Sector search strategies for odor trail tracking

Gautam Reddy[a,b], Boris I. Shraiman[c,1], and Massimo Vergassola[b,d,1]

[a]NSF–Simons Center for Mathematical & Statistical Analysis of Biology, Harvard University, Cambridge, MA 02138; [b]Department of Physics, University of California San Diego, La Jolla, CA 92093; [c]Kavli Institute for Theoretical Physics, University of California, Santa Barbara, CA 93106; and [d]Laboratoire de physique de l'École Normale Supérieure, CNRS, Paris Sciences et Lettres Research University, Sorbonne Université, Paris 75005, France

Edited by Eric Siggia, Center for Studies in Physics and Biology, The Rockefeller University, New York, NY; received April 27, 2021; accepted October 23, 2021

Ants, mice, and dogs often use surface-bound scent trails to establish navigation routes or to find food and mates, yet their tracking strategies remain poorly understood. Chemotaxis-based strategies cannot explain casting, a characteristic sequence of wide oscillations with increasing amplitude performed upon sustained loss of contact with the trail. We propose that tracking animals have an intrinsic, geometric notion of continuity, allowing them to exploit past contacts with the trail to form an estimate of where it is headed. This estimate and its uncertainty form an angular sector, and the emergent search patterns resemble a "sector search." Reinforcement learning agents trained to execute a sector search recapitulate the various phases of experimentally observed tracking behavior. We use ideas from polymer physics to formulate a statistical description of trails and show that search geometry imposes basic limits on how quickly animals can track trails. By formulating trail tracking as a Bellman-type sequential optimization problem, we quantify the geometric elements of optimal sector search strategy, effectively explaining why and when casting is necessary. We propose a set of experiments to infer how tracking animals acquire, integrate, and respond to past information on the tracked trail. More generally, we define navigational strategies relevant for animals and biomimetic robots and formulate trail tracking as a behavioral paradigm for learning, memory, and planning.

tracking | algorithm | behavior | optimization

Experimental studies demonstrate the ability of ants, dogs, humans, and rodents to track odor trails (1–6). Rodents accurately track trails in the dark, remaining close to the trail and casting when contact is lost (Fig. 1*A*) (5). Carpenter ants closely follow a trail while sampling it using a "crisscross" pattern with their two antennae (Fig. 1*B*) (1). Current models of this behavior rely on variants of chemotaxis (7) based on continuous estimates of the rising and falling odor gradients as the trail is crossed. One such strategy compares simultaneous odor concentrations detected by two spatially separated sensors (8). Yet, rats with a blocked nostril (5) and ants with a single antenna (1) are still able to track trails, although less accurately. An alternative chemotaxis strategy has the animal measuring odor gradients along its trajectory and turning when a significant decrease is perceived (5).

While chemotaxis-based strategies can allow for trail tracking when trails are continuous, they fail when trails are broken and gradients are absent, which is certainly relevant for animals tracking trails in the wild. In experiments with broken trails (1, 5), the absence of signal triggers casting, which is a fundamental feature shared with olfactory searches in a turbulent medium (9–11). Even though turbulent searches also feature sporadic cues, airborne odor signals tend to be localized in a cone, and even within the cone, the signal is highly fluctuating (12, 13). Therefore, beyond qualitative similarities between terrestrial trail tracking and airborne olfactory searches, the specific statistics of detections, geometric constraints, and behavioral patterns are distinct.

In contrast with chemotaxis-based algorithms, we propose an alternative framework built on the searcher exploiting past contacts with the trail to maintain an estimate of the trail's local heading and its uncertainty. A minimal memory of the approximate locations of the two most recent contacts suffices to delineate an angular sector of probable trail headings that radiates from the most recent detection point. The resulting "sector search" provides a quantitative description of trail-tracking behavior that unifies its various phases and yields specific experimental predictions.

Results

We first show that reinforcement learning (RL) based on the sector search idea can recapitulate natural behavior. An RL agent in this scheme learns to traverse the trail as quickly as possible while minimizing the probability of losing it (*Materials and Methods* has details). Our in silico RL experiments show that general aspects of animal tracking behavior naturally emerge (Fig. 1 *C* and *D*). Specifically, casts are observed around the most likely heading of the trail, and their amplitude is within the angular sector defined by the initial uncertainty σ of the trail's heading ϕ. The reason for the oscillatory pattern of casting is intuitive. Indeed, while moving along a path $\mathcal{C}$ without detecting the trail, the estimated heading's probability distribution $P(\phi)$ (Fig. 1*E*) is updated into $P_{\mathcal{C}}(\phi)$ as

$$P(\phi) \to P_{\mathcal{C}}(\phi) \propto \Gamma_{\mathcal{C}}(\phi) P(\phi) \qquad [1]$$

where $\Gamma_{\mathcal{C}}(\phi)$ is the probability of not detecting the trail headed along ϕ. Irrespective of the explicit form of $\Gamma_{\mathcal{C}}(\phi)$, the depletion of headings already explored generally leads to a bimodal posterior distribution, with the two modes at the edges of the angular sector (Fig. 1*E*). The search process is analogous to an agent "foraging" for the trail at two spatially separated patches. The

Significance

Surface-bound odor trail tracking is critical for the survival of terrestrial animals dependent on olfaction. Little is known about how animals track trails at the algorithmic level. In the present study, we propose that a tracking animal maintains a noisy estimate of where the trail is headed based on its past contacts with the trail. We show that virtual agents trained to exploit this strategy reproduce the tracking patterns of ants and rodents. The observed patterns emerge simply as a consequence of common geometric constraints, which also impose fundamental limits on how quickly an animal can track trails. A series of experiments is proposed to quantify how past experience and trail statistics shape tracking behavior.

Author contributions: G.R., B.I.S., and M.V. designed research, performed research, and wrote the paper.

The authors declare no competing interest.

This article is a PNAS Direct Submission.

See online for related content such as Commentaries.

[1]To whom correspondence may be addressed. Email: shraiman@ucsb.edu or massimo.vergassola@phys.ens.fr.

This article contains supporting information online at https://www.pnas.org/lookup/suppl/doi:10.1073/pnas.2107431118/-/DCSupplemental.

Published December 30, 2021.

BIOPHYSICS AND COMPUTATIONAL BIOLOGY

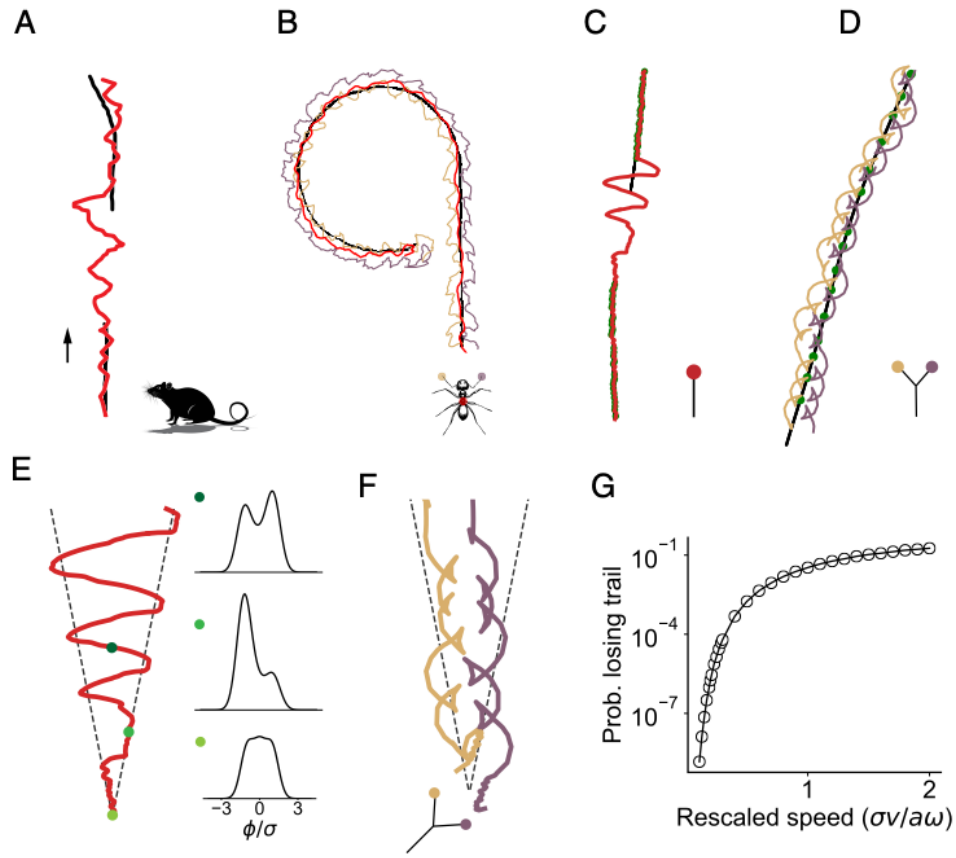

Fig. 1. Sample trail-tracking trajectories from previous experiments and our RL simulations. (*A*) A rat (head position in red) tracking a trail (in black). Note the wide casts on extended loss of contact with the trail. Data reproduced from ref. 5. (*B*) A carpenter ant tracking an odor trail (black) using a stereotyped crisscrossing strategy (1). (*C* and *D*) Sample trajectories obtained from RL for agents with one sensor (*C*) and two sensors (*D*) recapitulate experimentally observed tracking patterns in *A* and *B*. (*E*, *Left*) Search paths executed by RL agents with a single sensor upon loss of contact with the trail. (*E*, *Right*) The initial prior distribution (*E*, *Bottom Right*) over trail headings transforms into a bimodal posterior distribution (*E*, *Top Right* and *E*, *Middle Right*), which drives the oscillatory pattern of casting. (*F*) RL agents with two sensors show a characteristic crisscrossing pattern close to the last detection point. The search path is similar to the single-sensor agent at long distances (*SI Appendix*, Fig. S1*A*). (*G*) RL agents show a trade-off between tracking speed (rescaled by the sector angle σ, sensor size *a*, and sampling frequency ω) and the probability of losing the trail entirely.

emergence of oscillations is then understood in terms of marginal value theory (14, 15); we show using a minimal model of casting (*Materials and Methods*) that the turning point of a cast occurs when the marginal value of continuing on one side of the sector (i.e., without paying the cost of traveling) is outweighed by the probability of finding the trail on the opposite side.

We proceed now by establishing geometric limits on tracking speed. A typical RL curve for the probability of losing the trail vs. speed is shown in Fig. 1*G*. Its monotonicity epitomizes universal limits that "staying on the trail" imposes on tracking speed. Intuitively, searching slowly reduces the distance between detections (the interdetection interval [IDI]), decreasing the uncertainty in the estimate of the trail's heading and thus, the probability of losing the trail. However, these benefits come at the cost of slow forward progression along the trail. In contrast, moving quickly reduces the detection rate, leading to longer IDIs, increased uncertainty, and loss probability.

We quantify the above trade-off using simple scaling arguments. Suppose the tracking agent has a sensor of size *a*, samples at a frequency ω, and moves with a fixed forward speed *v*. As shown in Fig. 2*A*, the angle subtended by the detector at distance *r* from the last contact is a/r. The agent searching over an angular sector then scans at a rate $d\phi/dt = \omega a/r \simeq \omega a/vt$. Integrating the above expression for the angular rate, $\frac{v}{a\omega}\int d\phi = \int dt/t$, we obtain the typical time for searching over a sector angle σ: $t_c \sim \omega^{-1} e^{\sigma v/a\omega}$. The corresponding distance $L \sim vt_c$ is obtained using $r \sim vt$. The heading of the trail is known with uncertainty σ, which is the opening angle of the conical sector shown in Fig. 2*A*. Uncertainty is expected to depend on the distance L' from the previous detection as $\sigma(L') = (L'/\ell)^\gamma$, where ℓ and γ characterize the statistics of trails (below and Fig. 2*D*). Importantly, a stable strategy for long-term tracking requires that successive IDIs should on average be equal (i.e., $L = L'$). Combining $L = vt_c$ with $L' = \sigma^{1/\gamma}\ell$ and the expression for t_c, we finally obtain an upper bound on the tracking speed *v*:

$$\frac{v^{1+\gamma}}{a\ell^\gamma\omega^{1+\gamma}} = (\omega t_c)^{-\gamma}\log(\omega t_c) \leq \gamma^{-1}e^{-1}. \qquad [2]$$

Its maximum $v_{\max} \sim \omega(a\ell^\gamma)^{\frac{1}{1+\gamma}}$ defines the optimal stable tracking speed in terms of the tracker's sensory parameters and trail statistics. The basic element that leads to this bound is the geometric factor $1/r$ that underlies searching over an angular sector. The result from Eq. **2** that ωt_c is of order one ($e^{1/\gamma}$) explains experimental observations (Fig. 1) that tracking animals typically take only a few samples to reestablish contact with the trail.

The above argument implies that tracking speed depends on the trail statistics via the relation between uncertainty and the distance between points of contact. We use ideas from polymer physics to quantify how this relationship depends on geometric properties of the trails. Specifically, we ask how detecting the trail at a set of points $\boldsymbol{r}_0, \boldsymbol{r}_1, \boldsymbol{r}_2, \ldots$ (Fig. 2*B*) constrains the searcher's estimate of its future heading. We consider the case when the

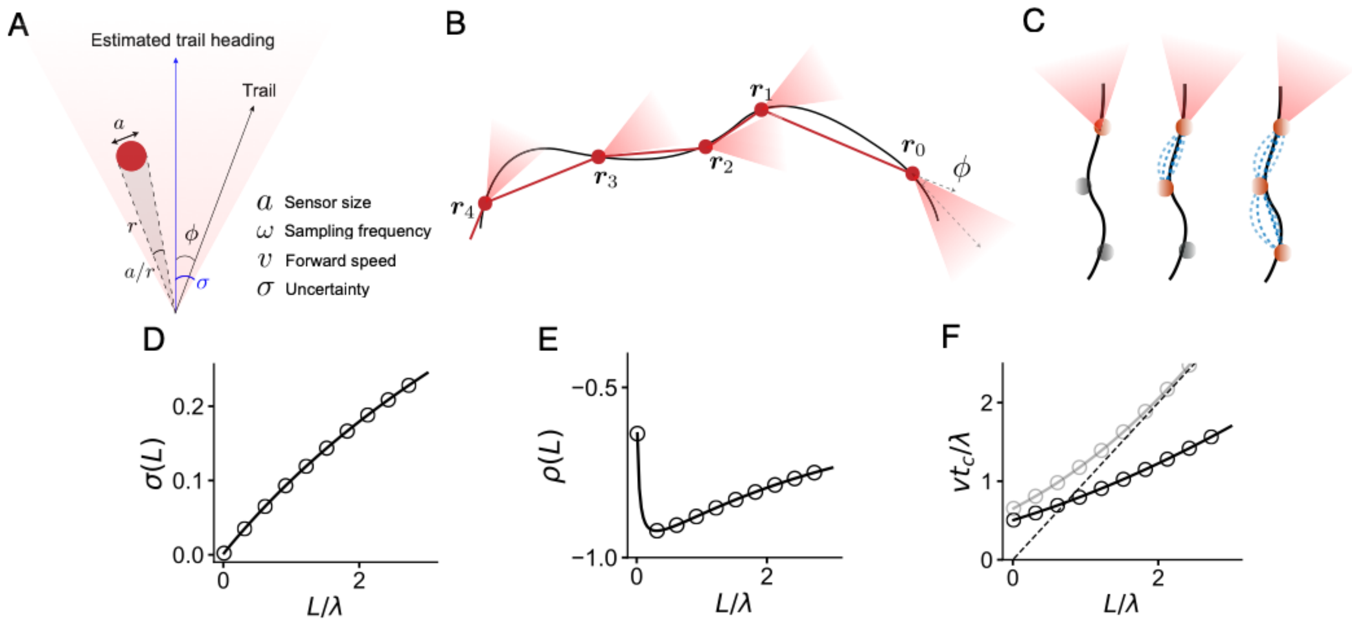

Fig. 2. History dependence and trail models. (*A* and *B*) Trail tracking naturally splits into distinct episodes punctuated by trail detections by the searcher. In each episode, we propose that the tracker searches for the trail using an estimate of the trail's heading updated based on the past points of contact with the trail and a model of trail statistics. We affix a polar coordinate system with the origin at the most recent contact point and the azimuthal angle defined relative to the estimated trail heading. The uncertainty σ fixes the angular width of the search. The searcher moves forward with a speed v while sampling at a frequency ω. A sensor of size a spans a/r radians at distance r, which determines the rate at which the angular space is searched. (*C*) To estimate where the trail is headed and its uncertainty from past contacts, the tracker can either use local anisotropy estimated from a single contact (*C*, *Left*) or extrapolate from previous points of contact using a model of trail statistics (*C*, *Center* and *C*, *Right*). In the latter case, the most likely trail paths (dashed blue lines) are similar to interpolated splines, which capture basic geometric notions of persistence in heading and curvature. (*D*) The uncertainty in trail heading (in radians) as a function of the distance, L, between points of contact for the GWLC model of trails discussed in the text. λ is the correlation length scale of the trail's curvature. *SI Appendix*, Fig. S3C illustrates the various scaling regimes exhibited by the GWLC model. (*E*) The correlation between trail heading at the most recent and second most recent points of contact for the GWLC model changes with the distance between these points, yet it is generally expected to be negative. (*F*) The expected search distance, vt_c, against the distance, L, between the previous two points of contact for $a/\lambda = 0.1$ and $v/a\omega = 5, 6.25$ (black and gray, respectively) (discussion above Eq. **2**). The point of intersection with the 45° dashed line is the condition for a stable tracking strategy. The gray curve corresponds to $v_{\max}$ beyond which tracking is unstable.

searcher keeps track of the two most recent points of contact with perfect memory of their location. A more extended memory is discussed further below; an imperfect memory adds to the uncertainty and can be easily accommodated within the framework developed below. Intuition for the two-point case is provided by the familiar "curve" tool in graphical design software, which draws a cubic spline through a set of prescribed points (Fig. 2*C*). The tool captures the simple intuition that tangents to a curve are continuous (i.e., the trail's heading has local persistence), which is a plausible, minimal assumption about trails. We show in *SI Appendix* that cubic spline interpolation corresponds to the most likely path (through a fixed set of points) in the so-called worm-like chain (WLC) ensemble (originally introduced for polymers) (16, 17). In this ensemble, the tangent direction undergoes diffusion with rate κ, and the uncertainty is then $\sigma = (\kappa L/3)^{\frac{1}{2}}$, which determines the two parameters: the scaling law, $\gamma = 1/2$, and the correlation length scale, $\ell = 3\kappa^{-1}$, in Eq. **2**. Actual trails could be smoother and have a well-defined curvature (the rate of change of heading) that persists on a characteristic length scale λ. We capture this ensemble of curves by introducing two additional parameters: persistence length λ and typical radius of curvature ξ (*Materials and Methods*). Uncertainty is then given by $\sigma \approx L/2\xi$ (hence, $\gamma = 1$ and $\ell = 2\xi$) at distances $L < \lambda$, while diffusive scaling is recovered at larger distances with an effective diffusivity $\kappa = 2\lambda\xi^{-2}$. This extended model defines a generalized worm-like chain (GWLC) ensemble with crossovers across the various regimes (*Materials and Methods*). In summary, the model leads to a "propagator," which encodes how information about past contacts is integrated to form an estimate of the trail's heading while taking into account geometric aspects of trails. A general feature is that the headings at two consecutive contacts are anticorrelated (Fig. 2*E*), which reflects the bending of the spline relative to the chord seen in Fig. 2*C*. We emphasize that although the general strategy of the agent depends on the statistical properties of the trail ensemble, the specific actions taken by the tracking agent along a particular trail, such as reorientation based on the most likely trail heading, will depend on the history of contact points via the propagator for the WLC (or GWLC) model.

Why and when do searchers need to cast? The question stems from our previous result that a few samples are typically sufficient to reestablish contact with the trail. To address it quantitatively, we consider again the setup of Eqs. **1** and **2**. The nondetection probability averaged over the ensemble of trails that pass through past contact points is

$$\Gamma_{\mathcal{C}} = \langle e^{-\omega \int_{\mathcal{C}} \frac{ds}{v} I_a(\boldsymbol{r}(s), \boldsymbol{y}(s))} \rangle_{\boldsymbol{y}}, \qquad [3]$$

where s parametrizes the searcher's path $\mathcal{C}$ and the Boolean indicator function I_a measures if the agent at $\boldsymbol{r}(s)$ is within sensing range a of the trail at $\boldsymbol{y}$ (i.e., the integral is the time spent in contact with the trail). Numerical simulations of the search show a power law scaling regime for $\Gamma_{\mathcal{C}}$, which is cut off at short distances by the initial surge and at long distances by trails escaping out of the casting envelope (Fig. 3 *A*–*C*). We proceed to explain these three regimes shown in Fig. 3*B*. Intuitively, at short radial distances $r < a/\sigma \sim v/\omega$ (the latter from Eq. **2**), the sensor covers the entire sector of likely headings, the searcher can just move forward, and $\Gamma_{\mathcal{C}} \propto e^{-\omega r/v}$ (Fig. 3*B*). Casting sets in if the searcher reaches, without detection, a distance $r \gtrsim$

https://doi.org/10.1073/pnas.2107431118