

SIMULATION OPTIMIZATION WITH NON-STATIONARY STREAMING INPUT DATA

Songhao Wang¹, Haowei Wang², Jianglin Xia¹, and Xiuqin Xu³

¹College of Business, Southern University of Science and Technology, Shenzhen, Guangdong, CHINA

²Rice-Rick Digitalization PTE. Ltd., SINGAPORE

³McKinsey & Company, SINGAPORE

ABSTRACT

Simulation optimization has become an emerging tool to design and analysis of real-world systems. In stochastic simulation, input distribution is a main driving force to account for system randomness. Most existing works on input modeling focus on stationary input distributions. In reality, however, input distributions could experience sudden disruptive changes due to external factors. In this work, we consider input modeling through non-stationary streaming input data, where the input data arrive sequentially across different decision stages. Both the parameters of the input distributions and the disruptive change points are unknown. We use a Markov Switching Model to estimate the non-stationary input distributions, and design a metamodel-based approach to solve the following optimization problem. The proposed metamodel and optimization algorithm can utilize the simulation results from all the past stages. A numerical study on an inventory system shows that our algorithm can solve the problem more efficiently compared to common approaches.

1 INTRODUCTION

The stochastic simulation model is a powerful tool for the analysis of complex systems, and simulation optimization refers to the procedures to find the best system designs through simulation models. Running a simulation model typically requires an input distribution as a driving force to account for the randomness in the system inputs. For instance, in an inventory model, the decision-makers search for the optimal policy to maximize the profits. Here, the distributions for the customer demands and the lead times are relevant input distributions. The simulation model generates sample demands and lead times to mimic the operations of the inventory systems.

However, the underlying true input distributions are difficult to know in reality. In practice, they can be estimated from a set of real-world input data (the random realizations of the unknown input distributions). Most of the works in simulation input modeling assume that only a fixed set of input data is provided before the decision-making process. Recently, there has been a growing interest in simulation problems with streaming input data, where additional input data become available during different decision-making stages (Liu et al. 2024; Wu et al. 2024; Song and Shanbhag 2019). In this setting, it would be ideal if the input distribution, as well as the simulation model, could be adapted to the new data. Such streaming input data can help facilitate the construction of more accurate and up-to-date simulation models (Wu et al. 2024). The online simulation optimization with streaming data is applicable when acquiring new data is efficient. For example, in an inventory problem, daily demand and lead time data can be easily collected during the sales season.

The common assumption on the streaming input data is that the input data are independent and identically distributed (i.i.d.). It can be expected that with additional data, the input modeling estimation will be gradually improved. In this context, Liu et al. (2024) and Wu et al. (2024) adopt Bayesian framework to update the input model parameters while Song and Shanbhag (2019) estimate the parameters through maximum likelihood estimator (MLE). Wu et al. (2024) additionally consider a setting when the

input distributions are endogenous (decision-dependent). Hence, the upcoming input data depend on the current decision and are no longer identically distributed.

In this work, we focus on a different scenario in which the streaming input data are non-stationary, i.e., the input distribution can change at certain time points due to external factors. Non-stationary input models arise when the systems experience exogenous disruptive events. For example, the natural disasters (Sydnor et al. 2017; Carvalho et al. 2021) and Covid-19 pandemic (Ivanov and Dolgui 2021) could bring disruptive changes, such as the change of distribution of customer demands, to the supply chain systems. As systems become increasingly complex, such as in the case of extensive international supply chains with numerous diverse and interconnected subsystems, they become susceptible to instability and encounter subtle disruptions that pose challenges for decision-makers to anticipate and address. It is of great importance to study systems' performance under these changes to prevent them from large losses or even breakdowns. In the work, we assume that input data are no longer stationary and there exist different regimes for the data. The distributions of input data from different regimes belong to the same parametric family but with different parameters. Both the distribution parameters and the change points are unknown and need to be inferred from input data. The non-stationary streaming input data could be treated as a set of time series data that exhibit regime switching over time. Figure 1 illustrates the regime-switching dynamics. Two switchings occur at time points t_1 and t_2 , respectively, among the three regimes. The data points within every regime are generated from a stationary normal distribution whose mean is indicated as the straight line and variance is 1. Once the regimes of any two data points are given, the observations become conditionally independent.




Figure 1: Illustration of Regime Switching.

We adopt the Markov Switching Model (MSM) to model the regime-switching dynamics. It builds a hidden Markov chain to model the regimes of each data point. Different regimes can switch to each other according to a transition matrix. MSM has been a popular model for non-stationary time series since the seminal work of Hamilton (1989). It has been increasingly applied in areas including biology (Albert 1991), engineering (Thyer and Kuczera 2000; Zucchini and Guttorp 1991), finance and marketing (Pagan and Schwert 1990; Ang and Bekaert 2002; Hamilton and Raj 2002). We refer to Bhar and Hamori (2004) and Frühwirth-Schnatter (2006) for detailed introductions to MSM. At the beginning of each time stage, a new set of input data arrives, which is then incorporated into the existing data stream from the previous stage. We then apply MSM to estimate the distribution parameters and the transition matrix, based on which, we could derive the predictive distribution of the input data in the next time stage. It can be shown that

the predictive distribution is a mixture of several stationary input distributions, corresponding to different regimes.

Given this predictive input distribution from the MSM, we aim to find the optimal solution based on simulation models for the forthcoming time stage. Essentially, we are addressing an online simulation optimization problem with streaming data to deliver optimal decisions across multiple time stages. There are two major goals we aim to achieve. The first is to enhance input data modeling using streaming non-stationary data, a goal accomplished through the utilization of the MSM. Another goal is to reuse the simulation results from previous time stages, especially when the simulation experiments are expensive and the simulation budget for each stage is limited. However, the input distributions are non-stationary and the model estimations can be different across different stages. The simulation outputs obtained under heterogeneous input distributions, although valuable, cannot be straightforwardly aggregated. Wu et al. (2024) design a moving average aggregation approach that discards the previous simulation outputs gradually to maintain an unbiased estimator for ranking and selection problems. In this work, we propose a Gaussian process (GP) metamodel-based approach, a well-accepted approach for expensive simulation optimization problems (Meng et al. 2022; Shahriari et al. 2015), to fully reuse the previous simulation outputs. We construct a GP model for the simulation outputs with respect to both the decision parameters and the input distribution parameters. The model is expected to learn the relationship based on all the simulation outputs from previous stages. To guide the search for the optimum of the upcoming stage, we further design an aggregated Expected Improvement (EI) searching criterion. Similar approaches have been used in Wang et al. (2020) but for a different setting: they work on optimization with input uncertainty based on fixed-size input data from a stationary distribution. A numerical experiment on an inventory problem shows that our approaches provide better estimation for the optimal solutions across different time stages compared to commonly used approaches which do not consider the non-stationarity of the input distribution.

The remainder of this paper is organized as follows. Section 2 introduces the MSM model, the basics of the GP model and GP-based optimization algorithm. Section 3 provides the problem formulation considered in this work. Section 4 presents the GP model and the optimization algorithm. Section 5 shows a numerical example and Section 6 concludes the paper.

2 BACKGROUND AND BASICS

In this section, we will review the background and basics of MSM, GP model and GP-based optimization algorithms.

2.1 Markov Switching Model with Two Regimes

The two key issues in using MSM is the parameter estimation and the prediction.

2.1.1 Model Parameters Estimation

We take a frequentist perspective to estimate the parameters through MLE. Denote the parameters collectively as $\vartheta = (\lambda_1, \lambda_2, p, q)$. After stage t , we observe ξ^t . We further denote S^t as the collection of all the hidden indicators $\{S_1, S_2, \dots, S_t\}$. The MLE $P(\xi^t | \vartheta)$ can be derived as:

$$P(\xi^t | \vartheta) = \sum_{S^t \in \mathcal{S}} P(\xi^t, S^t | \vartheta) = \sum_{S^t \in \mathcal{S}} P(\xi^t | S^t, \vartheta) P(S^t | \vartheta), \quad (1)$$

where $\mathcal{S}$ is the set of all possible sequences S^t . It basically consists of every sequence of length t where the first position is 1 and all other positions are 1 or 2. In (1), $P(\xi^t | S^t, \vartheta)$ can be easily computed as:

$$P(\xi^t | S^t, \vartheta) = P(\xi_1 | \lambda_1) \prod_{i=2}^t P(\xi_i | S_i, \xi^{i-1}, \vartheta) = \tilde{P}(\xi_1 | \lambda_1) \prod_{i=2}^t \tilde{P}(\xi_i | \lambda_{S_i}).$$

The second equality holds as ξ^t becomes independent given the indicators S^t . The distribution $\tilde{P}(\xi|\lambda_{S_i})$ is often called the emission distribution in MSM. In (1), $P(S^t|\vartheta)$ can be computed as follows:

$$P(S^t|\vartheta) = P(S^t|A) = \prod_{i=2}^t P(S_i|S_{i-1}, A) = \prod_{i=2}^t A_{S_{i-1}, S_i} = \prod_{j=1}^2 \prod_{k=1}^2 A_{j,k}^{N_{jk}(S^t)},$$

where $A_{j,k}$ is the j, k -th entry of the transition matrix A and $N_{jk}(S^t)$ counts the numbers of transitions from j to k in S^t :

$$N_{jk}(S^t) = \#\{S_{i-1} = j, S_i = k\}, \quad \forall j, k \in \{1, 2\}.$$

2.1.2 Prediction of $P(\xi_{t+1}|\xi^t)$ for a Two-Regime MSM

Given the model parameters, the predictive distribution $P(\xi_{t+1}|\xi^t)$ can be derived as:

$$\begin{aligned} P(\xi_{t+1}|\xi^t, \vartheta) &= \sum_{l=1}^2 P(\xi_{t+1}|\xi^t, S_{t+1} = l, \vartheta) P(S_{t+1} = l|\xi^t, \vartheta) \\ &= \sum_{l=1}^2 P(\xi_{t+1}|S_{t+1} = l, \vartheta) P(S_{t+1} = l|\xi^t, \vartheta) \\ &= \sum_{l=1}^2 \tilde{P}(\xi_{t+1}|\lambda_l) \left[\sum_{k=1}^2 P(S_{t+1} = l|S_t = k, \xi^t, \vartheta) P(S_t = k|\xi^t, \vartheta) \right] \\ &= \sum_{l=1}^2 \tilde{P}(\xi_{t+1}|\lambda_l) \left[\sum_{k=1}^2 A_{k,l} P(S_t = k|\xi^t, \vartheta) \right], \end{aligned} \tag{2}$$

where

$$\begin{aligned} P(S_t = k|\xi^t, \vartheta) &= \frac{P(\xi_t|S_t = k, \xi^{t-1}, \vartheta) P(S_t = k|\xi^{t-1}, \vartheta)}{\sum_{j=1}^2 P(\xi_t|S_t = j, \xi^{t-1}, \vartheta) P(S_t = j|\xi^{t-1}, \vartheta)} \\ &= \frac{\tilde{P}(\xi_t|\lambda_k) P(S_t = k|\xi^{t-1}, \vartheta)}{\sum_{j=1}^2 \tilde{P}(\xi_t|\lambda_j) P(S_t = j|\xi^{t-1}, \vartheta)}. \end{aligned}$$

Notice that,

$$P(S_t = k|\xi^{t-1}, \vartheta) = \sum_{i=1}^2 P(S_t = k|S_{t-1} = i, \vartheta) P(S_{t-1} = i|\xi^{t-1}, \vartheta) = \sum_{i=1}^2 A_{i,k} P(S_{t-1} = i|\xi^{t-1}, \vartheta).$$

Therefore, to evaluate $P(\xi_{t+1}|\xi^t, \vartheta)$, we need to compute $P(S_i|\xi^{i-1}, \vartheta)$ and $P(S_i|\xi^i, \vartheta)$ for all $i = 2, \dots, t$. This can be done recursively starting from the following initial condition:

$$P(S_1 = 1|\xi^1, \vartheta) = 1, \quad P(S_1 = 0|\xi^1, \vartheta) = 0.$$

Computing $P(S_t|\xi^t, \vartheta)$ is known as the filtering problem in MSM (see details for general settings in Frühwirth-Schnatter (2006)).

2.2 Basics of Stochastic GP Model and GP-Based Algorithms

In a stochastic GP model, the simulation output is modeled as:

$$Y(\mathbf{x}) = Z(\mathbf{x}) + \varepsilon(\mathbf{x}) = \mathbf{v}(\mathbf{x})^T \boldsymbol{\beta} + \boldsymbol{\eta}(\mathbf{x}) + \varepsilon(\mathbf{x}),$$

where $Z(\mathbf{x})$ is the noise-free value of the simulation output and $\varepsilon(\mathbf{x})$ is the stochastic simulation noise with mean zero and variance $\sigma_\varepsilon^2(\mathbf{x})$. The response $Z(\mathbf{x})$ is decomposed into a mean value $v(\mathbf{x})^T \beta$ and a second-order stationary GP, $\eta(\mathbf{x})$. In the mean function, $v(\mathbf{x})$ is a known regressor, reflecting our knowledge in the trend of the response, and β is the coefficient of the regressor that should be estimated from data. Without loss of generality, we assume $v(\mathbf{x}) = 0$ for zero mean when no additional information is available. The GP $\eta(\mathbf{x})$ has the following covariance structure: $c(\mathbf{x}_1, \mathbf{x}_2) := \text{cov}(\eta(\mathbf{x}_1), \eta(\mathbf{x}_2)) = \sigma_z^2 \text{corr}(\mathbf{x}_1, \mathbf{x}_2)$, where $\sigma_z^2 = \text{var}(\eta(\mathbf{x}))$ is the process variance. We model the spatial correlation with the Gaussian correlation function: $\text{corr}(\mathbf{x}_1, \mathbf{x}_2) = \exp \left\{ \sum_{j=1}^d \theta_j (\mathbf{x}_{1,j} - \mathbf{x}_{2,j})^2 \right\}$, where $\mathbf{x}_{i,j}$ is the j th coordinate of $\mathbf{x}_i$ and θ is the lengthscales parameter. This is a common choice in the GP modeling that provides reasonably good accuracy in many applications. See Williams and Rasmussen (2006) for details in the GP modeling.

At a new input point $\mathbf{x}_0$, the GP model can provide a prediction of $Z(\mathbf{x}_0)$ with a predictive variance to evaluate the uncertainty in the prediction:

$$\hat{Z}(\mathbf{x}_0) = c(\mathbf{x}_0)^T R^{-1} \mathcal{Y},$$

$$\hat{s}^2(\mathbf{x}_0) := \text{var}(\hat{Z}(\mathbf{x}_0)) = \sigma_z^2 - c(\mathbf{x}_0)^T R^{-1} c(\mathbf{x}_0),$$

where $\mathcal{Y} = (\bar{y}(\mathbf{x}_1), \dots, \bar{y}(\mathbf{x}_n))^T$ is the observations vector at all the past design points $\mathbf{x}_1, \dots, \mathbf{x}_n$. Due to the noise, the stochastic GP model is built with more than one observation at any single design point (suppose r replications) with $\bar{y}$ denoting the sample mean: $\bar{y}(\mathbf{x}_i) = \frac{1}{r} \sum_{j=1}^r y(\mathbf{x}_i, \xi_j)$, $i = 1, \dots, n$. $c(\mathbf{x}_0)$ is the $n \times 1$ covariance vector of $\mathbf{x}_0$ with existing design points whose i -th entry is $c(\mathbf{x}_0, \mathbf{x}_i)$, $R = R_z + R_\varepsilon$ is the covariance matrix of the design points, where R_z is the $n \times n$ covariance matrix for the spatial process whose (i, j) -th entry is $c(\mathbf{x}_i, \mathbf{x}_j)$ and R_ε is the $n \times n$ diagonal covariance matrix for the noises whose i -th diagonal entry is $\text{var}(\varepsilon(\mathbf{x}_i))$. When used in practice, R_ε is estimated through sample variance from the simulation results. The parameters θ and σ_z^2 can be obtained by maximum likelihood estimation (Ankenman et al. 2010).

The optimization approach based on the GP model is referred to as Bayesian optimization (BO) in the literature. It adopts the GP prediction to guide the search. As the GP model summarizes the information from all the previous results, this search is expected to find the optimal solutions efficiently and is widely used in scenarios with expensive experiments. The key ingredient for BO is an acquisition function constructed from the GP model. It generally evaluates the utility in running experiments at a design point. Thus, the design point that should be tried in the next iteration is the maximizer of the acquisition function. Commonly used acquisition functions include the expected improvement (EI) function, upper confidence bound (UCB) and knowledge gradient (KG). Readers could refer to Shahriari et al. (2015) for introductions to different types of acquisition functions. In addition to selecting the next sample point, we need to decide the number of simulation replications to run at this point. We take the simple way to run an equal number of replications in this work. More delicate approaches can be used to adaptively improve the model fitting or the optimizer identification. See Jalali et al. (2017) for numerical comparisons among different replication allocation strategies.

3 PROBLEM FORMULATION

Denote the decision variable as $\mathbf{x}$, a stochastic simulation optimization problem can be formulated as:

$$\min_{\mathbf{x} \in \mathcal{X}} g(\mathbf{x}, \lambda_S) := \mathbb{E}_{\xi \sim \tilde{P}(\xi | \lambda_S)} [y(\mathbf{x}, \xi)],$$

where $\mathcal{X}$ is the design space assumed to be compact, ξ represents the simulation input, $\tilde{P}$ is the input distribution parameterized by λ_S , and y is the simulation output. The subscript of λ , S , indicates the regime of the current input distribution. In this work, we assume there are two different regimes: $S \in \{1, 2\}$, i.e., the current input distribution can switch to another one due to external factors. As aforementioned, the

input distribution in each regime is assumed to be from a parametric family. For ease of exposition, we assume that at the beginning of decision stage t , we receive only one new input data point. Nonetheless, our approach can easily extend to situations with more than two possible regimes and a varying number of input data points received in each stage.

We denote all the input data that have been received before the end of stage t as $\xi^t := \{\xi_1, \xi_2, \dots, \xi_t\}$. We further define a hidden indicator variable S_i for $i \in \{1, \dots, t\}$. It takes the value 1 or 2, indicating which regime ξ_i is from. In the MSM, $\{S_i : i = 1, \dots, t\}$ forms a Markov chain with the following transition matrix:

$$A := \begin{bmatrix} p & 1-p \\ 1-q & q \end{bmatrix}.$$

By default, we set $S_1 = 1$ to start from the first regime. When $q = 1$, the model can be used to describe a situation where the input distribution will not revert back if a switching happens. We admit flexibility in non-zero value of $1 - q$ so that the two regimes communicate.

Given ξ^t at stage t and the simulator, we first estimate the model parameters $(\lambda_1, \lambda_2, p, q)$ through MLE and predict the distribution for ξ_{t+1} : $P(\xi_{t+1}|\xi^t)$ (see (2)). Then, we aim to solve the following online optimization problem via simulation:

$$\min_{\mathbf{x} \in \mathcal{X}} E_{\xi \sim P(\xi_{t+1}|\xi^t)}[y(\mathbf{x}, \xi)]. \quad (3)$$

The metamodel-based algorithm to solve (3) is provided in Section 4.

4 THE GP-BASED ALGORITHM

The proposed metamodel and algorithm will be presented in Section 4.1 and Section 4.2.

4.1 The Proposed GP Model

Denote the estimated model parameters after receiving the t -th input data point as: $\hat{\vartheta}^t = (\hat{\lambda}_1^t, \hat{\lambda}_2^t, \hat{p}^t, \hat{q}^t)$. From (2), we see that the input distribution for the next stage is a mixture of two stationary distributions: $\tilde{P}(\hat{\lambda}_1^t)$ and $\tilde{P}(\hat{\lambda}_2^t)$. For short, $\tilde{P}(\hat{\lambda}_i^t)$ represents $\tilde{P}(\xi|\hat{\lambda}_i^t)$, w_1^t and w_2^t represent the weights for the two components in the mixture in (2), respectively. Therefore, the input distribution $P(\xi_{t+1}|\xi^t)$ can be represented as:

$$P(\xi_{t+1}|\xi^t) = w_1^t \tilde{P}(\hat{\lambda}_1^t) + w_2^t \tilde{P}(\hat{\lambda}_2^t),$$

where the weights $w_i^t = P(S_{t+1} = i|\xi^t, \vartheta)$, $i = 1, 2$. Thus, the objective function in (3) is:

$$E_{\xi \sim P(\xi_{t+1}|\xi^t)}[y(\mathbf{x}, \xi)] = w_1^t E_{\xi \sim \tilde{P}(\hat{\lambda}_1^t)}[y(\mathbf{x}, \xi)] + w_2^t E_{\xi \sim \tilde{P}(\hat{\lambda}_2^t)}[y(\mathbf{x}, \xi)] = w_1^t g(x, \hat{\lambda}_1^t) + w_2^t g(x, \hat{\lambda}_2^t). \quad (4)$$

To find the optimal design point for the next stage, our simulation experiments will be run under the input distribution $\tilde{P}(\hat{\lambda}_1^t)$ or $\tilde{P}(\hat{\lambda}_2^t)$. As the estimated parameters generally have different values across different decision stages, the simulation results from previous stages cannot be directly used to fit the GP model which only captures the relation between $\mathbf{x}$ and Y . Also, the estimated parameter values between successive stages usually are close in terms of values, and thus the simulation results in the former stage can carry information that helps infer the simulation results in the later stage. Given a limited simulation budget, it would be extremely valuable if previous simulation results could be used for the following inference. Therefore, we propose to fit a universal GP model of Y with respect to both the decision vector $\mathbf{x}$ and the input parameter λ . Specifically, we propose to construct a stochastic GP model for the simulation output $y(x, \xi)$, $\xi \sim P(\lambda)$:

$$Y(\mathbf{x}, \lambda) = \mathbf{Z}(\mathbf{x}, \lambda) + \varepsilon(\mathbf{x}, \lambda) = v(\mathbf{x}, \lambda)^T \beta + \eta(\mathbf{x}, \lambda) + \varepsilon(\mathbf{x}, \lambda).$$

For the second-order stationary GP, $\eta(\mathbf{x}, \lambda)$, we use a mixed covariance structure: $c((\mathbf{x}_1, \lambda_1), (\mathbf{x}_2, \lambda_2)) := \text{cov}(\eta(\mathbf{x}_1, \lambda_2), \eta(\mathbf{x}_2, \lambda_2)) = \sigma_z^2 \text{corr}(\mathbf{x}_1, \mathbf{x}_2) \text{corr}(\lambda_1, \lambda_2)$, with $\text{corr}(\mathbf{x}_1, \mathbf{x}_2) = \exp\left\{\sum_{j=1}^d \theta_j^x (\mathbf{x}_{1,j} - \mathbf{x}_{2,j})^2\right\}$ and $\text{corr}(\lambda_1, \lambda_2) = \exp\left\{\sum_{j=1}^d \theta_j^\lambda (\lambda_{1,j} - \lambda_{2,j})^2\right\}$. The posterior distribution for $Z(\mathbf{x}, \lambda)$ can be used to estimate the surface of $g(\mathbf{x}, \lambda)$ and it can be derived as follows:

$$\widehat{Z}(\mathbf{x}, \lambda) = c(\mathbf{x}, \lambda)^T R^{-1} \mathcal{Y},$$

$$\text{var}(\widehat{Z}(\mathbf{x}, \lambda)) = \sigma_z^2 - c(\mathbf{x}, \lambda)^T R^{-1} c(\mathbf{x}, \lambda),$$

$$\text{cov}(\widehat{Z}(\mathbf{x}, \lambda), \widehat{Z}(\mathbf{x}', \lambda')) = c((\mathbf{x}, \lambda), (\mathbf{x}', \lambda')) - c(\mathbf{x}, \lambda)^T R^{-1} c(\mathbf{x}', \lambda'),$$

where $\mathcal{Y} = (\bar{y}(\mathbf{x}_1, \lambda_1), \dots, \bar{y}(\mathbf{x}_n, \lambda_n))^T$ is the observations vector at all the past design points $(\mathbf{x}_1, \lambda_1), \dots, (\mathbf{x}_n, \lambda_n)$.

4.2 Aggregated Expected Improvement (EI) Acquisition Function

Based on the above GP, we can proceed to model the surface of the objective in (4) with a weighted average of two sliced GP models as follows:

$$F_t(x) := w_1' Z(\mathbf{x}, \hat{\lambda}_1^t) + w_2' Z(\mathbf{x}, \hat{\lambda}_2^t) \sim \text{GP}(\mu_t, \sigma_t^2),$$

where μ_t is the predictive mean and σ_t is the predictive standard deviation. We have

$$\mu_t(x) = w_1' \widehat{Z}(\mathbf{x}, \hat{\lambda}_1^t) + w_2' \widehat{Z}(\mathbf{x}, \hat{\lambda}_2^t) \quad (5)$$

$$\sigma_t^2(x) = (w_1')^2 \text{var}(\widehat{Z}(\mathbf{x}, \hat{\lambda}_1^t)) + (w_2')^2 \text{var}(\widehat{Z}(\mathbf{x}, \hat{\lambda}_2^t)) + w_1' w_2' \text{cov}(\widehat{Z}(\mathbf{x}, \hat{\lambda}_1^t), \widehat{Z}(\mathbf{x}, \hat{\lambda}_2^t)). \quad (6)$$

The EGO algorithm (Jones et al. 1998) is one of the most popular GP-based optimization algorithms. It focuses on selecting the next point that can best improve the estimated optimal objective function value. It takes into consideration both the mean and the variance of the estimation and thus achieves a balance between exploration and exploitation. The EGO uses the expected improvement (EI) as the acquisition function, which is defined as

$$\text{EI}(x) = E[(T - F_t(x))^+] = \Delta \Phi\left(\frac{\Delta}{\sigma_t(x)}\right) + \sigma_t(x) \phi\left(\frac{\Delta}{\sigma_t(x)}\right), \quad (7)$$

where T is the approximated best current value for the surface to be optimized: $\Delta = T - \mu_t(x)$. We propose to estimate $T = \min\{\mu_t(x_1), \dots, \mu_t(x_n)\}$ where $x_1, \dots, x_n$ are previous evaluated points. $\mu_t(x)$ and $\sigma_t^2(x)$ are the posterior mean and variance of the GP model respectively, see (5) and (6). Φ and ϕ are the cdf and pdf of the standard normal random variable. The next point to be evaluated is selected with $\mathbf{x}_{n+1} = \arg \max_{x \in \mathcal{X}} \text{EI}(\mathbf{x})$. The first- and second-order derivatives of EI are easy to obtain, enabling the use of continuous second-order optimization methods to optimize (7). We then proceed to evaluate at $(\mathbf{x}_{n+1}, \hat{\lambda}_1^t)$ and $(\mathbf{x}_{n+1}, \hat{\lambda}_2^t)$ and then reconstruct the GP model in Section 4.1, based on which the subsequent point is again selected using the EI acquisition function. The whole algorithm is described in Algorithm 1.

5 EXPERIMENTS

In this section, we will use a classic (s, S) inventory problem (Fu and Healy 1997) to evaluate the empirical performance of the proposed algorithm under non-stationary input distributions.

Algorithm 1 The switching Gaussian Process based optimization Algorithm

-
- 1: In decision stage t : given the real-world data $\xi^t = \{\xi_1, \xi_2, \dots, \xi_t\}$, estimate a Markov Switching Model with parameters $(\hat{\lambda}_1^t, \hat{\lambda}_2^t, \hat{p}^t, \hat{q}^t)$, and calculate the weights w_1^t, w_2^t in the predictive distribution $P(\xi_{t+1}|\xi^t)$
 - 2: **GP model fitting**: Denote the design points that have been simulated in the past $t - 1$ stages as $\{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n\}$, where point $\mathbf{x}_i$ is simulated in stage t_i . Fit the GP model with design vector $\{(\mathbf{x}_1, \hat{\lambda}_1^{t_1}), \dots, (\mathbf{x}_n, \hat{\lambda}_1^{t_n}), (\mathbf{x}_1, \hat{\lambda}_2^{t_1}), \dots, (\mathbf{x}_n, \hat{\lambda}_2^{t_n})\}$ and the observation vector $\bar{Y}_n = [\bar{y}(\mathbf{x}_1, \hat{\lambda}_1^{t_1}), \dots, \bar{y}(\mathbf{x}_n, \hat{\lambda}_1^{t_n}), \bar{y}(\mathbf{x}_1, \hat{\lambda}_2^{t_1}), \dots, \bar{y}(\mathbf{x}_n, \hat{\lambda}_2^{t_n})]^T$.
 - 3: Set $i = 0$
 - 4: **while** $i < \text{specified \# of optimization steps in each time stage}$ **do**
 - 5: Select design point to be evaluated $\mathbf{x}_{n+1}$ using the expected improvement acquisition function
 - 6: Run simulation experiments at $(\mathbf{x}_{n+1}, \hat{\lambda}_1^t)$ and $(\mathbf{x}_{n+1}, \hat{\lambda}_2^t)$ with r replications and obtain the observed output mean $\bar{y}(\mathbf{x}_{n+1}, \hat{\lambda}_1^t)$ and $\bar{y}(\mathbf{x}_{n+1}, \hat{\lambda}_2^t)$, set $\bar{Y}_{n+1} = [\bar{Y}_n, \bar{y}(\mathbf{x}_{n+1}, \hat{\lambda}_1^t), \bar{y}(\mathbf{x}_{n+1}, \hat{\lambda}_2^t)]^T$;
 - 7: **Update**: update the stochastic GP model $Z(\mathbf{x}, \lambda)$ based on $\bar{Y}_{n+1}$ and recalculate the stochastic GP $F_t(\mathbf{x})$;
 - 8: **return** $\hat{\mathbf{x}}_t^* = \arg \min_{\mathbf{x}_1, \dots, \mathbf{x}_{n+1}} \mu_t(\mathbf{x})$
 - 9: Set $t = t + 1$ and return to Step 1.
-

5.1 The Inventory Problem

In an (s, S) inventory problem, a company manages the inventory of a single product periodically. There are two decision variables, the basic ordering level s and the order-up-to level S . For each period, such as each week, the company will manage inventory positions using the following rule: not order any new items if the inventory position is larger than s ; otherwise, the company will order and keep the inventory position to be S , i.e. order the quantity being the difference between S and the current position.

There are several hyperparameters about cost that need to be set in the inventory problem simulator. We set the fixed ordering cost = 100, unit cost = 1, holding cost = 1, and back-order cost = 100. All these hyperparameters used in this experiment are set the same as in Jalali et al. (2017).

5.2 Non-Stationary Demand Distributions

The majority of the works on the inventory problem (Wang et al. 2020; Jalali et al. 2017) focus on the stationary setting that the customer demands ξ , are independently identically distributed across different periods. However, in reality, the demand distributions are usually shifting due to external factors such as economic status, weather conditions, etc. In this work, we consider a non-stationary setting. We assume that the demand follows a Markov switching model with the emission distributions being exponential distributions:

$$\begin{aligned}\xi_t |_{S_i=1} &\sim \exp(\lambda = 1) \\ \xi_t |_{S_i=2} &\sim \exp(\lambda = 1/20).\end{aligned}$$

The transition matrix is

$$A := \begin{bmatrix} 0.7 & 0.3 \\ 0.2 & 0.8 \end{bmatrix}.$$

Figure 2 displays the time series plot of the random demands for the above switching dynamics. The distributions of the demands are significantly different in different regimes. Specifically, the demand distribution of Regime 2 has a much larger mean value and variance compared to those of Regime 1.

The candidate decision space is $\{x = [s, S] | s \in [1, 69], S \in [70, 250]\}$. The demand distributions are significantly different between different regimes, so are the optimal solutions and the optimal values: the optimal solution for Regime 1 is (1, 70) with the optimal cost being 38, while the optimal solution for Regime 2 is (63.8, 127.0) with the optimal cost being 147.




Figure 2: The time series plot of the random demands with the regimes reflected as the color in the background. The blue color corresponds to $S_t = 1$ and the red color corresponds to $S_t = 2$.

5.3 Comparison

Based on historical demand $\xi_1, \dots, \xi_{100}$ over the initial 100 time stages, we estimate a 2-regime Markov switching model with exponential emission for the demand. To start the GP-based optimization algorithm, we select 20 initial design points $\mathbf{x}$ at the beginning of the 101th stage. We proceed to optimize the inventory problem over an additional of 25 time stages. For each time stage t , we use the switching GP-based algorithm with 15 iterations to estimate the optimal solution, denoted $\hat{x}_t^*$. Denote the true minimizer at period t as $x_{S_t}^*$ and the corresponding optimal value as $g(x_{S_t}^*, \lambda_{S_t})$. The GAP value for time stage t is $g(\hat{x}_t, \lambda_{S_t}) - g(x_{S_t}^*, \lambda_{S_t})$.

We compare the proposed approach with the commonly used plug-in approach which uses the empirical distributions of historical demands ξ^t or a fitted parametric distribution (exponential distribution in our case) to drive the simulation and the EGO algorithm (Jones et al. 1998) for optimization. We consider two variations of the classic plug-in approaches which ignore the non-stationarity of the input data. The first one uses all the available input data to obtain the empirical distribution / the fitted exponential distribution, while the second one uses the historical data with a moving window of size 50: in stage t , it adopts the input data $\xi_{t-49}, \dots, \xi_t$ to fit the input distribution. In total, we have four benchmark approaches: “hist”, “hist_window”, “exp”, and “exp_window”. To ensure a fair comparison with the same evaluation budget, the alternative algorithms start with 40 initial design points and for each time stage t , we use the classic EGO algorithm with 30 iterations to estimate the optimal solution.

The experiments are repeated for 100 macroreplications. We plot the cumulative GAP of the five approaches ($\sum_{i=101}^t [g(\hat{x}_i, \lambda_{S_i}) - g(x_{S_i}^*, \lambda_{S_i})]$) for the 25 time stages in Figure 3 for this online optimization problem. It shows that the switching approach consistently outperforms the classic approaches with smaller cumulative GAP values across different time stages. We can see that the plug-in approaches using fitted exponential distributions (“exp” and “exp_window”) perform the worst. Although the input data are from exponential family, they are from two different exponential distribution. Hence, fitting the data with a single exponential distribution will be inappropriate. Adopting non-parametric distributions will slightly improve the input modeling. This can be verified through the performances of the “hist_window” and “hist”, which are very similar and better than those of “exp” and “exp_window”. However, they are still significantly worse than the proposed approach, which shows that using a moving window cannot remedy the impact from non-stationarity of the input distribution.




Figure 3: The line plot of cumulative GAP value over 100 macroreplications for the switching GP-based algorithm and the classic hist/exp approaches with EGO. The shadow areas represent the 90% confidence interval.

6 CONCLUSION

We study simulation optimization under non-stationary input distributions, where input data arrive continuously in a streaming fashion. A Markov Switching model is proposed to model the non-stationary input data and a GP model is constructed to aggregate the simulation outputs under different input distributions across time stages, given a limited simulation budget at each time stage. The expected improvement acquisition function is used to select the subsequent evaluation points. Numerical results show that our proposed procedures achieved better performance, compared to commonly used approaches which ignore the non-stationarity of the input data.

In summary, we address an online simulation optimization problem with streaming input data to deliver optimal decisions across multiple decision stages. We have made some simple assumptions in this paper, including a known and small number of regimes and a Markov structure of regime switching. In addition, we have not included the estimation error of the model parameters in the current paper. Interesting future research directions include thorough analysis of the input uncertainty and the asymptotic convergence of our framework.

ACKNOWLEDGMENTS

This work is supported by the National Natural Science Foundation (NNSF) of China under Grant 72101106 and the Shenzhen Science and Technology Program under Grant RCBS20210609103119020.

REFERENCES

- Albert, P. S. 1991. “A Two-State Markov Mixture Model for a Time Series of Epileptic Seizure Counts”. *Biometrics* 47(4):1371–1381.

- Ang, A. and G. Bekaert. 2002. "Regime Switches in Interest Rates". *Journal of Business & Economic Statistics* 20(2):163–182.
- Ankenman, B., B. L. Nelson, and J. Staum. 2010. "Stochastic Kriging for Simulation Metamodeling". *Operations Research* 58(2):371–382.
- Bhar, R. and S. Hamori. 2004. *Hidden Markov Models: Applications to Financial Economics*. New York: Springer.
- Carvalho, V. M., M. Nirei, Y. U. Saito, and A. Tahbaz-Salehi. 2021. "Supply Chain Disruptions: Evidence from the Great East Japan Earthquake". *The Quarterly Journal of Economics* 136(2):1255–1321.
- Frühwirth-Schnatter, S. 2006. *Finite Mixture and Markov Switching Models*. New York: Springer.
- Fu, M. C. and K. J. Healy. 1997. "Techniques for Optimization via Simulation: An Experimental Study on an (s, S) Inventory System". *Institute of Industrial Engineers Transactions* 29(3):191–199.
- Hamilton, J. D. 1989. "A New Approach to the Economic Analysis of Nonstationary Time Series and the Business Cycle". *Econometrica* 57(2):357–384.
- Hamilton, J. D. and B. Raj. 2002. "New Directions in Business Cycle Research and Financial Analysis". *Empirical Economics* 27:149–162.
- Ivanov, D. and A. Dolgui. 2021. "A Digital Supply Chain Twin for Managing the Disruption Risks and Resilience in the Era of Industry 4.0". *Production Planning & Control* 32(9):775–788.
- Jalali, H., I. Van Nieuwenhuysse, and V. Picheny. 2017. "Comparison of Kriging-Based Algorithms for Simulation Optimization with Heterogeneous Noise". *European Journal of Operational Research* 261(1):279–301.
- Jones, D. R., M. Schonlau, and W. J. Welch. 1998. "Efficient Global Optimization of Expensive Black-Box Functions". *Journal of Global Optimization* 13(4):455–492.
- Liu, T., Y. Lin, and E. Zhou. 2024. "Bayesian Stochastic Gradient Descent for Stochastic Optimization with Streaming Input Data". *SIAM Journal on Optimization* 34(1):389–418.
- Meng, Q., S. Wang, and S. H. Ng. 2022. "Combined Global and Local Search for Optimization with Gaussian Process Models". *INFORMS Journal on Computing* 34(1):622–637.
- Pagan, A. R. and G. W. Schwert. 1990. "Alternative Models for Conditional Stock Volatility". *Journal of Econometrics* 45(1-2):267–290.
- Shahriari, B., K. Swersky, Z. Wang, R. P. Adams and N. De Freitas. 2015. "Taking the Human Out of the Loop: A Review of Bayesian Optimization". *Proceedings of the IEEE* 104(1):148–175.
- Song, E. and U. V. Shanbhag. 2019. "Stochastic Approximation for Simulation Optimization under Input Uncertainty with Streaming Data". In *2019 Winter Simulation Conference (WSC)*, 3597–3608 <https://doi.org/10.1109/WSC40007.2019.9004677>.
- Sydnor, S., L. Niehm, Y. Lee, M. Marshall and H. Schrank. 2017. "Analysis of Post-Disaster Damage and Disruptive Impacts on the Operating Status of Small Businesses after Hurricane Katrina". *Natural Hazards* 85:1637–1663.
- Thyer, M. and G. Kuczera. 2000. "Modeling Long-Term Persistence in Hydroclimatic Time Series Using a Hidden State Markov Model". *Water Resources Research* 36(11):3301–3310.
- Wang, H., J. Yuan, and S. H. Ng. 2020. "Gaussian Process Based Optimization Algorithms with Input Uncertainty". *IIEE Transactions* 52(4):377–393.
- Williams, C. K. and C. E. Rasmussen. 2006. *Gaussian Processes for Machine Learning*. Cambridge: MIT Press.
- Wu, D., Y. Wang, and E. Zhou. 2024. "Data-Driven Ranking and Selection under Input Uncertainty". *Operations Research* 72(2):781–795.
- Zucchini, W. and P. Guttorp. 1991. "A Hidden Markov Model for Space-Time Precipitation". *Water Resources Research* 27(8):1917–1923.

AUTHOR BIOGRAPHIES

SONGHAO WANG is an assistant professor in the School of Business, Southern University of Science and Technology. He received his B.Sci. degree in mechanical engineering from University of Science and Technology of China in 2015 and his Ph.D. degree in industrial and systems engineering from NUS in 2020. His research interests include simulation optimization, statistical modeling and machine learning. His email address is wangsh2021@sustech.edu.cn.

HAOWEI WANG is currently a research scientist at Rice-Rick Digitalization PTE. Ltd. Before joining Rice-Rick, he received the B.Eng. degree in industrial engineering from Nanjing University in 2016 and the Ph.D. degree in industrial and systems engineering from National University of Singapore(NUS) in 2021. His research interest includes simulation optimization, Bayesian optimization under uncertainties. His email address is haowei_wang@ricerick.com.

JIANGLIN XIA is a Ph.D. student in the School of Business, Southern University of Science and Technology. He received his B.Sci. degree in information and computing science from Nanjing Agricultural University in 2020 and his M.Sci. degree in analytics from University of Southern California in 2022. His research interests include simulation optimization and Bayesian

optimization. His email address is 12331459@mail.sustech.edu.cn.

XIUQIN XU obtained her Ph.D. degree in Data Science from National University of Singapore in 2022. She has a Bachelor's degree in Financial Management with a minor in Applied Mathematics from Nankai University in 2016 and a Master's degree in Statistics from NUS in 2018. Her research interests include interpretable machine learning, probabilistic deep learning, functional data analysis, and high-dimensional non-stationary time series modeling. Her email address is xiuqin.xu@u.nus.edu.