



HAL
open science

Adaptive simulated annealing through alternating Renyi divergence minimization

Thomas Guilmeau, Emilie Chouzenoux, Víctor Elvira

► **To cite this version:**

Thomas Guilmeau, Emilie Chouzenoux, Víctor Elvira. Adaptive simulated annealing through alternating Renyi divergence minimization. ICASSP 2023 - IEEE International Conference on Acoustics, Speech and Signal Processing, Jun 2023, Rhodes (Grèce), Greece. 10.1109/ICASSP49357.2023.10096619 . hal-04094489

HAL Id: hal-04094489

<https://inria.hal.science/hal-04094489>

Submitted on 11 May 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

ADAPTIVE SIMULATED ANNEALING THROUGH ALTERNATING RÉNYI DIVERGENCE MINIMIZATION

Thomas Guilmeau[†], Emilie Chouzenoux[†], and Víctor Elvira^{* ‡}

[†] Université Paris-Saclay, CentraleSupélec, Inria, CVN, Gif-sur-Yvette, France

^{*} School of Mathematics, University of Edinburgh, United Kingdom

[‡] The Alan Turing Institute, United Kingdom

ABSTRACT

Simulated annealing is a popular approach to solve non-convex and black-box optimization problems. It consists in running a non-homogeneous Markov chain to sample from a sequence of Boltzmann probability distributions. This sequence is controlled by a cooling schedule, which governs the concentration of the mass of the Boltzmann distributions around the global minimizers. However, convergence is often slow, difficult to assess, and requires a fixed cooling schedule. We propose here a new simulated annealing algorithm with adaptive cooling schedule, which draws samples from variational approximations of the Boltzmann distributions. Our approach is theoretically sound and relies on an alternating Bregman proximal-gradient scheme minimizing a regularized Rényi divergence. Numerical experiments illustrate the performance of the method.

Index Terms— Adaptive simulated annealing, alternating Bregman proximal-gradient algorithm, Rényi divergence.

1. INTRODUCTION

In this work, we focus on the black-box minimization of a non-convex cost. In this setting, only the cost function can be evaluated, without access to any extra information (e.g. gradients). This can occur while optimizing hyper-parameters [1] or when dealing with categorical variables [2]. Finding the global minimum in such problems requires schemes able to explore the space and avoid local minima. Many available methods are reviewed in [3], including the natural evolution strategies (NES) [4] or the cross-entropy (CE) method [5].

Among these methods, simulated annealing (SA) [6] relies on Boltzmann distributions, which are probability densities constructed from the cost function and parametrized by a temperature. Low-temperature Boltzmann densities are concentrated around the global minimizers, but it is hard to sam-

ple from them. Therefore, SA schemes use decreasing temperatures to generate samples approximating these Boltzmann distributions, and thus approaching the global minimizers.

The cooling schedule (i.e., the sequence of temperatures) must be fixed in advance and decrease slowly enough to reach convergence [7, 6]. In particular, the cooling schedule cannot be adapted to the quality of the samples generated in past iterations. Many adaptive SA algorithms have thus been proposed [8, 9] to circumvent these issues, but mostly without the theoretical guarantees of the original SA scheme [6].

Most SA methods rely on Markov chain Monte Carlo schemes to draw samples, which allows to derive asymptotic convergence results, but can make convergence slow and difficult to assess [10]. Sampling from parametric densities, as in the NES [4] and CE [5] algorithms, provides an interesting perspective for SA schemes, as demonstrated in [11].

Our recent APSA algorithm [12] provides a methodology to adapt the cooling schedule in SA while generating parametric approximations of the Boltzmann distributions. To do so, it reformulates the problem into the minimization of the sum of the Kullback-Leibler (KL) divergence between the Boltzmann distribution and its parametric approximation, and a regularizer term. This objective function is then minimized using an alternating Bregman proximal algorithm, relying on numerical integration to adapt the temperature.

In this work, we propose a novel SA algorithm that leverages the core ideas of [12], with a new choice of divergence, regularizer, and optimization procedure. These novel features lead to an improved algorithm which (i) is more efficient, (ii) exploits available information about the global minimum of the objective, (iii) admits theoretical guarantees, and (iv) shows good empirical performance. Specifically, we propose to use a Rényi divergence instead of the KL divergence, and we introduce a novel and more principled regularizer. The parametric distributions are chosen in an exponential family. We propose an alternating scheme involving Bregman proximal gradient steps. This allows to use the geometry of the considered families [13]. We show in two benchmark cost functions that the resulting method exhibits good performance compared to existing black-box algorithms.

T.G. and E.C. acknowledge support from the ERC Starting Grant MAJORIS ERC-2019-STG-850925. The work of V. E. is supported by the *Agence Nationale de la Recherche* of France under PISCES (ANR-17-CE40-0031-01), the Leverhulme Research Fellowship (RF-2021-593), and by ARL/ARO under grant W911NF-22-1-0235.

The rest of the paper is organized as follows. Background is provided in Section 2. We detail our method in Section 3. Numerical experiments are presented in Section 4, before concluding in Section 5.

2. PRELIMINARIES

2.1. Optimization problem and notation

We consider the following optimization problem

$$\text{Find } x_* \in \mathcal{X} \text{ s.t. } f(x_*) = f_* = \min_{x \in \mathcal{X}} f(x). \quad (1)$$

Here above, we consider a space $\mathcal{X} \subset \mathbb{R}^d$ and a cost function $f : \mathcal{X} \rightarrow \mathbb{R}$. In the following, \mathcal{H} denotes a finite-dimensional Hilbert space with scalar product $\langle \cdot, \cdot \rangle$. $\mathcal{B}(\mathcal{X})$ denotes the Borel algebra of \mathcal{X} and given a measure ν on $(\mathcal{X}, \mathcal{B}(\mathcal{X}))$, we denote by $\mathcal{P}(\mathcal{X}, \nu)$ the set of probability measures on $(\mathcal{X}, \mathcal{B}(\mathcal{X}))$ which admits a density with respect to ν . Given a measurable function h and $p \in \mathcal{P}(\mathcal{X}, \nu)$, we write $p(h) = \int h(x)p(x)\nu(dx)$. For a function h , we denote its conjugate function by h^* .

2.2. Exponential and Boltzmann families

Definition 1. *The exponential family with sufficient statistics Γ is the family $\mathcal{Q} = \{q_\theta \in \mathcal{P}(\mathcal{X}, \nu), \theta \in \Theta\}$ such that*

$$q_\theta(x) = \exp(\langle \theta, \Gamma(x) \rangle - A(\theta)), \quad \forall x \in \mathcal{X}, \quad (2)$$

with A being the log-partition function, such that $\Theta = \text{dom } A \subset \mathcal{H}$, and which reads:

$$A(\theta) = \log \left(\int \exp(\langle \theta, \Gamma(x) \rangle) \nu(dx) \right), \quad \forall \theta \in \Theta. \quad (3)$$

Example 1. *Gaussian distributions with mean μ and covariance Σ form an exponential family [14], with ν being the Lebesgue measure and parameters $\theta = (\Sigma^{-1}\mu, -\frac{1}{2}\Sigma^{-1})^\top$.*

The Boltzmann distributions, denoted by π_β for $\beta > 0$, are central in SA schemes. π_β gives the probability of a state $x \in \mathcal{X}$ depending on $f(x)$ and the inverse temperature β .

Definition 2. *The family \mathcal{B} of Boltzmann distributions associated with f is the set of densities defined for $\beta > 0$ by*

$$\pi_\beta(x) = \exp(-\beta f(x) - B(\beta)), \quad \forall x \in \mathcal{X}, \quad (4)$$

where B is the log-partition function of π_β :

$$B(\beta) = \log \left(\int \exp(-\beta f(x)) \nu(dx) \right), \quad \forall \beta > 0. \quad (5)$$

Remark 1. *The Boltzmann family \mathcal{B} forms an exponential family, with sufficient statistics $-f$ and parameter β .*

The mass of the Boltzmann distributions concentrates on the global minimizers of f as β increases. Indeed, $\lim_{\beta \rightarrow +\infty} \pi_\beta(f) = f_*$ [6, Eq. (5.13)]. Therefore, points sampled from π_β with high β would be close from the global minimizer. However, sampling such points is difficult.

2.3. Rényi and KL divergences

Rényi and KL divergences are widely used as discrepancy measures between probability distributions [15, 16].

Definition 3. *The Rényi divergence with parameter $\alpha \in (0, 1) \cup (1, +\infty)$ and the KL divergence between p_1 and p_2 in $\mathcal{P}(\mathcal{X}, \nu)$ are respectively defined by*

$$RD_\alpha(p_1, p_2) = \frac{1}{\alpha - 1} \log \left(\int p_1(x)^\alpha p_2(x)^{1-\alpha} \nu(dx) \right),$$

$$KL(p_1, p_2) = \int \log \left(\frac{p_1(x)}{p_2(x)} \right) p_1(x) \nu(dx).$$

If one of the above is not well-defined, its value is $+\infty$.

The KL divergence is a limiting case of Rényi divergence [17], since $\lim_{\alpha \rightarrow 1, \alpha \leq 1} RD_\alpha(p_1, p_2) = KL(p_1, p_2)$.

3. PROPOSED APPROACH

3.1. A Rényi divergence minimization problem

SA methods create a sequence of Boltzmann distributions with increasing inverse temperature and approximate them with tractable distributions. In this work, we approximate the Boltzmann distributions $\pi_\beta \in \mathcal{B}$ by parametric distributions q_θ from an exponential family \mathcal{Q} . Therefore, we search for parameters β, θ such that q_θ and π_β are close and β is sufficiently large. To achieve this goal, we consider the minimization of the following function:

$$J_\alpha(\beta, \theta) = d_\alpha(\beta, \theta) + r(\beta), \quad \forall \beta > 0, \theta \in \Theta, \quad (6)$$

where $\alpha \in (0, 1)$, $d_\alpha(\beta, \theta) = RD_\alpha(\pi_\beta, q_\theta)$. We furthermore introduce the regularizer

$$r(\beta) = B(\beta) + \beta f_\epsilon, \quad \forall \beta > 0, \quad (7)$$

where $f_\epsilon = f_* + \epsilon$ for any small value $\epsilon > 0$. r is convex and differentiable on \mathbb{R}_{++} . It is lower bounded as it is equal, up to an additive constant, to $KL(\pi_{\beta_\epsilon}, \pi_\beta)$ where $\pi_{\beta_\epsilon}(f) = f_\epsilon$. It promotes the concentration of π_β around the global minimizers of f since $\nabla r(\beta) = 0$ if and only if $\pi_\beta(f) = f_\epsilon$.

Solving Problem (6) yields a distribution π_β concentrated around the global minimizers of f , along with its approximation q_θ from which samples are available. In the following, we show how to construct in an adaptive fashion a sequence $\{(\beta_k, \theta_k)\}_{k \in \mathbb{N}}$ converging to such a solution.

3.2. Proposed algorithm

We propose to minimize J_α by using an *alternating Bregman proximal gradient algorithm* where the Bregman divergences are induced by the log-partition functions A and B . These are equivalent to the KL divergences [13]:

$$d_A(\theta, \theta') = KL(q_{\theta'}, q_\theta), \quad \forall \theta, \theta' \in \Theta,$$

$$d_B(\beta, \beta') = KL(\pi_{\beta'}, \pi_\beta), \quad \forall \beta, \beta' > 0.$$

Such choices allow to decouple the updates of β and θ , so that each update reduces to a variational inference update. Indeed, in this framework, Boltzmann distributions are approached by a variational approximation from \mathcal{Q} , and proposals from \mathcal{Q} are also approximated by a Boltzmann distribution. Our choices also allow the use of the underlying geometries of \mathcal{Q} and \mathcal{B} , which often leads to more efficient schemes [18, 19]. Finally, the available studies of [20, 21, 22, 23] give our method a solid theoretical background.

We are now ready to present our approach, whose iteration $k \in \mathbb{N}$ consists in the following three steps with step-sizes $\eta_k, \tau_k \in (0, 1]$, and $\tilde{\eta}_k = \frac{1-\alpha}{\alpha}\eta_k$ with $\alpha \in (0, 1)$. First, we perform a Bregman gradient descent step on $d_\alpha(\cdot, \theta_{k-1})$ within the Bregman divergence d_B with step-size $\tilde{\eta}_k$:

$$\beta_{k-\frac{1}{2}} = \nabla B^*(\nabla B(\beta_{k-1}) - \tilde{\eta}_k \nabla_\beta d_\alpha(\beta_{k-1}, \theta_{k-1})). \quad (8)$$

Second, we perform a Bregman proximal step on r within the divergence d_B and with step-size $\tilde{\eta}_k$:

$$\beta_k = \arg \min_{\beta > 0} r(\beta) + \frac{1}{\tilde{\eta}_k} d_B(\beta, \beta_{k-\frac{1}{2}}). \quad (9)$$

Third, we perform a Bregman gradient descent step on $d_\alpha(\beta_k, \cdot)$ within the divergence d_A and with step-size τ_k :

$$\theta_k = \nabla A^*(\nabla A(\theta_{k-1}) - \tau_k \nabla_\theta d_\alpha(\beta_k, \theta_{k-1})). \quad (10)$$

We can compute for any $\beta > 0, \theta \in \Theta$,

$$\nabla_\beta d_\alpha(\beta, \theta) = \frac{\alpha}{1-\alpha} (\pi_\beta(-f) - \pi_{\beta, \theta}^{(\alpha)}(-f)), \quad (11)$$

$$\nabla_\theta d_\alpha(\beta, \theta) = q_\theta(\Gamma) - \pi_{\beta, \theta}^{(\alpha)}(\Gamma), \quad (12)$$

where $\pi_{\beta, \theta}^{(\alpha)} \in \mathcal{P}(\mathcal{X}, \nu)$ and has a density defined by

$$\pi_{\beta, \theta}^{(\alpha)}(x) \propto \left(\frac{\exp(-f(x))^\beta}{q_\theta(x)} \right)^\alpha q_\theta(x), \quad \forall x \in \mathcal{X}. \quad (13)$$

Then, using results from [13, 22, 14], we can show that Eq. (8)-(10) are equivalent for every $k \in \mathbb{N}$ to

$$\pi_{\beta_{k-\frac{1}{2}}} (f) = (1 - \eta_k) \pi_{\beta_{k-1}}(f) + \eta_k \pi_{\beta_{k-1}, \theta_{k-1}}^{(\alpha)}(f), \quad (14)$$

$$\pi_{\beta_k}(f) = \frac{1}{1 + \tilde{\eta}_k} \pi_{\beta_{k-\frac{1}{2}}}(f) + \frac{\tilde{\eta}_k}{1 + \tilde{\eta}_k} f_\epsilon, \quad (15)$$

$$q_{\theta_k}(\Gamma) = (1 - \tau_k) q_{\theta_{k-1}}(\Gamma) + \tau_k \pi_{\beta_k, \theta_{k-1}}^{(\alpha)}(\Gamma). \quad (16)$$

Let us now state our convergence result for the proposed iterative scheme, using denominations from [14].

Proposition 1. *Let $\epsilon > 0$. Assume that \mathcal{B} and \mathcal{Q} are regular and full, that $\tau_k, \eta_k \in (0, 1]$ for all $k \in \mathbb{N}$, and that the updates are well-defined. Then, the sequence $\{J_\alpha(\beta_k, \theta_k)\}_{k \in \mathbb{N}}$ is decreasing. If further, $\mathcal{B} \subset \mathcal{Q}$, then a fixed point of the algorithm (β_*, θ_*) is such that $\pi_{\beta_*}(f) = q_{\theta_*}(f) = f_\epsilon$.*

Proposition 1 establishes a monotonic decrease of the objective J_α . The second part of the result applies in particular when f is quadratic and \mathcal{Q} is the Gaussian family.

3.3. Our Rényi-based adaptive SA (RASA) scheme

We now discuss the practical implementation of our algorithm for black-box optimization, meaning that only evaluations of f are available. We propose to approximate $\pi_{\beta, \theta}^{(\alpha)}$ and π_β using adaptive importance sampling [24]. To this end, consider $\theta \in \Theta$ and denote by X_θ a set of N samples from q_θ . We define for $\beta > 0, \alpha > 0$ the empirical distributions

$$\pi_{X_\theta}^{(\beta, \alpha)} \propto \sum_{x \in X_\theta} \left(\frac{\exp(-f(x))^\beta}{q_\theta(x)} \right)^\alpha \delta_x. \quad (17)$$

The distribution above approximates $\pi_{\beta, \theta}^{(\alpha)}$, while $\pi_{X_\theta}^{(\beta, 1)}$ approximates π_β . This yields Algorithm 1, which approximates the idealized iterative algorithm described by Eq. (14)-(16), with a precision depending on the sample size N .

Algorithm 1: RASA algorithm

input: Choose $\beta_0 > 0, \theta_0 \in \Theta, N \in \mathbb{N}$, and $\{\tau_k, \eta_k\}_{k \in \mathbb{N}}$ in $(0, 1]$. Set $\tilde{\eta}_k = \frac{1-\alpha}{\alpha}\eta_k$ for every $k \in \mathbb{N}$.

for $k = 1, 2, \dots$ **do**

1 Sample N points $X_{\theta_{k-1}} = \{x_n\}_{n=1}^N$ from $q_{\theta_{k-1}}$.

2 Set $f_k = \min\{f(x), x \in X_{\theta_l}, 0 \leq l \leq k-1\}$.

3 If $k = 1$, set $\pi_{X_{\theta_{k-2}}}^{(\beta_{k-1}, 1)}(f) = \pi_{X_{\theta_{k-1}}}^{(\beta_{k-1}, 1)}(f)$.

4 Compute the following empirical moment:

$$\pi_{\beta_{k-\frac{1}{2}}} (f) = (1 - \eta_k) \pi_{X_{\theta_{k-2}}}^{(\beta_{k-1}, 1)}(f) + \eta_k \pi_{X_{\theta_{k-1}}}^{(\beta_{k-1}, \alpha)}(f).$$

5 Set β_k such that

$$\pi_{X_{\theta_{k-1}}}^{(\beta_k, 1)}(f) = \frac{1}{1 + \tilde{\eta}_k} \pi_{\beta_{k-\frac{1}{2}}}(f) + \frac{\tilde{\eta}_k}{1 + \tilde{\eta}_k} f_k.$$

6 Adapt the proposal by

$$q_{\theta_k}(\Gamma) = (1 - \tau_k) q_{\theta_{k-1}}(\Gamma) + \tau_k \pi_{X_{\theta_{k-1}}}^{(\beta_k, \alpha)}(\Gamma).$$

Note that $\beta \mapsto \pi_{X_{\theta_{k-1}}}^{(\beta, 1)}(f)$ is a decreasing continuous function. Thus, β_k can be found by a dichotomy search.

3.4. Discussion

In [12], KL divergences and quadratic regularizers are used. In contrast, our scheme relies on Rényi divergences benefiting from a useful symmetry property [17, Proposition 2] and a novel regularizer, leading to simpler β updates.

The adaptation step of θ of the MARS algorithm [11] is recovered in Algorithm 1 with $\alpha = 1$. However, a fixed cooling schedule was used in [11], while ours is fully adaptive.

Non-linear importance weights arise in Algorithm 1 from the densities $\pi_{\beta, \theta}^{(\alpha)}$ and the approximations $\pi_{X_\theta}^{(\beta, \alpha)}$. Such weights are known to prevent weight degeneracy issues [25].

4. NUMERICAL SIMULATIONS

4.1. Test problems and compared algorithms

We use two benchmarks with minimum value f_* reached at a unique x_* . Both are defined for any $x \in \mathbb{R}^d$, where $d \in \mathbb{N}$. We consider a banana-shaped Rosenbrock function and a multimodal Rastrigin function, defined respectively by

$$f(x - x_*) = \sum_{i=1}^{d-1} (10(x_{i+1} - 1 - (x_i - 1)^2)^2 + x_i^2) + f_*,$$

$$f(x - x_*) = 4d + \sum_{i=1}^d (0.4x_i^2 - 4 \cos(2\pi x_i)) + f_*.$$

The minimum values f_* are sampled uniformly in $[-1, 1]$, and the two functions are translated so that they reach the value f_* at x_* uniformly sampled in $[-1, 1]^d$.

We compare our method with other schemes updating approximating densities using sampling. Namely, we consider two SA schemes, the APSA scheme [12] and the MARS algorithm [11], as well as the cross-entropy (CE) algorithm [5].

All the algorithms are run with step-size $\tau_k = \frac{0.5}{k+1}$ for the parameters updates, and step-size $\eta_k = 0.9$ for the eventual temperature updates. The CE algorithm is run using the proportion $\rho = 0.5$. We simulated $N = 100$ samples at each iteration $k \in \mathbb{N}$ from Gaussian proposals with means μ_k . The algorithms are initialized with mean μ_0 uniformly sampled in $[-5, 5]^d$, covariance $\Sigma_0 = 10I$, and initial temperature $\beta_0 = 0.1$. At each iteration, the dichotomy search in RASA is performed in the interval $[0.1\beta_{k-1}, 1.5\beta_{k-1}]$.

4.2. Assessing our adaptive cooling schedule

We first compare the cooling schedules of the APSA, MARS, and RASA algorithms in dimension $d = 2$. Indeed, the APSA scheme is implemented with numerical integration steps which become intractable in higher dimensions.




Fig. 1. Plots of the quantities β_k averaged over 500 runs for the Rastrigin (left) and Rosenbrock (right) functions in dimension $d = 2$.

Figure 1 shows that the three methods have very different cooling schedules. As observed in [12], the cooling schedule of the APSA algorithm reaches a high stationary value very fast. The MARS algorithm follows a logarithmic cooling schedule, which is the slowest of the three. In contrast, the

cooling schedule of the APSA algorithm increases in a convex fashion, much faster than the logarithmic cooling schedule and without the stagnation of the APSA schedule.

4.3. Optimization performance in high dimension

We now compare our RASA algorithm with the MARS and CE algorithms for $d = 50$, where the task is to minimize the benchmark objective f . Note that the APSA method is intractable in this high-dimensional setting.




Fig. 2. Plots of the quantity $f(\mu_k) - f_*$ averaged over 500 runs for the Rastrigin (left) and Rosenbrock (right) functions in dimension $d = 50$.

We can see in Fig. 2 that RASA outperforms the MARS and CE algorithms on both benchmarks. This demonstrates the positive impact on optimization performance of our adaptive cooling schedule. Note that contrary to the CE algorithm, which only uses a ranking of the samples, the RASA and MARS algorithms use each sample's cost function value.

We also note that the lowest value of α yields the best performance. Indeed, transforming the weights as in Eq. (17) reduces weight degeneracy [25], which may lower the approximation error arising from sampling with a limited number of points in high dimension. Also, low values of α lead to approximations q_θ fitting the mode of π_β [15]. This shows the interest of the additional parameter α , coming from the use of a Rényi divergence instead of a KL divergence as in [11].

5. CONCLUSION

In this work, we presented a novel SA algorithm for black-box optimization. It approaches Boltzmann distributions by variational approximations and treats the Boltzmann distributions as variational approximations themselves. This creates an adaptive cooling schedule while avoiding some drawbacks of the commonly used Markov chains schemes. This procedure is interpreted as an alternating Bregman proximal-gradient algorithm, allowing to gain novel theoretical insights. Numerical experiments illustrate the good behavior of the resulting method. This work also opens potential research tracks. On a theoretical level, the links between our practical implementation and its exact counterpart remain to be established precisely. The use of more complex approximating distributions, such as mixtures, can also be investigated, in particular in relation to multi-modal objectives.

6. REFERENCES

- [1] S. Watanabe and J. Le Roux, “Black box optimization for automatic speech recognition,” in *Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP 2014)*, Firenze, Italy, 4-9 May 2014, pp. 3256–3260.
- [2] M. A. Abramson, T. J. Asaki, J. E. Dennis, K. R. O’Reilly, and R. L. Pingel, “Quantitative object reconstruction using Abel transform x-ray tomography and mixed variable optimization,” *SIAM Journal on Imaging Sciences*, vol. 1, no. 3, pp. 322–342, 2008.
- [3] L. M. Rios and N. V. Sahinidis, “Derivative-free optimization: a review of algorithms and comparison of software implementations,” *Journal of Global Optimization*, vol. 56, no. 3, pp. 1247–1293, 2013.
- [4] D. Wiestra, T. Schaul, T. Glasmachers, Y. Sun, J. Peters, and J. Schmidhubler, “Natural evolution strategies,” *Journal of Machine Learning Research*, vol. 15, no. 27, pp. 949–980, 2014.
- [5] D. Kroese, S. Porotsky, and R. Rubinstein, “The cross-entropy method for continuous multi-extremal optimization,” *Methodology and Computing in Applied Probability*, vol. 8, pp. 383–407, 2006.
- [6] H. Haario and E. Saksman, “Simulated annealing process in general state space,” *Advances in Applied Probability*, vol. 23, no. 4, pp. 866–893, 1991.
- [7] R. A. Holley, S. Kusuoka, and D. W. Stroock, “Asymptotics of the spectral gap with application to the theory of simulated annealing,” *Journal of Functional Analysis*, vol. 83, no. 2, pp. 333–347, 1989.
- [8] O. Molvalioglu and Z. Zabinsky, “Meta-control of an interacting-particle algorithm for global optimization,” *Nonlinear Analysis : Hybrid Systems*, vol. 4, pp. 659–671, 2010.
- [9] A. Beskos, A. Jasra, and A. Thiery, “On the convergence of adaptive sequential Monte Carlo methods,” *The Annals of Applied Probability*, vol. 26, no. 2, pp. 1111–1146, 2017.
- [10] E. Angelino, M. J. Johnson, and R. P. Adams, “Patterns of scalable Bayesian inference,” *Foundations and Trends® in Machine Learning*, vol. 9, no. 2-3, pp. 119–247, 2016.
- [11] J. Hu and P. Hu, “Annealing adaptive search, cross-entropy, and stochastic approximation in global optimization,” *Naval Research Logistics*, vol. 58, no. 5, pp. 457–477, 2011.
- [12] T. Guilmeau, E. Chouzenoux, and V. Elvira, “Proximal-based adaptive simulated annealing for global optimization,” in *Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP 2022)*, Singapore, Singapore, 22-27 May 2022, pp. 5453–5457.
- [13] F. Nielsen and R. Nock, “Entropies and cross-entropies of exponential families,” in *Proceedings of the IEEE International Conference on Image Processing (ICIP 2010)*, Hong Kong, China, 26 - 29 Sep. 2010, pp. 3621–3624.
- [14] O. Barndorff-Nielsen, *Information and Exponential Families in Statistical Theory*, John Wiley & Sons, Ltd, 2014.
- [15] Y. Li and R. Turner, “Rényi divergence variational inference,” in *Advances in Neural Information Processing Systems 29 (NIPS 2016)*, Barcelona, Spain, 5-10 December 2016, vol. 29, pp. 1073–1081.
- [16] K. Daudel, R. Douc, and F. Portier, “Infinite-dimensional gradient-based descent for alpha-divergence minimisation,” *The Annals of Statistics*, vol. 49, no. 4, pp. 2250–2270, 2021.
- [17] T. van Erven and P. Harremoës, “Rényi divergence and Kullback-Leibler divergence,” *IEEE Transactions on Information Theory*, vol. 60, no. 7, pp. 3797–3820, 2014.
- [18] M. D. Hoffman, D. M. Blei, C. Wang, and J. Paisley, “Stochastic variational inference,” *Journal of Machine Learning Research*, vol. 14, no. 4, pp. 1303–1347, 2013.
- [19] G. Raskutti and S. Mukherjee, “The information geometry of mirror descent,” *IEEE Transactions on Information Theory*, vol. 61, no. 3, pp. 1451–1457, 2015.
- [20] H. Bauschke, P. Combettes, and D. Noll, “Joint minimization with alternating Bregman proximity operators,” *Pacific Journal of Optimization*, vol. 2, 2006.
- [21] M. El Gheche, G. Chierchia, and J.-C. Pesquet, “Proximity operators of discrete information divergences,” *IEEE Transactions on Information Theory*, vol. 64, no. 2, pp. 1092–1104, 2018.
- [22] M. Teboulle, “A simplified view of first order methods for optimization,” *Mathematical Programming*, vol. 170, no. 1, pp. 67–96, 2018.
- [23] M. Ahookhosh, H. Le Ti Khanh, N. Gillis, and P. Patrinos, “Multi-block Bregman proximal alternating linearized minimization and its application to orthogonal nonnegative matrix factorization,” *Computational Optimization and Applications*, vol. 79, no. 3, pp. 681–715, 2021.
- [24] M. F. Bugallo, V. Elvira, L. Martino, D. Luengo, J. Míguez, and P. M. Djuric, “Adaptive importance sampling: The past, the present, and the future,” *IEEE Signal Processing Magazine*, vol. 34, no. 4, pp. 60–79, 2017.
- [25] E. Koblents and J. Míguez, “A population Monte Carlo scheme with transformed weights and its application to stochastic kinetic models,” *Statistics and Computing*, vol. 25, no. 2, pp. 407–425, 2013.