Olfactory navigation: how to make decisions using a broken signal

Nicola Rigolli (Nicodemo Magnoli, Agnese Seminara)

Background

During my third year of PhD I obtained some interesting results and I clearly defined the domain of my research: the general setting of my work is studying how biological systems adapt to decode the sparse information brought by turbulence; in particular I tried to understand how an animal can extract information from olfactory cues. My research stands at the interface between Fluid Dynamics, Machine Learning, Animal Behavior and Neuroscience.

Plume dynamics impacts on neural activity and affects animal behavior

Mice and many other animals locate resources using turbulent airborne odor plumes, these signals are stochastic and intermittent and challenging to interpret. In [1] we measured the Olfactory Bulb response to naturally fluctuating odor plumes and showed that plume dynamics shapes neural olfactory representation in mice. What are the features of the odor signal the brain extracts when exposed to natural olfactory scenes is unknown.

Simple odor features allow agents to make sense of turbulent flows

In my first and second year reports I already introduced the problem and presented some of the techniques we used to solve it. In particular I realized some state of the art direct numerical simulations with open source software Nek5000 [2] and I developed Machine Learning algorithms [3] to infer the location of an odor source measuring a fluctuating odor signal in time. We distinguished two categories of odor features that are relevant to obtain precise predictions: when the animal is closer to the source and the signal is continuous is convenient to rely on the intensity of the odor signal, while when the intermittency is high is better measuring the timing of odor encounters. Interestingly the combination of timing and intensity improves the performance of the algorithm in every search conditions [4].

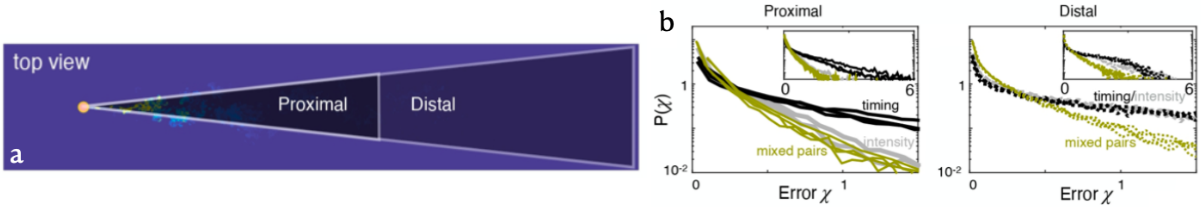

Figure 1: a) Top view of an odor field snapshot at the source height, space is divided in proximal and distal depending on the distance from the source. b) Distribution of test error for the proximal (left) and distal (right) problem showing intensity features (grey) outperform timing features (black) 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.

Algorithms and animals optimize strategies to navigate towards an odor source

Animals navigate the surrounding environment to find food or conspecifics, Reinforcement Learning algorithms are commonly used to solve navigation problems [5]-[6]. Interestingly dogs and rodents alternate between sniffing in the air and close to the ground while tracking an odor scent. In fact ground and airborne odor signals convey different information even if both signals are generated by a unique source of odors. Indeed, airborne odors cues are valuable as distal cues because they are transported rapidly over long distances by flows that are turbulent. The downside of airborne cues is that turbulence breaks odors plumes in discrete pockets that can only be detected sparsely. Furthermore, since local gradients are randomized in relation to the source direction at the timescales of olfactory searches, gradient-ascent navigation strategies are not possible. Conversely, odor cues close to the ground are smoother and more continuous than odors in the air. It is therefore likely that the relative value of sniffing closer to *vs* farther from the ground depends on the position of the searcher relative to the source. We created a realistic environment with direct numerical simulations and we look for the optimal solution to the navigation problem using Partially Observable Markov Decision Process (POMDP). The agent performing the olfactory search is given the choice between the actions of moving while sniffing on the ground or stopping and sniffing in the air. Solving the POMDP problem yields a policy of actions taken in response to a history of odor stimuli, which is encoded into a set of probabilistic beliefs on the location of the source. While the searcher could *a priori* reach the target by using ground cues only, we demonstrate that learned strategies generically feature alternation between airborne and ground odor cues. The emergence of this non-trivial behavior is rationalized as reflecting the exploitation-exploration trade-off that results from the need to balance gathering of information with the request to move rapidly to reach the target.

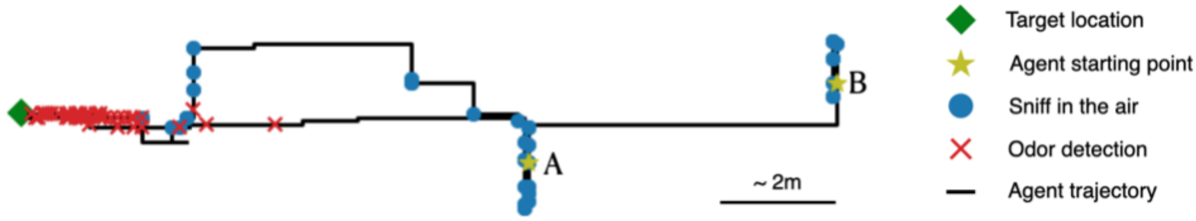

Figure 2: Two different trajectory by the POMDP. A is starting closer to the source but far away from the plume centerline, while B is starting far away: both the trajectories successfully reach the odor source.

References

[1] S. Lewis et al. Plume Dynamics Structure the Spatiotemporal Activity of Mitral/Tufted Cell Networks in the Mouse Olfactory Bulb. *Frontiers in Cellular Neuroscience 15 Journal* 2021.

[2] D. Fischer er al. NEK5000 web page, *Http://nek5000.mcs.anl.gov.* 2008.

[3] C. E. Rasmussen, C. K. I. Williams, Gaussian Processes for Machine Learning, the MIT Press, 2006

[4] N. Rigolli et al. Learning to predict target location with turbulent odor plume. *https://arxiv.org/abs/2106.08988* 2021.

[5] R. Sutton and A. Reinforcement Learning: an Introduction. MIT Press, 2nd edition, 2018

[6] G. Reddy et al. Olfactory sensing and navigation in turbulent environments. *Ann. Rev. of Condensed Matter Physics* 2021.

1 Scientific publications

- Plume Dynamics Structure the Spatiotemporal Activity of Mitral/Tufted Cell Networks in the Mouse Olfactory Bulb. Lewis Suzanne M., Xu Lai, Rigolli Nicola, Tariq Mohammad F., Suarez Lucas M., Stern Merav, Seminara Agnese, Gire David. *Frontiers in Cellular Neuroscience 15 (2021) Journal*

- Learning to predict target location with turbulent odor plume. Nicola Rigolli, Nicodemo Magnoli, Lorenzo Rosasco and Agnese Seminara. *under revision*

- Learning to integrate sensory modalities by alternating airborne and ground odor cues. Nicola Rigolli, Gautam Reddy, Massimo Vergassola, Agnese Seminara. *on going*

2 Attended courses

- Quantum Phase Transition (passed the exam)
- Computazione Quantistica (passed the exam)
- Advanced Computational Physics (passed the exam)
- La Matematica del Machine Learning (passed the exam (x2))
- Boulder School 2019: Theoretical Biophysics (passed the exam)

3 Conferences

- Stochastic Models and Experiments in Ecology and Biology, Venice 22-25 June 2021 - oral presentation
- Physics meets Biology IOP, online conference 26-28 July 2021 - poster presentation