Bridging Offline Design and Online Adaptation in Safe Learning-Enabled Systems



Nikolai Matni
PI, Penn

George Pappas
Co-PI, Penn

Benjamin Recht
Co-PI, UCB

Penn Engineering
UNIVERSITY *of* PENNSYLVANIA

NSF

Berkeley ENGINEERING

Are meaningful safety guarantees even possible?

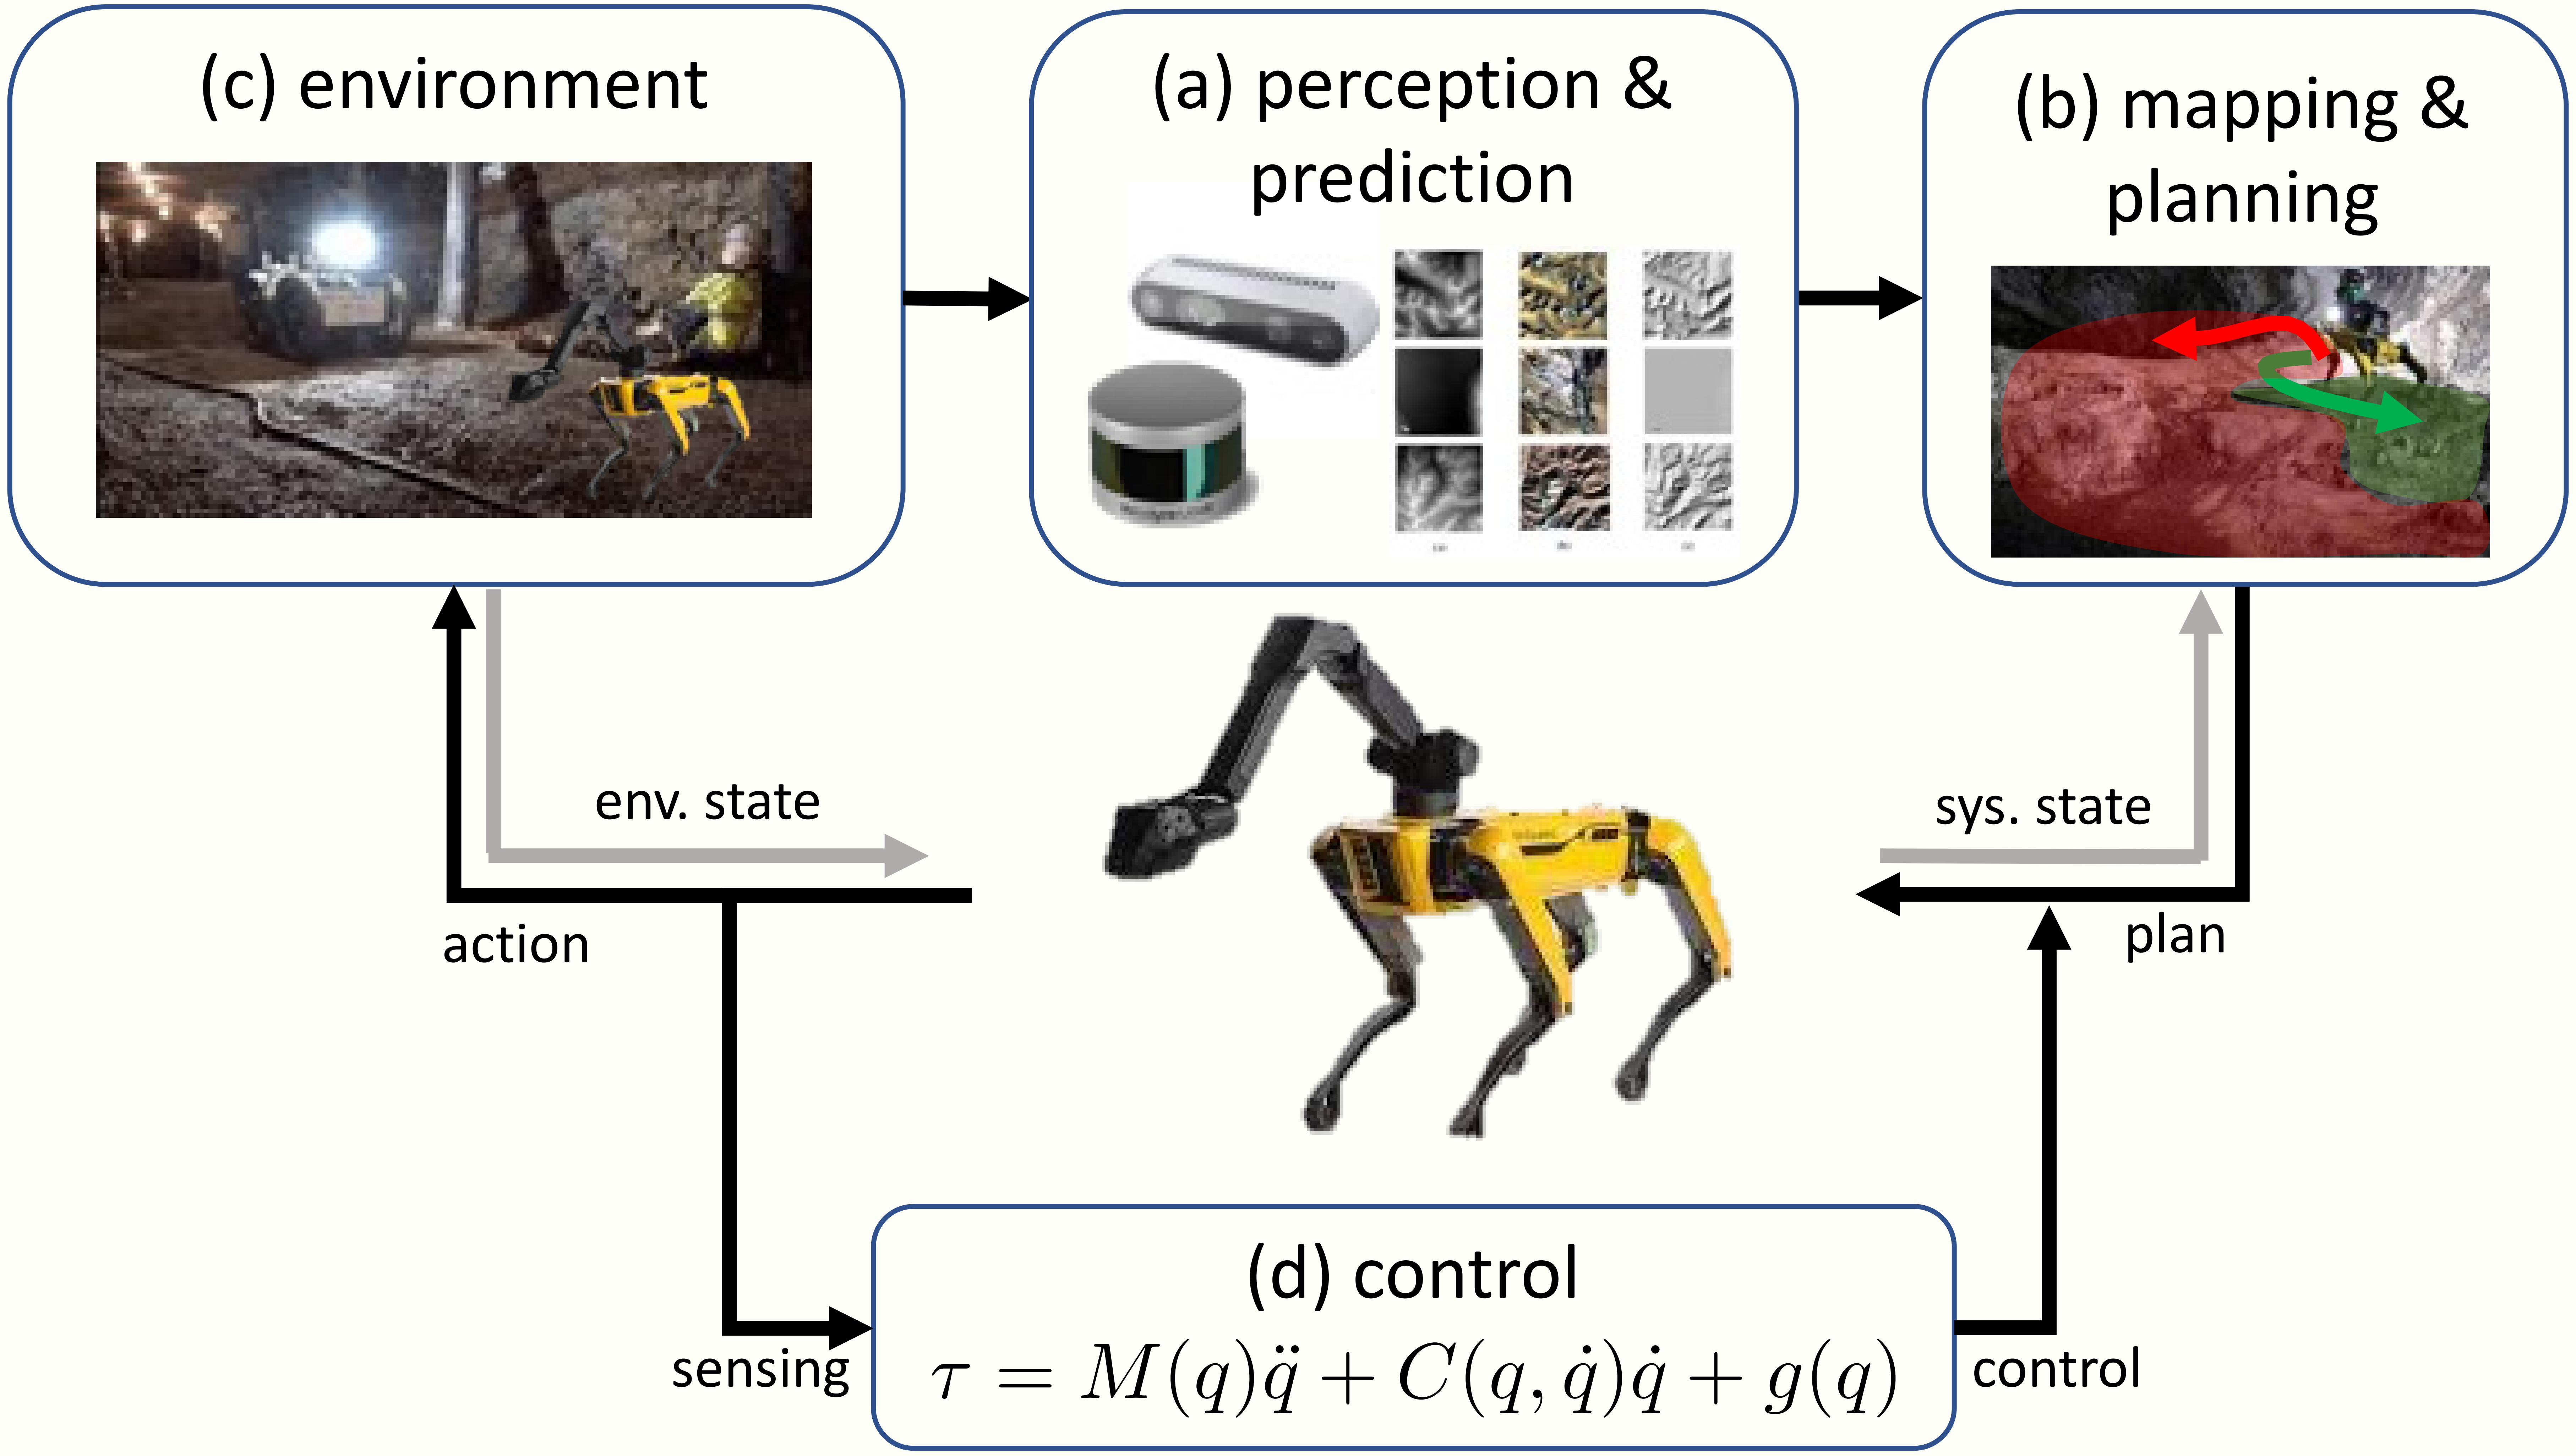

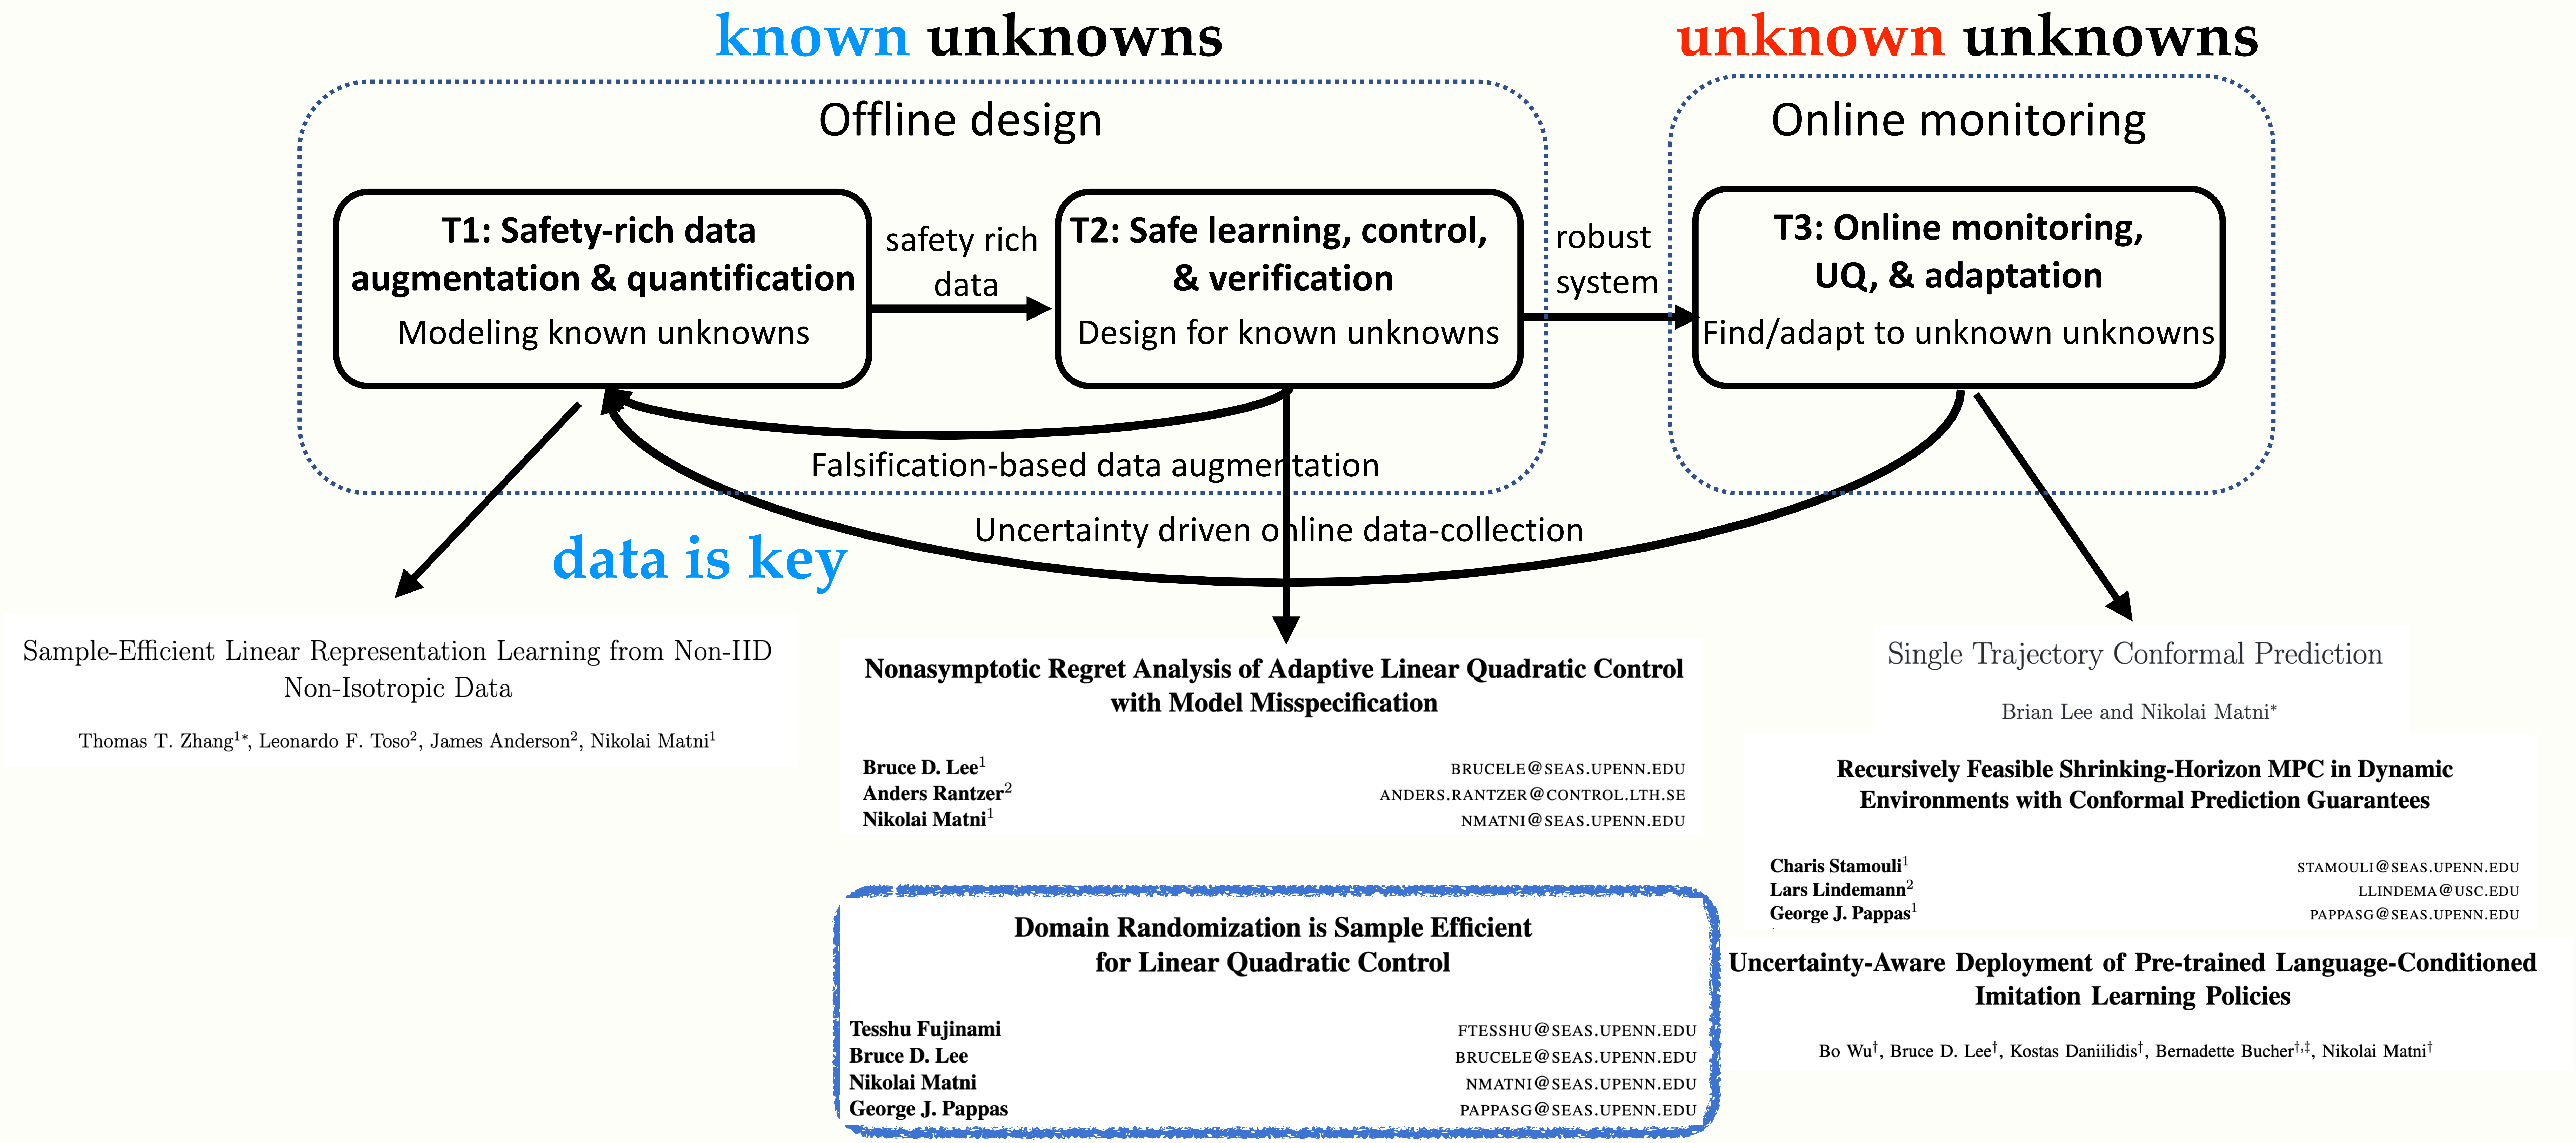
known unknowns
unknown unknowns
Offline design
Online monitoring
T1: Safety-rich data augmentation & quantification
Modeling known unknowns
safety rich data
T2: Safe learning, control, & verification
Design for known unknowns
robust system
T3: Online monitoring, UQ, & adaptation
Find/adapt to unknown unknowns
Falsification-based data augmentation
Uncertainty driven online data-collection
data is key
Sample-Efficient Linear Representation Learning from Non-IID Non-Isotropic Data
Thomas T. Zhang[1*], Leonardo F. Toso[2], James Anderson[2], Nikolai Matni[1]
Nonasymptotic Regret Analysis of Adaptive Linear Quadratic Control with Model Misspecification
Bruce D. Lee[1] BRUCELE@SEAS.UPENN.EDU
Anders Rantzer[2] ANDERS.RANTZER@CONTROL.LTH.SE
Nikolai Matni[1] NMATNI@SEAS.UPENN.EDU
Domain Randomization is Sample Efficient for Linear Quadratic Control
Tesshu Fujinami FTESSHU@SEAS.UPENN.EDU
Bruce D. Lee BRUCELE@SEAS.UPENN.EDU
Nikolai Matni NMATNI@SEAS.UPENN.EDU
George J. Pappas PAPPASG@SEAS.UPENN.EDU
Single Trajectory Conformal Prediction
Brian Lee and Nikolai Matni*
Recursively Feasible Shrinking-Horizon MPC in Dynamic Environments with Conformal Prediction Guarantees
Charis Stamouli[1] STAMOULI@SEAS.UPENN.EDU
Lars Lindemann[2] LLINDEMA@USC.EDU
George J. Pappas[1] PAPPASG@SEAS.UPENN.EDU
Uncertainty-Aware Deployment of Pre-trained Language-Conditioned Imitation Learning Policies
Bo Wu[†], Bruce D. Lee[†], Kostas Daniilidis[†], Bernadette Bucher[†,‡], Nikolai Matni[†]

Domain Randomization is Sample Efficient in Linear Quadratic Control


Tesshu Fujinami


Bruce Lee


Nikolai Matni


George Pappas

Penn Engineering
UNIVERSITY of PENNSYLVANIA

NSF

Domain Randomization: Key Step in Sim2Real Recipe

Rapid Motor Adaptation for Legged Robots

Ashish Kumar, UC Berkeley
Zipeng Fu, CMU
Deepak Pathak, CMU
Jitendra Malik, UC Berkeley/FAIR

Robotics: Science and Systems 2021

Reinforcement Learning for Versatile, Dynamic, and Robust Bipedal Locomotion Control

Zhongyu Li, Xue Bin Peng, Pieter Abbeel, Sergey Levine,
Glen Berseth, Koushil Sreenath

Berkeley UNIVERSITY OF CALIFORNIA · SFU SIMON FRASER UNIVERSITY · Université de Montréal · Mila

Deep Drone Racing: From Simulation to Reality with Domain Randomization

Antonio Loquercio, Elia Kaufmann, René Ranftl,
Alexey Dosovitskiy, Vladlen Koltun, Davide Scaramuzza

University of Zurich UZH, Department of Neuroinformatics · ETH zürich · University of Zurich UZH, Department of Informatics · intel

Problem Formulation

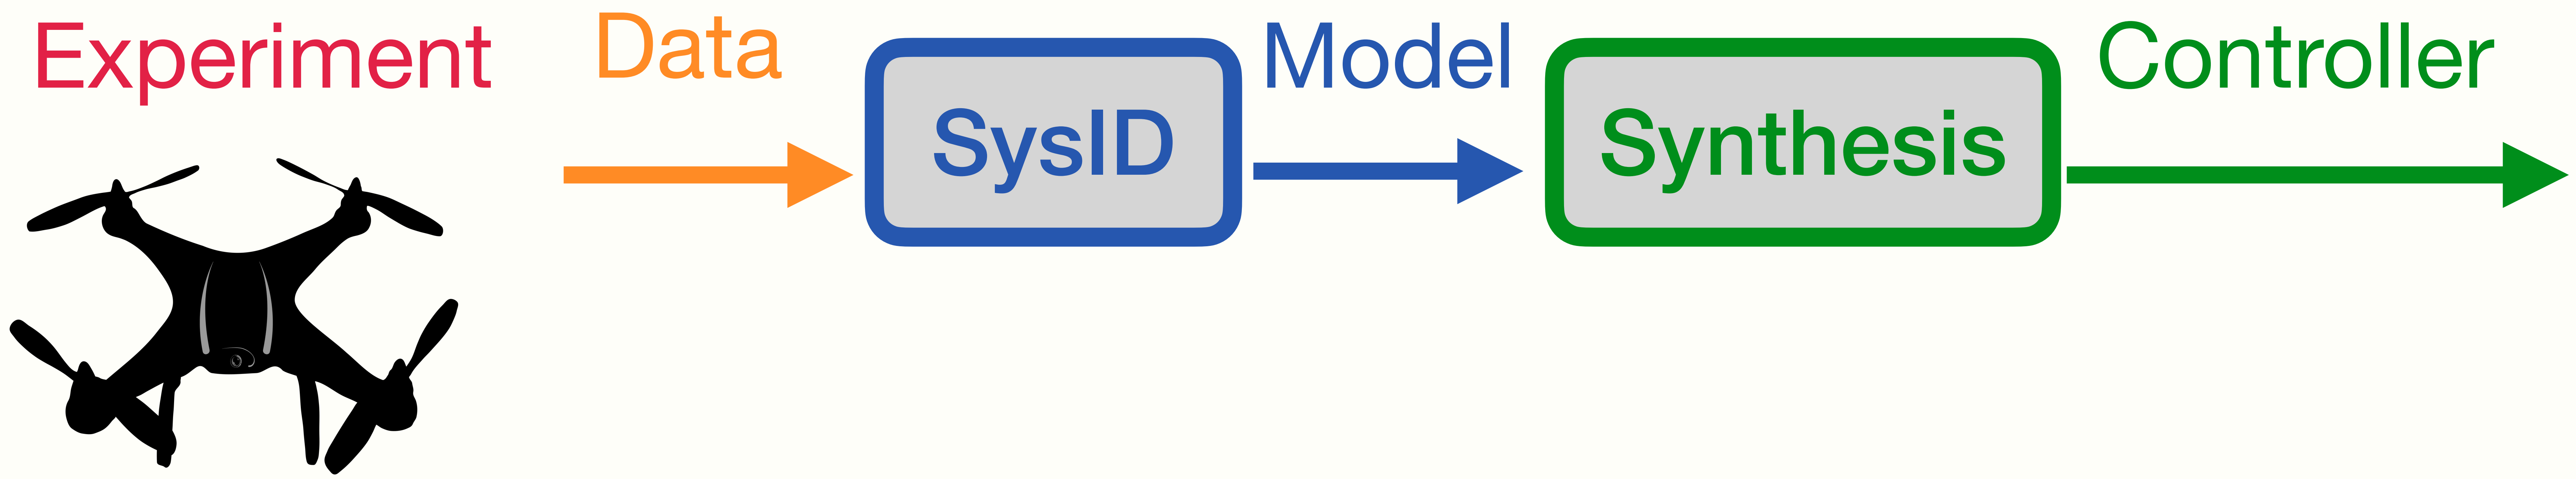

System dynamics: $X_{t+1} = f(X_t, U_t, \theta^\star) + W_t \quad \theta^\star$ unknown, $\quad W_t \sim \mathcal{N}(0, \sigma_w^2 I)$

Want policy $U_t = \pi_t(X_t)$ from class $\Pi^\star$ to minimize control objective:

$$J(\pi, \theta) = \mathbf{E}_\theta^\pi \left[\sum_{t=1}^{T} c_t(X_t, U_t) + c_{T+1}(X_{T+1}) \right]$$

1. Estimate $\hat{\theta}$ via least squares over data $\hat{\theta} = \arg\min_\theta \sum_{(X,U,X^+)\in D} \| X^+ - f(X, U; \theta) \|^2$
2. Synthesize a controller $\pi_\star(X_t, \hat{\theta})$

Linear System

$$X_{t+1} = \begin{bmatrix} 1.01 & 0.01 & 0 \\ 0.01 & 1.01 & 0.01 \\ 0 & 0.01 & 1.01 \end{bmatrix} X_t + U_t + W_t$$

Quadratic Cost

$$c_t(X_t, U_t) = \|X_t\|^2 + \|U_t\|^2$$

Certainty Equivalent Control

$$\pi_{\mathrm{CE}}(\hat{\theta}) \triangleq \underset{\pi \in \Pi_\star}{\operatorname{argmin}} J(\pi, \hat{\theta})$$

high-performing *when stabilizing* ✔

unreliable in low-data regime ✘

Robust Control

$$\pi_{\mathrm{robust}}(\hat{\theta}) \triangleq \underset{\pi \in \Pi_\star}{\operatorname{argmin}} \sup_{\theta \in G} J(\pi, \theta)$$

overly conservative performance ✘

reliable in low-data regime ✔

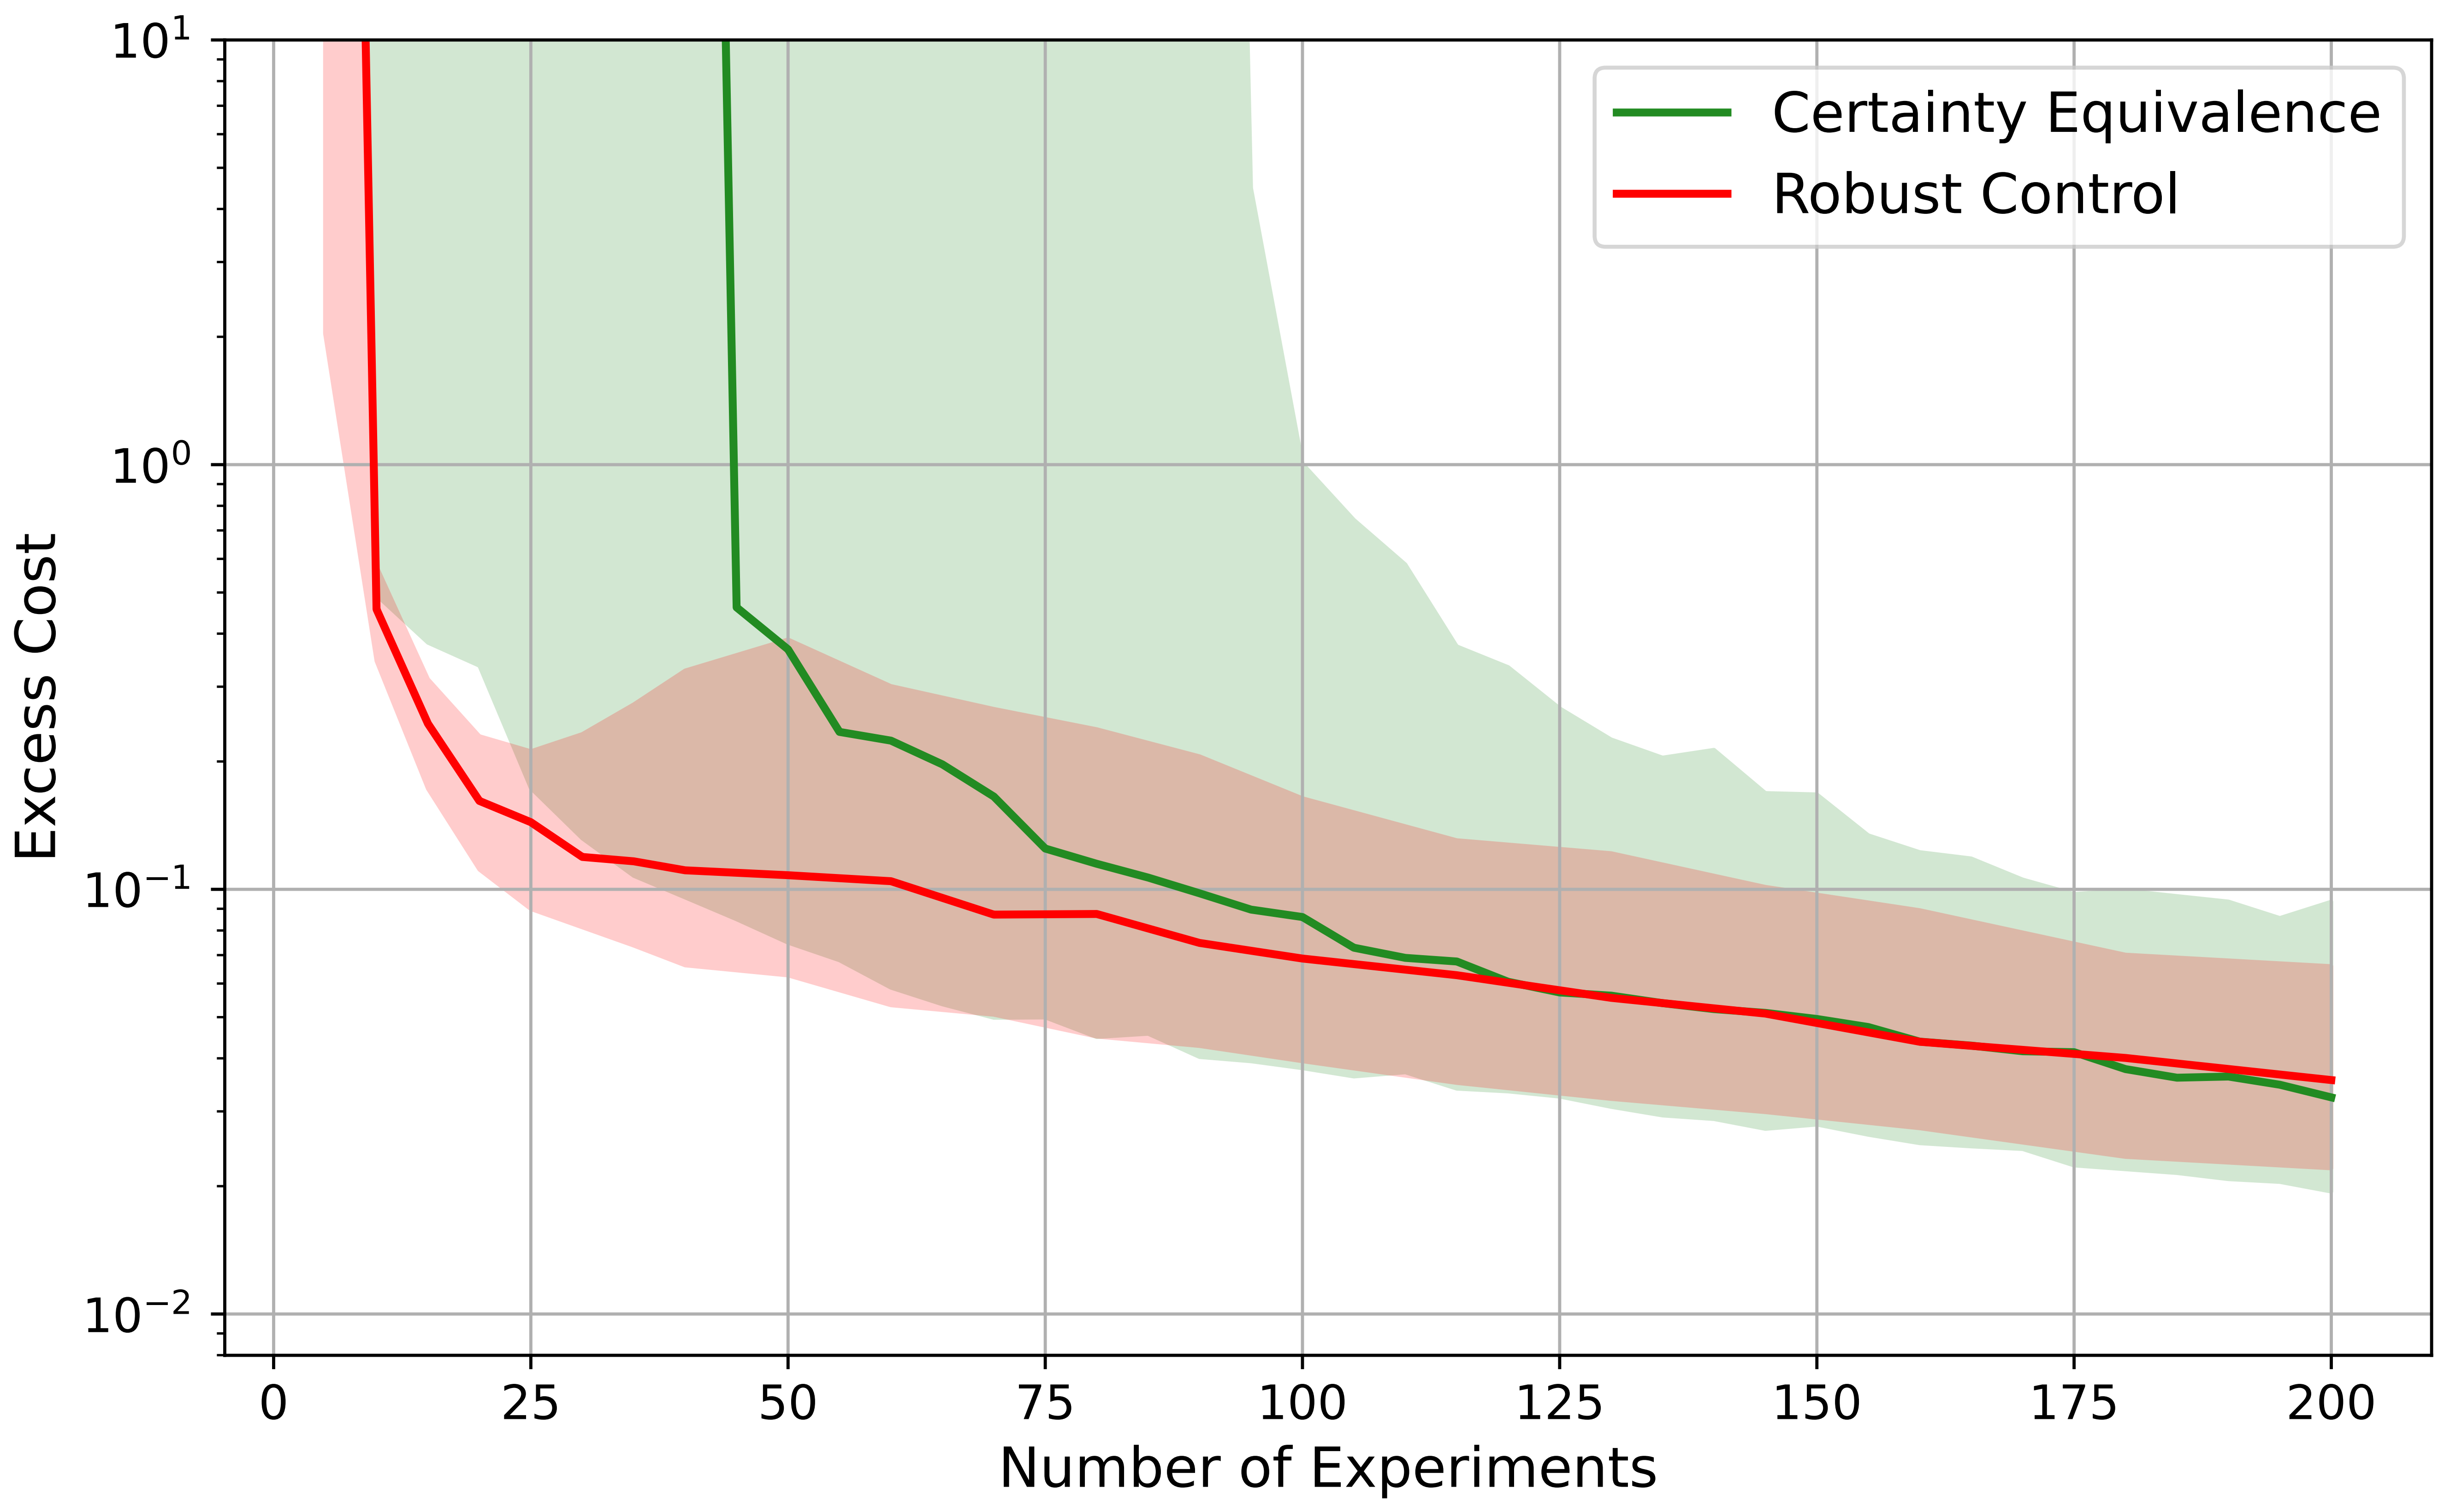

Domain Randomization

Input: Estimate $\hat{\theta}$

↳ **Set** $G = \{\theta : (\theta - \hat{\theta})^\top (N\mathsf{FI}(\hat{\theta}))^{-1} (\theta - \hat{\theta}) \leq \mathsf{conf}\}$

↳ **Define** $\mathscr{D}$ as a uniform dist. over G

↳ **Return** $\pi_{\mathsf{DR}}(\hat{\theta}) \triangleq \operatorname{argmin}_{\pi \in \Pi_\star} \mathbf{E}_{\Theta \sim \mathscr{D}} \left[J(\pi, \Theta) \right]$

reliable in low-data regime ✓

high performing controller ✓

Certainty Equivalent Control

$$\pi_{\mathsf{CE}}(\hat{\theta}) \triangleq \operatorname{argmin}_{\pi \in \Pi_\star} J(\pi, \hat{\theta})$$

high-performing *when stabilizing* ✓

unreliable in low-data regime ✗

Robust Control

$$\pi_{\mathsf{robust}}(\hat{\theta}) \triangleq \operatorname{argmin}_{\pi \in \Pi_\star} \sup_{\theta \in G} J(\pi, \theta)$$

overly conservative performance ✗

reliable in low-data regime ✓

Domain Randomization

$$\pi_{\mathrm{DR}}(\hat{\theta}) \triangleq \underset{\pi \in \Pi_\star}{\operatorname{argmin}} \, \mathbf{E}_{\boldsymbol{\Theta} \sim \mathscr{D}} \left[J(\pi, \boldsymbol{\Theta}) \right]$$

reliable in low-data regime ✓

high performing controller ✓

Certainty Equivalent Control

$$\pi_{\mathrm{CE}}(\hat{\theta}) \triangleq \underset{\pi \in \Pi_\star}{\operatorname{argmin}} J(\pi, \hat{\theta})$$

high-performing *when stabilizing* ✓

unreliable in low-data regime ✗

Robust Control

$$\pi_{\mathrm{robust}}(\hat{\theta}) \triangleq \underset{\pi \in \Pi_\star}{\operatorname{argmin}} \sup_{\theta \in G} J(\pi, \theta)$$

overly conservative performance ✗

reliable in low-data regime ✓

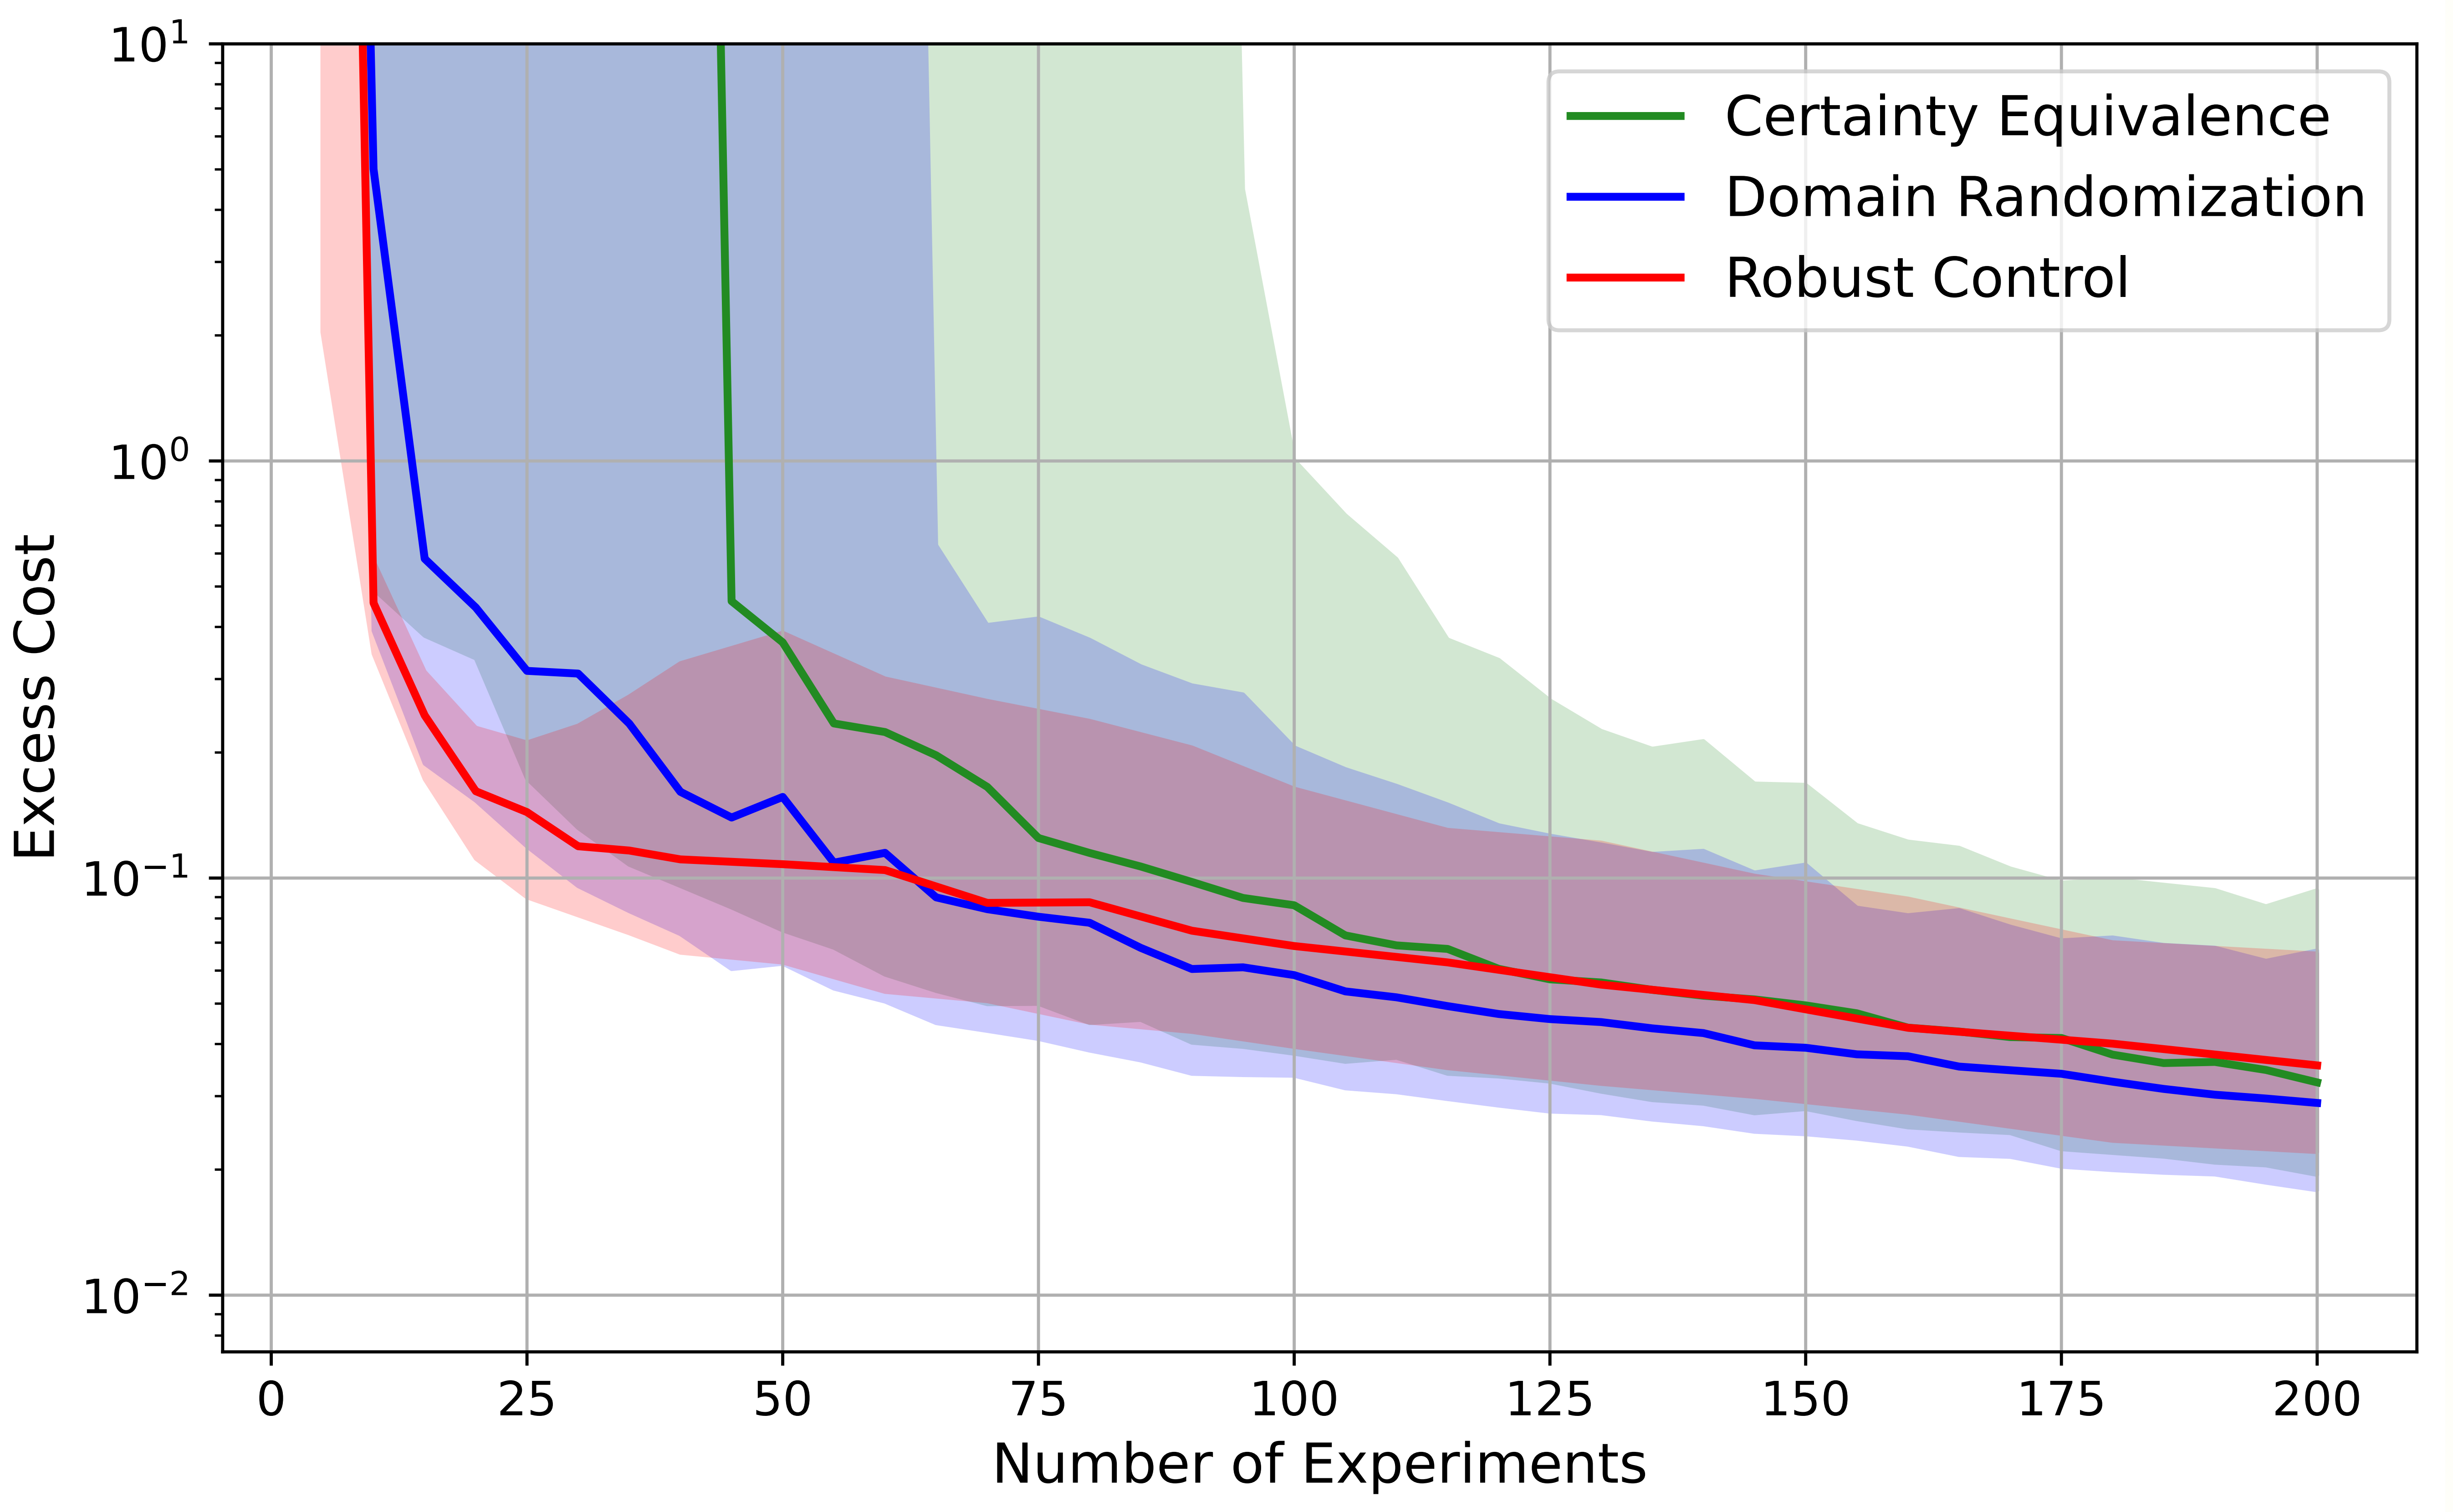

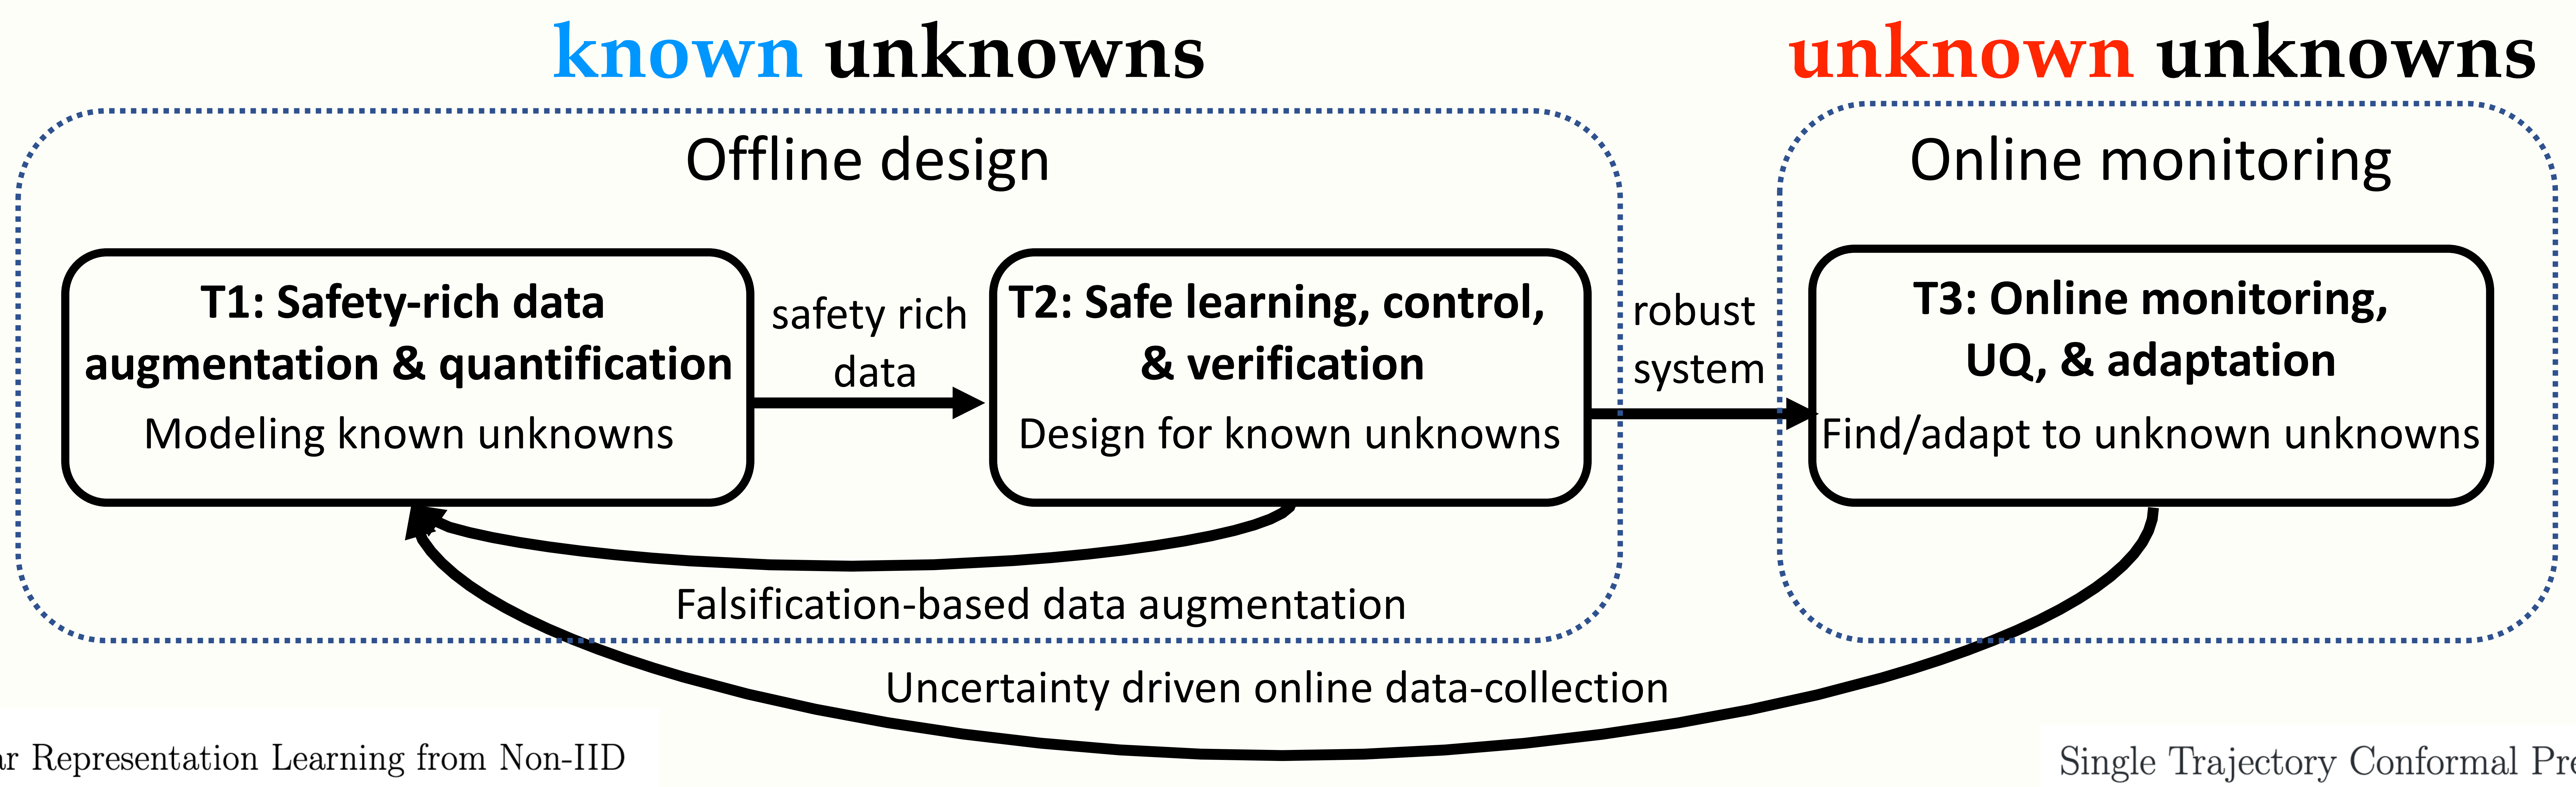

Sample-Efficient Linear Representation Learning from Non-IID Non-Isotropic Data

Thomas T. Zhang[1*], Leonardo F. Toso[2], James Anderson[2], Nikolai Matni[1]

Nonasymptotic Regret Analysis of Adaptive Linear Quadratic Control with Model Misspecification

Bruce D. Lee[1] BRUCELE@SEAS.UPENN.EDU
Anders Rantzer[2] ANDERS.RANTZER@CONTROL.LTH.SE
Nikolai Matni[1] NMATNI@SEAS.UPENN.EDU

Domain Randomization is Sample Efficient for Linear Quadratic Control

Tesshu Fujinami FTESSHU@SEAS.UPENN.EDU
Bruce D. Lee BRUCELE@SEAS.UPENN.EDU
Nikolai Matni NMATNI@SEAS.UPENN.EDU
George J. Pappas PAPPASG@SEAS.UPENN.EDU

Single Trajectory Conformal Prediction

Brian Lee and Nikolai Matni*

Recursively Feasible Shrinking-Horizon MPC in Dynamic Environments with Conformal Prediction Guarantees

Charis Stamouli[1] STAMOULI@SEAS.UPENN.EDU
Lars Lindemann[2] LLINDEMA@USC.EDU
George J. Pappas[1] PAPPASG@SEAS.UPENN.EDU

Uncertainty-Aware Deployment of Pre-trained Language-Conditioned Imitation Learning Policies

Bo Wu[†], Bruce D. Lee[†], Kostas Daniilidis[†], Bernadette Bucher[†,‡], Nikolai Matni[†]

Relevant Papers

- Fujinami, T., Lee, B.D., Matni, N. and Pappas, G.J., 2025. *Domain Randomization is Sample Efficient for Linear Quadratic Control*. arXiv preprint arXiv:2502.12310.
- Stamouli, C., Lindemann, L. and Pappas, G., 2024, June. *Recursively feasible shrinking-horizon mpc in dynamic environments with conformal prediction guarantees*. In 6th Annual Learning for Dynamics & Control Conference (pp. 1330-1342). PMLR.
- Wu, B., Lee, B.D., Daniilidis, K., Bucher, B. and Matni, N., 2024, October. *Uncertainty-aware deployment of pre-trained language-conditioned imitation learning policies*. In 2024 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS) (pp. 878-883). IEEE.
- Zhang, T.T., Toso, L.F., Anderson, J. and Matni, N., *Sample-Efficient Linear Representation Learning from Non-IID Non-Isotropic Data*. In The Twelfth International Conference on Learning Representations.
- Lee, B. and Matni, N., 2024. *Single trajectory conformal prediction*. arXiv preprint arXiv:2406.01570.

Penn Engineering
UNIVERSITY of PENNSYLVANIA

NSF

Berkeley ENGINEERING