$O(1/k)$ Finite-Time Bound for Non-Linear Two-Time-Scale Stochastic Approximation

Siddharth Chandak

October 27, 2025

Department of Electrical Engineering, Stanford University

Outline

- Framework
- Prior Works
- Our Result
- Key Proof Technique
- Open Question

Framework

Two-Time-Scale Iterations

- Coupled iterations updating on separate time-scales

$$x_{k+1} = x_k + \alpha_k(f(x_k, y_k) - x_k + M_{k+1})$$
$$y_{k+1} = y_k + \beta_k(g(x_k, y_k) - y_k + M'_{k+1})$$

- Want to solve $f(x, y) = x$ and $g(x, y) = y$ given noisy realizations
- M_{k+1} and M'_{k+1} are martingale difference noise sequences arising from noisy observations
- Timescales dictated by the different stepsizes α_k and β_k

Two-Time-Scale Iterations

$$\text{Faster: } x_{k+1} = x_k + \alpha_k(f(x_k, y_k) - x_k + M_{k+1})$$
$$\text{Slower: } y_{k+1} = y_k + \beta_k(g(x_k, y_k) - y_k + M'_{k+1})$$

- α_k is larger, or decays at a slower rate, e.g., $1/n^{0.6}$
- β_k is smaller, or decays at a faster rate, e.g., $1/n^{0.75}$
- Analysis
    - Faster time-scale: y_k considered quasi-static
    - Slower time-scale: x_k tracks $x^*(y_k)$, the fixed point for $f(\cdot, y_k)$

Why study two-time-scale iterations?

Many applications:

- Minimax optimization
  - Two-time-scale stochastic gradient descent ascent algorithm
- Constrained optimization
  - Two-time-scale Lagrangian optimization
  - Particularly useful in distributed settings where agents make local updates with global constraints
- Game Control
  - Players update on faster time-scale
  - Game manager updates game parameters on slower time-scale

... and obviously more in Reinforcement Learning

Applications in RL

- SSP Q Learning
  - An algorithm for control of average reward MDPs
- Off Policy TD Learning with function approximation
  - GTD, TDC, GTD2
- A special case is RL algorithms with Polyak averaging
  - The slower timescale is just an averaging step
  - Better statistical guarantees

Key Contractive Assumptions

- There exists $0 \leq \lambda < 1$ such that,

$$\|f(x_1, y) - f(x_2, y)\| \leq \lambda \|x_1 - x_2\|$$

  for all x_1, x_2, y

  - Unique fixed point $x^*(y)$ for each y, such that $f(x^*(y), y) = x^*(y)$

- There exists $0 \leq \mu < 1$ such that

$$\|g(x^*(y_1), y_1) - g(x^*(y_2), y_2)\| \leq \mu \|y_1 - y_2\|$$

  for all y_1, y_2

  - Unique fixed point y^* such that $g(x^*(y^*), y^*) = y^*$

Equivalent Root-finding Formulation

- Consider the following iterations:

$$x_{k+1} = x_k + \alpha_k(\tilde{f}(x_k, y_k) + M_{k+1})$$
$$y_{k+1} = y_k + \beta_k(\tilde{g}(x_k, y_k) + M'_{k+1})$$

- Want to solve $\tilde{f}(x, y) = 0$ and $\tilde{g}(x, y) = 0$
- The key assumption now is strong monotonicity[1]:
  - $-\tilde{f}(\cdot, y)$ is strongly monotone
  - Unique $x^*(y)$ such that $\tilde{f}(x^*(y), y) = 0$
  - $-\tilde{g}(x^*(\cdot), \cdot)$ is strongly monotone

[1] Mapping $T(\cdot)$ is strongly monotone if there exists $\lambda' > 0$ such that $\langle T(x_1) - T(x_2), x_1 - x_2 \rangle \geq \lambda' \|x_1 - x_2\|^2$

Standard Assumptions

$$x_{k+1} = x_k + \alpha_k(f(x_k, y_k) - x_k + M_{k+1})$$
$$y_{k+1} = y_k + \beta_k(g(x_k, y_k) - y_k + M'_{k+1})$$

- Functions $f(\cdot)$ and $g(\cdot)$ are Lipschitz in x and y
- M_{k+1} and M'_{k+1} are martingale difference sequences. Moreover,

$$\mathbb{E}[\|M_{k+1}\|^2 + \|M'_{k+1}\|^2 \mid \mathcal{F}_k] \leq \mathfrak{c}_1(1 + \|x_k\|^2 + \|y_k\|^2),$$

  for some $\mathfrak{c}_1$.

Prior Works

Mean Square Error Bounds

$$x_{k+1} = x_k + \alpha_k(f(x_k, y_k) - x_k + M_{k+1})$$
$$y_{k+1} = y_k + \beta_k(g(x_k, y_k) - y_k + M'_{k+1})$$

- Want bounds on:

$$\mathbb{E}\left[\|x_k - x^*(y_k)\|^2\right] \text{ and } \mathbb{E}\left[\|y_k - y^*\|^2\right]$$

Stepsize Choices

$$\text{Faster: } x_{k+1} = x_k + \alpha_k(f(x_k, y_k) - x_k + M_{k+1})$$
$$\text{Slower: } y_{k+1} = y_k + \beta_k(g(x_k, y_k) - y_k + M'_{k+1})$$

- Can divide prior works into two types based on stepsizes α_k and β_k
- **Case I:** $\lim_{k\uparrow\infty} \beta_k/\alpha_k = 0$
  - 'True' Time-Scale Separation
- **Case II:** $\beta_k = \alpha_k = \Theta(1/k)$
  - Also called 'single-time-scale analysis of multiple coupled sequences'

Case I: True Time-Scale Separation

$$\alpha_k = \frac{\alpha}{(k+K)^a} \text{ and } \beta_k = \frac{\beta}{k+K},$$

where $0 < a < 1$.

- For the general non-linear two-time-scale, the previous best bound was $\mathcal{O}(1/k^{2/3})$ achieved when $a = 2/3$ [Doan (2023)[2]]

[2] T. T. Doan, "Nonlinear two-time-scale stochastic approximation: Convergence and finite-time performance", (2023)

Case II: 'Single Time-Scale' Analysis

$$\alpha_k = \frac{\alpha}{k+K} \text{ and } \beta_k = \frac{\beta}{k+K},$$

where β/α is sufficiently small.

- **No previous bound** without additional assumptions
- Under the assumption that $x^*(y)$ is differentiable and smooth, [Shen and Chen (2022)[3]] achieve $\mathcal{O}(1/k)$
- By modifying the iteration with additional averaging steps, [Doan (2024)[4]] achieve $\mathcal{O}(1/k)$

---

[3] H. Shen, and T. Chen, "A Single-Timescale Analysis For Stochastic Approximation With Multiple Coupled Sequences", (2022)

[4] T. T. Doan, "Fast Nonlinear Two-Time-Scale Stochastic Approximation: Achieving O(1/k) Finite-Sample Complexity", (2024)

Our Results

Our Results

We improve the bounds in both cases

Case I: 'True' Time-Scale Separation:

- Achieve $\mathcal{O}(1/k^a)$ where a can be arbitrarily close to one

Case II: 'Single Time-Scale' Analysis

- Achieve $\mathcal{O}(1/k)$ without any additional assumptions

Case I

Theorem

Suppose

$$\alpha_k = \frac{\alpha}{(k+K)^a} \text{ and } \beta_k = \frac{\beta}{k+K},$$

where $0.5 < a < 1$ and β, K are sufficiently large. Then,

$$\mathbb{E}[\|x_k - x^*(y_k)\|^2 + \|y_k - y^*\|^2] \leq \frac{C}{(k+K)^a}$$

Case II

Theorem

Suppose

$$\alpha_k = \frac{\alpha}{k+K} \text{ and } \beta_k = \frac{\beta}{k+K},$$

where β/α is sufficiently small and β, K are sufficiently large. Then,

$$\mathbb{E}[\|x_k - x^*(y_k)\|^2 + \|y_k - y^*\|^2] \leq \frac{C}{k+K}$$

Key Proof Technique

An Important Observation

- Recall that the previous best bound was $\mathcal{O}(1/k^{2/3})$
- **Observation:** The reason for this weaker bound was the way the noise in the slower time-scale (M'_{k+1}) was handled

$$x_{k+1} = x_k + \alpha_k(f(x_k, y_k) - x_k + M_{k+1})$$
$$y_{k+1} = y_k + \beta_k(g(x_k, y_k) - y_k + M'_{k+1})$$

- In fact, [Chandak et al. (2025)[5]] obtained $\mathcal{O}(1/k)$ in absence of noise in the slower time-scale

$$x_{k+1} = x_k + \alpha_k(f(x_k, y_k) - x_k + M_{k+1})$$
$$y_{k+1} = y_k + \beta_k(g(x_k, y_k) - y_k)$$

- **Need to handle M'_{k+1} better**

[5] S. Chandak, S. U. Haque, N. Bambos, "Finite-Time Bounds for Two-Time-Scale Stochastic Approximation with Arbitrary Norm Contractions and Markovian Noise"

A Simple (but powerful) Technique

- Define an averaged noise sequence and an auxiliary iterate
- Averaged Noise Sequence:

$$U_{k+1} = (1 - \beta_k)U_k + \beta_k M'_{k+1}, \text{ with } U_0 = 0$$

- Auxiliary Iterates:

$$z_k = y_k - U_k$$

Implications: Decay Rate of averaged noise

- Suppose $\mathbb{E}\left[1 + \|x_i\|^2 + \|y_i\|^2\right] \leq \Gamma_1$ for all $i \leq k-1$ and some Γ_1, then
$$\mathbb{E}\left[\|U_m\|^2\right] \leq 2\mathfrak{c}_1 \Gamma_1 \beta_m,\ \forall m \leq k.$$
- The averaged noise sequence decays at a rate of β_k
- Will come back later to the the boundedness in expectation

Implications: An Iterate Easier to Analyze

- The iteration can be rewritten as:

$$x_{k+1} = x_k + \alpha_k(f(x_k, z_k) - x_k + M_{k+1} + d_k)$$
$$z_{k+1} = z_k + \beta_k(g(x_k, z_k) - z_k + e_k).$$

  Here, $\|d_k\|^2$ and $\|e_k\|^2$ are both upper bounded by $L^2\|U_k\|^2$.
- Will now study $\mathbb{E}\left[\|x_k - x^*(z_k)\|^2\right]$ and $\mathbb{E}\left[\|z_k - y^*\|^2\right]$
- The noise in slower time-scale is now e_k, and $\mathbb{E}[\|e_k\|^2]$ decays at a rate of β_k

Implications: Going Back to Original Iterates

- Bound on original iterates directly follows from bound on auxiliary iterates

$$\begin{aligned}
&\mathbb{E}\left[\|x_k - x^*(y_k)\|^2 + \|y_k - y^*\|^2\right] \\
&\leq 2\mathbb{E}\left[\|x_k - x^*(z_k)\|^2 + \|z_k - y^*\|^2\right] + C_1\mathbb{E}\left[\|U_k\|^2\right].
\end{aligned}$$

Boundedness in Expecation

- Recall: Suppose $\mathbb{E}\left[1 + \|x_i\|^2 + \|y_i\|^2\right] \leq \Gamma_1$ for all $i \leq k-1$ and some Γ_1, then

$$\mathbb{E}\left[\|U_m\|^2\right] \leq 2\mathfrak{c}_1 \Gamma_1 \beta_m, \; \forall m \leq k.$$

- Induction-based approach -
  - Choose approrpiate Γ_2
  - Base Case: Iterates bounded by Γ_2 at time $k = 0$
  - Suppose iterates bounded in expectation by Γ_2 at time $k-1$
  - Implies required bounds hold at time k
  - Implies iterates bounded in expectation by Γ_2 at time k

Why did I call the technique powerful?

This simple proof technique can be used in many settings

- Easy to extend to other noise sequences, e.g., Markov noise
- Expectation Bounds for SA under arbitrary norm contractions
    - Directly use $\|x_k - x^*\|$ as the Lyapunov function
- Sub-Gaussian concentration bounds for SA with Markov noise
- A key step in obtaining last-iterate bounds for non-expansive SA

Open Questions

Better Bounds in Linear SA

- When the functions f and g are linear:

$$\mathbb{E}[\|x_k - x^*(y_k)\|^2] = \frac{C}{(k+K)^a} \text{ and } \mathbb{E}[\|y_k - y^*\|^2] = \frac{C}{(k+K)},$$

is achieved when

$$\alpha_k = \frac{\alpha}{(k+K)^a} \text{ and } \frac{\beta}{k+K}$$

Extending to non-linear SA?

- A recent work [Han et al. (2024)[6]] obtain the same rate for non-linear SA but under the assumption of local linearity
  - Local linearity allows them to use the same kind of techniques as used in linear SA
- Also give *empirical* evidence that local linearity is necessary to achieve this

[6] Y. Han, X. Li, Z. Zhang, "Finite-Time Decoupled Convergence in Nonlinear Two-Time-Scale Stochastic Approximation", (2024)

Thank You!

Thank You!

The talk was based on

- Chandak, Siddharth, "$O(1/k)$ Finite-Time Bound for Non-Linear Two-Time-Scale Stochastic Approximation." *arXiv:2504.19375* (2025).