Almost Sure Convergence of Nonlinear Stochastic Approximation Under General Moment Conditions

Anh Duc Nguyen
Institute of Operations Research and Analytics
National University of Singapore
ducna@nus.edu.sg

Quang Nguyen
Department of Computer Science
University of Information Technology
Vietnam National University HCMC
23521285@gm.uit.edu.vn

Hoang Huy Nguyen
H. Milton Stewart School of Industrial
& Systems Engineering
Georgia Institute of Technology
hnguyen455@gatech.edu

Siva Theja Maguluri
H. Milton Stewart School of Industrial
& Systems Engineering
Georgia Institute of Technology
siva.theja@gatech.edu

Abstract

We study the almost sure convergence of the Stochastic Approximation algorithm with diminishing step sizes $\alpha_n = \mathcal{O}\left(n^{-\xi}\right)$ for some $\xi \in (0, 1]$ under a general noise moment assumption and a contractive operator. In particular, we show that for a martingale difference noise with p-th order integrability, we have almost sure convergence whenever $\xi \in (\max\{1/2, 1/p\}, 1]$. Our result generalizes (weighted) Law of Large Numbers and the almost sure convergence results in [1, 3, 6]. To establish such results, we introduce a state-dependent moving truncation coupled with a fine-grained Lyapunov drift analysis. This approach effectively manages the bias from truncated terms and addresses the challenges posed by multiplicative noise, allowing us to relax the stringent assumptions often found in the literature.

1 Introduction

Various learning and optimization problems in Machine Learning (ML) and Operations Research (OR) can be reduced to solving for a solution $x^\star$ of the fixed point equation $T(x) = x$, which can be solved using the Robbins–Monro Stochastic Approximation (SA) algorithm written as follows

$$x_{n+1} = (1 - \alpha_n)x_n + \alpha_n \left(T(x_n) + \varepsilon_{n+1}\right) \tag{1}$$

where $T : \mathbb{R}^d \to \mathbb{R}^d$ is a possibly nonlinear mapping, $x_n, \varepsilon_n \in \mathbb{R}^d$ denote the state and the noise respectively. Traditionally, the SA algorithm has been studied through the lens of almost sure convergence [3], that is, one provides a guarantee that for any sample path, the iterate will not blow up and eventually converge to the solution $x^\star$ with probability 1. Under this viewpoint, [26] shows that by letting the step size sequence $\{\alpha_n\}$ decays at an appropriate rate, the effect of noise will vanish and the iterate will eventually converge to the desired solution. And so, it is apparent that there is a trade-off between convergence and variance of the iterates: for a larger step size, the iterates would converge faster to the desired solution, but this risks magnifying the effect of the noise. Thus, one would expect the choice of step sizes to depend on the noise conditions. For unbiased, square-integrable noise, it is sufficient to choose non-summable but square-integrable step sizes [1, 6, 23]. However, for general noise conditions, it is unclear what step sizes to take in order to ensure almost sure convergence. More formally, we study the problem under the following noise assumptions. The first assumption is that the noise sequence is a martingale difference sequence.

Assumption 1 (MDS). *A sequence $\{\varepsilon_n\}_{n\geq 0}$ adapted to $\{\mathcal{F}_n\}_{n\geq 0}$ is a martingale difference sequence (MDS) if each ε_n is $\mathcal{F}_n$-measurable and $\mathbb{E}[\varepsilon_{n+1}|\mathcal{F}_n] = 0$ a.s. for all $n \geq 0$.*

Secondly, instead of assuming square-integrability, we generalize this notion of noise boundedness to the p-th moment. In particular, we have the following assumption.

39th Conference on Neural Information Processing Systems (NeurIPS 2025) Workshop: ML×OR.

Assumption 2 (p-th order integrability)**.** *There exists $a, b \geq 0$ such that for any non-negative integer n, we have $\mathbb{E}\left[\|\varepsilon_{n+1}\|^p \,|\mathcal{F}_n\right] \leq a + b\,\|x_n - x^\star\|^p$ for some $p \in (1, \infty)$.*

When $b > 0$, ε_n is called multiplicative noise, which appears in many ML applications [8, 9, 18, 29], and handling it is highly non-trivial [8, 18]. For the step size sequence $\{\alpha_n\}_{n\geq 0}$, we write $\alpha_n = \alpha(n+K)^{-\xi}$ where $\alpha, K > 0$, and our focus is the diminishing step sizes setting $\xi \in (0, 1]$. For $p = 2$, [1, 6, 23] show that when $\xi \in (1/2, 1]$ and under some negative drift condition, one can obtain almost sure convergence. Moreover, with i.i.d noise, when $\alpha_n = (n+1)^{-1}$ (i.e. $\xi = 1$) and $T(x) = 0 \; \forall \; x \in \mathbb{R}^d, p \to 1$, we recover the Strong Law of Large Numbers. Thus, we pose the following conjecture.

Conjecture: Let $(\mathcal{H}, \langle \cdot, \cdot \rangle)$ be a real Hilbert space with norm $\|\cdot\|$ and $T : \mathcal{H} \to \mathcal{H}$ be a contraction with constant $\rho \in [0, 1)$ and unique fixed point $x^\star$. Under Assumptions 1 and 2 and step sizes $\alpha_n = \alpha(n+K)^{-\xi}$ for $\xi \in (1/p, 1]$, x_n converges to $x^\star$ almost surely.

Contributions: In this work, we resolve the conjecture for $p \in (1, 2)$ and show almost sure convergence for ξ all the way to $1/p$. To the best of our knowledge, our result is the first of its kind for this regime, obtained under remarkably mild noise conditions. For $p \geq 2$, we further show that $\xi \in (1/2, 1]$ guarantees almost sure convergence, and we conjecture that this bound is not yet tight.

To establish such results, we introduce a novel drift and truncation technique, which allows us to overcome the absence of finite second moment when $p \in (1, 2)$ while harnessing the power of the Lyapunov drift method to obtain almost sure convergence for the nonlinear Stochastic Approximation iterates and even generalize the classical Strong Law of Large Numbers. A key technical innovation is the design of a state-dependent truncation threshold to handle the multiplicative noise and the bias term introduced by the truncation step. This approach significantly extends beyond the classical truncation arguments for Strong Law of Large Numbers in the literature [12].

Related works: The SA algorithm has been studied under many different settings, with a fairly complete theoretical picture emerging. Researchers have explored its behavior with constant step sizes [36, 35, 7, 33], diminishing step sizes [6], as well as more elaborate schemes such as two-timescale SA [10, 34] under some general negative drift conditions [15, 23], which can have further applications in Markov chain mixing [24, 30] or learning [25, 28]. More recently, the focus of the literature has shifted toward finite-time guarantees with Reinforcement Learning often serving as the motivating application [18, 5, 11, 8]. However, proving almost sure convergence ensures the optimal solution is reached for almost all instantiations and sample paths of random noise without blowing up, thus an important and classic problem [3, 4, 17, 21, 32]. Several previous works such as [3, 4, 2, 16] establish almost sure convergence via the so-called ODE method by investigating continuous-time dynamics $\dot{x} = T(x) - x$ under some Lyapunov stability assumptions. When the noise is square integrable (Assumption 2 for $p = 2$), [3, 4] show the behavior of the discrete-time system (1) is identical to that of the ODE whenever $\xi \in (1/2, 1]$. However, under our general noise setting (Assumption 2), it is unclear how to pick the appropriate step sizes using the said ODE method.

Some prior heavy-tail SA results like [19, 14, 20] focus on the scalar Robbins–Monro recursion with mainly additive noise and give conditions for convergence but do not treat multidimensional, nonlinear contractive maps with multiplicative noise. [31] analyzes strongly convex SGD under heavy-tailed noise with finite p-th moment ($p \in [1, 2)$) and establishes L^p rates as opposed to almost sure convergence. The works most closely related to ours are [22] and [16], which also assume bounded p-moment on the noise. However, their noise assumption is considerably more restrictive as they assume $b = 0$, known as the additive noise condition in the literature [8]. In contrast, our Assumption 2 allows for multiplicative noise, where the noise can scale with the iterate, significantly broadening the applicability of our analysis. Additionally, both require the stringent locally bounded $(2p-2)$-th moment noise assumption, which is challenging to verify in the general setting we consider. Furthermore, these studies focus on the case $p > 2$, whereas our work concentrates on the $p \in (1, 2)$ setting. Thus, our results complement the existing literature by addressing the gap left by these prior works in heavier tail regimes.

2 Almost Sure Convergence Analysis

2.1 Main Results

We now present our main theorem for almost sure convergence of stochastic approximation.

Theorem 2.1. *Assume that the noise follows Assumption 1 and 2. With step sizes $\alpha_n = \alpha(n+K)^{-\xi}$, if $\xi \in (\max\{1/p, 1/2\}, 1]$, then the stochastic approximation iterates* (1) *with a ρ-contractive operator T converge a.s., i.e. $x_n \xrightarrow{\text{a.s.}} x^\star$.*

We want to highlight that the condition $\xi > 1/p$ is, in fact, tight for $p \in (1, 2)$.

Theorem 2.2. *Fix $p \in (1, 2)$ and step sizes $\alpha_n = \alpha(n + K)^{-\xi}$ with $\xi \in (0, 1]$. There exist a noise process $(\varepsilon_n)_{n\geq 0}$ satisfying Assumptions* 1 *and* 2 *and a contraction T such that if $\xi \leq 1/p$, the stochastic approximation iterates* (1) *fail to converge a.s.*

We obtain the following corollary for weighted LLN by letting $T(x) = \mu$ in Theorem 2.1.

Corollary 2.1 (Weighted LLN)**.** *Let $\{X_n\}_{n\geq 1}$ be i.i.d. random variables with mean μ and $\mathbb{E}|X_1 - \mu|^p < \infty$ for some $p > 1$. Fix $\xi \in (0, 1]$. Consider the recursion $S_{n+1} = (1 - \alpha_n)S_n + \alpha_n X_{n+1}$. If $\xi > \max\{1/2, 1/p\}$, then $S_n \xrightarrow{\text{a.s.}} \mu$.*

Remark 2.1. *Weighted LLN with $p \in (1, 2)$ can be proved directly by expanding the recursion into a weighted average and applying truncation plus summation; see, e.g., [13]. However, that expansion-based argument hinges on linearity and does not extend to nonlinear SA. Our drift + truncation approach treats the LLN and SA settings in a unified way: it controls truncation-induced* bias *term via a moving threshold, yielding almost sure convergence with one-step drift bounds.*

2.2 Proof Sketch for Main Theorem 2.1

Since $\xi > \max\{1/p, 1/2\} \geq 1/2$, the step sizes satisfy: $\alpha_n \to 0$, $\sum_{n=0}^{\infty} \alpha_n = \infty$, $\sum_{n=0}^{\infty} \alpha_n^2 < \infty$, and $\prod_{j=k}^{\infty}(1 - \alpha_j) = 0$. These properties will be used throughout the proof sketch.

Robbins–Siegmund. An essential tool for establishing almost sure convergence in stochastic approximation is the Robbins–Siegmund theorem [27]. We mainly use its following specialization.

Lemma 2.3. *Let $\{V_n\}_{n\geq 0}$ be nonnegative and adapted random variables, and let $\{M_n\}_{n\geq 0}$, $\{N_n\}_{n\geq 0}$ be nonnegative adapted sequences with $M_n \leq 1$ for all $n \geq 0$. Suppose $\sum_n M_n = \infty$ a.s., $\sum_n N_n < \infty$ a.s., and $\mathbb{E}[V_{n+1}|\mathcal{F}_n] \leq (1 - M_n)V_n + N_n$ a.s. for all n. Then, $V_n \xrightarrow{\text{a.s.}} 0$.*

Truncation Approach for $p \in (1, 2)$. In this case, we circumvent the lack of a second moment by truncating the heavy tails and controlling the truncated second moment via the p-moment. Indeed, given any random variable Y with bounded p-moment, any truncation threshold $B > 0$ and any σ-algebra $\mathcal{F}$, $\mathbb{E}\left[Y^2 \mathbf{1}_{\{|Y|\leq B\}}|\mathcal{F}\right] \leq B^{2-p}\mathbb{E}\left[|Y|^p|\mathcal{F}\right]$ a.s.

State-dependent Truncation Threshold. While truncation approach is a standard technique in the Law of Large Numbers setting [12], our key novelty is to consider a time- and state-dependent truncation threshold $B_n = (n+K)^\kappa(1+V_n)^{1/2}$, with the quadratic Lyapunov $V_n = \|x_n - x^\star\|^2$ and an adjustable $\kappa > 0$. Intuitively, the state-dependent term $(1 + V_n)^{1/2}$ allows the one-step Lyapunov drift stays linear in V_n (see Remark 2.2), rather than generating higher-order terms that are harder to control. The time-dependent factor $(n + K)^\kappa$ serves as a tunable knob to ensure the variance and bias series are summable in the Robbins–Siegmund step.

We also define the truncated noise $\tilde{\varepsilon}_{n+1} = \varepsilon_{n+1}\mathbf{1}_{\{\|\varepsilon_{n+1}\|\leq B_n\}}$, the truncated recursion $\tilde{x}_{n+1} = (1 - \alpha_n)\tilde{x}_n + \alpha_n T(\tilde{x}_n) + \alpha_n\, \tilde{\varepsilon}_{n+1}$, and the truncated quadratic Lyapunov $W_n = \|\tilde{x}_n - x^\star\|^2$. Let $\Delta_n = \|x_n - \tilde{x}_n\|^2$ be the squared distance from the original iterate to the truncated iterate at step n, we note that from Cauchy-Schwarz

$$V_n = \|x_n - x^\star\|^2 \leq 2\left(\|\tilde{x}_n - x^\star\|^2 + \|x_n - \tilde{x}_n\|^2\right) \leq 2(W_n + \Delta_n). \tag{2}$$

One-step Drift. Since T is ρ-contractive and $\alpha_n \in [0, 1]$, we have $\|(1 - \alpha_n)x + \alpha_n T(x) - x^\star\| \leq (1 - \alpha_n(1 - \rho))\|x - x^\star\|$. Let $\gamma = (1 - \rho)/2 \in (0, 1/2)$. Since $\alpha_n \to 0$, there exists an index M such that $\alpha_n \leq 1$ for all $n \geq M$. We upper bound the truncated quadratic Lyapunov for $n \geq M$ as

$$\mathbb{E}[W_{n+1}|\mathcal{F}_n] \leq (1 - 2\gamma\alpha_n)W_n + \underbrace{\alpha_n^2\mathbb{E}[\|\tilde{\varepsilon}_{n+1} - \mathbb{E}[\tilde{\varepsilon}_{n+1}|\mathcal{F}_n]\|^2|\mathcal{F}_n]}_{\text{variance}} + \underbrace{\alpha_n(1 + \gamma^{-1})\|\mathbb{E}[\tilde{\varepsilon}_{n+1}|\mathcal{F}_n]\|^2}_{\text{bias}}. \tag{3}$$

Bias Bound. In classical L^2 drift analyses the *noise* is conditionally mean–zero, so only a variance term (order α_n^2) appears. However, as the truncated noise $\tilde{\varepsilon}_{n+1}$ is biased, an extra bias term $\alpha_n(1 + \gamma^{-1})\|\mathbb{E}[\tilde{\varepsilon}_{n+1}|\mathcal{F}_n]\|^2$ enters the one–step bound. One of our novelty is to, via a moving threshold, show that it is summable under the p–moment, enabling almost sure convergence of the iterates. For a constant C_1, we can bound the bias as

$$\begin{aligned}\|\mathbb{E}[\tilde{\varepsilon}_{n+1}|\mathcal{F}_n]\|^2 &\le B_n^{2(1-p)}\mathbb{E}[\|\varepsilon_{n+1}\|^p|\mathcal{F}_n]^2 \le C_1(n+K)^{2\kappa(1-p)}(1+V_n)^{1-p}(1+V_n^{p/2})^2\\ &\le 4C_1(n+K)^{2\kappa(1-p)}(1+V_n) \le 8C_1(n+K)^{2\kappa(1-p)}(1+W_n+\Delta_n), \quad (4)\end{aligned}$$

where the last inequality follows from (2) and the fact that

$$\begin{aligned}(1+V_n)^{1-p}(1+V_n^{p/2})^2 &\le 2(1+V_n)^{1-p}(1+V_n^p) = 2(1+V_n)^{1-p} + 2V_n^p(1+V_n)^{1-p}\\ &\le 2(1+V_n) + 2(1+V_n)^p(1+V_n)^{1-p} = 4(1+V_n).\end{aligned}$$

Variance Bound. Similarly, with some constant C_2, we can bound the variance term as follows

$$\begin{aligned}\mathbb{E}[\|\tilde{\varepsilon}_{n+1}\|^2|\mathcal{F}_n] &\le B_n^{2-p}\mathbb{E}[\|\varepsilon_{n+1}\|^p|\mathcal{F}_n] \le C_2(n+K)^{\kappa(2-p)}(1+V_n)^{1-p/2}(1+V_n^{p/2})\\ &\le 2C_2(n+K)^{\kappa(2-p)}(1+V_n) \le 4C_2(n+K)^{\kappa(2-p)}(1+W_n+\Delta_n). \quad (5)\end{aligned}$$

Plugging (4) and (5) in (3), for $\zeta_n = 8C_1\alpha_n(1+\gamma^{-1})(n+K)^{2\kappa(1-p)} + 4C_2\alpha_n^2(n+K)^{\kappa(2-p)}$,

$$\mathbb{E}[W_{n+1}|\mathcal{F}_n] \le (1-2\gamma\alpha_n+\zeta_n)W_n + \zeta_n(\Delta_n+1), \quad (6)$$

ζ_n **Bound.** Since $\alpha_n = \alpha(n+K)^{-\xi}$, we note that $\sum_{n=0}^{\infty}\zeta_n = 8C_1(1+\gamma^{-1})\alpha\sum_{n=0}^{\infty}(n+K)^{2\kappa(1-p)-\xi} + 4C_2\alpha^2\sum_{n=0}^{\infty}(n+K)^{\kappa(2-p)-2\xi}$. The conditions in Corollary 2.3 require that the sum $\sum_{n=0}^{\infty}\zeta_n$ is finite; this is ensured if the exponents satisfy $\kappa(2-p)-2\xi, 2\kappa(1-p)-\xi < -1$, i.e. whenever $\kappa \in \left(\frac{1-\xi}{2p-2}, \frac{2\xi-1}{2-p}\right)$ (this interval is nonempty as $\xi > 1/p$). We also note that $\zeta_n/\alpha_n = 8C_1(1+\gamma^{-1})(n+K)^{2\kappa(1-p)} + 4C_2\alpha(n+K)^{\kappa(2-p)-\xi} \to 0$. Hence, there exists N_0 such that $\zeta_n \le \gamma\alpha_n$ for all $n \ge N_0$. For $n \ge \max\{M, N_0\}$, we can further bound the one-step drift (6) as

$$\mathbb{E}[W_{n+1}|\mathcal{F}_n] \le (1-\gamma\alpha_n)W_n + \zeta_n(\Delta_n+1). \quad (7)$$

Conditional Borel–Cantelli. By Markov's inequality, there exists a constant $C > 0$ such that

$$\Pr(\|\varepsilon_{n+1}\| > B_n|\mathcal{F}_n) \le \frac{\mathbb{E}[\|\varepsilon_{n+1}\|^p|\mathcal{F}_n]}{B_n^p} \le C\frac{1+V_n^{p/2}}{(n+K)^{\kappa p}(1+V_n)^{p/2}} \le 2C(n+K)^{-\kappa p}.$$

Hence, $\sum_n \mathbb{P}(\|\varepsilon_{n+1}\| > B_n|\mathcal{F}_n) < \infty$ almost surely whenever $\kappa > 1/p$, and the conditional Borel–Cantelli lemma gives an a.s. finite index N such that $\mathbf{1}_{\{\|\varepsilon_{n+1}\|>B_n\}} = 0$ for all $n \ge N$.

By contraction, $\|x_{n+1} - \tilde{x}_{n+1}\| \le (1-2\gamma\alpha_n)\|x_n - \tilde{x}_n\|$. Iterating for $n \ge N$, we have

$$\Delta_n^{1/2} = \|x_n - \tilde{x}_n\| \le \|x_N - \tilde{x}_N\|\prod_{t=N}^{n-1}(1-2\gamma\alpha_t) \xrightarrow{\text{a.s.}} 0.$$

Hence, we have $\Delta_n \xrightarrow{\text{a.s.}} 0$, implying $\sup_n \Delta_n < \infty$. From that we have $\sum_{n=0}^{\infty}(\Delta_n+1)\zeta_n$ is finite.

Conditions on κ. We now need $\kappa \in \left(\frac{1-\xi}{2p-2}, \frac{2\xi-1}{2-p}\right)$ for ζ_n to be summable and $\kappa > 1/p$ for the conditional Borel–Cantelli. As $\xi > 1/p$, we have $\frac{1}{p} > \frac{1-\xi}{2(p-1)}$. Hence, we require $\kappa \in \left(\frac{1}{p}, \frac{2\xi-1}{2-p}\right)$.

Convergence of the Original Process. With the above conditions applying to (7), we obtain $W_n \xrightarrow{\text{a.s.}} 0$ or $\tilde{x}_n \xrightarrow{\text{a.s.}} x^\star$ for all $n \ge \max\{M, N_0\}$ by Lemma 2.3. Since M and N_0 are finite, the truncated process converges almost surely, i.e. $\tilde{x}_n \xrightarrow{\text{a.s.}} x^\star$ in general. As we have shown $\Delta_n \xrightarrow{\text{a.s.}} 0$, the original process also converges almost surely to the fixed point, i.e. $x_n \xrightarrow{\text{a.s.}} x^\star$ when $\xi > 1/p$.

Drift approach for $p \ge 2$. We have Assumption 2 holds for $p = 2$. Using the quadratic Lyapunov and Lemma 2.3, we obtain almost sure convergence whenever $\xi > 1/2$ [3, 4]. □

Remark 2.2. *While the fine-grained conditions of κ in $B_n = (n+K)^\kappa(1+V_n)^{1/2}$ are discussed above, we also want to highlight that the exponent $1/2$ (of $1+V_n$) is optimal to obtain the linearity of V_n, hence of W_n, in the one-step drift bound. Indeed, we can consider $B_n = (n+K)^\kappa(1+V_n)^\theta$ for any $\theta > 0$. Then, we consider the bias and variance control:*

- *Bias* (4) *scales as $B_n^{2(1-p)}\mathbb{E}[\|\varepsilon_{n+1}|^p|\mathcal{F}_n]^2 \approx (1+V_n)^{2\theta(1-p)}(1+V_n^{p/2})^2$, whose dominant power is $V_n^{p+2\theta(1-p)}$, so linearity in V_n requires $p+2\theta(1-p) \le 1$, i.e. $\theta \ge 1/2$.*
- *Variance* (5) *scales as $B_n^{2-p}\mathbb{E}[\|\varepsilon_{n+1}\|^p|\mathcal{F}_n] \approx (1+V_n)^{\theta(2-p)}(1+V_n^{p/2})$, whose dominant power $V_n^{\theta(2-p)+p/2}$, so linearity in V_n requires $\theta(2-p)+p/2 \le 1$, i.e. $\theta \le 1/2$.*

Thus, the linearity of V_n and W_n is obtained if and only if $\theta = 1/2$.

References

[1] Dimitri P. Bertsekas and John N. Tsitsiklis. *Neuro-Dynamic Programming*. Athena Scientific, 1st edition, 1996.

[2] Vivek Borkar, Shuhang Chen, Adithya Devraj, Ioannis Kontoyiannis, and Sean Meyn. The ode method for asymptotic statistics in stochastic approximation and reinforcement learning, 2024.

[3] Vivek S Borkar and Vivek S Borkar. *Stochastic approximation: a dynamical systems viewpoint*, volume 100. Springer, 2008.

[4] Robert Buche and Harold J. Kushner. Rate of convergence for constrained stochastic approximation algorithms. *SIAM Journal on Control and Optimization*, 40(4):1011–1041, 2002.

[5] Siddharth Chandak, Shaan Ul Haque, and Nicholas Bambos. Finite-time bounds for two-time-scale stochastic approximation with arbitrary norm contractions and markovian noise. *arXiv preprint arXiv:2503.18391*, 2025.

[6] Zaiwei Chen, Siva Theja Maguluri, Sanjay Shakkottai, and Karthikeyan Shanmugam. Finite-sample analysis of stochastic approximation using smooth convex envelopes, 2020.

[7] Zaiwei Chen, Shancong Mou, and Siva Theja Maguluri. Stationary behavior of constant stepsize sgd type algorithms: An asymptotic characterization, 2021.

[8] Zaiwei Chen and Martin Siva Theja Maguluri. Concentration of contractive stochastic approximation: Additive and multiplicative noise. *The Annals of Applied Probability*, 35(2):1298–1352, 2025.

[9] Zaiwei Chen, Kaiqing Zhang, Eric Mazumdar, Asuman E. Ozdaglar, and Adam Wierman. A finite-sample analysis of payoff-based independent learning in zero-sum stochastic games. In *NeurIPS*, 2023.

[10] Thinh T Doan. Fast nonlinear two-time-scale stochastic approximation: Achieving $o(1/k)$ finite-sample complexity. *IEEE Transactions on Automatic Control*, 2025.

[11] Alain Durmus, Eric Moulines, Alexey Naumov, and Sergey Samsonov. Finite-time high-probability bounds for polyak–ruppert averaged iterates of linear stochastic approximation. *Mathematics of Operations Research*, 50(2):935–964, 2025.

[12] Rick Durrett. *Probability: theory and examples*, volume 49. Cambridge university press, 2019.

[13] Istvan Fazekas, Przemysław Matuła, and Maciej Ziemba. A note on the weighted strong law of large numbers under general conditions. *Publicationes Mathematicae-Debrecen*, 9(34), 2017.

[14] CA Goodsell and DL Hanson. Almost sure convergence for the robbins–monro process. *The Annals of Probability*, pages 890–901, 1976.

[15] Shaan Ul Haque and Siva Theja Maguluri. Stochastic approximation with unbounded markovian noise: A general-purpose theorem. In Yingzhen Li, Stephan Mandt, Shipra Agrawal, and Emtiyaz Khan, editors, *Proceedings of The 28th International Conference on Artificial Intelligence and Statistics*, volume 258 of *Proceedings of Machine Learning Research*, pages 3718–3726. PMLR, 03–05 May 2025.

[16] Ruinan Jin, Difei Cheng, Hong Qiao, Xin Shi, Shaodong Liu, and Bo Zhang. Stochastic gradient descent in non-convex problems: Asymptotic convergence with relaxed step-size via stopping time methods. *arXiv preprint arXiv:2504.12601*, 2025.

[17] Rajeeva Laxman Karandikar and Mathukumalli Vidyasagar. Convergence rates for stochastic approximation: Biased noise with unbounded variance, and applications. *Journal of Optimization Theory and Applications*, 203(3):2412–2450, 2024.

[18] Sajad Khodadadian, Pranay Sharma, Gauri Joshi, and Siva Theja Maguluri. Federated reinforcement learning: Linear speedup under markovian sampling, 2022.

[19] TP Krasulina. The method of stochastic approximation for the determination of the least eigenvalue of a symmetrical matrix. *USSR Computational Mathematics and Mathematical Physics*, 9(6):189–195, 1969.

[20] Gang Li. *Almost sure convergence of stochastic approximation procedures*. State University of New York at Binghamton, 1993.

[21] Jun Liu and Ye Yuan. Almost sure convergence rates analysis and saddle avoidance of stochastic gradient methods. *Journal of Machine Learning Research*, 25(271):1–40, 2024.

[22] Panayotis Mertikopoulos, Nadav Hallak, Ali Kavis, and Volkan Cevher. On the almost sure convergence of stochastic gradient descent in non-convex problems. *Advances in Neural Information Processing Systems*, 33:1117–1128, 2020.

[23] Hoang Huy Nguyen and Siva Theja Maguluri. Stochastic approximation for nonlinear discrete stochastic control: Finite-sample bounds. *arXiv preprint arXiv:2304.11854*, 2023.

[24] Hoang Huy Nguyen, Sushil Mahavir Varma, and Siva Theja Maguluri. Finite-time behavior of erlang-c model: Mixing time, mean queue length and tail bounds, 2025.

[25] M. Raginsky, A. Rakhlin, and M. Telgarsky. Non-convex learning via stochastic gradient Langevin dynamics: A nonasymptotic analysis. *arXiv preprint arXiv:1702.03849*, 2017.

[26] H. Robbins and S. Monro. A stochastic approximation method. *Annals of Mathematical Statistics*, 22:400–407, 1951.

[27] Herbert Robbins and David Siegmund. A convergence theorem for non negative almost supermartingales and some applications. In *Optimizing methods in statistics*, pages 233–257. Elsevier, 1971.

[28] Luke Snow and Vikram Krishnamurthy. Finite-sample bounds for adaptive inverse reinforcement learning using passive langevin dynamics. *IEEE Transactions on Information Theory*, 71(6):4637–4670, 2025.

[29] R. Srikant and Lei Ying. Finite-time error bounds for linear stochastic approximation and td learning, 2019.

[30] Amirhossein Taghvaei and Prashant G. Mehta. On the lyapunov foster criterion and poincaré inequality for reversible markov chains. *IEEE Transactions on Automatic Control*, 67(5):2605–2609, 2022.

[31] Hongjian Wang, Mert Gurbuzbalaban, Lingjiong Zhu, Umut Simsekli, and Murat A Erdogdu. Convergence rates of stochastic gradient descent under infinite noise variance. *Advances in Neural Information Processing Systems*, 34:18866–18877, 2021.

[32] Simon Weissmann, Sara Klein, Waïss Azizian, and Leif Döring. Almost sure convergence rates of stochastic gradient methods under gradient domination. *arXiv preprint arXiv:2405.13592*, 2024.

[33] Lu Yu, Krishnakumar Balasubramanian, Stanislav Volgushev, and Murat A. Erdogdu. An analysis of constant step size sgd in the non-convex regime: Asymptotic normality and bias, 2020.

[34] Sihan Zeng, Thinh T Doan, and Justin Romberg. A two-time-scale stochastic optimization framework with applications in control and reinforcement learning. *SIAM Journal on Optimization*, 34(1):946–976, 2024.

[35] Yixuan Zhang, Dongyan Huo, Yudong Chen, and Qiaomin Xie. Prelimit coupling and steady-state convergence of constant-stepsize nonsmooth contractive sa. In *Abstracts of the 2024 ACM SIGMETRICS/IFIP PERFORMANCE Joint International Conference on Measurement and Modeling of Computer Systems*, pages 35–36, 2024.

[36] Yixuan Zhang and Qiaomin Xie. Constant stepsize q-learning: Distributional convergence, bias and extrapolation. *arXiv preprint arXiv:2401.13884*, 2024.

Optional Appendix

A Proof of Lemma 2.3

We consider the classical Robbins–Siegmund Theorem [27].

Theorem A.1 (Robbins–Siegmund [27])**.** *Let $\{V_n\}_{n\geq 0}$, $\{A_n\}_{n\geq 0}$, $\{B_n\}_{n\geq 0}$, $\{C_n\}_{n\geq 0}$ be sequences of nonnegative integrable random variables, adapted to $\{\mathcal{F}_n\}_{n\geq 0}$, with $\sum_{i=0}^{\infty} A_i < \infty$, $\sum_{i=0}^{\infty} C_i < \infty$ a.s., and $\mathbb{E}[V_{n+1}|\mathcal{F}_n] \leq (1+A_n)V_n - B_n + C_n$ a.s. for all $n \geq 0$. Then, almost surely, $\{V_n\}$ converges and $\sum_{i=0}^{\infty} B_i < \infty$.*

By setting $A_n = 0,\ B_n = M_n V_n,\ C_n = N_n$, we obtain the desired Lemma 2.3.

B Full Derivation of the One-step Drift (3)

Let $G_\alpha(x) := (1-\alpha)x + \alpha T(x)$. Since T is ρ-contractive, for all $\alpha \in [0,1]$, we have

$$\|G_\alpha(x) - x^\star\| \leq (1-\alpha(1-\rho))\|x - x^\star\|.$$

Recall that $W_n := \|\tilde{x}_n - x^\star\|^2$. Using the facts that $\tilde{\varepsilon}_{n+1} = (\tilde{\varepsilon}_{n+1} - \mathbb{E}[\tilde{\varepsilon}_{n+1}|\mathcal{F}_n]) + \mathbb{E}[\tilde{\varepsilon}_{n+1}|\mathcal{F}_n]$, we can expand

$$\begin{aligned}
\mathbb{E}[W_{n+1}|\mathcal{F}_n] &= \mathbb{E}\left[\|G_{\alpha_n}(\tilde{x}_n) - x^\star + \alpha_n \mathbb{E}[\tilde{\varepsilon}_{n+1}|\mathcal{F}_n] + \alpha_n\left(\tilde{\varepsilon}_{n+1} - \mathbb{E}[\tilde{\varepsilon}_{n+1}|\mathcal{F}_n]\right)\|^2 \middle| \mathcal{F}_n\right] \\
&= \|G_{\alpha_n}(\tilde{x}_n) - x^\star + \alpha_n \mathbb{E}[\tilde{\varepsilon}_{n+1}|\mathcal{F}_n]\|^2 + \alpha_n^2 \mathbb{E}\left[\|\tilde{\varepsilon}_{n+1} - \mathbb{E}[\tilde{\varepsilon}_{n+1}|\mathcal{F}_n]\|^2 \middle| \mathcal{F}_n\right],
\end{aligned}$$

since the mixed term with the zero–mean part vanishes. We further have that

$$\begin{aligned}
\mathbb{E}[W_{n+1}|\mathcal{F}_n] &= \|G_{\alpha_n}(\tilde{x}_n) - x^\star\|^2 + \alpha_n^2 \left\|\mathbb{E}[\tilde{\varepsilon}_{n+1}|\mathcal{F}_n]\right\|^2 + 2\alpha_n \langle G_{\alpha_n}(\tilde{x}_n) - x^\star, \mathbb{E}[\tilde{\varepsilon}_{n+1}|\mathcal{F}_n]\rangle \\
&\quad + \alpha_n^2 \mathbb{E}\left[\|\tilde{\varepsilon}_{n+1} - \mathbb{E}[\tilde{\varepsilon}_{n+1}|\mathcal{F}_n]\|^2 \middle| \mathcal{F}_n\right] \\
&\leq (1-\alpha_n(1-\rho))^2 W_n + \alpha_n^2 \left\|\mathbb{E}[\tilde{\varepsilon}_{n+1}|\mathcal{F}_n]\right\|^2 + 2\alpha_n \langle G_{\alpha_n}(\tilde{x}_n) - x^\star, \mathbb{E}[\tilde{\varepsilon}_{n+1}|\mathcal{F}_n]\rangle \\
&\quad + \alpha_n^2 \mathbb{E}\left[\|\tilde{\varepsilon}_{n+1} - \mathbb{E}[\tilde{\varepsilon}_{n+1}|\mathcal{F}_n]\|^2 \middle| \mathcal{F}_n\right]
\end{aligned}$$

With Young's inequality $2\alpha_n\langle a, b\rangle \leq \alpha_n\gamma\|a\|^2 + \alpha_n\gamma^{-1}\|b\|^2$ where $\gamma = (1-\rho)/2$, we then obtain

$$\begin{aligned}
\mathbb{E}[W_{n+1}|\mathcal{F}_n] &\leq (1-\alpha_n(1-\rho))^2 W_n + \alpha_n^2 \left\|\mathbb{E}[\tilde{\varepsilon}_{n+1}|\mathcal{F}_n]\right\|^2 + \alpha_n\gamma\|G_{\alpha_n}(\tilde{x}_n) - x^\star\|^2 \\
&\quad + \alpha_n\gamma^{-1}\|\mathbb{E}[\tilde{\varepsilon}_{n+1}|\mathcal{F}_n]\|^2 + \alpha_n^2 \mathbb{E}\left[\|\tilde{\varepsilon}_{n+1} - \mathbb{E}[\tilde{\varepsilon}_{n+1}|\mathcal{F}_n]\|^2 \middle| \mathcal{F}_n\right] \\
&\leq (1-\alpha_n(1-\rho))^2 W_n + \alpha_n^2 \left\|\mathbb{E}[\tilde{\varepsilon}_{n+1}|\mathcal{F}_n]\right\|^2 + \alpha_n\gamma(1-\alpha_n(1-\rho))^2 W_n \\
&\quad + \alpha_n\gamma^{-1}\|\mathbb{E}[\tilde{\varepsilon}_{n+1}|\mathcal{F}_n]\|^2 + \alpha_n^2 \mathbb{E}\left[\|\tilde{\varepsilon}_{n+1} - \mathbb{E}[\tilde{\varepsilon}_{n+1}|\mathcal{F}_n]\|^2 \middle| \mathcal{F}_n\right] \\
&= (1+\alpha_n\gamma)(1-\alpha_n(1-\rho))^2 W_n + \left(\alpha_n^2 + \alpha_n\gamma^{-1}\right)\left\|\mathbb{E}[\tilde{\varepsilon}_{n+1}|\mathcal{F}_n]\right\|^2 \\
&\quad + \alpha_n^2 \mathbb{E}\left[\|\tilde{\varepsilon}_{n+1} - \mathbb{E}[\tilde{\varepsilon}_{n+1}|\mathcal{F}_n]\|^2 \middle| \mathcal{F}_n\right].
\end{aligned}$$

Since $\alpha_n \to 0$, there exists an index M such that $\alpha_n \leq 1$ for all $n \geq M$. We then obtain the following inequalities for all $n \geq M$

$$\begin{aligned}
(1+\alpha_n\gamma)(1-\alpha_n(1-\rho))^2 &\leq 1 - 2\gamma\alpha_n, \\
\alpha_n^2 + \alpha_n\gamma^{-1} &\leq (1+\gamma^{-1})\alpha_n.
\end{aligned}$$

The first inequality can be obtained by noting that

$$\begin{aligned}
(1+\alpha_n\gamma)(1-\alpha_n(1-\rho))^2 &= (1+\alpha_n(1-\rho)/2)(1-\alpha_n(1-\rho))^2 \\
&= 1 + \alpha_n(1-\rho)/2 - 2\alpha_n(1-\rho) - \alpha_n^2(1-\rho)^2 \\
&\quad + \alpha_n^2(1-\rho)^2 + \alpha_n^3(1-\rho)^3/2 \\
&= 1 - 3\alpha_n(1-\rho)/2 + \alpha_n^3(1-\rho)^3/2 \\
&\leq 1 - 3\alpha_n(1-\rho)/2 + \alpha_n(1-\rho)/2 \\
&= 1 - (1-\rho)\alpha_n \\
&= 1 - 2\gamma\alpha_n,
\end{aligned}$$

where the second last line comes from the fact that $\alpha_n(1-\rho) \leq 1$. We then obtain the bound (3)

$$\mathbb{E}[W_{n+1}|\mathcal{F}_n] \leq (1-2\gamma\alpha_n)W_n + (1+\gamma^{-1})\alpha_n \left\|\mathbb{E}[\tilde{\varepsilon}_{n+1}|\mathcal{F}_n]\right\|^2 + \alpha_n^2 \mathbb{E}\left[\left\|\tilde{\varepsilon}_{n+1} - \mathbb{E}[\tilde{\varepsilon}_{n+1}|\mathcal{F}_n]\right\|^2 \middle| \mathcal{F}_n\right],$$

for all $n \geq M$.

C Tightness of $\xi > 1/p$ (Theorem 2.2)

Fix $p \in (1,2)$, $\xi \in (0, 1/p]$, $\alpha > 0$, and $K \geq 1$. Take $T \equiv 0$ (a contraction with $\rho = 0$, fixed point 0), so $x_{n+1} = (1-\alpha_n)x_n + \alpha_n \varepsilon_{n+1}$. Let the filtration be $\mathcal{F}_n = \sigma(\varepsilon_1, \ldots, \varepsilon_n)$. For each n, define independent mean-zero noises ε_{n+1} by

$$\varepsilon_{n+1} = \begin{cases} s_n & \text{w.p. } q_n \\ -s_n & \text{w.p. } q_n \\ 0 & \text{w.p. } 1-2q_n \end{cases}, \; s_n := \frac{4}{\alpha}(n+K)^{\xi}, \; q_n := c\,(n+K)^{-\xi p},$$

with any $c \in (0, 1/2]$. Then $\mathbb{E}[\varepsilon_{n+1}|\mathcal{F}_n] = 0$ (MDS) and

$$\mathbb{E}\left[|\varepsilon_{n+1}|^p|\mathcal{F}_n\right] = 2q_n s_n^p = 2c\left(\frac{4}{\alpha}\right)^p < \infty,$$

implying that $\sup_n \mathbb{E}|\varepsilon_{n+1}|^p < \infty$.

Let $I_n := \{|\varepsilon_{n+1}| = s_n\}$. The I_n are independent with $\mathbb{P}(I_n) = 2q_n = 2c(n+K)^{-\xi p}$. Since $\xi p \leq 1$, $\sum_n \mathbb{P}(I_n) = \infty$; by Borel–Cantelli (for independent events), I_n occurs infinitely often a.s.

On I_n, we have $\alpha_n s_n = \alpha(n+K)^{-\xi} \cdot \frac{4}{\alpha}(n+K)^{\xi} = 4$, hence

$$x_{n+1} - x_n = -\alpha_n x_n \pm 4.$$

For every real u, at least one of $|4-u|$ or $|-4-u|$ is ≥ 4; since the sign $\pm$ is an independent symmetric coin flip, we have

$$\mathbb{P}\left(|x_{n+1} - x_n| \geq 4 | \, \mathcal{F}_n, \; I_n\right) \geq \frac{1}{2}.$$

Define $J_n := I_n \cap \{|x_{n+1} - x_n| \geq 4\}$. Then $\sum_n q_n = \infty$ and

$$\mathbb{P}(J_n|\mathcal{F}_n) \;\geq \frac{1}{2}\,\mathbb{P}(I_n|\mathcal{F}_n) \;=\; q_n.$$

By Lévy's conditional Borel–Cantelli lemma, J_n occurs infinitely often a.s. Thus $|x_{n+1} - x_n| \geq 4$ infinitely often a.s., so (x_n) cannot converge.