

LETTERS

'Infotaxis' as a strategy for searching without gradientsMassimo Vergassola¹, Emmanuel Villermaux² & Boris I. Shraiman³

Chemotactic bacteria rely on local concentration gradients to guide them towards the source of a nutrient¹. Such local cues pointing towards the location of the source are not always available at macroscopic scales because mixing in a flowing medium breaks up regions of high concentration into random and disconnected patches. Thus, animals sensing odours in air or water detect them only intermittently as patches sweep by on the wind or currents^{2–6}. A macroscopic searcher must devise a strategy of movement based on sporadic cues and partial information. Here we propose a search algorithm, which we call 'infotaxis', designed to work under such conditions. Any search process can be thought of as acquisition of information on source location; for infotaxis, information plays a role similar to concentration in chemotaxis. The infotaxis strategy locally maximizes the expected rate of information gain. We demonstrate its efficiency using a computational model of odour plume propagation and experimental data on mixing flows⁷. Infotactic trajectories feature 'zigzagging' and 'casting' paths similar to those observed in the flight of moths⁸. The proposed search algorithm is relevant to the design of olfactory robots^{9–11}, but the general idea of infotaxis can be applied more broadly in the context of searching with sparse information.

Chemotactic search strategies based on local concentration gradients require the concentration to be high enough to ensure that its average difference measured at two nearby locations is larger than typical fluctuations^{1,12} (see also Supplementary Information). The signal-to-noise ratio depends of course on the averaging time and might be improved by waiting. However, the average concentration may be decaying rapidly (for example, exponentially) with distance away from the source, and in this weak signal-to-noise (dilute) case the waiting time becomes huge. An example of organisms performing olfactory search in a dilute limit is provided by moths which use pheromones to locate their mates^{2–6}. Moths are known to proceed upwind by way of counterturning patterns of extended ('casting') or limited ('zigzagging') crosswind width, thought to correlate with low and high rates of odour detection. A practical situation involving the challenge of searching in dilute conditions is encountered in the design of 'sniffers'^{9–11}—robots that track chemicals emitted by drugs, chemical leaks, explosives and mines. Existing methods apply to high-concentration conditions, where chemotactic^{13–16} or plume-tracking strategies^{17–21} might be used.

In the dilute limit, the searcher detects odour in a sporadic sequence of distinct events arising from its encounters with patches of fluid (or air) where turbulent mixing has failed to dissipate the advected odour down to a level below the detectability threshold^{22–24}. These detection events, or 'hits', are separated by wide 'voids' with no detectable signal. Because the probability of odour encounter depends on the distance from the source, the set of encounters that occurred at times $\{t_i\}$ along the search trajectory $\mathbf{r}(t)$ carries informa-

tion about the source location. We shall use \mathcal{T}_t to denote times and coordinates of these hits.

In the spirit of coding theory, the trace \mathcal{T}_t might be thought of as a message, sent by the source and transmitted to the searcher with strong noise due to the random nature of odour propagation in the turbulent medium. Decoding of the message is implemented using Bayes' formula to construct, given the received signal, the posterior probability distribution $P_t(\mathbf{r}_0)$ for the unknown location of the source \mathbf{r}_0 (see Methods, and similar independent arguments in ref. 25). The subscript t reminds us that \mathcal{T}_t and $P_t(\mathbf{r}_0)$ are dynamical objects, continuously updated with time. The specific decoding protocol depends of course on the nature of the detection events and the transmitting medium. For concreteness, we treat here two cases: (1) experimental time-course data for a mixing flow; and (2) a model where detectable 'particles' (which represent patches of detectable odours) are emitted by the source at rate R , have a finite lifetime τ , propagate with effective diffusivity D and are advected by a mean current or wind \mathbf{V} . The decoding protocol requires knowing the probability of odour encounters as a function of the distance to the source. This function can be computed analytically for model (2) or estimated from experimental data for case (1), as detailed in Supplementary Information. The latter method is quite general and might be applied to other cases as well.

Given a probability distribution $P(\mathbf{r}_0)$ for the location of the source, we can show (see Supplementary Information) that the expected search time $\langle T \rangle$ is bounded by $\langle T \rangle \geq e^{S-1}$, where S is Shannon's entropy for the distribution $S \equiv - \int d\mathbf{x} P(\mathbf{x}) \ln P(\mathbf{x})$ (refs 26, 27). The latter quantifies how spread-out the distribution is, and goes to zero when the position of the source is localized to one site, that is, is known. The rate of acquisition of information is quantified by the rate of reduction of entropy^{26,27} (see also Supplementary Information). The main problem for the searcher is that the real probability distribution is unknown (to it) and must be estimated from the available data: the history of its odour encounters. As information accumulates, the entropy of the estimated distribution decreases and with it the expected time to locate the source. The searcher is faced with conflicting choices of either proceeding with its current information (that is, going to the estimated most probable source location), or alternatively, pausing to gather more information and obtain a more reliable estimate of the source distribution. The problem of dealing with only partially reliable information is quite general, and has received a systematic formulation in learning theory in terms of the 'exploration versus exploitation trade-off' to be struck for effective learning²⁸. In the search context, 'exploitation' of the currently estimated $P_t(\mathbf{r}_0)$ by chasing locations of maximal estimated probability is very risky, because it can lead off the track. The most conservative 'exploration' approach is to accumulate information before taking any step. This strategy is safe but not productive,

¹CNRS URA 2171, Institut Pasteur, "In Silico Genetics", 75724 Paris Cedex 15, France. ²Université Aix Marseille 1, IRPHE, Technopole Chateau Gombert F-13384 Marseille, France. ³Kavli Institute for Theoretical Physics, University of California, Santa Barbara, California 93106, USA.

and is inferior to more active exploration—for example, systematic search in a particular sector²⁹.

To balance exploration and exploitation, we propose the following ‘infotaxis’ strategy. At each time step, the searcher chooses the direction that locally maximizes the expected rate of information acquisition. Specifically, the searcher chooses, among the neighbouring sites on a lattice and standing still, the move that maximizes the expected reduction in entropy of the posterior probability field.

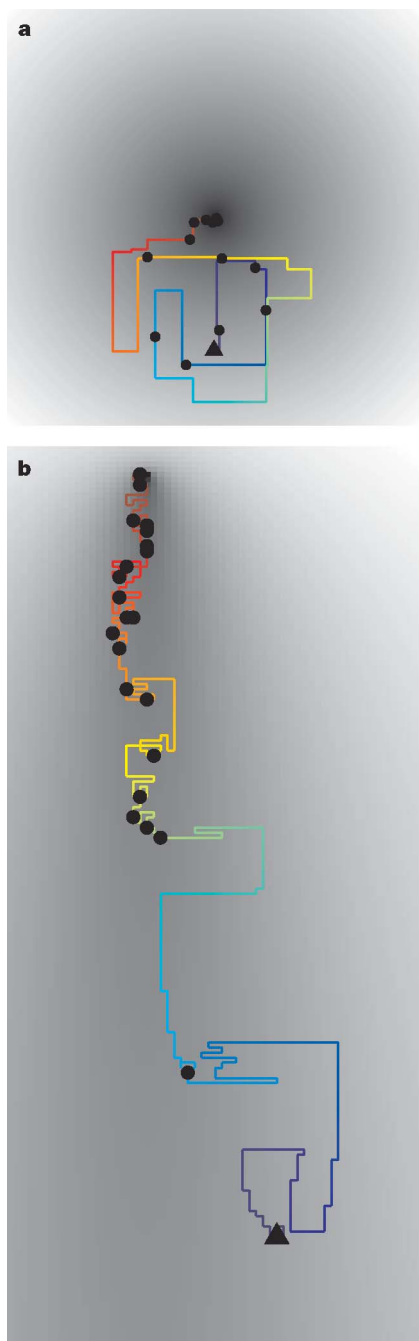


Figure 1 | Typical infotactic trajectories. **a**, Without wind; **b**, with wind. Simulations are performed for a model of odour spreading where detectable ‘particles’ are emitted at rate R , have a lifetime τ , propagate with diffusivity D (combining turbulent and molecular diffusion) and are advected by a mean wind \mathbf{V} . The wind in **b** is directed downwards. The greyscale represents the mean detection rate, decaying exponentially at large distances. In each panel, the searcher starts from the black filled triangle, the colour code on the trajectories is linear in the elapsed time, and odour detections are indicated by black filled circles. Note the long lags with no particle detections, characteristic of searches in dilute conditions.

Expectations are based on the information currently available, that is, the field $P_t(\mathbf{r}_0)$ itself. The intuitive idea is that entropy decreases (and thus information accumulates) faster close to the source because cues arrive at a higher rate, hence tracking the maximum rate of information acquisition will guide the searcher to the source much like concentration gradients in chemotaxis.

Suppose that the searcher has arrived at \mathbf{r} at time t , and gathered information is stored into the field $P_t(\mathbf{r}_0)$ having entropy S . The variation of entropy expected upon moving to one of the neighbouring points \mathbf{r}_j (or standing still) is:

$$\overline{\Delta S}(\mathbf{r} \rightarrow \mathbf{r}_j) = P_t(\mathbf{r}_j)[-S] + [1 - P_t(\mathbf{r}_j)][\rho_0(\mathbf{r}_j)\Delta S_0 + \rho_1(\mathbf{r}_j)\Delta S_1 + \dots] \quad (1)$$

The first term on the right-hand side corresponds to finding the source, that is, P_{t+1} becoming a δ -function and entropy becoming zero, which occurs with estimated probability $P_t(\mathbf{r}_j)$. The second term on the right-hand side corresponds to the alternative case when the source is not found at \mathbf{r}_j . Symbols $\rho_k(\mathbf{r}_j)$ denote the probability that k detections be made at \mathbf{r}_j during a time-step Δt , given by a Poisson law $\rho_k = h^k e^{-h}/k!$ for independent detections. The expected number of hits is estimated as $h(\mathbf{r}_j) \equiv \Delta t \int P_t(\mathbf{r}_0)R(\mathbf{r}_j|\mathbf{r}_0)d\mathbf{r}_0$, with $R(\mathbf{r}|\mathbf{r}_0)$ denoting the mean rate of hits at position \mathbf{r} if the source is located in \mathbf{r}_0 (see Methods). The symbols ΔS_k in equation (1) denote the change of entropy between the fields $P_{t+1}(\mathbf{r}_0)$ and $P_t(\mathbf{r}_0)$. Two effects contribute to ΔS_k : first, $P_{t+1}(\mathbf{r}_j) \equiv 0$ because the source was not found; and second, the estimated posterior probabilities are modified by the k cues received. The first term on the right-hand side of equation (1) is the exploitative term, weighing only the event that the source is found at the point \mathbf{r}_j and favouring motion to maximum likelihood points. The second term on the right-hand side of equation (1) is the information gain from receiving additional cues. It appears even when the searcher does not move, and thus represents conservative ‘exploration’. Thus we explicitly see that infotaxis naturally combines exploitative and exploratory tendencies (see Supplementary Information for details of this point and for quantitative comparisons among different strategies).

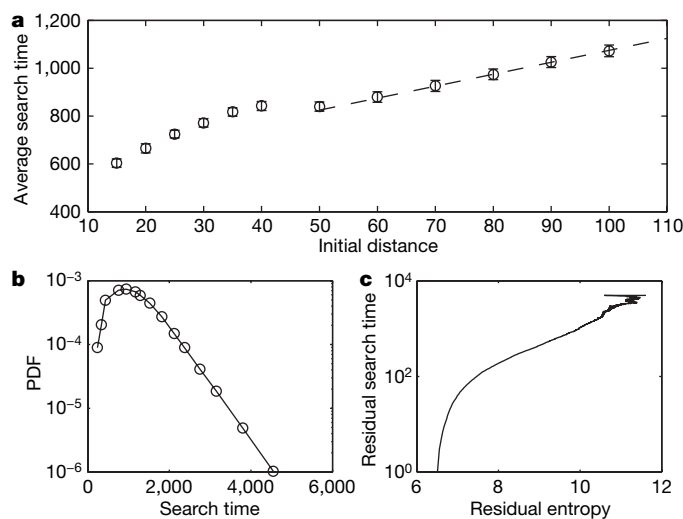


Figure 2 | Quantitative characterization of infotaxis searches. **a**, Scaling of the average search time with the initial distance to the source. The mean path of particles during their lifetime is 50. The linear scaling at large distances compares favourably with the exponential time needed to average out concentration noise. Error bars indicate s.d. **b**, The exponential decay of the search time probability distribution function (PDF), indicating that infotaxis is not plagued by strong fluctuations. **c**, The residual time to locate the source plotted against the entropy of the estimated source location PDF. The exponential dependence of the residual time indicates that reducing entropy is effective in ensuring a rapid search process.

Infotactic paths are illustrated in Fig. 1, which presents the result of a numerical simulation using the model of odour propagation described in Methods. In the absence of wind (Fig. 1a), $P_t(\mathbf{r})$ is rotationally symmetric around the starting point and the searcher starts spiralling around it (as is observed with sea urchin sperm³⁰). Interestingly, in the absence of hits the radius of the spiral increases in a scale invariant manner, making an approximately Archimedean spiral. As time progresses and information is gathered along the trajectory, the direction towards the source emerges as the preferential one, finally leading the searcher to the source. In the presence of wind (Fig. 1b), the search alternates phases of consistent progression upwind with phases of wider crosswind excursion and even downwind movements similar to the classical casting and zigzagging patterns observed during bird and moth flights⁸.

To quantify the performance of the proposed search algorithm, we examine (Fig. 2a) the scaling of the average search time with the initial distance to the source for the model without wind (the most difficult one). The linear dependence found for large initial distances should be contrasted with the exponential dependence needed to average out the concentration noise (see Supplementary Information). Furthermore, random walking searchers would often attain the boundaries of the box where the search is taking place. The corresponding probability distribution of search times would have a long tail, decaying as $1/T^2$, contrary to the exponential decay shown in Fig. 2b. Figure 2c shows the relation between search times and the entropy of the posterior field $P_t(\mathbf{r}_0)$, which supports the theoretical bound mentioned above.

Figure 3 presents an infotactic path generated in a simulation using experimental measurements of dye concentration in a turbulent flow⁷. Hits occur when the searcher encounters concentration above a threshold, which we chose sufficiently high to keep the number of hits low. Simulations indicate that the infotactic strategy is robust with respect to the searcher's model of the turbulent medium and to fluctuations and inhomogeneities of the medium. Indeed, even the simplistic hypothesis of time-independent odour encounters does not hinder the search. Modelling of the turbulent medium might be further improved by accounting for temporal correlations of odour plume encounter and velocity fluctuations (see Supplementary Information).

We have presented an olfactory search algorithm that works in the dilute limit corresponding to weak sources in realistic flows. These are the conditions encountered in applications of olfactory robots and by various living creatures. When comparing our results with the behaviour of living creatures, we need to bear in mind that similarity in trajectories does not imply identity of causal mechanisms and decision processes^{10,15}. Still, it is worth remarking that the olfactory search motion observed with moths and birds exhibits a pattern of extended crosswind casts and zigzags similar to the one generated in Fig. 1b by the general principle of local maximization of information gain. We note that the dilute limit also describes the case of 'odour' diffusion at the molecular scale provided that the searcher can detect single molecules. This situation may apply, for example, in the case of sea urchin sperm, which responds to a single molecule of resact, a chemoattractant peptide³⁰. It is difficult to imagine a single cell

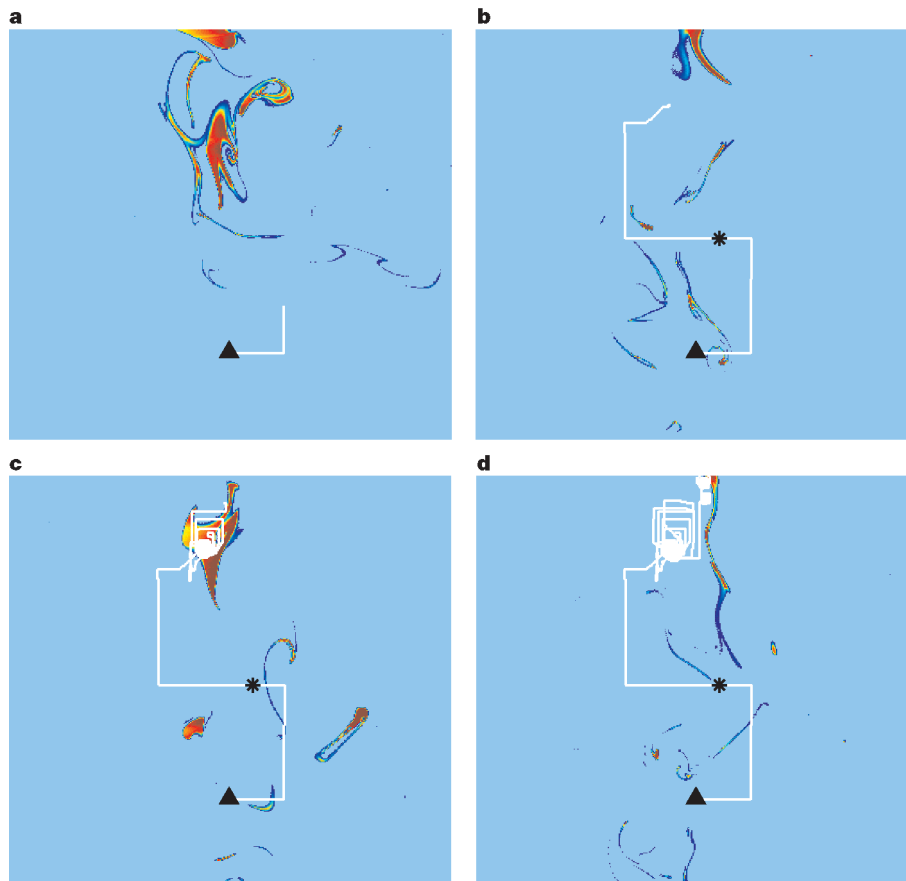


Figure 3 | Simulation of infotaxis using mixing flow experimental data⁷. Snapshots (false colours; red corresponds to high concentrations) of dye concentration levels are acquired from mixing flow experimental data, and the trajectory of the searcher is numerically simulated in the resulting sequence of fields. Snapshots and trajectories are shown at four successive times in panels a–d. Light blue regions correspond to concentrations below

the detection threshold. The black star denotes an isolated odour detection event. The infotactic strategy is robust with respect to fluctuations and inhomogeneities of the medium (for example, of the wind direction) and its modelling by the searcher. Robustness stems from the tracking of information rather than estimated maximum likelihood locations.

performing complex computations like those required for infotaxis, as described above. It will be interesting to explore heuristic approximations simple enough to be plausible for single cells. Finally, we note that the general information-theoretic approach described above in the olfactory search context applies more broadly to any situation where competing demands of exploration and exploitation must be efficiently balanced.

METHODS

Estimation of the posterior probability distribution. The probability distribution posterior to experiencing a trace \mathcal{T}_t of uncorrelated odour encounters is given by:

$$P_t(\mathbf{r}_0) = \frac{\mathcal{L}_{\mathbf{r}_0}(\mathcal{T}_t)}{\int \mathcal{L}_{\mathbf{x}}(\mathcal{T}_t) d\mathbf{x}} = \frac{\exp\left[-\int_0^t R(\mathbf{r}(t'))|\mathbf{r}_0| dt'\right] \prod_{i=1}^H R(\mathbf{r}(t_i)|\mathbf{r}_0)}{\int \exp\left[-\int_0^t R(\mathbf{r}(t')|\mathbf{x}) dt'\right] \prod_{i=1}^H R(\mathbf{r}(t_i)|\mathbf{x}) d\mathbf{x}} \quad (2)$$

Here, H is the number of hits along the trajectory, the t_{ts} are the corresponding times and $\mathcal{L}_{\mathbf{r}_0}(\mathcal{T}_t)$ is the likelihood of observing the trace \mathcal{T}_t of odour encounters for a source located at \mathbf{r}_0 . This expression is supplemented by the prescription that visited regions where the source was not found have zero probability. Note that $P_{t+\Delta t}(\mathbf{r}_0)$ factorizes as $P_t(\mathbf{r}_0)$ times a term that depends on the hits received in the Δt interval. Thus, keeping track of the whole trajectory and the history of detections is not required. The expression for $P_t(\mathbf{r}_0)$ is derived by taking the mean 'hit' rate during an infinitesimal interval dt to be $R(\mathbf{r}|\mathbf{r}_0)dt$ and the number of hits to be Poisson distributed. The function $R(\mathbf{r}|\mathbf{r}_0)$ appearing in equation (2) denotes the mean rate of hit encounters at position \mathbf{r} for a source located at \mathbf{r}_0 . For the model where detectable 'particles' are emitted by the source at rate R , have a finite lifetime τ , propagate with isotropic effective diffusivity D (which parameterizes the combined effect of turbulent and molecular diffusion) and are (possibly) advected by a mean current or wind \mathbf{V} , the function can be computed analytically, as described in Supplementary Information. The result for the three dimensional case is:

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

where a is the size of the searcher, and the coordinates are chosen to have the wind blowing along the y axis in the negative direction. A similar expression holds in two dimensions. In both cases, the rate function decreases exponentially at large distances (anisotropically in the presence of wind). For experimental data on mixing flows, the parameters in the rate function are estimated from the data as detailed in Supplementary Information.

Mixing flow experiment and simulation parameters. Parameters used in the simulations are as follows. Figs 1a and 2: emission rate of the source $R = 1$, particle lifetime $\tau = 2,500$ and diffusivity $D = 1$. The typical travel distance during a lifetime is approximately $\sqrt{D\tau} = 50$. The search space is a grid 512×512 , and equation (1) was evaluated for five possible actions at each time step (moves to the four neighbours and standing still). Figure 1b: wind blows downwards with unit speed, the emission rate $R = 1/2$ and the lifetime of particles $\tau = 150$, the initial vertical separation to the source. Figure 3: experimental data are generated injecting dye by a 8 mm tube in the far field of a jet, a large-scale (integral scale ~ 10 cm) sustained turbulent flow. The pictures in the figure show the random advection of the dye downstream of the injection point (for about three integral scales). Dye is passively transported by the flow, as odours are. Odour detection events correspond to encounters with dye concentrations above a threshold fixed at about five times the average intensity level. The mean velocity (pointing downward in Fig. 3) is about 4 cm s^{-1} (its measurement by olfactory robots might be realized using standard anemometers), and the root-mean-square velocity is about 30% of the mean velocity, giving a Reynolds number of $\sim 10^3$. Snapshots of the field are acquired at a frequency of 200 Hz. Coherent odour patches make the searcher spiral around the location of the encounters due to correlations among the detections. Yet, the resulting search process is just twice as long as for a model with independent hits, and might be further accelerated by accounting for time-correlations, as detailed in Supplementary Information.

Received 1 May; accepted 14 November 2006.

1. Berg, B. C. *Random Walks in Biology* (Princeton Univ. Press, Princeton, 1993).

2. Payne, T. L., Birch, M. C. & Kennedy, M. C. (eds) *Mechanisms in Insect Olfaction* (Clarendon, Oxford, 1986).
3. Murlis, J., Elkinton, J. S. & Cardé, R. T. Odor plumes and how insects use them. *Annu. Rev. Entomol.* **37**, 505–532 (1992).
4. Dusenbery, D. B. *Sensory Ecology: How Organisms Acquire and Respond to Information* (Freeman, New York, 1992).
5. Mafra-Neto, A. & Cardé, R. T. Fine-scale structure of pheromone plumes modulates upwind orientation of flying moths. *Nature* **369**, 142–144 (1994).
6. Hansson, B. S. (ed.) *Insect Olfaction* (Springer, Berlin, 1999).
7. Villermaux, E. & Duplat, J. Mixing as an aggregation process. *Phys. Rev. Lett.* **91**, 184501 (2003).
8. Kennedy, J. S. Zigzagging and casting as a preprogrammed response to wind-borne odour: A review. *Physiol. Entomol.* **27**, 58–66 (1983).
9. Russell, R. A. *Odor Detection by Mobile Robots* (World Scientific, Singapore, 1999).
10. Webb, B. Robots in invertebrate neuroscience. *Nature* **417**, 359–363 (2002).
11. Marques, L. & de Almeida, A. (eds) Special issue on mobile robots olfaction. *Auton. Robots* **20**, 183–287 (2006).
12. Berg, H. C. & Purcell, E. M. Physics of chemoreception. *Biophys. J.* **20**, 193–219 (1977).
13. Ishida, H., Kagawa, Y., Nakamoto, T. & Moriizumi, T. Odor-source localization in the clean room by an autonomous mobile sensing system. *Sens. Actuators B* **33**, 115–121 (1996).
14. Kuwana, Y., Nagasawa, S., Shimoyama, I. & Kanzaki, R. Synthesis of the pheromone oriented behaviour of silkworm moths by a mobile robot with moth antennae as pheromone sensors. *Bioelectron.* **14**, 195–202 (1999).
15. Grasso, F. W., Consi, T. R., Mountain, D. C. & Atema, J. Biomimetic robot lobster performs chemo-orientation in turbulence using a pair of spatially separated sensors: Progress and challenges. *Rob. Auton. Syst.* **30**, 115–131 (2000).
16. Russell, R. A., Bab-Hadiashar, A., Shepherd, R. L. & Wallace, G. G. A comparison of reactive robot chemotaxis algorithms. *Rob. Auton. Syst.* **45**, 83–97 (2003).
17. Belanger, J. H. & Arbas, E. Behavioral strategies underlying pheromone-modulated flights in moths: Lessons from simulation studies. *J. Comp. Physiol. A* **183**, 345–360 (1998).
18. Li, W., Farrell, J. A. & Cardé, R. T. Tracking of fluid-advected odor plumes: strategies inspired by insect orientation to pheromone. *Adapt. Behav.* **9**, 143–170 (2001).
19. Farrell, J. A., Pang, S. & Li, W. Plume mapping via hidden Markov methods. *IEEE Trans. Syst. Man Cybern. B* **33**, 850–863 (2003).
20. Farrell, J. A., Pang, S. & Li, W. Chemical plume tracing via an autonomous underwater vehicle. *IEEE J. Ocean. Eng.* **30**, 428–442 (2005).
21. Ishida, H., Nakayama, G., Nakamoto, T. & Moriizumi, T. Controlling a gas/odor plume-tracking robot based on transient responses of gas sensors. *IEEE Sensors J.* **5**, 537–545 (2005).
22. Murlis, J. & Jones, C. D. Fine-scale structure of odor plumes in relation to insect orientation to distant pheromone and other attractant sources. *Physiol. Entomol.* **6**, 71–86 (1981).
23. Shraiman, B. I. & Siggia, E. D. Scalar turbulence. *Nature* **405**, 639–646 (2000).
24. Falkovich, G., Gawędzki, K. & Vergassola, M. Particles and fields in fluid turbulence. *Rev. Mod. Phys.* **73**, 913–975 (2001).
25. Pang, S. & Farrell, J. A. Chemical plume source localization. *IEEE Trans. Syst. Man Cybern. B* (in the press).
26. Shannon, C. E. A mathematical theory of communication. *Bell Syst. Tech. J.* **27**, 379–423 623–656 (1948).
27. Cover, T. M. & Thomas, J. A. *Elements of Information Theory* (Series in Telecommunication, Wiley, New York, 1991).
28. Sutton, R. S. & Barto, A. G. *Reinforcement Learning: An Introduction* (MIT Press, Cambridge, Massachusetts, 1998).
29. Balkovsky, E. & Shraiman, B. I. Olfactory search at high Reynolds number. *Proc. Natl Acad. Sci. USA* **99**, 12589–12593 (2002).
30. Kaupp, U. B. et al. The signal flow and motor response controlling chemotaxis of sea urchin sperm. *Nature Cell Biol.* **5**, 109–117 (2003).

Supplementary Information is linked to the online version of the paper at www.nature.com/nature.

Acknowledgements This work was done during the visits of M.V. and E.V. to KITP, and was supported by the ARO. E.V. is a member of the Institut Universitaire de France.

Author Information Reprints and permissions information is available at www.nature.com/reprints. The authors declare no competing financial interests. Correspondence and requests for materials should be addressed to B.I.S. (shraiman@kitp.ucsb.edu).