Proceedings of the 2018 Winter Simulation Conference
M. Rabe, A. A. Juan, N. Mustafee, A. Skoogh, S. Jain, and B. Johansson, eds.

GREEN SIMULATION OPTIMIZATION USING LIKELIHOOD RATIO ESTIMATORS

David J. Eckman

School of Operations Research and Information Engineering
Cornell University
206 Rhodes Hall
Ithaca, NY 14853, USA

ABSTRACT

In the setting of repeated simulation experiments, reusing past outputs to make current and future experiments more efficient is referred to as green simulation. Green simulation estimators can be naturally extended to simulation optimization, with outputs from past iterations of a search being reused in subsequent iterations to estimate the objective and gradient. However, for simulation optimization searches that identify new designs based on past outputs, outputs from different iterations are conditionally dependent given the visited designs. This conditional dependence violates a key assumption used to establish the unbiasedness of green simulation likelihood ratio estimators. We explore the consequences of the resulting conditional bias on the behavior of gradient-based optimization algorithms that use green simulation estimators.

1 INTRODUCTION

In many practical settings, a stochastic simulation experiment can assist in making recurring operational decisions; e.g., estimating project completion times. The up-to-date progress of the project's tasks can be regarded as the current *design* and simulation experiments are then repeatedly performed over time to inform the decision-maker. For a design x, let Y be a random variable drawn according to a conditional likelihood $f(\cdot;x)$ and let $h(Y)$ represent the output of a simulation model evaluated on Y. The objective, or expected performance, at a design x is given by

$$\mu(x) = \mathbb{E}_x[h(Y)] = \int_{\mathscr{Y}} h(y)f(y;x)dy,$$

where the expectation $\mathbb{E}_x$ is taken with respect to the conditional likelihood $f(\cdot;x)$ and $\mathscr{Y}$ is the support of Y. Under this model, the design does not affect the structural parameters of the simulation model, only the distributional parameters; i.e., x is not an input to the function $h(\cdot)$.

In the setting of repeated simulation experiments, green simulation refers to the reuse of outputs from past simulation experiments to enhance future experiments (Feng and Staum 2017). Green simulation is appealing because it is often computationally cheaper to reuse preexisting simulation outputs than to take a greater number of replications in future simulation experiments. The idea of reusing simulation outputs has been studied in different contexts, such as metamodeling and off-policy methods for reinforcement learning. We focus on a particular manifestation of green simulation: the likelihood ratio method, also known as the score function method. For green simulation via the likelihood ratio method, simulation outputs from a past design are reweighted by the corresponding likelihood ratio to provide an unbiased estimate at the current design. This use of the likelihood ratio to reweight outputs is closely related to importance sampling in variance reduction.

Given a sequence of designs $X_1,\ldots,X_n$ with r replications taken at each design, the green simulation likelihood ratio estimators for the objective and gradient at a given design x are

$$\widehat{\mu}_{n,r}^{ILR}(x)=\frac{1}{n}\sum_{k=1}^{n}\left[\frac{1}{r}\sum_{j=1}^{r}h\left(Y_k^{(j)}\right)\frac{f\left(Y_k^{(j)};x\right)}{f\left(Y_k^{(j)};X_k\right)}\right],\quad \text{and} \tag{1}$$

$$\widehat{\nabla\mu}_{n,r}^{ILR}(x)=\frac{1}{n}\sum_{k=1}^{n}\left[\frac{1}{r}\sum_{j=1}^{r}h\left(Y_k^{(j)}\right)\frac{f\left(Y_k^{(j)};x\right)}{f\left(Y_k^{(j)};X_k\right)}\nabla_x\log f(Y_k^{(j)};x)\right], \tag{2}$$

respectively, where for $k=1,\ldots,n$, the observations $Y_k^{(1)},\ldots,Y_k^{(r)}$ are independent and identically distributed according to the conditional likelihoods $f(\cdot;X_k)$.

2 GREEN LIKELIHOOD RATIO ESTIMATORS IN SIMULATION OPTIMIZATION

In simulation optimization, a decision-maker wishes to find the design associated with the optimal objective function value. Since simulation optimization searches typically use simulation outputs to estimate the objective or gradient, it is enticing to use green simulation methods in estimating these quantities. However, for the vast majority of simulation optimization searches, the current design has been determined by the outputs from past designs. Hence, knowing the current and past designs provides information about the outputs that led the search to the current design. In other words, conditional on the identities of the designs visited by the search, the current design is not conditionally independent of the past outputs. This violates a crucial assumption used in Feng and Staum (2017) to show that green simulation likelihood ratio estimators are unbiased in the setting of repeated experiments. When used within a simulation optimization search, the estimators in (1) and (2) will therefore be conditionally biased.

This conditional bias could theoretically be eliminated if the denominator of the likelihood ratio in (1) and (2) were replaced with the conditional likelihood of an output given the identities of the visited designs. However, this would require knowledge of the distribution of the simulation outputs which, if known, would make optimization unnecessary.

3 EXPERIMENTAL RESULTS

To illustrate how the use of green simulation likelihood ratio estimators can affect gradient-based searches, we tested a stochastic approximation algorithm on a pair of one-dimensional minimization problems: one with a constant objective function and Bernoulli distributed outputs and another with a quadratic objective function and normally distributed outputs. In both cases, the search trajectories were smoother than would be expected if green simulation estimators were not used, a likely consequence of the similarities among the gradient estimators in (2) for different values of n. Furthermore, the green simulation estimates of the objective function were negatively biased throughout the entire search, with the bias vanishing over time. Increasing the number of simulation replications taken at each design also reduced the magnitude of the bias. For the quadratic problem, the green simulation estimates of the gradient around the true minimizer were negatively biased, causing the search to continue past the true minimizer—believing that there was further improvement to be found—before looping back. These examples suggest that using green simulation estimators could hinder the convergence of a search, but further study is needed to determine whether the efficiency gains outweigh any possibly adverse effects.

REFERENCES

Feng, M., and J. Staum. 2017. "Green Simulation: Reusing the Output of Repeated Experiments". *ACM Transactions on Modeling and Computer Simulation* 27(4): Article 23, 1–28.