

OPTIMIZING INPUT DATA ACQUISITION FOR RANKING AND SELECTION: A VIEW THROUGH THE MOST PROBABLE BEST

Taeho Kim
Eunhye Song

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

ABSTRACT

This paper concerns a Bayesian ranking and selection (R&S) problem under input uncertainty when all solutions are simulated with common input models estimated from data. We assume that there are multiple independent input data sources from which additional data can be collected at a cost to reduce input uncertainty. To optimize input data acquisition, we first show that the most probable best (MPB)—the solution with the largest posterior probability of being optimal (posterior preference)—is a strongly consistent estimator for the real-world optimum. We investigate the optimal asymptotic static sampling ratios from the input data sources that maximizes the exponential convergence rate of the MPB's posterior preference. We then create a sequential sampling rule that balances the simulation and input data collection effort. The proposed algorithm stops with posterior confidence in the solution quality.

1 INTRODUCTION

In many applications, decision-makers build simulation models to optimize the designs of complex real-world stochastic systems. The input models that generate random variates fed into these simulators are typically estimated from data observed from the systems. When the data are finite, the estimated input models do not match the real-world data-generating processes. Therefore, the stochastic variability in the simulation output is inflated by the estimation error in the input models, which is referred to as *input uncertainty*. Since the conclusions from the simulation experiment depend on the fidelity of the input models, input uncertainty must be properly accounted for in the simulation analysis to draw statistically correct conclusions.

This paper investigates a ranking and selection (R&S) problem under input uncertainty where all competing solutions' simulators share a collection of common estimated input models. In particular, we consider the case when the real-world stochastic system has several independent input-generating processes from which additional data can be acquired, albeit at a cost, to improve the input models. The objective of the R&S problem is to find the solution that has the optimal performance in the real-world stochastic system. However, due to input uncertainty, the conditional optimal solution given the estimated input models is not the real-world optimum in general. Nevertheless, as more input data are collected, one might expect the estimated input models become closer to the real-world distributions. In turn, the conditional optimum also converges to the true optimum. Thus, the decision-maker would want to carefully allocate the input data collection effort to most efficiently reduce input uncertainty. Meanwhile, the mapping between allocation of estimated input models and its conditional optimum must be learned via simulation. We assume that the simulation cost is non-negligible, yet significantly cheaper than the cost of input data. Thus, it is important to design an experiment that judiciously spends sampling effort for both input and simulation

data collection. We aim to propose a unified framework integrating optimal input data and simulation sampling strategies to discover the true optimum efficiently.

Many studies concerning simulation optimization under input uncertainty assume that the input data batch is fixed, and thus input uncertainty cannot be reduced further. In this case, one can adopt an appropriate risk measure and focus on finding the “optimal” solution under the risk measure to hedge against input uncertainty. For instance, Corlu and Biller (2015), Pearce and Branke (2017), and Ungredda et al. (2022) take the mean function as a risk measure (risk-neutral), and Xie and Zhou (2015), Zhu et al. (2020) use a conditional value at risk or value at risk. Gao et al. (2017) and Fan et al. (2020) find the solution that has the best worst-case performance within finite ambiguity set of the true input distribution.

There are some recent work that incorporate streaming input data in simulation optimization (Song and Shanbhag 2019; Liu et al. 2021; Wu and Zhou 2019), which are different from our problem setting. We actively choose which input data to acquire instead of passively collecting the streaming data.

A similar setting to our problem is considered in Wu and Zhou (2019), where the authors propose an optimal computing budget allocation (OCBA-IU) algorithm for fixed-budget setting. Xu et al. (2020) investigate the case when two input distributions are correlated. Both approaches apply central-limit-theorem and delta-method type results to decide the tradeoff between input data collection and simulation. On the other hand, Ungredda et al. (2022) takes a Bayesian optimization approach to quantify the value of information (VoI) of additional input data collection and simulation run, respectively, and propose a sequential sampling rule based on the VoIs.

Our work differs from these in several ways. First, we adopt the *most probable best* (MPB) proposed by Kim et al. (2021) as an estimator of the true optimum. The MPB is the solution with the largest posterior probability of being optimal (posterior preference). We show i) the MPB is strongly consistent; and ii) the posterior preference of the MPB converges to one almost surely as the input data size increases. The latter gives us a basis for adopting the posterior preference of the MPB to measure the model risk in the R&S problem caused by input uncertainty; the closer to one, the less the model risk there is. Motivated by this observation, we analyze the exponential convergence rate of the posterior preference of the MPB as a function of sampling ratios among the several input data sources, then devise sequential input data acquisition rule to achieve the fastest convergence rate. Meanwhile, simulation sampling is carefully balanced with the input data collection effort to learn solutions’ performances under different parameter values. Lastly, we propose a Bayesian stopping criterion to terminate both input data and simulation sampling when there is enough statistical confidence that the MPB indeed is the true optimal.

The rest of this paper is organized as follows. Section 2 presents some preliminaries on Bayesian input modeling and introduces the MPB. In Section 3, we formulate the input data acquisition problem based on the large-deviation analysis of the posterior preference of the MPB. In Section 4, an efficient estimation scheme to learn the MPB’s posterior preference is introduced. The integrated sequential sampling framework for input data acquisition and simulation is discussed in Section 5. Section 6 validates our algorithm with numerical experiments. Proofs of all theorems are omitted due to the page limit.

2 PROBLEM FORMULATION

Suppose there exist $L > 1$ independent input data sources. Let $f_{\theta^\ell}^\ell$ and θ^ℓ represent the density function and parameter vector of the ℓ -th input. We assume that the L parametric distribution families are known, however, the true parameter vector, $\theta_0 = (\theta_0^1, \theta_0^2, \dots, \theta_0^L)$, is unknown. All k solutions in comparison are simulated using the same joint input distribution. For each Solution i ($1 \leq i \leq k$), we define $y_i(\theta)$ as the mean response of Solution i given θ , i.e., $y_i(\theta) = E[Y_i(\theta)|\theta]$, where $Y_i(\theta)$ is the simulation output of Solution i when θ is the input model parameter. The true optimum is defined as

$$i_0 := \arg \min_{1 \leq i \leq k} y_i(\theta_0). \tag{1}$$

Both $\{y_i(\theta)\}_{1 \leq i \leq k}$ and θ_0 are unknown and must be estimated. We first introduce a Bayesian machinery to model uncertainty about θ_0 . Let $\theta = (\theta^1, \theta^2, \dots, \theta^L)$ be the Bayesian estimator for θ_0 , where $\theta^\ell \in \Theta^\ell$

for $1 \leq \ell \leq L$. Denote the data size from the ℓ -th data source by m_ℓ , and let $m = \sum_{\ell=1}^L m_\ell$. Then, the likelihood process for the m observed data is $f_{\boldsymbol{\theta}}(\mathcal{Z}_m) := \prod_{\ell=1}^L f_{\boldsymbol{\theta}^\ell}(\mathcal{Z}_{m_\ell}^\ell)$ where $\mathcal{Z}_m := \cup_{\ell=1}^L \mathcal{Z}_{m_\ell}^\ell$ and $\mathcal{Z}_{m_\ell}^\ell := \{Z_1^\ell, Z_2^\ell, \dots, Z_{m_\ell}^\ell\}$ is the size- m_ℓ independent and identically distributed (i.i.d.) collection of the ℓ -th input data. We denote the prior distribution of $\boldsymbol{\theta}^\ell$ by $\pi_0^\ell(\boldsymbol{\theta}^\ell)$. Then, for each $1 \leq \ell \leq L$, the posterior distribution of $\boldsymbol{\theta}^\ell$ can be written as

$$\pi_{m_\ell}^\ell(\boldsymbol{\theta}^\ell) = \pi_0^\ell(\boldsymbol{\theta}^\ell) f_{\boldsymbol{\theta}^\ell}^\ell(\mathcal{Z}_{m_\ell}^\ell) / \int_{\Theta^\ell} \pi_0^\ell(\boldsymbol{\theta}_1^\ell) f_{\boldsymbol{\theta}_1^\ell}^\ell(\mathcal{Z}_{m_\ell}^\ell) d\boldsymbol{\theta}_1^\ell,$$

and the joint posterior density of $\boldsymbol{\theta}$ is $\pi_m = \prod_{\ell=1}^L \pi_{m_\ell}^\ell(\boldsymbol{\theta}^\ell)$.

When $\boldsymbol{\theta}_0$ is unknown, one may consider a plug-in version of (1) by replacing $\boldsymbol{\theta}_0$ with its point estimate $\widehat{\boldsymbol{\theta}}$; here, we adopt the maximum a posteriori (MAP) estimator of π_m . However, the optimal solution for the plug-in version of (1) is suboptimal for (1) in general (Song and Nelson (2019)). Another source of uncertainty is simulation error. Even if we have correct $\boldsymbol{\theta}_0$, false selection of i_0 may occur since we estimate i_0 based on sample estimates of $\{y_i(\boldsymbol{\theta})\}$. Hence, it is essential to reduce both input uncertainty and simulation error to find i_0 .

To analyze the effect of input uncertainty to the R&S problem, we introduce the concept of *posterior preference*. First proposed by Kim et al. (2022), the posterior preference of Solution j is defined as

$$P_j(\pi_m) := \mathbb{P}_{\pi_m} \{y_j(\boldsymbol{\theta}) = \min_{1 \leq i \leq k} y_i(\boldsymbol{\theta})\}.$$

Namely, $P_j(\pi_m)$ quantifies how likely Solution j is optimal under π_m . Under the following assumption, the posterior preferences of all solutions sum to one.

Assumption 1 The true optimum, i_0 , is unique. Moreover, $\arg \min_{1 \leq i \leq k} y_i(\boldsymbol{\theta})$ is unique almost everywhere under π_0 .

Because $P_j(\pi_m)$ depends on the m observed data, its value would change when additional input data are collected. In the next section, we show that $P_{i_0}(\pi_m)$ converges to one almost surely with an exponential rate as m grows under some assumptions. This result combined with Assumption 1 implies that i_0 has the largest posterior preference for sufficiently large m . Therefore, it is sensible to adopt the solution with the largest posterior preference as an estimator for i_0 ; Kim et al. (2021) refer to this estimator as the *most probable best* (MPB). Mathematically, the MPB given π_m is defined as

$$i^*(\pi_m) := \arg \max_{1 \leq j \leq k} P_j(\pi_m).$$

In the next section, we show that $i^*(\pi_m)$ is indeed a strongly consistent estimator for i_0 under some assumptions.

We further define the *favorable set* of Solution i as $\Theta_i = \{\boldsymbol{\theta} : y_i(\boldsymbol{\theta}) \leq \min_{j \neq i} y_j(\boldsymbol{\theta})\}$, and refer to its complement as the *adversarial set* of Solution i . Note that the favorable and adversarial sets are deterministic regardless of π_m . By definition, we have $P_i(\pi_m) = \mathbb{P}_{\pi_m}(\Theta_i)$.

Since collecting additional input data is feasible in our case, the key question is how much data to observe from which sources to find i_0 most efficiently? Ideally, we would like to design a sequential input data sampling rule that makes $P_{i_0}(\pi_m)$ converge to one as fast as possible so that the MPB quickly converges to i_0 . Since i_0 is unknown to us, instead, we focus on maximizing the exponential convergence rate of $P_{i^*(\pi_m)}(\pi_m)$. Moreover, we stop additional input data collection if $P_{i^*(\pi_m)}(\pi_m)$ is above a target value as we have enough statistical confidence that the MPB is indeed i_0 .

We assume collecting data from any of the L input processes is equally costly in this paper while deferring a more general case for the future work.

3 ASYMPTOTIC PROPERTIES OF THE MOST PROBABLE BEST

This section is devoted to technical results associated with the MPB and its posterior preference. We begin with Assumptions 2 and 3 below, which summarize the additional conditions we impose on $f_{\boldsymbol{\theta}^\ell}^\ell$ and $\pi_0^\ell(\boldsymbol{\theta}^\ell)$ to study the asymptotic behavior of $i^*(\pi_m)$.

Assumption 2 For each i , the mean function $y_i(\boldsymbol{\theta})$ is continuous in $\boldsymbol{\theta}$.

The definition of the favorable set and Assumption 2 together imply that Θ_i is a closed set, whereas Θ_i^c is open for all i .

Assumption 3 For each $1 \leq \ell \leq L$, $\mathcal{F}_{\Theta^\ell}$ and \mathcal{F}_ℓ are σ -algebras of prior π_0^ℓ and $Z \sim f_{\theta_0^\ell}^\ell$, respectively, and the following holds:

- (a) Θ^ℓ is compact;
- (b) For all m_ℓ , the product likelihood, $f_{\theta^\ell}^\ell(\mathcal{Z}_{m_\ell}^\ell)$, is $\mathcal{F}_{\Theta^\ell} \times \mathcal{F}_\ell^{m_\ell}$ -measurable, where $\mathcal{F}_\ell^{m_\ell}$ is a product σ -algebra with respect to $\mathcal{Z}_{m_\ell}^\ell$;
- (c) The class, $\mathcal{H} := \left\{ \log \left(f_{\theta^\ell}^\ell / f_{\theta_0^\ell}^\ell \right) : \theta^\ell \in \Theta^\ell \right\}$, is $P_{\theta_0^\ell}$ -Glivenko-Cantelli;
- (d) $D_{\text{KL}}(\theta_0^\ell || \theta^\ell)$, is continuous in θ^ℓ and $0 < \inf_{\theta^\ell \in \Theta^\ell} \pi_0^\ell(\theta^\ell) \leq \sup_{\theta^\ell \in \Theta^\ell} \pi_0^\ell(\theta^\ell) < \infty$, where $D_{\text{KL}}(\theta_0^\ell || \theta^\ell) := E_{Z \sim f_{\theta_0^\ell}^\ell} \left[\log \left(f_{\theta_0^\ell}^\ell(Z) / f_{\theta^\ell}^\ell(Z) \right) \right]$ is the Kullback-Leibler (KL) divergence.

Given probability measure P , function class \mathcal{H} is said to be a *P-Glivenko-Cantelli* (GC) class if $\|P_m - P\|_{\mathcal{H}} := \sup_{f \in \mathcal{H}} |P_m f - P f| \rightarrow 0$, P-a.s where P_m is an empirical measure consisting of m i.i.d observations from P and $P f = \int f dP$. In other words, the strong law of large numbers holds uniformly over the GC class. Assumption 4 is made to investigate consistency and asymptotic analysis of $i^*(\pi_m)$.

Assumption 4 For each ℓ , $m_\ell \rightarrow \infty$ as $m \rightarrow \infty$, and $\lim_{m \rightarrow \infty} m_\ell / m = \beta_\ell$ for some $\beta_\ell \geq 0$.

Under these assumptions, Theorem 1 below stipulates strong consistency of $i^*(\pi_m)$.

Theorem 1 Under Assumptions 1 to 4, $i^*(\pi_m)$ converges to i_0 , P_{θ_0} -almost surely.

Recall that our goal is to increase $P_{i^*(\pi_m)}(\pi_m)$ (or equivalently, decrease $1 - P_{i^*(\pi_m)}(\pi_m)$) at the optimal rate. Analyzing $P_{i^*(\pi_m)}(\pi_m)$ for finite m is difficult. Instead, Theorem 2 investigates the asymptotic convergence rate of $1 - P_{i^*(\pi_m)}(\pi_m)$ when input data are collected from L processes according to static sampling ratios as stated in Assumption 4.

Theorem 2 Suppose Assumptions 1–4 hold. Then, we have

$$\lim_{m \rightarrow \infty} -\frac{1}{m} \log(1 - P_{i^*(\pi_m)}(\pi_m)) = \inf_{\boldsymbol{\theta} \in \Theta_0^c} \sum_{\ell=1}^L \beta_\ell D_{\text{KL}}(\theta_0^\ell || \theta^\ell). \tag{2}$$

Therefore, the MPB’s posterior preference converges to 1 at an exponential rate that depends on the sampling ratios, $\boldsymbol{\beta} = (\beta_\ell)_{1 \leq \ell \leq L}$. Thus, the optimal $\boldsymbol{\beta}$ that maximizes (2) is a solution to

$$\max_{\boldsymbol{\beta}=(\beta_1, \dots, \beta_L)} \inf_{\boldsymbol{\theta} \in \Theta_0^c} \sum_{\ell=1}^L \beta_\ell D_{\text{KL}}(\theta_0^\ell || \theta^\ell) \text{ subject to } \boldsymbol{\beta} \geq \mathbf{0}, \mathbf{1}^\top \boldsymbol{\beta} = 1, \tag{3}$$

where $\mathbf{0}$ and $\mathbf{1}$ are L -dimensional vectors of zeroes and ones, respectively. However, (3) depends on unknown quantities such as i_0 , Θ_{i_0} , and $\boldsymbol{\theta}_0$. Instead, we consider the following plug-in version by replacing i_0 , Θ_{i_0} , and $\boldsymbol{\theta}_0$ with $i^*(\pi_m)$, $\Theta_{i^*(\pi_m)}$, and $\hat{\boldsymbol{\theta}}$, respectively:

$$\max_{\boldsymbol{\beta}=(\beta_1, \dots, \beta_L)} \inf_{\boldsymbol{\theta} \in \Theta_{i^*(\pi_m)}^c} \sum_{\ell=1}^L \beta_\ell D_{\text{KL}}(\hat{\boldsymbol{\theta}}^\ell || \theta^\ell) \text{ subject to } \boldsymbol{\beta} \geq \mathbf{0}, \mathbf{1}^\top \boldsymbol{\beta} = 1. \tag{4}$$

Because $i^*(\pi_m)$ and $\Theta_{i^*(\pi_m)}^c$ are unknown in advance, we first need to estimate $i^*(\pi_m)$ and $\Theta_{i^*(\pi_m)}^c$ by simulation before solving (4). In Section 4, we introduce an R&S framework proposed by Kim et al. (2022) that efficiently learns $i^*(\pi_m)$ and $\Theta_{i^*(\pi_m)}^c$ by solving an optimal computing budget allocation (OCBA) problem. Since this procedure is designed for the case when $\boldsymbol{\theta}$ has a finite support, we assume the following to hold in the remainder of the paper.

Assumption 5 For each ℓ , the support of π_0^ℓ , \mathcal{S}_ℓ , is finite and $\theta_0^\ell \in \mathcal{S}_\ell$.

In the appendix, we discuss how to construct an approximate finite support when θ is continuous.

An additional benefit of Assumption 5 is that it reduces (4) to a linear program (LP). When θ is continuous, (4) is difficult to solve in general since we do not impose any structural property (e.g., linearity or convexity) of $\{\Theta_i\}$ in θ . With the discrete support, we can construct finite $\Theta_{i^*}^c(\pi_m) = \{\theta_1, \theta_2, \dots, \theta_J\}$, where J is the size of $\Theta_{i^*}^c(\pi_m)$, and the infimum in (4) can be replaced with inequalities over the J elements in $\Theta_{i^*}^c(\pi_m)$. Consequently, (4) can be rewritten as the following LP:

$$\max C \quad \text{subject to} \quad C \leq \mathbf{d}_j^\top \boldsymbol{\beta}, 1 \leq j \leq J, \quad \mathbf{1}^\top \boldsymbol{\beta} = 1, \quad \boldsymbol{\beta} \geq \mathbf{0}, \quad (5)$$

where $\mathbf{d}_j := (\text{D}_{\text{KL}}(\hat{\theta}^\ell || \theta_j^\ell))_{1 \leq \ell \leq L}$ for each θ_j in $\Theta_{i^*}^c(\pi_m)$.

4 LEARNING THE MPB AND ITS ADVERSARIAL SET

Throughout this section, we assume π_m is fixed and focus on learning $i^*(\pi_m)$ and $\Theta_{i^*}^c(\pi_m)$ for given π_m . Under Assumption 5, the support of θ is the Cartesian product of $\{\mathcal{S}_\ell\}_{1 \leq \ell \leq L}$, $\mathcal{S} := \mathcal{S}_1 \times \dots \times \mathcal{S}_L$, and the posterior preference for each solution simplifies to

$$P_j(\pi_m) = \sum_{\theta_s \in \mathcal{S}} \pi_m(\theta_s) 1\{j = i^s\},$$

where i^s is a conditional optimum at θ_s . Each i^s must be estimated via simulation. Denote the number of replications allocated to (i, θ_s) given simulation budget n as $N_i^n(\theta_s)$ and the sample mean of the $N_i^n(\theta_s)$ replication outputs as $\mu_{i,n}(\theta_s)$. Define the sample version of i^s as $i_n^s := \arg \min_{1 \leq j \leq k} \mu_{j,n}(\theta_s)$. The empirical posterior preference $P_{j,n}(\pi_m)$ is defined as $P_{j,n}(\pi_m) = \sum_{\theta_s \in \mathcal{S}} \pi_m(\theta_s) 1\{j = i_n^s\}$. From these quantities, we estimate the MPB and its favorable set as $i_n^*(\pi_m) := \arg \max_{1 \leq j \leq k} P_{j,n}(\pi_m)$ and $\hat{\Theta}_{i_n^*(\pi_m)} := \{\theta_s \in \mathcal{S} | i_n^*(\pi_m) = i_n^s\}$, respectively. To solve (5) correctly, our goal is to maximize the probability of the following event

$$F_{m,n} := \{i_n^*(\pi_m) = i^*(\pi_m), \hat{\Theta}_{i_n^*(\pi_m)} = \Theta_{i^*(\pi_m)}\}. \quad (6)$$

Equivalently, we aim to minimize $\text{P}(F_{m,n}^c)$. However, computing $\text{P}(F_{m,n}^c)$ is challenging for finite n .

Instead, we first consider an asymptotic budget allocation problem ($n \rightarrow \infty$) to determine the static sampling ratios for all (i, θ_s) that maximize the exponential convergence rate of $\text{P}(F_{m,n}^c)$ under Assumption 6 below.

Assumption 6 For each (i, θ) , the simulation output is normally distributed with mean $y_i(\theta)$ and known variance $\lambda_i^2(\theta)$.

For a nonnormal case, Assumption 6 can be justified by batching (Kim and Nelson 2006). For an unknown variance case, one can plug in a sample variance, and our theoretical results still hold. Let $\boldsymbol{\alpha} = \{\alpha_i(\theta_s)\}_{1 \leq i \leq k, \theta_s \in \mathcal{S}}$, where $\alpha_i(\theta_s)$ stands for the proportion of simulation budget allocated to (i, θ_s) . Under Assumption 6, Kim et al. (2022) provide an easy-to-compute lower bound of the large deviation rate (LDR) of $\text{P}(F_{m,n}^c)$ as stated in the following theorem.

Theorem 3 (Kim et al. 2022) Suppose Assumption 6 holds. Then, for fixed π_m , the LDR of $F_{m,n}^c$ can be bounded from below as

$$\liminf_{n \rightarrow \infty} -\frac{1}{n} \log \text{P}(F_{m,n}^c) \geq \min_{(i, \theta_s): i \neq i^s} W_{i,m}(\theta_s) G_i(\theta_s),$$

where $G_i(\theta_s)$ is an LDR of the incorrect pairwise comparison between (i, θ_s) and (i^s, θ_s) defined as

$$G_i(\theta_s) := \frac{(y_i(\theta_s) - y_{i^s}(\theta_s))^2}{2(\lambda_i^2(\theta_s)/\alpha_i(\theta_s) + \lambda_{i^s}^2(\theta_s)/\alpha_{i^s}(\theta_s))},$$

and the balance weight, $W_{i,m}(\boldsymbol{\theta}_s)$, is given as

$$W_{i,m}(\boldsymbol{\theta}_s) = \begin{cases} 1, & \text{if } \boldsymbol{\theta}_s \in \Theta_{i^*}(\boldsymbol{\pi}_m) \text{ or } i = i^*(\boldsymbol{\pi}_m), \boldsymbol{\theta}_s \in \Theta_{i^*}^c(\boldsymbol{\pi}_m), \\ \max(\{P_{i^*}(\boldsymbol{\pi}_m) - P_i(\boldsymbol{\pi}_m)\}/\pi_m(\boldsymbol{\theta}_s), 1), & \text{otherwise.} \end{cases} \quad (7)$$

Note that $W_{i,m}(\boldsymbol{\theta}_s)$ quantifies importance of learning $y_i(\boldsymbol{\theta}_s)$ in reducing the chance of $F_{m,n}^c$. Define $\underline{\text{LDR}}_m = \min_{(i,\boldsymbol{\theta}_s):i \neq i^*} W_{i,m}(\boldsymbol{\theta}_s)G_i(\boldsymbol{\theta}_s)$. From Theorem 3, we can formulate the following OCBA:

$$\max_{\boldsymbol{\alpha}} \underline{\text{LDR}}_m \text{ subject to } \sum_{(i,\boldsymbol{\theta}_s)} \alpha_i(\boldsymbol{\theta}_s) = 1, \alpha_i(\boldsymbol{\theta}_s) \geq 0, \forall (i, \boldsymbol{\theta}_s). \quad (8)$$

The optimality condition for (8) can be derived from the Karush-Kuhn-Tucker (KKT) conditions; Theorem 4 below summarizes the result.

Theorem 4 (Kim et al. 2022) The allocation $\boldsymbol{\alpha}^*$ is an optimal solution of (8), if and only if, $\boldsymbol{\alpha}^*$ satisfies:

- (Global balance condition) For all $\boldsymbol{\theta}_s$, we have $\alpha_{i^*}^2(\boldsymbol{\theta}_s)/\lambda_{i^*}^2(\boldsymbol{\theta}_s) = \sum_{i \neq i^*} \alpha_i^2(\boldsymbol{\theta}_s)/\lambda_i^2(\boldsymbol{\theta}_s)$;
- (Pairwise balance condition) $W_{i,m}(\boldsymbol{\theta}_s)G_i(\boldsymbol{\theta}_s) = W_{j,m}(\boldsymbol{\theta}_{s'})G_j(\boldsymbol{\theta}_{s'}), \forall (i, \boldsymbol{\theta}_s), (j, \boldsymbol{\theta}_{s'}) \in \{(i, \boldsymbol{\theta}_i) : i \neq i^*\}$.

The conditions in Theorem 4 depend on unknown means $\{y_i(\boldsymbol{\theta}_s)\}$. In Algorithm 2 in Section 5.3, we present a sequential sampling rule based on the plug-in version of Theorem 4, where each $y_i(\boldsymbol{\theta}_s)$ is replaced with $\mu_{i,n}(\boldsymbol{\theta}_s)$.

We close this section by introducing notation for our algorithm. The empirical allocation ratio $\alpha_{i,n}(\boldsymbol{\theta}_s)$ is defined as $\alpha_{i,n}(\boldsymbol{\theta}_s) := N_i^n(\boldsymbol{\theta}_s)/n$. Similarly, $G_{i,n}(\boldsymbol{\theta}_s)$ and $W_{i,m,n}(\boldsymbol{\theta}_s)$ are plug-in versions of $G_i(\boldsymbol{\theta}_s)$ and $W_{i,m}(\boldsymbol{\theta}_s)$, respectively.

5 OPTIMAL INPUT DATA ACQUISITION

In this section, we introduce our sequential sampling procedure that balances input data collection and simulation sampling. Algorithm 0 summarizes the overall procedure with the following subroutines:

- Stopping Rule: a criterion that determines when to stop collecting more input data;
- Data Collection vs Simulation (DCvS): a criterion to determine whether to collect input data or to simulate at some $(i, \boldsymbol{\theta})$. We collect input data if this condition is met;
- Input Data Collection Rule: a sequential sampling rule for input data acquisition;
- Simulation Sampling Rule: a sequential sampling rule for selecting the next $(i, \boldsymbol{\theta})$ to simulate.

Several sample statistics are updated throughout the algorithm; we categorize them into three groups:

- S.1** Updated with new input data: data size vector $(m_\ell)_{1 \leq \ell \leq L}$, posterior $\{\pi_m(\boldsymbol{\theta}_s)\}$, and MAP $\hat{\boldsymbol{\theta}}$
- S.2** Updated with new simulation sample: simulation sample sizes $\{N_i^n(\boldsymbol{\theta}_s)\}$, allocation ratio $\{\alpha_{i,n}(\boldsymbol{\theta}_s)\}$, and sample LDR $\{G_{i,n}(\boldsymbol{\theta}_s)\}$
- S.3** Updated in either case: empirical posterior preference $\{P_{j,n}(\boldsymbol{\pi}_m)\}$, MPB $i_n^*(\boldsymbol{\pi}_m)$, its estimated adversarial set $\hat{\Theta}_{i_n^*}^c(\boldsymbol{\pi}_m)$, and balance weights $\{W_{i,m,n}(\boldsymbol{\theta}_s)\}$

5.1 Data Collection versus Simulation

Recall that in Section 3, we discuss finding optimal $\boldsymbol{\beta}$ that maximizes the convergence rate of $P_{i^*}(\boldsymbol{\pi}_m) = P_{\pi_m}(\Theta_{i^*}(\boldsymbol{\pi}_m))$. Because we estimate $i^*(\boldsymbol{\pi}_m)$ by simulation, this objective is modified to incorporate simulation error. The OCBA discussed in Section 4 aims to maximize the probability of event $F_{m,n}$ defined in (6). If $P(F_{m,n})$ is close to 1, then input data acquisition may be optimized conditional on $F_{m,n}$ while ignoring its complement, $F_{m,n}^c$. Then, the goal is to maximize $P_{\pi_m}(\Theta_{i_n^*}(\boldsymbol{\pi}_m))1\{F_{m,n}\}$ instead of $P_{\pi_m}(\Theta_{i^*}(\boldsymbol{\pi}_m))$. After taking

Algorithm 0 Unified Framework

- 1: **Warm start:** Collect m_0 observations from each of the L input data sources and run n_0 replications at each $(i, \boldsymbol{\theta}_s)$.
 - 2: **while** Stopping Rule (14) is not met **do**
 - 3: **if** DCvS (12) is met **then**
 - 4: Allocate data collection budget based on Input Data Collection Rule (= Algorithm 1).
 - 5: **else**
 - 6: Find a pair $(i, \boldsymbol{\theta}_s)$ to be sampled based on Sequential Simulation Rule (= Algorithm 2).
 - 7: **end if**
 - 8: **end while**
-

the expectation over simulation error, we have

$$\mathbb{E} [\mathbb{P}_{\pi_m} (\Theta_{i_n^*}(\pi_m)) \mathbb{1} \{F_{m,n}\}] = \mathbb{E} [\mathbb{P}_{\pi_m} (\Theta_{i^*}(\pi_m)) \mathbb{1} \{F_{m,n}\}] = \mathbb{P}_{\pi_m} (\Theta_{i^*}(\pi_m)) \mathbb{P}(F_{m,n}), \quad (9)$$

where \mathbb{E} and \mathbb{P} are taken with respect to simulation error. The first equation holds from the definition of $F_{m,n}$, and the second is straightforward since $\mathbb{P}_{\pi_m} (\Theta_{i^*}(\pi_m))$ does not depend on n .

We take the view of increasing n just enough so that the overall convergence rate of (9) is not slowed down by n . With this motivation, the following theorem derives the LDR of (9) when the ratio of m to n is at least γ in the limit for some constant $\gamma \geq 0$.

Theorem 5 Suppose Assumptions 1–6 hold and $\liminf_{m,n \rightarrow \infty} m/n = \gamma$ for some $\gamma \geq 0$. Let $\underline{\text{LDR}} := \lim_{m \rightarrow \infty} \underline{\text{LDR}}_m$. Then, for fixed $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$ we have

$$\liminf_{m,n \rightarrow \infty, \liminf \frac{m}{n} = \gamma} -\frac{1}{n} \log (1 - \mathbb{P}_{\pi_m} (\Theta_{i^*}(\pi_m)) \mathbb{P}(F_{m,n})) \geq \min \left\{ \gamma \inf_{\boldsymbol{\theta} \in \Theta_{i_0}^c} \sum_{\ell=1}^L \beta_{\ell} \text{D}_{\text{KL}}(\theta_0^{\ell} || \boldsymbol{\theta}^{\ell}), \underline{\text{LDR}} \right\}. \quad (10)$$

Recall that $\underline{\text{LDR}}_m$ depends on the simulation allocation ratios, $\boldsymbol{\alpha}$, therefore, $\underline{\text{LDR}}$ in turn depends on $\boldsymbol{\alpha}$. Also, observe that $\underline{\text{LDR}}$ depends on $\lim_{m \rightarrow \infty} W_{i,m}(\boldsymbol{\theta}_s)$ by definition. We can further see that $\underline{\text{LDR}}$ does not depend on $\boldsymbol{\beta}$ from Definition (7) because the limit of each $W_{i,m}(\boldsymbol{\theta}_s)$ is uniquely determined regardless of $\boldsymbol{\beta}$. By viewing $\boldsymbol{\alpha}, \boldsymbol{\beta}$, and γ as decision variables, the lower bound in (10) can be maximized by solving:

$$\max_{\boldsymbol{\alpha}, \boldsymbol{\beta}, \gamma} \min \left\{ \underbrace{\gamma \inf_{\boldsymbol{\theta} \in \Theta_{i_0}^c} \sum_{\ell=1}^L \beta_{\ell} \text{D}_{\text{KL}}(\theta_0^{\ell} || \boldsymbol{\theta}^{\ell})}_{=g_1(\boldsymbol{\beta})}, \underline{\text{LDR}} \right\} = \max_{\gamma \geq 0} \min \{ \gamma g_1(\boldsymbol{\beta}^*), g_2(\boldsymbol{\alpha}^*) \}, \quad (11)$$

where $\boldsymbol{\beta}^*$ and $\boldsymbol{\alpha}^*$ are maximizers of g_1 and g_2 , respectively. Observe that (11) is maximized as long as $\gamma \geq g_2(\boldsymbol{\alpha}^*)/g_1(\boldsymbol{\beta}^*)$ while minimal m is spent when the equality holds. Therefore, our strategy is to increase m if m/n falls below $g_2(\boldsymbol{\alpha}^*)/g_1(\boldsymbol{\beta}^*)$. Of course $\boldsymbol{\alpha}^*$ and $\boldsymbol{\beta}^*$ are unknown. However, we can use the respective sample allocation ratios, $\{\alpha_{i,n}(\boldsymbol{\theta}_s)\}$ and $\{m_{\ell}/m\}$, to construct the following criterion: gather additional input data if

$$\min_{\boldsymbol{\theta} \in \widehat{\Theta}_{i_n^*}^c(\pi_m)} \sum_{\ell=1}^L m_{\ell} \text{D}_{\text{KL}}(\widehat{\boldsymbol{\theta}}^{\ell} || \boldsymbol{\theta}^{\ell}) < n \underline{\text{LDR}}_{m,n} := n \min_{i \neq i_n^s} W_{i,m,n}(\boldsymbol{\theta}_s) G_{i,n}(\boldsymbol{\theta}_s), \quad (12)$$

and run simulation, otherwise. Observe that (12) seeks to keep the minimum input data requirements by increasing m only when $m \leq g_2(\boldsymbol{\alpha}^*)/g_1(\boldsymbol{\beta}^*)n$ in an asymptotic sense. Hence, one may expect that $\gamma \approx m/n \approx g_2(\boldsymbol{\alpha}^*)/g_1(\boldsymbol{\beta}^*)$ would be achieved in the limit.

Algorithm 1 Optimal Input Data Acquisition

- 1: **Initialization:** $\Delta :=$ batch input data size, **S.1**, and **S.3**.
- 2: Evaluate $\widehat{\mathbf{d}}_j = (\text{D}_{\text{KL}}(\widehat{\theta}^\ell || \theta_j^\ell))_{1 \leq \ell \leq L}$ for each $\theta_j \in \widehat{\Theta}_{i_n^*(\pi_m)}^c$ for all $1 \leq j \leq J_n$ ($:= |\widehat{\Theta}_{i_n^*(\pi_m)}^c|$).
- 3: **if** $\widehat{\theta} \in \widehat{\Theta}_{i_n^*(\pi_m)}^c$ **then**
- 4: Return the uniform allocation rule $\mathbf{v}_* = (\Delta/L, \dots, \Delta/L)$.
- 5: **else**
- 6: Solve the following LP, and denote the optimizer as \mathbf{v}_* :

$$\max_{\mathbf{v}=(v_1, \dots, v_L)} C \text{ subject to } C \leq \widehat{\mathbf{d}}_j^\top (\mathbf{m} + \mathbf{v}), \forall 1 \leq j \leq J_n, \mathbf{1}^\top \mathbf{v} = \Delta, \mathbf{v} \geq \mathbf{0}. \quad (13)$$

- 7: **end if**
 - 8: Collect additional input data according to \mathbf{v}_* .
 - 9: Update **S.1** and **S.3**.
-

Algorithm 2 Simulation Sampling Rule

- 1: **Initialization:** All quantities in **S.1– S.3**.
 - 2: **while** $\min_{\theta \in \widehat{\Theta}_{i_n^*(\pi_m)}^c} \sum_{\ell=1}^L m_\ell \text{D}_{\text{KL}}(\widehat{\theta}^\ell || \theta^\ell) > n \min_{i \neq i_n^*} W_{i,m,n}(\theta_s) G_{i,n}(\theta_s)$ **do**
 - 3: Find $(i, \theta_s) := \arg \min_{i \neq i_n^*} W_{i,m,n}(\theta_s) G_{i,n}(\theta_s)$.
 - 4: **if** $(N_{i_n^*}^n(\theta_s) / \lambda_{i_n^*}(\theta_s))^2 > \sum_{i \neq i_n^*} (N_i^n(\theta_s) / \lambda_i(\theta_s))^2$ **then**
 - 5: Run a replication at (i, θ_s) .
 - 6: **else**
 - 7: Run a replication at (i_n^s, θ_s) .
 - 8: **end if**
 - 9: Update **S.2** and **S.3**.
 - 10: **end while**
-

5.2 Input Data Collection Problem

Once (12) is satisfied, the algorithm proceeds to the input data collection problem. In this step, we solve the empirical version of (5). Algorithm 1 formalizes this procedure when the input data are collected in batch of size Δ .

Notice that, if $\widehat{\theta} \in \widehat{\Theta}_{i_n^*(\pi_m)}^c$, then there exists j such that $\widehat{\mathbf{d}}_j = \mathbf{0}$, and (13) cannot give us any allocation rule \mathbf{v} because the optimal value is always zero regardless of \mathbf{v} . However, this event becomes rarer as m increases since $\pi_m(\widehat{\theta}) \rightarrow 1$ almost surely so that $\widehat{\theta} \in \widehat{\Theta}_{i_n^*(\pi_m)}^c$ holds eventually. Meanwhile, when $\widehat{\theta} \in \widehat{\Theta}_{i_n^*(\pi_m)}^c$, Algorithm 1 allocates Δ uniformly to all L input distributions.

5.3 Simulation Sampling Rule

If (12) is not satisfied, then we increase n by sequentially selecting the next (i, θ_s) to simulate until it is satisfied. Algorithm 2 presents the sequential sampling rule motivated by the pairwise and global balance conditions in Theorem 3. We find (i, θ_s) that deviates the most from the pairwise balance condition, and then determine whether to sample the selected (i, θ_s) or (i_n^s, θ_s) based on the global balance condition.

5.4 Stopping Criterion

When m and n are large enough to meet a precision requirement, we stop collecting input data and simulation sampling. In particular, we adopt a variation of the posterior-based PCS (pPCS) stopping rule discussed in Eckman and Henderson (2022).

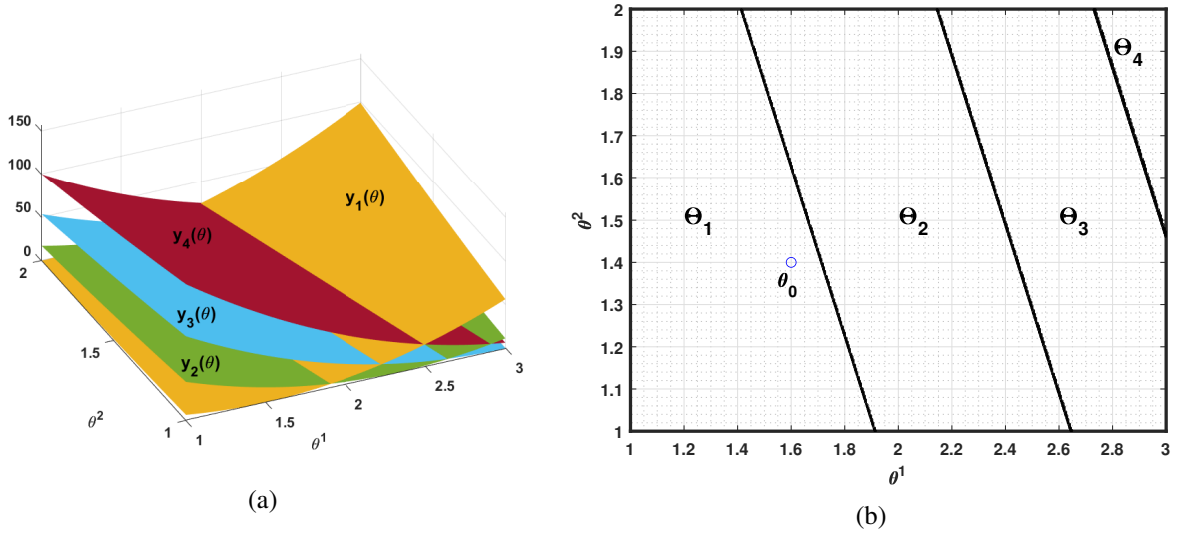


Figure 1: For Solutions 1–4, (a) shows the conditional mean surfaces and (b) shows the favorable sets; observe that Solution 1 is the optimum under θ_0 .

Consider a policy that terminates the procedure when the estimated MPB’s posterior preference is above user-specified threshold $1 - \alpha$ (e.g., $1 - \alpha = 0.95$); that is, $P_{i_n^*(\pi_m), n}(\pi_m) \geq 1 - \alpha$. However, this policy does not account for the simulation error in estimating $i_n^*(\pi_m)$. If n is small, the procedure may stop when $i_n^*(\pi_m)$ is a poor estimate of $i^*(\pi_m)$. To circumvent this issue, we modify the stopping rule as follows:

$$P_{i_n^*(\pi_m), n}(\pi_m) \geq 1 - \alpha \text{ and } \min_{\theta \in \hat{\Theta}_{i_n^*(\pi_m)}} \sum_{\ell=1}^L m_\ell \text{D}_{\text{KL}}(\hat{\theta}^\ell || \theta^\ell) < n \underline{\text{LDR}}_{m, n}. \quad (14)$$

Following the same logic in Section 5.1, the second condition of (14) guarantees that n is large enough relative to m .

6 NUMERICAL EXPERIMENTS

In this section, we present numerical results with a synthetic example. We assume that there are $k = 10$ solutions and two exponential input distributions with respective mean parameters θ^1 and θ^2 , where $\theta_0^1 = 1.6$ and $\theta_0^2 = 1.4$. The priors, $\pi_0^1(\theta^1)$ and $\pi_0^2(\theta^2)$, follow $U[1, 3]$ and $U[1, 2]$, respectively. Since these choices make θ continuous, we apply the grid approximation in the appendix to discretize its support. We adopt uniform grids $\{1 + 2j/(S_1 - 1)\}_{0 \leq j \leq S_1 - 1}$ and $\{1 + j/(S_2 - 1)\}_{0 \leq j \leq S_2 - 1}$ for θ^1 and θ^2 , respectively, and set $S = S_1 = S_2$ for simplicity. Further, for each (i, θ) , mean surface $y_i(\theta)$ is given as

$$y_i(\theta) = (\mathbf{a}^\top \theta - 10\sqrt{i})^2, \quad (15)$$

where $\mathbf{a} = [5 \ 5/2]^\top$, and its simulation error is normally distributed with variance $\lambda_i^2(\theta)$.

We observe the performance of our algorithm varying three factors: grid size S , data batch size Δ , and simulation variances $\{\lambda_i^2(\theta)\}$. The baseline case has $S = 11, \Delta = 50$, and $\lambda_i(\theta) \sim U[1, 3]$ for all $1 \leq i \leq k, \theta \in \mathcal{S}$. We observe six additional scenarios by varying $S \in \{31, 51\}$, $\Delta \in \{20, 100\}$, and $\{U[0.5, 1.5], U[2, 6]\}$ for the sampling distribution of $\lambda_i(\theta)$ while fixing all other factors to be the same as the baseline. For all scenarios, we examine two confidence levels for stopping criterion: $1 - \alpha \in \{0.95, 0.99\}$.

Given (15), only Solutions 1–4 can be a conditional optimum in $\Theta = [1, 3] \times [1, 2]$. Figure 1a shows mean surfaces $\{y_i(\theta)\}_{1 \leq i \leq 4}$. Figure 1b divides $[1, 3] \times [1, 2]$ into the favorable sets of Solutions 1–4, $\{\Theta_i\}_{1 \leq i \leq 4}$. Notice that θ_0 belongs to Θ_1 , which implies $i_0 = 1$. By design, Assumption 5 is satisfied for all $S \in \{11, 31, 51\}$. For all scenarios, we set $(m_0, n_0) = (20, 5)$ for initial sampling.

Table 1: Numerical performances under the baseline scenario ($S = 11, \Delta = 50, U[1, 3]$) and the additional six scenarios that differs from the baseline by each factor presented in the first column. All statistics are averaged from 1,000 macroruns with relative errors less than 8%.

Scenarios	$1 - \alpha = 0.95$				$1 - \alpha = 0.99$			
	(m_1, m_2)	n	PCS	$\text{freq}\{\hat{\boldsymbol{\theta}} \in \hat{\Theta}_{i_n^*}^c(\pi_m)\}$	(m_1, m_2)	n	PCS	$\text{freq}\{\hat{\boldsymbol{\theta}} \in \hat{\Theta}_{i_n^*}^c(\pi_m)\}$
Baseline	(437, 221)	1.21E+4	0.964	0.072	(756, 341)	1.45E+4	0.993	0.046
$S = 31$	(755, 459)	8.89E+4	0.984	0.063	(1524, 911)	1.41E+5	0.999	0.039
$S = 51$	(725, 411)	1.67E+5	0.978	0.065	(1502, 742)	2.10E+5	1.000	0.037
$\Delta = 20$	(400, 205)	1.28E+4	0.959	0.069	(727, 324)	1.62E+4	0.997	0.045
$\Delta = 100$	(485, 257)	1.17E+4	0.972	0.075	(800, 369)	1.38E+4	0.997	0.051
$U[0.5, 1.5]$	(447, 222)	7.73E+3	0.963	0.065	(774, 343)	8.38E+3	0.992	0.044
$U[2, 6]$	(455, 245)	3.32E+4	0.967	0.082	(773, 356)	4.25E+4	0.992	0.055

For performance measures, we recorded the input data size collected from each data source, (m_1, m_2) , and the total simulation effort, n , spent until Algorithm 0 terminates. Once the procedure stops, we estimated the empirical (frequentist’s) PCS, $P\{i_0 = i_n^*(\pi_m)\}$. Recall that for small m , $\hat{\boldsymbol{\theta}}$ may lie outside of the MPB’s favorable set in which case Algorithm 1 evenly distributes input data sampling effort to all L input sources. We report the frequency of such an event, $\text{freq}\{\hat{\boldsymbol{\theta}} \in \hat{\Theta}_{i_n^*}^c(\pi_m)\}$. Table 1 reports the performance measures averaged from 1,000 macroruns of each scenario.

For all scenarios, higher m and n are required to stop as $1 - \alpha$ increases, which is intuitive. Notice that the estimated frequentist’s PCS matches the Bayesian target, $1 - \alpha$ quite well in all cases. The frequency of uniform allocation (i.e., $\text{freq}\{\hat{\boldsymbol{\theta}} \in \hat{\Theta}_{i_n^*}^c(\pi_m)\}$) decreases as $1 - \alpha$ increases, indicating that $\hat{\boldsymbol{\theta}}$ is more likely to belong to the MPB’s favorable set as m grows. Hence, we can confirm that the uniform allocation has little impact on the large-sample behavior of the algorithm.

Next, we compare the performance of Algorithm 0 under the baseline scenario vs. the others. As the number of $\boldsymbol{\theta}^s$ on the grid increases, the number of constraints in (5) increases. As a result, optimal C of (5) becomes smaller, which results in (seemingly) slower convergence rate for $1 - P_{i(\pi_m)}(\pi_m)$. This explains why required m_1 and m_2 until termination increases as S increases. Moreover, for larger S , there are more “opportunities” for $F_{n,m}^c$ to occur and thus required n increases. When Δ increases, the ratios between m_1 and m_2 as well as n remain similar to the baseline case, which shows that the algorithm is pretty robust to the choice of Δ . Lastly, as the average simulation variance increases, higher n is required for the algorithm to terminate, however, acquired input data sizes are similar to the baseline. This makes sense as the input data acquisition strategy should not be affected by the simulation error variance if the MPB and its adversarial set are learned correctly. This implies that our Simulation Sampling Rule in Algorithm 2 is effective.

ACKNOWLEDGMENTS

This research is supported by National Science Foundation grant CAREER CMMI-2045400.

A APPENDIX: DISCRETE APPROXIMATION FOR A CONTINUOUS INPUT PARAMETER

In this section, we discuss discretizing the support of π_m using a grid approximation when $\boldsymbol{\theta}$ is continuous. For the ℓ -th input model, we denote the discretized support by $\mathcal{S}_\ell = \{\boldsymbol{\theta}_1^\ell, \boldsymbol{\theta}_2^\ell, \dots, \boldsymbol{\theta}_{S_\ell}^\ell\}$. Recall that we assume compactness of support in Assumption 3. In this case, one can generate a random sample of size $|\mathcal{S}_\ell|$ by sampling from prior π_0^ℓ as π_0^ℓ is always proper. One can also generate a Quasi Monte Carlo (QMC) sample or a uniform grid on the support can be adopted.

Once \mathcal{S}_ℓ is constructed, we assign a normalized weight to each $\boldsymbol{\theta}_s^\ell \in \mathcal{S}_\ell$ proportional to the posterior probability distribution function evaluated at $\boldsymbol{\theta}_s^\ell$. The resulting approximate probability mass function




Figure 2: For varying input data sizes, approximate pmfs of θ^1 and θ^2 respectively with eleven grid points are shown. The input-data-generating distributions are exponential with mean parameters θ_0^1 and θ_0^2 , respectively.

(pmf), $p_{m_\ell}^\ell$, satisfies

$$p_{m_\ell}^\ell(\theta_s^\ell) \propto \pi_{m_\ell}^\ell(\theta_s^\ell) \propto f_{\theta_0^\ell}^\ell(\mathcal{Z}_{m_\ell}^\ell) \pi_0^\ell(\theta_s^\ell). \tag{16}$$

From (16), we can compute $p_{m_\ell}^\ell(\theta_s^\ell)$ without the normalizing constant of $\pi_{m_\ell}^\ell$. Moreover, updating (16) when there are new input data is simple; we only need to multiply the likelihood function of the new data. The joint posterior pmf of L input parameters is denoted by p_m , which assigns probability mass $p_m(\boldsymbol{\theta}) = \prod_{\ell=1}^L p_{m_\ell}^\ell(\theta^\ell)$ to each $\boldsymbol{\theta}$ on the grid.

Proposition 1 below stipulates a limiting behavior of weights $\{p_{m_\ell}^\ell(\theta^\ell)\}_{\theta^\ell \in \mathcal{S}_\ell}$ as m_ℓ increases. Although we omit its proof, we provide an intuitive explanation here. As $m_\ell \rightarrow \infty$, one can show that $\pi_{m_\ell}^\ell(\theta^\ell) \approx \exp(-m_\ell D_{\text{KL}}(\theta_0^\ell || \theta^\ell))$ under Assumption 3. Therefore, p_m tends to assign the largest probability mass to the point in \mathcal{S}_ℓ minimizing the KL divergence from θ_0^ℓ , and the probability converges to one while those of other points converge to zero in the limit.

Proposition 1 For each ℓ , let $\theta_*^\ell := \arg \min_{\theta^\ell \in \mathcal{S}_\ell} D_{\text{KL}}(\theta_0^\ell || \theta^\ell)$. Under Assumption 3, $\lim_{m_\ell \rightarrow \infty} p_{m_\ell}^\ell(\theta^\ell) = 1\{\theta^\ell = \theta_*^\ell\}$ almost surely.

When the true parameter is included in the grid, i.e., $\theta_0^\ell \in \mathcal{S}_\ell$, for all ℓ , the probability mass on the grid concentrates on θ_0^ℓ as data size grows and i_0 is invariant with the grid approximation. However, we cannot guarantee $\theta_0^\ell \in \mathcal{S}_\ell$ in practice since all $\{\theta_0^\ell\}$ are unknown in advance. In this case, the approximation scheme will assign probability one on some $\boldsymbol{\theta}_* := (\theta_*^\ell)_{1 \leq \ell \leq L}$ in the limit. Hence, $\boldsymbol{\theta}_*$ can be viewed as a stand-in for $\boldsymbol{\theta}_0$ in the approximation scheme. Even if $\boldsymbol{\theta}_* \neq \boldsymbol{\theta}_0$, observe that the MPB converges to the same i_0 as long as $\boldsymbol{\theta}_* \in \Theta_{i_0}$. For ease of discussion, we assume that \mathcal{S}_ℓ contains $\boldsymbol{\theta}_0$ in this paper (Assumption 5); a more general case will be investigated in the journal version of the paper.

Recall that (4) replaces unknown $\boldsymbol{\theta}_0$ with its point estimator $\hat{\boldsymbol{\theta}}$. Depending on the prior and the density function, computing MAP may be expensive for continuous $\boldsymbol{\theta}$. With Approximation (16), we can replace the MAP with the approximate MAP, $\hat{\boldsymbol{\theta}}_* = (\theta_{s_\ell}^\ell)_{1 \leq \ell \leq L}$, where $s_\ell := \arg \max_{1 \leq s \leq S_\ell} p_{m_\ell}^\ell(\theta_s^\ell)$.

Figure 2 shows examples of the grid approximation of exponential distributions with mean parameters θ^1 and θ^2 discussed in Section 6 for $S = 11$. Observe that both pmfs concentrate on θ_0^1 and θ_0^2 , respectively, as the sample sizes increase as expected from Proposition 1.

REFERENCES

- 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, T. M. K. R. I. Moon, C. Macal, and M. D. Rossetti, 437–446. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.
- Eckman, D. J., and S. G. Henderson. 2022. "Posterior-Based Stopping Rules for Bayesian Ranking-and-Selection Procedures". *INFORMS Journal on Computing* 34(3):1711–1728.
- Fan, W., L. J. Hong, and X. Zhang. 2020. "Distributionally Robust Selection of the Best". *Management Science* 66(1):190–208.
- Gao, S., H. Xiao, E. Zhou, and W. Chen. 2017. "Robust Ranking and Selection with Optimal Computing Budget Allocation". *Automatica* 81:30–36.
- Kim, K.-K., T. Kim, and E. Song. 2021. "Selection of the Most Probable Best under Input Uncertainty". In *Proceedings of the 2021 Winter Simulation Conference*, edited by S. Kim, B. Feng, K. Smith, S. Masoud, and Z. Zheng. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.
- Kim, S.-H., and B. L. Nelson. 2006. "Selecting the Best System". In *Simulation*, edited by S. G. Henderson and B. L. Nelson, Volume 13 of *Handbooks in Operations Research and Management Science*, 501–534. Amsterdam, Netherlands: Elsevier.
- Kim, T., K.-K. Kim, and E. Song. 2022. "Selection of the Most Probable Best". <https://arxiv.org/abs/2207.07533>.
- Liu, T., Y. Lin, and E. Zhou. 2021. "A Bayesian Approach to Online Simulation Optimization with Streaming Input Data". In *Proceedings of the 2021 Winter Simulation Conference*, edited by S. Kim, B. Feng, K. Smith, S. Masoud, and Z. Zheng, 1–12. Institute of Electrical and Electronics Engineers, Inc.
- Pearce, M., and J. Branke. 2017. "Efficient Expected Improvement Estimation for Continuous Multiple Ranking and 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, 2161–2172. 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., and U. V. Shanbhag. 2019. "Stochastic Approximation for Simulation Optimization under Input Uncertainty with Streaming Data". In *Proceedings of the 2019 Winter Simulation Conference*, edited by N. Mustafee, K.-H. Bae, S. Lazarova-Molnar, M. Rabe, C. Szabo, P. Haas, and Y.-J. Son, 3597–3608. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.
- Ungredda, J., M. Pearce, and J. Branke. 2022. "Bayesian Optimisation vs. Input Uncertainty Reduction". *ACM Transactions on Modeling and Computer Simulation* 32(3):1–26.
- Wu, D., and E. Zhou. 2019. "Ranking and Selection under Input Uncertainty: Fixed Confidence and Fixed Budget". <https://arxiv.org/abs/1708.08526>.
- Xie, W., and E. Zhou. 2015. "Simulation Optimization when Facing 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, 3714–3724. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.
- Xu, J., P. Glynn, and Z. Zheng. 2020. "Joint Resource Allocation for Input Data Collection and Simulation". In *Proceedings of the 2020 Winter Simulation Conference*, edited by K.-H. G. Bae, B. Feng, S. Kim, S. Lazarova-Molnar, Z. Zheng, T. Roeder, and R. Thiesing, 2060–2070. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.
- Zhu, H., T. Liu, and E. Zhou. 2020. "Risk Quantification in Stochastic Simulation under Input Uncertainty". *ACM Transactions on Modeling and Computer Simulation* 30(1):1–24.

AUTHOR BIOGRAPHIES

TAEHO KIM is a postdoctoral fellow in the School of Industrial and Systems Engineering at Georgia Institute of Technology. He received his PhD degree in the Department of Mathematical Sciences at Korea Advanced Institute of Science and Technology. His research focuses on the data-driven simulation modeling of complex system and simulation optimization. His email address is thk5594@gmail.com and his webpage is at <https://sites.google.com/view/taehokim>.

EUNHYE SONG is a Coca-Cola Foundation Early Career Assistant Professor in the School of Industrial and Systems Engineering at Georgia Institute of Technology. She earned her PhD degree in Industrial Engineering and Management Sciences at Northwestern University in 2017. She was awarded the National Science Foundation CAREER Award in 2020. Her research interests include simulation design of experiments, uncertainty and risk quantification, and simulation optimization. Her email address is esong32@gatech.edu and her website can be found at <http://eunhyesong.info>.