

FIXED CONFIDENCE RANKING AND SELECTION UNDER INPUT UNCERTAINTY

Di Wu
Enlu Zhou

School of Industrial & Systems Engineering
Georgia Institute of Technology
755 Ferst Drive NW
Atlanta, GA 30332, USA

ABSTRACT

In stochastic simulation, input uncertainty (IU) is caused by the error in estimating input distributions using finite real-world data. When it comes to simulation-based Ranking and Selection (R&S), ignoring IU can lead to the failure of many existing procedures. In this paper, we study a new version of the fixed confidence R&S problem, where sequential input data can be acquired to reduce IU over time. To solve the problem, we first propose a moving average estimator for online estimation with sequential data. Then, a new procedure is designed by extending a Sequential Elimination framework. As is shown numerically, our procedure can effectively achieve the desired probability of correct selection, but there is plenty of room for improving its efficiency.

1 INTRODUCTION

In building a stochastic simulation, a set of input distributions need to be specified as an input to the simulation model. Then, random samples are drawn from the input distributions to mimic the uncertainty in the real-world system. As a result, there are two sources of uncertainty to account for when we use simulation to evaluate a system design's expected performance. One is the input uncertainty (IU) caused by the error in estimating the input distributions using finite real-world data; the other is the stochastic uncertainty (SU), which stems from using finite simulation runs to evaluate a design's performance. The fundamental difference between IU and SU is that, while SU can be controlled by increasing simulation effort, the only way to reduce IU is to collect more input data.

In simulation optimization, Ranking and Selection (R&S), or Discrete Optimization via Simulation (DOvS), refers to using simulation to identify the best design among a finite number of alternatives. Traditionally, R&S has been studied under the assumption that the input distributions are known accurately, and SU is the only source of uncertainty. One widely studied formulation of R&S is the fixed confidence R&S problem, where the goal is to achieve certain target selection guarantee (e.g., probability of correct selection) by using as little simulation effort as possible. In this setting, considerable amount of effort has been devoted to the indifference-zone (IZ) formulation, which dates back at least to Bechhofer (1954). An IZ procedure allows the user to specify the smallest difference in performance worth detecting, and it guarantees selecting the best design with (frequentist) probability higher than a prespecified level (e.g., 95%), provided that the difference between the top-two designs is sufficiently large. Numerous efficient IZ procedures have been proposed in simulation literature, including but are not limited to the KN procedure in Kim and Nelson (2001), the KVP and UVP procedures in Hong (2006), and the BIZ procedure in Frazier (2014). We refer the reader to Branke et al. (2005) and Kim and Nelson (2007) for excellent reviews of the development on this topic. In addition, the Bayesian approaches (see, e.g., Chick and Frazier (2012))

and the probably approximately correct (PAC) selection (see, e.g., Ma and Henderson (2017)) have also been studied in this stream of works.

In practice, directly applying traditional R&S procedures while ignoring IU can be misleading and may render selection guarantees invalid (see Song et al. (2015)). In light of such observations, recent effort has been made to account for IU when the input dataset is given and fixed. For instance, Corlu and Biller (2015) take a Bayesian perspective and select the design with the best performance averaged over the posterior distribution of input models; Song et al. (2015) and Corlu and Biller (2013) consider a fixed confidence formulation under an Indifference-Zone (IZ) setting, and both discover that a larger IZ parameter is required to maintain the desired statistical guarantee under IU; Fan et al. (2013) take a distributionally robust approach by comparing the designs based on their worst-case performance over a finite set of possible input distributions. Aside from selection procedures, Song and Nelson (2019) propose a comparison procedure which exploits the common input distribution effect to construct simultaneous confidence intervals for all designs' performance.

The aforementioned works all assume a static (fixed) input dataset, because in some applications the cost of collecting additional data is prohibitively high (e.g., collecting experimental data by running clinical trials). However, there are also cases where new data can be accessed at a reasonable pace and cost. Some motivating examples are as follows.

- *Ride-hailing platform.* A ride hailing platform wants to use simulation to find the best pricing strategy. Some of the input models about traveling time, user behavior, route selection, etc., can be updated using real-time data.
- *Online retailer.* An online retailer is interested in optimizing its supply chain network by comparing different ordering/inventory/fulfillment strategies through simulation. Upstream and downstream customer demand data can be collected in a timely manner to refine the demand distribution estimates.
- *Express delivery service.* An express delivery service provider intends to compare different routing and scheduling strategies via a simulation model. The input data on customer demand, traveling time, deliveryman's preference can be harvested on a daily basis.

In all of the above examples, a moderate amount of new data can be collected at a relatively low cost to continuously improve the accuracy of the simulation model.

2 PROBLEM FORMULATION

2.1 Basic Notations

Without loss of generality, the R&S problem studied in this paper is concerned with identifying the design with the highest expected performance among $K \geq 2$ alternatives. Denote by $\mathcal{J} := \{1, 2, \dots, K\}$ the enumeration of all designs. For a design $i \in \mathcal{J}$, let $h_i : \mathbb{R}^m \rightarrow \mathbb{R}$ be its performance measure function, and let $\xi \in \mathbb{R}^m$ be a random vector capturing the stochasticity in the system. Similar to Song and Nelson (2019), we study a case where all K designs share the same input distribution P^c ("c" means "correct"). The best design is defined as

$$b := \arg \max_{i \in \mathcal{J}} \mathbb{E}_{P^c} [h_i(\xi)],$$

where the expectation is assumed to be finite. We also assume that b is unique to avoid technicality. Furthermore, we consider a case where P^c is known to belong to a parametric family of distributions, $\{P_\theta \mid \theta \in \Theta \subseteq \mathbb{R}^d\}$, where the form of P_θ is known but the true parameter θ^c is not. For example, in an M/M/1 queue simulation model, the inter-arrival time is known to be exponentially distributed, but the arrival rate needs to be estimated from historical data. The parametric assumption on P^c can be justified by allowing a mixture of multiple parametric distributions (see the discussion in Cheng and Holland (1997)), provided that the parameter space is finite-dimensional. Generally speaking, there may be multiple independent

sources of IU (e.g., inter-arrival time and service time), in which case P_{θ^c} should be interpreted as a product measure. We will focus on the case of a single source of IU, where the input data, denoted by $\zeta_1, \zeta_2, \dots$, are independent and identically distributed (i.i.d.) samples from P_{θ^c} . The following notations will be used throughout the paper.

- (i) $H_i(\theta) := \mathbb{E}_{P_\theta}[h_i(\xi)]$, i.e., the true performance of design i under input distribution P_θ .
- (ii) $\delta_{ij}(\theta) := H_i(\theta) - H_j(\theta)$, i.e., the difference between designs i and j 's performances under input distribution P_θ .
- (iii) $\sigma_i^2(\theta) := \text{Var}_{P_\theta}[h_i(\xi)]$, i.e., the variance of design i 's simulation output under input distribution P_θ .
- (iv) The estimates of the above quantities will be denoted by $\hat{H}, \hat{\delta}$ and $\hat{\sigma}^2$, respectively.

Let $\hat{b}$ denote the estimated best design, and the probability of correct selection (PCS) is defined as $\mathbb{P}(\hat{b} = b)$. In general, the fixed confidence formulation of R&S aims to provide a statistical selection guarantee (e.g., 95% PCS) using minimal simulation effort or other resources. In the case without IU (i.e., θ^c is known exactly), a large body of literature studies the IZ formulation, which allows the user to specify the smallest difference in performance worth detecting. Most IZ procedures construct a continuation region for all pairs of designs (i, j) such that, if $\hat{\delta}_{ij}(\theta^c)$ escapes the region, then the sign of $\delta_{ij}(\theta^c)$ can be determined confidently based on which side $\hat{\delta}_{ij}(\theta^c)$ exits from. The key is to find a small continuation region for fast stopping without compromising the selection guarantee. Procedures of this type include the KN procedure (Kim and Nelson (2001)), the BIZ procedure (Frazier (2014)), the IZ-free procedures (Fan et al. (2016)) among several others.

2.2 Fixed Confidence R&S with Sequential Data

In the presence of IU, we consider a multi-stage scenario, where incremental data become available at each stage. Suppose that new batches of i.i.d. input data arrive sequentially, and our goal is to continuously reduce IU and SU in order to identify the best design with high confidence. More specifically, we would like our R&S procedure to run over a number of “stages”, where at each stage the following two steps are carried out.

- (i) Collect $k > 0$ additional new data samples to update the estimate of θ^c .
- (ii) For each design, run additional $R > 0$ replications under the new estimate of θ^c , and update the estimate of $H_i(\theta^c)$.

For simplicity, k and R are assumed to be fixed constants across different stages. We call a procedure *valid* if it selects the best design with a guaranteed PCS upon termination at a certain stage. The validness of a procedure hinges on three aspects.

1. Choice of estimator. What estimator is used to estimate θ^c ? The choice will affect the properties of the online estimator of $H_i(\theta^c)$, as well as the difficulty of designing a continuation region.
2. Online estimation. Although the estimate of θ^c gets increasingly accurate over the stages, the estimate of $H_i(\theta^c)$ cannot converge to its true value without reusing the simulation outputs from previous stages. How should we approach this online estimation problem?
3. Procedure design. The fixed confidence formulation essentially seeks to find a stopping time τ^* such that by the τ^* th stage, we can confidently determine which design is the best one. How can we design τ^* to make $\mathbb{E}[\tau^*]$ as small as possible?

The major challenge in this setting is that many existing procedures cannot be extended easily to handle IU. For example, most IZ procedures rely on a normality assumption on the simulation outputs, as this would admit the use of well-established tools associated with Brownian motion. While normality is

justified by batching and the Central Limit Theorem (CLT), the assumption typically fails when simulations are run under an estimate of θ^c , especially if such an estimate is updated in an online fashion. In this paper, we build our procedures on a Sequential Elimination framework (Even-Dar et al. (2002), Even-Dar et al. (2006), Glynn and Juneja (2015)), as it allows us to construct valid continuation regions even in the presence of IU.

Regarding the first aspect, we will restrict our discussion to a specific type of estimator. Let $\hat{\theta}_N$ be the estimate of θ^c using N data samples, and recall that $\{\zeta_1, \dots, \zeta_N\}$ are the input data. The following assumption is made to make the analysis more tractable.

Assumption 1. *The estimator of θ^c can be written in the form of $\hat{\theta}_N = \frac{1}{N} \sum_{j=1}^N G(\zeta_j)$, where $G: \mathbb{R}^m \rightarrow \mathbb{R}^d$ and $\mathbb{E}[G(\zeta_1)] = \theta^c$.*

Assumption 1 can often be satisfied through reparametrization. For example, the normal distribution can be reparametrized by its first two moments, and then $\hat{\theta}_N$ corresponds to the moment estimators. Under Assumption 1, our problem setting can be simplified considerably. During the n th stage, we collect k additional data samples $\{\zeta_j \mid j = (n-1)k + 1, \dots, nk\}$, and the sample mean of these incremental data samples is

$$D_n := \frac{1}{k} \sum_{j=(n-1)k+1}^{nk} G(\zeta_j),$$

which can be viewed as a single batched data sample with variance shrunk by a factor of k . It can be easily checked that

$$\frac{1}{n} \sum_{j=1}^n D_j = \frac{1}{nk} \sum_{j=1}^{nk} G(\zeta_j).$$

In other words, without loss of generality, we can assume that at the end of the n th stage, the estimator of θ^c takes the form of $\hat{\theta}_n = \frac{1}{n} \sum_{j=1}^n D_j$, where D_j are i.i.d. samples with $\mathbb{E}[D_1] = \theta^c$. Similarly, it suffices to consider $R = 1$, i.e., when only one additional simulation replication is run at each stage. From this point on, our problem setting is simplified as follows. During the n th stage,

- (i) first collect one data sample D_n , and compute $\hat{\theta}_n = \frac{n-1}{n} \hat{\theta}_{n-1} + \frac{1}{n} D_n$;
- (ii) then, for each design, run one more independent simulation replication under $\hat{\theta}_n$, and aggregate the simulation output with the previous ones.

3 PROCEDURE DESIGN

3.1 Moving Average Estimator

For the online estimation problem described in Section 2.2, a consistent estimator of $H_i(\theta^c)$ can be constructed in various ways. For instance, simply averaging all the simulation outputs $\{h_i(\xi_{in})\}_n$ usually ensures consistency. An alternative is to use a likelihood ratio estimator by reweighting the simulation outputs, but due to the correlation among $\{\hat{\theta}_n\}$, the resulting estimator will be biased (see Eckman and Feng (2018) for insights into this observation).

Since our ultimate goal is to solve the R&S problem, the main challenge lies in finding an estimator which facilitates the design of a valid procedure. Let $\hat{H}_{i,n}$ denote the estimate of $H_i(\theta^c)$ at the end of the n th stage. We construct an estimator by discarding the first (or the “oldest”) $n_\eta := \lfloor \eta n \rfloor, \eta \in (0, 1)$, simulation outputs and then averaging the rest, i.e.,

$$\hat{H}_{i,n} := \frac{1}{n - n_\eta} \sum_{r=n_\eta+1}^n h_i(\xi_{ir}), \quad i \in \mathcal{I}. \quad (1)$$

The estimator in (1) will be referred to as a *moving average estimator*, since it averages simulation outputs within a moving and expanding time window. The motivation is to throw away some of those “outdated”

simulation outputs which were generated under less accurate estimates of θ^c . We establish the asymptotic normality of $\hat{H}_{i,n}$ in the following theorem, where $\mathcal{N}$ denotes normal distribution, and “ $\Rightarrow$ ” denotes convergence in distribution.

Theorem 1 Let Assumption 1 hold. Further suppose that $\Sigma_G := \text{Cov}[D_1]$ exists and H_i is twice continuously differentiable at θ^c . Then, for any $\eta \in [0, 1)$,

$$\sqrt{n}[\hat{H}_{i,n} - H_i(\theta^c)] \Rightarrow \mathcal{N}(0, \tilde{\sigma}_{i,\infty}^2) \quad \text{as } n \rightarrow \infty,$$

where

$$\tilde{\sigma}_{i,\infty}^2 := w_\eta \nabla H_i(\theta^c)^\top \Sigma_G \nabla H_i(\theta^c) + \frac{1}{1-\eta} \sigma_i^2(\theta^c),$$

with ∇ being the gradient, and

$$w_\eta := \frac{2}{1-\eta} + \frac{2\eta \log \eta}{(1-\eta)^2}.$$

Theorem 1 is an interesting result in its own right. It shows that the limiting variance $\tilde{\sigma}_{i,\infty}^2$ is a weighted sum of variances caused by IU and SU, which are $\tilde{V}_I := \nabla H_i(\theta^c)^\top \Sigma_G \nabla H_i(\theta^c)$ and $\tilde{V}_S := \sigma_i^2(\theta^c)$, respectively. To interpret the weights, we look at the following cases.

- (i) Setting $\eta = 0$ gives $w_\eta = 2$, meaning that if we retain all the outputs, then the variance caused by IU, corresponding to $\tilde{V}_I$, will be doubled.
- (ii) Sending $\eta \rightarrow 1$ gives $w_\eta \rightarrow 1$. This loosely corresponds to the case of $\hat{H}_{i,n} = h_i(\xi_{in})$, where we only retain the single most recent output. As a result, $\hat{H}_{i,n}$ is free from the error accumulated over previous estimates of θ^c , and thus $\tilde{V}_I$ is not inflated. However, $\tilde{V}_S$ is inflated by a factor $1/(1-\eta) \rightarrow \infty$, since the effective number of outputs is not tending to ∞ as $n \rightarrow \infty$.
- (iii) A balance between IU and SU can be achieved by choosing the η that minimizes $\tilde{\sigma}_{i,\infty}^2$.

Roughly speaking, the parameter η captures a bias-variance tradeoff. On the one hand, discarding previous outputs helps reduce the bias $H_i(\hat{\theta}_n) - H_i(\theta^c)$ due to IU. On the other hand, the variance caused by SU gets inflated if we average over fewer simulation outputs. In Section 3.2, we will explain the role of moving average estimator in designing a valid procedure. At this point, an important note is that *simply ignoring IU and applying existing procedures may result in undershooting the PCS target*.

3.2 The SE-IU Procedure

Our procedure is a direct extension of a Sequential Elimination framework proposed by Even-Dar et al. (2002), Even-Dar et al. (2006), which is also discussed in Glynn and Juneja (2015) recently. This general paradigm has a simple structure and can be extended to handle IU. Given $\alpha \in (0, 1)$, the idea is to construct confidence bounds $\{c_{i,n}\}$ on $\hat{H}_{i,n}$ for each design i such that

$$\mathbb{P}\{|\hat{H}_{i,n} - H_i(\theta^c)| \leq c_{i,n}, \forall i, n\} \geq 1 - \alpha, \quad \alpha \in (0, 1), \quad (2)$$

where $c_{i,n} \rightarrow 0$ as $n \rightarrow \infty$. At each stage n , a design i gets eliminated if

$$\hat{H}_{i,n} + c_{i,n} < \max_{j \neq i} \{\hat{H}_{j,n} - c_{j,n}\}.$$

In other words, a design is eliminated if its upper confidence bound is below some other design's lower confidence bound. Then, on the event $\mathcal{E} := \{|\hat{H}_{i,n} - H_i(\theta^c)| \leq c_{i,n}, \forall i, n\}$, we have for any $i \neq b$,

$$\hat{H}_{b,n} + c_{b,n} - (\hat{H}_{i,n} - c_{i,n}) \geq \delta_{bi}(\theta^c) > 0, \quad \forall n.$$

Therefore, design b will never be eliminated on event $\mathcal{E}$. Since $c_{i,n} \rightarrow 0$, the procedure terminates almost surely, and b will survive all eliminations with probability at least $1 - \alpha$. By (2), this delivers the desired PCS guarantee.

The key to efficiently ruling out inferior designs is to find tight confidence bounds $\{c_{i,n}\}$ that satisfy (2). When there is no IU, this can be done easily using well-known concentration inequalities on $\hat{H}_i(\theta^c)$, such as the Chernoff bound and Hoeffding's inequality. In the presence of IU, these inequalities do not apply directly and a new concentration bound needs to be derived for $\hat{H}_{i,n}$. The following assumption will be useful to this end. Recall that a random variable X is sub-Gaussian with parameter σ if $\mathbb{E}[e^{sX}] \leq \exp(\sigma^2 s^2/2)$ for all $s \in \mathbb{R}$.

Assumption 2.

- (i) For all $1 \leq j \leq d$, the j th coordinate of D_n is sub-Gaussian with parameter v_j .
- (ii) For any $\theta \in \Theta$, if $\xi \sim P_\theta$, then $h_i(\xi)$ is sub-Gaussian.
- (iii) For all $u > 0$ and any design i , there exists a function $L_i(\cdot) > 0$ such that

$$|H_i(\theta_1) - H_i(\theta_2)| \leq L_i(u) \|\theta_1 - \theta_2\|, \quad \forall \theta_1, \theta_2 \in \{\theta \in \Theta \mid \|\theta - \theta^c\| \leq u\},$$

where $\|\cdot\|$ is the Euclidean norm.

- (iv) For any design i , $\sigma_i^2(\theta)$ is a continuous function of θ .

Assumption 2 (i) and (ii) may appear restrictive at first sight. For example, the input distributions in an M/M/1 queue simulation model are sub-exponential but not sub-Gaussian. However, we make these assumptions mainly to avoid unnecessary technicality, and in our future work we will show numerically that our procedure works even if these conditions are not met.

To get a sense of how to meet the guarantee in (2), consider constructing confidence bounds $c_{i,n}$ such that

$$\mathbb{P} \left\{ \{ |\hat{H}_{i,n} - H_i(\theta^c)| \leq c_{i,n}, \forall n \}^c \right\} \leq \beta \quad (3)$$

for some $\beta \in (0, 1)$, where “c” denotes set complement. One way is to consider an event

$$A_u := \{ \|\hat{\theta}_n - \theta^c\| \leq u, \forall n \}$$

for some $u > 0$. Then, the guarantee in (3) can be met if we can control

$$\mathbb{P} \{ A_u \cap \{ |\hat{H}_{i,n} - H_i(\theta^c)| \leq c_{i,n}, \forall n \}^c \} + \mathbb{P}(A_u^c),$$

since it is an upper bound on the left-hand side (LHS) of (3). Note that $\mathbb{P}(A_u^c)$ can be controlled by enlarging u . Meanwhile, on the event A_u , we have H_i being Lipschitz continuous and $\sigma_i(\hat{\theta}_n)$ being bounded, where it is possible to derive a concentration bound for $|\hat{H}_{i,n} - H_i(\theta^c)|$ through a decomposition,

$$|\hat{H}_{i,n} - H_i(\theta^c)| \leq |\hat{H}_{i,n} - H_i(\hat{\theta}_n)| + |H_i(\hat{\theta}_n) - H_i(\theta^c)|.$$

The rest is to combine all the bounds through a union bound, where the choice of the estimator $\hat{H}_{i,n}$ is crucial. For example, if we simply average all the simulation outputs, then the bound will be infinite due to cumulative bias. Using the moving average estimator, however, we are able to construct $c_{i,n}$ that satisfies (3) by virtue of a bias-variance tradeoff (as long as $\eta > 0$).

The upcoming SE-IU procedure relies on some key parameters including $\{v_j\}$, $\{\sigma_i\}$ and $\{L_i\}$. For now, we present an ideal version of the procedure by assuming full knowledge of these parameters, and defer implementation details to Section 3.3.

Procedure: SE-IU (ideal version)

- **Input.** $\alpha \in (0, 1), \eta \in (0, 1), K \geq 2, n_0 \geq 10$.

- **Step 1.** Solve the following equation in u , and let u^* be the solution.

$$\sum_{j=1}^d \frac{\exp\left(-\frac{(n_0+1)u^2}{2dv_j^2}\right)}{1 - \exp\left(-\frac{u^2}{2dv_j^2}\right)} = \frac{\alpha}{6}. \quad (4)$$

- **Step 2.** For each design, compute

$$\bar{\sigma}_i := \sup_{\substack{\|\theta - \theta^c\| \leq u^* \\ \theta \in \Theta}} \sigma_i(\theta), \quad \bar{L}_i := L_i(u^*),$$

as well as the constants

$$\kappa_{n_0} := \sum_{n=n_0+1}^{\infty} n^{-2}, \quad \beta_{n_0} := \sum_{n=n_0+1}^{\infty} (n - n_\eta)^{-2}.$$

Let $\bar{v} := \max_j v_j$. Run n_0 stages and set $n \leftarrow n_0 + 1$. Also set $S \leftarrow \{1, 2, \dots, K\}$.

- **Step 3.** Run an additional stage for all designs $1, 2, \dots, K$, and compute their estimates $\hat{H}_{i,n}$ using the moving average estimator in (1).
- **Step 4.** Compute the confidence bounds $c_{i,n}$ for each design i , where $c_{i,n} = t_{i,n} + r_{i,n}$ and

$$t_{i,n} := 2\bar{\sigma}_i \sqrt{\frac{\ln\left(\sqrt{\frac{6K\kappa_{n_0}}{\alpha}} n\right)}{n - n_\eta}}, \quad r_{i,n} := \bar{v}\bar{L}_i \sqrt{\frac{6d \ln\left(\left(\frac{6dK\beta_{n_0}}{\alpha}\right)^{\frac{1}{3}} (n - n_\eta)\right)}{n_\eta + 1}}.$$

For each $i \in S$, if

$$\hat{H}_{i,n} + c_{i,n} < \max_{j \neq i} (\hat{H}_{j,n} - c_{j,n}),$$

then set $S \leftarrow S \setminus \{i\}$. Go to Output if $|S| = 1$; otherwise, set $n \leftarrow n + 1$ and go to Step 3.

- **Output.** Select the only design in S as the best one.

Some important features of SE-IU are outlined as follows.

1. First, the width of $\{c_{i,n}\}$ is of order $O(\sqrt{\ln(n)/n})$, which is standard for Sequential Elimination procedures. However, the confidence bounds are widened compared with the case without IU, since $t_{i,n}$ and $r_{i,n}$ correspond to SU and IU, respectively.
2. Second, we do not eliminate any design in the first n_0 stages. In view of (4), a larger n_0 leads to a smaller u^* , which in turn gives us smaller $\bar{\sigma}_i, \bar{L}_i, \kappa_{n_0}$ and β_{n_0} , hence tighter confidence bounds. Also, equation (4) arises as a result of controlling a probability bound, and it always has a unique solution since the LHS is a continuous and monotone function of u with range $(0, \infty)$.
3. Third, the running time of SE-IU primarily depends on the parameters $v_j, \bar{\sigma}_i, \bar{L}_i$, and $\delta_{bi}(\theta^c)$. For instance, if $\bar{\sigma}_i$ is increased by a factor of $k > 1$, then it would take at least k^2 times as many stages to reach the same width of confidence bounds.

Let τ^* be the number of stages until the procedure terminates. A nice property of the Sequential Elimination framework is that it is automatically equipped with an upper bound on $\mathbb{E}[\tau^*]$.

Theorem 2 Let Assumption 2 hold. Then, the SE-IU procedure guarantees to select the best design with probability at least $1 - \alpha$. Furthermore,

$$\mathbb{E}[\tau^*] \leq 2 \sum_{i \neq b} \tau_i^* + 4(K-1)(\alpha + 2de^{-\mathcal{K}}(1 - e^{-\mathcal{K}})^{-2}),$$

where $\tau_i^* := \inf\{n > n_0 \mid 2(c_{b,n} + c_{i,n}) \leq \delta_{bi}(\theta^c)\}$ and $\mathcal{K} := \eta(u^*)^2/(2d\bar{v}^2)$.

The dominating term in Theorem 2's bound is $\sum_{i \neq b} \tau_i^*$, where each τ_i^* characterizes the difficulty in eliminating design i . For example, if design i has a large variance $\sigma_i^2(\theta^c)$ and a small gap $\delta_{bi}(\theta^c)$, then τ_i^* would be large, and it will take longer to eliminate design i . Given the same performance gap $\delta_{bi}(\theta^c)$, τ_i^* primarily depends on $\{c_{i,n}\}$, i.e., the width of the confidence bounds.

3.3 Implementation Guidance

We briefly discuss how to estimate the unknown quantities in SE-IU. One may start off by collecting a small size of input data for initial estimation. If IU or SU is high (relative to the estimates of $\delta_{ij}(\theta^c)$), then consider using a larger n_0 . The difficult parameters are $\bar{\sigma}_i$ and $\bar{L}_i$, which are the suprema of $\sigma_i(\cdot)$ and $\|\nabla H_i(\cdot)\|$ over a small neighborhood of θ^c . While one can attempt to maximize the corresponding likelihood ratio estimators, we suggest simply replacing them with estimates of $\sigma_i(\theta^c)$ and $\|\nabla H_i(\theta^c)\|$ for the following reasons: (i) estimates based on such maximization often suffer from high variance and severe overestimation; (ii) the Sequential Elimination framework is already conservative since it resorts to loose union bounds, so highly accurate estimates are often unnecessary.

4 NUMERICAL RESULTS

We test SE-IU on a simple quadratic example for illustration, in which the target PCS is 95%. The performance of the i th design is characterized by

$$h_i(\xi) = -(i - \xi)^2,$$

where $\xi \sim \mathcal{N}(\theta^c, \sigma^2)$ is the only source of SU. We assume that the standard deviation σ (not to be confused with $\sigma_i(\theta)$) is known, and the unknown input parameter is the mean θ^c . There are 20 designs to be compared, which are $\mathcal{J} = \{-10, -9, \dots, 8, 9\}$. The true parameter value is $\theta^c = 0$, so $i = 0$ is the best design. In addition, the batch size per stage is 1, meaning that at each stage an additional data sample is collected, and each design is simulated one more time; the number of initial stages is $n_0 = 1$. Recall that $H_i(\theta) := \mathbb{E}_{P_\theta}[h_i(\xi)]$ and $\sigma_i^2(\theta) := \text{Var}_{P_\theta}[h_i(\xi)]$ are the mean and the variance of design i 's simulation output, respectively. Figure 1 shows the magnitude of $H_i(\theta^c)$ and $\sigma_i(\theta^c)$ for all designs when $\sigma = 0.5$.




Figure 1: $H_i(\theta^c)$ and $\sigma_i(\theta^c)$ for all designs $i \in \mathcal{J}$, where $\theta^c = 0$ and $\sigma = 0.5$.

Next, the proposed moving average estimator in (1) has a parameter $\eta \in (0, 1)$, which is the fraction of historical simulation outputs to be discarded. Adjusting the value of η can balance IU and SU in a quite explicit and intuitive way. For instance, Figure 2 shows the continuation region of design -10 when η takes three values: 0.2, 0.5 and 0.8. In all three subfigures, the blue curve corresponds to $\{t_{i,n}\}$, the red curve corresponds to $\{r_{i,n}\}$, and the black curve is the combined confidence bands $\{c_{i,n}\}$. It can be seen that when $\eta = 0.2$, IU dominates SU, and the dominance reverses when $\eta = 0.8$. Taking $\eta = 0.5$ seems

to yield the best balance between IU and SU, where the width of $c_{i,n} = t_{i,n} + r_{i,n}$ is slightly smaller than the other two cases. Therefore, we set $\eta = 0.5$ in subsequent results.




Figure 2: Continuation regions for design -10 under different values of η .

Finally, we test SE-IU under $\sigma = 0.1, 0.2, \dots, 0.5$, and compare the corresponding expected running times until termination. The PCSs are all close to 1, so SE-IU can achieve the target PCS despite initial estimation error. In Figure 3, the expected stopping time, $\mathbb{E}[\tau^*]$, are estimated over 100 independent replications. It can be seen that $\mathbb{E}[\tau^*]$ increases roughly quadratically in σ , and the magnitude indicates that SE-IU needs further improvement in efficiency before it can be deemed a practical procedure. For a comparison, if there is no IU, then the classical SE procedure would take significantly shorter time to terminate. This can be expected from Figure 1, where the simulation noise $\sigma_i(\theta^c)$ is not too large relative to the performance differences $\delta_{bi}(\theta^c), i \neq b$.

5 CONCLUSION AND FUTURE DIRECTIONS

In this work, we study a new fixed confidence Ranking and Selection problem, where new input data can be collected sequentially to update the input models towards higher accuracy. To solve the problem, we propose an SE-IU procedure by extending a Sequential Elimination framework. Theoretical guarantees are provided on SE-IU's performance, and numerical results demonstrate that it is able to achieve the target PCS. However, we also observe that SE-IU suffers from conservativeness, as it overshoots the PCS




Figure 3: Expected running time under different values of σ .

target and requires a long running time. One reason behind this is that, accounting for IU and cumulative estimation error naturally inflates the continuation region; another reason is that the current SE-IU does not exploit the common input distribution effect and the common random numbers technique, which often induce positive correlations among the designs and may greatly sharpen the comparisons. Therefore, our next step is to take advantage of such correlations through pairwise comparisons or simultaneous multiple comparisons, which will most likely boost the efficiency of SE-IU.

REFERENCES

- Bechhofer, R. E. 1954. "A Single-sample Multiple Decision Procedure for Ranking Means of Normal Populations with Known Variances". *The Annals of Mathematical Statistics* 25(1):16–39.
- Branke, J., S. E. Chick, and C. Schmidt. 2005. "New Developments in Ranking and Selection: An Empirical Comparison of the Three Main Approaches". In *Proceedings of the 37th conference on Winter simulation*, edited by M. E. Kuhl, N. M. Steiger, F. B. Armstrong, and J. A. Joines, 708–717. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.
- Cheng, R. C., and W. Holland. 1997. "Sensitivity of Computer Simulation Experiments to Errors in Input Data". *Journal of Statistical Computation and Simulation* 57(1-4):219–241.
- Chick, S. E., and P. Frazier. 2012. "Sequential Sampling with Economics of Selection Procedures". *Management Science* 58(3):550–569.
- Corlu, C. G., and B. Biller. 2013. "A Subset Selection Procedure under Input Parameter Uncertainty". In *Proceedings of the 2013 Winter Simulation Conference*, edited by R. Pasupathy, S.-H. Kim, A. Tolk, R. Hill, and M. E. Kuhl, 463–473. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.
- Corlu, C. G., and B. Biller. 2015. "Subset Selection for Simulations Accounting for Input Uncertainty". In *Proceedings of the 2015 Winter Simulation Conference*, edited by L. Yilmaz, W. K. V. Chan, I. Moon, T. M. K. Roeder, C. Macal, and M. D. Rossetti, 437–446. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.
- Eckman, D. J., and M. B. Feng. 2018. "Green Simulation Optimization Using Likelihood Ratio Estimators". In *Proceedings of the 2018 Winter Simulation Conference*, edited by M. Rabe, A. A. Juan, N. Mustafee, A. Skoogh, S. Jain, and B. Johansson, 2049–2060. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.
- Even-Dar, E., S. Mannor, and Y. Mansour. 2002. "PAC Bounds for Multi-armed Bandit and Markov Decision Processes". In *International Conference on Computational Learning Theory*, 255–270. Heidelberg, Berlin: Springer.
- Even-Dar, E., S. Mannor, and Y. Mansour. 2006. "Action Elimination and Stopping Conditions for the Multi-armed Bandit and Reinforcement Learning Problems". *Journal of Machine Learning Research* 7:1079–1105.
- Fan, W., L. J. Hong, and B. L. Nelson. 2016. "Indifference-zone-free Selection of the Best". *Operations Research* 64(6):1499–1514.
- Fan, W., L. J. Hong, and X. Zhang. 2013. "Robust Selection of the Best". In *Proceedings of the 2013 Winter Simulation Conference*, edited by R. Pasupathy, S.-H. Kim, A. Tolk, R. Hill, and M. E. Kuhl, 868–876. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.
- Frazier, P. I. 2014. "A Fully Sequential Elimination Procedure for Indifference-zone Ranking and Selection with Tight Bounds on Probability of Correct Selection". *Operations Research* 62(4):926–942.
- Glynn, P., and S. Juneja. 2015, July. "Selecting the Best System and Multi-armed Bandits". *arXiv e-prints*:arXiv:1507.04564.

- Hong, L. J. 2006. "Fully Sequential Indifference-zone Selection Procedures with Variance-dependent Sampling". *Naval Research Logistics* 53(5):464–476.
- Kim, S.-H., and B. L. Nelson. 2001. "A Fully Sequential Procedure for Indifference-zone Selection in Simulation". *ACM Transactions on Modeling and Computer Simulation (TOMACS)* 11(3):251–273.
- 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. G. Henderson, B. Biller, M.-H. Hsieh, J. Shortle, J. D. Tew, and R. R. Barton, 162–172. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.
- Ma, S., and S. G. Henderson. 2017. "An Efficient Fully Sequential Selection Procedure Guaranteeing Probably Approximately Correct Selection". In *Proceedings of the 2017 Winter Simulation Conference*, edited by W. K. V. Chan, A. D'Ambrogio, G. Zacharewicz, N. Mustafee, G. Wainer, and E. Page, 2225–2236. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.
- Song, E., and B. L. Nelson. 2019. "Input-output Uncertainty Comparisons for Discrete Optimization via Simulation". *Operations Research* 67(2):562–576.
- Song, E., B. L. Nelson, and L. J. Hong. 2015. "Input Uncertainty and Indifference-zone Ranking & Selection". In *Proceedings of the 2015 Winter Simulation Conference*, edited by L. Yilmaz, W. K. V. Chan, I. Moon, T. M. K. Roeder, C. Macal, and M. D. Rossetti, 414–424. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.

AUTHOR BIOGRAPHIES

DI WU is a Ph.D. candidate in the H. Milton Stewart School of Industrial and Systems Engineering at Georgia Institute of Technology. He received his B.E. in electrical engineering from East China University of Science and Technology, China, in 2011, and his M.S. in control science and engineering from Tsinghua University, China, in 2014. His research interests include simulation optimization, risk quantification, and online learning. His e-mail address is dwu80@gatech.edu.

ENLU ZHOU is an Associate Professor in the H. Milton Stewart School of Industrial and Systems Engineering at Georgia Institute of Technology. She received the B.S. degree with highest honors in electrical engineering from Zhejiang University, China, in 2004, and received the Ph.D. degree in electrical engineering from the University of Maryland, College Park, in 2009. Her research interests include stochastic control and simulation optimization. Her email address is enlu.zhou@isye.gatech.edu and her web page is <http://enluzhou.gatech.edu/>.