



HAL
open science

Effectiveness and robustness of robot infotaxis for searching in dilute conditions

Eduardo Martin Moraud, Dominique Martinez

► **To cite this version:**

Eduardo Martin Moraud, Dominique Martinez. Effectiveness and robustness of robot infotaxis for searching in dilute conditions. *Frontiers in Neurorobotics*, 2010, 4 (1), pp.1-8. 10.3389/fnbot.2010.00001 . inria-00467858

HAL Id: inria-00467858

<https://inria.hal.science/inria-00467858v1>

Submitted on 30 Mar 2010

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.

Effectiveness and robustness of robot infotaxis for searching in dilute conditions

Eduardo Martin Moraud^{1,2} and Dominique Martinez ^{*1,3}

¹‘CORTEX’ Research Team, LORIA, France

²Advanced Concepts Team, ESA-ESTEC, The Netherlands

³UMR PISC, INRA, France

Abstract

Tracking scents and locating odor sources is a major challenge in robotics. The odor plume is not a continuous cloud but consists of intermittent odor patches dispersed by the wind. Far from the source, the probability of encountering one of these patches vanishes. In such dilute conditions, a good strategy is to first ‘explore’ the environment and gather information, then ‘exploit’ current knowledge and direct toward the estimated source location. Infotactic navigation has been recently proposed to strike the balance between exploration and exploitation. Infotaxis was tested in simulation and produced trajectories similar to those observed in the flight of moths attracted by a sexual pheromone. In this paper, we assess the performance of infotaxis in dilute conditions by combining robotic experiments and simulations. Our results indicate that infotaxis is both effective (seven detections on average were sufficient to reach the source) and robust (the source is found in presence of inaccurate modeling by the searcher). The biomimetic characteristic of infotaxis is also preserved when searching with a robot in a real environment.

Keywords: Robot, Infotaxis, Olfactory search, Chemical plume, Artificial ethology, Insect behavior

Running title: Robot infotaxis

Supplementary data: infotaxis.mpeg, robot_infotaxis.mpeg, pulsed_infotaxis.mpeg

*Correspondence: Dominique Martinez, ‘CORTEX’ Research Team, LORIA, Campus Scientifique, BP 239, 54506 Vandoeuvre-les-Nancy, France. dominique.martinez@loria.fr

1 Introduction

Olfactory cues are employed by many forms of life to locate food or find mates with a high degree of precision. Moths and bacteria are the most illustrative and well documented examples of navigation strategies under real world conditions. The latter rely on local concentration gradients to direct toward the source of a nutrient [Berg, 1975]. Male moths on the other hand are guided by pheromonal cues to locate their female [Baker et al., 1985, Birch et al., 1990, Mafra-Neto and Carde, 1994]. Upon sensing an odor signal, they surge upwind, since a good estimate of the source direction is provided by the direction of the flow. When odor information vanishes, they exhibit an extended cross-wind casting to perform a local search until the plume is reacquired. Considerable research has been carried out in an attempt to unravel the biological mechanisms that control some of these behaviors and apply them to robotics [Murlis et al., 1992, Vickers, 2000].

Computer or robot-based implementations of biomimetic strategies are relevant not only for testing hypotheses about animal behavior [Belanger and Arbas, 1998], but also for tackling practical problems for which pure engineering solutions are still missing, e.g. finding dangerous substances such as explosives or drugs, or exploring inhospitable environments. Previous robotic attempts have been mainly based on plume tracking –performing a local search within the plume– or chemotaxis –climbing a concentration gradient– [Kowadlo and Russell, 2008]. However, these strategies are effective in dense conditions only, i.e. close to the source where the odor plume can be considered as a continuous cloud. Far from the source, odor dispersal occurs mainly through advection and turbulence dominates. The local concentration gradient hence fluctuates constantly in magnitude and direction and does not always point to the source. The fine-scale structure of the chemical distribution consists of sparse, dispersed and sporadic patches of odors surrounded by wide voids, and the probability of encountering one of these patches decays exponentially with distance [Schraiman and Siggia, 2000]. Tracking scents and locating odor sources in such dilute conditions is a major challenge in current robotics.

Recently, Vergassola et al. proposed a strategy, named ‘infotaxis’, for searching in turbulent environments [Vergassola et al., 2007]. Infotaxis relies on Bayesian inference to maximize information gain, and exploits the scarcity of odor encounters. It involves a period of exploration during which the agent (robot) builds a probabilistic map of the source location, in a similar way to [Pang and Farrell, 2006]. As the agent accumulates information, the map becomes sharper and its entropy –which reflects the uncertainty about the location of the source– decreases. Because the expected search time is determined by the uncertainty of the belief, the robot navigates to maximize the expected reduction in entropy, and therefore the rate of information acquisition. Maximizing information gain entails a competition between two conflicting terms. The exploitative term drives the robot toward locations where the probability of finding the source is high. The explorative term favours motion to regions with lower probabilities of source discovery but high

rewards in terms of information gain. Interestingly, although animal patterns such as ‘casting’ or ‘zigzagging’ are not pre-programmed or imposed through explicit rules of movement, these behaviors do actually emerge naturally from the trade-off between exploration and exploitation at the core of the model [Martinez, 2007].

Promising results were achieved with infotaxis in simulation, even for environmental conditions that consider turbulence. Nevertheless, matching the complexity of the world in simulation has been shown to be extremely difficult [Webb, 2000]. A formal description of the instantaneous structure of the plume in a turbulent flow may for instance require simplifications or assumptions to make the problem tractable. The use of a robot, on the contrary, compels to consider and confront all factors in the environment, yielding complete results. We present hereafter a successful solution for implementing infotaxis within a real robotic system, and we assess its performance in terms of effectiveness and robustness under turbulent conditions. This framework is employed as a testbed to assert complete and rigorous evaluations under real conditions. In addition, we confront infotaxis in simulation to time-varying environments such as the ones used in biological experiments, and thereby further evaluate the biomimetic characteristics pointed out by [Vergassola et al., 2007]. Our evaluation is twofold, and consists of both quantitative analyses of the agent’s propensity to surge upwind or to cast cross-wind, and qualitative interpretations of what compels him to exhibit such behaviors.

2 Materials and Methods

Although infotaxis is fully described in [Vergassola et al., 2007], for the sake of completion the algorithm is sketched in the appendix. Infotaxis relies on the capacity to exploit the finest characteristics of the turbulent medium, i.e. discontinuous odor cues dispersed by the flow. A pre-requisite is that odor ‘cues’ are detected. They may refer either to brief and discrete odor patches or to odor filaments with extended spatiotemporal characteristics.

We tested extensively several gas sensors that are commercially available and none of them suited our needs in terms of response time and sensitivity. They saturate at medium concentrations and require a long-lasting phase of degasing before they can react again. To circumvent the problem, we used a temperature sensor that reacts quickly and does not saturate easily. Note that the transport model of heat is identical to the one of smell in environments where advection clearly dominates over diffusion (high Peclet numbers) [Schraiman and Siggia, 2000] so that the statistical model of the turbulent medium (only the time-averaged concentration is used) described in [Vergassola et al., 2007] may be used for updating the source distribution map. Such a model considers independent detections over time, regardless of previous events. Yet in reality an odor patch or filament covers a certain volume and generates correlated hits as it passes in front of the sensor. In order to ensure that consecutive detections are not overcounted, the posterior probability

distribution of the source is derived from a modified model of turbulent medium which accounts for correlated hits (see Appendix). The model is built from the time intervals of no-detection and from the transitions from no-detection to detection. In our implementation, the transitions occur whenever the sensor signal exceeds an adaptive threshold, whose role is to filter noisy oscillations due to wind or sensor fluctuations. The threshold is derived by averaging sensor readings over two time-steps (40 samples), and adding a constant term established empirically in the absence of stimuli (set to 25 under our environmental conditions - Fig. 1, bottom). Note nevertheless that since the duration of the detections is not taken into account, all the patches are equally considered, irrespective of their size.

Robot infotactic experiments were done with a Koala robot designed by K-Team SA, Switzerland and equipped with an on-board low-level CPU for motion control. The sensor output is sampled at 10Hz, amplified (Amplifier LMC6462 from National Semiconductor) and quantized with the analog to digital converter available on the robot (10 bits of resolution, 5 volts of dynamic range). The heat source had a power of 2000 watts and a fan created a wind speed of 2.5m/s. The wind is assumed to be constant and in the same direction at any time, i.e. local wind speed and direction needed not to be measured by the robot. In order to limit corruption of the temperature sensor by the additional airflow created from the movement of the robot, the motion was implemented as consecutive discrete steps and sensor readings were taken while the robot staying still. Steps of 20 cm were used so as to minimize the effect of discontinuities. The navigation experiments were performed in an arena of 5 meters long by 4 meters large resulting in a grid-based model of the environment of 25×20 points. In order to obtain statistically comparable results, all trials reported hereafter are initialized with the robot located at (10,2) and the source at (9,24). At every step, the agent updates its belief (probability map of the source distribution) according to the history of detection and non-detection events and chooses the best strategy in terms of entropy minimization among the five possible actions, i.e. making a move to one of the four neighboring steps or staying still. The robot is assumed to have reached its goal at one step from the source.

Complementary infotactic simulations were performed in Python. Continuous and pulsed sources were considered (see Appendix). For the continuous case, the parameters are: diffusivity $D = 1$, life-time of particle $\tau = 1.5$ and emission rate of the source $R = 2$, expressed in arbitrary units. Wind speed is set to $V = -2.5$ m/s and the size of the sensor is $a = 1$ cm. These parameter values were established empirically to match the real robotic environment. For the pulsed case, the same values of parameters were used and the source frequency was set to 0.2 Hz for slow pulses (pulse duration=0.2 sec, air gap between pulses = 1.3 sec) and 0.67 Hz for fast pulses (pulse duration=0.2 sec, air gap = 4.8 sec).




Figure 1: Robot infotaxis in action. **Top:** Snapshot of the robot at particular times during its path. **Middle:** Corresponding source distribution maps (belief functions). The blue (resp. red) color code corresponds to low (resp. high) probabilities. The path of the robot from start to current time is superimposed to the map as consecutive red dots when there is no detection, a green dot indicating a detection. Wind blows downwards. The source is at location (9,24), the robot starting point is (10,2). **Bottom:** Detection procedure. The sensor signal, recorded at 10 Hz during the robot path, is shown in red. An adaptive detection threshold, in blue, is derived from the smoothing average, in green. The five detected patches are indicated as black dots.

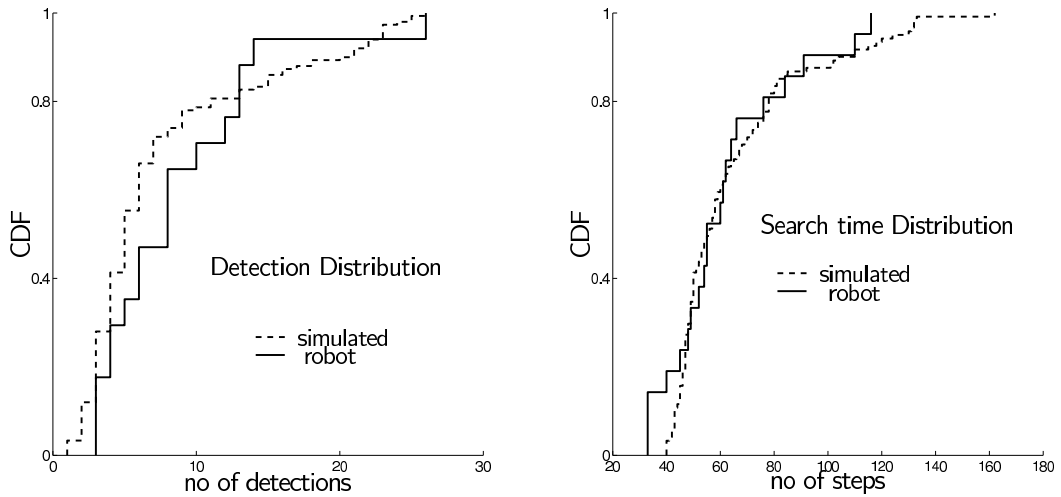


Figure 2: Robot vs simulated infotaxis. **Left:** Cumulative distributions of the number of detections for robot infotaxis (plain curve) and for simulated infotaxis (dashed curve). **Right:** Cumulative distributions of the number of steps (search time) for robot infotaxis (plain curve) and for simulated infotaxis (dashed curve).




Figure 3: **Left:** Examples of robot trajectories. The source is continuous and located at (9, 24). The robot starting point is (10, 2). Wind blows downwards. **Right:** track angle histogram. Track angles were computed as movement vectors with respect to the wind direction. A peak at zero degrees indicates a tendency to move upwind, as compared to movements perpendicular to the wind direction ($\pm 90^\circ$).

3 Results

To test the effectiveness of robot infotaxis for searching in dilute conditions, we performed 21 robot runs (see one example in Fig. 1) and compared them with 150 simulations of Infotaxis. For 20 trials out of 21, the robot was able to reach the source within a reasonable time limit of 150 steps, above which the robot is considered to be lost. This case occurred in one trial only, during which too many detections persuaded the agent that the source was already found, i.e. exploitation was predominantly compelling the robot to stay in its current location rather than exploring further and gather information. For the successful runs, the number of detections was low (7.53 ± 6.42 , mean \pm s.d.), reflecting the dilution condition of the experiments. The cumulative distribution of the number of detections for robot infotaxis, plotted in Fig. 2 (left), is not statistically different from the one obtained with simulated infotaxis (8.59 ± 5.85 , mean \pm s.d.; two-sample Kolmogorov-Smirnov test $p=0.26$). The search time cumulative distributions between simulated and robot infotaxis were also not different (Fig. 2 right, two-sample Kolmogorov-Smirnov test $p=0.73$). The search time distribution is well described by a gamma distribution with shape and scale parameters 8.5 and 7.5 respectively. From Fig. 3 (left), we note that robot trajectories are similar to those obtained in simulations, e.g. biomimetic patterns such as ‘extended crosswind casting’ or ‘zigzagging upwind’ typical of moth flight emerge naturally from the trade-off between exploration and exploitation [Vergassola et al., 2007]. The track angle histogram of robot infotaxis shown in Fig 3 (right) presents a peak at zero degrees ($p=0.001$, Rao’s circular test of non-uniformity), indicating a predominance of the robot to move upwind. Unimodal track angle histograms with mode at zero are also representative of moths flying upwind in turbulent plumes [Mafra-Neto and Carde, 1994, Lei et al., 2009].

Infotactic simulations and robotic experiments described above were done with a continuous source. Females of several moth species however are known to rhythmically extrude their pheromone glands [Baker et al., 1985]. To assess the capacity of infotaxis to cope with real conditions such as those faced in biology, we considered simulations with a pulsed source model that rhythmically releases odor patches in the environment. Note that there is a mismatch between the pulsed source generator and the continuous source model used to update the internal beliefs in Infotaxis. We performed 60 repeated simulation runs, 30 with fast pulses and 30 with slow pulses. For all trials, the agent was able to reach the source within a time limit of 200 steps. Typical infotactic trajectories under fast and slow pulsed conditions are shown in Fig. 4 (left) and the percentage of wind-oriented movements per trajectory in the two regimes in Fig. 4 (right). Percentages of downwind movements were low and not significantly different in both conditions, reflecting the high success rate of the searcher. We found however that the searcher moves mainly upwind in the fast pulsed condition and crosswind in the slow pulsed condition. We further investigated this behavioral difference by looking at the updates of the beliefs under both conditions. With fast pulses (Fig 5A), intervals of no-detection between pulses are




Figure 4: Left: Examples of simulated infotactic trajectories with incorrect modelling by the searcher. The searcher considers a continuous source model whereas the actual source is pulsed (fast and slow pulsed conditions). Source location indicated as a black dot is (20,49). The searcher initial location is (6,3). Detections are marked as red dots. Wind blows downwards. **Right:** Percent wind-oriented movements under the fast and slow pulsed conditions. Movements having no letters in common are significantly different at $P < 0.05$ (Kruskal-Wallis test followed by pairwise comparisons, $n=30$ trials per condition).

short enough to keep the searcher on exploitation. Each update sharpens the posterior distribution. The high probability bump emerging in the wind direction induces the agent to move upwind. With slow pulses (Fig 5B), unexpected long periods of time with no odor encounter broaden the posterior distribution, hence compelling the agent to counterturn and explore the environment in large spirals.

Three videos are appended as supplementary material and illustrate these evaluations. They cover the foundations of infotaxis, its robotic implementation and its efficiency when confronted to a pulsed source (infotaxis.mpeg, robot_infotaxis.mpeg, pulsed_infotaxis.mpeg).

4 Discussion

The fundamental aspect of infotaxis is to exploit the finest characteristics of the turbulent medium, i.e. discontinuous odor patches dispersed by the flow. A requirement is thus to be capable of resolving single cues and to exploit them to guide the search. Output neurons in the phomonal system of the moth have been shown to respond to pulses of pheromone delivered at a rate up to 10 Hz. When their capacity to follow pheromone pulses is pharmacologically disrupted, moths do not navigate successfully toward the source [Lei et al., 2009]. Likewise, robot infotaxis did not succeed in the search for odor sources because the transient response of gas sensors is too slow to track individual odor patches above 1 Hz (data not shown). As the diffusion model of heat is identical to the one of smell, we used instead a fast temperature sensor and a heat source. In our experiments, infotaxis led to effective searching and did not require fine tuning of parameters to work in the real environ-




Figure 5: Navigation patterns observed when infotaxis is confronted to a pulsed source (fast pulses in A, slow pulses in B). Top rows represent snapshots of the simulated environment at different times (pulsed source located in $\mathbf{r}_0 = (25, 2)$, wind blows downwards) and bottom rows are for the corresponding source distribution (belief function). False blue and red colors correspond to low and high probabilities, respectively. The path of the robot is superimposed to the map as consecutive red dots when there is no detection, and green dots for detections. **Fig. A:** High-frequency pulsed patches provoke new detections before the robot starts spiraling. High probabilities are frequently updated and assigned to upwind locations (at times t_1, t_3 and t_5) hence pushing the agent forward. **Fig. B:** Between pulses, long intervals of clean air –during which no detections arise– compel the agent to explore regions where previous detections were recorded. Probabilities updates take the form of concentric ellipses that spread as the robot navigates around them, as clearly seen at times t'_2, t'_3 and t'_4 . (Note that such behavior –although with much smaller radius– may also be recorded for the fast-pulsed case when the agent is close to the source, due to an excessive amount of detections that push him to switch to exploitation mode).

ment. Seven detections on average and as few as three in many trials were sufficient to reach the source, reflecting the dilute condition of the experiments. The number of detections as well as the search time were not statistically different between simulated and robot infotaxis (see Fig. 2).

We found however that the search was influenced by the frequency at which the source was pulsed. The searcher moved mainly upwind in the fast-pulsed condition and crosswind in the slow-pulsed case (Fig. 4, right). Such a frequency modulated behavior is in agreement with biological observations. Experiments with a puffing device revealed that upwind flights of moths were sustainable in fast but not slow pulsed plumes [Mafrá-Neto and Carde, 1994, Vickers and Baker, 1994]. It has been emphasized that a tempo of cues above a certain frequency is needed to ensure odor encounters before the moth undertakes a counter-turning or casting behavior.

Our infotactic evaluations illustrate that the behavior of the searcher not only depends on the detections made in the past, but also on the expectations derived from his internal belief. Note from Fig. 5A that every single detection updates the agent's belief in a way that pushes him forward as a first step. When the information provided by the cues cannot be further exploited to drive upwind, spiraling becomes a better strategy (as in Fig. 5B) unless a new cue is detected. The correlation between both the external detection rate and the internal expectations leads to efficient trajectories. Fast frequencies in Fig. 5A for instance bring new cues with a tempo that matches better the searcher's internal expectations (i.e. before he switches to exploration mode), and iteratively updates his belief so that the behavior is targeted toward the goal. Resulting trajectories are mainly straight and upwind, except close to the source where too many detections persuade the agent that the source is already found, hence producing short-length zigzags and very localized spirals. On the contrary, in the slow-pulsed condition, unexpected long time-periods without odor encounters result from a mismatch which compels the agent to counter-turn and explore in large spirals as the probability map gets broadened. In both cases, a continuous belief model is assumed.

The question of whether moths do actually employ infotactic strategies as described is yet of a more delicate nature. The assumption that the agent constructs a detailed grid-based map of his environment is very computationally expensive, and requires the robot to be acquainted with the size and shape of the arena. Biological experiments have reported however that similar principles account for navigation strategies in rats. These mammals employ internal spatial maps that combine both a topographical description of the environment (encoded by 'grid cells') and location specific information ('place cells') [Hafting et al., 2005], in a similar way to the grid-model of the arena used in Infotaxis which gets updated as localized cues are detected. Yet simpler descriptions of the environment could also be considered, for which infotactic strategies may prove equally efficient.

5 Conclusion

Previous experiments from the same laboratory revealed that a concentration gradient can be extracted from a turbulent plume in dense conditions when the robot moves slowly (2.5 cm/sec) and near to the source (search area=2.9 m²), see e.g. fig. 4 left in [Martinez et al., 2006]. To move the robot in the vicinity of the source, previous works considered the possibility of exploring the environment by using vision, in addition to olfaction, e.g. [Martinez and Perrinet, 2002]. The main limitation is that odor source candidates need to be identifiable from visual features. Here we tackled the more difficult problem of searching in dilute conditions, for which only the information released by the source and wind direction were used. We implemented infotaxis on a real robot and considered a larger search area (20 m²) and faster robot speed (14 cm/sec) than in our previous experiments.

The robustness of infotaxis was evaluated with respect to inaccurate modeling by the agent. The parameters employed internally to guide the search were not fine-tuned or adapted over time, and could differ from the instantaneous characteristics of the surrounding. Despite this discrepancy, the robot was able to reach the source within a reasonable time limit and produced very few downwind movements (Fig. 4, right). This emphasized the capability of the model to cope with the unpredictability imposed by real environments. Additional analyses with a pulsed source pointed out that frequencies at which cues are encountered account for the efficiency of the strategy followed, just as reported in the case of moths, which depends on how well detection rates match the internal map used by the searcher.

The extension of Infotaxis to cope with stereo sensing capabilities, just as in the case of insects, may prove beneficial in terms of effectiveness and ought to be considered as future work. Two sensors, employed in parallel to update the probability map, may indeed increase directionality. Further applications of Infotaxis to collective search have also been reported recently by [Masson et al., 2009] with impressive gains in search times.

References

- [Baker et al., 1985] Baker, T., Willis, M., Haynes, K., and Phelan, P. (1985). A pulsed cloud of sex pheromone elicits upwind flight in male moths. *Physiol. Entom.*, 10:257–265.
- [Belanger and Arbas, 1998] Belanger, J. and Arbas, E. (1998). Behavioral strategies underlying pheromone-modulated flight in moths: lessons from simulation studies. *J. Comp. Physiol. A*, 183:345–360.
- [Berg, 1975] Berg, H. (1975). Bacterial behaviour. *Nature*, 254:389–392.
- [Birch et al., 1990] Birch, M. C., Poppy, G. M., and Baker, T. (1990). Scents and eversible scent structures of male moths. *Physiol. Entom.*, 35:25 – 58.

- [Hafting et al., 2005] Hafting, T., Fyhn, M., Molden, S., Moser, M.-B., and Moser, E. I. (2005). Microstructure of a spatial map in the entorhinal cortex. *Nature*, 436:801–806.
- [Kowadlo and Russell, 2008] Kowadlo, G. and Russell, R. A. (2008). Robot odor localization: A taxonomy and survey. *Int. J. Rob. Res.*, 27(8):869–894.
- [Lei et al., 2009] Lei, H., Riffel, J., Gage, S., and Hildebrand, J. (2009). Contrast enhancement of stimulus intermittency in a primary olfactory network and its behavioral significance. *J. Biology*, 8(21).
- [Mafra-Neto and Carde, 1994] Mafra-Neto, A. and Carde, R. T. (1994). Fine-scale structure of pheromone plumes modulated upwind orientation of flying moths. *Nature*, 369:142 – 144.
- [Martinez, 2007] Martinez, D. (2007). On the right scent. *Nature*, 445:371–372.
- [Martinez and Perrinet, 2002] Martinez, D. and Perrinet, L. (2002). Cooperation between vision and olfaction in a koala robot. *Report on the 2002 Workshop on Neuromorphic Engineering*, pages 51–53.
- [Martinez et al., 2006] Martinez, D., Rochel, O., and Hugues, E. (2006). A biomimetic robot for tracking specific odors in turbulent plumes. *Autonomous Robots, Special Issue on Mobile Robot Olfaction*, 20(3):185–195.
- [Masson et al., 2009] Masson, J.-B., Bailly-Bechet, M., and Vergassola, M. (2009). Chasing information to search in ransom environments. *J. Phys. A: Math. Theor.*, 42:434009.
- [Murlis et al., 1992] Murlis, J., Elkinton, J. S., and Card, R. T. (1992). Odor plumes and how insects use them. *Annual Review of Entomology*, 37:479–503.
- [Pang and Farrell, 2006] Pang, S. and Farrell, J. (2006). Chemical plume source localization. *IEEE Transactions on Systems, Man, and Cybernetics, Part B: Cybernetics*, 36(5):1068–1080.
- [Schraiman and Siggia, 2000] Schraiman, B. I. and Siggia, E. D. (2000). Scalar turbulence. *Nature*, 405:639–646.
- [Smoluchowski, 1917] Smoluchowski, M. (1917). Versuch einer mathematischen theorie des koagulationskinetic kolloider losungen. *Z. Phys. Chem*, 92:129–168.
- [Thrun et al., 2005] Thrun, S., Burgard, W., and Fox, D. (2005). *Probabilistic Robotics (Intelligent Robotics and Autonomous Agents)*. The MIT Press.
- [Vergassola et al., 2007] Vergassola, M., Villermaux, E., and Shraiman, B. I. (2007). Infotaxis as a strategy for searching without gradients. *Nature*, 445(7126):406–409.

[Vickers, 2000] Vickers, N. (2000). Mechanisms of animal navigation in odor plumes. *Biol. Bull.*, 198:203–212.

[Vickers and Baker, 1994] Vickers, N. and Baker, T. (1994). reiterative responses to single strands of odor promote sustained upwind flight and odor source location by moths. *PNAS*, 91:5756–5760.

[Webb, 2000] Webb, B. (2000). What does robotics offer animal behaviour? *Animal Behaviour*, 60:545–558.

Acknowledgements

The authors acknowledge E. Villermaux for helpfull discussions and for suggesting the use of a turbulent plume of heat in robotic experiments.

Appendix: Infotaxis

Infotaxis is fully described in [Vergassola et al., 2007]. For completion, we detail its core modules in terms of probabilistic robotics [Thrun et al., 2005] as employed in our robot implementation. The model combines a belief function –the robot internal knowledge about his environment, updated as cues are encountered– along with decision-making –execution of an action that maximizes a reward–.

Statistical model of the odor plume and grid-based map of the environment:

In Infotaxis, the robot is provided with a statistical description of the odor plume that he uses to infer the probability that the source be located at any point of his internal grid-based probability map of the environment. The statistical description of the odor plume is derived from the resolution of the following advection-diffusion equation

$$D\nabla^2 U(\mathbf{r}) + \vec{V} \cdot \nabla U(\mathbf{r}) - \frac{1}{\tau} U(\mathbf{r}) - R\delta(\mathbf{r} - \mathbf{r}_0) = 0 \quad (1)$$

for an odor source located at r_0 and emitting ‘particles’ or patches’ at a rate R . The particles propagate with diffusivity D , have a mean lifetime described by τ and are advected by a mean current or wind V . $U(\mathbf{r})$ is the local concentration at location r and δ is the Dirac delta function. In such an environment, the mean frequency of odor encounters with a spherical sensor of radius ‘ a ’ follows the Smoluchowski’s [Smoluchowski, 1917] expression

$$R(\mathbf{r}) = 4\pi Da \cdot U(\mathbf{r}) \quad (2)$$

This model provides a framework by which to take into consideration the geometry of the environment when navigating. It can be easily solved through numerical

methods and makes it possible for autonomous robots to iteratively infer knowledge about their surrounding. In the *continuous* case, the solution to Eq. 1 writes:

$$U(\mathbf{r}|\mathbf{r}_0) = \frac{1}{4\pi D|\mathbf{r}-\mathbf{r}_0|} e^{-\frac{V}{2D}(y-y_0)} e^{-\frac{|\mathbf{r}-\mathbf{r}_0|}{\lambda}} \text{ where } \lambda = \sqrt{\frac{D\tau}{(1+\frac{V^2\tau}{4D})}} \quad (3)$$

For the non-continuous case, i.e. under the influence of a *pulsed* odor source $R(t)$ at location r_0 , we derived such function by solving the non-homogeneous diffusion-advection equation which, for the two-dimensional problem with mean vertical wind V_y , takes the form:

$$U(\mathbf{r}, t|\mathbf{r}_0) = e^{-\frac{V}{2D}(y-y_0) - \frac{1}{\tau}(1+\frac{V^2\tau}{4D})t} \left[\frac{e^{-\frac{|\mathbf{r}-\mathbf{r}_0|}{4Dt}}}{4\pi Dt} * R(t) e^{\frac{1}{\tau}(1+\frac{V^2\tau}{4D})t} \right] \quad (4)$$

Belief function:

Let us consider the trace $\Gamma_t = \{(r_1, t_1), (r_2, t_2) \dots (r_n, t_n)\}$ of the hits (odor encounters) experienced by the searcher at locations $x_1 \dots x_n$ and times $t_1 < \dots < t_n < t$ during the path from its start to the current time t . The belief function is given by the posterior probability for the source to be located in \mathbf{r}_0 given the trace Γ_t

$$P_t(\mathbf{r}_0|\Gamma_t) = \frac{\mathcal{L}(\Gamma_t|\mathbf{r}_0)}{\int \mathcal{L}(\Gamma_t|\mathbf{r}_x) d\mathbf{r}_x} \quad (5)$$

where $\mathcal{L}(\Gamma_t|\mathbf{r}_0)$ is the likelihood of experimenting the trace Γ_t for a source in \mathbf{r}_0 .

In our robot implementation of infotaxis, consecutive detections are not considered as they belong to a same patch and are correlated (see Materials and Methods). We therefore employ the following likelihood function (from Eq. 13 in supplementary materials of [Vergassola et al., 2007]) instead of the one considered in the original algorithm

$$\mathcal{L}(\Gamma_t|\mathbf{r}_0) = e^{-\sum_i \int_{V_i} R(\mathbf{r}(t')|\mathbf{r}_0) dt'} \prod_{i=1}^T R(\mathbf{r}(t_i), t_i|\mathbf{r}_0) \quad (6)$$

in which T represents the *transitions* from no-detection to detection, i.e. new patches, and the V_i 's are the time intervals of absence of detection. Note from Eq. 6 that the absence of correlations permits to update the probability map without storing the whole history. Indeed, $P_{t+\Delta t}(\mathbf{r}_0) = P_t(\mathbf{r}_0) \cdot \text{update}_{\Delta t}(\text{nparticules})$, where 'nparticules' is the number of detections that the searcher experienced during the short time interval Δt . Memory requirements are therefore kept to a minimum.

Decision-Making:

For the decision-making, the robot moves in the direction that minimizes its *local* uncertainty about the location of the source. The expected reduction of entropy

–reward function– for the robot moving from \mathbf{r}_i to \mathbf{r}_j , consists of two terms:

$$\Delta H(\mathbf{r}_i \rightarrow \mathbf{r}_j) = P_t(\mathbf{r}_j)(0 - H) \quad (7)$$

$$+ [1 - P_t(\mathbf{r}_j)] \Delta S \quad (8)$$

The first term (7) evaluates the reduction of entropy if the source is found at the next step. Reaching the source in \mathbf{r}_j occurs with estimated probability P_t and the entropy goes from H to 0. The second term (8) corresponds to the reduction of entropy if the source is not found. It occurs with probability $1 - P_t$ and ΔS represents the information gain in \mathbf{r}_j coming from expected odor encounters. The first term is seen as exploitative as it drives the searcher toward locations where the probability of finding the source is high. The second term is explorative as it compels the searcher toward regions with lower probabilities of source discovery but high information gains.

The *expected* reduction of entropy in the case where the source is not found derives from the probability sum of experiencing i new detections during the movement, where encounters are modeled by means of a Poisson-distributed random variable ρ_i

$$\Delta S = \rho_0 \Delta S_0 + \rho_1 \Delta S_1 + \rho_2 \Delta S_2 \dots \quad (9)$$

therefore accounting for all possible cases that new information is detected along the way (either 1, 2 or n encounters).