Chapter 10
A Review of Random Search Methods

Sigrún Andradóttir

Abstract This chapter provides a brief review of random search methods for simulation optimization. We start by describing the structure of random search when system performance is estimated via simulation. Next, we discuss methods for solving simulation optimization problems with discrete decision variables and one (stochastic) performance measure, with emphasis on simulated annealing. Finally, we expand our scope to address simulation optimization problems with continuous decision variables and/or multiple (stochastic) performance measures.

10.1 Introduction

This chapter describes the use of random search to optimize complex stochastic systems whose expected performance under any particular system design is unavailable in closed form, and instead must be estimated via computer simulation. Thus, if Θ denotes the set of all possible system designs and $f(x) = \mathsf{E}[Y(x,\xi)]$ denotes the expected system performance under each design $x \in \Theta$, then we aim to solve the optimization problem

$$\min_{x\in\Theta} f(x) \tag{10.1}$$

under the assumption that the values of the objective function f will be estimated using simulation. We will outline the generic structure of random search methods and provide a review of a representative set of specific random search approaches.

Random search methods involve repeatedly sampling and evaluating system designs based on the observed history (i.e., the designs that have been sampled so far and their estimated performance) in search of the best feasible design. They are well suited for solving simulation optimization problems where the objective function often has little known structure (and hence derivatives are unavailable) and the optimization procedure must identify improved solutions with guidance from the

S. Andradóttir (✉)
Georgia Institute of Technology, Atlanta, GA, USA
e-mail: sa@gatech.edu

M.C. Fu (ed.), *Handbook of Simulation Optimization*, International Series in Operations Research & Management Science 216, DOI 10.1007/978-1-4939-1384-8_10

estimated performance of the solutions considered so far. The scope of applicability of random search is broad, and includes deterministic and stochastic optimization problems with discrete and/or continuous decision parameters.

The outline of this chapter is as follows. In Sect. 10.2, we describe the structure of random search methods. In Sect. 10.3, we review a representative set of random search methods for solving discrete simulation optimization problems. In Sect. 10.4, we discuss example procedures for solving simulation optimization problems with either continuous decision variables or multiple performance measures. An earlier overview of random search for simulation optimization emphasizing desirable features such as convergence and efficiency can be found in Andradóttir [14].

10.2 Structure of Random Search Methods for Simulation Optimization

In this section, we describe the structure of random search methods applied to solve the simulation optimization problem (10.1). We will let $\mathscr{S}_n$ denote the sampling strategy used to select candidate system designs in iteration $n \geq 1$ of the algorithm, with M_n specifying the number of such designs. The sampling strategy can be updated adaptively as the algorithm learns from the results so far, and does not need to be chosen in advance of execution. Similarly, M_n is a parameter of the sampling strategy $\mathscr{S}_n$, and consequently does not need to be chosen in advance. 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). The following generic random search algorithm can also be found in Andradóttir [14].

Generic Random Search Algorithm for Simulation Optimization:

Step 0 (initialize): Choose the initial sampling strategy $\mathscr{S}_1$ and let $n = 1$.
Step 1 (sample): Select $x_n^{(1)}, \ldots, x_n^{(M_n)} \in \Theta$ according to the sampling strategy $\mathscr{S}_n$.
Step 2 (simulate): Estimate $f(x_n^{(i)})$, for $i = 1, \ldots, M_n$, using simulation.
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 $\mathscr{S}_{n+1}$. Let $n = n+1$ and go to Step 1.

The primary difference among random search methods involves the choice of the sampling strategies $\{\mathscr{S}_n\}$. Also, the generic algorithm does not include a stopping criterion. This is consistent with the fact that convergence results for random search are typically asymptotic in nature. In practice, it is of course necessary to augment the algorithm with a suitable stopping criterion.

In the remainder of this chapter, we will describe specific random search procedures for simulation optimization. As much of the random search literature

focuses on problems with discrete decision variables, we start by considering such problems in Sect. 10.3. Then we consider extensions to continuous feasible regions and multiple objectives in Sect. 10.4.

10.3 Discrete Simulation Optimization

In this section, we review example procedures for discrete simulation optimization problems with one (stochastic) performance measure. (In addition, there may be deterministic performance measures whose values do not need to be estimated via simulation. They can be incorporated into the constraint set Θ.) We start by discussing research on how the well-known simulated annealing algorithm can be applied to solve discrete simulation optimization problems in Sect. 10.3.1. Then we briefly review certain other random search procedures for solving the optimization problem (10.1) in Sect. 10.3.2.

10.3.1 Simulated Annealing

The simulated annealing algorithm dates back to the pioneering work by Metropolis et al. [61]. Since then, a large body of literature has appeared on simulated annealing, including important works by Kirkpatrick et al. [55], Mitra et al. [62], Hajek [42], and others. A basic version of the simulated annealing algorithm for solving a deterministic optimization problem of the form (10.1) is provided below. We will need the following notation:

- for all $x \in \Theta$, $N(x) \subset \Theta$ is a set of "neighbors" of x (alternatively, $N(x)$ is a "neighborhood" of x);
- for all $x \in \Theta$, $R(x, \cdot)$ is a probability distribution on $N(x)$;
- $\{T_n\}$ is a sequence of strictly positive numbers;
- w.p. $=$ with probability;
- $(y)^+ = \max\{0, y\}$ for all $y \in \mathbb{R}$.

Basic Simulated Annealing Algorithm for Deterministic Optimization:

Step 0 (initialize): Choose an initial system design $x_0 \in \Theta$ and let $n = 1$.
Step 1 (sample): Select a candidate solution $x'_n \in N(x_{n-1})$ according to the probability distribution $R(x_{n-1}, \cdot)$.
Step 2 (update): Let

$$x_n = \begin{cases} x'_n & \text{w.p. } \exp\left(-\frac{\left(f(x'_n) - f(x_{n-1})\right)^+}{T_n}\right), \\ x_{n-1} & \text{otherwise,} \end{cases}$$

$n = n + 1$, and go to Step 1.

It is clear that the Basic Simulated Annealing Algorithm described above fits within the framework outlined in Sect. 10.2 with $M_n = 1$ and $x_n^* = x_n$ for all $n \in \mathbb{N}$, except that a "simulate" step is of course not needed here, as the exact objective function values are available. Simulated annealing is designed to solve global optimization problems in the presence of (possibly multiple) local optimal solutions. In each iteration n, the algorithm generates a candidate solution x_n', and then decides whether to stay at the current solution x_{n-1} or move to the candidate solution. More specifically, if the candidate solution x_n' satisfies $f(x_n') \leq f(x_{n-1})$, so that x_n' is better than x_{n-1}, then x_n' becomes the new estimated optimal solution. On the other hand, if $f(x_n') > f(x_{n-1})$, so that the candidate solution is worse than x_{n-1}, then there is nevertheless a chance that x_n' will be chosen as a new estimate of the optimal solution. This "hill-climbing" feature is designed to allow the algorithm to escape from locally optimal solutions that are not globally optimal (i.e., solutions x satisfying $f(x) < f(x')$ for all $x' \in N(x)$ but $f(x) > \inf_{x' \in \Theta} f(x')$). The probability $\exp(-(f(x_n') - f(x_{n-1}))^+/T_n)$ of making such a hill-climbing move depends both on how inferior the candidate solution is relative to the current solution (i.e., on the magnitude of $(f(x_n') - f(x_{n-1}))^+$) and on the current "temperature" T_n, with hill-climbing moves being less likely for worse candidate solutions and for smaller temperatures. Hill-climbing decisions can be made by generating a uniform random number U_n on the interval $[0,1]$, accepting the candidate solution x_n' if $U_n \leq \exp(-(f(x_n') - f(x_{n-1}))^+/T_n)$, and rejecting it otherwise (in which case the current solution x_{n-1} remains the estimate of the optimal solution). In addition to requiring that the temperature T_n be strictly positive for all n, most of the simulated annealing literature assumes that $T_n \to 0$ as $n \to \infty$ at a logarithmic rate, see for example Hajek [42].

The previous discussion addressed the use of simulated annealing to solve deterministic optimization problems. In the remainder of this section, we will review simulated annealing algorithms designed to solve optimization problems with noisy objective function values, as is the case in simulation optimization.

Bulgak and Sanders [22] present a heuristic simulated annealing approach and use it to solve a buffer allocation problem. They deal with the noise in the objective function evaluations by using confidence intervals to ensure that the difference in performance is statistically significant when the candidate state has a better objective function estimate than the current state. Haddock and Mittenthal [43] also present a heuristic simulated annealing method together with numerical results. Their method employs a different "update" step than the Basic Simulated Annealing Algorithm described above (motivated by the work of Glauber [36]), and also uses a rapidly decreasing temperature sequence $\{T_n\}$ (their sequence decreases at an exponential rate, rather than at a logarithmic rate).

We now turn to simulated annealing approaches that are provably convergent when applied to solve discrete simulation optimization problems. For all a and $b \neq 0$, let $\mathcal{N}(a, b^2)$ denote a normal random variable with mean a and variance b^2. Gelfand and Mitter [34] show that if the transition matrix $R(\cdot,\cdot)$ is irreducible, the temperature sequence $\{T_n\}$ converges to zero, and the noise in the estimate of the

difference $f(x'_n) - f(x_{n-1})$ in performance between the candidate and current states in iteration n has a $\mathcal{N}(0, \sigma_n^2)$ distribution where $\sigma_n = o(T_n)$ as $n \to \infty$, then their simulated annealing procedure with noisy objective function estimates converges in probability to the set Θ^* of global optimal solutions provided that the same algorithm using exact objective function values converges to Θ^* in probability.

Gutjahr and Pflug [41] also present convergence results for the simulated annealing algorithm when the noise in the estimated objective function values needed in iteration n has the normal $\mathcal{N}(0, \sigma_n^2)$ distribution. They show that when $\sigma_n = O(n^{-\gamma})$, where $\gamma > 1$, then the simulated annealing algorithm with noisy objective function evaluations has the same asymptotic performance as when exact objective functions are available (and hence converges in probability to Θ^* if the temperature sequence $\{T_n\}$ is chosen properly). They also generalize their convergence result to noise distributions that are symmetric around zero and more peaked around zero than the $\mathcal{N}(0, \sigma_n^2)$ distribution satisfying $\sigma_n = O(n^{-\gamma})$, where $\gamma > 1$ (i.e., for all $\varepsilon > 0$, the noise is more likely to take values in $(-\varepsilon, \varepsilon)$ than the specified $\mathcal{N}(0, \sigma_n^2)$ random variable).

Gelfand and Mitter [34] and Gutjahr and Pflug [41] assume that the variance σ_n^2 in the objective function evaluations required in iteration n of the simulated annealing algorithm converges to zero as n grows. Thus, more precise estimates are required for larger values of n, which typically results in more computation time per iteration. Fox and Heine [33] provide convergence guarantees for simulated annealing applied to solve discrete simulation optimization problems that do not require a restrictive variance assumption. However, they assume that the objective function values are restricted to a finite set (they also consider relaxing this assumption). Each time an estimate of an objective function value $f(x)$ is needed, they generate a few more observations of $f(x)$ and average them with observations of $f(x)$ obtained in earlier iterations to obtain the desired estimate of $f(x)$. They show that this variant of simulated annealing converges in probability to Θ^* if the same algorithm with the exact objective function values converges in probability to Θ^*.

Alrefaei and Andradóttir [6] present two simulated annealing algorithms for discrete simulation optimization. Beyond using noisy objective function estimates, rather than the exact objective function values, these algorithms differ from the Basic Simulated Annealing Algorithm presented above in two important ways, namely the temperature sequence is constant (i.e., $T_n = T > 0$ for all n) and the choice of the estimate of the optimal solution x_n^* is decoupled from the sequence $\{x_n\}$ used to search the state space Θ for the optimal solution (see also the Generic Random Search Algorithm described in Sect. 10.2). Alrefaei and Andradóttir [6] prove that their algorithms converge almost surely to Θ^* and provide numerical results comparing their algorithms with each other and with the methods analyzed by Gelfand and Mitter [34], Gutjahr and Pflug [41], and Fox and Heine [33].

Alkhamis et al. [5] study a simulated annealing algorithm that employs confidence intervals to determine whether the difference between the estimated objective function values at the current and candidate solutions is statistically significant.

They prove convergence in probability to Θ^* when the noise the objective function evaluations in iteration n converges to zero sufficiently fast relative to the temperature sequence $\{T_n\}$, and also provide numerical results. Ahmed and Alkhamis [1] analyze a simulated annealing approach with a constant temperature and with decoupled sequences $\{x_n\}$ and $\{x_n^*\}$ (see Alrefaei and Andradóttir [6]) that uses the two-stage ranking and selection procedure by Dudewicz and Dalal [30] to decide how many objective function observations are collected from the current and candidate solutions in iteration n. Alkhamis and Ahmed [4] continue this work by combining the approach with constant temperature and decoupled $\{x_n\}$ and $\{x_n^*\}$ sequences of Ahmed and Alkhamis [1] with the confidence interval approach of Alkhamis et al. [5]. Wang and Zhang [76] study a simulated annealing approach where a hypothesis test is used to determine whether to stay at the current state or move to the candidate state.

Prudius and Andradóttir [70] study two simulated annealing algorithms for discrete simulation optimization with decreasing temperatures $\{T_n\}$ and with decoupled $\{x_n\}$ and $\{x_n^*\}$ sequences. The two algorithms differ in that one uses only the data collected on the objective function values at the current and candidate solutions x_{n-1} and x_n' in the current iteration n to decide on the next current point x_n (no averaging), whereas the other one uses data collected on the values of $f(x_{n-1})$ and $f(x_n')$ in iterations 1 through n to decide on x_n (averaging). Both algorithms are shown to converge almost surely, and numerical results show that using all available data on the objective function values (as in averaging) does not necessarily improve performance (because the associated reduction in noise is not necessarily beneficial).

As was mentioned at the beginning of this section, the literature on simulated annealing for deterministic optimization is vast, and several researchers have studied the application of simulated annealing to solve simulation optimization problems. In addition to the contributions reviewed so far in this section, other works on simulated annealing for noisy response functions include Painton and Diwekar [66], who incorporate a penalty function to account for noise in the objective function estimates, Rosen and Harmonosky [71], who combine simulated annealing with response surface methodology, and Branke et al. [21], who consider a different "update" step than the Basic Simulated Annealing Algorithm above under a known variance assumption.

10.3.2 Other Developments

In this section, we briefly review certain other random search methods that have been developed for solving discrete simulation optimization problems. Several of these techniques are reviewed in more detail in other chapters in this volume, viz. Chaps. 2, 11, and 12. Additional material on random search for simulation optimization can, for example, be found in the reviews by Jin and Branke [52] and Bianchi et al. [20].

Stochastic ruler methods constitute a class of random search methods for discrete simulation optimization. Like simulated annealing, they involve sampling a single candidate point x_n' in each iteration n and deciding whether to accept this point (so that $x_n = x_n'$) or reject it (so that $x_n = x_{n-1}$). However, unlike simulated annealing, this decision is not made by comparing estimated objective function values at the current and candidate points. Instead, estimated objective function values at the candidate point x_n' are compared with observations of the "stochastic ruler," which is a uniform random variable whose range covers (approximately) the range of the estimated objective function values. The original stochastic ruler method was proposed by Yan and Mukai [83] and proven to converge to Θ^* in probability. Alrefaei and Andradóttir [7, 8] have studied modified versions of the stochastic ruler method that involve less work per iteration (i.e., the maximum number of comparisons with the stochastic ruler in iteration n is constant, rather than diverging to infinity with n) and decoupling the $\{x_n\}$ and $\{x_n^*\}$ sequences (see Sect. 10.2). They prove that their approaches converge almost surely to Θ^* and provide numerical results comparing the approaches.

Most random search methods for simulation optimization compare system designs based on performance estimates, and can thus be regarded as *ascent/descent methods*. For example, Gong et al. [37] present and analyze a "stochastic comparison" method that resembles the stochastic ruler method of Yan and Mukai [83] except that each iteration involves comparisons between objective function estimates at the current and candidate points, rather than comparisons with a stochastic ruler. For other related work, see Andradóttir [11–13] and Homem-de-Mello [44].

Deterministic optimization features various methods that involve partitioning the feasible region, including branch-and-bound (see, e.g., Nemhauser and Wosley [63]) and random search (see, e.g., Pintér [69]). *Partitioning methods* have also been developed for simulation optimization. For example, Norkin et al. [64] present a *branch-and-bound* method that involves partitioning the feasible region Θ into subsets, estimating upper and lower bounds on the optimal objective function value within each subset, choosing the estimated optimal solution from the subset with the smallest upper bound, and further partitioning the "record subset" that has the smallest lower bound. The estimates of the upper and lower bounds improve when the subset remains in the partition, and eventually converge to the actual values of the upper and lower bounds. Moreover, the upper and lower bounds are tight for singletons. Norkin et al. [64] prove that their method converges almost surely to Θ^*, discuss the choice of upper and lower bounds in various settings, and provide illustrative examples. This work is continued by Norkin et al. [65], who provide additional analysis, discussion about bound estimation, and examples.

Shi and Ólafsson [73] present a *nested partitions* method for simulation optimization. Like branch-and-bound, their approach involves partitioning one subset ("the most promising region") in most iterations. The other subsets are then combined into one "surrounding region." Then sample designs are collected from each region and their performance estimated via simulation. The promising index is estimated for each region as the best estimated performance of the designs sampled from

the region, and the region with the best promising index will be the new most promising region. If the most promising region is a singleton, then it can clearly not be partitioned further. If it is the surrounding region, then the method backtracks to either its super-region or the entire feasible region. Otherwise, the most promising region is partitioned. Shi and Ólafsson [73] prove that their method converges almost surely to Θ^*. Pichitlamken and Nelson [68] present a nested partitions method that differs from the original method of Shi and Ólafsson [73] in that a ranking and selection method is used to select the best sampled solution from each subset, hill-climbing and restart steps are added, and the sequence $\{x_n^*\}$ is chosen differently. Moreover, Xu and Nelson [80] present and analyze a method that combines the branch-and-bound and nested partition approaches.

Hong and Nelson [46] present a random search algorithm for local simulation optimization named *COMPASS* (for "convergent optimization via most-promising-area stochastic search"). The feasible region is composed of vectors with integer elements, and the approach is local in that the aim is to identify a design with better performance than any point with Euclidian distance of one away from that design. COMPASS keeps track of all designs sampled so far and performance estimates at these designs. In each iteration, new designs are sampled from the portion of the feasible region that is closer to the design with the best estimated performance than to any other sampled point (again in Euclidian distance, with adjustments to ensure that the sampling region is bounded) and simulation results are obtained for the newly sampled points (and possibly also for previously sampled points). Hong and Nelson [46] prove convergence w.p.1 to a local optimal solution and provide numerical examples. This work is continued by Hong [45], Hong et al. [47], and Xu et al. [78, 79] who improve the efficiency of the original COMPASS approach (see also the discussion of Xu [77] in Sect. 10.4.1 below).

Andradóttir and Prudius [15] discuss the need for balancing exploration (global search), exploitation (local search), and estimation (of objective function values at promising points) within simulation optimization, and then present two versions of *BEESE* (for "balanced explorative and exploitative search with estimation"), a random search approach designed to achieve such a balance. The two approaches are called R-BEESE (for Random-BEESE) and A-BEESE (Adaptive-BEESE). Both methods switch between global search and local search for improved system designs, with R-BEESE doing so at random and A-BEESE doing so adaptively based on recent progress made via global and local search. Both methods also add an estimation component via resampling of the design with the best estimated performance and by ensuring that sufficient data has been collected at the estimated optimal solution (moreover, A-BEESE involves more local search than its deterministic variant A-BEES, which also adds an estimation component). Andradóttir and Prudius [15] prove that their methods converge almost surely to Θ^* and provide numerical results.

Model-based methods form another class of random search techniques. These methods maintain a probabilistic model on the solutions space Θ that is used to generate candidate solutions, whose estimated performance is in turn used to update the probabilistic model. The cross-entropy method of Rubinstein and Kroese [72]

involves the use of a parametric family of distributions and relies on the relationship between optimization and rare event simulation. In each iteration, solutions are sampled from the current model in the parametric family. The performance of these solutions is estimated, a sample quantile is computed, and the solutions whose estimated performance exceeds the sample quantile are used to obtain an updated model in the parametric family. The model update involves approximating an optimal importance sampling distribution on the set of solutions whose performance exceeds the sample quantile via the use of the Kullback–Leibler (cross-entropy) distance. Hu et al. [48] provide a model-based algorithm for simulation optimization that differs from the cross-entropy method primarily in the updating of the model. They provide convergence analysis that is applicable both when Θ is discrete and continuous.

10.4 Extensions

In this section, we review example procedures for solving simulation optimization problems with either continuous decision variables or multiple (stochastic) performance measures. We start by discussing research on continuous simulation optimization problems in Sect. 10.4.1. Then we review certain procedures for solving simulation optimization problems with multiple performance measures in Sect. 10.4.2.

10.4.1 Continuous Simulation Optimization

Several researchers have studied the use of random search methods to solve continuous simulation optimization problems. The simplest form of random search is *pure (non-adaptive) random search*, where solutions are sampled repeatedly from a fixed distribution on the feasible region Θ (e.g., the search does not utilize information gathered in previous iterations to guide the search for improved solutions). Baumert and Smith [19] present a pure random search approach that estimates the objective function value at each sampled solution x by averaging all observations that are within a certain distance from x. The sampled point with the best estimated objective function value is chosen as the estimated optimal solution. Baumert and Smith [19] discuss at what rate the distance should decrease in order for the method to converge in probability. Their work was continued by Andradóttir and Prudius [16] who provide further analysis of the (deterministic) shrinking ball method of Baumert and Smith [19], develop and analyze the stochastic shrinking ball method, and provide numerical results.

Chia and Glynn [24] study the rate of convergence of pure random search as a function of the number m of sampled points and number n of observations at each point, with the estimated optimal solution being the point with the best estimated

objective function value. They identify at what rates m and n should grow to achieve the best rate of convergence. Similarly, Ensor and Glynn [31] study the choice of m and n in grid search. Cheng [23] also studies the asymptotic behavior of pure random search, addresses implementation issues, and provides numerical examples. In related research, Yakowitz et al. [81] study how the number of points vs. number of observations per point should be selected in search approaches that use low-dispersion sequences to select points. They also discuss a sequential version of their approach and the use of different numbers of observations to estimate the objective function value at different sampled points.

Alexander et al. [3] develop a pure random search procedure that iteratively samples solutions from Θ and then compares the incumbent and sampled solutions using increasingly precise (as the number of iterations grows) estimates of the objective function values at these solutions; the point with the better estimate becomes the new estimate of the optimal solution. They show that their procedure is globally convergent w.p.1. Ghate and Smith [35] study a generalized simulated annealing procedure that also involves comparing estimated objective function values at the incumbent and sampled solutions in each iteration, with the estimate being more precise for larger numbers of iterations. They prove convergence in probability and provide numerical results. Various other authors also study methods that move between current and sampled solutions based on estimated objective function values at those points, see for example Gurin [38], Gurin and Rastrigin [39], Devroye [28], and Marti [59] (and Devroye [29] for related work with finite Θ).

Yakowitz and Lugosi [82] develop a method that in certain iterations samples new solutions from a fixed global distribution (as in pure random search) and ensures that every sampled point has a sufficient number of observations, and in other iterations it adaptively *resamples* previously sampled points. The estimate of the optimal solution is the most recently sampled point. They prove that their method is globally convergent in probability. Andradóttir and Prudius [16] present the Adaptive Search with Resampling (ASR) method and prove that it is globally convergent w.p.1. Their method includes both sampling and resampling steps (similar to the approach of Yakowitz and Lugosi [82]), but the search is adaptive, only promising sampled points are "accepted" for further consideration (and hence additional observations are not collected at points that are not promising), and the estimated optimal solution is the best point sampled so far. Numerical results suggest that the ASR method performs better in practice than the earlier approach. Hu and Andradóttir [49] improve further on the ASR method by allowing previously accepted points to be discarded once better points have been found. They prove that their Adaptive Search with Resampling and Discarding (ASRD) method is convergent w.p.1 and provide numerical results indicating that the addition of discarding leads to substantial improvements in performance.

Huang et al. [51], Sun et al. [74], and Xu [77] all use Kriging meta-models and random search to solve simulation optimization problems. More specifically, Huang et al. [51] propose the SKO (Sequential Kriging Optimization) approach, where each iteration starts with a kriging meta-model of the objective function, identifies a solution that maximizes an Expected Improvement (EI) function (described in

Sect. 2.7 of [51]), and then either terminates the search (if the EI is small) or updates the kriging model using the new data point. Sun et al. [74] propose and analyze the GPS (Gaussian Process-based Search) algorithm whose sampling strategy takes into account how likely feasible solutions are to improve on the current best estimate of the optimal solution based on the current kriging model. The GPS approach can be used for both continuous and discrete simulation optimization. In related work, Xu [77] presents the SKOPE (Stochastic Kriging for OPtimization Efficiency) sampling approach and integrates this approach with the AHA discrete simulation optimization method of Xu et al. [79].

We conclude this section by briefly mentioning other methods that can be used to solve continuous simulation optimization problems. Methods that involve partitioning the feasible region have been developed by Deng and Ferris [27] and Kabirian and Ólafsson [54]. More specifically, Deng and Ferris [27] adapt the DIRECT (DIviding RECTangles) algorithm of Jones et al. [53] to simulation optimization, and Kabirian and Olafsson [54] present and analyze a golden region search algorithm for continuous simulation optimization. Model-based methods (see Sect. 10.3.2, Rubinstein and Kroese [72], and Hu et al. [48]) can be used for both discrete and continuous simulation optimization. Finally, Ferris et al. [32] and Deng and Ferris [26] discuss continuous simulation optimization algorithms that involve successive quadratic approximations of the objective function.

10.4.2 Simulation Optimization with Multiple Objectives

In this section, we review certain random search approaches designed for solving simulation optimization problems with multiple (stochastic) performance measures. Ahmed et al. [2] consider a problem with a deterministic objective function (cost) and stochastic constraints (on system performance). They present a simulated annealing approach for solving such problems, where a hypothesis testing step is added after candidate generation to determine if the candidate solution is feasible with the desired confidence. Baeslar and Sepúlveda [17] present a *goal programming* framework to handle multiple stochastic performance measures. A goal value is specified for each performance measure and the original (multi-objective) optimization problem is translated into a (single-objective) optimization problem where a weighted sum of the deviations from the specified goals is minimized (possibly after normalization to address discrepancies between measurement units for the different performance measures). This optimization problem is then solved using a genetic algorithm. Baeslar and Sepúlveda [18] continue this research by applying their methodology to optimize a cancer treatment center.

Multiple authors have proposed random search methods aiming to identify *Pareto optimal* solutions to simulation optimization problems with multiple performance measure. More specifically, a solution x dominates another solution x' if no objective performs worse at x than at x', and at least one objective performs strictly better at x than at x'. A solution is Pareto optimal if it is not dominated by any other solution.

Gutjahr [40] presents a Stochastic Pareto Simulated Annealing (SPSA) based on the Pareto Simulated Annealing (PSA) approach of Czyzak and Jaszkiewicz [25], which maintains a search set and a solution set, and in each iteration a candidate solution is generated in the neighborhood of each element of the search set, the estimated performance of the candidate and current points are compared, and the search and solution sets are updated (the update of the search set involves hill-climbing with weights computed by the algorithm). Gutjahr [40] also specifies a Stochastic Pareto Ant Colony Approach (SP-ACO) and compares the two approaches with a brute-force approach. Other approaches for multi-objective simulation optimization based on simulated annealing include Alrefaei et al. [9, 10] and Mattila et al. [60]. Lee et al. [56] present an approach based on multiobjective evolutionary algorithms (MOEA) (see, e.g., Zhou et al. [84]) and use it to solve an aircraft spare parts allocation problem, and Lee et al. [57] study a multi-objective COMPASS approach (see Sect. 10.3.2 for discussion of the original COMPASS approach).

Another approach to handling multiple performance measures is to designate one as the objective and the others as constraints. Li et al. [58] combine COMPASS with a penalty-function approach for handling *constraints*, and prove almost sure convergence of the resulting approach. Vieira et al. [75] also present and analyze an adaptation of COMPASS designed to handle one constraint. Park and Kim [67] present the penalty function with memory (PFM) approach for handling stochastic constraints. This approach can be combined with a random search approach designed for solving unconstrained problems (the authors combine it with the nested partitions approach, see Sect. 10.3.2) and differs from the approach of Li et al. [58] primarily in that the penalty function does not diverge at feasible solutions. Hu and Andradóttir [50] combine the ASRD framework discussed in Sect. 10.4.1 with a penalty function approach and prove that their framework guarantees almost sure convergence when applied to solve simulation optimization problems with stochastic constraints.

Acknowledgements The author gratefully acknowledges support from the National Science Foundation under Grant CMMI–0856600. The author also thanks Liujia Hu for helpful comments.

References

1. M. A. Ahmed and T. M. Alkhamis. Simulation-based optimization using simulated annealing with ranking and selection. *Computers and Operations Research*, 29:387–402, 2002.
2. M. A. Ahmed, T. M. Alkhamis, and M.Hasan. Optimizing discrete stochastic systems using simulated annealing and simulation. *Computers and Industrial Engineering*, 32:823–836, 1997.
3. D. L. J. Alexander, D. W. Bulger, J. M. Calvin, H. E. Romejin, and R. L. Sherriff. Approximate implementations of pure random search in the presence of noise. *Journal of Global Optimization*, 31:601–612, 2005.
4. T. M. Alkhamis and M. A. Ahmed. Simulation-based optimization using simulated annealing with confidence interval. In R. G. Ingalls, M. D. Rossetti, J. S. Smith, and B. A. Peters, editors, *Proceedings of the 2004 Winter Simulation Conference*, 514–519, 2004.

5. T. M. Alkhamis, M. A. Ahmed, and V. K. Tuan. Simulated annealing for discrete optimization with estimation. *European Journal of Operational Research*, 116:530–544, 1999.
6. M. H. Alrefaei and S. Andradóttir. A simulated annealing algorithm with constant temperature for discrete stochastic optimization. *Management Science*, 45:748–764, 1999.
7. M. H. Alrefaei and S. Andradóttir. A modification of the stochastic ruler method for discrete stochastic optimization. *European Journal of Operational Research*, 133:160–182, 2001.
8. M. H. Alrefaei and S. Andradóttir. Discrete stochastic optimization using variants of the stochastic ruler method. *Naval Research Logistics*, 52:344–360, 2005.
9. M. H. Alrefaei and A. H. Diabat. A simulated annealing technique for multi-objective simulation optimization. *Applied Mathematics and Computation*, 215:3029–3035, 2009.
10. M. Alrefaei, A. Diabat, A. Alawneh, R. Al-Aomar, and M. N. Faisal. Simulated annealing for multi objective stochastic optimization. *International Journal of Science and Applied Information Technology*, 2:18–21, 2013.
11. S. Andradóttir. A method for discrete stochastic optimization. *Management Science*, 41:1946–1961, 1995.
12. S. Andradóttir. A global search method for discrete stochastic optimization. *SIAM Journal on Optimization*, 6:513–530, 1996.
13. S. Andradóttir. Accelerating the convergence of random search methods for discrete stochastic optimization. *ACM Transactions on Modeling and Computer Simulation*, 9:349–380, 1999.
14. S. Andradóttir. An overview of simulation optimization via random search. Chapter 20 in S. G. Henderson and B. L. Nelson, editors, *Handbooks in Operations Research and Management Science: Simulation*. North-Holland, 2006.
15. S. Andradóttir and A. A. Prudius. Balanced explorative and exploitative search with estimation for simulation optimization. *INFORMS Journal on Computing*, 21:193–208, 2009.
16. S. Andradóttir and A. A. Prudius. Adaptive random search for continuous simulation optimization. *Naval Research Logistics*, 57:583–604, 2010.
17. F. F. Baeslar and J. A. Sepúlveda. Multi-response simulation optimization using stochastic genetic search within a goal programming framework. In J. A. Joines, R. R. Barton, K. Kang, and P. A. Fishwick, editors, *Proceedings of the 2000 Winter Simulation Conference*, 788–794, 2000.
18. F. F. Baeslar and J. A. Sepúlveda. Multi-objective simulation optimization for a cancer treatment center. In B. A. Peters, J. S. Smith, D. J. Medeiros, and M. W. Rohrer, editors, *Proceedings of the 2001 Winter Simulation Conference*, 1405–1411, 2001.
19. S. Baumert and R. L. Smith. Pure random search for noisy objective functions. Technical Report, University of Michigan, 2002.
20. L. Bianchi, M. Dorigo, L. M. Gambardella, and W. J. Gutjahr. A survey on metaheuristics for stochastic combinatorial optimization. *Natural Computing*, 8:239–287, 2009.
21. J. Branke, S. Meisel, and C. Schmidt. Simulated annaeling in the presence of noise. *Journal of Heuristics*, 14:627–654, 2008.
22. A. A. Bulgak and J. L. Sanders. Integrating a modified simulated annealing algorithm with the simulation of a manufacturing system to optimize buffer sizes in automatic assembly systems. In M. Abrams, P. Haigh, and J. Comfort, editors, *Proceedings of the 1988 Winter Simulation Conference*, 684–690, 1988.
23. R. C. H. Cheng. Fitting statistical models of random search in simulation studies. *ACM Transactions on Modeling and Computer Simulation* 23:article 15, 2013.
24. Y. L. Chia and P. W. Glynn. Limit theorems for simulation-based optimization via random search. *ACM Transactions on Modeling and Computer Simulation*, 23:article 16, 2013.
25. P. Czyzak and A. J. Jaszkiewicz. Pareto simulated annealing – A metaheuristic technique for multiple-objective combinatorial optimization. *Journal of Multi-Criteria Decision Analysis*, 7:34–47, 1998.
26. G. Deng and M. C. Ferris. Adaptation of the UOBYQA algorithm for noisy functions. In L. F. Perrone, F. P. Wieland, J. Liu, B. G. Lawson, D. M. Nicol, and R. M. Fujimoto, editors, *Proceedings of the 2006 Winter Simulation Conference*, 312–319, 2006.

27. G. Deng and M. C. Ferris. Extension of the DIRECT optimization algorithm for noisy functions. In S. G. Henderson, B. Biller, M. H. Hsieh, J. Shortle, J. D. Tew, and R. R. Barton, editors, *Proceedings of the 2007 Winter Simulation Conference*, 497–504, 2007.
28. L. Devroye. On the convergence of statistical search. *IEEE Transactions on Systems, Man and Cybernetics*, 6:46–56, 1976.
29. L. Devroye. Probabilistic search as a strategy selection procedure. *IEEE Transactions on Systems, Man and Cybernetics*, 6:315–321, 1976.
30. E. J. Dudewicz and S. R. Dalal. Allocation of observations in ranking and selection with unequal variances. *Sankhya*, B37:28–78, 1975.
31. K. B. Ensor and P. W. Glynn. Stochastic optimization via grid search. *Lectures in Applied Mathematics*, 33:89–100, 1997.
32. M. C. Ferris, T. S. Munson, and K. Sinapiromsaran. A practical approach to sample-path simulation optimization. In J. A. Joines, R. R. Barton, K. Kang, and P. A. Fishwick, editors, *Proceedings of the 2000 Winter Simulation Conference*, 795–804, 2000.
33. B. L. Fox and G. W. Heine. Probabilistic search with overrides. *The Annals of Applied Probability*, 5:1087–1094, 1995.
34. S. B. Gelfand and S. K. Mitter. Simulated annealing with noisy or imprecise energy measurements. *Journal of Optimization Theory and Applications*, 62:49–62, 1989.
35. A. Ghate and R. L. Smith. Adaptive search with stochastic acceptance probabilities for global optimization. *Operations Research Letters*, 36:285–290, 2008.
36. R. J. Glauber. Time-dependent statistics of the Ising model. *Journal of Mathematical Physics*, 4:294–307, 1963.
37. W.-B. Gong, Y.-C. Ho, and W. Zhai. Stochastic comparison algorithm for discrete optimization with estimation. *SIAM Journal on Optimization*, 10:384–404, 1999.
38. L. S. Gurin. Random search in the presence of noise. *Engineering Cybernetics*, 4:252–260, 1966.
39. L. S. Gurin and L. A. Rastrigin. Convergence of the random search method in the presence of noise. *Automation and Remote Control*, 26:1505–1511, 1965.
40. W. J. Gutjahr. Two metaheuristics for multiobjective stochastic combinatorial optimization. *Stochastic Algorithms: Foundations and Applications*, Lecture Notes in Computer Science, 3777:116–125, 2005.
41. W. J. Gutjahr and G. Ch. Pflug. Simulated annealing for noisy cost functions. *Journal of Global Optimization*, 8:1–13, 1996.
42. B. Hajek. Cooling schedules for optimal annealing. *Mathematics of Operations Research*, 13:311–329, 1988.
43. J. Haddock and J. Mittenthal. Simulation optimization using simulated annealing. *Computers and Industrial Engineering*, 22:387–395, 1992.
44. T. Homem-de-Mello. Variable-sample methods for stochastic optimization. *ACM Transactions on Modeling and Computer Simulation*, 13:108–133, 2003.
45. L. J. Hong. Discrete optimization via simulation using coordinate search. In M. E. Kuhl, N. M. Steiger, F. B. Armstrong, and J. A. Joines, editors, *Proceedings of the 2005 Winter Simulation Conference*, 803–810, 2005.
46. L. J. Hong and B. L. Nelson. Discrete optimization via simulation using COMPASS. *Operations Research*, 54:115–129, 2006.
47. L. J. Hong, B. L. Nelson, J. Xu. Speeding up COMPASS for high-dimensional discrete optimization via simulation. *Operations Research Letters*, 38:550–555, 2010.
48. J. Hu, M. C. Fu, and S. I. Marcus. A model reference adaptive search method for stochastic global optimization. *Communications in Information and Systems*, 8:245–276, 2008.
49. L. Hu and S. Andradóttir. Adaptive random search with discarding for continuous simulation-based optimization. Submitted for publication.
50. L. Hu and S. Andradóttir. Simulation-based continuous optimization with stochastic constraints. Working paper.
51. D. Huang, T. T. Allen, W. I. Notz, and N. Zeng. Global optimization of stochastic black-box systems via sequential kriging meta-models. *Journal of Global Optimization*, 34:441–466, 2006.

52. Y. Jin and J. Branke. Evolutionary optimization in uncertain environments. *IEEE Transactions on Evolutionary Computation*, 9:303–317, 2005.
53. D. R. Jones, C. D. Perttunen, and B. E. Stuckman. Lipschitzian optimization without the Lipschitz constant. *Journal of Optimization Theory and Application*, 79:157–181, 1993.
54. A. Kabirian and S. Ólafsson. Continuous optimization via simulation using Golden Region search. *European Journal of Operational Research*, 208:19–27, 2011.
55. S. Kirkpatrick, C. D. Gelatt, and M. P. Vecchi. Optimization by simulated annealing. *Science*, 220:671–680, 1983.
56. L. H. Lee, E. P. Chew, S. Teng, and Y. Chen. Multi-objective simulation-based evolutionary algorithm for an aircraft spare parts allocation problem. *European Journal of Operational Research*, 189:476–491, 2008.
57. L. H. Lee, E. P. Chew, and H. Li. Multi-objective COMPASS for discrete optimization via simulation. In S. Jain, R. R. Creasey, J. Himmelspach, K. P. White, and M. Fu, editors, *Proceedings of the 2011 Winter Simulation Conference*, 4070–4079, 2011.
58. J. Li, A. Sava, and X. Xie. Simulation-based discrete optimization of stochastic discrete event systems subject to non closed-form constraints. *IEEE Transactions on Automatic Control*, 54:2900–2904, 2009.
59. K. Marti. Minimizing noisy objective functions by random search methods. *Zeitschrift fur Angewandte Mathematik and Mechanik*, 62:T377–T380, 1982.
60. V. Mattila, K. Virtanen, and R. P. Hämäläinen. A simulated annealing algorithm for noisy multiobjective optimization. *Journal of Multi-criteria Decision Analysis*, 20:255–276, 2013.
61. N. Metropolis, A. Rosenbluth, M. Rosenbluth, A. Teller, and E. Teller. Equation of state calculations by fast computing machines. *Journal of Chemical Physics*, 21:1087–1092, 1953.
62. D. Mitra, F. Romeo, and A. Sangiovanni-Vincentelli. Convergence and finite-time behavior of simulated annealing. *Advances in Applied Probability*, 18:747–771, 1986.
63. G. L. Nemhauser and L. A. Wosley. *Integer and Combinatorial Optimzation*. Wiley, 1999.
64. V. I. Norkin, Y. M. Ermoliev, and A. Ruszczyński. On optimal allocation of indivisibles under uncertainty. *Operations Research*, 46:381–395, 1998.
65. V. I. Norkin, G. Ch. Pflug, and A. Ruszczyński. A branch and bound method for stochastic global optimization. *Mathematical Programming*, A83:425–450, 1998.
66. L. Painton and U. Diwekar. Stochastic annealing for synthesis under uncertainty. *European Journal of Operational Research*, 83:489–502, 1995.
67. C. Park and S.-H. Kim. Handling stochastic constraints in discrete optimization via simulation. In S. Jain, R. R. Creasey, J. Himmelspach, K. P. White, and M. Fu, editors, *Proceedings of the 2011 Winter Simulation Conference*, 4217–4226, 2011.
68. J. Pichitlamken and B. L. Nelson. A combined procedure for optimization via simulation. *ACM Transactions on Modeling and Computer Simulation*, 13:155–179, 2003.
69. J. Pintér. Convergence qualification of adaptive partition algorithms in global optimization. *Mathematical Programming*, 56:343–360, 1992.
70. A. A. Prudius and S. Andradóttir. Averaging frameworks for simulation optimization with applications to simulated annealing. *Naval Research Logistics*, 59:411–429, 2012.
71. S. L. Rosen and C. M. Harmonosky. An improved simulated annealing simulation optimization method for discrete parameter stochastic systems. *Computers and Operations Research*, 32:343–358, 2005.
72. R. Y. Rubinstein and D. P. Kroese. *The Cross-Entropy Method: A Unified Approach to Combinatorial Optimization, Monte-Carlo Simulation, and Machine Learning*. Springer, 2004.
73. L. Shi and S. Ólafsson. Nested partitions method for stochastic optimization. *Methodology and Computing in Applied Probability*, 2:271–291, 2000.
74. L. Sun, L. J. Hong, and Z. Hu. Optimization via simulation using Gaussian process-based search. In S. Jain, R. R. Creasey, J. Himmelspach, K. P. White, and M. Fu, editors, *Proceedings of the 2011 Winter Simulation Conference*, 4139–4150, 2011.
75. H. Vieira Jr., K. H. Kienitz, M. C. N. Belderrain. Discrete-valued, stochastic-constrained simulation optimization with COMPASS. In S. Jain, R. R. Creasey, J. Himmelspach, K. P. White, and M. Fu, editors, *Proceedings of the 2011 Winter Simulation Conference*, 4196–4205, 2011.

76. L. Wang and L. Zhang. Stochastic optimization using simulated annealing with hypothesis test. *Applied Mathematics and Computation*, 174:1329–1342, 2006.
77. J. Xu. Efficient discrete optimization via simulation using stochastic kriging. In C. Laroque, J. Himmelspach, R. Pasupathy, O. Rose, and A. M. Uhrmacher, editors, *Proceedings of the 2012 Winter Simulation Conference*, 466–477, 2012.
78. J. Xu, L. J. Hong, and B. L. Nelson. Industrial strength COMPASS: A comprehensive algorithm and software for optimization via simulation. *ACM Transactions on Modeling and Computer Simulation* 20:article 3, 2010.
79. J. Xu, B. L. Nelson, and L. J. Hong. An adapative hyperbox algorithm for high-dimensional discrete optimization via simulation problems. *INFORMS Journal on Computing*, 25:133–146, 2013.
80. W. L. Xu and B. L. Nelson. Empirical stochastic branch-and-bound for optimization via simulation. In B. Johansson, S. Jain, J. Montoya-Torres, J. Hugan, and E. Yücesan, editors, *Proceedings of the 2010 Winter Simulation Conference*, 983–994, 2010.
81. S. Yakowitz, P. L'Ecuyer, and F. Vázquez-Abad. Global stochastic optimization with low-dispersion point sets. *Operations Research*, 48:939–950, 2000.
82. S. Yakowitz and E. Lugosi. Random search in the presence of noise, with application to machine learning. *SIAM Journal on Scientific and Statistical Computing*, 11:702–712, 1990.
83. D. Yan and H. Mukai. Stochastic discrete optimization. *SIAM Journal on Control and Optimization*, 30:594–612, 1992.
84. A. Zhou, B.-Y. Qu, H. Li, S.-Z. Zhao, P. N. Suganthan, and Q. Zhang. Multiobjective evolutionary algorithms: A survey of the state of the art. *Swarm and Evolutionary Computation*, 1:32–49, 2011.