



Olfactory navigation : how to make decisions using a sparse signal

Nicola Rigolli

► To cite this version:

Nicola Rigolli. Olfactory navigation : how to make decisions using a sparse signal. Biological Physics [physics.bio-ph]. Université Côte d'Azur; Università degli studi (Gênes, Italie), 2022. English. NNT : 2022COAZ4023 . tel-03789670

HAL Id: tel-03789670

<https://theses.hal.science/tel-03789670>

Submitted on 27 Sep 2022

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.



$$\rho \left(\frac{\partial v}{\partial t} + v \cdot \nabla v \right) = -\nabla p + \nabla \cdot T + f$$

$$e^{i\pi} + 1 = 0$$

THÈSE DE DOCTORAT

Navigation olfactive: comment prendre
des décisions avec un signal brisé

Nicola RIGOLLI

Université Côte d'Azur – Institute de physique de Nice (INPHYNI)

Università degli Studi di Genova

**Présentée en vue de l'obtention
du grade de docteur en Physique
de l'Université Côte d'Azur,**

Dirigée par : Agnese Seminara, Nicodemo
Magnoli

Soutenue le : 19/05/2022

Devant le jury, composé de :

Antonio Celani, International Centre for Theoretical
Physics, Trieste.

Claire Wyart, Institut du Cerveau et de la Moelle
épinière (ICM), Paris.

Greg Stephens, VU University, Amsterdam.

Mederic Argentina, Université Côte d'Azur.





UNIVERSITÉ UNIVERSITÀ
FRANCO ITALO
ITALIENNE FRANCESE



Doctoral Dissertation
Doctoral Program in Physics and Nanoscience (XXXIV cycle)

Olfactory navigation: how to make decisions using a sparse signal

Nicola Rigolli

Supervisor(s):

Prof. Agnese Seminara, Université Côte d'Azur

Prof. Nicodemo Magnoli, Università di Genova

Doctoral Examination Committee:

Prof. Antonio Celani , Referee, International Centre for Theoretical Physics, Trieste

Prof. Greg J Stephens, Referee, VU University, Amsterdam

Prof. Claire Wyart, Institut du Cerveau et de la Moelle épinière (ICM), Paris

Prof. Mederic Argentina, Université Côte d'Azur, Nice

Università degli Studi di Genova - Université Côte d'Azur

2022

Abstract

All living organisms are surrounded by fluids, either air or water, which create unique sensory landscapes. For example, chemical signals disperse in the flow by diffusion and advection and, when the flow is turbulent, odor breaks up in filaments and discrete patches of varying intensity. In my thesis I focused on olfactory navigation in turbulent environments and I aimed at understanding how organisms overcome uncertainties to make decisions. I developed three-dimensional direct numerical simulations of a turbulent channel flow to recreate a realistic environment for olfactory searches. I realized these state of the art simulations by customizing an open software called Nek5000, which solves the Navier-Stokes equations for the velocity field and the advection-diffusion equation, which regulates the evolution of the odor (passive scalar) in a fluid. After generating large fluid dynamics datasets of odorant evolution in a channel, I analyzed which features of the olfactory signal are more relevant to locate the odor source. Surprisingly, not only the signal, but also its absence can be informative to infer the distance from the odor source. Using supervised learning algorithms I showed that the intensity of odor concentration is an informative measure, but that when it is coupled to the temporal dynamics of the signal, it allows robust predictions in different conditions and at different ranges from the source. These theoretical results suggest that it is computationally advantageous to measure both odor intensity and timing. I analyzed a set of neural recording from awake mice, demonstrating that they are indeed able to store both quantities, and that the neural representation depends on the underlying flow. I then considered the problem of navigating to the source of the turbulent odor. Although animals (for example moths and crustaceans) robustly perform this task, the algorithms they use are not understood. I modeled olfactory navigation using the framework of Partially Observable Markov Decision Processes (POMDP) and I proposed a normative theory to explain the alternation between sniffing in the air and sniffing the ground, typical of mammals like rodents and dogs. Alternation stems from the physics of fluids, pre-

scribing that odor near the ground is more continuous than up in the air, but remains relatively close to the source. In contrast, at nose level the odor is transported quickly away from the source, but is more noisy and intermittent. An agent searching for the odor source should thus sniff in the air when it is far from the source to increase its chances of detecting the odor. Once the agent localizes the odor plume, it should continue the search sniffing the ground where the trail is less intermittent. The exact timing for alternation stems from marginal value theory. Finally, the commonly observed behavior of searchers proceeding in casts and surges emerges from this computational framework, and alternation naturally complements this dynamics to ensure optimal exploration.

Key words: turbulence, navigation, olfaction.

Tutti gli esseri viventi sono immersi in aria o in acqua, questi fluidi creano scenari sensoriali unici. Ad esempio, i segnali chimici si disperdono in un fluido per diffusione e avvezione e, quando il flusso è turbolento, la concentrazione degli odori si rompe in filamenti e in strutture isolate di varia intensità. Nel mio lavoro di tesi mi sono concentrato sulla navigazione olfattiva in ambienti turbolenti e ho cercato di capire come gli organismi superano l'incertezza per prendere decisioni. Ho sviluppato delle simulazioni numeriche (DNS) tridimensionali del flusso in un canale turbolento, per ricreare un ambiente realistico per la ricerca olfattiva. Ho realizzato queste simulazioni allo stato dell'arte adattando un software open source chiamato Nek5000, che risolve le equazioni di Navier-Stokes per il campo di velocità e l'equazione di avvezione-diffusione, che regola l'evoluzione dell'odore (scalare passivo) in un fluido. Dopo aver generato grandi dataset di evoluzione dell'odore in un canale, ho analizzato quali caratteristiche del segnale olfattivo sono più rilevanti per individuare la sorgente dell'odore. Sorprendentemente, non solo il segnale, ma anche la sua assenza può essere informativa per inferire la distanza dalla fonte di odore. Utilizzando algoritmi di Reinforcement Learning ho dimostrato che l'intensità della concentrazione di odore è una misura informativa, ma che accoppiata alla dinamica temporale del segnale consente previsioni robuste in condizioni diverse e a varie distanze dalla sorgente. Questi risultati teorici suggeriscono che è computazionalmente vantaggioso misurare sia l'intensità che la dinamica dell'odore. Ho analizzato una serie di registrazioni neurali di topi svegli, dimostrando che sono effettivamente in grado di memorizzare entrambe le quantità e che la rappresentazione neurale dipende dal flusso di odore. Ho quindi considerato il problema della navigazione verso la sorgente di odore turbolento. Sebbene gli animali (ad esempio falene e crostacei) svolgano in modo efficace questo compito, gli algoritmi che utilizzano non sono compresi. Ho modellato la navigazione olfattiva utilizzando il framework dei Partially Observable Markov Decision Processes (POMDP) e ho proposto una teoria normativa per spiegare l'alternanza tra annusare in aria e annusare il suolo, tipica di mammiferi come roditori e cani. L'alternanza deriva dalla fluidodinamica, infatti l'odore vicino al suolo è più continuo che in aria, ma rimane relativamente vicino alla sorgente. Al contrario, a livello del naso l'odore viene trasportato rapidamente lontano dalla fonte, ma è più rumoroso e intermittente. Un agente che cerca la fonte dell'odore dovrebbe quindi annusare l'aria quando è lontano dalla sorgente per aumentare le sue possibilità di rilevare l'odore. Una volta che l'agente ha localizzato l'odore, dovrebbe continuare la ricerca annusando il terreno dove la traccia è meno

intermittente. Ho dimostrato che i tempi di tale alternanza sono determinati da una condizione di marginalità fra esplorazione locale e spostamento verso un'altra regione di spazio. Infine, in questo framework computazionale, emerge naturalmente il comportamento comunemente osservato negli organismi che cercano una sorgente di odore, in cui l'agente procede in movimenti laterali intervallati da spostamenti in direzione opposta al flusso.

Parole chiave: turbolenza, navigazione, olfatto.

Tous les organismes vivants sont entourés de fluides, air ou eau, qui créent des paysages sensoriels uniques. Par exemple, les signaux chimiques se dispersent dans le fluide par diffusion et advection et lorsque l'écoulement est turbulent, la concentration d'odeurs se décompose en filaments et en taches discrètes d'intensité variable. Dans ma thèse, je me suis concentré sur la navigation olfactive dans des environnements turbulents et j'ai cherché de comprendre comment les organismes surmontent les incertitudes pour prendre des décisions. J'ai développé des simulations numériques (DNS) tridimensionnelles d'un écoulement turbulent dans un canal afin de recréer un environnement réaliste pour la recherche olfactive. J'ai réalisé ces simulations en personnalisant un open-source software appelé Nek5000, qui résout les équations de Navier-Stokes pour le champ de vitesse et l'équation d'advection-diffusion, qui régule l'évolution de l'odeur (scalaire passif) dans un fluide. Après avoir généré de grands ensembles de données de dynamique des fluides sur l'évolution des odeurs dans un canal, j'ai analysé quelles caractéristiques du signal olfactif sont les plus pertinentes pour localiser la source de l'odeur. Étonnamment, non seulement le signal, mais aussi son absence peut être informative pour déduire la distance de la source de l'odeur. En utilisant des algorithmes de Reinforcement Learning, j'ai montré que l'intensité de la concentration d'odeur est une mesure informative, mais que la dynamique temporelle du signal permet des prédictions robustes dans différentes conditions et à différentes distances de la source. Ces résultats théoriques suggèrent qu'il est avantageux sur le plan informatique de mesurer à la fois l'intensité et la dynamique de l'odeur. J'ai analysé un ensemble d'enregistrements neuronaux de souris éveillées, démontrant qu'elles sont bien capables de stocker les deux quantités, et que la représentation neuronale dépend du flux sous-jacent. J'ai alors considéré le problème de la navigation jusqu'à la source de l'odeur turbulente. Bien que les animaux (par exemple les papillons de nuit et les crustacés) effectuent cette tâche de manière robuste, les algorithmes qu'ils utilisent ne sont pas compris. J'ai modélisé la navigation olfactive en utilisant le cadre des Partially Observable Markov Decision Processes (POMDP) et j'ai proposé une théorie normative pour expliquer l'alternance entre renifler dans l'air et renifler le sol, typique des mammifères comme les rongeurs et les chiens. L'alternance découle de la physique des fluides, prescrivant que l'odeur près du sol est plus continue que dans l'air, mais reste relativement proche de la source. En revanche, au niveau du nez, l'odeur est transportée rapidement loin de la source, mais est plus bruyante et intermittente. Un agent recherchant la source de l'odeur doit donc renifler l'air lorsqu'il

est loin de la source pour augmenter ses chances de détecter l'odeur. Une fois que l'agent a localisé le panache d'odeur, il doit continuer la recherche en reniflant le sol où la piste est moins intermittente. Le moment exact de l'alternance découle de la Marginal Value Theory. Enfin, le comportement communément observé chez les organismes recherchant une source d'odeur, procédant en cast et surge émerge spontanément de ce cadre computationnel.

Mots clés: turbulence, navigation, olfaction.

Contents

1	Introduction	1
1.1	Bacteria chemo-taxis	3
1.2	Animal chemo-sensing	5
2	Channel turbulence with an odor source	10
2.1	A brief introduction to fluid dynamics and turbulence	10
2.2	Odor transport	12
2.3	Lagrangian representation	14
2.4	Theoretical description of odor molecules transported by turbulence	15
2.5	Channel turbulence	17
2.6	Direct numerical simulations of odor in a turbulent channel flow . .	19
2.6.1	Boundary conditions and odor source	21
2.6.2	Results	22
3	Olfactory predictions	25
3.1	Odor encoding: intensity vs timing	25
3.2	An introduction to supervised learning	26
3.3	Kernel ridge regression to infer odor source location	30
3.4	Intensity and timing features allow a robust prediction of the source position	33

3.5	Performance in space	40
3.5.1	Dependence on the sampling height	40
3.5.2	Dependence on the source distance	41
3.6	Modeling the predictive power of individual features	45
3.7	Conclusions	47
4	Neural recordings in the olfactory bulb of mice	51
4.1	A brief introduction to the olfactory system in mice	51
4.2	Monitoring neural response with calcium imaging in different flow conditions	54
4.2.1	Odor, sensor and flow characterization	55
4.2.2	Calcium Imaging	59
4.3	Results	60
4.3.1	Qualitative evidence that brain responds to olfactory stimuli	60
4.3.2	Correlation among odorant plume dynamics and glomerular activity	61
4.3.3	Responsivity and response power analysis	63
4.4	Conclusions	65
5	Navigation in turbulent odor plumes	68
5.1	Algorithms for olfactory navigation	68
5.1.1	Gradient climbing strategy	69
5.1.2	Bio-inspired algorithms for olfactory navigation	70
5.1.3	Infotaxis	73
5.1.4	Reinforcement Learning	74
5.2	Partially Observable Markov Decision Processes for olfactory searches	75
5.2.1	An introduction to POMDPs	75
5.2.2	Algorithms to solve POMDPs	77

5.2.3	POMDPs allow navigation in a turbulent environment . . .	79
5.2.4	A theoretical model	82
5.2.5	Parameters used in the POMDP	88
5.3	Alternation between detection in air and at the ground	89
5.3.1	Model	91
5.3.2	Results	93
5.3.3	Conclusions	96
Publication List		100
References		101

Chapter 1

Introduction

We are all familiar with the experience of having a fragmented knowledge of reality, our senses only allow a partial representation: sometimes we can adjust our limits (i.e. wearing glasses) and also overcome them (i.e. using a microscope). This limit is shared with all the other living organisms, quoting a thread on Twitter by the science writer Ed Yong [1]: "each [species] is trapped in a unique sensory bubble", different species experience the world in different ways. Sight, hearing, tactile perception, taste and smell are what we usually refer to as five senses; in Nature sensing goes far beyond human sensing: other living entities can perceive electrical and magnetic field, air moisture, light polarization, see at infrared or microwave frequencies, use sonar (echolocation), sense temperature and humidity variations, know their own position in space, have proprioception. I will quote again Ed Yong Twitter thread because it sums up the motivations of my work in a few vibrant and fascinating lines: "other organisms have different conceptions of heat, pain, darkness, and silence. Their senses help us to re-imagine our own surroundings. I find it breathtaking to know what birds hear in their own songs, what dogs smell on the streets, what insects feel as they stand on plants. To perceive the world through other senses is to find splendor in familiarity, wilderness in one's backyard, the sacred in the mundane. And being able to even contemplate another creature's sensory world at all is an utterly profound act. There's so much info out there that no animal could perceive it all and no animal needs to. Our senses filter in what we need. We must choose to learn about the rest ". Therefore where there is sensation there are biases, hence also animals do have biases in their world representation. Anthropomorphism is just another kind of bias: it affects our approach to behavioral studies, we often pretend

organisms we study are like us; but understanding the perception of the world of other species is challenging. During my PhD I did my tiny part in "learning about the rest", focusing on the diversity of solutions developed by distinct organisms to solve the concrete problem of finding a target of interest through odor cues it releases.

The classification of living systems in different domains is important because it highlights the complexity of different organisms, it indicates how long ago evolution diverged, and suggests the level of sophistication organisms can use to post-process the sensed information. Not only vertebrates with their developed nervous systems are capable of sensing and perception. Bacteria, as we will see later in this introduction, respond to a wide range of stimuli such as concentration of chemicals (chemotaxis), light (phototaxis), electric fields (galvanotaxis), magnetic fields (magnetotaxis), pH (pH-taxis), temperature (thermotaxis) [2]. Higher organisms face more complex, typically turbulent, environments. Plants are capable of sensing, e.g. gravity [3] and temperature [4], but how they integrate noisy information is not understood. Some organisms appear to be specifically adapted at sensing turbulent signals: i.e. moths do not track the odor dynamics when exposed to a laminar flow [5]. To disperse across space Fungi eject spores in the atmosphere [6] and planktonic larvae leave the reef where they were born to enter the large scale oceanic circulation [7]: how they make decisions and what parameters mostly affect their behavior is still unclear, yet they can clearly survive in a turbulent environment.

Chemical sensing is ubiquitous in organisms across the tree of life. The task of olfactory navigation however is constrained by the physics that dictates fluid flow and odor transport, which varies with the specific niche an organism live in. Hundreds of thousands of years of evolution have shaped the sensory systems of different species, sometimes converging and sometimes diverging. Thus studying the strategies of olfactory navigation contributes to the study of evolution, which is intimately connected to the concept of life itself.

As I mentioned above, living organisms sense various observables that enable a partial knowledge of the surrounding environment. Perception starts from the interaction of sensors and external stimuli. Information can be processed at different levels depending on whether organisms have a nervous system; and the complexity of data elaboration depends on the structure and complexity of the nervous system. After sensing, the organism will make decisions and react to the stimuli with one

or more actions, the ensemble of actions is named behavior. In my thesis I focused on distant chemo-sensing, I sought to understand how organisms locate a chemical source exploiting the odor signal it emits in the environment.

All creatures live surrounded by fluids, either air or water; fluid flow can shape a large variety of possible landscapes depending on the spatial and temporal scales where the organism operates. Chemicals disperse in the flow by diffusion and advection; at large scales, flows are often turbulent and far from the source, odorants break up in filaments and discrete patches of varying intensity. Odor fields at large scales are sparse and intermittent and it is difficult to trace the path of the molecules back to the source. Sparsity means that odor is organized in whiffs (a portion of air where odor stays above detection threshold for a certain time), alternated with blanks (fractions of time without odor). In this chaotic framework, where whiffs and blanks of different duration alternate, organisms have to survive and proliferate.

To conclude this overview I will present two extremely different scenarios for how organism locate and navigate to chemical sources. I will first present how unicellular and brain-less bacteria sense and behave in smooth chemical landscapes. I will then contrast this picture by considering a turbulent scenario and discuss the sensory systems of species living at large scales; as well as behavioral evidences in natural and laboratory settings.

1.1 Bacteria chemo-taxis

Bacteria are typically few micrometers long; at this scale diffusion dominates the transport of chemicals and this enables chemotaxis. As etymology suggests, chemo-taxis is an oriented movement by an agent in response to a chemical stimulus. This term assumed a more specific meaning and it refers to the motion towards well established chemical gradients that are present in diffusion dominated environments, at small scales [8]. Bacteria are the most widely studied model organisms for chemo-taxis, and it is now clear that they can sense and track nutrient gradients to flee from poisons and dangerous substances or to reach nutrients. Their chemo-tactic behavior is well-known and it has been modeled and reproduced using gradient ascent algorithms when the agent moves toward higher food concentration (positive chemotaxis) or gradient descent when they are fleeing from repellent (negative chemotaxis). Among the bacteria, *E.coli* is the most studied for its chemo-tactic




Fig. 1.1 (A) Swimming of *E. coli* in the absence of a gradient. Movement of *E. coli* cells in a uniform environment consists of smooth runs that last up to several seconds and are interrupted by short (0.1 s) tumbles. Runs result from the counterclockwise rotation of flagella, which results in formation of a propelling flagellar bundle behind the cell. Tumbles are caused by the clockwise rotation of one or several flagella, which destabilizes the bundle. Tumbles randomly reorient the cell body before the next run, with the angle of reorientation (indicated by red arrow) being dependent on the number of clockwise-rotating flagella. The resulting random walk ensures effective foraging in the environment, and may be further enhanced by occasional long runs (green) resulting from stochastic fluctuation in the pathway activity. (B) Chemotaxis in gradients. The chemotaxis strategy of *E. coli* and other bacteria is based on a biased random walk, whereby cells make temporal comparisons of chemoeffector concentrations during a run and suppress the onset of the next tumble if the level of positive stimulation increases. As a consequence, runs in the positive direction (i.e., up the chemoattractant gradient) are prolonged. Moreover, since on average fewer flagella participate in tumbles when cells are moving up the gradient, the degree of cell body reorientation during such tumbles is smaller. The magnitude of response to the gradient depends on the change in attractant concentration (Δc) experienced by the swimming cell during a run before the cell's memory is reset by the adaptation system, with the typical run time 1s and the corresponding measurement distance $20 \mu m$. This figure is taken from [10].

behavior. For this model system, the molecular machinery that implements gradient climbing is well known. Every cell has several flagella, from 4 to 8, that can rotate and produce two alternating movements: "runs" and "tumbles" [9], trajectories are made by forward discrete steps and short pauses, followed by random changes in direction (see Figure 1.1). Runs happen when flagella rotate counter-clockwise and they align in a single bundle increasing the velocity, tumbles occur when flagella rotate clockwise and they all point to different directions stopping the run of the bacterium. If the environment is homogeneous the frequency of changes in direction is about 1Hz; in the presence of chemical gradients it varies. Specifically, runs lengthen when bacteria sense an increase in chemicals. As a result, this two states alternation yields a biased random walk leading to the chemical source.

As animals have olfactory neural receptors to detect odorants, so bacteria need sensor to measure chemicals: *E.coli* is very sensitive (up to 3nM [10]), it measures concentration thanks to trans-membrane receptors called methyl-accepting chemotaxis proteins (MCPs), different bacteria encode different receptors. The binding of chemicals to the receptors is favored by a pair of enzymes, CheB and CheR, that regulate the receptor temporal dynamics. The activation of receptors induces a cascade mechanism: CheY, a diffusible messenger protein phosphorylate (it gains a phosphoryl group) and binds a motor protein that regulates the probability of the flagella to turn and align [11–13].

The physical limits of chemotaxis are due to noise in the binding of molecules diffusing to the receptors. These limits and the quantitative model of chemotaxis are treated in the seminal study of Berg and Purcell [14] and in [9]. As a consequence of noise induced by the diffusion process, measuring gradients is not trivial for a micron-size cell. In fact the small size of this biological sensor only allows point measures of nutrient concentration. Thus, to measure gradients, *E.coli* uses a short-term memory that allows the cell to compare instantaneous measures of odor. Although bacteria have no neurons and no long-term memory, the temporal gradient is computed by the bio-chemical process happening in the methylation sites [15].

1.2 Animal chemo-sensing

Let us now switch to an entirely different physical scenario, dominated by turbulence. In the following I present an overview of biological organisms, and illustrate what is known about what they sense and how they behave, as they search for an olfactory target in a turbulent flow.

Moths are a model organism for turbulent olfactory searches. Evolution optimized their sense of smell: the life of a moth lasts few days depending on the species and the sole goal of a male moth in life is to find a female and reproduce before dying. Moths can sense, with the two antennae, low pheromone concentrations (a few hundred molecules per cubic centimeter [17]) kilometers away, many obstacles like trees or hills can stand in between a male and the female and yet future generations spawn [18]. The pioneering work by Mafra-Neto and Cardé [5] showed that male *Cadra cautella* moths track successfully an odorant when the velocity field is intermittent, while they struggle at finding the source of an odor in a laminar flow.




Fig. 1.2 Tracks of male gypsy moths approaching a source releasing sex pheromone together with bubbles (that can be visualized and provide a sense of the distribution of odor stimuli). Panels (a) and (b) show two typical tracks, with thick/thin parts referring to periods within/outside the bubbles, i.e., likely associated with high-frequency/absence of odor detection. Thick parts are known as surge, as they show consistent progression up the wind (indicated by arrows). Conversely, thin lines show crosswind motion and are known as casting. The portion of a track in panel c shows that a change in wind direction is reflected in the orientation of behavioral moves. From [16].

During turbulent navigation, moths exhibit a peculiar behavior that alternates upwind surges with crosswind casts. Specifically after detection male moths move straight toward the source (surging), when the chemical signal is lost moths exhibit body turns and loops ending up in a zigzag trajectory called casting (see Figure 1.2). To control casts and surges, it has been hypothesized that some species may have an internal clock that produces regular variation in flight direction while the animal is progressing toward the source (counter-turning). However the principles that should govern the width, frequency and duration of the casts remain unknown.

Cast and surges are widely observed in other organisms as well. *Drosophila melanogaster*, also known as fruit fly, is another experimental model system. Fruit flies measure odorants with the two antennae not only when flying but also when walking. When they fly, they surge upwind when they detect odor, similar to moths. But in contrast with moths, flies succeed also in laminar flows [19, 20]. This is because walking flies experience a smoother olfactory landscape, due to the fact that

the flow near the ground slows down and creates a stagnant layer of air. Similar to what observed in flying individuals, walking *Drosophila* move crosswind and execute a local search to get inside the plume. Once inside a complex plume cone they execute random left/right fast movements, with frequent odor encounters biasing these displacements upwind. Walking *Drosophila* has been shown to use both timing and frequency of odor encounters to advance toward the source [21]; the duration of odor signals further helps the search [22]. Other insects that do not rely uniquely on the intensity of the odor signal are the American cockroaches, they have a pair of ON and OFF cells in their antennae that responds oppositely to changes in the concentration of target odorants: ON cells activate when concentration is rising, OFF cells activate when concentration is decreasing [23].

Chemo-sensing is ubiquitous and not only limited to insects and other animals exposed to air flows. Among marine organisms, arthropods and, in particular, crustacean are the most widely studied phylum for olfaction. Atema in 1996 [24] studied lobster behavior in turbulent water plumes, showing that they can find odor sources 2 m away in 30 s. Crustaceans use of bilateral antennulae to locate the source in a short time. Evidence suggests that they often move towards the time-averaged borders of the plume where there is a strong contrast in odor perception by the two antennulae as one is located in the interior of the plume and the other outside of the plume [25]. Lobsters are believed to orient by comparing bilateral differences in the timing of arrival of the odor pointing toward the antennula that encounters odor more frequently. Sharks have also been proposed to employ the same strategy [26]. Blue crabs are another example of successful turbulent navigation, they measure chemicals both with antennulae near the mouth, both with receptors on their 10 legs [27]. As the other mentioned animals, they also display some sort of alternation between surges and casts, and their behavior has been studied for different values of turbulence, specifically tuning fluctuations increasing the bed roughness [25, 28]. The overall advantage is that they can successfully tackle more or less turbulent conditions, and they adapt their speed and cast and surge to the details of the flow.

Finally, let us stress that quantifying animal behavior is challenging. Experiments show a lot of variability among different individuals within the same population. Designing experiments to track the most relevant observables is highly non trivial; the time and length scales vary considerably and the environment is complex.

Moreover, several steps separate the external physical dynamics from the process that eventually leads to acquisition of information by the organism. Some steps are actively controlled by the organism. For example, the action of sniffing is key in olfactory search. Physical mechanisms that bring odor molecules to the receptors have a key role in the process: sniffing replace the volume of air in contact with receptors with a new volume of air and it may increase the rate at which particles reach the receptors, an analogous to sniffing is the antennulae flicking by arthropods [29]. Other steps are not controlled by organisms. For example, transduction limits the sampling frequency, identity (how to distinguish different odorant patches) and quality of odor detection and it also influences the sensory threshold (i.e the number of molecules needed to bind to a receptor to enable a response): processing the signal produce a physiological delay to behavioral response [30, 31].

Thus to summarize, the mechanisms regulating bacterial chemotaxis are well understood: the genes responsible for sensing and flagellar movements have been identified and we have mathematical models and algorithms that reproduce and explain their behavior. On the other hand turbulent navigation is not understood. The overview presented above collects common observations about animal behavior and neural recordings. Three elements are particularly relevant for the rest of the thesis: many animals adapted to turbulence display a characteristic alternation between casts and surges; different features of the odor signal may be monitored in the brain; some animals perform searches both close to the ground and higher up in the bulk of the fluid.

In my thesis I address the complex problem of olfactory navigation combining fluid dynamics with machine learning, as well as neural recordings in mice. I aim at providing a quantitative framework to study prediction and navigation in turbulent environments. I will start with an overview of the fluid dynamics of turbulent odor transport in Chapter 2 and introduce the numerical simulations I developed to compose a realistic dataset of turbulent odor plumes. I will then proceed to develop supervised learning algorithms and models to identify the most accurate predictors of source location (Chapter 3). In Chapter 4, I will present a series of experiments showing that the neural activity in mice can reflect both the intensity and the temporal dynamics of an odor cue. In Chapter 5 I will discuss turbulent navigation and propose a normative theory to understand why animals often alternate between sensing odor

close to ground, and higher up in the air. I will complement Partially Observable Markov Decision Processes with simple analytical models to show the connection of alternation with Marginal Value Theory. My work resulted in three publications listed at page 100 and indicated with letters throughout the manuscript, to distinguish them from the rest of the references (page 101) indicated with numbers.

Chapter 2

Channel turbulence with an odor source

To understand olfactory navigation it is important to quantify the characteristics of the environment where organisms live. Air and water are often turbulent and the odor cues that reach the animals are sparse. In this chapter I characterize turbulent odor landscapes through direct numerical simulations. At the beginning, I briefly present the fundamental aspects of fluid dynamics and turbulence and I explain how odor particles are transported in a turbulent flow. In the second part of the chapter I describe how, during my PhD, I realized direct numerical simulations of a turbulent channel flow with an odor source. Datasets from these simulations have been used to obtain results in refs. [A, C] that I will present in the following chapters.

2.1 A brief introduction to fluid dynamics and turbulence

Fluid dynamics is described by mass conservation coupled to the Navier-Stokes equations, which express the conservation of momentum for Newtonian fluids:

$$\partial_t \mathbf{u} + \mathbf{u} \cdot \nabla \mathbf{u} = -\frac{1}{\rho} \nabla P + \nu \nabla^2 \mathbf{u} + \mathbf{g} \quad \nabla \cdot \mathbf{u} = 0 \quad (2.1)$$

where $\mathbf{u}$ is the velocity field, P is the pressure, $\mathbf{g}$ includes external accelerations acting on the continuum (i.e. gravity), ρ is fluid density and it is constant for Newtonian fluids, ν is the kinematic viscosity. To identify a solution to the Navier-Stokes equations, initial and boundary conditions on $\mathbf{u}$ are needed. The second equation is the continuity equation, which for constant fluid density results in a condition for flow incompressibility [32]. The first equation contains the evolution of the flow, the two terms on the left hand side correspond to inertial terms. $(\mathbf{u} \cdot \nabla) \mathbf{u}$ is the advective term and it introduces a non linearity in the equation, which can ultimately lead to turbulence and makes the Navier-Stokes equations extremely challenging both from the numerical and analytical point of view. The pressure and viscous terms are components of the stress tensor, representing the internal forces between adjacent portions of the fluid. Pressure is the Lagrange multiplier that ensures global incompressibility, thus it is non local. The laplacian term models viscosity which is responsible for dissipation, originating from the resistance opposed by adjacent portions of the fluid to move relative to each other (shear). Finally, $\mathbf{g}$ represents all the external forces acting on the fluid. To understand which term of Navier-Stokes equation has more importance in different flow conditions it is possible to define an adimensional parameter, the Reynolds number: $Re = UL/\nu$, where U and L are respectively the characteristic velocity and length scale of the considered flow and ν is the kinematic viscosity, whose value is $\nu_{\text{air}} = 1.5 \cdot 10^{-5} m^2/s$ in air, $\nu_{\text{water}} = 1 \cdot 10^{-6} m^2/s$ in water at atmospheric conditions. Dimensional analysis shows that for small Reynolds numbers ($Re < 1$), viscous forces are more important than inertial forces. At these Reynolds number, the flow is laminar (smooth), like if many parallel fluid layers were moving together in the same direction. When the Reynolds number increases, symmetries break, inertial forces become dominant and the fluid evolves chaotically. Eddies of different size appear: viscosity dissipates energy that is transferred from large to small eddies (energy cascade), and typically for Reynolds $\sim o(1000)$ or more, the flow transitions to a fully turbulent statistics, where it is useful to split the velocity field $\mathbf{u}$ in the average contribution $\bar{\mathbf{U}}$ and the velocity fluctuations $\mathbf{u}'$.

When a flow is turbulent (see Figure 2.1) all scales interact: the integral scale (L) determines the maximum distance at which two different points in the fluid can be correlated and the Kolmogorov scale (η) corresponds to the size of the smallest eddies. η is also called viscous scale and results from balance between the energy flux and dissipation. The range in between the integral and Kolmogorov scales,

where both energy injection and dissipation are negligible, is called inertial range, here energy is transferred across scales by eddies fragmentation.

An observable related to the inertial range is the Taylor scale λ :

$$\frac{1}{\lambda^2} = \frac{\langle (\partial_i \mathbf{u}_i)^2 \rangle}{u_{\text{RMS}}^2}, \quad (2.2)$$

where i runs over the three spatial coordinates, λ is a quantity commonly used in literature because it is easy to measure in experiments [33]. In the inertial range it is possible to show that the dimensional expectation for the energy spectrum is $E(k) \sim \varepsilon^{2/3} k^{-5/3}$. Here k is the wavenumber, and $\varepsilon = \frac{\nu}{2} \langle (du_i/dx_j + du_j/dx_i)^2 \rangle$ is the energy dissipation rate in units of kinetic energy per unit mass and time (indexes i, j run over the three spatial coordinates). Dimensional analysis shows that $\eta \sim (\nu^3/\varepsilon)^{1/4}$ and that $L/\eta \sim \text{Re}^{3/4}$: this relation is particularly useful to understand which scales should be fully resolved in fluid dynamic simulations for different Reynolds numbers.




Fig. 2.1 Turbulence visualization. Example of highly fluctuating velocity field in a channel flow. From [34].

2.2 Odor transport

As we discussed in the Introduction, locating an odor source in air or water can be beneficial for different reasons: finding food and mates, avoiding predators, recognizing other individuals and finding a new habitat or the way home. Odor is a passive scalar, odor molecules are carried by the flow and do not modify the velocity field and the statistical properties of the two fields are decoupled. Here I will present the

two mechanisms that contribute to odor transport: diffusion and advection.

Concentration is defined as the number of molecules in a fixed volume. Diffusion is the molecule motion due to Brownian agitation and its intensity is measured by the diffusion coefficient D . Particles realised from a punctual source spread in a time t over a sphere of radius $\sqrt{Dt}$ at a certain rate J . If odorant transport was purely diffusive, it would take too long for organisms living at scales larger than few centimeters to find an odor source, in fact they would have to wait 1 day before one part per thousand of the odor concentration at the source disperses up to 5 meters from it: scalar advection is the key for locating an odor source from far away. Concentration is a scalar field and the contribution of diffusion and advection is regulated by the conservation equation:

$$\partial_t c + \nabla \cdot (\mathbf{u}c) = D\nabla^2 c + s \quad (2.3)$$

where c is the odor concentration, $\mathbf{u}$ is the velocity field and s is the contribute of a source or a sink and D is diffusivity ($D_{\text{air}} = 1.6 \cdot 10^{-5} \text{m}^2/\text{s}$, $D_{\text{water}} = 1.6 \cdot 10^{-9} \text{m}^2/\text{s}$). This equation dictates how a concentration field evolves in time. Sources or sinks can add or remove particles. There are two additional contributions: advection, which is the velocity-coupled term on the left hand side and diffusion, regulated by the laplacian term on the right hand side. Equation (2.3) is also know as advection-diffusion equation. In a turbulent flow, a patch or a filament of concentration is advected and meanwhile it is stirred by the random velocity fluctuations: the shape of the patch becomes irregular and steep gradients appear. While the concentration field is advected diffusion also acts, but advection allows a much faster and larger scale transport than diffusion. It is possible to define a non-dimensional parameter that is the ratio of the advective and diffusive contributions, it is called Péclet number: $Pe = UL/D$. When Péclet is $\ll 1$ the diffusion dominates, whereas when Péclet is large the advection dominates [35]. Batchelor scale is where scalar fluctuations are smeared out by molecular diffusion. The Schmidt number is a non dimensional parameter that quantifies the ratio of viscosity and diffusivity is the Schmidt number $Sc = \nu/D = Pe/Re$, which in air is of the order 1 and in water rises to 1000. When $Sc \gg 1$ the Batchelor scale is much smaller than the Kolmogorov scale [36].




Fig. 2.2 A shark plume tracking a prey in turbulent advection condition ($Pe \gg 1$). Cue structure is characterized by chaotic spatial and temporal distributions of concentration filaments. The time-averaged field or filament properties are often inaccessible to the receiver. From [35]

2.3 Lagrangian representation

I referred to flow velocity as a continuous field $\mathbf{u}(\mathbf{x}, t)$: this is called Eulerian representation of a flow and it represents the point of view of an external observer monitoring a portion of space. It is possible to describe the velocity field from the perspective of an observer flowing within a fluid element, this is called Lagrangian representation and it describes the flow evolution in terms of ensemble of trajectories of fluid particles. A fluid particle is a point that moves with the local fluid velocity and its position can be defined by two equations:

$$\begin{aligned}\mathbf{X}_L(t_0, \mathbf{x}_0) &= \mathbf{x}_0 \\ \frac{\partial}{\partial t} \mathbf{X}_L(t, \mathbf{x}_0) &= \mathbf{u}(\mathbf{X}_L(t, \mathbf{x}_0))\end{aligned}$$

The first equation indicates that the particle position is at $\mathbf{x}_0$ at the reference time t_0 and the second equation describes the fluid particle motion due to the local fluid velocity. Particle dynamic equation can be integrated backward and forward in time to obtain the trajectory $\mathbf{X}_L(t, \mathbf{x}_0)$ of the fluid particle at any t [37]. This representation can be very useful to describe an odor evolving in a velocity field. In the next section I will show that tracing the trajectories of single fluid parcels back in time, we can compute odor concentration in space and time.

2.4 Theoretical description of odor molecules transported by turbulence

Trajectories followed by passive particles in turbulent flows are chaotic and mixing due to random velocity fluctuations occurs rapidly. The resulting odor probability distribution is considerably different from a standard Gaussian [38]. Interestingly, the statistics of the scalar reproduces the hallmarks of fluid turbulence, despite the fact that the advection-diffusion equation is linear. In particular, the probability distribution function of scalar differences between two points separated by a distance r has a shape that depends on r : when $r \sim L$, the probability distribution is a Gaussian and it shifts to an exponential when r decreases from the integral scale to the Kolmogorov scale η [38]. These non Gaussian distributions are characterized by a larger central peak and longer tails, due to the peculiar structure of turbulent scalar fields.

As formulated in [17], it is possible to derive the asymptotic scalings of important quantities for a scalar in a turbulent field, at large distance from the source. Consider a constantly emitting odor source of size a at a rate J of molecules per unit time. The velocity field is incompressible $\mathbf{u}(\mathbf{x}, t) = \mathbf{U} + \mathbf{v}(\mathbf{x}, t)$, where $\mathbf{U}$ is a constant longitudinal velocity and $\mathbf{v}$ are the turbulent fluctuations. The ratio v/U dictates the turbulence level and it is assumed to be small. Considering a 2-dimensional plane x, y , the time series of odor concentration in every location of the spatial domain will be defined by the solution to the Navier-Stokes (2.1) and advection-diffusion (2.3) equations, where the source s is a normalized top hat function ($h_a(\mathbf{x})$) non zero for $|\mathbf{x}| < a$.

With a simplified stochastic model [17], it is possible to calculate the probability distribution of the odorant concentration field c ; the average concentration C over periods of time when $c > 0$ and the intermittency factor χ , defined as the fraction of time when the concentration is nonzero. It is also possible to determine the probability distribution of the duration of whiffs $p(t_w)$ and blanks $p(t_b)$, where whiffs are periods in when the concentration exceeds the sensitivity threshold c_{thr} and blanks are periods when the concentration is below threshold.

According to the conditions we introduced above, we can write the solution of the advection-diffusion equation (2.3) in the Lagrangian formalism:

$$c(\mathbf{x}, t) = J \int_{-\infty}^t dt' \int d\mathbf{x}' h_a(\mathbf{x}') p(\mathbf{x}', t' | \mathbf{x}, t), \quad (2.4)$$

$p d\mathbf{x}'$ is the probability that a particle is around $\mathbf{x}'$ at t' given that it was in $\mathbf{x}$ at time t , where p is averaged over the molecular noise statistics and not on the flow velocity. It is possible to determine $c(\mathbf{x}, t)$ tracking back in time particle trajectories ending in $\mathbf{x}$ at time t . All these trajectories form a scalar structure, or "puff", of size $r(t')$ that grows moving back in time (see Figure 2.3). We seek to determine if an odor puff ever overlaps with the odor source. To this end we calculate the distance between the center of mass of the odor pocket and the source; if this distance is ever shorter than the size of the puff, $c(\mathbf{x}, t)$ will be non-zero, otherwise it will vanish. Whiffs are more intense if their size when overlapping with the source is small.




Fig. 2.3 Scheme of the Lagrangian approach. The concentration c at a given location $\mathbf{x}$ and time t is expressed in terms of the history of a Lagrangian puff, that is, an ensemble of particles transported by the turbulent flow, all starting at $\mathbf{x}$ at time t and dispersing backwards in time. The concentration c is determined by the size of the Lagrangian puff when it hits the source (if it does): (a) Average values $c \sim C$ correspond to the puff hitting the source with a typical value of the size; (b) intense concentrations c correspond to the puff hitting the source with unusually small sizes; (c) the concentration c vanishes if the puff never hits the source throughout its history. (d) The sketch of a time series. From left to right: blank, the concentration c vanishes; whiff, the puff hits the source with a small size and c passes the threshold of detection c_{thr} ; blank, turbulent diffusion enlarges the size of the puff and c decays below the threshold, then c vanishes because of the puff losing contact with the source. The red strips indicate the regions of the puff overlapping with the source as the puff is swept by the turbulent flow. Figure and caption from [17].

The relation between the odor signal and the turbulent flow can be expressed by two exponents: α , γ . α is related to the single particle dynamics: at short times t a particle disperses as $(kt)^{1/\alpha}$, where k is a constant. This exponent expresses the scal-

ing of the crosswind width of the detection region based on the downwind direction $y \sim x^{1/\alpha}$, when $\alpha = 1$ the detection region is conical. The second exponent is related to the scaling of 2-points scalar dispersion $(k't)^{1/\gamma}$, where k' is a constant. A further exponent is β and it is related to the growth of a puff $d \log r / dt = t^{-1} (k'r/r^\gamma)^\beta$.

Through Lagrangian calculations [17] it is possible to express:

$$\chi = \text{Prob}(c > 0) \sim \left(\frac{k'x^{1-\gamma}}{U} \right)^{(3-\alpha)/\gamma} f\left(\frac{Uy^\alpha}{kx} \right) \quad (2.5)$$

$$C = \langle c | c > 0 \rangle \sim \frac{J}{k} \left(\frac{k'x}{U} \right)^{-(3-\alpha)/\gamma}, \quad (2.6)$$

where f is a non-dimensional function that decays rapidly for large arguments. Interestingly, intermittency decreases crosswind, while C stays constant. Moving from the centerline of the conical plume toward the edges the signal is sparser but its intensity remains constant. Close to the source (small x) the intensity of the puffs increases and the frequency depends on the exponent γ .

Only rarely very intense puffs are released by the source, we can treat this process as Poissonian and write the probability distribution $p(c)$, assuming that $C \ll c \ll c_0$, where c_0 is the concentration at the source:

$$p(c) \sim \frac{\chi}{C} \left(\frac{c}{C} \right)^{-2 + \frac{\beta\gamma}{3-\alpha}} \exp \left[- \left(\frac{c}{C} \right)^{\frac{\beta\gamma}{3-\alpha}} \right] \quad (2.7)$$

It is also possible to compute the higher moments of c : $\langle c^n \rangle \sim \chi C^n$. [17] reports experimental values for the exponents and scalings for three standard flows: jet flow, Kraichnan flow and atmospheric boundary layer, in the last section I will compute the exponents for my simulation and compare it to jet and atmospheric boundary layer.

2.5 Channel turbulence

To design a realistic turbulent environment for olfactory navigation I developed fluid dynamics simulations based on the well-studied case of a channel flow [37]. This is a three dimensional flow, where typically the height $2H$ of the channel is the smaller

dimension, while channel length L and width W are much longer than the height. Two walls are placed at ground level $z = 0$ and at $z = 2H$. Fluid flows on average along the longitudinal direction x and it is statistically independent of y . Centering the origin of the coordinates at mid-height on the inlet of the channel, we will refer to the centerline as the one dimensional line spanning along x , at $y = 0$ and $z = H$. Velocities in the (x, y, z) coordinates are respectively U, V, W and their fluctuations are u', v', w' . At the inlet ($x = 0$) the flow develops and it is not stationary. The statistical properties of the channel flow are computed in the fully developed region (large x), where the velocity statistics do not depend on x and are statistically stationary: at large distances the only dependence of velocity statistics is on z . The bulk Reynolds number for channel flow, $Re = 2HU_b/\nu$, is based on the height of the channel and on the bulk velocity $U_b \equiv \frac{1}{H} \int_0^H \langle U \rangle dz$ and fully developed turbulence is generally associated to a Reynolds number larger than 3000.

Since $\langle V \rangle$ is zero and $\langle U \rangle$ is independent of x , the averaged continuity equation reduces to $d\langle W \rangle/dz = 0$ and at the ground the velocity is null $\langle W \rangle_{z=0}$: this implies $\langle W \rangle$ to be null for all the z including $\langle W \rangle_{z=2H}$. It is also possible to derive that, in a channel flow, the mean pressure gradient along x is uniform across the flow: it is the pressure drop between the inlet and the outlet to drive the flow. Total shear stress is $\tau = \rho\nu \frac{d\langle U \rangle}{dz} - \rho\langle uw \rangle$, in the fully developed region the effects of this negative pressure gradient are balanced by the shear stress gradient $d\tau/dz = dp/dx$. In channel flow, close to the wall, viscosity dominates and it generates the wall shear stress that defines viscous length and time scales; these two quantities are related by the friction velocity $u_\tau \equiv \sqrt{\tau/\rho}$, consequently viscous lengthscale is $\delta_\nu = \frac{\nu}{u_\tau}$. A non dimensional lengthscale normalized with distance from wall is usually indicated with a + superscript, hence for example with our notation z^+ is the nondimensional distance from wall and it is defined as $z^+ = \frac{z}{\delta_\nu}$. The region near the wall is dominated by viscosity, further from the wall the contribution of Reynolds stresses becomes dominant. The larger the Reynolds number the smaller the viscous region; at the same time when the Reynolds number is large a shorter viscous lengthscale must be resolved: fluid dynamics simulations become computationally expensive and experimental measurements very difficult to perform. Based on z^+ it is possible to define different regions: the first layer close to the substrate is called inner layer ($z^+ < 0.1$); a viscous sublayer, where shear stress is negligible, extends till $z^+ \sim 5$; the viscous wall region ends at $z^+ = 50$, beyond this value we have outer layer, where viscosity does not influence $\langle U \rangle$. Outside of the viscous wall region ($z^+ > 30$),

according to asymptotic assumptions, the velocity profile, from the wall to the center of the channel, follows a log law

$$\frac{\langle U \rangle}{u_\tau} = \frac{1}{\kappa} \ln z^+ + B,$$

where $\kappa = 0.41$ is the von Kármán constant and B is another constant, phenomenologically $B = 5.2$. In the next Section I will show that my simulations follow the log law for $z^+ > 30$ as reported in the literature.




Fig. 2.4 Turbulent channel flow and odor statistics from Nek5000 Direct Numerical Simulations. (a) Mean velocity profile in fully developed turbulent channel flow. (b) Velocity profile, near the wall up to mid-channel, follows the log law when $z^+ = z/\delta_v > 30$ as predicted by theory (c) Spectra of odor fluctuations compared to the $-5/3$ prediction for turbulent signals.

2.6 Direct numerical simulations of odor in a turbulent channel flow

To reproduce a realistic odor landscape and generate the dataset showed in Figure 2.5, we solve the Navier-Stokes (2.1) and the advection-diffusion equation for passive odor transport (2.3) at all relevant scales of motion from the Kolmogorov scale η to the integral scale ($L > 600\eta$), using Direct numerical simulations (DNS). We simulate a turbulent channel flow with a concentrated odor source and an obstacle that generates turbulence by customizing the open-source software Nek5000 [39] developed at Argonne National Laboratory, Illinois. Nek5000 employs a spectral element method (SEM) [40, 41] based on Legendre polynomials for discretization [42], and a 4th order Runge-Kutta scheme for time marching. The code is written in fortran77 and C and it uses MPI for parallelization.




Fig. 2.5 Turbulent odor cues are patchy and intermittent. Snapshot of streamwise velocity (a) in a vertical plain at mid channel; odor snapshot side view at mid channel (b) and top view at source height (c). White regions mark the cylindrical obstacle. Snapshots are obtained from direct numerical simulations of the Navier-Stokes equations and the equation for odor transport (see parameters summarized in Table 2.1). (d) Typical time courses of the odor cues at locations labeled with 1 and 2 in c, visualizing noise and sparsity, particularly at location 1. Adapted from [A].

Table 2.1 Parameters of the simulation. Length L , width W , height H of the computational domain; horizontal speed along the centerline U ; mean horizontal speed $U_b = \langle u \rangle$; Kolmogorov length scale $\eta = (\nu^3/\varepsilon)^{1/4}$ where ν is the kinematic viscosity and ε is the energy dissipation rate; mean size of gridcell Δx ; Kolmogorov timescale $\tau_\eta = \eta^2/\nu$; energy dissipation rate $\varepsilon = \nu/2\langle(\partial u_i/\partial x_j + \partial u_j/\partial x_i)^2\rangle$; Taylor microscale $\lambda = \sqrt{\langle u^2 \rangle / \langle (\partial u / \partial x)^2 \rangle}$; viscous lengthscale $\delta_\nu = \nu/u_\tau$ where the friction velocity is $u_\tau = \sqrt{\tau/\rho}$ and the wall stress is $\tau = \rho \nu du/dz|_{z=0}$; Reynolds number $Re = U(H/2)/\nu$ based on the centerline speed U and half height; Reynolds number $Re_\lambda = U\lambda/\nu$ based on the centerline speed and the Taylor microscale λ ; Schmidt number ($Sc = \nu/D$), magnitude of velocity fluctuations u' relative to the centerline speed; large eddy turnover time $T = H/2u'$.

L	W	H	U	U_b	η	Δx	τ_η	ε	λ	δ_ν	Re	Re_λ	Sc	u'/U	T
40	8	4	32	23	0.006	0.025	0.01	39	0.17	0.0035	16000	1360	1	11%	$64\tau_\eta$

The three dimensional channel is divided in $E = 160\,000$ discrete elements: $200 \times 40 \times 20$ (number of elements in length $\times$ width $\times$ height); within each element the solution is expanded in 8th grade tensor-product polynomials so that the domain is effectively discretized in 81 920 000 elements. The average spatial resolution is equal in each direction $\Delta x \approx 4\eta$. To tune the velocity fluctuations a cylindrical cap of height $= 160\eta$ is added on the ground; the cylinder spans the entire width of the channel. The mesh is adapted to fit the cylinder. Fluid flows from left to right and the obstacle generates turbulence in the channel, in particular the height of the cylinder tunes the velocity fluctuations. The spatial coordinates are $\mathbf{x} = (x, y, z)$. The velocity fluctuations are defined as $\delta u(\mathbf{x}, t) = u(\mathbf{x}, t) - \langle u(\mathbf{x}, t) \rangle$; their intensity is $u' = \sqrt{\langle (\delta u)^2 \rangle}$, where averages are intended in space and time. Table 2.1 summarizes the parameters that characterize turbulence.

Each simulation runs for 300 000 time steps where $\delta t = 10^{-2} \tau_\eta$ and follows from a severe Courant criterium with $U \Delta t / \Delta x < 0.4$ to ensure convergence of both the velocity and scalar fields. Snapshots of velocity and odor fields are saved at constant frequency $\omega = 1/\tau_\eta$. Each DNS requires 2 weeks of computational time using 320 cpus.

2.6.1 Boundary conditions and odor source

We impose a Poiseuille velocity profile at the inlet: $\mathbf{u} = (u, 0, 0)$ and $u = 6U_b(\zeta - \zeta^2)$, where $\zeta = z/H$ is the vertical coordinate normalized to the height of the channel and U_b is the mean speed. We set a no-slip condition $\mathbf{u} = 0$ at the ground and on the obstacle; on the remaining boundaries we impose the turbulent outflow condition defined in [43] that imposes a positive exit velocity to avoid potential negative flux and the consequent instability it generates.

More precisely, the divergence ramps up from zero to a positive value along the element closest to the boundary: $\nabla \cdot \mathbf{u} = C[1 - (z_\perp/\Delta x)^2]$, where $z_\perp$ is the distance from the boundary and $C = 2$ is the minimal value that ensures convergence. For the odor, we impose a Dirichlet condition ($\theta = 0$) at the ground, on the obstacle and at the inlet; while an outflow condition is set at the top, on the sides and at the outlet: $k(\nabla \theta) \cdot \mathbf{n} = 0$. We introduce a source located right above and downstream of the obstacle, at coordinates $x_s = 810\eta$, $y_s = 650\eta$, $z_s = 238\eta$; odor intensity at

the source is defined by a Gaussian distribution $q = e^{[(z_1-x_s)^2+(z_2-y_s)^2+(z_3-z_s)^2]/(2\sigma^2)}$, where $\sigma = 5\eta$.

Nek5000 is divided in core routines and a small subset of routines which are most naturally customized by the users. The most significant variations of the code I have implemented were aimed at adding the Gaussian odor source, setting boundary conditions and initial conditions and modify the mesh to add the cylindrical obstacle with a body fitted grid.

2.6.2 Results

The channel flow in my simulations is characterized by a mean flow in the longitudinal direction (Figure 2.6a); I identified a region, far from the obstacle, where the flow is stationary once turbulence is fully developed. To characterize the flow we show in Figure 2.4a that the mean velocity profile has the typical shape of a fully developed turbulent channel flow [37]. In 2.4b, that the mean flow follows the law of the wall for $z^+ > 30$ (introduced in the channel flow Section), recovering classical statistics for channel turbulence. Figure 2.4b allows visualization of different regions where viscosity *vs* Reynolds stresses are dominant. Figure 2.6b-c present respectively a 2-d visualization of velocity fluctuations u' and of the energy dissipation rate ϵ , computed according to formulas presented in Section 1 of this chapter.

The odor field is emitted from a concentrated source downstream from the obstacle; it develops as a meandering filament that fluctuates as it travels downstream and soon breaks into discrete pockets of odor (whiffs) separated by odor-less stretches (blanks) (Figure 2.5b,c,d). Odor cues appear intermittent, as seen from the sparsity of their time series (Figure 2.5e). Note that depending on the sampling location, odor may be more or less sparse (compare for example Figure 2.5e left and right). The spectrum of odor fluctuations is consistent with the $k^{-5/3}$ scaling typical of turbulent transport (Figure 2.4c), More precisely the Figure 2.4c shows the two dimensional spectra

$$E(k) = \frac{d}{dk} \left(\int_{|\mathbf{k}_{xy}| < k} |\hat{c}(\mathbf{k}_{xy})|^2 d^2k \right)$$

2.6. Direct numerical simulations of odor in a turbulent channel flow




Fig. 2.6 (a) Longitudinal velocity averaged over the width of the channel (y coordinate). Black continuous line represents the average value of the wall velocity u_τ , straight black line shows the region where the turbulence is fully developed and the average profile of u is steady. (b) Profile of the intensity of velocity fluctuations $u' = \sqrt{\langle (\delta u)^2 \rangle}$. (c) Energy dissipation rate $\epsilon = \nu/2 \langle (\partial u_i / \partial x_j + \partial u_j / \partial x_i)^2 \rangle$, black line shows the average profile of over the steady region.

normalized with the scalar variance σ_c^2 , where $\hat{c}(\mathbf{k}_{xy})$ is the two dimensional Fourier transform of the scalar concentration at the height of the source. The integral of the spectra is the scalar variance. The results are shown as a function of the nondimensional wavenumber $k\eta$ where η is the Kolmogorov scale. The $k^{-5/3}$ scaling holds for $k\eta \lesssim 0.1$, consistent with previous experimental results in channel flow [44]. Note that in this regime the Schmidt number plays a minor role, because the effects of diffusivity are felt at small scales [45], but the statistics of turbulent plumes from concentrated sources are dictated by the separation of Lagrangian particles across inertial scales hence will depend weakly on scales below the Kolmogorov and Batchelor scales [17, 46]. Using equations (2.5) and (2.6) and many samples of C and χ obtained from my simulations, I fit for the exponents of C and χ as a function of distance from the source and derive estimates for the single particle and pair separation exponents. I obtain values of the exponents that change with height, consistent with the fact that particle transport is affected by the presence of the wall. At heights up to the source, I obtain $\alpha = \gamma = 1$, matching estimates for the atmospheric boundary layer [17]. If the odor is observed from a plane that lies above the source, the predictions obtained for the cases discussed in [17] are not expected to hold anymore. This is because the Lagrangian trajectories starting at the source and ending above the source encounter strong vertical gradients, that are not

accounted for in the cases discussed in ref [17]. Indeed, the exponents vary in a non trivial way, and I will treat their theoretical modeling in future work.




Fig. 2.7 Left, fit of α exponent; right, fit of γ exponent based on the equations presented in Section 4 using data from my turbulent channel simulation. Horizontal axis presents different heights (z) from very close to the bottom of the channel to slightly above mid-height (blue dots), red stars indicate experimental exponents for jet flow (J) and atmospheric boundary layer (AB). Data to generate these plots have been randomly sampled over a 2-D conical region (see Figure 2.5c), first dataset has $N = 13500$ points and second one has $M = 5000$ points, in the plots I computed the average with error bar of every exponent over these datasets.

Chapter 3

Olfactory predictions

In this chapter we ask if measuring odor signal in time at fixed spatial positions can allow an agent to infer the correct location of an odor source. We also aim at identifying what salient features of turbulent odor signals are more relevant: specifically, comparing the predictive performance of quantities related to intensity *vs* timing of odor encounters. We compose a dataset of realistic odor fields at scales of several meters using the fluid dynamics simulations presented in the previous chapter. We then develop machine learning algorithms that predict the source location based on the odor fields. Finally, we present the results obtained, we rank the different measures and we discuss when intensity *vs* timing features have better predictive power, we show how the ranking varies in space, we reach some general conclusions and we develop a theoretical framework to verify how much information can provide every individual feature. Results presented in this chapter were condensed in a paper currently under revision at eLife [A].

3.1 Odor encoding: intensity *vs* timing

At large scale, fluid environments are often turbulent and macroscopic organisms detect odors in intermittent patches, that may be separated by extended regions with no odor. The concentration of an odor is a complex signal that evolves in space and time depending on the details of the flow. A finite portion of this temporal signal can be encoded in simpler quantities that we call features. There are two broad classes of features that quantify the dynamics of olfactory cues: those that depend on odor

intensity including e.g. odor gradients in space or time, and those that depend only on its timing, i.e. on whether the odor is on or off regardless of its concentration. To compute quantities that depend on odor intensity, an accurate representation of the odor is needed. In contrast, measuring the timing of odor detection simply requires a binary representation. As anticipated in the Introduction there are evidences that suggest animals are able to identify when they detect an odor as well as how intense it is; but whether they record and rely on both kinds of information is still unclear. In the previous chapter we showed that concentration and intermittency measures contain information about source location, here I will show that they can be used to infer source location or navigate to it.

To prove that distant odor source inference is possible and to classify which class of features is more precise we apply machine learning algorithms to a dataset from simulations introduced in the previous chapter, where the odor field is emitted from a concentrated source downstream of an obstacle; it develops as a meandering filament that fluctuates as it travels downstream and soon breaks into discrete pockets of odor (whiffs) separated by odor-less stretches (blanks). Note that depending on the sampling location, odor may be more or less sparse (compare for example Figure 2.5). We report here in Table 3.1, the most important dimensional parameters in air and water (see nondimensional values in Table 2.1).

Table 3.1 Dimensional parameters of the simulation, same as Table 2.1 in Chapter 2. Assuming a flow of mean velocity $U = 50$ cm/s in air and $U=12$ cm/s in water and considering the air/water kinematic viscosity (ν) we can obtain all relevant observables in dimensional units.

	L	W	H	U_b	η	Δx	τ_η	ε	z^+
air	9.50 m	1.90 m	0.96 m	36 cm/s	0.15 cm	0.6 cm	0.15 s	$6.3e-4 \text{ m}^2/\text{s}^3$	0.09 cm
water	2.66 m	0.53 m	0.27 m	8.6 cm/s	0.04 cm	0.2 cm	0.18 s	$3e-5 \text{ m}^2/\text{s}^3$	0.02 cm

3.2 An introduction to supervised learning

In this chapter we ask if odor cues bear information about source location meters away from the source. To find the answer, we develop supervised machine learning algorithms that learn the relationship between the input (odor) and the distance from the source (output) from a large dataset of examples (data from fluid dynamics simulation). Here is a review of some key ideas about supervised learning and we

refer to standard textbooks for further details e.g. [47].

The goal in supervised learning is to infer a function f given a training set $(\mathbf{x}_1, y_1), \dots, (\mathbf{x}_N, y_N)$ of input/output pairs, where $\mathbf{x}_i$ are d -dimensional vectors and y_i are scalars. A good function estimate should allow to *predict* the outputs associated to *new* input points. To measure how close the prediction $f(\mathbf{x})$ is to the correct output y , we consider the square loss $(f(\mathbf{x}) - y)^2$. Following a statistical learning framework, the data are assumed to be sampled according to a fixed, but unknown data distribution P . In this view, the ideal solution f^* should minimize the expected loss $\langle l(f(\mathbf{x}), y) \rangle = \langle (f(\mathbf{x}) - y)^2 \rangle$ over all data distributed according to P . In practice, only an empirical loss based on training data can be measured, and the search for a solution needs be restricted to a suitable class of hypothesis. Note that, the choice of the latter is critical since the nature of the function to be learnt is not known a priori. A basic choice is considering linear functions $f(\mathbf{x}) = \mathbf{w} \cdot \mathbf{x}$. In this case, minimizing the empirical loss reduces to linear least squares $\min_{\mathbf{w}} \frac{1}{N} \|Y - X \cdot \mathbf{w}\|^2$, where X is the matrix composed of the N training data input $X = (\mathbf{x}_1, \dots, \mathbf{x}_N)^T$ and Y is the vector composed of the N labels of the training set $Y = (y_1, \dots, y_N)^T$. The corresponding solution is easily shown to be $\mathbf{w} = (X^T X + \lambda n I)^{-1} X^T Y$. In order to ensure that the model generalise well we add a regularization term, depending on the parameter λ , it balances out data fitting vs regularity of the obtained solution.

Linear models have limited predictive power (we will show this applying a linear model to our data in the next section). To generalize the applicability of the algorithm we consider a space of hypothesis composed by linear combination of D non linear functions $\phi_i : R^d \rightarrow R$. The target function is defined by the D -dimensional vector of weights $w \in R^D$: $f(\mathbf{x}) = w^T \Phi(\mathbf{x})$, where $\Phi : R^d \rightarrow R^D$ and $\Phi(\mathbf{x}) = (\phi_1(\mathbf{x}), \dots, \phi_D(\mathbf{x}))^T$. The $n \times d$ data matrix X maps into the $n \times D$ matrix $\hat{\Phi} = (\Phi(\mathbf{x}_1), \dots, \Phi(\mathbf{x}_n))^T$; the output $n \times 1$ vector is Y . In other words, each d -dimensional input vector $\mathbf{x}_i$ is mapped into a D -dimensional vector $\Phi(\mathbf{x})$ that can be treated in the same way we presented above:

$$f(\mathbf{x}) = w^T \Phi(\mathbf{x}) = \sum_{j=1}^D w_j \phi_j(\mathbf{x})$$

$$w^* = (\hat{\Phi}_{D \times n}^T \hat{\Phi}_{D \times n} + \lambda n I)^{-1} \hat{\Phi}_{D \times n}^T Y_{n \times 1}$$

To obtain a function that describes training data precisely it is often useful in practice to choose a large D ; but inverting a large size matrix is computationally expensive. To overcome this technical limit we can apply the representer theorem (for the demonstration and more details see [48]). In a nutshell:

$$\begin{aligned} w &= \hat{\Phi}^T c = \sum_{i=1}^n c_i \Phi(\mathbf{x}_i) \\ c &= (\hat{\Phi} \hat{\Phi}^T + \lambda n I)^{-1} Y \\ \implies w^* &= \hat{\Phi}^T (\hat{\Phi} \hat{\Phi}^T + \lambda n I)^{-1} Y \end{aligned}$$

$$f(\mathbf{x}) = w^T \Phi(\mathbf{x})$$

$$f(\mathbf{x}) = \underbrace{\Phi(\mathbf{x})}_{d \times D}^T \underbrace{\hat{\Phi}^T}_{D \times n} (\underbrace{\hat{\Phi} \hat{\Phi}^T}_{n \times n} + \underbrace{\lambda n I}_{n \times n})^{-1} \underbrace{Y}_{n \times 1} \quad (3.1)$$

where now we have to invert an $n \times n$ matrix, rather than a $D \times D$ matrix.

Note that to compute f from equation (3.1) we do not need every function Φ , but only the product of functions $\Phi(\mathbf{x})^T \Phi(\mathbf{x}')$. We call these products kernels and we define them as

$$K(\mathbf{x}, \mathbf{x}') = \Phi(\mathbf{x})^T \Phi(\mathbf{x}')$$

$$(\hat{\Phi} \hat{\Phi}^T)_{ij} = \Phi(\mathbf{x}_i) \Phi(\mathbf{x}_j)^T := K(\mathbf{x}_i, \mathbf{x}_j)$$

$$\Phi^T(\mathbf{x})(\Phi(\mathbf{x}_1), \dots, \Phi(\mathbf{x}_n)) = K(\mathbf{x}, \mathbf{x}_i)$$

$$f(\mathbf{x}) = \Phi(\mathbf{x})^T \hat{\Phi}^T c = K(\mathbf{x}, \mathbf{x}_i) (K(\mathbf{x}_i, \mathbf{x}_j) + \lambda n I)^{-1} Y$$

The theory of the functional spaces generated by kernels is extremely powerful [48] because D does not have to be a finite number. We used a Gaussian kernel

$K(\mathbf{x}, \mathbf{x}') = e^{-\|\mathbf{x} - \mathbf{x}'\|^2 / 2\sigma^2}$ to treat the data from the simulations. λ , the regularization parameter, together with the kernel parameters (the Gaussian width σ in our case) are free parameters that need to be tuned.




Fig. 3.1 Visualisation of a typical cross validation procedure, exemplifying the need for regularization for this problem. Top: left and right color maps show the error on the training set (left) and validation set (right) as a function of the two hyperparameters λ (Tikhonov regularization parameter) and σ (width of the Gaussian kernel). Bottom right, test and training errors for $\lambda \rightarrow 0$ as a function of $1/\sigma$. For large values of $1/\sigma$ the solution overfits the data, for small values of $1/\sigma$ the solution does not overfit but is unstable. Bottom left, test and training errors for λ_{optimal} as a function of $1/\sigma$. For large values of $1/\sigma$ the solution overfits the data. There is an optimal value of $1/\sigma$ which minimizes error on both training and validation and is stable. From [A].

Kernel methods offer a number of advantages. They are nonlinear and non-parametric (there are no strong assumptions over the form of the function f), the

complexity of the model can adapt to the problem at hand and indeed the algorithm learns any kind of continuous function provided enough data. This can be contrasted to linear models that clearly cannot learn any nonlinear function. Moreover, by tuning the hyper-parameters λ, σ more or less complex shape can be selected. When λ is small we are simply fitting the data, possibly at the price of stability, whereas for large λ we are favoring simpler models. With small σ we allow highly varying functions, whereas with large enough σ we essentially recover linear models.

Indeed, the choice of these parameters is crucial as tested and visualized in Figure 3.1. Here it is shown that for $\lambda \rightarrow 0$, the solution incurs in the well known stability issues for large σ and overfitting for small σ . We note that ideally one would want to choose these hyper-parameters minimizing the test error, however this would lead to overoptimistic estimates of the prediction properties of the obtained model. Hence, we consider a hold-out cross validation protocol, where the training data are further split in a training and a validation sets. The new training set is used to compute solutions corresponding to different hyper-parameters. The validation set is used as a proxy for the test error to select the hyper-parameters with small corresponding error. The prediction properties of the model thus tuned is then assessed on the test set.

3.3 Kernel ridge regression to infer odor source location

In order to find what are the best predictors of source location and how ranking depends on the statistics of the odor, we need to detail more specifically the input and output of the algorithm.

From Nek5000 simulations we obtain the odor concentration field $c(\mathbf{z}, t)$ which varies stochastically in space and time as a result of turbulent transport. Here $\mathbf{z} = (z_1, z_2, z_3)$ is a location in the three dimensional space and t is time. To obtain a standard dataset we choose a 2D-plane fixing the height, each snapshot from the simulation has dimensions 1600×320 (number of points in the downwind direction $\times$ crosswind direction) and we consider only the conical region where odor can be detected, the “cone of detection” (Figure 3.2a). We define the odor plume as the region where the probability of detection computed over the entire simulation is

larger than 0.35. The initial evolution up to $300\tau_\eta$ is excluded from the analysis as odor has not yet reached a stationary state. At stationary state we save 2700 frames at frequency $\omega = 1/\tau_\eta$ per simulation. Thus at each spatial location we have the entire time evolution composed of 2700 time points at regular intervals of τ_η . We partition each simulation in fragments, called time series ($\mathbf{c}_i$), with M snapshots (duration $M\tau_\eta$). Most simulations are shown for $M = 100$, thus for each spatial location we have 27 time series of the same duration (except for results leading to Figure 3.5a.1-2, where we vary memory from $10\tau_\eta$ to $250\tau_\eta$ resulting in 270 to 10 time series per location respectively).

Thus each time series is a vector $\mathbf{c}_i = (c(\mathbf{z}_i, t_i), \dots, c(\mathbf{z}_i, t_{i+M}))$, where $t_{i+M} - t_i = M/\omega$ is the temporal span of the time series, or memory. From each time series $\mathbf{c}_i$ we calculate five features $x_i^1, \dots, x_i^5$, where x_i^1 is the temporal average of the concentration during whiffs in the time series $\mathbf{c}_i$; x_i^2 is its average slope (time derivative of odor upon detection, averaged across whiffs within $\mathbf{c}_i$); x_i^3 is the average duration of blanks (stretches of time when odor is below detection within $\mathbf{c}_i$); x_i^4 is the average duration of whiffs (stretches of time when odor is above threshold within $\mathbf{c}_i$); and x_i^5 is the intermittency factor (the fraction of time the time series $\mathbf{c}_i$ is above threshold). The first two features quantify intensity of the odor and rely on a precise representation of odor concentration, while the latter three quantify timing of odor encounters and are computed after binarizing the odor.

The threshold c_{thr} used for binarization is adaptive i.e. $c_{thr} = 0.5\langle c | c > 0 \rangle$, where the average is computed over each time series separately. The threshold thus varies from $c_{thr} = 0.5c_0$ at the source to $c_{thr} = 10^{-6}c_0$ at the farthest edges of the cone, where c_0 is the concentration at the source. The choice of an adaptive threshold was suggested in [49].

Our input $\mathbf{x}_i = (x_i^1, \dots, x_i^d)$ is composed of d -dimensional vectors of features and we will focus on $d = 1, 2, 5$. We seek to infer distance from the source, thus our output y is the coordinate of the sampling point $\mathbf{z}$ in the downwind direction, i.e. $y = z_1$, with the source placed at the origin (see sketch in Figure 3.2a). We also investigated what happens in the crosswind direction, $y = z_2$. We train the algorithm by providing N examples of input-output pairs $(\mathbf{x}_i, y_i)$ selected randomly from the full simulation, and obtain the function that connects input and output: $y \approx f(\mathbf{x})$.

The training set and test set are obtained by extracting $N = 5000$ (unless otherwise stated) and $N_t = 13500$ time series portions of duration $M\tau_\eta$. To select these M -long




Fig. 3.2 Individual features enable inference in two dimensions. (a) Sketch of the geometry. (b) Test error χ for inference using individual features as input. (c) Predicted vs actual distance for inference. Prediction for representative test points (grey circles); 30th to 70th percentile (patch, same color code as in (b)); trivial prediction $f(\mathbf{x}) = \langle y \rangle_{\text{test}}$ (solid horizontal line, corresponds to $\chi = 1$); exact prediction (bisector, corresponds to $\chi = 0$); dispersion away from the bisector visualizes the prediction error. Results are obtained with a supervised learning algorithm based on regularized empirical risk minimization. Each input datum x_i is one individual scalar feature computed from the time course of odor concentration measured at location $\mathbf{z}_i$ at 100 evenly spaced time points with sampling frequency $\omega = 1/\tau_\eta$, where τ_η is Kolmogorov time. The training/test set are composed of $N = 5000$ and $N_t = 13500$ data points respectively. From [A].

time series we extract random locations $\mathbf{z}_i$ to cover homogeneously the cone, i.e. with flat probability within the cone, and random initial times t_i , with the training in the first half of the time history and the test in the second half of the time history. Time series that remain entirely under threshold are excluded.

The parameters λ and σ are obtained through 4-folds cross validation: the training set is split in 4 equal parts, 3 are used for training and 1 for validation. The empirical risk is computed on the validation set and averaged over the 4 possible permutations, systematically varying the hyperparameters λ, σ . The couple of

3.4. Intensity and timing features allow a robust prediction of the source position

hyperparameters that minimize the empirical risk over the validation set is selected through grid search using an 8×8 regular grid and further refined with a 4×4 subgrid. Results are insensitive to further refinement because there is a large plateau around the minimum, as shown in Figure 3.1. The optimal hyperparameters are used to compute the best function. The error χ used throughout the manuscript is simply the normalized test error $\chi = \sum_{i=1}^{N_t} [y_i - f(\mathbf{x}_i)]^2 / \sum_{i=1}^{N_t} [y_i - \bar{y}]^2$. We implemented Kernel ridge regression, as discussed in Section 3.2, using FALKON [50], a fast algorithm for matrix inversion (the number of iterations is set to 5 and the number of Nystrom centers is equal to the number of points in the training set) and we used it both for training and test.

3.4 Intensity and timing features allow a robust prediction of the source position

To evaluate if it is possible to infer source location from odor time series, we pick the two dimensional plane that contains the source at height $H/4$. The first result is that individual features ($d = 1$) bear useful information for two-dimensional source localization even at several meters from the source. Performance is quantified by the normalized squared error χ averaged over the N_t points in the test set.




Fig. 3.3 Test error at source height, for prediction in the crosswind direction. Symbols as in Figure 2. From [A].

For this dataset, intensity features rank higher than timing features (Figure 3.2b-c), consistent with previous work [24] and predictions are more accurate in the crosswind than in the downwind direction (compare Figure 3.2b-c and 3.3a-b). For reference, a random guess with flat probability within the correct lower and upper bounds yields $\chi_{\text{random}} = 2$; whereas a target function $f_{\text{trivial}}(x) = \langle y \rangle_{\text{test}}$ that learns the average of the output over the test set yields $\chi_{\text{trivial}} = 1$.

We tested the performance of a linear least square algorithm as a benchmark to the Gaussian kernel. In Figure 3.4, we show that the choice of linear models has limited predictive power and does not allow to rank features.




Fig. 3.4 Linear least square algorithm. Prediction error for a least squares algorithm assuming the target function is a linear function of the input (no regularization). Performance is poor regardless of the input features and the dataset (dataset A to E are defined as for Figure 3.11). From [A].

Next, we analyze whether and how the sampling strategy affects performance and ranking of the features. Most results are shown for a memory of $100\tau_\eta \approx 15\text{ s}$. Performance improves with longer memory (Figure 3.5a.1-2), because this allows to better average out noise and obtain more stable estimates of the features. But improvement follows a slow power law so that waiting for example 20 times longer yields predictions only about twice as precise. On the other hand, waiting as little

3.4. Intensity and timing features allow a robust prediction of the source position

as $10\tau_\eta \approx 1.5$ seconds still allows to make predictions, albeit less precise. We then verify whether performance may improve with a larger training set.




Fig. 3.5 The sampling strategy affects performance but not ranking. Left panel shows results for prediction in downwind direction, right panel in crosswind direction. (a.1-2) Error χ as a function of memory in units of Kolmogorov times τ_η ; memory is defined as the duration of the time series of odor concentration $\mathbf{c}_i = (c(\mathbf{z}_i, t_i), \dots, c(\mathbf{z}_i, t_{i+M}))$ used to compute the five features $x_i^1, \dots, x_i^5$, i.e. $\text{memory} = t_{i+M} - t_i = M/\omega$. Red and pink: Performance using $\mathbf{x}_i = x_i^1$ (average concentration) and $\mathbf{x}_i = x_i^5$ (intermittency factor). The number of training points and the frequency of sampling are fixed, $N = 5000$ and $\omega = 1/\tau_\eta$. Dotted, dashed and solid grey lines are power laws with exponents $-1/5$, $-1/10$ and $-1/4$ respectively to guide the eye. (b.1-2) Error as a function of number of points in the training set N , with $N_t = 13500$ points in the test set, memory = $100\tau_\eta$ and $\omega = 1/\tau_\eta$. Color code as in (a). (c.1-2) Performance using the five individual features as input with $N = 5000$, $N_t = 13500$, memory = $100\tau_\eta$ sampling odor at frequency $\omega = 1/\tau_\eta$ (empty bars) and $\omega = 10/\tau_\eta$ (filled bars). From [A].

Because we infer distance from an individual (scalar) feature, the problem is one dimensional and we find that a small number of training points, which we indicate with N , is sufficient to reach a plateau in prediction performance (Figure 3.5b.1-2). We choose $N = 5000$ training points to ensure statistics is sufficient to converge both for this case and for models that pair more than one feature (discussed late in this section). Finally, sampling more frequently than once per Kolmogorov time does not essentially affect the results nor ranking (Figure 3.5c.1-2).

Another relevant parameter is the threshold, we showed it has little effect on the results (see Figure 3.6, left). Fixed thresholds were tested and discarded because

results depend sensibly on the threshold and the optimal threshold varies with the dataset in non-trivial ways (Figure 3.6, right). Furthermore, adaptive thresholds that are defined based on purely local information appear more plausible for a biological system that has no information on the intensity of the source.




Fig. 3.6 Fixed vs adaptive threshold. Test error using an adaptive threshold (left column) vs a fixed threshold (right column) for different simulations described in the main text. Adaptive thresholds are defined as a fraction of the local average concentration, $\langle c \rangle_{\text{local}}$, computed over the memory M/ω . Fixed thresholds are defined as a fraction of the global maximum concentration c_0 . Results are robust with respect to the choice of adaptive threshold, whereas they vary considerably with the choice of fixed thresholds. Large fixed thresholds (marked with yellow squares) prevent odor detection in dilute regions i.e. far from the source or from the substrate. From [A].

To summarize, I have showed that linear models are not sufficient to predict the correct location of the odor source taking as input intensity and timing measurements. On the other hand, Gaussian kernels allow robust predictions based on single features ($d = 1$). Interestingly, sampling at a higher frequency does not improve the performance or vary the ranking of the features; while increasing the memory is beneficial, but short memories are as well effective.

We now increase d to 2, i.e. we pair different features together, exploring all the possible combinations: the performance improves in some cases, but not always. In

3.4. Intensity and timing features allow a robust prediction of the source position

particular, it is possible to pair features within the same class (intensity or timing) or to take one feature from every category. The results are showed in Figure 3.7 for downwind, and Figure 3.8 for crosswind direction: it is clear that pairing two features of the same category results in little to no improvement. In contrast, combining one intensity and one timing feature improves performance considerably, up to 65%. This result can be understood by mapping the error done by individual features in space (Figure 3.9), showing that intensity and timing features are complementary, i.e. intensity features perform well in locations where timing features perform poorly. More precisely intensity features offer a good performance along the direction parallel to the cone edge, while timing features allow robust prediction along the opposite diagonal (see Figure 3.9).




Fig. 3.7 Pairing one timing feature and one intensity feature considerably improves performance. a) Error χ obtained with individual features (full bars) and pairs of features (empty bars). Grey and black indicate pairings of two intensity features and two timing features respectively; green indicates mixed pairs of one timing and one intensity feature. (b) Performance (left) and relative improvement over the best of the two paired features (right). Results for the median (bottom) and the 95th percentile (top). Within each table plot, rows from bottom to top and columns from left to right are labeled by the 5 individual features: A (average, x^1), S (slope x^2), B (blanks x^3), W (whiffs x^4), I (intermittency x^5). Results with individual features are shown on the diagonal; results pairing feature i and feature j are shown at position (i, j) . Mixed pairs provide both the best performance and the largest improvement over individual features. From [A].

3.4. Intensity and timing features allow a robust prediction of the source position




Fig. 3.8 Effect of pairing two individual features, for prediction in the crosswind direction. Symbols as in Figure 3.7. From [A].

When using two features, we increase dimensionality of the input space and expect that more data points may be needed to converge. Thus we test dependence of performance on the number of training points (N). We find that indeed, more data are needed for this model to converge, compared to the single model (compare Figure 3.10 and Figure 3.5b.1-2). To ensure robust convergence we set $N=5000$.




Fig. 3.9 Test error mapped in space. Test error mapped in space for our kernel ridge regression algorithm that takes individual observables in input and predicts distance in the downwind (left) and crosswind (right) direction for simulation *b* (see Figure 3.11). The test error is color coded from 0 (blue) to 2 (yellow) corresponding to the limits of perfect prediction and random prediction. From [A].

3.5 Performance in space

3.5.1 Dependence on the sampling height

We next seek to clarify whether the results depend on space. To this end we compose five different dataset, *a* to *e*, obtained by extracting odor snapshots from horizontal planes at source height (*b*), above the source (*c* to *e*), and below the source (*a*) (Figure 3.11a). From *a* to *e*, sparsity increases and intensity decreases (Figure 3.11b) simply because closer to the boundary, where fluid velocity is null, the air slows down and the odor accumulates. By analyzing performance across these dataset, we find that ranking of individual features shifts considerably. The two intensity features outperform all timing features when the dataset is not very sparse (dataset

a-b, Figure 3.11c, 3.11d right). In contrast, two timing features (intermittency factor and blank duration) outperform all others for the more sparse and less intense dataset *d-e* (Figure 3.11c, 3.11d left). Whiff duration performs poorly in *d-e* because intermittency is too severe and whiffs are short in duration thus bear little information (the average whiff duration is 1 to 7 time steps in over 90% of the time series). Although the ranking of individual features shifts with height, pairing one intensity and one timing feature remains the most successful strategy across all heights (Figure 3.11c, 3.11d). In contrast, combining all five features contributes little improvement (Figure 3.11c, dark green curve).




Fig. 3.10 Prediction error as a function of the number of points in the training set for individual features and pairs of features. Based on this analysis we chose $N=5000$. From [A].

3.5.2 Dependence on the source distance

Let us now focus on the plane at source height and separate locations based on their distance from the source. We assemble a distal dataset and a proximal dataset, composed of points that are further and closer than 2330η from the source respectively (Figure 3.12a). The odor is more intense and more sparse closer to the source and it becomes more dilute and less sparse with distance from the source (Figure 3.12b). Performance of individual features degrades with distance (Figure 3.12d). Intensity

features clearly outperform timing features at close range, as seen both from various percentiles of the test error (Figure 3.12d, left) as well as the full distribution (Figure 3.12c, left). The disparity between timing and intensity features disappears in the distal problem: the error distribution for all individual features is essentially superimposed except for the tails (Figure 3.12c, right and inset), which cause small differences in the median and other percentiles of the error (Figure 3.12d, right). Remarkably, mixed pairs outperform all individual features in both the distal and proximal problems (Figure 3.12c-d). In the aggregate, results demonstrate that, even within a single turbulent flow, ranking shifts considerably. Namely, measuring timing of odor encounters is most useful in regions where the odor is dilute, i.e. far from the source and from the substrate; whereas measuring intensity is most useful in concentrated conditions, i.e. close to the source or the substrate.




Fig. 3.11 Ranking shifts with height from the ground. (a) Datasets *a* to *e* correspond to data obtained at heights $z/H = 25\%$, 37.5% , 50% , 55% and 65% respectively. (b) Distribution of intensities (top) and intermittency factors (bottom) over the training set from *a* to *e* (left to right). Moving away from the boundary, the odor becomes less intense and more sparse. (c) Median performance as a function of average intermittency factor of the training set for individual intensity (grey) and timing (black) features, mixed pairs of one intensity and one timing feature (green) and all five features together (dark green). (d) Predicted vs actual distance, to visualize a representative subset of the results in (c), scale bar $10^3\eta$. Ranking depends sensibly on height: intensity features outperform timing features near the substrate, where there is more odor and it is more continuous; timing features outperform intensity features further from the substrate where there is less odor and it is more sparse; mixed pairs perform best across all conditions; combining five features provides little to no improvement over mixed pairs. From [A].




Fig. 3.12 Ranking depends on distance from the source. (a) At source height, the dataset is split in proximal (distance $< 2330\eta$) and distal (distance $> 2330\eta$). (b) Distributions of average odor intensity (left) and intermittency factor (right) over the training set; closer to the source, the odor is more intense and more sparse. (c) Distribution of test error for the proximal (left) and distal (right) problem showing intensity features (grey) outperform timing features (black) at close range, but not in the distal problem where differences in the error distribution are limited to the tails (see insets). Mixed pairs of features (green) outperform individual features either marginally (left) or considerably (right). (d) Percentiles of the error distribution in (c) for the proximal (left) and distal (right) problems confirming the picture emerged from (c). From [A].

3.6 Modeling the predictive power of individual features

All the results I have presented in this chapter are obtained applying kernel methods to the odor concentration fields from the fluid dynamics simulations: I used the normalized mean square error computed over the test set to rank the different features and to determine where they perform best. Here, I aim at analyzing the dataset with a theoretical approach. The goal is to calculate the predictive power of individual features at different heights. To compute the predictive power of an individual feature x as a function of distance from the source we define the normalized mean square error as:

$$\chi = \frac{\int_0^A (y - f(x))^2 \mu(x, y) dy}{\int_0^A (y - \bar{y})^2 p(y) dy} \quad (3.2)$$

where A is the length of the cone, $\mu(x, y)$ is the joint probability distribution of the input-output pair (x, y) and $p(y)$ is the prior on the output y (distance from the odor source, which is uniformly sampled). Here $f(x)$ is the model function obtained through the machine learning algorithm. The algorithm is designed so that $f(x)$ approximates the target function, i.e. the best possible function f^* to predict the unknown output y . If I was given the conditional probability $p(y|x)$, then the best predictor of y would simply be its expected value:

$$f(x) \approx f^*(x) = \langle y|x \rangle = \int_0^A y p(y|x) dy \quad (3.3)$$

To solve eqs (3.2),(3.3) we compute the joint, marginal and posterior distributions, given a prior on y and the likelihood $p(x|y)$:

$$\mu(x, y) = p(x|y)p(y) \quad (3.4)$$

$$p(x) = \int_0^\infty \mu(x, y) dy \quad (3.5)$$

$$p(y|x) = \frac{\mu(x, y)}{p(x)} \quad (3.6)$$

where the prior is $p(y) = 2y/A^2$ (thus $\bar{y} = 2A/3$ and the denominator in eq (3.2) is $A^2/18$).

The likelihood is dictated by the fluid dynamics of odor plumes from a concentrated source. Our features are sample averages of intensity or timing attributes of the odor signal, in the limit of large samples, they are normally distributed:

$$p(x|y) = \mathcal{N}(x - g(y), s(y)) \quad (3.7)$$

Assuming that we are in this limit, we can use our data to find empirical estimates of $g(y)$ and $s(y)$. This completes the theoretical framework: I solve numerically equations (3.2) to (3.6) with the assumption (3.7) and the empirical estimates for $g(y)$ and $s(y)$; the results are presented in in Figure 3.13 where we show that the predictive power of individual features qualitatively matches the results from machine learning (see Figure 3.11). Intensity features have a strong predictive power near the ground (location a-b); further away, at the height of the obstacle and above, their χ worsen and is comparable to the predictive power of timing features.




Fig. 3.13 Estimated predicted power of individual features using the theoretical framework outlined in the text with the empirical likelihood. Interestingly, the predictive power of the individual features is in agreement with what observed applying supervised learning (see Figure 3.11). The only feature with a surprising behavior is whiff duration, probably their distribution is not well approximated by a Gaussian. From [A].

To move beyond empirical estimates of the likelihood and generalize predictions to other kinds of flows we can leverage asymptotic arguments proposed in [17] and

illustrated in Chapter 2. Turbulent plumes at distances much larger than the size of the source $y/a \gg 1$, where a is the size of the source can be asymptotically described with a stochastic model that neglects high order correlations. An appealing property of these models is that the probability distribution of odor concentration can be written in terms of fundamental exponents characterizing the flow. Adapting these models we predict that $g(y)$ and $s(y)$ behave as power laws and we can calculate their exponents for features x_1 , x_3 and x_4 . However, the joint distribution (3.4) is sensitive to the prefactor as well. I plan to finalize this generalization by calculating exponents and prefactors, as a follow up of my work [A].

3.7 Conclusions

Our results demonstrate that within the cone of detection, the time course of an odor bears useful information for source localization even at meters from the source. We find that the concentration and the slope of a turbulent odor signal, averaged over a memory lag, are particularly useful to predict source location at close range or near the boundary. These features quantify the intensity of the odor and its variation. The primacy of the intensity features wanes in more challenging conditions, e.g. moving away from the source or away from the boundary. In these portions of space, where the odor is scarcer, features that quantify timing of odor detection become as effective as intensity features, or more effective. One of the best studied example of olfactory search in dilute conditions is arguably the case of insects. Interestingly, olfactory receptor neurons in insects appear to encode efficiently information about timing across a wide range of intensities [49, 51].

Note that while the statistics of an odor plume clearly depends on all details of the flow and the source, see e.g. [52, 17, 53], here we keep all of these parameters constant and demonstrate that even within a single flow, odor dynamics and the best predictors vary considerably in space. This begs the next question: do organisms switch between different modalities depending on attributes of odor dynamics, which will vary in space? In the next chapter I will present our analysis of neural data in the olfactory bulb of mice, showing that the neural representation of an odor depends on its intermittency.

We find that features within the same class are redundant whereas features from different classes are complementary. Indeed, features of the same class have similar patterns of performance in space, but each class has a distinct pattern (Figure 3.9). As a consequence, measuring both timing and intensity is beneficial, but using more than one feature to quantify either timing or intensity provides no advantage. Combining all features does not improve over the performance of mixed pairs, consistent with redundancy within each class. Note that there is no fundamental reason to expect features from the same class to be redundant, and further work with a larger library of features is needed to prove or disprove this notion.

Importantly, mixed time/intensity pairs of features outrank individual features robustly, i.e. in all portions of space, regardless of distance from the source and from the ground. This is in contrast with individual features and suggests relying on simultaneous timing and intensity features is advantageous when odors are sensed at various distances from the source and from the substrate. Interestingly, the coexistence of bursting olfactory neurons and canonical olfactory neurons in lobsters suggests these animals are in fact able to measure simultaneously timing and intensity [54, 55], which is consistent with the increased predictive power of the mixed pairs of features. Similarly, in mammals, optogenetic activation of the olfactory bulb [56] demonstrates that both kinds of measures guide behavior (lick vs no lick).

In this chapter, we have investigated the problem of predicting the location of a target from measures of the time course of a turbulent odor. Previous work explored a related question, i.e. how to best represent instantaneous snapshots of the odor to encode maximum information about source location [57]. The two approaches are not immediately comparable: first, [57] consider few snapshots of the odor, rather than measures of its time course. Second, maximizing information does not guarantee good predictions (to make predictions information needs to be extracted and processed, and importantly the focus is on new data that were not previously seen). We provide two comments that are relevant if information is the limiting factor for prediction accuracy: (i) binary representations were suboptimal in all conditions considered in [57, 58], i.e. at few tens of cm from the source. This is consistent with our results in concentrated conditions, where timing features -accessible through binary representations- are suboptimal. Our evidences suggest however that the

result may not hold in more dilute conditions, where the gap between binary and more accurate representations should become increasingly small. (ii) Individual snapshots of odor from [57, 58] contained 1 to 2 bits of information about source location, but allocating more resources to represent how the odor varies in time was found informative [57, 58]. Our mixed pairs of features at close range achieve precision of 5% to 6%, corresponding to coding for position with words of 4 to 4.3 bits. Our results thus confirm that memory is indeed useful, but the gain does not increase indefinitely with further memory.

I will discuss olfactory navigation in Chapter 5, it is interesting to anticipate that recent results investigated gradient descent algorithms using either concentration alone [59], or various measures of timing and intensity [60–62]. Overall, both intensity and timing appear to have a potential to lead to an odor source, consistent with our results on individual features. A combination of the two kinds of features was found beneficial in [62], consistent with our results on mixed pairs. Whether good predictors may be good variables for navigation in more general contexts remains to be understood.

Here we have analyzed the features that enable the most accurate prediction of source location. We add a few observations about the significance of the results for animal behavior. First: whether animals rely on features from either class will depend on what features best support behavior. It is often implicitly assumed that features that bear reliable information on source location are also the most useful for navigation. However, this connection between prediction and navigation is far from straightforward and more work is needed to establish whether accurate predictions imply efficient navigation. Second: animals are unlikely to have prior information on the details of the odor source, e.g. its intensity. Timing features are more robust than intensity features with respect to the intensity of the source and may thus be favored regardless of their performance, which was argued in [63]. In our work, timing features are precisely invariant with source intensity because we define the detection threshold adaptively. More realistic conditions have to be evaluated in future works, where dependence on source intensity emerges as a result of non-linearities that we did not model in this work. These effects emerge for example, close to a boundary which partially absorbs the odor [64], or in the case of fixed thresholds, although this dependence is weak in the far field where timing features are most useful [17]. Third:

we have focused on predicting source location from within the cone of detection, where an agent will detect the odor quite often. However, a crucial difficulty of turbulent navigation is to find the cone itself. We cannot address the problem of predicting source location from outside the cone because detections are so rare that we lack statistics. The distinction between inside and outside the cone of detection is key for navigation with sparse cues [65] and the distinction between looking for the odor plume and looking for the odor source within the plume will be treated extensively in Chapter 5.

Chapter 4

Neural recordings in the olfactory bulb of mice

In this chapter I will present how the olfactory system functions in mammals, with a focus on mice. Olfaction is one of the least understood sensory systems, but it is fundamental for many species to locate food and partners. After focusing on the biology of olfaction we will dig into the neuroscientific aspects and I will present some experimental evidences related to the theoretical predictions presented in the previous chapter: can animals measure intensity and timing features of an odor signal when exposed to a turbulent flow? In Lewis et al. [B] I collaborated with a group of experimentalists who recorded the neural activity in head-fixed mice, exposed to different odor laden air flows. I will present the experimental setup, characterize the flow and the odor and finally show that a consistent portion of mouse brain follows odor concentration dynamics, while, at the same time, the strongest responding areas are the best at following odor plume fluctuations. Results I report in this chapter are published on Frontiers in Cellular Neuroscience [B].

4.1 A brief introduction to the olfactory system in mice

Humans do not rely on olfaction as much as they do on vision and this is probably the reason why olfaction is the least understood of the five senses. As we stated in

the introduction to understand chemo-sensing we first need to know what an odor physically is and what sensors are able to detect odorants. Then we can focus on the odor-sensor interaction, and finally we want to understand how the brain processes the olfactory stimulus. Different molecules have different odors, they are bound by different odor receptors and they elicit different neural response: also small variations in molecular structure can lead to a wildly different odor perception. We can say that molecules have a smell themselves but to detect it a nose-like-sensor is needed and different organisms evolved different kinds of sensors. Chemical receptors are housed in organs that vary across species: mammals have a nose with two cavities called nostrils, some crustaceans and insects have antennae, cephalopods have suckers distributed all over their arms. Our focus is on distant chemo-reception and not on contact chemo-reception that is possible when the detector is "touching" the odor source. Shape, geometry and number of the sensors can vary widely. Smelling is often an active process where the mammal inhales air for breathing which subsequently reaches the olfactory receptors; crustaceans and insects actively move their antennae to enhance odor detection. Once odor molecules bind to the appropriate receptors a nervous signal is transmitted to the brain, we will discuss in detail how the olfactory system functions in mice: the same structure is well preserved across mammals.

In mammals odorant detection occurs in the nasal cavity: here a mucus layer embeds olfactory sensory neurons (OSNs), also called olfactory receptor neurons (ORNs), that are the proper detectors for the odors. Molecules of odorants, carried by the airflow, reach the olfactory sensory neuron cilia; here they are sorbed to the surface and diffuse into the bulk (Figure 4.1A). The diffusion across the mucus layer is mediated by olfactory binding proteins, abundant throughout the mucus. As they travel through mucus, odor molecules solubilize and bind the odorant receptors. This binding triggers a transduction pathway that controls the opening of plasma membrane ion channels and subsequent voltage changes in the OSN [16]. What varies among different species is the number of different olfactory receptors that spans from dozens to thousands. Different receptors have different transduction mechanisms [66, 67]. A receptor can be activated by different molecules and a single odorant can interact with different receptors [68]. The binding affinity between chemicals and receptors can vary significantly and a quantitative description has not been proposed yet. Unfortunately most experiments in the laboratory tested

4.1. A brief introduction to the olfactory system in mice




Fig. 4.1 Sketch of neural apparatus that sense and process olfactory stimuli. (A) Schematic of the nasal cavity and the neural components involved in olfaction in mice. Airflow brings odors over the olfactory epithelium during inhalation. OSNs are electrically activated by odors and signal to the brain through their axons converging on glomeruli in the OB (light blue). Processed information is carried from the OB to multiple brain areas by mitral-tufted cells (MTCs). (B) Common neural circuit motif in rodents and insects. OSNs expressing a particular receptor type out of a large repertoire (indicated by like colors) converge selectively in individual glomeruli in the OB, making connections with MTCs. Local circuit elements in the OB include inhibitory neurons (shown as a black circle) that receive excitation from MTCs and reciprocally inhibit them. MTCs project to multiple brain regions, including the piriform cortex, where they are thought to make dispersed, random, and sparse connections (shown as intersecting wires, with connections denoted by small circles). Figure from Reddy et al. [16]

how single odorants activate receptors but in nature smells are often a mixture of many chemicals. Recent studies showed that olfactory system is highly nonlinear in detecting odorant mixtures [69, 70]; furthermore another recent work shows that mice may leverage the temporal correlations among fluctuations in the different components of a mixture to identify its identity [71]. Information about correlation is treated further downstream of the OSN. Axel and Buck [72] showed that OSNs express only one type of receptor and neurons with the same receptor send their axons to a structure called glomerulus (see Figure 4.1B), so at the early stages the olfactory information upcoming from different receptors is kept separated. Decoding olfactory signals becomes a matter of combinatorial calculation based on which receptors get activated by a certain odorant [73]. Humans with around 1000 different receptors can discriminate among many odors although the exact figure is still debated, since combinations based on 1000 receptors are potentially unlimited.

The structure with glomeruli, that are made of mitral and tufted cells, and their connections is named olfactory bulb (OB) and it is downstream of the olfactory neuron receptors and upstream of the olfactory cortex, the amygdala, and the ventral striatum (see Figure 4.1B) [74]. Information is transmitted to different brain areas [75], but which of these regions are involved in odor-guided navigation is still unclear. Architectural complexity of the olfactory system increases when we consider the neural structure downstream of the olfactory bulb. Indeed, while individual glomeruli are relatively isolated units, synaptic connections downstream of glomeruli appear intricate and random: their functions should be normalization, redundancy reduction, and decorrelation of the signal, but how these are accomplished is not understood [76].

The focus of our work is the olfactory bulb (OB): we will track the neural response of mitral and tufted cells to single odors. We aim at understanding how glomerular activity depends on temporal dynamics of the odor.

4.2 Monitoring neural response with calcium imaging in different flow conditions

In Lewis et al [B] we showed for the first time that the rapid fluctuations present in natural olfactory scenes significantly affect the activity of glomerular mitral-tufted cell populations in the mouse olfactory bulb. The novelty of our approach laid in the very realistic experimental setting: a mouse was located in a 80 x 40 x 40 cm (length, width, height) acrylic wind tunnel where air velocity was controlled by a vacuum at the rear of the wind tunnel, posterior to the animal's location, adjusting the strength of the vacuum exhaust was possible to tune the airflow. The location of the mouse was fixed at mid height of the channel and also the odor port was fixed at around 13 cm from the mouse nose (Figure 4.2). The head of the mouse was surgical implanted [77] and its olfactory bulb activity was monitored by an electronic microscope that live records wide-field calcium images. In short we were able to measure odor responsivity of the glomerular cells, with an accuracy of an individual glomerulus. A sensor was placed near the right mouse nostril to measure the odor signal in time, we combined calcium imaging with odor tracing in awake mice and we observed how they process natural olfactory scenes.

4.2. Monitoring neural response with calcium imaging in different flow conditions




Fig. 4.2 Graphic detailing experimental setup. Air flows from left to right and transport odor to a head-fix mouse, setup for in-vivo recording experiments. The odor port is located ~ 13 cm upwind of the animal's nose. Ethanol odor concentration is measured using a modified, commercially available ethanol sensor placed ~ 4 mm from the outer edge of the mouse's nostril. Figure adapted from [B].

4.2.1 Odor, sensor and flow characterization

A recording session consisted of 40 trials with different flow conditions, concentration dynamics varied stochastically and was unique in every trial. The flow velocity was controlled manually, a priori experimentalists identified 3 different regimes low, medium and high flow (Figure 4.3A) based on absolute velocity computed at channel mid-height (20 cm above the ground level) using an anemometer. Absolute velocities of low, medium and high flow were respectively 18 ± 8 , 66 ± 3 , and 74 ± 2 cm/s. That correspond to Reynolds number of $2,400 \pm 1,000$, $8,800 \pm 400$, and $9,800 \pm 200$, respectively (mean $\pm$ st. dev). We will quantify precisely the flow regime analyzing the odor dynamic in the next paragraphs.

Odor was released from the odor port for an interval of 10 s and a detailed protocol was followed to prevent the animals to habituate to a pattern of odor presentation, thus responding potentially to flow change rather than to the odor itself [B]. A single session consisted of 40 trials of odor presentation. Three different odors were used in the experiments: a solution of water and ethanol, a benzaldehyde-ethanol mixture and a isoamyl acetate- ethanol; the benzaldehyde-ethanol mixture was then chosen as standard to perform the different experimental sessions. Odor concentration and volume released from the odor port remained constant across all flow conditions, making the plume dynamics the only source of variation. Odor signal was recorded by a modified ethanol sensor (metal oxide sensor - MOX) located 3.5-4 mm from the mice's right nostril [78]. This system was miniaturized and its measurements do not affect the airflow, this was an important improvement if compared to commonly




Fig. 4.3 Differences in airflow conditions. (A) example odor traces are depicted for each flow condition. (B) PDF of the deconvolved odor concentration magnitude sampled across one low flow and one high flow examples. The distribution shift to right increasing the flow velocity, in low flow condition odor distribution is symmetric. Figure adapted from [B].

used photo-ionization detectors (PID), that are active sensors and their sampling strongly affects plume dynamics. To validate and optimize the measurements of the ethanol sensor we compared its performance to a standard PID, results are presented here below.

It has been proved that MOX sensors like the one we used are able to capture the turbulent dynamics of an odor plume despite their slower recording dynamics [78, 79]. Sensor signal was acquired at 100 Hz and then low pass filtered at 30 Hz using a Kaiser window (Matlab function). The signal, s , was then normalized within each trial subtracting the mean and dividing by the standard deviation of the signal during the plume ($s_n = (s - \langle s \rangle) / \sigma_s$). To account for the response dynamics of the sensor we used a standard deconvolution procedure [78], in which the raw signal is convolved with an exponential kernel:

$$k_0(t) = e^{-t/\tau_{\text{decay}}} - e^{-t/\tau_{\text{rise}}}, \quad k(t) = \frac{k_0(t)}{\int_0^T k_0(\theta) d\theta}$$

where t was evenly discretized at the proper sampling rate for the length of a single trial (T), τ_{decay} and τ_{rise} are free parameters. s and k , are transformed into Fourier space using the Matlab Fourier transform function, and the ethanol signal was deconvolved in Fourier space by dividing $\hat{s}$ by $\hat{k}$. The inverse Fourier transform of the resulting deconvolution was taken and the deconvolved signal was $d = F^{-1}(\hat{s}/\hat{k})$.

4.2. Monitoring neural response with calcium imaging in different flow conditions

The deconvolved signal d was then normalized within each trial $d_n = (d - \langle d \rangle) / \sigma_d$. To optimize the deconvolution parameters for the ethanol we recorded a complete session of 40 trials with different airflow, adding both the ethanol sensor and a PID in the wind tunnel (no mice were present during this session). The PID was placed 4 mm from the ethanol sensor at the same location where the animal was usually head-fixed. To optimize parameters, the PID signal, p , was first downsampled to 100 Hz to match the sensor sampling rate. Next, the signal was normalized within each trial $p_n = (p - \langle p \rangle) / \sigma_p$. This normalized signal was then Fourier transformed and convolved with the kernel and back-transformed to obtain the convolved signal $c = F^{-1}(\hat{p} \cdot \hat{k})$. It was then normalized ($c_n = (c - \langle c \rangle) / \sigma_c$) and compared to the ethanol signal across a range of τ_{decay} and τ_{rise} parameter values. The Kernel parameters were chosen by minimizing the mean squared error between the s_n and c_n signals averaged across all trials within the paired recording session: $\min_{\tau_{\text{rise}}, \tau_{\text{decay}}} \|s_n - c_n\|^2$. The best parameters are $\tau_{\text{decay}} = 0.4629$, and $\tau_{\text{rise}} = 0.0001$ and this optimized kernel is used to deconvolve the raw ethanol signal in the experimental recording sessions. The deconvolved signal (see Figure 4.3 A) is significantly correlated with the PID signal (Figure 4.4A) as measured during plume presentations ($r_{d-p} = 0.61$, $p < 0.001$), which is a 0.22 improvement from the correlation between the raw ethanol sensor and PID signal ($r_{s-p} = 0.39$, $p < 0.001$), correlation coefficients and p-values were computed with *corrcoef* Matlab function. The deconvolution preserves the odor concentration dynamics across trials, but does not preserve the absolute value of odor concentration. Note that the deconvolved trace was downsampled from 100 Hz to 30 Hz for figures and analyses to match the calcium imaging trace by averaging all samples taken across each camera frame. An initial inflection of signal at plume onset can be observed in the deconvolved ethanol signal for some trials (Figure 4.3A). This peak at plume onset is not reported by the raw sensor signal or by the PID signal and is likely an artifact of the deconvolution. Since these experiments focus on how well mitral-tufted activity follows odor concentration dynamics during plume encounters, the first and last seconds of the 10 s plume were omitted when analyzing neural responses to plume dynamics, the time interval considered lasted then 8 s. The only exception is for the analysis of responsivity, which is based on the percentage of time for which a significant response is observed and so this thresholded measure does not directly consider signal magnitude. Therefore, any artifact of plume onset dynamics in the sensor signal due to the deconvolution of its slower dynamics do not




Fig. 4.4 (A) Comparison between the deconvolved ethanol signal and the PID signal during a low and a high flow trial. (B) Skewness and (C) asymmetry of the deconvolved ethanol signal vs the PID for each trial in different flow conditions. All points lie close to the bisector (purple line, labeled "exact") showing the deconvolution preserves the statistics of the raw signal. We can notice that low flow trials cluster in the area of skewness = 0, while for medium and high flow trials skewness is significantly higher than 0. Asymmetry is null for low flow and it is negative for high flow, it means odor concentration PDF is strongly peaked on the right (see Figure 4.3B). Figure adapted from [B].

affect correlations reported between stimulus and response.

Airflow is manually tuned and the timing of low/high flow switch impacted differently on single trial. Every trial has its own specific dynamics that follows a realization of the turbulent flow. In particular we collected information about odor intermittency measuring two different quantities: the skewness of the odor concentration (it corresponds to the 3rd order moment) and the asymmetry ($p(d_n > \bar{d}_n) - p(d_n < \bar{d}_n)$, Figure 4.4B-C). Both observables quantify how far is the signal from a symmetric distribution: in absence of fluctuations we expect the odor concentration to distribute symmetrically, when the fluctuations increase the peak of the distribution shifts to the right [80, 38].

4.2. Monitoring neural response with calcium imaging in different flow conditions

Low and high flow trials were separable using either of these measures. Intermittency increased with airflow moving from low to high flow conditions. To verify if flow intensity is responsible for the variation of skewness and asymmetry, a one-way analysis of variance (ANOVA) test was conducted for each parameter. We divided the variance of skewness/asymmetry among different flow conditions by the sum of the variances within each flow condition and obtained that there is a significant difference among flow condition for both the variables (we used the *anovan* and *multcompar* Matlab functions). This shows that variations in skewness and asymmetry well describe different flow regimes. In conclusion there was a significant difference between low and medium flow and a significant difference between low and high flow. No significant difference was found between medium and high flow. Therefore, only low and high flow conditions were selected when examining the effect of air flow on neural parameters.

4.2.2 Calcium Imaging

Fluorescent response of mitral and tufted cells in glomeruli was measured with calcium imaging (Figure 4.5A), regions of interests were manually selected by the experimentalist. Calcium imaging is a microscopy technique that measures the level of calcium in a cell or in a tissue through fluorescence: experimentalist exposed transgenic mice, whose calcium ions are labeled with fluorescent molecules, to 488 nm LED stimulation, images were automatically acquired by a camera at a fixed sampling rate (30 Hz). Acquired images were processed with a constrained non-negative matrix factorization algorithm, this algorithm factorizes the pixel matrix in a product of matrices with non-zero elements: this procedure helps to isolate single glomeruli, to monitor the global activity of mitral and tufted cells and to remove the correlated signal among neighboring glomeruli [B]. Glomerular activity was baseline normalized using the mean and standard deviation of a 5 s activity period prior to stimulus onset, the signal was deconvolved to recover the average activity rate of each glomerulus according to Stern et al. [81].




Fig. 4.5 (A) In vivo recording of glomerular activity measuring the variation in fluorescent response of mitral tufted cells in the glomeruli. (B) Population response of mitral and tufted cells in the olfactory bulb to odorant exposure. Simultaneously recording ethanol plume (top) and imaging of calcium signals from glomerular activity (bottom). Figure adapted from [B].

4.3 Results

4.3.1 Qualitative evidence that brain responds to olfactory stimuli

Now we have defined how we measured odor concentration and glomerular response we can evaluate how mouse brain respond to odor stimuli. As depicted in Figure 4.5B, the response of glomerular population can be more or less intense based on the variation of the ethanol stimulus.

We can calculate the mean response of many individual glomeruli during each recording session. This is an informative quantity about glomerular activity but its relation with the olfactory stimulus has to be clarified. In Figure 4.6A we report the calcium response of a single glomerulus (glomerulus 26) to different flow trials, while in Figure 4.6B we represented the sum of the mean responses for each glomerulus in low and high flow. We can notice that on average glomeruli were more active during low flow and that some glomeruli are in general more active than others. In low flow the dynamics is less intermittent but the mean odor concentration is twice the mean concentration in high flow condition (as showed by PID recordings, $z =$




Fig. 4.6 (A) The deconvolved calcium response of a single glomerulus (number 26) to low (gray) and high (black) flow trials across the recording session shows glomerular responses vary due to the unique odor concentration dynamics of each plume. (B) A sum of the mean responses for each glomerulus in low and high flow, are plotted as a stacked bar graph so that comparisons between mean responses can be made within and across glomeruli simultaneously. Mean responses are calculated for the deconvolution within each flow condition during the plume release and vary significantly between conditions [$t = 11.43$, $p < 0.001$, $ttest$ function in Matlab] with higher average responses in low flow. Figure adapted from [B].

6.65, $p < 0.001$, Matlab function $ztest$) and this could explain why activity is much higher when velocity is lower.

4.3.2 Correlation among odorant plume dynamics and glomerular activity

Output from calcium imaging is difficult to interpret *per se*, to understand how mouse brain records and processes olfactory signals, we calculated the correlation coefficients between the odor stimulus and the calcium imaging response. Mean value was subtracted to both deconvolved ethanol and calcium signal, then the two signals were cross-correlated using Matlab $xcorr()$ function within every single trial. The mean coefficient for each glomerulus is calculated by averaging across all trials within the session for that glomerulus. Within flow cross-correlations were calculated with the same method but averaged only across trials within the specified flow condition. To control that the cross-correlations are not random we compute for each glomerulus the cross-correlation of ethanol signal during a trial with the calcium




Fig. 4.7 Glomerular population tracks odor concentration dynamics. (A) Simultaneously recorded signals shown for two example glomeruli responding to odor in a high flow trial. The glomerulus on the left is failing at tracking the odor signal, while the glomerulus on the right is successful. (B) Left, the cross-correlation between the deconvolved ethanol and each glomerulus's deconvolved activity, calculated within each trial and then averaged across trials. Each row is a glomeruli and each time point represents the cross-correlation averaged value at the indicated lag. Glomeruli are sorted in order of decreasing magnitude of correlation coefficient. Right, same as left but glomerular responses are trial shuffled so that the glomerular response is compared to ethanol signal from different trials. Glomeruli are sorted to match their corresponding unshuffled cross-correlation in the right panel. (C) The cumulative correlation, a sum of correlation coefficients for each glomerulus in low (light blue) and high flow (dark blue), are plotted as a stacked bar graph so that comparisons between mean responses can be made within and across glomeruli simultaneously. The cumulative plotting shows variation in ability to detect changes in odor concentration dynamics across glomeruli both within and across flow conditions. On average, a glomerulus's tracking ability varies significantly between conditions [$t = 12.81$, $p < 0.001$, *ttest* function in Matlab], with most glomeruli having stronger correlation coefficients in high flow trials. Glomeruli that significantly correlate with plume dynamics in at least one condition are plotted in blues while those that do not are plotted in grays. Figure adapted from [B].

response in all the other trials, the obtained result is the baseline to distinguish what is random from what is relevant. The difference between the matched and shuffled coefficients suggests correlations are not solely a result of plume structure, but are driven by the temporal dynamics unique to each trial (see Figure 4.7B).

Comparison of correlation coefficients in the matched vs shuffled cross-correlations does not naturally divide the glomeruli into two sub-populations, but rather the strength of this relationship varies continuously across glomeruli. Therefore, instead of dividing glomeruli into sub-populations of tracking vs non-tracking Figure 4.7A, our analyses consider how the strength of odor concentration tracking compares to other properties of the glomerulus and its response. Cross-correlations show a relation between glomerular and ethanol signals during odor presentation with all glomeruli having significant correlation with the plume during odor presentation as compared to their respective null distributions from trial shuffled correlation analyses. Interestingly, the degree of dynamic tracking is moderated by plume dynamics and it becomes stronger on average during plumes with higher levels of intermittency: glomeruli are tracking more when the signal is sparse Figure 4.7C. Correlation coefficients (tracking) increased from the null expectations by 0.08 ± 0.07 within low flow, and 0.16 ± 0.11 within high flow, glomeruli were significantly better at tracking plume dynamics in high flow than they were in low [$t = 12.81$, $p < 0.001$, Matlab function *ttest*].

4.3.3 Responsivity and response power analysis

To confirm that plume dynamics structure mitral-tufted cell population activity we quantified glomerular responsivity and response power. Responsivity is defined as the proportion of trials to which a glomerulus responded to the odor, it is a thresholded measure that determines if a glomerulus is more active than expected by chance when exposed to a plume, but responsivity is not informative about a glomerulus tracking ability. To measure responsivity we split the glomerular trace in baseline activity (before odor injection) and odor response. The signal was first normalized by subtracting the mean and dividing by the standard deviation of the baseline activity within each trial. Next it is binarized, thresholding for time points where activity exceeded the 95% confidence interval of the glomerulus's baseline activity. Within each trial, if the number of events exceeded the null expectation (5% of the total number of time points during plume presentation rounded up to the nearest integer), the glomerulus was considered to be responsive during that trial. Glomeruli had significantly higher responsivity during low flow trials ($t = 12.1$, $p < 0.001$, Matlab *ttest* function) and higher correlation in high flow trials. However, within each class, glomeruli with higher responsivity to the plume were better at following changes




Fig. 4.8 Glomeruli that respond more reliably to plumes are more correlated with their dynamics. (A) Responsivity is plotted against tracking ability (cross-correlation values) for each glomerulus (every circle is a single glomerulus), separating low flow (light blue) and high flow (dark blue) conditions. To represent the population response, the average responsivity across all glomeruli (low flow average = yellow dot, high flow average = red dot) is plotted against average correlation with plume dynamics. Across glomeruli, responsivity is positively correlated with tracking ability as is illustrated by the lines of best fit. On average, higher flow predicts a decrease in average responsivity level but also predicts an increase in tracking ability. (B) Glomeruli are sorted by increasing tracking ability (from left to right) and we show the average responsivity for each glomerulus in low and high flow. The change in mean responsivity between flows is plotted for each glomerulus (red line). Glomeruli with higher response reliability are also the ones more sensitive to plume dynamics. Figure adapted from [B].

in odor concentration ($r = 0.76$, $p < 0.001$). Thus, the more reliably a glomerulus responded to plume dynamics, the more likely it was to better follow changes in odor concentration (Figure 4.8). This is not a perfect relationship as glomeruli that are responsive to the plume but not to its dynamics exist, but a glomerulus with higher responsivity is more likely to be correlated to plume dynamics than one with lower responsivity.

To measure the activity of neurons it is standard in Neuroscience to compute the response power in Fourier spectrum $P = \int_0^{5\text{Hz}} \|\hat{C}(f)\|^2 df$; to monitor glomerular dynamics we computed the Fast Fourier transform (FFT) in the range 0–5 Hz between baseline and odor-on intervals. We choose this interval because the majority of cumulative response power for both the ethanol (on average $86.7 \pm 4.9\%$) and the calcium signal ($88.5 \pm 7.9\%$) was within 0–5 Hz. Across all the recorded glomeruli, response power was not varying significantly between flow conditions (see red and yellow dots in Figure 4.9A). To examine the effect of flow conditions on the response power of cells that most strongly responded to the odor, we next analyzed only glomeruli

whose mean response power was above the 75th percentile (red and yellow border dots in Figure 4.9A). Under this condition, response power did change significantly between flows [$t = 5.52$, $p < 0.001$], with stronger response power during high flow conditions. Response power was correlated with how well a glomerulus follows changes in odor concentration, when averaged across all trials glomeruli with higher response power were significantly better at following plume dynamics ($r = 0.74$, $p < 0.001$) (Figure 4.9B). Thus, a glomerulus's response power predicts its ability to track plume dynamics and for stronger responders, this relationship is moderated by changes in dynamic regimes.

Results from responsivity and response power suggest that both the reliability and the temporal pattern of mitral-tufted cells activity is significantly moderated by odor concentration dynamics. Thus, the spatio-temporal dynamics of plumes play a role in structuring activity in the first olfactory relay of the mouse's brain during natural olfactory processing. Thus, a large fraction of the glomerular population follows fluctuations during plume encounters, and the degree of dynamic tracking is moderated by plume dynamics, becoming stronger on average during plumes with higher levels of intermittency. As noted previously, the average correlation between plume dynamics and mitral-tufted cells activity increased in high flow conditions, so although glomerular responses became less reliable as airflow increased, they became more correlated with plume dynamics (Figure 4.7).

4.4 Conclusions

Mice evolved to track odor plumes despite the stochasticity and complexity due to the sparse nature of the signal. As discussed in the previous chapter spatio-temporal cues present in natural odor scenes drive decision-making in olfactory search [5, 82], but how they moderate population activity in the olfactory bulb is unknown. The experiments we presented here are a first step in investigating the relationship between odor features and neural representation: intensity of the odor concentration seems to elicit a stronger glomerular activity but it is the timing of the odor signal to better correlate with the neural response when intermittency becomes a relevant factor. New experiments with different air velocity and the possibility of varying the distance between the head-fixed mouse and the odor port may bring additional




Fig. 4.9 Higher magnitude of glomerular response power (0–5 Hz) is associated with higher correlation with plume dynamics. (A) Response power 0–5 Hz spectrum is plotted against tracking ability (cross-correlation values) for each glomerulus (every circle is a single glomerulus), separating low flow (light blue) and high flow (dark blue) conditions. Glomeruli with stronger tracking have a greater increase in response power during plume presentations ($r = 0.74$, $p < 0.001$). When calculated within flow, this relationship is significant within high flow ($r = 0.73$, $p < 0.001$), but not within low flow ($r = 0.19$, $p = 0.05$). The average response across all glomeruli is plotted (low flow average = yellow dot, high flow average = red dot) to represent the population response. Mean response power of the glomerular population is not significantly different between low and high flow, but it becomes significant when calculated with glomeruli whose mean activity is in the 75th percentile. (B) Glomeruli are sorted by increasing tracking ability (from left to right) and we show the response power 0–5 Hz for each glomerulus in low and high flow. The change in mean power response between flows is plotted for each glomerulus (red line). As tracking ability increases, so does the change in response power between flow conditions. This is consistent with the significant change in mean response between flow condition observed in the 75th percentile, plotted as red/yellow circles in (A). Figure adapted from [B].

information. Inspired by the theoretical results from Chapter 3 it is possible to design new experiments to establish what features of the fluctuating odor plume are encoded in the mouse neural activity. Particularly, we could test the prediction that timing of odor fluctuations are particularly useful at larger distances from an odor source. Interestingly, the fact that correlations are higher in high flow suggests mitral-tufted activity may be more responsive to whiff and blank features as opposed to tracking fine fluctuations in odor concentration across more constant plume encounters.

Another possible direction for future work is to monitor the single mitral-tufted cells instead of individual glomeruli: the odor tracking could be an emergent collective activity of heterogeneously tuned cells. Future research across a variety of odor concentration dynamic regimes and odor mixtures at both the cellular and population level are needed to further investigate the degree to which bulbar responses are tuned

to features of odor concentration dynamics and how this tuning may impact optimal encoding of odor information. Sniff frequencies are known to influence bulbar oscillations, and thus if sniffing behavior varied significantly between flow conditions, this may have contributed to some of the observed differences in tracking behavior between low and high flow trials. A recent work [71] showed that mice discriminate temporal correlations of rapidly fluctuating odours at frequencies of up to 40 Hz and that the mammalian olfactory system can track unexpectedly fast temporal features in odour stimuli and this could be relevant not only for different sources discrimination but as effective mechanism to navigate across turbulence. Despite the importance of sub-sniff sampling we believe it is important to monitor sniff frequency in the future to observe how it affects inter-sniff and intra-sniff activity, and consequently how these changes relate to tracking behavior observed in mitral-tufted populations across flow conditions.

Chapter 5

Navigation in turbulent odor plumes

I treated extensively the transport of passive scalar in Chapter 2, it can be dominated by diffusion or by advection and the transported quantity does not interact with the flow. In Chapter 3, I showed that inference of the odor source location from a distant position is possible and I presented two classes of features that would be convenient to measure during olfactory searches. Finally, in this Chapter, I will present the behavior of an agent actively searching for an odor source in a flow: the agent can measure odor in time and it will orient and move depending on cues it can detect. Navigation problems can be solved with many different strategies depending on the scale of the search and on the sparsity of the information present in the flow. I will present algorithms that were proposed to understand and reproduce navigation by different organisms and introduce an original framework I developed to explain air/ground sniffing alternation observed in rodents and dogs. The results presented in Section 3 of this chapter are collected in a publication [C], currently under revision at eLife.

5.1 Algorithms for olfactory navigation

Organisms sense the surrounding environment measuring visual, mechanical and chemical signals; in most animals the brain processes the acquired information and makes decisions relying on sensations. Small, brain-less, organisms have no central entity governing decisions and behavior, still they navigate successfully. Many factors constrain behavior of different organisms, e.g. the length and time scales

of odor detection, odor sparsity, the capability to alter the flow in the proximity of the sensors and the possibility of *a posteriori* elaboration of the recorded signal. All these constraints play an important role in determining the optimal strategy for successful navigation. Hence reproducing and understanding olfactory searches is challenging.

5.1.1 Gradient climbing strategy

In the Introduction I described chemo-tactic behavior in bacteria and I illustrated the olfactory organs of different species of animals, they present multiple sensor appendices that are able to measure gradients in odor concentration. When the odorant signal is smooth, navigating to the source is not very difficult and gradient descent is a successful algorithm. Even in turbulent conditions close enough to the source, the odor is continuous, although noisy, and gradient descent can still work, although less efficiently [59].

Gradient ascent/descent is an elementary algorithm in which an agent, at every time step, can measure a certain quantity in its proximity and it moves to a position where this quantity is higher/lower. Often some noise (η) is empirically added to the algorithm to avoid the agent to conclude its search in local maximum/minimum: $\mathbf{x}_{t+1} = \mathbf{x}_t - \gamma_t \nabla F(\mathbf{x}, t) + \eta$. Here $\mathbf{x}_t$ is the current agent position, γ_t is the length step that can vary in time and be either discretized or continuous. F is the cost function to be minimized, i.e. it can be the odor concentration $c(\mathbf{x}, t)$ and the gradient can be calculated in space or time [83]. Gradient descent convergence is warranted when the function F is convex, if F is stochastic convergence is likely when the signal can be split in a convex average plus some noise, rates of convergence depend on the details of the signal [59].

As long as olfactory signal is continuous and the sensors measure the odor plume, gradient climbing is an effective strategy both at scales of dozens of centimeters. This strategy can reproduce the biased random walk of bacteria and also navigation of animals in proximity to the source. Interestingly, single, small size cells cannot measure spatial gradients across their size and their strategy must rely on temporal gradients as described for *E.coli* in the Introduction. Larger eukaryotic cells and

animals with external appendices like antennae or nostrils can measure spatial gradients; furthermore organism with a developed neural system can exploit memory over varying temporal windows. Gradient climbing strategies have been implemented also in robots performing olfactory navigation, the optimality depends on the flux they are exposed to [63].

5.1.2 Bio-inspired algorithms for olfactory navigation

Larger organisms have developed sensors and more computational resources (neurons) but the scales they navigate are larger and odor fluctuates dramatically, because advection dominates over diffusion. The sparsity of information makes it difficult to locate the odor source: odor is detected rarely when range increases. When an olfactory search happens at large scales, especially in the presence of a flow, gradients are not smooth and odor is encountered in intermittent patches of different size as presented in Chapter 2. In unpredictable environments strategies diversify and become more complex: as exposed in Chapter 3, an agent can rely on multiple measures, not only odor concentration but also timing features like intermittency or whiffs and blank duration are useful as well as hints from the velocity field. Algorithms inspired by animal behavior during olfactory searches in turbulent flows have been designed since more than thirty years for both air and water navigation and here I will present the most known.

One of the earliest algorithms was inspired by male moths seeking for a partner: a male moth flies downwind when detecting female moth pheromones and it flies crosswind when it has no-detections. This behavior is made of two distinct phases called casting and surge: moving straight in the direction of the source is surging, while looking for the plume crosswind after missing the contact is casting. To emulate surging and casting behavior by moths and other flying insects treated in the Introduction a class of algorithm name surge-cast algorithms have been proposed. A standard algorithm [84] consists of three independent phases: after detection the agent surges with a variable delay, after losing the plume the agent casts for a certain time and it is assumed to locate the source when it arrives inside a certain radius r (close to the source also visual information is relevant). The central question is how does the agent choose the length and duration of both surges and casts? Different mostly phenomenological rules have been proposed. Parameters can be varied to

explore a range of different strategies. This strategy could be used by organism that possess an internal clock that regulates transition from a state to the other and the casting and surging would be a stereotyped behavior with odor detections affecting the transition rates from different states.

The so-called p-moth algorithm gives more importance to odor detections: it consists in an agent moving along straight paths, when the odor sensors measure a high concentration the agent moves in that direction till another prolonged whiff is sensed, if no detections arrive to the sensor the agent will move zigzag [85]. Shigaki et al [86] introduced a time varying modification to the p-moth algorithm, the duration of surging is determined by an equation, when the agent detects odor for the first time surge lasts longer and duration decreases with the following detections. They also proposed a navigation algorithm inspired by silkworm moth (they walk instead of flying) based on supervised learning: the moving agent records an odor time series that uses as input for a previously learned model, that produce as output the optimal displacement at every time step [18].

As I anticipated in the Introduction, moths and other insects are a paradigm for olfactory navigation in the air, crabs and lobsters play a similar role for underwater navigation. In [62] the authors proposed different algorithms inspired by lobster behavior: the standard routine is that the agent has two antennule, if one of the two senses odor then the agent moves diagonally in that direction, if both antennule measure odor the agent moves straight upstream, when no detection occurs then the agent moves crosswind picking a random direction. Novelty in this paper is that authors included a timing measure: the instantaneous time passed since the last detection. This quantity is used when no odor is measured, the agent instead of moving randomly will move in the direction from where odor was last recorded in the previous 5s. A further step is done including chemo-sensing by the legs, a lower 2-D odor field is included in the algorithm and combined with the sensing from the antennule. Relying on multiple inputs (antennule and legs) as well as multiple measurements (concentration and timing) improved the performance of the searching agent. The limit of this algorithm is that concentration is measured instantaneously and only the timing quantity has a memory-like role: so it is unclear if the improvement observed by the combination of concentration and timing is due to the timing or to the memory. Interestingly agents able to track intermittency locate




Fig. 5.1 Left: trajectories of an agent, starting from different positions, navigating to the odor source moving in the direction the odor comes from. Right: trajectories of an agent moving upstream the flow whenever it encounters the odor signal (standard rheotaxis). Figure adapted from [89].

the source navigating along the edge of the turbulent plume instead of following the centerline [61] as it has been documented for blue crabs [87] and cockroaches [88].

Two recent papers by the Emonet lab [21, 89] presented interesting experimental and algorithmical results on walking fruit flies. The first illustrates the importance of pausing to acquire information and avoid energy waste; the second explores the common assumption that insects turn upwind when they encounter odor signal. Kadakia et al. show that fruit flies turn specifically in the direction where the odor comes from, so the odor and not the flow plays the dominant role. They compared trajectories of these slightly different algorithms (see Figure 5.1): flow tracking algorithms are very effective when the agent is near the centerline, but they are not optimal when the searcher exits the plume; odor direction tracking allows the agent not to lose the plume.

Other bio-inspired algorithms like the beetle-inspired method have been developed to reproduce animal behavior both in laminar and turbulent flows, especially for olfactory driven robot. An extensive review can be found in [90].

In a more elaborated algorithm the agent updates its estimate of plume's centerline position following Bayes rule (we will illustrate this probability rule at the end of the paragraph), the more it is confident to be close to the centerline the more it surges upwind. The authors include a finite memory, so that the estimate of centerline location is averaged on this temporal window [91]. This algorithm is included in a class called probabilistic algorithms, which are engineering-based

methods that construct and update iteratively a source probability map based on detection or no-detection. We will treat some of these methods in the next section keeping in mind that, compared to bio-inspired algorithms, they require large computational power. One limitation of bio-inspired algorithms is that they are always responding in the same way to the stimuli and they fail at generalizing to different environments and flows. Moreover, the rules to start or end casting and surging are largely phenomenological.

Interestingly, when the navigation involves complex organism, memory and adaptation have an important role as shown by Gire et al.[59]: after being trained for few days mice stop relying on olfactory cues and use their memory to find food locations in the shortest time. It is likely that experience and evolution play an important role in animal sensing, treating olfactory navigation as a learning problem seems reasonable because the searcher will set the navigational strategy based on some a priori knowledge stored in its brain.

5.1.3 Infotaxis

Gradients become useless at large scales. A successful algorithm in these sparse conditions **have** been proposed by Vergassola et al. [92]. Instead of looking for the maximum of the odorant concentration, *infotaxis* aims at maximizing the information gain (or minimize the entropy). At every step the algorithm updates the estimate ($\mathbf{b}$) of position of the target relative to the agent x , the searcher observes o (detection or no-detection) and moves to maximize the information gain; the searcher has a model $\Pr(o|x)$ of the odor transport. The policy for infotaxis is:

$$\pi_{\text{infotaxis}}(\mathbf{b}) = \arg \max_a \left\{ \sum_{o,x} \Pr(o|x, a) \mathbf{b}(x) [H(\mathbf{b}) - H(\mathbf{b}^{o,a})] \right\},$$

where $H(\mathbf{b}) = -\sum_{\mathbf{b}_i} P(\mathbf{b}_i) \log P(\mathbf{b}_i)$ is the Shannon entropy of $\mathbf{b}$. The search is guided by minimizing the entropy if the searcher is nearby the source or if it expects to receive odor cues because it explored efficiently. Infotaxis is an effective algorithm for a broad spectrum of phenomena where information is limiting [93, 94] including olfactory searches in sparse regimes and it elucidates the well-known trade off between exploration and exploitation, it offers a heuristic description that does not guarantee convergence for navigation problems. Infotaxis is a model based algorithm,

the agent knows the odor statistics but it does not know its location relative to the source. At every step the agent makes an observation which is binary: detection or no-detection. Based on successful/unsuccessful observation the agent's belief about its location is updated according to Bayes rule. While infotaxis works in practice, it remains a heuristic algorithm, and its connection to maximizing reward is not immediately clear. Infotaxis tells us that gathering information about source location is an effective strategy for navigating in a sparse environment, but how information is defined profoundly matters. In nature, organisms use multiple senses to extract information and they are familiar with different environmental conditions: this suggests that different algorithms may emerge to cope with different situations. To generalize infotaxis to other sensory modalities and to explore the adaptation of an organism to different conditions, a model that computes information about target location must be defined.

5.1.4 Reinforcement Learning

Reinforcement Learning (RL) is a branch of machine learning that generalizes the classical frameworks of dynamic programming and optimal control. In RL, agents learn how to optimize their actions by repeated trials and errors in order to maximize a cumulative future reward. RL became popular in recent years for astonishing results in robot control or in beating human players at go, chess and recently Texas hold'em. In Reinforcement Learning, the agent interacts with the environment and receives a reward based on the action taken; the feedback obtained from the interaction agent-environment during training defines an optimal policy. The problem of olfactory search described above fits intuitively the RL framework: the agent/organism has to make a decision based on cues from the environment (detection or no detection) and it has a limited amount of time to reach its target. Reward is discounted exponentially in time, so that the agent is encouraged to find the target as fast as possible. To find the best policy, the agent needs to balance exploration and exploitation; the amount of exploration is tuned by parameters which make the agent more or less greedy and consequently determine a convergence rate. It is important to remember that most RL problems have no analytical solution and convergence is showed empirically *a posteriori*.

Reinforcement Learning explores the interaction of an agent with the environment, based on what the agent knows of itself and the surrounding world. We can

distinguish different classes of RL. When the agent bases its choices on a model and it also knows the state of the environment optimal policies are obtained by planning, this is the classical framework of optimal control and Markov Decision Process (MDP). In this case, the agent knows its own position and the position of a target and it can use the map to trace the shortest path. If the state of the environment is unknown (in the previous example, the relative position of the agent to the target is unknown) the problem can be treated by hidden Markov models which we will discuss below. Finally, when knowledge of both model and environment is missing, we have full RL.

5.2 Partially Observable Markov Decision Processes for olfactory searches

5.2.1 An introduction to POMDPs

A powerful category of algorithms for olfactory searches are Partially Observable Markov Decision Processes (POMDPs) [95]. Their theoretical description has been developed in the 60s, but only recent increased computational power and novel algorithms allowed to solve complex problems. To understand POMDPs it is pedagogical to start from a description of a Markov Decision Process (MDP), it enables modeling the decision-making of an active agent [96, 97]. In this framework we start by introducing a state space, an action space and a reward function. At each time, the system is in the state s , and the agent takes action a . As a consequence, the agent receives a reward r which depends on the current state and action, and the system transitions to a new state s' . The dynamics of the state space is Markovian (the probability to jump to the next state depends only on the current state) and is defined entirely by the transition matrix, $T(s'|s, a)$, the agent jumps to state s' depending on the action taken a and given the current state s . The agent obtains a reward $r(s, a, s')$ after each transition.

The goal of the process is to find the unique optimal policy $\Pi^*(s)$ which picks actions that maximize the expected sum of future rewards $\langle r_0 + \gamma r_1 + \gamma^2 r_2 + \dots \rangle$, where $\gamma \in (0, 1)$ is the discount factor (varying γ makes the agent more or less greedy, indeed $\gamma = 0$ corresponds to maximizing the immediate reward, whereas $\gamma \rightarrow 1$

fixes a very long time horizon), and r_t is the expected reward t steps after the initial state. Different algorithms allow policy optimization, many of these involve solving for the value function, $V(s)$, which is the expected discounted sum of rewards from state s , conditional on policy $\Pi^*(s)$. The value function satisfies the central dynamic programming equation known as the Bellman equation [98]:

$$V(s) = \max_a \left(\sum_{s'} T(s'|a, s) \{r(s, a, s') + \gamma V(s')\} \right). \quad (5.1)$$

If the system dynamics is known and we have access to the transition matrix T the Bellman equation can be directly iterated to find the value function, otherwise MDP can be repeatedly simulated to obtain transition matrices (i.e through Monte Carlo simulations) [97]. In MDPs the agent has complete knowledge of its current state, in POMDPs an additional layer of complexity is present. Knowledge of the current state is partial: the agent, after picking an action, makes an observation o . Observations are the only way the agent can access to a probabilistic representation of its current state, partial knowledge can be due to the noisy environment or be related to the low precision of the sensors. The agent is assumed to have an internal model $P(o|s, a)$, after observation it compares the outcome with its knowledge of the reality and it builds a belief. For olfactory search the agent knows the source location, assumes a known probability of detection which depends on distance to the source: the position of the agent relative to the source is unknown. POMDPs map a sequence of observations and actions $o_{-1}, a_{-1}, o_{-2}, a_{-2}, \dots$ to the best action. The entire story needs not to be stored in a high dimensional memory but is instead encoded by the current posterior distribution over states, b , also known as the *belief vector*. The problem of solving for the optimal action is re-interpreted as the problem of solving for the policy $\Pi^*(b)$. The Bellman equation on states for MDPs translates into a Bellman equation on belief vectors for POMDPs:

$$V(b) = \max_a \left(\sum_{s, s'} b(s) T(s'|a, s) \times \left\{ r(s, a, s') + \gamma \sum_o P(o|s', a) V(b^{a,o}) \right\} \right), \quad (5.2)$$

where $b^{a,o}$ is the posterior belief state given the agent takes action a and observes o . Using Bayes' rule, $b^{a,o}$ is given by

$$b^{a,o}(s') = \frac{P(o|s',a) \sum_s T(s'|s,a)b(s)}{P(o|b,a)}, \quad (5.3)$$

where the normalizing factor is

$$P(o|b,a) = \sum_{s'} b(s') \sum_s T(s'|s,a)P(o|s,a) \quad (5.4)$$

Intuitively, the Bayes' rule takes into account the new information gained from the most recent observation and the information lost due to the state space dynamics, which are in turn influenced by the action.

5.2.2 Algorithms to solve POMDPs

To solve a POMDP we need to find the most accurate approximation of the value function $V(b)$, for all b ; the optimal policy then follows from choosing the action that yields the highest future expected reward given the current belief vector :

$$\Pi^*(b) = \arg \max_a \sum_{s,s'} b(s) T(s'|a,s) \times \left\{ r(s,a,s') + \gamma \sum_o P(o|s',a) V(b^{a,o}) \right\}. \quad (5.5)$$

The value function is on belief vectors that have dimension equal to the number of states and whose components take continuous values between 0 and 1. Computing the value function $V(b)$ exactly for all belief vectors for problems with a high number of states is clearly infeasible. Existing methods exploit a specific representation of the value function, which leads to the approximation discussed by [95].

In particular, it can be shown that the value function can be approximated arbitrarily well by a finite set H of hyperplanes [99], each of which is parameterized by $\alpha(s)$:

$$V(b) = \max_{\alpha \in H} \alpha \cdot b. \quad (5.6)$$

An initial set H is expanded using the Bellman equation (5.2). Using vector notation, we can write

$$V(b) = \max_a \left\{ r_a \cdot b + \gamma \sum_o P(o|b,a) V(b^{a,o}) \right\}, \quad (5.7)$$

where $r_a(s) \equiv \sum_{s'} T(s'|a, s)r(s, a, s')$. Let $\alpha^{a,o}(s)$ be defined as

$$\alpha^{a,o}(s) = \sum_{s'} \alpha(s') P(o|s', a) T(s'|s, a). \quad (5.8)$$

and using the belief update (5.3) it follows that

$$V(b) = \max_a \left\{ r_a \cdot b + \gamma \sum_o \max_{\alpha} b \cdot \alpha^{a,o} \right\}. \quad (5.9)$$

Given the previous set H , we can add a new α vector to it corresponding to belief vector b called the “backup” operation:

$$\text{backup}(H, b) = \arg \max_{\alpha_a^b} b \cdot \alpha_a^b, \quad (5.10)$$

$$\text{where } \alpha_a^b = r_a + \gamma \sum_o \arg \max_{\alpha^{a,o}} b \cdot \alpha^{a,o}. \quad (5.11)$$

In other words, given a previous set H and new belief vector, one can use the Bellman equation to update H and obtain a better approximation to the value function. The key computational advantage of using the above backup operation is that the $\alpha^{a,o}$ ’s can be pre-computed for the current H and re-used when backing up. The following issue to tackle is how to efficiently collect new belief vectors to update H and how to remove vectors from H that are no longer necessary. The difference between algorithms rises at these two stages. We use Perseus [100], which simulates random exploration of the agent. Specifically, at each step in a “training” episode, we start from an initial prior, pick actions uniform randomly and then sample observations from $P(o|b, a)$. The new belief vector obtained using Bayes’ rule is then used to backup H . Finally, after adding a new set of α vectors into H , it is efficient to prune the existing ones that are guaranteed to not be used. We prune the α vectors whose every component is smaller than those of another vector (see [95] for other heuristic pruning methods).

In this method three free parameters need to be tuned: the discount factor γ , the number of belief points sampled per each training episode (random exploration) and the total number of training episodes. The discount rate sets the planning horizon: a long term policy can lead to large rewards but it will also require long time and many computational resources to be learnt. In olfactory searches we set the discount rate so that the effective horizon $\frac{1}{1-\gamma}$ is of the same order as the typical number of steps

to get to the target. Increasing the other two parameters allows a higher accuracy in training and potentially a more effective policy, but the price to pay is an increase in training time: we defined optimal values for these parameters by selecting the shortest training that led to a consistent improvement in performance.

5.2.3 POMDPs allow navigation in a turbulent environment

First we design a POMDP to solve a simplified olfactory search, where we allow the agent to detect odor only in the air with a constant probability rate, while no detection is possible at ground level. In the next Section we will tackle a more realistic framework with odor statistics taken from direct numerical simulations described in Chapter 2.

Successful navigation is possible if the agent encounters enough detections on its trajectory. It moves along a trajectory $r_1, r_2, \dots, r_t$ while measuring the odor signal $o_1, o_2, \dots, o_t$. At each time step the agent choose among six actions: staying at the same location, moving in one of the four neighboring locations, or staying at the same location and sniffing in the air. Sniffing in the air leads to the only non trivial observation o_t that updates agent belief b_t , according to Bayes rule (5.3). The belief vector contains all the information the agent gathered up to time t about its relative distance from the source. Sniffing in the air has a cost, due to the fact that the agent has to pause to perform it; pausing slows down the search, and delays the unit reward upon finding the source, which drives navigation. Rewards are discounted at a rate λ , i.e., the expected long-term reward is $\langle e^{-\lambda T} \rangle_T$. T is the time taken to find the source and the expectation is over the prior knowledge available to the agent, its navigational strategy and the statistics of odor encounters. The value function $V(b_t)$ for the current belief b_t is calculated using Bellman's equation (5.2), which takes into account all possible future trajectories of an optimal agent.

$$V(b_t) = \max_a \left\{ \Gamma_a + \gamma(1 - \Gamma_a) \sum_{o_{t+1}} P(o_{t+1}|b_t, a) V(b_{t+1}) \right\}, \quad (5.12)$$

where Γ_a is the probability of finding the source immediately after taking action a ; $\gamma \equiv e^{-\lambda t_s}$, where t_s is the duration of the time step, and the probability $P(o_{t+1}|b_t, a)$ of observing o_{t+1} is determined by the physical environment and the signal detection threshold of the agent. Intuitively, the terms in the argument of the max function



Fig. 5.2 The probability per unit step of detecting an odor signal in the air in this simple model is non-zero just in a rectangle $x_{thr} \times y_{thr}$ where it has value $k_{prob} = 0.4$, the environment is defined by a much larger box.

in (5.12) represent the value of finding the source, detecting the odor signal or not detecting the odor signal, each event being weighted by the probability for the six possible actions. The optimal action is the one that maximizes the value, i.e., the parenthesis on the right-hand side of (5.12).

For simplicity, we discretize observations into detections and non-detections, which implies that the behavior depends solely on the probability per unit time of detecting the odor in the air. The agent likelihood model is here a constant probability of detection defined over a rectangle of dimensions $x_{thr} \times y_{thr}$ ($7.5m \times 0.75m$). Thus, the agent uses a Poissonian detection model, whose maps are an idealized model of an odor plume.

To solve (5.12) we use Perseus algorithm as explained in the previous Section. For each trial run, we begin with a uniform prior distribution ~ 4 times larger than the agent likelihood and we simulate the POMDP until the agent finds the source. If the agent does not find the source within 1000 steps, we interrupt the simulation. The time step, t_s , and the distance traveled at each step are set such that the agent sniffs three times per second and at every step it moves 12 cm.

In the representative example depicted in Figure 5.3, the agent begins with a uniform prior belief, much larger than the constant model, as shown in the top row of Figure 5.3. The agent makes its first action by sniffing in the air and does not detect an odor signal. Since odor is not detected, the likelihood that the agent is immediately downwind of the source is reduced, which leads to a posterior belief updated via Bayes' rule (second row, Figure 5.3). The agent proceeds by casting crosswind while occasionally pausing to sniff in the air (third row, Figure 5.3), after which it executes an upwind surge (fourth row, Figure 5.3). The decision to surge at that specific moment can be understood from examining the belief immediately before the surge: because the agent did not detect any odor over the entire cast-and-sniff sequence,

5.2. Partially Observable Markov Decision Processes for olfactory searches




Fig. 5.3 Progression of the belief the agent has about its own position relative to the source. From top, first panel: before starting the search the agent has a flat belief about its own position, much broader than the plume in air. Second: belief after a single sniff in the air and no detection. The lighter grey region corresponds to the extent of the model in air and indicates that because the agent did not detect the odor, it now believes it is likely *not* right downstream of the source. Third: As the agent casts, its belief about its own position translates sideways with it; additionally, at each sniff in the air with no detection, the belief gets depleted right downstream of the source, as in the panel right above. As a result, the cast-and-sniff cycle sweeps away a region of the belief as wide as the cast and as long as the plume. Fourth: as the agent surges upwind, its belief about its own position translates forward with it. Fifth: after detection, the belief shrinks to a narrow region around the actual position of the agent, which leads to the final phase of the search within the plume. Green (Purple) wedges indicate that the entropy of the belief decreases (value of the belief increases) as the agent narrows down its possible positions (and approaches the source).

the likelihood that the agent is located near the source, i.e., within the plume, is low (third row, Figure 5.3). At this point, it is more valuable to surge upwind rather than continuing to explore the same area. By surging forward, the agent is now more likely to encounter the plume, which enables it to effectively explore the remaining part of the belief. The key to the above argument is that the agent lacks knowledge of its position relative to the source, and it acts so as to narrow down its belief.

A repetition of the sequence of casting, alternation and surging follows as the agent steadily narrows down the belief, until it finally detects the odor (bottom panel, Figure 5.3). The detection shrinks the posterior to the likelihood rectangle and makes entropy plummet and leads the agent straight to the source.

5.2.4 A theoretical model

We develop a theoretical model in which the agent only detects odor by sniffing in the air. This matches the POMDP search conditions and it is a simplified analysis of the more complex problem of alternation I will discuss next. Here, search path is parameterized by the discrete locations at which the agent sniffs in the air. The prior distribution, $b(x, y)$, of the agent's location with respect to the source is assumed uniform with length L_x ($\gg x_{\text{thr}}$) (along the downwind direction) and width L_y ($\gg y_{\text{thr}}$), as in the example shown in Figure 5.3 (top). The probability of detecting an odor signal in a sniff, $r(x, y)$, depends on the extent of the plume via the parameters x and y . In our situation, the likelihood is $r(x, y) = \chi(x < x_{\text{thr}}) \chi(|y| < y_{\text{thr}})$. For the more complex case discussed next, we decouple by approximating the detection probability map (Figure 5.7) as $r(x, y) = f(x)g(y)$, where $f(x)$ is a constant when $0 < x < x_{\text{thr}}$ and 0 otherwise, $g(y)$ has a characteristic length-scale y_{thr} .

To localize the plume, the agent has to sufficiently explore, by sniffing in the air, patches of size $\sim x_{\text{thr}} \times y_{\text{thr}}$ within its prior. Since the prior's width is larger than the plume width ($L_y \gg y_{\text{thr}}$), the agent has to cast in order to determine how far it is from the plume's center-line. Each sniff effectively explores a patch of length $\sim x_{\text{thr}}$ immediately downwind of the source.

Therefore, a bout of casting across a width L_y while constantly sniffing in the air explores a region of size $\sim x_{\text{thr}} \times L_y$. There, the likelihood of containing the source is strongly depleted, which converts the initial prior into a posterior of reduced length $L_x - x_{\text{thr}}$. Since the agent now believes to be outside of the plume, it is convenient to continue the search surging upwind by x_{thr} , and exploring a new patch via casting. The process is repeated until the plume is detected.

The search process can therefore be split into distinct episodes where the agent cycles between sniffing while casting and surging upwind by $\sim x_{\text{thr}}$. We identify three main questions about the search, which we address in more detail below: 1) how wide should the agent cast?; 2) how long should the agent spend casting

before surging upwind?; 3) where should the agent sniff during the casting phase? Specifically, we highlight and quantify the various trade-offs associated with the cast-sniff-surge modes of exploration.

Since the rate map is uniform in x and has length x_{thr} , the agent surges exactly a distance x_{thr} . The search process is then decomposed into $N \sim L_x/x_{\text{thr}}$ distinct episodes. In each episode n ($n = 1, \dots, N$), the agent spends a time t_n deciding whether the source is within reach, i.e., closer than x_{thr} . The casting duration t_n is to be optimized. After t_n , since the agent has determined that the source is not yet within reach, it surges upwind and continues to the next episode $n + 1$. The process continues until the agent obtains a detection. The cumulative probability of not detecting the signal (conditional on the target being in that patch) after casting for time t , $c(t)$, depends on the sampling strategy during casting and is discussed further below.

The expected discounted reward at the beginning of the search is $V_1 \equiv \langle e^{-\lambda T} \rangle_T$, where T is the time taken to find the odor signal. We use dynamic programming to compute and optimize V_1 . V_1 is the sum of the expected reward if the agent finds the signal in the first patch within time t_1 and the expected reward after moving to the next patch if it does not. The information gained from the observation of not detecting a signal is taken into account in the latter term through a Bayesian update of the prior. However, we show that V_1 and the casting times, $t_1, t_2, \dots, t_N$, can be calculated using an equivalent, simpler expression which does not require Bayesian updates (next section). Specifically, denote V_n as the expected discounted reward at the beginning of the n th episode, i.e., before the cast and surge. V_1 is calculated using the recursive equation

$$V_n = \max_t \left\{ -\frac{1}{N} \int_0^t c'(s) e^{-\lambda s} ds + e^{-\lambda(t + \frac{x_{\text{thr}}}{v})} V_{n+1} \right\}. \quad (5.13)$$

The time t that maximizes the parenthesis determines the optimal duration t_n the agent should spend casting before surging upwind. The first and second terms in the parenthesis of (5.13) are the expected discounted rewards if the agent detects a signal during casting (and the search ends) or if it does not detect a signal, surges a distance x_{thr} and continues to the next episode, respectively. The factor $-c'$ in the first term is the probability density to make a detection at time t conditional on the target being in the current patch, which has probability $1/N$.

We first show that the duration t_n obeys a marginality condition. The agent should stop casting when the value of continuing to explore the current patch is just outweighed by the value of moving on and exploring the next patch. This intuition is quantified by optimizing for t in (5.13). Zeroing the time derivative of (5.13), we obtain that t_n is the value of t that satisfies the equality $-c'e^{-\lambda t}/N = \lambda e^{-\lambda(t+\frac{x_{\text{thr}}}{v})}V_{n+1}$. The left hand side is the rate of value acquisition upon staying in the current patch. The right hand side is the negative rate of value acquisition upon delaying departure, that is the rate of value acquisition upon anticipating departure. Thus by maximizing value we obtain that, at optimality, the added value of continuing to cast matches the added value of anticipating surge, i.e., marginality of the two actions as prescribed by marginal value theory [101]. The marginality condition leads to a relationship between the casting time and the value at the next episode

$$-c'(t_n) = N\lambda e^{-\lambda x_{\text{thr}}/v}V_{n+1}. \quad (5.14)$$

When $n = N$, the agent casts indefinitely, which gives $V_N = -\frac{1}{N} \int_0^\infty c'(s)e^{-\lambda s}ds$ from (5.13). The casting time for each episode is obtained using this boundary condition, (5.14), and $c(t)$, which we shall determine in the next paragraph. Note that we have ignored the possibility that at $n = N$, the agent turns back and moves downwind to re-explore earlier regions, which can be incorporated into this framework and leads to a different boundary condition. However, we do not take this into account since this extension only marginally affects the earlier stages of the search path and does not affect general conclusions.

We now optimize for the sampling strategy during casting, which in turn determines $c(t)$. The casting phase can be formulated as a decision-making process of deciding where to sniff next on the crosswind axis given the marginal distribution $\tilde{b}(y) = \int_0^x dx b(x, y)$. The next sniff location at a displacement Δy from the current location is obtained from the dynamic programming equation similar to (5.13), which relates the current value to the value of moving and sampling elsewhere.

$$V(\tilde{b}) = \max_{\Delta y} \left\{ \left[\Gamma_{\tilde{b}}(\Delta y) + (1 - \Gamma_{\tilde{b}}(\Delta y))V(\tilde{b}') \right] \times e^{-\lambda(|\Delta y|/v + t_{\text{sniff}})} \right\}, \quad (5.15)$$

where $\tilde{b}'$ is the posterior after sampling at the new location conditional on no detection, and $\Gamma_{\tilde{b}}(\Delta y)$ is the probability of detection. The two terms in the Bellman equation correspond to the cases when the agent detects a signal and does not detect

a signal respectively, which are discounted in proportion to the time taken to travel a distance $|\Delta y|$ and sniff in the air. Numerically solving (5.15) yields a sampling strategy and the corresponding $c(t)$. The optimized casting strategy is a zigzag (Figure 5.4a) which expands over time to the width of the prior. The probability of not detecting the signal decays exponentially with a rate depending on the optimization depth (Figure 5.4b). In the low-detection rate limit, we generically expect a constant detection rate (say $c(t) = e^{-\kappa t}$), consistent with the exponential decay observed in the simulations. The detection rate κ decreases with t_{sniff} (Figure 5.4c), which in turn translates to a decreased value (from (5.14)) and highlights the cost of pausing to sniff in the air. From (5.14), we then have

$$t_n = \kappa^{-1} \log \left(\frac{\kappa e^{\lambda x_{\text{thr}}/v}}{N \lambda V_{n+1}} \right). \quad (5.16)$$

We use (5.13) and (5.16) along with the boundary condition $V_N = \frac{\kappa}{N(\kappa + \lambda)}$ to solve for the casting times. The results show increasing casting times with episode index (Figure 5.4d). Intuitively, as the search progresses, the marginal cost for the agent to continue casting decreases due to its increasing confidence that it is in the right patch, driving the agent to spend more time casting before leaving the patch.

We verify predictions from the theory using simulations of the simplified POMDP illustrated in the previous section. Simulations confirm that the surge length and cast width are equal to the detection range and the prior width respectively (Figure 5.4e,f).

Derivation of equation (5.13)

We consider a scenario where a target is located at one of N possible patches, $n = 1, 2, \dots, N$ with probabilities $p_0 = (p_1, p_2, \dots, p_N)$ ($\sum_n p_n = 1$). Note that $p_n = 1/N$ for all n for the prior considered in the previous section. The agent starts at $n = 1$ and moves sequentially from $n = 1$ to $n = N$ while spending time t_n sampling in each patch. Moving from a patch to the next one takes time $\tau \equiv x_{\text{thr}}/v$. At $n = N$, the agent samples indefinitely, $t_N = \infty$. The agent receives reward of one when the target is found in a patch, which is discounted at rate λ . The value $V_1 \equiv \langle e^{-\lambda T} \rangle_T$, where T is the search time, is the expected discounted reward optimized w.r.t t_n 's. We derive two sets of recursive equations (with and without Bayesian updates) to calculate V_1 .




Fig. 5.4 (a) Optimized sniff locations during casting (conditional on no detection) show a zigzag of increasing amplitude. At each decision, (5.15) is expanded and optimized w.r.t the subsequent $n_{\text{steps}} = 10$ sniff locations $\Delta y_1, \Delta y_2, \dots, \Delta y_{n_{\text{steps}}}$ using standard black-box optimization methods. The agent then moves by Δy_1 and the procedure is repeated. (b) The probability of not detecting the signal against time, $c(t)$, decays exponentially with detection rate, κ , shown here for different values of the optimization depth n_{steps} . κ saturates beyond $n_{\text{steps}} = 9$. (c) κ monotonically decreases with the time per sniff, t_{sniff} , reflecting the cost of pausing to sniff the air. In panels (a), (b) and (c), we use $y/L_y = 1/20$, $L_y = 1$, $\lambda = 0.5$, $v = 1$ and $t_{\text{sniff}} = 0$ (for (a) and (b)). (d) Casting times (in units of $1/\lambda$) generally increase as the search progresses. Obtained using (5.16) for different values of κ/λ (colored lines). Here $N = 6$ and $\lambda x_{\text{thr}}/v = 0.05$. (e,f) The surge length and cast width from simulations of a simplified POMDP, where the agent can detect an odor signal only by sniffing in the air. Results for different prior and plume dimensions (blue stars) align with the theoretical prediction (red line) that the surge length and cast width are equal to the detection range in air, x_{thr} , and the prior width, L_y , respectively. Adapted from [C].

We show that both formulations lead to the same optimal casting times, however, the set of equations without Bayesian updates are much simpler to compute.

Suppose the cumulative probability of finding the target in time t *conditional* on the target being in that patch is $d(t)$. Note that $c(t) \equiv 1 - d(t)$ is used in the main text. Denote $r(t) \equiv \int_0^t d'(s) e^{-\lambda s} ds$. This is the expected discounted reward if the agent searches for time t in a patch that contains the target.

Say $N = 3$. Since V_1 is the expected discounted reward optimized over the casting times t_1, t_2 , we have

$$\begin{aligned} V_1 &= \max_{t_1, t_2} \left\{ p_1 r(t_1) + e^{-\lambda(t_1+\tau)} p_2 r(t_2) + e^{-\lambda(t_1+t_2+2\tau)} p_3 r(\infty) \right\} \\ &= \max_{t_1, t_2} \left\{ p_1 r(t_1) + e^{-\lambda(t_1+\tau)} \left(p_2 r(t_2) + e^{-\lambda(t_2+\tau)} (p_3 r(\infty)) \right) \right\} \\ &= \max_{t_1} \left\{ p_1 r(t_1) + e^{-\lambda(t_1+\tau)} \times \max_{t_2} \left\{ p_2 r(t_2) + e^{-\lambda(t_2+\tau)} (p_3 r(\infty)) \right\} \right\} \end{aligned} \quad (5.17)$$

The last equation above motivates a recursive equation for general N :

$$V_n = \max_{t_n} \left\{ p_n r(t_n) + e^{-\lambda(t_n+\tau)} V_{n+1} \right\}, \quad (5.18)$$

with boundary condition, $V_N = p_N r(\infty)$. Optimizing over t_n , we obtain the marginal value condition

$$p_n r'(t_n) = \lambda e^{-\lambda(t_n+\tau)} V_{n+1}. \quad (5.19)$$

If the rate of detection during casting is a constant κ , we have $d(t) = 1 - e^{-\kappa t}$, $r(t) = \kappa \int_0^t e^{-(\kappa+\lambda)s} ds = \frac{\kappa}{\kappa+\lambda} \left(1 - e^{-(\lambda+\kappa)t} \right)$ and $r'(t) = \kappa e^{-(\kappa+\lambda)t}$. Plugging this expression into (5.19), we get

$$t_n = \kappa^{-1} \log \left(\frac{p_n \kappa e^{\lambda \tau}}{\lambda V_{n+1}} \right) \quad (5.20)$$

Now, let's calculate V_1 using Bayesian updates and show that the optimal times exactly correspond to what we have in the previous equation. Denote $\tilde{V}_n(q)$ as the value at patch n for an arbitrary probability vector $q = (q_1, q_2, \dots, q_n)$. We now show that $V_1 = \tilde{V}_1(p_0)$, where $p_0 = (p_1, p_2, \dots, p_N)$ is the prior. We have

$$\tilde{V}_n(q) = \max_{t_n} \left\{ q_n r(t_n) + e^{-\lambda(t_n+\tau)} (1 - q_n d(t_n)) \tilde{V}_{n+1}(q') \right\}, \quad (5.21)$$

where q' is the posterior conditional on no detection. The two terms on the r.h.s correspond to the case when the agent finds the target in the patch before t_n (with probability $q_n d(t_n)$) and does not find it (with probability $1 - q_n d(t_n)$) respectively. Given the observation that the target is not found in patch n , the posterior probabilities,

q' , are obtained using Bayes' rule:

$$\begin{aligned} q'_m &= \frac{q_m}{1 - q_n d(t_n)}, \quad \text{for } m \neq n, \\ q'_n &= \frac{q_n(1 - d(t_n))}{1 - q_n d(t_n)}. \end{aligned} \quad (5.22)$$

We show that $\tilde{V}_1(p_0) = V_1$ for $N = 3$. The general case of starting from any patch, prior and number of patches (N) follows. Expanding (5.21) starting from $n = 1$,

$$\begin{aligned} \tilde{V}_1(p_0) &= \max_{t_1, t_2} \left\{ p_1 r(t_1) + e^{-\lambda(t_1+\tau)} (1 - p_1 d(t_1)) \right. \\ &\quad \times \left. \left(p'_2 r(t_2) + e^{-\lambda(t_2+\tau)} (1 - p'_2 d(t_2)) p''_3 r(\infty) \right) \right\}, \end{aligned} \quad (5.23)$$

where p'_2 is obtained from the first Bayesian update and p''_3 is obtained after the second Bayesian update. Using (5.22), we have $p'_2 = p_2/(1 - p_1 d(t_1))$, $p'_3 = p_3/(1 - p_1 d(t_1))$ and $p''_3 = p'_3/(1 - p'_2 d(t_2)) = p_3/(1 - p_1 d(t_1) - p_2 d(t_2))$.

Since $p''_3(1 - p'_2 d(t_2)) = p'_3$, simplifying (5.23), we get

$$\begin{aligned} \tilde{V}_1(p) &= \max_{t_1, t_2} \left\{ p_1 r(t_1) + e^{-\lambda(t_1+\tau)} (1 - p_1 d(t_1)) \left(p'_2 r(t_2) + e^{-\lambda(t_2+\tau)} p'_3 r(\infty) \right) \right\}, \\ &= \max_{t_1, t_2} \left\{ p_1 r(t_1) + e^{-\lambda(t_1+\tau)} \left(p_2 r(t_2) + e^{-\lambda(t_2+\tau)} p_3 r(\infty) \right) \right\}, \end{aligned} \quad (5.24)$$

where $p'_2 = p_2/(1 - p_1 d(t_1))$, $p'_3 = p_3/(1 - p_1 d(t_1))$ are used in the second step. This equation exactly corresponds to (5.17). The upshot is that the normalization factors from the Bayesian updates go through the parenthesis and cancel out. However, optimizing for t_n directly using (5.21) is difficult due to the dependence of q' on t_n .

5.2.5 Parameters used in the POMDP

The main figures represent results using: discount factor $\gamma = 0.99$, number of training episodes $i = 320$, number of belief points sampled per training episode $i' = 100$, likelihood in the air and at the ground is defined as shown in Figure 5.9, in Figure 5.2 likelihood in the air is defined as a rectangle with dimensions $x_{\text{thr}} \times y_{\text{thr}}$. Results in

Figure 5.10 are tested and averaged over three different starting positions ($x = 8$; $y = 0, 0.3, -0.5$), 8 different seeds, 50 different realizations for the same seed (trajectories differ for the history of detections according to the Poissonian model).

5.3 Alternation between detection in air and at the ground

In the previous section, we proved that POMDPs are an effective algorithm for olfactory navigation. Next, we want to explore the behavior of the agent in a more realistic setting where sniffing is allowed both in the air and near the ground. Introducing the possibility of multi-modal sensing will elucidate a commonly observed alternating behavior: indeed it has been reported that, during olfactory navigation, different species like rodents [102, 59] and dogs [103–106] can follow odor traces above the ground, but occasionally pause and rear with their nose sniffing up in air. Thus during navigation to an odor source animals alternate between these two different sensorimotor modalities: sniffing the ground and sniffing in the air. The reasons underlying this alternation are largely unknown [16], in fact this behavior could be explained by the novelty detection and information acquisition but also by a state of fear and anxiety [107]. Gire et al. observed a more frequent alternation in the early stage of a laboratory odor-guided experiment [59] consistent with the hypothesis that rearing is related to information gathering.

From a fluid-dynamical perspective, we expect odor molecules to be distributed differently in air and near the ground, a priori both sniffing the air and the ground could be beneficial. Airborne odors are valuable as distal cues because they are transported rapidly over long distances by flows that are often turbulent. Turbulent flows are sparse and highly intermittent and odor breaks in discrete pockets [108, 38, 45, 17]. This is the reason why local gradients are not helpful in navigating to the odor source [92]. In contrast, ground odor cues are smoother and more continuous than odors in the air [53, 80]. This is because fluid flowing close to the ground slows down due to viscous effects: and creates a region called boundary layer where the structure of the flow depends on the height from the ground [109]. To summarize: odors in the air are more sparse and difficult to measure during navigation than ground cues, yet they are available at longer ranges and they can be




Fig. 5.5 Alternation between different olfactory modalities is widespread in animal behavior. Left: a rodent rearing on hind legs and smelling with its nose high up in the air; a dog performing a similar behavior. Credit: irin-k/Shutterstock.com and Kasefoto/Shutterstock.com. Right: Side view of the direct numerical simulation of odor transport. Shades of blue represent the intensity of velocity fluctuations and are used to visualize the boundary layer near the bottom, where the velocity is reduced by the no-slip condition at the ground. Representative time courses of intense intermittent odor cues in air (sampled at 53 cm from the ground, locations marked with 1 and 2) vs smoother and dimmer cues near ground (sampled at 5 mm from the ground, locations marked with 1' and 2'). Different animals sniff at different heights, which alters details of the plumes but does not affect the general conclusions. Data obtained from direct numerical simulations of odor transport as described in Chapter 2. From [C].

measured earlier than signal at the ground.

In the next Section I will discuss normative theory that rationalizes this alternating behavior and the integration of airborne and ground-based olfactory modalities. In a nutshell, we use the same algorithm described in the previous Section and odor statistics from the simulations described in Chapter 2. In the next section we formalize the olfactory search problem with alternation as a POMDP and we use the algorithm Perseus as a solver. The searcher can choose among 6 actions: moving in one of the four directions while sniffing on the ground; pausing and sniff on the ground; pausing and sniff in the air. We show that the searcher can reach the target by using only one modality, i.e. either ground cues or airborne cues. However the most effective strategy stems from the alternation of the two modalities according to the expected position of the source. Alternation is more frequent far downwind of the source and is associated with casting. The emergence of this non-trivial behavior is rationalized as the need to gather information under strong uncertainty from distal airborne cues, which leads to better long-term reward compared to local exploration for the source or proximal ground cues. Alternation stems spontaneously from navigation when two possible sensor-modalities are allowed.




Fig. 5.6 Snapshots of odor plume obtained from direct numerical simulations of the Navier-Stokes equations in three spatial dimensions. Top view of the odor plume **(a)** at nose height and **(b)** at ground. **(c)** 10% isline of the probability to detect the odor $r(x,y)$ (defined as the probability that odor is above a fixed threshold of 0.14% with respect to the maximum concentration at the source) at the ground (grey) and at the nose height (black). From [C].

5.3.1 Model

Atmosphere and oceans are highly fluctuating environments, odor molecules are carried by the wind or by the currents and turbulent transport dominates over diffusion and defines the statistics of the odor signal. Consider an agent trying to locate a food source that emits odor at a constant rate: first it has to evaluate the direction of the field velocity and then it can move upwind towards the source alternating air and ground sniffs.

We previously mentioned that the no-slip region delimited by the boundary layer is responsible for the different statistics in the air and at the ground (see Figure 5.5). To obtain dimensional estimates of space and time scales we use inlet velocity of 25 cm/s and air viscosity ($\eta = 1.5 \cdot 10^{-5} \text{ m}^2/\text{s}$), as presented in Chapter 2. Odor is released from a spherical source of size 4 cm located 56 cm above the ground. At the height of the source, odor is transported several meters downwind in concentrated puffs, broken and deformed by turbulence: in the air intensity fluctuations of the odor signal and intermittency are high. As depicted in Figure 5.5 odor concentration at the ground is weaker but more continuous since molecules are stuck into the boundary, where turbulent fluctuations are weak; in fact odorant molecules often bind to surfaces, which act as odor sinks [110].




Fig. 5.7 The probability per unit step of detecting an odor signal in the air and at the ground obtained from direct numerical simulations of odor transport. These detection rate maps constitute the observation likelihood models used to train the POMDP. Note that the arena defined in the POMDP is larger than the rate maps shown here (see Figure 5.3 for instance). The detection rate is set zero beyond the bounds of the above rectangles. From [C].

I realized direct numerical simulations with total adsorption of the odor at the ground. While this is a limiting case, we expect similar results for the more realistic case of partial absorption (absorbed particles can be re-emitted in the bulk) as long as the statistics at nose and ground level present differences. As shown in Figure 5.6(a)-(b) the instantaneous snapshot of the odor shows a larger plume at nose level than at ground level, which yields a larger area of detection as shown in Figure 5.6(c).

Navigation proceeds in the same way as in the idealized model described in section 2.4. As in the previous case, there are six actions. However here detections are allowed both in the air and at the ground. Observations o_t are collected by both sniffing air and ground, and lead to update the agent's belief b_t , according to Bayes rule as explained in the previous Section. Two different likelihoods are used for Bayesian updates, corresponding to the rates of odor detection in air and on ground.

The rates of detection are obtained from realistic numerical simulations (Figure 5.6(c), Figure 5.7). Note that odor is detected according to a Poisson process, thus although the rates match those obtained from realistic simulations, the model lacks spatio-temporal correlations.




Fig. 5.8 Representative trajectories undertaken by an agent learning how to reach the source of a turbulent odor cue. **(a)** Top view of a representative trajectory at the end of training. **(b)** Three dimensional view of sample trajectory from panel **(a)**, superimposed to two snapshots of odor plumes near ground (shades of blue) and in the air (shades of red). Trajectories are obtained by training a POMDP, where the agent computes Bayesian updates of the belief using observations (odor detection or no detection) and their likelihood (detection rates from simulations of odor transport). Agents trained with this idealized model of odor plumes successfully track targets when tested in realistic conditions. From [C].

5.3.2 Results

The agent starts its search downwind from the source and learns to efficiently navigate the odor plume to reach the source and get rewarded. Interestingly the alternation between sniffing on the ground and the air (Figure 5.8) emerges spontaneously after training. Increasing the number of training episodes and consequently also the number of sampled points within an episode improves considerably the performance of the agent in locating the source (Figure 5.9), witnessing the emergence of an effective navigational strategy.

The trajectories learnt by the agent display a variety of behaviors reminiscent of those exhibited by animals, which include wide crosswind casts interleaved with upwind surges. Notably, the agent exhibits a recurring motif which cycles between moving to a new location and pausing to sniff in the air. The alternating behavior emerges directly as a consequence of the statistics of the physical environment in spite of pausing to sniff in the air, which leads to the cost of a stronger discount in the reward.




Fig. 5.9 To be able to compare performance across a range of different starting positions we define the additional time to target (time to reach target minus time to travel to the source in a straight line). Performance of POMDP improves with training episodes and saturates after a certain number of iterations. Training converges for different values of the discount factor γ , yellow is $\gamma = 0.90$, red is $\gamma = 0.95$ and blue is $\gamma = 0.99$. The gain in performance with increasing γ is indicative of the development of a long-term strategy. We use $\gamma = 0.99$ and number of training episodes $i = 320$ throughout this section. From [C].

When, where and why does the agent sniff in the air now that it has the option of avoiding pauses and sniff continuously on the ground? Trajectories shown in Figure 5.8 exhibit extensive alternation at the beginning of the search when the agent is far downwind compared to when it is close to the source. A quantitative analysis across training and test realizations confirms that the agent’s rate of sniffing in the air is significantly higher farther away from the source (Figure 5.10(a)). This observation is rationalized by the greater probability of detecting an odor signal in the air at distant locations (Figure 5.6(c)) despite the increased intermittency in the airborne signal (Figure 5.5). In spite of the added cost entailed by slowing down locomotion, sniffing in the air ultimately speeds up the localisation of the source (Figure 5.10(b)). This behavior is maintained across different training realizations and when the discount factor, γ , is reduced so that the delay incurs a greater cost. In sum, alternation emerges as a robust, functional aspect of an effective long-term strategy of olfactory search.

A striking feature of the trajectories in Figure 5.8 is the strong correlation between casting and sniffing the air, especially before the first detection is made. To quantify this effect, we categorize the agent’s behavior into casts and surges, and measure the rate of sniffing the air for both of these behaviors.

We find that the rate of sniffing in the air is typically an order of magnitude greater during casts as compared to surges (Figure 5.10(c)), indicating that alternation is tightly linked to the switch between casting and surging. Casting has been classically interpreted as a strategy for efficient exploration in an intermittent environment. The coupling between casting and alternation observed here suggests that sniffing in the air is an alternative mode of exploration which aids and complements casting when searching for a sparse cue.

Exploration dominates the first part of the search until the first detections which substantially reduce uncertainty (see entropy of the posterior distribution in Figure 5.10(d)). Before the first detection the behavior of the agent is substantially identical to the one illustrated in the previous Section.

Overall, we are led to the following picture of the search dynamics. At the beginning of the search, the agent has a broad prior that is much larger than the odor plume of size $\sim x_{\text{thr}} \times y_{\text{thr}}$, where x_{thr} is the plume length and y_{thr} is the plume width when sniffing in the air. The agent then has to identify and home into the $x_{\text{thr}} \times y_{\text{thr}}$ region that contains the odor plume. The bottleneck in this phase is the scarcity of odor detections, which require an efficient exploration strategy. Once the odor plume is detected, the agent knows it is near the source and the search is driven by surface-borne odor cues, while the frequency of sniffing in the air is significantly reduced. In short, our simulations show that the behavior can be split into two distinct phases: 1) an initial exploration phase accompanied by extensive casting and alternation, where the agent aims at localizing the plume (similar to the simplified model), and 2) odor-guided behavior in a regime relatively rich in cues, which enable the agent to precisely locate the source within the plume.

We conclude this Section noting that the above remarks are expected to hold more generally than in the specific setup of our simulations. The same behaviors are displayed by agents navigating a realistic plume despite their learning in a Poissonian model of odor detections (see [C]). This finding indicates the robustness of the learning scheme to inaccuracies in the model of the environment, which are inevitably present in any realistic situation. More specifically, the static information provided by the average detection rate map is found to be sufficient for navigation and alternation.

While more information on dynamical spatio-temporal correlations may help further performance improvement, the fundamental requirement for alternation is

the presence of wider detection rate maps in the air than on the ground. Thus, as long as this feature is preserved, we expect agents to display alternating behaviors and the two phases mentioned above. In particular, these properties should hold also when different models of odor transport are employed by the searcher and/or surface adsorption chemistry is more involved than pure adsorption.

5.3.3 Conclusions

Examining olfactory navigation and motivated by the goal of disentangling elementary components in the complexity of animal behavior, we have investigated a dynamics driven entirely by olfactory cues. In the model, an agent searches for a source of odors transported by a turbulent flow and at each step decides either to move while sniffing on the ground or pause, rear up and sniff in the air. The goal is to locate the source in the shortest possible time, which is the reward function used to identify effective policies of action using machine-learning methods. Analogously to dogs and rodents mentioned at the beginning of this Section, we obtain behavioral policies which feature alternation between the two modalities of sniffing on the ground *vs* in the air. The appeal of our approach is that we could identify the rationale for the observed alternation and its basic factors. On the one hand, movement and progression toward the source is halted during the rearing phase of sniffing in the air. On the other hand, odor sources create large turbulent plumes that reach larger distances in the air than on the ground. Therefore, sniffing in the air may have a higher chance of intersecting odor cues than on the ground. These two competing effects underlie the process of alternation and their balance determines the rate of switching between the two modalities, which depends on the distance as discussed in the next paragraph.

The effect of alternation is particularly pronounced at large distances from the source. There, due to turbulent mixing, the odor concentration drops substantially and no gradients are present [17]. In our realistic setting, where the searcher does start at large distances, the process can be qualitatively split in two phases : first, the agent needs to approach the source enough for an almost continuous odor plume to be present ; second, it needs to locate the source within the plume. The latter task, which is the regime that most laboratory experiments have considered so far [16], is much easier than the former as the rate of odor detection close to the source and within the conical plume is relatively high. Therefore, the task boils down to staying




Fig. 5.10 Empirical characterisation of the alternation between olfactory sensory modalities. **(a)** The agent sniffs more often in the air when it is far from the source, i.e., outside of the airborne plume. The rate of sniffing in the air is the fraction of times the agent decides to sniff in the air rather than move and sniff on the ground. The fraction is computed over the entire trajectory in the conditions identified in the different panels. Statistics is collected over different realizations of the training process and many trajectories, with different starting positions (see last Section for details). **(b)** The number of steps needed to reach the target minus the number of steps needed to travel from the starting position to the source in a straight line. The horizontal line marks the median, boxes mark 25th and 75th percentiles; red dot: outlier (value exceeds 75th percentile + 1.5 × interquartile range). Dashed lines mark 10th and 90th percentile. For reference, a straight line from the center of the belief to the source is 240 steps. Agents that are given the possibility to pause and sniff in the air are able to reach the target sooner than agents that can only sniff on the ground. **(c)** Agents sniff in the air once every 5 steps on average when they cast, whereas they only sniff in the air once every 60 steps while surging upwind. We consider the agent to be surging if it moves k consecutive steps upwind and casting if it moves k consecutive steps crosswind or sniffs in the air. We use $k = 3$, results shown hereafter do not depend strongly on this choice. **(d)** Entropy (cyan) and value (purple) of the belief vs time, along the course of one trajectory. The red dot indicates a detection, which provides considerable information about source location and thus makes entropy plummet and value increase. From [C].

close to the center of the conical plume, where the signal is highest. Conversely, the

bottleneck during the first, harder phase is the scarcity of information on the location of the source, which the agent tries to overcome by increasing its chances of odor detection. Slowing down its progression is thus the price that the agent pays in order to get oriented in the uncertain conditions typical of large distances to the source. The transition between the two search phases typically occurs after a handful of odor detections.

Note that we have focused here on the case of a stationary source, where odor statistics in the air and on the bottom layers are discriminated by the adsorption on the ground. In fact, at the onset of odor emission (and even in the absence of adsorption), plumes start out larger in the air than near the ground, simply because air travels more slowly near the ground. It follows from our results that alternation should be more frequent in the early stages of odor release in non-steady conditions. This prediction could be tested experimentally by switching on an odor source and monitoring the fraction of sniffing in the air as a function of the time elapsed since the switch and the onset of odor emission.

Partially Observed Markov Decision Processes (POMDP) [96, 97] can be applied to a broad class of decision problems, where agents need to accomplish a prescribed task by a series of actions taken with partial knowledge of the environment. Specifically, the agent combines external cues and its internal model of the world to infer a belief about the state of the environment. Here, we have given the agent the choice of multiple sensory modalities at each decision step, which allowed us to highlight the presence of alternation and establish its link with marginal value theory (MVT) [101].

MVT describes the behavior of an optimally foraging individual in a system with spatially separated resources. Due to the spatial separation, animals must spend time traveling between patches. Since organisms face diminishing returns, there is a moment the animal exhausts the patch and ought to leave. In MVT, the optimal departure time is determined as the time at which the marginal value of staying in a patch equals that of leaving and exploring another patch. In our setting, these patches correspond to regions of the agent's belief which are explored using a combination of casting and sniffing in the air. MVT thus determines when to stop cast-and-sniff exploration and surge towards the next patch in the belief.

While we considered two olfactory sensorimotor modalities, our methodology and results apply more broadly to distinct sensory systems and cues. If there is no

conflict in the acquisition and processing of multiple sensory cues, then it is clearly advantageous to combine them. Conversely, if their combination has some form of cost and a partial or total conflict exists, which we expect to be the generic case, then our results predict that there will be alternation and that it will follow the same logic identified here.

We conclude by noting that, in addition to the familiar cases of dogs and rodents mentioned in the introduction, other species can sense chemical cues both in the bulk and on surfaces, and may feature a similar phenomenology of alternation. In particular, a large body of experimental evidence has been collected for turbulent plume-tracking by aquatic organisms, as reviewed in [35]. Crustaceans sense chemical cues with their antennules floating in water and switch to sensing with their feet as they approach the target [111]. For example, lobsters were observed in dim light in a flume of dimensions 2.5m x 90cm x 20cm, as they left their shelter upon release of a turbulent plume of odor obtained from grounded mussel [112]. As the animals encountered the plume, they often displayed special behaviors, including raising up, sweeping their sensory legs on the bottom of the flume and increasing flicking of lateral antennules. Similar observations were made for blue crabs capturing live clams or tracking spouts releasing clam extract [113]. In these experiments blue crabs would occasionally lower their abdomen closer to the surface or extend their walking legs to raise above their normal height. Finally, pelagic marine mollusks *Nautilus pompilius* were observed to track the source of a turbulent plume by swimming at different heights, above and below the center of the plume. Interestingly, most animals sampled at higher heights beyond one meter from the source, and swam at lower heights when closer to the source [114]. These experiments indicate that animals may alternate between different heights, and that sampling at higher elevation may be particularly useful at larger distances, which is again in qualitative agreement with our results. The ensemble of these observations suggest that alternation between sensorimotor modalities is likely to be present in the behavior of aquatic organisms as well. The results presented may inspire more experiments, on dogs, rodents and aquatic organisms alike, with the goal of assessing quantitative aspects of the observed behaviors, testing our framework and advancing understanding of how sensorimotor modalities are integrated.

Publication List

[A] **N. Rigolli**, N. Magnoli, L. Rosasco, and A. Seminara. Learning to predict target location with turbulent odor plumes. *q-bio. arXiv:2106.08988*, 2021. (Accepted in eLife).

[B] S.M. Lewis, L. Xu, **N. Rigolli**, M. F. Tariq, L. M. Suarez, M. Stern, A. Seminara and D. H. Gire. Plume Dynamics Structure the Spatiotemporal Activity of Mitral/Tufted Cell Networks in the Mouse Olfactory Bulb. *Frontiers in Cellular Neuroscience*, 15:104, 2021.

[C] **N. Rigolli**, G. Reddy, A. Seminara and M. Vergassola. Alternation emerges as a multi-modal strategy for turbulent odor navigation. *bioRxiv 2021.12.14.472675*, 2021. (In revision, eLife).

References

- [1] Ed Yong. Thread on twitter about his new book: "an immense world", 2022. <https://twitter.com/edyong209/status/1483441016749412356>.
- [2] M. D. Manson. Bacterial motility and chemotaxis. volume 33 of *Advances in Microbial Physiology*, pages 277–346. Academic Press, 1992.
- [3] N. Levernier, O. Pouliquen, and Y. Forterre. An integrative model of plant gravitropism linking statoliths position and auxin transport. *Frontiers in Plant Science*, 12, 2021.
- [4] A. Angel, J. Song, C. Dean, and M. Howard. A polycomb-based switch underlying quantitative epigenetic memory. *Nature*, 476:105–108, 2011.
- [5] A. Mafra-Neto and R.T. Cardé. Fine-scale structure of pheromone plumes modulates upwind orientation of flying moths. *Nature*, 369:142–144, 1994.
- [6] D. Lagomarsino Oneto, J. Golan, A. Mazzino, A. Pringle, and A. Seminara. Timing of fungal spore release dictates survival during atmospheric transport. *Proceedings of the National Academy of Sciences*, 117(10):5134–5143, 2020.
- [7] B. Gaylord, J. Hodin, and M.C. Ferner. Turbulent shear spurs settlement in larval sea urchins. *Proceedings of the National Academy of Sciences*, 110(17):6901–6906, 2013.
- [8] J. Adler. The sensing of chemicals by bacteria. *Scientific American*, 234(4):40–47, 2014.
- [9] H.C. Berg. *E. coli in Motion*. Springer Science, 2008.
- [10] V. Sourjik and N. S. Wingreen. Responding to chemical gradients: bacterial chemotaxis. *Current opinion in cell biology*, 24(2):262–268, 2012.
- [11] G. H. Wadhams and J. P. Armitage. Making sense of it all: bacterial chemotaxis. *Nature reviews. Molecular cell biology*, 5(12):1024–1037, 2004.
- [12] J.J. Falke, R.B. Bass, S.B. Butler, S.A. Chervitz, and M.A. Danielson. The two-component signaling pathway of bacterial chemotaxis: a molecular view of signal transduction by receptors, kinases, and adaptation enzymes. *Annual Review of Cell and Developmental Biology*, 13:457–512, 1997.

-
- [13] D.J. Webre, P.M. Wolanin, and J.B. Stock. Bacterial chemotaxis. *Current Biology*, 13, no. 2:R47–R49, 2003.
- [14] H.C. Berg and E. M. Purcell. Physics of chemoreception. *Biophysics Journal*, 20:193–219, 1977.
- [15] N. Vladimirov and V. Sourjik. Chemotaxis: how bacteria use memory. *Biophysics Journal*, 390(10):1097–1104, 2009.
- [16] G. Reddy, V.N. Murthy, and M. Vergassola. Olfactory sensing and navigation in turbulent environments. *Annual Review of Condensed Matter Physics*, 13:191–213, 2022.
- [17] A. Celani, E. Villermaux, and M. Vergassola. Odor landscapes in turbulent environments. *Physical Review X*, 4:041015, 2014.
- [18] K. Okajima, S. Shigaki, T. Suko, D.N. Luong, R.C. Hernandez, Y. Hattori, K. Sanada, and D. Kurabayashi. A novel framework based on a data-driven approach for modelling the behaviour of organisms in chemical plume tracing. *Journal of the Royal Society Interface*, (18), 2021.
- [19] A. Borst and M. Heisemberg. Osmotropotaxis in drosophila melanogaster. *Journal of Computational Physiology*, 147:479–484, 1982.
- [20] S.A. Budick and M.H. Dickinson. Free-flight responses of *Drosophila melanogaster* to attractive odors. *Journal of Experimental Biology*, 209(15):3001–3017, 2006.
- [21] M. Demir, N. Kadakia, H.D. Anderson, D.A. Clark, and T. Emonet. Walking drosophila navigate complex plumes using stochastic decisions biased by the timing of odor encounters. *eLife*, 9:e57524, 2020.
- [22] E. Álvarez Salvado, A.M. Licata, E.G. Connor, M.K. McHugh, B.M.N. King, N. Stavropoulos, J.D. Victor, J.P. Crimaldi, and K.I. Nagel. Elementary sensory-motor transformations underlying olfactory navigation in walking fruit-flies. *eLife*, 7:e37815, 2018.
- [23] H. Tichy, A. Hinterwirth, and E. Gingl. Olfactory receptors on the cockroach antenna signal odour on and odour off by excitation. *The European journal of Neuroscience*, 22(12), 2005.
- [24] J. Atema. Eddy chemotaxis and odor landscapes: Exploration of nature with animal sensors. *The Biological Bulletin*, 191:129–138, 1996.
- [25] J. Leah Jackson, Donald R. Webster, S. Rahman, and Marc J. Weissburg. Bed roughness effects on boundary-layer turbulence and consequences for odor-tracking behavior of blue crabs (*callinectes sapidus*). *Limnology and Oceanography*, 52:1883–1897, 2007.

-
- [26] J.M. Gardiner and J. Atema. The function of bilateral odor arrival time differences in olfactory orientation of sharks. *Current Biology*, 20:1187–1191, 2010.
- [27] B.D. Dickman, D.R. Webster, J.L. Page, and M.J. Weissburg. Three-dimensional odorant concentration measurements around actively tracking blue crabs. *Limnology and Oceanography*, 7:96–108, 2009.
- [28] T. A. Keller and Marc J. Weissburg. Effects of odor flux and pulse rate on chemosensory tracking in turbulent odor plumes by the blue crab, *callinectes sapidus*. *The Biological Bulletin*, 207(1):44–55, 2004.
- [29] M.A.R. Koehl. The Fluid Mechanics of Arthropod Sniffing in Turbulent Odor Plumes. *Chemical Senses*, 31(2):93–105, 2005.
- [30] N.J. Vickers. Mechanisms of animal navigation in odor plumes. *The Biological Bulletin*, 198:203–212, 2000.
- [31] J. Crimaldi, H. Lei, A.T. Schaefer, M. Schmuker, B.H. Smith, A.C. True, J.V. Verhagen, and J.D. Victor. Active sensing in a dynamic olfactory world. *Journal of Computational Neuroscience*, 50:1–6, 2022.
- [32] P. K. Kundu, I. M. Cohen, and D. R. Dowling. *Fluid Mechanics*. Academic Press, 2016.
- [33] U. Frisch. *Turbulence. The legacy of A. N. Kolomogorov*. Cambridge: Cambridge University Press, 1995.
- [34] S. Pirozzoli, M. Bernardini, and P. Orlandi. Passive scalars in turbulent channel flow at high reynolds number. *J. Comput. Phys.*, 788:614–639, 2016.
- [35] D.R. Webster and M.J. Weissburg. The hydrodynamics of chemical cues among aquatic organisms. *Annual Review of Fluid Mechanics*, 41:73–90, 2009.
- [36] K. R. Sreenivasan. Turbulent mixing: A perspective. *Proceedings of the National Academy of Sciences*, 116(37):18175–18183, 2019.
- [37] Stephen B. Pope. *Turbulent Flows*. Cambridge: Cambridge University Press., 2000.
- [38] B. Shraiman and E. Siggia. Scalar turbulence. *Nature*, 405:639–646, 2000.
- [39] P. F. Fischer, J. W. Lottes, and J.G. Kerkemeier. Nek5000 web page, 2008. <http://nek5000.mcs.anl.gov>.
- [40] A.T. Patera. A spectral element method for fluid dynamics : laminar flow in a channel expansion. *Journal of Computational Physics*, 54:468–488, 1984.
- [41] S.A. Orszag. Spectral methods for problems in complex geometry. *Journal of Computational Physics*, 37:70–92, 1980.

-
- [42] L.W. Ho. *A Legendre spectral element method for simulation of incompressible unsteady viscous free-surface flows*. Ph.D. thesis, Massachusetts Institute of Technology, 1989.
- [43] Paul F. Fischer, Francis Loth, Seung E. Lee, Sang-Wook Lee, David S. Smith, and Hisham S. Bassiouny. Simulation of high-reynolds number vascular flows. *Computer Methods in Applied Mechanics and Engineering*, 196(31):3049–3060, 2007.
- [44] S. Saddoughi and S. V. Veeravalli. Local isotropy in turbulent boundary layers at high reynolds number. *Journal of Fluid Mechanics*, 268:333–372, 1994.
- [45] G. Falkovich, K. Gawedzki, and M. Vergassola. Particles and fields in fluid turbulence. *Reviews of Modern Physics*, 73:913, 2001.
- [46] J. Duplat, A. Jouary, and E. Villiermaux. Entanglement rules for random mixtures. *Phys. Rev. Lett.*, 105:034504, 2010.
- [47] T Hastie, R Tibshirani, and J Friedman. *The elements of statistical learning: Data mining, inference, and prediction*. New York: Springer, 2001.
- [48] B. Schölkopf and AJ. Smola. *Learning with Kernels: Support Vector Machines, Regularization, Optimization, and Beyond*. Adaptive Computation and Machine Learning. MIT Press, Cambridge, MA, USA, December 2002.
- [49] S. Gorur-Shandilya, M. Demir, J. Long, D.A. Clark, and T. Emonet. Olfactory receptor neurons use gain control and complementary kinetics to encode intermittent odorant stimuli. *eLife*, 6:e27670, 2017.
- [50] A. Rudi, L. Carratino, and L. Rosasco. Falkon: An optimal large scale kernel method, 2018. arXiv 1705.10958.
- [51] J.R. Martelli, C. Carlson and T. Emonet. Intensity invariant dynamics and odor-specific latencies in olfactory receptor neuron response. *The Journal of Neuroscience*, 33(15):6285– 6297, 2013.
- [52] K.A. Justus, J. Murlis, C. Jones, and R.T. Cardé. Measurement of odor-plume structure in a wind tunnel using a photoionization detector and a tracer gas. *Environmental Fluid Mechanics*, 2:115–142, 2002.
- [53] J.E. Fackrell and A.G. Robins. Concentration fluctuations and fluxes in plumes from point sources in a turbulent boundary layer. *Journal of Fluid Mechanics*, 117:1, 1982.
- [54] I.J. Park, Y.V. Bobkov, B.W. Ache, and J.C. Principe. Intermittency coding in the primary olfactory system: A neural substrate for olfactory scene analysis. *Journal of Neuroscience*, 34:941–952, 2014.
- [55] B.W. Ache, A.M. Hein, Y.V. Bobkov, and J.C. Principe. Smelling time: A neural basis for olfactory scene analysis. *Trends Neuroscience*, 39:649–655, 2016.

-
- [56] M. Smear, A. Resulaj, J. Zhang, T. Bozza, and D. Rinberg. Multiple perceptible signals from a single olfactory glomerulus. *Nature Neuroscience*, 16:1687–1691, 2013.
- [57] J.D. Victor, S.D. Boie, E.G. Connor, J.P. Crimaldi, Ermentrout G.B., and K.I. Nagel. Olfactory navigation and the receptor nonlinearity. *Journal of Neuroscience*, 39(19):3713–3727, 2019.
- [58] S.D. Boie, E.G. Connor, M. McHugh, K.I. Nagel, G.B. Ermentrout, J.P. Crimaldi, and J.D. Victor. Information-theoretic analysis of realistic odor plumes: What cues are useful for determining location? *PLoS Computational Biology*, 14:e1006275, 2018.
- [59] D.H. Gire, V. Kapoor, A. Arrighi-Allisan, A. Seminara, and V.N. Murthy. Mice develop efficient strategies for foraging and navigation using complex natural stimuli. *Current Biology*, 26(D1):1261–1273, 2016.
- [60] I.J. Park, A.M. Hein, Y.V. Bobkov, M.A. Reidenbach, B.W. Ache, and J.C. Principe. Neurally encoding time for olfactory navigation. *PLoS Computational Biology*, 12:1–16, 01 2016.
- [61] B.T. Michaelis, K.W. Leathers, Y.V. Bobkov, B.W. Ache, J.C. Principe, R. Baharloo, I.M. Park, and M.A. Reidenbach. Odor tracking in aquatic organisms: the importance of temporal and spatial intermittency of the turbulent plume. *Scientific Reports*, 10:7961, 2020.
- [62] K.W. Leathers, B.T. Michaelis, and M.A. Reidenbach. Interpreting the spatial-temporal structure of turbulent chemical plumes utilized in odor tracking by lobsters. *Fluids*, 5(2), 2020.
- [63] M. Schmuker, V. Bahr, and R. Huerta. Exploiting plume structure to decode gas source distance using metal-oxide gas sensors. *Sensors and Actuators B*, 235:636–646, 2016.
- [64] S. Gorur-Shandilya, C. Martelli, M. Demir, and T. Emonet. Controlling and measuring dynamic odorant stimuli in the laboratory. *The Journal of Experimental Biology*, 222:207787, 2019.
- [65] G. Reddy, B.I. Shraiman, and M. Vergassola. Sector search strategies for odor trail tracking. 119(1):e2107431118.
- [66] S. Pifferi, A. Menini, and T. Kurahashi. *The Neurobiology of Olfaction*. Boca Raton: CRC Press/Taylor Francis, 2010.
- [67] I. Glezer and B. Malnic. Olfactory receptor function. *Handbook of Clinical Neurology*, 167:67–78, 2019.
- [68] B. Malnic, J. Hirono, T. Sato, and L.B. Buck. Combinatorial receptor codes for odors. *Cell*, 96(5):713–723, 1999.

-
- [69] G. Reddy, J.D. Zak, M. Vergassola, and V.N. Murthy. Antagonism in olfactory receptor neurons and its implications for the perception of odor mixtures. *eLife*, 7:e34958, 2018.
- [70] V. Singh, N.R. Murphy, V. Balasubramanian, and J.D. Mainland. Competitive binding predicts nonlinear responses of olfactory receptors to complex mixtures. *Proceedings of the National Academy of Sciences*, 116(19):9598–9603, 2019.
- [71] T. Ackels, A. Erskine, D. Dasgupta, and A.T. Schaefer. Fast odour dynamics are encoded in the olfactory system and guide behaviour. *Nature*, 593:558–563, 2021.
- [72] L. Buck and R. Axel. A novel multigene family may encode odorant receptors: a molecular basis for odor recognition. *Cell*, 65(1):175–187, 1991.
- [73] C.Y. Su, K. Menuz, and J.R. Carlson. Olfactory perception: receptors, cells, and circuits. *Cell*, 139(1):45–59, 2009.
- [74] D.H. Brann and S.R. Datta. Finding the brain in the nose. *Annual Review of Neuroscience*, 43:277–295, 2020.
- [75] K.M. Igarashi, N. Ieki, M. An, Y. Yamaguchi, S. Nagayama, K. Kobayakawa, R. Kobayakawa, M. Tanifuji, H. Sakano, W.R. Chen, and K. Mori. Parallel mitral and tufted cell pathways route distinct odor information to different targets in the olfactory cortex. *Journal of Neuroscience*, 32 (23):7970–7985, 2012.
- [76] X. Zhang and S. Firestein. The olfactory receptor gene superfamily of the mouse. *Nature neuroscience*, 5(2):124–133, 2002.
- [77] R. Batista-Brito, M. Vinck, K.A. Ferguson, J. Chang, D. Laubender, and G. Lur. Developmental dysfunction of vip interneurons impairs cortical circuits. *Neuron*, 95:884–895, 2017.
- [78] M.F. Tariq, S.M. Lewis, A. Lowell, S. Moore, J.T. Miles, and D.J. Perkel. Using head-mounted ethanol sensors to monitor olfactory information and determine behavioral changes associated with ethanol-plume contact during mouse odor-guided navigation. *eNeuro*, 8:ENEURO.0285–20.2020, 2021.
- [79] D. Martinez, J. Burgués, , and S. Marco. Fast measurements with mox sensors: a least-squares approach to blind deconvolution. *Sensors*, 19:4029, 2019.
- [80] C. Nironi, P.o Salizzoni, M. Marro, P. Mejean, N. Grosjean, and L. Soulhac. Dispersion of a passive scalar fluctuating plume in a turbulent boundary layer. part i: Velocity and concentration measurements. *Boundary-Layer Meteorology*, 156:415–446, 2015.
- [81] M. Stern, E. Shea-Brown, and D. Witten. Inferring the spiking rate of a population of neurons from wide-field calcium imaging. *bioRxiv*, 2020.

-
- [82] N.J. Vickers and W. It. Moth flight behavior and responses of olfactory neurons are shaped by pheromone plume dynamics. *Chemical Senses*, 31(2):155–166, 2006.
- [83] N.Z. Shor. *Minimization Methods for Non-Differentiable Functions and Applications*. Berlin; New York: Springer-Verlag, 1985.
- [84] F. Van Bruegel and M. H. Dickinson. Plume-tracking behavior of flying drosophila emerges from a set of distinct sensory-motor reflexes. *Current Biology*, 24:274–286, 2014.
- [85] A.J. Lilienthal, D. Reimann, and A. Zell. Gas source tracing with a mobile robot using an adapted moth strategy. In *Autonome Mobile Systeme (AMS)*, 2003.
- [86] S. Shigaki, T. Sakurai, N. Ando, D. Kurabayashi, and R. Kanzak. Time-varying moth-inspired algorithm for chemical plume tracing in turbulent environment. *IEEE Robotics and Automation Letters*, 3(1):76–83, 2017.
- [87] M.J. Weissburg and R.K. Zimmer-Faust. Odor plumes and how blue crabs use them in finding prey. *Journal of Experimental Biology*, 197:349–375, 1994.
- [88] M. A. Willis, J.L. Avondet, and A.S. Finnell. Effects of altering flow and odor information on plume tracking behavior in walking cockroaches, *Periplaneta americana* (L.). *Journal of Experimental Biology*, 211(14):2317–2326, 07 2008.
- [89] N. Kadakia, M. Demir, B.T. Michaelis, M.A. Reidenbach, D.A. Clark, and T. Emonet. Odor motion sensing enables complex plume navigation. *bioRxiv*, 2021.
- [90] L. Wang. Robotic olfactory-based navigation with mobile robots. *PhD Dissertations and Master’s Theses, Embry-Riddle Aeronautical University Daytona Beach, Florida*, 2021.
- [91] R. Pang, F. van Breugel, M. Dickinson, J. A. Riffell, and A. Fairhall. History dependence in insect flight decisions during odor tracking. *PLoS Computational Biology*, 14(2):1–26, 2018.
- [92] M. Vergassola, E. Villermaux, and B.I. Shraiman. ‘Infotaxis’ as a strategy for searching without gradients. *Nature*, 445:406, 2007.
- [93] E.M. Moraud and D. Martinez. Effectiveness and robustness of robot infotaxis for searching in dilute conditions. *Frontiers in Neurobotics*, 4,1, 2010.
- [94] J.B. Masson. Olfactory searches with limited space perception. *Proceedings of the National Academy of Sciences*, 110(28):11261–11266, 2013.
- [95] G. Shani, J. Pineau, and R. Kaplow. A survey of point-based pomdp solvers. *Autonomous Agents and Multi-Agent Systems*, 27(1):1–51, 2013.

-
- [96] L. Kaelbling, M. Littman, and Cassandra A. Planning and acting in partially observable stochastic domains. *Artificial Intelligence*, 101:99134, 1998.
 - [97] R.S. Sutton and A.G. Barto. *Reinforcement Learning : an Introduction*. MIT Press, Cambridge, 2nd edition edition, 2018.
 - [98] R.E. Bellman. *Dynamic Programming*. Dover, 2003.
 - [99] E.J. Sondik. The optimal control of partially observable markov processes over the infinite horizon: Discounted costs. *Operations research*, 26(2):282–304, 1978.
 - [100] M.T.J. Spaan and N. Vlassis. Perseus: Randomized point-based value iteration for pomdps. *Journal of Artificial Intelligence Research*, 24:195–220, 2005.
 - [101] E.L. Charnov. Optimal foraging: the marginal value theorem. *Theoretical Population Biology*, 9:129–136, 1976.
 - [102] A.G. Khan, M. Sarangi, and U.S. Bhalla. Rats track odour trails accurately using a multi-layered strategy with near-optimal sampling. *Nature Communications*, 3:703, 2012.
 - [103] A. Thesen, J.B. Steen, and K.B. Døving. Behaviour of dogs during olfactory tracking. *The Journal of Experimental Biology*, 180:247–251, 1993.
 - [104] J.B. Steen, I. Mohus, T. Kvesetberg, and L. Walløe. Olfaction in bird dogs during hunting. *Acta Physiologica*, 157:115–119, 1996.
 - [105] P.G. Hepper and D.L. Wells. How many footsteps do dogs need to determine the direction of an odour trail? *Chemical Senses*, 30(4):291–298, 2005.
 - [106] J. Jinn, E.G. Connor, and L.F. Jacobs. How ambient environment influences olfactory orientation in search and rescue dogs. *Chemical senses*, 45(8):625–634, 2020.
 - [107] C. Lever, S. Burton, and J. O’Keefe. Rearing on hind legs, environmental novelty and the hippocampal formation. *Reviews in the Neurosciences*, 17:111–133, 2006.
 - [108] J. Murlis and C.D. Jones. Fine-scale structure of odor plumes in relation to insect orientation to distant pheromone and other attractant sources. *Physiological Entomology*, 6:71–86, 1981.
 - [109] J.D. Jr Anderson. Ludwig Prandtl’s boundary layer. *Physics Today*, 58:42, 2005.
 - [110] S. Gorur-Shandilya, C. Martelli, M. Demir, and T. Emonet. Controlling and measuring dynamic odorant stimuli in the laboratory. *Journal of Experimental Biology*, 222(23), 11 2019.

-
- [111] F.W. Grasso. Invertebrate-inspired sensory-motor systems and autonomous, olfactory-guided exploration. *The Biological Bulletin*, 200:160–168, 2001.
 - [112] P.A. Moore, N. Scholz, and J. Atema. Chemical orientation of lobsters, *homarus americanus*, in turbulent odor plumes. *Journal of Chemical Ecology*, 17:1293–1307, 1991.
 - [113] M.J. Weissburg, D.B. Dusenbery, H. Ishida, J. Janata, T. Keller, P.J.W. Roberts, and D.R. Webster. A multidisciplinary study of spatial and temporal scales containing information in turbulent chemical plume tracking. *Environmental Fluid Mechanics*, 2:65–94, 2002.
 - [114] J.A. Basil, R.T. Hanlon, S.I. Sheikh, and J. Atema. Three-dimensional odor tracking by *Nautilus pompilius*. *The Journal of Experimental Biology*, 203:1409–1414, 2000.

This work is licensed under the CC-BY-SA 4.0 International license. To view a copy of this license, visit <http://creativecommons.org/licenses/by/4.0/> or send a letter to Creative Commons, PO Box 1866, Mountain View, CA 94042, USA. 

Acknowledgements

Let me start with institutional, but no less genuine, acknowledgments: thanks to Università Italo-Francese for funding my travel expenses with the Vinci program. Thanks to the Referees of my thesis: Antonio Celani and Greg Stephens, it was an honor for me to receive your positive feedback. I am looking forward to having you here in Genova with Mederic Argentina and Claire Wyart to discuss this in person. Thank you for your availability: having such a high-quality committee is a privilege and a responsibility.

During the first pandemic wave, I was surfing the net and my attention got caught by a pretty Latvian magazine: the outlook was colorful and stylish, and the title of that volume "I'm training in uncertainty" perfectly suited the moment. I found that funny because it could have been the perfect title for my thesis. Anyway, life and thesis strictly merge during a Ph.D., here is the time to sum up the last four years and express my gratitude to the people I had around and who supported me during my training. A short disclaimer: I am mentioning the people I shared part of my academic life, as I do not expect my other beloved to read these few lines.

First of all, I want to thank Nico Magnoli, I vividly remember you reaching me in the old Difi library (sad face) just for telling me there was someone I had to meet: you knew about my interests and you introduced me to Agnese Seminara. You always show me your kind support in Science and in all the smaller, but no less important, details of life in academia. I hope you will recover soon and finish the books you dream of writing.

I sincerely feel no words would be enough to thank you, Agnese. You considered me a scientist from the very beginning and gave me the opportunity to do things, try myself, think, and learn. I always felt free to make my choices with the awareness that I was safely guided by you. Thanks for the uncountable things I learned from you: how to write and prepare good presentations, to have a proactive scientific approach, I wish one day to have your brilliant intuitions. Thanks for encouraging me to take part in schools and conferences: I had the possibility to meet excellent scientists and people from all over the world. Finally, thanks for showing me that the best reason for being a scientist and doing research is that Science is fun. I want to underline that having a female boss in academia helped me to interiorize a positive approach to gender equality.

Thanks to Daniele Lagomarsino Oneto, you are a friend and an inspiration. I appreciate so much your devotion to Science. Over these years the times you helped

or gave me tips are countless. I am happy to share the office with you and with Gianvito Losapio.

Thanks to Dario Vincenzi for being in my “suivi de these” with Christophe Raufaste. Your kindness, humor, and invitations from seminars to free dinner at the Italian consulate made my stay in Nice much more pleasant.

Thanks to Lorenzo Rosasco, working on the paper all together was fun and enriching. I admire the magic world you created at Malga.

Thanks to Chris Rycroft for inviting me to Harvard in the summer of 2019. You really are the kindest person I met during my journey, and thanks to all the students who made my short time there very joyful.

Thanks to David Gire for our collaboration and for the time at UW and Friday Harbor. I hope you will get better soon! Everyone can not wait to have you back in the lab and at our Zoom meetings.

Thanks to Massimo Vergassola and Gautam Reddy, I feel lucky to collaborate with both of you, every time we discuss I have the occasion to learn something new.

You may think it is a choice not to forget someone, but no. I am paranoid about thinking my personal information will stay indefinitely online in this manuscript. I will not mention any name from now on: if you know, you know!

Starting from the beginning, let me thank my mates from Physics in Genova. We shared the joys and obstacles from Bachelor’s to Ph.D. and it was a hell of a journey.

Thanks to my Ph.D. colleagues at Inphyni and in Nice, I can not believe such a long time has passed since the last time I was there. Beers, board games, Ramy, and football games in the lab will stay with me forever.

I want to thank the participants of Boulder Summer School 2019. It was an unforgettable experience: anytime I meet again someone of you I am surprised by the strength of our bond.

I thank Ph.D. students at Gire Lab both in the mice and octopus team: I stayed there only a few days and you made me feel at home, Friday Zoom calls are a pleasant routine.

I have to thank my grandma for hosting me and my research during the pandemic.

I thank all the Malga people, Malga became a safe space for me. I love to pass by for some chatting, quality coffee, or just for smiling and sharing some happiness.

I thank PhDs, postdocs, and professors from Dicca. Villa Cambiaso is a gorgeous building, but walls stay walls and the people making them alive are much more important.

Finally, thanks to my sister, this is a check to know if you are reading my thesis.

Declaration

I hereby declare that, the contents and organization of this dissertation constitute my own original work and does not compromise in any way the rights of third parties, including those relating to the security of personal data.

Nicola Rigolli
2022

* This dissertation is presented in partial fulfillment of the requirements for **Ph.D. degree** in the Graduate School of Università degli Studi di Genova.