University of Stuttgart
Germany

SimTech

PAUL G. ALLEN SCHOOL
OF COMPUTER SCIENCE & ENGINEERING

ist

Annual Conference on Artificial Intelligence and Statistics, May 2, 2026

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]Allen School of Computer Science & Engineering
University of Washington

Identification from a finite hypothesis class

- Prediction/Control requires an accurate model

$$x(t+1) = A_* x(t) + B_* u(t) + w(t), \qquad w(t) \overset{\text{i.i.d.}}{\sim} \mathcal{N}(0, \Sigma_w)$$

- Hypothesis class: $\{(A_i, B_i)\}_{i=0}^{N} = \Big\{$ 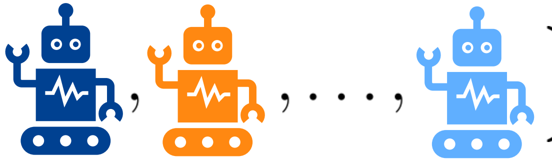 $\Big\}$

Input sequence $\{u_t\}_{t=1}^{T}$

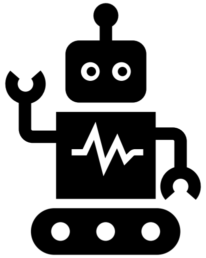

State sequence $\{x_t\}_{t=1}^{T}$

How can we identify the true system from a finite hypothesis class?		
Non-i.i.d. data	Probabilistic noise	Finite samples

Which questions do we answer?

How can we generate data for sample-efficient identification?

What can we gain from active learning over random excitations?

Can we unify sample complexity analysis across excitation strategies?

Designing sample-efficient excitation strategies

Algorithm 1 Sequential identification algorithm

Require: Hypothesis class $\{\theta_i\}_{i=0}^{N}$, desired confidence δ

1: **for** epoch $= 1, 2, \ldots$ **do**

(1) Data collection — Experiment Design

2: Sample model index k using exponential weighting

3: Compute optimal excitation for θ_k (using CE)

4: Collect data with ε-greedy excitation

(2) Identification — Identification

5: Compute empirical prediction errors for all candidates

6: **if** hypothesis test on errors succeeds **then**

7: Stop and **return** estimate

Theorem: Correctness & Optimality

With probability $1 - \delta$, Algorithm 1 terminates with **correct estimate** and achieves **optimal sample complexity** (as $\delta \to 0$)

How much do we gain from active learning?

Low gain regime

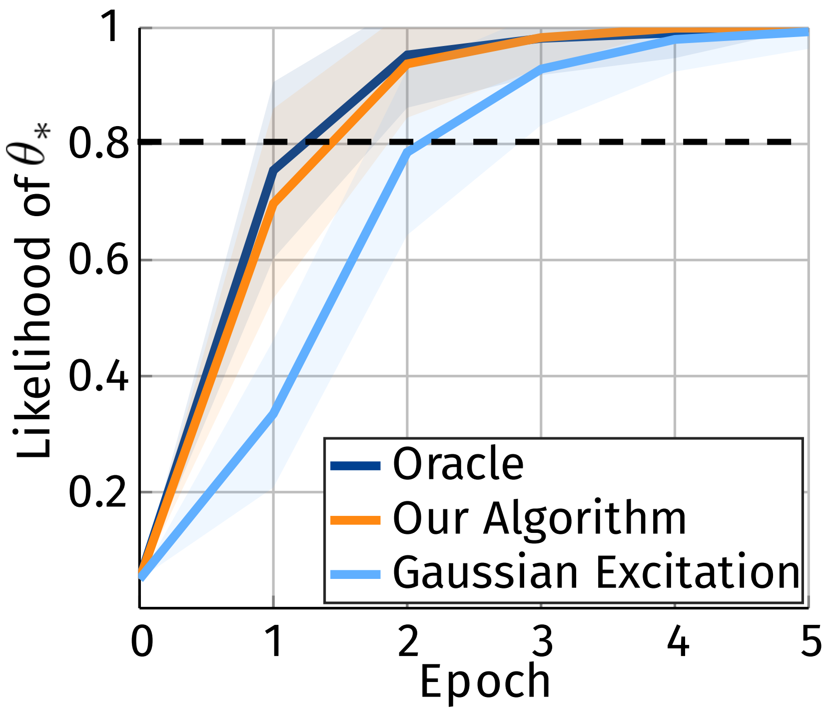

How much do we gain from active learning?

Low gain regime

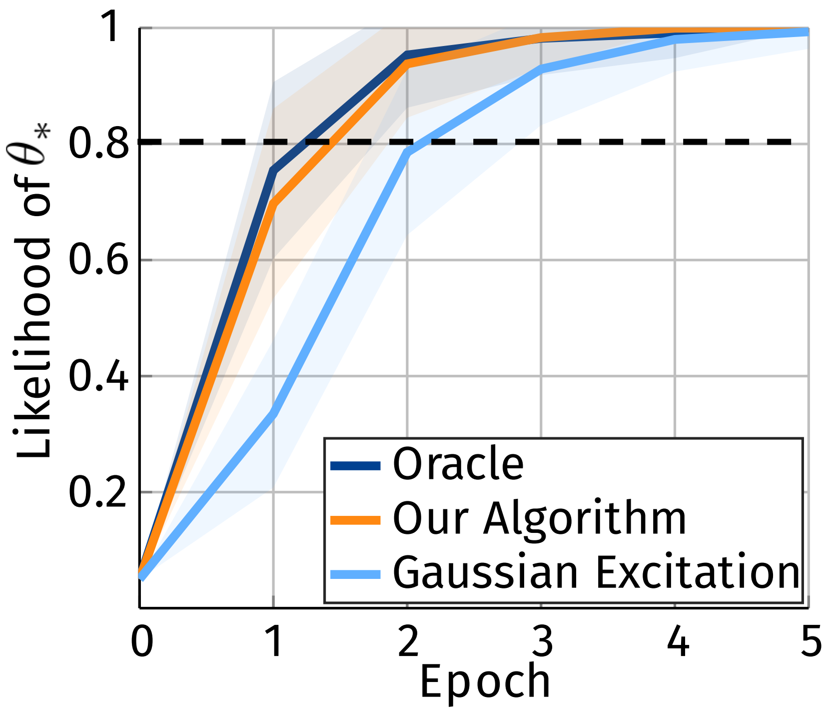

High gain regime

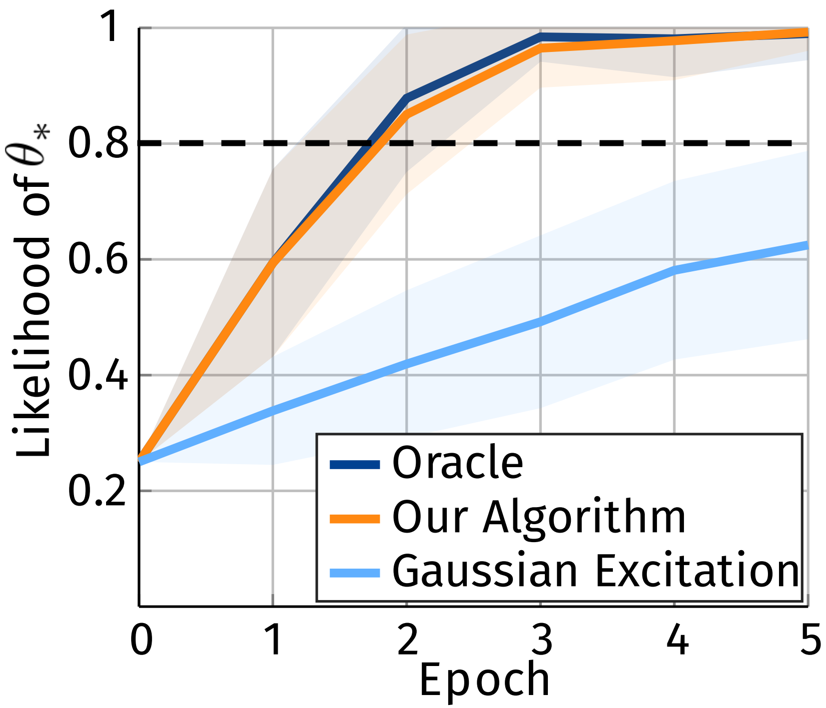

Why do gains vanish?

- Problem structure determines limits of experiment design
- Our theory provides the explanation

What we learned

Provably optimal active learning algorithm

Unified sample complexity analysis of excitation schemes

Characterization of fundamental limits in active learning

Poster Session 1, Poster 190

Nicolas Chatzikiriakos

Institute for Systems Theory and Automatic Control

University of Stuttgart

nicolas.chatzikiriakos@ist.uni-stuttgart.de

We acknowledge the support DFG under Germany's Excellence Strategy - EXC 2075–390740016