Proceedings of the 2008 Winter Simulation Conference
S. J. Mason, R. R. Hill, L. Mönch, O. Rose, T. Jefferson, J. W. Fowler eds.

THE MATHEMATICS OF CONTINUOUS-VARIABLE SIMULATION OPTIMIZATION

Sujin Kim

Department of Industrial and Systems Engineering
10 Kent Ridge Crescent, National University of Singapore
Singapore, 119260, SINGAPORE

Shane G. Henderson

School of Operations Research and Information Engineering
Cornell University
Ithaca, NY 14853, U.S.A.

ABSTRACT

Continuous-variable simulation optimization problems are those optimization problems where the objective function is computed through stochastic simulation and the decision variables are continuous. We discuss verifiable conditions under which the objective function is continuous or differentiable, and outline some key properties of two classes of methods for solving such problems, namely sample-average approximation and stochastic approximation.

1 INTRODUCTION

Consider optimization problems of the form

$$\min_{x \in D} f(x), \tag{1}$$

where $f(x) = Ef(x,\xi)$ and the domain, D, is a convex subset of $\mathbb{R}^d$. We assume that $E|f(x,\xi)| < \infty$ for all $x \in D$. We also assume that $Ef(x,\xi)$ cannot be computed exactly, and is instead estimated through simulation. The domain D is assumed deterministic and known. In other words, we do not consider the case where the domain D needs to be estimated via constraint functions that are also evaluated through simulation. We also assume that D is convex, so the variables x take values in a continuum, excluding the trivial case where D consists of a single point.

This simulation-optimization problem has received a great deal of attention in the literature, perhaps owing to its generality, the many problems of practical interest that can be cast as such problems, and, certainly not least, its intellectual challenges. For various introductions and perspectives on the problem see Kim (2006) and earlier simulation-optimization tutorials at the Winter Simulation Conference, the optimization-related chapters in Henderson and Nelson (2006), Chapters VII and VIII of Asmussen and Glynn (2007), Fu (2002), Andradóttir (1998) and Chapter 12 of Law (2007).

Our goal in this paper is to review some key mathematical concepts that often play an important role in the analysis and solution of simulation-optimization problems.

We begin by discussing verifiable conditions under which the function $f(\cdot)$ is continuous, along with closely related properties like uniform continuity and differentiability. This is important for two primary reasons. First, any method for solving practical optimization problems of the form (1) is approximate, in the sense that it cannot guarantee that it will return a point x that is an *exact* local minimum of $f(\cdot)$. Instead, all that can be hoped for is that x will be *close* to a local minimum, x^* say, of $f(\cdot)$. We would, of course, like $f(x)$ to be close to $f(x^*)$. A simple sufficient condition for this to hold is that $f(\cdot)$ be continuous at x^*. Second, many optimization methods include a local-search component whereby a current guess for the optimal solution is replaced by the best of some "nearby" solutions. Here, "nearby" means solutions that are close as defined by some metric which is often the Euclidean metric. Again continuity of $f(\cdot)$ is desirable because, without it, one cannot expect the simulation estimates of $f(\cdot)$ to be continuous, and then it is not clear that there is value in searching in a neighborhood of the current solution.

We then turn our attention to 2 general approaches to simulation optimization. First, we describe key concepts related to a collection of methods known variously as sample-average approximation, the stochastic counterpart method, retrospective optimization, and sample-path optimization. Second, we discuss what is perhaps the most famous, in academic circles at least, simulation-optimization algorithm, namely stochastic approximation.

Any paper-length coverage of simulation optimization will necessarily be incomplete. So what do we *not* cover here? We do not cover results related to meta-heuristics such as genetic algorithms and tabu search, e.g., Glover, Kelly, and Laguna (1996). We do not discuss results used to prove that various search algorithms converge to local or global optima, e.g., Hong and Nelson (2006). We do not discuss the very important and applicable area of meta-

978-1-4244-2708-6/08/$25.00 ©2008 IEEE

modeling, including the very promising area of kriging, e.g., Barton and Meckesheimer (2006) and Ankenman, Nelson, and Staum (2008). We do not cover the frameworks of the cross-entropy method, e.g., Rubinstein and Kroese (2004), or model-reference adaptive search Hu, Fu, and Marcus (2007). Finally, we do not discuss results related to the difficulty of solving global optimization problems, e.g., Calvin (2004).

In summary, we believe the primary contribution of this paper is to survey some key ideas related to establishing structural properties of the true function we wish to optimize, as well as to review some standard approaches to developing optimization algorithms.

2 CONTINUITY AND DIFFERENTIABILITY

In the previous section we discussed the important role that continuity of $f(\cdot)$ can play in simulation optimization. But how can we establish that $f(\cdot)$ is continuous in x, when it is not directly observable?

A very general approach to proving that $f(\cdot)$ is continuous is coupling theory. For introductions to coupling theory, see Lindvall (1992) and Thorisson (2000). Let x_1 and x_2 be two points in the domain D, and let ξ_1 and ξ_2 be random objects that have the same distribution as ξ. Then $f(x_1)-f(x_2)=Ef(x_1,\xi_1)-Ef(x_2,\xi_2)$. If ξ_1 and ξ_2 are constructed on the same probability space, then

$$f(x_1)-f(x_2)=E[f(x_1,\xi_1)-f(x_2,\xi_2)]. \tag{2}$$

Properties of the unobservable $f(x_1)-f(x_2)$ can now be deduced from corresponding properties of $f(x_1,\xi_1)-f(x_2,\xi_2)$. One is free to choose the joint distribution of (ξ_1,ξ_2) in any way, and the idea in coupling is to choose the joint distribution to one's advantage. For example, the method of common random numbers (CRN) arises when we take $\xi_1=\xi_2$.

As an example, recall that we want to show that $f(\cdot)$ is continuous. In view of (2), one might postulate that this will occur whenever $f(\cdot,\xi)$ is continuous for any choice of ξ. It turns out that something like this is indeed the case as we will show shortly, but some additional conditions need to be imposed.

Example 1. Take the domain $D=[0,1/2]$, and $\xi\sim U(0,1)$. For $x>0$ define

$$f(x,\xi)=\begin{cases}\xi/x^2 & 0<\xi\le x\\ 2/x-\xi/x^2 & x\le\xi\le 2x\\ 0 & 2x\le\xi<1,\end{cases}$$

and for $x=0$ define $f(x,\xi)=0$ for all ξ. One can check that $f(\cdot,\xi)$ is continuous for all ξ. Furthermore $f(x)=Ef(x,\xi)=1$ for all $x\in(0,1]$ and $f(0)=Ef(0,\xi)=0$, so that $f(\cdot)$ is not continuous at 0.

The problem in this example is essentially a lack of uniform integrability. Some additional regularity is needed to ensure the validity of the interchange of limit and expectation in

$$\lim_{y\to x}f(y)=\lim_{y\to x}Ef(y,\xi)=E\lim_{y\to x}f(y,\xi)=Ef(x,\xi)=f(x).$$

Suppose that ξ takes values in some set H. Proposition 1 below gives sufficient conditions for the interchange. Weaker sufficient conditions are known, but it is not clear that those weaker conditions are broadly applicable in the simulation context. The result is standard, but we include a proof as we believe it to be instructive.

Proposition 1. *Suppose that the family of functions $\mathscr{F}=(f(\cdot,\xi):\xi\in H)$ is equicontinuous at $x\in D$, i.e., for all $\varepsilon>0$, there exists $\delta=\delta(x,\varepsilon)>0$ such that $\|y-x\|<\delta$ implies that $|f(y,\xi)-f(x,\xi)|<\varepsilon$ for all $y\in D$ and all $\xi\in H$. Then $f(\cdot)$ is continuous at x.*

Proof. Let $\varepsilon>0$ and $x\in D$ be given. Choose δ as in the statement of the proposition. Let $\xi_1,\xi_2,\ldots$ be i.i.d., distributed according to ξ. The strong law of large numbers then implies that for any particular $y\in D$,

$$f_n(y)=\frac{1}{n}\sum_{i=1}^{n}f(y,\xi_i)\to f(y)=Ef(y,\xi)$$

as $n\to\infty$ a.s. Equicontinuity of $\mathscr{F}$ implies that $|f_n(y)-f_n(x)|\le\varepsilon$ for any $n\ge1$ and any y such that $\|y-x\|\le\delta$. Hence, for such y,

$$\begin{aligned}|f(y)-f(x)|&\le|f(y)-f_n(y)|+|f_n(y)-f_n(x)|\\&\quad+|f_n(x)-f(x)|\\&\le|f(y)-f_n(y)|+\varepsilon+|f_n(x)-f(x)|.\end{aligned}$$

This holds for any $n\ge1$, so taking the limit of the right-hand side, the strong law applied at both x and y implies that $|f(y)-f(x)|\le\varepsilon$. □

The proof above can be easily modified to give the following result, yielding stronger versions of continuity for $f(\cdot)$. A function $g(\cdot)$ is uniformly continuous if, for any $\varepsilon>0$, there exists $\delta>0$ such that $|g(x)-g(y)|<\varepsilon$ whenever $\|x-y\|<\delta$. Also, a function $g(\cdot)$ is Lipschitz continuous if there exists a finite constant ℓ (the Lipschitz constant) such that $|g(x)-g(y)|\le\ell\|x-y\|$ for all x,y.

Proposition 2. *If $\mathscr{F}$ is equicontinuous at all $x\in D$, then $f(\cdot)$ is continuous on D. If $\mathscr{F}$ is* uniformly *equicontinuous (δ does not depend on x), then $f(\cdot)$ is uniformly continuous. If each*

$f(\cdot,\xi)$ is Lipschitz continuous with Lipschitz constant $L(\xi)$, and $\ell = EL(\xi) < \infty$, then $f(\cdot)$ is Lipschitz with Lipschitz constant ℓ.

A related important question is whether the function $f(\cdot)$ is differentiable or not. Again, this question can be partially answered using a suitable coupling. The following argument is essentially that given in Section 1.3 of Glasserman (1991). For simplicity, assume that $d = 1$, i.e., that x is a scalar.

Let x be an interior point of D. Suppose that $f(\cdot,\xi)$ is differentiable at x almost surely. In other words, we fix x and look at the set of ξ for which the function $f(\cdot,\xi)$ is differentiable at x. The assumption here is that this set has probability 1. Hence, with probability 1,

$$\lim_{h\to 0} \frac{f(x+h,\xi)-f(x,\xi)}{h} \tag{3}$$

exists and equals $f'(x,\xi)$ say. If the family of random variables

$$\left\{ \frac{f(x+h,\xi)-f(x,\xi)}{h} : |h| \le \delta \right\} \tag{4}$$

is uniformly integrable for some $\delta > 0$, then one can pass expectations through to the limit in (3) to obtain

$$\begin{aligned} Ef'(x,\xi) &= E\lim_{h\to 0} \frac{f(x+h,\xi)-f(x,\xi)}{h} \\ &= \lim_{h\to 0} E\frac{f(x+h,\xi)-f(x,\xi)}{h} \\ &= \lim_{h\to 0} \frac{Ef(x+h,\xi)-Ef(x,\xi)}{h} \\ &= \lim_{h\to 0} \frac{f(x+h)-f(x)}{h} \end{aligned} \tag{5}$$

thereby establishing that $f'(x)$ exists and equals $Ef'(x,\xi)$. The key then is to establish uniform integrability of the family (4) for some $\delta > 0$. The following result, proved in Section 8.5 of Dieudonné (1960) and restated on p. 15 of Glasserman (1991), is very useful in this regard.

Theorem 3 (Generalized Mean-Value Theorem)**.** *Let g be a continuous real-valued function on the closed interval $[a,b]$ that is differentiable everywhere except possibly on a set D of at most countably many points. Then, for all $x, x+h \in [a,b]$,*

$$\left| \frac{g(x+h)-g(x)}{h} \right| \le \sup_{y\in[a,b]\setminus D} |g'(y)|.$$

The idea is to apply this result to $f(\cdot,\xi)$, assuming that $f(\cdot,\xi)$ satisfies the conditions of the theorem. Here the interval $[a,b]$ is taken to be $[x-\delta, x+\delta]$, and the set D depends on ξ (i.e., is random). We then have the bound

$$\left| \frac{f(x+h,\xi)-f(x,\xi)}{h} \right| \le \sup_{y\in[x-\delta,x+\delta]\setminus D} |f'(y,\xi)|$$

for almost all ξ. Notice that the bound is observable, as it is defined on sample paths. If the bound has finite expected value, then the family (4) is uniformly integrable and the interchange (5) is valid, and so $f'(x)$ exists and is defined as $Ef'(x,\xi)$.

The interchange argument above is the fundamental idea underlying infinitesimal perturbation analysis and its variants, which involve obtaining estimates of the derivative $f'(x)$ from the same sample paths used to estimate $f(x)$ itself. See Ho and Cao (1991), Glasserman (1991), and Fu and Hu (1997) for (much) more on this idea, and Fu (2006) for a recent introduction to the general area of gradient estimation in simulation.

Example 2. Suppose that a single ambulance serves calls from a base located at the point $x = (x_1, x_2)$ in the unit square $[0,1]^2$. Calls arrive according to a homogeneous Poisson process at rate λ and are distributed randomly (and not necessarily uniformly) over the square, independent of all else. Calls that arrive when the ambulance is busy are queued, and answered in order of arrival. The ambulance travels in a straight line (i.e., in Euclidean fashion) at constant rate v from its current location to its destination. The ambulance spends a time V serving a call at the destination, independent of all else, where we assume for convenience that V is bounded above (but this is easily relaxed). We assume that treatment is complete at the scene for all calls, so patients are not transported to a hospital. This is merely to keep things simple - all of the results to follow are easily modified to allow hospital transports for some patients. After completing a call, the ambulance returns to base unless there are queued calls, in which case the ambulance travels directly to the next call. When a call arrives while the ambulance is returning to base the ambulance immediately responds to the call from its current location. The ambulance receives calls from $t = 0$ to $t = 8$, and then stops receiving new calls. Calls that are not yet complete at $t = 8$ are served by the ambulance before it returns to base at the end of the day. The ambulance starts out idle and at its base at time $t = 0$.

We are interested in selecting the base location x that minimizes the average response time of the calls received over a large number of i.i.d. days, where response time is defined as the time from when the call was received to the time when the ambulance arrives at the location of the call. Related performance measures, like the fraction of calls with response times under a threshold can be handled using a similar approach. If $R_i(x)$ is the sum of the response times on the ith day when the ambulance base lies at location x and N_i is the number of calls received on day i, then the

average response time over n days of operations is

$$\frac{R_1(x)+R_2(x)+\cdots+R_n(x)}{N_1+N_2+\cdots+N_n} \to \frac{ER_1(x)}{EN_1}$$

as $n \to \infty$ a.s. We can compute $EN_1 = 8\lambda$, so our problem reduces to minimizing $f(x) = Ef(x,\xi) = ER_1(x)$ over all possible base locations $x = (x_1, x_2)$. Here the random object ξ contains all information relating to call arrival times and locations over the interval $[0,8]$.

Using the general techniques described above we can prove the following.

Proposition 4. *The function $f(\cdot)$ is Lipschitz continuous with Lipschitz constant $8\lambda/v$.*

Proof. We give a proof sketch rather than a complete proof. Fix any sample path of the process. Let N be the number of calls received in the 8 hour period. Let $L(t)$ be the location of the ambulance at time t when the base location is x, and let $\tilde{L}(t)$ be the corresponding location of the ambulance at time t when the base location is $x+h$ for some vector h. (Notice that here we are using the CRN coupling so that the call times, locations and service times are identical for the 2 base locations.) Then $\|L(t) - \tilde{L}(t)\| \le \|h\|$ at all times $t \ge 0$. This can be shown by induction on the times at which events (call arrivals, ambulance departs from a base or call, ambulance reaches base or a call, ambulance completes service at the scene) occur because the ambulance follows the same sequence of events. As a consequence, the difference in response times for the two base locations is never more than $\|h\|/v$. (Recall that the ambulance travels at rate v.) So if R and $\tilde{R}$ are the sum of the response times when the base is located at x and $x+h$ respectively, then $|R-\tilde{R}| \le N\|h\|/v$. So then

$$\begin{aligned} |f(x+h) - f(x)| &= |ER - E\tilde{R}| \\ &\le E|R-\tilde{R}| \\ &\le 8\lambda\|h\|/v \end{aligned}$$

and the result follows. □

Example 3. Consider the same situation as in Example 2 but this time suppose that there are 2 ambulances with bases at locations $x(1)$ and $x(2)$. (This can be generalized to more than 2 ambulances but for simplicity we stick with 2 ambulances.) Let x denote the 4-dimensional random variable that gives both base locations. We would like to show that the "long-run" average response time is continuous in x. As with the 1-ambulance case, it suffices to show that ER, the expected sum of the response times of all calls received in 1 day, is continuous in x.

Unfortunately, in contrast to the 1-ambulance case, the sample paths do not change in a smooth fashion as we vary x. To see why, consider a single sample path and 2 base locations. As we perturb the base locations, the set of calls that each ambulance responds to could change: Ambulance 1 will be assigned some of the calls previously assigned to Ambulance 2 and vice versa. This reassignment means that the ambulances do not follow the same paths, and the response times can be dramatically different as a result of the perturbation of base locations. Nevertheless, it is possible to show that ER is, indeed, continuous in x, despite these sample-path discontinuities. Our proof of the following result again uses coupling. As we will see more precisely in the proof below, while the sample paths have discontinuities, they are "locally continuous" with probability 1, and this is the key observation that allows us to prove continuity.

Proposition 5. *Assume that call locations have a density with respect to Lebesgue measure. Then, the function $f(\cdot)$ is continuous in x in the two-ambulance case.*

Proof. We give a proof sketch. We adopt the same notation as in the proof of Proposition 4. We again adopt the usual CRN coupling, and then

$$\begin{aligned} |ER - E\tilde{R}| &\le E|R-\tilde{R}| \\ &= E[|R-\tilde{R}|I(B^c)] + E[|R-\tilde{R}|I(B)], \end{aligned} \quad (6)$$

where B denotes the event that the assignment of calls to ambulances is different for the original (x) and perturbed ($x+h$) base locations, and B^c is the complement of B. On B^c, the sum of the response times is Lipschitz continuous in h, as observed in the proof of Proposition 4, and so the first term in (6) is bounded by

$$E[NI(B^c)\|h\|/v] \le 8\lambda\|h\|/v.$$

Hence, the first term in (6) is Lipschitz continuous in x.

To deal with the second term in (6) we apply dominated convergence. We have

$$|R-\tilde{R}|I(B) \le R+\tilde{R}.$$

We assumed the time spent at the scene of a call is bounded. This, together with the fact that we are operating within a bounded region (the unit square), means that we can bound ER and $E\tilde{R}$ as follows. Let s be a bound on the time spent at the scene of a call. The time to travel to any call is bounded by $\sqrt{2}/v$. So the response time to the first call is bounded by $\sqrt{2}/v$, and the response time for the second call is bounded by the first call's response time and service time plus $\sqrt{2}/v$. (This worst case for response time can only arise when the first and second calls arrive simultaneously and the calls occur at diagonally opposite corners.) In general

the response time for the kth call to arrive is bounded by

$$(k-1)s + k\sqrt{2}/v \le k(s+\sqrt{2}/v).$$

Hence, both R and $\tilde{R}$ are bounded by

$$(s+\sqrt{2}/v)\sum_{k=1}^{N} k = (s+\sqrt{2}/v)N(N+1)/2.$$

Now, N is Poisson distributed and so has finite second moment, so the right-hand side of this expression multiplied by 2 gives a uniform (in x) bound on $R+\tilde{R}$ that has finite mean. Dominated convergence will then complete the proof, provided that $|R-\tilde{R}|I(B) \to 0$ as $\|h\| \to 0$ a.s.

First assume that the ambulances are never simultaneously at exactly the same place. Consider the first time τ at which a call is assigned to a different ambulance under base locations $x+h$ than under base locations x. As $\|h\| \to 0$, the ambulance locations and status at time τ using bases at $x+h$ converge to their respective values at time τ for bases at x. So the only way that the call could change its assigned ambulance in the limit as $\|h\| \to 0$ is if it was located at a point that is *exactly* the same distance from both ambulances at the time of arrival, or it arrived at the *exact* time point at which an ambulance changed status from busy to available, or available to busy, under base locations x. Both of these possibilities occur with probability 0 under our assumptions. Therefore $I(B) = 0$ for $\|h\|$ sufficiently small with probability 1 if ambulances are never simultaneously at the same place.

Unfortunately, ambulances can be at the same location, but as we explain this happens at a set of times with measure 0 unless the base locations $x(1)$ and $x(2)$ coincide. Suppose the ambulance bases are distinct, i.e., $x(1) \neq x(2)$. Calls happen one at a time in a Poisson process with probability 1, and we assumed that call locations have a density, so with probability 1, ambulances can only be in the same location when their paths cross while at least 1 of them is traveling. Our assumptions ensure that this happens at a set of times with measure 0. Hence, when the base locations do not coincide, $I(B) = 0$ for $\|h\|$ sufficiently small with probability 1.

Finally, suppose that the ambulance base locations coincide, i.e., $x(1) = x(2)$. Consider the first time τ that a call is assigned to a different ambulance under base locations $x+h$ than under base locations x. The gain in response time is of order $\|h\|$, and as $\|h\| \to 0$ the ambulance locations will essentially be identical in the sample path and its perturbed version until the next call arrives. (The ambulance *identities* will be different, but since ambulances are indistinguishable in this model, this does not matter.) By induction on the set of events on the path, we find that even though $I(B)$ does not converge to 0, the difference in the sums of response times, $R-\tilde{R}$, does converge to 0. So $|R-\tilde{R}|I(B) \to 0$ as $\|h\| \to 0$ a.s. in this case as well. □

In Proposition 4 we were able to show *Lipschitz* continuity, while in Proposition 5 we were only able to show continuity. One might hope to strengthen these results to *differentiability* of $f(\cdot)$. Indeed, the main device in the proof of Proposition 5 is the decomposition (6) onto complementary events, where the sample paths are continuous on 1 event, and discontinuous on the other. This is a fundamental idea in smoothed perturbation analysis for obtaining derivative estimates from sample paths; see p. 27 of Fu and Hu (1997). This type of argument may also prove useful for analyzing other situations where the sample paths are not continuous.

Another interpretation of the proof of Proposition 5 that formalizes the "local continuity" comment above is given in the following proposition.

Proposition 6. *Fix $x \in D$. Suppose that, with probability 1, $f(\cdot,\xi)$ is continuous at x. Suppose also that the family of random variables*

$$(f(y,\xi) - f(x,\xi) : |y-x| \le \delta)$$

is uniformly integrable for some $\delta > 0$. Then $f(\cdot) = Ef(\cdot,\xi)$ is continuous at x.

Proof. Consider any sequence $\{x_n\}$ in D with $x_n \to x$ as $n \to \infty$. Our continuity assumption then gives that $f(x_n,\xi) - f(x,\xi) \to 0$ as $n \to \infty$. The uniform integrability assumption then ensures that we can pass expectations through this limit. □

3 SAMPLE-AVERAGE APPROXIMATION

One standard technique for solving the problem (1) is the sample average approximation (SAA) approach. This method approximates the original simulation optimization problem (1) with a deterministic optimization problem in the following manner. Let N be a deterministic positive integer, and suppose that we generate an independent random sample $\xi_1, .., \xi_N$. For a fixed x, define the sample mean over $(f(x,\xi_i) : 1 \le i \le N)$ as

$$\bar{f}_N(x) = \frac{1}{N}\sum_{i=1}^{N} f(x,\xi_i).$$

The SAA problem corresponding to (1) is

$$\min_{x \in D} \bar{f}_N(x), \tag{7}$$

i.e., we minimize the *sample* average. Once the sample is fixed, $\bar{f}_N(x)$ becomes deterministic. Its values and gradient (assuming it is differentiable) can be computed for a given value of the parameter x. Consequently, the SAA problem (7) becomes a deterministic optimization problem and one can solve it using any convenient optimization algorithm. The algorithm can exploit the IPA gradients, which are exact gradients of $\bar{f}_N(x)$.

Various forms of this method have been used by different researchers. The *stochastic counterpart method* by Rubinstein and Shapiro (1993) uses likelihood ratios to obtain the approximate optimization problem. In Robinson (1996) this approach is called the *sample path method*. Healy and Schruben (1991) refer to it as *retrospective optimization*. Chen and Schmeiser (2001) developed retrospective approximation algorithms for stochastic root finding problems. For a general introduction to the SAA approach, the reader is referred to Shapiro (2003).

To solve the SAA problem (7), one can use the steepest descent method with a line search. However, when the computational budget is limited and the evaluation of the sample function is expensive, one may want to draw samples conservatively and then use a more sophisticated search algorithm. In Bastin, Cirillo, and Toint (2006), the number of samples is adaptively drawn based on the statistical inference of the simulation error, which limits the number of samples in early iterations. The convergence of the algorithm is guaranteed by a trust region technique, where the improvement of the solution is ensured by locally minimizing a quadratic model of the objective function. When a considered problem has a specific structure (for example, a two-stage stochastic linear programming problem), an efficient SAA method can be developed by using a good deterministic algorithm, which exploits the structure of the optimization problem (Shapiro 2003;Verweij et al. 2003). Polak and Royset (2008) and Pasupathy (2008) look at the question of how to progressively increase the size of the SAA sample in conjunction with a numerical optimization algorithm that converges at a specific rate, in order to obtain fast overall rates of convergence as a function of computational effort.

Generally, the SAA problem (7) is close to the original problem (1) when N is large. By the strong law of large numbers (SLLN), we have that $\bar{f}_N(x)$ converges to $f(x)$ w.p.1 as $N \to \infty$. But, the point-wise convergence does not guarantee the convergence of the SAA method. We next look at conditions under which we can prove that the SAA method will converge; that is, under which the optimal solutions of the SAA problem (7) approach the set of optimal solutions of the original problem as N grows.

First, we introduce some definitions. For every x in the domain D, we define the set $\mathscr{C}(x)$ as follows. For x in the interior of D, $\mathscr{C}(x) = \{0\}$. For x on the boundary of D, $\mathscr{C}(x)$ is the convex cone generated by the outward normals of the faces on which x lies. A *first-order critical point* x of the function f satisfies

$$-\nabla f(x) = z \text{ for some } z \in \mathscr{C}(x).$$

A first-order critical point is either a point where the gradient $\nabla f(x)$ is zero, or a point on the boundary of D where the gradient "points towards the interior of D". Let $S(f,D)$ be the set of first-order critical points of f in D. We define the distance from a point x to a set S to be

$$d(x,S) = \inf_{y \in S} \|x - y\|.$$

Let $\hat{v}_N$ and v^* denote the optimal objective values of the SAA problem (7) and the true problem (1), respectively. Throughout this section, we assume that D is a non-empty convex compact set.

Theorem 7. *Suppose that*

- *(i)* *the set $S(f,D)$ of optimal solutions of the true problem (1) is non-empty and contained in D,*
- *(ii)* *the function $f(\cdot)$ is finite-valued and continuous on D,*
- *(iii)* *$\bar{f}_N$ converges to f uniformly on D, that is*

$$\sup_{x \in D} \left| \frac{1}{N} \sum_{i=1}^{N} f(x, \xi_i) - E f(x, \xi) \right| \to 0$$

as $N \to \infty$ a.s., and

- *(iv)* *the set $S(\bar{f}_N, D)$ of optimal solutions of the SAA problem (7) is non-empty w.p.1 for large enough N and contained in D.*

Let $\hat{x}_N$ be an optimal solution of the SAA problem (7). Then $\hat{v}_N \to v^*$ and $d(\hat{x}_N, S(f,D)) \to 0$ a.s. as $N \to \infty$.

The above theorem can be proved based on epi-convergence theory. For a proof, see Proposition 6 in Shapiro (2003).

The key in establishing the above convergence results is a uniform version of the strong law of large numbers (ULLN), which is the assumption *(iii)* in the above proposition. In general, the point-wise law of large numbers (LLN) does not imply the ULLN. However, in a convex problem, the point-wise LLN ensures the ULLN holds.

Let H be a closed support of ξ.

Proposition 8. *Suppose that*

- *(i)* *for every $z \in H$, the function $f(\cdot, z)$ is convex on D,*

(ii) *the law of large numbers (LLN) holds point-wise, that is $\bar{f}_N(x)$ converges to $f(x)$ as $N \to \infty$ a.s., for any fixed $x \in x$, and*
(iii) *$f(\cdot)$ is finite valued on a neighborhood of D.*

Then $\bar{f}_N$ converges to f uniformly on D, as $N \to \infty$ a.s.

Proof. We only provide a sketch of the proof here. Under the assumptions in the above proposition, it can be shown that the sample mean $\bar{f}_N$ epi-converges to the function f. But epi-convergence implies uniform convergence in convex programming, and hence the result in the above proposition follows. For a complete proof, see Proposition 2 and Corollary 3 in Shapiro (2003). □

A large class of simulation optimization problems are convex problems (e.g., two-stage stochastic linear programs with recourse), and Proposition 8 shows that in the convex case the SAA method converges under very mild conditions. However, we often have non-convex stochastic problems in practice. For example, in the ambulance base location problem in Section 2, the sum of the response times $R(\cdot)$ can be non-convex with a certain call location distribution.

The following proposition provides relatively simple conditions for the ULLN without the convexity assumption. Indeed, one can easily show that the one-ambulance location problem in Section 2 satisfies the conditions in the proposition below. We say that $f(x,\xi)$ is *dominated* by an integrable function $h(\cdot)$ if $Eh(\xi) < \infty$ and for every $x \in D$, $|f(x,\xi)| \leq h(\xi)$ a.s.

Proposition 9. *Suppose that*

(i) *for every $z \in H$, the function $f(\cdot,z)$ is continuous on D, and*
(ii) *$f(x,\xi)$ is dominated by an integrable function.*

Then $f(\cdot)$ is finite valued and continuous on D and $\bar{f}_N$ converges to f uniformly on D, as $N \to \infty$ a.s.

Proof. Under the condition (ii), the family of path-wise functions $\{f(x,\xi) : x \in D\}$ is uniformly integrable. This, together with the condition (i), ensures the continuity of the function f. Then by establishing the uniform integrability of the family of random variables

$$\left\{ \sup_{x \in V_N} |f(x,\xi) - f(\hat{x}_N,\xi)| : N \geq 1 \right\},$$

where V_N is a properly defined neighborhood of x_N, we can obtain the ULLN on $\bar{f}_N$. For a complete proof, see Proposition 7 in Shapiro (2003). □

Theorem 7 ensures that, if $\hat{x}_N$ solves the SAA problem (7) exactly, then $\hat{x}_N$ converges to the set of optimizers of the limit function f. Moreover, if the true problem (1) has a unique optimal solution x^*, then $\hat{x}_N \to x^*$. However, in the non-convex case, the best that we can hope for from a computational point of view is that $\hat{x}_N$ is a first-order critical point for the SAA problem. The following results ensure the convergence of the first-order critical points to those of the true problem.

Theorem 10. *Suppose that D is convex, and*

(i) *for every $z \in H$, the function $f(\cdot,z)$ is continuously differentiable on a neighborhood of D, and*
(ii) *the gradient components $\frac{\partial}{\partial x_i} f(x,z)(i = 1,...,d)$ are dominated by an integrable function.*

Let $\hat{x}_N \in S(\bar{f}_N, D)$ be the set of first-order critical points of $\bar{f}_N$ on D. Then f is continuously differentiable over x and $d(\hat{x}_N, S(f,D)) \to 0$ as $N \to \infty$ a.s.

Proof. First apply Proposition 9 to the each component of the gradient and then apply Theorem 7 to the sample gradient function. □

Theorem 10 shows that $\hat{x}_N$ converges to the set of first-order critical points of f as $N \to \infty$. This does not guarantee that the sequence $\{\hat{x}_1, \hat{x}_2, \hat{x}_3, \cdots\}$ converges almost surely. In general, we cannot guarantee this because, when there are multiple critical points, the particular critical point chosen depends, among other things, on the optimization algorithm that is used.

There exists a well-developed statistical inference of estimators derived by the SAA method. That inference is incorporated into validation analysis and error bounds for obtained solutions. If the true problem has a unique optimal solution x^*, then under a set of conditions an optimal solution $\hat{x}_N$ for the SAA problem converges to x^* at a stochastic rate of $O_p(N^{-1/2})$ and the bias $E[\hat{v}_N] - v^* = E[f(\hat{x}_N)] - f(x^*)$ is of order $o(N^{-1/2})$ (Shapiro 1993; Shapiro 2003). We further discuss the assessment of the solution quality in the context of the SAA method in Section 5.

There are some difficulties in applying the SAA method when the the sample function $f(\cdot,\xi)$ is discontinuous (e.g., the two-ambulance location problem). All of the results above assume at least the continuity of the sample function, so we may not be able to apply the above convergence results directly to a discontinuous problem. Even if the SAA problem converges to the true problem uniformly, the SAA method can perform poorly in finite time. Since the sample function is discontinuous, the resulting SAA problem can be a non-convex and discontinuous optimization problem, which is generally hard to solve.

4 STOCHASTIC APPROXIMATION

Another standard method to solve the problem (1) is that of stochastic approximation (SA). This technique is analogous to the steepest descent gradient search method in deterministic optimization, except here the gradient does not have an analytic expression and must be estimated. Since the basic stochastic algorithms were introduced by Robbins and Monro (1951) and Kiefer and Wolfowitz (1952), a huge amount of work has been devoted to this area.

The general form of the SA algorithm is a recursion where an approximation x_n for the optimal solution is updated to x_{n+1} using an estimator $g_n(x_n)$ of the gradient (or a sub-gradient) of the objective function $f(\cdot)$ at x_n. For a minimization problem, the recursion is of the form

$$x_{n+1} = \Pi_D(x_n - a_n g_n(x_n)), \tag{8}$$

where Π_D denotes a projection of points outside D back into D, and $\{a_n\}$ is a sequence of positive real numbers such that

$$\sum_{n=1}^{\infty} a_n = \infty \text{ and } \sum_{n=1}^{\infty} a_n^2 < \infty. \tag{9}$$

The original Robbins-Monro algorithm is designed for single-dimensional root-finding problems. The objective function is assumed to be monotonically increasing and unknown to the experimenter, but its estimates are available. In general, an SA algorithm equipped with an unbiased estimator of the gradient is called a Robbins-Monro algorithm. If unbiased estimation of the gradient is not possible, we may appeal to finite difference (FD) schemes. The idea is to use the difference quotient of $f(\cdot,\xi)$ as an estimate of the gradient. The resulting stochastic approximation procedure is called the Kiefer-Wolfowitz algorithm (Kiefer and Wolfowitz 1952). In the Kiefer-Wolfowitz algorithm, the computational effort per iteration increases linearly in d. For example, the central FD gradient estimates require estimation of $2d$ parameter values. To deal with high-dimensional problems, Spall (2000) developed the simultaneous perturbation stochastic algorithm (SPSA). All parameter components in the gradient estimate are randomly perturbed, and only two simulations of the objective function are required.

In the presence of non-convexity, the SA algorithm may only converge to a local minimum. Theorem 11 below is an immediate specialization of [Theorem 2.1, p. 127] Kushner and Yin (2003), which gives conditions under which x_n converges to a local minimizer x^* a.s. as $n \to \infty$. The recursive structure of the algorithm provides an environment where we can apply martingale ideas to analyze the convergence of the algorithm.

Let $(\mathscr{G}_n : n \geq 0)$ be a filtration, where the initial guess x_0 is measurable with respect to $\mathscr{G}_0$ and $g_n(x_n)$ is measurable with respect to $\mathscr{G}_{n+1}$ for all n.

Theorem 11. *Let $f : \mathbb{R}^d \to \mathbb{R}$ be $\mathscr{C}^1$. Suppose that for $n \geq 0$, $x_{n+1} = \Pi_D(x_n - a_n g_n(x_n))$ with the following additional conditions.*

- *(i)* *The conditions (9) hold.*
- *(ii)* $\sup_n E\|g_n(x_n)\|^2 < \infty$.
- *(iii)* *There is a sequence of random variables $\{\beta_n\}$ such that*

$$E[g_n(x_n)|\mathscr{G}_n] = \nabla f(x_n) + \beta_n, \text{ for all } n \geq 0,$$

where

$$\sum_{i=1}^{\infty} |a_i \beta_i| < \infty \text{ w.p. 1.}$$

Then,

$$d(x_n, S(f,D)) \to 0$$

as $n \to \infty$ a.s. Moreover, suppose that $S(f,D)$ is a discrete set. Then, on almost all sample paths, x_n converges to a unique point in $S(f,D)$ as $n \to \infty$.

The limiting points in $S(f,D)$ can be random. If there exists a unique optimal solution, then x_n converges to that point. Assumption *(iii)* implies that the bias sequence $\{\beta_n\}$ is asymptotically negligible. In the Kiefer-Wolfowitz algorithm, β_n represents FD bias. In the classical SA methods, the function $g_n(\cdot)$ does not depend on n. In this case, the sufficient condition for the condition *(ii)* is that $\sup_{x \in D} E||g_n(x)||^2 < \infty$.

Consider the one-ambulance location problem, and assume that the function $f(\cdot)$ is twice continuously differentiable. Since it is not possible to directly compute the gradient of the sum of response times $R(\cdot)$, we can consider a Kiefer-Wolfowitz type algorithm to find the optimal base location. Assumption *(ii)* in the theorem above can be easily verified by using the Lipschitz continuity of $R(\cdot)$. From the second order Taylor series approximation of f, the bias of the FD gradient estimator is of order $O(c_n)$, where $\{c_n\}$ is the sequence of the difference parameters in the FD estimator. Therefore, Assumption *(iii)* is satisfied as long as the difference parameter c_n goes to zero fast enough with n. For example, if $a_n = O(n^{-1})$, $c_n = O(n^{-\gamma})$ with $\gamma > 0$ satisfies the condition.

The stochastic approximation procedure is easy to implement and exhibits good performance with appropriately chosen step sizes. In the classical Robbins-Monro SA algorithm, the sequence $\{a_n\}$ is usually chosen to be of the form $a_n = a/n$ for all n, where a is a positive scalar. With the optimal value of a, the distance from x_n to the set of optimizers is $O_p(n^{-1/2})$. The optimal choice of a depends on the Hessian matrix of $f(\cdot)$ at the optimal solution, which is usually unknown. The algorithm can perform poorly with a bad choice of a, but unfortunately directly determining

the required Hessian matrix is often difficult or impossible in practice.

Various adaptive Newton-type procedures have been developed in which the Hessian and the gradient are adaptively updated as the number of iterations grows. In Ruppert (1985), a direct estimate of the gradient is available and the Hessian is estimated using finite differences of the gradient estimates. At each iteration, this approach uses $O(d)$ evaluations of the gradient. In Spall (2000), the Hessian and gradients are estimated based on the simultaneous perturbation methodology. In high-dimensional problems, such simultaneous changes admit an efficient implementation by greatly reducing the number of objective function evaluations or gradient evaluations. At each iteration, the Hessian estimates are projected on the set of positive definite and symmetric matrices so that the algorithm converges to a minimum. Bhatnagar (2005) developed three-time scale SPSA-type algorithms. At each iteration, the algorithms update the Hessian and the gradients simultaneously, using no more than four samples at each iteration. The resulting Hessian estimates can be easily modified to make them symmetric matrices.

The concept of iterate averaging provides a great improvement in the theory of the SA method. Ruppert (1991) and Polyak and Juditsky (1992) showed that for a Robbins-Monro type algorithm, the averaged sequence $\bar{x}_n = \sum_{i=1}^{n} x_i/n$ is more robust than the final value x_n. Under certain conditions, notably that the step-size a_n decrease to 0 at a slower rate than n^{-1}, the average $\bar{x}_n$ achieves the optimal asymptotic convergence rate without needing to know or estimate the Hessian matrix. Dippon and Renz (1997) applied the Polyak-Ruppert averaging idea to Kiefer-Wolfowitz type algorithms. The iterate averaging scheme has been widely used in the context of the SA method, and its advantages have been observed in many empirical results. On the other hand, it has been reported that the asymptotic promise of iterate averaging can be difficult to realize in practical finite-sample problems (Spall 2000).

Juditsky et al. (2007) considered SA methods based on a weighted-averaging scheme for a class of convex stochastic optimization problems, and compared them with the SAA method. Convexity of the function $f(\cdot)$ allows one to improve the estimates for x^* by taking a weighted average of the trajectory $\{x_1, x_2, ..., x_n\}$. Although the asymptotic convergence rate of the algorithm is not optimal when the objective function is strongly convex, the algorithm is more robust than the the classical SA algorithm. The numerical experiments demonstrate that the method can perform well even with a simple constant step size policy, and can significantly outperform an implementation of the SAA method.

5 SOLUTION QUALITY

Assessing the quality of the obtained solution is important in optimization algorithms, and there are several approaches for testing solution quality in simulation optimization. One approach uses the limiting distribution of the solution.

For example, in the Robbins-Monro algorithm, under certain conditions, $\bar{x}_n$ is asymptotically normally distributed, i.e.,

$$\sqrt{n}(\bar{x}_n - x^*) \Rightarrow N(0, \Sigma(x^*)) \text{ as } n \to \infty,$$

where x^* is an optimal solution. In theory, this allows us to construct confidence regions for x^*. However, the covariance matrix of the limiting normal distribution $\Sigma(x^*)$ depends on the Hessian of the objective function at x^*, which is challenging to estimate. One way to produce confidence regions without having to explicitly estimate the covariance matrix of the limiting normal distribution is through multiple replications (Hsieh and Glynn 2002). The idea is similar to the batch means method; first simulate m independent replications of the estimator $\bar{x}_n$, and then use them to estimate x^* and the covariance matrix $\Sigma(x^*)$.

Another approach to determining the solution quality is via confidence intervals on the optimality gap,

$$\mu = Ef(\hat{x}, \xi) - v^*,$$

where $\hat{x}$ is a candidate solution and $v^* = f(x^*)$ for some optimal solution x^*. The analysis of the optimality gap is well studied in the context of the SAA method (Norkin, Pflug, and Ruszczyński 1998, Mak, Morton, and Wood 1999). By interchanging minimization and expectation, we can obtain a lower bound on v^*,

$$E\hat{v}_N = E\left[\min_{x \in D} \bar{f}_N(x)\right] \leq \min_{x \in D} E\left[\bar{f}_N(x)\right] = v^*.$$

It can also be shown that $E\hat{v}_N \leq E\hat{v}_{N+1}$ for all N. Thus, the lower bound $\hat{v}_N$ provides a better estimate for v^* on average as the sample size grows. On the other hand, we have $v^* \leq Ef(\hat{x}, \xi)$ from the sub-optimality of $\hat{x}$. Then $Ef(\hat{x}, \xi) - E\hat{v}_N$ is an upper bound for the optimality gap for $\hat{x}$. We can estimate this quantity by

$$\hat{\mu}_N = \bar{f}_N(\hat{x}) - \min_{x \in D} \bar{f}_N(x). \tag{10}$$

Under the conditions in Theorem 7, $\hat{\mu}_N$ converges to μ w.p.1. In general, we cannot expect the asymptotic normality of the estimator $\hat{\mu}_N$ because of the minimization in the right-hand side of the equation (10). However, again by using the multiple replications idea, we can construct a confidence interval on the optimality gap μ. One drawback of the multiple replication procedure is that it needs to solve the

SAA problem at least 30 times to obtain a statistically valid confidence interval. Bayraksan and Morton (2007) developed single- and double-replication procedures, which require less computational effort but still produce confidence intervals that attain a desirable level of coverage.

ACKNOWLEDGMENTS

We thank Sam Ehrlichman for very helpful conversations relating to the ambulance example. This research was partially supported by National Science Foundation Grants DMI 0400287 and CMMI 0800688.

REFERENCES

Andradóttir, S. 1998. Simulation optimization. In *Handbook of Simulation: Principles, Methodology, Advances, Applications and Practice*, ed. J. Banks, 307–333. New York: Wiley.

Ankenman, B., B. L. Nelson, and J. Staum. 2008. Stochastic kriging for simulation metamodeling. Northwestern University IEMS Working paper 08-01.

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., and M. Meckesheimer. 2006. Metamodel-based simulation optimization. In *Handbook of Simulation*, ed. S. G. Henderson and B. L. Nelson, 535–574. Amsterdam: Elsevier.

Bastin, F., C. Cirillo, and P. L. Toint. 2006. Application of an adaptive Monte-Carlo algorithm to mixed logit estimation. *Transportation Research B* 40:577–593.

Bayraksan, G., and D. P. Morton. 2007. Assessing solution quality in stochastic programs. *Mathematical Programming Series B* 108:495–514.

Bhatnagar, S. 2005. Adaptive three-timescale stochastic approximation. *ACM Transactions on Modeling and Computer Simulation* 15 (1): 74–107.

Calvin, J. M. 2004. Lower bound on complexity of optimization of continuous functions. *Journal of Complexity* 20 (5): 773–795.

Chen, H., and B. W. Schmeiser. 2001. Stochastic root finding via retrospective approximation. *IIE Transactions* 33:259–275.

Dieudonné, J. A. 1960. *Foundations of modern analysis*. New York: Academic Press.

Dippon, J., and J. Renz. 1997. Weighted means in stochastic approximation for minima. *SIAM Journal on Control and Optimization* 35:1811–1827.

Fu, M. C. 2002. Optimization for simulation: theory vs. practice. *INFORMS Journal on Computing* 14:192–215.

Fu, M. C. 2006. Gradient estimation. In *Simulation*, ed. S. G. Henderson and B. L. Nelson, Handbooks in Operations Research and Management Science, 575–616. Amsterdam: Elsevier.

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.

Glover, F., J. P. Kelly, and M. Laguna. 1996. New advances and applications of combining simulation and optimization. In *Proceedings of the 1996 Winter Simulation Conference*, 144–152. Washington, DC, USA: IEEE.

Healy, K., and L. W. Schruben. 1991. Retrospective simulation response optimization. In *Proceedings of the 1991 Winter Simulation Conference*, ed. B. L. Nelson, W. D. Kelton, and G. Clark, 901–907. Piscataway NJ: IEEE.

Henderson, S. G., and B. L. Nelson. (Eds.) 2006. *Simulation*. Handbooks in Operations Research and Management Science. Amsterdam: Elsevier.

Ho, Y. C., and X. R. Cao. 1991. *Perturbation analysis of discrete-event systems*. Boston MA: Kluwer Academic.

Hong, L. J., and B. L. Nelson. 2006. Discrete optimization via simulation using COMPASS. *Operations Research* 54 (1): 115–129.

Hsieh, M., and P. W. Glynn. 2002. Confidence regions for stochastic alogorithms. In *Proceedings of the 2002 Winter Simulation Conference*, ed. E. Yucesan, C. H. Chen, J. L. Snowdon, and J. M. Charnes, 370–376. Piscataway NJ: IEEE.

Hu, J., M. C. Fu, and S. I. Marcus. 2007. A Model Reference Adaptive Search Method for Global Optimization. *Operations Research* 55 (3): 549–568.

Juditsky, A., G. Lan, A. Nemirovski, and A. Shapiro. 2007. Stochastic approximation approach to stochastic programming. *SIAM Journal on Optimization*. Submitted.

Kiefer, J., and J. Wolfowitz. 1952. Stochastic estimation of the maximum of a regression function. *Annals of Mathematical Statistics* 23:462–466.

Kim, S. 2006. Gradient-based simulation optimization. In *Proceedings of the 2006 Winter Simulation Conference*, ed. L. F. Perrone, F. P. Wieland, J. Liu, B. G. Lawson, D. M. Nicol, and R. M. Fujimoto, 159–167. Piscataway NJ: IEEE.

Kushner, H. J., and G. G. Yin. 2003. *Stochastic approximation and recursive algorithms and applications*. 2nd ed. New York: Springer-Verlag.

Law, A. M. 2007. *Simulation modeling and analysis*. 4th ed. New York: McGraw-Hill.

Lindvall, T. 1992. *Lectures on the coupling method*. New York: Wiley.

Mak, W. K., D. P. Morton, and R. K. Wood. 1999. Monte Carlo bounding techniques for determining solution quality in stochastic programs. *Operations Research Letters* 24:47–56.

Norkin, V. I., G. C. Pflug, and A. Ruszczyński. 1998. A branch and bound method for stochastic global optimization. *Mathematical Programming* 83:425–450.

Pasupathy, R. 2008. On choosing parameters in retrospective approximation algorithms for stochastic root finding and simulation-optimization. *Operations Research*. To appear.

Polak, E., and J. O. Royset. 2008. Efficient sample sizes in stochastic nonlinear programming. *Journal of Computational and Applied Mathematics* 217 (2): 301–310.

Polyak, B. T., and A. B. Juditsky. 1992. Acceleration of stochastic approximation by averaging. *SIAM Journal on Control and Optimization* 30:838–855.

Robbins, H., and S. Monro. 1951. A stochastic approximation method. *Annals of Mathematical Statistics* 22:400–407.

Robinson, S. M. 1996. Analysis of sample-path optimization. *Mathematics of Operations Research* 21:513–528.

Rubinstein, R. Y., and D. P. Kroese. 2004. *The cross-entropy method: A unified approach to combinatorial optimization, Monte-Carlo simulation and machine learning*. New York: Springer-Verlag.

Rubinstein, R. Y., and A. Shapiro. 1993. *Discrete event systems: Sensitivity analysis and stochastic optimization by the score function method*. Chichester: Wiley.

Ruppert, D. 1985. A Newton-Raphson version of the multivariate Robbins-Monro procedure. *Annals of Statistics* 13:236–245.

Ruppert, D. 1991. Stochastic approximation. In *Handbook of Sequential Analysis*, ed. B. K. Ghosh and P. K. Sen, 503–529. New York: Marcel Dekker.

Shapiro, A. 1993. Asymptotic behavior of optimal solutions in stochastic programming. *Mathematics of Operations Research* 18:829–845.

Shapiro, A. 2003. Monte Carlo sampling methods. In *Stochastic Programming*, ed. A. Ruszczynski and A. Shapiro, Handbooks in Operations Research and Management Science. Elsevier.

Spall, J. C. 2000. Adaptive stochastic approximation using a simultaneous perturbation gradient approximation. *IEEE Transactions on Automatic Control* 45:1839–1853.

Thorisson, H. 2000. *Coupling, stationarity, and regeneration*. New York, NY: Springer.

Verweij, B., S. Ahmed, A. J. Kleywegt, G. Nemhauser, and A. Shapiro. 2003. The sample average approximation method applied to stochastic routing problems: A computational study. *Computational Optimization and Applications* 24:289–333.

AUTHOR BIOGRAPHIES

SUJIN KIM is an assistant professor in the Department of Industrial and Systems Engineering at the National University of Singapore. She received her Ph.D. degree in Operations Research from Cornell University in 2006. Before she joined the National University of Singapore, she was a visiting assistant professor in the Department of Industrial Engineering at Purdue University. Her research concerns simulation methodology and stochastic simulation-based optimization, with applications in electric power and health service systems. Her e-mail address is ⟨iseks@nus.edu.sg⟩.

SHANE G. HENDERSON is an associate professor in the School of Operations Research and Information Engineering at Cornell University. He is the simulation area editor at *Operations Research*, and an associate editor for the *ACM Transactions on Modeling and Computer Simulation* and *Operations Research Letters*. He co-edited the handbook *Simulation* as part of Elsevier's series of Handbooks in Operations Research and Management Science, and also co-edited the Proceedings of the 2007 Winter Simulation Conference. He likes cats but is allergic to them. His research interests include discrete-event simulation and simulation optimization, and he has worked for some time with emergency services. His web page can be found via <http://www.orie.cornell.edu>.