Safe Learning and Control with $\mathcal{L}_1$-Adaptation

Naira Hovakimyan

Mechanical Science and Engineering
University of Illinois at Urbana-Champaign

May 9, 2021

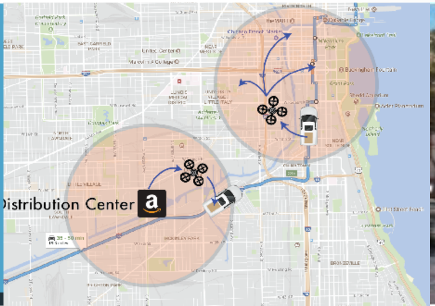

Aerospace History at a Glance

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

Building a Bridge

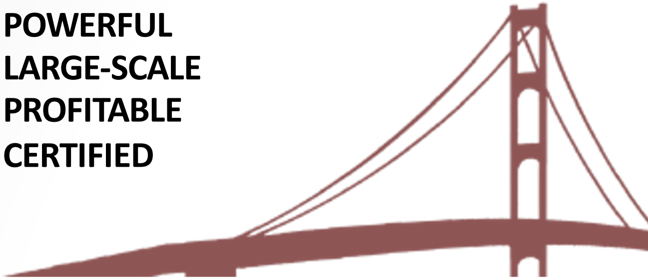

- AUTONOMOUS
- COLLABORATIVE
- INTELLIGENT
- LOW-COST
- RECONFIGURABLE
- MULTI-AGENT
- HUMAN-CENTERED

Talent
Technology
Expertise

Data-Driven

Safety-Critical Systems

Learjet 45

Aerial robots

Self-driving cars (credit: Daniel Lu)

- Difficult to ***safely*** obtain a large amount of data required by e.g. Reinforceent Learning methods to produce good control policies.
- Accidents can be expensive and sometimes even deadly.

Learning-Based Control Setup

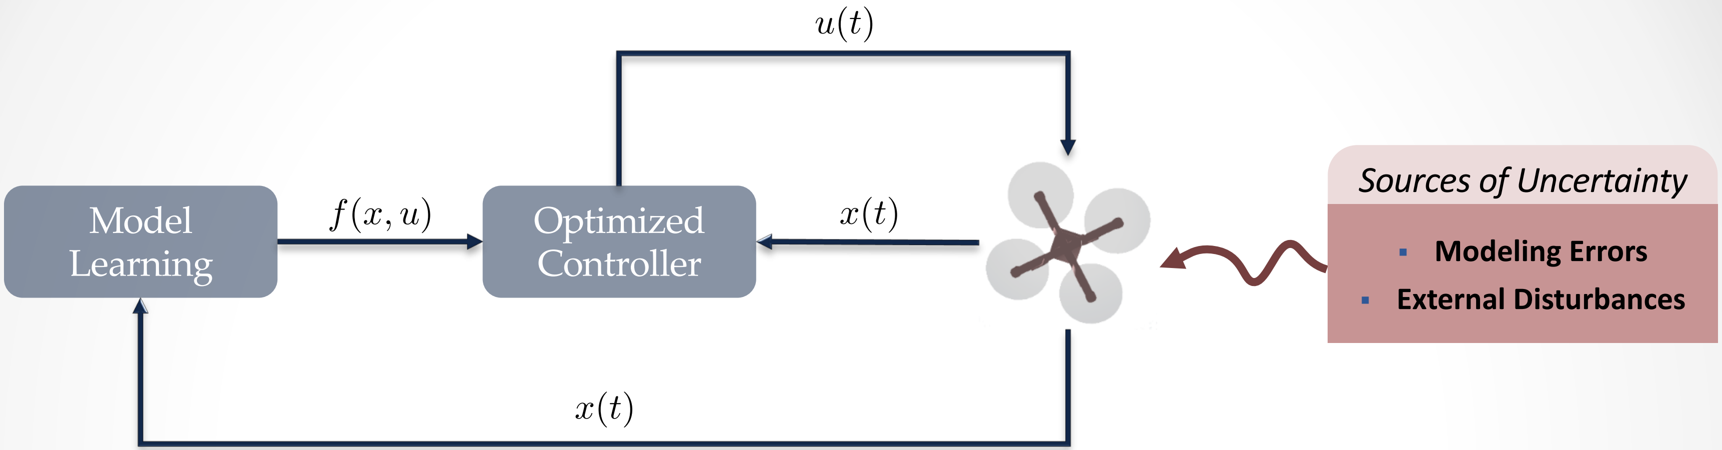

Loss of performance and stability guarantees for robots under uncertainty!

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

Safe Learning and Control Setup

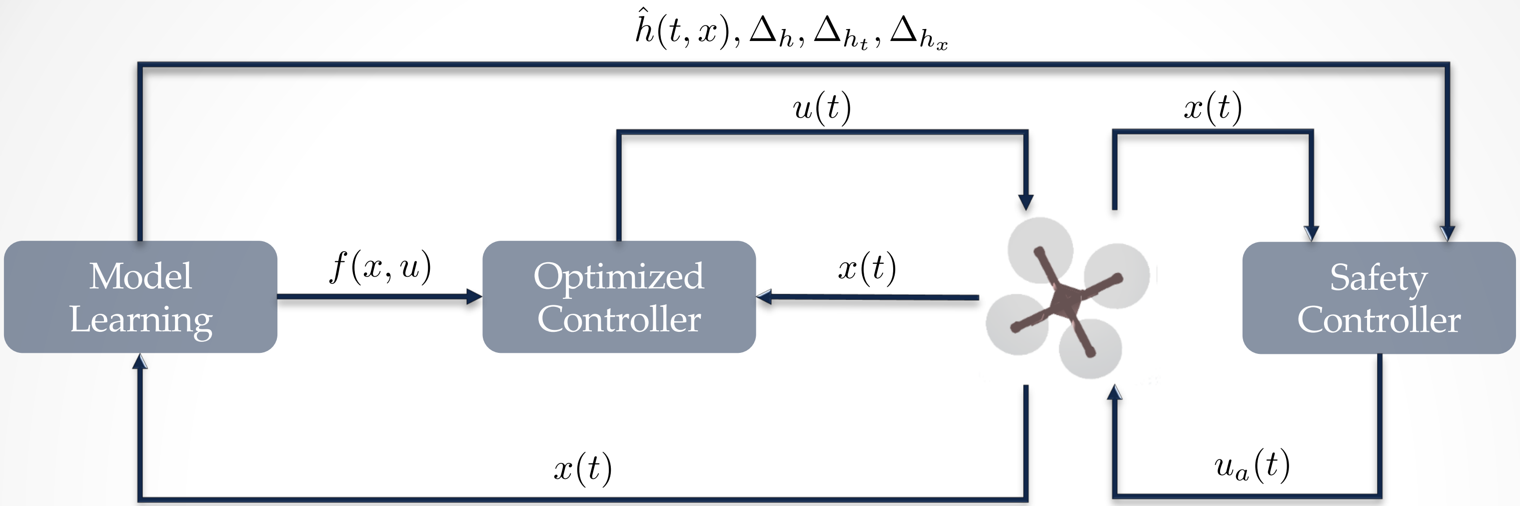

$$\dot{x}(t) = \boxed{f(x,u)} + \boxed{h(t,x)} \qquad \boxed{\Delta_h, \Delta_{h_t}, \Delta_{h_x}}$$

Nominal Model

Uncertainty

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:***

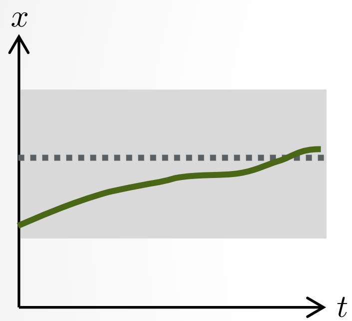

Transient performance

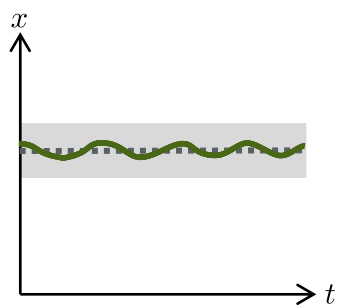

Steady-state performance

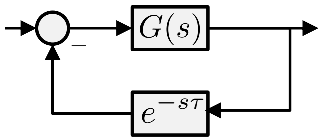

Time-delay margin

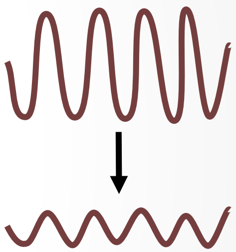

Disturbance rejection

nhovakim@Illinois.edu

$\mathcal{L}_1$-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.

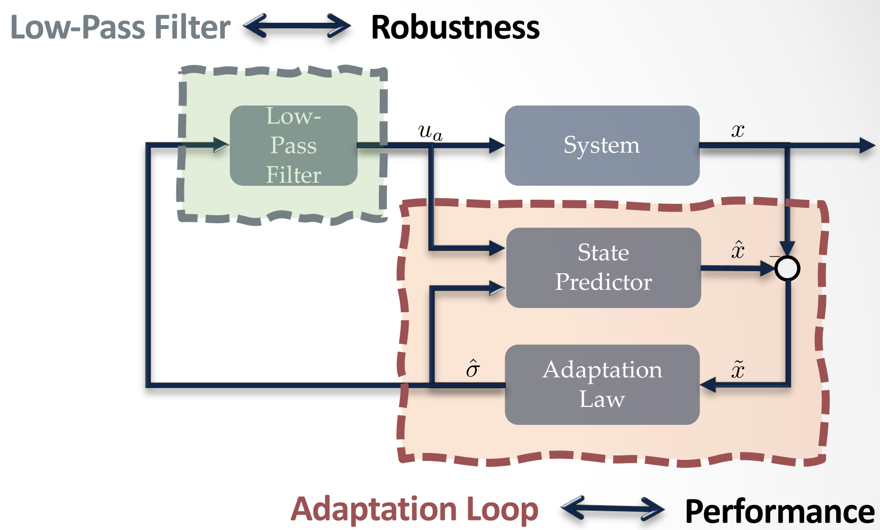

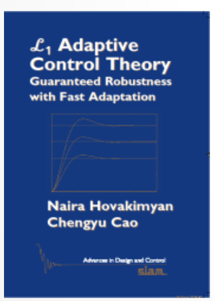

intelinair

Statoil

nhovakim@Illinois.edu

Safe Learning and Control Setup

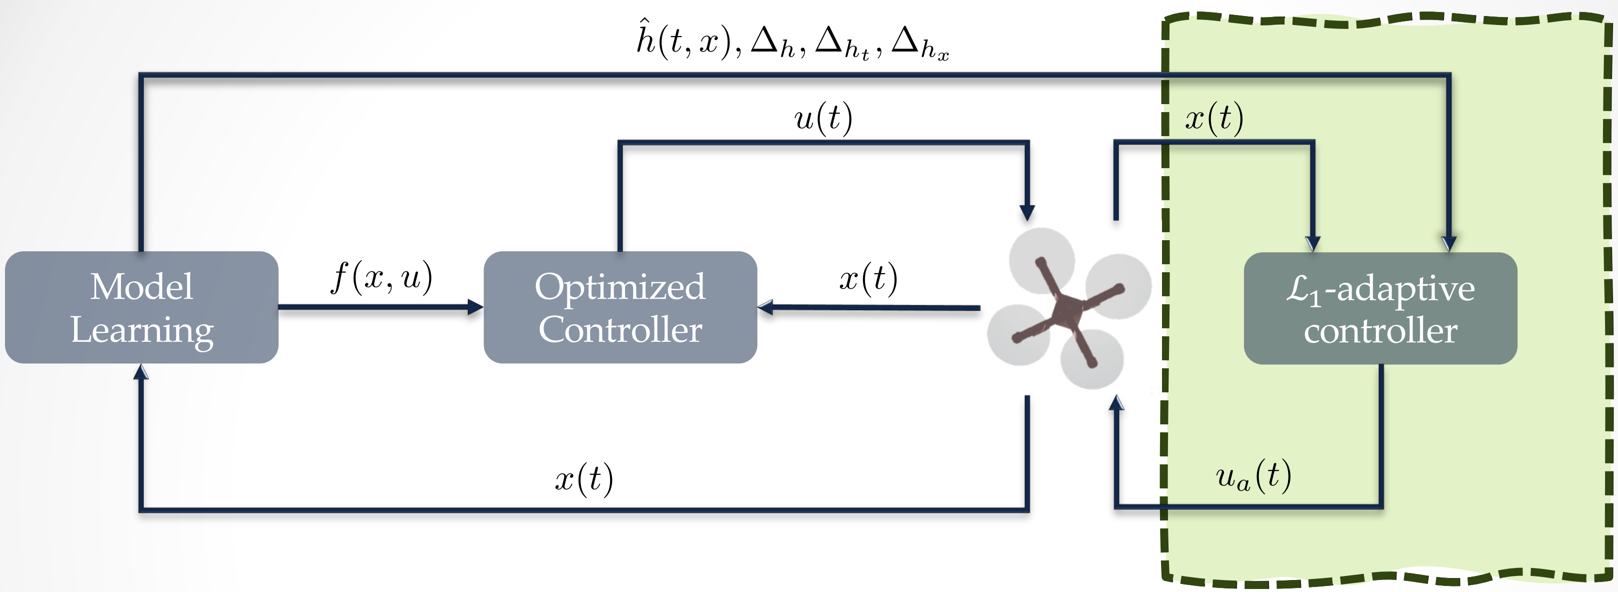

- **Retaining the key features of performance and robustness guarantees of $\mathcal{L}_1$- adaptive controllers**
- **Benefitting from the versatility offered by machine learning methods**

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

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

Joint papers on highly unstable aircraft configurations

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

Tests on Learjet and F16 at the Edwards AFB

Safe Learning and Control

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

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.

nhovakim@Illinois.edu

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.

nhovakim@Illinois.edu

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

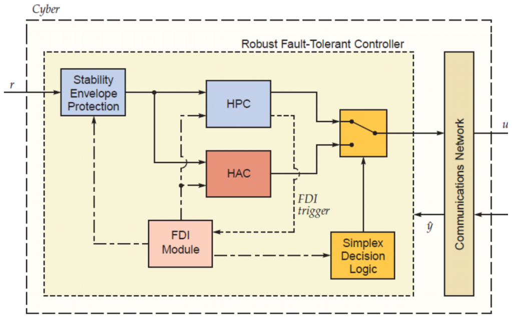

Fault-tolerant software

Drone—based package delivery

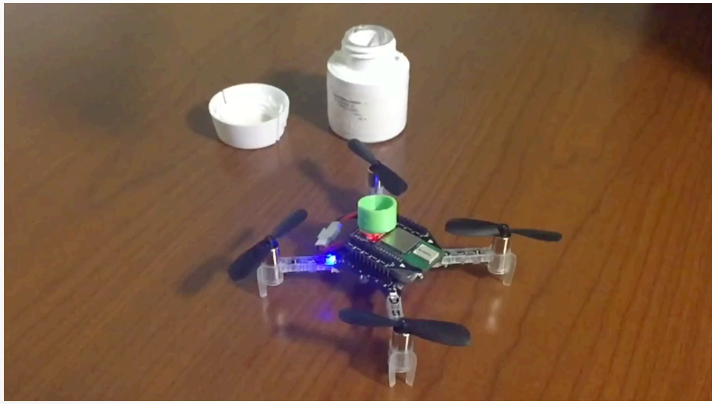

Indoor aerial vehicles

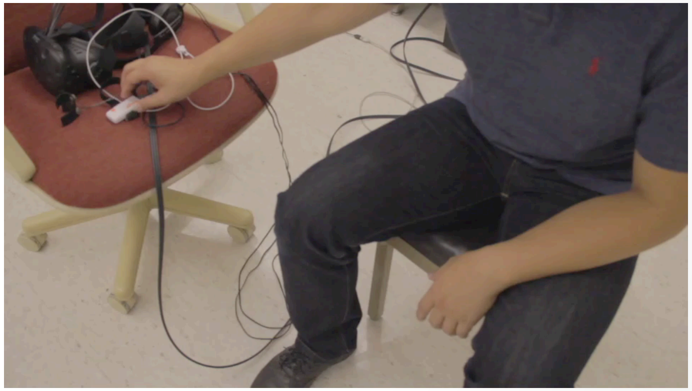

Human-centered (perceived) safety

nhovakim@Illinois.edu

Complex Environments

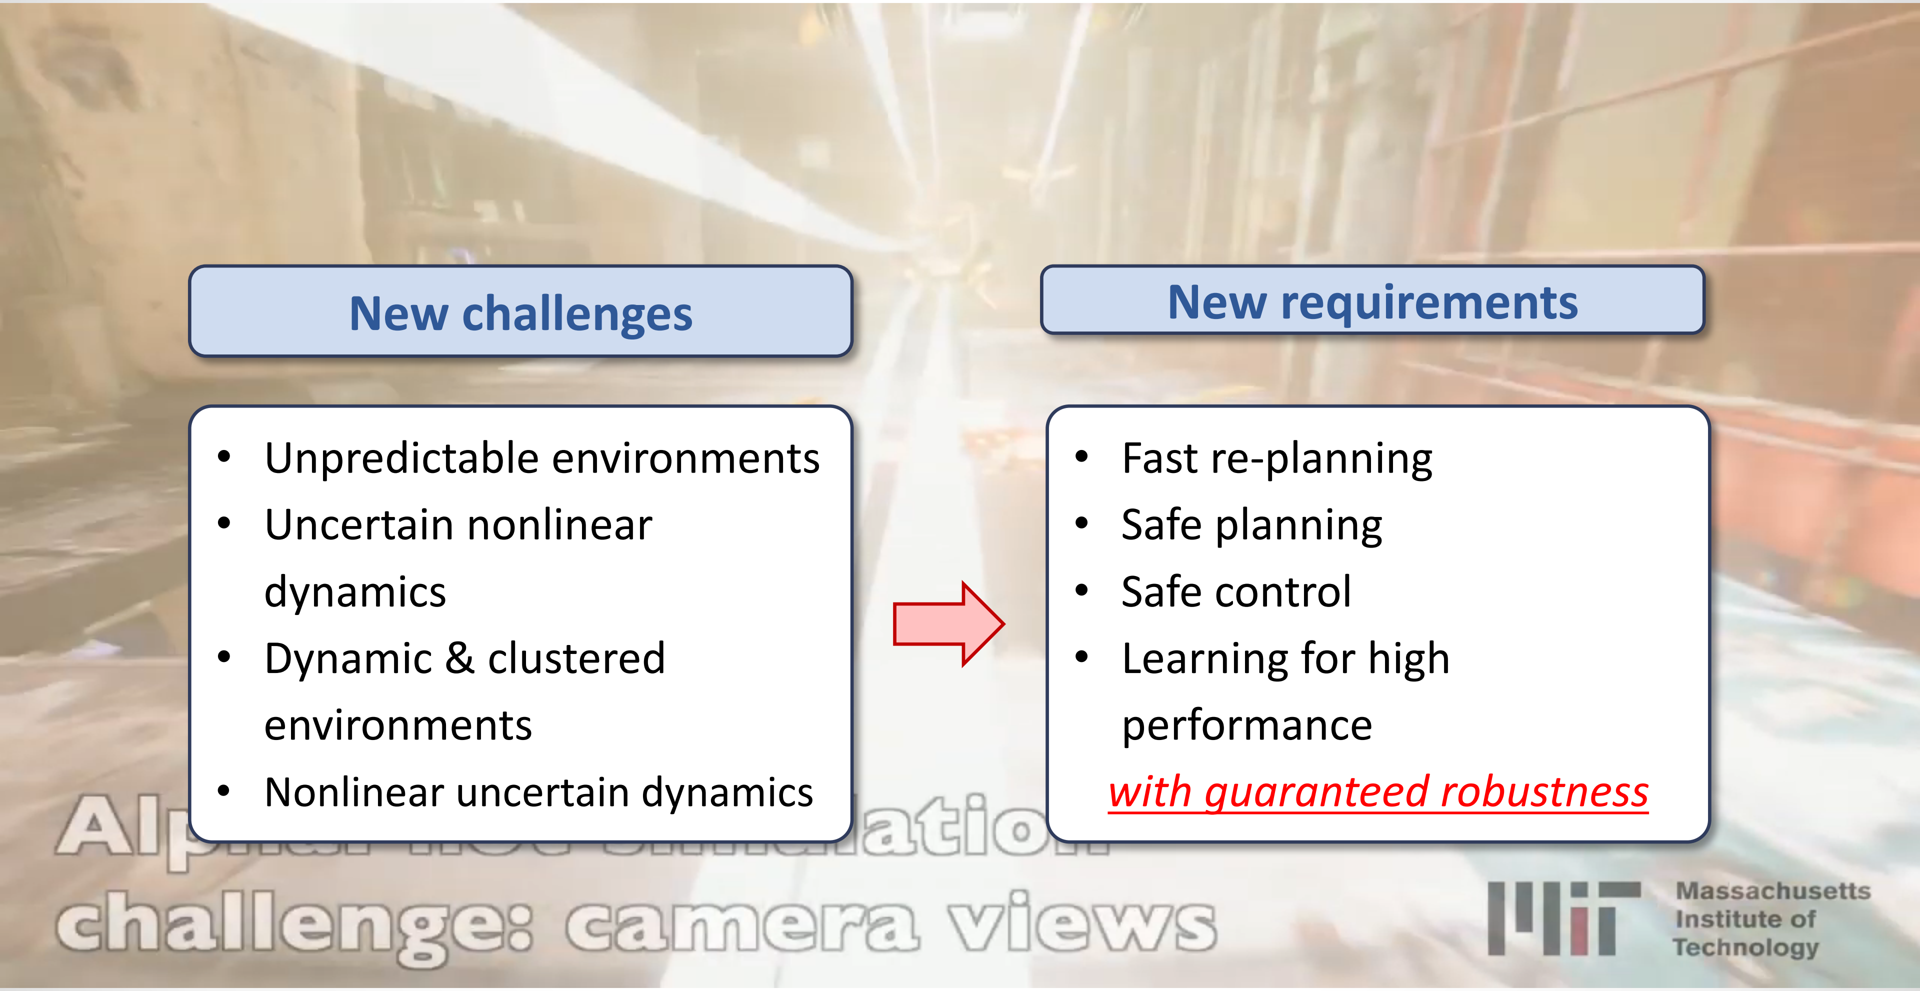

New challenges

- Unpredictable environments
- Uncertain nonlinear dynamics
- Dynamic & clustered environments
- Nonlinear uncertain dynamics

New requirements

- Fast re-planning
- Safe planning
- Safe control
- Learning for high performance

with guaranteed robustness

Guerra, Tal, Murali, Ryou, and Karaman. "FlightGoggles: Photorealistic Sensor Simulation for Perception-driven Robotics using Photogrammetry and Virtual Reality." International Conference on Intelligent Robots and Systems (IROS), 2019.

Planning and Control Pipeline

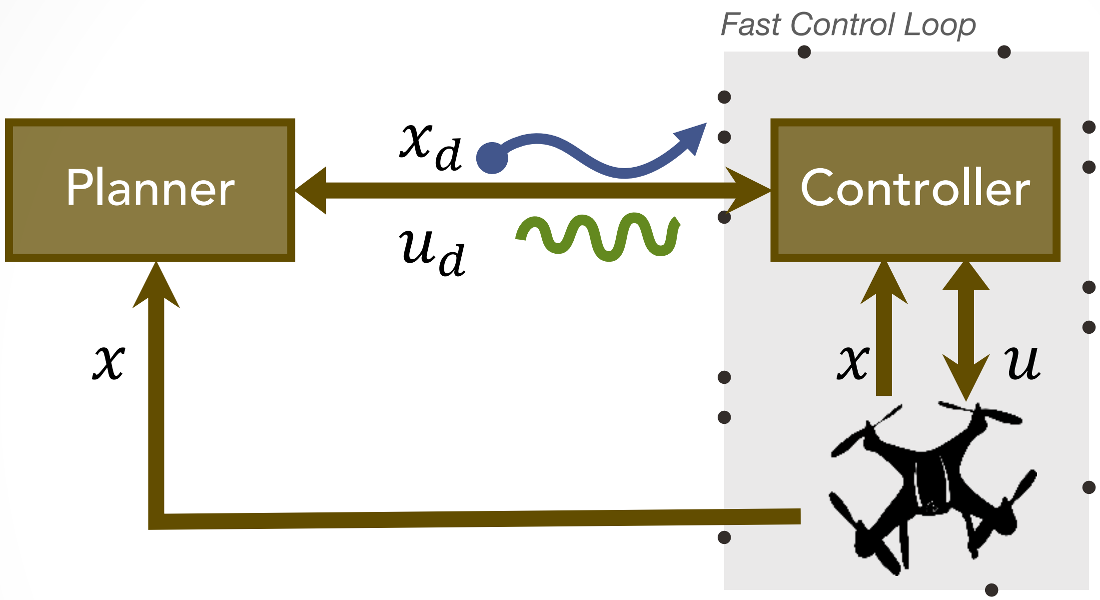

Standard architecture

- Planner generates desired x_d and open loop u_d
- Controller attempts to track using feedback law $u = k(u_d, x_d, x)$

nhovakim@Illinois.edu

Learning-Based Control Pipeline

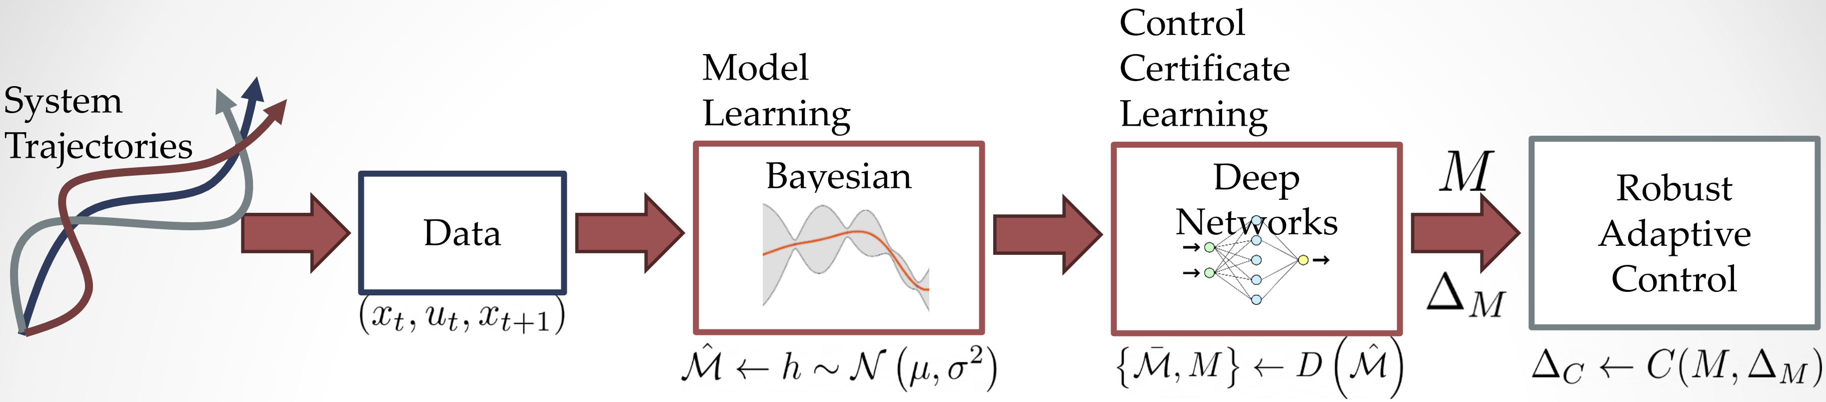

A bottom-up approach

- Foundation based on Control theory: Safety as a requirement
- Build-up based on ML: Performance as a luxury

Unified in an architecture that decouples safety from learning

- Learning does not affect safety, only performance

Learning-Based Control Pipeline

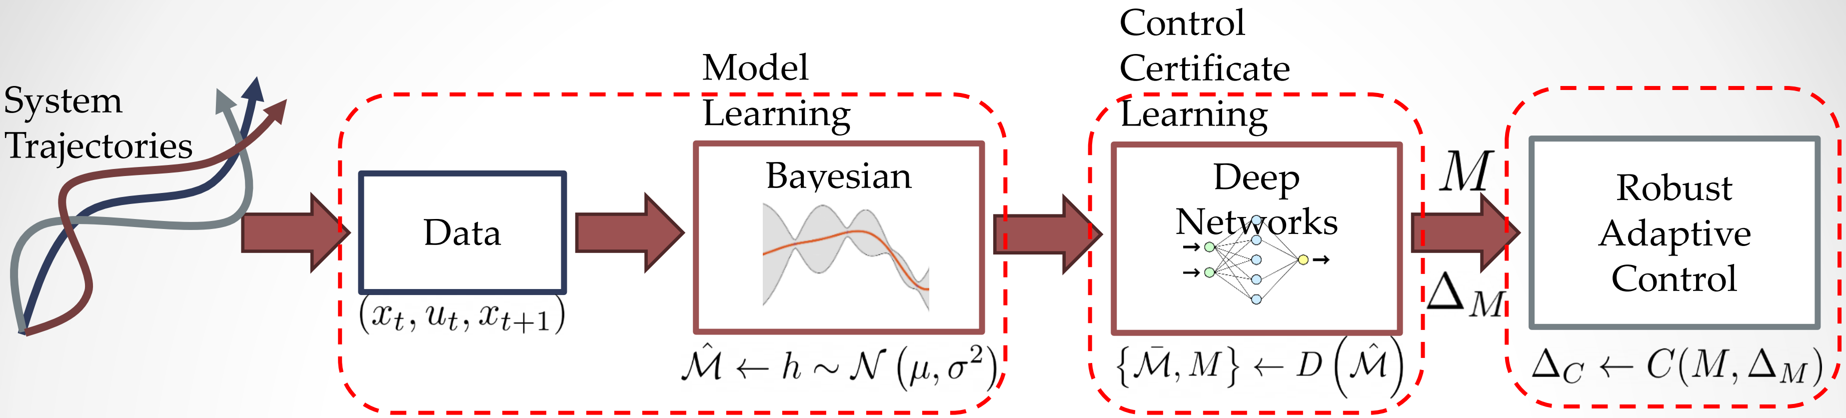

The building blocks

- $\mathcal{L}_1$ - MPPI and Riemannian-energy $\mathcal{L}_1$ ($\mathcal{RL}_1$) Control: Safety
- Bayesian learning of dynamics - $\mathcal{RL}_1 - GP$ control: Performance
- Deep incrementally stabilizing control: Usability

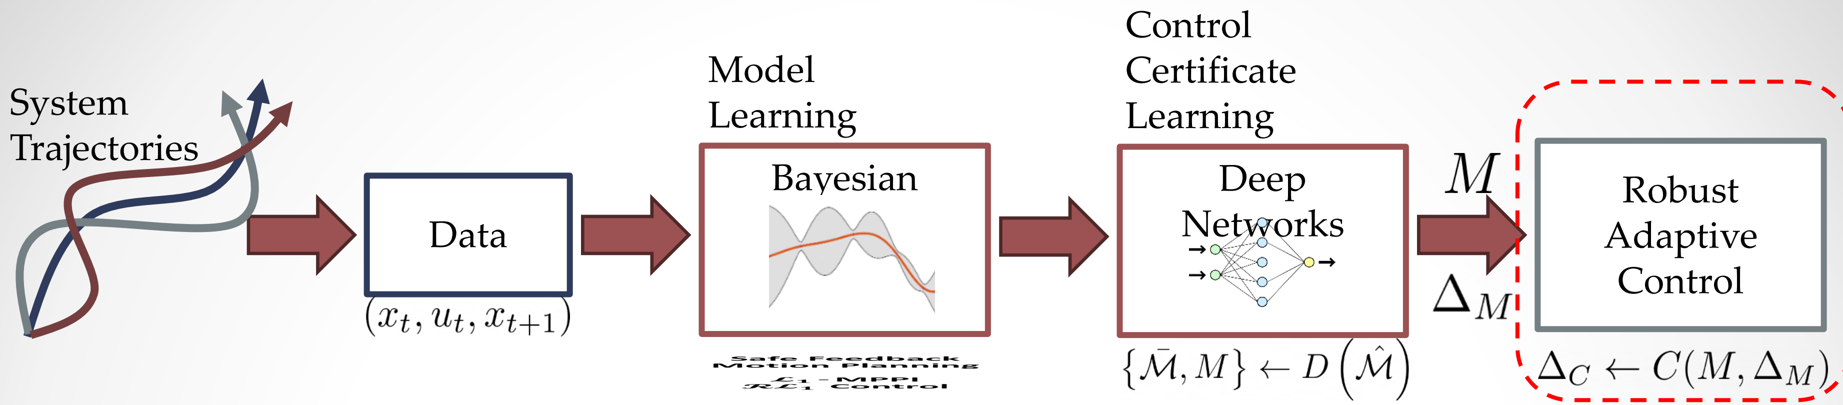

Safe Feedback Motion Planning

$\mathcal{L}_1$ - MPPI

Contraction $\mathcal{L}_1$ Control

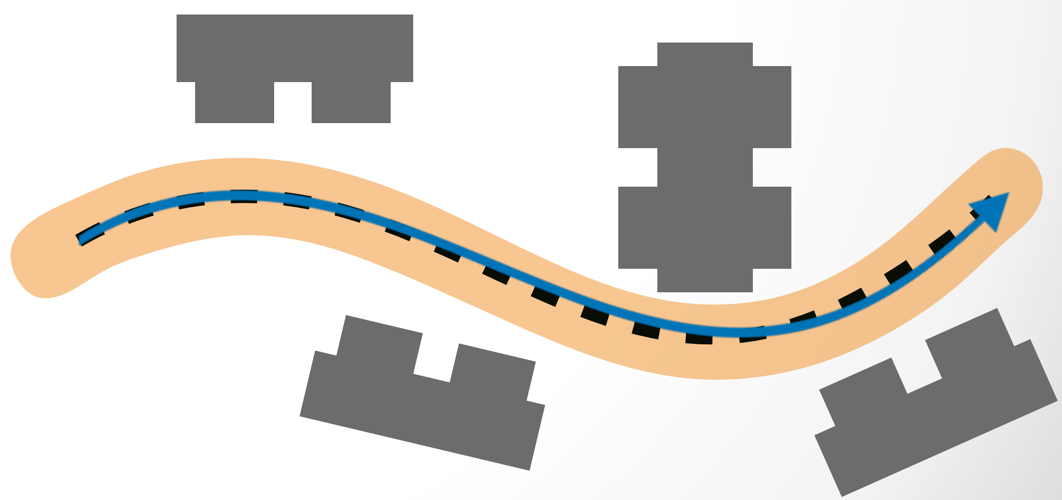

nhovakim@Illinois.edu

$\mathcal{L}_1$-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 $\boldsymbol{\mathcal{L}_1}$ **augmentation.**

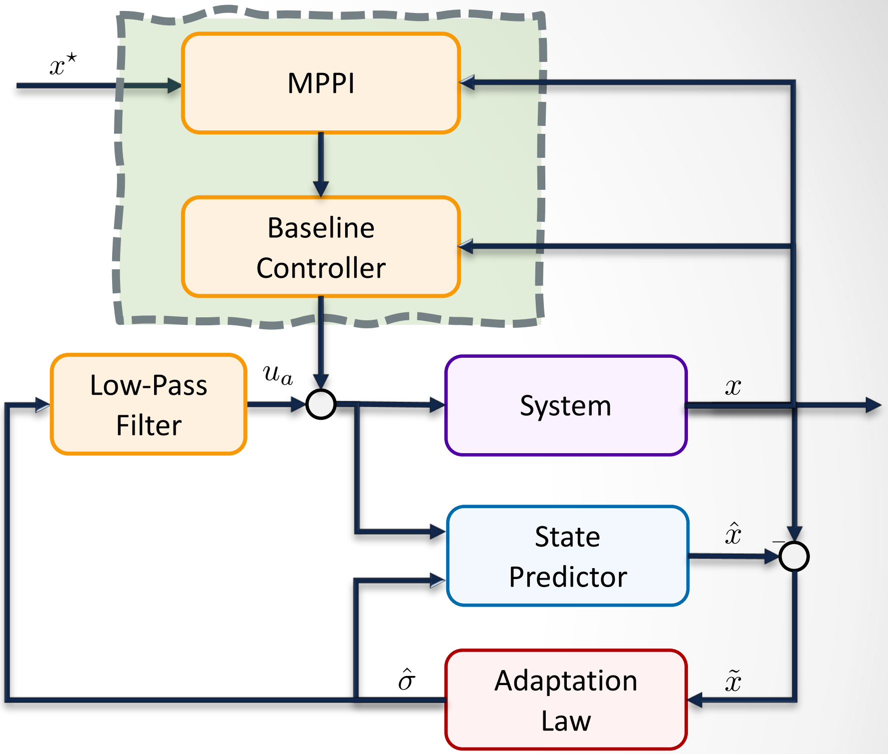

Pravitra, J., Ackerman, K. A., Cao, C., Hovakimyan, N., and Theodorou, E. A. L1-Adaptive MPPI Architecture for Robust and Agile Control of Multirotors. International Conference on Intelligent Robots and Systems, 2020.

Goggle Environment

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

✗ : crash ✓ : complete

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.

Quantifiable safety certificates for nonlinear systems that enable ***safe motion planning***.

Challenges in Safe Planning

- Planning under uncertainty is of vital importance for **safety critical applications**
- A planner-agnostic approach to certify **safe tubes** around desired trajectories, that the UAV is **guaranteed to remain inside**
- Tuning knobs that allow **tubes** to be made arbitrarily **small** as a **trade-off between performance and robustness.**

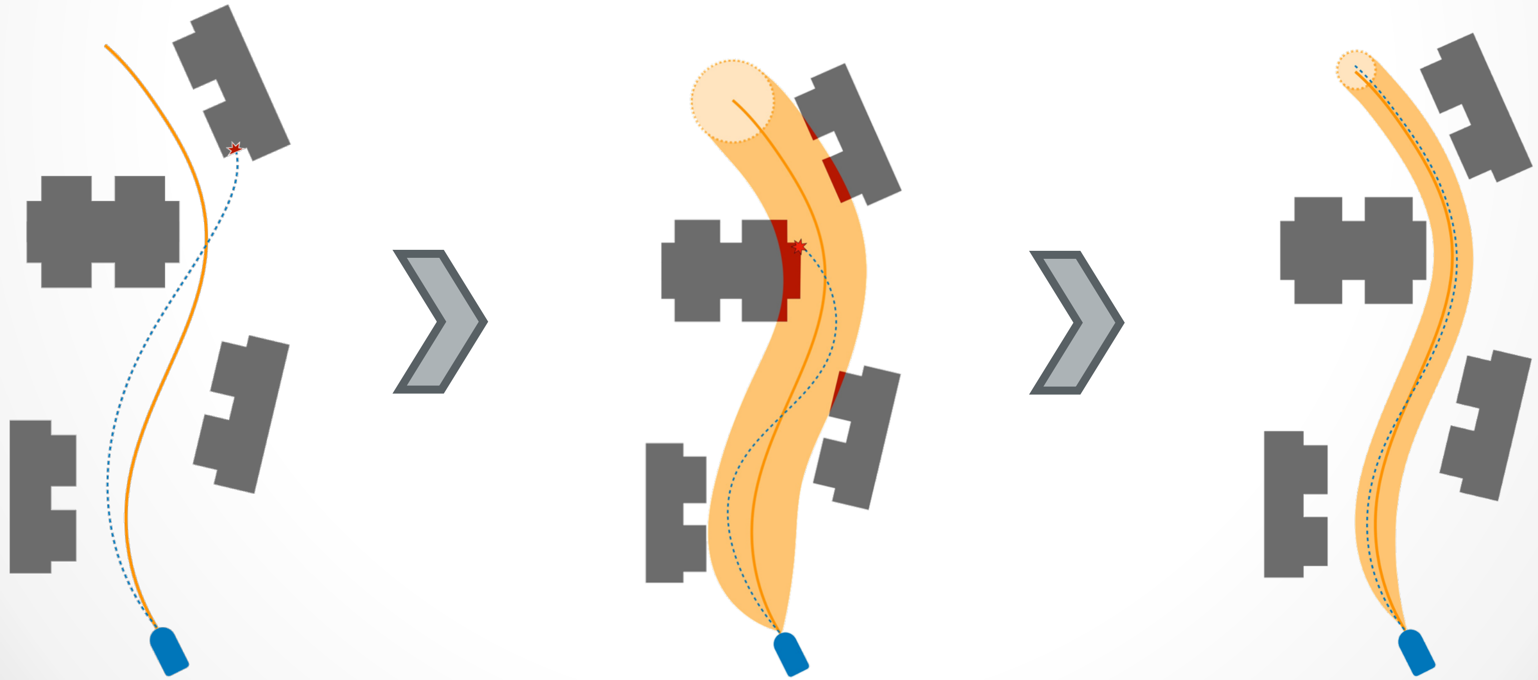

Contraction $\mathcal{L}_1$- Architecture: Safe Planning

- Contraction theory enables consideration of **nonlinear reference systems**
  - applies to **a large class of nonlinear systems.**

- **$\mathcal{L}_1$ adaptive control** compensates for uncertainties with **guaranteed robustness.**

- Combined architecture allows for publishing **<u>certificates of performance and robustness</u>**.

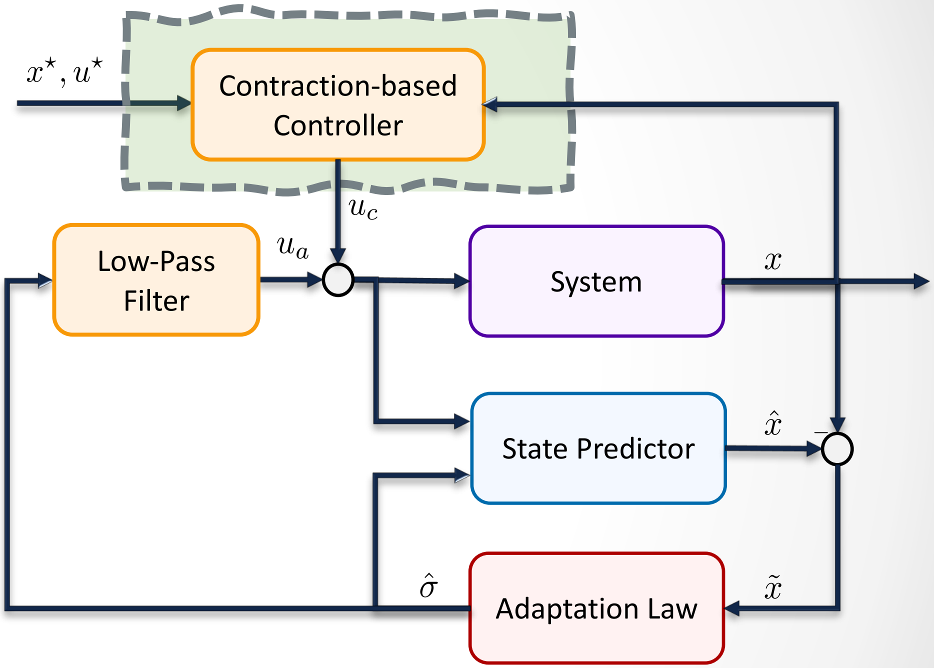

Lakshmanan, A., Gahlawat, A., & Hovakimyan, N. (2020). "Safe feedback motion planning: A contraction theory and L1-adaptive control based approach." Conference on Decision and Control (CDC) 2020.

Quantifiable safety certificates for nonlinear systems for ***safe learning and control***.

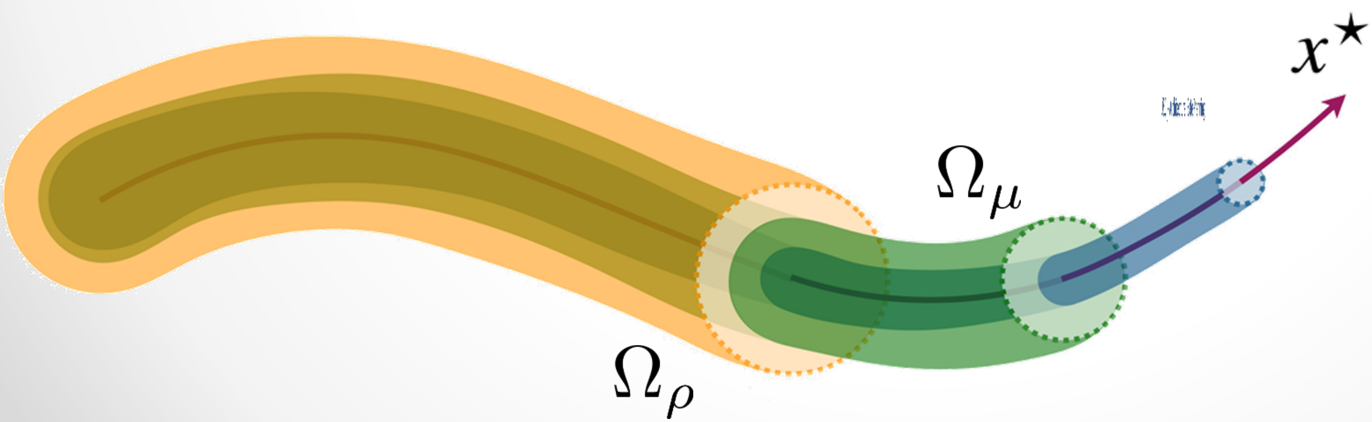

Ω_c : decaying error

$\Omega_\mu \propto$ filter bandwidth

$\Omega_\rho \propto$ adaptation rate

nhovakim@Illinois.edu

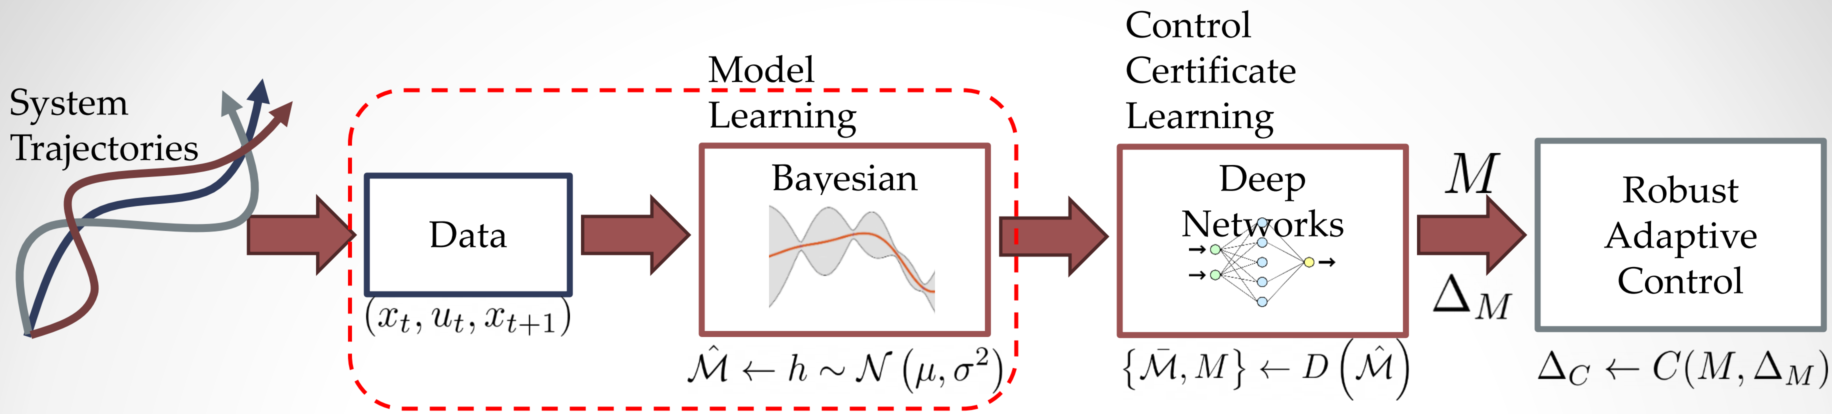

Safe Simultaneous Learning and Control

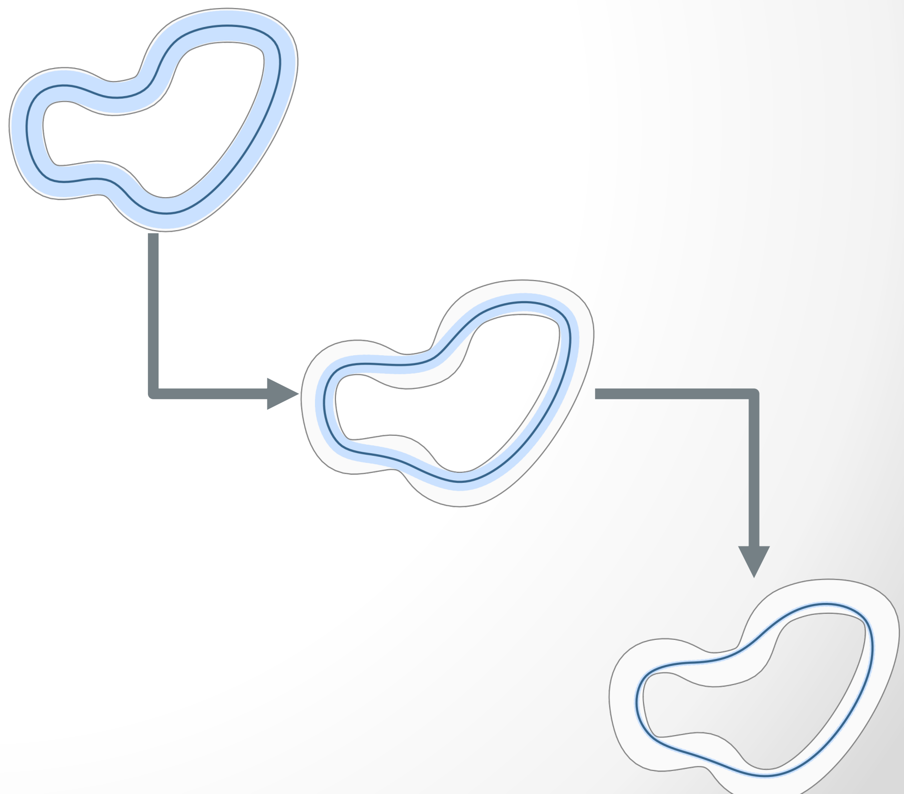

Contraction $\mathcal{L}_1$- $\mathcal{GP}$

nhovakim@Illinois.edu

$\mathcal{L}_1$-$\mathcal{GP}$: Safe Bayesian Learning

- Gaussian process ($\mathcal{GP}$) learns model uncertainties
  - **No need for persistent excitation**
  - No need for parametric models
  - The kernel function **integrates prior model knowledge**
  - **High probability error bounds** on the prediction*
- $\mathcal{L}_1$ augmentation
  - Ensures safety during learning
  - Vanishes once learning is complete

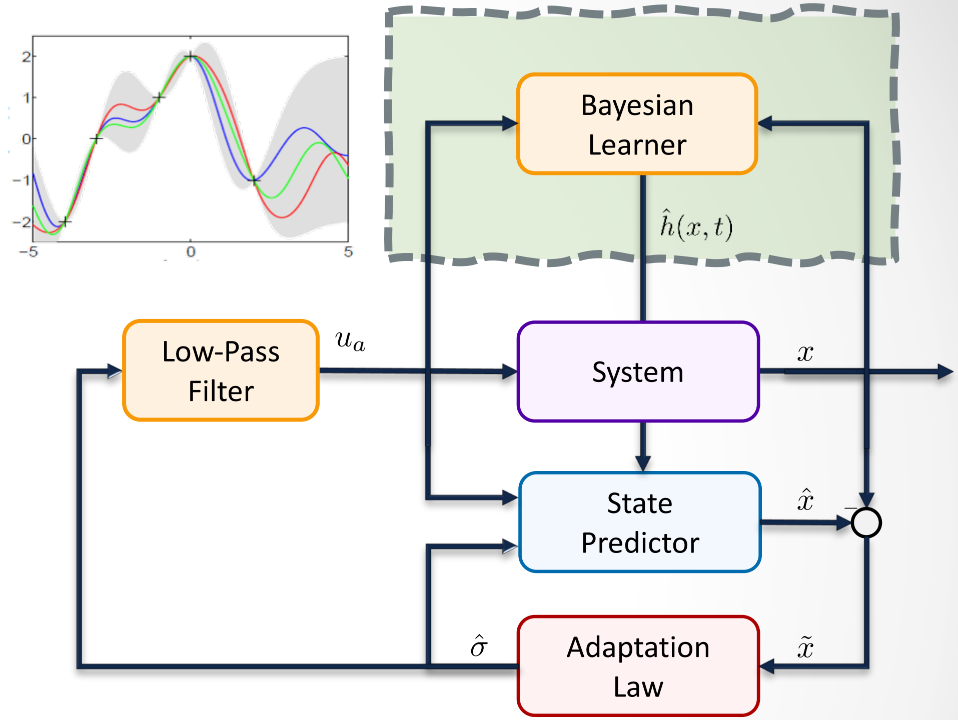

*A. Lederer, J. Umlauft, and S. Hirche. "Uniform error bounds for Gaussian process regression with application to safe control." In the Advances in Neural Information Processing Systems." 2019

A. Gahlawat, P. Zhao, A. Patterson, N. Hovakimyan, E. Theodorou, "$\mathcal{L}_1$-$\mathcal{GP}$: $\mathcal{L}_1$ Adaptive Control with Bayesian Learning." In the Learning for Dynamics and Control (L4DC) Conference, Berkeley, CA, 2020

Safe Simultaneous Learning, Planning and Control

Gaussian process

- Learning with minimal data

$\mathcal{L}_1$ controller

- Safety during learning, robustness, performance

Contraction theory

- Safety tubes

Contraction $\mathcal{L}_1$- $\mathcal{GP}$ Architecture: Safe Learning and Control

- **Safety certificates** in the form of tubes from the $\mathcal{RL}_1$- $\mathcal{GP}$ 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

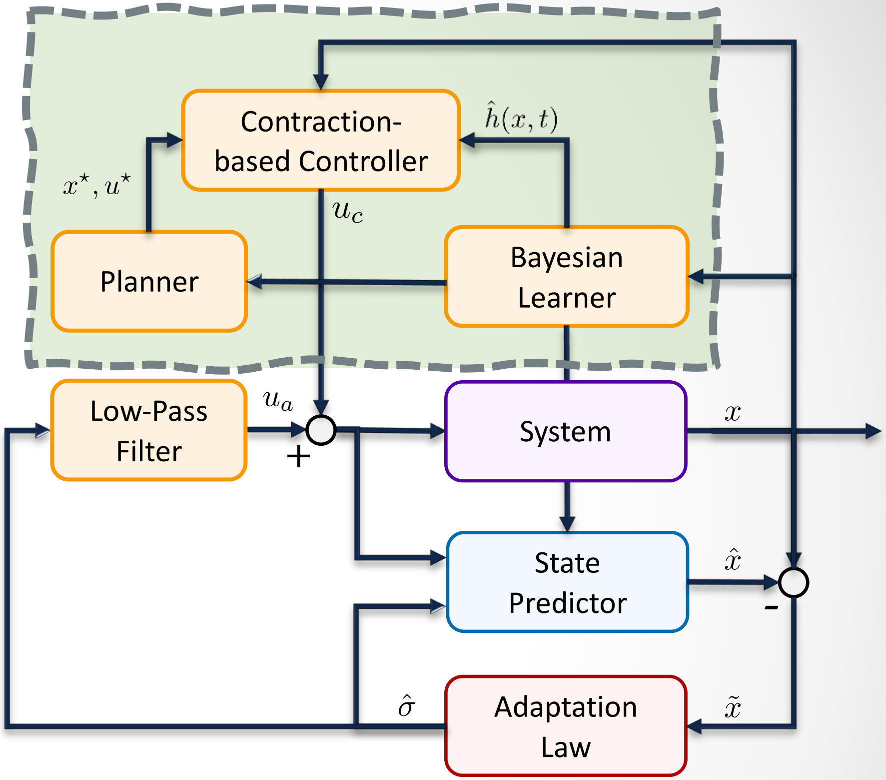

Gahlawat, A., Lakshmanan, A., Song, L., Patterson, A., Wu, Z., Hovakimyan, N., & Theodorou, E. "Contraction $\mathcal{L}_1$ Adaptive Control with Gaussian Processes ." Learning for Dynamics and Control (L4DC) Conference (To Appear). 2021.

Contraction $\mathcal{L}_1$- $\mathcal{GP}$ Architecture: Safe Learning and Control

Learning **kicks in** ⟹ tube **shrinks** ⟹ **better** planning

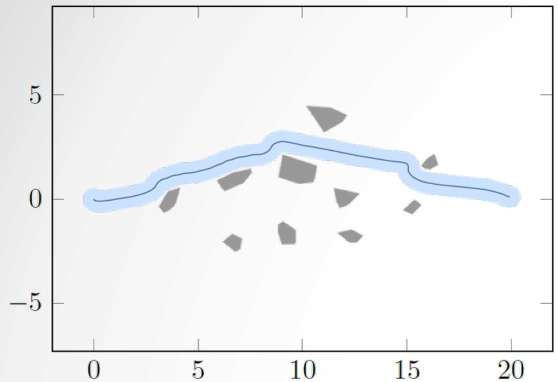

(a) Episode 1: $\omega = 90$ rad/s, $\Gamma = 7e10$, $N = 0$.
Traverse time: 27secs.

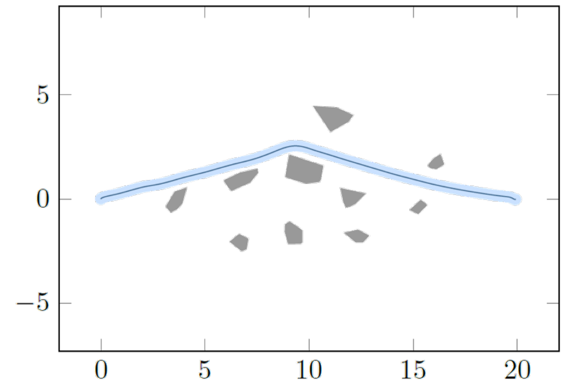

(b) Episode 2: $\omega = 30$ rad/s, $\Gamma = 2e6$, $N = 25$.
Traverse time: 16secs.

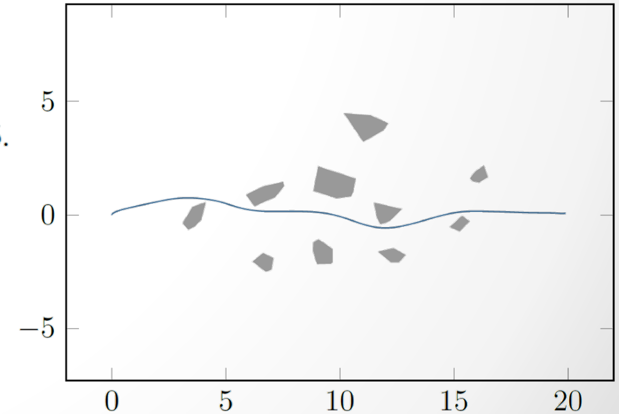

(c) Episode 3: $\omega = 30$ rad/s, $\Gamma = 2e6$, $N = 100$.
Traverse time: 14secs.

Contraction $\mathcal{L}_1$- $\mathcal{GP}$ Architecture: Safe Learning and Control

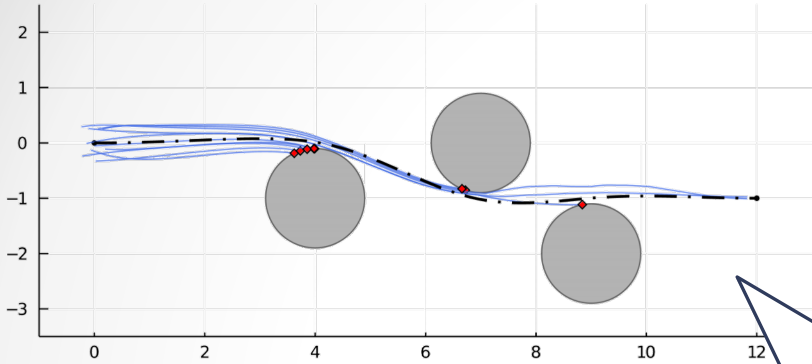

(a) Contraction-based feedback

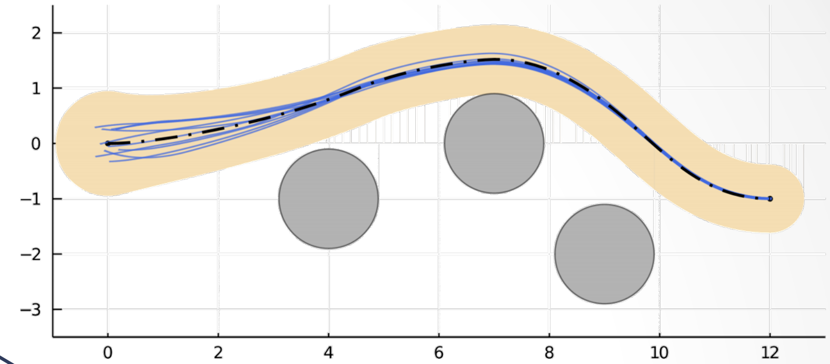

) $\mathcal{RL}_1$

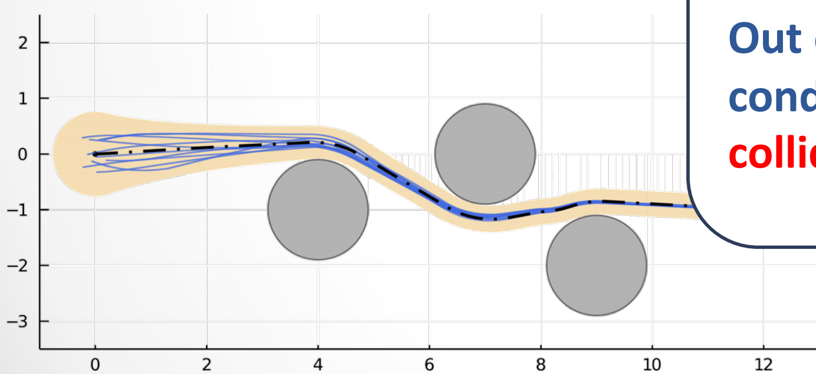

(c) $\mathcal{RL}_1$-$\mathcal{GP}$

No safety guarantees!

Out of ten random initia conditions, 8 trajectorie collide with obstacles.

Uses a conservative knowledge of the uncertainty.

teed!

As the uncertainty is learned, performance is improved without sacrificing robustness.

Contraction $\mathcal{L}_1$- $\mathcal{GP}$ Architecture: Safe Learning and Control

Learning **kicks in** ⟹ tube **shrinks** ⟹ **better** planning

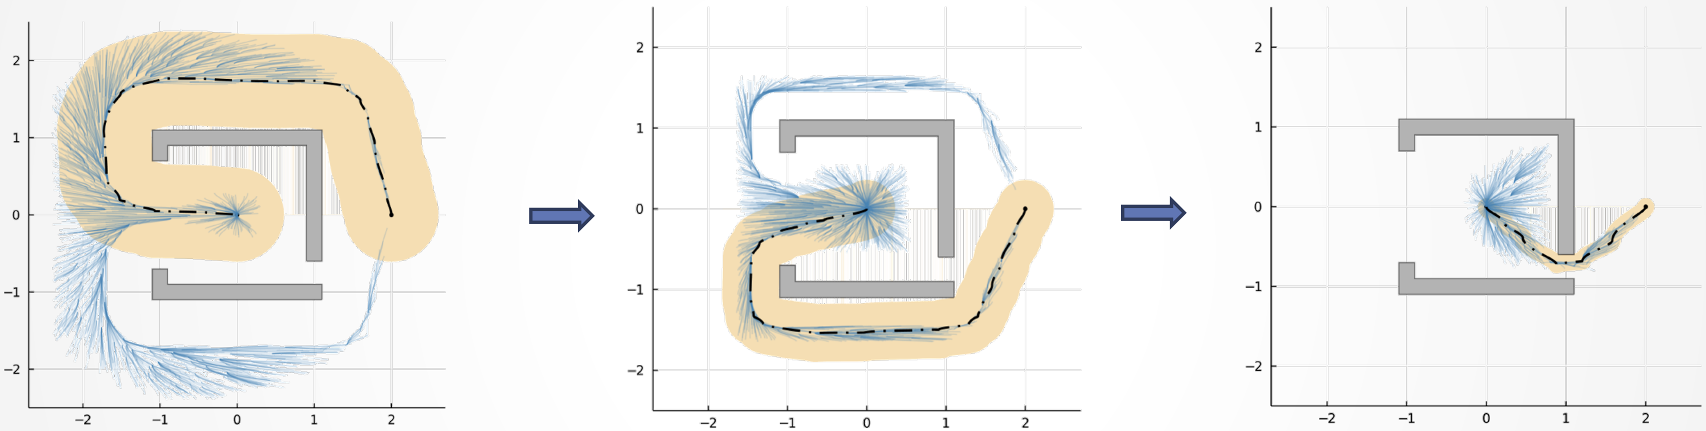

(a) only a deterministic knowledge of the uncertainty

(b) model learned with N = 25 dataset

(c) model learned with N = 100 dataset

When learned model updates, no retuning required for $\mathcal{RL}_1$ ⟹ **Safety is decoupled from learning!**

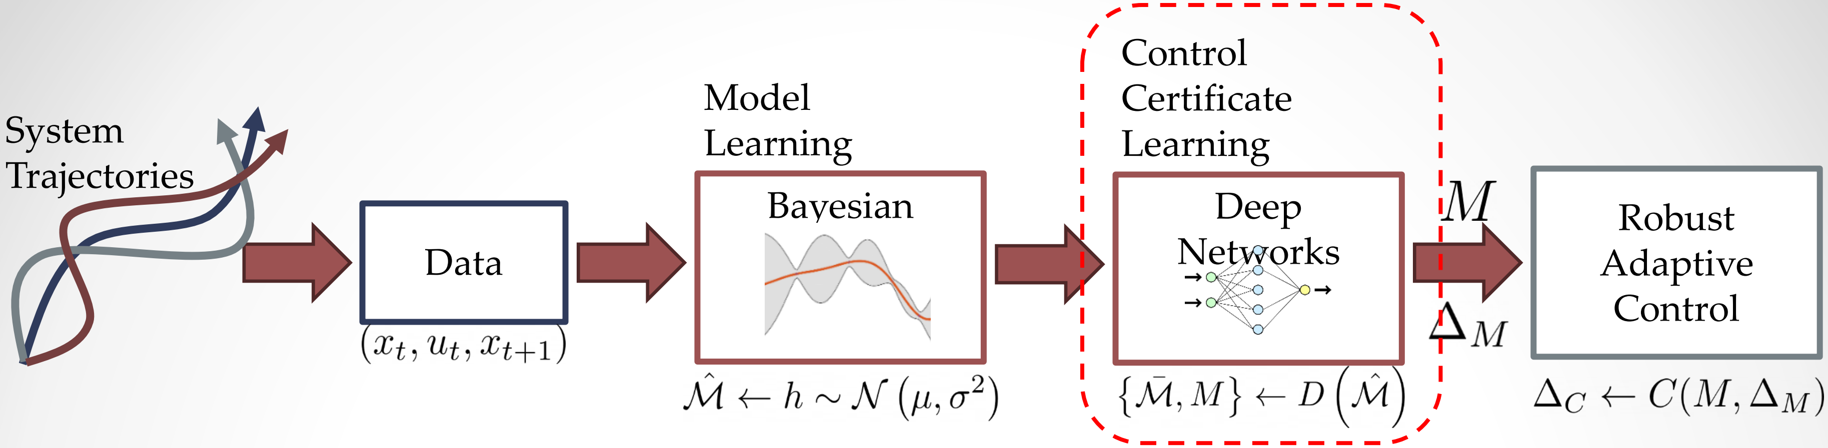

Learning the Control Certificates

DISCO

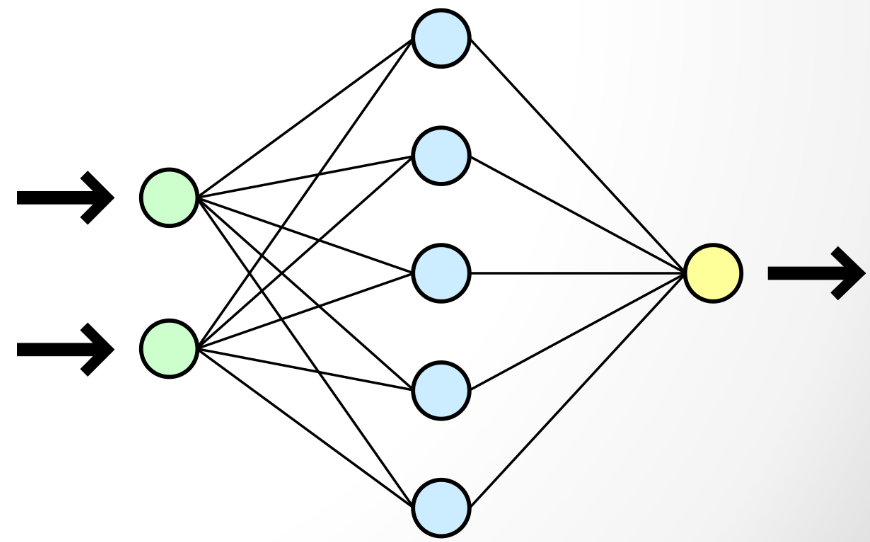

Learning accurate models by themselves is **not** enough!

The accuracy of the learned models is **not sufficient** to guarantee that controllers can be synthesized.

Deep Incrementally Stabilizing Control

Our approach: Optimization-based relaxation of the feasibility problem

Design spaces of valid CCMs/ILFs and feedback laws

- Extract models that are analytically verified to be IES
- A projection-based approach

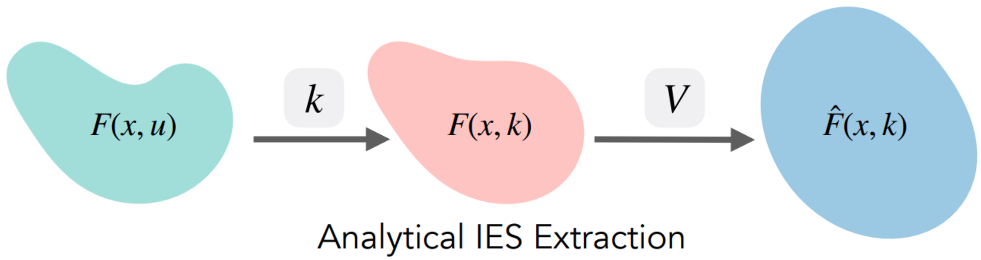

Parameterize the space of valid candidates using DNNs

- DNN opt. minimizes model-mismatch between learned and extracted models

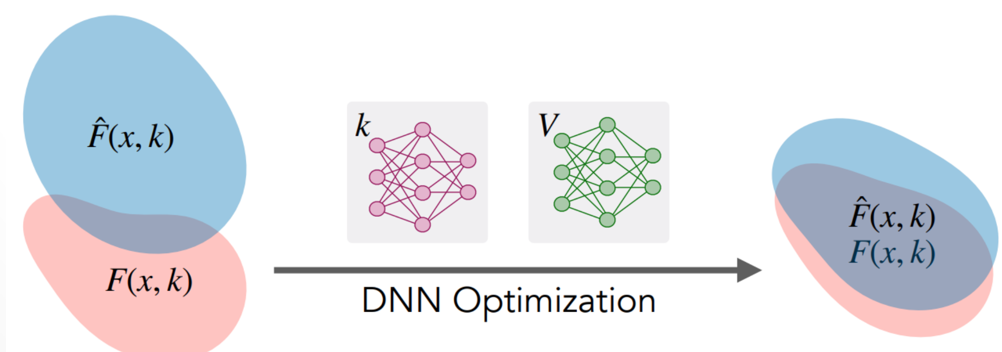

*Gahlawat, A., Lakshmanan, A., Vlahov, B. Gandhi, M. Song, L, Hovakimyan, N., and Theodorou, E. "Deep Incrementally Stabilizing Control. *Under Review.*

Control Pipeline

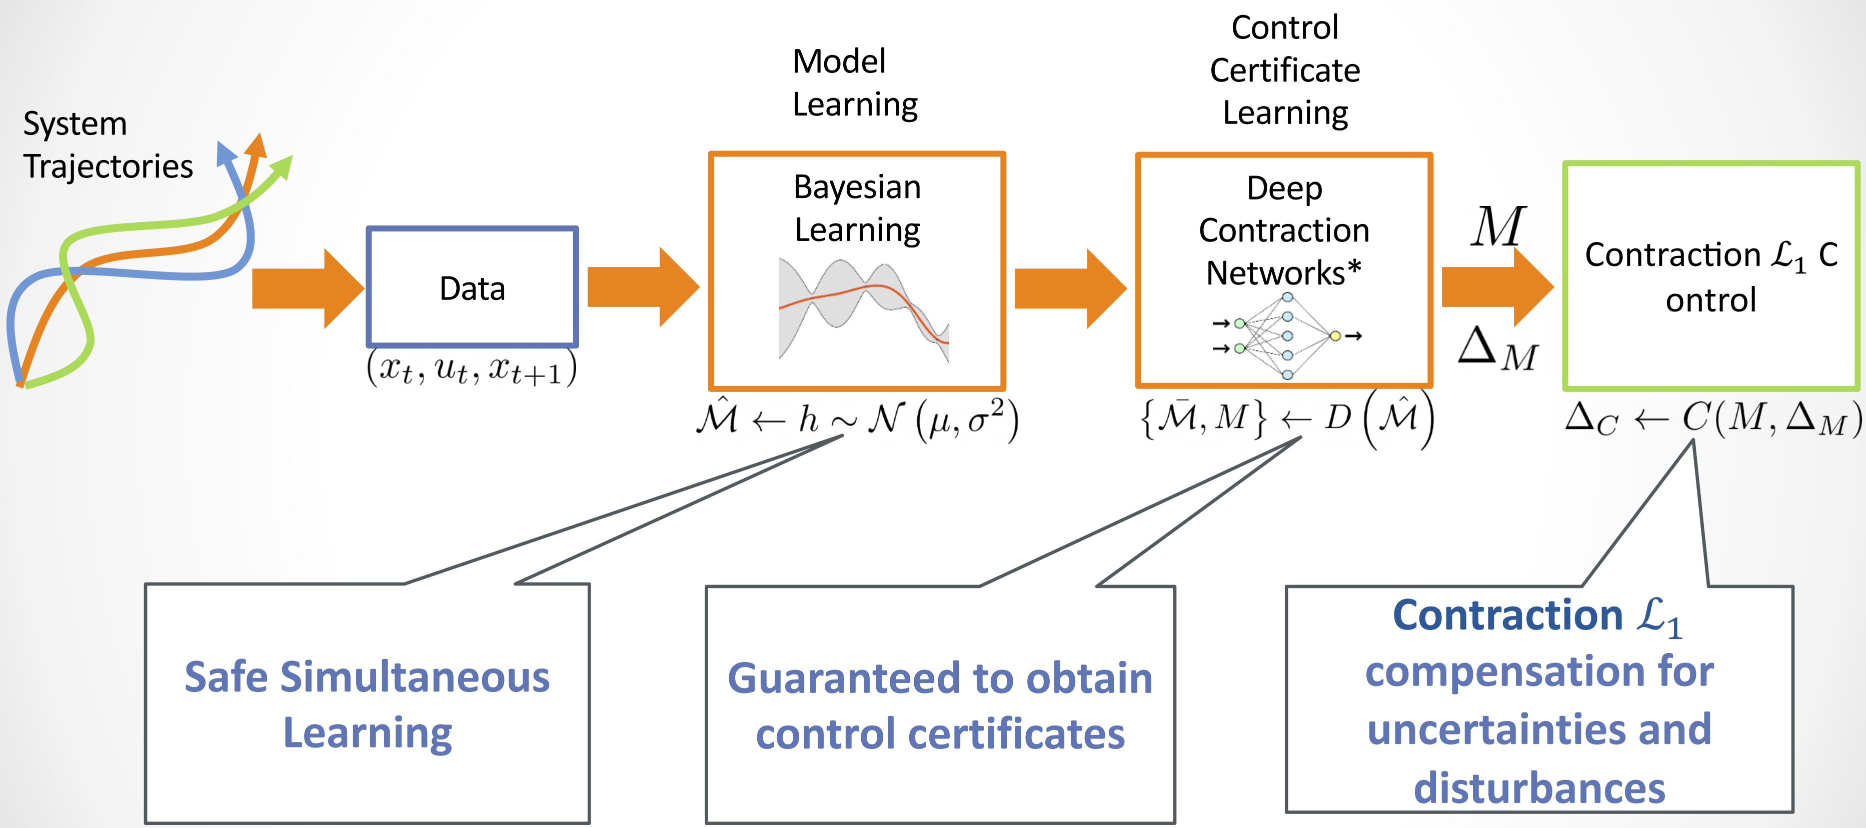

*Gahlawat, A., Lakshmanan, A., Vlahov, B. Gandhi, M. Song, L, Hovakimyan, N., and Theodorou, E. "Deep Incrementally Stabilizing Control. "*In preparation.*

A few ***real-world problems*** that need the robustness and performance guarantees provided by our safe learning and control framework.

nhovakim@Illinois.edu

Learn-to-Fly (L2F)

Conventional aircraft development process

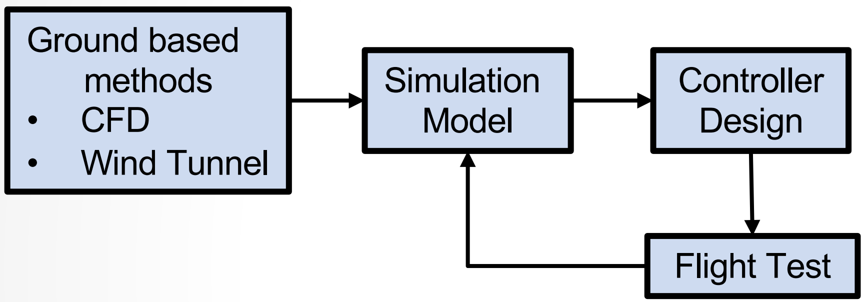

L2F concept*

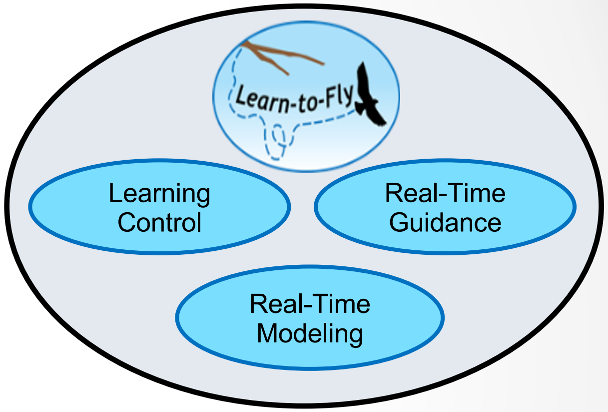

*E. H. Heim, et al. "NASA's Learn-to-Fly Project Overview." In *2018 Atmospheric Flight Mechanics Conference*, 2018.

Learn-to-Fly (L2F)

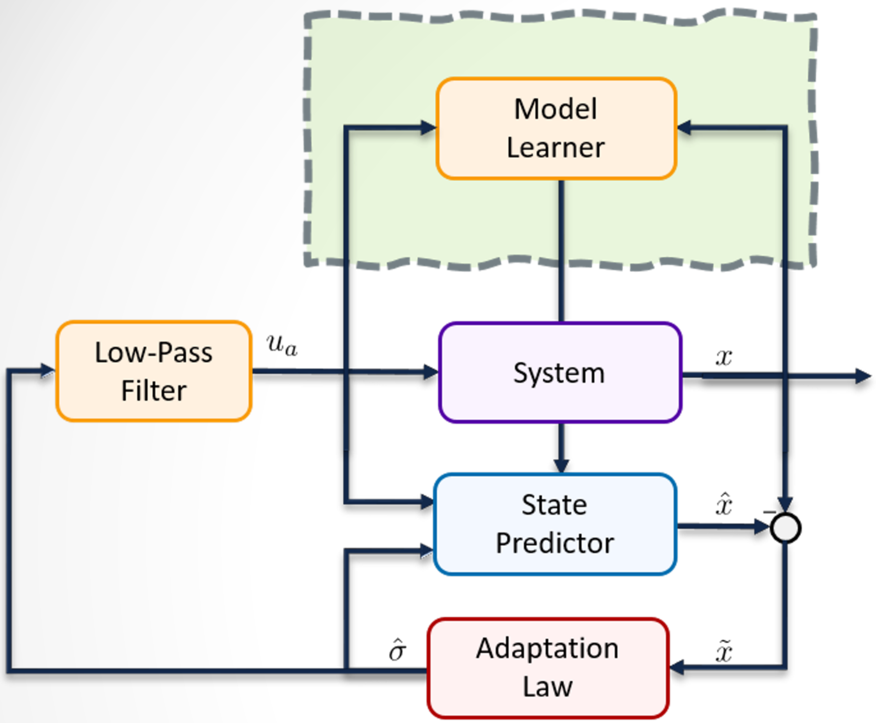

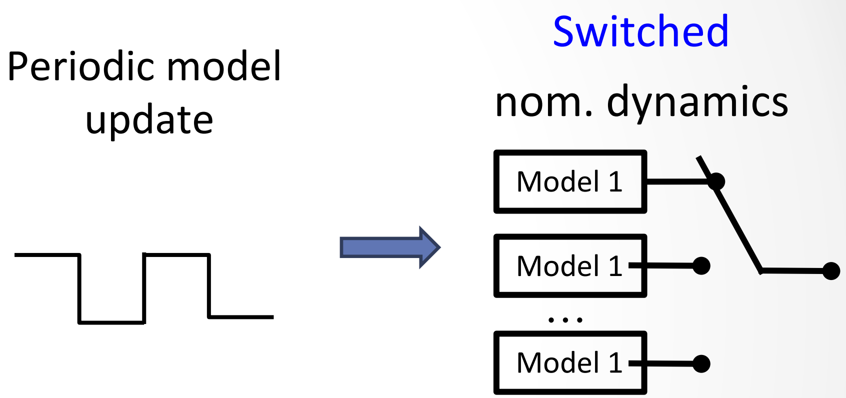

$$\dot{x}(t) = A_p x(t) + B_p\left(\omega_p u(t) + f_p(t,x)\right)$$

- Learned model is used to update the nominal dynamics
- Adaptive control handles the discrepancy between the learned model and the actual dynamics
- Periodic model update leads to switched nominal dynamics

S. Snyder, P. Zhao and N. Hovakimyan. L1 adaptive control for switching reference systems: Application to flight control. *IFAC-PapersOnLine*, 52(16), pp.718-723, 2019.

S. Snyder, P. Zhao*, N. Hovakimyan. L1 adaptive control for learn-to-fly validated by flight tests. *In preparation*.

Flight Tests at NASA LaRC

- Flight tests on NASA Woodstock and E-1 aircraft performed at NASA LaRC

This video demonstrates the flight test of a Learn-to-Fly system supported by an L1 adaptive augmentation on the E-1 aircraft.

Unstable aircraft emulation by activating the hidden feedback

- Challenging for the pilot to control
- L2F able to achieve level flight after initial learning phase

S. Snyder, P. Zhao*, N. Hovakimyan. L1 adaptive control for learn-to-fly validated by flight tests. *In preparation*.

Deep Perceptual Adaptive Control

- Uncertainty-driven safety guarantees for the full autonomy stack from perception to low-level control
- Quantify performance as a function of available hardware and software resources
- Dynamic resource allocation:
  - **Fast environment**: Perception, high-rate planning, and adaptation prioritized.
  - **Slow environment**: Safe model learning.

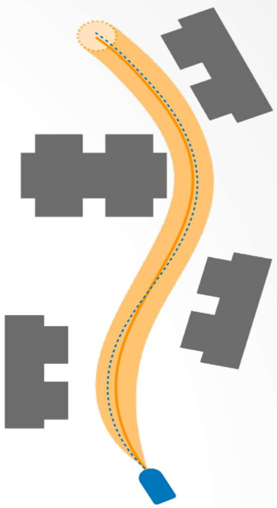

A. Lakshmanan, A. Gahlawat and N. Hovakimyan, Safe Feedback Motion Planning: A Contraction Theory and L1-Adaptive Control Based Approach, In Proceedings of Control and Decision Conference, Jeju Island, Korea, 2020, accepted.

NASA's Urban Aerial Mobility platforms

P. Drews, G. Williams, B. Goldfain, E. A., Theodorou, and J. M. Rehg, "Vision-based high-speed driving with a deep dynamic observer." IEEE Robotics and Automation Letters, 4(2), 1564-1571, 2019

Virtual Sully: Towards Full Pilotless Autonomy

- Airbus 320 lost both engines shortly after takeoff due to bird strike.
- Captain Sullenberger made the **correct decision** by landing on the Hudson river instead of attempting return to LaGuardia
- Long years of experience and training of neurons in human brain lead to correct decisions (**what can we learn from neuroscience?**)

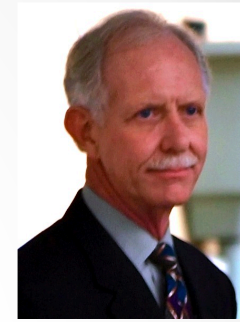

Credits: (left) Greg L., (right) I. Talyar

Mission Objectives ⇄ Task and Motion Planning ⇄ Tracking & Stabilizing Control

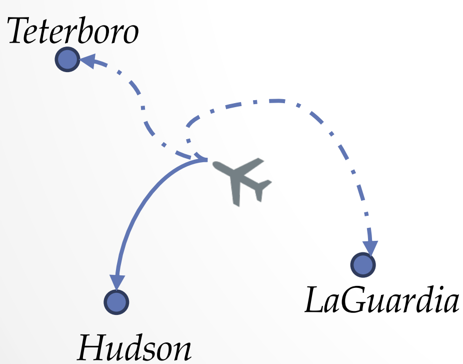

Safest landing option

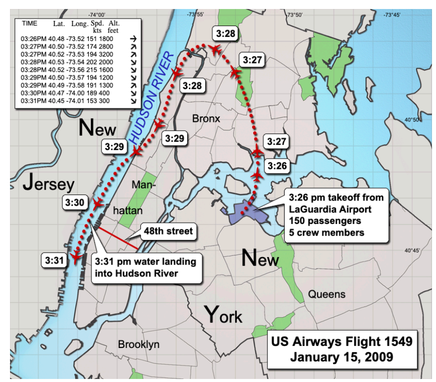

Safe path to the Hudson river

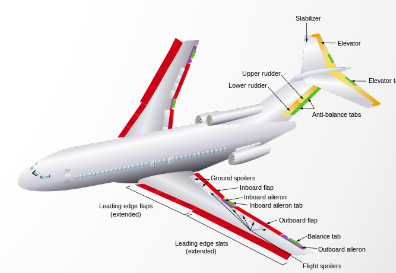

Safe landing with just the available control surfaces

nhovakim@Illinois.edu

Conclusions

- Developed advanced **safe trajectory tracking** control algorithms for complex and uncertain environments.
- **Contraction-based control** ensures exponential convergence of the nominal system to some planned trajectory.
- Relies on **fast adaptation** to immediately compensate for disturbances and **slow adaptation** to learn the unmodeled dynamics over time.
- Performance guarantees (as tubes) and robustness margins (as filter design) can be tuned to **ensure safety during the entire learning process.**

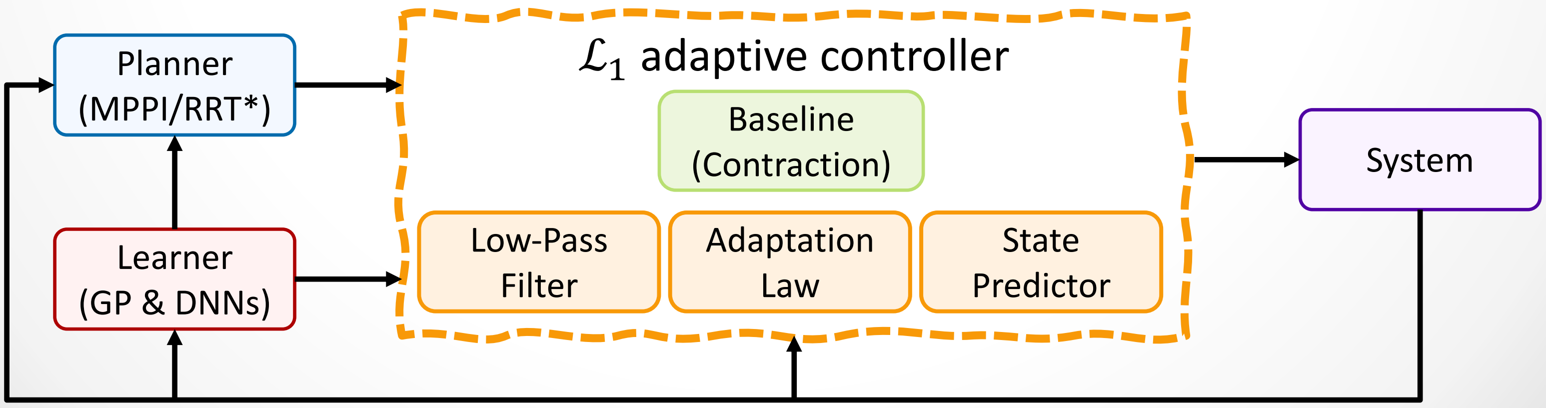

Current collaborators

Evangelos Theodorou

Petros Voulgaris

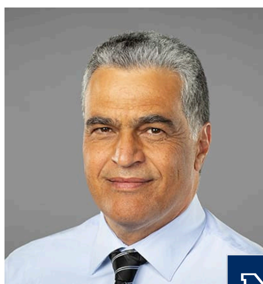

Xiaofeng Wang

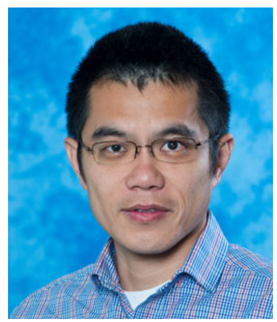

Lui Sha

Marco Pavone

Srinivasa Salapaka

nhovakim@Illinois.edu

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 (Seagate)
Hui Sun (Marvell Semiconductor)
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 (University of Nevada)
Alexandre Barbosa (Amazon Robotics)
Kasey Ackerman (NASA)
Javier Puig Navarro (NASA)
Arun Lakshmanan (Optimum Ride)
Gabriel Barsi Haberfeld (Apple)

Current Postdocs:
Aditya Gahlawat
Hunmin Kim
Pan Zhao
Yanbing Mao

Current Graduate Students:
Andrew Patterson, Yikun Cheng
Hyungsoo Kang, Neng Wan
Wenbin Wan, Zhuohuan Wu
Sitao Zhang, Erin Swansen
Lin Song, Ziyao Guo
Vivek Sharma, Min Jun Sung
Yuliang Gu, Jing Wu
Chuyuan Tao

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.