

Finite-time Analysis for the Knowledge-Gradient Policy

Yingfei Wang ^{*} Warren Powell [†]

June 16, 2016

Abstract

We consider sequential decision problems in which we adaptively choose one of finitely many alternatives and observe a stochastic reward. We offer a new perspective of interpreting Bayesian ranking and selection problems as adaptive stochastic multi-set maximization problems and derive the first finite-time bound of the knowledge-gradient policy for adaptive submodular objective functions. In addition, we introduce the concept of prior-optimality and provide another insight into the performance of the knowledge gradient policy based on the submodular assumption on the value of information. We demonstrate submodularity for the two-alternative case and provide other conditions for more general problems, bringing out the issue and importance of submodularity in learning problems. Empirical experiments are conducted to further illustrate the finite time behavior of the knowledge gradient policy.

1 Introduction

We consider sequential decision problems in which at each time step, we choose one of finitely many alternatives and observe a random reward. The rewards are independent of each other and follow some unknown probability distribution. One goal can be to identify the alternative with the best expected performance within a limited measurement budget, which is the objective of Bayesian ranking and selection problems. Ranking and selection problems are examples of sequential decision making problems with partial information that address the exploration-exploitation trade-off. Since the learner does not know the true distribution of each alternative, it needs to explore the

^{*}Department of Computer Science, Princeton University, Princeton, NJ 08540, USA, yingfei@cs.princeton.edu

[†]Department of Operations Research and Financial Engineering, Princeton University, Princeton, NJ 08540, USA, powell@princeton.edu

choices that might give good rewards in the future as well as exploit the alternatives that appear to be better based on previous observations.

Ranking and selection problems arise in many settings. We may have to choose a type of material that has the best performance, the features in a laptop or car that produce the highest sales, or the molecular combination that produces the most effective drug. Often, the cost of a measurement may be substantial. Laboratory or field experiments may take a day or several weeks. For this reason, we assume we have a limited budget for making measurements.

Raiffa and Schlaifer established the Bayesian framework for R&S problems [33]. Several two-stage and sequential procedures exist for selecting the best alternative. Branke et. al made a thorough comparison of several fully sequential sampling procedures [6]. They indicate that the optimal computing budget allocation (OCBA) [9, 11, 21] and value of information procedures (VIP) [12] perform quite well and better than a deterministic or two-stage policy [10]. Another single-step Bayesian look-ahead policy first introduced by [19] and then further studied by [15] is called the “knowledge-gradient policy” (KG). It chooses to measure the alternative that maximizes the single-period expected value of information. Whereas the above mentioned policies assumed an independent normal or one-dimensional Wiener process prior on the alternatives’ true means, Frazier et. al modified the knowledge-gradient policy to handle correlated multivariate normal belief on the mean values of these rewards [13].

A similar field is the multi-armed bandit problem, which were originally studied under Bayesian assumptions [17]. A widely used class of policies for multi-armed bandit problems is called *upper confidence bounding* policies (UCB). Different UCB-type variants have been developed for many types of reward distributions and have provable logarithmic regret bounds [28, 1, 4, 27, 7]. By contrast, knowledge gradient policies, which enjoy some nice theoretical properties, have never been characterized by the type of regret bounds for which UCB policies are famous.

This paper makes the following contributions: We first establish the connection between Bayesian ranking and selection problem and adaptive stochastic multi-set function maximization problems where each multi-set corresponds to a set of selected alternatives. The multi-set representation captures our ability to evaluate the same alternative more than once. This new perspective offers a new line of analysis for the properties of value-of-information policies. We derive the first finite-time bound for the knowledge gradient policy for R&S problems under the assumption that the utility function is adaptive submodular. However, pathwise adaptive submodularity can fail in offline learning settings when the utility function itself involves a maximum. To this end, instead of the pathwise behavior analyses of the utility function, we further study its average behavior by taking expectations over the observations given any fixed sample allocation, resulting in a well-known quantity: the value of information. As a result, we introduce the concept of the prior-value of a policy and analyze the prior-optimality of the KG policy to provide another insight into its performance based on the submodular assumption

of the value of information that is weaker than adaptive submodularity. To accomplish this, we build on the general structure of the analysis of greedy algorithms given in [31] and [18]. We demonstrate submodularity for the two-alternative case and provide other conditions for more general problems, filling in a gap in the analysis of the knowledge gradient policy. Finally, we propose experiments to illustrate our theoretical analysis on the finite time behavior of the knowledge gradient policy. We further compare the KG policy with other policies with or without theoretical guarantees. Aside from the fact that the KG policy performs competitively with or significantly better than other policies especially in early iterations, we draw the conclusion that there is no universal best policy for all problem classes, which means that theoretical guarantees are not by themselves reliable indicators of which policy is best for a particular problem class and empirical experiments are needed to better understand their finite time performance.

This paper is organized as follows. In section 2, we lay out the mathematical models for Bayesian ranking and selection problems. In section 3, we describe the knowledge gradient policies. In section 4, we provide finite-time analyses of the knowledge gradient policy from two directions: the posterior optimality and the prior optimality. In section 5, we analyze the submodularity of the two-alternative case and provide other conditions for more general problems, bringing out the issue and importance of submodularity in leaning problems. Finally, in section 6 we present finite-time performance results and analyses of various policies for R&S problems.

2 Ranking and Selection Problems

Suppose we have a collection \mathcal{X} of M alternatives (where M might be quite large), each of which can be measured sequentially to estimate its unknown mean μ_x . We assume normally distributed measurement noise with known variance σ_W^2 . We first introduce the model for independent normal beliefs. We begin with a normally distributed Bayesian prior belief on the sampling means that is independent across alternatives, $\mu_x \sim \mathcal{N}(\theta_x^0, \sigma_x^0)$. At the n th iteration, we use some measurement policy π to choose one alternative x^n and observe $W_{x^n}^{n+1} \sim \mathcal{N}(\mu_{x^n}, \sigma_W)$.

For convenience, we introduce the σ -algebras \mathcal{F}^n for any $n = 0, 1, \dots, N - 1$ which is formed by the previous n measurement choices and outcomes, $x^0, W^1, \dots, x^{n-1}, W^n$. We define $\theta_x^n = \mathbb{E}[\mu_x | \mathcal{F}^n]$ and $(\sigma_x^n)^2 = \text{Var}[\mu_x | \mathcal{F}^n]$. Then conditionally on \mathcal{F}^n , $\mu_x \sim \mathcal{N}(\theta_x^n, \sigma_x^n)$. Let $\beta_x^n = \frac{1}{(\sigma_x^n)^2}$ be the conditional precision of μ_x and our state of knowledge be $S^n = (\theta_x^n, \beta_x^n)_{x \in \mathcal{X}}$. We will use \mathcal{F}^n and S^n interchangeably. After the n th measurement we update our beliefs using Bayes' rule:

$$\theta_x^{n+1} = \begin{cases} \frac{\beta_x^n \theta_x^n + \beta^W W^{n+1}}{\beta_x^n + \beta^W} & \text{if } x^n = x \\ \theta_x^n & \text{otherwise,} \end{cases} \quad \beta_x^{n+1} = \begin{cases} \beta_x^n + \beta^W & \text{if } x^n = x \\ \beta_x^n & \text{otherwise,} \end{cases}$$

where $\beta^W = 1/\sigma_W^2$.

We may impose correlated beliefs between alternatives in order to strengthen the effect of each measurement. Starting from a prior distribution $\mathcal{N}(\theta^0, \Sigma^0)$ and after measurement W^{n+1} of alternative x , a posterior distribution on the beliefs are calculated by:

$$\theta^{n+1} = \Sigma^{n+1} ((\Sigma^n)^{-1} \theta^n + \beta^W W^{n+1} e_x), \quad (1)$$

$$\Sigma^{n+1} = ((\Sigma^n)^{-1} + \beta^W e_x e_x^T)^{-1}, \quad (2)$$

where e_x is the vector with 1 in the entry corresponding to alternative x and 0 elsewhere. $S^n = (\theta^n, \Sigma^n)$ is then our state of knowledge in this case.

A decision function $X^\pi(S^n)$ is defined as a mapping from the knowledge state to \mathcal{X} . We refer to the decision function X^π and the policy π interchangeably.

If we are limited to N measurements, the objective is to maximize the expected reward of the final recommended alternative:

$$\max_{\pi \in \Pi} \mathbb{E}[\mu_{x^\pi}], \quad (3)$$

where $x^\pi = \arg \max_{x \in \mathcal{X}} \theta_x^N$ and $x^n = X^\pi(S^n)$ for $0 \leq n < N$.

3 Knowledge Gradient

For R&S problems, the knowledge gradient is a policy that at the n th iteration chooses its $(n+1)$ st measurement from \mathcal{X} to maximize the single-period expected increase in value [15, 13]. To be more specific, the value of being in state S^n is $\max_{x \in \mathcal{X}} \theta_x^n$. If we choose to measure $x^n = x$ right now, allowing us to observe W_x^{n+1} , then we transition to a new state of knowledge $S^{n+1} = (\theta^{n+1}, \Sigma^{n+1})$. At iteration n , θ_x^{n+1} is a random variable since we do not yet know what W^{n+1} is going to be. We would like to choose x at iteration n which maximizes the expected value of $\max_{x \in \mathcal{X}} \theta_x^{n+1}$. We can think of this as choosing an alternative to maximize the incremental value, given by

$$\nu_x^{\text{KG},n} = \mathbb{E}[\max_{x'} \theta_{x'}^{n+1} - \max_{x'} \theta_{x'}^n | x^n = x, S^n]. \quad (4)$$

The knowledge gradient policy $X^{\text{KG}}(S^n)$ is defined by

$$X^{\text{KG}}(S^n) = \arg \max_{x \in \mathcal{X}} \nu_x^{\text{KG},n}. \quad (5)$$

The knowledge gradient policy can handle the presence of a variety of belief models such as (generalized) linear [30, 36] or nonparametric [29, 5].

The knowledge gradient policy has some nice properties. For Bayesian ranking and selection problems, the knowledge gradient policy is optimal (by definition) if the measurement budget $N = 1$. The knowledge gradient is guaranteed to find the best alternative as the measurement budget N tends to infinity. If there are only two choices,

the knowledge gradient policy is optimal for any measurement budget. The knowledge gradient policy is the only stationary policy that is both myopically and asymptotically optimal. However, the KG has not enjoyed the finite-time bounds that have been popular in the UCB policies.

4 Finite-time Analysis of the Knowledge Gradient Policy

We follow the general structure of the analysis of greedy approximation [31] to develop the first finite-time bound for the knowledge gradient policy for R&S problems as follows. In Section 4.1, by interpreting the Bayesian R&S problems as the adaptive stochastic multi-set maximization problems, we show that the KG policy inherits precisely the performance guarantees of the greedy algorithm for classic submodular maximization problems if the utility function is adaptive submodular. We theoretically analyze the adaptive submodular assumption and point out that it can fail in the ranking and selection problems. In such cases, instead of the pathwise behavior analyses of the utility function, we study its average behavior by taking expectation over the observations in Section 4.2. In Section 4.3, we analyze the *prior-optimality* which provides another insight into the performance of the KG policy based on the submodularity of a well-understood quantity: value of information.

It is important to note that both the submodular maximization reduction and the theoretical analyses on the prior-optimality are not limited to the specific setup of Gaussian noise in observations and Gaussian prior structure. The theoretical guarantees are more generally applicable to any prior and measurement noise model as long as the adaptive submodular assumption or the submodular value of information assumption holds.

4.1 The Reduction of R&S to Adaptive Stochastic Multi-set Maximization

We first introduce the adaptive stochastic maximization problem. Let E be a finite set of items. Each item $e \in E$ maps to a random outcome of a measurement $\Phi(e)$ in a set O of possible values. We define a realization as a function $\phi : E \mapsto O$ representing the observation of each item in the ground set. Under Bayesian interpretation, we assume that there is a known prior probability distribution $p(\phi) := \mathbb{P}(\Phi = \phi)$ over all possible realizations. The adaptive stochastic optimization problem consists of sequentially picking an item $e \in E$, revealing its outcome $\Phi(e)$ and picking the next item. After each pick, the observations so far can be represented as a partial realization ψ . A partial realization ψ is consistent with realization ϕ , denoted as $\phi \sim \psi$, if all the items selected in ψ have the same outcomes as in ϕ . We use $\text{dom}(\psi)$ to refer to the items observed

in ψ . We use the notation $Z^\pi(\phi)$ to denote the set of items chosen by policy π under realization ϕ .

We wish to maximize some utility function $f : 2^E \times O^E \mapsto \mathbb{R}$ that depends on which items we pick and which states they are in. The expected utility of a policy π is $f_{\text{avg}}(\pi) := \mathbb{E}[f(Z^\pi(\Phi), \Phi)]$ where the expectation is taken over the prior distribution $p(\phi)$. The goal of adaptive stochastic set maximization problem is to find an optimal policy π^* that maximizes its expected utility under a cardinality constraint,

$$\pi^* \in \arg \max_{\pi} f_{\text{avg}}(\pi), \text{ subject to } |Z^\pi(\phi)| \leq N,$$

where N is the measurement budget.

It is not obvious to treat the ranking and selection problem in an adaptive stochastic multi-set maximization way of thinking. To see this, define the ground set $E = \mathcal{X}$. The outcomes are real numbers with $O = \mathbb{R}$. Each alternative $e = x$ can be selected multiple times. After each selection, its random outcome $\Phi(e) = W_x \in O$ is revealed.

Since the true values μ_x are random variables, we can let φ be a sample realization of the truth with a (correlated) prior distribution $p(\varphi) = \mathcal{N}(\theta^0, \Sigma^0)$. We use the notation $\phi \in \Phi$ to denote an realization of the random observations in our problem. The prior probability distribution over the realizations is determined by $p(\varphi)$ and the noise distribution $\mathcal{N}(0, \sigma_W)$. For example, if in the ranking and selection problems each alternative can only be selected once, $\phi : E \mapsto O$. For multi-selections, one way of defining the realization is by first making replicas of each item to construct E' and then selecting each $e' \in E'$ at most once.

Consider any sampling allocation $z = (z_x)_{x \in \mathcal{X}}$, by which we measure alternative x for $z_x \in \mathbb{N}$ times. We use Z to represent its corresponding multi-set. We use $Z^\pi(\phi) : \Phi \mapsto (\mathcal{X} \times \mathbb{N})$ to refer to the alternatives selected by π under realization ϕ . Let θ^n be our vector of estimates of the means after n measurements according to allocation Z under realization ϕ , where $|Z| = n$. θ^n can be obtained according to the updating equation (1) and (2), and does not depend on the order of the allocations. It can thus be denoted as $\theta^n(Z, \phi) : (\mathcal{X} \times \mathbb{N}) \times \Phi \mapsto \mathbb{R}^M$. The next lemma states the equivalence of $\mathbb{E}[\mu_{x^\pi}]$ and $\mathbb{E}[\max_x \theta_x^N]$. Hence, the utility function $f : (\mathcal{X} \times \mathbb{N}) \times \Phi \mapsto \mathbb{R}$ can be defined as $\max_x \theta_x^n(Z, \phi)$ and $f_{\text{avg}}(\pi) := \mathbb{E}[\max_x \theta_x^N(Z^\pi(\Phi), \Phi)]$. The R&S objective (3) can then be re-written as

$$\pi^* \in \arg \max_{\pi} f_{\text{avg}}(\pi), \text{ subject to } |Z^\pi(\phi)| \leq N.$$

LEMMA 4.1 ([32]). *Let π be a policy, and let $x^\pi = \arg \max_x \theta_x^N$ be the alternative selected by the policy. Then*

$$\mathbb{E}[\mu_{x^\pi}] = \mathbb{E}[\max_x \theta_x^N].$$

The definition of the knowledge gradient $\nu_x^{KG,n}$ coincides with the *Conditional Expected Marginal Benefit* $\Delta(e|\psi)$ defined by [18]:

$$\Delta(e|\psi) := \mathbb{E} \left[f(\text{dom}(\psi) \cup \{e\}, \Phi) - f(\text{dom}(\psi), \Phi) | \Phi \sim \psi \right].$$

The knowledge gradient policy is thus in fact the adaptive greedy policy with uniform item costs, with a slight difference in the ability of selecting each item more than once. We generalize the definition of adaptive monotonicity and adaptive submodularity for set functions given by [18] to multi-set functions as follows.

DEFINITION 4.2 (Adaptive Monotonicity). *A function $f : (\mathcal{X} \times \mathbb{N}) \times \Phi \mapsto \mathbb{R}$ is adaptive monotone with respect to distribution $p(\phi)$ if the conditional expected marginal benefit of any item is nonnegative: for all ψ and all $x \in \mathcal{X}$.*

$$\Delta(x|\psi) \geq 0.$$

DEFINITION 4.3 (Adaptive Submodularity). *A function $f : (\mathcal{X} \times \mathbb{N}) \times \Phi \mapsto \mathbb{R}$ is adaptive submodular with respect to distribution $p(\phi)$ if for all ψ and ψ' such that $\text{dom}(\psi) \subseteq \text{dom}(\psi')$ and both ψ, ψ' are consistent with some realization ϕ (i.e. $\psi \subseteq \psi'$), we have the conditional expected marginal benefit of any fixed item $x \in \mathcal{X}$ does not increase as more items are selected and observed,*

$$\Delta(x|\psi) \geq \Delta(x|\psi').$$

Let π^* be the optimal policy to R&S problems. If $f := \max_x \theta_x^n(Z, \phi)$ is adaptive monotone and adaptive submodular with respect to the prior distribution $p(\phi)$, then

$$f_{\text{avg}}(\text{KG}) > (1 - e^{-1})f_{\text{avg}}(\pi^*).$$

We next show that the instances generated by ranking and selection problems are adaptive monotone.

LEMMA 4.4. *In ranking and selection problems, the utility function $\max_x \theta_x$ is adaptive monotone with any Gaussian prior.*

Proof. For any ψ , let $n = |\psi|$. Then for any item $x \in \mathcal{X}$, $\Delta(x|\psi)$ can be rewritten as $\mathbb{E}[\max_{x'} \theta_{x'}^{n+1} - \max_{x'} \theta_{x'}^n | x^n = x, \mathcal{F}^n] = \nu_x^{\text{KG}, n}$. Since for any x , $\theta_x^{n+1} = \theta_x^n + \tilde{\sigma}(\Sigma^n, x^n)Z^{n+1}$, where $\tilde{\sigma}(\Sigma, x) = \frac{\Sigma e_x}{\sqrt{1/\beta^W + \Sigma_{xx}}}$ and the random variable Z^{n+1} is standard normal when conditioned on \mathcal{F}^n [13]. Hence we have $\mathbb{E}[\theta_{x'}^{n+1} | x^n = x, \mathcal{F}^n] = \theta_{x'}^n$ for any x' . By Jensen's inequality, we have $\Delta(x|\psi) = \nu_x^{\text{KG}, n} \geq 0$. \square

Even though intuition suggests that the utility function should be adaptive submodular in the amount of information collected, as we collect more information it is natural to expect that the marginal value of this information should decrease, yet it is not always the case as shown in the next lemma. The proof can be found in Appendix A.1.

LEMMA 4.5. *For any independent normal prior distribution $p(\varphi)$ and nondegenerated noise distribution (i.e. $\sigma^W \neq 0$), there exists ψ, ψ' and $x \in \mathcal{X}$ such that $\psi \subseteq \psi'$ and $\Delta(x|\psi) < \Delta(x|\psi')$.*

It can be seen that the adaptive submodular assumption can fail in the ranking and selection problems with the special utility function $f = \max_x \theta_x^n(Z, \phi)$ that involves maximization itself. Hence, instead of the above pathwise behavior analyses of the utility function, we would like to study its average behavior by taking the expectation over the observations given any fixed sample allocation Z in the next section.

4.2 The Value of Information

We define the pathwise value of information $\hat{v}(Z, \phi)$ as the incremental improvement over the best expected value that can be obtained without measurement, which is $\max_{x \in \mathcal{X}} \theta_x^0$,

$$\hat{v}(Z, \phi) := \max_{x \in \mathcal{X}} \theta_x^n(Z, \phi) - \max_{x \in \mathcal{X}} \theta_x^0.$$

The value of information $v(Z)$ is then defined to be

$$v(Z) := \mathbb{E}_\Phi[\hat{v}(Z, \Phi)],$$

where the expectation is taken over the prior distribution $p(\phi)$.

The value of information has a long history spanning the literatures of several disciplines. Stigler considers the value of information in economics when buyers search for the best price [35]. Howard laid the groundwork for the value of information in a decision-theoretic context and spawned a great deal of work in this area [22]. Yokota and Thompson gives a first comprehensive review of value of information analyses related to health risk management [37]. Raiffa and Schlaifer poses the Bayesian R&S problem and defines the associated value of information [33], which marked the beginning of a number of literature on the value of information within Bayesian R&S and the budgeted learning problem [20, 26, 9, 12, 15].

Since the value of information is a multi-set function, we first generalize the definitions and properties of submodular set functions described by [31] to submodular multi-set functions.

DEFINITION 4.6. *Given a finite set E , a real-valued function g on the set of multi-sets over E is called submodular if for all multi-sets S and T whose elements belong to E ,*

$$\rho_x(S) \geq \rho_x(T), \forall S \subseteq T, \forall x \in E,$$

where $\rho_x(S) \triangleq g(S \cup \{x\}) - g(S)$ is the incremental value of adding element x to the multi-set S .

Proposition 1. *Each of the following statements is equivalent and defines a submodular multi-set function (S pathwise and T are multi-sets on E , $x, y \in E$):*

1. $\rho_x(S) \geq \rho_x(T), \forall S \subseteq T$ and $\forall x$.
2. $\rho_x(S) \geq \rho_x(S \cup \{y\}), \forall S, x, y$.

3. $g(T) \leq g(S) + \sum_{x \in T-S} \rho_x(S) - \sum_{x \in S-T} \rho_x(S \cup T - \{x\}), \forall S, T.$
4. $g(T) \leq g(S) + \sum_{x \in T-S} \rho_x(S), \forall S \subseteq T.$

This proposition follows from a similar proof of Proposition 2.1 in [31]. For completeness we provide the proof in Appendix A.2.

It is obvious that if $\theta_x^n(Z, \phi)$ is adaptive monotone or adaptive submodular with respect to $p(\phi)$, then so does $\hat{v}(Z, \phi)$. It is also easy to show that if $\theta_x^n(Z, \phi)$ is adaptive monotone or adaptive submodular with respect to $p(\phi)$, then by the law of total expectation, i.e. $\mathbb{E}[\mathbb{E}[U|V]] = \mathbb{E}[U]$ for any random variables U and V , the value of information $v(Z)$ is monotone or submodular. We close this section by showing the monotonicity of the multi-set function v and leave the analyses of submodularity in Section 5.

LEMMA 4.7. (Monotonicity of the value of information)

For any sampling allocation Z_1 and Z_2 , if $Z_1 \subseteq Z_2$, then $v(Z_1) \leq v(Z_2)$.

Proof. We prove the monotonicity of v by showing $v(Z) \leq v(Z \cup \{x^{n+1}\})$ for any allocation Z (with $\sum_{x \in \mathcal{X}} z_x = n$) and any additional measurement x^{n+1} . By the tower property,

$$\begin{aligned}
v(Z \cup \{x^{n+1}\}) - v(Z) &= \mathbb{E}_{\Phi}[\hat{v}(Z \cup \{x^{n+1}\})] - \mathbb{E}[\hat{v}(Z)] \\
&= \mathbb{E}_{\Phi}[\max_{x \in \mathcal{X}} \theta_x^{n+1}(Z \cup \{x^{n+1}\})] - \mathbb{E}[\max_{x \in \mathcal{X}} \theta_x^n(Z)] \\
&= \mathbb{E}_{\Phi}[\mathbb{E}[\max_x \theta_x^{n+1}(Z \cup \{x^{n+1}\}) - \max_x \theta_x^n(Z) | \Phi \sim \psi_Z]] \\
&= \mathbb{E}_{\Phi}[\nu_x^{\text{KG},n}],
\end{aligned}$$

where ψ_Z is the partial realization with $\text{dom}(\psi_Z) = Z$. The lemma follows from the adaptive monotonicity, $\nu_x^{\text{KG},n} \geq 0$. □

4.3 Guarantees on the Prior-optimality of the Knowledge Gradient Policy

There are two ways to evaluate the value of a policy. The first, which we call the *posterior view*, conditions on the allocation $Z = Z^\pi(\Phi)$ that would have occurred under policy π for each sample path $\phi \in \Phi$. This is the more conventional approach for evaluating policies. The second, which we call the *prior view*, starts by characterizing the value of an arbitrary allocation Z (before we have seen any sample realizations).

More formally, the classical way to estimate the value of a policy is to calculate the incremental improvement over what we could do before we collect any information, is given by

$$f'_{\text{avg}}(\pi) = \mathbb{E}[f(Z^\pi(\Phi), \Phi)] - \max_x \theta_x^0.$$

We let $\mathbb{P}(\pi \rightsquigarrow Z)$ be the probability that policy π produces allocation Z . Since with a fixed budget of N measurements, there are only finite choices of possible allocations, using the tower property, we can condition on the allocation $Z^\pi = Z$ which gives us

$$f'_{\text{avg}}(\pi) = \sum_{Z \in \mathcal{Z}^N} \mathbb{P}(\pi \rightsquigarrow Z) \left(\mathbb{E}[\max_x \theta_x^n(Z^\pi(\Phi), \Phi) | Z^\pi = Z] - \max_x \theta_x^0 \right).$$

We note that in this method for evaluating a policy (which is the standard method), we only consider allocations Z that are actually produced by policy π for the outcomes in ϕ . This approach makes it much more difficult to understand the relationship between the allocation Z and the value of a policy.

For this reason, we adopt a different method of evaluating a policy which we term the *prior view*. Since this idea is new, we define it formally as follows

DEFINITION 4.8 (The prior-value of a policy). *Let \mathcal{Z}^n be the set of all possible allocations with a limited budget n . The value of a policy π with N measurements is defined as*

$$\begin{aligned} F^\pi &= \sum_{Z \in \mathcal{Z}^N} \mathbb{P}(\pi \rightsquigarrow Z) \left(\mathbb{E}_\Phi[\max_x \theta_x^n(Z, \Phi)] - \max_x \theta_x^0 \right) \\ &= \sum_{Z \in \mathcal{Z}^N} \mathbb{P}(\pi \rightsquigarrow Z) v(Z). \end{aligned}$$

In this view, we use the prior probability of an outcome $p(\phi)$ instead of the posterior $p(\phi | Z^\pi(\phi) = Z)$ which is conditioned on an allocation Z . The value of this approach is that it writes the value of a policy directly as a function of $v(Z)$, making it easier to study the effect of the properties of $v(Z)$ on the value of a policy. Intuitively, since a policy could generate different allocations Z for different sample realizations, it is natural to define the value of a policy π as the weighted sum of the expected value of information based on all possible allocations Z and the weight should be the probability of occurrence of Z based on policy π .

We make the following assumption which is weaker than the adaptive submodularity assumption and will analyze it further in Section 5.

Assumption 1. *The value of information v is a submodular multi-set function on the set of alternatives \mathcal{X} with respect to the prior distribution $p(\phi)$.*

Let π^* be the optimal sequential policy under a budget of N measurements in the sense that the prior-value of π^* is the largest. We call it *prior-optimality*. In what follows, we first bound KG's sub-prior-optimality in Proposition 4.12:

$$F^{\pi^*} \leq F^{\text{KG}^{[n]} @ \pi^*} \leq F^{\text{KG}^{[n-1]}} + N(F^{\text{KG}^{[n]}} - F^{\text{KG}^{[n-1]}}), \quad n = 1, 2, \dots, N.$$

Then we derive the worst-case bound for the KG policy in Theorem 4.14:

$$\frac{F^{\text{KG}}}{F^{\pi^*}} \geq 1 - \left(\frac{N-1}{N}\right)^N \geq \frac{e-1}{e} \approx 0.632.$$

Besides the *posterior optimality* bound obtained from adaptive stochastic multi-set maximization, the *prior-optimality* provides another insight into the performance of the KG policy based on a well-understood quantity: value of information.

DEFINITION 4.9 (Policy concatenation). [18] *A concatenated policy $\pi = \pi_1 @ \pi_2$ is constructed by running π_1 to completion, and then running policy π_2 from a fresh start ignoring all the information collected while running π_1 .*

To be more specific, suppose π_i has a budget of n_i , $i = 1, 2$, the first phase is to run π_1 for n_1 iterations starting from S^0 and we get a sample realization including decisions and their corresponding measurements. The second phase is to run π_2 for n_2 measurements starting from S^0 and we get another sample realization. Thus the sample realization of the concatenated process is all the decisions and their corresponding measurements collected in two phases. Note here, when running the second policy, we ignore all the information collected during running the first one, but when calculating the value of $\pi_1 @ \pi_2$, $F^{\pi_1 @ \pi_2}$, we use all the information collected in two phases.

DEFINITION 4.10 (Policy truncation). [18] *For a policy π , define the j -truncation $\pi^{[j]}$ of π as the policy that runs exactly $(j + 1)$ steps under π 's decision rule and $\pi^{\{j\}}$ as the single step policy that randomly chooses an alternative according to the probability distribution of policy π 's decision for the $(j + 1)$ -th step.*

We now show that the value of π_1 is no larger than the value of $\pi_1 @ \pi_2$.

LEMMA 4.11. $F^{\pi_1} \leq F^{\pi_2 @ \pi_1}$ for all policies π_1 and π_2 under any prior and probability distribution that describes a measurement.

Proof. We first show that $F^{\pi_1 @ \pi_2} = F^{\pi_2 @ \pi_1}$. In a concatenated policy, the two phases are independent since no information is shared among the two phases. Hence for a given allocation pair (Z_1, Z_2) where $Z_1 \in \mathcal{Z}^{n_1}$, $Z_2 \in \mathcal{Z}^{n_2}$, we have

$$\begin{aligned} \mathbb{P}(\pi_1 @ \pi_2 \rightsquigarrow (Z_1, Z_2)) &= \mathbb{P}(\pi_1 \rightsquigarrow Z_1) \mathbb{P}(\pi_2 \rightsquigarrow Z_2) \\ &= \mathbb{P}(\pi_2 \rightsquigarrow Z_2) \mathbb{P}(\pi_1 \rightsquigarrow Z_1) \\ &= \mathbb{P}(\pi_2 @ \pi_1 \rightsquigarrow (Z_2, Z_1)). \end{aligned}$$

$F^{\pi_1 @ \pi_2} = F^{\pi_2 @ \pi_1}$ follows immediately from taking the sum over all possible pairs of (Z_1, Z_2) such that $Z_2 \cup Z_1 = Z$ for any fixed allocation Z .

Therefore $F^{\pi_1} \leq F^{\pi_1 @ \pi_2}$ holds if and only if $F^{\pi_1} \leq F^{\pi_2 @ \pi_1}$. We then finish this proof by showing $F^{\pi_1} \leq F^{\pi_1 @ \pi_2}$. We write $F^{\pi_1 @ \pi_2} - F^{\pi_1}$ as a telescoping sequence

$$\begin{aligned}
& F^{\pi_1 @ \pi_2} - F^{\pi_1} \\
= & \sum_{Z \in \mathcal{Z}^{n_1+n_2}} v(Z) \mathbb{P}(\pi_1 @ \pi_2 \rightsquigarrow Z) - \sum_{Z_1 \in \mathcal{Z}^{n_1}} v(Z_1) \mathbb{P}(\pi_1 \rightsquigarrow Z_1) \\
= & \sum_{Z \in \mathcal{Z}^{n_1+n_2}} \sum_{Z_1 \cup Z_2 = Z} v(Z) \mathbb{P}(\pi_1 \rightsquigarrow Z_1) \mathbb{P}(\pi_2 \rightsquigarrow Z_2) \\
- & \sum_{Z_1 \in \mathcal{Z}^{n_1}} \sum_{Z_2 \in \mathcal{Z}^{n_2}} v(Z_1) \mathbb{P}(\pi_1 \rightsquigarrow Z_1) \mathbb{P}(\pi_2 \rightsquigarrow Z_2) \\
= & \sum_{Z_1 \in \mathcal{Z}^{n_1}} \sum_{Z_2 \in \mathcal{Z}^{n_2}} \left[v(Z_1 \cup Z_2) - v(Z_1) \right] \mathbb{P}(\pi_1 \rightsquigarrow Z_1) \mathbb{P}(\pi_2 \rightsquigarrow Z_2) \\
\geq & 0,
\end{aligned}$$

where the second equality holds due to the same reason as in the proof above for $F^{\pi_1 @ \pi_2} = F^{\pi_2 @ \pi_1}$ and the third equality is just the same summation in different orders. The last inequality holds because of the monotonicity of multi-set function v . \square

Based on the monotonicity of v and a similar argument as in Proposition 4.11, F is non-decreasing with respect to the number of measurements. Thus the more measurements, the better the policy. Hence π^* has exactly N measurements. We have the following sub-optimality bound on KG 's prior-value. For a proof see Appendix A.3.

PROPOSITION 4.12. *Let $\rho^{KG,n} = F^{KG^{[n]}} - F^{KG^{[n-1]}}$, then*

$$\begin{aligned}
F^{\pi^*} \leq F^{KG^{[n-1] @ \pi^*}} & \leq F^{KG^{[n-1]}} + N\rho^{KG,n} \\
& = \sum_{i=0}^{n-1} \rho^{KG,i} + N\rho^{KG,n}, \quad n = 0, 1, \dots, N-1.
\end{aligned} \tag{6}$$

We now derive a bound for the adaptive greedy policy by applying linear programming to the problem of minimizing $\frac{F^{KG}}{F^{\pi^*}}$ subject to the inequalities (6), which is a worst-case analysis. The following lemma states the linear program and its solution. We use it afterwards to establish the bounds.

LEMMA 4.13. *Given $N \in \mathbb{Z}_+$, consider the following linear program*

$$\begin{aligned}
& \min \sum_{i=0}^{N-1} a_i, \\
& \sum_{i=0}^{t-1} a_i + Na_t \geq 1, \quad t = 0, 1, \dots, N-1.
\end{aligned}$$

Then under these N constraints, $\min \sum_{i=0}^{N-1} a_i = 1 - \alpha^N$, where $\alpha = \frac{N-1}{N}$.

The proof of this lemma can be found in [31].

We have the following results, which generalizes the classic result of the greedy algorithm that achieves $(1 - 1/e)$ -approximation to prior-optimality for ranking and selection problems.

THEOREM 4.14. *Assume we have a budget of N measurements. Let π^* denote the optimal sequential policy for the ranking and selection problem, then we have*

$$\frac{F^{KG}}{F^{\pi^*}} \geq 1 - \left(\frac{N-1}{N}\right)^N.$$

Proof. By Proposition 4.12, we have $F^{\pi^*} \leq \sum_{i=0}^{n-1} \rho^{KG,i} + N\rho^{KG,n}$, $n = 0, 1, \dots, N-1$. Divide by F^{π^*} on both sides of this inequality, we have

$$1 \leq \sum_{i=0}^{n-1} \frac{\rho^{KG,i}}{F^{\pi^*}} + N \frac{\rho^{KG,n}}{F^{\pi^*}}, n = 0, 1, \dots, N-1.$$

Let $a_i = \frac{\rho^{KG,i}}{F^{\pi^*}}$, and then these inequalities are identical to the constraints in Lemma 4.13. We notice that

$$\min \sum_{i=0}^{N-1} a_i = \min \sum_{i=0}^{N-1} \frac{\rho^{KG,i}}{F^{\pi^*}} \leq \sum_{i=0}^{N-1} \frac{\rho^{KG,i}}{F^{\pi^*}} = \frac{F^{KG}}{F^{\pi^*}}.$$

By Lemma 4.13, we have $\min \sum_{i=0}^{N-1} a_i = 1 - \alpha^N$, so $\frac{F^{KG}}{F^{\pi^*}} \geq 1 - \alpha^N = 1 - \left(\frac{N-1}{N}\right)^N$. \square

5 Analysis of Submodularity of the Value of Information

The finite-time bounds obtained in the previous sections assume that the value of information is submodular. In general, submodularity does not hold for arbitrary value functions. In this section, we analyze the submodularity of the two-alternative case for independent beliefs.

While submodularity is a property for multi-set functions, we can extend it to any continuous function by making it possible for the increment to take any positive value. It could be easily extended to any continuous function. This allows us to use results from real analysis to study submodularity.

DEFINITION 5.1. *A function $f : \mathbb{R}^n \mapsto \mathbb{R}$ is submodular if for all $x, y \in \mathbb{R}^n$, $x_i \leq y_i$ and $\delta \in \mathbb{R}_+^n$,*

$$f(x + \delta) - f(x) \geq f(y + \delta) - f(y).$$

We show that submodularity of \mathcal{C}^2 functions is directly related to its second derivatives and cross-derivatives (the proof is given in Appendix A.4):

THEOREM 5.2. C^2 function $f: \mathbb{R}^n \rightarrow \mathbb{R}$ is submodular if and only if every element of its Hessian is non-positive.

The concavity of the value of information has been studied extensively by [14]. In this section, we only study the cross-derivatives of the value of information.

Let $M = 2$ and the measurement allocation $z = (z_1, z_2)$. The value of information $v(z) = s(z)f(-\frac{|\theta_1^0 - \theta_2^0|}{s(z)})$, where $s(z) = \sqrt{\tilde{\sigma}_1^2(z_1) + \tilde{\sigma}_2^2(z_2)}$, $\tilde{\sigma}_i^2(z_i) = \frac{\sigma_i^{2,0} z_i}{\sigma_W^2 / \sigma_i^{2,0} + z_i}$, $f(a) = a\Phi(a) + \phi(a)$, Φ and ϕ are the standard normal cumulative distribution and density respectively [14].

Although the value of information is not concave in general in the two-alternative case, v is concave on the region where all z_i 's are large enough (see Theorem 2 in [14]).

We directly calculate the first derivative and cross-derivative of v as

$$\begin{aligned} \frac{\partial v}{\partial z_1} &= \frac{\tilde{\sigma}_1(z_1)\tilde{\sigma}'_1(z_1)}{s(z)} \left[f\left(-\frac{|\theta_1^0 - \theta_2^0|}{s(z)}\right) + |\theta_1^0 - \theta_2^0| \frac{\Phi\left(-\frac{|\theta_1^0 - \theta_2^0|}{s(z)}\right)}{s(z)} \right], \\ \frac{\partial^2 v}{\partial z_1 \partial z_2} &= \frac{\tilde{\sigma}_1(z_1)\tilde{\sigma}'_1(z_1)\tilde{\sigma}_2(z_2)\tilde{\sigma}'_2(z_2)}{s^3(z)} \phi\left(-\frac{|\theta_1^0 - \theta_2^0|}{s(z)}\right) \left(\frac{|\theta_1^0 - \theta_2^0|^2}{\tilde{\sigma}_1^2(z_1) + \tilde{\sigma}_2^2(z_2)} - 1 \right). \end{aligned}$$

THEOREM 5.3. The value of information is submodular when $M = 2$ and $\theta_1^0 = \theta_2^0$.

Proof. Concavity of $v(z)$ is proven in Remark 2 by [14]. Since $\theta_1^0 = \theta_2^0$, $|\theta_1^0 - \theta_2^0| = 0$ and thus $\frac{\partial^2 v}{\partial z_1 \partial z_2} \leq 0$. Therefore, v is submodular in this case. \square

$\frac{\partial^2 v}{\partial z_1 \partial z_2} \leq 0$ is equivalent to $|\theta_1^0 - \theta_2^0|^2 \leq \tilde{\sigma}_1^2(z_1) + \tilde{\sigma}_2^2(z_2)$. Rewriting this inequality, we get

$$\frac{1}{\frac{1}{\sigma_1^{2,0}} + \frac{z_1}{\sigma_W^2}} + \frac{1}{\frac{1}{\sigma_2^{2,0}} + \frac{z_2}{\sigma_W^2}} \leq \sigma_1^{2,0} + \sigma_2^{2,0} - |\theta_1^0 - \theta_2^0|^2. \quad (7)$$

We need $\sigma_1^{2,0} + \sigma_2^{2,0} - |\theta_1^0 - \theta_2^0|^2 \geq 0$, which can be achieved by setting our prior variance large enough or using a uniform prior over all alternatives. This is very reasonable when we have very little information about our problem domain.

Inequality equation (7) defines a region in the $z_1 - z_2$ plane. Specifically, this region has the hyperbolic line $\frac{1}{\frac{1}{\sigma_1^{2,0}} + \frac{z_1}{\sigma_W^2}} + \frac{1}{\frac{1}{\sigma_2^{2,0}} + \frac{z_2}{\sigma_W^2}} = \sigma_1^{2,0} + \sigma_2^{2,0} - |\theta_1^0 - \theta_2^0|^2$ as its boundary and contains infinity. In particular, when z_1 and z_2 are large enough (or equivalently when our measurement is accurate enough), the value of information is submodular.

Since there is no closed-form expression for the value of information under arbitrary allocations, we cannot verify submodularity in a simple way for problems with more than two alternatives and for correlated beliefs. Instead, it can be checked using numerical approximation and is easy to guarantee by running repeated experiments and averaging to reduce measurement noise. A necessary condition is the concavity of the value of information for measuring a fixed alternative x for n times, which can be checked exactly.

Intuitively, we may expect that the marginal value of information should decline as we make more observations. But it is not always the case. It is shown that the value of information for measuring a single alternative may form an S-curve which is concave when there are many measurements, but may be convex at the beginning [14]. The S-curve behavior arises when the measurement noise is large and thus a single measurement simply contains too little information, leading to algorithmic difficulties and apparent paradoxes. This issue is not related to any specific policy, but rather is an inherent property of learning problems. Although the value of information is not necessarily concave, it can be made concave by measuring each alternative enough times or (equivalently) using sufficiently precise measurements.

6 Computational Experiments

Since the seminal paper by [28], there has been a long history in the optimal learning literature of designing algorithms with provable asymptotic or finite-time bounds [2, 8, 34, 4, 16, 3]. But none of these bounds are tight in finite time and different bounds can be based on different metrics. Hence, empirical experiments are needed to better understand the finite time performance of each policy. To this end, we propose experiments to illustrate the finite time behavior of both KG and other optimal learning policies. We consider the following learning settings that arise a lot in black box Bayesian optimization.

Equal-prior: $M = 100$. The true values μ_x are uniformly distributed over $[0, 60]$ and measurement noise $\sigma_W = 100$. $\theta_x^0 = 30$ and $\sigma_x^0 = 10$ for every x .

Asymmetric unimodular function (AUF): x is a controllable parameter ranging from 21 to 120. The objective function is $F(x, \xi) = \theta_1 \min(x, \xi) - \theta_2 x$, where θ_1, θ_2 and the distribution of the random variable ξ are all unknown. The aim is to solve $\max_x \mathbb{E}F(x, \xi)$ while learning θ_1, θ_2 and the parameters that determine the distribution of ξ . The true distribution of ξ is taken as a normal distribution with mean 60 and standard deviation 18 (corresponding to a 30% noise ratio).

Goldstein-Price’s function with additive noise:

$$f(x, y, \phi) = [1 + (x + y + 1)^2(19 - 14x + 3x^2 - 14y + 6xy + 3y^2)] \cdot [30 + (2x - 3y)^2(18 - 32x + 12x^2 + 48y - 36xy + 27y^2)] + \phi,$$

where $-3 \leq x \leq 3$, $-3 \leq y \leq 3$ and are uniformly discretized into 13×13 alternatives.

In order to obtain the prior distribution, we follow [24] and [23] to use Latin hypercube designs for initial fit. For independent beliefs, we adopt a uniform prior with the same mean value θ_x^0 and standard deviation σ_x^0 for all alternatives. For correlated beliefs, we use a constant mean value θ_x^0 for all alternatives and a prior covariance matrix of the form

$$\Sigma_{xx'}^0 = \sigma e^{-\sum_{i=1}^d \lambda_i (x_i - x'_i)^2},$$

where each arm x is a d -dimensional vector and σ, λ_i are constant. We adopt the rule of thumb by [24] for the default number ($10 \times p$) of points, where p is the number of parameters to be estimated. In addition, as suggested by [23], to estimate the random errors, after the first $10 \times p$ points are evaluated, we add one replicate at each of the locations where the best p responses are found. Maximum likelihood estimation is then used to estimate the parameters based on the points in the initial design.

The policies considered in this section is described as follows.

EXPL: A pure exploration strategy that tests each alternative equally often.

EXPT: A pure exploitation strategy, $X^{\text{EXPT},n}(S^n) = \arg \max_x \hat{\mu}_x^n$.

Interval Estimation (IE): [25]

$$X^{\text{IE},n}(S^n) = \arg \max_x \theta_x^n + z_{\alpha/2} \sigma_x^n.$$

Kriging: [23]

Let $x^* = \arg \max_x (\theta_x^n + \sigma_x^n)$, then

$$X^{\text{Kriging},n}(S^n) = \arg \max_x (\theta_x^n - \theta_{x^*}^n) \Phi\left(\frac{\theta_x^n - \theta_{x^*}^n}{\sigma_x^n}\right) + \sigma_x^n \phi\left(\frac{\theta_x^n - \theta_{x^*}^n}{\sigma_x^n}\right),$$

where ϕ and Φ are the standard normal density and cumulative distribution functions.

UCB-E: [2]

$$X^{\text{UCB-E},n}(S^n) = \arg \max_x \hat{\mu}_x^n + \sqrt{\frac{\alpha}{N_x^n}},$$

where $\hat{\mu}_x^n, N_x^n$ are the sample mean of μ_x and number of times x has been measured up to time n . The quantity $\hat{\mu}_x^0$ is initialized by measuring each alternative once.

SR: [2] Let $A_1 = \mathcal{X}$, $\log(M) = \frac{1}{2} + \sum_{i=2}^M \frac{1}{i}$,

$$n_m = \left\lceil \frac{1}{\log(M)} \frac{n - M}{M + 1 - m} \right\rceil.$$

For each phase $m = 1, \dots, M - 1$:

1. For each $x \in A_m$, select alternative x for $n_m - n_{m-1}$ rounds.
2. Let $A_{m+1} = A_m \setminus \arg \min_{x \in A_m} \hat{\mu}_x$.

6.1 Finite Time Performance of Different Policies

Although the theoretical analysis in the previous section is to bound the performance of the knowledge gradient policy to the optimal policy (in theory), the optimal sequential policy is impossible to find in practice. To this end, we compare the value of KG to the expected value of the best alternative $\max_x \mu_x$. Define the opportunity cost (OC^π) of any policy π at any time step n as:

$$\text{OC}^\pi = \max_x \mu_x - \mu_{\hat{x}^n},$$

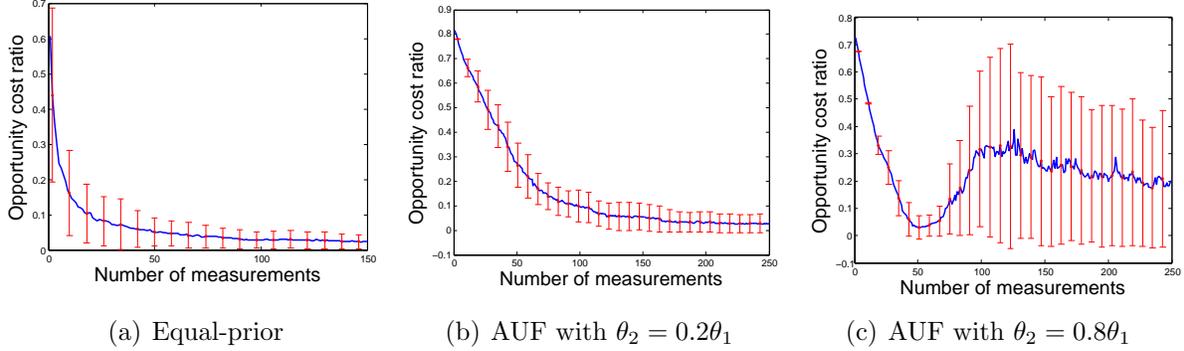


Figure 1: Opportunity cost ratio.

where $\tilde{x}^n = \arg \max_x \theta_x^n$. We illustrate the finite time behavior of the KG policy under Equal-prior and AUF with independent normal beliefs. We run KG and calculate the opportunity cost ratio $= \frac{\max_x \mu_x - \mu_{\tilde{x}^n}}{\max_x \mu_x}$ in each iteration. We report the mean with 90% confidence interval averaged over 1000 experiments in Figure 1.

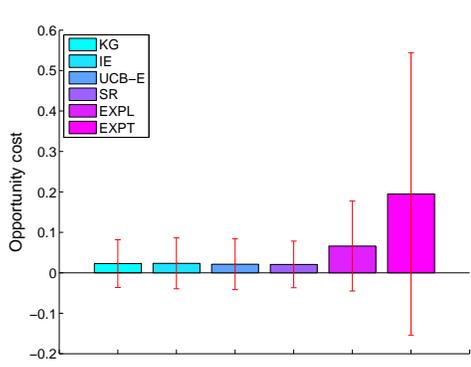
We next compare the performance of KG, IE with tuning, UCB-E with tuning, SR, EXPL and EXPT. Figure 2 shows the performance in problem classes AUF and Goldstein with independent beliefs under a measurement budget five times the number of alternatives. We run each policy for 1000 times. In each run, we pre-generate all the observations and share across different policies. We illustrate in the first column of Figure 2 the mean opportunity cost and the standard deviation of each policy over 1000 runs after the measurement budget is exhausted.

In order to give a comprehensive comparison based on different metrics, we also calculate the probability that the final recommendation of each policy is the optimal one and the probability that the opportunity cost of each policy is the lowest, as illustrated in the figures on the right hand side of Figure 2.

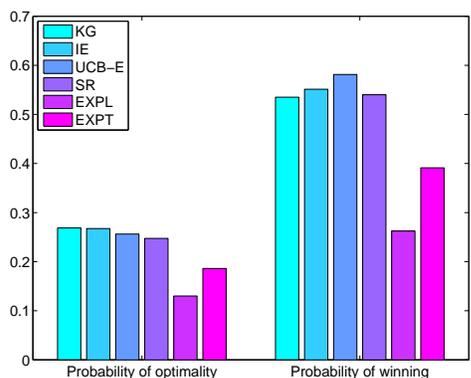
The three criteria characterize the behavior of policies from different perspectives. One observation is that there is no universal best policy for all problem classes or under all criteria, which means that theoretical guarantees are not by themselves reliable indicators of which policy is best for a particular problem class.

We also exploit correlated beliefs between alternatives in order to strengthen the effect of each measurement so that one measurement of some alternative can provide information for other alternatives.

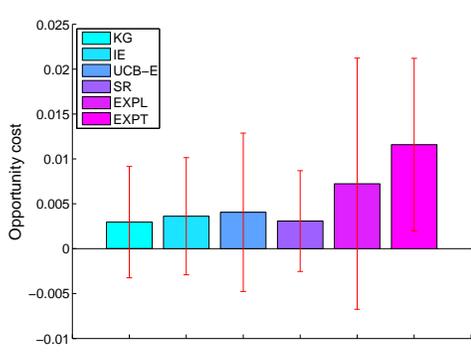
First, we present the OC of different policies after each iteration under AUF ($\theta_2 = 0.5\theta_1$) in Figure 3. We tune z_α for IE and α for UCB for $N = 400$ measurements and the optimal values are $z_\alpha = 0.969$ and $\alpha = 6.657$. Since UCB-E needs to measure each alternative once, we omit the OC for its first 100 (which is the number of alternatives) steps. KG uses independent beliefs while KGCB, IE and Kriging start from MLE fitted correlated beliefs. When incorporating correlated beliefs, a measurement of one alternative tells us something about other alternatives. As a result, KGCB learns faster



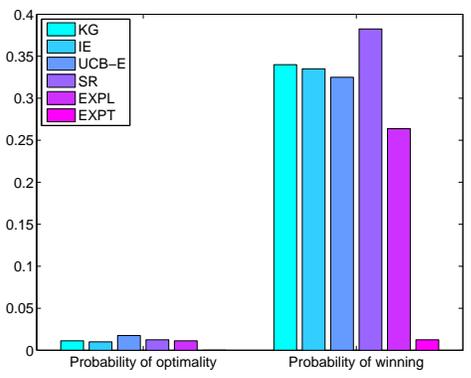
(a) AUF: Opportunity cost



(b) AUF: Probability of optimality/winning



(c) Goldstein: Opportunity cost



(d) Goldstein: Probability of optimality/winning

Figure 2: Comparisons for AUF and Goldstein. (a) and (c) depict the mean opportunity cost with error bars indicating the standard deviation of each policy. The first bar group in (b) and (d) demonstrates the probability that the final recommendation of each policy is the optimal one. The second bar group in (b) and (d) illustrates the probability that the opportunity cost of each policy is the lowest.

than KG.

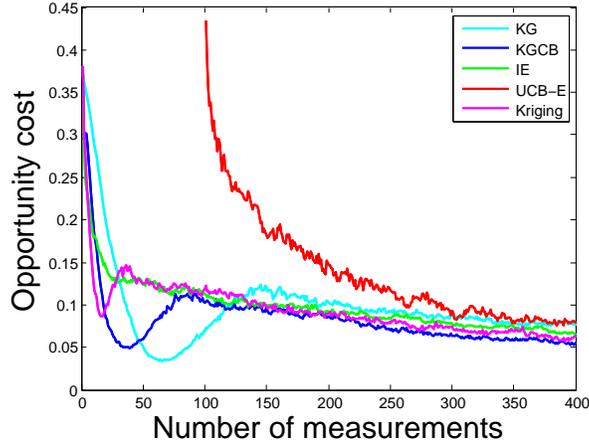


Figure 3: OC obtained by each policy after each measurement under AUF ($\theta_2 = 0.5\theta_1$).

7 Conclusion

In this paper, we offer a new perspective of interpreting ranking and selection problems as adaptive stochastic multi-set maximization problems. We present the first finite-time bounds for the knowledge gradient on both the posterior optimality and the prior optimality. The prior view provides a cleaner relationship between the performance of the policy and the sample taken, making it possible to relate the value of information to the submodularity of the sample. We analyze the submodularity of the two-alternative case and provide other conditions for more general problems, bringing out the issue and importance of submodularity in learning problems. We propose experiments to further illustrate the finite time behavior of the knowledge gradient policy as well as other policies with or without theoretical guarantees.

A Proofs

A.1 Proof of Lemma 4.5

For any ψ with $|\psi| = n$, we consider the resulting knowledge state $S^n = (\theta_x^n, \beta_x^n)_{x \in \mathcal{X}}$. Since $\sigma^W \neq 0$, there exists such ψ that $\max_x \theta_x^n > \max_{x \neq x'} \theta_x^n$ with positive probability. Now consider another realization ψ' with $\text{dom}(\psi') = \text{dom}(\psi) \cup \{x_2\}$, where x_2 is the second largest alternative of θ_x^n . We denote the observation of x_2 in ψ' as W_2 and the resulting S^{n+1} as $(\theta_x^{n+1}, \beta_x^{n+1})_{x \in \mathcal{X}}$ according to Bayes' rule. The knowledge gradient $\Delta(x|\psi) = \nu_x^{\text{KG},n}$ can be analytically expressed by

$$\nu_x^{\text{KG},n} = \tilde{\sigma}_x^n f(\zeta_x^n),$$

where $\tilde{\sigma}_x^n = \sqrt{(\beta_x^n)^{-1} - (\beta_x^n + \beta^W)^{-1}}$, $\zeta_x^n = -\left| \frac{\theta_x^n - \max_{x' \neq x} \theta_{x'}^n}{\tilde{\sigma}_x^n} \right|$ and $f(\zeta) = \zeta \Phi(\zeta) + \phi(\zeta)$. $\Phi(\zeta)$ and $\phi(\zeta)$ are, respectively, the cumulative standard normal distribution and the standard normal density [15]. We first notice that $f'(\zeta) = \Phi(\zeta) \geq 0$ for any $\zeta \in \mathbb{R}$ so that $f(\zeta)$ is non-decreasing. We next compare $\nu_{x_1}^{\text{KG},n}$ and $\nu_{x_1}^{\text{KG},n+1}$ for $x_1 = \arg \max_x \theta_x^n$. According to Bayes' rule, the precision β of x_2 changes only when x_2 is measured. So we have $\tilde{\sigma}_{x_1}^n = \tilde{\sigma}_{x_1}^{n+1}$. Similarly we have all the θ_x^{n+1} unchanged except for alternative x_2 . By some algebra, it can be shown that for any W_2 such that $\theta_{x_2}^n < W_2 \leq \frac{\beta_{x_2}^n}{\beta^W} (\theta_{x_1}^n - \theta_{x_2}^n) + \theta_{x_1}^n$, we have $\nu_{x_1}^{\text{KG},n} < \nu_{x_1}^{\text{KG},n+1}$. Since $\theta_{x_1}^n > \theta_{x_2}^n$ by construction, such W_2 can be obtained with positive probability.

A.2 Proof of Proposition 1

In this appendix, we prove the properties of submodular multi-set functions. We prove the equivalence by showing $2) \Rightarrow 1) \Rightarrow 3) \Rightarrow 4) \Rightarrow 2)$.

- $2) \Rightarrow 1)$. Take $S \subseteq T$ and $T - S = \{x_1, x_2, \dots, x_r\}$. Then from 3) we have $\rho_x(S) \geq \rho_x(S \cup \{x_1\})$, $\rho_x(S \cup \{x_1\}) \geq \rho_x(S \cup \{x_1, x_2\})$, ..., $\rho_x(S \cup \{x_1, x_2, \dots, x_{r-1}\}) \geq \rho_x(T)$. Summing these r inequalities yields 1).
- $1) \Rightarrow 3)$. For arbitrary S and T with $T - S = \{x_1, x_2, \dots, x_r\}$ and $S - T = \{y_1, y_2, \dots, y_q\}$, from 1) we have

$$\begin{aligned} g(S \cup T) - g(S) &= \sum_{t=1}^r [g(S \cup \{x_1, \dots, x_t\}) - g(S \cup \{x_1, \dots, x_{t-1}\})] \\ &= \sum_{t=1}^r \rho_{x_t}(S \cup \{x_1, \dots, x_{t-1}\}) \\ &\leq \sum_{t=1}^r \rho_{x_t}(S) = \sum_{x \in T-S} \rho_x(S). \end{aligned} \tag{8}$$

And

$$\begin{aligned}
g(S \cup T) - g(T) &= \sum_{t=1}^q [g(T \cup \{y_1, \dots, y_t\}) - g(T \cup \{y_1, \dots, y_{t-1}\})] \\
&= \sum_{t=1}^q \rho_{y_t}(T \cup \{y_1, \dots, y_t\} - \{y_t\}) \\
&\geq \sum_{t=1}^q \rho_{y_t}(T \cup S - \{y_t\}) = \sum_{x \in S-T} \rho_x(S \cup T - \{x\}). \quad (9)
\end{aligned}$$

Subtracting equation (9) from equation (8) we get 3).

- 3) \Rightarrow 4). If $S \subseteq T$, $S - T = \emptyset$, and therefore the last term in 3) vanishes.
- 4) \Rightarrow 2). Substitute $T = S \cup \{x, y\}$ into 4) to obtain

$$g(S \cup \{x, y\}) \leq g(S) + \rho_x(S) + \rho_y(S) = \rho_x(S) + g(S \cup \{y\}).$$

Rearrange this inequality, we get

$$\rho_x(S \cup \{y\}) = g(S \cup \{x, y\}) - g(S \cup \{y\}) \leq \rho_x(S).$$

A.3 Proof of Proposition 4.12

Let $z^*(Z, \pi, \Phi)$ be the next adaptive greedy choice that maximizes the expected marginal increment given that policy π has generated Z . We first show that

$$F^{\pi_2 @ \pi_1} \leq F^{\pi_2} + n_1 \sum_{Z \in \mathcal{Z}^n} \mathbb{P}(\pi_2 \rightsquigarrow Z) \left(\mathbb{E}[\hat{v}(Z \cup \{z^*(Z, \pi_2, \Phi)\}, \Phi)] - v(Z) \right)$$

for all policies π_1 with a measurement budget n_1 and π_2 with a budget n_2 under any prior and probability distribution that describes a measurement.

Proof. Let $\pi^{[j]}$ denote the first j measurement decisions under some policy π . First of all we break $F^{\pi_2 @ \pi_1} - F^{\pi_2}$ into n_1 consecutive differences,

$$F^{\pi_2 @ \pi_1} - F^{\pi_2} = \sum_{j=1}^{n_1} \left(F^{\pi_2 @ \pi_1^{[j]}} - F^{\pi_2 @ \pi_1^{[j-1]}} \right).$$

Similar to what we did in the last lemma, for each difference we have

$$\begin{aligned}
& F^{\pi_2 @ \pi_1^{[j]}} - F^{\pi_2 @ \pi_1^{[j-1]}} \\
= & \sum_{Z_1 \in \mathcal{Z}^{n_2+j}} \mathbb{P}(\pi_2 @ \pi_1^{[j]} \rightsquigarrow Z_1) v(Z_1) - \sum_{Z_2 \in \mathcal{Z}^{n_2+j-1}} \mathbb{P}(\pi_2 @ \pi_1^{[j-1]} \rightsquigarrow Z_2) v(Z_2) \\
= & \sum_{Z_1 \in \mathcal{Z}^{n_2+j}} \sum_{Z_2 \in \mathcal{Z}^{n_2+j-1}, Z_2 \cup Z_3 = Z_1} \mathbb{P}(\pi_2 @ \pi_1^{[j-1]} \rightsquigarrow Z_2) \mathbb{P}(\pi_1^{\{j\}} \rightsquigarrow Z_3 | \pi_2 @ \pi_1^{[j-1]} \rightsquigarrow Z_2) v(Z_1) \\
& - \sum_{Z_2 \in \mathcal{Z}^{n_2+j-1}} \sum_{Z_3 \in \mathcal{Z}^1} \mathbb{P}(\pi_2 @ \pi_1^{[j-1]} \rightsquigarrow Z_2) \mathbb{P}(\pi_1^{\{j\}} \rightsquigarrow Z_3 | \pi_2 @ \pi_1^{[j-1]} \rightsquigarrow Z_2) v(Z_2) \\
= & \sum_{Z_2 \in \mathcal{Z}^{n_2+j-1}} \sum_{Z_3 \in \mathcal{Z}^1} \mathbb{P}(\pi_2 @ \pi_1^{[j-1]} \rightsquigarrow Z_2) \mathbb{P}(\pi_1^{\{j\}} \rightsquigarrow Z_3 | \pi_2 @ \pi_1^{[j-1]} \rightsquigarrow Z_2) (v(Z_2 \cup Z_3) - v(Z_2)).
\end{aligned}$$

Now we consider all possible pair (Z_4, Z_5) such that $Z_4 \in \mathcal{Z}^{n_2}$, $Z_5 \in \mathcal{Z}^{j-1}$ and $Z_4 \cup Z_5 = Z_2$. Notice that the policy $\pi_2 @ \pi_1^{[j]}$ employs a fresh start at the time n_2 , therefore the events before and after time n_2 are independent. Then we have

$$\begin{aligned}
& \sum_{Z_2 \in \mathcal{Z}^{n_2+j-1}} \sum_{Z_3 \in \mathcal{Z}^1} \mathbb{P}(\pi_2 @ \pi_1^{[j-1]} \rightsquigarrow Z_2) \mathbb{P}(\pi_1^{\{j\}} \rightsquigarrow Z_3 | \pi_2 @ \pi_1^{[j-1]} \rightsquigarrow Z_2) (v(Z_2 \cup Z_3) - v(Z_2)) \\
= & \sum_{Z_2 \in \mathcal{Z}^{n_2+j-1}} \sum_{Z_4 \cup Z_5 = Z_2} \sum_{Z_3 \in \mathcal{Z}^1} \mathbb{P}(\pi_2 \rightsquigarrow Z_4) \mathbb{P}(\pi_1^{[j-1]} \rightsquigarrow Z_5) \mathbb{P}(\pi_1^{\{j\}} \rightsquigarrow Z_3 | \pi_2 @ \pi_1^{[j-1]} \rightsquigarrow Z_2) \\
& \times (v(Z_2 \cup Z_3) - v(Z_2)).
\end{aligned}$$

Based on the submodular property of function v , we have

$$v(Z_2 \cup Z_3) - v(Z_2) \leq v(Z_4 \cup Z_3) - v(Z_4).$$

Then from the definition of z^* , we have

$$\begin{aligned}
v(Z_4 \cup Z_3) - v(Z_4) &= \mathbb{E}[\hat{v}(Z_4 \cup Z_3, \Phi) - \hat{v}(Z_4, \Phi)] \\
&= \mathbb{E}_{\Phi} [\mathbb{E}[\hat{v}(Z_4 \cup Z_3, \Phi) - \hat{v}(Z_4, \Phi) | Z^{\pi_2}(\Phi) = Z_4]] \\
&\leq \mathbb{E}_{\Phi} [\mathbb{E}[\hat{v}(Z_4 \cup \{z^*(Z_4, \pi_2, \Phi)\}, \Phi) - \hat{v}(Z_4, \Phi) | Z^{\pi_2}(\Phi) = Z_4]] \\
&= \mathbb{E}_{\Phi} [\hat{v}(Z_4 \cup \{z^*(Z_4, \pi_2, \Phi)\}, \Phi)] - v(Z_4).
\end{aligned}$$

Combining the last two inequalities, we have

$$\begin{aligned}
& \sum_{Z_2 \in \mathcal{Z}^{n_2+j-1}} \sum_{Z_4 \cup Z_5 = Z_2} \sum_{Z_3 \in \mathcal{Z}^1} \mathbb{P}(\pi_2 \rightsquigarrow Z_4) \mathbb{P}(\pi_1^{[j-1]} \rightsquigarrow Z_5) \mathbb{P}(\pi_1^{\{j\}} \rightsquigarrow Z_3 | \pi_2 @ \pi_1^{[j-1]} \rightsquigarrow Z_2) \\
& \times (v(Z_2 \cup Z_3) - v(Z_2)) \\
\leq & \sum_{Z_2 \in \mathcal{Z}^{n_2+j-1}} \sum_{Z_4 \cup Z_5 = Z_2} \sum_{Z_3 \in \mathcal{Z}^1} \mathbb{P}(\pi_2 \rightsquigarrow Z_4) \mathbb{P}(\pi_1^{[j-1]} \rightsquigarrow Z_5) \mathbb{P}(\pi_1^{\{j\}} \rightsquigarrow Z_3 | \pi_2 @ \pi_1^{[j-1]} \rightsquigarrow Z_2) \\
& \times (\mathbb{E}\hat{v}(Z_4 \cup \{z^*(Z_4, \pi_2, \Phi)\}, \Phi) - v(Z_4)) \\
= & \sum_{Z_2 \in \mathcal{Z}^{n_2+j-1}} \sum_{Z_4 \cup Z_5 = Z_2} \mathbb{P}(\pi_2 \rightsquigarrow Z_4) \mathbb{P}(\pi_1^{[j-1]} \rightsquigarrow Z_5) (\mathbb{E}\hat{v}(Z_4 \cup \{z^*(Z_4, \pi_2, \Phi)\}, \Phi) - v(Z_4)) \\
= & \sum_{Z_4 \in \mathcal{Z}^{n_2}} \sum_{Z_5 \in \mathcal{Z}^{j-1}} \mathbb{P}(\pi_2 \rightsquigarrow Z_4) \mathbb{P}(\pi_1^{[j-1]} \rightsquigarrow Z_5) (\mathbb{E}\hat{v}(Z_4 \cup \{z^*(Z_4, \pi_2, \Phi)\}, \Phi) - v(Z_4)) \\
= & \sum_{Z_4 \in \mathcal{Z}^{n_2}} \mathbb{P}(\pi_2 \rightsquigarrow Z_4) (\mathbb{E}\hat{v}(Z_4 \cup \{z^*(Z_4, \pi_2, \Phi)\}, \Phi) - v(Z_4)),
\end{aligned}$$

and this ends the proof. \square

Set $\pi_1 = \pi^*$ and $\pi_2 = \text{KG}^{[n-1]}$ in Lemma 4.11 and the above proposition then what left to show is that

$$F^{\text{KG}^{[n]}} - F^{\text{KG}^{[n-1]}} \geq \sum_{Z \in \mathcal{Z}^n} \mathbb{P}(\pi_2 \rightsquigarrow Z) (\mathbb{E}\hat{v}(Z \cup \{z^*(Z, \text{KG}^{[n-1]}, \Phi)\}, \Phi) - v(Z)).$$

From the definition, the left hand side of the last equation:

$$\begin{aligned}
F^{\text{KG}^{[n]}} - F^{\text{KG}^{[n-1]}} &= \sum_{Z_1 \in \mathcal{Z}^{n+1}} \mathbb{P}(\text{KG} \rightsquigarrow Z_1) v(Z_1) - \sum_{Z_2 \in \mathcal{Z}^n} \mathbb{P}(\text{KG} \rightsquigarrow Z_2) v(Z_2) \\
&= \sum_{Z_2 \in \mathcal{Z}^n} \sum_{Z_3 \in \mathcal{Z}^1} \mathbb{P}(\text{KG} \rightsquigarrow Z_2) \mathbb{P}(\text{KG} \rightsquigarrow Z_3 | \text{KG} \rightsquigarrow Z_2) v(Z_2 \cup Z_3) \\
&\quad - \sum_{Z_2 \in \mathcal{Z}^n} \mathbb{P}(\text{KG} \rightsquigarrow Z_2) v(Z_2).
\end{aligned}$$

Now it is enough to show that

$$\begin{aligned}
& \sum_{Z_3 \in \mathcal{Z}^1} \mathbb{P}(\text{KG} \rightsquigarrow Z_3 | \text{KG} \rightsquigarrow Z_2) v(Z_2 \cup Z_3) - v(Z_2) \\
\geq & \mathbb{E}\hat{v}(Z_2 \cup \{z^*(Z_2, \text{KG}^{[n-1]}, \Phi)\}, \Phi) - v(Z_2).
\end{aligned}$$

We could group together the partial realizations ψ that lead to the same single step optimal decision $z^*(Z_2, \text{KG}^{[n-1]}, \Phi)$, and then the last inequality follows from the adaptive greedy nature of the KG policy.

A.4 Proof of Theorem 5.2

First of all, we consider the case when f is a two dimensional function and the four points we pick form a rectangle. Assume $f(x, y)$ is submodular. For any given point (x_0, y_0) , we have $f(x_0+t+s, y_0) - f(x_0+t, y_0) \leq f(x_0+s, y_0) - f(x_0, y_0)$ and $f(x_0+t, y_0) - f(x_0, y_0) \leq f(x_0+t, y_0+s) - f(x_0, y_0+s)$ for any $s, t > 0$. From the first inequality we get $f_{xx}(x_0, y_0) \leq 0$ directly. From the second inequality, we have $f_x(x_0, y_0) \leq f_x(x_0, y_0+s)$, and finally $f_{x,y}(x_0, y_0) \leq 0$. On the other hand, if we have $f_{xy} \leq 0$, $f_{xx} \leq 0$, for any (x, y) , then due to the fact that $f(x_0+t, y_0+s) - f(x_0+t, y_0) - (f(x_0, y_0+s) - f(x_0, y_0)) = \int_{x_0}^{x_0+t} \int_{y_0}^{y_0+s} f_{xy}(u, v) du dv \leq 0$, $f(x_0+t+s, y_0) - f(x_0+t, y_0) - (f(x_0+s, y_0) - f(x_0, y_0)) = stf_{xx}(x_0 + \xi, y_0) \leq 0$, we obtain the submodularity.

We next consider the general case when f is n dimensional and the four points only form a parallelogram. Since the difference between the two marginal values can be decomposed into summation of several marginal value differences whose reference points form rectangles that parallel to coordinate planes, the result for the general case is straightforward from the two dimensional case.

References

- [1] R. AGRAWAL, *Sample mean based index policies with $o(\log n)$ regret for the multi-armed bandit problem*, Adv. in Appl. Probab., (1995), pp. 1054–1078.
- [2] J.-Y. AUDIBERT, S. BUBECK, ET AL., *Best arm identification in multi-armed bandits*, COLT 2010-Proceedings, (2010).
- [3] J.-Y. AUDIBERT, R. MUNOS, AND C. SZEPESVÁRI, *Exploration-exploitation tradeoff using variance estimates in multi-armed bandits*, Theor. Comput. Sci., 410 (2009).
- [4] P. AUER, N. CESA-BIANCHI, AND P. FISCHER, *Finite-time analysis of the multi-armed bandit problem*, Mach. Learn., 47 (2002), pp. 235–256.
- [5] E. BARUT AND W. B. POWELL, *Optimal learning for sequential sampling with non-parametric beliefs*, J. Global Optim., (2013), pp. 1–27.
- [6] J. BRANKE, S. E. CHICK, AND C. SCHMIDT, *Selecting a selection procedure*, Manag. Sci., 53 (2007), pp. 1916–1932.
- [7] S. BUBECK, N. CESA-BIANCHI, AND G. LUGOSI, *Bandits with heavy tail*, arXiv preprint arXiv:1209.1727, (2012).
- [8] O. CAPPÉ, A. GARIVIER, O.-A. MAILLARD, R. MUNOS, G. STOLTZ, ET AL., *Kullback–leibler upper confidence bounds for optimal sequential allocation*, Ann. Statist., 41 (2013), pp. 1516–1541.

- [9] C.-H. CHEN, H.-C. CHEN, AND L. DAI, *A gradient approach for smartly allocating computing budget for discrete event simulation*, in 28th Proc. Winter Simul., IEEE Computer Society, 1996, pp. 398–405.
- [10] C.-H. CHEN, D. HE, AND M. FU, *Efficient dynamic simulation allocation in ordinal optimization*, IEEE Trans. Automat. Control, 51 (2006), pp. 2005–2009.
- [11] C.-H. CHEN, J. LIN, E. YÜCESAN, AND S. E. CHICK, *Simulation budget allocation for further enhancing the efficiency of ordinal optimization*, Discrete Event Dyn. Syst., 10 (2000), pp. 251–270.
- [12] S. E. CHICK, *New two-stage and sequential procedures for selecting the best simulated system*, Oper. Res., 49 (2001), pp. 732–743.
- [13] P. FRAZIER, W. POWELL, AND S. DAYANIK, *The knowledge-gradient policy for correlated normal beliefs*, INFORMS J. Comput., 21 (2009), pp. 599–613.
- [14] P. I. FRAZIER AND W. B. POWELL, *Paradoxes in learning and the marginal value of information*, Decis. Anal., 7 (2010), pp. 378–403.
- [15] 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.
- [16] A. GARIVIER AND E. MOULINES, *On upper-confidence bound policies for non-stationary bandit problems*.
- [17] J. C. GITTINS, *Bandit processes and dynamic allocation indices*, J. R. Stat. Soc. Ser. B. Stat. Methodol., (1979).
- [18] D. GOLOVIN AND A. KRAUSE, *Adaptive submodularity: A new approach to active learning and stochastic optimization.*, in COLT, 2010, pp. 333–345.
- [19] S. S. GUPTA AND K. J. MIESCKE, *Bayesian look ahead one-stage sampling allocations for selection of the best population*, J. Statist. Plann. Inference, 54 (1996), pp. 229–244.
- [20] I. GUTTMAN, G. C. TIAO, ET AL., *A bayesian approach to some best population problems*, Ann. Math. Statist., 35 (1964), pp. 825–835.
- [21] D. HE, S. E. CHICK, AND C.-H. CHEN, *Opportunity cost and OCBA selection procedures in ordinal optimization for a fixed number of alternative systems*, IEEE Trans. Syst. , Man, and Cybern., Part C: Applications and Reviews, 37 (2007), pp. 951–961.

- [22] R. HOWARD, *Information value theory*, IEEE Trans Syst. Sci. and Cybern., 2 (1966), pp. 22–26.
- [23] D. HUANG, T. T. ALLEN, W. I. NOTZ, AND N. ZENG, *Global optimization of stochastic black-box systems via sequential kriging meta-models*, J. Global Optim., 34 (2006), pp. 441–466.
- [24] D. R. JONES, M. SCHONLAU, AND W. J. WELCH, *Efficient global optimization of expensive black-box functions*, J. Global Optim., 13 (1998), pp. 455–492.
- [25] L. P. KAEHLING, *Learning in embedded systems*, 1993.
- [26] A. KAPOOR AND R. GREINER, *Learning and classifying under hard budgets*, Springer, 2005.
- [27] R. KLEINBERG, A. NICULESCU-MIZIL, AND Y. SHARMA, *Regret bounds for sleeping experts and bandits*, Mach. Learn., 80 (2010), pp. 245–272.
- [28] T. L. LAI AND H. ROBBINS, *Asymptotically efficient adaptive allocation rules*, Adv. Appl. Math., 6 (1985), pp. 4–22.
- [29] M. R. MES, W. B. POWELL, AND P. I. FRAZIER, *Hierarchical knowledge gradient for sequential sampling*, J. Mach. Learn. Res., 12 (2011), pp. 2931–2974.
- [30] D. M. NEGOESCU, P. I. FRAZIER, AND W. B. POWELL, *The knowledge-gradient algorithm for sequencing experiments in drug discovery*, INFORMS J. Comput., 23 (2011), pp. 346–363.
- [31] G. L. NEMHAUSER, L. A. WOLSEY, AND M. L. FISHER, *An analysis of approximations for maximizing submodular set functions*, Math. Program., 14 (1978), pp. 265–294.
- [32] W. B. POWELL AND I. O. RYZHOV, *Optimal learning*, vol. 841, 2012.
- [33] H. RAIFFA AND R. SCHLAIFER, *Applied statistical decision theory*, Harvard Business School Publications, (1961).
- [34] N. SRINIVAS, A. KRAUSE, S. M. KAKADE, AND M. SEEGER, *Gaussian process optimization in the bandit setting: No regret and experimental design*, arXiv preprint arXiv:0912.3995, (2009).
- [35] G. J. STIGLER, *The economics of information*, J. political Econ., (1961), pp. 213–225.
- [36] Y. WANG, C. WANG, AND W. POWELL, *The knowledge gradient for sequential decision making with stochastic binary feedbacks*, in 33rd Proc. ICML, 2016.

- [37] F. YOKOTA AND K. M. THOMPSON, *Value of information literature analysis: a review of applications in health risk management*, Med. Decis. Mak., 24 (2004), pp. 287–298.