



*Research article***Relative entropy minimization over Hilbert spaces via Robbins-Monro****Gideon Simpson^{1*} and Daniel Watkins²**¹ Department of Mathematics, Drexel University, Philadelphia, PA 19104, USA² Department of Ocean, Earth, and Atmospheric Sciences, Oregon State University, Corvallis, OR 97331, USA*** Correspondence:** Email: grs53@drexel.edu; Tel: +12158952668; Fax: +12158951582.

Abstract: One way of getting insight into non-Gaussian measures is to first obtain good Gaussian approximations. These best fit Gaussians can then provide a sense of the mean and variance of the distribution of interest. They can also be used to accelerate sampling algorithms. This begs the question of how one should measure optimality, and how to then obtain this optimal approximation. Here, we consider the problem of minimizing the distance between a family of Gaussians and the target measure with respect to relative entropy, or Kullback-Leibler divergence. As we are interested in applications in the infinite dimensional setting, it is desirable to have convergent algorithms that are well posed on abstract Hilbert spaces. We examine this minimization problem by seeking roots of the first variation of relative entropy, taken with respect to the mean of the Gaussian, leaving the covariance fixed. We prove the convergence of Robbins-Monro type root finding algorithms in this context, highlighting the assumptions necessary for convergence to relative entropy minimizers. Numerical examples are included to illustrate the algorithms.

Keywords: Robbins-Monro; relative entropy; Hilbert space**Mathematics Subject Classification:** 65K10, 62L20, 60G15, 65C05

1. Introduction

In [13–16, 19], it was proposed that insight into a probability distribution, μ , posed on a Hilbert space, $\mathcal{H}$, could be obtained by finding a best fit Gaussian approximation, ν . This notion of best, or optimal, was with respect to the relative entropy, or Kullback-Leibler divergence:

$$\mathcal{R}(\nu||\mu) = \begin{cases} \mathbb{E}^{\nu} \left[\log \frac{d\nu}{d\mu} \right], & \nu \ll \mu, \\ +\infty, & \text{otherwise.} \end{cases} \quad (1.1)$$

Having a Gaussian approximation provides qualitative insight into μ , as it provides a concrete notion of the mean and variance of the distribution. Additionally, this optimized distribution can be used in algorithms, such as random walk Metropolis, as a preconditioned proposal distribution to improve performance. Such a strategy can benefit a number of applications, including path space sampling for molecular dynamics and parameter estimation in statistical inverse problems.

Observe that in the definition of $\mathcal{R}$, (1.1), there is an asymmetry in the arguments. Were we to work with $\mathcal{R}(\mu||\nu)$, our optimal Gaussian would capture the first and second moments of μ , and in some applications this is desirable. However, for a multimodal problem (consider a distribution with two well separated modes), this would be inadequate; our form attempts to match individual modes of the distribution by a Gaussian. For a recent review of the $\mathcal{R}(\nu||\mu)$ problem, see [4], where it is remarked that this choice of arguments is likely to underestimate the dispersion of the distribution of interest, μ . The other ordering of arguments has been explored, in the finite dimensional case, in [2, 3, 10, 18].

To be of computational use, it is necessary to have an algorithm that will converge to this optimal distribution. In [15], this was accomplished by first expressing $\nu = N(m, C(p))$, where m is the mean and p is a parameter inducing a well defined covariance operator, and then solving the problem,

$$(m, p) \in \operatorname{argmin} \mathcal{R}(N(m, C(p))||\mu), \quad (1.2)$$

over an admissible set. The optimization step itself was done using the Robbins-Monro algorithm (RM), [17], by seeking a root of the first variation of the relative entropy. While the numerical results of [15] were satisfactory, being consistent with theoretical expectations, no rigorous justification for the application of RM to the examples was given.

In this work, we emphasize the study and application of RM to potentially infinite dimensional problems. Indeed, following the framework of [15, 16], we assume that μ is posed on the Borel σ -algebra of a separable Hilbert space $(\mathcal{H}, \langle \bullet, \bullet \rangle, \|\bullet\|)$. For simplicity, we will leave the covariance operator C fixed, and only optimize over the mean, m . Even in this case, we are seeking $m \in \mathcal{H}$, a potentially infinite-dimensional space.

1.1. Robbins-Monro

Given the objective function $f : \mathcal{H} \rightarrow \mathcal{H}$, assume that it has a root, x_* . In our application to relative entropy, f will be its first variation. Further, we assume that we can only observe a noisy version of f , $F : \mathcal{H} \times \mathcal{X} \rightarrow \mathcal{H}$, such that for all $x \in \mathcal{H}$,

$$f(x) = \mathbb{E}[F(x, Z)] = \int_{\mathcal{X}} F(x, z) \mu_Z(dz), \quad (1.3)$$

where μ_Z is the distribution associated with the random variable (r.v.) Z , taking values in the auxiliary space $\mathcal{X}$. The naive Robbins-Monro algorithm is given by

$$X_{n+1} = X_n - a_{n+1} F(X_n, Z_{n+1}), \quad (1.4)$$

where $Z_n \sim \mu_Z$, are independent and identically distributed (i.i.d.), and $a_n > 0$ is a carefully chosen sequence. Subject to assumptions on f , F , and the distribution μ_Z , it is known that X_n will converge to x_* almost surely (a.s.), in finite dimensions, [5, 6, 17]. Often, one needs to assume that f grows at most linearly,

$$\|f(x)\| \leq c_0 + c_1 \|x\|, \quad (1.5)$$

in order to apply the results in the aforementioned papers. The analysis in the finite dimensional case has been refined tremendously over the years, including an analysis based on continuous dynamical systems. We refer the reader to the books [1, 8, 11] and references therein.

1.2. Trust regions and truncations

As noted, much of the analysis requires the regression function f to have, at most, linear growth. Alternatively, an *a priori* assumption is sometimes made that the entire sequence generated by (1.4) stays in a bounded set. Both assumptions are limiting, though, in practice, one may find that the algorithms converge.

One way of overcoming these assumptions, while still ensuring convergence, is to introduce trust regions that the sequence $\{X_n\}$ is permitted to explore, along with a “truncation” which enforces the constraint. Such truncations distort (1.4) into

$$X_{n+1} = X_n - a_{n+1}F(X_n, Z_{n+1}) + a_{n+1}P_{n+1}, \quad (1.6)$$

where P_{n+1} is the projection keeping the sequence $\{X_n\}$ within the trust region. Projection algorithms are also discussed in [1, 8, 11].

We consider RM on a possibly infinite dimensional separable Hilbert space. This is of particular interest as, in the context of relative entropy optimization, we may be seeking a distribution in a Sobolev space associated with a PDE model. A general analysis of RM with truncations in Hilbert spaces can be found in [20]. The main purpose of this work is to adapt the analysis of [12] to the Hilbert space setting for two versions of the truncated problem. The motivation for this is that the analysis of [12] is quite straightforward, and it is instructive to see how it can be easily adapted to the infinite dimensional setting. The key modification in the proof is that results for Banach space valued martingales must be invoked. We also adapt the results to a version of the algorithm where there is prior knowledge on the location of the root. With these results in hand, we can then verify that the relative entropy minimization problem can be solved using RM.

1.2.1. Fixed trust regions

In some problems, one may have *a priori* information on the root. For instance, we may know that $x_\star \in U_1$, some open bounded set. In this version of the truncated algorithm, we have two open bounded sets, $U_0 \subsetneq U_1$, and $x_\star \in U_1$. Let $\sigma_0 = 0$ and $X_0 \in U_0$ be given, then (1.6) can be formulated as

$$\tilde{X}_{n+1} = X_n - a_{n+1}F(X_n, Z_{n+1}) \quad (1.7a)$$

$$X_{n+1} = \begin{cases} \tilde{X}_{n+1} & \tilde{X}_{n+1} \in U_1 \\ X_0^{(\sigma_n)} & \tilde{X}_{n+1} \notin U_1 \end{cases} \quad (1.7b)$$

$$\sigma_{n+1} = \begin{cases} \sigma_n & \tilde{X}_{n+1} \in U_1 \\ \sigma_n + 1 & \tilde{X}_{n+1} \notin U_1 \end{cases} \quad (1.7c)$$

We interpret $\tilde{X}_{n+1}$ as the proposed move, which is either accepted or rejected depending on whether or not it will remain in the trust region. If it is rejected, the algorithm restarts at $X_0^{(\sigma_n)} \in U_0$. The restart points, $\{X_0^{(\sigma_n)}\}$, may be random, or it may be that $X_0^{(\sigma_n)} = X_0$ is fixed. The essential property is that the

algorithm will restart in the interior of the trust region, away from its boundary. The r.v. σ_n counts the number of times a truncation has occurred. Algorithm (1.7) can now be expressed as

$$\begin{aligned} X_{n+1} &= X_n - a_{n+1}F(X_n, Z_n + 1) + P_{n+1} \\ P_{n+1} &= \{X_0^{(\sigma_n)} - \tilde{X}_{n+1}\}1_{\tilde{X}_{n+1} \notin U_1}. \end{aligned} \quad (1.8)$$

1.2.2. Expanding trust regions

In the second version of truncated Robbins-Monro, define the sequence of open bounded sets, U_n such that:

$$U_0 \subsetneq U_1 \subsetneq U_2 \subsetneq \dots, \quad \cup_{n=0}^{\infty} U_n = \mathcal{H}. \quad (1.9)$$

Again, letting $X_0 \in U_0$, $\sigma_0 = 0$, the algorithm is

$$\tilde{X}_{n+1} = X_n - a_{n+1}F(X_n, Z_{n+1}) \quad (1.10a)$$

$$X_{n+1} = \begin{cases} \tilde{X}_{n+1} & \tilde{X}_{n+1} \in U_{\sigma_n} \\ X_0^{(\sigma_n)} & \tilde{X}_{n+1} \notin U_{\sigma_n} \end{cases} \quad (1.10b)$$

$$\sigma_{n+1} = \begin{cases} \sigma_n & \tilde{X}_{n+1} \in U_{\sigma_n} \\ \sigma_n + 1 & \tilde{X}_{n+1} \notin U_{\sigma_n} \end{cases} \quad (1.10c)$$

A consequence of this formulation is that $X_n \in U_{\sigma_n}$ for all n . As before, the restart points may be random or fixed, and they are in U_0 . This would appear superior to the fixed trust region algorithm, as it does not require knowledge of the sets. However, to guarantee convergence, global (in $\mathcal{H}$) assumptions on the regression function are required; see Assumption 2 below. (1.10) can be written with P_{n+1} as

$$\begin{aligned} X_{n+1} &= X_n - a_{n+1}F(X_n, Z_n + 1) + P_{n+1} \\ P_{n+1} &= \{X_0^{(\sigma_n)} - \tilde{X}_{n+1}\}1_{\tilde{X}_{n+1} \notin U_{\sigma_n}} \end{aligned} \quad (1.11)$$

1.3. Outline

In Section 2, we state sufficient assumptions for which we are able to prove convergence in both the fixed and expanding trust region problems, and we also establish some preliminary results. In Section 3, we focus on the relative entropy minimization problem, and identify what assumptions must hold for convergence to be guaranteed. Examples are then presented in Section 4, and we conclude with remarks in Section 5.

2. Convergence of Robbins-Monro

We first reformulate (1.8) and (1.11) in the more general form

$$X_{n+1} = \underbrace{X_n - a_{n+1}f(X_n) - a_{n+1}\delta M_{n+1}}_{=\tilde{X}_{n+1}} + a_{n+1}P_{n+1}, \quad (2.1)$$

where δM_{n+1} , the noise term, is

$$\begin{aligned} \delta M_{n+1} &= F(X_n, Z_{n+1}) - f(X_n) \\ &= F(X_n, Z_{n+1}) - \mathbb{E}[F(X_n, Z_{n+1}) \mid X_n]. \end{aligned} \quad (2.2)$$

A natural filtration for this problem is $\mathcal{F}_n = \sigma(X_0, Z_1, \dots, Z_n)$. X_n is $\mathcal{F}_n$ measurable and the noise term can be expressed in terms of the filtration as $\delta M_{n+1} = F(X_n, Z_{n+1}) - \mathbb{E}[F(X_n, Z_{n+1}) | \mathcal{F}_n]$.

We now state our main assumptions:

Assumption 1. f has a zero, $x_\star$. In the case of the fixed trust region problem, there exist $R_0 < R_1$ such that

$$U_0 \subseteq B_{R_0}(x_\star) \subset B_{R_1}(x_\star) \subseteq U_1.$$

In the case of the expanding trust region problem, the open sets are defined as $U_n = B_{r_n}(0)$ with

$$0 < r_0 < r_1 < r_2 < \dots < r_n \rightarrow \infty. \quad (2.3)$$

These sets clearly satisfy (1.9).

Assumption 2. For any $0 < a < A$, there exists $\delta > 0$:

$$\inf_{a \leq \|x - x_\star\| \leq A} \langle x - x_\star, f(x) \rangle \geq \delta.$$

In the case of the fixed truncation, this inequality is restricted to $x \in U_1$. This is akin to a convexity condition on a functional $\mathcal{F}$ with $f = D\mathcal{F}$.

Assumption 3. $x \mapsto \mathbb{E}[\|F(x, Z)\|^2]$ is bounded on bounded sets, with the restriction to U_1 in the case of fixed trust regions.

Assumption 4. $a_n > 0$, $\sum a_n = \infty$, and $\sum a_n^2 < \infty$

Theorem 2.1. Under the above assumptions, for the fixed trust region problem, $X_n \rightarrow x_\star$ a.s. and σ_n is a.s. finite.

Theorem 2.2. Under the above assumptions, for the expanding trust region problem, $X_n \rightarrow x_\star$ a.s. and σ_n is a.s. finite.

Note the distinction between the assumptions in the two algorithms. In the fixed truncation algorithm, Assumptions 2 and 3 need only hold in the set U_1 , while in the expanding truncation algorithm, they must hold in all of $\mathcal{H}$. While this would seem to be a weaker condition, it requires identification of the sets U_0 and U_1 for which the assumptions hold. Such sets may not be readily identifiable, as we will see in our examples.

We first need some additional information about f and the noise sequence δM_n .

Lemma 2.1. Under Assumption 3, f is bounded on U_1 , for the fixed trust region problem, and on arbitrary bounded sets, for the expanding trust region problem.

Proof. Trivially,

$$\|f(x)\| = \|\mathbb{E}[F(x, Z)]\| \leq \mathbb{E}[\|F(x, Z)\|] \leq \sqrt{\mathbb{E}[\|F(x, Z)\|^2]},$$

and the results follows from the assumption. $\square$

Proposition 2.1. For the fixed trust region problem, let

$$M_n = \sum_{i=1}^n a_i \delta M_i.$$

Alternatively, in the expanding trust region problem, for $r > 0$, let

$$M_n = \sum_{i=1}^n a_i \delta M_i 1_{\|X_{i-1} - x_\star\| \leq r}.$$

Under Assumptions 3 and 4, M_n is a martingale, converging in $\mathcal{H}$, a.s.

Proof. The following argument holds in both the fixed and expanding trust region problems, with appropriate modifications. We present the expanding trust region case. The proof is broken up into 3 steps:

1. Relying on Theorem 6 of [7] for Banach space valued martingales, it will be sufficient to show that M_n is a martingale, uniformly bounded in $L^1(\mathbb{P})$.
2. In the case of the expanding truncations,

$$\begin{aligned} \mathbb{E}[\|\delta M_i 1_{\|X_{i-1} - x_\star\| \leq r}\|^2] &\leq 2\mathbb{E}[\|F(X_{i-1}, Z_i) 1_{\|X_{i-1} - x_\star\| \leq r}\|^2] + 2\mathbb{E}[\|f(X_{i-1}) 1_{\|X_{i-1} - x_\star\| \leq r}\|^2] \\ &\leq 2 \sup_{\|x - x_\star\| \leq r} \mathbb{E}[\|F(x, Z)\|^2] + 2 \sup_{\|x - x_\star\| \leq r} \|f(x)\|^2 \end{aligned}$$

Since both of these terms are bounded, independently of i , by Assumption 3 and Lemma 2.1, this is finite.

3. Next, since $\{\delta M_i 1_{\|X_{i-1} - x_\star\| \leq r}\}$ is a martingale difference sequence, we can use the above estimate to obtain the uniform $L^2(\mathbb{P})$ bound,

$$\mathbb{E}[\|M_n\|^2] = \sum_{i=1}^n a_i^2 \mathbb{E}[\|\delta M_i 1_{\|X_{i-1} - x_\star\| \leq r}\|^2] \leq \sup_i \mathbb{E}[\|\delta M_i 1_{\|X_{i-1} - x_\star\| \leq r}\|^2] \sum_{i=1}^{\infty} a_i^2 < \infty$$

Uniform boundedness in L^2 , gives boundedness in L^1 , and this implies a.s. convergence in $\mathcal{H}$.

□

2.1. Finite truncations

In this section we prove results showing that only finitely many truncations will occur, in either the fixed or expanding trust region case. Recall that when a truncation occurs, the equivalent conditions hold: $P_{n+1} \neq 0$; $\sigma_{n+1} = \sigma_n + 1$; and $\tilde{X}_{n+1} \notin U_1$ in the fixed trust region algorithm, while $\tilde{X}_{n+1} \notin U_{\sigma_n}$ in the expanding trust region case.

Lemma 2.2. In the fixed trust region algorithm, if Assumptions 1, 2, 3, and 4 hold, then the number of truncations is a.s. finite; a.s., there exists N , such that for all $n \geq N$, $\sigma_n = \sigma_N$.

Proof. We break the proof up into 7 steps:

1. Pick ρ and ρ' such that

$$R_0 < \rho' < \rho < R_1 \quad (2.4)$$

Let $\bar{f} = \sup \|f(x)\|$, with the supremum over U_1 ; this bound exists by Lemma 2.1. Under Assumption 2, there exists $\delta > 0$ such that

$$\inf_{R_0/2 \leq \|x - x_\star\| \leq R_1} \langle x - x_\star, f(x) \rangle = \delta. \quad (2.5)$$

Having fixed $\rho, \rho', \bar{f}$, and δ , take $\epsilon > 0$ such that:

$$\epsilon < \min \left\{ \rho' - R_0, \frac{R_1 - \rho'}{2 + \bar{f}}, \frac{\rho' - R_0}{\bar{f}}, \frac{R_0}{2}, \frac{\delta}{2\bar{f}}, \frac{\delta}{\bar{f}^2}, \rho - \rho' \right\}. \quad (2.6)$$

Having fixed such an ϵ , by the assumptions of this lemma and Proposition 2.1, a.s., there exists n_ϵ such that for any $n, m \geq n_\epsilon$, both

$$\left\| \sum_{k=n}^m a_k \delta M_k \right\| \leq \epsilon, \quad a_n \leq \epsilon. \quad (2.7)$$

2. Define the auxiliary sequence

$$X'_n = X_n - \sum_{k=n+1}^{\infty} a_k \delta M_k. \quad (2.8)$$

Using (2.1), we can then write

$$X'_{n+1} = X'_n - a_{n+1} f(X_n) + a_{n+1} P_{n+1}. \quad (2.9)$$

By (2.7), for any $n \geq n_\epsilon$,

$$\|X'_n - X_n\| \leq \epsilon \quad (2.10)$$

3. We will show $X'_n \in B_{\rho'}(x_\star)$ for all n large enough. The significance of this is that if $n \geq n_\epsilon$, and $X'_n \in B_{\rho'}(x_\star)$, then no truncation occurs. Indeed, using (2.6)

$$\begin{aligned} \|\tilde{X}_{n+1} - x_\star\| &\leq \|X'_n - x_\star\| + \|X_n - X'_n\| + a_{n+1} \bar{f} + \|a_{n+1} \delta M_{n+1}\| \\ &< \rho' + \epsilon + \epsilon \bar{f} + \epsilon < R_1, \Rightarrow \tilde{X}_{n+1} \in U_1. \end{aligned} \quad (2.11)$$

Consequently, $P_{n+1} = 0$, $X_{n+1} = \tilde{X}_{n+1}$, and $\sigma_{n+1} = \sigma_n$. Thus, establishing $X'_n \in B_{\rho'}(x_\star)$ will yield the result.

4. Let

$$N = \inf\{n \geq n_\epsilon \mid \tilde{X}_{n+1} \notin U_1\} + 1 \quad (2.12)$$

This corresponds to the first truncation after n_ϵ . If the above set is empty, for that realization, no truncations occur after n_ϵ , and we are done. In such a case, we may take $N = n_\epsilon$ in the statement of the lemma.

5. We now prove by induction that in the case that (2.12) is finite, $X'_n \in B_{\rho'}(x_\star)$ for all $n \geq N$. First, note that $X_N \in B_{R_0}(x_\star) \subset B_\rho(x_\star)$. By (2.6) and (2.10),

$$\|X'_N - x_\star\| \leq \|X_N - x_\star\| + \|X'_N - X_N\| < R_0 + \epsilon < \rho', \Rightarrow X'_N \in B_{\rho'}(x_\star).$$

Next, assume $X'_N, X'_{N+1}, \dots, X'_n$ are all in $B_{\rho'}(x_\star)$. Using (2.11), we have that $P_{N+1} = \dots = P_{n+1} = 0$ and $\sigma_N = \dots = \sigma_n = \sigma_{n+1}$. Therefore,

$$\begin{aligned} \|X'_{n+1} - x_\star\|^2 &= \|X'_n - x_\star\|^2 - 2a_{n+1} \langle X'_n - x_\star, f(X_n) \rangle + a_{n+1}^2 \|f(X_n)\|^2 \\ &\leq \|X'_n - x_\star\|^2 - 2a_{n+1} \langle X'_n - x_\star, f(X_n) \rangle + a_{n+1} \epsilon \bar{f}^2 \end{aligned} \quad (2.13)$$

We now consider two cases of (2.13) to conclude $\|X'_{n+1} - x_\star\| < \rho'$.

6. In the first case, $\|X'_n - x_\star\| \leq R_0$. By Cauchy-Schwarz and (2.6)

$$\|X'_{n+1} - x_\star\|^2 < R_0^2 + 2\epsilon R_0 \bar{f} + \epsilon^2 \bar{f}^2 = (R_0 + \epsilon \bar{f})^2 < (\rho')^2.$$

In the second case, $R_0 < \|X'_n - x_\star\| < \rho'$. Dissecting the inner product term in (2.13) and using Assumption 2 and (2.10),

$$\begin{aligned} \langle X'_n - x_\star, f(X_n) \rangle &= \langle X_n - x_\star, f(X_n) \rangle + \langle X'_n - X_n, f(X_n) \rangle \\ &\geq \langle X_n - x_\star, f(X_n) \rangle - \bar{f} \epsilon \end{aligned} \quad (2.14)$$

Conditions (2.6) and (2.10) yield the following upper and lower bounds:

$$\begin{aligned} \|X_n - x_\star\| &\geq \|X'_n - x_\star\| - \|X'_n - X_n\| \geq R_0 - \epsilon > \frac{1}{2}R_0, \\ \|X_n - x_\star\| &\leq \|X'_n - x_\star\| + \|X'_n - X_n\| \leq \rho' + \epsilon < \rho < R_1. \end{aligned}$$

Therefore, (2.5) applies and $\langle X_n - x_\star, f(X_n) \rangle \geq \delta$. Using this in (2.14), and condition (2.6),

$$\langle X'_n - x_\star, f(X_n) \rangle \geq \delta - \bar{f} \epsilon > \frac{1}{2} \delta.$$

Substituting this last estimate back into (2.13), and using (2.6),

$$\|X'_{n+1} - x_\star\|^2 < (\rho')^2 - a_{n+1}(\delta - \epsilon \bar{f}^2) < (\rho')^2.$$

This completes the inductive step.

7. Since the auxiliary sequence remains in $B_{\rho'}(x_\star)$ for all $n \geq N > n_\epsilon$, (2.11) ensures $\tilde{X}_{n+1} \in B_{R_1}(x_\star)$, $P_{n+1} = 0$, and $\sigma_{n+1} = \sigma_N$, a.s.

□

To obtain a similar result for the expanding trust region problem, we first relate the finiteness of the number of truncations with the sequence persisting in a bounded set.

Lemma 2.3. *In the expanding trust region algorithm, if Assumptions 1, 3, and 4 hold, then the sequence remains in a set of the form $B_R(0)$ for some $R > 0$ if and only if the number of truncations is finite, a.s.*

Proof. We break this proof into 4 steps:

1. If the number of truncations is finite, then there exists N such that for all $n \geq N$, $\sigma_n = \sigma_N$. Consequently, the proposed moves are always accepted, and $X_n \in U_{\sigma_n} = U_{\sigma_N}$ for all $n \geq N$. Since $X_n \in U_{\sigma_n} \subset U_{\sigma_N}$ for $n < N$, $X_n \in U_{\sigma_N}$ for all n . By Assumption 3, $B_R(0) = B_{r_{\sigma_N}}(0) = U_{\sigma_N}$ is the desired set.
2. For the other direction, assume that there exists $R > 0$ such that $X_n \in B_R(0)$ for all n . Since the r_n in (2.3) tend to infinity, there exists N_1 , such that $R < R + 1 < r_{N_1}$. Hence, for all $n \geq N_1$,

$$B_R(0) \subset B_{R+1}(0) \subset U_n \quad (2.15)$$

Let $\bar{f} = \sup \|f(x)\|$, with the supremum over $B_R(0)$. Let $\tilde{R}$ be sufficiently large such that $B_{R+1}(0) \subset B_{\tilde{R}}(x_*)$. Lastly, using Proposition 2.1 and Assumption 4, a.s., there exists N_2 , such that for all $n \geq N_2$

$$\|a_n \delta M_n 1_{\|X_n - x_*\| \leq \tilde{R}}\| < \frac{1}{2}, \quad a_n < \frac{1}{2(1 + \bar{f})} \quad (2.16)$$

Since $X_n \in B_R(0) \subset B_{\tilde{R}}(x_*)$, the indicator function in (2.16) is always one, and $\|a_n \delta M_n\| < 1/2$.

3. Next, let

$$N = \inf\{n \geq 0 \mid \sigma_n \geq \max\{N_1, N_2\}\} \quad (2.17)$$

If the above set is empty, then $\sigma_n < \max\{N_1, N_2\}$ for all n , and the number of truncations is a.s. finite. In this case, the proof is complete.

4. If the set in (2.17) is not empty, then $N < \infty$. Take $n \geq N$. As $X_n \in B_R(0)$, and since $n \geq \sigma_n \geq \max\{N_1, N_2\}$, (2.16) applies. Therefore,

$$\begin{aligned} \|\tilde{X}_{n+1}\| &\leq \|X_n\| + \|\tilde{X}_{n+1} - X_n\| \\ &\leq \|X_n\| + a_{n+1} \|f(X_n)\| + \|a_{n+1} \delta M_{n+1}\| \\ &< R + \frac{1}{2} + \frac{1}{2} < R + 1. \end{aligned} \quad (2.18)$$

Thus, $\tilde{X}_{n+1} \in B_{R+1}(0) \subset U_{N_1}$, $\sigma_n \geq N_1$, and $U_{N_1} \subset U_{\sigma_n}$. Therefore, $\tilde{X}_{n+1} \in U_{\sigma_n}$. No truncation occurs, and $\sigma_n = \sigma_{n+1}$. Since this holds for all $n \geq N$, $\sigma_n = \sigma_N$, and the number of truncations is a.s. finite.

□

Next, we establish that, subject to an additional assumption, the sequence remains in a bounded set; the finiteness of the truncations is then a corollary.

Lemma 2.4. *In the expanding trust region algorithm, if Assumptions 1, 2, 3, and 4 hold, and for any $r > 0$, there a.s. exists $N < \infty$, such that for all $n \geq N$,*

$$P_{n+1} 1_{\|X_n - x_*\| \leq r} = 0,$$

then $\{X_n\}$ remains in a bounded open set, a.s.

Proof. We break this proof into 7 steps:

1. We begin by setting some constants for the rest of the proof. Fix $R > 0$ sufficiently large such that $B_R(x_\star) \supset U_0$. Next, let $\bar{f} = \sup \|f(x)\|$ with the supremum taken over $B_{R+2}(x_\star)$. Assumption 2 ensures there exists $\delta > 0$ such that

$$\inf_{R/2 \leq \|x - x_\star\| \leq R+2} \langle x - x_\star, f(x) \rangle = \delta. \quad (2.19)$$

Having fixed R , $\bar{f}$, and δ , take $\epsilon > 0$ such that:

$$\epsilon < \min \left\{ 1, \frac{1}{\bar{f}}, \frac{\delta}{2\bar{f}}, \frac{\delta}{\bar{f}^2}, \frac{R}{2} \right\}. \quad (2.20)$$

By the assumptions of this lemma and Proposition 2.1 there exists, a.s., $n_\epsilon \geq N$ such that for all $n \geq n_\epsilon$,

$$\left\| \sum_{i=n+1}^{\infty} a_i \delta M_i 1_{\|X_{i-1} - x_\star\| \leq R+2} \right\| \leq \epsilon, \quad (2.21a)$$

$$P_{n+1} 1_{\|X_n - x_\star\| \leq R+2} = 0, \quad (2.21b)$$

$$a_{n+1} \leq \epsilon \quad (2.21c)$$

2. Define the modified sequence for $n \geq n_\epsilon$ as

$$X'_n = X_n - \sum_{k=n+1}^{\infty} a_k \delta M_k 1_{\|X_{k-1} - x_\star\| \leq R+2}, \Rightarrow \|X'_n - X_n\| \leq \epsilon. \quad (2.22)$$

Using (2.1), we have the iteration

$$X'_{n+1} = X'_n - a_{n+1} \delta M_{n+1} 1_{\|X_n - x_\star\| > R+2} - a_{n+1} f(X_n) + a_{n+1} P_{n+1}. \quad (2.23)$$

3. Let

$$N = \inf \{n \geq n_\epsilon \mid \sigma_{n+1} \neq \sigma_n\} + 1, \quad (2.24)$$

the first time after n_ϵ that a truncation occurs.

If the above set is empty, no truncations occur after n_ϵ . In this case, $\sigma_n = \sigma_{n_\epsilon} \leq n_\epsilon < \infty$ for all $n \geq n_\epsilon$. Therefore, for all $n \geq n_\epsilon$, $X_n \in U_{\sigma_n} \subset U_{\sigma_{n_\epsilon}}$. Since $U_{\sigma_n} \subset U_{\sigma_{n_\epsilon}}$ for all $n < n_\epsilon$ too, the proof is complete in this case.

4. Now assume that $N < \infty$. We will show that $\{X'_n\}$ remains in $B_{R+1}(x_\star)$ for all $n \geq N$. Were this to hold, then for $n \geq N$,

$$\begin{aligned} \|X_n - x_\star\| &\leq \|X'_n - x_\star\| + \left\| \sum_{i=n+1}^{\infty} a_i \delta M_i 1_{\|X_{i-1} - x_\star\| \leq R+2} \right\| \\ &< R + 1 + \epsilon < R + 2, \end{aligned} \quad (2.25)$$

having used (2.21) and (2.22). For $n < N$, $X_n \in U_{\sigma_n} \subset U_{\sigma_N} = B_{r_N}(0)$. Therefore, for all n , $X_n \in B_{\tilde{R}}(0)$ where $\tilde{R} = \max\{r_N, \|x_\star\| + R + 2\}$.

5. We prove $X'_n \in B_{R+1}(x_\star)$ by induction. First, since $\epsilon < 1$ and $X_N \in U_0 \subset B_R(x_\star)$,

$$\|X'_N - x_\star\| \leq \|X'_N - X_N\| + \|X_N - x_\star\| < \epsilon + R < R + 1.$$

Next, assume that $X'_N, X'_{N+1}, \dots, X'_n$ are all in $B_{R+1}(x_\star)$. By (2.25), $X_n \in B_{R+2}(x_\star)$. Since $P_{n+1} \mathbf{1}_{\|X_n - x_\star\| \leq R+2} = 0$, we conclude $P_{n+1} = 0$. The modified iteration (2.23) simplifies to have

$$X'_{n+1} = X'_n - a_{n+1}f(X_n),$$

and

$$\begin{aligned} \|X'_{n+1} - x_\star\|^2 &= \|X'_n - x_\star\|^2 - 2a_{n+1} \langle X'_n - x_\star, f(X_n) \rangle \\ &\quad + a_{n+1}^2 \|f(X_n)\|^2 \\ &< \|X'_n - x_\star\|^2 - 2a_{n+1} \langle X'_n - x_\star, f(X_n) \rangle \\ &\quad + a_{n+1} \epsilon \bar{f}^2. \end{aligned} \tag{2.26}$$

6. We now consider two cases of (2.26). First, assume $\|X'_n - x_\star\| \leq R$. Then (2.26) can immediately be bounded as

$$\|X'_{n+1} - x_\star\|^2 < R^2 + 2\epsilon R \bar{f} + \epsilon^2 \bar{f}^2 = (R + \epsilon \bar{f})^2 < (R + 1)^2,$$

where we have used condition (2.20) in the last inequality.

7. Now consider the case $R < \|X'_n - x_\star\| < R + 1$. Using (2.20), the inner product in (2.26) can first be bounded from below:

$$\begin{aligned} \langle X'_n - x_\star, f(X_n) \rangle &= \langle X_n - x_\star, f(X_n) \rangle + \langle X'_n - X_n, f(X_n) \rangle \\ &\geq \langle X_n - x_\star, f(X_n) \rangle - \epsilon \bar{f} > \langle X_n - x_\star, f(X_n) \rangle - \frac{1}{2}\delta. \end{aligned}$$

Next, using (2.20)

$$\|X_n - x_\star\| \geq \|X'_n - x_\star\| - \|X_n - X'_n\| > R - \epsilon > R - \frac{1}{2}R = \frac{1}{2}R$$

Therefore, $\frac{1}{2}R < \|X_n - x_\star\| < R + 2$, so (2.19) ensures $\langle X_n - x_\star, f(X_n) \rangle \geq \delta$ and

$$\langle X'_n - x_\star, f(X_n) \rangle > \delta - \frac{1}{2}\delta = \frac{1}{2}\delta.$$

Returning to (2.26), by (2.20),

$$\|X'_{n+1} - x_\star\|^2 \leq (R + 1)^2 - a_{n+1}(\delta - \epsilon \bar{f}^2) < (R + 1)^2.$$

This completes the proof of the inductive step in this second case, completing the proof. $\square$

Corollary 2.1. *For the expanding trust region algorithm, if Assumptions 1, 2, 3, and 4 hold, then the number of truncations is a.s. finite.*

Proof. The proof is by contradiction. We break the proof into 4 steps:

1. Assuming that there are infinitely many truncations, Lemma 2.3 implies that the sequence cannot remain in a bounded set. Then, continuing to assume that Assumptions 1, 2, 3, and 4 hold, the only way for the conclusion of Lemma 2.4 to fail is if the assumption on $P_{n+1}1_{\|X_n - x_\star\| \leq r}$ is false. Therefore, there exists $r > 0$ and a set of positive measure on which a subsequence, $P_{n_k+1}1_{\|X_{n_k} - x_\star\| \leq r} \neq 0$. Hence $X_{n_k} \in B_r(x_\star)$, and $P_{n_k+1} \neq 0$. So truncations occur at these indices, and $\tilde{X}_{n_k+1} \notin U_{\sigma_{n_k}}$.
2. Let $\bar{f} = \sup \|f(x)\|$ with the supremum over the set $B_r(x_\star)$ and let $\epsilon > 0$ satisfy

$$\epsilon < (\bar{f} + 1)^{-1}. \quad (2.27)$$

By our assumptions of the lemma and Proposition 2.1, there exists n_ϵ such that for all $n \geq n_\epsilon$

$$\|a_{n+1}\delta M_{n+1}1_{\|X_n - x_\star\| \leq r}\| \leq \epsilon, \quad a_{n+1} \leq \epsilon \quad (2.28)$$

Along the subsequence, for all $n_k \geq n_\epsilon$,

$$\|a_{n_k+1}\delta M_{n_k+1}1_{\|X_{n_k} - x_\star\| \leq r}\| = \|a_{n_k+1}\delta M_{n_k+1}\| \leq \epsilon. \quad (2.29)$$

3. Furthermore, for $n_k \geq n_\epsilon$:

$$\begin{aligned} \|\tilde{X}_{n_k+1} - x_\star\| &\leq \|X_{n_k} - x_\star\| + a_{n_k+1}\|f(X_{n_k})\| + \|a_{n_k+1}\delta M_{n_k+1}\| \\ &< r + \epsilon\bar{f} + \epsilon < r + 1, \Rightarrow \tilde{X}_{n_k+1} \in B_{r+1}(x_\star), \end{aligned} \quad (2.30)$$

where (2.27) has been used in the last inequality.

4. By the definition of the U_n , there exists an index M such that $U_M \supset B_{r+1}(x_\star)$. Let

$$N = \inf\{n \geq n_\epsilon \mid \sigma_n \geq M\}. \quad (2.31)$$

This set is nonempty and $N < \infty$ since we have assumed there are infinitely many truncations. Let $n_k \geq N$. Then $\sigma_{n_k} \geq M$ and $U_{\sigma_{n_k}} \supset B_{r+1}(x_\star)$. But (2.30) then implies that $\tilde{X}_{n_k+1} \in U_{\sigma_{n_k}}$, and no truncation will occur; $P_{n_k+1} = 0$, providing the desired contradiction.

□

2.2. Proof of convergence

Using the above results, we are able to prove Theorems 2.1 and 2.2. Since the proofs are quite similar, we present the more complicated expanding trust region case.

Proof. We split this proof into 6 steps:

1. First, by Corollary 2.1, only finitely many truncations occur. By Lemma 2.3, there exists $R > 0$ such that $X_n \in B_R(0)$ for all n . Consequently, there is an r such that $X_n \in B_r(x_\star)$ for all n .

2. Next, we fix constants. Let $\bar{f} = \sup \|f(x)\|$ with the supremum taken over $B_r(x_\star)$. Fix $\eta \in (0, 2R)$, and use Assumption 2 to determine $\delta > 0$ such that

$$\inf_{\eta/2 \leq \|x - x_\star\| \leq r} \langle x - x_\star, f(x) \rangle = \delta \quad (2.32)$$

Take $\epsilon > 0$ such that:

$$\epsilon < \min \left\{ 1, \frac{\eta}{2}, \frac{\delta}{2\bar{f}}, \frac{\delta}{2\bar{f}^2} \right\} \quad (2.33)$$

Having set ϵ , we again appeal to Assumption 4 and Proposition 2.1 to find n_ϵ such that for all $n \geq n_\epsilon$:

$$\left\| \sum_{i=n+1}^{\infty} a_i \delta M_i 1_{\|X_{i-1} - x_\star\| \leq r} \right\| = \left\| \sum_{i=n+1}^{\infty} a_i \delta M_i \right\| \leq \epsilon, \quad a_{n+1} \leq \epsilon \quad (2.34)$$

3. Define the auxiliary sequence,

$$X'_n = X_n - \sum_{i=n+1}^{\infty} a_i \delta M_i 1_{\|X_{i-1} - x_\star\| \leq r} = X_n - \sum_{i=n+1}^{\infty} a_i \delta M_i. \quad (2.35)$$

Since there are only finitely many truncations, there exists $N \geq n_\epsilon$, such that for all $n \geq N$, $P_{n+1} = 0$, as the truncations have ceased. Consequently, for $n \geq N$,

$$X'_{n+1} = X'_n - a_{n+1} f(X_n) \quad (2.36)$$

By (2.34) and (2.35), for $n \geq N$, $\|X_n - X'_n\| \leq \epsilon$. Since $\epsilon > 0$ may be arbitrarily small, it will be sufficient to prove $X'_n \rightarrow x_\star$.

4. To obtain convergence of X'_n , we first examine $\|X'_{n+1} - x_\star\|$. For $n \geq N$,

$$\|X'_{n+1} - x_\star\|^2 \leq \|X'_n - x_\star\|^2 - 2a_{n+1} \langle X'_n - x_\star, f(X_n) \rangle + a_{n+1} \epsilon \bar{f}^2, \quad (2.37)$$

Now consider two cases of this expression. First, assume $\|X'_n - x_\star\| \leq \eta$. In this case, using (2.33),

$$\begin{aligned} -2a_{n+1} \langle X'_n - x_\star, f(X_n) \rangle + a_{n+1} \epsilon \bar{f}^2 &\leq a_{n+1} (2\eta \bar{f} + \epsilon \bar{f}^2) \\ &< a_{n+1} (4R\bar{f} + \bar{f}^2) = a_{n+1} B. \end{aligned} \quad (2.38)$$

where $B > 0$ is a constant depending only on R and $\bar{f}$. For $\|X'_n - x_\star\| > \eta$, using (2.33)

$$\begin{aligned} \langle X'_n - x_\star, f(X_n) \rangle &= \langle X_n - x_\star, f(X_n) \rangle + \langle X'_n - X_n, f(X_n) \rangle \\ &\geq \langle X_n - x_\star, f(X_n) \rangle - \epsilon \bar{f} \\ &> \langle X_n - x_\star, f(X_n) \rangle - \frac{1}{2} \delta. \end{aligned} \quad (2.39)$$

By (2.33),

$$\|X_n - x_\star\| \geq \|X'_n - x_\star\| - \|X_n - X'_n\| > \eta - \epsilon > \frac{1}{2} \eta$$

Since $\|X_n - x_\star\| < r$ too, (2.32) and (2.39) yield the estimate

$$\langle X'_n - x_\star, f(X_n) \rangle > \delta - \epsilon \bar{f} > \frac{1}{2} \delta$$

Thus, in this regime, using (2.33),

$$\begin{aligned} -2a_{n+1} \langle X'_n - x_\star, f(X_n) \rangle + a_{n+1} \epsilon \bar{f}^2 &\leq -a_{n+1}(\delta - \epsilon \bar{f}^2) \\ &< -\frac{1}{2} \delta a_{n+1} = -A a_{n+1} \end{aligned} \quad (2.40)$$

where $A > 0$ is a constant depending only on δ .

Combining estimates (2.38) and (2.40), we can write for $n \geq N$

$$\|X'_{n+1} - x_\star\|^2 < \|X'_n - x_\star\|^2 - a_{n+1} A 1_{\|X'_n - x_\star\| > \eta} + a_{n+1} B 1_{\|X'_n - x_\star\| \leq \eta}. \quad (2.41)$$

5. We now show that $\|X'_n - x_\star\| \leq \eta$ i.o. The argument is by contradiction. Let $M \geq N$ be such that for all $n \geq M$, $\|X'_n - x_\star\| > \eta$. For such n ,

$$\begin{aligned} \eta^2 &< \|X'_{n+1} - x_\star\|^2 < \|X'_n - x_\star\|^2 - a_{n+1} A \\ &< \|X'_{n-1} - x_\star\|^2 - a_{n+1} A - a_n A \\ &< \dots < \|X'_M - x_\star\|^2 - A \sum_{i=M}^n a_{i+1}. \end{aligned} \quad (2.42)$$

Using Assumption 4 and taking $n \rightarrow \infty$, we obtain a contradiction.

6. Finally, we prove convergence of $X'_n \rightarrow x_\star$. Since $X'_n \in B_\eta(x_\star)$ i.o., let

$$N' = \inf\{n \geq N \mid \|X'_n - x_\star\| < \eta\}. \quad (2.43)$$

For $n \geq N'$, we can then define

$$\varphi(n) = \max\{p \leq n \mid \|X'_p - x_\star\| < \eta\}. \quad (2.44)$$

For all such n , $\varphi(n) \leq n$, and $X'_{\varphi(n)} \in B_\eta(x_\star)$.

We claim that for $n \geq N'$,

$$\|X'_{n+1} - x_\star\|^2 < \|X'_{\varphi(n)} - x_\star\|^2 + B a_{\varphi(n)+1} < \eta^2 + B a_{\varphi(n)+1}.$$

First, if $n = \varphi(n)$, this trivially holds in (2.41). Suppose now that $n > \varphi(n)$. Then for $i = \varphi(n) + 1, \varphi(n) + 2, \dots, n$, $\|X'_i - x_\star\| > \eta$. Consequently,

$$\begin{aligned} \|X'_{n+1} - x_\star\|^2 &< \|X'_n - x_\star\|^2 < \|X'_{n-1} - x_\star\|^2 < \dots \\ &< \|X'_{\varphi(n)+1} - x_\star\|^2 < \|X'_{\varphi(n)} - x_\star\|^2 + B a_{\varphi(n)+1} \\ &< \eta^2 + B a_{\varphi(n)+1} \end{aligned}$$

As $\varphi(n) \rightarrow \infty$,

$$\limsup_{n \rightarrow \infty} \|X'_{n+1} - x_\star\|^2 \leq \eta^2$$

Since η may be arbitrarily small, we conclude that

$$\limsup_{n \rightarrow \infty} \|X'_{n+1} - x_\star\| = \lim_{n \rightarrow \infty} \|X'_{n+1} - x_\star\| = 0,$$

completing the proof. □

3. Minimization of relative entropy

Recall from the introduction that our distribution of interest, μ , is posed on the Borel subsets of Hilbert space $\mathcal{H}$. We assume that $\mu \ll \mu_0$, where $\mu_0 = N(m_0, C_0)$ is some reference Gaussian. Thus, we write

$$\frac{d\mu}{d\mu_0} = \frac{1}{Z_\mu} \exp\{-\Phi_\mu(u)\}, \quad (3.1)$$

where $\Phi_\nu : X \rightarrow \mathbb{R}$, X a Banach space, a subspace of $\mathcal{H}$, of full measure with respect to μ_0 , a Gaussian on $\mathcal{H}$, assumed to be continuous. $Z_\mu = \mathbb{E}^{\mu_0}[\exp\{-\Phi_\mu(u)\}] \in (0, \infty)$ is the partition function ensuring we have a probability measure.

Let $\nu = N(m, C)$, be another Gaussian, equivalent to μ_0 , such that we can write

$$\frac{d\nu}{d\mu_0} = \frac{1}{Z_\nu} \exp\{-\Phi_\nu(v)\}, \quad (3.2)$$

Assuming that $\nu \ll \mu$, we can write

$$\mathcal{R}(\nu|\mu) = \mathbb{E}^\nu[\Phi_\mu(u) - \Phi_\nu(u)] + \log(Z_\mu) - \log(Z_\nu) \quad (3.3)$$

The assumption that $\nu \ll \mu$ implies that ν and μ are equivalent measures. As was proven in [16], if $\mathcal{A}$ is a set of Gaussian measures, closed under weak convergence, such that at least one element of $\mathcal{A}$ is absolutely continuous with respect to μ , then any minimizing sequence over $\mathcal{A}$ will have a weak subsequential limit.

If we assume, for this work, that $C = C_0$, then, by the Cameron-Martin formula (see [9]),

$$\Phi_\nu(u) = -\langle u - m, m - m_0 \rangle_{\mathcal{H}^1} - \frac{1}{2} \|m - m_0\|_{\mathcal{H}^1}^2, \quad Z_\nu = 1. \quad (3.4)$$

Here, $\langle \bullet, \bullet \rangle_{\mathcal{H}^1}$ and $\|\bullet\|_{\mathcal{H}^1}$ are the inner product and norms of the Cameron-Martin Hilbert space, denoted $\mathcal{H}^1$,

$$\langle f, g \rangle_{\mathcal{H}^1} = \langle C_0^{-1/2} f, C_0^{-1/2} g \rangle, \quad \|f\|_{\mathcal{H}^1}^2 = \langle f, f \rangle_{\mathcal{H}^1}. \quad (3.5)$$

Convergence to the minimizer will be established in $\mathcal{H}^1$, and $\mathcal{H}^1$ will be the relevant Hilbert space in our application of Theorems 2.1 and 2.2 to this problem.

Letting $\nu_0 = N(0, C_0)$ and $\nu \sim \nu_0$, we can then rewrite (3.3) as

$$J(m) \equiv \mathcal{R}(\nu|\mu) = \mathbb{E}^{\nu_0}[\Phi_\mu(v + m)] + \frac{1}{2} \|m - m_0\|_{\mathcal{H}^1}^2 + \log(Z_\mu) \quad (3.6)$$

The Euler-Lagrange equation associated with (3.6), and the second variation, are:

$$J'(m) = \mathbb{E}^{\nu_0}[\Phi'_\mu(v + m)] + C_0^{-1}(m - m_0), \quad (3.7)$$

$$J''(m) = \mathbb{E}^{\nu_0}[\Phi''_\mu(v + m)] + C_0^{-1}. \quad (3.8)$$

3.1. Application of Robbins-Monro

In [15], it was suggested that rather than try to find a root of (3.7), the equation first be preconditioned by multiplying by C_0 ,

$$C_0 \mathbb{E}^{v_0} [\Phi'_\mu(v + m)] + (m - m_0), \quad (3.9)$$

and a root of this mapping is sought, instead. Defining

$$f(m) = C_0 \mathbb{E}^{v_0} [\Phi'_\mu(v + m)] + (m - m_0), \quad (3.10a)$$

$$F(m, v) = C_0 \Phi'_\mu(v + m) + (m - m_0). \quad (3.10b)$$

The Robbins-Monro formulation is then

$$m_{n+1} = m_n - a_{n+1} F(m_n, v_{n+1}) + P_{n+1}, \quad (3.11)$$

with $v_n \sim v_0$, i.i.d.

We thus have

Theorem 3.1. Assume:

- There exists $v = N(m, C_0) \sim \mu_0$ such that $v \ll \mu$.
- Φ'_μ and Φ''_μ exist for all $u \in \mathcal{H}^1$.
- There exists $m_\star$, a local minimizer of J , such that $J'(m_\star) = 0$.
- The mapping

$$m \mapsto \mathbb{E}^{v_0} \left[\left\| \sqrt{C_0} \Phi'_\mu(m + v) \right\|^2 \right] \quad (3.12)$$

is bounded on bounded subsets of $\mathcal{H}^1$.

- There exists a convex neighborhood $U_\star$ of $m_\star$ and a constant $\alpha > 0$, such that for all $m \in U_\star$, for all $u \in \mathcal{H}^1$,

$$\langle J''(m)u, u \rangle \geq \alpha \|u\|_{\mathcal{H}^1}^2 \quad (3.13)$$

Then, choosing a_n according to Assumption 4,

- If the subset $U_\star$ can be taken to be all of $\mathcal{H}^1$, for the expanding truncation algorithm, $m_n \rightarrow m_\star$ a.s. in $\mathcal{H}^1$.
- If the subset $U_\star$ is not all of $\mathcal{H}^1$, then, taking U_1 to be a bounded (in $\mathcal{H}^1$) convex subset of $U_\star$, with $m_\star \in U_1$, and U_0 any subset of U_1 such that there exist $R_0 < R_1$ with

$$U_0 \subset B_{R_0}(x_\star) \subset B_{R_1}(x_\star) \subset U_1,$$

for the fixed truncation algorithm, $m_n \rightarrow m_\star$ a.s. in $\mathcal{H}^1$.

Proof. We split the proof into 2 steps:

1. By the assumptions of the theorem, we clearly satisfy Assumptions 1 and 4. To satisfy Assumption 3, we observe that

$$\mathbb{E}^{v_0} [\|F(m, v)\|_{\mathcal{H}^1}^2] \leq 2\mathbb{E}^{v_0} \left[\left\| \sqrt{C_0} \Phi'_\mu(m + v) \right\|^2 \right] + 2\|m - m_0\|_{\mathcal{H}^1}^2.$$

This is bounded on bounded subsets of $\mathcal{H}^1$.

2. Per the convexity assumption, (3.13), implies Assumption 2, since, by the mean value theorem in function spaces,

$$\begin{aligned}\langle m - m_\star, f(m) \rangle_{\mathcal{H}^1} &= \langle m - m_\star, C_0 [J'(m_\star) + J''(\tilde{m})(m - m_\star)] \rangle_{\mathcal{H}^1} \\ &= \langle m - m_\star, J''(\tilde{m})(m - m_\star) \rangle \geq \alpha \|m - m_\star\|_{\mathcal{H}^1}^2\end{aligned}$$

where $\tilde{m}$ is some intermediate point between m and $m_\star$. This completes the proof. $\square$

While condition (3.13) is sufficient to obtain convexity, other conditions are possible. For instance, suppose there is a convex open set $U_\star$ containing $m_\star$ and constant $\theta \in [0, 1)$, such that for all $m \in U_\star$,

$$\inf_{\substack{u \in \mathcal{H} \\ u \neq 0}} \frac{\langle \mathbb{E}^{v_0}[\Phi_\mu''(v + m)]u, u \rangle}{\|u\|^2} \geq -\theta\lambda_1^{-1}, \quad (3.14)$$

where λ_1 is the principal eigenvalue of C_0 . Then this would also imply Assumption 2, since

$$\begin{aligned}\langle m - m_\star, f(m) \rangle_{\mathcal{H}^1} &= \langle m - m_\star, C_0 [J'(m_\star) + J''(\tilde{m})(m - m_\star)] \rangle_{\mathcal{H}^1} \\ &= \langle m - m_\star, J''(\tilde{m})(m - m_\star) \rangle \\ &\geq \|m - m_\star\|_{\mathcal{H}^1}^2 + \langle m - m_\star, \mathbb{E}^{v_0}[\Phi_\mu''(v + \tilde{m})](m - m_\star) \rangle \\ &\geq \|m - m_\star\|_{\mathcal{H}^1}^2 - \theta\lambda_1^{-1} \|m - m_\star\|^2 \\ &\geq (1 - \theta) \|m - m_\star\|_{\mathcal{H}^1}^2.\end{aligned}$$

We mention (3.14) as there may be cases, shown below, for which the operator $\mathbb{E}^{v_0}[\Phi_\mu''(v + m)]$ is obviously nonnegative.

4. Examples

To apply the Robbins-Monro algorithm to the relative entropy minimization problem, the Φ_μ functional of interest must be examined. In this section we present a few examples, based on those presented in [15], and examine when the assumptions hold. The one outstanding assumption that we must make is that, *a priori*, μ_0 is an equivalent measure to μ .

4.1. Scalar problem

Taking $\mu_0 = N(0, 1)$, the standard unit Gaussian, let $V : \mathbb{R} \rightarrow \mathbb{R}$ be a smooth function such that

$$\frac{d\mu}{d\mu_0} = \frac{1}{Z_\mu} \exp\{-\epsilon^{-1}V(x)\} \quad (4.1)$$

is a probability measure on $\mathbb{R}$. For these scalar cases, we use x in place of v . In the above framework,

$$\begin{aligned}F(x, \xi) &= \epsilon^{-1}V'(x + \xi) - \xi, \\ f(x) &= \epsilon^{-1}\mathbb{E}[V'(x + \xi)]m \\ \Phi'_\mu(x) &= \epsilon^{-1}V'(x), \\ \Phi''_\mu(x) &= \epsilon^{-1}V''(x)\end{aligned}$$

and $\xi \sim N(0, 1) = v_0 = \mu_0$.

4.1.1. Globally convex case

Consider the case that

$$V(x) = \frac{1}{2}x^2 + \frac{1}{4}x^4. \quad (4.2)$$

In this case

$$\begin{aligned} F(x, \xi) &= \epsilon^{-1} \left(x + \xi + (x + \xi)^3 \right) + x, \\ f(x) &= \epsilon^{-1} (4x + x^3) + x, \\ \mathbb{E}[\Phi''_{\mu}(x + \xi)] &= \epsilon^{-1} (4 + 3x^2), \\ \mathbb{E}[\Phi'_{\mu}(x + \xi)^2] &= \epsilon^{-1} (22 + 58x^2 + 17x^4 + x^6). \end{aligned}$$

Since $\mathbb{E}[\Phi''_{\mu}(x + \xi)] \geq 4\epsilon^{-1}$, all of our assumptions are satisfied and the expanding truncation algorithm will converge to the unique root at $x_{\star} = 0$ a.s. See Figure 1 for an example of the convergence at $\epsilon = 0.1$, $U_n = (-n - 1, n + 1)$, and always restarting at 0.5.

We refer to this as a “globally convex” problem since $\mathcal{R}$ is globally convex about the minimizer.




Figure 1. Robbins-Monro applied to a globally convex scalar problem associated with (4.2) with $\epsilon = 0.1$ and expanding trust regions $U_n = (-1 - n, 1 + n)$.

4.1.2. Locally convex case

In contrast to the above problem, some mimizers are only “locally” convex. Consider the case the double well potential

$$V(x) = \frac{1}{4}(4 - x^2)^2 \quad (4.3)$$

Now, the expressions for RM are

$$\begin{aligned} F(x, \xi) &= \epsilon^{-1} \left((x + \xi)^3 - 4(x + \xi) \right) + x, \\ f(x) &= \epsilon^{-1} (x^3 - x) + x, \\ \mathbb{E}[\Phi''_{\mu}(x + \xi)] &= \epsilon^{-1} (3x^2 - 1), \\ \mathbb{E}[\Phi'_{\mu}(x + \xi)^2] &= \epsilon^{-1} (1 + x^2)(7 + 6x^2 + x^4). \end{aligned}$$

In this case, $f(x)$ vanishes at 0 and $\pm\sqrt{1-\epsilon}$, and J'' changes sign from positive to negative when x enters $(-\sqrt{(1-\epsilon)/3}, \sqrt{(1-\epsilon)/3})$. We must therefore restrict to a fixed trust region if we want to ensure convergence to either of $\pm\sqrt{1-\epsilon}$.

We ran the problem at $\epsilon = 0.1$ in two cases. In the first case, $U_1 = (0.6, 3.0)$ and the process always restarts at 2. This guarantees convergence since the second variation will be strictly positive. In the second case, $U_1 = (-0.5, 1.5)$, and the process always restarts at -0.1 . Now, the second variation can change sign. The results of these two experiments appear in Figure 2. For some random number sequences the algorithm still converged to $\sqrt{1-\epsilon}$, even with the poor choice of trust region.




Figure 2. Robbins-Monro applied to the nonconvex scalar problem associated with (4.3). Figure (a) shows the result with a well chosen trust region, while (b) shows the outcome of a poorly chosen trust region.

4.2. Path space problem

Take $\mu_0 = N(m_0(t), C_0)$, with

$$C_0 = \left(-\frac{d^2}{dt^2} \right)^{-1}, \quad (4.4)$$

equipped with Dirichlet boundary conditions on $\mathcal{H} = L^2(0, 1)$.^{*} In this case the Cameron-Martin space $\mathcal{H}^1 = H_0^1(0, 1)$, the standard Sobolev space equipped with the Dirichlet norm. Let us assume $m_0 \in H^1(0, 1)$, taking values in $\mathbb{R}^d$.

Consider the path space distribution on $L^2(0, 1)$, induced by

$$\frac{d\mu}{d\mu_0} = -\frac{1}{Z_\mu} \exp \left\{ -\Phi_\mu(v) \right\}, \quad \Phi_\mu(u) = \epsilon^{-1} \int_0^1 V(v(t)) dt, \quad (4.5)$$

where $V : \mathbb{R}^d \rightarrow \mathbb{R}$ is a smooth function. We assume that V is such that this probability distribution exists and that $\mu \sim \mu_0$, our reference measure.

We thus seek an $\mathbb{R}^d$ valued function $m(t) \in H^1(0, 1)$ for our Gaussian approximation of μ , satisfying the boundary conditions

$$m(0) = m_-, \quad m(1) = m_+. \quad (4.6)$$

^{*}This is the covariance of the standard unit Brownian bridge, $Y_t = B_t - tB_1$.

For simplicity, take $m_0 = (1-t)m_- + tm_+$, the linear interpolant between $m_{\pm}$. As above, we work in the shifted coordinated $x(t) = m(t) - m_0(t) \in H_0^1(0, 1)$.

Given a path $v(t) \in H_0^1$, by the Sobolev embedding, v is continuous with its L^∞ norm controlled by its H^1 norm. Also recall that for $\xi \sim N(0, C_0)$, in the case of $\xi(t) \in \mathbb{R}$,

$$\mathbb{E} [\xi(t)^p] = \begin{cases} 0, & p \text{ odd,} \\ (p-1)!! [t(1-t)]^{\frac{p}{2}}, & p \text{ even.} \end{cases} \quad (4.7)$$

Letting $\lambda_1 = 1/\pi^2$ be the ground state eigenvalue of C_0 ,

$$\begin{aligned} \mathbb{E}[\|\sqrt{C_0}\Phi'_\mu(v + m_0 + \xi)\|^2] &\leq \lambda_1 \mathbb{E}[\|\Phi'_\mu(v + m_0 + \xi)\|^2] \\ &= \lambda_1 \epsilon^{-2} \int_0^1 \mathbb{E}[|V'(v(t) + m_0(t) + \xi(t))|^2] dt. \end{aligned}$$

The terms involving $v+m_0$ in the integrand can be controlled by the L^∞ norm, which in turn is controlled by the H^1 norm, while the terms involving ξ can be integrated according to (4.7). As a mapping applied to x , this expression is bounded on bounded subsets of H^1 .

Minimizers will satisfy the ODE

$$\epsilon^{-1} \mathbb{E}[V'(x + m_0 + \xi)] - x'' = 0, \quad x(0) = x(1) = 0. \quad (4.8)$$

4.3. Globally convex example

With regard to convexity about a minimizer, $m_\star$, if, for instance, V'' were pointwise positive definite, then the problem would satisfy (3.14), ensuring convergence. Consider the quartic potential V given by (4.2). In this case,

$$\Phi(v) = \epsilon^{-1} \int_0^1 \frac{1}{2} v(t)^2 + \frac{1}{4} v(t)^4 dt, \quad (4.9)$$

and

$$\begin{aligned} \Phi'(v + m_0 + \xi) &= \epsilon^{-1} \left[(v + m_0 + \xi) + 3(v + m_0 + \xi)^3 \right] \\ \Phi''(v + m_0 + \xi) &= \epsilon^{-1} \left[1 + 3(v + m_0 + \xi)^2 \right], \\ \mathbb{E}[\Phi'(v + m_0 + \xi)] &= \epsilon^{-1} \left[v + m_0 + (v + m_0)^3 + 3t(1-t)(v + m_0) \right] \\ \mathbb{E}[\Phi''(v + m_0 + \xi)] &= \epsilon^{-1} \left[1 + 3(v + m_0)^2 + 3t(1-t) \right] \end{aligned}$$

Since $\Phi''(v + m_0 + \xi) \geq \epsilon^{-1}$, we are guaranteed convergence using expanding trust regions. Taking $\epsilon = 0.01$, $m_- = 0$ and $m_+ = 2$, this is illustrated in Figure 3, where we have also solved (4.8) by ODE methods for comparison. As trust regions, we take

$$U_n = \{m \in H_0^1(0, 1) \mid \|x\|_{H^1} \leq 10 + n\}, \quad (4.10)$$

and we always restart at the zero solution Figure 3 also shows robustness to discretization; the number of truncations is relatively insensitive to Δt .




Figure 3. The mean paths computed for (4.9) at different resolutions, along with the truncation sequence.

4.4. Locally convex example

For many problems of interest, we do not have global convexity. Consider the double well potential (4.3), but in the case of paths,

$$\Phi(u) = \epsilon^{-1} \int_0^1 \frac{1}{4} (4 - v(t)^2)^2 dt. \quad (4.11)$$

Then,

$$\begin{aligned} \Phi'(v + m_0 + \xi) &= \epsilon^{-1} \left[(v + m_0 + \xi)^3 - 4(v + m_0 + \xi) \right] \\ \Phi''(v + m_0 + \xi) &= \epsilon^{-1} \left[3(v + m_0 + \xi)^2 - 4 \right], \\ \mathbb{E}[\Phi'(v + m_0 + \xi)] &= \epsilon^{-1} \left[(v + m_0)^3 + 3t(1-t)(v + m_0) - 4(v + m_0) \right] \\ \mathbb{E}[\Phi''(v + m_0 + \xi)] &= \epsilon^{-1} \left[3(v + m_0)^2 + 3t(1-t) - 4 \right] \end{aligned}$$

Here, we take $m_- = 0$, $m_+ = 2$, and $\epsilon = 0.01$. We have plotted the numerically solved ODE in Figure 4. Also plotted is $\mathbb{E}[\Phi''(v_\star + m_0 + \xi)]$. Note that $\mathbb{E}[\Phi''(v_\star + m_0 + \xi)]$ is not sign definite, becoming as small as -400 . Since C_0 has $\lambda_1 = 1/\pi^2 \approx 0.101$, (3.14) cannot apply.

Discretizing the Schrödinger operator

$$J''(v_\star) = -\frac{d^2}{dt^2} + \epsilon^{-1} \left(3(v_\star(t) + m_0(t))^2 + 3t(1-t) - 4 \right), \quad (4.12)$$

we numerically compute the eigenvalues. Plotted in Figure 5, we see that the minimal eigenvalue of $J''(m_\star)$ is approximately $\mu_1 \approx 550$. Therefore,

$$\langle J''(x_\star)u, u \rangle \geq \mu_1 \|u\|_{L^2}^2 \Rightarrow \langle J''(x)u, u \rangle \geq \alpha \|u\|_{H^1}^2, \quad (4.13)$$

for all v in some neighborhood of $v_\star$. For an appropriately selected fixed trust region, the algorithm will converge.

However, we can show that the convexity condition is not global. Consider the path $m(t) = 2t^2$, which satisfies the boundary conditions. As shown in Figure 5, this path induces negative eigenvalues.




Figure 4. The numerically computed solution to (4.8) in the case of the double well, (4.11), m_* , and the associated $\mathbb{E}^{v_0}[\Phi''(m_* + \xi)]$.




Figure 5. The numerically computed spectrum for (4.12), associated with the m_* shown in Figure 4. Also shown is the numerically computed spectrum for the path $m(t) = 2t^2$, which introduces negative eigenvalues.

Despite this, we still observe convergence. Using the fixed trust region

$$U_1 = \{x \in H_0^1(0, 1) \mid \|x\|_{H^1} \leq 100\}, \quad (4.14)$$

we obtain the results in Figure 6. Again, the convergence is robust to discretization.




Figure 6. The mean paths computed for (4.11) at different resolutions, along with the truncation sequence.

5. Discussion

We have shown that the Robbins-Monro algorithm, with both fixed and expanding trust regions, can be applied to Hilbert space valued problems, adapting the finite dimensional proof of [12]. We have also constructed sufficient conditions for which the relative entropy minimization problem fits within this framework.

One problem we did not address here was how to identify fixed trust regions. Indeed, that requires a tremendous amount of *a priori* information that is almost certainly not available. We interpret that result as a local convergence result that gives a theoretical basis for applying the algorithm. In practice, since the root is likely unknown, one might run some numerical experiments to identify a reasonable trust region, or just use expanding trust regions. The practitioner will find that the algorithm converges to a solution, though perhaps not the one originally envisioned. A more sophisticated analysis may address the convergence to a set of roots, while being agnostic as to which zero is found.

Another problem we did not address was how to optimize not just the mean, but also the covariance in the Gaussian. As discussed in [15], it is necessary to parameterize the covariance in some way, which will be application specific. Thus, while the form of the first variation of relative entropy with respect to the mean, (3.7), is quite generic, the corresponding expression for the covariance will be specific to the covariance parameterization. Additional constraints are also necessary to guarantee that the parameters always induce a covariance operator. We leave such specialization as future work.

Acknowledgments

This work was supported by US Department of Energy Award DE-SC0012733. This work was completed under US National Science Foundation Grant DMS-1818716. The authors would like to thank J. Lelong for helpful comments, along with anonymous reviewers whose reports significantly impacted our work.

Conflict of interest

The authors declare that there is no conflicts of interest in this paper.

References

1. A. E. Albert and L. A. Gardner Jr., *Stochastic approximation and nonlinear regression*, The MIT Press, 1967.
2. C. Andrieu and É. Moulines, *On the ergodicity properties of some adaptive MCMC algorithms*, Ann. Appl. Probab., **16** (2006), 1462–1505.
3. C. Andrieu and J. Thoms, *A tutorial on adaptive MCMC*, Stat. Comput., **18** (2008), 343–373.
4. D. M. Blei, A. Kucukelbir and J. D. McAuliffe, *Variational Inference: A Review for Statisticians*, J. Am. Stat. Assoc., **112** (2017), 859–877.
5. J. R. Blum, *Approximation methods which converge with probability one*, Annals of Mathematical Statistics, **25** (1954), 382–386.
6. J. R. Blum, *Multidimensional stochastic approximation methods*, Annals of Mathematical Statistics, **25** (1954), 737–744.
7. S. D. Chatterji, *Martingale convergence and the Radon-Nikodym theorem in Banach spaces*, Math. Scand., **22** (1968), 21–41.
8. H.-F. Chen, *Stochastic approximation and its applications*, 2002.
9. G. Da Prato, *An Introduction to Infinite-Dimensional Analysis*, Springer, 2006.
10. H. Haario, E. Saksman and J. Tamminen, *An adaptive Metropolis algorithm*, Bernoulli, **7** (2001), 223–242.
11. H. J. Kushner and G. Yin, *Stochastic Approximation and Recursive Algorithms and Applications*, Springer, 2003.
12. J. Lelong, *Almost sure convergence of randomly truncated stochastic algorithms under verifiable conditions*, Stat. Probabil. Lett., **78** (2008), 2632–2636.
13. Y. Lu, A. Stuart and H. Weber, *Gaussian approximations for probability measures on $\mathbb{R}^d$* , SIAM/ASA Journal on Uncertainty Quantification, **5** (2017), 1136–1165.
14. Y. Lu, A. Stuart and H. Weber, *Gaussian approximations for transition paths in brownian dynamics*, SIAM J. Math. Anal., **49** (2017), 3005–3047.
15. F. J. Pinski, G. Simpson, A. M. Stuart, et al. *Algorithms for Kullback–Leibler Approximation of Probability Measures in Infinite Dimensions*, SIAM J. Sci. Comput., **37** (2015), A2733–A2757.
16. F. J. Pinski, G. Simpson, A. M. Stuart, et al. *Kullback–Leibler Approximation for Probability Measures on Infinite Dimensional Spaces*, SIAM J. Math. Anal., **47** (2015), 4091–4122.
17. H. Robbins and S. Monro, *A stochastic approximation method*, The Annals of Mathematical Statistics, (1950), 400–407.
18. G. O. Roberts and J. S. Rosenthal, *Coupling and Ergodicity of Adaptive Markov Chain Monte Carlo Algorithms*, J. Appl. Probab., **44** (2007), 458–475.

19. D. Sanz-Alonso and A. M. Stuart, *Gaussian approximations of small noise diffusions in kullback–leibler divergence*, Commun. Math. Sci., **15** (2017), 2087–2097.
20. G. Yin and Y. M. Zhu, *On H -valued Robbins-Monro processes*, J. Multivariate Anal., **34** (1990), 116–140.



AIMS Press

©2019 the Author(s), licensee AIMS Press. This is an open access article distributed under the terms of the Creative Commons Attribution License (<http://creativecommons.org/licenses/by/4.0>)