

A HYBRID OF SHRINKING BALL METHOD AND OPTIMAL LARGE DEVIATION RATE ESTIMATION IN CONTINUOUS CONTEXTUAL SIMULATION OPTIMIZATION WITH SINGLE OBSERVATION

Xiao Jin

Centre for Next Generation Logistics
National University of Singapore
Singapore, 117602, SINGAPORE

Yichi Shen

Institute of Operations Research and Analytics
National University of Singapore
Singapore, 117602, SINGAPORE

Loo Hay Lee
Ek Peng Chew
Christine A. Shoemaker

Industrial Systems Engineering and Management
National University of Singapore
Singapore, 117576, SINGAPORE

ABSTRACT

We propose a new method for solving continuous contextual simulation optimization with a single observation. By adopting the estimation on the large deviation rate in the contextual ranking and selection problem, we transfer the old theorem to the continuous setting using a shrinking ball inspired construct. Through the estimation of the rate, the new method is expected to achieve the optimal performance in this new problem setting. Brief numerical experiments are conducted and show significant advantages of our method against the uniform sampling scheme.

1 INTRODUCTION

Due to the emerging demand for implementing the simulation methodology in operational problems, the conditional analysis has become more important to simulation study in recent years (Nelson 2016). Take the AGV routing problem in an inventory as an example. We may want to choose the best routing policy that can maximize the utility of AGV using simulation. However, the optimal solution may be different under different demand patterns. In this case, we may want to find optimal one conditioned on each demand scenario. In a conditional analysis of simulation, contextual optimization is one of the essential problems, where the goal is to find the optimal decision conditioned in each scenario. This type of analysis share some features with the contextual multi-armed bandit (cMAB) problem (Zhou 2015) but is unique in only finding the optimal for each scenario rather than an accumulative gain. In particular, when there are a finite number of scenarios and each decision and scenario can be observed multiple times, the problem is called “contextual ranking and selection” (Gao et al. 2019). From the previous studies, the budget allocation with the optimal large deviation rate is derived (Jin et al. 2019). However, for a large portion of real problems, the number of scenarios is large and even continuous. When the scenario is continuous, a finite number of observations is unable to cover all scenarios and decisions. On one hand, we can assume a strong structure (e.g. linear) between the performance and the scenario (Shen, Hong, and Zhang 2017)(Gao, Li, and Du

2019). This, however, causes a modeling risk of assuming a structure distinct from the reality. On the other hand, we can assume a weak structure (e.g. continuity). In this case, surrogate/extrapolation is needed to infer the performance at nearby scenarios. Since the optimal budget allocation is achieved in the discrete scenario, by augment it with the surrogate method, we hope to transform the optimality of the solution in the old problem to the new continuous setting.

There are two fundamental differences between continuous and discrete scenarios. The first one is that in a discrete case, each scenario can be observed multiple times. Therefore, a sample is defined in each scenario. In the continuous scenario case, however, it is less efficient to sample multiple times at the same scenario. However, this would incur a situation where at some scenarios, not all decision is observed. If we want to transform the original formulation into this new setting, the first task is to find a method that can define a local sample at an arbitrary scenario and decision. This is where the shrinking ball method (Linz et al. 2017) comes into the picture. The shrinking ball method is originally a random search method (Kiatsupaibul et al. 2018) in stochastic optimization. The key idea is to estimate the value at a point based on the sample average within a ball of a controlled radius. If the radius keeps decreasing while the sample points within the ball increases, the estimation at that point would converge to the real value. An important feature of the shrinking ball is that it provides a method to define a local sample. Even though the conventional shrinking ball only has a ball at an observed point, we shall borrow this neighborhood construct and augment it to apply at an arbitrary point. The second fundamental difference is that in the continuous case, not only do we need to efficiently reduce the randomness caused by simulation observation (as is done in discrete case), but also to reduce the uncertainty caused by being too far from the observed points. This raises the trade-off problem between “randomness” and “uncertainty”. The optimal allocation in the discrete scenario is meant only for the efficient “randomness” reduction. To include the “uncertainty” factor, we need to augment the original large deviation rate formulation.

In the following sections, we shall first define the continuous contextual simulation optimization rigorously. Then, step by step, we establish an adaptive sampling method by hybridizing the estimation on the large deviation rate and a shrinking ball construct.

2 CONTEXTUAL SIMULATION OPTIMIZATION

The original objective is to find the decision with the optimal expected performance conditioned on each scenario. We denote the set including all possible scenarios as X and each element $x \in X$ represents a scenario. We assume that these scenarios can be parameterized inside a euclidean space (e.g. $X \subset \mathbb{R}^m$). We also assume that based on historical data, we have a good knowledge on the distribution of scenarios. Therefore, x is well defined on a probability space $(X, \mathcal{F}_X, \mu_X)$. Suppose there are finite number of decisions. We index them by $i \in I$. Apart from decision i and scenario x , other factors that have an impact on the simulation output is denoted by $\omega \in \Omega$ and we assume that they are well defined on a probability space $(\Omega, \mathcal{F}_\Omega, \mu_\Omega)$. For each scenario x , the outcome of a decision i with the impact of factor ω is quantified by a real-valued function $G_i(x, \omega)$ (e.g. $G_i : X \times \Omega \mapsto \mathbb{R}$). To emphasize the impact of randomness ω to the outcome, $G_i(x, \omega)$ is also expressed in the following way:

$$G_i(x, \omega) = \tilde{G}_i(x) = g_i(x) + \varepsilon_i(x),$$

where $\varepsilon_i(x)$ is the randomness caused by ω reflected on the simulation output. The performance of each decision i conditioned on each scenario x is indicated by the expected outcome $g_i(x) = E_\omega[G_i(x, \omega)]$. The objective of optimization is to find the best conditional decision with respect to the ranking of $g_i(x)$:

$$i^*(x) = \arg \min_{i \in I} g_i(x)$$

The simulation model can calculate each value of $G_i(x, \omega)$ at a cost of fixed computational effort/time. However, due to the fact that there is randomness ω , it is impossible to achieve a complete knowledge of function $g_i(x)$ (thus $i^*(x)$) in finite time and the estimation is made. In practice, we have a finite

computational budget which translates into N number of budgeted simulation observations. We allocate the budget and then, base on the observations to infer $g_i(x)$ statistically at an arbitrary (i, x) . From a set of observations (e.g. $\{(i, x, \tilde{G}_i(x))_j\}$), we construct our estimation on $g_i(x)$, which is denoted as $\hat{g}_i(x)$. By comparing estimations $\hat{g}_i(x)$ at every scenario x , we have the estimated optimal decision:

$$\hat{i}^*(x) = \arg \min_{i \in I} \hat{g}_i(x)$$

To quantify the effectiveness of applying estimated optimal decision, we adopt the Total Decision Loss (TDL) as the measure:

$$TDL = \int P\{i^*(x) \neq \hat{i}^*(x)|x\}d\mu_X(x) \tag{1}$$

Notice that TDL is affected by three factors:

1. $\mu_X(x)$: The distribution of scenarios
2. $i^*(x)$: The real optimal policy determined by $\{g_i(x)\}$
3. $\hat{i}^*(x)$: The estimated optimal policy determined by i) generated observations $\{\tilde{G}_{i_j}(x_j)\}$ and ii) the inference of $\{\hat{g}_i(x)\}$ by $\{\tilde{G}_{i_j}(x_j)\}$

The focus of this work is to derive an efficient budget allocation scheme which is an adaptive rule of generating observations $\{\tilde{G}_{i_j}(x_j)\}$. To achieve a fair comparison between allocation rules, we need to isolate the impact of other factors and set them the same. Therefore,

1. For the distribution of scenarios, $\mu_X(x)$ is assumed known ahead.
2. For the real optimal policy, we shall benchmark our allocation by the same test problems.
3. For the inference from $\{\tilde{G}_{i_j}(x_j)\}$ to $\{\hat{g}_i(x)\}$, we apply the same shrinking ball construct.

Before we begin devising our rule, we have following assumptions to mildly narrow our discussion:

Assumption 1 $X \subset \mathbb{R}^m$ is a compact probability space $(X, \mathcal{F}_X, \mu_X)$ with a PDF $f_X(x)$.

This assumption guarantees that a proper probability measure can be defined on X such that we could adaptively sample on scenarios that are more likely to occur.

Assumption 2 The expected performance $g_i(x)$ is Lipschitz continuous for all $i \in I$.

This assumption makes sure that the estimation constructed by neighborhood information has a limited bias. This is essential for future analysis on the consistency of the proposed algorithm.

Assumption 3 Each $\varepsilon_{i_j}(x_j)$ follows normal distribution. Each of two $\varepsilon_{i_j}(x_j)$ are mutually independent. The variance of $\varepsilon_i(x)$, denoted $\sigma_i^2(x)$, is uniformly bounded (e.g. $\sigma_i^2(x) < M < \infty$).

This assumption secures the Law of Large Numbers so that the randomness induced by ε vanishes as observation increases.

3 SEQUENTIAL BUDGET ALLOCATION

There are in total N budget to be allocated. The budget is allocated one at a time. Each point observation consists of two parts: the allocation and the observed result. The allocation is specified by the point selected (e.g. (i, x)) while observation is afterwards generated by random sampling ω and calculating $\tilde{G}_i(x, \omega)$. A complete observation is formulated as (i, x, \tilde{G}) . We label each observation using $j = 1, 2, \dots, N$.

The stage is indexed by $t = 0, 1, \dots, T$ where $t = 0$ represents the initial stage. The purpose of initialization is to guarantee that each i has enough sample points to establish estimation. On the initial stage, a total number of N_0 samples are distributed to all i for some scenarios x . Then, the remaining budget (e.g. $T := N - N_0$) is allocated sequentially. On each stage, an additional budget is allocated and a new point observed. To specify the stage, if the stage information is to be indicated, we label them with $(\cdot)^{(t)}$ (e.g. the observation at stage t would be $(i, x, \tilde{G})^{(t)}$).

On each stage, the historical observations are denoted $F^{(t)} = \{F^{(0)}, (i, x, \tilde{G})^{(1)}, (i, x, \tilde{G})^{(2)}, \dots, (i, x, \tilde{G})^{(t)}\}$ where $F^{(0)}$ is the observations on the initial stage. The σ -field defined on the set including all possible $F^{(t)}$ at stage t is denoted $\mathcal{F}^{(t)}$. The set of scenario and decision of observed points on stage t is denoted $\Gamma^{(t)}$ (e.g. $\Gamma^{(t)} = \{(i, x)_1^{(0)}, \dots, (i, x)_{N_0}^{(0)}, (i, x)^{(1)}, (i, x)^{(2)}, \dots, (i, x)^{(t)}\}$). The complete allocation rule consists of T stages of sequential budget allocation. Conditionally, since on stage t , only $F^{(t)}$ is visible, we need to devise a rule determining $(i, x)^{(t+1)}$ that is measurable $\mathcal{F}^{(t)}$.

To determine a $\mathcal{F}^{(t)}$ measurable $(i, x)^{(t+1)}$ on each stage, we use the method of optimizing acquisition function. An acquisition function is a measure of impact brought by a new allocation $(i, x)^{(t+1)}$ conditioned on the current observations $F^{(t)}$. To specify the partial information within $F^{(t)}$ associated with each i , we denote it as $F_i^{(t)}$. To emphasize the conditional information on $F^{(t)}$, the acquisition function is denoted as $a(i, x|F^{(t)})$. On stage t , the allocation for the next stage is determined by the following optimization:

$$(i, x)^{(t+1)} = \arg \min a(i, x|F^{(t)}) \tag{2}$$

Directly composing a mapping from a data set to functions is difficult. A way of simplifying the problem is to construct estimation on $g_i(x)$ under different i independently. The benefit of this is that we could utilize the continuity feature of $g_i(x)$ (suggested in Assumption 2) to estimate the performance at an unobserved (i, x) . We shall use surrogate method to estimate $g_i(x)$ and then use the surrogate to synthesize the acquisition function. In the following sections, we shall first illustrate our general idea of using large deviation theorem to compose the high level rule. Then, due to the challenge of single observation, we introduce shrinking ball method to fill the technical gap.

3.1 Large Deviation Estimation

Our goal of budget allocation is to minimize TDL in the long run. The main challenge is that X is continuous. In this case, it is impossible for any of x to receive observations on each i . This means that the estimation on $g_i(x)$ at a given x have to be inferred from its neighborhood. The exact method of including the neighborhood information shall be presented later. We set aside this challenge first and focus on the situation where each i, x can be observed multiple times.

Without loss of generality, we temporarily fix x and omit x from discussion. This would yield a conventional Ranking and Selection (R&S) problem. Suppose for each of i , \tilde{G}_i are extensively simulated (e.g. each i has many sample points). We use Sample Average Approximation (SAA) to estimate $g_i(x)$ (e.g. $\hat{g}_i = \frac{1}{N_i} \sum_{j=1}^{N_i} \tilde{G}_j$). From previous studies, it has been proved that, the Probability of False Selection (PFS) (e.g. $PFS = P\{i^* \neq \hat{i}^*\}$), which is essentially TDL conditioned on a fixed x is bounded:

$$PFS \leq \exp[N \max_{i \neq i^*} (-H(\alpha_{i^*}, \alpha_i))] = \max_{i \neq i^*} \exp[-NH(\alpha_{i^*}, \alpha_i)] \tag{3}$$

N indicates the total budget allocated. α_i is the ratio of budget allocated to i (e.g. $\alpha_i = \frac{N_i}{N}$). The function $H(\alpha_{i^*}, \alpha_i)$ intuitively indicates the exponential rate of probability of falsely identifying a non-optimal i as the optimal. $H(\alpha_{i^*}, \alpha_i)$ are decided solely by the budget allocation ratio among different i . In addition, it has the following asymptotic property led by Large Deviation Theorem:

$$-\lim_{N \rightarrow \infty} \frac{1}{N} \ln PFS = \min_{i \neq i^*} H(\alpha_{i^*}, \alpha_i), \tag{4}$$

provided the limit on the left exists. If $\hat{g}_i(x)$ weakly converges to a normal distribution, the left size of (4) exists and H would have an analytical form (Glynn and Juneja 2004):

$$H(\alpha_{i^*}, \alpha_i) = \frac{(g_{i^*} - g_i)^2}{2(\sigma_{i^*}^2/\alpha_{i^*} + \sigma_i^2/\alpha_i)}. \tag{5}$$

Based on these results, we could regard the right side of (3) as an upper bound of the probability of PFS, and a better approximation to the left side as N keeps increasing. Plug (5) into (3), we use the following term to approximate PFS:

$$PFS \approx \max_{i \neq i^*} \exp[-H(N_{i^*}, N_i)] \quad (6)$$

Now we include x back. We use (6) to approximate PFS and plug it into the original definition (1). This yields the following approximation:

$$TDL \approx \int f_X(x) \max_{i \neq i^*} \exp[-H(N_{i^*}, N_i)] dx,$$

Since our goal is to minimize TDL, the acquisition function should be the approximated TDL after allocating an additional budget to (i, x) . However, the calculation on integral can be challenging. Moreover, since only one observation is allowed, we should target on scenario that currently has the largest PFS. Based on these, the extra one observation should be able to minimize the maximum PFS across all scenarios after the observation:

$$a'(i, x) = \begin{cases} \max \left\{ f_X(x) \max_{j \neq i} \exp[-H(N_i + 1, N_j)], \right. \\ \left. \max_{y \neq x} f_Y(y) \max_{k \neq i^*(y)} \exp[-H(N_{i^*(y)}, N_k)] \right\} & \text{if } i = i^*(x) \\ \max \left\{ f_X(x) \max \{ \max_{j \neq i, i^*(x)} \exp[-H(N_{i^*}, N_j)], \exp[-H(N_{i^*}, N_i + 1)] \} \right. \\ \left. \max_{y \neq x} f_Y(y) \max_{k \neq i^*(y)} \exp[-H(N_{i^*(y)}, N_k)] \right\} & \text{if } i \neq i^*(x) \end{cases} \quad (7)$$

Equation (7) is an idealistic acquisition function where the real parameter is known and each (i, x) can be observed multiple times. However, the exact $g_i(x)$, $\sigma_i(x)$ and $i^*(x)$ are unknown, thus need to be estimated. Moreover, it is impossible to have observations of on every (i, x) . We need to design a construct such that a local sample can be defined for an arbitrary (i, x) to infer $g_i(x)$ and $\sigma_i(x)$.

3.2 Shrinking Ball as Surrogate

The shrinking ball/box idea is adopted to construct a local sample. At an arbitrary (i, x) , it defines the local sample as the observations within a specified radius. At stage t , we use $B_i^{(t)}(x)$ to denote the ball of neighborhood within radius $r_i^{(t)}(x)$ for each i :

$$B_i^{(t)}(x) = \{(l, y) | l = i, \|y - x\|_p \leq r_i^{(t)}(x)\}$$

The observed points within the region of $B_i^{(t)}(x)$ is denoted $b_i^{(t)}(x)$:

$$b_i^{(t)}(x) = \{(l, y)_j | (l, y)_j \in B_i^{(t)}(x)\}$$

Based on this neighborhood construct, the estimation of $g_i(x)$ conditioned on observations on stage t is:

$$\hat{g}_i(x | F_i^{(t)}) = \frac{1}{|b_i^{(t)}(x)|} \sum_{j: (l, y)_j \in B_i^{(t)}(x)} \tilde{G}_j$$

And the estimation on $\sigma_i(x)$ is:

$$\hat{\sigma}_i^2(x | F_i^{(t)}) = \frac{1}{|b_i^{(t)}(x)| - 1} \sum_{j: (l, y)_j \in B_i^{(t)}(x)} [\tilde{G}_j - \hat{g}_i(x | F_i^{(t)})]^2$$

Therefore, by (5) the estimated rate function is:

$$\hat{H}(N_{i^*}, N_i) = \frac{(\hat{g}_{i^*} - \hat{g}_i)^2}{2(\hat{\sigma}_{i^*}^2/N_{i^*} + \hat{\sigma}_i^2/N_i)}.$$

There are two sources of error in these estimators. The first one is the bias caused by sampling at adjacent points. This type of error decreases if the ball keeps shrinking. The second one is the sample variance induced by ε . This type of error decreases if the sample size keeps increasing. To achieve a consistent estimation on $g_i(x)$, one needs to control the size of the ball so that it keeps shrinking while maintaining increase of the sample size within the ball.

We now specify the rule that governs the radius $r_i^{(t)}(x)$. The basic idea is to make sure that, neither do $r_i^{(t)}(x)$ shrink too slow to efficiently decrease the “uncertainty” (i.e. the bias led by observations in the neighborhood), nor too fast to let “randomness” decrease (caused by decrease of sample size within ball). To make sure that $r_i^{(t)}(x)$ decreases, we let it controlled by the following rule:

$$r_i^{(t)}(x) = \min\{r_i^{(t-1)}(x), \min\{r : |b_i^{(t)}(x)| \geq \phi_i^{(t)}(x), \forall x\}\}$$

By this rule, on each stage, $r_i^{(t)}(x)$ is the smallest radius that allows the ball to contain at least $\phi_i^{(t)}(x)$ number of samples, while guarantees decreasing (e.g. $r_i^{(t)}(x) \leq r_i^{(t-1)}(x)$). To make sure that the size of the local sample increases, $\phi_i^{(t)}(x)$ is applied to regulate the shrinking speed. As a rule of thumb, we let $\phi_i^{(t)}(x)$ be a counting of successful shrinking for each x , which is:

$$\phi_i^{(t)}(x) = 2 + |\{k < t | r_i^{(k-1)}(x) < r_i^{(k)}(x), k \in \mathbb{N}^+\}|$$

The additional 2 in $\phi_i^{(t)}(x)$ is to guarantee that each ball has at least two sample points to calculate $\hat{\sigma}_i(x)$ when algorithm initialized.

3.3 Modification of Uncertainty

In addition to the shrinking ball surrogate and the large deviation estimation, we need to include a factor that describes the “uncertainty” level to encourage sample on less observed scenarios. For example, if two scenarios have the same estimation on $\hat{g}_i(x)$ and $\hat{\sigma}_i(x)$, the one which is more distant to the existing observations (e.g. larger ball radius) should be sampled more. This is because that if an observation is distant from the current sample, it means the inference made is poor. In this regard, we introduce function $\psi_i^{(t)}(x)$ to reflect this information. In this study, we let the uncertainty modifier be:

$$\psi_i^{(t)}(x) = \frac{1}{|r_i^{(t)}(x)|^2}. \tag{8}$$

3.4 Acquisition Function and Optimization

Since “uncertainty” always exists no matter how many observations we acquire, the uncertainty modifier is applied to the estimated rate function. Based on the large deviation rate estimation (7), the shrinking ball construct and the uncertainty modifier (8), we establish the final acquisition function:

$$a(i, x|F^{(t)}) = \begin{cases} \max \left\{ f_X(x) \max_{j \neq i} \exp[-\psi_i^{(t)}(x) \hat{H}(|b_i^{(t)}(x)| + 1, |b_j^{(t)}(x)|)], \right. \\ \left. \max_{y \neq x} f_Y(y) \max_{k \neq i^*(y)} \exp[-\psi_i^{(t)}(y) \hat{H}(|b_{i^*(y)}^{(t)}(y)|, |b_j^{(t)}(y)|)] \right\} & \text{if } i = \hat{i}^* \\ \max \left\{ f_X(x) \max \left\{ \max_{j \neq i, \hat{i}^*} \exp[-\psi_j^{(t)}(x) \hat{H}(|b_{\hat{i}^*}^{(t)}(x)|, |b_j^{(t)}(x)|)], \right. \right. \\ \left. \left. \exp[-\psi_i^{(t)}(x) \hat{H}(|b_{\hat{i}^*}^{(t)}(x)|, |b_i^{(t)}(x)| + 1)] \right\}, \right. \\ \left. \max_{y \neq x} f_Y(y) \max_{k \neq i^*(y)} \exp[-\psi_i^{(t)}(y) \hat{H}(|b_{i^*(y)}^{(t)}(y)|, |b_j^{(t)}(y)|)] \right\} & \text{if } i \neq \hat{i}^* \end{cases} \quad (9)$$

Although (9) has an explicit form, finding the optimal solution on each stage is still a challenging job. For a fixed x , the optimization is simple. However, since X is continuous and $f_X(x)$ is a part of the formulation, the optimization could be hard to solve. To simplify the problem, we adopt a random sampling rule. This means on each stage, we uniformly sample a fixed number Q of x s to form a discretized X (i.e. $\hat{X}^{(t)} = \{x_1^{(t)}, x_2^{(t)}, \dots, x_Q^{(t)}\}$) and optimize each of the sub-problem conditioned on each x . Then we find the best (i, x) among the solutions of the sub-problems to get an estimation on the optimal solution. This means we use the following rule as an approximation to (2):

$$(i, x)^{(t+1)} = \arg \min_{x \in \hat{X}^{(t)}} a(i, x|F^{(t)})$$

To summarize, the proposed algorithm ‘‘Hybrid of Shrinking ball method and optimal Large deviation rate estimation’’ (HSL) works as follows. In the initial stage, N_0 number of scenarios are independently sampled from the scenario set. For each sampled scenario, we simulate, on each decision, the corresponding performance once. On each stage afterward, we add one simulation observation on the point (one scenario and decision combination) that has the minimal acquisition function value, which anticipates the improvement of TDL convergence rate if the point is sampled, based on the previous simulated performance within the shrinking ball. The pseudocode of HSL is presented in Algorithm 1:

Algorithm 1: HSL

Input: $f_X(x)$, I , X , T , Q , N_0

Output: $\hat{i}^*(x)$

Initialization: $t = 0$, generate initial $\Gamma^{(0)} = \{\{i, x_1\}_{i \in I}, \{i, x_2\}_{i \in I}, \dots, \{i, x_{N_0}\}_{i \in I}\}$ from Θ ; Observe \tilde{G} at each x for all i and initialize the data set

$F^{(0)} = \left\{ \{(i, x_1, \tilde{G}_i(x_1))\}_{i \in I}, \{(i, x_2, \tilde{G}_i(x_2))\}_{i \in I}, \dots, \{(i, x_{N_0}, \tilde{G}_i(x_{N_0}))\}_{i \in I} \right\}$; initialize $b_i^{(t)}(x)$,

$\phi_i^{(t)}(x)$, $r_i^{(t)}(x)$, $\psi_i^{(t)}(x)$, $\hat{g}_i^{(t)}(x)$ and $\hat{\sigma}_i^{(t)}(x)$;

while $t < T$ **do**

Create $\hat{X}^{(t)}$;

$(i, x)^{(t+1)} = \arg \min_{x \in \hat{X}^{(t)}} a(i, x|F^{(t)})$;

Observe \tilde{G} at $(i, x)^{(t+1)}$;

Update $\Gamma^{(t+1)} = \Gamma^{(t)} \cup (i, x)^{(t+1)}$;

Update $F_{t+1} = F_t \cup \{(i, x, \tilde{G}_i(x))^{(t+1)}\}$;

Update $b_i^{(t)}(x)$, $\phi_i^{(t)}(x)$, $r_i^{(t)}(x)$, $\psi_i^{(t)}(x)$, $\hat{g}_i^{(t)}(x)$ and $\hat{\sigma}_i^{(t)}(x)$;

$t++$;

end

Calculate $\hat{i}^*(x)$ based on $\hat{g}_i^{(T)}(x)$;

4 NUMERICAL EXPERIMENTS

The goal of numerical experiments is to test whether HSL can let TDL efficiently decreases. To highlight the problem features, we use the following benchmark problem:

$$g_i(x) = \begin{cases} \frac{10}{x^2} \sin e^x & i = 1 \\ 0 & i = 2 \end{cases}, \quad x \in [1, 3]$$

The graph of $g_i(x)$ is shown in Figure. 1. The reason for choosing this function is that it has three features that challenge the correct identification. Consider the function from the perspective of $g_1(x)$, three features are:

- $\sin(\cdot)$: The sin function oscillates, creating frequent switch of optimal decision;
- e^θ : The exponential function creates an imbalanced frequency of switching of optimal decisions;
- $\frac{1}{\theta^2}$: It creates different amplitude, leading to the imbalanced difficulty level of correct identifying the optimal decision

If HSL gains an efficiency advantage even at the presence of these features, it implies that our algorithm works.

The $\varepsilon_i(x)$ is set to be normally distributed with the standard deviation set to be monotone increasing. The monotone increasing of the variance indicates that the right side of the problem is more difficult to identify the optimal decision. The Q is set to be 100 to obtain full coverage of the whole space. $f_X(x)$ is set to be constant throughout the space.

The benchmark algorithm applied is the uniform sampling where the allocation of (i, x) is decided by uniform sampling from $\hat{X}^{(t)}$. In this benchmark algorithm, we use the same shrinking policy and shrinking ball construct to do surrogate so that the efficiency of budget allocation can be compared.

In the first case, we set the variance of ε to be large:

$$\sigma_i(x) = x, \quad \forall i.$$

Figure. 2 shows a comparison between HSL and uniform sampling after 500 observations (N_0 is set to be 30). We see that HSL adaptively allocates more budget at points where g_1 and g_2 switch ranks. This is what we are expecting as the goal of contextual simulation optimization is not to achieve a good estimation throughout the space but to identify the conditional optimal decisions effectively.

In the second case, we set the variance of ε to be smaller:

$$\sigma_i(x) = \frac{1}{2}x, \quad \forall i.$$

Figure. 3 shows a comparison of correct identification as observation increases. After 120 observations, HSL correctly identifies the optimal decision almost everywhere. Moreover, since the right side of the scenario space is more difficult to differentiate, HSL can adaptively allocate more budget to the right side of the space. Then, we plot the TDL against the total observations to see the expected performance of the HSL. From the plot (Figure. 4), HSL gains a clear advantage to uniform sampling. To achieve 0.1 TDL, HSL only needs approximately 170 observations while uniform sampling costs almost 400 observations. For a simulation that costs several hours to run a single observation, the cost saved is tremendous.

5 CONCLUSION

In conclusion, we propose a hybrid method of the shrinking ball and estimation on the large deviation rate that intends to achieve the optimal convergence rate in the new problem setting. Numerical experiments show a significant efficiency boost with respect to the decreasing of Total Decision Loss. It paves the way for the theoretical study of the optimal convergence rate in continuous contextual simulation optimization.

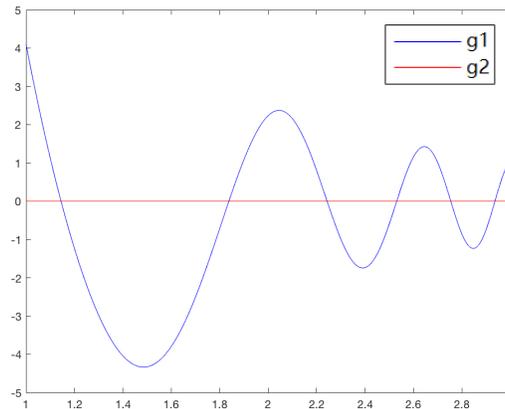


Figure 1: The graph of the benchmark problem.

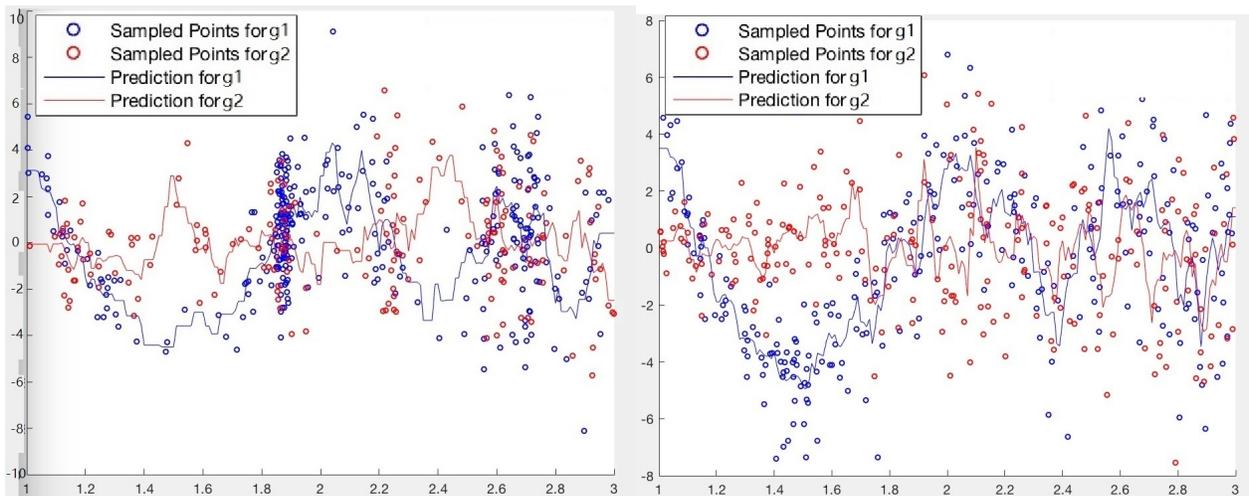


Figure 2: A comparison of surrogate after 500 observations between HSL (Left) and uniform sampling (right). HSL attains TDL of 0.09 while the uniform sampling attains only 0.29.

REFERENCES

- Gao, S., J. Du, and C.-H. Chen. 2019. “Selecting the Optimal System Design Under Covariates”. In *2019 IEEE 15th International Conference on Automation Science and Engineering (CASE)*, 547–552. Institute of Electrical and Electronics Engineers, Inc.
- Gao, S., C. Li, and J. Du. 2019. “Rate Analysis for Offline Simulation Online Application”. In *Proceedings of the 2019 Winter Simulation Conference*, edited by S. L.-M. M. R. C. S. P. H. N. Mustafee, K.-H.G. Bae and Y.-J. Son, 3468–3479. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.
- Glynn, P., and S. Juneja. 2004. “A Large Deviations Perspective on Ordinal Optimization”. In *Proceedings of the 2004 Winter Simulation Conference*, edited by J. S. S. R. G. Ingalls, M. D. Rossetti and B. A. Peters, Volume 1, 585. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.
- Jin, X., H. Li, and L. H. Lee. 2019. “Optimal Budget Allocation in Simulation Analytics”. In *2019 IEEE 15th International Conference on Automation Science and Engineering (CASE)*, 178–182. Institute of Electrical and Electronics Engineers, Inc.

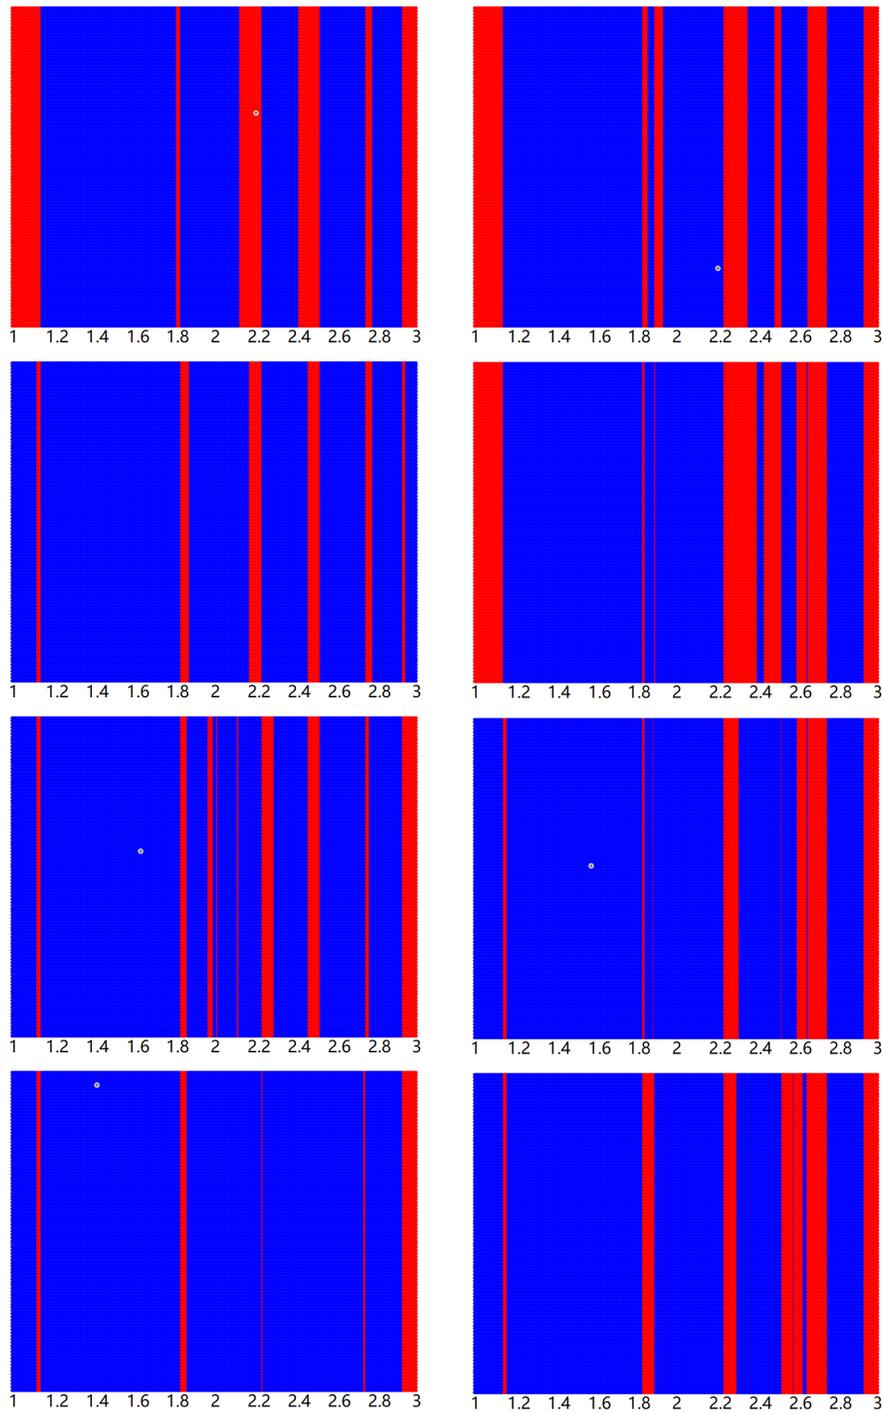


Figure 3: The comparison of region where false identifying the optimal decisions occurs (red) as observations increases. From up to down, it shows the quality of identification after 30,60,90,120 observations. The HSL (left) identifies the correct decision effectively.

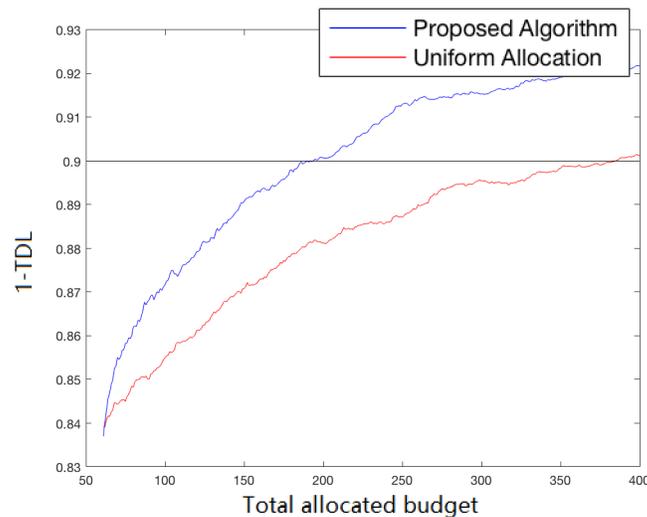


Figure 4: A comparison of (1-TDL) between HSL (blue) and uniform sampling (red) as observation increases. The plot is the averaged results from 300 independent trails. To attain 0.1 TDL (0.9 on the graph), HSL spends less than half of the budget needed for uniform sampling.

- Kiatsupaibul, S., R. L. Smith, and Z. B. Zabinsky. 2018. “Single Observation Adaptive Search for Continuous Simulation Optimization”. *Operations Research* 66(6):1713–1727.
- Linz, D. D., Z. B. Zabinsky, S. Kiatsupaibul, and R. L. Smith. 2017. “A Computational Comparison of Simulation Optimization Methods Using Single Observations within a Shrinking Ball on Noisy Black-box Functions with Mixed Integer and Continuous Domains”. In *Proceedings of the 2017 Winter Simulation Conference*, edited by G. Z. N. M. G. W. W. K. V. Chan, A. D’Ambrogio and E. Page, 2045–2056. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.
- Nelson, B. L. 2016. “Some Tactical Problems in Digital Simulation’for the Next 10 Years”. *Journal of Simulation* 10(1):2–11.
- Shen, H., L. J. Hong, and X. Zhang. 2017. “Ranking and Selection with Covariates”. In *Proceedings of the 2017 Winter Simulation Conference*, edited by G. Z. N. M. G. W. W. K. V. Chan, A. D’Ambrogio and E. Page, 2137–2148. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.
- Zhou, Li 2015. “A Survey on Contextual Multi-armed Bandits”. <https://arxiv.org/pdf/1508.03326.pdf>, accessed 2015.08.03.

ACKNOWLEDGMENTS

Mr. Xiao Jin have been supported by Singapore A*STAR IAF-PP fund (Grant No. A1895a0033) under the project “Digital Twin for Next Generation Warehouse”. Dr. Yichi Shen is supported by Institute of Operations Research and Analytics.

AUTHOR BIOGRAPHIES

XIAO JIN is a Ph.D. candidate at Industrial Systems Engineering and Management, National University of Singapore. His research interests are simulation optimization, simulation analytics and parallel computing. His email address is isejinx@nus.edu.sg.

YICHI SHEN is currently a postdoctoral research fellow in the Institute of Operations Research and Analytics, National University of Singapore. He obtained his Ph.D. in Industrial and System Management from Nanyang Technological University, Singapore and holds the M.Sc. and B.Sc. in Mathematics from Nanjing University, China.

His current research focuses on developing efficient and cost-effective algorithms for the computational optimization of stochastic expensive black-box functions and he is passionate about applying the research to solving real life problems. His research interests also cover Operations Research, Data Analysis, Stochastic Models and Smart Manufacturing. His e-mail address is orasy@nus.edu.sg.

LOO HAY LEE is an Associate Professor in the Department of Industrial Systems Engineering and Management at the National University of Singapore. His research focuses on the simulation-based optimization, maritime logistics which includes port operations and the modeling and analysis for the logistics and supply chain system. He has co-authored a book: *Stochastic Simulation Optimization-Optimal Computing Budget Allocation* (World Scientific) and co-edited the book *Advances in Maritime Logistics and Supply Chain Systems*. He has won the Best Paper Award at the INFOCOMP 2011 and co-lead a team with A/Prof Chew Ek Peng to win the grand prize for the next generation container port challenge in 2013 by proposing a revolutionary double storey container terminal design, called SINGA Port. He has also served as the assistant dean in charge of research for the Faculty of Engineering of NUS from 2001 to 2003, and is now the deputy head of research for the department. His email address is iselelh@nus.edu.sg.

EK PENG CHEW is an Associate Professor in the Department of Industrial Systems Engineering and Management at the National University of Singapore. He also holds positions as the Deputy Head (Undergraduate Studies), Director for the Centre of Excellence in Modeling and Simulation for Next Generation Ports, and Co-Director for the Centre for Next Generation Logistics. He was a Visiting Scholar and a Visiting Professor, respectively, at the Georgia Institute of Technology and University of British Columbia in 2006. His current research areas are in port logistics and maritime transportation, simulation optimization and inventory management. He is serving as Editor-in-Chief of the *Asia Pacific Journal of Operational Research*, member of the Editorial Board Editors of the *Transportation Research Part B*, member in the Editorial of *Flexible Services and Manufacturing Journal*. He has also co-edited a book in “*Advances in Maritime Logistics and Supply Chain Systems*”, a special issue for *OR Spectrum* on “IT-based planning and control of seaport container terminals and freight transportation systems” and three special issues for *Flexible Services and Manufacturing Journal* on “Maritime Container Logistics and Onshore Transportation Systems (Part 1, Part 2 and Part 3)”. He and Dr. Lee Loo Hay led a team of multidisciplinary researchers and experience practitioners to win the Next Generation Container Port Challenge with a grand prize of US\$1 mil by proposing a revolutionary double-storey container terminal, called the SINGA port. His email address is isecep@nus.edu.sg

CHRISTINE A. SHOEMAKER is Distinguished Professor in the Industrial Systems Engineering and Management Department and Civil and Environmental Engineering Department at the National University of Singapore. She is Joseph P. Ripley Professor Emerita at Cornell University in USA. Prof Shoemaker is Member of the USA National Academy of Engineering, and she is a (professional society elected) Fellow in SIAM (applied mathematics), a Fellow in INFORMS (operations research), a Fellow in AGU (geophysics including hydrology) and a Fellow and Distinguished Member of ASCE (Civil and Environmental Engineering). She has also won a number of other awards. Prof. Shoemaker received her PhD in mathematics (on dynamic programming applied to environmental problems) supervised by Richard Bellman. Her current research has two main components which are: 1) development of computationally efficient optimization algorithms that can be used to solve problems where the objective function is multimodal (i.e. has multiple local minima) and is expensive to compute (so relatively few objective functions can be made) and no derivatives are available. In this research, a radial basis function surrogate is used and both serial and parallel algorithms have been developed for continuous and integer variables and for single and multiobjective problems. 2) application of optimization methods to difficult real problems described by nonlinear simulation models (deterministic or stochastic, PDE or discrete time dynamic models) that describe environmental problems, especially in the area of water quality or water resources (including lakes, watersheds, and groundwater). The methods developed in 1) above are applied to these problems and more recently to hyperparameter optimization in deep learning. Her email address is shoemaker@nus.edu.sg.