A Distributionally Robust Adaptive Approach to Data-driven Control

Aditya Gahlawat, Sambhu Karumanchi

Advanced Controls Research Laboratory
PI: Naira Hovakimyan

Mechanical Science and Engineering
University of Illinois at Urbana-Champaign

NASA ULI Seminar, September 8, 2023

Data-driven Control Pipeline

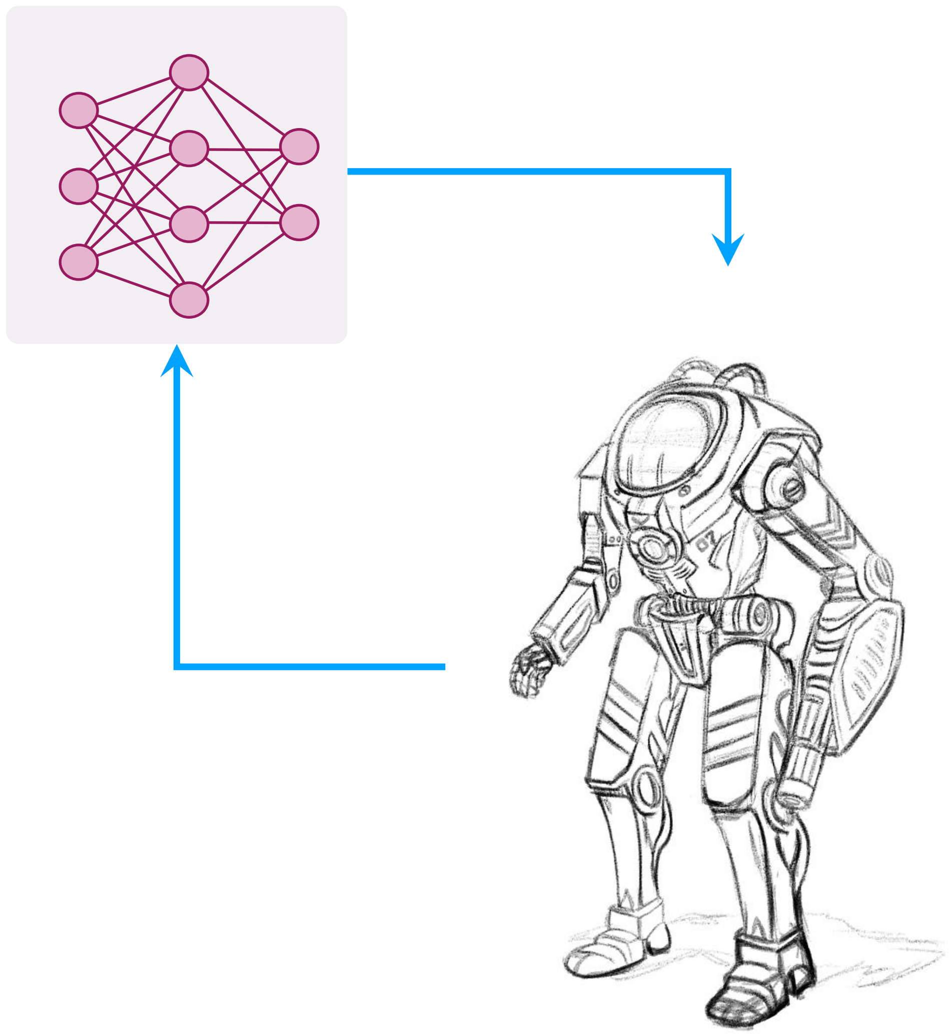

Data-driven Control Design

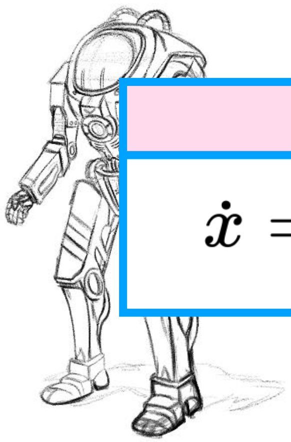

System

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

→

Data

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

↓

Regression

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

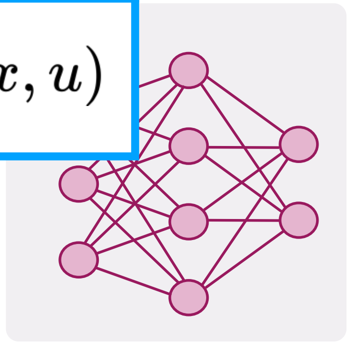

←

Control Design

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

↑

Distribution Dependent Pipeline!

Data Distribution Dependence

Data

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

←

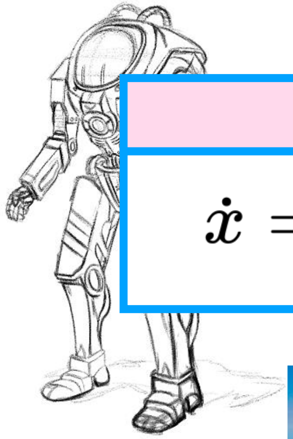

System

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

$$\cup$$

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

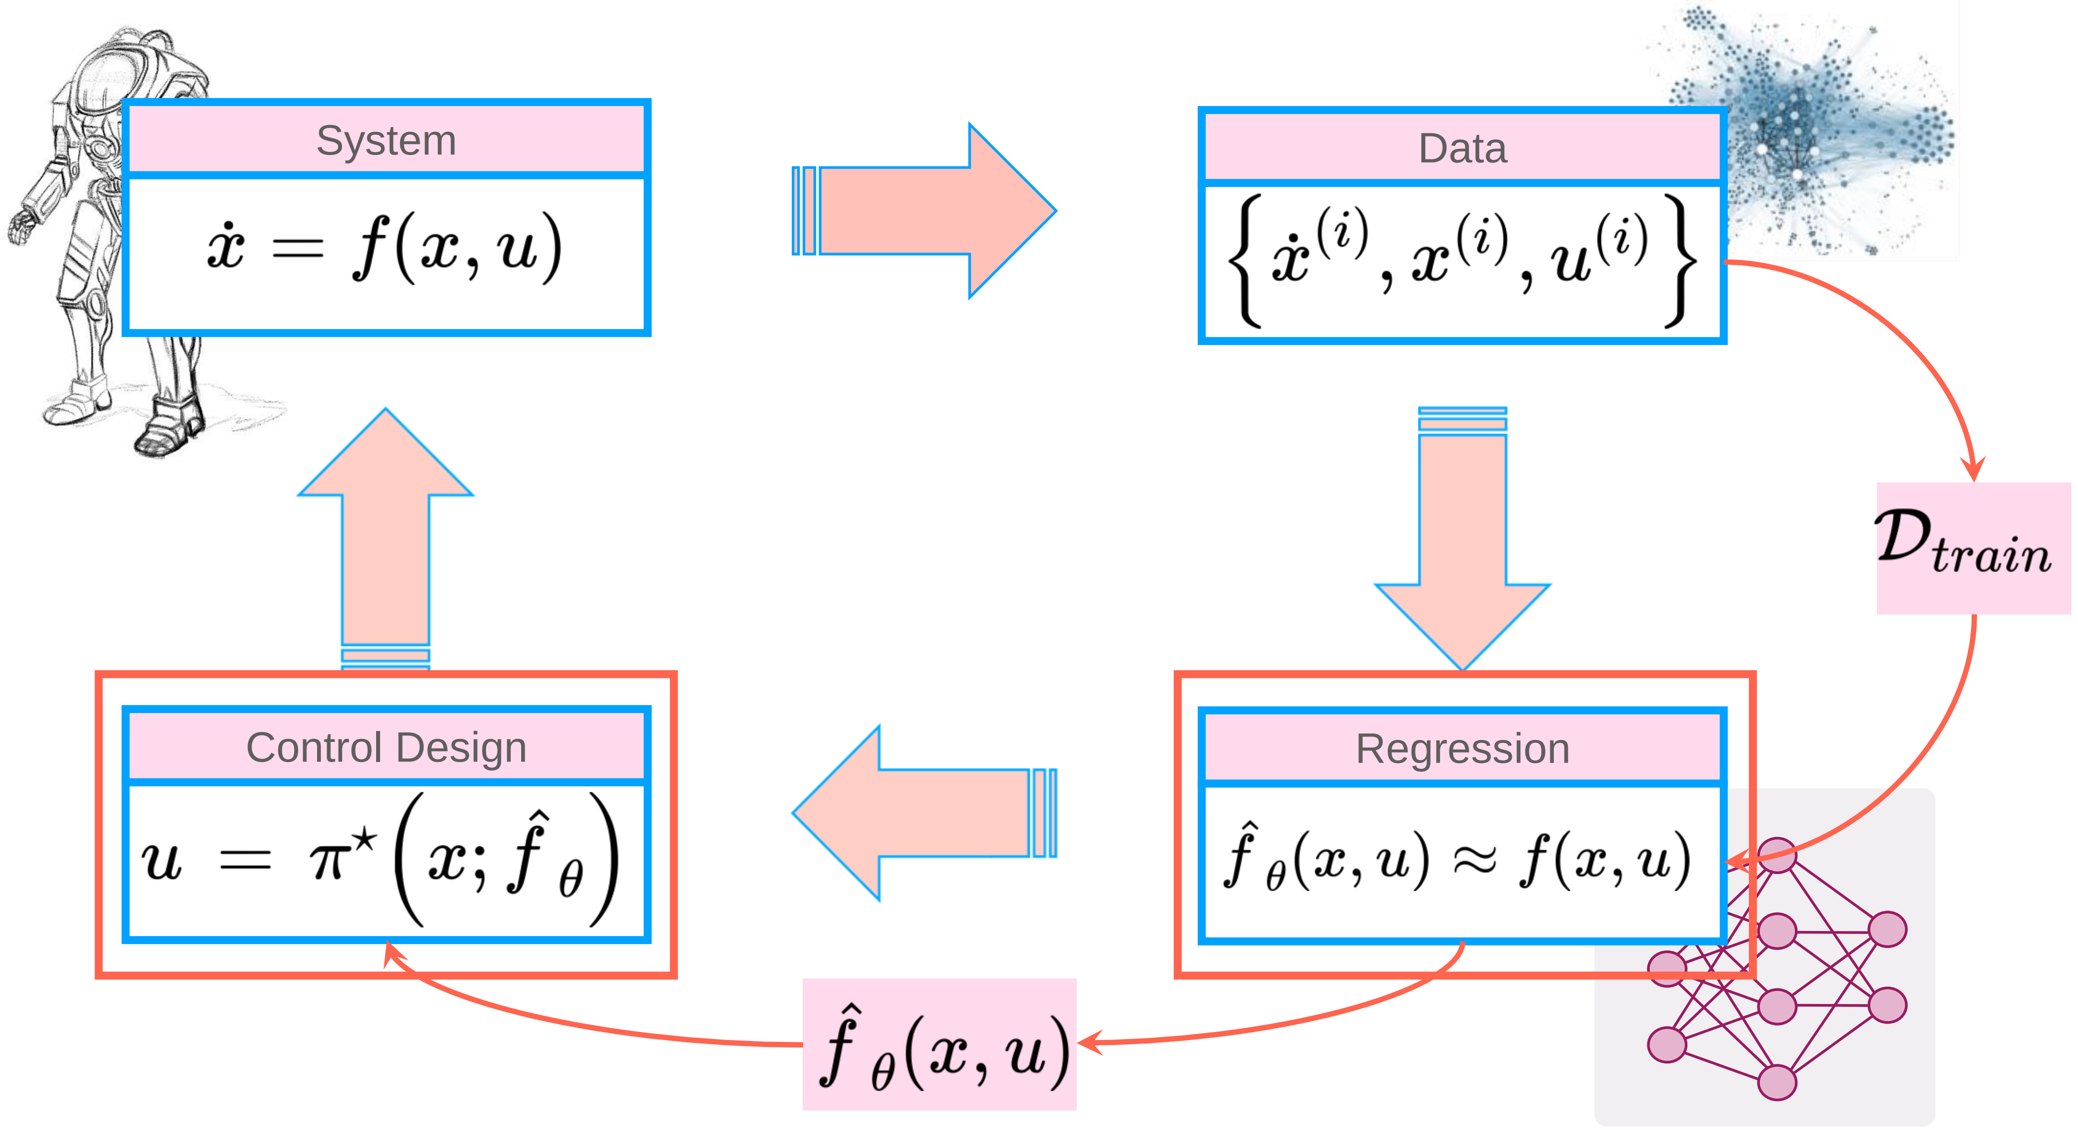

The Role of Training Distribution

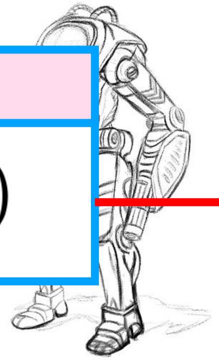

System

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

Data

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

Regression

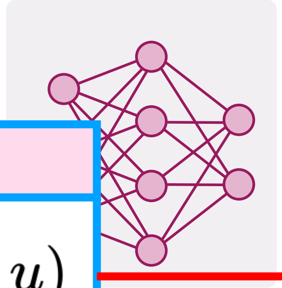

$$\hat{f}_\theta(x, u) \approx f(x, u)$$

Control Design

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

$\mathcal{D}_{true} \neq \mathcal{D}_{train}$ **Distribution Shift!**

Epistemic & Aleatoric

$\mathcal{D}_{train}$

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

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

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

Robustness & Distribution Shifts

Epsitemic & 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

  - We can mitigate the distribution shift
    - We can control the size of the guaranteed set $\Delta(\mathcal{D}_{train})$

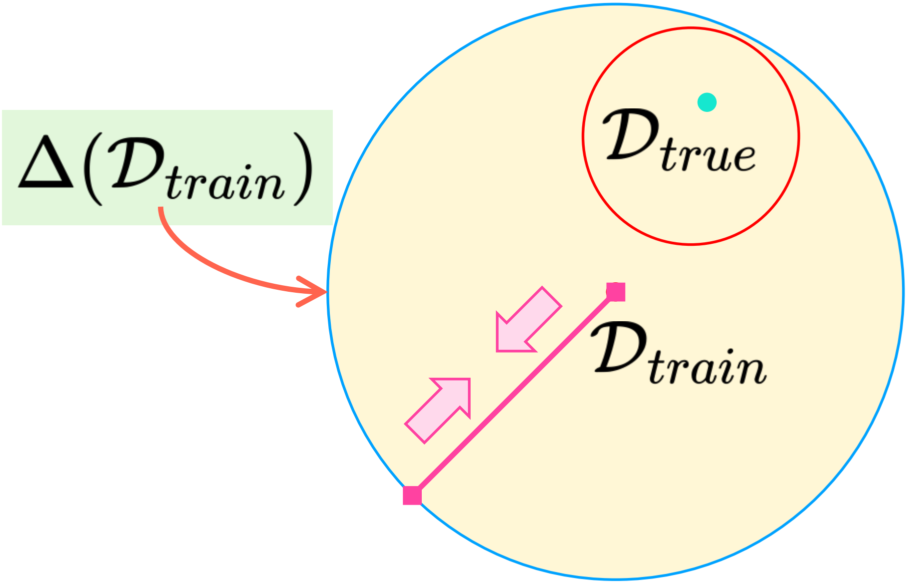

$$\min_{\Gamma \in \Delta(\mathcal{D}_{train})} (\mathbb{P}_\Gamma\{\text{Control objective is met}\}) \geq 1-\delta$$

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

Distributionally Robust Adaptive Control (DRAC)

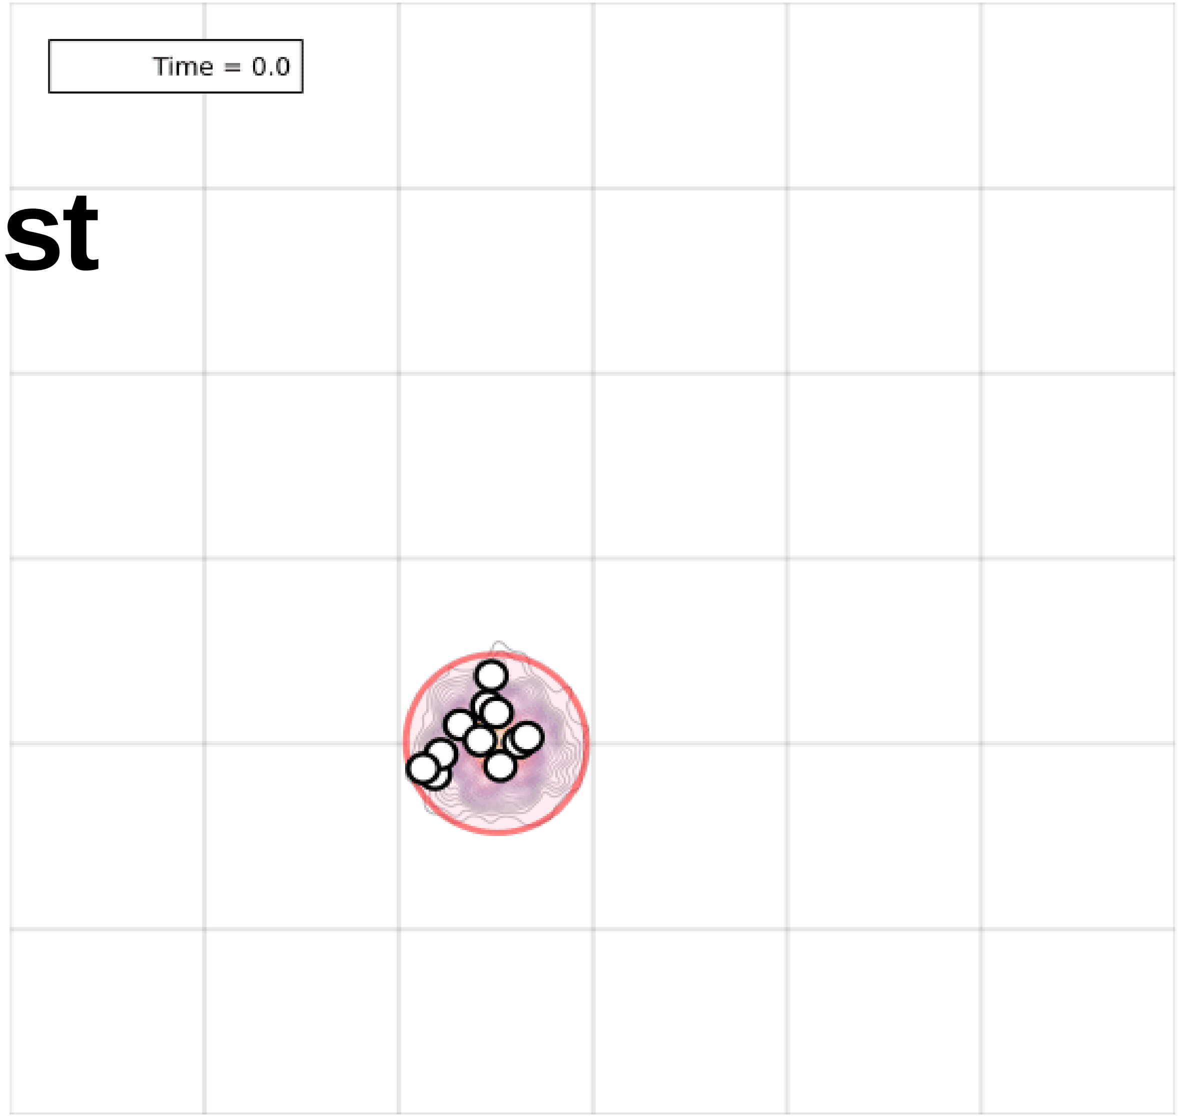

DRAC

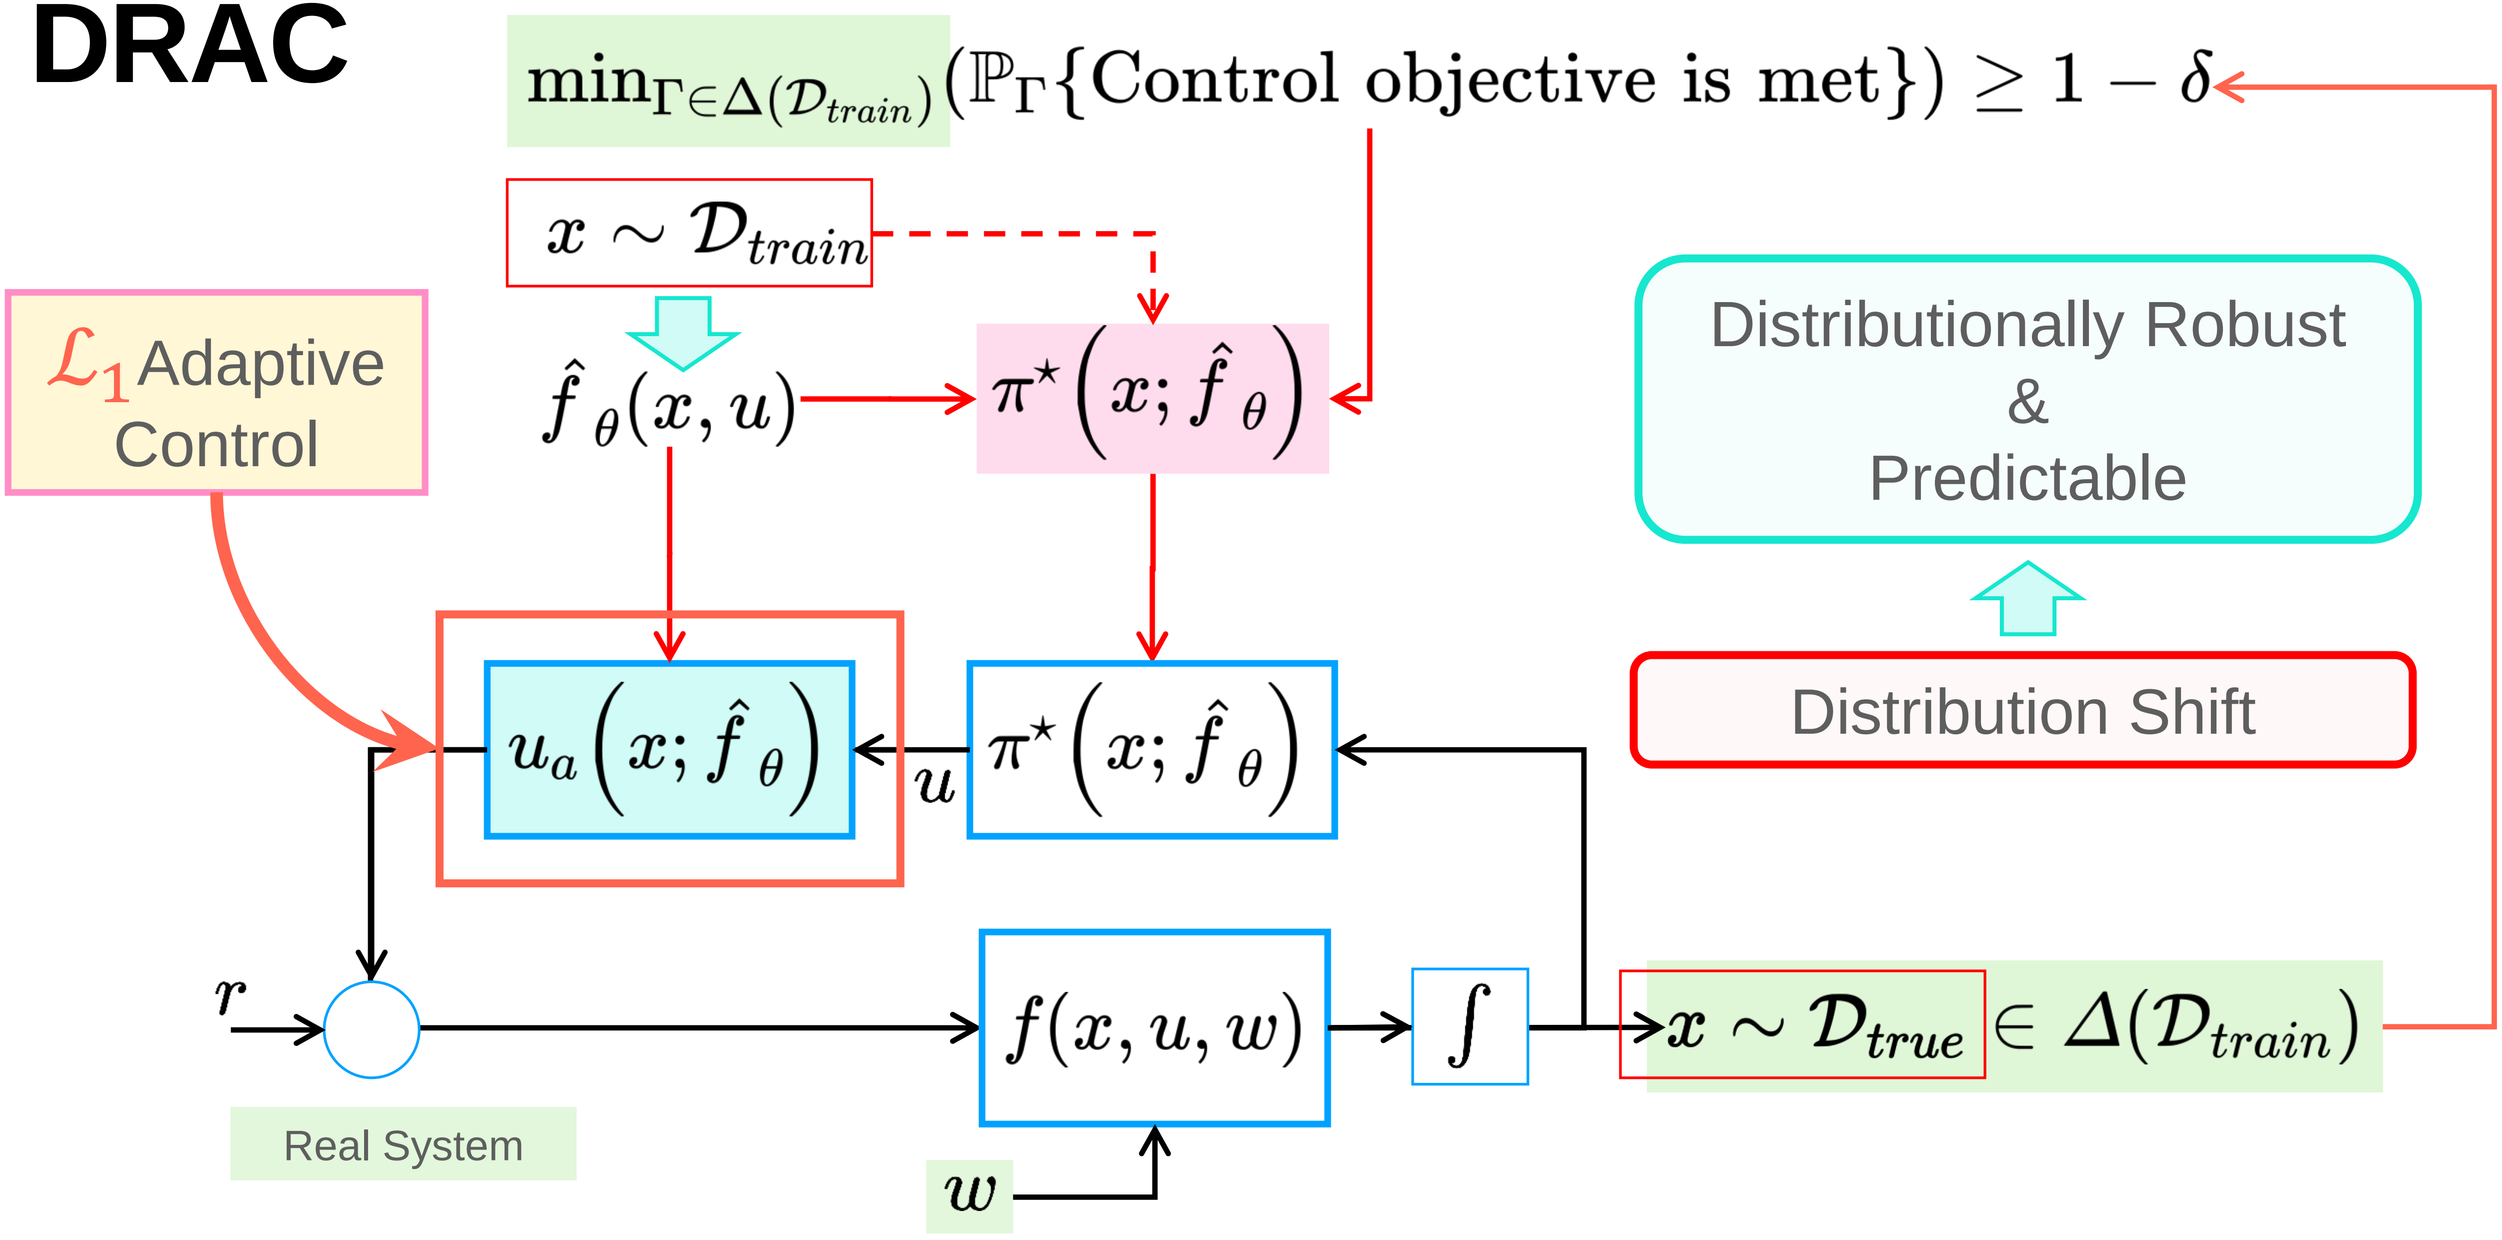

Control augmentation u_a to guarantee certificates of distributional robustness

$\mathcal{L}_1$ 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.

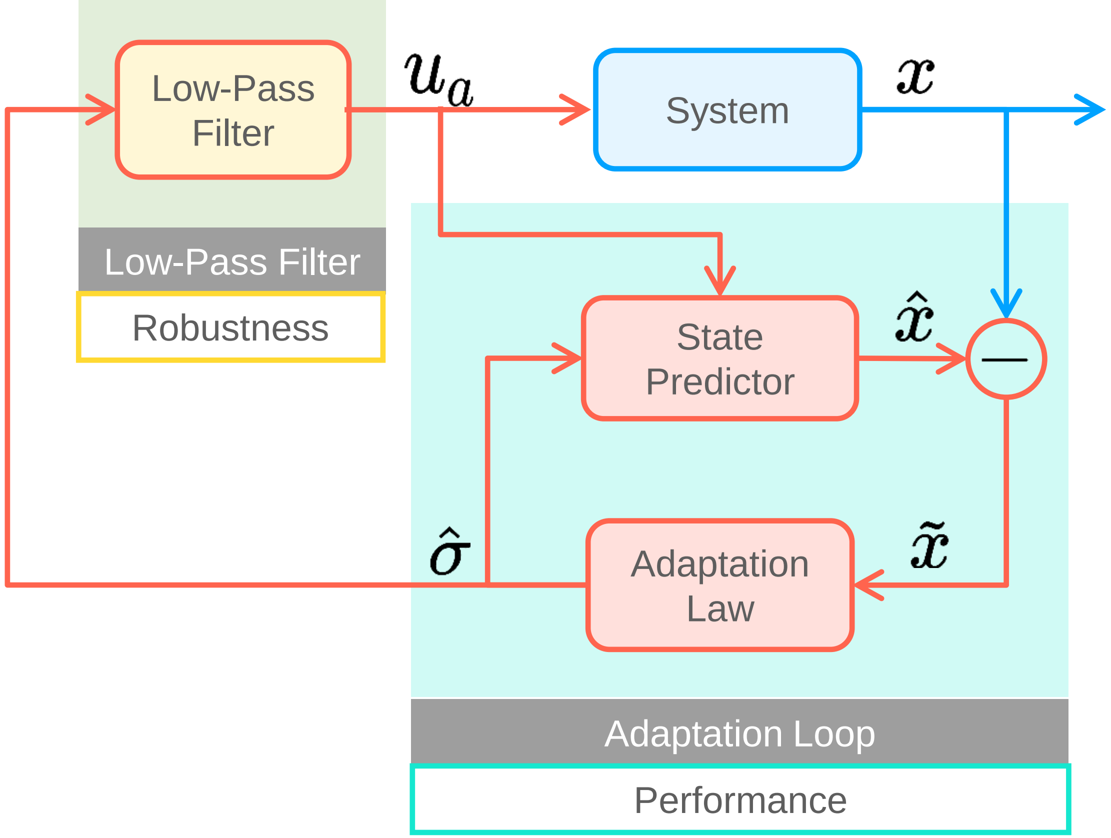

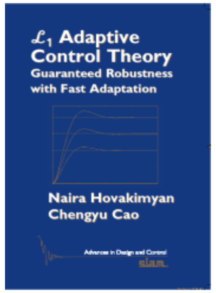

intelinair

Statoil

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

First papers appeared at ACC

Joint papers on highly unstable aircraft configurations

Raytheon

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

Distributionally Robust Control

Safe Learning and Control

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

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

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

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

The Systems

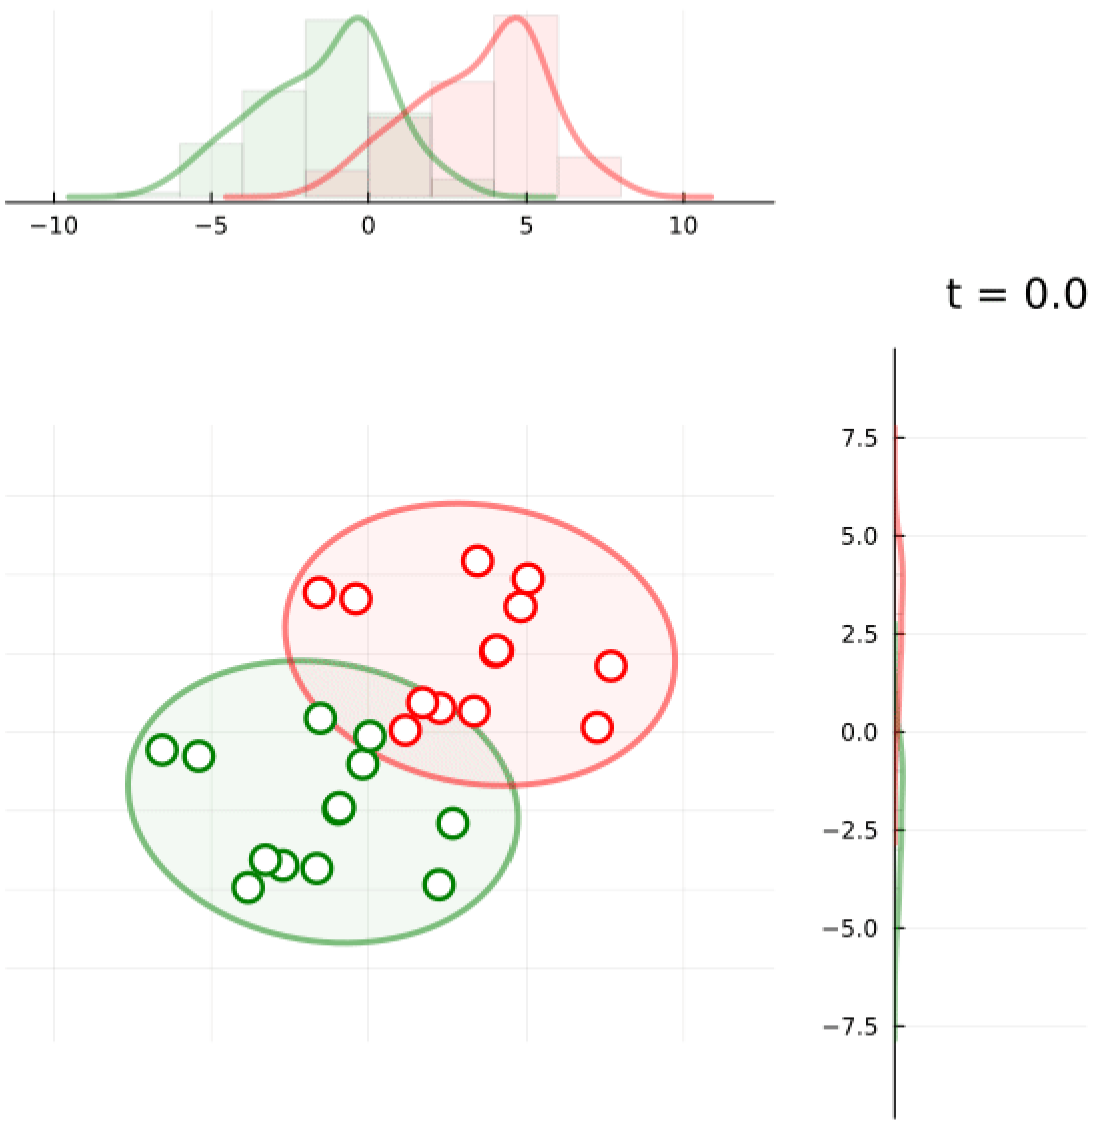

Nonlinear Itô Processes

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]
    - Modelling uncertain (learned) systems

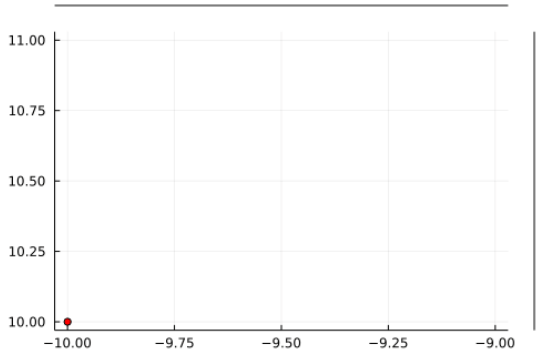

[1] Skorokhod, A. "Random Processes with Independent Increments, Nauka, Moscow, 1964." (1991

Nonlinear Itô Processes

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

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)
    - twice-weakly differentiable and locally essentially bounded: ReLU DNN

Systems

True System

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

$\mathcal{D}_{true}$

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

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

Distribution Shift

Nominal System

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

$\mathcal{D}_{learn}$

True system - epistemic uncertainties

Independent Brownian motions

$W_t \qquad W_t^\star$

True System Measure (Distribution)

$\mathbb{Q}_t$

Nominal System Measure (Distribution)

$\mathbb{Q}_t^\star$

The Goals

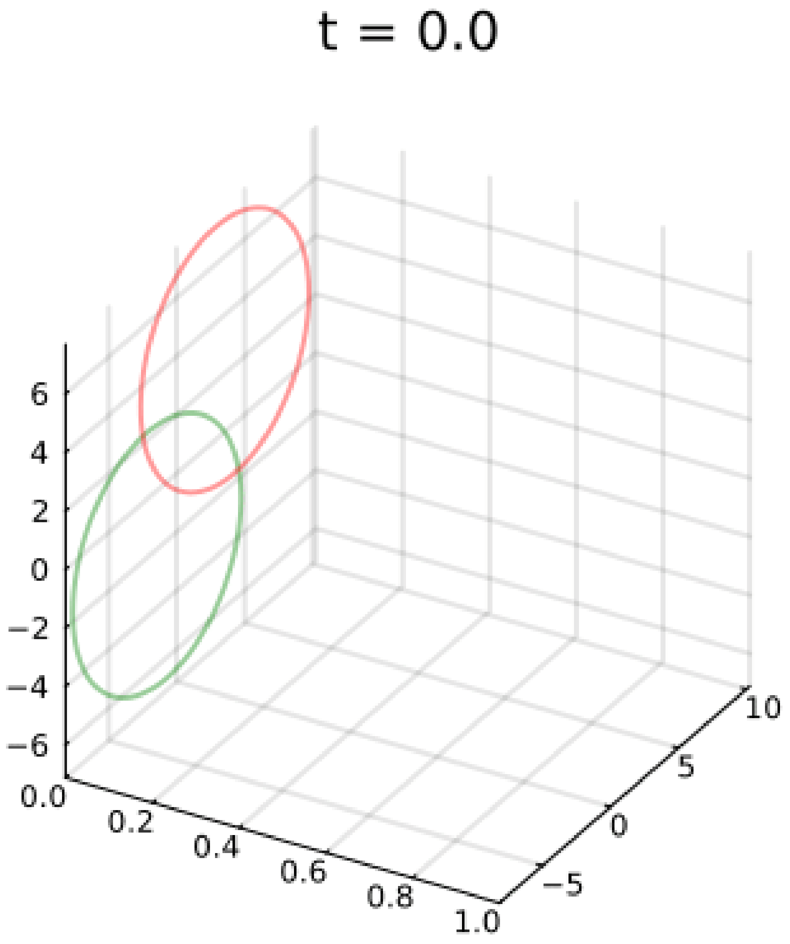

Goals

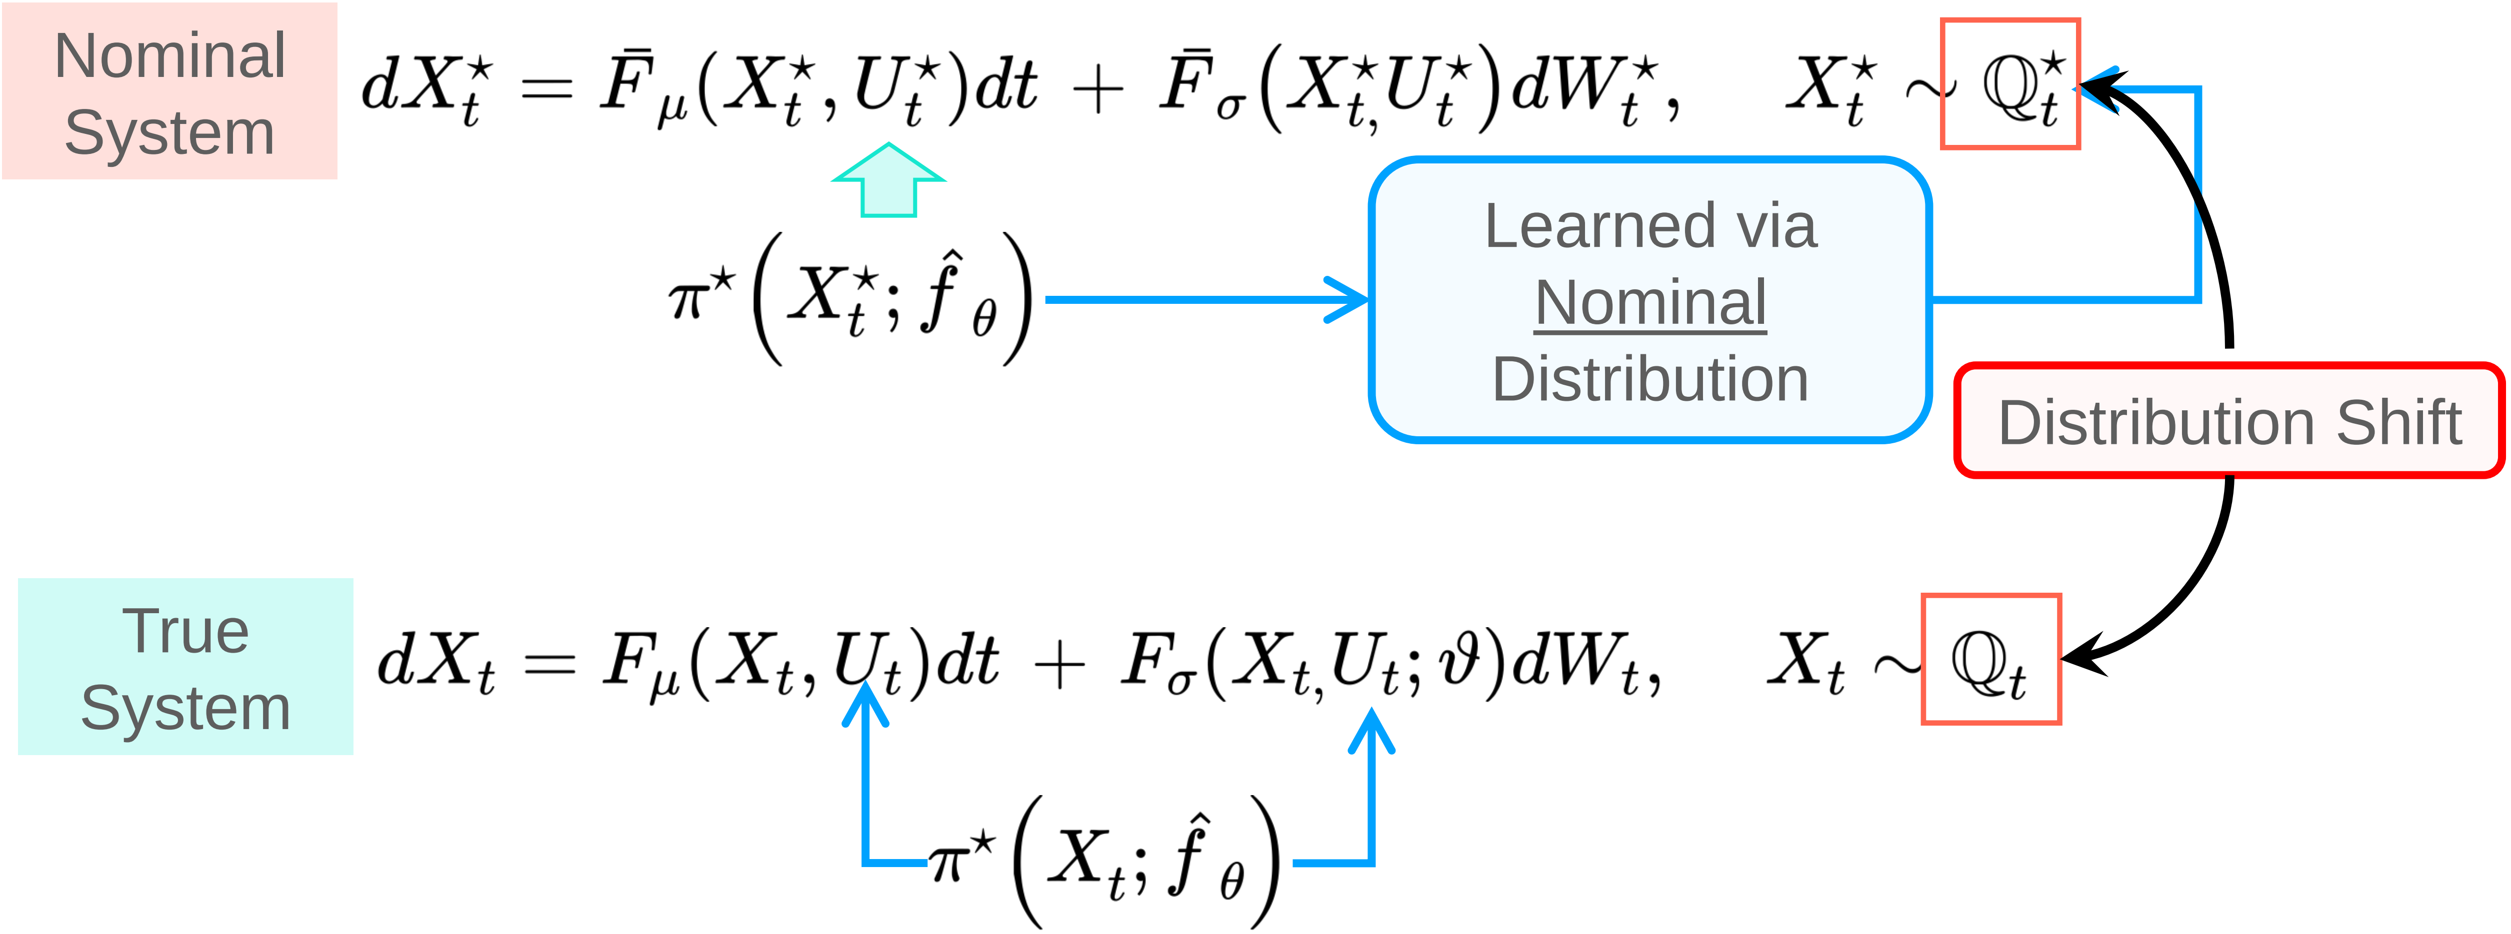

- Learned controller on true system: Distribution shift
  - Guarantees of safety and predictability: Invalid

Goals

Nominal System

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

True System

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

Bounded

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

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

Wasserstein Distance

Ambiguity Set

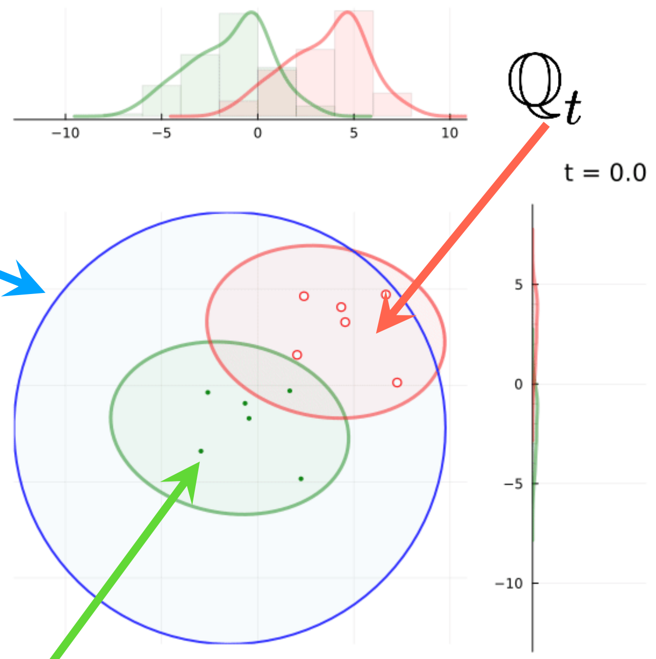

$\mathbb{Q}_t$

$\mathbb{Q}_t^\star$

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

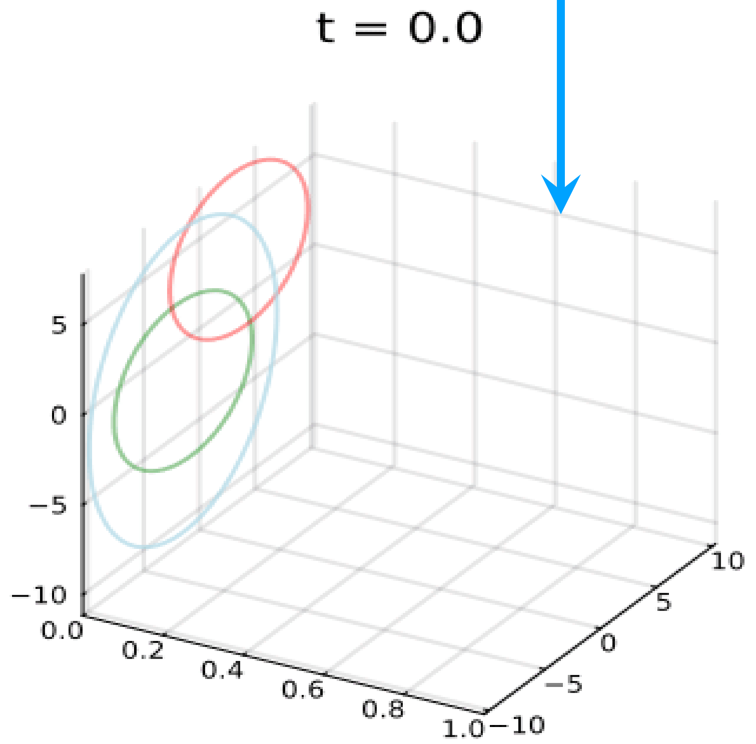

Set of path measures (Girsanov)
Distributions on
$\mathcal{C}([0,T];\mathbb{R}^n)$

Nominal Solution

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

$X_t \sim \mathbb{Q}_t$

True Solution

Controller

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

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

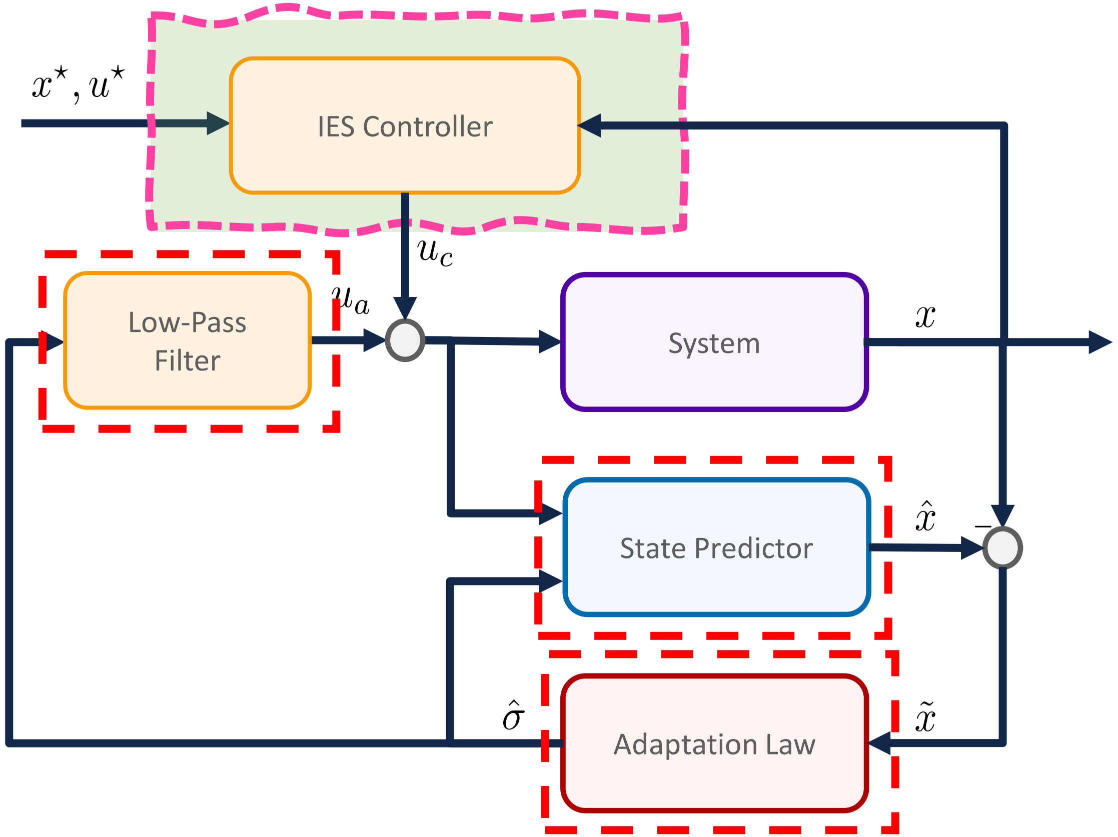

Predictor: System driven by colored noise

Numerical Experimentation

Angular rate dynamics of a quadrotor

Drift and diffusion uncertainties

Divergence of nominal and true distributions

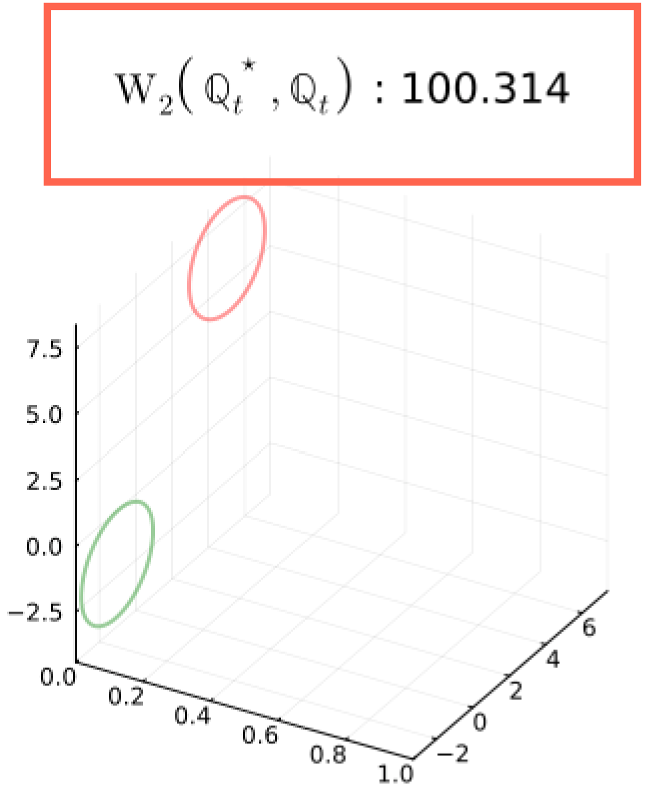

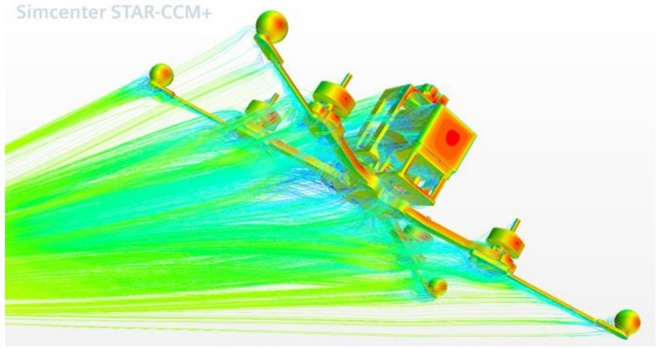

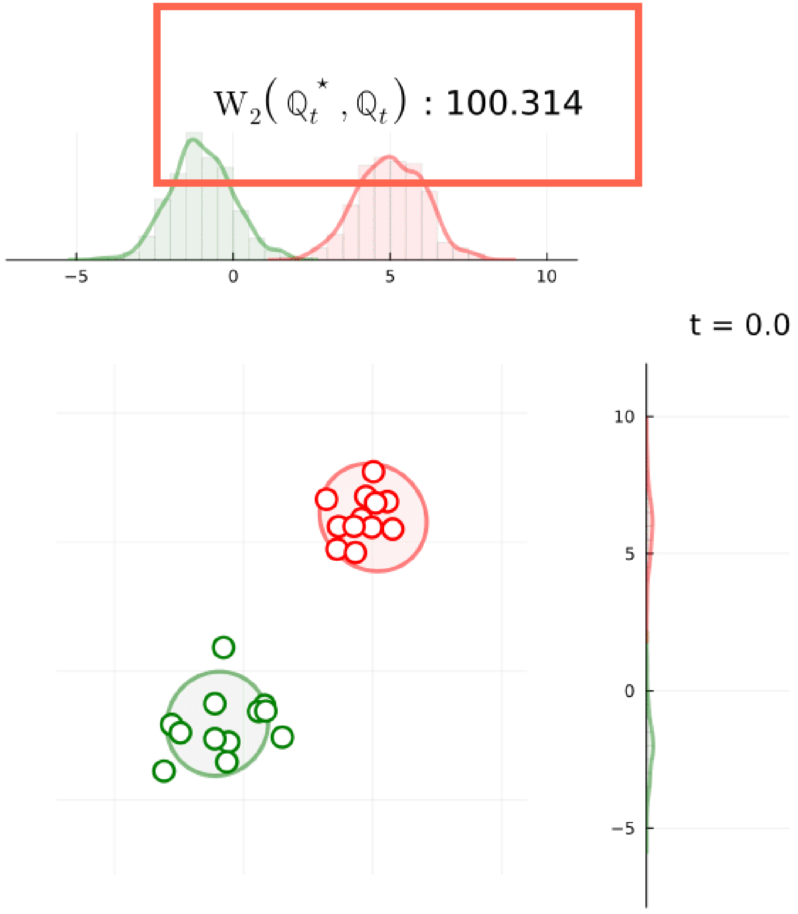

Numerical Experimentation

DRAC control

Independent Brownian motions → Convergence up to a nonzero limit

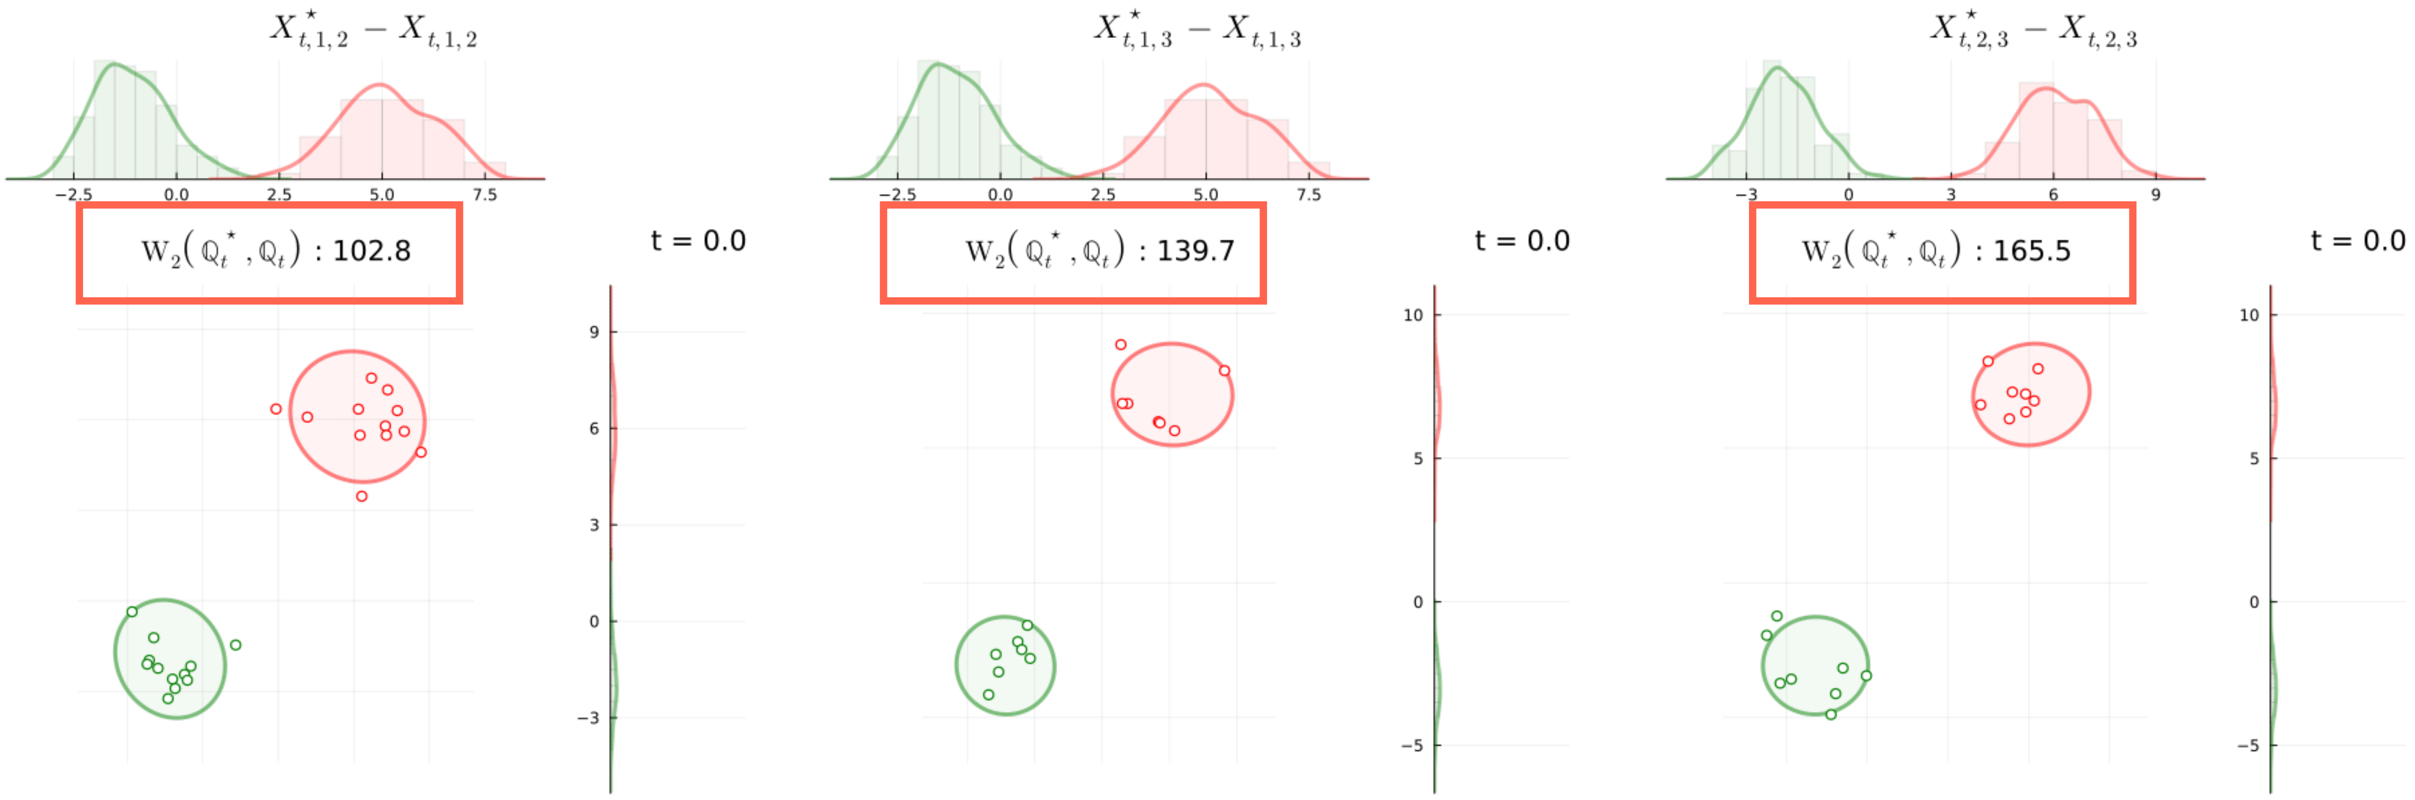

Boundedness and Convergence of distributions in the Wasserstein metric

Numerical Experimentation

DRAC control

Independent Brownian motions → Convergence up to a nonzero limit

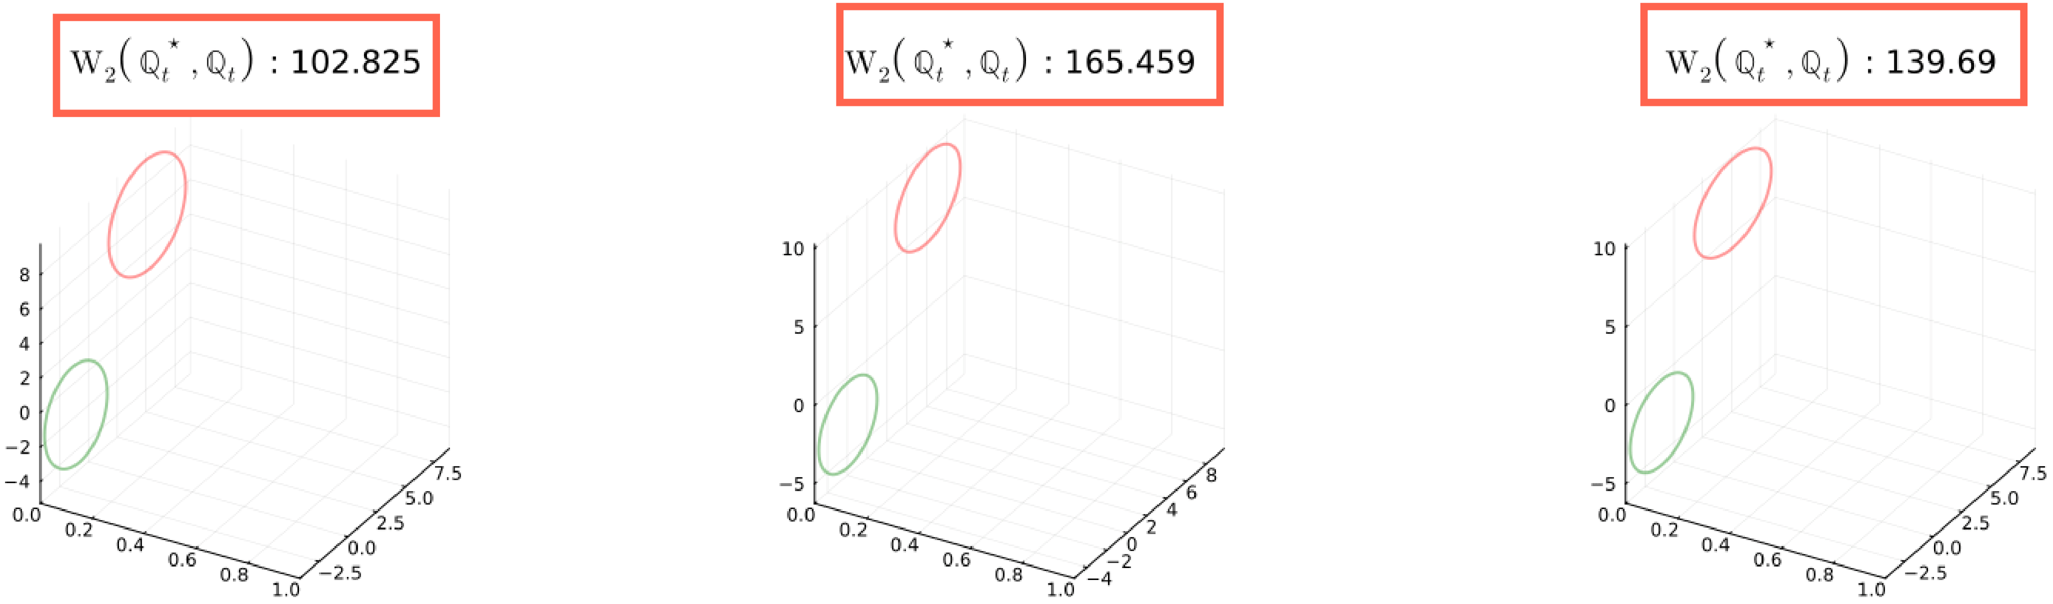

Boundedness and Convergence of distributions in the Wasserstein metric

On-going work

- Further experimentation of DRAC
- 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
- Ongoing: Distributionally robust planning and control