

Optimal policies for Bayesian olfactory search in turbulent flowsR. A. Heinonen¹, L. Biferale¹, A. Celani², and M. Vergassola^{3,4}¹*Department of Physics and INFN, University of Rome “Tor Vergata,” Via della Ricerca Scientifica 1, 00133 Rome, Italy*²*Quantitative Life Sciences, Abdus Salam International Centre for Theoretical Physics, 34151 Trieste, Italy*³*Laboratoire de Physique, École Normale Supérieure, CNRS, PSL Research University, Sorbonne University, 75005 Paris, France*⁴*Department of Physics, University of California San Diego, La Jolla, California 92093, USA*

(Received 14 July 2022; accepted 27 March 2023; published 16 May 2023)

In many practical scenarios, a flying insect must search for the source of an emitted cue which is advected by the atmospheric wind. On the macroscopic scales of interest, turbulence tends to mix the cue into patches of relatively high concentration over a background of very low concentration, so that the insect will detect the cue only intermittently and cannot rely on chemotactic strategies which simply climb the concentration gradient. In this work we cast this search problem in the language of a partially observable Markov decision process and use the Perseus algorithm to compute strategies that are near-optimal with respect to the arrival time. We test the computed strategies on a large two-dimensional grid, present the resulting trajectories and arrival time statistics, and compare these to the corresponding results for several heuristic strategies, including (space-aware) infotaxis, Thompson sampling, and QMDP. We find that the near-optimal policy found by our implementation of Perseus outperforms all heuristics we test by several measures. We use the near-optimal policy to study how the search difficulty depends on the starting location. We also discuss the choice of initial belief and the robustness of the policies to changes in the environment. Finally, we present a detailed and pedagogical discussion about the implementation of the Perseus algorithm, including the benefits—and pitfalls—of employing a reward-shaping function.

DOI: [10.1103/PhysRevE.107.055105](https://doi.org/10.1103/PhysRevE.107.055105)**I. INTRODUCTION**

Certain flying insects depend on a remarkable ability to use olfactory cues to locate distant sources. Two salient, well-studied examples are female mosquitoes, which use a combination of carbon dioxide and other cues such as odors and heat [1,2] to find their human hosts, and moths, which track potential mates using emitted sex pheromones to which they are extremely sensitive.

In experiments, mosquitoes immediately begin flying upwind in the presence of fluctuating CO₂ plumes [3–5] from a distance of up to tens of meters away from the source [6], from which distances visual cues are not useful [7,8]. Male moths exhibit similar upwind search behavior when exposed to pheromones from females [9,10], and their maximum effective search range is even farther, on the order of 100 m [11]. Thus the relevant length scales for the search are macroscopic. On such length scales, effects due to turbulence dominate the transport of passive scalars; the turbulent transport induces a wildly fluctuating, intermittent concentration landscape. A fixed location tens of meters from a source may go long times—tens of seconds—without a detectable increase in local concentration [12]. Therefore, the insect cannot quickly estimate the concentration gradient, and simple search strategies like chemotaxis which rely on moving up the local gradient have minimal efficacy. A searching insect must then make effective use of the limited information it can glean from intermittent detections [13]. Information-theoretic studies suggest that the insect is best served by measuring concentrations as coarsely as possible, i.e., a binary low or high signal relative to some threshold [14,15].

In addition to insects, many mammals, such as dogs and rodents, alternate between sniffing the ground and sniffing the air when performing olfactory navigation [16–19]. This behavior implies that the mammals integrate airborne cues into their search and that they may depend in part on the same fundamental strategies as insects.

Both moths and mosquitoes exhibit similar behavior during their upwind flight, including a tendency towards zigzagging motion across the mean wind [20] or “casting”; fruit flies also exhibit similar behavior in response to attractive odor stimuli [21,22]). This suggests that good strategies for olfactory search are universal and above all depend on the physics underlying the dispersal of the attractant cues. Inspired by this observation, Balkovsky and Shraiman [23] proposed *cast-and-surge*, a simple heuristic policy for olfactory search. Cast-and-surge combines cross-wind casting, which helps to locate the downwind axis of the source where detections are most probable, with “surging,” direct upwind motion immediately after a detection.

The cast-and-surge decision-making algorithm depends only on the detection history. Another class of policies are *model-based*, in the sense that they take as input some hard-wired model for the statistics of detection, which one can imagine may be instinctual to an insect. An important example is *infotaxis* [24], which chooses actions by maximizing the information about the source location which, based on the information the agent currently has, is expected to be gained. Trajectories generated using infotaxis encompass a number of behaviors, including casting.

While a zoology of model-based heuristics exists, when one has an exact or approximate expression for the detection

probability, a natural question to ask is what kind of policy is *optimal*—that is, results in a minimum mean arrival time—given the statistics? What behaviors are seen in this optimal policy? And how do the various heuristic policies compare?

The mathematical language of partially observable Markov decision processes (POMDPs) allows us to formalize the search problem and specify the optimal strategy as the solution to the Bellman equation, a nonlinear functional equation. However, solving for the optimal policy in a POMDP is known to be a difficult computational challenge, and scalability to large problem sizes is a key issue. Recently, a promising effort [25] was made to solve the problem using a variant of deep reinforcement learning to obtain an approximate solution of the Bellman equation; however, it was left uncertain how effectively this approach scaled to large problem sizes, and the authors did not include a mean wind.

In this work, we use a POMDP algorithm called Perseus [26], coupled with reward potential shaping [27], to approximate the solution to the Bellman equation and obtain near-optimal policies for olfactory search on a large grid with thousands of points. Perseus was previously used for the search problem in Ref. [13], but the results presented were limited and purely qualitative. Presently we build model environments with three different characteristic emission rates, compute near-optimal policies for each environment, and compare the arrival-time statistics of the near-optimal policies with those of a few interesting heuristic policies (including two versions of infotaxis). We find that the near-optimal policy found by our implementation of Perseus successfully outperforms all the tested heuristics in each environment. This establishes a reasonably scalable baseline for the POMDP solution of this problem, allows for the study of the behaviors and statistics which characterize optimal search, and opens the possibility to apply the same algorithm to more complex search problems, for example, the case of multiple sources. We also suggest that the present work may have applications to robotics for the purpose of detection of hazardous chemicals or explosives; previous studies saw the design of robots for olfactory search using reactive policies [28,29] and using infotaxis [30,31].

The paper is organized as follows. In Sec. II we detail the search problem and its assumptions, and we introduce the simple mathematical model for detections which we use in this work. In Sec. III we review the POMDP formalism, cast the search problem in this language, and describe both the heuristic policies and the methods used to solve the POMDP directly. Important details include the choice of initial condition and the use of a reward-shaping function to accelerate convergence. In Sec. IV we present example trajectories and arrival time statistics for the various policies, including mean arrival times for several test problems and detailed probability density functions (pdfs). We also test the robustness of the policies under changes in the model environment and use the near-optimal policies to obtain the approximate best mean arrival time as a function of starting position. Finally, in Sec. V we summarize our results and discuss avenues for future research.

II. PROBLEM DESCRIPTION

We will consider the problem of a model insect, or agent, searching for a stationary source located somewhere upwind,

using some olfactory cue. We will constrain the motion of the agent to a two-dimensional plane, and we will discretize space into a rectangular grid (see Fig. 1). While this constraint is primarily made for computational ease (since very large POMDPs are extremely difficult to solve), it also models the fact that for problems of interest the source will typically be close to the ground. We will assume the agent begins its search after a detection event (see Sec. III B for more details on the initialization) and stops when it reaches the grid point corresponding to the source location. The goal is to find a search strategy which minimizes the mean arrival time to the source.

At each time step, the agent will make either a detection or a nondetection of the cue, which may be interpreted as the insect observing a concentration above threshold. Due to the random nature of turbulent mixing, detection will occur with some space-dependent probability, which will be small far enough downwind from the source and nearly vanishing upwind of the source [32]. We assume the Markov property: the detection events are independent in time and space. In principle, we could allow for multiple detections per time step, but in the most interesting case, detections are rare enough that more than one detection in a small time interval is exceedingly unlikely (unless the agent is very close to the source).

Physically, the discretization may be understood as assuming the agent flies at a fixed speed v and measures the local concentration by integrating over a characteristic sampling time Δt (the decision time step). The grid spacing is then $\Delta x = v\Delta t$.

We seek an expression for the mean rate of detections in a statistically steady turbulent flow, as a function of spatial position. We model the turbulent environment with an effective turbulent diffusivity D (assumed to be much larger than the collisional viscosity). We impose a mean wind $V\hat{x}$ and fix a source with emission rate S at the origin, $\mathbf{r}_0 = 0$. The advection-diffusion equation is then

$$\partial_t c + V\partial_x c = D\nabla^2 c + S\delta(\mathbf{r}) - c/\tau, \quad (1)$$

where c is the concentration field and τ is a particle lifetime that can be identified as a turbulent mixing time or coherence time. A simple dimensional estimate for the turbulent diffusivity is $D \sim \ell\tilde{v}$, where \tilde{v} is the rms fluctuation velocity and ℓ is the turbulent mixing length. In this framework, we then have $\tau \sim \ell/\tilde{v}$. In three spatial dimensions, the stationary solution of Eq. (1) is

$$c(\mathbf{r}) = \frac{S}{4\pi D|\mathbf{r}|} \exp\left(\frac{V\mathbf{r} \cdot \hat{x}}{2D} - \frac{|\mathbf{r}|}{\lambda}\right) \quad (2)$$

with $\lambda \equiv \sqrt{D\tau/(1 + V^2\tau/4D)}$. Using Smoluchowski's expression for the rate of encounters of a sphere of radius a with molecules diffusing with diffusivity D [33], we obtain an estimated number of encounters in a time interval Δt ,

$$h(\mathbf{r}) = 4\pi aD\Delta t c(\mathbf{r}), \quad (3)$$

where a is a characteristic length scale of the searcher. We then imagine that the animal, which is constrained to the plane of the source, treats the number of detections per time step as a Poisson variable with rate h .

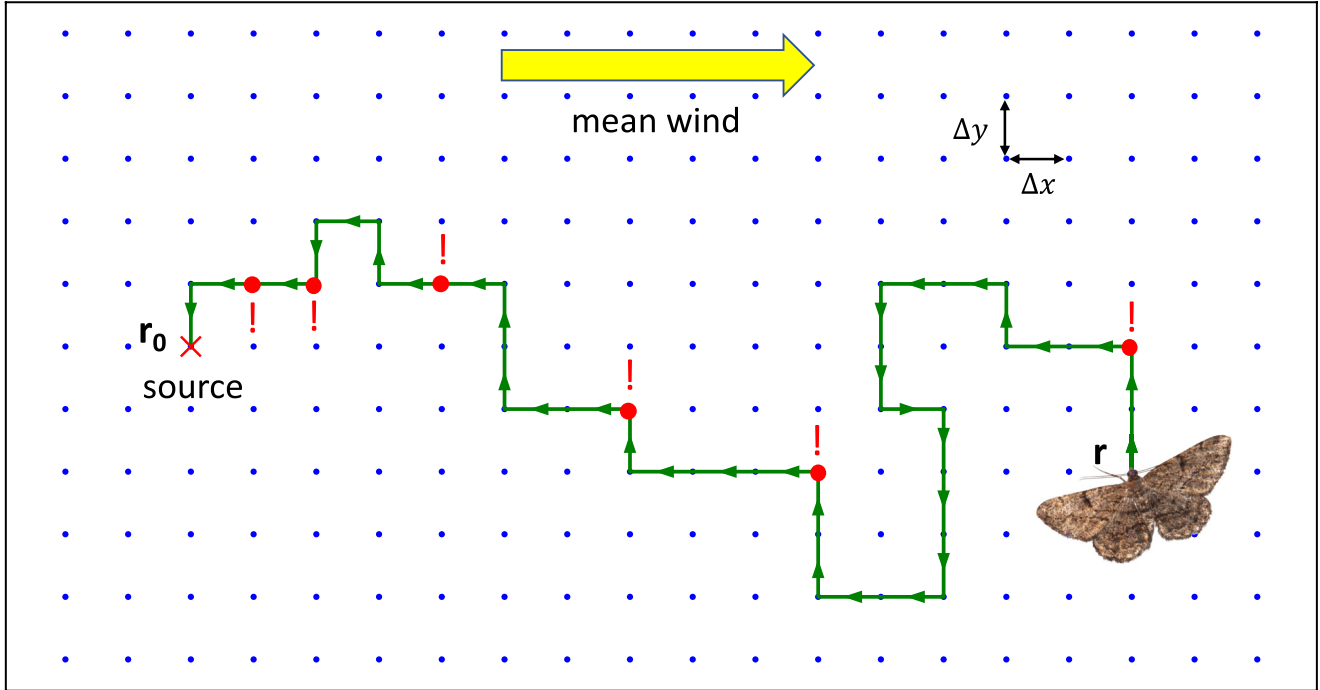


FIG. 1. Basic schematic of the search POMDP. The model insect searches for the source (red X) on a 2D grid, in the presence of a mean wind. At each time step, it either detects the cue (red circle with “!”) or makes no detection, with a specified probability; these observations are used to update the belief of where the source is. It then moves to an adjacent grid point. The search terminates when it finds the source or when some maximum search time is exceeded.

In this work the agent will search in a toy environment wherein the detections are generated artificially, drawn from the distribution specified by the diffusive model. (Understanding the robustness of policies trained using the present model when applied to more realistic environments remains one open important problem that goes well beyond the scope of this paper.) But, of course, the model is a simplification of the turbulence physics; see, for example, [32] for a more sophisticated treatment. In reality, the dynamics of the odor molecules will be neither purely diffusive nor purely ballistic, and moreover detection events will have a nonzero spatiotemporal correlation. On the other hand, the present diffusive model has seen significant use in past work (e.g., [24,25,34]) and in any case leads to a good benchmark search problem which is far from trivial to solve. Above all, the model captures the key phenomenological features that make the problem difficult and interesting [32]: detections are *stochastic* (with some space-dependent probability) and *rare* enough that fine-grained information about the local concentration field is not very useful. As long as correlations are neglected, differences between the toy model and real data will be quantitative rather than qualitative.

After one introduces a grid spacing Δx , the model can be parametrized by three nondimensional quantities: the nondimensional emission rate $\bar{S} \equiv aS\Delta t/\Delta x$, the nondimensional mean wind $V\Delta x/D$, and the nondimensional coherence time $V^2\tau/D$. The values we use for these quantities are shown in Table I.

A final, important assumption is that the agent has knowledge (say, through instinct) of the detection statistics implied

by Eqs. (2) and (3). This will be necessary to perform Bayesian inference.

III. METHODS

A. POMDP setup

We now cast the search problem in the language of a partially observable Markov decision process (POMDP) [35,36]. The fundamental ingredients of a POMDP are a state space S , a set of actions A , a set of observations O , and a reward function

$$R : S \times A \rightarrow \mathbb{R}.$$

At each time step, the *agent* is in some state $s \in S$ and selects an action $a \in A$, which causes the agent to transition from state s to s' with a specified probability $\Pr(s'|s, a)$.

For our purposes, the agent is a model insect living on an $N_x \times N_y$ grid with spacings Δx and Δy (see Fig. 1). The source is fixed at a point \mathbf{r}_0 . The state of the agent is its relative position with respect to the source $\mathbf{s} \equiv \mathbf{r} - \mathbf{r}_0$, with \mathbf{r} the agent location. The actions are to simply to move to an adjacent grid point, $A = \{(\Delta x, 0), (-\Delta x, 0), (0, \Delta y), (0, -\Delta y)\}$, so that after taking action \mathbf{a} the new state is $\mathbf{s}' = \mathbf{s} + \mathbf{a}$ with probability 1—unless it attempts to leave the grid, in which case the agent is unmoved, or if it has found the source. The state of occupying the same grid point as the source is treated as an *absorbing* state, which is to say that no action will change the state. (More details are given in the Appendix A 1.)

The agent also receives a reward $R(s, a)$ for taking an action, which is *discounted* by a factor $0 < \gamma < 1$. We set the

TABLE I. Parameters and hyperparameters, relating to the turbulence physics, the POMDP definition and grid, and the Perseus algorithm.

Category	(Hyper)parameter	Description	Value(s) used
Environment	$V \Delta x / D$	Nondimensional mean wind	2
	$V^2 \tau / D$	Nondimensional turb. coherence time	150
	$\bar{S} \equiv a \Delta t S / \Delta x$	Nondimensional emission rate	0.25, 2.5, 25
	$\Delta x / \Delta y$	Grid-spacing ratio	1
	N_x	Grid points along wind axis	81
	N_y	Grid points along cross-wind axis	41
POMDP solution	\mathbf{r}_0	Source position	(10,20) (grid spacing units)
	$ \mathcal{B} $	Number of Perseus beliefs	45 000
	$g(D)$	Reward-shaping function	$0.001D^2, 0.1D$
	π_B	Policy for belief collection	Infotaxis
	T_{wait}	Max. time to wait for first detection	1000
	γ	Discount factor	0.96, 0.98

reward to be unity for finding the source, and zero otherwise (other choices are easily seen to be equivalent, provided $\gamma < 1$; see Appendix A 2). This sets up an optimization problem, namely, to craft a *policy* for choosing actions which maximizes the expected total reward

$$\mathbb{E}[R_{\text{tot}}] = \sum_{t=0}^{\infty} \gamma^t \mathbb{E}[R(s_t, a_t)] = \mathbb{E}[\gamma^{T-1}], \quad (4)$$

where T is the arrival time to the source. The discount factor helps to regularize POMDPs by reducing the influence of times far in the future, and its value sets the extent to which the agent should prioritize immediate rewards vis-à-vis future rewards. This preference for short- or long-term rewards is quantified by a characteristic time called the *horizon* $\sim 1/\log \gamma^{-1} \simeq 1/(1-\gamma)$. Rewards which are earned at times in the future beyond the horizon are suppressed in the decision-making process. We refer the reader to Appendix A 2 for further discussion.

So far, we have described only a basic Markov decision process (MDP). The challenge of partial observability is that the agent does not have access to its state, and instead maintains a probability distribution over S called a *belief*, $b(s)$. This is of course relevant to the present search problem because the agent does not know where the source is. The belief lives in a $|\mathcal{S}| - 1$ -dimensional simplex: the set of vectors in S with nonnegative components summing to one. At each time step, after taking its action, the agent makes an observation $o \in \mathcal{O}$ which is used to update the belief using Bayes' rule

$$b_{o,a}(s') = \frac{\Pr(o|s', a) \sum_s b(s) \Pr(s'|s, a)}{\sum_{s,s'} \Pr(o|s', a) b(s) \Pr(s'|s, a)}, \quad (5)$$

where $b_{o,a}$ denotes the updated belief after taking action a and making observation o . Note that in our problem the transition probability is deterministic and, excluding the aforementioned edge cases where the agent has already found the source or tries to exit the grid, we can write $\Pr(s'|s, a) = \delta_{s',s+a}$. Also, the observation likelihood is independent of the previous action taken by the agent and depends only on its relative displacement from the origin, $\Pr(o|s, a) = \Pr(o|s)$. However, we will usually give expressions pertaining to POMDPs in a general form when possible.

The likelihood is set by the diffusive model of Sec. II as follows. We define three possible observations. First, the agent may discover that it has found the source, which occurs with probability $p_s = \delta_{s,0}$, where δ is a Kronecker function. Otherwise, the agent may observe either a detection with probability

$$\Pr(o = \text{det}|s) = \{1 - \exp[-h(s)]\}(1 - p_s) \quad (6)$$

or a nondetection with probability

$$\Pr(o = \text{nondet}|s) = \exp[-h(s)](1 - p_s). \quad (7)$$

That is, the number of detections in a time step is treated as a Poisson process with rate h , and any number of detections ≥ 1 is considered equivalent. The factor $1 - p_s$ enforces the fact that the agent observes the source if it finds it. Note that the observations are defined so that $\sum_{o \in \mathcal{O}} \Pr(o|s, a) = 1$. The detection likelihood for our choice of model parameters, with $\bar{S} = 2.5$, is shown in Fig. 2.

Now, the problem becomes that of finding a good policy $\pi : b \mapsto a$ mapping each belief to an action which yields a maximal expected total reward (i.e., a short mean arrival

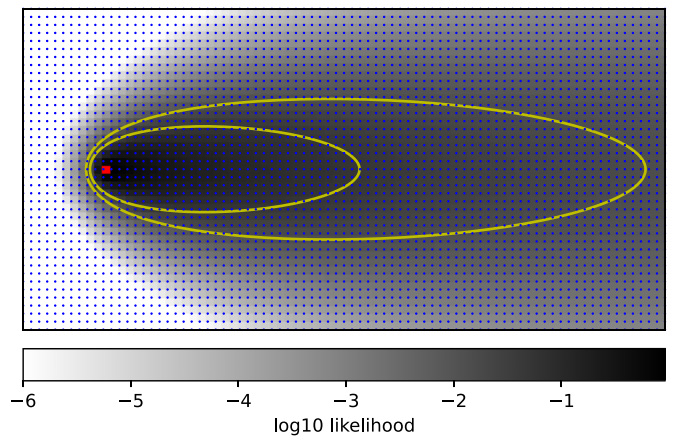


FIG. 2. Plot showing the grid world (blue dots) overlaid with the log detection likelihood for our choice of model parameters when $\bar{S} = 2.5$. The source, in red, has zero detection likelihood because it triggers a special observation. The agent always starts its search within two selected likelihood isocurves, shown in yellow.

time), conditioned on that belief. Explicitly, under a given policy π , we may define the *value* V_π of a belief as the total expected reward that can be accrued by following π :

$$V_\pi(b) = \mathbb{E} \left[\sum_{t=0}^{\infty} \gamma^t \sum_{s \in S} R(s, \pi(b_t)) b_t(s) \mid b_0 = b \right]. \quad (8)$$

We will define V^* as the value function under the *optimal* policy π^* . V^* can be shown to satisfy the *Bellman equation*

$$V^*(b) = \max_{a \in A} \left[\sum_{s \in S} b(s) R(s, a) + \gamma \sum_{o \in O} \Pr(o|b, a) V^*(b_{o,a}) \right], \quad (9)$$

where $\Pr(o|b, a) = \sum_{s \in S} \Pr(o|s, a) b(s)$. Once a solution to the Bellman equations is found, the optimal policy consists in a greedy selection of the action that maximizes the r.h.s. of (9). The argument of the maximum of the r.h.s. of (9) is simply the sum of the immediate expected reward for taking the action a and the discounted expected reward for all future actions. Many solution methods for POMDPs are based on “value iteration” on the Bellman equation, which is to say one computes

$$V^{n+1}(b) = \max_{a \in A} \left[\sum_{s \in S} b(s) R(s, a) + \gamma \sum_{o \in O} \Pr(o|b, a) V^n(b_{o,a}) \right] \quad (10)$$

until V^n converges. However, due to the large size of the belief simplex, it is challenging to obtain an approximation which is good on a sufficiently large subspace of the belief simplex, and convergence may be slow. This is the “curse of dimensionality” and the fundamental issue making POMDPs hard.

B. Initial belief

While Bayesian inference suffices to specify the evolution of the agent’s belief, we still need to set the initial belief b_0 that the agent holds when it starts searching (the prior). A naive choice, common for many POMDPs, would be to start from a uniform belief on the grid. However, we argue this is unphysical, as insects in nature generally do not start searching unless they have detected a cue. Moreover, when we have tested a uniform prior, we find that the resulting policies have the agent tending to explore the full extent of the box in order to locate the boundaries.

A second idea, then, is to bias the uniform prior by enforcing artificially a detection at time $t = 0$. This approach was taken in, for example, Ref. [25]. However, for our choice of parameters, detections are relatively rare except very close to the source. Thus, under this prescription, the agent will have the strong initial impression that it is within a few grid points away from the source, which is far from the ground truth.

Instead, we have elected to use a third approach, wherein the agent starts with a uniform prior and then waits in place for up to T_{wait} time steps, continually updating its belief, until it makes a detection. Only after this detection does the agent begin searching. Thus the agent’s initial belief is itself a random variable which carries some amount of useful information about where the source might be. This is intended to model the reasonable hypothesis that the insect knows the source is




FIG. 3. Plot showing a typical initial belief in an environment with emission rate $\bar{S} = 2.5$. The agent location is indicated with a red square and the source location is indicated with a red X, and positions are measured in units of the grid spacing.

unlikely to be very close when it receives its first detection signal. A typical initial belief is shown in Fig. 3.

To be clear, our choice of the initial belief is nothing more than a physical model, and it is not obvious that it should be preferred to the second approach (with a detection at time $t = 0$). We will briefly explore this alternate initialization in Appendix C.

C. POMDP solution

POMDPs are in principle exactly solvable by dynamic programming [37], but due to the curse of dimensionality such an approach is usually so computationally expensive as to be intractable. Instead, approximation methods are generally preferred.

In this work, we use the Perseus algorithm [26] to find near-optimal policies which approximately solve the Bellman equation. Perseus is a *point-based value iteration* algorithm [38,39] for POMDPs, which means it involves collecting an initial large sample of beliefs \mathcal{B} and then performing value iteration on those beliefs in order to obtain an approximation to the optimal policy. At each stage n in the value iteration [cf. Eq. (10)], V^n is approximated by a piecewise-linear and convex function, represented by a collection \mathcal{A}^n of hyperplanes in the belief simplex called α -vectors. We have

$$V^n(b) = \max_{\alpha \in \mathcal{A}^n} \alpha \cdot b, \quad (11)$$

where the dot product is over the states $s \in S$. Each α -vector has an action associated with it, such that the (near-)optimal action for each belief is that associated with the maximizing α .

The assembly of the belief set \mathcal{B} can in principle be performed using any policy. We use infotaxis (see Sec. III D) in this work.

We also accelerate the convergence of the algorithm using a *reward-shaping function* [27]. One can show (see Appendix A 3 for details) that adding a shaping function

$$F(s, a) = \phi(s) - \gamma \sum_{s'} p(s'|s, a) \phi(s') \quad (12)$$

to the reward will not change the optimal policy, for any state-dependent function $\phi(s)$. We will take

$$\phi(s) = -g[D(s)], \quad (13)$$

where g is some monotonically increasing function [with $g(0) = 0$] and D is the distance to the source, according to the metric induced by the state and action spaces—here the Manhattan distance. The point of this choice is to incentivize the agent to move closer to the source. We tested several such g in this work and found that, on problems of this size, a good choice of the shaping improves both the speed of convergence of Perseus and the performance of the resulting policies (see the Appendix B for more details). In particular, we are unable to achieve comparable performance in the absence of reward-shaping function. In contrast, on smaller problems with $O(100)$ points, we found that reward shaping was unnecessary and even counterproductive.

D. Heuristic strategies

As an alternative to trying to find a (near-)Bellman-optimal policy, one can instead propose a heuristic policy for a POMDP. Whereas solving the POMDP directly is a “black-box” approach that tries to directly maximize the reward over the entire horizon via (approximate) dynamic programming, a heuristic policy prescribes a simple, interpretable rule to choose an action, often by considering only a single time step in the future. Somewhat remarkably, there are a number of heuristics which can be effective for the search problem, despite its difficulty; here we present a few (see [40] for a review).

1. QMDP

Every POMDP has an underlying (fully observable) MDP, for which the optimal policy π^* is generally much easier to specify. One can then calculate the value of taking action a in state s , the so-called Q -function

$$Q(s, a) = \mathbb{E} \left[r(s, a) + \sum_{t=1}^{\infty} \gamma^t r(s_t, \pi^*(s_t)) \right]. \quad (14)$$

For our search problem, the MDP-optimal policy is just to take the path of minimal distance to the source, so we have $Q_{\text{MDP}}(s, a) = \gamma^{D(s')}$, where s' is the state resulting from taking action a in state s and $D(s)$ is again the gridwise distance to the source (Manhattan distance). This observation motivates the QMDP policy, which selects the action which maximizes the expectation of Q_{MDP} (conditioned on the belief):

$$\pi_{\text{QMDP}}(b) = \arg \max_{a \in A} \sum_{s \in S} Q_{\text{MDP}}(s, a) b(s). \quad (15)$$

Because it tries to directly minimize the expected time to reach the source, the QMDP policy tends to be exploitative. We will see that it is effective for this problem only when the emission rate is relatively high.

2. Infotaxis and space-aware infotaxis

The fundamental challenge of a POMDP is to make good action choices when faced with uncertainty. If the belief were perfectly informative (a δ distribution), then we would have a fully observable MDP, and the problem would be relatively trivial. This motivates an approach that tries to directly maximize the information content in the belief, or equivalently to

minimize the Shannon entropy. Let

$$H[b] = - \sum_{s \in S} b(s) \log b(s),$$

with the logarithm expressed in some units of choice. Then we can craft a policy which chooses the action maximizing the immediate expected information gain,

$$\pi_{\text{info}}(b) = \arg \min_{a \in A} \sum_{o \in O} \Pr(o|b, a) H[b_{o,a}]. \quad (16)$$

In the context of olfactory search, this policy is called *infotaxis* in analogy to chemotaxis [24].

Since one of the possible observations is to find the source, which would collapse the belief into a δ distribution, the infotactic policy naturally balances the immediate reward of finding the source with longer-term rewards associated with exploration. However, the probability of finding the source immediately is usually small, so the explorative component tends to dominate; indeed, we will see that infotaxis has an often excessive tendency towards safety. Thus, a number of variations and improvements have been proposed. One recent and promising variant, dubbed “space-aware infotaxis” (SAI), essentially combines infotaxis with the QMDP policy [25]. Explicitly, the policy is

$$\pi_{\text{SAI}}(b) = \arg \min_{a \in A} \sum_{o \in O} \Pr(o|b, a) \log_2 \left(\sum_{s \in S} D(s) b_{o,a}(s) + 2^{H_2[b_{o,a}]} - 1 + \frac{1}{2} \right). \quad (17)$$

We have chosen the base-2 logarithm and measured the entropy in bits to be consistent with the original authors, but our implementation differs very slightly in that we have reversed the sign in front of $1/2$, which ensures that the contribution to the outer sum from finding the source is nonsingular.

The second term of the summand is a crude estimation of the expected time to learn the location of the source (i.e., by checking 2^{H_2} cells), so the SAI policy is an attempt to directly minimize the total time to find the source. More generally, it balances infotaxis’s tendency towards exploration with QMDP’s tendency towards exploitation, and we will see it performs quite well.

3. Thompson sampling

A classical heuristic for decision problems with partial information is, at each time step, to estimate the true state by *sampling* from the current belief (posterior) and then to choose the action with maximal expected reward according to that sample. This strategy is usually referred to as *Thompson sampling* [41–43].

We adapt Thompson sampling to the search problem in the following way. At each time step, we sample a possible source location from the current belief, and then take the action which brings the agent closest to that location (if there is more than one such action, we choose from these at random).

We also find it useful to generalize this policy by introducing a persistence time $\tau \geq 1$: rather than sample a new location at every time step, the agent follows the sampled location for τ time steps (or until it reaches the sampled




FIG. 4. Comparison of performance of Perseus (using the same policies as those chosen in Sec. IV A) vis-à-vis heuristic policies on the three environments tested. Problems P1–P4 represent starting from the single points $(25, -4)$, $(35, -4)$, $(45, -4)$, and $(55, -4)$ in order of increasing distance from the source, and problem E is the ensemble of starting points described in Sec. A 6. The performance on each problem is measured by the mean excess arrival time $\langle \tilde{T} \rangle \equiv \langle T \rangle - \langle T_{\text{MDP}} \rangle$, normalized by the minimum time $\langle T_{\text{MDP}} \rangle$. Error bars are suppressed for visual clarity; uncertainties on the means were at most 3.3%, as measured by the standard error. In the lower right corner, we show the test problems: the ensemble E (for $\bar{S} = 2.5$) as green circles, and P1–P4 as blue squares.

location), and only then does it resample. The benefit of pursuing a sample for an extended time rather than resampling at every time step is linked to the need for “deep exploration” when navigating problems with sparse reward structure [43,44].

That Thompson sampling depends on moving towards random locations suggests it should have a tendency towards exploration. Indeed, we will see it is generally a safe policy that, for sufficiently large τ , performs especially well in environments with low emission rate and thus low information.

IV. RESULTS

Once the Perseus algorithm has reached convergence, we freeze the policy and test it against the heuristic baselines described in Sec. III D. We present results for three different model environments, which have emission rates $\bar{S} = 0.25, 2.5$, and 25 but are otherwise identical. We suggest that the most direct ways to change the character of the problem are to either alter the emission rate or move towards the windless, isotropic limit $V\lambda/D \rightarrow 0$, which we will not study here as it was examined thoroughly in Ref. [25] and will be further studied in Ref. [45]. The windy case is also arguably more relevant to insect behavior. The agent’s model for the environment, by which it performs Bayesian updates, is exact, and detection events are drawn at random from the distribution defined by Eqs. (6) and (7). Results are shown for several choices of the shaping function $g[D(s)]$.

The discount factor is a very important hyperparameter, and we find values in the range $0.95 \geq \gamma \geq 0.99$ work best. In the environment with $\bar{S} = 0.25$, we set $\gamma = 0.96$, and in the

other two cases we set $\gamma = 0.98$. Further optimization may be possible but is not the goal of this study. For additional details on γ dependence, as well as other convergence properties of Perseus, refer to the Appendix.

The main results are summarized in Fig. 4, where the performance of all policies are tested on both individual starting points and an ensemble of starting points. The individual points are Problems P1–P4 and represent starting from the single points $(25, -4)$, $(35, -4)$, $(45, -4)$, and $(45, -4)$, where the coordinates are relative to the source and given in units of the grid spacing. Performance is measured by the excess arrival time $\langle \tilde{T} \rangle \equiv \langle T \rangle - \langle T_{\text{MDP}} \rangle$, which we then normalize by the minimum time $\langle T_{\text{MDP}} \rangle$. For the single-point problems, we exclude “failed” trials with $T \geq 10000$ from the mean (in no case, except for QMDP when $\bar{S} = 0.25$ or $\bar{S} = 25$, was the failure rate significant), whereas, for the ensembles, we include the case $T \geq 1000$ (which all are registered as 1000 due to the time limit imposed). Anecdotally, failures are usually due to the agent somehow becoming trapped and seem to be correlated with extremely rare events, i.e., when the agent makes a detection in a region where the likelihood is very low.

We see that Perseus clearly outperforms all tested heuristics, and that the best-performing heuristic depends strongly on the emission rate, making it difficult to select *a priori* an optimal baseline. We can also conclude that space-aware infotaxis is a strictly better policy than its “vanilla” counterpart, that QMDP is inferior except when the emission rate is large, and that Thompson sampling performs best with a persistence time $\tau > 1$ and when the emission rate is small. We briefly remark that in certain cases, especially when $\bar{S} = 25$, the

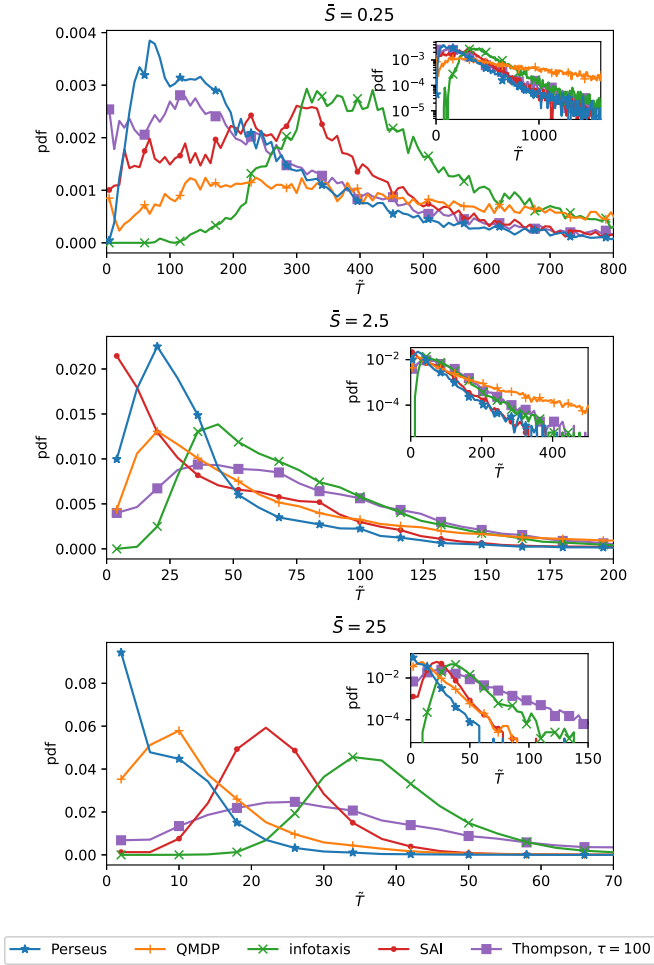


FIG. 5. Comparison of excess arrival time pdfs, in each environment, from the point $(45, -4)$ for a selected Perseus policy and several heuristics. The insets show the same pdfs on a log-linear scale to emphasize the tails.

normalized mean arrival time decreases with starting distance from the source. This is not a contradiction and simply means that the excess arrival time is increasing sublinearly with distance.

A. Excess arrival time pdfs from a single starting point

Perseus was run on each environment using several choices of the reward-shaping function. We then chose, for each emission rate, one Perseus policy and performed 20 000 Monte Carlo trials from a single starting point, P3, $(45, -4)$. We did the same for the heuristic policies. Here the agent was allowed up to 10 000 time steps to reach the source, to better resolve the tails of the distributions.

We selected for testing Perseus policies that (1) outperformed heuristics on the ensemble averaging, (2) outperformed heuristics on the four individual points, and (3) were evolved for as many iterations as possible.

The resulting pdfs are shown in Fig. 5. Note that their “wiggly” appearance is due to finite sample size effects. These pdfs help illustrate the strengths and weaknesses of the heuristic policies. We see that infotaxis clearly tends to be too

safe, with a tail that decays quickly for large arrival time but with a relatively small probability of rapidly arriving to the source. Infotaxis performs best at intermediate emission rate but is clearly outperformed in all cases by SAI. SAI, too, is especially good at intermediate emission rate and is closely competitive with Perseus. In fact, for intermediate emission rate, SAI’s arrival time pdf peaks at the minimum arrival time; the Perseus policy achieves a lower mean arrival time only due to having substantially less probability of a long arrival time ($\tilde{T} \gtrsim 60$). QMDP is again seen to be inferior except when the emission rate is large. The excessive greediness of this policy can be seen by the very heavy, slowly decaying tails when $\bar{S} = 0.25$ or 2.5. Finally, Thompson sampling shines at low emission rate, achieving results competitive with Perseus, but is otherwise relatively mediocre.

B. Sample trajectories and policy similarity for $\bar{S} = 2.5$

In Fig. 6 we show some sample trajectories in the intermediate-emission-rate environment ($\bar{S} = 2.5$). In Fig. 7 we compare policies to Perseus in this environment by estimating the pdfs for angular differences in chosen actions. To be precise, we let the agent search using the Perseus policy, and at each time step we also compute the actions which would have been chosen by the heuristic policies, given the same belief. These actions are converted into a polar angle $\theta \in \{0, \pi/2, \pi, 3\pi/2\}$, and we record the angular differences (modulo 2π) $\Delta\theta = \theta_i - \theta_{\text{Perseus}}$ between the Perseus action and the heuristic actions for each policy i . The starting points were selected randomly in the usual way, and 1000 Monte Carlo trials were performed. Unsurprisingly, of the four heuristics tested, SAI was most similar to the near-optimal policy, with more than 50% of the actions being identical.

The trajectory plots are of course qualitative, and a careful quantitative analysis of the behaviors which typify different policies is beyond the scope of the present work. However, it is worth commenting briefly on some of the observed features. In general, each policy tends to result in a diverse set of behaviors, and policies are not always easily distinguishable by the naked eye. Nevertheless, a few traits are apparent. The Perseus agents tend to prefer to move upwind initially before beginning cross-wind motion. If some time passes without a detection, there is often a tendency to return in the downwind direction, a behavior which helps prevent the agent from overshooting the source. Infotaxis agents often move in broadly arcing trajectories, including outward-moving spirals. Spiraling behavior under infotaxis is well known, having been observed in the original paper [24] as well as subsequent work [25,34,46], and it can be connected to the theory of search games: one can show that the optimal minimax trajectory for searching for a single point in a plane is an exponential spiral [47]. SAI trajectories are often similar to the Perseus trajectories, sharing the tendency to initially move upwind (a consequence of the distance term in the objective function). QMDP tends to be a greedier policy; QMDP agents have a tendency to either strike downwind or search semiexhaustively on small scales. Thompson sampling agents exhibit meandering trajectories characteristic of the randomness of the underlying policy; this randomness is also consistent with

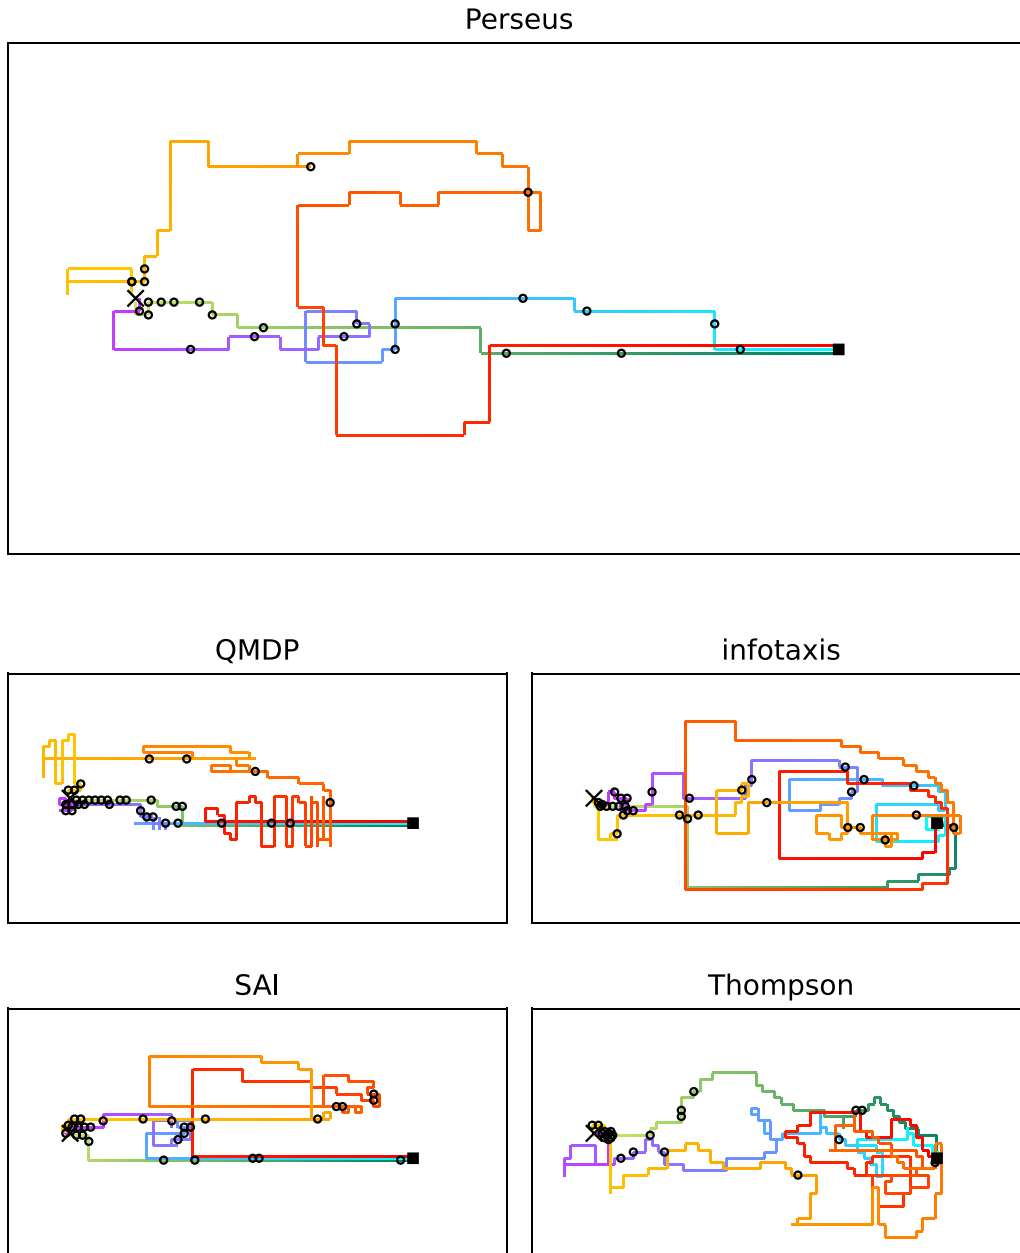


FIG. 6. Sample trajectories from the same starting point using various policies in the $\bar{S} = 2.5$ environment. The source is indicated with a black X, and the starting point with a red square. Detection events are indicated with black circles. We have deliberately chosen a trajectory with an especially small arrival time (green), a trajectory close to the median (light blue-purple), and a trajectory with an especially long arrival time (red-orange). The color gradients indicate the passage of time so that there is no ambiguity when the trajectories self-intersect. The trajectories have been offset in space slightly for visual clarity. We show trajectories for the near-optimal Perseus policies, QMDP (with $\gamma = 0.98$), infotaxis, SAI, and Thompson sampling (with $\tau = 10$).

the relatively small observed probability that a Thompson agent will chose the same action as a near-optimal Perseus agent. As a final remark, we found that all policies tested in the present work sometimes yielded behaviors which resemble casting.

C. Dependence of problem difficulty on starting point

How does the mean arrival time of a (near-)optimal policy depend on the starting position? Trivially, the arrival time is

bounded below by the MDP optimum, which measures the component of the problem difficulty due to the time needed to reach the source. The mean excess arrival time $\langle \tilde{T} \rangle$ then measures, in a sense, the component of the difficulty due to partial observability, or the time required to gather information and determine where the source is. It is not immediately obvious how $\langle \tilde{T} \rangle$ should depend on the starting position.

To help answer this question, in Fig. 8 we plot the mean excess arrival time using Perseus, when $\bar{S} = 2.5$, as a function of the downwind distance at fixed cross-wind distance, and




FIG. 7. Pdfs for angular differences between Perseus policy and other policies, when $\bar{s} = 2.5$. Here the Thompson sampling policy used $\tau = 10$ and the QMDP policy used $\gamma = 0.98$.

vice versa. The averages were taken over 10^4 Monte Carlo trials, and arrival times greater than or equal to 5000 were suppressed from the calculation. Somewhat surprisingly, the problem difficulty does not appear to depend strongly on the starting cross-wind distance, as long as the agent starts within a few cross-wind diffusion lengths from the symmetry axis. Instead, the excess arrival time mostly depends on the




FIG. 8. Dependence of the mean excess arrival time on the initial Manhattan distance from the source, using the near-optimal Perseus policy, for $\bar{s} = 2.5$. We compare curves at fixed downwind distances $x - x_0$ to the curve at fixed $y = y_0$ (i.e., where the agent starts on the symmetry axis). The inset shows typical entropy values H_{typ} for the fixed downwind distance curves.

downwind distance, scaling approximately linearly therewith. From Fig. 8 it is clear that starting farther downwind from the source generally makes the problem more difficult than starting farther off-axis.

This behavior can be qualitatively explained by the following argument. The problem of finding the source can be solved by locating the symmetry axis (say, by casting) and then proceeding upwind. Assuming that τ is large, a detection at (x, y) (measured with respect to the source) means that $y^2 < 2\lambda x$ with high probability, since outside of this parabola, detections are exponentially suppressed with decay length scale λ . The farther away from the source the agent starts, the wider this parabola is, and the wider the agent must cast in order to make a detection (which of course consumes more time). On the other hand, the time spent locating the symmetry axis in a casting-based strategy should *not* depend on its initial cross-wind distance from the axis.

We also note a precipitous drop in the excess mean arrival times with cross-wind distance starting around $y = 10$. This can be explained by a drop in the entropy of a typical starting belief when one starts sufficiently far from the symmetry axis. If the agent starts at \mathbf{r}_i , the time to the first hit, call it k , obeys a geometric distribution with parameter $\ell_0 = \Pr(o = \text{hit} | \mathbf{r}_i - \mathbf{r}_0)$, and we have $\mathbb{E}[k] = 1/\ell_0$. If $\ell(s)$ is the probability of detection as a function of states, the belief will then accrue, through Bayesian updates, $k - 1$ factors of $1 - \ell(s)$ (one for each nondetection) and one factor of $\ell(s)$. Thus a typical initial belief is

$$b(s) = \frac{[1 - \ell(s)]^{1/\ell_0 - 1} \ell(s)}{\sum_s [1 - \ell(s)]^{1/\ell_0 - 1} \ell(s)}. \quad (18)$$

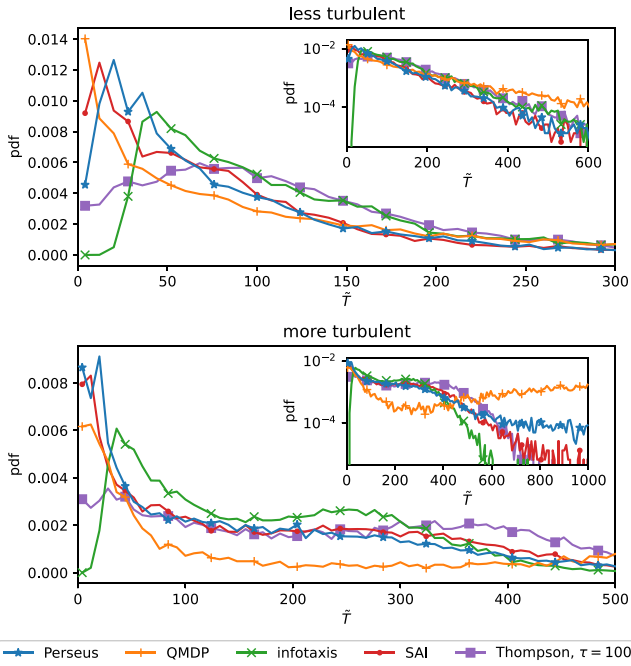


FIG. 9. Arrival time pdfs from $(45, -4)$ for the previously obtained near-optimal policy for $\bar{S} = 2.5$ and for heuristics, but now searching in environments which are less (top) and more (bottom) turbulent than the agent believes.

The entropies of the typical beliefs corresponding to the starting positions of the constant downwind distance curves of Fig. 8 are shown in the inset of that figure. The reason for the entropy drop is restriction of the support of the initial belief to a progressively smaller area: as k gets larger, the initial belief is confined to the maximum likelihood curve $\ell(s) = 1/k$ with increasingly narrow characteristic width $w \propto 1/\sqrt{k-1}$.

D. Robustness of policies to changes in environment

Because POMDP is a model-based approach, we have assumed that the insect has some instinctual knowledge of the turbulent environment. Up until now, this knowledge has been an exact model of the detection statistics. In this section, we relax this strong assumption and experiment with a scenario where the agent has an *imperfect* model of the environment. In particular, the true physical parameters will be different from those the agent uses to update its belief and those which were used to construct a near-optimal policy. We present results for two cases: one, a more turbulent environment where $D \rightarrow 2D$ and $V \rightarrow V/2$, and two, a less turbulent environment where $D \rightarrow D/2$ and $V \rightarrow 2V$, relative to the parameters we have used previously. The agent will use the old parameters to update its belief. In both cases, we set $\bar{S} = 2.5$.

In Fig. 9 we show excess arrival time pdfs obtained using the same methods as in Sec. IV A, but now for these two scenarios. The mean excess arrival times are shown in Table II, with the previous results where the model is exact shown for comparison. In Table III we show failure rates for the problems.

In the more turbulent environment, the Perseus policy performs poorly compared to most other heuristics, in particular

TABLE II. Mean excess arrival times from $(45, -4)$ (with standard error shown) for $\bar{S} = 2.5$ on three problems: the original problem where the agent’s model for the environment is *exact* (E), a scenario where the true detection statistics are reflective of a *more turbulent* environment (MT) than the agent believes, and a scenario where the true detection statistics are reflective of a *less turbulent* environment (LT) than the agent believes.

Policy	$\langle \tilde{T} \rangle$ (E)	$\langle \tilde{T} \rangle$ (MT)	$\langle \tilde{T} \rangle$ (LT)
Perseus	39.1 ± 0.3	349.1 ± 6.2	91.2 ± 1.6
QMDP	97.9 ± 1.4	1852.1 ± 11.1	231.4 ± 4.4
Infotaxis	75.5 ± 0.3	174.5 ± 0.9	120.1 ± 5.9
SAI	43.8 ± 0.3	179.4 ± 1.2	79.6 ± 0.6
Thompson ($\tau = 10$)	77.0 ± 0.3	262.1 ± 1.3	105.2 ± 0.5

suffering a high rate of failure and a fat tail, leading to a large mean arrival time. Infotaxis has the best performance here, due to its rapidly decaying tail, which is consistent with its being a “safe” policy. Moreover, whereas the failure rates were negligibly small previously, Perseus and QMDP now have substantial probabilities of failing (i.e., taking $\geq 10^4$ time steps to reach the source).

The problem of searching in an environment that is less turbulent than believed is evidently substantially easier. In the less turbulent environment, the Perseus policy performs very adequately, scoring the second best performance behind SAI. Perseus also has a substantially smaller failure rate on this problem than in the more turbulent environment. This underscores an obvious drawback of training a policy to be optimal for a given environment: if the environment changes, the policy may be too overtuned to perform adequately. However, these results suggest that if the environmental parameters are unknown, for the sake of robustness it may be preferable to train a policy to be optimal assuming a more turbulent environment.

V. DISCUSSION

We have studied a search problem relevant to the behavior of a number of flying insects. We have shown that solving for near-optimal search policies on a problem space with several thousand points is feasible using an existing POMDP algorithm, Perseus, when accelerated with a good choice of reward shaping. This approach yielded policies which outperformed all tested heuristics in terms of the mean arrival time, over a wide range of emission rate regimes. We are thus optimistic that more sophisticated search problems, such as

TABLE III. Same as Table II, but showing failure rates for the problems.

Policy	Failure rate (E)	Failure rate (MT)	Failure rate (LT)
Perseus	5×10^{-5}	0.0264	0.0017
QMDP	5×10^{-5}	0.00935	0.0096
Infotaxis	0	10^{-4}	0
SAI	0	0	0
Thompson ($\tau = 10$)	0	0	0

ones which take spatiotemporal correlations between detections into account, could be amenable to a direct POMDP solution approach, despite necessitating a larger POMDP state space. Future work will investigate such problems.

We also studied which heuristic strategies perform best in environments with different characteristic concentration levels (emission rates). In particular, we found a randomized search algorithm—Thompson sampling—to be well suited for very dilute environments, space-aware infotaxis to be excellent at a somewhat higher concentration, and QMDP to be effective only on the easiest problems with substantial detection rates. It should be noted, however, that these conclusions may be sensitive to the choice of prior.

An advantage of certain heuristics, especially Thompson sampling and variants of infotaxis, over a near-optimal policy is their being more flexible when the agent’s model of the environment is imperfect. Finding policies which are effective in a variety of environments is an interesting avenue of future research; in that case, model-free approaches may be preferable to POMDP.

We found that a variety of behaviors emerge from different policies, but classifying these carefully is highly nontrivial since behaviors depend on the observation history and reflect correlations between actions over relatively long timescales. For instance, we tried measuring the fraction of time each policy spent moving cross-wind, upwind, or downwind as a simple metric but did not find it informative. Thus, we defer a serious quantitative study of behaviors to future efforts.

It should be noted that while we found increases to the emission rate led to a reduction in typical arrival times, this trend cannot continue indefinitely. If the emission rate is sufficiently large, then the likelihood of detection will begin to saturate near unity and a binary detection scheme will cease to be informative or useful. Instead, in this regime one would expect gradient-based (chemotactic) strategies to once again be effective.

Finally, we used a near-optimal policy to study the spatial dependence of the mean excess arrival time, a proxy for the intrinsic problem difficulty as a function of starting point. This dependence is strongly anisotropic: the mean excess arrival time increases monotonically as the agent starts farther downwind, but has a strongly nonmonotonic dependence on cross-wind distance. Moving only a few λ off-axis has virtually no effect on the problem difficulty and may even make it slightly easier, which may be related to why cast-and-surge is an effective search strategy.

The approach proposed in this work can be extended in several directions. First, and most importantly, it would be extremely interesting to validate it on realistic data from direct numerical simulations of emission from a point source in 2D or 3D turbulent flows, where a Markovian model for observations will necessarily be incomplete. A study of this kind is currently ongoing. Second, similar techniques can be used to attack multisource problems and/or multiagent problems. It is increasingly urgent to identify clear setup with high quality and quantity of data for training and validations of data-driven algorithms, and the one here studied is certainly a good paradigmatic candidate.

In this work, we kept the discount factor γ as close to one as possible and aimed to minimize the mean arrival time

T . Strictly speaking, however, we were actually maximizing γ^T ; in general the closer γ is to 1, the more we care about optimizing the tail of the arrival time pdf, which is to say avoiding very long arrival times. We have also noticed that in the presence of a nonzero failure rate, it is not always obvious which policy is “best”—depending on one’s tolerance of failure, a lower mean arrival time with higher failure rate may be preferable to a “safer” policy which almost always finds the source. These ideas can be generalized by the notion of *risk-sensitive* problems [48], where the objective function is transformed in a way that reflects the agent’s aversion or attraction to risky behavior. Obtaining optimal policies for the search problem subject to risk sensitivity is another subject for future study.

A forthcoming study [45] will compare the performance of the deep-RL method proposed in [25] to the approach of the present work on problems of this size.

ACKNOWLEDGMENTS

R.A.H. gratefully acknowledges useful discussions with Aurore Loisy, Fabio Bonaccorso, Michele Buzzicotti, and Antonio Costa. R.A.H. and L.B. received funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation program (Grant No. 882340). A.C. acknowledges funding from the European Union’s Horizon 2020 research and innovation program under the Marie Skłodowska-Curie Grant No. 956457.

APPENDIX A: DETAILED METHODS

1. POMDP implementation

Defining the POMDP in a careful way is important to keep the problem tractable and to avoid issues with boundary artifacts. Hence, perhaps at the risk of seeming pedantic, we will try to be precise as possible in what follows.

Let G be the $N_x \times N_y$ grid world. The POMDP state space may be defined to be the Cartesian product of the possible agent locations \mathbf{r} and the possible source locations \mathbf{r}_0 , $G \times G$, and we have a belief b on the $G \times G$ -simplex which is the joint distribution of the agent location belief and the source location belief. This belief simplex is intractably large, $(|G|^2 - 1)$ -dimensional. However, we can exploit to our advantage (1) the sparsity of the agent location belief, which is a δ distribution since the agent knows where it is, and (2) the fact that the observation likelihoods depend only on the displacement $\mathbf{s} = \mathbf{r} - \mathbf{r}_0$ between the agent and source, and not on \mathbf{r} or \mathbf{r}_0 independently.

We map the state (living on $G \times G$) to a smaller state space, the $(2N_x - 1) \times (2N_y - 1)$ -dimensional grid spanning $-L_x \leq x \leq L_x$ and $-L_y \leq y \leq L_y$, with spacing Δx and Δy , which we call G' . G' is the space of all possible displacements, given that the agent and source are both on G . We thus form a new belief b' of displacements living on the G' -simplex. b' is found by embedding the belief of source locations (the nonzero slice of the full belief) using the rule

$$b'(\mathbf{s}) = b(\mathbf{r} - \mathbf{s}), \quad \mathbf{r} - \mathbf{s} \in G \quad (\text{A1})$$

$$0, \quad \text{otherwise}$$

for all $\mathbf{s} \in G'$, where, in a slight abuse of notation, b is the belief of the source location. This embedding can be inverted to recover b from a given b' using the known agent location \mathbf{r} .

Let us summarize the procedure. As the agent moves, it maintains a belief of the location of the source location b . b is updated after taking an action \mathbf{a} and making an observation o according to Bayes' rule

$$b(\mathbf{r}_0|o, \mathbf{a}) = \frac{b(\mathbf{r}_0)\Pr(o|\mathbf{r}', \mathbf{r}_0)}{\sum_{\mathbf{r}_0} b(\mathbf{r}_0)\Pr(o|\mathbf{r}', \mathbf{r}_0)}, \quad (\text{A2})$$

where \mathbf{r}' is the position of the searcher after previously being at \mathbf{r} and taking action \mathbf{a} . Explicitly,

$$\begin{aligned} \mathbf{r}' &= \mathbf{r}, & \mathbf{r} + \mathbf{a} \notin G \text{ or } \mathbf{r} = \mathbf{r}_0 \\ \mathbf{r}' &= \mathbf{r} + \mathbf{a}, & \text{otherwise.} \end{aligned} \quad (\text{A3})$$

Finally, the belief is embedded into the G' -simplex to form b' . The computed, near-optimal policies in this work use b' as input. For the purposes of computing policies, the boundaries of G' (which are generally not encountered by the agent) are chosen to be doubly periodic for simplicity.

The primary reason for choosing this somewhat complicated representation is that it avoids boundary artifacts while maintaining a relatively compact dimensionality (many fewer dimensions than the full belief on $G \times G$). Naive implementations which depend on propagating the belief of the relative position of the agent can introduce such artifacts when the agent moves, causing probability mass to exit the domain and be lost. In our representation, we propagate the belief of where the *source* is; this propagation does not introduce artifacts because the source is static, and therefore the transition matrix for the source location is trivial: $\Pr(\mathbf{r}'_0|\mathbf{r}_0, \mathbf{a}) = \delta_{\mathbf{r}_0, \mathbf{r}'_0}$ for any $\mathbf{a} \in A$.

2. Alternate reward structures

A more general way to model the search problem as a POMDP would give a (discounted) penalty at each time step until the agent finds the source, and then supply some onetime reward $R \geq 0$ (in the case $R = 0$ then we are directly minimizing the arrival time). But as long as $\gamma < 1$, this is equivalent to the reward we used in the present work. If we let the arrival time be T , the reward is then

$$\begin{aligned} \mathbb{E} \left[\sum_{t=0}^{T-1} (-1)\gamma^t + R\gamma^T \right] &= \mathbb{E} \left[-\frac{1-\gamma^T}{1-\gamma} + R\gamma^T \right] \\ &= \frac{1}{1-\gamma} [(1+R)\mathbb{E}[\gamma^T] - 1]. \end{aligned} \quad (\text{A4})$$

Thus maximizing the expected reward for this reward structure is equivalent to maximizing the expected reward for the structure with no penalty per unit time.

It is an important technical point that, while our ultimate goal is to reach the source in a minimal time, we are not directly minimizing the mean arrival time $\langle T \rangle$, but rather maximizing a proxy $\langle \gamma^T \rangle$. Note that in the limit $\gamma \rightarrow 1$, the two objectives are equivalent (provided that the typical T does not diverge in this limit), as can be shown by simple Taylor expansion of γ^T .

We are thus effectively treating γ as a hyperparameter of the algorithm, and we will simply select the one that yields the

best empirical performance, as measured by the mean arrival time, on a given problem setup. In contrast, the more typical viewpoint would be to consider γ to be a parameter which reflects the environment and/or the agent and its priorities, each value of which defines a separate problem with a different optimal policy.

The basic reason for our choice to use $\gamma < 1$ is that most available POMDP algorithms, Perseus included, require use of a discount factor in order to converge. A key underlying fact is that when $\gamma < 1$, the r.h.s. of the Bellman equation acts as a contraction operator on the value function, which guarantees the convergence of iterative solution techniques to a unique fixed point [49].

The undiscounted case $\gamma = 1$ results in direct minimization of the mean arrival time $\mathbb{E}[T]$ and is thus an interesting alternative problem. However, there are few algorithms which can deal with the absence of the discount factor, the DRL approach of Ref. [25] being a rare example. We have directly verified [45] that introducing γ does not substantially increase the mean arrival time, relative to that of the solution to the undiscounted problem.

3. Reward shaping

In many (fully observable) MDPs, it is often possible to speed up convergence to an optimal policy by adding a well-chosen *shaping function* to the reward. In particular, there exist *potential shaping functions* such that the MDP under the transformed reward has precisely the same optimal policy as the original MDP [27]. This idea can be generalized to POMDPs in a straightforward manner.

We start by introducing the function $Q(b, a)$ which expresses the value (expected total reward) of taking action a when the agent has belief b . In particular, we have

$$V^*(b) = \max_{a \in A} Q(b, a) \quad (\text{A5})$$

and

$$\pi^*(b) = \arg \max_a Q(b, a). \quad (\text{A6})$$

Importantly, the optimal policy is unchanged under any transformation $Q \rightarrow Q + \varphi(b)$. Intuitively, this means the component of the value of a state which is intrinsic to the state, independent of the choice of action, should not affect the policy. Q satisfies its own Bellman equation

$$Q(b, a) = \sum_{s \in S} b(s)R(s, a) + \gamma \sum_{o \in O} \Pr(o|b, a) \max_{a' \in A} Q(b_{o,a}, a'). \quad (\text{A7})$$

Letting $\hat{Q}(b, a) = Q(b, a) + \varphi(b)$ and substituting, we have

$$\begin{aligned} \hat{Q}(b, a) &= \sum_s b(s)R(s, a) + F(b, a) \\ &\quad + \gamma \sum_o \Pr(o|b, a) \max_{a'} \hat{Q}(b_{o,a}, a'), \end{aligned} \quad (\text{A8})$$

with

$$F(b, a) = -\varphi(b) + \gamma \sum_o \Pr(o|b, a)\varphi(b_{o,a}), \quad (\text{A9})$$

which is a new Bellman equation for \hat{Q} .

As an important special case, we can restrict φ to a linear functional $\varphi(b) = \sum_s b(s)\phi(s)$. (This special case is especially useful for us because we seek a piecewise linear approximation for V .) Then the introduction of the potential is equivalent to modifying the reward $R(s, a) \rightarrow R(s, a) + F(s, a)$, where

$$F(s, a) = \phi(s) - \gamma \sum_{s'} p(s'|s, a)\phi(s'). \quad (\text{A10})$$

Thus, adding any function of this form to the reward in a POMDP will not change the optimal policy. This flexibility in defining the reward is akin to a kind of gauge invariance. Note that if \hat{V}^* is the value function under the shaped reward, we also have the simple identity

$$\hat{V}^*(b) = V^*(b) + \varphi(b). \quad (\text{A11})$$

How do we choose a good shaping function? One can argue that the “best” potential, which would accelerate value iteration as much as possible, would in fact be, up to an additive constant, the (negative) optimal value function, which is suggested by Eq. (A11). To see this, consider value iteration in the presence of a generic shaping $\varphi(b)$:

$$\begin{aligned} V^{n+1}(b) = \max_{a \in A} & \left[\sum_{s \in S} b(s)R(s, a) - \varphi(b) \right. \\ & + \gamma \sum_{o \in O} \Pr(o|b, a)\varphi(b_{o,a}) \\ & \left. + \gamma \sum_o \Pr(o|b, a)V^n(b_{o,a}) \right]. \end{aligned} \quad (\text{A12})$$

Suppose $V^0(b) = 0$ for some b (if $\gamma < 1$, we can always define the reward in such a way that 0 is a safe choice for the initialization of the value function), and suppose we choose $\varphi(b) = -V^*(b)$. Then value iteration would yield

$$\begin{aligned} V^1(b) = \max_{a \in A} & \left[\sum_{s \in S} b(s)r(s, a) - V^*(b) \right. \\ & \left. + \gamma \sum_{o \in O} \Pr(o|b, a)V^*(b_{o,a}) \right] = 0, \end{aligned} \quad (\text{A13})$$

where we have used the Bellman equation. The value at b has already converged, and in particular the maximizing action is none other than $\pi^*(b)$. Thus, a single-value iteration would instantly give the correct action in a neighborhood of b , which would reduce the problem to simply sampling enough b .

Of course, we do not have access to the optimal value function (this would defeat the purpose of value iteration!), so we must make an inspired guess for the reward-shaping function, which is structurally similar to the true value. For the search problem, a natural choice is to shape the reward to encourage moving toward the source by setting, as we have in this work,

$$\varphi(b) = - \sum_s b(s)g[D(s)], \quad (\text{A14})$$

where g is monotonically increasing and $g(0) = 0$, and D is the Manhattan distance to the source.

The importance of ensuring the shaping function is potential cannot be overstated. A nonpotential choice like rewarding the agent for making detections, while having intuitive merit, would explicitly destroy the optimality of the solution with respect to the arrival time.

4. Perseus algorithm

Value iteration in Perseus is accomplished through the *backup* operation. With the aid of Eq. (11), we can express the value iteration as

$$V^{n+1}(b) = \max_{a \in A} \left[b \cdot R_a + \gamma \sum_{o \in O} \max_i b \cdot g_{a,o}^i \right], \quad (\text{A15})$$

where $R_a = R(s, a)$ and

$$g_{a,o}^i = \sum_{s' \in S} \Pr(o|s', a)\Pr(s'|s, a)\alpha_i(s'). \quad (\text{A16})$$

There is a $g_{a,o}^i$ for each α -vector, which can be computed once and stored. The backup operation is then

$$\text{backup}(b) = \arg \max_{\{\alpha'_a\}_{a \in A}} b \cdot \alpha'_a, \quad (\text{A17})$$

where

$$\alpha'_a = R_a + \gamma \sum_{o \in O} \arg \max_{\{g_{b,a}^i\}_i} b \cdot g_{b,a}^i. \quad (\text{A18})$$

The backup thus produces a new α -vector α' so that $V^{n+1}(b) = b \cdot \alpha'$, and the optimizing action is found during the argmax in Eq. (A17).

The backup operator forms the basis of an array of different “point-based” algorithms [39], which differ in how they sample beliefs from the simplex and in what order they are backed up. In the original Perseus algorithm, backups are performed in a random order, but we perform them in the order of decreasing *Bellman error*—the so-called “prioritized” version of Perseus. The Bellman error of a belief b is defined as

$$\epsilon(b) = \max_{a \in A} \left[b \cdot R_a + \sum_{o \in O} \Pr(o|b, a)V^n(b_{o,a}) \right] - V^n(b). \quad (\text{A19})$$

Prioritizing Perseus in this way was shown to accelerate convergence in [50].

Our implementation of the Perseus algorithm proceeds as follows. First, we assemble a large collection of beliefs \mathcal{B} , by exploring the environment according to some policy (after initializing the agent in the same way that we do during evaluation), updating the belief using Bayes’ theorem, and adding the belief to \mathcal{B} . When the agent finds the source, we restart and repeat, until we have enough beliefs. The original paper [26] suggested using a uniform random policy, but we find it is much better to employ a heuristic, as suggested in [39]. Intuitively, the most useful beliefs to sample are those which are likely to be encountered when taking optimal actions [51], and it is generally understood that the subspace of these “reachable beliefs” is much smaller than the whole simplex. Thus using a good heuristic which is reasonably close to optimal is a far more efficient way to sample beliefs; we use infotaxis. We found that $|\mathcal{B}| = O(10^4)$ beliefs were required to obtain good

results, with clear loss in performance when fewer beliefs were sampled.

Next, we initialize the policy to a single α -vector: $\mathcal{A}^0 = \{\mathbf{0}\}$, the zero-vector. This was chosen to guarantee $V^0(b) \leq V^*(b) \forall b$ [note that reward shaping does not affect this choice, as long the potential $\varphi(b)$ is non-negative, in view of Eq. (A11)]. We assign one of the four actions to this vector; it does not matter which.

Finally, we perform some number of iterations until a convergence or performance criterion is met. An iteration consists of the following steps:

- (1) Compute, for each $b \in \mathcal{B}$, $\epsilon(b)$.
- (2) Initialize \mathcal{B}' to \mathcal{B} and \mathcal{A}^{n+1} to \emptyset .
- (3) Find $\alpha = \text{backup}(b)$ for the $b \in \mathcal{B}'$ with largest ϵ .
- (4) If $\alpha \cdot b \geq V^n(b)$, then add α to \mathcal{A}^{n+1} . Otherwise, add the $\alpha \in \mathcal{A}^n$ which previously gave the maximum value for b .
- (5) Set $\mathcal{B}' \leftarrow \{b \in \mathcal{B}' : \alpha \cdot b < V^n(b)\}$, where α is the vector added to \mathcal{A}^{n+1} in step 4.
- (6) If $\mathcal{B}' \neq \emptyset$, go to step 3.

Perseus is not the only algorithm for POMDP planning. We note in particular the existence of ‘‘heuristic search value iteration’’ (HSVI) [52,53] and SARSOP [51], both of which involve building a tree of beliefs reachable from some (single) initial belief, and both of which come with more rigorous guarantees of convergence than Perseus. While HSVI and SARSOP have been shown to outperform Perseus on certain problems, for the present problem we view the restriction to a single initial belief as a limitation. We will test SARSOP on the olfactory search problem in Ref. [45].

Finally, we note that it may be possible to improve the performance of Perseus by projecting the beliefs onto a space of smaller dimension, using a form of non-negative matrix factorization (see, for example, [54]).

5. Other heuristic strategies

In addition to those detailed in Sec. III D, there are a litany of other heuristic strategies which we could have considered. Among these, we tested the *most likely state* policy which strikes towards the location with maximum belief,

$$\pi_{\text{MLS}}(b) = \pi_{\text{MDP}}^*[\arg \max_{s \in \mathcal{S}} b(s)], \quad (\text{A20})$$

and *action voting*, which selects the action most likely to be optimal according to the underlying MDP,

$$\pi_{\text{AV}}(b) = \arg \max_{a \in \mathcal{A}} \sum_{s \in \mathcal{S}} b(s) \delta(\pi_{\text{MDP}}^*(s), a), \quad (\text{A21})$$

where $\delta(\cdot, \cdot)$ is a Kronecker delta.

We found these policies to be too myopic or greedy and generally inferior to the others when applied to the search problem, so we did not present results for them in this work.

6. Testing policies

In determining whether or not a policy is good, we must first specify the problem space we are interested in solving: namely, where the agent starts its search. We limit ourselves to problems that are neither unreasonably hard nor unreasonably easy in the following way: the agent always begins its search somewhere between two isocurves of detection likelihood

$\ell = \text{Pr}(o|s)$. Specifically, we (somewhat arbitrarily) select the range $0.006\bar{S} < \ell < 0.02\bar{S}$, where, as a reminder, \bar{S} is the nondimensionalized emission rate of the source (see Fig. 2). The scaling with \bar{S} ensures that the curves are approximately invariant when the emission rate is changed. For consistency, this condition is imposed on the agent’s starting location both when we are testing the policy and when we are collecting beliefs in the initial phase of Perseus. During training, at each iteration of Perseus 10, we evaluate the resulting policy in a couple of different ways. First, we randomly select (without replacement) an ensemble of 100 points lying within the acceptable isocurves, and trial the policy starting from each point ten times. (For each S , the set of starting points is held constant across policies and iterations for consistency.) This leads to 1000 Monte Carlo trials whose arrival times and rewards are then averaged, to get a sense of the overall performance of the policy.

Because the ensemble averaging is poorly controlled, we also evaluate the policies on a small set of four fixed starting points at different distances from the source, using 1000 trials each. For a more detailed view of the statistics, for a few select policies we extend the number of trials to 10000, which we use to construct the arrival time pdfs.

It should be pointed out that this prescription for evaluating performance is *ad hoc* and in particular diverges somewhat from the rigorous definition of optimality for a POMDP. Formally, the optimal policy is optimal when conditioned on the prior, so the source should be drawn from the initial belief for a rigorous evaluation of optimality. We take this approach in Appendix C, which loosely follows the initialization used in [25]. The careful reader should understand the approach taken in the main body of the present work as the specification of an interesting problem based on phenomenology, and the use of a POMDP solver as a heuristic to find an empirically best solution to this problem.

APPENDIX B: CONVERGENCE OF PERSEUS

In this section we show how a few key features of the Perseus policies evolve from iteration to iteration, while comparing different choices of the shaping function. These features include the performance on test problems and the Bellman error. We also compare performance for different choices of γ .

1. Mean arrival times for an ensemble of starting points

In Fig. 10 we show the evolution, over Perseus iterations, of the mean arrival time for an ensemble of 100 randomly selected starting points. The agent searches from each starting point ten times, yielding ten estimates for the mean arrival time; the standard error over these ten estimates is used as our error bar. Note that, to limit computation time, the agent’s search was limited to 1000 time steps.

For each emission rate, we show as a baseline the corresponding results for the heuristic which performed best at that emission rate. These are, respectively, Thompson sampling with $\tau = 100$, space-aware infotaxis, and QMDP. We also tested Thompson sampling with $\tau = 1$ and $\tau = 10$; increasing τ beyond 100 had essentially no effect since the agent will

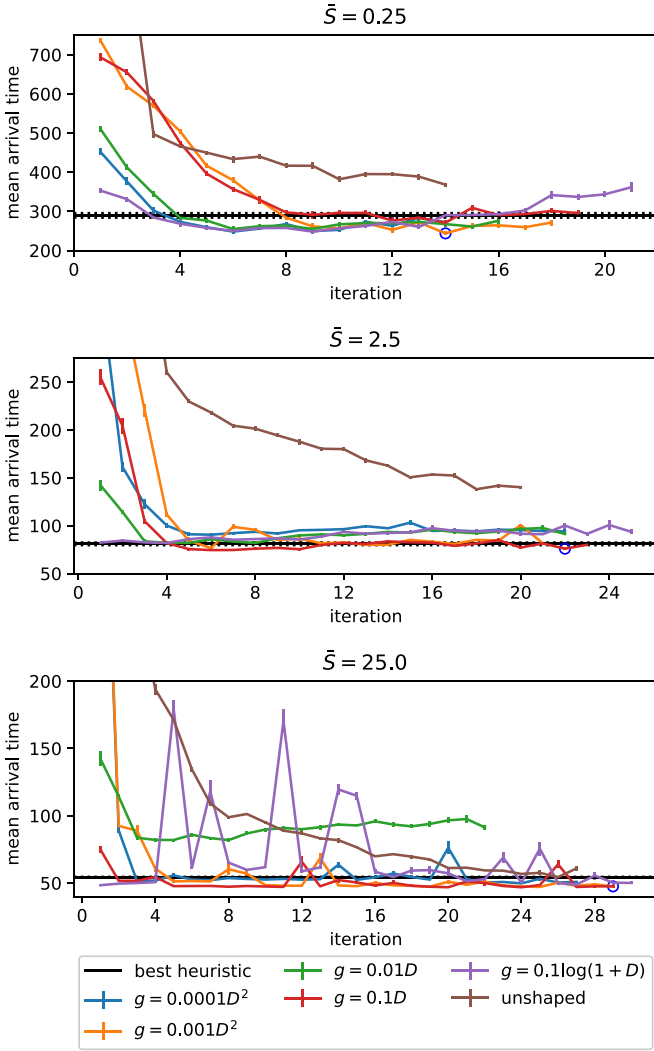


FIG. 10. Mean arrival times from ensemble of starting points in the each environment, for several choices of reward shaping. Error bars represent the standard error of the mean as estimated by the variance across the ten trials. The policy chosen for testing is circled.

almost never have to travel farther than 100 units to reach a sampled point.

An immediate takeaway is that, with a good choice of reward shaping, significantly fewer iterations are required to achieve good performance on the search problem, relative to the unshaped baseline. In fact, on the relatively large grid studied here, we have found that no number of iterations seem to suffice for unshaped Perseus to “catch up” with Perseus using a good shaping.

Another observation is that introducing a reward-shaping function can apparently reduce the stability of the policy from iteration to iteration, as evidenced by significant fluctuations in policy performance which were occasionally observed (the logarithmic shaping in the $\bar{S} = 25$ environment provides an extreme example of this). This behavior is more evident when γ is larger and is usually intensified when the shaping is increased in magnitude. This, along with the concomitant need for additional hyperparameter tuning, appears to be the main drawback of using reward shaping.




FIG. 11. RMS Bellman training and validation errors, i.e., errors on \mathcal{B} and those encountered during testing, respectively, for $\bar{S} = 2.5$.

2. Bellman error

To claim that the policies obtained using Perseus are near-optimal, one should confirm that the Bellman error [Eq. (A19)] is decreasing from iteration to iteration. In Fig. 11 we show, for $\bar{S} = 2.5$, the rms Bellman error on the belief set \mathcal{B} and on beliefs encountered during testing. Borrowing from the lexicon of machine learning, we call these “training” and “validation” errors.

3. Dependence on γ

On the left side of Fig. 12 we show how the performance of Perseus, in the absence of a reward-shaping function, depends on γ , for $\bar{S} = 2.5$. We plot the mean arrival times for the ensemble of starting points as a function of iteration for a few γ , corresponding to a range of horizons 12.5–200. Note that performance is comparable for most of the γ 's, but $\gamma = 0.99$ is too large and leads to poor stability or convergence.

In the right side of the same figure, we show the same, but with a reward shaping $g = 0.1D$ turned on. While an excellent choice when $\gamma = 0.98$, this shaping function leads to poor performance on all other γ 's. Thus it is clear that the best choice of shaping depends strongly on the choice of γ ; one should select a γ first.

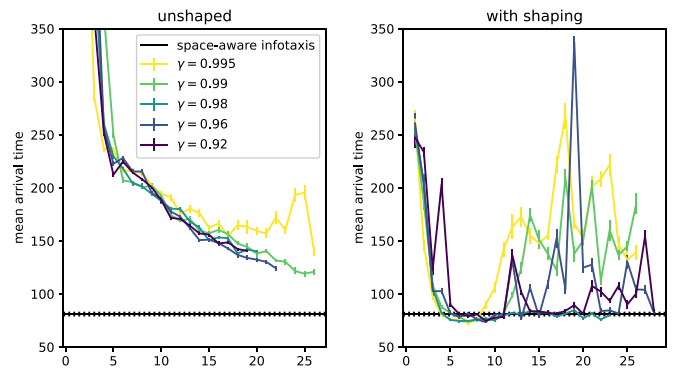


FIG. 12. Evolution of mean arrival time from ensemble of starting points using Perseus, for $\bar{S} = 2.5$, and for several choices of γ . On the left we do not use a reward-shaping function, and on the right we use $g = 0.1D$.

TABLE IV. Performance statistics of various policies with the alternate initialization described in this Appendix, with $\bar{S} = 0.25$. Note that T_{\max} was set to 2500. The MDP optimum (i.e., the mean Manhattan distance) is shown for comparison. The best-performing policies have been highlighted in bold font.

Policy	$\mathbb{E}[T T < T_{\max}]$	T_{90}	$\Pr(T > 2\langle T \rangle_{\text{perseus}})$	$\Pr(T \geq T_{\max})$
MDP optimum	21.6 ± 0.1	—	—	—
Perseus	203.2 ± 1.8	539	0.162	0.0011
Infotaxis	197.6 ± 1.8	567	0.165	$< 5 \times 10^{-5}$
SAI	136.8 ± 1.2	469	0.127	0.043
QMDP	22.2 ± 5.8	2500	0.858	0.857
Thompson, $\tau = 1$	483.8 ± 3.9	1837	0.455	0.065
Thompson, $\tau = 10$	347.0 ± 2.8	884	0.299	0.0086
Thompson, $\tau = 100$	351.3 ± 2.5	795	0.311	0.0019

APPENDIX C: ALTERNATE INITIALIZATION

Previously in this work, we have explored an initialization of the belief where the agent waits in place until it receives a detection. Another initialization, as in Ref. [25], is to force a detection at time zero. Under such a prescription, the initial belief will always be the same, and it will simply be proportional to the likelihood of detection. In this Appendix, we briefly present performance results for the various policies using this approach.

In Tables IV–VI we show several properties of the arrival time distributions. To wit, we show the mean and several properties of the tails: the time T_{90} corresponding to the 90th percentile of the pdf, the probability that the arrival time exceeds twice the mean obtained under Perseus, and the failure rate $\Pr(T \geq T_{\max})$ for $T_{\max} = 2500$. To be clear, we have followed Ref. [25] and, for each Monte Carlo trial, drawn the location of the source randomly from the initial belief, which is a more precise test of optimality.

The results are qualitatively very different from those obtained with the initialization we used in the main text. Across all three environments, one of the versions of infotaxis yielded the best results among heuristics, and Perseus performed equally well or slightly worse. In the dilute $\bar{S} = 0.25$ environment, infotaxis and Perseus are nearly indistinguishable in terms of their performance—infotaxis has a slightly lower mean arrival time but a slightly fatter tail. We find that SAI, while having a low mean arrival time when the agent does not fail, nevertheless suffers from a relatively high

failure rate (more than 4%) in this environment and thus is (arguably) not very competitive. QMDP fails more than 85% of the time. Thompson sampling, the best heuristic in the corresponding environment in the main text, is less performant with this initialization. It is not terribly surprising that the performance Thompson sampling might depend on the prior, since the strategy depends on sampling from this prior.

In the $\bar{S} = 2.5$ environment, the performance of SAI and Perseus are now nearly indistinguishable. Infotaxis also performs well, and the other heuristic policies are not competitive.

Finally, in the $\bar{S} = 25$ environment, SAI was slightly better than Perseus, and infotaxis lagged somewhat further behind. QMDP, which was the best heuristic for the high emission problem under the initialization in the main text, still suffers from a rather high failure rate.

These results illustrate the fact that the performance of various policies on the search problem depends strongly on the choice of prior, so it may be interesting to study which prior comports best with observed insect behavior. It is remarkable how well infotaxis and SAI perform in this setting, and the results here strongly suggest that, depending on the emission rate, one or both of these heuristics are very close to optimal under this prior.

Other solvers, such as SARSOP, may be better suited for this initialization than Perseus due to there being only a single initial belief; this will be checked in [45].

TABLE V. Same as Table IV, but with $\bar{S} = 2.5$. The best-performing policies have been highlighted in bold font.

Policy	$\mathbb{E}[T T < T_{\max}]$	T_{90}	$\Pr(T > 2\langle T \rangle_{\text{perseus}})$	$\Pr(T \geq T_{\max})$
MDP optimum	23.3 ± 0.1	—	—	—
Perseus	64.4 ± 0.4	143	0.131	3×10^{-4}
Infotaxis	66.8 ± 0.5	155	0.163	$< 5 \times 10^{-5}$
SAI	62.1 ± 0.4	145	0.139	1.5×10^{-4}
QMDP	86.7 ± 0.8	2500	0.444	0.281
Thompson, $\tau = 1$	138.9 ± 1.4	301	0.359	0.0027
Thompson, $\tau = 10$	102.0 ± 0.7	212	0.284	$< 5 \times 10^{-5}$
Thompson, $\tau = 100$	124.3 ± 0.7	241	0.396	$< 5 \times 10^{-5}$

TABLE VI. Same as Table IV, but with $\bar{S} = 25$. The best-performing policies have been highlighted in bold font.

Policy	$\mathbb{E}[T T < T_{\max}]$	T_{90}	$\Pr(T > 2\langle T \rangle_{\text{perseus}})$	$\Pr(T \geq T_{\max})$
MDP optimum	29.8 ± 0.1	—	—	—
Perseus	43.9 ± 0.2	76	0.070	4.5×10^{-4}
Infotaxis	48.5 ± 0.2	92	0.125	0.0018
SAI	41.6 ± 0.2	76	0.044	$< 5 \times 10^{-5}$
QMDP	51.2 ± 0.3	100	0.138	0.031
Thompson, $\tau = 1$	70.9 ± 0.5	122	0.281	5×10^{-5}
Thompson, $\tau = 10$	58.07 ± 0.2	99	0.171	$< 5 \times 10^{-5}$
Thompson, $\tau = 100$	82.7 ± 0.3	136	0.398	$< 5 \times 10^{-5}$

- [1] C. J. McMeniman, R. A. Corfas, B. J. Matthews, S. A. Ritchie, and L. B. Vosshall, Multimodal integration of carbon dioxide and other sensory cues drives mosquito attraction to humans, *Cell* **156**, 1060 (2014).
- [2] R. T. Cardé, Multi-cue integration: How female mosquitoes locate a human host, *Curr. Biol.* **25**, R793 (2015).
- [3] T. Healy and M. Copland, Activation of *Anophelesgambiae* mosquitoes by carbon dioxide and human breath, *Med. Vet. Entomol.* **9**, 331 (1995).
- [4] M. Geier, O. J. Bosch, and J. Boeckh, Influence of odour plume structure on upwind flight of mosquitoes towards hosts, *J. Exp. Biol.* **202**, 1639 (1999).
- [5] T. Dekker and R. T. Cardé, Moment-to-moment flight manoeuvres of the female yellow fever mosquito (*Aedes aegypti* L.) in response to plumes of carbon dioxide and human skin odour, *J. Exp. Biol.* **214**, 3480 (2011).
- [6] M. Gillies, The role of carbon dioxide in host-finding by mosquitoes (Diptera: Culicidae): A review, *Bull. Entomol. Res.* **70**, 525 (1980).
- [7] W. Bidlingmayer and D. Hem, The range of visual attraction and the effect of competitive visual attractants upon mosquito (Diptera: Culicidae) flight, *Bull. Entomol. Res.* **70**, 321 (1980).
- [8] F. van Breugel, J. Riffell, A. Fairhall, and M. H. Dickinson, Mosquitoes use vision to associate odor plumes with thermal targets, *Curr. Biol.* **25**, 2123 (2015).
- [9] 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).
- [10] J. Kennedy, A. Ludlow, and C. Sanders, Guidance of flying male moths by wind-borne sex pheromone, *Physiol. Entomol.* **6**, 395 (1981).
- [11] 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).
- [12] E. Yee, P. Kosteniuk, G. Chandler, C. Biloft, and J. Bowers, Statistical characteristics of concentration fluctuations in dispersing plumes in the atmospheric surface layer, *Boundary-Layer Meteorol.* **65**, 69 (1993).
- [13] G. Reddy, V. N. Murthy, and M. Vergassola, Olfactory sensing and navigation in turbulent environments, *Annu. Rev. Condens. Matter Phys.* **13**, 191 (2022).
- [14] S. D. Boie, E. G. Connor, M. McHugh, K. I. Nagel, G. B. Ermentrout, J. P. Crimaldi, and J. D. Victor, Information-theoretic analysis of realistic odor plumes: What cues are useful for determining location? *PLoS Comput. Biol.* **14**, e1006275 (2018).
- [15] J. D. Victor, S. D. Boie, E. G. Connor, J. P. Crimaldi, G. B. Ermentrout, and K. I. Nagel, Olfactory navigation and the receptor nonlinearity, *J. Neurosci.* **39**, 3713 (2019).
- [16] A. Thesen, J. B. Steen, and K. Doving, Behaviour of dogs during olfactory tracking, *J. Exp. Biol.* **180**, 247 (1993).
- [17] A. G. Khan, M. Sarangi, and U. S. Bhalla, Rats track odour trails accurately using a multi-layered strategy with near-optimal sampling, *Nat. Commun.* **3**, 703 (2012).
- [18] D. H. Gire, V. Kapoor, A. Arrighi-Allisan, A. Seminara, and V. N. Murthy, Mice develop efficient strategies for foraging and navigation using complex natural stimuli, *Curr. Biol.* **26**, 1261 (2016).
- [19] J. Jinn, E. G. Connor, and L. F. Jacobs, How ambient environment influences olfactory orientation in search and rescue dogs, *Chem. Senses* **45**, 625 (2020).
- [20] J. Kennedy, Zigzagging and casting as a programmed response to wind-borne odour: A review, *Physiol. Entomol.* **8**, 109 (1983).
- [21] S. A. Budick and M. H. Dickinson, Free-flight responses of *Drosophila melanogaster* to attractive odors, *J. Exp. Biol.* **209**, 3001 (2006).
- [22] F. van Breugel and M. H. Dickinson, Plume-tracking behavior of flying *Drosophila* emerges from a set of distinct sensory-motor reflexes, *Curr. Biol.* **24**, 274 (2014).
- [23] E. Balkovsky and B. I. Shraiman, Olfactory search at high Reynolds number, *Proc. Natl. Acad. Sci. U. S. A.* **99**, 12589 (2002).
- [24] M. Vergassola, E. Villermaux, and B. I. Shraiman, 'Infotaxis' as a strategy for searching without gradients, *Nature (London)* **445**, 406 (2007).
- [25] 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).
- [26] M. T. Spaan and N. Vlassis, Perseus: Randomized point-based value iteration for POMDPs, *J. Artif. Intell. Res.* **24**, 195 (2005).
- [27] A. Y. Ng, D. Harada, and S. Russell, Policy invariance under reward transformations: Theory and application to reward shaping, in *Proceedings of the Sixteenth International Conference on Machine Learning (ICML 1999)*, edited by I. Bratko and S. Dzeroski, Vol. 99 (Morgan Kaufmann, San Francisco, 1999), pp. 278–287.

- [28] Y. Kuwana, S. Nagasawa, I. Shimoyama, and R. Kanzaki, Synthesis of the pheromone-oriented behaviour of silkworm moths by a mobile robot with moth antennae as pheromone sensors, *Biosens. Bioelectron.* **14**, 195 (1999).
- [29] P. Pyk, S. Bermúdez i Badia, U. Bernardet, P. Knüsel, M. Carlsson, J. Gu, E. Chanie, B. S. Hansson, T. C. Pearce, and P. F. J. Verschure, An artificial moth: Chemical source localization using a robot based neuronal model of moth optomotor anemotactic search, *Auton. Robots* **20**, 197 (2006).
- [30] J.-B. Masson, Olfactory searches with limited space perception, *Proc. Natl. Acad. Sci. U. S. A.* **110**, 11261 (2013).
- [31] D. Martinez and E. M. Moraud, Reactive and cognitive search strategies for olfactory robots, *Neuromorphic Olfaction* **5**, 153 (2013).
- [32] A. Celani, E. Villermaux, and M. Vergassola, Odor Landscapes in Turbulent Environments, *Phys. Rev. X* **4**, 041015 (2014).
- [33] M. v. Smoluchowski, Versuch einer mathematischen Theorie der Koagulationskinetik kolloider Lösungen, *Z. Phys. Chem.* **92**, 129 (1918).
- [34] J. Masson, M. B. Bechet, and M. Vergassola, Chasing information to search in random environments, *J. Phys. A: Math. Theor.* **42**, 434009 (2009).
- [35] K. J. Astrom, Optimal control of Markov decision processes with incomplete state estimation, *J. Math. Anal. Appl.* **10**, 174 (1965).
- [36] L. P. Kaelbling, M. L. Littman, and A. R. Cassandra, Planning and acting in partially observable stochastic domains, *Artif. Intell.* **101**, 99 (1998).
- [37] A. R. Cassandra, Exact and approximate algorithms for partially observable Markov decision processes, Ph.D. thesis, Brown University, 1998.
- [38] J. Pineau, G. Gordon, and S. Thrun, Anytime point-based approximations for large POMDPs, *J. Artif. Intell. Res.* **27**, 335 (2006).
- [39] G. Shani, J. Pineau, and R. Kaplow, A survey of point-based POMDP solvers, *Auton. Agents Multi-Agent Syst.* **27**, 1 (2013).
- [40] J. L. Fernández, R. Sanz, R. G. Simmons, and A. R. Diéguez, Heuristic anytime approaches to stochastic decision processes, *J. Heuristics* **12**, 181 (2006).
- [41] W. R. Thompson, On the likelihood that one unknown probability exceeds another in view of the evidence of two samples, *Biometrika* **25**, 285 (1933).
- [42] W. R. Thompson, On the theory of apportionment, *Am. J. Math.* **57**, 450 (1935).
- [43] D. J. Russo, B. Van Roy, A. Kazerouni, I. Osband, and Z. Wen, A tutorial on Thompson sampling, *Found. Trends Machine Learn.* **11**, 1 (2018).
- [44] I. Osband, B. Van Roy, D. J. Russo, and Z. Wen, Deep exploration via randomized value functions, *J. Mach. Learn. Res.* **20**, 1 (2019).
- [45] A. Loisy and R. A. Heinonen, Deep reinforcement learning for the olfactory search POMDP: A quantitative benchmark, *Eur. Phys. J. E* **46**, 17 (2023).
- [46] 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).
- [47] 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, 2003).
- [48] R. A. Howard and J. E. Matheson, Risk-sensitive Markov decision processes, *Manag. Sci.* **18**, 356 (1972).
- [49] E. V. Denardo, Contraction mappings in the theory underlying dynamic programming, *SIAM Rev.* **9**, 165 (1967).
- [50] G. Shani, R. I. Brafman, and S. E. Shimony, Prioritizing point-based POMDP solvers, in *European Conference on Machine Learning*, edited by J. Fürnkranz, T. Scheffer, and M. Spiliopoulou (Springer, Berlin, Heidelberg, 2006), pp. 389–400.
- [51] 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*, edited by O. Brock, J. Trinkle, and F. Ramos, Vol. 2008 (MIT Press, Cambridge, 2008).
- [52] T. Smith and R. Simmons, Heuristic search value iteration for POMDPs, in *UAI'04 Proceedings of the 20th Conference on Uncertainty in Artificial Intelligence*, edited by D. M. Chickering and J. Y. Halpern (Arlington AUAI Press, 2004), pp. 520–527.
- [53] T. Smith and R. Simmons, Point-based POMDP algorithms: Improved analysis and implementation, in *UAI'05 Proceedings of the 20th Conference on Uncertainty in Artificial Intelligence*, edited by F. Bacchus and T. Jaakkola (Arlington AUAI Press, 2005), pp. 542–549.
- [54] X. Li, W. K. Cheung, J. Liu, and Z. Wu, A novel orthogonal NMF-based belief compression for POMDPs, in *Proceedings of the 24th International Conference on Machine Learning*, edited by Z. Ghahramani (Association for Computing Machinery, New York, 2007), pp. 537–544.