The bias of constant step size stochastic approximation

(via Stein's method)

Nicolas Gast Sebastian Allmeier

Inria, Univ. Grenoble Alpes

Toulouse, workshop RL4SN, June 2024

Stochastic approximation algrithms

Robin-Monroe 51, Kushner-Yin 03, Benaïm 06, Borkar 08...

Many learning algorithm (SGD, TD-learning, Q-learning, Policy-gradient...) can be written as:

$$\theta_{n+1} = \theta_n + \alpha \Big(f(\theta_n, Y_n) + M_{n+1} \Big).$$

Parameter $\in \mathbb{R}^d$

step-size (constant or decreasing)

Stochastic approximation algrithms

Robin-Monroe 51, Kushner-Yin 03, Benaïm 06, Borkar 08...

Many learning algorithm (SGD, TD-learning, Q-learning, Policy-gradient...) can be written as:

$$\theta_{n+1} = \theta_n + \alpha \Big(f(\theta_n, Y_n) + M_{n+1} \Big).$$

Parameter $\in \mathbb{R}^d$

step-size (constant or decreasing)

e.g. Gradient of the loss function

Stochastic approximation algrithms

Robin-Monroe 51, Kushner-Yin 03, Benaïm 06, Borkar 08...

Many learning algorithm (SGD, TD-learning, Q-learning, Policy-gradient...) can be written as:

$$\theta_{n+1} = \theta_n + \alpha \left(f(\theta_n, Y_n) + M_{n+1} \right).$$

Parameter $\in \mathbb{R}^d$

step-size (constant or decreasing)

(optional) Markov chain

Noise

e.g. Gradient of the loss function

Stochastic approximation algrithms

Robin-Monroe 51, Kushner-Yin 03, Benaïm 06, Borkar 08...

Many learning algorithm (SGD, TD-learning, Q-learning, Policy-gradient...) can be written as:

$$\theta_{n+1} = \theta_n + \alpha \Big(f(\theta_n, Y_n) + M_{n+1} \Big).$$

Parameter $\in \mathbb{R}^d$

step-size

What is the behavior of θ_n (as n goes to infinity)?

e.g. Gradient of the loss function

The ODE method

e.g., Borkar–Meyn 2000

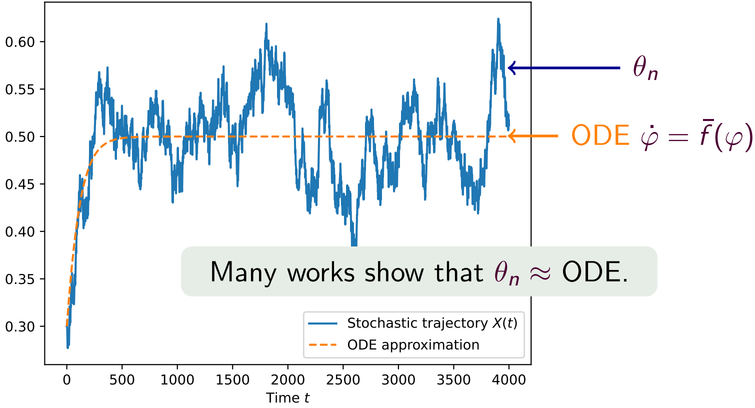

Many works show that $\theta_n \approx$ ODE.

The ODE method

e.g., Borkar–Meyn 2000

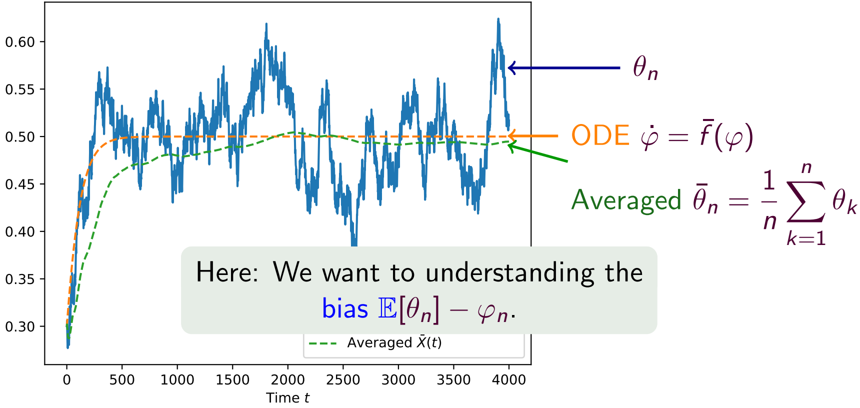

Here: We want to understanding the bias $\mathbb{E}[\theta_n] - \varphi_n$.

Results in a nutshell

Allmeier, Gast 2024, https://arxiv.org/abs/2405.14285

1. (related work): $|\theta_n - \varphi_n| = O(\sqrt{\alpha})$

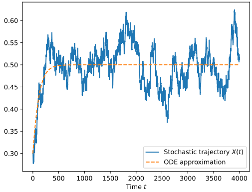

.

Results in a nutshell

Allmeier, Gast 2024, https://arxiv.org/abs/2405.14285

1. (related work): $|\theta_n - \varphi_n| = O(\sqrt{\alpha})$ 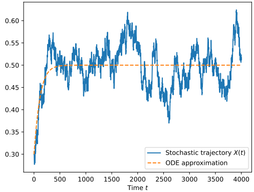

.

2. If $\bar{f}$ is *differentiable*: $|\bar{\theta}_n - \varphi_n| = V\alpha + O(\alpha^2)$ 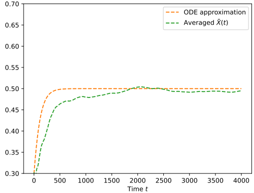

.

3. V can be characterized and computed

Outline

Outline

Stochastic approximation with constant step-size and Markovian noise

$$\theta_{n+1} = \theta_n + \alpha \Big(f(\theta_n, Y_n) + M_{n+1} \Big).$$

Assumptions:

1. $\mathbf{P}\left[Y_n = y' \mid Y_n = y, \theta_n = \theta\right] = K_{y,y'}(\theta)$ with $K(\theta)$ unichain for all θ.
   - This allows to define $\bar{f}(\theta) = \sum_y \pi_y(\theta) f(\theta, y)$.

Stochastic approximation with constant step-size and Markovian noise

$$\theta_{n+1} = \theta_n + \alpha\left(f(\theta_n, Y_n) + M_{n+1} \right).$$

Assumptions:

1. $\mathbf{P}\left[Y_n = y' \mid Y_n = y, \theta_n = \theta\right] = K_{y,y'}(\theta)$ with $K(\theta)$ unichain for all θ.
   - This allows to define $\bar{f}(\theta) = \sum_y \pi_y(\theta) f(\theta, y)$.
2. The ODE $\dot{\vartheta} = \bar{f}(\vartheta)$ has a unique, exponentially stable, attractor θ^*.

Stochastic approximation with constant step-size and Markovian noise

$$\theta_{n+1} = \theta_n + \alpha\Big(f(\theta_n, Y_n) + M_{n+1} \Big).$$

Assumptions:

1. $\mathbf{P}\left[Y_n = y' \mid Y_n = y, \theta_n = \theta\right] = K_{y,y'}(\theta)$ with $K(\theta)$ unichain for all θ.
   - This allows to define $\bar{f}(\theta) = \sum_y \pi_y(\theta) f(\theta, y)$.
2. The ODE $\dot{\vartheta} = \bar{f}(\vartheta)$ has a unique, exponentially stable, attractor θ^*.
3. All functions are (four times) differentiable.
4. θ_n lives in a compact and $\mathbb{E}\left[M_{n+1} \mid \mathcal{F}_n\right] = 0$.

The bias is asymptotically equal to $V\alpha$

Theorem (Allmeier, G. 2024)

Under assumptions (1) to (4), there exists V such that

$$\mathbb{E}\left[\bar{\theta}_n\right] = \theta^* + V\alpha + O(\frac{1}{n} + \alpha^2).$$

Moreover, $\exists C'$ s.t. for all ε: $\limsup_{n\to\infty} \mathbf{P}\left[\bar{\theta}_n - (\theta^* + \alpha V) \geq C'\alpha^{5^4}/\varepsilon\right] \leq \varepsilon.$

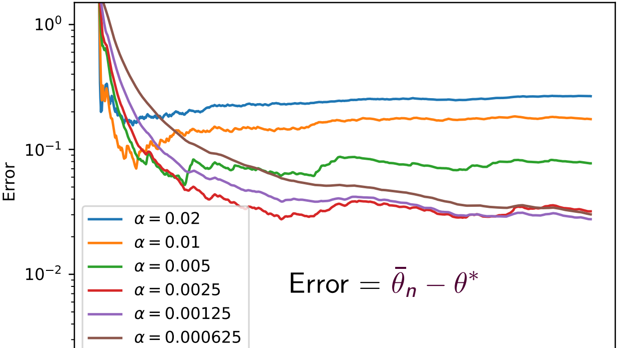

We can extrapolate V by using two step-sizes α and 2α

$$\bar{\theta}_n^{(\alpha)} = \theta^* + V\alpha + O(\alpha^2)$$

$$\bar{\theta}_n^{(2\alpha)} = \theta^* + V2\alpha + O(\alpha^2)$$

We can extrapolate V by using two step-sizes α and 2α

$$\bar{\theta}_n^{(\alpha)} = \theta^* + V\alpha + O(\alpha^2)$$
$$\bar{\theta}_n^{(2\alpha)} = \theta^* + V2\alpha + O(\alpha^2)$$

Hence:

$$2\bar{\theta}_n^{(\alpha)} - \bar{\theta}_n^{(2\alpha)} = \theta^* + O(\alpha^2).$$

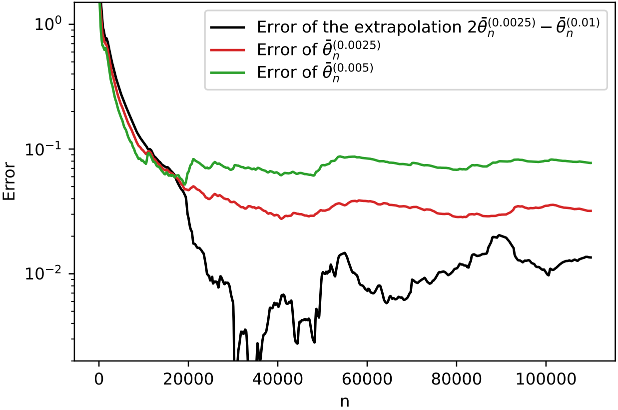

Outline

To compare θ_n and ODE, we study infinitesimal changes

$$\theta_{n+1} = \theta_n + \alpha \left(f(\theta_n, Y_n) + M_{n+1} \right).$$ Stochastic

$$\varphi_{n+1} = \varphi_n + \alpha \bar{f}(\varphi_n)$$ Deterministic

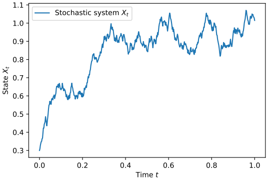

θ_n

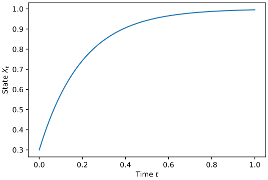

$\varphi_n(\theta_0)$

We want to compare:

$$\mathbb{E}\left[\theta_n\right] - \varphi_n(\theta_0)$$

To compare θ_n and ODE, we study infinitesimal changes

$$\theta_{n+1} = \theta_n + \alpha \left(f(\theta_n, Y_n) + M_{n+1} \right).$$ Stochastic

$$\varphi_{n+1} = \varphi_n + \alpha \bar{f}(\varphi_n)$$ Deterministic

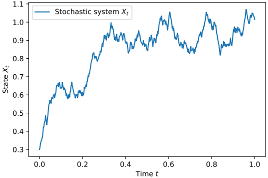

θ_n

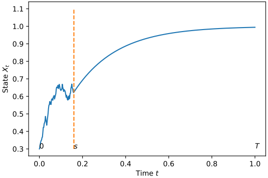

$\varphi_{n-k}(\theta_k)$

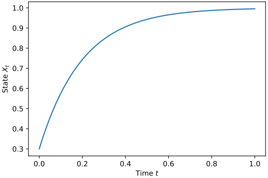

$\varphi_n(\theta_0)$

We want to compare:

$$\mathbb{E}\left[\theta_n\right] - \varphi_n(\theta_0)$$

To compare θ_n and ODE, we study infinitesimal changes

$$\theta_{n+1} = \theta_n + \alpha \left(f(\theta_n, Y_n) + M_{n+1} \right). \qquad \text{Stochastic}$$

$$\varphi_{n+1} = \varphi_n + \alpha \bar{f}(\varphi_n) \qquad \text{Deterministic}$$

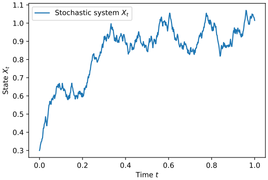

θ_n

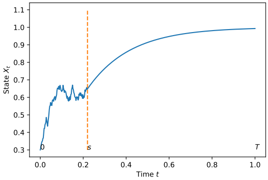

$\varphi_{n-k}(\theta_k)$

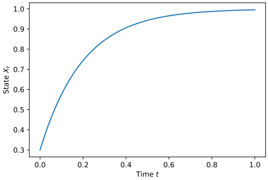

$\varphi_n(\theta_0)$

We want to compare:

$$\mathbb{E}\left[\theta_n\right] - \varphi_n(\theta_0) = \sum_{k=0}^{n-1} \mathbb{E}\left[\varphi_{n-k-1}(\theta_{k+1}) - \varphi_{n-k}(\theta_k)\right]$$

To compare θ_n and ODE, we study infinitesimal changes

$$\theta_{n+1} = \theta_n + \alpha \left(f(\theta_n, Y_n) + M_{n+1} \right).$$ Stochastic

$$\varphi_{n+1} = \varphi_n + \alpha \bar{f}(\varphi_n)$$ Deterministic

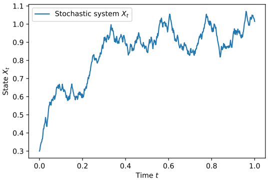

θ_n

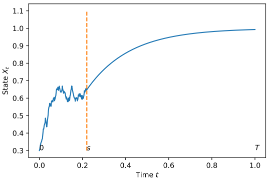

$\varphi_{n-k}(\theta_k)$

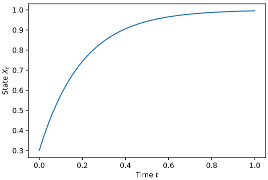

$\varphi_n(\theta_0)$

We want to compare:

$$\mathbb{E}\left[\theta_\infty\right] - \varphi_\infty(\theta_0) = \sum_{k=0}^{\infty} \mathbb{E}\left[\varphi_{n-k-1}(\theta_{k+1}) - \varphi_{n-k}(\theta_k)\right]$$

To compare θ_n and ODE, we study infinitesimal changes

$$\theta_{n+1} = \theta_n + \alpha \left(f(\theta_n, Y_n) + M_{n+1} \right). \qquad \text{Stochastic}$$

$$\varphi_{n+1} = \varphi_n + \alpha \bar{f}(\varphi_n) \qquad \text{Deterministic}$$

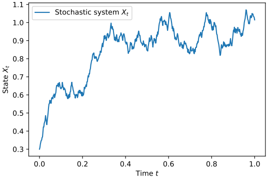

θ_n

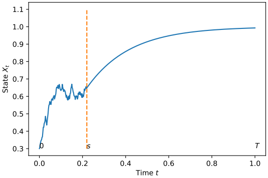

$\varphi_{n-k}(\theta_k)$

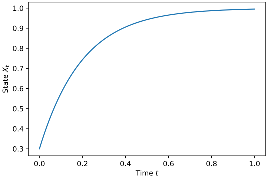

$\varphi_n(\theta_0)$

We want to compare:

$$\mathbb{E}\left[\theta_\infty\right] - \varphi_\infty(\theta_0) = \sum_{k=0}^{\infty} \mathbb{E}\left[\varphi_{n-k-1}(\theta_{k+1}) - \varphi_{n-k}(\theta_k)\right]$$

What is $\mathbb{E}\left[\varphi_{n-k-1}(\theta_{k+1}) - \varphi_{n-k}(\theta_k)\right]$?

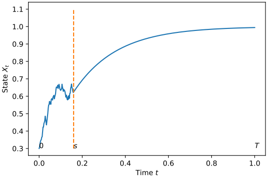

By definition:

$$\varphi_{n-k-1}(\theta_{k+1}) = \varphi_{n-k-1}\left(\theta_k + \alpha(f(\theta_k, Y_k) + M_{k+1})\right)$$
$$\varphi_{n-k}(\theta_k) = \varphi_{n-k-1}(\theta_k + \alpha\bar{f}(\theta_k))$$

What is $\mathbb{E}\left[\varphi_{n-k-1}(\theta_{k+1}) - \varphi_{n-k}(\theta_k)\right]$?

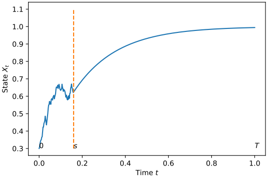

By definition:

$$\varphi_{n-k-1}(\theta_{k+1}) = \varphi_{n-k-1}\left(\theta_k + \alpha(f(\theta_k, Y_k) + M_{k+1})\right)$$
$$\varphi_{n-k}(\theta_k) = \varphi_{n-k-1}(\theta_k + \alpha\bar{f}(\theta_k))$$

Denoting $\Delta_n := f(\theta_k, Y_k) + M_{k+1} - \bar{f}(\theta_k)$, the difference equals:

$$\underbrace{\alpha D\varphi_{n-k-1}(\theta_k)\Delta_k}_{=:(A)} + \underbrace{\alpha^2 D^2\varphi_{n-k-1}(\theta_k)\Delta_k^2}_{=:(B)} + O(\alpha^3)$$

Term (B) is "easy to analyze"

Lemma (Exponentially small derivatives)

By exponential stabiliy:

$$\left\|D^i\phi_n\right\| \leq ce^{-c'\alpha n}$$

Hence:

$$\sum_{k=0}^{\infty}\mathbb{E}\left[(B)\right] = \alpha^2\sum_{k=0}^{\infty}\mathbb{E}\left[\underbrace{D^2\varphi_{n-k-1}(\theta_k)}_{e^{-\Omega(\alpha(n-k))}}\underbrace{\Delta_k^2}_{\text{Bounded}}\right]$$
$$= O(\alpha).$$

Term (A) needs “averaging”

Lemma (Averaging)

By using an averaging principle (Poisson equation):

$$\sum_{k=0}^{T}\Big(f(\theta_k, Y_k) - \bar{f}(\theta_k)\Big) = O(\alpha + \frac{1}{T}).$$

Hence:

$$\begin{aligned}\sum_{k=0}^{\infty}\mathbb{E}\left[(A)\right] &= \alpha^2\sum_{k=0}^{\infty}\mathbb{E}\left[D\varphi_{n-k-1}(\theta_k)f(\theta_k, Y_k) - \bar{f}(\theta_k)\right]\\ &= O(\alpha).\end{aligned}$$

The detailed proof is in the paper

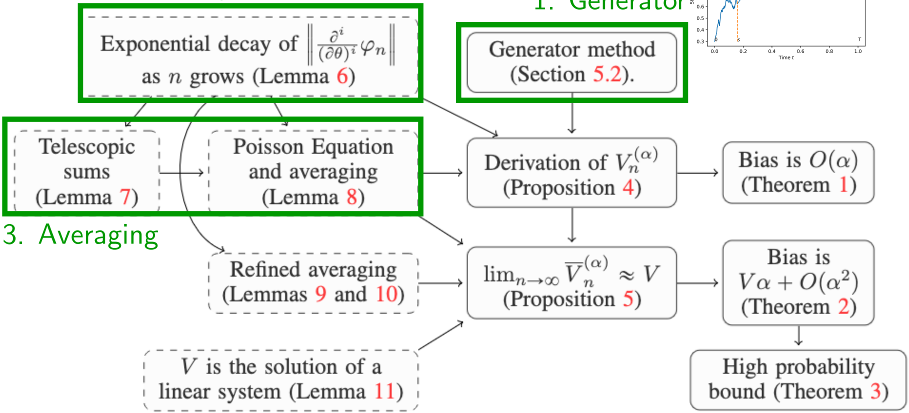

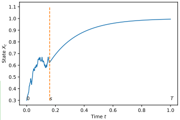

Outline

Conclusion and discussion

Stochastic approximation with constant step size and Markovian noise

- Methodology to characterize and compute the bias.
- Can be used for extrapolation.

Conclusion and discussion

Stochastic approximation with constant step size and Markovian noise

- Methodology to characterize and compute the bias.
- Can be used for extrapolation.

Methodology similar to the refined mean field ideas (Stein's method).

- **Main limit:** Dynamics needs to be smooth.

Slides and references: `http://polaris.imag.fr/nicolas.gast`

References

Results on which this talk is based:

- Bias and Refinement of Multiscale Mean Field Models. Allmeier, Gast, 2022. Sigmetrics 2023.
- Computing the Bias of Constant-step Stochastic Approximation with Markovian Noise. Allmeier, Gast. https://www.arxiv.org/abs/2405.14285

Q-learning and bias:

- Bias and Extrapolation in Markovian Linear Stochastic Approximation with Constant Stepsizes. Dongyan Huo, Yudong Chen, Qiaomin Xie. Sigmetrics 2023.

Related refined mean-field approximation papers:

- Mean Field and Refined Mean Field Approximations for Heterogeneous Systems: It Works! by Allmeier and Gast. SIGMETRICS 2022.
- A Refined Mean Field Approximation by Gast and Van Houdt. SIGMETRICS 2018.