Controlled Sensing for Multi-hypothesis Testing with Co-dependent Actions

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

Abstract—**Multi-hypothesis testing, which is widely used in many domains for discerning the true model governing the data, is often studied in a fixed sample-size setting. In such settings, the data-acquisition and decision-making processes are decoupled and the data-acquisition policies are pre-specified. Motivated by the advantages of sequential sampling, this paper treats the inherently coupled problems of data-acquisition and decision-making for multi-hypothesis testing, where data-acquisition can be abstracted as selecting one possible sensing action from a finite set. It aims to devise the quickest detection strategy by characterizing the minimum number of samples required to make a reliable decision as well as designing the dynamic attendant decision rules for selecting the best actions. The setting in which the available control actions are co-dependent is considered, which is a major distinction from the existing literature. Specifically, the existing data-adaptive approaches lose their optimality guarantees for this problem as they fail to account for such dependence. A novel sampling strategy that incorporates the dependence of the control actions into its decision rules is proposed, and its optimality properties are established.**

I. Introduction

Multiple hypothesis testing is a data-driven statistical approach to discerning the underlying statistical model of the data. Such testing problems are classical inference problems and have been investigated extensively in the literature [1] and [2]. They are often performed in frameworks with pre-specified data size, where the data is provided by an unknown data-acquisition process, and the objective is forming the most reliable decision. However, it is well understood that these tests can have often substantially lower sample complexity, and subsequently lower data-acquisition and computational complexity, when performed in *sequential* settings. In sequential settings, specifically, the data-size is data-driven, and the data is collected sequentially over time until a stochastic stopping time at which a reliable decision can be formed.

Sequential strategies work based on the premise that collecting the samples sequentially guides the data-acquisition process to progressively identify and take the most informative actions. The primary challenge in designing the sequential methods pertains to jointly co-designing the decision-making process in conjunction with the underlying data-acquisition process. The data acquisition process consists in two main components. The first component is the data-driven stopping time of the process, which determines the size of the data. The second one pertains to the process of collecting the data prior to the stopping time. This component manifests in network settings in which there exist multiple information sources each generating an information stream over time, where one faces selecting the sequence of information sources to collect data from over time. Throughout this paper we refer to the selection of the information sources over time as the control actions.

The settings in which the control actions are statistically independent over time is well-investigated [3]–[8], for which there exists a large body of decisions rules that are variations of, or inspired by, the original work of Chernoff [9] for *controlled sensing*. In many emerging domains, however, the statistical independence of the actions cannot be sustained. For instance, in networks of interconnected information sources, whose data bears statistical correlation due to an underlying physical interaction or coupling, the assumption that a control action at any time does not impact future ones is not necessarily valid. Representative application domains include sensor management [10], inspection and classification [11], medical diagnosis [12], search [13], channel coding with feedback [14], universal source coding [15], [16], and outlier detection [17].

Motivated by such scenarios, in this paper we consider a general setting for sequential hypothesis testing in which the control actions are statistically co-dependent. The problem is considered in a fully sequential setting in which, data is collected sequentially one-at-a-time until a reliable decision about the true model can be made. Furthermore, it is assumed that at each time instant, there exist a set of possible control actions to be taken. The goal is to design a sequential and data-adaptive strategy that consists of a causal control policy which is a function of all the previous actions and collected data. Such strategies are specified by the number of samples to be collected, as well as the actions to be taken at each time. When there exists only one action, determining the optimal sampling strategy reduces to minimizing the (average) number of samples. This can be effectively facilitated via sequential hypothesis testing, which is well-investigated [18]–[22].

However, incorporating control actions into the decisions is less-investigated. One directly applicable approach to treat such coupled sampling and decision-making processes is *controlled sensing*, originally developed by Chernoff for binary composite hypothesis testing under independent control actions [9]. Chernoff's rule decides in favor of the action with the best *immediate* return according to proper information measures, and achieves optimal performance in the asymptote

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

978-1-5386-4780-6/18/$31.00©2018 IEEE

of diminishing rates of erroneous decisions. In [7] the Chernoff rule is modified to relax the assumption that the hypotheses should be uniformly distinguishable in the multi-hypothesis setting by introducing a randomization policy at certain time instants. The results are extended to the setting in which the available data belongs to a discrete alphabet and follows a stationary Markov model [8].

Despite their discrepancies in settings and approaches, all the aforementioned studies on controlled sensing assume that the available control actions are independent, or they follow a first-order stationary Markov process. This is in contrast to the setting of this paper, in which the control actions follow a known general dependency kernel. This paper devises a sequential sampling strategy, in which the dependence of control actions plays a significant role in designing the control policy. Specifically, the devised control policy, unlike the Chernoff rule, incorporates such dependence into the decision rules via accounting for the impact of each action on the future ones.

II. Data Model and Problem Formulation

A. Data Model

Consider a multi-hypothesis testing problem, consisting of L hypotheses denoted by $\mathcal{H} \triangleq \{\mathsf{H}_1, \dots, \mathsf{H}_L\}$. There exist n possible control actions. We denote the sample taken under action A_i by $X_i \in \mathbb{R}$, for $i \in \{1, \dots, n\}$. The samples are correlated, inducing correlation among this associated actions. We denote the joint cumulative distribution function (cdf) of $\mathcal{X} \triangleq \{X_1, \dots, \mathcal{X}_n\}$ under hypothesis H_ℓ by F_ℓ. For convenience in notations, we assume that the distributions of the random variables under each hypothesis $\mathsf{H}_\ell \in \mathcal{H}$ are absolutely continuous with respect to a common distribution and have well-defined probability density functions (pdfs). For every non-empty set $B \subseteq \{1, \dots, n\}$, we denote the joint pdf of $X_B \triangleq \{X_i : i \in B\}$ under H_ℓ by $f_\ell(\cdot; B)$. We also define $\mathsf{T} \in \mathcal{H}$ as the true hypothesis and denote the prior probability that hypothesis H_ℓ is true by ϵ_ℓ, i.e.,

$$\epsilon_\ell \triangleq \mathbb{P}(\mathsf{T} = \mathsf{H}_\ell) , \quad \text{for } \ell \in \{1, \dots, L\} , \tag{1}$$

where $\sum_{\ell=1}^L \epsilon_\ell = 1$.

B. Sampling Model

We consider a fully sequential data acquisition mechanism, in which we select one control action at each time and collect the sample generated under that action. It is assumed that each control action can be selected only once[1]. The goal is to design an optimal sequence of control actions, such that, with the minimum number of samples, the true model $\mathsf{T} \in \mathcal{H}$ can be determined. Samples are collected sequentially such that at any time t and based on the information accumulated up to t, the sampling procedure takes one of the following actions.

A$_1$) *Exploration:* Due to lack of sufficient confidence, making any decision is deferred and one more sample is taken by selecting one of the control actions that has not been taken before. Thus, the next control action should be specified.

A$_2$) *Detection:* Data collection process is terminated and a reliable decision about the true model of the data is formed. Hence, the stopping time and the final decision rule upon stopping should be specified.

The sampling process can be expressed uniquely by its stopping time, the final decision rule, and the data-adaptive control policy. In order to formalize the detection rule, we define $\tau_n \in \mathbb{N}$ as the stochastic stopping time, at which the sampling process is terminated and a decision is formed, and define $\delta_n \in \{1, \dots, L\}$ as the decision rule at the stopping time, where $\delta_n = \ell$ indicates a decision in favor of H_ℓ, for $\ell \in \{1, \dots, L\}$. Furthermore, in order to characterize the information-gathering process (exploration), we define the control policy $\psi_n : \{1, \dots, \tau\} \rightarrow \mathcal{A} \triangleq \{1, \dots, n\}$, where $\psi_n(t)$ specifies the control action at time t. Accordingly, we define ψ_n^t as the *ordered* set of control actions up to time t, i.e.,

$$\psi_n^t \triangleq \{\psi_n(1), \dots, \psi_n(t)\} , \tag{2}$$

and denote the sample taken under action $\psi_n(t)$ at time t by $Y_t \triangleq X_{\psi_n(t)}$. Also, the sequence of samples accumulated up to time t is denoted by

$$Y^t \triangleq \{Y_1, \dots, Y_t\} . \tag{3}$$

Furthermore, we define φ_n^t as the set of control actions that can still be taken at time t, i.e.,

$$\varphi_n^t \triangleq \mathcal{A} \setminus \{\psi_n(1), \dots, \psi_n(t-1)\} . \tag{4}$$

The information accumulated sequentially can be abstracted by the filtration $\{\mathcal{F}_t : t = 1, 2, \dots\}$, where

$$\mathcal{F}_t \triangleq \sigma(Y^t; \psi_n^t) . \tag{5}$$

The stopping time, the selection rule, and the control policy are $\mathcal{F}_t$-measurable functions. We define the $\mathcal{F}_t$-measurable tuple $\Phi_n \triangleq (\tau_n, \delta_n, \psi_n^\tau)$ to uniquely describe the sampling strategy.

Finally, we define the following information measures that are instrumental to formalizing and analyzing various decision rules. Specifically, for any $k \neq \ell$, any given ψ_n^t, and $B \subseteq \varphi_n^{t+1}$ we define

$$J_{\ell k}(B, \psi_n^t) \triangleq D_{\rm KL}\big(f_\ell(X_B; B \mid \mathcal{F}_t) \| f_k(X_B; B \mid \mathcal{F}_t)\big) , \tag{6}$$

where $D_{\rm KL}(f\|g)$ denotes the Kullback-Leibler (KL) divergence from a statistical model with pdf g to a model with pdf f.

C. Problem Formulation

The sequence of coupled information-gathering strategy and decision-making processes is uniquely specified by Φ_n. Designing the optimal sampling strategy for achieving the quickest reliable decision involves striking a balance between the *quality* and *agility* of the process as two opposing measures. The agility of the process is captured by the average

[1]This assumption is for convenience in notations, and can be relaxed with proper adjustment in the analysis.

delay in reaching a decision, i.e., $\mathbb{E}\{\tau\}$, and the decision quality is captured by the frequency of erroneous decisions. Such error rates can be captured by

$$\mathsf{P}_n^\ell \triangleq \mathbb{P}(\delta_n \neq \ell \mid \mathsf{T} = \mathsf{H}_\ell) , \quad \forall \ell \in \{1, \dots, L\} . \quad (7)$$

There exists an inherent tension between the accuracy and the agility of the process, since improving one penalizes the other one. Hence, to formalize the quickest reliable decision, we control the quality of the decision and minimize the average number of samples over all possible strategies. Hence, the optimal sampling strategy of interest is the solution to

$$\begin{array}{ll} \inf_{\Phi_n} & \mathbb{E}\{\tau_n\} \\ \text{s.t.} & \mathsf{P}_n^\ell \leq \alpha_n^\ell , \quad \forall \ell \in \{1, \dots, L\} \end{array} , \quad (8)$$

where $\alpha_n^\ell \in (0,1)$ controls the probability of error when H_ℓ is the true model. The values of α_n^ℓ, for $\ell \in \{1, \dots, L\}$, are selected such that problem (8) is feasible almost surely.

D. Assumptions

In order to formalize the performance bounds and analyze the proposed rules, we need to make some assumptions about the candidate models F_ℓ for $\ell \in \{1, \dots, L\}$. Let us define

$$Z_\ell(t) \triangleq \log f_\ell(Y^t; \psi_n^t) , \quad (9)$$

as the log-likelihood of the samples collected up to time t under hypothesis $\mathsf{H}_\ell \in \mathcal{H}$. Then, for any $\ell, k \in \{1, \dots, L\}$ where $\ell \neq k$, we define

$$\Lambda_{\ell k}(t) \triangleq Z_\ell(t) - Z_k(t) , \quad (10)$$

as the log-likelihood ratio (LLR) processes of the samples collected up to time t. When the samples are independent and identically distributed (i.i.d.), LLR processes are random walks with fixed expected step-sizes. When the collected samples are non-i.i.d., they lose this property, and the ensuing uncontrolled fluctuations in the step-sizes hinder our ability to analyze the problem in (8). In order to obtain lower bounds for the moments of a stopping time and to prove asymptotic optimality of the proposed rules, some restrictions should be imposed on these fluctuations. For this purpose, we use the notion of r-quick convergence [22]. Specifically, corresponding to any subsequence of control actions $\{\psi(t)\}_{t=1}^m$ and sequence of samples collected under these actions $\{Y_t\}_{t=1}^m$, we define $I_{\ell k}(\psi^m)$ as the convergence limit of the normalized LLRs under H_ℓ as follows

$$\frac{1}{m}\Lambda_{\ell k}(m) \to I_{\ell k}(\psi^m) , \quad (11)$$

when the convergence is r-quick as $m, n \to \infty$. The convergence in (11) is r-quick if and only if $\mathbb{E}_\ell\{T_{\ell k}^r(h)\} < \infty$ for any $h > 0$ and $r > 0$, where

$$T_{\ell k}(h) \triangleq \sup \left\{ t : \left| \frac{1}{m}\Lambda_{\ell k}(m) - I_{\ell k}(\psi^m) \right| \geq h \right\} .$$

We also define

$$I_{\ell k}^* \triangleq \sup_{\psi^m \subseteq \{1, \dots, n\}} I_{\ell k}(\psi^m) , \quad (12)$$

as the largest value of these information measures for each pair of hypotheses attained by a subsequence.

III. Data-Adaptive Sampling

In this section, we offer a coupled sampling and decision-making strategy $\Phi_n = (\tau_n, \delta_n, \psi_n^\tau)$ as a solution to (8) and analyze its optimality properties in Section IV.

A. Stopping Time and Decision Rules

The detection action consists of determining the stopping time of the sampling process and the final decision rule. To determine these two rules, for a set of thresholds $\{\gamma_{\ell k}\}$ we define

$$\nu_\ell \triangleq \inf \left\{ t : Z_\ell(t) \geq \max_{k \neq \ell} [Z_k(t) + \gamma_{\ell k}] , \text{or } t = n \right\} , \quad (13)$$

for any $\ell \in \{1, \dots, L\}$, and set

$$\tau_n^* \triangleq \min\{\nu_1, \dots, \nu_L\} . \quad (14)$$

Furthermore, at the stopping time we make a decision about the true model according to

$$\delta_n^* \triangleq \underset{\ell \in \{1, \dots, L\}}{\text{argmin}} \quad \nu_\ell . \quad (15)$$

The following theorem establishes that by selecting suitable set of thresholds $\{\gamma_{\ell k}\}$, the error probability constraints in (8) are satisfied by the likelihood ratio test (LRT) described above.

Theorem 1: When problem (8) is feasible, for any data adaptive sampling strategy $\Phi_n = (\tau_n^*, \delta_n^*, \psi_n^\tau)$ with the stopping time and decision rule specified in (13)–(15), we have $\mathsf{P}_n^\ell \leq \alpha_n^\ell$, for all $\ell \in \{1, \dots, L\}$, provided that

$$\sum_{k \neq \ell} \exp\{-\gamma_{k\ell}\} \leq \alpha_n^\ell , \quad \forall \ell \in \{1, \dots, L\} . \quad (16)$$

The following corollary provides one possible choice for selecting the thresholds.

Corollary 1: By setting

$$\gamma_{\ell k} = \log \frac{L}{\alpha_n^k} , \quad \forall \ell, k \in \{1, \dots, L\} , \quad (17)$$

the sampling strategy $\Phi_n = (\tau_n^*, \delta_n^*, \psi_n^\tau)$ satisfies $\mathsf{P}_n^\ell \leq \alpha_n^\ell$ for all $\ell \in \{1, \dots, L\}$.

Based on (13)–(15), the sampling process continues as long as none of the hypotheses is a strong candidate for the true model, and terminates when a sufficiently reliable candidate is found and makes a decision in favor of that model.

Detection rule in (15) specifies the decision at the stopping time. Prior to that, we need to dynamically characterize $\psi_n(t)$. In other words, based on (13) and (14), as long as the LLRs associated with none of the hypotheses are large enough, we need to take at least one more sample. In the next subsection we characterize the control policy $\psi_n(t)$ which identifies the action to be taken at each time.

B. Control Policy

Prior to the stopping time τ_n, at any time t the sampling process should dynamically identify one of the remaining control actions that provides the most relevant information about the true hypothesis. When the samples under different actions are statistically independent, the control policy at the current time instant has no impact on the future ones. In such a setting, the Chernoff rule [9], which selects the best immediate action, admits optimality properties under certain conditions. Specifically, at any time t the Chernoff rule first forms the maximum likelihood (ML) decision about the true hypothesis. Next, at time $(t+1)$ based on this decision, it selects the most informative action based on some information measures. The Chernoff rule minimizes the average delay in the asymptote of small rate of erroneous decisions, if all the actions are *independent* [9] and [7]. In this paper, however, the available actions are co-dependent. Therefore, the Chernoff rule, which ignores such dependence and the impact of the current decision on the future ones, fails to exploit it.

We propose a control policy which incorporates the impact of all future actions into the decision rules. To this end, we denote the ML decision about the true hypothesis at time t by $\delta_{\rm ML}(t)$, i.e.,

$$\delta_{\rm ML}(t) \triangleq \underset{\ell\in\{1,\dots,L\}}{\rm argmax} \; f_\ell(Y^t;\psi_n^t) \ . \tag{18}$$

Next, based on this decision we set the control policy. This step is the main distinction between the Chernoff rule and the proposed control policy of this paper. The Chernoff rule at any time selects the action that maximize the reliability of the decision about the model at that time, while our proposed rule considers the impact of the next action on all the future ones. We define $\psi_{\rm ch}(t)$ as the control policy of the Chernoff rule at time t, and accordingly define the *ordered* set $\psi_{\rm ch}^t \triangleq \{\psi_{\rm ch}(1),\dots,\psi_{\rm ch}(t)\}$. The Chernoff rule selects the action that maximizes the distance between f_ℓ and its closest alternative when the ML decision is in favor of H_ℓ, i.e.,

$$\psi_{\rm ch}(t) \triangleq \arg\max_{i\in\varphi_n^t} \; \min_{k\neq\ell} \; J_{\ell k}\big(\{i\},\psi_{\rm ch}^{t-1}\big) \ . \tag{19}$$

In our proposed control policy, we incorporate the impact of $\psi_n(t)$ on all future actions. To this end, new information measures are introduced to facilitate selecting the next action, which is the one that maximizes the combination of immediate information and future expected information gain. For this purpose, at time t and for each action $i\in\varphi_n^t$ we define the set $\mathcal{R}_t^i$ as the set of all subsets of φ_n^t that contain i, i.e.,

$$\mathcal{R}_t^i \triangleq \{\mathcal{S} \ : \ \mathcal{S}\subseteq\varphi_n^t \text{ and } i\in\mathcal{S}\} \ . \tag{20}$$

Corresponding to the samples collected under the actions in the set $\mathcal{S}\in\mathcal{R}_t^i$ under $\mathsf{H}_\ell\in\mathcal{H}$ we define the following information measures:

$$M_\ell^i(t,\mathcal{S}) \triangleq \min_{k\neq\ell} \; J_{\ell k}(\mathcal{S},\psi^{t-1}) \ . \tag{21}$$

The terms $M_\ell^i(t,\mathcal{S})$ capture the information content of $|\mathcal{S}|$ samples. Hence, the normalized terms $\frac{M_\ell^i(t,\mathcal{S})}{|\mathcal{S}|}$ account for the *average* information content per sample. Based on these normalized measures, a candidate decision is to select the control action that maximizes the *average* information over all possible future control actions. Therefore, when the ML decision about the true hypothesis is H_ℓ, the optimal control policy is the solution of the following optimization problem over all combinations of the remaining actions:

$$\psi_n^*(t) \triangleq \arg\max_{i\in\varphi_n^t} \; \max_{\mathcal{S}\in\mathcal{R}_t^i} \; \frac{M_\ell^i(t,\mathcal{S})}{|\mathcal{S}|} \ , \tag{22}$$

In this control policy, an ML decision about the true hypothesis is formed based on the collected data, and the action that maximizes the average information over all possible sequences of future control actions is selected. We note that the sets $\mathcal{S}$ are selected such that they contain the control action i, which is a candidate to be taken at time t, and possibly additional control actions that will be taken in the future.

Since under some of the control actions different distributions may be non-distinguishable, similar to [7] we introduce randomized decisions at certain exponentially-separated time instants. Specifically, at time instants $t=\lfloor a^s\rfloor$, for some $a>1$, which is close to 1 and $s\in\mathbb{N}$, we select one control action from φ_n^t randomly. The randomized actions guarantee that the ML decision converges to the true model of the data in a polynomially-bounded time [7].

IV. Performance Analysis

In this section, we focus of feasible problems of form (8), and analyze the asymptotic performance of the proposed sequence of strategies $\{\Phi_n\}_{n\in\mathbb{N}}$ in the asymptote of small error rates. To this end, we define

$$\bar\alpha_\ell \triangleq -\lim_{n\to\infty}\frac{1}{n}\log\alpha_n^\ell \ , \tag{23}$$

and assume that $\bar\alpha_\ell>0$ for all $\ell\in\{1,\dots,L\}$. Then, as n grows, the error rates approach zero. A counter example is provided in [23], which shows that the Chernoff rule loses its optimality properties when the control actions are co-dependent. In this section, we prove that the proposed strategy attains the same asymptotic optimality properties under co-dependent actions as that of the Chernoff rule under independent actions.

Based on the measures defined in (11) and (12), we provide the optimality guarantee of the proposed strategy for the problem in (8). First, in the following theorem we provide the performance bounds of any feasible solution to problem (8).

Theorem 2: Under the assumptions in (11) and (12), in the asymptote of large n and for any $m\in\mathbb{N}$, any feasible solution to problem (8) satisfies

$$\frac{\mathbb{E}_\ell\{\tau^m\}}{n^m} \geq \left[\max_{k\neq\ell}\frac{\bar\alpha_k}{I_{\ell k}^*}\right]^m\cdot\big(1+o(1)\big) \ . \tag{24}$$

Next, we show that the proposed strategy $\Phi_n^*=(\tau_n^*,\delta_n^*,\psi_n^{\tau *})$ achieves asymptotic optimality under this setting.

Theorem 3: Under the assumptions in (11) and (12), if $0<\frac{\alpha_n^\ell}{\alpha_n^k}<\infty$ for any $\ell,k\in\{1,\dots,L\}$, the proposed strategy

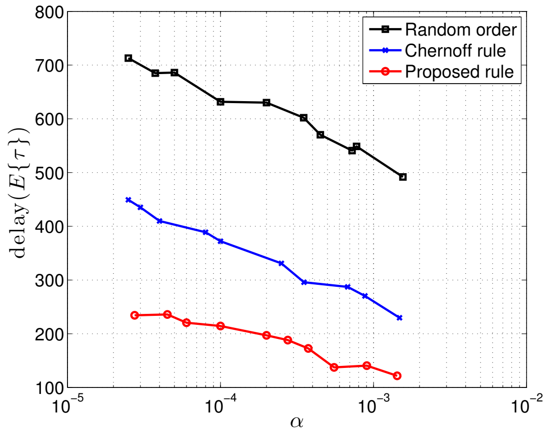

Fig. 1. Average delay of different approaches versus error probability.

$\Phi_n^* = (\tau_n^*, \delta_n^*, \psi_n^{\tau^*})$ specified by (14), (15), and (22) is an optimal solution to problem (8) with respect to any moment of the stopping time in the asymptote of large n, i.e., for any $m \in \mathbb{N}$

$$\frac{\mathbb{E}_\ell\{(\tau_n^*)^m\}}{n^m} \le \left[\max_{k \neq \ell} \frac{\bar{\alpha}_k}{I_{\ell k}^*}\right]^m \cdot \big(1 + o(1)\big) \ . \tag{25}$$

Theorems 2 and 3 prove the optimality of the proposed strategy in the asymptote of large n.

V. Simulation Results

In this section, we evaluate the performance of the proposed sampling strategy by comparing it with the existing approaches through simulations. For this purpose, we use a random control policy and the Chernoff rule as the benchmark methods. We consider zero-mean Gaussian distributions for data, and test for two different covariance matrices. We also set $\epsilon_\ell = 0.5$ and $\alpha_n^\ell = \alpha$ for $\ell \in \{1, 2\}$. In Fig. 1, we compare the performance of different approaches in terms of the average decision delay for making a final decision with the same quality. It is observed that the proposed sampling procedure outperforms both the pre-specified and the Chernoff rule in terms of the reliability-agility trade-off.

VI. Conclusion

We have considered the problem of controlled sensing for multi-hypothesis testing when the actions are co-dependent. The objective is to determine the true hypothesis with the desired reliability by taking the minimum average number of samples. After discussing the widely used Chernoff rule and its shortcomings, we have designed a sequential and data-adaptive sampling strategy, consisting of a stopping time, a final decision rule, and a control policy. The proposed sampling strategy, which judiciously incorporates the dependence of the actions into its decision rules, involves dynamically deciding whether to terminate the sampling process, or to continue collecting further evidence, and prior to terminating the process specifies the best control action at each time instant. We have established the optimality properties of the proposed sampling strategy and verified its superior performance through simulations.

References

[1] A. Anandkumar, L. Tong, and A. Swami, "Detection of Gauss-Markov random fields with nearest-neighbor dependency," *IEEE Transactions on Information Theory*, vol. 55, no. 2, pp. 816–827, Feb. 2009.

[2] H. V. Poor, *An Introduction to Signal Detection and Estimation*. New York: Springer, 1988.

[3] S. Bessler, "Theory and applications of the sequential design of experiments, k-actions and infinitely many experiments: Part I theory," Department of Statistics, Stanford University, Tech. Rep., Mar 1960.

[4] A. E. Albert, "The sequential design of experiments for infinitely many states of nature," *The Annals of Mathematical Statistics*, vol. 32, no. 3, pp. 774–799, Sep. 1961.

[5] W. J. Blot and D. A. Meeter, "Sequential experimental design procedures," *Journal of the American Statistical Association*, vol. 68, no. 343, pp. 586–593, Sep. 1973.

[6] K. Cohen and Q. Zhao, "Active hypothesis testing for anomaly detection," *IEEE Transactions on Information Theory*, vol. 61, no. 3, pp. 1432–1450, Mar. 2015.

[7] S. Nitinawarat, G. K. Atia, and V. V. Veeravalli, "Controlled sensing for multihypothesis testing," *IEEE Transactions on Automatic Control*, vol. 58, no. 10, pp. 2451–2464, May 2013.

[8] S. Nitinawarat and V. V. Veeravalli, "Controlled sensing for sequential multihypothesis testing with controlled Markovian observations and non-uniform control cost," *Sequential Analysis*, vol. 34, no. 1, pp. 1–24, Feb. 2015.

[9] H. Chernoff, "Sequential design of experiments," *The Annals of Mathematical Statistics*, vol. 30, no. 3, pp. 755–770, Sep. 1959.

[10] A. O. Hero and D. Cochran, "Sensor management: Past, present, and future," *IEEE Sensors Journal*, vol. 11, no. 12, pp. 3064–3075, Dec. 2011.

[11] G. A. Hollinger, B. Englot, F. S. Hover, U. Mitra, and G. S. Sukhatme, "Active planning for underwater inspection and the benefit of adaptivity," *International Journal of Robotics Research*, vol. 32, no. 1, pp. 3–18, Jan. 2013.

[12] S. M. Berry, B. P. Carlin, J. J. Lee, and P. Müller, *Bayesian adaptive methods for clinical trials*, ser. Chapman & Hall/CRC Biostatistics Series. CRC Press, Boca Raton, FL, 2011, vol. 38, with a foreword by David J. Spiegelhalter.

[13] R. D. Nowak, "The geometry of generalized binary search," *IEEE Transactions on Information Theory*, vol. 57, no. 12, pp. 7893–7906, Dec. 2011.

[14] M. V. Burnashev, "Data transmission over a discrete channel with feedback. random transmission time," *Problemy Peredachi Informatsii*, vol. 12, no. 4, pp. 10–30, 1976.

[15] N. Iri and O. Kosut, "Third-order coding rate for universal compression of markov sources," in *Proc. IEEE International Symposium on Information Theory*, Jun. 2015, pp. 1996–2000.

[16] ——, "A new type size code for universal one-to-one compression of parametric sources," in *Proc. IEEE International Symposium on Information Theory (ISIT)*, Jul. 2016, pp. 1227–1231.

[17] 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.

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

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

[20] V. P. Dragalin, A. G. Tartakovsky, and V. V. Veeravalli, "Multihypothesis sequential probability ratio tests–part I: Asymptotic optimality," *IEEE Transactions on Information Theory*, vol. 45, pp. 2448–2461, 1999.

[21] ——, "Multihypothesis sequential probability ratio tests–part II: Accurate asymptotic expansions for the expected sample size," *IEEE Transactions on Information Theory*, vol. 46, no. 4, pp. 1366–1383, Jul. 2000.

[22] A. G. Tartakovsky, "Asymptotic optimality of certain multihypothesis sequential tests: Non-i.i.d. case," *Statistical Inference for Stochastic Processes*, vol. 1, no. 3, pp. 265–295, 1998.

[23] J. Heydari, A. Tajer, and H. V. Poor, "Quickest detection of Markov networks," *arXiv:1711.04268v2*, Dec. 2017.