Safe Learning for Autonomous Systems

via $\mathcal{L}_1$ Adaptive Control, Contraction Theory, and Machine Learning

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



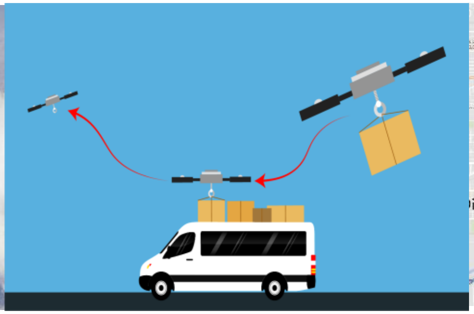
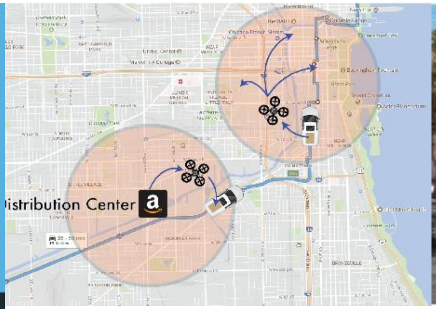

UNIVERSITY OF ILLINOIS URBANA-CHAMPAIGN

Evolution of Aerospace Industry over the Last Century

- CONSISTENT
- SAFETY CRITICAL
- POWERFUL
- LARGE-SCALE
- PROFITABLE
- CERTIFIED

Building a Bridge

- AUTONOMOUS
- COLLABORATIVE
- INTELLIGENT
- LOW-COST
- RECONFIGURABLE
- MULTI-AGENT
- HUMAN-CENTERED

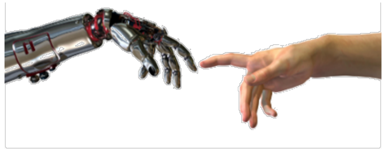

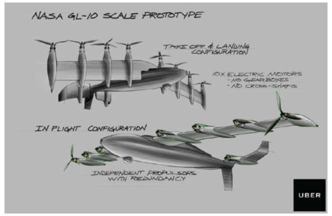

Talent
Technology
Expertise

AI
Cloud
Culture

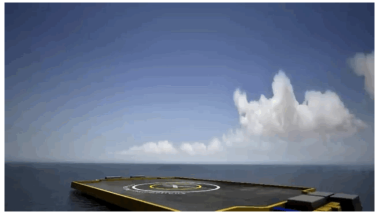

Recent Advances in Learning-Based Control

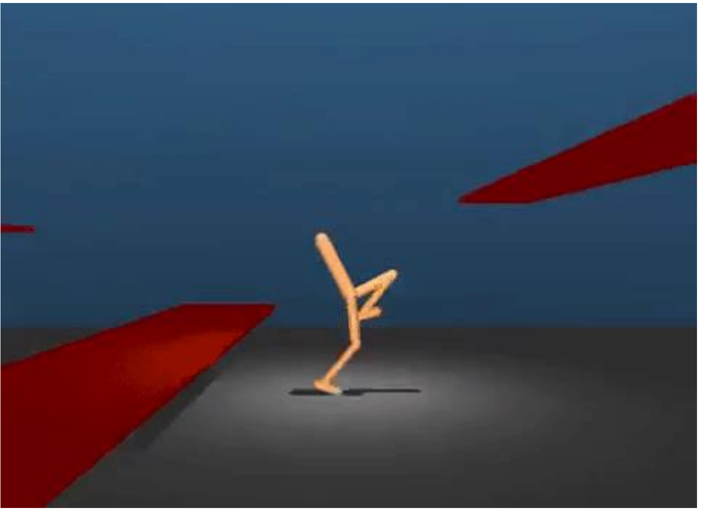

Source: Deepmind (Alphabet Inc.)

Heess et al. "Emergence of Locomotion Behaviours in Rich Environments." arXiv preprint arXiv:1707.02286 (2017)

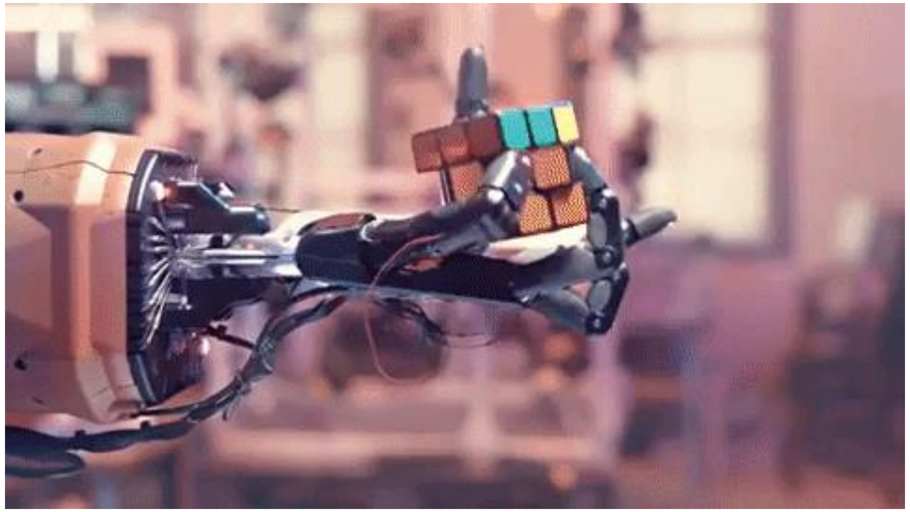

Source: OpenAI Inc.

Akkaya et al. "Solving rubik's cube with a robot hand." arXiv preprint arXiv:1910.07113 (2019).

Source: Skydio

But the Tesla accidents…

Source: The Inrtercept

Source: Autoweek

Source: ABC News

Safety-Critical Systems

Learjet 45

Aerial robots

Self-driving cars (credit: Daniel Lu)

- Difficulty in ***safely*** obtaining a large amount of data for training.
  - ➢ The architecture must be designed appropriately.
- Accidents can be expensive and sometimes even deadly.
  - ➢ Safety needs to be ***guaranteed***.

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

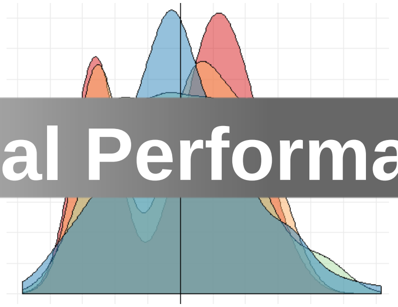

Bridging the divide

Learning-Based Control Setup

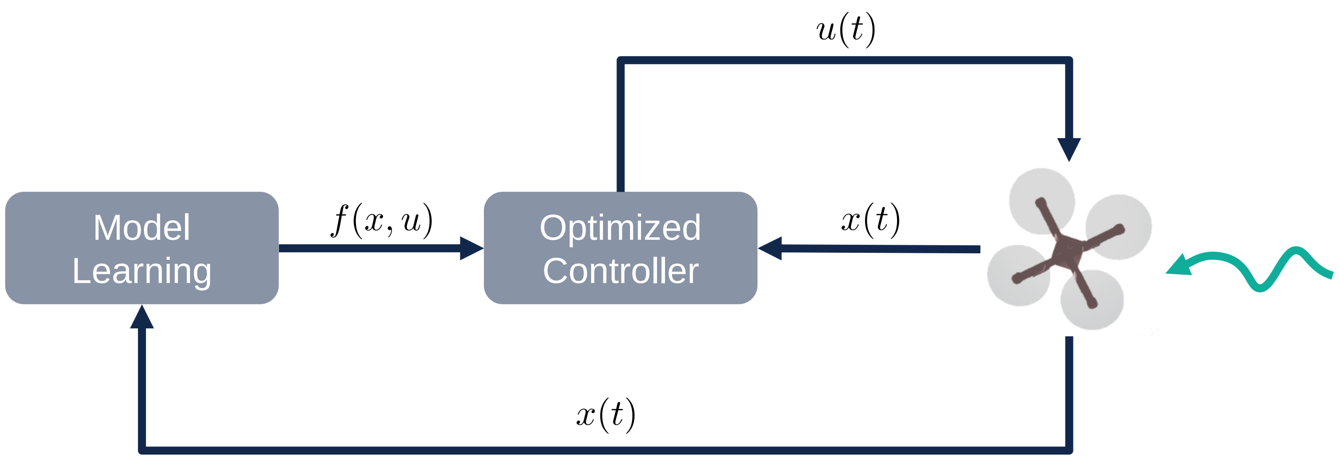

Sources of Uncertainty

- Modeling Errors
- External Disturbances

Loss of performance and stability guarantees for systems under uncertainty!

Safety ***must*** be built into control architecture

Safe Learning and Control Setup

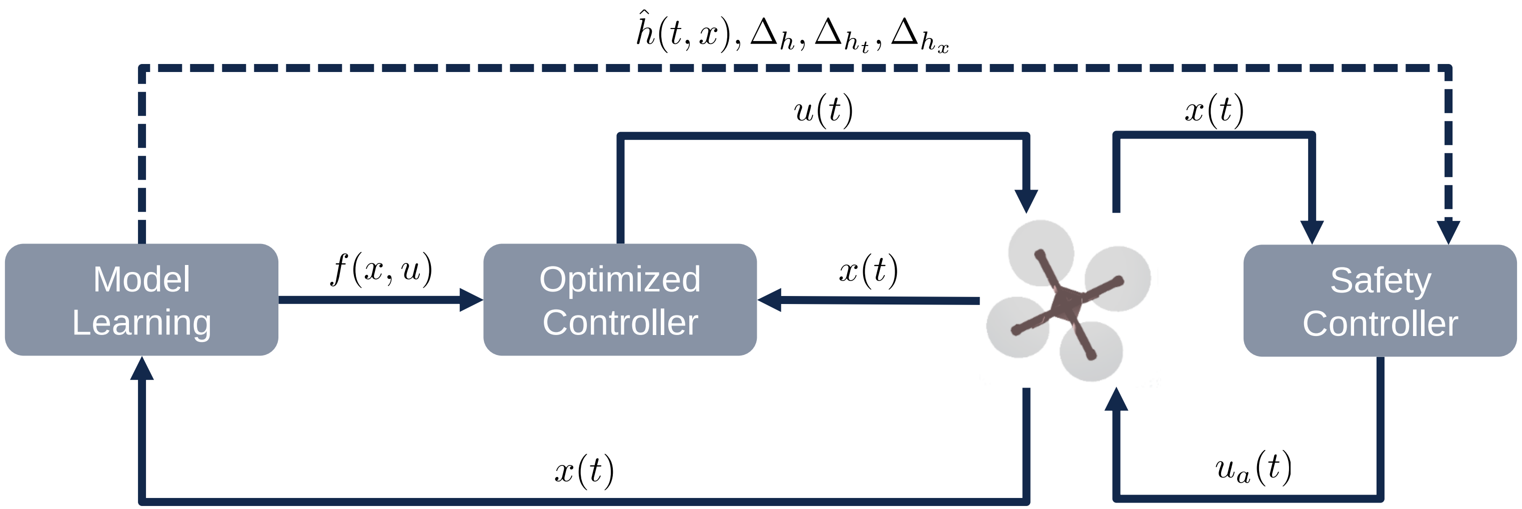

$$\dot{x}(t) = \boxed{f(x,u)} + \boxed{h(t,x)}$$

Nominal Model **Uncertainty**

$$\boxed{\Delta_h, \Delta_{h_t}, \Delta_{h_x}}$$

Bounds on the uncertainty
and its growth in t and x.

Safe Learning and Control Setup

The safety controller must provide ***certificates of performance and robustness:***

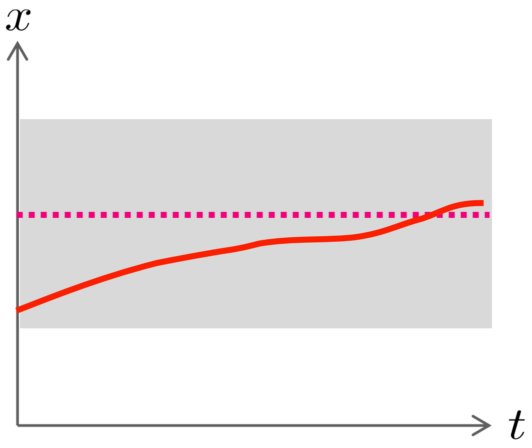

Transient performance

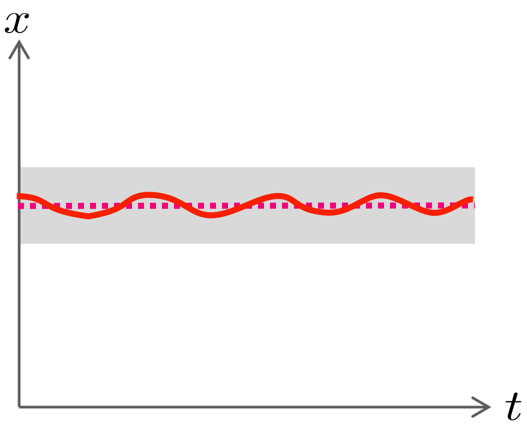

Steady-state performance

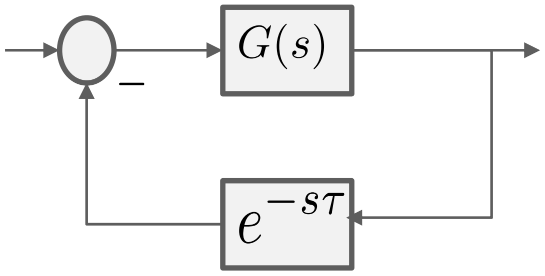

Time-delay margin

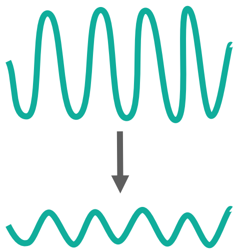

Disturbance rejection

$\mathcal{L}_1$-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, etc.

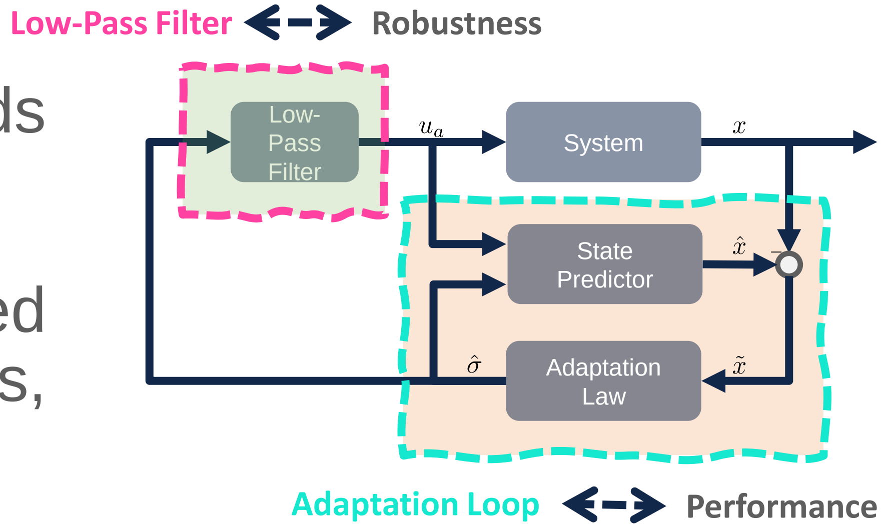

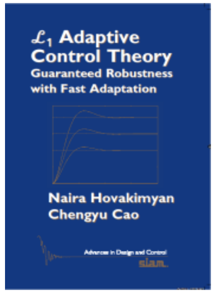

Safe Learning and Control Setup

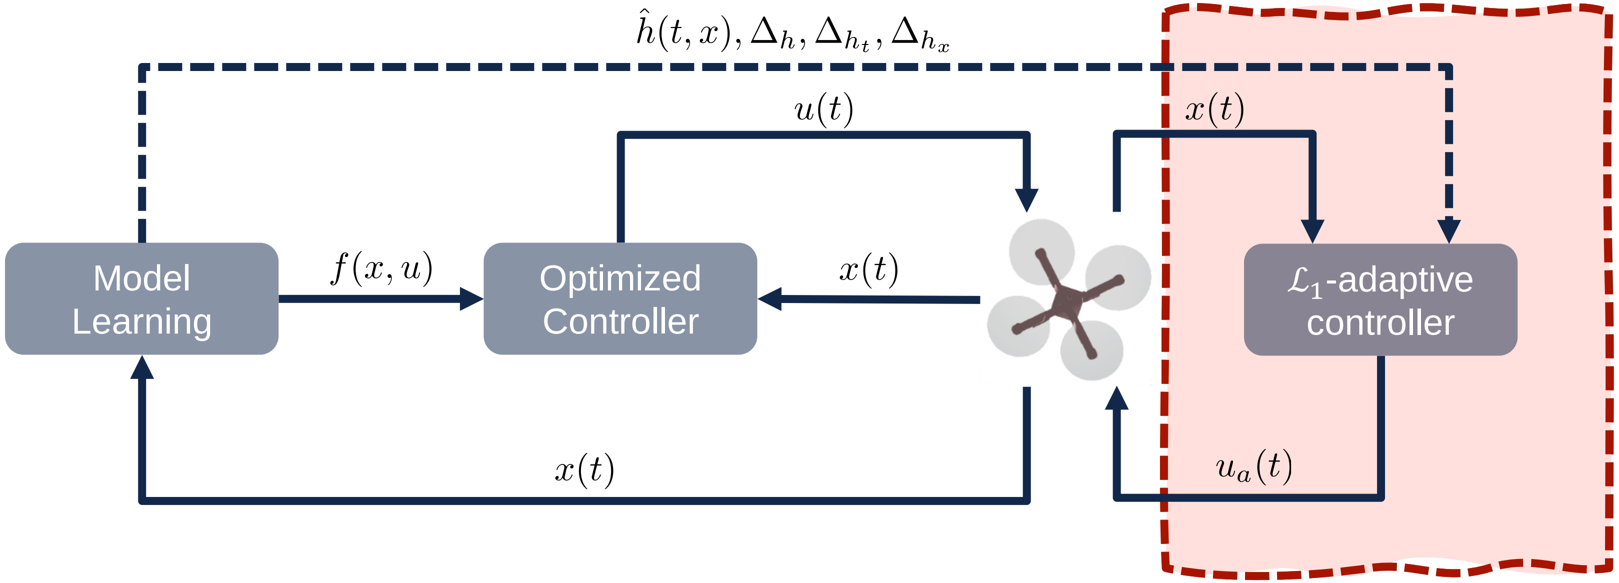

- Retaining the key features of performance and robustness guarantees of $\mathcal{L}_1$- adaptive controllers
- Benefitting from the versatility offered by machine learning methods

Guarantees Provided by $\mathcal{L}_1$ Theory

- $x^\star$: Desired system trajectory
- x_r: Reference system (best performance achieved using the $\mathcal{L}_1$ framework with known uncertainties, *but non-implementable*)
- x : Actual system trajectory

$$\|x_r(t) - x(t)\| \leq \rho_a(\Gamma^{-1})$$

Performance: Adaptation rate Γ

Decoupling

$$\|x^\star(t) - x_r(t)\| \leq \sqrt{ke^{-2\lambda t}\|x^\star(0) - x(0)\| + \rho_r(\omega^{-1})}$$

Robustness: Filter bandwidth ω

Exponential decay of error due to the difference in initial conditions

A Timeline of $\mathcal{L}_1$-Adaptive Control Theory

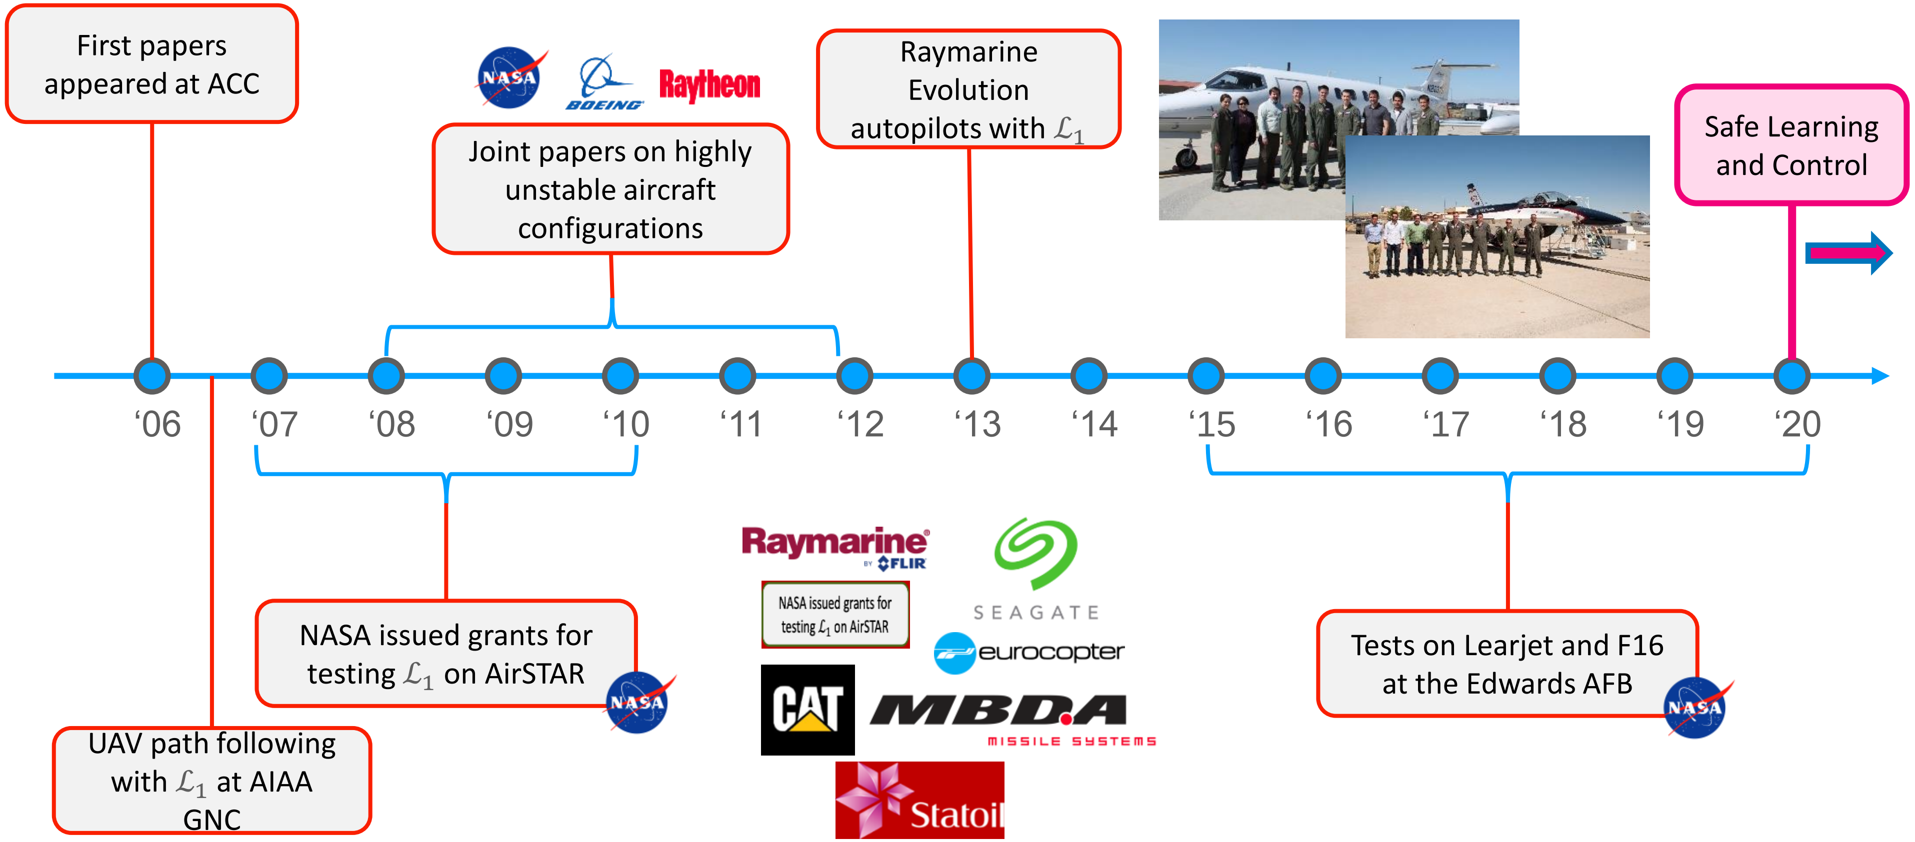

Robust Flight Control: Learjet at Edwards AFB

- Recovered flying qualities of different Variable Stability System configurations
- Restored handling qualities to a safe and consistent level despite the off-nominal dynamics
- The controller was shown to be easily adjusted to improve handling qualities.

Ackerman, et al. "Evaluation of an $\mathcal{L}_1$ Flight Control Law on Calspan's Variable-Stability Learjet." AIAA Journal of Guidance, Control and Dynamics, vol. 40, No. 4, pp. 1051-1060, 2017.

AVIATION WEEK NETWORK

USAF Test Pilots School Vets Safer Adaptive Flight Controller

April 06, 2015

Credit: Calspan

"The $\mathcal{L}_1$ controller is designed to automatically intervene in the case of control problems, immediately reconfiguring the flight control system to compensate for degraded flying qualities from mechanical failure or battle damage to a control surface, or even the unintended result of shifting center-of-gravity inflight for better cruise performance. Acting as a backup to the standard flight control system, the $\mathcal{L}_1$ is **designed to provide safe, predictable, reliable** and repeatable responses that would free up pilots to deal with the emergency and further compensate for reduced performance."

Guy Norris, Aviation Week and Space Technology.
Published on April 6, 2015.

2016 Flight Test of F-16


TECHNOLOGY

Adaptive Advance

Flight tests of adaptive controller could pave way for backup protection against control loss

U.S. AIR FORCE/CALSPAN

Lifting Body Incident (1967) and the 2018 Flight Test

Puig-Navarro et al. “An $\mathcal{L}_1$ Adaptive Stability Augmentation System Designed to MIL-HDBK-1797 Level 1 Specifications.” *In Proceedings of* AIAA Guidance, Navigation and Control Conference, San Diego, CA, 2019.

AEROSPACE & DEFENSE TECHNOLOGY

The Engineer's Guide to Design & Manufacturing Advances

APRIL 3, 2018 | AEROSPACE

Flight Control System Virtually Eliminates Pilot Error

The $\mathcal{L}_1$ adaptive flight control system installed in an aircraft. (Maj. Miguel J. Carreras)

"$\mathcal{L}_1$ adaptive control has overcome some of the major limitations of conventional adaptive control systems by providing **predictable robustness guarantees in the presence of a large class of uncertainties**. It has the potential to revolutionize aircraft safety by greatly diminishing the possibility of pilot error during high workload maneuvers."

Published on April 3, 2018.

... and many other applications

Time-critical coordination

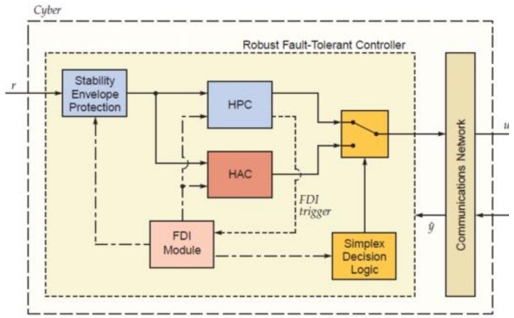

Fault-tolerant software

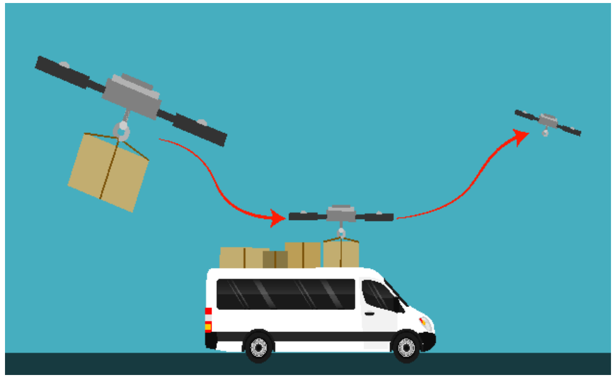

Drone—based package delivery

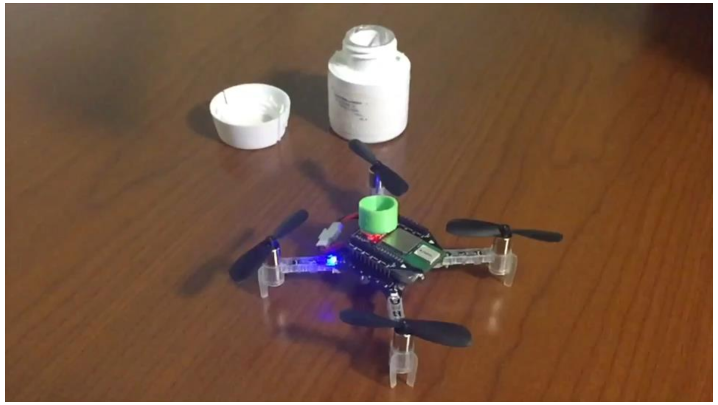

Indoor aerial vehicles

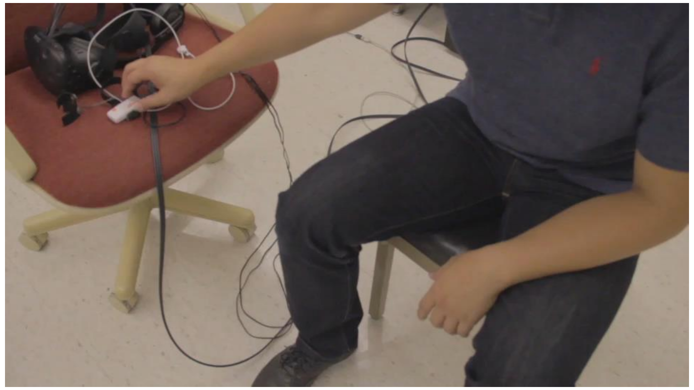

Human-centered (perceived) safety

Complex Dynamics and Environments

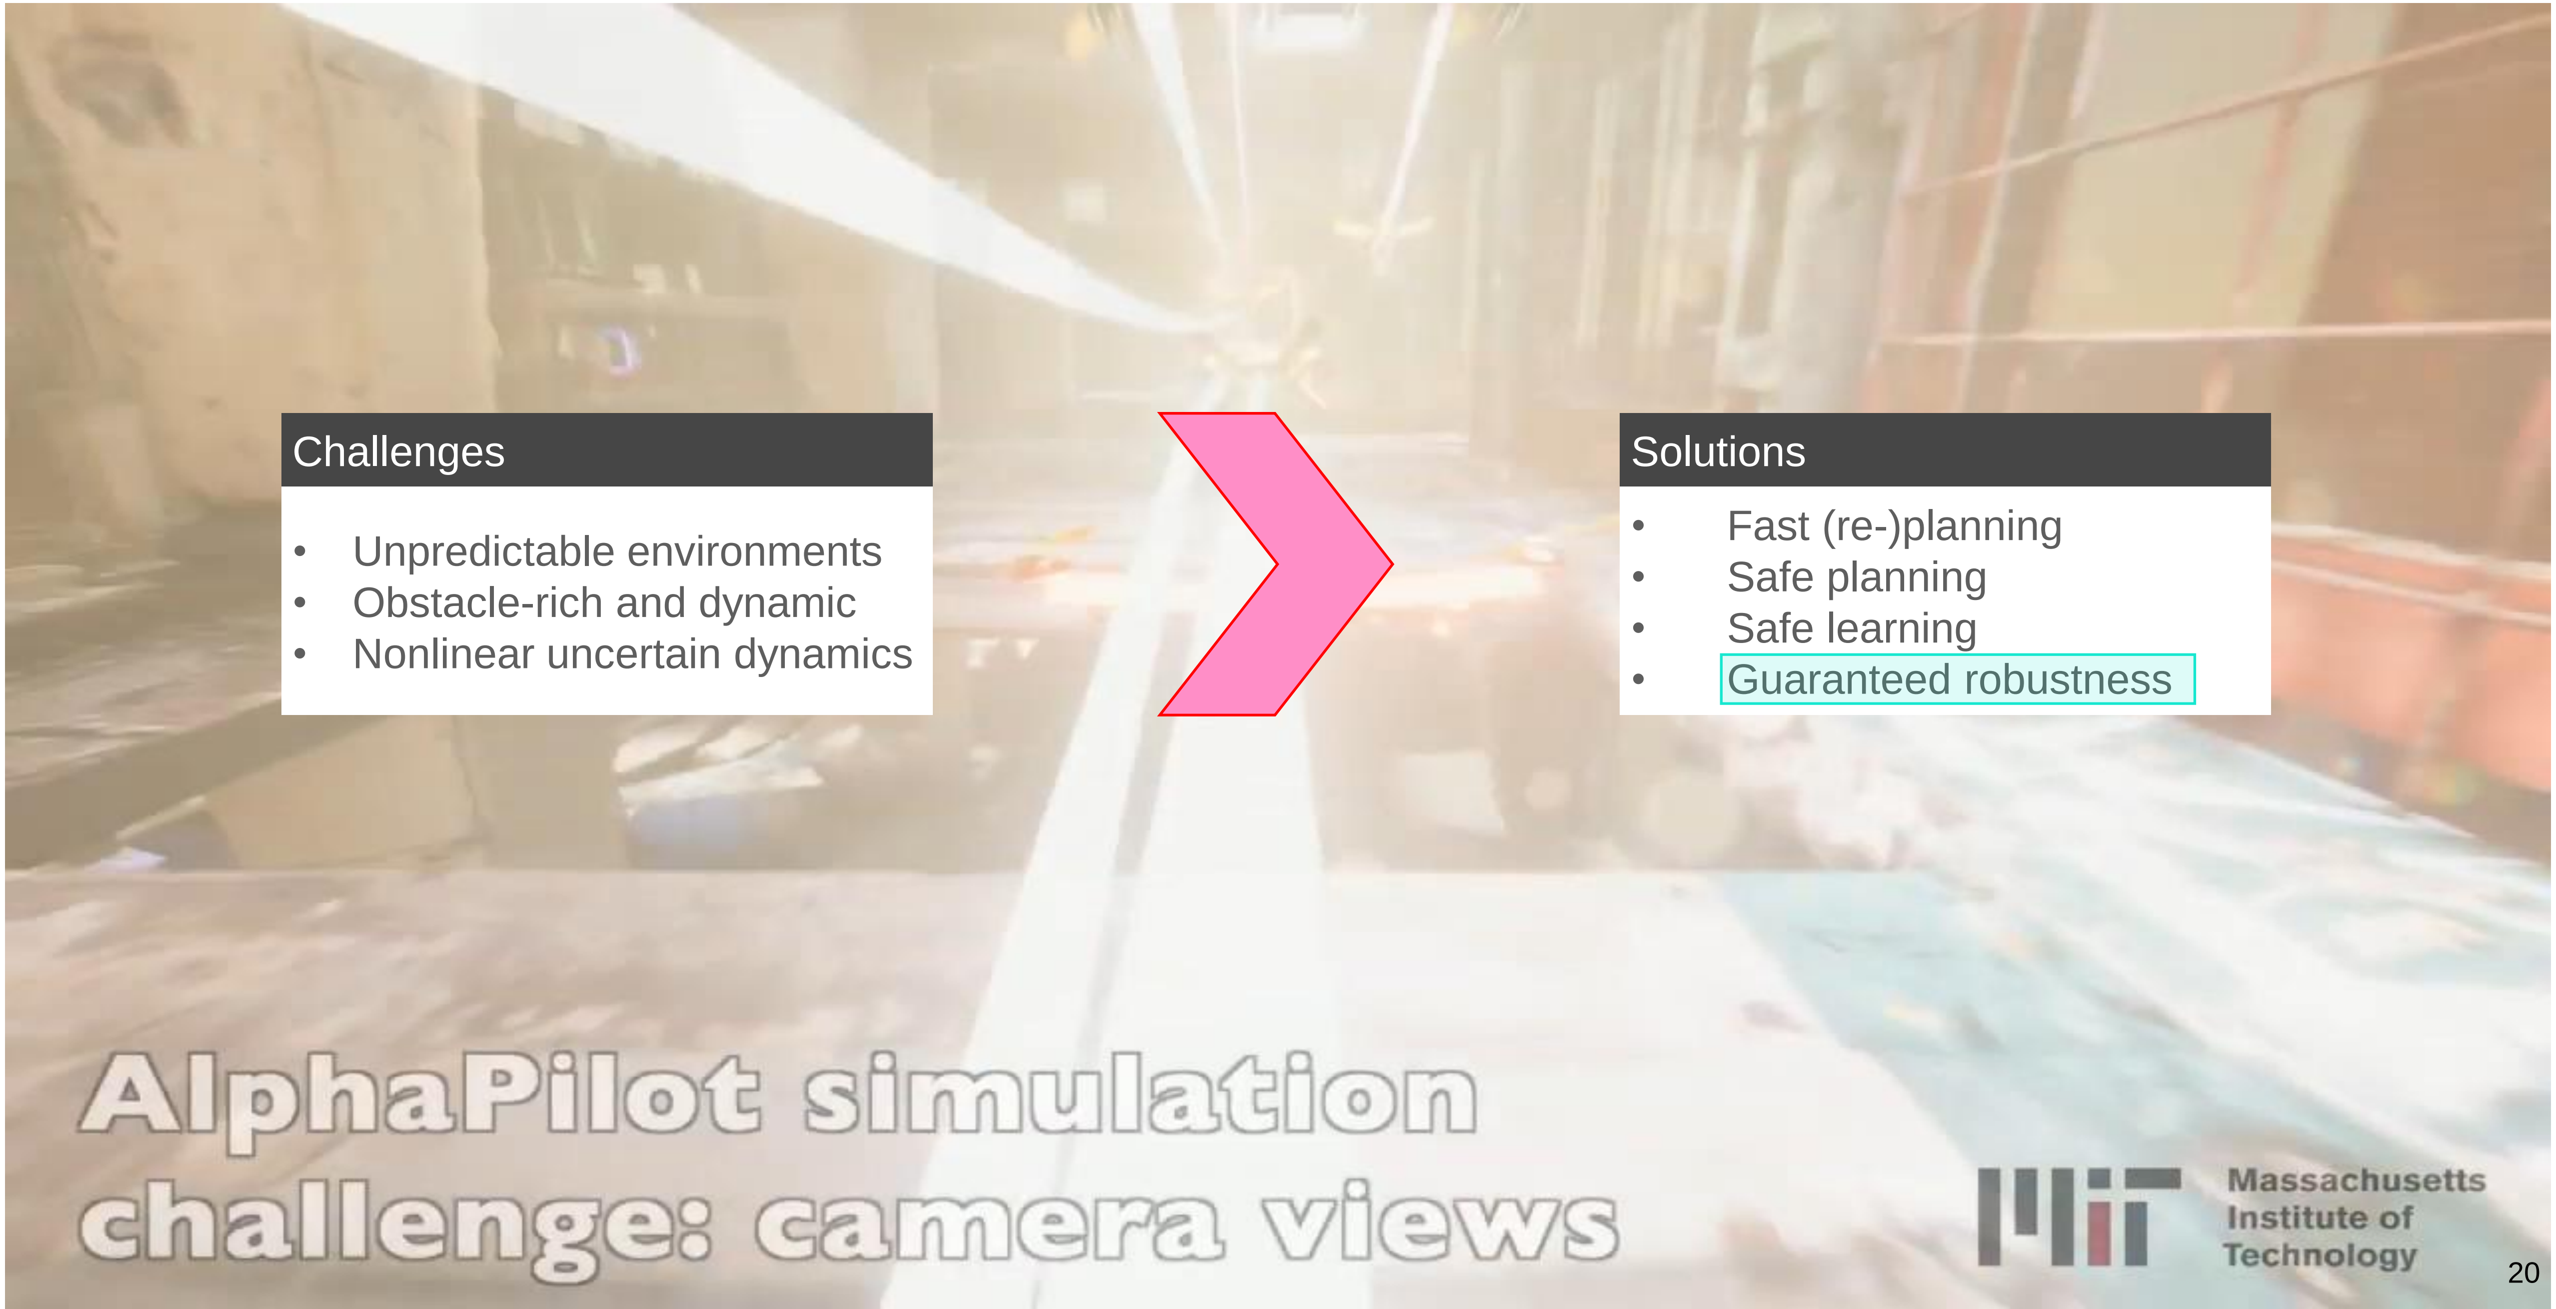

Challenges

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

Solutions

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

Safe Learning & Control: Framework

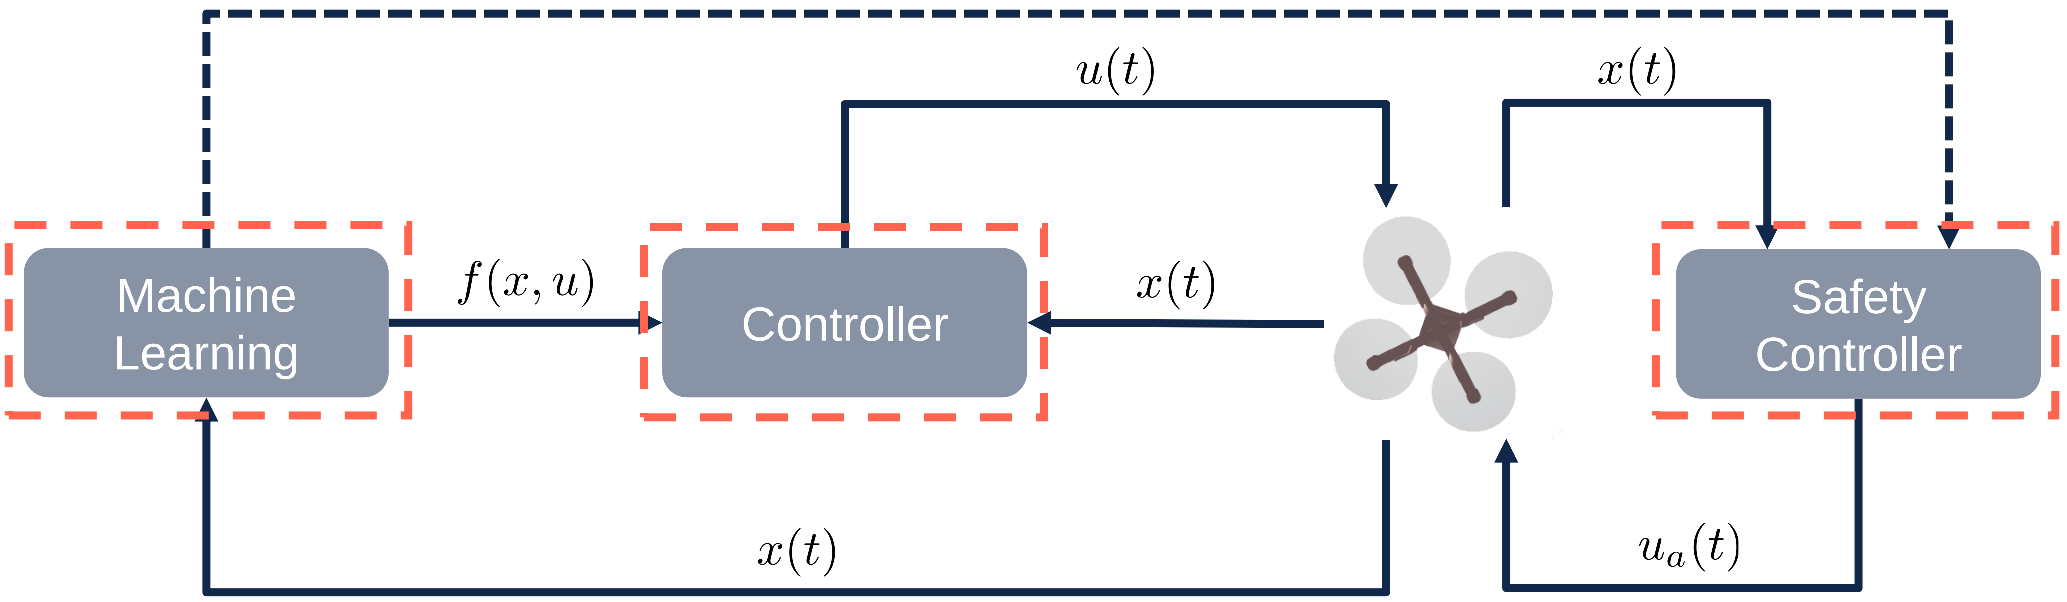

- Nominal Control Design – Foundation for desired properties
- Robust Adaptive Augmentation – Build-up for safety
- Learned Models – Performance and robustness

Motivation

- Recent developments:
  - Contraction theory: Differential geometry-based simplified synthesis for nonlinear models
  - Machine learning: Use of data and computation to accurately model complex systems
  - Autonomous UAS: Low-cost platforms with ever-improving autonomy capabilities.

We need to rethink and view robust adaptive augmentation through the lens of the recent developments

- Robustness and $\mathcal{L}_1$ adaptive augmentation of learned systems and controllers
- Safe use of data-dependent ML models
- Certificates of guaranteed performance

Outline

- Robust Control Contraction Metrics (RCCM)
- $\mathcal{CL}_1$-$\mathcal{GP}$: Contraction $\mathcal{L}_1$ adaptive control with Bayesian learning
- $\mathcal{L}_1$-RG: Adaptive reference governors for constrained uncertain systems
- $\mathcal{L}_1$-MPPI: Fast and safe re-planning with stochastic optimization
- $\mathcal{L}_1$-RL: Robustifying RL policies with $\mathcal{L}_1$ adaptive augmentation
- Learn-to-Fly (L2F): $\mathcal{L}_1$ adaptive control for safe learning on the fly
- Robust adaptive control of linear parameter-varying (LPV) systems
- Distributionally robust adaptive control (DRAC) of stochastic systems
- ILF-$\mathcal{L}_1$: Hardware experimental verification on a quadrotor
- Difftune: Data-driven tuning controllers

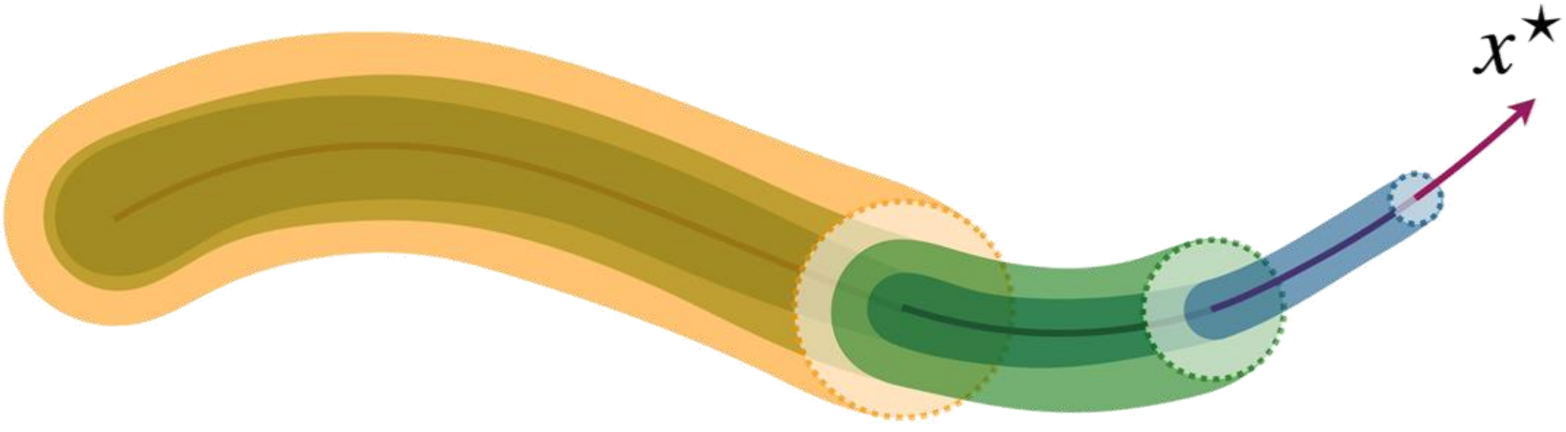

Safe Learning-based Control with $\mathcal{L}_1$ Adaptive Control

Contraction Theory-based Control

- Consider the nominal uncertainty-free dynamics $\dot{x} = F(x,u) = f(x) + B(x)u$ and a desirable pair $(x^\star, u^\star)$ that satisfies $\dot{x}^\star = F(x^\star, u^\star)$
- Contraction theory allows the synthesis of feedback $u = k(x, x^\star, u^\star)$ such that $x \to x^\star$
  - Incremental stability

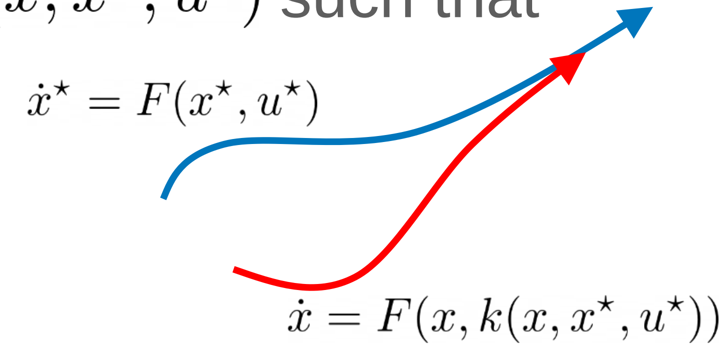

- Control contraction metrics (CCMs) parametrize Lyapunov functions on the differential space

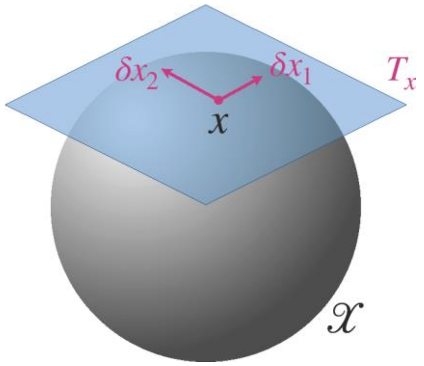

- A uniformly positive and bounded $M(x) \in \mathbb{S}^n$ is a CCM if

$$\underline{\alpha}\mathbb{I}_n \preceq M(x) \preceq \mathbb{I}_n\overline{\alpha},$$

$$\delta_x^\top M(x)B(x) = 0 \Rightarrow$$

$$\delta_x^\top \left(\partial_f M(x) + \left[M(x)\frac{\partial f(x)}{\partial x}\right]_{\mathbb{S}} + 2\lambda M(x)\right)\delta_x \leq 0$$

Convex Conditions!

where δ_x is the differential.

Robust Control Contraction Metrics

- Standard CCM-based controllers can be conservative for disturbance rejection
  - Disturbance not explicitly considered during synthesis
- For disturbed system

$$\dot{x} = f(x) + B(x)u + B_w(x)w$$
$$z = g(x, u)$$

Bounded disturbances

Output variables

we develop robust CCMs to minimize the universal $\mathcal{L}_\infty$ gain, α

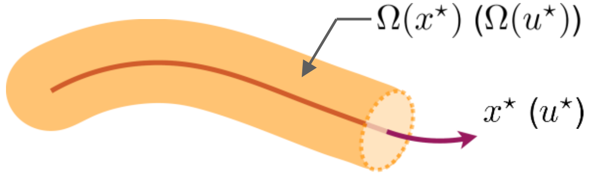

$$\|z - z^\star\|_{\mathcal{L}_\infty} \leq \alpha \|w - w^\star\|_{\mathcal{L}_\infty}$$

  - Disturbance rejection explicitly considered during synthesis
  - Provides certificate tubes for both states and control inputs
  - Can be synthesized via solving convex optimization (i.e., LMI) problems
  - Proves to yield tighter state tubes than the CCM-based approach [1], under certain conditions

Zhao, Lakshmanan, Ackerman, Gahlawat, Pavone, and Hovakimyan. Tube-certified trajectory tracking for nonlinear systems with robust control contraction metrics. *IEEE RA-L*, 2022.

[1] Singh, Landry, Majumdar, Slotine, and Pavone. Robust feedback motion planning via contraction theory. *IJRR*, 2019.

RCCM on a 3D Quadrotor

- To show the effectiveness of our proposed RCCM, we verify the controller on a 3D quadrotor in a cluttered environment

- Dynamics
  - 10 states, 4 inputs
  - Wind disturbances

- Comparative Controllers
  - Purely CCM-based approach [1]
  - RCCM-P (ours, optimizing the tube size for position states only)

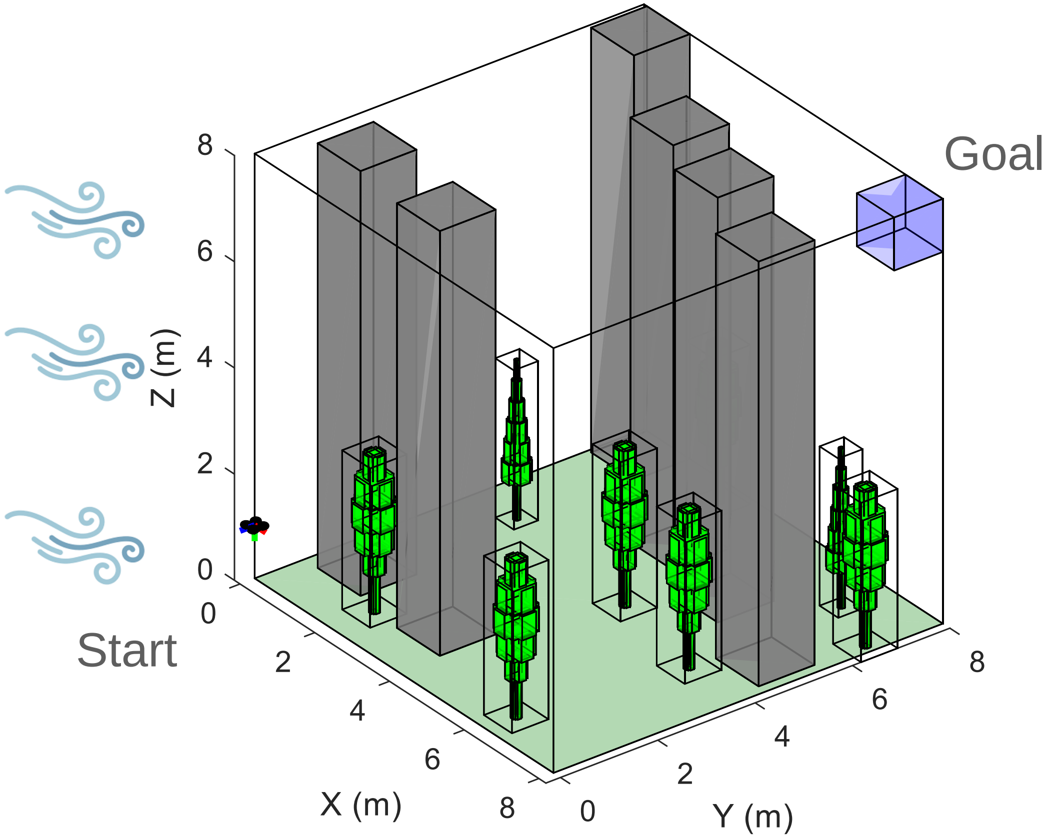

[1] Singh, Landry, Majumdar, Slotine, and Pavone. Robust feedback motion planning via contraction theory. *IJRR*, 2019.

RCCM on a 3D Quadrotor

CCM [1]

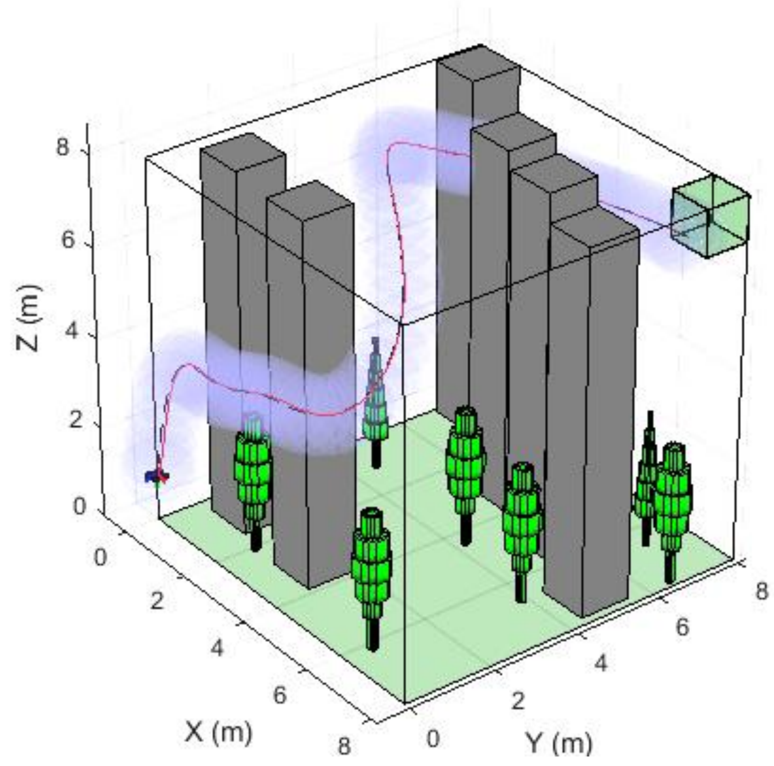

RCCM-P (ours)

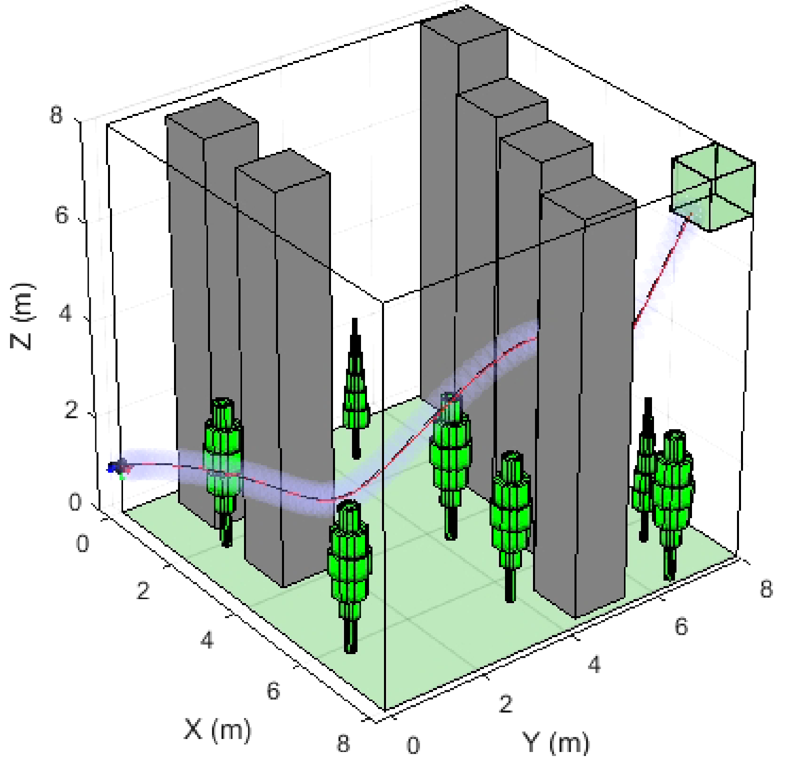

---- Planned   —— Actual

- Tube size for X-Y-Z: 0.768 m
- Travel time: 13.92 seconds

Under Wind Disturbance up to $1\ \mathrm{m}/s^2$

- Tube size for X-Y-Z: 0.316 m
- Travel time: 9.68 seconds

[1] Singh, Landry, Majumdar, Slotine, and Pavone. Robust feedback motion planning via contraction theory. *IJRR*, 2019.

Zhao, Lakshmanan, Ackerman, Gahlawat, Pavone, and Hovakimyan. Tube-certified trajectory tracking for nonlinear systems with robust control contraction metrics. *IEEE RA-L*, 2022.

$\mathcal{CL}_1$-$\mathcal{GP}$: Contraction $\mathcal{L}_1$ with Bayesian Learning

- RCCM allows to synthesize robust baseline controllers
- We can use ML in the form of Gaussian process regression to learn epistemic uncertainties and further improve performance
- To counter learning transients and to provide guarantees of safety, we bring in $\mathcal{L}_1$ adaptive control

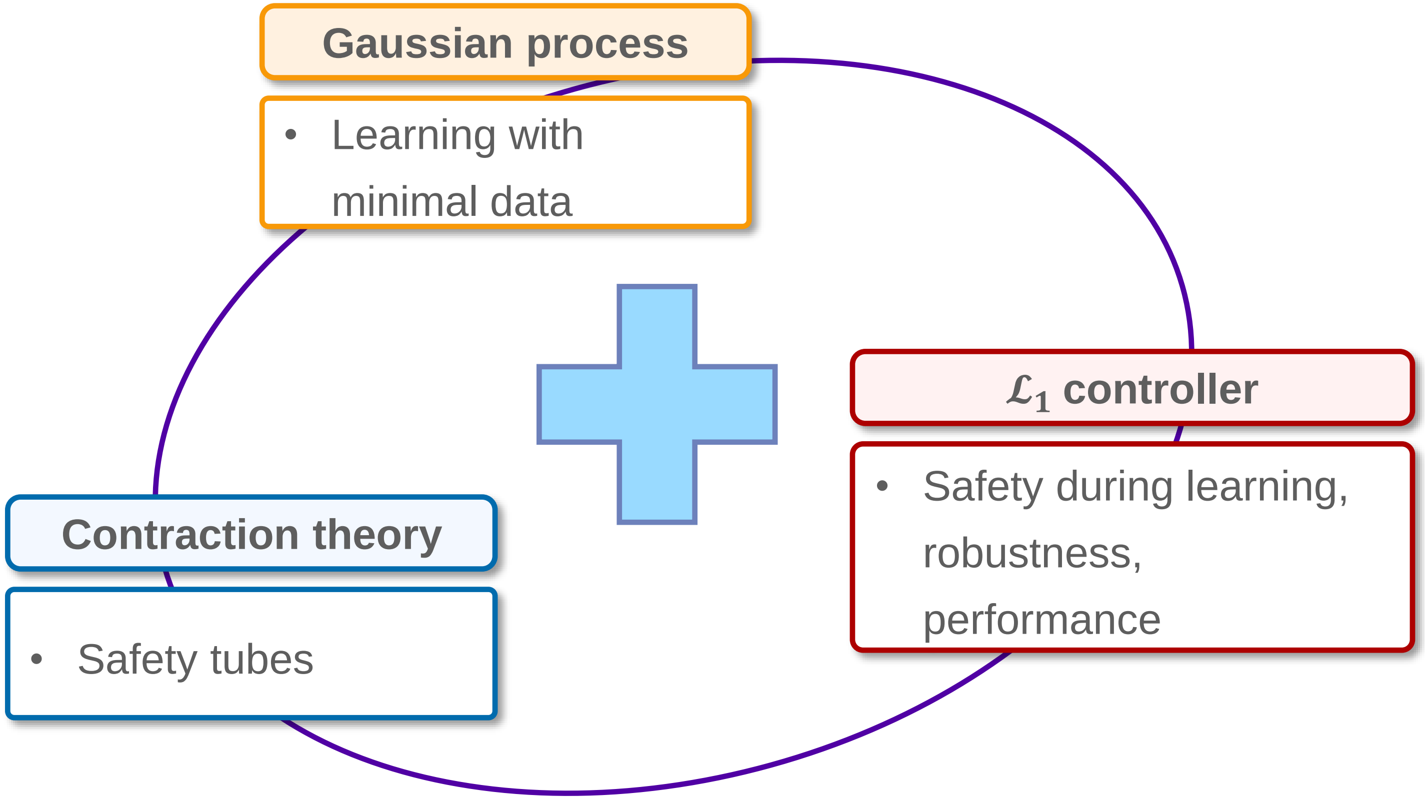

Gahlawat, et al. Contraction $\mathcal{L}_1$-Adaptive Control using Gaussian Processes. *Annual Conference on Learning for Dynamics and Control (L4DC)*, 2022.

$\mathcal{CL}_1$-$\mathcal{GP}$: Contraction $\mathcal{L}_1$ with Bayesian Learning

- Safety certificates in the form of tubes from the $\mathcal{CL}_1$-$\mathcal{GP}$ framework which enables safety during learning

- Natural framework for learning using $\mathcal{GP}$:
  - guaranteed performance during the learning transients
  - improved performance of the $\mathcal{L}_1$ adaptive controller, i.e., smaller tubes
  - Improved quality of the planned trajectory

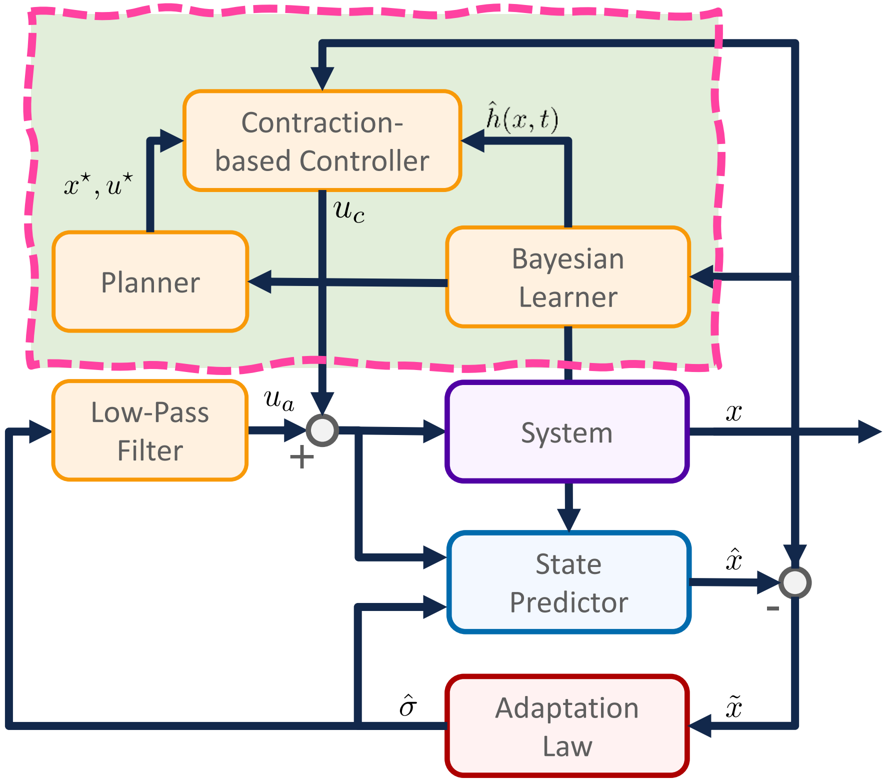

Gahlawat, et al. Contraction $\mathcal{L}_1$-Adaptive Control using Gaussian Processes. *Annual Conference on Learning for Dynamics and Control (L4DC)*, 2022.

$\mathcal{CL}_1$-$\mathcal{GP}$: Contraction $\mathcal{L}_1$ with Bayesian Learning

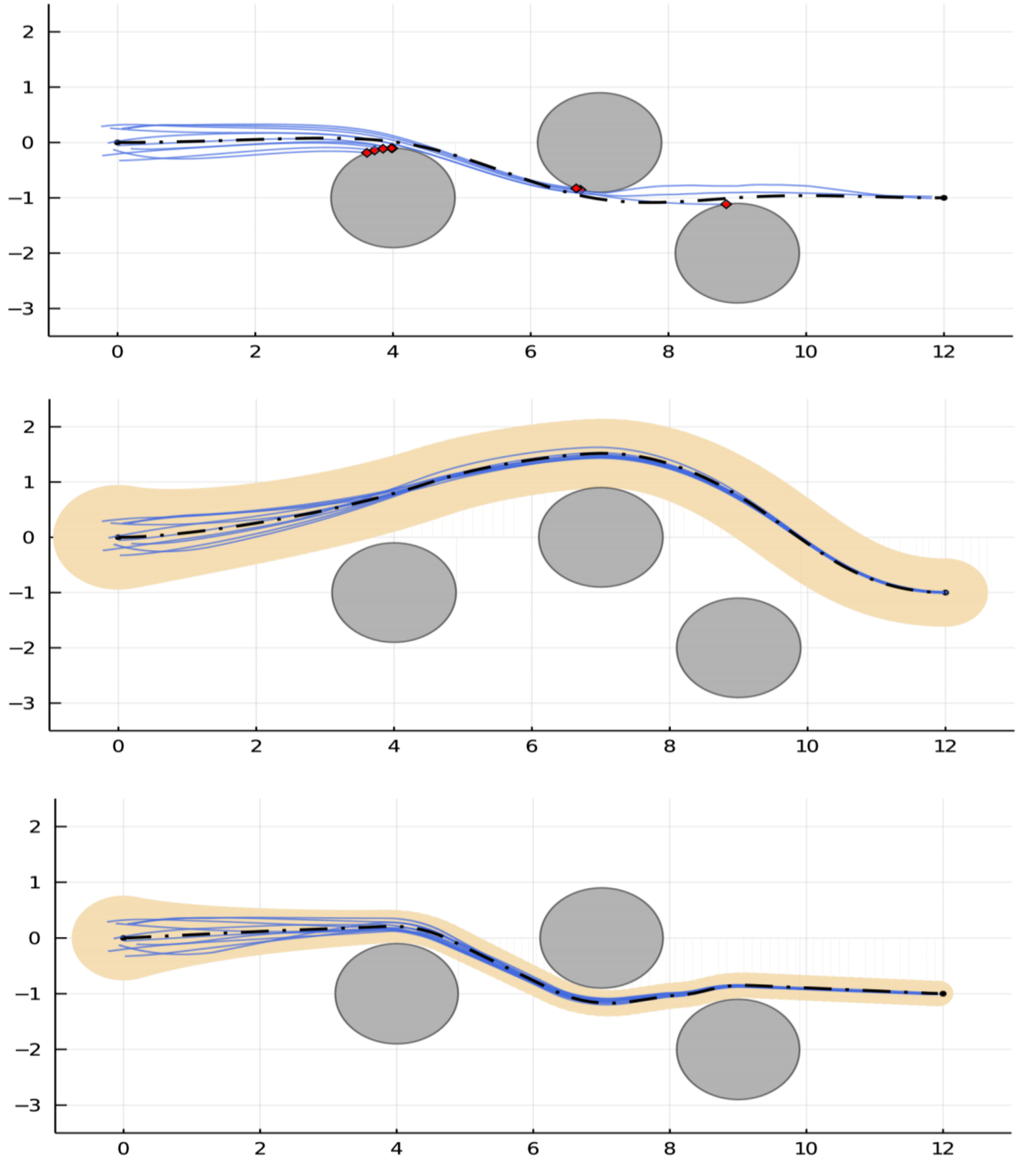

CCM only feedback → No safety guarantees

Out of 10 random initial conditions, 8 end in collision

Contraction $\mathcal{L}_1$ → Safety guaranteed

No learning → Safe but conservative

$\mathcal{CL}_1$-$\mathcal{GP}$ → Safety & Performance

As the uncertainty is learned → Performance improvement without sacrificing robustness

Gahlawat, et al. Contraction $\mathcal{L}_1$-Adaptive Control using Gaussian Processes. *Annual Conference on Learning for Dynamics and Control (L4DC)*, 2021.

$\mathcal{L}_1$-MPPI: Fast and Safe Re-Planning

- Fast and robust re-planning is needed for mission success in complex, dynamic and uncertain environments.
- Model predictive path integral (MPPI) control provides a framework for solving nonlinear MPC with complex constraints in near real-time.
- Robustness against dynamic uncertainties and disturbances is achieved through an $\mathcal{L}_1$ augmentation.

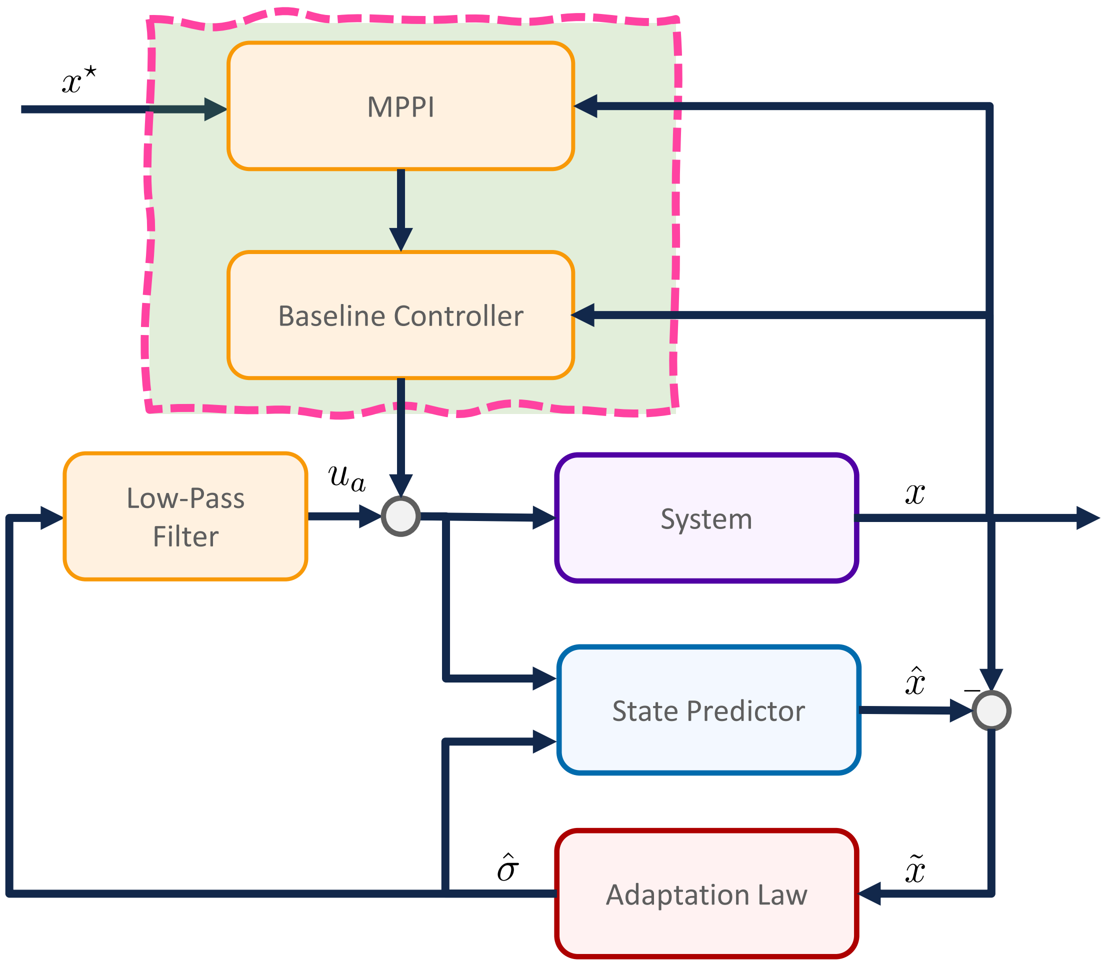

Pravitra, J., Ackerman, K. A., Cao, C., Hovakimyan, N., and Theodorou, E. A. L1-Adaptive MPPI Architecture for Robust and Agile Control of Multirotors. International Conference on Intelligent Robots and Systems, 2020.

$\mathcal{L}_1$-MPPI: Fast and Safe Re-Planning

Nonlinear system

$$\dot{x}_{t+1} = f(x_t, u_t + \epsilon) + h(x_t . u_t)$$

Nominal system — Unknown nonlinearity, disturbance

MPPI

- Optimal control computation forward in time for nominal system
- Update the control sequence iteratively by **sampling** (thousands of) trajectories

$\mathcal{L}_1$ adaptive control

- Robustifies against the **discrepancy**
- Guaranteed robustness and transient and steady-state bounds for performance

- Can leverage **parallel sampling** for modern GPUs
- Can handle **complex** (possibly non-differentiable) dynamics and cost functions

$\mathcal{L}_1$-MPPI: Fast and Safe Re-Planning

System: Control of a quadrotor (12 states, 4 control inputs) with state-dependent uncertainties and external disturbances in Flight Goggles.

Case	$\mathcal{L}_1$ off	$\mathcal{L}_1$ on
1)	✓	✓
2)	✗	✓
3)	✓	✓
4)	✗	✓
5)	✗	✓

Pravitra, Ackerman, Cao, Hovakimyan, Theodorou. $\mathcal{L}_1$-*Adaptive MPPI Architecture for Robust and Agile Control of Multirotors*, IROS, 2020.

Robustifying Reinforcement Learning (RL) with $\mathcal{L}_1$ Augmentation

Cheng, Zhao, Wang, Block, Hovakimyan. Improving the Robustness of Reinforcement Learning Policies with $\mathcal{L}_1$-*Adaptive Control*, Robotics and Automation Letters, 2020.

Distributionally Robust Adaptive Control (DRAC)

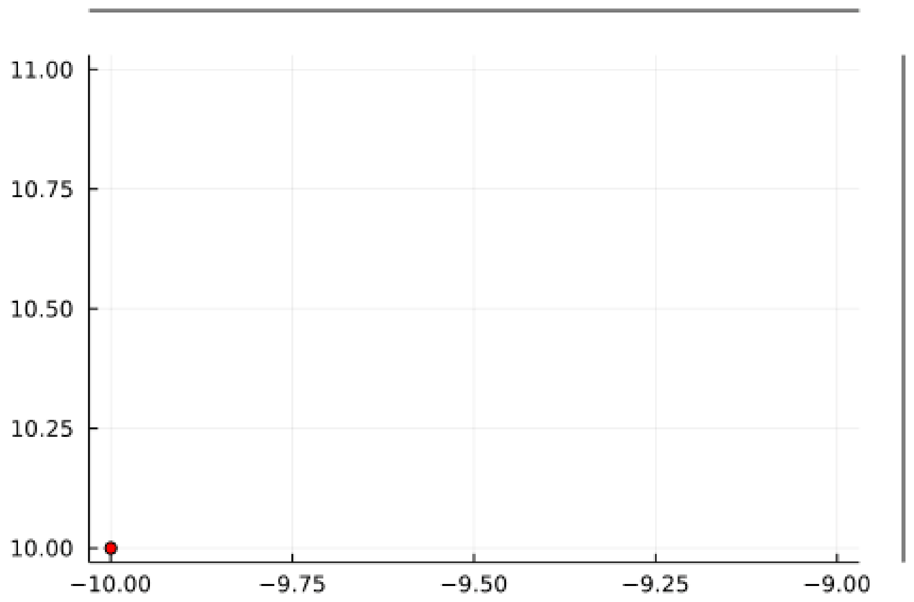

Motivation

- Real-world systems have inherent uncertainties that are best represented by stochastic systems;
- A natural setting to capture the uncertainty-aware ML-driven developments: developed in data-driven environments;
- Can bridge the design space and the test space, as the synthesis solutions get validated in Monte-Carlo environments, results are distributional in nature;
- Many robotic solutions today are developed using the language of stochastic optimal control, and hence the benchmarking and comparisons are easier to do.

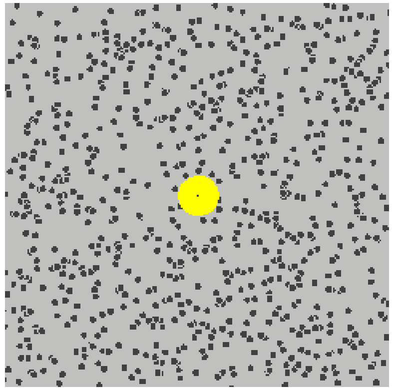

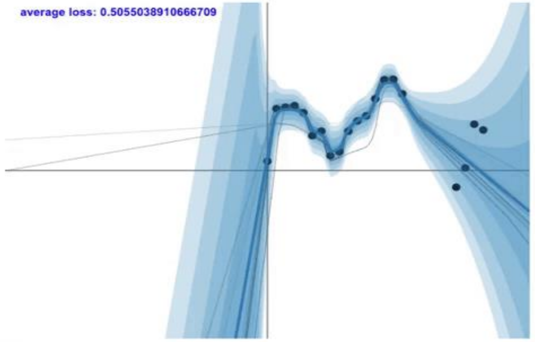

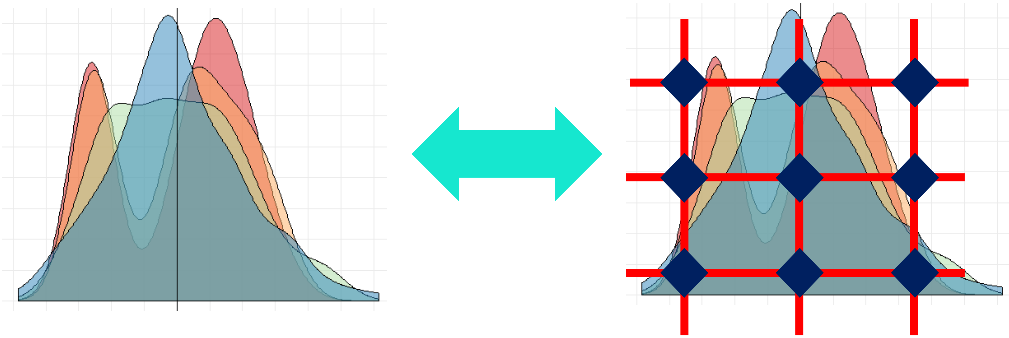

The Systems

Most processes evolve in a nonlinear fashion and are subject to random perturbations

- E.g. wind disturbances, thermal effects

$$dX = [f(X) + g(X)U]dt + p(X)d\boxed{W} \longleftarrow \text{Brownian Motion}$$

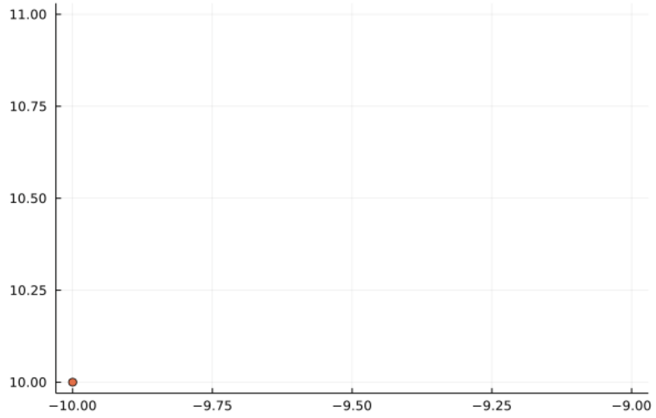

Nonlinear process continuously perturbed by random disturbance: Itô Diffusion Process (SDE)

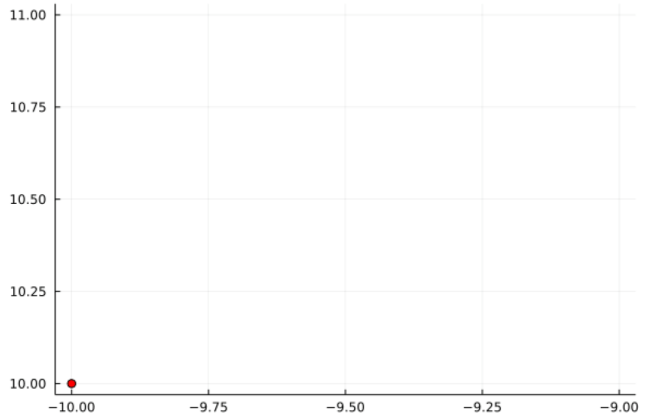

- However, we do not have full understanding of the vector field itself
  - Epistemic uncertainties
- Thus, we consider the following uncertain SDE

$$dX = [f(X) + g(X)U + h(X)] + [p(X) + q(X)]dW$$

Unknowns → Epistemic

The Systems

We consider the following true dynamics given by a nonlinear Ito diffusion SDE

$$dX_t = [f(X_t) + g(X_t)U_t + h(X_t)]dt + [p(X_t) + q(X_t)]dW_t, \quad X_t \sim \nu_t$$

Unknowns → Epistemic

Brownian Motion → Aleatoric

True State Distribution

We consider the following epistemic uncertainty-free nominal dynamics

$$dX_t^* = [f(X_t^*) + g(X_t^*)U_t^*]dt + [p(X_t^*)]dW_t^*, \quad X_t^* \sim \nu_t^*$$

Nominal State Distribution

Robustness Paradigm

- For linear, parameterized and deterministic models the error in the estimates of the parameters leads to the difference between the true solutions/trajectories and the nominal/learned trajectories

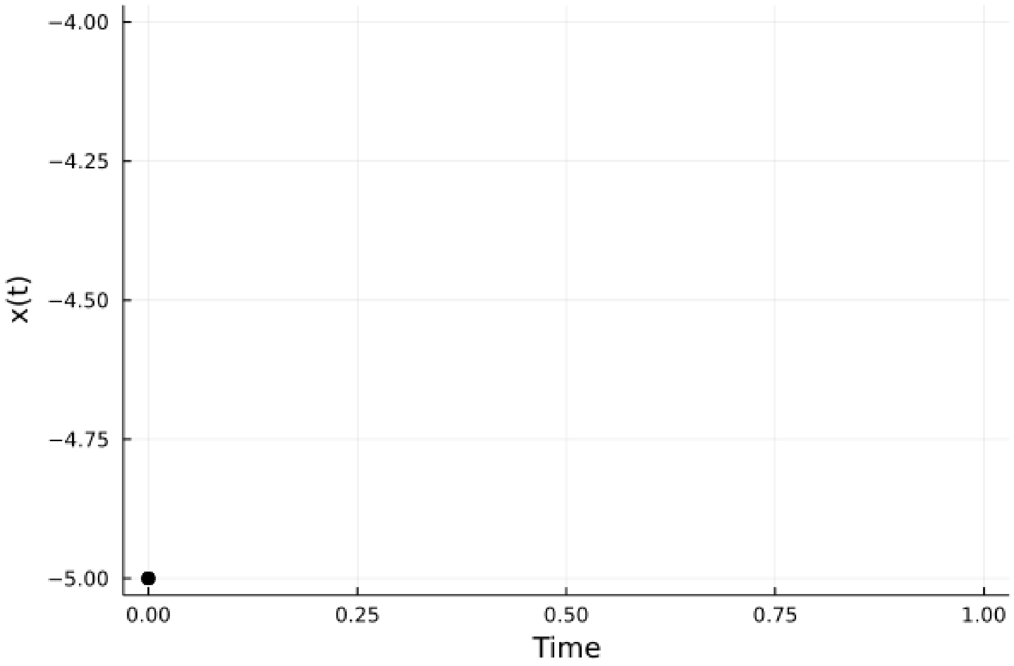

- Robustness to the estimates of the parameters produces certificates quantifying the difference between true and nominal solutions

Robustness Paradigm

- Looking at the solutions of nonlinear stochastic models is not helpful since even the same SDE will produce different solutions each time it is evolved: property of stochastic processes.

- Instead, we should look at the time-evolving distributions, the samples of which are the solutions themselves.

- We can say that the two SDEs are equivalent if the temporal trajectories of their respective distributions are the same.
    - The solutions are the samples from the same distribution.

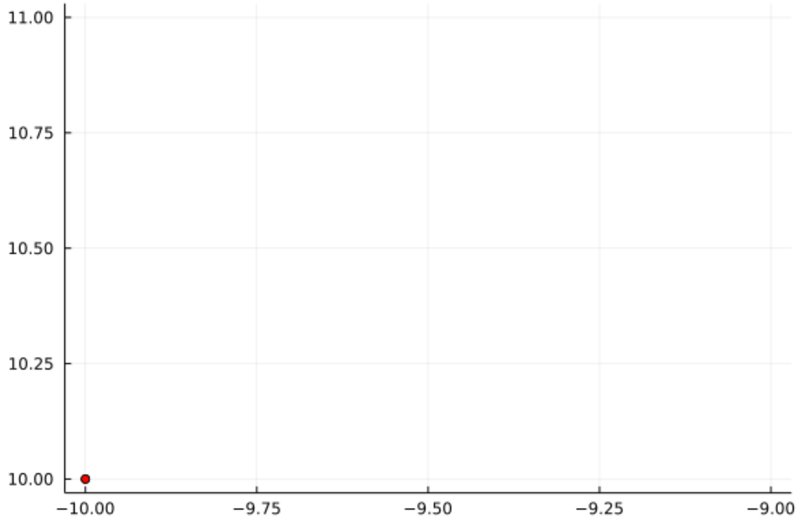

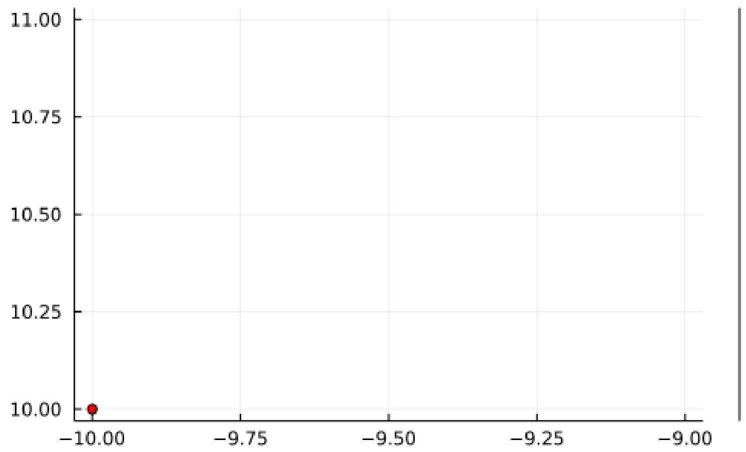

Distributional Robustness

- Safe use of machine learning
  - Safe predictive control

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

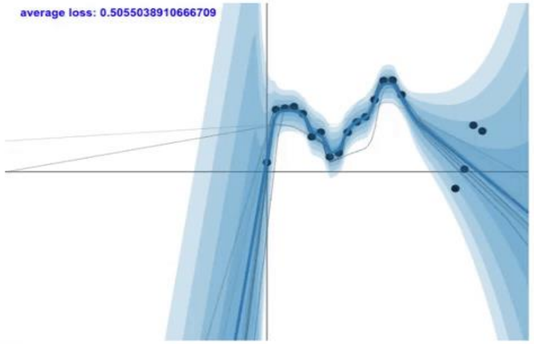

- Design principles guided by distributional guarantees independently verified by e.g. Monte-Carlo methods
  - Design space = Test space
  - Easier feedback between the spaces

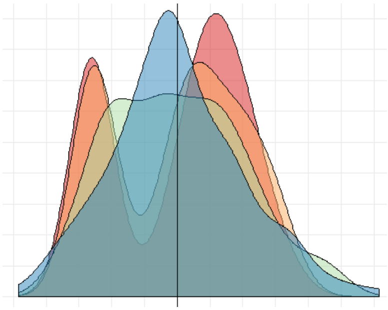

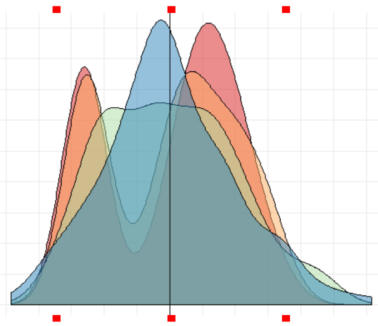

The Goals

True $dX_t = [f(X_t) + g(X_t)U_t + h(X_t)]dt + [p(X_t) + q(X_t)]dW_t, \quad X_t \sim \nu_t$

Nominal $dX_t^* = [f(X_t^*) + g(X_t^*)U_t^*]dt + [p(X_t^*)]dW_t^*, \quad X_t^* \sim \nu_t^*$

- We want to design feedback $U_t = K(X_t, X_t^*, U_t^*)$ such that
- True distribution ν_t remains uniformly bounded around the nominal distribution ν_t^*
  - Uniform bound incorporated by high-level planner to ensure safe operation
- Bound in the sense of Wasserstein metric
  - Optimal transport theory
  - A metric on the space of distributions (distance and shape)

The Goals: Pictorial Depiction

Ambiguity tube $\Omega(\rho)$: induced by ambiguity sets $\mathcal{A}(\rho, \nu_t^\star)$

- For each $t \geq 0$ $\underbrace{\mathbb{W}_2(\nu_t, \nu_t^\star) \leq \rho}_{\text{Wasserstein Distance}} \Rightarrow \underbrace{\nu_t \in \mathcal{A}(\rho, \nu_t^\star)}_{\text{Ambiguity Set}}$

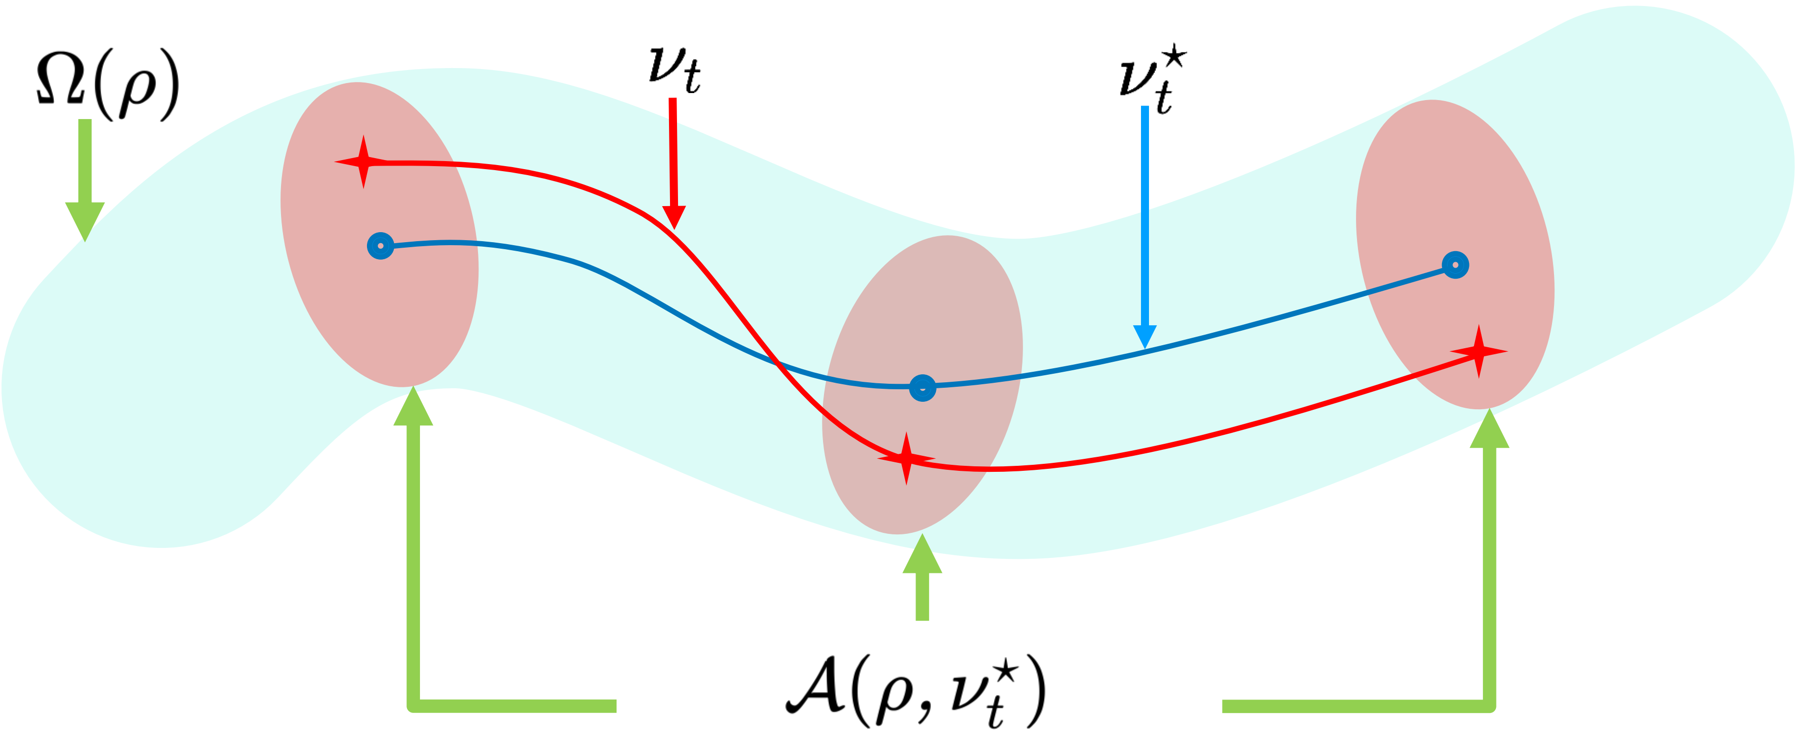

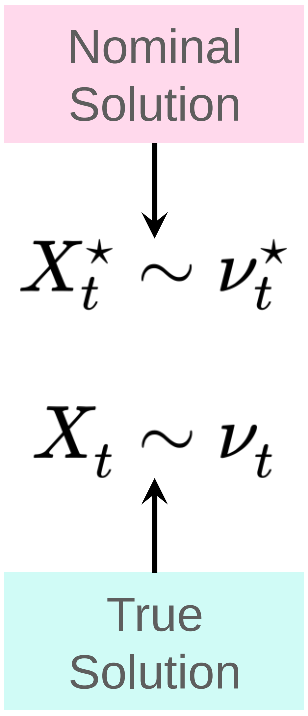

Controller

True

$$dX_t = [f(X_t) + g(X_t)(U_t + h(X_t))]dt + [p(X_t) + g(X_t)q(X_t)]dW_t, \quad X_t \sim \nu_t$$

Nominal

$$dX_t^* = [f(X_t^*) + g(X_t^*)U_t^*]dt + [p(X_t^*)]dW_t^*, \quad X_t^* \sim \nu_t^*$$

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

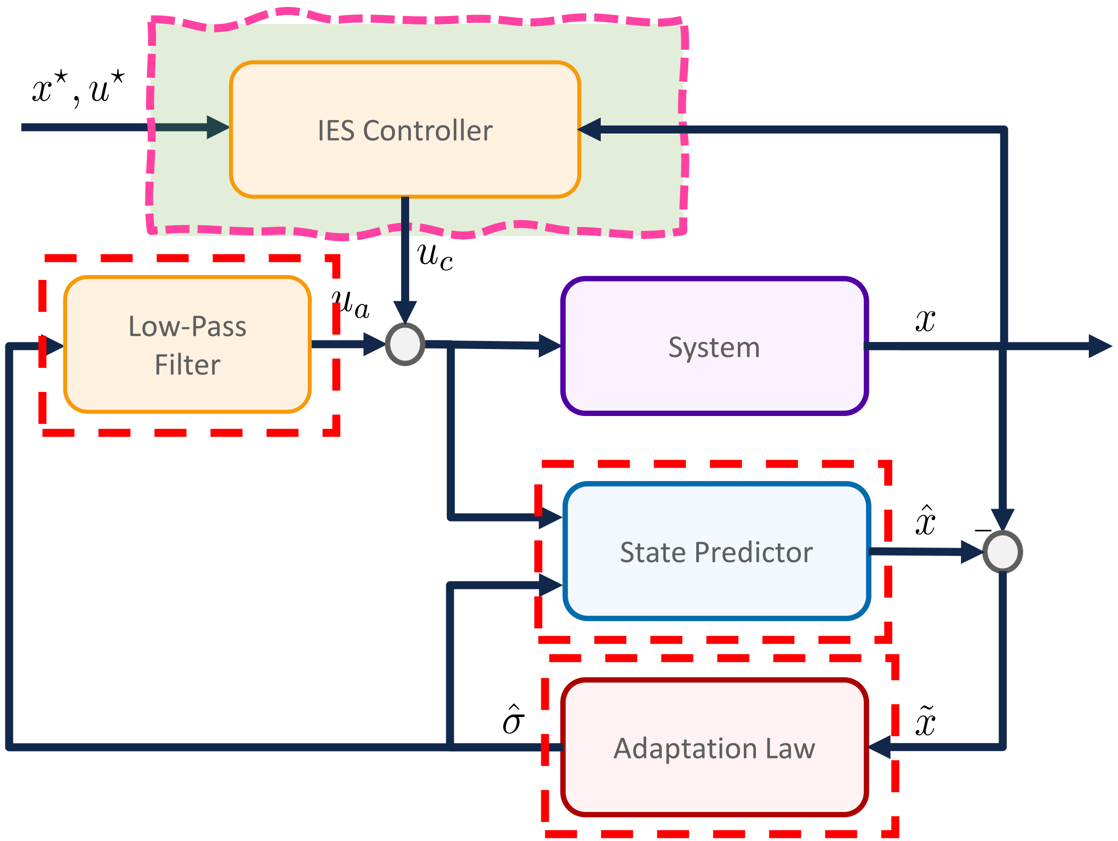

Simulation Results

We consider the stochastic version of a feedback linearizable system from [1]

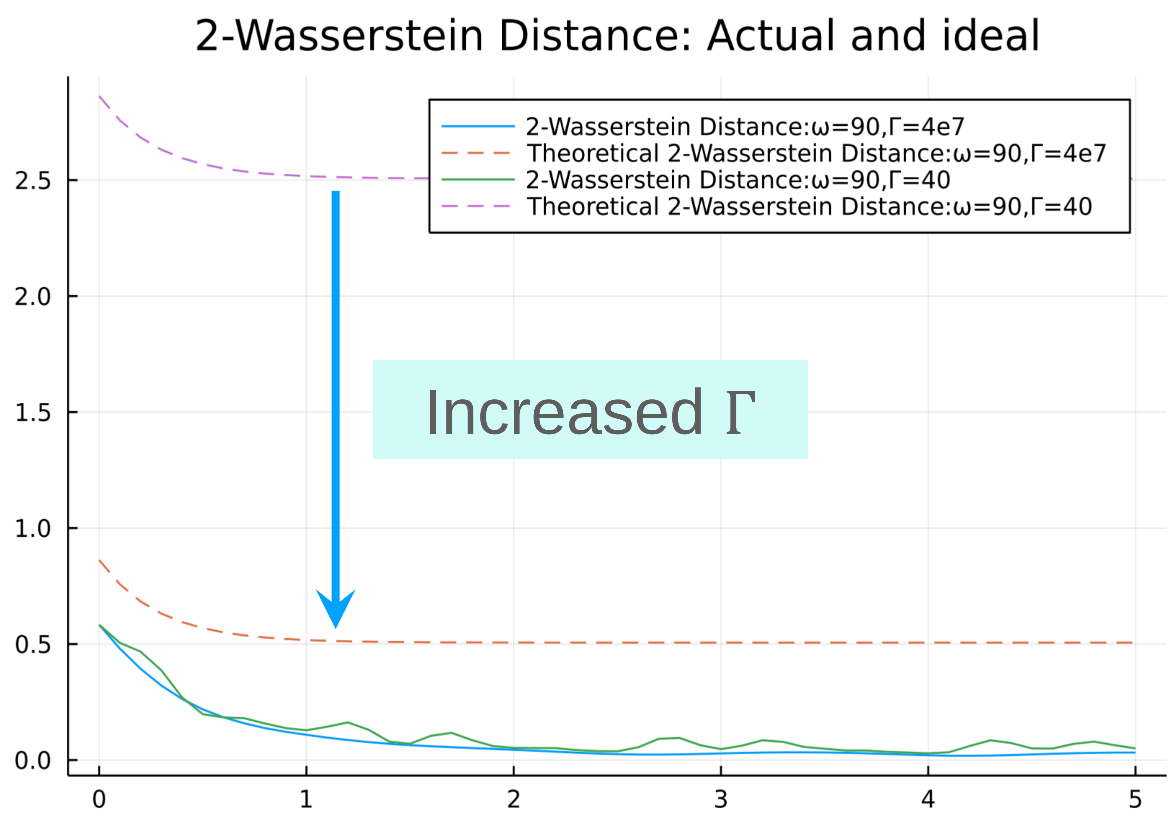

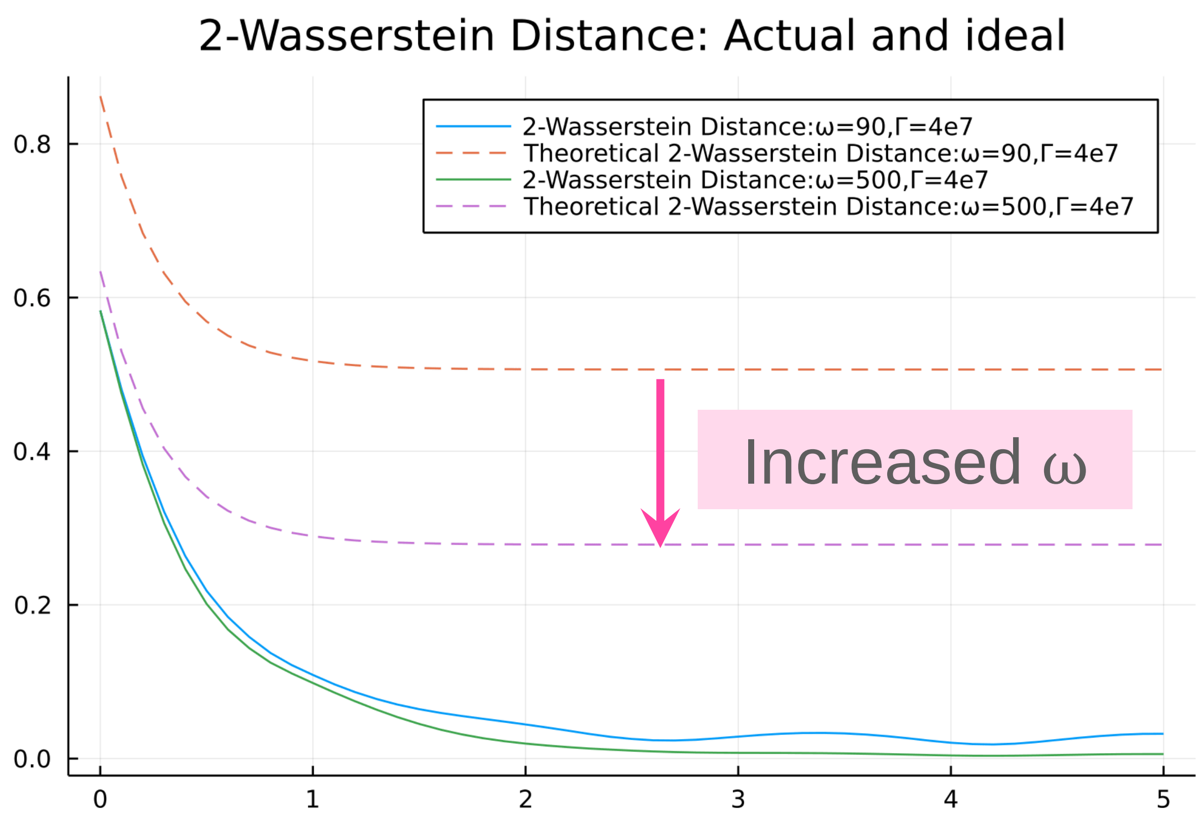

- Increasing the adaptation-rate Γ reduces the bounds
- Similarly, increasing the filter-bandwidth ω reduces the bounds

[1] Lakshmanan, Gahlawat, and Hovakimyan. *IEEE CDC*, 2020.

ILF-$\mathcal{L}_1$ for Autonomous Drones

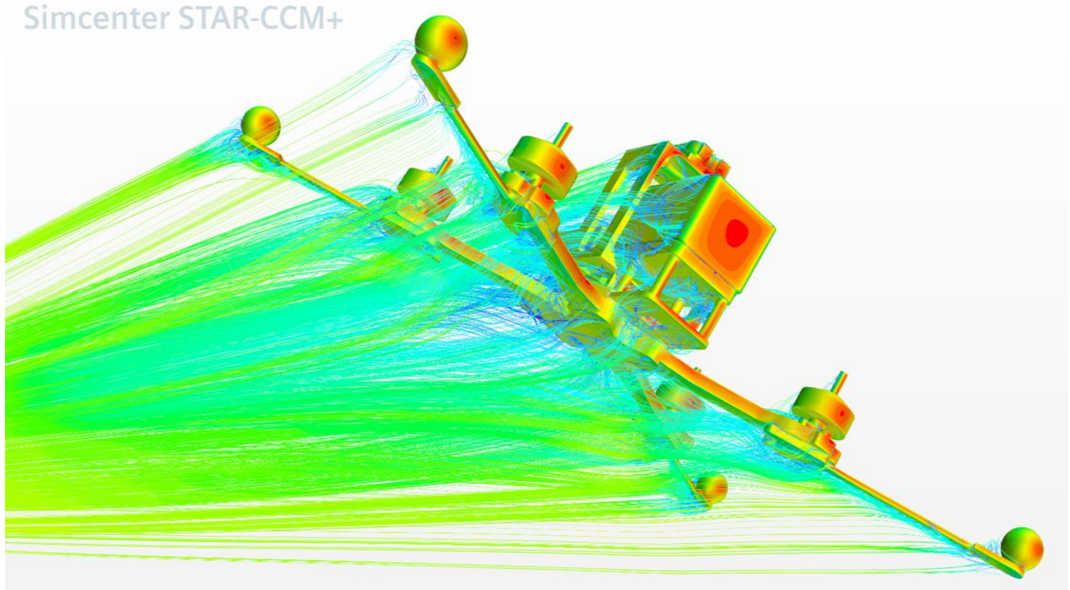

Challenges

- Underactuated dynamics

- Coupled translational and rotational motions

- Uncertainties and disturbances
  - Aerodynamic drag
  - Varying payload/moment of inertia/center of gravity
  - Wind

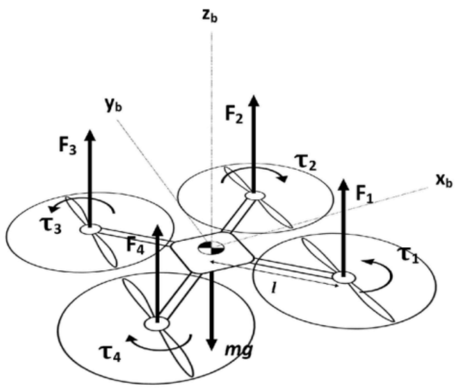

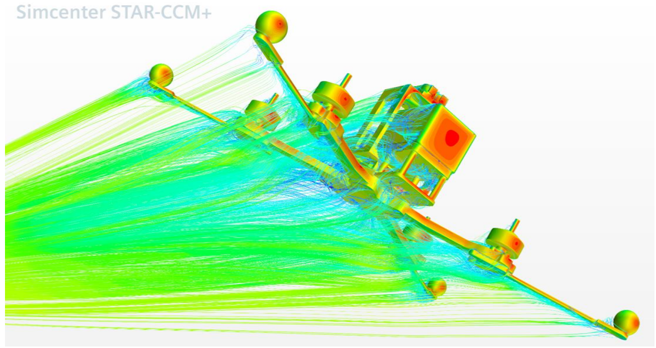

Geometric Control + $\mathcal{L}_1$ Augmentation

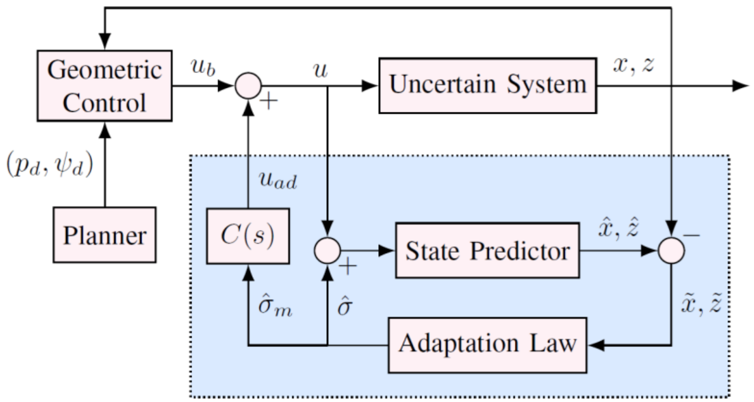

Z. Wu, S. Cheng, K. A. Ackerman, A. Gahlawat, A. Lakshmanan, P. Zhao and N. Hovakimyan, "$\mathcal{L}_1$ Adaptive Augmentation for Geometric Tracking Control of Quadrotors", accepted by ICRA 2022.

- Modelling all possible uncertainties and disturbances using first principles
- $\mathcal{L}_1$ adaptive control compensates for the uncertainties and disturbances
  - Safety guaranteed in the sense of tubes around desired trajectories
- Low-level control: All computations are onboard (440 Hz)
- Experiments on a custom-built quadrotor

Ground Effects

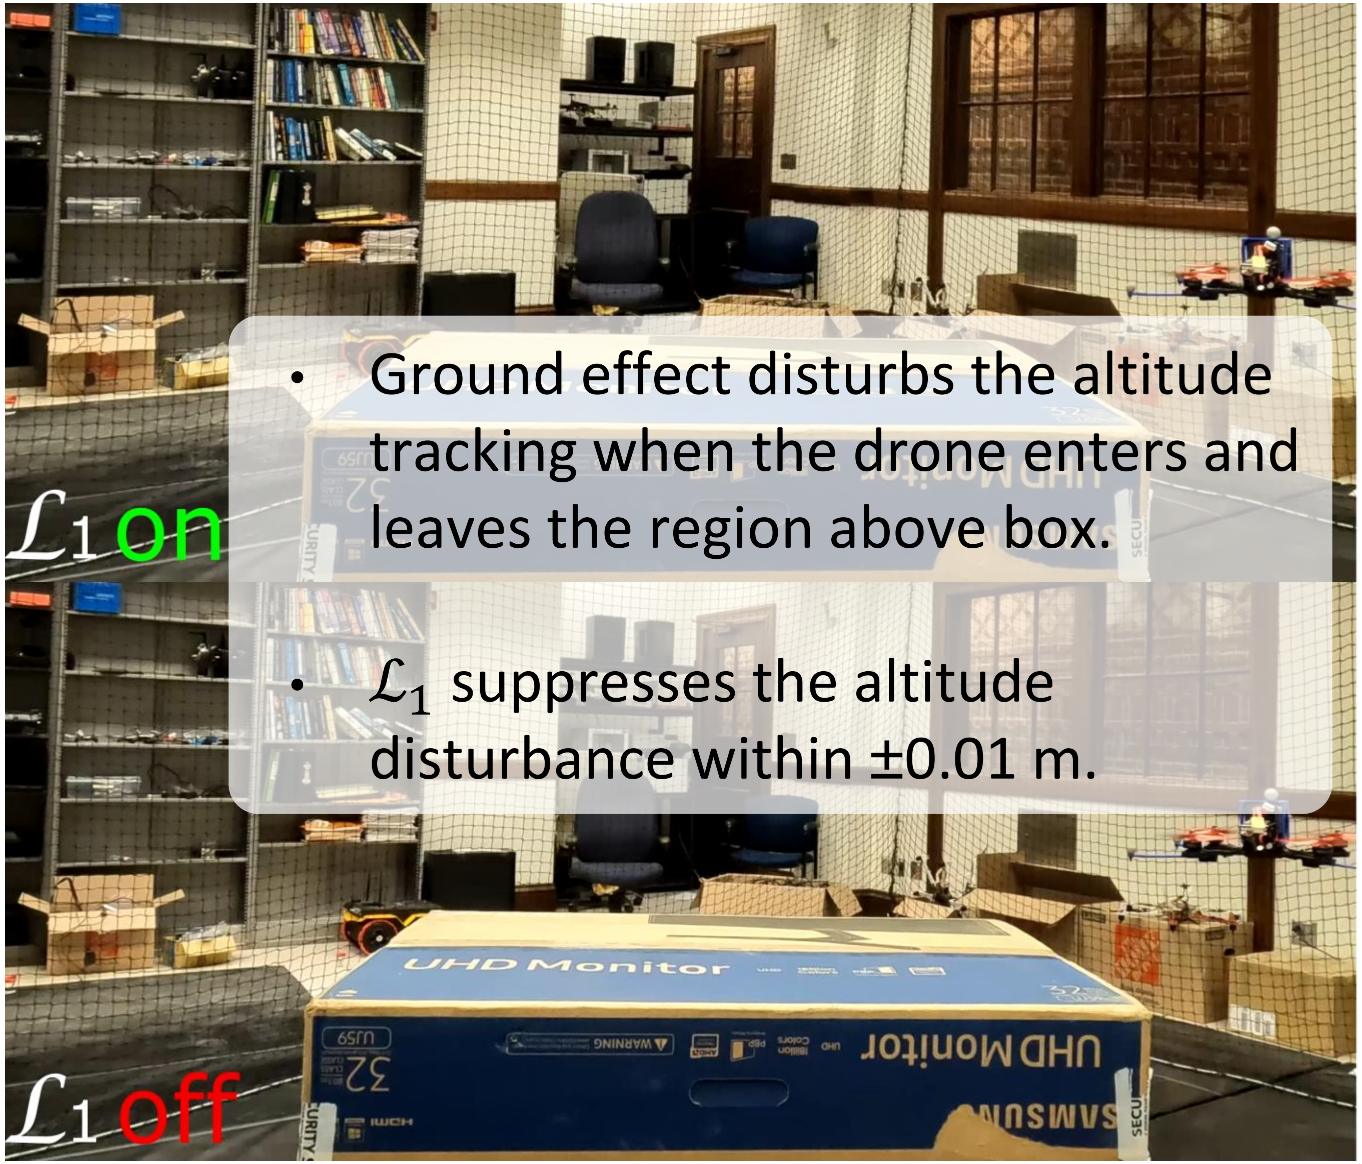

- Ground effect disturbs the altitude tracking when the drone enters and leaves the region above box.
- $\mathcal{L}_1$ suppresses the altitude disturbance within ±0.01 m.

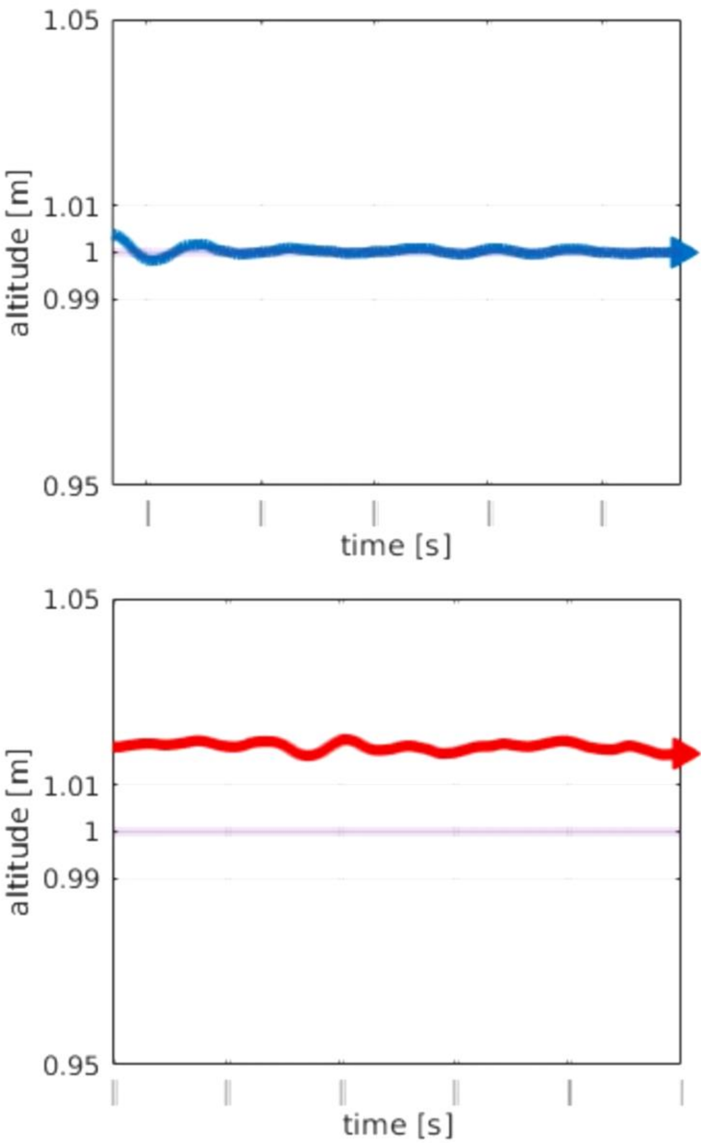

Slosh payload

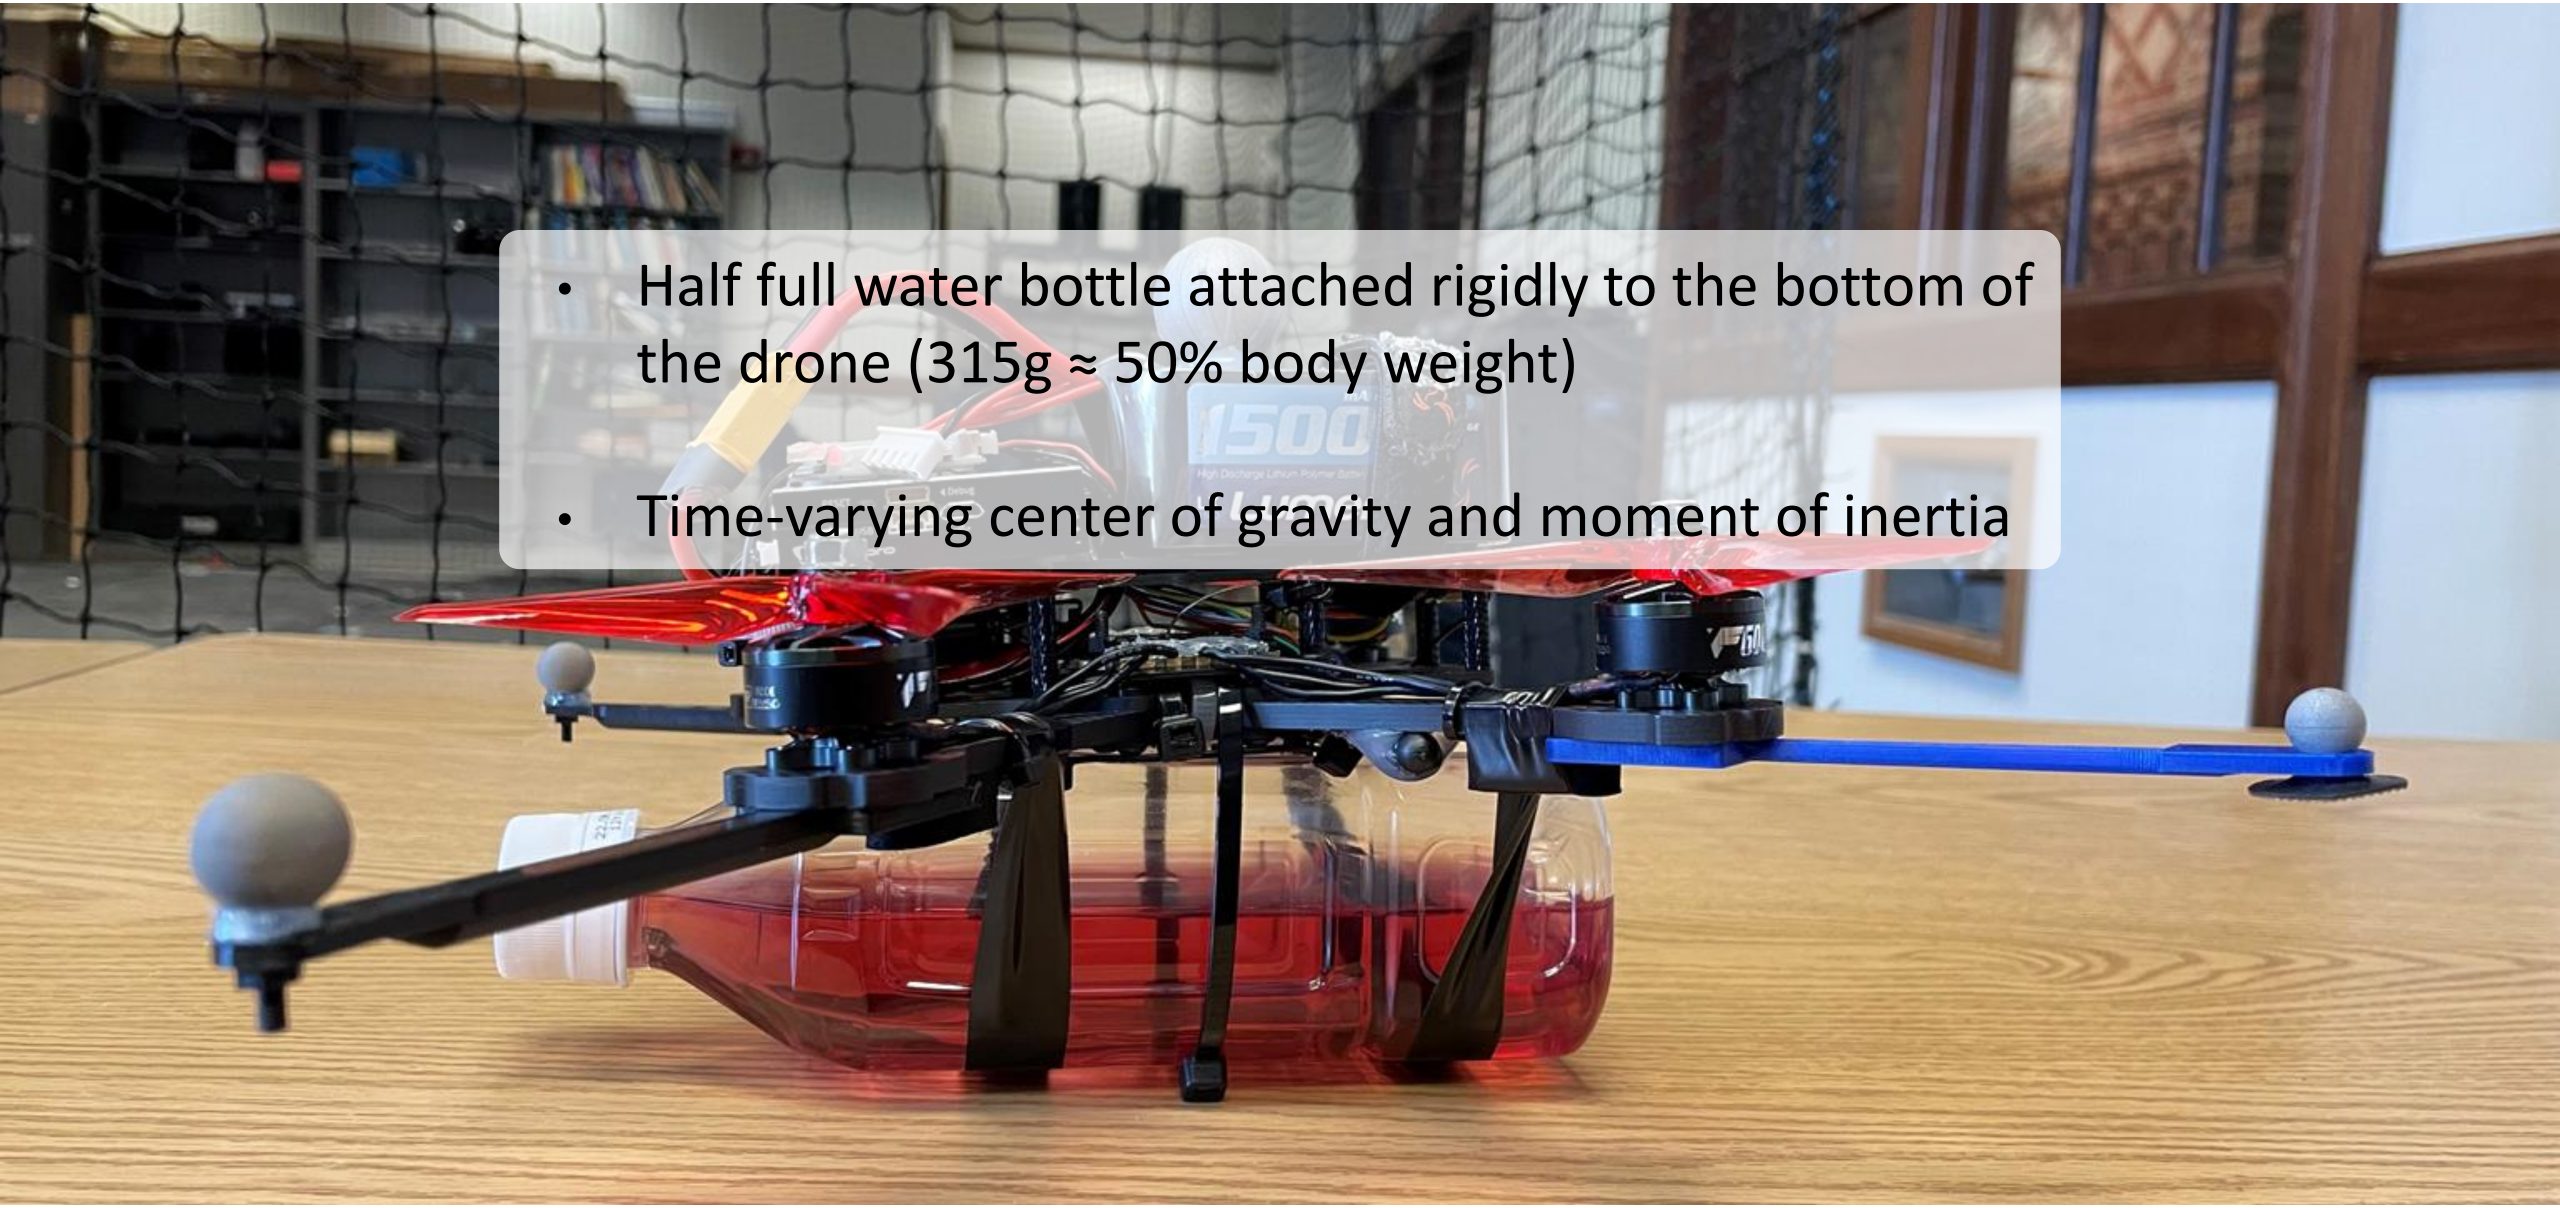

- Half full water bottle attached rigidly to the bottom of the drone (315g ≈ 50% body weight)

- Time-varying center of gravity and moment of inertia

Slosh payload

- Half full water bottle attached rigidly to the bottom of the drone (315g ≈ 50% body weight)
- Time-varying center of gravity and moment of inertia

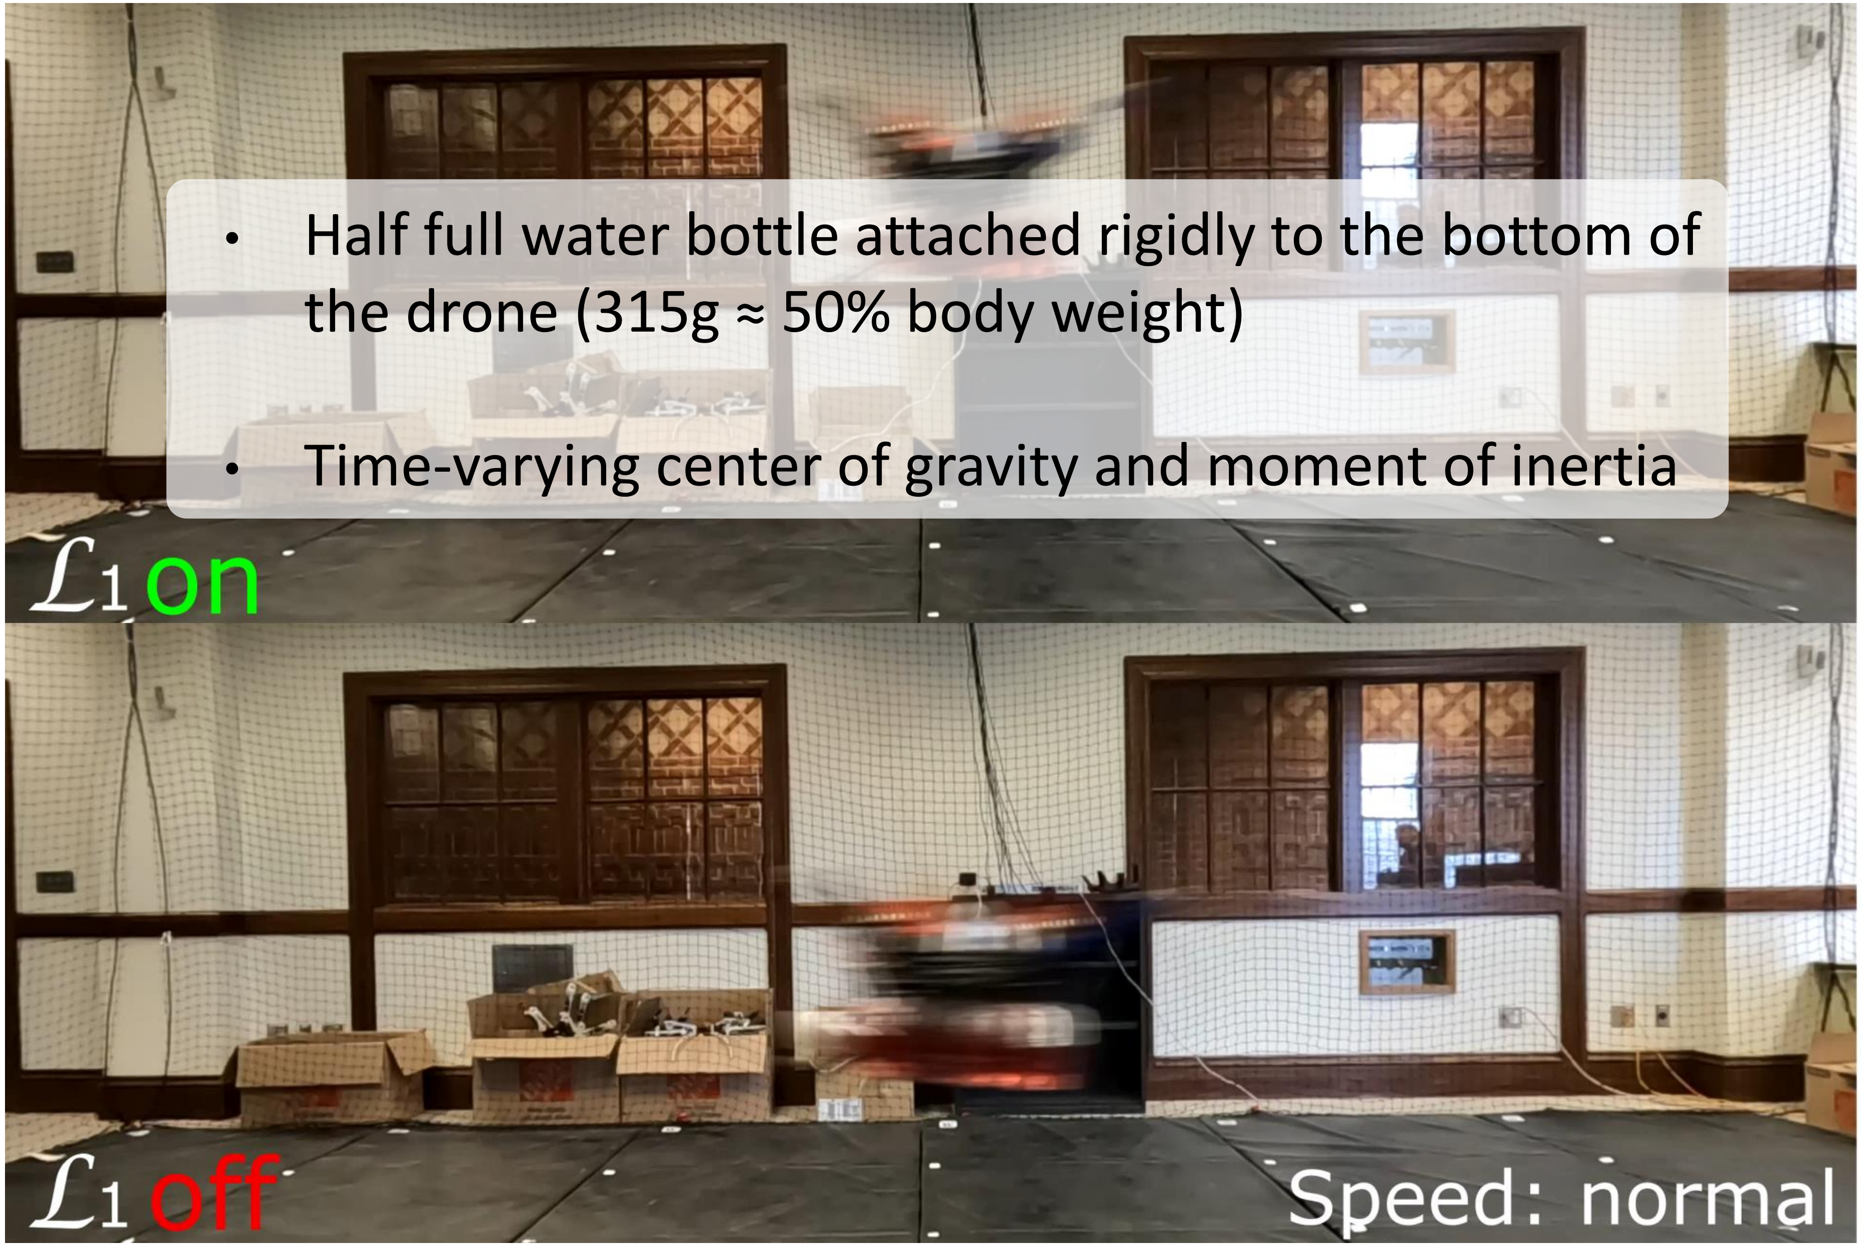

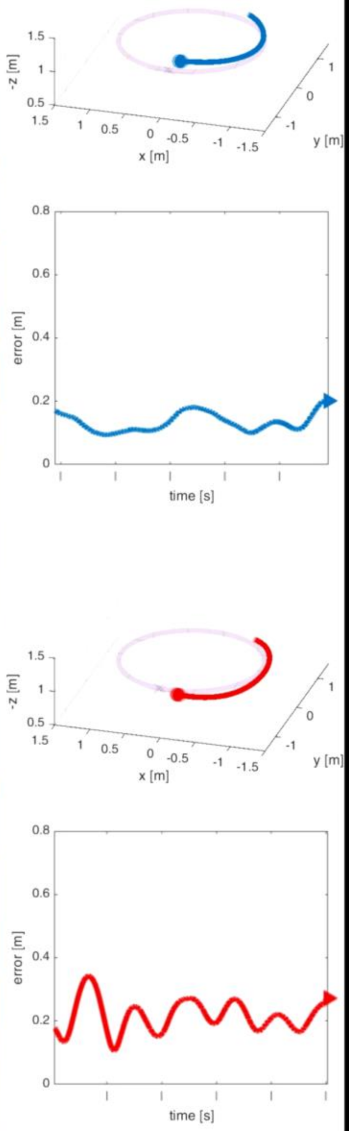

Chipped Propeller

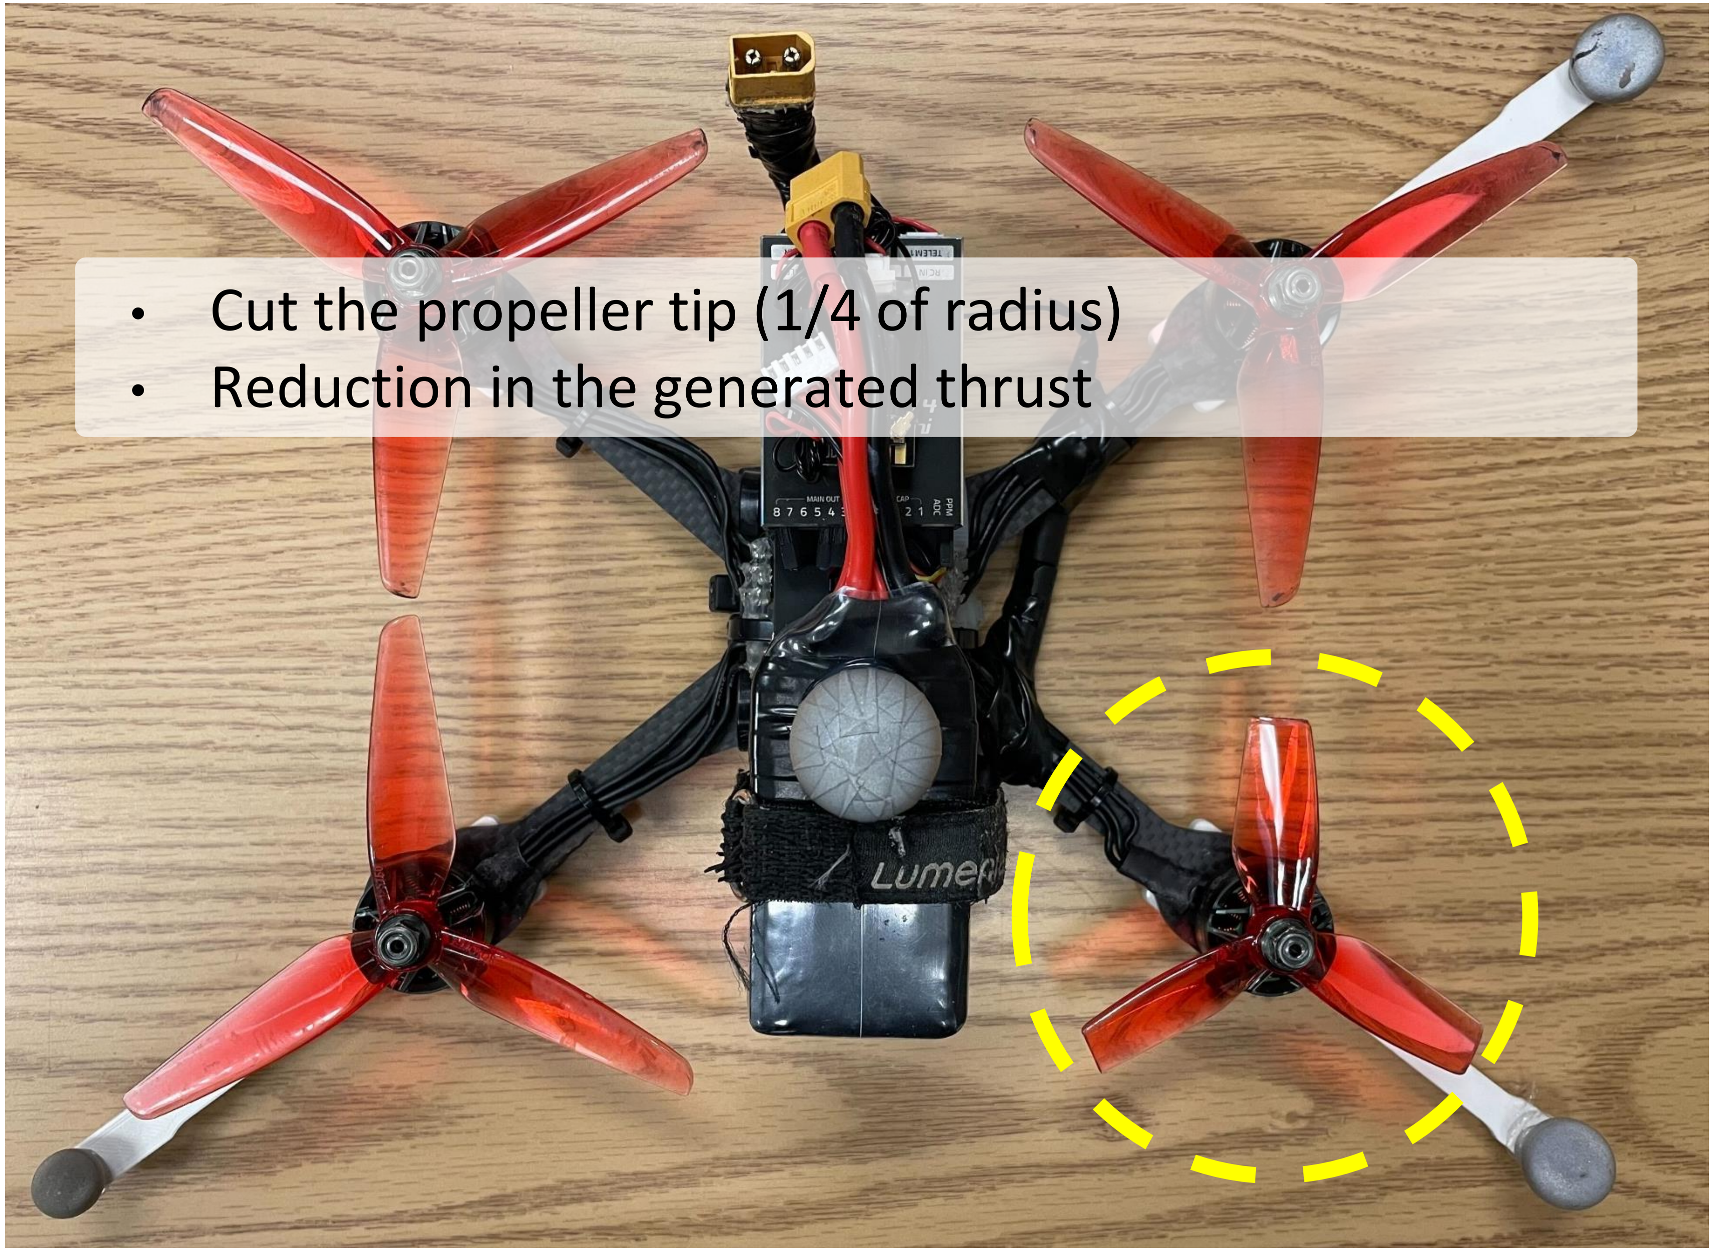

- Cut the propeller tip (1/4 of radius)
- Reduction in the generated thrust

Chipped Propeller

- Cut the propeller tip (1/4 of radius)
- Reduction in the generated thrust
- Baseline controller cannot even stabilize the drone.
- $\mathcal{L}_1$ keeps the drone stabilized.

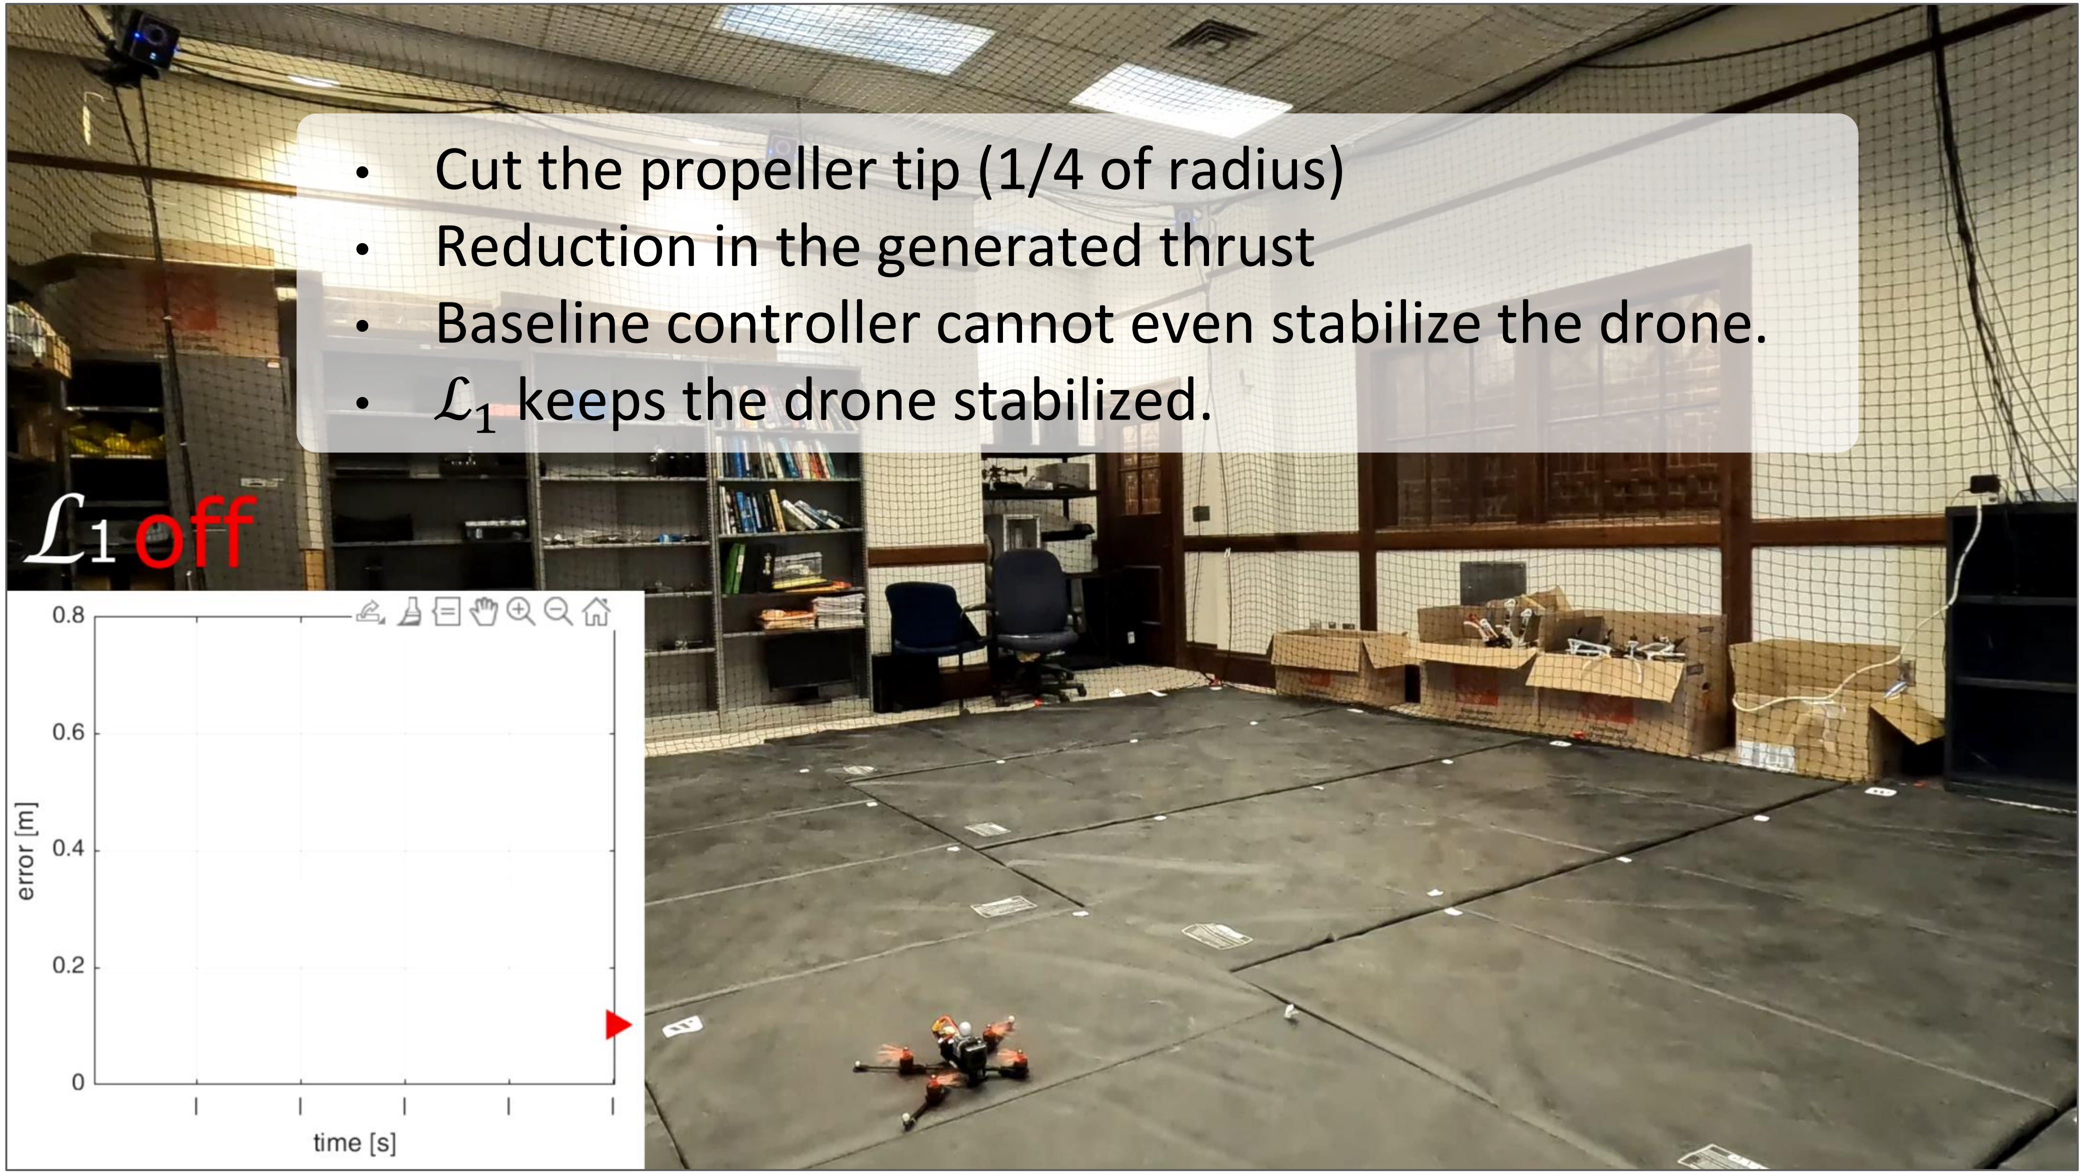

Downwash

- Downwash from a heavier drone hovering 0.6 m on top
- $\mathcal{L}_1$ shows advantageous transient response compared with PID.

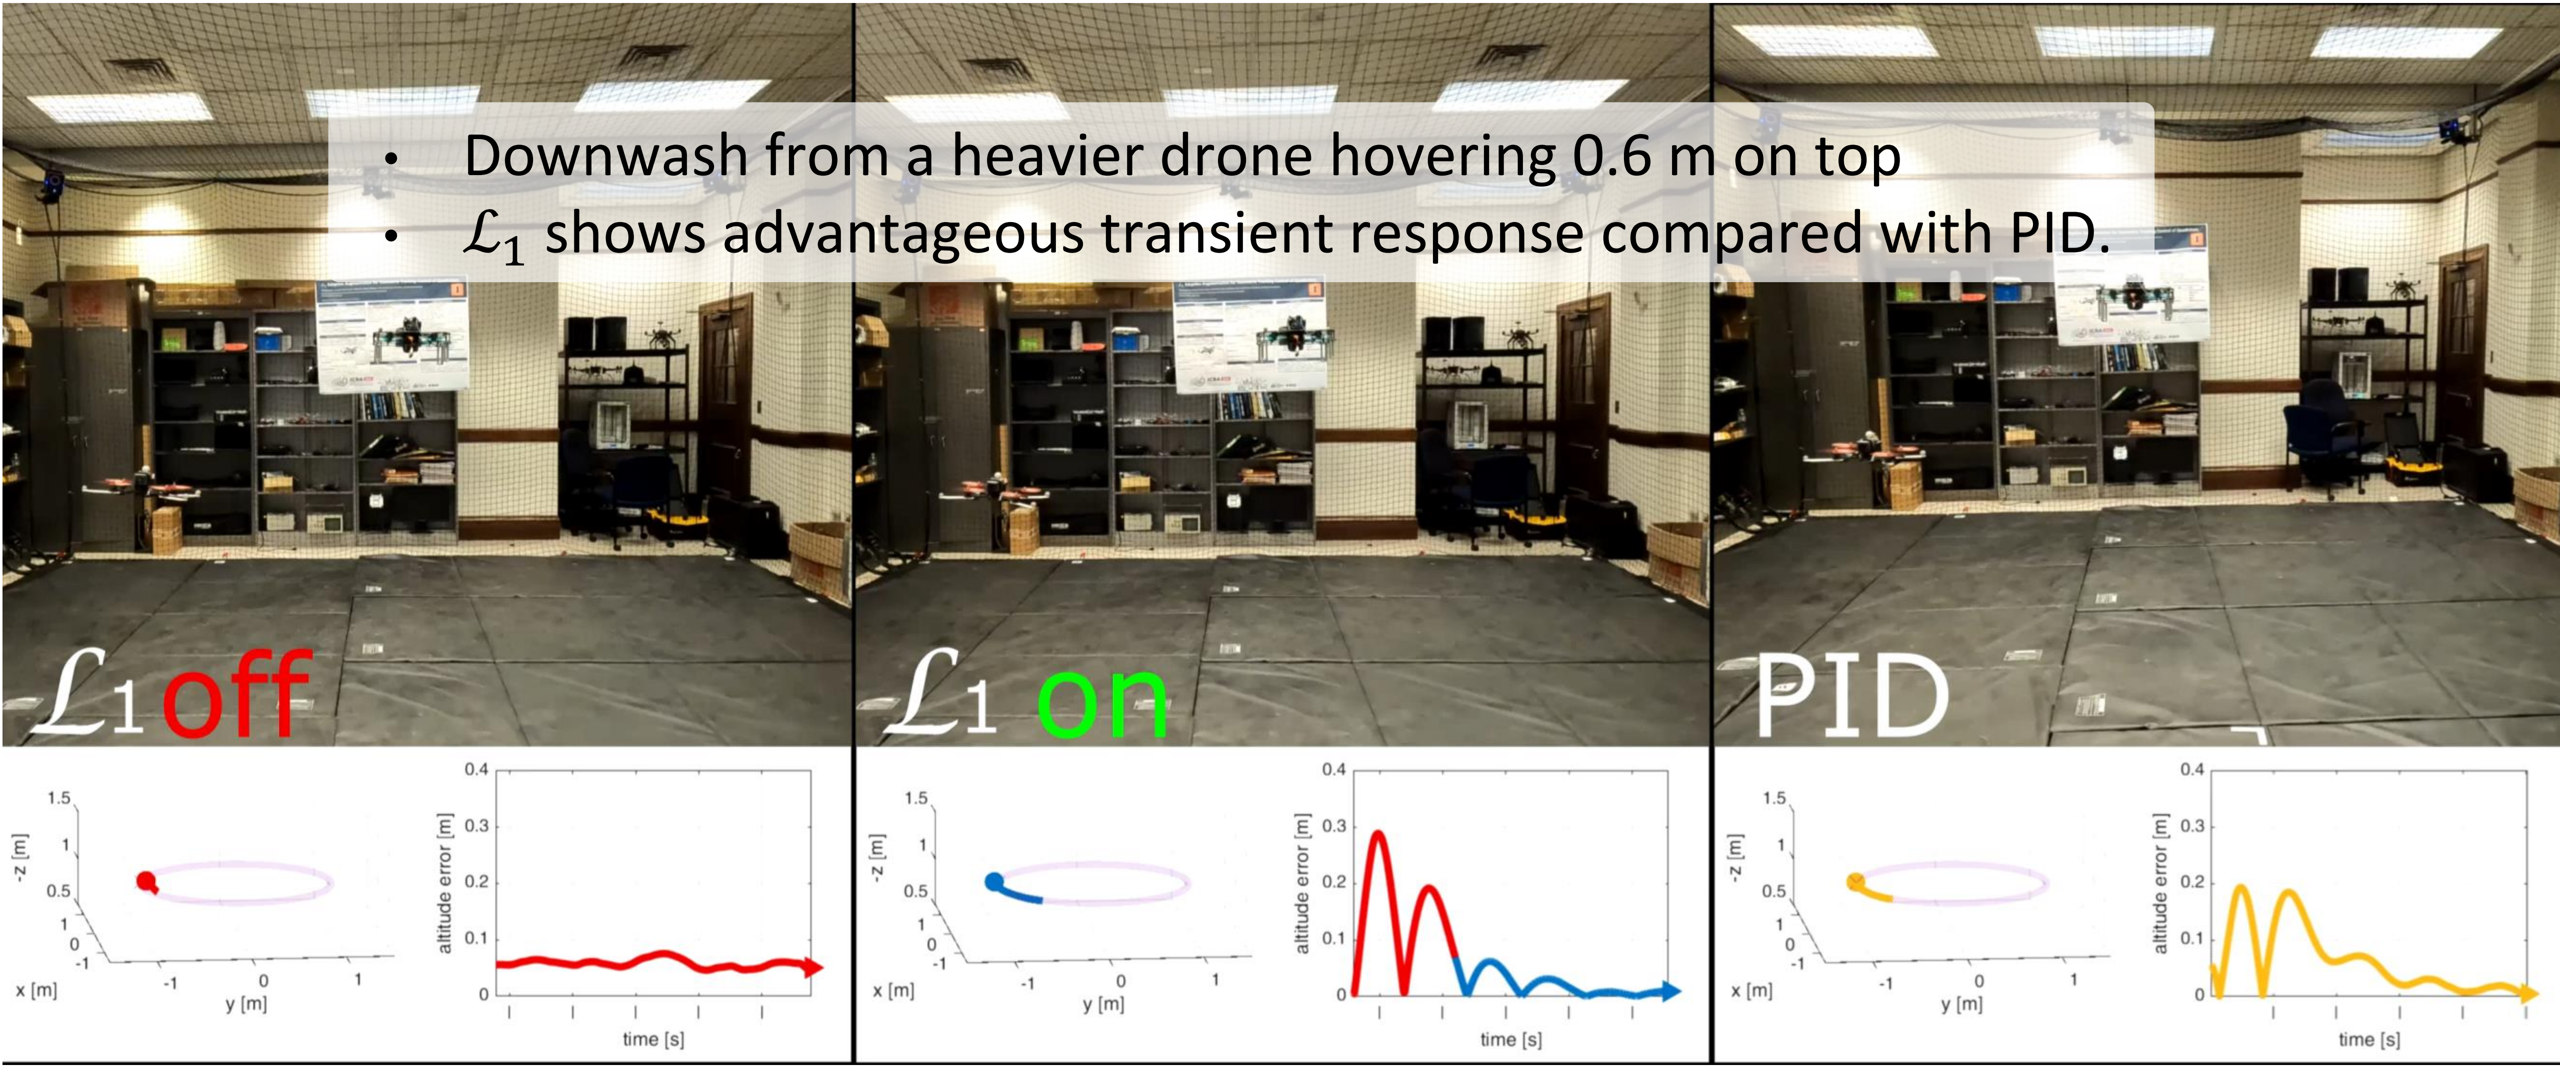

Quadrotor on top: 1400 g | Quadrotor at bottom: 640 g

Fly in a Tunnel

- Tunnel size: 1.2 m (W) x 0.6 m (H) x 0.6 m (D)
- Complex air flow indicated by tufts
- $\mathcal{L}_1$ shows consistent altitude tracking compared with PID.

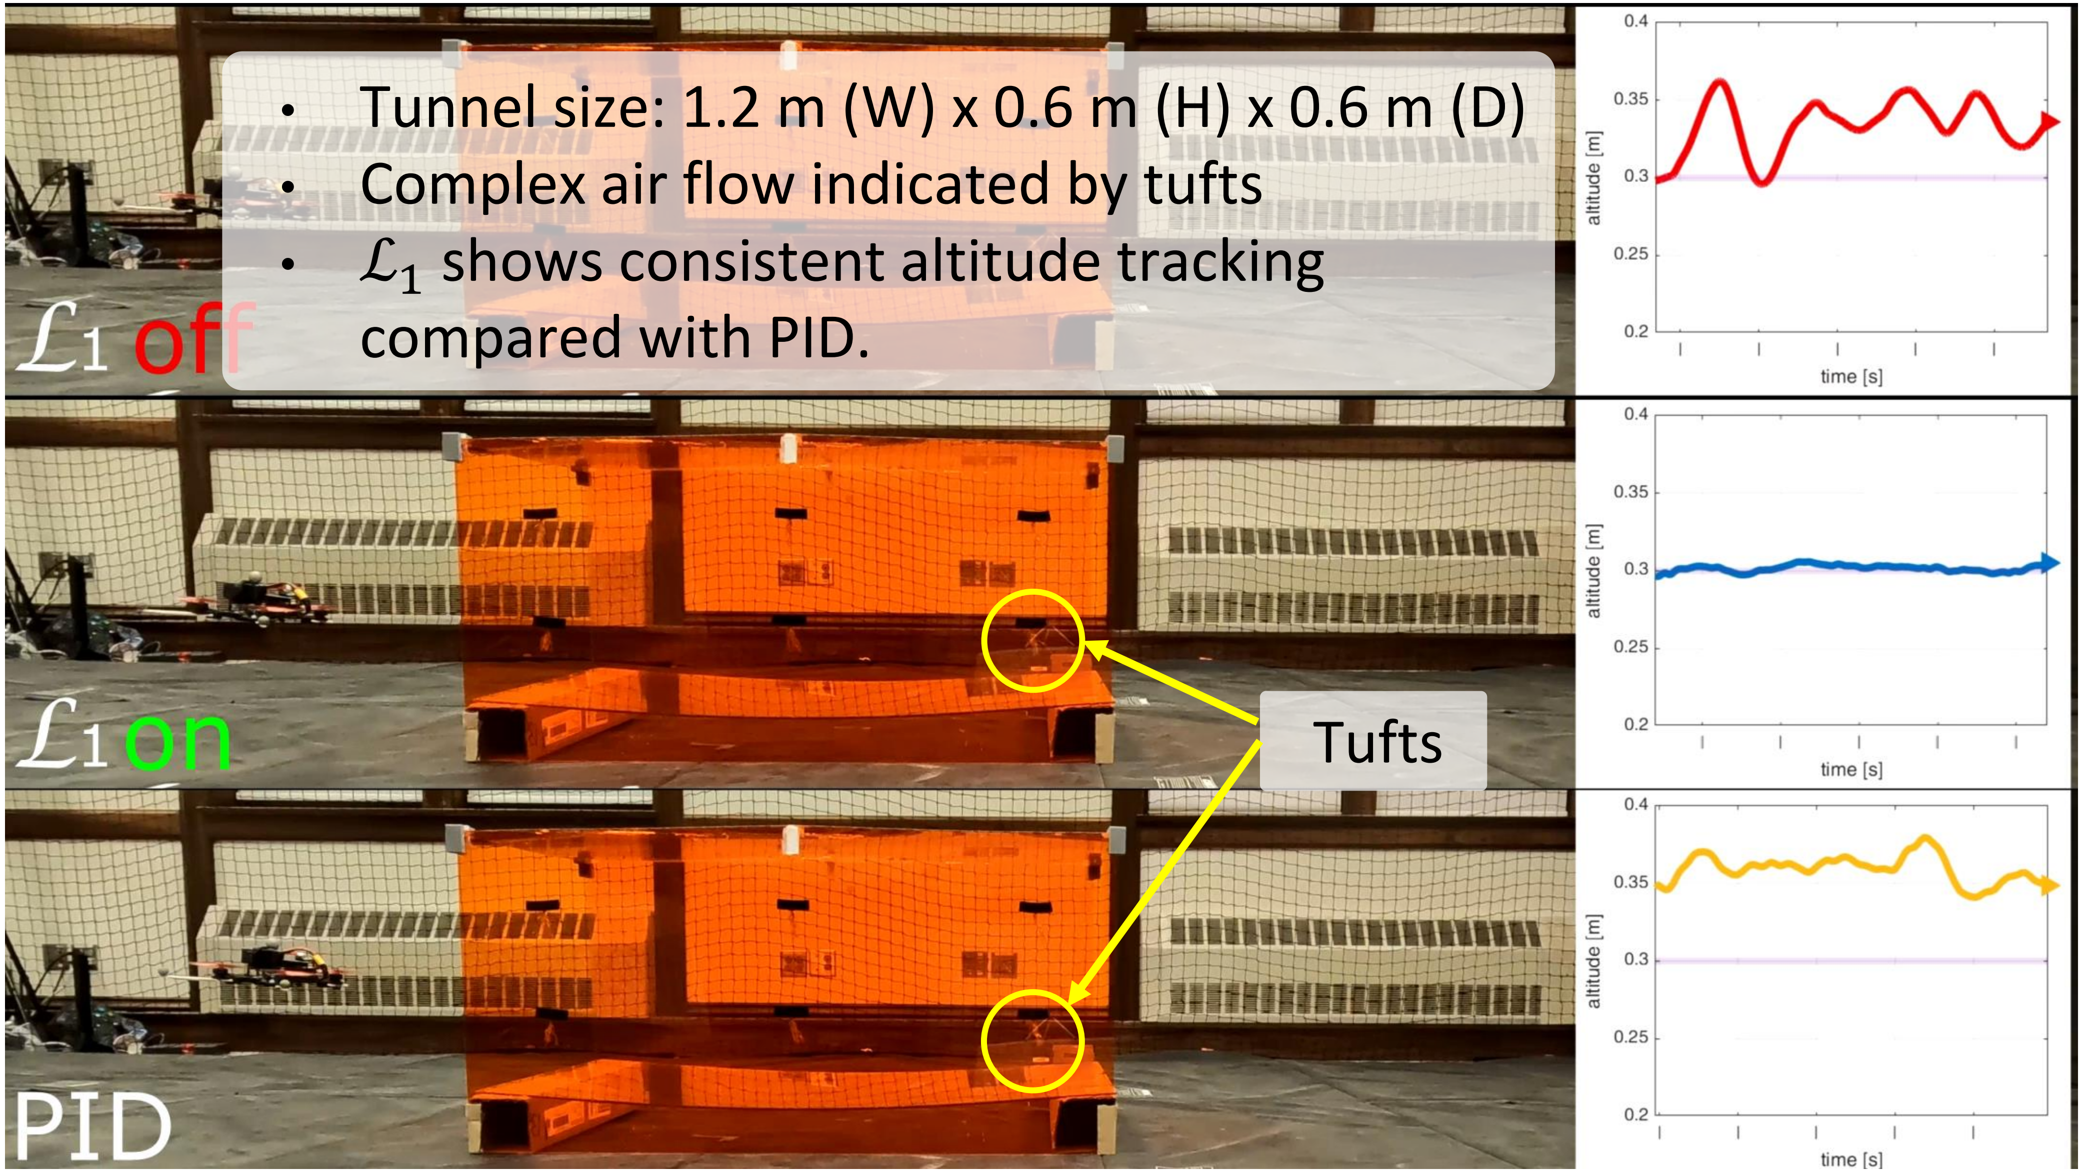

Hanging Off-Center Weights

- Hanging weights beneath the motor (off the drone’s geometric center)
- $\mathcal{L}_1$ shows the fastest transient performance compared with PID.

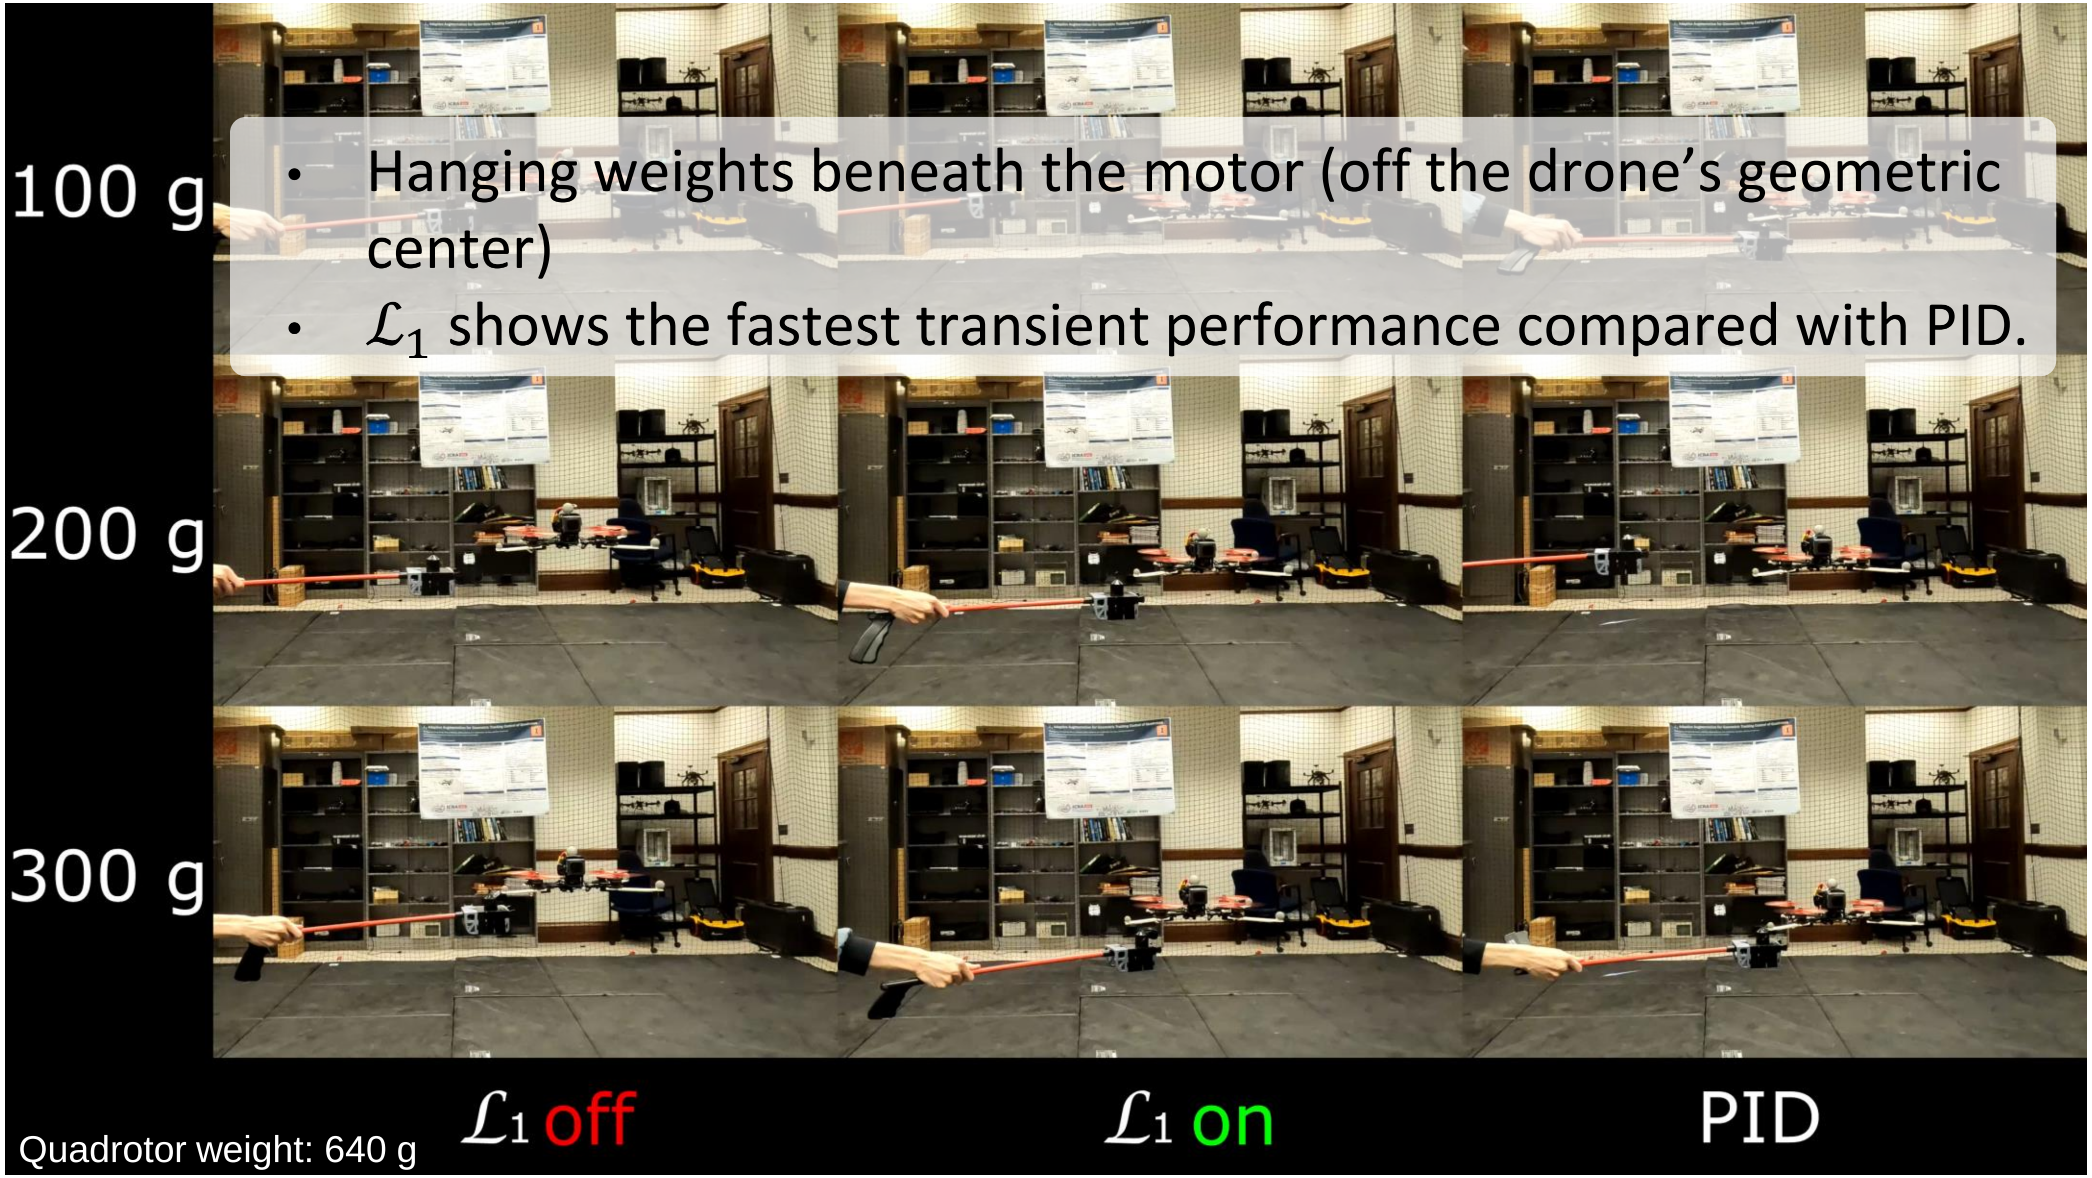

Benchmark Experiments with Slung Weights

Consistent performance without retuning $\mathcal{L}_1$ parameters.

× indicates failed trials due to instability.

Z. Wu*, S. Cheng*, P. Zhao, A. Gahlawat, K. A. Ackerman, A. Lakshmanan, C. Yang, J. Yu, and N. Hovakimyan, “L1Quad: L1 adaptive augmentation of geometric control for agile quadrotors with performance guarantees,” arXiv preprint arXiv:2302.07208, 2023.

NASA ULI on Assured Autonomy (2022 - 2025)

Robust and Resilient Autonomy for Advanced Air Mobility

Team: PI: Naira Hovakimyan, 15 Co-Is from 4 universities and 2 companies

LOCKHEED MARTIN

SIERRA NEVADA CORPORATION

Highlights

- Learning-enabled autonomy with principled ways to deal with off-nominal situations
- Fast code-level verification
- Runtime fault diagnosis and reachability analysis
- Flight tests at an FAA-designated UAS test site

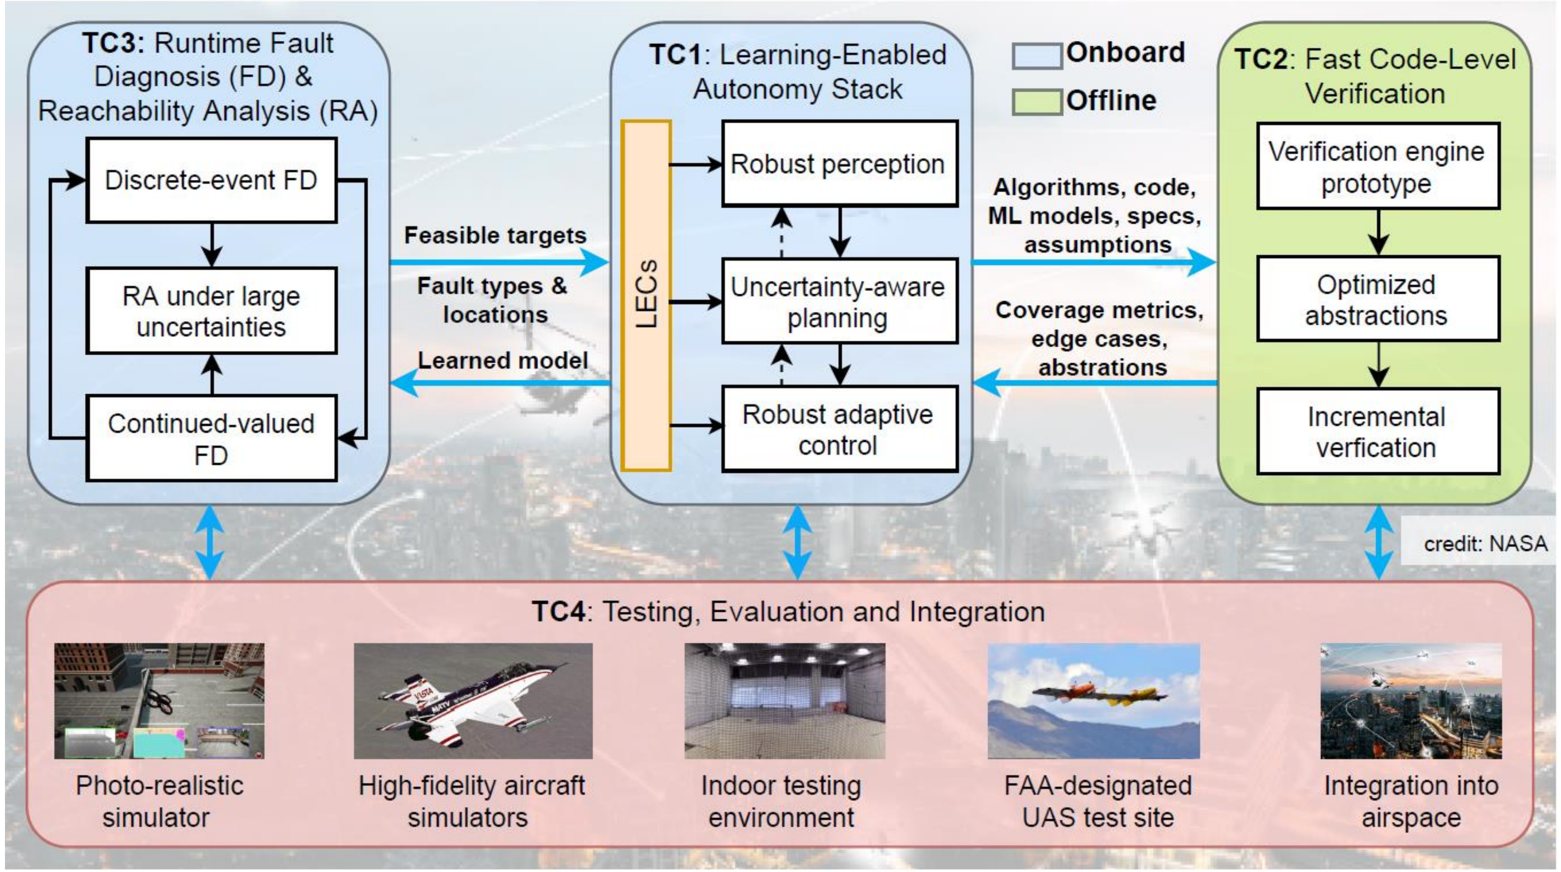

In conclusion

We wish to bridge the divide

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

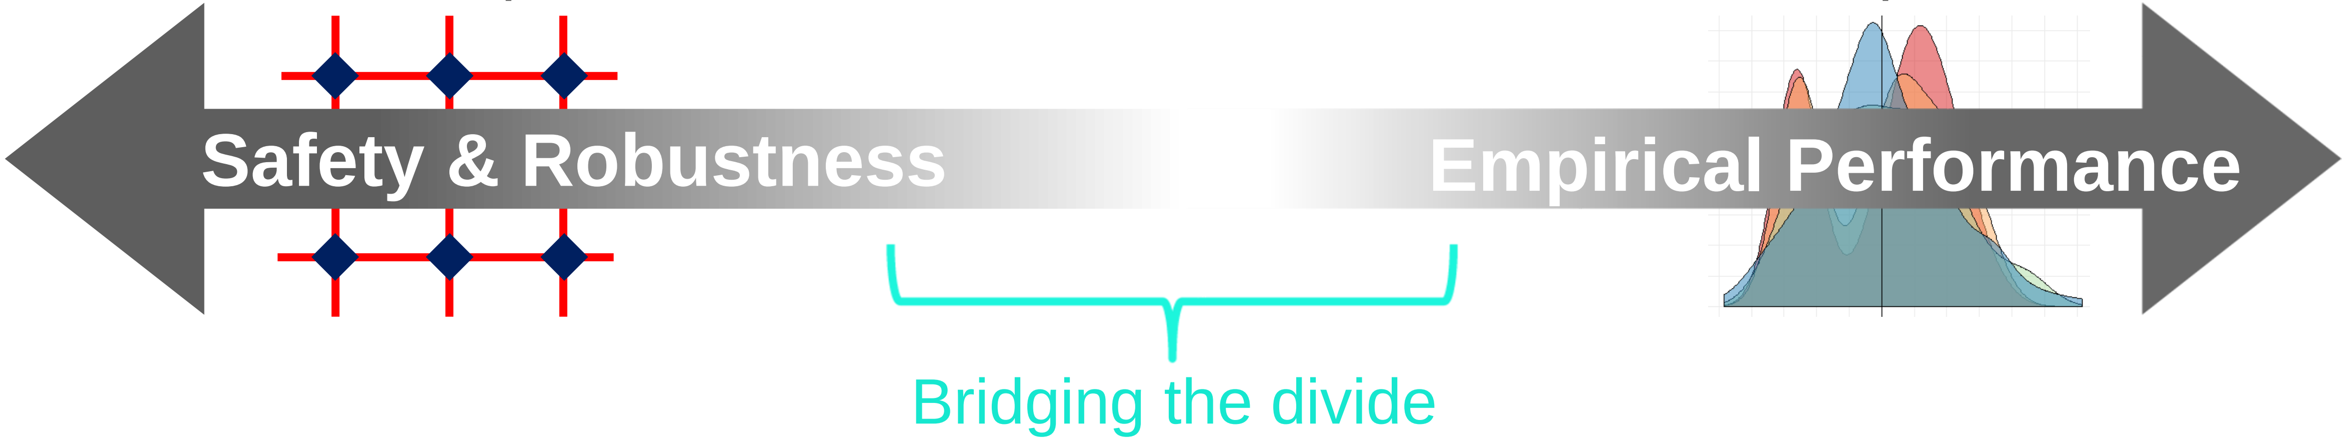

But, we also want to learn the fundamental limitations.

Understanding the limitations

- Ignoring **the fundamental limitations** can have catastrophic consequences, e.g. X-15 crash.
  - **Gunter Stein: Respect the Unstable. IEEE CDC Bode Lecture 1989**
- Control Engineering
  - Well established **conservation laws**, e.g. Bode Integral
  - **"Conservation laws make us humble" – Gunter Stein (T+S=1)**

- Modern ML tools
  - Need to understand the fundamental limitations $\Longrightarrow$ safe and intelligent aerial co-robots.

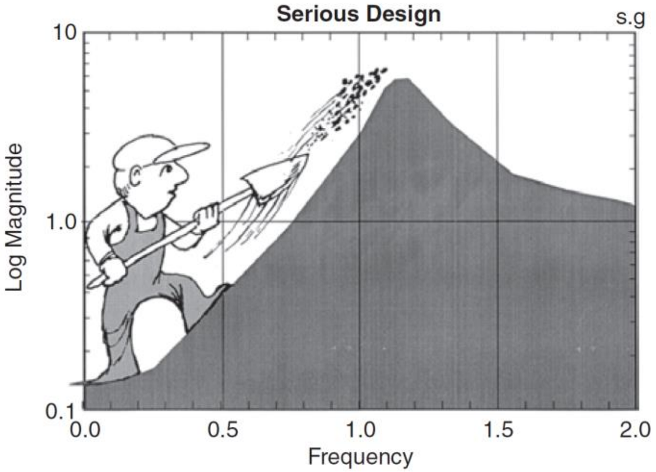

Control

ML

Future Directions

- Substantial efforts underway towards building a perceptual control framework through integration of $\mathcal{L}_1$ adaptive controllers with only-vision based feedback (in collaboration with Shenlong Wang)

- Ongoing efforts towards V&V of $\mathcal{L}_1$ adaptive control

- Distributionally Robust Adaptive Control framework
    - Robust planning and control for uncertain systems and environments
    - Theoretical guarantees for model-based stochastic optimal control
    - Robust deep learning architectures
    - Robust generative modeling

- Benchmarking and experimentation

Our Group

Alumni (Postdocs and Ph.D.):
Chengyu Cao (University of Connecticut)
Xiaofeng Wang (University of South Carolina)
Lili Ma (WIT, Boston, MA)
Vijay Patel (Indian Ministry of Defense)
Vahram Stepanyan (NASA Ames)
Jiang Wang (Apple)
Amanda Dippold (Howard Community College)
Dapeng Li (Acker.com)
Enric Xargay (Barcelona, Spain)
Zhiyuan Li (DJI)
Evgeny Kharisov (Waymo)
Hui Sun (Apple)
Venanzio Cichella (University of Iowa)
Syed Bilal Mehdi (Zoox)
Hanmin Lee (Korean Defense Agency)
Ronald Choe (Singapore Defense Agency)
Hamid Jafarnejadsani (Stevens Inst. of Tech.)
Steve Snyder (NASA)
Thiago Marinho (Waymo, Alphabet)
Hyung-Jin Yoon (Tennessee Tech University)
Alexandre Barbosa (Amazon Robotics)
Kasey Ackerman (NASA)
Javier Puig Navarro (NASA)
Arun Lakshmanan (Apple)
Gabriel Barsi Haberfeld (Apple)
Andrew Patterson (NASA)
Wenbin Wan (University of New Mexico)

Current Postdocs:
Aditya Gahlawat
Sheng Cheng
Pan Zhao
Tigran Bakaryan

Current Graduate Students:
Yikun Cheng, Hyungsoo Kang, Neng Wan, Zhuohuan Wu, Erin Swansen, Lin Song, Ziyao Guo
Vivek Sharma, Min Jun Sung
Yuliang Gu, Jing Wu, Chuyuan Tao, Michael Aramyan, Sambhu Karamanas

Sponsors

Source: New York Times

Learjet 1

F-16

Learjet 2

Research supported by NSF, AFOSR, NASA, ARO, USSOCOM, ONR, Boeing, IntelinAir and transitioned into commercial products by Raymarine, Caterpillar, Raytheon, JOUAV, StatOil, IntelinAir, among many others.