Green Simulation Optimization Using Likelihood Ratio Estimators

David J. Eckman

Cornell University
Operations Research and Info Eng
dje88@cornell.edu

M. Ben Feng

University of Waterloo
Statistics and Actuarial Science
ben.feng@uwaterloo.ca

Winter Simulation Conference
December 11, 2018

Model

A simulation model $h : \mathcal{Y} \mapsto \mathbb{R}$ maps a vectors of inputs, y, to a scalar output $h(y)$.

The expected performance of a design x, is given by

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

where the random vector $Y|x \sim f(\cdot; x)$.

The design affects the simulation output only through the likelihood of the inputs.

This model is not always suitable, but sometimes it is possible to "push out" any dependence of $h(\cdot)$ on x to $f(\cdot; x)$.

Examples

Example: An $M/D/1$ queue with mean interarrival time x.

- Y: vector of arrival times
- $h(Y)$: associated average waiting time
- $\mu(x)$: expected waiting time

Examples

Example: An $M/D/1$ queue with mean interarrival time x.

- Y: vector of arrival times
- $h(Y)$: associated average waiting time
- $\mu(x)$: expected waiting time

Example: A stochastic activity network with mean task durations x_i.

- Y: vector of task lengths
- $h(Y)$: associated project completion time
- $\mu(x)$: expected project completion time

Unbiased Estimators of $\mu(x)$

1. **Standard Monte Carlo**

   Take r independent replications at design x and average the outputs:

$$\widehat{\mu}_r^{SMC}(x) = \frac{1}{r}\sum_{j=1}^{r} h\left(Y^{(j)}\right).$$

Unbiased Estimators of $\mu(x)$

1. **Standard Monte Carlo**

   Take r independent replications at design x and average the outputs:

$$\widehat{\mu}_r^{SMC}(x) = \frac{1}{r}\sum_{j=1}^{r} h\left(Y^{(j)}\right).$$

2. **Likelihood Ratio Method** (importance sampling)

   Take r independent replications at design $\widetilde{x} \neq x$ and average the weighted outputs:

$$\widehat{\mu}_r^{LR}(x) = \frac{1}{r}\sum_{j=1}^{r} h\left(\widetilde{Y}^{(j)}\right) \underbrace{\frac{f\left(\widetilde{Y}^{(j)}; x\right)}{f\left(\widetilde{Y}^{(j)}; \widetilde{x}\right)}}_{\text{likelihood ratio}}.$$

Green Simulation

Setting: Repeated experiments with a sequence of random designs

$$\underbrace{X_1, X_2, \ldots, X_{n-1},}_{\text{past}} \underbrace{X_n}_{\text{current}} .$$

A design may represent exogenous conditions (e.g., economic, weather).

Assumption

The current design is independent of outputs of past designs, i.e., no feedback loop.

Main idea: *Reuse* simulation outputs from past designs to estimate the expected performance of the current design.

Green Likelihood Ratio Estimators

Suppose we have taken r independent replications from each design $X_1, \ldots, X_n$.

Green individual likelihood ratio (ILR) estimators for any point $x \in \mathcal{X}$ are given by

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

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

where $Y_k^{(j)}$ are i.i.d. $\sim f(\cdot; X_k)$ for all $j = 1, \ldots, r$ and $k = 1, \ldots, n$.

Conditional on $X_1, \ldots, X_n$, the green ILR estimators are **unbiased**.

Green Simulation Optimization

Consider the optimization problem:

$$\min_{x \in \mathcal{X}} \mu(x) = \mathbb{E}_x[h(Y)].$$

A design now represents a vector of decision variables.

An algorithm searches over the domain, $\mathcal{X}$, visiting random designs $X_1, X_2, \ldots$.

- Uses estimates of the objective function and/or gradient at the current design X_n to identify the next design X_{n+1}.
- E.g., stochastic approximation, SPSA, and simulated annealing.

Main idea: Use green simulation estimates of these quantities.

- I.e., $\widehat{\mu}^{ILR}_{n,r}(X_n)$ and $\widehat{\nabla\mu}^{ILR}_{n,r}(X_n)$.

Green Simulation Optimization

Advantages:

- Computationally cheap to reuse outputs in this way.
  - Recalculate the likelihood ratio and score each iteration.
- The green ILR estimator may have a smaller variance than the SMC estimator.

Green Simulation Optimization

Advantages:

- Computationally cheap to reuse outputs in this way.
  - Recalculate the likelihood ratio and score each iteration.
- The green ILR estimator may have a smaller variance than the SMC estimator.

Complications:

1. Correlated estimators
2. Conditionally dependent outputs
3. Conditionally biased estimators

"How do these issues manifest themselves in a search?"

Correlated Estimators

Several forms of correlation:

1. The estimators $\widehat{\mu}^{ILR}_{n,r}(X_n)$ and $\widehat{\mu}^{ILR}_{n',r}(X_{n'})$ contain similar terms.
2. The estimators $\widehat{\nabla\mu}^{ILR}_{n,r}(X_n)$ and $\widehat{\nabla\mu}^{ILR}_{n',r}(X_{n'})$ contain similar terms.
   - Gradient-based search trajectories may be **smoother**.
3. The estimators $\widehat{\mu}^{ILR}_{n,r}(X_n)$ and $\widehat{\nabla\mu}^{ILR}_{n,r}(X_n)$ contain similar terms.

Conditionally Dependent Outputs

In most simulation optimization algorithms, *the current design is determined by the outputs of past designs.*

- The independence assumption of repeated experiments is violated.

Example: Gradient-based searches (with or without green simulation).

- Knowing the designs X_{n-1} and X_n reveals additional information about the outputs $h(Y_{n-1}^{(1)}), \ldots, h(Y_{n-1}^{(r)})$ used to estimate $\nabla \mu(X_{n-1})$.

Conditional on X_{n-1}, $h(Y_{n-1}^{(j)})$ and X_n are **conditionally dependent**.

- Conditional on the visited designs $X_1, \ldots X_n$, the outputs $h(Y_k^{(j)})$ and $h(Y_k^{(j')})$ are conditionally dependent for all $k < n$ and $j \neq j'$.

Conditionally Biased Estimators

From the conditional dependence of the outputs,

$$\mathbb{E}\left[\widehat{\mu}_{n,r}^{ILR}(x) \mid X_1 = x_1, \ldots, X_n = x_n\right] \neq \mu(x), \text{ and}$$
$$\mathbb{E}\left[\widehat{\nabla\mu}_{n,r}^{ILR}(x) \mid X_1 = x_1, \ldots, X_n = x_n\right] \neq \nabla\mu(x).$$

Conditional on the visited designs $X_1, \ldots, X_n$, the estimators $\widehat{\mu}_{n,r}^{ILR}(x)$ and $\widehat{\nabla\mu}_{n,r}^{ILR}(x)$ are **conditionally biased**.

Biased estimates of the objective function and gradient *at the current design* could adversely affect simulation optimization algorithms.

Example: Minimize a Quadratic

Minimize $\mu(x) = \mathbb{E}_x[h(Y)]$ where $Y|x \sim \mathcal{N}(x, \sigma^2)$ and $h(y) = y^2$.

- $\mu(x) = \sigma^2 + x^2$ with a global minimizer $x^* = 0$.

Example: Minimize a Quadratic

Minimize $\mu(x) = \mathbb{E}_x[h(Y)]$ where $Y|x \sim \mathcal{N}(x, \sigma^2)$ and $h(y) = y^2$.

- $\mu(x) = \sigma^2 + x^2$ with a global minimizer $x^* = 0$.

Green ILR estimators are given by

$$\widehat{\mu}_{n,r}^{ILR}(x) = \frac{1}{n}\sum_{k=1}^{n}\left[\frac{1}{r}\sum_{j=1}^{r}\left(Y_k^{(j)}\right)^2 \exp\left(\frac{X_k^2 - 2Y_k^{(j)}(X_k - x) - x^2}{2\sigma^2}\right)\right], \text{ and}$$

$$\widehat{\nabla\mu}_{n,r}^{ILR}(x) = \frac{1}{n}\sum_{k=1}^{n}\left[\frac{1}{r}\sum_{j=1}^{r}\left(Y_k^{(j)}\right)^2 \exp\left(\frac{X_k^2 - 2Y_k^{(j)}(X_k - x) - x^2}{2\sigma^2}\right)\left(\frac{Y_k^{(j)} - x}{\sigma^2}\right)\right].$$

Ran 100 iterations of stochastic approximation with $X_{n+1} = X_n - 0.1\widehat{\nabla\mu}_{n,r}^{ILR}(X_n)$.

Search Trajectories

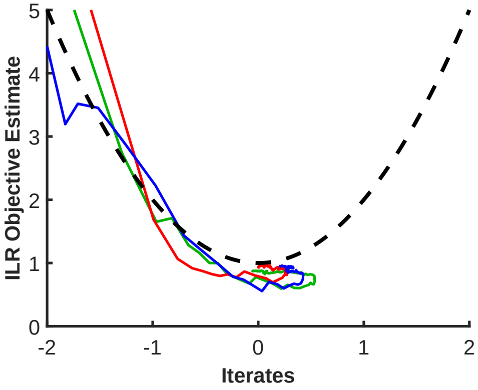

Figure: *

$r = 5$ reps/design

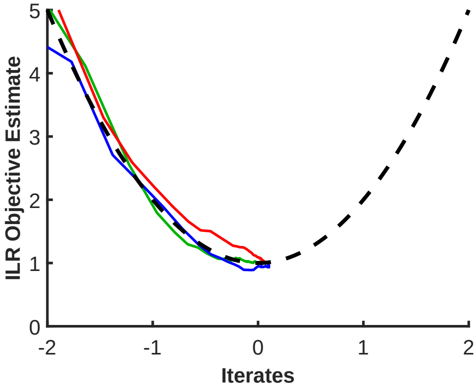

Figure: *

$r = 50$ reps/design

Conditional Biases

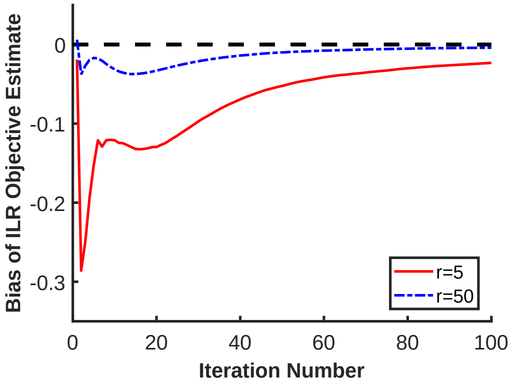

Figure: *

Bias of Objective Function Estimator

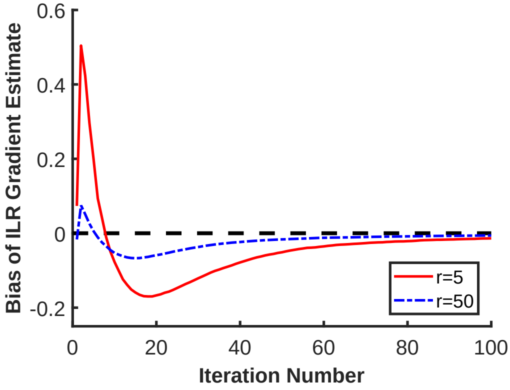

Figure: *

Bias of Gradient Estimator

Ongoing Work

Are these issues of *practical* significance?

Can the use of green simulation estimators cause a simulation optimization algorithm to fail to converge?

Under what conditions are green simulation estimators *asymptotically* unbiased and consistent as $r \to \infty$ or $n \to \infty$?

Acknowledgments

This material is based upon work supported by the National Science Foundation under grant DGE-1650441.

Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.