

TWENTY QUESTIONS WITH NOISE: BAYES OPTIMAL POLICIES FOR ENTROPY LOSS

BRUNO JEDYNAK,* *Johns Hopkins University*

PETER I. FRAZIER,** *Cornell University*

RAPHAEL SZNITMAN,*** *Johns Hopkins University*

Abstract

We consider the problem of twenty questions with noisy answers, in which we seek to find a target by repeatedly choosing a set, asking an oracle whether the target lies in this set, and obtaining an answer corrupted by noise. Starting with a prior distribution on the target's location, we seek to minimize the expected entropy of the posterior distribution. We formulate this problem as a dynamic program and show that any policy optimizing the one-step expected reduction in entropy is also optimal over the full horizon. Two such Bayes optimal policies are presented: one generalizes the probabilistic bisection policy due to Horstein and the other asks a deterministic set of questions. We study the structural properties of the latter, and illustrate its use in a computer vision application.

Keywords: Twenty questions; dynamic programming; bisection; search; object detection; entropy loss; sequential experimental design; Bayesian experimental design

2010 Mathematics Subject Classification: Primary 60J20

Secondary 62C10; 90B40; 90C39

1. Introduction

In this paper we consider the problem of finding a target $X^* \in \mathbb{R}^d$ by asking a knowledgeable oracle questions. Each question consists in choosing a set $A \subseteq \mathbb{R}^d$, querying the oracle whether X^* lies in this set, and observing the associated response. While this is closely related to the popular game of ‘twenty questions’, we consider here the case where answers from the oracle are corrupted with noise from a known model. This game appears naturally in a number of problems in stochastic search, stochastic optimization, and stochastic root finding. In this paper we present an illustrative application in computer vision.

We consider a Bayesian formulation of this problem using entropy loss. In dimension $d = 1$ we seek to minimize the expected entropy of the posterior after a fixed number of questions. After formulating the problem in Section 2, we show in Section 3 that any policy myopically maximizing the expected one-step reduction in entropy is also optimal in a fully sequential sense (Theorems 1 and 2), and to follow such a policy it is sufficient to query sets A whose posterior probability of containing X^* is a specific value given in Theorem 2. We then

Received 1 March 2011; revision received 16 June 2011.

* Postal address: Department of Applied Mathematics and Statistics, Johns Hopkins University, Whitehead 208-B, 3400 North Charles Street, Baltimore, MD 21218, USA. Email address: bruno.jedynak@jhu.edu

** Postal address: School of Operations Research and Industrial Engineering, Cornell University, 232 Rhodes Hall, Ithaca, NY 14850, USA.

Research supported in part by AFOSR YIP FA9550-11-1-0083.

*** Postal address: Johns Hopkins University, Hackerman Hall, 3400 North Charles Street, Baltimore, MD 21218, USA.

Research supported in part by NIH grant R01 EB 007969-01.

provide two specific Bayes optimal policies. The first, described in Section 4.1, poses questions about intervals, $A = (-\infty, x)$. The second, which we call the dyadic policy and describe in Section 4.2, poses questions about more general sets. We also provide further analysis of this second policy: a law of large numbers and a central limit theorem for the posterior entropy (Theorem 3), and an explicit characterization of the expected number of size-limited noise-free questions required to find the target after noisy questioning ceases (Theorem 4). In Section 5 we consider a modified version of the entropic loss in $d = 2$ dimensions, and show that a simple modification of the dyadic policy is asymptotically Bayes optimal for this loss function (Theorem 5). In Section 5 we also provide a central limit theorem for the posterior entropy under this policy (Theorem 6). In Section 6 we provide an illustrative application in computer vision. Concluding remarks are given in Section 7.

When the noise corrupting the oracle's responses is of a special form, that of a symmetric channel, the Bayes optimal policy for $d = 1$ with questions A restricted to be intervals (described in Section 4.1) takes a particularly natural form: choose $A = (-\infty, x)$, where x is the median of the posterior distribution. This policy, called the probabilistic bisection strategy, was first proposed in [12] (republished in [13]). This policy was recently shown to be optimal in the binary symmetric case by one of the authors in [31]. Burnašhev and Zigangirov [4] introduced a similar procedure that measures on either side of the median of the posterior over a discrete set of points, and showed that its error probability decays at an asymptotically optimal rate. For a review of these two procedures, see [5]. Both Karp and Kleinberg [14] and Ben-Or and Hassidim [1] also considered a noisy binary search problem with constant error probability over a discrete set of points, and gave optimality results for policies similar to measuring at the median of the posterior. In [14], this is part of a larger analysis in which the error probability may vary. Nowak [19], [20] analyzed noise-tolerant versions of generalized binary search for searching in a space of hypotheses. A parallel line of research has considered the case when the oracle is adversarial rather than stochastic, and is surveyed in [21].

When the questions are restricted to be intervals, the problem that we consider is similar to the stochastic root-finding problem considered in the seminal paper [24] and generalized to multiple dimensions in [3]. In the stochastic root-finding problem, one chooses a sequence of points x_1, x_2, \dots to query, and observes the corresponding values $f(x_1), f(x_2), \dots$ of some decreasing function f at x , obscured by noise. The goal in this problem is to find the root of f . Procedures include the stochastic approximation methods of [3] and [24], as well as the Polyak–Ruppert averaging introduced independently in [22] and [25]. Asymptotic rates of convergence of these procedures are well understood; see [15, Chapter 10]. Our problem and the stochastic root-finding problem are similar because if X^* is the root of f then querying whether X^* is in $(-\infty, x)$ can be recast as querying whether $f(x) < 0$. The problems differ because the noise in observing whether $f(x) < 0$ depends upon x and is generally larger when $f(x)$ is closer to 0, while in our formulation we assume that the distribution of the oracle's response depends only on whether X^* is in the queried subset or not.

Both our problem and stochastic root-finding lie within the larger class of problems in sequential experimental design, in which we choose at each point in time which experiment to perform in order to optimize some overall value of the information obtained. The study of this area began with Robbins [23], who introduced the multi-armed bandit problem, later studied in [2], [11], [16], [32], and [33], among others. For a self-contained discussion of sequential experimental design in a Bayesian context, see [7].

2. Formulation of the problem

Nature chooses a continuous random variable X^* with density p_0 with respect to the Lebesgue measure over \mathbb{R}^d . The fact that X^* is continuous will turn out to be important and the arguments presented below do not generalize easily to the case where X^* is a discrete random variable.

To discover X^* , we can sequentially ask N questions. Asking the n th question, $0 \leq n \leq N - 1$, involves choosing a Lebesgue measurable set $A_n \subset \mathbb{R}^d$ and evaluating: Does X^* belong to A_n ? To avoid technical issues below, we require that A_n is the union of at most J_n half-open intervals, where J_0, J_1, \dots is a fixed sequence of natural numbers. The answer, denoted Z_n , is the indicator function of the event $\{X^* \in A_n\}$. However, Z_n is not openly communicated to us. Instead, Z_n is the input of a memoryless noisy transmission channel from which we observe the output Y_{n+1} . Here Y_{n+1} is a random variable which can be discrete or continuous, univariate or multivariate. The memoryless property of the channel expresses the fact that Y_{n+1} depends on Z_n , but not on previous questions or answers. As a consequence, repeatedly answering the same question may not provide the same answer each time. Moreover, we assume that the distribution of Y_{n+1} given Z_n does not depend on n . There is a measure μ on the space in which Y_{n+1} takes value, and the density with respect to μ of Y_{n+1} given Z_n is

$$\frac{P(Y_{n+1} \in dy \mid Z_n = z)}{d\mu} = \begin{cases} f_1(y) & \text{if } z = 1, \\ f_0(y) & \text{if } z = 0. \end{cases} \tag{1}$$

If Y_{n+1} is discrete then we take μ to be a discrete measure, while if Y_{n+1} is continuous we take μ to be the Lebesgue measure. We require that the Shannon entropy of the conditional distribution $P(Y_{n+1} \in \cdot \mid Z_n = z)$ be finite for both $z = 0$ and $z = 1$. At any time step n , we characterize what we know about X^* by computing the conditional density p_n of X^* given the history of previous measurements $D_n = (A_m, Y_{m+1})_{m=0}^{n-1}$. Following the terminology of Bayesian statistics, we call p_n the posterior density. The study of the stochastic sequences of densities p_n , under different policies, constitutes the main mathematical contribution of this paper. For an event A , we will use the notation

$$p_n(A) = \int_A p_n(x) dx.$$

The posterior density p_{n+1} of X^* after observing D_{n+1} is elegantly described as a function of p_n, f_0, f_1 , the n th question A_n , and the answer to this question Y_{n+1} .

Lemma 1. *On the event $A_n = A$ and $Y_{n+1} = y$, the posterior density on X^* is*

$$p_{n+1}(u) = \frac{1}{\mathcal{Z}} (f_1(y) \mathbf{1}_{\{u \in A\}} + f_0(y) \mathbf{1}_{\{u \notin A\}}) p_n(u),$$

where

$$\mathcal{Z} = f_1(y) p_n(A) + f_0(y) (1 - p_n(A)). \tag{2}$$

Proof. On the event $A_n = A$ and $Y_{n+1} = y$, the posterior density $p_{n+1}(u) = P(X^* \in du \mid D_{n+1})/d\lambda = P(X^* \in du \mid D_n, A_n = A, Y_{n+1} = y)/d\lambda$, where λ is the Lebesgue measure, can be written using Bayes' formula as

$$\begin{aligned} & \frac{1}{\mathcal{Z}} \frac{P(Y_{n+1} \in dy \mid D_n, A_n = A, X^* = u) P(X^* \in du \mid D_n, A_n = A)}{d\mu} \frac{1}{d\lambda} \\ &= \frac{1}{\mathcal{Z}} (f_1(y) \mathbf{1}_{\{u \in A\}} + f_0(y) \mathbf{1}_{\{u \notin A\}}) p_n(u), \end{aligned}$$

where \mathcal{Z} is the normalizing constant,

$$\mathcal{Z} = \int_u (f_1(y) \mathbf{1}_{\{u \in A\}} + f_0(y) \mathbf{1}_{\{u \notin A\}}) p_n(u) du.$$

Later, we will take conditional expectations given the density p_n . Formally, these conditional expectations are taken with respect to the sigma-algebra generated by the stochastic process $\{p_n(u) : u \in I\}$. Because $p_n(u)$ for each u is a function of D_n by the recursive expression in Lemma 1, this sigma-algebra is a subset of the sigma-algebra generated by D_n .

We will measure the quality of the information gained about X^* from these N questions using the Shannon differential entropy. The Shannon differential entropy (see [6, Chapter 9]), or simply ‘the entropy’ of p_n , $H(p_n)$, is defined as

$$H(p_n) = - \int_{-\infty}^{+\infty} p_n(x) \log p_n(x) dx,$$

where \log is the logarithm to base 2. In particular, we consider the problem of finding a sequence of N questions such that the expected entropy of X^* after observing the N th answer is minimized.

We will write this problem more formally as the infimum over policies of the expectation of the posterior entropy, but before doing so we must formally define a policy. Informally, a policy is a method for choosing the questions A_n as a function of the observations available at time n . The technical assumption that each question A_n is a union of only finitely many intervals ensures the Borel measurability of $H(p_N)$ under each policy.

First, A_n is the union of at most J_n half-open intervals, and so may be written as

$$A_n = \bigcup_{j=1}^{J_n} [a_{n,j}, b_{n,j}),$$

where $a_{n,j} \leq b_{n,j}$ are elements of $\overline{\mathbb{R}} = \mathbb{R} \cup \{-\infty, +\infty\}$. If $a_{n,j} = -\infty$ then the corresponding interval is understood to be open on the left. Here J_0, J_1, \dots, J_{N-1} is any fixed sequence of natural numbers that is the same for all policies. If A_n comprises strictly less than J_n intervals then we may take $a_{n,j} = b_{n,j}$ for some j . When A_n is written in this way, the space in which A_n takes values may be identified with the space $\mathbb{A}_n = \{(a_j, b_j) : j = 1, \dots, J_n, a_j \leq b_j\}$, which is a closed subset of $\overline{\mathbb{R}}^{2J_n}$.

Then, with fixed p_0 , p_n may be identified with the sequence $((a_{m,j}, b_{m,j})_{j=1}^{J_m}, Y_{m+1})_{m=0}^{n-1}$, which takes values in the space $\mathbb{S}_n = (\mathbb{A}_0 \times \dots \times \mathbb{A}_{n-1}) \times \mathbb{R}^n$. Furthermore, the function $p_n \mapsto H(p_n)$ may be written as a measurable function from \mathbb{S}_n to \mathbb{R} .

After having identified possible values for A_n with points in \mathbb{A}_n and possible values for p_n with points in \mathbb{S}_n , we define a policy π to be a sequence of functions $\pi = (\pi_0, \pi_1, \dots)$, where $\pi_n : \mathbb{S}_n \mapsto \mathbb{A}_n$ is a measurable function. We let Π be the space of all such policies. Any such policy π induces a probability measure on $((a_{n,j}, b_{n,j})_{j=1}^{J_n}, Y_{n+1})_{n=0}^{N-1}$. We let E^π indicate the expectation with respect to this probability measure. In a slight abuse of notation, we will sometimes talk of $p \in \mathbb{S}_n$ and $A \in \mathbb{A}_n$, by which we mean the density p associated with a vector in \mathbb{S}_n , or the set A associated with a vector in \mathbb{A}_n .

With this definition of a policy π , the associated measure E^π , and the space of all policies Π , the problem under consideration may be written as

$$\inf_{\pi \in \Pi} E^\pi [H(p_N)]. \tag{3}$$

Any policy attaining the infimum is called optimal. We consider this problem for the general case in Section 3, and for the specific cases of $d = 1$ and $d = 2$ in Sections 4 and 5, respectively. In Section 5 we also consider a modification of this objective function that separately considers the entropy of the marginal posterior distribution, and ensures that both entropies are small. This prevents a policy from obtaining optimality by learning one coordinate of X^* without learning the other.

3. Entropy loss and channel capacity

In this section we consider the problem (3) of minimizing the expected entropy of the posterior over \mathbb{R}^d . We present general results characterizing optimal policies, which will be used to create specific policies in Sections 4 and 5.

We first present some notation that will be used within our results. Let φ be the function with domain $[0, 1]$ defined by

$$\varphi(u) = H(uf_1 + (1 - u)f_0) - uH(f_1) - (1 - u)H(f_0);$$

$\varphi(u)$ is a mutual information for each u (see (7) and (10) below). The associated channel capacity C is

$$C = \sup_{u \in [0, 1]} \varphi(u).$$

Below, in Theorem 1, we show that this maximum is attained in $(0, 1)$. Let $u^* \in (0, 1)$ be a point attaining this maximum, so $\varphi(u^*) = C$.

We show that an optimal policy consists of choosing each A_n so that $p_n(A_n) = u^*$. When the A_n are chosen in this way, the expected entropy decreases arithmetically by the constant C at each step. Moreover, if the communication channel is symmetric in the sense that $\varphi(1 - u) = \varphi(u)$ for all $0 \leq u \leq 1$, then $u^* = \frac{1}{2}$. In the noiseless case, or even the case where the supports of f_0 and f_1 do not overlap, the model is symmetric, $C = 1$, and the obvious bisection policy is optimal.

Optimal policies constructed by choosing $p_n(A_n) = u^*$ are greedy policies (or ‘knowledge-gradient’ policies, as defined in [9]), since they make decisions that would be optimal if only one measurement remained, i.e. if N were equal to $n + 1$. Such greedy policies are usually used only as heuristics, and so it is interesting that they are optimal in this problem.

Our analysis relies on dynamic programming. To support this analysis, we define the value function

$$V(p, n) = \inf_{\pi \in \Pi} E^\pi[H(p_N) \mid p_n = p], \quad p \in \mathbb{S}_n, n = 0, \dots, N.$$

Standard results from controlled Markov processes show that this value function satisfies Bellman’s recursion (see Section 3.7 of [8]),

$$V(p, n) = \inf_{A \in \mathbb{A}_n} E[V(p_{n+1}, n + 1) \mid A_n = A, p_n = p], \quad p \in \mathbb{S}_n, n < N, \quad (4)$$

where the expectation is taken over Y_{n+1} , and any policy attaining the minimum of (4) is optimal (see Section 2.3 of [8]). In general, the results of [8] for general Borel models imply only that $V(\cdot, n): \mathbb{S}_n \mapsto \mathbb{R}$ is universally measurable, and do not imply Borel measurability. However, we show below in Theorem 2 that, in our case, $V(\cdot, n): \mathbb{S}_n \mapsto \mathbb{R}$ is a Borel-measurable function.

As a preliminary step toward solving Bellman’s recursion, we present the following theorem, which shows that minimizing the expected entropy of the posterior one step into the future can be accomplished by choosing A_n as described above. Furthermore, it shows that the expected reduction in entropy is the channel capacity C .

Theorem 1. *We have*

$$\inf_{A \in \mathbb{A}_n} E[H(p_{n+1}) \mid A_n = A, p_n] = H(p_n) - C, \tag{5}$$

where the expectation is taken over Y_{n+1} . Moreover, there exists a point $u^* \in (0, 1)$ such that $\varphi(u^*) = C$, and the minimum in (5) is attained by choosing A such that $p_n(A) = u^*$.

Proof. We first rewrite the expected entropy as

$$E[H(p_{n+1}) \mid A_n = A, p_n] = H(p_n) - I(X^*, Y_{n+1} \mid A_n = A, p_n),$$

where $I(X^*, Y_{n+1} \mid A_n = A, p_n)$ is the mutual information between the conditional distributions of X^* and Y_{n+1} (see [6, Chapter 2]), and we recall that the entropy of X^* given $A_n = A$ and p_n is exactly $H(p_n)$. This leads to

$$\inf_{A \in \mathbb{A}_n} E[H(p_{n+1}) \mid A_n = A, p_n] = H(p_n) - \sup_{A \in \mathbb{A}_n} I(X^*, Y_{n+1} \mid A_n = A, p_n). \tag{6}$$

Temporarily fixing A , we expand the mutual information as

$$I(X^*, Y_{n+1} \mid A_n = A, p_n) = H(Y_{n+1} \mid A_n = A, p_n) - H(Y_{n+1} \mid X^*, A_n = A, p_n), \tag{7}$$

where $H(\cdot \mid \cdot)$ is the conditional entropy, as defined in [6, Chapter 2]. Using (2),

$$H(Y_{n+1} \mid A_n = A, p_n) = H(p_n(A)f_1 + (1 - p_n(A))f_0). \tag{8}$$

Also,

$$\begin{aligned} H(Y_{n+1} \mid X^*, A_n = A, p_n) &= \int_u p_n(u)H(Y_{n+1} \mid X^* = u, A_n = A, p_n) \, du \\ &= \int_{u \in A} p_n(u)H(f_1) \, du + \int_{u \notin A} p_n(u)H(f_0) \, du \\ &= H(f_1)p_n(A) + H(f_0)(1 - p_n(A)). \end{aligned} \tag{9}$$

The difference between (8) and (9) is $\varphi(p_n(A))$, and so

$$I(X^*, Y_{n+1} \mid A_n = A, p_n) = \varphi(p_n(A)). \tag{10}$$

This and (6) together show that

$$\sup_{A \in \mathbb{A}_n} I(X^*, Y_{n+1} \mid A_n = A, p_n) = \sup_{A \in \mathbb{A}_n} \varphi(p_n(A)) = \sup_{u \in [0,1]} \varphi(u) = C.$$

This shows (5), and that the infimum in (5) is attained by any set A with $\varphi(p_n(A)) = C$. It remains only to show the existence of a point $u^* \in (0, 1)$, with $\varphi(u^*) = C$.

First, φ is a continuous function, so its maximum over the compact interval $[0, 1]$ is attained. If the maximum is attained in $(0, 1)$ then we simply choose u^* to be this point. Now consider the case when the maximum is attained at $u \in \{0, 1\}$. Because φ is a mutual information for each u , it is nonnegative. Also, $\varphi(0) = \varphi(1) = 0$. Thus, if the maximum is attained at $u \in \{0, 1\}$ then $\varphi(u) = 0$ for all u , and we can choose u^* in the open interval $(0, 1)$.

We are now ready to present the main result of this section, which gives a simple characterization of optimal policies.

Theorem 2. Any policy that chooses each A_n to satisfy $p_n(A_n) = u^* \in \arg \max_{u \in [0,1]} \varphi(u)$ is optimal. In addition, for each n , the value function $V(\cdot, n): \mathbb{S}_n \mapsto \mathbb{R}$ is Borel measurable and is given by

$$V(p_n, n) = H(p_n) - (N - n)C. \tag{11}$$

Proof. It is enough to show that, for each $n = 0, 1, \dots, N$, the value function is given by (11), and that the described policy achieves the minimum in Bellman’s recursion (4). Measurability of $V(\cdot, n): \mathbb{S}_n \mapsto \mathbb{R}$ then follows from the fact that $p_n \mapsto H(p_n)$ is Borel measurable when written as a function from \mathbb{S}_n to \mathbb{R} . We proceed by backward induction on n . The value function clearly has the claimed form at the final time $n = N$. Now, fix any $n < N$ and assume that the value function is of the form claimed for $n + 1$. Then, Bellman’s recursion and the induction hypothesis show that

$$\begin{aligned} V(p_n, n) &= \inf_{A \in \mathbb{A}_n} E[V(p_{n+1}, n + 1) \mid A_n = A, p_n] \\ &= \inf_{A \in \mathbb{A}_n} E[H(p_{n+1}) - (N - n - 1)C \mid A_n = A, p_n] \\ &= \inf_{A \in \mathbb{A}_n} E[H(p_{n+1}) \mid A_n = A, p_n] - (N - n - 1)C \\ &= H(p_n) - C - (N - n - 1)C \\ &= H(p_n) - (N - n)C, \end{aligned} \tag{12}$$

where we have used (5) in Theorem 1 to obtain the fourth equality. Theorem 1 also shows that the infimum in (12) is attained when A satisfies $p_n(A) = u^*$, and so the described policy achieves the minimum in Bellman’s recursion.

We offer the following interpretation of the optimal reduction in entropy shown in Theorem 2. First, the entropy of a random variable uniformly distributed over $[a, b]$ is $\log(b - a)$. The quantity $2^{H(X)}$ for a continuous random variable X can then be interpreted as the length of the support of a uniform random variable with the same entropy as X . We refer to this quantity more simply as the ‘length of X ’. If the prior distribution of X^* is uniform over $[0, 1]$ then the length of X^* under p_0 is 1 and Theorem 2 shows that the expected length of X^* under p_N is no less than 2^{-CN} , where this bound on the expected length can be achieved using an optimal policy.

We conclude this section by discussing u^* and C in a few specific cases. In general, there are no simple expressions for u^* and C . However, in certain symmetric cases the following proposition shows that $u^* = \frac{1}{2}$.

Proposition 1. *If the channel has the symmetry*

$$\varphi(u) = \varphi(1 - u) \quad \text{for all } 0 \leq u \leq 1 \tag{13}$$

then $\frac{1}{2} \in \arg \max_{u \in [0,1]} \varphi(u)$ and we may take $u^* = \frac{1}{2}$. Furthermore, if

$$H(uf_1 + (1 - u)f_0) = H(uf_0 + (1 - u)f_1) \quad \text{for all } u \in [0, 1] \tag{14}$$

then this is sufficient to guarantee (13).

Proof. Let u' be a maximizer of $\varphi(u)$. It might be equal to u^* , or if there is more than one maximizer, it might differ. Note that $\frac{1}{2} = \frac{1}{2}u' + \frac{1}{2}(1 - u')$. The function φ is concave (see [6, Theorem 2.7.4, Chapter 2]), implying that $\varphi(\frac{1}{2}) \geq \frac{1}{2}\varphi(u') + \frac{1}{2}\varphi(1 - u')$. Now, using $\varphi(u') = \varphi(1 - u')$, we obtain $\varphi(\frac{1}{2}) \geq \varphi(u')$, which shows that $\frac{1}{2} \in \arg \max_{u \in [0,1]} \varphi(u)$. If (14) is met then $H(f_0) = H(f_1)$ by taking $u = 0$, and (13) follows directly from the definition of φ .

TABLE 1: Channel capacity and the value u^* at which the channel capacity is achieved. The binary symmetric case is treated in [31].

Channel	Model			Channel capacity	u^*
Binary symmetric	0		1	$1 - h(\varepsilon)$	$\frac{1}{2}$
	f_0	$1 - \varepsilon$	ε		
	f_1	ε	$1 - \varepsilon$		
Binary erasure	0		1	$1 - \varepsilon$	$\frac{1}{2}$
	f_0	$1 - \varepsilon$	0		
	f_1	0	$1 - \varepsilon$		
Z	0		1	$h(u^*(1 - \varepsilon)) - u^*h(\varepsilon)$	$\frac{1/(1 - \varepsilon)}{1 + e^{h(\varepsilon)/(1 - \varepsilon)}}$
	f_0	1	0		
	f_1	ε	$1 - \varepsilon$		
Multivariate normal	$f_0 \sim N(m_0, \Sigma)$			Not analytical	$\frac{1}{2}$
Symmetric	$f_1 \sim N(m_1, \Sigma)$			Not analytical	$\frac{1}{2}$

A few simple channels with expressions for u^* and C are presented in Table 1. We use the notation $B(u)$ for a Bernoulli random variable with parameter u and $h(u)$ for $H(B(u))$, the entropy of this random variable. In the multivariate normal case, $u^* = \frac{1}{2}$ follows from Proposition 1 because $uf_1 + (1 - u)f_0$ is the multivariate normal density with mean $um_1 + (1 - u)m_0$ and variance Σ , and the entropy of a multivariate normal distribution does not depend on its mean, implying that (14) is satisfied.

4. One-dimensional optimal policies

We now present two specific policies in dimension $d = 1$ that satisfy the sufficient conditions for optimality given in Theorem 2: the probabilistic bisection policy and the dyadic policy. After defining these two policies in Sections 4.1 and 4.2, we study the sequence of entropies $(H(p_n) : n \geq 1)$ that they generate, focusing on the dyadic policy. In addition to Theorem 2, which shows that $E^\pi[H(p_n)] = H(p_0) - nC$ for any optimal policy π , the analysis of the dyadic policy in Section 4.2 provides a strong law of large numbers and a central limit theorem for $H(p_n)$. In further analysis of the dyadic policy, in Section 4.3 we analyze the number of size-limited noise-free questions required to find X^* after noisy questioning with the dyadic policy ceases, which is a metric that is important in the application discussed in Section 6.

To support the analyses in Sections 4.1 and 4.2, we first give here a general expression for the one-step change in entropy, $H(p_{n+1}) - H(p_n)$, under any policy π satisfying $p_n(A_n) = u^*$.

Lemma 2. *We have*

$$\begin{aligned}
 &H(p_{n+1}) - H(p_n) \\
 &= -D\left(B\left(\frac{u^* f_1(y)}{Z}\right), B(u^*)\right) \\
 &\quad + \frac{u(1 - u^*)}{Z}(f_1(y) - f_0(y))(H(p_n^+) - \log u^* - H(p_n^-) + \log(1 - u^*)), \quad (15)
 \end{aligned}$$

where D is the Kullback–Leibler divergence.

Proof. First, we define two densities:

$$p_n^+(x) = \begin{cases} \frac{p_n(x)}{u^*} & \text{if } x \in A_n, \\ 0 & \text{if } x \in \bar{A}_n, \end{cases} \quad p_n^-(x) = \begin{cases} \frac{p_n(x)}{1-u^*} & \text{if } x \in \bar{A}_n, \\ 0 & \text{if } x \in A_n, \end{cases}$$

where \bar{A}_n is the complement of A_n . Their respective entropies are

$$H(p_n^+) = \log u^* - \frac{1}{u^*} \int_{A_n} p_n(x) \log p_n(x) \, dx,$$

$$H(p_n^-) = \log(1-u^*) - \frac{1}{1-u^*} \int_{\bar{A}_n} p_n(x) \log p_n(x) \, dx,$$

and $H(p_n) = u^*H(p_n^+) + (1-u^*)H(p_n^-) + h(u^*)$.

Using Lemma 1, for a given observation $Y_{n+1} = y$, we have

$$\begin{aligned} H(p_{n+1}) &= \log Z - p_{n+1}(A_n) \log f_1(y) - p_{n+1}(\bar{A}_n) \log f_0(y) \\ &\quad - \frac{1}{Z} f_1(y) \int_{A_n} p_n(x) \log p_n(x) \, dx - \frac{1}{Z} f_0(y) \int_{\bar{A}_n} p_n(x) \log p_n(x) \, dx \\ &= \log Z - \frac{1}{Z} u f_1(y) \log f_1(y) - \frac{1}{Z} (1-u) f_0(y) \log f_0(y) \\ &\quad - \frac{1}{Z} u^* f_1(y) (\log u^* - H(p_n^+)) - \frac{1}{Z} (1-u^*) f_0(y) (\log(1-u^*) - H(p_n^-)). \end{aligned}$$

Expanding and rearranging, we obtain (15).

Note also that, under an optimal policy, the density of Y_{n+1} is the mixture of densities $u^* f_1 + (1-u^*) f_0$ according to Lemma 1, and the random variables Y_1, Y_2, \dots are independent and identically distributed (i.i.d.).

4.1. Probabilistic bisection policy

We first consider the case when questions are limited to intervals $A = (-\infty, a)$, $a \in \mathbb{R}$. This limitation appears naturally in applications, such as stochastic root finding [24] and signal estimation [5]. In this case, an optimal policy consists of choosing a_n such that $\int_{-\infty}^{a_n} p_n(x) \, dx = u^*$. Such an a_n always exists, but is not necessarily unique.

When the model is symmetric in the sense of Proposition 1, $u^* = \frac{1}{2}$, and a_n is the median of p_n . This policy of measuring at the median of the posterior is the probabilistic bisection policy introduced in [12]. Thus, the optimal policy with interval questions and general channels is a generalization of the probabilistic bisection policy, and we continue to refer to it as the probabilistic bisection policy even when $u^* \neq \frac{1}{2}$.

We briefly consider the behavior of $(H(p_n) : n \geq 1)$ under the probabilistic bisection policy. We assume a binary symmetric channel with noise parameter ε . Recall that $u^* = \frac{1}{2}$ in this case, and

$$D\left(B\left(\frac{f_1(Y_{n+1})}{2Z}\right), B\left(\frac{1}{2}\right)\right) = 1 - h(\varepsilon). \tag{16}$$

Moreover,

$$H(p_{n+1}) - H(p_n) = h(\varepsilon) - 1 + \left(\frac{1}{2} - \varepsilon\right) W_{n+1} (H(p_n^+) - H(p_n^-)),$$

where the W_n are i.i.d. Rademacher random variables. In this situation, even when p_0 is the density of the uniform distribution over the interval $[0, 1]$, the behavior of the process $H(p_n)$ can be complicated. A simulation of $H(p_n)$ is presented in Figure 1. The high degree of




FIGURE 1: The process $H(p_n)$ for the binary symmetric channel. Here p_0 is Uniform($[0, 1]$). *Top*: the questions are chosen by the probabilistic bisection policy. *Top left*: $\epsilon = 0.2$ and $C = 0.28$. *Top right*: $\epsilon = 0.4$ and $C = 0.03$. *Bottom*: the questions are chosen according to the dyadic policy. *Bottom left*: binary symmetric channel $\epsilon = 0.2$. *Bottom right*: normal channel, with $f_0 \sim N(-1, 1)$, $f_1 \sim N(1, 1)$, and $C = 0.47$.

variation of $H(p_n)$ around its mean value evident in Figure 1 may be disadvantageous in some applications. We do not pursue the probabilistic bisection policy further in this paper.

4.2. Dyadic policy

Consider now the situation where all sets in \mathbb{A}_n are available as questions, and p_0 is piecewise constant with finite support. Let $I = \{I_k : k = 0, \dots, K - 1\}$ be a finite partition of the support of p_0 into intervals such that p_0 is constant and strictly positive in each of these intervals. We assume that each interval I_k is closed on the left and open on the right, so $I_k = [a_k, b_k)$ with $a_k \in \mathbb{R}$ and $b_k \in \mathbb{R}$. This assumption is without loss of generality, because if it is not met, we can alter the prior density p_0 on a set of Lebesgue measure 0 (which does not change the corresponding prior probability measure) to meet it. We also assume that the constants J_n used to construct \mathbb{A}_n satisfy $J_n \geq 2^{n+1}K$. If this restriction is not met then we are free to increase J_n in most applications.

For each $k = 0, \dots, K - 1$, we partition I_k into two intervals, $A_{0,2k}$ and $A_{0,2k+1}$:

$$A_{0,2k} = [a_{0,2k}, b_{0,2k}) = [a_k, a_k + u^*(b_k - a_k)),$$

$$A_{0,2k+1} = [a_{0,2k+1}, b_{0,2k+1}) = [a_k + u^*(b_k - a_k), b_k).$$

With this partition, the mass $p_0(A_{0,2k}) = u^* p_0(I_k)$. The question asked at time 0 is

$$A_0 = \bigcup_{k=0}^{K-1} A_{0,2k},$$

and $p_0(A_0) = u^*$.




FIGURE 2: Illustration of the dyadic policy when p_0 is uniform on $[0, 1]$ and $u^* = \frac{5}{8}$. The prior is displayed above illustrations of the sets $A_{n,k}$ for $n = 0, 1, 2$. Each question A_n is the union of the dark gray subsets $A_{n,k}$ for that value of n .

We use a similar procedure recursively for each $n = 0, 1, \dots$ to partition each $A_{n,k}$ into two intervals, $A_{n+1,2k}$ and $A_{n+1,2k+1}$, and then construct the question A_{n+1} from these partitions. Let $K_n = 2^{n+1}K$, and, for $k = 0, \dots, K_n - 1$, define

$$A_{n+1,2k} = [a_{n+1,2k}, b_{n+1,2k}) = [a_{n,k}, a_{n,k} + u^*(b_{n,k} - a_{n,k})),$$

$$A_{n+1,2k+1} = [a_{n+1,2k+1}, b_{n+1,2k+1}) = [a_{n,k} + u^*(b_{n,k} - a_{n,k}), b_{n,k}).$$

Then, from these, the question to be asked at time $n + 1$ is

$$A_{n+1} = \bigcup_{k=0}^{K_n-1} A_{n+1,2k}.$$

This construction is illustrated in Figure 2.

Observe that $p_{n+1}(A_{n+1,2k}) = u^* p_{n+1}(A_{n,k})$ implies that

$$p_{n+1}(A_{n+1}) = \sum_{k=0}^{K_n-1} u^* p_{n+1}(A_{n,k}) = u^*$$

because $\{A_{n,k} : k = 0, \dots, K_n - 1\}$ is a partition of the support of p_0 . Thus, this construction satisfies $p_n(A_{n+1}) = u^*$, and is optimal. In addition, the sets A_0, \dots, A_{n-1} are constructed without knowledge of the responses, and, thus, this policy is nonadaptive. This is useful in applications, allowing multiple questions to be asked simultaneously. We call this policy the dyadic policy because each question is constructed by dividing the previous question’s intervals into two pieces.

We now provide an analysis that leads to a law of large numbers and a central limit theorem for $H(p_n)$ under this policy when n is large. Under the dyadic policy, we have

$$H(p_n^+) = H(p_n) + \log u^* \quad \text{and} \quad H(p_n^-) = H(p_n) + \log(1 - u^*),$$

which implies, using (15), that

$$H(p_{n+1}) - H(p_n) = -D\left(B\left(\frac{u^* f_1(Y_{n+1})}{u^* f_1(Y_{n+1}) + (1 - u^*) f_0(Y_{n+1})}\right), B(u^*)\right), \quad (17)$$

where Y_n is, as already stated, a sequence of i.i.d. random variables with density $u^* f_1 + (1 - u^*) f_0$. We read from (17) that $H(p_n)$ is, in this case, a sum of i.i.d. random variables. Moreover, each one is bounded above and below. Indeed,

$$0 \leq D\left(B\left(\frac{u^* f_1(Y_{n+1})}{u^* f_1(Y_{n+1}) + (1 - u^*) f_0(Y_{n+1})}\right), B(u^*)\right) \leq \max(D(B(0), B(u^*)), D(B(1), B(u^*))),$$

implying the bound

$$\min(\log(u^*), \log(1 - u^*)) \leq H(p_{n+1}) - H(p_n) \leq 0. \tag{18}$$

For the binary symmetric channel, (17) reduces to a constant, as noted in (16). This proves the following theorem.

Theorem 3. *For any piecewise constant p_0 , using the dyadic policy,*

$$\lim_{n \rightarrow \infty} \frac{H(p_n)}{n} = -C \text{ almost surely (a.s.)} \tag{19}$$

and

$$\lim_{n \rightarrow \infty} \frac{H(p_n) + nC}{\sqrt{n}} \stackrel{D}{=} N(0, \sigma^2), \tag{20}$$

where σ^2 is the variance of the increment $H(p_{n+1}) - H(p_n)$, which can be computed from the distribution given in (17). A degenerate situation occurs for the binary symmetric channel with noise ε . In this case, the sequence $H(p_n) = H(p_0) - nC$ is constant.

The dyadic policy is illustrated in the two bottom diagrams of Figure 1, where $H(p_n)$ is plotted as a function of n . The binary symmetric channel model with $\varepsilon = 0.2$ is shown in the bottom-left diagram. The sequence $H(p_n)$ is constant, in sharp contrast with the behavior of $H(p_n)$ for the same model under the probabilistic bisection policy, shown in the top-left diagram. Finally, a normal channel is presented in the bottom-right diagram.

4.3. Expected number of noise-free questions

In this section we consider an alternative to entropy for measuring performance, which arises in the example considered in Section 6. We suppose that, in addition to the noisy questions previously discussed, we also have the ability to ask a noise-free oracle whether X^* lies in a given set, where the sets about which we can ask noise-free questions come from some restricted class, e.g. their size is below a threshold. In Section 6, the sets about which we can ask noise-free questions correspond to pixels in an image. We suppose that after a fixed number N of noisy questions, we query sets using the noise-free questions until we find X^* . The loss function that arises naturally in this situation is the expected number of noise-free questions until X^* is found.

Given a posterior p_N resulting from the first stage of noisy questions, the optimal way in which to ask the noise-free questions is to first sort the available sets about which noise-free questions can be asked, in decreasing order of their probability of containing X^* under p_N . Then, query these sets in this order until X^* is found. Observing that X^* is not in a particular set alters the probability of the other sets, but does not change the order of these probabilities. Thus, it is sufficient to ask the noise-free questions in an order that depends only upon p_N , and no subsequent information.

We give below in Theorem 4 an explicit expression for the expected number of noise-free questions required after the dyadic policy completes. Before giving this expression in Theorem 4, we have the following preliminary result. In both this result and Theorem 4, we assume the dyadic policy, a uniform p_0 , and a binary symmetric channel with noise parameter ε .

Proposition 2. *For each $k \in \{0, \dots, 2^N - 1\}$, let M_k be the number of noisy questions A_n whose answer has indicated $X^* \in A_{N-1,k}$, either via $A_{N-1,k} \subseteq A_n$ and $Y_n = 1$, or $A_{N-1,k} \subsetneq I \setminus A_n$ and $Y_n = 0$. Then, the density of $A_{N-1,k}$ under p_N is*

$$|I|^{-1} 2^N (1 - \varepsilon)^{M_k} \varepsilon^{N - M_k}.$$

Furthermore, for each $m \in \{0, \dots, N\}$, the number of k with $M_k = m$ is deterministic, and is equal to $\binom{N}{m}$.

Proof. In the proof, we refer to $A_{N-1,k}$ as B_k . During noisy questioning, each time an answer indicates $X^* \in B_k$, we multiply the posterior density on B_k by $2(1 - \varepsilon)$, and each time an answer indicates $X^* \notin B_k$, we multiply by 2ε . Since the prior density was $|I|^{-1}$, the posterior density on B_k after all N measurements is $|I|^{-1} 2^N (1 - \varepsilon)^{M_k} \varepsilon^{N - M_k}$.

For each $k \in \{0, \dots, 2^N - 1\}$, let $b_{kn} = \mathbf{1}_{\{B_k \subseteq A_n\}}$, and define the binary sequence $b_k = (b_{k1}, \dots, b_{kN})$. By construction of the sets B_k under the dyadic policy, each b_k is unique. Since there are 2^N possible binary sequences of N bits, and 2^N sets B_k , the mapping between B_k and b_k is a bijection.

Consider a sequence of answers to noisy questions, Y_1, \dots, Y_N . For each b_k , define a subset $D_k = \{n \in \{1, \dots, N\} : b_{kn} = Y_n\}$. Each b_k defines a unique subset D_k . Since there are 2^N subsets and 2^N sequences b_k , each subset $D \subseteq \{1, \dots, N\}$ is equal to some D_k . Thus, the mapping between b_k and D_k is a bijection.

Because $M_k = |D_k|$, the number of k with $M_k = m$ is equal to the number of subsets $D \subseteq \{1, \dots, N\}$ of size m . This number is exactly $\binom{N}{m}$.

Proposition 2 shows that the number of sets $A_{N-1,k}$ with any given posterior density $|I|^{-1} 2^N (1 - \varepsilon)^m \varepsilon^{N - m}$ is deterministic. Figure 3(a) shows this posterior probability distribution, after sorting the sets according to their density, for $N = 5$ and $\varepsilon = 0.3$. The expectation under p_N of the number of noise-free questions required to find X^* depends only upon this sorted posterior probability density, and is thus also deterministic. We now give an expression for this expectation in Theorem 4.

Theorem 4. *In each interval $A_{N-1,k}$ for $k = 0, \dots, 2^N - 1$, assume that there are ℓ disjoint, equally sized sets about which we can ask noise-free questions. Then the expectation under p_N of the number of noise-free questions required to find X^* is*

$$\sum_{m=0}^N \binom{N}{m} (1 - \varepsilon)^m \varepsilon^{N - m} \left[\frac{\binom{N}{m} + 1/\ell}{2} + \sum_{m'=m+1}^N \binom{N}{m'} \right] \ell.$$

Proof. First, if we have a collection of disjoint subsets C_1, \dots, C_K , each with probability $1/K$ of containing X^* , and we query each subset in order of increasing index until we find X^* , then we ask k questions when $X^* \in C_k$ and the expected number of questions asked is $\sum_{k=1}^K k P(X^* \in C_k) = \sum_{k=1}^K k/K = (K + 1)/2$. Under p_N , Proposition 2 shows that X^* has probability

$$\binom{N}{m} (1 - \varepsilon)^m \varepsilon^{N - m} \tag{21}$$

of being in a subset $A_{N-1,k}$ with $M_k = m$, because there are $\binom{N}{m}$ such intervals, each of




FIGURE 3: (a) The posterior density p_N for the binary symmetric channel with the dyadic policy, with subsets $A_{N-1,k}$ sorted in order of decreasing posterior density $p_N(x)$, and $N = 5$. (b) The expected number of noise-free questions as a function of N , for a fixed collection of 2^{16} subsets about which noise-free questions may be asked. In both (a) and (b), $\varepsilon = 0.3$.

size $2^{-N}|I|$, and each of density $|I|^{-1}2^N(1 - \varepsilon)^m\varepsilon^{N-m}$. Then, because the number of noise-free questions available in each $A_{N-1,k}$ is ℓ , the expected number of noise-free questions, conditioned on X^* being in a subset $A_{N-1,k}$ with $M_k = m$, is

$$\frac{\binom{N}{m}\ell + 1}{2} + \sum_{m'=m+1}^N \binom{N}{m'}\ell. \tag{22}$$

Here, the first term is the number of questions asked in subsets with $M_k = m$, and the second term is the number of questions asked in subsets with $M_k > m$, which had a strictly higher density $p_N(x)$ and were queried earlier. The result follows by combining (21) and (22) and summing over k .

Using Theorem 4, we consider the effect of varying N . Suppose that the sets about which noise-free questions may be asked are pixels in an image, as in the example in Section 6. Take $I = [0, 1]$, and suppose that each pixel is of size 2^{-L} and occupies a region $[k2^{-L}, (k+1)2^{-L}]$ for some $k = 0, \dots, 2^L$. If sets $A_{N-1,k}$ must contain an integer numbers of pixels then we may naturally consider any N between 0 and L . For any such N , the number of pixels ℓ in a subset $A_{N-1,k}$ is $\ell = 2^{L-N}$. In this setting, the expected number of noise-free questions asked as a function of N is shown in Figure 3(b) for $L = 16$ and $\varepsilon = 0.3$. It can be seen from the figure that there is a dramatic decrease in the expected number of noise-free questions as the number of noisy questions N increases.

5. Optimal policies in two dimensions with entropy loss

We now consider the case $d = 2$, in which X^* is a two-dimensional random variable, $X^* = (X_1^*, X_2^*)$, with joint density p_0 . To minimize the expected entropy $E[H(p_N)]$ of the two-dimensional posterior distribution on X^* at time N , Theorem 2 from Section 3 shows that it is optimal to use any policy satisfying $p_n(A_n) = u^*$.

While the objective function $E[H(p_N)]$ is natural in dimension $d = 1$, it has a drawback in $d = 2$ and higher dimensions. This is well illustrated using an example. Assume that X_1^* and

X_2^* are independent and uniformly distributed over intervals of lengths s_1 and s_2 , respectively. Then $H(p) = \log(s_1 s_2)$. In this case, $H(p)$ can be arbitrarily small even if the entropy of one of the marginal densities remains large, e.g. $s_2 = 1$.

This leads us to consider objective functions without this drawback. For example, we might wish to solve $\inf_{\pi} E^{\pi}[\max(H_1(p_N), H_2(p_N))]$, where $H_1(p_N) = H(\int p_N(\cdot, u_2) du_2)$ and $H_2(p_N) = H(\int p_N(u_1, \cdot) du_1)$ are the entropies of the marginals. However, solving this problem directly seems out of reach. Instead, we focus on reducing $E^{\pi}[\max(H_1(p_N), H_2(p_N))]$ at an asymptotically optimal rate by solving

$$V(p) = \inf_{\pi} \liminf_{N \rightarrow \infty} \frac{1}{N} E^{\pi}[\max(H_1(p_N), H_2(p_N)) \mid p_0 = p]. \tag{23}$$

We use the \liminf to include policies for which the limit might not exist. Throughout this section, we assume that both $H_1(p_0)$ and $H_2(p_0)$ are finite.

For further simplification, we assume that questions concern only one coordinate. That is, the sets queried are either of type 1, $A_n = B \times \mathbb{R}$, where B is a finite union of intervals of \mathbb{R} , or, alternatively, of type 2, $A_n = \mathbb{R} \times B$. In each case, we assume that the response passes through a memoryless noisy channel with densities $f_0^{(1)}$ and $f_1^{(1)}$ for questions of type 1, and $f_0^{(2)}$ and $f_1^{(2)}$ for questions of type 2. Let C_1 and C_2 be the channel capacities for questions of type 1 and 2, respectively. We also assume that p_0 is a product of its marginals. This guarantees that p_n for all $n > 0$ remains a product of its marginals and that only one marginal distribution is modified at each point in time. This is shown by the following lemma.

Lemma 3. *Assume that $p_n(u_1, u_2) = p_n^{(1)}(u_1)p_n^{(2)}(u_2)$ and that we choose a question of type 1 with $A_n = B \times \mathbb{R}$. Then, given $Y_{n+1} = y$,*

$$p_{n+1}(u_1, u_2) = \frac{1}{Z_1} (f_1^{(1)}(y) \mathbf{1}_{\{u_1 \in B\}} + f_0^{(1)}(y) \mathbf{1}_{\{u_1 \notin B\}}) p_n^{(1)}(u_1) p_n^{(2)}(u_2),$$

where $Z_1 = f_1^{(1)}(y) p_n^{(1)}(B) + f_0^{(1)}(y) (1 - p_n^{(1)}(B))$.

Similarly, if we choose a question of type 2 with $A_n = \mathbb{R} \times B$ then

$$p_{n+1}(u_1, u_2) = \frac{1}{Z_2} (f_1^{(2)}(y) \mathbf{1}_{\{u_2 \in B\}} + f_0^{(2)}(y) \mathbf{1}_{\{u_2 \notin B\}}) p_n^{(2)}(u_2) p_n^{(1)}(u_1),$$

where $Z_2 = f_1^{(2)}(y) p_n^{(2)}(B) + f_0^{(2)}(y) (1 - p_n^{(2)}(B))$.

Proof. The proof is straightforward using Bayes formula, and is similar to the proof of Lemma 1 in the one-dimensional case.

In the two-dimensional setting, any policy can be understood as making two decisions at each time n . The first decision is which coordinate to query, that is, whether to ask a question of type 1 or type 2. Given this choice, the second decision is which question of this type to ask, which corresponds to a finite union of intervals of \mathbb{R} . As before, these decisions may depend only upon the information gathered by time n , for which the corresponding sigma-algebra is \mathcal{F}_n . For $N > 0$, let S_N be the number of questions of type 1 answered by time N . That is, S_N is the number of $n \in \{0, \dots, N - 1\}$ such that A_n is of the form $A_n = B \times \mathbb{R}$. We take $S_0 = 0$.

We first present a lower bound on the expected decrease in the entropy of each marginal posterior distribution.

Lemma 4. *Under any valid policy π ,*

$$E^{\pi}[H_1(p_n)] \geq H_1(p_0) - C_1 E^{\pi}[S_n], \quad E^{\pi}[H_2(p_n)] \geq H_2(p_0) - C_2 (n - E^{\pi}[S_n]).$$

Proof. Define $M_n^{(1)} = H_1(p_n) + C_1 S_n$ and $M_n^{(2)} = H_2(p_n) + C_2(n - S_n)$. We will show that $M^{(1)}$ and $M^{(2)}$ are submartingales. Focusing first on $M^{(1)}$, we calculate

$$E^\pi [M_{n+1}^{(1)} \mid \mathcal{F}_n] = E^\pi [H_1(p_{n+1}) \mid \mathcal{F}_n] + C_1 S_{n+1}$$

since S_{n+1} is \mathcal{F}_n -measurable. We consider two cases. First, if $S_{n+1} = S_n$ (which occurs if A_n is of type 2) then $H_1(p_{n+1}) = H_1(p_n)$ and the \mathcal{F}_n -measurability of $H_1(p_n)$ implies that $E^\pi [M_{n+1}^{(1)} \mid \mathcal{F}_n] = M_n^{(1)}$. Second, if $S_{n+1} = S_n + 1$ (which occurs if A_n is of type 1) then Theorem 2 implies that $E^\pi [H_1(p_{n+1}) \mid \mathcal{F}_n] \geq H_1(p_n) - C_1$. Hence,

$$E^\pi [M_{n+1}^{(1)} \mid \mathcal{F}_n] \geq C_1(S_n + 1) + H_1(p_n) - C_1 = M_n^{(1)},$$

which shows that $M_n^{(1)}$ is a submartingale. The proof is similar for $M_n^{(2)}$.

Now, because $M_n^{(1)}$ is a submartingale, $E^\pi [M_n^{(1)}] \geq M_0^{(1)}$, which implies that $E^\pi [H_1(p_n)] \geq H_1(p_0) - C_1 E^\pi [S_n]$. Proceeding similarly for $M_n^{(2)}$ completes the proof.

Consider the following policy, notated π^* . At step n , choose the type of question at random, choosing type 1 with probability $C_2/(C_1 + C_2)$ and type 2 with probability $C_1/(C_1 + C_2)$. Then, in the dimension chosen, choose the subset to be queried according to the one-dimensional dyadic policy.

We show below in Theorem 5 that π^* is optimal for the objective function (23). Before presenting this result, which is the main result of this section, we present an intermediate result concerning the limiting behavior of π^* . This intermediate result is essentially a strong law of large numbers for the objective function (23).

Lemma 5. *Let $T_N = (1/N) \max(H_1(p_N), H_2(p_N))$. Under π^* , as $N \rightarrow \infty$,*

$$T_N \rightarrow -\frac{C_1 C_2}{C_1 + C_2} \quad \text{a.s.} \tag{24}$$

Moreover, there is a constant K such that $|T_N| < K$ for all N .

Proof. Recall that S_N is the number of questions of type 1 answered by time N , so $S_N/N \rightarrow C_2/(C_1 + C_2)$ a.s. The law of large numbers established in (19) for the one-dimensional posterior shows that $H_1(p_N)/S_N \rightarrow -C_1$ a.s. Combining these two facts shows that $H_1(p_N)/N \rightarrow -C_1 C_2/(C_1 + C_2)$ a.s. By a similar argument, $H_2(p_N)/N \rightarrow -C_1 C_2/(C_1 + C_2)$ a.s., which shows (24).

We now show the bound on $|T_N|$. Using π^* , according to (18),

$$H_1(p_N) = H_1(p_0) + \sum_{n=1}^N Z_n,$$

where the Z_n are independent bounded random variables and

$$|Z_n| \leq |\min(\log(u), \log(1 - u))| = \beta.$$

As a consequence, for any $N \geq 1$, $|H_1(p_N)/N| \leq |H_1(p_0)| + \beta$. The same is true for $H_2(p_N)$, which proves that there is a constant K such that $|T_N| < K$.

We now present the main result of this section.

Theorem 5. *The policy π^* is optimal with respect to (23). Moreover, the optimal value is, for any p_0 with $H(p_0) < \infty$,*

$$V(p_0) = -\frac{C_1 C_2}{C_1 + C_2}. \tag{25}$$

Proof. First we show that the value in (25) constitutes a lower bound for $V(p_0)$. Second, we show that (25) is an upper bound on $V(p_0)$ using the properties of the policy π^* presented in Lemma 5.

For the lower bound,

$$\begin{aligned} V(p_0) &\geq \inf_{\pi} \liminf_{N \rightarrow \infty} \frac{1}{N} \max(\mathbb{E}^{\pi}[H_1(p_N)], \mathbb{E}^{\pi}[H_2(p_N)]) \\ &\geq \inf_{\pi} \liminf_{N \rightarrow \infty} \frac{1}{N} \max(H_1(p_0) - \mathbb{E}[S_N]C_1, H_2(p_0) - (N - \mathbb{E}[S_N])C_2) \\ &= \inf_{0 \leq a \leq 1} \max(-aC_1, -(1-a)C_2) \\ &= -\frac{C_1 C_2}{C_1 + C_2}. \end{aligned}$$

We obtained the first line using Jensen’s inequality, the second line using Lemma 4, the third line by choosing $a = \liminf_{n \rightarrow \infty} \mathbb{E}[S_N]/N$, and the fourth line by recalling that $C_1 > 0$ and $C_2 > 0$.

Now, for the upper bound,

$$\begin{aligned} V(p_0) &\leq \liminf_{N \rightarrow \infty} \mathbb{E}^{\pi^*} \left[\max\left(\frac{H_1(p_N)}{N}, \frac{H_2(p_N)}{N}\right) \right] \\ &= \mathbb{E}^{\pi^*} \left[\max\left(\liminf_{N \rightarrow \infty} \frac{H_1(p_N)}{N}, \liminf_{N \rightarrow \infty} \frac{H_2(p_N)}{N}\right) \right] \\ &= -\frac{C_1 C_2}{C_1 + C_2}. \end{aligned}$$

The uniform bound on T_N from Lemma 5 is sufficient to justify the exchange between the limit and the expected value in going from the first to the second line.

We remark as an aside that in the case where $C_1 = C_2$, this policy is also optimal for the value function (3) since it verifies (11).

We conclude this section by providing a central limit theorem for the objective under this policy π^* .

Theorem 6. *Under π^* ,*

$$\lim_{n \rightarrow \infty} \frac{1}{\sqrt{n}} \left[\max(H_1(p_n), H_2(p_n)) + \frac{C_1 C_2}{C_1 + C_2} n \right] \stackrel{D}{=} \frac{\max(\sigma_1 \sqrt{C_2} Z_1, \sigma_2 \sqrt{C_1} Z_2)}{\sqrt{C_1 + C_2}}. \tag{26}$$

Here, Z_1 and Z_2 are independent standard normal random variables, and σ_i^2 is the variance of the increment of $H_i(p_{n+1}) - H_i(p_n)$ when measuring type i , whose distribution is given by (17).

Proof. For $i = 1, 2$, let $S_{n,i}$ be the number of questions of type i answered by time n , so $S_{n,1} = S_n$ and $S_{n,2} = n - S_n$. Let $t_{s,i} = \inf\{n : S_{n,i} = s\}$ for $s = 0, 1, \dots$. Then $t_{0,i} = 0$ and $\{t_{s,i} : s = 1, 2, \dots\}$ are the times when questions of type i are answered. Thus, each stochastic process $\{H_i(p_{t_{s,i}}) : s = 0, 1, \dots\}$ for $i = 1, 2$ has a distribution identical to that of the entropy of

the one-dimensional posterior under the dyadic policy. In addition, the two stochastic processes are independent.

The central limit theorem established in (20) shows that

$$\lim_{s \rightarrow \infty} \frac{H_i(p_{t_{s,i}}) + sC_i}{\sqrt{s}} \stackrel{D}{=} \sigma_i Z_i,$$

where each Z_i is a standard normal random variable and Z_1 is independent of Z_2 .

From the definition of $t_{s,i}$,

$$\lim_{s \rightarrow \infty} \frac{H_i(p_{t_{s,i}}) + sC_i}{\sqrt{s}} \stackrel{D}{=} \lim_{n \rightarrow \infty} \frac{H_i(p_n) + S_{n,i}C_i}{\sqrt{S_{n,i}}}.$$

Let $j = 1$ when $i = 2$, and $j = 2$ when $i = 1$. Then $\lim_{n \rightarrow \infty} S_{n,i}/n = C_j/(C_1 + C_2)$ a.s. and

$$\lim_{n \rightarrow \infty} \frac{H_i(p_n) + S_{n,i}C_i}{\sqrt{S_{n,i}}} \stackrel{D}{=} \lim_{n \rightarrow \infty} \frac{H_i(p_n) + nC_1C_2/(C_1 + C_2)}{\sqrt{n}} \sqrt{\frac{C_1 + C_2}{C_j}}.$$

These three facts imply that

$$\lim_{n \rightarrow \infty} \frac{H_i(p_n) + nC_1C_2/(C_1 + C_2)}{\sqrt{n}} \stackrel{D}{=} \sqrt{\frac{C_j}{C_1 + C_2}} \sigma_i Z_i.$$

This proves (26) for the limit.

6. $\mathbb{L}\mathbb{T}\mathbb{E}\mathbb{X}$ character localization

In this section we present an application of the dyadic policy to a well-established problem in computer vision: object localization. While the probabilistic bisection policy has already been applied in computer vision, see [10] and [27], the dyadic policy has not, and we feel that it offers considerable promise in this application area.

In the object localization problem, we are given an image and a known object, and must output parameters that describe the *pose* of the object in the image. In the simplest case, the pose is defined by a single pixel, but more complex cases can include, e.g. a rotation angle, a scale factor, or a bounding box. Machine learning techniques have led to the development of classifiers that, given a specific pose, provide accurate answers to the binary question ‘Is the object in this pose?’. In our model, we assume these classifiers act as oracles, i.e. are perfect, even though they may occasionally classify incorrectly in practice. Classifiers such as support vector machines [28] and boosting [26] are combined with discriminant features (see, e.g. [18]) to provide the most accurate algorithms (see [29] and [30]). To find the object’s pose within an image, classifiers are evaluated at nearly every possible pose, which is computationally costly. We demonstrate that the use of the dyadic policy rather than this brute force approach considerably reduces the computational cost. Although a detailed comparison would be beyond the scope of the illustrative example we present here, the branch and bound algorithm used in [17] is an alternative methodology for reducing computational cost in object localization.

6.1. $\mathbb{L}\mathbb{T}\mathbb{E}\mathbb{X}$ character images, noisy queries, and model estimation

The task we consider is localizing a specific $\mathbb{L}\mathbb{T}\mathbb{E}\mathbb{X}$ character in a binary image. In this setting, an image is a binary matrix $I \in \{0, 1\}^{m \times m}$, where the image has m rows and m columns. A $\mathbb{L}\mathbb{T}\mathbb{E}\mathbb{X}$ character is another smaller binary image $J \in \{0, 1\}^{j \times j}$, where $j < m$. We present experiments where the character of interest, or pattern, is the letter ‘T’. We assume that the pattern is always




FIGURE 4: (a) Example of an image containing the character ‘T’. (b)–(d) Examples of subset-based questions. The gray sectors in images (b)–(d) represent the queried regions A_2^1 , A_2^2 , and A_1^3 , respectively.

present in the image, and fully visible (i.e. not occluded by other objects or only partially visible in the image). The goal is to find the location $X^* = (X_1^*, X_2^*)$ of the pixel at the upper-left corner of the pattern within the image.

We generated 1000 images, each of size 256×256 pixels. Each image has a black background (i.e. pixel values of zero), and contains a single fully visible ‘T’ at a random location in the image. This ‘T’ is a binary image of size 32×32 pixels (see Figure 4(a)). Noise is added to the image by flipping each pixel value independently with probability 0.1. We then randomly assign each image into one of two sets of approximately equal size: one for training and the other for testing. The training set is used to learn the noise model as described below, and the testing set is used to evaluate the performance of the algorithm.

In this task, querying a set A corresponds to asking whether the upper-left corner of the ‘T’ resides in this set. We use a simple image-processing technique to provide a noisy answer to this question. The technique we use is chosen for its simplicity, and other more complex image-processing techniques might produce more informative responses, improving the overall performance of the algorithm.

In describing this technique, we first observe that all the images are of size 256×256 pixels and so any pixel coordinate can be represented in base 2 using two 8-bit strings, or octets. For example, the pixel with column–row location (32,14) is represented by (00100000, 00001110). We define 16 sets of pixels. Let A_1^i , $i = 1, \dots, 8$, be the set of pixels whose column pixel coordinate has a 1 for its i th bit. Similarly, let A_2^i , $i = 1, \dots, 8$, be the set of pixels whose row pixel coordinate has a 1 for its i th bit. Figure 4(b)–(d) respectively show the sets A_1^1 , A_1^2 , and A_2^3 . For any given image I and set A_j^i , we define the response

$$y(A_j^i) = \sum_{x \in A_j^i} I(x) - \sum_{x \notin A_j^i} I(x), \tag{27}$$

where $I(x) \in \{0, 1\}$ is the binary image’s value at pixel x . The motivation for using the response defined by (27) is that $y(A_j^i)$ is more likely to be large when A_j^i contains the ‘T’.

Although the response $y(A_j^i)$ is entirely determined by the image I and the location of the ‘T’ within it, our algorithm models the response using a noise model of the form (1). For simplicity, we assume that both the density f_1 of $y(A)$ when A contains the ‘T’, and the density f_0 of $y(A)$ when A does not contain the ‘T’, are normal with respective distributions $N(\mu, \sigma^2)$ and $N(-\mu, \sigma^2)$. The training set is used to estimate these parameters, leading to $\mu = 64.76$ and $\sigma = 105.7$. Because the model is symmetric, $u^* = 0.5$. The channel capacity is estimated with Monte Carlo integration to be $C = 0.23$.

6.2. Prior, posterior, and algorithm

We let $X^* = (X_1^*, X_2^*)$, $X_1^* \in [0, 255]$ and $X_2^* \in [0, 255]$, with p_0 uniform over the domain of X^* . Since the sets A_j^i constrain only one coordinate, the posterior over X^* is a product distribution, as was discussed in Section 5. The posterior for each coordinate $j = 1, 2$ was computed in Lemma 3. We now specialize to the model at hand using the notation ‘ α ’ to define equality up to a term that does not depend on x_j :

$$p_8^{(j)}(x_j) \propto \prod_{i=1}^8 (f_1(y_j^i) \mathbf{1}_{\{x_j \in A_j^i\}} + f_0(y_j^i) \mathbf{1}_{\{x_j \notin A_j^i\}}),$$

$$\log p_8^{(j)}(x_j) \propto \sum_{\{i: x_j \in A_j^i\}} \log \frac{f_1(y_j^i)}{f_0(y_j^i)} \propto \sum_{\{i: x_j \in A_j^i\}} y_j^i.$$

The algorithm has two phases: (i) the noisy query phase and (ii) the noise-free query phase. The noisy query phase comes first, and uses the dyadic policy to obtain a posterior distribution on X^* . The implementation of this noisy query phase is facilitated by the nonadaptive nature of the dyadic policy’s questions, which allows us to compute the answers to the questions all at once. The noise-free query phase then uses the posterior resulting from the first phase, together with a sequence of size-limited noise-free questions, to determine the exact location of X^* .

Noisy query phase. Given an image I , we begin by computing $y(A_j^i) = y_j^i$ for each $j = 1, 2$ and $i = 1, \dots, 8$. We then compute $\ell(x)$ for each pixel x , which is proportional to the logarithm of the posterior density at x ,

$$\ell(x) = \sum_{\{i: x \in A_1^i\}} y_1^i + \sum_{\{i: x \in A_2^i\}} y_2^i.$$

The top row of Figure 5 shows example images from our test set, while the bottom row of Figure 5 shows the corresponding ℓ -images, in which the value of $\ell(x)$ is plotted for each pixel.




FIGURE 5: Pixel reordering: example images from the test set (*top*) and corresponding ℓ -images (*bottom*). Dark regions indicate pixels more likely to contain the character, while light regions are less likely.

Dark regions of the ℓ -image indicate pixels with large $\ell(x)$, which are more likely to contain the ‘T’.

Noise-free query phase. We sort the pixels in decreasing order of $\ell(x)$. We then sequentially perform noise-free evaluations at each pixel x in this order until the true pixel location X^* is found. To perform a noise-free evaluation at a given pixel, we compare the ‘T’ pattern with the 32×32 pixel square from the image with upper-left corner at x to see if they match. When X^* is found, we stop and record the number of noise-free evaluations performed.

6.3. Results

We validated the algorithm above by evaluating it on the test set described in Section 6.1. To do this, we ran the algorithm on each image and recorded the number of noise-free evaluations required to locate the target character. The results described below (i) demonstrate that the dyadic policy significantly reduces the number of noise-free evaluations required to locate the ‘T’ character, and (ii) allows us to visualize the results summarized in (11), (19), and (20) within the context of this application.

Recall that each image has $256 \times 256 = 65\,536$ pixels. Over 500 test images, the mean, median, and standard deviation of the number of noise-free evaluations are 2021.5, 647, and 4066.9, respectively. This corresponds to a speed-up factor of 15 over an exhaustive (and typical) search policy. Figure 6(a) shows the sample distribution of the number of noise-free evaluations. We also computed the entropy of the posterior distribution after the 16 noisy questions are answered. According to (11), $E[H(p_{16})] = H(p_0) - 16C = 16 - 16(0.23) = 12.32$, which is in agreement with the empirically observed value $E[H(p_{16})] = 12.3$ (with standard deviation 0.716). We also visualized the convergence of the entropy for *each* image, as predicted by the law of large numbers in (19). In Figure 6(b), we plot $H(p_n)/n$, $n = 0, \dots, 16$, for each image in our test set. The empirical variance at $n = 16$ is very small. Finally, according to (20), the distribution of $(H(p_n) - (H(p_0) - nC))/\sqrt{n}$ should be approximately normal. In Figure 7(a) we present the histogram and in Figure 7(b) we present a normal Q-Q plot, demonstrating close agreement with the normal distribution.




FIGURE 6: Noise-free evaluations and convergence in entropy. (a) The distribution of the number of noise-free evaluations needed to locate the target character. (b) Plot of $H(p_n)/n$ as a function of n . Each line corresponds to one image, with $H(p_n)/n$ plotted over $n = 1, \dots, 16$. $H(p_n)/n$ converges to $1 - C$.




FIGURE 7: Central limit theorem. (a) Distribution of $(H(p_n) - (H(p_0) - nC))/\sqrt{n}$, with mean -0.01 . The distribution is close to Gaussian as the Q-Q plot (b) shows.

7. Conclusion

We have considered the problem of twenty questions with noisy responses, which arises in stochastic search, stochastic optimization, computer vision, and other application areas. By considering the entropy as our objective function, we obtained sufficient conditions for Bayes optimality, which we then used to show optimality of two specific policies: probabilistic bisection and the dyadic policy. This probabilistic bisection policy generalizes a previously studied policy, while we believe that the dyadic policy has not been previously considered.

The dyadic policy asks a deterministic set of question, despite being optimal among fully sequential policies. This lends it to applications that allow multiple questions to be asked simultaneously. The structure of this policy also lends itself to further analysis. We provided a law of large numbers, a central limit theorem, and an analysis of the number of noise-free questions required after noisy questioning ceases. We also showed that a generalized version of the dyadic policy is asymptotically optimal in two dimensions for a more robust version of the entropy loss function. We then demonstrated the use of this policy on an example problem from computer vision.

A number of interesting and practically important questions present themselves for future work. First, our optimality results assume the entropy as the objective, but in many applications other objectives are more natural, e.g. the expected number of noise-free questions as in Section 4.3, or the mean-squared error. Second, our results assume that noise is added by a memoryless transmission channel. In many applications, however, the structure of the noise depends upon the questions asked, which calls for generalizing the results herein to this more complex style of noise dependence. We feel that these and other questions will be fruitful areas for further study.

References

- [1] BEN-OR, M. AND HASSIDIM, A. (2008). The Bayesian learner is optimal for noisy binary search (and pretty good for quantum as well). In *2008 49th Ann. IEEE Symp. Foundations of Computer Science*. IEEE Computer Society Press, Washington, DC, pp. 221–230.
- [2] BERRY, D. A. AND FRISTEDT, B. (1985). *Bandit Problems*. Chapman & Hall, London.
- [3] BLUM, J. R. (1954). Multidimensional stochastic approximation methods. *Ann. Math. Statist.* **25**, 737–744.

- [4] BURNAŠEV, M. V. AND ZIGANGIROV, K. Š. (1974). A certain problem of interval estimation in observation control. *Problemy Peredachi Informatsii* **10**, 51–61.
- [5] CASTRO, R. AND NOWAK, R. (2008). Active learning and sampling. In *Foundations and Applications of Sensor Management*, Springer, pp. 177–200.
- [6] COVER, T. M. AND THOMAS, J. A. (1991). *Elements of Information Theory*. John Wiley, New York.
- [7] DEGROOT, M. H. (1970). *Optimal Statistical Decisions*. McGraw Hill, New York.
- [8] DYNKIN, E. B. AND YUSHKEVICH, A. A. (1979). *Controlled Markov Processes*. Springer, New York.
- [9] FRAZIER, P. I., POWELL, W. B. AND DAYANIK, S. (2008). A knowledge-gradient policy for sequential information collection. *SIAM J. Control Optimization* **47**, 2410–2439.
- [10] GEMAN, D. AND JEDYNAK, B. (1996). An active testing model for tracking roads in satellite images. *IEEE Trans. Pattern Anal. Machine Intelligence* **18**, 1–14.
- [11] GITTINS, J. C. (1989). *Multi-Armed Bandit Allocation Indices*. John Wiley, Chichester.
- [12] HORSTEIN, M. (1963). Sequential decoding using noiseless feedback. *IEEE Trans. Inf. Theory* **9**, 136–143.
- [13] HORSTEIN, M. (2002). Sequential transmission using noiseless feedback. *IEEE Trans. Inf. Theory* **9**, 136–143.
- [14] KARP, R. M. AND KLEINBERG, R. (2007). Noisy binary search and its applications. In *Proc. 18th Ann. ACM-SIAM Symp. Discrete Algorithms*, ACM, New York, pp. 881–890.
- [15] KUSHNER, H. J. AND YIN, G. G. (2003). *Stochastic Approximation and Recursive Algorithms and Applications*, 2nd edn. Springer, New York.
- [16] LAI, T. L. AND ROBBINS, H. (1985). Asymptotically efficient adaptive allocation rules. *Adv. Appl. Math.* **6**, 4–22.
- [17] LAMPERT, C. H., BLASCHKO, M. B. AND HOFMANN, T. (2009). Efficient subwindow search: a branch and bound framework for object localization. *IEEE Trans. Pattern Anal. Machine Intelligence* **31**, 2129–2142.
- [18] LOWE, D. G. (2004). Distinctive image features from scale-invariant keypoints. *Internat. J. Comput. Vision* **60**, 91–110.
- [19] NOWAK, R. (2008). Generalized binary search. In *2008 46th Ann. Allerton Conf. Commun., Control, and Computing*, pp. 568–574.
- [20] NOWAK, R. (2009). Noisy generalized binary search. *Adv. Neural Inf. Processing Systems* **22**, 1366–1374.
- [21] PELC, A. (2002). Searching games with errors—fifty years of coping with liars. *Theoret. Comput. Sci.* **270**, 71–109.
- [22] POLYAK, B. T. (1990). A new method of stochastic approximation type. *Automat. Remote Control* **51**, 937–946.
- [23] ROBBINS, H. (1952). Some aspects of the sequential design of experiments. *Bull. Amer. Math. Soc.* **58**, 527–535.
- [24] ROBBINS, H. AND MONRO, S. (1951). A stochastic approximation method. *Ann. Math. Statist.* **22**, 400–407.
- [25] RUPPERT, D. (1988). Efficient estimators from a slowly convergent Robbins-Monro procedure. Tech. Rep. 781, School of Operations Research and Industrial Engineering, Cornell University.
- [26] SCHAPIRE, R. E. (1990). The strength of weak learnability. *Machine Learning* **5**, 197–227.
- [27] SZNITMAN, R. AND JEDYNAK, B. (2010). Active testing for face detection and localization. *IEEE Trans. Pattern Anal. Machine Intelligence* **32**, 1914–1920.
- [28] VAPNIK, V. N. (1995). *The Nature of Statistical Learning Theory*. Springer, New York.
- [29] VEDALDI, A., GULSHAN, V., VARMA, M. AND ZISSERMAN, A. (2009). Multiple kernels for object detection. In *Proc. Internat. Conf. Computer Vision*, pp. 606–613.
- [30] VIOLA, P. AND JONES, M. J. (2004). Robust real-time face detection. *Internat. J. Comput. Vision* **57**, 137–154.
- [31] WAEBER, R., FRAZIER, P. I. AND HENDERSON, S. G. (2011). A Bayesian approach to stochastic root finding. In *Proc. 2011 Winter Simulation Conference*, eds S. Jain et al., IEEE.
- [32] WHITTLE, P. (1981). Arm-acquiring bandits. *Ann. Prob.* **9**, 284–292.
- [33] WHITTLE, P. (1988). Restless bandits: activity allocation in a changing world. In *A Celebration of Applied Probability* (J. Appl. Prob. Spec. Vol. **25A**), ed. J. Gani, Applied Probability Trust, Sheffield, pp. 287–298.