Systems & Control Letters 75 (2015) 69–76

Contents lists available at ScienceDirect

Systems & Control Letters

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

Performance bounds for mismatched decision schemes with Poisson process observations

Chetan D. Pahlajani [a,*], Jianxin Sun [b], Ioannis Poulakakis [b], Herbert G. Tanner [b]

[a] Mathematics, Indian Institute of Technology Gandhinagar, Chandkheda, Ahmedabad 382424, India
[b] Department of Mechanical Engineering, University of Delaware, Newark, DE 19716, USA

ARTICLE INFO

Article history:
Received 16 February 2014
Received in revised form
10 November 2014
Accepted 13 November 2014
Available online 4 December 2014

Keywords:
Decision making
Poisson process
Model mismatch
Nuclear detection

ABSTRACT

This paper develops a framework for analyzing the performance loss in fixed time interval decision algorithms that are based on observations of time-inhomogeneous Poisson processes, when some parameters characterizing the observation process are not known exactly. Key to the development is the formulation of an analytically computable performance metric which can be used in lieu of the true, but intractable, error probabilities. The proposed metric is obtained by identifying analytical upper bounds on the error probabilities in terms of the uncertain parameters. Using these tools, it is shown that performance degrades gracefully as long as the true values of the parameters remain within a neighborhood of the nominal values used in decision making. The results find direct application to problems of detecting illicit nuclear materials in transit.

1. Introduction

Many physical processes of interest are characterized by sequences of discrete events occurring randomly in time, modeled mathematically as *point processes* [1,2]; an important class of the latter is the collection of Poisson processes, used to capture the underlying physics in queueing theory [1], optical communications [3], neuroscience [4], and nuclear detection [5]. Problems of decision making between two hypotheses on the basis of Poisson (and more general point) process observations have long been studied [1,2,6–9]. For the Poisson case, the optimal Neyman–Pearson rule is known to be given by a Likelihood Ratio Test (LRT), where the decision is based on comparing a likelihood ratio formed by the observations against a suitable threshold. The functional form of the likelihood ratio is determined by the intensities of the Poisson process under the two hypotheses.

In many situations, however, these intensities are subject to uncertainty due to incomplete knowledge of the model. For instance, Poisson process intensities under the two hypotheses may be specified in terms of a parameter vector whose exact value – assumed to be the same under both hypotheses – may not be accurately known. One approach to ensuring acceptable performance of decision algorithms over a *range* of parameter values is to apply robust techniques [7,10,11]. Then, to identify the parameters most crucial for robustness, one needs to understand the relative impact of parameter uncertainty on decision-making performance. The challenge now is that performance is measured by error probabilities, the analytical computation of which is extremely difficult, if not impossible. This observation sets the stage for the present research, which aims at establishing an alternate analytically tractable performance metric which can shed light on the above problem. We note that although the setting described is, at first glance, reminiscent of composite hypothesis testing [11], there are some subtle differences. Indeed, the parameter vector in the framework above takes the *same*, albeit imperfectly known, value under both hypotheses; Remark 1 describes how this is a natural assumption in some instances. In contrast, composite hypothesis testing typically assumes that the parameter vector takes *different* values under the two hypotheses in disjoint subsets of parameter space.

The mathematical models and techniques described above find natural application in the field of *nuclear detection*. A particularly challenging instance of the problem of nuclear detection is that of detecting illicit Special Nuclear Material (SNM) in transit [5,12–14]. Assuming that the moving target is identified, one is asked to decide whether that target is a carrier of an SNM radiation source, using radiation count data from a spatially dispersed network of inexpensive Geiger counters or scintillators. The critical question is whether the photons recorded by the counters are solely due

* Corresponding author.
E-mail addresses: cdpahlajani@iitgn.ac.in (C.D. Pahlajani), jxsun@udel.edu (J. Sun), poulakas@udel.edu (I. Poulakakis), btanner@udel.edu (H.G. Tanner).

http://dx.doi.org/10.1016/j.sysconle.2014.11.006
0167-6911/

to ubiquitous background radiation or whether they also contain emissions from a moving source. Since both background and source photon arrivals at a sensor can be modeled by Poisson processes, one is faced with a problem of detecting a Poisson signal buried inside another signal of similar nature and magnitude, within a small time interval. Furthermore, one of these processes is actually time-inhomogeneous, since the perceived source intensity incident at a sensor varies with the inverse square of the distance between source and sensor [5].

This decision problem has been studied in a fixed interval framework [15,16], i.e., when data is collected by sensors over a fixed time interval, at the end of which a decision is made. The likelihood ratio has been identified in terms of the problem parameters, including the motion of the source [15]. Chernoff upper bounds [17–19] on the error probabilities for the corresponding LRT have been computed [16], identifying the analytical dependence of the bounds on the problem parameters. To fully exploit these insights in a field-deployable nuclear detector network system, however, one needs to recognize and account for the presence of model uncertainty, a dominant source of which is *radiation clutter* [13]: the myriad "nuisance sources" and spatiotemporal environmental variations whose cumulative effect is to create a dynamic and imperfectly modeled background.

In this paper, we study the effect of imperfectly known intensities on a class of decision problems for Poisson processes which include, as special cases, several scenarios encountered in nuclear detection. Working with a parametrized family of models, where each value of the parameter vector corresponds to a specific choice of intensities, we obtain Chernoff upper bounds on the error probabilities for decision schemes with *mismatch* [20,21]. By the latter, we mean that the decision rule is an LRT based on some nominal model which may be different from the true model governing the stochastic processes of interest. The Chernoff bounds, or equivalently, the exponents in the bounds, now furnish a performance measure which can be analytically characterized in terms of the problem parameters under the true and nominal models. Further, the exponents are seen to vary smoothly when the true model is a sufficiently small perturbation about the nominal one, implying that at least locally, performance degrades gracefully as parameters deviate from their nominal (known) values.

2. Background

We start with a binary hypothesis testing problem. The probabilistic setup consists of a measurable space $(\Omega, \mathscr{F})$ supporting a k-dimensional counting process $\boldsymbol{N}_t \triangleq (N_t(1), \ldots, N_t(k))$, $t \in [0, T]$, together with probability measures $\mathbb{P}_0$ and $\mathbb{P}_1$, with $\mathbb{P}_1 \ll \mathbb{P}_0$, i.e., $\mathbb{P}_1$ is absolutely continuous with respect to $\mathbb{P}_0$. Here, $\mathbb{P}_j$ denotes the probability measure under hypothesis H_j, $j \in \{0, 1\}$. We assume that the components $N_t(i)$, $i \in \{1, \ldots, k\}$, of $\boldsymbol{N}_t$ are independent Poisson processes under each $\mathbb{P}_j$, $j \in \{0, 1\}$, having intensity $\beta_i(t)$ with respect to $\mathbb{P}_0$, and intensity $\beta_i(t)+\nu_i(t)$ with respect to $\mathbb{P}_1$. The functions $\beta_i(\cdot)$ and $\nu_i(\cdot)$ are assumed to be positive, continuous, and bounded away from zero. The problem is to decide, based on the observed sample path of $\boldsymbol{N}_t$ over $t \in [0, T]$, between hypotheses H_0 and H_1.

Let $\mu_i(t)$ be the ratio of intensities for $N_t(i)$ under hypothesis H_1 versus H_0, i.e., $\mu_i(t) \triangleq 1 + \nu_i(t)/\beta_i(t)$. With $(\tau_n(i) : n \geq 1)$ denoting the jump times of $N_t(i)$, and after defining

$$L_t(i) \triangleq \exp\left(-\int_0^t \nu_i(s)ds\right) \prod_{n=1}^{N_t(i)} \mu_i(\tau_n(i)), \tag{1}$$

let $\{L_t : t \in [0, T]\}$ be the stochastic process

$$L_t \triangleq \prod_{i=1}^{k} L_t(i). \tag{2}$$

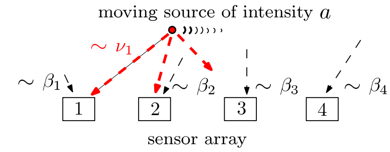

Fig. 1. Setup for a basic networked fixed-interval moving source detection scenario. Sensors are indexed by $\{1, 2, \ldots\}$ and receive photons that can be attributed either to background (thin dashed arrows) or to source radiation (thick red dashed arrows). Background intensity at sensor i location is characterized by β_i, and the intensity of the source is determined by the parameter a. The intensity of this source ν_i, as perceived at a sensor i, depends on the distance between sensor and source, r_i. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

By convention, $\prod_{n=1}^{0}(\cdot) = 1$. The optimal Neyman–Pearson test for deciding between H_0 and H_1 is an LRT given by comparing L_T to a suitably chosen threshold $\gamma > 0$, deciding H_1 if $L_T \geq \gamma$, and H_0 if $L_T < \gamma$ [15]. The performance of the LRT can be measured in terms of the corresponding error probabilities; that is, the probability of false alarm $P_F \triangleq \mathbb{P}_0(L_T \geq \gamma)$ and the probability of miss $P_M \triangleq \mathbb{P}_1(L_T < \gamma)$. More often than not, computing P_F and P_M is analytically intractable, thereby motivating the need for easily computable upper bounds that can be used as proxies for the corresponding probabilities at the expense of some sharpness. It can be shown [16] that if one explicitly computes

$$\Lambda(p) \triangleq \log \mathbb{E}_0[L_T^p] = \sum_{i=1}^{k} \int_0^T \left[\mu_i(s)^p - p\mu_i(s) + p - 1\right] \beta_i(s)ds,$$

for $p \in \mathbb{R}$, then P_F and P_M admit the Chernoff bounds

$$P_F \leq \exp\left(\inf_{p>0} [\Lambda(p) - p\log\gamma]\right), \quad P_M \leq \exp\left(\inf_{p<1} [\Lambda(p) + (1-p)\log\gamma]\right). \tag{3}$$

The availability of the bounds (3) in analytical form greatly facilitates the implementation of the test in many practical situations. For example, these bounds can be used [16] to devise a procedure for selecting the threshold γ so that the LRT $\{L_T \geq \gamma\}$ conforms with desired performance requirements, typically characterized by the probability of false alarm P_F being less than or equal to a desired level α.

To motivate the general treatment which follows, we begin with a concrete example of using the framework described above to detect a moving nuclear source (see Fig. 1). At the initial time $t = 0$, a moving vehicle (target) which may be a source of SNM with minimum activity $a > 0$, is identified. The target's trajectory over a fixed time interval $[0, T]$ is assumed to be known. This target is within sensing range of a spatially dispersed network of k radiation sensors, some of which may be mobile. For $i \in \{1, \ldots, k\}$, $N_t(i)$ represents the number of counts registered at sensor i up to and including time $t \in [0, T]$, while $\beta_i(t)$ represents the intensity at time t due to background at the spatial location of sensor i. In keeping with the inverse square fall-off with distance for source intensity – as is common in the relevant literature [5] – we take $\chi > 0$ as a sensor-specific cross section coefficient and $r_i(t)$ to be the distance at time t between the target and sensor i, and define the *perceived* source intensity at sensor i at time t as

$$\nu_i(t) = \frac{\chi a}{2\chi + r_i(t)^2}. \tag{4}$$

The goal is to decide, at the fixed time T, whether the counts recorded by the collection of sensors correspond solely to background radiation (hypothesis H_0), or whether they also contain