

SIMULATION-BASED OPTIMIZATION OVER DISCRETE SETS WITH NOISY CONSTRAINTS

Yao Luo
Eunji Lim

Industrial Engineering
University of Miami
Coral Gables, FL 33124 USA

ABSTRACT

We consider a constrained optimization problem over a discrete set where noise-corrupted observations of the objective and constraints are available. The problem is challenging because the feasibility of a solution cannot be known for certain, due to the noisy measurements of the constraints. To tackle this issue, we propose a new method that converts constrained optimization into the unconstrained optimization problem of finding a saddle point of the Lagrangian. The method applies stochastic approximation to the Lagrangian in search of the saddle point. The proposed method is shown to converge, under suitable conditions, to the optimal solution almost surely (a.s.) as the number of iterations grows. We present the effectiveness of the proposed method numerically in two settings: (1) inventory control in a periodic review system, and (2) staffing in a call center.

1 INTRODUCTION

We consider the simulation-based optimization problem of the form:

$$\begin{aligned} \min_{\theta \in C \cap \mathbb{Z}^d} \quad & f^0(\theta) \\ \text{s/t} \quad & f^i(\theta) \leq 0, \quad 1 \leq i \leq r, \end{aligned} \tag{1}$$

where C is a nonempty closed convex subset of \mathbb{R}^d , and $f^i : \mathbb{Z}^d \rightarrow \mathbb{R}$ ($0 \leq i \leq r$) has no analytic form and thus must be computed only through simulation at each θ in \mathbb{Z}^d . Such functions f^i often arise in the setting of a complex stochastic system where one performance measure is described by f^0 and the other performance measures are denoted by f^i ($1 \leq i \leq r$). One such example is an inventory system with discrete demand processes, where the goal is to minimize the average ordering and holding costs per unit time while achieving a prescribed level of customer service.

The problem of minimizing an unconstrained function over a discrete set has received a considerable amount of attention from the research community, and a number of methods are proposed in the literature; see Goldsman and Nelson (1994), Andradóttir (1995), Yan and Mukai (1992), Ho, Sreenivas, and Vakili (1992), Shi and Olafsson (2000), Hong and Nelson (2006), Kleywegt, Shapiro, and Homem-de-mello (2001), Gelfand and Mitter (1989), Glover (1989), and Liepins and Hilliard (1989). For more comprehensive survey, see Nelson (2010), Henderson and Nelson (2006), and Fu (2002).

However, when the simulation-based optimization problem has stochastic constraints, a limited number of methods are available in the literature. In the presence of one stochastic constraint, Andradóttir et al. (2005) and Andradóttir and Kim (2010) propose a two-phase procedure where the first phase identifies all feasible solutions or near-feasible solutions with a pre-specified probability of correct identification, and the second phase solves the problem with the solutions identified in the first phase. Batur and Kim (2001) apply a ranking and selection procedure to identify a feasible solution while Pujowidianto et al. (2009)

address how to allocate computer time in an optimal way among the solutions in order to maximize the probability of correctly identifying the optimal solution. These methods have a requirement that all the solutions be simulated at least once, so they are more appropriate to the setting where the domain of f^0 is finite and contains a small number of elements.

In the presence of multiple constraints, the idea of replacing a constrained optimization problem with an unconstrained one by adding a penalty function to the objective function has been investigated; Li et al. (2009) combine a penalty function type method with a random search scheme, and Whitney et al. (2001) incorporate a penalty function type method into a gradient-based search scheme. However, the convergence of these methods is not guaranteed or is based on restrictive assumptions that are difficult to verify.

Another way of transforming a constrained problem to an unconstrained one is to use the Lagrangian function. The idea of using the Lagrangian function has already been adopted when the decision variables are continuous (see Kushner and Sanvicente (1975) and p. 177 of Kushner and Clark (1978) for example), but this idea has never been explored in the setting of discrete decision variables. This paper explores this idea in the discrete setting and studies the effectiveness of this approach. One of our motivations is that in the deterministic optimization problem, the Lagrangian method has certain advantages over the penalty function type methods because most penalty function type methods suffer from numerical instabilities as the controlling parameter becomes too large or too small, see Murray (1967) for example.

Our proposed approach, therefore, observes that the minimizer of the constrained problem (1) can be found by finding the saddle point of the corresponding Lagrangian $L : \mathbb{Z}^d \times \mathbb{R}_+^r \rightarrow \mathbb{R}$ defined by $L(\theta, \lambda) = f^0(\theta) + \sum_{i=1}^r \lambda^i f^i(\theta)$ for $\theta \in \mathbb{Z}^d$ and $\lambda = (\lambda^1, \dots, \lambda^r) \in \mathbb{R}_+^r$; i.e., if L has a saddle point (θ_*, λ_*) satisfying $L(\theta_*, \lambda) \leq L(\theta_*, \lambda_*) \leq L(\theta, \lambda_*)$ for all $\theta \in C \cap \mathbb{Z}^d$ and $\lambda \in \mathbb{R}_+^r$, then θ_* is a minimizer of f^0 subject to the constraints $f^i \leq 0$ for $1 \leq i \leq r$. From this observation, we reformulate (1) as

$$\max_{\lambda \in \mathbb{R}_+^r} \min_{\theta \in C \cap \mathbb{Z}^d} L(\theta, \lambda) \quad (2)$$

and propose a gradient-based method to solve (2). In particular, we wish to apply stochastic approximation to L (and hence update θ in the steepest descent direction and λ in the steepest ascent direction in each iteration) in search of the saddle point of L . Stochastic approximation requires gradient estimates of L in each iteration, but θ is integer-valued; therefore, the gradient of L with respect to θ cannot be defined in a traditional way. To overcome this obstacle, we extend L from $\mathbb{Z}^d \times \mathbb{Z}_+^r$ to $\mathbb{R}^d \times \mathbb{R}_+^r$ by extending f^i ($0 \leq i \leq r$) from a discrete domain to a continuous domain and using the extended f^i 's to construct the extended L , and compute the gradient of the extended L in the usual way. The gradient of the extended L is then used in each iteration of stochastic approximation in search of the solution to (2). The main theorem of this paper (Theorem 1) proves that this procedure is convergent to the optimal solution of the original problem (1) a.s. under suitable conditions.

We illustrate the effectiveness of our proposed method numerically in two settings: (1) inventory control in a periodically-reviewed single-item inventory system, and (2) staffing in a call center that handles multiple types of calls while maintaining a satisfactory level of customer service. The proposed method performs displays good performance when compared with alternative approaches.

The main advantages of the proposed method can be summarized as follows: (1) it is designed to handle stochastic constraints when the number of feasible solutions is large or infinite, (2) it is shown to be convergent to the optimal solution a.s. under certain technical conditions, and (3) it shows promising numerical performance in various examples.

This paper is organized as follows. In Section 2, we introduce some definitions. Section 3 describes our proposed method formally and states the main theorem (Theorem 1) of this paper. Numerical results are presented in Section 4.

2 DEFINITIONS

In this section, we introduce some definitions that will be used throughout this paper. For a positive integer m , \mathbb{Z}^m , \mathbb{R}^m , and \mathbb{R}_+^m denote the set of m -dimensional integer vectors, the set of m -dimensional real vectors, and the set of m -dimensional nonnegative real vectors, respectively. We view vectors as columns and write x^T to denote the transpose of a vector $x \in \mathbb{R}^m$. For $x \in \mathbb{R}^m$, we write its j th component as x^j , so $x = (x^1, \dots, x^m)$. By $\|x\|$, we denote $((x^1)^2 + \dots + (x^m)^2)^{1/2}$. For a subset I of $\{1, \dots, m\}$, χ_I is an m -dimensional vector whose j th entry is 1 when j belongs to I and 0 otherwise ($1 \leq j \leq m$). We denote $\chi_{\{j\}}$ by e_j ($1 \leq j \leq m$).

For $x \in \mathbb{R}$, $\lfloor x \rfloor$ and $\lceil x \rceil$ denote the largest integer less than or equal to x and the smallest integer greater than or equal to x , respectively. For $x = (x^1, \dots, x^m) \in \mathbb{R}^m$, $\lfloor x \rfloor$ and $\lceil x \rceil$ denote $(\lfloor x^1 \rfloor, \dots, \lfloor x^m \rfloor)$ and $(\lceil x^1 \rceil, \dots, \lceil x^m \rceil)$, respectively. We denote the closest integer point to $x \in \mathbb{R}^m$ by $[x]$.

For a function $g : \mathbb{R}^m \rightarrow \mathbb{R}$, a vector $\xi \in \mathbb{R}^m$ is said to be a subgradient of g at $x \in \mathbb{R}^m$ if $g(y) \geq g(x) + \xi^T(y - x)$ for all $y \in \mathbb{R}^m$. If g is convex, then a subgradient of g at x exists at every $x \in \mathbb{R}^m$ (see Theorem 23.4 in p. 217 of Rockafellar (1970)).

For a function $g : \mathbb{R}^m \rightarrow \mathbb{R}$, the partial derivative of g with respect to the j th component at $x \in \mathbb{R}^m$ is denoted by $\partial g(x) / \partial x^j$ for $1 \leq j \leq m$ if the partial derivative exists.

For $x \in \mathbb{R}$, $\max(0, x)$ is x if $x \geq 0$, and zero otherwise.

3 PROBLEM FORMULATION

3.1 General Approach

We consider the following problem

$$\begin{aligned} \min_{\theta \in C \cap \mathbb{Z}^d} \quad & f^0(\theta) \\ \text{s/t} \quad & f^i(\theta) \leq 0, \quad 1 \leq i \leq r, \end{aligned} \tag{3}$$

where C is a nonempty closed convex subset of \mathbb{R}^d and we can observe $f^i : \mathbb{Z}^d \rightarrow \mathbb{R}$ ($0 \leq i \leq r$) via simulation at each $\theta \in \mathbb{Z}^d$.

To convert (3) into an unconstrained problem, we consider the Lagrangian $L : \mathbb{Z}^d \times \mathbb{R}_+^r \rightarrow \mathbb{R}$ defined as follows:

$$L(\theta, \lambda) = f^0(\theta) + \sum_{i=1}^r \lambda^i f^i(\theta)$$

for $\theta \in \mathbb{Z}^d$ and $\lambda = (\lambda^1, \dots, \lambda^r) \in \mathbb{R}_+^r$. We observe that if $\theta_* \in C \cap \mathbb{Z}^d$ and $\lambda_* = (\lambda_*^1, \dots, \lambda_*^r) \in \mathbb{R}_+^r$ satisfy

- i) θ_* minimizes $L(\theta, \lambda_*)$ over $\theta \in C \cap \mathbb{Z}^d$, and
- ii) For $1 \leq i \leq r$, $\lambda_*^i > 0$ implies $f^i(\theta_*) = 0$ and $\lambda_*^i = 0$ implies $f^i(\theta_*) \leq 0$,

then θ_* is an optimal solution to (3). To see why this is true, we note that i) and ii) imply that $\theta_* \in \mathcal{F}$, where $\mathcal{F} \triangleq \{\theta \in C \cap \mathbb{Z}^d : f^i(\theta) \leq 0 \text{ for } 1 \leq i \leq r\}$. In addition, for any $\theta \in \mathcal{F}$, we have $f^0(\theta_*) + \sum_{i=1}^r \lambda_*^i f^i(\theta_*) \leq f^0(\theta) + \sum_{i=1}^r \lambda_*^i f^i(\theta)$, and hence, $f^0(\theta_*) \leq f^0(\theta) + \sum_{i=1}^r \lambda_*^i (f^i(\theta) - f^i(\theta_*)) \leq f^0(\theta)$.

Furthermore, it can be easily seen that i) and ii) are equivalent to the condition

$$L(\theta_*, \lambda) \leq L(\theta_*, \lambda_*) \leq L(\theta, \lambda_*)$$

for all $\theta \in C \cap \mathbb{Z}^d$ and $\lambda \in \mathbb{R}_+^r$; i.e., (θ_*, λ_*) is a saddle-point of L .

Therefore, it is reasonable to attempt to find the optimal solution to (3) by solving

$$\max_{\lambda \in \mathbb{R}_+^r} \min_{\theta \in C \cap \mathbb{Z}^d} L(\theta, \lambda). \tag{4}$$

In our proposed method, we search for the solution to (4) by updating θ and λ iteratively using a gradient-based method. One obstacle to this approach is that L has a discrete input variable θ , so it is impossible to define the gradient of L with respect to θ in a traditional way. To overcome this, we extend L from $\mathbb{Z}^d \times \mathbb{R}_+^r$ to $\mathbb{R}^d \times \mathbb{R}_+^r$ and compute the gradient of the extended L in the usual way. In order to extend L , we extend f^i ($0 \leq i \leq r$) from \mathbb{Z}^d to \mathbb{R}^d and use the extended functions to extend L . In particular, we denote the extension of f^i by \hat{f}^i ($\hat{f}^i : \mathbb{R}^d \rightarrow \mathbb{R}$) for $0 \leq i \leq r$ and define $\hat{L} : \mathbb{R}^d \times \mathbb{R}_+^r \rightarrow \mathbb{R}$ by

$$\hat{L}(\theta, \lambda) = \hat{f}^0(\theta) + \sum_{i=1}^r \lambda^i \hat{f}^i(\theta)$$

for $\theta \in \mathbb{R}^d$ and $\lambda = (\lambda^1, \dots, \lambda^r) \in \mathbb{R}_+^r$.

Our proposed method then solves

$$\max_{\lambda \in \mathbb{R}_+^r} \min_{\theta \in C \cap \mathbb{R}^d} \hat{L}(\theta, \lambda) \tag{5}$$

in the hope of solving (4). We observe that if $(\hat{\theta}_*, \hat{\lambda}_*)$ solves (5), then $\hat{\theta}_*$ solves the following problem:

$$\begin{aligned} \min_{\theta \in C \cap \mathbb{R}^d} \quad & \hat{f}^0(\theta) \\ \text{s/t} \quad & \hat{f}^i(\theta) \leq 0, \quad 1 \leq i \leq r \end{aligned} \tag{6}$$

(see Theorem 2.18 in p. 48 of Zangwill (1969)), which can be viewed as a relaxed version of (3). Thus the remaining question is the relationship between the solution to (6) and the solution to (3). The following proposition confirms their relationship in the case where the solution to (6) is an integer point; see Luo and Lim (2011) for a proof.

Proposition 1 Suppose that there exists a solution to (6), say $\hat{\theta}_*$, and that $\hat{\theta}_* \in \mathbb{Z}^d$. Then $\hat{\theta}_*$ is a solution to (3). Therefore, if $(\hat{\theta}_*, \hat{\lambda}_*)$ is a saddle point of \hat{L} , i.e.,

$$\hat{L}(\hat{\theta}_*, \lambda) \leq \hat{L}(\hat{\theta}_*, \hat{\lambda}_*) \leq \hat{L}(\theta, \hat{\lambda}_*)$$

for all $\theta \in C \cap \mathbb{R}^d$ and $\lambda \in \mathbb{R}_+^r$, and $\hat{\theta}_*$ is an integer point, then $\hat{\theta}_*$ is an optimal solution to (3).

Given the relationship between the saddle point of \hat{L} and the solution to (3), the proposed method applies stochastic approximation to \hat{L} in order to search for the saddle point. Denoting the n th estimator of the saddle point of L by (θ_n, λ_n) , we update θ_n and λ_n as follows. Given $(\theta_1, \lambda_1), \dots, (\theta_n, \lambda_n)$, we observe a quantity $-D_n(\theta_n, \lambda_n)$ that guides us towards the steepest descent direction in θ ($-D_n(\theta_n, \lambda_n)$ can be interpreted as the negative of the derivative of \hat{L} with respect to θ if \hat{L} is differentiable in θ or the negative of the subgradient of \hat{L} in θ if \hat{L} is convex in θ). We then update θ_n by the recursion

$$\theta_{n+1} = \Pi_C(\theta_n - c_n D_n(\theta_n, \lambda_n)),$$

where $\Pi_C(\theta)$ is the closest point in C to $\theta \in \mathbb{R}^d$ with respect to the norm $\|\cdot\|$. $(c_n : n \geq 1)$ is a sequence of positive real numbers.

Using the fact that

$$\partial \hat{L}(\theta, \lambda) / \partial \lambda^i = \hat{f}^i(\theta)$$

for $1 \leq i \leq r$ provided that the partial derivative exists, we update λ_n by the recursion

$$\lambda_{n+1}^i = \max(0, \lambda_n^i + c_n F^i(\theta_n))$$

for $1 \leq i \leq r$, where $F^i(\theta_n)$ is an observation of \widehat{f}^i at θ_n .

Finally, we assume that there exists a known bound K for a saddle point of \widehat{L} ; i.e., there exists a positive constant K such that $|\widehat{\theta}_*^i| \leq K$ for $1 \leq i \leq d$ and $|\widehat{\lambda}_*^i| \leq K$ for $1 \leq i \leq r$, where $(\widehat{\theta}_*, \widehat{\lambda}_*)$ is a saddle point of \widehat{L} . With this additional information, we project (θ_n, λ_n) onto $\mathcal{B} \triangleq \{(\theta, \lambda) \in \mathbb{R}^d \times \mathbb{R}^r : |\theta^i| \leq K \text{ for } 1 \leq i \leq d, |\lambda^i| \leq K \text{ for } 1 \leq i \leq r\}$ and the projected point is $(\theta_{n+1}, \lambda_{n+1})$.

Our proposed method takes the following form in general:

Algorithm 1: General Form of the Proposed Method

Step 0. Initialize: Select a starting point $(\theta_0, \lambda_0) \in \mathbb{Z}^d \times \mathbb{R}^r$. Set $n = 0$.

Step 1. Update θ_n and λ_n : Generate observations $D_n(\theta_n, \lambda_n)$ and $F^i(\theta_n)$ for $1 \leq i \leq r$ and set

$$\begin{aligned} \bar{\theta}_{n+1} &= \Pi_C(\theta_n - c_n D_n(\theta_n, \lambda_n)), \\ \bar{\lambda}_{n+1}^i &= \max(0, \lambda_n^i + c_n F^i(\theta_n)) \end{aligned}$$

for $1 \leq i \leq r$.

Step 2. Project $(\bar{\theta}_{n+1}, \bar{\lambda}_{n+1})$ onto \mathcal{B} and set the projected point equal to $(\theta_{n+1}, \lambda_{n+1})$.

Step 3. Let $n = n + 1$ and go to Step 1.

To investigate the asymptotic behavior of $((\theta_n, \lambda_n) : n \geq 1)$, we focus on the case where \widehat{f}^i ($0 \leq i \leq r$) is convex. In this case, differentiability of \widehat{f}^i is not necessary; only D_n needs to be an unbiased estimate of a subgradient of \widehat{L} in θ and F^i needs to be an unbiased estimate of \widehat{f}^i for $1 \leq i \leq r$. Below are the assumptions of the main theorem (Theorem 1). For a proof of Theorem 1, see Luo and Lim (2011).

- A1. $(c_n : n \geq 1)$ is a sequence of positive numbers satisfying $\sum_{n=1}^{\infty} c_n = \infty$ and $\sum_{n=1}^{\infty} c_n^2 \leq \infty$.
- A2. \widehat{f}^0 is strictly convex and \widehat{f}^i ($1 \leq i \leq r$) is convex.
- A3. There exists $\eta \in C$ such that $\widehat{f}^i(\eta) < 0$ for $1 \leq i \leq r$. The optimal value of (6) is finite.
- A4. $\mathbb{E}[D_n(\theta_n, \lambda_n) | \mathcal{F}_n]$ is a subgradient of \widehat{L} at (θ_n, λ_n) as a function of θ ; i.e.,

$$\widehat{L}(\theta, \lambda_n) \geq \widehat{L}(\theta_n, \lambda_n) + \mathbb{E}[D_n(\theta_n, \lambda_n) | \mathcal{F}_n]^T (\theta - \theta_n) \quad (7)$$

for all $\theta \in \mathbb{R}^d$. In addition, we have

$$\mathbb{E}[F^i(\theta_n) | \mathcal{F}_n] = \widehat{f}^i(\theta_n) \quad (8)$$

for $1 \leq i \leq r$ and $n \geq 1$,

$$\mathbb{E}[\|D_n(\theta_n, \lambda_n) - \mathbb{E}[D_n(\theta_n, \lambda_n) | \mathcal{F}_n]\|^2 | \mathcal{F}_n] < \sigma^2, \quad (9)$$

for $n \geq 1$, and

$$\mathbb{E}[(F^i(\theta_n) - \mathbb{E}[F^i(\theta_n) | \mathcal{F}_n])^2 | \mathcal{F}_n] < \sigma^2 \quad (10)$$

for $1 \leq i \leq r$ and $n \geq 1$ for some positive constant σ^2 , where \mathcal{F}_n is the σ -field generated by $(\theta_1, \lambda_1), \dots, (\theta_n, \lambda_n)$.

Theorem 1 Under A1–A4, i) there exists a saddle point $(\widehat{\theta}_*, \widehat{\lambda}_*)$ of \widehat{L} , ii) $\widehat{\theta}_*$ is unique, and iii) $\theta_n \rightarrow \widehat{\theta}_*$ a.s. as $n \rightarrow \infty$. By Proposition 1, if $\widehat{\theta}_*$ is an integer point, then $\widehat{\theta}_*$ is a solution to (3).

3.2 Extension via Piecewise Linear Interpolation

The continuous extensions \widehat{f}^i ($0 \leq i \leq r$) introduced in Section 3.1 can be chosen arbitrarily. However to make our procedure more concrete, we introduce one possible way of extending functions from \mathbb{Z}^d to \mathbb{R}^d . In particular, we consider extending a function $h : \mathbb{Z}^d \rightarrow \mathbb{R}$ via the piecewise linear interpolation over a particular partition of \mathbb{R}^d as follows. For $\theta \in \mathbb{R}^d$, we let $p = \lfloor \theta \rfloor$ and $q = (q^1, \dots, q^d) = \theta - p$. σ is the permutation of $(1, \dots, d)$ such that $\sigma(j)$ is the index of the j th largest of q^1, \dots, q^d (if $q^{\sigma(j)} = q^{\sigma(k)}$ for some j and k , then let $\sigma(j) > \sigma(k)$ when $j > k$). We set $U_0 = \emptyset$ and $U_j = \{\sigma(1), \dots, \sigma(j)\}$ for $1 \leq j \leq d$. We define $\widetilde{h} : \mathbb{R}^d \rightarrow \mathbb{R}$ by

$$\begin{aligned} \widetilde{h}(\theta) &= (1 - q^{\sigma(1)})h(p) + (q^{\sigma(1)} - q^{\sigma(2)})h(p + \chi_{U_1}) \\ &\quad + \dots + (q^{\sigma(d-1)} - q^{\sigma(d)})h(p + \chi_{U_{d-1}}) + q^{\sigma(d)}h(p + \chi_{U_d}). \end{aligned} \quad (11)$$

By construction, $\widetilde{h}(\theta) = h(\theta)$ for $\theta \in \mathbb{Z}^d$, so \widetilde{h} is a continuous extension of h over \mathbb{R}^d .

Even though \widetilde{h} is not differentiable at some points in \mathbb{R}^d , a subgradient of \widetilde{h} can be easily computed when \widetilde{h} is convex. We define $\varphi\widetilde{h}(\theta) = (\varphi\widetilde{h}^j(\theta) : j = 1, \dots, d)$ by

$$\varphi\widetilde{h}^j(\theta) = \widetilde{h}(p + \chi_{U_k}) - \widetilde{h}(p + \chi_{U_{k-1}}) \quad (12)$$

for $\theta \in \mathbb{R}^d$, where $q^j = q^{\sigma(k)}$.

The following propositions prove that $\varphi\widetilde{h}(\theta)$ is a subgradient of h at $\theta \in \mathbb{R}^d$ when \widetilde{h} is convex; see Luo and Lim (2011) for proofs.

Proposition 2 Let $h : \mathbb{Z}^d \rightarrow \mathbb{R}$ be given and $\widetilde{h} : \mathbb{R}^d \rightarrow \mathbb{R}$ be defined by (11). For any $\theta \in \mathbb{R}^d$ and $\delta > 0$, i) there exists $\theta_\delta \in \mathbb{R}^d$ such that $\|\theta - \theta_\delta\| \leq \delta$, ii) \widetilde{h} is differentiable at θ_δ , iii) $\varphi\widetilde{h}(\theta) = \varphi\widetilde{h}(\theta_\delta)$, and iv) $\varphi\widetilde{h}(\theta_\delta)$ is the gradient of \widetilde{h} at θ_δ .

Proposition 3 Let $h : \mathbb{Z}^d \rightarrow \mathbb{R}$ be given and \widetilde{h} and $\varphi\widetilde{h}$ be defined by (11) and (12), respectively. If \widetilde{h} is convex, then $\varphi\widetilde{h}(\theta)$ is a subgradient of \widetilde{h} at $\theta \in \mathbb{R}^d$.

We are now ready to discuss how this strategy of constructing an extension can be adopted in the proposed method. We define the extensions $\widetilde{f}^i : \mathbb{R}^d \rightarrow \mathbb{R}$ ($0 \leq i \leq r$) via the linear interpolation by

$$\begin{aligned} \widetilde{f}^i(\theta) &= (1 - q^{\sigma(1)})f^i(p) + (q^{\sigma(1)} - q^{\sigma(2)})f^i(p + \chi_{U_1}) \\ &\quad + \dots + (q^{\sigma(d-1)} - q^{\sigma(d)})f^i(p + \chi_{U_{d-1}}) + q^{\sigma(d)}f^i(p + \chi_{U_d}) \end{aligned} \quad (13)$$

for $\theta \in \mathbb{R}^d$, where $p, q, \sigma, U_0, \dots, U_d$ are defined as before. Using these functions, a continuous extension $\widetilde{L} : \mathbb{R}^d \times \mathbb{R}^r \rightarrow \mathbb{R}$ of L is defined as

$$\widetilde{L}(\theta, \lambda) = \widetilde{f}^0(\theta) + \sum_{i=1}^r \lambda^i \widetilde{f}^i(\theta)$$

for $\theta \in \mathbb{R}^d$ and $\lambda \in \mathbb{R}_+^r$.

We next define $\varphi\widetilde{f}^i(\theta) = (\varphi\widetilde{f}_j^i(\theta) : j = 1, \dots, d)$ for $0 \leq i \leq r$ by

$$\varphi\widetilde{f}_j^i(\theta) = \widetilde{f}^i(p + \chi_{U_k}) - \widetilde{f}^i(p + \chi_{U_{k-1}}) \quad (14)$$

for $\theta \in \mathbb{R}^d$, where $q^j = q^{\sigma(k)}$. Propositions 2 and 3 justify our choice of $\varphi\widetilde{f}^i$ as a subgradient of \widetilde{f}^i .

With the \tilde{f}^i s as the extended functions, our proposed method proceeds as follows. We denote the n th estimator of the saddle point of \tilde{L} by (θ_n, λ_n) . Given $(\theta_1, \lambda_1), \dots, (\theta_n, \lambda_n)$, we observe f^i ($0 \leq i \leq r$) at $p_n + \chi_{U_k}$ for $0 \leq k \leq d$, and presume that

$$Y_n^i(k) = f^i(p_n + \chi_{U_k}) + \varepsilon_n(k), \tag{15}$$

where $p_n = \lfloor \theta_n \rfloor$, $q_n = (q_n^1, \dots, q_n^d) = \theta_n - p_n$, σ_n is the permutation of $(1, \dots, d)$ such that $\sigma_n(j)$ is the index of the j th largest of q_n^1, \dots, q_n^d (if $q_n^{\sigma_n(j)} = q_n^{\sigma_n(k)}$ for some j and k , then let $\sigma_n(j) > \sigma_n(k)$ when $j > k$), $U_0 = \emptyset$, $U_k = \{\sigma_n(1), \dots, \sigma_n(k)\}$ for $1 \leq k \leq d$, and $(\varepsilon_n(k) : 0 \leq k \leq d, n \geq 1)$ are mean zero random variables. We then update θ_n and λ_n by the recursion

$$\begin{aligned} \theta_{n+1} &= \Pi_C(\theta_n - c_n D_n(\theta_n, \lambda_n)), \\ \lambda_{n+1}^i &= \max\left(0, \lambda_n^i + c_n \left((1 - q_n^{\sigma(1)}) Y_n^i(0) + \dots + q_n^{\sigma(d)} Y_n^i(d) \right)\right) \end{aligned}$$

for $1 \leq i \leq r$, where the j th element of $D_n(\theta_n, \lambda_n)$ is $Y_n^0(k) - Y_n^0(k-1) + \sum_{i=1}^r \lambda_n^i (Y_n^i(k) - Y_n^i(k-1))$ with $q_n^j = q_n^{\sigma(j)}$. We then project $(\theta_{n+1}, \lambda_{n+1})$ onto \mathcal{B} and the projected point becomes $(\theta_{n+1}, \lambda_{n+1})$.

We observe that under the assumption that the $\varepsilon_n(k)$ s are mean zero random variables, (14) and (15) imply

$$\mathbb{E}[D_n(\theta_n, \lambda_n) | (\theta_1, \lambda_1), \dots, (\theta_n, \lambda_n)] = \varphi \tilde{f}(\theta_n) + \sum_{i=1}^r \lambda_n^i \varphi \tilde{f}^i(\theta_n).$$

Below is the proposed method when we adopt the above procedure.

Algorithm 2: Proposed algorithm with Extensions via Linear Interpolation

Step 0. Initialize: Select a starting point $(\theta_0, \lambda_0) \in \mathbb{Z}^d \times \mathbb{R}^r$. Set $n = 0$.

Step 1. Update θ_n and λ_n : Generate an observation $Y_n^i(k)$ of f^i at $p_n + \chi_{U_k}$ for $0 \leq i \leq r$ and $0 \leq k \leq d$, where p_n and the U_k s are defined as before. Set

$$\begin{aligned} \bar{\theta}_{n+1} &= \Pi_C(\theta_n - c_n D_n(\theta_n, \lambda_n)), \\ \bar{\lambda}_{n+1}^i &= \max\left(0, \lambda_n^i + c_n \left((1 - q_n^{\sigma(1)}) Y_n^i(0) + \dots + q_n^{\sigma(d)} Y_n^i(d) \right)\right), \end{aligned}$$

where the j th element of $D_n(\theta_n, \lambda_n)$ is $Y_n^0(k) - Y_n^0(k-1) + \sum_{i=1}^r \lambda_n^i (Y_n^i(k) - Y_n^i(k-1))$ with $q_n^j = q_n^{\sigma(j)}$.

Step 2. Project $(\bar{\theta}_{n+1}, \bar{\lambda}_{n+1})$ onto \mathcal{B} and set the projected point equal to $(\theta_{n+1}, \lambda_{n+1})$.

Step 3. Let $n = n + 1$ and go to Step 1.

To analyze the behavior of $((\theta_n, \lambda_n) : n \geq 1)$ generated from Algorithm 2, we shall impose some assumptions. In particular, we require:

A5. \tilde{f}^0 is strictly convex and \tilde{f}^i ($1 \leq i \leq r$) is convex.

A6. There exists $\eta \in \mathbb{R}^d$ such that $\tilde{f}^i(\eta) < 0$ for $1 \leq i \leq r$. Let $f_* = \min_{\theta \in C \cap \mathbb{R}^d} \tilde{f}^0$ subject to $\tilde{f}^i(\theta) \leq 0$ for $1 \leq i \leq r$ and assume f_* is finite.

A7. The $\varepsilon_n(j)$ s are random variables satisfying

$$\mathbb{E}[\varepsilon_n(j) | \mathcal{F}_n] = 0$$

and

$$\mathbb{E}[\varepsilon_n(j)^2 | \mathcal{F}_n] < \kappa^2$$

for $0 \leq j \leq d$, $n \geq 1$, and some positive constant κ^2 , where \mathcal{F}_n is the σ -field generated by $(\theta_1, \lambda_1), \dots, (\theta_n, \lambda_n)$.

From Theorem 1, we have the following corollary.

Corollary 2 Under A1 and A5–A7, i) there exists a saddle point $(\tilde{\theta}_*, \tilde{\lambda}_*)$ of \tilde{L} , ii) $\tilde{\theta}_*$ is unique, and iii) $\theta_n \rightarrow \tilde{\theta}_*$ a.s. as $n \rightarrow \infty$. By Proposition 1, if $\tilde{\theta}_*$ is an integer point, then $\tilde{\theta}_*$ is the solution to (3).

4 NUMERICAL RESULTS

In this section, we investigate the performance of the proposed algorithms in the following two settings: (1) inventory control in a periodically-reviewed single-item inventory system, and (2) staffing in a call center that handles multiple types of calls while maintaining a satisfactory level of customer service.

We compare the proposed method to the methods proposed by Whitney et al. (2001) and Ahmed et al. (1997).

The method proposed in Whitney et al. (2001) incorporates a penalty function type method into the objective function. At the n th iteration of the method proposed by Whitney et al. (2001), we generate $\Delta_n = (\Delta_n^1, \dots, \Delta_n^d)$, where the Δ_n^i s are independent and identically distributed (iid) random variables taking values $+1$ or -1 with equal probability. Denoting the n th estimator of the optimal solution to (3) by θ_n , we then generate an observation Y_+^i of f^i at $\theta_n + \Delta_n$ and an observation Y_-^i of f^i at $\theta_n - \Delta_n$ for $0 \leq i \leq r$, and update θ_n by the recursion

$$\theta_{n+1} = \Pi_{C \cap \mathcal{B}_\theta} (\theta_n - [a_n H_n(\theta_n)]),$$

where the j th component of $H_n(\theta_n)$ is

$$((Y_+^0 + b_n \sum_{i=1}^r \max(0, Y_+^i)) - (Y_-^0 + b_n \sum_{i=1}^r \max(0, Y_-^i))) / (2\Delta_n^j)$$

for $1 \leq j \leq d$ and $\Pi_{C \cap \mathcal{B}_\theta}$ is the projection onto the set $C \cap \mathcal{B}_\theta$ with $\mathcal{B}_\theta \triangleq \{\theta \in \mathbb{R}^d : |\theta^j| \leq K \text{ for } 1 \leq j \leq d\}$. ($a_n : n \geq 1$) and ($b_n : n \geq 1$) are sequences of positive real numbers.

On the other hand, at the n th iteration of the method proposed by Ahmed et al. (1997), we set the neighborhood $\mathcal{N}(\theta_n)$ of θ_n as

$$\mathcal{N}(\theta_n) = \{\theta \in \mathbb{Z}^d : \|\theta - \theta_n\| = 1\}, \tag{16}$$

and choose a candidate for θ_{n+1} , say θ'_n , from $\mathcal{N}(\theta_n)$ with equal probability. We then generate iid observations of f_i ($1 \leq i \leq r$) at θ'_n and compute the sample mean and standard deviation, say \hat{f}^i and $\hat{\sigma}^i$, of those observations. We consider θ'_n feasible if

$$\hat{f}^i - t_{n-1, 1-\alpha} \hat{\sigma}^i \leq 0$$

for $1 \leq i \leq r$, where $t_{n-1, 1-\alpha}$ is the upper $1 - \alpha$ critical point for the t distribution with $n - 1$ degrees of freedom. If θ'_n is considered feasible by this criterion, we generate an observation Y_n of f^0 at θ_n and an observation Y'_n of f^0 at θ'_n , and accept θ'_n as θ_{n+1} if $Y'_n \leq Y_n$ or $\exp(-(Y'_n - Y_n)/T_f) > U_n$, where T_f is a positive constant and U_n is a random variable uniformly distributed between 0 and 1. If $Y'_n > Y_n$ and $\exp(-(Y'_n - Y_n)/T) \leq U_n$, then θ_n is chosen as θ_{n+1} . If θ'_n is not considered feasible, then θ_n is chosen as θ_{n+1} . We repeat this procedure M times after which T_f is replaced by $T_f R$, where $R < 1$ is a positive constant. The procedure is repeated M times again until T_f is replaced by $T_f R$ again. This process is repeated until a stopping criterion is satisfied.

The subsequent sections report the performance of the proposed method and other competing methods.

4.1 Inventory Control in a Periodic Review System

We consider a finite-horizon, periodically-reviewed, single-item inventory system with integer-valued iid demands and full backlogging. Orders are received at the beginning of each period, the demand for the

period arrives next, and we review the inventory position (= on hand stock minus backorders plus any outstanding orders) to make an ordering decision. The ordering decisions are made according to the (s, S) policy. If the inventory position is less than s , an order for the amount of S minus the inventory position is placed. Otherwise, no action is taken. The order lead time is assumed to be zero. When an order of x units is placed, the ordering cost of $K + cx$ is incurred, where K is the fixed setup cost per order and c is the unit cost. A holding cost of h per unit per period is charged against any unit left at the end of each time period. The service level is measured using the fill rate, which is defined as the fraction of demand that is met directly from stock on hand. By $f(s, S)$, we denote the average ordering and holding costs per period over 1,000 time periods when the inventory position at the beginning of the first period is initialized at S and the system is governed by the (s, S) policy. By $g(s, S)$, we denote the fill-rate over 1,000 time periods when the inventory position at the beginning of the first period is initialized at S and the system is governed by the (s, S) policy. Our goal is to determine the values s and S , say s_* and S_* , that minimize $f(s, S)$ subject to the constraint that $g(s, S)$ is greater than or equal to a prescribed level β .

We apply Algorithm 2, and the methods proposed by Whitney et al. (2001) and Ahmed et al. (1997) to find (s_*, S_*) . Whenever we observe f at each point in $\mathcal{F}_1 \triangleq \{(x, y) \in \mathbb{Z}^2 : 1 \leq x \leq 100, 1 \leq y \leq 100, x \leq y\}$ in all three methods, the inventory system is simulated over 1,000 time periods, the ordering and holding costs are averaged over the 1,000 time periods, and the average of 20 iid such replications is used as an observation of f . Likewise, whenever we observe g at each point in \mathcal{F}_1 , the inventory system is simulated over 1,000 time periods, the demand which is met directly from stock over the 1,000 periods is divided by the total demand over the 1,000 periods, and the average of 20 iid such replications is used as an observation of g . θ_1 is set as $(100, 100)$ for all the three methods and $\lambda_1 = 275$ is used for Algorithm 2. The parameters used are $c_n = 500/(0.1N/60 + n)$ for the first 10% of the total iterations available, $c_n = 50/(0.1N/60 + n)$ for the rest of the iterations available, $a_n = 300/(0.1N/40 + n)$, $b_n = 400$, $\alpha = 0.95$, $T_f = 100$, $R = 0.6$, $M = 10$, $K = 100$, $c = 3$, $h = 3$, $\beta = 0.95$, and demand in each time period follows a Poisson distribution with a mean of 30.

To compare the estimates of (s_*, S_*) produced by the proposed method and other methods to the true values, s_* and S_* are estimated by evaluating $f(s, S)$ and $g(s, S)$, using the average of 100 iid replications at each point in \mathcal{F}_1 and selecting the values s and S that minimize the estimated f value while the estimated g value is greater than or equal to β . The “true” optimal solution estimated this way is $(s_*, S_*) = (18, 60)$.

Table 1: Averages and standard deviation of $\theta_{l(N)+1}$.

N	Algorithm 2		Whitney et al. (2001)		Ahmed et al. (1997)	
	Mean	Std	Mean	Std	Mean	Std
1000	(15, 64)	8.8	(23, 62)	15.7	(88, 107)	13.3
2000	(17, 61)	3.7	(18, 61)	10.7	(75, 106)	19.9
4000	(18, 60)	1.4	(16, 60)	3.7	(34, 77)	21.1
8000	(18, 60)	0.5	(17, 60)	1.5	(24, 66)	17.0
12000	(18, 60)	0.3	(16, 60)	0.9	(21, 63)	14.6
16000	(18, 60)	0.3	(16, 60)	0.8	(19, 61)	13.5
20000	(18, 60)	0.3	(16, 60)	0.8	(18, 61)	12.8
(s_*, S_*)	(18, 60)		(18, 60)		(18, 60)	

Denoting the number of simulation runs made at iteration n by t_n and fixing the total number N of simulation runs available, we compute $\theta_{l(N)+1}$ where $l(N)$ is the maximum number of iterations given the N simulation runs available; i.e., $l(N)$ is the largest integer satisfying $t_1 + \dots + t_{l(N)} \leq N$. Thus, $\theta_{l(N)+1}$ is the best estimate of (s_*, S_*) given the computational budget N . We notice that both f and g can be simultaneously computed in a single simulation run, so $t_n = 60, 40$, and 20 for Algorithm 2, the method by

Whitney et al. (2001), and the method by Ahmed et al. (1997), respectively. Table 1 reports the averages (Mean) of $\theta_{l(N)+1}$ generated by Algorithm 2, the method by Whitney et al. (2001), and the method by Ahmed et al. (1997) based on 200 independent copies of $\theta_{l(N)}$ for each value of N . To measure how the distribution of $\theta_{l(N)+1}$ is spread out, the average of the sample standard deviation of $\theta_{l(N)+1}^1$ and $\theta_{l(N)+1}^2$ is reported in Table 1.

4.2 Staffing in a Call Center

We consider a call center which handles three types of calls; calls that request technical support, calls that ask for sales information, and calls that wish to check order status. An incoming call is one of the three types with probability 0.5, 0.3, and 0.2, respectively. Calls arrive at the call center according to a Poisson process with rate λ per minute. Calls that enter the center form a single queue of infinite capacity and are served on a first come first serve basis. The call center opens at 8 AM and closes at 6 PM. After 6 PM, all remaining calls should be handled before they exit the system. Thus, each simulation run starts and ends with an empty system.

If a customer requests technical support, they must select one of the three products (products 1, 2, and 3) that they wish technical support for. We assume that the percentages of requests for the three products are 25%, 34%, and 41%, respectively. The request for product i is served by a staff member of type i for $1 \leq i \leq 3$, and the service time per customer requested by a staff member of type i follows a triangular distribution with lower limit 3, upper limit 18, and mode 6. Staff members of type 4 are available to handle the calls for all three products. They serve a customer only when there are no staff members of types 1, 2, and 3 available. The service time per customer requested by a staff member of type 4 follows a triangular distribution with lower limit 3, upper limit 18, and mode 6.

If a customer asks for sales information, then the customer is serviced by a staff member of type 5. The service time per customer requested by a staff member of type 5 follows a triangular distribution with lower limit 4, upper limit 45, and mode 15.

If a customer wishes to check order status, the request is handled by an automatic phone system, and there is no limit on the number of such calls that the automatic phone system can handle. The service time spent on the automated system follows a triangular distribution with lower limit 2, upper limit 4, and mode 3. After this automated service, 15% of the customers ask for a salesperson and wait on line until served by a staff member of type 5. The service time per customer requested by a staff member of type 5 in this case follows a triangular distribution with lower limit 4, upper limit 45, and mode 15.

Each staff member serves calls on a first come first serve basis. All service times are independent of each other and independent of the arrival process.

We denote the number of staff members of type i by θ^i ($1 \leq i \leq 5$). By $f(\theta^1, \theta^2, \theta^3, \theta^4, \theta^5)$, we denote the daily average operating costs, given the staffing level $(\theta^1, \theta^2, \theta^3, \theta^4, \theta^5)$. By $g(\theta^1, \theta^2, \theta^3, \theta^4, \theta^5)$, we denote the fraction of calls waiting less than 90 seconds in the queue before they initiate their service, given the staffing level $(\theta^1, \theta^2, \theta^3, \theta^4, \theta^5)$. The goal is to find the numbers of staff members of types 1, 2, 3, 4, and 5 minimizing $f(\theta^1, \theta^2, \theta^3, \theta^4, \theta^5)$ while ensuring that $g(\theta^1, \theta^2, \theta^3, \theta^4, \theta^5) \geq 0.8$. The operating costs consist of labor costs, which are \$100 per day for staff members of types 1, 2, 3, and 5, and \$200 per day for a staff member of type 4.

We apply Algorithm 2, and the methods proposed by Whitney et al. (2001) and Ahmed et al. (1997) to find the optimal values of θ^i , say θ_*^i ($1 \leq i \leq 5$). Whenever we observe f at each point in $\mathcal{F}_2 \triangleq \{\theta \in \mathbb{Z}^5 : 1 \leq \theta^i \leq 100 \text{ for } 1 \leq i \leq 5\}$ in all three methods, we simulate the system over 10 days, and compute the average of the operating costs over the 10-day time horizon, and use the average of 10 iid such replications as an observation of f . Likewise, whenever we observe the fraction of calls that wait less than 90 seconds in the queue, we simulate the system over 10 days, divide the number of calls that waited less than 90 seconds in the queue by the total number of calls over the 10-day time horizon, and use the average of 10 iid such replications as an observation of g . θ_1 is set as $(50, 50, 50, 50, 50)$ for all the three methods and $\lambda_1 = 10,000$ is used for Algorithm 2. The parameters used are $c_n = 0.3/(0.1N/60 + n)$ for

the first 50% of the total iterations available, $c_n = 0.1/(0.1N/60 + n)$ for the rest of the iterations available, $a_n = 0.5/(0.1N/20 + n)$, $b_n = 20000$, $\alpha = 0.95$, $T_f = 50$, $R = 0.4$, $M = 10$, and $\lambda = 10$.

The optimal policy $(\theta_*^1, \theta_*^2, \theta_*^3, \theta_*^4, \theta_*^5)$ is estimated by evaluating $f(\theta)$ and the fraction of calls waited less than 90 seconds, using the average of 50 iid observations at each $\theta \in \mathcal{F}_2$. The “true” optimal solution estimated this way is (8, 9, 11, 1, 39).

Table 2: Averages and standard deviation of $\theta_{l(N)+1}$.

N	Algorithm 2		Whitney et al. (2001)		Ahmed et al. (1997)	
	Mean	Std	Mean	Std	Mean	Std
1000	(25, 27, 25, 3, 43)	3.2	(15, 15, 16, 4, 40)	7.3	(26, 11, 10, 10, 40)	9.8
1500	(17, 19, 17, 1, 42)	2.8	(11, 12, 15, 3, 41)	6.3	(25, 10, 9, 10, 39)	9.2
2000	(12, 14, 14, 1, 43)	2.2	(9, 10, 13, 3, 41)	5.7	(24, 9, 8, 9, 39)	8.7
3000	(9, 11, 13, 1, 42)	1.1	(7, 9, 12, 3, 41)	4.9	(23, 9, 7, 9, 39)	8.2
4000	(8, 11, 12, 1, 42)	0.7	(7, 8, 12, 3, 40)	4.3	(22, 8, 7, 8, 39)	7.9
5000	(8, 11, 13, 1, 42)	0.5	(7, 8, 11, 3, 40)	4.3	(22, 8, 7, 8, 39)	7.8
6000	(8, 11, 12, 1, 42)	0.4	(7, 8, 11, 2, 40)	4.2	(21, 8, 7, 8, 39)	7.7
7000	(8, 11, 12, 1, 42)	0.3	(8, 8, 12, 2, 40)	4.2	(21, 7, 7, 8, 39)	7.6
$(\theta_*^1, \theta_*^2, \theta_*^3, \theta_*^4, \theta_*^5)$	(8, 9, 11, 1, 39)		(8, 9, 11, 1, 39)		(8, 9, 11, 1, 39)	

Denoting the number of simulation runs made at iteration n by t_n and fixing the total number N of simulation runs available, we compute $\theta_{l(N)+1}$ where $l(N)$ is the maximum number of iterations given the N simulation runs available; i.e., $l(N)$ is the largest integer satisfying $t_1 + \dots + t_{l(N)} \leq N$. Thus, $\theta_{l(N)+1}$ is the best estimate of $(\theta_*^1, \theta_*^2, \theta_*^3, \theta_*^4, \theta_*^5)$ given the computational budget N . We note that both f and g can be simultaneously computed in a single simulation run, so $t_n = 60, 20$, and 10 for Algorithm 2, the method by Whitney et al. (2001), and the method by Ahmed et al. (1997), respectively. Table 2 reports the averages (Mean) of $\theta_{l(N)+1}$ generated by Algorithm 2, the method by Whitney et al. (2001), and the method by Ahmed et al. (1997) based on 50 independent copies of $\theta_{l(N)}$ for each value of N . To measure how the distribution of $\theta_{l(N)+1}$ is spread out, the average of the sample standard deviation of $\theta_{l(N)+1}^1, \theta_{l(N)+1}^2, \theta_{l(N)+1}^3, \theta_{l(N)+1}^4$, and $\theta_{l(N)+1}^5$ is reported in Table 2.

REFERENCES

- Ahmed, M. A., T. M. Alkhamis, and M. Hasan. 1997. “Optimizing discrete stochastic systems using simulated annealing and simulation”. *Computers and Industrial Engineering* 32:823–836.
- Andradóttir, S. 1995. “A method for discrete stochastic optimization”. *Management Science* 41 (12): 1946–1961.
- Andradóttir, S., D. Goldsman, and S. H. Kim. 2005, December. “Finding the best in the presence of a stochastic constraint”. In *Proceedings of the 2005 Winter Simulation Conference*, edited by M. E. Kuhl, N. M. Steiger, F. B. Armstrong, and J. A. Joines, 732–738. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.
- Andradóttir, S., and S. H. Kim. 2010. “Fully sequential procedures for comparing constrained systems via simulation”. *Naval Res. Logistics* 57:403–421.
- Batur, D., and S. H. Kim. 2001. “Finding feasible systems in the presence of constraints on multiple performance measures”. *ACM Transactions on Modeling and Computer Simulation* 47:800–816.
- Fu, M. C. 2002. “Optimization for simulation: Theory vs. practice”. *INFORMS Journal on Computing* 14:192–215.

- Gelfand, S. B., and S. K. Mitter. 1989. "Simulated annealing with noisy or imprecise energy measurements". *Journal of Optimization Theory and Applications* 62 (1): 49–62.
- Glover, F. 1989. "Tabu search—Part I". *ORSA Journal on Computing* 1:190–206.
- Goldsman, D., and B. L. Nelson. 1994, December. "Ranking, selection and multiple comparisons in computer simulations". In *Proceedings of the 1994 Winter Simulation Conference*, 192–199. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.
- Henderson, S. G., and B. L. Nelson. 2006. *Simulation. Handbooks in Operations Research and Management Science*, Volume 13. Amsterdam, The Netherlands: Elsevier.
- Ho, Y. C., R. Sreenivas, and P. Vakili. 1992. "Ordinal optimization of discrete event dynamic systems". *Journal of Discrete Event Dynamic Systems* 2 (2): 61–88.
- Hong, L. J., and B. L. Nelson. 2006. "Discrete optimization via simulation using COMPASS". *Operations Research* 54:115–129.
- Kleywegt, A. J., A. Shapiro, and T. Homem-de-mello. 2001. "The sample average approximation method for stochastic discrete optimization". *SIAM Journal on Optimization* 12 (2): 479–502.
- Kushner, H. J., and D. C. Clark. 1978. *Stochastic approximation methods for constrained and unconstrained systems*. New York: Springer-Verlag.
- Kushner, H. J., and E. Sanvicente. 1975. "Stochastic Approximation of Constrained Systems With System and Constraint Noise". *Automatica* 11:375–380.
- Li, J., A. Sava, and X. Xie. 2009. "Simulation-based discrete optimization of stochastic discrete event systems subject to non closed form constraints". *IEEE Transactions on Automatic Control* 54:2900–2904.
- Liepins, G. E., and M. R. Hilliard. 1989. "Genetic algorithms: foundations and applications". *Annals of Operations Research* 21:31–58.
- Luo, Y., and E. Lim. 2011. "Simulation-based optimization over discrete sets with noisy constraints". *Submitted*.
- Murray, W. 1967. "Ill-conditioning in barrier and penalty functions arising in constrained nonlinear programming". In *Proceedings of the Sixth International Symposium on Mathematical Programming*.
- Nelson, B. L. 2010. "Optimization via Simulation over Discrete Decision Variables". In *TutORials in Operations Research*, edited by J. J. Hasenbein, Volume 7, 193–207. INFORMS.
- Pujowidianto, N. A., L. H. Lee, C. H. Chen, and C. M. Yap. 2009, December. "Optimal Computing Budget Allocation for Constrained Optimization". In *Proceedings of the 2009 Winter Simulation Conference*, edited by M. D. Rossetti, R. R. Hill, B. Johansson, A. Dunkin, and R. G. Ingalls, 584–589. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.
- Rockafellar, R. T. 1970. *Convex Analysis*. Princeton, New Jersey: Princeton University Press.
- Shi, L., and S. Olafsson. 2000. "Nested partitions method for stochastic optimization". *Methodology and Computing in Applied Probability* 2:271–291.
- Whitney, J. E., L. I. Solomon, and S. D. Hill. 2001, December. "Constrained Optimization over Discrete Sets via SPSA with Application to Non-separable Resource Allocation". In *Proceedings of the 2001 Winter Simulation Conference*, edited by B. A. Peters, J. S. Smith, D. J. Medeiros, and M. W. Rohrer, 313–317. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.
- Yan, D., and H. Mukai. 1992. "Stochastic discrete optimization". *SIAM Journal of Control and Optimization* 30:594–612.
- Zangwill, W. I. 1969. *Nonlinear Programming: A Unified Approach*. Englewood Cliffs, New Jersey: Prentice-Hall, Inc.

AUTHOR BIOGRAPHIES

YAO LUO is currently a Ph.D. student in the Department of Industrial Engineering at the University of Miami. He received a B.En. in Industrial Engineering from Shanghai Jiao Tong University. His research interests are in simulation optimization and supply chain management. His email address is y.luo6@umiami.edu.

EUNJILIM is an Assistant Professor in the Department of Industrial Engineering at the University of Miami. Her research interest includes stochastic optimization and simulation. Her e-mail address is lim@miami.edu.