Green Simulation Optimization Using Likelihood Ratio Estimators

David J. Eckman (Cornell University, ORIE) and M. Ben Feng (University of Waterloo, SAS)

Green Simulation for Repeated Experiments

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 represents exogenous conditions (e.g., economic, weather).

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 $h(\cdot)$ represents the simulation model and $Y|x \sim f(\cdot; x)$ is a vector of random inputs.

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

Example: The expected waiting time of an $M/D/1$ queue with mean interarrival time x.

Main Idea of Green Simulation

Reuse outputs from **past** designs to estimate the performance of the **current** design.

- Under this model, one can use the likelihood ratio method (importance sampling).

Green individual likelihood ratio (ILR) estimators for any point x are given by

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

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

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

Key Assumption: The current design is independent of outputs of past designs.

- Under this assumption, 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 $\mathcal{X}$, visiting random designs $X_1, X_2, \ldots$ and using 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.

Idea: Use green simulation estimates of these quantities, i.e., $\hat{\mu}^{ILR}_{n,r}(X_n)$ and $\widehat{\nabla\mu}^{ILR}_{n,r}(X_n)$.

Advantages:

- Computationally cheap to reuse outputs in this way.
- The green ILR estimator may have a smaller variance than the standard Monte Carlo estimator.

Complications

- **Correlated estimates:** estimators contain similar terms for different iterations.
- **Dependence:** conditional on X_{n-1}, $Y_{n-1}^{(j)}$ and X_n are conditionally dependent.
- **Bias:** given the visited designs $X_1, \ldots, X_n$, the estimators $\hat{\mu}^{ILR}_{n,r}(x)$ and $\widehat{\nabla\mu}^{ILR}_{n,r}(x)$ are **biased**.

Experimental Results

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

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

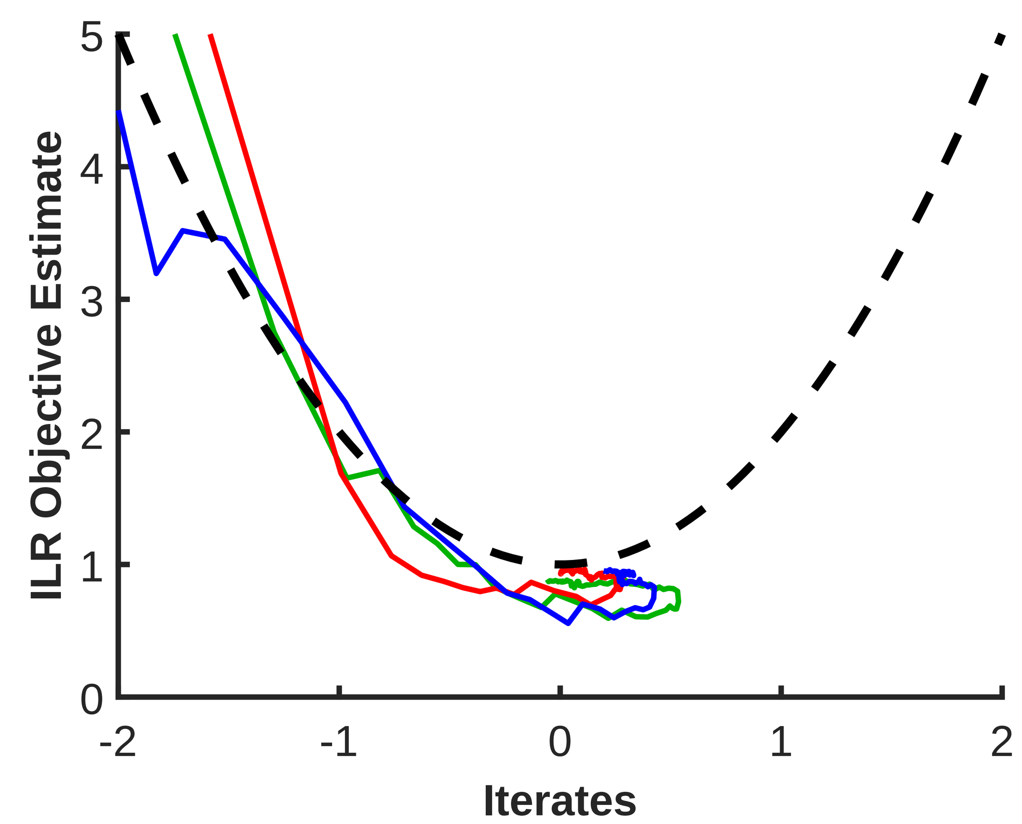

$r = 5$ reps/design

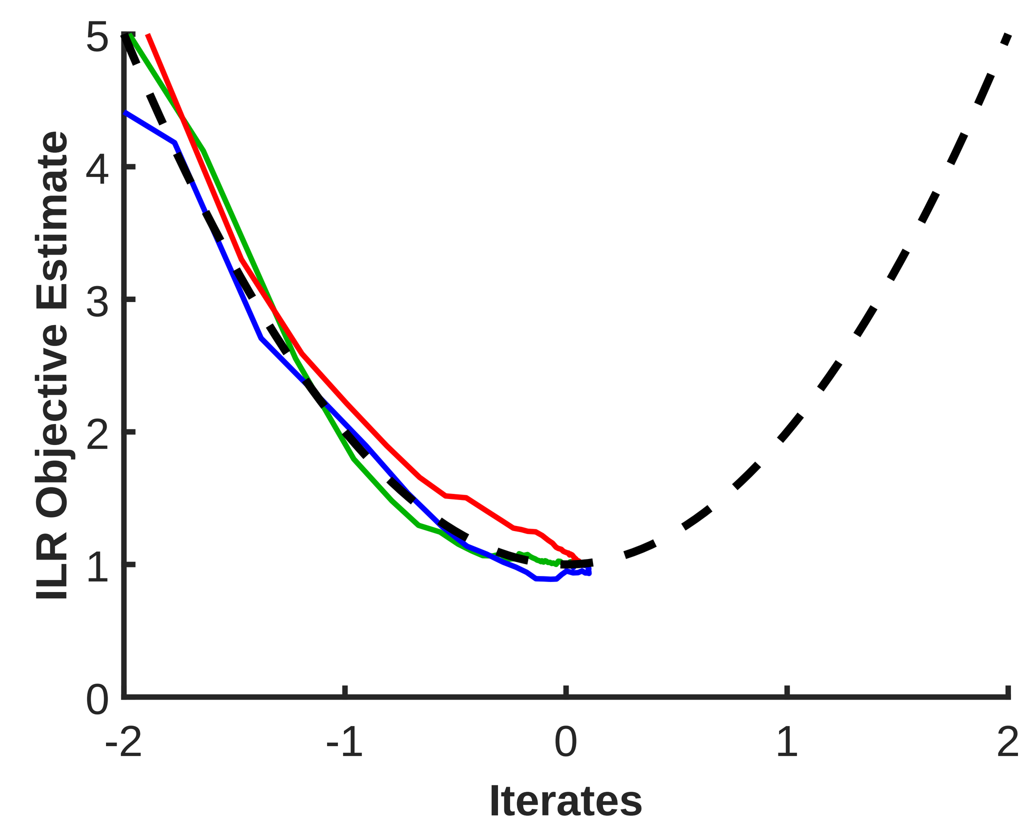

$r = 50$ reps/design

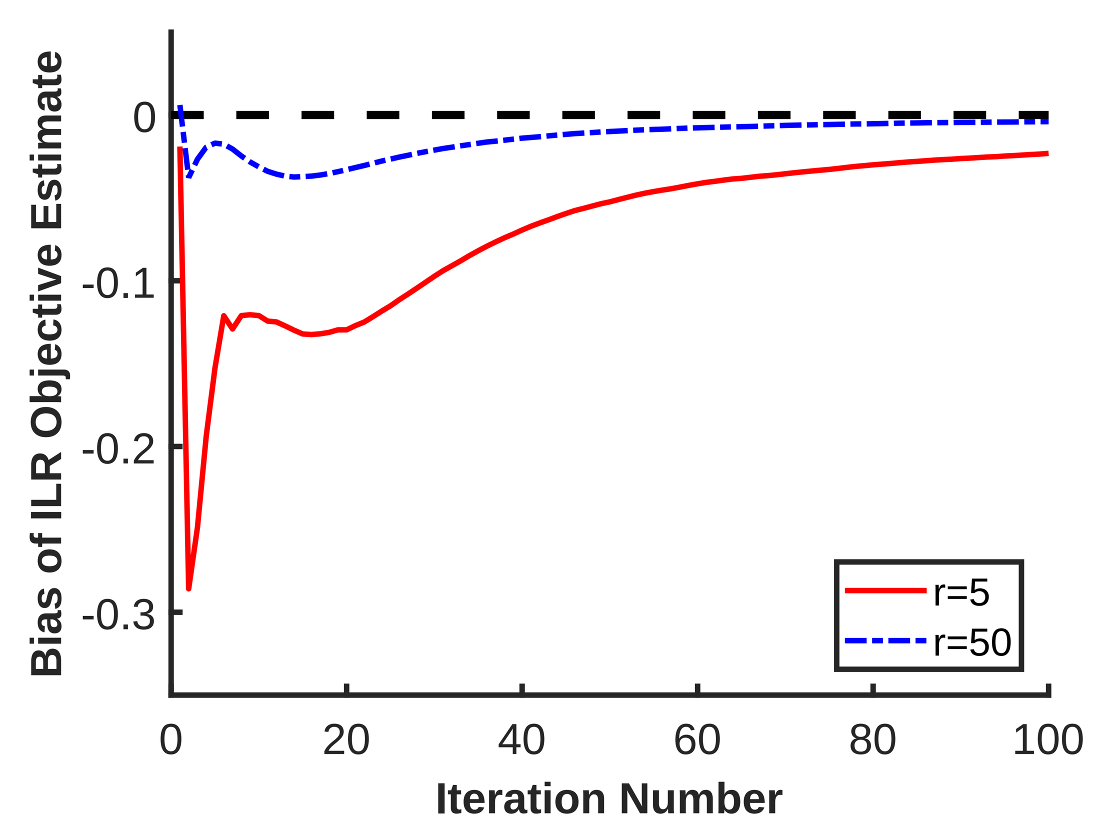

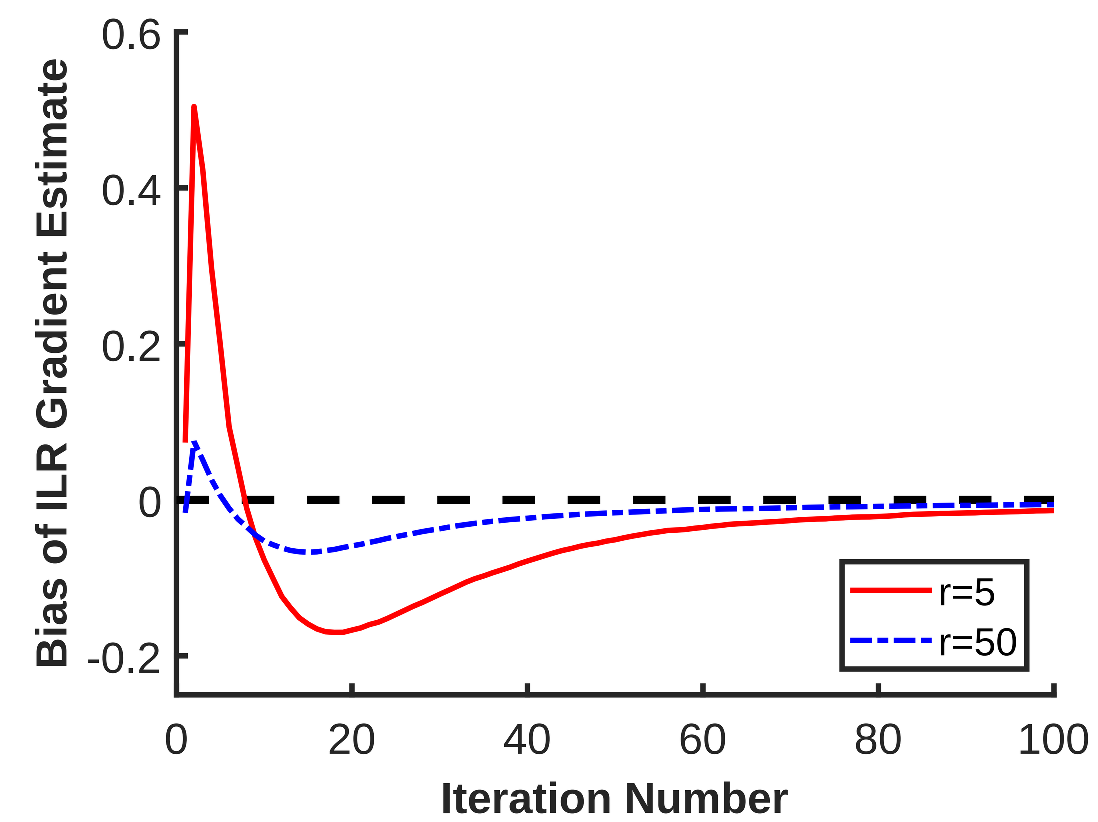

Ongoing Work

- Are these issues of *practical* significance? Can the use of green simulation estimators cause a simulation optimization algorithm to fail to converge when it otherwise would?
- Under what conditions are green simulation estimators *asymptotically* unbiased and consistent?

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