Information Fusion 28 (2016) 1–9

Contents lists available at ScienceDirect

Information Fusion

journal homepage: www.elsevier.com/locate/inffus

A study of cognitive strategies for an autonomous search

Branko Ristic*, Alex Skvortsov, Ajith Gunatilaka

Defence Science and Technology Organisation, Melbourne, Australia

ARTICLE INFO

Article history:
Received 4 February 2015
Received in revised form 5 May 2015
Accepted 14 June 2015
Available online 2 July 2015

Keywords:
Autonomous search
Turbulent flow
Sensor management
Sequential Monte Carlo method
Dispersion modelling

ABSTRACT

Cognitive search is a collective term for search strategies based on information theoretic rewards required in sequential decision making under uncertainty. The paper presents a comparative study of cognitive search strategies for finding an emitting source of unknown strength using sparse sensing cues in the form of occasional non-zero sensor measurements. The study is cast in the context of an emitting source of particles transported by turbulent flow. The search algorithm, which sequentially estimates the source parameters and the reward function for motion control, has been implemented using the sequential Monte Carlo method. The distribution of the search time has been explained by the inverse Gaussian distribution.

1. Introduction

Searching strategies for finding a source of emissions (e.g. of particles, gas, odour, energy, or even propaganda) are of great importance in many aspects of life. For example, in the context of national security, there could be a need to find a source of accidental or deliberate release of toxic gases or biochemical particles into the atmosphere [1]. Furthermore, a recovery and rescue mission could be tasked to localise a lost piece of equipment emitting weak signals [2]. For the sake of understanding nature, scientists endeavour to explain the foraging behaviour of animals in their search for food or a mate [3]. Similar problems are encountered in molecular biology (protein searching on DNA sequence), meteorology (finding strong sources of carbon emissions), ecology (sources of pollution), etc.

Typically, the searcher is mobile and capable of sensing emissions from the source. The sensing cues, in the form of the non-zero sensor measurements, are often sporadic, fluctuating and discontinuous, due to the turbulent transport through the medium in a large search domain [4]. The objective of search is to find the emitting source in the shortest possible time.

The earliest theoretical studies of search strategies were conducted during WWII, when the US navy was looking for the most efficient flight paths of aircraft hunting submarines [5]. These classical approaches relied on predetermined systematic paths, such as the parallel sweep or Archimedean spiral [2]. Search patterns of animals, on the other hand, have been known to be random, rather than deterministic paths. At first glance, diffusion may seem to be the appropriate random search strategy for biological systems (e.g. foraging animals, protein searching on DNA sequence) [6]. However, for this strategy, the mean time to find the source tends to infinity. Viswanathan et al. [3] demonstrated that albatrosses fly fractal search patterns, called Lévy flights. This discovery led to several papers demonstrating Lévy walks/flight as the optimal search strategy for foraging animals (deer, bees, etc.).

Search motion patterns, however, seem to depend on the ratio between the search domain and the sensing range or the density of targets. Humphries et al. [7] demonstrated that Lévy behaviour occurs only in environments where the prey is sparsely distributed, whilst Brownian motion is optimal if the prey is abundant. An alternative to Lévy strategies is the intermittent search: a combination of fast and non-receptive displacement phase (long jumps within the search domain, with no sensing) with a slow (diffusive-like) search phase, characterised by sensing and reaction [8]. Bénichou et al. provide both a theoretical study and an experimental verification of the intermittent search strategy [9]. In their study, random walk (diffusion) was assumed during the slow search phase thus ignoring the sensing data. Sensory cues, however, play an important role in this phase of the search. Bacteria like *Escherichia coli*, for example, direct their motion (swim) towards a higher or lower gradient of chemical concentration. This type of slow search, referred to as *chemotaxis* [6], is of great interest in biology and chemistry. Chemotaxis, however, is restricted to concentration fields with well defined gradients. In many practical situations, where mixing of turbulent flow breaks up the chemical signal into random and disconnected patches,

* Corresponding author at: DSTO, LD, Bld 51, 506 Lorimer Street, Melbourne, Vic 3207, Australia. Tel.: +61 3 9626 8370; fax: +61 3 9626 8473.

E-mail address: branko.ristic@dsto.defence.gov.au (B. Ristic).

http://dx.doi.org/10.1016/j.inffus.2015.06.008
1566-2535/

the chemotaxis strategy fails. Vergassola et al. [10] proposed, for this case, a search strategy (referred to as *infotaxis*) based on information theoretic principles: infotaxic strategy maximises the expected rate of entropy reduction. This search strategy made a significant impact, resulting in several papers studying its properties and proposing modifications [11–15]. Alternative information theoretic search strategies (collectively referred to as *cognitive* strategies) have been developed and exploited in the context of searching for radioactive sources [16,17] and a chemical source in the presence of obstacles [18].

Autonomous cognitive search strategies are formulated as a partially-observed Markov decision process (POMDP) [19]. The elements of a POMDP include: an information state, a set of admissible actions and a reward function. Adopting the Bayesian framework for estimation of source parameters, the posterior probability density function (PDF) acts as the information state in the POMDP. Current knowledge about source parameters (its location, strength) is fully specified by the posterior PDF. The set of admissible actions is the set of locations where the searcher can move next. Finally, the reward function maps each admissible action into a non-negative real number, which represents the measure of its corresponding expected knowledge-gain. Optimal strategy selects at each time the action with the highest reward. Admissible actions can be one or multiple steps ahead. The infotaxi strategy [10] was formulated as a one-step ahead strategy, with five possible actions for motion on a square lattice: stay in the current node, or move right, left, up or down to the neighbouring node. Sequential estimation in [10] was carried out in the Bayesian framework, using an approximate grid-based nonlinear filtering technique [20] on a two-dimensional parameter space (source position in x–y plane). The source strength (release rate) is assumed known, as well as the environmental parameters. The reward function is based on the expected variation of entropy.

In this paper we make the following contributions. First, we extend the estimation parameter space to include the source strength (in addition to source location), since in practice, this parameter is unknown. Bayesian estimation via approximate grid-based numerical integration, used in [10], is known to suffer from the curse of dimensionality. It provides good approximation only if the grid is sufficiently dense, which comes at increased computational complexity. Another disadvantage of grid-based methods is that the state space is predefined, and therefore cannot be partitioned unevenly to give greater resolution in high probability density regions [21,22]. This brings us to the second contribution: the search (including the sequential estimation of source parameters and the computation of the reward function) is developed and implemented in the sequential Monte Carlo framework as a particle filter [21,22]. Information theoretic sensor management favours the use of Rényi divergence as the reward function [23]. The paper compares (in the context of search) the infotaxic reward with a particular Rényi-type divergence, known as the Bhattacharyya distance [24]. Finally we postulate and demonstrate that under persistent sensory cues, the PDF of search time is inverse Gaussian. Numerical analysis in the paper is carried out both by simulations and using real data.

2. Mathematical formulations

A source of constant emission of particles, characterised by the emission-rate Q_0, is located at $\mathbf{r}_0 = (x_0, y_0)^\intercal \in \mathcal{A}$, where $\mathcal{A} \subset \mathbb{R}^2$ is designated the "open-field" two-dimensional search area. The unknown source parameter vector is defined as $\mathbf{x} = [Q_0\ \mathbf{r}_0^\intercal]^\intercal$.

2.1. Measurement model

The model of turbulent transport of particles through the medium is adopted from [10]. Suppose the average lifetime of particles, propagating with the isotropic diffusivity D, is τ. The particles can be advected by wind or current, with the speed of advection V, and the mean direction (adopted by convention) to coincide with the direction of the x axis. A spherical sensor of size a, installed on a searching platform located at $\mathbf{r} = (x, y)^\intercal$, will experience a series of encounters with emitted particles at the rate:

$$R(\mathbf{r}|\mathbf{x}) = \frac{Q_0}{\ln\left(\frac{\lambda}{a}\right)} \exp\left[\frac{(x - x_0)V}{2D}\right] \cdot K_0\left(\frac{\sqrt{(x - x_0)^2 + (y - y_0)^2}}{\lambda}\right) \tag{1}$$

where K_0 is the modified Bessel function of order zero and

$$\lambda = \sqrt{\frac{D\tau}{1 + \frac{V^2\tau}{4D}}} \tag{2}$$

Fig. 1 shows the rate $R(\mathbf{r}|\mathbf{x})$ over a square search area $\mathcal{A}$ spanning from $-L$ to $+L$, with $L = 15$, in both x and y direction.[1] The source parameters are: $Q_0 = 1$, $x_0 = -10$, $y_0 = 3$; the dispersion particle parameters $\tau = 250$ and $D = 1$ and the sensor size $a = 1$. Fig. 1(a) and (b) corresponds to $V = 1$ and $V = 0$, respectively.

The stochastic process of sensor encounters with emitted particles is modelled by a Poisson distribution: the probability that the sensor at location $\mathbf{r}$ encounters $z \in \mathbb{Z}^+$ particles during a time interval t_0 is then:

$$\mathcal{P}(z; \mu) = \frac{\mu^z}{z!} e^{-\mu} \tag{3}$$

where $\mu = R(\mathbf{r}|\mathbf{x})\, t_0$ is the mean concentration. Assuming that the size of the sensor a, as well as the environmental parameters τ, D and V, are known, then (3) represents the full specification of the likelihood function of $\mathbf{x}$, given a measurement z taken at location $\mathbf{r}$. Let us denote this likelihood by $g(z(\mathbf{r})|\mathbf{x})$.

2.2. Cognitive search formulation

Autonomous cognitive search strategies are formulated as POMDPs [19], whose elements are: the information state, the set of admissible actions and the reward function. We adopt the Bayesian framework for estimation of source parameter vector $\mathbf{x}$. The information state at discrete-time k is then represented by the posterior PDF $p(\mathbf{x}|\mathcal{T}_{1:k})$, where $\mathcal{T}_{1:k} = \{z_i(\mathbf{r}_i)\}_{1 \leqslant i \leqslant k}$ is the trace of visited locations with collected sensor measurements. After collecting the next measurement z_{k+1} at location r_{k+1}, the posterior PDF is updated according to the Bayes rule:

$$p(\mathbf{x}|\mathcal{T}_{1:k+1}) = \frac{g(z_{k+1}(\mathbf{r}_{k+1})|\mathbf{x})\ p(\mathbf{x}|\mathcal{T}_{1:k})}{p(z_{k+1}(\mathbf{r}_{k+1})|\mathcal{T}_{1:k})}, \tag{4}$$

where

$$p(z_{k+1}(\mathbf{r}_{k+1})|\mathcal{T}_{1:k}) = \int g(z_{k+1}(\mathbf{r}_{k+1})|\mathbf{x})\ p(\mathbf{x}|\mathcal{T}_{1:k}) d\mathbf{x}. \tag{5}$$

The initial (prior) PDF at $k = 0$, i.e. $\pi(\mathbf{x}) \equiv p(\mathbf{x}|\emptyset)$ is assumed known.

The admissible set of actions at time k, when the searcher is at location $\mathbf{r}_k$, consists of the applicable motion controls and is denoted by $\mathcal{V}_k$. For convenience, similar to [11], we consider a one-step ahead searcher motion control along the square lattice of mesh size v. Two admissible sets of actions are considered: $\mathcal{V}'_k = \{\cdot, \uparrow, \downarrow, \leftarrow, \rightarrow\}$, and $\mathcal{V}''_k = \{\cdot, \uparrow, \nearrow, \rightarrow, \searrow, \downarrow, \swarrow, \leftarrow, \nwarrow\}$. The searcher can move only inside the search area $\mathcal{A}$. Consequently, if the current searcher location $\mathbf{r}_k$ is on the edge of $\mathcal{A}$, the admissible set of actions will contain a reduced set of options compared to $\mathcal{V}'_k$ or $\mathcal{V}''_k$.

[1] All physical quantities are in arbitrary units (a.u.).