

Chasing information to search in random environments

J-B Masson¹, M Bailly Bechet² and M Vergassola¹

¹ Institut Pasteur, CNRS URA 2171, Unit In Silico Genetics, 75724 Paris Cedex 15, France

² Université C Bernard Lyon 1, CNRS, UMR5558, Laboratoire de Biométrie et Biologie Evolutive, 69622, Villeurbanne, France

E-mail: jbmasson@pasteur.fr

Received 15 February 2009, in final form 16 May 2009

Published 13 October 2009

Online at stacks.iop.org/JPhysA/42/434009

Abstract

We discuss search strategies for finding sources of particles transported in a random environment and detected by the searcher(s). The mixing of the particles in the environment is supposed to be strong, so that strategies based on concentration-gradient ascent are not viable. These dilute conditions are common in natural environments typical of searches performed by insects and birds. The sparseness of the detections constitutes the major stumbling block in developing efficient olfactory robots to detect mines, chemical leaks, etc. We first discuss a search strategy, 'infotaxis', recently introduced for the search of a single source by a single robot. Decisions are made by locally maximizing the rate of acquisition of information on the location of the source and they balance exploration and exploitation. We present numerical simulations demonstrating the efficiency of the method and, most importantly, its robustness to lack of detailed modeling of the transport of particles in the random environment. We then introduce a novel formulation of infotaxis for collective searches where a swarm of robots is available and must be coordinated. Gains in the search time are impressive and the method can be further generalized to deal with conflicts arising in the identification of multiple sources.

PACS numbers: 87.10.Vg, 02.60.Pn

(Some figures in this article are in colour only in the electronic version)

1. Introduction

The issue addressed here is, can (and how) a source of particles be located when the only clues of its presence are rare detections of the particles transported in the random environment?

The affirmative answer to the first question is provided by animals, such as birds and insects, capable of solving this problem. Male moths, for instance, locate females for mating by tracing the pheromones that they emit [1–4]. Entomologists distinguish two phases in

the trajectories characteristic of searching moths: ‘casting’ (when moths make large lateral excursions across the wind direction) and ‘zigzag’ (consistent motion upwind) [5]. As for birds, recent evidence for the role of olfaction in the foraging of albatrosses is presented in [6].

Entomology answers the first question yet it leaves open the second: how is this challenging goal achieved? Indeed, finding a source of food or pheromones at the macroscopic scales typical of insects and birds is a real challenge. Odors, pheromones etc, emanating from the source will tend to disperse laterally and downwind of the odor source and acquire an irregular and patchy concentration distribution due to turbulent transport [7]. High concentration regions are broken up into random patches because of the strong mixing typical of natural environments, e.g. the atmosphere. Some information is clearly hidden in the history of detection of the patches, yet we are dealing with a notoriously strongly fluctuating process [8, 9] and it is not obvious how these clues can be extracted. In particular, classical arguments presented in [10] show that no reliable strategy of chemotaxis [11] (ascent of the gradients of concentration) can be employed. Dilute conditions would in fact require extremely long waiting times to estimate local gradients in a reliable way [10], whilst moths respond to pheromones in fractions of a second.

No quantitative insight is available on the strategies of decision employed by insects and birds to move along their characteristic zigzagging and casting patterns and how they process the signal of their history of detections, e.g. of pheromones. This issue also bears upon important technological issues: the challenge of searching in dilute conditions is encountered in the design of sniffers [12–14]—robots that track chemicals emitted by drugs, chemical leaks, explosives and mines, trying to locate their source. In the absence of better strategies, the existing methods mostly rely upon chemotactic gradient ascent [15–18] or plume-tracking strategies [19–22]. However, as already mentioned, gradient ascent methods are constrained to work at high concentration. This severely limits the searches to short distances from the source and the same holds for plume-tracking (plumes mix and decay moving away from the source). To summarize, there is dire need to develop alternative search strategies effective in dilute conditions, which is the problem we shall discuss hereafter.

2. Building a probability map for the location of the source

The trace \mathcal{T}_t of detections, e.g. of odors, experienced by the searcher carries clues on the location of the source yet they are almost lost in the environmental strong noise. This is quite an unusual situation as communication signals are usually sent in channels with very low levels of noise. Here, the transmission channel is out of our control and we cannot dispense with its noisy characteristics. We shall therefore still treat the trace \mathcal{T}_t as a message sent from the source to the searcher. The message is input in Bayes’ formula, as usual in message decoding (see [23]), to construct the posterior probability $P_t(\mathbf{r}_0)$ for the unknown position of the source \mathbf{r}_0 . Both \mathcal{T}_t and $P_t(\mathbf{r}_0)$ are time-varying quantities that are constantly updated. The posterior distribution depends on odor transport in the environment via the rate of detection $R(\mathbf{r}|\mathbf{r}_0)$ at position \mathbf{r} conditional to the presence of the source at \mathbf{r}_0 . This is analogous to the properties of the channel in communication signals [23].

A very simple model of turbulent transport employed in [24] features particles emitted by a source at a rate R , having a lifetime τ , propagating with isotropic diffusivity D and possibly advected by a mean wind \mathbf{V} . This yields the following explicit expression for the conditional rate function in two dimensions (2D):

$$R(\mathbf{r}|\mathbf{r}_0) = \frac{R}{\ln\left(\frac{\lambda}{a}\right)} \exp\left(-\frac{V(y-y_0)}{2D}\right) K_0\left(\frac{|\mathbf{r}-\mathbf{r}_0|}{\lambda}\right), \quad (1)$$

where K_0 is the modified Bessel function of order zero and the correlation length $\lambda = \sqrt{\frac{D\tau}{1 + \frac{v^2\tau}{4D}}}$. An analogous expression holds in 3D [24]. We stress that no perfect model of the environment is realistically available in concrete applications and that the rate function $R(\mathbf{r}|\mathbf{r}_0)$ will only be an approximation to the actual transport process. This is the spirit of the simplistic model (1) and it will be crucial to check that any method of search be robust with respect to the choice of $R(\mathbf{r}|\mathbf{r}_0)$. Along the same lines, the simplest assumption for the time correlation of successive detections is their independence. This implies that the stochastic process of detections is Poissonian and that the probability distribution at time t posterior to experiencing a trace \mathcal{T}_t is

$$P_t(\mathbf{r}_0) = \frac{L_{\mathbf{r}_0}(\mathcal{T}_t)}{\int d\mathbf{x} L_{\mathbf{x}}(\mathcal{T}_t)} = \frac{e^{-\int_0^t R(\mathbf{r}(t')|\mathbf{r}_0) dt'} \prod_{i=1}^H R(\mathbf{r}(t_i)|\mathbf{r}_0)}{\int d\mathbf{x} e^{-\int_0^t R(\mathbf{r}(t')|\mathbf{x}) dt'} \prod_{i=1}^H R(\mathbf{r}(t_i)|\mathbf{x})}. \quad (2)$$

Here, H is the number of detections that the searcher has experienced, the t_i s are the times of these hits and $L_{\mathbf{r}_0}$ denotes the likelihood of experiencing the trace \mathcal{T}_t conditional to the presence of the source at \mathbf{r}_0 . Note that the absence of correlations permits the probability map to be updated without storing the whole history. Indeed, $P_{t+\delta t}(\mathbf{r}_0) = P_t(\mathbf{r}_0) e^{-R(\mathbf{r}(t+\delta t)|\mathbf{r}_0)\delta t} R^\eta(\mathbf{r}(t+\delta t)|\mathbf{r}_0)/Z_{t+\delta t}$, where η is the number of detections that the searcher experienced during the short time interval δt and $Z_{t+\delta t}$ is the constant normalizing $P_{t+\delta t}$ with respect to its argument \mathbf{r}_0 .

3. A simplified search problem

The aim of this section is to derive a lower bound on the expected search time as a function of the Shannon entropy S of the probability distribution map for the location of the source (supposed given and reliable). The rationale is to demonstrate that reduction of S is a necessary condition for a rapid search.

Since we are looking for a lower bound, we shall dismiss neighbor constraints and let the searcher be allowed to move freely from one point to any other. ‘Points’ should in fact be rather thought of as finite regions of space having the size of the searcher. The probability that the source is found at the j th visited point is p_j and $\sum p_j = 1$. The Shannon entropy [23] of the distribution is $S = -\sum p_j \ln p_j$. The expected search time is $T = \sum_j j p_j$. Indeed, if the source is found at the first point (probability p_1) then the searcher stops. If the source is not found, the map is updated by taking into account this additional information, i.e. $p_1 \mapsto 0$ and $p_{j \neq 1} \mapsto p_{j \neq 1}/(1 - p_1)$. The same process is repeated at each point visited and we end up with $T = 1 \cdot p_1 + (1 - p_1) \left[2 \frac{p_2}{1 - p_1} + \left(1 - \frac{p_2}{1 - p_1} \right) \left(3 \frac{p_3}{1 - p_1 - p_2} + \dots \right) \right]$. Reorganizing the various terms yields the aforementioned expression of T .

To obtain the lower bound, we minimize T with respect to all distributions p_j having entropy S :

$$T' = \sum_j j p_j + \alpha \left(\sum_j p_j - 1 \right) + \frac{1}{\beta} \left(\sum_j p_j \ln p_j + S \right), \quad (3)$$

where α and β are Lagrange multipliers enforcing the normalization of the p_j s and the value S of the entropy. The probability distribution corresponding to the minimum of (3) has the Gibbs form $p_j \propto \exp(-\beta j)$. Inserting this expression into (3), it is easy to check that the relation between T and S , if both are large compared to unity, reads $T = \sum_j j p_j \approx \exp(S - 1)$. This shows that the lower bound scales as the ‘effective’ number of points to be sought and that it increases rapidly with the entropy. This rapid increase is the motivation to infotaxis [24], which is based on chasing information, i.e. rapidly reducing S , as we shall describe below.

4. Infotaxis

The major problem for the searcher is that the probability map for the location of the source is unknown and that the posterior $P_t(\mathbf{r}_o)$ built in section 2 is just an estimation. The estimate is more and more reliable as detections are gathered but, especially during the initial stages of the search, the estimated map can be flat and biased. This is a classical problem encountered when trying to learn about random environments in the absence of any oracle supervising the learning. Two conflicting tendencies must be faced. On one hand, new actions should be tried and the available phase space must be explored. On the other hand, this should not be done blindly, e.g. the searcher in our case wants to minimize the search time. It is therefore sensible that exploration be weighted by an estimation of possible benefits based on the exploitation of previous actions. Greedy strategies push exploitation to its limit, namely choosing moves toward locations of maximal estimated probability to find the source. The general lesson of ‘reinforcement learning’ is that neither pure exploitation nor pure exploration is viable and a tradeoff should be struck [25]. As the name suggests, the weight of ‘good’ actions should be reinforced whilst ‘bad’ actions should be avoided. The problem is that no general method is available for the reinforcement procedure and that effective balances between exploitation and exploration ought to be found on a case-by-case basis.

Let us show that decisions aimed at maximizing the rate of acquisition of information on the location of the source indeed balance exploitation and exploration for our search problem [24]. We shall suppose that the searcher has reached the point \mathbf{r} at time t , its posterior distribution is $P_t(\mathbf{r}_o)$ and its Shannon entropy is S . The expected variation of entropy upon moving to one of the neighbors \mathbf{r}' (or standing still at \mathbf{r}) is

$$\overline{\Delta S}(\mathbf{r} \rightarrow \mathbf{r}') = -P_t(\mathbf{r}')S + (1 - P_t(\mathbf{r}')) \left[\sum_{k=0}^{\infty} \rho_k(\mathbf{r}') \Delta S_k \right]. \quad (4)$$

The first term on the right-hand side corresponds to the event of finding the source at the neighboring \mathbf{r}' point. If so happens, $P_{t+\delta t}$ becomes a δ function and the entropy vanishes. The second term takes into account the alternative case: source not found at \mathbf{r}' . The probability of having k hits during the δt time step is denoted $\rho_k(\mathbf{r}')$. In the Poissonian model of detections $\rho_k = h^k \exp(-h)/k!$, where h is the average number of hits. Its estimate, on the basis of what the searcher knows, is $h(\mathbf{r}') = \delta t \int P_t(\mathbf{r}_o) R(\mathbf{r}'|\mathbf{r}_o) d\mathbf{r}_o$ where $R(\mathbf{r}'|\mathbf{r}_o)$ is the average hit rate introduced in section 2 and δt is the time step. Finally, ΔS_k is the variation of entropy between the fields $P_{t+\delta t}$ and P_t . The field $P_{t+\delta t}$ involved in the calculation of ΔS_k is constructed from P_t assuming that k detections have taken place and updating the field as described in section 2.

Equation (4) shows that chasing information ensures a balance between exploration and exploitation. The first term alone would correspond to a criterion of maximum likelihood. Note that S would then drop out and no sense of the current state of uncertainty would be kept. The second term itself is a mixture of explorative and exploitative tendencies but it dismisses the possibility of chasing the source and considers just additional clues. As shown in [24], a strategy based on the second term alone performs much better than pure exploitation yet it is over-conservative and is outperformed by infotaxis.

Having identified a suitable quantity to balance exploitation and exploration, the decision-making process simply picks the action that locally maximizes the expected rate of information acquisition. More specifically, at each time step the searcher chooses the neighboring \mathbf{r}' which gives the minimum value of $\overline{\Delta S}$ in (4). Stochastic weighted choices are also possible but we have not observed dramatic changes.

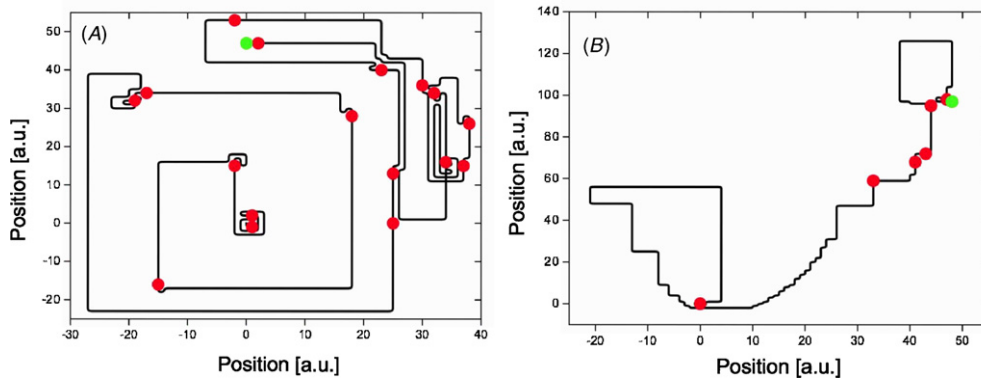


Figure 1. (A) An infotactic trajectory without wind (black line) toward a source (green circle). The source emission rate $R = 1$, the particles lifetime $\tau = 2500$ au and the diffusivity of the particles $D = 1$ au. Red circles are detections experienced by the searcher during its motion. (B) An infotactic trajectory in the presence of wind blowing vertically in the descending direction with unit intensity. Parameters are $\tau = 150$ au, $R = 1$ and $D = 1$ au.

Finally, one may wonder whether a free energy formulation, reinforcing paths toward maximum likelihood locations, would not be even more advantageous. The problem is how to choose an appropriate ‘temperature’ on a general basis. Numerical results (not presented here) show that for a given choice of the initial distance to the source and parameters of the environment, values of the temperature can be chosen so that the search will be slightly more efficient than infotaxis. The gain mostly arises from the last phase, when the searcher is close to the source and its wandering is slightly reduced. However, the choice of the appropriate temperature is far from universal and the search process can be dramatically slowed down for a different choice of the initial distance from the source. An adaptive choice of temperature should be considered but it is not clear how a good cooling could be realized and how universal such a prescription could be. To summarize, infotaxis appears the best option which remains quite general and does not require any tuning of parameters to ensure rapid and reliable searches.

5. Numerical results

Performances of the infotaxis method are conveniently assessed by numerical simulations. A couple of instances of infotactic trajectories are displayed in figure 1. In the absence of wind (figure 1(A)) only detections can break the spherical symmetry: as more and more detections are experienced, the direction to the source emerges in the posterior P_t and this information is exploited to point toward the source. The search is easier and more guided in the presence of wind (figure 1(B)). Similarly to flights of birds and moths [5], the infotactic searcher alternates ‘casting’ periods, moving perpendicularly to the wind or even in the same direction, and ‘zigzagging’ periods. The latter tend to be more frequent after detections. Conversely, casting and backtracking are more frequent at the beginning of the trajectory. This is just a general tendency though, as it can be seen in figure 1 that even near the end of the trajectory, the searcher went too far and had to backtrack to get to the source.

In the case where no wind and no detections are experienced, the infotactic searcher moves in two dimensions (2D) along an Archimedean spiral [29] (see figure 2(A)). The underlying mesh was chosen hexagonal to ensure the maximum of isotropy on a discrete lattice. The




Figure 2. (A) An infotactic trajectory in the absence of any source. Note that it is an Archimedean spiral (on a hexagonal lattice). Parameters defined in figure 1 are $R = 1$, $\tau = 400$ au and $D = 1$ au. (B) The spiral arm length as a function of the correlation length $\lambda = \sqrt{D\tau}$ of the concentration field. The (red) line is the linear fit.

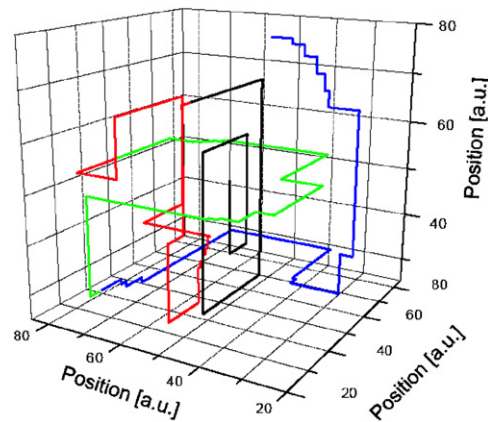


Figure 3. An infotactic trajectory in the absence of any source in three dimensions. Parameters are the same as in figure 2. Colors are used to help visualization of the three-dimensional structure of the trajectory.

spiral arm scales linearly with the correlation length of the source λ (figure 2(B)). This sensibly reduces the chance of missing the source, an event which was indeed never observed in our simulations. The same type of trajectory in 3D is shown in figure 3. Interestingly, the searcher seems to spiral in a plane, then change plane, spiral again there and so on. Determining the functional form of the corresponding curve (on the lattice and/or in the continuum) remains an open problem.

The dependence of the average search time on the initial distance from the source (figure 4(A)) shows a region of linear behavior for large distances (roughly between λ and 2λ , where λ is the correlation length of the concentration field). This evolution becomes quadratic for even larger distance ($d > 2\lambda$), reflecting the Archimedean spiraling that takes place at very large distances, where detections are almost absent. For comparison, the search time for a gradient-climbing strategy would depend exponentially on the distance to the source. This dependence follows from classical criteria [10] for gradient detection, applied to an exponentially decaying concentration field. The probability distribution function (PDF) of the search time for a given initial distance exhibits an exponential decay for large times (figure 4) proving that fluctuations are limited. Furthermore, even though a proof that infotaxis never




Figure 4. (A) The average search time versus the initial distance from the source. The lifetime of the particles is $\tau = 2500$ au, the emission rate of the source is $R = 1$ and the diffusivity of the particle $D = 1$ au. The red line is a linear fit of the curve from distance 50 to 110 au. The green line is a parabolic fit of the curve from distance 120 to 200 au. (B) Probability distribution function of the search times for three initial distances, 30, 90 and 150 au (black, red and green).

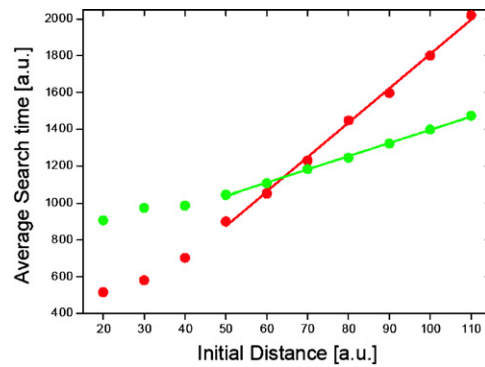


Figure 5. The average search time versus the initial distance from the source for two infotactic searchers with poor modeling of the environmental transport. Parameters are as in figure 4. The correlation length of the detection probability field is under- (over-) estimated by a factor two for the red (green) circle curve. Straight lines are the corresponding linear fits.

misses a source, extensive simulations never exhibited such a case, even for searches starting at initial distances around 10λ .

Another major asset of infotaxis is its robustness with respect to the model of the environment which is employed. The robustness of the method with respect to the choice of $R(r|r_0)$ was tested in [24] using experimental images of dye diffusion in a jet flow. Here, we demonstrate more quantitatively this robustness by testing it in simulations with well-controlled models of the environment. Particles are transported in the environment according to the model (1), yet the searcher is supposed to have bad estimates of the correlation length λ . Namely, in calculating the posterior distribution it employs the same functional form but with values λ' either underestimated ($\lambda' = \lambda/2$) or overestimated by a factor two. As can be seen in figure 5, both cases show similar characteristics and, even though less efficient, they still remain very effective and never miss the source. These numerical simulations provide further support to the conclusion that infotaxis is well able to manage mistakes in the modeling of the environment.

6. More is better

Infotaxis was introduced in [24] for the case of a single searcher and a unique source. In fact, the cost of olfactory robots being moderate, there is a great deal of interest in developing collective strategies, i.e. having a group of robots working in collaboration [26–28]. This is *a priori* poised to facilitate the task, namely, reduce the search time. However, empirical forms of collaboration often end up with rather inefficient interactions among the robots of the swarm. Conversely, the general principle underlying infotaxis naturally lends to collaborations among multiple searchers and the purpose of this section will be to demonstrate this point.

In the simplest formulation of the problem, a group of searchers tries to localize a single source. The searchers build up collectively a shared map $P_t^{\text{sh}}(\mathbf{r}_0)$ for the probability of the source location. Sharing is quite realistic as robots are normally equipped with wireless devices and detections possibly experienced by members of the swarm in a short time interval are rapidly communicated. The probability field shared by all the searchers is updated as follows:

$$P_{t+\delta t}^{\text{sh}}(\mathbf{r}_0) = \frac{P_t^{\text{sh}}(\mathbf{r}_0) \prod_{i=1}^n [e^{-R(\mathbf{r}_i(t+\delta t)|\mathbf{r}_0)\delta t} R^{\eta_i}(\mathbf{r}_i(t+\delta t)|\mathbf{r}_0)]}{Z_{t+\delta t}}. \quad (5)$$

Here, η_i is the number of hits that the i th searcher experienced between t and $t + \delta t$, $Z_{t+\delta t}$ is the normalizing constant and n is the number of searchers in the swarm.

Two possible options of decision-making are now available. In the first one, decisions are taken individually by each searcher: the i th searcher uses the shared field P_t^{sh} to estimate the probability of the various events but only evaluates the consequences of its own actions. Namely, the i th searcher chooses the move minimizing $\overline{\Delta S}(\mathbf{r}_i \rightarrow \mathbf{r}'_i) = -S P_t^{\text{sh}}(\mathbf{r}'_i) + (1 - P_t^{\text{sh}}(\mathbf{r}'_i)) [\sum_{k=0}^{\infty} \rho_k(\mathbf{r}'_i) \Delta S_k]$ among the sites \mathbf{r}'_i neighbors to its current position \mathbf{r}_i . To summarize, searchers share information but do not coordinate decisions.

In the second option, decision making is coordinated and it evaluates all joint moves available to the whole swarm. The expected variation of entropy is evaluated for all combinations as

$$\overline{\Delta S}(\{\mathbf{r}_i\} \rightarrow \{\mathbf{r}'_i\}) = -S P_t^*(\{\mathbf{r}'_i\}) + (1 - P_t^*(\{\mathbf{r}'_i\})) \sum_{k_1=0}^{\infty} \dots \sum_{k_n=0}^{\infty} \Delta S_{\{k_i\}} \prod_{j=1}^n \rho_{k_j}(\mathbf{r}'_j), \quad (6)$$

where $P_t^* = \sum_j P_t(\mathbf{r}'_j) \prod_{k \neq j} (1 - P_t(\mathbf{r}'_k))$ is the probability that the source is found by one of the searchers and the notation $\{\bullet\}$ denotes the whole ensemble of the n possible variables, e.g. k_i ($i = 1, \dots, n$). The reduction of entropy is by definition more effective than for the first choice but the expression (6) makes it clear that the price to be paid is the rapidly increasing number of operations per time step. The number of possible actions for n searchers in 2D is 5^n (on a square mesh) so that already for $n > 4$ their exhaustive evaluation is rather cumbersome. Full collaboration can still be maintained by making two simplifications. First, the sums on the indices k in (6) are truncated to unity. This is not a harsh simplification because probabilities of multiple detections typically are $< 10^{-4}$. Second, we perform a fast-cooling simulated annealing procedure to identify the coordinated move to be chosen rather than making an exhaustive evaluation. In practice, the cooling procedure was tuned so as to reach convergence after 50 steps.

To evaluate the performance of collective infotaxis, numerical results were compared to those for a group of independent infotactic searchers (ignoring each other). Note that for collective searches the whole geometry of the initial positions of the swarm ought to be specified. To limit this space, we focused on symmetric configurations and initially placed the searchers on a circle of radius d , at the center of which stands the source. The search

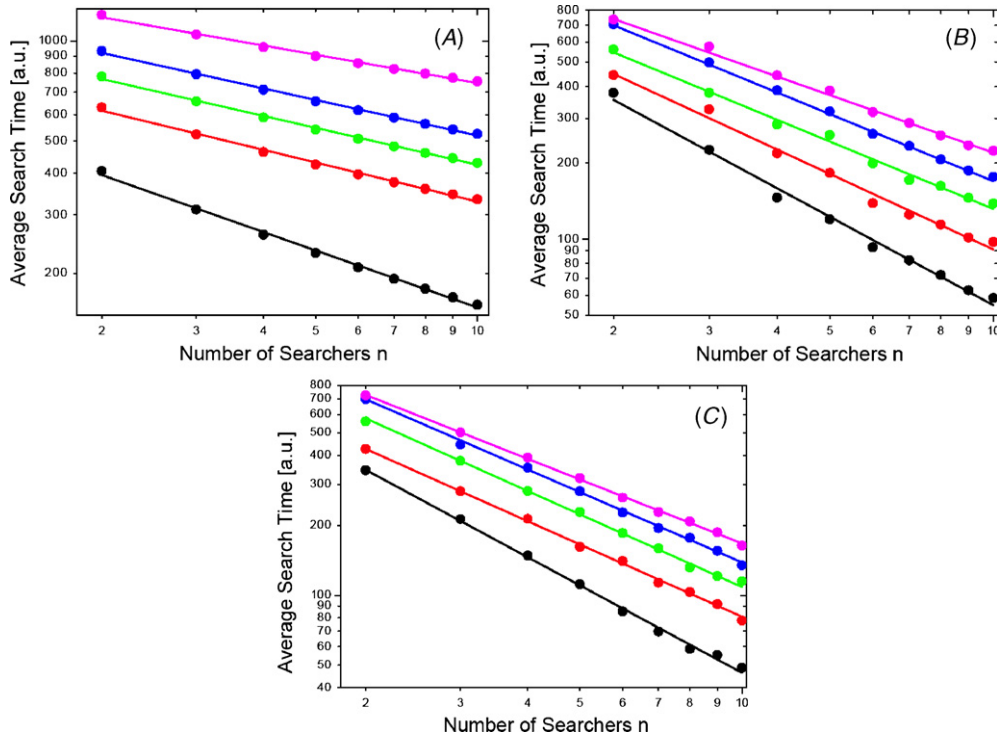


Figure 6. Dependence of the average search time on the number of searchers n . All plots are in log–log scale, circles are results from numerical simulations and lines are power-law fits. Colors indicate the initial distance from the source (black, red, green, blue and magenta correspond to $d = 30, 60, 90, 120, 150$ au, respectively). Independent searchers are plotted in **A**, infotactic searchers not sharing decisions in **B** and fully collaborative infotactic searchers in **C**.

stops when one of the searchers has found the source. In the case of completely independent searchers it was assumed that every searcher had an initial detection at $t = 0$ to make the comparison to the collective case quite fair.

The dependence of the average search time on the number of searchers n for different initial distances d to the source is shown in figure 6. Both forms of collaborative infotaxis perform vastly better than independent infotactic searchers. In all cases we considered, the data are well fitted by a power law d^β . The exponents β for the extreme distances $d = 30$ and $d = 150$ are $(-0.57, -1.16, -1.25)$ and $(-0.28, -0.76, -0.91)$ for independent, partially and fully collaborating searchers. Note that the exponents for the latter are slightly superior to partially collaborating searchers, but they are both much larger than the values for independent searchers. The PDFs of the search times (figure 7) also exhibit exponential decays that are much more rapid.

A relevant conclusion drawn from the previous results is that gains in the search time are most conspicuous for fully collaborative swarms, yet even just sharing the posterior probability map leads to spectacular effects. This is important as the experimental implementation of full collaboration might be hampered by the single-step complexity of operations. Conversely, partially collaborating swarms are easy to implement. Intermediate options, such as decision sharing between close robots, are possible depending on technical characteristics and specificities of the robots available.

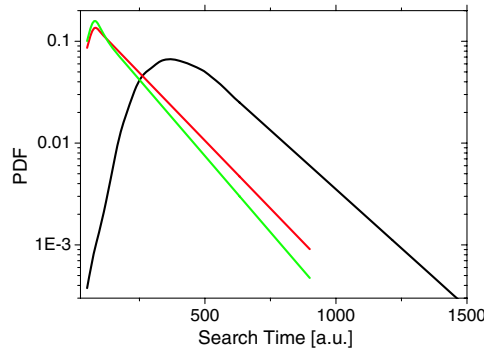


Figure 7. The probability distribution function of the search time for $d = 60$ au and four searchers. Black indicates independent searchers, red indicates infotactic searchers not sharing decisions and green indicates fully collaborative infotactic searchers. Parameters of the simulations are $R = 1$, $\tau = 2500$ au and $D = 1$ au.




Figure 8. The correlation C defined in the text versus the normalized initial distance between two searchers z/λ . Black and red colors are for partial and full collaboration, respectively. Circles are the results of simulations, lines are just visual aids. Note the change of sign occurring at $\simeq 2\lambda$.

Sharing of information and decisions generates non-trivial effective interactions among the searchers. In particular, the interactions mean that searchers tend to repel each other when their distance is inferior to $\simeq 2\lambda$ and tend to attract when their distance is superior to $\simeq 2\lambda$. We recall that λ denotes the correlation length of the detection probability field as a function of the distance to the source. For the sake of clarity, we focus on the case of two searchers. We initially placed them at the same distance from the source, their mutual distance is denoted $z(t)$ and their respective positions are followed for the first $t' = 100$ time steps (so as to remain far from the source). A statistic appropriate to our purposes is constructed as $C = \sum_{s=1}^{t'} \sigma_s / t'$, where $\sigma_s = \pm 1$ depending on whether the distance $z(t)$ increases or reduces in the current time step ($\sigma_s = 0$ if $z(t)$ remains the same). The dependence of C on the initial distance between the two searchers is displayed in figure 8. Obviously, for independent searchers C is the sum of roughly independent variables and it remains close to zero. Interestingly, both types of infotactic collaborations show the same general trend of repulsion/attraction that we previously announced. The repulsion observed in figure 8 and the preference to remain at




Figure 9. Trajectories of eight infotactic searchers initially located in a small circle of radius $d = 5$ au. No source is present and parameters of the simulation are $R = 1$, $\tau = 2500$ au and $D = 1$ au. The simulation was performed on a hexagonal lattice.

distances scaling with λ is very sensible. This is indeed the most effective way to identify a region of typical size λ (where the source is located and its emissions are frequent and easily detectible).

The behavior of the correlation C can also be studied during the period just before one of the searchers finds the source. We build up the function in the same manner as before but starting 100 time steps before the end and summing up to the end of the search. Again, independent searchers have values very close to zero. Both types of infotactic collaboration have negative C . The average value for full collaboration is -0.148 and for partial collaboration is -0.110 . Negative values indicate that searchers tend to get close to each other when approaching the source. This is again quite sensible since, if enough information has been gathered to locate the source reliably, the searchers should indeed tend to converge toward the same region.

To summarize, interactions induced by sharing information (and decisions) permit the searchers to explore the space with steps of typical size comparable to the correlation length of the field to be explored. Furthermore, when enough information has been gathered and the source is localized reliably, interactions make the whole group converge toward the source. For more than a pair of searchers, the patterns get of course more complex, as can be appreciated in figure 9. The aforementioned tendencies persist but different subgroups can co-exist. In figure 9, one can see for instance a group of searchers moving away from each other, but also a subgroup of three searchers moving toward each other before separating. Identifying appropriate quantitative characterizations of the geometry of these collective trajectories is an open problem.

6.1. Multiple sources

Multiple sources can generate conflicts and ambiguities both in decoding and decisions and their search is harder than for a single source. In the following, we shall deal with two sources but the methods are quite general. The goal of the search will be to get as fast as possible to one of the sources. The most difficult case (where conflicts are hampering searches the most)

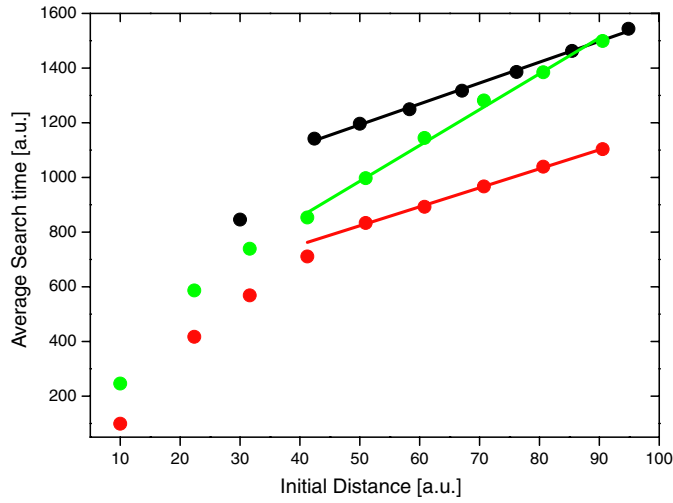


Figure 10. The average search time versus the initial distance to two sources for individual infotaxis searching for a single source (black), for two collaborative infotactic searchers chasing a single source (green) and, finally, for the mean-field multisource method proposed here (red).

is when the distance between the sources is comparable to their correlation length. Conflicts arise from the fact that a searcher approaching the sources ends up over-exploring the inter-source region. Trying to interpret its detections with a model featuring a single source, it gets stuck by contradictory clues that make it move alternatively toward one or the other source.

An obvious way out would be to construct a probability map for the positions of two sources. However, even in two dimensions this requires storing an array of size N^4 (N being the number of points along each direction) and this is not a viable solution for realistic sizes. The method presented hereafter is based on a mean-field type method where the map of the multiple sources is approximated using combinations of the individual maps of a swarm of n robots. The cost in memory is nN^2 (in 2D), which can be easily accommodated.

Let us denote s the number of sources. Every searcher has its own map of probability and decisions are again taken individually by maximizing the reduction of entropy of the individual probability map. The update of the map of the i th searcher is also done as for individual searchers, i.e.

$$P_{t+\delta t}^i(\mathbf{r}_0) = \frac{P_t^i(\mathbf{r}_0) e^{-h(\mathbf{r}_i(t+\delta t)|\mathbf{r}_0)\delta t} h^{\eta_i}(\mathbf{r}_i(t+\delta t)|\mathbf{r}_0)}{Z_{t+\delta t}} \quad (7)$$

except that the expected hit rate now reads

$$h(\mathbf{r}_i|\mathbf{r}_0) = R(\mathbf{r}_i|\mathbf{r}_0) + \frac{s-1}{n-1} \sum_{j=1; j \neq i}^n \int d\mathbf{x} P_t^j(\mathbf{x}) R(\mathbf{r}_i|\mathbf{x}). \quad (8)$$

This might be thought of as a mean field approximation exploiting the information gathered by the other $n - 1$ searchers. It obviously reduces to the single-searcher case for $s = 1$.

The validity of the new method is confirmed in figure 10 by measuring the search times for the case where two sources are present. Extensive simulations show that, as long as $n \geq s$, the process is efficient. In figure 11 one can see the plot of the search time versus the number of searchers. The evolution is again a scaling law of exponent $\beta \simeq 0.95$, validating the method.

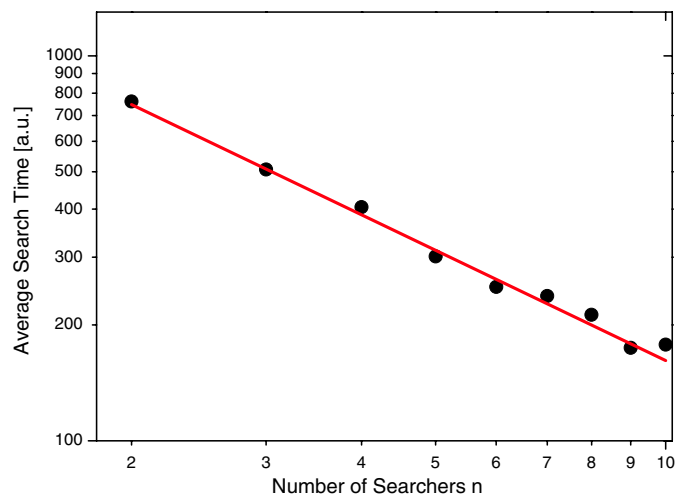


Figure 11. Dependence of the average search time on the number of searchers n in the presence of two sources sought by the mean-field multisource method presented in the text.

7. Conclusion

We have confirmed the validity of infotaxis for a single source and a single searcher by quantitatively testing its robustness with respect to inappropriate models of transport in the environment. We have presented novel formulations of the method for collective searches involving swarms of searchers. The value of the coordination ensured by information and decision sharing among members of the swarm has been demonstrated by extensive numerical simulations. Gains in the search times are impressive and we have also shown how the swarm can interact for effectively searching multiple sources. The method seems mature to be implemented on real robots and to provide the search strategy for the next generation of sniffers.

Acknowledgments

This research was partially supported by the National Science Foundation under Grant No. PHY05-51164.

References

- [1] Payne T L, Birch M C and Kennedy M C (eds) 1986 *Mechanisms in Insect Olfaction* (Oxford: Clarendon)
- [2] Murlis J, Elkinton J S and Cardé R T 1992 *Ann. Rev. Entomol.* **37** 505–32
- [3] Mafrá-Neto A and Cardé R T 1994 *Nature* **369** 142–4
- [4] Hansson B S (ed) 1999 *Insect Olfaction* (Berlin: Springer)
- [5] Kennedy J S 1983 *Physiol. Entomol.* **27** 58–66
- [6] Nevitt G A, Losekoot M and Weimerskirch H 2008 *Proc. Natl Acad. Sci.* **105** 4576–81
- [7] Murlis J and Jones C D 1981 *Physiol. Entomol.* **6** 71–86
- [8] Shraiman B I and Siggia E D 2000 *Nature* **405** 639–46
- [9] Falkovich G, Gawędzki K and Vergassola M 2001 *Rev. Mod. Phys.* **73** 913–75
- [10] Berg H C and Purcell E M 1977 *Biophys. J.* **20** 193–219
- [11] Berg H C 1993 *Random Walks in Biology* (Princeton, NJ: Princeton University Press)

- [12] Russell R A 1999 *Odor Detection by Mobile Robots* (Singapore: World Scientific)
- [13] Webb B 2002 *Nature* **417** 359–63
- [14] Settles G S 2005 *J. Fluids Eng.* **127** 189–218
- [15] Ishida H, Kagawa Y, Nakamoto T and Moriizumi T 1996 *Sensors Actuators B* **33** 115–21
- [16] Kuwana Y, Nagasawa S, Shimoyama I and Kanzaki R 1999 *Biosensors Bioelectron.* **14** 195–202
- [17] Grasso F W, Consi T R, Mountain D C and Atema J 2000 *Robot. Auton. Syst.* **30** 115–31
- [18] Russell R A, Bab-Hadiashar A, Shepherd R L and Wallace G G 2003 *Robot. Auton. Syst.* **45** 83–97
- [19] Belanger J H and Arbas E 1998 *J. Comp. Physiol. A* **183** 345–60
- [20] Li W, Farrell J A and Cardé R T 2001 *Adaptive Behav.* **9** 143–70
- [21] Farrell J A, Pang S and Li W 2005 *IEEE J. Oceanic Eng.* **30** 428–42
- [22] Ishida H, Nakayama G, Nakamoto T and Moriizumi T 2005 *IEEE Sensors J.* **5** 537–45
- [23] Cover T M and Thomas J A 1991 *Elements of Information Theory (Wiley Series in Telecommunication)* (New York: Wiley)
- [24] Vergassola M, Villermaux E and Shraiman B I 2007 *Nature* **445** 406–9
- [25] Sutton R S and Barto A G 1998 *Reinforcement Learning: An Introduction* (Cambridge, MA: MIT Press)
- [26] Hayes A T, Martinoli A and Goodman R M 2002 *IEEE Sensors J.* **2** 260–71
- [27] Hayes A T, Martinoli A and Goodman R M 2003 *Robotica* **21** 260–71
- [28] Jatmiko W, Sekiyama K and Fukuda T 2006 *Proc. 8th Symp. Distributed Autonomous Robotic Systems (DARS 2006) (IEEE)* pp 71–80
- [29] Barbieri C Infotassi: un algoritmo di ricerca senza gradienti *Tesi di Laurea* Università di Roma La Sapienza