

Universal scheme for optimal search and stop

SIRIN NITINAWARAT¹ and VENUGOPAL V. VEERAVALLI²

¹Senior Systems Engineer, Qualcomm Technologies, Inc., 5775 Morehouse Drive, San Diego, CA 92121, USA. E-mail: sirin.nitinawarat@gmail.com

²Professor, Coordinated Science Laboratory, Department of Electrical and Computer Engineering, University of Illinois at Urbana-Champaign, Urbana, IL, 61801, USA. E-mail: vvv@illinois.edu

The problem of universal search and stop using an adaptive search policy is considered. When the unique target location is searched, the observation is distributed according to the target distribution, otherwise it is distributed according to the absence distribution. A universal scheme for search and stop is proposed using only the knowledge of the absence distribution, and its asymptotic performance is analyzed. The universal test is shown to yield a vanishing error probability, and to achieve the *optimal* reliability when the target is present, universally for every target distribution. Consequently, it is established that the knowledge of the target distribution is only useful for improving the reliability for detecting that the target is missing. It is also shown that a multiplicative gain for the search reliability equal to the number of searched locations is achieved by allowing adaptivity in the search.

Keywords: search and stop; sequential design of experiments; sequential hypothesis testing

1. Introduction

We study the problem of universal search and stop using an adaptive policy to sequentially search among a finite number of locations. The unique target could be missing or present in one of the search locations, and when present it is stationary. When the target is present and when the unique target location is searched, the observation is assumed to be distributed according to the target distribution, otherwise it is distributed according to the absence distribution. We assume that only the absence distribution is known, and the target distribution can be arbitrarily distinct from the absence distribution. We also assume a “noisy” situation in which both the target and absence distributions have full supports. This means that even when the search has covered all the locations, it is not possible to determine the target location or if it is missing with perfect certainty. An adaptive search policy specifies the current search location based on the past observations and past search locations. At the stopping time, the target location is determined or it is decided that it is missing. *The overall goal is to achieve a certain level of accuracy for the final decision using the fewest number of observations.* A canonical example for this universal search and stop problem is environment monitoring in a sensor network where the absence distribution could be from a pure noisy observation collected at most sensors covering “uninteresting” regions and the target distribution represents the distribution of the observation at a unique sensor effected by the “interested” event or phenomenon. Our main interest is in the asymptotic regime of vanishing error probability (for the final decision), whereupon the number of searches (up to the stopping time) has to go to infinity, yielding asymptotically large number of recurrent searches at each location. In terms of the theoretical contributions of this work, the results in this paper should be

regarded as a contribution to the long-studied area of search theory (see, e.g., [1–3,10,13,21]), in particular, searching for a stationary target in discrete time and space with a discrete search effort (cf. [3], Section 4.2).

Since we assume that the observations at all locations without the target follow the same absence distribution, a desirable goal of the search at each location should be to decide if the target is there or not (elsewhere or missing entirely). To this end, a universal sequential test for two hypotheses could be used at each location to collect multiple subsequent observations that will eventually lead to a binary outcome of whether the target is there or not. To improve reliability for this binary decision at a particular search location, one could use a test that takes more observations at that location. If we insist on using the mentioned sequential binary test at each location as an “inner” test, then it is convenient to select the current search location based on the past binary outcomes of the subsequent binary tests, each resulting from multiple subsequent observations at a particular location, instead of all the past outcomes of the individual searches, often taken multiple times at each of the locations. With this imposition, the search and stop problem could be conceptually reduced to the problem of constructing an “outer” test for a sequential design of such inner experiments. This intuitive decomposition leads to our proposed universal test for search and stop.

Universal sequential testing for two hypotheses was first considered for certain parametric families of distributions for continuous observation spaces in [8,9,14,20], the latest of which employed the concept of time-dependent thresholding. Here in Section 4.1, we look at a non-parametric family of distributions for a *finite* observation space, for which we propose a universal test using a suitable time-dependent threshold and analyze its performance.

The sequential design of experiments with a uniform experimental cost was first considered in [4,7] under a certain positivity assumption for the model, which was successfully dispensed with later in [17,18]. A generalization of the model with a more complicated memory structure for the experimental outcomes and with a non-uniform experimental cost was studied in [19].

We show that when the target is present, the proposed universal test based on the aforementioned decomposition yields a vanishing error probability, and achieves the *optimal* reliability, in terms of a suitable exponent for the error probability with respect to the expected stopping time, *universally* for every target distribution. Consequently, we establish that the knowledge of the target distribution is only useful for improving the reliability for detecting that the target is missing. We also show that a multiplicative gain for the search reliability equal to the number of searched locations is achieved by allowing adaptivity in the search. We also would like to stress that our results and proof methods hold only in the scope of discrete observations each coming from a fixed finite set, wherein a certain probability mass function of the finite set could be unknown. This setting is quite different from the typical non-parametric setting in statistics, wherein continuous observations are often considered. Consequently, since our universal optimality claim mainly concerns the asymptotic regime of vanishing probability of error, in this regime, the negate of the exponent of the error is typically much larger than the finite size of the observation space.

Another line of research that is related to the current work is adaptive sensing for sparse recovery [5,6,12]. Instead of focusing on the case of the single signal support, corresponding to the unique target as in our setting, these authors considered general sparse sets of the signal support. On the other hand, they considered only the Gaussian case and did not cover the fully non-parametric universal setting (albeit for only finite observation spaces) like in our study.

We review the pertinent result on the sequential design of experiments in Section 2.1. The general model for universal search and stop is set up in Section 3. We present the universal sequential test for search and stop and state the main results pertaining to its asymptotic performance in Section 4.

2. Preliminaries

Throughout the paper, random variables (r.v.s) are denoted by capital letters, and their realizations are denoted by the corresponding lower-case letters. All r.v.s are assumed to take values in finite sets, and all logarithms are the natural ones. For a finite set \mathcal{X} , and a probability mass function (p.m.f.) p on \mathcal{X} we write $X \sim p$ to denote that the r.v. X is distributed according to p .

The following technical facts will be useful; their derivations can be found in [11], Chapter 11. Consider random variables $Y^n = (Y_1, \dots, Y_n)$ which are independent and identically distributed (i.i.d.) according to a p.m.f. p on a finite set \mathcal{Y} , that is, $Y_i \sim p, i = 1, \dots, n$. Let $y^n = (y_1, \dots, y_n) \in \mathcal{Y}^n$ be a sequence with an empirical distribution $\gamma = \gamma^{(n)}$ on \mathcal{Y} . It follows that the probability of such sequence y^n , under the i.i.d. assumption according to the p.m.f. p , is

$$\begin{aligned} p(y^n) &= \prod_{y \in \mathcal{Y}} p(y)^{|\{i=1, \dots, n, Y_i=y\}|} \\ &= e^{-n(\sum_{y \in \mathcal{Y}} |\{i=1, \dots, n, Y_i=y\}|/n \log(1/p(y)))} \\ &= e^{-n(\sum_{y \in \mathcal{Y}} \gamma(y) \log(\gamma(y)/p(y)) + \sum_{y \in \mathcal{Y}} \gamma(y) \log(1/\gamma(y)))} \\ &= e^{-n[D(\gamma \| p) + H(\gamma)]}, \end{aligned} \tag{2.1}$$

where $D(\gamma \| p)$ and $H(\gamma)$ are the relative entropy of γ and p , and entropy of γ , defined as

$$D(\gamma \| p) \triangleq \sum_{y \in \mathcal{Y}} \gamma(y) \log \frac{\gamma(y)}{p(y)},$$

and

$$H(\gamma) \triangleq - \sum_{y \in \mathcal{Y}} \gamma(y) \log \gamma(y),$$

respectively. Consequently, it holds that for each y^n , the p.m.f. p that maximizes $p(y^n)$ is $p = \gamma$, and the associated maximal probability of y^n is

$$\gamma(y^n) = e^{\lceil -nH(\gamma) \rceil}. \tag{2.2}$$

Next, for each $n \geq 1$, the number of all possible empirical distributions from a sequence of length n in \mathcal{Y}^n is upper bounded by $(n + 1)^{|\mathcal{Y}|}$ (cf. [11], Theorem 11.1.1). In particular, using this and the fact that the number of all sequences of length n with a feasible empirical distribution γ is

upper bounded by $e^{nH(\gamma)}$ (cf. [11], Theorem 11.1.3), we get that for any $\varepsilon > 0$, it holds that the probability of the i.i.d. sequence Y^n under p satisfies

$$\begin{aligned} \mathbb{P}[D(\gamma \| p) \geq \varepsilon] &= \sum_{\substack{\gamma: \text{feasible empirical} \\ \text{distribution in } \mathcal{Y}^n \\ \text{with } D(\gamma \| \pi) \geq \varepsilon}} \sum_{\substack{y^n \text{ with empirical} \\ \text{distribution} = \gamma}} e^{-n[D(\gamma \| p) + H(\gamma)]} \\ &\leq (n + 1)^{|\mathcal{Y}|} e^{nH(\gamma)} e^{-n[\varepsilon + H(\gamma)]} \\ &= (n + 1)^{|\mathcal{Y}|} e^{-n\varepsilon}. \end{aligned} \tag{2.3}$$

We now review the relevant result on the model-based sequential design of experiments with a varying experimental cost. This result will be key to our proposed universal test for search and stop.

2.1. Sequential design of experiments with varying experimental cost

Consider the problem of the sequential design of experiments to facilitate the eventual testing for H hypotheses. We assume a (conditionally) memoryless model for the experimental outcome conditioned on the currently chosen experiment. In particular, under the i th hypothesis, $i \in \{1, \dots, H\} = [H]$, and conditioned on the current experiment $u_t = u \in \mathcal{U}$, at time $t = 1, 2, \dots$, the current outcome of the experiment, denoted by Z_t , is assumed to be conditionally independent of all past outcomes and past experiments Z^{t-1}, U^{t-1} , and to be conditionally distributed according to a p.m.f. p_i^u on \mathcal{Z} . There is a cost function $c : [H] \times \mathcal{U} \rightarrow \mathbb{R}^+$, and the current experiment u_t is assumed to incur a cost of $c(i, u_t)$ under the i th hypothesis. We assume that for every $i = 1, \dots, H, u \in \mathcal{U}, z \in \mathcal{Z}, p_i^u(z) > 0, c(i, u) > 0$. A test consists of an adaptive policy ϕ that chooses each experiment as a suitable (possibly randomized) function of past experiments and their outcomes, a stopping time τ , and a final decision rule δ that outputs a guess of the true hypothesis in $[H]$. In particular, a policy ϕ specifies the experiments $U_t, t = 1, 2, \dots$, according to conditional distributions $q(u_1), q(u_t | z^{t-1}, u^{t-1}), t = 2, 3, \dots$. The joint distribution of all experiments and their outcomes up to time n under hypothesis $i = 1, \dots, H$, is given by $q(u_1) p_i^{u_1}(z_1) [\prod_{t=2}^n q(u_t | z^{t-1}, u^{t-1}) p_i^{u_t}(z_t)]$. The goal is to design a test to optimize the tradeoff between the cost accumulated up to the stopping time, as measured by $\sum_{t=1}^{\tau} c(i, U_t), i \in [H]$, and the accuracy of the final decision, as measured by $P_{\max} \triangleq \max_{i=1, \dots, H} \mathbb{P}_i[\delta(Z^\tau) \neq i]$. The problem is model-based: all the (conditional) distributions $p_i^u, i \in [H], u \in \mathcal{U}$, and the cost function c are assumed to be known.

For each hypothesis $i \in [H]$, let

$$q_i^*(u) \triangleq \operatorname{argmax}_q \frac{\min_{j \neq i} \sum_u q(u) D(p_i^u \| p_j^u)}{\sum_u q(u) c(i, u)}. \tag{2.4}$$

Then an asymptotically optimal test can be specified based on these distributions as follows. At each time $t \geq 1$, the maximum likelihood (ML) estimate of the true hypothesis \hat{i} can be computed based on past experiments and their outcomes u^{t-1}, z^{t-1} using the model $p_i^u, i \in [H], u \in \mathcal{U}$

(ties are broken arbitrarily). For $b > 0$, during the sparse occasions $t = \lfloor e^{b\ell} \rfloor, \ell = 0, 1, \dots$, the experiment is selected to explore all possible options in \mathcal{U} in a round-robin manner independently of \hat{i} : for $\mathcal{U} = \{u_1, \dots, u_{|\mathcal{U}|}\}$,

$$u_t = u_{(\ell \bmod |\mathcal{U}|)+1}. \tag{2.5}$$

At all other times, the current (random) experiment is selected as $U_t \sim q_{\hat{i}}^*$. Denote the joint distribution under the i th hypothesis of all experiments and their outcomes up to time t (induced by the control policy) by $p_i(z^t, u^t)$. This policy affords a suitable tradeoff between exploiting the control policy (2.4) that best facilitates detection and exploring among all possible experiments (2.5) to ensure the consistency of the ML estimate of the hypothesis in closed loop [18,19]. Had we not performed the sparse exploration in (2.5), the ML estimate of the hypothesis might not converge to the true hypothesis or not quickly enough. This is because a particular experiment may only be useful to differentiate *some* hypotheses from the others, but not to differentiate among *all* hypotheses. We have to be careful when we close the loop between the ML estimation and using it to select the instantaneous experiment.

For a threshold $a' > 1$, the test stops at time τ^* and decides in favor of the ML hypothesis according to the rule δ^* , where

$$\tau^* \triangleq \operatorname{argmin}_t \min_{j \neq \hat{i}} \frac{p_i(z^t, u^t)}{p_j(z^t, u^t)} > a', \quad \delta^*(z^{\tau^*}, u^{\tau^*}) = \hat{i}. \tag{2.6}$$

Note that as the $q_i^*, i \in [H]$, are, in general, not point-mass distributions, in addition to the realization of all experimental outcomes z^t , we also need to account for the realization of the experiments u^t as well in the instantaneous computation of the ML hypothesis and checking the stopping criterion (2.6). If the experiments have been chosen deterministically at all times, we can just use the joint distributions of all experimental outcomes $p_i(z^t), i \in [H], t = 1, 2, \dots$, in these computations. The resulting test is asymptotically optimal and its performance is characterized in Proposition 2.1 as follows.

Proposition 2.1 ([19])¹. *For $b > 0$, in (2.5) chosen to be sufficiently small, and as $a' \rightarrow \infty$, the test in (2.4), (2.5), (2.6) yields a vanishing error probability $P_{\max}^* \rightarrow 0$, and satisfies for each $i = 1, \dots, H$, that*

$$\mathbb{E}_i \left[\sum_{t=1}^{\tau^*} c(i, U_t) \right] = \frac{-\log P_{\max}^*}{\max_q \frac{\min_{j \neq i} \sum_u q(u) D(p_i^u \| p_j^u)}{\sum_u q(u) c(i, u)}} (1 + o(1)).$$

¹The result in [19] was proven for the model in which the cost function depends only on the experiment; however, the proof generalizes to the current setting when the cost function also depends on the hypothesis.

In addition, the proposed test is asymptotically optimal simultaneously under all hypotheses in the sense that any sequence of tests (ϕ, τ, δ) that achieve $P_{\max} \rightarrow 0$ must satisfy

$$\mathbb{E}_i \left[\sum_{t=1}^{\tau} c(i, U_t) \right] \geq \frac{-\log P_{\max}}{\max_q \frac{\min_{j \neq i} \sum_u q(u) D(p_i^u \| p_j^u)}{\sum_u q(u) c(i, u)}} (1 + o(1)),$$

for every $i = 1, \dots, H$.

3. Model for search and stop

Consider searching for a single target located in one of M locations. At each time $k \geq 1$, if a location without the target is searched, then the observation $Y_k \in \mathcal{Y}$ is assumed to be conditionally independent of all past observations and past search locations, and to be conditionally distributed according to the absence distribution π . The distribution π represents pure noise, and we shall assume that this distribution is *known* to the searcher. On the other hand, if the target location is searched, then the observation would be conditionally distributed according to the “target” distribution μ (same no matter where the target is) on \mathcal{Y} (and would be conditionally independent of past observations and search locations). We assume that both μ and π have full supports on \mathcal{Y} . Other than this assumption, the searcher has no knowledge of the target distribution and it could be arbitrarily close to the absence distribution.

We also allow for the possibility that the target is absent. In this latter case, the observations at all locations are distributed according to π . Denote the search location at time $k \geq 1$ by $U_k \in [M]$, which is allowed to be any function of all past observations $Y^{k-1} = (Y_1, \dots, Y_{k-1})$ and past search locations $U^{k-1} = (U_1, \dots, U_{k-1})$.

It is interesting to note that the most basic search problem with an overlook probability $\alpha > 0$ (see, e.g., Chapters 4, 5 of [21] and Section 4.2 of [3]) that is same over all locations, corresponds to a special case of our general model wherein $\mathcal{Y} = \{0, 1\}$, $\mu(0) = \alpha$, $\pi(0) = 1$. In contrast, our model allows for any general (finite) observation set \mathcal{Y} , but assumes that both μ and π have full supports. The degeneracy in the model for the classic search problem as mentioned affords the construction of a search plan that is more efficient than that for our model (with the assumption of the full support). The main concern for the classic search problem has been to come up with the search plan that is absolutely optimal (non-asymptotically), whereas our main concern is to construct a *universal* test that is asymptotically efficient in the regime of vanishing error probability.

We seek to design a universal sequential test to search the target (or to decide that it is missing). Precisely speaking, a test consists of a sequential search policy, a stopping rule and a final decision rule. The stopping rule defines a stopping time, denoted by N , which is the number of searches taken until the final decision is made. At the stopping time, the final decision for the target location is made based on the decision rule $\delta : \mathcal{Y}^N \times [M]^N \rightarrow \{0, 1, \dots, M\}$, where the 0 output corresponds to the final decision that the target is missing. The overall goal is to achieve a certain level of accuracy for the final decision using the fewest number of observations, universally for all $\mu \neq \pi$.

3.1. Fundamental performance limit

When both μ and π are known, the search and stop problem falls under the umbrella of the sequential design of experiments with a uniform experimental cost [7]. In particular, there are $M + 1$ hypotheses: $0, 1, \dots, M$, where the null (0th) hypothesis corresponds to the case that the target is missing. Each i th hypothesis, $i = 1, \dots, M$, corresponds to each possible location of the present target. The experiment set corresponds to $\mathcal{U} = [M]$, and each experiment corresponds to where to search next. The model $p_i^u(y), i = 0, \dots, M$, for the sequential design of experiments can be identified as

$$\begin{aligned} p_i^u &= \mu, & u &= i, & p_i^u &= \pi, & u &\neq i, & i &= 1, \dots, M, \\ p_0^u &= \pi, & u &= 1, \dots, M. \end{aligned} \tag{3.1}$$

Then in this idealistic situation when the probabilistic model (both μ and π) is known, by particularizing the characterization of the asymptotically optimal performance in Proposition 2.1 to our search and stop problem using (3.1) and $c(i, u) = 1$, for each $i = 0, \dots, M$, and for each $u \in \mathcal{U}$, we get that as the error probability $P_{\max} = \max_{i=0, \dots, M} \mathbb{P}_i[\delta^*(Y^{N^*}, U^{N^*}) \neq i]$ is driven to zero, the optimal asymptotes of $\mathbb{E}_i[N^*], i = 0, \dots, M$, can be characterized as follows.

Proposition 3.1. *There exists a sequence of tests to search the target with a stopping time N^* that satisfy $P_{\max}^* \rightarrow 0$ and yield*

$$\mathbb{E}_i[N^*] = \begin{cases} \frac{-\log P_{\max}^*}{D(\pi \parallel \mu)/M} (1 + o(1)), & i = 0, \\ \frac{-\log P_{\max}^*}{D(\mu \parallel \pi)} (1 + o(1)), & i = 1, \dots, M. \end{cases} \tag{3.2}$$

Furthermore, the asymptotic performance in (3.2) (each term in the denominators) is optimal for every $i = 0, \dots, M$ simultaneously. In particular, any sequence of tests with a stopping time N that achieve $P_{\max} \rightarrow 0$ must simultaneously satisfy

$$\begin{aligned} \mathbb{E}_0[N] &\geq \frac{-\log P_{\max}}{D(\pi \parallel \mu)/M} (1 + o(1)), \\ \mathbb{E}_i[N] &\geq \frac{-\log P_{\max}}{D(\mu \parallel \pi)} (1 + o(1)), & i &= 1, \dots, M. \end{aligned}$$

Of course, the asymptotic performance in Proposition 3.1 is idealistic, as it requires the knowledge of μ (with π being already known). When μ is not known, since μ can be arbitrarily close to π , this asymptotic performance cannot be achieved universally. Nevertheless, our main contribution (Theorem 4.1) described below shows that one can design a universal test (without the knowledge of μ) that drives the error probability to zero and achieves the optimal exponent of $D(\mu \parallel \pi)$ under all the *non-null* hypotheses universally for any $\mu \neq \pi$.

4. Proposed universal scheme for search and stop and its performance

Since we assume that the observations at all locations without the target follow the same absence distribution, a desirable goal of the search at each location should be to determine if the target is there or not (elsewhere or missing entirely). To this end, a universal sequential test for two hypotheses can be used at each location to collect multiple subsequent observations that will eventually lead to a binary outcome (say 1 if it is guessed that the target is there, and 0 otherwise). We now discuss this universal sequential testing for two hypotheses.

4.1. Universal sequential testing for two hypotheses at each location

Consider sequential testing between the null hypothesis H_0 (for the case that the target is not at the currently searched location) with i.i.d. observations $Y_k \in \mathcal{Y}, k = 1, 2, \dots$, according to a p.m.f. π on \mathcal{Y} , and the alternative hypothesis H_1 (for the case that the target is at the currently searched location) with i.i.d. $Y_k, k = 1, 2, \dots$, according to a p.m.f. $\mu \neq \pi$, where only π is known, and nothing is known about μ (except that both μ and π have full supports on \mathcal{Y}).

The goal for the binary testing at a particular searched location is to construct a test consisting of a stopping time N^b and a binary decision rule δ_b to efficiently decide whether the target is there at the searched location or not (elsewhere or missing entirely). For a parameter $a > 1$, our goal is to construct a universal sequential test for the two hypotheses with the maximal error probability being $\max(\mathbb{P}_0[\delta_b(N^b) = 1], \mathbb{P}_1[\delta_b(N^b) = 0]) \leq \frac{1}{a}$. To this end, we shall employ a sequential binary test defined in terms of the following (Markov) time:²

$$\tilde{N}^b \triangleq \operatorname{argmin}_{n \geq 1} [nD(\gamma \parallel \pi) > (\log a + n^{2/3} + |\mathcal{Y}| \log(n + 1))], \tag{4.1}$$

where γ denotes the empirical distribution of the observation sequence (y_1, \dots, y_n) . The test stops at this time or $\lfloor a(\log a)^{\rho_1} \rfloor$ for some $\rho_1 > 1$, depending on which one is smaller, that is, it stops at time N^b , where

$$N^b \triangleq \min(\tilde{N}^b, \lfloor a(\log a)^{\rho_1} \rfloor). \tag{4.2}$$

Correspondingly, the final decision is made according to

$$\delta_b(Y^{N^b}) = \begin{cases} 1, & \text{if } \tilde{N}^b \leq a(\log a)^{\rho_1}, \\ 0, & \text{if } \tilde{N}^b > a(\log a)^{\rho_1}. \end{cases} \tag{4.3}$$

The stopping time N^b in (4.2) means that for a parameter $a > 1$, the search at each particular location would take no more than $\lfloor a(\log a)^{\rho_1} \rfloor$ subsequent observations before deciding whether

²The term $n^{2/3}$ on the right-hand side of (4.1) below is necessary to obtain the upper bound of $\frac{1}{a}$ in (4.4). With only the $\log(n + 1)$ term, we can only achieve an upper bound of $\frac{K}{a}$, where $K > 1$, for an arbitrarily large coefficient in front of the $\log(n + 1)$ term. See also equation (A.3) in the proof of Lemma 4.1.

the target is there or not. In particular, the decision is made based on whether the time-dependent threshold $\log a + n^{2/3} + |\mathcal{Y}|\log(n + 1)$ in (4.1) is crossed by $nD(\gamma\|\pi)$, which is supposed to be large when the target is at the searched location, after the time $n = \lfloor a(\log a)^{\rho_1} \rfloor$ or not. As a grows, the error in the binary decision making is smaller, but the number of subsequent observations taken at the location gets larger as well.

Lemma 4.1. *With μ and π having full supports on \mathcal{Y} , for some $\rho_1 > 1$ and for every $a > 1$, the sequential test in (4.1), (4.2), (4.3) yields that*

$$\alpha_a \triangleq \mathbb{P}_0[\delta_b(Y^{N^b}) = 1] \leq \frac{1}{a}. \tag{4.4}$$

In addition, for any $v < 1$, $\mu \neq \pi$ and every $a \geq a^(v, \mu, \pi)$, the test also yields that*

$$c_a \triangleq \mathbb{E}_1[N^b] \leq \mathbb{E}_1[\tilde{N}^b] \leq \frac{\log a}{vD(\mu\|\pi)}, \tag{4.5}$$

$$\kappa_a \triangleq \mathbb{E}_0[N^b] \leq a(\log a)^{\rho_1}, \tag{4.6}$$

$$\beta_a \triangleq \mathbb{P}_1[\delta_b(Y^{N^b}) = 0] \leq \frac{1}{vD(\mu\|\pi)a(\log a)^{(\rho_1-1)}}. \tag{4.7}$$

The proof of Lemma 4.1 will be given in Section A.1.

4.2. Proposed universal test for search and stop

If we use the mentioned sequential binary test at each location as the “inner” test, then it is convenient to select the current search location based on the past binary outcomes of the subsequent binary tests (instead of the past \mathcal{Y} -ary outcomes of the individual searches, often taken multiple times at each of the locations). With this imposition, the search and stop problem can be reduced to a problem of constructing an “outer” test for the sequential design of such inner experiments, each of which has a binary outcome.

Mathematically speaking, we have reduced the original problem of sequential design of \mathcal{Y} -ary-output experiments specified by the (conditional) distributions as in (3.1) to one of the sequential design of binary-output experiments specified as

$$\mu_b(0) = 1 - \mu_b(1) = \beta_a, \quad \pi_b(1) = 1 - \pi_b(0) = \alpha_a, \tag{4.8}$$

and

$$\begin{aligned} p_i^u &= \mu_b, & u &= i, & p_i^u &= \pi_b, & u &\neq i, & i &= 1, \dots, M, \\ p_0^u &= \pi_b, & u &= 1, \dots, M, \end{aligned} \tag{4.9}$$

where α_a, β_a are as defined in (4.4), (4.7), respectively. On the other hand, each binary-output experiment will not have the same cost as for the original \mathcal{Y} -ary-output experiment. In particular,

the cost of each binary-output experiment can be specified as

$$\begin{aligned} c(i, u) &= c_a, & u = i, & & c(i, u) &= \kappa_a, & u \neq i, & & i = 1, \dots, M, \\ c(0, u) &= \kappa_a, & u = 1, \dots, M, \end{aligned} \tag{4.10}$$

where c_a, κ_a are as defined in (4.5) and (4.6), respectively.

There is still a large gap in turning the motivation described above into a “working” test for search and stop. To this end, there are two major challenges. First, the optimal test for the sequential design of experiments in (2.4), (2.5), (2.6), achieving the performance stated Proposition 3.1, requires precise knowledge of the model. In contrast, the induced model for the sequential design of binary-output experiments in (4.8), (4.9), (4.10) is a complicated function of the inner threshold a for the sequential binary test in (4.1), (4.2), (4.3). Only an estimate of this “true” induced model is available through the bounds for $\alpha_a, \beta_a, c_a, \kappa_a$ stated in (4.4), (4.7), (4.5), (4.6) of Lemma 4.1, respectively. Second, as the threshold a' for the optimal test in (2.4), (2.5), (2.6) increases, the model for the sequential design of experiments in Proposition 3.1 remains fixed. In contrast, in our proposed test, the “outer” threshold for the test for the sequential design of binary-output experiments increases together with the inner threshold a , the latter of which determines the induced model in (4.8), (4.9), (4.10). Consequently, the analysis leading to Proposition 3.1 does not apply to our proposed test. Our main technical contributions are precisely, first, to overcome these challenges through the proposed test described below in Section 4.2, employing an outer threshold, which is chosen to be an appropriate function of the inner threshold, and, second, to provide the analysis for its performance, stated in Theorem 4.1 below.

As mentioned in the previous subsection, the “true” induced model for the sequential design of the binary-output experiments in (4.8), (4.9), (4.10) is a complicated function of the inner threshold a and is not available to us. Nevertheless, Lemma 4.1 yields that for a sufficiently large (as a function of ν , in (4.5), (4.7) and μ, π),

$$\alpha_a, \beta_a \leq \frac{1}{a}. \tag{4.11}$$

Our idea would be to use a mismatched model defined in terms of $\bar{\mu}_b, \bar{\pi}_b$, where

$$\bar{\mu}_b(0) = 1 - \bar{\mu}_b(1) = \bar{\pi}_b(1) = 1 - \bar{\pi}_b(0) = \frac{1}{a} \tag{4.12}$$

to perform the sequential design of the binary-output experiments. Specifically, instead of (4.8), (4.9), consider the following mismatched model for the sequential design of binary-output experiments

$$\begin{aligned} \bar{p}_i^u &= \bar{\mu}_b, & u = i, & & \bar{p}_i^u &= \bar{\pi}_b, & u \neq i, & & i = 1, \dots, M, \\ \bar{p}_0^u &= \bar{\pi}_b, & u = 1, \dots, M. \end{aligned} \tag{4.13}$$

Heuristically speaking, by (4.11), this mismatched model is “more noisy” than the true model (for large a); hence, the test designed based on this mismatched model should be conservative enough to work well for the true model as well. This intuition will be proven to be correct.

With the mismatched model specified in (4.13), we can now describe our universal test as follows. At each time $t \geq 1$, we compute the estimate of the true hypothesis \hat{i} based on past searched locations and their binary outcomes u^{t-1}, z^{t-1} using the (mismatched) model $\bar{p}_i^u, i = 0, \dots, M, u \in [M]$ in (4.13). Denote $N(i, 1), N(i, 0), i \in [M]$, as the number of times the i th location was searched and the sequential binary test in (4.1), (4.2), (4.3) decides that the target is there, and that the target is not there, respectively. By the reciprocity of $\bar{\mu}_b$ and $\bar{\pi}_b$, in (4.12), the computation of this estimate can be simplified as

$$\hat{i} = \begin{cases} \operatorname{argmax}_{i \in [M]} N(i, 1) - N(i, 0), & \text{if } \max_{i \in [M]} N(i, 1) - N(i, 0) > 0, \\ 0, & \text{if } \max_{i \in [M]} N(i, 1) - N(i, 0) \leq 0. \end{cases} \tag{4.14}$$

The estimation in (4.14) is quite intuitive, as the difference between the numbers of “searched-and-found” and “searched-and-not-found” at the i th location: $N(i, 1) - N(i, 0), i \in [M]$, should approximate the likelihood that the target is there. When all these numbers are negative, it is most likely that the target is missing.

For $b > 0$, during the sparse occasions $t = \lfloor e^{b\ell} \rfloor, \ell = 0, 1, \dots$, the experiment is selected to explore all locations in a round-robin manner as

$$u_t = (\ell \bmod M) + 1 \tag{4.15}$$

independently of \hat{i} . At all the other times, if $\hat{i} \neq 0$, we shall search at the \hat{i} th location, i.e.,

$$u_t = \hat{i}, \quad \text{if } \hat{i} \neq 0. \tag{4.16}$$

If $\hat{i} = 0$, we search among all locations with equal frequency, namely,

$$u_t = (i_{t'} \bmod M) + 1, \tag{4.17}$$

where $i_{t'}$ was the search location at the last time $t' < t$ such that $\hat{i} = 0$. Denote the joint (mismatched) distribution under the i th hypothesis of all binary searched outcomes up to time t (induced by the above control policy) by $\bar{p}_i(z^t)$. The test stops at time τ and decides in favor of the current estimate of the hypothesis as:

$$\tau \triangleq \operatorname{argmin}_t \left[\left(\min_{\substack{j=0, \dots, M, \\ j \neq \hat{i}}} \frac{\bar{p}_i(z^t)}{\bar{p}_j(z^t)} \right) > e^{a^{\rho_2} (\log a)^{\rho_1}} \right], \quad \delta(z^\tau) = \hat{i}, \tag{4.18}$$

where a is the inner threshold for the binary test in (4.1), (4.2), (4.3) and some $\rho_2 > 1$. As clarified at the end of the paragraph preceding Proposition 2.1, since the search policy in (4.15), (4.16), (4.17) specifies the search location as a deterministic function of the current estimate of the hypothesis at all times, it suffices to work with the joint distribution $\bar{p}_i(z^t)$ instead of $\bar{p}_i(z^t, u^t)$, that is, u^t can be written as a deterministic function of z^t .

Using (4.12), (4.13), we can simplify (4.18) as

$$\tau = \operatorname{argmin}_t \left[\min_{j \neq \hat{i}} (S(\hat{i}) - S(j)) > \frac{a^{\rho_2} (\log a)^{\rho_1}}{\log(a - 1)} \right], \quad \delta(z^\tau) = \hat{i}, \tag{4.19}$$

where $S(i), i = 0, \dots, M$, represent indices of the corresponding hypotheses defined as

$$S(i) = \begin{cases} 0, & i = 0, \\ N(i, 1) - N(i, 0), & i = 1, \dots, M. \end{cases} \tag{4.20}$$

Note that the total number of \mathcal{Y} -ary-output observations N used to produce the search result is related to the stopping time τ above as

$$N = \sum_{t=1}^{\tau} N_t^b,$$

where each $N_t^b, t = 1, \dots, \tau$, is the number of observations taken at each location until the sequential test in (4.1), (4.2), (4.3) produces a binary result Z_t . Consequently, we get from successive uses of the property of conditional expectation and (4.10) that under the true hypothesis $i = 0, \dots, M$, it holds that

$$\begin{aligned} \mathbb{E}_i[N] &= \mathbb{E}_i \left[\sum_{t=1}^{\tau-1} N_t^b + \mathbb{E}_i[N_{\tau}^b | Y^{(\sum_{t=1}^{\tau-1} N_t^b)}] \right] \\ &= \mathbb{E}_i \left[\sum_{t=1}^{\tau-1} N_t^b + \mathbb{E}_i[N_{\tau}^b | U_{\tau}] \right] \\ &= \mathbb{E}_i \left[\sum_{t=1}^{\tau-1} N_t^b + c(i, U_{\tau}) \right], \quad \text{by (4.10), (4.5), (4.6)} \\ &= \mathbb{E}_i \left[\sum_{t=1}^{\tau} c(i, U_t) \right]. \end{aligned}$$

4.3. Performance of proposed test

Theorem 4.1. *For any $v < 1$ in (4.5), (4.7) and for $b > 0$ used in (4.15) chosen to be sufficiently small, as $a \rightarrow \infty$, the test in (4.14), (4.15), (4.16), (4.17), (4.18) yields a vanishing error probability $P_{\max} = \max_{i=0, \dots, M} \mathbb{P}_i[\delta \neq i] \rightarrow 0$ and also satisfies simultaneously for each hypothesis $i = 1, \dots, M$, with a present target that*

$$\mathbb{E}_i[N] = \mathbb{E}_i \left[\sum_{t=1}^{\tau} c(i, U_t) \right] \leq \frac{-\log P_{\max}}{vD(\mu || \pi)} (1 + o(1)), \quad i = 1, \dots, M, \tag{4.21}$$

universally for every $\mu \neq \pi$.

Remark 4.1. Compared to the idealistically optimal performance (when μ is known) in Proposition 3.1, it is interesting to note that our universal test is universally asymptotically optimal,

except only when the target is missing. In other words, the knowledge of the target distribution is only useful in improving reliability for detecting that the target is missing. This consequence of our result is directly relevant to practical settings, wherein the knowledge of the target distribution μ would be lacking before the target is found.

4.4. Comparison with universal non-adaptive scheme for search and stop

Our main result in Theorem 4.1 illustrates that one can construct a test with adaptive search policy, using only the knowledge of π , that yields a vanishing error probability and achieves the exponent of $D(\mu\|\pi)$ universally for every $\mu \neq \pi$ when the target is present. A natural question that arises is how much can be gained by employing such an adaptive search policy beyond a non-adaptive one. A non-adaptive search policy $\bar{\phi}$ has to specify the sequence of search locations at the outset and cannot adapt to the outcomes of the instantaneous searches. By the symmetry of the problem, there is no reason for a non-adaptive search policy to favor any location. Consequently, the only non-adaptive search policy that should be considered in the universal setting is the one that searches all locations with equal frequency:

$$u_k = (k \bmod M) + 1, \quad k \geq 0. \quad (4.22)$$

We denote this non-adaptive search policy by $\bar{\phi}^*$. With this search policy, an efficient universal test has been constructed in [16], which we now describe.

For each time $k = \ell M$, $\ell = 1, 2, \dots$, let γ_i , $i = 1, \dots, M$ denote the empirical distribution of the observations when the i th location is searched, namely, $\gamma_i = (y_i, y_{M+i}, \dots, y_{(\ell-1)M+i})$, $i = 1, \dots, M$. Next, denote the estimate of the target location \hat{i} as

$$\hat{i} = \operatorname{argmax}_{i \in [M]} D(\gamma_i \|\pi). \quad (4.23)$$

With the non-adaptive search policy $\bar{\phi}^*$, consider the stopping rule defined in terms of the following Markov time:

$$\bar{N}' \triangleq M \times \operatorname{argmin}_{\ell \geq 1} \left[\left(D(\gamma_{\hat{i}} \|\pi) - \max_{j \neq \hat{i}} D(\gamma_j \|\pi) \right) > \log \bar{a} + M |\mathcal{Y}| \log(\ell + 1) \right]. \quad (4.24)$$

The test stops at time \bar{N} , where

$$\bar{N} \triangleq \min(\bar{N}', \lfloor \bar{a} \log \bar{a} \rfloor). \quad (4.25)$$

Correspondingly, the final decision is made according to

$$\bar{\delta}(Y^{\bar{N}}) = \begin{cases} \hat{i}, & \text{if } \bar{N}' \leq \bar{a} \log \bar{a}, \\ 0, & \text{if } \bar{N}' > \bar{a} \log \bar{a}. \end{cases} \quad (4.26)$$

The performance of this test with the non-adaptive search scheme follows from the result in [16].

Proposition 4.1 ([16]). *With the non-adaptive search policy $\bar{\phi}^*$ in (4.22), the test in (4.23), (4.24), (4.25), (4.26) yields a vanishing error probability $P_{\max} \rightarrow 0$ and also satisfies*

$$\mathbb{E}_i[\bar{N}] \leq \frac{-\log P_{\max}}{D(\mu\|\pi)/M}(1 + o(1)), \quad i = 1, \dots, M, \tag{4.27}$$

universally for every $\mu \neq \pi$.

In summary, adaptivity offers a multiplicative gain of M for search reliability beyond non-adaptive searching. This gain increases with the size of the area to be searched.

Remark 4.2. Note that for the test in (4.22), (4.23), (4.24), (4.25), (4.26), the maximal error probability P_{\max} decays only subexponentially with $\mathbb{E}_0[\bar{N}]$. The additional exponential decay cannot be obtained on top of the ones for all the other hypotheses with a present target (cf. (4.27)), because we are considering the universal setting in which only π is known (contrary to the result in Proposition 3.1 with both μ and π being known). For any sequence of universal tests, a bad μ can always be selected to be sufficiently close to π to render such a subexponential decay.

Appendix

A.1. Proof of Lemma 4.1

The proof relies on the following lemmas.

Lemma A.1 ([15]). *For any p.m.f.s μ, π on \mathcal{Y} with full supports and with $B(\mu, \pi)$ being the Bhattacharyya distance between μ and π defined as $B(\mu, \pi) \triangleq -\log(\sum_{y \in \mathcal{Y}} \mu(y)^{1/2} \pi(y)^{1/2})$ [11], it holds that*

$$2B(\mu, \pi) = \min_q (D(q\|\mu) + D(q\|\pi)), \tag{A.1}$$

where the minimum above is over all p.m.f.s on \mathcal{Y} .

Lemma A.2. *Under the alternative hypothesis, it holds for every $n \geq 1$, that*

$$\mathbb{P}_1[\tilde{N}^b \geq n] \leq a e^{-(n-1)2B(\mu, \pi)} n^{2|\mathcal{Y}|} e^{(n-1)^{2/3}}.$$

Proof.

$$\begin{aligned} \mathbb{P}_1[\tilde{N}^b \geq n] &\leq \mathbb{P}_1[(n-1)D(\gamma\|\pi) \leq \log a + (n-1)^{2/3} + |\mathcal{Y}|\log(n)] \\ &= \mathbb{P}_1\left[D(\gamma\|\mu) \geq -\frac{(\log a + (n-1)^{2/3} + |\mathcal{Y}|\log n)}{n-1} + D(\gamma\|\mu) + D(\gamma\|\pi)\right] \\ &\leq \mathbb{P}_1\left[D(\gamma\|\mu) \geq -\frac{(\log a + (n-1)^{2/3} + |\mathcal{Y}|\log n)}{n-1} + 2B(\mu, \pi)\right] \\ &\leq a e^{-(n-1)2B(\mu, \pi)} n^{2|\mathcal{Y}|} e^{(n-1)^{2/3}}, \end{aligned}$$

where the second inequality follows from Lemma A.1 and the last inequality follows from (2.3). \square

First, we prove (4.4). It follows from (4.2), (4.3) that

$$\begin{aligned} \mathbb{P}_0[\delta = 1] &= \mathbb{P}_0[N^b = \tilde{N}^b] \\ &\leq \mathbb{P}_0[\tilde{N}^b \leq a(\log a)^{\rho_1}] \leq \sum_{n=1}^{\infty} \mathbb{P}_0[\tilde{N}^b = n] \\ &\leq \sum_{n=1}^{\infty} \mathbb{P}_0[nD(\gamma \parallel \pi) > (\log a + n^{2/3} + |\mathcal{Y}| \log(n + 1))] \\ &\leq \sum_{n=1}^{\infty} \frac{1}{a} e^{-n^{2/3}} \tag{A.2} \\ &\leq \frac{1}{a}, \tag{A.3} \end{aligned}$$

where (A.2) follows from (2.3). In addition (A.3) follows from the fact that

$$\begin{aligned} \sum_{n=1}^{\infty} e^{-n^{2/3}} &\leq \sum_{n=1}^5 e^{-n^{2/3}} + \int_5^{\infty} e^{-x^{2/3}} dx \\ &= \sum_{n=1}^5 e^{-n^{2/3}} + 3 \int_{5^{1/3}}^{\infty} e^{-y^2} y^2 dy \\ &= \sum_{n=1}^5 e^{-n^{2/3}} - \frac{3}{2} \int_{5^{1/3}}^{\infty} y d(e^{-y^2}) \\ &= \sum_{n=1}^5 e^{-n^{2/3}} - \frac{3}{2} \left[-\frac{5^{1/3}}{e^{5^{2/3}}} - \int_{5^{1/3}}^{\infty} e^{-y^2} dy \right] \\ &= \sum_{n=1}^5 e^{-n^{2/3}} + \frac{3}{2} \frac{5^{1/3}}{e^{5^{2/3}}} + \frac{3\sqrt{\pi}}{4} \operatorname{erfc}(5^{1/3}) \leq 1. \end{aligned}$$

Equation (4.6) follows trivially from the definition of N^b in (4.2).

It now remains to prove the inequality in (4.5), as the inequality in (4.7) would follow from it and from the definition of N^b in (4.2) upon noting that

$$\begin{aligned} \beta_a &\triangleq \mathbb{P}_1[\delta_b(Y^{N^b}) = 0] = \mathbb{P}_1[\tilde{N}^b > a(\log a)^{\rho_1}] \\ &\leq \frac{\mathbb{E}_1[\tilde{N}^b]}{a(\log a)^{\rho_1}} \leq \frac{1}{vD(\mu \parallel \pi)a(\log a)^{(\rho_1-1)}}. \end{aligned}$$

To this end, it suffices to show that under the alternative hypothesis and as $a \rightarrow \infty$,

$$\frac{\mathbb{E}_1[\tilde{N}^b]}{\log a} \rightarrow \frac{1}{D(\mu\|\pi)}. \tag{A.4}$$

First observe that under the alternative hypothesis \mathbb{P}_1 , by the strong law of large numbers, we have that for every $y \in \mathcal{Y}$, $\frac{1}{n} \sum_{k=1}^n \mathbb{I}[Y_k = y]$ converges to $\mu(y)$ a.s. Consequently, we have that $\gamma^{(n)} \rightarrow \mu$ a.s. Next, we get from Lemma 4.1 that for every $a > 1$,

$$\mathbb{P}_1[\tilde{N}^b = \infty] \leq \lim_{n \rightarrow \infty} \mathbb{P}_1[\tilde{N}^b \geq n] = 0. \tag{A.5}$$

It then follows from (A.5) and (4.1) that

$$D(\gamma^{(\tilde{N}^b)}\|\pi) > \frac{\log a + (\tilde{N}^b)^{2/3} + |\mathcal{Y}| \log(\tilde{N}^b + 1)}{\tilde{N}^b}, \tag{A.6}$$

$$D(\gamma^{(\tilde{N}^b-1)}\|\pi) \leq \frac{\log a + (\tilde{N}^b - 1)^{2/3} + |\mathcal{Y}| \log(\tilde{N}^b)}{\tilde{N}^b - 1}. \tag{A.7}$$

Next, by observing that for any distribution q , $D(q\|\pi) \leq \log(\frac{1}{\min_y \pi(y)})$, we get from (A.6) that

$$\begin{aligned} \mathbb{P}_1[\tilde{N}^b \leq n] &\leq \mathbb{P}_1[\tilde{N}^b D(\gamma^{(\tilde{N}^b)}\|\pi) > \log a; \tilde{N}^b \leq n] \\ &\leq \mathbb{P}_1\left[n \log\left(\frac{1}{\min_y \pi(y)}\right) > \log a\right] \\ &= 0, \quad \text{for every } n < \frac{\log a}{\log(1/\min_y \pi(y))}, \end{aligned}$$

thereby yielding that $\tilde{N}^b \rightarrow \infty$ a.s. as $a \rightarrow \infty$ under \mathbb{P}_1 . Consequently, we conclude from the continuity of $D(\cdot\|\pi)$ that under \mathbb{P}_1 , $D(\gamma^{(\tilde{N}^b)}\|\pi), D(\gamma^{(\tilde{N}^b-1)}\|\pi) \rightarrow D(\mu\|\pi)$ a.s. as $a \rightarrow \infty$. We now get from this, (A.6), and (A.7) that

$$\frac{\tilde{N}^b}{\log a} \xrightarrow{\text{a.s.}} \frac{1}{D(\mu\|\pi)}. \tag{A.8}$$

To go from convergence a.s. (A.8) to convergence in mean (A.4), it now suffices to prove that the sequence of r.v.s $\frac{\tilde{N}^b}{\log a}$ is uniformly integrable as $a \rightarrow \infty$. To this end, for any $\eta > 0$, sufficiently large, we shall upper bound the following quantity using Lemma A.2 as follows.

$$\begin{aligned} \mathbb{E}_1\left[\frac{\tilde{N}^b}{\log a} \mathbb{I}_{\{\tilde{N}^b/\log a \geq \eta\}}\right] &\leq \mathbb{E}_1\left[\frac{(\tilde{N}^b - \lfloor \eta \log a \rfloor) + \eta \log a}{\log a} \mathbb{I}_{\{\tilde{N}^b \geq \lfloor \eta \log a \rfloor\}}\right] \\ &\leq \frac{1}{\log a} \mathbb{E}_1[(\tilde{N}^b - \lfloor \eta \log a \rfloor) \mathbb{I}_{\{\tilde{N}^b - \lfloor \eta \log a \rfloor \geq 0\}}] \end{aligned}$$

$$\begin{aligned}
 & + \frac{\eta \log a}{\log a} \mathbb{P}_1[\tilde{N}^b \geq \lfloor \eta \log a \rfloor] \\
 = & \frac{1}{\log a} \sum_{\ell=1}^{\infty} \mathbb{P}_1[\tilde{N}^b \geq \lfloor \eta \log a \rfloor + \ell] \\
 & + \eta \mathbb{P}_1[\tilde{N}^b \geq \lfloor \eta \log a \rfloor] \\
 \leq & \frac{a}{\log a} \sum_{\ell=1}^{\infty} e^{-(\eta \log a + \ell - 2)2B(\mu, \pi) + (\eta \log a + \ell)^{2/3}} (\lfloor \eta \log a \rfloor + \ell)^{2|\mathcal{Y}|} \\
 & + \eta a e^{-(\eta \log a - 2)2B(\mu, \pi) + (\eta \log a)^{2/3}} (\lfloor \eta \log a \rfloor)^{2|\mathcal{Y}|} \\
 \leq & \frac{a}{\log a} \sum_{\ell=1}^{\infty} e^{-(\eta \log a + \ell - 4)B(\mu, \pi)} (\lfloor \eta \log a \rfloor + \ell)^{2|\mathcal{Y}|} \\
 & + \eta a e^{-(\eta \log a - 4)B(\mu, \pi)} (\lfloor \eta \log a \rfloor)^{2|\mathcal{Y}|},
 \end{aligned} \tag{A.9}$$

for any $\eta > \frac{1}{B(\mu, \pi)}$ and a sufficiently large such that $(\eta \log a)B(\mu, \pi) \geq (\eta \log a)^{2/3}$.

Continuing from (A.9), upon noting that for a sufficiently large, it holds that $\lfloor \eta \log a \rfloor + \ell \leq 2\lfloor \eta \log a \rfloor \ell$, we get

$$\begin{aligned}
 \mathbb{E}_1 \left[\tilde{N}^b / \log a \mathbb{I}_{\{\tilde{N}^b / \log a \geq \eta\}} \right] & \leq \frac{a}{\log a} \sum_{\ell=1}^{\infty} e^{-(\eta \log a + \ell - 4)B(\mu, \pi)} (2\lfloor \eta \log a \rfloor \ell)^{2|\mathcal{Y}|} \\
 & + \eta a e^{-(\eta \log a - 4)B(\mu, \pi)} (\lfloor \eta \log a \rfloor)^{2|\mathcal{Y}|} \\
 & = \frac{a}{\log a} (2\lfloor \eta \log a \rfloor)^{2|\mathcal{Y}|} e^{-\eta B(\mu, \pi) \log a} \\
 & \quad \times \left(e^{4B(\mu, \pi)} \sum_{\ell=1}^{\infty} e^{-B(\mu, \pi) \ell} \ell^{2|\mathcal{Y}|} \right) \\
 & + \eta a (\lfloor \eta \log a \rfloor)^{2|\mathcal{Y}|} e^{-\eta B(\mu, \pi) \log a} \times e^{4B(\mu, \pi)},
 \end{aligned}$$

which vanishes as $a \rightarrow \infty$, for any $\eta > \frac{1}{B(\mu, \pi)}$, thereby establishing the uniform integrability.

A.2. Proof of Theorem 4.1

The proof of Theorem 4.1 relies on the following two lemmas.

Lemma A.3. *For any $\lambda > 2$, when the parameter b used in (4.15) is selected to be sufficiently close to 0, it holds for any true non-null hypothesis $i = 1, \dots, M$, and any $\varepsilon > 0$, that the first*

time T from which the estimate \hat{i} in (4.14) always equals the true hypothesis i , satisfies

$$\mathbb{P}_i[T > \varepsilon n] = O(n^{-\lambda}). \tag{A.10}$$

Proof. We first note that for any other hypothesis $j \neq i, j = 0, \dots, M$, and $n' \geq \lfloor \varepsilon n \rfloor$,

$$\mathbb{P}_i \left[\sum_{t=1}^{n'} \log \left(\frac{\bar{p}_i^{U_t}(Z_t)}{\bar{p}_j^{U_t}(Z_t)} \right) \leq 0 \right] \leq \mathbb{E}_i \left[e^{-1/2(\sum_{t=1}^{n'} \log(\bar{p}_i^{U_t}(Z_t)/\bar{p}_j^{U_t}(Z_t)))} \right], \tag{A.11}$$

where $\bar{p}_i^u, i = 0, \dots, M, u = 1, \dots, M$ are as defined in (4.13).

First, note that for $j = 1, \dots, M, j \neq i$, and for any time t when $U_t = s, s \neq i, s \neq j$, we get from (4.12), (4.13) that

$$\mathbb{E}_i \left[e^{-1/2(\log(\bar{p}_i^{U_t}(Z_t)/\bar{p}_j^{U_t}(Z_t)))} | U_t = s \right] = \mathbb{E}_i \left[e^{-1/2(\log 1)} | U_t = s \right] = 1 \quad \text{a.s.} \tag{A.12}$$

On the other hand, for the time t when $U_t = i$ or $U_t = j$, we get from (4.12), (4.13) that

$$\begin{aligned} & \mathbb{E}_i \left[e^{-1/2(\log(\bar{p}_i^{U_t}(Z_t)/\bar{p}_j^{U_t}(Z_t)))} | U_t = i \right] \\ &= \beta_a e^{-1/2 \log(1/(a-1))} + (1 - \beta_a) e^{-1/2 \log(a-1)} \\ &= \frac{\beta_a(a-1) + (1 - \beta_a)}{\sqrt{(a-1)}} \end{aligned} \tag{A.13}$$

$$\begin{aligned} &= \frac{\beta_a a + 1 - 2\beta_a}{\sqrt{(a-1)}} \leq \frac{2}{\sqrt{(a-1)}} < 1; \\ & \mathbb{E}_i \left[e^{-1/2(\log(\bar{p}_i^{U_t}(Z_t)/\bar{p}_j^{U_t}(Z_t)))} | U_t = j \right] \\ &= \alpha_a e^{-1/2 \log(1/(a-1))} + (1 - \alpha_a) e^{-1/2 \log(a-1)} \\ &\leq \frac{2}{\sqrt{(a-1)}} < 1, \end{aligned} \tag{A.14}$$

for $a > 5$, and where the inequalities in (A.13) and (A.14) follow from (4.11).

Similarly, we get from (4.12), (4.13) that for any time t when $U_t = s \neq i$,

$$\mathbb{E}_i \left[e^{-1/2(\log(\bar{p}_i^{U_t}(Z_t)/\bar{p}_0^{U_t}(Z_t)))} | U_t = s \right] = 1 \quad \text{a.s.} \tag{A.15}$$

On the other hand, for the time t when $U_t = i$, we get from (4.12), (4.13) that

$$\mathbb{E}_i \left[e^{-1/2(\log(\bar{p}_i^{U_t}(Z_t)/\bar{p}_0^{U_t}(Z_t)))} | U_t = i \right] \leq \frac{2}{\sqrt{(a-1)}} < 1, \tag{A.16}$$

for $a > 5$.

Consequently, we get from (4.15), (A.11), (A.13), (A.14), and (A.16) by successive uses of the smoothing property of conditional expectation that

$$\begin{aligned} \mathbb{P}_i[T > \varepsilon n] &\leq \sum_{n'=\lfloor \varepsilon n \rfloor}^{\infty} \mathbb{P}_i[\bar{p}_i(Z^{n'}) \leq \bar{p}_j(Z^{n'})] + \mathbb{P}_i[\bar{p}_i(Z^{n'}) \leq \bar{p}_0(Z^{n'})] \\ &\leq \sum_{n'=\lfloor \varepsilon n \rfloor}^{\infty} \left(\frac{2}{\sqrt{a-1}} \right)^{(2 \log n')/(Mb)} + \left(\frac{2}{\sqrt{a-1}} \right)^{(\log n')/(Mb)} \\ &= O(n^{-\lambda}), \end{aligned}$$

for $a > 5$ and for b used in (4.15) chosen sufficiently close to 0. □

For $\nu < 1, \rho_1 > 1$ as in Lemma 4.1 and a sufficiently large, let

$$\bar{c}(i, u) \triangleq \begin{cases} \frac{\log a}{\nu D(\mu \parallel \pi)}, & i = 1, \dots, M, u = i, \\ a(\log a)^{\rho_1}, & i = 1, \dots, M, u \neq i \end{cases} \tag{A.17}$$

that are the upper bounds for c_a, κ_a , as in Lemma 4.1, respectively. Let us consider the “true” model for the sequential design of binary-output experiments specified as (4.8), (4.9), induced by using the sequential binary test (4.1), (4.2), (4.3) as the “inner” test at each location. Also consider the “mismatched” model (for a sufficiently large) as in (4.12), (4.13) satisfying (4.11) (cf. Lemma 4.1). In addition, for $i = 1, \dots, M$, let

$$\begin{aligned} \bar{d}_i^a &= \frac{\mu_b(0) \log(\bar{\mu}_b(0)/\bar{\pi}_b(0)) + \mu_b(1) \log(\bar{\mu}_b(1)/\bar{\pi}_b(1))}{(\log a)/(\nu D(\mu \parallel \pi))} \\ &= \frac{\beta_a \log(1/(a-1)) + (1-\beta_a) \log(a-1)}{(\log a)/(\nu D(\mu \parallel \pi))} \\ &\rightarrow \nu D(\mu \parallel \pi), \quad \text{as } a \rightarrow \infty. \end{aligned} \tag{A.18}$$

Then, we have the following lemma.

Lemma A.4. *When the causal control policy (4.14), (4.15), (4.16), (4.17) is applied perpetually, it holds for any true non-null hypothesis $i \in [M]$, any other hypothesis $j = 0, 1, \dots, M, j \neq i$, any small $\varepsilon' > 0$, any $\lambda > 2$, and for all n sufficiently large (for $b > 0$ in (4.15) chosen sufficiently small) that*

$$\mathbb{P}_i \left[\log \left(\frac{\bar{p}_i(Z^n)}{\bar{p}_j(Z^n)} \right) < \left(\sum_{t=1}^n \bar{c}(i, U_t) \right) (\bar{d}_i^a - \varepsilon') \right] = O(n^{-\lambda}). \tag{A.19}$$

Proof. We first note that with $\underline{c} = \frac{\log a}{vD(\mu||\pi)}$, and, hence, $\bar{c}(i, u) \geq \underline{c}$, we get that

$$\begin{aligned} & \mathbb{P}_i \left[\log \left(\frac{\bar{p}_i(Z^n)}{\bar{p}_j(Z^n)} \right) < \left(\sum_{t=1}^n \bar{c}(i, U_t) \right) (\bar{d}_i^a - \varepsilon') \right] \\ & \leq \mathbb{P}_i \left[\sum_{t=1}^n \left(\log \left(\frac{\bar{p}_i^{U_t}(Z_t)}{\bar{p}_j^{U_t}(Z_t)} \right) - \mathbb{E}_i \left[\log \left(\frac{\bar{p}_i^{U_t}(Z_t)}{\bar{p}_j^{U_t}(Z_t)} \right) \middle| U_t \right] \right) < -n\underline{c} \frac{\varepsilon'}{2} \right] \tag{A.20} \\ & \quad + \mathbb{P}_i \left[\sum_{t=1}^n \left(\mathbb{E}_i \left[\log \left(\frac{\bar{p}_i^{U_t}(Z_t)}{\bar{p}_j^{U_t}(Z_t)} \right) \middle| U_t \right] - \bar{d}_i^a \bar{c}(i, U_t) \right) < -n\underline{c} \frac{\varepsilon'}{2} \right]. \end{aligned}$$

The proof that the probability of the first term on the right-hand side of (A.20) goes to zero exponentially fast in n follows from observing that the sequence

$$M_n = \sum_{t=1}^n \left(\log \left(\frac{\bar{p}_i^{U_t}(Z_t)}{\bar{p}_j^{U_t}(Z_t)} \right) - \mathbb{E}_i \left[\log \left(\frac{\bar{p}_i^{U_t}(Z_t)}{\bar{p}_j^{U_t}(Z_t)} \right) \middle| U_t \right] \right)$$

is a martingale and can be carried out by invoking the Chernoff bounding argument similar to the argument leading to equation (5.10) in [7]. In addition, we note that for a sufficiently large,

$$\begin{aligned} \min_{j \neq i} \min_{k=1, \dots, M} \mathbb{E}_i \left[\log \left(\frac{\bar{p}_i^{U_t}(Z_t)}{\bar{p}_j^{U_t}(Z_t)} \right) \middle| U_t = k \right] & \geq 0; \\ \max_{u=1, \dots, M} \bar{c}(i, u) & \leq a(\log a)^{\rho_1}. \end{aligned} \tag{A.21}$$

Next, for the T in Lemma A.3, we get from (4.15), (A.18), (A.17) that for all $t \geq T, t \neq \lfloor e^{b\ell} \rfloor, \ell = 0, 1, \dots$, that

$$\begin{aligned} \mathbb{E}_i \left[\log \left(\frac{\bar{p}_i^{U_t}(Z_t)}{\bar{p}_j^{U_t}(Z_t)} \right) \middle| U_t \right] & = \mathbb{E}_i \left[\log \left(\frac{\bar{p}_i^{U_t}(Z_t)}{\bar{p}_j^{U_t}(Z_t)} \right) \middle| U_t = i \right] \\ & = \bar{d}_i^a \left(\frac{\log a}{vD(\mu||\pi)} \right) \\ & = \bar{d}_i^a \bar{c}(i, U_t). \end{aligned} \tag{A.22}$$

Consequently, from (A.21) and by selecting ε in Lemma A.3 sufficiently small, we have that the second term on the right-hand side of (A.20) can be upper bounded according to Lemma A.3 as

$$\begin{aligned} & \mathbb{P}_i \left[\sum_{t=1}^n \left(\mathbb{E}_i \left[\log \left(\frac{\bar{p}_i^{U_t}(Z_t)}{\bar{p}_j^{U_t}(Z_t)} \right) \middle| U_t \right] - \bar{d}_i^a \bar{c}(i, U_t) \right) < -n\underline{c} \frac{\varepsilon'}{2} \right] \\ & \leq \mathbb{P}_i \left[\left(T + \frac{\log n}{b} \right) a(\log a)^{\rho_1} > n\underline{c} \frac{\varepsilon'}{2} \right] \\ & \leq \mathbb{P}_i \left[T > \frac{\varepsilon}{2} n \right] = O(n^{-\lambda}), \end{aligned} \tag{A.23}$$

thereby completing the proof of Lemma A.4. \square

To prove Theorem 4.1, we first shall prove that the stopping and final decision rules in (4.18) yield that

$$P_{\max} \leq \frac{M}{e^{a^{\rho_2}(\log a)^{\rho_1}}}. \quad (\text{A.24})$$

To this end, we consider two separate cases: when the true hypothesis is a non-null hypothesis, and when the true hypothesis is the null hypothesis.

First consider the case when the true hypothesis is a non-null hypothesis, say $i \in [M]$. For any $t \geq 1$ and a realization z^t of the binary search results of the search policy in (4.15), (4.16), (4.17), as in the paragraph preceding (4.14), we let $N(i, 1)$, $N(i, 0)$, $i \in [M]$, denote the number of times (up to time t) that the i th location is searched and the sequential binary test in (4.1), (4.2), (4.3) decides that the target is there, and that the target is not there, respectively. Then, we get from (4.12) that for any other non-null hypothesis $j \in [M]$, $j \neq i$, it holds that

$$\frac{\bar{p}_j(z^t)}{\bar{p}_i(z^t)} = \left(\frac{1/a}{1-1/a} \right)^{N(j,0)} \left(\frac{1-1/a}{1/a} \right)^{N(j,1)} \left(\frac{1-1/a}{1/a} \right)^{N(i,0)} \left(\frac{1/a}{1-1/a} \right)^{N(i,1)}. \quad (\text{A.25})$$

Now, for each such $j \neq i$, consider a pair of distributions $\tilde{\mu}_b, \tilde{\pi}_b$ (which are functions of both i, j) defined according to

$$\begin{aligned} \tilde{\mu}_b(0) &= (1 - \alpha_a) \frac{1/a}{1-1/a} = (1 - \pi_b(1)) \frac{1/a}{1-1/a}; \\ \tilde{\pi}_b(1) &= (1 - \beta_a) \frac{1/a}{1-1/a} = (1 - \mu_b(0)) \frac{1/a}{1-1/a}. \end{aligned} \quad (\text{A.26})$$

From (4.11), and (A.26), we get by an easy calculation that

$$\begin{aligned} \frac{1 - \tilde{\mu}_b(0)}{\pi_b(1)} &\geq \frac{1-1/a}{1/a} \geq 1; \\ \frac{1 - \tilde{\pi}_b(1)}{\mu_b(0)} &\geq \frac{1-1/a}{1/a} \geq 1, \end{aligned} \quad (\text{A.27})$$

for a large. Now consider another probability distribution, $\tilde{p}_j(z^t)$, defined as a function of both j and i , based on the same search policy according to

$$\tilde{p}_j^u = \tilde{\mu}_b, \quad u = j, \quad \tilde{p}_j^u = \tilde{\pi}_b, \quad u = i, \quad \tilde{p}_j^u = \pi_b, \quad u \neq i, j. \quad (\text{A.28})$$

Then, it holds that

$$\begin{aligned} \frac{\tilde{p}_j(z^t)}{\tilde{p}_i(z^t)} &= \left(\frac{\tilde{\mu}_b(0)}{1 - \pi_b(1)} \right)^{N(j,0)} \left(\frac{1 - \tilde{\mu}_b(0)}{\pi_b(1)} \right)^{N(j,1)} \\ &\quad \times \left(\frac{1 - \tilde{\pi}_b(1)}{\mu_b(0)} \right)^{N(i,0)} \left(\frac{\tilde{\pi}_b(1)}{1 - \mu_b(0)} \right)^{N(i,1)}. \end{aligned} \quad (\text{A.29})$$

Consequently, we get from (A.25), (A.29), (A.26) and (A.27) that

$$\frac{\tilde{p}_j(z^t)}{p_i(z^t)} \geq \frac{\bar{p}_j(z^t)}{\bar{p}_i(z^t)}. \tag{A.30}$$

Similarly, by observing that

$$\frac{\bar{p}_0(z^t)}{\bar{p}_i(z^t)} = \left(\frac{1/a}{1-1/a}\right)^{N(i,1)} \left(\frac{1-1/a}{1/a}\right)^{N(i,0)}, \tag{A.31}$$

and define $\tilde{p}_0(z^t)$ according to

$$\tilde{p}_0^u = \tilde{\pi}_b, \quad u = i, \quad \tilde{p}_0^u = \pi_b, \quad u \neq i, \tag{A.32}$$

we get from (A.32), the second equality of (A.26), the second inequality of (A.27) and (A.31) that

$$\frac{\tilde{p}_0(z^t)}{p_i(z^t)} = \left(\frac{\tilde{\pi}_b(1)}{1-\mu_b(0)}\right)^{N(i,1)} \left(\frac{1-\tilde{\pi}_b(1)}{\mu_b(0)}\right)^{N(i,0)} \geq \frac{\bar{p}_0(z^t)}{\bar{p}_i(z^t)}. \tag{A.33}$$

Hence, under the non-null hypothesis i , the error probability incurred by the rules (4.18) can be upper bounded as

$$\begin{aligned} \mathbb{P}_i[\delta(Z^\tau) \neq i] &= \sum_{j=0}^M \sum_{\substack{t=1 \\ j \neq i}}^{\infty} \mathbb{P}_i[\delta(Z^t) = j, \tau = t] \\ &\leq \sum_{j=0}^M \sum_{\substack{t=1 \\ j \neq i}}^{\infty} \mathbb{P}_i\left[\frac{\tilde{p}_j(Z^t)}{p_i(Z^t)} > e^{a^{\rho_2}(\log a)^{\rho_1}}, \tau = t\right] \\ &< \sum_{\substack{j=0 \\ j \neq i}}^M \frac{\sum_{t=1}^{\infty} \tilde{\mathbb{P}}_j[\tau = t]}{e^{a^{\rho_2}(\log a)^{\rho_1}}} \\ &\leq \frac{M}{e^{a^{\rho_2}(\log a)^{\rho_1}}}, \end{aligned} \tag{A.34}$$

where the first inequality above follows from the stopping rule in (4.18), (A.30) and (A.33), and the second inequality follows from a change of measure argument.

The error probability under the null hypothesis can be analyzed in a similar manner. In particular, for any non-null hypothesis $j \in [M]$, we have that

$$\frac{\bar{p}_j(z^t)}{\bar{p}_0(z^t)} = \left(\frac{1/a}{1-1/a}\right)^{N(j,0)} \left(\frac{1-1/a}{1/a}\right)^{N(j,1)}. \tag{A.35}$$

By defining $\tilde{p}_j(z^t)$ according to

$$\tilde{p}_j^u = \tilde{\mu}_b, \quad u = j, \quad \tilde{p}_j^u = \pi_b, \quad u \neq j, \quad (\text{A.36})$$

we get from the first equality of (A.26), the first inequality of (A.27), and (A.35) that

$$\frac{\tilde{p}_j(z^t)}{p_0(z^t)} = \left(\frac{\tilde{\mu}_b(0)}{1 - \pi_b(1)} \right)^{N(j,0)} \left(\frac{1 - \tilde{\mu}_b(0)}{\pi_b(1)} \right)^{N(j,1)} \geq \frac{\bar{p}_j(z^t)}{\bar{p}_0(z^t)}. \quad (\text{A.37})$$

Using (A.37) and the arguments similar to the one leading to (A.34), we get that

$$\mathbb{P}_0[\delta(Z^\tau) \neq 0] \leq \frac{M}{e^{a^{\rho_2}(\log a)^{\rho_1}}}, \quad (\text{A.38})$$

thereby, together with (A.34), yielding (A.24).

Next, for a non-null hypothesis $i = 1, \dots, M$, and any other hypothesis $j = 0, 1, \dots, M, j \neq i$, let τ_j denote the smallest time for which $\log\left(\frac{\bar{p}_i(Z^t)}{\bar{p}_j(Z^t)}\right) > a^{\rho_2}(\log a)^{\rho_1}$ for all $t \geq \tau_j$. Now for any $\delta > 0$ and $A > \frac{a^{\rho_2}(\log a)^{\rho_1}}{\bar{d}_i^a - \delta}$, where \bar{d}_i^a is as in (A.18) and with $\tilde{c} = a(\log a)^{\rho_1}$, it follows that for a large,

$$\begin{aligned} \mathbb{P}_i \left[\sum_{t=1}^{\tau_j} \bar{c}(i, U_t) > A \right] &\leq \mathbb{P}_i \left[\sum_{t=1}^{\tau_j-1} \bar{c}(i, U_t) > A - \tilde{c}; (\tau_j - 1) \geq \left\lfloor \frac{A}{2\tilde{c}} \right\rfloor \right] \\ &\leq \mathbb{P}_i \left[\sum_{t=1}^{\tau_j-1} \bar{c}(i, U_t) > \frac{a^{\rho_2}(\log a)^{\rho_1}}{(\bar{d}_i^a - \delta)} - \tilde{c}; (\tau_j - 1) \geq \left\lfloor \frac{A}{2\tilde{c}} \right\rfloor \right] \\ &\leq \sum_{n=\lfloor A/(2\tilde{c}) \rfloor}^{\infty} \mathbb{P}_i \left[\sum_{t=1}^n \bar{c}(i, U_t) > \frac{\log(\bar{p}_i(Z^n)/\bar{p}_j(Z^n))}{(\bar{d}_i^a - \delta/2)} \right] \end{aligned} \quad (\text{A.39})$$

$$\leq \sum_{n=\lfloor A/(2\tilde{c}) \rfloor}^{\infty} O(n^{-\lambda}) = O\left(\left(\frac{A}{2\tilde{c}}\right)^{-\lambda+1}\right), \quad (\text{A.40})$$

where (A.39) follows from the fact that for a large, $\frac{\tilde{c}}{a^{\rho_2}(\log a)^{\rho_1}} \rightarrow 0$, and (A.40) follows from Lemma A.4. Consequently, we get from (A.40) that for any $j \neq i$,

$$\begin{aligned} \mathbb{E}_i \left[\sum_{t=1}^{\tau_j} \bar{c}(i, U_t) \right] &\leq \frac{a^{\rho_2}(\log a)^{\rho_1}}{\bar{d}_i^a - \delta} \left(1 + \int_{\frac{a^{\rho_2}(\log a)^{\rho_1}}{\bar{d}_i^a - \delta}}^{\infty} O\left(\left(\frac{A}{2\tilde{c}}\right)^{-\lambda+1}\right) dA \right) \\ &\leq \frac{a^{\rho_2}(\log a)^{\rho_1}}{\bar{d}_i^a - \delta} \left(1 + 2\tilde{c} O\left(\left(\frac{a^{\rho_2}(\log a)^{\rho_1}}{\bar{d}_i^a - \delta}\right)^{-\lambda+2}\right) \right) \end{aligned} \quad (\text{A.41})$$

$$\begin{aligned} &\leq \frac{a^{\rho_2}(\log a)^{\rho_1}}{\bar{d}_i^a - \delta} \left(1 + 2a(\log a)^{\rho_1} O\left(\left(\frac{a^{\rho_2-1}}{\bar{d}_i^a - \delta}\right)^{-\lambda+2}\right) \right) \\ &= \frac{a^{\rho_2}(\log a)^{\rho_1}}{\bar{d}_i^a - \delta} (1 + o(1)), \end{aligned}$$

for λ sufficiently large so that $(\lambda - 2)(\rho_2 - 1) > 1$.

Last, it follows from (4.18) that $\tau \leq \max_{j \neq i} \tau_j$. Consequently, we get from (A.41), (A.18) by virtue of the fact that $\bar{c}(i, u) \geq c(i, u)$ (cf. Lemma 4.1 and (A.17)), that

$$\begin{aligned} \mathbb{E}_i \left[\sum_{t=1}^{\tau} c(i, U_t) \right] &\leq \frac{a^{\rho_2}(\log a)^{\rho_1}}{\bar{d}_i^a - \delta} (1 + o(1)) \\ &= \frac{a^{\rho_2}(\log a)^{\rho_1}}{\nu D(\mu \|\pi) - \delta} (1 + o(1)), \end{aligned}$$

thereby, together with (A.24), yielding (4.21) and, hence, completing the proof, as δ and ν can be arbitrarily close to 0 and 1, respectively.

Acknowledgements

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, by the U.S. Defense Threat Reduction Agency through subcontract 147755 at the University of Illinois from prime award HDTRA1-10-1-0086, and by the National Science Foundation under Grant NSF CCF 11-11342.

References

- [1] Ahlswede, R. and Wegener, I. (1987). *Search Problems. Wiley-Interscience Series in Discrete Mathematics and Optimization*. Chichester: Wiley. [MR0929026](#)
- [2] Alpern, S. and Gal, S. (2003). *The Theory of Search Games and Rendezvous. International Series in Operations Research & Management Science*, 55. Boston, MA: Kluwer Academic. [MR2005053](#)
- [3] Benkoski, S.J., Monticino, M.G. and Weisinger, J.R. (1991). A survey of the search theory literature. *Naval Res. Logist.* **38** 469–494.
- [4] Bessler, S.A. (1960). Theory and applications of the sequential design of experiments, k -actions and infinitely many experiments: Part I – Theory. Tech. Report **55**. Dept. Statistics, Stanford Univ.
- [5] Castro, R.M. (2014). Adaptive sensing performance lower bounds for sparse signal detection and support estimation. *Bernoulli* **20** 2217–2246. [MR3263103](#)
- [6] Castro, R.M. and Tánzos, E. (2015). Adaptive sensing for estimation of structured sparse signals. *IEEE Trans. Inform. Theory* **61** 2060–2080. [MR3332997](#)
- [7] Chernoff, H. (1959). Sequential design of experiments. *Ann. Math. Statist.* **30** 755–770. [MR0108874](#)
- [8] Chernoff, H. (1961). Sequential tests for the mean of a normal distribution. In *Proc. 4th Berkeley Sympos. Math. Statist. and Prob., Vol. 1* 79–91. Berkeley, CA: Univ. California Press. [MR0131941](#)

- [9] Chernoff, H. (1965). Sequential tests for the mean of a normal distribution. III. (Small t). *Ann. Math. Statist.* **36** 28–54. [MR0170442](#)
- [10] Chudnovsky, D.V. and Chudnovsky, G.V., eds. (1989). *Search Theory: Some Recent Developments. Lecture Notes in Pure and Applied Mathematics* **112**. New York: Dekker. [MR0977561](#)
- [11] Cover, T.M. and Thomas, J.A. (2006). *Elements of Information Theory*, 2nd ed. Hoboken, NJ: Wiley. [MR2239987](#)
- [12] Haupt, J., Castro, R.M. and Nowak, R. (2011). Distilled sensing: Adaptive sampling for sparse detection and estimation. *IEEE Trans. Inform. Theory* **57** 6222–6235. [MR2857969](#)
- [13] Koopman, B.O. (1980). *Search and Screening. General Principles with Historical Applications*, 2nd ed. Elmsford, NY: Pergamon Press. [MR0642358](#)
- [14] Lai, T.L. (1988). Nearly optimal sequential tests of composite hypotheses. *Ann. Statist.* **16** 856–886. [MR0947582](#)
- [15] Li, Y., Nitinawarat, S. and Veeravalli, V.V. (2014). Universal outlier hypothesis testing. *IEEE Trans. Inform. Theory* **60** 4066–4082. [MR3225950](#)
- [16] Li, Y., Nitinawarat, S. and Veeravalli, V.V. (2014). Universal sequential outlier hypothesis testing. In *Proceedings of the International Symposium on Information Theory* 3205–3209. Honolulu, HI, USA: IEEE.
- [17] Naghshvar, M. and Javidi, T. (2013). Active sequential hypothesis testing. *Ann. Statist.* **41** 2703–2738. [MR3161445](#)
- [18] Nitinawarat, S., Atia, G.K. and Veeravalli, V.V. (2013). Controlled sensing for multihypothesis testing. *IEEE Trans. Automat. Control* **58** 2451–2464. [MR3106054](#)
- [19] Nitinawarat, S. and Veeravalli, V.V. (2015). Controlled sensing for sequential multihypothesis testing with controlled Markovian observations and non-uniform control cost. *Sequential Anal.* **34** 1–24. [MR3306137](#)
- [20] Schwarz, G. (1962). Asymptotic shapes of Bayes sequential testing regions. *Ann. Math. Statist.* **33** 224–236. [MR0137226](#)
- [21] Stone, L.D. (2007). *Theory of Optimal Search*, 2nd ed. *Topics in Operations Research Series*. Hanover, MD: INFORMS.

Received July 2015 and revised November 2015