Learning-Based Planning & Control with Persistent Safety for UAS via $\mathcal{L}_1$ Adaptive Control and Machine Learning

Naira Hovakimyan and Aditya Gahlawat
Mechanical Science and Engineering
University of Illinois at Urbana-Champaign

Annual AFOSR Review
Grant# FA9550-21-1-0411

September 12, 2023

Complex Dynamics and Environments

Challenges

- Unpredictable environments
- Obstacle-rich and dynamic
- Nonlinear uncertain dynamics

Solutions

- Fast (re-)planning
- Safe planning
- Safe learning
- Guaranteed robustness

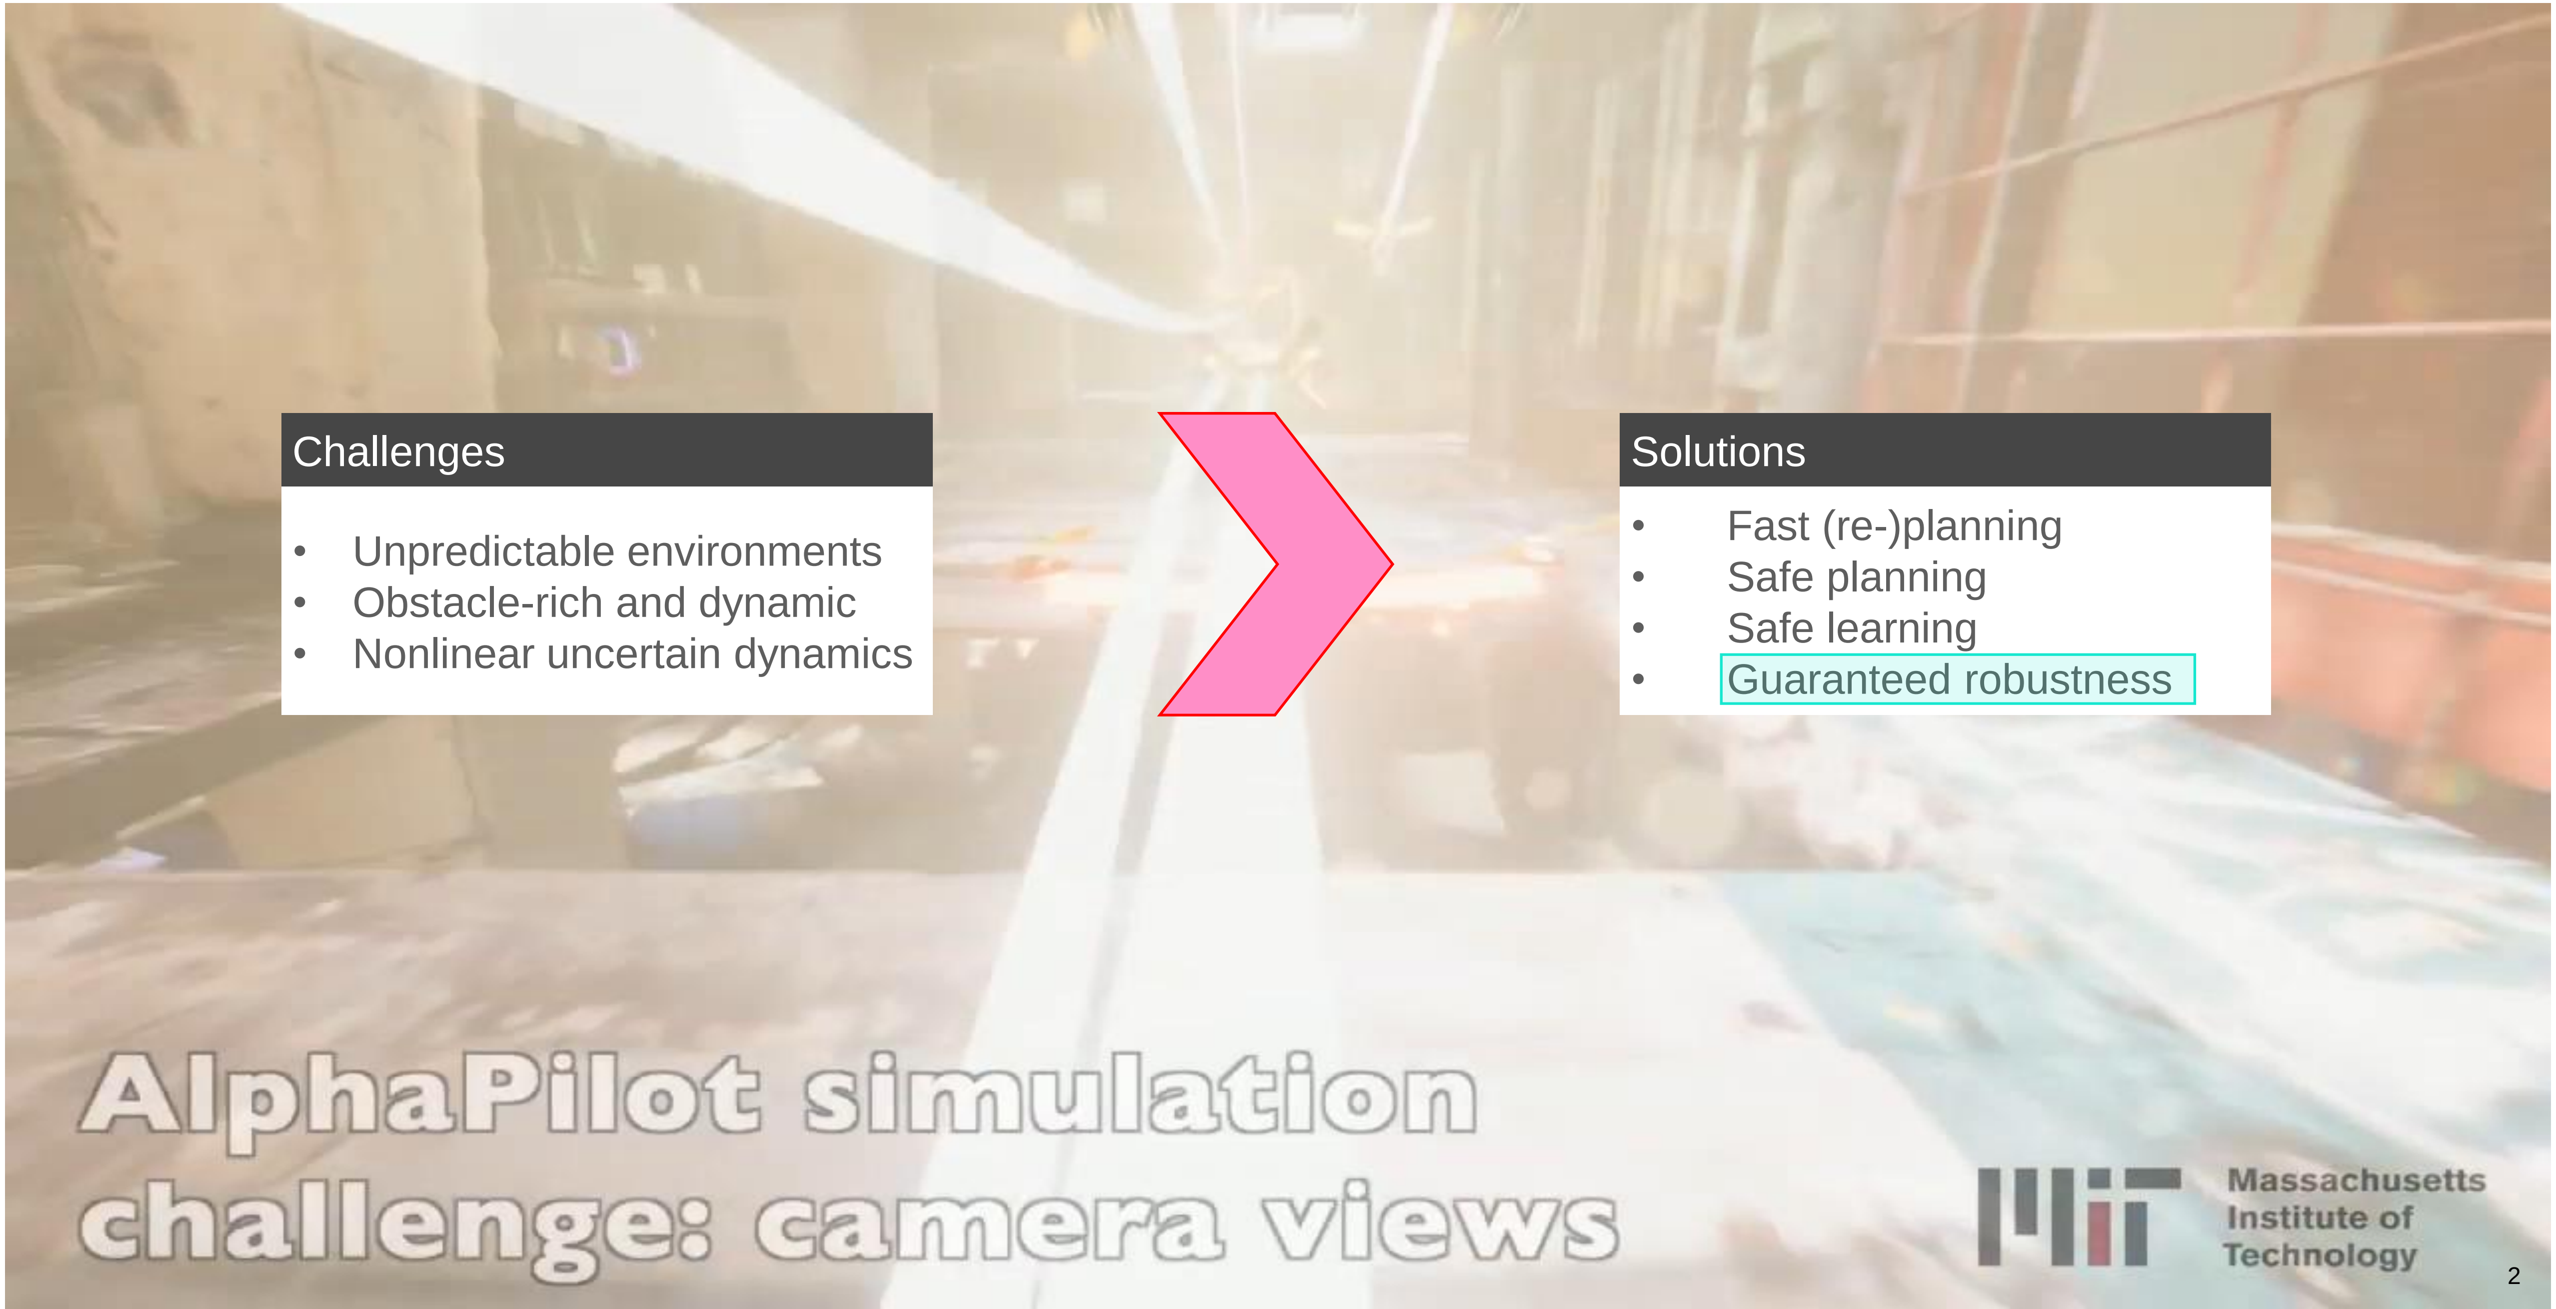

Massachusetts Institute of Technology

Challenges and the Tools

Complex Dynamics | Uncertain Models | Uncertain Environments

Control theoretic tools

- Structured models
- Parametric uncertainties
- Deterministic representations

Data-driven ML tools

- General models
- Unstructured uncertainties
- Stochastic representations

Safety & Robustness

Empirical Performance

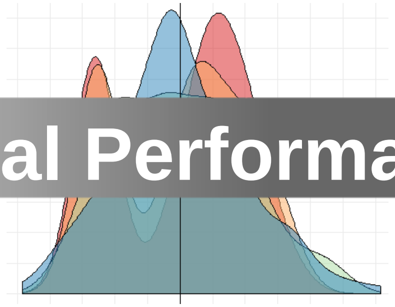

Bridging the divide

Data-driven
Control Pipeline

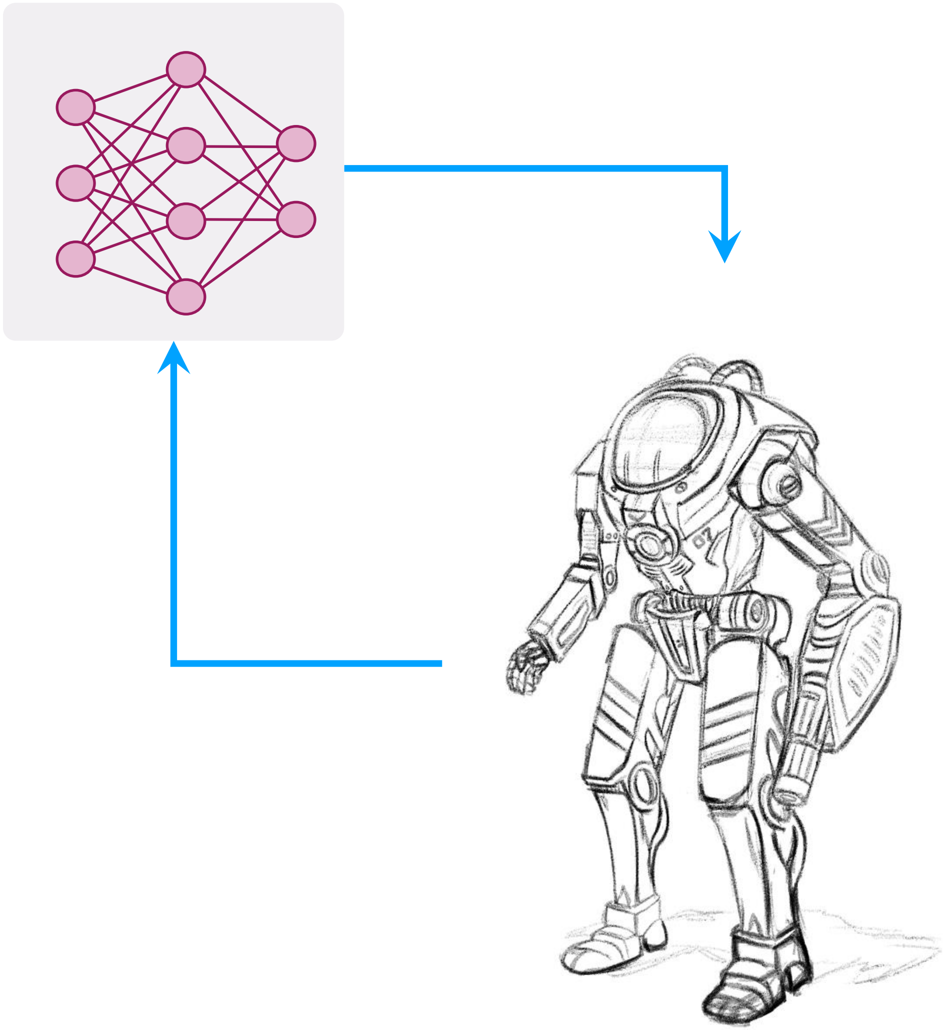

Data-driven Control Design

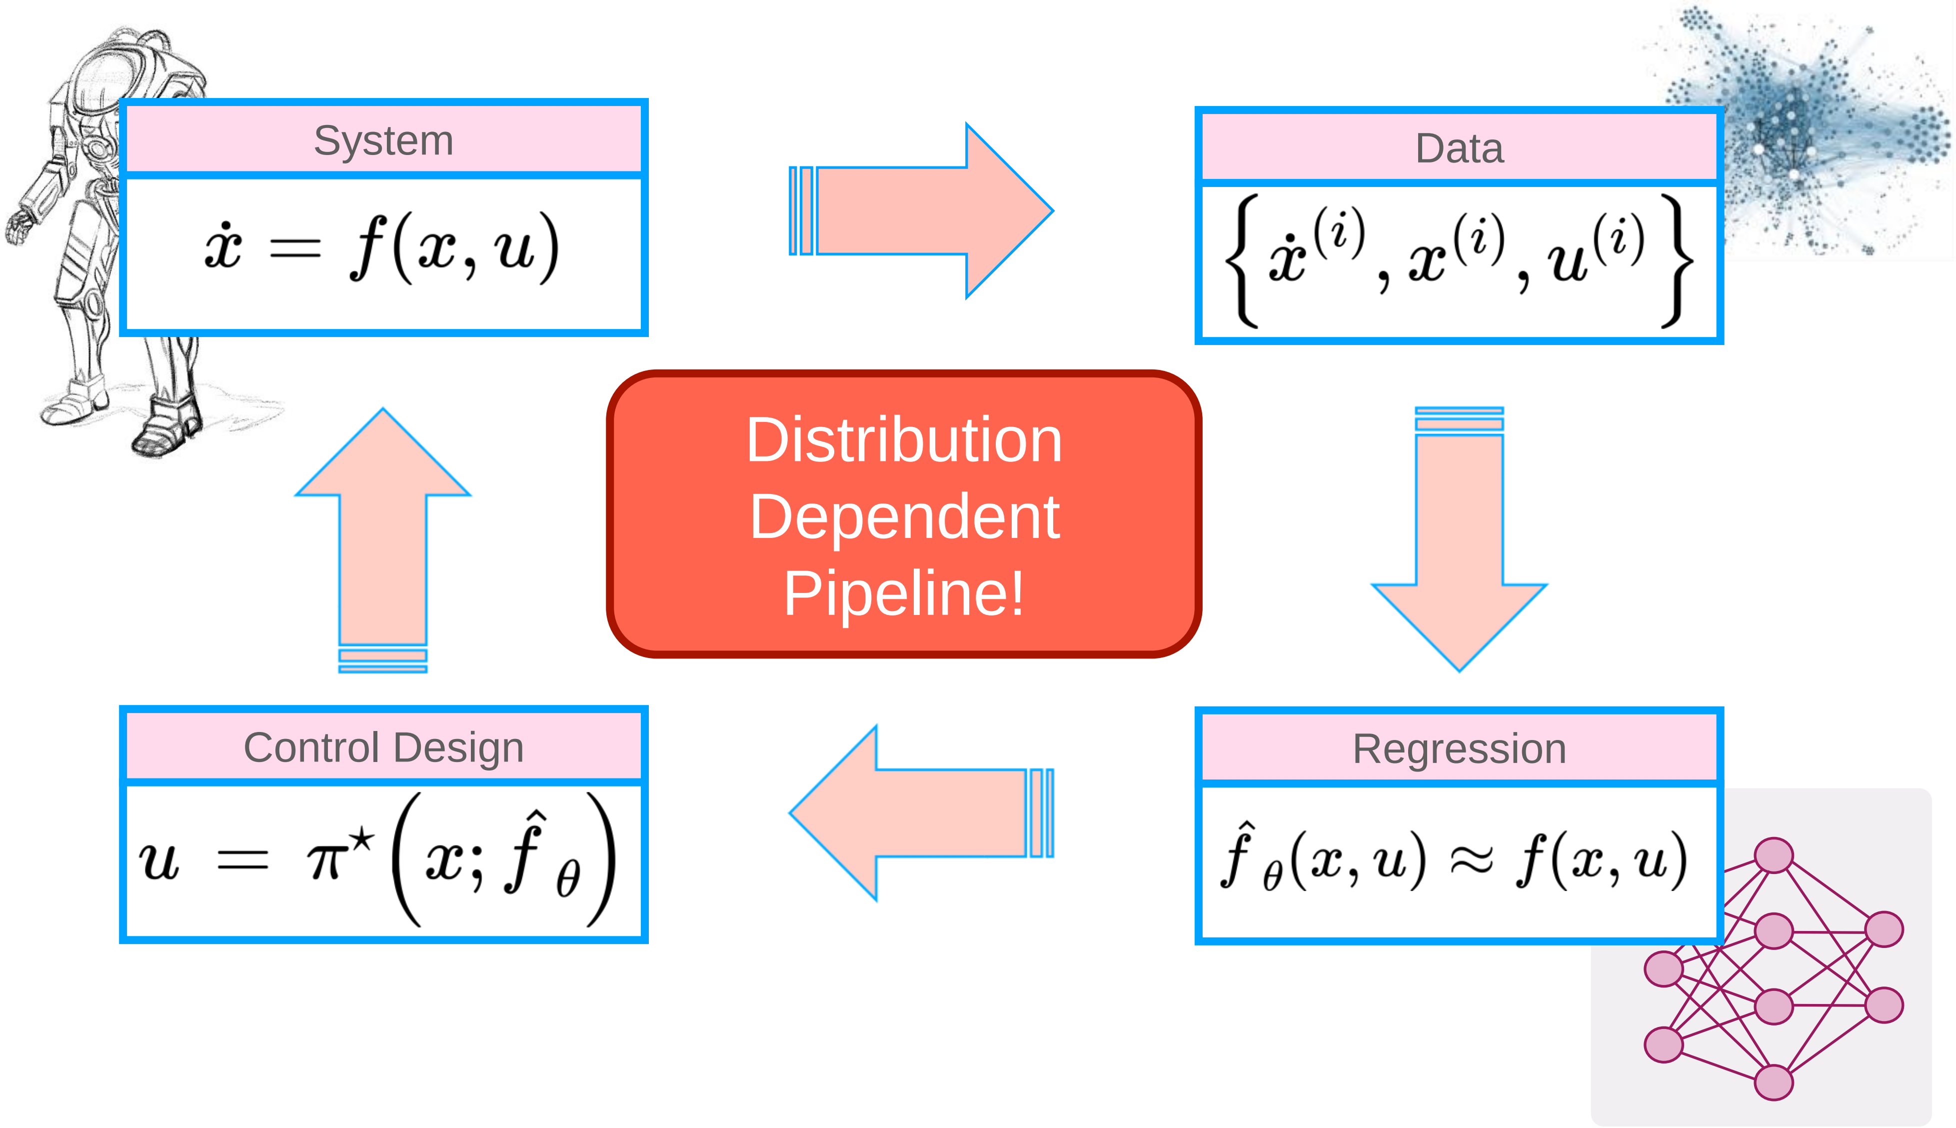

Data Distribution Dependence

Data
$\left\{\boxed{\dot{x}^{(i)}, x^{(i)}}, \boxed{u^{(i)}}\right\}$

$\leftarrow$

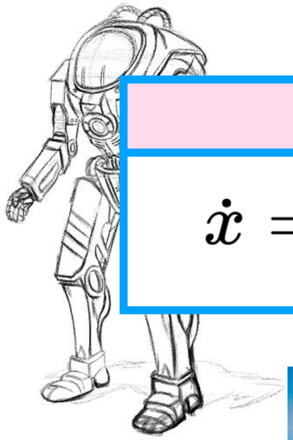

System
$\dot{x} = f(x, u)$

$\cup$

Data Logging Policy
$u = \pi^d(x)$

Human or Compute heavy (MPC)
Imitation Learning

$$\left\{\dot{x}^{(i)}, x^{(i)}, u^{(i)}\right\} \sim \mathcal{D}_{train}$$

- Training data distribution: True dynamics & data logging policy
- Training data has an associated distribution

The Role of Training Distribution

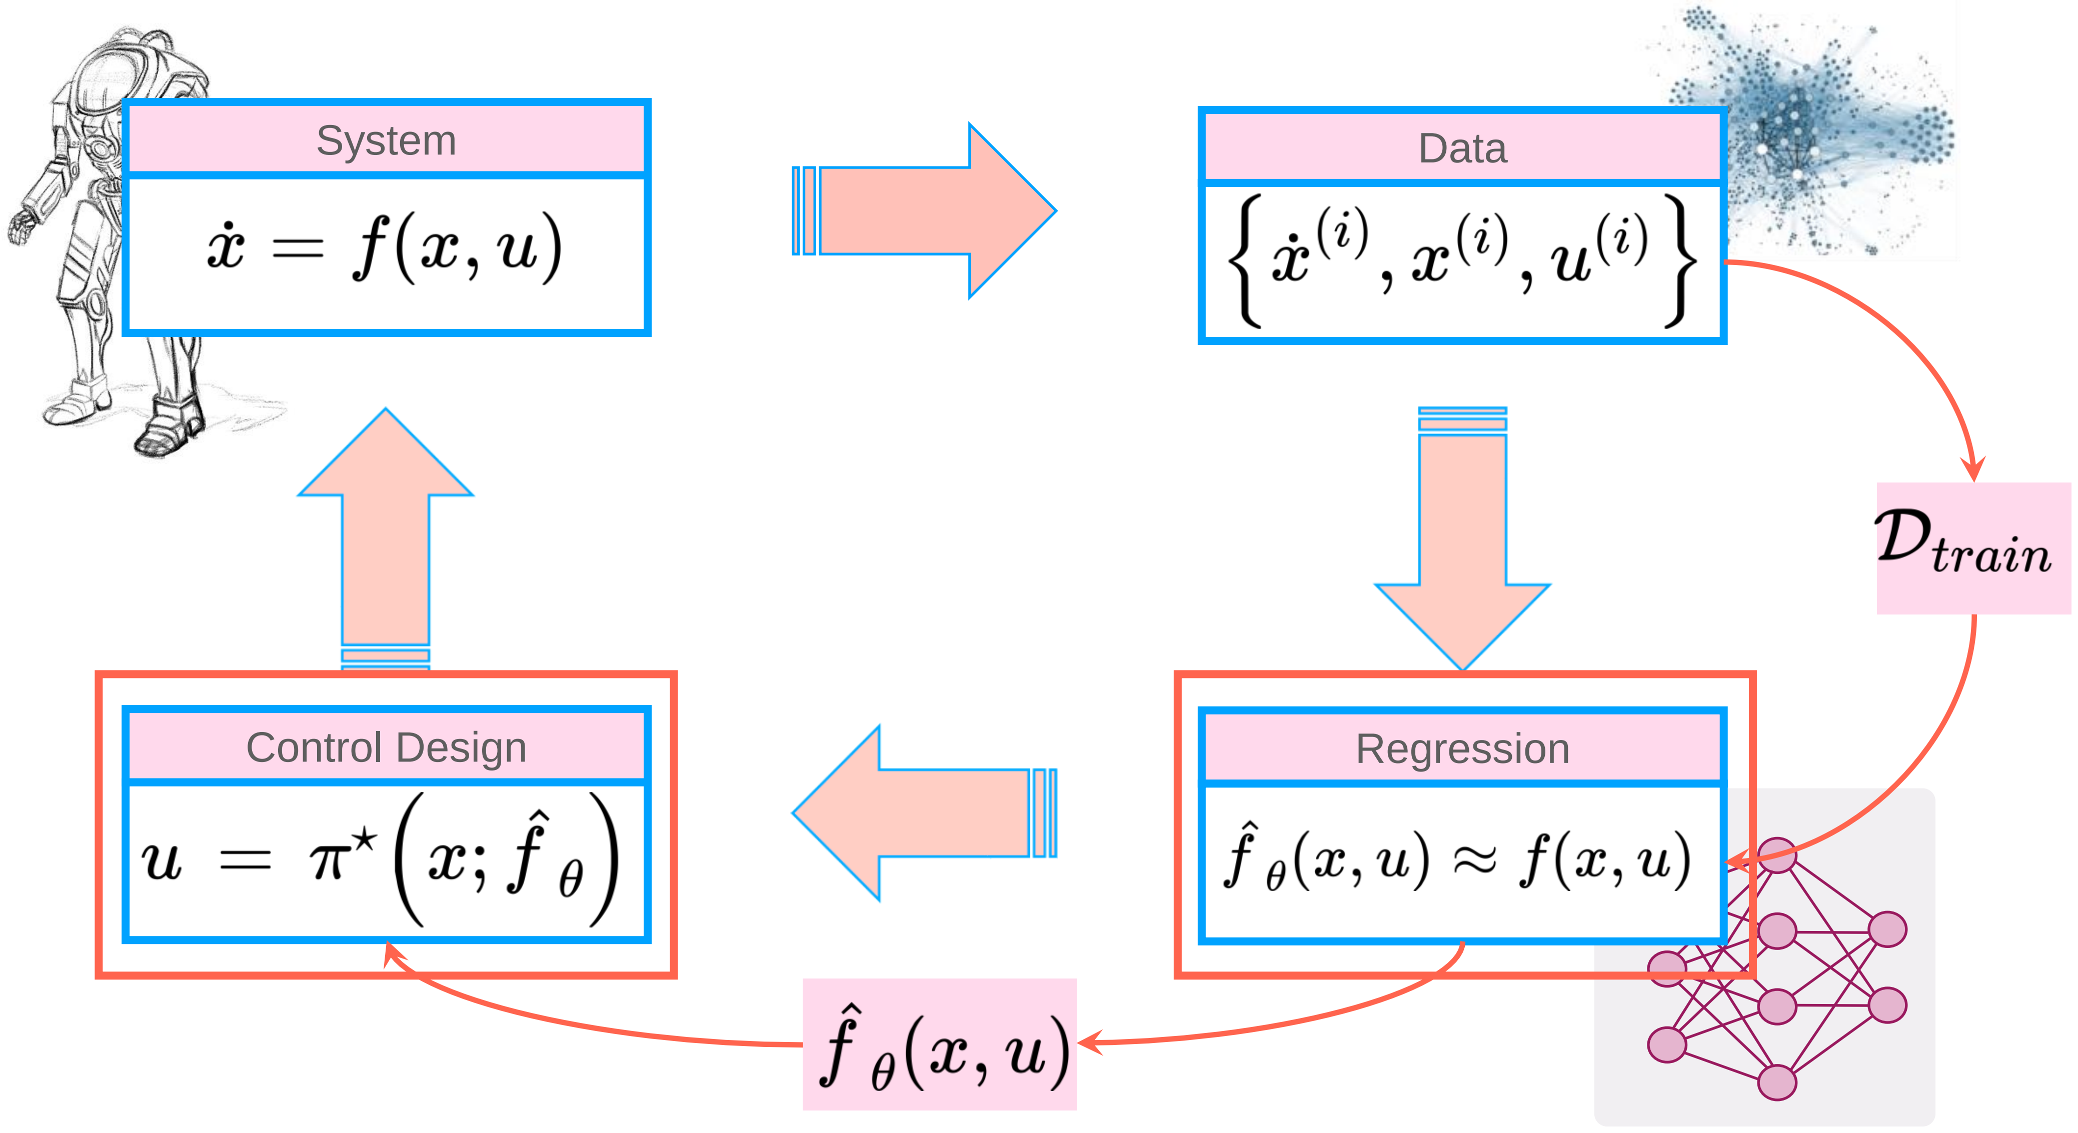

The Role of Training Distribution

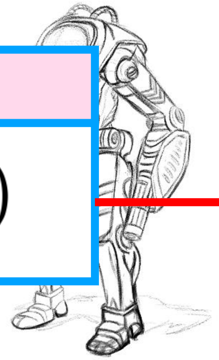

System

$$\dot{x} = f(x,u)$$

Data

$$\left\{\dot{x}^{(i)}, x^{(i)}, u^{(i)}\right\}$$

Regression

$$\hat{f}_\theta(x,u) \approx f(x,u)$$

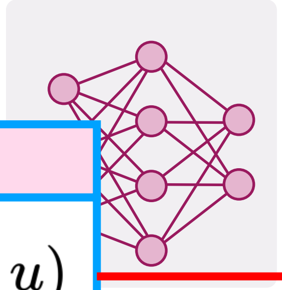

Control Design

$$u = \pi^\star\left(x; \hat{f}_\theta\right)$$

$\mathcal{D}_{true} \neq \mathcal{D}_{train}$

Distribution Shift!

Epistemic & Aleatoric

$\mathcal{D}_{train}$

$$\min_\theta\left(\mathbb{E}_{\mathcal{D}_{train}}\left\|\hat{f}_\theta(x,u) - f(x,u)\right\|\right)$$

$$\mathbb{P}_{\cancel{\mathcal{D}_{train}}}\left\{\text{Control objective is met}\right\} \geq 1 \;\; \delta$$

$x \sim \mathcal{D}_{true}$

Robustness & Distribution Shifts

Epistemic & Aleatoric ❯ $\hat{f}_\theta \neq f, \quad (x,u) \sim \mathcal{D}_{true}$ ❯ $\mathcal{D}_{true} \neq \mathcal{D}_{train}$

- We want to be robust to epistemic and aleatoric uncertainties such that
  - We can quantify the distribution shift

$$\mathcal{D}_{true} \in \Delta(\mathcal{D}_{train})$$

  - True distribution always lies within a known ball $\Delta(\mathcal{D}_{train})$

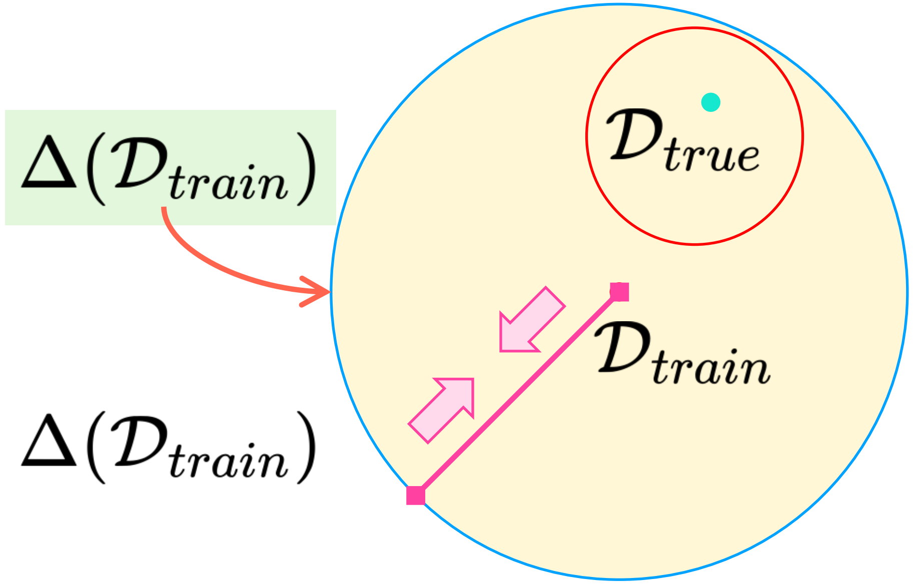

Agnostic to $\mathcal{D}_{train}$

  - We can mitigate the distribution shift
    - We can control the size of the guaranteed set

$$\min_{\Gamma \in \Delta(\mathcal{D}_{train})} (\mathbb{P}_\Gamma\{\text{Control objective is met}\}) \geq 1 - \delta$$

Robustness Certificates

Why certificates $\Delta(\mathcal{D}_{train})$ in the space of distributions?

- Upstream nominal controllers designed with certificates of distributional rob.
  - Available data with associated distribution

$$\mathbb{P}_{\mathcal{D}}\{\text{Control objective is met}\} \geq 1 - \delta$$

- Availability of data with true distribution is difficult to justify
  - Expensive and unsafe $\mathcal{D}_{true}$ ❯ $\mathcal{D}_{train}$
  - Only training data is available: from past operation, sim etc.
- Instead, if we can produce certificates of distributional robustness

$$\mathcal{D}_{true} \in \Delta(\mathcal{D}_{train})$$

  - Robust nominal control → Distributionally robust control, learning, and optimization $\min_{\Gamma \in \Delta(\mathcal{D}_{train})}(\mathbb{P}_{\Gamma}\{\text{Control objective is met}\}) \geq 1 - \delta$

Distributional Robustness

- Safe use of machine learning
  - Safe predictive control

- Natural ability to consider epistemic and aleatoric uncertainties
  - Systems and our understanding of them are stochastic

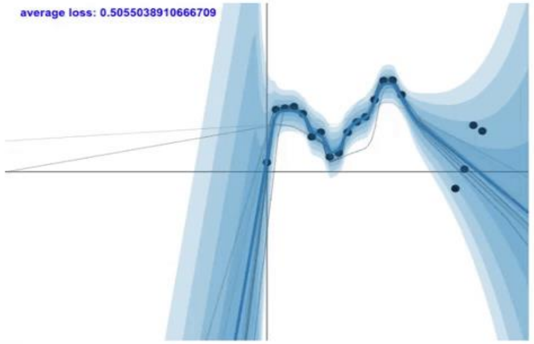

- Design principles guided by distributional guarantees independently verified by e.g. Monte-Carlo methods
  - Design space = Test space
  - Easier feedback between the spaces

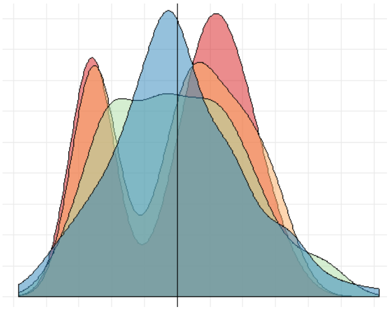

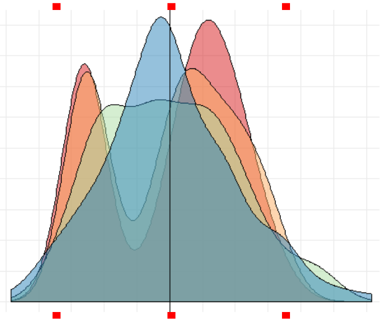

DR Control

- Distributionally robust control
  - Samples from the true distribution, only finite samples required [1,2]
  - Known distribution (measure), state constraint satisfaction [3]
- (Un)stable deterministic systems → robustness to stochastic disturbances
  - Stochastic disturbances enter unstable deterministic systems through uniformly bounded input operators, incremental stability extended to stochastic systems [4,5]
  - Stable systems with uniformly bounded input operators for stochastic disturbances, control of higher moments [6], robustness in the sense of the Wasserstein metric [7]

[1] Yang, Insoon. "Wasserstein distributionally robust stochastic control: A data-driven approach." *IEEE Transactions on Automatic Control* 66.8 (2020): 3863-3870.
[2] Hakobyan, Astghik, and Insoon Yang. "Distributionally Robust Differential Dynamic Programming with Wasserstein Distance." *IEEE Control Systems Letters* (2023).
[3] Van Parys, Bart PG, et al. "Distributionally robust control of constrained stochastic systems." *IEEE Transactions on Automatic Control* 61.2 (2015): 430-442.
[4] Tsukamoto, Hiroyasu, and Soon-Jo Chung. "Robust controller design for stochastic nonlinear systems via convex optimization." *IEEE Transactions on Automatic Control* 66.10 (2020): 4731-4746.
[5] Pham, Quang-Cuong, Nicolas Tabareau, and Jean-Jacques Slotine. "A contraction theory approach to stochastic incremental stability." *IEEE Transactions on Automatic Control* 54.4 (2009): 816-820.
[6] Mazumdar, Eric, et al. "High confidence sets for trajectories of stochastic time-varying nonlinear systems." *2020 59th IEEE Conference on Decision and Control (CDC)*. IEEE, 2020.
[7] Bouvrie, Jake, and Jean-Jacques Slotine. "Wasserstein contraction of stochastic nonlinear systems." *arXiv preprint arXiv:1902.08567* (2019).

Distributionally Robust Adaptive Control (DRAC)

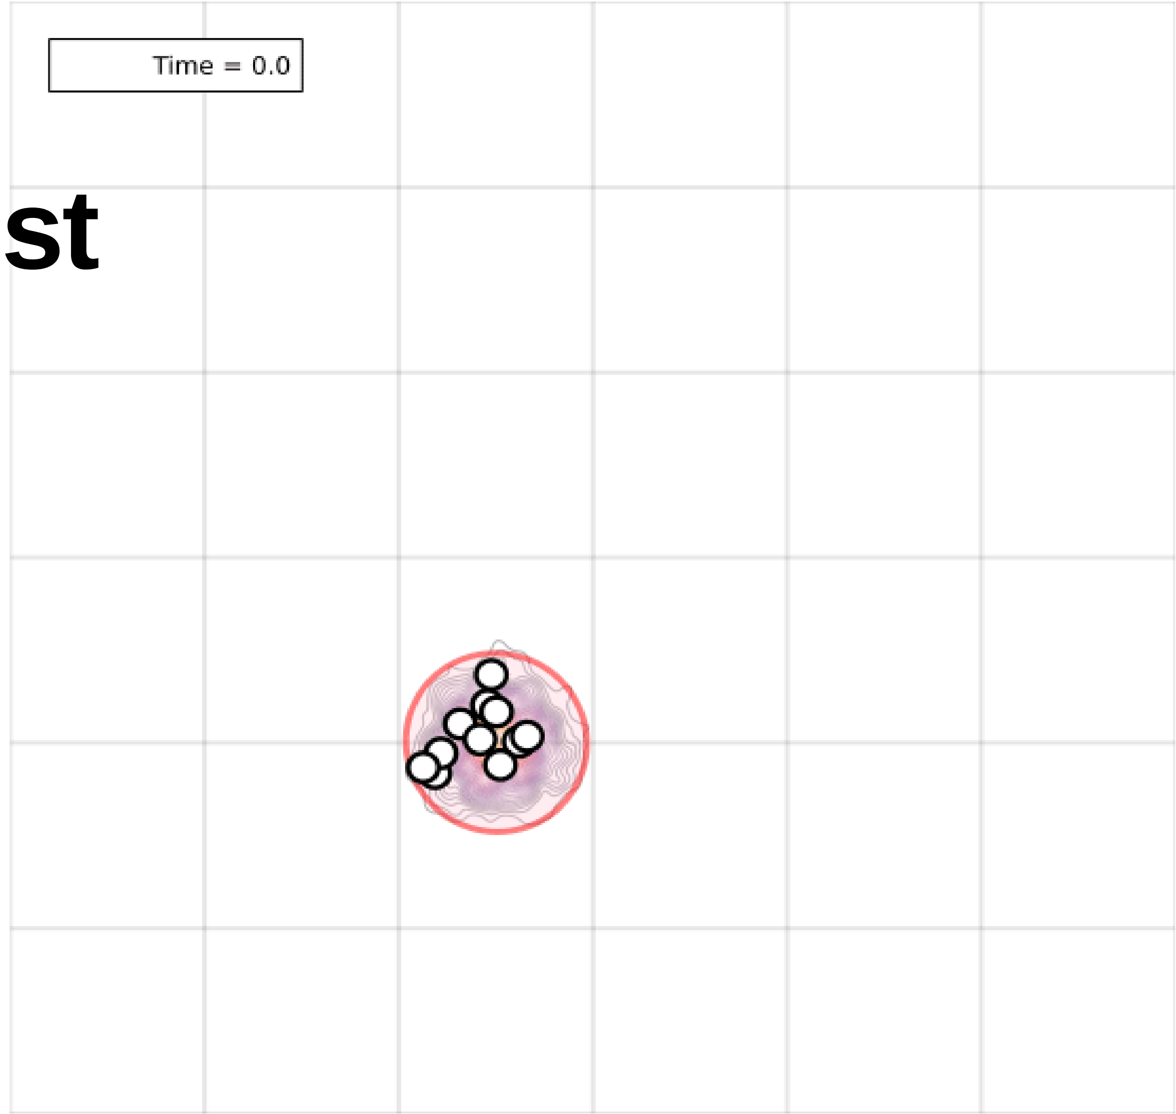

DRAC

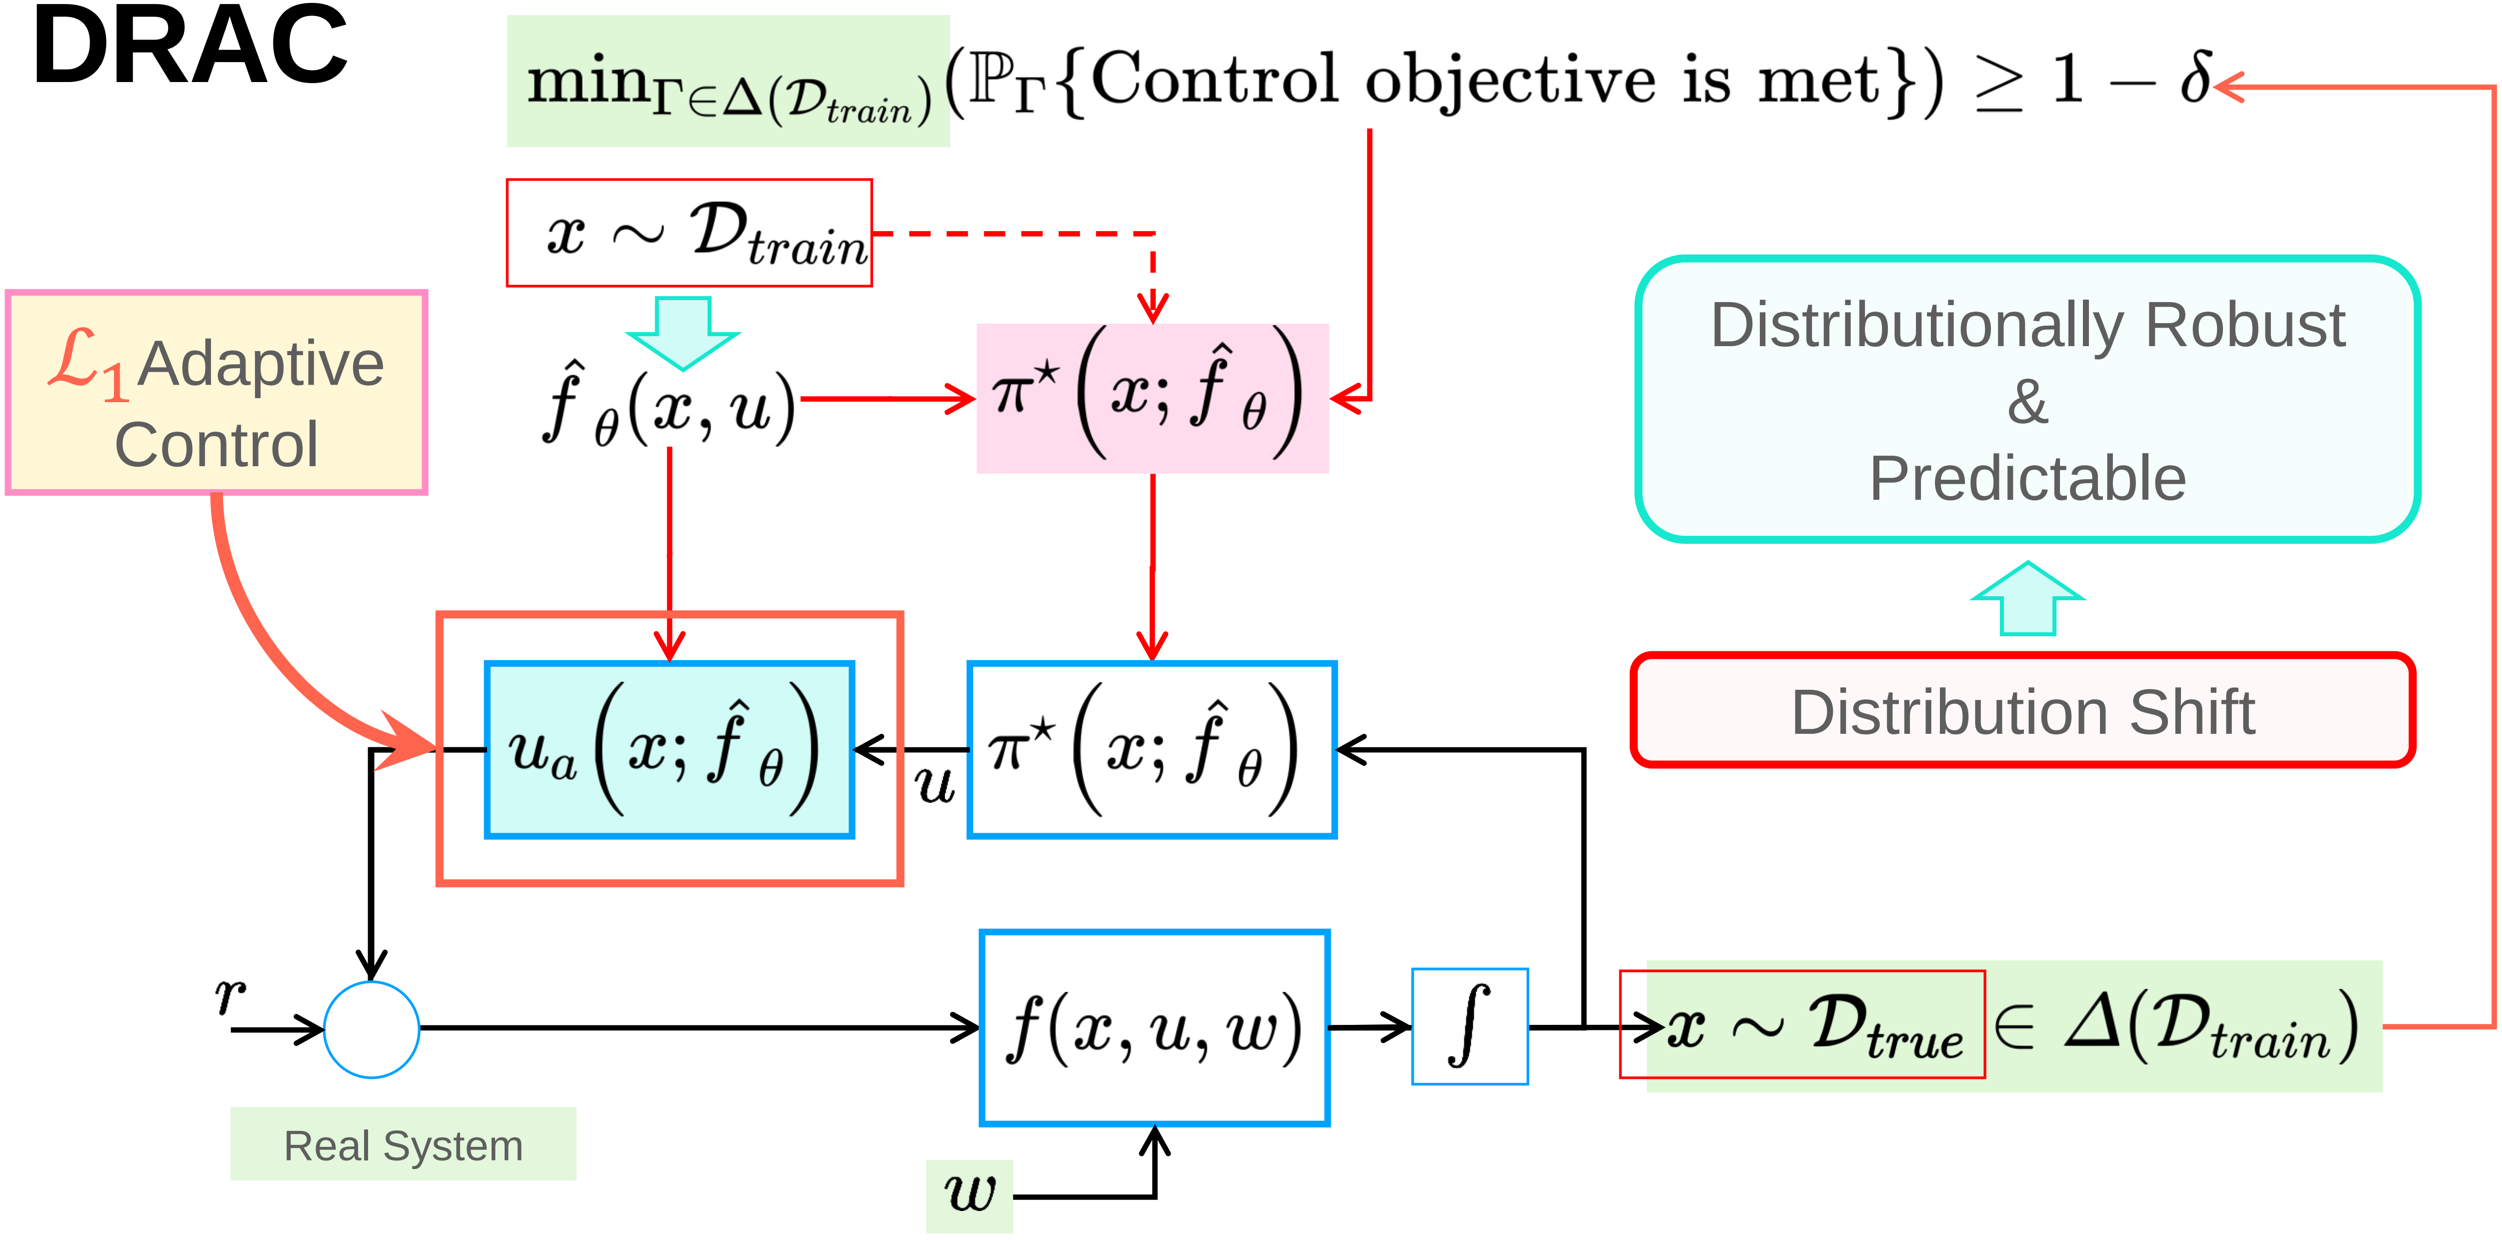

Control augmentation u_a to guarantee certificates of distributional robustness

$\mathcal{L}_1$ Adaptive Control Architecture

- Guaranteed uniform performance bounds and robustness margins
- Validated for manned and unmanned aerial vehicles, oil drilling operations, hydraulic pumps, etc.
- Commercialized by various industries, including Raymarine, Caterpillar, JOUAV Automation Tech, etc.

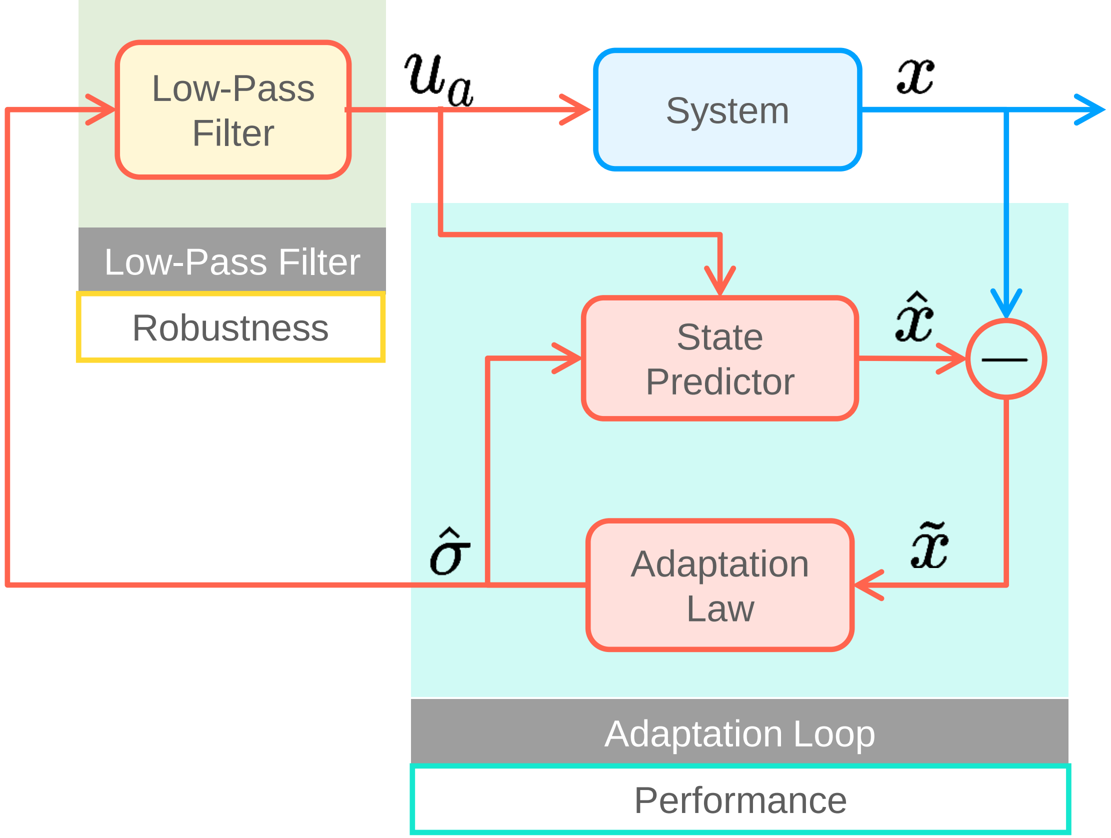

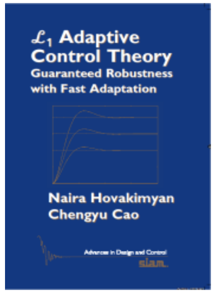

intelinair

Statoil

$\mathcal{L}_1$ Adaptive Control: Timeline

‘06 ‘07 ‘08 ‘09 ‘10 ‘11 ‘12 ‘13 ‘14 ‘15 ‘16 ‘17 ‘18 ‘19 ‘20 ‘21 ‘22

First papers appeared at ACC

UAV path following with $\mathcal{L}_1$ at AIAA GNC

NASA issued grants for testing $\mathcal{L}_1$ on AirSTAR

NASA Boeing Raytheon

Joint papers on highly unstable aircraft configurations

Raymarine Evolution autopilots with $\mathcal{L}_1$

Raymarine by FLIR

NASA issued grants for testing $\mathcal{L}_1$ on AirSTAR

SEAGATE

eurocopter

CAT

MBDA MISSILE SYSTEMS

Statoil

Tests on Learjet and F16 at the Edwards AFB

Distributionally Robust Control

Safe Learning and Control

AFOSR support 2003-2023

$\mathcal{L}_1$ Adaptive Control: Guarantees

$\mathcal{L}_1$ adaptive control provides *certificates* of *performance and robustness*

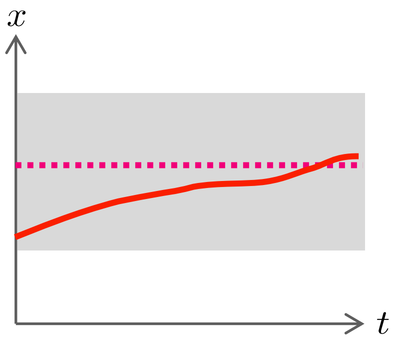

Transient performance

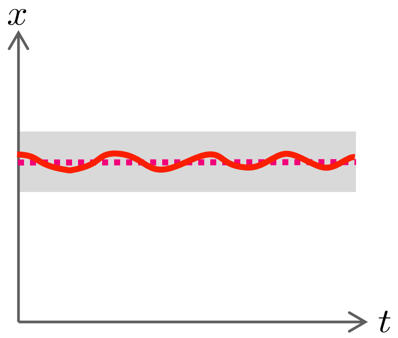

Steady-state performance

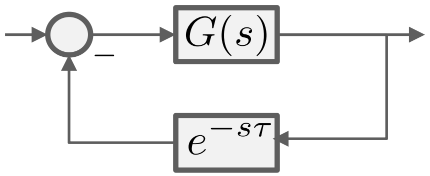

Time-delay margin

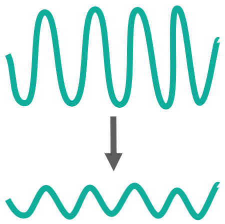

Disturbance rejection

And yet, *as the world around it changes*, new forms of guarantees are needed.

Data-driven systems,
Stochastic representations

Distributional
Guarantees

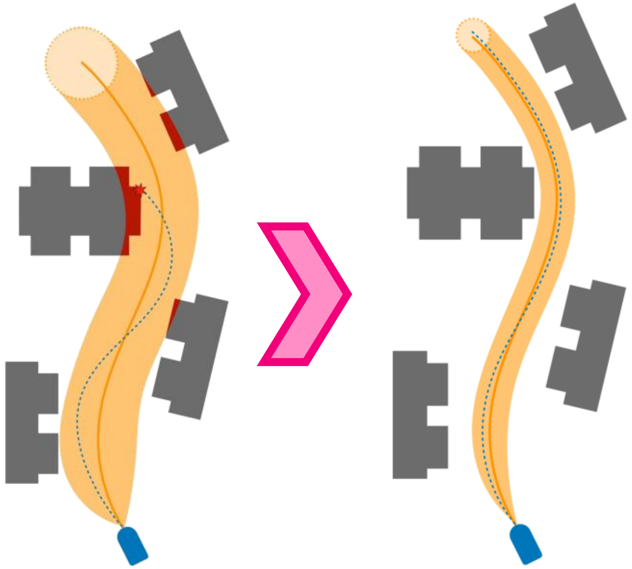

The Systems

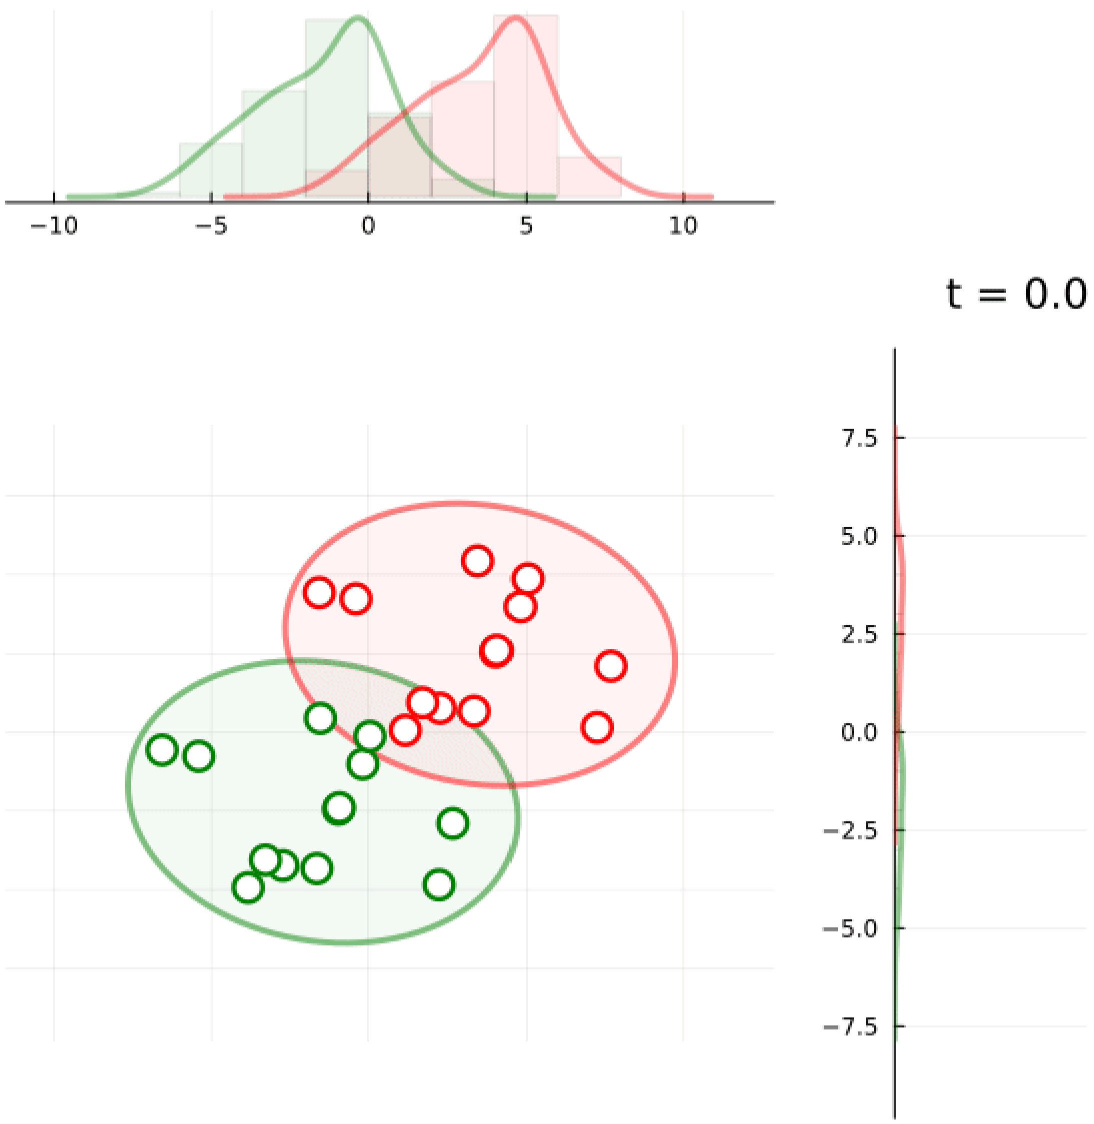

Nonlinear Itô Processes

True System $$dX_t = F_\mu(X_t, U_t)dt + F_\sigma(X_t, U_t; \vartheta)dW_t, \qquad X_t \sim \mathbb{Q}_t$$

- W_t: Brownian motion
  - Gaussian Markov Process
  - Stationary independent increments: Lévy process
  - Continuous and nowhere differentiable, almost surely
  - Motivation: Every almost surely continuous process with independent increments is Gaussian [1]

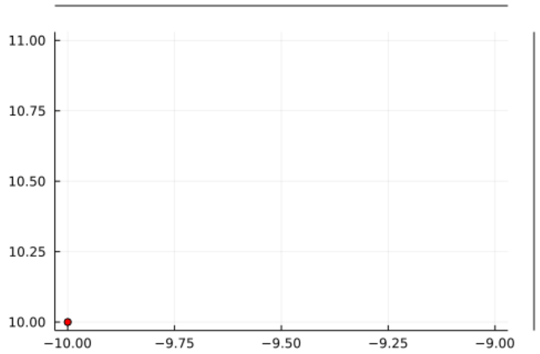

[1] Skorokhod, A. "Random Processes with Independent Increments, Nauka, Moscow, 1964." (1991

Nonlinear Itô Processes

True System $$dX_t = F_\mu(X_t, U_t)dt + F_\sigma(X_t, U_t; \vartheta)dW_t, \quad X_t \sim \mathbb{Q}_t$$

Uncertain drift $$F_\mu(X_t, U_t)dt = f(X_t) + g(X_t)(U_t + h(X_t)) + l(X_t)$$

- Known drift component
- Matched and unmatched uncertainties
  - Locally Lipschitz, linear growth

Nonlinear Itô Processes

True System $dX_t = F_\mu(X_t, U_t)dt + F_\sigma(X_t, U_t; \vartheta)dW_t, \quad X_t \sim \mathbb{Q}_t$

Uncertain diffusion $F_\sigma(X_t, U_t)dt = [\vartheta g(X_t)(U_t + h(X_t)) \quad p(X_t) + q(X_t)]$

- Known diffusion component: uniformly bounded
- Drift uncertainty
  - sublinear growth, Holder continuous $\alpha \leq \frac{1}{2}$
  - Robust approaches fail due to the growth
- Control channel noise parameter $\vartheta \in \mathbb{R}$
  - Stronger results if $h \in S^{2,\infty}_{loc}(\mathbb{R}^m)$ (Sobolev space)
    - ∃ locally essentially bounded weak derivatives up to order 2 [1]

[1] Krylov, Nikolaj Vladimirovič. *Controlled diffusion processes*. Vol. 14. Springer Science & Business Media, 2008.

Systems

True System $$dX_t = F_\mu(X_t, U_t)dt + F_\sigma(X_t, U_t; \vartheta)\boxed{dW_t}, \quad X_t \sim \boxed{\mathbb{Q}_t}$$

$\mathcal{D}_{true}$

$$F_\mu(X_t, U_t)dt = f(X_t) + g(X_t)(U_t + h(X_t)) + l(X_t)$$

$$F_\sigma(X_t, U_t)dt = [\vartheta g(X_t)(U_t + h(X_t)) \quad p(X_t) + q(X_t)]$$

Distribution Shift

Nominal System $$dX_t^\star = \bar{F}_\mu(X_t^\star, U_t^\star)dt + \bar{F}_\sigma\left(X_t^\star, U_t^\star\right)\boxed{dW_t^\star}, \quad X_t^\star \sim \boxed{\mathbb{Q}_t^\star}$$

$\mathcal{D}_{learn}$

Nominal system: No epistemic uncertainties, only aleatoric

Independent Brownian motions	True System Measure (Distribution)	Nominal System Measure (Distribution)
W_t $\quad W_t^\star$	$\mathbb{Q}_t$	$\mathbb{Q}_t^\star$

The Goals

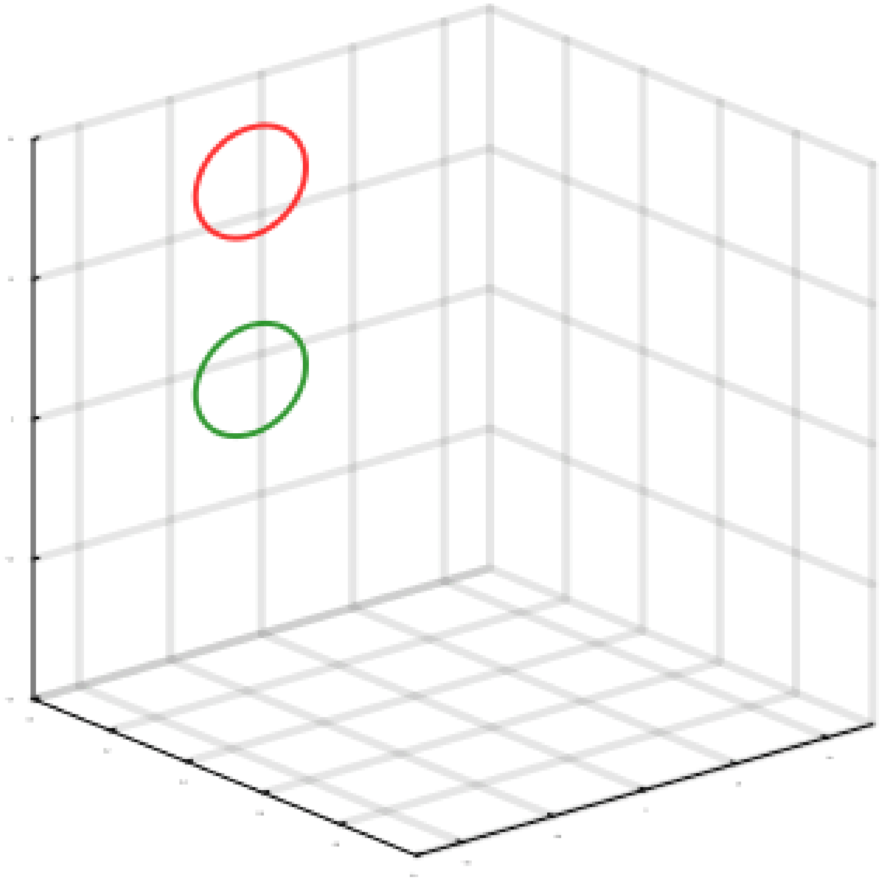

Goals

Nominal System

$$dX_t^\star = \bar{F}_\mu(X_t^\star, U_t^\star)dt + \bar{F}_\sigma(X_t^\star, U_t^\star)dW_t^\star, \qquad X_t^\star \sim \boxed{\mathbb{Q}_t^\star}$$

$$\pi^\star\left(X_t^\star; \hat{f}_\theta\right)$$

Learned via <u>Nominal</u> Distribution

Distribution Shift

True System

$$dX_t = F_\mu(X_t, U_t)dt + F_\sigma(X_t, U_t; \vartheta)dW_t, \qquad X_t \sim \boxed{\mathbb{Q}_t}$$

$$\pi^\star\left(X_t; \hat{f}_\theta\right)$$

- Learned controller on true system: Distribution shift
  - Guarantees of safety and predictability: Invalid

Goals

Any Distribution

Nominal System

$$dX_t^\star = \bar{F}_\mu(X_t^\star, U_t^\star)dt + \bar{F}_\sigma\big(X_t^\star, U_t^\star\big)dW_t^\star, \quad X_t^\star \sim \boxed{\mathbb{Q}_t^\star}$$

Bounded

True System

$$dX_t = F_\mu(X_t, U_t)dt + F_\sigma(X_t, U_t; \vartheta)dW_t, \quad X_t \sim \boxed{\mathbb{Q}_t}$$

$$\pi^\star\left(X_t; \hat{f}_\theta\right) + \boxed{\pi_a\left(X_t; \hat{f}_\theta\right)}$$

- We want to design a feedback augmentation such that
- True distribution $\mathbb{Q}_t$ remains uniformly bounded around the nominal distribution $\mathbb{Q}_t^\star$
    - Robustness bounds used upstream for DR planning and control
- Bound in the sense of Wasserstein metric
    - Optimal transport theory
    - A metric on the space of distributions (distance and shape)

The Goals: Pictorial Depiction

- For each $t \geq 0$ $\underbrace{\mathbb{W}_p\left(\mathbb{Q}_t, \mathbb{Q}_t^\star\right) \leq \rho}_{\text{Wasserstein Distance}} \Rightarrow \underbrace{\mathbb{Q}_t^\star \in \mathcal{A}\left(\rho, \mathbb{Q}_t^\star\right)}_{\text{Ambiguity Set}}$

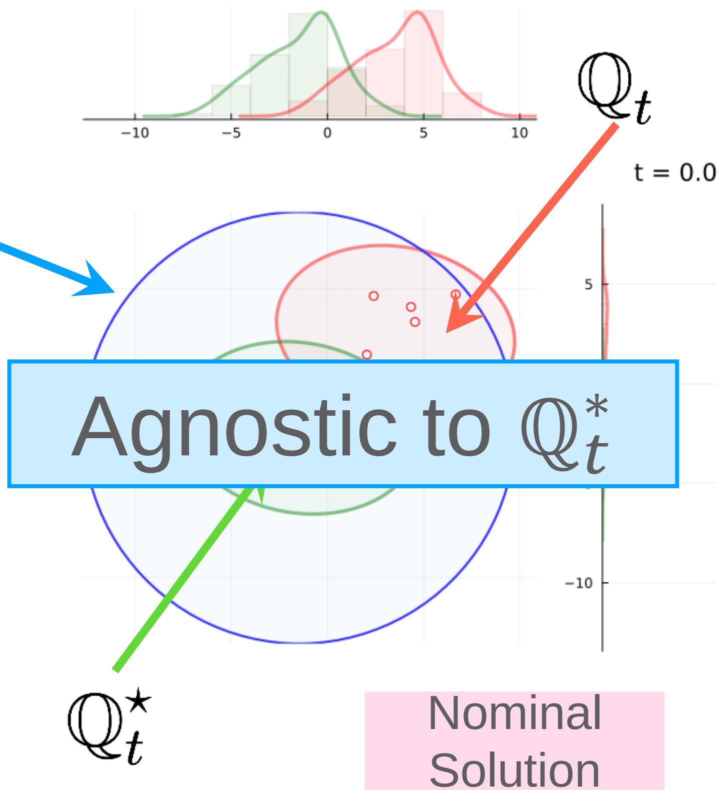

Nominal Solution

$X_t^\star \sim \mathbb{Q}_t^\star$

$X_t \sim \mathbb{Q}_t$

True Solution

Ambiguity tube of distributions $\Omega(\rho)$

Set of path measures (Kolmogorov extension) Distributions on

$\mathcal{C}([0, T]; \mathbb{R}^n)$

Sample trajectories via Girsanov

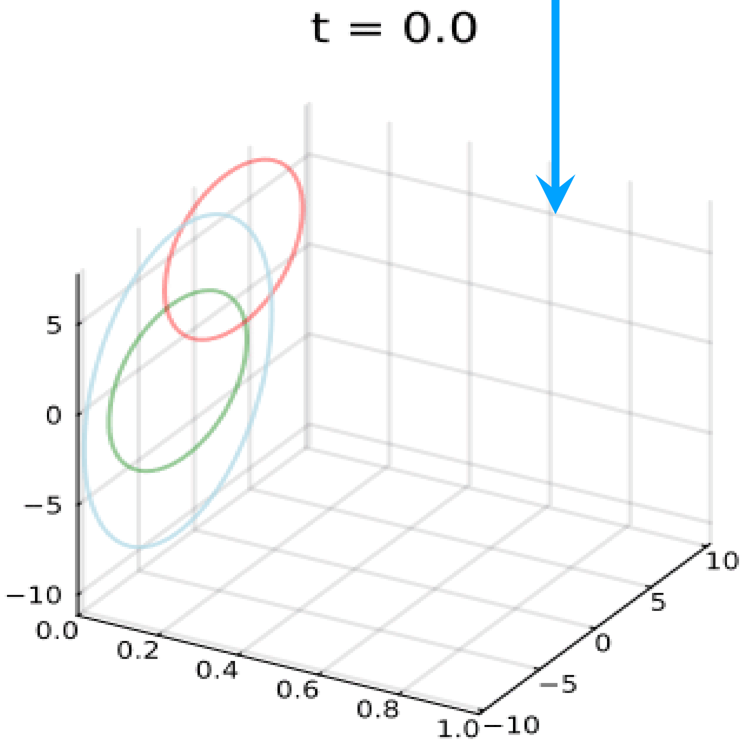

Nominal Stability Assumptions

$\equiv 0$

Nominal System

$$dX_t^\star = \bar{F}_\mu(X_t^\star, U_t^\star)dt + \boxed{\cancel{\bar{F}_\sigma\left(X_t^\star, U_t^\star\right)dW_t^\star}}, \qquad X_t^\star \sim \mathbb{Q}_t^\star$$

Nominal Deterministic Sub-system

$$\dot{x}^\star = \bar{F}_\mu(x^\star, u^\star) \qquad \boxed{\pi^\star\left(x^\star; \hat{f}_\theta\right)} \qquad x^\star(0) \sim \mathbb{Q}_0^\star$$

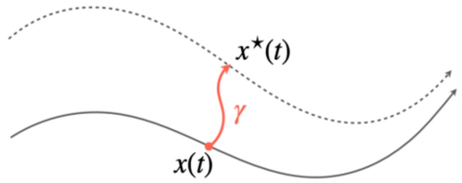

The nominal deterministic subsystem with the nominal controller

- Incrementally exponentially stable (IES)

$$\|x_1^\star(t) - x_2^\star(t)\| \leq Ce^{-\lambda t}\|x_1^\star(0) - x_2^\star(0)\|, \quad \forall t \geq 0 \;\; \forall \;\text{ feasible } (x_1^\star, x_2^\star)$$

Furthermore, ∃ positive scalars α_1, α_2, and λ, and function $V : \mathbb{R}^n \times \mathbb{R}^n \to \mathbb{R}$

$$\alpha_1\|x - x^\star\| \leq V(x, x^\star) \leq \alpha_2\alpha_1\|x - x^\star\|$$

$$L_{F(x^\star, u^\star)}V(x, x^\star) + L_{F(x, u_c)}V(x, x^\star) \leq -2\lambda V(x, x^\star)$$

Incremental Lyapunov function (ILF)

Certificate for IES [1]

[1] Angeli, David. "A Lyapunov approach to incremental stability properties." *IEEE Transactions on Automatic Control* 47.3 (2002): 410-421

Controller

Nominal System

$$dX_t^\star = \bar{F}_\mu(X_t^\star, U_t^\star)dt + \bar{F}_\sigma(X_t^\star, U_t^\star)dW_t^\star, \quad X_t^\star \sim \mathbb{Q}_t^\star$$

True System

$$dX_t = F_\mu(X_t, U_t)dt + F_\sigma(X_t, U_t; \vartheta)dW_t, \quad X_t \sim \mathbb{Q}_t$$

The controller has the architecture of an $\mathcal{L}_1$ adaptive controller

The controller has three main components

State Predictor

Adaptation Law

Low-Pass Filter

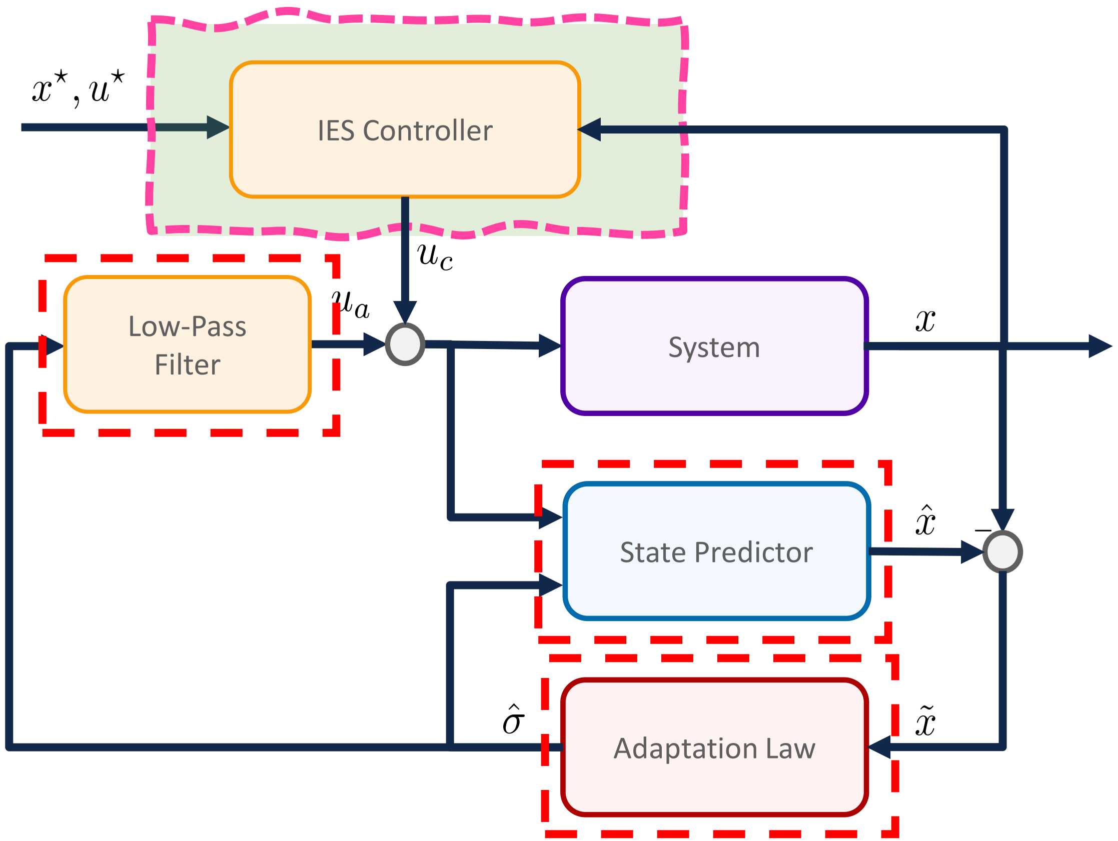

Controller

$$\pi^\star\left(X_t;\hat{f}_\theta\right) \xrightarrow{U_t^\star} \boxed{\pi_a\left(X_t;\hat{f}_\theta\right)} \xrightarrow{} U_{a,t}$$

True System

$$dX_t = F_\mu(X_t, U_t)dt \, + \, F_\sigma(X_t, U_t; \vartheta)dW_t, \qquad X_t \sim \mathbb{Q}_t$$

State Predictor

Adaptation Laws

Control Law

$$d\hat{X}_t = \left[A_s\left(\hat{X}_t - X_t\right) + \bar{F}_\mu(X_t, U_t^\star + U_{a,t}) + g(X_t)\hat{\sigma}_m(t) + \hat{\sigma}_{um}(t)\right]$$

$$\hat{X}_t \sim \hat{\mathbb{Q}}_t$$

- Components:
  - Known deterministic subsystem
  - Cannot estimate diffusion due to independent Brownian motion→fundamental limitation
- Deterministic dynamics, stochastic state
  - State-feedback injection
  - Predictor driven by a state induced by its dynamics and Brownian motion
    - Colored noise
- Adaptive estimates

Controller

$$\pi^\star\left(X_t; \hat{f}_\theta\right) \xrightarrow{U_t^\star} \boxed{\pi_a\left(X_t; \hat{f}_\theta\right)} \quad U_{a,t}$$

True System $\quad dX_t = F_\mu(X_t, U_t)dt + F_\sigma(X_t, U_t; \vartheta)dW_t, \quad X_t \sim \mathbb{Q}_t$

State Predictor | Adaptation Laws | Control Law

$$\begin{bmatrix} \hat{\sigma}_m(t) \\ \hat{\sigma}_{um}(t) \end{bmatrix} \doteq \hat{\sigma}(t) = \hat{\sigma}(iT_s) = -\bar{g}(X_{iT_s})^{-1}\Phi^{-1}\mu(iT_s),$$
$$(t, i) \in [iT_s, (i+1)T_s) \times \mathbb{N}$$

- Piecewise constant adaptive law
  - T_s: Sampling period
  - Quality of adaptive estimates $\propto \frac{1}{T_s}$

  Can run at a rate up to the digital hardware limit
- Entities $\bar{g}$, Φ^{-1}, μ are (partially) computed before runtime
  - Minimal computation to produce adaptive estimates
- Numerically stable implementation → avoids stiffness

Controller

$$\pi^\star\left(X_t;\hat{f}_\theta\right) \xrightarrow{U_t^\star} \boxed{\pi_a\left(X_t;\hat{f}_\theta\right)} \xrightarrow{} U_{a,t}$$

True System $\quad dX_t = F_\mu(X_t, U_t)dt + F_\sigma(X_t,U_t;\vartheta)dW_t, \qquad X_t \sim \mathbb{Q}_t$

State Predictor

Adaptation Laws

Control Law

$$U_{a,t} = (\mathcal{K}_\omega \hat{\sigma}_m(\cdot))(t) = -\omega \int_0^t e^{-\omega(t-s)} \hat{\sigma}_m(s)ds$$

- Feedback operator
  - Linear bounded operator $L_\infty^{loc}(\mathbb{R}_{\geq 0}, \mathbb{R}^m) \to L_\infty^{loc}(\mathbb{R}_{\geq 0}, \mathbb{R}^m)$
  - First order low-pass filter with bandwidth ω
- $\hat{\sigma}_m(t)$: adaptive estimate of the matched uncertainty

Main Result

Theorem: Suppose for any p, the initial nominal and true measures $\mathbb{Q}_0^*$ and $\mathbb{Q}_0$ satisfy

$$\mathbb{W}_p(\mathbb{Q}_0^\star, \mathbb{Q}_0) \leq \infty$$

Given any ϵ and ρ_a, define

$$\rho_r = \left(\left(\frac{\alpha_2}{\alpha_1}\right)^p \mathbb{W}_p(\mathbb{Q}_0^\star, \mathbb{Q}_0) + \sum_{i=1}^{2} \zeta_i(\omega) + \zeta_3(T_s) + \sum_{i=1}^{3} C_i \right)^{\frac{1}{p}}, \quad \rho = \rho_r + \rho_a,$$

where, α_i are the bounds on the ILF $V(.,.)$, C_i, $i \in \{1, \dots, 3\}$, are constants dependent on the unmatched uncertainty $l(x)$, diffusion terms $p(x)$ and $q(x)$, and the control noise parameter ϑ, respectively. Moreover, $\zeta_i(\omega)$,$i \in \{1,2\}$, and $\zeta_3(\omega)$ are functions of the filter bandwidth ω, and the adaptive law's sampling period T_s, respectively.

Then, there exists a filter-bandwidth ω_0 and adaptation-rate Γ_0 such that for $\omega \geq \omega_0$ and $\Gamma \geq \Gamma_0$,

$$\mathbb{W}_p(\mathbb{Q}_t^\star, \mathbb{Q}_t) \leq \rho, \quad \forall t \geq 0$$ UB

$$\mathbb{W}_p(\mathbb{Q}_t^\star, \mathbb{Q}_t) \leq D\mathbb{W}_p(\mathbb{Q}_0^\star, \mathbb{Q}_0)e^{-\lambda T} + \sum_{i=1}^{2} \zeta_2(\omega) + \zeta_3(\omega) + \sum_{i=1}^{3} C_i, \ \forall t \geq 0 \geq T > 0$$ UUB

where $\zeta_i(\omega) \propto \frac{1}{\omega}$, $i \in \{1,2\}$, $\zeta_3(\omega) \propto T_s$, and D is a constant.

Main Result

$$\rho = \boxed{\rho_r} + \boxed{\rho_a}$$

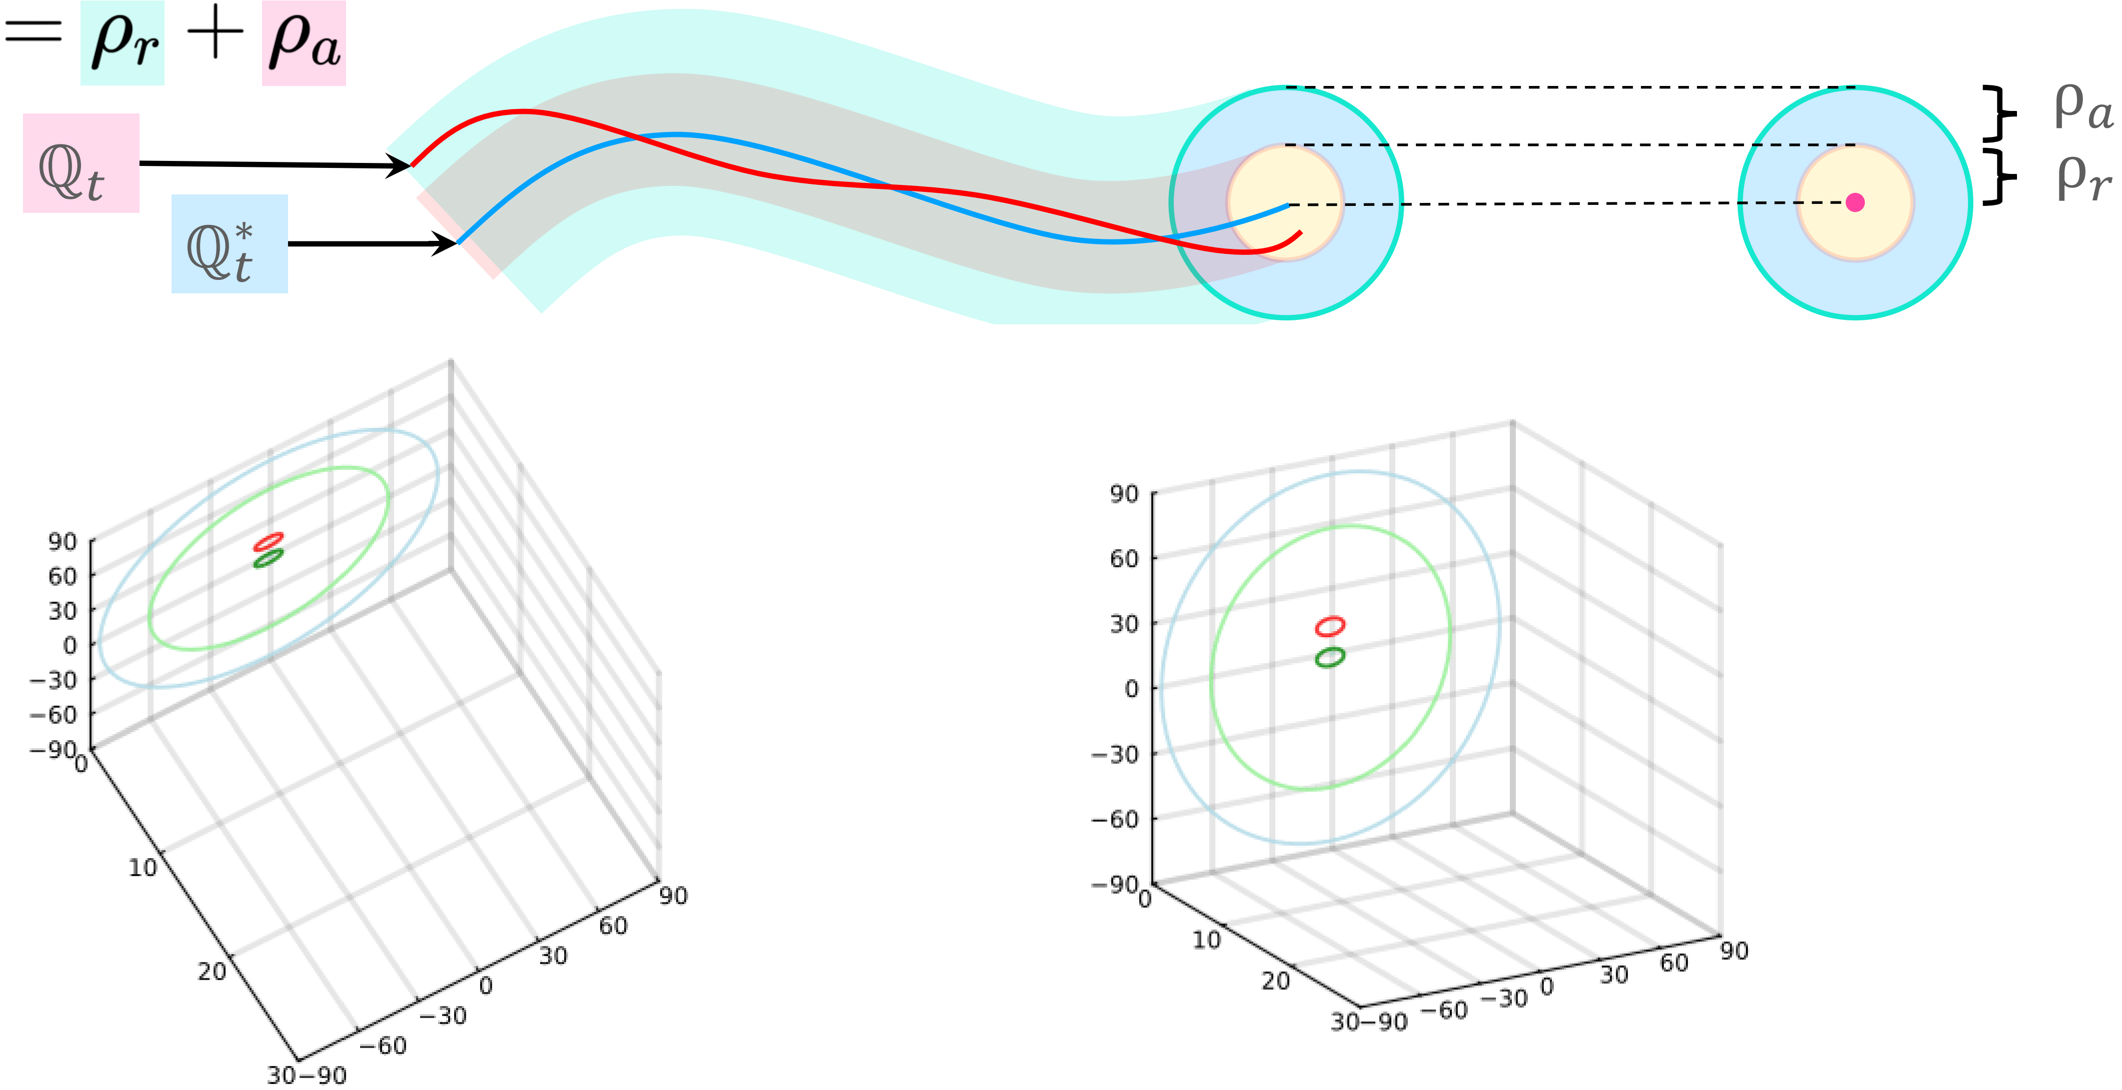

Main Result

$$\rho = \rho_r + \rho_a$$

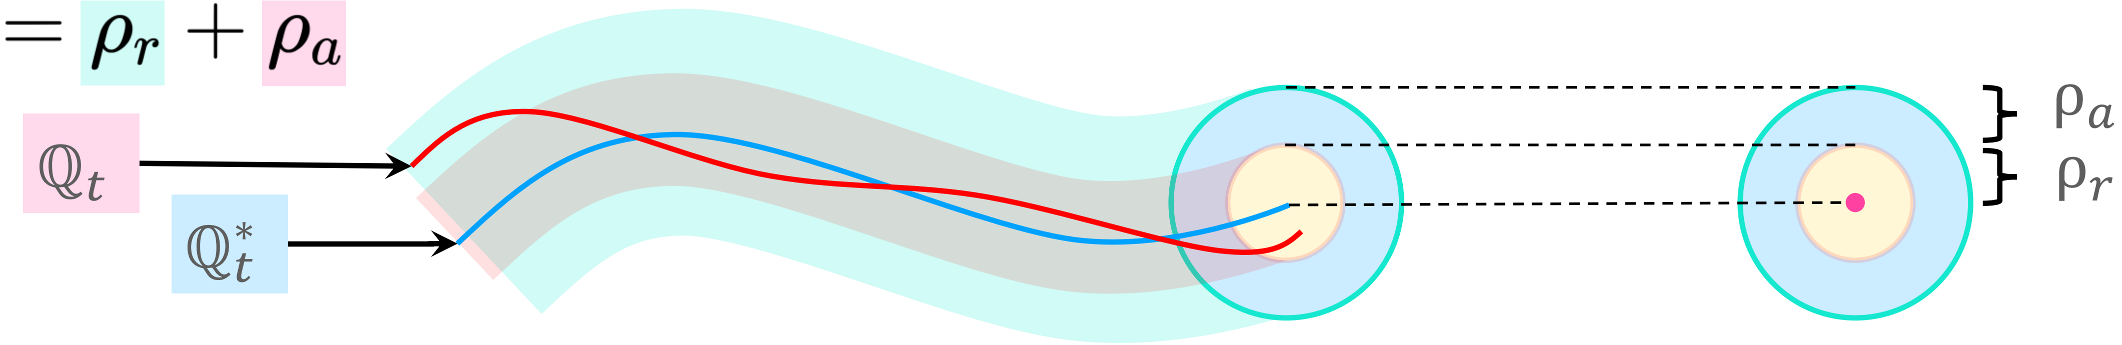

$$\mathbb{W}_p(\mathbb{Q}_t^\star, \mathbb{Q}_t) \leq D\mathbb{W}_p(\mathbb{Q}_0^\star, \mathbb{Q}_0)e^{-\lambda T} + \sum_{i=1}^{2} \zeta_2(\omega) + \zeta_3(\omega) + \sum_{i=1}^{3} C_i, \ \forall t \geq 0 \geq T > 0$$

- Using the transient bounds of $\mathcal{L}_1$ adaptive controller
- Bound on the p^{th}-Wasserstein due to Burkholder-Davis-Gundy inequality [1]
  - Bound on supremum of martingales via their quadratic variation
- Term C_1 due to unmatched uncertainty
- Term C_2 due to independence of Brownian motions [2]
  - Lower bounded by a strictly positive $C > 0$
  - Not present in deterministic systems

[1] Ren, Yao-Feng. "On the Burkholder–Davis–Gundy inequalities for continuous martingales." *Statistics & probability letters* 78.17 (2008): 3034-3039
[2] Pham, Quang-Cuong, Nicolas Tabareau, and Jean-Jacques Slotine. "A contraction theory approach to stochastic incremental stability." *IEEE Transactions on Automatic Control* 54.4 (2009): 816-820.

Main Result

$$\rho = \rho_r + \rho_a$$

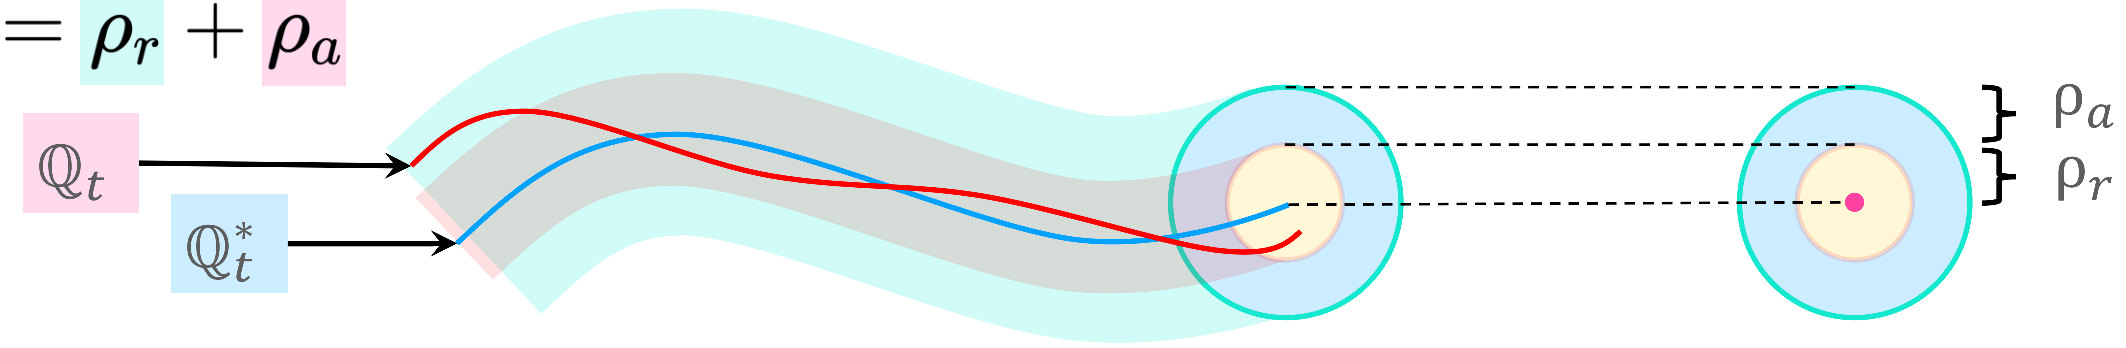

$$\mathbb{W}_p(\mathbb{Q}_t^\star, \mathbb{Q}_t) \leq D\mathbb{W}_p(\mathbb{Q}_0^\star, \mathbb{Q}_0)e^{-\lambda T} + \sum_{i=1}^{2} \zeta_2(\omega) + \zeta_3(\omega) + \sum_{i=1}^{3} C_i, \ \forall t \geq 0 \geq T > 0$$

- Term C_3 due to noisy control channel
  - Can obtain stronger results when $h \in S^{2,\infty}_{loc}(\mathbb{R}^m)$ (Generalized Ito Lemma)
  - Feedback operator is Frechet Differentiable
  - Ito lemma for weakly differentiable functions [1]

[1] Krylov, Nikolaj Vladimirovič. *Controlled diffusion processes*. Vol. 14. Springer Science & Business Media, 2008.

Main Result

$$\rho = \rho_r + \rho_a$$

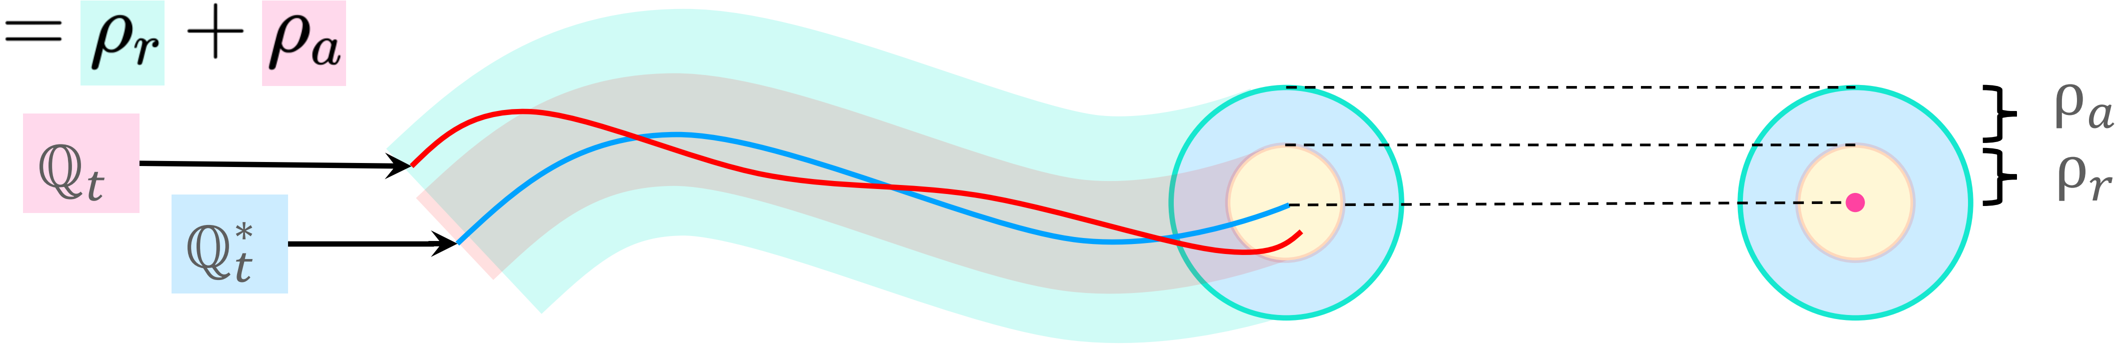

$$\mathbb{W}_p(\mathbb{Q}_t^\star, \mathbb{Q}_t) \leq D\mathbb{W}_p(\mathbb{Q}_0^\star, \mathbb{Q}_0)e^{-\lambda T} + \sum_{i=1}^{2} \zeta_2(\omega) + \zeta_3(\omega) + \sum_{i=1}^{3} C_i, \ \forall t \geq 0 \geq T > 0$$

In the case when W_t^* and $W_t \equiv 0$ (deterministic)

- Wasserstein → Euclidean norm
- We recover $\mathcal{L}_1$ guarantees for nonlinear deterministic systems [1]

[1] Lakshmanan, Arun, Aditya Gahlawat, and Naira Hovakimyan. "Safe feedback motion planning: A contraction theory and $\mathcal{L}$ 1-adaptive control based approach." *2020 59th IEEE Conference on Decision and Control (CDC)*. IEEE, 2020.

Numerical Experimentation

Angular rate dynamics of a quadrotor

Drift and diffusion uncertainties

Divergence of nominal and true distributions

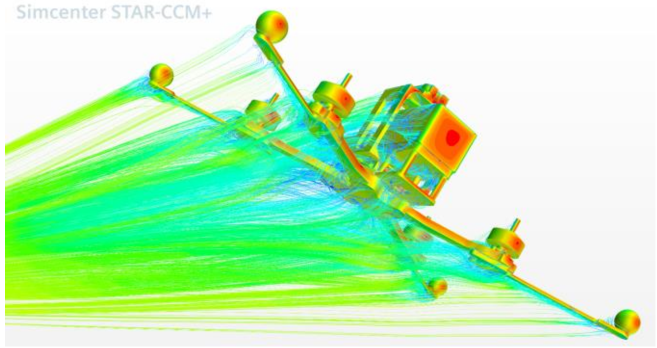

$$W_2(\mathbb{Q}_t^{\star}, \mathbb{Q}_t) : 99.259$$

$$W_2(\mathbb{Q}_t^{\star}, \mathbb{Q}_t) : 100.314$$

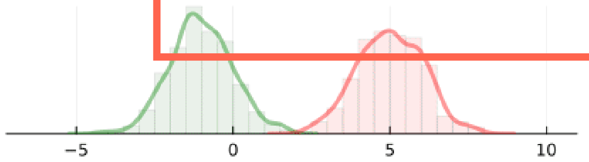

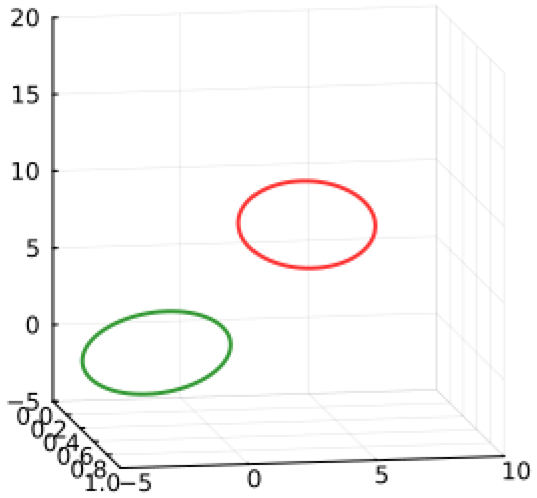

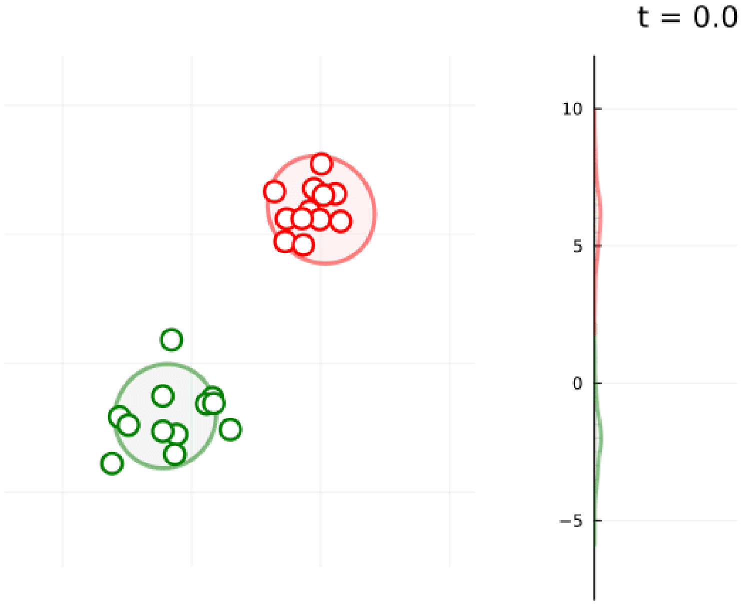

Numerical Experimentation

DRAC control

Independent Brownian motions → Convergence up to a nonzero limit

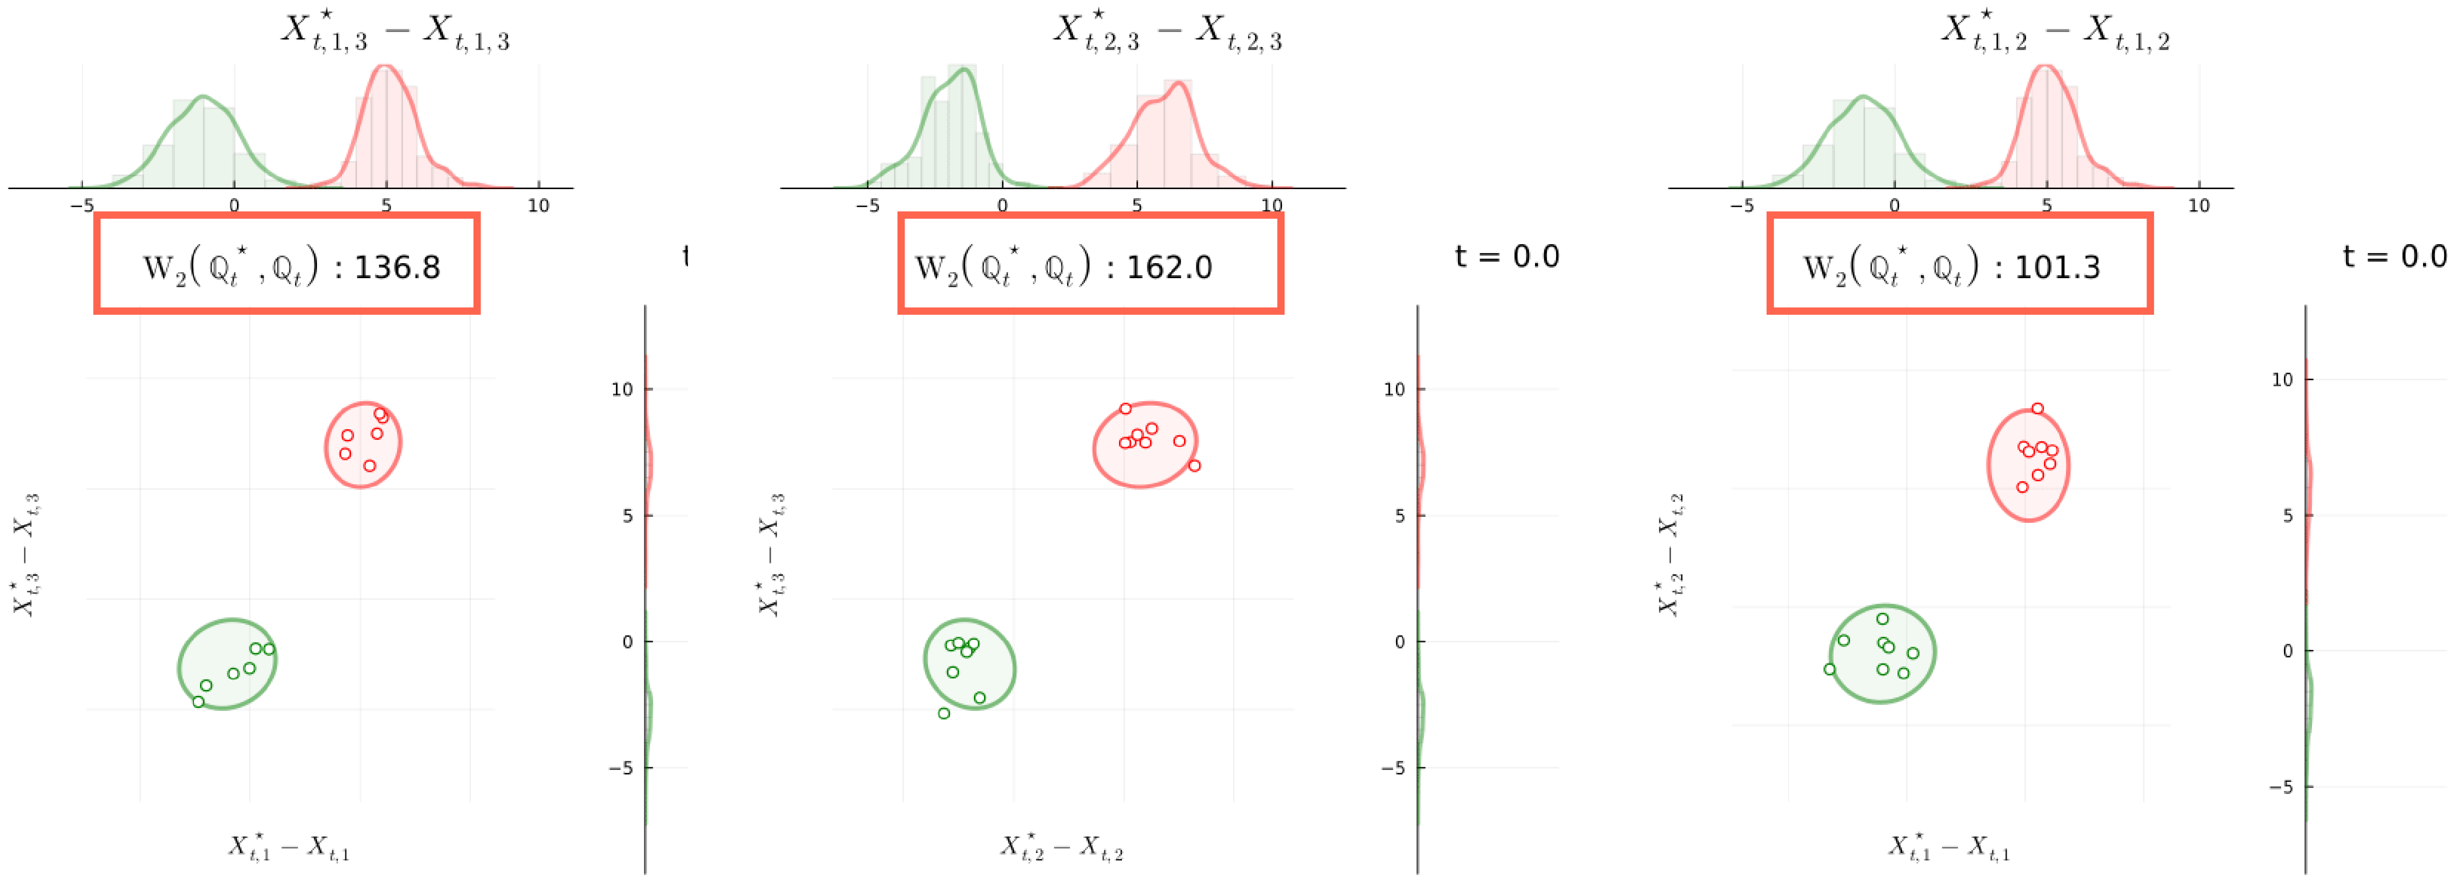

Boundedness and Convergence of distributions in the Wasserstein metric

Numerical Experimentation

DRAC control

Independent Brownian motions → Convergence up to a nonzero limit

$\mathrm{W}_2(\mathbb{Q}_t^\star, \mathbb{Q}_t)$: 162.023

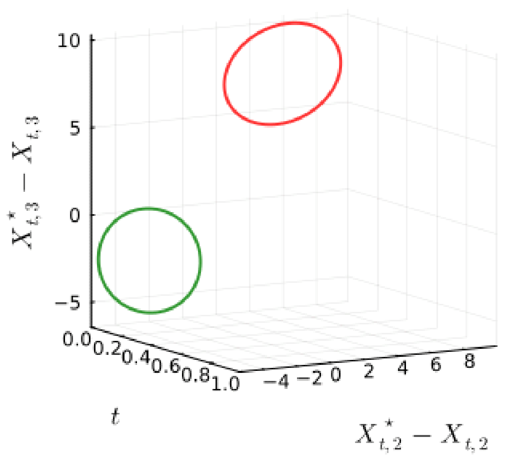

$\mathrm{W}_2(\mathbb{Q}_t^\star, \mathbb{Q}_t)$: 101.314

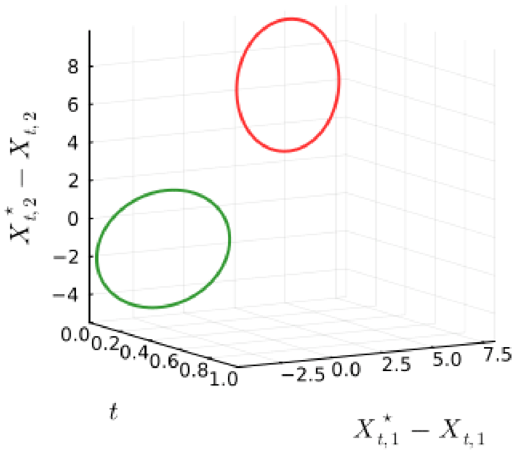

$\mathrm{W}_2(\mathbb{Q}_t^\star, \mathbb{Q}_t)$: 136.801

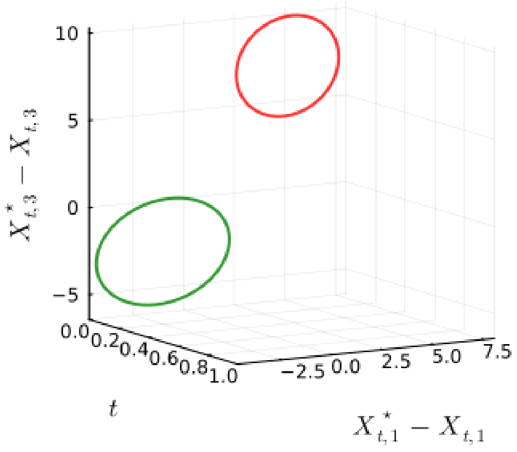

Boundedness and Convergence of distributions in the Wasserstein metric

Continuation

- Extension to general Lévy processes
  - Non-Gaussian
  - Continuous in probability
- Further experimentation of DRAC
  - Learned systems subject to distribution shifts
- V&V of controlled systems with learned components in the loop
  - Distributional certificates
  - Deep learned dynamics and controllers
  - Learned sensing (perception)
- Propagation of robust data-driven certificates through the complete control pipeline
  - Distributional robustness as a language for sensing, planning, and control to communicate
- Distributionally robust planning and control

Ongoing Projects

- Social Information Dynamics and Control (AFOSR)
- NASA ULI on Robust Resilient Autonomy (AVIATE Center, UIUC)
- Two NSF projects on safe learning and robotics, one pending
- Industry developments at Lockheed Martin with Air Force Academy vehicles
- Potential opportunities at AFRL with Boeing (Archer)

- **Acknowledgments:**
  - Sitao Zhang (MS student) and Sambhu Karamanas (Ph.D. student)
  - Evangelos Theodorou (GaTech; introduction to AlphaPilot drone racing environment)
  - Chris Elliott (Lockheed Martin)
  - Irene Gregory (NASA LaRC)