

A Model Reference Adaptive Search Method for Global Optimization

Jiaqiao Hu

Department of Applied Mathematics and Statistics, State University of New York at Stony Brook,
Stony Brook, New York 11794, jqhu@ams.sunysb.edu

Michael C. Fu

Robert H. Smith School of Business, and Institute for Systems Research, University of Maryland,
College Park, Maryland 20742, mfu@rsmith.umd.edu

Steven I. Marcus

Department of Electrical and Computer Engineering, and Institute for Systems Research, University of Maryland,
College Park, Maryland 20742, marcus@umd.edu

Model reference adaptive search (MRAS) for solving global optimization problems works with a parameterized probabilistic model on the solution space and generates at each iteration a group of candidate solutions. These candidate solutions are then used to update the parameters associated with the probabilistic model in such a way that the future search will be biased toward the region containing high-quality solutions. The parameter updating procedure in MRAS is guided by a sequence of implicit probabilistic models we call reference models. We provide a particular algorithm instantiation of the MRAS method, where the sequence of reference models can be viewed as the generalized probability distribution models for estimation of distribution algorithms (EDAs) with proportional selection scheme. In addition, we show that the model reference framework can also be used to describe the recently proposed cross-entropy (CE) method for optimization and to study its properties. Hence, this paper can also be seen as a study on the effectiveness of combining CE and EDAs. We prove global convergence of the proposed algorithm in both continuous and combinatorial domains, and we carry out numerical studies to illustrate the performance of the algorithm.

Subject classifications: programming: nondifferentiable, nonlinear.

Area of review: Stochastic Models.

History: Received March 2005; revisions received September 2005, May 2006; accepted May 2006.

1. Introduction

Global optimization problems arise in a wide range of applications and are often extremely difficult to solve. Following Zlochin et al. (2004), we classify the solution methods for both continuous and combinatorial problems as being either *instance-based* or *model-based*. In instance-based methods, searches for new candidate solutions depend directly on previously generated solutions. Some well-known instance-based methods are simulated annealing (SA) (Kirkpatrick et al. 1983), genetic algorithms (GAs) (Srinivas and Patnaik 1994), tabu search (Glover 1990), and the recently proposed nested partitions (NP) method (Shi and Ólafsson 2000). In model-based algorithms, new solutions are generated via an intermediate probabilistic model that is updated or induced from the previous solutions. The model-based search methods are a class of solution techniques introduced fairly recently. In general, most of the algorithms that fall in this category share a similar framework and usually involve the following two phases:

(1) Generate candidate solutions (e.g., random samples) according to a specified probabilistic model.

(2) Update the probabilistic model, on the basis of the data collected in the previous step, to bias the future search toward “better” solutions.

Some examples of model-based methods are ant colony optimization (ACO) (Dorigo and Gambardella 1997), the cross-entropy (CE) method (Rubinstein and Kroese 2004, De Boer et al. 2005), and the estimation of distribution algorithms (EDAs) (Mühlenbein and Paaß 1996). Among the above approaches, those that are most relevant to our work are the CE method and the EDAs.

The CE method was motivated by an adaptive algorithm for estimating probabilities of rare events (Rubinstein 1997) in stochastic networks. It was later realized (Rubinstein 1999, 2001) that the method can be modified to solve combinatorial and continuous optimization problems. The CE method starts with a family of parameterized probability distributions on the solution space and tries to find the parameter of the distribution that assigns maximum probability to the set of (near) optimal solutions. Implicit in CE is an optimal (importance sampling) distribution concentrated only on the set of optimal solutions (i.e., zero variance), and the key idea is to use an iterative scheme

to successively estimate the optimal parameter that minimizes the Kullback-Leibler (KL) divergence between the optimal distribution and the family of parameterized distributions. In the context of estimation of rare event probabilities, Homem-de-Mello (2007) shows the convergence of an adaptive version of CE to an estimate of the optimal (possibly local) CE parameter with probability one. Rubinstein (1999) shows the probability one convergence of a modified version of the CE method to the optimal solution for combinatorial optimization problems.

The EDAs were first introduced in the field of evolutionary computation in Mühlenbein and Paaß (1996). They inherit the spirit of the well-known GAs but eliminate the crossover and the mutation operators to avoid the disruption of partial solutions. In EDAs, a new population of candidate solutions is generated according to the probability distribution *induced* or *estimated* from the promising solutions selected from the previous generation. Unlike CE, EDAs often take into account the interaction effects between the underlying decision variables needed to represent the individual candidate solutions, which are often expressed explicitly through the use of different probabilistic models. We refer the reader to Larrañaga et al. (1999) for a review of the way in which different probabilistic models are used as EDAs instantiations. The convergence of a class of EDAs, under the infinite population assumption, to the global optimum can be found in Zhang and Mühlenbein (2004).

In this paper, we introduce a new randomized method, called model reference adaptive search (MRAS), for solving both continuous and combinatorial optimization problems. MRAS resembles CE in that it works with a family of parameterized distributions on the solution space. The motivation behind the method is to use a sequence of intermediate reference distributions to facilitate and guide the updating of the parameters associated with the family of parameterized distributions during the search process (see Wolpert 2004 for a similar “iterative focusing” idea in the context of finding bounded rational equilibria of common-interest games). At each iteration of MRAS, candidate solutions are generated from the distribution (among the prescribed family of distributions) that possesses the minimum KL-divergence with respect to the reference model corresponding to the previous iteration. These candidate solutions are, in turn, used to construct the next distribution by minimizing the KL-divergence with respect to the current reference model, from which future candidate solutions will be generated. In contrast, as mentioned previously, the CE method targets a single optimal (importance sampling) distribution to direct its parameter updating, where new parameters are obtained as the solutions to the problems of estimating the optimal reference parameters.

We propose an instantiation of the MRAS method, where the sequence of reference distributions can be viewed as

the generalized probability distribution models for EDAs with proportional selection scheme (Zhang and Mühlenbein 2004) and can be shown to converge to a degenerate distribution concentrated only on the set of optimal solutions. An attractive feature is that the sequence of reference distributions depends directly on the performance of the candidate solutions; thus the method automatically takes into account the correlations between the underlying decision variables so that the random samples generated at each stage can be efficiently utilized. However, in EDAs (with proportional selection scheme), these distribution models are directly constructed to generate new candidate solutions, which is in general a difficult and computationally intensive task, whereas in our approach, the sequence of reference models is only used implicitly to guide the parameter updating procedure; thus there is no need to build them explicitly. We show that for a class of parameterized probability distributions, the so-called natural exponential family (NEF), the proposed algorithm converges to an optimal solution with probability one.

In sum, our main contributions in this paper include: (i) We introduce a new framework for global optimization, which allows considerable flexibility in the choices of the reference models. Consequently, by carefully analyzing and selecting the sequence of reference models, one can design different (perhaps even more efficient) versions of our proposed approach. (ii) We propose an instantiation of the MRAS method, which incorporates the key ideas of CE and EDAs. We analyze its convergence properties in both continuous and combinatorial domains, and we test its performance on some widely used benchmark problems. (iii) We also explore the relationship between CE and MRAS. In particular, we show that the CE method can also be interpreted as an instance of the model reference framework. Based on this observation, we establish and discuss some important properties of CE. (iv) In addition, we extend the recent results on the convergence of quantile estimates in Homem-de-Mello (2007) to the cases where the random samples are drawn from a sequence of distributions rather than a single fixed distribution.

The rest of this paper is organized as follows. In §2, we describe the MRAS method and motivate the use of a specific sequence of distributions as the reference models. In §3, we describe the idealized version of MRAS and establish its global convergence properties. In §4, we provide a unified view of CE and MRAS and study the properties of the CE method. In §5, we implement the Monte Carlo version of MRAS and prove its probability one convergence. Illustrative numerical studies on both continuous and combinatorial optimization problems are given in §6. Some possible future research topics are outlined in §7. The proofs of the main results are given in the appendix, and the proofs of all other intermediate results can be found in an online supplement to this paper at <http://or.journal.informs.org>.

2. The Model Reference Adaptive Search Method

We consider the following global optimization problem:

$$x^* \in \arg \max_{x \in \mathcal{X}} H(x), \quad \mathcal{X} \subseteq \mathfrak{R}^n, \quad (1)$$

where the solution space \mathcal{X} is a nonempty set in \mathfrak{R}^n , and $H(\cdot): \mathcal{X} \rightarrow \mathfrak{R}$ is a deterministic function that is bounded from below, i.e., $\exists \mathcal{M} > -\infty$ such that $H(x) \geq \mathcal{M} \forall x \in \mathcal{X}$. Throughout this paper, we assume that problem (1) has a unique global optimal solution, i.e., $\exists x^* \in \mathcal{X}$ such that $H(x) < H(x^*) \forall x \neq x^*, x \in \mathcal{X}$.

In addition to the general regularity conditions stated above, we also make the following assumptions on the objective function.

ASSUMPTION A1. For any given constant $\xi < H(x^*)$, the set $\{x: H(x) \geq \xi\} \cap \mathcal{X}$ has a strictly positive Lebesgue or discrete measure.

ASSUMPTION A2. For any given constant $\delta > 0$, $\sup_{x \in A_\delta} H(x) < H(x^*)$, where $A_\delta := \{x: \|x - x^*\| \geq \delta\} \cap \mathcal{X}$, and we define the supremum over the empty set to be $-\infty$.

Intuitively, Assumption A1 ensures that any neighborhood of the optimal solution x^* will have a positive probability of being sampled. For ease of exposition, A1 restricts the class of problems under consideration to either continuous or discrete problems; however, we remark that the work of this paper can be easily extended to problems with mixture of both continuous and discrete variables. Because $H(\cdot)$ has a unique global optimizer, Assumption A2 is satisfied by many functions encountered in practice. Note that both Assumptions A1 and A2 hold trivially when \mathcal{X} is (discrete) finite and the counting measure is used.

The MRAS method approaches the above problem as follows. It works with a family of parameterized distributions $\{f(\cdot, \theta), \theta \in \Theta\}$ on the solution space, where Θ is the parameter space. Assume that at the k th iteration of the method, we have a sampling distribution $f(\cdot, \theta_k)$. We generate candidate solutions from this sampling distribution. The performances of these randomly generated solutions are then evaluated and used to calculate a new parameter vector $\theta_{k+1} \in \Theta$ according to a specified parameter updating rule. The above steps are performed repeatedly until a termination criterion is satisfied. The idea is that if the parameter updating rule is chosen appropriately, the future sampling process will be more and more concentrated on regions containing high-quality solutions.

In MRAS, the parameter updating is determined by another sequence of distributions $\{g_k(\cdot)\}$, called the *reference distribution*. In particular, at each iteration k , we look at the projection of $g_k(\cdot)$ on the family of distributions $\{f(\cdot, \theta), \theta \in \Theta\}$ and compute the new parameter vector θ_{k+1} that minimizes the Kullback-Leibler (KL) divergence

$$\begin{aligned} \mathcal{D}(g_k, f(\cdot, \theta)) &:= E_{g_k} \left[\ln \frac{g_k(X)}{f(X, \theta)} \right] \\ &= \int_{\mathcal{X}} \ln \frac{g_k(x)}{f(x, \theta)} g_k(x) \nu(dx), \end{aligned}$$

where ν is the Lebesgue/counting measure defined on \mathcal{X} , $X = (X_1, \dots, X_n)$ is a random vector taking values in \mathcal{X} , and $E_{g_k}[\cdot]$ denotes the expectation taken with respect to $g_k(\cdot)$. Intuitively, $f(\cdot, \theta_{k+1})$ can be viewed as a compact representation (approximation) of the reference distribution $g_k(\cdot)$; consequently, the feasibility and effectiveness of the method will, to some large extent, depend on the choices of reference distributions.

There are many different ways to construct the sequence of reference distributions $\{g_k(\cdot)\}$. Here, we propose to use the following simple iterative scheme. Let $g_0(x) > 0 \forall x \in \mathcal{X}$ be an initial probability density/mass function (p.d.f./p.m.f.) on the solution space \mathcal{X} . At each iteration $k \geq 1$, we compute a new p.d.f./p.m.f. by tilting the old p.d.f./p.m.f. $g_{k-1}(x)$ with the performance function $H(x)$ (for simplicity, here we assume that $H(x) > 0 \forall x \in \mathcal{X}$), i.e.,

$$g_k(x) = \frac{H(x)g_{k-1}(x)}{\int_{\mathcal{X}} H(x)g_{k-1}(x) \nu(dx)} \quad \forall x \in \mathcal{X}. \quad (2)$$

By doing so, we are assigning more weight to the solutions that have better performance. One direct consequence of this is that each iteration of (2) improves the expected performance. To be precise,

$$E_{g_k}[H(X)] = \frac{E_{g_{k-1}}[(H(X))^2]}{E_{g_{k-1}}[H(X)]} \geq E_{g_{k-1}}[H(X)].$$

Furthermore, it is possible to show that the sequence $\{g_k(\cdot), k = 0, 1, \dots\}$ will converge to a distribution that concentrates only on the optimal solution for arbitrary $g_0(\cdot)$. So, we will have $\lim_{k \rightarrow \infty} E_{g_k}[H(X)] = H(x^*)$. The above idea has been used previously, for example, in EDAs with proportional selection schemes (cf. Zhang and Mühlenbein 2004), and in randomized algorithms for solving Markov decision processes (Chang et al. 2007). However, in those approaches, the construction of $g_k(\cdot)$ in (2) needs to be carried out explicitly to generate new samples; moreover, because $g_k(\cdot)$ may not have any structure, sampling from it could be computationally expensive. In MRAS, these difficulties are circumvented by projecting $g_k(\cdot)$ on the family of parameterized distributions $\{f(\cdot, \theta)\}$. On the one hand, $f(\cdot, \theta_k)$ often has some special structure and therefore could be much easier to handle; on the other hand, the sequence $\{f(\cdot, \theta_{k+1}), k = 0, 1, \dots\}$ may retain some nice properties of $\{g_k(\cdot)\}$ and also converge to a degenerate distribution concentrated on the optimal solution.

3. The MRAS₀ Algorithm (Exact Version)

We now present a particular algorithm instantiation of MRAS that uses the reference distributions proposed in §2. Throughout the analysis, we use $P_{\theta_k}(\cdot)$ and $E_{\theta_k}[\cdot]$ to denote the probability and expectation taken with respect to the p.d.f./p.m.f. $f(\cdot, \theta_k)$, and $I_{\{A\}}$ to denote the indicator function, i.e.,

$$I_{\{A\}} := \begin{cases} 1 & \text{if event } A \text{ holds,} \\ 0 & \text{otherwise.} \end{cases}$$

Thus, under our notational convention,

$$P_{\theta_k}(H(X) \geq \gamma) = \int_{\mathcal{X}} I_{\{H(x) \geq \gamma\}} f(x, \theta_k) \nu(dx) \quad \text{and}$$

$$E_{\theta_k}[H(X)] = \int_{\mathcal{X}} H(x) f(x, \theta_k) \nu(dx).$$

3.1. Algorithm Description

Algorithm MRAS₀—Exact Version

• **Initialization:** Specify $\rho \in (0, 1]$, a small number $\varepsilon \geq 0$, a strictly increasing function $S(\cdot): \mathfrak{R} \rightarrow \mathfrak{R}^+$, and an initial p.d.f./p.m.f. $f(x, \theta_0) > 0 \forall x \in \mathcal{X}$. Set the iteration counter $k = 0$.

• **Repeat until a specified stopping rule is satisfied:**

Step 1. Calculate the $(1 - \rho)$ -quantile

$$\gamma_{k+1} := \sup_l \{l: P_{\theta_k}(H(X) \geq l) \geq \rho\}.$$

Step 2. **If** $k = 0$, **then** set $\bar{\gamma}_{k+1} = \gamma_{k+1}$.

elseif $k \geq 1$

if $\gamma_{k+1} \geq \bar{\gamma}_k + \varepsilon$, **then** set $\bar{\gamma}_{k+1} = \gamma_{k+1}$.

else set $\bar{\gamma}_{k+1} = \bar{\gamma}_k$.

endif

endif

Step 3. Compute the parameter vector θ_{k+1} as

$$\theta_{k+1} := \arg \max_{\theta \in \Theta} E_{\theta_k} \left[\frac{[S(H(X))]^k}{f(X, \theta_k)} I_{\{H(X) \geq \bar{\gamma}_{k+1}\}} \ln f(X, \theta) \right]. \quad (3)$$

Step 4. Set $k = k + 1$.

The MRAS₀ algorithm requires specification of a parameter ρ , which determines the proportion of samples that will be used to update the probabilistic model. At successive iterations of the algorithm, a sequence $\{\gamma_k, k = 1, 2, \dots\}$, i.e., the $(1 - \rho)$ -quantiles with respect to the sequence of p.d.f.'s/p.m.f.'s $\{f(\cdot, \theta_k)\}$, are calculated at Step 1 of MRAS₀. These quantile values are then used in Step 2 to construct a sequence of nondecreasing thresholds $\{\bar{\gamma}_k, k = 1, 2, \dots\}$; only those candidate solutions that have performances better than these thresholds will be used in parameter updating (cf. (3)). As we will see, the theoretical convergence of MRAS₀ is unaffected by the value of the parameter ρ . We use ρ in our approach to concentrate the computational effort on the set of elite/promising samples, which is a standard technique employed in most of the population-based approaches, like GAs and EDAs.

During the initialization step of MRAS₀, a small number ε and a strictly increasing function $S(\cdot): \mathfrak{R} \rightarrow \mathfrak{R}^+$ are also specified. The function $S(\cdot)$ is used to account for the cases where the values of $H(x)$ are negative for some x , and the parameter ε ensures that each strict increment in the sequence $\{\bar{\gamma}_k\}$ is lower bounded, i.e.,

$$\inf_{\substack{\bar{\gamma}_{k+1} \neq \bar{\gamma}_k \\ k=1,2,\dots}} (\bar{\gamma}_{k+1} - \bar{\gamma}_k) \geq \varepsilon.$$

We require ε to be strictly positive for continuous problems and nonnegative for discrete (finite) problems.

In continuous domains, the division by $f(x, \theta_k)$ in the performance function in Step 3 is well defined if $f(x, \theta_k)$ has infinite support (e.g., normal p.d.f.), whereas in discrete/combinatorial domains, the division is still valid as long as each point x in the solution space has a positive probability of being sampled. Additional regularity conditions on $f(x, \theta_k)$ in §5 will ensure that Step 3 of MRAS₀ can be used interchangeably with the following equation:

$$\theta_{k+1} = \arg \max_{\theta \in \Theta} \int_{\mathcal{X}} [S(H(x))]^k I_{\{H(x) \geq \bar{\gamma}_{k+1}\}} \ln f(x, \theta) \nu(dx).$$

The following lemma shows that there is a sequence of reference models $\{g_k(\cdot), k = 1, 2, \dots\}$ implicit in MRAS₀, and the parameter θ_{k+1} computed at Step 3 indeed minimizes the KL-divergence $\mathcal{D}(g_{k+1}, f(\cdot, \theta))$.

LEMMA 1. *The parameter θ_{k+1} computed at the k th iteration of the MRAS₀ algorithm minimizes the KL-divergence $\mathcal{D}(g_{k+1}, f(\cdot, \theta))$, where*

$$g_{k+1}(x) := \frac{S(H(x)) I_{\{H(x) \geq \bar{\gamma}_{k+1}\}} g_k(x)}{E_{g_k} [S(H(X)) I_{\{H(X) \geq \bar{\gamma}_{k+1}\}}]} \quad \forall x \in \mathcal{X}, k = 1, 2, \dots, \quad \text{and}$$

$$g_1(x) := \frac{I_{\{H(x) \geq \bar{\gamma}_1\}}}{E_{\theta_0} [I_{\{H(X) \geq \bar{\gamma}_1\}} / f(X, \theta_0)]}.$$

PROOF. The proof appears in the online appendix. \square

3.2. Global Convergence

Global convergence of the MRAS₀ algorithm clearly depends on the choice of the parameterized distribution family. Throughout this paper, we restrict our analysis and discussions to a particular family of distributions called the natural exponential family (NEF) (cf., e.g., Morris 1982) for which the global convergence properties can be established. We start by stating the definition of NEF and some regularity conditions.

DEFINITION 1. A parameterized family of p.d.f.'s/p.m.f.'s $\{f(\cdot, \theta), \theta \in \Theta \subseteq \mathfrak{R}^m\}$ on \mathcal{X} is said to belong to the NEF if there exist functions $h(\cdot): \mathfrak{R}^n \rightarrow \mathfrak{R}$, $\Gamma(\cdot): \mathfrak{R}^n \rightarrow \mathfrak{R}^m$, and $K(\cdot): \mathfrak{R}^m \rightarrow \mathfrak{R}$ such that

$$f(x, \theta) = \exp\{\theta^T \Gamma(x) - K(\theta)\} h(x) \quad \forall \theta \in \Theta, \quad (4)$$

where $K(\theta) = \ln \int_{x \in \mathcal{X}} \exp\{\theta^T \Gamma(x)\} h(x) \nu(dx)$, and the superscript T denotes the vector transposition. For the case where $f(\cdot, \theta)$ is a p.d.f., we assume that $\Gamma(\cdot)$ is a continuous mapping.

Many common p.d.f.'s/p.m.f.'s belong to the NEF, e.g., Gaussian, Poisson, binomial, geometric, and certain multivariate forms of them.

ASSUMPTION A3. *There exists a compact set Π such that the level set $\{x: H(x) \geq \bar{\gamma}_1\} \cap \mathcal{X} \subseteq \Pi$, where $\bar{\gamma}_1 = \sup_l \{l: P_{\theta_0}(H(X) \geq l) \geq \rho\}$ is defined as in the MRAS₀ algorithm.*

ASSUMPTION A4. *The maximizer of Equation (3) is an interior point of Θ for all k .*

ASSUMPTION A5. *$\sup_{\theta \in \Theta} \|\exp\{\theta^T \Gamma(x)\} \Gamma(x) h(x)\|$ is integrable/summable with respect to x , where θ , $\Gamma(\cdot)$, and $h(\cdot)$ are defined as in Definition 1.*

Assumption A3 restricts the search of the MRAS₀ algorithm to some compact set; it is satisfied if the function $H(\cdot)$ has compact level sets or the solution space \mathcal{X} is compact. In actual implementation of the algorithm, Step 3 of MRAS₀ is often posed as an unconstrained optimization problem, i.e., $\Theta = \mathfrak{R}^m$, in which case Assumption A4 is automatically satisfied. It is also easy to verify that Assumption A5 is satisfied by most NEFs.

To show the convergence of MRAS₀, we will need the following key observation.

LEMMA 2. *If Assumptions A3–A5 hold, then we have*

$$E_{\theta_{k+1}}[\Gamma(X)] = E_{g_{k+1}}[\Gamma(X)] \quad \forall k = 0, 1, \dots,$$

where $E_{\theta_{k+1}}[\cdot]$ and $E_{g_{k+1}}[\cdot]$ denote the expectations taken with respect to $f(\cdot, \theta_{k+1})$ and $g_{k+1}(\cdot)$, respectively.

PROOF. The proof appears in the online appendix. \square

We have the following convergence result for the MRAS₀ algorithm.

THEOREM 1. *Let $\{\theta_k, k = 1, 2, \dots\}$ be the sequence of parameters generated by MRAS₀. If $\varepsilon > 0$ and Assumptions A1–A5 are satisfied, then*

$$\lim_{k \rightarrow \infty} E_{\theta_k}[\Gamma(X)] = \Gamma(x^*), \quad (5)$$

where the limit is component-wise.

PROOF. We prove Theorem 1 in the appendix. \square

REMARK 1. Note that for many NEFs used in practice, $\Gamma(\cdot)$ is a one-to-one mapping, in which case the convergence result (5) can be equivalently written as $\Gamma^{-1}(\lim_{k \rightarrow \infty} E_{\theta_k}[\Gamma(X)]) = x^*$. Also, note that for some particular p.d.f.'s/p.m.f.'s, the solution vector x itself will be a component of $\Gamma(x)$ (e.g., multivariate normal distribution). Under these circumstances, we can interpret (5) as $\lim_{k \rightarrow \infty} E_{\theta_k}[X] = x^*$. Another special case of particular interest is when the components of the random vector $X = (X_1, \dots, X_n)$ are independent, i.e., each has a univariate p.d.f./p.m.f. of the form

$$f(x_i, \vartheta_i) = \exp(x_i \vartheta_i - K(\vartheta_i)) h(x_i), \quad \vartheta_i \in \mathfrak{R} \quad \forall i = 1, \dots, n.$$

In this case, because the distribution of the random vector X is simply the product of the marginal distributions, we will clearly have $\Gamma(x) = x$. Thus, (5) is again equivalent to $\lim_{k \rightarrow \infty} E_{\theta_k}[X] = x^*$, where $\theta_k := (\vartheta_1^k, \dots, \vartheta_n^k)$, and ϑ_i^k is the value of ϑ_i at the k th iteration.

REMARK 2. Note that, as mentioned in §2, for problems with finite solution spaces, Assumptions A1 and A2 are automatically satisfied. Furthermore, if we take the input parameter $\varepsilon = 0$, then Step 2 of MRAS₀ is equivalent to $\bar{\gamma}_{k+1} = \max_{1 \leq i \leq k+1} \gamma_i$. Thus, $\{\bar{\gamma}_k\}$ is nondecreasing, and each strict increment in the sequence is bounded from below by

$$\min_{\substack{H(x) \neq H(y) \\ x, y \in \mathcal{X}}} |H(x) - H(y)|.$$

Therefore, the $\varepsilon > 0$ assumption in Theorem 1 can be relaxed to $\varepsilon \geq 0$.

We now address some of the special cases discussed in Remark 1.

COROLLARY 1 (MULTIVARIATE NORMAL). *For continuous optimization problems in \mathfrak{R}^n , if multivariate normal p.d.f.'s are used in MRAS₀, i.e.,*

$$f(x, \theta_k) = \frac{1}{\sqrt{(2\pi)^n |\Sigma_k|}} \exp\left(-\frac{1}{2}(x - \mu_k)^T \Sigma_k^{-1} (x - \mu_k)\right), \quad (6)$$

where $\theta_k := (\mu_k; \Sigma_k)$, $\varepsilon > 0$, and Assumptions A1–A4 are satisfied, then

$$\lim_{k \rightarrow \infty} \mu_k = x^* \quad \text{and} \quad \lim_{k \rightarrow \infty} \Sigma_k = 0_{n \times n},$$

where $0_{n \times n}$ represents an n -by- n zero matrix.

PROOF. By Lemma 2, it is easy to show that

$$\mu_{k+1} = E_{g_{k+1}}(X) \quad \forall k = 0, 1, \dots,$$

and

$$\Sigma_{k+1} = E_{g_{k+1}}[(X - \mu_{k+1})(X - \mu_{k+1})^T] \quad \forall k = 0, 1, \dots$$

The rest of the proof amounts to showing that

$$\lim_{k \rightarrow \infty} E_{g_k}(X) = x^* \quad \text{and}$$

$$\lim_{k \rightarrow \infty} E_{g_k}[(X - \mu_k)(X - \mu_k)^T] = 0_{n \times n},$$

which is the same as the proof of Theorem 1. \square

REMARK 3. Corollary 1 shows that in the multivariate normal case, the sequence of parameterized p.d.f.'s will converge to a degenerate p.d.f. concentrated only on the optimal solution. In this case, the parameters are updated as

$$\mu_{k+1} = \frac{E_{\theta_k}[\{[S(H(X))]^k / f(X, \theta_k)\} I_{\{H(X) \geq \bar{\gamma}_{k+1}\}} X]}{E_{\theta_k}[\{[S(H(X))]^k / f(X, \theta_k)\} I_{\{H(X) \geq \bar{\gamma}_{k+1}\}}]} \quad (7)$$

and

$$\Sigma_{k+1} = \frac{E_{\theta_k}[\{[S(H(X))]^k / f(X, \theta_k)\} I_{\{H(X) \geq \bar{\gamma}_{k+1}\}} (X - \mu_{k+1})(X - \mu_{k+1})^T]}{E_{\theta_k}[\{[S(H(X))]^k / f(X, \theta_k)\} I_{\{H(X) \geq \bar{\gamma}_{k+1}\}}]}, \quad (8)$$

where $f(x, \theta_k)$ is given by (6). Note that when the solution space \mathcal{X} is a (simple) constrained region in \mathfrak{R}^n , one straightforward approach is to use the acceptance-rejection method (cf. Kroese et al. 2006), and it is easy to verify that the parameter updating rules remain the same.

COROLLARY 2 (INDEPENDENT UNIVARIATE). *If the components of the random vector $X = (X_1, \dots, X_n)$ are independent, each has a univariate p.d.f./p.m.f. of the form*

$$f(x_i, \vartheta_i) = \exp(x_i \vartheta_i - K(\vartheta_i))h(x_i), \quad \vartheta_i \in \mathfrak{R} \quad \forall i = 1, \dots, n,$$

$\varepsilon > 0$, and Assumptions A1–A5 are satisfied, then

$$\lim_{k \rightarrow \infty} E_{\theta_k} [X] = x^*, \quad \text{where } \theta_k := (\vartheta_1^k, \dots, \vartheta_n^k).$$

4. An Alternative View of the Cross-Entropy Method

In this section, we give an alternative interpretation of the CE method for optimization and discuss its similarities and differences with the MRAS₀ algorithm. Specifically, we show that the CE method can also be viewed as a search strategy guided by a sequence of reference models. From this particular point of view, we establish some important properties of the CE method.

The deterministic version of the CE method for solving (1) can be summarized as follows.

Algorithm CE₀: Deterministic Version of the CE Method

Step 1. Choose the initial p.d.f./p.m.f. $f(\cdot, \theta_0)$, $\theta_0 \in \Theta$. Specify the parameter $\rho \in (0, 1]$ and a nondecreasing function $\varphi(\cdot): \mathfrak{R} \rightarrow \mathfrak{R}^+ \cup \{0\}$. Set $k = 0$.

Step 2. Calculate the $(1 - \rho)$ -quantile γ_{k+1} as

$$\gamma_{k+1} := \sup\{l: P_{\theta_k}(H(X) \geq l) \geq \rho\}.$$

Step 3. Compute the new parameter

$$\theta_{k+1} := \arg \max_{\theta \in \Theta} E_{\theta_k} [\varphi(H(X))I_{\{H(X) \geq \gamma_{k+1}\}} \ln f(X, \theta)].$$

Step 4. If a stopping rule is satisfied, then terminate; otherwise set $k = k + 1$ and go to Step 2.

In CE₀, choosing $\varphi(H(x)) = 1$ gives the standard CE method, whereas choosing $\varphi(H(x)) = H(x)$ (if $H(x) \geq 0 \forall x \in \mathcal{X}$) gives an extended version of the standard CE method (cf. De Boer et al. 2005).

One resemblance between CE and MRAS₀ is the use of the parameter ρ and the $(1 - \rho)$ -quantile in both algorithms. However, the fundamental difference is that in CE, the problem of estimating the optimal value of the parameter is broken down into a sequence of simple estimation problems, in which the parameter ρ assumes a crucial role. Because a small change in the values of ρ may disturb the whole estimation process and affect the quality of the resulting estimates, the convergence of CE cannot be always guaranteed unless the value of ρ is chosen sufficiently small (cf. De Boer et al. 2005, Homem-de-Mello 2007; also see Example 4.1 below), whereas the theoretical convergence of MRAS₀ is unaffected by the parameter ρ .

The following lemma provides a unified view of MRAS and CE; it shows that by appropriately defining a sequence of implicit reference models $\{g_k^{ce}(\cdot): k = 1, 2, \dots\}$, the CE method can be recovered, and the parameter updating in CE is guided by this sequence of models.

LEMMA 3. *The parameter θ_{k+1} computed at the k th iteration of the CE₀ algorithm minimizes the KL-divergence $\mathcal{D}(g_{k+1}^{ce}, f(\cdot, \theta))$, where*

$$g_{k+1}^{ce}(x) := \frac{\varphi(H(x))I_{\{H(x) \geq \gamma_{k+1}\}}f(x, \theta_k)}{E_{\theta_k}[\varphi(H(X))I_{\{H(X) \geq \gamma_{k+1}\}}]} \quad \forall x \in \mathcal{X}, k = 0, 1, \dots \quad (9)$$

PROOF. Similar to the proof of Lemma 1. \square

The key observation to note is that in contrast to MRAS₀, the sequence of reference models in CE depends explicitly on the family of parameterized p.d.f.’s/p.m.f.’s $\{f(\cdot, \theta_k)\}$ used. Because $g_{k+1}^{ce}(\cdot)$ is obtained by tilting $f(\cdot, \theta_k)$ with the performance function, it improves the expected performance in the sense that

$$E_{g_{k+1}^{ce}}[\varphi(H(X))I_{\{H(X) \geq \gamma_{k+1}\}}] = \frac{E_{\theta_k}[(\varphi(H(X))I_{\{H(X) \geq \gamma_{k+1}\}})^2]}{E_{\theta_k}[\varphi(H(X))I_{\{H(X) \geq \gamma_{k+1}\}}]} \geq E_{\theta_k}[\varphi(H(X))I_{\{H(X) \geq \gamma_{k+1}\}}].$$

Thus, it is reasonable to expect that the projection of $g_{k+1}^{ce}(\cdot)$ on $\{f(\cdot, \theta): \theta \in \Theta\}$ (i.e., $f(\cdot, \theta_{k+1})$) also improves the expected performance. This result is formalized in the following theorem, whose proof is given in the online appendix.

THEOREM 2. *For the CE₀ algorithm, we have*

$$E_{\theta_{k+1}}[\varphi(H(X))I_{\{H(X) \geq \gamma_{k+1}\}}] \geq E_{\theta_k}[\varphi(H(X))I_{\{H(X) \geq \gamma_{k+1}\}}] \quad \forall k = 0, 1, \dots$$

In the standard CE method, Theorem 2 implies the monotonicity of the sequence $\{\gamma_k: k = 1, 2, \dots\}$.

LEMMA 4. *For the standard CE method (i.e., CE₀ with $\varphi(H(x)) = 1$), we have*

$$\gamma_{k+2} \geq \gamma_{k+1} \quad \forall k = 0, 1, \dots$$

PROOF. By Theorem 2, we have

$$E_{\theta_{k+1}}[I_{\{H(X) \geq \gamma_{k+1}\}}] \geq E_{\theta_k}[I_{\{H(X) \geq \gamma_{k+1}\}}],$$

i.e.,

$$P_{\theta_{k+1}}(H(X) \geq \gamma_{k+1}) \geq P_{\theta_k}(H(X) \geq \gamma_{k+1}) \geq \rho.$$

The result follows by the definition of γ_{k+2} (see Step 2 of the CE₀ algorithm). \square

Note that because $\gamma_k \leq H(x^*)$ for all k , Lemma 4 implies that the sequence $\{\gamma_k: k = 1, 2, \dots\}$ generated by the standard CE method converges. However, depending on the p.d.f.’s/p.m.f.’s and the parameter ρ used, the sequence $\{\gamma_k\}$ may not converge to $H(x^*)$ or even to a small neighborhood of $H(x^*)$ (cf. Examples 4.1 and 4.2 below).

Similar to MRAS₀, when $f(\cdot, \theta)$ belongs to the natural exponential families, the following lemma relates the sequence $\{f(\cdot, \theta_k), k = 1, 2, \dots\}$ to the sequence of reference models $\{g_k^{ce}(\cdot): k = 1, 2, \dots\}$.

LEMMA 5. Assume that:

(1) There exists a compact set $\bar{\Pi}$ such that the level set $\{x: H(x) \geq \gamma_k\} \cap \mathcal{X} \subseteq \bar{\Pi}$ for all $k = 1, 2, \dots$, where $\gamma_k = \sup_l \{l: P_{\theta_{k-1}}(H(X) \geq l) \geq \rho\}$ is defined as in the CE₀ algorithm.

(2) The parameter θ_{k+1} computed at Step 3 of the CE₀ algorithm is an interior point of Θ for all k .

(3) Assumption A5 is satisfied.

Then,

$$E_{\theta_{k+1}}[\Gamma(X)] = E_{g_k^{ce}}[\Gamma(X)] \quad \forall k = 0, 1, \dots$$

PROOF. Similar to the proof of Lemma 2. \square

The above lemma indicates that the behavior of the sequence of p.d.f.'s/p.m.f.'s $\{f(\cdot, \theta_k)\}$ is closely related to the properties of the sequence of reference models. To understand this, consider the particular case where $\Gamma(x) = x$. If the CE method converges to the optimal solution in the sense that $\lim_{k \rightarrow \infty} E_{\theta_k}[H(X)] = H(x^*)$, then we must have $\lim_{k \rightarrow \infty} E_{\theta_k}[X] = x^*$ because $H(x) < H(x^*) \forall x \neq x^*$. Thus, by Lemma 5, a necessary condition for this convergence is $\lim_{k \rightarrow \infty} E_{g_k^{ce}}[X] = x^*$. However, unlike MRAS₀, where the convergence of the sequence of reference models to an optimal degenerate distribution is guaranteed, the convergence of the sequence $\{g_k^{ce}(\cdot): k = 1, 2, \dots\}$ relies on the choices of the families of distributions $\{f(\cdot, \theta)\}$ and the values of the parameter ρ used (cf. (9)). We now illustrate this issue by two simple examples.

EXAMPLE 4.1 (THE STANDARD CE METHOD). Consider maximizing the function $H(x)$ given by

$$H(x) = \begin{cases} 0, & x \in \{(0, 1), (1, 0)\}, \\ 1, & x = (0, 0), \\ a, & x = (1, 1), \end{cases} \quad (10)$$

where $a > 1$, and $x := (x_1, x_2) \in \mathcal{X} := \{(0, 0), (0, 1), (1, 0), (1, 1)\}$.

If we take $0.25 < \rho \leq 0.5$ and an initial p.m.f.

$$f(x, \theta_0) = p_0^{x_1} (1 - p_0)^{1-x_1} q_0^{x_2} (1 - q_0)^{1-x_2} \quad \text{with } \theta_0 = (p_0, q_0) = (0.5, 0.5),$$

then because $P_{\theta_0}(x \in \{(0, 0), (1, 1)\}) = 0.5 \geq \rho$, we have $\gamma_1 = 1$. It is also straightforward to see that

$$g_1^{ce}(x) = \begin{cases} 0.5, & x = (0, 0) \text{ or } (1, 1), \\ 0, & \text{otherwise,} \end{cases}$$

and the parameter θ_1 computed at Step 3 (with $\varphi(H(x)) = 1$) of CE₀ is given by $\theta_1 = (0.5, 0.5)$. Proceeding iteratively, we have $\gamma_k = 1$ and $g_k^{ce}(x) = g_1^{ce}(x) \forall k = 1, 2, \dots$, i.e., the algorithm does not converge to a degenerate distribution at the optimal solution.

On the other hand, if we choose $\rho \leq 0.25$, then it turns out that $\gamma_k = a$ and

$$g_k^{ce}(x) = \begin{cases} 1, & x = (1, 1), \\ 0, & \text{otherwise,} \end{cases}$$

for all $k = 1, 2, \dots$, which means the algorithm converges to the optimum.

EXAMPLE 4.2 (THE EXTENDED VERSION OF THE CE METHOD). Consider solving problem (10) by CE₀ with the performance function $\varphi(H(x)) = H(x)$. We use the same family of p.m.f.'s as in Example 4.1 with the initial parameter $\theta_0 = (1/(1+a), 1/(1+a))$. If the values of ρ are chosen from the interval $(1/(1+a)^2, (a^2+1)/(1+a)^2)$, then we have $\theta_k = (1/(1+a), 1/(1+a))$, $\gamma_k = 1$, and

$$g_k^{ce}(x) = \begin{cases} \frac{a}{1+a}, & x = (0, 0), \\ \frac{1}{1+a}, & x = (1, 1), \\ 0, & \text{otherwise,} \end{cases}$$

for all $k = 1, 2, \dots$.

On the other hand, if we choose $\rho = 0.5$ and $\theta_0 = (0.5, 0.5)$, then it is easy to verify that $\lim_{k \rightarrow \infty} \gamma_k = a$ and

$$\lim_{k \rightarrow \infty} g_k^{ce}(x) = \begin{cases} 1, & x = (1, 1), \\ 0, & \text{otherwise.} \end{cases}$$

5. The MRAS₁ Algorithm (Monte Carlo Version)

The MRAS₀ algorithm describes the idealized situation where quantile values and expectations can be evaluated exactly. In practice, we will usually resort to its stochastic counterpart, where only a finite number of samples are used and expected values are replaced with their corresponding sample averages. For example, Step 3 of MRAS₀ will be replaced with

$$\tilde{\theta}_{k+1} = \arg \max_{\theta \in \Theta} \frac{1}{N} \sum_{i=1}^N \frac{[S(H(X_i))]^k}{f(X_i, \tilde{\theta}_k)} I_{\{H(X_i) \geq \tilde{\gamma}_{k+1}\}} \ln f(X_i, \theta), \quad (11)$$

where X_1, \dots, X_N are i.i.d. random samples generated from $f(x, \tilde{\theta}_k)$, $\tilde{\theta}_k$ is the estimated parameter vector computed at the previous iteration, and $\tilde{\gamma}_{k+1}$ is a threshold determined by the sample $(1-\rho)$ -quantile of $H(X_1), \dots, H(X_N)$.

However, the theoretical convergence can no longer be guaranteed for a simple stochastic counterpart of MRAS₀. In particular, the set $\{x: H(x) \geq \tilde{\gamma}_{k+1}, x \in \{X_1, \dots, X_N\}\}$ involved in (11) may be empty because all the random samples generated at the current iteration may be much worse than those generated at the previous iteration. Thus, we can only expect the algorithm to converge if the expected

values in the MRAS₀ algorithm are closely approximated. Obviously, the quality of the approximation will depend on the number of samples to be used in the simulation, but it is difficult to determine in advance the appropriate number of samples. A sample size too small will cause the algorithm to fail to converge and result in poor quality solutions, whereas a sample size too large may lead to high computational cost.

As mentioned earlier, the parameter ρ , to some extent, will affect the performance of the algorithm. Large values of ρ mean that almost all samples generated, regardless of their performances, will be used to update the probabilistic model, which could slow down the convergence process. On the other hand, because a good estimate will necessarily require a reasonable number of valid samples, the quantity ρN (i.e., the approximate number of samples that will be used in parameter updating) cannot be too small. Thus, small values of ρ will require a large number of samples to be generated at each iteration and may result in significant simulation efforts. For a given problem, although it is clear that we should avoid those values of ρ that are either too close to one or too close to zero, it may be difficult to determine a priori which ρ gives a satisfactory performance.

To address the above difficulties, we adopt the same idea as in Homem-de-Mello (2007) and propose a modified Monte Carlo version of MRAS₀ in which the sample size N is adaptively increasing and the parameter ρ is adaptively decreasing.

5.1. Algorithm Description

Roughly speaking, the MRAS₁ algorithm is essentially a Monte Carlo version of MRAS₀ except that the parameter ρ and the sample size N may change from one iteration to another. The rate of increase in the sample size is controlled by an extra parameter $\alpha > 1$, specified during the initialization step. For example, if the initial sample size is N_0 , then after k increments, the sample size will be approximately $\lceil \alpha^k N_0 \rceil$.

At each iteration k , random samples are drawn from the density/mass function $\tilde{f}(\cdot, \tilde{\theta}_k)$, which is a mixture of the initial density/mass $f(\cdot, \theta_0)$ and the density/mass calculated from the previous iteration $f(\cdot, \tilde{\theta}_k)$ (cf. Auer et al. 2002 for a similar idea in the context of multiarmed bandit models). We assume that $f(\cdot, \theta_0)$ satisfies the following condition:

ASSUMPTION A3'. *There exists a compact set Π_ε such that $\{x: H(x) \geq H(x^*) - \varepsilon\} \cap \mathcal{X} \subseteq \Pi_\varepsilon$. Moreover, the initial density/mass function $f(x, \theta_0)$ is bounded away from zero on Π_ε , i.e., $f_* := \inf_{x \in \Pi_\varepsilon} f(x, \theta_0) > 0$.*

Algorithm MRAS₁—Monte Carlo version

• **Initialization:** Specify $\rho_0 \in (0, 1]$, an initial sample size $N_0 > 1$, $\varepsilon \geq 0$, $\alpha > 1$, a mixing coefficient $\lambda \in (0, 1]$, a strictly increasing function $S(\cdot): \Re \rightarrow \Re^+$, and an initial p.d.f. $f(x, \theta_0) > 0 \forall x \in \mathcal{X}$. Set $\theta_0 \leftarrow \theta_0, k \leftarrow 0$.

• **Repeat until a specified stopping rule is satisfied:**

Step 1. Generate N_k i.i.d. samples $X_1^k, \dots, X_{N_k}^k$ according to $\tilde{f}(\cdot, \tilde{\theta}_k) := (1 - \lambda)f(\cdot, \tilde{\theta}_k) + \lambda f(\cdot, \theta_0)$.

Step 2. Compute the sample $(1 - \rho_k)$ -quantile $\tilde{\gamma}_{k+1}(\rho_k, N_k) := H_{(\lceil (1-\rho_k)N_k \rceil)}$, where $\lceil a \rceil$ is the smallest integer greater than a , and $H_{(i)}$ is the i th order statistic of the sequence $\{H(X_i^k), i = 1, \dots, N_k\}$.

Step 3. **If $k = 0$ or $\tilde{\gamma}_{k+1}(\rho_k, N_k) \geq \bar{\gamma}_k + \varepsilon/2$, then**

3(a). Set $\bar{\gamma}_{k+1} \leftarrow \tilde{\gamma}_{k+1}(\rho_k, N_k), \rho_{k+1} \leftarrow \rho_k, N_{k+1} \leftarrow N_k$.

else, find the largest $\bar{\rho} \in (0, \rho_k)$ such that

$\tilde{\gamma}_{k+1}(\bar{\rho}, N_k) \geq \bar{\gamma}_k + \varepsilon/2$.

3(b). **If such a $\bar{\rho}$ exists, then set**

$\bar{\gamma}_{k+1} \leftarrow \tilde{\gamma}_{k+1}(\bar{\rho}, N_k), \rho_{k+1} \leftarrow \bar{\rho},$

$N_{k+1} \leftarrow N_k$.

3(c). **else (if no such $\bar{\rho}$ exists), set $\bar{\gamma}_{k+1} \leftarrow \bar{\gamma}_k,$**

$\rho_{k+1} \leftarrow \rho_k, N_{k+1} \leftarrow \lceil \alpha N_k \rceil$.

endif

Step 4. Compute $\tilde{\theta}_{k+1}$ as

$$\tilde{\theta}_{k+1} = \arg \max_{\theta \in \Theta} \frac{1}{N_k} \sum_{i=1}^{N_k} \frac{[S(H(X_i^k))]^k}{\tilde{f}(X_i^k, \tilde{\theta}_k)} \cdot I_{\{H(X_i^k) \geq \bar{\gamma}_{k+1}\}} \ln f(X_i^k, \theta). \quad (12)$$

Step 5. Set $k \leftarrow k + 1$.

In practice, the initial density $f(\cdot, \theta_0)$ can be chosen according to some prior knowledge of the problem structure (cf. §6.2); however, if nothing is known about where the good solutions are, one simple choice of $f(\cdot, \theta_0)$ is the uniform distribution. Intuitively, mixing in the initial density forces the algorithm to explore the entire solution space and to maintain a global perspective during the search process. Also note that if $\lambda = 1$, then random samples will always be drawn from the initial density, in which case MRAS₁ becomes a pure random sampling approach.

At Step 2, the sample $(1 - \rho_k)$ -quantile $\tilde{\gamma}_{k+1}$ is calculated by first ordering the sample performances $H(X_i^k), i = 1, \dots, N_k$, from smallest to largest, $H_{(1)} \leq H_{(2)} \leq \dots \leq H_{(N_k)}$, and then taking the $\lceil (1 - \rho_k)N_k \rceil$ th order statistic. We use the function $\tilde{\gamma}_{k+1}(\rho_k, N_k)$ to emphasize the dependencies of $\tilde{\gamma}_{k+1}$ on both ρ_k and N_k , so that different sample quantile values used during one iteration can be distinguished by their arguments.

Step 3 of MRAS₁ is used to extract a sequence of non-decreasing thresholds $\{\bar{\gamma}_k, k = 1, 2, \dots\}$ from the sequence of sample quantiles $\{\tilde{\gamma}_k\}$, and to determine the appropriate values of ρ_{k+1} and N_{k+1} to be used in subsequent iterations. This step is carried out as follows. At each iteration k , we first check whether the inequality $\tilde{\gamma}_{k+1}(\rho_k, N_k) \geq \bar{\gamma}_k + \varepsilon/2$ is satisfied, where $\bar{\gamma}_k$ is the threshold value used in the previous iteration. If the inequality holds, then it means that both the current ρ_k value and the current sample size N_k are satisfactory; thus we proceed to Step 3(a) and update the parameter vector $\tilde{\theta}_{k+1}$ in Step 4 by using $\tilde{\gamma}_{k+1}(\rho_k, N_k)$. Otherwise, it indicates that either ρ_k is too large or the

sample size N_k is too small. To determine which, we fix the sample size N_k and check if there exists a smaller $\bar{\rho} < \rho_k$ such that the above inequality can be satisfied with the new sample $(1 - \bar{\rho})$ -quantile. If such a $\bar{\rho}$ does exist, then the current sample size N_k is still deemed acceptable, and we only need to decrease the ρ_k value. Accordingly, the parameter vector is updated in Step 4 by using the sample $(1 - \bar{\rho})$ -quantile. On the other hand, if no such $\bar{\rho}$ can be found, then the parameter vector is updated by using the threshold $\bar{\gamma}_k$ calculated during the previous iteration and the sample size N_k is increased by a factor α .

We make the following assumption about the parameter vector $\tilde{\theta}_{k+1}$ computed at Step 4.

ASSUMPTION A4'. The parameter vector $\tilde{\theta}_{k+1}$ computed at Step 4 of MRAS₁ is an interior point of Θ for all k .

It is important to note that the set $\{x: H(x) \geq \bar{\gamma}_{k+1}, x \in \{X_1^k, \dots, X_{N_k}^k\}\}$ could be empty if Step 3(c) is visited. If this happens, the right-hand side of (12) will be equal to zero, so any $\theta \in \Theta$ is a maximizer, and we define $\tilde{\theta}_{k+1} := \tilde{\theta}_k$ in this case.

5.2. Global Convergence

In this section, we discuss the convergence properties of the MRAS₁ algorithm for NEFs. To be specific, we will explore the relations between MRAS₁ and MRAS₀ and show that with high probability, the gaps (e.g., approximation errors incurred by replacing expected values with sample averages) between the two algorithms can be made small enough such that the convergence analysis of MRAS₁ can be ascribed to the convergence analysis of the MRAS₀ algorithm; thus, our analysis relies heavily on the results obtained in §3.2. Throughout this section, we denote by $P_{\tilde{\theta}_k}(\cdot)$ and $E_{\tilde{\theta}_k}[\cdot]$ the respective probability and expectation taken with respect to the p.d.f./p.m.f. $f(\cdot, \tilde{\theta}_k)$, and $\tilde{P}_{\tilde{\theta}_k}(\cdot)$ and $\tilde{E}_{\tilde{\theta}_k}[\cdot]$ the respective probability and expectation taken with respect to $\tilde{f}(\cdot, \tilde{\theta}_k)$. Note that because the sequence $\{\tilde{\theta}_k\}$ results from random samples generated at each iteration of MRAS₁, these quantities are also random.

Let $\tilde{g}_{k+1}(\cdot)$, $k = 0, 1, \dots$, be defined by

$$\tilde{g}_{k+1}(x) := \begin{cases} \frac{[[S(H(x))]^k / \tilde{f}(x, \tilde{\theta}_k)] I_{\{H(x) \geq \bar{\gamma}_{k+1}\}}}{\sum_{i=1}^{N_k} [[S(H(X_i^k))]^k / \tilde{f}(X_i^k, \tilde{\theta}_k)] I_{\{H(X_i^k) \geq \bar{\gamma}_{k+1}\}}} & \text{if } \{x: H(x) \geq \bar{\gamma}_{k+1}, x \in \{X_1^k, \dots, X_{N_k}^k\}\} \neq \emptyset, \\ \tilde{g}_k(x), & \text{otherwise,} \end{cases} \quad (13)$$

where $\bar{\gamma}_{k+1}$ is given by

$$\bar{\gamma}_{k+1} := \begin{cases} \tilde{\gamma}_{k+1}(\rho_k, N_k) & \text{if Step 3(a) is visited,} \\ \tilde{\gamma}_{k+1}(\bar{\rho}, N_k) & \text{if Step 3(b) is visited,} \\ \bar{\gamma}_k & \text{if Step 3(c) is visited.} \end{cases}$$

The following lemma shows the connection between $f(\cdot, \tilde{\theta}_{k+1})$ and $\tilde{g}_{k+1}(\cdot)$, whose proof is similar to the proof of Lemma 2, and is thus omitted here.

LEMMA 6. If Assumptions A4' and A5 hold, then the parameter $\tilde{\theta}_{k+1}$ computed at Step 3 of MRAS₁ satisfies

$$E_{\tilde{\theta}_{k+1}}[\Gamma(X)] = E_{\tilde{g}_{k+1}}[\Gamma(X)] \quad \forall k = 0, 1, \dots$$

Note that the region $\{x: H(x) \geq \bar{\gamma}_{k+1}\}$ will become smaller and smaller as $\bar{\gamma}_{k+1}$ increases. Lemma 6 shows that the sequence of sampling p.d.f.'s/p.m.f.'s $\{f(\cdot, \tilde{\theta}_{k+1})\}$ is adapted to this sequence of shrinking regions. For example, consider the case where $\{x: H(x) \geq \bar{\gamma}_{k+1}\}$ is convex and $\Gamma(x) = x$. Because $E_{\tilde{g}_{k+1}}[X]$ is the convex combination of $X_1^k, \dots, X_{N_k}^k$, the lemma implies that $E_{\tilde{\theta}_{k+1}}[X] \in \{x: H(x) \geq \bar{\gamma}_{k+1}\}$. Thus, it is natural to expect that the random samples generated at the next iteration will fall in the region $\{x: H(x) \geq \bar{\gamma}_{k+1}\}$ with large probabilities (e.g., consider the normal p.d.f. where its mode is equal to its mean). In contrast, if we use a fixed sampling distribution for all iterations as in pure random sampling (i.e., the $\lambda = 1$ case), then sampling from this sequence of shrinking regions could become a substantially difficult problem in practice.

Next, we present a useful intermediate result, which shows the convergence of the quantile estimates when random samples are generated from a sequence of different distributions.

LEMMA 7. For any given $\rho^\dagger \in (0, 1)$, let γ_k^\dagger be the set of $(1 - \rho^\dagger)$ -quantiles of $H(X)$ with respect to the p.d.f./p.m.f. $f(\cdot, \theta_k)$, and let $\tilde{\gamma}_k^\dagger(\rho^\dagger, N_k)$ be the corresponding sample quantile of $H(X_1^k), \dots, H(X_{N_k}^k)$, where $\tilde{f}(\cdot, \tilde{\theta}_k)$ and N_k are defined as in MRAS₁, and $X_1^k, \dots, X_{N_k}^k$ are i.i.d. with common density $\tilde{f}(\cdot, \tilde{\theta}_k)$. Then, the distance from $\tilde{\gamma}_k^\dagger(\rho^\dagger, N_k)$ to γ_k^\dagger tends to zero as $k \rightarrow \infty$ w.p.1.

PROOF. The proof appears in the online appendix. \square

We are now ready to state the main theorem.

THEOREM 3. Let $\varepsilon > 0$, and define the ε -optimal set $\mathcal{C}_\varepsilon := \{x: H(x) \geq H(x^*) - \varepsilon\} \cap \mathcal{X}$. If Assumptions A1, A3', A4', and A5 are satisfied, then there exists a random variable \mathcal{H} such that w.p.1., $\mathcal{H} < \infty$, and

- (1) $\bar{\gamma}_k > H(x^*) - \varepsilon \quad \forall k \geq \mathcal{H}$.
- (2) $E_{\tilde{\theta}_{k+1}}[\Gamma(X)] \in \text{CONV}\{\Gamma(\mathcal{C}_\varepsilon)\} \quad \forall k \geq \mathcal{H}$, where $\text{CONV}\{\Gamma(\mathcal{C}_\varepsilon)\}$ indicates the convex hull of the set $\Gamma(\mathcal{C}_\varepsilon)$.

Furthermore, let β be a positive constant satisfying the condition that the set $\{x: S(H(x)) \geq 1/\beta\}$ has a strictly positive Lebesgue/counting measure. If Assumptions A1, A2, A3', A4', and A5 are all satisfied and $\alpha > (\beta S^*)^2$, where $S^* := S(H(x^*))$, then

- (3) $\lim_{k \rightarrow \infty} E_{\tilde{\theta}_k}[\Gamma(X)] = \Gamma(x^*)$ w.p.1.

PROOF. See the appendix. \square

REMARK 4. Roughly speaking, the second result can be understood as finite time ε -optimality. To see this, consider the special case where $H(x)$ is locally concave on

the set \mathcal{C}_ε . Let $x, y \in \mathcal{C}_\varepsilon$ and $\eta \in [0, 1]$ be arbitrary. By the definition of concavity, we will have $H(\eta x + (1 - \eta)y) \geq \eta H(x) + (1 - \eta)H(y) \geq H(x^*) - \varepsilon$, which implies that the set \mathcal{C}_ε is convex. If in addition $\Gamma(x)$ is also convex and one-to-one on \mathcal{C}_ε (e.g., multivariate normal p.d.f.), then $\text{CONV}\{\Gamma(\mathcal{C}_\varepsilon)\} = \Gamma(\mathcal{C}_\varepsilon)$. Thus, it follows that $\Gamma^{-1}(E_{\tilde{\theta}_{k+1}}[\Gamma(X)]) \in \mathcal{C}_\varepsilon \forall k \geq \mathcal{H}$ w.p.1.

The following results are now immediate.

COROLLARY 3 (MULTIVARIATE NORMAL). For continuous optimization problems in \mathfrak{R}^n , if multivariate normal p.d.f.'s are used in MRAS₁, i.e.,

$$f(x, \tilde{\theta}_k) = \frac{1}{\sqrt{(2\pi)^n |\tilde{\Sigma}_k|}} \exp\left(-\frac{1}{2}(x - \tilde{\mu}_k)^T \tilde{\Sigma}_k^{-1} (x - \tilde{\mu}_k)\right),$$

$\varepsilon > 0$, $\alpha > (\beta S^*)^2$, and Assumptions A1, A2, A3', and A4' are satisfied, then

$$\lim_{k \rightarrow \infty} \tilde{\mu}_k = x^* \quad \text{and} \quad \lim_{k \rightarrow \infty} \tilde{\Sigma}_k = 0_{n \times n} \quad \text{w.p.1.}$$

COROLLARY 4 (INDEPENDENT UNIVARIATE). If the components of the random vector $X = (X_1, X_2, \dots, X_n)$ are independent, each with a univariate p.d.f./p.m.f. of the form

$$f(x_i, \vartheta_i) = \exp(x_i \vartheta_i - K(\vartheta_i))h(x_i), \quad \vartheta_i \in \mathfrak{R} \quad \forall i = 1, \dots, n,$$

$\varepsilon > 0$, $\alpha > (\beta S^*)^2$, and Assumptions A1, A2, A3', A4', and A5 are satisfied, then

$$\lim_{k \rightarrow \infty} E_{\tilde{\theta}_k}[X] = x^* \quad \text{w.p.1,} \quad \text{where } \tilde{\theta}_k := (\vartheta_1^k, \dots, \vartheta_n^k).$$

6. Numerical Examples

In this section, we illustrate the performance of MRAS₁ on both continuous and combinatorial optimization problems. In the former case, we test the algorithm on various benchmark problems that are well known in global optimization and compare its performance to both the standard CE method and the SA algorithm. In the latter case, we apply the algorithm to several asymmetric traveling salesman problems (ATSP), which are typical representatives of NP-hard combinatorial optimization problems.

Because all examples considered in this section are minimization problems, whereas MRAS was presented in a maximization context, the following modifications are required: (i) $S(\cdot)$ needs to be initialized as a strictly decreasing function instead of strictly increasing. Throughout this section, we take $S(H(x)) := \exp\{-rH(x)\}$, where r is a positive constant. (ii) The sample $(1 - \rho)$ -quantile $\tilde{\gamma}_{k+1}$ will now be calculated by first ordering the sample performances $H(X_i^k), i = 1, \dots, N_k$, from largest to smallest, and then taking the $\lceil (1 - \rho)N_k \rceil$ th order statistic. (iii) We need to replace the “ \geq ” operator with “ \leq ” operator in (12). (iv) The inequalities at Step 3 need to be replaced

with $\tilde{\gamma}_{k+1}(\rho_k, N_k) \leq \bar{\gamma}_k - \varepsilon/2$ and $\tilde{\gamma}_{k+1}(\bar{\rho}, N_k) \leq \bar{\gamma}_k - \varepsilon/2$, respectively.

In actual implementation of MRAS₁ and the CE method, a smoothed parameter updating procedure (cf. De Boer et al. 2005) is used, i.e., first a smoothed parameter vector $\hat{\theta}_{k+1}$ is computed at each iteration k according to

$$\hat{\theta}_{k+1} := v\tilde{\theta}_{k+1} + (1 - v)\hat{\theta}_k \quad \forall k = 0, 1, \dots, \quad \text{and } \hat{\theta}_0 := \tilde{\theta}_0,$$

where $\tilde{\theta}_{k+1}$ is the parameter vector computed at Step 4 of MRAS₁, and $v \in (0, 1]$ is a smoothing parameter; then $f(x, \hat{\theta}_{k+1})$ (instead of $f(x, \tilde{\theta}_{k+1})$) is used in Step 1 to generate new samples. Although this modification will not affect the theoretical convergence results, it may substantially improve the numerical performance of the algorithm (cf. Rubinstein and Kroese 2006 for a discussion).

Another practical issue is that to obtain a valid estimate $\tilde{\theta}_{k+1}$ at each iteration of MRAS₁, we must make sure that enough samples are used in parameter updating. This can be achieved by using an additional parameter N_{\min} , and performing the update (12) only when the number of the elite samples (i.e., those samples having performances better than the threshold $\tilde{\gamma}_{k+1}$) is greater than N_{\min} . In effect, this is equivalent to searching $\bar{\rho}$ from (ρ_{\min}, ρ_k) instead of $(0, \rho_k)$ at Step 3 of MRAS₁, where $\rho_{\min} := N_{\min}/N_k \rightarrow 0$ as $k \rightarrow \infty$.

6.1. Continuous Optimization

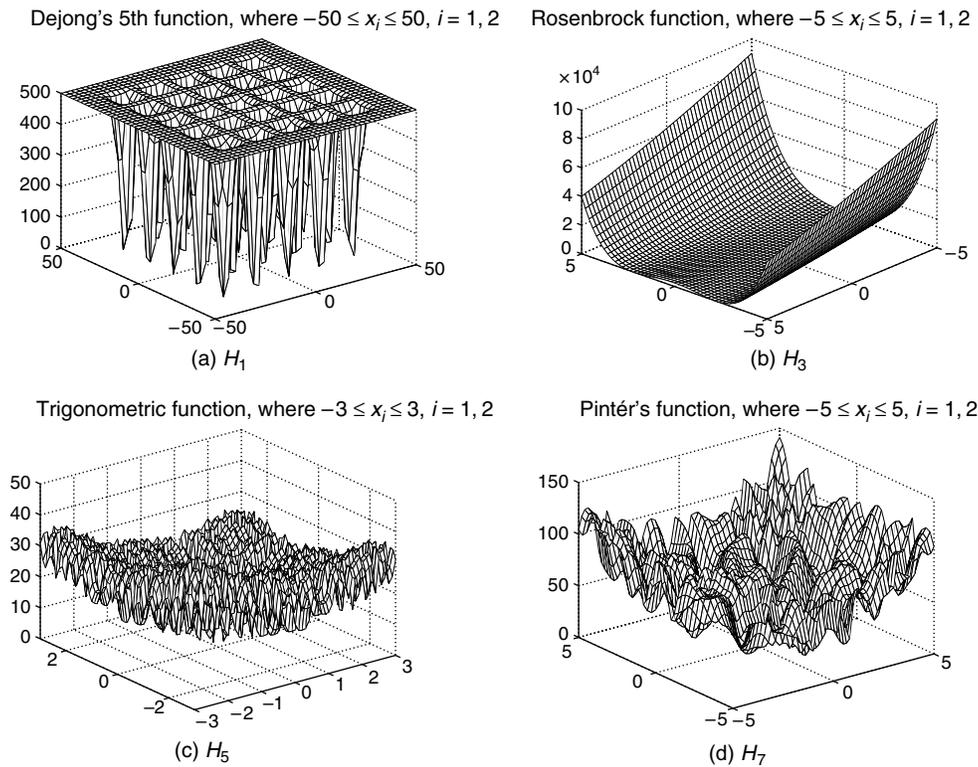
In our preliminary experiments, we take the family of parameterized p.d.f.'s to be multivariate normal p.d.f.'s. Initially, a mean vector μ_0 and a covariance matrix Σ_0 are specified; then at each iteration k of the algorithm, new parameters $\tilde{\mu}_{k+1}$ and $\tilde{\Sigma}_{k+1}$ are updated according to the respective stochastic counterparts of Equations (7) and (8). By Corollary 3, the sequence of mean vectors $\{\tilde{\mu}_k\}$ will converge to the optimal solution x^* , and the sequence of covariance matrices $\{\tilde{\Sigma}_k\}$ to the zero matrix.

The following benchmark problems, which have been previously studied, e.g., in Corana et al. (1987), Pintér (1996), Yao and Liu (1996), and Kroese et al. (2006), are used in our experiments. Functions H_1 and H_2 are low-dimensional problems that have only a few local optima; however, the minima are separated by plateaus and are relatively far apart. Functions H_3 and H_4 are 20-dimensional badly-scaled problems. Functions H_5 and H_6 are highly multimodal, and the number of local optima increases exponentially with the problem dimension. Function H_7 is both badly scaled and highly multimodal. The graphical representations of some of these functions in two dimensions are plotted in Figure 1.

- (1) Dejong's 5th function ($n = 2$),

$$H_1(x) = \left[0.002 + \sum_{j=1}^{25} \frac{1}{j + \sum_{i=1}^2 (x_i - a_{j,i})^6} \right]^{-1},$$

Figure 1. Selected test problems in two dimensions, (a) H_1 : Dejong’s 5th; (b) H_3 : Rosenbrock; (c) H_5 : Trigonometric; (d) H_7 : Pintér.



where $a_{j,1} = \{-32, -16, 0, 16, 32, -32, -16, 0, 16, 32, -32, -16, 0, 16, 32, -32, -16, 0, 16, 32, -32, -16, 0, 16, 32, -32, -16, 0, 16, 32\}$,

$a_{j,2} = \{-32, -32, -32, -32, -32, -16, -16, -16, -16, -16, 0, 0, 0, 0, 0, 0, 16, 16, 16, 16, 16, 16, 32, 32, 32, 32, 32\}$; with 24 local minima and one global minimum at $x^* = (-32, -32)^T, H_1(x^*) \approx 0.998$.

(2) Shekel’s function ($n = 4$),

$$H_2(x) = \sum_{i=1}^5 ((x - a_i)^T(x - a_i) + c_i)^{-1},$$

where $a_1 = (4, 4, 4, 4)^T, a_2 = (1, 1, 1, 1)^T, a_3 = (8, 8, 8, 8)^T, a_4 = (6, 6, 6, 6)^T, a_5 = (3, 7, 3, 7)^T$, and $c = (0.1, 0.2, 0.2, 0.4, 0.4)$, $x^* \approx (4, 4, 4, 4)^T, H_2(x^*) \approx -10.153$.

(3) Rosenbrock function ($n = 20$),

$$H_3(x) = \sum_{i=1}^{n-1} 100(x_{i+1} - x_i^2)^2 + (x_i - 1)^2,$$

where $x^* = (1, \dots, 1)^T, H_3(x^*) = 0$.

(4) Powel singular function ($n = 20$),

$$H_4(x) = \sum_{i=2}^{n-2} [(x_{i-1} + 10x_i)^2 + 5(x_{i+1} - x_{i+2})^2 + (x_i - 2x_{i+1})^4 + 10(x_{i-1} - x_{i+2})^4],$$

where $x^* = (0, \dots, 0)^T, H_4(x^*) = 0$.

(5) Trigonometric function ($n = 20$),

$$H_5(x) = 1 + \sum_{i=1}^n 8 \sin^2(7(x_i - 0.9)^2) + 6 \sin^2(14(x_i - 0.9)^2) + (x_i - 0.9)^2,$$

where $x^* = (0.9, \dots, 0.9)^T, H_5(x^*) = 1$.

(6) Griewank function ($n = 20$),

$$H_6(x) = \frac{1}{4,000} \sum_{i=1}^n x_i^2 - \prod_{i=1}^n \cos\left(\frac{x_i}{\sqrt{i}}\right) + 1,$$

where $x^* = (0, \dots, 0)^T, H_6(x^*) = 0$.

(7) Pintér’s function ($n = 20$),

$$H_7(x) = \sum_{i=1}^n ix_i^2 + \sum_{i=1}^n 20i \sin^2(x_{i-1} \sin x_i - x_i + \sin x_{i+1}) + \sum_{i=1}^n i \log_{10}(1 + i(x_{i-1}^2 - 2x_i + 3x_{i+1} - \cos x_i + 1)^2),$$

where $x_0 = x_n, x_{n+1} = x_1, x^* = (0, \dots, 0)^T, H_7(x^*) = 0$.

We have experimented with different sets of parameters in MRAS₁. We found empirically that the performance of the algorithm is primarily determined by the values of parameters r and v , but is insensitive to the choices of the initial solutions (i.e., mean vector and covariance matrix), provided that the initial sampling variance is large enough.

Table 1. Performance of different algorithms on benchmark problems $H_1 - H_7$ based on 100 independent replications (standard errors are in parentheses).

Test problems	MRAS ₁		CE ($\nu = 0.7$)		CE ($\nu = 0.2$)		SA	
	\bar{H}_i^*	M_ε	\bar{H}_i^*	M_ε	\bar{H}_i^*	M_ε	\bar{H}_i^*	M_ε
H_1	0.998 (3.8e-07)	100	2.22 (0.23)	61	0.998 (4.3e-09)	100	10.12 (0.92)	12
H_2	-10.15 (6.6e-07)	100	-8.38 (0.30)	72	-9.12 (0.11)	1	-6.62 (0.35)	1
H_3	11.64 (5.4e-02)	0	74.68 (19.30)	0	22.63 (4.86)	0	248.5 (23.59)	0
H_4	3.2e-10 (1.8e-11)	100	1.9e+04 (2.8e+03)	0	2.5e-06 (7.5e-08)	100	68.19 (2.94)	0
H_5	1.45 (6.4e-02)	47	1.00 (00e-00)	100	1.00 (4.6e-09)	100	75.69 (4.94)	0
H_6	4.7e-03 (5.8e-04)	55	1.5e-04 (1.0e-04)	98	2.2e-04 (1.3e-04)	97	0.12 (9.7e-03)	0
H_7	4.9e-08 (7.1e-09)	100	4.75 (1.07)	0	2.1e-03 (7.5e-05)	0	1.1e+03 (93.4)	0

Roughly speaking, parameters r and ν essentially serve as trade-offs between explorative search and exploitative search. Smaller values of r and ν help to maintain the search of the algorithm at a more scattered level, whereas larger values of r and ν will, in general, lead to a more rapid decrease in the variance of the underlying sampling distribution, so the search will become more concentrated around the mean of the sampling distribution. Thus, as a general guideline, if explorative search is considered more desirable (especially for high-dimensional problems), then small values of r and ν are often preferred. For all seven test problems, our numerical results are based on the following parameter setting: $\varepsilon = 10^{-5}$, initial sample size $N_0 = 1,000$, $\rho_0 = 0.1$, $\lambda = 0.01$, $\alpha = 1.1$, $r = 10^{-4}$, smoothing parameter $\nu = 0.2$, and $N_{\min} = 5n$. The initial mean vector μ_0 is an n -by-1 vector with each component randomly selected from the interval $[-50, 50]$ according to the uniform distribution, and Σ_0 is an n -by- n diagonal matrix with all diagonal elements equal to 500.

For comparison purposes, we also applied the CE method and the SA algorithm to the above test functions. For CE, we have used the univariate normal p.d.f. with parameter values suggested in Kroese et al. (2006): sample size $N = 2,000$, $\rho = 0.01$, and smoothing parameter $\nu = 0.7$. Again, the initial mean vector μ_0 is randomly selected from $[-50, 50]^n$ according to the uniform distribution, and Σ_0 is an n -by- n diagonal matrix with all elements equal to 500. We found that the above parameters work well for some functions, but in some other cases the variance matrices in CE may converge too quickly to the zero matrix, which freezes the algorithm at some low-quality solutions. To address this issue, for each problem, we also tried CE with different values of the smoothing parameter. In the numerical results reported below, we have used a smaller smoothing parameter value $\nu = 0.2$, which gives reasonable performance for all test cases. For SA, we have used the parameters suggested in Corana et al. (1987): initial temperature $T = 50,000$, temperature reduction factor $r_T = 0.85$, the search neighborhood of a point x is taken to be $\mathcal{N}(x) = \{y: \max_{1 \leq i \leq n} |x_i - y_i| \leq 1\}$, and the initial solution is uniformly selected from $[-50, 50]^n$.

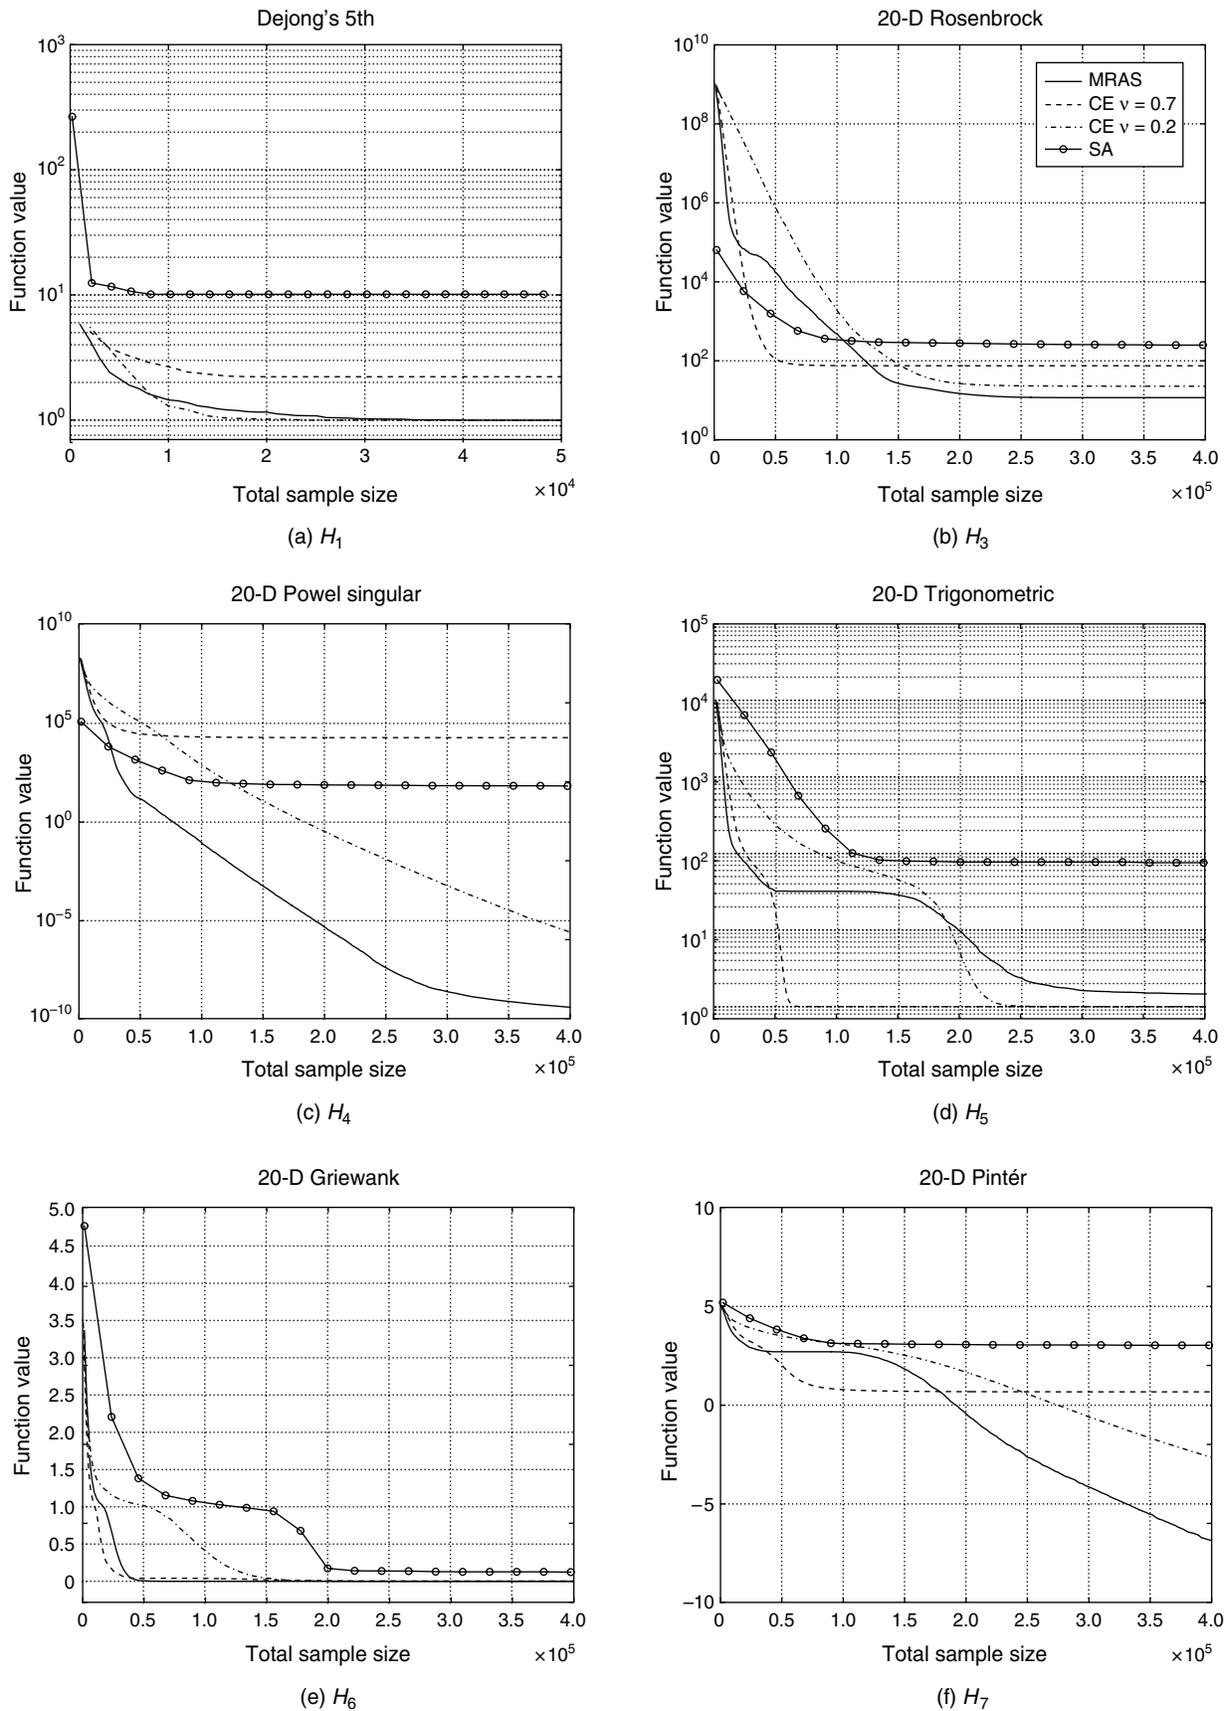
For each problem, we performed 100 independent replications of all three algorithms, and numerical results are

reported in Table 1, where \bar{H}_i^* is the averaged value of the function $H_i(\cdot)$ at the best solution visited by the algorithm, with standard error in parentheses, and M_ε indicates the number of replications in which an ε -optimal solution was found out of 100 trials. We also plotted in Figure 2 the average function values of the current best solution given the number of samples generated for selected benchmark problems. The performance comparison is based on the same amount of computational effort, where for each algorithm the total number of function evaluations (i.e., sample size) is set to 50,000 for H_1 and H_2 , and 400,000 for high-dimensional cases. Here we choose to use the total number of function evaluations to estimate the computational effort used by different algorithms because the computational time of each algorithm is dominated by the time spent in evaluating the objective function.

Functions H_1 and H_2 have only a few local minima, and because SA combines local search, it may quickly locate one of them. However, as we can see, SA stops making improvement even during the early search phase. This is caused by the plateaus surrounding the local minima, which makes it very difficult for SA to escape local optima. In contrast, because both MRAS₁ and CE are population based, they show more robustness in dealing with local optima. We see that CE ($\nu = 0.7$) does not always converge to the global optimal solution, but it still performs much better than SA does. For H_1 , MRAS₁ and CE ($\nu = 0.2$) converge in a similar pattern and consistently find ε -optimal solutions in all simulation runs. Note that decreasing the value of the smoothing parameter slows down the convergence of CE. In particular, for the H_2 case, although better average function values are achieved in CE, no ε -optimal solutions were found within the allowed simulation budget.

For H_3 , none of these three algorithms found ε -optimal solutions. However, Figure 2(b) indicates that both MRAS₁ and CE perform better than SA when the total sample size is large enough. CE with $\nu = 0.2$ converges slowly, but slightly outperforms CE ($\nu = 0.7$) after about 150,000 function evaluations. MRAS₁ performs the best, it has a similar convergence rate as CE ($\nu = 0.7$), and finds better solutions than the other algorithms do. On H_4 , MRAS₁ is clearly superior to both CE and SA. It converges to the global

Figure 2. Average performance of MRAS₁, CE, and SA on selected benchmark problems.



optimal solution in all 100 runs at an exponential rate. The performance of SA is similar to the H_3 case, whereas the performance of CE ($\nu = 0.7$) is even worse than that of SA. As we can see, the algorithm frequently gets trapped at solutions that are far from optimal. CE with $\nu = 0.2$ yields much better performance.

H_5 and H_6 are highly multimodal functions. CE ($\nu = 0.7$) works better than both MRAS₁ and SA. It not only converges the fastest but also finds ε -optimal solutions in almost all runs. SA finds no ε -optimal solutions in any of the runs. MRAS₁ consistently outperforms SA and converges to the optimal solution in 50% of the total simulation runs in both cases. Initially, MRAS₁ converges very fast to good values near the optimum, then it proceeds at a slower rate and spends most of the time in fine-tuning the solution. The behavior of MRAS₁ can be explained by looking at the parameter updating Equations (7) and (8). Because the values of H_5 and H_6 at local minima near the optimum are very close to each other, the parameter updating in MRAS₁ is dominated by the density function in the denominator, especially when the iteration counter k is small.

H_7 contains both a badly scaled quadratic term and some badly scaled noise terms. For this function, SA does not seem to be competitive at all. Similar to the H_3 and H_4 cases, CE ($\nu = 0.7$) converges the fastest but stagnates at some nonoptimal solutions in all runs. Using $\nu = 0.2$ in CE greatly improves the solution quality, but slows down the convergence speed. The initial behavior of MRAS₁ is similar to the H_5 and H_6 cases, but the algorithm outperforms CE ($\nu = 0.7$) after about 170,000 function evaluations and then approaches the optimum at an exponential rate.

The above comparison seems to suggest that MRAS₁ is better adapted to the optimization of badly scaled multimodal problems, whereas CE works best on problems that are well scaled and contain a large number of local optima. Of course, a more comprehensive numerical study needs to be carried out to confirm this finding.

6.2. Combinatorial Optimization

In this section, we present the performance of MRAS₁ on various ATSP problems taken from the website <http://www.iwr.uni-heidelberg.de/groups/comopt/software/TSPLIB95>.

For each ATSP problem with N_c cities, an N_c -by- N_c distance matrix G is given, whose (i, j) th element $G_{i,j}$ represents the distance from city i to city j . The goal is to find the shortest path that visits all the cities and returns to the starting city. Mathematically, the problem can be formulated as follows:

$$\min_{x \in \mathcal{L}} H(x) := \min_{x \in \mathcal{L}} \left\{ \sum_{i=1}^{N_c-1} G_{x_i, x_{i+1}} + G_{x_{N_c}, x_1} \right\}, \quad (14)$$

where $x := (x_1, x_2, \dots, x_{N_c}, x_1)$ is an admissible tour and \mathcal{L} is the set of all admissible tours.

We use the same technique as in Rubinstein (2001) and De Boer et al. (2005) for solving these problems, i.e., we

associate for each distance matrix G an initial state transition matrix \tilde{P}_0 , whose (i, j) th element specifies the probability of transitioning from city i to city j . Thus, at each iteration of MRAS, the following two steps are fundamental: (i) generate random (admissible) tours according to the transition matrix and evaluate the performance of each sample tour; and (ii) update the transition matrix based on the sample tours generated from the previous step.

The detailed discussion of how to generate admissible tours can be found, e.g., in De Boer et al. (2005). We now briefly address the issue of how to update the transition matrix. At each iteration k of MRAS₁, the p.m.f. $f(\cdot, \tilde{P}_k)$ on \mathcal{L} is parameterized by the transition matrix \tilde{P}_k and is given by

$$f(x, \tilde{P}_k) = \prod_{l=1}^{N_c} \sum_{i,j} \tilde{P}_k(i, j) I_{\{x \in \mathcal{X}_{i,j}(l)\}},$$

where $\mathcal{X}_{i,j}(l)$ is the set of all tours in \mathcal{L} such that the l th transition is from city i to city j . It is straightforward to show that the new transition matrix \tilde{P}_{k+1} is updated in (12) as

$$\begin{aligned} \tilde{P}_{k+1}(i, j) &= \frac{\sum_{l=1}^{N_c} [S(H(X_l^k))]^k / \tilde{f}(X_l^k, \tilde{P}_k) I_{\{H(X_l^k) \leq \bar{\gamma}_{k+1}\}} I_{\{x_l^k \in \mathcal{X}_{i,j}\}}}{\sum_{l=1}^{N_c} [S(H(X_l^k))]^k / \tilde{f}(X_l^k, \tilde{P}_k) I_{\{H(X_l^k) \leq \bar{\gamma}_{k+1}\}}} \end{aligned}, \quad (15)$$

where $X_1^k, \dots, X_{N_c}^k$ are the i.i.d. sample tours generated from $\tilde{f}(\cdot, \tilde{P}_k)$, $\bar{\gamma}_{k+1}$ is defined as in (13), and $\mathcal{X}_{i,j}$ represents the set of tours in which the transition from city i to city j is made.

For each of the cases, we performed 30 independent replications of the algorithm. In Table 2, N_{total} is the total number of tours generated (mean and standard error are reported), H_{best} is the length of the shortest path, H_* and H^* are the worst and best solutions obtained out of 30 trials, δ_* and δ^* are the respective relative errors for H_* and H^* , and δ is the relative error (mean and standard error are reported). For all cases, $\varepsilon = 1$, the initial samples $N_0 = 1,000$, $\rho_0 = 0.1$, $\lambda = 0.02$, $\alpha = 1.5$, $r = 0.1$, smoothing parameter $\nu = 0.5$, and the initial transition matrix \tilde{P}_0 is initialized as a stochastic matrix whose (i, j) th entry is proportional to the inverse of the (i, j) th entry of G , i.e., $\tilde{P}_0(i, j) \propto 1/G_{i,j}$ and $\sum_j \tilde{P}_0(i, j) = 1 \forall i$. We stop the algorithm as soon as either one of the following two conditions is satisfied at iteration k : (i) $\max_{1 \leq i \leq 5} |\bar{\gamma}_k - \bar{\gamma}_{k-1}| = 0$; (ii) $N_k > 10N_c^2$. The performance of MRAS₁ is similar to that of CE (cf. Rubinstein 2001, De Boer et al. 2005). We see that the algorithm yields very good solutions by using only a small number of tours.

7. Conclusions and Future Work

We have introduced a randomized optimization technique called model reference adaptive search (MRAS) for solving both continuous and discrete optimization problems.

Table 2. Performance of MRAS on various ATSP problems based on 30 independent replications.

ATSP	N_c	N_{total} (Std. err.)	H_{best}	H_*	H^*	δ_*	δ^*	δ (Std. err.)
ftv33	34	7.41e+04 (3.44e+03)	1,286	1,364	1,286	0.061	0.000	0.023 (0.004)
ftv35	36	1.05e+05 (5.03e+03)	1,473	1,537	1,475	0.043	0.001	0.012 (0.002)
ftv38	39	1.19e+05 (4.90e+03)	1,530	1,598	1,530	0.044	0.000	0.017 (0.003)
p43	43	1.25e+05 (6.29e+03)	5,620	5,638	5,620	0.003	0.000	0.001 (1.4e−4)
ry48p	48	2.75e+05 (1.07e+04)	14,422	14,944	14,446	0.036	0.002	0.018 (0.001)
ft53	53	2.98e+05 (8.71e+03)	6,905	7,352	6,964	0.065	0.008	0.032 (0.003)
ft70	70	5.16e+05 (2.35e+04)	38,673	40,154	38,744	0.038	0.002	0.022 (0.002)

The method offers an alternative framework for global optimization based on which one can design and implement other efficient algorithms. More specifically, we proposed an instantiation of MRAS called MRAS₁, shown its probability one convergence under some mild regularity conditions, and tested its performance on some widely used benchmark problems. We also studied the relationship between CE and MRAS and established some properties of the CE method. This paper can also be seen as a study on the effectiveness of combining CE and EDAs.

The MRAS₁ algorithm demonstrated promise on some preliminary examples, but practical implementation issues remain. For example, selection of the input parameters in our numerical experiments was based mainly on trial and error. For a given problem, how to determine a priori the most appropriate values of these parameters is an open issue. Designing an adaptive scheme to choose these parameters during the search process may also enhance the convergence rate of the algorithm.

A more important line of research is to extend the MRAS method to stochastic optimization problems, where the function values can only be observed in the presence of noise. Denoting $\tilde{H}(x)$ as the random observation of the true function value $H(x)$ made at point x , the stochastic version of problem (1) can be formulated as

$$x^* \in \arg \max_{x \in \mathcal{X}} E[\tilde{H}(x)], \quad x \in \mathcal{X} \subseteq \mathfrak{R}^n, \quad (16)$$

where $E(\cdot)$ is the expectation with respect to the probability distribution of the observation noise. Because an unbiased estimate of $E[\tilde{H}(x)]$ is

$$\frac{1}{M} \sum_{i=1}^M \tilde{H}_i(x),$$

where $\tilde{H}_i(x)$, $i = 1, \dots, M$, are i.i.d. observations made at x , it would be natural to generalize the performance function $[S(H(x))]^k$ in MRAS to

$$S_k(\tilde{H}(x)) := \prod_{i=1}^k S(\tilde{H}_i(x)). \quad (17)$$

Clearly, for the deterministic case (i.e., no observation noise), we will have the original performance function. In particular, if we take $S(\cdot)$ to be an exponential function

(e.g., $S(H(x)) = e^{H(x)}$), then Equation (17) can be written as

$$S_k(\tilde{H}(x)) := \exp\left(\sum_{i=1}^k \tilde{H}_i(x)\right).$$

Therefore, by the strong law of large numbers, it is possible to show that MRAS with the generalized performance function will converge w.p.1 to an optimal solution of (16). However, for this generalized performance function, we need to keep track of all the past observations made at all points visited thus far, which could be computationally difficult to handle when the solution space is large or uncountable. Construction of a practically efficient generalization of MRAS with provable convergence is addressed in Hu et al. (2005).

8. Electronic Companion

An electronic companion to this paper is available as part of the online version that can be found at <http://or.journal.informs.org>.

Appendix. Proofs of Theorems

PROOF OF THEOREM 1. In Lemma 2, we have already established a relationship between reference models $\{g_k(\cdot)\}$ and the sequence of sampling distributions $\{f(\cdot, \theta_k)\}$. Therefore, proving Theorem 1 amounts to showing that $\lim_{k \rightarrow \infty} E_{g_k}[\Gamma(X)] = \Gamma(x^*)$.

Recall from Lemma 1 that $g_{k+1}(\cdot)$ can be expressed recursively as

$$g_{k+1}(x) := \frac{S(H(x))I_{\{H(x) \geq \bar{\gamma}_{k+1}\}}g_k(x)}{E_{g_k}[S(H(X))I_{\{H(X) \geq \bar{\gamma}_{k+1}\}}]} \quad \forall x \in \mathcal{X}, k = 1, 2, \dots$$

Thus,

$$\begin{aligned} E_{g_{k+1}}[S(H(X))I_{\{H(X) \geq \bar{\gamma}_{k+1}\}}] &= \frac{E_{g_k}[[S(H(X))]^2 I_{\{H(X) \geq \bar{\gamma}_{k+1}\}}]}{E_{g_k}[S(H(X))I_{\{H(X) \geq \bar{\gamma}_{k+1}\}}]} \\ &\geq E_{g_k}[S(H(X))I_{\{H(X) \geq \bar{\gamma}_{k+1}\}}]. \end{aligned} \quad (18)$$

Because $\bar{\gamma}_k \leq H(x^*) \forall k$, and each strict increment in the sequence $\{\bar{\gamma}_k\}$ is lower bounded by the quantity $\varepsilon > 0$, there exists a finite \mathcal{N} such that $\bar{\gamma}_{k+1} = \bar{\gamma}_k \forall k \geq \mathcal{N}$. Before we proceed any further, we need to distinguish between two cases: $\bar{\gamma}_{\mathcal{N}} = H(x^*)$ and $\bar{\gamma}_{\mathcal{N}} < H(x^*)$.

Case 1. If $\bar{\gamma}_{\mathcal{N}} = H(x^*)$ (note that because $\rho > 0$, this could only happen when the solution space is discrete), then from the definition of $g_{k+1}(\cdot)$ (see Lemma 1), we obviously have

$$g_{k+1}(x) = 0 \quad \forall x \neq x^*$$

and

$$g_{k+1}(x^*) = \frac{[S(H(x^*))]^k I_{\{H(x)=H(x^*)\}}}{\int_{\mathcal{X}} [S(H(x))]^k I_{\{H(x)=H(x^*)\}} \nu(dx)} = 1 \quad \forall k \geq \mathcal{N}.$$

Hence, it follows immediately that

$$E_{g_{k+1}}[\Gamma(X)] = \Gamma(x^*) \quad \forall k \geq \mathcal{N}.$$

Case 2. If $\bar{\gamma}_{\mathcal{N}} < H(x^*)$, then from (18), we have

$$E_{g_{k+1}}[S(H(X))I_{\{H(X) \geq \bar{\gamma}_{k+2}\}}] \geq E_{g_k}[S(H(X))I_{\{H(X) \geq \bar{\gamma}_{k+1}\}}] \quad \forall k \geq \mathcal{N} - 1, \quad (19)$$

i.e., the sequence $\{E_{g_k}[S(H(X))I_{\{H(X) \geq \bar{\gamma}_{k+1}\}}], k = 1, 2, \dots\}$ converges.

Now we show that the limit of the above sequence is $S(H(x^*))$. To do so, we proceed by contradiction, and assume that

$$S_* := \lim_{k \rightarrow \infty} E_{g_k}[S(H(X))I_{\{H(X) \geq \bar{\gamma}_{k+1}\}}] < S^* := S(H(x^*)). \quad (20)$$

Define the set \mathcal{A} as

$$\mathcal{A} := \{x: H(x) \geq \bar{\gamma}_{\mathcal{N}}\} \cap \{x: S(H(x)) \geq (S^* + S_*)/2\} \cap \mathcal{X}.$$

Because $S(\cdot)$ is strictly increasing, its inverse $S^{-1}(\cdot)$ exists. Thus, \mathcal{A} can be reformulated as

$$\mathcal{A} = \{x: H(x) \geq \max\{\bar{\gamma}_{\mathcal{N}}, S^{-1}((S^* + S_*)/2)\}\} \cap \mathcal{X}.$$

Because $\bar{\gamma}_{\mathcal{N}} < H(x^*)$, \mathcal{A} has a strictly positive Lebesgue/discrete measure by A1.

Note that $g_k(\cdot)$ can be rewritten as

$$g_k(x) = \prod_{i=1}^{k-1} \frac{S(H(x))I_{\{H(x) \geq \bar{\gamma}_{i+1}\}}}{E_{g_i}[S(H(X))I_{\{H(X) \geq \bar{\gamma}_{i+1}\}}]} \cdot g_1(x).$$

Because

$$\lim_{k \rightarrow \infty} \frac{S(H(x))I_{\{H(x) \geq \bar{\gamma}_{k+1}\}}}{E_{g_k}[S(H(X))I_{\{H(X) \geq \bar{\gamma}_{k+1}\}}]} = \frac{S(H(x))I_{\{H(x) \geq \bar{\gamma}_{\mathcal{N}}\}}}{S_*} > 1 \quad \forall x \in \mathcal{A},$$

we conclude that

$$\lim_{k \rightarrow \infty} g_k(x) = \infty \quad \forall x \in \mathcal{A}.$$

Thus, by Fatou’s lemma, we have

$$1 = \liminf_{k \rightarrow \infty} \int_{\mathcal{X}} g_k(x) \nu(dx) \geq \liminf_{k \rightarrow \infty} \int_{\mathcal{A}} g_k(x) \nu(dx) \geq \int_{\mathcal{A}} \liminf_{k \rightarrow \infty} g_k(x) \nu(dx) = \infty,$$

which is a contradiction. Hence, it follows that

$$\lim_{k \rightarrow \infty} E_{g_k}[S(H(X))I_{\{H(X) \geq \bar{\gamma}_{k+1}\}}] = S^*. \quad (21)$$

To show that $\lim_{k \rightarrow \infty} E_{g_k}[\Gamma(X)] = \Gamma(x^*)$, we now bound the difference between $E_{g_k}[\Gamma(X)]$ and $\Gamma(x^*)$. Note that $\forall k \geq \mathcal{N}$, we have

$$\begin{aligned} \|E_{g_k}[\Gamma(X)] - \Gamma(x^*)\| &\leq \int_{\mathcal{X}} \|\Gamma(x) - \Gamma(x^*)\| g_k(x) \nu(dx) \\ &= \int_{\mathcal{C}} \|\Gamma(x) - \Gamma(x^*)\| g_k(x) \nu(dx), \end{aligned} \quad (22)$$

where $\mathcal{C} := \{x: H(x) \geq \bar{\gamma}_{\mathcal{N}}\} \cap \mathcal{X}$ is the support of $g_k(\cdot)$ $\forall k \geq \mathcal{N}$.

By the assumption on $\Gamma(\cdot)$ in Definition 1, for any given $\zeta > 0$, there exists a $\delta > 0$ such that $\|x - x^*\| < \delta$ implies $\|\Gamma(x) - \Gamma(x^*)\| < \zeta$. With A_δ defined from Assumption A2, we have from (22),

$$\begin{aligned} \|E_{g_k}[\Gamma(X)] - \Gamma(x^*)\| &\leq \int_{A_\delta \cap \mathcal{C}} \|\Gamma(x) - \Gamma(x^*)\| g_k(x) \nu(dx) \\ &\quad + \int_{A_\delta^c \cap \mathcal{C}} \|\Gamma(x) - \Gamma(x^*)\| g_k(x) \nu(dx) \\ &\leq \zeta + \int_{A_\delta \cap \mathcal{C}} \|\Gamma(x) - \Gamma(x^*)\| g_k(x) \nu(dx) \quad \forall k \geq \mathcal{N}. \end{aligned} \quad (23)$$

The rest of the proof amounts to showing that the second term in (23) is also bounded. Clearly, the term $\|\Gamma(x) - \Gamma(x^*)\|$ is bounded on the set $A_\delta \cap \mathcal{C}$. We only need to find a bound for $g_k(x)$.

By Assumption A2, we have

$$\sup_{x \in A_\delta \cap \mathcal{C}} H(x) \leq \sup_{x \in A_\delta} H(x) < H(x^*).$$

Define $S_\delta := S^* - S(\sup_{x \in A_\delta} H(x))$. Because $S(\cdot)$ is strictly increasing, we have $S_\delta > 0$. Thus, it follows that

$$S(H(x)) \leq S^* - S_\delta \quad \forall x \in A_\delta \cap \mathcal{C}. \quad (24)$$

On the other hand, from (19) and (21), there exists $\bar{\mathcal{N}} \geq \mathcal{N}$ such that $\forall k \geq \bar{\mathcal{N}}$,

$$E_{g_k}[S(H(X))I_{\{H(X) \geq \bar{\gamma}_{k+1}\}}] \geq S^* - \frac{1}{2}S_\delta. \quad (25)$$

Observe that $g_k(x)$ can be alternatively expressed as

$$g_k(x) = \prod_{i=\bar{\mathcal{N}}}^{k-1} \frac{S(H(x))I_{\{H(x) \geq \bar{\gamma}_{i+1}\}}}{E_{g_i}[S(H(X))I_{\{H(X) \geq \bar{\gamma}_{i+1}\}}]} \cdot g_{\bar{\mathcal{N}}}(x) \quad \forall k \geq \bar{\mathcal{N}}.$$

Thus, it follows from (24) and (25) that

$$g_k(x) \leq \left(\frac{S^* - S_\delta}{S^* - S_\delta/2} \right)^{k-\bar{N}} \cdot g_{\bar{N}}(x) \quad \forall x \in A_\delta \cap \mathcal{C}, \forall k \geq \bar{N}.$$

Therefore,

$$\begin{aligned} & \|E_{g_k}[\Gamma(X)] - \Gamma(x^*)\| \\ & \leq \zeta + \sup_{x \in A_\delta \cap \mathcal{C}} \|\Gamma(x) - \Gamma(x^*)\| \int_{A_\delta \cap \mathcal{C}} g_k(x) \nu(dx) \\ & \leq \zeta + \sup_{x \in A_\delta \cap \mathcal{C}} \|\Gamma(x) - \Gamma(x^*)\| \left(\frac{S^* - S_\delta}{S^* - S_\delta/2} \right)^{k-\bar{N}} \quad \forall k \geq \bar{N} \\ & = \left(1 + \sup_{x \in A_\delta \cap \mathcal{C}} \|\Gamma(x) - \Gamma(x^*)\| \right) \zeta \quad \forall k \geq \hat{N}, \end{aligned}$$

where \hat{N} is given by

$$\hat{N} := \max \left\{ \bar{N}, \left\lceil \bar{N} + \ln \zeta / \ln \left(\frac{S^* - S_\delta}{S^* - S_\delta/2} \right) \right\rceil \right\}.$$

Because ζ is arbitrary, we have

$$\lim_{k \rightarrow \infty} E_{g_k}[\Gamma(X)] = \Gamma(x^*).$$

Finally, the proof is completed by applying Lemma 2 to both Cases 1 and 2. \square

PROOF OF THEOREM 3. (1) The first part of the proof is an extension of the proofs given in Homem-de-Mello (2007). First, we claim that given ρ_k and $\bar{\gamma}_k$, if $\bar{\gamma}_k \leq H(x^*) - \varepsilon$, then $\exists \bar{\mathcal{H}} < \infty$ w.p.1 and $\bar{\rho} \in (0, \rho_k)$ such that $\tilde{\gamma}_{k'+1}(\bar{\rho}, N_{k'}) \geq \bar{\gamma}_k + \varepsilon/2 \forall k' \geq \bar{\mathcal{H}}$. To show this, we proceed by contradiction.

Let $\rho_k^* := \tilde{P}_{\bar{\theta}_k}(H(X) \geq \bar{\gamma}_k + 2\varepsilon/3)$. If $\bar{\gamma}_k \leq H(x^*) - \varepsilon$, then $\bar{\gamma}_k + 2\varepsilon/3 \leq H(x^*) - \varepsilon/3$. By Assumptions A1 and A3', we have

$$\rho_k^* \geq \tilde{P}_{\bar{\theta}_k} \left(H(X) \geq H(x^*) - \frac{\varepsilon}{3} \right) \geq \lambda^{\mathcal{C}}(\varepsilon, \theta_0) > 0, \quad (26)$$

where $\mathcal{C}(\varepsilon, \theta_0) = \int_{\mathcal{X}} I_{\{H(x) \geq H(x^*) - \varepsilon/3\}} f(x, \theta_0) \nu(dx)$ is a constant.

Now assume that $\exists \rho \in (0, \rho_k^*)$ such that $\gamma_{k+1}(\rho, \tilde{\theta}_k) < \bar{\gamma}_k + 2\varepsilon/3$, where $\gamma_{k+1}(\rho, \tilde{\theta}_k)$ is the $(1 - \rho)$ -quantile of $H(X)$ with respect to $f(\cdot, \tilde{\theta}_k)$. By the definition of quantiles, we have

$$\begin{aligned} & \tilde{P}_{\tilde{\theta}_k}(H(X) \geq \gamma_{k+1}(\rho, \tilde{\theta}_k)) \geq \rho \quad \text{and} \\ & \tilde{P}_{\tilde{\theta}_k}(H(X) \leq \gamma_{k+1}(\rho, \tilde{\theta}_k)) \geq 1 - \rho > 1 - \rho_k^*. \end{aligned} \quad (27)$$

It follows that $\tilde{P}_{\tilde{\theta}_k}(H(X) \leq \gamma_{k+1}(\rho, \tilde{\theta}_k)) \leq \tilde{P}_{\tilde{\theta}_k}(H(X) < \bar{\gamma}_k + 2\varepsilon/3) = 1 - \rho_k^*$ by the definition of ρ_k^* , which contradicts Equation (27); thus, we must have that if $\bar{\gamma}_k \leq H(x^*) - \varepsilon$, then

$$\gamma_{k+1}(\rho, \tilde{\theta}_k) \geq \bar{\gamma}_k + \frac{2\varepsilon}{3} \quad \forall \rho \in (0, \rho_k^*).$$

Therefore, by (26), $\exists \bar{\rho} \in (0, \min\{\rho_k, \lambda^{\mathcal{C}}(\varepsilon, \theta_0)\}) \subseteq (0, \rho_k)$ such that $\gamma_{k+1}(\bar{\rho}, \tilde{\theta}_k) \geq \bar{\gamma}_k + 2\varepsilon/3$ whenever $\bar{\gamma}_k \leq H(x^*) - \varepsilon$. By Lemma 7, the distance from the sample $(1 - \bar{\rho})$ -quantile $\tilde{\gamma}_{k+1}(\bar{\rho}, N_k)$ to the set of $(1 - \bar{\rho})$ -quantiles $\gamma_{k+1}(\bar{\rho}, \tilde{\theta}_k)$ goes to zero as $k \rightarrow \infty$ w.p.1; thus, $\exists \bar{\mathcal{H}} < \infty$ w.p.1 such that $\tilde{\gamma}_{k'+1}(\bar{\rho}, N_{k'}) \geq \bar{\gamma}_k + \varepsilon/2 \forall k' \geq \bar{\mathcal{H}}$.

Note that from the MRAS₁ algorithm, if neither Step 3(a) nor 3(b) is visited at the k th iteration, we will have $\rho_{k+1} = \rho_k$ and $\bar{\gamma}_{k+1} = \bar{\gamma}_k$. Thus, whenever $\bar{\gamma}_k \leq H(x^*) - \varepsilon$ w.p.1, Step 3(a)/3(b) will be visited after a finite number of iterations. Furthermore, because the total number of visits to Steps 3(a) and 3(b) is finite (i.e., bounded by $2[H(x^*) - \mathcal{M}]/\varepsilon$, where, recall that \mathcal{M} is a lower bound for $H(x)$), we conclude that there exists $\mathcal{H} < \infty$ w.p.1, such that

$$\bar{\gamma}_k > H(x^*) - \varepsilon \quad \forall k \geq \mathcal{H} \text{ w.p.1.}$$

(2) From the MRAS₁ algorithm, it is easy to see that $\bar{\gamma}_{k+1} \geq \bar{\gamma}_k \forall k = 0, 1, \dots$. By part (1), we have $\bar{\gamma}_{k+1} \geq H(x^*) - \varepsilon \forall k \geq \mathcal{H}$ w.p.1. Thus, by the definition of $\tilde{g}_{k+1}(x)$ (cf. (13)), it follows immediately that if $\{x: H(x) \geq \bar{\gamma}_{k+1}, x \in \{X_1^k, \dots, X_{N_k}^k\}\} \neq \emptyset$, then the support of $\tilde{g}_{k+1}(x)$ satisfies $\text{supp}\{\tilde{g}_{k+1}\} \subseteq \mathcal{C}_\varepsilon \forall k \geq \mathcal{H}$ w.p.1; otherwise if $\{x: H(x) \geq \bar{\gamma}_{k+1}, x \in \{X_1^k, \dots, X_{N_k}^k\}\} = \emptyset$, then $\text{supp}\{\tilde{g}_{k+1}\} = \emptyset$. We now discuss these two cases separately.

Case 1. If $\text{supp}\{\tilde{g}_{k+1}\} \subseteq \mathcal{C}_\varepsilon$, then we have $\{\Gamma(\text{supp}\{\tilde{g}_{k+1}\})\} \subseteq \{\Gamma(\mathcal{C}_\varepsilon)\}$. Because $E_{\tilde{g}_{k+1}}[\Gamma(X)]$ is the convex combination of $\Gamma(X_1^k), \dots, \Gamma(X_{N_k}^k)$, it follows that

$$E_{\tilde{g}_{k+1}}[\Gamma(X)] \in \text{CONV}\{\Gamma(\text{supp}\{\tilde{g}_{k+1}\})\} \subseteq \text{CONV}\{\Gamma(\mathcal{C}_\varepsilon)\}.$$

Thus, by Assumptions A4', A5, and Lemma 6,

$$E_{\tilde{\theta}_{k+1}}[\Gamma(X)] \in \text{CONV}\{\Gamma(\mathcal{C}_\varepsilon)\}.$$

Case 2. If $\text{supp}\{\tilde{g}_{k+1}\} = \emptyset$ (note that this could only happen if Step 3(c) is visited), then from the algorithm, there exists some $\hat{k} < k + 1$ such that $\bar{\gamma}_{k+1} = \bar{\gamma}_{\hat{k}}$ and $\text{supp}\{\tilde{g}_{\hat{k}}\} \neq \emptyset$. Without loss of generality, let \hat{k} be the largest iteration counter such that the preceding properties hold. Because $\bar{\gamma}_{\hat{k}} = \bar{\gamma}_{k+1} > H(x^*) - \varepsilon \forall k \geq \mathcal{H}$ w.p.1, we have $\text{supp}\{\tilde{g}_{\hat{k}}\} \subseteq \mathcal{C}_\varepsilon$ w.p.1. By following the discussions in Case 1, it is clear that

$$E_{\tilde{\theta}_{\hat{k}}}[\Gamma(X)] \in \text{CONV}\{\Gamma(\mathcal{C}_\varepsilon)\} \quad \text{w.p.1.}$$

Furthermore, because $\tilde{\theta}_{\hat{k}} = \tilde{\theta}_{\hat{k}+1} = \dots = \tilde{\theta}_{k+1}$ (see the discussions in §5.1), we will again have

$$E_{\tilde{\theta}_{k+1}}[\Gamma(X)] \in \text{CONV}\{\Gamma(\mathcal{C}_\varepsilon)\} \quad \forall k \geq \mathcal{H} \text{ w.p.1.}$$

Define $\hat{g}_{k+1}(x)$ as

$$\hat{g}_{k+1}(x) := \frac{[S(H(x))]^k I_{\{H(x) \geq \bar{\gamma}_k\}}}{\int_{\mathcal{X}} [S(H(x))]^k I_{\{H(x) \geq \bar{\gamma}_k\}} \nu(dx)} \quad \forall k = 1, 2, \dots,$$

where $\bar{\gamma}_k$ is defined as in MRAS₁. Note that because $\bar{\gamma}_k$ is a random variable, $\hat{g}_{k+1}(x)$ is also a random variable. It follows that

$$E_{\hat{g}_{k+1}}[\Gamma(X)] = \frac{\int_{\mathcal{X}} [\beta S(H(x))]^k I_{\{H(x) \geq \bar{\gamma}_k\}} \Gamma(x) \nu(dx)}{\int_{\mathcal{X}} [\beta S(H(x))]^k I_{\{H(x) \geq \bar{\gamma}_k\}} \nu(dx)}.$$

Let $\omega = (X_1^0, \dots, X_{N_0}^0, X_1^1, \dots, X_{N_1}^1, \dots)$ be a particular sample path generated by the algorithm. For each ω , the sequence $\{\bar{\gamma}_k(\omega), k = 1, 2, \dots\}$ is non decreasing and each strict increase is lower bounded by $\varepsilon/2$. Thus, $\exists \tilde{N}(\omega) > 0$ such that $\bar{\gamma}_{k+1}(\omega) = \bar{\gamma}_k(\omega) \forall k \geq \tilde{N}(\omega)$. Now define $\Omega_1 := \{\omega: \lim_{k \rightarrow \infty} \bar{\gamma}_k(\omega) = H(x^*)\}$. By the definition of $\tilde{g}_{k+1}(\cdot)$ (cf. (13)), for each $\omega \in \Omega_1$, we clearly have $\lim_{k \rightarrow \infty} E_{\tilde{g}_k(\omega)}[\Gamma(X)] = \Gamma(x^*)$; thus, it follows from Lemma 6 that $\lim_{k \rightarrow \infty} E_{\tilde{\theta}_k(\omega)}[\Gamma(X)] = \Gamma(x^*) \forall \omega \in \Omega_1$. The rest of the proof amounts to showing that the result also holds almost surely (a.s.) on the set Ω_1^c .

Because $\lim_{k \rightarrow \infty} \bar{\gamma}_k(\omega) = \bar{\gamma}_{\tilde{N}}(\omega) < H(x^*) \forall \omega \in \Omega_1^c$, we have by Fatou's lemma,

$$\begin{aligned} \liminf_{k \rightarrow \infty} \int_{\mathcal{X}} [\beta S(H(x))]^k I_{\{H(x) \geq \bar{\gamma}_k\}} \nu(dx) \\ \geq \int_{\mathcal{X}} \liminf_{k \rightarrow \infty} [\beta S(H(x))]^k I_{\{H(x) \geq \bar{\gamma}_k\}} \nu(dx) > 0 \\ \forall \omega \in \Omega_1^c, \end{aligned} \tag{28}$$

where the last inequality follows from the fact that $\beta S(H(x)) \geq 1 \forall x \in \{x: H(x) \geq \max\{S^{-1}(1/\beta), \bar{\gamma}_{\tilde{N}}\}\}$ and Assumption A1.

Because $f(x, \theta_0) > 0 \forall x \in \mathcal{X}$, we have $\mathcal{X} \subseteq \text{supp}\{\tilde{f}(\cdot, \tilde{\theta}_k)\} \forall k$; thus,

$$\begin{aligned} E_{\tilde{g}_{k+1}}[\Gamma(X)] &= \frac{\tilde{E}_{\tilde{\theta}_k}[\beta^k \tilde{S}_k(H(X)) I_{\{H(X) \geq \bar{\gamma}_k\}} \Gamma(X)]}{\tilde{E}_{\tilde{\theta}_k}[\beta^k \tilde{S}_k(H(X)) I_{\{H(X) \geq \bar{\gamma}_k\}}]} \\ &\forall k = 1, 2, \dots, \end{aligned}$$

where $\tilde{S}_k(H(x)) := [S(H(x))]^k / \tilde{f}(x, \tilde{\theta}_k)$. We now show that $E_{\tilde{g}_{k+1}}[\Gamma(X)] \rightarrow E_{\tilde{g}_{k+1}}[\Gamma(X)]$ a.s. on Ω_1^c as $k \rightarrow \infty$. Because we are only interested in the limiting behavior of $E_{\tilde{g}_{k+1}}[\Gamma(X)]$, it is sufficient to show that

$$\begin{aligned} \frac{(1/N_k) \sum_{i=1}^{N_k} \beta^k \tilde{S}_k(H(X_i^k)) I_{\{H(X_i^k) \geq \bar{\gamma}_{k+1}\}} \Gamma(X_i^k)}{(1/N_k) \sum_{i=1}^{N_k} \beta^k \tilde{S}_k(H(X_i^k)) I_{\{H(X_i^k) \geq \bar{\gamma}_k\}}} \\ \rightarrow E_{\tilde{g}_{k+1}}[\Gamma(X)] \quad \text{a.s. on } \Omega_1^c, \end{aligned}$$

where and hereafter, whenever $\{x: H(x) \geq \bar{\gamma}_{k+1}, x \in \{X_1^k, \dots, X_{N_k}^k\}\} = \emptyset$, we define $0/0 = 0$.

For brevity, we use the following shorthand notations:

$$\begin{aligned} \hat{Y}^k &:= \tilde{E}_{\tilde{\theta}_k}[\beta^k \tilde{S}_k(H(X)) I_{\{H(X) \geq \bar{\gamma}_k\}}], \\ \hat{Y}_\Gamma^k &:= \tilde{E}_{\tilde{\theta}_k}[\beta^k \tilde{S}_k(H(X)) I_{\{H(X) \geq \bar{\gamma}_k\}} \Gamma(X)], \\ \hat{Y}_i^k &:= \beta^k \tilde{S}_k(H(X_i^k)) I_{\{H(X_i^k) \geq \bar{\gamma}_{k+1}\}}, \\ \hat{Y}_i^k &:= \beta^k \tilde{S}_k(H(X_i^k)) I_{\{H(X_i^k) \geq \bar{\gamma}_k\}}. \end{aligned}$$

We also let $\mathcal{T}_\varepsilon := \lceil 2[H(x^*) - \mathcal{M}]/\varepsilon \rceil$. Note that the total number of visits to Steps 3(a) and 3(b) of MRAS₁ is bounded by \mathcal{T}_ε , thus for any $k > \mathcal{T}_\varepsilon$, the total number of visits to Step 3(c) is greater than $k - \mathcal{T}_\varepsilon$.

We have

$$\begin{aligned} \frac{(1/N_k) \sum_{i=1}^{N_k} \beta^k \tilde{S}_k(H(X_i^k)) I_{\{H(X_i^k) \geq \bar{\gamma}_{k+1}\}} \Gamma(X_i^k)}{(1/N_k) \sum_{i=1}^{N_k} \beta^k \tilde{S}_k(H(X_i^k)) I_{\{H(X_i^k) \geq \bar{\gamma}_{k+1}\}}} - E_{\tilde{g}_{k+1}}[\Gamma(X)] \\ = \left(\frac{(1/N_k) \sum_{i=1}^{N_k} \bar{Y}_i^k \Gamma(X_i^k)}{(1/N_k) \sum_{i=1}^{N_k} \bar{Y}_i^k} - \frac{(1/N_k) \sum_{i=1}^{N_k} \hat{Y}_i^k \Gamma(X_i^k)}{(1/N_k) \sum_{i=1}^{N_k} \hat{Y}_i^k} \right) \\ + \left(\frac{(1/N_k) \sum_{i=1}^{N_k} \hat{Y}_i^k \Gamma(X_i^k)}{(1/N_k) \sum_{i=1}^{N_k} \hat{Y}_i^k} - \frac{\hat{Y}_\Gamma^k}{\hat{Y}^k} \right). \end{aligned}$$

Because for each $\omega \in \Omega_1^c$, $\bar{\gamma}_{k+1}(\omega) = \bar{\gamma}_k(\omega) \forall k \geq \tilde{N}(\omega)$, it is straightforward to see that the first term

$$\begin{aligned} \frac{(1/N_k) \sum_{i=1}^{N_k} \bar{Y}_i^k \Gamma(X_i^k)}{(1/N_k) \sum_{i=1}^{N_k} \bar{Y}_i^k} - \frac{(1/N_k) \sum_{i=1}^{N_k} \hat{Y}_i^k \Gamma(X_i^k)}{(1/N_k) \sum_{i=1}^{N_k} \hat{Y}_i^k} = 0 \\ \forall k \geq \tilde{N}(\omega), \forall \omega \in \Omega_1^c. \end{aligned} \tag{29}$$

To show that the second term also converges to zero, we denote by \mathcal{V}_k the event $\mathcal{V}_k = \{\bar{\gamma}_k > H(x^*) - \varepsilon\}$. For any $\zeta > 0$, we also let \mathcal{C}_k be the event

$$\mathcal{C}_k = \left\{ \left| (1/N_k) \sum_{i=1}^{N_k} \hat{Y}_i^k - \hat{Y}^k \right| > \zeta \right\}.$$

We have

$$\begin{aligned} P(\mathcal{C}_k \text{ i.o.}) &= P(\{\mathcal{C}_k \cap \mathcal{V}_k\} \cup \{\mathcal{C}_k \cap \mathcal{V}_k^c\} \text{ i.o.}) \\ &= P(\mathcal{C}_k \cap \mathcal{V}_k \text{ i.o.}) \\ &\text{because } P(\mathcal{V}_k^c \text{ i.o.}) = 0 \text{ by part (1)}. \end{aligned} \tag{30}$$

It is easy to see that conditional on $\tilde{\theta}_k$ and $\bar{\gamma}_k, \hat{Y}_1^k, \dots, \hat{Y}_{N_k}^k$ are i.i.d. and $E[\hat{Y}_i^k | \tilde{\theta}_k, \bar{\gamma}_k] = \hat{Y}^k \forall i$. Furthermore, by Assumption A3', conditional on the event \mathcal{V}_k , the support $[a_k, b_k]$ of the random variable \hat{Y}_i^k satisfies $[a_k, b_k] \subseteq [0, (\beta S^*)^k / \lambda f_*]$. Therefore, we have from the Hoeffding inequality (Hoeffding 1963),

$$\begin{aligned} P(\mathcal{C}_k | \mathcal{V}_k, \tilde{\theta}_k = \theta, \bar{\gamma}_k = \gamma) \\ = P\left(\left| (1/N_k) \sum_{i=1}^{N_k} \hat{Y}_i^k - \hat{Y}^k \right| > \zeta \mid \mathcal{V}_k, \tilde{\theta}_k = \theta, \bar{\gamma}_k = \gamma \right) \\ \leq 2 \exp\left(\frac{-2N_k \zeta^2}{(b_k - a_k)^2} \right) \\ \leq 2 \exp\left(\frac{-2N_k \zeta^2 [\lambda f_*]^2}{(\beta S^*)^{2k}} \right) \quad \forall k = 1, 2, \dots \end{aligned} \tag{31}$$

Because

$$\begin{aligned} P(\mathcal{C}_k \cap \mathcal{V}_k) &= \int_{\theta, \gamma} P(\mathcal{C}_k \cap \mathcal{V}_k | \tilde{\theta}_k = \theta, \bar{\gamma}_k = \gamma) f_{\tilde{\theta}_k, \bar{\gamma}_k}(d\theta, d\gamma) \\ &\leq \int_{\theta, \gamma} P(\mathcal{C}_k | \mathcal{V}_k, \tilde{\theta}_k = \theta, \bar{\gamma}_k = \gamma) f_{\tilde{\theta}_k, \bar{\gamma}_k}(d\theta, d\gamma), \end{aligned}$$

where $f_{\tilde{\theta}_k, \tilde{\gamma}_k}(\cdot, \cdot)$ is the joint distribution of random variables $\tilde{\theta}_k$ and $\tilde{\gamma}_k$, we have by (31),

$$\begin{aligned} P(\mathcal{E}_k \cap \mathcal{V}_k) &\leq 2 \exp\left(\frac{-2N_k \zeta^2 [\lambda f_*]^2}{(\beta S^*)^{2k}}\right) \\ &\leq 2 \exp\left(\frac{-2(\alpha^{k-\mathcal{J}_\varepsilon} N_0) \zeta^2 [\lambda f_*]^2}{(\beta S^*)^{2k}}\right) \quad \forall k \geq \mathcal{J}_\varepsilon \\ &= 2 \exp\left(\frac{-2N_0 \zeta^2 \lambda^2 f_*^2}{\alpha^{\mathcal{J}_\varepsilon}} \left(\frac{\alpha}{(\beta S^*)^2}\right)^k\right) \quad \forall k \geq \mathcal{J}_\varepsilon. \end{aligned}$$

Because $\alpha/(\beta S^*)^2 > 1$ (by assumption), it follows that

$$\lim_{k \rightarrow \infty} P(\mathcal{E}_k \cap \mathcal{V}_k) = 0.$$

Furthermore, because $e^{-x} < 1/x \forall x > 0$, we have

$$P(\mathcal{E}_k \cap \mathcal{V}_k) < \frac{\alpha^{\mathcal{J}_\varepsilon}}{N_0 \zeta^2 \lambda^2 f_*^2} \left(\frac{(\beta S^*)^2}{\alpha}\right)^k \quad \forall k \geq \mathcal{J}_\varepsilon,$$

and because $(\beta S^*)^2/\alpha < 1$, we have

$$\sum_{k=0}^{\infty} P(\mathcal{E}_k \cap \mathcal{V}_k) < \mathcal{J}_\varepsilon + \frac{\alpha^{\mathcal{J}_\varepsilon}}{N_0 \zeta^2 \lambda^2 f_*^2} \sum_{k=\mathcal{J}_\varepsilon}^{\infty} \left(\frac{(\beta S^*)^2}{\alpha}\right)^k < \infty.$$

Finally, by the Borel-Cantelli lemma and (30),

$$P(\mathcal{E}_k \text{ i.o.}) = P(\mathcal{E}_k \cap \mathcal{V}_k \text{ i.o.}) = 0.$$

Because this holds for any $\zeta > 0$, we have $(1/N_k) \cdot \sum_{i=1}^{N_k} \hat{Y}_i^k \rightarrow \hat{Y}^k$ w.p.1.

By following the same argument as before, we can also show that $(1/N_k) \sum_{i=1}^{N_k} \hat{Y}_i^k \Gamma(X_i^k) \rightarrow \hat{Y}_\Gamma^k$ w.p.1. Because $\liminf_{k \rightarrow \infty} \hat{Y}^k > 0 \forall \omega \in \Omega_1^c$ (i.e., (28)), we have

$$\frac{(1/N_k) \sum_{i=1}^{N_k} \hat{Y}_i^k \Gamma(X_i^k)}{(1/N_k) \sum_{i=1}^{N_k} \hat{Y}_i^k} \rightarrow \frac{\hat{Y}_\Gamma^k}{\hat{Y}^k} \quad \text{as } k \rightarrow \infty \text{ a.s. on } \Omega_1^c.$$

By the definition of $\tilde{g}_{k+1}(\cdot)$, the above result together with (29) suggests that

$$E_{\tilde{g}_k}[\Gamma(X)] \rightarrow E_{\tilde{g}_k}[\Gamma(X)] \quad \text{as } k \rightarrow \infty \text{ a.s. on } \Omega_1^c.$$

Thus, in conclusion, we have

$$E_{\tilde{g}_k}[\Gamma(X)] \rightarrow E_{\tilde{g}_k}[\Gamma(X)] \quad \text{as } k \rightarrow \infty \text{ w.p.1.}$$

On the other hand, by Assumptions A1 and A2, and following the proof of Theorem 1, it is not difficult to show that

$$E_{\tilde{g}_k}[\Gamma(X)] \rightarrow \Gamma(x^*) \quad \text{as } k \rightarrow \infty \text{ w.p.1.}$$

Hence, by Lemma 6, we have

$$\lim_{k \rightarrow \infty} E_{\tilde{\theta}_k}[\Gamma(X)] = \lim_{k \rightarrow \infty} E_{\tilde{g}_k}[\Gamma(X)] = \Gamma(x^*) \quad \text{w.p.1. } \square$$

Acknowledgments

This work was supported in part by the National Science Foundation under grants DMI-9988867 and DMI-0323220, and by the Air Force Office of Scientific Research under grants F496200110161 and FA95500410210. The authors thank the associate editor and referees for their detailed comments and suggestions that led to a substantially improved paper, and they thank one of the referees for pointing out the related work of Wolpert.

References

- Auer, P., N. Cesa-Bianchi, Y. Freund, R. E. Schapire. 2002. The non-stochastic multiarmed bandit problem. *SIAM J. Comput.* **32** 48–77.
- Chang, H. S., M. C. Fu, J. Hu, S. I. Marcus. 2007. An asymptotically efficient simulation-based algorithm for finite horizon stochastic dynamic programming. *IEEE Trans. Automat. Control.* **52** 89–94.
- Corana, A., M. Marchesi, C. Martini, S. Ridella. 1987. Minimizing multimodal functions of continuous variables with the simulated annealing algorithm. *ACM Trans. Math. Software* **13** 262–280.
- De Boer, P. T., D. P. Kroese, S. Mannor, R. Y. Rubinstein. 2005. A tutorial on the cross-entropy method. *Ann. Oper. Res.* **134** 19–67.
- Dorigo, M., L. M. Gambardella. 1997. Ant colony system: A cooperative learning approach to the traveling salesman problem. *IEEE Trans. Evolutionary Comput.* **1** 53–66.
- Glover, F. 1990. Tabu search: A tutorial. *Interfaces* **20** 74–94.
- Hoeffding, W. 1963. Probability inequalities for sums of bounded random variables. *J. Amer. Statist. Assoc.* **58**(301) 13–30.
- Homem-de-Mello, T. 2007. A study on the cross-entropy method for rare event probability estimation. *INFORMS J. Comput.* Forthcoming.
- Hu, J., M. C. Fu, S. I. Marcus. 2005. Simulation optimization using model reference adaptive search. *Winter Simulation Conf.* Orlando, FL, 811–818.
- Kirkpatrick, S., C. D. Gelatt, M. P. Vecchi. 1983. Optimization by simulated annealing. *Science* **220** 671–680.
- Kroese, D. P., S. Porotsky, R. Y. Rubinstein. 2006. The cross-entropy method for continuous multiextremal optimization. *Methodology Comput. Appl. Probab.* **8** 383–407.
- Larrañaga, P., R. Etxeberria, J. A. Lozano, B. Sierra, I. Iñza, J. M. Peña. 1999. A review of the cooperation between evolutionary computation and probabilistic graphical models. *Proc. Second Sympos. Artificial Intelligence. Adaptive Systems. CIMA 99. Special Session on Distributions and Evolutionary Computation.* Havana, Cuba, 314–324.
- Morris, C. N. 1982. Natural exponential families with quadratic variance functions. *Ann. Statist.* **10** 65–80.
- Mühlenbein, H., G. Paaß. 1996. From recombination of genes to the estimation of distributions: I. Binary parameters. H.-M. Voigt, W. Ebeling, I. Rechenberg, H.-P. Schwefel, eds. *Parallel Problem Solving from Nature-PPSN IV.* Springer Verlag, Berlin, Germany, 178–187.
- Pintér, J. D. 1996. *Global Optimization in Action.* Kluwer Academic Publishers, Dordrecht, The Netherlands.
- Rubinstein, R. Y. 1997. Optimization of computer simulation models with rare events. *Eur. J. Oper. Res.* **99** 89–112.
- Rubinstein, R. Y. 1999. The cross-entropy method for combinatorial and continuous optimization. *Methodology Comput. Appl. Probab.* **2** 127–190.
- Rubinstein, R. Y. 2001. Combinatorial optimization, ants and rare events. S. Uryasev, P. M. Pardalos, eds. *Stochastic Optimization: Algorithms and Applications.* Kluwer Academic Publishers, Dordrecht, The Netherlands, 304–358.

- Rubinstein, R. Y., D. P. Kroese. 2004. *The Cross-Entropy Method: A Unified Approach to Combinatorial Optimization, Monte-Carlo Simulation, and Machine Learning*. Springer, New York.
- Shi, L., S. Ólafsson. 2000. Nested partitions method for global optimization. *Oper. Res.* **48** 390–407.
- Srinivas, M., L. M. Patnaik. 1994. Genetic algorithms: A survey. *Computer* **27** 17–26.
- Wolpert, D. H. 2004. Finding bounded rational equilibria part I: Iterative focusing. T. Vincent, ed. *Proc. Eleventh Internat. Sympos. on Dynamic Games Appl.* Tucson, AZ, 18–21.
- Yao, X., Y. Liu. 1996. Fast evolutionary programming. *Proc. 5th Annual Conf. Evolutionary Programming*. MIT Press, Cambridge, MA, 451–460.
- Zhang, Q., H. Mühlenbein. 2004. On the convergence of a class of estimation of distribution algorithm. *IEEE Trans. Evolutionary Comput.* **8** 127–136.
- Zlochin, M., M. Birattari, N. Meuleau, M. Dorigo. 2004. Model-based search for combinatorial optimization: A critical survey. *Ann. Oper. Res.* **131** 373–395.