European Journal of Operational Research 215 (2011) 218–226

Contents lists available at ScienceDirect

European Journal of Operational Research

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

Innovative Applications of O.R.

A sequential perspective on searching for static targets

Kurt E. Wilson, Roberto Szechtman *, Michael P. Atkinson

Operations Research Department, Naval Postgraduate School, Monterey, CA 93943, United States

ARTICLE INFO

Article history:
Received 4 January 2010
Accepted 25 May 2011
Available online 13 June 2011

Keywords:
Applied probability
OR in military
Sequential analysis

ABSTRACT

We present a sequential approach to detect static targets with imperfect sensors, which range from tower-mounted cameras to satellites. The scenario is operationally relevant to many military, homeland security, search and rescue, environmental engineering, counter-narcotics, and law enforcement applications. The idea is to stop the search as soon as there is enough probabilistic evidence about the targets' locations, given an operator-prescribed error tolerance, knowledge of the sensors' parameters, and a sequence of detection signals from the sensors. By stopping the search as soon as possible, we promote efficiency by freeing up sensors and operators to perform other tasks. The model we develop has the added benefits of decreasing operator workload and providing negative information as a search progresses.

Published by Elsevier B.V.

1. Introduction

Today's operational planners and sensor operators face numerous challenges inherent to the complex environments that shape their space of operation. These challenges are further magnified by scarce resources, imperfect information, and operator task overload.

The time critical nature of the command decisions that serve as milestones throughout the Find, Fix, Track, Target, and Engage (F^2T^2E) process further exacerbate the situation. Defense planners must strive to develop and incorporate new, efficient procedures to allocate scarce resources in many different complex environments. Any efficiency gained within the F^2T^2E chain, however small, may have a compound effect over time on overall operational readiness because this will free up assets to perform other time-sensitive, critical sensing actions, as well as decrease operator workload and capitalize upon negative information. Such negative information could be utilized to find where targets are *not* located, and may help determine areas to set up certain operations or paths through the environment that are free of hostile forces.

In this article we consider a scenario with multiple fixed-sensors and multiple static targets in discrete-time and discrete-space. The sensors may range from tower-mounted cameras, to Unmanned Aircraft Systems (UASs), to satellites, and the targets under consideration do not react to any sensing action. The scenario is operationally relevant to many military, homeland security, search and rescue (SAR), environmental engineering, counter-narcotics, and law enforcement applications. UASs have been used in Iraq and Afghanistan to search for Improvised Explosive Devices (IEDs), insurgent safe houses, suspected weapons caches, and mortar points of origin [11,12]. Other relevant applications include searching for downed aircraft or life rafts, detecting illegal drug harvesting and processing operations, patrolling border infiltration points, and tracking flora and fauna counts in biological environments.

We formulate a model to locate static targets of interest (TOIs) hidden within an area of interest (AOI). As in reality, our model allows the analyst to contend with the fact that the search sensors are imperfect; i.e., the sensors may declare fewer or more targets than are actually present on a particular search attempt. The idea is to stop the search as soon as there is enough probabilistic evidence about the TOIs' locations, given an operator-prescribed error tolerance, knowledge of the sensors' parameters, and a sequence of detection signals from the sensors.

The AOI for the scenario is comprised of a grid of discrete, non-overlapping area-cells (ACs). The area-cells might be defined by geo-political borders, terrain features, or some arbitrary grid system of tactical significance to the operator, and need not be uniform in size nor shape. Each cell is characterized by the number of sensors (known), the sensors' operational parameters (known), and the number of targets (unknown).

The sensor parameters are the conditional probabilities of returning each possible detection signal given each possible number of actual TOIs in that area-cell. More specifically, when the operator makes an investigation into an area-cell, the sensor returns a detection signal corresponding to the number of TOIs seen by the sensor with some probability that depends on the actual

* Corresponding author. Tel.: +1 831 656 3311.
E-mail addresses: kurt.e.wilson@afghan.swa.army.mil (K.E. Wilson), rszechtm@nps.edu (R. Szechtman), mpatkins@nps.edu (M.P. Atkinson).

0377-2217/$ - see front matter Published by Elsevier B.V.
doi:10.1016/j.ejor.2011.05.045

(i.e., the *ground truth*, which is unknown) number of TOIs in that area-cell.

To efficiently determine the targets' locations subject to the operator-prescribed error tolerance, we develop a sequential eliminating procedure [14]. A sequential eliminating procedure attempts to isolate, from among several candidate configurations, one particular desired configuration—the objective. During a particular stage of a sequential eliminating procedure, all candidate configurations are examined and ranked in order of their likelihood of producing the sequence of observed signals up to that stage.

Any configuration whose likelihood, when compared with the configuration of maximum likelihood, exceeds a particular threshold (which depends on the user's error tolerance) is permanently eliminated from the set of candidates. If no configuration exceeds the threshold during a particular stage, then all those configurations remain in the set of candidates. The procedure advances to the next stage, using the updated candidate set. The process continues until only one configuration remains in the candidate set, and that configuration is declared the winner. In our case, the configurations are the ways the TOIs can be located in the area-cells of the AOI, and the winner is the *determined configuration*. We designate the actual location of the targets in the area-cells to be the *ground truth configuration* (GTC).

Search theory [1] traces its roots to the pioneering work of Koopman [5]. For the search scenario we focus on, the objective is to locate targets within a finite number of cells [1]. In this case, searcher success is achieved by either detecting the targets, or, if the targets are not detected, by correctly guessing the cells containing the targets. Tognetti [17] and Kadane [4] treat the scenario of whereabouts search against a stationary target. Washburn [19] is a classical reference in search theory.

Siegmund's 1985 book [14] is the classic reference in sequential analysis, and deals primarily with sequential hypothesis testing and related problems of estimation. In many of these cases, a fixed-sample solution exists, but one can employ sequential methods to achieve greater efficiency in the solution. Siegmund presents a sequential test with the same power as a fixed-sample test and requires fewer observations [14]. Therefore, the sequential test has a reasonable claim to be regarded as more efficient [14].

While the work to date in selection using sequential eliminating procedures [9,10,20] has focused on isolating the *best* system – usually the one with a maximum unknown parameter value – our goal is instead to isolate one determined configuration. The desire is for the determined configuration to be the ground truth configuration. That is it correctly specifies the number of TOIs in each area-cell. We show that our sequential model provides determined configurations efficiently, while guaranteeing to meet the user-prescribed error tolerance.

Compared to existing sensor employment models (e.g. [6,13,7,15]), our approach does not consider moving targets, does not dynamically allocate the sensors (i.e., no decision is taken as to where the sensors look in each stage), and does not find optimal search paths (with or without restrictions on searching area-cells within a vicinity of the last searched area-cell). Delving into the last point, most recent models (e.g. [6,13,7]) employ optimization techniques (deterministic or stochastic) that yield search paths that are optimal in a certain sense (often, but not always, maximizing the expected number of detected targets) for a prescribed number of time periods or search effort. However, these models scale poorly and become intractable even for a relatively small number of TOIs, area-cells, and time periods under consideration. This occurs because the computational cost grows at least exponentially in the number of variables ([18]), which generally is #TOIs $\times$ #ACs $\times$ #time periods. Some heuristics (e.g. [16]) have been proposed to overcome this difficulty, but their performance cannot be theoretically guaranteed. While our approach does provide some benefits over existing methods, it too has limitations with regard to the the size of problems it can be applied to. The algorithm can only be used in situations where there are a handful of TOIs (e.g. less than 5) because it must consider all possible configurations of TOIs in the area cells. There are many applications where this is the case (e.g. SAR scenarios, searching for an insurgent safe house) and our algorithm would provide an appropriate and effective approach.

Once again, this article presents a sequential perspective on imperfect sensor employment applicable to static targets that is easy to implement, is computationally tractable when there are a small number of TOIs for a larger class of problems than the optimization approaches currently being employed, and stops when the user prescribed probabilistic guarantees are met (and thus the number of time periods or search effort is an output of the model).

The article is organized as follows. Section 2 introduces the notation and key definitions. In Section 3 we present the sequential approach. Section 4 shows numerical illustrations of the model, and Section 5 closes the paper with the main conclusions.

2. Notation

In this section we introduce the notation that will be employed throughout this article.

A: Number of area-cells.
M: Number of targets of interest.
m_i: Number of targets of interest in area-cell i. This value is unknown, and is what the analyst wishes to determine for $i = 1,\ldots,A$.
m: The true configuration, $m = (m_1,m_2,\ldots,m_A)$.
$\mathcal{C}$: The set of feasible target configurations, formed by the elements $t = (t_1,\ldots,t_A)$ non-negative and integer such that $\sum_{i=1}^{A} t_i = M$.
K: The number of feasible target configurations $K = |\mathcal{C}| = \binom{M+A-1}{M}$.
$p_i(d|t_i)$: Conditional probability that the sensor in area-cell i returns a signal "d targets" given that t_i TOIs are present there.
S_i: The sensor present in AC_i is completely characterized by the $(M+1) \times (M+1)$ matrix S_i. The value in the tth row and dth column of S_i is the probability $p_i(d|t)$. The matrix S_i is stochastic, so the elements of each row constitute a probability mass function.
$X_{i,1},X_{i,2},\ldots$: Sequence of signals returned by the sensor in AC_i, independently and identically distributed (IID) with probability mass function $p_i(\cdot|m_i)$.
$\ell(x_{i,1},\ldots,x_{i,n}\,;\,t_i)$: For n IID signals from the sensor in AC_i, $X_{i,1} = x_{i,1}, X_{i,2} = x_{i,2},\ldots,X_{i,n} = x_{i,n}$, the likelihood of having t_i targets in area-cell i is

$$\ell(x_{i,1},\ldots,x_{i,n};t_i) = \prod_{j=1}^{n} p_i(x_{i,j}|t_i).$$

3. Model

We describe the model in more detail in section 3.1. Our objective is to analyze the scenario of multiple targets in an area of interest with many cells. This produces many potential configurations. This case is difficult to analyze because knowledge about the presence/absence of targets in an area-cell yields insight about the presence/absence of targets in other area-cells; i.e., the number of targets in each area-cell is not independent. To gain insight into the problem, in Section 3.2 we examine the scenario where there are