PROCEEDINGS A

royalsocietypublishing.org/journal/rspa

Research

Cite this article: Loisy A, Eloy C. 2022 Searching for a source without gradients: how good is infotaxis and how to beat it. *Proc. R. Soc. A* **478**: 20220118.
https://doi.org/10.1098/rspa.2022.0118

Received: 17 February 2022
Accepted: 12 May 2022

Subject Areas:
statistical physics, artificial intelligence, applied mathematics

Keywords:
olfactory search, infotaxis, partially observable Markov decision process, reinforcement learning

Authors for correspondence:
Aurore Loisy
e-mail: aurore.loisy@irphe.univ-mrs.fr
Christophe Eloy
e-mail: christophe.eloy@centrale-marseille.fr

Electronic supplementary material is available online at https://doi.org/10.6084/m9.figshare.c.6043560.

THE ROYAL SOCIETY PUBLISHING

Searching for a source without gradients: how good is infotaxis and how to beat it

Aurore Loisy and Christophe Eloy

Aix Marseille University, CNRS, Centrale Marseille, IRPHE, Marseille, France

AL, 0000-0002-8089-8636; CE, 0000-0003-4114-7263

Infotaxis is a popular search algorithm designed to track a source of odour in a turbulent environment using information provided by odour detections. To exemplify its capabilities, the source-tracking task was framed as a partially observable Markov decision process consisting in finding, as fast as possible, a stationary target hidden in a two-dimensional grid using stochastic partial observations of the target location. Here, we provide an extended review of infotaxis, together with a toolkit for devising better strategies. We first characterize the performance of infotaxis in domains from one dimension to four dimensions. Our results show that, while being suboptimal, infotaxis is reliable (the probability of not reaching the source approaches zero), efficient (the mean search time scales as expected for the optimal strategy) and safe (the tail of the distribution of search times decays faster than any power law, though subexponentially). We then present three possible ways of beating infotaxis, all inspired by methods used in artificial intelligence: tree search, heuristic approximation of the value function, and deep reinforcement learning. The latter is able to find, without any prior human knowledge, the (near) optimal strategy. Altogether, our results provide evidence that the margin of improvement of infotaxis towards the optimal strategy gets smaller as the dimensionality increases.

1. Introduction

Tracking down a source of odour in a turbulent environment is a task performed by numerous animals [1–7]. In the hope of uncovering the search algorithms used in nature, Vergassola *et al.* [8] formulated what we will refer to as the 'source-tracking problem'. The

modelling of this biology-inspired problem relies on physics. Its solutions, however, call for methods of operations research, automated planning and artificial intelligence. Our first objective here is to make this problem and its solutions accessible to these different communities, which are not always aware of each other's work.

A source that emits a chemical substance, such as an odour, in a quiescent fluid generates a smooth concentration field that decays with the distance to the emission point. This source is easily tracked, because one only needs to follow concentration gradients: this behaviour, ubiquitous in biology, is called chemotaxis. Source-tracking becomes much harder in a turbulent medium, because the concentration field consists of disconnected patches of high concentration, randomly distributed and separated by voids. An animal able to detect concentration levels would receive a highly intermittent signal consisting of sharp peaks separated by long periods with no measurable concentration [7,9]. Gradient-based strategies are therefore doomed to fail in turbulence.

To mimic these searching conditions in a computation-friendly environment, Vergassola *et al.* [8] designed the source-tracking problem. It consists of minimizing the number of steps (cumulated cost) to reach a source (target) hidden in a discrete grid. At each step, the searcher (agent) moves to a neighbour cell (action) and receives a stochastic sensory signal (observation, called 'hits'). These observations model odour detection events, and provide a noisy information about the distance to the source (stochastic partial information). The agent knows how observations are generated (the model), and has a perfect memory of past observations and actions, therefore it can maintain a probability distribution over source locations (belief) updated after each observation using Bayes' inference theorem. Solving this problem consists in finding the optimal strategy (policy), defined as the mapping from belief to action, that minimizes the expected number of steps to reach the source.

In the language of artificial intelligence and related disciplines, the source-tracking problem is a partially observable Markov decision process (POMDP), which is framed as a belief-MDP (an MDP—Markov decision process—where states are replaced by belief states). It actually belongs to a narrower class of problems, called partially observable stochastic shortest path problems, for which a few formal mathematical results exist, namely: the existence of a deterministic policy that is optimal within the class of Markov policies, and the pointwise convergence of value iteration to the unique bounded fixed point of the dynamic programming operator [10,11].

Yet computing the optimal policy is not feasible: POMDP exact solvers take prohibitively large amounts of computation time for any but the smallest problems [12–15], and the continuous nature of the belief space prevents the use of exact methods for tabular MDPs (e.g. value iteration). Approximate solution methods exist [14–17], but scalability remains an issue. An alternative is deep reinforcement learning: originally designed to solve MDPs using a neural network approximation of the value function [18], existing algorithms can account for partial observability using the belief-MDP formulation.

Another approach is to handcraft heuristic policies using intuition and knowledge. Infotaxis is a such a heuristic: it states that the agent should choose the action from which it expects the greatest information gain about the source location. The physical intuition behind this algorithm is, quoting the authors, that 'information accumulates faster close to the source because cues arrive at a higher rate, hence tracking the maximum rate of information acquisition will guide the searcher to the source much like concentration gradients in chemotaxis' [8].

The infotaxis policy is more precisely defined as the policy that maximizes the expected gain of information, defined as the decrement of Shannon entropy of the belief, over a one-step horizon. In other words, infotaxis is an information-greedy policy. The idea of information maximization (or, equivalently, uncertainty minimization) using a probabilistic formulation can be traced back to Cassandra *et al.* [19], who proposed it as a heuristic for robot navigation in uncertain environments, and it has since become a central concept for robotic exploration algorithms [20].

Infotaxis is generally believed to be robust and has become a popular search algorithm. It has been implemented in robots [21–23] and has inspired various extensions: infotaxis on different

types of lattices [24,25], continuous-space infotaxis [26], mapless infotaxis [23], collective infotaxis [24], socialtaxis [27], entrotaxis [28] and energy-constrained proportional betting [29]. Its relevance to biological searches remains speculative, but some attempts have been made to explain the trajectories of moths [8,30] and of the worm *C. elegans* [31] by infotactic searches.

Yet, the supposedly good performance of infotaxis is somewhat surprising: there is, *a priori*, no reason why greedily reducing uncertainty should minimize the time to reach the source (unless the two were linearly related; it is easy to check that they are not). Besides, infotaxis was designed to bring the agent close to the source in a robust manner. It is known not to be optimal with respect to the search time, and that better performance can be achieved by shifting it towards a more exploitative behaviour. This latter approach is not satisfactory though, because it introduces a tunable parameter whose appropriate value is far from universal [23,24].

Several papers implemented (variants of) infotaxis and confirmed its ability to lead to the source in various searching conditions [24–26,28,32–34]. A number of studies also evaluated its robustness to model uncertainty [24,28,32,34]. Nevertheless, the performance of the original infotaxis algorithm has never been assessed in a *systematic* manner. In addition, most of these prior studies applied infotaxis to searches in two-dimensional domains: very few considered the three-dimensional problem [26,33], and none tested the algorithm in one dimension. The questions of how infotactic trajectories are affected by dimensionality and whether infotaxis generalizes well to high dimensional spaces remain vastly unexplored.

In this paper, our main contributions are the following. (1) We generalize the source-tracking problem to any space dimension and we reformulate it using the POMDP framework, which allows us to formally write the optimal strategy as the solution of a recurrence equation (known as the Bellman optimality equation). (2) We introduce a protocol for a rigorous evaluation of intofaxis performance (in the sense that it is not affected by finite-size effects or by an arbitrarily chosen initial distance to the source) and we propose an efficient method to compute the distribution of arrival times. (3) Using this protocol, we compute and analyse the performance of infotaxis as a function of the problem dimensionality (from one dimension to four dimensions) and of the relevant parameters that govern the physics of odour propagation and detection. (4) We evaluate to what extent three well-established techniques from artificial intelligence are able to improve on infotaxis, or even to yield optimal strategies: (i) tree search based on an existing heuristic (N-step infotaxis), (ii) knowledge-based approximation of the optimal value function (space-aware infotaxis) and (iii) deep reinforcement learning.

The lack of directionality in the information provided by observations makes the source-tracking problem particularly challenging. The hardest scenario is the one where no mean flow is present to break the radial symmetry of the problem. We focus on this worst-case scenario, which is expected to provide lower bounds on performance for situations where a mean flow is present, as is often the case in real applications. Our results provide extensive evidence that, while being suboptimal, infotaxis has three important properties: (i) reliability (the probability of never finding the source approaches zero), (ii) efficiency (the mean search time scales as expected for the optimal policy) and (iii) safety (in the sense that arrival times are not plagued by large fluctuations). The performance of space-aware infotaxis and of (near) optimal policies obtained by reinforcement learning strongly suggest that while infotaxis is vastly suboptimal in one dimension, the margin of improvement towards the optimal policy gets tighter as the dimensionality increases.

The remainder of this paper is organized as follows. The source-tracking problem and its POMDP formulation are presented in §2, together with our policy evaluation protocol. Infotaxis and its generalization to multi-step anticipation are described and evaluated in §3. Space-aware infotaxis is introduced in §4, where it is shown to do better than infotaxis without any tunable parameter. Section 5 is concerned with deep reinforcement learning and presents results for (near) optimal policies. The reasons behind the good performance of infotaxis are discussed in §6. Conclusions are drawn in §7.

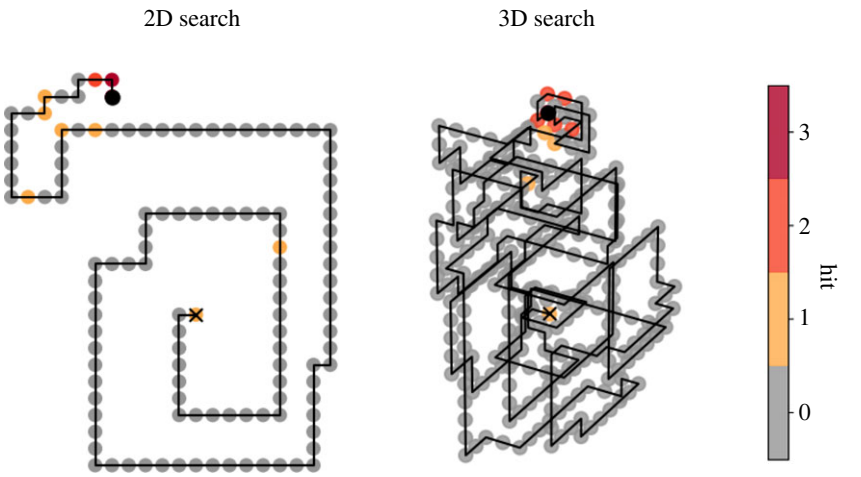

Figure 1. Examples of trajectories in the source-tracking problem, where a searcher (agent) must find a source (hidden target) using information provided by odour detections called 'hits' (partial observations). The start of the search is indicated by a cross, and the end of the search (source location) is depicted by a black dot. These trajectories were generated using the infotaxis policy. Videos are provided in the electronic supplementary material. (Online version in colour.)

2. The source-tracking partially observable Markov decision process

(a) Overview

The source-tracking problem, illustrated in figure 1, is a POMDP in which the agent seeks to identify the true location of an imperfectly observed stationary target (a source of odour). The environment is an n-dimensional Cartesian grid and the source, invisible to the agent, is located in one of the cells. At each step, the agent moves to a neighbour cell and receives an observation, which provides a noisy measurement of its distance to the source. The search continues until the agent moves to the cell occupied by the source. The agent has a perfect memory and a perfect knowledge of the process that generates observations. How should the agent behave in order to reach the source in the smallest possible number of steps?

We denote $\boldsymbol{x}^s$ and $\boldsymbol{x}^a$ the source and the agent locations (n-tuples of integers), respectively. Allowed actions are moves to a neighbour cell along the grid axes, for example, in three dimensions, the set of allowed actions is 'north', 'south', 'east', 'west', 'top', 'bottom'. Staying in the same location is not allowed (unlike in [8], we made this choice because preliminary tests showed that standing still is almost never beneficial). After moving to a cell, either the source is found (event F) and the search is over, or the source is not found (complementary event $\bar{F}$) and the agent receives a stochastic sensor measurement h (hits), which represents the integer number of odour particles detected by the agent. Hits are received with conditional probability $\Pr(h|\boldsymbol{x}^a, \boldsymbol{x}^s)$ (the probabilistic model of detection will be presented in §2c). We encompass the presence/absence of the source and the number of hits in a single observation variable o. Possible observations are $o \in \{F, (\bar{F}, 0), (\bar{F}, 1), (\bar{F}, 2), \ldots\}$.

We denote $p(\boldsymbol{x})$ the discrete probability distribution of the source being in each grid cell, that is $p(\boldsymbol{x}) = \Pr(\boldsymbol{x}^s = \boldsymbol{x})$. After each action and observation, $p(\boldsymbol{x})$ can be updated using Bayesian inference (§2d). The agent has access to its position $\boldsymbol{x}^a$ and to the distribution $p(\boldsymbol{x})$. In the POMDP terminology, this defines a belief state $s = [\boldsymbol{x}^a, p(\boldsymbol{x})]$. Examples of belief states are shown in figure 2. Finding the source is a special belief state s^Ω where the source position is known and matches the agent position: $s^\Omega = [\boldsymbol{x}^a, \delta(\boldsymbol{x} - \boldsymbol{x}^a)]$. The agent's behaviour is described by a policy, denoted π, which maps each belief state to an action. For a deterministic policy (as seeked here, since the optimal policy is deterministic), the action chosen is $a = \pi(s)$.

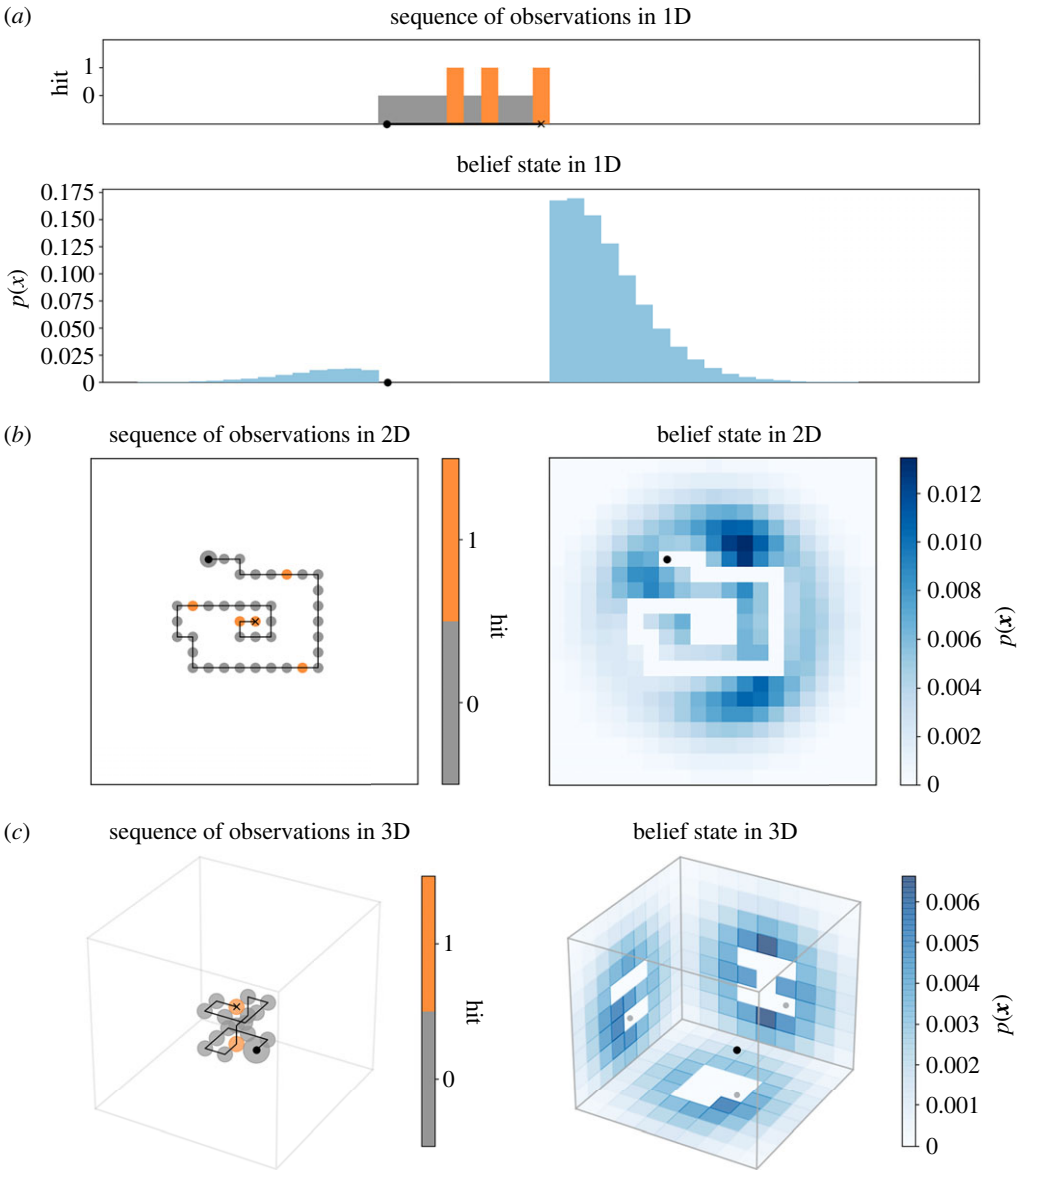

Figure 2. Examples of observation sequences and corresponding belief states in (*a*) one dimension, (*b*) two dimensions and (*c*) three dimensions. The belief state consists of the agent location (depicted by a black dot) and of a discrete probability distribution over the possible source locations. It accounts for all past observations made by the agent (the agent's initial position is shown by a cross in observation maps). In (*c*), the source probability distribution is shown in cut-off planes containing the agent. (Online version in colour.)

The search proceeds as follows:

— Initially

(i) The belief state is $s_0 = [x_0^a, p_0(x)]$, where the agent location x_0^a is at the centre of the domain and where the prior distribution of source location $p_0(x)$ is drawn randomly from the set of priors (details are provided in appendix A).
(ii) The source location x^s is drawn randomly according to $p_0(x)$.

— At the tth step of the search

(i) Knowing the current belief state $s_t = [x_t^a, p_t(x)]$, the agent chooses an action according to some policy $\pi: a_t = \pi(s_t)$.

(ii) The agent moves deterministically to the neighbour cell associated with a_t. This move is associated with a unit cost. The agent's position is updated to x_{t+1}^a.

(iii) The agent receives an observation o_t and the source location distribution is updated, using Bayes' rule, to $p_{t+1}(x) = \text{Bayes}(p_t(x), x_{t+1}^a, o_t)$ (the Bayes operator will be made explicit in §2d).

(a) If $s_{t+1} = s^\Omega$, that is $x_{t+1}^a = x^s$, the search terminates and the agent receives no more costs.

(b) Otherwise, the search continues to step $t+1$.

Each episode (each search) is a sequence like this

$$s_0, a_0, o_0, s_1, a_1, o_1, \ldots, s_{T-1}, a_{T-1}, o_{T-1}, s^\Omega.$$

The cumulated cost of an episode is equal to the number of steps T to termination (which can be infinite if the source is never found). A step-by-step illustration depicting how a search proceeds is provided in the electronic supplementary material.

The performance of a policy π is measured by $\mathbb{E}_{p_0,\pi}[T]$, the expected number of steps to reach the source, where the expectation is over all possible sequences generated following policy π, over all possible priors p_0, and, implicitly, over all possible source locations that are distributed according to these priors. The optimal policy is the policy that minimizes $\mathbb{E}_{p_0,\pi}[T]$. We expand on the optimal policy in the next section.

Since the search does not stop until the source is found, any policy that fails to ensure termination with probability one has an infinite cumulated cost. While it is trivial to show that a policy which guarantees termination exists (any policy that exhaustively searches the grid), the policies we will consider (including infotaxis) do not have such a guarantee. Therefore, we will evaluate the performance of a policy based on Pr(failure), the probability of never finding the source, and on Mean(T), the mean number of steps to reach the source, restricted to episodes where the source is ultimately found.

Policy evaluation is performed by averaging over a large number of episodes. We use a hybrid Bayesian/Monte Carlo method, which allows faster convergence than traditional Monte Carlo simulations. Details on our approach are given in appendix B and the electronic supplementary material.

(b) Optimal policy

Solving the source-tracking problem means finding the optimal policy π^* that minimizes the duration of the search

$$\pi^* = \arg\min_\pi \mathbb{E}_{p_0,\pi}[T]. \tag{2.1}$$

The optimal policy can, at least formally, be determined from the solution of a recurrence equation known as the Bellman optimality equation as follows.

The optimal value function $v^*(s)$ of a belief state s is defined as the minimum, over all policies, of the expected number of steps remaining to find the source when starting from belief state s

$$v^*(s) = \min_\pi v^\pi(s) \quad \text{where } v^\pi(s) = \mathbb{E}_\pi[T - t | s_t = s]. \tag{2.2}$$

It satisfies the Bellman optimality equation

$$v^*(s) = \min_a \sum_{s'} \Pr(s'|s,a)[1 + v^*(s')] \quad \forall s \neq s^\Omega, \tag{2.3}$$

where $\Pr(s'|s,a)$ is the probability of transitioning from belief state s to the next belief state s' after executing action a, and where $v^*(s^\Omega) = 0$ (by definition of the terminal state s^Ω). Possible

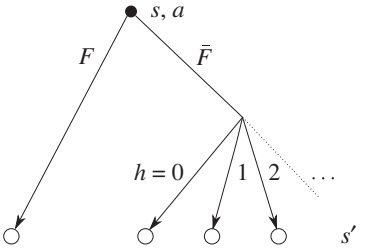

Figure 3. Tree of possible successor belief states s' starting from a belief state s and executing action a. Transitions from s to s' are determined by the observations: either finding the source (F) or not finding the source ($\bar{F}$) and receiving h hits.

transitions from s to s' correspond to possible observations: either finding the source (F) or not finding the source ($\bar{F}$) and receiving h hits, as illustrated in figure 3.

Given $v^*(s)$, the optimal policy consists in choosing the action that minimizes the expected number of remaining steps $v^*(s')$

$$\pi^*(s) = \arg\min_a \sum_{s'} \Pr(s'|s,a)[1 + v^*(s')]. \tag{2.4}$$

When the number of states s is finite (and small enough), equation (2.3) can be solved using dynamic programming [18,35]. In the source-tracking problem, however, s contains a discrete probability distribution, of which there are infinitely many. More precisely, there are as many belief states as there are possible histories (sequences of actions and observations, which can be infinitely long). Even if a maximum duration $t_{\max}$ is imposed, the number of possible beliefs grows as $(A \times O)^{t_{\max}}$ where O is the number of possible observations (at least three) and A is the number of possible actions (twice the number of dimensions). If we consider a two-dimensional domain of 10 by 10 cells: $A = 4$, $O = 3$. With $t_{\max} = 100$, the number of possible beliefs is larger than the number of atoms in the observable universe, which prevents from computing the optimal policy exactly. It can, however, be approximated (§5).

A lower bound for $\mathbb{E}_{p_0,\pi^*}[T]$ can be obtained by considering an omniscient agent that has access to the source location. In that case the optimal policy is simply to follow the shortest path to the source, and T is the Manhattan distance $\|\boldsymbol{x}^s - \boldsymbol{x}^a\|_1$ between the agent and the source. The lower bound is obtained by taking the expectation of T over the initial source distribution

$$\mathbb{E}_{p_0,\pi^*}[T] \geq \mathbb{E}_{p_0}\left[\sum_{\boldsymbol{x}} p_0(\boldsymbol{x})\|\boldsymbol{x} - \boldsymbol{x}_0^a\|_1\right]. \tag{2.5}$$

Note that this lower bound is not tight since it assumes full observability, however, it should be approached as the rate of information received by the agent increases (that is, as the source rate of emission increases).

An upper bound for $\mathbb{E}_{p_0,\pi^*}[T]$ can be obtained by considering an agent without sensors (no information provided by hits) that searches the domain exhaustively following a predetermined trajectory. The optimal trajectory in this scenario is an optimization problem on its own [36]. However, a simple yet efficient way to cover space is by 'spiralling' outward from the starting point (visiting all cells within a Chebyshev distance d of the agent's initial position, then covering all cells at a distance $d + 1$, and so on, avoiding as much as possible visiting the same cell twice). In one dimension, such spirals are not efficient (the constant back-and-forth motion of the agent would result in a quadratic scaling of the search time with the domain size). An alternative simple trajectory consists in going to one end of the grid and then to the other end. In either case, we call this policy $\pi^{\text{exhaustive}}$ and we can write

$$\mathbb{E}_{p_0,\pi^*}[T] \leq \mathbb{E}_{p_0,\pi^{\text{exhaustive}}}[T] = \sum_{\boldsymbol{x}} p_0(\boldsymbol{x})r(\boldsymbol{x}), \tag{2.6}$$

where $r(\boldsymbol{x})$ is the time step at which location $\boldsymbol{x}$ is visited following the prescribed trajectories (the first visited cell corresponds to $r=1$, the next one to $r=2$, etc.).

(c) Probabilistic model of detections

We now need to specify the model used to generate observations (hits), that is, to specify $\Pr(h|\boldsymbol{x}^a,\boldsymbol{x}^s)$. This model is based on a physical modelling of dispersion and detection in a turbulent flow as follows.

The source emits detectable odour particles with a finite lifetime that disperse in the ambient turbulent environment and can be detected by the searcher. The detection events, or 'hits', are distributed according to a Poisson's Law

$$\Pr(h|\mu)=\frac{\mu^h\exp(-\mu)}{h!}, \tag{2.7}$$

where μ is the mean number of hits. It is a function of the Euclidean distance $d=\|\boldsymbol{x}^s-\boldsymbol{x}^a\|_2$ between the agent and the source such that

$$\Pr(h|\boldsymbol{x}^a,\boldsymbol{x}^s)=\Pr(h|\mu(d)) \quad \text{with } d=\|\boldsymbol{x}^s-\boldsymbol{x}^a\|_2. \tag{2.8}$$

The derivation of $\mu(d)$ for an arbitrary number of dimensions n is given in the electronic supplementary material. The resulting expressions are provided below:

$$n=1: \quad \mu(d)=R\Delta t\frac{\lambda}{\lambda-a}\exp\left(\frac{-d}{\lambda}\right) \tag{2.9a}$$

$$n=2: \quad \mu(d)=R\Delta t\frac{1}{\ln(\lambda/a)}K_0\left(\frac{d}{\lambda}\right) \tag{2.9b}$$

$$n=3: \quad \mu(d)=R\Delta t\frac{a}{d}\exp\left(\frac{-d}{\lambda}\right) \tag{2.9c}$$

and

$$n=4: \quad \mu(d)=R\Delta t\left(\frac{a}{\lambda}\right)^2\frac{\lambda}{d}K_1\left(\frac{d}{\lambda}\right), \tag{2.9d}$$

and more generally for $n\geq 3$

$$\mu(d)=R\Delta t\left(\frac{a}{\lambda}\right)^{n-2}\left(\frac{\lambda}{d}\right)^{n/2-1}\frac{(n-2)}{\Gamma(n/2)}\frac{K_{n/2-1}(d/\lambda)}{2^{n/2-1}}, \tag{2.9e}$$

where a is the agent radius, λ is the dispersion lengthscale of the particles in the medium, R is the source emission rate, Δt is the duration of a sensor measurement, Γ is the gamma function and K_ν is the modified Bessel function of the second kind of order ν. Note that we recover the expressions provided in [8] for $n=2$ and $n=3$.

(d) Update by Bayesian inference

Each observation provides some information about the source location, which can be accounted for using Bayesian inference. In the case of a sequential process such as ours, Bayes' rule can be applied after each observation to maintain an up-to-date belief $p(\boldsymbol{x})$ that encompasses all the information gathered so far. The update after observing o_t in $\boldsymbol{x}^a_{t+1}$ reads

$$p_{t+1}(\boldsymbol{x})=\text{Bayes}(p_t(\boldsymbol{x}),\boldsymbol{x}^a_{t+1},o_t), \tag{2.10}$$

where $\text{Bayes}(p(\boldsymbol{x}),\boldsymbol{x}^a,o)$ is the operator that maps the prior p_t to the posterior p_{t+1} through Bayes' rule

$$\text{Bayes}(p(\boldsymbol{x}),\boldsymbol{x}^a,o)=\frac{\Pr(o|\boldsymbol{x}^a,\boldsymbol{x})p(\boldsymbol{x})}{\sum_{\boldsymbol{x}'}\Pr(o|\boldsymbol{x}^a,\boldsymbol{x}')p(\boldsymbol{x}')}, \tag{2.11}$$

and where $\Pr(o|\boldsymbol{x}^a,\boldsymbol{x})$ is called the evidence in Bayesian terminology.

Let us now go through the update rule for each observation. If $o = F$, the source has been found in $\boldsymbol{x}^a$, and the posterior distribution is simply a Dirac distribution

$$\text{Bayes}(p(\boldsymbol{x}), \boldsymbol{x}^a, F) = \delta(\boldsymbol{x} - \boldsymbol{x}^a). \tag{2.12}$$

Otherwise, $o = (\bar{F}, h)$, meaning that the source has not been found and that h hits were perceived. The posterior distribution after not finding the source is a simple renormalization

$$\text{Bayes}(p(\boldsymbol{x}), \boldsymbol{x}^a, \bar{F}) = \begin{cases} 0 & \text{if } \boldsymbol{x} = \boldsymbol{x}^a, \\ \dfrac{p(\boldsymbol{x})}{\sum_{\boldsymbol{x}' \neq \boldsymbol{x}^a} p(\boldsymbol{x}')} & \text{otherwise.} \end{cases} \tag{2.13}$$

The posterior after a hit h is

$$\text{Bayes}(p(\boldsymbol{x}), \boldsymbol{x}^a, h) = \frac{\Pr(h|\boldsymbol{x}^a, \boldsymbol{x}) p(\boldsymbol{x})}{\sum_{\boldsymbol{x}'} \Pr(h|\boldsymbol{x}^a, \boldsymbol{x}') p(\boldsymbol{x}')}. \tag{2.14}$$

The full update after observing $o = (\bar{F}, h)$ is therefore given by the successive application of each partial update

$$\text{Bayes}(p(\boldsymbol{x}), \boldsymbol{x}^a, o) = \text{Bayes}(\text{Bayes}(p(\boldsymbol{x}), \boldsymbol{x}^a, \bar{F}), \boldsymbol{x}^a, h). \tag{2.15}$$

A step-by-step illustration of Bayesian updates performed during a search is provided in the electronic supplementary material.

(e) Parameters of the problem

The source-tracking problem is parameterized by

(i) the space resolution Δx (the size of the agent step, also the linear size of a grid cell);
(ii) the time resolution Δt (how often does the agent make a decision, also the integration time for the sensors); or alternatively the agent speed v and then $\Delta t = \Delta x / v$;
(iii) the probabilistic law for hits encounters, itself parameterized by $(a, \lambda, R, \Delta t)$, where a is the agent radius, λ is the characteristic lengthscale of dispersion and R is the source emission rate;
(iv) the initial conditions (prior source distribution and agent's position);
(v) the grid size N (linear size of the domain).

The number of parameters can be greatly reduced by assuming that the search is initialized by a (non-zero) hit. The advantages are twofold: (i) the definition of the start is not arbitrary but instead corresponds to the moment when the agent is informed that there is source (as opposed to nothing) in the neighbourhood and where tracking it down becomes meaningful, and (ii) the grid size can be chosen large enough such that the domain boundaries play virtually no role in the search (this is because the source probability distribution after a hit decays exponentially with the distance to the agent). Note that mimicking an open world is highly desirable since the physical problem of dispersion is modelled for an infinite space and any finite-size effect in this context would be physically irrelevant. The details of this initialization procedure are provided in appendix A. We will also assume that the size of the agent step is equal to its body diameter $\Delta x = 2a$.

With our initialization protocol, the source-tracking problem involves only four parameters, which govern the physics of propagation and detection (Δx, λ, R, Δt). From those, we can construct two independent dimensionless numbers, we chose

$$\mathcal{L} = \frac{\lambda}{\Delta x} \quad \text{and} \quad \mathcal{I} = R \Delta t, \tag{2.16}$$

which characterize the problem size and the source intensity, respectively.

Together with the problem dimensionality (one dimension, two dimensions, etc.), $\mathcal{L}$ and $\mathcal{I}$ are the relevant physical parameters of the source-tracking POMDP.

(f) Source-tracking is hard

To get a sense of the difficulty of the source-tracking problem, it is worth reporting the performance of simple (and somewhat naive) heuristic policies. Examples of such policies are the usual greedy policy (which maximizes the probability of finding the source at the next step), and policies that would be optimal in the absence of uncertainty (such as the most likely state and the voting policies, both proposed by Cassandra *et al.* [19], and a policy that minimizes the expected distance to the source at the next step).

The description and the performance of these four policies are provided in the electronic supplementary material. Reliability is very poor for all of them, with Pr(failure) typically larger than 1% or even 10%. The greedy policy, which is locally optimal, leads the agent astray when few detections occur. Policies that would be optimal in the absence of uncertainty are plagued by the emergence of loops in the agent's trajectories. These loops arise because of the high degree of symmetry in $p(\boldsymbol{x})$, due to the absence of directionality in the information provided by hits. Videos illustrating the behaviour of these different policies are provided in the electronic supplementary material.

3. The infotaxis policy

Infotaxis [8] is a heuristic solution to the source-tracking problem that greedily minimizes uncertainty on the source location: at each step, an infotactic agent chooses the action that maximizes the expected information gain. In this section, we describe the infotaxis policy and evaluate its performance.

(a) Description of infotaxis

The (Shannon) entropy of a belief state $s = [\boldsymbol{x}^a, p(\boldsymbol{x})]$ is defined by

$$H(s) = -\sum_{\boldsymbol{x}} p(\boldsymbol{x}) \log_2 p(\boldsymbol{x}), \tag{3.1}$$

and is a measure of how uncertain is the source location (we use the logarithm with base 2, as is standard in information theory where H is measured in bits of information). In particular, the entropy of a Dirac distribution is zero, so $H(s^{\Omega}) = 0$. Note that H is independent of the agent's position $\boldsymbol{x}^a$.

The expected entropy upon taking action a in belief state s, denoted $H(s|a)$, is the expected entropy of successor belief states s'

$$H(s|a) = \sum_{s'} \Pr(s'|s,a) H(s'), \tag{3.2}$$

where $\Pr(s'|s,a)$ is the probability of transitioning from belief state s to next belief state s' after executing action a (cf. figure 3). An explicit example showing how $H(s|a)$ is calculated in practice is provided in the electronic supplementary material.

The information gain associated with action a in belief state s is then given by

$$G(s,a) = H(s) - H(s|a). \tag{3.3}$$

Note that the information gain is also, by definition [37], the mutual information between the current belief state s and the possible observations after action a. The infotaxis policy, denoted $\pi^{\text{infotaxis}}$, consists in choosing the action that maximizes the information gain (or equivalently, the action that minimizes the expected entropy) at the next step, that is,

$$\pi^{\text{infotaxis}}(s) = \arg\min_a G(s,a) = \arg\min_a H(s|a). \tag{3.4}$$

Combining equation (3.2) with equation (3.4), the infotaxis policy can be written as

$$\pi^{\text{infotaxis}}(s) = \arg\min_a \sum_{s'} \Pr(s'|s,a)H(s'), \tag{3.5}$$

which makes clear that it shares the same structure as the optimal policy, given by equation (2.4) and which can be rewritten as

$$\pi^*(s) = 1 + \arg\min_a \sum_{s'} \Pr(s'|s,a)v^*(s'). \tag{3.6}$$

If $v^*(s)$, the expected time to find the source starting from belief state s with the optimal policy, was a linear function of $H(s)$, then infotaxis would be optimal. A trivial counter-example shows that this is not the case: when the source location is known, v^* is equal to the Manhattan distance to the source whereas $H = 0$ for all source locations.

(b) Performance of infotaxis

We now provide an assessment of the infotaxis policy. The performance of infotaxis has been evaluated in $n = \{1, 2, 3\}$ dimensions and for a wide range of problem sizes $\mathcal{L}$ and source intensities $\mathcal{I}$, defined by equation (2.16). A few additional simulations have also been performed in four dimensions. The results are summarized in figure 4. Missing data correspond to cases we cannot simulate, either because they require too much memory or because they are too expensive computationally. Videos showing infotactic searches are provided in the electronic supplementary material.

The first quantity of interest is the probability of failing to find the source, Pr(failure), which is plotted in the $(\mathcal{L}, \mathcal{I})$ plane in figure 4*a*. In our simulations, cases with Pr(failure) $\leq 10^{-6}$ have negligible probability of failure within our numerical accuracy (black dots). Conversely, cases with Pr(failure) $> 10^{-6}$ (grey to orange dots) exhibit a measurable probability of never finding the source, which is due to the agent entering an infinite loop at some point during the search (which then never terminates). It is clear from figure 4*a* that infotaxis does not always guarantee termination, however Pr(failure) $< 10^{-3}$ (black to grey shades) for almost all cases considered. The exceptions (orange shades) correspond to situations with small $\mathcal{L}$ and large $\mathcal{I}$: these are searches where hits provide so much information that the source location can be identified exactly from a distance. In this situation, the entropy of the belief state is zero and cannot be reduced any further: the infotactic agent is then 'lost'. This marginal issue could be easily fixed by requiring the agent to go to the source if its location is perfectly known.

The second quantity of interest is the mean number of steps to reach the source, provided that the source is found with sufficiently high probability (we set this threshold arbitrarily to 0.999). Mean(T) is plotted in the $(\mathcal{L}, \mathcal{I})$ plane using a colourmap in figure 4*b*, where we also show grey contours for the mean cumulated number of hits (excluding the initial hit). A more quantitative representation of Mean(T) is presented in figure 5, where we also show the theoretical lower and upper bounds we obtained for the optimal policy (equations (2.5) and (2.6), respectively).

The upper bound was obtained by considering a search where the agent explores the space exhaustively by following n-dimensional spiralling trajectories, or, in one dimension, by going to one end and then to the other end of the domain (unless the one-dimensional spiral is better, which is the case only for very small domains). This upper bound therefore scales as $\mathcal{L}^n$ (with n the dimensionality), as illustrated by figure 5*a*. Cases where Mean(T) is larger than this upper bound are highlighted in figure 4*b* by outer circles: the mean search time achieved by infotaxis is never above the upper bound for the optimal policy, except in one dimension for small $\mathcal{L}$ and $\mathcal{I}$ (a rather marginal configuration).

Besides, it has been shown that in the absence of detections (succession of zero hits), infotactic trajectories are Archimedean spirals in two dimensions [8] and an approximate generalization of those in three dimensions [24,26]. We recover such trajectories, they are shown in appendix C. However, in the one-dimensional case, we find that infotaxis does not generate one-dimensional

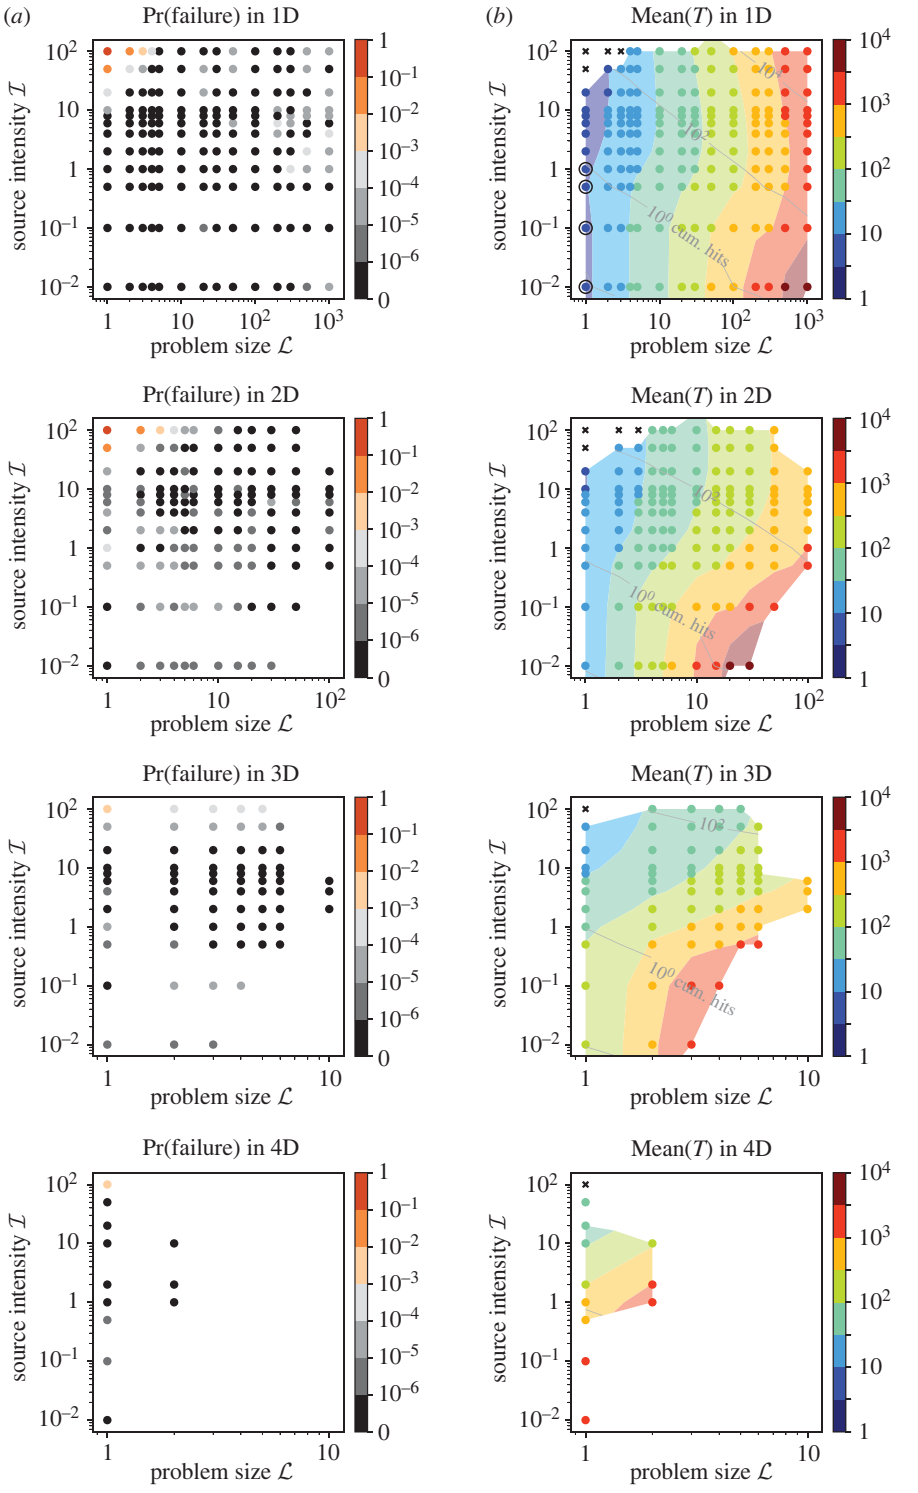

Figure 4. Performance of infotaxis for source tracking in one dimension, two dimensions, three dimensions and four dimensions (rows), for a wide range of problem sizes $\mathcal{L}$ and source intensities $\mathcal{I}$: (*a*) probability of never finding the source and (*b*) mean number of steps to find the source. In (*b*), the black crosses depict cases where $\Pr(\text{failure}) > 10^{-3}$, the grey contour lines indicate the mean cumulated number of hits gathered along the search (the detailed data are provided in the electronic supplementary material), and the outer circles depict cases where $\text{Mean}(T)$ is larger than the upper bound we obtained for the optimal policy. (Online version in colour.)

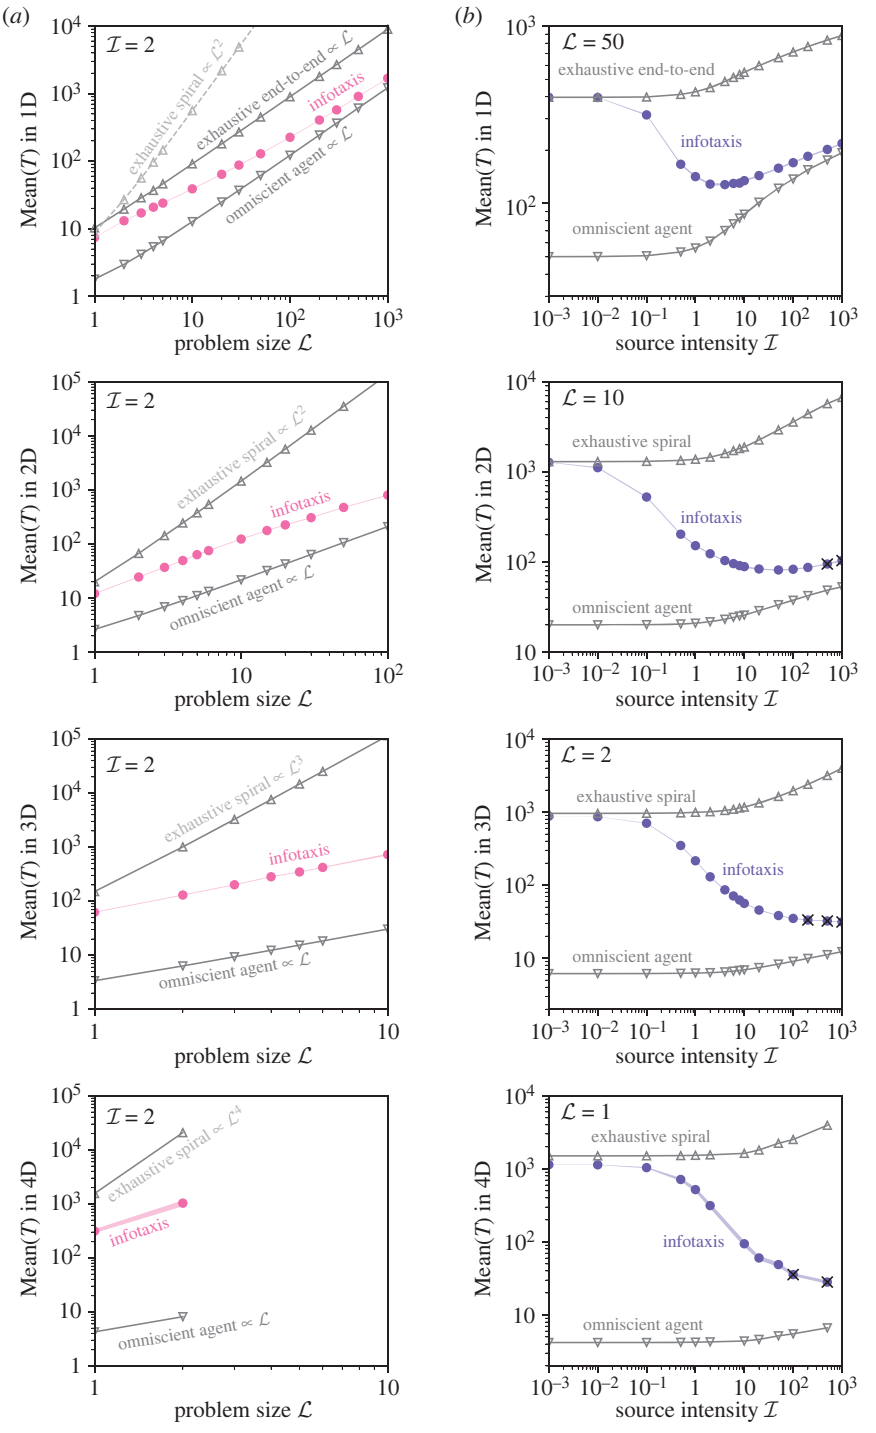

Figure 5. Mean number of steps to find the source with infotaxis in one dimension, two dimensions, three dimensions and four dimensions (rows) as a function of (*a*) the problem size $\mathcal{L}$ and (*b*) the source intensity $\mathcal{I}$. Infotaxis is depicted by dots (crossed dots indicate $\Pr(\text{failure}) > 10^{-3}$), while downward (resp. upward) triangles indicate the lower (resp. upper) bound for the optimal policy computed according to equation (2.5) (resp. equation (2.6)). The lower bound is strict and assumes an omniscient agent. The upper bound corresponds to an exhaustive search following an outward spiral, except in one dimension where an 'end-to-end' trajectory is better (unless the domain is very small). The (thin) shaded area between infotaxis dots show 95% confidence intervals. (Online version in colour.)

spirals but instead 'end-to-end' trajectories in the absence of hits. These trajectories are essentially the ones we assumed to derive our upper bounds: the latter are therefore a very good estimate of Mean(T) in the limit of no information $\mathcal{I} \to 0$, as can be seen from figure 5*b*.

The lower bound was obtained by considering an agent that knows the source location, and hence scales linearly in $\mathcal{L}$. We expect the optimal policy to approach this bound when $\mathcal{I}$ becomes large (though this bound is not tight). From figure 5*b*, this is clearly the case for the infotaxis policy. Note that at high $\mathcal{I}$, infotactic agents tend to get stuck (cf. figure 4*a*), which limits the maximal value of $\mathcal{I}$ we can reach.

Even though Mean(T) is the quantity to minimize in the source-tracking problem, a strategy exhibiting huge fluctuations in search times may be too hazardous to be used in practice. We found that the standard deviation of the arrival times distribution, Std(T), is of the same order of Mean(T) with very little variations across the entire parameter range, we explored: Std(T)/Mean(T) $\in [0.6, 1.1]$ in one dimension, $[0.7, 1.9]$ in two dimensions, $[1.0, 2.7]$ in three dimensions, and around 3 in four dimensions (though the parameter range is limited in this case). The detailed data are provided in the electronic supplementary material.

The distributions of arrival times for all considered parameters are presented in figure 6. Note that, compared to prior work, our methodology allowed us to properly sample the tail of the distributions up to $O(10)$ standard deviations. Our data (left column of figure 6) establish that the tail is subexponential, unlike what was previously believed [8,24,32,33]. By plotting the same data using a log-log scale, it can be shown that the tail decays faster than any power law. We found that the arrival time distributions approach lognormal distributions as the number of space dimensions increases. The lognormal probability density function reads

$$f(T) = \frac{1}{\sigma\tau\sqrt{2\pi}} \exp\left(-\frac{(\ln T - \mu)^2}{2\sigma^2}\right), \tag{3.7}$$

where $\mu = \text{Mean}(\ln T)$ and $\sigma = \text{Std}(\ln T)$. The distribution of $\tau = (\ln T - \mu)/\sigma$ is plotted in figure 6 (central column). If T was perfectly lognormally distributed, τ would follow a standard normal distribution (which is shown by the dashed curve). Lognormal distributions generally result from multiplicative random processes, by application of the central limit theorem in the log domain. Yet, lognormal distributions often appear as a good empirical fit for lifetime or reliability analyses, or more generally for distributions of a positive variable.

To summarize, we computed the statistics of infotactic searches in one dimension, two dimensions, three dimensions and four dimensions, for a wide range of dimensionless problem sizes $\mathcal{L}$ and cue emission intensities $\mathcal{I}$. Our data show that infotaxis is *reliable* (the probability of failure is less than 10^{-3}), *efficient* (the mean search time scales as expected for the optimal policy) and *safe* (the tail of the distribution of search times decays faster than any power law, though subexponentially).

(c) *N*-step infotaxis

Infotaxis is based on a one-step anticipation of possible outcomes of each action. In this section, we evaluate N-step infotaxis, the generalization of infotaxis to an anticipation over an arbitrary number of steps.

N-step infotaxis considers all possible outcomes of all possible sequences of N_{steps} actions, and maximizes the cumulated information gain over those steps. The detailed algorithm, which relies on an exhaustive tree search, is given in the electronic supplementary material. N-step infotaxis is of very limited practical interest, because its cost is exponential in N_{steps}. Yet it is fundamentally interesting to determine whether a significant increase of computational power leads to valuable improvements of the search.

The performance of N-steps infotaxis as a function of the number of anticipated steps is reported in figure 7 for searches in one dimension, two dimensions and three dimensions, for $\mathcal{L} = 1$, 5, or 10 and $\mathcal{I} = 2$. Results are quite variable depending on the case considered. In one dimension, improvements are significant for the smallest domain (40% reduction in the mean

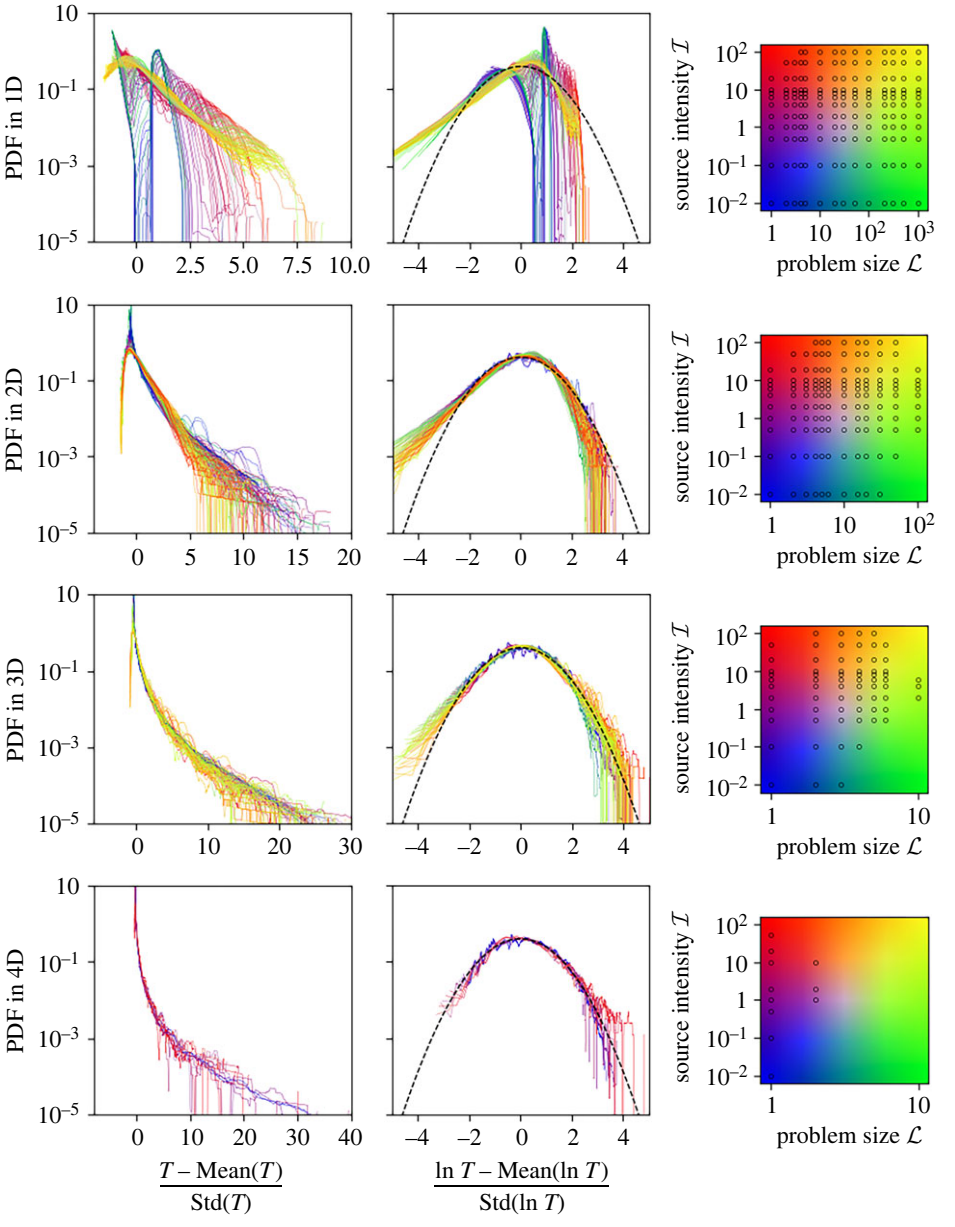

Figure 6. Distributions of arrival times with infotaxis for source tracking in one dimension, two dimensions, three dimensions and four dimensions (rows) for a wide range of problem sizes $\mathcal{L}$ and source intensities $\mathcal{I}$ (colour-coded according to the map on the right). The same data are plotted in both columns using different scales to evidence the distribution features. In the left column, exponential decay would appear as linear. In the right column, a standardized lognormal distribution would appear as a parabola (dashed curve), which corresponds to a standard Gaussian in the log-lin space. (Online version in colour.)

search time for $\mathcal{L} = 1$) while more modest for larger domains (15 and 7% for $\mathcal{L} = 5$ and 10, respectively, though a plateau is not reached yet). In two dimensions, the search time is reduced by 3% for the smallest domain ($\mathcal{L} = 1$), and by 11% for a larger one ($\mathcal{L} = 5$). The two-dimensional case with $\mathcal{L} = 1$ also shows an example where N-step infotaxis ($N_{\text{steps}} = 5$) performs worse than infotaxis. In three dimensions, improvements plateau at 3%, though we tested only the smallest domain size.

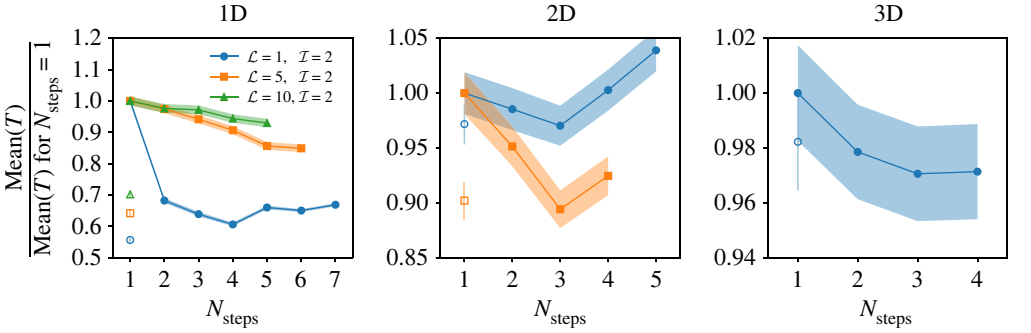

Figure 7. Performance of N-step infotaxis in one dimension, two dimensions and three dimensions (columns): mean number of steps to find the source as a function of the number of anticipated steps N_{steps}, normalized by its value for one-step infotaxis (filled symbols). The shaded areas represent 95% confidence intervals. For comparison, we also show the performance of our 'space-aware infotaxis' with open symbols (cf. §4). (Online version in colour.)

Overall, the search duration first decreases as N_{steps} increases, and then either reaches a plateau or increases again as N_{steps} is increased further. This non-monotonic behaviour is consistent with prior experimental results in two dimensions [22], and is further evidence of the absence of a linear correlation between the entropy and the remaining time to find the source. Importantly this minimum (or this plateau) does not generally correspond to the optimal performance: it is clear in one dimension that our own heuristic, 'space-aware infotaxis' (depicted by open symbols, and which will be presented in the next section) performs better.

4. A better heuristic: space-aware infotaxis

Infotaxis is inherently risk-averse, which explains why it is so reliable. But this behaviour is not optimal: it is indeed known that shifting infotaxis towards more exploitation using a tunable parameter can reduce search duration [23,24]. This is however not satisfactory, since the value of this parameter has to be optimized for each set of $(\mathcal{L},\mathcal{I})$. Here, we present a parameter-free policy that achieves better performance than infotaxis.

(a) Description of space-aware infotaxis

Entropy is a quantity that does not contain any information about the source location relative to agent. To make infotaxis more exploitative, we will build a policy based on entropy H, which measures uncertainty, and on a spatial metric, denoted D, which quantifies the distance between the agent and the source.

The quantity we want to minimize should naturally balance H and D, and be related to the time remaining to find the source. We chose the following expression:

$$J(s) = \log_2\left(D(s) + 2^{H(s)-1} - \frac{1}{2}\right), \tag{4.1}$$

and $J(s^\Omega) = 0$. It is a function of $H(s)$, the entropy of belief state s, defined in equation (3.1), and of $D(s)$, the mean Manhattan distance between the agent and the source

$$D(s) = \sum_{\boldsymbol{x}} p(\boldsymbol{x}) \|\boldsymbol{x} - \boldsymbol{x}^a\|_1, \tag{4.2}$$

where $p(\boldsymbol{x})$ is the probability of the source being in $\boldsymbol{x}$ and $\boldsymbol{x}^a$ is the agent location. This seemingly *ad hoc* expression of J has been constructed from the following hypotheses.

Consider a distribution $p(\boldsymbol{x})$ with entropy H. Assuming $p(\boldsymbol{x})$ is uniform, this corresponds to an effective number of cells $N_{\text{eff}} = 2^H$. The expected time to find a uniformly distributed source by

visiting exhaustively N_{eff} cells is

$$\mathbb{E}[T] = \sum_{i=0}^{N_{\text{eff}}-1} \frac{i}{N_{\text{eff}}} = \frac{N_{\text{eff}}-1}{2} = 2^{H-1} - \frac{1}{2}. \tag{4.3}$$

This expression does not take into account any spatial constraint and is derived assuming the agent can jump from cell to cell for a unit cost. On the other hand, if the source location is known, the minimal time needed to reach a source located at a distance D is D. Assuming the effects of distance and of uncertainty combine linearly, this gives the expression $D + 2^{H-1} - \frac{1}{2}$ as an estimate of the number of time steps remaining to find the source.

An additional consideration is the mathematical properties that we would like J to satisfy. For any belief-MDP with a cost to minimize, the optimal value function $v^*(s)$ of a belief state s, defined by equation (2.3), is a concave function of that belief state [38,39]. Since J is designed to be an approximation of the optimal value function v^*, we applied a logarithm function to ensure that J is a concave function of $p(\boldsymbol{x})$. The particular choice of the logarithm is motivated by the fact that infotaxis is recovered if one enforces $D = 0$. Note that the entropy H is indeed a concave function of $p(\boldsymbol{x})$.

We now introduce 'space-aware infotaxis' (SAI), the policy that minimizes the expected J at the next step. It reads

$$\pi^{\text{SAI}}(s) = \arg\min_a J(s|a), \tag{4.4}$$

where $J(s|a)$ is the expected J upon executing action a in belief state s, which is given by

$$J(s|a) = \sum_{s'} \Pr(s'|s,a) J(s') \tag{4.5}$$

where the sum is taken over all successor belief states s' (cf. figure 3). Note that SAI looks one step ahead into possible futures, the same way infotaxis does, therefore its computational cost is essentially the same as infotaxis (the only extra computation required is that of D).

(b) Performance of space-aware infotaxis

The performance of SAI is plotted in figure 8 as a function of the dimensionless problem size $\mathcal{L}$ and source intensity $\mathcal{I}$ (as defined by equation (2.16)). Video examples of its behaviour are provided in the electronic supplementary material. We show the probability of failure and the relative difference in the mean number of steps between SAI and infotaxis, defined such that a negative value signifies an improvement: for example a relative difference of -10% means that SAI finds the source in 10% less time than infotaxis, on average. Note that we also tested a version of SAI based on the Euclidean norm (rather than the Manhattan norm used here), and obtained marginally worse performance (see the electronic supplementary material).

In one dimension, SAI never fails and beats infotaxis everywhere by roughly 20–50% in the lower left quadrant of the $(\mathcal{L}, \mathcal{I})$ parameter space, while improvements are more modest in the upper right quadrant. This impressive gain is due to the markedly different behaviour of SAI in the absence of hits: while infotaxis tends to go to one end of the domain before turning back and going to the other end, SAI goes back and forth, exploring the domain further each time, as depicted in appendix C. This behaviour illustrates the more exploitative nature of SAI compared with infotaxis.

In two dimensions, SAI works only in the presence of cues ($\mathcal{I} \gtrsim 1$) but is able to reduce the time to find the source by roughly 5–15%. In three dimensions, again SAI works only provided that $\mathcal{I} \gtrsim 1$, and overall performs slightly better that infotaxis (roughly 2–3% improvement overall). We did not evaluate SAI in four dimensions due to the error bar on infotaxis data that prevents meaningful comparisons.

The fact that SAI is not reliable for small source intensities is not an important issue, since in this case the search is essentially performed without cues, therefore one could simply use the exhaustive spiral search (which is what infotaxis reduces to in this limit). The lack of reliability

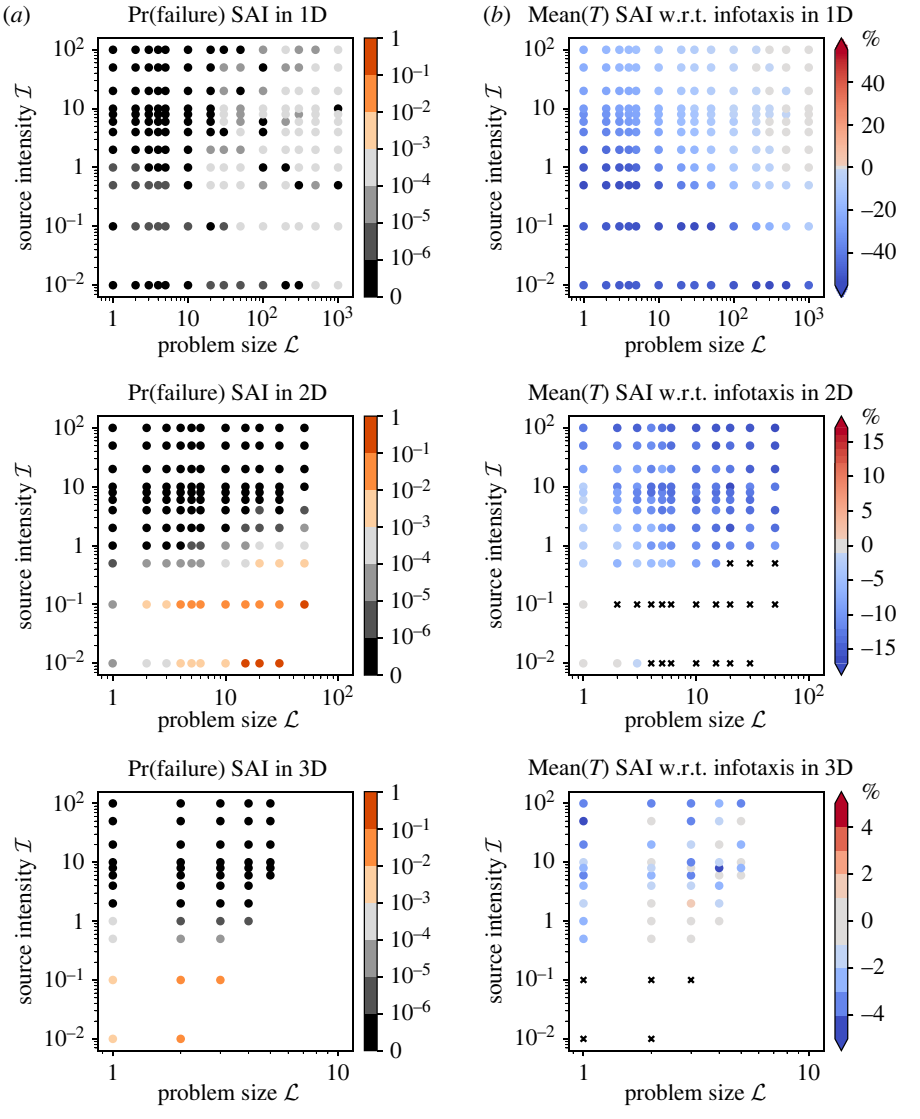

Figure 8. Performance of space-aware infotaxis for source tracking in one dimension, two dimensions and three dimensions (rows), for a wide range of problem sizes $\mathcal{L}$ and source intensities $\mathcal{I}$: (*a*) probability of never finding the source and (*b*) relative difference in the mean number of steps to find the source compared with infotaxis (−10% means that SAI finds the source in 10% less time than infotaxis). In (*b*), the black crosses depict cases where $\Pr(\text{failure}) > 10^{-3}$. (Online version in colour.)

of SAI in this case is due to the fact the SAI agent starts with a spiralling trajectory, but after a long period without hits, comes back to the centre of the domain and remains trapped there, as illustrated in appendix C.

Importantly, if the source location is known, SAI will direct the agent towards it (in an optimal manner). By contrast, an infotactic agent will be lost in this situation (the entropy being zero, it can not be reduced any further). This explains why SAI is reliable for the high source intensities where infotaxis was shown to fail (cf. figure 4*a*).

The distributions of arrival times generated by SAI and by infotaxis exhibit the same features. In particular, the standard deviation is in the order of the mean: for SAI we have $\text{Std}(T)/\text{Mean}(T) \in [0.6, 1.3]$ in one dimension, $[0.7, 1.6]$ in two dimensions and $[1.0, 2.3]$ in three dimensions; these

values are comparable with those obtained with infotaxis. The standardized distributions of arrival times are similar to those obtained with infotaxis (electronic supplementary material).

We now discuss the performance of space-aware infotaxis in the context of prior work. A variant of infotaxis, called entrotraxis, was shown to do better than infotaxis by maximizing the uncertainty on the next observation when the source strength is unknown [28]. We tested it, and found this is not the case in our set-up. Another study suggested to replace the entropy gain by the Bhattacharyya distance [32], but improvement, when existent, is marginal. Finally, the J functional of space-aware infotaxis ressembles the free energy functional proposed by Masson [23], but the free energy involves a parameter, the temperature, that has to be tuned while J is parameter-free. Besides, we tested our heuristic over the entire parameter space, from one dimension to three dimensions, whereas previous heuristics have only been compared with infotaxis in a limited number of cases.

5. Towards the optimal policy: deep reinforcement learning

In this section, we show how deep reinforcement learning can be applied to the source-tracking problem. Our goal here is not to perform a complete study of the capabilities of reinforcement learning to solve the source-tracking problem, which could easily be an entire study on its own. Instead, we wish to report on our attempt to solve the source-tracking POMDP using deep reinforcement learning in order to approach at best the optimal policy, as well as to provide a comparison with heuristic policies such as infotaxis and space-aware infotaxis.

(a) Reinforcement learning algorithm

Our reinforcement learning algorithm is based on DQN [40], an extension of traditional Q-learning to deep neural networks. It is adapted here to take advantage of the known model that determines the transitions between a belief state and all its possible successors.

We first recall that (truly) solving the source-tracking problem means finding the optimal policy, as explained in §2b. For that, one needs to compute the optimal value function $v^*(s)$, which is the function that satisfies equation (2.3) for all belief states s. If belief states could be tabulated, v^* could be obtained by classical dynamic programming using a value iteration algorithm [18,35]. Here, we must resort to an approximate method instead.

We approximate $v^*(s)$ by a parameterized function $\hat{v}(s;\boldsymbol{w})$ with weight vector $\boldsymbol{w}$. In practice here, $\hat{v}$ is the function computed by a multi-layer neural network, and $\boldsymbol{w}$ is the vector containing all 'synaptic' weights. The Bellman optimality equation for the approximate value function is

$$\hat{v}(s;\boldsymbol{w}^*) = \min_a \sum_{s'} \Pr(s'|s,a)[1+\hat{v}(s';\boldsymbol{w}^*)] \quad \forall s \neq s^{\Omega}, \tag{5.1}$$

with $\hat{v}(s^{\Omega};\boldsymbol{w}^*)=0$, and the problem becomes that of finding the weights $\boldsymbol{w}^*$ that allow equation (5.1) to be satisfied 'at best', that is, the weights that minimize the residual error. This residual error, called the Bellman optimality error, reads

$$L(\boldsymbol{w}) = \mathbb{E}_s\left[\min_a \sum_{s'} \Pr(s'|s,a)[1+\hat{v}(s';\boldsymbol{w})] - \hat{v}(s;\boldsymbol{w})\right]^2, \tag{5.2}$$

where the expectation is taken over belief states s visited when following the policy $\hat{\pi}$ derived from $\hat{v}$ and defined by

$$\hat{\pi}(s;\boldsymbol{w}) = \arg\min_a \sum_{s'} \Pr(s'|s,a)[1+\hat{v}(s';\boldsymbol{w})]. \tag{5.3}$$

The functional $L(\boldsymbol{w})$ is, in the language of deep neural networks, known as the 'loss function', and 'training' the network then refers to the iterative update (through stochastic gradient descent) of the weights $\boldsymbol{w}$ so as to minimize this loss function.

It is known that basic training algorithms (such as Q-learning) are unstable or even diverge when nonlinear function approximators are used to represent value functions. This major

issue was overcome with DQN, a reinforcement learning algorithm that assembled various stabilization techniques (experience replay, delayed target network) to facilitate convergence, and whose capabilities were demonstrated by achieving super-human performance on classic Atari video games such as Pong, Breakout and Space Invaders [40].

Most reinforcement learning algorithms, including Q-learning and DQN, are model-free, and hence rely on the action-value function (also known as the 'Q function'). The source-tracking problem is, however, model-based, because the probability of transitioning from a belief state s to a successor belief state s' is known exactly. This allows us to work directly with the value function $\hat{v}$, and to perform a full backup (that is, to compute the sum over s' in equations (5.1)–(5.3)) rather than a sample backup (based on a single successor belief state randomly sampled).

Our reinforcement learning algorithm is therefore identical to DQN [40], except that the network is trained to approximate the value function rather than the action-value function, and uses full backups rather than sample ones. Further technical information (neural network architecture, hyperparameters, etc.) is provided in the electronic supplementary material.

(b) Performance of (near) optimal policies

In order to obtain the best possible policy, we kept on increasing the size (width and depth) of the neural network until the performance of the learned policy stopped improving. This is a good indication (though not a proof) that the learned policy is optimal or very close to optimal. Due to the high cost of training and of ensuring near-optimality, we trained policies only for one- and two-dimension domains. Importantly, neural networks were trained from scratch (random initial weights) and all the results we present are very reproducible (e.g. for different random initializations). Several videos that illustrate the behaviour of these (near) optimal policies are available in the electronic supplementary material.

The trajectories obtained in the absence of hits ($\mathcal{I} \to 0$) are shown in appendix C. In one dimension, the (near) optimal trajectory is qualitatively similar (but not identical) to the one generated by space-aware infotaxis with a 'back and forth' motion extending further each time. In two dimensions, the (near) optimal trajectory is essentially an exhaustive spiral, as we obtained with infotaxis and space-aware infotaxis. Its geometry is however different, since it approximates a continuous Archimedean spiral defined with a Euclidean norm, rather than one defined with Chebyshev norm (as was the case with infotaxis and space-aware infotaxis). Repeated trainings produced slightly different trajectories, though all looked very similar to the one shown in appendix C. In continuous space, the Euclidean Archimedean spiral can be proven to be the optimal path for $\mathcal{I} \to 0$ using the fact that p_0 (the initial source probability distribution) is a monotonically decreasing function of the initial distance to the agent. Infotaxis has been shown to follow such a spiral in continuous space [26], hence the differences we observe here between (space-aware) infotaxis and the reinforcement learning policy are only due to our discrete setting.

The (near) optimal performance of learned policies is reported in figure 9 for searches in one dimension and two dimensions with varying source intensities $\mathcal{I}$, and, to some extent, varying problem sizes $\mathcal{L}$ (these dimensionless parameters are defined by equation (2.16)). Importantly, all these learned policies have Pr(failure) $< 10^{-6}$. For comparison, we also show the performance of infotaxis, N-step infotaxis and space-aware infotaxis on the same figure. In one dimension, the mean time to find the source can be reduced by roughly 20–50% compared with infotaxis. In two dimensions, this reduction is roughly around 5–10%, and up to 16%.

The most surprising result is maybe the remarkable performance of space-aware infotaxis: for all cases considered here, the performance attained with our heuristic is very close to the (near) optimal one obtained by deep reinforcement learning. It will be very interesting to evaluate to what extent this result generalizes to larger domain sizes and to three dimensions.

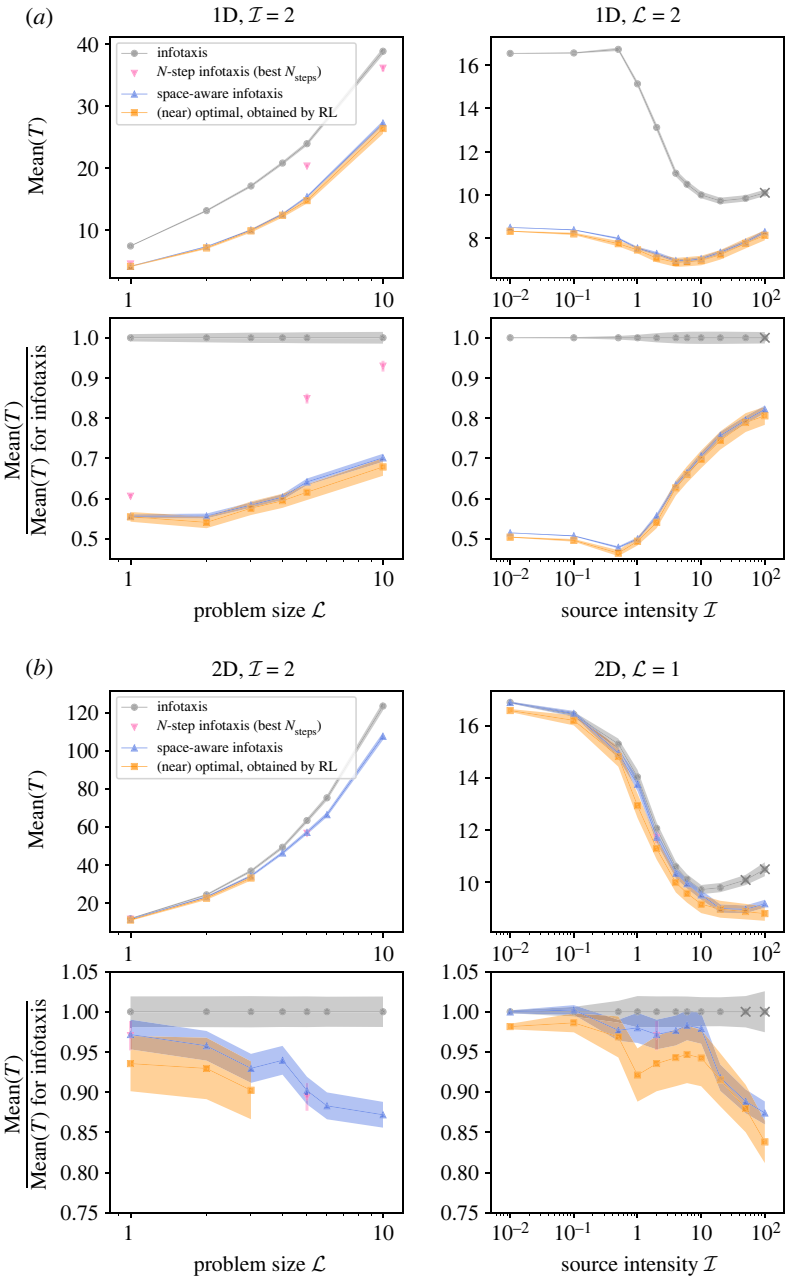

Figure 9. Performance of (near) optimal policies obtained by deep reinforcement learning (RL) in (*a*) one dimension and (*b*) two dimensions: mean number of steps to find the source as a function of the problem size $\mathcal{L}$ (left) and the source intensity $\mathcal{I}$ (right). All RL policies have $\Pr(\text{failure}) < 10^{-6}$. For comparison, we also show the performance of infotaxis, space-aware infotaxis and N-step infotaxis (when available). Crosses indicate that $\Pr(\text{failure}) > 10^{-3}$. Shaded areas show 95% confidence intervals. (Online version in colour.)

6. Discussion: Why does infotaxis work so well?

We have shown in this paper that while infotaxis is generally suboptimal, it performs remarkably well in a wide range of searching conditions. This is somewhat suprising, because the entropy of

a belief state is not linearly related to its optimal value (§3). In this section, we discuss possible explanations for this good performance.

A theoretical argument was proposed in [8], where the authors show that given a probability distribution $p(\boldsymbol{x})$ with entropy H, the expected optimal search time satisfies $\mathbb{E}_{\pi^*}[T|p_0 = p(\boldsymbol{x})] \geq 2^{H-1}$, and conclude that reducing entropy is a necessary condition to an efficient search. This bound is, according to the authors, a lower bound because it assumes the agent can jump to any cell for a unit cost. This result, however, is derived assuming a frozen distribution $p(\boldsymbol{x})$ and therefore only holds in the absence of cues gathered along the search ($\mathcal{I} \to 0$). This is the same assumption we actually used to derive our upper bound in equation (2.6) based on an exhaustive search, and also for our estimate of the remaining time for space-aware infotaxis in equation (4.3). It can be easily verified, from our data but also from fig. 2*c* in [8], that it is not a valid lower bound for a search with cues. Besides, based on this argument, one should seek to minimize the expectation of 2^{H-1} at the next step, rather than that of H (since these expressions are not linearly related, their expectations may rank actions differently). We actually tested this idea, and found that it does not perform as well as infotaxis.

We believe that the correct theoretical argument lies in the concavity of the optimal value function. This argument was provided in early artificial intelligence papers on POMDPs [19,41], where the idea of entropy minimization was proposed for robotic navigation. First recall that solving the source-tracking problem is equivalent to computing its optimal value function v^*. The optimal policy is then to choose the action that minimizes the expected optimal value at the next step. It can be mathematically proven that the optimal value function is a concave function of the belief state, in our case $p(\boldsymbol{x})$ [38,39]. The concavity of the optimal value function implies that the lowest values correspond to belief states located in the corners of the belief space, where $p(\boldsymbol{x}) = \delta(\boldsymbol{x} - \boldsymbol{x}')$: these are also the belief states with the lowest entropy.

Quoting Kaelbling *et al.* [41] for an intuitive explanation: 'The [concavity] of the optimal value function makes intuitive sense when we think about the value of belief states. States that are in the 'middle' of the belief space have high entropy—the agent is very uncertain about the real underlying state of the world. In such belief states, the agent cannot select actions very appropriately [$\cdots$]. In low-entropy belief states, which are near the corners of the simplex, the agent can take actions more likely to be appropriate for the current state of the world [$\cdots$]. This has some connection to the notion of 'value of information', where an agent can incur a cost to move it from a high-entropy to a low-entropy state; this is only worthwhile when the value of the information (the difference in value between the two states) exceeds the cost of gaining the information.'

As we have seen in this paper, beating infotaxis seems to be increasingly difficult as the dimensionality increases. While this could be inherent to the greater difficulty of finding good approximate solutions in higher dimensions, this result is well explained by the above argument: higher dimensionality implies larger uncertainty ($H \sim n \log N$, with n the dimension and N the linear domain size), meaning that optimal actions are those which help disambiguate the true source location (those which reduce entropy) for a larger fraction of the entire search.

7. Conclusion

In this paper, we demonstrated that infotaxis is generally suboptimal and can be beaten by other heuristics or by reinforcement learning, in particular, for searches in lower-dimensional spaces (and so even if an anticipation over multiple steps rather than one step is allowed). Yet, it remains a strong contender because of its generality: indeed, infotaxis performs well over a huge range of parameters without any tuning. More precisely, we have shown that (i) the probability of failing to find the source is negligible (infotaxis is *reliable*), (ii) the mean search time scales with physical parameters as one would expect for the optimal policy (infotaxis is *efficient*) and (iii) the tail of the distribution of search times decays faster than any power law, though subexponentially (infotaxis is not plagued by large fluctuations and hence is *safe*). Finally, we have shown that infotaxis can be made more efficient if the uncertainty measure (entropy) is balanced by a distance measure. We

called this parameter-free policy 'space-aware infotaxis'. Overall, this new heuristic reduces the mean time to find the source by 10–50% in one dimension, 5–15% in two dimensions and 2–3% in three dimensions.

To provide a firmer answer to the question of how good infotaxis is, one needs to compare it with the optimal policy. The nature and the size of the source-tracking problem do not allow the computation of exact solutions, however finding approximately optimal solutions is a task well suited for deep reinforcement learning. Our learning algorithm is a model-based version of DQN [40] where the value function is approximated by a deep neural network. We used large neural networks in order to approach at best the optimal policy. Because of the computational cost, we only trained policies in one dimension and two dimensions. Our results demonstrate that policies superior to infotaxis can be learned from scratch purely from experience. We found that, for the cases considered, the mean time to find the source can be reduced by 20–50% in one dimension and 5–15% in two dimensions compared with infotaxis. Almost identical improvements were obtained by our space-aware infotaxis in those cases. Overall, this strongly suggests that (i) while infotaxis is vastly suboptimal in one dimension, the margin of improvement towards the optimal policy gets tighter as the dimensionality increases and (ii) space-aware infotaxis is probably a very good approximation of the optimal policy.

The source-tracking problem was inspired by olfactory searches performed by moths, which use pheromones to locate their mates [2,3]. A similar navigation task is also faced by crustaceans looking for food on the sea floor [4,5]. It is not plausible that these animals perform infotaxis as presented here [42]: infotaxis involves complex computations (due to the one-step lookahead into possible futures) and relies on unrealistic assumptions (physical space representation, perfect memory, model knowledge). Developing heuristic approximations of infotaxis with much weaker requirements is possible [23], and it will be interesting to see whether such memory-based strategies can be discovered by minimal recurrent neural networks through reinforcement learning, as recently achieved for chemotaxis [43].

Reinforcement learning has recently emerged as a powerful tool for finding non-trivial solutions to complex navigation problems in turbulent flows [44–48], and provides a complementary approach to physics-based heuristics [49]. While all these navigation problems are POMDPs, they have never been explicitely framed as such and hence never exploited existing mathematical results and algorithms developed for this class of problems. In addition, learned policies are most often memoryless, while it is clear that navigation strategies shaped by Darwinian evolution rely on memory, either in a physical (molecular) form for single cells (e.g. bacteria performing chemotaxis [50,51]) or in an abstract one for higher organisms with cognitive capabilities (from the worm *C. elegans* to humans).

We believe that combining ideas from infotaxis (as a viable alternative to physical gradient climbing), Bayesian inference (for memory encoding) and POMDPs (which provide the mathematical framework) while leveraging the power of deep reinforcement learning is a promising route for uncovering search and navigation strategies used by living organisms and for adapting those to robotic applications. Through this paper, we made a first step in this direction by bringing together the relevant concepts and methods, and we hope to trigger further interest from all relevant communities on this highly interdisciplinary topic.

Data accessibility. This study has been performed with the open-source code OTTO, which is freely available on GitHub at http://github.com/C0PEP0D/otto. The data that support the findings of this study are openly available at https://doi.org/10.5281/zenodo.6125391 [52]. Additional data are provided in the electronic supplementary material [53].

Authors' contributions. A.L.: conceptualization, formal analysis, investigation, methodology, software, visualization, writing—original draft, writing—review and editing; C.E.: conceptualization, formal analysis, funding acquisition, investigation, methodology, software, writing—review and editing.

Both authors gave final approval for publication and agreed to be held accountable for the work performed therein.

Conflict of interest declaration. We declare we have no competing interests.

Funding. This project has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement no. 834238).
Acknowledgements. Centre de Calcul Intensif d'Aix-Marseille is acknowledged for granting access to its high-performance computing resources.

Appendix A. Initialization protocol

The physical model used in the source-tracking problem assumes an infinite, open domain. Because tracking can only start once the searcher knows that a source exists (that is, after a detection event has occurred), we propose to start the search with a non-zero hit and give a criteria to define a numerical domain 'large enough' to mimic an open domain.

To simplify the derivation, we will assume in the following a continuous space. Consider a spherical coordinates system (in n dimensions) centred on the agent. We denote $\Pr(r)\,\mathrm{d}r$ the probability of the source being in a spherical shell of radius r and thickness $\mathrm{d}r$. Before the search starts, we assume a uniform prior over the entire space and excluding the agent volume

$$\Pr(r) = \begin{cases} 0 & \text{for } r \in [0, a] \\ cS_n(r) & \text{for } r > a \end{cases}, \tag{A 1}$$

where a is the agent radius, c is a constant and $S_n(r)$ is the surface area of the n-ball of radius r

$$S_n(r) = \frac{2\pi^{n/2} r^{n-1}}{\Gamma(n/2)}, \tag{A 2}$$

with Γ the gamma function. After a initial hit h_{init}, using Bayes' rule, we can write

$$\Pr(r|h_{\text{init}}) = \frac{\Pr(h_{\text{init}}|r)\Pr(r)}{\Pr(h_{\text{init}})} \quad \text{for } r > a, \tag{A 3}$$

with the normalization constant given by

$$\Pr(h_{\text{init}}) = \int_a^\infty \Pr(h_{\text{init}}|r)\Pr(r)\,\mathrm{d}r. \tag{A 4}$$

For $h_{\text{init}} \neq 0$, $\Pr(r|h_{\text{init}})$ decays exponentially fast with r for large r. We can define a cut-off radius R_c such that the probability of the source being outside a ball of radius R_c is at most ε_{out} for any h_{init} (with ε_{out} small). This reads

$$R_c = \max_{h_{\text{init}} \neq 0} r_c(h_{\text{init}}), \tag{A 5}$$

where $r_c(h_{\text{init}})$ is implicitly defined by

$$\int_a^{r_c(h_{\text{init}})} \Pr(r|h_{\text{init}})\,\mathrm{d}r = 1 - \varepsilon_{\text{out}}. \tag{A 6}$$

The initialization procedure for the search is therefore as follows.

(i) The initial hit h_{init} is drawn at the beginning of each episode from the corresponding probability distribution (excluding zero hit)

$$\Pr(h_{\text{init}}|h_{\text{init}} \neq 0) = \begin{cases} 0 & \text{if } h_{\text{init}} = 0 \\ \dfrac{\Pr(h_{\text{init}})}{Z} & \text{otherwise} \end{cases}, \tag{A 7}$$

with Z the normalization constant such that $\sum_{h_{\text{init}}} \Pr(h_{\text{init}}|h_{\text{init}} \neq 0) = 1$.

(ii) The linear size of the grid, N, is set to

$$N = 2\,\mathrm{ceil}(R_c) + 1. \tag{A 8}$$

(iii) The source distribution is set to uniform: $p_{\text{init}}(\boldsymbol{x}) = 1/N^n$.

(iv) The agent's position $\boldsymbol{x}_0^a$ is set to the centre of the domain.

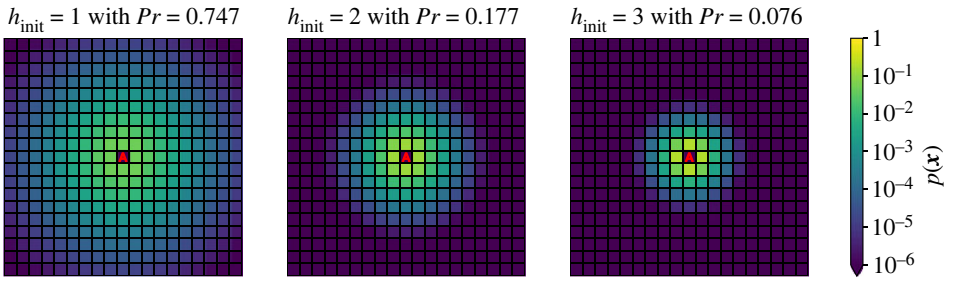

Figure 10. Initialization procedure: a set of initial beliefs p_0, here in two dimensions for $\mathcal{L} = 1$ and $\mathcal{I} = 2$, with the corresponding values of the initial hit and probabilities of occurring shown on top. 'A' indicates the agent's position, at the centre of the domain. (Online version in colour.)

(v) The initial source distribution p_0 is computed using $p_0 = \text{Bayes}(p_{\text{init}}(\boldsymbol{x}), \boldsymbol{x}_0^a, (\bar{F}, h_{\text{init}}))$, and the initial belief state is then $s_0 = [\boldsymbol{x}_0^a, p_0(\boldsymbol{x})]$.

Following this procedure, the set of initial beliefs p_0 is the set generated by $h_{\text{init}} = \{1, 2, \ldots\}$ and their probabilities of occurring are given by equation (A 7). An example of such a set is shown in figure 10. We used $\varepsilon_{\text{out}} = 10^{-3}$ in all our simulations. As a rule of thumb, this gives $N \approx 15\mathcal{L}$. More information on how N depends on the parameters is provided in the electronic supplementary material.

Appendix B. Policy evaluation

Policy evaluation is performed by generating a large number of episodes and computing the resulting distribution of arrival times T, denoted here as $f(T)$.

The convergence of $f(T)$ with the number of episodes can be vastly improved by realizing that $p(\boldsymbol{x})$, interpreted as the agent's belief in the context of decision-making, is also the true (in the Bayesian sense) probability distribution of sources that could have generated the sequence of observations. In this probabilistic approach, each episode can be continued until the probability of having found the source is equal to one (within numerical accuracy $\varepsilon_{\text{stop}} = 10^{-6}$) or until the agent is stuck in an infinite loop, and the hits are drawn at each step according to the distribution

$$\Pr(h|\boldsymbol{x}^a) = \sum_x \Pr(h|\boldsymbol{x}^a, \boldsymbol{x}) p(\boldsymbol{x}), \tag{B 1}$$

such that episodes can be generated independently of the true source location $\boldsymbol{x}^s$. We refer to this alternative framework as 'hybrid Bayesian/Monte Carlo' (as opposed to 'full Monte Carlo'). A video illustrating how the search proceeds in this framework, together with numerical proofs of its correctness and efficiency, are provided in the electronic supplementary material. This approach is particularly advantageous to sample rare events (such as failing to find the source) and more generally to sample heavy tailed distributions (as is $f(T)$).

The number of episodes is chosen such that the mean of the distribution is well-converged (95% confidence interval less than ±2%). Typically, we use at least 16 000 episodes in one dimension, 6400 episodes in two dimensions and 25 600 episodes in three dimensions. In four dimensions, we determined the mean with less accuracy due to the high cost of the simulations: we used typically 4096 episodes (yielding a 95% confidence interval less than ±10%) and down to 1024 episodes for the most expensive cases.

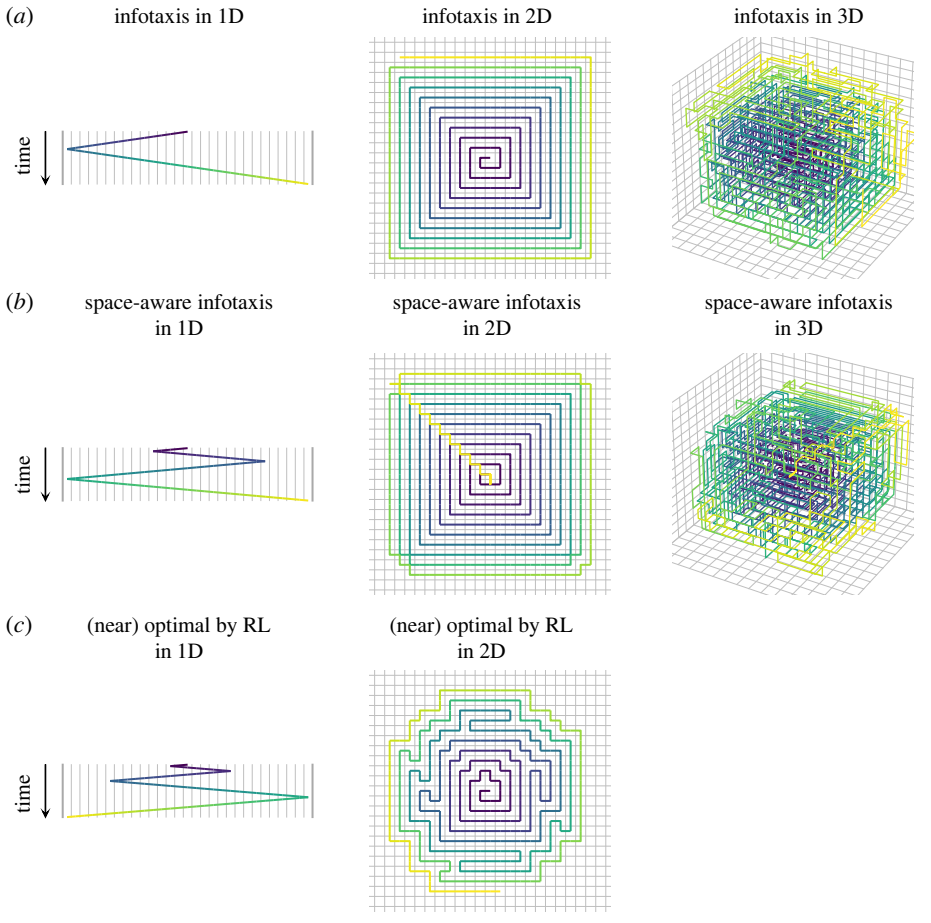

Figure 11. Trajectories in the absence of cues (succession of zero hits in the limit $\mathcal{I} \to 0$) for various policies: infotaxis (*a*), space-aware infotaxis (*b*), near-optimal by deep reinforcement learning (*c*). Time is colour-coded from blue (start) to yellow (end). The grid is shown in grey. (Online version in colour.)

Appendix C. Zero-hit trajectories

In the limit of a vanishing source intensity ($\mathcal{I} \to 0$), the search is performed without cues and the trajectories are deterministic, because they are generated by a succession of zero hits. These trajectories are shown in figure 11 for infotaxis, space-aware infotaxis and the near-optimal policy obtained by deep reinforcement learning. They have been calculated using $\mathcal{I} = 10^{-6}$ and $\mathcal{L} = 2$ in one dimension, $\mathcal{L} = 1.5$ in two dimensions and $\mathcal{L} = 1$ in three dimensions.

References

1. Vickers NJ. 2000 Mechanisms of animal navigation in odor plumes. *Biol. Bull.* **198**, 203–212. (doi:10.2307/1542524)
2. Hansson BS, ed. 1999 *Insect olfaction*. Berlin, Germany: Springer.
3. Murlis J, Elkinton JS, Carde RT. 1992 Odor plumes and how insects use them. *Annu. Rev. Entomol.* **37**, 505–532. (doi:10.1146/annurev.en.37.010192.002445)
4. Koehl MAR, Koseff JR, Crimaldi JP, McCay MG, Cooper T, Wiley MB, Moore PA. 2001 Lobster sniffing: antennule design and hydrodynamic filtering of information in an odor plume. *Science* **294**, 1948–1951. (doi:10.1126/science.1063724)

5. Reidenbach MA, Koehl MAR. 2011 The spatial and temporal patterns of odors sampled by lobsters and crabs in a turbulent plume. *J. Exp. Biol.* **214**, 3138–3153. (doi:10.1242/jeb.057547)
6. Kiørboe T. 2008 *A mechanistic approach to plankton ecology*. Princeton, NJ: Princeton University Press.
7. Reddy G, Murthy VN, Vergassola M. 2022 Olfactory sensing and navigation in turbulent environments. *Annu. Rev. Condens. Matter Phys.* **13**, 191–213. (doi:10.1146/annurev-conmatphys-031720-032754)
8. Vergassola M, Villermaux E, Shraiman BI. 2007 Infotaxis as a strategy for searching without gradients. *Nature* **445**, 406–409. (doi:10.1038/nature05464)
9. Celani A, Villermaux E, Vergassola M. 2014 Odor landscapes in turbulent environments. *Phys. Rev. X* **4**, 041015–041015. (doi:10.1103/PhysRevX.4.041015)
10. Patek SD. 2001 On partially observed stochastic shortest path problems. In *Proc. of the 40th IEEE Conf. on Decision and Control*, vol. 5, pp. 5050–5055. IEEE. (doi:10.1109/CDC.2001.981011)
11. Patek SD. 2007 Partially observed stochastic shortest path problems with approximate solution by neurodynamic programming. *IEEE Trans. Syst. Man Cybern.—Part A: Syst. Hum.* **37**, 710–720. (doi:10.1109/TSMCA.2007.902662)
12. Chatterjee K, Chmelik M, Gupta R, Kanodia A. 2016 Optimal cost almost-sure reachability in POMDPs. *Artif. Intell.* **234**, 26–48. (doi:10.1016/j.artint.2016.01.007)
13. Shani G, Pineau J, Kaplow R. 2013 A survey of point-based POMDP solvers. *Auton. Agents Multi-Agent Syst.* **27**, 1–51. (doi:10.1007/s10458-012-9200-2)
14. Kochenderfer MJ, Wheeler TA, Wray KH. 2022 *Algorithms for decision making*. Cambridge, MA: MIT Press.
15. Kurniawati H. 2022 Partially observable Markov decision processes and robotics. *Annu. Rev. Control Robot. Auton. Syst.* **5**, 253–277. (doi:10.1146/annurev-control-042920-092451)
16. Hauskrecht M. 2000 Value-function approximations for partially observable Markov decision processes. *J. Artif. Intell. Res.* **13**, 33–94. (doi:10.1613/jair.678)
17. Ross S, Pineau J, Paquet S, Chaib-draa B. 2008 Online planning algorithms for POMDPs. *J. Artif. Intell. Res.* **32**, 663–704. (doi:10.1613/jair.2567)
18. Sutton RS, Barto AG. 2018 *Reinforcement learning: an introduction*, 2nd edn. Cambridge, MA: MIT Press.
19. Cassandra AR, Kaelbling LP, Kurien JA. 1996 Acting under uncertainty: discrete Bayesian models for mobile-robot navigation. In *Proc. of IEEE/RSJ Int. Conf. on Intelligent Robots and Systems. IROS '96*, vol. 2, pp. 963–972.Osaka, Japan: IEEE.
20. Thrun S, Burgard W, Fox D. 2006 *Probabilistic robotics*. Cambridge, MA: MIT Press.
21. Moraud EM, Martinez D. 2010 Effectiveness and robustness of robot infotaxis for searching in dilute conditions. *Front. Neurorobotics* **4**, 1–8. (doi:10.3389/fnbot.2010.00001)
22. Lochmatter T. 2010 Bio-inspired and probabilistic algorithms for distributed odor source localization using mobile robots. PhD thesis, EPFL Lausanne.
23. Masson JB. 2013 Olfactory searches with limited space perception. *Proc. Natl Acad. Sci. USA* **110**, 11 261–11 266. (doi:10.1073/pnas.1221091110)
24. Masson JB, Bailly Bechet M, Vergassola M. 2009 Chasing information to search in random environments. *J. Phys. A: Math. Theor.* **42**, 434009. (doi:10.1088/1751-8113/42/43/434009)
25. Rodriguez JD, Gomez-Ullate D, Mejia-Monasterio C. 2014 Geometry-induced fluctuations of olfactory searches in bounded domains. *Phys. Rev. E* **89**, 042145. (doi:10.1103/PhysRevE.89.042145)
26. Barbieri C, Cocco S, Monasson R. 2011 On the trajectories and performance of infotaxis, an information-based greedy search algorithm. *EPL* **94**, 20005. (doi:10.1209/0295-5075/94/20005)
27. Karpas ED, Shklarsh A, Schneidman E. 2017 Information socialtaxis and efficient collective behavior emerging in groups of information-seeking agents. *Proc. Natl Acad. Sci. USA* **114**, 5589–5594. (doi:10.1073/pnas.1618055114)
28. Hutchinson M, Oh H, Chen WH. 2018 Entrotaxis as a strategy for autonomous search and source reconstruction in turbulent conditions. *Inf. Fusion* **42**, 179–189. (doi:10.1016/j.inffus.2017.10.009)
29. Chen C, Murphey TD, MacIver MA. 2020 Tuning movement for sensing in an uncertain world. *eLife* **9**, e52371. (doi:10.7554/eLife.52371)
30. Voges N, Chaffiol A, Lucas P, Martinez D. 2014 Reactive searching and infotaxis in odor source localization. *PLoS Comput. Biol.* **10**, e1003861. (doi:10.1371/journal.pcbi.1003861)

31. Calhoun AJ, Chalasani SH, Sharpee TO. 2014 Maximally informative foraging by *Caenorhabditis Elegans*. *eLife* **3**, e04220. (doi:10.7554/eLife.04220)
32. Ristic B, Skvortsov A, Gunatilaka A. 2016 A study of cognitive strategies for an autonomous search. *Inf. Fusion* **28**, 1–9. (doi:10.1016/j.inffus.2015.06.008)
33. Eggels AW, Kunnen RPJ, Koren B, Tijsseling AS. 2017 Infotaxis in a turbulent 3D channel flow. *J. Comput. Appl. Math.* **310**, 44–58. (doi:10.1016/j.cam.2016.06.004)
34. Rodriguez JD, Gomez-Ullate D, Mejia-Monasterio C. 2017 On the performance of blind-infotaxis under inaccurate modeling of the environment. *Eur. Phys. J. Spec. Top.* **226**, 2407–2420. (doi:10.1140/epjst/e2017-70067-1)
35. Bellman RE. 1957 *Dynamic programming*. Princeton, NJ: Princeton University Press.
36. Baeza-Yates RA, Culberson JC, Rawlins GJE. 1993 Searching in the plane. *Inf. Comput.* **106**, 234–252. (doi:10.1006/inco.1993.1054)
37. MacKay DJC. 2005 *Information theory, inference and learning algorithms*. Cambridge, UK: Cambridge University Press.
38. Sondik EJ. 1971 The optimal control of partially observable Markov decision processes. PhD thesis, Stanford University.
39. Smallwood RD, Sondik EJ. 1973 The optimal control of partially observable Markov processes over a finite horizon. *Oper. Res.* **21**, 1071–1088. (doi:10.1287/opre.21.5.1071)
40. Mnih V*et al.* 2015 Human-level control through deep reinforcement learning. *Nature* **518**, 529–533. (doi:10.1038/nature14236)
41. Kaelbling LP, Littman ML, Cassandra AR. 1998 Planning and acting in partially observable stochastic domains. *Artif. Intell.* **101**, 99–134. (doi:10.1016/S0004-3702(98)00023-X)
42. Cardé RT. 2021 Navigation along windborne plumes of pheromone and resource-linked odors. *Annu. Rev. Entomol.* **66**, 317–336. (doi:10.1146/annurev-ento-011019-024932)
43. Hartl B, Hübl M, Kahl G, Zöttl A. 2021 Microswimmers learning chemotaxis with genetic algorithms. *Proc. Natl Acad. Sci. USA* **118**, e2019683118. (doi:10.1073/pnas.2019683118)
44. Reddy G, Celani A, Sejnowski TJ, Vergassola M. 2016 Learning to soar in turbulent environments. *Proc. Natl Acad. Sci. USA* **113**, E4877–E4884. (doi:10.1073/iti1816113)
45. Reddy G, Wong-Ng J, Celani A, Sejnowski TJ, Vergassola M. 2018 Glider soaring via reinforcement learning in the field. *Nature* **562**, 236–239. (doi:10.1038/s41586-018-0533-0)
46. Colabrese S, Gustavsson K, Celani A, Biferale L. 2017 Flow navigation by smart microswimmers via reinforcement learning. *Phys. Rev. Lett.* **118**, 158004. (doi:10.1103/PhysRevLett.118.158004)
47. Gustavsson K, Biferale L, Celani A, Colabrese S. 2017 Finding efficient swimming strategies in a three-dimensional chaotic flow by reinforcement learning. *Eur. Phys. J. E* **40**, 110. (doi:10.1140/epje/i2017-11602-9)
48. Alageshan JK, Verma AK, Bec J, Pandit R. 2020 Machine learning strategies for path-planning microswimmers in turbulent flows. *Phys. Rev. E* **101**, 043110. (doi:10.1103/PhysRevE.101.043110)
49. Monthiller R, Loisy A, Koehl MAR, Favier B, Eloy C. 2021 Surfing on turbulence. (https://arxiv.org/abs/2110.10409)
50. Macnab RM, Koshland DE. 1972 The gradient-sensing mechanism in bacterial chemotaxis. *Proc. Natl Acad. Sci. USA* **69**, 2509–2512. (doi:10.1073/pnas.69.9.2509)
51. Lan G, Tu Y. 2016 Information processing in bacteria: memory, computation, and statistical physics: a key issues review. *Rep. Prog. Phys.* **79**, 052601. (doi:10.1088/0034-4885/79/5/052601)
52. Loisy A, Eloy C. 2022 Data from: Searching for a source without gradients: how good is infotaxis and how to beat it. *Zenodo*. (doi:10.5281/zenodo.6125392)
53. Loisy A, Eloy C. 2022 Searching for a source without gradients: how good is infotaxis and how to beat it. FigShare. (https://doi.org/10.6084/m9.figshare.c.6043560)