Proceedings of the 2023 Winter Simulation Conference
C. G. Corlu, S. R. Hunter, H. Lam, B. S. Onggo, J. Shortle, and B. Biller, eds.

EFFICIENT INPUT UNCERTAINTY QUANTIFICATION FOR REGENERATIVE SIMULATION

Linyun He

H.Milton Stewart School of Industrial and Systems Engineering
Georgia Institute of Technology
755 Ferst Dr.NW
Atlanta, GA 30332, USA

ABSTRACT

The initial bias in steady-state simulation can be characterized as the bias of a ratio estimator if the simulation model has a regenerative structure. This work tackles input uncertainty quantification for a regenerative simulation model when its input distributions are estimated from finite data. Our aim is to construct a bootstrap-based confidence interval (CI) for the true simulation output mean performance that provides a correct coverage with significantly less computational cost than the traditional methods. Exploiting the regenerative structure, we propose a k-nearest neighbor (kNN) ratio estimator for the steady-state performance measure at each set of bootstrapped input models and construct a bootstrap CI from the computed estimators. Asymptotically optimal choices for k and bootstrap sample size are discussed. We further improve the CI by combining the kNN and likelihood ratio methods. We empirically compare the efficiency of the proposed estimators with the standard estimator using queueing examples.

1 INTRODUCTION

In stochastic simulation, input models are often estimated by fitting distributions to a limited number of real-world observations. Input uncertainty refers to the variability in the simulation output contributed by such estimation errors. In this paper, we focus on input uncertainty quantification (IUQ) for a *steady-state* simulation model. We assume that the distribution families of the input models are known, but the parameters are unknown and estimated from the data. We apply the parametric bootstrap method to find a CI that covers the true steady-state performance measure.

In steady-state simulation, the simulation output often has initial bias due to finite run time even if a warm-up period is implemented. Our goal in this paper is to *directly tackle the issue of the initial bias of a steady-state simulation in IUQ experiment design* when the simulation model has a regenerative structure and thus the steady-state output mean can be written as a ratio estimator. At each bootstrapped parameter, we divide up a single replication run into renewal cycles and construct the ratio estimator by pooling the sample paths generated within the renewal cycles from its k-nearest neighboring parameters. We propose the asymptotically optimal choice for k as well as the bootstrap sample size. To address the issue that the kNN estimation does not scale well in the parameter dimension, we propose a second ratio estimator that combines the kNN estimator with the likelihood ratio (LR) method to reduce the effect of the dimensionality. We empirically compare the two proposed estimators' performances against the standard ratio estimator to demonstrate their finite sample efficiencies.

2 PROPSED ESTIMATORS AND CONVERGENCE

Let $\theta^c \in \mathbb{R}^d$ denote the unknown true parameter of the input model and $\widehat{\theta}$ be the maximum likelihood estimator (MLE) computed from observed data. Let $\tilde{f}\left(\widehat{\theta}\middle|\theta^c\right)$ be the sampling distribution of $\widehat{\theta}$. Given

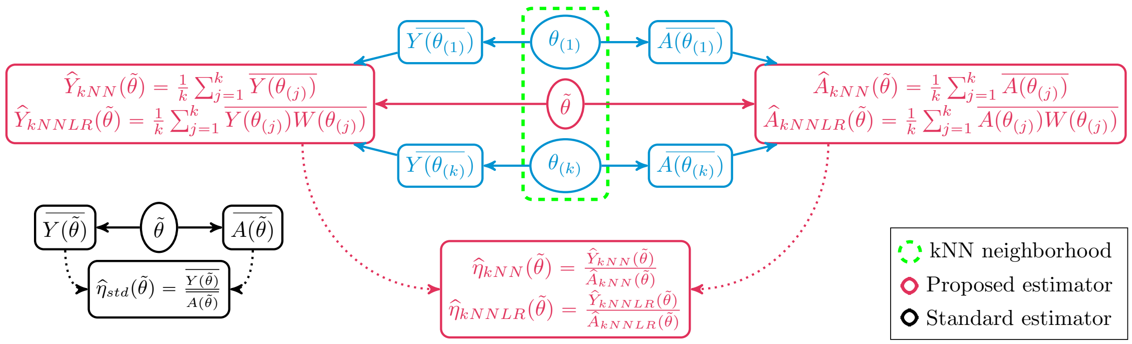

Figure 1: Schematic illustration of standard ratio estimator and the proposed kNN-based estimator.

a generic input parameter vector θ, let $Y_j(\theta) \in \mathbb{R}$ and $A_j(\theta) \in \mathbb{R}$ denote the cumulative reward and the length of the jth renewal cycle, respectively. The paired sequence $\{(Y_j(\theta), A_j(\theta))\}_{j\geq 1}$ is i.i.d. conditional on θ. The renewal reward theorem stipulates that the long-run reward is $\eta(\theta) = \mathbb{E}[Y(\theta)|\theta]/\mathbb{E}[A(\theta)|\theta]$, where the expectation is taken with respect to the inner-level simulation error run with input parameter θ.

The standard regenerative simulation estimator of $\eta(\theta)$ replaces the expectations with their respective sample averages (Glynn 2006), that is, $\widehat{\eta}_{std}(\theta) = \overline{Y_j(\theta)}/\overline{A_j(\theta)}$. To construct the CI, we first bootstrap a size-n parameter set $\{\theta_i\}_{1\leq i\leq n}$ using $\widehat{\theta}$. By running r regenerative cycles generated at each θ_i, we can obtain the ratio estimators $\widehat{\eta}_{std}(\theta_i)$ for all $1 \leq i \leq n$. Let $\widehat{q}_{\alpha,n}(\cdot)$ be the empirical α-quantile computed from a size-n sample, we can construct the bootstrap CI as $[\widehat{q}_{\alpha/2,n}(\widehat{\eta}_{std}(\theta)), \widehat{q}_{1-\alpha/2,n}(\widehat{\eta}_{std}(\theta))]$, which is expected to have $1-\alpha$ coverage as r tends to infinity.

To enhance the computational efficiency, we propose two kNN based estimators and describe the new experiment design in the following. We first bootstrap size-$\tilde{n}$ parameter set $\{\tilde{\theta}_i\}_{1\leq i\leq \tilde{n}}$, where $\tilde{\theta}_i \overset{i.i.d.}{\sim} \tilde{f}(\tilde{\theta}|\widehat{\theta})$, however, we do not run simulations at these parameters. Instead, we generate a second size-n set of parameters $\{\theta_j\}_{1\leq j\leq n}$ to drive the simulations. The two sample sizes, n and $\tilde{n}$, can differ and the two sampling distributions may not be identical. Figure 1 illustrates how the two proposed ratio estimators are constructed by pooling simulation outputs generated at $\{\theta_i\}_{1\leq i\leq n}$. Given a $\tilde{\theta}$, let $\theta_{(i)}$ be its ith nearest neighbor among $\{\theta_j\}_{1\leq j\leq n}$. We compute $\overline{Y\left(\theta_{(i)}\right)}$ and $\overline{A\left(\theta_{(i)}\right)}$ for $1 \leq i \leq k$ and take their averages to obtain $\widehat{Y}_{kNN}(\tilde{\theta})$ and $\widehat{A}_{kNN}(\tilde{\theta})$, respectively, and define $\widehat{\eta}_{kNN}(\tilde{\theta}) \triangleq \widehat{Y}_{kNN}(\tilde{\theta})/\widehat{A}_{kNN}(\tilde{\theta})$. We show the bias and mean squared error (MSE) of $\widehat{\eta}_{kNN}(\tilde{\theta})$ is $\mathscr{O}\left((k/n)^{2/d}\right)$ and $\mathscr{O}\left((k/n)^{4/d}\right) + \mathscr{O}(1/k) + \mathscr{O}(1/rk)$, respectively. Based on all $\{\widehat{\eta}_{kNN}(\tilde{\theta}_i)\}_{1\leq i\leq \tilde{n}}$, we can construct the empirical bootstrap CI.

Note that the convergence rates of the bias and the MSE for $\widehat{\eta}_{kNN}(\tilde{\theta})$ depend on the parameter dimension d. This dependence is due to the bias introduced by the kNN method. To eliminate such bias, we combine the kNN method with the LR method. Specifically, we propose another estimator $\widehat{Y}_{kNNLR}(\tilde{\theta}) = \frac{1}{k}\sum_{i=1}^{k}\overline{Y\left(\theta_{(i)}\right)W\left(\theta_{(i)}\right)}$, where $W\left(\theta_{(i)}\right)$ denotes the empirical LR between $\theta_{(i)}$ and $\tilde{\theta}$. We accordingly compute $\widehat{A}_{kNNLR}(\tilde{\theta})$ and write $\widehat{\eta}_{kNNLR}(\tilde{\theta}) \triangleq \widehat{Y}_{kNNLR}(\tilde{\theta})/\widehat{A}_{kNNLR}(\tilde{\theta})$. We show that both the asymptotic bias and MSE of $\widehat{\eta}_{kNNLR}(\tilde{\theta})$ is of order $\mathscr{O}(1/rk)$ for the exponential family input models.

We apply all three estimators to an M/M/1/10 queueing model and control the total simulation budget to compare the performance. The CI constructed by $\widehat{\eta}_{std}$ tends to show over-coverage and performs best when matching $n = r$. The $\widehat{\eta}_{kNN}$ CI exhibits under-coverage when n is small, which we improve by carefully choosing the sampling distribution of $\{\theta_j\}_{1\leq j\leq n}$. The $\widehat{\eta}_{kNNLR}$ CI shows robust performance across n, regardless of the sampling distribution of $\{\theta_j\}_{1\leq j\leq n}$.

REFERENCES

Glynn, P. W. 2006. "Simulation Algorithms for Regenerative Processes". In *Handbooks in Operations Research and Management Science*, edited by S. G. Henderson and B. L. Nelson, Volume 13, 477–500. Elsevier.