

BIASED GRADIENT ESTIMATORS IN SIMULATION OPTIMIZATION

David J. Eckman

Department of Industrial Engineering
and Management Sciences
Northwestern University
2145 Sheridan Road
Evanston, IL 60208, USA

Shane G. Henderson

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

ABSTRACT

Within the simulation community, the prevailing wisdom seems to be that when solving a simulation optimization problem, biased gradient estimators should not be used to guide a local-search algorithm. On the contrary, we argue that for certain problems, biased gradient estimators may still provide useful directional information. We focus on the infinitesimal perturbation analysis (IPA) gradient estimator, which is biased when an interchange of differentiation and expectation fails. Although a local-search algorithm guided by biased gradient estimators will likely not converge to a local optimal solution, it might be expected to reach a neighborhood of one. We test such a gradient-based search on an ambulance base location problem, demonstrating its effectiveness in a non-trivial example, and present some supporting theoretical results.

1 INTRODUCTION

The problem of simulation optimization, as explored here, entails minimizing a function $f(x) = \mathbb{E}h(x, Y)$ where $x \in \mathbb{R}^d$ is a vector of decision variables, Y is a random element that contains all stochastic primitives, and the function $h(\cdot, \cdot)$ encompasses the simulation model logic that translates the decision vector x and random inputs Y to a real-valued simulation output. The function $f(\cdot)$ that we want to minimize is only observed indirectly, through the simulation model $h(\cdot, \cdot)$.

Local-search algorithms for this problem iteratively attempt to improve an incumbent solution x by identifying a “nearby” solution, \tilde{x} say, such that $f(\tilde{x}) < f(x)$. The new solution \tilde{x} then becomes the incumbent and this process is repeated. Such algorithms benefit from gradient information when it is available. Ideally, we would know the gradient at the incumbent solution, $\nabla f(x)$, but invariably we cannot observe it directly. Often it is possible to compute a gradient of $h(\cdot, Y)$ for any, or almost any, fixed Y at the solution x with respect to x . Denote this gradient by $\nabla h(x, Y)$. If

$$\mathbb{E}(\nabla h(x, Y)) = \nabla \mathbb{E}(h(x, Y)), \quad (1)$$

then we have an *unbiased estimator* of the gradient of f at x , since the right-hand side (RHS) of Equation (1) equals $\nabla f(x)$. Unbiased estimators of the gradient are valuable, in that they drive gradient-based local-search methods such as stochastic approximation; see, e.g., Chapter VIII of Asmussen and Glynn (2007). When a gradient estimator $\nabla h(x, Y)$ is available, it is known within the simulation literature as an *infinitesimal perturbation analysis* (IPA) gradient estimator. IPA gradient estimators have been the focus of extensive investigation; see Ho et al. (1983), Glasserman (1991), and Fu and Hu (1997) for central references and Fu (2006), Asmussen and Glynn (2007), Kim et al. (2015), Fu (2015), and Fu and Henderson (2017) for related reviews. IPA gradient estimators can be viewed as the result of “differentiating the simulation source code,” and can be obtained either through special code or through the use of automatic differentiation (Griewank 1989) of the simulation source code.

When the interchange in Equation (1) does not hold, the IPA gradient estimator $\nabla h(x, Y)$ is *biased*; in expectation, its direction and/or magnitude does not agree with $\nabla f(x)$. While not precise, a rule of thumb is that the interchange in Equation (1) holds at a solution x if and only if the simulation output $h(\cdot, Y)$ is differentiable in its first argument for almost all Y and some additional regularity conditions (uniform integrability) hold. Thus, IPA gradient estimators are almost always biased whenever $h(x, Y)$ is discontinuous in x . Such discontinuities are extremely difficult to avoid in practical simulation modeling, and so IPA gradient estimators are typically biased. The method of smoothed perturbation analysis (Gong and Ho 1987; Zazanis and Suri 1994; Fu and Hu 1997) can be used in some cases to recover an unbiased gradient estimator, but it does not solve the problem in many, if not most, simulation models.

The folklore in simulation modeling seems to be that biased gradient estimators should not be used to attempt to solve simulation optimization problems. In this paper, we argue that such a perspective is too limiting. We contend that progress may be made on *many*, but certainly not *all*, problems by exploiting biased IPA gradient estimators. When we say that “progress may be made,” we take a practical perspective that softens the goal of reaching a locally optimal solution to reaching a *neighborhood* of one. Provided that the neighborhood is not too large, this softened goal is adequate in virtually any application. After all, simulation models are *just models*, and are rarely, if ever, perfectly calibrated to reality as evidenced by the tremendous, recent interest in model error and input uncertainty in simulation modeling, e.g., Song et al. (2014), Barton et al. (2014), and Lam (2016). Related softened goals have been discussed before, especially in the context of ordinal optimization (Ho et al. 2008). Related bias issues can arise with likelihood-ratio gradient estimators but we do not consider such estimators here.

It is important to differentiate the setting we consider from stochastic approximation algorithms using gradient estimates obtained through finite differences, as first explored in Kiefer and Wolfowitz (1952). Gradient estimators in “Kiefer-Wolfowitz” algorithms employ an increasing sample size and decreasing finite difference to carefully control the magnitude of the bias in the estimators as the algorithm proceeds; see Chapter VIII of Asmussen and Glynn (2007) for an in-depth discussion. In contrast, in our setting we do not attempt to achieve vanishing bias as a search proceeds.

Our objective in this paper is to *challenge the folklore that biased IPA gradient estimators should be ignored and instead promote their use in simulation optimization with the relaxed goal of identifying a near-optimal solution, rather than establishing convergence to an optimal solution*. Biased IPA gradient estimators will certainly not be useful in *all* cases, but perhaps they can be useful in *many*.

In this paper we advance this view primarily through an extended example, with some supporting theoretical results. In Section 2 we provide context for the application of gradient estimators in a search through two examples: one in which the IPA gradient estimator is unbiased and another in which it is so biased as to be useless for the purposes we envisage. In Section 3 we explore an example that is more practical and more representative of the applications we have in mind; the IPA gradient estimators are biased, but still provide a useful direction for a local-search algorithm when it is far from a local minimizer. We provide some limited theoretical support for the use of biased gradient estimators in Section 4 that is an outgrowth of work in Jian (2017). Section 5 develops a simple local-search algorithm that exploits biased gradient estimators and demonstrates that it can achieve robust performance on a practical example. Section 6 concludes and discusses some of the many directions for future research.

2 TOY EXAMPLES

Our first example is one in which IPA gradient estimators are unbiased—this can be viewed as the “ideal” case. Our second example is one in which IPA gradient estimators are biased to such an extent that they are useless for optimization. We present these examples to help provide context and clarify ideas before moving on to a more advanced example.

2.1 Continuous Newsvendor

A newsvendor is trying to determine how many newspapers $x \in \mathbb{R}$ to stock to meet a random demand $Y \in \mathbb{R}$. For simplicity we assume both quantities are continuous in nature, though in the classical newsvendor interpretation of this problem both are integers. Let c denote the per-unit cost of newspapers and r denote the per-unit revenue for sold newspapers. For a given stock of x newspapers and demand Y , the associated purchase cost is cx and the associated revenue is $r \min\{x, Y\}$. Our goal is to minimize the negative of the expected profit, i.e., to minimize $f(x) = \mathbb{E}(cx - r \min\{x, Y\})$. In terms of our earlier notation, $h(x, Y) = cx - r \min\{x, Y\}$.

Since x is one-dimensional, the IPA gradient estimate for a realized demand $Y = y$ is just the derivative with respect to x , which equals $c - r$ if $x < y$, c if $x > y$, and is undefined if $x = y$. The IPA gradient estimator is therefore

$$\nabla h(x, Y) = c - r \mathbb{I}(Y > x),$$

except if $Y = x$, where $\mathbb{I}(\cdot)$ is the indicator function. We arbitrarily define $\nabla h(x, Y) = 0$ on the event $\{Y = x\}$, so that $\nabla h(x, Y)$ is well defined.

Is the IPA gradient estimator unbiased? In other words, does the interchange in Equation (1) hold here? The short answer is yes, under an assumption that $P(Y = x) = 0$, which holds, for example, if Y has a density. To establish this there are more general techniques that can be employed (Glasserman 1991), but a direct calculation is as follows, proceeding from the definition of the derivative as a limit. Let $\nabla h(x, y)$ be the derivative of $h(\cdot, y)$ with respect to x , holding y fixed, if it exists. Then

$$\begin{aligned} \nabla h(x, y) &= \lim_{\delta \rightarrow 0} \frac{c(x + \delta) - r \min\{x + \delta, y\} - cx + r \min\{x, y\}}{\delta} \\ &= c - r \lim_{\delta \rightarrow 0} \frac{\min\{x + \delta, y\} - \min\{x, y\}}{\delta} \\ &= c - r \lim_{\delta \rightarrow 0} G(x, \delta, y), \end{aligned}$$

say. It is straightforward to show that $G(x, \delta, y) - \mathbb{I}(y > x) = g(x, \delta, y)$, where $g(x, \delta, y)$ is a function that satisfies $|g(x, \delta, y)| \leq \mathbb{I}(x - |\delta| \leq y \leq x + |\delta|)$. Fix some $\delta_0 > 0$, so that for all $|\delta| < \delta_0$, $|g(x, \delta, y)| \leq \mathbb{I}(x - \delta_0 \leq y \leq x + \delta_0)$. The dominated convergence theorem now allows us to conclude that

$$\begin{aligned} \mathbb{E} \nabla h(x, Y) &= c - r \mathbb{E} \left[\mathbb{I}(Y > x) + \lim_{\delta \rightarrow 0} g(x, \delta, Y) \right] \\ &= c - r P(Y > x) - r \lim_{\delta \rightarrow 0} \mathbb{E} g(x, \delta, Y) \\ &= c - r P(Y > x) - 0, \end{aligned}$$

provided that $P(Y = x) = 0$. The RHS above is $\nabla f(x)$, thereby establishing the unbiasedness of the IPA gradient estimator for the continuous newsvendor problem.

The essential ingredient in the analysis above is that the function $h(x, y)$ is appropriately differentiable in x for each y . It is this condition that fails for many practical simulation models, as illustrated in our next example.

2.2 Bus Scheduling

Bus passengers arrive to a terminal according to a Poisson process in time at rate λ customers per hour over the interval $[0, 1]$, where time is measured in hours. A bus is scheduled to depart at time 1, and one other bus departure is to be scheduled, at time x say. Buses are assumed to have infinite capacity, so that all passengers arriving before time x leave on the first bus and all those arriving after time x leave on the second. We want to select the time x so as to minimize the expected sum of passenger waiting times.

Here the random element Y consists of the number, N , of arriving passengers and the times T_1, T_2, \dots, T_N at which they arrive. This problem can be found in Ross (1996), p. 68, and was discussed in Kim et al. (2015).

Using standard properties of Poisson arrivals, one can show that $f(x) = \frac{\lambda}{2}(x^2 + (1-x)^2)$ which is a quadratic function that is easily minimized. We have that $\nabla f(x) = \lambda(2x - 1)$, so the optimal solution is $x = 1/2$, as intuition suggests.

What if we tackled this problem by simulating passenger arrivals and computing the IPA gradient estimator? Take a fixed realization $y = (n, t_1, t_2, \dots, t_n)$ of $Y = (N, T_1, T_2, \dots, T_N)$ with $n \geq 1$. The (sample) average waiting time is

$$h(x, y) = \sum_{i=1}^n [(x - t_i)\mathbb{I}(t_i \leq x) + (1 - t_i)\mathbb{I}(t_i > x)],$$

with derivative (with respect to x)

$$\nabla h(x, y) = \sum_{i=1}^n \mathbb{I}(t_i \leq x),$$

provided that $t_i \neq x$ for all $i = 1, \dots, n$, i.e., no passenger arrives at exactly the time when the first bus departs. One might hope that since this event has probability 0—just as in the newsvendor problem—the IPA gradient estimator would be unbiased. Unfortunately, this is not the case, as can be seen by noting that $\mathbb{E}\nabla h(x, Y) > 0$ for all $x > 0$, yet $\nabla f(x)$ is negative for $x < 1/2$. This complication arises because $h(x, y)$ is discontinuous in x at points x that coincide with passenger arrival times within y .

In the context of simulation optimization, the IPA gradient estimator, being nonnegative for all $x \geq 0$, would encourage a local-search algorithm to keep reducing the bus departure time. In the extreme case, the departure time would be reduced to $x = 0$, which is a maximizer of $f(x)$, not a minimizer! This example suggests that we must exercise caution in blindly applying IPA gradient estimators for simulation optimization; they can lead us astray in some problem instances. If the sample functions are discontinuous, as in this example, then is the situation hopeless? Our next example suggests not.

3 AMBULANCE BASES

Calls to an ambulance service arrive according to a Poisson process in time at rate $\lambda = 1/25$ calls per minute. Call locations are uniformly distributed in a square with side length 20 kilometers. The closest free ambulance is dispatched to the call, and travels to the call in Manhattan fashion (traveling first in the x direction, followed by the y direction) at a constant speed of $s = 60$ kilometers per hour. It then spends a random amount of time at the scene of the call that is exponentially distributed with mean 10 minutes, and afterwards returns to its base in Manhattan fashion. For simplicity, hospitals are not modeled, and ambulances returning to base are *not* diverted to new calls. Calls arising when all ambulances are busy are queued and answered in first-in-first-out order. Call arrivals, locations, and scene times are mutually independent. There are two ambulance bases, with one fixed at the location $(15, 15)$. We want to select the coordinates of the other base $x = (x_1, x_2)$ so as to minimize the expected sum of the response times—defined as the time from a call arising to an ambulance arriving at scene—over seven simulated days, starting from both ambulances being available and at base, with no queued calls.

Figure 1 provides a contour plot of the estimated objective function based on the average of 40 simulation replications. The estimated objective function value ranges from 9000 minutes at the (approximate) minimizer at the point $(9, 10.5)$ to 30,000 minutes near the upper-left and lower-right corners of the square. The halfwidths of 95% confidence intervals range from about 600 near the minimizer to about 6000 near the corners of the plot. It is perhaps surprising that the minimizer of the estimated function lies at the point $(9, 10.5)$ and not closer to the point $(5, 5)$ which would yield symmetry of the positions of the two bases in the square. To see why symmetry is not the right goal, consider the case where the fixed base is instead at the point $(19, 19)$. Would one then want the second base at the point $(1, 1)$? The approximate minimizer




Figure 1: Contour plot of the estimated objective based on 40 replications with arrows indicating the averaged IPA gradient estimates.

(9, 10.5) does not lie on the diagonal where $x_1 = x_2$ as we might expect given the symmetry in the call location distribution; we believe this to be an artifact of random sampling of the call locations.

Figure 1 also features a quiver plot of the average IPA gradient estimates (to be derived soon) from the 40 replications. The average IPA gradient estimates appear to emanate from the center of the square, with larger magnitudes near the edges, reflecting the curvature of the objective function. The average IPA gradient estimates also appear to be very nearly orthogonal to the contour lines—as the exact gradients of a smooth function would be—with slight departures from orthogonality illustrating the estimated (directional) bias of the IPA gradient estimators. The quiver and contour plots together intimate that the biases of the IPA gradient estimators are minor, suggesting that a gradient-based local-search algorithm may perform well for this problem.

Figure 2 depicts a realization of the sample-path function, $h(x, Y)$, from which the IPA gradient estimates are computed. The sample-path function appears to reflect the general structure of the objective function, with smaller values near the center of the square. Although the contours give the impression that the sample-path function is continuous in x , it is not. As one perturbs the location of the second base, some calls that were previously handled by the fixed base at (15, 15) are now handled by the base at x , and vice versa. This swapping of calls changes not only the response time of the first swapped call, but also disrupts how the ambulances respond to successive calls. These disruptions yield discontinuities in the sample-path function that are similar in nature to those observed in the bus scheduling problem.

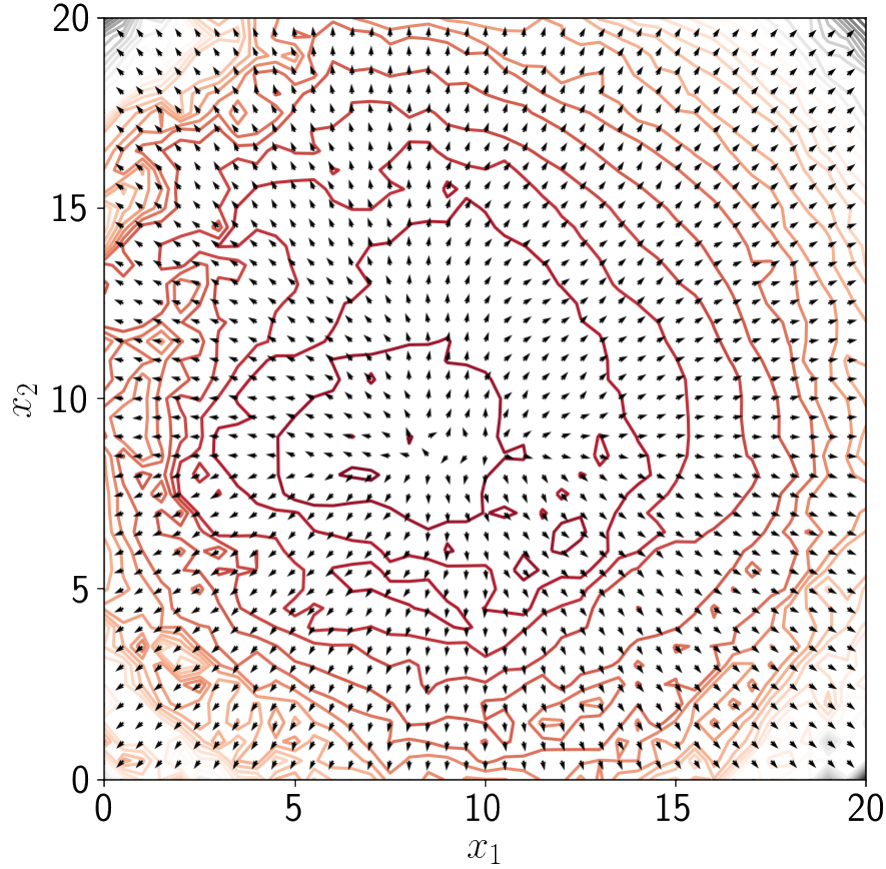


Figure 2: Contour plot of a sample-path function with arrows indicating the normalized IPA gradient estimates.

We further illustrate the discontinuity of the sample-path function by plotting its cross-section along the diagonal $x_1 = x_2$ in Figure 3. The interval $[0, 20]$ is discretized into subintervals of size 0.02 to provide a fine-resolution approximation of the discontinuous sample-path function. The plotting software connects consecutive sample-path function values, giving the impression of continuity, but the sample-path function is not continuous. The sample-path function is noticeably smoother around $x_1 = x_2 = 15$ where the two bases coincide. An intuitive explanation is that when the ambulance bases are near each other, the response time to a given call does not greatly depend on which ambulance responds. On the other hand, when the second ambulance base is located far from the first, the identity of the responding ambulance has a larger impact on the associated response time.

Are the discontinuities in the sample-path functions sufficiently disruptive as to render IPA gradient estimates useless? We compute the IPA gradient estimates as follows. Fix a base location $x = (x_1, x_2)$. Let N be the number of calls received over the course of a single replication and let S be the subset of those N calls that are handled by the ambulance based at Location x . Let $\{(X_1(i), X_2(i)) : 1 \leq i \leq N\}$ be the sequence of call locations. The response time of the i th such call, $R(i)$, consists of the time $W(i)$ spent waiting for previous calls to be completed plus $D(i)$, the driving time from the ambulance base to the call, for $i = 1, 2, \dots, N$. Here $W(i) = 0$ if an ambulance is able to respond immediately to the i th call, i.e., the i th call is not queued. Thus, the sum of the response times is

$$\sum_{i=1}^N R(i) = \sum_{i=1}^N (W(i) + D(i)),$$




Figure 3: Sample-path function along the diagonal $x_1 = x_2$.

and for $j = 1, 2$,

$$\frac{\partial}{\partial x_j} \sum_{i=1}^N R(i) = \sum_{i \in S} \left(\frac{\partial W(i)}{\partial x_j} + \frac{\partial D(i)}{\partial x_j} \right). \quad (2)$$

The response times for calls answered by the ambulance based at $(15, 15)$ do not figure in the derivative calculation. For calls answered by the ambulance based at x , the partial derivatives of $D(i)$ in Equation (2) can be obtained from the fact that $D(i) = \sum_{j=1}^2 |X_j(i) - x_j|/s$. For $j = 1, 2$, the partial derivative for $i \in S$ is

$$\frac{\partial D(i)}{\partial x_j} = \frac{\text{sgn}(x_j - X_j(i))}{s}, \quad (3)$$

provided that $X_j(i) \neq x_j$, where $\text{sgn}(\cdot)$ is the signum function.

The first term in Equation (2) can be found by analogy to the waiting times in a single-server queue, as given in, e.g., Section 2.2 of Fu and Hu (1997). For $j = 1, 2$, the partial derivative for $i \in S$ is

$$\frac{\partial W(i)}{\partial x_j} = \begin{cases} 0 & \text{if } W(i) = 0, \\ \frac{\partial W(i-)}{\partial x_j} + \frac{\partial [V(i-) - U(i-)]}{\partial x_j} & \text{if } W(i) > 0. \end{cases}$$

In this expression, $i-$ is the index of the call that immediately precedes i in S , and (arbitrarily) equals -1 if i is the first index in S . Hence, $V(i-)$ is the service time of the preceding call in S , where by “service time” we mean the time to drive to the call plus the time spent at scene plus the time spent returning to base, and $U(i-)$ is the time between the arrival of the $(i-)$ th call and the i th call. (Define $V(-1) = U(-1) = 0$.) The continuous nature of all of these random variables means that we can ignore edge cases where there are ties in times and so forth, since such events happen with probability 0. For example, we have already used this fact in our choice of the cases $W(i) = 0$ and $W(i) > 0$ for the derivative recursion, since the derivative is not well defined if the ambulance arrives back at its base x at precisely the arrival time of the next call served by that same ambulance. Moreover, we are ignoring the situation where the set of indices in S changes as we perturb x , as happens when the two ambulances are both free and equidistant from a new call.

The partial derivative of $V(i-)$ is equal to twice the partial derivative of the driving time $D(i-)$, because the ambulance must drive to the scene and back to base. The partial derivative of $U(i-)$ is zero, because ambulance base locations do not affect interarrival times and we have assumed that the set of indices in S does not change with an infinitesimal perturbation in x , which happens with probability 1. Thus,

$$\frac{\partial W(i)}{\partial x_j} = \begin{cases} 0 & \text{if } W(i) = 0, \\ \frac{\partial W(i-)}{\partial x_j} + 2\frac{\partial D(i-)}{\partial x_j} & \text{if } W(i) > 0. \end{cases} \quad (4)$$

Combining Equations (2), (3), and (4) completes the specification of the derivative calculations.

We compute these IPA gradient estimates for the sample-path function depicted in Figure 2 and superimpose a quiver plot of the normalized IPA gradient estimates. The IPA gradient estimates are normalized to better visualize the directions in which they point near the optimal solution. We do not expect perfect alignment between the IPA gradient estimates and the contours of the sample-path function owing to the presence of discontinuities in the sample-path function. Nevertheless, the direction of the IPA gradient estimates are very well oriented relative to the contours. In some areas of the feasible region where the contour lines indicate local minimizers, the IPA gradient estimates notably still point away from the center of the square. This phenomenon could help a gradient-based local-search algorithm to continue moving past local minimizers and towards the global minimizer.

4 CONVERGENCE TO A NEIGHBORHOOD

We envisage using IPA gradient estimates in a local-search algorithm that uses gradient information; an example of such an algorithm is given in Section 5. In this section we do not specify the local-search algorithm. Instead we state results that apply to *any* local-search algorithm. These results lend support to the idea of using biased gradient estimators in simulation-optimization search algorithms, and make clear that one should accordingly expect convergence to a *neighborhood* of a locally optimal solution, rather than convergence to a locally optimal solution itself. For simplicity, we provide these results under quite stringent conditions on the function f , namely that f is strongly convex.

Definition 1 A differentiable function $g : \mathbb{R}^d \rightarrow \mathbb{R}$ is σ -strongly convex for some $\sigma > 0$ if, for all $x_1, x_2 \in \mathbb{R}^d$,

$$g(x_2) \geq g(x_1) + \nabla g(x_1)^T (x_2 - x_1) + \frac{\sigma}{2} \|x_2 - x_1\|_2^2.$$

We also make strong assumptions about the quality of IPA gradient estimates relative to the true gradients. Under these conditions the results are straightforward to establish. Such a simplified setting also helps us to make our key point with the minimum of preamble and related machinery. Our key point is this:

When a local search algorithm is guided by biased gradient estimators, we should aspire to convergence to a neighborhood of a locally optimal solution, but not convergence to a local optimum.

As noted in the introduction, we believe this to be a reasonable goal, or even the “right” goal in simulation-optimization problems.

Recall that $\nabla h(x, Y)$ is the IPA gradient estimator at a point x for a single realization Y of the random elements in the simulation. Let

$$\nabla f_n(x) = \frac{1}{n} \sum_{i=1}^n \nabla h(x, Y_i)$$

be the sample average of n such terms, where the random elements Y_1, Y_2, \dots, Y_n are i.i.d. and distributed as Y . Then $\nabla f_n(x)$ is our (average) IPA gradient estimator of $\nabla f(x)$. We assume that the local-search algorithm terminates with a point, X , such that $\|\nabla f_n(X)\| = \varepsilon_1$. Here we view ε_1 as the norm of the

estimated gradient at the (random) point X at which the local-search algorithm terminates and not as an error tolerance. The reason is that we cannot guarantee the existence of a (random) point X_ε such that $\|\nabla f_n(X_\varepsilon)\| \leq \varepsilon$ for arbitrary $\varepsilon > 0$, since the estimated (random) function $f_n = n^{-1} \sum_{i=1}^n h(\cdot, Y_i)$ may have no stationary points. Thus, ε_1 is a random variable.

We require that

1. $\sup_x \|\nabla f_n(x) - \mathbb{E}\nabla h(x, Y)\| \leq \varepsilon_2(n)$ for some $\varepsilon_2(n) > 0$ almost surely, and
2. $\sup_x \|\mathbb{E}\nabla h(x, Y) - \nabla f(x)\| \leq \varepsilon_3$ for some $\varepsilon_3 > 0$.

The first condition above provides a uniform bound on the simulation error in the IPA gradient estimates relative to the mean of the IPA gradient estimates. The error bound $\varepsilon_2(n)$ is a function of n because the simulation error is expected to converge to 0 as $n \rightarrow \infty$ almost surely. It is also a random variable, since the left-hand side of the inequality is random. The assumption that the bound holds uniformly in x is a strong one that is adopted for its simplicity and not for its generality. Related uniform error bounds have been used, for example, in the analysis of sample-average approximation (Shapiro 2003). The second condition above provides a uniform bound on the bias of IPA gradient estimators relative to the true gradients $\nabla f(x)$. We do not present any methods for verifying this condition here; such methods are the topic of current research.

We are now in a position to state a result about convergence.

Proposition 1 Suppose that f is differentiable and σ -strongly convex, and let x^* be its unique minimizer. Suppose further that Conditions 1 and 2 hold, and let $\varepsilon_1 \geq \|\nabla f_n(X)\|$. Then, almost surely,

$$f(X) - f(x^*) \leq \frac{(\varepsilon_1 + \varepsilon_2(n) + \varepsilon_3)^2}{2\sigma}.$$

Proof. The result is a direct consequence of Inequality (4.12) in Bottou et al. (2018) which asserts that $f(x) - f(x^*) \leq (2\sigma)^{-1} \|\nabla f(x)\|_2^2$ for any $x \in \mathbb{R}^d$. Indeed, if Y is independent of X and distributed as a generic random input to the simulation, then almost surely,

$$\begin{aligned} \|\nabla f(X)\| &= \|(\nabla f(X) - \mathbb{E}[\nabla h(X, Y)|X]) + (\mathbb{E}[\nabla h(X, Y)|X] - \nabla f_n(X)) + \nabla f_n(X)\| \\ &\leq \|\nabla f(X) - \mathbb{E}[\nabla h(X, Y)|X]\| + \|\mathbb{E}[\nabla h(X, Y)|X] - \nabla f_n(X)\| + \|\nabla f_n(X)\| \\ &\leq \varepsilon_3 + \varepsilon_2(n) + \varepsilon_1. \end{aligned}$$

□

Proposition 1 shows that if the local search terminates at a point X with an estimated gradient that is small in norm (ε_1), the estimated gradients are close to their expected values (ε_2), and the bias in the estimated gradients is small (ε_3), then the point X is close to optimal. The right-hand side of the error bound in Proposition 1 is a random variable, which is perhaps unconventional. It is possible to develop related results where the error bound is deterministic, but the bound then holds only with some probability.

Proposition 1 supports our contention that a local-search algorithm guided by biased gradient estimators can approach a neighborhood of a locally optimal solution. The result says nothing about convergence to a locally optimal solution. Such a goal is challenging when a local search is guided by biased gradient estimators. Kiefer-Wolfowitz algorithms (Kiefer and Wolfowitz 1952; Asmussen and Glynn 2007) employ biased finite-difference gradient estimators to obtain convergence to a local minimizer, but they use increasing sample sizes as the algorithm progresses to ensure that the bias vanishes as the algorithms proceed. We make no attempt to reduce bias, instead contenting ourselves with convergence to a neighborhood of a local minimizer.

5 A LOCAL-SEARCH ALGORITHM

We describe gradient-based local-search algorithms that can easily incorporate IPA gradient estimates and evaluate the performance of one such algorithm on the ambulance base location problem introduced in Section 3.

We characterize a gradient-based local-search algorithm as one that starts at a solution x_0 and moves from solution to solution, guided by estimates of the gradient of the objective function. Specifically, on each iteration, the algorithm takes n replications at the current solution, computes a gradient estimate, and takes a step in the direction of the negative gradient estimate. In general the number of replications, n , can vary with the iteration, but for simplicity we take it as fixed across iterations. Mathematically, we have in mind a rudimentary version of stochastic approximation for which the sequence of iterates is given by either

$$x_{k+1} = x_k - a_k \nabla f_n(x_k) \quad \text{or} \quad x_{k+1} = x_k - a_k \frac{\nabla f_n(x_k)}{\|\nabla f_n(x_k)\|} \quad (5)$$

for $k = 0 = 1, 2, \dots$, where a_0, a_1, \dots is a sequence of positive scalar constants. In the former case, the step size is affected by the magnitude of the gradient estimate, whereas for the latter only the *direction* of the gradient estimate is used.

Local-search algorithms of this form employ a variety of natural stopping rules. In our simulation experiments, we fix an overall simulation budget—measured in the number of replications—and run the algorithm until the budget is exhausted. This setup facilitates aggregating results from across multiple macroreplications (runs) of the algorithm and also allows us to study its long-term behavior; e.g., do the iterates eventually remain within a neighborhood of a local optimal solution? Alternatively, the search could terminate when the norm of the estimated gradient is sufficiently small, or when the estimated objective function value at the next iterate fails to indicate an improvement relative to the current iterate. We did not pursue such alternative stopping rules.

To test the effectiveness of gradient-based local-search algorithms on the ambulance base location problem, we consider an extreme case in which only a single replication is taken at each solution, i.e., $n = 1$. The algorithm employs common random numbers across solutions; thus when using IPA gradient estimates defined in Equation (2), it effectively behaves like a (deterministic) gradient-descent algorithm on the fixed sample-path function. The search uses the normalized IPA gradient estimate, as described on the right-hand side of Equation (5), with $a_k = 0.5$ for all k . We run the algorithm with a budget of 60 replications, so that it can theoretically reach any solution in the feasible region from any initial solution within the simulation budget, provided the directions of the gradient estimates are appropriately oriented.

The left panel of Figure 4 shows the contours of a fixed sample-path function and the trajectories of the search when starting from solutions at the corners of the square and using the corresponding IPA gradient estimates. Even though the local-search algorithm uses a single replication to estimate the gradient, it appears to perform well, with all four trajectories moving towards the center of the square and remaining within a neighborhood of the point (9,9). While the search trajectories tend to follow the local geometry of the sample-path function, the IPA gradient estimators do not always appear to align with the function's contours. For instance, as the trajectory originating at (0,20) approaches the point (4.5,14), it takes a couple of steps down and to the right while the contours suggest descent up and to the right. In addition, the trajectories originating at (0,0) and (20,0) pass over local minimizers early on and continue on towards the center of the square.

We repeat the experiment detailed above for 10 independent sample-path functions, tracking the trajectories of a local-search algorithm when starting from the corners of the square. The final solutions x_{60} from each of these 40 macroreplications are shown in a scatter plot in the right panel of Figure 4. The final solutions tend to concentrate in the center of the square and are all of high-quality, as can be inferred from the contour plots of the objective function. Other experimental results indicate that the local-search algorithm is remarkably consistent across macroreplications in terms of its rapid progress in finding near-optimal solutions, often within the first 20 iterations.

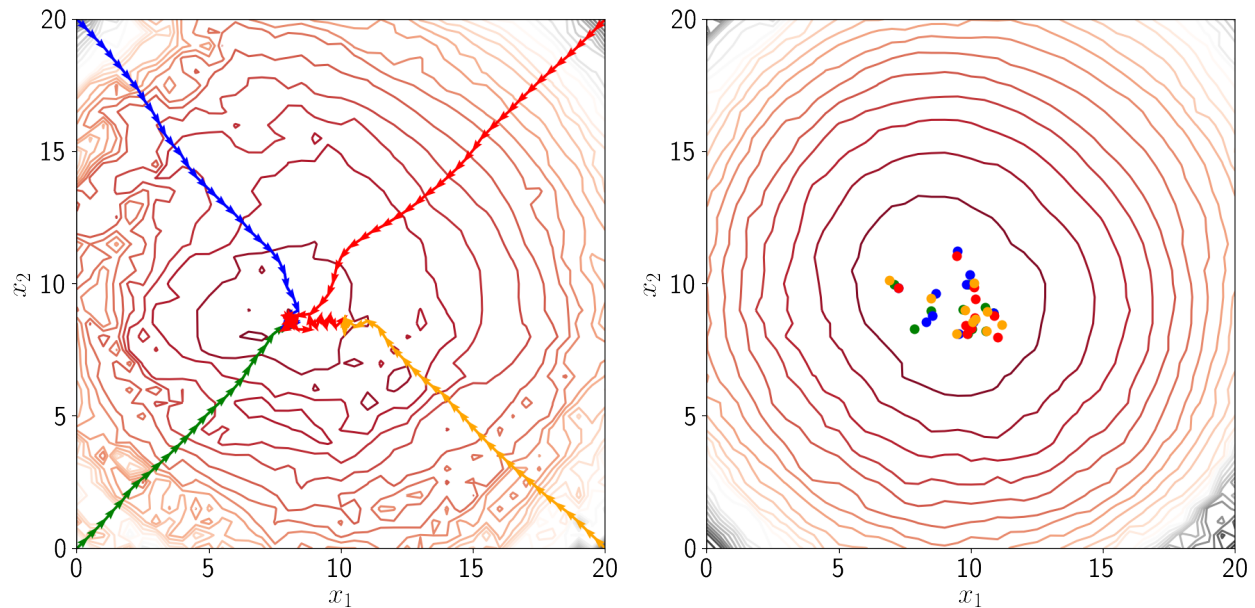


Figure 4: A local-search algorithm starting from four initial solutions. Contours of the sample-path function with four trajectories (left) and contours of the estimated objective function with final solutions from 40 macroreplications (right). The color of each final solution in the right panel indicates the starting point of the search as in the left panel.

6 CONCLUSION

We have argued that biased IPA gradient estimators might be useful in simulation optimization, provided one softens one's goal from convergence to a local minimizer to convergence to a *neighborhood* of a local minimizer. The ambulance base problem considered here, while highly stylized, supports this overall conclusion. Still, one example does not prove a general point. Therefore, much remains to be explored:

1. Trying this idea on multiple test problems might lead to additional insight as to when it can be expected to work well and when not. Doing so might even lead to practical guidelines for its applicability.
2. The ideas in Section 4 provide some theoretical support for this approach. Can more be done to understand the behavior of local-search algorithms that used biased gradient estimators, e.g., convergence rates to a neighborhood?
3. As seen in Section 4, the bias in IPA gradient estimators plays a central role in the quality of solutions obtained via gradient-based search algorithms. Understanding the nature of that bias in classes of applications, and perhaps even quantifying it, would add insight. The theory of conditional Monte Carlo as found in Fu and Hu (1997) seems especially relevant.

ACKNOWLEDGMENTS

This work was partially supported by the Army Research Office under grant W911NF-17-1-0094 and the National Science Foundation under grant TRIPODS+X DMS-1839346.

REFERENCES

Asmussen, S., and P. W. Glynn. 2007. *Stochastic Simulation: Algorithms and Analysis*, Volume 57 of *Stochastic Modeling and Applied Probability*. New York: Springer.

- Barton, R. R., B. L. Nelson, and W. Xie. 2014. "Quantifying Input Uncertainty via Simulation Confidence Intervals". *INFORMS Journal on Computing* 26(1):74–87.
- Bottou, L., F. E. Curtis, and J. Nocedal. 2018. "Optimization Methods for Large-Scale Machine Learning". *SIAM Review* 60(2):223–311.
- Fu, M. C. 2006. "Gradient Estimation". In *Simulation*, edited by S. G. Henderson and B. L. Nelson, Handbooks in Operations Research and Management Science, 575–616. Amsterdam: Elsevier.
- Fu, M. C. 2015. "Stochastic Gradient Estimation". In *Handbook of Simulation Optimization*, edited by M. C. Fu, Volume 216 of *International Series in Operations Research & Management Science*, Chapter 5, 105–148. Springer New York.
- Fu, M. C., and S. G. Henderson. 2017. "History of Seeking Better Solutions, AKA Simulation Optimization". In *Proceedings of the 2017 Winter Simulation Conference*, edited by W. K. V. Chan, A. D'Ambrogio, G. Zacharewicz, N. Mustafee, G. Wainer, and E. Page, 131–157. Piscataway NJ: IEEE.
- Fu, M. C., and J.-Q. Hu. 1997. *Conditional Monte Carlo: Gradient Estimation and Optimization Applications*. Boston: Kluwer.
- Glasserman, P. 1991. *Gradient Estimation Via Perturbation Analysis*. The Netherlands: Kluwer.
- Gong, W.-B., and Y.-C. Ho. 1987. "Smoothed (Conditional) Perturbation Analysis of Discrete Event Dynamical Systems". *IEEE Transactions on Automatic Control* 32(10):858–866.
- Griewank, A. 1989. "On Automatic Differentiation". In *Mathematical Programming: Recent Developments and Applications*, 83–108: Kluwer Academic Publishers.
- Ho, Y. C., X. R. Cao, and C. G. Cassandras. 1983. "Infinitesimal and Finite Perturbation Analysis for Queueing Networks". *Automatica* 19:439–445.
- Ho, Y. C., Q. C. Zhao, and Q. S. Jia. 2008. *Ordinal Optimization: Soft Optimization for Hard Problems*. Springer.
- Jian, N. 2017. *Exploring and Exploiting Structure in Large Scale Simulation Optimization*. Ph. D. thesis, Operations Research and Information Engineering, Cornell University, Ithaca NY.
- Kiefer, J., and J. Wolfowitz. 1952. "Stochastic Estimation of the Maximum of a Regression Function". *Annals of Mathematical Statistics* 23:462–466.
- Kim, S., R. Pasupathy, and S. G. Henderson. 2015. "A Guide to Sample Average Approximation". In *Handbook of Simulation Optimization*, edited by M. C. Fu, Volume 216 of *International Series in Operations Research & Management Science*, Chapter 8, 207–244. Springer New York.
- Lam, H. 2016. "Advanced Tutorial: Input Uncertainty and Robust Analysis in Stochastic Simulation". In *Proceedings of the 2016 Winter Simulation Conference*, edited by T. M. K. Roeder, P. I. Frazier, R. Szechtman, E. Zhou, T. Huschka, and S. E. Chick, 178–192. Piscataway, NJ: IEEE.
- Ross, S. M. 1996. *Stochastic Processes*. 2nd ed. New York: Wiley.
- Shapiro, A. 2003. "Monte Carlo Sampling Methods". In *Stochastic Programming*, edited by A. Ruszczyński and A. Shapiro, Handbooks in Operations Research and Management Science, 353–425. Amsterdam: Elsevier.
- Song, E., B. L. Nelson, and C. D. Pegden. 2014. "Advanced Tutorial: Input Uncertainty Quantification". In *Proceedings of the 2014 Winter Simulation Conference*, edited by A. Tolk, D. Diallo, I. O. Ryzhov, L. Yilmaz, S. Buckley, and J. A. Miller, 162–176. Piscataway NJ: IEEE.
- Zazanis, M., and R. Suri. 1994. "Perturbation Analysis of the GI/GI/1 Queue". *Queueing Systems* 18:199–248.

AUTHOR BIOGRAPHIES

DAVID J. ECKMAN is a postdoctoral research fellow in the Department of Industrial Engineering and Management Sciences at Northwestern University. His research interests deal with optimization and output analysis for stochastic simulation models. His e-mail address is david.eckman@northwestern.edu.

SHANE G. HENDERSON is a professor in the School of Operations Research and Information Engineering at Cornell University. His research interests include discrete-event simulation and simulation optimization, and he has worked for some time with emergency services and bike-sharing applications. He co-edited the Proceedings of the 2007 Winter Simulation Conference. His web page is <http://people.orie.cornell.edu/~shane>. and his email address is sgh9@cornell.edu.