Editors' Suggestion

Optimal trajectories for Bayesian olfactory search in turbulent flows: The low information limit and beyond

R. A. Heinonen [1,*] L. Biferale [1] A. Celani [2] and M. Vergassola [3]

[1] Department of Physics and INFN, University of Rome "Tor Vergata", Via della Ricerca Scientifica 1, 00133 Rome, Italy

[2] Quantitative Life Sciences, The Abdus Salam International Centre for Theoretical Physics, 34151 Trieste, Italy

[3] Laboratoire de physique, École Normale Supérieure, CNRS, PSL Research University, Sorbonne University, Paris 75005, France

(Received 10 October 2024; accepted 11 March 2025; published 9 April 2025; corrected 26 August 2025)

In turbulent flows, tracking the source of a passive scalar cue requires exploiting the limited information that can be gleaned from rare, stochastic encounters with the cue. When crafting a search policy, the most challenging and important decision is what to do in the absence of an encounter. In this work, we perform high-fidelity direct numerical simulations of a turbulent flow with a stationary source of tracer particles, and obtain quasioptimal policies (in the sense of minimal average search time) with respect to the empirical encounter statistics. We study the trajectories under such policies and compare the results to those of the infotaxis heuristic. In the presence of a strong mean wind, the optimal motion in the absence of an encounter is zigzagging (akin to the well-known insect behavior "casting") followed by a return to the starting location. The zigzag motion generates characteristic $t^{1/2}$ scaling of the rms displacement envelope. By passing to the limit where the probability of detection vanishes, we connect these results to the classical linear search problem and derive an estimate of the tail of the arrival time pdf as a stretched exponential, which agrees with Monte Carlo results. We also discuss what happens as the wind speed decreases.

DOI: 10.1103/PhysRevFluids.10.044601

I. INTRODUCTION

A wide variety of animals, from plankton [1] to fruit flies [2] to dogs [3], combine olfactory sensing of odors with complex search behaviors which enable them to track the sources of the odors [4]. A famous, striking example is the ability of moths to locate, starting from $\gtrsim 100$ m away, potential mates using a single sex pheromone cue to which they are sensitive via olfactory organs [5–7]. In many settings, including those relevant to flying insects [8], this olfactory search behavior occurs in a fully turbulent flow. Turbulent conditions are also relevant to chemical-sensing robots, which may be used to locate hazardous substances, fires, or explosive devices [9–12].

The consequences of turbulence in the olfactory search problem, by now well studied [4,13–15], cannot be overstated. Odors may be modeled as passive scalars; the turbulent transport of passive scalars exhibits a remarkable degree of spatiotemporal intermittency, and the concentration fluctuations exhibit strongly non-Gaussian statistics [16]. As odors are emitted by a stationary and continuous source and mix into the surrounding flow, they form a characteristic concentration profile

*Contact author: robin.alrik.heinonen@edu.unige.it

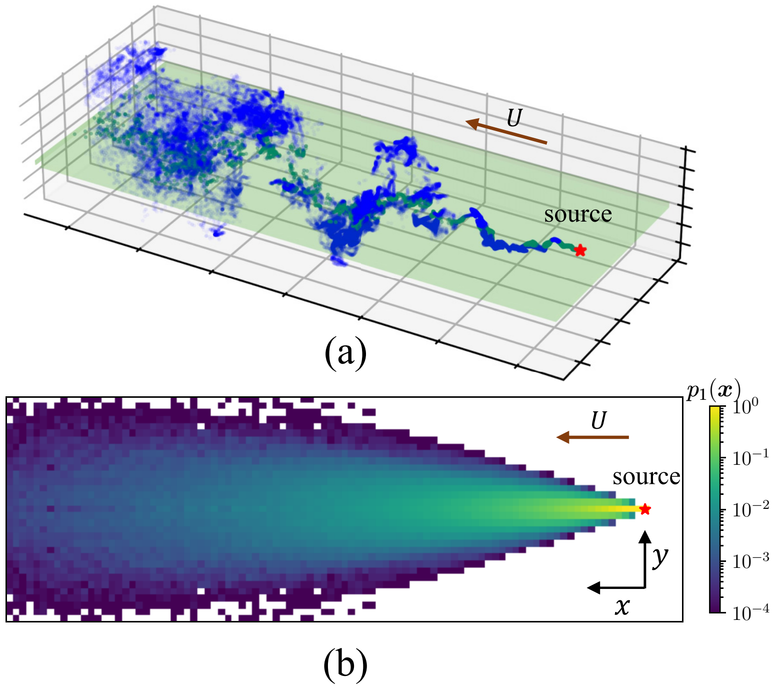

FIG. 1. Upper panel (a) snapshot from the DNS showing particles (blue) and the 2D slab (green) on which the data are coars grained. Lower panel (b) likelihood $p_1(\boldsymbol{x})$ of an encounter obtained from the DNS for $U/u_{\rm rms} \simeq 6.1$.

called a plume, whose shape depends on details of the flow. Within the plume, the concentration fluctuates wildly on account of the turbulence, forming a complex and patchy landscape (see Fig. 1 for a visualization from simulation). If a searcher is sensitive to an odor above some detection threshold, its encounters with the odor are randomized and typically rare (except when the searcher is close to the source) [15]. Moreover, the odor typically manifests in isolated patches rather than a continuous field, rendering simple gradient-based search strategies effectively useless.

How should a searcher, be it an animal or robot, respond to these limited, random encounters in order to locate the source in minimal average time? Looking to biology can provide some inspiration. During upwind flight, a diverse array of flying insects exhibit a behavior called "casting," which consists of broad zigzags perpendicular to the wind direction. Wind tunnel experiments suggest that this behavior is a very generic response to loss of contact with the odor cue [17,18]. (In biological literature, "zigzagging" often refers to a specific insect behavior consisting of small-scale crosswind motion combined with upwind flight. We will frequently use the word "zigzagging" as simply a qualitative description of trajectories, without intending reference to this behavior.)

For robotics applications, this observation may be used to formulate biomimetic strategies, as in [14]. Alternatively, if the searcher is equipped with a model for the spatially dependent likelihood of an encounter, the search problem may be formalized as a Bayesian partially observable Markov decision process (POMDP) [19]. This formulation allows the optimal, i.e., minimum average time, search strategy to be described by a dynamic programming (Bellman) equation, which may be approximately solved [20–22].

The model likelihood is the key physics input into the problem. Previous work on the POMDP approach imposed a model likelihood *ad hoc*, with the effects of turbulence captured by an effective diffusivity only. The realism of such a model is highly dubious in light of the complexity of turbulent transport; it is certain to fail, for example, in the large wind regime where ballistic transport dominates. (In the absence of a good model, the search problem could instead be attacked by model-free reinforcement learning techniques, as in Refs. [23,24].)

In order to maximize the realism of the model, in this work we study optimal trajectories with respect to likelihoods taken directly from simulation data. In particular, we examine the trajectories which result from the absence of an encounter; because encounters are rare, in a typical search, most of the search time will be spent in this low-information phase.

The optimal trajectories which we compute are complex and depend strongly on the strength of the mean wind. We focus on the case where the mean wind is strong: in this case the optimal motion is to zigzag upwind similarly to casting. This motion exhibits characteristic diffusive scaling in both directions. At some critical time, if the searcher has still not found the source, it should return to its starting location. We explain these results using a simple model that assumes encounters are very rare except close to the source. The diffusive scaling is explained by a reduction to the classical linear search problem and is used to predict the arrival time pdf.

II. METHODS

From the POMDP point of view, the searcher is a Bayesian agent living on a gridworld which maintains a posterior probability density (or "belief") $b(\boldsymbol{x})$ over the possible relative positions of the source with respect to the agent; critically, this relative position is unknown. The searcher makes an observation at every timestep, which we define as measuring if the local concentration is above or below some threshold c_{thr}. (In accordance with usual POMDP practice, the observation rate ν_Ω is the same as the action rate, i.e., $\nu_\Omega = v/\Delta x$, where v is the speed of the searcher and Δx is the grid spacing. In all figures, we will express time in units of ν_Ω^{-1} and distance in units of Δ.) The belief is updated in two ways. When the searcher moves ("takes an action" in POMDP jargon), the belief is transported according to $b(\boldsymbol{x}) \to b(\boldsymbol{x} - \delta\boldsymbol{x})$, where $\delta\boldsymbol{x}$ is the searcher's displacement under the action. Second, when the searcher makes an observation $\Omega = \theta(c(\boldsymbol{x}) - c_{\mathrm{thr}})$, it updates its belief according to Bayes' theorem [25]:

$$b(\boldsymbol{x}) \to b(\boldsymbol{x}) p_\Omega(\boldsymbol{x})/Z, \tag{1}$$

where Z is the appropriate normalization constant and we have introduced the notation $p_\Omega(\boldsymbol{x}) \equiv \Pr(\Omega|\boldsymbol{x})$. The likelihood $p_\Omega(\boldsymbol{x})$ must therefore be assumed to be known to the searcher. In the case where the mean wind $U\hat{x}$ is strong compared to flow velocity fluctuations, it can be modeled as proposed in [15]. The Lagrangian particle motion is predominantly ballistic on the short times relevant to this limit, which facilitates analysis of the tail of the concentration distribution. In particular, the probability that the instantaneous concentration is higher than c_{thr} is

$$p_1(\boldsymbol{x}) \propto \theta(x) x^{-\alpha} \exp\left[-\left(\frac{Uy}{u_{\mathrm{rms}} x}\right)^2\right] \exp(-x/x_D), \tag{2}$$

where u_{rms} is the typical fluctuation speed of the flow, θ is a Heaviside function, we have supposed the source is in $x = 0$, y is the distance from the plume centerline, $\alpha \geqslant 0$ is a free parameter that depends on the environment, and x_D is a characteristic decay length which depends on c_{thr}. This result is only valid sufficiently far from the source.

On the other hand, when the mean wind is small, long diffusive timescales become relevant to the Lagrangian particle motion and one instead expects

$$p_1(\boldsymbol{x}) \propto |\boldsymbol{x}|^{-1} \exp\left(-|\boldsymbol{x}|/\lambda\right) \exp(Ux/2D), \tag{3}$$

with D the effective (Taylor) diffusivity, $\lambda = \sqrt{D\tau/(1 + U^2\tau/4D)}$, and $\tau = \tau(c_{\mathrm{thr}})$ is a typical time for a fluid puff, initially at the source, to grow large enough that the concentration is less than c_{thr}.

Rather than impose one of these models directly, we used empirical models taken from simulations. We performed high-fidelity direct numerical simulations (DNS) of a turbulent flow ($\mathrm{Re}_\lambda \simeq 150$), while tracking the positions of tracer particles which are emitted by a stationary

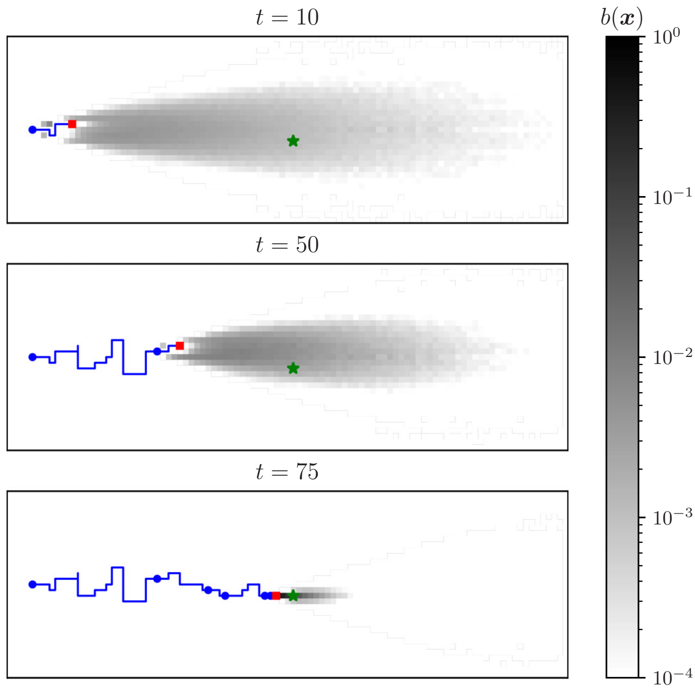

FIG. 2. Snapshots from a trajectory produced from a quasioptimal policy. The source is shown as a green star and the searcher as a red square. Each encounter is shown as a blue circle. The searcher's belief $b(\boldsymbol{x})$ on the position of the source is shown in grayscale. The searcher found the source a couple timesteps after $t = 75$.

point source. The system is advected by a uniform mean flow $U\hat{x}$; we repeated the simulation for $U/u_{\rm rms} \in \{0, 1.0, 2.1, 4.1, 6.1\}$. The tracers are taken as proxies for a passive scalar cue; this is a good approximation at high Péclet number. After coarse graining the particle data onto a 2D grid of about $\simeq 3250$ cells, we imposed a threshold concentration $c_{\rm thr} = 200$ particles/cell and computed the likelihood $p_1(\boldsymbol{x})$ that this threshold is instantaneously exceeded (see Fig. 1). For the largest values of U, the data are a good fit to Eq. (2) with $\alpha = 1$, and for the smallest values of U the data fit Eq. (3) (see Fig. 3 of the Supplemental Material [26]).

For each flow, we then used the SARSOP algorithm [27] to extract a quasioptimal search policy from the underlying dynamic programming equation [19,28] (Supplemental Material, Eq. (3) [26]). To study trajectories, we generated large ensembles of trajectories via Monte Carlo trials; for comparison, we did the same for the popular heuristic policy "infotaxis" [29], which seeks to minimize the entropy of the posterior. During search, observations were randomly drawn from the empirical likelihood. This artificially suppressed correlations between observations. It must be noticed that the no-hit trajectories under a policy that is optimized with respect to the uncorrelated statistics, as done here, will be entirely unchanged by the presence of correlations in the data during the search: only the arrival time statistics can change, since the optimal no-hit trajectory is a function of the likelihood and prior only. What happens when optimizing the policy including awareness of correlations is a technically different problem that will be discussed in a forthcoming paper.

An important detail is that all trajectories begin with an encounter at time zero, modeling the fact that the searcher has no incentive to search in the absence of an encounter. This induces a prior $b_0(\boldsymbol{x}) \propto p_1(\boldsymbol{x})$. To ensure consistency, we begin each trajectory by randomly selecting a source position from the prior. We also constrain the searcher to move in two dimensions, which both eases the computational burden of the POMDP and simplifies the analysis of the results. An example quasioptimal trajectory is shown in Fig. 2.

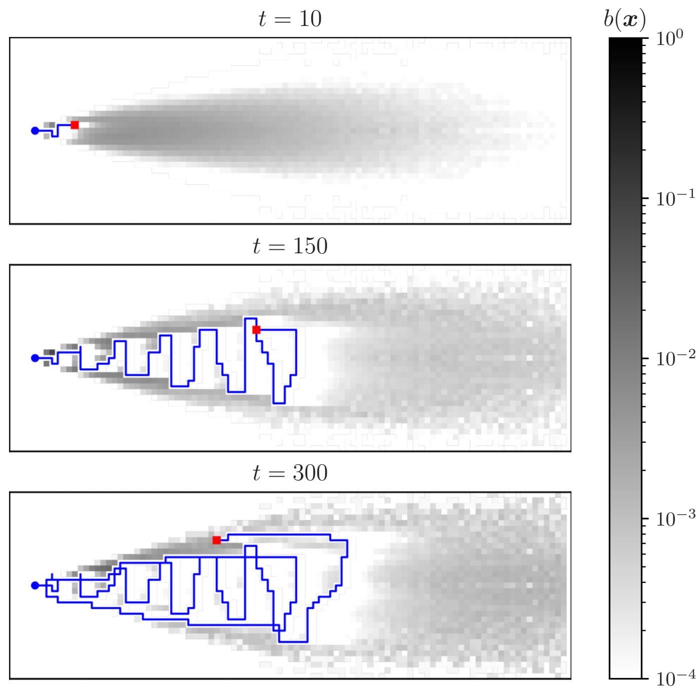

FIG. 3. Quasioptimal no-hit trajectories, for $U/u_{\mathrm{rms}} \simeq 6.1$. The searcher is shown as a red square, the searcher's belief is shown in grayscale, and the starting point is a blue circle.

III. RESULTS

Our results can be organized and understood through the following principle: because encounters are rare, the most important optimization the searcher must perform is to decide how it should move when it does not smell anything, and the single most likely trajectory is the one generated in the absence of any encounters after time zero. We will call this deterministic trajectory the "no-hit trajectory." To support this claim, we evaluated the probability $P_{\mathrm{nh}}(t)$ that, at time t, the searcher has not encountered the cue since time 0, given that it has not yet found the source. This is plotted for the largest wind speed in the upper panel of Fig. 4; the results for other wind speeds are very similar and shown in Figs. 4 and 5 in the Supplemental Material [26]. For both quasioptimal and infotactic motion, $P_{\mathrm{nh}}(t) \gtrsim 0.5$ for all times t, so the no-hit trajectory plays an outsize role.

We find that when the mean wind is strong, the quasioptimal no-hit trajectory begins by moving slowly upwind while zigzagging in the crosswind direction (Fig. 3). This motion strongly resembles casting. However, at some critical time, the searcher ceases this behavior and returns downwind to the vicinity of its starting position. From then on, the motion is dominated by upwind and downwind motion. Downwind returns have also been occasionally observed in wind-tunnel experiments after a long time without an encounter [30,31], but this behavior is far less well known. In contrast, infotaxis produces almost pure zigzagging motion in the absence of an encounter, and essentially never executes a downwind return.

Other examples of no-hit trajectories are supplied in Secs. VI C and D of the Supplemental Material [26]. As the wind speed decreases, the downwind return occurs earlier and earlier, and zigzagging quickly becomes irrelevant to the optimal motion. Instead, at moderate wind speed, a quasioptimal searcher exhibits repeated up- and downwind motion that begins to resemble a complex, asymmetric spiral. We suggest that the curves tend to approximately lie on isolines of the prior, which would imply the trajectories are essentially performing gradient ascent of the likelihood

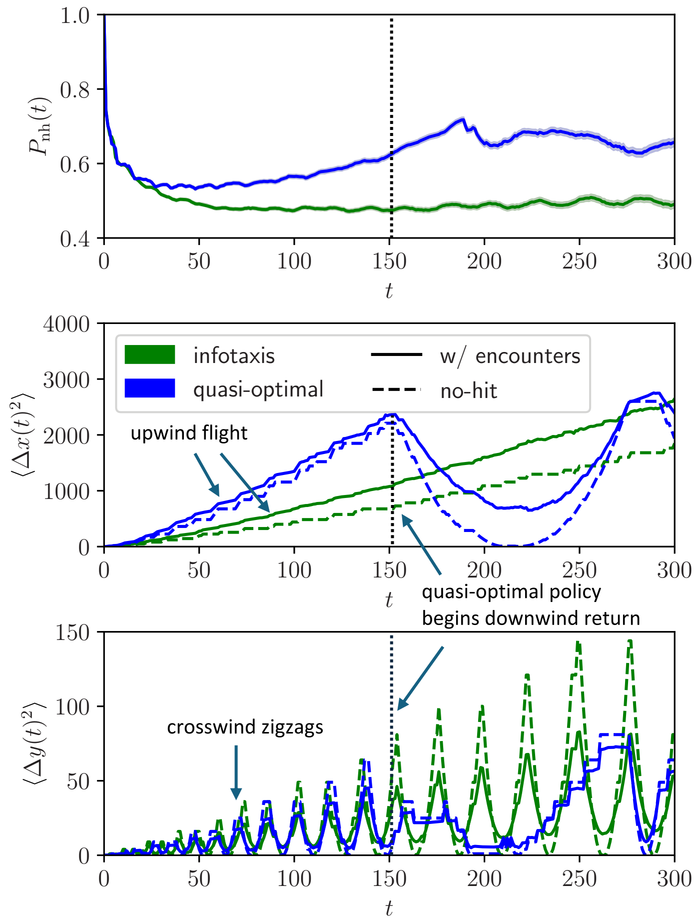

FIG. 4. Search trajectory statistics for $U/u_{\mathrm{rms}} \simeq 6.1$. Top panel: probability $P_{\mathrm{nh}}(t)$ that, at time t, the searcher has not encountered the cue since time 0, given that it has not yet found the source; error bars are shown as the standard error. Lower panels: mean squared displacement (MSD) in the x and y directions.

profile. When U and u_{rms} are comparable, the problem has very little symmetry and, accordingly, the trajectories are especially complex. At $U = 0$, both the quasioptimal and infotaxis trajectories resemble Archimedean spirals, consistent with previous studies [20,29,32].

These behaviors generate clear signatures in the mean squared displacement (MSD) of the searcher, $\langle \Delta x(t)^2 \rangle \equiv \langle (x(t) - x(0))^2 \rangle$ and $\langle \Delta y(t)^2 \rangle \equiv \langle (y(t) - y(0))^2 \rangle$, which we plot in Fig. 4 (lower panels). These are averaged over an ensemble of 10^4 trajectories, and compared with the squared displacements resulting from a no-hit trajectory. The results are very similar, suggesting that motion in absence of an encounter is dominating the MSD signal. For the larger values of U, the MSD in y exhibits fast oscillations, the amplitude of which grows like t. The MSD in x also grows approximately linearly. These results imply that the rms displacement envelopes in both x and y exhibit diffusive scaling: $x, y \sim t^{1/2}$.

In the same figure, we also show the MSD resulting from an infotactic trajectory. While both infotaxis and quasioptimal policy show the same behavior at small times, there is a clear qualitative difference at larger times: for the quasioptimal policy, at some critical time, the oscillations and linear scaling cease, and from $t \simeq 150 - 200$ the MSD in x decays rapidly like t^2, consistent with a ballistic downwind return. The time elapsed before returning becomes smaller as U decreases

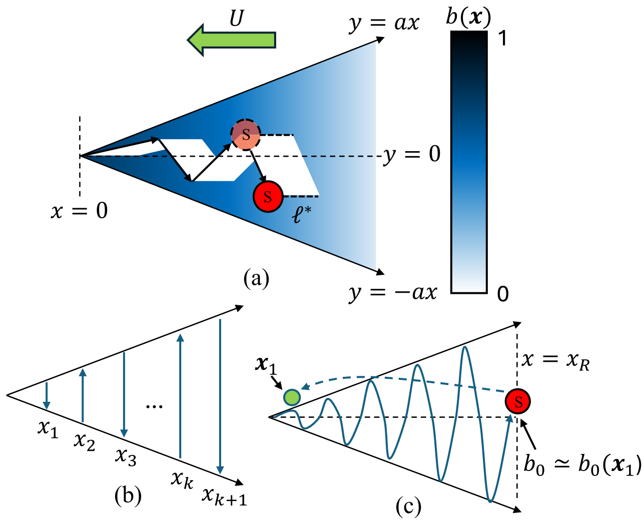

FIG. 5. Upper panel (a) in the simplified model, the prior decays as a function of x and is only supported on a cone symmetric about the x axis. The searcher (red circle, labeled "S") sweeps out probability mass in the upwind direction as it moves, over a characteristic length ℓ^*. Lower left panel (b) a zigzag trajectory may be approximated by a sequence of crosswind motions at upwind distances x_k, and the decision problem reduces to selecting the x_k. Lower right panel (c) when the searcher has swept out the cone up to an upwind distance $x_R \sim x_D$, the local value of the prior b_0 is comparable to that near the starting point and just outside the cone, $\boldsymbol{x}_1$. The searcher should return to $\boldsymbol{x}_1$.

(see Secs. VI A, C, and D in the Supplemental Material [26]), and eventually the zigzag oscillations disappear completely. The infotactic policy, on the other hand, always produces an almost pure zigzag motion, except at the smallest nonzero value of U.

IV. LOW-INFORMATION LIMIT

The results at large wind speed can be understood by means of a simplified model. In particular, we again assume the searcher initializes its belief as the prior $b_0(\boldsymbol{x})$ induced by the encounter at time zero, i.e., proportional to Eq. (2) for large enough x. However, consistent with the dominance of the no-hit trajectory, we assume successive encounters only occur when the searcher is close to the source—the likelihood is very close to unity in a strip directly downwind of the source (see yellow region in Fig. 1, lower panel), which is where $\langle c \rangle > c_{\text{thr}}$. Empirically, this strip has negligible length except when U is large; in the Supplemental Material (Sec. IV) we show that $\ell^* \propto U^{2/3}$ for large U [26].

Explicitly, we model the likelihood of an encounter after time $t = 0$ as

$$p_1(\boldsymbol{x}) = \begin{cases} 1, & 0 < x < \ell^*,\ y = 0 \\ 0, & \text{otherwise,} \end{cases} \tag{4}$$

where x and y are, respectively, the upwind and crosswind position of the source with respect to the searcher, and ℓ^* is the characteristic length of the strip. The searcher uses this likelihood to update its belief. We will call this simplified model the low-information limit. A similar model was used to describe sector search in Ref. [33].

In this limit, the result of the Bayesian update [Eq. (1)] is that the searcher will simply zero out its prior in a ℓ^* neighborhood (and renormalize) as it moves [see Fig. 5(a)]. The optimal trajectory will then be the one that most efficiently explores the prior's support, so the geometry of the prior plays an essential role.

We now determine the optimal trajectories in the low-information limit. Due to the Gaussian factor in Eq. (2), the prior will be strongly confined to a cone C given by $|y| < ax$ with $a \sim u_{\rm rms}/U$, so it is useful to approximate the confinement as exact: suppose the prior depends only on x inside the cone, and it is zero outside [see Fig. 5(a)]. We seek to minimize the expected time of arrival

$$J = \frac{1}{2a}\int_{-a}^{a} du \int_{0}^{\infty} dx\, p(x)T(x,u),$$

where $u \equiv y/x$, $T(\cdot)$ is the time that the searcher first sweeps out a given point in space, and $p(x) \equiv xb_0(x)$.

As a first observation, if $p(x)$ decays monotonically, then the optimal trajectory should (on average) visit states in order of increasing x [34]. Second, the searcher should make as few turning points in y as possible, since no probability mass is swept out if $dy/dt = 0$. Therefore, it is reasonable to assume the only turning points occur at the edges of the cone. Such a trajectory will necessarily look like an upwind zigzag, but we have not yet determined the scalings of x and y with time.

A zigzag trajectory can be (for large enough x) modeled by a sequence of purely crosswind moves at positions $x_1, x_2, \ldots$ [see Fig. 5(b)]. A key observation is that the searcher should not move further upwind than ℓ^* during any crosswind traversal, i.e., we have the constraint $x_{k+1} - x_k \leqslant \ell^*$ for each k. Otherwise, there will be points not swept out by the trajectory, and because the trajectory is monotonic in x, J will be infinite.

In the absence of the inequality constraint, the decision problem of choosing the optimal sequence $\{x_k\}$ is exactly equivalent to the linear search problem (LSP), a well-studied decision problem which asks for the turning points of a trajectory on a line which minimizes the expected arrival time to a target drawn from some known probability density [35–38]. This is shown in the Supplemental Material (Sec. V), after which we prove that if $p(x)$ decays exponentially or slower—as it does in our case—then $x_{k+1} - x_k$ asymptotically diverges for large k [26]. Thus, the inequality constraint must saturate far enough upwind, i.e., $x_{k+1} - x_k = \ell^*$, and one can show this gives rise to diffusive scaling as observed in the MSD data.

One may parametrize the resulting trajectory by a characteristic zigzag frequency ω and a distance traveled per half-period λ. Note that the envelope of the trajectory defines a cone of exploration with half opening angle $\varphi = \arctan(2v/\lambda\omega)$, (here, v is again the speed of the searcher). The preceding arguments suggest that, for a quasioptimal policy, $\lambda \sim \ell^*$ and the opening angle should correspond to that of the belief cone, i.e., $\tan\varphi \sim u_{\rm rms}/U$. This, in turn, implies $\omega \sim vU/\ell^* u_{\rm rms}$, leading to a model for the trajectory of the form (ignoring numerical factors)

$$x(t) \sim \sqrt{\frac{\ell^* vUt}{u_{\rm rms}}};\ y(t) \sim \sqrt{\frac{\ell^* v u_{\rm rms} t}{U}} \cos\sqrt{\frac{vUt}{\ell^* u_{\rm rms}}}, \tag{5}$$

with asymptotically constant speed $\sim v$ in an rms sense. Empirically, the quasioptimal trajectory selects larger λ but approximately equal ω as compared to infotaxis; this is a greedier choice that means the no-hit trajectory explores a narrower cone before the downwind return begins (the cones of exploration are compared in Figs. 9 and 10 in the Supplemental Material [26]).

We therefore conclude that the zigzag motion with $t^{1/2}$ scaling is a consequence of the fact that (i) the likelihood decays exponentially with x (and not faster), ii) the likelihood decays faster than exponential outside the cone, and (iii) the likelihood is small outside of a strip in front of the source. However, the argument [in particular, point (ii)] rests on the approximation $\exp[-(u_{\rm rms}y/Ux)^2] \ll \exp(-x/x_D)$. When the searcher reaches an upwind distance $x_R \sim x_D$, this approximation breaks down, and the local prior probability will be comparable to that just outside the cone, near the starting point [see Fig. 5(c)]. This means that there is no continued benefit to search upwind, and the searcher should instead return downwind—as we indeed observe in the quasioptimal policy.

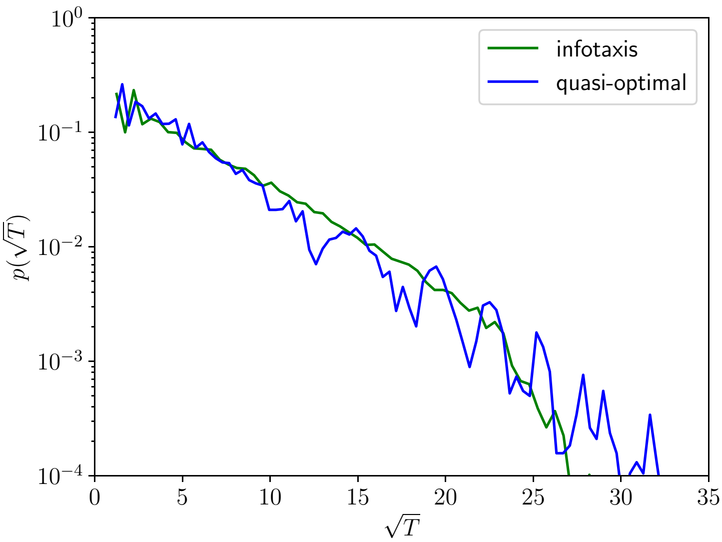

FIG. 6. Pdf of the square root of the arrival time $\sqrt{T}$, for $U/u_{\mathrm{rms}} \simeq 6.1$, obtained using 10^5 Monte Carlo trials.

We cannot analytically predict the optimal trajectories at finite U [and moreover the approximation Eq. (4) becomes untenable], but the low-information limit still grants some insights. Using the analysis of Ref. [15], one expects $x_D \sim u_{\mathrm{rms}}^{-1}(SU/c_{\mathrm{thr}})^{1/2}$ (where S is the emission rate) when U is large enough that ballistic single- and two-particle dispersion scaling still holds on relevant timescales. This is in basic agreement with our observation that downwind returns occur earlier as U decreases. Eventually, as U becomes comparable to u_{rms}, diffusive motion begins to dominate trajectories of Lagrangian particles on timescales of interest. In this regime, the likelihood is no longer well confined to a cone and gains some support even upwind of the source (this is visible in our simulations at $U/u_{\mathrm{rms}} \simeq 1.0$). This complexification of the prior geometry leads to concomitantly complex optimal trajectories (see Secs. VI A, C, and D in the Supplemental Material [26]).

Finally, a useful consequence of the $t^{1/2}$ scaling at large U is that it allows us to predict the tail of the arrival time pdf analytically. In the low-information limit, the arrival time is dominated by the time of the first encounter after $t = 0$, which is the time it takes to reach a ℓ^* neighborhood of the source. Therefore, we can estimate the arrival time pdf $p(T)$ as

$$p(T) = \int d^2\boldsymbol{x}\, b_0(\boldsymbol{x})\delta(T - T^*(\boldsymbol{x})), \tag{6}$$

where $T^*(\boldsymbol{x})$ is the time that the optimal trajectory first reaches $\boldsymbol{x}$. $b_0(\boldsymbol{x})$ is proportional to the likelihood, which is given by Eq. (2) for large U, and we have from Eq. (5) that $T^*(x, y) \sim (u_{\mathrm{rms}}/U)(x^2/v\ell^*)$. This gives

$$p(T) \propto T^{-\alpha/2} \exp(-k\sqrt{T}), \tag{7}$$

with $k \sim x_D^{-1}\sqrt{Uv\ell^*/u_{\mathrm{rms}}}$.

In Fig. 6 we plot the pdf of $\sqrt{T}$ for $U/u_{\mathrm{rms}} \simeq 6.1$, where T is the time of arrival. There is a broad range $0 \leqslant T \lesssim 500$ where the pdf appears to decay exponentially, implying $p(T) \sim \exp(-k\sqrt{T})$ as predicted. The agreement with exponential decay is more precise for infotaxis; we could have anticipated this since infotaxis was observed to obey $t^{1/2}$ scaling for the entire trajectory, whereas this scaling breaks down for the quasioptimal policy once it begins its downwind return.

V. CONCLUSION

In summary, our results suggest that the ubiquity of casting as a response to loss of plume contact may be due to its being the optimal no-hit trajectory in the limit of low information, in the presence of strong mean wind. The same limit makes nontrivial quantitative predictions about the temporal scaling of the trajectories and the arrival time pdfs, which agree with results obtained using quasioptimal policies extracted from the dynamic programming equation.

Finally, while it is not the goal of this paper to precisely model the behavior of any specific animal, it is nevertheless useful to discuss the relationship between our results and insect experiments. To our knowledge, Ref. [31] provides the most detailed quantitative experimental results characterizing moth behavior after plume loss during upwind olfactory search. These experiments found, for one, that crosswind excursion distances increase roughly linearly with the period number, which agrees with our results and simply reflects the conical shape of the plume in the strong wind regime. On the other hand, that study also found that the upwind distance traveled decreases monotonically with period, whereas in our study this was found to be constant under a quasioptimal strategy. There are a number of possible explanations for this discrepancy—it could be related to the fact that the experiments were performed in three dimensions whereas we were constrained to two, or the fact that our searchers did not have access to instantaneous anemometric cues, or perhaps moths simply do not implement a quasioptimal Bayesian strategy.

Downwind returns in the context of plume loss were reported in Refs. [30,31], but essentially no quantitative data on the behavior were provided. In Ref. [31], the behavior is reported anecdotally; in Ref. [30], it is described as the consistent behavior of moths exhibited after a long time without contact with the plume. This latter result is qualitatively consistent with our findings on Bayesian searchers in an environment with large mean wind. Additional experiments would be necessary to establish quantitative agreement. It would be important for any such experiment to distinguish between downwind returns after plume loss from downwind flight executed in the service of anemometry (Refs. [39,40] discuss downwind flight in this context). One would expect anemometric downwind flight to occur much earlier in the trajectory and to not be correlated with plume loss.

ACKNOWLEDGMENTS

We thank Aurore Loisy, Massimo Cencini, Lorenzo Piro, and Fabio Bonaccorso for fruitful discussions and acknowledge useful interactions at the 2023 Les Houches summer school on turbulence. This work received funding from the European Union's Horizon 2020 Program under Grant Agreement No. 882340. We also acknowledge financial support under the National Recovery and Resilience Plan (NRRP), Mission 4, Component 2, Investment 1.1, Call for tender No. 104 published on 2.2.2022 by the Italian Ministry of University and Research (MUR), funded by the European Union – NextGenerationEU– Project Title Equations informed and data-driven approaches for collective optimal search in complex flows (CO-SEARCH), Contracts No. 202249Z89M, No. CUP B53D23003920006, and No. E53D23001610006.

DATA AVAILABILITY

The data that support the findings of this article are openly available [41,42], embargo periods may apply.

[1] F. Roozen and M. Lürling, Behavioural response of daphnia to olfactory cues from food, competitors and predators, J. Plankton Res. **23**, 797 (2001).

[2] E. Álvarez-Salvado, A. M. Licata, E. G. Connor, M. K. McHugh, B. M. King, N. Stavropoulos, J. D. Victor, J. P. Crimaldi, and K. I. Nagel, Elementary sensory-motor transformations underlying olfactory navigation in walking fruit-flies, Elife **7**, e37815 (2018).

[3] A. Thesen, J. B. Steen, and K. Doving, Behaviour of dogs during olfactory tracking, J. Exp. Biol. **180**, 247 (1993).

[4] G. Reddy, V. N. Murthy, and M. Vergassola, Olfactory sensing and navigation in turbulent environments, Annu. Rev. Condens. Matter Phys. **13**, 191 (2022).

[5] R. Cardé and T. Hagaman, Behavioral responses of the gypsy moth in a wind tunnel to air-borne enantiomers of disparlure, Environ. Entomol. **8**, 475 (1979).

[6] J. Kennedy, A. Ludlow, and C. Sanders, Guidance of flying male moths by wind-borne sex pheromone, Physiol. Entomol. **6**, 395 (1981).

[7] J. Elkinton, C. Schal, T. Onot, and R. Cardé, Pheromone puff trajectory and upwind flight of male gypsy moths in a forest, Physiol. Entomol. **12**, 399 (1987).

[8] J. Murlis, J. S. Elkinton, and R. T. Carde, Odor plumes and how insects use them, Annu. Rev. Entomol. **37**, 505 (1992).

[9] D. Martinez, O. Rochel, and E. Hugues, A biomimetic robot for tracking specific odors in turbulent plumes, Auton. Rob. **20**, 185 (2006).

[10] D. Martinez and E. M. Moraud, Reactive and cognitive search strategies for olfactory robots, *Neuromorphic Olfaction*, edited by K. C. Persaud, S. Marco, and A. Gutiérrez-Gálvez (CRC Press/Taylor & Francis, Boca Raton, FL, 2013), Chap. 5.

[11] H. G. Nguyen, A. Nans, K. Talke, P. Candela, and H. R. Everett, Automatic behavior sensing for a bomb-detecting dog, in *Proceedings of the SPIE 9468, Unmanned Systems Technology XVII, Baltimore, MD*, Vol. 94680K (SPIE, Bellingham, Washington, 2015), pp. 138–146.

[12] T. Wiedemann, A. J. Lilienthal, and D. Shutin, Analysis of model mismatch effects for a model-based gas source localization strategy incorporating advection knowledge, Sensors **19**, 520 (2019).

[13] E. Yee, P. Kosteniuk, G. Chandler, C. Biltoft, and J. Bowers, Statistical characteristics of concentration fluctuations in dispersing plumes in the atmospheric surface layer, Boundary Layer Meteorol. **65**, 69 (1993).

[14] E. Balkovsky and B. I. Shraiman, Olfactory search at high Reynolds number, Proc. Natl. Acad. Sci. USA **99**, 12589 (2002).

[15] A. Celani, E. Villermaux, and M. Vergassola, Odor landscapes in turbulent environments, Phys. Rev. X **4**, 041015 (2014).

[16] B. I. Shraiman and E. D. Siggia, Scalar turbulence, Nature (London) **405**, 639 (2000).

[17] J. Kennedy, Zigzagging and casting as a programmed response to wind-borne odour: A review, Physiol. Entomol. **8**, 109 (1983).

[18] S. A. Budick and M. H. Dickinson, Free-flight responses of Drosophila melanogaster to attractive odors, J. Exp. Biol. **209**, 3001 (2006).

[19] L. P. Kaelbling, M. L. Littman, and A. R. Cassandra, Planning and acting in partially observable stochastic domains, Artif. Intell. **101**, 99 (1998).

[20] A. Loisy and C. Eloy, Searching for a source without gradients: how good is infotaxis and how to beat it, Proc. R. Soc. A. **478**, 20220118 (2022).

[21] R. A. Heinonen, L. Biferale, A. Celani, and M. Vergassola, Optimal policies for Bayesian olfactory search in turbulent flows, Phys. Rev. E **107**, 055105 (2023).

[22] A. Loisy and R. A. Heinonen, Deep reinforcement learning for the olfactory search POMDP: A quantitative benchmark, Eur. Phys. J. E **46**, 17 (2023).

[23] S. H. Singh, F. van Breugel, R. P. Rao, and B. W. Brunton, Emergent behaviour and neural dynamics in artificial agents tracking odour plumes, Nat Mach Intell **5**, 58 (2023).

[24] M. Rando, M. James, A. Verri, L. Rosasco, and A. Seminara, Q-learning with temporal memory to navigate turbulence, eLife **13**, RP102906 (2024).

[25] G. E. Box and G. C. Tiao, *Bayesian Inference in Statistical Analysis* (John Wiley & Sons, New York, 1992).

[26] See Supplemental Material at http://link.aps.org/supplemental/10.1103/PhysRevFluids.10.044601 for detailed information about methods for the DNS and POMDP solution, as well as additional calculations, plots, and movies, which includes Refs. [43,44].

[27] H. Kurniawati, D. Hsu, and W. S. Lee, SARSOP: Efficient point-based POMDP planning by approximating optimally reachable belief spaces, in *Robotics: Science and Systems IV* (MIT Press, Cambridge, 2008).

[28] E. J. Sondik, The optimal control of partially observable Markov processes over the infinite horizon: Discounted costs, Oper. Res. **26**, 282 (1978).

[29] M. Vergassola, E. Villermaux, and B. I. Shraiman, 'Infotaxis' as a strategy for searching without gradients, Nature (London) **445**, 406 (2007).

[30] M. A. Willis and E. A. Arbas, Odor-modulated upwind flight of the sphinx moth, *Manduca sexta* L, J. Comp. Physiol. A **169**, 427 (1991).

[31] L. Kuenen and R. T. Carde, Strategies for recontacting a lost pheromone plume: Casting and upwind flight in the male gypsy moth, Physiol. Entomol. **19**, 15 (1994).

[32] C. Barbieri, S. Cocco, and R. Monasson, On the trajectories and performance of infotaxis, an information-based greedy search algorithm, Europhys. Lett. **94**, 20005 (2011).

[33] G. Reddy, B. I. Shraiman, and M. Vergassola, Sector search strategies for odor trail tracking, Proc. Natl. Acad. Sci. USA **119**, e2107431118 (2022).

[34] To see this, note that if $x_2 > x_1$ but $\int_{-a}^{a} du\, T(x_1, u) > \int_{-a}^{a} du\, T(x_2, u)$, a new assignment of arrival times $\tilde{T}$ with $\tilde{T}(x_1) = T(x_2)$, $\tilde{T}(x_2) = T(x_1)$, and $\tilde{T}(x) = T(x)$ otherwise will result in a smaller value of J.

[35] A. Beck, On the linear search problem, Isr. J. Math. **2**, 221 (1964).

[36] A. Beck and M. Beck, Son of the linear search problem, Isr. J. Math. **48**, 109 (1984).

[37] A. Beck and M. Beck, The linear search problem rides again, Isr. J. Math. **53**, 365 (1986).

[38] S. Alpern and S. Gal, *The Theory of Search Games and Rendezvous*, International Series in Operations Research & Management Science, Vol. 55 (Springer Science & Business Media, New York, 2003).

[39] S. Guichard, D. Kriticos, A. Leriche, S. Worner, J. Kean, and D. M. Suckling, Evidence of active or passive downwind dispersal in mark–release–recapture of moths, Entomol. Exp. Appl. **134**, 160 (2010).

[40] S. D. Stupski and F. van Breugel, Wind gates olfaction-driven search states in free flight, Curr. Biol. **34**, 4397 (2024).

[41] L. Biferale, F. Bonaccorso, and R. A. Heinonen, Turb-Odor dataset, Smart-Turb database (2025), https://smart-turb.roma2.infn.it/init/routes/#/logging/view_dataset/9/tabmeta.

[42] R. A. Heinonen, OlfactorySearch software, GitHub (2025), https://github.com/rheinonen/OlfactorySearch.

[43] G. Shani, J. Pineau, and R. Kaplow, A survey of point-based POMDP solvers, Auton. Agents Multi-Agent Syst. **27**, 1 (2013).

[44] M. T. Spaan and N. Vlassis, Perseus: Randomized point-based value iteration for POMDPs, J. Artif. Intell. Res. **24**, 195 (2005).

Correction: Erroneous values for wind speed, U/u_{rms}, have been fixed and Figure 6 has been replaced.