Statistical inference for Linear Stochastic Approximation with Markovian Noise

HSE UNIVERSITY

Sergey Samsonov[1], Marina Sheshukova[1], Eric Moulines[2,3], Alexey Naumov[1,4]

[1]HSE University [2]CMAP, École Polytechnique, [3]MBZUAI, [4]Steklov Mathematical Institute of Russian Academy of Sciences

Problem Setting

- Goal: Solve approximately a linear system

$$\bar{\mathbf{A}}\theta^\star = \bar{\mathbf{b}}, \qquad \theta^\star \in \mathbb{R}^d. \tag{1}$$

- We observe noisy samples $\{(\mathbf{A}(Z_k), \mathbf{b}(Z_k))\}_{k\in\mathbb{N}}$ taking values in a measurable space $(\mathsf{Z}, \mathcal{Z})$ such that

$$\mathbb{E}_\pi[\mathbf{A}(Z_k)] = \bar{\mathbf{A}}, \qquad \mathbb{E}_\pi[\mathbf{b}(Z_k)] = \bar{\mathbf{b}},$$

  where π is the stationary distribution of an ergodic Markov chain $\{Z_k\}_{k\in\mathbb{N}}$ with transition kernel P.
- We solve (1) using LSA algorithm. Given a step size sequence $\{\alpha_k\}$ and an initial point θ_0, define

$$\theta_k = \theta_{k-1} - \alpha_k\{\mathbf{A}(Z_k)\theta_{k-1} - \mathbf{b}(Z_k)\}, \qquad \bar{\theta}_n = \frac{1}{n}\sum_{k=0}^{n-1}\theta_k.$$

  The Polyak–Ruppert average $\bar{\theta}_n$ reduces variance and ensures stability.
- Under suitable assumptions (see [3]), the LSA iterate is asymptotically normal:

$$\sqrt{n}(\bar{\theta}_n - \theta^\star) \xrightarrow{d} \mathcal{N}(0, \Sigma_\infty), \qquad \Sigma_\infty = \bar{\mathbf{A}}^{-1}\Sigma_\varepsilon\bar{\mathbf{A}}^{-\top}.$$

- We study the projected statistic

$$\sqrt{n}\, u^\top(\bar{\theta}_n - \theta^\star), \qquad u \in \mathbb{S}^{d-1},$$

  and aim to derive:
  - Non-asymptotic normal approximation bounds (Berry–Esseen type);
  - Bootstrap-based confidence regions for the projection.

Assumptions

Introduce centered "noise at optimum point": $\varepsilon(Z_k) = \tilde{\mathbf{A}}(Z_k)\theta^\star - \tilde{\mathbf{b}}(Z_k)$, where $\tilde{\mathbf{A}}(Z_k) = \mathbf{A}(Z_k) - \bar{\mathbf{A}}$, $\tilde{\mathbf{b}}(Z_k) = \mathbf{b}(Z_k) - \bar{\mathbf{b}}$.

A1 P admits π as a unique invariant distribution and is uniformly geometrically ergodic, that is, there exists $t_{\mathrm{mix}} \in \mathbb{N}$, such that for any $k \in \mathbb{N}$, it holds that

$$\Delta(\mathrm{P}^k) := \sup_{z,z'\in\mathsf{Z}} \mathsf{d}_{\mathrm{tv}}(\mathrm{P}^k(z,\cdot), \mathrm{P}^k(z',\cdot)) \le (1/4)^{\lceil k/t_{\mathrm{mix}}\rceil}. \tag{2}$$

A2 Matrix $-\bar{\mathbf{A}}$ is Hurwitz. Moreover, $\|\varepsilon\|_\infty = \sup_{z\in\mathsf{Z}}\|\varepsilon(z)\| < +\infty$, and the mapping $z \to \mathbf{A}(z)$ is bounded, that is,

$$\mathrm{C}_{\mathbf{A}} = \sup_{z\in\mathsf{Z}}\|\mathbf{A}(z)\| \vee \sup_{z\in\mathsf{Z}}\|\tilde{\mathbf{A}}(z)\| < \infty. \tag{3}$$

Moreover, we assume that $\lambda_{\min}(\Sigma_\varepsilon) > 0$, where

$$\Sigma_\varepsilon = \mathbb{E}_\pi[\varepsilon(Z_0)\{\varepsilon(Z_0)\}^\top] + 2\sum_{\ell=1}^{\infty}\mathbb{E}_\pi[\varepsilon(Z_0)\{\varepsilon(Z_\ell)\}^\top].$$

is the noise covariance matrix under the stationary distribution π;

A3 Step sizes $\{\alpha_k\}_{k\in\mathbb{N}}$ have a form $\alpha_k = c_0/(k+k_0)^\gamma$, where $\gamma \in [1/2; 1)$.

Contributions

- Non-asymptotic bound for projected PR-averaged LSA iterates $\sqrt{n}\, u^\top(\bar{\theta}_n - \theta^\star)$ under Markov noise, with rate $\mathcal{O}(n^{-1/4})$ in Kolmogorov distance.
- We obtain the first non-asymptotic analysis of the *multiplier subsample bootstrap* [4] for LSA with Markov noise. The coverage error decays as $\mathcal{O}(n^{-1/10})$. As a byproduct, we recover the $\mathcal{O}(n^{-1/8})$ rate for estimating asymptotic variance via the overlapping batch mean estimator [5].
- Application: Applied to temporal difference (TD) learning for policy evaluation in reinforcement learning.

Gaussian Approximation

Denote by Φ the c.d.f. of a standard Gaussian random variable and set

$$\mathsf{d}_K(X) = \sup_{x\in\mathbb{R}}|\mathbb{P}(X \le x) - \Phi(x)|.$$

Goal: Non-asymptotic CLT for the projected statistic $\sqrt{n}\, u^\top(\bar{\theta}_n - \theta^\star)$.

- **Decomposition:** $\sqrt{n}u^\top(\bar{\theta}_n - \theta^\star) = n^{-1/2}u^\top M + D$, where M is a martingale term (via Poisson equation), with variance $\sigma_n(u)$, and D is a small remainder controlledvia concentration results for additive functionals of Markov chains
- **Martingale CLT:** Apply quantitative results of [1, 2] for M in Kolmogorov distance.

Theorem 1. Under A1–A3, for any $u \in \mathbb{S}_{d-1}$,

$$d_K\left(\frac{\sqrt{n}\, u^\top(\bar{\theta}_n - \theta^\star)}{\sqrt{u^\top\Sigma_\infty u}}, \mathcal{N}(0,1)\right) \lesssim_{\log n} B_n, \quad B_n \asymp n^{-1/4} + n^{-(\gamma-1/2)} + n^{\gamma-1}.$$

Optimized rate: Balancing terms gives $\gamma = 3/4$, hence $d_K\left(\frac{\sqrt{n}\, u^\top(\bar{\theta}_n - \theta^\star)}{\sqrt{u^\top\Sigma_\infty u}}, \mathcal{N}(0,1)\right) \lesssim_{\log n} n^{-1/4}$.

Multiplier Subsampling Bootstrap

Goal: Confidence intervals for $u^\top\theta^\star$ using $\{u^\top\theta_k\}_{k=0}^{n-1}$.

MSB procedure:

- Block length b_n, local averages $\bar{\theta}_{b_n,t} = \frac{1}{b_n}\sum_{\ell=t}^{t+b_n-1}\theta_\ell$
- MSB estimator: $\bar{\theta}_{n,b_n}(u) = \frac{\sqrt{b_n}}{\sqrt{n-b_n+1}}\sum_{t=0}^{n-b_n} w_t(\bar{\theta}_{b_n,t} - \bar{\theta}_n)^\top u, \quad w_t \sim \mathcal{N}(0,1)$
- Under $\mathbb{P}^{\mathsf{b}}$, the MSB estimator is Gaussian: $\bar{\theta}_{n,b_n}(u) \sim \mathcal{N}(0, \hat{\sigma}^2_\theta(u))$ with variance $\hat{\sigma}^2_\theta(u) = \frac{b_n}{n-b_n+1}\sum_{t=0}^{n-b_n}\left((\bar{\theta}_{b_n,t} - \bar{\theta}_n)^\top u\right)^2$

Concentration of OBM estimator:

- We can show that $\hat{\sigma}^2_\theta(u) = \hat{\sigma}^2_\varepsilon(u) + t(b_n, n, \gamma)$, where $\hat{\sigma}^2_\varepsilon(u)$ is the overlapping batch mean estimator of $u^\top\Sigma_\infty u$.
- Using the concentration results for the overlapping batch mean estimator, we have:

Proposition 1. With $b_n \asymp n^{3/4}$ and aggressive steps $\alpha_k \asymp (k+k_0)^{-1/2-\varepsilon}$:

$$\left|\hat{\sigma}^2_\theta(u) - u^\top\Sigma_\infty u\right| \lesssim_{\log n} n^{-1/8+\varepsilon/2} \quad \text{(w.h.p.)}.$$

Key idea: Bootstrap distribution approximates real distribution: $\mathbb{P}^{\mathsf{b}}(\bar{\theta}_{n,b_n}(u) \le x) \approx \mathbb{P}\left(\sqrt{n}(\bar{\theta}_n - \theta^\star)^\top u \le x\right)$

Theorem 2. For $b_n = \lceil n^{4/5}\rceil$, $\alpha_k = c_0/(k_0+k)^{3/5}$

$$\sup_x\left|\mathbb{P}\left(\sqrt{n}(\bar{\theta}_n - \theta^\star)^\top u \le x\right) - \mathbb{P}^{\mathsf{b}}(\bar{\theta}_{n,b_n}(u) \le x)\right| \lesssim_{\log n} n^{-1/10}$$

References

[1] E. Bolthausen.
Exact Convergence Rates in Some Martingale Central Limit Theorems.
The Annals of Probability, 10(3):672 – 688, 1982.

[2] Xiequan Fan.
Exact rates of convergence in some martingale central limit theorems.
Journal of Mathematical Analysis and Applications, 469(2):1028–1044, 2019.

[3] G. Fort.
Central limit theorems for stochastic approximation with controlled Markov chain dynamics.
ESAIM: PS, 19:60–80, 2015.

[4] Ruru Ma and Shibin Zhang.
Multiplier subsample bootstrap for statistics of time series.
J. Statist. Plann. Inference, 233:Paper No. 106183, 15, 2024.

[5] Xi Chen Wanrong Zhu and Wei Biao Wu.
Online Covariance Matrix Estimation in Stochastic Gradient Descent.
Journal of the American Statistical Association, 118(541):393–404, 2023.