




PAC-Bayesian Optimal Control with Stability and Generalization Guarantees

Mahrokh Ghoddousi Boroujeni



Mahrokh Boroujeni

Atinary Technologies

Clara Galimberti

SUPSI

Andreas Krause

ETH

Giancarlo Ferrari-Trecate

EPFL

A PAC-Bayesian framework for optimal control with stability guarantees

MG Boroujeni, CL Galimberti, A Krause, G Ferrari-Trecate

IEEE 63rd Conference on Decision and Control (CDC), 8237-8244

*Outstanding Student
Paper Award*

PAC-Bayesian Optimal Control with Stability and Generalization Guarantees



MG Boroujeni, CL Galimberti, A Krause, G Ferrari-Trecate

arXiv preprint arXiv:2512.02858

Flexibility & Success



Computer vision




Flexibility & Success



Computer vision



Natural language processing



Flexibility & Success



Computer vision



Natural language processing



Self-driving vehicles



Flexibility & Success



Computer vision



Natural language processing




Self-driving vehicles

Limitations & Challenges



Overfitting



Flexibility & Success



Computer vision



Natural language processing



Self-driving vehicles

Limitations & Challenges



Overfitting



Hallucinations in LLMs



Generate Kalman's
Theorem on hallucinations

Theorem (Kalman, 1961*).

A language model exhibits bounded hallucination if and only if its internal representations are observable and its attention dynamics are controllable with respect to the semantic state evolution. Formally, let $\mathcal{M}$ denote the transformer architecture with state transition operator A_{attn} and output decoder C_{tok} . Then the hallucination error ε_t remains bounded for all t if the pair $(A_{\text{attn}}, C_{\text{tok}})$ is observable and $(A_{\text{attn}}, B_{\text{prompt}})$ is controllable under a bounded input distribution.



Flexibility & Success



Computer vision



Natural language processing



Self-driving vehicles

Limitations & Challenges



Overfitting



Hallucinations in LLMs



Lack of robustness



Flexibility & Success



Computer vision



Natural language processing



Self-driving vehicles

Limitations & Challenges



Overfitting



Hallucinations in LLMs



Lack of robustness




Flexibility $\neq$ Reliability

Data-driven control calls for reliability guarantees
beyond training data



Flexibility $\neq$ Reliability

Closed-loop stability




- System trajectories remain bounded
- Converge to desired working points




Flexibility $\neq$ Reliability

Closed-loop stability



- System trajectories remain bounded
- Converge to desired working points


Generalization






Flexibility $\neq$ Reliability

Closed-loop stability



- System trajectories remain bounded
- Converge to desired working points

Generalization




- Performance on unseen realizations
- Robust to uncertainty



Flexibility $\neq$ Reliability


Closed-loop stability



- System trajectories remain bounded
- Converge to desired working points



Generalization




- Performance on unseen realizations
- Robust to uncertainty

Goal of the talk: develop a reliable learning framework
by ensuring *stability* and *generalization*

How to Learn Reliable Controllers?

Learn **flexible** controllers, from **finite** data


Closed-loop stability



Methodology: Parameterization of stability-preserving neural controllers

- Enables learning flexible controllers
- Guarantees closed-loop stability

Generalization



Methodology: PAC-Bayesian learning

- Bound true control performance from finite data
- Incorporate prior knowledge


EPFL Stochastic Nonlinear Optimal Control (SNOC)

System

- Nonlinear, discrete time, stable/pre-stabilized, stochastic

Process noise

w →




EPFL Stochastic Nonlinear Optimal Control (SNOC)

System

- Nonlinear, discrete time, stable/pre-stabilized, stochastic

Controller design

- Optimize complex costs while preserving stability
- Perform well on average w.r.t noise realizations



SNOC:

$$K(\cdot) \in \operatorname{argmin}_{w_{0:T}} \mathbb{E} [L(K, w_{0:T})]$$

s.t. closed-loop stability

EPFL Stochastic Nonlinear Optimal Control (SNOC)


System

- Nonlinear, discrete time, stable/pre-stabilized, stochastic

Controller design

- Optimize complex costs while preserving **stability**
- Perform well on average w.r.t noise realizations

$$\left(\sum_{t=0}^{\infty} |w_t|^p \right)^{\frac{1}{p}} < \infty \iff \left(\sum_{t=0}^{\infty} |x_t|^p \right)^{\frac{1}{p}} < \infty$$



SNOC:

$$K(\cdot) \in \operatorname{argmin}_{w_{0:T}} \mathbb{E} [L(K, w_{0:T})]$$

s.t. closed-loop stability


EPFL Stochastic Nonlinear Optimal Control (SNOC)

System

- Nonlinear, discrete time, stable/pre-stabilized, stochastic

Controller design

- Optimize complex costs while preserving stability
- Perform well on average w.r.t noise realizations



Goal

- Leverage neural controllers
- Provide performance and stability guarantees



SNOC:

$$K(\cdot) \in \operatorname{argmin}_{w_{0:T}} \mathbb{E} [L(K, w_{0:T})]$$

s.t. closed-loop stability

SNOC:

$$K^\theta(\cdot) \in \operatorname{argmin}_{\theta} \mathbb{E}_{w_{0:T}} [L(\theta, w_{0:T})]$$

s.t. closed-loop stability




True cost



Perform well on average



Generally intractable




SNOC:

$$K^\theta(\cdot) \in \operatorname{argmin}_{\theta} \mathbb{E}_{w_{0:T}} [L(\theta, w_{0:T})]$$

s.t. closed-loop stability

True cost

- ✓ Perform well on average
- ✗ Generally intractable



Empirical SNOC:

$$K^\theta(\cdot) \in \operatorname{argmin}_{\theta} \frac{1}{|S|} \sum_{s=1}^{|S|} [L(\theta, w_{0:T}^s)]$$

s.t. closed-loop stability

Empirical cost

- ✓ Tractable
- ✗ Might overfit

Dataset of
sampled
process
noise


SNOC:

$$K^\theta(\cdot) \in \operatorname{argmin}_{\theta} \mathbb{E}_{w_{0:T}} [L(\theta, w_{0:T})]$$

s.t. closed-loop stability

True cost

- ✓ Perform well on average
- ✗ Generally intractable



Empirical SNOC:

$$K^\theta(\cdot) \in \operatorname{argmin}_{\theta} \frac{1}{|S|} \sum_{s=1}^{|S|} [L(\theta, w_{0:T}^s)]$$

s.t. closed-loop stability

Empirical cost

- ✓ Tractable
- ✗ Might overfit

Dataset of
sampled
process
noise


?

How to ensure generalization to unseen noise?

Similar to *overfitting* in supervised learning



PAC-Bayesian inference^[1]



?

How to ensure generalization to unseen noise?

SNOC:


$$K^\theta(\cdot) \in \operatorname{argmin}_{\theta} \mathbb{E}_{w_{0:T}} [L(\theta, w_{0:T})]$$

s.t. closed-loop stability



True cost

- ✓ Perform well on average
- ✗ Generally intractable



Empirical SNOC:

$$K^\theta(\cdot) \in \operatorname{argmin}_{\theta} \frac{1}{|\mathcal{S}|} \sum_{s=1}^{|\mathcal{S}|} [L(\theta, w_{0:T}^s)]$$

s.t. closed-loop stability



Empirical cost

- ✓ Tractable
- ✗ Might overfit

Dataset of
sampled
process
noise



How to ensure generalization to unseen noise?




PAC-Bayesian inference


Propose a SNOC framework allowing upper-bounding
the true cost given the empirical cost

EPFL Comparable Problem in ML

	Optimal control	Supervised learning
Data 	Process noise sequences	(feature, target) pairs
True cost	$\mathbb{E}_{w_{0:T}} [L(\theta, w_{0:T})]$	$\mathbb{E}_{(x,y)} l(h(x), y)$
Empirical cost	$\frac{1}{ \mathcal{S} } \sum_{s=1}^{ \mathcal{S} } [L(\theta, w_{0:T}^s)]$	$\frac{1}{ \mathcal{S} } \sum_{s=1}^{ \mathcal{S} } l(h(x^s), y^s)$
Goal 	Upper-bound the true cost given the empirical cost	
Tool 	  PAC-Bayes bounds ^[1]	

[1] P. Alquier, "User-friendly introduction to PAC-Bayes bounds," Foundations and Trends® in Machine Learning, 2024.





Prior distribution
over controllers


Inference

Cost function

Which controllers are
good?

Data = $\emptyset$:

- No info $\equiv$ uniform
- Existing info $\Rightarrow$ informed prior



Prior distribution
over controllers

Dataset of
sampled
process noise

Inference


Cost function

Which controllers are
good?

Data = $\emptyset$:

- No info $\equiv$ uniform
- Existing info $\Rightarrow$ informed prior

EPFL Probabilistic Inference over Controllers



Which controllers are good?

Data = $\emptyset$:

- No info $\equiv$ uniform
- Existing info $\Rightarrow$ informed prior



Data:

- Update belief about good controllers

Theorem: Consider a dataset $\mathbb{S}$, a prior $\mathcal{P}$, and a bounded cost $L \in [0, C]$. For every posterior $\mathcal{Q}$, $\lambda > 0$, confidence level $\delta \in (0, 1)$, and $\theta \sim \mathcal{Q}$, the inequality

$$\underbrace{\mathbb{E}_{w_{0:T}} [L(\theta, w_{0:T})]}_{\text{True cost}} \leq \underbrace{\frac{1}{|\mathbb{S}|} \sum_{s=1}^{|\mathbb{S}|} [L(\theta, w_{0:T}^s)]}_{\text{Empirical cost}} + \underbrace{\frac{1}{\lambda} \ln \frac{d\mathcal{Q}}{d\mathcal{P}}(\theta)}_{\text{Distance(prior, posterior)}} + \underbrace{\frac{1}{\lambda} \ln \frac{1}{\delta} + \frac{\lambda C^2}{8|\mathbb{S}|}}_{\text{Other terms}}$$

holds with probability $\geq 1 - \delta$ over sampling $\mathbb{S}$ and $\theta \sim \mathcal{Q}$.

Theorem: Consider a dataset $\mathbb{S}$, a prior $\mathcal{P}$, and a bounded cost $L \in [0, C]$. For every posterior $\mathcal{Q}$, $\lambda > 0$, confidence level $\delta \in (0, 1)$, and $\theta \sim \mathcal{Q}$, the inequality

$$\mathbb{E}_{w_{0:T}} [L(\theta, w_{0:T})] \leq \frac{1}{|\mathbb{S}|} \sum_{s=1}^{|\mathbb{S}|} [L(\theta, w_{0:T}^s)] + \frac{1}{\lambda} \ln \frac{d\mathcal{Q}}{d\mathcal{P}}(\theta) + \frac{1}{\lambda} \ln \frac{1}{\delta} + \frac{\lambda C^2}{8|\mathbb{S}|}$$

= 0 when $\mathcal{Q} = \mathcal{P}$

holds with probability $\geq 1 - \delta$ over sampling $\mathbb{S}$ and $\theta \sim \mathcal{Q}$.

Theorem: Consider a dataset $\mathbb{S}$, a prior $\mathcal{P}$, and a bounded cost $L \in [0, C]$. For every posterior $\mathcal{Q}$, $\lambda > 0$, confidence level $\delta \in (0, 1)$, and $\theta \sim \mathcal{Q}$, the inequality

$$\mathbb{E}_{w_{0:T}} [L(\theta, w_{0:T})] \leq \frac{1}{|\mathbb{S}|} \sum_{s=1}^{|\mathbb{S}|} [L(\theta, w_{0:T}^s)] + \frac{1}{\lambda} \ln \frac{d\mathcal{Q}}{d\mathcal{P}}(\theta) + \cancel{\frac{1}{\lambda} \ln \frac{1}{\delta}} + \cancel{\frac{\lambda C^2}{8|\mathbb{S}|}}$$

holds with probability $\geq 1 - \delta$ over sampling $\mathbb{S}$ and $\theta \sim \mathcal{Q}$.

$\rightarrow 0$ when $\begin{cases} |\mathbb{S}| \rightarrow \infty \\ \lambda \rightarrow \infty \text{ slower than } |\mathbb{S}| \end{cases}$

Theorem: Consider a dataset $\mathbb{S}$, a prior $\mathcal{P}$, and a bounded cost $L \in [0, C]$. For every posterior $\mathcal{Q}$, $\lambda > 0$, confidence level $\delta \in (0, 1)$, and $\theta \sim \mathcal{Q}$, the inequality

$$\underbrace{\mathbb{E}_{w_{0:T}} [L(\theta, w_{0:T})]}_{\text{True cost}} \leq \underbrace{\frac{1}{|\mathbb{S}|} \sum_{s=1}^{|\mathbb{S}|} [L(\theta, w_{0:T}^s)]}_{\text{Empirical cost}} + \frac{1}{\lambda} \ln \frac{d\mathcal{Q}}{d\mathcal{P}}(\theta) + \frac{1}{\lambda} \ln \frac{1}{\delta} + \frac{\lambda C^2}{8|\mathbb{S}|}$$

holds with probability $\geq 1 - \delta$ over sampling $\mathbb{S}$ and $\theta \sim \mathcal{Q}$.



Upper-bound the *true* cost given the *empirical* cost for a single controller drawn from the posterior

Theorem: Consider a dataset $\mathbb{S}$, a prior $\mathcal{P}$, and a bounded cost $L \in [0, C]$. For every posterior $\mathcal{Q}$, $\lambda > 0$, confidence level $\delta \in (0, 1)$, and $\theta \sim \mathcal{Q}$, the inequality

$$\mathbb{E}_{w_{0:T}} [L(\theta, w_{0:T})] \leq \frac{1}{|\mathbb{S}|} \sum_{s=1}^{|\mathbb{S}|} [L(\theta, w_{0:T}^s)] + \frac{1}{\lambda} \ln \frac{d\mathcal{Q}}{d\mathcal{P}}(\theta) + \frac{1}{\lambda} \ln \frac{1}{\delta} + \frac{\lambda C^2}{8|\mathbb{S}|}$$

holds with probability $\geq 1 - \delta$ over sampling $\mathbb{S}$ and $\theta \sim \mathcal{Q}$.

- Distribution-independent bound
- Holds for arbitrary cost, prior, posterior, confidence level, and parameter λ

Theorem: Consider a dataset $\mathbb{S}$, a prior $\mathcal{P}$, and a bounded cost $L \in [0, C]$. For every posterior $\mathcal{Q}$, $\lambda > 0$, confidence level $\delta \in (0, 1)$, and $\theta \sim \mathcal{Q}$, the inequality

$$\mathbb{E}_{w_{0:T}} [L(\theta, w_{0:T})] \leq \frac{1}{|\mathbb{S}|} \sum_{s=1}^{|\mathbb{S}|} [L(\theta, w_{0:T}^s)] + \frac{1}{\lambda} \ln \frac{d\mathcal{Q}}{d\mathcal{P}}(\theta) + \frac{1}{\lambda} \ln \frac{1}{\delta} + \frac{\lambda C^2}{8|\mathbb{S}|}$$

holds with probability $\geq 1 - \delta$ over sampling $\mathbb{S}$ and $\theta \sim \mathcal{Q}$.

- Holds for arbitrary cost, prior, posterior, confidence level, and parameter λ

Idea: compute the posterior minimizing the expected upper bound

$$\mathcal{Q}^* := \arg \min_{\mathcal{Q}} \mathbb{E}_{\theta \sim \mathcal{Q}} \left[\frac{1}{|\mathbb{S}|} \sum_{s=1}^{|\mathbb{S}|} [L(\theta, w_{0:T}^s)] + \frac{1}{\lambda} \ln \frac{d\mathcal{Q}}{d\mathcal{P}}(\theta) \right]$$

Idea: compute the posterior minimizing the expected upper bound

$$\mathcal{Q}^* := \arg \min_{\mathcal{Q}} \mathbb{E}_{\theta \sim \mathcal{Q}} \left[\frac{1}{|\mathcal{S}|} \sum_{s=1}^{|\mathcal{S}|} [L(\theta, w_{0:T}^s)] + \frac{1}{\lambda} \ln \frac{d\mathcal{Q}}{d\mathcal{P}}(\theta) \right]$$

Closed-form solution: Gibbs distribution

Controller parameters

$$\mathcal{Q}^*(\theta) \propto \mathcal{P}(\theta) e^{-\lambda \frac{1}{|\mathcal{S}|} \sum_{s=1}^{|\mathcal{S}|} [L(\theta, w_{0:T}^s)]}$$



↑
↓
↓

Prior
Empirical cost

Idea: compute the posterior minimizing the expected upper bound

$$Q^* := \arg \min_Q \mathbb{E}_{\theta \sim Q} \left[\frac{1}{|\mathcal{S}|} \sum_{s=1}^{|\mathcal{S}|} [L(\theta, w_{0:T}^s)] + \frac{1}{\lambda} \ln \frac{dQ}{dP}(\theta) \right]$$

Closed-form solution: Gibbs distribution



Closed-form solution: Gibbs distribution

$$Q^*(\theta) \propto \mathcal{P}(\theta) e^{-\lambda \frac{1}{|\mathbb{S}|} \sum_{s=1}^{|\mathbb{S}|} [L(\theta, w_{0:T}^s)]}$$

Theorem: Consider a dataset $\mathbb{S}$, a prior $\mathcal{P}$, and a bounded cost $L \in [0, C]$. For every posterior $\mathcal{Q}$, $\lambda > 0$, confidence level $\delta \in (0, 1)$, and $\theta \sim \mathcal{Q}$, the inequality

$$\mathbb{E}_{w_{0:T}} [L(\theta, w_{0:T})] \leq \frac{1}{|\mathbb{S}|} \sum_{s=1}^{|\mathbb{S}|} [L(\theta, w_{0:T}^s)] + \frac{1}{\lambda} \ln \frac{d\mathcal{Q}}{d\mathcal{P}}(\theta) + \frac{1}{\lambda} \ln \frac{1}{\delta} + \frac{\lambda C^2}{8|\mathbb{S}|}$$

holds with probability $\geq 1 - \delta$ over sampling $\mathbb{S}$ and $\theta \sim \mathcal{Q}$.

Closed-form solution: Gibbs distribution

$$Q^*(\theta) \propto \mathcal{P}(\theta) e^{-\lambda \frac{1}{|\mathbb{S}|} \sum_{s=1}^{|\mathbb{S}|} [L(\theta, w_{0:T}^s)]}$$

Corollary: Consider a dataset $\mathbb{S}$, a prior $\mathcal{P}$, and a bounded cost $L \in [0, C]$, and optimal posterior Q^* , $\lambda > 0$, confidence level $\delta \in (0, 1)$, and $\theta \sim Q^*$, the inequality


$$\mathbb{E}_{w_{0:T}} [L(\theta, w_{0:T})] \leq \frac{-1}{\lambda} \ln \left(\mathbb{E}_{\theta \sim \mathcal{P}} e^{\frac{-\lambda}{|\mathbb{S}|} \sum_{s=1}^{|\mathbb{S}|} L(\theta, w_{0:T}^s)} \right) + \frac{1}{\lambda} \ln \frac{dQ}{d\mathcal{P}}(\theta) + \frac{1}{\lambda} \ln \frac{1}{\delta} + \frac{\lambda C^2}{8|\mathbb{S}|}$$

holds with probability $\geq 1 - \delta$ over sampling $\mathbb{S}$ and $\theta \sim Q^*$.

✓ Tighter than the generic bound

✓ Computationally efficient variant in the paper

EPFL Summary: PAC-Bayesian Inference of Controllers




PAC-Bayes vs. empirical cost minimization:

- ✓ Incorporate prior knowledge
- ✓ Mitigate overfitting
- ✓ Allows deriving generalization bounds

Closed-form solution: Gibbs distribution



Sampling from the Gibbs distribution is intractable




[3] Q. Liu and D. Wang, "Stein variational gradient descent: A general purpose Bayesian inference algorithm," NeurIPS, 2016.

[4] D. Rezende and S. Mohamed, "Variational inference with normalizing flows," ICML, 2015.

Approximating the Optimal Posterior

Closed-form solution: Gibbs distribution  intractable to sample

Approximate inference
E.g., SVGD ^[3], Normalizing Flows ^[4], ...




SVG

Normalizing Flows


[3] Q. Liu and D. Wang, "Stein variational gradient descent: A general purpose Bayesian inference algorithm," NeurIPS, 2016.

[4] D. Rezende and S. Mohamed, "Variational inference with normalizing flows," ICML, 2015.


EPFL Closed-Loop Stability




EPFL Closed-Loop Stability



EPFL Closed-Loop Stability




How to define distributions **only over stabilizing controllers**?



[1] D. Youla, H. Jabr, and J. Bongiorno, "Modern Wiener-Hopf design of optimal controllers—part ii: the multivariable case," IEEE Transactions on Automatic Control, 1976.
 [2] L. Furieri, C. L. Galimberti, and G. Ferrari-Trecate, "Learning to Boost the Performance of Stable Nonlinear Systems," IEEE Open Journal of Control Systems, 2024.


EPFL Parametrization of All Stabilizing Controllers

Performance Boosting
Framework^[2]



EPFL Parametrization of All Stabilizing Controllers

Performance Boosting
Framework^[2]




Theorem: Assume that the open-loop system is ℓ_p -stable,

($\Rightarrow$) If $\mathcal{M} \in \mathcal{L}_p$, then the closed-loop system is ℓ_p -stable

($\Leftarrow$) If there is a causal controller $\mathbf{K}^*$ such that the closed-loop system is ℓ_p -stable, then $\exists \mathcal{M} \in \mathcal{L}_p$ such that $\mathbf{K} = \mathbf{K}^*$

All stable behaviour can be learned by parametrizing the $\mathcal{L}_p$ -stable operator $\mathcal{M}(\cdot)$

Models of stable operators $\forall \theta$?



All stable behaviour can be learned by parametrizing the $\mathcal{L}_p$ -stable operator $\mathcal{M}(\cdot)$

Models of stable operators $\forall \theta$?

Recurrent Equilibrium Networks:
Flexible Dynamic Models with Guaranteed Stability
and Robustness

Max Revay, Ruigang Wang, Ian R. Manchester

2024 IEEE 63rd Conference on Decision and Control (CDC)
December 16-19, 2024, MiCo, Milan, Italy

Neural Distributed Controllers with Port-Hamiltonian Structures

Muhammad Zakwan and Giancarlo Ferrari-Trecate

$\mathcal{L}_2$

Direct Parameterization of Lipschitz-Bounded Deep Networks

Ruigang Wang¹ Ian R. Manchester¹

Efficiently Modeling Long Sequences with Structured State Spaces

Albert Gu, Karan Goel, and Christopher Ré

Resurrecting Recurrent Neural Networks for Long Sequences


Antonio Orvieto^{1*} Samuel L. Smith² Albert Gu² Anushan Fernando² Caglar Gulcehre² Razvan Pascanu²
Soham De²

EPFL Summary: Parametrizing Stabilizing Controllers




Parameterize $\mathcal{M}(\cdot)$ as a NN
with parameters θ
CL-stability $\forall \theta \in \mathbb{R}^d$

Learning the controller
 $\equiv$
Learning $\mathcal{M}(\cdot)$




Performance Boosting Framework^[2]




Free $\mathcal{L}_p$ operator

IMC architecture parametrizing *all and only* stabilizing controllers in terms of a freely chosen operator $\mathcal{M}(\cdot)$

EPFL Closed-Loop Stability



How to define distributions **only over stabilizing controllers**?




[1] D. Youla, H. Jabr, and J. Bongiorno, "Modern Wiener-Hopf design of optimal controllers—part ii: the multivariable case," IEEE Transactions on Automatic Control, 1976.

[2] L. Furieri, C. L. Galimberti, and G. Ferrari-Trecate, "Learning to Boost the Performance of Stable Nonlinear Systems," IEEE Open Journal of Control Systems, 2024.


Experiments



EPFL Scalar LTI System – Setup



- Linear system, linear controller
- Stability constraint: range on the gain
- Compare against:
 - **Empirical controller**
 - Minimize empirical cost
 - **Benchmark controller** (ideal)
 - Exploits knowledge of noise mean and a very large dataset



- Prior over the **gain**

Gaussian(LQR, variance)


Optimal solution for
noise-free system

Account for
uncertainty

- Prior over the **bias**


- Case 1: minimal information  Uniform
- Case 2: knowledge of noise mean  Gaussian

EPFL Scalar LTI System – Results



- Both posteriors are concentrated near **benchmark** : Good performance for low and high uncertainty
- Empirical** overfits : Outperform **empirical**
- Accommodate prior knowledge


EPFL Scalar LTI System – Validating the Upper Bound






Each point is the true cost of a controller sampled from the optimal posterior

1. All bounds are valid
2. Bounds are tighter for larger datasets
3. Bounds are tighter for more informative priors

- 2 robots: point-mass dynamics, nonlinear drag
- Noise on initial conditions
- **Goal:** CL stability on targets, avoid collisions & obstacles
 - Highly nonlinear cost
- Pre-stabilized system: robots go to targets, but might collide
- Performance-Boosting neural controller with ~900 parameters






Controller	Cost		Collisions	
	Train	Test	Train	Test
Empirical	21.79	23.14	0.0%	5.8%
Ours	21.86	22.29	0.0%	5.0%


- Can learn complex neural controllers
- Better generalization to OOS noise realizations

- **Conclusions:**
 - PAC-Bayesian framework for SNOC problems




- Conclusions:

- PAC-Bayesian framework for SNOC problems






- Control design by minimizing the expected upper bound to avoid overfitting



Conclusions:

- PAC-Bayesian framework for SNOC problems
- Control design by minimizing the expected upper bound to avoid overfitting
- Ensure closed-loop stability






Conclusions:

- PAC-Bayesian framework for SNOC problems
 - Control design by minimizing the expected upper bound to avoid overfitting
- Ensure closed-loop stability

Future work:

- PAC-Bayesian generative policies
- Transfer learning for designing informed priors



Thank you!

Mahrokh Boroujeni

