Quickest Search and Learning over Multiple Sequences

Javad Heydari and Ali Tajer
Electrical, Computer, and Systems Engineering Department
Rensselaer Polytechnic Institute
Troy, NY 12180

***Abstract*—Consider a set of random sequences, each consisting of independent and identically distributed random variables. Each sequence is generated according to one of the two possible distributions F_0 or F_1 with unknown prior probabilities $(1-\epsilon)$ and ϵ, respectively. The objective is to design a sequential decision-making procedure that identifies a sequence generated according to F_1 with the fewest number of measurements. Earlier analyses of this search problem have demonstrated that the optimal design of the sequential rules strongly hinge on the exact value of ϵ. Such information, however, might not be available in certain applications, especially in anomaly detection where the anomalous sequences occur with unpredicted patterns. Motivated by this premise, this paper designs a sequential inference mechanism that forms two coupled decisions for identifying a sequence of interest, and also learning the value of ϵ. The paper devises three strategies that place different levels of emphasis on each of these inference goals.**

I. Introduction

Quickest search over a set of data streams aims to identify one stream that exhibits desired statistical features in a real-time and data-adaptive fashion. Quickest search aims to strike a balance between the quality and the agility of the search process, and it arises in many application domains such as detecting free spectrum bands in spectrum sensing and monitoring computer networks for detecting faults or security breaches [1]. The significance of searching over data streams is expected to grow well into the future due to the advances in sensing and data acquisition technologies, which generate and process large volumes of data. Due to the variety of information sources and the large scale of the data sets, there might exist certain levels of uncertainty in the data models, which makes the design of the search mechanisms more challenging.

The problem of quickest search was first formalized and analyzed in [2] as an extension of sequential binary hypothesis testing [3] and [4]. In [2] it is assumed that *multiple* data streams are available such that they are generated according to one of the two known distributions F_0 and F_1 independently of each other, with *known* prior probabilities $(1-\epsilon)$ and ϵ, respectively. The ultimate goal of quickest search is to identify one sequence generated according to the desired distribution F_1 with the fewest number of measurements. Other variations of this search problem under different settings and objectives are also studied in [5]–[9]. In [2], [5] and [6] it is shown that the structure of the optimal decision rules for performing quickest search strongly depends on the known prior probability ϵ. In reality, however, this parameter might not be known, especially when F_1 captures anomalous behaviors with unpredicted patterns and rates of occurrence.

In this paper we formalize the quickest search problem over multiple data streams when ϵ is unknown, and analyze three specific problems. First we consider a purely sequential detection problem in which the objective is to identify one sequence generated according to F_1 with minimal delay, when facing uncertainty about ϵ. Next we consider a purely sequential estimation problem, in which the objective is to form a reliable estimate for ϵ with the fewest number of measurements made across the sequences. Finally, we consider a sequential problem that pursues both inference goals, in which a sequence generated according to F_1 is identified, and also a reliable estimate about ϵ is produced. In analyzing all these three problems, besides characterizing the stopping rules and the detectors and estimators, we also need to design an optimal information-gathering process, the role of which is to dynamically decide about abandoning a sequence and identifying the next sequence to make a measurement from.

Other studies of the quickest search problem under different settings and objectives relevant to the scope of this paper include the scanning problem, in which a finite number of sequences are available and *exactly one* sequence is generated according to the desired statistical feature, which is different from the model in this paper [9]–[12]. In another direction, the set of available sequences contains multiple sequences with the desired distribution and the goal is to identify *all* of them [13]–[15]. The sequential estimation of a single parameter by observing independent and identically distributed (i.i.d.) random variables is studied in [16]. It is further extended in [17] to a setting in which multiple unknown parameters are available, and at each time instance one of a finite set of actions can be taken for sampling. In [17] each action depends only on one of the unknown parameters, while [18] generalizes the results to the actions depending on common unknown parameters.

The remainder of the paper is organized as follows. Section II provides the data model and the sampling model, and formalizes the search problem of interest. The optimal inference rules are characterized in Section III, and the sampling procedures and the associated stopping rules for purely detection and purely estimation routines are characterized in Section IV and Section V, respectively. The combined inference procedure is provided and analyzed in Section VI.

This research was supported in part by the U. S. National Science Foundation under Grant ECCS-1455228 and the CAREER Award ECCS-1554482.

Section VII provides simulation results, and concluding remarks are provided in Section VIII.

II. Problem Formulation

A. Data Model

Consider an *ordered* set of n sequences $\{\mathcal{X}^1,\dots,\mathcal{X}^n\}$, where each sequence consists of independent and identically distributed (i.i.d.) real-valued observations $\mathcal{X}^i \triangleq \{X_1^i, X_2^i,\dots\}$. Each sequence, independently of the rest, is generated according to one of the two possible distributions, hence, obeying one of the two hypotheses

$$\begin{array}{ll} \mathsf{H}_0: & X_j^i \sim F_0 \\ \mathsf{H}_1: & X_j^i \sim F_1 \end{array} \quad \text{for } j=1,2,\dots\ , \tag{1}$$

where F_0 and F_1 denote cumulative distribution functions (cdfs) and are assumed to be known. The distribution F_0 captures the statistical behavior of the normal sequences, and the distribution of the abnormal sequences is F_1. We further assume that well-defined probability density functions (pdfs) corresponding to F_0 and F_1 exist and are denoted by f_0 and f_1, respectively. We assume that each sequence is abnormal with the prior probability $\epsilon \in (0,1)$. Specifically, by defining T_i as the *true* model underlying sequence $\mathcal{X}^i$, for $i \in \{1,\dots,n\}$ we have

$$\mathbb{P}(\mathsf{T}_i = \mathsf{H}_1) \ = \ \epsilon\ . \tag{2}$$

Furthermore, we assume that ϵ is known only stochastically with the continuous pdf g with support $[0,1]$.

B. Sampling Model

The objective of the search process is to identify *one* abnormal sequence with the fewest expected number of measurements. As shown in [2], [5], and [6], when ϵ is known precisely, the optimal decision rules strongly depend on the value of ϵ. Hence, the quickest search objective is strongly coupled with concurrently forming a reliable estimate for ϵ. The sampling procedure sequentially examines the sequences according to their order by taking one measurement at-a-time until sufficiently reliable detection and estimation decisions can be made. By denoting the index of the observed sequence and its sample at time $t \in \mathbb{N}$ by s_t and Y_t, respectively, we can abstract the information accumulated sequentially by the filtration $\{\mathcal{F}_t\ :\ t=1,2,\dots\}$ where $\mathcal{F}_t \triangleq \sigma(Y_1,\dots,Y_t)$. The sampling process starts from the first sequence, i.e., $s_1 = 1$, and at time t takes one of the following three actions based on the information accumulated up to time t, i.e., $\mathcal{F}_t$.

- A_1) *Decision:* stops taking more samples and declares one of the sequences observed up to time t as an abnormal one and forms an estimate for ϵ;
- A_2) *Observation:* due to lack of sufficient confidence to make a decision or form a reliable estimate for ϵ, one more sample is taken from the same sequence, i.e., $s_{t+1} = s_t$; or
- A_3) *Exploration:* sequence s_t is discarded and the sampling procedure switches to the next sequence and takes one observation from the new sequence, i.e., $s_{t+1} = s_t + 1$.

In order to formalize the sampling procedure we define τ as the stopping time of the procedure, that is the time instant at which action A_1 (decision) is performed. We denote the detection rule at the stopping time by $\delta(\tau) \in \{1,\dots,s_\tau\}$ and the estimate formed for ϵ by $\hat{\epsilon}(\tau)$. To characterize dynamic switching between observation and exploration actions we define the binary function $\psi : \{1,\dots,\tau-1\} \rightarrow \{0,1\}$ such that at time $t \in \{1,\dots,\tau-1\}$ if the decision is in favor of performing observation (A_2) we set $\psi(t) = 0$, while $\psi(t) = 1$ indicates a decision in favor of exploration (A_3). Hence, $\forall t \in \{1,\dots,\tau-1\}$:

$$\psi(t) = \left\{ \begin{array}{ll} 0 & \text{action A}_2 \\ 1 & \text{action A}_3 \end{array} \right. . \tag{3}$$

A sequential decision is completely characterized by the combination $\Phi \triangleq (\tau, \delta(\tau), \hat{\epsilon}(\tau), \psi(1),\dots,\psi(\tau-1))$.

C. Problem Formulation

The optimal search procedure can be found by determining Φ. The natural performance measures for evaluating the efficiency of any sampling strategy Φ include the quality of the final detection, which is captured by (i) the frequency of the erroneous decisions, i.e., $\mathsf{P}_\Phi(\tau) \triangleq \mathbb{P}(\mathsf{T}_{\delta(\tau)} = \mathsf{H}_0 | \mathcal{F}_\tau)$ and (ii) the mean squared error risk of estimation, i.e., $\mathsf{R}_\Phi(\tau) \triangleq \mathbb{E}[[\epsilon - \hat{\epsilon}(\tau)]^2 | \mathcal{F}_\tau]$, and (iii) the delay in reaching a decision, i.e., τ. By integrating these three figures of merit into one cost function, the stochastic aggregate cost function for a given Φ at time t is given by

$$\mathsf{J}_\Phi(t) \triangleq c_{\rm d} \cdot \mathsf{P}_\Phi(t) + c_{\rm e} \cdot \mathsf{R}_\Phi(t) + c_{\rm s} \cdot t\ , \tag{4}$$

where $c_{\rm d}$, $c_{\rm e}$, and $c_{\rm s}$ are positive constants that balance the quality and the agility of the search process. In the following sections, under different settings we characterize the stopping time, the switching rules, the final decision rules, and the associated performance guarantees

III. Optimal Inference Rules

In this section we start by characterizing the detection and estimation rules $\delta(\tau)$ and $\hat{\epsilon}(\tau)$ for any given stopping time τ, and switching sequence $\{\psi(t)\ :\ t=1,\dots,\tau-1\}$. Then, we focus on two special cases of the search problem, where one places the emphasis on the detection subproblem by setting $c_{\rm e} = 0$, and the other focuses on the estimation subproblem by setting $c_{\rm d} = 0$. Based on the insights from these special cases, we finally treat the problem in its general form.

It is shown in [19] that in a sequential setting, for any given sampling strategy and stopping rule there exists a fixed set of final decisions that are optimal. Furthermore, the quadratic estimation cost $\mathsf{R}_\Phi(t)$ is independent of the detection rule, and also the detection cost is independent of the estimate of the prior probability. Hence, for any given stopping time τ and sequence of switching functions $\{\psi(t)\ :\ t=1,\dots,\tau-1\}$, detection and estimation decision rules can be decoupled. The following theorem formalizes these results.

Theorem 1. *The sequential strategy for optimizing the cost function* (4) *can be decoupled in a way that the detection* *and estimation rules are decoupled, and both of them are independent of the stopping time and the switching rule.*

This theorem facilitates characterizing the detection and estimation decisions. To proceed, we define π_t^i as the posterior probability that the sequence $\mathcal{X}^i$ is abnormal given the information up to time t, i.e., $\pi_t^i \triangleq \mathbb{P}(\mathsf{T}_i = \mathsf{H}_1 \mid \mathcal{F}_t)$. By defining $\kappa_t^i(\epsilon) \triangleq \mathbb{P}(\mathsf{T}_i = \mathsf{H}_1 \mid \mathcal{F}_t, \epsilon)$, it can be readily verified that

$$\kappa_{t+1}^i(\epsilon) = \begin{cases} \frac{\kappa_t^i(\epsilon) f_1(Y_{t+1})}{\kappa_t^i(\epsilon) f_1(Y_{t+1}) + (1-\kappa_t^i(\epsilon)) f_0(Y_{t+1})} & \text{if } s_t = i \\ \kappa_t^i(\epsilon) & \text{if } s_t \neq i \end{cases} ,$$

and $$\pi_t^i = \int \kappa_t^i(\epsilon) g(\epsilon \mid \mathcal{F}_t)\, d\epsilon ,$$

where $g(\epsilon \mid \mathcal{F}_t)$ is the posterior pdf of ϵ. Also, at any time t, we denote the posterior average of ϵ by

$$\hat{\epsilon}(t) \triangleq \mathbb{E}[\epsilon \mid \mathcal{F}_t] . \tag{5}$$

Based on these definitions, the next theorem provides the optimal decision rules for any given stopping time and switching sequence.

Theorem 2 (Decision Rules)**.** *For a given stopping time τ and switching sequence:*

1) *The optimal detection rule, which minimizes $\mathsf{J}_\Phi(\tau)$, is*

$$\delta(\tau) = \arg\max_{i \in \{s_1, \ldots, s_\tau\}} \pi_\tau^i . \tag{6}$$

2) *The optimal estimation rule, which minimizes $\mathsf{J}_\Phi(\tau)$, is*

$$\hat{\epsilon}(\tau) = \mathbb{E}[\epsilon \mid \mathcal{F}_\tau] . \tag{7}$$

Given the decision rules in Theorem 2, we next characterize the optimal stopping time τ, and the associated optimal switching rules $\{\psi(t) : t = 1, \ldots, \tau - 1\}$. For this purpose, we first consider the sequential search and sequential estimation problems in two different settings. By leveraging the insight gained from these special cases we solve the joint search and estimation problem.

IV. Emphasis on Detection ($c_{\rm e} = 0$)

We first consider a purely sequential detection setting, in which the estimation quality is unintegrated by setting $c_{\rm e} = 0$. This problem when ϵ is known and the objective is to minimize the *average* delay is studied in [2], [6] and [7], where the analyses provide the optimal stopping time and switching rules. In this section, we provide stopping and switching rules that enjoy certain optimality guarantees and facilitate generalizing the results to the general case of $c_{\rm e} \neq 0$. To this end, we propose a stopping and switching rules, the combination of which accepts asymptotic optimality guarantees. Specifically, we define

$$\tau_{\rm d}^* \triangleq \inf \left\{ t : \max_i \; \pi_t^i \geq 1 - c_{\rm sd} \right\} , \tag{8}$$

where $c_{\rm sd} \triangleq \frac{c_{\rm s}}{c_{\rm d}}$, as the stopping time of the sampling process. According to this stopping rule, the sampling process continues until one is confident enough that one of the sequences is generated by F_1.

Next, we characterize the switching rule prior to the stopping time in order to dynamically decide between the *exploration* and *observation* actions. Specifically, at any time $t \in \{1, \ldots, \tau_{\rm d}^* - 1\}$ we set the switching rule to discard sequence s_t and switch to $s_t + 1$ when $\pi_t^{s_t} < \hat{\epsilon}(t)$. Hence,

$$\psi_{\rm d}(t) = \begin{cases} 1 & \text{if } \pi_t^{s_t} < \hat{\epsilon}(t) \\ 0 & \text{if } \pi_t^{s_t} \geq \hat{\epsilon}(t) \end{cases} . \tag{9}$$

This switching rule, when combined with the stopping time given in (8), achieves the asymptotic *pointwise* optimality (APO), as formalized in the following theorem.

Theorem 3. *Consider a sequential strategy $\Phi_{\rm d}$ with the detection and estimation rules given in* (5) *and* (6)*, the stopping time defined in* (8)*, and the switching rule in* (9)*. For any other sampling strategy $\hat{\Phi}$, we have*

$$\lim_{c_{\rm sd} \to 0} \mathbb{P}\left\{ \frac{\mathsf{J}_{\Phi_{\rm d}}(\tau_{\rm d}^*)}{\mathsf{J}_{\hat{\Phi}}(\hat{\tau})} \leq 1 + \Delta \right\} = 1 , \quad \forall \Delta > 0 . \tag{10}$$

Proof: The proof can be carried out in two steps, which we briefly discuss. In the first step we assume that L sequences are visited during the sampling process and τ_ℓ for $\ell \in \{1, \ldots, L\}$ is the number of samples taken from sequence ℓ. It can be shown that τ_ℓ for $\ell \in \{1, \ldots, L-1\}$ are exponentially bounded and the delay is dominated by the number of samples taken from the sequence that will be declared as abnormal. By defining $D_{\rm KL}(f_1 \| f_0)$ as the Kullback-Leibler divergence between f_0 and f_1, in the second step we show that

$$\frac{1}{\tau_{\rm d}^*} \log \mathsf{P}_{\Phi_{\rm d}}(\tau_{\rm d}^*) \to -D_{\rm KL}(f_1 \| f_0) , \quad \text{as } c_{\rm sd} \to 0 , \tag{11}$$

which can be concluded from the first step and the fact that

$$\frac{1}{\tau_L} \log \mathsf{P}_{\Phi}(\tau_L) \to -D_{\rm KL}(f_1 \| f_0) , \quad \text{as } c_{\rm sd} \to 0 . \tag{12}$$

Then, the optimality property is concluded from [16]. ■

V. Emphasis on Estimation ($c_{\rm d} = 0$)

The problem of sequential estimation from one sequence is studied in [16], and the extension to controlled sequential estimation is studied in [17] and [18]. In this section we show how the results of [18] can be leveraged to solve the sequential estimation of the quickest search problem. To this end, we denote the Fisher information corresponding to each switching rule $\psi(t) = \ell$, for $\ell \in \{0, 1\}$, by

$$I_\ell(\epsilon) \triangleq \mathbb{E}\Big[\frac{-\partial^2}{\partial \epsilon^2} \mathsf{LL}(Y_{t+1} | \epsilon, \psi(t) = \ell) \Big] , \tag{13}$$

where $\mathsf{LL}(Y_{t+1} | \epsilon, \psi(t) = \ell)$ is the log-likelihood of observing Y_{t+1} at time $t+1$ when $\psi(t) = \ell$. Based on this definition, we characterize a stopping time and a switching rule that achieve asymptotic optimality when $c_{\rm se} \triangleq \frac{c_{\rm s}}{c_{\rm e}}$ tends to zero. Specifically, we define the stopping time as

$$\tau_{\rm e}^* \triangleq \inf \Big\{ t \; : \; \mathsf{R}_\Phi(t) \leq t \cdot c_{\rm se} \Big\} . \tag{14}$$

According to this stopping rule, when the estimation cost $\mathsf{R}_\Phi(t)$ falls below the total sampling cost $t \cdot c_{\rm se}$ the sampling process terminates. Furthermore, the switching rule should select the action between *observation* and *exploration* that minimizes the estimation variance, which according to the Cramer-Rao bound is lower bounded by the inverse of the Fisher information value. Hence, we first compute $\hat{\epsilon}_{\rm ML}(t)$ as the maximum likelihood (ML) estimate of ϵ at time t based on the observations up to time t. Then, we select the action with the largest Fisher information value for that estimate, i.e.,

$$\psi_{\rm e}(t) = \begin{cases} 1 & \text{if } I_1(\hat{\epsilon}_{\rm ML}(t)) > I_0(\hat{\epsilon}_{\rm ML}(t)) \\ 0 & \text{if } I_1(\hat{\epsilon}_{\rm ML}(t)) \leq I_0(\hat{\epsilon}_{\rm ML}(t)) \end{cases} . \tag{15}$$

This switching rule ensures that the sampling process takes the action that minimizes the variance of estimation. While ignoring the impact of the current decision on the future ones, it can be shown that in the large sample regimes, it is asymptotically optimal.

Theorem 4. *Let $\Phi_{\rm e}$ be the sampling strategy characterized by the stopping time and the switching rule given in* (14) *and* (15)*, respectively. Then, when $c_{\rm se}$ approaches zero, for any other sampling strategy $\hat{\Phi}$ we have*

$$\lim_{c_{\rm se}\to 0} \mathbb{P}\left\{ \frac{\mathsf{J}_{\Phi_{\rm e}}(\tau^*_{\rm e})}{\mathsf{J}_{\hat{\Phi}}(\hat{\tau})} \leq 1+\Delta \right\} = 1 \ , \quad \forall \Delta > 0 \ . \tag{16}$$

This switching rule ensures that a sufficient number of samples is taken from the current sequence so that its distribution is distinguishable before switching to the next sequence.

Remark 1. *When distributions F_0 and F_1 are distinguishable enough, the switching is more inclined to explore more sequences and have a more reliable estimate of ϵ. On the other hand, if F_0 and F_1 are less-distinguishable, the switching rule tends to continue taking samples from the same sequence in order to be more confident about its distribution before switching to the following sequence.*

VI. Balance Between Detection and Estimation

With the insights gained from the previous two sections, in this section we treat the quickest search problem of interest in its general form, which involves forming reliable decisions for both estimation and detection routines. We first characterize a stopping time by noting that in the detection problem the sampling process terminates when $(1-\max_i \ \pi^i_t)$ falls below $c_{\rm sd}$, i.e., the relative cost of one new sample, while in the estimation problem, it stops when the normalized estimation cost $\left(\frac{\mathsf{R}_\Phi(t)}{t}\right)$ is smaller than the relative cost of one new sample $c_{\rm se}$. Hence, for the general quickest search problem we define the stopping time as

$$\tau^* \triangleq \inf\left\{ t \ : \ c_{\rm d}(1-\max_i \ \pi^i_t) + c_{\rm e}\frac{\mathsf{R}_\Phi(t)}{t} \leq c_{\rm s} \right\} . \tag{17}$$

While it is a combination of the stopping rules in the previous settings, we will show that it can be also obtained directly by optimizing the total Bayesian cost given in (4).

Theorem 5. *Let Φ^* be the sampling strategy with the stopping time give in* (17) *and the optimal switching sequence $\{\psi^*(t) : t = 1,2,\dots\}$, and $\hat{\Phi}$ be any arbitrary sampling strategy with the same switching rule and any other stopping time $\hat{\tau}$. For all $\hat{\Phi}$ and $\hat{\tau}$ we have*

$$\lim_{c_{\rm sd},c_{\rm se}\to 0} \mathbb{P}\left\{ \frac{\mathsf{J}_{\Phi^*}(\tau^*)}{\mathsf{J}_{\hat{\Phi}}(\hat{\tau})} \leq 1+\Delta \right\} = 1 \ , \quad \forall \Delta > 0 \ . \tag{18}$$

Proof: The detection cost depends only on one sequence, while the estimation cost relies on all the observed sequences. When the sampling cost is substantially smaller than the costs of estimation and incorrect detection decisions, we know that

$$\frac{1}{t}\log \mathsf{P}_\Phi(t) \to -\gamma D_{\rm KL}(f_1 \| f_0) \ , \quad \text{as } t\to\infty \ , \tag{19}$$

$$\text{and} \ , \qquad t\mathsf{R}_\Phi(t) \to V(\epsilon) \ , \quad \text{as } t\to\infty \ , \tag{20}$$

for some $0 < V(\epsilon) < \infty$ and $\gamma \in (0,1)$. In case (19) and (20) are exact, i.e., they are true for any t, the stopping time would be the first time t for which we have

$$c_{\rm d}\mathsf{P}_\Phi(t-1) + c_{\rm e}\mathsf{R}_\Phi(t-1) + c_{\rm s}(t-1) \leq c_{\rm d}\mathsf{P}_\Phi(t) + c_{\rm e}\mathsf{R}_\Phi(t) + c_{\rm s}\cdot t \ , \tag{21}$$

which by replacing P_Φ and R_Φ from (19) and (20) we have

$$c_{\rm d}\mathsf{P}_\Phi(t)(1-\mathsf{P}_\Phi(t)^{\frac{1}{t}}) + c_{\rm e}\frac{\mathsf{R}_\Phi(t)}{t} \leq c_{\rm s} \ , \tag{22}$$

and since $(1-\mathsf{P}_\Phi(t)^{\frac{1}{t}}) \to 1$, (22) is equivalent to

$$c_{\rm d}(1-\max_i \ \pi^i_t) + c_{\rm e}\frac{\mathsf{R}_\Phi(t)}{t} \leq c_{\rm s} \ . \tag{23}$$

The remainder of the proof, which involves showing that for the asymptotic convergence in (19) and (20) we have asymptotic optimality, follows the same line of argument as in [16]. ∎

Next, we characterize the switching rules for dynamically deciding between *exploration* and *observation* actions. First, we note that, taking a new sample from any sequence reduces the average estimation cost, while the detection error probability depends on the number of samples taken from the sequence declared as the abnormal sequence. Also, taking many samples from one sequence cannot improve the estimation cost significantly. Hence, for the switching rule, at the beginning we apply the rule given in the purely estimation setting (15). When the estimation cost becomes sufficiently small, we apply the switching rule given in the purely detection setting (9). Based on this, we set the switching rule as follows:

$$\psi(t) = \begin{cases} \psi_{\rm e}(t) & \text{if } \frac{\mathsf{R}_\Phi(t)}{t} > c_{\rm se} \\ \psi_{\rm d}(t) & \text{if } \frac{\mathsf{R}_\Phi(t)}{t} \leq c_{\rm se} \end{cases} . \tag{24}$$

This switching rule, at the beginning of the sampling process, is more focused on forming a reliable estimate for ϵ. When the estimation cost is sufficiently small, it gradually shifts the focus to forming a reliable detection decision. The following theorem formalizes the asymptotic optimality properties of the sampling strategy characterized by the stopping rule in (17) and the switching rule in (24).

Theorem 6. *Let Φ^* be the sampling strategy with the stopping time give in* (17) *and a switching rule given in* (24)*, and $\hat{\Phi}$*

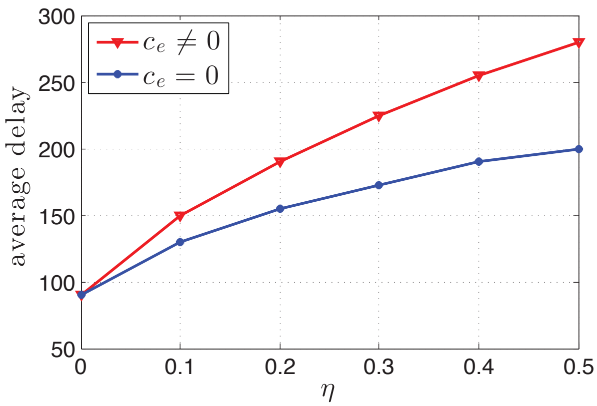

Fig. 1. Average delay versus η.

be any arbitrary sampling strategy with the stopping time $\hat{\tau}$. Then, in the asymptote of $\frac{c_s}{c_e} \to 0$, $\frac{c_s}{c_d} \to 0$, and $c_e = O(c_d)$ we have

$$\lim_{c_{sd}, c_{se} \to 0} \mathbb{P}\left\{ \frac{\mathsf{J}_{\Phi^*}(\tau^*)}{\mathsf{J}_{\hat{\Phi}}(\hat{\tau})} \leq 1 + \Delta \right\} = 1 \ , \quad \forall \Delta > 0 \ . \tag{25}$$

VII. Simulation

In order to evaluate the performance of the proposed adaptive strategies, in this section we compare the average delay of the quickest linear search process with and without emphasis on the estimation of the prior probability. For this purpose, we assume zero-mean Gaussian distributions with variance values 1 and 2 for F_0 and F_1, respectively. We also assume that prior probability has a uniform distribution over $(0.3 - \eta/2, 0.3 + \eta/2)$, where η is a constant in the range $\eta \in [0, 0.6]$. For a general setting with $c_s = 0.001$, $c_d = 0.1$, and $c_e = 1$, the average delay is compared with that of the setting in which $c_e = 0$ and have the same rate of erroneous detection decisions. Figure 1 shows that by increasing the uncertainty of the prior probability, more samples are required even when the estimation cost of the prior probability is not a concern. Also, it is observed that for reliable estimation of the prior probability (when $c_e \neq 0$), we require to take more measurements compared to the case that estimation cost is unintegrated from the total cost ($c_e = 0$).

VIII. Conclusion

We have analyzed the problem of quickest search and learning over multiple sequences, in which each sequence is generated according to one of the normal and abnormal distributions with unknown prior probabilities. The main objective is to identify one abnormal sequence with the fewest number of measurements, which is known to depend strongly on the unknown prior probability. Hence, achieving the detection objective, also necessitates producing a reliable estimate for the prior probabilities. For this purpose, we have characterized the optimal detection and estimation rules, and have designed asymptotically optimal sequential mechanisms that at each time dynamically decide which sequence should be sampled. First, we have considered a purely detection setting in which the estimation of the prior probability is not a concern and have shown that the decision rules reduce to comparing the posterior probability values with two thresholds. In the next setting, we have focused on the reliable estimation of the prior probability and have shown that the optimal procedure selects the sequence that maximizes the Fisher information value and stops when the cost of estimation falls below the total cost of sampling. Finally, we have combined the results of the first two settings to characterize the sampling strategy for the quickest search problem in its general form.

References

[1] A. Tajer, V. V. Veeravalli, and H. V. Poor, "Outlying sequence detection in large data sets: A data-driven approach," *IEEE Signal Processing Magazine*, vol. 31, no. 5, pp. 44–56, Sep. 2014.

[2] L. Lai, H. V. Poor, Y. Xin, and G. Georgiadis, "Quickest search over multiple sequences," *IEEE Transactions on Information Theory*, vol. 57, no. 8, pp. 5375–5386, Aug. 2011.

[3] A. Wald, "Sequential tests of statistical hypotheses," *The Annals of Mathematical Statistics*, vol. 16, no. 2, pp. 117–186, Jun. 1945.

[4] H. V. Poor and O. Hadjiliadis, *Quickest Detection*. Cambridge, UK: Cambridge University Press, Dec. 2009.

[5] J. Heydari, A. Tajer, and H. V. Poor, "Quickest linear search over correlated sequences," *IEEE Transactions on Information Theory*, vol. 62, no. 10, pp. 5786–5808, Oct. 2016.

[6] J. Heydari and A. Tajer, "Quickest linear search over correlated sequences," in *Proc. IEEE International Symposium on Information Theory*, Hong Kong, Hong Kong, Jun. 2015, pp. 765–769.

[7] A. Tajer and H. V. Poor, "Quick search for rare events," *IEEE Transactions on Information Theory*, vol. 59, no. 7, pp. 4462–4481, Jul. 2013.

[8] M. L. Malloy and R. D. Nowak, "Sequential testing for sparse recovery," *IEEE Transactions on Information Theory*, vol. 60, no. 12, pp. 7862–7873, Dec. 2014.

[9] Y. Li, S. Nitinawarat, and V. V. Veeravalli, "Universal sequential outlier hypothesis testing," in *Proc. IEEE International Symposium on Information Theory*, Honolulu, HI, Jun.-Jul. 2014, pp. 3205–3209.

[10] K. Zigangirov, "On a problem in optimal scanning," *Theory of Probability and Its Applications*, vol. 11, no. 2, pp. 294–298, Jan. 1966.

[11] V. Dragalin, "A simple and effective scanning rule for a multi-channel system," *Metrika*, vol. 43, no. 1, pp. 165–182, Dec. 1996.

[12] S. Nitinawarat and V. V. Veeravalli, "Universal scheme for optimal search and stop," in *Proc. Information Theory and Applications Workshop*, Feb. 2015, pp. 322–328.

[13] K. Cohen, Q. Zhao, and A. Swami, "Optimal index policies for anomaly localization in resource-constrained cyber systems," *IEEE Transactions on Signal Processing*, vol. 62, no. 16, pp. 4224–4236, Aug. 2014.

[14] S. Zou, Y. Liang, H. V. Poor, and X. Shi, "Nonparametric detection of anomalous data streams," *arXiv:1405.2294v2*, Dec. 2016.

[15] Y. Li, S. Nitinawarat, and V. V. Veeravalli, "Universal outlier hypothesis testing," *IEEE Transactions on Information Theory*, vol. 60, no. 7, pp. 4066–4082, Jul. 2014.

[16] P. J. Bickel and J. A. Yahav, "Asymptotically pointwise optimal procedures in sequential analysis," in *Proc. of the Fifth Berkeley Symposium on Mathematical Statistics and Probability, Volume 1: Statistics*. Berkeley, Calif.: University of California Press, 1967, pp. 401–413.

[17] V. J. Yohai, "Asymptotically optimal bayes sequential design of experiments for estimation," *The Annals of Statistics*, vol. 1, no. 5, pp. 822–837, Sep. 1973.

[18] G. Atia and S. Aeron, "Asymptotic optimality results for controlled sequential estimation," in *Proc. 51st Annual Allerton Conference on Communication, Control, and Computing*, Oct. 2013, pp. 1098–1105.

[19] K. J. Arrow, D. Blackwell, and M. A. Girshick, "Bayes and minimax solutions of sequential decision problems," *Econometrica*, vol. 17, no. 3/4, pp. 213–244, 1949.