

# Distributionally Robust Adaptive Control (DRAC)

*Robust Adaptive Control* Meets *Probability Measures*

**Naira Hovakimyan** and **Aditya Gahlawat**

Mechanical Science and Engineering  
University of Illinois at Urbana-Champaign

Annual AFOSR Review

**Grant# FA9550-21-1-0411**

August 27, 2024



UNIVERSITY OF  
**ILLINOIS**  
URBANA-CHAMPAIGN

# Complex Dynamics and Environments

## Challenges

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



## Solutions

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

**AlphaPilot simulation  
challenge: camera views**



Massachusetts  
Institute of  
Technology

# Challenges and the Tools

Complex Dynamics

Uncertain Models

Uncertain Environments

## Control theoretic tools

- Structured models
- (Parametric) uncertainties  $\in \mathbb{R}^m$
- Limited representations

## Data-driven ML tools

- General models
- Non-parametric uncertainties  $\in ?$
- Stochastic representations



The diagram illustrates the divide between control theory and machine learning tools. On the left, a large pink arrow points left, labeled "Safety & Robustness". Above this arrow are three blue diamonds connected by a red line, with vertical red lines extending from each diamond. On the right, a large pink arrow points right, labeled "Empirical Performance". Above this arrow is a plot of several overlapping bell curves in blue, orange, and green. A teal bracket connects the two arrows, with the text "Bridging the divide" below it.


**Safety & Robustness**

**Empirical Performance**

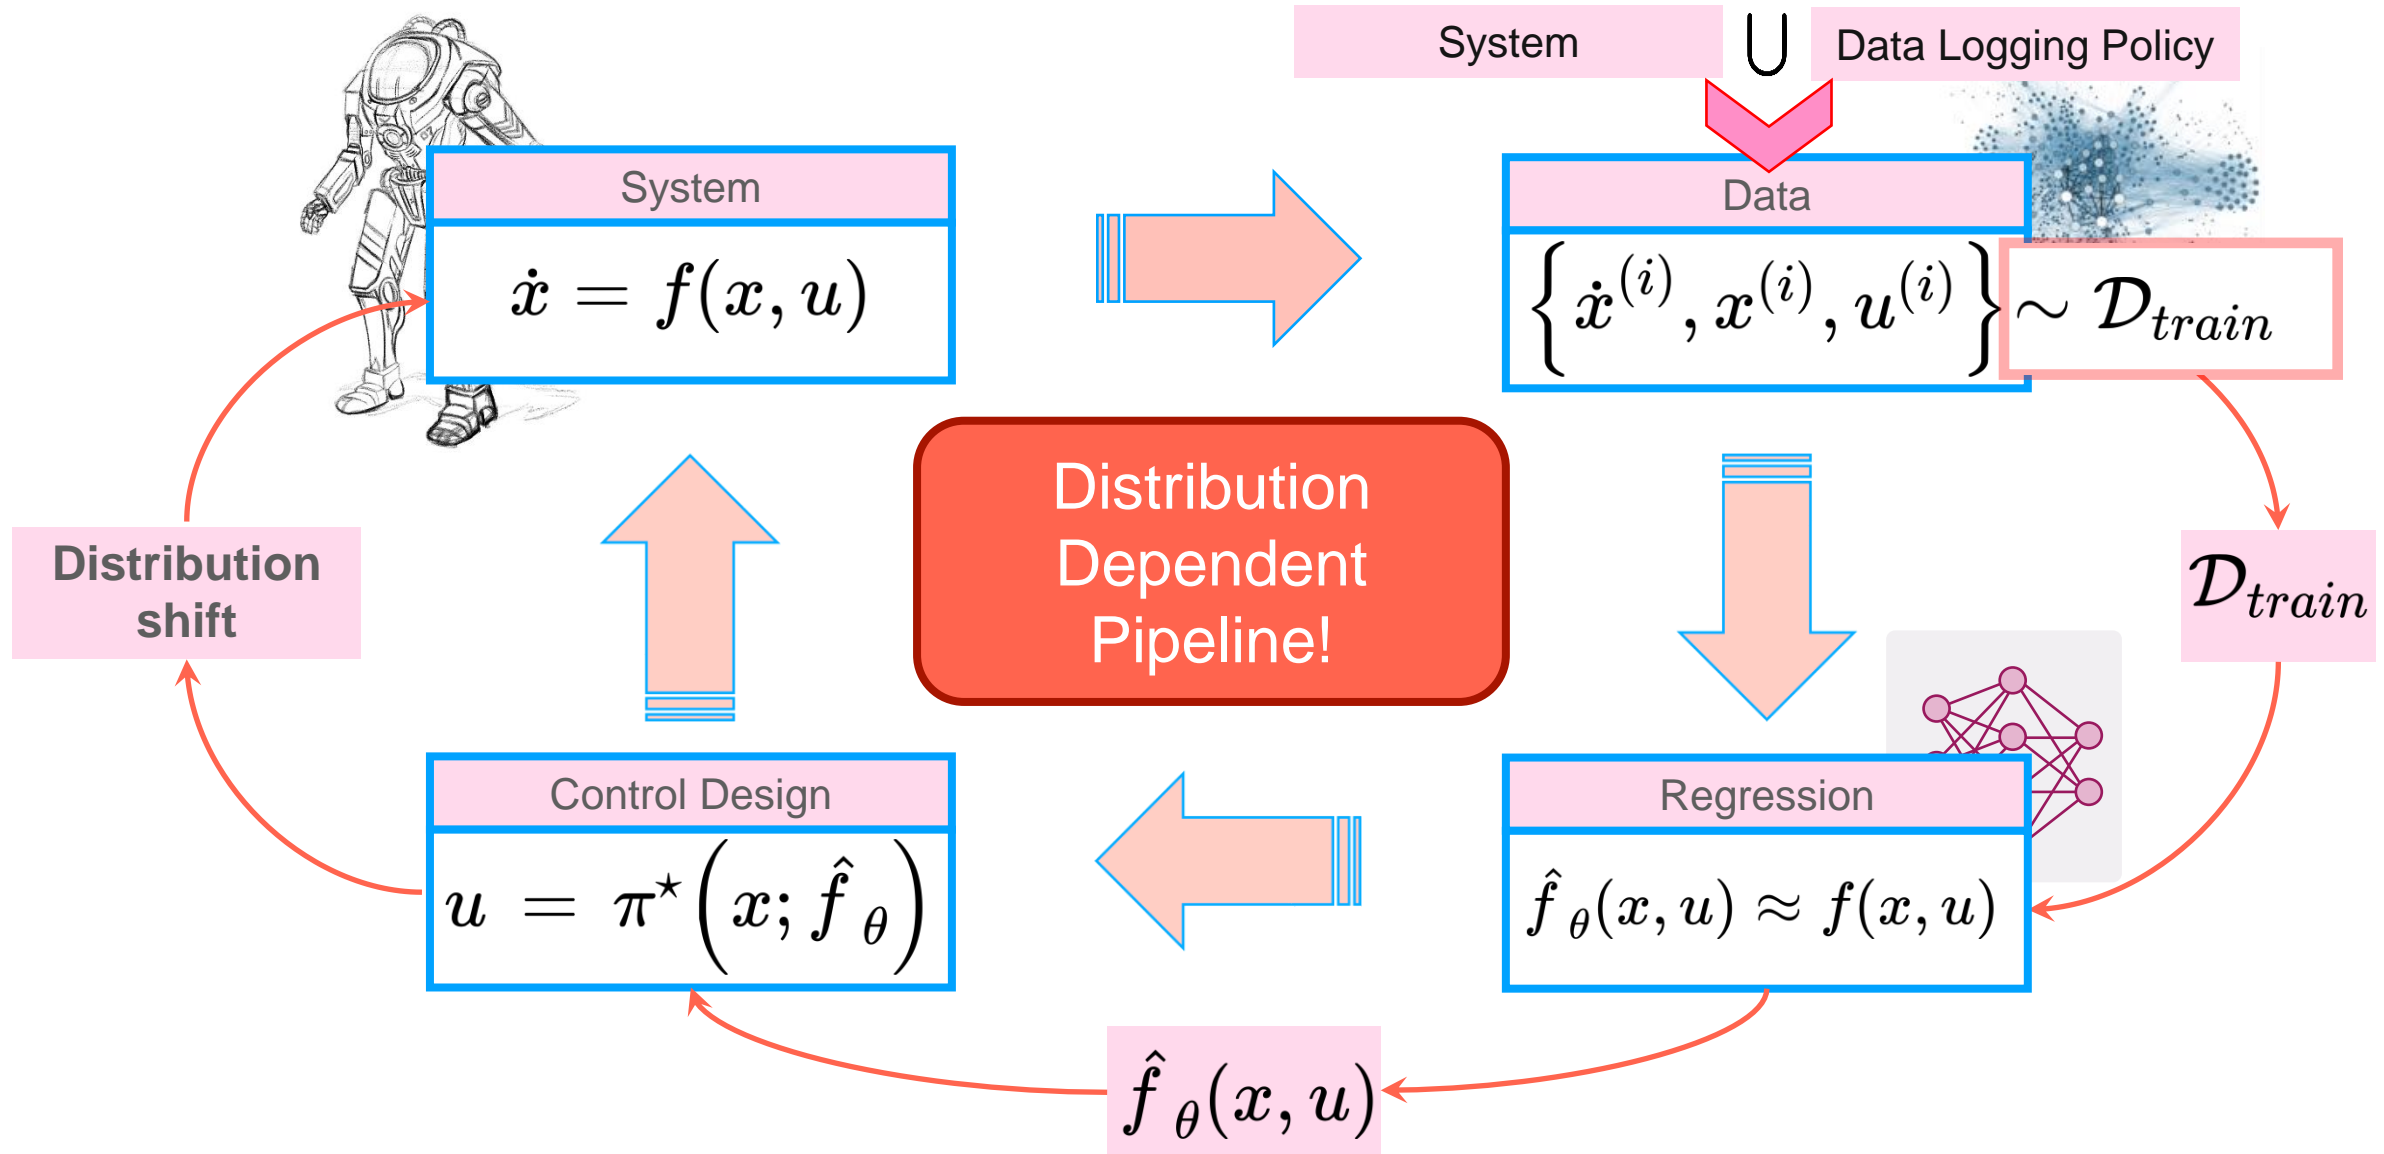
Bridging the divide

# Data-driven Learning & Control

An *ill-posed* consolidation




# Data-driven Control Design



# Machine Learning & Control

- A **general consolidation framework** of ML and Control must contend with a **language barrier**
- **Prototypical ML pipeline**  $\dot{x} = f(x, u)$   $\hat{f}_\theta(x, u) \approx f(x, u)$

Empirical  
Distribution




# Guarantees of Machine Learning?


- What **guarantees** can we expect from the machine learned models?

- **Trained model**  $\hat{f}_{\theta^*} \leftarrow \theta^* = \operatorname{argmin}_{\theta} \mathbb{E}_{(x_+, x, u) \sim \mathcal{D}_{train}} \operatorname{Loss} \left( \hat{f}_{\theta}(x, u), x_+ \right)$

- **Learned model's certificate:**

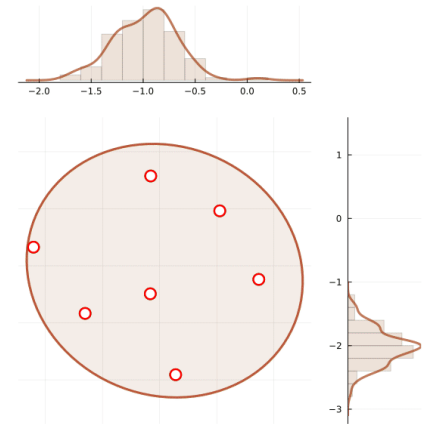

$$\underbrace{\mathcal{R}(\mathcal{D}_{train}, \hat{f}; \theta^*)}_{\text{Model Risk}} = \mathbb{E}_{(x_+, x, u) \sim \mathcal{D}_{train}} \operatorname{Loss} \left( \hat{f}_{\theta^*}(x, u), x_+ \right) \leq \delta$$


- **Distribution dependent** performance guarantees.
- Can be extended to **Wasserstein ambiguity sets** around empirical distributions  $\mathcal{D}_{train}$  [1]



# Control of Probability Measures


- Using what ML can provide, how to ensure **predictable and safe** use in control?
- The guarantees are defined **distributionally**  $\mathcal{R}(\mathcal{D}_{train}, \hat{f}; \theta^*)$ 
  - Not finite-dimensional vector spaces, but **probability measures** on them
- Therefore, we need **control architectures** that satisfy:
  - A priori **uniform guarantees on the space of distributions** (probability measures), and;
  - Applicable to a **general class** of **nonlinear** systems, subject to both **epistemic** and **aleatoric** disturbances.



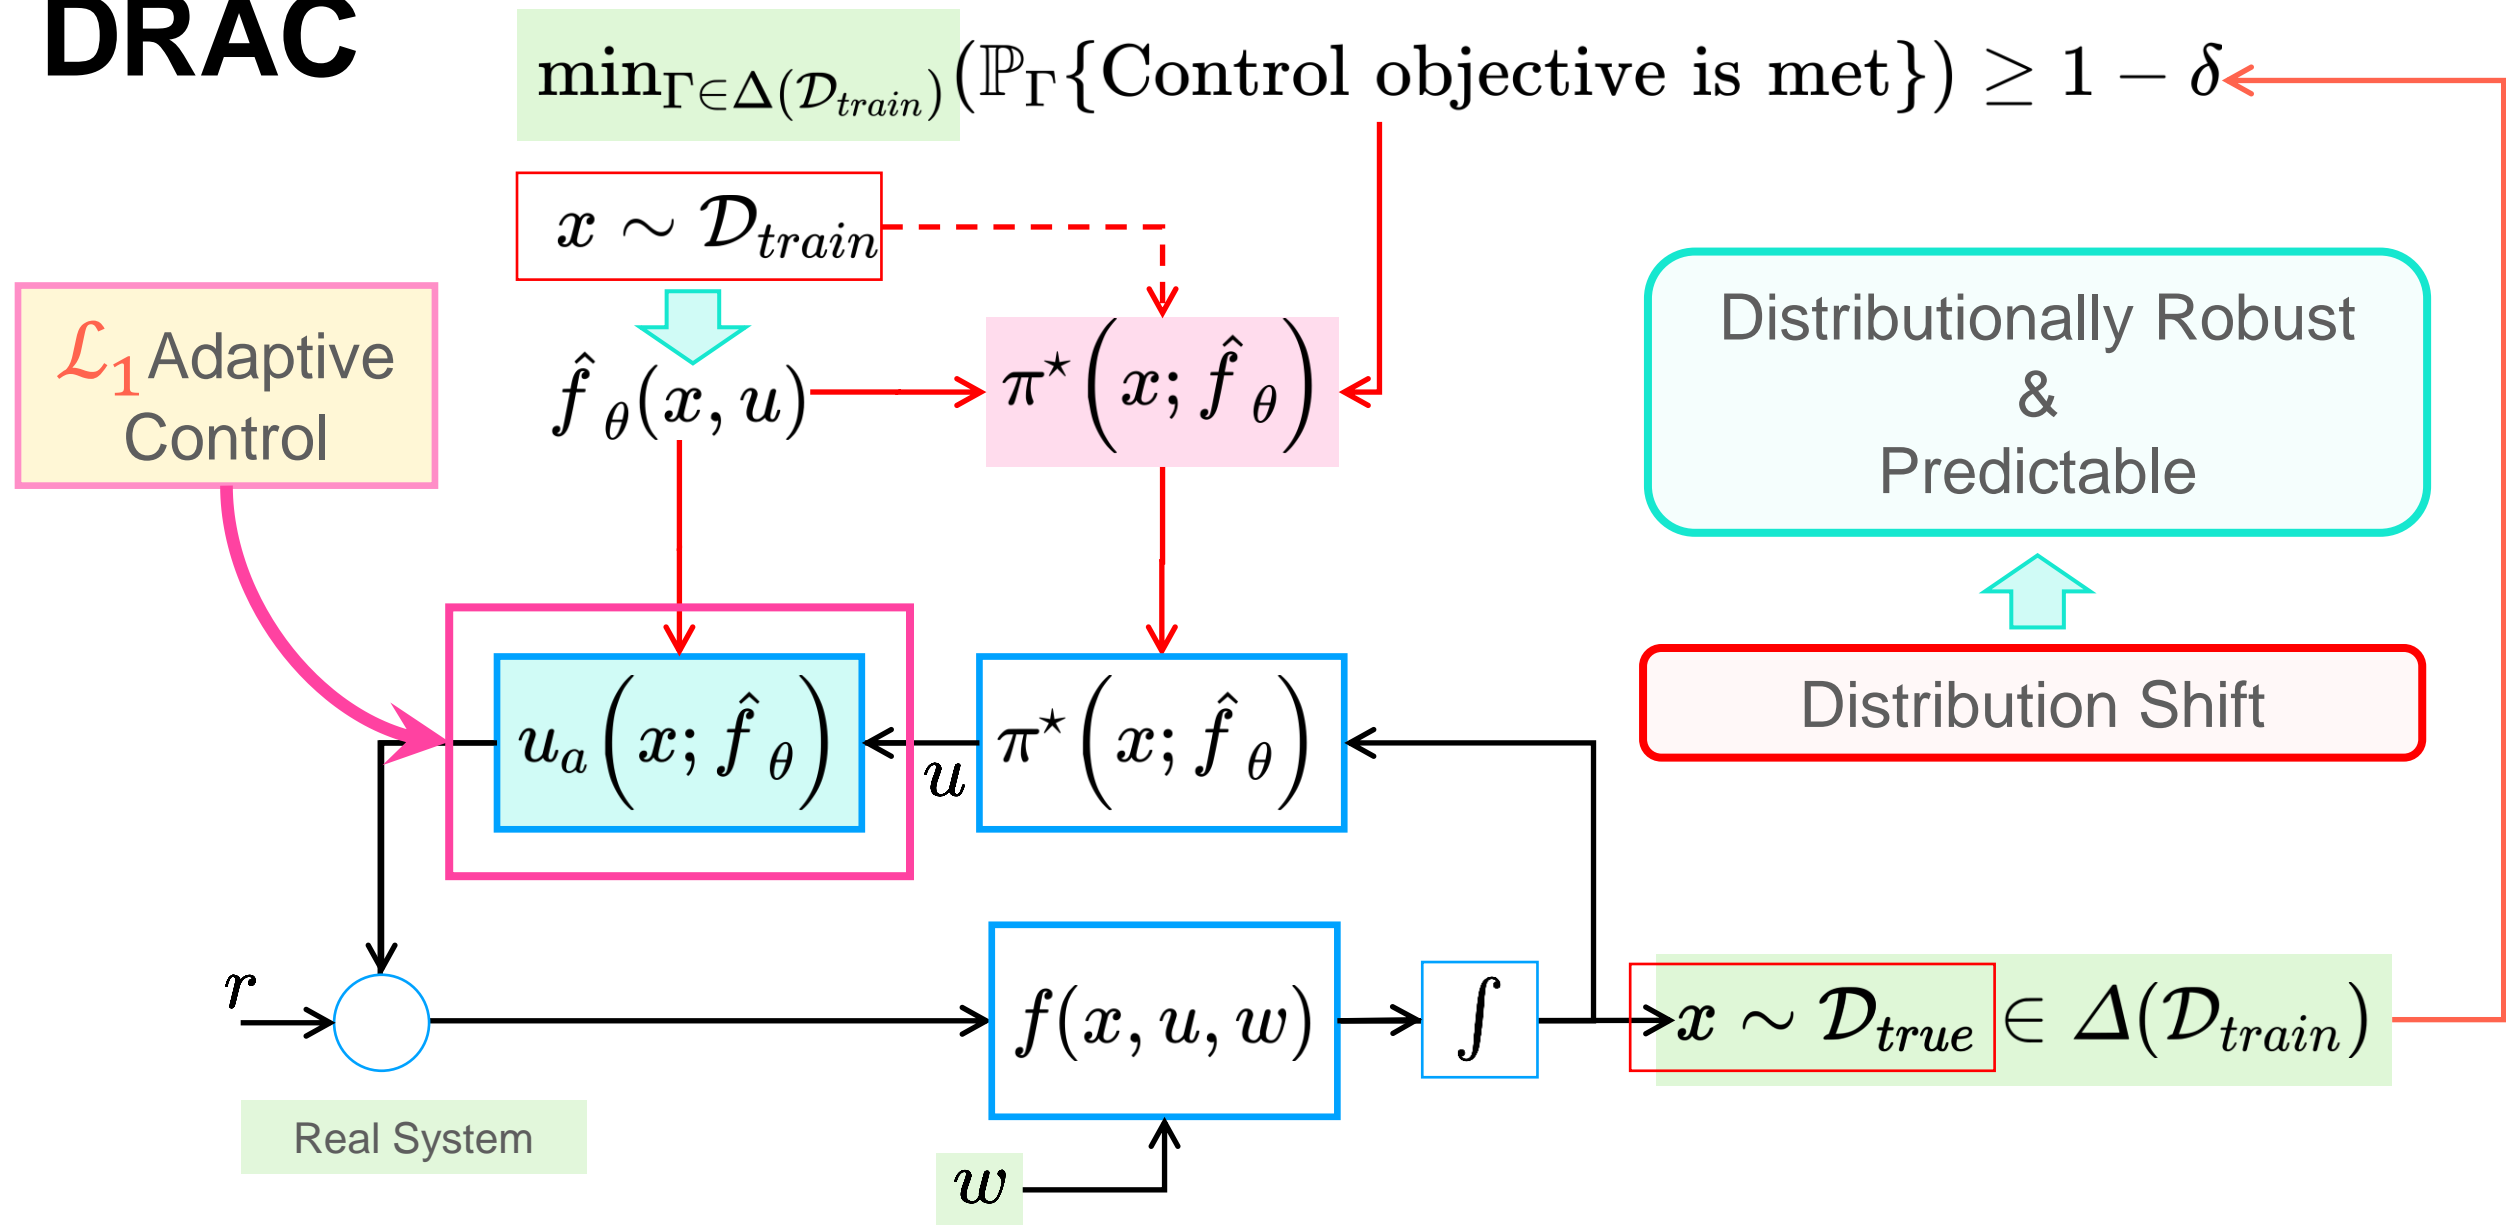


# Distributionally Robust Adaptive Control (DRAC)

*Robust Adaptive Control  
Meets  
Probability Measures*




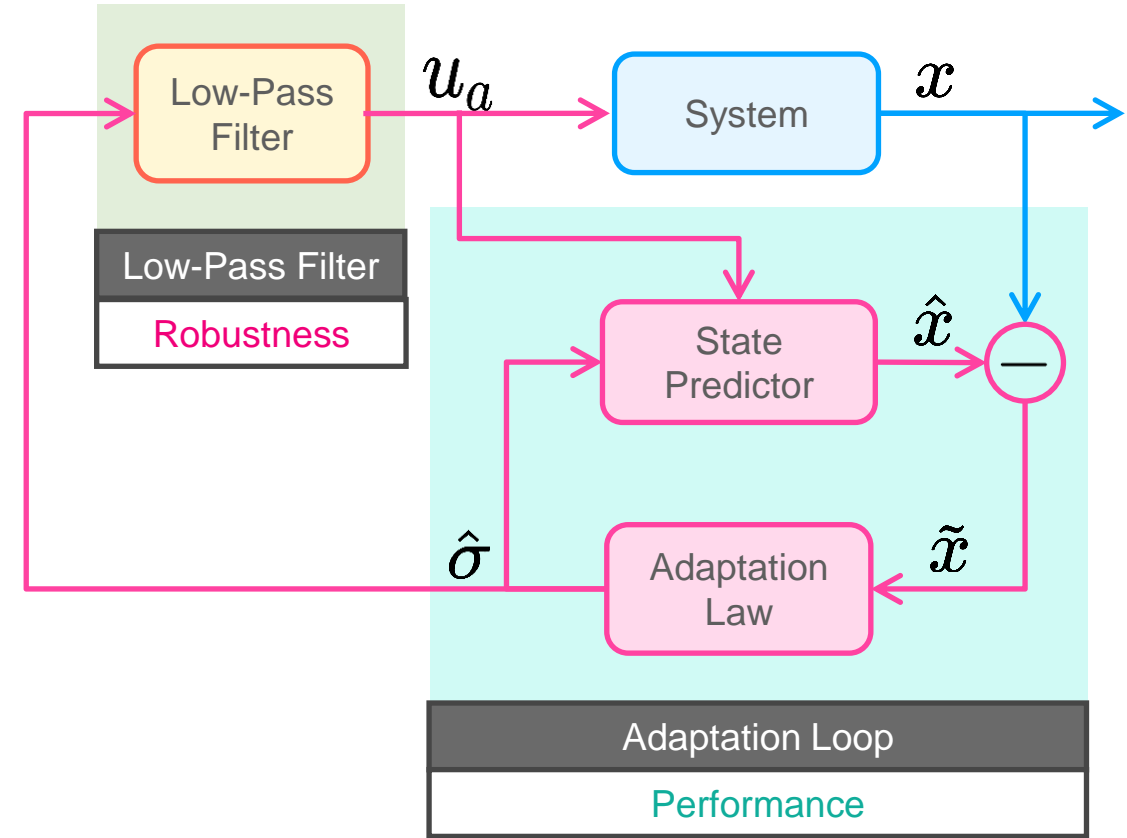
# DRAC



Control augmentation  $u_a$  to guarantee certificates of **distributional robustness**

# $\mathcal{L}_1$ Adaptive Control Architecture

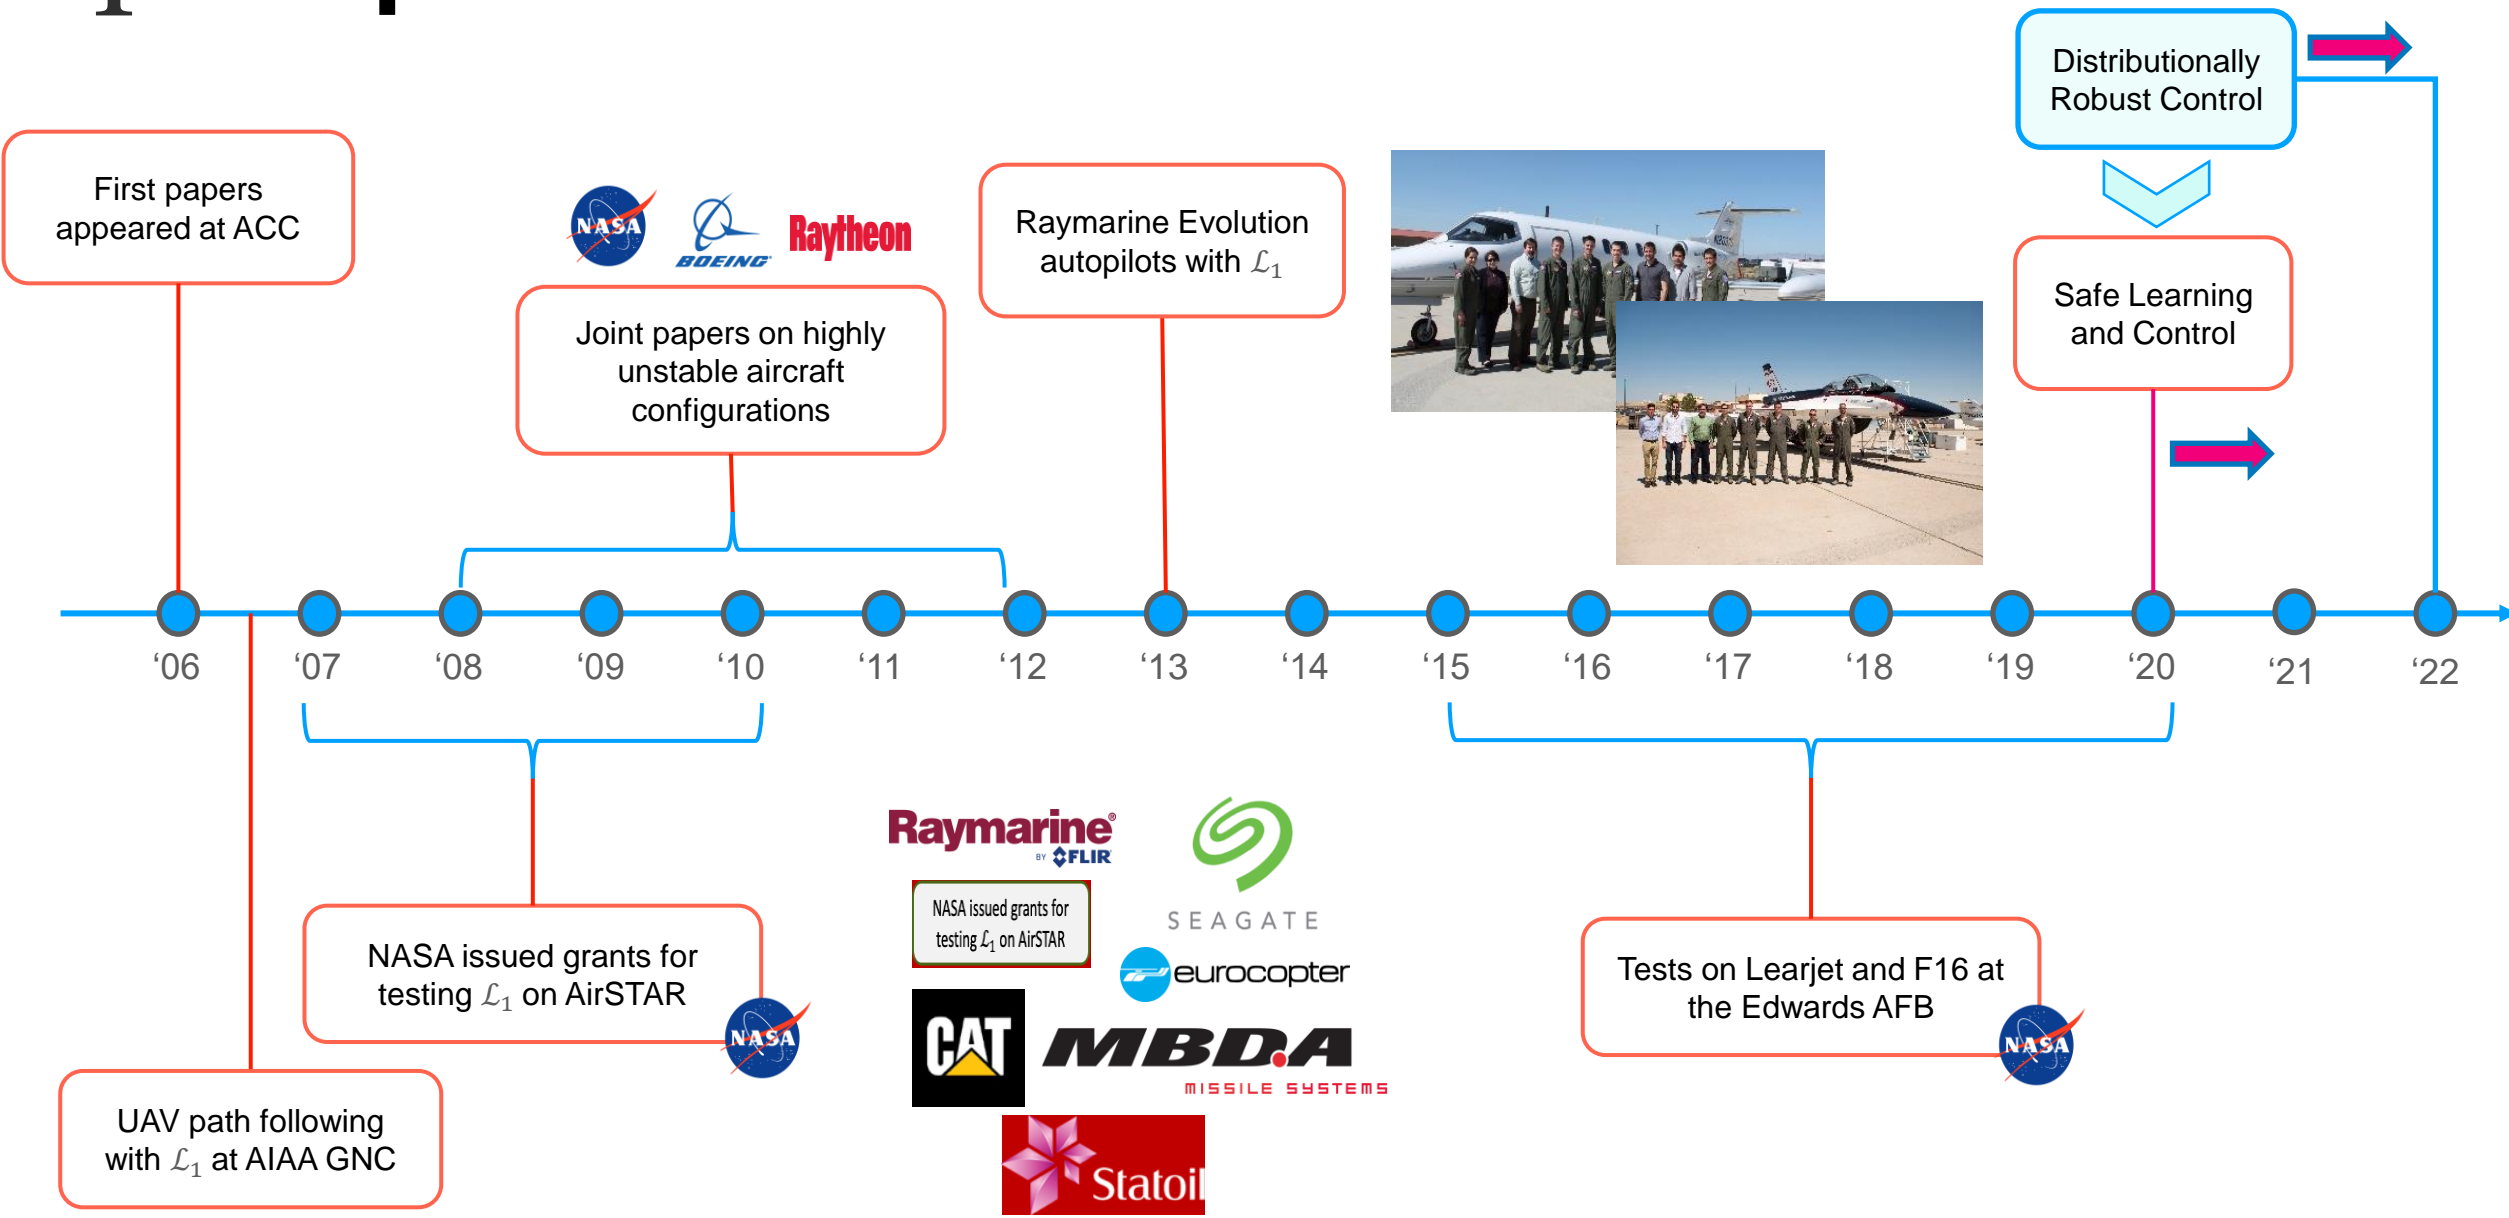
- Guaranteed **uniform performance** bounds and **robustness margins**
- Validated for manned and unmanned aerial vehicles, oil drilling operations, hydraulic pumps, etc.
- Commercialized by various industries, including Raymarine, Caterpillar, JOUAV Automation Tech, etc.



intelinair







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



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

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




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




Data-driven systems,  
Stochastic representations

Distributional  
Guarantees



# The Systems




# Itô Stochastic Differential Equations

True  
System

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

- Evolution description of general class of dynamics
  - Nonlinear and uncertain
- Representation: Sample path-wise (trajectories)
- Representation: Distributions (of trajectories)
- Incomplete knowledge of vector fields  $F_\mu$  and  $F_\sigma$ 
  - Epistemic uncertainties
- Inherent (unlearnable) randomness  $W_t$ 
  - Aleatoric uncertainties




# Nonlinear Itô Processes

True  
System

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

- Strong Solution of a Stochastic Differential Equation (SDE)
  - Nowhere differentiable ← Stochastic Calculus (Itô, Stratanovich, and Doeblin)
- Brownian motion  $W_t$ : Almost surely continuous Lévy process
  - Consequence of Central Limit Theorem/Donsker's Theorem
- Drift  $F_\mu$ : Average (mean) behavior of sample paths
- Diffusion  $F_\sigma$ : Randomness of sample paths via correlated and state multiplicative injection of Brownian motion





# Nonlinear Itô Processes

True  
System

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

Uncertain drift  $F_\mu(X_t, U_t) = \bar{F}_\mu(X_t, U_t) + \Delta F_\mu(X_t, U_t)$

- Known drift component
- Drift uncertainty
  - Matched and unmatched uncertainties
  - Affine in control  $U_t$

Uncertain diffusion  $F_\sigma(X_t, U_t) = \bar{F}_\sigma(X_t, U_t) + \Delta F_\sigma(X_t, U_t)$

- Known diffusion component
- Diffusion uncertainty
  - Matched and unmatched uncertainties
  - Noisy control channel  $U_t dW_t$

Drift only  
=  
Deterministic  
Nonlinear ODEs

# Systems

True System

$$dX_t = \underbrace{F_\mu(X_t, U_t)}_{\bar{F}_\mu(X_t, U_t) + \cancel{\Delta F_\mu(X_t, U_t)}} dt + \underbrace{F_\sigma(X_t, U_t)}_{\bar{F}_\sigma(X_t, U_t) + \cancel{\Delta F_\sigma(X_t, U_t)}} dW_t$$

$\mathcal{D}_{true}$

$X_t \sim \mathbb{Q}_t$

Distribution Shift

Nominal System

$$dX_t^* = \bar{F}_\mu(X_t^*, U_t^*) dt + \bar{F}_\sigma(X_t^*, U_t^*) dW_t^*$$

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

$\mathcal{D}_{train}$

Nominal system: No **epistemic uncertainties**, only **aleatoric**

Independent  
Brownian motions

$W_t$

$W_t^*$

True System  
Probability Measure  
(Distribution)

$\mathbb{Q}_t$

Nominal System  
Probability Measure  
(Distribution)

$\mathbb{Q}_t^*$

# Goals

Nominal  
System

$$dX_t^* = \bar{F}_\mu(X_t^*, U_t^*)dt + \bar{F}_\sigma(X_t^*, U_t^*)dW_t^*,$$

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

Learned via  
Training  
Distribution

$$\pi^*(X_t^*; \hat{f}_\theta)$$

Distribution Shift

True  
System

$$dX_t = F_\mu(X_t, U_t)dt + F_\sigma(X_t, U_t; \vartheta)dW_t,$$

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

$$\pi^*(X_t; \hat{f}_\theta)$$

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

# Goals

Nominal  
System

$$dX_t^* = \bar{F}_\mu(X_t^*, U_t^*)dt + \bar{F}_\sigma(X_t^*, U_t^*)dW_t^*,$$

Any Distribution

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

Bounded

True  
System


$$dX_t = F_\mu(X_t, U_t)dt + F_\sigma(X_t, U_t; \vartheta)dW_t,$$

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

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

- We want to design a **feedback augmentation** such that
- True distribution  $\mathbb{Q}_t$  remains uniformly bounded around the nominal distribution  $\mathbb{Q}_t^*$ 
  - 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)

# Goals: Uniform Bounds

$$\begin{aligned} X_t &\sim \boxed{Q_t} \\ X_t^* &\sim \boxed{Q_t^*} \end{aligned}$$


- We need to define the **uniform bound** between  $Q_t^*$  and  $Q_t$
- For **deterministic** systems, one may interpret true and nominal solutions

as  $x, x^* \in (\mathcal{X}, d)$ ,  $\mathcal{X} = \mathcal{C}([0, T]; \mathbb{R}^n)$ ,  $d = \sup_{[0, T]} \|\cdot - \cdot\|$

- Polish space for  $T < \infty$
- Deterministic  $\mathcal{L}_1$  provides bounds on  $x, x^* \in (\mathcal{X}, d)$  of the form:

$$d(x, x^*) = \sup_{t \in [0, T]} \|x(t) - x^*(t)\| \leq \rho$$

# Goals: Uniform Bounds



- For **stochastic** systems, we may consider  $X_t, X_t^*$  as **random variables** on  $(\mathbb{R}^n, \|\cdot\|)$

$$X_t, X_t^* : [0, T] \rightarrow (\mathbb{R}^n, \|\cdot\|)$$

- Alternatively, we can consider strong solutions  $X_t, X_t^*$  as **random variables** on  $(\mathcal{X}, d)$

$$X_{[0,T]}, X_{[0,T]}^* : \boxed{\Omega} \rightarrow (\mathcal{X}, d)$$

Sample  
Space



# Goals: Uniform Bounds

$$X_{[0,T]}^* \sim \boxed{\mathbb{Q}_{[0,T]}^*} \quad X_{[0,T]} \sim \boxed{\mathbb{Q}_{[0,T]}}$$

Laws: Probability Measures on



$$\mathcal{B}(\mathcal{X}, d) = \underbrace{\mathcal{B}}_{\substack{\text{Borel} \\ \sigma\text{-Algebra}}} \left( \mathcal{C}([0, T]; \mathbb{R}^n), \sup_{[0,T]} \|\cdot - \cdot\| \right)$$

$$\forall B, B^* \in \mathcal{B}(\mathcal{X}, d)$$

$$\mathbb{Q}_{[0,T]}^*(B^*) = \Pr \{ X_{[0,T]}^* \in B^* \}$$

$$\mathbb{Q}_{[0,T]}(B) = \Pr \{ X_{[0,T]} \in B \}$$

$$X_{[0,T]}, X_{[0,T]}^* : \Omega \rightarrow (\mathcal{X}, d)$$




# Goals: Uniform Bounds

- Laws: Probability Measures on  $\mathcal{B}(\mathcal{X}, d)$
- Exploiting two facts for an alternative representation

$$X_{[0,T]}^* \sim \mathbb{Q}_{[0,T]}^*$$


$$X_{[0,T]} \sim \mathbb{Q}_{[0,T]}$$

Fact 1 Finite-dimensional Cylinder Sets [1]



$$\mathcal{B}\left(\mathcal{C}([0, T]; \mathbb{R}^n), \sup_{[0, T]} \|\cdot - \cdot\|\right)$$

=



$$B_i \in \mathcal{B}(\mathbb{R}^n), i \in \{1, \dots, k\}$$

=

$$\sigma\text{-alg. : Borel sets} \in \underbrace{\mathbb{R}^n \times \dots \times \mathbb{R}^n}_{k \in \mathbb{N}}$$

$$0 < t_1 < t_2 < \dots < t_k \leq T$$

$$\forall k \in \mathbb{N}$$



# Goals: Uniform Bounds

- Laws: Probability Measures on  $\mathcal{B}(\mathcal{X}, d)$
- Exploiting two facts for an alternative representation

Fact 2

Kolmogorov Extension [1]

$$X_{[0,T]}^* \sim \mathbb{Q}_{[0,T]}^*$$

$$X_{[0,T]} \sim \mathbb{Q}_{[0,T]}$$

Cylinder Set:  $C_k = \{B_{\{1 \dots k\}} \in \mathbb{R}^n \times \dots \times \mathbb{R}^n, 0 < T_1 < \dots < T_k \leq T\}$

Finite-dimensional distributions (FDD):


$$\mathbb{Q}_{T_1 \dots T_k}(C_k) = \Pr\{X_{T_1} \in B_1, \dots, X_{T_k} \in B_k\}$$

Kolmogorov Extension Theorem:

$$\mathbb{Q}_{[0,T]}(C_k) = \mathbb{Q}_{T_1 \dots T_k}(C_k)$$

- Cylinder sets are a determining class [2] for the laws  $\mathbb{Q}_{[0,T]}^*$  and  $\mathbb{Q}_{[0,T]}$

$$\mathbb{Q}_{[0,T]}^* = \mathbb{Q}_{[0,T]} \text{ on cylinder sets} \Rightarrow \mathbb{Q}_{[0,T]}^* = \mathbb{Q}_{[0,T]} \text{ on } \mathcal{B}(\mathcal{X}, d)$$




[1] Oksendal, B. "Stochastic Differential Equations: An Introduction with Applications", Springer, 2003.

[2] Dudley, R. M. "Real Analysis and Probability", Cambridge University Press, 2002.


# Goals: Uniform Bounds

- The laws  $\mathbb{Q}_{[0,T]}^*$  and  $\mathbb{Q}_{[0,T]}$ : Probability measures on  $\mathcal{B} \left( \mathcal{C}([0, T]; \mathbb{R}^n), \sup_{[0,T]} \|\cdot - \cdot\| \right)$
- Consequences of finite-dimensional cylinder sets and the Kolmogorov extension theorem:
  - The 'distance' between  $\mathbb{Q}_{[0,T]}^*$  and  $\mathbb{Q}_{[0,T]}$  determined by their actions on finite-dimensional cylinder sets

$$\mathbb{Q}_{[0,T]}(C_k) = \Pr \{X_{T_i} \in B_i\}$$




$$\mathbb{Q}_{[0,T]}^*(C_k) = \Pr \{X_{T_i}^* \in B_i\}$$



# Goals: Uniform Bounds

- The ‘distance’ between the laws  $\mathbb{Q}_{[0,T]}$  and  $\mathbb{Q}_{[0,T]}^*$



- Comparator:** Action of the finite-dimensional measures  $\mathbb{Q}_{T_1 \dots T_k}$  and  $\mathbb{Q}_{T_1 \dots T_k}^*$  on cylinder sets  $C_k = \{B_{\{1 \dots k\}} \in \mathcal{B}(\mathbb{R}^n) \times \dots \times \mathcal{B}(\mathbb{R}^n), 0 < T_1 < \dots < T_k \leq T\}$

$p^{th}$  order Wasserstein metric:

$$\mathbb{W}_p^k(\mathbb{Q}_{T_1 \dots T_k}, \mathbb{Q}_{T_1 \dots T_k}^*) = \left( \inf_{\pi \in \Pi(\mathbb{Q}, \mathbb{Q}^*)} \int_{\mathbb{R}^{nk}} \|x - y\|_{\mathbb{R}^{nk}}^p \pi(dx, dy) \right)^{\frac{1}{p}}$$

- Minimization over all couplings between  $\mathbb{Q}_{T_1 \dots T_k}^*$  and  $\mathbb{Q}_{T_1 \dots T_k}$  [1]
- Earth mover’s distance:** Least ‘energy’ to move pile  $\mathbb{Q}_{T_1 \dots T_k}^*$  to pile  $\mathbb{Q}_{T_1 \dots T_k}$
- Metric** on the **space of probability measures** with finite moment of order  $p$ .
- No** requirement on **absolute continuity** of measures

# Goals: Uniform Bounds

- Design **DRAC** augmentation  $u_a$  to guarantee certificates of **distributional robustness** of the form:


*Uniform  
Bound*

$$\mathfrak{D}_p \left( \mathbb{Q}_{[0,T]}, \mathbb{Q}_{[0,T]}^* \right) \doteq \sup_{k \in \mathbb{N}} \mathbb{W}_p^k \left( \mathbb{Q}_{T_1 \dots T_k}, \mathbb{Q}_{T_1 \dots T_k}^* \right) \leq \rho \in \mathbb{R}_{>0}$$


*Known a priori*

$$\Upsilon \left( \rho, \mathbb{Q}_{[0,T]}^* \right) \doteq \left\{ \text{Prob. Meas. } \xi \text{ on } \mathcal{C}([0, T]; \mathbb{R}^n) : \mathfrak{D}_p \left( \xi, \mathbb{Q}_{[0,T]}^* \right) \leq \rho \right\}$$

*Ambiguity Tube*



*Ambiguity Sets*



# Nominal Stability Assumptions

Nominal System

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




Nominal Deterministic Sub-system

$$\dot{x}^* = \bar{F}_\mu(x^*, u^*) \quad \pi^*(x^*; \hat{f}_\theta)$$

Ordinary Differential Equation

$$x^*(0) \sim \mathbb{Q}_0^*$$



The nominal deterministic subsystem with the nominal controller

- Incrementally exponentially stable (IES)

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

Furthermore,  $\exists$  positive scalars  $\alpha_1$ ,  $\alpha_2$ , and  $\lambda$ , and function  $V : \mathbb{R}^n \times \mathbb{R}^n \rightarrow \mathbb{R}$

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

Incremental Lyapunov function (ILF)

Certificate for IES [1]

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

# Controller

Nominal  
System

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

True  
System

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

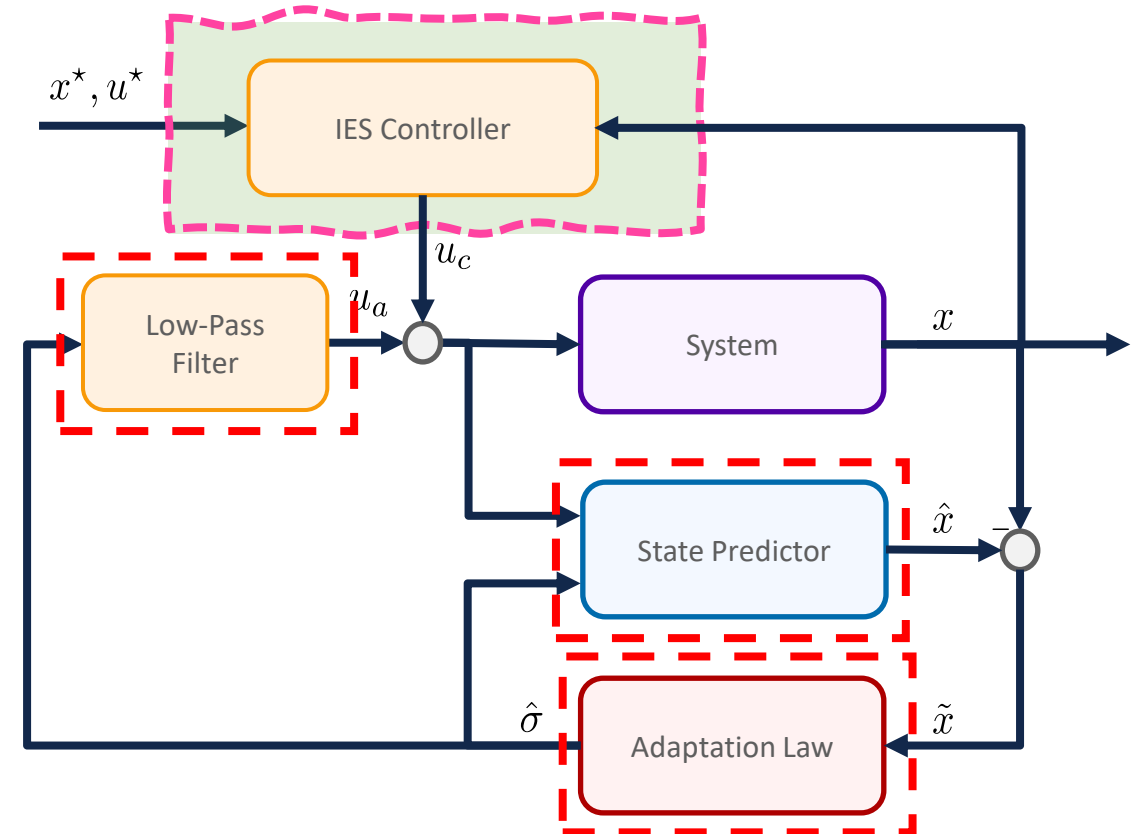
The controller has the architecture of an  $\mathcal{L}_1$  adaptive controller

The controller has three main components

State Predictor

Adaptation Law

Low-Pass Filter





# Controller

True System

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

## $\mathcal{L}_1$ DRAC Controller


- Retains identical architecture of  $\mathcal{L}_1$  for **deterministic** nonlinear systems
  - Standard** implementation
- Stochasticity **inherited** through feedback
  - Colored noise feedback** injection:  $X_t$  driven by  $W_t$



**Implementation:** Feedback space (standard)



**Guarantees:** Space of probability measures (distributions) on state space



Does not need/assume **unrealistic** distributional feedback

# Controller

True System

$$\pi^*(X_t; \hat{f}_\theta) \xrightarrow{U_t^*} \pi_a(X_t; \hat{f}_\theta) \xrightarrow{U_{a,t}} dX_t = F_\mu(X_t, U_t)dt + F_\sigma(X_t, U_t; \vartheta)dW_t, \quad X_t \sim \mathbb{Q}_t$$

State Predictor

Adaptation Laws

Control Law

$$d\hat{X}_t = \left[ -\lambda_s \left( \hat{X}_t - X_t \right) + \bar{F}_\mu(X_t, U_t^* + U_{a,t}) + \hat{\Sigma}_t \right] dt \quad \hat{X}_t \sim \hat{\mathbb{Q}}_t$$

- Components:
  - Known deterministic subsystem
  - State-feedback injection
  - Adaptive estimates
- Deterministic dynamics ← Exogenous colored noise input  $X_t$
- No access to the driving Brownian motion  $W_t$ 
  - Best estimate is zero:  $W_t - W_s \sim \mathcal{N}(0, t - s), 0 \leq s \leq t$



# Controller

True System

$$\pi^*(X_t; \hat{f}_\theta) \xrightarrow{U_t^*} \pi_a(X_t; \hat{f}_\theta) \xrightarrow{U_{a,t}} dX_t = F_\mu(X_t, U_t)dt + F_\sigma(X_t, U_t; \vartheta)dW_t, \quad X_t \sim \mathbb{Q}_t$$

State Predictor

Adaptation Laws

Control Law

$$\hat{X}_t - X_t \doteq \tilde{X}_t = \tilde{X}_0 - \lambda_s \int_0^t \tilde{X}_\nu d\nu + \underbrace{\int_0^t \hat{\Sigma}_\nu d\nu - \int_0^t (\Delta F_\mu d\nu + F_\sigma dW_\nu)}_{\text{Estimation Error}}$$

Learning Signal

Estimation Error

$$\approx \underbrace{\int_0^t (\hat{\Sigma}_\nu - \Delta F_\mu) d\nu}_{\text{Mean Error}} + \underbrace{F_\sigma \sqrt{t}}_{\text{Error Std. Dev.}} \mathcal{N}(0, 1)$$

- **Learning signal:** Stable LTI system
  - Driven by stochastic estimation error
- **Estimation error:** Drift uncertainty + **complete diffusion** vector field
  - Cannot disambiguate between epistemic and aleatoric uncertainties

**Quadratic Variation**  $dW \approx (dt)^{1/2}$

# Controller

True System

$$\pi^*(X_t; \hat{f}_\theta) \xrightarrow{U_t^*} \pi_a(X_t; \hat{f}_\theta) \xrightarrow{U_{a,t}} dX_t = F_\mu(X_t, U_t)dt + F_\sigma(X_t, U_t; \vartheta)dW_t, \quad X_t \sim \mathbb{Q}_t$$

State Predictor

Adaptation Laws

Control Law

$$\tilde{X}_t = \tilde{X}_{iT_s} - \lambda_s \int_{iT_s}^{(i+1)T_s} \tilde{X}_\nu d\nu + \int_{iT_s}^{(i+1)T_s} \hat{\Sigma}_\nu d\nu - \underbrace{\int_{iT_s}^{(i+1)T_s} (\Delta F_\mu d\nu + F_\sigma dW_\nu)}_{\text{Estimation Error}}$$

- Piecewise constant Learning signal

- Sampling period  $T_s$

- Estimation error:  $\approx \int_{iT_s}^{(i+1)T_s} (\hat{\Sigma}_\nu - \Delta F_\mu) d\nu + F_\sigma \sqrt{T_s} \mathcal{N}(0, 1)$

Mean Error

Error Std. Dev.

$\propto T_s$

$\propto \sqrt{T_s}$

= Deterministic  $\mathcal{L}_1$

# Controller

True System

$$\pi^*(X_t; \hat{f}_\theta) \xrightarrow{U_t^*} \boxed{\pi_a(X_t; \hat{f}_\theta)} \xrightarrow{U_{a,t}} dX_t = F_\mu(X_t, U_t)dt + F_\sigma(X_t, U_t; \vartheta)dW_t, \quad X_t \sim \mathbb{Q}_t$$

State Predictor

Adaptation Laws

Control Law

$$\hat{\Sigma}_t \doteq \hat{\Sigma}_{T_s} = \Phi(\lambda_s, T_s) \tilde{X}_{iT_s}, \quad (t, i) \in [iT_s, (i+1)T_s) \times \mathbb{N}$$

- Piecewise constant adaptation law
  - Quality of (moments of ) adaptive estimates  $\uparrow$  as  $\sqrt{T_s} \downarrow$  ( $dW \approx (dt)^{1/2}$ )
- Quantity  $\Phi$  (partially) computed before runtime
  - Minimal computation to produce adaptive estimates (linear operations)
- Numerically stable implementation  $\rightarrow$  avoids stiffness

# Controller

True System

$$\pi^*(X_t; \hat{f}_\theta) \xrightarrow{U_t^*} \boxed{\pi_a(X_t; \hat{f}_\theta)} \xrightarrow{U_{a,t}} dX_t = F_\mu(X_t, U_t)dt + F_\sigma(X_t, U_t; \vartheta)dW_t, \quad X_t \sim \mathbb{Q}_t$$

State Predictor

Adaptation Laws

Control Law

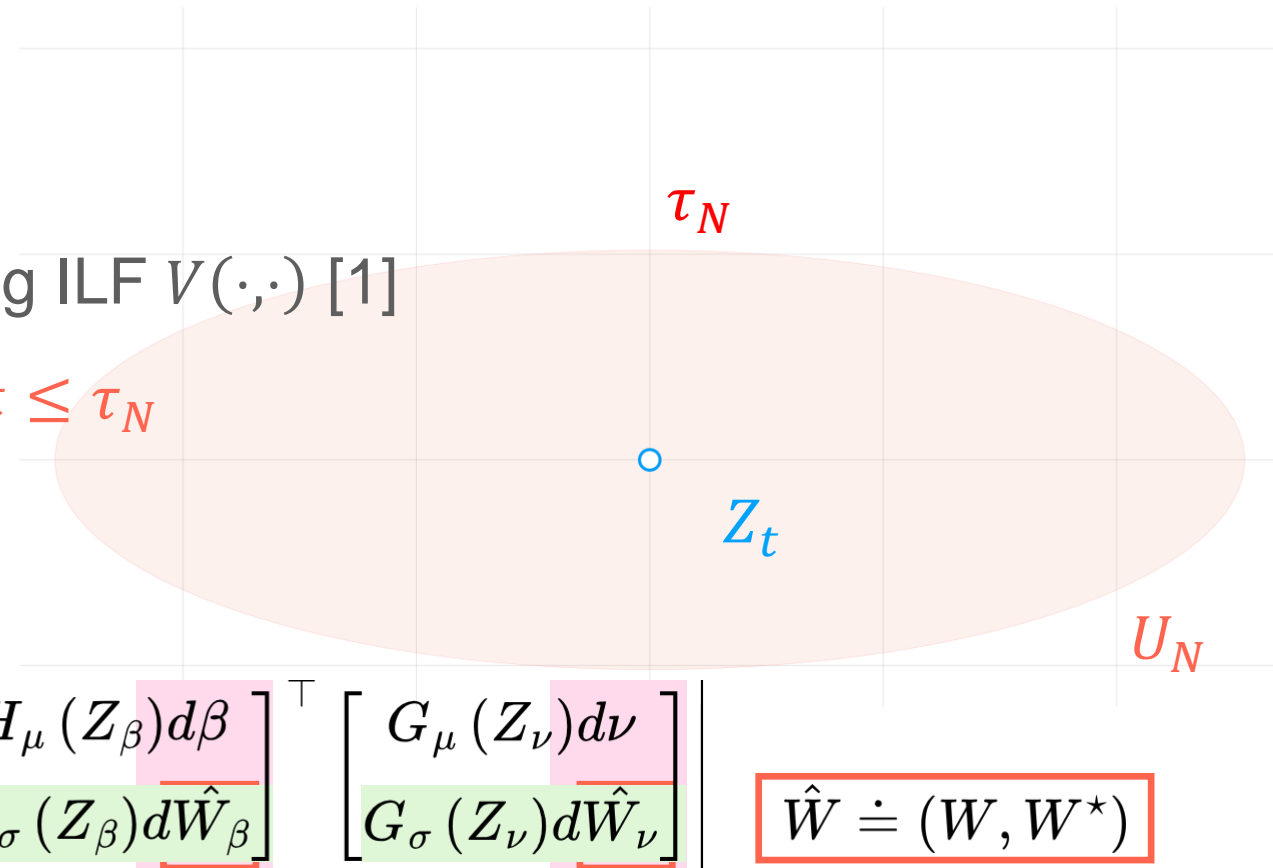
$$U_{a,t} = -\omega \int_0^t e^{-\omega(t-\nu)} \hat{\Sigma}_\nu^\parallel d\nu$$

- **Feedback** operator
  - Absolutely continuous for adaptive estimates driven by strong solutions
- First order low-pass filter with bandwidth  $\omega$
- $\hat{\Sigma}_t^\parallel$ : matched component of the adaptive estimate  $\hat{\Sigma}_t$

# Main Results

## Analysis Sketch


- Khasminskii condition-type analysis using ILF  $V(\cdot, \cdot)$  [1]
- Analysis of joint process  $Z_t \doteq (X_t, X_t^*), \forall t \leq \tau_N$ 
  - $\tau_N =$  **First-exit time** of  $Z_t$  from  $U_N$
- **Goal:**  $\mathbb{E}_{z_0}[\|X_t - X_t^*\|^p]^{\frac{1}{p}} \leq \rho, \forall t \leq \tau_N$
- Bounds on moments of 
$$\sup_{t \in [0, \tau_N]} \left| \int_0^t \int_0^\nu \begin{bmatrix} H_\mu(Z_\beta) d\beta \\ H_\sigma(Z_\beta) d\hat{W}_\beta \end{bmatrix}^\top \begin{bmatrix} G_\mu(Z_\nu) d\nu \\ G_\sigma(Z_\nu) d\hat{W}_\nu \end{bmatrix} \right|$$
- **Nested** Lebesgue-Ito integrals
- Composition of **local martingales** and finite-variation processes
- **Novel bounds** due to:
  - Burkholder-Davis-Gundy inequality
  - $L_p$ -estimates of (semi)-martingales



# Main Results

## Analysis Sketch

$$\lim_{N \rightarrow \infty}$$



- Uniqueness and existence of **strong solutions** over  $[0, T]$ ,  $\forall T < \infty$
- Extension via Fatou's lemma

$$\mathbb{E}_{z_0} \left[ \|X_{t^*} - X_{t^*}^*\|^p \right]^{\frac{1}{p}} \leq \rho, \quad \forall t^* \in [0, T]$$

- $t^* \in [0, T]$  allowed to be **random**  $\leftarrow$  **strongly Markov** over stopping times  $\tau_N$
- Chapman-Kolmogorov to formulate the **joint measure** over cylinder sets
  - **Initial** distribution **coupling**
  - **Transition probability**
- Probability measure + conditional expectation  $\Rightarrow$  **Wasserstein metric**


$$\pi_0(dz_0) \prod_{i=1}^k \mathbb{T}(dz_i, t_i : z_{i-1}, t_{i-1})$$

# Main Results

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

Bounds guaranteed by  $\mathcal{L}_1$  DRAC on

$$\mathfrak{D}_p \left( \mathbb{Q}_{[0,T]}, \mathbb{Q}_{[0,T]}^* \right) \doteq \sup_{k \in \mathbb{N}} \mathbb{W}_p^k \left( \mathbb{Q}_{T_1 \dots T_k}, \mathbb{Q}_{T_1 \dots T_k}^* \right)$$



**UB**  $\mathfrak{D}_p \left( \mathbb{Q}_{[0,T]}, \mathbb{Q}_{[0,T]}^* \right) \leq \boxed{\rho}$

**UUB**  $\mathfrak{D}_p \left( \mathbb{Q}_{[T',T]}, \mathbb{Q}_{[T',T]}^* \right) \leq \boxed{\rho'(T')} \doteq \gamma \boxed{\mathbb{W}_p(\mathbb{Q}_{T'}, \mathbb{Q}_{T'}^*)} e^{-\lambda(T')} + \boxed{\zeta(\omega, T_s)} + \boxed{\kappa}, \forall t \in [T', T]$

The **uniform (ultimate)** bounds  $\rho$  and  $\rho'$  are determined by

- **Wasserstein** distance between the **initial distributions**
- Bounds that are  $\propto \frac{1}{\omega}$  (filter bandwidth) and  $\propto T_s$  (sampling period)
- Constant bounded away from zero due to:
  - Unmatched uncertainties
  - **Independence** of driving Brownian motions  $W_t$  and  $W_t^*$

# Generality of Results $\mathcal{L}_1$ guarantees of uniform bounds

## Deterministic

On **trajectories**

$$\text{Dist.}(x, x^*)^p \doteq \sup_{t \in [0, T]} \|x(t) - x^*(t)\|^p \leq \rho^p$$

Continuous

Finite-dimensional

$$\text{Dist.}_k(x, x^*)^p \doteq \sup_{k \in \mathbb{N}} \|x(t_{1..k}) - x^*(t_{1..k})\|_\infty^p \leq \rho^p$$

$W_t^*$  and  $W_t \equiv 0$  (**deterministic**): Wasserstein  $\rightarrow$  Euclidean metric



- 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."

## Stochastic

On **probability measures** of trajectories

$$\underbrace{W_p\left(\mathbb{Q}_{[0, T]}, \mathbb{Q}_{[0, T]}^*\right)^p}_{\inf_{\pi \in \Pi\left(\mathbb{Q}_{[0, T]}, \mathbb{Q}_{[0, T]}^*\right)} \int_{\mathcal{C}([0, T]; \mathbb{R}^n)} \text{Dist.}(x, x^*)^p \pi(dx, dx^*) \leq \rho^p$$



$$\sup_{k \in \mathbb{N}} \underbrace{W_p^k\left(\mathbb{Q}_{T_1 \dots T_k}, \mathbb{Q}_{T_1 \dots T_k}^*\right)^p}_{\inf_{\pi \in \Pi\left(\mathbb{Q}_{T_1 \dots T_k}, \mathbb{Q}_{T_1 \dots T_k}^*\right)} \int_{\mathbb{R}^{nk}} \text{Dist.}_k(x, x^*)^p \pi(dx, dx^*)} \leq \rho^p$$



$$\inf_{\pi \in \Pi\left(\mathbb{Q}_{T_1 \dots T_k}, \mathbb{Q}_{T_1 \dots T_k}^*\right)} \int_{\mathbb{R}^{nk}} \text{Dist.}_k(x, x^*)^p \pi(dx, dx^*)$$



# Numerical Experimentation

Angular rate dynamics of a **quadrotor**

Simcenter STAR-CCM+



# Numerical Experimentation



## DRAC control


- Marginal Gaussian & KDE fits
- **Distributions diverge** in the absence of  $\mathcal{L}_1$  DRAC

$$W_2(Q_t^*, Q_t) : 21.0$$





$$t = 0.0$$


  $\mathcal{L}_1$  DRAC **OFF**  
vs.  
 Nominal




$$W_2(Q_t^*, Q_t) : 15.7$$



$$W_2(Q_t^*, Q_t) : 17.3$$



$$W_2(Q_t^*, Q_t) : 18.4$$





# Numerical Experimentation

## DRAC control

- Marginal Gaussian & KDE fits
- **Independent** Brownian motions
  - Convergence up to a **nonzero** limit


$$W_2(Q_t^*, Q_t) : 21.0$$

$$t = 0.0$$


  $\mathcal{L}_1$  DRAC **ON**  
vs.  
 Nominal




$$X_{t,1,2}^* - X_{t,1,2}$$




$$W_2(Q_t^*, Q_t) : 15.6$$




$$X_{t,1,3}^* - X_{t,1,3}$$




$$W_2(Q_t^*, Q_t) : 17.2$$



$$X_{t,2,3}^* - X_{t,2,3}$$



$$W_2(Q_t^*, Q_t) : 18.5$$



# Numerical Experimentation

DRAC control


$t = 0.0$

ON




$W_2(Q_t^*, Q_t) : 21.0$

OFF




$W_2(Q_t^*, Q_t) : 21.0$




$\mathcal{L}_1$  DRAC ON



$\mathcal{L}_1$  DRAC OFF



# Predictable & Certifiable Data-driven Control

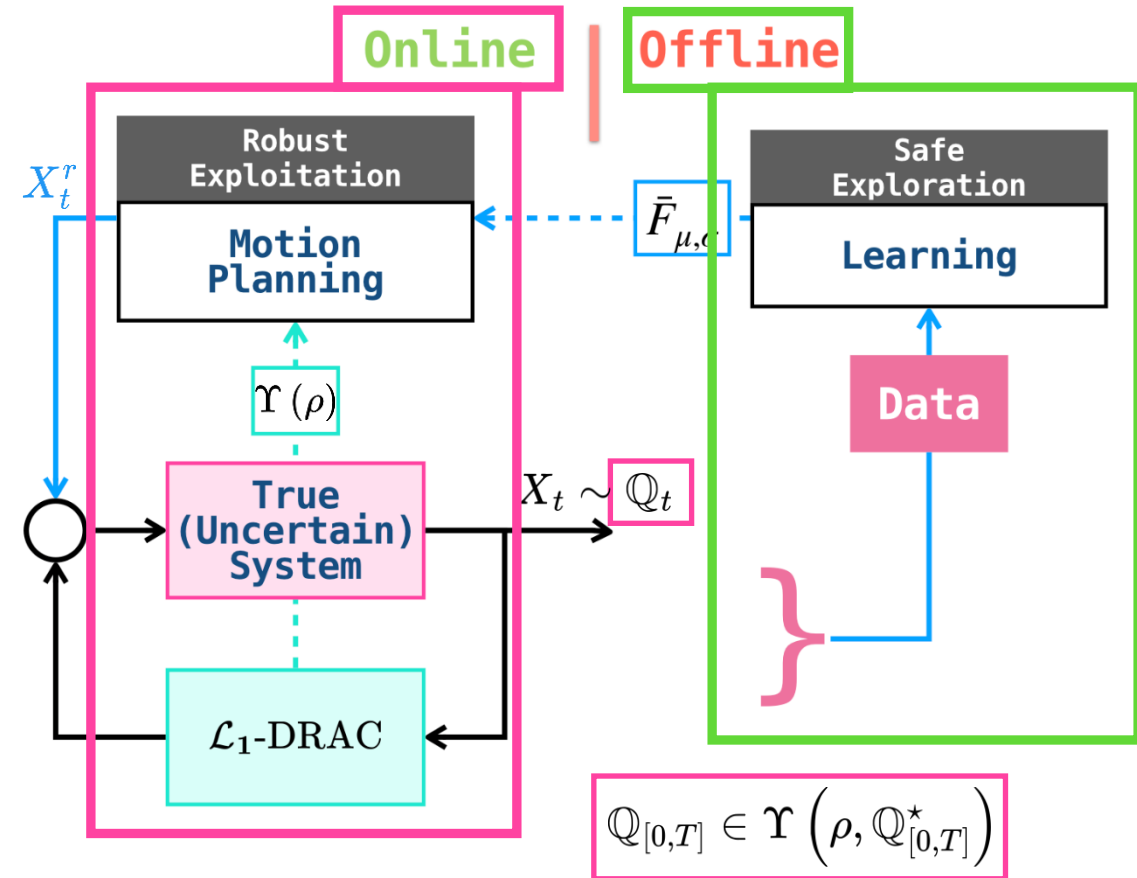


# Framework

- A **decoupled** framework
  - **Safety** decoupled from **performance**

## Safety

- Planning and Control
  - Online
  - **Persistent**
- ## Performance
- Data-driven learning
  - Offline and opportunistic
  - **Intermittent**



# Dist. Robust Planning

$$\min_{u_b \in \mathcal{U}_{[0,T]}} \max_{\bar{\mathbb{Q}}_{[0,T]} \in \Upsilon(\mathbb{Q}_{[0,T]}^*, \rho)} \mathbb{E}_{\bar{X}_{[0,T]} \sim \bar{\mathbb{Q}}_{[0,T]}} \left[ \int_0^T C(\bar{X}_t, u_{b(t)}) dt \right]$$

- **Min-max** optimal control over the **ambiguity set** of distributions

- Hedging against the effects of **worst distribution**  $\in \Upsilon$

- **Generalization** of standard min-max optimal control problem

- Solutions via **distributionally robust optimization** (DRO) [1]

- Theoretically sound and computationally tractable

- **Challenge:** Determination of ambiguity sets

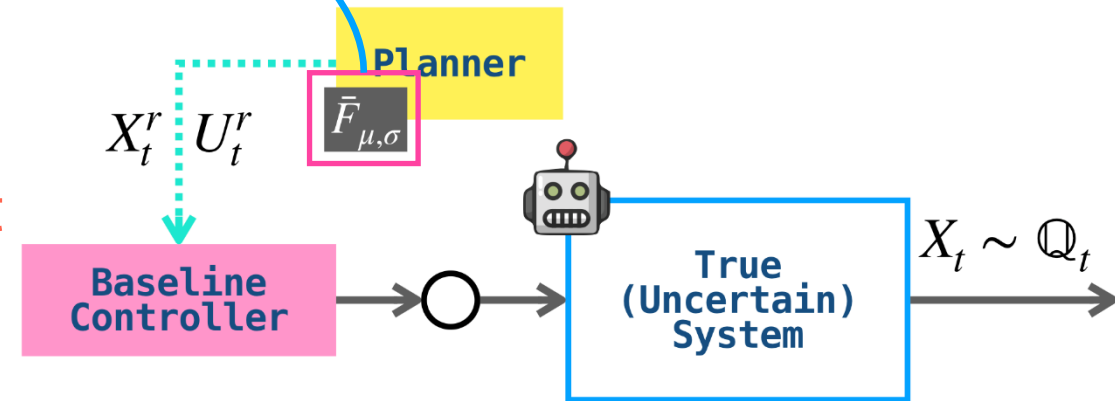
- $\neg$  **Finite samples** from **true distributions**

- Feedback  $\leftrightarrow$  only one sample from true time-varying conditional  $\mathbb{Q}_t$

- $\neg$  **Knowledge** of **moments** of  $\mathbb{Q}_t$

- Only finiteness of the second moment via well-posedness

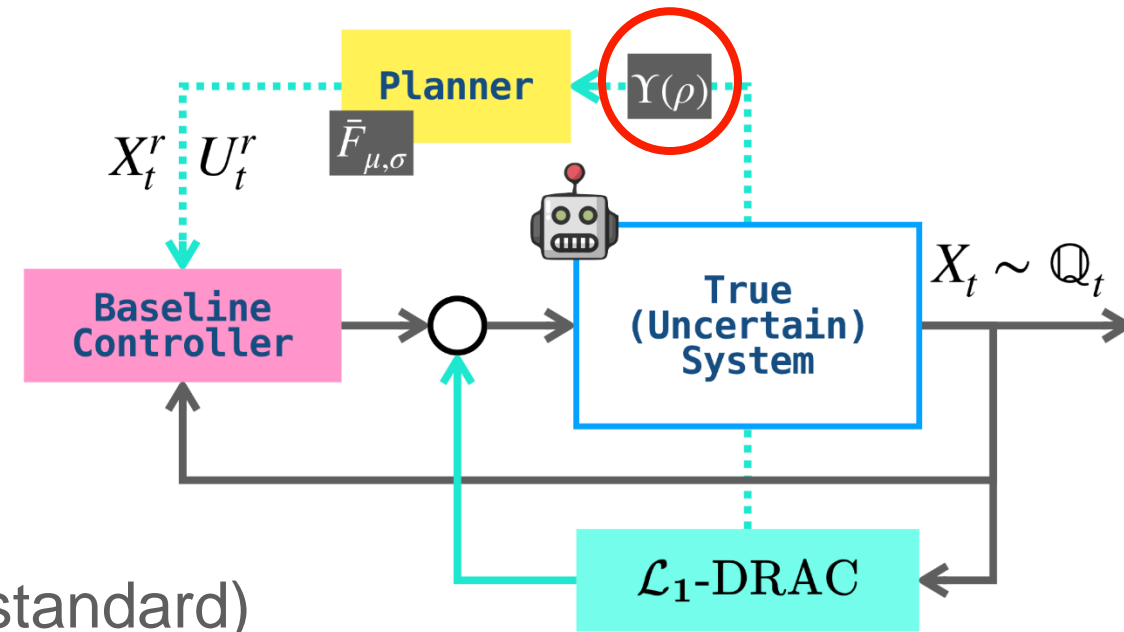
- $\neg$  **Compactness** of  $\text{supp}(\mathbb{Q}_t)$



# Dist. Robust Planning

$$\min_{u_b \in \mathcal{U}_{[0,T]}} \max_{\bar{Q}_{[0,T]} \in \Upsilon(Q_{[0,T]}^*, \rho)} \mathbb{E}_{\bar{X}_{[0,T]} \sim \bar{Q}_{[0,T]}} \left[ \int_0^T C(\bar{X}_t, u_{b(t)}) dt \right]$$

- **Challenge:** Determination of ambiguity sets
  - **Solution:** Use the uniform guarantees provided by  $\mathcal{L}_1$  DRAC
- **Bi-directional communication** between high and low-level components
  - **High  $\rightarrow$  Low-level:** Reference commands (standard)
  - **Low  $\rightarrow$  High-level:** Ambiguity sets/tubes



Planner  $\rightarrow$  **steers** uncertain  $Q_{[0,T]}$

- Via nominal  $Q_{[0,T]}^*$  and knowledge of  $Y(\rho)$





# Robust Exploitation

Predictable use of data-driven learned components

- Learned model's certificate:

$$\mathcal{R}(\mathcal{D}_{train}, \hat{f}; \theta^*) = \mathbb{E}_{(x_+, x, u) \sim \mathcal{D}_{train}} \text{Loss}(\hat{f}_{\theta^*}(x, u), x_{\perp}) \leq \delta$$

- Performance guarantee  $\delta$ 
  - over the training distribution (empirical probability measure)
- Guarantees for probability measures “around”  $\mathcal{D}_{train}$ ?

Worst-case risk:  $\hat{\mathcal{R}}_{\epsilon}(\mathcal{D}_{train}, \hat{f}; \theta^*) \doteq \sup_{\hat{\mathcal{D}} \in \mathbb{B}(\mathcal{D}_{train}, \epsilon)} \mathcal{R}(\hat{\mathcal{D}}, \hat{f}; \theta^*) \leq \hat{\delta}_{\epsilon}$

- Performance guarantee  $\hat{\delta}_{\epsilon} (> \delta)$  on the set of distributions  $\mathbb{B}(\mathcal{D}_{train}, \epsilon)$ 
  - $\mathbb{B}(\mathcal{D}_{train}, \epsilon)$  = Wasserstein ball of radius  $\epsilon$  centered on  $\mathcal{D}_{train}$
- Consequence of Kantorovich-Rubinstein duality [1]
  - Worst-case risk upper bounded by Lipschitz regularized nominal risk

# Robust Exploitation

Data-driven Learning

$$\hat{\mathcal{R}}_{\epsilon}(\mathcal{D}_{train}, \hat{f}; \theta^*) \doteq \sup_{\hat{\mathcal{D}} \in \mathbb{B}(\mathcal{D}_{train}, \epsilon)} \mathcal{R}(\hat{\mathcal{D}}, \hat{f}; \theta^*) \leq \hat{\delta}_{\epsilon}$$

Certificates of learning Over Set of Probability Measures

DR Planning +  $\mathcal{L}_1$  DRAC


$$\mathbb{Q}_{[0,T]} \in \Upsilon(\mathbb{Q}_{[0,T]}^*, \rho)$$

Certificates of Robustness Over Set of Probability Measures

Same Language

- Lack of a ‘language barrier’
  - Certificates in the space of probability measures
- No constraints on learning
- Raised abstractions of robustness

Example setup



# Continuation

- Inclusion of perception and other high-dimensional sensing
  - Natural via high-probability finite-sample guarantees of Wasserstein metrics
- Local high-probability results
  - Application oriented
  - Larger class of dynamics: well-posedness up to first exit times
- Stochastic planning enabled by  $\mathcal{L}_1$  DRAC's uniform bounds
  - Nominal system planning via Girsanov change of measures and Feynman-Kac formulae for nominal system
  - Extended to uncertain systems via  $\mathcal{L}_1$  DRAC
  - Robust planning for uncertain systems for the cost of nominal planning
- Reachability analysis in the space of probability measures
  - Off-the-shelf tools for known systems with the robustness guarantees of  $\mathcal{L}_1$  DRAC

# Ongoing Related Projects

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

