Deployable Robotics (Part I)

Russ Tedrake

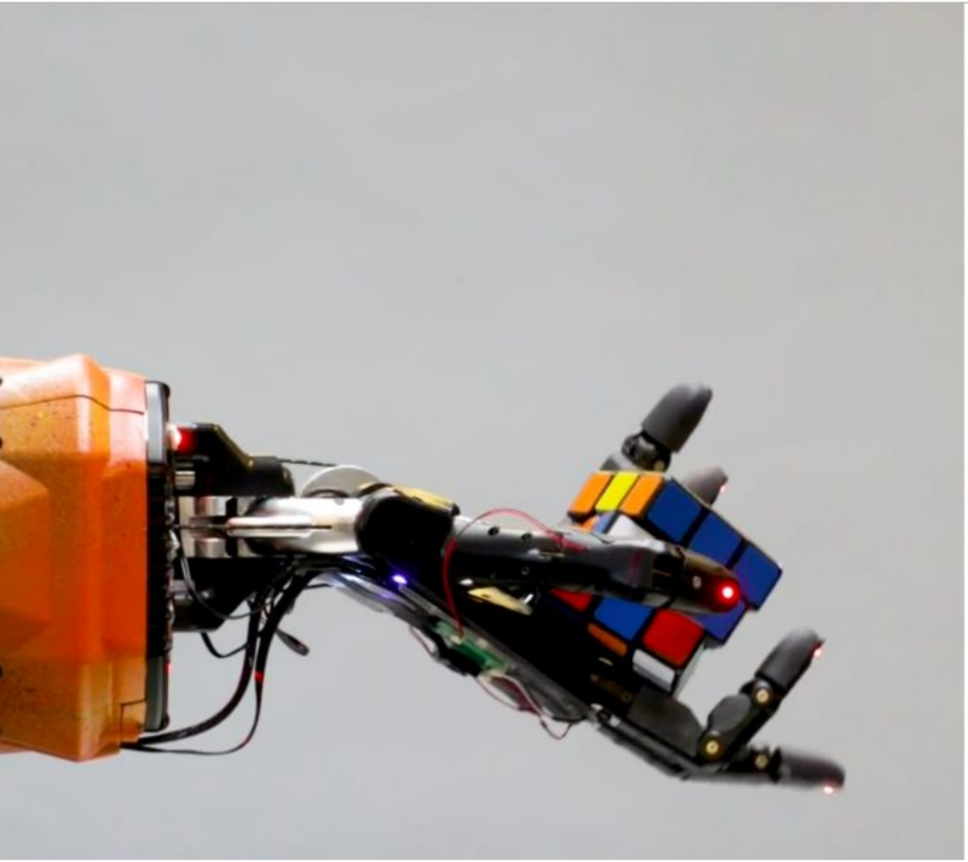

The New York Times

If a Robotic Hand Solves a Rubik's Cube, Does It Prove Something?

A five-fingered feat could show important progress in A.I. research. It is also a stunt.

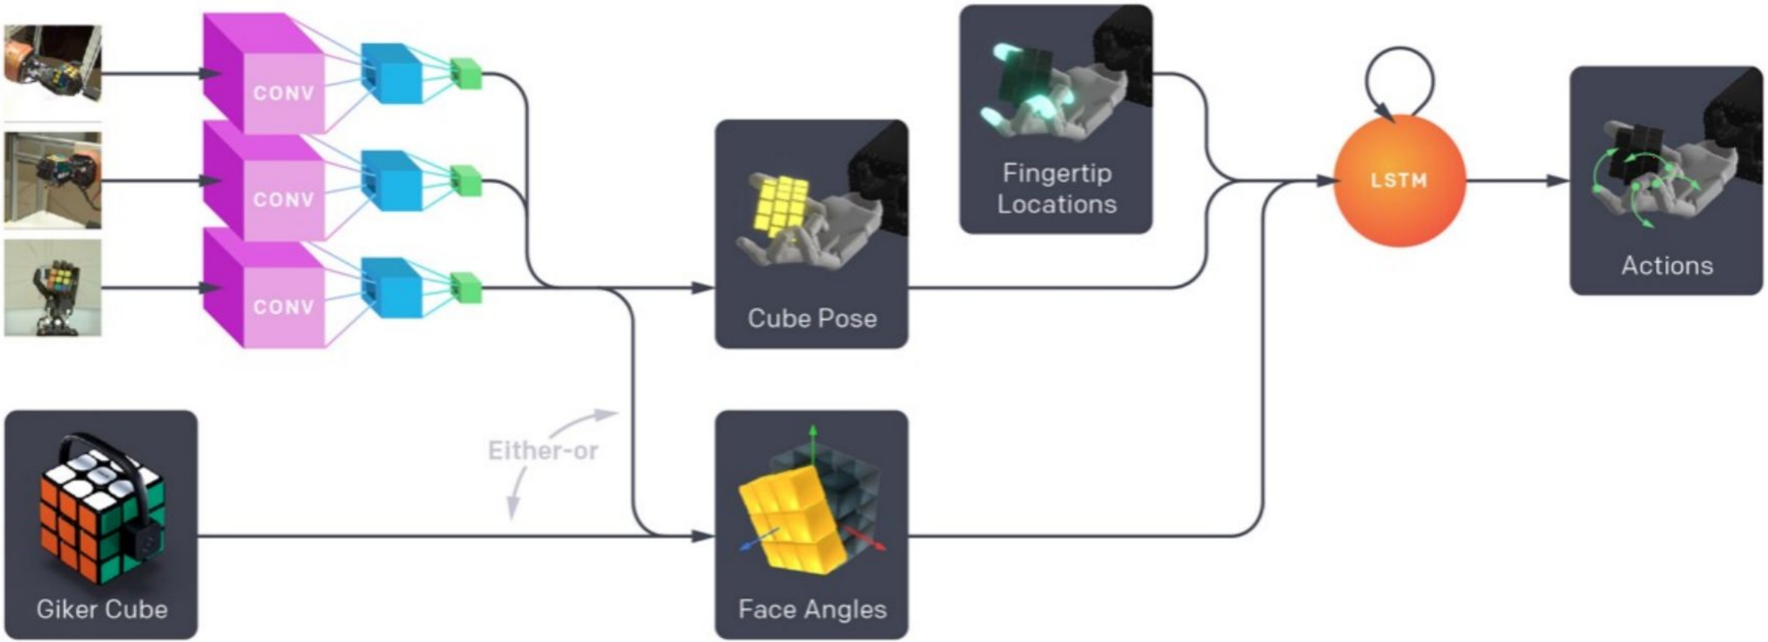

"For the Rubik's cube task, we use 8 × 8 = 64 NVIDIA V100 GPUs and 8 × 115 = 920 worker machines with 32 CPU cores each. … The cumulative amount of experience ... is roughly **13 thousand years**."

Solving Rubik's Cube with a Robot Hand by OpenAI, arXiv:1910.07113

Table 6: Performance of different policies on the Rubik's cube for a fixed fair scramble goal sequence. We evaluate each policy on the real robot (N=10 trials) and report the mean $\pm$ standard error and median number of successes (meaning the total number of successful rotations and flips). We also report two success rates for applying half of a fair scramble ("half") and the other one for fully applying it ("full"). For ADR policies, we report the entropy in nats per dimension (npd). For "Manual DR", we obtain an upper bound on its ADR entropy by running ADR with the policy fixed and report the entropy once the distribution stops changing (marked with an "*").

Policy	**Sensing**		**ADR Entropy**	**Successes (Real)**		**Success Rate**	
	Pose	**Face Angles**		**Mean**	**Median**	**Half**	**Full**
Manual DR	Vision	Giiker	-0.569^{*} npd	1.8 ± 0.4	2.0	0 %	0 %
ADR	Vision	Giiker	-0.084 npd	3.8 ± 1.0	3.0	0 %	0 %
ADR (XL)	Vision	Giiker	0.467 npd	17.8 ± 4.2	12.5	30 %	10 %
ADR (XXL)	Vision	Giiker	**0.479 npd**	$\mathbf{26.8 \pm 4.9}$	**22.0**	**60 %**	**20 %**
ADR (XXL)	Vision	Vision	**0.479 npd**	12.8 ± 3.4	10.5	20 %	0 %

Solving Rubik's Cube with a Robot Hand by OpenAI, arXiv:1910.07113

Pablo’s lectures

Introduction to (robust) control and Lyapunov; we’ll do a bit more of that here...

Billion dollar question:

What will be the *epistemology* of deployable ML?

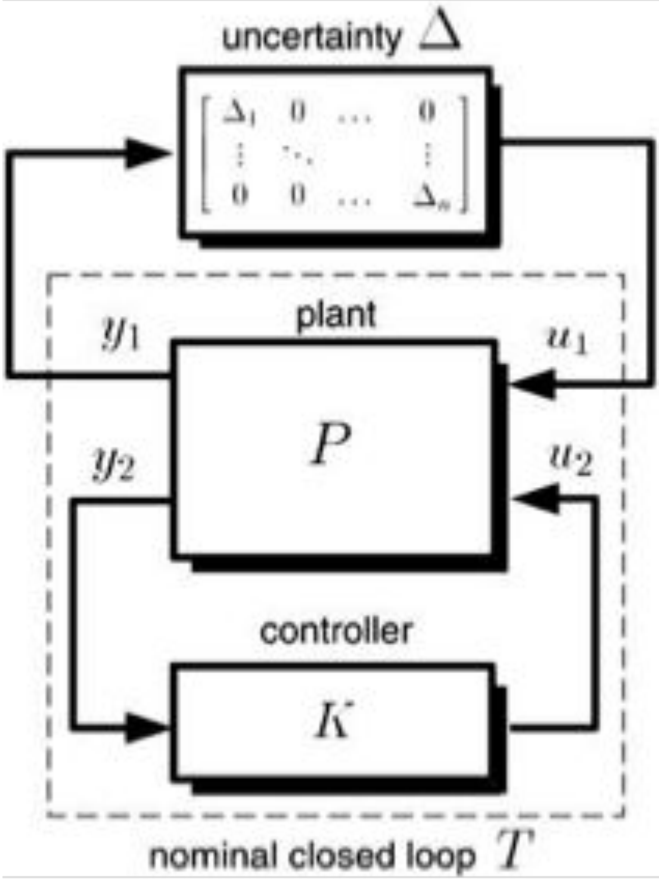

Challenge #1: System Complexity

Deploying Autonomous/Learning Systems

The complexity of perception breaks our existing tools…

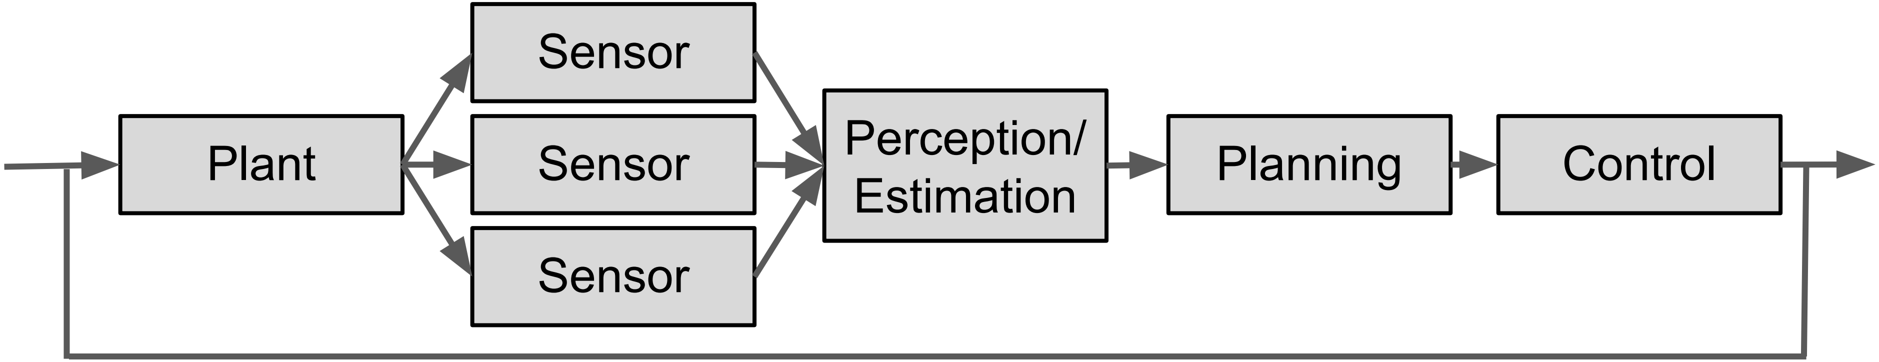

- Sensors include cameras ⇒ sensor model is a photo-realistic rendering engine
- Perception components (especially) include deep neural networks; but verifying planning algorithms also nontrivial.
- Plant model has to capture distributions over natural scenes (numbers/types of objects, material properties, lighting conditions)

Challenge #2: Distributional Robustness and Black Swans

What do you see in this Picture?

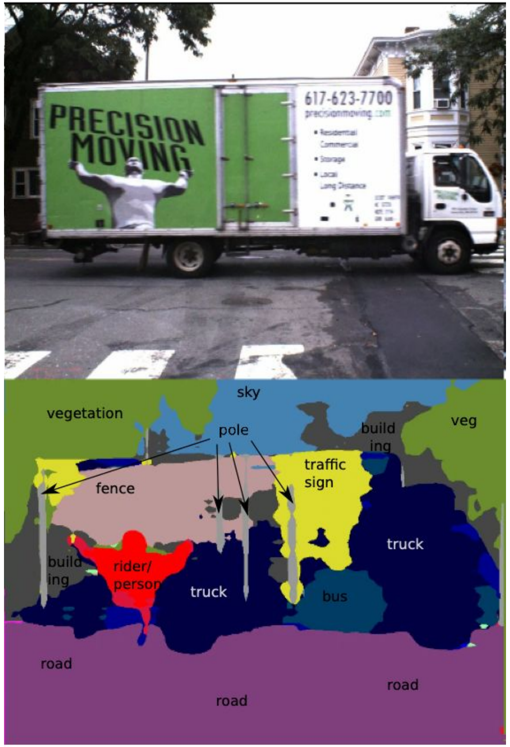

Courtesy: John Leonard

Coolidge Corner, Brookline MA

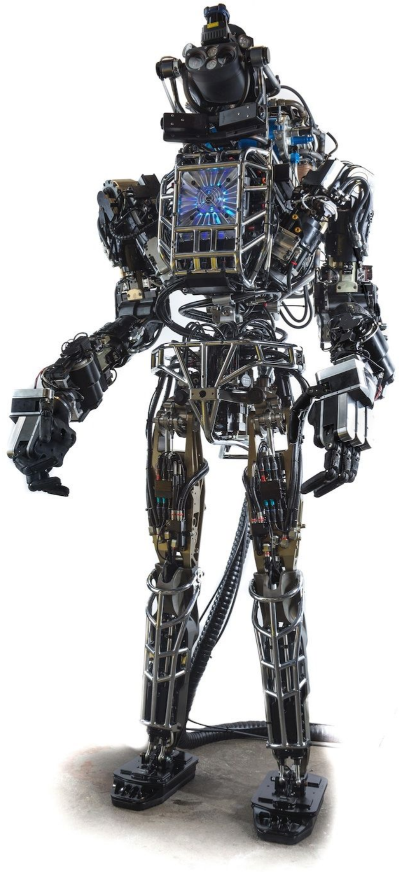


DANGER
HIGH
VOLTAGE

My lesson in robustness

In a garage at MIT just days before the competition…

Now passionate to understand how to get robustness from these complex systems.

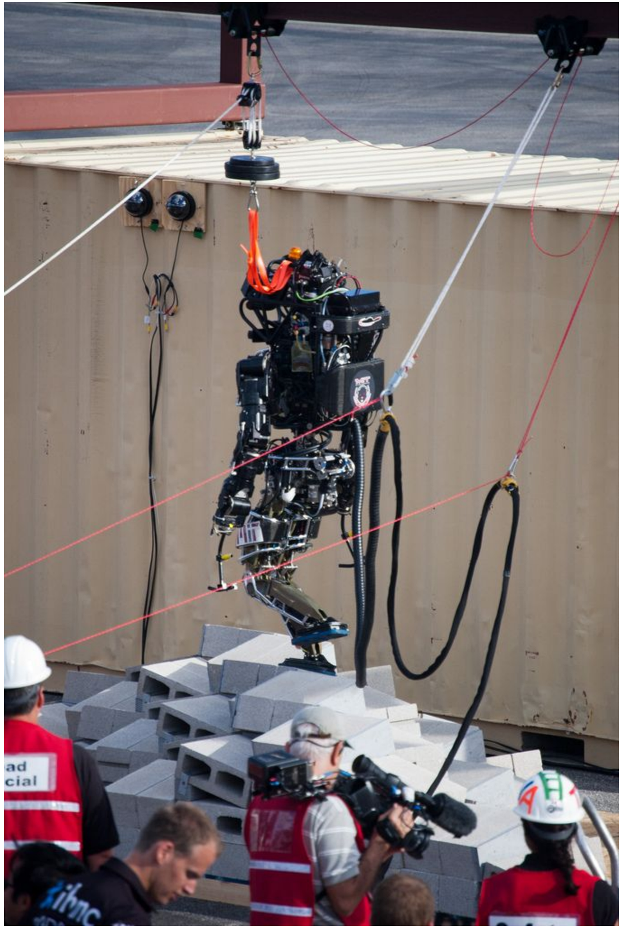

Challenge #3: High expectations

(how safe do we have to be to deploy?)

RAND > Press Room > News Releases > 2016 >

Autonomous Vehicles Cannot Be Test-Driven Enough Miles to Demonstrate Their Safety; Alternative Testing Methods Needed

FOR RELEASE
Tuesday
April 12, 2016

Autonomous vehicles would have to be driven hundreds of millions of miles and, under some scenarios, hundreds of billions of miles to create enough data to clearly demonstrate their safety, according to a new RAND report.

Media Resources

RAND Office of Media Relations

(703) 414-4795
(310) 451-6913
media@rand.org

Researcher Spotlight

Analysis (via a very simple coin flipping model): To estimate to within 20% of assumed rate (1.09/100 million), with 95% confidence, requires ~ 8.8 billion miles.

Releases of technology in the airline industry

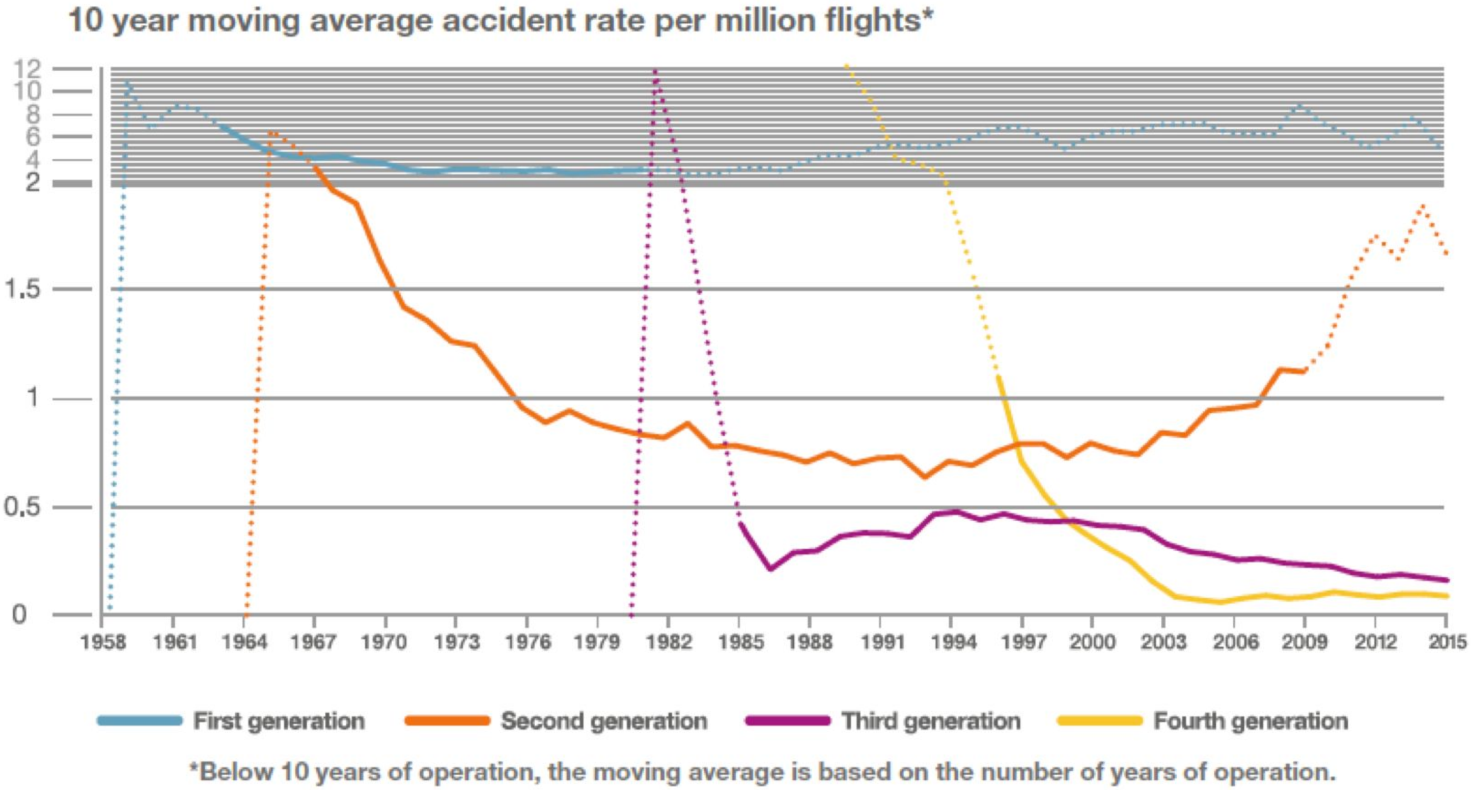

"A Statistical Analysis of Commercial Aviation Accidents:1958-2015" by Airbus

Goal: Toolkit for reasoning about about **uncertainty** in **closed-loop systems**

Uncertainty representations

In controls (polytopic/ellipsoidal, etc)

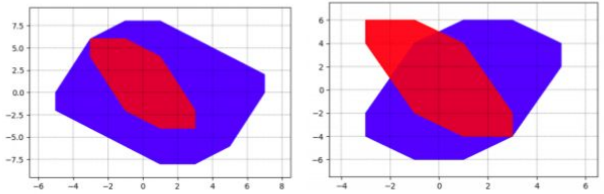

Fig. 1. Example 1: Zonotope Containment Problem: [left] $\mathbb{Z}_l \subseteq \mathbb{Z}_r$, [Right] $\mathbb{Z}_l \not\subseteq \mathbb{Z}_r^*$, where the last column of G_r is dropped.

Domain randomization in reinforcement learning

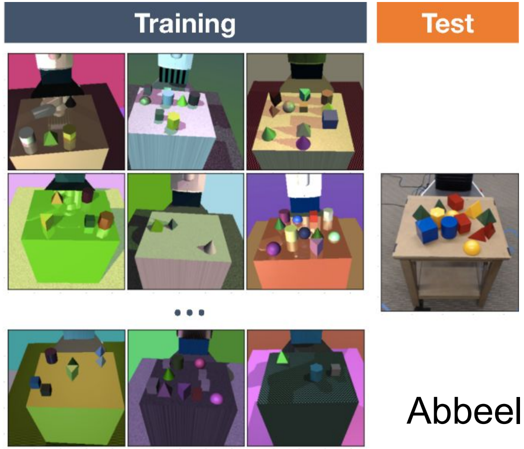

Abbeel et al.

27 Feb 2018 | 16:48 GMT

Creating Driving Tests for Self-Driving Cars

Volvo-backed Zenuity wants to prove that autonomous vehicles can drive more safely than humans

By **Erik Coelingh and Jonas Nilsson**

Developing autonomous systems in the real world.

IEEE Spectrum

Common Lyapunov Functions and Invariant Sets

On the board.

Relevant course notes (from 6.832) are here.

Region of attraction for the (time-reversed) van der Pol oscillator

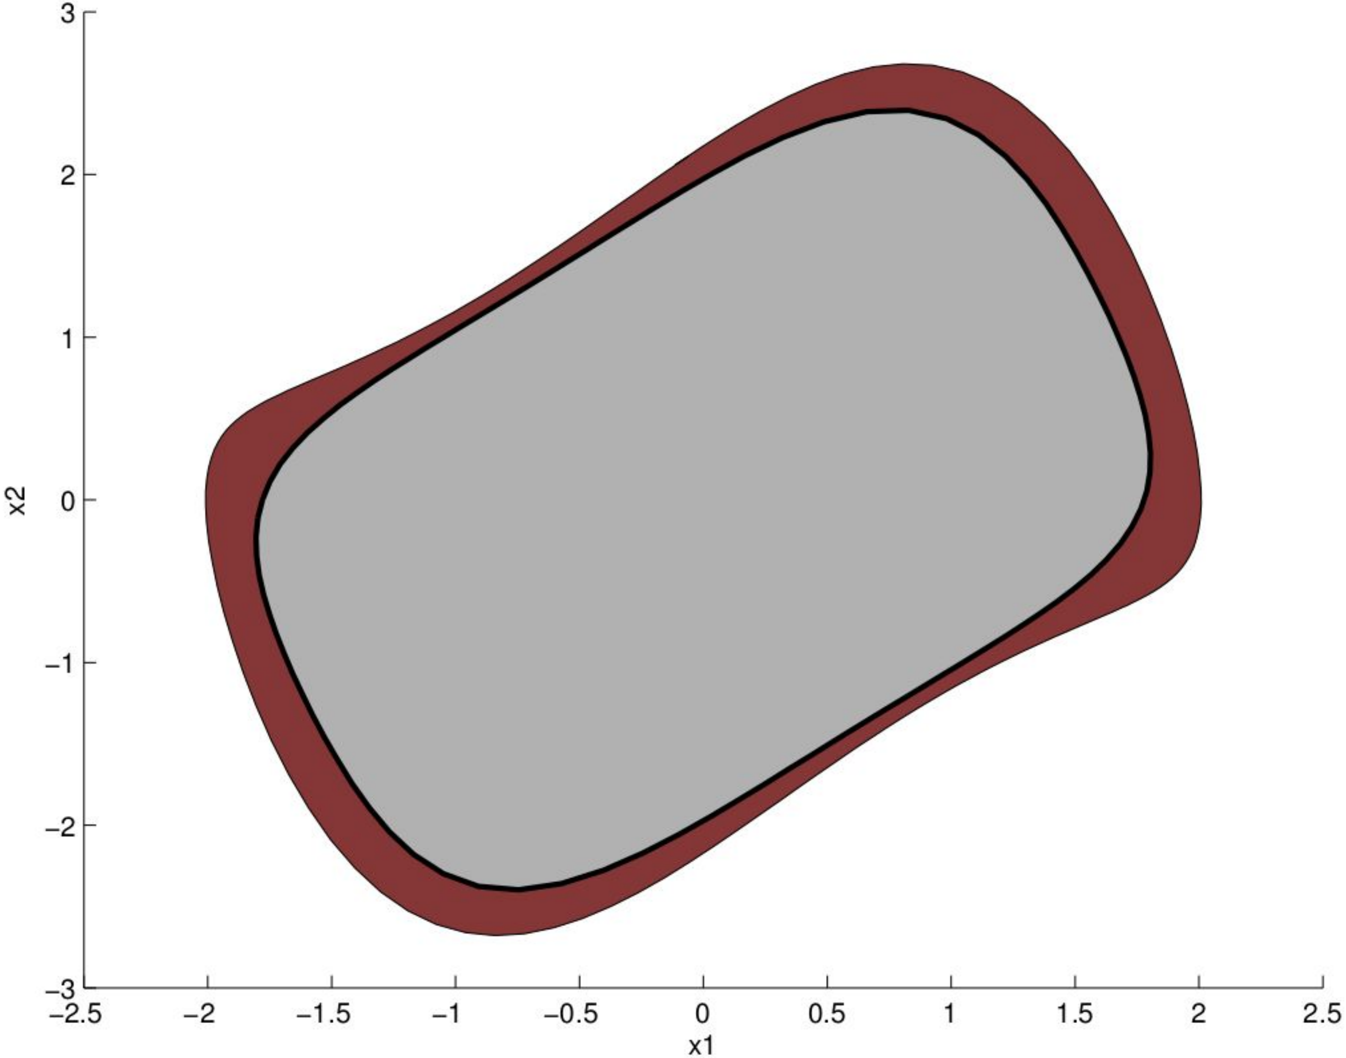

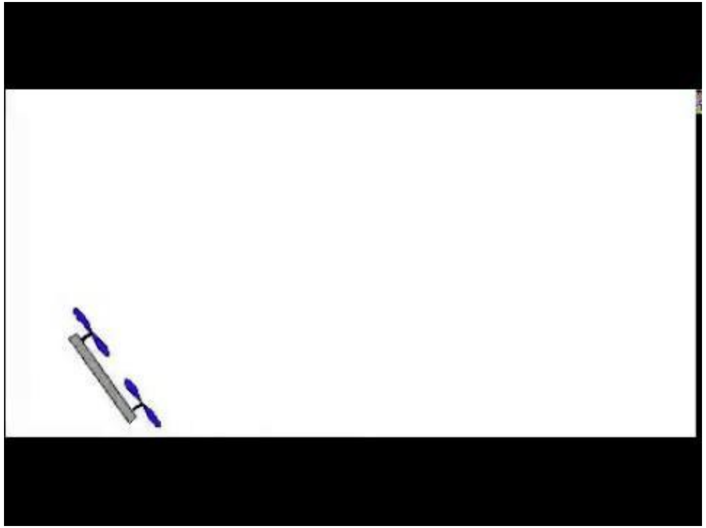

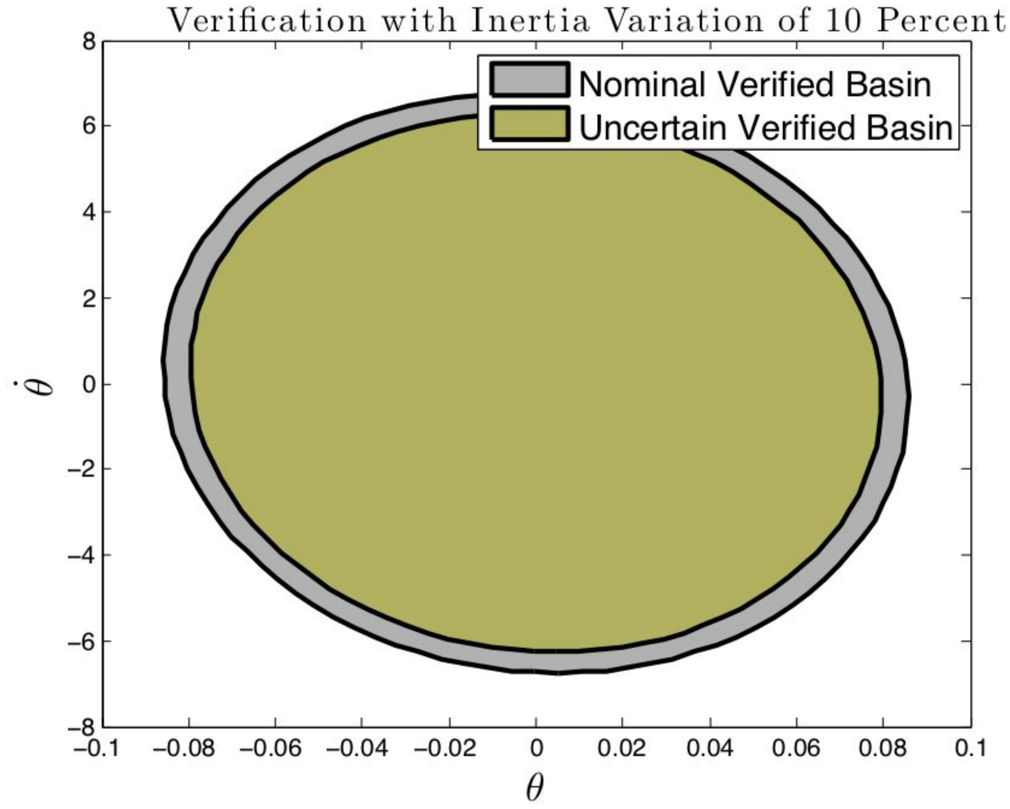
Verification with Inertia Variation of 10 Percent
Nominal Verified Basin
Uncertain Verified Basin
8
6
4
2
0
−2
−4
−6
−8
−0.1
−0.08
−0.06
−0.04
−0.02
0
0.02
0.04
0.06
0.08
0.1
$\dot{\theta}$
θ

Can we make a control system for a fixed-wing airplane to land on a perch like a bird?

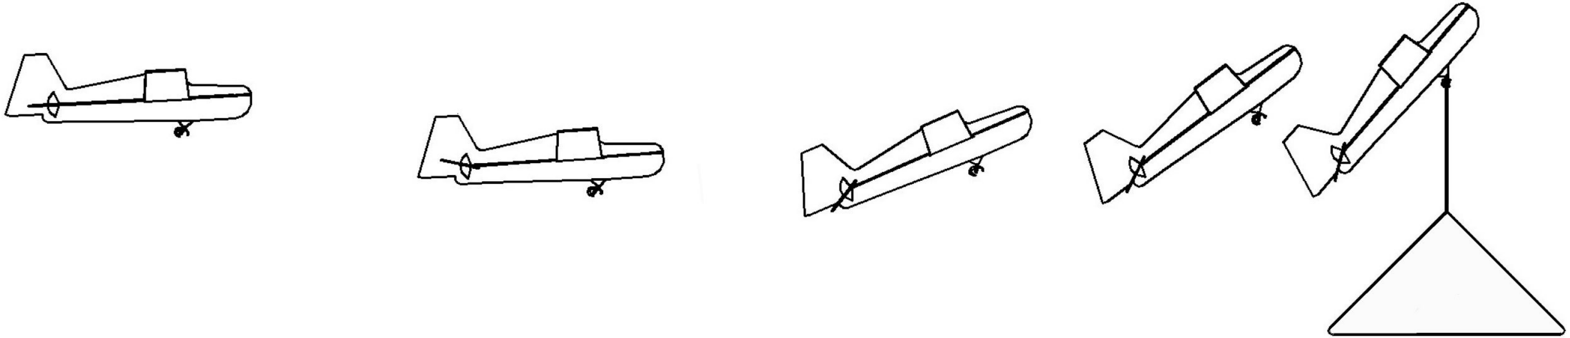

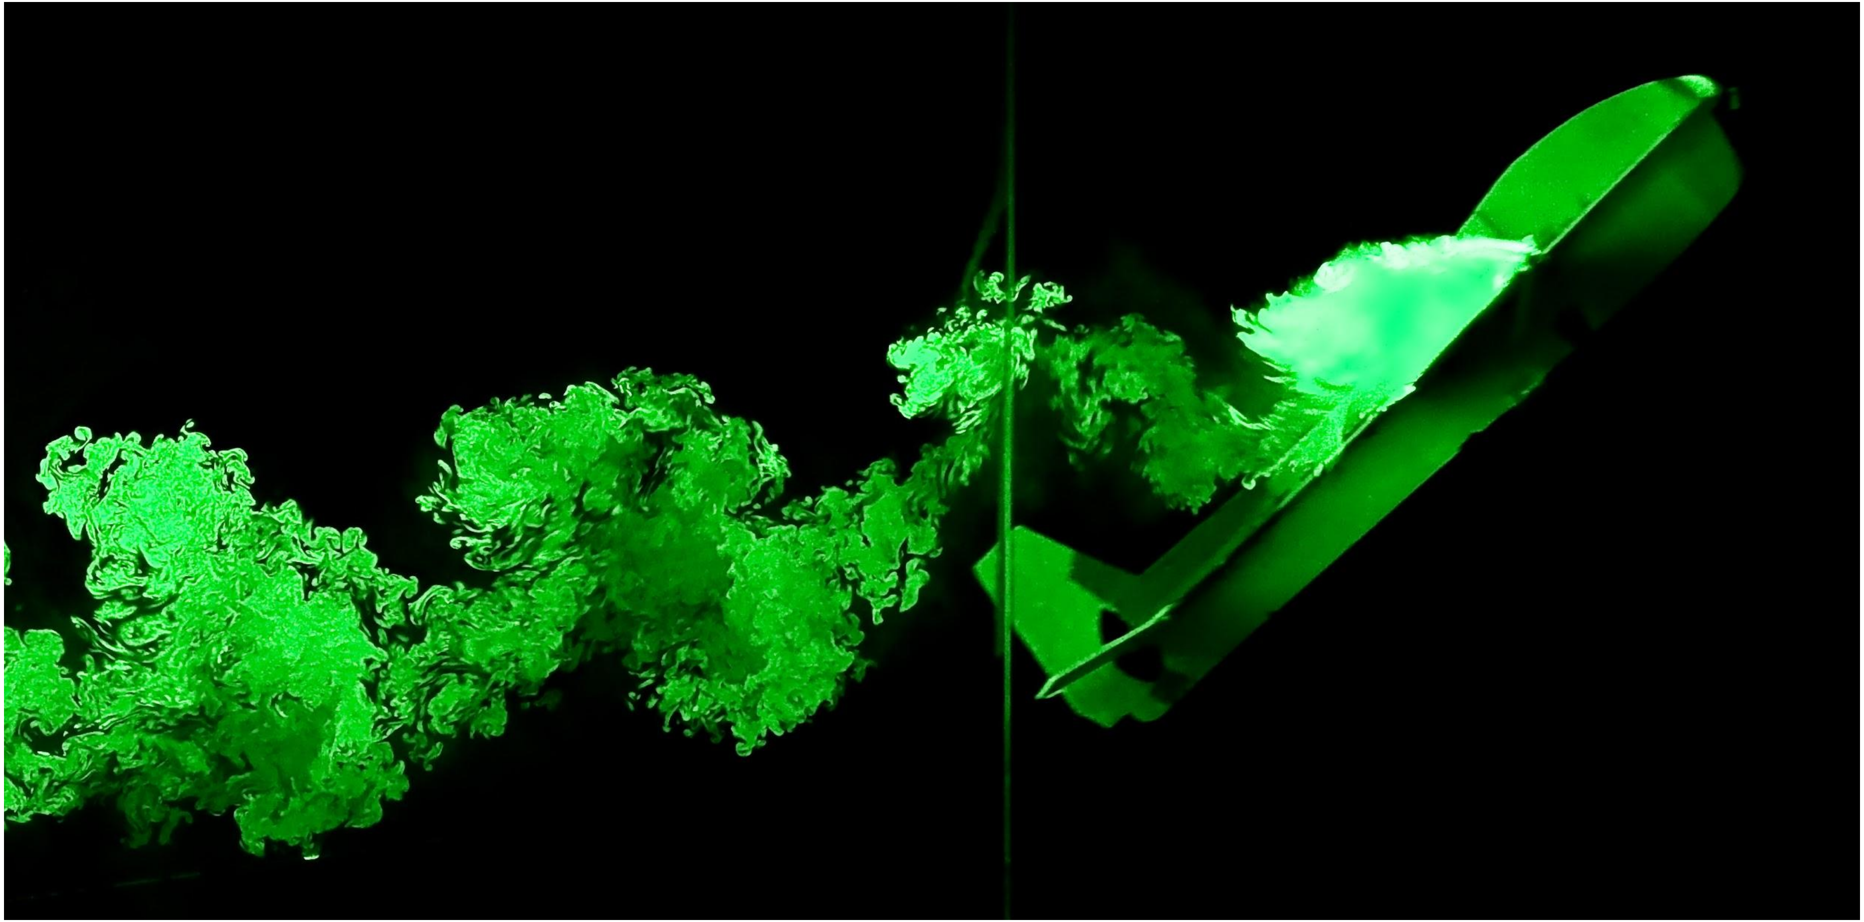

MIT Robot Locomotion Group

Lift Coefficient

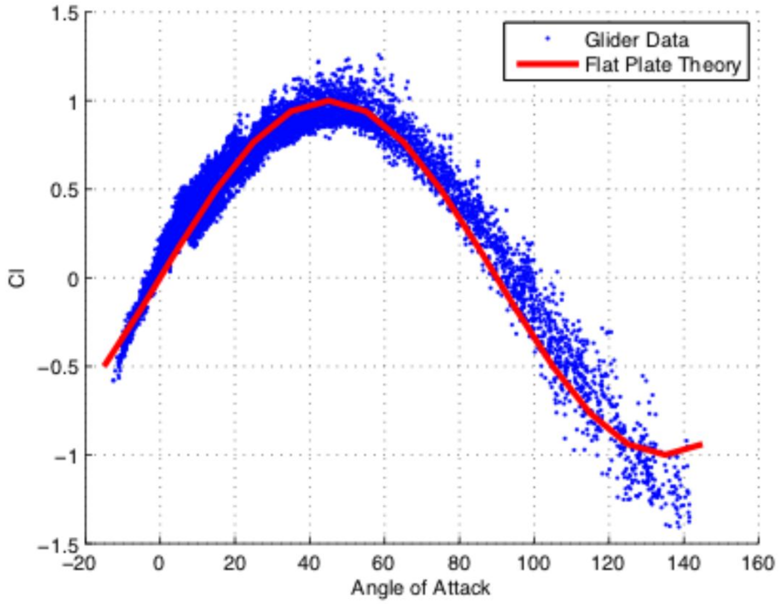

Drag Coefficient

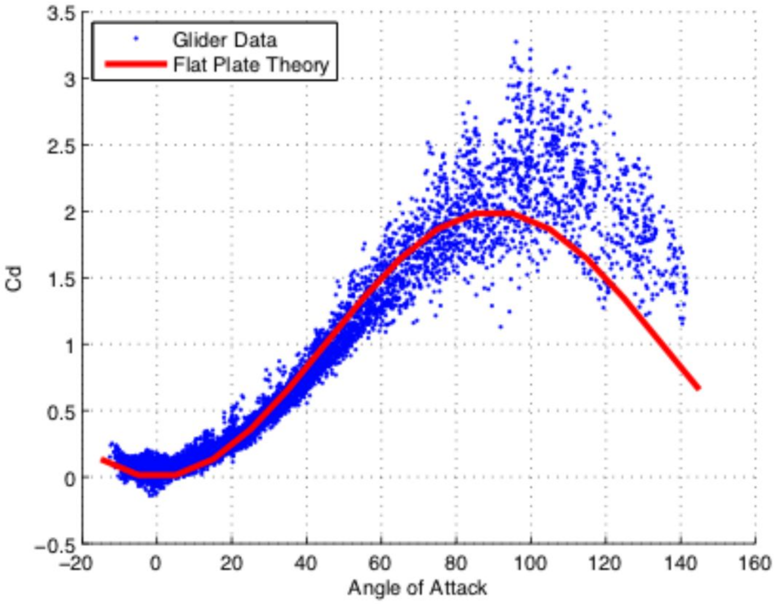

$$\dot{x} = f(x, u)$$

Nonlinear (post-stall) dynamics described well by **polynomial** diff eq.

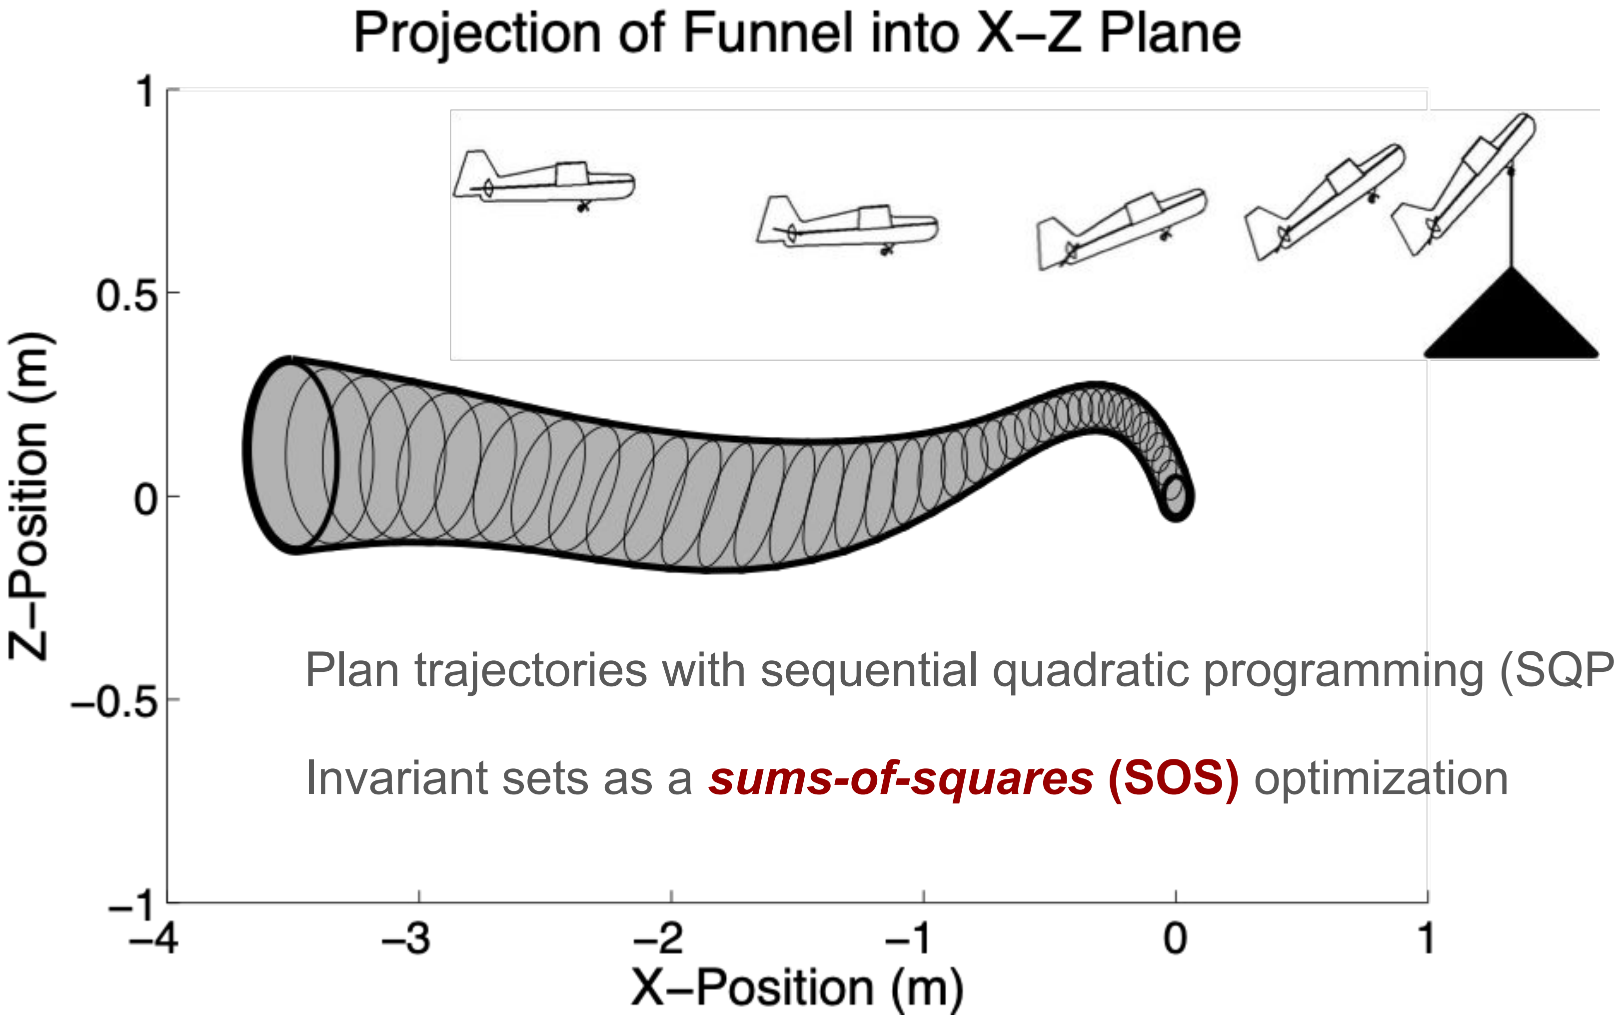

Plan trajectories with sequential quadratic programming (SQP)

Invariant sets as a ***sums-of-squares*** **(SOS)** optimization

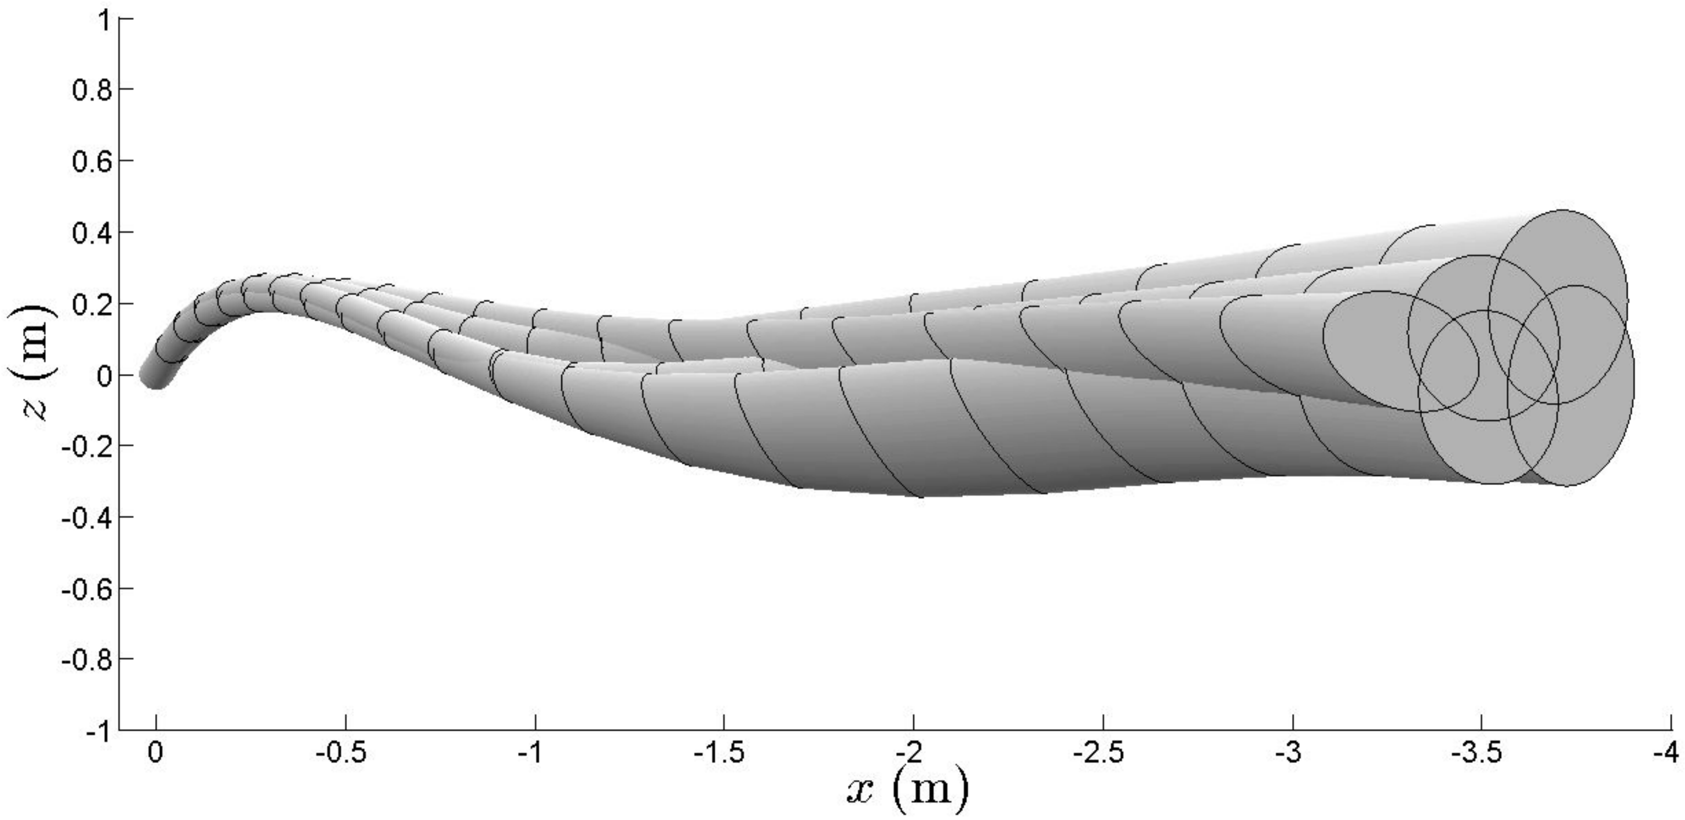

1
0.8
0.6
0.4
0.2
0
-0.2
-0.4
-0.6
-0.8
-1
z (m)
0
-0.5
-1
-1.5
-2
-2.5
-3
-3.5
-4
x (m)

Wind disturbances: colored “noise” drawn from ***ellipsoidal*** uncertainty set

Robust control via ***bilinear SOS*** alternations

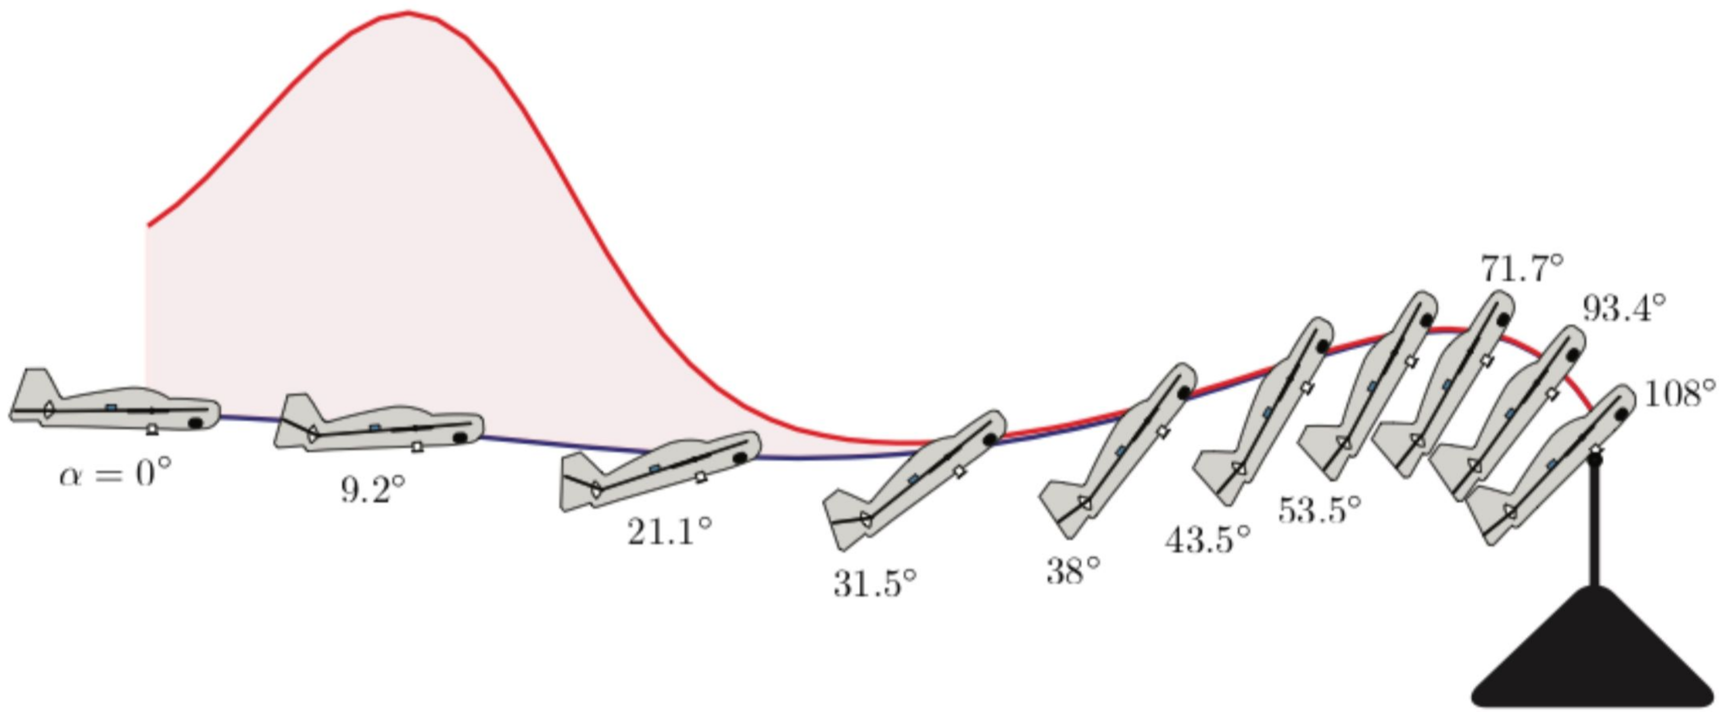
$\alpha = 0°$
9.2°
21.1°
31.5°
38°
43.5°
53.5°
71.7°
93.4°
108°

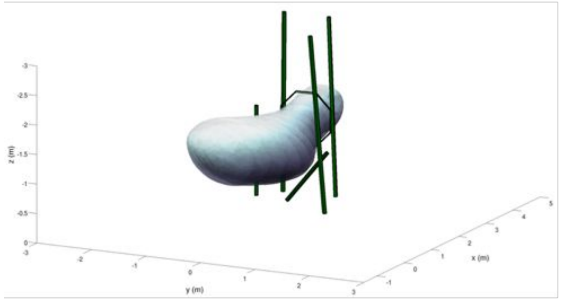

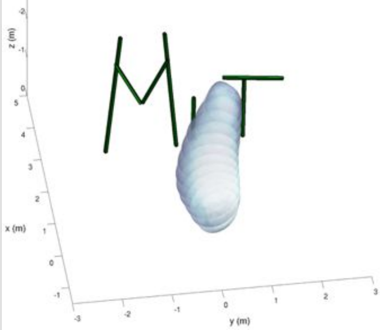

ONR MURI: Provable-safe high-speed flight through forests

Some final thoughts (on the board)