Linear System Identification: Sample Complexity Lower Bounds and Optimal Algorithms

Yassir Jedra and Alexandre Proutiere

KTH Royal Institute of Technology

Talk is mainly based on parts of Yassir's thesis:

- Sample Complexity Lower Bounds for Linear System Identification
  CDC, 2019
- Finite-time Identification of Stable Linear Systems: Optimality of the Least Squares Estimator
  CDC, 2020
- Finite-time Identification of Linear Systems: Fundamental Limits and Optimal Algorithms
  IEEE TAC, 2023
- Minimal Expected Regret in Linear Quadratic Control
  AISTATs, 2022

Linear System Identification

Linear time invariant systems. $x_{t+1} = Ax_t + Bu_t + \eta_{t+1}$

Unknown $A \in \mathbb{R}^{d\times d}$, known $B \in \mathbb{R}^{d\times m}$
$(\eta_t)_{t\geq 1}$ i.i.d. $\sim \mathcal{N}(0, I_d)$ or i.i.d. isotropic[1] vectors with independent coordinates of ψ_2-norm[2] less than K.

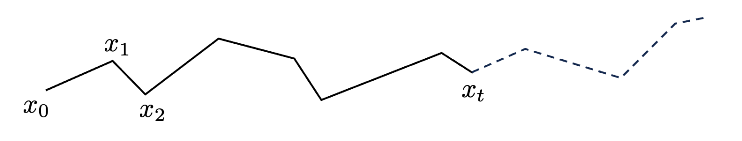

From the observation of a finite system trajectory, learn A with minimal error.
The trajectory can be **uncontrolled** or **controlled**.

[1] Isotropic means $\mathbb{E}[\eta_t \eta_t^\top] = I_d$.
[2] $\|\eta_t\|_{\psi_2} = \inf\{K \geq 0 : \mathbb{E}[\exp(\eta_t^2/K^2)] \leq 2\}$

Fixed budget vs. fixed confidence settings

Fixed budget. Fixed trajectory length.

- Algorithm: an adaptive control policy[3] + an estimator $\hat{A}_t$;
- Sample complexity $\tau_A := \inf\{t' : \forall t \geq t', \mathbb{P}_A(\|\hat{A}_t - A\| > \varepsilon) < \delta\}$;
- Objective: devise an algorithm with minimal sample complexity.

Fixed confidence. Target accuracy and confidence levels (ε, δ).

- Algorithm: an adaptive control policy[3] + a stopping rule + an estimator $\hat{A}$;
- Sample complexity $\mathbb{E}_A[\tau_A]$ where τ_A is the time when the algorithm stops;
- Objective: devise an (ε, δ)-PAC algorithm with minimal sample complexity $\mathbb{E}_A[\tau_A]$;
  (ε, δ)-PAC: $\mathbb{P}_A(\|\hat{A} - A\| > \varepsilon) < \delta$.

[3] Only in controlled / active learning scenario.

Minimax vs instance-specific guarantees

Minimax sample complexity. $G(\mathcal{A}, \varepsilon, \delta) = \inf_{\text{algo}} \sup_{A \in \mathcal{A}} \tau_A$ where $\mathcal{A}$ is a wide set of possible systems.
Minimax optimality: an algorithm with $\tau_A \leq G(\mathcal{A}, \varepsilon, \delta)$ for all $A \in \mathcal{A}$.

Instance-specific sample complexity. $g(A, \varepsilon, \delta) = \inf_{\text{algo} \in \Pi} \tau_A$ where Π is a wide set of (reasonably adaptive) algorithms.
Instance-specific optimality: an algorithm in Π with $\tau_A \leq g(A, \varepsilon, \delta)$ for all A.

Learning: be specific!

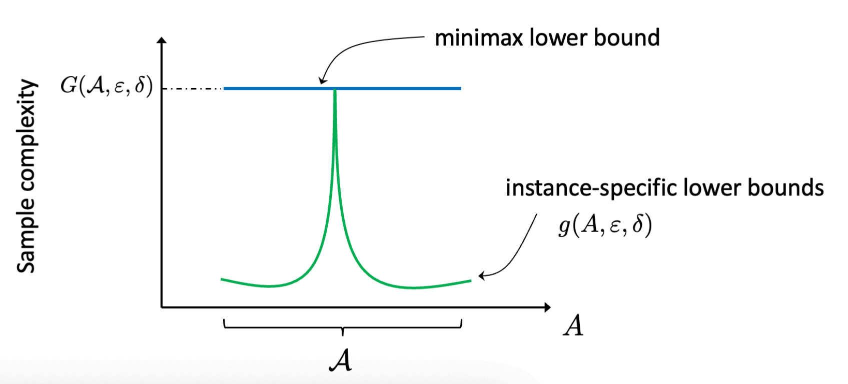

Minimax lower bounds come from wacky and meaningless examples.
Minimax optimality just states that the algorithm performs ok in the worst possible system, it does not say whether the algorithm *learns* and *adapts* to the system.

Research agenda

Part I. Derive instance-specific sample complexity lower bounds for system identification
How do they depend on A?
How do they scale in $\varepsilon, \delta, d, m$?

Part II. Devise instance-specific optimal algorithm
Is the LSE optimal?
What stopping rule works in the fixed confidence setting?
How to perform optimal exploration/excitation?

Related work – Uncontrolled systems

Control community. Mostly asymptotic results, e.g. **Ljung**'76, LSE and Prediction error methods. For finite time analysis, see **Weyer et al.**'96, **Tikku-Poola**'93, **Daleh**'93, etc.

Learning approach. (Given accuracy and confidence levels ε, δ)

- **Faradonbeh-Tewari-Michailidis**'18: $\tau_A \approx \frac{1}{\varepsilon^2}\log(1/\delta)^3 C(A) d\log(d)$
- **Simchowitz-Mania-Tu-Jordan-Recht**'18:

$$\tau_A \approx \frac{1}{\lambda_{\min}(\Gamma_k(A))\varepsilon^2}(d\log(d/\delta) + \log\det(\Gamma_t(A)\Gamma_k(A)^{-1}))$$

  if (τ_A, k) satisfies:

$$\tau_A \geq ck(d\log(d/\delta) + \log\det(\Gamma_t(A)\Gamma_k(A)^{-1}))$$

- **Sarkar-Rakhlin**'19: $\tau_A \approx C(A,d)\frac{1}{\varepsilon^2}\log(\frac{1}{\delta})$

We aim for an optimal and explicit dependence in δ, ε, and A.

Related work – Controlled systems

Control community. **Mehra**'76, **Goodwin-Payne**'77, **Jansson-Hjalmarsson**'05, **Rojas et al.**'07, etc.

Learning approach. (Given accuracy and confidence levels ε, δ)

- **Wagenmaker-Jamieson**'20: more later on this paper ...

Objectives of this talk

Simplifying (as much as possible) and explaining learning tools towards a finite-time analysis of linear sysID

1. Lower bounds (**information theory**)
2. Performance analysis of algorithms (**concentration of random vectors and matrices**)

Part I. Instance-specific sample complexity lower bounds

A generic inference problem

Familly of stochastic generative parametrized models $\mathcal{A}$
Observations O sampled under some unknown model $A \in \mathcal{A}$
Inference algorithm $\pi : O \mapsto \hat{A}$, performance metric p_A^π

What is the fundamental inference limit $\sup_\pi p_A^\pi$?

Example 1: $\mathcal{A} = \mathbb{R}^{d \times d}$, $O = (x_0, \dots, x_t)$ where $x_{t+1} = Ax_t + \eta_{t+1}$,
$p_A^\pi = \tau_A^{-1}$ (sample complexity to get $\mathbb{P}_A(\|\hat{A} - A\| > \varepsilon) < \delta$).

Example 2: $\mathcal{A} = \mathbb{R}^{d \times d}$, $O = (x_0, u_0, \dots, u_{t-1}, x_t)$ where $x_{t+a} = Ax_t + Bu_t + \eta_{t+1}$.
The adaptive control $(u_t)_{t \geq 0}$, adapted to the natural filtration, can be part of π.

The change-of-measure argument

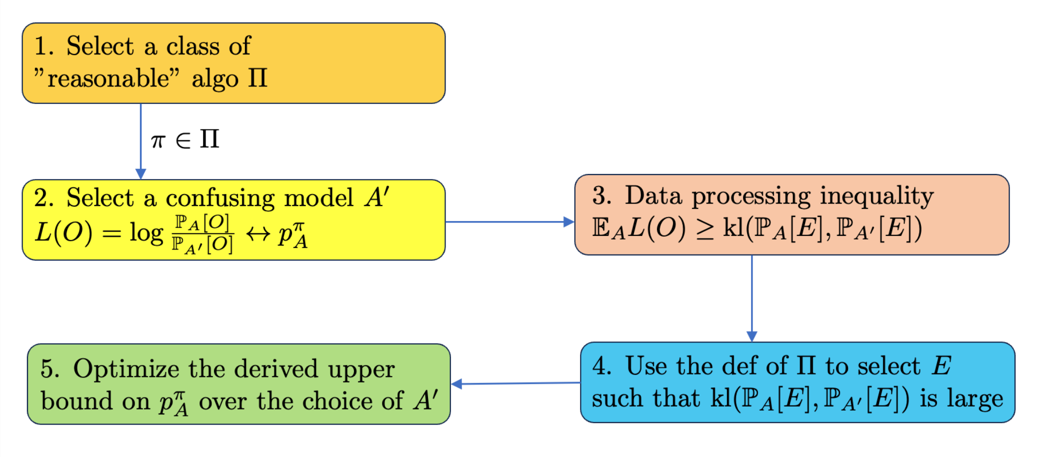

From change-of-measure to optimal algorithms: success stories

Lai–Robbins'85: Regret minimization in stochastic Multi-Armed Bandits

Recent success stories:

- Regret in infinite bandits (NIPS 2013)
- Regret in unimodal bandits (ICML 2014)
- Regret in Lipschitz bandits (COLT 2014)
- Regret in combinatorial bandits (NIPS 2015, NeurIPS 2023)
- Clustering in SBM (COLT 2014, NIPS 2014, NIPS 2015, NIPS 2016, NeurIPS 2019)
- Regret in bandits with generic structure (NIPS 2017)
- Regret in MDP with generic structure (NeurIPS 2018)
- Best arm identification in linear bandits (NeurIPS 2020)
- Regret in muti-agent bandits (AISTATs 2020)
- Best arm identification in structured bandits (NeurIPS 2021)
- Best policy identification in MDPs (ICML 2021, NeurIPS 2021)
- System identification in linear systems (IEEE TAC 2023)
- Networked bandits (ICML 2023, NeurIPS 2023)

Uncontrolled systems: $x_{t+1} = Ax_t + \eta_{t+1}$; fixed budget

Class Π of reasonable algorithms. An algorithm is (ε, δ)-locally stable in A if there exists τ such that $\forall t \geq \tau$ and $\forall A' \in B(A, 6\sqrt{2}\varepsilon)$, $\mathbb{P}_{A'}[\|\hat{A}_t - A'\| \leq \varepsilon] \geq 1 - \delta$.

Finite-time gramian. For any $t \geq 0$, $\Gamma_t(A) = \sum_{k=0}^{t} A^k (A^k)^\top$

Theorem 1 *For any A, for all $\varepsilon > 0$, $\delta \in (0,1)$, the sample complexity τ_A of any algorithm (ε, δ)-locally stable in A satisfies:*

$$\lambda_{\min}\Big(\sum_{s=1}^{\tau_A - 1} \Gamma_{s-1}(A)\Big) \geq \frac{1}{2\epsilon^2} \log(\frac{1}{2.4\delta}).$$

Uncontrolled systems: $x_{t+1} = Ax_t + \eta_{t+1}$; fixed budget

Class Π of reasonable algorithms. An algorithm is (ε, δ)-locally stable in A if there exists τ such that $\forall t \geq \tau$ and $\forall A' \in B(A, 6\sqrt{2}\varepsilon)$, $\mathbb{P}_{A'}[\|\hat{A}_t - A'\| \leq \varepsilon] \geq 1 - \delta$.

Finite-time gramian. For any $t \geq 0$, $\Gamma_t(A) = \sum_{k=0}^{t} A^k (A^k)^\top$

Theorem 2 *For any stable A, for all $\varepsilon \in (0, \|\Gamma_\infty(A)\|^{-3}/12)$, $\delta \in (0, 1/2)$, under any algorithm (ε, δ)-locally stable in A, we have: for $c > 0$ (universal constant)*

$$\lambda_{\min}\Big(\sum_{s=1}^{\tau_A - 1} \Gamma_{s-1}(A) \Big) \geq \frac{c}{\epsilon^2} \left(\log(\frac{1}{\delta}) + d \right).$$

Uncontrolled systems: $x_{t+1} = Ax_t + \eta_{t+1}$; fixed budget

A looser but explicit lower bound (exact for scalar systems).

$$\phi_a(t) = \sum_{s=1}^{t-1} \sum_{k=0}^{s-1} a^{2k} = \begin{cases} t-1 & \text{if} \quad a = 0, \\ \frac{a^{2t} + t(1-a^2) - 1}{(1-a^2)^2} & \text{if} \quad a \neq 1, \\ \frac{t(t-1)}{2} & \text{if} \quad a = 1. \end{cases}$$

For any estimator (ε, δ)-locally stable in A: $\phi_{|\lambda_d(A)|}(\tau_A) \geq \frac{c}{\epsilon^2}(\log(\frac{1}{\delta}) + d)$, where $\lambda_d(A)$ is the complex eigenvalue of A with smallest amplitude.

Proof of Theorem 1 (1/2)

1. Let π be (ε, δ)-locally stable in A, and let τ_A denote its sample complexity.

2. Select a confusing model A'. Let O denote the observation up to time τ_A, then

$$\mathbb{E}_A\left[\ln \frac{\mathbb{P}_A[O]}{\mathbb{P}_{A'}[O]}\right] = \frac{1}{2}\mathrm{tr}\Big((A - A')^\top (A - A') \sum_{s=1}^{\tau_A - 1} \underbrace{\mathbb{E}_A\left[x_s x_s^\top\right]}_{=\Gamma_{s-1}(A)}\Big).$$

(here τ_A is the performance metric – sample complexity)

3. Data processing inequality. For any event E O-measurable,

$$\frac{1}{2}\mathrm{tr}\Big((A - A')^\top (A - A') \sum_{s=1}^{\tau_A - 1} \Gamma_{s-1}(A)\Big) \geq \mathrm{kl}(\mathbb{P}_A[E], \mathbb{P}_{A'}[E]).$$

Proof of Theorem 1 (2/2)

4. Select the event E. For any $\forall A'$: $2\varepsilon < \|A' - A\| < 6\sqrt{2}\varepsilon$, define $E = \{\|\hat{A} - A\| \leq \varepsilon\}$. Since π is (ε, δ)-locally stable in A, $\mathbb{P}_A[E] \geq 1 - \delta$ and $\mathbb{P}_{A'}[E] \leq \delta$. Thus,

$$\mathrm{tr}\Big((A - A')^\top (A - A') \sum_{s=1}^{\tau_A - 1} \Gamma_{s-1}(A)\Big) \geq 2\mathrm{kl}(1 - \delta, \delta).$$

5. Optimizing over A' such that $2\varepsilon < \|A' - A\| < 6\sqrt{2}\varepsilon$. $M = (A - A')^\top (A - A')$.

$$\begin{aligned} \min_{M \succeq 0} \quad & \mathrm{tr}\Big(\sum_{s=1}^{\tau_A} \Gamma_{s-1}(A) M\Big) \\ s.t. \quad & \sigma_{\max}(M) \geq 4\epsilon^2 \end{aligned} \tag{1}$$

The value of the optimization problem is: $4\varepsilon^2 \lambda_{\min}\Big(\sum_{s=1}^{\tau_A - 1} \Gamma_{s-1}(A)\Big)$. □

Proof of Theorem 2 (1/2)

We need more confusing models! Let us look at the optimization problem (1) but with constraints $\|M\| \geq 18\varepsilon^2$. Its solution set is of the form:

$$\mathcal{C}^\star = \{A + 3\sqrt{2}\varepsilon v u^\top : v \in S^{d-1}\},$$

where u is the eivenvector corresponding to the smallest eigenvalue of $\sum_{s=1}^{\tau_A - 1} \Gamma_{s-1}(A)$.

Based on results on the unit sphere (**Wyner**'67), there is a packing $\mathcal{P} \subset \mathcal{C}^\star$ such that:

(i) for all $A' \in \mathcal{P}$, $\|A' - A\| = 3\varepsilon$;

(ii) for all $A_1, A_2 \in \mathcal{P}$, $3\varepsilon \leq \|A_1 - A_2\| < 6\sqrt{2}\varepsilon$;

(iii) $|\mathcal{P}| \geq \left(\frac{2}{\sqrt{3}}\right)^d$.

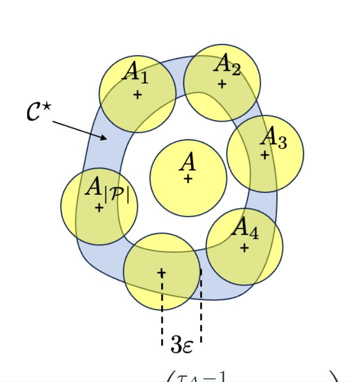

$$\forall A' \in \mathcal{P}, \qquad \mathrm{tr}\Big((A - A')^\top (A - A') \sum_{s=1}^{\tau_A - 1} \Gamma_{s-1}(A)\Big) = 18\varepsilon^2 \lambda_{\min}\left(\sum_{s=1}^{\tau_A - 1} \Gamma_{s-1}(A)\right).$$

Proof of Theorem 2 (2/2)

3. Data processing inequality. $\mathcal{P} = \{A_1, \dots, A_{|\mathcal{P}|}\}$. For all disjoint events E_i O-measurable:

$$\frac{1}{|\mathcal{P}|}\sum_{i=1}^{|\mathcal{P}|} \mathrm{tr}\left((A_i - A)^\top (A_i - A) \sum_{s=0}^{\tau_A - 2} \Gamma_s(A_i)\right) \geq \mathrm{kl}\left(\frac{1}{|\mathcal{P}|}\sum_{i=1}^{|\mathcal{P}|} \mathbb{P}_{A_i}(E_i),\ \frac{\mathbb{P}_A\left(\bigcup_{i=1}^{|\mathcal{P}|} E_i\right)}{|\mathcal{P}|}\right).$$

4. Let the events be $E_i = \{\|A_i - \hat{A}\| \leq \varepsilon\}$. Since π is (ε, δ)-locally stable in A,

$$\frac{1}{|\mathcal{P}|}\sum_{i=1}^{|\mathcal{P}|} \mathrm{tr}\left((A_i - A)^\top (A_i - A) \sum_{s=0}^{\tau_A - 2} \Gamma_s(A_i)\right) \geq 2\log(\frac{|\mathcal{P}|}{\delta}) \sim d + \log(1/\delta).$$

5. From $\Gamma_s(A_i)$ to $\Gamma_s(A)$. Perturbation of the gramians:

Proposition. Let A be stable. For all matrix Δ such that $4\|\Delta\| \leq \|\Gamma_\infty(A)\|^{-3/2}$, for all $s \geq 0$,

$$\|\Gamma_s(A + \Delta) - \Gamma_s(A)\| \leq 16\|\Delta\|\|\Gamma_\infty(A)\|^3.$$

□

Uncontrolled systems: $x_{t+1} = Ax_t + \eta_{t+1}$; fixed confidence

An algorithm (a stopping time τ and an estimator) is (ε, δ)-PAC if for all A, it stops almost surely in finite time, and $\mathbb{P}_A[\|\hat{A}_\tau - A\| \leq \varepsilon] \geq 1 - \delta$.

Theorem 3 *The sample complexity of an (ε, δ)-PACs satisfies:*

$$\lim_{\varepsilon,\delta \to 0} \frac{2\varepsilon^2 \lambda_{\min}\left(\sum_{s=1}^{\mathbb{E}_A[\tau]-1} \Gamma_{s-1}(A)\right)}{\log(1/3\delta)} \geq 1.$$

Why only asymptotic? Change-of-measure yields $\lambda_{\min}(\mathbb{E}_A[\sum_{s=0}^{\tau-1} x_s x_s^\top) \geq \frac{1}{2\varepsilon^2} \log(1/3\delta)$. Wald's lemma does not work for Markov processes, but asymptotically as $\mathbb{E}_A[\tau]$ grows large (**Moustakides**'99).

Controlled systems: $x_{t+1} = Ax_t + Bu_t + \eta_{t+1}$; fixed budget

Class Π of reasonable algorithms. An algorithm is (ε, δ)-locally stable in A if there exists τ such that $\forall t \geq \tau$ and $\forall A' \in B(A, 6\sqrt{2}\varepsilon)$, $\mathbb{P}_{A'}[\|\hat{A}_t - A'\| \leq \varepsilon] \geq 1 - \delta$.

Finite-time control gramians. for any $t \geq 0$, $G_t^u(A, B) = \sum_{s=0}^{t-1} x_s^u (x_s^u)^\top$ where $x_t^u = A^{t-1}Bu_0 + \ldots + ABu_{t-2} + Bu_{t-1}$.

Theorem 4[a] *For any A, for all $\varepsilon > 0$, $\delta \in (0,1)$, the sample complexity τ_A of any algorithm (ε, δ)-locally stable in A satisfies:*

$$\sup_{u \in \mathcal{U}} \lambda_{\min}\Big(G^u_{\tau_A - 1}(A, B) + \sum_{s=1}^{\tau_A - 1} \Gamma_{s-1}(A)\Big) \geq \frac{1}{4\epsilon^2} \log(\frac{1}{2.4\delta}).$$

[a] Adapted from **Wagenmaker–Jamieson**'20.

Controlled systems: two remarks

1. The lower bound specifies the optimal excitation! The deterministic control achieving the $\sup$ is optimal. But it depends on (A, B).

2. Dimension-dependent lower bound. Should the power of the control input be bounded, perturbation bounds of $G_t^u(A, B)$ may be established. Leading to:

Let $\mathcal{U}_\gamma$ denote the set of control signals with power bounded by γ. For any stable A, for all $\varepsilon > 0$, $\delta \in (0,1)$, under any algorithm (ε, δ)-locally stable in A, we have: for some $c > 0$ (universal constant)

$$\sup_{u \in \mathcal{U}_\gamma} \lambda_{\min}\Big(G^u_{\tau_A - 1}(A, B) + \sum_{s=1}^{\tau_A - 1} \Gamma_{s-1}(A) \Big) \geq \frac{c}{\epsilon^2} \left(\log(\frac{1}{\delta}) + d \right).$$

Part II. Instance-specific optimal algorithms
A. Uncontrolled systems / fixed budget

Least Squares Estimator

LSE: $A_t = \operatorname*{arg\,min}_{A \in \mathbb{R}^{d\times d}} \sum_{s=0}^{t} \|x_{s+1} - A x_s\|^2 = \left(\sum_{s=0}^{t} x_{s+1} x_s^\top\right)\left(\sum_{s=0}^{t} x_s x_s^\top\right)^\dagger$

Estimation error: $A_t - A = \left(\sum_{s=0}^{t} \eta_{s+1} x_s^\top\right)\left(\sum_{s=0}^{t} x_s x_s^\top\right)^\dagger$

Truncated block Toeplitz matrix:

$$\Gamma = \begin{bmatrix} I_d & & & & \\ A & I_d & & O & \\ & & \ddots & & \\ A^{t-2} & \dots & A & I_d & \\ A^{t-1} & \dots & \dots & A & I_d \end{bmatrix}$$

When $\rho(A) < 1$, $\|\Gamma\| \leq \mathcal{J}(A) = \sum_{s\geq 0} \|A^s\|$.

Optimality of the LSE

Theorem 5 *Let A be stable. For any $0 < \delta < 1$, and any $\varepsilon > 0$, we have: $\mathbb{P}\left(\|A_t - A\| > \varepsilon\right) < \delta$, as long as the following condition holds*

$$\lambda_{\min}\left(\sum_{s=0}^{t-1}\Gamma_s(A)\right) \geq c\max\left\{\frac{1}{\varepsilon^2}, \|\Gamma\|^2\right\}\left(\log(\frac{1}{\delta}) + d\right),$$

for some universal constant c.

For ε small enough, the LSE is optimal.
This proves a conjecture in **Simchowitz-Mania-Tu-Jordan-Recht**'18.

Corollary. For any $\delta \in (0,1)$ and ε small enough, the sample complexity of the LSE satisfies:

$$\tau_A \leq \frac{1}{\lambda_{\min}\left(\Gamma_\infty(A)\right)\varepsilon^2}\left(\log(\frac{1}{\delta}) + d\right).$$

Performance of the LSE for marginally stable systems

Jordan decomposition of A: SJS^{-1}, $C_A = \|S\|\|S^{-1}\|$, p size of the largest block of J.

Theorem 6 *Let A be marginally stable ($\rho(A) \leq 1$). For any $0 < \delta < 1$, and any $\varepsilon > 0$, the sample complexity of the LSE satisfies:*

$$\tau_A \leq \frac{c}{\varepsilon^2}\left(\log(\frac{1}{\delta}) + d\log(C_A) + dp\log(\frac{dp}{\varepsilon})\right)$$

for some universal constant c.

Main ingredients of the proof

$$\|A_t - A\| \leq \underbrace{\left\| \left(\sum_{s=0}^{t-1} \eta_s x_s^\top\right)\left(\sum_{s=0}^{t-1} x_s x_s^\top\right)^{-\frac{1}{2}}\right\|}_{\text{self-normalized process}} \underbrace{\left\|\left(\sum_{s=0}^{t-1} x_s x_s^\top\right)^{-\frac{1}{2}}\right\|}_{\text{smallest eigen value of the covariate matrix}}$$

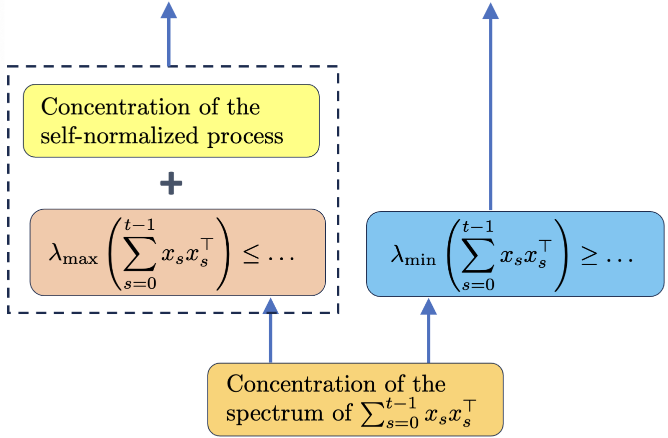

Spectrum of the covariates matrix

Theorem 7 *Let* $\varepsilon > 0$. *Let* $M = \left(\sum_{s=0}^{t-1} \Gamma_s(A)\right)^{-\frac{1}{2}}$ *and* $X^\top = (x_1, \ldots, x_t)$. *Then:*

$$\frac{1}{\|M\|}(1 - K^2\varepsilon) \leq s_d(X) \leq \cdots \leq s_1(X) \leq (1 + K^2\varepsilon)\frac{1}{s_d(M)}$$

holds with probability at least

$$1 - 2\exp\left(-c_1\varepsilon^2 \frac{1}{\|M\|^2\|\Gamma\|^2} + c_2 d\right),$$

for some universal constants $c_1, c_2 > 0$.

Concentration of the spectrum of the covariates matrix

Both $s_1(M)$ and $s_d(M)$ scale as $1/\sqrt{t}$, hence all singular values of the covariates matrix X scale as $\sqrt{t}$.

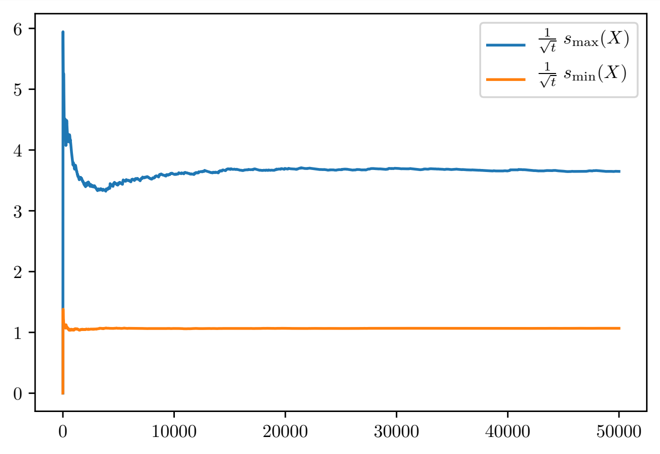

Proof of Theorem 7

1. **Approximate isometries.** Assume that $\|(XM)^\top XM - I_d\| \leq \max(\varepsilon, \varepsilon^2)$.
Then $\frac{1}{s_1(M)}(1-\varepsilon) \leq s_d(X) \leq \cdots \leq s_1(X) \leq (1+\varepsilon)\frac{1}{s_d(M)}$.

2. **Concentration on the set of isometries.**

Lemma 1 $\|(XM)^\top XM - I_d\| > \max(\varepsilon, \varepsilon^2)K^2$ *holds with probability at most*

$$2\exp\left(-c_1\varepsilon^2 \frac{1}{\|M\|^2\|\Gamma\|^2} + c_2 d\right)$$

for some positive absolute constants c_1, c_2.

3. **Concluding the proof.** Combine the two aforementioned results. $\square$

Proof of Lemma 1 (1/2)

1. $\|(XM)^\top XM - I_d\|$ **as the supremum of a chaos process.**

A *chaos* process is $(\xi^\top W\xi)_{W\in\mathcal{W}}$ with W is deterministic in $\mathbb{R}^{d\times d}$, ξ is random with independent coordinates. If ξ is isotropic, $\mathbb{E}[\xi^\top W\xi] = \mathrm{tr}W$.

Here $W = \Gamma^\top \sigma_{Mu}\sigma_{Mu}^\top \Gamma$. Specifically,

$$\|(XM)^\top XM - I_d\| = \sup_{u\in S^{d-1}} \left| \|\sigma_{Mu}{}^\top \Gamma\xi\|_2^2 - 1 \right|.$$

where $\xi^\top = (\eta_2^\top, \ldots, \eta_{t+1}^\top)$, and

$$\sigma_{Mu} = \begin{bmatrix} Mu & & & O \\ & Mu & & \\ & & \ddots & \\ O & & & Mu \end{bmatrix},$$

Proof of Lemma 1 (2/2)

2. Hanson-Wright inequality. Let $B \in \mathbb{R}^{m \times d}$, and $\xi \in \mathbb{R}^d$ be a random vector with zero-mean, unit-variance, sub-gaussian independent coordinates. Then for all $\varepsilon > 0$,

$$\mathbb{P}\left[\left|\|B\xi\|_2^2 - \|B\|_F^2\right| > \varepsilon \|B\|_F^2\right] \leq 2\exp\left(-c \min\left(\frac{\varepsilon^2}{K^4}, \frac{\varepsilon}{K^2}\right) \frac{\|B\|_F^2}{\|B\|^2}\right),$$

where c is an absolute positive constant and $K = \|\xi\|_{\psi_2}$.

Applying it to $B = \Gamma^\top \sigma_{Mu}$, one gets $\left|\|\sigma_{Mu}^\top \Gamma \xi\|_2^2 - 1\right| > \rho$ holds with probability at most

$$2\exp\left(-c \min\left(\frac{\rho^2}{K^4}, \frac{\rho}{K^2}\right) \frac{1}{\|M\|^2 \|\Gamma\|^2}\right).$$

3. ϵ-net argument. The concentration of the supremum over u is obtained using classical ϵ-net arguments.

□

Proof of Theorem 6 (1/3)

1. Upper bound on $\lambda_{\max}(\sum_{s=0}^{t-1} x_s x_s^\top)$. We establish that: $\mathcal{J}_t(A) = \sum_{s=0}^{t} \|A^s\| \leq (t+1)^p C_A$, and that

$$\lambda_{\max}\left(\sum_{s=0}^{t-1} \Gamma_s(A)\right) \leq (t+1)^{2p} C_A^2.$$

We conclude using the concentration result on the spectrum of $\sum_{s=0}^{t-1} x_s x_s^\top$.

2. Lower bound on $\lambda_{\min}(\sum_{s=0}^{t-1} x_s x_s^\top)$. A generic decomposition (works for controlled systems): if $x_t = y_t + z_t$, for all $\lambda > 0$:

$$\sum_{s=0}^{t} x_s x_s^\top \succeq \sum_{s=0}^{t} z_s z_s^\top - \left(\sum_{s=0}^{t} y_s z_s^\top\right)^\top \left(\sum_{s=0}^{t} y_t y_t^\top + \lambda I_d\right)^{-1} \left(\sum_{s=0}^{t} y_s z_s^\top\right) - \lambda I_d.$$

Use concentration of self-normalized processes[4] to upper bound the middle term.

[4] **Peña-Lai-Shao**'09

Proof of Theorem 6 (2/3)

Apply it to $y_s = Ax_{s-1}$ and $z_s = \eta_s$ to get:

$$\mathbb{P}\left(\lambda_{\min}\left(\sum_{s=0}^{t} x_s x_s^\top\right) \gtrsim t - d\log\left(\frac{\lambda_{\max}\left(\sum_{s=0}^{t-1}(Ax_s)(Ax_s)^\top\right)}{\lambda} + 1\right) - \log(1/\delta) - \lambda\right) \geq 1-\delta$$

provided that the following condition holds $T \gtrsim d + \log(1/\delta)$.

$+$ upper bound on $\lambda_{\max}(\sum_{s=0}^{t-1} x_s x_s^\top)$ yields:

$$\text{When } t \gtrsim d\log\left(C_A^5 T^{3p}\right) + \log(1/\delta), \qquad \mathbb{P}\left(\lambda_{\min}\left(\sum_{t=0}^{T} x_s x_s^\top\right) \gtrsim t\right) \geq 1-\delta$$

3. Self-normalized process.

$$\mathbb{P}\left[\left\|\left(\sum_{s=0}^{t}\eta_s x_s^\top\right)^\dagger \left(\sum_{s=0}^{t} x_s x_s^\top + \Lambda\right)^{-1/2}\right\|^2 \lesssim \log\left(\frac{\det^{1/2}(\sum_{s=0}^{t} x_s x_s^\top + \Lambda)}{\det^{1/2}(\Lambda)\delta}\right) + d\right] \geq 1-\delta.$$

$+$ upper bound on $\lambda_{\max}(\sum_{s=0}^{t} x_s x_s^\top)$ to control $\det(\sum_{s=0}^{t} x_s x_s^\top + \Lambda)$.

Proof of Theorem 6 (3/3)

4. **Putting things together.** We obtain:

$$\|A_t - A\|^2 \lesssim \frac{d\log(C_A^5(t+1)^{3p}+t)+d+\log(1/\delta)}{t}$$

provided that $t \gtrsim d\log(C_A) + dp\log(dp) + \log(1/\delta)$.

Renormalizing with ε, we get finally $\mathbb{P}[\|A_t - A\| < \varepsilon] \geq 1-\delta$ whenever

$$t \leq \frac{c}{\varepsilon^2}\left(\log(\frac{1}{\delta}) + d\log(C_A) + dp\log(\frac{dp}{\varepsilon})\right).$$

□

Part II. Instance-specific optimal algorithms
B. Uncontrolled systems / fixed confidence

Why the fixed-confidence setting?

After t samples, an optimal estimator (alone) has, w.p. $1-\delta$, an error equal to

$$\|A_t - A\| = c\sqrt{\frac{\log(\frac{1}{\delta}) + d}{t\lambda_{\min}(\Gamma_\infty(\boldsymbol{A}))}}.$$

It depends on the unknown A.

So when do you know when to stop gathering samples? **Add a stopping rule.**
(This is also very useful in adaptive control)

An optimal stopping rule for the LSE

Objective. Stop as soon as (ε, δ)-PAC guarantees are achieved.

Algorithm.

1. Stopping rule (based on the GLRT, **Chernof**'59): stop at τ where

$$\tau = \inf\left\{ t \geq 1 : \ \lambda_{\min}\left(\sum_{s=0}^{t-1} x_s x_s^\top\right) > \beta(\varepsilon, \delta, t) \vee \mu \right\}$$

$$\beta(\varepsilon, \delta, t) = \frac{(2\sigma)^2}{(1-\alpha)\varepsilon^2} \log\left(\frac{5^d \det(\sum_{s=0}^{t-1} x_s x_s^\top + \frac{\alpha\mu}{1-\alpha} I_d)^{\frac{1}{2}}}{\delta \det(\frac{\alpha\mu}{1-\alpha} I_d)^{\frac{1}{2}}} \right)$$

   for some tuned $\mu > 0$ and $\alpha \in (0, 1)$.

2. Estimator: the LSE A_τ.

Optimality of (Chernoff's stopping rule + LSE)

Define the increasing function ϕ_A as $\phi_A(t) = \lambda_{\min}\left(\sum_{s=1}^{t-1} \Gamma_{s-1}(A)\right)$.

Theorem 8 *For all $\varepsilon > 0$ and all $\delta \in (0,1)$, the LSE combined with the stopping rule τ is (ε, δ)-PAC and the expected sample complexity satisfies:*

$$\phi_A(\mathbb{E}[\tau]) \lesssim K^4 \max\left\{\frac{1}{\varepsilon^2}, \mathcal{J}_A^2\right\}\left(C_{A,d,\varepsilon,K}\, d + \log\left(\frac{1}{\delta}\right)\right),$$

where $C_{A,d,\varepsilon,K} \lesssim \log\left(\frac{K^6 \mathcal{J}_A^2 \|\Gamma_\infty(A)\|^2 d}{\varepsilon^2}\right)$. where $\lesssim$ hides constant that can only depend on μ and α.

Part II. Instance-specific optimal algorithms
C. Controlled systems / fixed budget

Based on **Wagenmaker-Jamieson**
"Active Learning for Identification of Linear Dynamical Systems", COLT 2020.

Sample complexity lower bound

$G_t^u(A,B) = \sum_{s=0}^{t-1} x_s^u (x_s^u)^\top$ where $x_t^u = A^{t-1}Bu_0 + \ldots + ABu_{t-2} + Bu_{t-1}$.

The sample complexity τ_A of any algorithm (ε, δ)-locally stable in A satisfies:

$$\sup_{u \in \mathcal{U}} \lambda_{\min}\Big(G^u_{\tau_A - 1}(A,B) + \sum_{s=1}^{\tau_A - 1} \Gamma_{s-1}(A)\Big) \geq \frac{1}{4\epsilon^2} \log(\frac{1}{2.4\delta}).$$

For k-periodic inputs with average power γ, $u \in \mathcal{U}_\gamma(k)$, define $\Gamma_k^u(A,B) := \lim_{t\to\infty} \frac{1}{\gamma^2 t} \sum_{s=0}^{t-1} x_s^u (x_s^u)^\top$. The lower bound becomes:

$$\tau_A \geq \frac{1}{8\varepsilon^2} \frac{1}{\max_{u \in \mathcal{U}_\gamma} \lambda_{\min}(\gamma^2 \Gamma^u_\infty(A,B) + \Gamma_\infty(A))} \log(\frac{1}{2.4\delta}),$$

where

$$\max_{u \in \mathcal{U}_\gamma} \lambda_{\min}(\gamma^2 \Gamma^u_\infty(A,B) + \Gamma_\infty(A)) := \lim_{k\to\infty} \max_{u \in \mathcal{U}_\gamma(k)} \lambda_{\min}(\gamma^2 \Gamma^u_k(A,B) + \Gamma_k(A))$$

A certainty equivalence-based algorithm

Main idea. The optimal inputs depends on A. Replace A by its current LSE $\hat{A}$ and solve $\max_{u \in \mathcal{U}_\gamma} \lambda_{\min}(\gamma^2 \Gamma^u_\infty(\hat{A}, B) + \Gamma_\infty(\hat{A}))$ to drive the excitation.

Algorithm.

Initialization. Set $T_0 = 100$, $k_0 = 1$
Observe T_0 without excitation
$\hat{A}_0 \leftarrow$ LSE based on data up to T_0
$u^1 \leftarrow \arg\max_{u \in \mathcal{U}_{\gamma/\sqrt{2}}(2k_0)} \lambda_{\min}(\gamma^2 \Gamma^u_{2k_0}(\hat{A}_0, B) + \Gamma_{2k_0}(\hat{A}_0))$

For $i = 1, 2, \ldots$, **do**

1. $T_i \leftarrow 3T_{i-1}$, $T \leftarrow T + T_i$, $k_{i+1} = 2k_i$
2. Observe T_i samples with $u_t = u^i_t + \eta^u_t$ with $\eta^u_t \sim \mathcal{N}(0, \frac{\gamma^2}{m} I)$
3. $\hat{A}_i \leftarrow$ LSE based on data up to T
4. $u^{i+1} \leftarrow \arg\max_{u \in \mathcal{U}_{\gamma/\sqrt{2}}(k_{i+1})} \lambda_{\min}(\gamma^2 \Gamma^u_{k_{i+1}}(\hat{A}_i, B) + \Gamma_{k_{i+1}}(\hat{A}_i))$

Asymptotic optimality of the algorithm

Theorem 9 *For $\varepsilon > 0$ small enough, the sample complexity of the previous algorithm satisfies: for some universal constant $C > 0$,*

$$\lim_{\delta \to 0} \frac{\tau_A}{\log(\frac{1}{\delta})} \leq \frac{C}{\varepsilon^2} \frac{1}{\max_{u \in \mathcal{U}_\gamma} \lambda_{\min}(\gamma^2 \Gamma_\infty^u(A, B) + \Gamma_\infty(A))}$$

Remark. What do we win over pure but colored noise? $\mathrm{sp}(A) = \lambda \in \mathbb{R}_+^d$, colored noise input $u_t \sim \mathcal{N}(0, \Sigma^\star)$ where $\Sigma^\star$ is optimized.

Sample complexity with colored noise:	$\Theta(\gamma^2 / \|\mathbf{1} - \lambda\|_1)$
Minimal sample complexity:	$\Theta(\gamma^2 / \|\mathbf{1} - \lambda\|_2^2)$

Example: $\lambda_i = 1 - 1/d$, a factor d won with the optimal excitation vs colored noise.

Conclusions

- **Minimal sample complexity to get (ε, δ)-PAC guarantees for linear sysID:**

$$\text{without excitation} \qquad \Theta\left(\frac{d + \log(1/\delta)}{\varepsilon^2 \lambda_{\min}(\Gamma_\infty(A))}\right)$$

$$\text{with excitation} \qquad \Theta\left(\frac{d + \log(1/\delta)}{\varepsilon^2 \max_{u \in \mathcal{U}_\gamma} \lambda_{\min}(\gamma^2 \Gamma_\infty^u(A, B) + \Gamma_\infty(A))}\right)$$

- **"Optimal" algorithms:** LSE + certainty equivalence for active excitation
- **Next challenges:**
  - Towards truly optimal and computationally efficient active learning algorithms
  - Non fully observable states
  - Non-linear system identification
  - A systematic understanding of lower bounds in the moderate confidence regime

Q&A

For more questions and feedback, please contact us:

jedra@mit.edu, jedra@kth.se

alepro@kth.se