Physics Letters A 380 (2016) 1698–1705

Contents lists available at ScienceDirect

Physics Letters A

www.elsevier.com/locate/pla

Multi-agent search for source localization in a turbulent medium

Hadi Hajieghrary [a], M. Ani Hsieh [a], Ira B. Schwartz [b]

[a] Scalable Autonomous Systems Lab, Mechanical Engineering & Mechanics, Drexel University, Philadelphia, PA 19104, USA
[b] Nonlinear Systems Dynamics Section, Plasma Physics Division, Code 6792, U.S. Naval Research Laboratory, Washington, DC 20375, USA

ARTICLE INFO

Article history:
Received 20 June 2015
Received in revised form 7 March 2016
Accepted 9 March 2016
Available online 15 March 2016
Communicated by F. Porcelli

Keywords:
Multi-agent systems
Information theory
Distributed control

ABSTRACT

We extend the gradient-less search strategy referred to as "infotaxis" to a distributed multi-agent system. "Infotaxis" is a search strategy that uses sporadic sensor measurements to determine the source location of materials dispersed in a turbulent medium. In this work, we leverage the spatio-temporal sensing capabilities of a mobile sensing agents to optimize the time spent finding and localizing the position of the source using a multi-agent collaborative search strategy. Our results suggest that the proposed multi-agent collaborative search strategy leverages the team's ability to obtain simultaneous measurements at different locations to speed up the search process. We present a multi-agent collaborative "infotaxis" strategy that uses the relative entropy of the system to synthesize a suitable search strategy for the team. The result is a collaborative information theoretic search strategy that results in control actions that maximize the information gained by the team, and improves estimates of the source position.

1. Introduction

Nature has been optimizing search strategies in complex uncertain environments for billions of years. For example, the efficiency of male moths searching for females is quite remarkable. In spite of a slew of serious obstructions in signal detection and processing, the olfactory pheromone system usually guarantees a successful encounter [1]. Far from the pheromone emitting female, odor plumes consist of sparsely distributed pheromone patches due to turbulence [2], leading to rare, intermittent detections [3,4]. These and similar search behaviors in biological systems give insight into the mechanisms linking perception to action, and for developing more effective search strategies in general. The understanding of how ants and honeybees locate and return to their respective colonies when foraging for food has led to new probabilistic graph search strategies like ant and bee colony optimization [5,6]. These successes have further increased interest in developing similar strategies for autonomous mobile robots for search and rescue applications like the detection of gas leaks and exploration of buildings on fire [7], tracking of hazardous chemical plumes [8], and multi-robot exploration of unknown environments [9,10].

Existing strategies for detecting, tracking, and localizing gas/odor/radiation sources with mobile robots include building of flow field maps [11–13], estimating concentration gradients [14–19], and using gradient-free search algorithms [16,17]. However, the variation in chemical concentrations from a source in a flow environment is heavily dependent on the Reynolds numbers. In general, gradient-based approaches generally work better in lower Reynolds regimes since the variation in chemical concentrations are generally smoother [20]. However, as the Reynolds number increases, dispersion of the chemical source becomes increasingly dominated by turbulent mixing which renders many gradient-based strategies impractical [21–25].

To address some of these challenges, various bio-inspired strategies based on bacteria [15–17], insects [26–28], and crabs [29] have also been proposed. However, these bio-inspired approaches are mostly ad-hoc, focused on developing novel sensor technology [30,27], or are equivalent to coverage and gradient based search strategies for single robots [16,17,26–28,20]. Alternatives to these existing strategies include a new class of reactive search strategies. These strategies do not rely on continuous or smooth concentration gradients and can adapt to past sensory information and action [4,11,31]. The so called "infotaxis" strategy maximizes the information gain about the location of the source in a turbulent medium [4]. While these strategies are more sophisticated, they are also computationally expensive and have almost exclusively focused on adaptive behavior in the context of single agent search strategies.

The main contribution of this work is the extension of the single-agent information theoretic search strategy coined as "infotaxis" to a multi-agent robotic system [4]. As such, we present two collaborative information maximizing search strategies for the multi-robot team and compare their performances to the single agent "infotaxis" strategy. The paper is organized as follows: We briefly summarize the single agent "infotaxis" strategy and lay out

E-mail address: Hadi.Hajieghrary@Drexel.edu (H. Hajieghrary).

http://dx.doi.org/10.1016/j.physleta.2016.03.013

our assumptions in Section 2. We present the multi-agent "infotaxis" search strategy in Section 3.1 and present our simulation results in Section 4. The comparison and discussion of the single- and multi-agent strategies are presented in Section 5. We conclude with a summary of our results in Section 6.

2. Background

The main objective is to extend the single-agent "infotaxis" search strategy presented in [4] to a multi-robot system. We begin by outlining our assumptions and briefly summarizing the single-agent strategy.

2.1. Assumptions

Given an obstacle free workspace in two dimensions (2D) denoted by $\mathcal{W}$, we assume $\mathcal{W}$ is discretized into uniform grid cells. A cell is *occupied* if the agent, or robot, is located within it and agents are only allowed to move from their current cell to any of the eight adjacent cells, i.e., up, down, left, right, and diagonally across the four corners. As such, each agent is effectively modeled as a single massless point particle with omnidirectional kinematics. Each robot is equipped with a binary chemical sensor that is capable to detect the presence (or absence) of chemical at the agent's current location in $\mathcal{W}$. We also assume every agent has the ability to localize within $\mathcal{W}$, i.e., know its position, and can measure the magnitude and direction of the local flow field at its current location. Lastly, we assume each agent can communicate with every other member of the team. While inter-agent information exchange can be done asynchronously, we assume no delays in communication since the focus of this work is to compare the performance of the multi-agent strategy with the single-agent system rather than to study the effects of communication delays on the search strategy.

In general, the expected rate of positive material plume detection in an environment depends on the spatial distance to the source, the dynamics of the surrounding flow field, the geometry of the environment, and many other factors. Due to the complexity of the dispersion dynamics of biochemical and/or radiological material in a turbulent medium, similar to [4], we model the rate of detecting the presence of a material plume as a Poisson distribution. The statistical model for positive material plume detection in a turbulent medium given in [4] is briefly summarized below.

The mean rate of positive detection at position $\mathbf{r}$ for a source located at $\mathbf{r_0}$ in 2D is given by:

$$R(\mathbf{r}|\mathbf{r_0}) = \frac{\mathcal{R}}{\ln(\frac{\lambda}{a})} e^{\frac{(y_0-y)V}{2D}} K_0\left(\frac{|\,\mathbf{r}-\mathbf{r_0}\,|}{\lambda}\right), \quad \text{with} \tag{1}$$
$$\lambda = \sqrt{\frac{D\tau}{1+\left(\frac{V^2\tau}{4D}\right)}},$$

where $\mathcal{R}$ is the emission rate of the source, τ is the finite lifetime of the chemical patch before its concentration falls off the detectable range, D is the isotropic effective diffusivity of the medium, and V is the mean velocity of the wind, and $K_0(\cdot)$ is the modified Bessel function of the second kind. Similar to [4], we assume a strong background directional flow that is predominantly in the $-y$ direction.

Let $P_t(\mathbf{r_0})$ denote the estimated probability distribution that describes the possible locations of the source in $\mathcal{W}$. This probability distribution function at time t represents the information gathered through a series of uncorrelated positive sensor measurements or positive odor encounters. In general, $P_t(\mathbf{r_0})$ can be computed using Bayes' rule:

$$P_t(\mathbf{r_0}) = \frac{P_{\mathbf{r_0}}(\mathbf{z}_{1:t})}{\int_{\mathcal{W}} P_{\mathbf{r}}(\mathbf{z}_{1:t})\,\mathrm{d}\mathbf{r}}, \tag{2}$$

where $\mathbf{r}$ denotes a position in the workspace $\mathcal{W}$, $\mathbf{z}_{1:t}$ denotes the history of odor encounters, and $P_{\mathbf{r_0}}(\mathbf{z}_{1:t})$ is the likelihood of obtaining such a history of sensor measurements if the source is located at $\mathbf{r_0}$. The expected rate of positive sensor measurements at any given location in $\mathcal{W}$ is a function of the relative position of the agent with respect to the source. Assuming that the detection of the plume at every location in $\mathcal{W}$ is independent of its neighboring positions, we use Poisson's law to estimate the number of detections at each step during the exploration as in [4]. We note that the assumption of independence for the detection probability holds since the location of the source is unknown. As such, $P_{\mathbf{r_0}}(\mathbf{z}_{1:t})$ is given by an exponential distribution of the form

$$P_{\mathbf{r_0}}(\mathbf{z}_{1:t}) = \exp\left(\left[-\int_0^t R(\mathbf{r}(t')|\mathbf{r_0})\,\mathrm{d}t'\right]\prod_{i=1}^{h} R(\mathbf{r}(t_i)|\mathbf{r_0})\right), \tag{3}$$

where $\mathbf{r}(t)$ denotes the positions where the agent obtained positive sensor readings, i.e., detection of the presence of the chemical/plume, and h is the number of such detections.

2.2. Single-agent search strategy

To maximize the expected rate of information gain for the source location, the single-agent search strategy in [4] is designed to maximize the expected reduction of the entropy of $P_t(\mathbf{r})$. Since the entropy of $P_t(\mathbf{r})$ is defined as $S_t = -\int_{\mathcal{W}} P_t(\mathbf{r})\log(P_t(\mathbf{r}))\,\mathrm{d}\mathbf{r}$, the expected change in entropy for an agent moving from its current location $\mathbf{r}$ to another location $\mathbf{r}_j$ in $\mathcal{W}$ at current time step t is given by:

$$\begin{aligned}&\mathbf{E}[\![\Delta S_t(\mathbf{r}\mapsto\mathbf{r}_j)]\!]\\ &\quad= P_t(\mathbf{r}_j)[-S_t] + [1-P_t(\mathbf{r}_j)][(1-\rho(\mathbf{r}_j))\Delta S_0 + \rho(\mathbf{r}_j)\Delta S_1].\end{aligned} \tag{4}$$

The first term in (4) corresponds to the change in entropy should the agent find the source at the next step. The second term of (4) accounts for the likelihood that the source is not at r_j and calculates the mean value of the information gained from additional encounters. Under these circumstances, we consider two possible outcomes: the agent obtains a positive or negative sensor measurement at the new position. In (4), $\rho(\mathbf{r}_j)$ denotes the probability of detection made at the next step. The expected number of positive measurements at the new position is calculated based on the current estimate of $P_t(\mathbf{r})$. Here, the observation model is the expected rate of positive sensor measurements at position $\mathbf{r}$ if the source is located at $\mathbf{r_0}$ given by $R(\mathbf{r}(t)|\mathbf{r_0})$. Given the estimate of the source location, the expected number of hits at any location in the workspace is given by:

$$h(\mathbf{r}_j) = \int P_t(\mathbf{r}_j) R(\mathbf{r}_j|\mathbf{r_0})\,\mathrm{d}\mathbf{r_0}, \tag{5}$$

where the probability of a single positive detection follows the Poisson law $\rho(\mathbf{r}_j) = h(\mathbf{r}_i)\exp(-h(\mathbf{r}_j))$. In (4), ΔS_0 denotes the entropy change of the probability density function of the source location if the agent receives no new positive sensor measurements at the next time step as it moves to the neighboring cell. The change in entropy of the probability density function if a detection is made at the new position is given by ΔS_1. The result is a search strategy that favors agent motions that maximize the likelihood of finding the source location. We refer the interested reader to [4] for the complete details.