Proceedings of the 2020 Winter Simulation Conference
K.-H. Bae, B. Feng, S. Kim, S. Lazarova-Molnar, Z. Zheng, T. Roeder, and R. Thiesing, eds.

SIMULATION OPTIMIZATION BY REUSING PAST REPLICATIONS: DON'T BE AFRAID OF DEPENDENCE

Tianyi Liu

School of Industrial and Systems Engineering
Georgia Institute of Technology
Atlanta, GA 30318, USA

ABSTRACT

The main challenge of simulation optimization is the limited simulation budget because of the high computational cost of simulation experiments. One approach to overcome this challenge is to reuse simulation outputs from previous iterations in the current iteration of the optimization procedure. However, due to the dependence among iterations, simulation replications from different iterations are not independent, which leads to the lack of theoretical justification for the good empirical performance. In this paper, we fill this gap by theoretically studying the stochastic gradient descent method with reusing past simulation replications. We show that reusing past replications does not change the convergence of the algorithm, which implies the bias of the gradient estimator is asymptotically negligible. Moreover, we show that reusing past replications reduces the conditional variance of gradient estimators, which implies that the algorithm can use larger step size sequences to achieve faster convergence.

1 Problem and Main Result

We consider the following simulation optimization problem:

$$\min_{\theta\in\Theta} H(\theta) = \mathbb{E}_{\xi\sim f(\cdot;\theta)}[h(\xi)], \tag{1}$$

where θ is the decision variable, Θ is the solution space, H is the expected performance of the system, h is a sample performance function, and $f(\cdot;\theta)$ is the probability density function (pdf) of a family of parametric distributions. Stochastic gradient descent (SGD) has been widely used to solve the simulation optimization problem above in two major steps: (1). Carry out simulations to estimate the gradient $\nabla H(\theta) = \mathbb{E}_{\xi\sim f(\cdot;\theta)}[h(\xi)\nabla_\theta \ln f(\xi;\theta)]$. Here, the gradient estimator is often chosen as $\widetilde{\nabla H}(\theta_n) = \frac{1}{B}\sum_{i=1}^{B} h(\xi^i)\ln f(\xi^i;\theta)$, where $\xi^1,\ldots,\xi^B$ are independent and identically distributed (i.i.d.) samples drawn from $f(\xi;\theta)$ and B is the sample size. (2). Search the solution space following the update $\theta_{n+1} = \mathrm{Proj}_\Theta\left(\theta_n - \alpha_n\widetilde{\nabla H}(\theta_n)\right)$, where $\alpha_n > 0$ is the step size and $\mathrm{Proj}_\Theta(\theta)$ is a projection operator that projects the iterate of θ to Θ. We refer to this algorithm as vanilla stochastic gradient descent (VSGD).

One major challenge here is that the simulation alone is computationally expensive in many practical applications. As a direct result, we cannot choose a large sample size B to control the variance of the gradient estimator $\widetilde{\nabla H}(\theta_n)$, and $\mathrm{Var}\left(\widetilde{\nabla H}(\theta_n)\right) = \mathscr{O}(\frac{1}{B})$ under some minor conditions. Therefore, SGD might spend much time bouncing around, leading to slower convergence and worse performance.

To address this issue, one simple yet efficient method is to reuse simulation outputs. Specifically, we can reuse the simulation outputs from previous iterations by the principle of importance sampling to obtain a new weighed set of simulation outputs under the current solution, i.e., $\widehat{\nabla H}(\theta_n) =$

$\frac{1}{KB}\sum_{m=n-K+1}^{n}\sum_{i=1}^{B}\omega(\xi_m^i,\theta_n|\theta_m)h(\xi_m^i)\ln f(\xi_m^i;\theta_n)$, where $\{\xi_m^i, i=1,\ldots,B\}\overset{\text{i.i.d}}{\sim} f(\cdot;\theta_m)$, and $\omega(\xi,\theta_1|\theta_2) = f(\xi;\theta_1)/f(\xi;\theta_2)$ is the likelihood ratio for $m=n-K+1,...,n$, where $K\geq 1$ is the number of reused past iterations. This leads to a new gradient estimator that uses more simulation replications than the naive gradient estimator. Then the update of SGD with reusing past replications (RSGD) is: $\theta_{n+1} = \text{Proj}_\Theta\left(\theta_n - \alpha_n\widehat{\nabla H}(\theta_n)\right)$.

Ideally, since we reuse past simulation outputs, RSGD should perform better than VSGD under limited simulation budget (i.e., small B). However, reusing past replications is not a free lunch. The dependence between iterations makes the gradient estimator $\widehat{\nabla H}(\theta_n)$ biased. Specifically, $\xi_{n-1}^i|\theta_{n-1}$ and $\xi_{n-1}^i|(\theta_{n-1},\theta_n)$ are not identically distributed since θ_n can provide additional information for ξ_{n-1}^i. That is, $\xi_{n-1}^i|\theta_{n-1}^i$ and θ_n are not independent. As a result, we have the following inequality.

$$\mathbb{E}\left[h(\xi_m^i)\ln f(\xi_m^i;\theta_n)\frac{f(\xi_m^i;\theta_n)}{f(\xi_m^i;\theta_m)}\right] = \mathbb{E}\left[\mathbb{E}\left[h(\xi_m^i)\ln f(\xi_m^i;\theta_n)\frac{f(\xi_m^i;\theta_n)}{f(\xi_m^i;\theta_m)}\middle|\theta_n\right]\right] \neq \mathbb{E}_{\theta_n}\left[h(\xi)\right].$$

Despite the bias, we empirically observe that reusing past replications works well and outperforms VSGD under many different circumstances. Thus, we have the following conjecture.

The bias in the gradient estimator with past replications is asymptotically negligible.

We theoretically justify this conjecture by considering the limit ordinary differential equation (ODE) of the solution trajectory of VSGD and RSGD. We show that the continuous-time interpolation of the solution trajectory of RSGD has the same limit ODE of VSGD as shown in the following theorem.

Theorem 1 Under some regularity conditions, for a fixed $K>0$, the solution trajectory of RSGD (or VSGD when $K=1$), $\{\theta_n\}$, converges almost surely to some limit set of the following ODE in Θ.

$$\dot{\theta}(t) = \nabla_\theta H(\theta(t)) + Z(\theta(t)), \tag{2}$$

where Z is the minimum force needed to keep the solution $\theta(t)$ in Θ.

Note that ODE (2) does not depend on the number of reused iterations K. Therefore, VSGD satisfies the same limit ODE. Moreover, the following corollary shows that the bias in the gradient estimator with past replications almost does not change asymptotically. Thus, it is negligible and does not affect the asymptotic convergence (limit ODE) of RSGD.

Corollary 2 Denote the continuous-time interpolation of the bias as $\Gamma(t)$, then $\Gamma(t)$ has zero asymptotic rate of change almost surely if for any $T>0$,

$$\mathbb{P}\left(\limsup_n \max_{j\geq n}\max_{0\leq t\leq T}|\Gamma(jT+t)-\Gamma(jT)|=0\right)=1.$$

Given Theorem 1, we can reuse the past replications without worrying about the convergence. We then further study the variance reduction effect of reusing past replications. More specifically, we consider the reduction of the conditional variance given the past solution trajectory, and have the following theorem.

Theorem 3 When K is large enough, we have for any $n>0$, $i\leq d_\theta$,

$$\text{Var}[\theta_{n+1}^{(i),\text{RSGD}}|\mathscr{F}_n^{\text{RSGD}}] \leq \text{Var}[\theta_{n+1}^{(i),\text{VSGD}}|\mathscr{F}_n^{\text{VSGD}}],$$

almost surely, where $\theta^{(i)}$ is the ith dimension of θ.

Theorem 3 suggests that reusing past replications is guaranteed to reduce the variance conditioned on the past history if the number of reused iterations is large enough. Therefore, RSGD has a much smoother trajectory than VSGD, which is verified in our numerical experiment.