

A TWO-TIME-SCALE ADAPTIVE SEARCH ALGORITHM FOR GLOBAL OPTIMIZATION

Qi Zhang
Jiaqiao Hu

Department of Applied Mathematics and Statistics
State University of New York at Stony Brook
Stony Brook, NY 11794, USA

ABSTRACT

We study a random search algorithm for solving deterministic optimization problems in a black-box scenario. The algorithm has a model-based nature and finds improved solutions by sampling from a distribution model over the feasible region that gradually concentrates its probability mass around high quality solutions. In contrast to many existing algorithms in the class, which are population-based, our approach combines random search with a two-time-scale stochastic approximation idea to address a certain ratio bias inherent in these algorithms and uses only a single candidate solution per iteration. We prove global convergence of the algorithm and carry out numerical experiments to illustrate its performance.

1 INTRODUCTION

We consider the following optimization problem:

$$x^* \in \arg \max_{x \in \mathbb{X}} H(x), \quad (1)$$

where the feasible region \mathbb{X} is a nonempty compact subset of \mathfrak{R}^d , and $H : \mathbb{X} \rightarrow \mathfrak{R}$ is a bounded, deterministic objective function. We assume that the optimal solution x^* is unique and that an upper bound H^* of the objective function is known. We consider solving the problem in a black-box setting where an explicit analytical expression of H may not be known but its exact function value can be observed (e.g., through computer simulation) for a specified solution. In addition, no structural properties of H , such as convexity, continuity, or differentiability, are assumed available.

A variety of approaches have been proposed in the literature for solving such problems. Among them, random search methods are especially prominent. Typically, these methods rely only on the objective function values rather than structural information, and thus can be easily implemented and applied to a broad class of problems. Based on the way new candidate solutions are generated, random search algorithms can be divided into two classes, called instance-based and model-based (Zlochin, Birattari, Meuleau, and Dorigo 2004). In instance-based methods, new candidate solutions are constructed directly based on the previously generated ones. Examples include simulated annealing (Kirkpatrick, Gelatt, and Vecchi 1983), genetic algorithms (Goldberg 1989), tabu search (Glover 1990), nested partitions (Shi and Ólafsson 2000), and evolutionary programming (Eiben and Smith 2003). The other class, model-based search, generates solutions from a sequence of probabilistic models on the solution space (Hu 2014). Specific model-based algorithms include the Cross-Entropy (CE) method (Rubinstein 1997, Rubinstein 1999, Rubinstein and Kroese 2004), estimation of distribution algorithms (Larrañaga and Lozano 2002), and model reference adaptive search (Hu, Fu, and Marcus 2007).

The focus of this paper is on a particular model-based algorithm called Annealing Adaptive Search (AAS) (Zabinsky 2003), which was originally introduced in Romeijn and Smith (1994) to investigate the

behavior of simulated annealing. However, AAS is idealized and cannot be carried out exactly because it requires sampling candidate solutions from a sequence of Boltzmann distributions that depends on the objective function H ; an issue that is known to be very difficult. To resolve this implementation difficulty, Hu and Hu (2011) proposed a Model-based Annealing Random Search (MARS) algorithm that samples solutions from a sequence of parameterized surrogate distributions. This, in effect, converts a difficult optimization problem into an estimation problem for finding the surrogate distribution parameters that provide the best approximation to the corresponding Boltzmann distributions.

However, due to a certain ratio bias inherent to the MARS algorithm, its theoretical convergence requires the number of sampled candidate solutions to increase polynomially with the number of algorithm iterations. This may have an adverse impact on the algorithm's practical performance, especially on high-dimensional problems, where a large number of iterations is often needed to obtain a solution with an acceptable accuracy. In this paper, we aim to address this limitation of MARS by proposing a novel two-time-scale implementation of the algorithm called TS-MARS. The idea is to incorporate two separate but nested stochastic approximation type of recursions into the algorithm. One recursion is used to iteratively update the parameters of the surrogate model, whereas the other replaces the sample average estimators used in the original MARS by a recursive estimator to eliminate the ratio bias effect during the course of the iterations. A similar approach has been independently developed for the CE method in Joseph and Bhatnagar (2016) but requires the estimation of distribution quantiles. In contrast to existing algorithms in the model-based class, which are mostly population-based, TS-MARS only requires a single candidate solution to be generated at each iteration. We show that properly designed time scales (step-size schedules) of the two recursions can lead to the desired global convergence of the algorithm. Preliminary empirical tests also indicate that the TS-MARS is promising and could result in significant computational savings compared to MARS.

The rest of the paper is organized as follows. In Section 2, we motivate our work and describe the proposed TS-MARS algorithm. The convergence result of TS-MARS is presented in Section 3. Illustrative numerical studies on a set of four benchmark problems are given in Section 4. Finally, we summarize the paper in Section 5.

2 A TWO-TIME SCALE MARS ALGORITHM

The search for improved solutions in model-based methods is based on constructing and sampling from a sequence of convergent distribution models $\{g_k\}$ on \mathbb{X} that iteratively focuses on the set of optimal solutions. The various model-based algorithm mentioned in Section 1 correspond to different forms of $\{g_k\}$. However, all algorithms face the same problem that g_k often depends explicitly on the objective function H , which may not be available in any explicit form. In addition, sampling exactly from even a known g_k is in general intractable. One natural approach to resolve these difficulties is to specify a family of parameterized distributions $\{f_\theta, \theta \in \Theta\}$ and then project $\{g_k\}$ onto the family to obtain a sequence of surrogate distributions that retain the desired (convergence) properties of $\{g_k\}$. The projection is carried out at each step by finding an optimal parameter θ_k that minimizes the Kullback-Leibler (KL) divergence between g_k and the parameterized distribution f_θ (Rubinstein and Kroese 2004), i.e.,

$$\theta_k = \arg \min_{\theta \in \Theta} \mathfrak{D}(g_k, f_\theta) := \arg \min_{\theta \in \Theta} E_{g_k} \left[\ln \frac{g_k(X)}{f_\theta(X)} \right], \quad (2)$$

where $E_g[\cdot]$ denotes the expectation with respect to g . As in Hu, Fu, and Marcus (2007), we call g_k the reference distribution since it is only used implicitly to guide the construction of sampling distribution f_θ . Noted that f_θ does not depend on H directly, and because it is chosen by the user, sampling from f_θ is much easier than from g_k . From another viewpoint, this projection procedure reformulates the original problem (1) as an estimation problem for finding a parameter sequence $\{\theta_k\}$ on the parameter space Θ that provides the best approximation to the $\{g_k\}$ sequence. In practice, $\{f_\theta, \theta \in \Theta\}$ is often chosen from the so-called natural exponential family (NEF) due to its flexibility and many appealing theoretical properties.

Definition 1 A parameterized family of density/mass functions $\{f_\theta, \theta \in \Theta \subseteq \mathfrak{R}^k\}$ on \mathbb{X} is called a natural exponential family if there exist mappings $\Gamma(\cdot) : \mathfrak{R}^d \rightarrow \mathfrak{R}^k$ and $K(\cdot) : \mathfrak{R}^k \rightarrow \mathfrak{R}$ such that $f_\theta(x) = \exp(\theta^T \Gamma(x) - K(\theta))$, where $K(\theta) = \ln \int_{\mathbb{X}} \exp(\theta^T \Gamma(x)) dx$.

In the MARS algorithm, the reference model is taken as the convex combination of the Boltzmann distribution and the current surrogate distribution f_{θ_k} . When NEF is used, this choice of the reference distribution naturally leads to a gradient-descent type of recursion for solving the optimization problem (2). To be precise, denote the Boltzmann distribution with temperature T_k by

$$g_{k+1}(x) = \frac{e^{H(x)/T_k}}{\int_{\mathbb{X}} e^{H(x)/T_k} dx}, \quad \forall x \in \mathbb{X}.$$

Then, the reference distribution $\hat{g}_{k+1}(x)$ used in MARS is given by

$$\hat{g}_{k+1}(x) = \alpha_k g_{k+1}(x) + (1 - \alpha_k) f_{\theta_k}(x), \quad \alpha_k \in [0, 1] \quad \forall k. \quad (3)$$

When θ_{k+1} is obtained by minimizing $\mathfrak{D}(\hat{g}_{k+1}, f_\theta)$, the first order necessary condition for optimality implies that

$$m(\theta_{k+1}) = m(\theta_k) - \alpha_k \nabla_\theta \mathfrak{D}(g_{k+1}, f_\theta)|_{\theta=\theta_k} \quad (4)$$

$$= m(\theta_k) + \alpha_k (E_{g_{k+1}}[\Gamma(X)] - m(\theta_k)), \quad (5)$$

where $m(\theta) := \nabla_\theta K(\theta) = E_\theta[\Gamma(X)]$ is called the mean parameter function of the NEF, which is a one-to-one transformation of θ .

In actual implementation of MARS, the expectation $E_{g_{k+1}}[\Gamma(X)]$ in (5) is estimated by sample averages:

$$E_{g_{k+1}}[\Gamma(X)] = \frac{\int_{\mathbb{X}} \exp(H(x)/T_k) \Gamma(x) dx}{\int_{\mathbb{X}} \exp(H(x)/T_k) dx} = \frac{E[\exp(H(X)/T_k) \Gamma(X) f_{\theta_k}^{-1}(X)]}{E[\exp(H(X)/T_k) f_{\theta_k}^{-1}(X)]} \quad (6)$$

$$\approx \frac{\frac{1}{N_k} \sum_{i=1}^{N_k} \exp(H(X_i)/T_k) \Gamma(X_i) f_{\theta_k}^{-1}(X_i)}{\frac{1}{N_k} \sum_{i=1}^{N_k} \exp(H(X_i)/T_k) f_{\theta_k}^{-1}(X_i)}, \quad (7)$$

where $\{X_i : i = 1, \dots, N_k\}$ is the population of candidate solutions generated from f_{θ_k} at the k th iteration. Although such an approach is straightforward to implement, the resulting estimator (7) is biased for any finite sample size N_k because (6) comprises ratio of expectations. Consequently, the convergence analysis of MARS requires increasing N_k polynomially to reduce the bias of the estimator. Moreover, it can be seen that the use of $f_{\theta_k}^{-1}(\cdot)$ is artificial because the integrals defining $E_{g_{k+1}}[\Gamma(X)]$ do not involve f_{θ_k} at all. In other words, any other sampling distribution can be used in lieu of f_{θ_k} without changing the value of the expectation. From a practical point of view, the division by f_{θ_k} in (7) may cause numerical stability issues, especially on high-dimensional problems when f_{θ_k} becomes a product of many marginal density/mass functions.

2.1 Algorithm Description

In this section, we address the limitations of MARS by describing a two-time-scale version of the algorithm. For simplicity, we focus on continuous optimization problems with box constraints, i.e., $\mathbb{X} = \{(x_1, \dots, x_d)^T : a_i \leq x_i \leq b_i, i = 1, \dots, d\}$ and choose the parameterized distribution f_{θ_k} at each iteration as an independent multivariate normal density with mean vector $\mu_k = (\mu_{k,1}, \dots, \mu_{k,d})$ and diagonal covariance matrix $\Sigma_k = \text{diag}(\sigma_{k,1}^2, \dots, \sigma_{k,d}^2)$ (cf. e.g., Hu, Fu, and Marcus 2007). By putting f_{θ_k} into the form of NEF, it is easy to verify that its sufficient statistic $\Gamma(x)$ and parameter θ_k are given by

$$\Gamma(x) = (x_1, \dots, x_d, x_1^2, \dots, x_d^2)^T \quad \text{and} \quad \theta_k = \left(\frac{\mu_k}{\sigma_k^2}, -\frac{1}{2\sigma_k^2} \right)^T,$$

where $\sigma_k^2 = (\sigma_{k,1}^2, \dots, \sigma_{k,d}^2)$, and the division is component-wise.

We consider the following distribution model φ_{k+1} formed by tilting the surrogate density f_{θ_k} with the exponential term $e^{H(x)/T_k}$:

$$\varphi_{k+1}(x) := \frac{e^{H(x)/T_k} f_{\theta_k}(x)}{\int_{\mathbb{X}} e^{H(x)/T_k} f_{\theta_k}(x) dx}, \quad \forall x \in \mathbb{X}. \quad (8)$$

This construction assigns more weights to solutions with higher objection function values so that φ_{k+1} will become more concentrated on regions containing high quality solutions as the temperature T_k decreases. Note that φ_{k+1} has a similar form as the reference distribution used in the CE method (Rubinstein and Kroese 2004, Hu and Hu 2009) but avoids the need for estimating the distribution quantiles. The reference model $\hat{\varphi}_{k+1}$ used in the proposed TS-MARS algorithm has the same form as in (3) with φ_{k+1} replacing g_{k+1} . In particular, when θ_{k+1} is obtained by minimizing $\mathfrak{D}(\hat{\varphi}_{k+1}, f_{\theta})$, the formula for updating the mean parameter function $m(\theta_k)$, i.e., (5), becomes

$$\begin{aligned} m^{(1)}(\theta_{k+1}) &= m^{(1)}(\theta_k) + \alpha_k (E_{\varphi_{k+1}}[X] - m^{(1)}(\theta_k)), \\ m^{(2)}(\theta_{k+1}) &= m^{(2)}(\theta_k) + \alpha_k (E_{\varphi_{k+1}}[X^2] - m^{(2)}(\theta_k)), \end{aligned}$$

where $m^{(1)}(\theta_k)$ and $m^{(2)}(\theta_k)$ are the two components of $m(\theta_k)$.

Thus, the problem reduces to how to efficiently estimate the two expectations $E_{\varphi_{k+1}}[X]$ and $E_{\varphi_{k+1}}[X^2]$. We use the former case to illustrate the idea. From (8), $E_{\varphi_{k+1}}[X]$ can be written as

$$E_{\varphi_{k+1}}[X] = \frac{E_{\theta_k}[\exp(H(X)/T_k)X]}{E_{\theta_k}[\exp(H(X)/T_k)]}, \quad (9)$$

where $E_{\theta_k}[\cdot]$ denotes the expectation with respect to f_{θ_k} . As discussed earlier, a straightforward application of sample average approximation would result in a ratio bias, so we instead consider an alternative recursive estimator of the form

$$U_{k+1}^{(1)} = U_k^{(1)} + \beta_k \left(e^{\frac{H(X_k)}{T_k}} X_k - e^{\frac{H(X_k)}{T_k}} U_k^{(1)} \right), \quad (10)$$

where X_k is a candidate solution sampled from f_{θ_k} and β_k is the updating step size. Since both $e^{\frac{H(X_k)}{T_k}} X_k$ and $e^{\frac{H(X_k)}{T_k}}$ are the respective unbiased estimators for $E_{\theta_k}[e^{\frac{H(X)}{T_k}} X]$ and $E_{\theta_k}[e^{\frac{H(X)}{T_k}}]$, the sequence $\{U_k^{(1)}\}$ generated by (10) can be seen to approach the solution set of the mean ODE

$$\frac{dU^{(1)}(t)}{dt} = E_{\theta(t)}[e^{\frac{H(X)}{T(t)}} X] - E_{\theta(t)}[e^{\frac{H(X)}{T(t)}}] U^{(1)}(t), \quad (11)$$

where $U^{(1)}(t)$, $\theta(t)$ and $T(t)$ are the continuous-time interpolations of $\{U_k^{(1)}\}$, $\{\theta_k\}$ and $\{T_k\}$, respectively. Note that directly analyzing (11) can be difficult due to its nonlinear, time-varying nature. However, if both $\theta(t)$ and $T(t)$ were held constant, i.e., $\theta(t) \equiv \theta$ and $T(t) \equiv T$, then the complexity of (11) could be dramatically reduced, giving rise to a linear time-invariant ODE, whose unique globally asymptotic equilibrium is given by $\frac{E_{\theta}[\exp(H(X)/T)X]}{E_{\theta}[\exp(H(X)/T)]}$, which has the same form as (9). Consequently, the idea is to use (10) as an intermediate estimation scheme to resolve the ratio bias issue in estimating $E_{\varphi_k}[X]$. Similarly, $E_{\varphi_{k+1}}[X^2]$ can be estimated by using the following iterative procedure analogous to (10):

$$U_{k+1}^{(2)} = U_k^{(2)} + \beta_k \left(e^{\frac{H(X_k)}{T_k}} X_k^2 - e^{\frac{H(X_k)}{T_k}} U_k^{(2)} \right).$$

This leads to the following two-time-scale algorithm we propose.

Two-time-scale Model-based Annealing Random Search (TS-MARS)

Step 0: Initialize T_0 , θ_0 , $U_0^{(1)}$, $U_0^{(2)}$ and η_0 . Specify annealing parameters $\{\lambda_k\}$ and two step-size sequences $\{\alpha_k\}$ and $\{\beta_k\}$. Set iteration counter $k = 0$.

Step 1: Generate a candidate solution X_k from f_{θ_k} .

Step 2: Compute $U_{k+1}^{(1)}$ and $U_{k+1}^{(2)}$ as

$$U_{k+1}^{(1)} = U_k^{(1)} + \beta_k \left(e^{\frac{H(X_k) - H^*}{T_k}} X_k - e^{\frac{H(X_k) - H^*}{T_k}} U_k^{(1)} \right), \quad (12)$$

$$U_{k+1}^{(2)} = U_k^{(2)} + \beta_k \left(e^{\frac{H(X_k) - H^*}{T_k}} X_k^2 - e^{\frac{H(X_k) - H^*}{T_k}} U_k^{(2)} \right). \quad (13)$$

Step 3: Update η_{k+1} by

$$\eta_{k+1} = \eta_k + \alpha_k \left(\begin{bmatrix} U_{k+1}^{(1)} \\ U_{k+1}^{(2)} \end{bmatrix} - \eta_k \right). \quad (14)$$

Step 4: Obtain a new parameter $\theta_{k+1} = m^{-1}(\eta_{k+1})$ and calculate a new temperature

$$T_{k+1} = T_k - \lambda_k T_k. \quad (15)$$

Set $k = k + 1$ and go back to Step 1.

Recall that H^* used at Step 2 is an upper bound of the objective function H . The negative difference $H(X_k) - H^*$ guarantees the iterates generated by (12) and (13) to remain bounded in all iterations (see Lemma 1 in Section 3 below). The step-size sequences $\{\alpha_k\}$, $\{\beta_k\}$ and $\{\lambda_k\}$ should be chosen to ensure that $U_k^{(1)}$ and $U_k^{(2)}$ are updated on a faster timescale than T_k and η_k , i.e., $\alpha_k = o(\lambda_k)$, $\lambda_k = o(\beta_k)$. so that when viewed from (12) and (13), the distribution and temperature parameters on the slower timescales would appear to be fixed. Under such choices of time-scales, we show in Section 3 that $U_k^{(1)}$ and $U_k^{(2)}$ indeed track $E_{\phi_k}[X]$ and $E_{\phi_k}[X^2]$ as desired.

3 GLOBAL CONVERGENCE OF TS-MARS

Since TS-MARS randomly selects candidate solutions from the sampling distributions, we let $\mathfrak{F}_k := \sigma\{X_1, \dots, X_{k-1}\}$ be the sequence of increasing σ -fields generated by the solutions sampled up to iteration $k-1$. Thus, conditional on \mathfrak{F}_k , the parameter θ_k is completely determined. We will occasionally use φ_{T_k, θ_k} to denote the distribution model φ_{k+1} defined in (8) to emphasize its dependency on T_k and θ_k . Finally, to simplify exposition, we combine iterations (12) and (13) into a single recursion in our analysis,

$$U_{k+1} = U_k + \beta_k \left(e^{\frac{H(X_k) - H^*}{T_k}} \Gamma(X_k) - e^{\frac{H(X_k) - H^*}{T_k}} U_k \right), \quad (16)$$

where $U_k := (U_k^{(1)}, U_k^{(2)})^T$.

We make the following assumptions on the objective function and the step-size parameters used in TS-MARS. In particular, A1 and A2 are directly taken from Hu, Fu, and Marcus (2007) and A3-A4 are standard in studying two-time-scale stochastic approximation algorithms (cf. e.g., Borkar (1996)).

Assumptions:

- A1.** For any constant $\varepsilon < H(x^*)$, the set of solutions with function values greater than or equal to ε has a strictly positive Lebesgue measure, i.e., $\nu(\{x \in \mathbb{X} : H(x) \geq \varepsilon\}) > 0$.
- A2.** For any given $\delta > 0$, $\sup_{x \in A_\delta} H(x) < H(x^*)$, where $A_\delta := \{x \in \mathbb{X} : \|x - x^*\| \geq \delta\}$.

A3. The step-size sequences $\{\alpha_k\}$, $\{\beta_k\}$ and $\{\lambda_k\}$ satisfy

- (a) $\alpha_k > 0 \forall k$, $\sum_{k=0}^{\infty} \alpha_k = \infty$, and $\sum_{k=0}^{\infty} \alpha_k^2 < \infty$;
- (b) $\beta_k > 0 \forall k$, $\sum_{k=0}^{\infty} \beta_k = \infty$, and $\sum_{k=0}^{\infty} \beta_k^2 < \infty$;
- (c) $\lambda_k \in (0, 1)$.

A4. The decay rates of the gain sequences satisfy

$$\alpha_k = o(\lambda_k), \quad \lambda_k = o(\beta_k).$$

The next result shows that all random quantities generated at Steps 2 and 3 of TS-MARS stay bounded w.p.1. The proof follows from a straightforward inductive argument by using the convex combination structure of (12)–(14) and the compactness of \mathbb{X} .

Lemma 1 If A3(a)(b) hold, then $P(\|U_k\| < \infty, \forall k) = 1$ and $P(\|\eta_k\| < \infty, \forall k) = 1$.

Next, we establish that for the class of optimization problems characterized by A1 and A2, the sequence of distribution models $\{\varphi_k\}$ converges weakly to a limiting distribution concentrated at x^* .

Lemma 2 If A1-A4 are satisfied, then

$$E_{\varphi_k}[\Gamma(X)] \rightarrow \Gamma(x^*) \text{ as } k \rightarrow \infty,$$

where the limit is component-wise.

Proof sketch: Let $B := \max\{\|x\| : x \in \mathbb{X}\}$ and denote the complement of A_δ by A_δ^c . Note that for any given $\xi > 0$, if $x \in A_\delta^c$ with $\delta := \frac{\xi}{2B+1}$, then we have

$$\begin{aligned} \|\Gamma(x) - \Gamma(x^*)\| &\leq \|x - x^*\| + \|x^2 - (x^*)^2\| \\ &= \|x - x^*\|(\|x + x^*\| + 1) < \frac{\xi}{2B+1}(2B+1) = \xi. \end{aligned}$$

Next, we define $\bar{H} := \sup_{x \in A_\delta} H(x)$ and $\varepsilon := \frac{\bar{H} + H(x^*)}{2}$. It is clear that $\bar{H} < H(x^*)$ by A2, and A1 implies that the set $B_\varepsilon := \{x \in \mathbb{X} : H(x) > \varepsilon\}$ has a positive Lebesgue measure. Thus,

$$\begin{aligned} \|E_{\varphi_{k+1}}[\Gamma(X)] - \Gamma(x^*)\| &\leq E_{\varphi_{k+1}}[\|\Gamma(X) - \Gamma(x^*)\|] \\ &= \int_{A_\delta^c} \|\Gamma(x) - \Gamma(x^*)\| \varphi_{k+1}(x) dx + \int_{A_\delta} \|\Gamma(x) - \Gamma(x^*)\| \varphi_{k+1}(x) dx \\ &\leq \xi + \sup_{x \in \mathbb{X}} \|\Gamma(x) - \Gamma(x^*)\| \frac{\int_{A_\delta} e^{\frac{H(x)}{T_k}} e^{\theta_k^T \Gamma(x)} dx}{\int_{\mathbb{X}} e^{\frac{H(x)}{T_k}} e^{\theta_k^T \Gamma(x)} dx} \\ &\leq \xi + \sup_{x \in \mathbb{X}} \|\Gamma(x) - \Gamma(x^*)\| \frac{\int_{A_\delta} e^{\frac{\bar{H}}{T_k}} e^{\sum_{i=1}^d \left(\frac{\mu_{k,i}}{\sigma_{k,i}^2} x_i - \frac{1}{2\sigma_{k,i}^2} x_i^2 \right)} dx}{\int_{B_\varepsilon} e^{\frac{\varepsilon}{T_k}} e^{\sum_{i=1}^d \left(\frac{\mu_{k,i}}{\sigma_{k,i}^2} x_i - \frac{1}{2\sigma_{k,i}^2} x_i^2 \right)} dx} \\ &\leq \xi + \sup_{x \in \mathbb{X}} \|\Gamma(x) - \Gamma(x^*)\| e^{\frac{1}{T_k} (\frac{\bar{H} - H(x^*)}{2} + 3B^2 \sum_{i=1}^d \frac{T_k}{\sigma_{k,i}^2})} \frac{v(A_\delta)}{v(B_\varepsilon)}. \end{aligned}$$

By A3(c), we have $\lim_{k \rightarrow \infty} T_k = 0$. In addition, A4 indicates that the speed T_k converges to zero is faster than the variation rate of η_k . This implies $T_k/\sigma_{k,i}^2 \rightarrow 0$ for all i . Hence the exponential term in the last equation goes to zero as $k \rightarrow \infty$. On the other hand, $\Gamma(x)$ is bounded and B_ε has a positive Lebesgue measure by A1. Therefore, since ξ is arbitrary, we have $\|E_{\varphi_{k+1}}[\Gamma(X)] - \Gamma(x^*)\| \rightarrow 0$ as $k \rightarrow \infty$. \square

Note that the step size determines the updating rate of a recursion. A4 implies that η_k is updated at the slowest rate, whereas U_k varies at the fastest speed. Intuitively, the fast component sees the slow component as quasi-static while the slow component sees the fast one as being converged. Lemma 3 makes this intuition precise.

Lemma 3 If A3 and A4 hold, then $\|U_k - E_{\varphi_{T_k, \theta_k}}[\Gamma(X)]\| \rightarrow 0$ w.p.1.

Proof sketch: Rewrite equations (15) and (14) using time steps $\{\beta_k\}$ as

$$\begin{aligned} T_{k+1} &= T_k - \beta_k T_k', \\ \eta_{k+1} &= \eta_k + \beta_k (U_{k+1}' - \eta_k'), \end{aligned}$$

where $T_k' := \frac{\lambda_k}{\beta_k} T_k$, $U_{k+1}' := \frac{\alpha_k}{\beta_k} U_{k+1}$ and $\eta_k' := \frac{\alpha_k}{\beta_k} \eta_k$. We can then construct piece-wise constant interpolations of $\{T_k\}$, $\{U_k\}$ and $\{\eta_k\}$ in the same way as discussed in Kushner and Yin (1997). Let $t_0 = 0$ and $t_n = \sum_{i=0}^{n-1} \beta_i$. Define the continuous-time interpolation $U^0(t)$ with interpolation intervals $\{\beta_k\}$ by setting $U^0(t) = U_k$ for $t \in [t_k, t_{k+1})$. Define $T^0(t)$ and $\eta^0(t)$ analogously to $U^0(t)$ but using T_k and η_k respectively. Then, the shifted processes are given by

$$\begin{aligned} U^k(t) &= U^0(t + t_k), \\ T^k(t) &= T^0(t + t_k) - T^0(t_k), \\ \eta^k(t) &= \eta^0(t + t_k) - \eta^0(t_k). \end{aligned}$$

A standard approximation argument using the ODE method (cf. e.g., Kushner and Yin (1997), Kushner and Clark (1978), Borkar (2008)) shows that the limit of a convergent subsequence of the shifted process $\{U^k(\cdot), T^k(\cdot), \eta^k(\cdot)\}$ satisfies the following mean ODEs,

$$\begin{aligned} \dot{U}(t) &= \int e^{\frac{H(x) - H^*}{T(t)}} \Gamma(x) f_{\theta(t)}(x) dx - \int e^{\frac{H(x) - H^*}{T(t)}} f_{\theta(t)}(x) dx \cdot U(t), \\ \dot{T}(t) &= 0, \\ \dot{\eta}(t) &= 0, \end{aligned}$$

where $\theta(t) = m^{-1}(\eta(t))$. For $T(t) \equiv T$ and $\eta(t) \equiv \eta$, the equation for $U(t)$ becomes a linear time-invariant ODE and has a unique global asymptotically stable equilibrium $E_{\varphi_{T, m^{-1}(\eta)}}[\Gamma(X)]$. This implies $(U_k, T_k, \eta_k) \rightarrow \{(E_{\varphi_{T, m^{-1}(\eta)}}[\Gamma(X)], T, \eta) : T \in (0, \infty), \eta \in \mathbb{R}^{2d}\}$ since the continuous interpolations characterize the asymptotic behavior of the corresponding discrete time sequences.

Note that

$$\|U_k - E_{\varphi_{T_k, \theta_k}}[\Gamma(X)]\| \leq \|U_k - E_{\varphi_{T, \theta}}[\Gamma(X)]\| + \|E_{\varphi_{T, \theta}}[\Gamma(X)] - E_{\varphi_{T_k, \theta_k}}[\Gamma(X)]\|,$$

where $\theta = m^{-1}(\eta)$. The first term on the right-hand side goes to zero as a result of the limiting ODEs. In addition, it can be verified that $E_{\varphi_{T, \theta}}[\Gamma(X)]$ is Lipschitz continuous in both T and θ . Thus, the second term on the right-hand side of the above equation also vanishes as k goes to infinity. Hence, we obtain $\|U_k - E_{\varphi_{T_k, \theta_k}}[\Gamma(X)]\| \rightarrow 0$ w.p.1. \square

Lemma 3 shows that U_k indeed tracks $E_{\varphi_{T_k, \theta_k}}[\Gamma(X)]$ asymptotically, this gives rise to our main convergence theorem.

Theorem 1 Assume A1-A4 hold, then

$$\eta_k \rightarrow \Gamma(x^*) \text{ as } k \rightarrow \infty \text{ w.p.1.}$$

Proof sketch: Rewrite equation (14) as

$$Y_{k+1} = Y_k - V_k,$$

where $Y_k := \eta_k - \Gamma(x^*)$, and $V_k := \alpha_k(Y_k + \Gamma(x^*) - U_{k+1})$. Let $M_k := E[V_k | \mathfrak{F}_k]$ and $Z_k := V_k - M_k$. Note that proving the claim is equivalent to showing $Y_k \rightarrow 0$ w.p.1. as $k \rightarrow \infty$. This desired result can be justified by verifying conditions (i)-(iv) in Evans and Weber (1986). The detailed derivations are very similar to those of Hu and Hu (2011) and are thus omitted. \square

Theorem 1 shows that $\lim_{k \rightarrow \infty} m(\theta_k) := \lim_{k \rightarrow \infty} E_{\theta_k}[\Gamma(X)] = \Gamma(x^*)$ w.p.1. Since we use normal densities as the parameterized family, it immediately follows that

$$\begin{aligned} \lim_{k \rightarrow \infty} E_{\theta_k}[X] &= x^*, \\ \lim_{k \rightarrow \infty} \text{Var}_{\theta_k}[X] &= \lim_{k \rightarrow \infty} (E_{\theta_k}[X^2] - (E_{\theta_k}[X])^2) = (x^*)^2 - (x^*)^2 = 0, \end{aligned}$$

where the limits and square operations are component-wise. Hence the sampling distribution sequence $\{f_{\theta_k}\}$ will converge to a Dirac delta function with all of its mass concentrated on the global optimizer x^* .

4 NUMERICAL EXAMPLES

To illustrate the effectiveness of the proposed algorithm, we consider some simple experiments on the following four benchmark test functions and compare the performance of TS-MARS with that of the MARS algorithm in Hu and Hu (2011).

1. Weighted Sphere function ($d = 10$, $-10 \leq x_i \leq 10$, $i = 1, \dots, d$)

$$H_1(x) = - \sum_{i=1}^d i(x_i - 1)^2 - 1$$

where $x^* = (1, \dots, 1)^T$, $H_1(x^*) = -1$.

2. Griewank function ($d = 10$, $-10 \leq x_i \leq 10$, $i = 1, \dots, d$)

$$H_2(x) = -\frac{1}{4000} \sum_{i=1}^d (x_i - 1)^2 + \prod_{i=1}^d \cos\left(\frac{x_i - 1}{\sqrt{i}}\right) - 1$$

where $x^* = (1, \dots, 1)^T$, $H_2(x^*) = 0$.

3. Trigonometric function ($d = 10$, $-10 \leq x_i \leq 10$, $i = 1, \dots, d$)

$$H_3(x) = - \sum_{i=1}^d [8 \sin^2(7(x_i - 0.9)^2) + 6 \sin^2(14(x_i - 0.9)^2) + (x_i - 0.9)^2]$$

where $x^* = (0.9, \dots, 0.9)^T$, $H_3(x^*) = 0$.

4. Levy function ($d = 10$, $-10 \leq x_i \leq 10$, $i = 1, \dots, d$)

$$H_4(x) = -10 \sin^2(\pi(x_1 - 1)) - \sum_{i=1}^{d-1} 100(x_i - 1)^2 (1 + 10 \sin^2(\pi(x_{i+1} - 1))) - 100(x_d - 2)^2 - 1$$

where $x^* = (1, \dots, 1, 2)^T$, $H_4(x^*) = -1$.

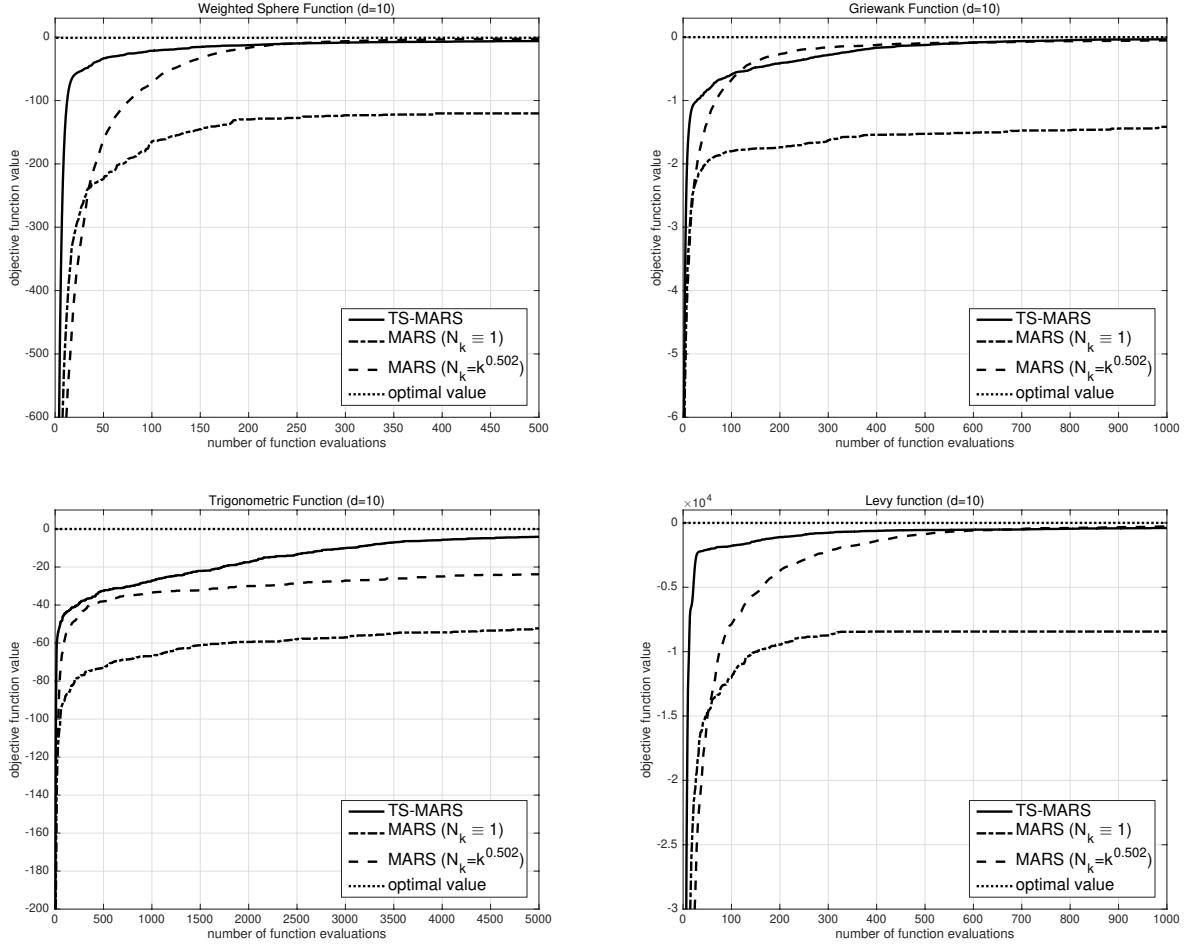


Figure 1: Averaged performance of TS-MARS and MARS on test functions H_1 to H_4 .

In all cases, the normal density in TS-MARS is initialized by choosing its mean vector uniformly from the domain $[-10, 10]^d$ and its covariance matrix as a $d \times d$ diagonal matrix with all diagonal entries equal to 100. The gain sequences for η_k and U_k are set to $\alpha_k = 20/(k+500)^{0.8}$ and $\beta_k = 5/(k+500)^{0.51}$, both satisfying the conditions in A3. It is important to be able to explore the entire domain in the early stage, so we choose a relatively large initial temperature $T_0 = 100$. The annealing schedule $\{T_k\}$ determines the varying speed of the distribution model φ_k , and we use an adaptive step-size $\lambda_k = (0.001 + |H(x_k^*)|)/(k^{0.79} + 100|H(x_k^*)|)$ to control its decay rate, where x_k^* denotes the current best solution found at the k th iteration of the algorithm. The same set of parameter values are used in implementing the MARS algorithm. In addition, we consider two different sample sizes in MARS: a constant sample size $N_k \equiv 1$ and a polynomially increasing sample size $N_k = \lfloor k^{0.502} \rfloor$ (Hu and Hu 2011), where $\lfloor a \rfloor$ denotes the integer part of a .

For each test function, we performed 50 independent replication runs of both algorithms. The results are presented in Figure 1, which plots the averaged current best objective function values as a function of the number of function evaluations used. The figure indicates the convergence of TS-MARS even when a single candidate solution is used at each iteration. In addition, we see that TS-MARS also shows good performance relative to MARS with polynomially increasing sample size, with its performance consistently dominating the latter in three out of the four testing cases. On the Griewank function H_2 , TS-MARS shows a faster initial improvement but is slightly outperformed by MARS when the number of function evaluations is between 100 and 500. We conjecture that this temporary degradation in performance of

TS-MARS is primarily due to the high nonlinearity of the test function, so that using a single candidate solution at each step may result in large variances of the estimation recursions (12) and (13), leading to slow convergence. In all cases, MARS with constant sample size $N_k \equiv 1$ does not seem to work well and may frequently stagnate at solutions that are far from optimal due to the ratio bias.

5 CONCLUSIONS

In this paper, we have presented a novel algorithm called TS-MARS for solving black-box optimization problems. The algorithm can be viewed as a variant of the MARS algorithm but employs a two-time-scale stochastic approximation idea to overcome the ratio bias issue in MARS. In contrast to most existing model-based methods, which are population-based, such a two-time-scale procedure enables the underlying algorithm to use a single candidate solution at each iteration while retaining the desired global convergence properties. Our preliminary empirical results indicate that the new algorithm is promising and may outperform the original MARS algorithm.

ACKNOWLEDGMENTS

This work was supported by the National Science Foundation under Grant CMMI-1634627.

REFERENCES

Borkar, V. S. 1996. "Stochastic Approximation with Two Time Scales". *Systems Control Letters* 29:291–294.

Borkar, V. S. 2008. *Stochastic Approximation: A Dynamical Systems Viewpoint*. Cambridge University Press.

Eiben, A. E., and J. E. Smith. 2003. *Introduction to Evolutionary Computing*. Berlin Heidelberg: Springer-Verlag.

Evans, S. N., and N. C. Weber. 1986. "On the Almost Sure Convergence of A General Stochastic Approximation Procedure". *Bull Australian Math Soc* 34:335–342.

Glover, F. 1990. "Tabu Search: A Tutorial". *Interfaces* 20:74–94.

Goldberg, D. E. 1989. *Genetic Algorithms in Search, Optimization, and Machine Learning*. Addison-Wesley Longman Publishing Co., Inc. Boston, MA.

Hu, J. 2014. "Model-based Stochastic Search Methods". In *Chapter 12 in Handbook of Simulation Optimization*, edited by M. C. Fu. New York: Springer.

Hu, J., M. C. Fu, and S. I. Marcus. 2007. "A Model Reference Adaptive Search Algorithm for Global Optimization". *Operations Research* 55:549–568.

Hu, J., and P. Hu. 2009. "On the Performance of the Cross-Entropy Method". In *Proceedings of the 2009 Winter Simulation Conference*, edited by M. D. Rossetti, R. R. Hill, B. Johansson, A. Dunkin, and R. Ingalls, 459–468. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.

Hu, J., and P. Hu. 2011. "Annealing Adaptive Search, Cross-Entropy, and Stochastic Approximation in Global Optimization". *Naval Research Logistics* 58:457–477.

Joseph, A. G., and S. Bhatnagar. 2016. "A Randomized Algorithm for Continuous Optimization". In *Proceedings of the 2016 Winter Simulation Conference*, edited by T. M. K. Roeder, P. I. Frazier, R. Szechtman, and E. Zhou, 907–918. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.

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

Kushner, H. J., and D. S. Clark. 1978. *Stochastic Approximation Methods for Constrained and Unconstrained Systems*. New York: Springer.

Kushner, H. J., and G. G. Yin. 1997. *Stochastic Approximation and Recursive Algorithms and Applications*. New York: Springer.

Larrañaga, P., and J. A. Lozano. 2002. *Estimation of Distribution Algorithms: A New Tool for Evolutionary Computation*. Kluwer Academic Publishers, Boston, MA.

Romeijn, H., and R. Smith. 1994. "Simulated Annealing and Adaptive Search in Global Optimization". *Probability in the Engineering and Informational Sciences* 8:571–590.

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

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

Rubinstein, R. Y., and D. P. Kroese. 2004. *The Cross-Entropy Method: A Unified Approach to Combinatorial Optimization, Monte-Carlo Simulation, and Machine Learning*. New York: Springer-Verlag.

Shi, L., and S. Ólafsson. 2000. "Nested Partitions Method for Global Optimization". *Operations Research* 48:390–407.

Zabinsky, Z. B. 2003. *Stochastic Adaptive Search for Global Optimization*. Norwell, MA: Kluwer Academic Publishers.

Zlochin, M., M. Birattari, N. Meuleau, and M. Dorigo. 2004. "Model-based Search for Combinatorial Optimization: A Critical Survey". *Annals of Operations Research* 131:373–395.

AUTHOR BIOGRAPHIES

QI ZHANG is a Ph.D. Student in the Department of Applied Mathematics and Statistics at the State University of New York, Stony Brook. He received B.S. and M.Phil. degrees in Mathematics from the Chinese University of Hong Kong. His research interests are Markov decision process, simulation optimization, stochastic approximation, and reinforcement learning. His e-mail address is qi.zhang.1@stonybrook.edu.

JIAQIAO HU is an Associate Professor in the Department of Applied Mathematics and Statistics at the State University of New York, Stony Brook. He received a B.S. degree in automation from Shanghai Jiao Tong University, an M.S. degree in applied mathematics from the University of Maryland, Baltimore County, and a Ph.D. degree in electrical engineering from the University of Maryland, College Park. His research interests include Markov decision processes, applied probability, and simulation optimization. His e-mail address is jqhu@ams.sunysb.edu.