CONTROLLED SENSING FOR HYPOTHESIS TESTING

Sirin Nitinawarat, George K. Atia and Venugopal V. Veeravalli

Department of Electrical and Computer Engineering
and
Coordinated Science Laboratory
University of Illinois at Urbana-Champaign
Urbana, IL 61801-2307, USA
Emails: {nitinawa, atia1, vvv}@illinois.edu

ABSTRACT

In this paper, the problem of multiple hypothesis testing with observation control is considered. The structure of the optimal controller under various asymptotic regimes is studied. First, a setup with a fixed sample size is considered. In this setup, the asymptotic quantity of interest is the optimal exponent for the maximal error probability. For the case of binary hypothesis testing, it is shown that the optimal error exponent corresponds to the maximum Chernoff information over the choice of controls. It is also shown that a pure stationary control policy, i.e., a fixed policy which does not depend on specific realizations of past measurements and past controls (open-loop), is asymptotically optimal even among the class of all causal control policies. We also derive lower and upper bounds for the optimal error exponent for the case of multiple hypothesis testing. Second, a sequential setup is considered wherein the controller can also decide when to stop taking observations. In this case, the objective is to minimize the expected stopping time subject to the constraints of vanishing error probabilities under each hypothesis. A sequential test is proposed for testing multiple hypotheses and is shown to be asymptotically optimal.

1. INTRODUCTION

The topic of controlled sensing for inference in uncertain environments deals primarily with adaptively managing and controlling multiple degrees of freedom in an information-gathering systems, ranging from the sensing modality to the physical control of sensors to fulfill the goal of a given inference task. It has immediate implications in various applications, including infrastructure monitoring systems, surveillance systems, sensor networks and social networks. The main challenge lies in understanding the tradeoff between the value of information and the cost of acquiring this information for making decisions regarding the underlying inference problem.

In this paper, we focus on the inference problem of hypothesis testing in a non-Bayesian setting, and our goal is an asymptotically optimal joint-design of a control policy and a decision rule to decide among the various hypotheses. In particular, we consider a Markovian model of simple hypothesis testing of multiple hypotheses with observation control. Prior to making a decision about the hypothesis, the decision-maker can choose among different actions, which in turn affect the quality of the observations. *Due to space constraints, we state all results without their proofs.*

In section 2, we focus on the setup with a fixed sample size. In this setup we further consider two possible control policies, namely, open-loop control and causal control. While in the former the control cannot depend on the measurements, in the latter, the control can be a function of past measurements and past controls.

For a binary hypothesis testing problem *without control*, the characterization of the optimal exponent for the maximal error probability (maximized over all hypotheses) in terms of the Chernoff information is well-known (see, e.g., [3]). However, in the presence of control, it is not clear what the optimal error exponent is and what the optimal controller should be. In Section 2, we first give a complete characterization of the optimal exponent for the maximal error probability and for the case of binary hypothesis testing. Interestingly, we show that a pure stationary control suffices to achieve the optimal error exponent among the class of causal controls. We also provide upper and lower bounds for the optimal error exponent for the case of general multiple hypothesis testing (with more than two hypotheses). It remains open to characterize the optimal error exponent among the class of all causal controls.

In section 3, we consider a sequential setup wherein the controller can adaptively choose to stop taking measurements.

This work was supported by the Air Force Office of Scientific Research (AFOSR) under the Grant FA9550-10-1-0458 through the University of Illinois at Urbana-Champaign, and by the U.S. Defense Threat Reduction Agency through subcontract 147755 at the University of Illinois from prime award HDTRA1-10-1-0086.

978-1-4673-0046-9/12/$26.00 ©2012 IEEE

In this case, a sequential test is fully described by a control policy, a stopping rule and a final decision rule. In contrast to the earlier fixed sample size setup in Section 2, the error exponent is no longer a good metric since the expected stopping time will surely depend on the underlying hypothesis.

In [8], the problem of binary sequential hypothesis testing *without control* was considered. Therein, the optimal expected values of the stopping time are characterized subject to the constraints of vanishing probabilities of error under each hypothesis. An extension to the case of multiple hypothesis testing was considered in [1] in which the authors proposed a sequential test which was later shown to satisfy certain asymptotic optimality conditions [4, 5].

The problem of sequential binary hypothesis testing with observation control was considered in [2] and an asymptotically optimal (in the aforementioned sense) sequential test was presented. A Bayesian version of this sequential problem (with the observation control) was considered by the authors in [6] in the non-asymptotic regime. Since the optimal policy is generally hard to characterize, they identified certain conditions under which the optimal control is shown to be an open-loop control. In this paper, we extend the result in [2] to the case of multiple hypothesis testing. Unlike the fixed sample size setup, an open-loop control may no longer be optimal for the case of binary hypothesis testing. In addition, randomization is necessary to attain the optimal tradeoff except for the case of binary hypothesis testing.

2. THE FIXED SAMPLE SIZE SETUP

In this section, we consider the setup in which the sample size does not depend upon specific realizations of the measurements and the controls.

Consider a hypothesis testing problem with M hypotheses: $H_i,\ i \in \mathcal{M} \triangleq \{0, \ldots, M-1\}$, where at each time the measurement takes values in $\mathcal{X}$ and the control takes values in $\mathcal{U}$. Both the alphabets of the observation $\mathcal{X}$ and of the control $\mathcal{U}$ are assumed to be finite. Under each hypothesis $H_i, i = 0, \ldots, M-1$, and at each time k, conditioning on the current control $u_k = u$, the current observation X_k is assumed to be conditionally independent of past measurements and past controls $\left(x^{k-1}, u^{k-1}\right) \triangleq \left(\left(x_1, \ldots, x_{k-1}\right), \left(u_1, \ldots, u_{k-1}\right)\right)$ and to be conditionally distributed according to a probability mass function (pmf) $P_i^u(x)$.

We consider two classes of control polices based on two possible information patterns. The first is an open-loop control where the (possibly randomized) control sequence $(U_1, \ldots, U_n)$ is assumed to be mutually independent of the measurements $(X_1, \ldots, X_n)$, and the second is a causal control where at each time k, the control U_k can be any (possibly randomized) function of past measurements and past controls, i.e., $U_k,\ k = 2, 3, \ldots, n$, is described by an arbitrary conditional pmf $q_k\left(u_k|x^{k-1}, u^{k-1}\right)$, and U_1 is distributed according to a pmf $q_1(u_1)$. If all these (conditional) pmfs are point-mass distributions, i.e., the current control is a deterministic function of past measurements and past controls, then the resulting control is a *pure* control policy. Under the aforementioned (conditionally) memoryless assumption, the joint pmf of (X^n, U^n) under each hypothesis H_i, denoted by $\mathbb{P}_i\left(x^n, u^n\right),\ i = 0, \ldots, M-1$, can be written as

$$\mathbb{P}_i\left(x^n, u^n\right) \triangleq q_1\left(u_1\right) \prod_{k=1}^{n} P_i^{u_k}(x_k) \prod_{k=2}^{n} q_k\left(u_k|x^{k-1}, u^{k-1}\right). \tag{1}$$

Note that for an open-loop control, each $q_k\left(u_k|x^{k-1}, u^{k-1}\right)$, $k = 2, \ldots, n$, is independent of x^{k-1}.

After n observations, a decision is made according to the rule $\delta : \mathcal{X}^n \times \mathcal{U}^n \to \mathcal{M}$ with the maximal error probability:

$$e\left(\{q_k\}_{k=1}^n, \{P_i^u\}_{i\in\mathcal{M}}^{u\in\mathcal{U}}, \delta\right) \triangleq \max_{i\in\mathcal{M}} \mathbb{P}_i\left\{\delta\left(X^n, U^n\right) \neq i\right\}.$$

Note that for a pure control, u^n is either fixed (pure open-loop control) or is a deterministic function of the measurements x^n (pure causal control). Consequently, when a pure control is adopted, it suffices to consider a decision rule that is a function only of the measurements x^n, i.e., $\delta\left(x^n, u^n\right) = \delta\left(x^n\right)$.

The asymptotic quantities of our interest will be the largest exponent for the maximal error probability achievable by an open-loop control, denoted by β_{OL}, and by a causal control, denoted by β_C, respectively. In particular,

$$\beta_{OL} \triangleq \overline{\lim_n} \sup_{\delta,\ q(u^n)} -\frac{1}{n}\log\left(e\left(q\left(u^n\right), \{P_i^u\}_{i\in\mathcal{M}}^{u\in\mathcal{U}}, \delta\right)\right);$$

$$\beta_C \triangleq \overline{\lim_n} \sup_{\substack{\delta,\ q_1(u_1),\\ \{q_k(u_k|x^{k-1},u^{k-1})\}_{k=2}^n}} -\frac{1}{n}\log\left(e\left(\{q_k\}_{k=1}^n, \{P_i^u\}_{i\in\mathcal{M}}^{u\in\mathcal{U}}, \delta\right)\right).$$

It follows immediately from these definitions that $\beta_{OL} \leq \beta_C$, as the information pattern associated with the causal control is more informative than that associated with the open-loop control.

2.1. The Case of Binary Hypothesis Testing ($M = 2$)

We start with an auxiliary result that involves a different setup in which a constraint is put on the probability of error under the null hypothesis H_0, i.e., only control policies and decision rules that satisfy

$$\mathbb{P}_0\left\{\delta \neq 0\right\} \leq \epsilon, \tag{2}$$

for a fixed $\epsilon,\ 0 < \epsilon < 1$, will be considered. Under this constraint, we are interested in the largest error exponents under the alternative hypothesis H_1, achievable by an open-loop control, denoted by $\alpha_{OL}(\epsilon)$, and by a causal control, denoted by $\alpha_C(\epsilon)$, respectively. In particular,

$$\alpha_{OL}(\epsilon) \triangleq \overline{\lim_n} \sup_{\delta,\ q(u^n)} -\frac{1}{n}\log\left(\mathbb{P}_1\left\{\delta \neq 1\right\}\right); \tag{3}$$

$$\alpha_C(\epsilon) \triangleq \overline{\lim_n} \sup_{\substack{\delta,\, q_1(u_1),\\ \{q_k(u_k|x^{k-1},u^{k-1})\}_{k=2}^n}} -\frac{1}{n}\log\left(\mathbb{P}_1\{\delta \neq 1\}\right). \tag{4}$$

where the suprema in (3) and (4) are over all corresponding control policies and decision rules satisfying (2).

For two pmfs q and p on $\mathcal{X}$, the Kullback-Leibler (KL) divergence of q and p, denoted by $D(q\|p)$, is defined as

$$D(q\|p) \triangleq \begin{cases} \infty, & \text{if } \exists x \in \mathcal{X},\ p(x) = 0,\ q(x) > 0, \\ \sum_{x:\ p(x)>0} q(x)\log\left(\frac{q(x)}{p(x)}\right), & \text{otherwise.} \end{cases}$$

Proposition 1 *For $M = 2$ and any ϵ, $0 < \epsilon < 1$, it holds that*

$$\alpha_{OL}(\epsilon) = \alpha_C(\epsilon) = \max_{u\in\mathcal{U}} D\left(P_0^u,\ P_1^u\right). \tag{5}$$

Remark 1 *The above characterization of $\alpha_{OL}(\epsilon)$ can be derived from a result of Tsitsiklis [7] in a different setup; our main contribution is the characterization of $\alpha_C(\epsilon)$.*

Remark 2 *It follows from Proposition 1 and Stein's Lemma (see, e.g., [3]) that to achieve the optimal error exponent, it suffices to use a stationary control sequence $u_k = u^* = \mathrm{argmax}_{u\in\mathcal{U}} D\left(P_0^u, P_1^u\right),\ k = 1, \ldots, n$. In particular, information from the past and randomization are superfluous for attaining the best error exponent.*

Our main contribution for the case of binary hypothesis testing is the following characterizations of the optimal exponents for the maximal error probability achievable by an open-loop control and by a causal control. Its derivation relies on Proposition 1 and is omitted.

For any $u \in \mathcal{U}$, and any $s \in [0,1]$, we consider the following pmf [1]

$$B_s^u(x) \triangleq \frac{P_0^u(x)^s P_1^u(x)^{1-s}}{\sum_x P_0^u(x)^s P_1^u(x)^{1-s}}. \tag{6}$$

We also let

$$s^*(u) \triangleq \operatorname*{argmax}_{s\in[0,1]} -\log\left(\sum_x P_0^u(x)^s P_1^u(x)^{1-s}\right).$$

Theorem 1 *For $M = 2$, it holds that*

$$\begin{aligned} \beta_{OL} &= \beta_C \\ &= \max_{u\in\mathcal{U}} \max_{s\in[0,1]} -\log\left(\sum_x P_0^u(x)^s P_1^u(x)^{1-s}\right) \qquad (7)\\ &= \max_{u\in\mathcal{U}} D\left(B_{s^*(u)}^u \| P_0^u\right) = \max_{u\in\mathcal{U}} D\left(B_{s^*(u)}^u \| P_1^u\right). \qquad (8) \end{aligned}$$

[1]To get continuity in the parameter $s \in [0,1]$, we use a convention that when $s = 0$ or 1, $0^0 = 0$.

Remark 3 *For each fixed $u \in \mathcal{U}$, the quantity*

$$\max_{s\in[0,1]} -\log\left(\sum_x P_0^u(x)^s P_1^u(x)^{1-s}\right)$$

is called the "Chernoff information" of P_0^u and P_1^u. It has an operational significance (see, e.g., [3]) of being the largest exponent for the maximal error probability for an independent and identically distributed (i.i.d.) model of observations P_0^u and P_1^u (for that fixed u) under the respective hypotheses. Consequently, Theorem 1 (cf. (7)) states that under observation control, the optimal error exponent is the maximum Chernoff information over the choice of controls.

Remark 4 *Similar to Remark 2 pertaining to Proposition 1, it follows from Theorem 1 and the result on the Chernoff information for i.i.d. observations that the above optimal error exponent is achievable by a stationary (fixed) control sequence $u_k = u^*$ which is the maximizer of the right-side of (7) (or, identically, that of the two quantities in (8)).*

2.2. The Case of Multiple Hypothesis Testing $(M > 2)$

Our last result for the fixed sample size setup is a full (partial) characterization of the optimal exponent for the maximal error probability achievable by an open-loop (a causal) control, respectively.

Theorem 2 *For $M > 2$, it holds that*[2]

$$\begin{aligned} &\max_{q(u)} \min_{i\neq j} \max_{s\in[0,1]} -\sum_{u\in\mathcal{U}} q(u)\log\left(\sum_x P_i^u(x)^s P_j^u(x)^{1-s}\right) \\ &\quad = \beta_{OL} \leq \beta_C \leq \\ &\min_{\substack{i\neq j\\ 0\leq i<j\leq M-1}} \max_{u\in\mathcal{U}} \max_{s\in[0,1]} -\log\left(\sum_x P_i^u(x)^s P_j^u(x)^{1-s}\right). \end{aligned}$$

3. THE SEQUENTIAL SETUP

In the previous section we considered tests with fixed sample size. In this section, we consider a different setup where new observation arrives at each time instant and where the controller also has to choose when to stop taking observations (stopping time). In this case, the goal is to design a sequential test, which consists of a rule to select control actions, a stopping rule, and a final decision rule to optimize the trade-off between reliability (in terms of probabilities of error) and delay.

In [2] Chernoff considered the problem of composite binary hypothesis testing with observation control. He presented a simple procedure for this problem which was shown to be asymptotically optimal. In this paper, we extend this result to the case where $M > 2$.

[2]The first equality regarding β_{OL} can be already inferred from a result in [7].

3.1. Model

Let $\mathcal{F}_k$ be the σ-field generated by (X^k, U^k). An admissible sequential test $\gamma = (\phi, N, \delta)$ consists of a causal observation control policy ϕ, defined through (conditional) pmfs akin to the paragraph preceding (1), an $\mathcal{F}_k$-stopping time N and the final decision rule δ. At each time instant the controller has to decide whether to take more observations or to stop and make a decision about the unknown hypothesis. If the controller decides to take an observation, a positive cost c is incurred and the controller selects one of the possible actions in $\mathcal{U}$. Clearly, there is a tradeoff between the performance of the test and the expected time to reach the final decision (hence a cost or a delay). Our goal is to design an efficient test to optimize the aforementioned tradeoff. More specifically, we are interested in designing a test and studying its behavior in terms of probabilities of error and delay as the cost of observations c approaches 0, i.e., when the probabilities of error $\mathbb{P}_i(\delta(X^n, U^n) \neq i)$ are small and the expected stopping time $\mathbb{E}_i[N]$ is large under each hypothesis $i \in \mathcal{M}$. Next, we propose a sequential test for multiple hypotheses and analyze its asymptotic performance.

3.2. The Sequential Test

At time k, we compute the most likely hypothesis, denoted by $\hat{i}_k$, and the *closest* alternative hypothesis $\tilde{i}_k$, given all observations upto time k as follows

$$\hat{i}_k = \underset{i \in \mathcal{M}}{\operatorname{argmax}} \mathbb{P}_i(x^k, u^k), \text{ and } \tilde{i}_k = \underset{i \in \mathcal{M} \setminus \{\hat{i}_k\}}{\operatorname{argmax}} \mathbb{P}_i(x^k, u^k) \tag{9}$$

If we decide to continue taking observations at time k, then we select a control $u_{k+1} \in \mathcal{U}$ sampled from a distribution

$$q(u) \triangleq \mathbb{P}\{U_{k+1} = u | \hat{I}_k = \hat{i}_k\}$$

obtained as a solution to the following maximin problem

$$\max_{q(u)} \min_{i \in \mathcal{M}, i \neq \hat{i}_k} \sum_{u \in \mathcal{U}} q(u) D(P^u_{\hat{i}_k}, P^u_i). \tag{10}$$

The stopping rule is then defined as follows. We stop taking observations at time $k = n$ if the following condition is satisfied

$$\log \frac{\mathbb{P}_{\hat{i}_n}(x^n, u^n)}{\mathbb{P}_{\tilde{i}_n}(x^n, u^n)} \geq -\log c. \tag{11}$$

If we stop at time $k = n$, the decision rule chooses

$$\delta(x^n, u^n) = \hat{i}_n. \tag{12}$$

Theorem 3 *The proposed sequential test which has the control policy in (10), the stopping rule in (11), and the final decision rule in (12) achieves, for each hypothesis H_i, $i \in \mathcal{M}$, the following probabilities of error*

$$\mathbb{P}_i(\delta(X^N, U^N) \neq i) = O(c), \tag{13}$$

and achieves, for each hypothesis H_i, $i \in \mathcal{M}$, the following values of the expected stopping time

$$\mathbb{E}_i[N] \leq \frac{-\log c}{\max_{q(u) \in \mathcal{U}} \min_{j \in \mathcal{M}, j \neq i} \sum_{u \in \mathcal{U}} q(u) D(P^u_i, P^u_j)} + o(-\log c), \forall i \in \mathcal{M}. \tag{14}$$

Furthermore, this test is asymptotically optimal, i.e., any sequential test which achieves probabilities of error as in (13) under all hypotheses will also have an expected stopping time larger than (14).

Remark 5 *Note that the maximizing distributions of the expression in the denominator on the right side of (14) may vary across the different hypotheses H_i, $i \in \mathcal{M}$. It then follows from this observation that for the sequential setup, unlike for the fixed sample size setup, an open-loop control may* not *be asymptotically optimal for the case of binary hypothesis testing. In addition, randomization is necessary to attain the optimal performance in (14) except for the case of $M = 2$.*

4. REFERENCES

[1] C.W. Baum and V. V. Veeravalli, "A Sequential Procedure for Multiple Hypothesis Testing," *IEEE Trans. Inf. Theory,* vol. 40, pp. 1994-2007, 1994.

[2] H. Chernoff, "Sequential Design of Experiments," *Ann. of Math. Stat.*, vol. 30, pp. 755-770, 1959.

[3] T. M. Cover and J. A. Thomas, *Elements of Information Theory,* John Wiley and Sons, Inc., New York, 2006.

[4] V. Dragalin, A. G. Tartakovsky and V. V. Veeravalli, "Multihypothesis Sequential Probability Ratio Tests, Part I: Asymptotic Optimality," *IEEE Trans. Inf. Theory,* vol. 45, pp. 2448-2462, 1999.

[5] V. Dragalin, A. G. Tartakovsky and V. V. Veeravalli, "Multihypothesis Sequential Probability Ratio Tests, Part II: Accurate Asymptotic Expansions for the Expected Sample Size," *IEEE Trans. Inf. Theory,* vol. 46, pp. 1366-1383, 2000.

[6] M. Naghshvar and T. Javidi, "Active M-ary Sequential Hypothesis Testing," *in the Proc. 2010 IEEE internat. Symp. Inf. Theory,* pp. 1623-1627, Texas, Austin, USA, June 13-18, 2010.

[7] J. Tsitsiklis, "Decentralized Detection by a Large Number of Sensors," *Math. Ctrl., Signals and Syst.,* vol. 1, pp. 167-182, 1988.

[8] A. Wald, *Sequential Analysis,* John Wiley and Sons, Inc., New York, 1947.