Stochastic approximation with set-valued maps and Markov noise : *Convergence analysis and applications*

Vinayaka G. Yaji and Shalabh Bhatnagar

Department of Computer Science and Automation,
Indian Institute of Science.

vgyaji@gmail.com — 9902584170

Introduction

Originally conceived as a tool for statistical computation, stochastic approximation schemes today find applications in varied fields of engineering, such as, adaptive signal processing, stochastic optimization, game theory and machine learning to name a few. In control systems engineering too, stochastic approximation is the main paradigm for on-line algorithms for system identification and adaptive control. Stochastic approximation schemes are designed to operate in uncertain environments and are iterative in nature which are the key traits that make them attractive for adaptive schemes.

In our work we investigate the asymptotic behavior of stochastic approximation schemes where the drift function is set-valued depending additionally on an iterate-dependent Markov noise term. Such schemes arise for example, when one is trying to minimize a convex function which is an expectation and is estimated via Markov chain Monte Carlo methods. We adopt the dynamical systems approach to analyze the asymptotic behavior of such schemes. We consider two variants of the same, namely:

(1) *Single time scale stochastic recursive inclusion (SRI) with Markov noise:*
  - Convergence analysis involves showing that the iterates track the flow of a differential inclusion and then invoking the limit-set theorem to characterize the limit sets of the recursion in terms of the dynamics of the differential inclusion.
  - Applications include controlled stochastic approximation, subgradient descent, approximate drift problem and analysis of discontinuous dynamics, all in the the presence of Markov noise.

(2) *Two time scale SRI with Markov noise:*
  - Iterates are updated along a slower and faster timescale induced due to a clever choice of step size regimes.
  - Slower timescale iterates appear static w.r.t the faster timescale and faster timescale iterates appear to have equilibrated w.r.t. the slower timescale recursion.
  - Applications include actor-critic algorithms in reinforcement learning and solving nested optimization problems, for example: computation of saddle points.

In the analysis above, the iterates are assumed to lie in compact set which is sample path dependent. This *stability* assumption is essential but is often hard to verify. We also consider the *analysis of standard SRI in the absence of a stability guarantee*. Our contributions are:

- Extension of the lock-in probability bound to the case with set-valued maps.
- Using the lock-in probability result we show that a feedback mechanism which involves resetting the iterates at regular time intervals, stabilizes the recursion when the mean field possesses a global attractor.

Single timescale SRI with Markov noise [1]

- *Recursion and assumptions:* Given $X_0 \in \mathbb{R}^d$, $S_0 \in \mathcal{S}$, a compact metric space, X_n's are generated according to the recursion

$$X_{n+1} - X_n - a(n)M_{n+1} \in a(n)H(X_n, S_n), \quad (1)$$

and, $\mathbb{P}(S_{n+1} \in A | X_m, S_m, m \le n) = \Pi(X_n, S_n)(A)$, for every $A \subseteq \mathcal{S}$ measurable, where,

(A1) $H : \mathbb{R}^d \times \mathcal{S} \to \{\text{subsets of } \mathbb{R}^d\}$ such that $H(x, s)$ is non-empty, convex and compact; $\sup_{z \in H(x,s)} \|z\| \le K(1 + \|x\|)$ and the set-valued map H, has a closed graph.

(A2) $\Pi : \mathbb{R}^d \times \mathcal{S} \to \mathcal{P}(\mathcal{S})$ is continuous, where $\mathcal{P}(S)$ denotes the metric space of probability measures on $\mathcal{S}$ with the metric of weak convergence.

(A3) $\{a(n)\}_{n\ge0}$ is the step size sequence satisfying the standard *Robbins-Munro* conditions.

(A4) $\{M_n\}_{n\ge1}$ denotes the additive noise terms satisfying a condition which ensures that their eventual contribution is negligible over any time interval.

(A5) The iterates remain stable, i.e., $\mathbb{P}(\sup_{n\ge0} \|X_n\| < \infty) = 1$.

- *Mean field and its properties:* As a consequence of (A2) and the compactness of $\mathcal{S}$, the Markov chain with transition probability kernel $\Pi(x, \cdot)(\cdot)$ admits at least one stationary distribution for every x. Let $D(x)$ denote the set of stationary distributions for every x. The differential inclusion (DI) whose flow the iterates generated as in recursion (1) are expected track is given by,

$$\frac{dx}{dt} \in \hat{H}(x) := \cup_{\mu \in D(x)} \int_{\mathcal{S}} H(x, s)\mu(ds), \quad (2)$$

  - The set-valued map $H(x, \cdot)$ is a measurable set-valued map. For measurable set-valued maps there exists a measurable selection, i.e., $h : \mathcal{S} \to \mathbb{R}^d$ s.t., $h(s) \in H(x, s)$ for every $s \in \mathcal{S}$.
  - As a consequence of (A1), we have that for every $\mu \in D(x)$, every measurable selection is integrable and the set valued integral is defined as the collection of the integrals of all the measurable selections (a.k.a. Aumann's integral).
  - Further we show that the set-valued map $\hat{H}$, is of the Marchaud type, which guarantees the existence of solutions for DI (2).

- *Convergence guarantee:* Given that there exists $A \subseteq \mathbb{R}^d$, a global attractor for the flow of DI (2), then the iterates X_n as in recursion (1) converge to A almost surely.
  - A global attractor is set to which all the solutions of DI (2) converge to from any initial condition.
  - The main component of the proof of the above convergence result is an *asymptotic pseudotrajectory* argument, which involves showing that every limit point the tail of the linearly interpolated trajectory is a solution of DI (2).

- *Application:* Let $\mu \in \mathcal{P}(\mathcal{S})$ and $J_\mu : \mathbb{R}^d \to \mathbb{R}$, such that $J_\mu(x) := \int_{\mathcal{S}} J(x, s)\mu(ds)$, where for every $s \in \mathcal{S}$, $J(\cdot, s)$ is convex. Let Π be as in (A2) s.t., $D(x) = \{\mu\}$. Then iterates in recursion (1) with $H(x, s) := -\partial J(x, s) := \{-g : \forall y,\ J(y, s) \ge J(x, s) + \langle g, y - x\rangle\}$, track the flow of DI (2). The set-valued map $\hat{H}$ satisfies the inclusion, $\hat{H}(x) \subseteq -\partial J_\mu(x)$. Hence a global attractor of DI, $\frac{dx}{dt} \in -\partial J_\mu(x)$, is also a global attractor of DI (2) and our convergence result guarantees that the iterates converge to set of points $\{x^* : 0 \in \partial J_\mu(x^*)\}$.

Two timescale SRI with Markov noise [2]

- *Recursion and assumptions:* Given $X_0 \in \mathbb{R}^{d_1}, Y_0 \in \mathbb{R}^{d_2}$, and for every $i \in \{1, 2\}$, $S_0^{(i)} \in \mathcal{S}^{(i)}$, a compact metric space, X_n, Y_n are generated according to the recursion,

$$Y_{n+1} - Y_n - b(n)M^{(2)}_{n+1} \in b(n)H_2(X_n, Y_n, S^{(2)}_n), \quad (3a)$$

$$X_{n+1} - X_n - a(n)M^{(1)}_{n+1} \in a(n)H_1(X_n, Y_n, S^{(1)}_n), \quad (3b)$$

and for every $i \in \{1, 2\}$, $\mathbb{P}(S^{(i)}_{n+1} \in A | X_m, Y_m, S^{(i)}_m, m \le n) = \Pi^{(i)}(X_n, Y_n, S^{(i)}_n)(A)$, for every $A \subseteq \mathcal{S}^{(i)}$ measurable, where,

(B1) for every $i \in \{1, 2\}$, $H_i : \mathbb{R}^{d_1+d_2} \times \mathcal{S}^{(i)} \to \{\text{subsets of } \mathbb{R}^{d_i}\}$ such that $H(x, y, s^{(i)})$ is non-empty, convex and compact; $\sup_{z \in H_i(x,y,s^{(i)})} \|z\| \le K(1 + \|x\| + \|y\|)$ and the set-valued map H_i, has a closed graph.

(B2) for every $i \in \{1, 2\}$, $\Pi^{(i)} : \mathbb{R}^{d_1+d_2} \times \mathcal{S}^{(i)} \to \mathcal{P}(\mathcal{S}^{(i)})$ is continuous.

(B3) $\{a(n)\}_{n\ge0}$ and $\{b(n)\}_{n\ge0}$ are step size sequences satisfying $\lim_{n\to\infty} \frac{b(n)}{a(n)} = 0$, in addition to the standard *Robbins-Munro* conditions.

(B4) for every $i \in \{1, 2\}$, $\{M^{(i)}_n\}_{n\ge1}$ denotes the additive noise terms satisfying a condition which ensures that their eventual contribution is negligible over any time interval.

(B5) the iterates remain stable, i.e., $\mathbb{P}\left(\sup_{n\ge0} (\|X_n\| + \|Y_n\|) < \infty\right) = 1$.

- *Overview of the analysis:*
  - Assumption *(B3)* tells that eventually the time step taken by recursion (3a) is smaller than the time step taken by recursion (3b). Hence recursion (3a) is called the slower timescale recursion and the recursion (3b) is called the faster timescale recursion.
  - With respect to the faster timescale (3b), the slower timescale recursion (3a) appears to be static and one would expect that the family of DIs,

$$\frac{dx}{dt} \in \hat{H}_1(x, y_0) := \cup_{\mu \in D^{(1)}(x,y_0)} \int_{\mathcal{S}^{(1)}} H_1(x, y_0, s^{(1)})\mu(ds^{(1)}), \quad (4)$$

  obtained by fixing some $y_0 \in \mathbb{R}^{d_2}$ to describe the behavior of the faster timescale recursion (3b).

(B6) For every $y \in \mathbb{R}^{d_2}$, let $\lambda(y)$ denote the global attractor of DI (4) with $y_0 = y$. We assume that the set-valued map, $y \to \lambda(y)$ is upper semicontinuous and possesses the linear growth property.

  - With respect to the slower timescale recursion (3a), the faster time scale recursion will appear to have equilibrated. Further the Markov noise terms average the set-valued drift function H_2 with respect to the stationary distributions. The set-valued map that the slower timescale recursion is expected to track which captures both the equilibration of the faster timescale and the averaging by the Markov noise terms, is given by,

$$\frac{dy}{dt} \in \hat{H}_2(y) := \cup_{\mu \in D(y)} \int_{\mathbb{R}^{d_1} \times \mathcal{S}^{(2)}} H_2(x, y, s^{(2)})\mu(dx, ds^{(2)}), \text{ where,} \quad (5)$$

$$D(y) := \Big\{ \mu \in \mathcal{P}(\mathbb{R}^{d_1} \times \mathcal{S}^{(2)}) : \mathrm{supp}(\mu_{\mathbb{R}^{d_1}}) \subseteq \lambda(y) \text{ and for every } A \in \mathscr{B}(\mathcal{S}^{(2)}),$$
$$\mu_{\mathcal{S}^{(2)}}(A) = \int_{\mathcal{S}^{(2)}} \Pi^{(2)}(x, y, s^{(2)})(A)\mu(dx, ds^{(2)}) \Big\}. \quad (6)$$

- *Convergence guarantee:* If $\mathcal{Y} \subseteq \mathbb{R}^{d_2}$ is a global attractor of the flow of DI (5), then the iterates (X_n, Y_n) converge to $\cup_{y\in\mathcal{Y}} (\lambda(y) \times \{y\})$ almost surely as $n \to \infty$.

Analysis of SRI in the absence of a stability guarantee [3]

- *Recursion and assumptions:* Given $X_0 \in \mathbb{R}^d$, X_n's are generated according to the recursion,

$$X_{n+1} - X_n - a(n)M_{n+1} \in a(n)H(X_n), \text{ where,} \quad (7)$$

(C1) $H : \mathbb{R}^d \to \left\{\text{subsets of } \mathbb{R}^d\right\}$ is a Marchaud map.

(C2) $\{a(n)\}_{n\ge0}$ is the step size sequence satisfying the standard *Robbins-Munro* conditions.

(C3) $\{M_n\}_{n\ge1}$ is an $\mathbb{R}^d$-valued, martingale difference sequence with respect to the filtration $\{\mathscr{F}_n := \sigma(X_m, M_m,\ m \le n)\}$. Furthermore, $\{M_n\}_{n\ge1}$ are such that, $\|M_{n+1}\| \le K(1 + \|x_n\|)$ *a.s.*, for every $n \ge 0$, for some constant $K > 0$.

- *Lock-in probability bound:* We prove that in the absence of stability, there exists $\tilde{K} > 0$, such that for large enough $n_0 \in \mathbb{N}$, the probability of converging to a attractor A, given that at iteration n_0 we observe the iterate in $\mathcal{O}' \subseteq \mathbb{R}^d$, a precompact neighborhood of the attractor, is lower bounded by $1 - 2de^{-\tilde{K}/b(n_0)}$, where $\{b(n)\}$ is a positive sequence of real numbers converging to zero. Formally,

$$\mathbb{P}(X_n \to A \text{ as } n \to \infty | X_{n_0} \in \mathcal{O}') \ge 1 - 2de^{-\tilde{K}/b(n_0)}.$$

- *Applications:* We use the above bound to obtain alternate condition for convergence to a local attractor in the absence of stability. Further we develop a stabilization mechanism involving resetting the iterates at regular time intervals which stabilizes the scheme when the mean field possess a global attractor. The key idea involves lower bounding the probability of having no future resets given that n_0 number of resets have occurred by $1 - 2de^{-\tilde{K}/b(n_0)}$.

References

[1] V. Yaji and S. Bhatnagar, "Stochastic recursive inclusions with non-additive iterate-dependent Markov noise," *arXiv preprint arXiv:1607.04735*, 2016.

[2] ——, "Stochastic recursive inclusions in two timescales with non-additive iterate dependent Markov noise," *CoRR*, vol. abs/1611.05961, 2016.

[3] ——, "Analysis of stochastic approximation schemes with set-valued maps in the absence of a stability guarantee and their stabilization," *CoRR*, vol. abs/1701.07590, 2017.

Stochastic approximation with set-valued maps and Markov noise

Vinayaka Yaji

March 14, 2017

Department of Computer Science and Automation,
Indian Institute of Science.
Email: vgyaji@gmail.com
Advisor: Dr. Shalabh Bhatnagar

Introduction

- Borkar[1] analyzed the iterative scheme of the form

$$X_{n+1} - X_n - a(n)M_{n+1} = a(n)h(X_n, S_n),$$

  where, $a(n)$ is the step size, M_{n+1} is the martingale noise, h is the drift function and S_n is the Markov noise with $S_n \in \mathcal{S}$, a compact metric space.

- $\mathbb{P}(S_{n+1} \in A | S_m, X_m, m \leq n) = \Pi(S_n, X_n)(A)$; where $\Pi : \mathbb{R}^d \times \mathcal{S} \to \mathcal{P}(\mathcal{S})$ assumed to be continuous;
- Markov chain with TPK $\Pi(\cdot, x)(\cdot)$ admits at least one stationary distribution for every x; $D(x)$.

[1] Vivek S Borkar. "Stochastic approximation with controlled Markov noise". In: *Systems & control letters* 55.2 (2006), pp. 139–145.

- Under the additional assumption of stability ($\sup_{n\geq 0} \|X_n\| < \infty$), iterates track the flow of the differential inclusion (DI),

$$\frac{dx}{dt} \in \hat{h}(x), \tag{1}$$

where $\hat{h}(x) := \cup_{\mu \in D(x)} \int_{\mathcal{S}} h(x,s)\mu(ds)$.
- Convergence guarantee: If $x^* \in \mathbb{R}^d$ is globally asymptotically stable for DI (1), then $X_n \to x^*$, *a.s.*

Set-valued drift function: Motivating cases

- Controlled stochastic approximation

$$X_{n+1} - X_n - a(n)M_{n+1} = a(n)h(X_n, S_n, U_n), \qquad (2)$$

where U_n is the control variable taking value in a compact metric space $\mathcal{U}$.
Recursion (2) can be equivalently written as,

$$X_{n+1} - X_n - a(n)M_{n+1} \in a(n)H(X_n, S_n),$$

where, $H(x, s) := \text{conv}\left(\{h(X_n, S_n, u) \ : \ u \in \mathcal{U}\}\right)$.

- Subgradient descent:
  - Assume that for every x, $D(x) = \mu \in \mathcal{P}(\mathcal{S})$.
  - Let $J : \mathbb{R}^d \times \mathcal{S} \to \mathbb{R}$ s.t., $J(\cdot, s)$ is convex.
  - Consider the recursion given by,

    $$X_{n+1} - X_n - a(n)M_{n+1} \in -a(n)\partial J(X_n, S_n),$$

    where,
    $\partial J(x, s) := \{g \;:\; J(y, s) \geq J(x, s) + \langle g, y - x \rangle, \text{ for every } y\}$.
  - Does the above recursion minimize the convex function $x \to \int_S J(x, s)\mu(ds)$?

Stochastic recursive inclusions (SRI) with Markov noise[2]

- We analyze the recursion given by,

$$X_{n+1} - X_n - a(n)M_{n+1} \in a(n)H(X_n, S_n), \quad (3)$$

where, $H : \mathbb{R}^d \times \mathcal{S} \to \{\text{subsets of } \mathbb{R}^d\}$ is such that, $H(x, s)$ is non-empty, convex and compact; H is upper semicontinuous and $\|H(x, s)\| \leq K(1 + \|x\|)$.

[2] Vinayaka Yaji and Shalabh Bhatnagar. "Stochastic Recursive Inclusions with Non-Additive Iterate-Dependent Markov Noise". In: *arXiv preprint arXiv:1607.04735* (2016).

- Conversion to the single-valued case:

$$H(x,s) \subseteq \underbrace{H^{(l)}(x,s)}_{\substack{\text{continuous,}\\ \cap_{l\geq 1} H^{(l)}(x,s)=H(x,s)}} \rightarrow \underbrace{h^{(l)}(x,s,u)}_{\{h^{(l)}(x,s,u):\|u\|\leq 1\}=H^{(l)}(x,s)}$$

Using the above, recursion (3) is written as,

$$X_{n+1} - X_n - a(n)M_{n+1} = a(n)h^{(l)}(X_n, S_n, U_n).$$

- Identifying the limiting trajectory

$$\mathbf{x}^*(t) = \int_0^t \left[\int_{S\times U} h^{(l)}(x,s,u)\gamma^{(l)}(t;ds,du) \right] dt,$$

where, $\gamma^{(l)} : [0,\infty) \rightarrow \mathcal{P}(\mathcal{S} \times U)$ and for every $t \geq 0$, $\gamma_{\mathcal{S}}^{(l)}(t) \in D(\mathbf{x}(t))$.

- The iterates X_n track the flow of the DI[3],

$$\frac{dx}{dt} \in \hat{H}(x) := \cup_{\mu \in D(x)} \int_{\mathcal{S}} H(x, s)\mu(ds). \tag{4}$$

- Convergence guarantee: If x^* is a global attractor of DI (4), then $X_n \to x^*$ a.s.

[3] Shoumei Li, Yukio Ogura, and Vladik Kreinovich. *Limit theorems and applications of set-valued and fuzzy set-valued random variables*. Vol. 43. Springer Science & Business Media, 2013.

Two timescale SRI with Markov noise[4]

- We analyze the recursion given by,

$$Y_{n+1} - Y_n - b(n)M^{(2)}_{n+1} \in b(n)H_2(X_n, Y_n, S^{(2)}_n), \qquad (5a)$$

$$X_{n+1} - X_n - a(n)M^{(1)}_{n+1} \in a(n)H_1(X_n, Y_n, S^{(1)}_n), \qquad (5b)$$

  where $a(n), b(n)$ are step-sizes satisfying $\lim_{n\to\infty} \frac{b(n)}{a(n)} = 0$.
- (5a) will appear to be static w.r.t. (5b).
- (5b) will appear to have equilibrated w.r.t (5a).

[4] Vinayaka Yaji and Shalabh Bhatnagar. "Stochastic Recursive Inclusions in two timescales with non-additive iterate dependent Markov noise". In: *CoRR* abs/1611.05961 (2016).

- $\lambda(y)$ denotes the globally attracting set for the flow of DI,

$$\frac{dx}{dt} \in \hat{H}_1(x, y)$$

where, $\hat{H}_1(x, y) = \cup_{\mu \in D^{(1)}(x,y)} \int_{\mathcal{S}^{(1)}} H_{1,(x,y)}(s^{(1)})\mu(ds^{(1)})$.

- $\mathcal{Y}$ denotes the globally attracting set for the flow of DI,

$$\frac{dx}{dt} \in \hat{H}_2(y),$$

where $\hat{H}_2(y) := \cup_{\mu \in D(y)} \int_{\mathbb{R}^{d_1} \times \mathcal{S}^{(2)}} H_{2,y}(x, s^{(2)})\mu(dx, ds^{(2)})$, with

$$D(y) := \Big\{ \mu \in \mathcal{P}(\mathbb{R}^{d_1} \times \mathcal{S}^{(2)}) : \text{supp}(\mu_{\mathbb{R}^{d_1}}) \subseteq \lambda(y), \\ \mu_{\mathcal{S}^{(2)}}(A) = \int_{\mathcal{S}^{(2)}} \Pi^{(2)}(x, y, s^{(2)})(A)\mu(dx, ds^{(2)}) \Big\}$$

- Convergence guarantee: Under stability, $(X_n, Y_n) \to \cup_{y \in \mathcal{Y}}(\lambda(y) \times \{y\})$ a.s.

Analysis in the absence of a stability guarantee[5]

- We consider the recursion given by,

$$X_{n+1} - X_n - a(n)M_{n+1} \in a(n)H(X_n)$$

- In the absence of a stability guarantee, we show that,

$$\mathbb{P}(X_n \to A \text{ as } n \to \infty | X_{n_0} \in \mathcal{O}') \geq 1 - 2de^{-\tilde{K}/b(n_0)}$$

  where A is local attractor, $\mathcal{O}'$ is precompact neighborhood of A and $b(n_0) \to 0$ as $n_0 \to \infty$.

- Using the above we design a feedback mechanism, which stabilizes the scheme in the presence of a global attractor, thus guaranteeing convergence.

[5] Vinayaka G. Yaji and Shalabh Bhatnagar. "Analysis of stochastic approximation schemes with set-valued maps in the absence of a stability guarantee and their stabilization". In: *CoRR* abs/1701.07590 (2017).

Questions ?