Institut für Systemtheorie und Regelungstechnik
Universität Stuttgart

PAUL G. ALLEN SCHOOL OF COMPUTER SCIENCE & ENGINEERING

# High Effort, Low Gain: Fundamental Limits of Active Learning for Linear Dynamical Systems


Nicolas Chatzikiriakos[1], Kevin Jamieson[2], Andrea Iannelli[1]

[1] Institute for Systems Theory and Automatic Control, University of Stuttgart
[2] Paul G. Allen School of Computer Science & Engineering, University of Washington


## Active Learning in Dynamical Systems

Learning the behavior of an unknown dynamical system is a central problem in many applications

- RL: Exploration required for good performance, but how to explore efficiently is unclear
- Robotics: Dynamics of robots can be hard to model
- Airplane control: Different modes can be (de-) activated, and the current configuration needs to be detected

**Problem:** Obtaining **data** for learning often is **expensive** or induces a loss in performance

➡ Active Learning methods select the excitation input to **accelerate** the **learning** process

### Limitations

- Isotropic Gaussian excitations are default exploration strategy, benefit of active learning methods often unclear
- Sample complexity lower bounds are restricted to specific input classes
- No modular framework to establish sample complexity upper bounds for general inputs
- Multi-model perspective has not been analyzed

### Contributions

- Instance-dependent sample complexity lower bounds uncovering the potential **benefit of active learning**
- Tailored persistency of excitation (PE) notion, characterizing inputs yielding informative data
- **Modular framework** to provide sample complexity upper bounds for any PE input sequence
- Provably optimal sequential active learning algorithm with **close to oracle performance**

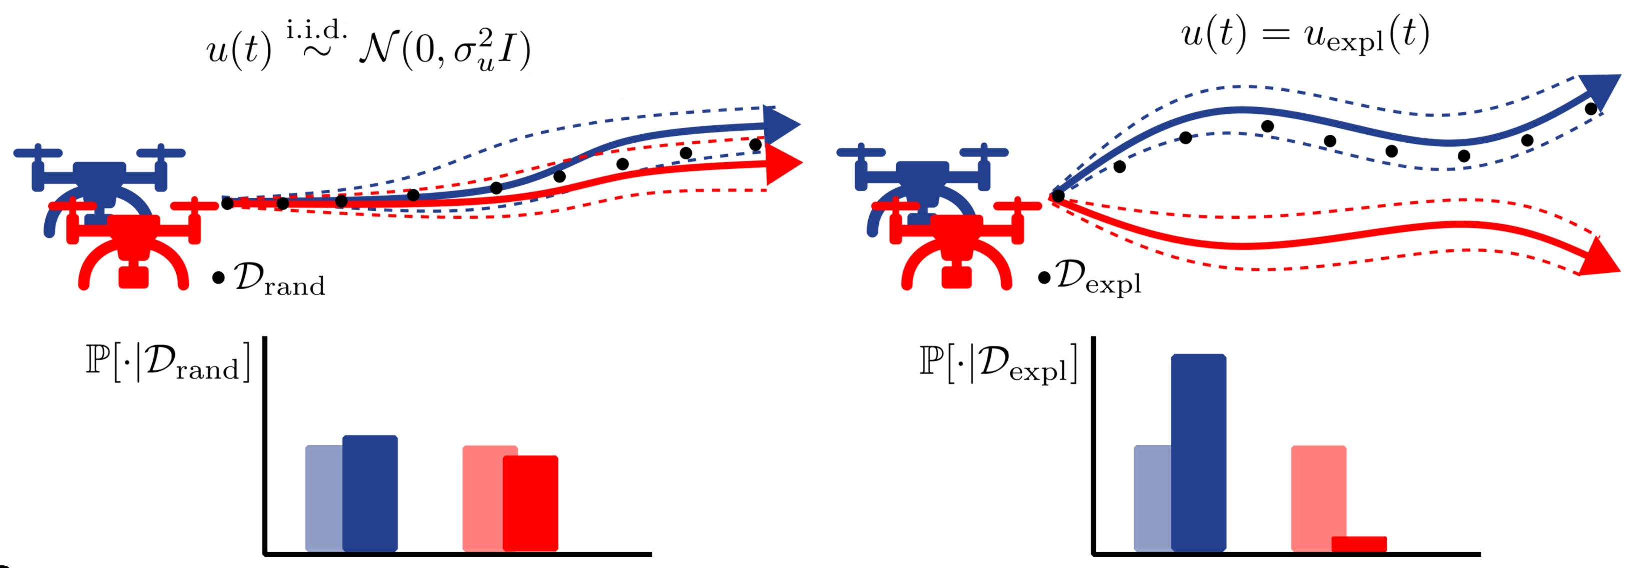


**Setup:**

- Unknown LTI system $x(t+1) = A_* x(t) + B_* u(t) + w(t), \quad \theta_* = (A_*, B_*)$
- Process noise $w(t) \overset{\text{i.i.d.}}{\sim} \mathcal{N}(0, \Sigma_w)$
- Finite hypothesis class $(A_*, B_*) \in \mathcal{S} = \{(A_0, B_0), \dots, (A_N, B_N)\}$
- Bounded input energy $\mathbb{E}\left[\sum_{t=0}^{T-1} \|u(t)\|^2\right] \leq \gamma_u^2 T$

- **Goal:** Identify the true system from data as fast as possible
- **Key Fact:** For simplicity, assume $x(0) = 0$. Define $U = \begin{bmatrix} u(0)^\top & \dots & u(t)^\top \end{bmatrix}^\top$, then

$$\mathbb{E}\left[\sum_{s=0}^{t} \|\Delta A_i x(t) + \Delta B_i u(t)\|^2_{\Sigma_w^{-1}}\right] = U^\top W_i(t) U + \text{Excitation due to noise.}$$

- $W_i$ captures the difficulty of distinguishing between $\theta_i$ and $\theta_*$ when applying input sequence $U$
- $\Delta A_i = A_* - A_i, \, \Delta B_i = B_* - B_i$

## Sample Complexity Lower Bound

- **$\delta$-correct algorithm:** Any algorithm that yields $\mathbb{P}[\hat{\theta}_t \neq \theta_*] \leq \delta$, for $t \geq \bar{T}$.
- **Sample complexity lower bound:** For any excitation input sequence $U$ and any $\delta$-correct algorithm, it holds that

$$\min_{i \in [1,N]} \mathbb{E}[U^\top W_i(\bar{T}) U] + \text{Excitation due to noise} \geq 2 \log\left(\frac{1}{2.4\delta}\right)$$

Further the lower bound is minimized by

$$U^* \in \underset{U^\top U \leq \gamma_u^2 \bar{T}}{\arg\max} \min_{i \in [1,N]} U^\top W_i U + \text{Excitation due to noise}$$

and when applying $U^*$ any $\delta$-correct algorithm satisfies

$$\min_{p \in \Delta_p} \gamma_u^2 \bar{T} \lambda_{\max}\left(\sum_{i=1}^{N} p_i W_i(\bar{T})\right) + \text{Excitation due to noise} \geq 2 \log\left(\frac{1}{2.4\delta}\right)$$

Finally, for $u(t) \overset{\text{i.i.d.}}{\sim} \mathcal{N}(0, \frac{\gamma_u^2}{n_u} I)$ and any $\delta$-correct algorithm it holds that

$$\min_{p \in \Delta_N} \gamma_u^2 \bar{T} \lambda_{\text{mean}}\left(\sum_{i=1}^{N} p_i W_i(\bar{T})\right) + \text{Excitation due to noise} \geq 2 \log\left(\frac{1}{2.4\delta}\right)$$

➡ The matrices $W_i$ are the key quantities influencing the sample complexity lower bound

➡ Condition number of $W_i$ determines the potential benefit of active learning

## Sequential Active Learning Algorithm

- Optimal oracle excitation requires exact model knowledge ➡ Sequential algorithm

**Algorithm 2** Input design subroutine

**Require:** $\mathcal{S}$, epoch length $\tau$, prediction errors $\varepsilon_{\theta_i}((k-1)\tau)$, state $x$, scaling $\rho_k \in [0,1]$

Compute weights using exponential weighting $w_{k+1}(i) = \exp\left(-\eta \varepsilon_{\theta_i}((k-1)\tau)\right)$

Sample $\hat{i}_k \sim p_k(i)$, where $p_k(i) = w_k(i) / \sum_{j=0}^{N} w_k(j)$

Set $\hat{\theta}_k = \theta_{\hat{i}_k}$ and compute optimal input sequence $U^*_{\hat{\theta}_k}(x)$ by solving (28)

Define excitation according to $\rho_k$: $u^*_k(t) = \sqrt{1-\rho_k} u^*_{\hat{\theta}_k}(t) + \sqrt{\rho_k} u_\eta(t)$, $u_\eta(t) \overset{\text{i.i.d.}}{\sim} \mathcal{N}(0, \frac{\gamma_u^2}{n_u} I_{n_u})$

**return** excitation sequence $u^*_k$

- **Optimality of Alg. 2:** Let $\rho_k \overset{k \to \infty}{\longrightarrow} 0$ slow enough. Then, for any $\delta \in (0,1)$ the inputs derived by Alg. 2 yield $\mathbb{P}[\hat{\theta}_k \neq \theta_*] \leq \delta$ after at most $k_\delta$ epochs with

$$\lim_{\delta \to 0} \frac{c' \log\left(\frac{N}{\delta}\right)}{k_\delta} \leq \gamma_u^2 \min_{p \in \Delta_N} \lambda_{\max}\left(\sum_{i=1}^{N} p_i W_i(\tau)\right) + \text{Excitation due to noise}$$

- Input derived by Alg. 2 are optimal as $\delta \to 0$
- Sample complexity upper bound follows directly from PE

## PE and Sample Complexity Upper Bound

- **Persistency of Excitation (PE):** An input sequence is PE for $(\theta_*, \mathcal{S})$ if

$$\frac{1}{\tau} \sum_{t=0}^{\tau-1} \mathbb{E}\left[\|\Delta A_i x(t) + \Delta B_i u(t)\|^2_{\Sigma_w^{-1}}\right] \geq c_u(\tau) \gamma_u^2 + \text{Excitation due to noise}$$

for a constant $c_u(\tau)$.

- Isotropic Gaussian excitation is PE with $c_u^{\text{rand}} = \min_{p \in \Delta_N} \lambda_{\text{mean}}\left(\sum_{i=1}^{N} p_i W_i(\tau)\right)$
- $U^*$ is PE with $c_u^{\text{opt}} = \min_{p \in \Delta_N} \lambda_{\max}\left(\sum_{i=1}^{N} p_i W_i(\tau)\right)$

**Algorithm 1** Sequential identification algorithm

**Require:** $\mathcal{S}$, epoch length $\tau$, desired confidence $\delta$

**for** $k = 1, 2, \dots$ **do**

  Collect data using PE excitation input $\{u(t)\}_{t=(k-1)\tau}^{k\tau-1}$ with coefficients $c_{u_k}(\tau)$, $c_w(\tau)$

  Compute $\varepsilon_{\theta_i}(k\tau)$ for all $i \in [0, N]$

  **if** $\exists \hat{\theta} \in \mathcal{S} : \varepsilon_{\theta_i}(k\tau) - \varepsilon_{\hat{\theta}}(k\tau) > 2 \log\left(\frac{N}{\delta}\right)$ for all $\theta_i \in \mathcal{S} \setminus \hat{\theta}$ **then**

    Stop and **return** estimate $\hat{\theta}$

- Prediction error $\varepsilon_{\theta_i}(t) = \sum_{s=0}^{t-1} \|x(s+1) - A_i x(s) - B_i u(s)\|^2_{\Sigma_w^{-1}}$
- **Sample complexity upper bound:** Algorithm 1 yields an estimate $\hat{\theta}$ satisfying $\mathbb{P}[\hat{\theta} \neq \theta_*] \leq \delta$ and terminates no later than when $k$ satisfies

$$\sum_{j=1}^{k} c_{u_j}(\tau) \gamma_u^2 + \text{Excitation due to noise} \geq c \log\left(\frac{N}{\delta}\right).$$

## Numerical Experiments

- Comparison of two problem setups with $\Sigma_w = I, \, \gamma_u = 1, \, n_x = 3, \, n_u = 2, \, \eta = 0.01$

  (a) Structured hypothesis class: Experiment design yields large benefits ($W_i$ ill-conditioned)

  (b) Randomly generated hypothesis class: Isotropic random excitation is closer to optimal

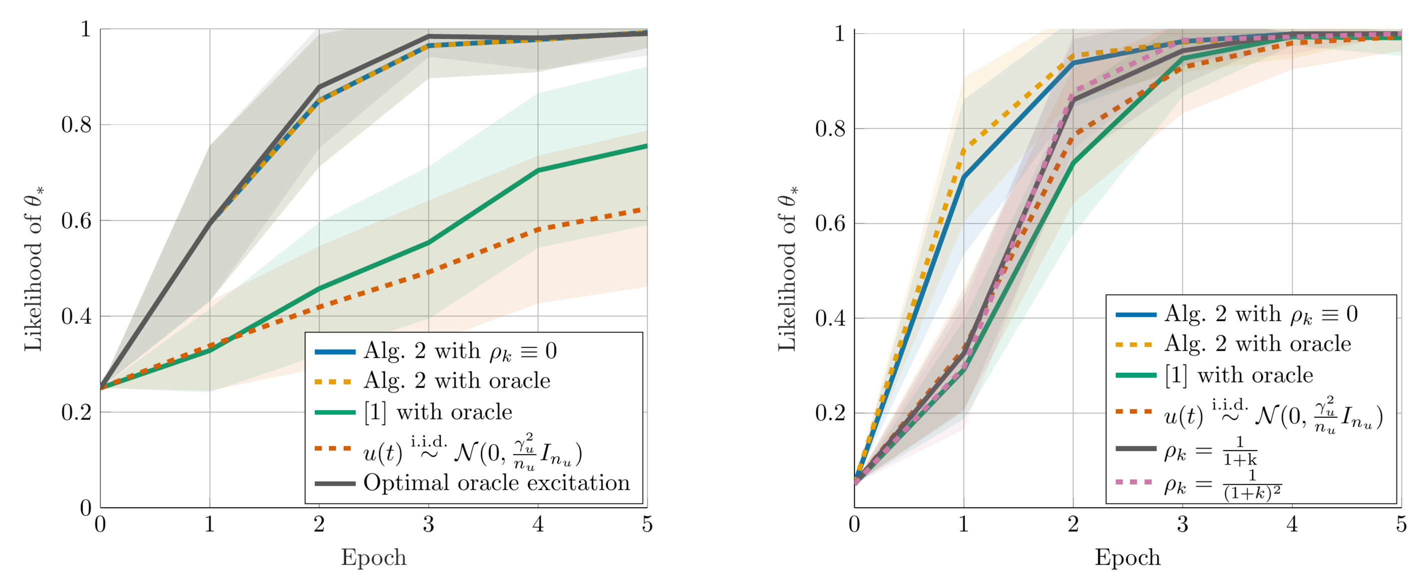


(a) Structured set $\mathcal{S}$ (b) Random set $\mathcal{S}$

- Proposed algorithm achieves close to optimal performance without oracle knowledge
- Approach in [1] neglects finite hypothesis class $\mathcal{S}$, yielding significantly suboptimal performance
- CE ($\rho_k = 0$) empirically performs best
- Gap between isotropic Gaussian and optimal input increases when $W_i$ is ill-conditioned


[1] Wagenmaker, A., Jamieson, K. (2020). Active Learning for Identification of Linear Dynamical Systems. *Proceedings of Thirty Third Conference on Learning Theory, 2020*


## Conclusion & Outlook

**Summary**

- Benefit of experiment design is problem-dependent and is captured in matrices $W_i$
- Sequential identification algorithm that allows for modular analysis for any PE input
- Sequential active learning algorithm with
  - sample complexity guarantees
  - empirically close to oracle performance using CE

**Outlook**

- Analysis of computational complexity
- Extensions to task-optimal active learning
- Connecting to online settings, e.g., RL


**Institute for Systems Theory and Automatic Control**
University of Stuttgart
www.ist.uni-stuttgart.de

**Contact:**
Nicolas Chatzikiriakos
nicolas.chatzikiriakos@ist.uni-stuttgart.de


SimTech

University of Stuttgart Germany