

BISECTION SEARCH WITH NOISY RESPONSES*

ROLF WAEBER[†], PETER I. FRAZIER[†], AND SHANE G. HENDERSON[†]

Abstract. Bisection search is the most efficient algorithm for locating a unique point $X^* \in [0, 1]$ when we are able to query an oracle only about whether X^* lies to the left or right of a point x of our choosing. We study a noisy version of this classic problem, where the oracle's response is correct only with probability p . The probabilistic bisection algorithm (PBA) introduced by Horstein [*IEEE Trans. Inform. Theory*, 9 (1963), pp. 136–143] can be used to locate X^* in this setting. While the method works extremely well in practice, very little is known about its theoretical properties. In this paper, we provide several key findings about the PBA, which lead to the main conclusion that the expected absolute residuals of successive search results, i.e., $\mathbb{E}[|X^* - X_n|]$, converge to 0 at a geometric rate.

Key words. probabilistic bisection, noisy bisection, geometric rate of convergence, sequential analysis, Bayesian performance analysis

AMS subject classifications. 93E20, 68W40, 62C10, 60G35

DOI. 10.1137/120861898

1. Introduction. The goal of a bisection search is to locate a unique unknown point $X^* \in [0, 1]$. To obtain information about X^* , one queries an oracle as to whether X^* lies to the left or to the right of a given point x . If the oracle always answers the question correctly, the well-known bisection algorithm, which halves the search space at each iterate, can be used to locate X^* efficiently. Conversely, if the oracle's responses are noisy with probability $p \in (1/2, 1)$ of being correct, the deterministic bisection algorithm will fail almost surely, as a single wrong answer will divert the search from the right path. To account for noise, Horstein [11] introduced the probabilistic bisection algorithm (PBA), which updates a probability density according to Bayes' rule at each iteration such that the posterior probability density reflects one's current belief about the location of X^* . After the n th iteration, the median X_n of the resulting posterior density f_n provides a new estimate of the point X^* and is used as the reference point for the $(n + 1)$ th query to the oracle.

Examples of PBA application include the following:

- (i) Transmission over a noisy channel with noiseless feedback [11]: a real number $X^* \in [0, 1]$ should be transmitted from a sender to a receiver. Only one bit of information (0's or 1's) can be sent at each iteration, and the signal is sometimes wrong due to corruption by noise. In addition, a noiseless feedback loop informs the sender of what has been recorded by the receiver after each iteration. In this setting, the PBA can be used to efficiently transmit the number X^* .
- (ii) Boundary detection with an airborne radar [4]: an airplane equipped with a scanner flies over a predetermined geographical area several times to locate an edge, such as a coast line. At each pass-over, the scanner receives an input

*Received by the editors January 10, 2012; accepted for publication (in revised form) February 21, 2013; published electronically May 28, 2013.

<http://www.siam.org/journals/sicon/51-3/86189.html>

[†]School of Operations Research and Information Engineering, Cornell University, Ithaca, NY 14850 (rw339@cornell.edu, pf98@cornell.edu, sgh9@cornell.edu). The first and second authors were partially supported by the Air Force Office of Scientific Research Young Investigator Program under contract FA9550-11-1-0083. The third author was partially supported by National Science Foundation grants CMMI-0800688 and CMMI-1200315.

as to whether the scanned point is water surface or solid ground, but the signal can be wrong. The PBA can be used to determine which point should be scanned at each time so that a good estimate of the edge can be obtained.

- (iii) Stochastic root-finding [21]: stochastic root-finding algorithms aim to solve the equation $g(x) = 0$ for some unknown function g that can only be observed with noise. A one-dimensional stylized version of this problem, where g is a step function with a single jump at some point X^* , can be solved efficiently using the PBA.
- (iv) Zone-detection on a hard disk [22]: a hard disk stores each block of data in one of the disk's several zones, which have different transfer rates. The performance of a file system can be improved by accounting for such differences explicitly, but in order to do so, one must be able to identify where each zone begins and ends on the disk. Reading a small collection of data at any location on the disk provides a noisy observation of the transfer rate and thus the zone identity of that section. The PBA can then efficiently determine the exact zone borders.

Discretized versions of the PBA, which divide the domain $[0, 1]$ into a finite number of intervals, have been studied extensively [3, 20, 19, 8, 13, 1, 2, 4, 5, 17, 18]. However, very little is known about the original PBA with continuous search space $[0, 1]$. Castro and Nowak [4] conclude the following in their review article: "The Probabilistic Bisection Algorithm seems to work extremely well in practice, but it is hard to analyze and there are few theoretical guarantees for it, especially pertaining error rates of convergence."

In this paper we provide such convergence guarantees for the PBA. The main result shows that the expected absolute residuals $\mathbb{E}[|X^* - X_n|]$ converge to 0 at least at a geometric rate; i.e., there exists a constant $c > 1$ such that $\mathbb{E}[|X^* - X_n|] = o(c^{-n})$.¹ This implies that the rate of convergence of the bisection search with noisy responses is faster than any polynomial rate and is hence comparable to the rate of the noise-free bisection search, which is $O(2^{-n})$. Our main result is shown in the Bayesian setting, where X^* is modeled as an absolutely continuous random variable with known probability density f_0 . Since we are considering residuals under the expectation operator, our result provides an average-case performance guarantee for the PBA. A consequence of this main result is that the PBA is a consistent method for locating X^* . This means that the sequence $(X_n)_n$ generated by the PBA converges almost surely to X^* .

The most popular discretized version of the PBA, first introduced in Burnashev and Zigangirov [3], is called the BZ algorithm. The algorithm splits the search domain $[0, 1]$ into a finite number of intervals and aims to locate the interval that contains the point X^* . It is known [3] that the BZ algorithm converges geometrically in the number of points queried for the probabilistic setting considered in this paper. The current paper confirms that a similar rate of convergence holds for the original PBA (without discretization) and effectively closes a gap between the theoretical understanding of the original continuous-space algorithm and that of the corresponding discrete-space version. Although the PBA and the BZ algorithm are conceptually similar, the proof techniques used to analyze the PBA are quite different from the proof techniques usually used to study the BZ algorithm. Such new proof techniques become necessary because the BZ algorithm only samples at breakpoints of the predefined intervals,

¹ $f(x) = o(g(x))$ means that $\lim_{x \rightarrow \infty} |f(x)/g(x)| = 0$, and $f(x) = O(g(x))$ means that $\limsup_{x \rightarrow \infty} |f(x)/g(x)| < \infty$.

whereas the PBA can sample on the whole domain $[0, 1]$.

There are two reasons for preferring the PBA over the BZ algorithm. First, the PBA is a consistent algorithm, in the sense that it maintains a best estimate X_n of the sought-after point X^* , and X_n converges to X^* almost surely as $n \rightarrow \infty$ (see Corollary 5.10). The BZ algorithm, on the other hand, requires a prespecified precision (the discretization grid), beyond which no better accuracy can be expected. Since the sequence of estimates X_n does not converge to X^* almost surely, the BZ algorithm is not consistent. While one can specify any strictly positive precision, this precision must be specified in advance, which can be inconvenient. Although it might be possible to modify the BZ algorithm to make it consistent, for example, by refining the discretization grid during a run, no such extension has been considered in the literature to the best of our knowledge. The second reason for preferring the PBA to the BZ algorithm is that its implementation is easier (see also [4] and [5]). For example, at each step the BZ algorithm requires an additional coin flip to decide which point on the discrete grid should be queried next. Such “splitting” between the discretization points is not necessary for the PBA.

While there are reasons for preferring the PBA over the BZ algorithm, and the current paper closes a gap between the theoretical understanding of the two algorithms, there are probabilistic settings beyond the scope of this paper in which the BZ algorithm has been analyzed but the PBA has not. For example, the BZ algorithm has been analyzed by [5] when the probability of receiving a correct response from the oracle depends on the query point. The analysis of the original PBA for such settings is still an open research direction.

In addition to the main convergence results, we show that the PBA is optimal in reducing the expected posterior entropy. This result has been proven recently in [12] using concepts from information theory, in particular, the mutual information of the responses and X^* . In this paper, we adopt a more direct approach, showing that the PBA minimizes expected posterior entropy using fewer concepts from information theory. To do so, we formulate a dynamic program corresponding to the objective of expected posterior entropy, and we solve this dynamic program analytically.

The outline of the paper is as follows. Section 2 introduces the exact problem statement. Section 3 provides the updating mechanism and the PBA. Section 4 shows optimality in terms of minimizing expected posterior entropy. Section 5 presents the main result of the paper. Section 6 summarizes and discusses possible future research directions. The main proofs are given in section 5. All other proofs are given in the appendix.

Throughout the paper we use the following standard abbreviations: iid for independent and identically distributed, pdf for probability density function, and cdf for cumulative distribution function. We also use the abbreviations PBA for probabilistic bisection algorithm and DP for dynamic program. Furthermore $\mathbb{1}\{\cdot\}$ denotes the indicator function, which is 1 if the argument is true and 0 otherwise.

2. Problem statement. Let X^* be a unique but *unknown* point in $[0, 1]$. At time $n = 0$ (the beginning of the experiment) we do not know which exact value X^* has attained, but we assume it is a realization of an absolutely continuous random variable with density f_0 . The density f_0 has domain $[0, 1]$ and is known. The oracle, on the other hand, knows the exact realization of X^* . At each iteration we are allowed to query the oracle as to whether X^* is to the left or to the right of a point x . We denote by X_n the queried point at iteration n , which can depend on previously

received replies of the oracle. The oracle answers this question with

$$Z_n|X^*, X_n = (\mathbb{1}\{X^* \geq X_n\} - \mathbb{1}\{X^* < X_n\})(2Q_n - 1),$$

where $(Q_n)_n$ is a sequence of iid Bernoulli(p) random variables and p is a fixed constant in $(1/2, 1)$. We often simply write $Z_n(X_n)$ instead of $Z_n|X^*, X_n$. By querying the oracle at X_n we receive the answer $Z_n(X_n) \in \{-1, +1\}$. The answer $Z_n(X_n) = +1$ indicates that X^* is to the right of X_n , whereas the answer $Z_n(X_n) = -1$ indicates that X^* is to the left of X_n . With probability $1 - p$ this indication is wrong. If we query exactly at the point X^* , i.e., $X_n = X^*$, then the oracle is more likely to indicate that X^* is to the right of X_n , but this does not affect our analysis because the event $\{X^* = X_n\}$ occurs in finite time with probability 0.

We assume that the parameter p , the probability of a correct answer, does not depend on X_n and is known. In many applications these are realistic assumptions since the reliability of a given device is known and does not depend on the queried point, e.g., the error probability of a scanner in an airplane or the error probability of transmitting a signal over a noisy channel. In other applications, such as stochastic root-finding, p is unknown and varies with x . See [21] for a detailed discussion of this setting.

The assumption $p \in (1/2, 1)$ is without loss of generality: if $p \in (0, 1/2)$, then we can consider $-Z_n(X_n)$ as the reply from the oracle; the case $p \in \{0, 1\}$ corresponds to the noise-free bisection search, and convergence properties are well known; and the case $p = 1/2$ corresponds to a noninformative oracle; i.e., querying the oracle does not improve our knowledge about X^* . Hence the only interesting case is $p \in (1/2, 1)$.

The task is to select points X_0, X_1, X_2, \dots at which the oracle should be queried in order to learn about the location of X^* as efficiently as possible. The PBA is a fully sequential method, meaning that the decision of where to sample at time n depends on the information available at time n . Formally, $(X_n)_n$ is a predictable stochastic process with respect to the filtration generated by $(X_n)_n$ and $(Z_n(X_n))_n$, i.e., $X_n \in \mathcal{F}_{n-1} := \sigma(X_m, Z_m(X_m) : 0 \leq m \leq n-1)$ for $n \in \mathbb{N}$. Furthermore, at each iteration a current best estimate \hat{X}_n of X^* should be provided. This estimate \hat{X}_n is also an \mathcal{F}_{n-1} -measurable random variable. Finally, the considered performance measure is the expected L^1 -loss $\mathbb{E}[|X^* - \hat{X}_n|]$ as a function of time n .

3. The probabilistic bisection algorithm. The PBA aims to locate X^* for the above setting. In order to learn about X^* , the density f_0 is updated in each step according to Bayes' rule. After n iterations the posterior density f_n is the conditional distribution of X^* given the query history $(X_j)_{j=0}^{n-1}$ and replies $(Z_j(X_j))_{j=1}^{n-1}$. First, we give the updating process after measuring at a generic point $x \in [0, 1]$, and then we introduce the PBA, which provides a tool for choosing the point X_n at each iteration.

3.1. Updating process. At time $n = 0$ our knowledge about X^* is reflected by the density f_0 . If we have no prior knowledge of X^* , then a natural choice of f_0 is the uniform $U[0, 1]$ distribution, i.e., $f_0(y) = \mathbb{1}\{y \in [0, 1]\}$. After n iterations, we can measure at any point x in the interior of f_n and receive the noisy response $Z_n(x)$. The density f_n can then be updated using Bayes' rule as follows.

LEMMA 3.1. *The domain of the prior density function f_0 is $[0, 1]$. The sequence of posterior densities $(f_n)_n$ is given by the following iterative process, where x is a*

point in the interior of f_n at which the oracle is called at step n :

$$(3.1) \quad \text{If } Z_n(x) = +1, \text{ then } f_{n+1}(y) = \begin{cases} \gamma(x)^{-1} p f_n(y) & \text{if } y \geq x, \\ \gamma(x)^{-1} (1-p) f_n(y) & \text{if } y < x; \end{cases}$$

$$(3.2) \quad \text{if } Z_n(x) = -1, \text{ then } f_{n+1}(y) = \begin{cases} (1-\gamma(x))^{-1} (1-p) f_n(y) & \text{if } y \geq x, \\ (1-\gamma(x))^{-1} p f_n(y) & \text{if } y < x, \end{cases}$$

where $\gamma(x) = \mathbb{P}(Z_n(x) = +1 | \mathcal{F}_{n-1}) = (1 - F_n(x))p + F_n(x)(1 - p)$ and F_n denotes the cdf of the density f_n .

The proof of Lemma 3.1 is given in the appendix.

The updating of the density is very natural: querying at the point x divides the posterior distribution into two regions. The posterior probability mass in the region where X^* is believed to be, as indicated by the noisy response of the oracle, is increased, and the probability mass in the other region, where X^* is believed not to be, is decreased. We will see that if we always measure at the median of f_n (as the PBA does), the probability mass will eventually concentrate at the point X^* .

The updating mechanism in Lemma 3.1 defines a stochastic sequence of pdfs $(f_n)_n$ and hence a sequence of random probability measures on $[0, 1]$. We denote the corresponding probability measure as $\mathbb{P}_n(\cdot)$, i.e., $\mathbb{P}_n(B) := \int_B f_n(y) dy$ for $B \in \mathcal{B}([0, 1])$, where $\mathcal{B}([0, 1])$ is the Borel σ -field on $[0, 1]$, and denote the expectation under this probability measure as $\mathbb{E}_n[\cdot]$. To prove our main results we study the stochastic sequence of probability measures $\mathbb{P}_n(\cdot)$ induced by the PBA. The random measure $\mathbb{P}_n(\cdot)$ and the random density f_n are \mathcal{F}_{n-1} -measurable mappings for $n \in \mathbb{N}$.

3.2. The probabilistic bisection algorithm. The PBA uses the updating mechanism described in Lemma 3.1 to learn about X^* . The key rule of the policy is that it always measures at the median of the current posterior distribution function. More specifically, the PBA works as follows:

1. Choose a prior density function f_0 that is positive on $[0, 1]$.
2. For $n = 0$ to $N - 1$, do the following:
 - (a) Calculate the next measurement point, $X_n = F_n^{-1}(1/2)$, where F_n is the cdf of f_n . Note that X_n is uniquely defined since f_n is a density with domain $[0, 1]$.
 - (b) Query the oracle at the point X_n to obtain the random variable $Z_n(X_n) \in \{-1, +1\}$.
 - (c) Update the density f_{n+1} using inputs p , f_n , X_n , and $Z_n(X_n)$ and the updating formulas (3.1) and (3.2). (Since X_n is the median of f_n , the multiplicative constant in the updating is 2 for both cases.)
3. Return $\hat{X}_N = F_N^{-1}(1/2)$ as the current best estimate of X^* .

Note that X_n can only assume a finite number of values for any finite $n \in \mathbb{N}$ (a maximum of 2^n possible values), X^* is a random variable with density f_0 , and hence the event $\{X^* = X_n\}$ has probability 0 for every finite time $n \in \mathbb{N}$. The final estimate \hat{X}_N is the median of f_N , which is optimal if we want to minimize the expected absolute error under the probability measure $\mathbb{P}_N(\cdot)$; i.e., \hat{X}_N is the solution to the optimization problem $\min_x \mathbb{E}_N[|X^* - x|]$. In this case, each median X_n is the best estimate of X^* after n calls to the oracle. This allows us to drop the notation \hat{X}_N and to simply focus on the sequence of medians $(X_n)_n$ from here onward. At any time during a sample run a best estimate of X^* is maintained and it is expected that X_n approaches X^* the more queries we ask the oracle. In Corollary 5.10 we confirm this by showing that $X_n \rightarrow X^*$ as $n \rightarrow \infty$ almost surely, a property that is in general not satisfied by discretized versions of the PBA.

Figure 3.1 shows a sample path of the density f_n after $n = 0, 1, 2, 3, 50, 100$ calls to the oracle where $p = 0.6$ and $X^* = 0.372$. In this example, the prior density f_0 is that of a $U[0, 1]$ random variable, i.e., $f_0(y) = \mathbb{1}\{y \in [0, 1]\}$. The piecewise constant line depicts the posterior density f_n . In each subfigure the point X^* is shown at the top and X_n is shown on the x -axis. Above every plot the (noisy) answer of the oracle is given. The posterior density appears to converge to a point mass at X^* , and Corollary 5.9 confirms that this happens.

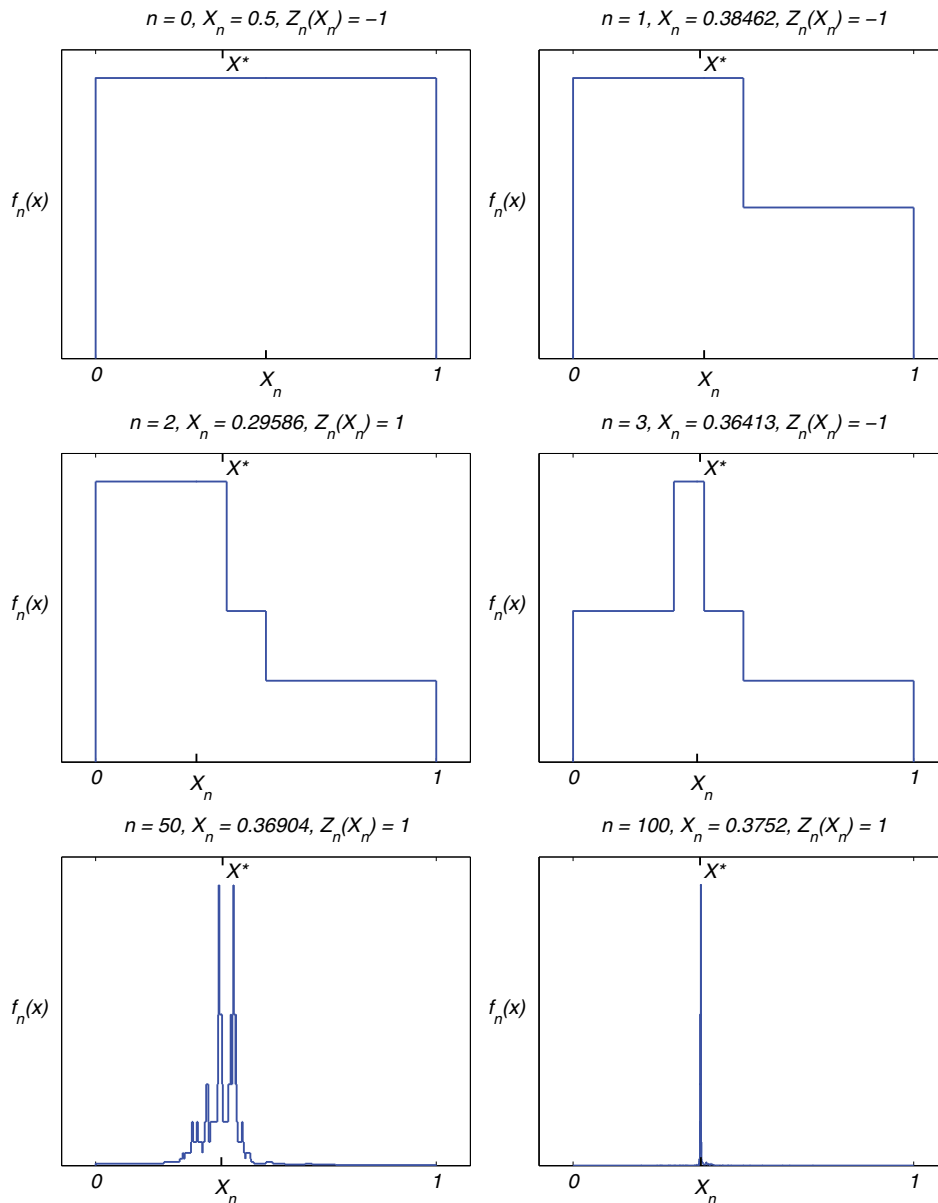


FIG. 3.1. The density f_n at time points $n = 0, 1, 2, 3, 50, 100$ on a sample path. We chose f_0 to be the uniform density over $[0, 1]$ in this example.

4. Optimality in reducing the expected posterior entropy. In this section, we show that the PBA is optimal in reducing the expected posterior entropy. This result has recently been proven in [12] using the mutual information of the responses $(Z_n(X_n))_n$ and X^* . In the appendix, we provide a different and more direct proof of this result that borrows fewer concepts from information theory. This proof relies solely on the dynamic programming principle.

The optimality result, stated in Theorem 4.1, uses the entropy to measure the information content of the density f_n . For a random variable Y with density f the entropy is defined as $H(f) := \mathbb{E}[-\log_2 f(Y)]$. The entropy is the predominant measure of uncertainty in information theory; see, for example, [6]. Using this measure of uncertainty and given a fixed simulation budget $N \in \mathbb{N}$, the optimality analysis seeks a policy π that minimizes the expected entropy of the posterior distribution at time N . Here, a policy refers to the allocation rule of the measurements X_0, \dots, X_N , where X_n has to be \mathcal{F}_{n-1} -measurable. A generic policy is denoted π , and the space of all possible policies is denoted Π . This optimization problem can be solved using a dynamic programming approach. The value function of the DP for fixed $N \in \mathbb{N}$ is defined by

$$(4.1) \quad V_n(f_n) := \inf_{\pi \in \Pi} \mathbb{E}^\pi[H(f_N)|f_n] \quad \text{for } n = 0, 1, \dots, N.$$

Any policy π induces, together with the input density f_0 and the parameter p , a distribution on $(X_j, Z_j(X_j))_{j=0}^{N-1}$ and through it a distribution on the sequence of pdfs $(f_j)_{j=0}^N$. It is under this distribution that \mathbb{E}^π is taken, and any policy π^* attaining the infimum is called optimal, that is, $\mathbb{E}^{\pi^*}[H(f_N)|f_0] = \inf_{\pi \in \Pi} \mathbb{E}^\pi[H(f_N)|f_0]$.

The value function (4.1) satisfies Bellman’s recursion,

$$(4.2) \quad V_n(f_n) = \inf_{\pi \in \Pi} \mathbb{E}^\pi[V_{n+1}(f_{n+1})|f_n] = \inf_{x \in [0,1]} \mathbb{E}[V_{n+1}(f_{n+1})|X_n = x, f_n],$$

where the last equation follows from the fact that the control of a policy $\pi \in \Pi$ is the point at which to query the oracle. The DP formulated in (4.1) can be solved explicitly.

THEOREM 4.1. *For $N \in \mathbb{N}$, the PBA, which always measures at the median of f_n for $n = 0, \dots, N - 1$, minimizes the expected entropy of the density f_N . Furthermore, the expected posterior entropy at time N using the PBA is*

$$(4.3) \quad \begin{aligned} V_n(f_n) &= \mathbb{E}[H(f_N)|f_n] \\ &= H(f_n) - (N - n)(1 + p \log_2 p + (1 - p) \log_2(1 - p)) \end{aligned}$$

for $n = 0, \dots, N$.

The key step in the proof of Theorem 4.1 is the analysis of the knowledge-gradient policy for the DP formulated in (4.1). A knowledge-gradient policy is a policy that acts optimally if there is only *one* measurement remaining, i.e., when $n = N - 1$. See [9] for more details on knowledge-gradient policies. For this knowledge-gradient policy the value attained by the infimum is equal to the entropy of f_n minus an additional amount which may be interpreted as the maximum information content of a single measurement. The fact that this amount does not depend on f_n is important in proving that the knowledge-gradient policy is in fact the optimal policy in general when more than just one measurement is remaining. The next proposition shows that the PBA is indeed the knowledge-gradient policy for the problem stated in (4.1).

PROPOSITION 4.2. *For any $N \in \mathbb{N}$,*

$$\begin{aligned} & \inf_{x \in [0,1]} \mathbb{E}[V_N(f_N) | X_{N-1} = x, f_{N-1}] \\ &= \inf_{x \in [0,1]} \mathbb{E}[H(f_N) | X_{N-1} = x, f_{N-1}] \\ &= H(f_{N-1}) - p \log_2 p - (1-p) \log_2(1-p) - 1, \end{aligned}$$

and the infimum is achieved by choosing X_{N-1} to be the median of f_{N-1} .

We provide proofs of Theorem 4.1 and Proposition 4.2 in the appendix.

5. Main result: Geometric rate of L^1 -convergence. In this section we present and prove the main result of this paper, which is that the expected absolute residuals of the PBA converge to 0 at a rate of $o(c^{-n})$ for some $c > 1$. This, in particular, implies that the asymptotic rate of convergence is faster than any polynomial rate and is comparable to the rate of convergence of the noise-free bisection algorithm which has rate $O(2^{-n})$. Such a geometric rate of convergence is known to hold for discretized versions (BZ algorithm) of the PBA; see [3, 4, 5]. But, to the best of our knowledge, it is a new result for the original PBA.

THEOREM 5.1. *There exists a constant $c(p) > 1$ such that $\mathbb{E}[|X^* - X_n|] = o(c(p)^{-n})$, where $(X_n)_n$ is the sequence of query points generated by the PBA.*

Before developing the proof of Theorem 5.1, we first discuss the constant $c(p)$, introduce some simplified notation, and provide a sketch of the proof.

The constant $c(p)$ can be any fixed value in the open interval $(1, C(p))$, where $\ln(C(p))$ is the smaller solution to the quadratic equation (5.2) given in Lemma 5.2. For the most part, it suffices to know that $C(p)$ is a constant depending only on the parameter p and $C(p) > 1$. From the rate of convergence of the noise-free bisection algorithm we know that $C(p) \leq 2$. In fact, $C(p)$ is often much smaller than 2 and is usually quite close to 1. This, however, does not necessarily imply that the rate of convergence of the PBA is much slower (in terms of the constant $c(p)$) than the rate of convergence of the noise-free bisection algorithm since our result provides only a lower bound on the rate of convergence; i.e., we show that $\lim_{n \rightarrow \infty} c(p)^n \mathbb{E}[|X^* - X_n|] = 0$. We leave for future work the problem of identifying the exact rate of convergence.

We now introduce some simplified notation. Define $D(p) = 1/2(\ln(2p) + \ln(2(1-p)))$, which is a constant that depends only on p . Note that $D(p) < 0$, since $p \in (1/2, 1)$. From now on we will often simply write c , C , and D when the context allows and keep in mind that all these constants depend only on the parameter p .

Sketch of proof of Theorem 5.1. The proof of Theorem 5.1 consists of two major steps. Each is formulated in the next subsection as a separate proposition. In Proposition 5.6 we show that the stochastic process $(c^n \mathbb{E}_n[|X^* - X_n|])_n$ converges to 0 in probability. We then show the uniform integrability of this process in Proposition 5.7, and Theorem 5.1 follows from the fact that a sequence of uniformly integrable random variables converges in L^1 if and only if it converges in probability.

The key to proving these two propositions is to analyze the stochastic process $\mathbb{E}_n[|X^* - X_n|]$. We now give an intuitive outline of why this process converges at a geometric rate. All the arguments are made precise in the next subsection.

Using integration by parts, it holds that

$$\begin{aligned} \mathbb{E}_n[|X^* - X_n|] &= \int_0^1 \mathbb{P}_n(|X^* - X_n| > h) dh \\ &\leq h + \mathbb{P}_n(|X^* - X_n| > h) \end{aligned}$$

for any $h \in (0, 1)$. The inequality holds since $\mathbb{P}_n(|X^* - X_n| > h) \leq 1$ and is decreasing in h . It is then enough to show that the process $\mathbb{P}_n(|X^* - X_n| > h)$ converges to 0 at a geometric rate and consider the case $h \rightarrow 0$. For now fix an $h \in (0, 1)$. At time n there exists an integer K_n such that $X_n \in [(K_n - 1)h, K_n h)$ and

$$\mathbb{P}_n(|X^* - X_n| > h) \leq \mathbb{P}_n(X^* \in [0, (K_n - 1)h)) + \mathbb{P}_n(X^* \in [K_n h, 1]).$$

We then focus on the process $(A_n)_n$, where $A_n := \mathbb{P}_n(X^* \in [0, (K_n - 1)h))$ (the analysis of the process $\mathbb{P}_n(X^* \in [K_n h, 1])$ follows analogously). After querying the oracle at X_n , the quantity A_n is multiplied by either $2p$ or $2(1 - p)$. Also, since X_n is the median of f_n , either multiplication happens with probability $1/2$. If we (for now) ignore the fact that K_n depends on n , then A_n behaves like a geometric random walk with drift e^D and hence converges to 0 at a geometric rate. This is the basic argument of why the geometric rate of convergence holds. Most of the proof is then devoted to the fact that K_n depends on n and is a stochastic process itself. It turns out that A_n is not a true geometric random walk (which can already be seen since A_n is always smaller than $1/2$), but that A_n can be dominated by a collection of dependent geometric random walks, and each of these random walks has drift e^D . Using this dominating argument and results from random walk theory, we can then show that the geometric rate of convergence indeed holds for $\mathbb{P}(|X^* - X_n| > h)$. By letting $h \rightarrow 0$ this geometric rate also holds for $\mathbb{E}_n[|X^* - X_n|]$, and for $\mathbb{E}[|X^* - X_n|]$ by applying the tower property of conditional expectations.

5.1. Proof of the main result. We start with a lemma which is an application of random walk theory. This lemma defines the constant C and will also be useful for later proofs.

LEMMA 5.2. *Let $(R_n)_n$ be a random walk with starting point $R_0 \leq \ln(1/2)$ and iid increments $(\psi_n)_n$, i.e., $R_n = R_0 + \sum_{j=1}^n \psi_j$, and let $\mathbb{P}(\psi_j = \ln(2p)) = \mathbb{P}(\psi_j = \ln(2(1 - p))) = 1/2$. Then*

$$(5.1) \quad \mathbb{P}(e^{R_n} > C^{-n}/2) \leq C^{-2n}$$

for all $n \in \mathbb{N}$. Here, $C = e^{\tilde{u}}$, where \tilde{u} is the smaller solution to

$$(5.2) \quad \left(\frac{u + D}{\ln(2p) - \ln(2(1 - p))} \right)^2 - u = 0.$$

Furthermore, $\tilde{u} > 0$.

Equation (5.2) is a quadratic equation, and it is possible to write down an explicit formula for C . However, the explicit form of C is cumbersome and not informative and hence is omitted. The proof of Lemma 5.2 is given in the appendix.

The next result, which studies the stochastic process $A_n := \mathbb{P}_n(X^* \in [0, a))$ for some $a \in [0, 1]$, is a first key ingredient in showing the geometric rate of convergence. Define $a \wedge b = \min\{a, b\}$.

PROPOSITION 5.3. *Let C be the constant defined in Lemma 5.2. For $a \in [0, 1]$ define $A_n := \mathbb{P}_n(X^* \in [0, a)) = \int_0^a f_n(y) dy$. Then*

$$\mathbb{P}(A_n \wedge (1 - A_n) > C^{-n}/2) \leq C^{-2n}$$

for all $n \in \mathbb{N}$.

Proof. The claim holds trivially for $a = 0$ or $a = 1$ since for all $n \in \mathbb{N}$ the probability measure $\mathbb{P}_n(\cdot)$ has a density.

Now fix an arbitrary $a \in (0, 1)$ and consider the stochastic process $(A_n)_n$. If $A_n \leq 1/2$, then $X_n \geq a$, and A_n will be multiplied by either $2p$ or $2(1 - p)$ in the next iteration, behaving like an iteration of a geometric random walk. If, on the other hand, $A_n > 1/2$, then $X_n \in [0, a)$ and A_n does not behave like an iteration of a geometric random walk anymore, but $(1 - A_n)$ does. We next make this argument precise. To simplify notation we take logarithms and consider the process $(\ln(A_n))_n$. The stochastic driver of this process is the sequence of responses from the oracle $(Z_n(X_n))_n$. If we condition on the available information up to time $n - 1$, then, by Lemma 3.1, $\gamma(X_n) = \mathbb{P}(Z_n(X_n) = +1 | \mathcal{F}_{n-1}) = 1/2$ for the PBA. Moreover, the only random source that drives the stochastic process $(Z_n(X_n) | \mathcal{F}_n)_n$ is the sequence $(Q_n)_n$, a sequence of iid Bernoulli(p) random variables that is used by the oracle to provide noisy responses, and hence the sequence $(Z_n(X_n) | \mathcal{F}_{n-1})_n$ is itself a sequence of independent random variables. At time n , the random variable $\ln(A_{n+1}) | \mathcal{F}_n$ can be constructed as follows: if $\ln(A_n) \leq \ln(1/2)$, then

$$\ln(A_{n+1}) | \mathcal{F}_n = \ln(A_n) + \begin{cases} \ln(2(1 - p)) & \text{if } Z_n(X_n) = +1, \\ \ln(2p) & \text{if } Z_n(X_n) = -1, \end{cases}$$

and if $\ln(A_n) > \ln(1/2)$, then

$$\ln(1 - A_{n+1}) | \mathcal{F}_n = \ln(1 - A_n) + \begin{cases} \ln(2p) & \text{if } Z_n(X_n) = +1, \\ \ln(2(1 - p)) & \text{if } Z_n(X_n) = -1. \end{cases}$$

Now consider the process $M_n := \ln(A_n) \wedge \ln(1 - A_n)$. The only times the dynamics of $(M_n)_n$ are different from a random walk are when the process $(M_n)_n$ crosses the boundary $\ln(1/2)$, i.e., at times when there is a switch from the process $(\ln(A_n))_n$ to the process $(\ln(1 - A_n))_n$ in the definition of $(M_n)_n$. To overcome this difficulty we construct a true random walk $(S_n)_n$ that is *coupled* with $(M_n)_n$ and dominates $(M_n)_n$.

We first define the coupling sequence

$$W_n = \begin{cases} Z_n(X_n) & \text{if } \ln(A_n) > \ln(1/2), \\ -Z_n(X_n) & \text{if } \ln(A_n) \leq \ln(1/2) \end{cases}$$

and then the process

$$S_{n+1} = S_n + \begin{cases} \ln(2p) & \text{if } W_n = +1, \\ \ln(2(1 - p)) & \text{if } W_n = -1 \end{cases}$$

for $n \in \mathbb{N}$ and starting point $S_0 = M_0$. The process $(S_n)_n$ is a random walk with iid increments $(\xi_n)_n$ and $\mathbb{P}(\xi_n = \ln(2p)) = \mathbb{P}(\xi_n = \ln(2(1 - p))) = 1/2$.

The processes $(M_n)_n$ and $(S_n)_n$ have the same starting point and are driven by the same sequence of random variables $(W_n)_n$. Assume that $M_0 = \ln(1 - A_0)$, and define $\tau := \inf \{n | 1 - A_n \geq 1/2\}$ (if $M_0 = \ln(A_0)$, then the definition of τ and the following arguments can be adapted accordingly). For $n < \tau$ it holds that $M_n = S_n$. At time τ the processes $\ln(1 - A_n)_n$ and $(S_n)_n$ increase by $\ln(2p)$. On the other hand, the process $(M_n)_n$ switches from being defined by $\ln(1 - A_n)$ to being defined by $\ln(A_n)$ and may increase or decrease, i.e.,

$$M_\tau - M_{\tau-1} = \ln(A_\tau) - M_{\tau-1} \leq \ln(1 - A_\tau) - M_{\tau-1} = S_\tau - S_{\tau-1},$$

and hence $M_\tau \leq S_\tau$. (See Figure 5.1.) After time τ this argument carries over in the following sense: each time $(S_n)_n$ decreases by $\ln(2(1 - p))$, $(S_n)_n$ also decreases




FIG. 5.1. The process $(S_n)_n$ (circles) dominates the process $(M_n)_n$ (squares) for all $n \in \mathbb{N}$. The process $(S_n)_n$ is a random walk with negative drift, so by the law of large numbers $S_n \rightarrow -\infty$ almost surely as $n \rightarrow \infty$, as indicated by the arrow on the right side of the figure. (Both processes are defined in discrete time. We draw a dashed line between time steps for better visibility.)

by $\ln(2(1-p))$. However, when $(S_n)_n$ increases by $\ln(2p)$, then $(M_n)_n$ increases by a quantity less than or equal to $\ln(2p)$ (the increase might also be negative). It follows that $M_n \leq S_n$, and hence $A_n \wedge (1 - A_n) \leq e^{S_n}$ for all $n \in \mathbb{N}$. Then

$$\begin{aligned} \mathbb{P}(A_n \wedge (1 - A_n) > C^{-n}/2) &\leq \mathbb{P}(e^{S_n} > C^{-n}/2) \\ &\leq C^{-2n}, \end{aligned}$$

where the last inequality follows from Lemma 5.2 since $(S_n)_n$ is a random walk as considered in that lemma. \square

We can now use the previous result to bound the probability of observing a large posterior probability mass outside a small neighborhood of the current best estimate X_n .

PROPOSITION 5.4. *Let C be the constant defined in Lemma 5.2. Then*

$$\mathbb{P}(\mathbb{P}_n(|X^* - X_n| > h) > C^{-n}) \leq h^{-1}C^{-2n}$$

for all $h \in (0, 1)$ and $n \in \mathbb{N}$.

Proof. Fix an arbitrary $h \in (0, 1)$ and denote $\bar{K} = \lfloor h^{-1} \rfloor$. Define intervals $I(k) := [(k-1)h, kh)$ for $k = 1, \dots, \bar{K}$ and $I(\bar{K} + 1) := [\bar{K}h, 1]$. These $\bar{K} + 1$ intervals are pairwise disjoint and cover the domain $[0, 1]$. Further define the stochastic processes

$$A_n(k) := \mathbb{P}_n \left(X^* \in \bigcup_{j=1}^k I(j) \right)$$

for $k = 1, \dots, \bar{K} + 1$ and the trivial process $A_n(0) = 0$ for all $n \in \mathbb{N}$.

At time $n \in \mathbb{N}$ let K_n be the index such that $X_n \in I(K_n)$. Then

$$\begin{aligned} \mathbb{P}_n(|X^* - X_n| > h) &\leq \mathbb{P}_n(X^* \in [0, (K_n - 1)h)) + \mathbb{P}_n(X^* \in [K_n h, 1]) \\ &= A_n(K_n - 1) + (1 - A_n(K_n)) \\ &= [A_n(K_n - 1) \wedge (1 - A_n(K_n - 1))] \\ &\quad + [A_n(K_n) \wedge (1 - A_n(K_n))], \end{aligned}$$

where the last equation holds since $X_n \in I(K_n)$ implies $A_n(K_n - 1) \leq 1/2$ and $1 - A_n(K_n) \leq 1/2$. The index K_n is a random variable taking values in $\{1, \dots, \bar{K} + 1\}$, and hence

$$\begin{aligned} \mathbb{P}_n(|X^* - X_n| > h) &\leq \max_{k \in \{1, \dots, \bar{K} + 1\}} ([A_n(k - 1) \wedge (1 - A_n(k - 1))] \\ &\quad + [A_n(k) \wedge (1 - A_n(k))]) \\ &\leq \max_{k \in \{1, \dots, \bar{K}\}} 2[A_n(k) \wedge (1 - A_n(k))], \end{aligned}$$

since $A_n(0) = 0$ and $A_n(\bar{K} + 1) = 1$ for all $n \in \mathbb{N}$. Then

$$\begin{aligned} \mathbb{P}(\mathbb{P}_n(|X^* - X_n| > h) > C^{-n}) &\leq \mathbb{P}\left(\max_{k \in \{1, \dots, \bar{K}\}} 2[A_n(k) \wedge (1 - A_n(k))] > C^{-n}\right) \\ &\leq \mathbb{P}\left(\max_{k \in \{1, \dots, \bar{K}\}} [A_n(k) \wedge (1 - A_n(k))] > C^{-n}/2\right) \\ &= \mathbb{P}\left(\bigcup_{k=1}^{\bar{K}} \{[A_n(k) \wedge (1 - A_n(k))] > C^{-n}/2\}\right) \\ &\leq \sum_{k=1}^{\bar{K}} \mathbb{P}([A_n(k) \wedge (1 - A_n(k))] > C^{-n}/2) \\ &\leq \bar{K}C^{-2n}. \end{aligned}$$

The last inequality follows by Proposition 5.3 since the processes $(A_n(k))_n$ are exactly of the form required for that proposition. Note that $\bar{K} = \lfloor h^{-1} \rfloor \leq h^{-1}$, and the claim follows. \square

The next proposition provides an upper bound on $\mathbb{P}(c^n \mathbb{E}_n[|X^* - X_n|] > \varepsilon)$ for $\varepsilon > 0$ and large n . This result is the last step before we can prove convergence in probability and uniform integrability of the stochastic process $(c^n \mathbb{E}_n[|X^* - X_n|])_n$, and is also interesting by itself. It provides a large deviation result for the stochastic process $c^n \mathbb{E}_n[|X^* - X_n|]$. In contrast to Theorem 5.1, it provides a finite time guarantee for large $n \in \mathbb{N}$, instead of an asymptotic convergence guarantee.

PROPOSITION 5.5. *Let C be the constant defined in Lemma 5.2, let $c \in (1, C)$, and let $\varepsilon > 0$. Then*

$$\mathbb{P}(c^n \mathbb{E}_n[|X^* - X_n|] > \varepsilon) \leq C^{-n} \text{ for } n \geq \max(0, \tilde{N}(\varepsilon, c, C)),$$

where

$$(5.3) \quad \tilde{N}(\varepsilon, c, C) = \frac{\ln(2/\varepsilon)}{\ln(C/c)}.$$

Proof. Fix $n \geq \max(0, \tilde{N}(\varepsilon, c, C))$. Consider

$$c^n \mathbb{E}_n[|X^* - X_n|] = c^n \int_0^1 \mathbb{P}_n(|X^* - X_n| > h) dh,$$

which follows from integration by parts of the right-hand side. The random function $\mathbb{P}_n(|X^* - X_n| > h)$ is nonincreasing in h and $\mathbb{P}_n(|X^* - X_n| > h) \leq 1$ for all $h \in (0, 1)$.

Then for any $h \in (0, 1)$,

$$\begin{aligned} c^n \mathbb{E}_n[|X^* - X_n|] &\leq c^n (h + (1 - h)\mathbb{P}_n(|X^* - X_n| > h)) \\ &\leq c^n (h + \mathbb{P}_n(|X^* - X_n| > h)). \end{aligned}$$

So we can choose $h = C^{-n} \in (0, 1)$ and get

$$c^n \mathbb{E}_n[|X^* - X_n|] \leq c^n (C^{-n} + \mathbb{P}_n(|X^* - X_n| > C^{-n})).$$

Note that on the event $\{\mathbb{P}_n(|X^* - X_n| > C^{-n}) \leq C^{-n}\}$,

$$\begin{aligned} c^n \mathbb{E}_n[|X^* - X_n|] &\leq c^n (2C^{-n}) \\ &= 2(c/C)^n \\ &\leq \varepsilon, \end{aligned}$$

where the last inequality follows since $n \geq \tilde{N}(\varepsilon, c, C)$.

Then $\mathbb{P}(c^n \mathbb{E}_n[|X^* - X_n|] \leq \varepsilon) \geq \mathbb{P}(\mathbb{P}_n(|X^* - X_n| > C^{-n}) \leq C^{-n})$ and

$$\begin{aligned} \mathbb{P}(c^n \mathbb{E}_n[|X^* - X_n|] > \varepsilon) &\leq \mathbb{P}(\mathbb{P}_n(|X^* - X_n| > C^{-n}) > C^{-n}) \\ &\leq C^{-n}, \end{aligned}$$

where the last step follows by Proposition 5.4. \square

Now we are ready to prove convergence in probability and uniform integrability of the process $(c^n \mathbb{E}[|X^* - X_n|])_n$ and finally prove Theorem 5.1.

PROPOSITION 5.6. *Let C be the constant defined in Lemma 5.2. Then $\mathbb{E}_n[|X^* - X_n|] = o_p(c^{-n})$ for all $c \in (1, C)$.²*

Proof. Choose arbitrary $c \in (1, C)$, which exists since $C > 1$. Fix $\varepsilon > 0$. Then, by Proposition 5.5, $\mathbb{P}(c^n \mathbb{E}_n[|X^* - X_n|] > \varepsilon) \leq 2C^{-n}$ for large n , i.e., for $n > \tilde{N}(\varepsilon, c, C)$. Thus,

$$\lim_{n \rightarrow \infty} \mathbb{P}(c^n \mathbb{E}_n[|X^* - X_n|] > \varepsilon) = 0,$$

which holds for any chosen $\varepsilon > 0$, and hence $(c^n \mathbb{E}_n[|X^* - X_n|])_n$ converges to 0 in probability. \square

PROPOSITION 5.7. *Let C be the constant defined in Lemma 5.2. Then the stochastic process $(c^n \mathbb{E}_n[|X^* - X_n|])_n$ is uniformly integrable for all $c \in (1, C)$.*

Proof. By definition, a sequence of random variables $(Y_n)_n$ is uniformly integrable if $\sup_{n \in \mathbb{N}} \mathbb{E}[|Y_n| \mathbb{1}_{\{|Y_n| > t\}}] \rightarrow 0$ as $t \rightarrow \infty$.

Choose arbitrary $c \in (1, C)$ and consider $\tilde{N}(1, c, C) = (\ln 2)/(\ln(C/c))$, which is strictly positive (the function $\tilde{N}(\varepsilon, c, C)$ is defined in Proposition 5.5). Note that $\tilde{N}(t, c, C) \leq \tilde{N}(1, c, C)$ for $t \geq 1$. Define $T(c, C) := c^{\tilde{N}(1, c, C)} > 1$ and consider arbitrary $t \geq T(c, C) > 1$. It follows that $\mathbb{P}(c^n \mathbb{E}_n[|X^* - X_n|] > t) = 0$ for $n \leq \tilde{N}(1, c, C)$, since $\mathbb{E}_n[|X^* - X_n|] \leq 1$ and $c^n \leq t$ for $n \leq \tilde{N}(1, c, C)$. By Proposition 5.5, $\mathbb{P}(c^n \mathbb{E}_n[|X^* - X_n|] > t) \leq C^{-n}$ for $n \geq \tilde{N}(1, c, C) \geq \tilde{N}(t, c, C)$. Hence $\mathbb{P}(c^n \mathbb{E}_n[|X^* - X_n|] > t) \leq C^{-n}$ for all $n \in \mathbb{N}$ and all $t > T(c, C)$. Using $\mathbb{E}_n[|X^* - X_n|] \leq 1$ shows

² $f(x) = o_p(g(x))$ means $f(x)/g(x) \rightarrow 0$ in probability as $x \rightarrow \infty$.

that for all $n \in \mathbb{N}$,

$$\begin{aligned} \mathbb{E} [c^n \mathbb{E}_n [|X^* - X_n|] \mathbb{1}_{\{c^n \mathbb{E}_n [|X^* - X_n|] > t\}}] &\leq c^n \mathbb{E} [\mathbb{1}_{\{c^n \mathbb{E}_n [|X^* - X_n|] > t\}}] \\ &= c^n \mathbb{E} [\mathbb{1}_{\{c^n > t\}} \mathbb{1}_{\{c^n \mathbb{E}_n [|X^* - X_n|] > t\}}] \\ &= c^n \mathbb{1}_{\{c^n > t\}} \mathbb{P}(c^n \mathbb{E}_n [|X^* - X_n|] > t) \\ &\leq c^n \mathbb{1}_{\{c^n > t\}} C^{-n} \\ &= (c/C)^n \mathbb{1}_{\{n > \log_c t\}}. \end{aligned}$$

Now we take on both sides the supremum over $n \in \mathbb{N}$:

$$\begin{aligned} \sup_{n \in \mathbb{N}} \mathbb{E} [c^n \mathbb{E}_n [|X^* - X_n|] \mathbb{1}_{\{c^n \mathbb{E}_n [|X^* - X_n|] > t\}}] &\leq \sup_{n \in \mathbb{N}} (c/C)^n \mathbb{1}_{\{n > \log_c t\}} \\ &= (c/C)^{\log_c t}, \end{aligned}$$

and uniform integrability follows by letting t go to $+\infty$. \square

Proof of Theorem 5.1. By Propositions 5.6 and 5.7 we can choose an arbitrary constant $c \in (1, C)$ such that the sequence $(c^n \mathbb{E}_n [|X^* - X_n|])_n$ converges to 0 in probability and is uniformly integrable. Then

$$\mathbb{E} [c^n \mathbb{E}_n [|X^* - X_n|]] \rightarrow 0 \text{ as } n \rightarrow \infty,$$

since convergence in probability and uniform integrability is a necessary and sufficient condition for convergence in L^1 ; see, for example, [7, Theorem 4.5.2]. Finally, by the tower property of conditional expectation, $c^n \mathbb{E} [|X^* - X_n|] = \mathbb{E} [c^n \mathbb{E}_n [|X^* - X_n|]]$ and hence $\mathbb{E} [|X^* - X_n|] = o(c^{-n})$. \square

5.2. Consistency and robustness. Almost immediate consequences of the preceding analysis are that the posterior absolute residuals converge to 0 almost surely and that the posterior density f_n converges to a point mass at X^* . Hence the PBA is a consistent method for locating X^* , a property that is in general not satisfied by discretized versions of the PBA.

THEOREM 5.8. $\mathbb{E}_n [|X^* - X_n|] \rightarrow 0$ almost surely as $n \rightarrow \infty$.

COROLLARY 5.9. With probability 1 the posterior distribution F_n converges weakly to a point mass at X^* , i.e., $\lim_{n \rightarrow \infty} F_n(x) = \mathbb{1}_{\{x \geq X^*\}}$ for all $x \neq X^*$ almost surely.

COROLLARY 5.10. The sequence of medians $(X_n)_n$ generated by the PBA converges to X^* almost surely, i.e., $\mathbb{P}(\lim_{n \rightarrow \infty} X_n = X^*) = 1$.

The proof of Theorem 5.8 is given in the appendix.

As a final remark, we show that in some cases the geometric rate still holds even if the density of the average-case performance measure is different from the density used in the updating process of the PBA. Suppose that the random variable X^* has a density g_0 on $[0, 1]$, and let $(X_n)_n$ be the sequence of medians generated by the PBA using some other initial prior density f_0 (f_0 has to be positive on $[0, 1]$). Then a sufficient condition for the geometric rate of convergence of the expected absolute residuals to still hold is that the likelihood ratio between g_0 and f_0 is bounded; that is, there exists a constant $L \in \mathbb{R}$ such that $g_0(x)/f_0(x) \leq L$ for all $x \in [0, 1]$. In this

case,

$$\begin{aligned} \mathbb{E}[|X^* - X_n|] &= \int_0^1 g_0(x) \mathbb{E}[|x - X_n| \mid X^* = x] dx \\ &= \int_0^1 f_0(x) \frac{g_0(x)}{f_0(x)} \mathbb{E}[|x - X_n| \mid X^* = x] dx \\ &\leq L \int_0^1 f_0(x) \mathbb{E}[|x - X_n| \mid X^* = x] dx = L \mathbb{E}[|X_f^* - X_n|], \end{aligned}$$

where $X_f^* \sim f_0$, and thus Theorem 5.1 implies $\mathbb{E}[|X^* - X_n|] = o(c^{-n})$. In the case that the performance measure has an unbounded likelihood ratio with respect to f_0 , for example, when g_0 is a point mass at a given point, it remains an open research question whether or not the geometric rate of convergence still holds.

6. Summary and future research. The PBA, initially introduced in [11], provides an efficient method for locating an unknown point $X^* \in [0, 1]$ for the noisy bisection setting in which the only way to learn about X^* is by querying an oracle as to whether X^* is to the left or to the right of a prescribed point x , and the oracle provides the correct answer only with probability p . We have shown that the PBA is optimal in reducing expected posterior entropy, that it is a consistent method for locating X^* , and that the rate of convergence of the expected absolute residuals is at least geometric. This shows that the convergence rate of the noisy bisection search on continuous search space is comparable to the convergence rate of the noise-free bisection search and to the convergence rate of the discretized noisy bisection search.

Important future research directions regarding the PBA include the following:

- (i) Further investigation of its finite time properties. Proposition 5.5 provides a large deviation guarantee of the absolute residuals for large n , and Theorem 4.1 shows the optimality of reducing expected posterior entropy for all sample sizes $n \in \mathbb{N}$. Our main result, on the other hand, provides an asymptotic rate of convergence guarantee. It would be very informative and important for concrete applications to better understand the finite time properties of the PBA.
- (ii) Finding a worst-case guarantee for the rate of convergence. Our main result provides a rate of convergence guarantee for the expected absolute residuals under a prior density f_0 and thus is an average-case performance guarantee. It would be interesting to know whether or not a similar geometric rate holds in a worst-case scenario, that is, for any fixed $X^* \in [0, 1]$. Since the symmetry in the studied geometric random walks dominating $(A_n \wedge (1 - A_n))_n$ (see the proof of Proposition 5.3) no longer holds when conditioning on a fixed value $X^* \in [0, 1]$, the proof techniques used in this paper, unfortunately, cannot easily be extended to such a worst-case analysis.
- (iii) Analyzing the PBA where p , the probability of a correct reply from the oracle, is unknown and varies with x . References [4] and [5] study the discretized version of the PBA in the setting where p can vary with x but a lower bound on p is known. More specifically, they show that if $|p(x) - 1/2| \geq d|x - X^*|^{\kappa-1}$ for all $x \in [0, 1]$, where $|p(x) - 1/2| \leq \delta$ ($\kappa > 1, d > 0$ and $\delta > 0$ are all constants), then the expected residuals of the BZ algorithm converge at the rate $O((\log n/n)^{\kappa/(2\kappa-2)})$. A similar rate might also hold for the original PBA.
- (iv) Extending the PBA to higher dimensions. The method of centers of gravity,

developed independently in [14] and [16], generalizes the noise-free bisection to higher dimensions. Reference [15] provides a discussion of complexity and efficiency results of the method of centers of gravity and the subsequent ellipsoid method for deterministic optimization problems. A similar multivariate extension of the PBA seems plausible. Major challenges are the proper updating and tracking of the posterior density. This multivariate extension would be very useful for many applications, such as simulation-optimization methods.

Appendix. Additional proofs.

Proof of Lemma 3.1. Conditional on X^* and \mathcal{F}_{n-1} , the random variable $Z_n(x)$ assumes the value +1 with the following probabilities:

$$\begin{aligned} \mathbb{P}(Z_n(x) = +1 | X^* \geq x, \mathcal{F}_{n-1}) &= p, \\ \mathbb{P}(Z_n(x) = +1 | X^* < x, \mathcal{F}_{n-1}) &= 1 - p. \end{aligned}$$

The conditional distribution of the event $\{Z_n(x) = +1\}$ given \mathcal{F}_{n-1} is then computed as

$$\begin{aligned} \mathbb{P}(Z_n(x) = +1 | \mathcal{F}_{n-1}) &= \mathbb{P}(X^* \geq x | \mathcal{F}_{n-1}) \mathbb{P}(Z_n(x) = +1 | X^* \geq x, \mathcal{F}_{n-1}) \\ &\quad + \mathbb{P}(X^* < x | \mathcal{F}_{n-1}) \mathbb{P}(Z_n(x) = +1 | X^* < x, \mathcal{F}_{n-1}) \\ \text{(A.1)} \quad &= (1 - F_n(x))p + F_n(x)(1 - p) = \gamma(x), \end{aligned}$$

where the first equation follows from the law of total probability. The result now follows from Bayes' rule. That is, on the event $\{Z_n(x) = +1\}$ we have

$$\begin{aligned} f_{n+1}(y) &= \frac{\mathbb{P}(Z_n(x) = +1 | \mathcal{F}_{n-1}, X^* = y) f_n(y)}{\mathbb{P}(Z_n(x) = +1 | \mathcal{F}_{n-1})} \\ &= \begin{cases} \gamma(x)^{-1} p f_n(y) & \text{if } y \geq x, \\ \gamma(x)^{-1} (1 - p) f_n(y) & \text{if } y < x. \end{cases} \end{aligned}$$

The expression (3.2) for $f_{n+1}(x)$ on the event $\{Z_n(x) = -1\}$ is derived similarly. \square

Proof of Proposition 4.2. The definition of entropy and the tower property of conditional expectation imply

$$\mathbb{E}[H(f_N) | X_{N-1} = x, f_{N-1}] = \mathbb{E}[-\log_2 f_N(X^*) | X_{N-1} = x, f_{N-1}].$$

Using the updating equations described in Lemma 3.1 for the query $X_{N-1} = x$, we can decompose the random variable $-\log_2 f_N(X^*) | X_{N-1} = x, f_{N-1}$ into a sum of three terms for both possible responses of the oracle:

$$\begin{aligned} &\text{If } Z_{N-1}(X_{N-1}) = +1, \text{ then} \\ &\quad -\log_2 f_N(X^*) | X_{N-1} = x, f_{N-1} \\ &\quad = -\log_2 f_{N-1}(X^*) - \log_2 \gamma(x)^{-1} - \begin{cases} \log_2 p & \text{if } X^* \geq x, \\ \log_2(1 - p) & \text{if } X^* < x; \end{cases} \\ &\text{if } Z_{N-1}(X_{N-1}) = -1, \text{ then} \\ &\quad -\log_2 f_N(X^*) | X_{N-1} = x, f_{N-1} \\ &\quad = -\log_2 f_{N-1}(X^*) - \log_2(1 - \gamma(x))^{-1} - \begin{cases} \log_2(1 - p) & \text{if } X^* \geq x, \\ \log_2 p & \text{if } X^* < x. \end{cases} \end{aligned}$$

By the linearity of the expectation operator we can calculate the expected value of each of the three terms separately. The first term is independent of the oracle's response and simply recovers the entropy at time $N - 1$,

$$\mathbb{E}[-\log_2 f_{N-1}(X^*)|X_{N-1} = x, f_{N-1}] = H(f_{N-1}).$$

To evaluate the second term we use the fact that $\mathbb{P}(Z_{N-1}(x) = +1|f_{N-1}) = (1 - F_{N-1}(x))p + F_{N-1}(x)(1 - p) = \gamma(x)$ as was shown in (A.1). The expectation of the second term is

$$\begin{aligned} \mathbb{E}[\mathbb{1}\{Z_{N-1}(x) = +1\} \log_2 \gamma(x) + \mathbb{1}\{Z_{N-1}(x) = -1\} \log_2(1 - \gamma(x))|f_{N-1}] \\ = \gamma(x) \log_2 \gamma(x) + (1 - \gamma(x)) \log_2(1 - \gamma(x)). \end{aligned}$$

The third term is equal to $\log_2 p$ when the oracle answers the query correctly and to $\log_2(1 - p)$ otherwise. This is independent of the measurement point x . Hence the expectation of the third term equals $-p \log_2 p - (1 - p) \log_2(1 - p)$. Combining these three terms together and noting that the first and third terms do not depend on the measurement location x yields

$$\begin{aligned} \inf_{x \in [0,1]} \mathbb{E}[H(f_N)|X_{N-1} = x, f_{N-1}] &= H(f_{N-1}) - p \log_2 p - (1 - p) \log_2(1 - p) \\ &+ \inf_{x \in [0,1]} [\gamma(x) \log_2 \gamma(x) + (1 - \gamma(x)) \log_2(1 - \gamma(x))]. \end{aligned}$$

The inner expression over which we take the infimum depends on x only through $\gamma(x)$, the probability of observing $Z_{N-1}(x) = +1$, which can take values in $[0, 1]$. Consider the function $g(\gamma) = \gamma \log_2 \gamma + (1 - \gamma) \log_2(1 - \gamma)$, which is strictly convex and has a global minimum at $\gamma = 1/2$. Further, $\gamma(x) = 1/2$ when $F_{N-1}(x) = 1/2$, which shows that the optimal choice of x is the median of the pdf f_{N-1} . Finally, combining all three terms yields

$$\begin{aligned} \mathbb{E}[H(f_N)|F_{N-1}(X_{N-1}) = 1/2, f_{N-1}] \\ = H(f_{N-1}) - p \log_2 p - (1 - p) \log_2(1 - p) - 1, \end{aligned}$$

and this finishes the proof. \square

Proof of Theorem 4.1. We show for each $n = 0, 1, \dots, N$ that the value function is as claimed in (4.3), and that the median achieves the minimum in Bellman's recursion (4.2). This is sufficient to show the claim.

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

$$\begin{aligned} V_n(f_n) &= \inf_{x \in [0,1]} \mathbb{E}[V_{n+1}(f_{n+1})|X_n = x, f_n] \\ &= \inf_{x \in [0,1]} \mathbb{E}[H(f_{n+1}) - (N - n - 1)(1 + p \log_2 p + (1 - p) \log_2(1 - p))|X_n = x, f_n] \\ &= \inf_{x \in [0,1]} \mathbb{E}[H(f_{n+1})|X_n = x, f_n] - (N - n - 1)(1 + p \log_2 p + (1 - p) \log_2(1 - p)). \end{aligned}$$

Finally, Proposition 4.2 shows that the infimum is achieved at the median $x = \inf\{x : F_n(x) \geq 1/2\}$, and that the resulting value is

$$V_n(f_n) = H(f_n) - (N - n)(1 + p \log_2 p + (1 - p) \log_2(1 - p)),$$

as stated in the theorem. \square

Proof of Lemma 5.2. Let us first focus on the definition of C . The reason for defining C in this way will become clear toward the end of the proof. Consider the function

$$U(u) = \left(\frac{u + D}{\ln(2p) - \ln(2(1-p))} \right)^2,$$

and note that

1. U is convex and nonnegative;
2. $U(|D|) = 0$, because $D < 0$.

These two properties imply that there exists a unique $\tilde{u} \in (0, |D|)$ such that $U(\tilde{u}) = \tilde{u}$. Then define $C := e^{\tilde{u}}$ and consequently $1 < C < e^{|D|}$.

Now we return to the random walk $(R_n)_n$. For any $n \in \mathbb{N}$,

$$\begin{aligned} \mathbb{P}(e^{R_n} > C^{-n}/2) &= \mathbb{P}(R_n > \ln(C^{-n}/2)) \\ &= \mathbb{P}\left(R_0 + \sum_{j=1}^n \psi_j > \ln(2^{-1}C^{-n})\right) \\ &\leq \mathbb{P}\left(\ln(1/2) + \sum_{j=1}^n \psi_j > \ln(1/2) - n \ln C\right) \\ &= \mathbb{P}\left(\sum_{j=1}^n \psi_j > -n \ln C\right) \\ &= \mathbb{P}\left(\sum_{j=1}^n \psi_j - nD > -n \ln C - nD\right) \\ &= \mathbb{P}(\bar{\psi}_n - D > -\ln C - D), \end{aligned}$$

where $\bar{\psi}_n := (1/n) \sum_{j=1}^n \psi_j$ and $\mathbb{E}[\psi_j] = 1/2(\ln(2p) + \ln(2(1-p))) = D$. The increments ψ_j are iid and bounded, and $C < e^{|D|}$, which implies that $-\ln C - D > 0$, so we can apply Hoeffding's bound:³

$$\mathbb{P}(e^{R_n} > C^{-n}/2) \leq \exp\left(-2 \left(\frac{\ln C + D}{\ln(2p) - \ln(2(1-p))}\right)^2 n\right).$$

Now by definition of C ,

$$\left(\frac{\ln C + D}{\ln(2p) - \ln(2(1-p))}\right)^2 = \ln C,$$

and hence $\mathbb{P}(e^{R_n} > C^{-n}/2) \leq C^{-2n}$, which holds for any chosen $n \in \mathbb{N}$. \square

Proof of Theorem 5.8. Consider arbitrary $\varepsilon > 0$. Proposition 5.5 shows that $\mathbb{P}(c^n \mathbb{E}_n[|X^* - X_n|] > \varepsilon) \leq C^{-n}$ for $n > \hat{N} = \max(0, \tilde{N}(\varepsilon, c, C))$, and then

$$\sum_{n=0}^{\infty} \mathbb{P}(c^n \mathbb{E}_n[|X^* - X_n|] > \varepsilon) \leq \hat{N} + \frac{C^{-\hat{N}+1}}{C-1} < \infty.$$

³Let X_1, \dots, X_n be iid bounded random variables, that is, $\mathbb{P}(X_i \in [a, b]) = 1$. Then for the empirical mean $\bar{X} = (1/n) \sum_{i=1}^n X_i$ we have the inequality $\mathbb{P}(\bar{X} - \mathbb{E}[\bar{X}] \geq t) \leq \exp(-\frac{2t^2n}{(b-a)^2})$ for $t \geq 0$. See [10].

By the Borel-Cantelli lemma it follows that $\mathbb{P}(c^n \mathbb{E}_n[|X^* - X_n|] > \varepsilon \text{ i.o.}) = 0$.⁴ Since this holds for any $\varepsilon > 0$ it follows that $c^n \mathbb{E}_n[|X^* - X_n|] \rightarrow 0$, and hence $\mathbb{E}_n[|X^* - X_n|] \rightarrow 0$, almost surely as $n \rightarrow \infty$. \square

REFERENCES

- [1] M. BEN-OR AND A. HASSIDIM, *Quantum Search in an Ordered List via Adaptive Learning*, preprint, arXiv:quant-ph/0703231v2, 2007.
- [2] M. BEN-OR AND A. HASSIDIM, *The Bayesian learner is optimal for noisy binary search*, in Proceedings of the 49th Annual IEEE Symposium on Foundations of Computer Science, IEEE, Piscataway, NJ, 2008, pp. 221–230.
- [3] M. BURNASHEV AND K. ZIGANGIROV, *An interval estimation problem for controlled observations*, Problemy Peredachi Informatsii, 10 (1974), pp. 51–61.
- [4] R. CASTRO AND R. NOWAK, *Active learning and sampling*, in Foundations and Applications of Sensor Management, A. O. Hero III, D. A. Castañón, D. Cochran, and K. Kastella, eds., Springer, New York, 2008, pp. 177–200.
- [5] R. M. CASTRO AND R. D. NOWAK, *Minimax bounds for active learning*, IEEE Trans. Inform. Theory, 54 (2008), pp. 2339–2353.
- [6] T. M. COVER AND J. A. THOMAS, *Elements of Information Theory*, John Wiley & Sons, New York, 1991.
- [7] R. DURRETT, *Probability: Theory and Examples*, 3rd ed., Duxbury Advanced Series, Duxbury Press, Belmont, CA, 2005.
- [8] U. FEIGE, P. RAGHAVAN, D. PELEG, AND E. UPFAL, *Computing with noisy information*, SIAM J. Comput., 23 (1994), pp. 1001–1018.
- [9] P. I. FRAZIER, W. B. POWELL, AND S. DAYANIK, *A knowledge-gradient policy for sequential information collection*, SIAM J. Control Optim., 47 (2008), pp. 2410–2439.
- [10] W. HOEFFDING, *Probability inequalities for sums of bounded random variables*, J. Amer. Statist. Assoc., (1963), pp. 13–30.
- [11] M. HORSTEIN, *Sequential transmission using noiseless feedback*, IEEE Trans. Inform. Theory, 9 (1963), pp. 136–143.
- [12] B. JEDYNAK, P. I. FRAZIER, AND R. SZNITMAN, *Twenty questions with noise: Bayes optimal policies for entropy loss*, J. Appl. Probab., 49 (2012), pp. 114–136.
- [13] R. KARP AND R. KLEINBERG, *Noisy binary search and its applications*, in Proceedings of the 18th Annual ACM-SIAM Symposium on Discrete Algorithms, SIAM, Philadelphia, ACM, New York, 2007, pp. 881–890.
- [14] A. I. LEVIN, *On an algorithm for the minimization of convex functions*, Dokl. Akad. Nauk, 160 (1965), pp. 1244–1247 (in Russian). English translation in Dokl. Math., 6 (1965), pp. 286–290.
- [15] A. S. NEMIROVSKI AND D. B. YUDIN, *Problem Complexity and Method Efficiency in Optimization*, John Wiley & Sons, New York, 1983.
- [16] D. J. NEWMAN, *Location of the maximum on unimodal surfaces*, J. ACM, 12 (1965), pp. 395–398.
- [17] R. NOWAK, *Generalized binary search*, in Proceedings of the 46th Annual Allerton Conference on Communication, Control, and Computing, IEEE, Piscataway, NJ, 2008, pp. 568–574.
- [18] R. NOWAK, *Noisy generalized binary search*, in Advances in Neural Information Processing Systems 22, Y. Bengio, D. Schuurmans, J. Lafferty, C. K. I. Williams, and A. Culotta, eds., Neural Information Processing Systems Foundation, La Jolla, CA, 2009, pp. 1366–1374.
- [19] A. PELC, *Searching with known error probability*, Theoret. Comput. Sci., 63 (1989), pp. 185–202.
- [20] R. RIVEST, A. MEYER, D. KLEITMAN, K. WINKLMANN, AND J. SPENCER, *Coping with errors in binary search procedures*, J. Comput. System Sci., 20 (1980), pp. 396–404.
- [21] R. WAEBER, P. I. FRAZIER, AND S. G. HENDERSON, *A Bayesian approach to stochastic root finding*, in Proceedings of the 2011 Winter Simulation Conference, S. Jain, R. R. Creasey, J. Himmelsbach, K. P. White, and M. Fu, eds., IEEE, Piscataway, NJ, 2011, pp. 4038–4050.
- [22] S. ZANGENEHPUR, *Method and Apparatus for Positioning Head of Disk Drive Using Zone-Bit-Recording*, U.S. Patent 5257143, 1993.

⁴i.o. stands for infinitely often, i.e., $\{c^n \mathbb{E}_n[|X^* - X_n|] > \varepsilon \text{ i.o.}\} = \bigcap_{n=0}^{\infty} \bigcup_{j=n}^{\infty} \{c^j \mathbb{E}_j[|X^* - X_j|] > \varepsilon\}$.