *Combustion and Flame* 152.3 (2008), pp. 415–432. ISSN: 0010-2180.
- [24] Nicolas Dumont. “Méthodes numériques et modèle réduit de chimie tabulée pour la propagation d’incertitudes de cinétique chimique”. PhD thesis. Centrale Paris, 2019.
- [25] Bradley Efron, Trevor J. Hastie, Iain M. Johnstone, and Robert Tibshirani. “Least angle regression”. In: *Annals of Statistics* 32 (2004), pp. 407–499.
- [26] Gerhard Wanner Ernst Hairer. *Solving Ordinary Differential Equations II. Stiff and differential-algebraic problems*. Second Revised. Springer series in Computational mathematics 14. Springer-Verlag, 2000.
- [27] L. Fejér. “Mechanische Quadraturen mit positiven Cotesschen Zahlen”. In: *Mathematische Zeitschrift* 37.1 (1933), pp. 287–309.
- [28] B. Fiorina, R. Mercier, G. Kuenne, A. Ketelheun, A. Avdić, J. Janicka, D. Geyer, A. Dreizler, E. Alenius, C. Duwig, P. Trisjono, K. Kleinheinz, S. Kang, H. Pitsch, F. Proch, F. Cavallo Marincola, and A. Kempf. “Challenging modeling strategies for LES of non-adiabatic turbulent stratified combustion”. In: *Combustion and Flame* 162.11 (2015), pp. 4264–4282.
- [29] Rodney O. Fox. *Computational Models for Turbulent Reacting Flows*. Cambridge Series in Chemical Engineering. Cambridge: Cambridge University Press, 2003. ISBN: 9780521659079.
- [30] Jan N. Fuhg, Amélie Fau, and Udo Nackenhorst. “State-of-the-Art and Comparative Review of Adaptive Sampling Methods for Kriging”. In: *Archives of Computational Methods in Engineering* 28.4 (2021), pp. 2689–2747.

- [31] S. Geman and D. Geman. “Stochastic Relaxation, Gibbs Distributions, and the Bayesian Restoration of Images”. In: *IEEE Transactions on Pattern Analysis and Machine Intelligence* PAMI-6.6 (1984), pp. 721–741.
- [32] Thomas Gerstner and Michael Griebel. “Numerical integration using sparse grids”. In: *Numerical Algorithms* 18.3 (1998), p. 209.
- [33] R. L. Gordon, A. R. Masri, S. B. Pope, and G. M. Goldin. “A numerical study of auto-ignition in turbulent lifted flames issuing into a vitiated co-flow”. In: *Combustion Theory and Modelling* 11.3 (2007), pp. 351–376. ISSN: 1364-7830.
- [34] R. L. Gordon, S. H. Starner, A. R. Masri, and R. W. Bilger. “Further Characterisation of Lifted Hydrogen and Methane Flames issuing into a Vitiated Coflow”. In: *Proceedings of the 5th Asia-Pacific Conference on Combustion*. University of Adelaide, 2005, pp. 333–336.
- [35] Wang Han, Venkat Raman, and Zheng Chen. “LES/PDF modeling of autoignition in a lifted turbulent flame: Analysis of flame sensitivity to differential diffusion and scalar mixing time-scale”. In: *Combustion and Flame* 171 (Sept. 2016), pp. 69–86.
- [36] Charles R. Harris, K. Jarrod Millman, Stéfan J. van der Walt, Ralf Gommers, Pauli Virtanen, David Cournapeau, Eric Wieser, Julian Taylor, Sebastian Berg, Nathaniel J. Smith, Robert Kern, Matti Picus, Stephan Hoyer, Marten H. van Kerkwijk, Matthew Brett, Allan Haldane, Jaime Fernández del Río, Mark Wiebe, Pearu Peterson, Pierre Gérard-Marchant, Kevin Sheppard, Tyler Reddy, Warren Weckesser, Hameer Abbasi, Christoph Gohlke, and Travis E. Oliphant. “Array programming with NumPy”. In: *Nature* 585.7825 (2020), pp. 357–362.
- [37] W. K. Hastings. “Monte Carlo sampling methods using Markov chains and their applications”. In: *Biometrika* 57.1 (1970), pp. 97–109.
- [38] Edgar G. Hertwich and Glen P. Peters. “Carbon Footprint of Nations: A Global, Trade-Linked Analysis”. In: *Environmental Science & Technology* 43.16 (Aug. 2009), pp. 6414–6420.
- [39] Z. Hong, D. F. Davidson, E. A. Barbour, and R. K. Hanson. “A new shock tube study of the $H+O_2 \rightarrow OH+O$ reaction rate using tunable diode laser absorption of H_2O near 2.5microm”. In: *Proceedings of the Combustion Institute* 33.1 (2011), pp. 309–316.
- [40] Zekai Hong, David F. Davidson, and Ronald K. Hanson. “An improved H_2/O_2 mechanism based on recent shock tube/laser absorption measurements”. In: *Combustion and Flame* 158.4 (2011), pp. 633–644.
- [41] Matthias Ihme. “Pollutant formation and noise emission in turbulent non-premixed flames”. PhD thesis. Stanford University, 2007.
- [42] Matthias Ihme, Chong M. Cha, and Heinz Pitsch. “Prediction of local extinction and re-ignition effects in non-premixed turbulent combustion using a flamelet/progress variable approach”. In: *Proceedings of the Combustion Institute* 30.1 (2005), pp. 793–800.
- [43] Matthias Ihme and Yee Chee See. “Prediction of autoignition in a lifted methane/air flame using an unsteady flamelet/progress variable model”. In: *Combustion and Flame* 157.10 (2010), pp. 1850–1862.
- [44] William Stanley Jevons. *The Coal Question*. 2nd. Macmillan & Co., 1865.
- [45] Weiqi Ji, Zhuyin Ren, Youssef Marzouk, and Chung K. Law. “Quantifying kinetic uncertainty in turbulent combustion simulations using active subspaces”. In: *Proceedings of the Combustion Institute* 37.2 (2019), pp. 2175–2182. ISSN: 1540-7489.

- [46] J Jimnez, A Linan, MM Rogers, and FJ Higuera. “A priori testing of subgrid models for chemically reacting non-premixed turbulent shear flows”. In: *Journal of Fluid Mechanics* 349 (1997), pp. 149–171.
- [47] L. V. Kantorovich. “Mathematical Methods of Organizing and Planning Production”. In: *Management Science* 6.4 (1960), pp. 366–422.
- [48] Sergei P. Karkach and Vladimir I. Osharov. “Ab initio analysis of the transition states on the lowest triplet H₂O₂ potential surface”. In: *The Journal of Chemical Physics* 110.24 (1999), pp. 11918–11927.
- [49] John Kay and Mervyn King. *Radical Uncertainty, Decision-Making Beyond the Numbers*. 1st ed. W. W. Norton & Company, 2020.
- [50] A. M. Kempf, B. J. Geurts, and J. C. Oefelein. “Error analysis of large-eddy simulation of the turbulent non-premixed sydney bluff-body flame”. In: *Combustion and Flame* 158.12 (2011), pp. 2408–2419.
- [51] MC Kennedy and A O’Hagan. “Predicting the output from a complex computer code when fast approximations are available”. In: *Biometrika* 87.1 (2000), pp. 1–13.
- [52] Mohammad Khalil, Guilhem Lacaze, Joseph C. Oefelein, and Habib N. Najm. “Uncertainty quantification in LES of a turbulent bluff-body stabilized flame”. In: *Proceedings of the Combustion Institute* 35.2 (2015), pp. 1147–1156.
- [53] Stephen J. Klippenstein, Lawrence B. Harding, Michael J. Davis, Alison S. Tomlin, and Rex T. Skodje. “Uncertainty driven theoretical kinetics studies for CH₃OH ignition: HO₂+CH₃OH and O₂+CH₃OH”. In: *Proceedings of the Combustion Institute* 33.1 (2011), pp. 351–357.
- [54] Alexander A. Konnov. “Remaining uncertainties in the kinetic mechanism of hydrogen combustion”. In: *Combustion and Flame* 152.4 (2008), pp. 507–528. ISSN: 0010-2180.
- [55] Alexander A. Konnov. “Yet another kinetic mechanism for hydrogen combustion”. In: *Combustion and Flame* 203 (2019), pp. 14–22. ISSN: 0010-2180.
- [56] E. O. Kreyszig. “Advanced Engineering Mathematics”. In: 4th. ISBN 0-471-02140-7: Wiley, 1979, p. 880.
- [57] D. G. Krige. “A statistical approach to some basic mine valuation problems on the Witwatersrand”. In: *Journal of the Southern African Institute of Mining and Metallurgy* 52.6 (1951), pp. 119–139.
- [58] Loic Le Gratiet and Josselin Garnier. “Recursive co-kriging model for design of computer experiments with multiple levels of fidelity”. In: 4.5 (2014), pp. 365–386.
- [59] P. L’Ecuyer. “History of uniform random number generation”. In: *2017 Winter Simulation Conference (WSC)*. 2017, pp. 202–230. ISBN: 1558-4305.
- [60] Pierre L’Ecuyer, Christian Lécot, and Bruno Tuffin. “A Randomized Quasi-Monte Carlo Simulation Method for Markov Chains”. In: *Operations Research* 56.4 (2008), pp. 958–975.
- [61] Christiane Lemieux. *Monte Carlo and Quasi-Monte Carlo Sampling*. Springer Series in Statistics. Springer, 2009.
- [62] Guotu Li, Mohamed Iskandarani, Matthieu Le Hénaff, Justin Winokur, Olivier P. Le Maître, and Omar M. Knio. “Quantifying initial and wind forcing uncertainties in the Gulf of Mexico”. In: *Computational Geosciences* 20.5 (2016), pp. 1133–1153.

- [63] Juan Li, Zhenwei Zhao, Andrei Kazakov, and Frederick L. Dryer. “An updated comprehensive kinetic model of hydrogen combustion”. In: *International Journal of Chemical Kinetics* 36.10 (2004), pp. 566–575.
- [64] Riccardo Longo, Magnus Fürst, Aurélie Bellemans, Marco Ferrarotti, Marco Derudi, and Alessandro Parente. “CFD dispersion study based on a variable Schmidt formulation for flows around different configurations of ground-mounted buildings”. In: *Building and Environment* 154 (2019), pp. 336–347.
- [65] D. Lucor, J. Meyers, and P. Sagaut. “Sensitivity analysis of large-eddy simulations to subgrid-scale-model parametric uncertainty using polynomial chaos”. In: *Journal of Fluid Mechanics* 585 (2007), pp. 255–279.
- [66] Olivier Le Maître and Omar M Knio. *Spectral Methods for Uncertainty Quantification*. 1st ed. Scientific Computation. <https://doi.org/10.1007/978-90-481-3520-2>: Springer, Dordrecht, 2010.
- [67] Harry Markowitz. “Portfolio Selection”. In: *The Journal of Finance* 7.1 (1952), pp. 77–91.
- [68] Epaminondas Mastorakos. “Ignition of turbulent non-premixed flames”. In: *Progress in Energy and Combustion Science* 35.1 (2009), pp. 57–97.
- [69] Georges Matheron. *Le krigeage universel*. Vol. 1. Cahiers du Centre de morphologie mathématique de Fontainebleau. École nationale supérieure des mines de Paris, 1969.
- [70] Georges Matheron. “Principles of geostatistics”. In: *Economic Geology* 58.8 (1963), pp. 1246–1266.
- [71] M. D. McKay, R. J. Beckman, and W. J. Conover. “A Comparison of Three Methods for Selecting Values of Input Variables in the Analysis of Output From a Computer Code”. In: *Technometrics* 42.1 (Feb. 2000), pp. 55–61.
- [72] Michael E. Mueller, Gianluca Iaccarino, and Heinz Pitsch. “Chemical kinetic uncertainty quantification for Large Eddy Simulation of turbulent nonpremixed combustion”. In: *Proceedings of the Combustion Institute* 34.1 (2013), pp. 1299–1306.
- [73] Michael E. Mueller and Venkat Raman. “Model form uncertainty quantification in turbulent combustion simulations: Peer models”. In: *Combustion and Flame* 187 (2018), pp. 137–146.
- [74] Bertrand Naud, Ricardo Novella, JoséManuel Pastor, and Johannes F. Winklinger. “RANS modelling of a lifted H₂/N₂ flame using an unsteady flamelet progress variable approach with presumed PDF”. In: *Combustion and Flame* 162.4 (Apr. 2015), pp. 893–906.
- [75] Ron-Ho Ni. “A Multiple-Grid Scheme for Solving the Euler Equations”. In: *AIAA Journal* 20.11 (1982), pp. 1565–1571.
- [76] Franck Nicoud, Hubert Baya Toda, Olivier Cabrit, Sanjeeb Bose, and Jungil Lee. “Using singular values to build a subgrid-scale model for large eddy simulations”. In: *Physics of fluids* 23.8 (2011), p. 085106.
- [77] Yi-Shuai Niu, Luc Vervisch, and Pham Dinh Tao. “An optimization-based approach to detailed chemistry tabulation: Automated progress variable definition”. In: *Combustion and Flame* 160.4 (2013), pp. 776–785.
- [78] S. S. Patwardhan, Santanu De, K. N. Lakshmisha, and B. N. Raghunandan. “CMC simulations of lifted turbulent jet flame in a vitiated coflow”. English. In: *Proceedings of the Combustion Institute* 32.Part 2 (2009), 1705–1712. ISSN: 1540-7489.

- [79] Karl Pearson. “On lines and planes of closest fit to systems of points in space”. In: *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science* 2.11 (1901), pp. 559–572.
- [80] Norbert Peters. *Turbulent Combustion*. Cambridge Monographs on Mechanics. Cambridge university press, 2000.
- [81] Brian D. Phenix, Joanna L. Dinaro, Menner A. Tatang, Jefferson W. Tester, Jack B. Howard, and Gregory J. McRae. “Incorporation of parametric uncertainty into complex kinetic mechanisms: Application to hydrogen oxidation in supercritical water”. In: *Combustion and Flame* 112.1 (1998), pp. 132–146.
- [82] Heinz Pitsch. “Unsteady flamelet modeling of differential diffusion in turbulent jet diffusion flames”. In: *Combustion and Flame* 123.3 (2000), pp. 358–374.
- [83] Heinz Pitsch and Matthias Ihme. “An unsteady/flamelet progress variable method for LES of nonpremixed turbulent combustion”. In: *43rd AIAA Aerospace Sciences Meeting and Exhibit*. AIAA. 2005, p. 557.
- [84] Thierry Poinsot. “Boundary conditions for direct simulations of compressible viscous flows”. In: *Journal of Computational Physics* 101 (1992), pp. 104–129.
- [85] Dimitris Nicolas Politis and Joseph P. Romano. “The stationary bootstrap”. In: *Journal of the American Statistical Association* 89 (1994), pp. 1303–1313.
- [86] M T Reagan, H N Najm, B J Debusschere, O P Le Maître, O M Knio, and R G Ghanem. “Spectral stochastic uncertainty quantification in chemical systems”. In: 8.3 (2004), pp. 607–632.
- [87] Matthew T. Reagan, Habib N. Najm, Roger G. Ghanem, and Omar M. Knio. “Uncertainty quantification in reacting-flow simulations through non-intrusive spectral projection”. In: *Combustion and Flame* 132.3 (2003), pp. 545–555.
- [88] Hannah Ritchie, Max Roser, and Pablo Rosado. *Energy*. 2022.
- [89] Murray Rosenblatt. “Remarks on a Multivariate Transformation”. In: *The Annals of Mathematical Statistics* 23.3 (1952), pp. 470–472.
- [90] A. Saltelli, S. Tarantola, and K. P. S. Chan. “A Quantitative Model-Independent Method for Global Sensitivity Analysis of Model Output”. In: *Technometrics* 41.1 (Feb. 1999), pp. 39–56.
- [91] Thomas J. Santner, Brian J. Williams, and William I. Notz. *The Design and Analysis of Computer Experiments*. 2nd ed. 2197-568X. <https://doi.org/10.1007/978-1-4939-8847-1>; Springer, New York, NY, 2018.
- [92] S. G. Saytzev and R. I. Soloukhin. “Study of combustion of an adiabatically-heated gas mixture”. In: *Symposium (International) on Combustion* 8.1 (1961), pp. 344–347.
- [93] R. Schöbi, P. Kersaudy, B. Sudret, and J. Wiart. *Combining Polynomial Chaos Expansions and Kriging*. Tech. rep. Orange Labs, 80305 Issy-lesMoulineaux, France: Chair of Risk, Safety and Uncertainty Quantification [ETH Zurich], 2014.
- [94] Roland Schöbi and Bruno Sudret. “PCE-based imprecise Sobol’ indices”. In: *Proceedings of the 12th International Conference on Structural Safety and Reliability (ICOSSAR 2017)*. TU Verlag, 2017, pp. 513–520.
- [95] Roland Schöbi, Bruno Sudret, and Joe Wiart. “Polynomial-Chaos-Based Kriging”. In: *International Journal of Uncertainty Quantification* 5.2 (Feb. 2015), pp. 171–193.
- [96] Thilo Schonfeld and Michael Rudgyard. “Steady and Unsteady Flow Simulations Using the Hybrid Flow Solver AVBP”. In: *AIAA Journal* 37.11 (1999), pp. 1378–1385.

- [97] O. Schulz, T. Jaravel, T. Poinso, B. Cuenot, and N. Noiray. “A criterion to distinguish autoignition and propagation applied to a lifted methane–air jet flame”. In: *Proceedings of the Combustion Institute* 36.2 (2017), pp. 1637–1644.
- [98] David A. Sheen and Hai Wang. “The method of uncertainty quantification and minimization using polynomial chaos expansions”. In: *Combustion and Flame* 158.12 (2011), pp. 2358–2374.
- [99] David A. Sheen, Xiaoqing You, Hai Wang, and Terese Løvås. “Spectral uncertainty quantification, propagation and optimization of a detailed kinetic model for ethylene combustion”. In: *Proceedings of the Combustion Institute* 32.1 (2009), pp. 535–542.
- [100] Andrei V. Smirnov, S. Shi, and Ismail B. Celik. “Random Flow Generation Technique for Large Eddy Simulations and Particle-Dynamics Modeling”. In: *Journal of Fluids Engineering-transactions of The Asme* 123 (2001), pp. 359–371.
- [101] Gregory P Smith. “GRI-Mech 3.0”. In: [http://www. me. berkley. edu/gri_mech/](http://www.me.berkeley.edu/gri_mech/) (1999).
- [102] S. A. Smolyak. “Quadrature and interpolation formulas for tensor products of certain classes of functions”. In: *Doklady Akademii Nauk SSSR* 148 (1963), pp. 1042–1045.
- [103] I. M. Sobol’. “Global sensitivity indices for nonlinear mathematical models and their Monte Carlo estimates”. In: *Mathematics and Computers in Simulation* 55.1 (2001), pp. 271–280.
- [104] Fabio Stagnari, Angelica Galieni, Sara D’Egidio, Giancarlo Pagnani, and Michele Pisante. “Sustainable Soil Management”. In: *Innovations in Sustainable Agriculture*. Ed. by Muhammad Farooq and Michele Pisante. Cham: Springer International Publishing, 2019, pp. 105–131. ISBN: 978-3-030-23169-9.
- [105] I. Stanković, E. Mastorakos, and B. Merci. “LES-CMC Simulations of Different Auto-ignition Regimes of Hydrogen in a Hot Turbulent Air Co-flow”. In: *Flow, Turbulence and Combustion* 90.3 (2013), pp. 583–604.
- [106] M. Stone. “An Asymptotic Equivalence of Choice of Model by Cross-Validation and Akaike’s Criterion”. In: *Journal of the Royal Statistical Society: Series B (Methodological)* 39.1 (1977), pp. 44–47.
- [107] M. Stone. “Cross-Validatory Choice and Assessment of Statistical Predictions”. In: *Journal of the Royal Statistical Society: Series B (Methodological)* 36.2 (1974), pp. 111–133.
- [108] Boxin Tang. “Orthogonal Array-Based Latin Hypercubes”. In: *Journal of the American Statistical Association* 88.424 (Dec. 1993), pp. 1392–1397.
- [109] Robert Tibshirani. “Regression Shrinkage and Selection Via the Lasso”. In: *Journal of the Royal Statistical Society: Series B (Methodological)* 58.1 (1996), pp. 267–288.
- [110] A. N. Tikhonov and V. B. Glasko. “Use of the regularization method in non-linear problems”. In: *USSR Computational Mathematics and Mathematical Physics* 5.3 (1965), pp. 93–107.
- [111] Charles Tong. “Refinement strategies for stratified sampling methods”. In: *Reliability Engineering & System Safety* 91.10 (2006), pp. 1257–1265.
- [112] Lloyd N. Trefethen. “Is Gauss Quadrature Better than Clenshaw–Curtis?” In: *SIAM Review* 50.1 (2008), pp. 67–87.
- [113] Bruno Tuffin. “Randomization of Quasi-Monte Carlo Methods for Error Estimation: Survey and Normal Approximation*”. In: 10.3-4 (2004), pp. 617–628.

- [114] Arnold Tukker and Bart Jansen. “Environmental Impacts of Products: A Detailed Review of Studies”. In: *Journal of Industrial Ecology* 10.3 (2006), pp. 159–182.
- [115] L. N. Vaserstein. “Markov Processes over Denumerable Products of Spaces, Describing Large Systems of Automata”. In: *Probl. Peredachi Inf.* 5.3 (1969), pp. 64–72.
- [116] Kevin Verleysen, Alessandro Parente, and Francesco Contino. “How does a resilient, flexible ammonia process look? Robust design optimization of a Haber-Bosch process with optimal dynamic control powered by wind”. In: *Proceedings of the Combustion Institute* (2022).
- [117] Ronan Vicquelin. “Tabulated chemistry for turbulent combustion modeling and simulation”. PhD thesis. Centrale Paris, 2010.
- [118] C Wall, BJ Boersma, and P Moin. “An evaluation of the assumed beta probability density function subgrid-scale model for large eddy simulation of nonpremixed, turbulent combustion with heat release”. In: *Physics of Fluids* 12.10 (2000), pp. 2522–2529.
- [119] Hai Wang and David A. Sheen. “Combustion kinetic model uncertainty quantification, propagation and minimization”. In: *Progress in Energy and Combustion Science* 47 (2015), pp. 1–31. ISSN: 0360-1285.
- [120] Z. J. Wu, A. R. Masri, and R. W. Bilger. “An experimental investigation of the turbulence structure of a lifted H-2/N-2 jet flame in a vitiated co-flow”. In: *Flow, Turbulence and Combustion* 76.1 (2006), pp. 61–81. ISSN: 1386-6184.
- [121] H. Xu, U. Lee, and M. Wang. “Life-cycle energy use and greenhouse gas emissions of palm fatty acid distillate derived renewable diesel”. In: *Renewable and Sustainable Energy Reviews* 134 (2020), p. 110144.
- [122] R. A. Yetter, H. Rabitz, and R. M. Hedges. “A combined stability-sensitivity analysis of weak and strong reactions of hydrogen/oxygen mixtures”. In: *International Journal of Chemical Kinetics* 23.3 (1991), pp. 251–278.
- [123] Chun Sang Yoo, Jacqueline H. Chen, and Ramanan Sankaran. “A 3D DNS study of the stabilization of a turbulent lifted hydrogen/air jet flame in an autoignitive heated coflow”. In: *Fall Meeting of the Western States Section of the Combustion Institute* (2007).
- [124] Sajjad Yousefian, Gilles Bourque, and Rory F. D. Monaghan. “Uncertainty Quantification of NO_x Emission Due to Operating Conditions and Chemical Kinetic Parameters in a Premixed Burner”. In: *Journal of Engineering for Gas Turbines and Power* 140.12 (Oct. 2018), pp. 121005–121005–11.
- [125] Hui Zou and Trevor Hastie. “Regularization and variable selection via the elastic net”. In: *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 67.2 (2005), pp. 301–320.
- [126] I. Gy. Zsély, J. Zádor, and T. Turányi. “Uncertainty analysis of updated hydrogen and carbon monoxide oxidation mechanisms”. In: *Proceedings of the Combustion Institute* 30.1 (2005), pp. 1273–1281.

Annex A

The Boiled-down Konnov (2019) mechanism

Table A.1: H₂/O₂ combustion mechanism. $k = AT^\beta \exp(-E_a/RT)$, UF is the uncertainty factor of the pre-exponential constant A .

Index	Reaction	A	β	E _a	UF
1	H+H+M=H ₂ +M Enhanced third-body efficiencies (relative to Ar): O ₂ =0.0, N ₂ =0.0, H=0.0, H ₂ O=14.3	7.00E+17	-1	0	2
2	H+H+H ₂ =H ₂ +H ₂	1.00E+17	-0.6	0	2.5
3	H+H+N ₂ =O ₂ +N ₂	5.40E+18	-1.3	0	3.2
4	H+H+H=H ₂ +H	3.20E+15	0	0	3.2
5	O+O+M=O ₂ +M Enhanced third-body efficiencies (relative to Ar): O=28.8, O ₂ =8.0, N ₂ =2.0, H ₂ O=5.0	1.00E+17	-1	0	2
6	O+H+M=OH+M Enhanced third-body efficiencies: H ₂ O=5.0	6.75E+18	-1	0	3
7	H ₂ O+M=H+OH+M Enhanced third-body efficiencies (relative to Ar): H ₂ O=0.0, H ₂ =3.0, N ₂ =2.0, O ₂ =1.5, He=1.1	6.06E+27	-3.312	120770	2
8	H ₂ O+H ₂ O=H+OH+H ₂ O	1.00E+26	-2.44	120160	2
9	H+O ₂ (+M)=HO ₂ (+M) Low pressure limit Fcent = 0.5 Enhanced third-body efficiencies (relative to N ₂): Ar=0.72, H ₂ O=16.6, O ₂ =1.0, H ₂ =1.5, He=0.57	4.66E+12	0.44	0	1.2
		1.23E+19	-1.2	0	
10	H ₂ O ₂ (+M)=OH+OH(+M) Low pressure limit Fcent = 0.42 Enhanced third-body efficiencies (relative to Ar): H ₂ O=7.5, H ₂ O ₂ =7.7, O ₂ =1.2, N ₂ =1.5, He=0.65, H ₂ =3.7	2.00E+12	0.9	48750	1.5
		2.49E+24	-2.3	48750	
11	O+H ₂ =OH+H	5.08E+04	2.67	6292	1.3
12	H+O ₂ =OH+O	1.04E+14	0	15286	1.1
13	H ₂ +OH=H ₂ O+H	2.14E+08	1.52	3450	1.5
14	OH+OH=H ₂ O+O	2.67E+06	1.82	-1647	1.4
15	HO ₂ +O=OH+O ₂	2.85E+10	1	-723.9	2
16	H+HO ₂ =OH+OH	7.08E+13	0	300	2
17	H ₂ O+O=H+HO ₂	2.20E+08	2	61600	5
18	H ₂ +O ₂ =H+HO ₂	7.40E+05	2.43	53500	1.5
19	HO ₂ +OH=H ₂ O+O ₂	7.00E+12	0	-1093	2
20	HO ₂ +OH=H ₂ O+O ₂	4.50E+14	0	10930	2
21	HO ₂ +HO ₂ =H ₂ O ₂ +O ₂	1.03E+14	0	11040	2.5
22	HO ₂ +HO ₂ =H ₂ O ₂ +O ₂	1.94E+11	0	-1409	1.4
23	H ₂ O ₂ +H=HO ₂ +H ₂	5.02E+06	2.07	4300	3
24	H ₂ O ₂ +H=H ₂ O+OH	2.03E+07	2.02	2620	3
25	H ₂ O ₂ +O=HO ₂ +OH	9.55E+06	2	3970	3
26	H ₂ O ₂ +OH=HO ₂ +H ₂ O	1.74E+12	0	318	1.5
27	H ₂ O ₂ +OH=HO ₂ +H ₂ O	7.59E+13	0	7269	1.5
28	H+O ₂ +H=H ₂ +O ₂	8.80E+22	-1.835	800	2
29	H+O ₂ +H=OH+OH	4.00E+22	-1.835	800	2
30	H+O ₂ +O=OH+O ₂	7.35E+22	-1.835	800	2
31	H+O ₂ +OH=H ₂ O+O ₂	2.56E+22	-1.835	800	2

Annex B

The UFPV table




Figure B.1: Slices of the filtered mass fraction of H_2 , $\widetilde{Y}_{\text{H}_2}$, as a function of the filtered mixture fraction \tilde{z} and of the filtered normalized progress variable \tilde{C} , for different strain rates a and different mixture fraction segregation values S_Z




Figure B.2: Slices of the filtered mass fraction of O_2 $\widetilde{Y}_{\text{O}_2}$ as a function of the filtered mixture fraction \widetilde{z} and of the filtered normalized progress variable \widetilde{C} , for different strain rates a and different mixture fraction segregation values S_z

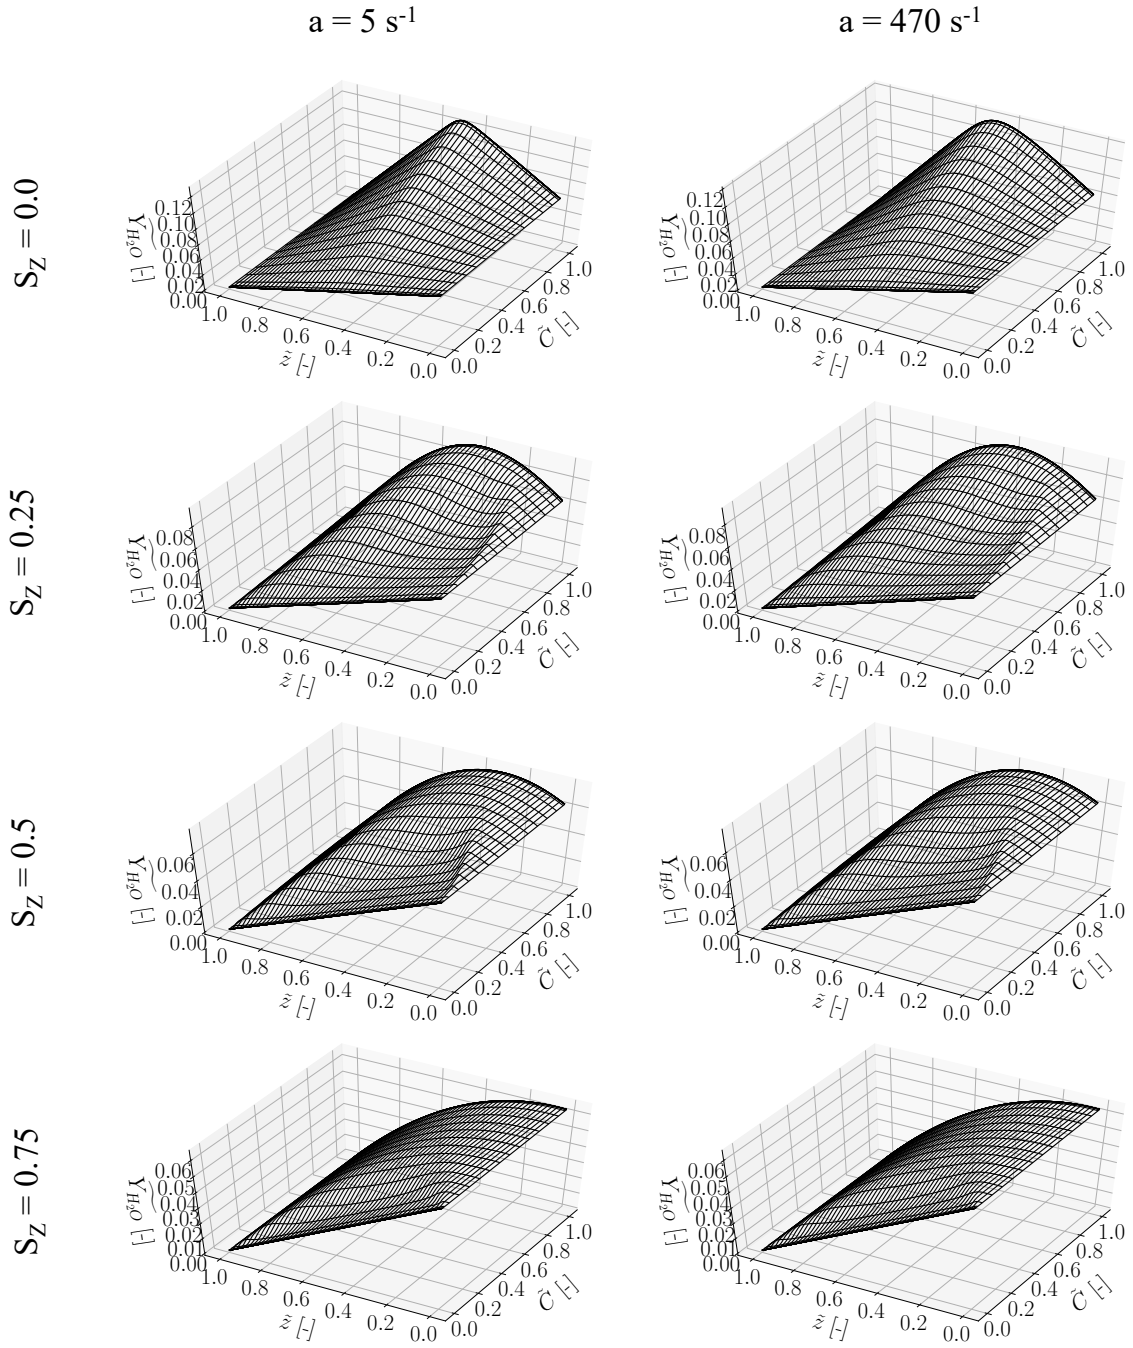


Figure B.3: Slices of the filtered mass fraction of H_2O $\widetilde{Y}_{\text{H}_2\text{O}}$ as a function of the filtered mixture fraction \tilde{z} and of the filtered normalized progress variable \tilde{C} , for different strain rates a and different mixture fraction segregation values S_z




Figure B.4: Slices of the filtered mass fraction of OH \widetilde{Y}_{OH} as a function of the filtered mixture fraction \tilde{z} and of the filtered normalized progress variable \tilde{C} , for different strain rates a and different mixture fraction segregation values S_Z . Note the evolution of the vertical scale




Figure B.5: Slices of the filtered mass fraction of HO_2 \tilde{Y}_{HO_2} as a function of the filtered mixture fraction \tilde{z} and of the filtered normalized progress variable \tilde{C} , for different strain rates a and different mixture fraction segregation values S_Z . Note the evolution of the vertical scale




Figure B.6: Slices of the filtered mass fraction of H_2O_2 $\widetilde{Y}_{\text{H}_2\text{O}_2}$ as a function of the filtered mixture fraction \widetilde{z} and of the filtered normalized progress variable \widetilde{C} , for different strain rates a and different mixture fraction segregation values S_z . Note the evolution of the vertical scale




Figure B.7: Slices of the filtered temperature \tilde{T} as a function of the filtered mixture fraction \tilde{z} and of the filtered normalized progress variable \tilde{C} , for different strain rates a and different mixture fraction segregation values S_z




Figure B.8: Slices of the filtered specific heat capacity at constant pressure \tilde{c}_p as a function of the filtered mixture fraction \tilde{z} and of the filtered normalized progress variable \tilde{C} , for different strain rates a and different mixture fraction segregation values S_z




Figure B.9: Slices of the filtered progress variable source term $\widetilde{\dot{\omega}}_{Y_c}$ as a function of the filtered mixture fraction \tilde{z} and of the filtered normalized progress variable \tilde{C} , for different strain rates a and different mixture fraction segregation values S_z . Note the evolution of the vertical scale

Annex C

Profiles of the simulated Cabra flame




Figure C.1: Axial profile of the mixture fraction Z in the LES, compared to the measurements from [11]




Figure C.2: Axial profile of the mass fraction of H_2 Y_{H_2} in the LES, compared to the measurements from [11]




Figure C.3: Axial profile of the mass fraction of O_2 Y_{O_2} in the LES, compared to the measurements from [11]




Figure C.4: Axial profile of the mass fraction of H_2O Y_{H_2O} in the LES, compared to the measurements from [11]




Figure C.5: Axial profile of the mass fraction of OH Y_{OH} in the LES, compared to the measurements from [11]




Figure C.6: Axial profile of the temperature T in the LES, compared to the measurements from [11]




Figure C.7: Radial profiles of the mixture fraction Z in the LES, compared to the measurements from [11]

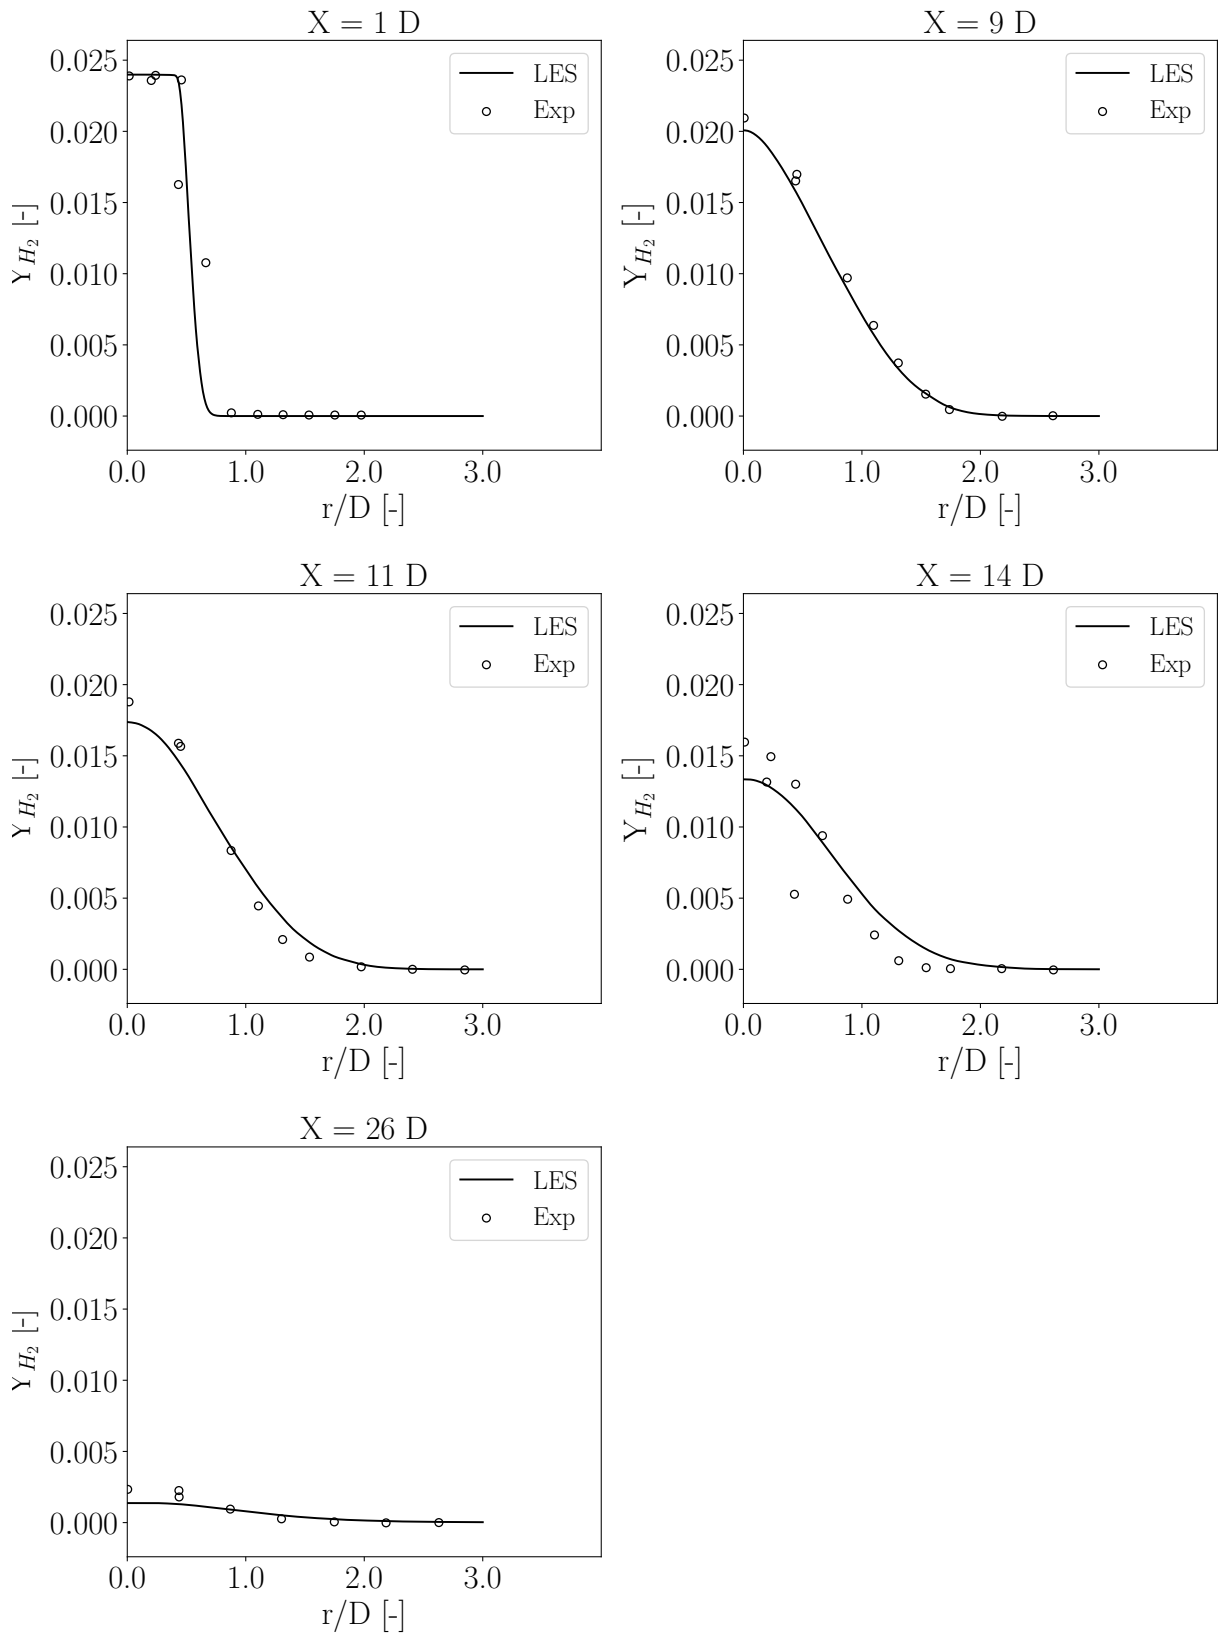


Figure C.8: Radial profiles of the mass fraction of H₂ Y_{H_2} in the LES, compared to the measurements from [11]

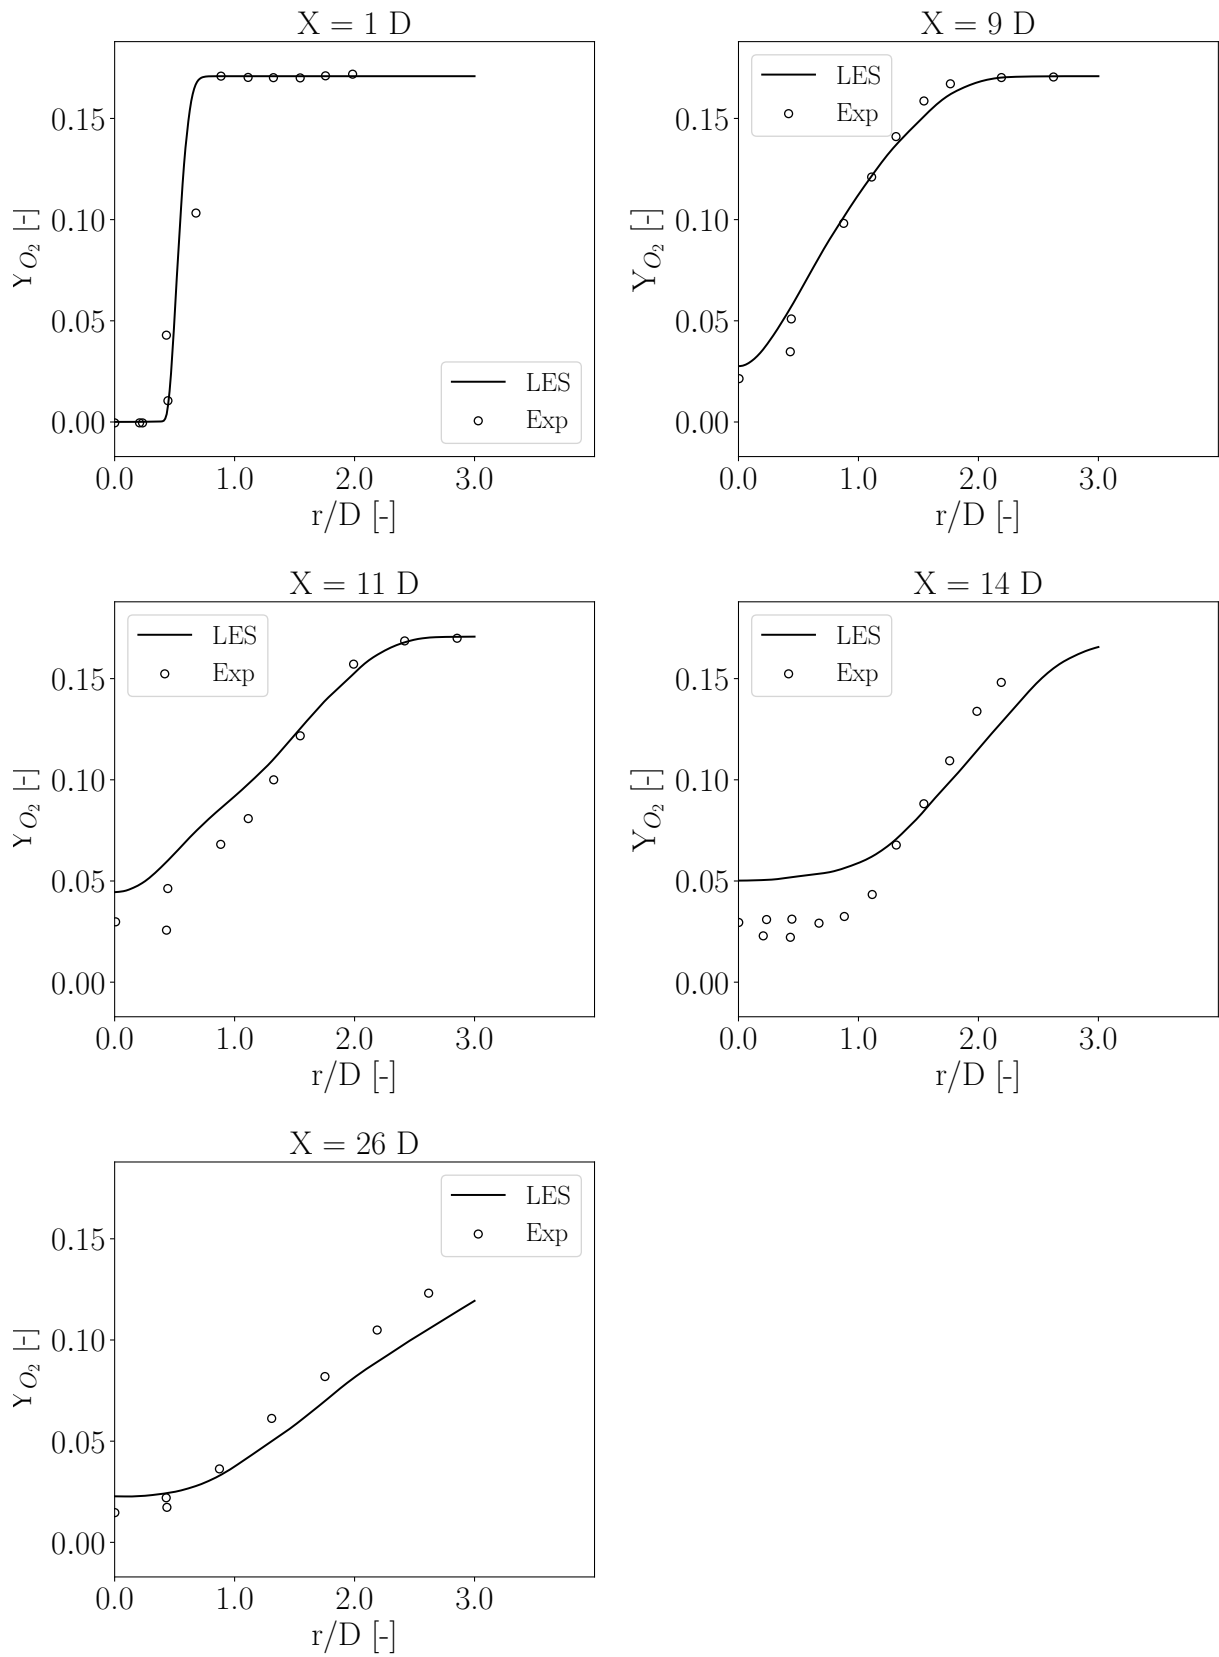


Figure C.9: Radial profiles of the mass fraction of O_2 Y_{O_2} in the LES, compared to the measurements from [11]

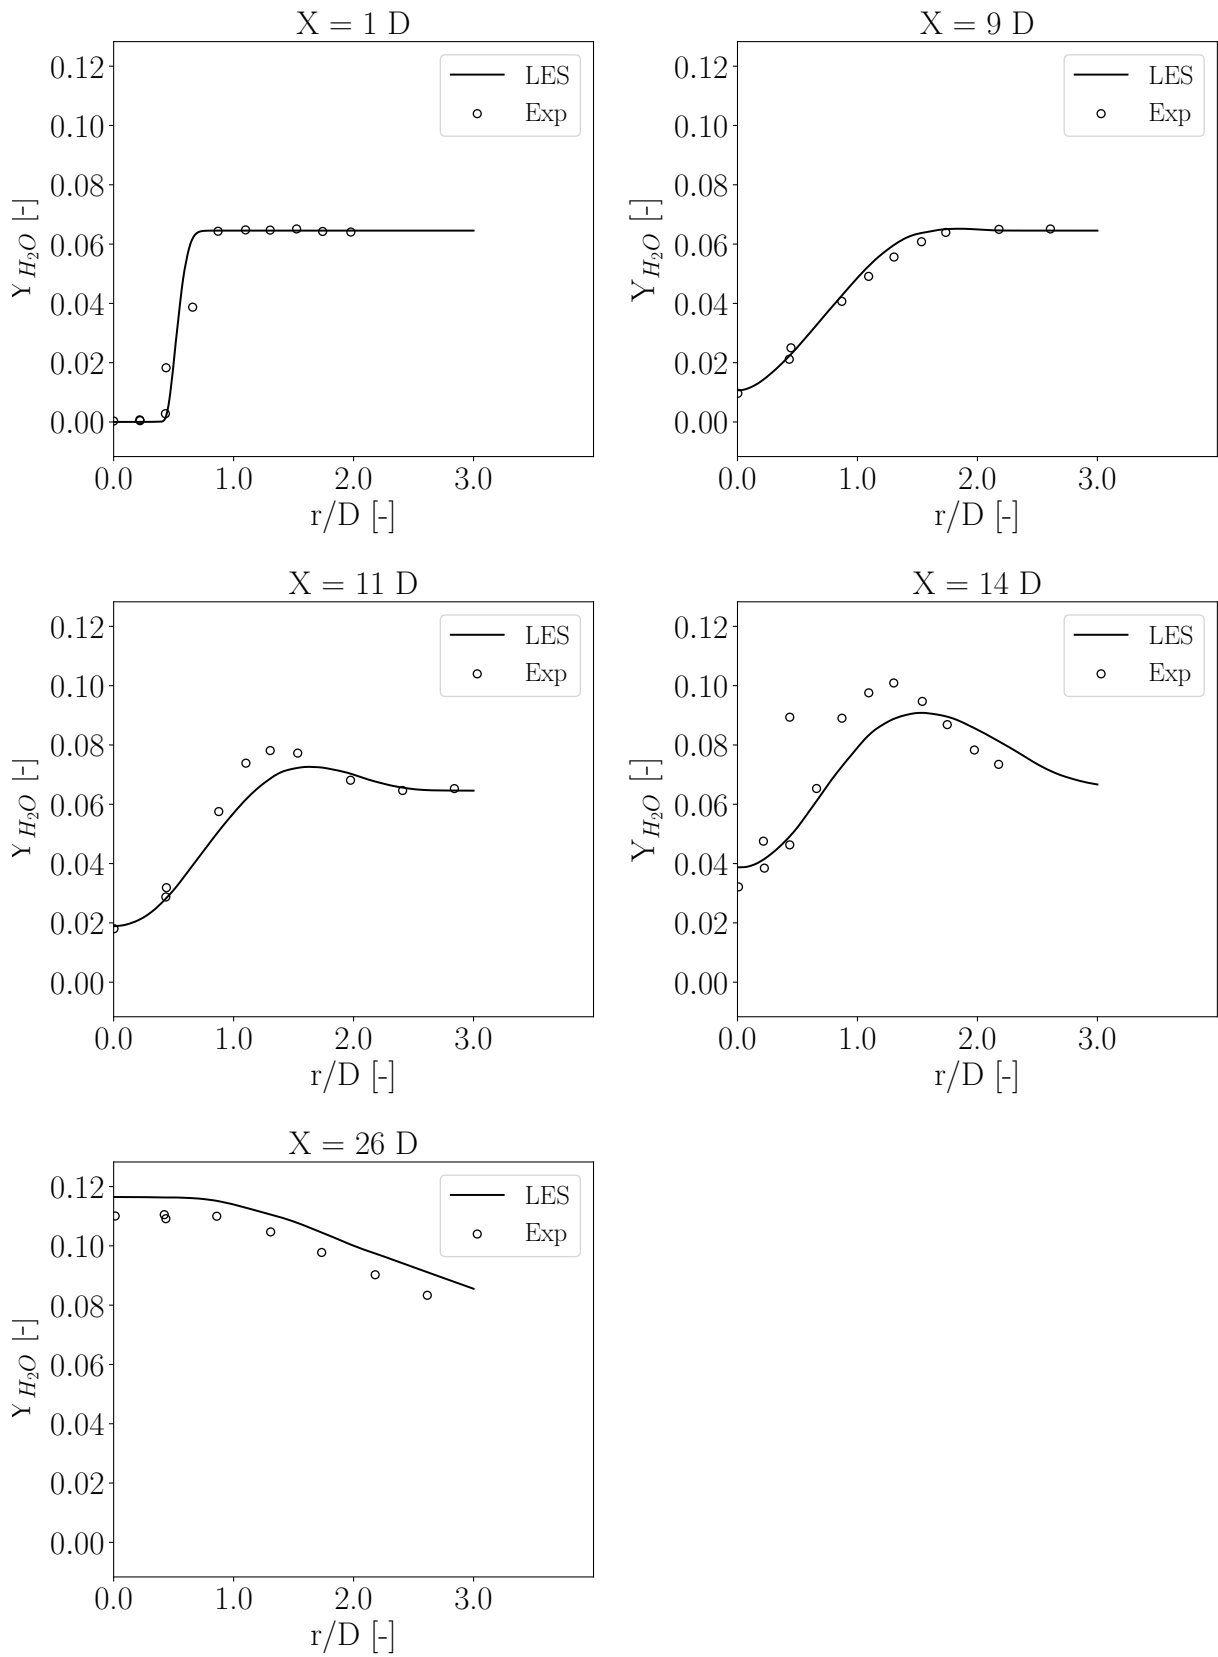


Figure C.10: Radial profiles of the mass fraction of H_2O Y_{H_2O} in the LES, compared to the measurements from [11]

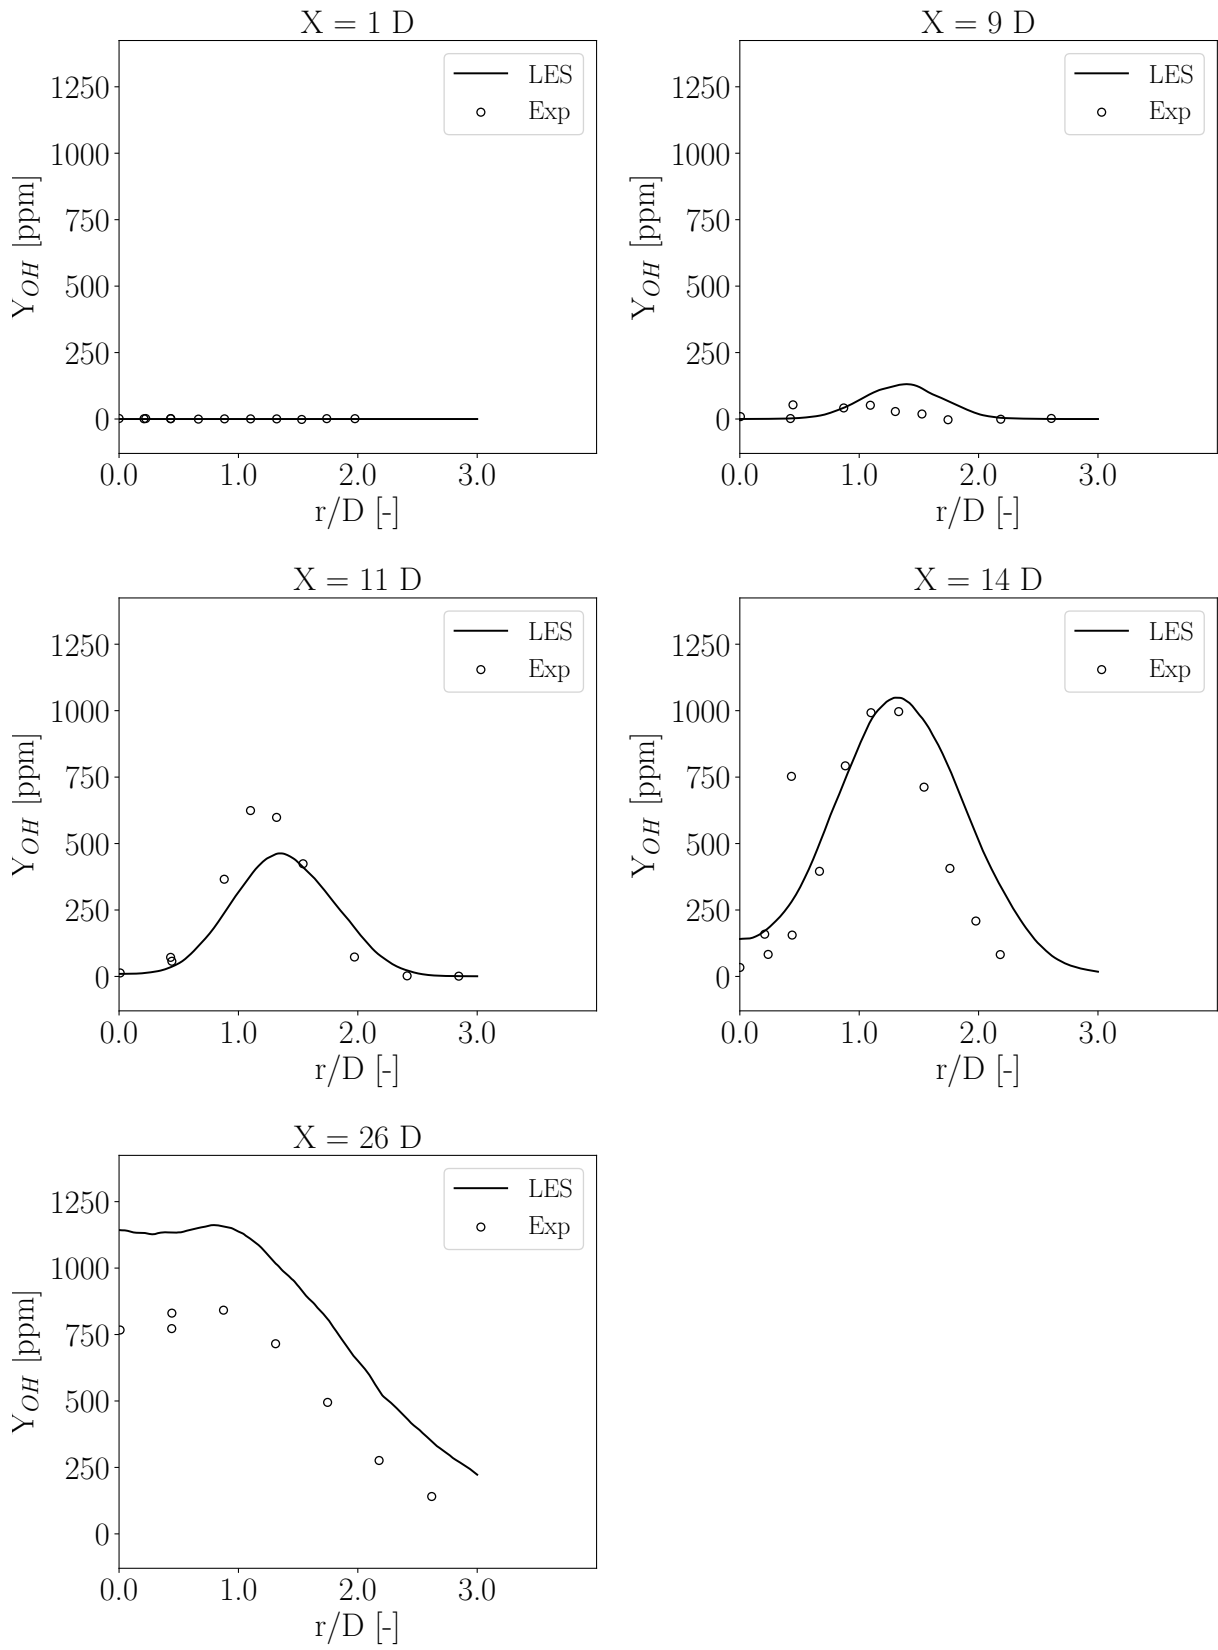


Figure C.11: Radial profiles of the mass fraction of OH Y_{OH} in the LES, compared to the measurements from [11]




Figure C.12: Radial profiles of the temperature T in the LES, compared to the measurements from [11]

Annex D

Convergence study for the LES of the Cabra flame

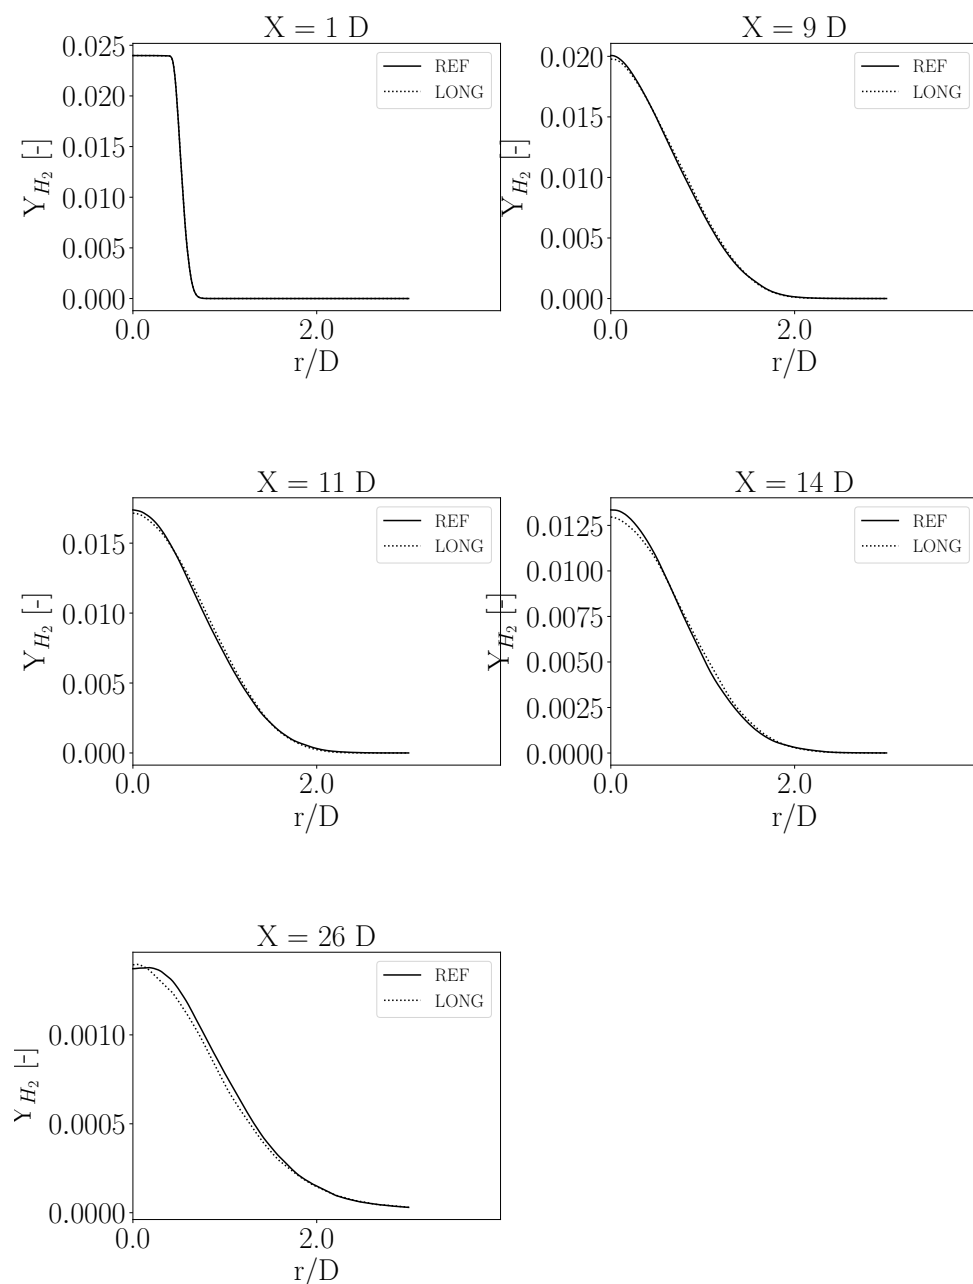


Figure D.1: Temporal convergence of the hydrogen mass fraction radial profiles

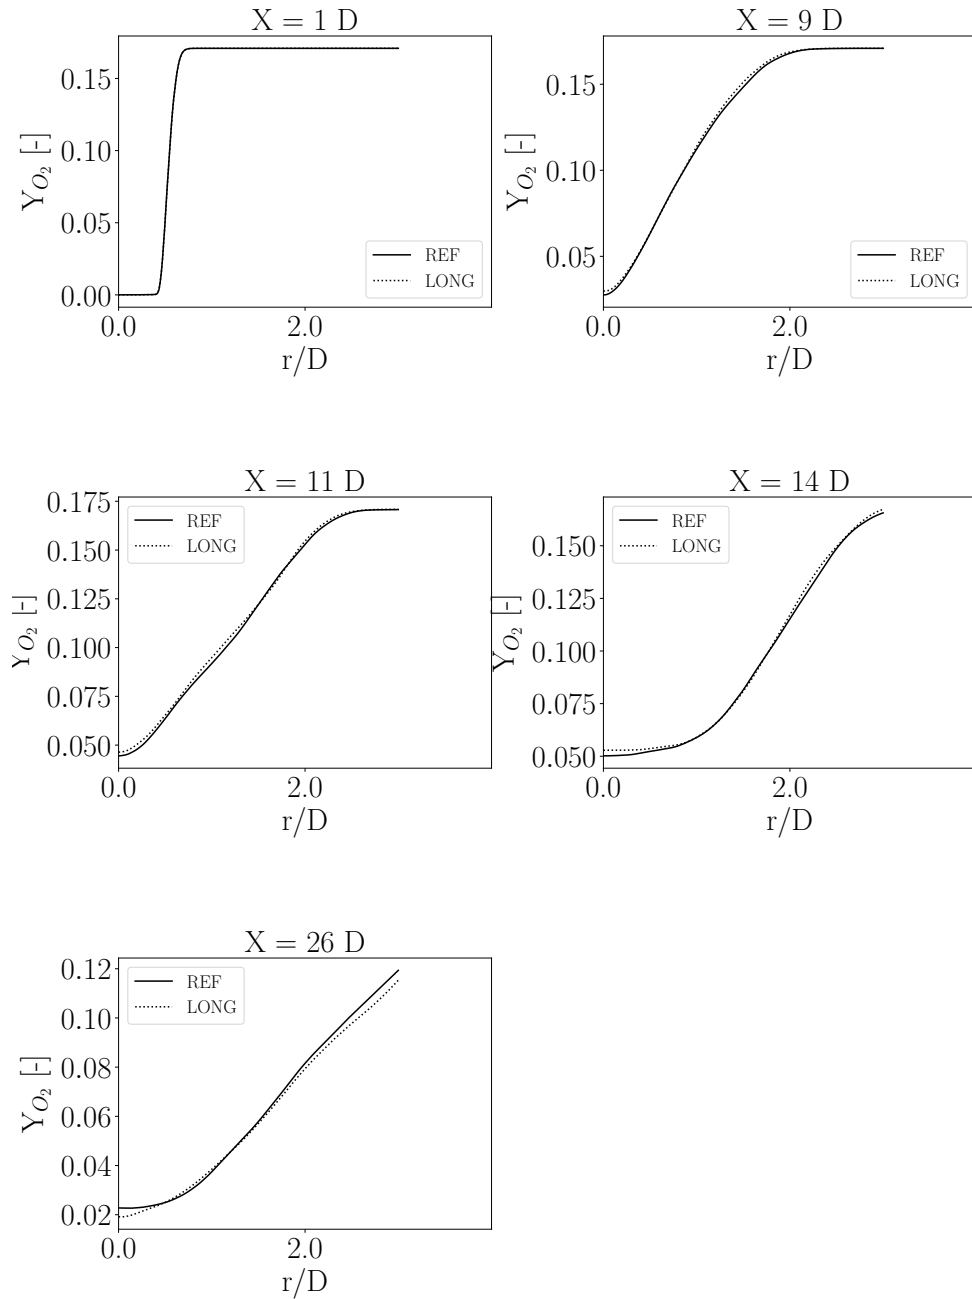


Figure D.2: Temporal convergence of the oxygen mass fraction radial profiles

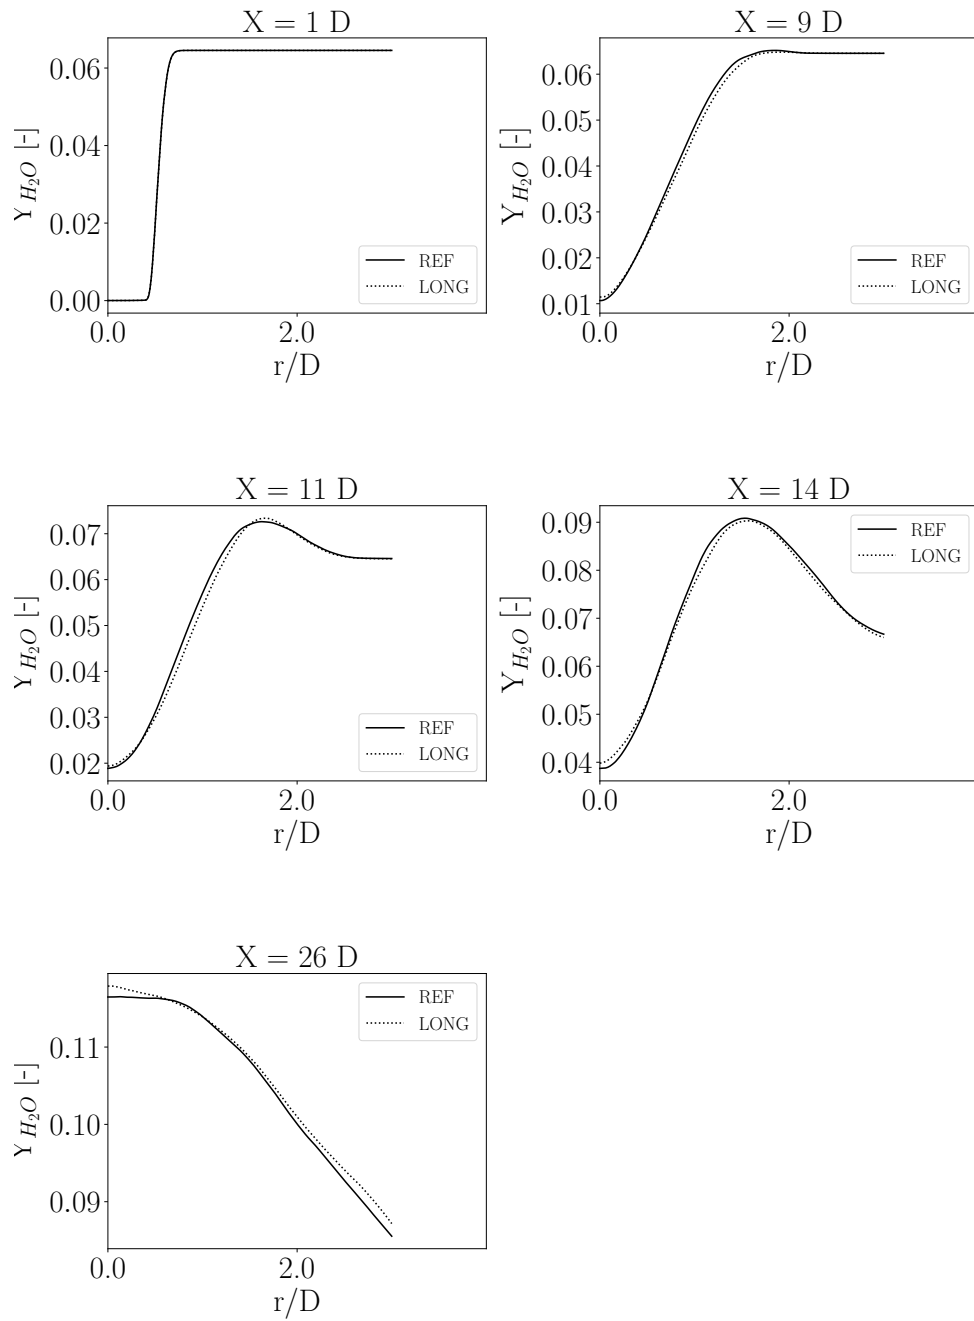


Figure D.3: Temporal convergence of the water vapor mass fraction radial profiles




Figure D.4: Temporal convergence of the hydroxyl mass fraction radial profiles

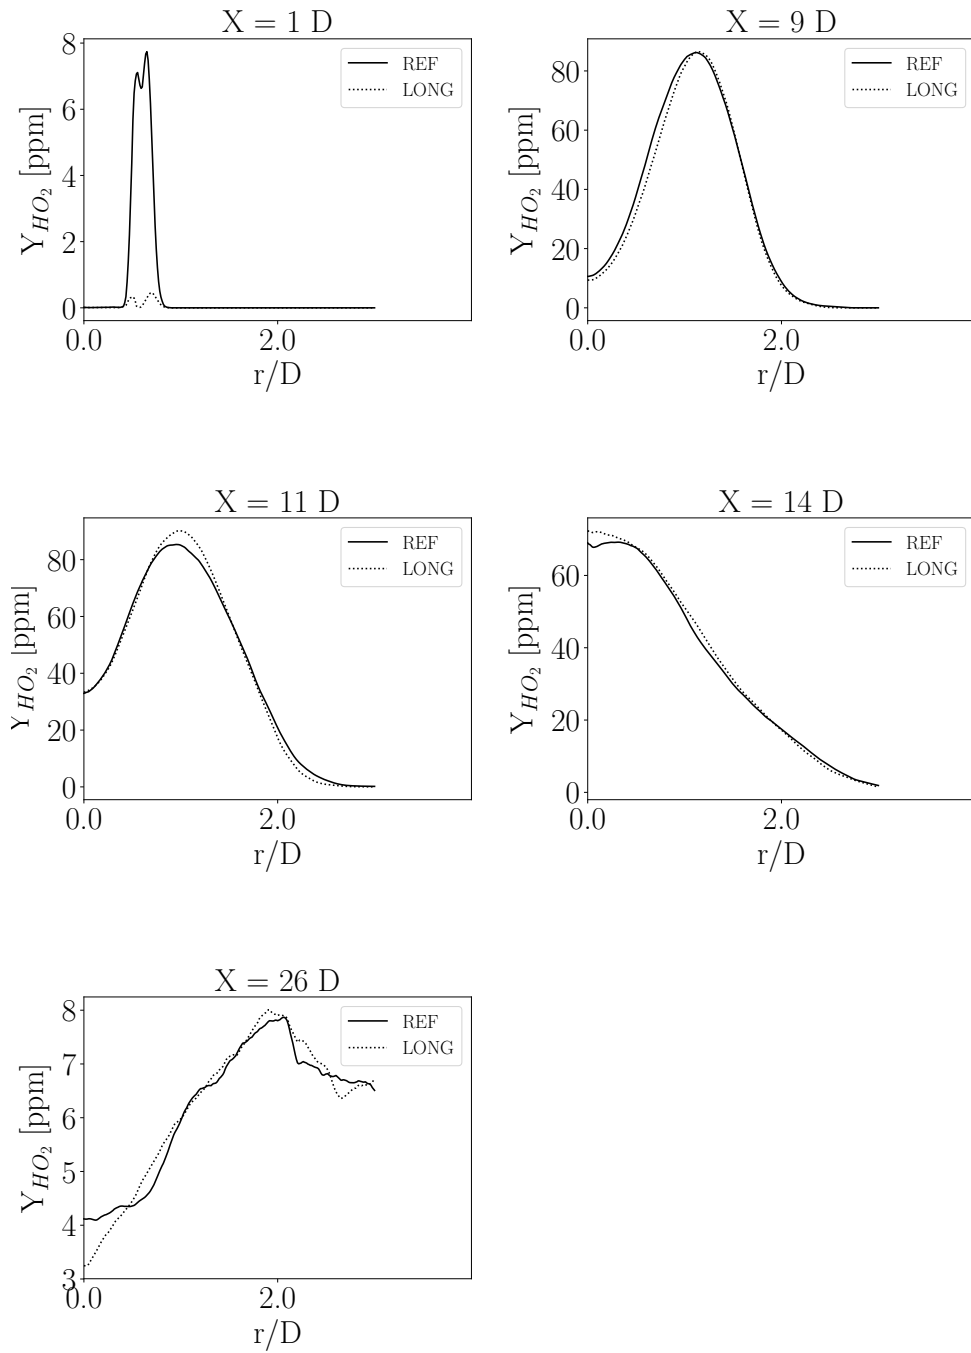


Figure D.5: Temporal convergence of the HO₂ mass fraction radial profiles

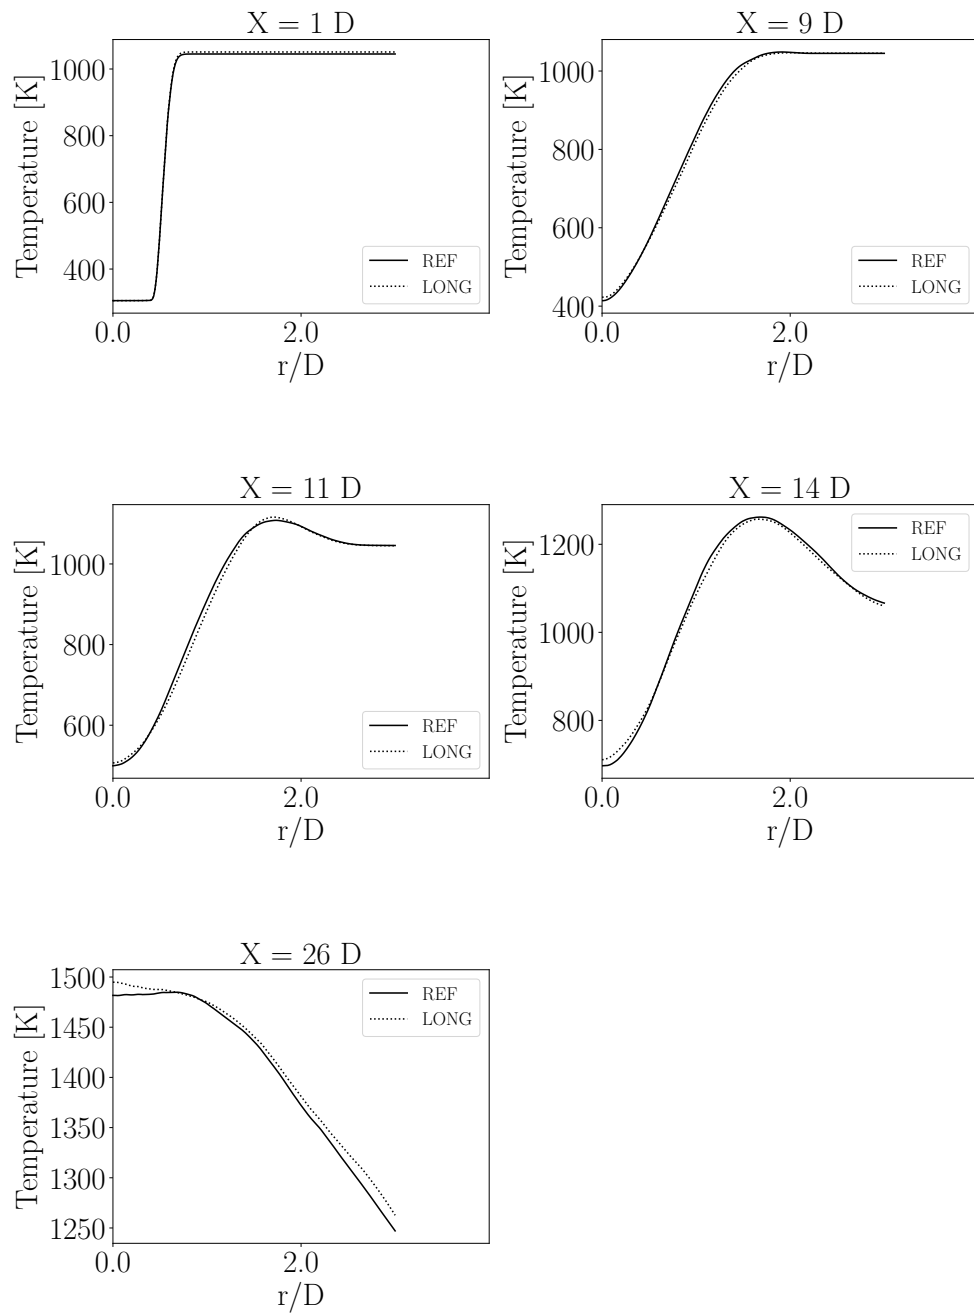


Figure D.6: Temporal convergence of the temperature radial profiles

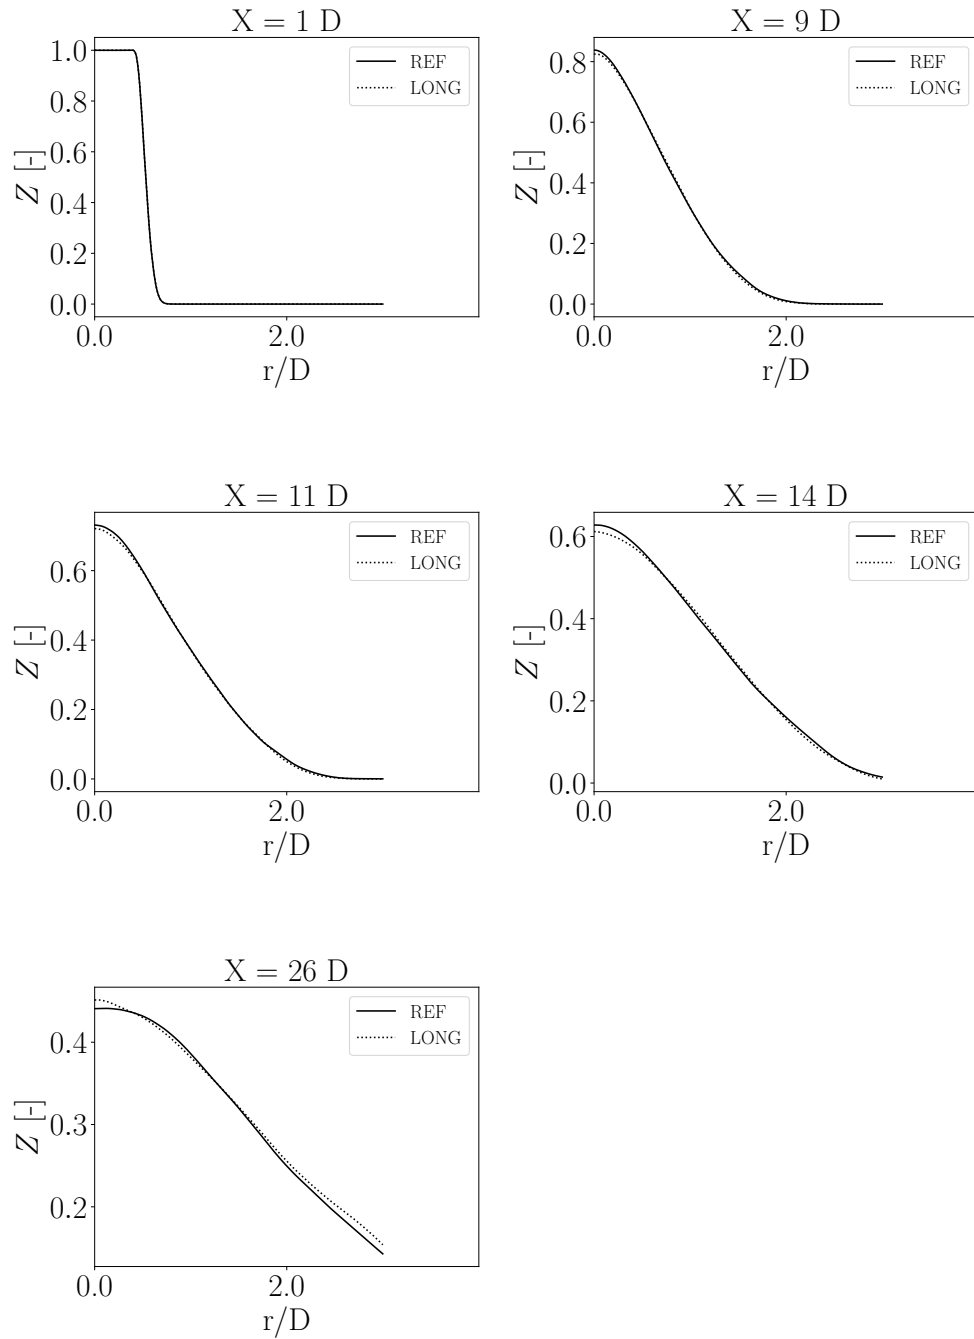


Figure D.7: Temporal convergence of the mixture fraction radial profiles

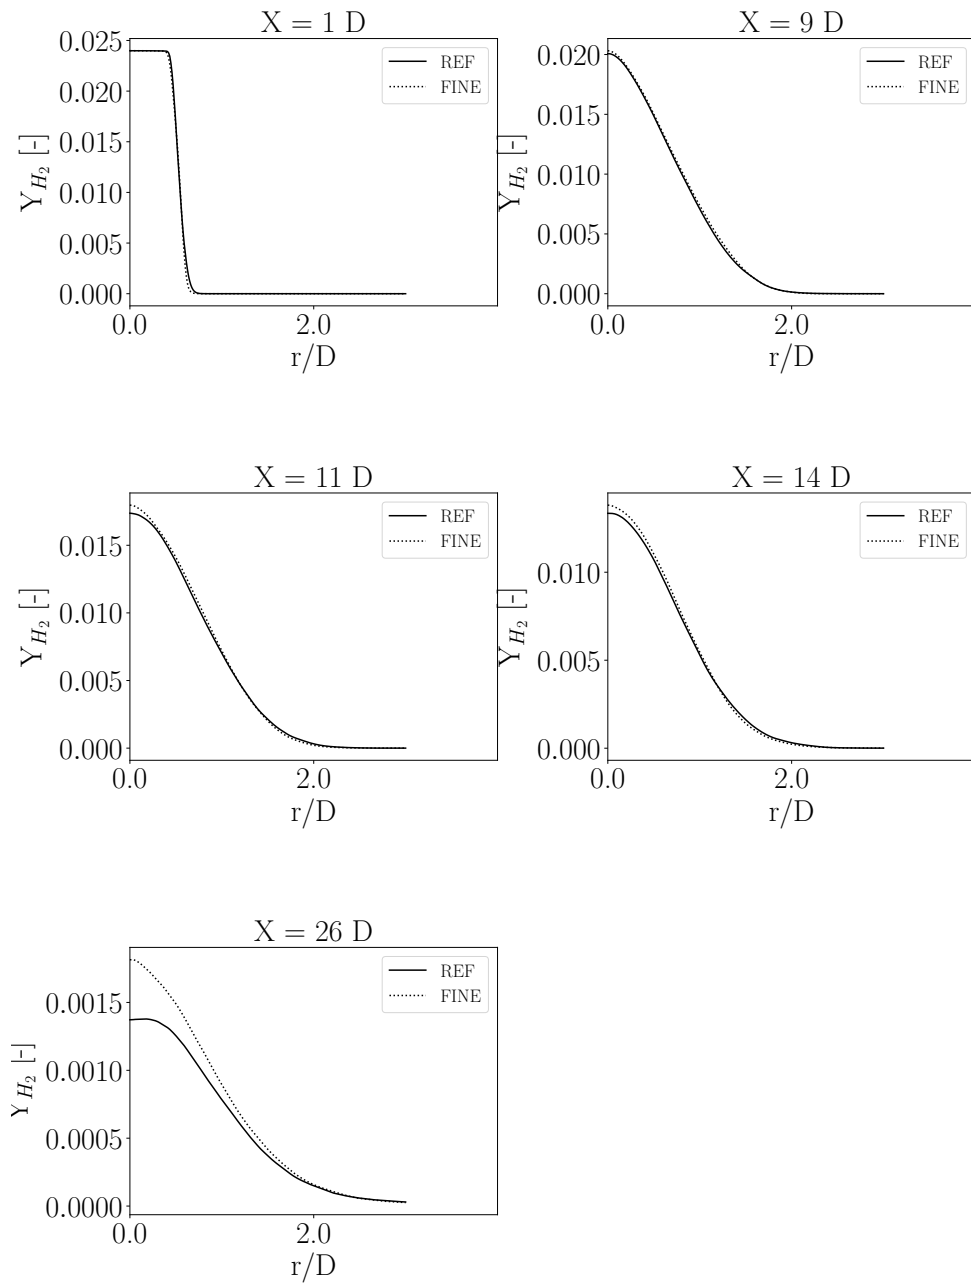


Figure D.8: Comparison of the hydrogen mass fraction radial profiles for the original and finer mesh

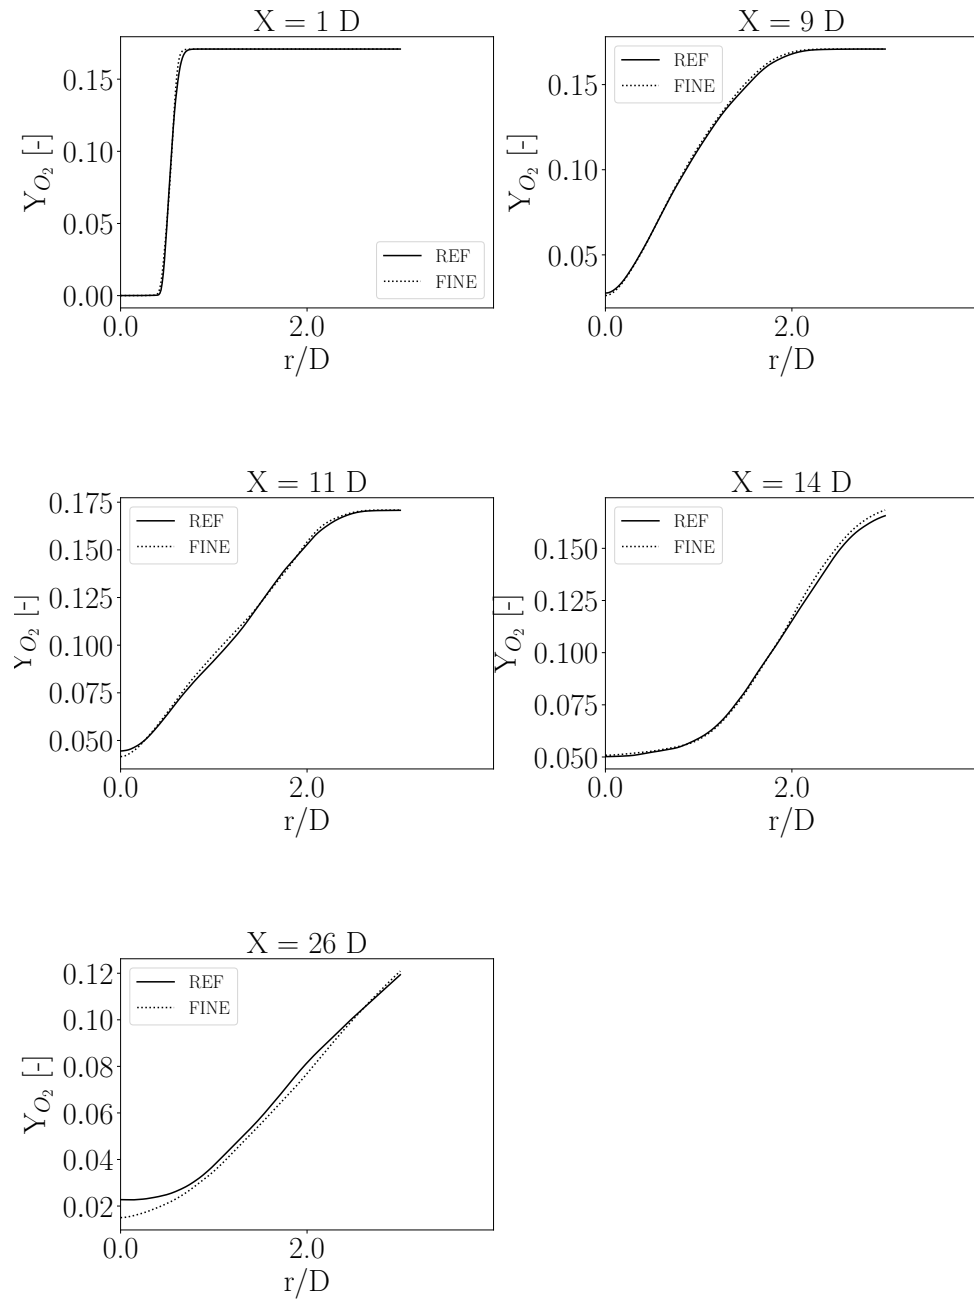


Figure D.9: Comparison of the oxygen mass fraction radial profiles for the original and finer mesh

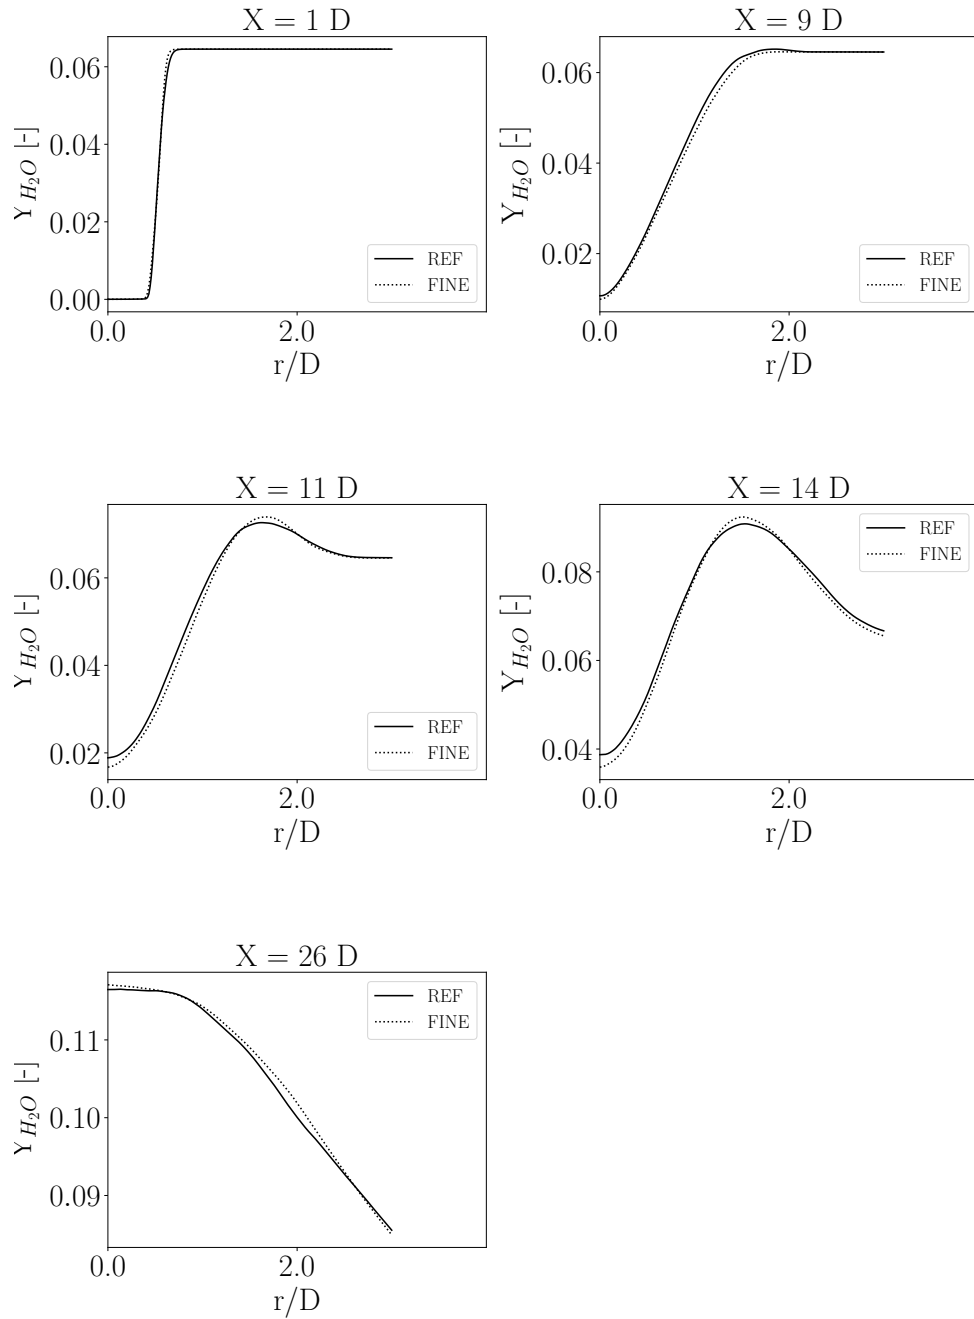


Figure D.10: Comparison of the water vapor mass fraction radial profiles for the original and finer mesh

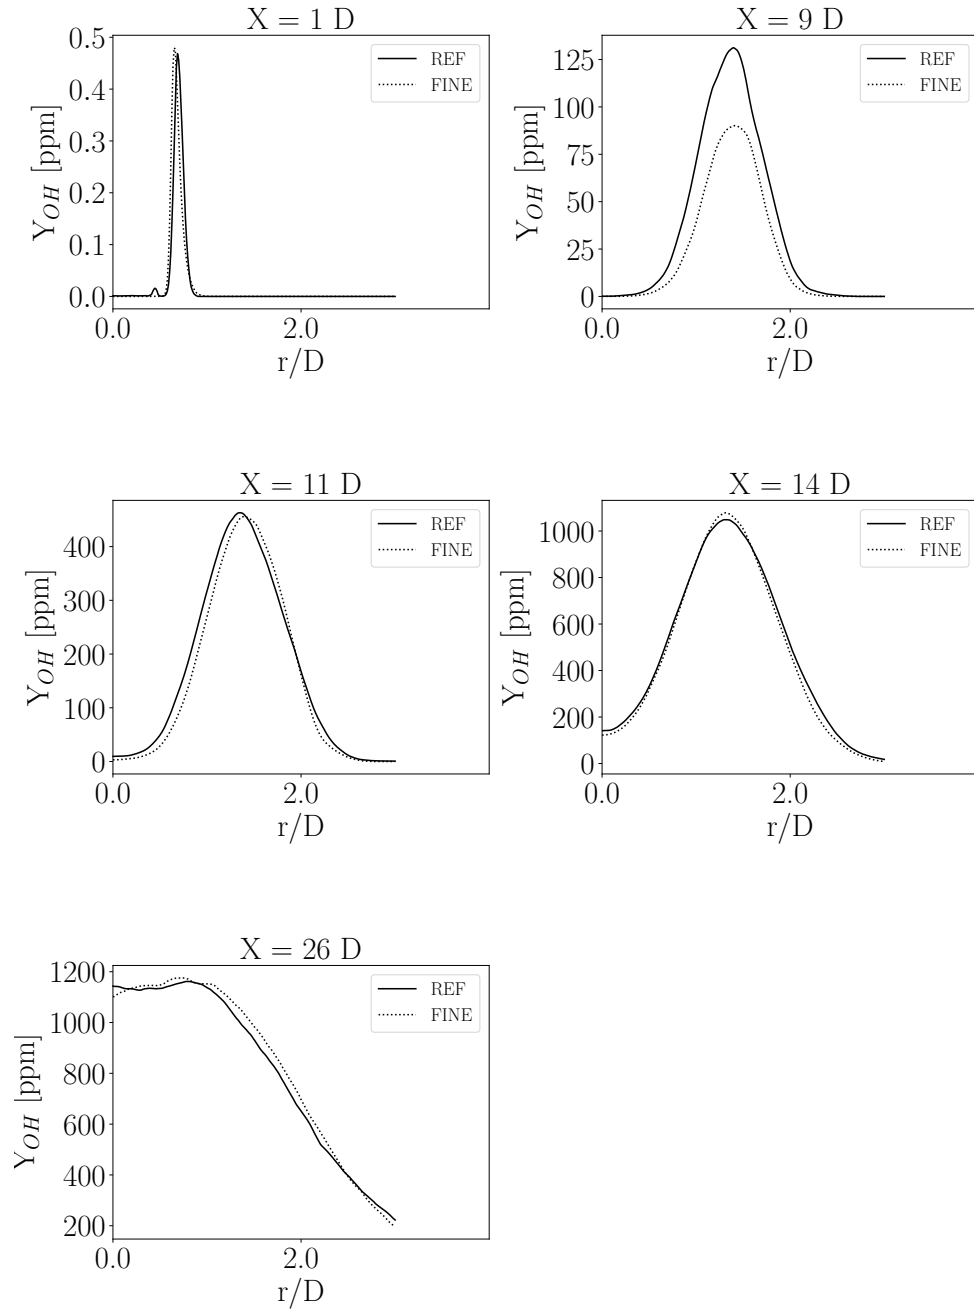


Figure D.11: Comparison of the hydroxyl mass fraction radial profiles for the original and finer mesh




Figure D.12: Comparison of the HO₂ mass fraction radial profiles for the original and finer mesh

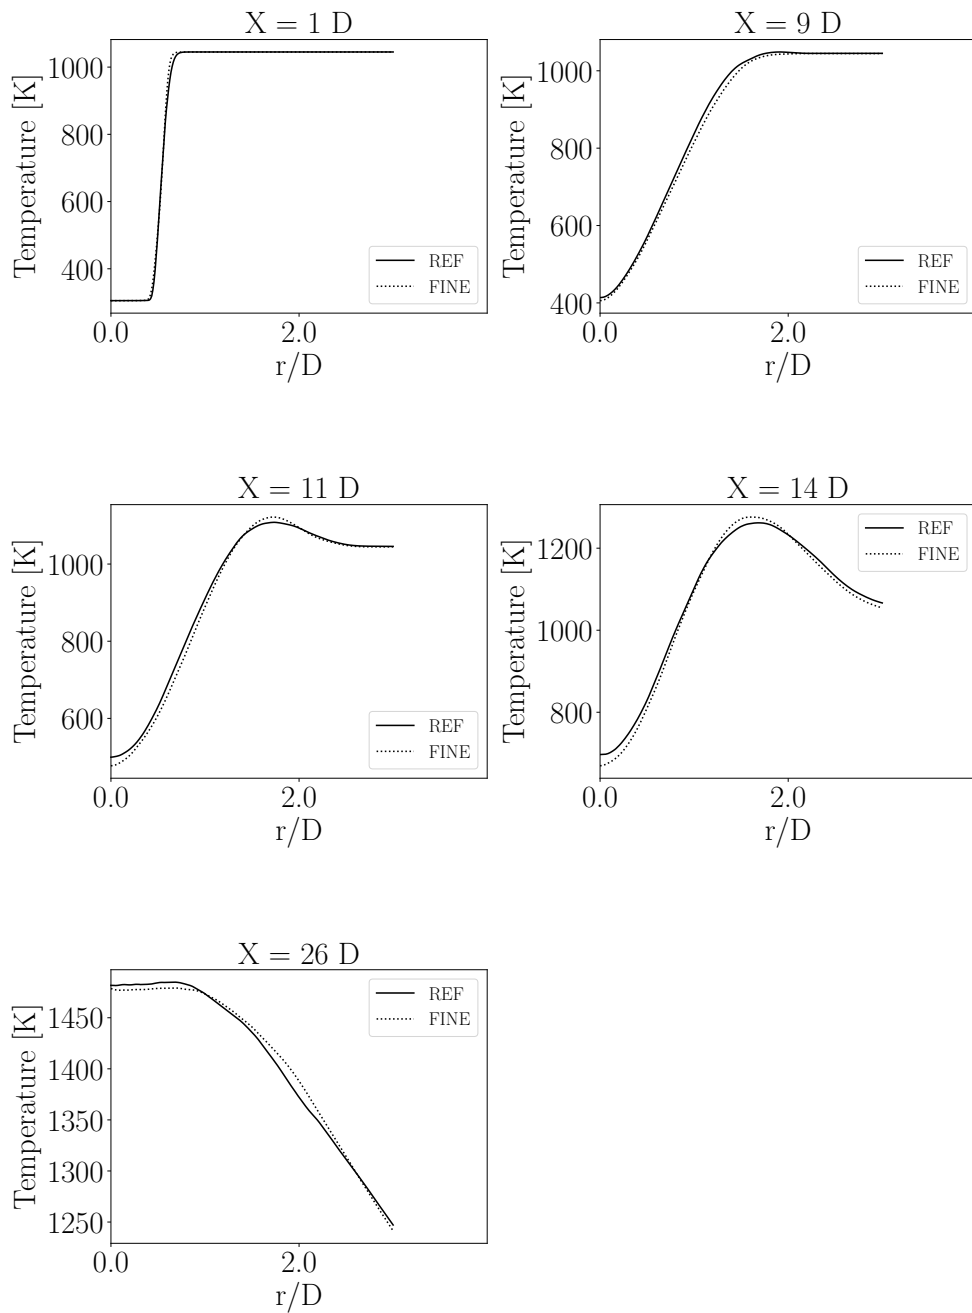


Figure D.13: Comparison of the temperature radial profiles for the original and finer mesh




Figure D.14: Comparison of the mixture fraction radial profiles for the original and finer mesh

Annex E

Surrogate evaluations for the prediction of the Cabra flame's lift-off height

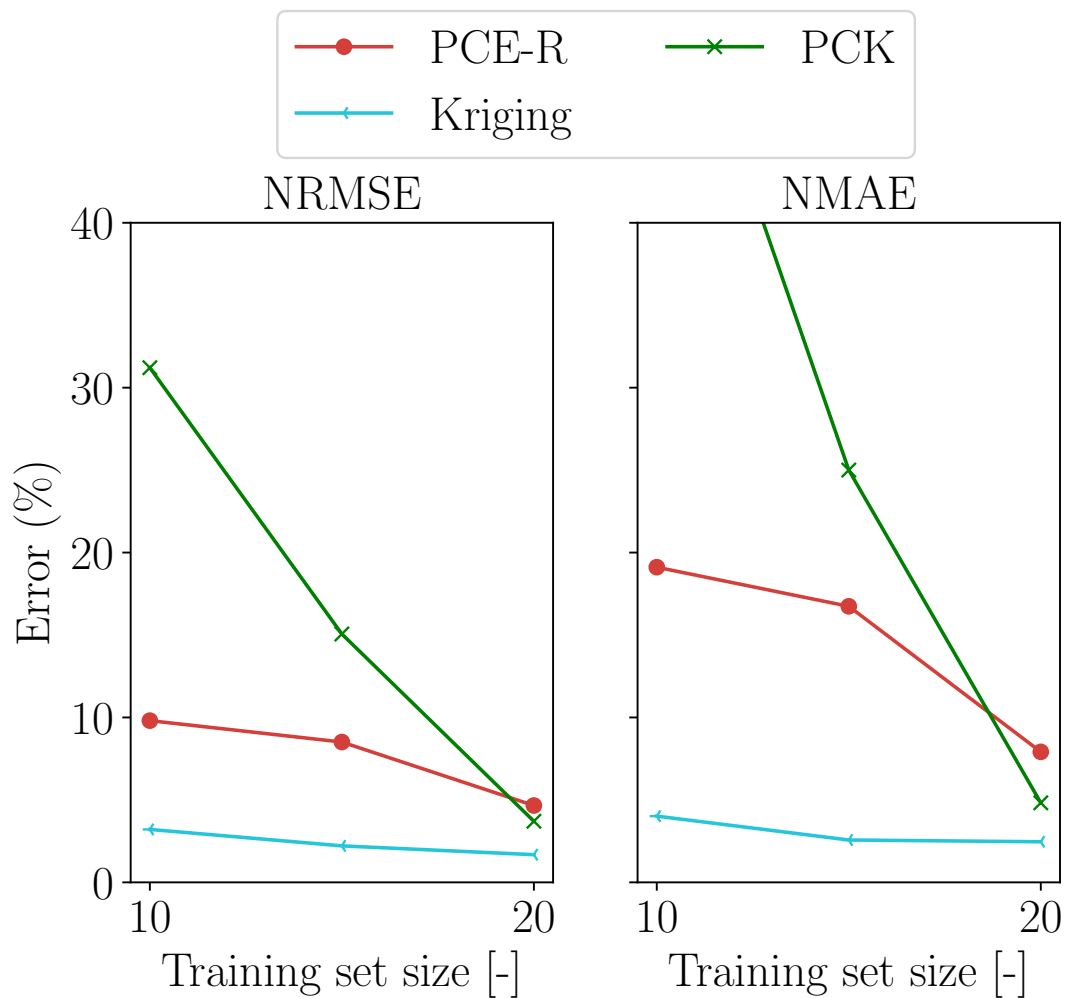


Figure E.1: Evolution of the error of the tested surrogates for the lift-off height of the Cabra flame in $\text{Span}(\xi_T, \xi_{A_{12}})$




Figure E.2: Evolution of the error of the tested surrogates for the lift-off height of the Cabra flame in $\text{Span}(\xi_T, \xi_{A_{12}}, \xi_{C_D})$

Annex F

Averaged fields of the uncertain Cabra flame's training observations

This annex shows the fields of every field taken into account in the field prediction study of Chapter 5. These are the same training samples used in the lift-off height prediction study of Chapter 4, except that it uses only the Y_{OH} field. The part that treats the 2D uncertain space $\text{Span}(\xi_T, \xi_{A_{12}})$ uses only the samples S0 to S19, and the part that treats the 3D uncertain space $\text{Span}(\xi_{A_{12}}, \xi_T, \xi_{C_D})$ uses all 39 samples.




Figure F.1: Averaged Y_{H_2} fields of the training samples

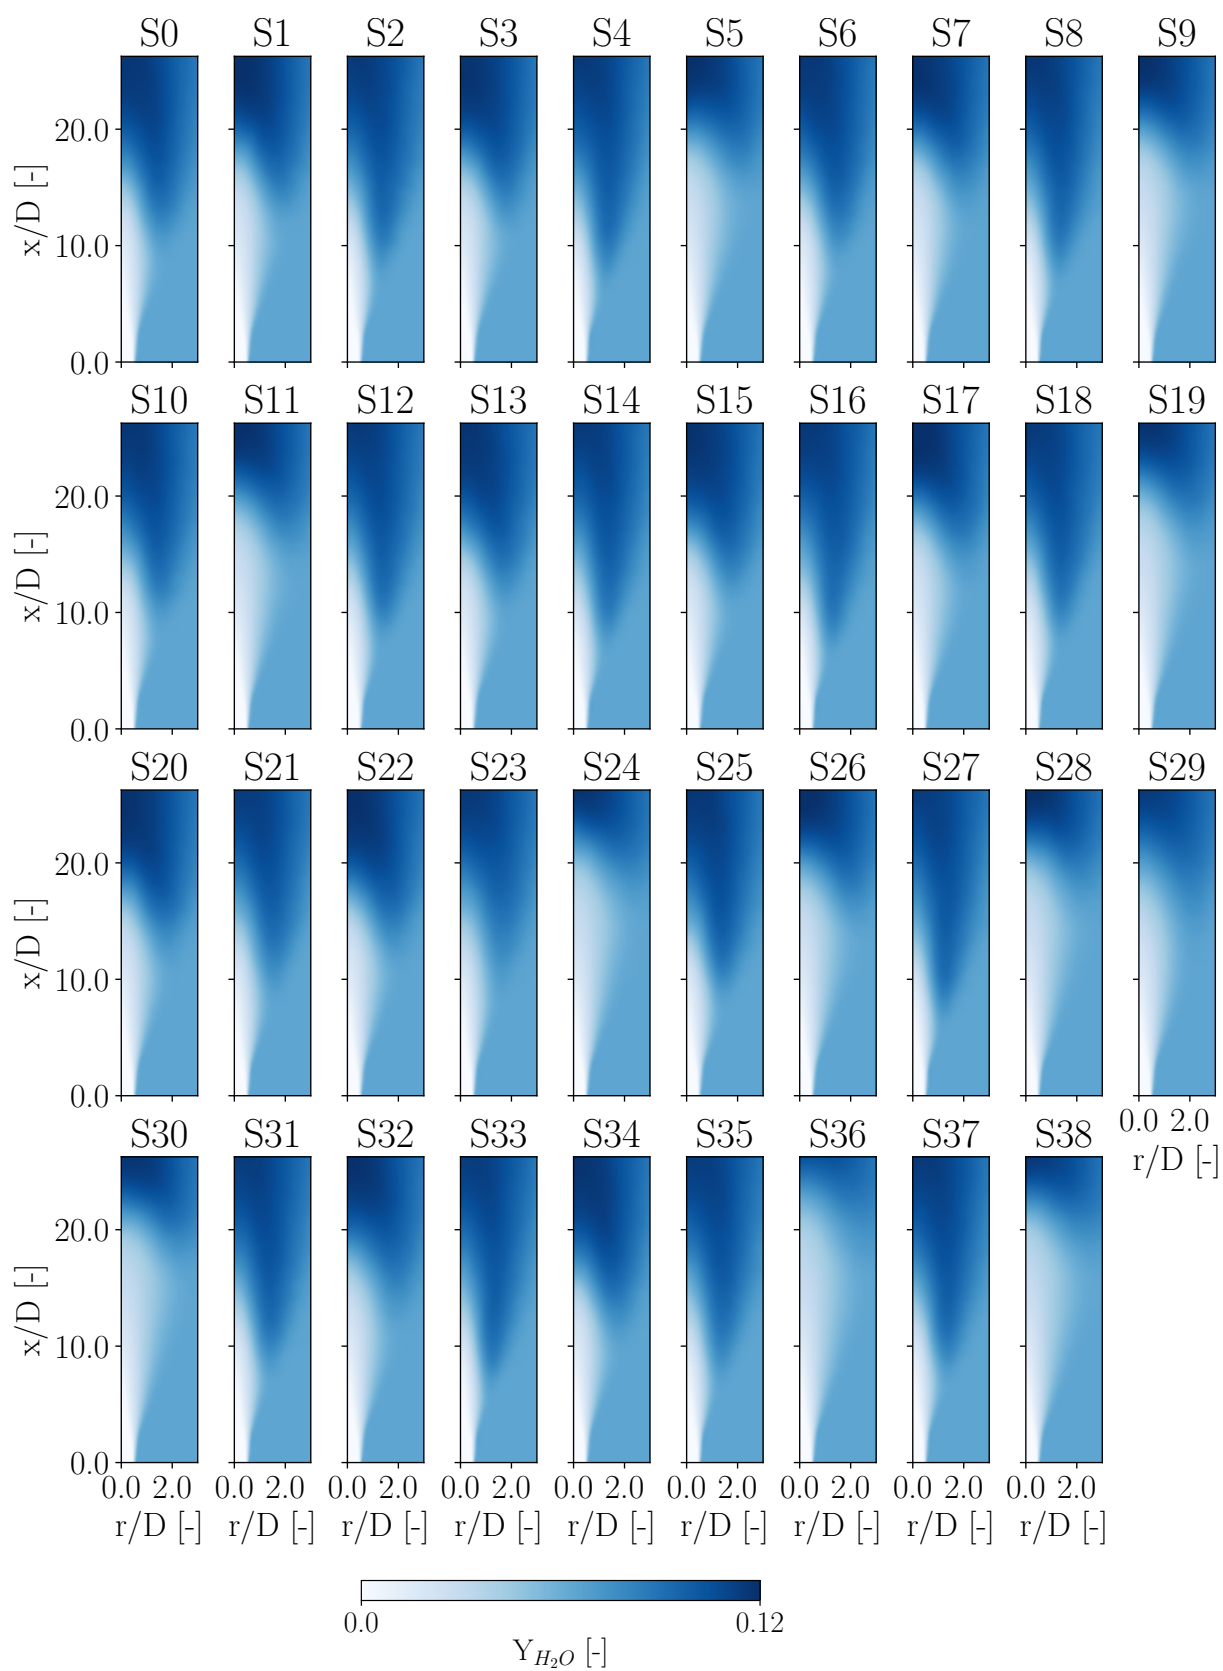


Figure F.2: Averaged Y_{H_2O} fields of the training samples

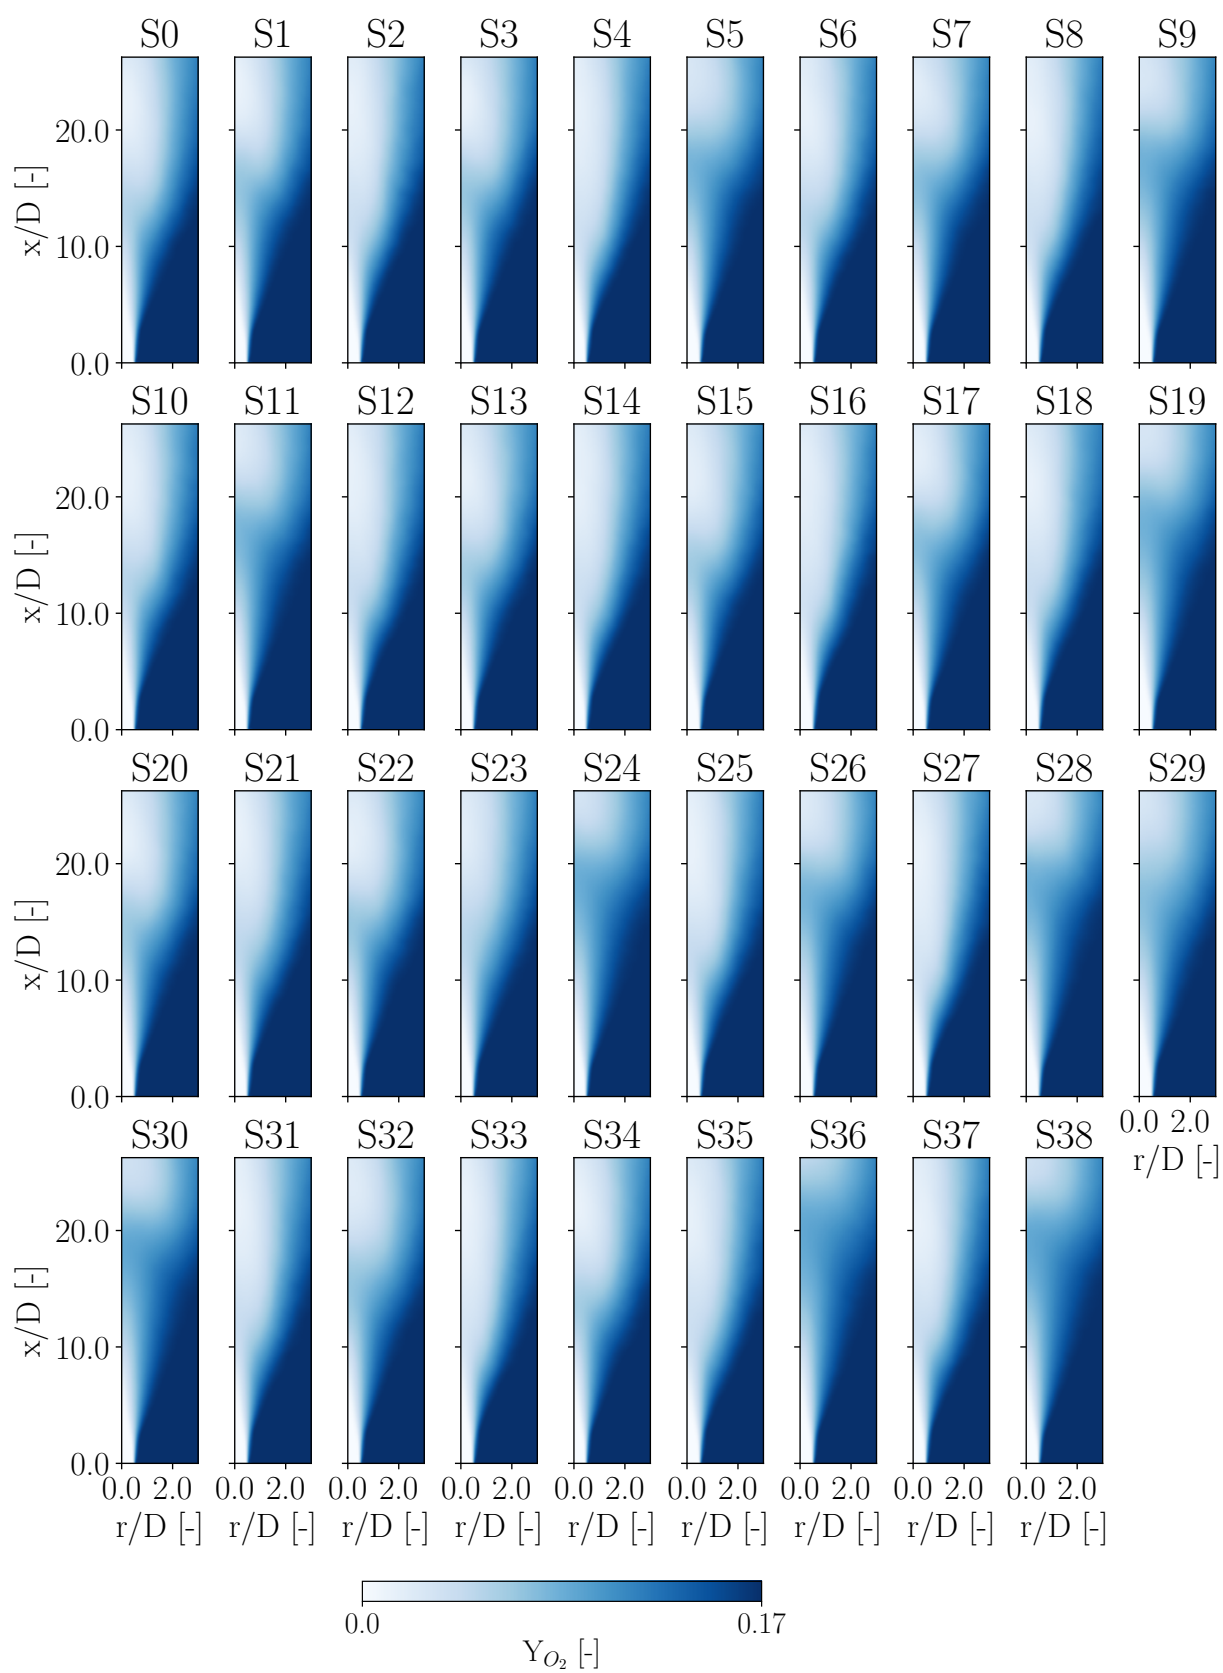


Figure F.3: Averaged Y_{O_2} fields of the training samples




Figure F.4: Averaged Y_{OH} fields of the training samples




Figure F.5: Averaged temperature fields of the training samples




Figure F.6: Averaged Z fields of the training samples

Annex G

Modes retained for the prediction of the Cabra flame's averaged fields




Figure G.1: Mean and retained modes for Y_{H_2}




Figure G.2: Mean and retained modes for Y_{H_2O}




Figure G.3: Mean and retained modes for Y_{O_2}




Figure G.4: Mean and retained modes for Y_{OH}




Figure G.5: Mean and retained modes for the gas temperature




Figure G.6: Mean and retained modes for Z

Annex H

Prediction of the Cabra flame validation samples' averaged fields




Figure H.1: Validation samples **predicted** using the PCA modes and their coefficients predicted by their surrogates (**right**) compared with the LES (left): Y_{H_2} field




Figure H.2: Validation samples **predicted** using the PCA modes and their coefficients predicted by their surrogates (**right**) compared with the LES (left): Y_{H_2O} field




Figure H.3: Validation samples **predicted** using the PCA modes and their coefficients predicted by their surrogates (**right**) compared with the LES (left): Y_{O_2} field

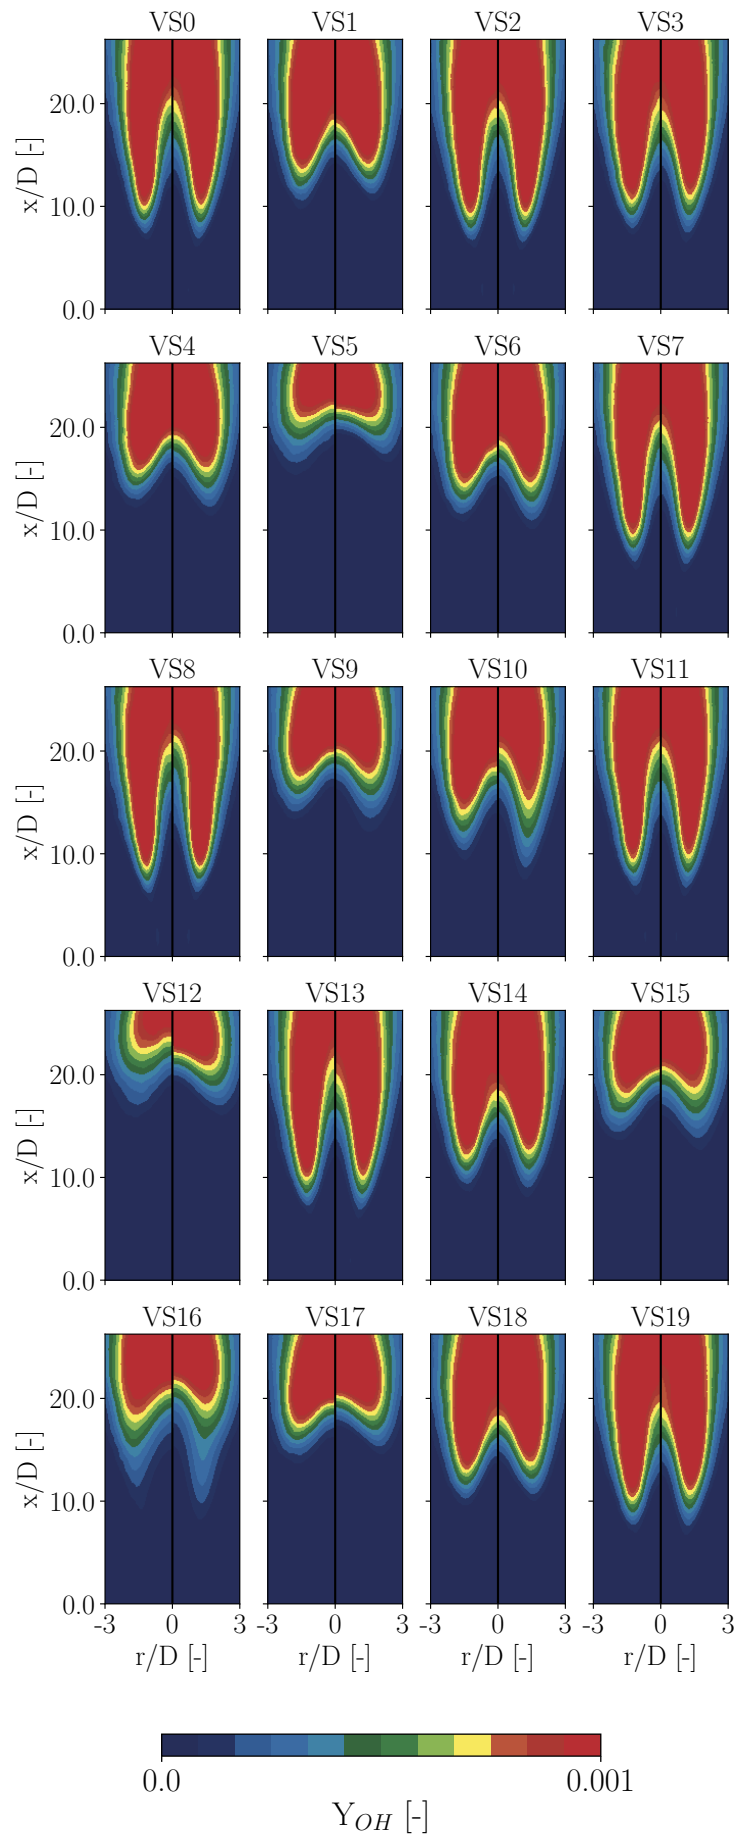


Figure H.4: Validation samples **predicted** using the PCA modes and their coefficients predicted by their surrogates (**right**) compared with the LES (left): Y_{OH} field




Figure H.5: Validation samples **predicted** using the PCA modes and their coefficients predicted by their surrogates (**right**) compared with the LES (left): Gas temperature field




Figure H.6: Validation samples **predicted** using the PCA modes and their coefficients predicted by their surrogates (**right**) compared with the LES (left): Z field

Annex I

Prediction of the Cabra flame validation samples' averaged profiles




Figure I.1: Nominal and 95% confidence interval of the H_2 mass fraction's axial profile in the LES, compared to the measurements from [11]




Figure I.2: Nominal and 95% confidence interval of the H_2O mass fraction's axial profile in the LES, compared to the measurements from [11]

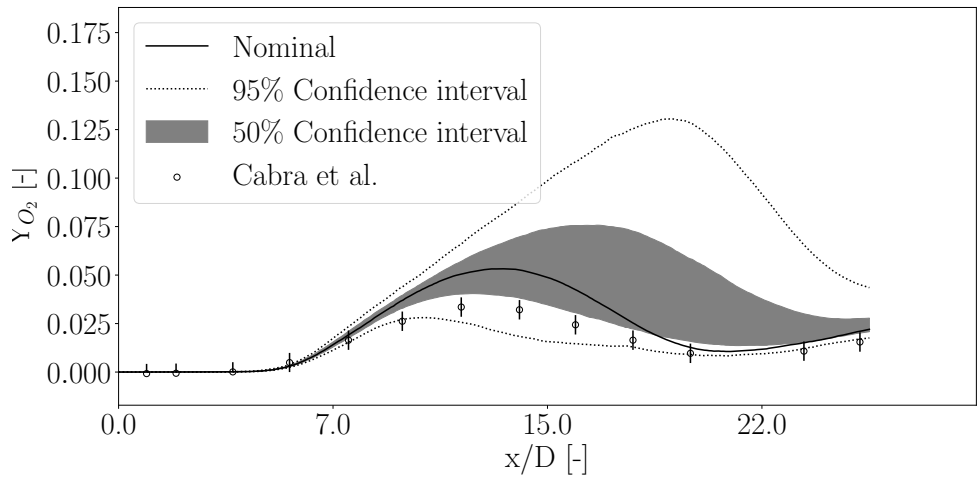


Figure I.3: Nominal and 95% confidence interval of the O₂ mass fraction's axial profile in the LES, compared to the measurements from [11]




Figure I.4: Nominal and 95% confidence interval of the OH mass fraction's axial profile in the LES, compared to the measurements from [11]




Figure I.5: Nominal and 95% confidence interval of the gas temperature's axial profile in the LES, compared to the measurements from [11]




Figure I.6: Nominal and 95% confidence interval of the mixture fraction's axial profile in the LES, compared to the measurements from [11]




Figure I.7: Nominal and 95% confidence interval of the H₂ mass fraction's radial profiles in the LES, compared to the measurements from [11]




Figure I.8: Nominal and 95% confidence interval of the O_2 mass fraction's radial profiles in the LES, compared to the measurements from [11]




Figure I.9: Nominal and 95% confidence interval of the H_2O mass fraction's radial profiles in the LES, compared to the measurements from [11]




Figure I.10: Nominal and 95% confidence interval of the OH mass fraction's radial profiles in the LES, compared to the measurements from [11]

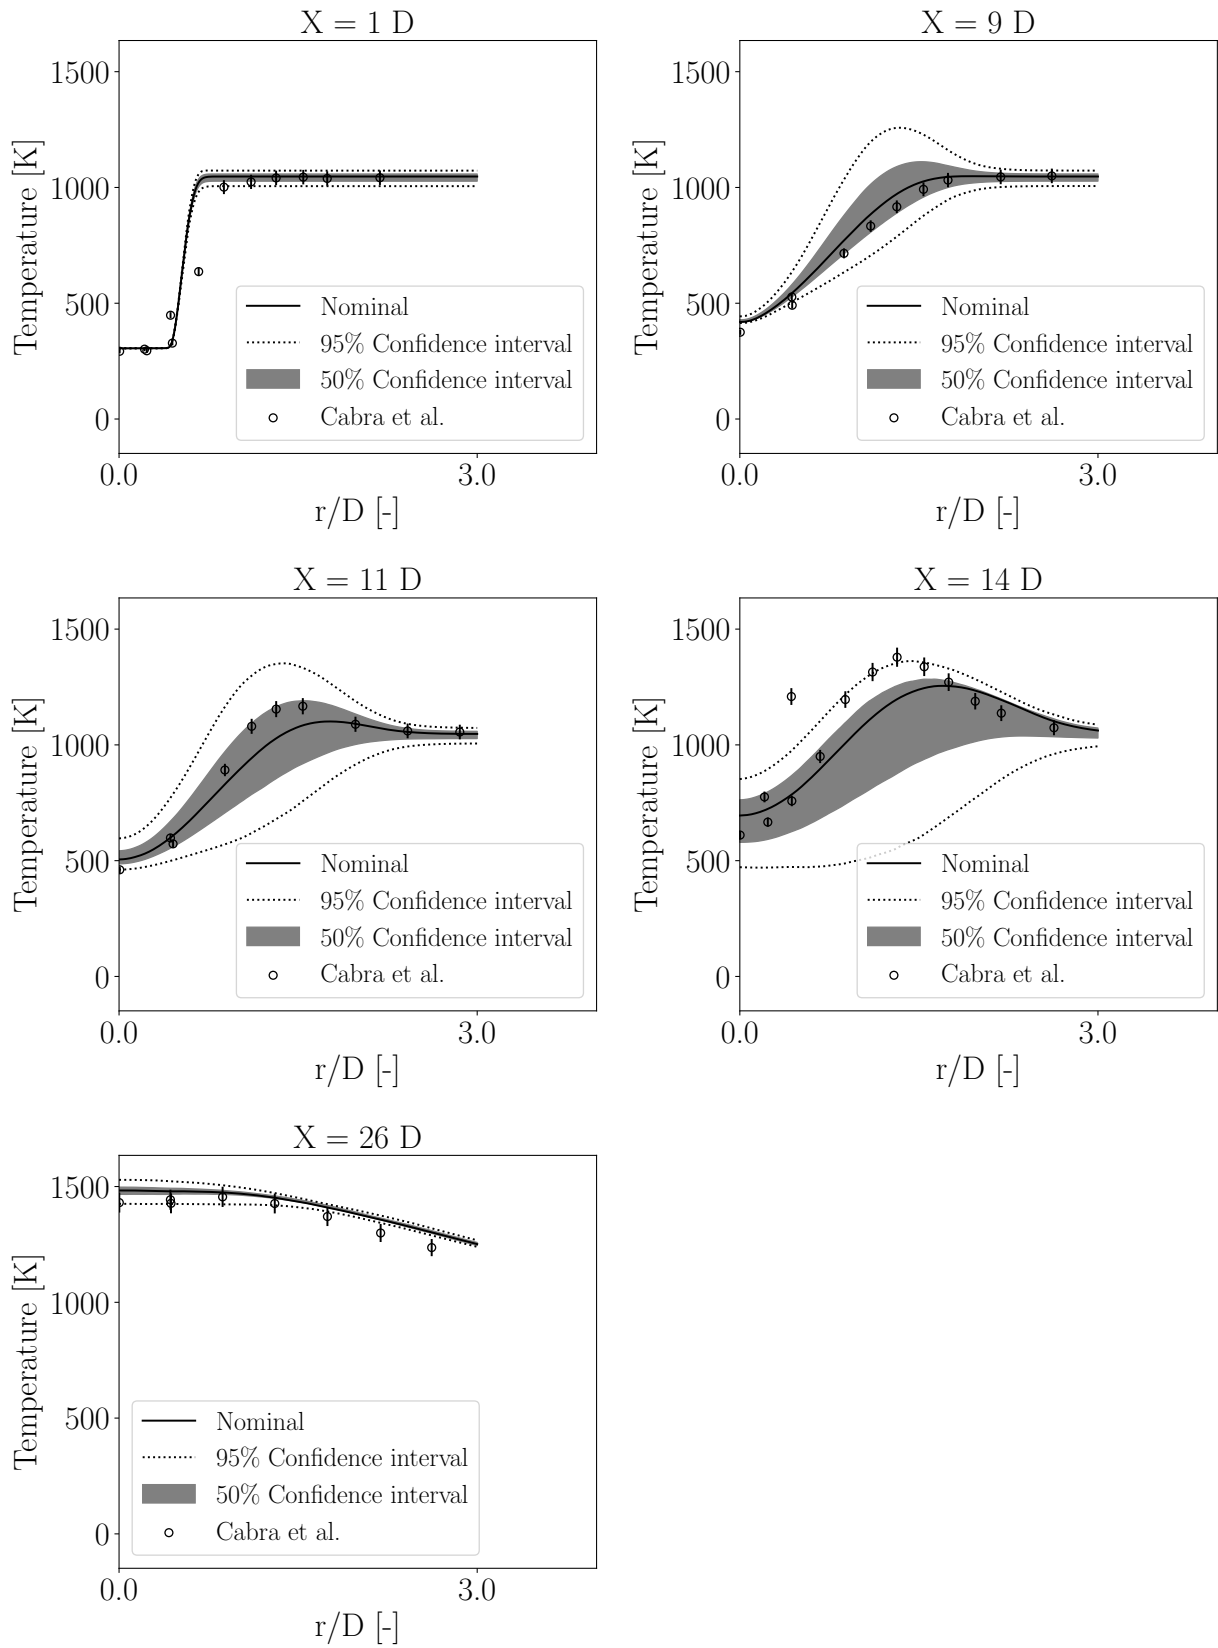


Figure I.11: Nominal and 95% confidence interval of the gas temperature's radial profiles in the LES, compared to the measurements from [11]

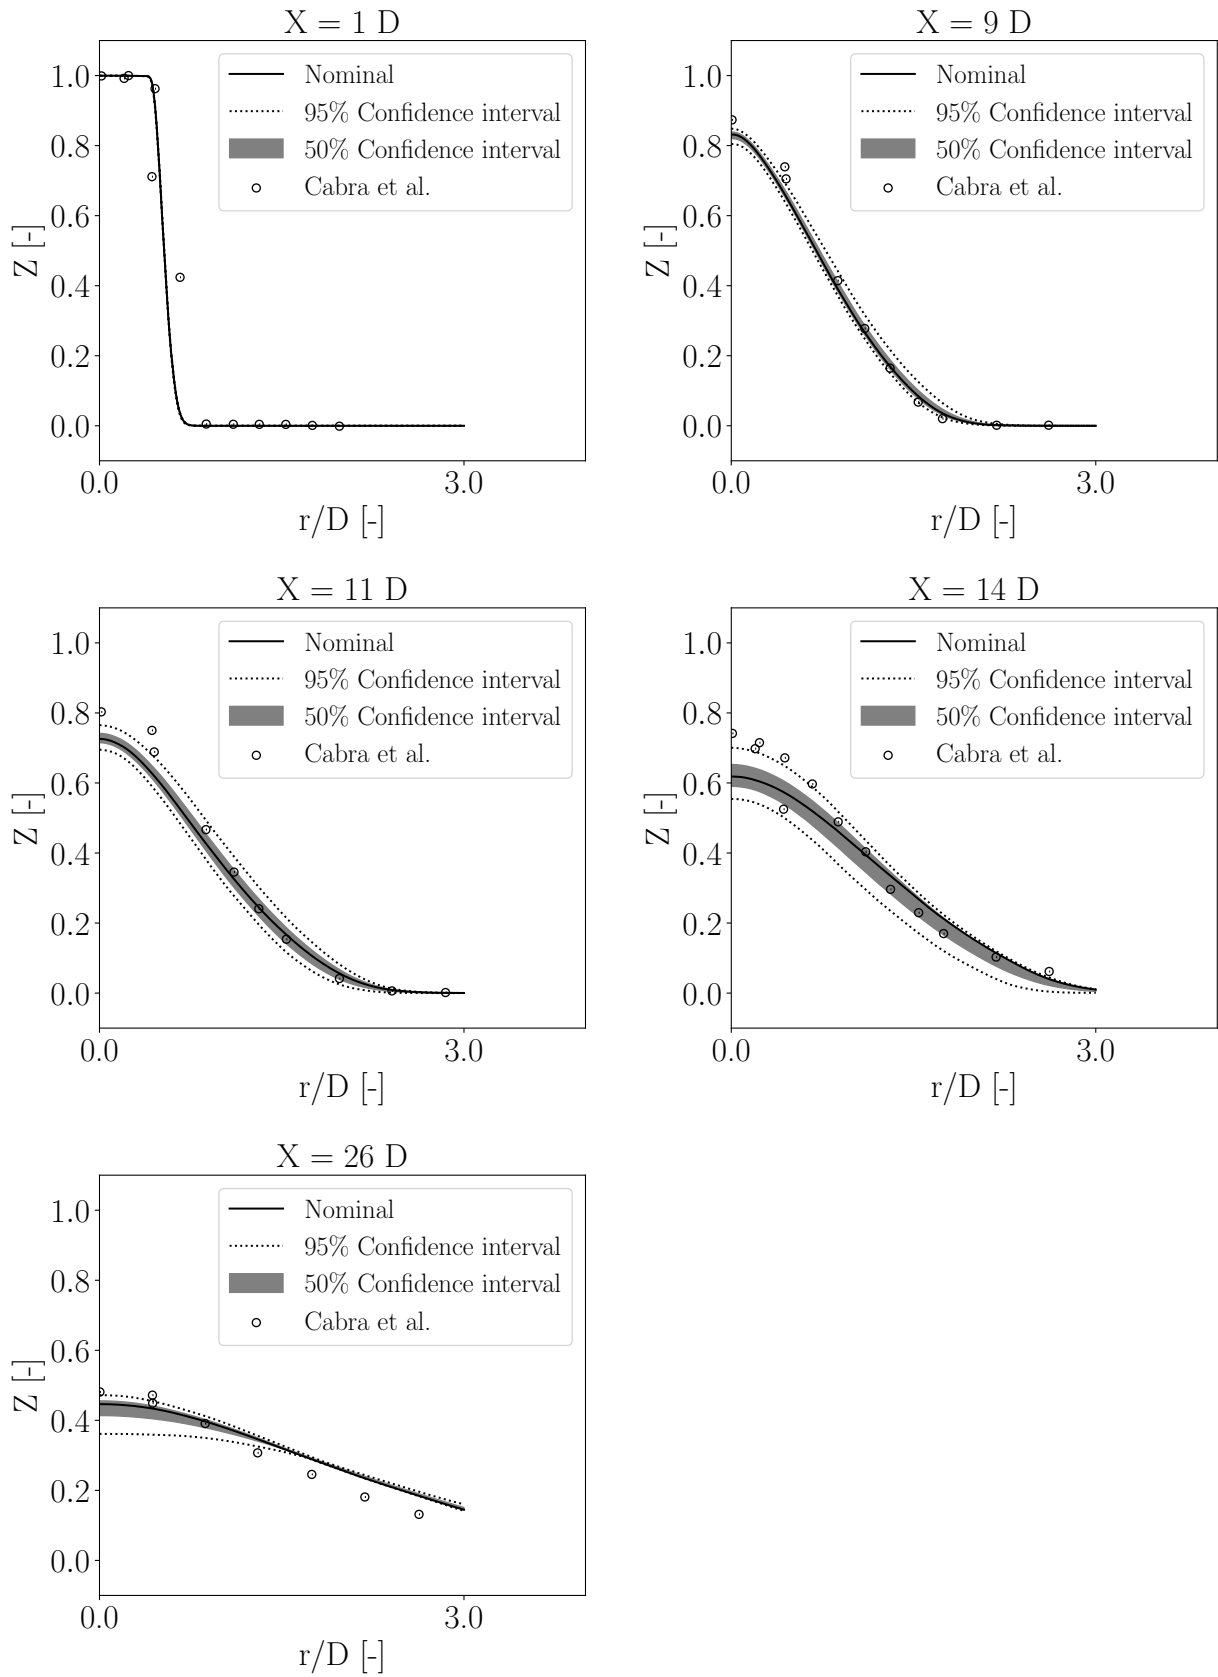


Figure I.12: Nominal and 95% confidence interval of the mixture fraction's radial profiles in the LES, compared to the measurements from [11]