TECHNICAL RESEARCH REPORT

A Model Reference Adaptive Search Algorithm for Global Optimization

by Jiaqiao Hu, Michael C. Fu, Steven I. Marcus

ISR
INSTITUTE FOR SYSTEMS RESEARCH

A Model Reference Adaptive Search Algorithm
for Global Optimization

Jiaqiao Hu
Department of Electrical and Computer Engineering & Institute for Systems Research,
University of Maryland, College Park, MD 20742, jqhu@glue.umd.edu

Michael C. Fu
Robert H. Smith School of Business & Institute for Systems Research,
University of Maryland, College Park, MD 20742, mfu@rhsmith.umd.edu

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

Abstract

This paper presents a randomized algorithm called Model Reference Adaptive Search (MRAS) for solving global optimization problems. The algorithm generates at each iteration a group of candidate solutions according to a parameterized probabilistic model. 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 reference models that will eventually converge to a model producing only the optimal solutions. In fact, we show that the model reference framework can also be used to describe the recently proposed cross-entropy (CE) method for optimization and study its properties. Some global convergence properties of MRAS are established in both continuous and combinatorial domains. Numerical studies are also carried out to demonstrate the effectiveness of the algorithm.

Keywords: global optimization, combinatorial optimization, cross-entropy method (CE), estimation of distribution algorithm (EDA).

---

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 new 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 (random samples, trajectories) according to a specified probabilistic model (e.g., a parameterized probability distribution on the solution space).

2. Update the parameters associated with the probabilistic model, on the basis of the data collected in the previous step, in order to bias the future search toward "better" solutions.

Some well-established *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. 2004), 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 cross-entropy (CE) method and the estimation of distribution algorithms (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 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 and Rubinstein (2003) show 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 estimation of distribution algorithms (EDAs) were first introduced in the field of evolutionary computation in Mühlenbein and Paaß (1996). It inherits the spirit of the well-known genetic algorithms (GAs), but eliminates the crossover and the mutation operators in order to avoid the disruption of partial solutions. In EDAs, a new population of candidate solutions are 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 propose a new randomized algorithm, called model reference adaptive search (MRAS), for solving both continuous and combinatorial optimization problems. MRAS resembles CE and EDAs in that it works with a family of parameterized distributions on the solution space. However, the algorithm has a particular working mechanism which, to the best of our knowledge, has never been proposed before. 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. In contrast, as mentioned previously, the CE method uses 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. The sequence of reference distributions in MRAS can be viewed as the generalized probability distribution models for EDAs with proportional selection scheme (Zhang and Mühlenbein 2004; see also Chang et al. 2003 in the context of Markov decision processes) and can be shown to converge to a degenerate distribution concentrated only on the set of optimal solutions. However, in EDAs (with proportional selection scheme), these distribution models are directly constructed in order to generate new candidate solutions; this step often constitutes the most crucial and difficult part of the method, whereas in our approach, the sequence of reference models are only used *implicitly* to guide the parameter updating procedure; thus there is no need to build them explicitly. 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 that has the minimum KL-divergence with respect to the current reference model, from which future candidate solutions will be generated. We show that for a class of parameterized probability distributions, the so-called Natural Exponential Family (NEF), the algorithm converges to an optimal solution with probability one. In sum, compared with CE and EDAs, MRAS has several advantages:

(1) The general idea behind the method offers a flexible framework for constructing and studying different

algorithms. For instance, by using different sequences of reference models, one can possibly design modified versions of our proposed approach. In fact, we show that this framework can also be used to interpret the CE method and analyze its properties.

(2) The "guided" model-construction procedure used in MRAS enables the method to avoid the (possibly) high computational cost of the direct model-construction procedures employed in EDAs.

(3) Since the sequence of reference distributions depend directly on the performance of the candidate solutions, MRAS automatically takes into account the correlations between the underlying decision variables, so that the random samples generated at each stage can be efficiently utilized.

The rest of the paper is organized as follows. In Section 2, we give the necessary background and discuss the general ideas behind MRAS. In Section 3, we describe the deterministic version of the algorithm and present its global convergence properties. In Section 4, we provide a unified view of CE and MRAS, and establish some important properties of the CE method. In Section 5, we present the Monte Carlo version of MRAS and show that the proposed method converges to an optimal solution with probability one. Illustrative numerical studies on both continuous and combinatorial optimization problems are given in Section 6. Finally some future research topics are outlined in Section 7.

2 Background

We consider the following optimization problem:

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

where $\mathcal{X}$ is the solution space, and $H(\cdot) : \mathcal{X} \to \Re$ 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., there exists $x^* \in \mathcal{X}$ such that $H(x) < H(x^*)$ for all $x \neq x^*$, $x \in \mathcal{X}$. Our primary concern in this paper is on *unconstrained* or *partially constrained* optimization problems; assume that random sampling can be done easily on $\mathcal{X}$, at least for a class of distributions of interest. For *constrained* optimization problems, we refer the reader to Kroese et al. (2004) for a discussion of how to convert them to *unconstrained* problems.

To explain the main idea behind MRAS, we consider the following naive *model-based* approach for solving (1). 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. by tilting the old p.d.f. $g_{k-1}(x)$ with the performance function $H(x)$ (for simplicity, here we assume $H(x) > 0 \, \forall x \in \mathcal{X}$), i.e.,

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

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, let $X = (X_1, \ldots, X_n)$ be a random variable taking values in $\mathcal{X}$. To reduce the notational burden, henceforth X will be used to denote a random variable having the distribution under which the expectation is indicated. Thus, $E_{g_k}[H(X)] = \int_{\mathcal{X}} H(x)g_k(dx)$ and $E_{g_{k-1}}[H(X)] = \int_{\mathcal{X}} H(x)g_{k-1}(dx)$. Then we have

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

Furthermore, it is possible to show that the sequence of p.d.f.'s $\{g_k(\cdot),\ k = 0, 1, \ldots\}$ will converge to a p.d.f. that concentrates only on the set of optimal solutions for arbitrary $g_0(\cdot)$. So we will have $\lim_{k \to \infty} E_{g_k}[H(X)] = H(x^*)$.

However, the above approach is generally of little practical use, due to the following reasons: (*i*) It is usually not possible to enumerate all the points in the solution space in order to perform the update (2); furthermore, if it were possible, the optimal solution could be immediately identified simply by checking which point has the best performance value. (*ii*) The p.d.f. $g_k(x)$ constructed at each iteration may not have any structure, and therefore may be very difficult to handle.

To overcome the above difficulties, we consider the Monte Carlo (sampling) version of the above approach and at the same time restrict ourselves to a family of parameterized p.d.f.'s $\{f(\cdot, \theta)\}$, where θ is the parameter vector. In particular, at each iteration k of the algorithm, we look at the projection of $g_k(\cdot)$ on the family of p.d.f.'s $\{f(\cdot, \theta)\}$ and compute the parameter vector θ_k that minimizes the Kullback-Leibler (KL) divergence

$$\mathcal{D}(g_k, f(\cdot, \theta)) := E_{g_k}\left[\ln \frac{g_k(X)}{f(X, \theta)}\right] = \int_{x \in \mathcal{X}} \ln \frac{g_k(x)}{f(x, \theta)} g_k(dx).$$

The benefits of the above consideration are twofold: on the one hand, $f(\cdot, \theta_k)$ often has some special structure and therefore could be much easier to handle than $g_k(\cdot)$. On the other hand, the sequence $\{f(\cdot, \theta_k)\}$ may retain some nice properties of $\{g_k(\cdot)\}$ and converge to a degenerate p.d.f. concentrated on the set of optimal solutions.

First, however, we present the deterministic version of the MRAS algorithm, because it serves as a starting point for deriving the results of Sections 4 and 5.

3 Model Reference Adaptive Search (Deterministic Version)

For ease of exposition, we put our main focus on continuous problems and leave the discussion for discrete/combinatorial problems as special cases. 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 $I_{\{\cdot\}}$ 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,

$$P_{\theta_k}(H(X) \geq \gamma) = \int_{x \in \mathcal{X}} I_{\{H(x) \geq \gamma\}} f(dx, \theta_k) \text{ and } E_{\theta_k}[H(X)] = \int_{x \in \mathcal{X}} H(x) f(dx, \theta_k).$$

3.1 Algorithm Description

The MRAS_0 algorithm requires specification of a parameter ρ, which determines the approximate proportion of samples that will be used to update the probabilistic model. At successive iterations of the algorithm, a sequence $\{\gamma_k, k = 1, 2, \ldots\}$, i.e., the $(1 - \rho)$-quantiles with respect to the sequence of p.d.f's $\{f(\cdot, \theta_k)\}$, are calculated at step 1 of MRAS_0. These quantile values are then used in step 2 to construct a sequence of non-decreasing thresholds $\{\bar{\gamma}_k, k = 1, 2, \ldots\}$; and only those candidate solutions that have performances better than these thresholds will be used in parameter updating (cf. equation (3)). As we will see, the theoretical convergence of MRAS_0 is unaffected by the value of the parameter ρ. The purpose of ρ in our approach is 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 MRAS_0, a small number ε and a continuous and strictly increasing function $S(\cdot) : \Re \to \Re^+$ 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.,

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

We require ε to be strictly positive for continuous problems, and non-negative for discrete 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 Section 5 will ensure that step 3 of MRAS$_0$ can be used interchangeably with the following equation:

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

We now show that there is a sequence of reference models $\{g_k(\cdot), \ k = 1, 2, \ldots\}$ implicit in MRAS$_0$, and the parameter θ_{k+1} computed at step 3 indeed minimizes the KL-divergence $\mathcal{D}(g_{k+1}, f(\cdot, \theta))$.

Lemma 3.1 *The parameter θ_{k+1} computed at the kth iteration of the MRAS$_0$ algorithm minimizes the KL-divergence $\mathcal{D}(g_{k+1}, f(\cdot, \theta))$, where*

$$g_{k+1}(x) := \frac{S(H(x))I_{\{H(x) \geq \bar{\gamma}_{k+1}\}} g_k(x)}{E_{g_k} \left[S(H(X))I_{\{H(X) \geq \bar{\gamma}_{k+1}\}} \right]} \quad \forall x \in \mathcal{X}, \quad k = 1, 2, \ldots, \quad \text{and} \quad g_1(x) := \frac{I_{\{H(x) \geq \bar{\gamma}_1\}}}{E_{\theta_0} \left[\frac{I_{\{H(X) \geq \bar{\gamma}_1\}}}{f(X,\theta_0)} \right]}.$$

Proof: For brevity, define $\widehat{S}_k(H(x)) := \frac{[S(H(x))]^k}{f(x,\theta_k)}$. We have

$$g_1(x) = \frac{I_{\{H(x) \geq \bar{\gamma}_1\}}}{E_{\theta_0} \left[\frac{I_{\{H(X) \geq \bar{\gamma}_1\}}}{f(X,\theta_0)} \right]} = \frac{I_{\{H(x) \geq \bar{\gamma}_1\}}}{E_{\theta_0} \left[\widehat{S}_0(H(X)) I_{\{H(X) \geq \bar{\gamma}_1\}} \right]}.$$

When $k \geq 1$, we have from the definition of $g_k(\cdot)$ above,

$$
\begin{aligned}
g_2(x) &= \frac{S(H(x))I_{\{H(x) \geq \bar{\gamma}_2\}} g_1(x)}{E_{g_1}\left[S(H(X))I_{\{H(X) \geq \bar{\gamma}_2\}}\right]} \\
&= \frac{S(H(x))I_{\{H(x) \geq \bar{\gamma}_2\}} I_{\{H(x) \geq \bar{\gamma}_1\}}}{E_{\theta_1}\left[\widehat{S}_1(H(X))I_{\{H(X) \geq \bar{\gamma}_2\}} I_{\{H(X) \geq \bar{\gamma}_1\}}\right]} \\
&= \frac{S(H(x))I_{\{H(x) \geq \bar{\gamma}_2\}}}{E_{\theta_1}\left[\widehat{S}_1(H(X))I_{\{H(X) \geq \bar{\gamma}_2\}}\right]},
\end{aligned}
$$

where the last equality follows from the fact that the sequence $\{\bar{\gamma}_k, \ k = 1, 2, \ldots\}$ is non-decreasing. Proceeding iteratively, it is easy to see that

$$
g_{k+1}(x) = \frac{[S(H(x))]^k I_{\{H(x) \geq \bar{\gamma}_{k+1}\}}}{E_{\theta_k}\left[\widehat{S}_k(H(X))I_{\{H(X) \geq \bar{\gamma}_{k+1}\}}\right]}, \ \forall \, k = 0, 1, \ldots.
$$

Thus, the KL-divergence between $g_{k+1}(\cdot)$ and $f(\cdot, \theta)$ can be written as

$$
\begin{aligned}
\mathcal{D}\left(g_{k+1}, f(\cdot, \theta)\right) &= E_{g_{k+1}}\left[\ln g_{k+1}(X)\right] - E_{g_{k+1}}\left[\ln f(X, \theta)\right] \\
&= E_{g_{k+1}}\left[\ln g_{k+1}(X)\right] - \frac{E_{\theta_k}\left[\widehat{S}_k(H(X))I_{\{H(X) \geq \bar{\gamma}_{k+1}\}} \ln f(X, \theta)\right]}{E_{\theta_k}\left[\widehat{S}_k(H(X))I_{\{H(X) \geq \bar{\gamma}_{k+1}\}}\right]}, \ \forall \, k = 0, 1, \ldots.
\end{aligned}
$$

The result follows by observing that minimizing $\mathcal{D}\left(g_{k+1}, f(\cdot, \theta)\right)$ with respect to θ is equivalent to maximizing the quantity $E_{\theta_k}\left[\widehat{S}_k(H(X))I_{\{H(X) \geq \bar{\gamma}_{k+1}\}} \ln f(X, \theta)\right]$. ∎

3.2 Global Convergence

Global convergence of the MRAS$_0$ algorithm clearly depends on the choice of parameterized p.d.f./p.m.f. family. The algorithm may not be computationally tractable for some choices. Throughout this paper, we restrict our analysis and discussions to a particular family of p.d.f.'s/p.m.f.'s called the natural exponential family (NEF), for which the global convergence properties can be established. We start by stating the definition of NEF and some regularity conditions.

Definition 3.1 *A parameterized family of p.d.f's $\{f(\cdot, \theta), \ \theta \in \Theta \subseteq \Re^m\}$ on $\mathcal{X}$ is said to belong to the natural exponential family (NEF) if there exist functions $h(\cdot) : \Re^n \to \Re$, $\Gamma(\cdot) : \Re^n \to \Re^m$, and $K(\cdot) : \Re^m \to \Re$ such that*

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

where $K(\theta) = \ln \int_{x \in \mathcal{X}} \exp\left\{\theta^T \Gamma(x)\right\} h(x) dx$, and the superscript "T" denotes the vector transposition.

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.

Assumptions:

A1. *There exists a compact set $\Pi \subseteq \mathcal{X}$ such that the level set $\{x : H(x) \geq \bar{\gamma}_1\} \subseteq \Pi$, where $\bar{\gamma}_1 = \sup_l \{l : P_{\theta_0}(H(X) \geq l) \geq \rho\}$ is defined as in the MRAS$_0$ algorithm.*

A2. *For any given constant $\xi < H(x^*)$, the set $\{x : H(x) \geq \xi\}$ has a strictly positive Lebesgue measure.*

A3. *For any given constant $\delta > 0$, $\sup_{x \in A_\delta} H(x) < H(x^*)$, where $A_\delta := \{x : \|x - x^*\| \geq \delta\}$.*

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

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 3.1.*

A6. *$\Gamma(\cdot) : \Re^m \to \Re^n$ given in Definition 3.1 is a continuous mapping.*

Remark 1: Assumptions A1–A3 are regularity conditions imposed on the optimization problem to be solved, whereas assumptions A4–A6 are restrictions imposed on the parameterized family of p.d.f.'s. A1 is satisfied if the function $H(\cdot)$ has compact level sets or the solution space $\mathcal{X}$ is compact. Intuitively, assumption A2 ensures that any neighborhood of the optimal solution x^* will have a positive probability of being sampled; it is satisfied if the objective function $H(\cdot)$ is continuous at x^*. Since $H(\cdot)$ has a unique global optimizer, A3 is satisfied by many functions encountered in practice, and is guaranteed to hold if $\mathcal{X}$ itself is compact. In actual implementation of the algorithm, step 3 of MRAS$_0$ is often posed as an unconstrained optimization problem, i.e., $\Theta = \Re^m$, in which case A4 is automatically satisfied. It is also easy to verify that A5 and A6 are satisfied by most NEFs.

To show the convergence of MRAS$_0$, we will need the following key observation.

Lemma 3.2 *If assumptions A1, A2, and A4–A6 hold, then we have*

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

where recall that $E_{\theta_{k+1}}[\cdot]$ and $E_{g_{k+1}}[\cdot]$ denote the expectations with respect to the p.d.f's $f(\cdot, \theta_{k+1})$ and $g_{k+1}(\cdot)$, respectively.

Remark 2: We only prove Lemma 3.2 in continuous domains. The proof for discrete/combinatorial domains follows easily.

Proof: Define $J_k(\theta, \bar\gamma_{k+1}) := \int_{x \in \mathcal{X}} [S(H(x))]^k I_{\{H(x) \geq \bar\gamma_{k+1}\}} \ln f(x, \theta) dx$. Since $f(\cdot, \theta)$ belongs to the NEF, we can write

$$J_k(\theta, \bar\gamma_{k+1}) = \int_{x \in \mathcal{X}} [S(H(x))]^k I_{\{H(x) \geq \bar\gamma_{k+1}\}} \ln h(x) dx + \int_{x \in \mathcal{X}} [S(H(x))]^k I_{\{H(x) \geq \bar\gamma_{k+1}\}} \theta^T \Gamma(x) dx$$

$$- \int_{x \in \mathcal{X}} [S(H(x))]^k I_{\{H(x) \geq \bar\gamma_{k+1}\}} \ln \left[\int_{x \in \mathcal{X}} \exp\left(\theta^T \Gamma(x)\right) h(x) dx \right] dx.$$

Thus the gradient of $J_k(\theta, \bar\gamma_{k+1})$ with respect to θ can be expressed as

$$\nabla_\theta J_k(\theta, \bar\gamma_{k+1}) = \int [S(H(x))]^k I_{\{H(x) \geq \bar\gamma_{k+1}\}} \Gamma(x) dx - \frac{\int e^{\theta^T \Gamma(x)} \Gamma(x) h(x) dx}{\int e^{\theta^T \Gamma(x)} h(x) dx} \int [S(H(x))]^k I_{\{H(x) \geq \bar\gamma_{k+1}\}} dx,$$

where the validity of the interchange of derivative and integral above is guaranteed by assumptions A5 and the dominated convergence theorem; see e.g., Rubinstein and Shapiro (1993) for further details.

By assumptions A1, A6, and the non-decreasing property of the sequence $\{\bar\gamma_k\}$, it turns out that the above gradient $\nabla_\theta J_k(\theta, \bar\gamma_{k+1})$ is finite and thus well-defined. Moreover, for continuous optimization problems, it can be seen from the MRAS$_0$ algorithm that $\bar\gamma_{k+1} < H(x^*)$, $\forall \, k = 0, 1, \ldots$. Thus by A2, the set $\{x : H(x) \geq \bar\gamma_{k+1}\}$ will have a strictly positive Lebesgue measure. It follows that we must have $\int_{x \in \mathcal{X}} [S(H(x))]^k I_{\{H(x) \geq \bar\gamma_{k+1}\}} dx > 0$.

By setting $\nabla_\theta J_k(\theta, \bar\gamma_{k+1}) = 0$, it immediately follows that

$$\int_{x \in \mathcal{X}} \frac{[S(H(x))]^k I_{\{H(x) \geq \bar\gamma_{k+1}\}} \Gamma(x)}{\int [S(H(x))]^k I_{\{H(x) \geq \bar\gamma_{k+1}\}} dx} dx = \int_{x \in \mathcal{X}} \frac{e^{\theta^T \Gamma(x)} h(x) \Gamma(x)}{\int e^{\theta^T \Gamma(x)} h(x) dx} dx,$$

and by definitions of $g_{k+1}(\cdot)$ (cf. proof of Lemma 3.1) and $f(\cdot, \theta)$, we have

$$E_{g_{k+1}}[\Gamma(X)] = E_\theta[\Gamma(X)]. \tag{5}$$

By assumption A4, since θ_{k+1} is the optimal solution of the problem

$$\arg\max_\theta J_k(\theta, \bar{\gamma}_{k+1}),$$

it must satisfy equation (5). Therefore we conclude that

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

∎

We have the following convergence result for the MRAS_0 algorithm.

Theorem 3.1 (Continuous Optimization) *Let* $\{\theta_k, \ k = 1, 2, \dots\}$ *be the sequence of parameters generated by* MRAS_0*. If* $\varepsilon > 0$ *and assumptions* $A1-A6$ *are satisfied, then*

$$\lim_{k \to \infty} E_{\theta_k}[\Gamma(X)] = \Gamma(x^*). \tag{6}$$

Remark 3: The convergence result in Theorem 3.1 is much stronger than it appears to be. For example, when $\Gamma(x)$ is a one-to-one function (which is the case for many NEFs encountered in practice), the convergence result (6) can be equivalently written as $\Gamma^{-1}\left(\lim_{k \to \infty} E_{\theta_k}[\Gamma(X)]\right) = x^*$. Also note that the limit in equation (6) is component-wise. 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 disregard the redundant components and interpret equation (6) as $\lim_{k \to \infty} E_{\theta_k}[X] = x^*$. 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. of the form

$$f(x_i, \vartheta_i) = \exp(x_i \vartheta_i - K(\vartheta_i)) h(x_i), \ \vartheta_i \in \Re, \forall \ i = 1, \dots, n.$$

In this case, since the p.d.f. of the random vector X is simply the product of the marginal p.d.f.'s, we will clearly have $\Gamma(x) = x$. Thus, equation (6) is again equivalent to $\lim_{k \to \infty} E_{\theta_k}[X] = x^*$, where $\theta_k := (\vartheta_1^k, \dots, \vartheta_n^k)$, and ϑ_i^k is the value of ϑ_i at the kth iteration.

Lemma 3.2 already established a relationship between the reference models $\{g_k\}$ and the sequence of actual models $\{f(\cdot, \theta_k)\}$. Therefore, proving Theorem 3.1 amounts to showing $\lim_{k \to \infty} E_{g_k}[\Gamma(X)] = \Gamma(x^*)$.

Proof: Recall from Lemma 3.1 that g_{k+1} can be expressed recursively as

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

Thus

$$E_{g_{k+1}}\left[S(H(X)) I_{\{H(X) \geq \bar{\gamma}_{k+1}\}}\right] = \frac{E_{g_k}\left[[S(H(X))]^2 I_{\{H(X) \geq \bar{\gamma}_{k+1}\}}\right]}{E_{g_k}\left[S(H(X)) I_{\{H(X) \geq \bar{\gamma}_{k+1}\}}\right]}$$

$$\geq E_{g_k}\left[S(H(X)) I_{\{H(X) \geq \bar{\gamma}_{k+1}\}}\right]. \tag{7}$$

Let $\mathcal{K} \geq 1$, and consider the sequence $\{\bar{\gamma}_k, \ 1 \leq k \leq \mathcal{K}\}$. Clearly we have $\bar{\gamma}_1 \leq \bar{\gamma}_2 \leq \cdots \leq \bar{\gamma}_{\mathcal{K}}$. Now denote $M_{\mathcal{K}}$ as the number of jumps in the sequence $\{\bar{\gamma}_k\}$, i.e., the number of times such that $\bar{\gamma}_{k+1} > \bar{\gamma}_k, \quad \forall \ 1 \leq$

8

$k < \mathcal{K}$. Since each increment in the sequence $\{\bar{\gamma}_k\}$ is lower bounded by the quantity $\varepsilon > 0$, it follows that $\lim_{\mathcal{K} \to \infty} M_{\mathcal{K}} < \infty$. Therefore, there exists $\mathcal{N} > 0$ finite such that

$$\bar{\gamma}_{k+1} = \bar{\gamma}_k, \quad \forall\, k \geq \mathcal{N}.$$

From equation (7), it follows that

$$E_{g_{k+1}}\left[S(H(X))I_{\{H(X)\geq\bar{\gamma}_{k+2}\}}\right] \geq E_{g_k}\left[S(H(X))I_{\{H(X)\geq\bar{\gamma}_{k+1}\}}\right], \quad \forall\, k \geq \mathcal{N} - 1, \tag{8}$$

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

$$S_* := \lim_{k\to\infty} E_{g_k}\left[S(H(X))I_{\{H(X)\geq\bar{\gamma}_{k+1}\}}\right] < S^* := S(H(x^*)). \tag{9}$$

Define the set $\mathcal{A}$ as

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

From MRAS$_0$, it is not difficult to see that $\bar{\gamma}_{\mathcal{N}} < S^*$; thus the set $\mathcal{A}$ is non-empty. Furthermore, since $S(\cdot)$ is continuous and strictly increasing, its inverse $S^{-1}(\cdot)$ exists. Thus $\mathcal{A}$ can be reformulated as

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

By A2, $\mathcal{A}$ has a strictly positive Lebesgue measure.

Notice that g_k can be rewritten as

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

Since $\lim_{k\to\infty} \frac{S(H(x))I_{\{H(x)\geq\bar{\gamma}_{k+1}\}}}{E_{g_k}\left[S(H(X))I_{\{H(X)\geq\bar{\gamma}_{k+1}\}}\right]} = \frac{S(H(x))I_{\{H(x)\geq\bar{\gamma}_{\mathcal{N}}\}}}{S_*} > 1, \; \forall x \in \mathcal{A}$, we conclude that

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

Thus, by Fatou's lemma, we have

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

which is a contradiction. Hence, it follows that

$$\lim_{k\to\infty} E_{g_k}\left[S(H(X))I_{\{H(X)\geq\bar{\gamma}_{k+1}\}}\right] = S^*. \tag{10}$$

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

$$\|E_{g_k}[\Gamma(X)] - \Gamma(x^*)\| \leq \int_{x\in\mathcal{X}} \|\Gamma(x) - \Gamma(x^*)\| g_k(x)dx$$

$$= \int_{\mathcal{C}} \|\Gamma(x) - \Gamma(x^*)\| g_k(x)dx, \tag{11}$$

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

By assumption A6, 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 A3, we have from equation (11),

$$\|E_{g_k}[\Gamma(X)] - \Gamma(x^*)\| \leq \int_{A_\delta^c \cap \mathcal{C}} \|\Gamma(x) - \Gamma(x^*)\| g_k(x) dx + \int_{A_\delta \cap \mathcal{C}} \|\Gamma(x) - \Gamma(x^*)\| g_k(x) dx$$

$$\leq \zeta + \int_{A_\delta \cap \mathcal{C}} \|\Gamma(x) - \Gamma(x^*)\| g_k(x) dx, \quad \forall\, k \geq \mathcal{N}. \tag{12}$$

The rest of the proof amounts to showing that the second term in equation (12) is also bounded. Clearly by A1 and A6, 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 A3, we have

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

Define $S_\delta := S^* - S(\sup_{x \in A_\delta} H(x))$. Since $S(\cdot)$ is strictly increasing, we have $S_\delta > 0$. It is easy to see that

$$S(H(x)) \leq S^* - S_\delta, \quad \forall x \in A_\delta \cap \mathcal{C}. \tag{13}$$

From equations (8) and (10), there exists $\bar{\mathcal{N}} \geq \mathcal{N}$ such that $\forall\, k \geq \bar{\mathcal{N}}$

$$E_{g_k}\left[S(H(X)) I_{\{H(X) \geq \bar{\gamma}_{k+1}\}}\right] \geq S^* - \frac{1}{2} S_\delta. \tag{14}$$

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

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

Thus, it follows from equations (13) and (14) that

$$g_k(x) \leq \left(\frac{S^* - S_\delta}{S^* - \frac{1}{2} S_\delta}\right)^{k - \bar{\mathcal{N}}} \cdot g_{\bar{\mathcal{N}}}(x), \quad \forall\, x \in A_\delta \cap \mathcal{C}, \quad \forall\, k \geq \bar{\mathcal{N}}.$$

Therefore,

$$\|E_{g_k}[\Gamma(X)] - \Gamma(x^*)\| \leq \zeta + \sup_{x \in A_\delta \cap \mathcal{C}} \|\Gamma(x) - \Gamma(x^*)\| \int_{A_\delta \cap \mathcal{C}} g_k(x) dx$$

$$\leq \zeta + \sup_{x \in A_\delta \cap \mathcal{C}} \|\Gamma(x) - \Gamma(x^*)\| \left(\frac{S^* - S_\delta}{S^* - \frac{1}{2} S_\delta}\right)^{k - \bar{\mathcal{N}}}, \quad \forall\, k \geq \bar{\mathcal{N}}$$

$$\leq \zeta + \sup_{x \in A_\delta \cap \mathcal{C}} \|\Gamma(x) - \Gamma(x^*)\| \zeta, \quad \forall\, k \geq \widehat{\mathcal{N}}$$

$$= \left(1 + \sup_{x \in A_\delta \cap \mathcal{C}} \|\Gamma(x) - \Gamma(x^*)\|\right) \zeta, \quad \forall\, k \geq \widehat{\mathcal{N}},$$

where $\widehat{\mathcal{N}}$ is defined as $\widehat{\mathcal{N}} := \max\left\{\bar{\mathcal{N}}, \lceil \bar{\mathcal{N}} + \ln \zeta / \ln\left(\frac{S^* - S_\delta}{S^* - \frac{1}{2} S_\delta}\right) \rceil\right\}$.

Since ζ is arbitrary, we have

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

The proof is completed by applying Lemma 3.2. ∎

We now address some of the special cases discussed in Remark 3.

Corollary 3.1 (Multivariate Normal) *If multivariate normal p.d.f.'s are used in MRAS$_0$, i.e.,*

$$f(x, \theta_k) = \frac{1}{\sqrt{(2\pi)^n |\Sigma_k|}} \exp\left(-\frac{1}{2}(x - \mu_k)^T \Sigma_k^{-1}(x - \mu_k)\right), \tag{15}$$

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

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

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

Proof: By Lemma 3.2, it is easy to show that

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

and

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

The rest of the proof amounts to showing that

$$\lim_{k\to\infty} E_{g_k}(X) = x^*, \quad and \quad \lim_{k\to\infty} E_{g_k}\left[(X - \mu_k)(X - \mu_k)^T\right] = 0_{n\times n},$$

which is the same as the proof of Theorem 3.1. ∎

Remark 4: Corollary 3.1 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

$$\mu_{k+1} = \frac{E_{\theta_k}\left[\{[S(H(X))]^k / f(X, \theta_k)\} I_{\{H(X) \geq \bar{\gamma}_{k+1}\}} X\right]}{E_{\theta_k}\left[\{[S(H(X))]^k / f(X, \theta_k)\} I_{\{H(X) \geq \bar{\gamma}_{k+1}\}}\right]}, \tag{16}$$

and

$$\Sigma_{k+1} = \frac{E_{\theta_k}\left[\{[S(H(X))]^k / f(X, \theta_k)\} I_{\{H(X) \geq \bar{\gamma}_{k+1}\}} (X - \mu_{k+1})(X - \mu_{k+1})^T\right]}{E_{\theta_k}\left[\{[S(H(X))]^k / f(X, \theta_k)\} I_{\{H(X) \geq \bar{\gamma}_{k+1}\}}\right]}, \tag{17}$$

where $f(x, \theta_k)$ is given by (15).

Corollary 3.2 (Independent Univariate) *If the components of the random vector $X = (X_1, \ldots, X_n)$ are independent, each has a univariate p.d.f. of the form*

$$f(x_i, \vartheta_i) = \exp(x_i \vartheta_i - K(\vartheta_i)) h(x_i), \quad \vartheta_i \in \Re, \forall\, i = 1, \ldots, n,$$

$\varepsilon > 0$, and assumptions A1$-$A6 are satisfied, then

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

Note that the convergence of MRAS$_0$ for discrete optimization problems with *infinite countable* domains can be shown similarly by following the proof of Theorem 3.1, provided that all assumptions A1$-$A6 except A2 are satisfied. However unlike the continuous case, there is one more complication, because the limit of the sequence $\{\bar{\gamma}_k\}$ (i.e., $\bar{\gamma}_{\mathcal{N}}$ in the proof of Theorem 3.1) may be equal to $H(x^*)$. The same issue also arises in problems with *finite* solution spaces. In order to avoid unnecessary repetitions, we will only discuss this issue for the latter case, as we believe that these types of problems are more representative. It is worth mentioning that for problems with *finite* solution spaces, assumptions A1 and A3 are automatically satisfied. Furthermore, the input parameter ε need not be strictly positive as in the continuous case.

Theorem 3.2 (Discrete Optimization with Finite Domain) *If the solution space $\mathcal{X}$ is finite and assumptions A4 and A5 are satisfied, then*

$$\lim_{k\to\infty} E_{\theta_k}\left[\Gamma(X)\right] = \Gamma(x^*).$$

Proof: First note that if we take the input parameter $\varepsilon = 0$, then step 2 of MRAS$_0$ is equivalent to

$$\bar{\gamma}_{k+1} = \max_{1\le i \le k+1} \gamma_i.$$

Thus $\{\bar{\gamma}_k\}$ is non-decreasing and each strict increase in the sequence is bounded from below by

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

On the other hand, if we take $\varepsilon > 0$, then each strict increase in $\{\bar{\gamma}_k\}$ is lower bounded by ε (see the proof of Theorem 3.1). In either case, the total number of jumps in $\{\bar{\gamma}_k\}$ can only be finite. Therefore there exists $\mathcal{N} > 0$ finite such that

$$\bar{\gamma}_{k+1} = \bar{\gamma}_k, \quad \forall\, k \ge \mathcal{N}.$$

If $\bar{\gamma}_\mathcal{N} > \max_{x\ne x^*} H(x)$, then from the definition of $g_{k+1}(\cdot)$ (see Lemma 3.1), we obviously have

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

and

$$g_{k+1}(x^*) = \frac{[S(H(x^*))]^k I_{\left\{H(x)>\max_{x\ne x^*} H(x)\right\}}}{\sum_x [S(H(x))]^k I_{\left\{H(x)>\max_{x\ne x^*} H(x)\right\}}} = 1 \ \ \forall\, k \ge \mathcal{N}.$$

Hence it follows immediately from Lemma 3.2 that

$$E_{\theta_{k+1}}\left[\Gamma(X)\right] = E_{g_{k+1}}\left[\Gamma(X)\right] = \Gamma(x^*) \ \ \forall\, k \ge \mathcal{N}.$$

The proof for $\bar{\gamma}_\mathcal{N} \le \max_{x\ne x^*} H(x)$ case is similar to the proof for the continuous case and is omitted here. ∎

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 MRAS method. 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 CE$_0$: Deterministic Version of the CE Method

1. *Choose the initial p.d.f./p.m.f. $f(\cdot, \theta_0)$, $\theta_0 \in \Theta$. Specify the parameter $\rho \in (0,1]$, a non-decreasing function $\varphi(\cdot): \Re \to \Re^+ \cup \{0\}$, and a positive integer d. Set $k = 0$.*

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

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

3. *Compute the new parameter*

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

12

4. If for some k,

$$\gamma_{k+1} = \gamma_k = \cdots = \gamma_{k+1-d},$$

then terminate the algorithm; otherwise set $k = k+1$ and go to Step 2.

In CE_0, 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. e.g., de Boer et al. 2004).

One resemblance between CE and MRAS 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. Since 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. 2004, Homem-de-Mello and Rubinstein 2003; also example 4.1 below), whereas the theoretical convergence of MRAS 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 $\{g_k^{ce}(\cdot) : k = 1, 2, \ldots\}$, the CE method can be recovered, and the parameter updating in CE is guided by this sequence of models.

Lemma 4.1 *The parameter θ_{k+1} computed at the kth iteration of the CE_0 algorithm minimizes the KL-divergence $\mathcal{D}\left(g_{k+1}^{ce}, f(\cdot, \theta)\right)$, where*

$$g_{k+1}^{ce}(x) := \frac{\varphi(H(x))I_{\{H(x)\geq\gamma_{k+1}\}}f(x,\theta_k)}{E_{\theta_k}\left[\varphi(H(X))I_{\{H(X)\geq\gamma_{k+1}\}}\right]} \quad \forall x \in \mathcal{X}, \quad k = 0, 1, \ldots. \tag{18}$$

Proof: Similar to the proof of Lemma 3.1. ∎

The key observation to note is that in contrast to MRAS, 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 (cf. Lemma 3.1). Since $g_{k+1}^{ce}(\cdot)$ is obtained by tilting $f(\cdot, \theta_k)$ with the performance function, it improves the expected performance in the sense that

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

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.

Theorem 4.1 *For the CE_0 algorithm, we have*

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

Proof: Define $\widehat{g}_{k+2}^{ce}(\cdot)$ as

$$\widehat{g}_{k+2}^{ce}(x) := \frac{\varphi(H(x))I_{\{H(x)\geq\gamma_{k+1}\}}f(x,\theta_{k+1})}{E_{\theta_{k+1}}\left[\varphi(H(X))I_{\{H(X)\geq\gamma_{k+1}\}}\right]} \quad \forall\, x \in \mathcal{X}, \ k = 0, 1, \ldots.$$

We have from the definition of $g_{k+1}^{ce}(\cdot)$,

$$\mathcal{D}(g_{k+1}^{ce}, \widehat{g}_{k+2}^{ce}) = E_{g_{k+1}^{ce}}\left[\ln\frac{g_{k+1}^{ce}(X)}{\widehat{g}_{k+2}^{ce}(X)}\right]$$

$$= E_{g_{k+1}^{ce}}\left[\ln\frac{f(X,\theta_k)}{f(X,\theta_{k+1})}\right] + \ln\frac{E_{\theta_{k+1}}[\varphi(H(X))I_{\{H(X)\geq\gamma_{k+1}\}}]}{E_{\theta_k}[\varphi(H(X))I_{\{H(X)\geq\gamma_{k+1}\}}]}.$$

Since θ_{k+1} minimizes the K-L divergence $\mathcal{D}(g_{k+1}^{ce}, f(\cdot, \theta))$ (cf. Lemma 4.1), it follows that

$$
\begin{aligned}
0 &\leq \mathcal{D}(g_{k+1}^{ce}, f(\cdot, \theta_k)) - \mathcal{D}(g_{k+1}^{ce}, f(\cdot, \theta_{k+1})) \\
&\leq \mathcal{D}(g_{k+1}^{ce}, f(\cdot, \theta_k)) - \mathcal{D}(g_{k+1}^{ce}, f(\cdot, \theta_{k+1})) + \mathcal{D}(g_{k+1}^{ce}, \widehat{g}_{k+2}^{ce}) \\
&= E_{g_{k+1}^{ce}}\left[\ln \frac{f(X, \theta_{k+1})}{f(X, \theta_k)}\right] + E_{g_{k+1}^{ce}}\left[\ln \frac{f(X, \theta_k)}{f(X, \theta_{k+1})}\right] + \ln \frac{E_{\theta_{k+1}}[\varphi(H(X))I_{\{H(X) \geq \gamma_{k+1}\}}]}{E_{\theta_k}[\varphi(H(X))I_{\{H(X) \geq \gamma_{k+1}\}}]} \\
&= \ln \frac{E_{\theta_{k+1}}[\varphi(H(X))I_{\{H(X) \geq \gamma_{k+1}\}}]}{E_{\theta_k}[\varphi(H(X))I_{\{H(X) \geq \gamma_{k+1}\}}]}
\end{aligned}
$$

Therefore

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

∎

In the standard CE method, Theorem 4.1 implies the monotonicity of the sequence $\{\gamma_k : k = 1, 2, \ldots\}$.

Lemma 4.2 *For the standard CE method (i.e., CE_0 with $\varphi(H(x)) = 1$), we have*

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

Proof: By Theorem 4.1, 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}. ∎

Note that since $\gamma_k \leq H(x^*)$ for all k, Lemma 4.2 implies that the sequence $\{\gamma_k : k = 1, 2, \ldots\}$ 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 MRAS (cf. Lemma 3.2), when $f(\cdot, \theta)$ belongs to the natural exponential families, the following lemma relates the sequence $\{f(\cdot, \theta_k), \ k = 1, 2, \ldots\}$ to the sequence of reference models $\{g_k^{ce}(\cdot) : k = 1, 2, \ldots\}$.

Lemma 4.3 *Assume that:*

1. *There exists a compact set $\bar{\Pi} \subseteq \mathcal{X}$ such that the level set $\{x : H(x) \geq \gamma_k\} \subseteq \bar{\Pi}$ for all $k = 1, 2, \ldots$, where $\gamma_k = \sup_l\{l : P_{\theta_{k-1}}(H(X) \geq l) \geq \rho\}$ is defined as in the CE_0 algorithm.*

2. *The parameter θ_{k+1} computed at step 3 of the CE_0 algorithm is an interior point of Θ for all k.*

3. *Assumptions A2, A5, and A6 are satisfied.*

Then

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

The above lemma indicates that the behavior of the sequence of p.d.f's/p.m.f's $\{f(\cdot, \theta_k)\}$ are 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 $\lim_{k \to \infty} E_{\theta_k}[H(X)] = H(x^*)$, then we must have $\lim_{k \to \infty} E_{\theta_k}[X] = x^*$, since $H(x) < H(x^*) \, \forall \, x \neq x^*$. Thus, a necessary condition for this convergence is $\lim_{k \to \infty} E_{g_k^{ce}}[X] = x^*$. However, unlike MRAS$_0$, where the convergence of the sequence of reference models to an optimal degenerate p.d.f. is guaranteed, the convergence of the sequence

14

$\{g_k^{ce}(\cdot) : k = 1, 2, \ldots\}$ relies on the choices of the families of p.d.f's/p.m.f's $\{f(\cdot, \theta_k)\}$ and the values of the parameter ρ used (cf. equation (18)). We now illustrate this issue by two simple examples.

Example 4.1 (The Standard CE Method) Consider maximizing the function $H(x)$ given by

$$H(x) = \begin{cases} 0 & x \in \{(0,1), (1,0)\}, \\ 1 & x = (0,0), \\ a & x = (1,1), \end{cases} \tag{19}$$

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

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

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

then since $P_{\theta_0}(x \in \{(0,0), (1,1)\}) = 0.5 \ge \rho$, 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 CE_0 is given by $\theta_1 = (0.5, 0.5)$. Proceeding iteratively, we have $\gamma_k = 1$ and $g_k^{ce}(x) = g_1^{ce}(x)$ for all $k = 1, 2, \ldots$, i.e., the algorithm does not converge to a degenerate distribution at the optimal solution.

On the other hand, if we choose $\rho \le 0.25$, 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 $k = 1, 2, \ldots$, which means the algorithm converges to the optimum.

Example 4.2 (The Extended Version of the CE Method) Consider solving problem (19) by CE_0 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 $\theta_0 = (\frac{1}{1+a}, \frac{1}{1+a})$. If the values of ρ are chosen from the interval $(\frac{1}{(1+a)^2}, \frac{a^2+1}{(1+a)^2})$, then we have $\theta_k = (\frac{1}{1+a}, \frac{1}{1+a})$, $\gamma_k = 1$, 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 $k = 1, 2, \ldots$.

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

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

5 Model Reference Adaptive Search (Monte Carlo Version)

The $MRAS_0$ 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 $MRAS_0$ will be replaced with

$$\widetilde{\theta}_{k+1} = \arg\max_{\theta \in \Theta} \frac{1}{N} \sum_{i=1}^{N} \frac{[S(H(X_i))]^k}{f(X_i, \widetilde{\theta}_k)} I_{\{H(X_i) \ge \widetilde{\gamma}_{k+1}\}} \ln f(X_i, \theta), \tag{20}$$

15

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

However, the theoretical convergence can no longer be guaranteed for a simple stochastic counterpart of MRAS_0. In particular, the set $\{x : H(x) \geq \overline{\gamma}_{k+1}\}$ involved in equation (20) may be empty, since 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 MRAS_0 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, whether "good" or "bad", will be used to update the probabilistic model, which could slow down the convergence process. On the other hand, since a good estimate will necessarily require a reasonable amount of valid samples, the quantity ρN (i.e., the approximate amount 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 1 or too close to 0, to determine a priori which ρ gives a satisfactory performance may be difficult.

In order to address the above difficulties, we adopt the same idea as in Homem-de-Mello and Rubinstein (2003) and propose a modified Monte Carlo version of MRAS_0 in which the sample size N is adaptively increasing and the parameter ρ is adaptively decreasing.

5.1 Algorithm Description

Algorithm MRAS$_1$: Model Reference Adaptive Search – 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 continuous and strictly increasing function $S(\cdot) : \Re \to \Re^+$, and an initial p.d.f. $f(x, \theta_0) > 0 \ \forall x \in \mathcal{X}$. Set $\widetilde{\theta}_0 \leftarrow \theta_0$, $k \leftarrow 0$.

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

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

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

    3. **If** $k = 0$ **or** $\widetilde{\gamma}_{k+1}(\rho_k, N_k) \geq \overline{\gamma}_k + \frac{\varepsilon}{2}$, **then**

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

        else, find the largest $\overline{\rho} \in (0, \rho_k)$ such that $\widetilde{\gamma}_{k+1}(\overline{\rho}, N_k) \geq \overline{\gamma}_k + \frac{\varepsilon}{2}$.

        3b. **If** such a $\overline{\rho}$ exists, **then** set $\overline{\gamma}_{k+1} \leftarrow \widetilde{\gamma}_{k+1}(\overline{\rho}, N_k)$, $\rho_{k+1} \leftarrow \overline{\rho}$, $N_{k+1} \leftarrow N_k$.

        3c. **else** (if no such $\overline{\rho}$ exists), set $\overline{\gamma}_{k+1} \leftarrow \overline{\gamma}_k$, $\rho_{k+1} \leftarrow \rho_k$, $N_{k+1} \leftarrow \lceil \alpha N_k \rceil$.

        endif

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

$$\widetilde{\theta}_{k+1} = \arg\max_{\theta \in \Theta} \frac{1}{N_k} \sum_{i=1}^{N_k} \frac{[S(H(X_i^k))]^k}{\widetilde{f}(X_i^k, \widetilde{\theta}_k)} I_{\{H(X_i^k) \geq \overline{\gamma}_{k+1}\}} \ln f(X_i^k, \theta). \qquad (21)$$

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

Roughly speaking, the MRAS_1 algorithm is essentially a Monte Carlo version of MRAS_0 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,

16

if the initial sample size is N_0, then after k increases, 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 $\widetilde{f}(\cdot, \widetilde{\theta}_k)$, which is a mixture of the initial density $f(\cdot, \theta_0)$ and the density calculated from the previous iteration $f(\cdot, \widetilde{\theta}_k)$ (cf. e.g., Auer et al. 2002 for a similar idea in the context of multiarmed bandit models). The initial density $f(\cdot, \theta_0)$ can be chosen according to some prior knowledge of the problem structure (cf. Section 6.2); however, if nothing is known about where the good solutions are, this density should be chosen in such a way that each region in the solution space will have an (approximately) equal probability of being sampled. For instance, when $\mathcal{X}$ is finite, 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, MRAS$_1$ becomes a pure random sampling method.

At step 2, the sample $(1 - \rho_k)$-quantile $\widetilde{\gamma}_{k+1}$ is calculated by first ordering the sample performances $H(X_i^k)$, $i = 1, \ldots, N_k$ from smallest to largest, $H_{(1)} \leq H_{(2)} \leq \cdots \leq H_{(N_k)}$, and then taking the $\lceil (1 - \rho_k) N_k \rceil$th order statistic. We use the function $\widetilde{\gamma}_{k+1}(\rho_k, N_k)$ to emphasize the dependencies of $\widetilde{\gamma}_{k+1}$ 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 MRAS$_1$ is used to extract a sequence of non-decreasing thresholds $\{\bar{\gamma}_k, k = 1, 2 \ldots\}$ from the sequence of sample quantiles $\{\widetilde{\gamma}_k\}$, 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 $\widetilde{\gamma}_{k+1}(\rho_k, N_k) \geq \bar{\gamma}_k + \frac{\varepsilon}{2}$ 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 $3a$ and update the parameter vector $\widetilde{\theta}_{k+1}$ in step 4 by using $\widetilde{\gamma}_{k+1}(\rho_k, N_k)$. 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 $\widetilde{\theta}_{k+1}$ computed at step 4:

Assumption A4′. The parameter vector $\widetilde{\theta}_{k+1}$ computed at step 4 of MRAS$_1$ is an interior point of Θ for all k.

It is important to note that the set $\{x : H(x) \geq \bar{\gamma}_{k+1}, x \in \{X_1^k, \ldots, X_{N_k}^k\}\}$ could be empty if step $3c$ is visited. If this happens, the right hand side of equation (21) will be equal to zero, so any $\theta \in \Theta$ is a maximizer, and we define $\widetilde{\theta}_{k+1} := \widetilde{\theta}_k$ in this case.

5.2 Global Convergence

In this section, we discuss the convergence properties of the MRAS$_1$ algorithm for natural exponential families (NEFs). To be specific, we will explore the relations between MRAS$_1$ and MRAS$_0$ 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 MRAS$_1$ can be ascribed to the convergence analysis of the MRAS$_0$ algorithm; thus, our analysis relies heavily on the results obtained in Section 3.2. Throughout this section, we denote by $P_{\widetilde{\theta}_k}(\cdot)$ and $E_{\widetilde{\theta}_k}[\cdot]$ the respective probability and expectation taken with respect to the p.d.f. $f(\cdot, \widetilde{\theta}_k)$, and $\widetilde{P}_{\widetilde{\theta}_k}(\cdot)$ and $\widetilde{E}_{\widetilde{\theta}_k}[\cdot]$ the respective probability and expectation taken with respect to $\widetilde{f}(\cdot, \widetilde{\theta}_k)$. Note that since the sequence $\{\widetilde{\theta}_k\}$ results from random samples generated at each iteration of MRAS$_1$, these quantities are also random.

Lemma 5.1 *For any given $\rho^\dagger \in (0,1)$, let $\gamma_k^\dagger$ be the set of $(1-\rho^\dagger)$-quantiles of $H(X)$ with respect to the p.d.f. $\widetilde{f}(\cdot,\widetilde{\theta}_k)$, and let $\widetilde{\gamma}_k^\dagger(\rho^\dagger, N_k)$ be the corresponding sample quantile of $H(X_1^k),\ldots,H(X_{N_k}^k)$, where $\widetilde{f}(\cdot,\widetilde{\theta}_k)$ and N_k are defined as in $MRAS_1$, and $X_1^k,\ldots,X_{N_k}^k$ are i.i.d. with common density $\widetilde{f}(\cdot,\widetilde{\theta}_k)$. Then the distance from $\widetilde{\gamma}_k^\dagger(\rho^\dagger, N_k)$ to $\gamma_k^\dagger$ tends to zero as $k \to \infty$ w.p.1.*

Proof: We prove Lemma 5.1 in the Appendix.   ∎

Let $\widetilde{g}_{k+1}(x)$, $k = 0,1,\ldots$, be defined by

$$\widetilde{g}_{k+1}(x) := \begin{cases} \dfrac{[S(H(x))]^k/\widetilde{f}(x,\widetilde{\theta}_k)]I_{\{H(x)\geq\bar{\gamma}_{k+1}\}}}{\sum_{i=1}^{N_k}[[S(H(X_i^k))]^k/\widetilde{f}(X_i^k,\widetilde{\theta}_k)]I_{\{H(X_i^k)\geq\bar{\gamma}_{k+1}\}}} & \text{if } \{x : H(x) \geq \bar{\gamma}_{k+1},\ x \in \{X_1^k,\ldots,X_{N_k}^k\}\} \neq \emptyset, \\ \widetilde{g}_k(x) & \text{otherwise,} \end{cases} \quad (22)$$

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

Similar to Lemma 3.2, the following lemma shows the connection between $f(x,\widetilde{\theta}_{k+1})$ and $\widetilde{g}_{k+1}(x)$. The proof is simple and is omitted.

Lemma 5.2 *If assumptions A4$'$ and A5 hold, then the parameter $\widetilde{\theta}_{k+1}$ computed at step 3 of $MRAS_1$ satisfies*

$$E_{\widetilde{\theta}_{k+1}}\left[\Gamma(X)\right] = E_{\widetilde{g}_{k+1}}\left[\Gamma(X)\right], \quad \forall\, k = 0,1,\ldots,$$

Remark 5: Note that the region $\{x : H(x) \geq \bar{\gamma}_{k+1}\}$ will become smaller and smaller as $\bar{\gamma}_{k+1}$ increases. Lemma 5.2 shows that the sequence of sampling p.d.f's $f(\cdot,\widetilde{\theta}_{k+1})$ 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 $\Gamma(x) = x$. Since $E_{\widetilde{g}_{k+1}}[X]$ is the convex combination of $X_1^k,\ldots,X_{N_k}^k$, the lemma implies that $E_{\widetilde{\theta}_{k+1}}[X] \in \{x : H(x) \geq \bar{\gamma}_{k+1}\}$. 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.

To present the main theorem, we require one more assumption.

Assumption B1. *There exists a compact set Π_ε such that $\{x : H(x) \geq H(x^*) - \varepsilon\} \subseteq \Pi_\varepsilon$. Moreover, $f(x,\theta_0)$ is bounded away from zero on Π_ε, i.e., $f_* := \inf_{x \in \Pi_\varepsilon} f(x,\theta_0) > 0$.*

Theorem 5.1 (Continuous Optimization) *Let $\varepsilon > 0$, and define the ε-optimal set $\mathcal{O}_\varepsilon := \{x : H(x) \geq H(x^*) - \varepsilon\}$. If assumptions A2, A4$'$, A5, and B1 are satisfied, then w.p.1. there exists a random variable $\mathcal{K} > 0$ such that*

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

  2. *$E_{\widetilde{\theta}_{k+1}}\left[\Gamma(X)\right] \in CONV\left\{\Gamma(\mathcal{O}_\varepsilon)\right\}$, $\forall\, k \geq \mathcal{K}$ w.p.1, where $CONV\left\{\Gamma(\mathcal{O}_\varepsilon)\right\}$ indicates the convex hull of the set $\Gamma(\mathcal{O}_\varepsilon)$.*

Furthermore, let β be a positive constant satisfying the condition that the set $\{x : S(H(x)) \geq \frac{1}{\beta}\}$ has a strictly positive Lebesgue measure. If assumptions A2, A3, A4$'$, A5, A6, and B1 are satisfied and $\alpha > (\beta S^*)^2$, where $S^* := S(H(x^*))$, then

  3. *$\lim_{k\to\infty} E_{\widetilde{\theta}_k}\left[\Gamma(X)\right] = \Gamma(x^*)$ w.p.1.*

Remark 6: 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 $\mathcal{O}_\varepsilon$. Let $x, y \in \mathcal{O}_\varepsilon$ 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 $\mathcal{O}_\varepsilon$ is convex. If in addition $\Gamma(x)$ is also convex and one-to-one on $\mathcal{O}_\varepsilon$ (e.g. multivariate normal p.d.f.), then $CONV\{\Gamma(\mathcal{O}_\varepsilon)\} = \Gamma(\mathcal{O}_\varepsilon)$. Thus it follows that $\Gamma^{-1}(E_{\widetilde{g}_{k+1}}[\Gamma(X)]) \in \mathcal{O}_\varepsilon, \ \forall\, k \geq \mathcal{K} \ w.p.1$.

Proof: **(1)** The first part of the proof is similar to the proofs given in Homem-de-Mello and Rubinstein (2003). First we claim that given ρ_k and $\bar{\gamma}_k$, if $\bar{\gamma}_k \leq H(x^*) - \varepsilon$, then there exist $\bar{\mathcal{K}} < \infty$ w.p.1 and $\bar{\rho} \in (0, \rho_k)$ such that $\widetilde{\gamma}_{k'+1}(\bar{\rho}, N_{k'}) \geq \bar{\gamma}_k + \frac{\varepsilon}{2}$ for all $k' \geq \bar{\mathcal{K}}$. To show this, we proceed by contradiction.

Let $\rho_k^* := \widetilde{P}_{\widetilde{\theta}_k}\left(H(X) \geq \bar{\gamma}_k + \frac{2\varepsilon}{3}\right)$. If $\bar{\gamma}_k \leq H(x^*) - \varepsilon$, then $\bar{\gamma}_k + \frac{2\varepsilon}{3} \leq H(x^*) - \frac{\varepsilon}{3}$. By assumptions A2 and B1, we have

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

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

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

$$\widetilde{P}_{\widetilde{\theta}_k}\left(H(X) \geq \gamma_{k+1}(\rho, \widetilde{\theta}_k)\right) \geq \rho, \ \text{ and}$$
$$\widetilde{P}_{\widetilde{\theta}_k}\left(H(X) \leq \gamma_{k+1}(\rho, \widetilde{\theta}_k)\right) \geq 1 - \rho > 1 - \rho_k^*. \tag{23}$$

It follows that $\widetilde{P}_{\widetilde{\theta}_k}\left(H(X) \leq \gamma_{k+1}(\rho, \widetilde{\theta}_k)\right) \leq \widetilde{P}_{\widetilde{\theta}_k}\left(H(X) < \bar{\gamma}_k + \frac{2\varepsilon}{3}\right) = 1 - \rho_k^*$ by the definition of ρ_k^*, which contradicts equation (23); thus we must have

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

Therefore, there exists $\bar{\rho} \in (0, \min\{\rho_k, \lambda \mathcal{C}(\varepsilon, \theta_0)\})$ such that $\gamma_{k+1}(\bar{\rho}, \widetilde{\theta}_k) \geq \bar{\gamma}_k + \frac{2\varepsilon}{3}$ whenever $\bar{\gamma}_k \leq H(x^*) - \varepsilon$. By Lemma 5.1, the distance from the sample $(1-\bar{\rho})$-quantile $\widetilde{\gamma}_{k+1}(\bar{\rho}, N_k)$ to the set of $(1-\bar{\rho})$-quantiles $\gamma_{k+1}(\bar{\rho}, \widetilde{\theta}_k)$ goes to zero as $k \to \infty$ w.p.1, thus there exists a finite $\bar{\mathcal{K}}$ w.p.1 such that $\widetilde{\gamma}_{k'+1}(\bar{\rho}, N_{k'}) \geq \bar{\gamma}_k + \frac{\varepsilon}{2}$ for all $k' \geq \bar{\mathcal{K}}$.

Notice that from the MRAS$_1$ algorithm, if neither step 3a or 3b is visited at the kth 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 3a/3b will be visited after a finite number of iterations. Furthermore, since the total number of visits to steps 3a and 3b is finite (i.e., bounded by $\frac{2[H(x^*) - \mathcal{M}]}{\varepsilon}$, where recall that $\mathcal{M}$ is a lower bound for $H(x)$), we conclude that there exists $\mathcal{K} > 0$ w.p.1, such that

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

(2) From the MRAS$_1$ algorithm, it is easy to see that $\bar{\gamma}_{k+1} \geq \bar{\gamma}_k, \ \forall\, k = 0, 1, \ldots$. By part **(1)**, we have $\bar{\gamma}_{k+1} \geq H(x^*) - \varepsilon, \ \forall\, k \geq \mathcal{K}$ w.p.1. Thus, from the definition of $\widetilde{g}_{k+1}(x)$ (cf. equation (22)), it follows immediately that if $\{x : H(x) \geq \bar{\gamma}_{k+1}, \ x \in \{X_1^k, \ldots, X_{N_k}^k\}\} \neq \emptyset$, then the support of $\widetilde{g}_{k+1}(x)$ satisfies $supp\{\widetilde{g}_{k+1}\} \subseteq \mathcal{O}_\varepsilon \ \forall\, k \geq \mathcal{K}$ w.p.1; otherwise if step 3c is visited and $\{x : H(x) \geq \bar{\gamma}_{k+1}, \ x \in \{X_1^k, \ldots, X_{N_k}^k\}\} = \emptyset$, then $supp\{\widetilde{g}_{k+1}\} = \emptyset$. We now discuss these two cases separately.

Case 1. If $supp\{\widetilde{g}_{k+1}\} \subseteq \mathcal{O}_\varepsilon$, then we have

$$\{\Gamma(supp\{\widetilde{g}_{k+1}\})\} \subseteq \{\Gamma(\mathcal{O}_\varepsilon)\}.$$

Since $E_{\widetilde{g}_{k+1}}[\Gamma(X)]$ is just the convex combination of $\Gamma(X_1^k), \ldots, \Gamma(X_{N_k}^k)$, it follows that

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

Thus by A4′, A5, and Lemma 5.2,

$$E_{\widetilde{\theta}_{k+1}}\left[\Gamma(X)\right] \in CONV\left\{\Gamma(\mathcal{O}_\varepsilon)\right\}.$$

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

$$E_{\widetilde{\theta}_{\widehat{k}}}\left[\Gamma(X)\right] \in CONV\left\{\Gamma(\mathcal{O}_\varepsilon)\right\}, \ w.p.1.$$

Furthermore, since $\widetilde{\theta}_{\widehat{k}} = \widetilde{\theta}_{\widehat{k}+1} = \cdots = \widetilde{\theta}_{k+1}$ (see discussions in Section 5.1), we will again have

$$E_{\widetilde{\theta}_{k+1}}\left[\Gamma(X)\right] \in CONV\left\{\Gamma(\mathcal{O}_\varepsilon)\right\}, \ \forall\, k \geq \mathcal{K} \ w.p.1.$$

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

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

where $\bar{\gamma}_k$ is defined as in MRAS$_1$. Note that since $\bar{\gamma}_k$ is a random variable, $\widehat{g}_{k+1}(x)$ is also a random variable. It follows that

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

Let $\omega = (X_1^0, \ldots, X_{N_0}^0, X_1^1, \ldots, X_{N_1}^1, \ldots)$ be a particular sample path generated by the algorithm. For each ω, the sequence $\{\bar{\gamma}_k(\omega), \ k = 1, 2, \ldots\}$ is non-decreasing and each strict increase is lower bounded by $\frac{\varepsilon}{2}$. Thus it is obvious that there exists $\widetilde{\mathcal{N}}(\omega) > 0$ such that $\bar{\gamma}_{k+1}(\omega) = \bar{\gamma}_k(\omega) \ \forall\, k \geq \widetilde{\mathcal{N}}(\omega)$. Now define $\Omega_1 := \{\omega : \lim_{k \to \infty} \bar{\gamma}_k(\omega) = H(x^*)\}$. By the definition of $\widetilde{g}_{k+1}(x)$ (cf. equation (22)), for each $\omega \in \Omega_1$ we clearly have $\lim_{k \to \infty} E_{\widetilde{g}_k(\omega)}\left[\Gamma(X)\right] = \Gamma(x^*)$; thus, it follows from Lemma 5.2 that $\lim_{k \to \infty} E_{\widetilde{\theta}_k(\omega)}\left[\Gamma(X)\right] = \Gamma(x^*)$, $\forall\, \omega \in \Omega_1$. The rest of the proof amounts to showing that the result also holds almost surely (a.s.) on the set Ω_1^c.

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

$$\liminf_{k \to \infty} \int_{\mathcal{X}} [\beta S(H(x))]^k I_{\{H(x) \geq \bar{\gamma}_k\}} dx \geq \int_{\mathcal{X}} \liminf_{k \to \infty} [\beta S(H(x))]^k I_{\{H(x) \geq \bar{\gamma}_k\}} dx > 0, \ \forall\, \omega \in \Omega_1^c, \qquad (24)$$

where the last inequality follows from the fact that $\beta S(H(x)) \geq 1 \ \forall\, x \in \left\{x : H(x) \geq \max\{S^{-1}(\frac{1}{\beta}), \bar{\gamma}_{\widetilde{\mathcal{N}}}\}\right\}$ and assumption A2.

Since $f(x, \theta_0) > 0$ for all $x \in \mathcal{X}$, clearly $\mathcal{X} \subseteq supp\{\widetilde{f}(\cdot, \widetilde{\theta}_k)\} \ \forall\, k$; thus

$$E_{\widehat{g}_{k+1}}\left[\Gamma(X)\right] = \frac{\widetilde{E}_{\widetilde{\theta}_k}\left[\beta^k \widetilde{S}_k(H(X)) I_{\{H(X) \geq \bar{\gamma}_k\}} \Gamma(X)\right]}{\widetilde{E}_{\widetilde{\theta}_k}\left[\beta^k \widetilde{S}_k(H(X)) I_{\{H(X) \geq \bar{\gamma}_k\}}\right]}, \ \forall\, k = 1, 2, \ldots,$$

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

$$\frac{\frac{1}{N_k} \sum_{i=1}^{N_k} \beta^k \widetilde{S}_k(H(X_i^k)) I_{\left\{H(X_i^k) \geq \bar{\gamma}_{k+1}\right\}} \Gamma(X_i^k)}{\frac{1}{N_k} \sum_{i=1}^{N_k} \beta^k \widetilde{S}_k(H(X_i^k)) I_{\left\{H(X_i^k) \geq \bar{\gamma}_{k+1}\right\}}} \longrightarrow E_{\widehat{g}_{k+1}}\left[\Gamma(X)\right] \ \text{a.s. on } \Omega_1^c,$$

20

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

For brevity, we use the following shorthand notations:

$$\widehat{Y}^k := \widetilde{E}_{\widetilde{\theta}_k}[\beta^k \widetilde{S}_k(H(X))I_{\{H(X) \geq \bar{\gamma}_k\}}], \quad \widehat{Y}_\Gamma^k := \widetilde{E}_{\widetilde{\theta}_k}[\beta^k \widetilde{S}_k(H(X))I_{\{H(X) \geq \bar{\gamma}_k\}}\Gamma(X)],$$

$$\bar{Y}_i^k := \beta^k \widetilde{S}_k(H(X_i^k))I_{\{H(X_i^k) \geq \bar{\gamma}_{k+1}\}}, \quad \widehat{Y}_i^k := \beta^k \widetilde{S}_k(H(X_i^k))I_{\{H(X_i^k) \geq \bar{\gamma}_k\}}.$$

We also let $\mathcal{T}_\varepsilon := \lceil \frac{2[H(x^*) - \mathcal{M}]}{\varepsilon} \rceil$. Note that the total number of visits to step $3a$ and $3b$ of MRAS$_1$ is bounded by $\mathcal{T}_\varepsilon$, thus for any $k > \mathcal{T}_\varepsilon$, the total number of visits to step $3c$ is greater than $k - \mathcal{T}_\varepsilon$.

We have

$$\frac{\frac{1}{N_k} \sum_{i=1}^{N_k} \beta^k \widetilde{S}_k(H(X_i^k))I_{\{H(X_i^k) \geq \bar{\gamma}_{k+1}\}}\Gamma(X_i^k)}{\frac{1}{N_k} \sum_{i=1}^{N_k} \beta^k \widetilde{S}_k(H(X_i^k))I_{\{H(X_i^k) \geq \bar{\gamma}_{k+1}\}}} - E_{\widehat{g}_{k+1}}[\Gamma(X)] =$$

$$\left(\frac{\frac{1}{N_k} \sum_{i=1}^{N_k} \bar{Y}_i^k \Gamma(X_i^k)}{\frac{1}{N_k} \sum_{i=1}^{N_k} \bar{Y}_i^k} - \frac{\frac{1}{N_k} \sum_{i=1}^{N_k} \widehat{Y}_i^k \Gamma(X_i^k)}{\frac{1}{N_k} \sum_{i=1}^{N_k} \widehat{Y}_i^k} \right) + \left(\frac{\frac{1}{N_k} \sum_{i=1}^{N_k} \widehat{Y}_i^k \Gamma(X_i^k)}{\frac{1}{N_k} \sum_{i=1}^{N_k} \widehat{Y}_i^k} - \frac{\widehat{Y}_\Gamma^k}{\widehat{Y}^k} \right).$$

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

$$\frac{\frac{1}{N_k} \sum_{i=1}^{N_k} \bar{Y}_i^k \Gamma(X_i^k)}{\frac{1}{N_k} \sum_{i=1}^{N_k} \bar{Y}_i^k} - \frac{\frac{1}{N_k} \sum_{i=1}^{N_k} \widehat{Y}_i^k \Gamma(X_i^k)}{\frac{1}{N_k} \sum_{i=1}^{N_k} \widehat{Y}_i^k} = 0, \quad \forall \; k \geq \widetilde{\mathcal{N}}(\omega), \; \forall \; \omega \in \Omega_1^c. \tag{25}$$

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 $\mathcal{C}_k = \{|\frac{1}{N_k} \sum_{i=1}^{N_k} \widehat{Y}_i^k - \widehat{Y}^k| > \zeta\}$. We have

$$P(\mathcal{C}_k \; i.o.) = P(\{\mathcal{C}_k \cap \mathcal{V}_k\} \cup \{\mathcal{C}_k \cap \mathcal{V}_k^c\} \; i.o.)$$
$$= P(\mathcal{C}_k \cap \mathcal{V}_k \; i.o.), \text{ since } P(\mathcal{V}_k^c \; i.o.) = 0 \text{ by part } \mathbf{(1)}. \tag{26}$$

It is easy to see that conditional on $\bar{\gamma}_k = \gamma$, $\widehat{Y}_1^k, \ldots, \widehat{Y}_{N_k}^k$ are i.i.d. and $E[\widehat{Y}_i^k | \bar{\gamma}_k = \gamma] = \widehat{Y}^k(\gamma) \; \forall \; i$. Furthermore, by assumption B1, conditional on the event $\mathcal{V}_k$, the support $[a_k, b_k]$ of the random variable $\widehat{Y}_i^k$ satisfies $[a_k, b_k] \subseteq \left[0, \frac{(\beta S^*)^k}{\lambda f_*}\right]$. Therefore, we have from the Hoeffding inequality (Hoeffding 1963),

$$P(\mathcal{C}_k \, | \mathcal{V}_k, \bar{\gamma}_k = \gamma) = P\left(|\frac{1}{N_k} \sum_{i=1}^{N_k} \widehat{Y}_i^k - \widehat{Y}^k| > \zeta \, \Big| \mathcal{V}_k, \bar{\gamma}_k = \gamma \right),$$
$$\leq 2\exp\left(\frac{-2N_k \zeta^2}{(b_k - a_k)^2} \right),$$
$$\leq 2\exp\left(\frac{-2N_k \zeta^2 [\lambda f_*]^2}{(\beta S^*)^{2k}} \right) \quad \forall \; k = 1, 2 \ldots. \tag{27}$$

Since

$$P(\mathcal{C}_k \cap \mathcal{V}_k) = \int_\gamma P(\mathcal{C}_k \cap \mathcal{V}_k \, | \bar{\gamma}_k = \gamma) f_{\bar{\gamma}_k}(d\gamma),$$
$$= \int_{\mathcal{V}_k} P(\mathcal{C}_k \, | \mathcal{V}_k, \bar{\gamma}_k = \gamma) f_{\bar{\gamma}_k}(d\gamma),$$

where $f_{\bar{\gamma}_k}(\cdot)$ is the distribution of the random variable $\bar{\gamma}_k$, we have by equation (27),

$$P(\mathcal{C}_k \cap \mathcal{V}_k) \le 2 \exp\left(\frac{-2N_k \zeta^2 [\lambda f_*]^2}{(\beta S^*)^{2k}}\right),$$

$$\le 2 \exp\left(\frac{-2(\alpha^{k-\mathcal{T}_\varepsilon} N_0)\zeta^2 [\lambda f_*]^2}{(\beta S^*)^{2k}}\right) \quad \forall\, k \ge \mathcal{T}_\varepsilon,$$

$$= 2 \exp\left(\frac{-2N_0 \zeta^2 \lambda^2 f_*^2}{\alpha^{\mathcal{T}_\varepsilon}}\left[\frac{\alpha}{(\beta S^*)^2}\right]^k\right) \quad \forall\, k \ge \mathcal{T}_\varepsilon,$$

Since $\frac{\alpha}{(\beta S^*)^2} > 1$ (by assumption), it follows that

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

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

$$P(\mathcal{C}_k \cap \mathcal{V}_k) < \frac{\alpha^{\mathcal{T}_\varepsilon}}{N_0 \zeta^2 \lambda^2 f_*^2}\left[\frac{(\beta S^*)^2}{\alpha}\right]^k \quad \forall\, k \ge \mathcal{T}_\varepsilon,$$

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

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

Finally by the Borel-Cantelli lemma and equation (26),

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

Since this holds for any $\zeta > 0$, we have $\frac{1}{N_k}\sum_{i=1}^{N_k} \widehat{Y}_i^k \to \widehat{Y}^k$ w.p.1.

By following the same argument as before, we can also show that $\frac{1}{N_k}\sum_{i=1}^{N_k} \widehat{Y}_i^k \Gamma(X_i^k) \to \widehat{Y}_\Gamma^k$ w.p.1. Note that since $\lim_{k\to\infty} \widehat{Y}^k > 0 \;\forall\, \omega \in \Omega_1^c$ (i.e., equation (24)), it follows that

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

By the definition of $\widetilde{g}_{k+1}(\cdot)$, the above result together with equation (25) suggests that

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

Thus, in conclusion, we have

$$E_{\widetilde{g}_k}[\Gamma(X)] \to E_{\widehat{g}_k}[\Gamma(X)] \text{ as } k \to \infty \text{ w.p.1.}$$

Also, by assumptions A2, A3, A5, A6, B1, and following the proof of Theorem 3.1, it is not difficult to show that

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

Hence by Lemma 5.2, we have

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

$\blacksquare$

The following results are now immediate.

Corollary 5.1 (Multivariate Normal) *If multivariate normal p.d.f.'s are used in MRAS$_1$, i.e.,*

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

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

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

Corollary 5.2 (Independent Univariate) *If the components of the random vector $X = (X_1, X_2, \ldots, X_n)$ are independent, each with a univariate p.d.f. of the form*

$$f(x_i, \vartheta_i) = \exp(x_i \vartheta_i - K(\vartheta_i)) h(x_i), \ \vartheta_i \in \Re, \forall \ i = 1, \ldots, n,$$

$\varepsilon > 0$, $\alpha > (\beta S^*)^2$, *and assumptions A2, A3, A4', A5, and B1 are satisfied, then*

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

Theorem 5.2 (Discrete Optimization with Finite Domains) *If the solution space $\mathcal{X}$ is finite, $\alpha > 1$, and assumptions A4', and A5 are satisfied, then*

$$\lim_{k \to \infty} E_{\widetilde{\theta}_k}[\Gamma(X)] = \Gamma(x^*) \quad w.p.1.$$

Proof: Replace the Lebesgue measure with the discrete measure in Theorem 5.1, take $\beta = \frac{1}{S^*}$, and use the same argument as before. ∎

6 Numerical Examples

In this section, we illustrate the performance of the MRAS method for both continuous and combinatorial optimization problems. In the former case, we test the algorithm on various functions that are well-known in global optimization and compare its performance with that of the standard CE method. 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.

Remark 7: It is not our primary intention here to compare our algorithm with the CE method and EDAs. A comprehensive comparison of different methods is beyond the scope of this paper. Our main goal here is to propose a novel algorithm with provable convergence, and show that the algorithm is promising in solving some difficult optimization problems. The performance of the CE method on continuous functions can be found in, e.g., Kroese et. al (2004) and Rubinstein (1999). Its performance on various ATSP instances can be found in, e.g., de Boer et. al (2004) and Rubinstein (2001).

We now discuss some implementation issues of the MRAS$_1$ algorithm.

1. Since all examples considered in this section are minimization problems, whereas MRAS was presented in a maximization context, the following modifications are required:

   - $S(\cdot)$ needs to be initialized as a strictly decreasing function instead of strictly increasing. Throughout this section, we take

   $$S(H(x)) := \exp\{-rH(x)\}, \text{ where } r \text{ is a positive constant.}$$

23

- The sample $(1-\rho)$-quantile $\widetilde{\gamma}_{k+1}$ will now be calculated by first ordering the sample performances $H(X_i^k)$, $i = 1, \ldots, N_k$ from largest to smallest, and then taking the $\lceil(1-\rho)N_k\rceil$th order statistic.

- We need to replace the "$\geq$" operator with "$\leq$" operator in equation (21).

- The inequalities at step 3 need to be replaced with

$$\widetilde{\gamma}_{k+1}(\rho_k, N_k) \leq \bar{\gamma}_k - \frac{\varepsilon}{2}, \text{ and } \widetilde{\gamma}_{k+1}(\bar{\rho}, N_k) \leq \bar{\gamma}_k - \frac{\varepsilon}{2},$$

respectively.

2. The sequence $\{f(x, \widetilde{\theta}_k)\}$ may converge too quickly to a degenerate p.d.f., which would cause the algorithm to get trapped in local optimal solutions. To prevent this from happening, a smoothed parameter updating procedure (cf. e.g., De Boer et al. 2004, Rubinstein 1999) is used in actual implementation, i.e., first a smoothed parameter vector $\widehat{\theta}_{k+1}$ is computed at each iteration k according to

$$\widehat{\theta}_{k+1} := \upsilon\widetilde{\theta}_{k+1} + (1-\upsilon)\widehat{\theta}_k, \quad \forall\, k = 0, 1, \ldots, \text{ and } \widehat{\theta}_0 := \widetilde{\theta}_0,$$

where $\widetilde{\theta}_{k+1}$ is the parameter vector computed at step 4 of MRAS$_1$, and $\upsilon \in (0, 1]$ is the smoothing parameter; then $f(x, \widehat{\theta}_{k+1})$ (instead of $f(x, \widetilde{\theta}_{k+1})$) is used in step 1 to generate new samples. It is important to note that this modification will not affect the theoretical convergence of our approach.

3. In practice, different stopping criteria can be used. The simplest method is to stop the algorithm when a predefined maximum number of iterations is reached, or when the total computational budget is exhausted. In the numerical experiments, a mixed stopping rule is used: We stop the algorithm either when no significant improvement in $\bar{\gamma}_k$ is obtained for several consecutive iterations or when the sample size at a single iteration exceeds some predefined threshold, i.e., as soon as either one of the following two conditions is satisfied at iteration k:

   (1) $\max_{1 \leq i \leq d} |\bar{\gamma}_k - \bar{\gamma}_{k+i}| \leq \tau$;

   (2) $N_k > N_{\max}$;

   where $\tau > 0$ is a predefined tolerance level, d is a positive integer, and $N_{\max}$ is the maximum number of samples allowed per iteration.

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 $\widetilde{\mu}_{k+1}$ and $\widetilde{\Sigma}_{k+1}$ are updated according to the respective stochastic counterparts of equations (16) and (17). By Corollary 5.1, the sequence of mean vectors $\{\widetilde{\mu}_k\}$ will converge to the optimal solution x^*, and the sequence of covariance matrices $\{\widetilde{\Sigma}_k\}$ to the zero matrix. Throughout this section, we will use $\widetilde{\mu}_k$ to represent the current best solution found at iteration k.

The following five functions $\{H_i,\ i = 1, \ldots, 5\}$ are used to test the algorithm.

(1) Quadratic function

$$H_1(x) = \sum_{i=1}^{3} x_i^2, \text{ where } x = (x_1, x_2, x_3).$$

The function has a unique global minimum $f(0, 0, 0) = 0$.

(2) Two-dimensional Rosenbrock function

$$H_2(x) = 100(x_1^2 - x_2)^2 + (1 - x_1^2), \text{ where } x = (x_1, x_2).$$

24

The function has the reputation of being difficult to minimize and is widely used to test the performance of different optimization algorithms. It has a global minimum $f(1,1) = 0$.

(3) Shekel's Foxholes

$$H_3(x) = \frac{1}{0.002 + \sum_{j=1}^{25} \frac{1}{j + \sum_{i=1}^{2}(x_i - a_{j,i})^6}},$$

where $a_{j,1} = \{-32, -16, 0, 16, 32, -32, -16, 0, 16, 32, -32, -16, 0, 16, 32, -32, -16,$
$0, 16, 32, -32, -16, 0, 16, 32\}$,
$a_{j,2} = \{-32, -32, -32, -32, -32, -16, -16, -16, -16, -16, 0, 0, 0, 0, 0, 16, 16, 16, 16,$
$16, 32, 32, 32, 32, 32\}$, and $x = (x_1, x_2)$. The function has 24 local minima and one global minimum at $f(-32, -32) \approx 0.998004$.

(4) Corana's Parabola

$$H_4(x) = \sum_{i=1}^{4} \begin{cases} 0.15[0.05\,\mathrm{sgn}(z_i) - z_i]^2 d_j & \text{if } |x_i - z_i| < 0.05, \\ d_i x_i^2 & \text{otherwise,} \end{cases}$$

where

$$z_i = 0.2 \left\lfloor \left| \frac{x_i}{2} \right| + 0.49999 \right\rfloor \mathrm{sgn}(x_i),$$

$d = \{1, 1000, 10, 100\}$, and $x = (x_1, x_2, x_3, x_4)$. In the region $-1000 < x_i < 1000$, $i = 1, 2, 3, 4$, the above function has more than 10^{20} local minima, which is very difficult to minimize. It has a global minimum $f(0, 0, 0, 0) = 0$.

(5) Goldstein-Price function

$$H_5(x) = (1 + (x_1 + x_2 + 1)^2(19 - 14x_1 + 3x_1^2 - 14x_2 + 6x_1 x_2 + 3x_2^2))$$
$$(30 + (2x_1 - 3x_2)^2(18 - 32x_1 + 12x_1^2 + 48x_2 - 36x_1 x_2 + 27x_2^2)),$$

where $x = (x_1, x_2)^T$. The function has four local minima and a global minimum $f(0, -1) = 3$.

For all five problems, the same set of parameters is used to test MRAS: $\varepsilon = 10^{-5}$, initial sample size $N_0 = 100$, $\rho_0 = 0.2$, $\lambda = 0.02$, $\alpha = 1.5$, $r = 0.1$, the stopping control parameters $d = 5$, $\tau = 10^{-5}$, $N_{max} = 50000$, and the smoothing parameter $\upsilon = 0.5$. The initial mean vector μ_0 is a d-by-1 vector of all 10s, and Σ_0 is a d-by-d diagonal matrix with all diagonal elements equal to 200, where d is the dimension of the problem.

H_i	N_{total} (std)	ρ_{final} (std)	$\widetilde{H}_i^*$ (std)	$H_i(x^*)$	M_ε
H_1	4.38e+03(6.77e+01)	0.13(6.21e-03)	9.86e-09(1.12e-09)	0	50
H_2	1.21e+04(4.89e+02)	0.04(2.40e-03)	2.29e-09(3.13e-10)	0	50
H_3	2.17e+04(7.16e+02)	0.02(1.11e-03)	2.40(4.15e-01)	0.998	37
H_4	7.43e+03(1.61e+02)	0.14(4.19e-03)	0.00(0.00e-00)	0	50
H_5	5.81e+03(1.40e+02)	0.11(5.88e-03)	3.00(5.30e-10)	3	50

Table 1: Performance of MRAS on five test functions, based on 50 independent simulation runs. The standard errors are in parentheses.

Table 1 shows the performance of the algorithm on the five test functions. For each function, we performed 50 independent simulation runs of the algorithm, and the means and standard errors are reported in the table, where N_{total} is the total number of function evaluations, ρ_{final} is the final value of ρ, and $\widetilde{H}_i^*$ is the value of the function $H_i(\cdot)$ at the final value of the estimated optimum. The optimal value $H_i(x^*)$ is included for reference, and M_ε indicates the number of runs that an ε-optimal solution was found out of 50 trials. The algorithm performs quite well in most cases, except for H_3, where only 37 ε-optimal solutions were

found. H_3 represents a class of continuous optimization problems that are extremely difficult to solve for most *model-based* sampling approaches. A graphical representation of the function H_3 is given in Figure 1. Notice that the function values at the 25 "holes" (local minima) are very close to each other; thus in order to locate the global optimal solution, the algorithm must make sure that samples are drawn from the right "hole", and there must be enough samples to fall in this "hole" to guarantee that the parameter vectors are updated in the right direction.
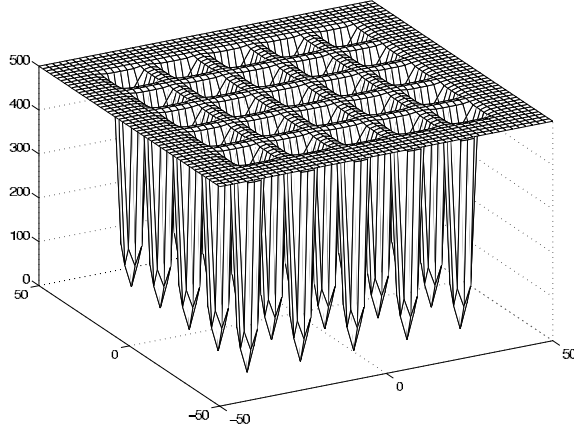

Figure 1: Shekel's Foxholes, where $-50 \leq x_i \leq 50$, $i = 1, 2$.

For comparison purposes, we also applied the CE method to the above five test functions, where we have used the multivariate normal p.d.f. with independent components (cf. e.g., Kroese et al. 2004 for detailed algorithm description and implementation issues). We have tested different sets of parameters (i.e., different (N, ρ) combinations); the results reported in Table 2 are based on the following "good" parameter settings: sample size $N = 1000$ (recall that the CE method is non-adaptive, so the same number of samples will be generated at each iteration), $\rho = 0.005$, smoothing parameter $\upsilon = 0.7$, and the algorithm is stopped either when there exists $k > 0$ such that $\max_{1 \leq i \leq 5} |\hat{\gamma}_k - \hat{\gamma}_{k+i}| \leq 10^{-5}$ or when the total number of samples generated exceeds 2×10^5, where $\hat{\gamma}_k$ is the sample $(1 - \rho)$-quantile generated at the kth iteration of CE.

Again, the mean vector μ_0 is initialized as a d-by-1 vector of all 10s and the variances are taken to be a d-by-1 vector with all elements equal to 200.

H_i	N_{total} (std)	$\widetilde{H}_i^*$ (std)	$H_i(x^*)$	M_ε
H_1	1.69e+04(1.73e+02)	4.94e-05(5.13e-06)	0	7
H_2	1.72e+04(1.18e+02)	1.92e-05(2.93e-06)	0	24
H_3	1.05e+04(1.08e+02)	8.83(2.54e-01)	0.998	0
H_4	5.84e+04(6.06e+03)	1.35e-03(4.21e-04)	0	38
H_5	1.89e+05(4.77e+03)	3.00(5.63e-05)	3	0

Table 2: Performance of the standard CE method on five test functions, based on 50 independent simulation runs. The standard errors are in parentheses.

From Tables 1 and 2, we see that MRAS uses fewer samples than CE does, but produces more accurate solutions. In general, the sequence $\{\hat{\gamma}_k\}$ generated by CE may often converge quickly to a small neighborhood of $H(x^*)$; however, since no sample performances are used in parameter updating, (i.e., the top $\rho\%$ samples are all considered to be of the same importance regardless of their sample performances), the future search will be biased toward the region that has been sampled most. In particular, for the H_3 case, since the function values at different local minima are very close to each other, even if the "hole" with the global minimum has been sampled during the search process, CE still cannot distinguish the global minimum from

the other local minima; instead CE will easily get stuck in the "hole" that has been sampled the most. As a result, we see that the algorithm gets trapped in local minima in all 50 trials. In contrast, the parameter updating procedure in MRAS is weighted by the performance function so that better samples will have more positive influence on the updating process. Consequently, the searches in MRAS will be biased toward the region containing more promising samples.

Table 3 gives the performance of CE and MRAS on function H_3 using different sample sizes and ρ values (all other parameters are the same as before). Test results indicate that increasing the samples size in CE has little effect on the quality of the resultant solutions. We see that the algorithm consistently gets stuck in local minima in repeated experiments. On the other hand, for MRAS with $N_0 = 200$, ε-optimal solutions were found in more than 90% of the total simulation runs; whereas for the $N_0 \geq 500$ cases, ε-optimal solutions were found in all 50 runs.

method	parameters	N_{total} (std)	$\widetilde{H}_3^*$ (std)	M_ε
CE	$N=1000, \rho=0.1$	1.47e+04(1.39e+02)	18.29(0.18)	0
	$N=1000, \rho=0.01$	1.13e+04(1.08e+02)	11.90(0.27)	0
	$N=2000, \rho=0.1$	2.91e+04(2.07e+02)	18.30(0.09)	0
	$N=2000, \rho=0.01$	2.25e+04(1.70e+02)	12.27(0.19)	0
	$N=2000, \rho=0.005$	2.14e+04(2.25e+02)	8.43(0.21)	0
	$N=5000, \rho=0.1$	7.19e+04(3.47e+02)	18.30(8.07e-11)	0
	$N=5000, \rho=0.01$	5.70e+04(3.78e+02)	12.52(0.14)	0
	$N=5000, \rho=0.001$	4.87e+04(6.67e+02)	5.61(0.32)	0
	$N=10000, \rho=0.1$	1.42e+05(6.10e+02)	18.30(4.80e-11)	0
	$N=10000, \rho=0.01$	1.12e+05(6.19e+02)	12.67(1.96e-12)	0
	$N=10000, \rho=0.001$	1.01e+05(1.39e+03)	4.80(0.32)	0
MRAS	$N_0=200, \rho_0=0.2$	2.27e+04(6.77e+02)	1.14(0.06)	45
	$N_0=200, \rho_0=0.1$	2.17e+04(7.14e+02)	1.08(0.05)	47
	$N_0=500, \rho_0=0.2$	3.01e+04(6.67e+02)	0.998(3.41e-11)	50
	$N_0=500, \rho_0=0.1$	2.76e+04(8.70e+02)	0.998(3.92e-11)	50
	$N_0=1000, \rho_0=0.2$	5.62e+04(8.30e+02)	0.998(3.41e-11)	50
	$N_0=1000, \rho_0=0.1$	4.31e+04(8.46e+02)	0.998(3.81e-11)	50

Table 3: Performance of CE and MRAS on test function H_3, based on 50 independent simulation runs. The standard errors are in parentheses. The optimum $H_3(x^*) \approx 0.998004$.

We also applied MRAS to a 10-dimensional trigonometric function and a 10-dimensional Rosenbrock function. These functions are used to test the performance of CE in e.g., Rubinstein (1999), Kroese et al. (2004).

(6) Trigonometric function

$$H_6(x) = \sum_{i=1}^{10} 8\sin^2\left(7(x_i - 0.9)^2\right) + 6\sin^2\left(14(x_i - 0.9)^2\right) + (x_i - 0.9)^2, \text{ where } x = (x_1,\ldots,x_{10})^T.$$

The function has a global minimum at $x^* = (0.9,\ldots,0.9)^T$, and $H_6(x^*) = 0$.

(7) Rosenbrock function

$$H_7(x) = \sum_{i=1}^{9} 100(x_{i+1} - x_i^2)^2 + (x_i - 1)^2, \text{ where } x = (x_1,\ldots,x_{10})^T.$$

The function has a global minimum at $x^* = (1,\ldots,1)^T$, and $H_7(x^*) = 0$.

In our preliminary experiments with function H_6, we use the same set of parameters as before. The

numerical results are recorded in Table 4, where we have used four different sets of (N_0, ρ_0) combinations. Test results indicate that the algorithm may frequently get trapped in local optimal solutions. We believe this is due to the fact that the covariance matrix often converges too quickly to a zero matrix.

parameters	N_{total} (std)	$\widetilde{H}_6^*$ (std)	M_ε
$N_0{=}200$, $\rho_0{=}0.1$	5.62e+04(4.35e+03)	1.05(0.17)	20
$N_0{=}200$, $\rho_0{=}0.2$	4.67e+04(2.41e+03)	1.02(0.16)	16
$N_0{=}500$, $\rho_0{=}0.1$	7.48e+04(3.86e+03)	0.29(0.06)	28
$N_0{=}500$, $\rho_0{=}0.2$	6.22e+04(3.00e+03)	0.64(0.11)	23

Table 4: Performance of MRAS ($r = 0.1$, $v = 0.5$) on test function H_6, based on 50 independent simulation runs. The standard errors are in parentheses.

In order to reduce the convergence speed of the covariance matrix, we reran the algorithm with smaller values of r and v. Table 5 shows the performance of the algorithm with $r = 0.01$ and $v = 0.2$. We see that the algorithm successfully escaped the local optima in all simulation runs. However, decreasing the values of r and v also slows down the convergence of the algorithm; note that the total number of function evaluations in all cases almost increased by a factor of 10 as compared to Table 4. For a given problem, how to choose the most appropriate parameters in MRAS is still an open issue. As a general guideline, the higher the dimension of the problem, the more easily the algorithm gets trapped in local optima. Thus, if high quality solutions are considered more desirable, small values of r and v are often preferred.

parameters	N_{total} (std)	$\widetilde{H}_6^*$ (std)	M_ε
$N_0{=}200$, $\rho_0{=}0.1$	5.82e+05(4.61e+04)	3.27e-07(2.64e-08)	50
$N_0{=}200$, $\rho_0{=}0.2$	4.24e+05(4.60e+04)	3.62e-07(2.64e-08)	50
$N_0{=}500$, $\rho_0{=}0.1$	5.97e+05(4.80e+04)	2.97e-07(2.30e-08)	50
$N_0{=}500$, $\rho_0{=}0.2$	5.42e+05(3.08e+04)	2.62e-07(1.90e-08)	50

Table 5: Performance of MRAS ($r = 0.01$, $v = 0.2$) on test function H_6, based on 50 independent simulation runs. The standard errors are in parentheses.

The same set parameters are also used to solve the function H_7, and the numerical results are reported in Table 6. In Figure 2, we also plot some typical sample path of MRAS$_1$ on H_6 and H_7, which clearly shows the fast (approximately exponential) convergence rate of the proposed method (note the log scale used in the plot).

parameters	N_{total} (std)	$\widetilde{H}_7^*$ (std)	M_ε
$N_0{=}200$, $\rho_0{=}0.1$	2.69e+05(1.31e+04)	2.39e-08(2.85e-09)	50
$N_0{=}200$, $\rho_0{=}0.2$	2.62e+05(1.51e+04)	2.25e-08(2.16e-09)	50
$N_0{=}500$, $\rho_0{=}0.1$	3.34e+05(1.36e+04)	1.79e-08(1.82e-09)	50
$N_0{=}500$, $\rho_0{=}0.2$	3.61e+05(1.57e+04)	2.92e-08(3.43e-09)	50

Table 6: Performance of MRAS ($r = 0.01$, $v = 0.2$) on test function H_7, based on 50 independent simulation runs. The standard errors are in parentheses.

6.2 Combinatorial Optimization

In this section, we present the performance of MRAS on various ATSP problems. All test cases are taken from the URL
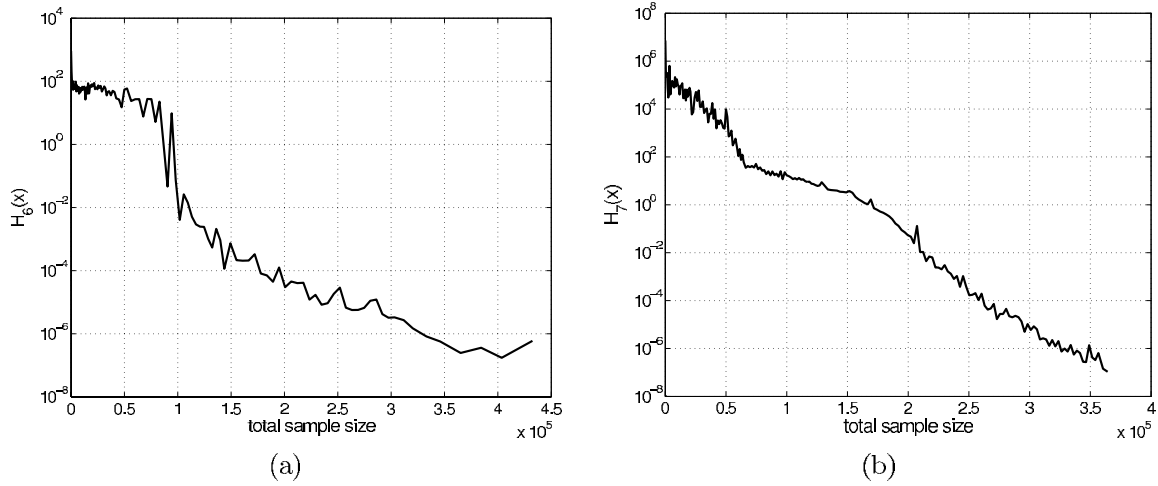
http://www.iwr.uni-heidelberg.de/groups/comopt/software/TSPLIB95.

Figure 2: Typical performance of MRAS on H_6 and H_7 ($N_0 = 500$, $\rho_0 = 0.1$, $r = 0.01$, $\upsilon = 0.2$), (a) 10-D trigonometric function; (b) 10-D Rosenbrock function.

For each ATSP problem with N_c number of cities, a 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:

$$\min_{x \in \mathcal{X}} H(x) := \min_{x \in \mathcal{X}} \left\{ \sum_{i=1}^{N_c-1} G_{x_i, x_{i+1}} + G_{x_{N_c}, x_1} \right\}, \tag{28}$$

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

We use the same technique as in Rubinstein (2001) and De Boer et al. (2004) for solving these problems, i.e., we associate for each distance matrix G an initial state transition matrix $\widetilde{P}_0$, 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:

- Generating random (admissible) tours according to the transition matrix and evaluate the performance of each sample tour.

- Updating 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 in e.g., De Boer et al. (2004). We now briefly address the issue of how to update the transition matrix. At each iteration k of MRAS, the p.d.f. $f(\cdot, \widetilde{P}_k)$ on $\mathcal{X}$ is parameterized by the transition matrix $\widetilde{P}_k$ and is given by

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

where $\mathcal{X}_{i,j}(l)$ is the set of all tours in $\mathcal{X}$ such that the lth transition is from city i to city j. It is straightforward to show that the new transition matrix $\widetilde{P}_{k+1}$ is updated in equation (21) as

$$\widetilde{P}_{k+1}(i,j) = \frac{\sum_{l=1}^{N_k} \widetilde{S}_k(H(X_l^k)) I_{\left\{ H(X_l^k) \leq \bar{\gamma}_{k+1} \right\}} I_{\left\{ X_l^k \in \mathcal{X}_{i,j} \right\}}}{\sum_{l=1}^{N_k} \widetilde{S}_k(H(X_l^k)) I_{\left\{ H(X_l^k) \leq \bar{\gamma}_{k+1} \right\}}}, \tag{29}$$

where $X_1^k, \ldots, X_{N_k}^k$ are the i.i.d. sample tours generated from $\widetilde{f}(\cdot, \widetilde{P}_k)$, $\bar{\gamma}_{k+1}$ is defined as in equation (22), and $\mathcal{X}_{i,j}$ represents the set of tours in which the transition from city i to city j is made.

The performance of the algorithm on various ATSP problems is reported in Table 7. For each of the 7 instances, we performed 10 independent runs of the algorithm. In Table 7, N_{total} is the total number of tours generated (mean and standard error reported), H_{best} is the length of the shortest path, H_* and H^* are the worst and best solutions obtained out of 10 trials, δ_* and δ^* are the respective relative errors for H_* and H^*, and δ is the relative error (mean and standard error reported). For all cases, $\varepsilon = 1$, the initial samples $N_0 = 1000$, $\rho_0 = 0.1$, $\lambda = 0.02$, $\alpha = 1.5$, $r = 0.1$, the stopping control parameters $d = 5$, $\tau = 0$, $N_{max} = 10N_c^2$, smoothing parameter $\upsilon = 0.5$, and the initial transition matrix $\widetilde{P}_0$ 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., $\widetilde{P}_0(i,j) \propto \frac{1}{G_{i,j}}$ and $\sum_j \widetilde{P}_0(i,j) = 1\ \forall i$.

ATSP	N_c	N_{total} (std err)	H_{best}	H_*	H^*	δ_*	δ^*	δ (std err)
ftv33	34	7.95e+04(3.25e+03)	1286	1364	1286	0.061	0.000	0.023(0.008)
ftv35	36	1.02e+05(3.08e+03)	1473	1500	1475	0.018	0.001	0.008(0.002)
ftv38	39	1.31e+05(4.90e+03)	1530	1563	1530	0.022	0.000	0.008(0.003)
p43	43	1.02e+05(4.67e+03)	5620	5637	5620	0.003	0.000	0.001(2.5e-4)
ry48p	48	2.62e+05(1.59e+04)	14422	14810	14446	0.027	0.002	0.012(0.003)
ft53	53	2.94e+05(1.58e+04)	6905	7236	6973	0.048	0.010	0.029(0.005)
ft70	70	4.73e+05(2.91e+04)	38673	39751	38744	0.028	0.002	0.017(0.003)

Table 7: Performance of MRAS on various ATSP problems based on 10 independent replications. The standard errors are in parentheses.

7 Conclusions and Future Work

We have proposed a randomized optimization technique called Model Reference Adaptive Search (MRAS) for solving both continuous and discrete optimization problems. Highlights of the method include the following:

(1) It is generic, requiring only a few mild regularity conditions on the underlying problem.

(2) It is convergent w.p.1 to the set of ε-optimal solutions in a finite number of iterations and asymptotically to a global optimal solution.

(3) It is insensitive to the choices of the initial solutions (parameter vectors), provided that the initial sampling variance is large enough.

(4) It offers an alternative *general framework* for global optimization, based on which one can design and implement other efficient algorithms.

The MRAS$_1$ algorithm demonstrated great 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 update 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 $\widetilde{H}(x)$ 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

$$x^* \in \underset{x \in \mathcal{X}}{\arg\max}\, E\left[\widetilde{H}(x)\right], \quad x \in \mathcal{X} \subseteq \Re^n, \tag{30}$$

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

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

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

$$S_k(\widetilde{H}(x)) := \prod_{i=1}^{k} S(\widetilde{H}_i(x)). \tag{31}$$

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 (31) can be written as

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

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 problem (30). Construction of a practically efficient generalization of MRAS with provable convergence is addressed in Hu et al. (2005).

References

P. Auer, N. Cesa-Bianchi, Y. Freund, and R. E. Schapire 2002. "The Nonstochastic Multiarmed Bandit Problem." *SIAM J. Comput. Vol. 32, No. 1, pp. 48 − 77.*

H. S. Chang, M. C. Fu, and S. I. Marcus 2003. "An Asymptotically Efficient Simulation-Based Algorithm for Finite Horizon Stochastic Dynamic Programming." *Proceedings of the 42nd IEEE Conference on Decision and Control, 3818 − 3823*

P. T. De Boer, D. P. Kroese, S. Mannor, and R. Y. Rubinstein 2004. "A Tutorial on the Cross-Entropy Method." *Annals of Operation Research. To appear.*

M. Dorigo and L. M. Gambardella 1997. "Ant Colony System: A Cooperative Learning Approach to the Traveling Salesman Problem." *IEEE Trans. on Evolutionary Computation Vol. 1, No. 1.*

F. Glover 1990. "Tabu Search: A Tutorial." *Interfaces. Vol. 20, No. 4. pp. 74 − 94.*

W. Hoeffding 1963. "Probability Inequalities for Sums of Bounded Random Variables." *Journal of the American Statistical Association, Vol. 58, No. 301.*

T. Homem-de-Mello and R. Y. Rubinstein 2003. "Rare Event Estimation for Static Models via Cross-Entropy and Importance Sampling." http://iew3.technion.ac.il/CE/pubs.php

J. Hu, M. C. Fu, and S. I. Marcus 2005. "A Model Reference Adaptive Search Algorithm for Stochastic Global Optimization." *Working paper.*

S. Kirkpatrick, C. D. Gelatt, and M. P. Vecchi 1983. "Optimization by Simulated Annealing." *Science, Vol. 220, pp. 671 − 680.*

D. P. Kroese, R. Y. Rubinstein and S. Porotsky 2004. "The Cross-Entropy Method for Continuous Multi-extremal Optimization." *Operations Research. Under review.*

P. Larrañaga, R. Etxeberria, J. A. Lozano, B. Sierra, I. Iñza, and J. M. Peña 1999. "A Review of the Cooperation Between Evolutionary Computation and Probabilistic Graphical Models." *Proceedings of the Second Symposium on Artificial Intelligence. Adaptive Systems. CIMAF 99. Special Session on Distributions and Evolutionary Computation,* $314 - 324$.

H. Mühlenbein and G. Paaß. 1996. "From Recombination of Genes to the Estimation of Distributions: *I.* Binary Parameters." *In Hans-Michael Voigt, Werner Ebeling, Ingo Rechenberg, and Hans-Paul Schwefel, editors, Parallel Problem Solving from Nature - PPSN IV, pp.* $178 - 187$, *Berlin, Springer Verlag.*

R. Y. Rubinstein 1997. "Optimization of Computer Simulation Models with Rare Events." *European Journal of Operations Research,* $99 : 89 - 112$.

R. Y. Rubinstein 1999. "The Cross-Entropy Method for Combinatorial and Continuous Optimization." *Methodology and Computing in Applied Probability,* $2 : 127 - 190$.

R. Y. Rubinstein 2001. "Combinatorial Optimization, Ants and Rare Events." *In S. Uryasev and P. M. Pardalos, editors, Stochastic Optimization: Algorithms and Applications, pp.* $304 - 358$ *Kluwer.*

R. Y. Rubinstein and D. P. Kroese 2004. "The Cross-Entropy Method: a unified approach to combinatorial optimization, Monte-Carlo simulation, and machine learning." *New York, Springer Science+Business Media, Inc.*

R. Y. Rubinstein and A. Shapiro 1993. "Discrete Event Systems: sensitivity analysis and stochastic optimization by the score function method." *John Wiley & Sons.*

L. Shi and Ólafsson 2000. "Nested Partitions Method for Global Optimization." *Operations Research,* $(48)3$.

M. Srinivas and L. M. Patnaik 1994. "Genetic Algorithms: a survey." *IEEE Computer, Vol. 27, No. 6, pp.* $17 - 26$.

M. Zlochin, M. Birattari, N. Meuleau, and M. Dorigo 2004. "Model-based Search for Combinatorial Optimization: A critical survey" *Annals of Operations Research, Vol.* $131, pp.$ $373 - 395$.

Q. Zhang, H. Mühlenbein 2004. "On the Convergence of a Class of Estimation of Distribution Algorithm." *IEEE Trans. on Evolutionary Computation Vol. 8, No. 2, pp.* $127 - 136$.

A   Appendix

Proof of Lemma 5.1:   Our proof is based on the proof of Lemma $A1$ in Rubinstein and Shapiro (1993). Notice that for given $\rho^\dagger$ and $\widetilde{f}(\cdot, \widetilde{\theta}_k)$, $\gamma_k^\dagger$ can be obtained as the optimal solution of the following problem (cf. Homem-de-Mello and Rubinstein 2003)

$$\min_{v \in \mathcal{V}} \ell_k(v), \tag{32}$$

where $\mathcal{V} = [0, H(x^*)]$, $\ell_k(v) := \widetilde{E}_{\widetilde{\theta}_k} \phi(H(X), v)$, and

$$\phi(H(x), v) := \begin{cases} (1 - \rho^\dagger)(H(x) - v) & \text{if } v \leq H(x), \\ \rho^\dagger(v - H(x)) & \text{if } v \geq H(x). \end{cases}$$

Similarily, the sample quantile $\widetilde{\gamma}_k^\dagger(\rho^\dagger, N_k)$ can be expressed as the solution to the sample average approximation of (32),

$$\min_{v \in \mathcal{V}} \bar{\ell}_k(v), \tag{33}$$

where $\bar{\ell}_k(v) := \frac{1}{N_k} \sum_{j=1}^{N_k} \phi(H(X_j^k), v)$ and $X_1^k, \ldots, X_{N_k}^k$ are i.i.d. with density $\widetilde{f}(\cdot, \widetilde{\theta}_k)$.

Since the function $\phi(H(x), v)$ is bounded and continuous on $\mathcal{V}$ for all $x \in \mathcal{X}$, it is not difficult to show that $\ell_k(v)$ is continuous on $\mathcal{V}$ (cf. Rubinstein and Shapiro 1993).

Now consider a point $v \in \mathcal{V}$ and let $B_i \subseteq \mathcal{V}$ be a sequence of open balls containing v such that $B_{i+1} \subseteq B_i \, \forall i$ and $\lim_{L \to \infty} \cap_{i=1}^L B_i = v$. Define the function

$$b_i(H(x)) := \sup \{|\phi(H(x), u) - \phi(H(x), v)| : \, u \in B_i\}.$$

We have from the dominated convergence theorem

$$\lim_{i \to \infty} \widetilde{E}_{\widetilde{\theta}_k}[b_i(H(X))] = \widetilde{E}_{\widetilde{\theta}_k}[\lim_{i \to \infty} b_i(H(X))] = 0 \quad \forall \, k = 1, 2, \ldots, \tag{34}$$

where the last equality follows from the fact that $\phi(H(x), v)$ is continuous on $\mathcal{V}$.

Since

$$|\bar{\ell}_k(u) - \bar{\ell}_k(v)| \leq \frac{1}{N_k} \sum_{j=1}^{N_k} |\phi(H(X_j^k), u) - \phi(H(X_j^k), v)|,$$

it follows that

$$\sup_{u \in B_i} |\bar{\ell}_k(u) - \bar{\ell}_k(v)| \leq \frac{1}{N_k} \sum_{j=1}^{N_k} b_i(H(X_j^k)). \tag{35}$$

We now show that $\frac{1}{N_k} \sum_{j=1}^{N_k} b_i(H(X_j^k)) \to \widetilde{E}_{\widetilde{\theta}_k}[b_i(H(X))]$ as $k \to \infty$ w.p.1.

Let M be an upperbound for $b_i(H(x))$, and let $\mathcal{T}_\varepsilon := \lceil \frac{2[H(x^*) - \mathcal{M}]}{\varepsilon} \rceil$, where $\mathcal{M}$ is a lower bound for the function $H(x)$, and ε is defined as in the MRAS$_1$ algorithm. Note that the total number of visits to step 3a and 3b of MRAS$_1$ is bounded by $\mathcal{T}_\varepsilon$, thus for any $k > \mathcal{T}_\varepsilon$, the total number of visits to step 3c is greater than $k - \mathcal{T}_\varepsilon$. Since conditional on $\widetilde{\theta}_k$, $\frac{1}{N_k} \sum_{j=1}^{N_k} b_i(H(X_j^k))$ is an unbiased estimate of $\widetilde{E}_{\widetilde{\theta}_k}[b_i(H(X))]$, by the Hoeffding inequality (cf. equation (27)), for any $\zeta > 0$,

$$P\left(|\frac{1}{N_k} \sum_{j=1}^{N_k} b_i(H(X_j^k)) - \widetilde{E}_{\widetilde{\theta}_k}[b_i(H(X))]| > \zeta \, \Big| \widetilde{\theta}_k = \theta\right) \leq 2\exp\left(\frac{-2N_k\zeta^2}{M^2}\right) \quad \forall \, k.$$

Therefore,

$$\begin{aligned} P\left(|\frac{1}{N_k} \sum_{j=1}^{N_k} b_i(H(X_j^k)) - \widetilde{E}_{\widetilde{\theta}_k}[b_i(H(X))]| > \zeta\right) &\leq 2\exp\left(\frac{-2N_k\zeta^2}{M^2}\right) \quad \forall \, k, \\ &\leq 2\exp\left(\frac{-2\alpha^{k-\mathcal{T}_\varepsilon} N_0 \zeta^2}{M^2}\right) \quad \forall \, k > \mathcal{T}_\varepsilon, \\ &\longrightarrow 0 \quad \text{as } k \to \infty, \text{ since } \alpha > 1. \end{aligned}$$

Furthermore, it is easy to see that

$$\sum_{k=1}^{\infty} P\left(\left|\frac{1}{N_k}\sum_{j=1}^{N_k} b_i(H(X_j^k)) - \widetilde{E}_{\widetilde{\theta}_k}[b_i(H(X))]\right| > \zeta\right) \leq 2\sum_{k=1}^{\infty} \exp\left(\frac{-2\alpha^{k-\mathcal{T}_\varepsilon}N_0\zeta^2}{M^2}\right) < \infty.$$

By the Borel-Cantelli lemma,

$$P\left(\left|\frac{1}{N_k}\sum_{j=1}^{N_k} b_i(H(X_j^k)) - \widetilde{E}_{\widetilde{\theta}_k}[b_i(H(X))]\right| > \zeta \ \ i.o.\right) = 0.$$

This implies that $\frac{1}{N_k}\sum_{j=1}^{N_k} b_i(H(X_j^k)) \to \widetilde{E}_{\widetilde{\theta}_k}[b_i(H(X))]$ as $k \to \infty$ w.p.1. Note that by using a similar argument as above, we can also show that $\bar{\ell}_k(v) \to \ell_k(v)$ w.p.1 as $k \to \infty$.

The above result together with (34) and (35) implies that for any $\delta > 0$, there exists a small neighborhood B_v of v such that

$$\sup\{|\bar{\ell}_k(u) - \bar{\ell}_k(v)| : \ u \in B_v\} < \delta \ \ \text{w.p.1 for } k \text{ sufficiently large.}$$

Since this holds for all $v \in \mathcal{V}$, we have $\mathcal{V} \subseteq \cup_{v \in \mathcal{V}} B_v$, and because $\mathcal{V}$ is compact, there exists a finite subcover $B_{v_1}, \ldots, B_{v_m}$ such that

$$\sup\{|\bar{\ell}_k(u) - \bar{\ell}_k(v_j)| : \ u \in B_{v_j}\} < \delta \ \ \text{w.p.1 for } k \text{ sufficiently large, and } \mathcal{V} \subseteq \cup_{j=1}^m B_{v_j}.$$

Furthermore, by the continuity of $\ell_k(v)$, these open balls can be chosen in such a way that

$$\sup\{|\ell_k(u) - \ell_k(v_j)| : \ u \in B_{v_j}\} < \delta \ \ \forall \, j = 1, \ldots, m.$$

Since $\bar{\ell}_k(v_j) \to \ell_k(v_j)$ w.p.1 as $k \to \infty$ for all $j = 1, \ldots, m$,

$$|\bar{\ell}_k(v_j) - \ell_k(v_j)| < \delta \ \ \text{w.p.1 for } k \text{ sufficiently large, } \forall \, j = 1, \ldots, m.$$

For any $v \in \mathcal{V}$, without lost of generality assume $v \in B_{v_j}$, we have w.p.1 for k sufficiently large

$$|\bar{\ell}_k(v) - \ell_k(v)| \leq |\bar{\ell}_k(v) - \bar{\ell}_k(v_j)| + |\ell_k(v) - \ell_k(v_j)| + |\bar{\ell}_k(v_j) - \ell_k(v_j)| < 3\delta,$$

which implies that $\bar{\ell}_k(v) \to \ell_k(v)$ uniformly w.p.1 on $\mathcal{V}$.

The rest of the proof follows from Theorem $A1$ in Rubinstein and Shapiro (1993) (pp. 69), which basically states that if $\bar{\ell}_k(v) \to \ell_k(v)$ uniformly w.p.1, then the distance from $\widetilde{\gamma}_k^\dagger(\rho^\dagger, N_k)$ to $\gamma_k^\dagger$ tends to zero w.p.1 as $k \to \infty$. ∎