Proceedings of the 2019 Winter Simulation Conference
N. Mustafee, K.-H.G. Bae, S. Lazarova-Molnar, M. Rabe, C. Szabo, P. Haas, and Y.-J. Son, eds.

A SPECTRAL INDEX FOR SELECTING THE BEST ALTERNATIVE

Guowei Sun

Department of Mathematics
Institute for Systems Research
University of Maryland
College Park, MD 20742, USA

Yunchuan Li

Department of Electrical and Computer Engineering
Institute for Systems Research
University of Maryland
College Park, MD, 20742, USA

Michael Fu

Robert H. Smith School of Business
Institute for Systems Research
University of Maryland
College Park, MD 20742, USA

ABSTRACT

This paper considers the problem of choosing the best design alternative under a small simulation budget where making inferences about all alternatives from a single observation could enhance the probability of correct selection. We propose a new selection rule exploiting the relative similarity information between pairs of alternatives and show its improvement on selection performance, evaluated by the probability of correct selection, compared to selection based on collected sample averages. We illustrate the effectiveness by applying our selection index on simulated ranking and selection problems using two well-known budget allocation policies.

1 INTRODUCTION

In the simulation optimization setting, stochastic ranking and selection (R&S) refers to the problem of selecting the best alternative through simulating samples of the performance measures on a finite set of candidate alternatives (Fu et al. 2008). Given a fixed simulation budget, the goal is to design a selection algorithm to achieve some objective, such as maximizing probability of correct selection (PCS) or minimizing opportunity cost (OC). It is also known as the best arm identification problem in the bandit community. Most of the current research focuses on designing a budget allocation policy to efficiently collect performance samples. A few notable allocation policies include static allocation rules, such as OCBA (optimal computing budget allocation) (Chen and Lee 2010), dynamic allocation rules balancing the exploration-exploitation trade-off using the expected improvement and knowledge gradient (Kim and Nelson 2007; Frazier et al. 2009), and procedures that focus on designing stopping rules for saving resources (Luo et al. 2015). In some scenarios, it makes sense to optimize for different objectives such as opportunity cost or probability of selecting a good subset (Chen et al. 2008). In this work, we focus on probability of correct selection (PCS), which is the most popular criterion. For an overview of ranking and selection problems, we refer the readers to Kim and Nelson (2007) and Branke et al. (2007).

Much of the R&S literature treats the alternatives as independent of each other, and the samples collected are only used to make inferences about properties of the corresponding alternative. Common random numbers used to induce correlation in the performance samples can lead to better pairwise comparison accuracy (see, e.g., Fu et al. (2007)), but the belief in alternative performances are usually still independent. However, in many applications it makes sense to use samples collected on one alternative to infer about

other alternatives. As a motivating example, consider the problem of selecting processing designs for manufacturing semiconductor chips. The values of some governing variables for processing steps such as oxidation, etching, and ion injection must be specified (Koehler and Owen 1996). Two designs with very similar input values on those steps can naturally be expected to have similar performance; and if one of the two is believed to be much inferior compared to some other designs, so should the other.

There are mainly two lines of research trying to address this phenomenon. One is to build a parametric model incorporating the input information into the objective functions and learn the model parameters. One popular choice is to assume the expected performance is a linear function of some input feature, and the simulation budget can be allocated to efficiently learn the feature parameters (Powell and Ryzhov 2012). This approach was applied to a drug discovery problem and showed promising results (Negoescu et al. 2011). Similar approaches were also proposed in the closely related multi-armed bandit setting and achieved good empirical performance even when the linear assumption is not met (Filippi et al. 2010). In this approach, a new performance sample on one alternative could provide information about all alternatives by updating the estimation of parameter values. Another line of research follows the Bayesian approach and treats the similarities between alternatives as correlations in prior beliefs. Frazier et al. (2009) model the similarities as assumed known correlations in a correlated normal prior, allocating the simulation budget using the knowledge gradient policy. Qu et al. (2015) further extended the approach by assuming the correlations are unknown and can be learned using a conjugate Bayesian learning model where the correlations have a Wishart distribution. A related line of work is stochastic kriging, which constructs a metamodel to predict the expected performance on all alternatives using similarity as correlation (Ankenman et al. 2010).

Despite the richness of literature on R&S, an important part of the problem remains relatively neglected: the final selection rule. In Peng et al. (2016), the authors proposed to select the alternative with maximum integrated posterior PCS rather than the one with the maximum sample average. Selection based on quantile estimates has also been proposed largely to avoid sensitivity to outliers and to address a different design objective than PCS (Gibbons et al. 1999). In our work, we assume the sample data are obtained through some sampling allocation policy and focus on how to make a final selection based on the collected data.

Our approach is largely motivated by research outside the operations research and statistics communities. The term *spectral methods* refers to a large family of methods constructing a similarity graph on the entity of interest to bring the vague similarity information into rigorous mathematical formulation to improve decision making. In image processing, a normalized graph cut approach for segmenting objects can be proven to be equivalent to a spectral clustering on the pixel values (Shi and Malik 2000). In semi-supervised learning, a similarity graph is constructed to transfer information on nodes of the graph to make classifications (Zhu 2005). Spectral clustering is a novel clustering approach capable of discovering structures within data that cannot be detected by traditional K-means or K-mods clustering methods (von Luxburg 2007).

We propose a spectral selection index that is a transform of the collected data using a similarity graph constructed from known information about all alternatives. This approach addresses the problem of making inferences about all alternatives from new observations from the perspective of *selection* rather than that of *allocation*. It provides an alternative to the parametric and Bayesian approaches. More interestingly, our approach does not depend on the data collection process, and therefore can be implemented with any existing simulation budget allocation policy. Our approach also has a provable improvement guarantee that is missing in many existing R&S procedures. Numerical experiments show that our approach will give good improvement under mild assumptions on the similarity graph.

This paper is organized as follows. Section 2 motivates our approach with a toy R&S problem. Section 3 introduces relevant graph-theoretic results and formulates our approach. Section 4 establishes a performance guarantee of the proposed selection index. We illustrate the performance of our approach on two synthetic problems in Section 5. Finally, we conclude in Section 6.

2 A TOY MOTIVATING PROBLEM

Consider a toy ranking and selection problem with a fixed allocation policy: the three candidate alternatives have normal random rewards with means $1,0,0$ and the same standard deviation of 10. Each alternative is allocated one simulation replication and the selection rule is to choose the alternative with the maximum observed value. The probability of correctly selecting the first alternative can be computed to be 0.362, only slightly better than a uniform random selection. However, if some prior information indicates that alternatives 2 and 3 should have similar performance, we may construct pairwise similarity between alternatives

$$s_{12} = 0, \quad s_{13} = 0, \quad s_{23} = 1.$$

Let y_1, y_2, y_3 be the obtained samples in this allocation policy. Motivated by the fact that similar alternative should have similar performance, we propose a new index z which minimizes the following expression

$$(z_1 - y_1)^2 + (z_2 - y_2)^2 + (z_3 - y_3)^2 + s_{23} \cdot (z_2 - z_3)^2 + s_{12} \cdot (z_1 - z_2)^2 + s_{13} \cdot (z_1 - z_3)^3,$$

where the first three terms force z_i, $i \leq 3$, to be close to the observed information y_i, $i \leq 3$, and the remaining terms force z_2 and z_3 to be closer as they have a non-zero similarity measure. Minimizing the above objective we obtain $z_1 = y_1$, $z_2 = \frac{2}{3}y_2 + \frac{1}{3}y_3$, $z_3 = \frac{2}{3}y_3 + \frac{1}{3}y_2$, which is a new set of indices that are weighted averages of their similar neighbours. Selecting based on the new index z, the probability of correct selection is found to be 0.472. We call such a smoothed index the *Spectral Index*, as a similarity matrix (or graph) is used for computing the new index to incorporate the known information to facilitate better selection.

3 SIMILARITY GRAPHS IN R&S

Let $\mathscr{A} = \{1, 2, \ldots, k\}$ be the set of alternatives. The reward distributions are unknown with expected values $(\mu_1, ..., \mu_k)$ and standard deviations $(\sigma_1, \ldots, \sigma_k)$. Without loss of generality, we assume that $\mu_1 \geq \mu_2 \geq \cdots \geq \mu_k$. The optimal alternative is defined to be the one with the maximum expected performance and denoted by i^*. Therefore in our setup, $i^* \equiv 1$. In this section, we discuss how to use a *spectral approach* to rigorously incorporate pairwise similarity information into an R&S procedure.

3.1 Graph Notation

Let $S \in R_+^{k \times k}$ be the *similarity matrix* where the element $s_{ij} \geq 0$ denotes the similarity measurement between alternatives i and j, $\forall i,\ j \in \mathscr{A}$. The pairwise similarity information can be represented by a similarity graph $G = (\mathscr{A}, S)$, where vertex i represents alternative i and the edges between vertices i and j are weighted by the similarity between the two connecting nodes s_{ij}. We will later show that the choice of the diagonal elements of S is arbitrary and does not affect the final outcome; therefore for simplicity, we assume that $s_{ii} = 0$, $\forall i \leq k$, meaning all diagonal elements of S (or self similarity) are 0. For the weighted graph G, the degree of vertex $i \in \mathscr{A}$ is defined by

$$d_i = \sum_{j=1}^{k} s_{ij},$$

which is a measurement of the total connectivity of this vertex and represents how similar it is to all other alternatives.

Definition 1 (Degree Matrix) The degree matrix D of the similarity matrix S is a diagonal matrix with degrees $d_1, ..., d_k$ as its diagonal elements.

Definition 2 (Graph Laplacian) The *unnormalized* graph Laplacian matrix is defined as

$$L = D - S.$$

The graph Laplacian is a key concept in spectral methods. Here, we presented the unnormalized graph Laplacian. Other alternatives involve normalizing L using degree matrix D in various ways, such as $L_{sym} = D^{-1/2}LD^{-1/2}$ or $L_{rw} = D^{-1}L$. The normalized Laplacians have been shown to lead to better performance in various tasks, such as spectral clustering (von Luxburg 2007) and semi-supervised learning (Johnson and Zhang 2007) . We use the unnormalized Laplacian, since it will provide an intuitive explanation to how similarity would affect the selection procedure. Notice that the diagonal elements of S will cancel out in the computation of L. In our approach this means that self-similarity will have no effect on the final outcome.

Lemma 1 (Proposition 1 of von Luxburg (2007)) The graph Laplacian L has two important properties:

1. L is symmetric and positive semi-definite.
2. The smallest eigenvalue of L is 0.

We refer readers to von Luxburg (2007) for the detailed proof.

3.2 Computing Similarity Graphs

Similar to *spectral methods*, another technique that uses a similarity matrix to represent the affinity on a finite set of entities is kriging or Gaussian process regression. Both are Bayesian approaches assuming a correlated normal prior with correlations determined by a similarity or kernel matrix. The posteriors are updated using observed samples. In this part, we introduce some basic methods for constructing similarity graphs.

The first type of matrix is based on expert knowledge. A similarity score s_{ij} can be manually assigned for all pairs if k is moderate in size. Such examples appear in linguistics (Downey et al. 2017) and gene enrichment analysis (Efron 2012).

More generally, the similarity could be computed from some features about the alternatives. Let $x(i)$ be an m-dimensional feature vector describing the characters of alternative i. In our motivating example, $x(i)$ could be a vector representing oxidization rates, etching time length and ion injection density for a manufacturing design. Some popular choices are the ε, Gaussian and exponential similarity graphs, which can be computed respectively by:

1. ε-graph:
$$s_{ij} = \delta \mathbb{1}_{\{\|x(i)-x(j)\|_2 \leq \varepsilon\}},$$
where $\mathbb{1}_{\{\cdot\}}$ is the indicator function.
2. Gaussian graph:
$$s_{ij} = e^{-\sum_{n=1}^{m} \theta_n (x_n(i)-x_n(j))^2}.$$
3. Exponential graph:
$$s_{ij} = e^{-\sum_{n=1}^{m} \beta_n \|x_n(i)-x_n(j)\|_2}.$$

The $\delta, \varepsilon, \theta, \beta$ are all input hyper-parameters determining the magnitude of the similarity. The ε-graph will return a sparse similarity matrix S, meaning one alternative is only connected with its few closest neighbors on the graph. The sparsity will also lead to faster numerical computations. Though a feature vector is used to describe each alternative and compute their similarities, there is no assumption on the parametric form of the expected performance relating to the features. In Section 5, the proposed graphs will be evaluated through simulation.

4 SPECTRAL SELECTION INDEX

We assume the allocation is given by a policy Π. Let y_{ij}, $\forall i \in \mathscr{A}$, $j \leq n_i$ denote the samples collected in the simulation process, where n_i is the number of simulations allocated to alternative i after exhausting the

total budget N. Denote $\bar{\mathbf{y}}$ as the vector of sample averages $(\bar{y}_1,...,\bar{y}_k)$. Instead of choosing the alternative with maximum sample average, we propose a spectral index $\mathbf{z} = (z_1,\ldots,z_k)$ as our selection criteria. Use $i_N^{\bar{\mathbf{y}}}$ and $i_N^{\mathbf{z}}$ to denote the final selected alternative given all observed data using the sample average and our proposed selection index, respectively.

4.1 Smooth Index on Similarity Graph

The usage of sample averages as selection criteria is intuitive, since they generally provide an unbiased estimator of the true expected performance. However, as our goal is to maximize PCS, comparing *relative* performance is more critical than finding accurate estimates. This is also the basis of ordinal optimization (Ho et al. 2008; Chen et al. 2013).

Given a similarity graph G and motivated by the fact that two similar alternatives should have similar selection index, we try to find an index $\mathbf{z} : (\mathscr{A}, \bar{\mathbf{y}}) \to R^k$ that is a smooth function on G. Similar to the approach in semi-supervised learning (Zhu 2005), we propose a spectral index $\mathbf{z}$ that is the solution to the following optimization problem:

$$\min_{\mathbf{x}\in R^k} \sum_{i=1}^{k} |x_i - \bar{y}_i|^p + \frac{\lambda}{2} \sum_{1\le i,j\le k} s_{ij}|x_i - x_j|^q. \tag{1}$$

The first term forces x_i to be close to the sample averages $\bar{y}_i$, and the second term forces two alternatives with larger similarity to have closer index values. λ is a positive regularization coefficient that controls the weight between the two terms. p,q are integers specifying the norms to use when enforcing the smoothness. Kyng et al. (2015) provide algorithms for solving such an optimization problem with various choices of p,q. In this work, we set $p = q = 2$, both for computational convenience and for developing intuitive explanations. We have the following theorem for compact representation of the optimization problem.

Theorem 1 When $p = q = 2$ in Equation (1), the optimization problem can be expressed with the graph Laplacian matrix L as

$$\mathbf{z} = \arg\min_{\mathbf{x}\in\mathbb{R}^k} (\mathbf{x} - \bar{\mathbf{y}})' (\mathbf{x} - \bar{\mathbf{y}}) + \lambda \mathbf{x}' L \mathbf{x}, \tag{2}$$

where x' denotes the transpose of x.

Proof. We can write the second term in Equation (2) using summations as

$$\begin{aligned}\mathbf{x}' L \mathbf{x} &= \sum_{i,j\le k} x_i L_{ij} x_j = \frac{1}{2}\left\{\sum_{i\le k} L_{ii}x_i^2 + \sum_{j\le k} L_{jj}x_j^2 + \sum_{i\ne j, i,j\le k} 2L_{ij}x_i x_j\right\} \\ &= \frac{1}{2}\left\{\sum_{i\le k} d_i x_i^2 + \sum_{j\le k} d_j x_j^2 - \sum_{i\ne j} 2s_{ij}x_i x_j\right\} = \sum_{i,j\le k} s_{ij}(x_i - x_j)^2.\end{aligned}$$

The $(\mathbf{x} - \bar{\mathbf{y}})'(\mathbf{x} - \bar{\mathbf{y}})$ terms matches trivially with the first term in Equation (1). □

Taking the derivative of the optimization objective and setting it to zero yields

$$\mathbf{z} = (I + \lambda L)^{-1} \bar{\mathbf{y}}, \tag{3}$$

where I is the identity matrix of rank k.

Proposition 1 $I + \lambda L$ is symmetric and positive definite with minimum eigenvalue 1.

Proof. This is a direct consequence of Lemma 1 . □

Proposition 1 implies that our approach has stable numerical properties, as the condition number of our linear system computation will not be too large.

The computation of $\mathbf{z}$ involves taking the inverse of the matrix $I+\lambda L$, which could be expensive for problems with a large number of alternatives. The fact that the minimum eigenvalue of the inverted matrix is 1 usually means the condition number of the inverse computation will not be too large, giving stable computing results. The solution in Equation (3) can also be obtained using an iterative approach. Denote the t-th iterate by $\mathbf{z}^{(t)}$ and its i-th element by $z_i^{(t)}$. We propose Algorithm 1 to calculate $\mathbf{z}$ iteratively.

Algorithm 1 (Iterative Updates) The spectral index z can be computed by:

1. Set $\mathbf{z}^{(0)} = \bar{\mathbf{y}}$,
2. For $t \leq T$, $\mathbf{z}^{(t)} = \mathbf{z}^{(t-1)} + a\left(\bar{\mathbf{y}} - (I+\lambda L)\mathbf{z}^{(t-1)}\right)$,

where a is a fixed stepsize.

Theorem 2 (Convergence of Algorithm 1) Let δ be the largest eigenvalue of L, then we have $\mathbf{z}^{(t)} \to \mathbf{z}$ as $t \to \infty$ if $0 < a < \frac{1}{|1+\lambda\delta|}$.

Proof. By Proposition 1, $\mathbf{z}$ defined in Equation (3) is unique. Let $\mathbf{r}^{(t)}$ be the residual at the t-th iteration in Algorithm 1, i.e.,

$$\mathbf{r}^{(t)} = (I+\lambda L)\mathbf{z}^{(t)} - \bar{\mathbf{y}}.$$

It suffices to prove that $\|\mathbf{r}^{(t)}\| \to \mathbf{0}$. We have

$$\begin{aligned}\mathbf{r}^{(t+1)} &= (I+\lambda L)\mathbf{z}^{(t+1)} - \bar{\mathbf{y}} = (I+\lambda L)\mathbf{z}^{(t)} - \bar{\mathbf{y}} - a(I+\lambda L)((I+\lambda L)\mathbf{z}^{(t)} - \bar{\mathbf{y}}) \\ &= (I - a(I+\lambda L))\,\mathbf{r}^{(t)} = (I - a(I+\lambda L))^{t+1}\,\mathbf{r}^{(0)}.\end{aligned}$$

By Proposition 1 and the assumption that $0 < a < \frac{1}{|1+\lambda\delta|}$, the eigenvalues of the matrix $I - a(I+\lambda L)$ are less than 1. Therefore,

$$\|\mathbf{r}^{(t+1)}\| \leq \alpha^{t+1}\|\mathbf{r}\|^{(0)},$$

where α is the largest eigenvalue of $I - a(I+\lambda L)$, and we have $0 < \alpha < 1$. Thus, $\|\mathbf{r}^{(t)}\| \to 0$, i.e., $\mathbf{r}^{(t)} \to \mathbf{0}$. □

Remark 1 This approach for solving z will not only provide an alternative to the matrix inversion approach, but also provide a mechanism for proving performance guarantees of our method. Notice that in each time step t, the update for an alternative $i \in \mathscr{A}$ is

$$z_i^{(t+1)} = a\bar{y}_i + (1-a)z_i^{(t)} + a\lambda \sum_{j \neq i} s_{ij}\left(z_j^{(t)} - z_i^{(t)}\right). \tag{4}$$

The update is discounting the observed information $\bar{y}_i$ and mixing in information from alternative j according to the similarity s_{ij}, the regularization coefficient λ and update step size a. Algorithm 1 and Theorem 2 lead to the following result.

Corollary 1 (Weighted Averages of Observed Information) The index $\mathbf{z}$ satisfies the following weighted averaging property

$$z_i = \frac{\bar{y}_i + \lambda \sum_{j \neq i} s_{ij} z_j}{1 + \lambda d_i}. \tag{5}$$

Proof. In Equation (4), as $t \to \infty$, we know that $z_i^{(t+1)} = z_i^{(t)} = z_i, \forall i \leq k$ from Theorem 2. Rearranging the terms yields

$$a(1+\lambda d_i)z_i = a\bar{y}_i + a\lambda \sum_{j \neq i} s_{ij} z_j,$$

which is equivalent to Equation (5). □

Corollary 1 will be the key to our performance proofs. The expression also provides two key intuitions on the z index: (1) a vertex with a larger degree d_i will be less affected by its actual observations $\bar{y}_i$, (2) the weighted averaging is the source of potential *PCS* gain: a sub-optimal alternative with unusually high observed performance could be "dragged down" by its neighbours on the graph. In the next part, we define a class of graphs that will lead to PCS improvement.

4.2 Performance Improvement

With z computed either from Equation (3) or Algorithm 1, we compare the following two rules for making the final selection:

1. Using sample averages $\bar{\mathbf{y}}$: $i_N^{\bar{\mathbf{y}}} = \arg\max_{i \in \mathscr{A}} \{\bar{y}_1, ..., \bar{y}_k\}$
2. Using spectral index $\mathbf{z}$: $i_N^{\mathbf{z}} = \arg\max_{i \in \mathscr{A}} \{z_1, ..., z_k\}$

Let $PCS(\bar{\mathbf{y}})$ and $PCS(\mathbf{z})$ denote the PCS for each respective selection rule.

Definition 3 (Aligned Graph) A graph G is aligned if the similarities $\{s_{ij}\}_{i,j \in \mathscr{A}}$ are monotonically decreasing with $|\mu_i - \mu_j|, \forall i \neq j; i, j \in \mathscr{A}$ and $d_i = d_1, \forall i$.

Though the similarities are functions of alternative feature vectors $x(i), i \in \mathscr{A}$, rather than expected performances $\mu_i, i \in \mathscr{A}$, we can still compare the similarity score between alternatives and their true underlying gap of expected performance. An aligned graph would use the prior known feature information $x(i)$ to correctly capture the relative closeness of alternatives: if i and j have smaller gap in expected performance, their similarity would be greater. It also requires the degree for all vertices to be the same, which could be achieved by normalizing an existing graph Laplacian L (von Luxburg 2007). It may be a strong assumption to expect that such a graph can be constructed without knowing the true performance, but they will provide nice theoretical results and shed light on the intuition of our approach. The family of graphs that are aligned will have provable nice performance improvement.

Theorem 3 (Order Preserving Updates) For an aligned graph G and performance sample averages with correct ordering, i.e., $\bar{y}_1 \geq \bar{y}_2 \geq \cdots \geq \bar{y}_k \geq 0$, the spectral index defined in Equation (3) preserves correct ordering, i.e., $z_1 \geq z_2 \geq, ..., \geq z_k \geq 0$.

Proof. We establish the proof by induction.

At iteration 0, $\mathbf{z}^{(0)} = \bar{\mathbf{y}}$. By assumption, we have $z_1^{(0)} \geq z_2^{(0)} \geq, ..., \geq z_k^{(0)} \geq 0$.

At iteration $t > 0$, assume $z_1^{(t-1)} \geq z_2^{(t-1)} \geq, \ldots, \geq z_k^{(t-1)} \geq 0$ holds. To prove Theorem 3, using Algorithm 1, we only need to prove that $z_i^{(t)} \geq z_{i+1}^{(t)}$ holds $\forall 1 \leq i \leq k-1$. Rearranging Equation (4), we write out z_i^t and z_{i+1}^t as

$$z_i^{(t)} = a\bar{y}_i + (1 - a - a\lambda d_i) z_i^{(t-1)} + a\lambda \sum_{m \neq i, i+1} \left(s_{i,m} z_m^{(t-1)} \right) + a\lambda s_{i,i+1} z_{i+1}^{(t-1)},$$

$$z_{i+1}^{(t)} = a\bar{y}_{i+1} + (1 - a - a\lambda d_{i+1}) z_{i+1}^{(t-1)} + a\lambda \sum_{m \neq i, i+1} \left(s_{i+1,m} z_m^{(t-1)} \right) + a\lambda s_{i,i+1} z_i^{(t-1)}.$$

Under the condition that $d_i = d_{i+1} = d \ \forall i$, taking the difference yields

$$z_i^{(t)} - z_{i+1}^{(t)} = a(\bar{y}_i - \bar{y}_{i+1}) + (1 - a - a\lambda d - a\lambda s_{i,i+1})(z_i^{(t-1)} - z_{i+1}^{(t-1)}) + a\lambda \left(\sum_{m \neq i,i+1} s_{i,m} z_m^{(t-1)} - \sum_{m \neq i,i+1} s_{i+1,m} z_m^{(t-1)} \right).$$

To show that $z_t^{(t)} - z_{i+1}^{(t)} \geq 0$, under the condition that $\bar{y}_i \geq \bar{y}_{i+1}$, $z_i^{(t-1)} \geq z_{i+1}^{(t-1)}$ and a is sufficiently small, we only need to show that

$$\sum_{m \neq i,i+1} s_{i,m} z_m^{(t-1)} - \sum_{m \neq i,i+1} s_{i+1,m} z_m^{(t-1)} \geq 0.$$

Intuitively it makes sense as the former is a weighted average with higher weights placed on larger terms of $z_m^{(t-1)}$ with an aligned graph. With $d_i = d_{i+1}$ and $s_{i,i+1} = s_{i+1,i}$, we know that $\sum_{m \neq i,i+1} s_{i,m} = \sum_{m \neq i,i+1} s_{i+1,m}$, which could be written as

$$\sum_{m<i} (s_{i,m} - s_{i+1,m}) = \sum_{m>i+1} (s_{i+1,m} - s_{i,m}) \geq 0.$$

Then, using the assumption that $z_1^{(t-1)} \geq z_2^{(t-1)} \geq \cdots \geq z_k^{(t-1)}$, we have

$$\sum_{m<i} (s_{i,m} - s_{i+1,m}) z_m^{(t-1)} \geq \sum_{m<i} (s_{i,m} - s_{i+1,m}) z_{i-1}^{(t-1)} \geq \sum_{m>i+1} (s_{i+1,m} - s_{i,m}) z_{i+2}^{(t-1)} \geq \sum_{m>i+1} (s_{i+1,m} - s_{i,m}) z_m^{(t-1)}.$$

Rearranging the terms would yield the result

$$\sum_{m \neq i,i+1} s_{i,m} z_m^{(t-1)} > \sum_{m \neq i,i+1} s_{i+1,m} z_m^{(t-1)}.$$

We now conclude that $z_i^{(t)} > z_{i+1}^{(t)}$. □

In the proof of Theorem 3 we assumed the step size a in Algorithm 1 is sufficiently small, meaning during such implementations, the correct ordering will be preserved during all update steps. Theorem 3 establishes the correctness of spectral selection in cases with $n_i \to \infty$ and $\bar{y}_i \to \mu_i$, where selecting using z will return the true optimal alternative. In cases where sample averages are incorrectly ordered, we expect the extra information from the aligned graph will correct the ordering in the final spectral indices. Therefore, our approach should work at least as well as sample averages, and we have the following conjectures.

Conjecture 1 (Selection Fixing) With an aligned graph, $P\{i_N^z = 1 | i_N^{\bar{y}} \neq 1\} > 0$ if all reward distributions have unbounded support.

If the sample average for a sub-optimal alternative is the largest among all averages due to random sampling, its neighbors on the aligned similarity graph G will be able to negate the random error, thus making the spectral index more robust. This is the statistical intuition behind our approach: the spectral index z is a smoothed version of sample averages $\bar{y}$ using a similarity graph with better robustness.

Conjecture 2 (PCS Improvement) With an aligned graph, $PCS(z) \geq PCS(\bar{y})$.

In circumstances where an aligned graph is not available, such as due to lack of information, we still expect the spectral index to improve PCS, i.e., an aligned graph is not necessary to achieve a better selection performance.

Conjecture 3 (Unaligned Graph) For a given R&S problem with μ,σ and simulation budget N, for any budget allocation policy Π, there exist unaligned graphs G such that $PCS(z) \geq PCS(\bar{y})$.

We validate Conjecture 1 to Conjecture 3 in our numerical experiments.

4.3 Information Collection on Graph

The constructed similarity graph motivates us to compute a selection index that utilizes the relative similarity information and given performance samples. However, it is expected that a graph (possibly unaligned) should also provide insight on how to allocate the simulation resources. A vertex $i \in \mathscr{A}$ with a higher degree d_i is *better connected* on the graph; therefore, one sample on this alternative would provide more information compared to a sample on a vertex with a small degree on the graph. This issue is partially addressed as active learning on graphs (Zhu 2005) or experimental design for kriging approaches (Koehler and Owen 1996). It is possible to develop variants of the Knowledge Gradient or Expected Improvement dynamic allocation policies to consider the similarity graph on alternatives. We defer this to future work.

5 NUMERICAL EXPERIMENTS

We test our approach on two synthetic R&S problems implemented with both the equal allocation and OCBA allocation policy. The final selection is made using both the sample averages and our proposed index, and the PCS is computed for each policy. When implementing the OCBA policy, the allocation is computed with the known parameters μ_i and σ_i, as our goal is to illustrate the benefit of employing the spectral index for selection given an allocation policy.

In the first numerical experiment, we simulate a problem with the so-called *least favorable configuration* and test our spectral index using a manually constructed aligned graph. The effect of different values of hyper-parameter λ is also illustrated. The second simulation is performed using a commonly used test problem, and we show that non-aligned graphs could give a spectral index that outperforms selection based on sample averages.

5.1 Least Favorable Configuration

We refer to the setting where all suboptimal alternatives have identical expected performance as the *least favorable configuration*, in the sense that none of the suboptimal ones can be easily identified (Fu et al. 2008). We simulate 5 alternatives each with i.i.d. normal reward samples with mean $\mu_i = I_{\{i=1\}}$ and $\sigma_i = 4, \forall i \leq 5$, where I denotes the indicator function. For a total simulation budget N ranging from 50 to $1,500$, we use both the equal allocation and the OCBA policy to allocate N and compute the PCS over $10,000$ simulation replications. The similarity graph is constructed as

$$s_{ij} = \begin{cases} 1, & i \neq j, 2 \leq i, j \leq 5 \\ 0, & \text{otherwise} \end{cases}. \tag{6}$$

This similarity graph reflects the belief that the 4 sub-optimal alternatives are closer to each other, whereas the optimal one is somewhat isolated. We test our approach using λ values of $0.1, 0.5, 1, 2$. Notice that the graph is an aligned graph according to Definition 3; thus z is guaranteed to give a better *PCS*.

In Figure 1, we can see that the spectral index indeed dominates selection using sample averages for both policies over all simulation budgets tested. With an aligned graph, larger values of λ means stronger belief in the relative relationships in alternatives and indeed leads to better performance.

5.2 Non-Aligned Graph

For more complicated R&S problems, it might be difficult to construct an aligned graph as in the first simulation experiment. However, it is often reasonable to believe the expected performances are smooth with respect to (w.r.t.) to some features.

Consider a problem with 10 alternatives with normal rewards, with means $\mu_i = \frac{1}{4}(i - 5.75)^2, \forall i \in \{1, 2, ..., 10\}$ and a common variance 10. We test the following similarity graphs:

1. *ε-Graph*: $s_{ij} = 0.4 \times \mathbb{1}_{|i-j| \leq 3}$. Only consider the effect of very close neighbours.

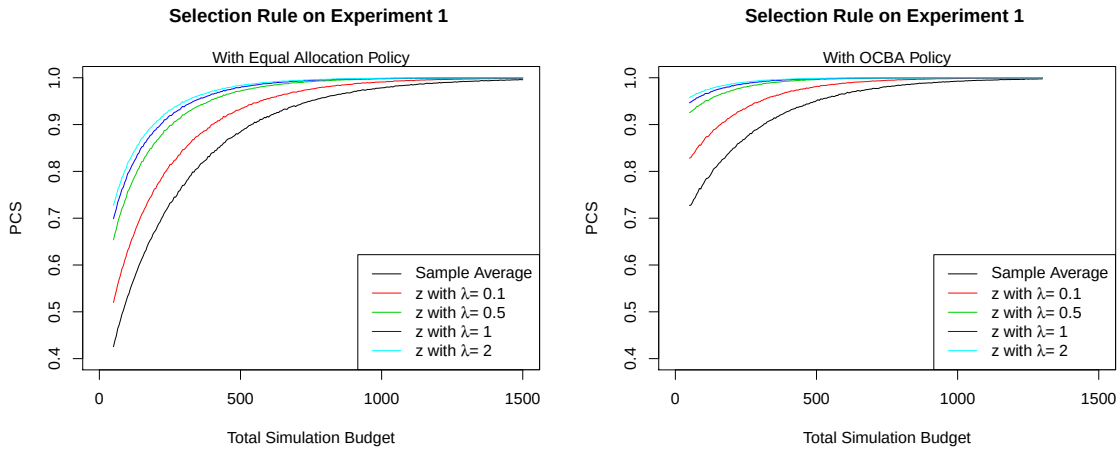

Figure 1: PCS for equal allocation and OCBA policies.

2. *Exponential Graph*: $s_{ij} = e^{-|i-j|}, \forall i, j \leq 10$
3. *Gaussian Graph*: $s_{ij} = e^{-(i-j)^2}, \forall i, j \leq 10$

The configuration of the test problem is shown in Figure 2 . On the left, we can see that the expected performance value μ is a smooth function of the alternative index, mimicking a smooth objective function w.r.t. some input features. On the right, we present the similarities between the optimal alternative with all other alternatives computed using the three graphs and present it with the gap in expected performance. The similarity is not monotonically decreasing; therefore, none of the three graphs is aligned.

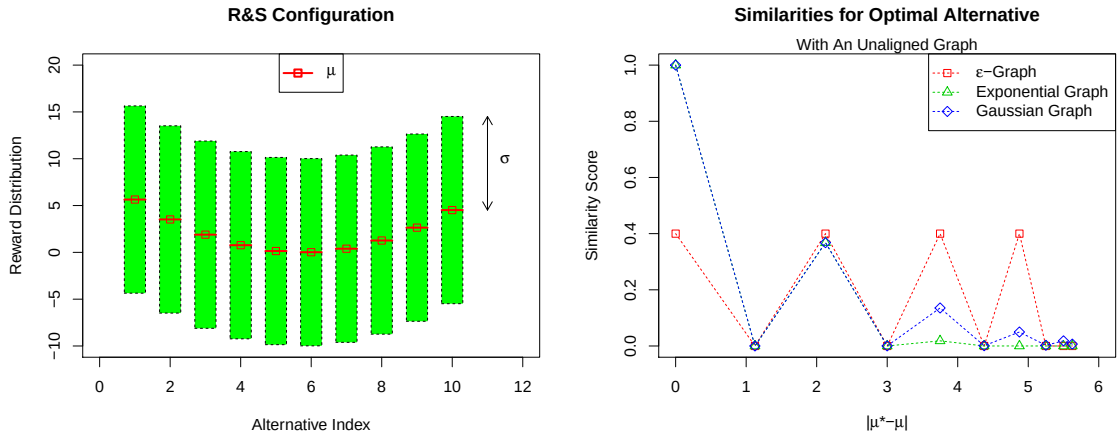

Figure 2: R&S test with non-aligned graphs.

The *PCS* is estimated based on $1{,}000$ simulation replications for budgets ranging from 50 to $3{,}000$ with both equal allocation and OCBA allocation policy. We test our approach using $\lambda = 0.2$ and 0.3. Though none of the three tested graphs are aligned, they still improve the selection performance for both allocation policies, as shown in Figure 3.

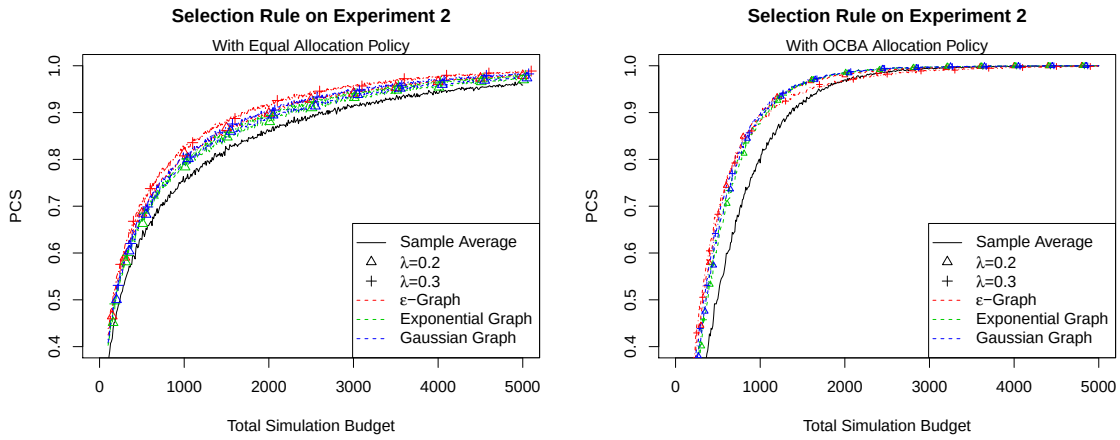

Figure 3: PCS with non-aligned graphs.

6 CONCLUDING REMARKS

We proposed a performance index that incorporates the similarity information between pairs of alternatives in R&S and proved that the proposed index will give a better selection performance if reasonable similarity information is available. Numerical experiments showed promising results on both theoretically good problem instances and more general problem instances. The main contribution is to provide an easy way for making inferences about all alternatives with samples from one alternative. The proposed approach is very intuitive and computationally trivial compared to many proposed Bayesian approaches. The proposed index does not affect the allocation policy procedures and therefore could be easily combined with any allocation policy, as illustrated with our numerical experiments.

ACKNOWLEDGMENTS

This work was supported in part by the National Science Foundation under Grant CMMI-1434419 and by the Defense Advanced Research Projects Agency (DARPA) under Grant N660011824024. The views, opinions and/or findings expressed are those of the authors and should not be interpreted as representing the official views or policies of the Department of Defense or the U.S. Government.

REFERENCES

Ankenman, B., B. L. Nelson, and J. Staum. 2010. "Stochastic Kriging for Simulation Metamodeling". *Operations Research* 58(2):371–382.

Branke, J., S. E. Chick, and C. Schmidt. 2007. "Selecting a Selection Procedure". *Management Science* 53(12):1916–1932.

Chen, C.-H., D. He, M. Fu, and L. H. Lee. 2008. "Efficient Simulation Budget Allocation for Selecting an Optimal Subset". *INFORMS Journal on Computing* 20(4):579–595.

Chen, C.-H., Q.-S. Jia, and L.-H. Lee. 2013. *Stochastic Simulation Optimization for Discrete Event Systems: Perturbation Analysis, Ordinal Optimization and Beyond*. New York City, New York: World Scientific.

Chen, C.-h., and L. H. Lee. 2010. *Stochastic Simulation Optimization: An Optimal Computing Budget Allocation*. 1st ed. River Edge, New Jersey: World Scientific.

Downey, S., G. Sun, and P. Norquest. 2017. "alineR: an R Package for Optimizing Feature-Weighted Alignments and Linguistic Distances". *The R Journal* 9(1):138–152.

Efron, B. 2012. *Large-scale Inference: Empirical Bayes Methods for Estimation, Testing, and Prediction*, Volume 1. Cambridge, UK: Cambridge University Press.

Filippi, S., O. Cappe, A. Garivier, and C. Szepesvári. 2010. "Parametric Bandits: The Generalized Linear Case". In *Advances in Neural Information Processing Systems 23*, edited by J. D. Lafferty, C. K. I. Williams, J. Shawe-Taylor, R. S. Zemel, and A. Culotta, 586–594. Vancouver, Canada: Curran Associates, Inc.

Frazier, P., W. Powell, and S. Dayanik. 2009. "The Knowledge-Gradient Policy for Correlated Normal Beliefs". *INFORMS Journal on Computing* 21(4):599–613.

Fu, M. C., C.-H. Chen, and L. Shi. 2008. "Some Topics for Simulation Optimization". In *Proceedings of the 2008 Winter Simulation Conference*, edited by S. Mason, R. Hill, L. Moench, and O. Rose, 27–38. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.

Fu, M. C., J.-Q. Hu, C.-H. Chen, and X. Xiong. 2007. "Simulation Allocation for Determining the Best Design in The Presence of Correlated Sampling". *INFORMS Journal on Computing* 19(1):101–111.

Gibbons, J. D., I. Olkin, and M. Sobel. 1999. *Selecting and Ordering Populations: A New Statistical Methodology*. Society for Industrial and Applied Mathematics.

Ho, Y.-C., Q.-C. Zhao, and Q.-S. Jia. 2008. *Ordinal Optimization: Soft Optimization for Hard Problems*. New York City, New York: Springer Science & Business Media.

Johnson, R., and T. Zhang. 2007, December. "On the Effectiveness of Laplacian Normalization for Graph Semi-supervised Learning". *Journal of Machine Learning Research* 8:1489–1517.

Kim, S.-H., and B. L. Nelson. 2007. "Recent Advances in Ranking and Selection". In *Proceedings of the 2007 Winter Simulation Conference*, edited by S. Henderson, B. Biller, M. hua Hsieh, and J. Shortle, 162–172. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.

Koehler, J., and A. Owen. 1996. "Computer Experiments". *Handbook of Statistics* 13:261–308.

Kyng, R., A. Rao, S. Sachdeva, and D. A. Spielman. 2015. "Algorithms for Lipschitz Learning on Graphs". In *Proceedings of The 28th Conference on Learning Theory*, edited by P. Grünwald, E. Hazan, and S. Kale, 1190–1223. Paris, France.

Luo, J., L. J. Hong, B. L. Nelson, and Y. Wu. 2015. "Fully Sequential Procedures for Large-Scale Ranking-and-Selection Problems in Parallel Computing Environments". *Operations Research* 63(5):1177–1194.

Negoescu, D. M., P. I. Frazier, and W. B. Powell. 2011. "The Knowledge-Gradient Algorithm for Sequencing Experiments in Drug Discovery". *INFORMS Journal on Computing* 23(3):346–363.

Peng, Y., C.-H. Chen, M. C. Fu, and J.-Q. Hu. 2016. "Dynamic Sampling Allocation and Design Selection". *INFORMS Journal on Computing* 28(2):195–208.

Powell, W. B., and I. O. Ryzhov. 2012. *Optimal Learning*, Volume 841. Hoboken, New Jersey: John Wiley & Sons.

Qu, H., I. O. Ryzhov, M. C. Fu, and Z. Ding. 2015. "Sequential Selection with Unknown Correlation Structures". *Operations Research* 63(4):931–948.

Shi, J., and J. Malik. 2000. "Normalized Cuts and Image Segmentation". *IEEE Transactions on Pattern Analysis and Machine Intelligence* 22(8):888–905.

von Luxburg, U. 2007. "A Tutorial on Spectral Clustering". *Statistics and Computing* 17(4):395–416.

Zhu, X. 2005. "Semi-Supervised Learning Literature Survey". Technical Report 1530, Computer Sciences, University of Wisconsin-Madison, Madison, Wisconsin.

AUTHOR BIOGRAPHIES

GUOWEI SUN is a Ph.D. candidate in Applied Mathematics, Statistics and Scientific Computation at the Department of Mathematics at the University of Maryland, College Park. His email is gwsun@math.umd.edu.

YUNCHUAN LI is a Ph.D. candidate in the Department of Electrical and Computer Engineering at the University of Maryland, College Park. His research focuses on optimization and control with applications to operations research problems. His email is yli93@umd.edu.

MICHAEL C. FU holds the Smith Chair of Management Science in the Robert H. Smith School of Business, with a joint appointment in the Institute for Systems Research, A. James Clark School of Engineering, University of Maryland, College Park. He has a Ph.D. in applied math from Harvard and degrees in math and EECS from MIT. He attended his first WSC in 1988 and served as WSC2011 Program Chair, NSF Operations Research Program Director, *Management Science* Stochastic Models and Simulation Department Editor, and *Operations Research* Simulation Area Editor. He is a Fellow of **INFORMS** and **IEEE**. His e-mail address is mfu@umd.edu, and his Web page is https://www.rhsmith.umd.edu/directory/michael-fu.