

SIMULATION OPTIMIZATION: A TUTORIAL OVERVIEW AND RECENT DEVELOPMENTS IN GRADIENT-BASED METHODS

Marie Chau

Department of Mathematics
University of Maryland
College Park, MD 20742, USA

Michael C. Fu

Robert H. Smith School of Business
University of Maryland
College Park, MD 20742, USA

Huashuai Qu

Department of Mathematics
University of Maryland
College Park, MD 20742, USA

Ilya O. Ryzhov

Robert H. Smith School of Business
University of Maryland
College Park, MD 20742, USA

ABSTRACT

We provide a tutorial overview of simulation optimization methods, including statistical ranking & selection (R&S) methods such as indifference-zone procedures, optimal computing budget allocation (OCBA), and Bayesian value of information (VIP) approaches; random search methods; sample average approximation (SAA); response surface methodology (RSM); and stochastic approximation (SA). In this paper, we provide high-level descriptions of each of the approaches, as well as some comparisons of their characteristics and relative strengths; simple examples will be used to illustrate the different approaches in the talk. We then describe some recent research in two areas of simulation optimization: stochastic approximation and the use of direct stochastic gradients for simulation metamodels. We conclude with a brief discussion of available simulation optimization software.

1 INTRODUCTION

Simulation modeling is a powerful tool for capturing the complexity of real-world systems. Optimization software packages now have the ability to consider the number of variables in the millions. However, the combination of arguably the two most useful paradigms in the operations research/management science (OR/MS) toolbox — simulation and optimization — is still in its relative infancy, which makes it an extremely fertile area for research. We give a tutorial overview of the most well-known simulation optimization methods, all of which address, in some form, the stochastic optimization problem

$$\min_{x \in \Theta} f(x), \quad (1)$$

where $f(x) = E[Y(x, \xi)]$ is a performance measure, $Y(x, \xi)$ is a sample performance, ξ denotes the stochastic effects, x denotes the decision variables, and Θ is the feasible region, which can be discrete, continuous, or a mixture of both. Here we focus on expected value performance measures, although there has been much research over the last decade or so on other performance measures such as quantiles, which are commonly used for metrics such as Value-at-Risk in the financial services industry.

Optimization problems are generally classified according to the types of values the decision variables x can take, the primary dichotomy being discrete versus continuous. If the decision variables are discrete, they could be ordered or unordered. In the former case, the usual setting is the set of integers. In the latter

case, the decision variables are often referred to as “categorical” variables, e.g., in a queueing system, it could be the choice of a service discipline (such as different types of priorities versus first-come, first-served). The simplest type of categorical variable would be a binary variable, often found in mathematical programming formulations, which is used to model a “yes-no” decision. Another dichotomy involves the size of the feasible region: finite versus infinite for discrete problems, bounded versus unbounded for continuous problems. Statistical ranking & selection (R&S) procedures, where here is meant to include efficient simulation budget allocation methods such as optimal computing budget allocation (OCBA), address discrete problems where Θ is not too large. Response surface methodology and stochastic approximation generally address continuous-variable problems, although in the latter case there have been some algorithms for the discrete setting. Random search methods, robust metaheuristics, and sample average approximation (also known as sample path optimization, stochastic counterpart) are generally applicable to both domains, including in principle mixed problems. We summarize each of the main approaches — indifference-zone procedures, Bayesian value-of-information approaches, optimal computing budget allocation (OCBA), random search methods, sample average approximation (SAA), response surface methodology (RSM), and stochastic approximation (SA), with more details provided for the R&S and SA procedures. We then describe some recent research in stochastic approximation and the use of direct stochastic gradients for simulation metamodels. Due to space limitations, the treatment is necessarily at a higher descriptive level, and technical details, including theoretical proofs, can be found in the references provided.

The remainder of this paper is organized as follows. The main simulation optimization methodologies are summarized in the subsections of Section 2, where longer descriptions are provided only for the R&S and SA methods, due to space limitations; Section 2.2 is based heavily on material in Chau and Fu (2014). Section 3 describes recent advances in SA methods and the use of direct stochastic gradients for simulation metamodeling, which can be used in response surface methodology. Section 4 provides some concluding remarks, including a brief discussion of available simulation optimization software. For a far more detailed description of many of the techniques described here, the reader is encouraged to consult the two books, *Handbook of Simulation Optimization* (Fu 2014a) and *Optimal Learning* (Powell and Ryzhov 2012); see also Fu (1994) and Fu (2002).

2 SIMULATION OPTIMIZATION METHODS

2.1 Ranking and Selection

Ranking and selection (R&S) can be viewed as the most fundamental class of simulation optimization problems. This class streamlines many computational and modeling challenges: the feasible region Θ is typically taken to be a small, finite set (representing a collection of promising “alternatives,” e.g., simulation designs), and the goal is to identify $x \in \{1, \dots, M\}$ with the “best” (smallest or largest) value $\mu_x \equiv f(x)$. The problem would be trivial if we knew these values, but we suppose that they can only be estimated using unbiased, i.i.d. observations $Y(x, \xi)$, and our simulation budget limits the number of observations that can be collected (e.g. some finite number N).

The R&S model emphasizes the essential issue of *exploration vs. exploitation*: we may wish to use part of our simulation budget to learn about a solution x that appears to be suboptimal, on the chance that it may turn out to be better than we thought. At the same time, there is no point in learning about solutions that are obviously poor. Thus, the goal of simulation optimization methods is to allocate the simulation budget efficiently, by making a tradeoff between the estimated quality of an alternative and the uncertainty of the estimate. This problem presents two research challenges: first, we need a statistical model that quantifies our uncertainty about a decision, and second, we need an algorithmic procedure to make the exploration/exploitation tradeoff.

There are two types of statistical models for R&S, frequentist and Bayesian. Frequentist models assume no prior knowledge about f and construct estimates based purely on the observed simulation output. For example, if we consider an additive noise model $Y(x, \xi) = f(x) + \xi$, where ξ is a zero mean

(perhaps normally distributed) random variable, and if we further suppose that each simulation produces an independent noise, we can use the simplest possible model from classical statistics, given by

$$\bar{\mu}_x^n = \frac{1}{N_x^n} \sum_{i=1}^{N_x^n} Y(x, \xi_{xi}), \quad (\bar{\sigma}_x^n)^2 = \frac{1}{N_x^n - 1} \sum_{i=1}^{N_x^n} (Y(x, \xi_{xi}) - \bar{\mu}_x^n)^2,$$

where N_x^n is the number of times that we have simulated alternative x out of the first n simulations, and (ξ_{xi}) is a sequence of i.i.d. noise terms. The estimator $\bar{\mu}_x^n$ is a random variable whose mean is the unknown (but fixed) quantity $f(x)$. Allocating more simulations to alternative x improves the precision of $\bar{\mu}_x^n$ at the expense of other alternatives.

By contrast, the Bayesian approach views the unknown performance μ_x as a random variable whose distribution encodes our own uncertainty about the exact value. For example, the statement $\mu_x \sim \mathcal{N}(\theta_x, \sigma_x^2)$ means that we believe μ_x to be equal to θ_x , but we are uncertain about this belief to a degree represented by σ_x^2 . This belief may come from the decision-maker's domain knowledge about the specific application being modeled, or it may be constructed from a few preliminary simulations of different alternatives. Either way, by making additional assumptions on the distribution of the simulation output, we can sometimes characterize the process by which our beliefs evolve using simple recursive equations. For example, if we assume that μ_x is normal, and furthermore that the noise $\xi \sim \mathcal{N}(0, \sigma_\xi^2)$ for some known σ_ξ^2 , the conditional distribution of μ_x given $Y(x, \xi)$ is again normal with parameters

$$\theta'_x = \frac{\sigma_x^{-2} \theta_x + \sigma_\xi^{-2} Y(x, \xi)}{\sigma_x^{-2} + \sigma_\xi^{-2}}, \quad (\sigma'_x)^{-2} = \sigma_x^{-2} + \sigma_\xi^{-2}.$$

This property, whereby the initial (or *prior*) and modified (or *posterior*) belief distributions come from the same family, is known as *conjugacy*, and is attractive because it allows us to capture an entire probability distribution over the possible values of μ_x using only a small number of parameters. The book by DeGroot (1970) provides a catalogue of conjugate Bayesian models for most well-known probability distributions.

2.1.1 Indifference-zone Methods

How can we evaluate whether the simulation budget is spent efficiently? One standard algorithmic approach uses the indifference-zone (IZ) criterion. Denote by $\mu_{(j)}$ the j th smallest value, so that $\mu_{(1)} \leq \dots \leq \mu_{(M)}$. For a fixed *indifference-zone parameter* value δ , the relation

$$P\left(\mu_{(1)} = \arg \min_x \bar{\mu}_x^{N_x^n} \mid \mu_{(1)} \leq \mu_{(2)} - \delta\right) \geq 1 - \alpha \quad (2)$$

indicates that our simulation procedure is valid, i.e., we select the best (lowest-mean) alternative with high probability as long as its value is sufficiently distinct from the next-best. Indifference-zone methods allocate simulations in a way that ensures (2) is achieved. Numerous such methods have been developed, dating back to early work by Bechhofer (1954) and Paulson (1964); see Bechhofer et al. (1995), Swisher et al. (2003), Kim and Nelson (2006), Kim (2014), or the Hong and Nelson (2009) tutorial for an overview.

In this space, we briefly describe the KN procedure by Kim and Nelson (2001), which first runs n_0 preliminary simulations of every alternative and calculates

$$\bar{S}_{xy}^2 = \frac{1}{n_0 - 1} \sum_{i=1}^{n_0} (Y(x, \xi_{xi}) - Y(y, \xi_{yi}) - [\bar{\mu}_x^{n_0} - \bar{\mu}_y^{n_0}])^2,$$

the usual frequentist estimator of the difference of the true values of alternatives x and y . We then define $K_{xy} = \lfloor \frac{h^2 \bar{S}_{xy}^2}{\delta^2} \rfloor$ for some suitably chosen constant h (depending on the desired performance level α), and then calculate $K_x = \max_{y \neq x} K_{xy}$. The quantity $K = \max_x K_x$ represents the maximum number of simulations needed to achieve the desired performance; if $n_0 > K$, no further simulation is needed. Otherwise, we use the remainder of the simulation budget to simulate alternatives whose values appear to be close to the best.

In the n th step of the procedure, we simulate from alternative x as long as, for all alternatives y that were simulated in the previous step, we have

$$\bar{\mu}_x^n \leq \bar{\mu}_y^n + \max \left\{ 0, \frac{\delta}{2cn} \left(\frac{h^2 \bar{S}_{xy}^2}{\delta^2} - n \right) \right\},$$

i.e., our estimate of μ_x is still sufficiently close to our estimates of μ_y to prevent us from ruling x out as the best. In this case, we stop when the procedure tells us to simulate only a single alternative. The parameter c should be suitably chosen to maintain theoretical performance (Kim and Nelson 2001).

The KN procedure possesses several characteristics common to many IZ methods. We begin by taking a first-stage sample in which every alternative receives equally many simulations. After that, however, we focus more on alternatives that show high potential to be good, and screen out alternatives that are unpromising. The procedure is said to be “fully sequential,” since the conditions for screening out alternatives in the n th iteration depend on the estimates at that time, and thus change as we run more simulations.

2.1.2 Bayesian Value of Information

Value of information procedures (VIPs) are based on Bayesian statistical models, and thus require the decision-maker to specify a prior distribution of belief on each performance value μ_x . Under normality assumptions, this amounts to providing θ_x^0 and σ_x^0 for each x ; these can be constructed by running a first-stage sample (e.g., simulating each alternative n_0 times), as in an IZ method. Once these parameters are given, however, conjugacy provides us with a way to update our beliefs very quickly, leading to sequences (θ_x^n, σ_x^n) of belief parameters characterizing our estimates and uncertainty at the n th stage of sampling. VIPs leverage the uncertainty in the belief distribution to quantify the *potential* of a new simulation to improve the quality of our solution with respect to some performance measure.

One very straightforward performance measure is simply $\min_x \theta_x^N$, the estimated smallest value among the alternatives after the simulation budget has been consumed. This metric is based on the intuitive idea that, if we have already used all of our allotted simulations, we should simply report $\arg \min_x \theta_x^N$ as being the best alternative. While this decision may not actually be the best, we are not likely to do better than this if there is no longer any way to learn. Below, we describe a VIP for this objective under independent normal belief distributions; however, Chick (2006), Powell and Ryzhov (2012), or the WSC tutorial Frazier (2012) can be seen for a much more detailed technical overview.

Essentially, a VIP first calculates the expected *improvement* in the desired performance measure that can be achieved by running one more simulation of alternative x . In this respect, the procedure is analogous to a steepest ascent algorithm (and is sometimes referred to by the name “knowledge gradient” for this reason). We assume the basic normal model with independent observations, and further simplify it by assuming that only one alternative can be simulated in each stage of the procedure (in contrast with the KN algorithm presented earlier). In this setting, the expected improvement quantity can be written as

$$v_x^n = E \left[\min_y \theta_y^n - \min_y \theta_y^{n+1} \mid x \text{ is simulated in the } n\text{th stage} \right]. \quad (3)$$

Note that (3) is a conditional expectation. First, we are implicitly conditioning on our beliefs at the n th stage; that is, θ_x^n and σ_x^n will be known at the time the calculation is made, and so $\min_y \theta_y^n$ is non-random from the point of view of this expectation. Second, we are also conditioning on the next decision. The quantity v_x^n assumes that we have committed to simulating x , but have not yet observed the outcome. The expectation in (3) is thus taken over the conditional distribution of $\min_y \theta_y^{n+1}$ given the n th-stage beliefs. Fortunately, this distribution is actually quite simple, owing to the fact that x is the only alternative simulated in the n th stage. As a consequence, we can write

$$\min_y \theta_y^{n+1} = \min \left\{ \min_{y' \neq y} \theta_{y'}^n, \theta_x^n + \tilde{\sigma}_x^n Z \right\},$$

where Z is a standard normal random variable and $(\tilde{\sigma}_x^n)^2 = (\sigma_x^n)^2 - (\sigma_x^{n+1})^2$ is the uncertainty reduction achieved by simulating x . We omit the derivation of this fact; for our purposes, it is enough to note that the randomness of $\min_y \theta_y^{n+1}$ originates from a single scalar random variable, making (3) into an expected value of a function of Z . The work by Gupta and Miescke (1996) gives the closed-form solution of (3) as

$$v_x^n = \tilde{\sigma}_x^n g \left(- \left| \frac{\theta_x^n - \min_{y \neq x} \theta_y^n}{\tilde{\sigma}_x^n} \right| \right),$$

where $g(z) = z\Phi(z) + \phi(z)$ and ϕ and Φ are the respective standard normal pdf and cdf. This quantity can be computed quickly for any x , leading to a simple decision rule: simulate $\arg \max_x v_x^n$ at each time step.

The VIP considers the estimated value of x relative to the others (represented by the “best-of-the-rest” value $\min_{y \neq x} \theta_y^n$), but scales the gap between these values by the uncertainty reduction σ_x^n . Potentially, we could measure an alternative that appears to be quite suboptimal (significantly higher than the estimated best), as long as σ_x^n was high enough. On the other hand, it is not worth simulating even a very good alternative with very low σ_x^n ; essentially, this means that we are already highly confident of its performance and have no need for further information about it.

2.1.3 Optimal Computing Budget Allocation (OCBA)

The OCBA methodology has gained attention in the literature due to its flexibility, as it can be adapted to both frequentist and Bayesian models, as well as both IZ-style objectives (related to the probability that the estimated-best alternative really is the best) and the “economic” objectives used by VIP. To focus the discussion, suppose that our objective is to maximize the probability of correct selection

$$P \left(\mu_{(1)} = \arg \min_x \bar{\mu}_x^{k_x} \right) \quad (4)$$

subject to the constraint $k_1 + \dots + k_M = N$, using frequentist estimates of the parameter values. We can view this problem as that of finding an optimal deterministic policy, in which the computing budget is pre-allocated once, before any simulations are made, but in a way that optimizes the final performance metric. This view is generally used to derive OCBA procedures; however, once the optimal k_x values have been calculated, the procedure can easily be made fully sequential by randomizing the simulation decision in each time stage (so alternative x is simulated w.p. $\frac{k_x}{k_1 + \dots + k_M}$). If the optimal allocation is difficult to compute, we can substitute an approximation k_1^n, \dots, k_M^n , making the procedure adaptive.

We briefly describe one OCBA rule for the objective in (4) stated in Chen, Fu, and Shi (2008); interested readers are referred to the monograph by Chen and Lee (2010) for more details. At the n th stage of sampling, let $x^n = \arg \min_x \bar{\mu}_x^n$ be the current-best alternative. The approximation to the optimal allocation k maximizing (4) can be expressed via the ratio

$$\frac{k_x^n}{k_y^n} = \left(\frac{\tilde{\sigma}_x^n \bar{\mu}_{x^n}^n - \bar{\mu}_x^n}{\tilde{\sigma}_y^n \bar{\mu}_{x^n}^n - \bar{\mu}_x^n} \right)^2, \quad x \neq y \neq x^n; \quad k_{x^n}^n = \tilde{\sigma}_{x^n}^n \sqrt{\sum_{x \neq x^n} \left(\frac{k_x^n}{\tilde{\sigma}_x^n} \right)^2}.$$

As $n \rightarrow \infty$, the estimators of the means and variances will converge to their true values, and the estimated allocations will likewise converge to the true optimal k_x values. See Glynn and Juneja (2004) for a proof of optimality for these allocations using large deviations theory.

2.1.4 Comparison of Approaches

The three surveyed classes of R&S procedures are based on very different ideas, and each type has numerous generalizations and extensions. A study by Branke, Chick, and Schmidt (2007) conducted an extensive empirical comparison of these methods, although each area has experienced new developments since the time of the study. Below, we give a summary of the main conceptual and empirical strengths of each class:

- **Indifference-zone methods** retain their computational and theoretical advantages in a wide variety of settings, e.g., common random numbers, problems with constraints, and subset selection procedures. IZ methods sometimes “over-deliver” or produce better performance than what is actually required, as was observed empirically by Branke, Chick, and Schmidt (2007); Frazier (2014) proposes a Bayesian-inspired method to correct this behavior.
- **Value of information procedures** can be extended to handle common random numbers unknown sampling noise, and non-normal sampling distributions, but are mainly attractive for their modeling flexibility. The VIP approach is more powerful when used in conjunction with Bayesian models with “correlated beliefs,” meaning that a simulation of a single alternative can provide information about other, “similar” alternatives (Frazier, Powell, and Dayanik 2009). The computational advantages of these methods are heavily tied to the use of a conjugate Bayesian model, though recent work by Qu, Ryzhov, and Fu (2012) suggests ways to approximate conjugate models when none is available.
- **Optimal computing budget allocation** methods have the ability to work with both frequentist and Bayesian models and can handle common random numbers. There is substantial work on improving the quality of the approximations used in the OCBA calculations (Chen et al. 2000, Chen, He, and Fu 2006). Recent work has considered constrained problems (Pujowidianto et al. 2009) as well as multi-objective simulation optimization (Lee et al. 2010). The large deviations characterization of OCBA (Glynn and Juneja 2004) has also led to new computational work on allocations that maximize theoretical rates of convergence (Hunter et al. 2011), which are generally very difficult to show for other R&S procedures.

2.2 Stochastic Approximation

Stochastic approximation (SA) is a recursive algorithm that can be viewed as the stochastic counterpart to steepest descent in deterministic optimization. SA was introduced by Robbins and Monro (1951) to solve noisy root-finding problems and was later applied to the setting of stochastic optimization by solving for the zero of the gradient. The gradient-free setting was addressed by Kiefer and Wolfowitz (1952). SA is an online method that requires very little memory and is currently one of the most widely applicable and most useful methods for simulation optimization.

Consider the stochastic optimization problem (1), where $\Theta \subseteq \mathbb{R}^d$ is a continuous parameter space. In this case, the objective is to find a sequence $\{x_n\}$ that converges to a unique (local) optimum

$$x^* = \arg \min_{x \in \Theta} f(x), \quad (5)$$

by using the iteration

$$x_{n+1} = \Pi_{\Theta} \left(x_n - a_n \widehat{\nabla} f(x_n) \right), \quad (6)$$

where $\Pi_{\Theta}(x)$ is a projection of x back into the feasible region Θ if $x \notin \Theta$, $a_n > 0$ is the step size or gain size, and $\widehat{\nabla} f(x_n)$ is an estimate of the gradient $\nabla f(x_n)$. The projection operator is only required in the constrained setting. Moreover, the minimization problem in (1) and (5) could easily be changed to maximization by changing the sign of a_n in (6). The two classical methods, Robbins-Monro (RM) and Kiefer-Wolfowitz (KW), estimate $\nabla f(x_n)$ using unbiased direct gradient estimates and finite difference gradient estimates, respectively. Under certain conditions, RM and KW have the respective asymptotic convergence rates $O(n^{-1/2})$ and $O(n^{-1/3})$. Although, RM is superior to KW in terms of asymptotic convergence rate, it is not always applicable in practice since direct gradients are not always available. In the simulation context, direct gradient estimation techniques include perturbation analysis, the likelihood ratio/score function method, and weak derivatives; see Fu (2006) and Fu (2014b) for more details and references.

Now, we describe the two classical methods, Robbins-Monro (RM) and Kiefer-Wolfowitz (KW), followed by three of the most useful modifications/variants: simultaneous perturbation stochastic approximation (SPSA), Kesten’s rule, and iterate averaging. Recent SA advances will be discussed in Section 3.

2.2.1 The Robbins-Monro (RM) Algorithm

The RM algorithm was introduced to solve the root-finding problem $M(x) = \alpha$ for $x \in \mathbb{R}$, where $M(x)$ is a monotone function and $\alpha \in \mathbb{R}$. However, it was later applied to a specific case of root-finding in the stochastic optimization setting, where the objective is to optimize a stochastic objective function $f(x)$ by setting $M(x) = \nabla f(x)$ and $\alpha = 0$. RM solves this problem iteratively as in (6) by replacing $\widehat{\nabla}f(x_n)$ with an unbiased estimator, and the output is taken as the last iterate, x_N^* , where N is the stopping time. However in RM, the direct gradient measurements are still approximations to the actual gradient because of the presence of noise ($\widehat{\nabla}f(x_n) = \nabla f(x_n) + \varepsilon_n$, where ε_n is noise with zero mean). Given the appropriate parameters, this algorithm converges asymptotically at a rate of $O(n^{-1/2})$.

The most well-known conditions are restrictions on the gain sequence $\{a_n\}$. Generally, $a_n \rightarrow 0$ but $\sum_{n=1}^{\infty} a_n = \infty$, which prevents the step size from converging to zero too quickly, so the iterates are able to make progress to x^* and do not get stuck at a poor estimate. The usual form is $a_n = \frac{\theta_a}{(n+A)^\alpha}$, where $\theta_a > 0$, $A \geq 0$ and $\frac{1}{2} < \alpha \leq 1$, with $A = 0$ and $\alpha = 1$ as a commonly used choice. The objective function f is assumed to have a global minimum with a bounded derivative.

2.2.2 The Kiefer-Wolfowitz (KW) Algorithm

The KW stochastic approximation algorithm is referred to as a gradient-free or stochastic zeroth-order method in the following chapter, since it only requires noisy measurements of the function and does not require additional information on the system dynamics or input distributions. The original KW iterative scheme

$$x_{n+1} = x_n - a_n \frac{Y(x_n + c_n, \xi_n^+) - Y(x_n - c_n, \xi_n^-)}{2c_n}, \quad (7)$$

estimates the gradient using a symmetric finite difference gradient estimate, and under certain conditions, KW can achieve an asymptotic convergence rate of $O(n^{-1/3})$. In addition, common random numbers (CRN) can be employed to decrease the variance of estimates, and KW can achieve an asymptotic convergence rate of $O(n^{-1/2})$ in certain settings.

Although the KW algorithm converges asymptotically, its finite-time performance is dependent on the choice of tuning sequences, $\{a_n\}$ and $\{c_n\}$. If the current x_n is in a relatively flat region of the function and the a_n is small, then the convergence will be slow. On the other hand, if the x_n is located in a very steep region of the function and $\{a_n\}$ is large, then the iterates will experience a long oscillation period. If $\{c_n\}$ is too small, the gradient estimates using finite differences could be extremely noisy.

2.2.3 Simultaneous Perturbation Stochastic Approximation (SPSA)

Simultaneous perturbation stochastic approximation (SPSA) specifically addresses multivariate optimization problems (Spall 1992). Similar to KW-type algorithms, SPSA only requires the objective function values to approximate the underlying gradient and is therefore easy to implement. However, SPSA only requires two function evaluations at each iteration regardless of the dimension of the parameter space Θ , which could potentially reduce the computational cost significantly in high-dimensional problems. SPSA perturbs the vector x randomly in all directions simultaneously (hence, the name of the method) and the i th component of the gradient estimate has the form

$$\widehat{f}_i(x_n) = \frac{Y(x_n + c_n \Delta_n, \xi_n^+) - Y(x_n - c_n \Delta_n, \xi_n^-)}{2c_n \Delta_{n,i}}, \quad (8)$$

where $\Delta_n = (\Delta_{n,1}, \dots, \Delta_{n,d}) \in \mathbb{R}^d$ and generally assumed to be i.i.d. and independent across components, $c_n \in \mathbb{R}^+$ is the finite difference step size, and ξ_n^\pm denotes the randomness. Observe that the numerator in (8) involves two function estimates and is identical for all i ; therefore, the cost of the full gradient (aside from generating Δ_n) is independent of dimension.

The optimal convergence rate for SPSA is $O(n^{-1/3})$. The perturbation sequence $\{\Delta_n\}$, where $\Delta_n = (\Delta_{n,1}, \dots, \Delta_{n,d})$ with $\{\Delta_{n,i}\}$ independent must have mean zero (i.e., $E[\Delta_n] = 0$), and finite inverse moments (i.e., $E[|\Delta_{n,i}|^{-1}] < \infty$ for $i = 1, \dots, d$). As a result, the Gaussian distribution is not applicable. Instead, the most common distribution used is the symmetric Bernoulli taking a positive and negative value (i.e., ± 1) with probability 0.5. In addition, an appropriately scaled x_n is approximately normal for large n , and the relative efficiency of SPSA depends on the geometric shape of $f(x)$, choice of $\{a_n\}$ and $\{c_n\}$, distribution of $\{\Delta_{n,i}\}$, and noise level.

2.2.4 Kesten’s Rule

It is well-known that the classical SA algorithms are extremely sensitive to the step-size sequence $\{a_n\}$. Therefore, it could be advantageous to consider adaptive step sizes that adjust based on the ongoing performance of the algorithm, in hopes of adapting them to the characteristics of the function at the current location of the iterate and proximity of the current iterate to the optimum. Kesten’s rule (Kesten 1958) decreases the step size only when there is a directional change in the iterates. The notion behind this adaptive step size is that, if the iterates continue in the same direction, there is reason to believe they are approaching the optimum and the pace should not be decreased in order to accelerate the convergence. If the errors in the estimate values change signs, it is an indication that either the step size is too large and the iterates are experiencing long oscillation periods or the iterates are in the vicinity of the true optimum; either way, the step size should be reduced to a more appropriate step size or to hone in on x^* . Kesten’s rule can be applied to both RM and KW and still guarantee convergence in probability, as long as $\{a_n\}$ satisfies certain conditions.

2.2.5 Iterate Averaging

Iterate averaging approaches SA from a different angle. Instead of fine-tuning the step sizes to adapt to the function characteristics, iterate averaging takes bigger steps (i.e., a_n larger than $O(n^{-1})$) for the estimates to oscillate around the optimum, so the average of the iterates will result in a good approximation to the true optimum. The idea is simple, and yet can be very effective. It is easy to see that for this method to be successful, it is essential for the iterates to surround the optimum in a balanced manner, and the domain for which the iterates oscillate shrinks as n increases. Averaging trajectories reduces the sensitivity to the initial step size choice. The algorithm follows recursion (6) for the RM case; however, instead of taking the last iterate x_N as the output, the optimum is estimated by $x_N^* = \frac{1}{N} \sum_{n=1}^N x_n$, which is an average of N iterates, where N is the stopping time. Under “classic” assumptions, iterate averaging achieves the same convergence rate as the RM method. Furthermore, $\sqrt{n}(x_n^* - x^*)$ is asymptotically normal with mean zero and the smallest covariance matrix, which is the inverse of the average Fisher information matrix (Polyak and Juditsky 1992). A constant step size can be applied and yields convergence in distribution (Kushner and Yin 2003). A variation of this method is called the “sliding window” average, which is based on the last m iterates. In general, averaging iterates leads to more robustness with respect to step-size sequence because of the reduced sensitivity, while converging at the same optimal asymptotic rate as RM.

2.3 Random Search Methods

Random search is covered in four different chapters in the simulation optimization handbook (Fu 2014a). Chapters 10 (Andradóttir 2014) and 11 (Zabinsky 2014) are dedicated to the topic on a more general level, focusing primarily on the unconstrained continuous setting, whereas Chapters 2 (Hong, Nelson, and Xu 2014) and 12 (Hu 2014) treat some specialized settings. Specifically Chapter 2 presents some locally and globally convergent algorithms for discrete optimization on a large countable (possibly infinite) set, and Chapter 12 presents iterative model-based algorithms for which a probability distribution over the solution space is updated based on the performance of solutions sampled from the probability distribution. The following is a generic random search algorithm adapted from Andradóttir (2014), where at iteration n , M_n is the number of designs sampled, \mathcal{S}_n is the sampling strategy used to select the M_n candidate solutions,

and $N_n^{(i)}$ is the number of simulations used to estimate the i th candidate performance.

Step 0 (initialize): Choose the initial sampling strategy \mathcal{S}_1 and let $n = 1$.

Step 1 (sample): Select $x_n^{(1)}, \dots, x_n^{(M_n)} \in \Theta$ according to the sampling strategy \mathcal{S}_n .

Step 2 (simulate): Estimate $f(x_n^{(i)})$, for $i = 1, \dots, M_n$, via simulated $Y(x_n^{(i)}, \xi_{n,j}^{(i)}), j = 1, \dots, N_n^{(i)}$.

Step 3 (update): Use the simulation results obtained in Step 2 to compute an estimate of the optimal solution x_n^* and to choose an updated sampling strategy \mathcal{S}_{n+1} . Let $n = n + 1$ and go to Step 1.

Generally, $f(x_n^{(i)})$ is estimated using the sample mean. Random search methods differ according to the sampling strategies $\{\mathcal{S}_n\}$ used, which can be updated adaptively or chosen in advance; M_n can also be adaptive. “The sampling strategy can be used both to identify new, promising system designs, and also to obtain improved objective function estimates for previously sampled designs (by re-sampling them)” (Andradóttir 2014). See Nelson (2010) and Hong, Nelson, and Xu (2014) for the discrete setting, including the Convergent Optimization via Most-Promising-Area Stochastic Search (COMPASS) algorithm introduced by Hong and Nelson (2006), which is the basis of an open-source solver described in the concluding remarks.

2.4 Sample Average Approximation

Sample average approximation (SAA) tackles (1) by converting it into a deterministic problem to be solved using an appropriate numerical procedure. First, select and fix $\xi_1, \xi_2, \dots, \xi_n$, where ξ_i follows the same distribution as ξ for all i . Then the problem becomes

$$\min_{x \in \Theta} f_n(x), \quad \text{where } f_n(x) = \frac{1}{n} \sum_{i=1}^n Y(x, \xi_i). \quad (9)$$

SAA is conceptually straightforward and simple to implement. According to Kim, Pasupathy, and Henderson (2014), SAA is most effective when f is continuous, which is the same requirement for infinitesimal perturbation analysis (IPA). Two decisions must be made at the start of SAA: sample size n and a method used to solve the generated sample path problem (9). The sample size can be chosen to ensure an ε -optimal solution for problem (1), i.e., if x_ε is an ε -optimal solution of (1) if $f(x_\varepsilon) \leq \theta^* + \varepsilon$, where θ^* is the true optimal objective value for (1). Suppose that the feasible set Θ is finite, then for $\varepsilon > 0, \varepsilon' \in [0, \varepsilon), \alpha \in (0, 1)$, and for sample size n satisfying $n \geq \frac{2\sigma^2}{(\varepsilon - \varepsilon')^2} \ln\left(\frac{|\Theta|}{\alpha}\right)$, then the ε' -optimal solution of the SAA problem (9) is an ε -optimal solution of the true solution of (1) with probability $1 - \alpha$, where σ^2 is the variance of $Y(x, \xi)$. In the literature, asymptotic convergence rates for SAA refers to the convergence of the objective function $f(x_n)$ to $f(x^*)$, in contrast to the classical results that focus on asymptotic convergence properties of the estimated optimum x^* . Under certain conditions, SAA can achieve an asymptotic convergence rate of $n^{-1/2}$, along with asymptotic normality; see also Shapiro (2013). In general, Kim, Pasupathy, and Henderson (2014) assert that the implementation of SAA is inferior to that of SA, mainly because it uses a fixed sample size regardless of the progression or lack thereof; however, a recent modification of SAA called retrospective approximation (RA) described there (including reference to the original papers) circumvents this issue and proposes to generate and solve a sequence of sample-path problems as opposed to solving a single sample path problem. Also, for problems where $Y(x, \xi)$ has a closed form but its expectation does not, SAA has been very effective.

2.5 Response Surface Methodology

Introduced by Box and Wilson (1951) in the same year as the two stochastic approximation algorithms, response surface methodology (RSM) is an iterative search procedure that combines statistical design of experiments, regression analysis, linear search, and deterministic optimization. After the preliminary transformation (to stabilize variance), scaling (to standardize across dimensions), and screening (to reduce dimensionality) of the input variables (called factors in the experimental design terminology), there are

two main phases to RSM. In Phase I, which is the iterative portion of RSM, a linear regression model is generally used to estimate a search direction to explore. Once a relatively flat area is found, RSM proceeds to Phase II where a higher-order — usually quadratic — model is fitted, which is generally used to estimate the optimum. The following is a generic RSM procedure adapted from Barton (2013); see Barton and Meckesheimer (2006) and Kleijnen (2014) for more detailed discussion.

1. **Scale** x variables and **transform** Y values.
2. **Screen** for important components of the x vector.
3. **Phase I:**
 - a. Select an experimental design appropriate for fitting a first-order regression model.
 - b. Conduct the experiment, fit the appropriate model, assess significance and fit.
 - c. If the first-order model is satisfactory, then identify a search direction and conduct a sequence of experiments in this direction, and return to step 3a; else proceed to Phase II.
4. **Phase II:**
 - a. Select an experiment design appropriate for fitting a second-order regression model (could be an augmentation of an existing first-order design).
 - b. Conduct the experiment, fit the appropriate model, assess significance and fit.
 - c. If the second-order model is satisfactory, then identify whether an optimum, saddle point or a ridge has been identified, and return the optimum settings or a ridge of near-optimum settings in the original units of x ; else identify a search direction and conduct a sequence of experiments in this direction, and return to step 4a.

2.6 Robust Metaheuristics

Robust metaheuristics comprise techniques that have found success in deterministic optimization, primarily focusing on combinatorial optimization problems. Into this category we include such well-known approaches as tabu search, scatter search, simulated annealing, genetic algorithms, as well as the nested partitions method, as categorized by Ólafsson (2006), where more details on metaheuristics can be found. Most of the commercial software packages, e.g., OptQuest and SimRunner, rely on some combination of these approaches, possibly with some random search, where the details of any specific algorithm are proprietary.

3 RECENT DEVELOPMENTS

We describe some recent research aimed at improving SA and RSM methods for simulation optimization, made possible by the availability of direct gradient estimates in the simulation setting.

3.1 Secant-Tangents AveRaged Stochastic Approximation (STAR-SA)

The Secant-Tangents AveRaged (STAR) stochastic approximation algorithm (Chau, Qu, and Fu 2014, Chau, Fu, and Qu 2014) estimates the gradient using a hybrid estimator, which is a convex combination of a symmetric finite difference and an average of two direct gradient estimators:

$$\widehat{\nabla}f(x_n) = \alpha_n \frac{Y(x_n + c_n, \varepsilon_n^+) - Y(x_n - c_n, \varepsilon_n^-)}{2c_n} + (1 - \alpha_n) \left(\frac{Y'(x_n + c_n, \delta_n^+) + Y'(x_n - c_n, \delta_n^-)}{2} \right), \quad (10)$$

where ε_n^\pm and δ_n^\pm denote the randomness (i.e., $f(x_n \pm c_n) = E[Y(x_n \pm c_n, \varepsilon_n^\pm)]$ and $f'(x_n \pm c_n) = E[Y'(x_n \pm c_n, \xi_n^\pm)]$), $\alpha_n \in [0, 1]$ for all n , $c_n \rightarrow 0$ and $\alpha_n \rightarrow 0$ as $n \rightarrow \infty$. The STAR gradient estimate requires function and gradient estimates on two points, $x_n \pm c_n$ for each $\widehat{\nabla}f(x_n)$. In a setting where direct gradients are available, if the direct gradient is very noisy relative to the function estimates, it is difficult to decide between implementing RM or KW, even though RM converges faster asymptotically. Since the performance of neither algorithm is always superior to the other, the STAR gradient incorporates both. The weights of the convex combination play a critical role in the performance of STAR-SA and can be chosen to minimize

the variance of the gradient estimate such that it is less than the variance of both the symmetric finite difference gradient estimate and direct gradient estimate. The extension to higher dimensions using SPSA is described in Chau, Fu, and Qu (2014).

3.2 Use of Direct Gradients for Enhanced Metamodels

We summarize some recent work in using direct gradients to enhance the accuracy of simulation metamodels that can be used in simulation optimization, e.g., in sequential RSM. Technical details can be found in Fu and Qu (2014), Qu and Fu (2012), Qu and Fu (2014), Chen, Ankenman, and Nelson (2013).

3.2.1 Direct Gradient Augmented Regression (DiGAR)

As described in Section 2.5, the central component of RSM is the fitting of the local response surface using a metamodel, and the most common procedure used is regression. Direct Gradient Augmented Regression (DiGAR) is an approach that augments traditional regression to provide a better metamodel. In the simulation optimization setting, simulation replications are often computationally expensive, so it is desirable to use as few of them as possible for each value of the input variable. When applying RSM for sequential search, the direction of improved performance is perhaps the most critical output of the fitted metamodel in Phase I. In several numerical experiments reported in Fu and Qu (2014) for an $M/M/1$ queue with relatively small number of simulation replications per design point, the slope of the standard linear regression model often gave the wrong sign, whereas the DiGAR model always estimated the sign correctly. Thus, where applicable, DiGAR should provide a better metamodel for simulation optimization using RSM.

3.2.2 Augmented Stochastic Kriging

Stochastic kriging is another multivariate metamodeling technique that was introduced by Ankenman, Nelson, and Staum (2010). Without providing the technical details (i.e., not even the mathematical model), we summarize two different ways that stochastic kriging can be augmented using direct gradient estimates:

- Stochastic Kriging with Gradients (SKG), introduced by Chen, Ankenman, and Nelson (2013), parallels DiGAR, in that it models the added gradient analogously with a separate term. Under certain conditions, it can be shown that the integrated mean-squared error is reduced, and experimental results indicate that infinitesimal perturbation analysis (IPA) generally outperforms the likelihood ratio/score function method when incorporated with SKG.
- Gradient Extrapolated Stochastic Kriging (GESK) uses the gradients in a totally different manner by generating new output data. Rather than modeling the gradient directly as in DiGAR and SKG, it *extrapolates* in the neighborhood of the original design points, i.e., additional response data are generated via linear extrapolation. Different extrapolation techniques can be applied, and multiple points can also be added around each original design point. The GESK model requires the choice of step sizes for the extrapolated points; large step sizes allow better coverage but at the cost of additional bias since the linearity is less likely to hold further from the original point. Thus, there is a basic bias-variance tradeoff to consider. In Qu and Fu (2014), this tradeoff is analyzed in a simplified setting, leading to conditions under which improvement can be guaranteed.

4 CLOSING REMARKS

Simulation optimization is one of the most exciting research areas, due to its mathematical foundations and practical applications. It offers opportunities and challenges for researchers in many fields, from OR/MS to mathematics to statistics to computer science to economics to most engineering fields. For some current thoughts on the state of the art in research and practice, see Fu et al. (2014). Other areas not discussed

here include stochastic constraints (Homem-de-Mello and Bayraksan 2014) and Markov decision processes (Gosavi 2014), which add dynamics (time) explicitly to the optimization problem given by (1), where the optimal solution is given by a *policy* rather than a static set of values.

Simulation optimization software includes integrated commercial products such as ProModel's SimRunner and WITNESS Optimizer, the commercial standalone solver OptQuest, and the open-source software Industrial Strength COMPASS (ISC) available for free download at <http://www.iscompass.net> (Xu, Nelson, and Hong 2010). The first three handle both continuous and discrete problems, whereas ISC is designed for discrete problems. According to Nelson (2010), "The ISC algorithm divides the optimization process into three phases: a global phase that explores the entire feasible region and identifies several areas that possibly contain competitive locally optimal solutions, a local phase that searches for a locally optimal solution in each of the regions identified in the global phase, and a cleanup phase that selects the best solution among all solutions identified as locally optimal and estimates its true expected performance to within a user-provided precision. The ISC uses a niching genetic algorithm for the global phase, the COMPASS algorithm for the local phase, and an R&S procedure for the cleanup phase. It also defines meaningful and testable transition rules between the phases."

REFERENCES

- Andradóttir, S. 2014. "A Review of Random Search Methods". In *Handbook of Simulation Optimization*, edited by M. C. Fu, Chapter 10. Springer.
- Ankenman, B., B. L. Nelson, and J. Staum. 2010. "Stochastic Kriging for Simulation Metamodeling". *Operations Research* 58 (2): 371–382.
- Barton, R. R. 2013. "Response Surface Methodology". In *Encyclopedia of Operations Research and Management Science* (3rd ed.), edited by S. I. Gass and M. Fu, 1307–1313. Springer.
- Barton, R. R., and M. Meckesheimer. 2006. "Metamodel-Based Simulation Optimization". In *Handbooks in Operations Research and Management Science: Simulation*, edited by S. G. Henderson and B. L. Nelson, Chapter 18, 535–574. Elsevier.
- Bechhofer, R. 1954. "A Single-Sample Multiple Decision Procedure for Ranking Means of Normal Populations with known Variances". *The Annals of Mathematical Statistics* 25 (1): 16–39.
- Bechhofer, R. E., T. J. Santner, and D. M. Goldsman. 1995. *Design and Analysis of Experiments for Statistical Selection, Screening and Multiple Comparisons*. New York: J.Wiley & Sons.
- Box, G. E. P., and K. B. Wilson. 1951. "On the Experimental Attainment of Optimum Conditions". *Journal of the Royal Statistical Society. Series B (Methodological)* 13 (1): 1–45.
- Branke, J., S. E. Chick, and C. Schmidt. 2007. "Selecting a Selection Procedure". *Management Science* 53 (12): 1916–1932.
- Chau, M., M. Fu, and H. Qu. 2014. "A New Multidimensional Hybrid Stochastic Approximation Algorithm". *Working paper, University of Maryland, College Park*.
- Chau, M., and M. C. Fu. 2014. "An Overview of Stochastic Approximation". In *Handbook of Simulation Optimization*, edited by M. C. Fu, Chapter 6. Springer.
- Chau, M., H. Qu, and M. Fu. 2014. "A New Hybrid Stochastic Approximation Algorithm". *Proceedings of the 12th International Workshop on Discrete Event Systems*: 241–246.
- Chen, C.-H., M. C. Fu, and L. Shi. 2008. "Simulation and Optimization". In *TutORials in Operations Research*, edited by Z.-L. Chen and S. Raghavan. Springer.
- Chen, C.-H., D. He, and M. C. Fu. 2006. "Efficient Dynamic Simulation Allocation in Ordinal Optimization". *IEEE Transactions on Automatic Control* 51 (12): 2005–2009.

- Chen, C.-H., and L. H. Lee. 2010. *Stochastic Simulation Optimization: An Optimal Computing Budget Allocation*. World Scientific.
- Chen, C.-H., J. Lin, E. Yücesan, and S. E. Chick. 2000. "Simulation Budget Allocation for Further Enhancing the Efficiency of Ordinal Optimization". *Discrete Event Dynamic Systems* 10 (3): 251–270.
- Chen, X., B. Ankenman, and B. L. Nelson. 2013. "Enhancing Stochastic Kriging Metamodels with Gradient Estimators". *Operations Research* 61 (2): 512–528.
- Chick, S. E. 2006. "Subjective Probability and Bayesian Methodology". In *Handbooks of Operations Research and Management Science, vol. 13: Simulation*, edited by S. G. Henderson and B. L. Nelson, 225–258. North-Holland Publishing, Amsterdam.
- DeGroot, M. H. 1970. *Optimal Statistical Decisions*. John Wiley and Sons.
- Frazier, P. I. 2012. "Optimization via Simulation with Bayesian Statistics and Dynamic Programming". In *Proceedings of the 2012 Winter Simulation Conference*, edited by C. Laroque, J. Himmelspach, R. Pasupathy, O. Rose, and A. M. Uhrmacher, 1–16. Piscataway NJ: IEEE, Inc.
- Frazier, P. I. 2014. "A Fully Sequential Elimination Procedure for Indifference-Zone Ranking and Selection with Tight Bounds on Probability of Correct Selection". *Operations Research* 62 (4): 926–942.
- Frazier, P. I., W. B. Powell, and S. Dayanik. 2009. "The Knowledge-Gradient Policy for Correlated Normal Rewards". *INFORMS Journal on Computing* 21 (4): 599–613.
- Fu, M. C. 1994. "Optimization via Simulation: A Review". *Annals of Operations Research* 53: 199–248.
- Fu, M. C. 2002. "Optimization for Simulation: Theory vs. Practice (Feature Article)". *INFORMS Journal on Computing* 14 (3): 192–215.
- Fu, M. C. 2006. "Gradient Estimation". In *Handbooks in Operations Research and Management Science: Simulation*, edited by S. G. Henderson and B. L. Nelson, Chapter 19, 575–616. Elsevier.
- Fu, M. C. 2014a. *Handbook of Simulation Optimization*. Springer.
- Fu, M. C. 2014b. "Stochastic Gradient Estimation". In *Handbook of Simulation Optimization*, edited by M. C. Fu, Chapter 5. Springer.
- Fu, M. C., G. Bayraksan, S. G. Henderson, B. L. Nelson, I. O. Ryzhov, and B. Thengvall. 2014. "Simulation Optimization: A Panel on the State of the Art in Research and Practice". In *Proceedings of the 2014 Winter Simulation Conference*, edited by A. Tolk, S. D. Diallo, I. O. Ryzhov, L. Yilmaz, S. Buckley, and J. A. Miller, 478–489. Piscataway NJ: IEEE, Inc.
- Fu, M. C., and H. Qu. 2014. "Regression Models Augmented with Direct Stochastic Gradient Estimators". *INFORMS Journal on Computing* 26 (3): 484–499.
- Glynn, P., and S. Juneja. 2004. "A large deviations perspective on ordinal optimization". In *Proceedings of the 2004 Winter Simulation Conference*, edited by R. Ingalls, M. D. Rossetti, J. S. Smith, and B. A. Peters, 577–585. Piscataway NJ: IEEE, Inc.
- Gosavi, A. 2014. "Solving Markov Decision Processes via Simulation". In *Handbook of Simulation Optimization*, edited by M. C. Fu, Chapter 13. Springer.
- Gupta, S. S., and K. J. Miescke. 1996. "Bayesian look ahead one-stage sampling allocations for selection of the best population". *Journal of Statistical Planning and Inference* 54 (2): 229–244.
- Homem-de-Mello, T., and G. Bayraksan. 2014. "Stochastic Constraints and Variance Reduction Techniques". In *Handbook of Simulation Optimization*, edited by M. C. Fu, Chapter 9. Springer.
- Hong, L. J., and B. L. Nelson. 2006. "Discrete Optimization via Simulation Using COMPASS". *Operations Research* 54:115–129.
- Hong, L. J., and B. L. Nelson. 2009. "A Brief Introduction to Optimization via Simulation". In *Proceedings of the 2009 Winter Simulation Conference*, edited by M. Rosetti, R. Hill, B. Johansson, A. Dunkin, and R. Ingalls, 75–85. Piscataway NJ: IEEE, Inc.

- Hong, L. J., B. L. Nelson, and J. Xu. 2014. “Discrete Optimization via Simulation”. In *Handbook of Simulation Optimization*, edited by M. C. Fu, Chapter 2. Springer.
- Hu, J. 2014. “Model-Based Stochastic Search Methods”. In *Handbook of Simulation Optimization*, edited by M. C. Fu, Chapter 12. Springer.
- Hunter, S. R., C.-H. Chen, R. Pasupathy, N. A. Pujowidianto, L. H. Lee, and C. P. Yap. 2011. “Optimal Sampling Laws for Constrained Simulation Optimization on Finite Sets: The Bivariate Normal Case”. In *Proceedings of the 2011 Winter Simulation Conference*, edited by S. Jain, R. R. Creasey, J. Himmelspach, K. P. White, and M. C. Fu, 4294–4302. Piscataway NJ: IEEE, Inc.
- Kesten, H. 1958. “Accelerated Stochastic Approximation”. *The Annals of Mathematical Statistics* 29 (1): 41–59.
- Kiefer, K., and J. Wolfowitz. 1952. “Stochastic Estimation of the Maximum of a Regression Function”. *The Annals of Mathematical Statistics* 23 (3): 462–466.
- Kim, S., R. Pasupathy, and S. G. Henderson. 2014. “A Guide to Sample Average Approximation”. In *Handbook of Simulation Optimization*, edited by M. C. Fu, Chapter 8. Springer.
- Kim, S.-H. 2014. “Statistical Ranking and Selection”. In *Encyclopedia of Operations Research and Management Science*, edited by S. I. Gass and M. C. Fu. Springer.
- Kim, S.-H., and B. L. Nelson. 2001. “A Fully Sequential Procedure for Indifference-zone Selection in Simulation”. *ACM Transactions on Modeling and Computer Simulation* 11 (3): 251–273.
- Kim, S.-H., and B. L. Nelson. 2006. “Selecting the Best System”. In *Handbooks of Operations Research and Management Science, vol. 13: Simulation*, edited by S. G. Henderson and B. L. Nelson, 501–534. North-Holland Publishing, Amsterdam.
- Kleijnen, J. P. C. 2014. “Response Surface Methodology”. In *Handbook of Simulation Optimization*, edited by M. C. Fu, Chapter 4. Springer.
- Kushner, H., and G. Yin. 2003. *Stochastic Approximation and Recursive Algorithms and Applications*. New York, NY: Springer.
- Lee, L. H., E. P. Chew, S. Teng, and D. Goldsman. 2010. “Finding the Non-dominated Pareto Set for Multi-objective Simulation Models”. *IIE Transactions* 42 (9): 656–674.
- Nelson, B. L. 2010. “Optimization via Simulation Over Discrete Decision Variables”. In *TutORials in Operations Research*, edited by J. J. Hasenbein, Chapter 7, 193–207. Hanover, MD: INFORMS.
- Ólafsson, S. 2006. “Metaheuristics”. In *Handbooks in Operations Research and Management Science: Simulation*, edited by S. G. Henderson and B. L. Nelson, Chapter 21, 633–654. Elsevier.
- Paulson, E. 1964. “A Sequential Procedure for Selecting the Population with the Largest Mean from k Normal Populations”. *The Annals of Mathematical Statistics* 35 (1): 174–180.
- Polyak, B., and A. Juditsky. 1992. “Acceleration of Stochastic Approximation by Averaging”. *SIAM Journal on Control and Optimization* 30 (4): 838–855.
- Powell, W. B., and I. O. Ryzhov. 2012. *Optimal Learning*. John Wiley and Sons.
- Pujowidianto, N. A., L. H. Lee, C.-H. Chen, and C. M. Yap. 2009. “Optimal Computing Budget Allocation for Constrained Optimization”. In *Proceedings of the 2009 Winter Simulation Conference*, edited by M. D. Rosetti, R. R. Hill, B. Johansson, A. Dunkin, and R. G. Ingalls, 584–589. Piscataway NJ: IEEE, Inc.
- Qu, H., and M. Fu. 2012. “On Direct Gradient Enhanced Simulation Metamodels”. In *Proceedings of the 2012 Winter Simulation Conference*, 478–489. Piscataway NJ: IEEE, Inc.
- Qu, H., and M. C. Fu. 2014. “Gradient Extrapolated Stochastic Kriging”. *ACM Transactions on Modeling and Computer Simulation*, forthcoming.

- Qu, H., I. O. Ryzhov, and M. C. Fu. 2012. “Ranking and Selection with Unknown Correlation Structures”. In *Proceedings of the 2012 Winter Simulation Conference*, edited by C. Laroque, J. Himmelspach, R. Pasupathy, O. Rose, and A. M. Uhrmacher, 144–155. Piscataway NJ: IEEE, Inc.
- Robbins, H., and S. Monro. 1951. “A Stochastic Approximation Method”. *The Annals of Mathematical Statistics* 22:400–407.
- Shapiro, A. 2013. “Sample Average Approximation”. In *Encyclopedia of Operations Research and Management Science* (3rd ed.), edited by S. I. Gass and M. C. Fu, 1350–1355. Springer.
- Spall, J. 1992. “Multivariate Stochastic Approximation Using a Simultaneous Perturbation Gradient Approximation”. *IEEE Transactions on Automatic Control* 37 (3): 332–341.
- Swisher, J. R., S. H. Jacobson, and E. Yücesan. 2003. “Discrete-event Simulation Optimization Using Ranking, Selection, and Multiple Comparison Procedures: A Survey”. *ACM Transactions on Modeling and Computer Simulation* 13 (2): 134–154.
- Xu, J., B. L. Nelson, and L. J. Hong. 2010. “Industrial Strength COMPASS: A Comprehensive Algorithm and Software for Optimization via Simulation”. *ACM Transactions on Modeling and Computer Simulation* 20 (3): 1–29.
- Zabinsky, Z. B. 2014. “Stochastic Adaptive Search Methods: Theory and Implementation”. In *Handbook of Simulation Optimization*, edited by M. C. Fu, Chapter 11. Springer.

AUTHOR BIOGRAPHIES

MARIE CHAU is a Ph.D. candidate in Applied Mathematics, Statistics, and Scientific Computation at the University of Maryland. She has a Master’s in applied math and degrees in math, finance and economics from UMCP. Her research interests lie in stochastic optimization and financial engineering. Her email address is mchau@math.umd.edu.

MICHAEL C. FU is Ralph J. Tyser Professor of Management Science in the Robert H. Smith School of Business, with a joint appointment in the Institute for Systems Research and affiliate faculty appointment in the Department of Electrical and Computer Engineering, all at the University of Maryland. His research interests include simulation optimization and applied probability, with applications in supply chain management and financial engineering. He has a Ph.D. in applied math from Harvard and degrees in math and EECS from MIT. He served as WSC 2011 Program Chair, NSF Operations Research Program Director, *Management Science* Stochastic Models and Simulation Department Editor, and *Operations Research* Simulation Area Editor. He is a Fellow of **INFORMS** and **IEEE**. His email address is mfu@umd.edu.

HUASHUAI QU is a Ph.D. candidate in Applied Mathematics, Statistics, and Scientific Computation at the University of Maryland, where he received a Gold Medal in Teaching Award from the Mathematics Department in 2010. His research interests lie in the broad areas of optimal learning and simulation optimization. He received the 2012 INFORMS Computing Society Student Paper Award and the Best Theoretical Paper Award at WSC 2012. His email address is huashuai@math.umd.edu.

ILYA O. RYZHOV is an Assistant Professor in the Robert H. Smith School of Business at the University of Maryland. He received a Ph.D. in Operations Research and Financial Engineering from Princeton University. His research deals with optimal learning and the broader area of stochastic optimization, with applications in disaster relief, energy, and revenue management. He was a recipient of WSC’s Best Theoretical Paper Award in 2012. His work has appeared in *Operations Research*, and he is a co-author of the book *Optimal Learning*, published in 2012 by John Wiley & Sons. His email address is iryzhov@rhsmith.umd.edu.