Proceedings of the
44th IEEE Conference on Decision and Control, and
the European Control Conference 2005
Seville, Spain, December 12-15, 2005

Stability of Stochastic Approximation under Verifiable Conditions

Christophe Andrieu
University of Bristol,
c.andrieu@bris.ac.uk

Éric Moulines
ENST, PARIS
moulines@tsi.enst.fr

Pierre Priouret
Universit Pierre & Marie Curie,
PARIS Cdex 05
priouret@ccr.jussieu.fr

Abstract

In this paper we address the problem of the stability and convergence of the stochastic approximation procedure

$$\theta_{n+1} = \theta_n + \gamma_{n+1}[h(\theta_n) + \xi_{n+1}].$$

The stability of such sequences $\{\theta_n\}$ is known to heavily rely on the behaviour of the mean field h at the boundary of the parameter set and the magnitude of the stepsizes used. The conditions typically required to ensure convergence, and in particular the boundedness or stability of $\{\theta_n\}$, are either too difficult to check in practice or not satisfied at all. The most popular technique to circumvent the stability problem consists of constraining $\{\theta_n\}$ to a compact subset $\mathscr{K}$ in the parameter space. This is obviously not a satisfactory solution as the choice of $\mathscr{K}$ is a delicate one. In the present contribution we first prove a "deterministic" stability result which relies on simple conditions on the sequences $\{\xi_n\}$ and $\{\gamma_n\}$. We then propose and analyze an algorithm based on projections on adaptive truncation sets which ensures that the aforementioned conditions required for stability are satisfied. We focus in particular on the case where $\{\xi_n\}$ is a so-called Markov state-dependent noise.

1. Introduction

In many contexts it is of interest to find the roots of possibly non linear equations of the form

$$h(\theta) = 0, \qquad \theta \in \Theta, \tag{1}$$

for some mapping $h : \Theta \to \mathbb{R}^{n_\theta}$, where $\Theta \subset \mathbb{R}^{n_\theta}$ for some integer n_θ. Most of the methods for solving the previous equation are iterative, *i.e.* produce a sequence of iterates $\{\theta_n, n \geq 0\}$ which eventually converges to the set of solutions of Eq. (1),

$$\mathscr{S} := \{\theta \in \Theta, h(\theta) = 0\}. \tag{2}$$

Stochastic Approximation (SA) is a class of algorithms to solve Eq. (1) in the situation where only noisy measurements of h are available. In its simplest form, the Robbins-Monro algorithm produces a sequence $\{\theta_n, n \geq 0\}$ defined recursively as follows,

$$\theta_0 \in \Theta, \qquad \theta_{n+1} = \theta_n + \gamma_{n+1}\zeta_{n+1}, \quad n \geq 1, \tag{3}$$

where $\{\gamma_n, n \geq 1\}$ is a sequence of stepsizes which satisfies standard conditions (say $\gamma_n \downarrow 0$ and $\sum_{n\geq 1} \gamma_n = \infty$) and for any $n \geq 1$, ζ_n is a noisy measurement of $h(\theta_n)$. It is useful to introduce the sequence $\{\xi_n, n \geq 1\}$ defined as

$$\zeta_{n+1} = h(\theta_n) + \xi_{n+1}, \tag{4}$$

which will be referred to as the *noise sequence*. Convergence of SA has been studied under various sets of assumptions for the mean field h and the noise sequence $\{\xi_n, n \geq 1\}$ since the early work by [14]; see *e.g.* [2],[13], [15], [12] and the references therein. Essentially, convergence of the SA sequence can be established toward an *attractive* subset provided that the sequence $\{\theta_n, n \geq 0\}$ is with probability 1 (hereafter w.p. 1) in a compact subset of Θ and is w.p. 1 infinitely often in the *domain of attraction* of this attractive subset. Showing in practice that $\{\theta_n, n \geq 0\}$ satisfies these boundedness and recurrence conditions proves to be a difficult task. The available results hold under conditions which are still restrictive, despite recent advances (see [1], [4], [3] and references therein). This major drawback has motivated the design of modified Robbins-Monro recursions. Probably the most widely used method in practice consists of constraining the sequence $\{\theta_n, n \geq 0\}$ to some compact set $\mathscr{K} \subset \Theta$ by means of a reprojection onto $\mathscr{K}$. This method has been thoroughly investigated in [15] (see also [5] and the references therein). Although relatively easy to implement, and sound when constraints about the system considered are available *a priori*, this approach becomes impractical and questionable in many situations of interest.
Our contributions to solve the stability and convergence problems are here twofold:

First we establish and prove in Section 2 a general result of stability, Theorem 1, for deterministic sequences of the form given by Eqs. (3)-(4). This key deterministic result assumes the existence of a global Lyapunov function for the mean field h and mild general assumptions about the noise and stepsize sequences. In contrast with previous results, the conditions required on the growth of the Lyapunov functions and the mean field h when θ approaches the boundaries of the parameter set Θ are minimal. As a consequence the result is applicable to quite general settings. We then show that, under the conditions that guarantee stability, the convergence of the deterministic sequence Eq. (3)-(4) is ensured (see Theorem 2).

Our second contribution here consists of proposing a SA algorithm (Section 3) for which the aforementioned noise and stepsize conditions are satisfied w.p. 1. There are many different applications of stochastic approximations which imply markedly different types of assumptions on the noise sequence $\{\xi_n\}$. Whereas our deterministic stability and convergence results mentioned above can be applied quite generally, we focus in this paper on the subtle Markov state dependent noise (see [15, Chapter 6, Section 6.6] and Section 3 in this paper), for which the availability of algorithms whose convergence can be established under general but nevertheless verifiable assumptions is still missing. The proposed algorithm is a modification of the classical Robbins-Monro procedure described in Eq. (3)-(4), based on truncations on adaptive truncation sets, in the spirit of the seminal works [8] and [7].

The convergence of SA with adaptive truncation sets has been considered under various conditions on the noise sequence $\{\xi_n\}$. These include state-independent noise conditions (see for example [9, Section 2.4, pp. 42-44]) but also state-dependent martingale differences ([17], [11], [6], [9, Section 2.5, pp. 49-57]) or state-dependent ϕ-mixing processes ([6], [9, Section 2.5, pp. 49]). However the application of this strategy to the Markovian state dependent case requires even more care, and it is therefore not surprising to find that the results on the topic are scarce, and have been obtained under conditions that are more stringent than those considered in the present paper; see [18], [10] and for the special case of ARMAX models, [9, Chapter 6]. As we shall see our procedure differs in some respects from the original procedure proposed by [8] and [7], and offers additional degrees of freedom. Our technique of proof for the stability relies on a novel approach and offers as a byproduct an explicit bound for the tail probability of the number of reprojections, which is found to be super-exponential under mild technical conditions.

2. Key deterministic results

In this section we establish both stability and convergence results for deterministic recursions of the type described in Eqs. (3)-(4). Before stating our first assumptions, some definitions and notation are needed. Let d be a positive integer. An element v of $\mathbb{R}^d$ is denoted by its column vector v and its transpose is denoted by v^{T}. For elements v, w of $\mathbb{R}^d$, we denote $< v, w >$ their inner product, so that $|v| = \sqrt{< v, v >}$ denotes the norm of v. Our first assumption is the existence of a global Lyapunov function w for the mean field h. Denoting $\mathscr{W}_M := \{\theta \in \Theta, w(\theta) \leq M\} \subset \Theta$ we assume,

(A1) Θ is an open subset of $\mathbb{R}^{n_\theta}$, $h : \Theta \to \mathbb{R}^{n_\theta}$ is continuous and there exists a continuously differentiable function $w : \Theta \to [0, \infty)$ such that

(i) There exists $M_0 > 0$ such that

$$\mathscr{L} := \left\{\theta \in \Theta, \left\langle \nabla w(\theta), h(\theta) \right\rangle = 0 \right\} \subset \{\theta \in \Theta, w(\theta) < M_0\},$$

(ii) There exists $M_1 \in (M_0, \infty]$ such that $\mathscr{W}_{M_1}$ is a compact set,

(iii) For any $\theta \in \Theta \setminus \mathscr{L}$, $\left\langle \nabla w(\theta), h(\theta) \right\rangle < 0$,

(iv) The closure of $w(\mathscr{L})$ has an empty interior.

Our approach to prove our stability and convergence results can be decomposed into two distinct steps. In the first step (this section), we establish deterministic conditions on a noise sequence $\{\xi_n\}$ and a stepsize sequence $\{\rho_n\}$ upon which a deterministic sequence $\{\theta_n\}$ defined as

$$\theta_0 \in \Theta \quad \theta_{n+1} = \theta_n + \rho_{n+1}[h(\theta_n) + \xi_{n+1}] \quad \text{for} \quad n \geq 0, \tag{5}$$

has the following properties: (i) it remains in a compact subset of Θ (see Theorem 1) and (ii) provided that $\{\theta_n\}$ remains in a compact subset of Θ, converges to $\mathscr{L}$ (Theorem 2). In a second step - which is probabilistic in nature and depends on how the noise is generated - we develop a general algorithm for the case where $\{\xi_n\}$ follows a Markovian state-dependent dynamic which allows one to show that the required condition on $\{\xi_n\}$ is satisfied w.p. 1.

Theorem 1 shows that under (A1) and mild additional conditions on $\{\xi_n\}$ and $\{\rho_n\}$, the sequence defined in Eq. (5) remains in a compact subset of Θ.

Theorem 1. *Assume (A1). For any $M \in (M_0, M_1]$ there exist $\delta_0 > 0$ and $\lambda_0 > 0$ such that, for all $n \geq 1$, all $\theta_0 \in \mathscr{W}_{M_0}$, all sequences $\{\rho_k\}$ of non negative integers and all sequences $\{\xi_k\}$ of n_θ-dimensional vectors satisfying*

$$\sup_{1 \leq k \leq n} \rho_k \leq \lambda_0 \quad \text{and} \sup_{1 \leq k \leq n} \left| \sum_{j=1}^{k} \rho_j \xi_j \right| \leq \delta_0,$$

we have for $k \in \{1, \ldots, n\}$, $w(\theta_k) \leq M$, where $\theta_k = \theta_{k-1} + \rho_k h(\theta_{k-1}) + \rho_k \xi_k$.

In the next proposition we show that whenever $\{\theta_k\}$ stays in a compact subset of Θ, then under mild additional assumptions it converges to $\mathscr{L}$. The key result of this section is the following theorem, adapted here from [11, Theorem 2] (see [9] for a similar result). For an integer d and A a subset of $\mathbb{R}^d$, we define $d(x,A) = \inf\{y \in A, |x-y|\}$. For any set $A \subset \Theta$ and any $\delta > 0$, we define $A_\delta := \{\theta \in \Theta, d(\theta,A) \leq \delta\}$; for any function $\phi : \Theta \to \mathbb{R}$, we define $\|\phi\|_A := \sup_{\theta\in A} |\phi(\theta)|$.

Theorem 2. *Assume (A1). Let $\mathscr{K}$ be a compact subset of Θ such that $\mathscr{L} \cap \mathscr{K} \neq \emptyset$. Let $\{\rho_k\}$ be a monotone non-increasing sequence of positive numbers such that $\rho_0 \leq \lambda_0$ (where λ_0 is given in Theorem 1),*

$$\sum_{k=1}^{\infty} \rho_k = \infty \quad \text{and} \quad \lim_{k\to\infty} \rho_k = 0.$$

Let $\{\xi_n\}$ be a sequence in $\mathbb{R}^{n_\theta}$ satisfying $\limsup_{k\to\infty} \sup_{l\geq k} \left|\sum_{i=k}^{l} \rho_i \xi_i\right| = 0$. Assume that the sequence defined by $\theta_k = \theta_{k-1} + \rho_k h(\theta_{k-1}) + \rho_k \xi_k$, is such that $\{\theta_k\} \subset \mathscr{K}$. Then, $\limsup_{k\to\infty} d(\theta_k, \mathscr{L} \cap \mathscr{K}) = 0$.

Note that the boundedness is here one of the required assumption. It is therefore natural to try to apply Theorem 1. This is what motivates the next section, where we describe a modification of the stochastic approximation algorithm which ensures that the conditions of Theorem 1 are satisfied. We consider here the Markov state dependent noise as it covers many applications of interest, encompasses the exogeneous scenario and as we shall see leads to general and verifiable conditions.

3. Markov state-dependent noise

In this section, we describe our stochastic approximation procedure with adaptive truncation sets and introduce the relevant notation required in the Markovian state dependent noise scenario (see [15, Section 6.6, p 159] for a detailed description and numerous examples). We first introduce a version without truncations of the algorithm in this setting (Subsection 3.2), and describe our adaptive procedure in terms of this plain algorithm in Subsection 3.1.

3.1. Non-homogeneous chain

Let $\rho = \{\rho_n\}$ be a monotone non-increasing sequence with $\rho_0 \leq 1$, define the product space $\bar{\mathsf{X}} := \mathsf{X} \cup \{x_c\} \times \bar{\Theta} := \Theta \cup \{\theta_c\}$, where $\theta_c \notin \Theta$ and $x_c \notin \mathsf{X}$ are two arbitrary cemetery points, and define the *non-homogeneous* Markov chain $\{Y_n^\rho := (X_n, \theta_n)\}$ on $\bar{\mathsf{X}} \times \bar{\Theta}$ as follows. Set $\theta_0 = \theta \in \Theta$, $X_0 = x \in \mathsf{X}$, and for $n \geq 0$, set

$$\theta_{n+1} = \theta_n + \rho_{n+1} H(\theta_n, X_{n+1}), X_{n+1} \sim P_{\theta_n}(X_n, \cdot) \qquad (6)$$

if $\theta_n \in \Theta$ and $\theta_{n+1} = \theta_c$ and $X_{n+1} = x_c$ if $\theta_n \notin \Theta$. Consider the following assumptions

(A2) For any $\theta \in \Theta$, the Markov kernel P_θ has a single stationary distribution π_θ, $\pi_\theta P_\theta = \pi_\theta$. In addition $H : \Theta \times \mathsf{X} \to \Theta$ is measurable, for all $\theta \in \Theta$, $\int_{\mathsf{X}} |H(\theta,x)| \pi_\theta(dx) < \infty$.

The existence and uniqueness of the invariant distribution can be guaranteed under classical irreducibility and recurrence conditions (see *e.g.* [16, Chapter 9,10]). We denote $h(\theta) := \int_{\mathsf{X}} H(\theta,x)\pi_\theta(dx)$ the mean-field associated to this stochastic approximation procedure and define the noise sequence $\{\xi_n = H(\theta_{n-1}, X_n) - h(\theta_{n-1})\}$. Following [2], we will often write $H_\theta(x)$ as an equivalent expression for $H(\theta,x)$, h_θ for $h(\theta)$, etc...

We denote $\mathscr{F} = \{\mathscr{F}_n, n \geq 0\}$ the natural filtration of this Markov chain, with $\mathscr{F}_n := \sigma((X_l,\theta_l), l \in \{0,\ldots,n\})$ and $\mathbb{P}^\rho_{x,\theta}$ the probability measure on the canonical space $\left((\mathsf{X}\times\Theta)^{\mathbb{N}}, (\mathscr{B}(\mathsf{X}) \otimes \mathscr{B}(\Theta))^{\otimes\mathbb{N}}\right)$ generated by the non-homogeneous Markov chain $\{Y_n^\rho\}$ started from the initial conditions $(X_0,\theta_0) = (x,\theta) \in \mathsf{X}\times\Theta$ and using the sequence ρ.

3.2. Homogeneous chain

Let $\{\mathscr{K}_q, q \geq 0\}$ be a sequence of compact subsets of Θ such that

$$\bigcup_{q\geq 0} \mathscr{K}_q = \Theta, \quad \text{and} \quad \mathscr{K}_q \subset \operatorname{int}(\mathscr{K}_{q+1}), \quad q \geq 0, \qquad (7)$$

where $\operatorname{int}(A)$ denotes the interior of set A. Let $\gamma = \{\gamma_k\}$ and $\varepsilon = \{\varepsilon_k\}$ be two monotone non-increasing sequences of positive numbers and let K be a subset of X. Let $\Phi : \mathsf{X}\times\Theta \to \mathsf{K}\times\mathscr{K}_0$ be a measurable function and $\phi : \mathbb{Z}^+ \to \mathbb{Z}$ be a function such that $\phi(k) > -k$ for any k. Our stochastic approximation algorithm with adaptive truncation sets is defined as an *homogeneous* Markov chain on $\mathsf{Z} := \mathsf{X}\times\Theta\times\mathbb{N}\times\mathbb{N}\times\mathbb{N}$

$$\{Z_n := (X_n, \theta_n, \kappa_n, \varsigma_n, \nu_n)\} \in \mathsf{Z}^{\mathbb{N}}, \qquad (8)$$

with the following transition at iteration $n+1$,

- If $\nu_n = 0$, then draw $(X_{n+1},\theta_{n+1}) \sim Q_{\gamma_{\varsigma_n}}(\Phi(X_n,\theta_n);\cdot)$; otherwise draw $(X_{n+1},\theta_{n+1}) \sim Q_{\gamma_{\varsigma_n}}(X_n,\theta_n;\cdot)$.
- If $|\theta_{n+1} - \theta_n| \leq \varepsilon_{\varsigma_n}$ and $\theta_{n+1} \in \mathscr{K}_{\kappa_n}$, then set: $\kappa_{n+1} = \kappa_n$, $\varsigma_{n+1} = \varsigma_n + 1$ and $\nu_{n+1} = \nu_n + 1$; otherwise, set $\nu_{n+1} = 0$, $\kappa_{n+1} = \kappa_n + 1$, $\varsigma_{n+1} = \varsigma_n + \phi(\nu_n)$.

In words, κ, ς and ν are counters: κ is the index of the current active truncation set; ν counts the number of iterations since the last reinitialization; ς is the current index in the sequences $\{\gamma_n\}$ and $\{\varepsilon_n\}$, and therefore defines the current proposal kernel Q_γ. The event $\{\nu_n = 0\}$ means that a reinitialization occurs and the condition on ϕ ensures that the algorithm is reinitialized with a value for γ_{ς_n} smaller than that used the last time such an event occurred. Various choices for the function ϕ can be considered. For example, the choice $\phi(k) = 1$ for all $k \in \mathbb{N}$ coincides with the procedure proposed in [7]: in this case $\varsigma_n = n$. Another sensible choice consists of setting $\phi(k) = 1 - k$ for all $k \in \mathbb{N}$, in which case the number of iterations between two successive reinitialisations is not taken into account.

The intuitive motivation for this modification of the original stochastic approximation recursion lies in Theorem 1. Indeed, in order to ensure the stability of the algorithm it is required that the sizesteps be not too large and that the average effect of the noise be small in order for the drift $h(\theta)$ to dominate, and confine the recursion to a compact set. The reprojections act as a -drastic- drift towards the center of Θ when $\{\theta_n\}$ grows too rapidly and allow one to reinitialize the algorithm with a smaller sizestep and weaker noise inside a "ring" of the type $\{\theta \in \Theta : w(\theta) \in (M_0, M_1]\}$ (M_0 and M_1 are defined in (A1)) where the drift is strictly positive. The fact that M_0 and M_1 are unknown *a priori* is the reason for the adaptive truncations, which ensure that one eventually selects $\mathscr{K}_q$ large enough in order to have $\mathscr{L} \cap \mathscr{K}_q \neq \varnothing$. As we shall see the limitation imposed on the increments of the sequence $\{\theta_n\}$ is required in order to ensure some type of homogeneity of the chain $\{\xi_n\}$, and therefore ergodicity properties of the noise sequence $\{\xi_n\}$.

We now introduce some further notation and briefly state our main result. For μ a probability on Z, we denote $\bar{\mathbb{P}}_\mu$ (resp. $\bar{\mathbb{E}}_\mu$) the probability (resp. the expectation) on the canonical space $(\mathsf{Z}^{\mathbb{N}}, \mathscr{B}(\mathsf{Z})^{\otimes \mathbb{N}})$ associated to the Markov chain $\{Z_n\}$ with initial distribution μ. For $z \in \mathsf{Z}$ we set $\bar{\mathbb{P}}_z := \bar{\mathbb{P}}_{\delta_z}$, $\bar{\mathbb{E}}_z := \bar{\mathbb{E}}_{\delta_z}$and for $(x,\theta) \in \mathsf{X} \times \Theta$

$$\bar{\mathbb{P}}_{x,\theta} := \bar{\mathbb{P}}_{x,\theta,0,0,0} \quad \text{and} \quad \bar{\mathbb{E}}_{x,\theta} := \bar{\mathbb{E}}_{x,\theta,0,0,0}. \tag{9}$$

This probability measure depends upon the deterministic sequences $\gamma = \{\gamma_n\}$ and $\varepsilon = \{\varepsilon_n\}$; this will be implicit hereafter in order to alleviate notation. We define recursively $\{T_n, n \geq 0\}$ the sequence of successive reinitialisation times

$$T_{n+1} = \inf\{k \geq T_n + 1,\ \nu_k = 0\}, \quad \text{with} \quad T_0 = 0, \tag{10}$$

where by convention $\inf\{\emptyset\} = \infty$. In the following sections we prove that under (A1), some regularity conditions on the family of transition probabilities $\{P_\theta, \theta \in \Theta\}$ and the sequences γ and ε then

$$\inf_{(x,\theta) \in \mathsf{K} \times \mathscr{K}_0} \bar{\mathbb{P}}_{x,\theta}\left(\sup_{n \geq 0} \kappa_n < \infty\right) = \inf_{(x,\theta) \in \mathsf{K} \times \mathscr{K}_0} \bar{\mathbb{P}}_{x,\theta}\left(\bigcup_{n=0}^{\infty} \{T_n = \infty\}\right) = 1,$$

i.e., the number of reinitializations of the procedure described above is finite $\bar{\mathbb{P}}_{x,\theta}$-a.e., for every $(x,\theta) \in \mathsf{K} \times \mathscr{K}_0$. Convergence will then follow using Theorem 2 for example.

4. Bound on $\bar{\mathbb{P}}_{x,\theta}(T_n < \infty)$

In this section we establish in Proposition 3 a bound on $\bar{\mathbb{P}}_{x,\theta}(T_n < \infty)$ in terms of the fluctuations of the noise sequence of the algorithm between successive reinitializations. Let $\mathscr{K}$ be a compact subset of Θ and let $\varepsilon = \{\varepsilon_n\}$ be a non-increasing sequence of positive numbers. We introduce $\sigma(\mathscr{K}, \varepsilon) = \sigma(\mathscr{K}) \wedge \nu(\varepsilon)$ where

$$\sigma(\mathscr{K}) = \inf\{k \geq 1, \theta_k \notin \mathscr{K}\},$$
$$\nu(\varepsilon) = \inf\{k \geq 1, |\theta_k - \theta_{k-1}| \geq \varepsilon_k\},$$

and for a sequence $\mathbf{a} = \{a_k\}$ and an integer l, we define $\mathbf{a}^{\leftarrow l} = \{a_k^{\leftarrow l}\}$ as $a_k^{\leftarrow l} = a_{k+l}$. Define, for any compact set $\mathscr{K} \subset \Theta$, $\varepsilon = \{\varepsilon_k\}$, $\rho = \{\rho_k\}$ and $1 \leq l \leq n$ the partial sum

$$S_{l,n}(\varepsilon, \rho, \mathscr{K}) := 1_{\{\sigma(\mathscr{K},\varepsilon) \geq n\}} \sum_{k=l}^{n} \rho_k (H(\theta_{k-1}, X_k) - h(\theta_{k-1})), \tag{11}$$

and for any $\delta \geq 0$ and any $M \in (M_0, M_1]$,

$$A(\delta, \varepsilon, M, \rho) := \sup_{\theta \in \mathscr{K}_0} \sup_{x \in \mathsf{K}} \left\{ \mathbb{P}^{\rho}_{\Phi(x,\theta)} \left[\sup_{k \geq 1} |S_{1,k}(\varepsilon, \rho, \mathscr{W}_M)| > \delta \right] + \mathbb{P}^{\rho}_{\Phi(x,\theta)} [\nu(\varepsilon) < \sigma(\mathscr{W}_M)] \right\}, \tag{12}$$

where $\mathscr{K}_0$ is defined in Eq. (7), $\mathscr{W}_M$, M_0 and M_1 are defined in (A1).

Proposition 3. *Assume (A1) and that $\mathscr{K}_0 \subset \mathscr{W}_{M_0}$ (where M_0 is defined in (A1)). Then for any $M \in (M_0, M_1]$ there exist an integer n_0 and a constant $\delta_0 > 0$ such that, for any $n > n_0$, we have*

$$\sup_{(x,\theta) \in \mathsf{K} \times \mathscr{K}_0} \bar{\mathbb{P}}_{x,\theta}[T_n < \infty] \leq \prod_{l=n_0}^{n-1} \sup_{q \geq l} A(\delta_0, \varepsilon^{\leftarrow q}, M, \gamma^{\leftarrow q}),$$

where T_n is defined in Eq. (10).

Corollary 4. *Assume (A1) and that $\mathscr{K}_0 \subset \mathscr{W}_{M_0}$ (where M_0 is defined in (A1)). Then for any $M \in (M_0, M_1]$ and $n \geq n_0$, there exists a constant $C < \infty$ such that for any $m \geq n$,*

$$\bar{\mathbb{P}}_{x,\theta}\left[\sup_{k \geq 1} \kappa_k \geq m\right] \leq C \left(\sup_{q \geq n} A(\delta_0, \varepsilon^{\leftarrow q}, M, \gamma^{\leftarrow q})\right)^m,$$

where $\{\kappa_k\}$ is the counter corresponding to the number of reinitialisations defined in Eq. (8).

5. Control of the fluctuations

Our aim is now to find a bound for $A(\delta,\varepsilon,M,\rho)$ defined in Eq. (12), which requires the following conditions to hold. Define, for $V:\mathsf{X}\to[1,\infty)$ and $g:\mathsf{X}\to\mathbb{R}^{n_\theta}$ the norm

$$\|g\|_V=\sup_{x\in\mathsf{X}}\frac{|g(x)|}{V(x)}. \tag{13}$$

Consider the following assumptions

(**A3**) For any $\theta\in\Theta$, the Poisson equation $g-P_\theta g=H_\theta-\pi_\theta(H_\theta)$ has a solution g_θ. There exist a function $W:\mathsf{X}\to[1,\infty]$ such that $\{x\in\mathsf{X},W(x)<\infty\}\neq\emptyset$, constants $\alpha\in(0,1]$, $p\geq 2$ such that for any compact subset $\mathscr{K}\subset\Theta$,

- (i) $\sup_{\theta\in\mathscr{K}}\|H_\theta\|_W<\infty$, $\sup_{\theta\in\mathscr{K}}\|g_\theta\|_W<\infty$, $\sup_{\theta\in\mathscr{K}}\|P_\theta g_\theta\|_W<\infty$, and
$$\sup_{(\theta,\theta')\in\mathscr{K}}|\theta-\theta'|^{-\alpha}\times\{\|g_\theta-g_{\theta'}\|_W+\|P_\theta g_\theta-P_{\theta'}g_{\theta'}\|_W\}<\infty.$$
- (ii) there exist constants $\{C_k,k\geq 0\}$ such that, for any $k\in\mathbb{N}$, for any sequence $\rho=\{\rho_k\}$ and for any $x\in\mathsf{X}$,
$$\sup_{\theta\in\mathscr{K}}\mathbb{E}^{\rho}_{x,\theta}[W^p(X_k)1_{\{\sigma(\mathscr{K})\geq k\}}]\leq C_kW^p(x),$$
- (iii) there exist $\varepsilon>0$ and a constant C such that for any sequence $\rho=\{\rho_k\}$ and for any $x\in\mathsf{X}$,
$$\sup_{\theta\in\mathscr{K}}\mathbb{E}^{\rho}_{x,\theta}[W^p(X_k)1_{\{\sigma(\mathscr{K})\wedge\nu_\varepsilon\geq k\}}]\leq CW^p(x).$$
where $\nu_\varepsilon=\inf\{k\geq 1,|\theta_k-\theta_{k-1}|>\varepsilon\}$.

Proposition 5. *Assume (A3). Let $\mathscr{K}$ be a compact subset of Θ and let $\rho=\{\rho_k\}$ and $\varepsilon=\{\varepsilon_k\}$ be two non-increasing sequences of positive numbers such that $\lim_{k\to\infty}\varepsilon_k=0$. Then, for p as defined in (A3),*

1. *There exists a constant C such that, for any $(x,\theta)\in\mathsf{X}\times\mathscr{K}$ and any integer l, any $\delta>0$*
$$\mathbb{P}^{\rho}_{x,\theta}\left(\sup_{n\geq l}|S_{l,n}(\varepsilon,\rho,\mathscr{K})|\geq\delta\right)\leq C\delta^{-p}\left\{\left(\sum_{k=l}^{\infty}\rho_k^2\right)^{p/2}+\left(\sum_{k=l}^{\infty}\rho_k\varepsilon_k^{\alpha}\right)^{p}\right\}W^p(x). \tag{14}$$
2. *There exists a constant C such that, for any $(x,\theta)\in\mathsf{X}\times\mathscr{K}$,*
$$\mathbb{P}^{\rho}_{x,\theta}(\nu(\varepsilon)<\sigma(\mathscr{K}))\leq C\left\{\sum_{k=1}^{\infty}(\varepsilon_k^{-1}\rho_k)^p\right\}W^p(x). \tag{15}$$

(**A4**) The sequences $\gamma=\{\gamma_k\}$ and $\varepsilon=\{\varepsilon_k\}$ are non-increasing, positive, and satisfy, $\sum_{k=0}^{\infty}\gamma_k=\infty$, $\lim_{k\to\infty}\varepsilon_k=0$ and
$$\sum_{k=1}^{\infty}\left\{\gamma_k^2+\gamma_k\varepsilon_k^{\alpha}+(\varepsilon_k^{-1}\gamma_k)^p\right\}<\infty,$$
where p and α are defined in (A3).

Theorem 6. *Assume (A1) to (A4). Then, for any subset $\mathsf{K}\subset\mathsf{X}$ such that $\sup_{x\in\mathsf{K}}W(x)<\infty$, $\mathscr{K}_0\subset\mathscr{W}_{M_0}$ (where M_0 is defined in (A1)) and any $\rho\in(0,1)$, there exists a constant $C<\infty$ such that, for all $(x,\theta)\in\mathsf{X}\times\Theta$,*

$$\bar{\mathbb{P}}_{x,\theta}\left[\sup_{n\geq 1}\kappa_n\geq k\right]\leq C\rho^k.$$

Hence, under the stated conditions, the tail probability of the number of reinitialization decreases faster than any exponential and $\sup_{n\geq 1}\kappa_n$ is finite $\bar{\mathbb{P}}_{x,\theta}$-a.s. Combining this result with Theorem 2, it is possible to obtain the following global convergence result.

Theorem 7. *Assume (A1) to (A4). Let $\mathsf{K}\subset\mathsf{X}$ be such that $\sup_{x\in\mathsf{K}}W(x)<\infty$ and that $\mathscr{K}_0\subset\mathscr{W}_{M_0}$ (where M_0 is defined in (A1)), and let $\{Z_n\}$ be as defined by Eq. (8). Then, for all $(x,\theta)\in\mathsf{X}\times\Theta$, we have $\lim_{k\to\infty}d(\theta_k,\mathscr{L})=0$, $\bar{\mathbb{P}}_{x,\theta}$-a.s.*

6. Drift conditions

In this section, we give conditions which imply (A3) in terms of a minorisation of the Markov kernel on a small set and a drift condition toward this small set (see [16] for the definitions and main results). Denote, for $V:\mathsf{X}\to[1,\infty)$, $\mathscr{L}_V:=\{g:\mathsf{X}\to\mathbb{R}^{n_\theta},\sup_{x\in\mathsf{X}}\|g\|_V<\infty\}$ where $\|\cdot\|_V$ is defined in Eq. (13).

(**DRI**) For any $\theta\in\Theta$, P_θ is ψ-irreducible and aperiodic [1]. In addition there exist a function $V:\mathsf{X}\to[1,\infty)$, constants $p\geq 2$ and $\beta\in[0,1]$ such that for any compact subset $\mathscr{K}\subset\Theta$,

- (DRI1) there exist constants $0<\lambda<1$, b, κ, $\delta>0$ and a probability measure ν such that $\forall x\in\mathsf{C}$, $\forall A\in\mathscr{B}(\mathsf{X})$,
$$\sup_{\theta\in\mathscr{K}}P_\theta V^p(x)\leq\lambda V^p(x)+b1_{\mathsf{C}}(x) \tag{16}$$
$$\inf_{\theta\in\mathscr{K}}P_\theta(x,A)\geq\delta\nu(A). \tag{17}$$
- (DRI2) there exists C such that, for all $x\in\mathsf{X}$,
$$\sup_{\theta\in\mathscr{K}}|H_\theta(x)|\leq CV(x),$$
$$\sup_{(\theta,\theta')\in\mathscr{K}}|\theta-\theta'|^{-\beta}|H_\theta(x)-H_{\theta'}(x)|\leq CV(x).$$

[1] We use in this article the standard terminology and the notations introduced in [16, Chapter 4,5]

(DRI3) there exists C such that, for all $(\theta, \theta') \in \mathscr{K} \times \mathscr{K}$,

$$\|P_\theta g - P_{\theta'} g\|_V \le C \, \|g\|_V \, |\theta - \theta'|^\beta \qquad \forall g \in \mathscr{L}_V, \tag{18}$$

$$\|P_\theta g - P_{\theta'} g\|_{V^p} \le C \, \|g\|_{V^p} \, |\theta - \theta'|^\beta, \quad \forall g \in \mathscr{L}_{V^p}. \tag{19}$$

Assumption (DRI1) is classical in the Markov chain literature; it implies the existence of a stationary distribution π_θ for all $\theta \in \Theta$ and V^p-uniform ergodicity, *i.e.* for each $\theta \in \Theta$ there exist constants $C_\theta < \infty$ and $\rho_\theta \in [0,1)$, such that for any function $f \in \mathscr{L}_{V^p}$ and any integer $k > 0$

$$\|P_\theta^k f - \pi_\theta(f)\|_{V^p} \le C_\theta \rho_\theta^k \|f\|_{V^p}.$$

Note that the constants C_θ and ρ_θ may be bounded over the compact sets of Θ, *i.e.* for each $\mathscr{K} \subset \Theta$, there exists $\bar{C} < \infty$ and $\bar{\rho} \in [0,1)$, such that $\sup_{\theta \in \mathscr{K}} C_\theta \le \bar{C}$ and $\sup_{\theta \in \mathscr{K}} \rho_\theta \le \bar{\rho}$. The regularity of the kernels $\theta \to P_\theta$ expressed in V and V^p norm is naturally less classical. The main result of this section is:

Proposition 8. *Assume (DRI). Then (A2) and (A3) are satisfied and for any* $0 < \alpha < \beta$*,*

$$\sup_{(\theta,\theta') \in \mathscr{K} \times \mathscr{K}} |\theta - \theta'|^{-\alpha} \, |h(\theta) - h(\theta')| < \infty. \tag{20}$$

References

[1] J. Abounadi, D. Bertsekas, and V. Borkar. Stochastic approximation for nonexpansive maps: Application to q-learning algorithms. *SIAM J. on Control and Optimization*, 41:1–22, 2002.

[2] A. Benveniste, M. Métivier, and P. Priouret. *Adaptive Algorithms and Stochastic Approximations*. Springer-Verlag, New York, 1990.

[3] V. Borkar. Stability of annealing schemes and related processes. *Systems Control Lett.*, 41:325–331, 2000.

[4] V. Borkar and S. Meyn. The o.d.e. method for convergence of stochastic approximation and reinforcement learning. *SIAM J. Control and Optimization*, 38:447–469, 2000.

[5] R. Buche and H. Kushner. Rate of convergence for constrained stochastic approximation algorithms. *SIAM J. on Control and Optimization*, 40:1011–1041, 2001.

[6] H. Chen. Stochastic approxmation with state-dependent noise. *Science in China Series E*, 43(5):531–541, 2000.

[7] H. Chen, L. Guo, and A. Gao. Convergence and robustness of the Robbins- Monro algorithm truncated at randomly varying bounds. *Stochastic Processes and their Applications*, 27:217–231, 1988.

[8] H. Chen and Y.-M. Zhu. Stochastic approximation procedures with randomly varying truncations. *Scientia Sinica 1*, 29:914–926, 1986.

[9] H.-F. Chen. *Stochastic approximation and its applications*, volume 64 of *Nonconvex Optimization and its Applications*. Kluwer Academic Publishers, Dordrecht, 2002.

[10] B. Delyon. Stochastic approximation with decreasing gain: Convergence and asymptotic theory. Technical report, Universit de Rennes, 2000.

[11] B. Delyon, M. Lavielle, and E. Moulines. On a stochastic approximation version of the EM algorithm. *Annals of Statistics*, 27(1), 1999.

[12] M. Duflo. *Random Iterative Systems*. Applications of mathematics. Springer-Verlag, 1997.

[13] L. Gerencer. Rate of convergence of recursive estimators. *SIAM J. on Control and Optimization*, 30:1200–1227, 1992.

[14] H. Kushner and D. Clark. *Stochastic Approximation for Constrained and Unconstrained Systems*. Springer-Verlag, Berlin, Heidelberg, 1978.

[15] H. Kushner and G. Yin. *Stochastic Approximation Algorithms and Applications*. Applications of Mathematics. Springer-Verlag, New-York, 1997.

[16] S. Meyn and R. Tweedie. *Markov Chains and Stochastic Stability.* Communication and Control Engineering series. Springer-Verlag, London, 1993.

[17] V. Tadic. Stochastic gradient with random truncations. *European J. of Operational Research*, 101:261–284, 1997.

[18] V. Tadic. Stochastic approximations with random truncations, state dependent noise and discontinuous dynamics. *Stochastics and Stochastics reports*, 64:283–326, 1998.