Provably-Convergent Bayesian Source Seeking with Mobile Agents in Multimodal Fields

Vivek Mishra VMISHRA9@GATECH.EDU
Department of Electrical and Computer Engineering
Georgia Institute of Technology
Atlanta, GA, USA

Raul Astudillo RASTUDIL@CALTECH.EDU
Department of Computing and Mathematical Sciences
California Institute of Technology
Pasadena, CA, USA

Peter Frazier PF98@CORNELL.EDU
School of Operations Research and Information Engineering
Cornell University
Ithaca, NY, USA

Fumin Zhang FUMIN@GATECH.EDU
Department of Electrical and Computer Engineering
Georgia Institute of Technology
Atlanta, GA, USA

Abstract

We consider source-seeking tasks, where the goal is to locate a source using a mobile agent that gathers potentially noisy measurements from the emitted signal. Such tasks are prevalent, for example, when searching radioactive or chemical sources using mobile sensors that track wind-carried particles. In this work, we propose an iterative Bayesian algorithm for source seeking, especially well-suited for challenging environments characterized by multimodal signal intensity and noisy observations. At each step, this algorithm computes a Bayesian posterior distribution characterizing the source's location using prior physical knowledge of the observation process and the accumulated data. Subsequently, it decides where the agent should move and observe next by following a search strategy that implicitly considers paths to the source's most likely location under the posterior. We show that the trajectory of an agent executing the proposed algorithm converges to the source's location asymptotically with probability one. We validate the algorithm's convergence through simulated experiments of an agent seeking a chemical plume in a turbulent environment.

Keywords: Source Seeking, Target Localization, Sensor Data Acquisition, Mobile Robots

1. Introduction

Source-seeking tasks involve locating a source using a mobile agent that collects potentially noisy measurements of the emitted signal. Such scenarios arise, for example, when identifying radioactive or chemical sources with sensors that trace wind-borne particles. Autonomous source-seeking robots become indispensable in hazardous environmental con-

ditions or when communication is limited. Their effective deployment can bolster search and rescue operations, localize chemical or nuclear threats, and enhance surveillance.

While existing source-seeking methods perform well under unimodal signal intensities and non-zero measurements —often leveraging stochastic gradient ascent (Zhang et al., 2007; Cochran and Krstic, 2009; Liu and Krstic, 2010; Atanasov et al., 2012; Ramírez-Llanos and Martínez, 2016; Azuma et al., 2012; Oliveira et al., 2014) or strategies quantifying local intensity fluctuations to adjust the agent's trajectory (Matveev et al., 2011; Mellucci et al., 2016; Ranó et al., 2017; Gillespie et al., 2017)—they struggle when these conditions are not met. Multimodal signal intensities and zero measurements can occur, for instance, when detecting a radioactive or chemical source masked by background signals or when the source particles traverse turbulent mediums, like wind or ocean currents. Although there are algorithms tailored for turbulent fields (see, e.g., Chang et al. 2013; Tian et al. 2015; Charrow et al. 2014), most are dependent on non-zero measurements for path determination. This becomes problematic, particularly in scenarios where maintaining the agent's position over extended durations is challenging. Additional related work is discussed in Appendix A.

In this study, we introduce a novel source-seeking algorithm, termed the *max probability algorithm*. This algorithm is especially well-suited for stochastic observations (including zero measurements) of signals whose mean intensity field may be multimodal. Our algorithm is developed under a Bayesian framework, similar to the one pursued by Mishra and Zhang (2016). In our empirical evaluation, we show that our algorithm outperforms the *expected rate algorithm* proposed by Mishra and Zhang (2016). Moreover, in contrast with this prior work, we prove that our algorithm is asymptotically consistent under sufficient conditions. Our main assumption can be stated roughly as follows: Given an infinite set of observations at a specific location, the presence or absence of the source at said location can be determined. This condition is met in our applications of interest (detecting a radioactive or chemical source) when the observation intensity at the source's location is above a known threshold, while observation intensities at other locations are below this threshold.

Our algorithm operates in an iterative manner. At each iteration, it builds a Bayesian posterior distribution characterizing the source's location using prior physical knowledge of the observation process and the data collected so far. Subsequently, this posterior distribution informs a search strategy that selects the agent's next intended location.

Similar iterative strategies have been proposed in the literature (Vergassola et al., 2007b; Hajieghrary et al., 2015; Mishra and Zhang, 2016). However, to our knowledge, convergence guarantees for any of these algorithms have not been reported. In contrast, we show that the trajectory of an agent guided by our algorithm converges to the source's location with probability one. Given that our search strategy does not explicitly rely on the field's gradients and operates within potentially multimodal intensity fields, establishing this convergence is non-trivial. Our proof relies on a martingale argument and a mechanism guaranteeing that the agent will not change its direction too often.

Our contributions can be summarized as follows:

- We propose a novel Bayesian source-seeking algorithm that is especially well-suited for stochastic observations of a potentially multimodal mean intensity field.
- We prove that the trajectory of an agent guided by our algorithm converges to the desired source's location asymptotically with probability one.

- We conduct numerical experiments demonstrating the superior performance of our algorithm over standard benchmarks from the literature.

The rest of the paper is organized as follows. In Section 2, we provide a mathematical formulation of the source-seeking problem. The max probability algorithm is presented in Section 3. We provide simulation results in Section 4. Finally, we conclude while offering directions for future research in Section 5.

2. Problem formulation

Consider a search space represented by a finite set $\mathbb{X} \subset \mathbb{R}^d$. While in real-world scenarios the search space might be continuous, our algorithm can be applied by adopting a suitable discretization. At each step, an agent records a noisy measurement $y \in \mathbb{R}^l$. The distribution of this measurement is characterized by the agent's present location $x \in \mathbb{X}$ and an environmental state $\theta \in \Theta$, where Θ denotes all potential states. More specifically, the likelihood of observing y at a location x is $y \mid x, \theta \sim L(\,\cdot\,; R(x;\theta))$, wherein $\{L(\,\cdot\,;\lambda)\}_{\lambda\in\Lambda}$ is a known family of densities or probability mass functions parameterized by $\lambda \in \Lambda$ and $\{R(\,\cdot\,;\theta)\}_{\theta\in\Theta}$ is a known family of functions parameterized by $\theta \in \Theta$. We assume that each observation is conditionally independent of past observations given x and θ. Finally, we adopt a Bayesian framework and assume that θ is drawn from a prior distribution p with support Θ.

The goal is to find a location $x \in \mathbb{X}$ such that $R(x;\theta) \in S^*$ for some fixed known set $S^* \subset \Lambda$. As we discuss below, this choice is motivated by chemical plume search applications, where the source-seeking task can be considered solved whenever we find a location whose signal intensity is above a known threshold ν, in which case $S^* = [\nu, \infty)$. Let $X^*(\theta) = \{x \in \mathbb{X} : R(x;\theta) \in S^*\}$. We refer to $X^*(\theta)$ as the target set. Thus, our goal is to find a point in the target set. Formally, if x_n is the agent's location at iteration n, we wish to develop a search strategy such that $\mathbf{P}(\lim_{n\to\infty} x_n \in X^*(\theta)) = 1$, where $\mathbf{P}$ denotes the probability operator.

3. The max probability algorithm

Our algorithm has two main components: (1) the posterior distribution over θ induced by the prior and the measurements collected so far, and (2) the search strategy, which depends on the posterior distribution and dictates the next location for the agent to visit. Let $x_1, \ldots, x_n \in \mathbb{X}$ be the locations visited by the agent up to time n, and let $y_1, \ldots, y_n \in \mathbb{R}^l$ be the corresponding measurements at these locations. We denote the posterior distribution of θ given the measurements collected up to time n by p_n. The computation of the posterior distribution is standard, so we defer its derivation to Appendix C.

To support the description of our algorithm, we introduce the following notation. For any given $x \in \mathbb{X}$, we denote the set of locations that are accessible from x in a single time step by $\mathcal{I}(x)$. In addition, for any given $x, x' \in \mathbb{X}$, we let $\rho(x, x')$ be the length of the shortest path between x and x', accounting for barriers. We assume that $\rho(x, x') < \infty$ for all $x, x' \in \mathbb{X}$, i.e., that each location is reachable from each other location. Finally, we denote by $q_n(x)$ to the posterior probability of x being in the target set, i.e., $q_n(x) = \mathbf{P}_n(x \in X^*(\theta))$, where $\mathbf{P}_n$ denotes the conditional probability given the measurements collected up to time n. We note that q_n can be computed from p_n.

We are now in a position to formally describe our algorithm, which we term the "max probability algorithm". This algorithm has a single hyperparameter, $\alpha \in (0, 1)$. Importantly, any value of α within this range ensures convergence. Let $\widehat{x}_0 \in \operatorname{argmax}_{x \in \mathbb{X}} q_0(x)$ and $k_0 = 0$. We define the sequences $\{\widehat{x}_n\}_{n=1}^{\infty}$ and $\{k_n\}_{n=1}^{\infty}$ recursively by $\widehat{x}_n \in \operatorname{argmax}_{x \in \mathbb{X}} q_n(x)$ and $k_n = n$, if $q_n(\widehat{x}_{n-1}) < \alpha q_{k_{n-1}}(\widehat{x}_{n-1})$, and $\widehat{x}_n = \widehat{x}_{n-1}$ and $k_n = k_{n-1}$, otherwise. Let x_0 be the agent's initial location. At every time step, our algorithm chooses the next location to visit as $x_{n+1} \in \operatorname{argmin}_{x \in \mathcal{I}(x_n)} \rho(x, \widehat{x}_n)$. In general problems, x_{n+1} can be computed using Dijkstra's algorithm. Algorithm 1 summarizes this description. In Appendix D, we prove that this algorithm is asymptotically consistent, i.e., it finds the source with probability one as the number of iterations goes to infinity.

Algorithm 1 Max probability algorithm

Require: Prior p_0, and agent's initial location x_0.
  Compute q_0 from p_0
  Set $k = 0$ and $\widehat{x} \in \operatorname{argmax}_{x \in \mathbb{X}} q_0(x)$.
  for $n = 0, 1, \ldots$ **do**
    if $q_n(\widehat{x}) < \alpha\, q_k(\widehat{x})$ **then**
      Set $k = n$ and $\widehat{x} \in \operatorname{argmax}_{x \in \mathbb{X}} q_n(x)$
    end if
    Send agent to $x_{n+1} \in \operatorname{argmin}_{x \in \mathcal{I}(x_n)} \rho(x, \widehat{x})$
    Observe measurement y_{n+1} at location x_{n+1}
    Compute posterior distribution p_{n+1}
    Compute q_{n+1} from p_{n+1}
  end for

4. Simulation results

We present simulation results for two instances of the 2-dimensional chemical plume search problem described in Appendix B. An experiment with a single source is discussed in Section 4.1. An additional experiment with two sources is discussed in Appendix E. In both cases, we compare the performance of the max probability, expected rate (Mishra and Zhang, 2016), and infotaxis (Vergassola et al., 2007b) algorithms.

Let θ_i be the location of i-th source and $r_i(x; \theta_i)$ be the rate at which plume particles from this source are detected when the agent is at location x. Based on Vergassola et al. (2007a), we model $r_i(x; \theta_i)$ using the advection-diffusion differential equation, i.e., $r_i(x; \theta_i) = f_i(x - \theta_i)$, where

$$f_i(z) = \frac{s_i}{\log \frac{\lambda_i}{a}} e^{-\frac{\langle z, V \rangle}{2D}} K_0 \left(\frac{\max\{b, ||z||_2\}}{\lambda_i} \right),$$

s_i is the rate at which the plume source releases the plume particles in the environment, a is the size of the sensor detecting plume particles, V is the average wind velocity, D is the diffusivity of the plume particles, K_0 is the modified Bessel function of order zero, and λ_i is the average distance traveled by particle in its lifetime τ_i, which satisfies

$$\lambda_i = \sqrt{\frac{\tau_i D}{1 + \frac{\tau_i ||V||_2^2}{4D}}}.$$

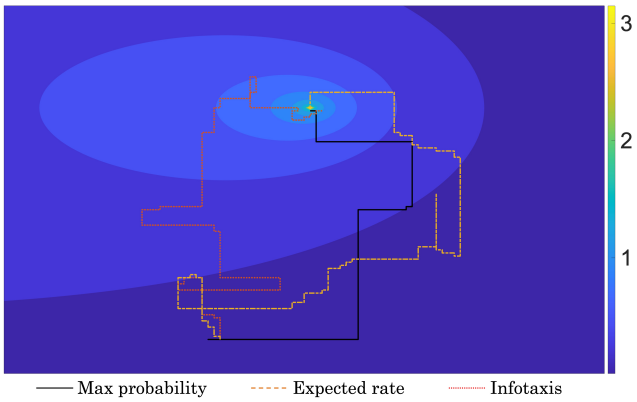

Figure 1: Sample paths of the max probability (black), expected rate (red), and infotaxis (yellow) algorithms for the unimodal source-seeking problem described in Section 4.1. The source's location and the agent's initial location are the same for the three algorithms.

The model presented in Vergassola et al. (2007a) assumes that $b = 0$, though this creates the possibility of an infinite rate when the $x = \theta_i$. To avoid this, we set $b = 10^{-308}$ in our simulations, which guarantees a finite rate of detection in our MATLAB implementation.

4.1 Unimodal case - single source in the search area

In this problem, we consider a single source in a $4m \times 4m$ 2-dimensional search space discretized into a 200×200 uniform grid. Since there is a single source, $\theta = \theta_1$ and $R(x;\theta) = r_1(x;\theta_1)$. The parameters of f_1 are as follows. The plume source is emitting plumes at a rate of $s_1 = 1\,\text{plumes}/s$. The plume particles have a diffusivity of $D = 1\,m^2/s$, and the average travel distance of plume particles is $\lambda_1 = 1.9\,m$. To focus on settings where the search is the most challenging, we choose a low magnitude for the velocity $V = [-1, 0]^\top\ m/s$. A robot of size $0.1\,m^2$ is deployed in the search space to find the source of the plume particles. We take $S^* = [\nu, \infty)$, where $\nu = f_1(0)$ is the intensity at the source, which ensures that $X^*(\theta) = \{\theta\}$. Figure 1 depicts the paths of three agents executing the max probability, expected rate, and infotaxis algorithms to localize a plume source. The max probability algorithm uses the lowest number of steps to find the source compared to the expected rate and infotaxis algorithms.

To understand the behavior of the algorithms under different configurations of the source's location and agent's initial location, we perform 3000 simulation runs, each with a configuration chosen uniformly at random over the search space. For each run, the number of steps to find the source taken by each algorithm is normalized by the length of the shortest path between the initial location of the agent and the source's location. The results are shown in Figure 2. As shown, the max probability algorithm outperforms its competitors.

Finally, we investigate the robustness of the max probability algorithm with respect to its only hyperparameter $\alpha \in (0, 1)$. To do so, we perform 3000 simulation runs using each of the values $\alpha \in \{0.05, 0.1, 0.3, 0.5, 0.7, 0.9, 0.95, 0.99\}$. The results are depicted in Figure 3, which shows that α has very little effect on the number of steps taken.

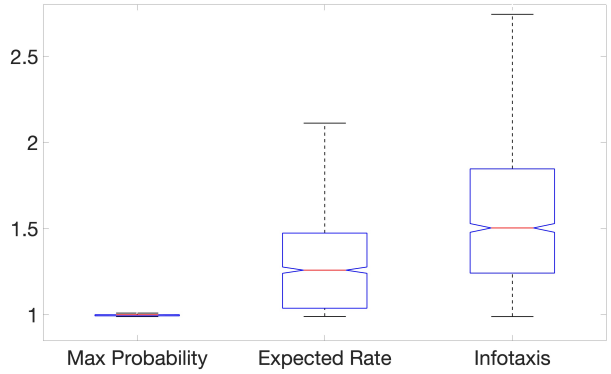

Figure 2: Box plot depicting the normalized number of steps taken by the search algorithms to find the source in the unimodal source-seeking problem described in Section 4.1 across 3000 simulation runs. The red line at the center of each box represents the median. The lower and upper horizontal edges of the box represent the 25^{th} and 75^{th} percentiles, respectively. Finally, the black horizontal lines at the lower and upper ends of the whiskers of each box represent the minimum and the maximum values, respectively.

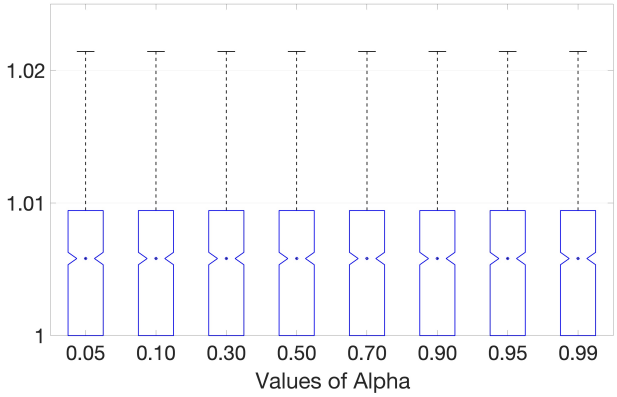

Figure 3: Box plot depicting the normalized number of steps taken by the max probability algorithm under different values of α.

5. Conclusion

This work introduced the max probability source-seeking algorithm. This algorithm maintains a posterior distribution characterizing the source's location, which is used iteratively to decide the location the agent should visit next. Our algorithm and two baselines from the literature were evaluated in two simulated chemical plume source-seeking tasks, showing that the former converges faster and more reliably to the source. Finally, we showed that our algorithm is asymptotically consistent. A promising direction for future research is to extend our algorithm and analysis to more complex settings, such as problems with multiple agents and moving sources.

Acknowledgments

This work is supported by AFOSR grant FA9550-19-1-0283; ONR grants N00014-19-1-2556 and N00014-19-1-2266; NSF grants OCE-1559475, CNS-1828678, and S&AS-1849228; NRL grants N00173-17-1-G001 and N00173-19-P-1412 ; and NOAA grant NA16NOS0120028.

References

The theory of search. II. Target detection. *Operations research*, 4(5):503–531, 1956.

Nikolay Atanasov, Jerome Le Ny, Nathan Michael, and George J Pappas. Stochastic source seeking in complex environments. In *Robotics and Automation (ICRA), 2012 IEEE International Conference on*, pages 3013–3018. IEEE, 2012.

Shun-ichi Azuma, Mahmut Selman Sakar, and George J Pappas. Stochastic source seeking by mobile robots. *IEEE Transactions on Automatic Control*, 57(9):2308–2321, 2012.

O Bénichou, M Coppey, M Moreau, PH Suet, and R Voituriez. Optimal search strategies for hidden targets. *Physical review letters*, 94(19):198101, 2005.

Dongsik Chang, Wencen Wu, Donald R Webster, Marc J Weissburg, and Fumin Zhang. A bio-inspired plume tracking algorithm for mobile sensing swarms in turbulent flow. In *Robotics and Automation (ICRA), 2013 IEEE International Conference on*, pages 921–926. IEEE, 2013.

Benjamin Charrow, Nathan Michael, and Vijay Kumar. Cooperative multi-robot estimation and control for radio source localization. *The International Journal of Robotics Research*, 33(4):569–580, 2014.

Erhan Çinlar. *Probability and stochastics*, volume 261. Springer.

Jennie Cochran and Miroslav Krstic. Nonholonomic source seeking with tuning of angular velocity. *IEEE Transactions on Automatic Control*, 54(4):717–731, 2009.

Peter I Frazier. Bayesian optimization. In *Recent Advances in Optimization and Modeling of Contemporary Problems*, pages 255–278. INFORMS, 2018.

James Gillespie, Inaki Ranó, Nazmul Siddique, José Santos, and Mehdi Khamassi. Reinforcement learning for bio-inspired target seeking. In *Conference Towards Autonomous Robotic Systems*, pages 637–650. Springer, 2017.

Hadi Hajieghrary, Alex Fabregat Tomás, and M Ani Hsieh. An information theoretic source seeking strategy for plume tracking in 3d turbulent fields. In *Safety, Security, and Rescue Robotics (SSRR), 2015 IEEE International Symposium on*, pages 1–8. IEEE, 2015.

Bernard O Koopman. The theory of search. III. The optimum distribution of searching effort. *Operations research*, 5(5):613–626, 1957.

Bernard Osgood Koopman. The theory of search. I. Kinematic bases. *Operations research*, 4(3):324–346, 1956.

Andreas Krause and Carlos Guestrin. Near-optimal observation selection using submodular functions. In *AAAI*, volume 7, pages 1650–1654, 2007.

Harold J Kushner. A new method of locating the maximum point of an arbitrary multipeak curve in the presence of noise. *Journal of Basic Engineering*, 86(1):97–106, 1964.

Shu-Jun Liu and Miroslav Krstic. Stochastic source seeking for nonholonomic unicycle. *Automatica*, 46(9):1443–1453, 2010.

Alexey S Matveev, Hamid Teimoori, and Andrey V Savkin. Navigation of a unicycle-like mobile robot for environmental extremum seeking. *Automatica*, 47(1):85–91, 2011.

Chiara Mellucci, Prathyush P Menon, Christopher Edwards, and Peter Challenor. Source seeking using a single autonomous vehicle. In *American Control Conference (ACC), 2016*, pages 6441–6446. IEEE, 2016.

Vivek Mishra and Fumin Zhang. A stochastic optimization framework for source seeking with infotaxis-like algorithms. In *2016 IEEE 55th Conference on Decision and Control (CDC)*, pages 6845–6850. IEEE, 2016.

Jonas Močkus. On Bayesian methods for seeking the extremum. In *Optimization Techniques IFIP Technical Conference*, pages 400–404. Springer, 1975.

Tiago Roux Oliveira, Nerito Oliveira Aminde, and Liu Hsu. Monitoring function based extremum seeking control for uncertain relative degrees with light source seeking experiments. In *Decision and Control (CDC), 2014 IEEE 53rd Annual Conference on*, pages 3456–3462. IEEE, 2014.

Eduardo Ramírez-Llanos and Sonia Martínez. Constrained source seeking for mobile robots via simultaneous perturbation stochastic approximation. In *Decision and Control (CDC), 2016 IEEE 55th Conference on*, pages 6851–6856. IEEE, 2016.

I Ranó, M Khamassi, and K Wong-Lin. A drift diffusion model of biological source seeking for mobile robots. In *Robotics and Automation (ICRA), 2017 IEEE International Conference on*, pages 3525–3531. IEEE, 2017.

Amarjeet Singh, Andreas Krause, and William J Kaiser. Nonmyopic adaptive informative path planning for multiple robots. In *Proceedings of the 21st International Joint Conference on Artificial Intelligence*, pages 1843–1850, 2009.

Yu Tian, Wei Li, and Fumin Zhang. Moth-inspired plume tracing via autonomous underwaer vehicle with only a pair of separated chemical sensors. In *OCEANS'15 MTS/IEEE Washington*, pages 1–8. IEEE, 2015.

Massimo Vergassola, Emmanuel Villermaux, and Boris I Shraiman. Supplementary materials for infotaxis: as a strategy for searching without gradients. *Nature*, 445(7126): 406–409, 2007a.

Massimo Vergassola, Emmanuel Villermaux, and Boris I Shraiman. 'infotaxis' as a strategy for searching without gradients. *Nature*, 445(7126):406–409, 2007b.

Ziyu Wang, Masrour Zoghi, Frank Hutter, David Matheson, and Nando De Freitas. Bayesian optimization in high dimensions via random embeddings. In *Twenty-Third International Joint Conference on Artificial Intelligence*, 2013.

Chunlei Zhang, Daniel Arnold, Nima Ghods, Antranik Siranosian, and Miroslav Krstic. Source seeking with non-holonomic unicycle without position measurement and with tuning of forward velocity. *Systems & control letters*, 56(3):245–252, 2007.

Tianpeng Zhang, Victor Qin, Yujie Tang, and Na Li. Source seeking by dynamic source location estimation. In *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 2598–2605. IEEE, 2021.

AG Zhilinskas. Single-step Bayesian search method for an extremum of functions of a single variable. *Cybernetics and Systems Analysis*, 11(1):160–166, 1975.

Appendix A. Additional related work

The task of identifying an item with distinct attributes within a collection has been extensively studied across various disciplines, each operating under its own assumptions. Consequently, our study intersects with several existing lines of work from the literature.

Our problem setting is closest to that considered by the robotics (Zhang et al., 2007; Liu and Krstic, 2010; Atanasov et al., 2012; Ramírez-Llanos and Martínez, 2016; Mishra and Zhang, 2016; Ranó et al., 2017; Gillespie et al., 2017; Zhang et al., 2021), control (Cochran and Krstic, 2009; Matveev et al., 2011; Azuma et al., 2012; Oliveira et al., 2014; Mellucci et al., 2016), and physics communities (Bénichou et al., 2005; Vergassola et al., 2007b), where a mobile agent equipped with a sensor searches for a source emitting a signal. Consistent with our work, these studies emphasize the need for the robot to move across the search space efficiently using a real-time search strategy. However, the majority of these strategies are guided by gradients, which, as discussed above, can be problematic for turbulent fields. Two exceptions areVergassola et al. (2007b) and Mishra and Zhang (2016), whose proposed algorithms we evaluate in our experiments. Finally, the algorithms proposed by these communities either lack convergence guarantees or rely on assumptions, such as the unimodality of the underlying field (Zhang et al., 2007, 2021), which are often violated in practice.

Emerging from the seminal work of Bernard Koopman (Koopman, 1956; koo, 1956; Koopman, 1957), the operations research community has developed the field of optimal search. The objective, akin to ours, remains the localization of a *target* within a search space. This community's pursuits primarily revolve around understanding how much effort (e.g., number of samples or time) should be spent at each location of the search space to maximize the probability of finding the target within a budget, or to minimize the expected effort required to find the target. However, unlike the work discussed in the previous paragraph, most of this work does not consider the cost of moving from one location to another through the search space. One exception is Wang et al. (2013), which proves that computing the optimal path is NP-hard even in simple settings. Like ours, work in this area uses probabilistic models of the target location. The uncertainty provided by such models is used to determine the effort that should be allocated at each region of the search space.

Unlike our work, however, work in this area has focused on binary response models and also often assumes there are no false positives.

The machine learning and statistics communities have also studied problems where the goal is to allocate search or sampling effort to find a specific item as efficiently as possible. Within this literature, our work is particularly close to Bayesian optimization, a framework for global optimization of functions with time-consuming evaluations (Kushner, 1964; Zhilinskas, 1975; Močkus, 1975; Frazier, 2018). In source-seeking tasks, the source corresponds to the location where the signal's mean field achieves its maximum value. Thus, Bayesian optimization could, in principle, be repurposed for such tasks. However, most Bayesian optimization algorithms do not take into account the cost of moving from the current point to the next one. Moreover, most work in Bayesian optimization uses Gaussian processes as probabilistic models. In contrast, we use a probabilistic model specifically designed for source-seeking tasks, which leverages physical knowledge of the observation process.

Within research pursued by the machine learning community, our work is also closely related to informative path planning (Krause and Guestrin, 2007; Singh et al., 2009), which focuses on finding paths that maximize the amount of information gathered by a mobile robot. In contrast, our algorithm is specifically tailored to guide the robot to find the source.

Appendix B. An example in chemical plume search

To illustrate our framework, we discuss an example focused on locating a chemical plume that emits particles within an environment with turbulent flows.

Let I denote the number of sources (i.e., chemical plumes) in the search space. For each source $i = 1, \dots, I$, we denote its location by θ_i and the rate at which particles from this source are detected when the agent is at location x by $r_i(x; \theta_i)$. The rate at which plume particles from all sources are detected when the robot is at location x is given by the sum of the rates of all sources, i.e., $R(x; \theta) = \sum_{i=1}^{I} r_i(x; \theta_i)$, where $\theta = (\theta_1, \dots, \theta_I)$.

To complete our model description, we must define the likelihood of measurements recorded by the agent. We assume that observations follow a Poisson distribution with rate λ. Hence, the likelihood is given by $L(y; \lambda) = (\lambda^y e^{-\lambda})/y!$ for $y \in \mathbb{N}_0$. This Poisson likelihood aligns well with the behavior observed in turbulent fields Vergassola et al. (2007a).

Suppose we wish to find an area of the search space where the intensity is above a dangerous level ν. We can model this by setting $S^* = [\nu, \infty)$, in which case $X^*(\theta) = \{x : R(x; \theta) \geq \nu\}$. In addition, consider a situation where the environment has a dominant "major" source accompanied by several weaker "background" sources. The major source's intensity is known to exceed ν close to its origin and decreases rapidly with increasing distance. In contrast, the combined intensity from background sources is well below ν at all locations in the search space. Under these conditions, $X^*(\theta)$ will be constituted by a small region in the neighborhood of the major source. In practice, finding a location $x \in X^*(\theta)$, as guaranteed by our algorithm, typically implies successful identification of the major source.

Appendix C. Posterior distribution

Let $x_1, \ldots, x_n \in \mathbb{X}$ denote the locations visited by the agent up to time n, and let $y_1, \ldots, y_n \in \mathbb{R}^l$ denote the corresponding measurements at these locations. We use the compact notation $x_{1:n} = (x_1, \ldots, x_n)$ and $y_{1:n} = (y_1, \ldots, y_n)$. Let $p_n(\theta) = p(\theta \mid x_{1:n}, y_{1:n})$ denote the posterior distribution of θ given the measurements collected up to time n. By Bayes rule we have

$$p_n(\theta) = \frac{p(\theta \mid x_{1:n}, y_{1:n-1}) p(y_n | \theta, x_{1:n}, y_{1:n-1})}{p(y_n | x_{1:n}, y_{1:n-1})}.$$

The first term in the numerator of the above equation is equivalent to $p(\theta | x_{1:n-1}, y_{1:n-1})$ because moving to a new location does not provide information about the source's location until a measurement is observed. Moreover, since measurements are assumed to be conditionally independent given the source's location,

$$\begin{aligned} p(y_n | \theta, y_{1:n-1}, x_{1:n}) &= p(y_n | x_n; \theta) \\ &= L(y_n; R(x_n; \theta)). \end{aligned}$$

Finally, the denominator $p(y_n | x_{1:n}, y_{1:n-1})$ is equal to

$$\int_{\theta'} p(\theta' | x_{1:n-1}, y_{1:n-1}) L(y_n; R(x_n; \theta'))\, d\theta'.$$

Thus, the posterior distribution of θ at time n can be rewritten as

$$p_n(\theta) = \frac{p(\theta | x_{1:n-1}, y_{1:n-1}) L(y_n; R(x_n; \theta))}{\int_{\theta'} p(\theta' | x_{1:n-1}, y_{1:n-1}) L(y_n; R(x_n; \theta'))\, d\theta'}.$$

The above equation also gives rise to a recursive expression to compute p_n in terms of p_{n-1} as follows

$$p_n(\theta) = \frac{p_{n-1}(\theta) L\left(y_n; R(x_n; \theta)\right)}{\int_{\theta'} p_{n-1}(\theta') L(y_n; R(x_n; \theta'))\, d\theta'}.$$

Appendix D. Convergence analysis

In this section, we prove that the location of an agent following the max probability algorithm converges almost surely to a location in $X^*(\theta)$. Our analysis relies on the following two assumptions.

Assumption 1. *$X^*(\theta)$ is non-empty almost surely.*

Assumption 2. *There is a sequence of functions $\{E_n\}_{n=1}^{\infty}$ such that, for any $\lambda \in \Lambda$, if $\{y_n\}_{n=1}^{\infty}$ is a sequence of i.i.d. random variables with common distribution $L(\,\cdot\,; \lambda)$, then $\{E_n(y_1, \ldots, y_n)\}_{n=1}^{\infty}$ converges almost surely to λ.*

Assumption 1 simply guarantees that the source-seeking task is feasible almost surely. We now explain how Assumption 2 is used at an intuitive level. Let x be a location visited infinitely often by the agent. Observations at x are conditionally i.i.d. given θ with common conditional distribution $L\left(\,\cdot\,; R(x; \theta)\right)$. Then, asymptotically as the number of observations

at x goes to infinity, we can estimate $R(x;\theta)$ with arbitrarily high precision by Assumption 2. Since, by definition, $x \in \mathbb{X}^*(\theta)$ if and only if $R(x;\theta) \in S^*$, we can then determine if x is in the target set.

Our main result is stated in Theorem 1. We prove this result through a series of lemmas. Before proving these lemmas, we establish the following notation. We denote the σ-algebra generated by the measurements collected up to time n by $\mathcal{F}_n$. We also let $\mathcal{F}_\infty$ denote the minimal σ-algebra generated by $\{\mathcal{F}_n\}_{n=1}^\infty$.

Lemma 1. *For each $x \in \mathbb{X}$, $\lim_{n\to\infty} q_n(x) = \mathbf{E}[\mathbf{1}\{x \in X^*(\theta)\}|\mathcal{F}_\infty]$ almost surely.*

Proof Fix any $x \in \mathbb{X}$ and observe that $\mathbf{1}\{x \in X^*(\theta)\}$ is an integrable random variable. Thus, by Theorem 4.7 in Çinlar, $\{\mathbf{E}[\mathbf{1}\{x \in X^*(\theta)\}|\mathcal{F}_n]\}_{n=1}^\infty$ is a uniformly integrable martingale converging to $\mathbf{E}[\mathbf{1}\{x \in X^*(\theta)\}|\mathcal{F}_\infty]$ almost surely. This ends the proof since $q_n(x) = \mathbf{E}[\mathbf{1}\{x \in X^*(\theta)\}|\mathcal{F}_n]$ by definition.

Let $N_n(x) = \sum_{n'=1}^n \mathbf{1}\{x_{n'} = x\}$ be the number of times a location x has been visited by the agent by time n. We also define $N_\infty(x) = \lim_{n\to\infty} N_n(x)$. The following lemma shows that if $N_\infty(x) = \infty$, i.e., if a location x is visited by the agent infinitely often, then we can determine if the location is in the target set almost surely.

Lemma 2. *For each $x \in \mathbb{X}$, there exists an event A_x such that $\mathbf{P}(A_x) = 1$ and $\lim_{n\to\infty} q_n(x) = \mathbf{1}\{x \in X^*(\theta)\}$ on the event $\{N_\infty(x) = \infty\} \cap A_x$.*

Proof Fix $x \in \mathbb{X}$ and let $\{y'_n\}_{n=1}^\infty$ be a sequence of random variables independent of each other and independent of $\{y_n\}_{n=1}^\infty$, whose common conditional distribution given θ is $L(\,\cdot\,; R(x;\theta))$. Define the sequence $\{m_n\}_{n=1}^\infty$ recursively by $m_1 = \inf\{m : x_m = x\}$ and $m_n = \inf\{m > m_{n-1} : x_m = x\}$ for $n > 1$. For each n, we define $z_n = y_{m_n}$ if $m_n < \infty$ and $z_n = y'_n$ otherwise. Observe that, conditioned on θ, $\{z_n\}_{n=1}^\infty$ is a sequence of i.i.d. random variables with common distribution $L(\,\cdot\,; R(x;\theta))$. By Assumption 2, there exists an event A_x such that $\mathbf{P}(A_x) = 1$ and $\lim_{n\to\infty} E_n(z_1,\ldots,z_n) = R(x;\theta)$ on A_x.

Let E_∞ be defined by $E_\infty = \lim_{n\to\infty} E_n(z_1,\ldots,z_n)$ if $N_\infty = \infty$ and $\lim_{n\to\infty} E_n(z_1,\ldots,z_n)$ exists, and $E_\infty = 0$ otherwise. Observe that E_∞ is $\mathcal{F}_\infty$-measurable and $E_\infty = R(x;\theta)$ on the event $\{N_\infty(x) = \infty\} \cap A_x$. Consequently, the statements below hold on the event $\{N_\infty(x) = \infty\} \cap A_x$.

We have $\mathbf{1}\{E_\infty \in S^*\} = \mathbf{1}\{R(x;\theta) \in S^*\}$. Moreover, $\mathbf{1}\{R(x;\theta) \in S^*\} = \mathbf{1}\{x \in X^*(\theta)\}$ by definition. Hence, $\mathbf{1}\{E_\infty \in S^*\} = \mathbf{1}\{x \in X^*(\theta)\}$, which in turn implies that $\mathbf{E}[\mathbf{1}\{E_\infty \in S^*\} \mid \mathcal{F}_\infty] = \mathbf{E}[\mathbf{1}\{x \in X^*(\theta)\} \mid \mathcal{F}_\infty]$. In addition, since $\mathbf{1}\{E_\infty \in S^*\}$ is $\mathcal{F}_\infty$-measurable, $\mathbf{E}[\mathbf{1}\{E_\infty \in S^*\} \mid \mathcal{F}_\infty] = \mathbf{1}\{E_\infty \in S^*\}$. From the last two equations we get $\mathbf{E}[\mathbf{1}\{x \in X^*(\theta)\} \mid \mathcal{F}_\infty] = \mathbf{1}\{x \in X^*(\theta)\}$, which concludes the proof by virtue of Lemma 1.

Lemma 3. *If $q_n(x) < \frac{1}{c}$, where $c = |\mathbb{X}|$, then $x \notin \operatorname{argmax}_{x'\in\mathbb{X}} q_n(x')$.*

Proof Observe that $\sum_{x'\in\mathbb{X}} q_n(x') \geq 1$ by Assumption 1. Thus, there exists $\tilde{x} \in \mathbb{X}$ such that $\frac{1}{c} \leq q_n(\tilde{x})$. Since, $q_n(x) < \frac{1}{c} \leq q_n(\tilde{x})$, it follows that $x \notin \operatorname{argmax}_{x'\in\mathbb{X}} q_n(x')$.

Let $T_\infty = \{x : \widehat{x}_n = x \text{ for infinitely many } n\}$. The following lemma shows that a location whose probability of being in the target set vanishes as the number of iterations goes to infinity cannot be in T_∞.

Lemma 4. *If $\lim_{n\to\infty} q_n(x) = 0$, then $x \notin T_\infty$.*

Proof We prove this by contradiction. Suppose $\lim_{n\to\infty} q_n(x) = 0$ for some $x \in T_\infty$. Since $\lim_{n\to\infty} q_n(x) = 0$, there exists N_1 large enough such that $q_n(x) < \frac{1}{c}$ for all $n > N_1$. From Lemma 3 we know that $x \notin \operatorname{argmax}_{x'\in\mathbb{X}} q_n(x')$ for all $n > N_1$. Since $x \in T_\infty$, it follows from this that there exists N_2 such that $\widehat{x}_{N_2} = x \in \operatorname{argmax}_{x'\in\mathbb{X}} q_{N_2}(x')$ and $q_n(x) \geq \alpha q_{N_2}(x)$ for all $n > N_2$. Moreover, since $x \in \operatorname{argmax}_{x'\in\mathbb{X}} q_{N_2}(x')$, it follows from Lemma 3 that $q_{N_2}(x) > 0$. Thus, $\liminf_{n\to\infty} q_n(x) \geq \alpha q_{N_2}(x) > 0$, which contradicts that $\lim_{n\to\infty} q_n(x) = 0$.

Lemma 5. *Almost surely, there exists N large enough such that $\widehat{x}_n = x$ and $q_n(x) \geq \alpha q_{k_{n-1}}(x)$ for all $n > N$, where x is the only element of T_∞. In particular, T_∞ has cardinality 1 almost surely.*

Proof T_∞ is non-empty by the pigeonhole principle. Fix $x \in T_\infty$. From Lemmas 1 and 4, we deduce $\lim_{n\to\infty} q_n(x)$ exists and is strictly positive almost surely. Let $q_\infty(x) = \lim_{n\to\infty} q_n(x)$ and define $\epsilon = \frac{1-\alpha}{1+\alpha} q_\infty(x)$. There exists N_1 large enough such that $q_n(x) \in [q_\infty(x) - \epsilon, q_\infty(x) + \epsilon]$ for all $n > N_1$. We now consider two cases. In the first case, there is no $n > N_1$ such that $\widehat{x}_n = x$ and $q_n(\widehat{x}_{n-1}) < \alpha q_n(\widehat{x}_{n-1})$. Since $x \in T_\infty$, this necessarily implies there is a $N_2 \leq N_1$ such that $\widehat{x}_{N_2} = x \in \operatorname{argmax}_{x'\in\mathbb{X}} q_{N_2}(x')$ and $q_n(x) \geq \alpha q_{N_2}(x)$ for all $n > N_2$. In the second case, there exists $n > N_1$ such that $\widehat{x}_n = x$ and $q_n(\widehat{x}_{n-1}) < \alpha q_n(\widehat{x}_{n-1})$. Let N_2 denote such a n. Since $N_2 > N_1$, we have

$$\begin{aligned} q_n(x) &\geq q_\infty(x) - \epsilon \\ &= q_\infty(x) - \frac{1-\alpha}{1+\alpha} q_\infty(x) \\ &= \frac{2\alpha}{1+\alpha} q_\infty(x) \\ &= \alpha(q_\infty(x) + \epsilon) \\ &\geq \alpha q_{N_2}(x). \end{aligned}$$

for any $n > N_2$. If we let $N = N_2$, we obtain, in any of the two cases described above, that $\widehat{x}_n = x$ and $q_n(x) \geq \alpha q_{k_{n-1}}(x)$ for all $n > N$. It follows that $T_\infty = \{x\}$, which concludes the proof.

We are now ready to prove our main result.

Theorem 1. $\lim_{n\to\infty} x_n = x$ *for some* $x \in X^*(\theta)$ *almost surely.*

Proof Fix $x \in \mathbb{X}$ and consider the event where the random variable N from Lemma 5 is finite and $T_\infty = \{x\}$. Call this event B_x. On this event, using Lemma 5, $\widehat{x}_n = x$ for all $n > N$. Thus, $x_{n+1} \in \operatorname{argmin}_{x'\in\mathcal{I}(x_n)} \rho(x', x)$ for all $n > N$. Since each location is reachable from each other location in $\mathbb{X}$, it must be the case that $x_n = x$ for all n large enough. By Lemma 2, $\lim_{n\to\infty} q_n(x) = \mathbf{1}\{x \in X^*(\theta)\}$ on the event $\{N_\infty(x) = \infty\} \cap A_x \cap B_x$. Moreover, since $x \in T_\infty$, $\lim_{n\to\infty} q_n(x) > 0$ by Lemma 4. Hence, $\mathbf{1}\{x \in X^*(\theta)\} > 0$, i.e., $x \in X^*(\theta)$. Observe that this analysis also implies that $B_x \subset \{N_\infty(x) = \infty\}$. Thus, we have shown that $x_n = x$ for all n large enough and $x \in X^*(\theta)$ on the event $A_x \cap B_x$. Finally, recall that $\cup_{x\in\mathbb{X}} B_x$ is an almost sure event by Lemma 5. Moreover, A_x is an almost sure event for each $x \in \mathbb{X}$. Hence, $\cup_{x\in\mathbb{X}} A_x \cap B_x$ is an almost sure event. We conclude that $\lim_{n\to\infty} x_n = x$ for some $x \in X^*(\theta)$ almost surely.

Appendix E. Additional numerical experiment

Here, we present simulation results for a problem with two sources in a $4m \times 4m$ 2-dimensional search space, discretized into a 80×80 uniform grid. We set the problem up so that the emission rate of the first source is significantly larger than the emission rate of the second source by taking $s_1 = 1$ plumes/s and $s_2 = 0.25$ plumes/s. We also make the lifetime of particles generated by the first source larger than that of the second source by setting $\tau_1 = 2500\,s$ and $\tau_2 = 1000\,s$. The remainder of the parameters of f_1 and f_2, which are common for both sources, are taken like in the unimodal case. The goal is to find the major source, i.e., θ_1. The value of the intensity at this location is $f_1(0) + f_2(||\theta_1 - \theta_2||_2)$. Because of the way in which problem parameters were chosen, this is the unique location whose intensity is larger than $f_1(0)$. This occurs because the first source is much stronger than the second, they are well-separated when generating instances, and because intensity falls off quickly with distance from a source. Thus, it is sufficient to take $S^* = [f_1(0), \infty)$, which ensures that $X^*(\theta) = \{\theta_1\}$.

To evaluate the performance of the three algorithms, we performed 300 simulation runs. The locations of the sources were chosen uniformly at random for each run with the constraint that the distance between the two sources was at least 25 cells. Out of 300 simulation runs, 100 were chosen such that the agent's initial location was closer to the highest intensity source, 100 were chosen such that the agent's initial location was closer to the lowest intensity source, and 100 were chosen agent's initial location was approximately the same distance to both sources. In all cases, it was also guaranteed that the agent's initial location was at least 20 cells away from both sources. Unlike in the experiment in Section 4.1, here we cap the maximum number of steps taken by the agent to 300 to reduce the computational cost of the simulations. In the next two subsections, we present the statistical analysis of consistency and rate of convergence of the three algorithms.

E.1 Consistency of convergence

Here, the agent is considered to have found the source if it is within one cell from the source. Figure 4a depicts the percentage of trajectories that converge to the highest intensity source within 300 steps for each algorithm. These percentages are 98%, 88%, and 84% for the max probability, expected rate, and infotaxis algorithms, respectively. Thus, the max probability algorithm converges more consistently than the other two algorithms.

We also analyze how often the algorithms converge to the highest intensity source, in line with the goal of the search, and how often they fail to do this and instead converge to the lowest intensity source. Figure 4b shows the percentage of trajectories that converged to the lowest intensity source for each algorithm. None of the trajectories of the max probability algorithm converged to the lowest intensity source. A manual inspection shows that the 2% of the trajectories of the max probability algorithm that did not converge to the highest intensity source were closer to it than to the lowest intensity source. On the other hand, 9.3% and 6.7% of the trajectories of the expected rate and infotaxis algorithms converged to the lowest intensity source, respectively. These trajectories were typically such that the agent's initial location was closer to the low-intensity source than to the high-intensity source. Thus, these results suggest that the expected rate and infotaxis algorithms are more prone to converging to the low-intensity source.

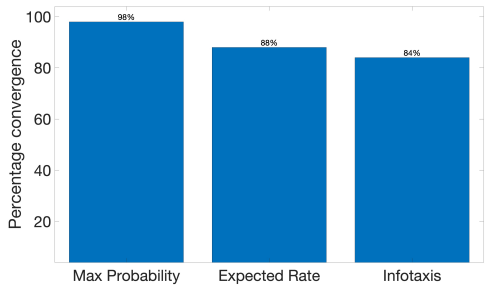

(a) Percentage of trajectories converging to the global maxima for each algorithm in the problem described in Section E.

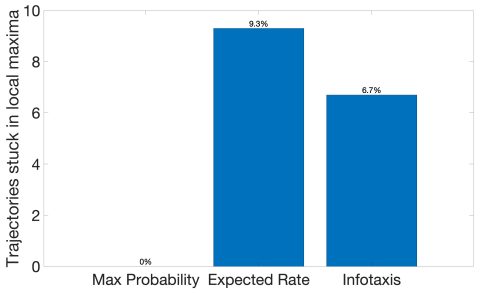

(b) Percentage of trajectories stuck in the local maxima for each algorithm in the problem described in Section E.

Figure 4

E.2 Rate of convergence

Figure 5a shows the median of the normalized distance between the agent and the high-intensity source after a given number of iterations. While the three algorithms eventually reach a median normalized distance of zero, the max probability and expected rate algorithms do so faster than the infotaxis algorithm. Figure 5b depicts the 95^{th} percentile of the normalized distance between the agent's location and the highest intensity source after a given number of iterations. This value is consistently lower for the max probability algorithm than the expected rate algorithm, which is, in turn, lower than the corresponding value of the infotaxis algorithm. Moreover, only the value corresponding to the max probability algorithm goes to zero within 300 iterations. From these results, we conclude that the max probability algorithm converges faster and more reliably than the expected rate and infotaxis algorithms.

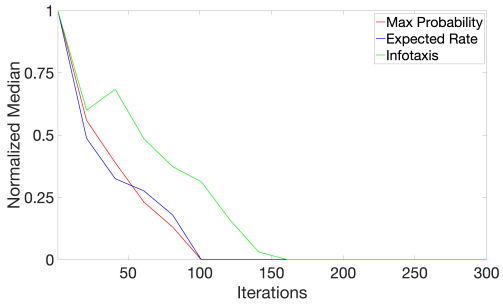

(a) Median of the normalized distance between the agent and the high-intensity source for the max probability (red), expected rate (blue), and infotaxis (green) algorithms.

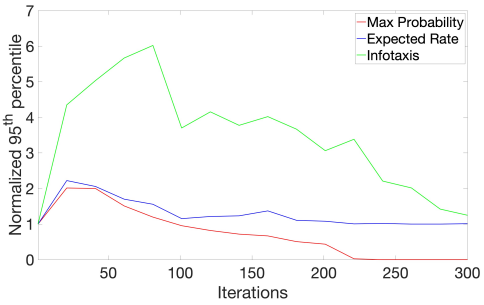

(b) 95^{th} percentile of the normalized distance between the agent and the high-intensity source for the max probability (red), expected rate (blue), and infotaxis (green) algorithms.

Figure 5