

On the trajectories and performance of Infotaxis, an information-based greedy search algorithm

C. BARBIERI², S. COCCO^{1,2}, R. MONASSON^{1,3}

¹ *Simons Center for Systems Biology, Institute for Advanced Study, Einstein Drive, Princeton, NJ 08540, USA*

² *Lab. Physique Statistique de l'ENS, CNRS and Univ. Paris 6 - 24 rue Lhomond, 75231 Paris, France*

³ *Lab. Physique Théorique de l'ENS, CNRS and Univ. Paris 6 - 24 rue Lhomond, 75231 Paris, France*

PACS 05.40.-a – Fluctuation phenomena, random processes, noise, and Brownian motion

PACS 02.50.Tt – Inference Methods

PACS 87.19.lt – Sensory systems: visual, auditory, tactile, taste, and olfaction

Abstract. - We present a continuous-space version of Infotaxis, a search algorithm where a searcher greedily moves to maximize the gain in information about the position of the target to be found. Using a combination of analytical and numerical tools we study the nature of the trajectories in two and three dimensions. The probability that the search is successful and the running time of the search are estimated. A possible extension to non-greedy search is suggested.

Introduction. – Reaching a target with limited information is a fundamental task for living organisms. Small organisms, such as bacteria and eukaryotic cells, are thought to estimate and ascend the gradient of nutrient concentration, a process called chemotaxis [1, 6]. At the scale of a larger organism the Reynolds number is higher [4]: most biologically relevant chemical fields become turbulent and dilute. As a result, the trajectories of insects following odor traces appear much more complex than those of smaller organisms [12]. The modeling of search processes in presence of noisy information is important not only for biology, but also for robotics [8, 9].

Assume the search has proceeded for some time, and the searcher has received some hits, *i.e.* has detected some molecule of odor sent by the target, along the trajectory. How should the searcher move next? The timing and locations of the hits, as well as the absence of hits along the remaining parts of the trajectory all provide useful information about the location \mathbf{y} of the target. In Bayesian terms, this defines a posterior probability $P_t(\mathbf{y})$ for the position of the target. Going towards the maximum of P_t is not an optimal strategy as the maximum does generally not coincide with the target, especially in the initial stage of the search process. Recently, Vergassola, Villermaux and Shraiman proposed an alternative strategy, called Infotaxis [11]. The searcher moves to maximize the (expected) gain in information about the location of the target, that is, the loss in the entropy of the distribution $P_t(\mathbf{y})$. As

the search goes on, the entropy typically decreases, until the source is finally located. The strategy naturally balances the needs for exploration (harvesting more information about the target location) and exploitation (going towards the maximum of P_t). Infotaxis was implemented and tested on two-dimensional square lattices¹: the target was almost always found, and the distribution of search time appeared to decay exponentially.

Yet some important questions about Infotaxis remain open. First, how well does the algorithm perform in three dimensions? In addition to its practical interest, this question arises naturally in the context of Brownian motion theory. Purely random walks are space-filling in two dimensions, and transient in higher-dimensional spaces. Finding a target in three dimensions is therefore much harder, and constitutes a real test for the capabilities of Infotaxis. Secondly, how dependent of the underlying lattice are the results reported in [11]? Realistic descriptions of animal behavior or implementations in biomimetic robots require us to consider continuous spaces. In addition the presence of a lattice introduces anisotropies, while the odor propagation model used in [11] was isotropic. Thirdly, two-dimensional trajectories seem to exhibit spiral-like shapes. How precisely can we characterize those spirals, and what are their counterparts in three dimensions?

¹Few short trajectories were obtained on small three-dimensional lattices in [10].

In this letter, we derive the equation of motion for the Infotaxis searcher in the continuous space. We then introduce an algorithm to solve this equation². The performances of Infotaxis, *i.e.* the probability of success and the distribution of the search times are studied in $D = 2$ and 3 dimensions. The spiral- and coil-like shape of the search trajectories for, respectively, $D = 2$ and 3, and the pinning of the trajectory around the received hits are investigated analytically and numerically. We show that the motion of a searcher receiving an average and deterministic signal is a good predictor of the typical properties of the motion in the presence of stochastic hits. Finally we discuss a possible extension to a non-greedy search strategy, which could help reduce pinning effects.

Equation of motion for the searcher. – A point-source in \mathbf{y}^* emits particles, which diffuse in space, and have a finite lifetime. In the stationary regime, the probability per unit of time to encounter a particle in \mathbf{x} is denoted by $R(\mathbf{y}^* - \mathbf{x})$ [11][Supplementary information]. Function R has an integrable divergence at the origin ($R(u) \sim -\log u$ in $D = 2$, $\sim \frac{1}{u}$ in $D = 3$ dimensions, when $u \rightarrow 0$), and exponentially decreasing tails for large distances u . In the following distances are measured in the unit of the decay length of R . The unit of time is the inverse of the rate of emission of particles by the source, divided by the (dimensionless) linear size, a , of the searcher for $D = 3$, or multiplied by $\log(1/a)$ for $D = 2$ [11].

Let $\mathbf{x}(t)$ be the position of the searcher at time t . We denote by N_H the number of particles detected (called hits) at earlier times, $0 \leq t_i \leq t$, with $i = 1, \dots, N_H$. Based on those hits, the searcher can draw a probabilistic map over the possible locations \mathbf{y} of the source. In the Bayesian framework, the posterior probability density $P_t(\mathbf{y})$ for the location of the source is the (normalized) product of the probabilities of having detected the N_H particles at locations $\mathbf{x}(t_i)$, times the probability of not having detected any particle at other locations along the trajectory, times the prior probability density P_0 over \mathbf{y} ,

$$P_t(\mathbf{y}) \propto \prod_{i=1}^{N_H} R(\mathbf{y} - \mathbf{x}(t_i)) e^{-\int_0^t dt' R(\mathbf{y} - \mathbf{x}(t'))} P_0(\mathbf{y}) . \quad (1)$$

Hence, P_t diverges where the hits have been received and vanishes in the other places along the trajectory. In the following we will use brackets to denote averages over this posterior distribution,

$$\langle f(\mathbf{y}) \rangle_{\mathbf{y};t} = \int d\mathbf{y} P_t(\mathbf{y}) f(\mathbf{y}) . \quad (2)$$

Assume now that the searcher stays in $\mathbf{x}(t)$ during an infinitesimal time δt . The number of hits, n , received during this time interval is a stochastic variable equal to zero or one, with probabilities $p(0|\mathbf{y}) = 1 - \delta t R(\mathbf{y} - \mathbf{x})$ and $p(1|\mathbf{y}) = \delta t R(\mathbf{y} - \mathbf{x})$, depending on the location \mathbf{y} of the

source. In the language of information theory, the particle emission and detection system can be thought as a noisy channel, and n is the output message associated to the input codeword \mathbf{y} . The mutual information δI between n and \mathbf{y} is

$$\delta I = \sum_{n=0,1} \left\langle P(n|\mathbf{y}) \log \left(\frac{P(n|\mathbf{y})}{\langle P(n|\mathbf{y}') \rangle_{\mathbf{y}';t}} \right) \right\rangle_{\mathbf{y};t} = -\delta t V_t(\mathbf{x}(t)) \quad (3)$$

up to $O(\delta t^2)$, where

$$V_t(\mathbf{x}) = \left\langle R(\mathbf{y} - \mathbf{x}) \log \left(\frac{\langle R(\mathbf{y}' - \mathbf{x}) \rangle_{\mathbf{y}';t}}{R(\mathbf{y} - \mathbf{x})} \right) \right\rangle_{\mathbf{y};t} , \quad (4)$$

is the entropy rate of the posterior distribution P_t . Infotaxis stipulates that δI should be maximized, or, equivalently that V_t should be minimized, *i.e.* made as negative as possible. In other words, interpreting V_t as a potential, the searcher should descend the gradient of V_t . A natural equation of motion is then

$$\gamma(t) \dot{\mathbf{x}}(t) = -\nabla_{\mathbf{x}} V_t(\mathbf{x}(t)) , \quad (5)$$

where $\gamma(t)$ plays the role of a friction coefficient. A possible choice is $\gamma(t) = |\nabla_{\mathbf{x}} V_t(\mathbf{x}(t))|/v_0$, to keep the modulus of the velocity fixed and equal to v_0 . This is close to the lattice version of [11], where the searcher could either stay immobile or move by one lattice site, and the velocity could take only one non-zero value. In general, $\gamma(t)$ can be any positive function of the time. In the following, we will first consider the case $\gamma(t) = \gamma$ constant. We will later discuss to what extent the trajectories and the performances change when γ varies with the time.

Numerical integration. – The equation of motion (5) is highly non-linear and depends on the whole history of the search process through the posterior probability (1). To solve equation (5) numerically we discretize the time with a step Δt . The positions \mathbf{x}_ℓ of the searcher at the instants $t_\ell = \ell \Delta t$, where ℓ is a positive integer, are memorized, as well as the occurrences (times) of the hits. The amplitude of the ℓ^{th} elementary move is estimated from (5): $\mathbf{x}_{\ell+1} - \mathbf{x}_\ell = -\frac{\Delta t}{\gamma(t_\ell)} \nabla_{\mathbf{x}} V_t(\mathbf{x}_\ell)$. The calculation of the gradient of the potential requires to estimate averages over the space \mathbf{y} with measure P_t (2). To do so, we use the importance sampling Monte Carlo method [7]. The term inside the bracket in (4), and its gradient with respect to \mathbf{x} are exponentially decreasing functions of the radial distance $u = |\mathbf{x}_\ell - \mathbf{y}|$. We thus perform the change of variable $u = u_0(1 - v)/v$, where $v \in]0; 1]$, and u_0 is a scale parameter. We draw N_{MC} values of the new variable, v_a , $a = 1, \dots, N_{MC}$, uniformly and at random, and calculate the corresponding u_a . Angular variables Ω_a are uniformly sampled on the unit sphere or circle to obtain the points $\mathbf{y}_a = \mathbf{x}_\ell + u_a \Omega_a$ in the original space. The corresponding probabilities $P_t(\mathbf{y}_a)$ are computed using Simpson's method to calculate the integral over time in (1).

The algorithm of [11] stored and updated P_t , which made the computational time linear in t and in the size

²The code is publicly available from <http://www.lps.ens.fr/~barbieri>.




Fig. 1: Two-dimensional trajectories for $\gamma = .01$ (A) and $\gamma = .1$ (D), after one hit is received at time $t = 0$ ($P_0 = R$). B. squared distance to the origin, $d(t)^2$, vs. time t for four values of γ . C. spacing b between turnings vs. γ . The dotted line has slope $\frac{1}{2}$. Integration is done with $N_{MC} = 10^4$ points.

of the lattice. Our procedure recalculates P_t at each time step, which makes the computational time quadratic in t , and independent of the *a priori* infinite size of the space. Errors on P_t do not accumulate with time, and the accuracy is directly controlled by the number of Monte Carlo sampling points, N_{MC} . In addition, the map P_t is guaranteed to be accurately determined where it really matters, *i.e.* in the vicinity of searcher.

Initial stage of the search. – We first focus on the initial stage of the search process, which strongly depends on the *a priori* distribution, P_0 . A possible choice is the *one-hit* prior, $P_0 = R$, which means that the search starts only when a first hit is received at time $t = 0$ [11]. In two dimensions and before subsequent hits are detected, the search trajectories are Archimedean spirals [5,10] for a large range of values of γ (Fig. 1A). The squared distance of the searcher to the origin at time t , $d(t)^2$, increases linearly with t , and is, to a large extent, independent of γ (Fig. 1B). The spacing b between successive turnings is independent of time and increases as $\sqrt{\gamma}$ (Fig. 1C); the velocity $|\dot{\mathbf{x}}| \sim \frac{1}{b}$ decreases with the friction γ as expected.

The increase of b can be intuitively understood: the larger γ , the longer the searcher spends along the trajectory without receiving hits, and the more likely is the source to be located far away. For values of $\gamma (> .08)$ such that b would exceed the length ($= 1$) over which R decreases, the spiral nature of trajectories breaks down (Fig. 1D). We have run simulations with a modified potential V_t , where the arguments $\mathbf{y} - \mathbf{x}$ and $\mathbf{y}' - \mathbf{x}$ of the functions R in (4) were divided by a large factor ($= 10$). The regular spirals then disappeared, and looked like the trajectory in Fig. 1D, even for small values of γ .

Simulations with other prior distributions show that the

long distance behavior of P_0 is critical to the existence of spiral trajectories, while the behavior of P_0 close to the origin is irrelevant. Choosing $P_0(\mathbf{y}) \sim \exp(-|\mathbf{y}|/y_0)$ we obtain spirals as long as y_0 is not too large. The reason is that the spacing b is proportional to y_0 : spirals explore as much space as allowed by the prior. Spirals breakdown for the (γ -dependent) value of y_0 such that b exceeds 1.

Search trajectories in three dimensions display a more complex structure than their two-dimensional counterparts. Figure 2 shows the motion of the searcher with the *one-hit* prior, $P_0 = R$. Roughly speaking, the trajectory is constituted of subsequent shells of increasing radii, which are densely covered before a new shell is built. The distance to the origin, $d(t)$, is compatible with $t^{1/3}$ at large times, but grows faster at smaller times (Fig. 2). To better understand how trajectories develop in three dimensions, we have resorted to a small- \mathbf{x} expansion of the potential V_t . The relevant contributions to the equation of motion are, up to cubic order,

$$\gamma \dot{\mathbf{x}}(t) = \alpha_1(t) \mathbf{x}(t) + \alpha_2(t) \int_0^t dt' \mathbf{x}(t') + \int_0^t dt' \mathbf{x}(t') \left[\beta_1(t) |\mathbf{x}(t')|^2 + \beta_2(t) \mathbf{x}(t') \cdot \int_0^t dt'' \mathbf{x}(t'') \right] \quad (6)$$

where the coefficients $\alpha_i(t)$ have explicit analytical expressions. When \mathbf{x} is very small, only the linear terms matter. As $\alpha_1 \simeq \frac{\sqrt{2} \log t}{3e t} > 0$ for large t , the trajectory tends to follow a line radiating from the origin. However, the straight line is unstable against local bending since $\alpha_2 \simeq -\frac{3\sqrt{3} \log t}{e^2 t^2} < 0$. The trajectory thus acquires a spiral shape confined within the plane spanned by $\mathbf{x}(0)$ and $\dot{\mathbf{x}}(0)$. The presence of cubic terms ($\beta_1, \beta_2 > 0$), which are not constrained to lie in this plane, eventually lead to a cross-over from the quasi-bidimensional spiral to a fully three-dimensional trajectory (Fig. 2).

Replacing coefficients α_1, α_2 with the smaller values $\alpha_1 = \frac{a_1}{t}, \alpha_2 = -\frac{a_2}{t^2}$, with $a_1, a_2 > 0$, allows for an exact resolution of (7). A spiral is found when $a_1 < a_2$, with a radius and an angle growing as, respectively, t^ω and $\eta \log t$, where $\omega = \frac{a_1 - \gamma}{2\gamma}$ and $\eta = \sqrt{4a_2\gamma - (\gamma + a_1)^2}$. As γ increases so does the angular velocity ($\propto \eta$), while the growth exponent ω diminishes: spirals stop growing if γ is too large, a fact reminiscent of Fig.1D. Numerical resolution of (7) shows that this scenario is qualitatively unchanged when the logarithmic factors in α_1, α_2 are taken into account.

Pinning after a hit. – Examples of trajectories where the searcher receives hits at times $t_i > 0$ are shown in Fig. 3. Generally speaking, the trajectories are denser when more hits are received. After each hit i , the posterior distribution P_t is considerably reinforced in $\mathbf{x}(t_i)$ and its neighborhood. The searcher remains in the immediate vicinity for a certain time, t_w , until the search resumes. Informally speaking, the searcher makes sure that the source is not in $\mathbf{x}(t_i)$ before looking elsewhere. Imagine that the

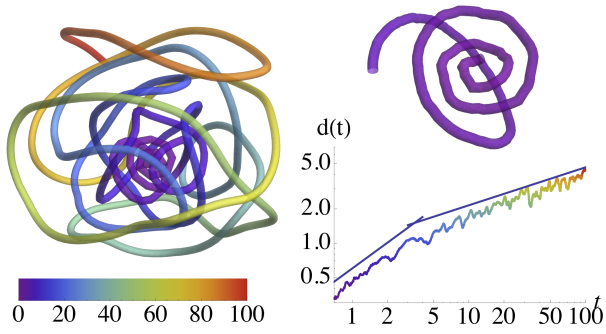


Fig. 2: A three-dimensional trajectory in the absence of hit for $\gamma = .01$ (left), with its quasi two-dimensional initial portion (top); the time axis is color coded. Bottom: distance to the origin, $d(t)$, compared to the power laws $t^{.75}$, then $t^{1/3}$.

searcher has not moved at all for a period of time t_w after the hit at time t_i . The posterior distribution is then

$$P_{t_i+t_w}(\mathbf{y}) \propto P_{t_i^-}(\mathbf{y}) R(\mathbf{y} - \mathbf{x}(t_i)) e^{-t_w R(\mathbf{y} - \mathbf{x}(t_i))}. \quad (7)$$

Assuming the posterior distribution right before the hit, $P_{t_i^-}$, is smooth in the vicinity of $\mathbf{x}(t_i)$, the potential $V_{t_i+t_w}(\mathbf{x}(t_i) + \mathbf{u})$ is a function of the small displacement $u = |\mathbf{u}|$ and t_w only. There are a local maximum in $u = 0$, since $\alpha_1 > 0$ in (7), and a global minimum in $u_m(t_w) > 0$. For $t_w < .4$, $u_m < .01$ is smaller than the error on the position deriving from our Monte Carlo integration, while for $t_w > .5$, $u_m > .1$, and the displacement of the searcher is easily seen.

The pinning effect is an important feature of the continuous formulation of Infotaxis, and was not observed on a lattice. The reason is that, at each time-step, a whole lattice site probability was set to zero in [11], which created a strong repulsive effect for the searcher and prevented pinning. In the continuous space, however, the trajectory has zero measure and the repulsion is too weak to override the pinning. As the searcher gets closer to the source, the average delay τ between successive hits gets smaller. When $\tau \simeq t_w$, the searcher could, in principle, come to a complete halt. The distance from the source for which this happens, d_{halt} , depends on the dimension: $d_{\text{halt}} = .1$ for $D = 2$, $d_{\text{halt}} = .3$ for $D = 3$.

Performances. – We now introduce a source and observe the trajectory of the searcher reacting to hits (Fig. 3). We are interested in the probability that the search process is successful as a function of the initial distance to the source, d_0 . If the searcher reaches the neighborhood of the source of radius d_{halt} defined above the source is declared found. If the searcher misses this neighborhood and reaches a distance $d_{\text{fail}} \gg 1$ to the source such that new hits are highly unlikely, the searcher is declared to be lost. Examples of successful and unsuccessful trajectories are shown in Fig. 3.

Figure 4 (left) shows that the probability of success in

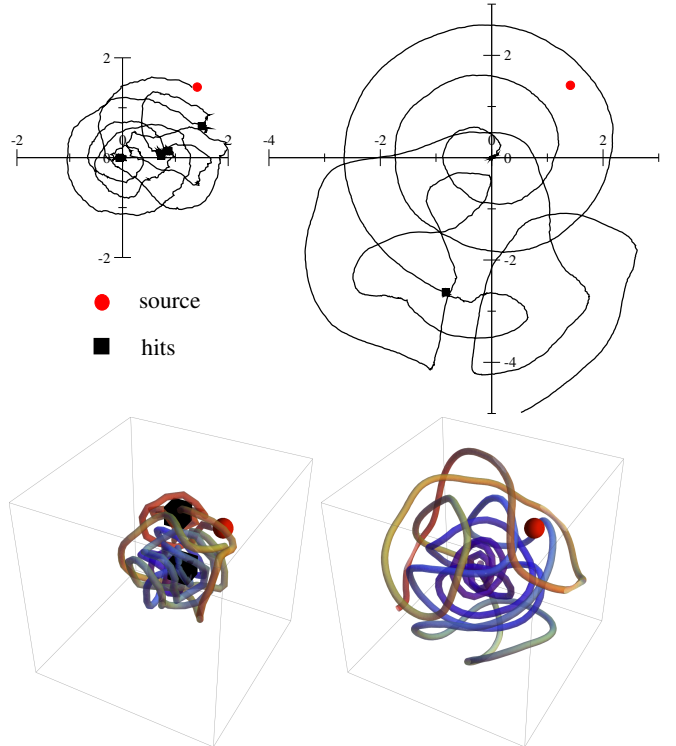


Fig. 3: Examples of search trajectories with hits in $D = 2$ (top, $\gamma = .02$) and $D = 3$ (bottom, $\gamma = .01$) dimensions. Trajectories on the left find the source, while searches on the right are not successful. The initial distance to the source is $d_0 = 2$. Red disks or spheres represent points at distance $< d_{\text{halt}}$ to the source. Black squares or cubes locate the hits; their size correspond to the amplitude of the erratic motion of the searcher during the pinning after a hit.

dimension $D = 2$ is compatible with unity for all distances d_0 smaller than a few units, in agreement with the findings of [11]. On the contrary, in dimension $D = 3$, the probability of success is definitely smaller than one, and is about .8 for distances d_0 ranging from 1 to 3 and for $\gamma = .01$. We have also measured the probability of success at fixed d_0 over a range of values of γ and found little variation (Fig. 4, right).

Figure 5A shows that the distribution of the search times t_s of successful runs has a positive skew and a roughly exponential tail not only in dimension $D = 2$ [11] but also in dimension $D = 3$. The exponential nature of the tail was checked for various values of the distance d_0 . Figure 5B shows that the average time to find the source, t_s , decreases with γ . The CPU time scales as $A(t_s/\Delta t)^2$, where $\Delta t = \gamma$ is chosen to obtain numerical stability, *i.e.* small enough local moves at any step ℓ . Hence, the CPU time is much larger for small friction (Fig. 5C). We find $A \simeq 3$ ms on one core of a 2.4 GHz Intel Core 2 Quad desktop computer, and for $N_{MC} = 10^4$ Monte Carlo steps. The CPU time can be decreased by choosing a smaller value for N_{MC} , or by increasing γ (and Δt) with time.

Case of a time-dependent friction $\gamma(t)$. – The results reported so far correspond to constant frictions




Fig. 4: Probability of success of Infotaxis as a function of the initial distance to the source, d_0 (left), and of the friction γ (right). Top points correspond to $D = 2$ dimensions, bottom points to $D = 3$. The numbers of runs is of about 200 for each point. All probabilities were obtained with $d_{\text{fail}} = 8$.

γ . We have generated trajectories with various time-dependent functions $\gamma(t)$, *e.g.* slowly increasing to reduce the scaling of the CPU time with t (Fig. 6A), or with the fixed-velocity requirement (Fig. 6B). From a qualitative point of view, we observe no drastic difference with the constant γ case, provided that $\gamma(t)$ does not exceed the maximal value ($\simeq .08$) at which $b \sim 1$ and spirals break down. For instance, the distance between turns in the spiral region of the trajectory in Fig. 6A can be deduced from the value of $b(\gamma(t))$ in Fig. 1C. In the fixed-velocity case of Fig. 6B, before the first hit is detected, $\gamma(t)$ shows regular oscillations and the trajectory has a spiral-like shape with b corresponding to the maximum value of $\gamma(t)$ ($\simeq .011$). Larger fluctuations are visible during the erratic motion after each hit, but $\gamma(t)$ remains smaller than the arrest value $\gamma \simeq .08$.

We have seen in Fig. 4 that the probability of success is essentially independent of d_0 and γ . This key result can be understood as follows. Far away from the source, hits are very unlikely and the trajectory develops as a spiral. As the area spanned by the trajectory is roughly independent of γ (Fig. 1B), the time it takes for the searcher to reach the D -dimensional sphere of radius 1 and centered in the source will be roughly independent of γ and will strongly increase with d_0 , while the time spent in the sphere and the number of hits received will be essentially independent of both γ and d_0 . We conclude that varying γ has little consequence on the performance of Infotaxis, as long as the spiral-like motion is possible.

Motion in presence of the 'average' signal. –

Certain characteristics of the search time distribution, such as the typical (most probable) value of t_s , can be assessed from the study of an abstract searcher receiving an average signal rather than discrete and sparse hits. To define an average posterior density $P_t^{av}(\mathbf{y})$ we remark that $P_t(\mathbf{y})$ contains the product of N_H stochastic R factors over the hits in (1). It is natural to define $P_t^{av}(\mathbf{y})$ through the

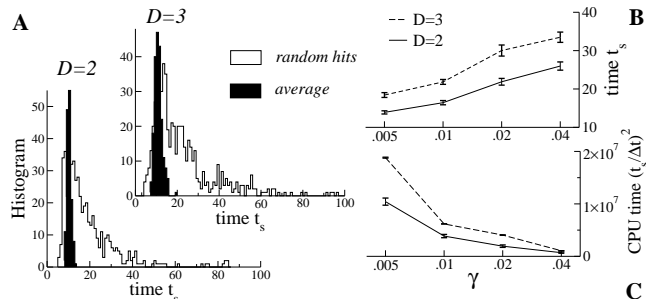


Fig. 5: **A.** histograms of the search times t_s in $D = 2$ ($\gamma = .02$) and $D = 3$ ($\gamma = .01$) dimensions for an initial distance $d_0 = 2$ to the source. Full histograms correspond to the average trajectories, contour histograms to trajectories with random hits. **B.** Average search time t_s as a function of γ . **C.** total CPU time as a function of γ , calculated as $(t_s/\Delta t)^2$.

average value of $\log P_t(\mathbf{y})$ over the hits:

$$P_t^{av}(\mathbf{y}) \propto e^{-\int_0^t dt' R(\mathbf{y}-\mathbf{x}(t')) + R(\mathbf{y}^*-\mathbf{x}(t')) \log R(\mathbf{y}-\mathbf{x}(t'))} P_0(\mathbf{y}) \quad (8)$$

up to a multiplicative normalization constant; here, \mathbf{y}^* denotes the location of the source as usual. We define the average search trajectory as the solution of equation (5) with the average $\langle \cdot \rangle$ (2) calculated over the measure (8). Average search trajectories are obviously smoother than trajectories with random hits, but have common features, such as a possible return towards the origin after having been close to the source.

Figure 5A shows that the search time for the average motion is in very good agreement with the typical search time for trajectories with random hits. This coincidence has been observed for all the frictions γ and distance d_0 to the source we have tested. Note that, while the average motion is fully deterministic, some noise is introduced to break the rotational invariance and select the initial phase of the spiral; this noise is responsible for the small width of the full histograms shown in Fig. 5A.

Figure 7 shows the average search trajectory and a random trajectory in $D = 2$ dimensions, together with the entropies of their posterior distributions. In the random case, the entropy abruptly decreases right after a hit, then increases until the next hit is received (due to the exponential time decay in P_t (1)). As for the average motion, the entropy shows weak oscillations (due to the spiral motion) superimposed to a smooth trend, which decreases as the searcher gets close to the source.

Conclusion. – In this letter we have presented a continuous-space version of Infotaxis, and have analyzed its behavior in two and three dimensions. When the initial distance to the source, d_0 , is of the order of the decay length of R , the probability that Infotaxis finds the target is essentially equal to unity in $D = 2$, and is smaller ($\simeq .8$) in $D = 3$ dimensions. The probability of success is roughly independent of $\gamma(t)$ in (5), while the search time

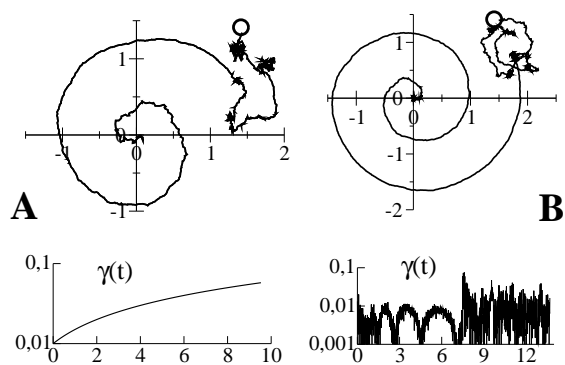


Fig. 6: **A.** A search trajectory for the time-dependent $\gamma(t)$ shown below. **B.** A trajectory with fixed velocity $v_0 = 2$, and $\gamma(t) = |\nabla_{\mathbf{x}} V_t|/v_0$. The source is located in $(\sqrt{2}, \sqrt{2})$ (circle).

and the CPU time strongly depend on the friction. The quadratic increase of the CPU time and the presence of the pinning effect make the computational cost increase as the searcher gets closer to the source. Note that the CPU time could be made linear in t_s (instead of quadratic) if the integral in (1) were restricted to the recent past, *i.e.* if the search had a finite memory.

The pinning of the trajectory by the hits is a direct consequence of the greedy nature of Infotaxis: the searcher moves to maximize the immediate gain in information, irrespectively of what could be gained on a longer time horizon. To overstep this greedy strategy consider the expected gain in information, $I[\mathbf{x}(t'); t < t' < t + \tau]$, when the searcher plans to move along a portion of trajectory $\mathbf{x}(t')$ during the current time t and the time $t + \tau$. The best portion of trajectory is determined through the maximization of I , minus a quadratic term $\propto \gamma \dot{\mathbf{x}}^2$ penalizing large velocities. While this variational calculation appears intractable for general τ , a systematic expansion in powers of τ is possible. To the lowest orders in τ the equation of motion becomes

$$\tau^2 (\nabla_{\mathbf{x}} \nabla_{\mathbf{x}} V_t(\mathbf{x})) \ddot{\mathbf{x}} + \gamma \dot{\mathbf{x}} = -\nabla_{\mathbf{x}} V_t(\mathbf{x}) \quad (9)$$

up to $O(\tau)$ corrections to the friction γ and to the force on the right hand side. The introduction of a finite-time horizon, τ , gives birth to an inertial term, with an effective mass tensor proportional to the curvature matrix of the potential (4). This inertial motion could help reduce the pinning following a hit, and avoid the slowing down of the searcher close to the source. The analysis of (9) and of the search trajectories is left for a future work.

Last of all, the shapes of the trajectories observed in two and three dimensions result from a trade-off between the self-repulsion of the trajectory (the searcher does not come again close a point where the source was not detected) and the confinement due to the hits or to the prior (the source is likely to be close to a detection). This trade-off is present in physical systems such as polyelectrolytes

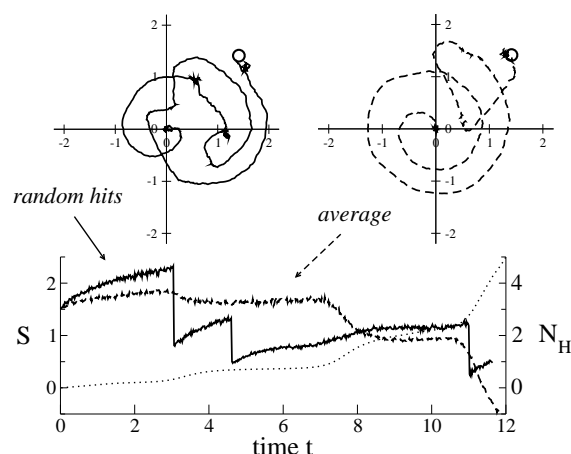


Fig. 7: Entropy $S(t)$ (bottom, left scale) for one trajectory $\mathbf{x}(t)$ obtained with random hits (top left, full curve, 3 hits are received) and the average trajectory (top right, dashed curve). The dotted line shows the average number of hits N_H (right scale) received along the average trajectory. The source is located in $(\sqrt{2}, \sqrt{2})$ (circle).

(charged polymers) confined in a volume or on a surface [3]. It is however unclear how far the analogy between the out-of-equilibrium process generated by Infotaxis and equilibrium polyelectrolytes could be pursued [2].

Acknowledgment: we thank J-F. Joanny, S. Leibler, A. Libchaber, T. Maggs, M. Vergassola for useful discussions. This work was partially funded by the ANR 06-JCJC-051 and 09-BLAN-6011 grants.

REFERENCES

- [1] Adler, J., *Science* **153**, 708 (1966).
- [2] Amit, D., Parisi, G. and Peliti, L., *Phys. Rev. B* **27**, 1635 (1983); Peliti, L. and Pietronero, L., *Riv. Nuovo Cimento* **10**, 1 (1987).
- [3] Angelescu, D.G., Linse, P., Nguyen, T.T. and Bruinsma, R.F., *Eur. Phys. J. E* **25**, 323 (2008).
- [4] Balkovsky, E. and Shraiman, B.I., *Proc. Nat. Acad. Sci.* **99**, 12589 (2002).
- [5] Barbieri, C., *Infotassi: un algoritmo di ricerca senza gradienti*, Master Thesis, University of Rome I (2007).
- [6] Berg, H.C., *Annual Review of Biophysics and Bioengineering* **4**, 119 (1975)
- [7] Hammersley, J.M. and Handscomb, D.C., *Monte Carlo Methods*, Taylor & Francis (1964).
- [8] Lochmatter, T. and Martinoli, A., *Experimental Robotics* **54**, 473 (2009).
- [9] Moraud, E. M. and Martinez D., *Frontiers in Neurobotics* **4**, 1 (2010).
- [10] Masson, J.-B., Bailly Bechet, M. and Vergassola, M., *J. Phys. A* **42**, 434009 (2009).
- [11] Vergassola, M., Villermaux, E. and Shraiman, B.I., *Nature* **445**, 406 (2007).
- [12] Vickers, N.J. and Baker, T.C., *Proc. Nat. Acad. Sci.* **91**, 5756 (1994).