Navigation in Safety-Critical Environments

Cherie Ho

8/19/2019

air lab

Carnegie Mellon University
The Robotics Institute

We're increasingly deploying robots to safety-critical environments.

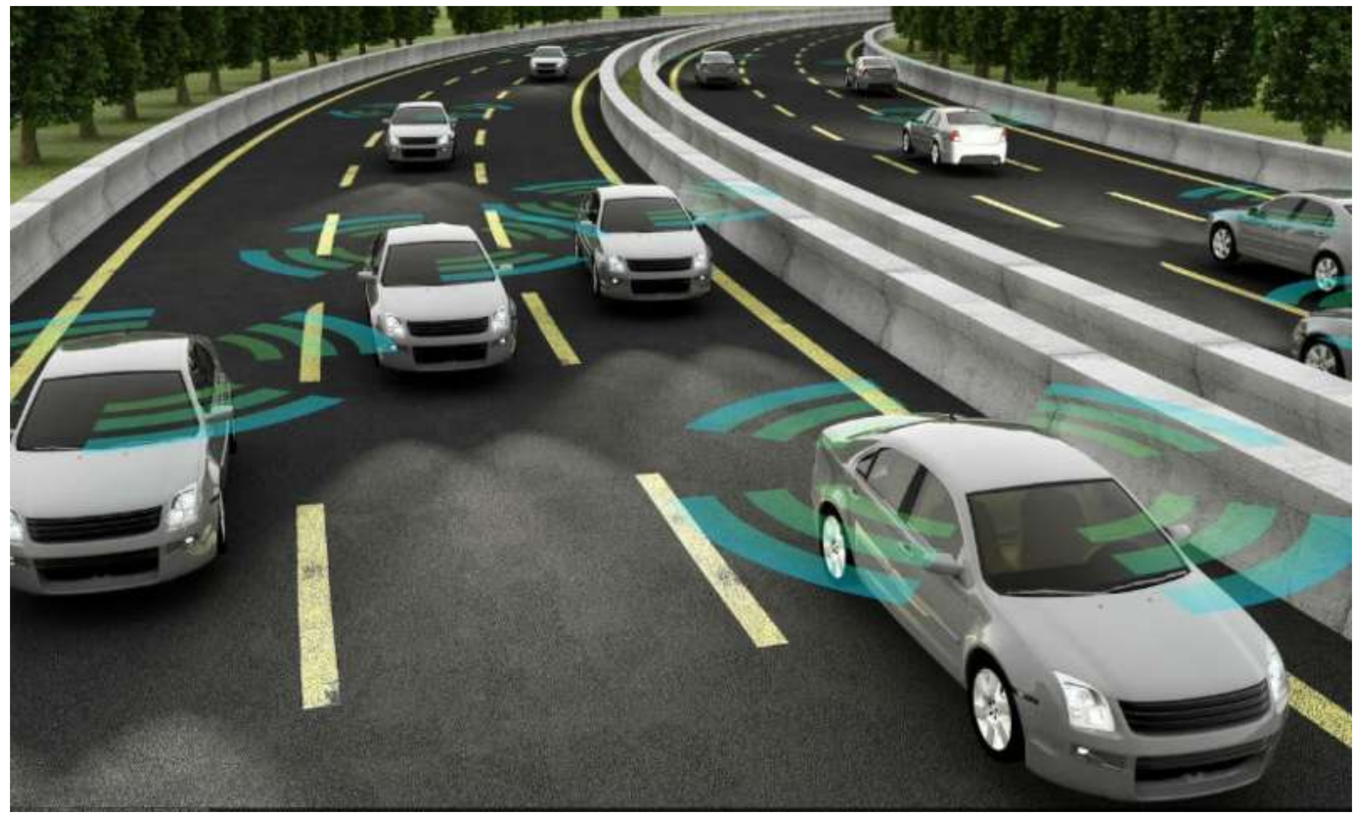

air lab

Requirement: Safety in a Predictable Manner

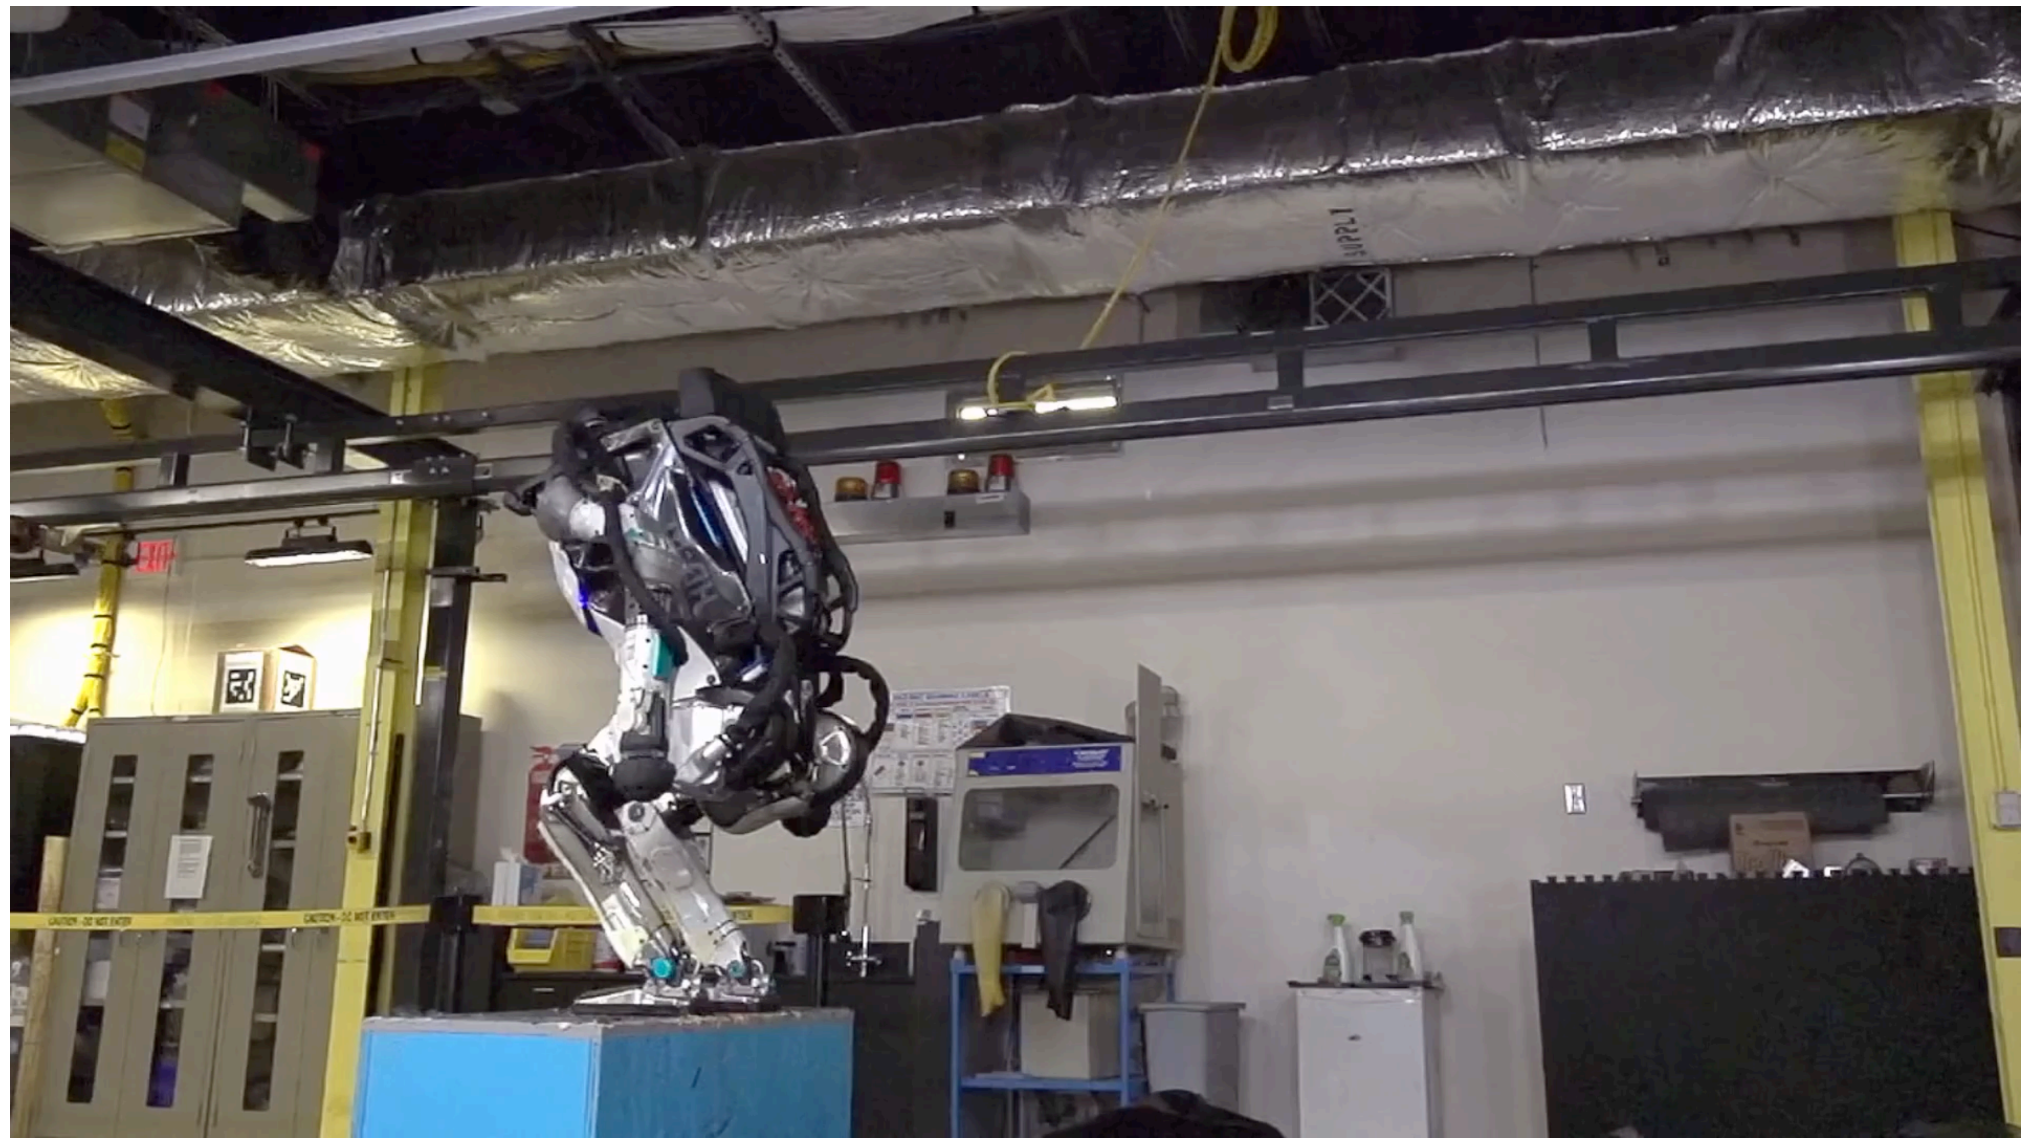

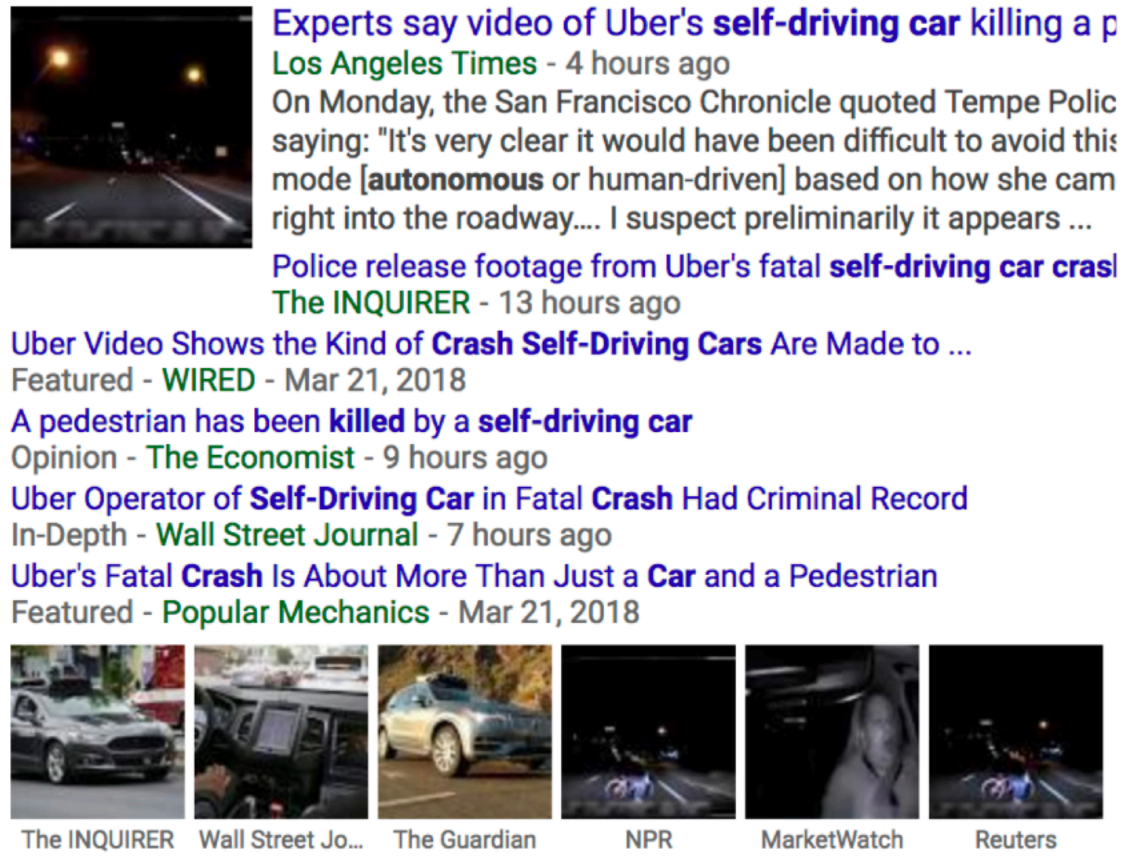
Experts say video of Uber's **self-driving car** killing a p
Los Angeles Times - 4 hours ago
On Monday, the San Francisco Chronicle quoted Tempe Polic
saying: "It's very clear it would have been difficult to avoid this
mode [**autonomous** or human-driven] based on how she cam
right into the roadway.... I suspect preliminarily it appears ...
Police release footage from Uber's fatal **self-driving car** **crasl**
The INQUIRER - 13 hours ago
Uber Video Shows the Kind of **Crash Self-Driving Cars** Are Made to ...
Featured - WIRED - Mar 21, 2018
A pedestrian has been **killed** by a **self-driving car**
Opinion - The Economist - 9 hours ago
Uber Operator of **Self-Driving Car** in Fatal **Crash** Had Criminal Record
In-Depth - Wall Street Journal - 7 hours ago
Uber's Fatal **Crash** Is About More Than Just a **Car** and a Pedestrian
Featured - Popular Mechanics - Mar 21, 2018
The INQUIRER | Wall Street Jo... | The Guardian | NPR | MarketWatch | Reuters

air lab

Case Study: AACUS Helicopter

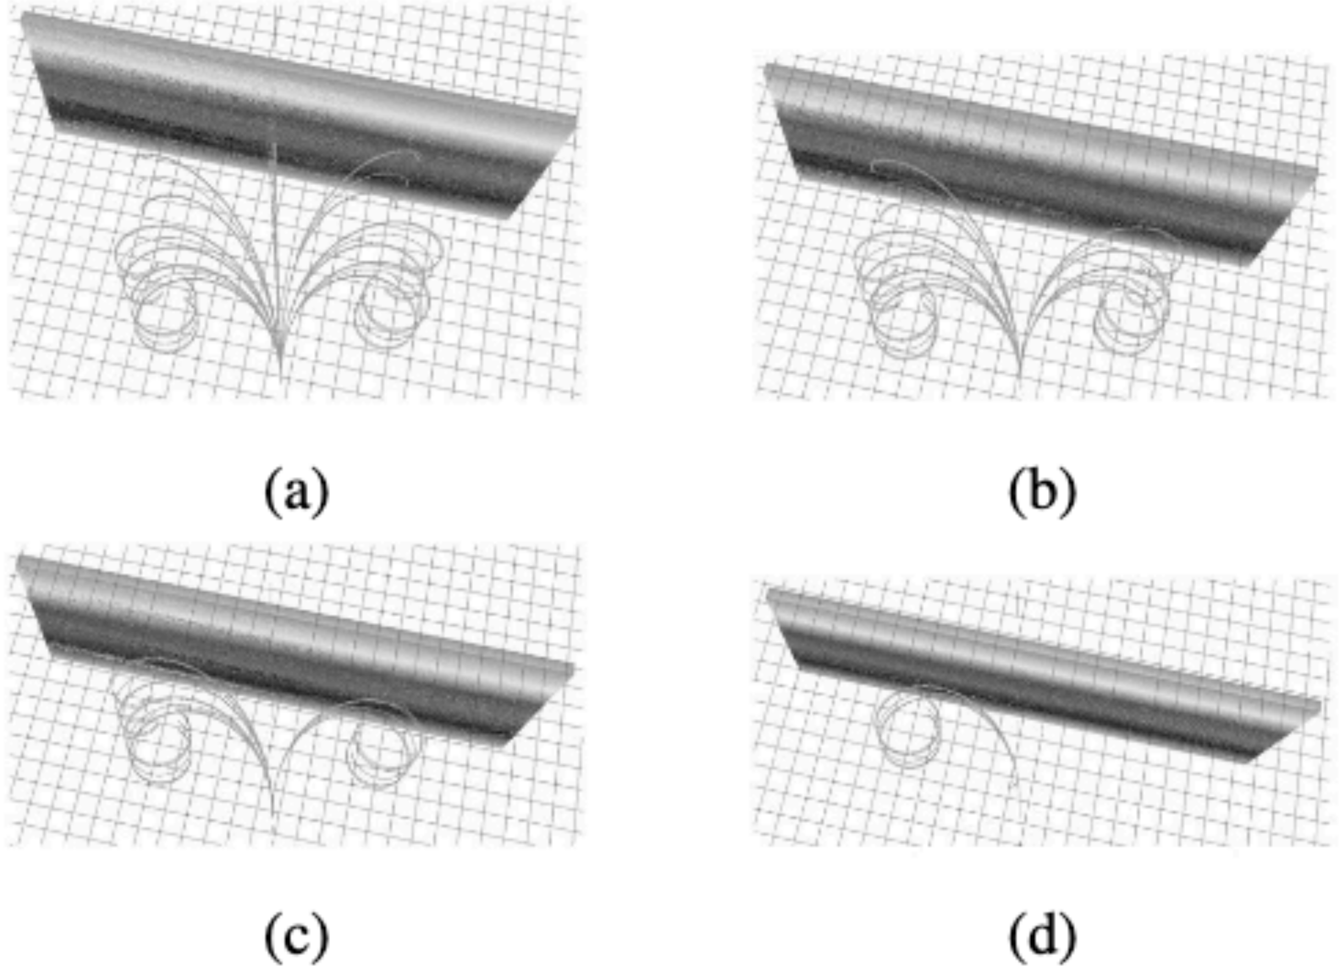

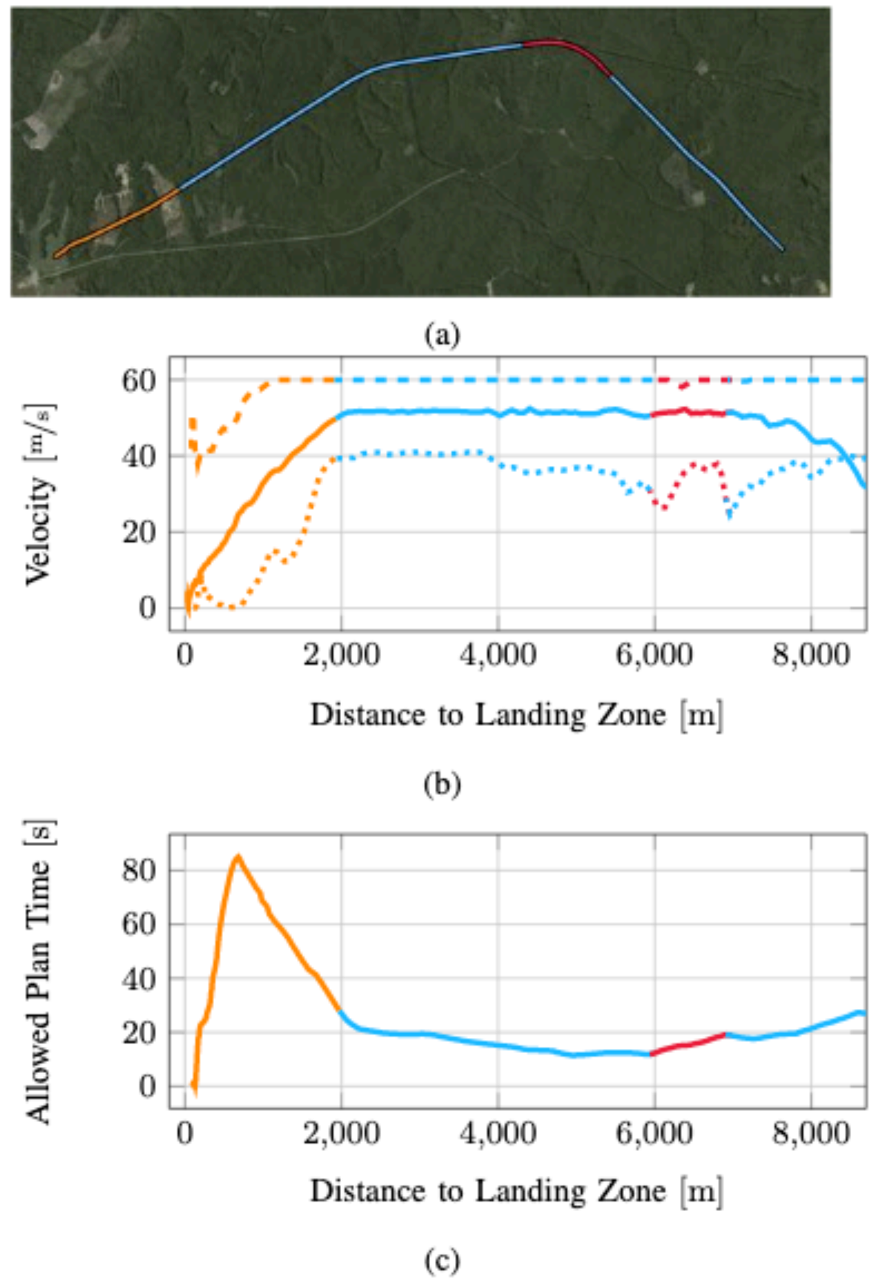

Trajectory Executive: guaranteed safe Emergency Maneuver Library

high performance whilst remaining safe

air lab

The Planner Ensemble and Trajectory Executive: A High Performance Motion Planning System with Guaranteed Safety (2014)
Emergency Maneuver Library – Ensuring Safe Navigation in Partially Known Environments (2015)
High Performance and Safe Flight of Full-Scale Helicopters from Takeoff to Landing with an Ensemble of Planners (2019)

From how much can we get away with

to safety-first design

Ensuring safety is hard

- **Unknown actions by other agents (humans)**
- **Unmodeled disturbances (wind)**

Ensuring safety in the field is even harder

- **Real-time**
- **Sensor and actuator noise**
- **Rapidly-changing partially known surroundings**
- **Unknown obstacle locations**

air lab

Safe Control Techniques

- Hamilton-Jacobi (HJI) Reachability
- Control Barrier Function
- Safety Envelopes

Key Tools:

- Reachability
- Lyapunov functions

Reachability analysis: avoidance

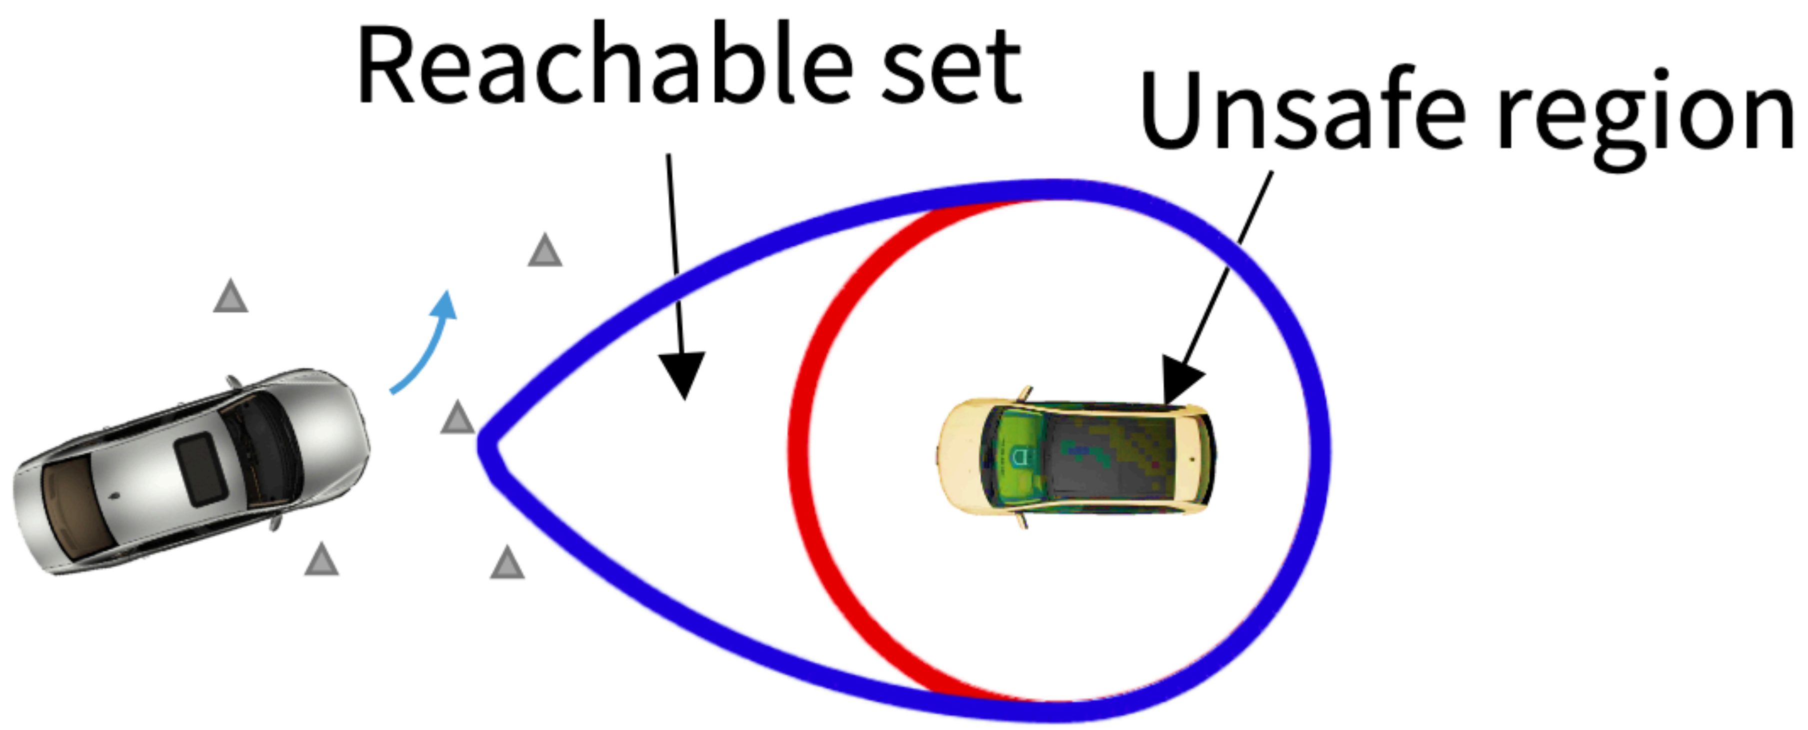

Assumptions:
- Model of robot
- Unsafe region: e.g., obstacle

Control policy

Backward reachable set
(States leading to danger)

Adapted from Stanford AA203 Optimal and Learning-based Control https://stanfordasl.github.io/aa203/pdfs/lecture/lecture_6.pdf
& Somil Bansal's Introduction to Reachability https://people.eecs.berkeley.edu/~somil/Papers/Introduction_to_Reachability_to_Share.pdf

air lab

A optimal control problem cursed by dimensionality.

- Avoiding danger

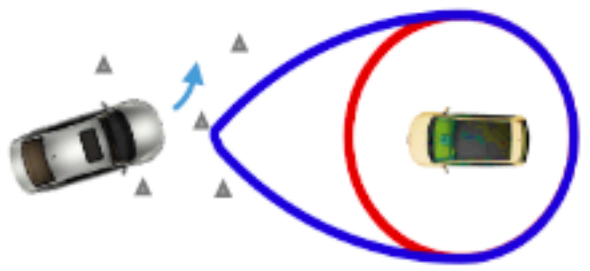

  - BRS definition

$$\mathcal{A}(t) = \{\bar{\mathbf{x}}: \exists\Gamma[\mathbf{u}](\cdot), \forall \mathbf{u}(\cdot), \dot{\mathbf{x}} = f(\mathbf{x}, \mathbf{u}, \mathbf{d}), \mathbf{x}(t) = \bar{\mathbf{x}}, \mathbf{x}(0) \in \mathcal{T}\}$$

  - Value function

$$J(\mathbf{x}, t) = \min_{\Gamma[\mathbf{u}]} \max_{\mathbf{u}} h(\mathbf{x}(0))$$

  - HJI

$$\frac{\partial J}{\partial t} + \max_{\mathbf{u}} \min_{\mathbf{d}} \left[\left(\frac{\partial J}{\partial \mathbf{x}} \right)' f(\mathbf{x}, \mathbf{u}, \mathbf{d}) \right] = 0$$

  - Optimal control

$$\mathbf{u}^* = \arg\max_{\mathbf{u}} \min_{\mathbf{d}} \left(\frac{\partial J}{\partial \mathbf{x}} \right)' f(\mathbf{x}, \mathbf{u}, \mathbf{d})$$

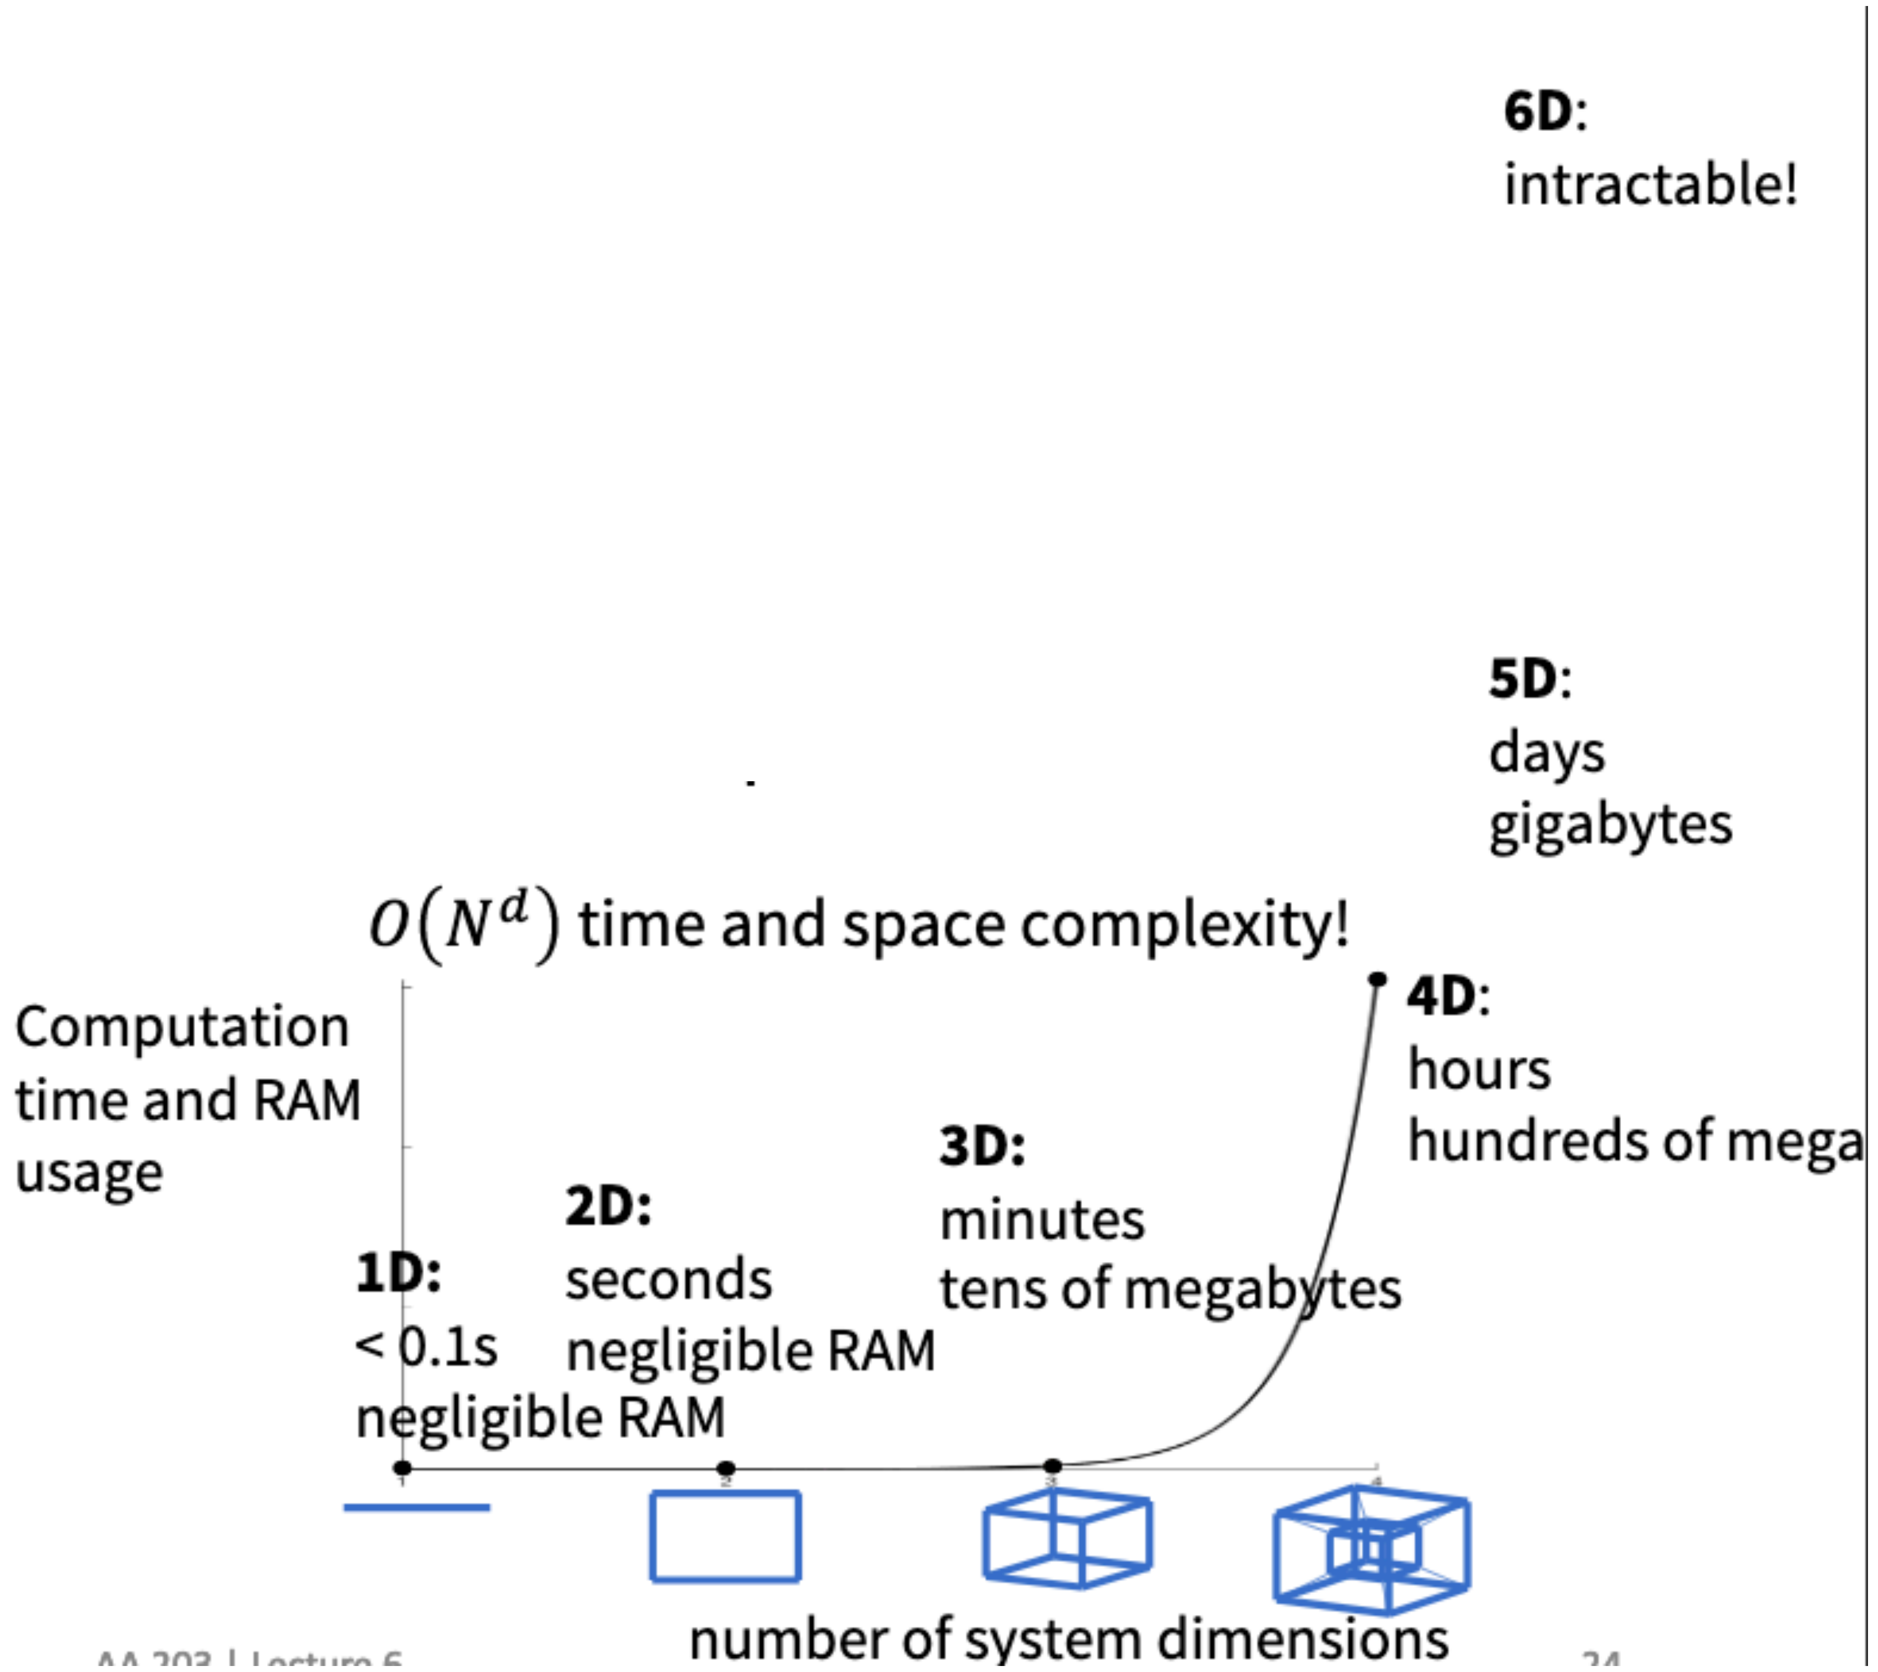

Adapted from Stanford AA203 Optimal and Learning-based Control https://stanfordasl.github.io/aa203/pdfs/lecture/lecture_6.pdf
& Somil Bansal's Introduction to Reachability https://people.eecs.berkeley.edu/~somil/Papers/Introduction_to_Reachability_to_Share.pdf

air lab

Reachability - Applications

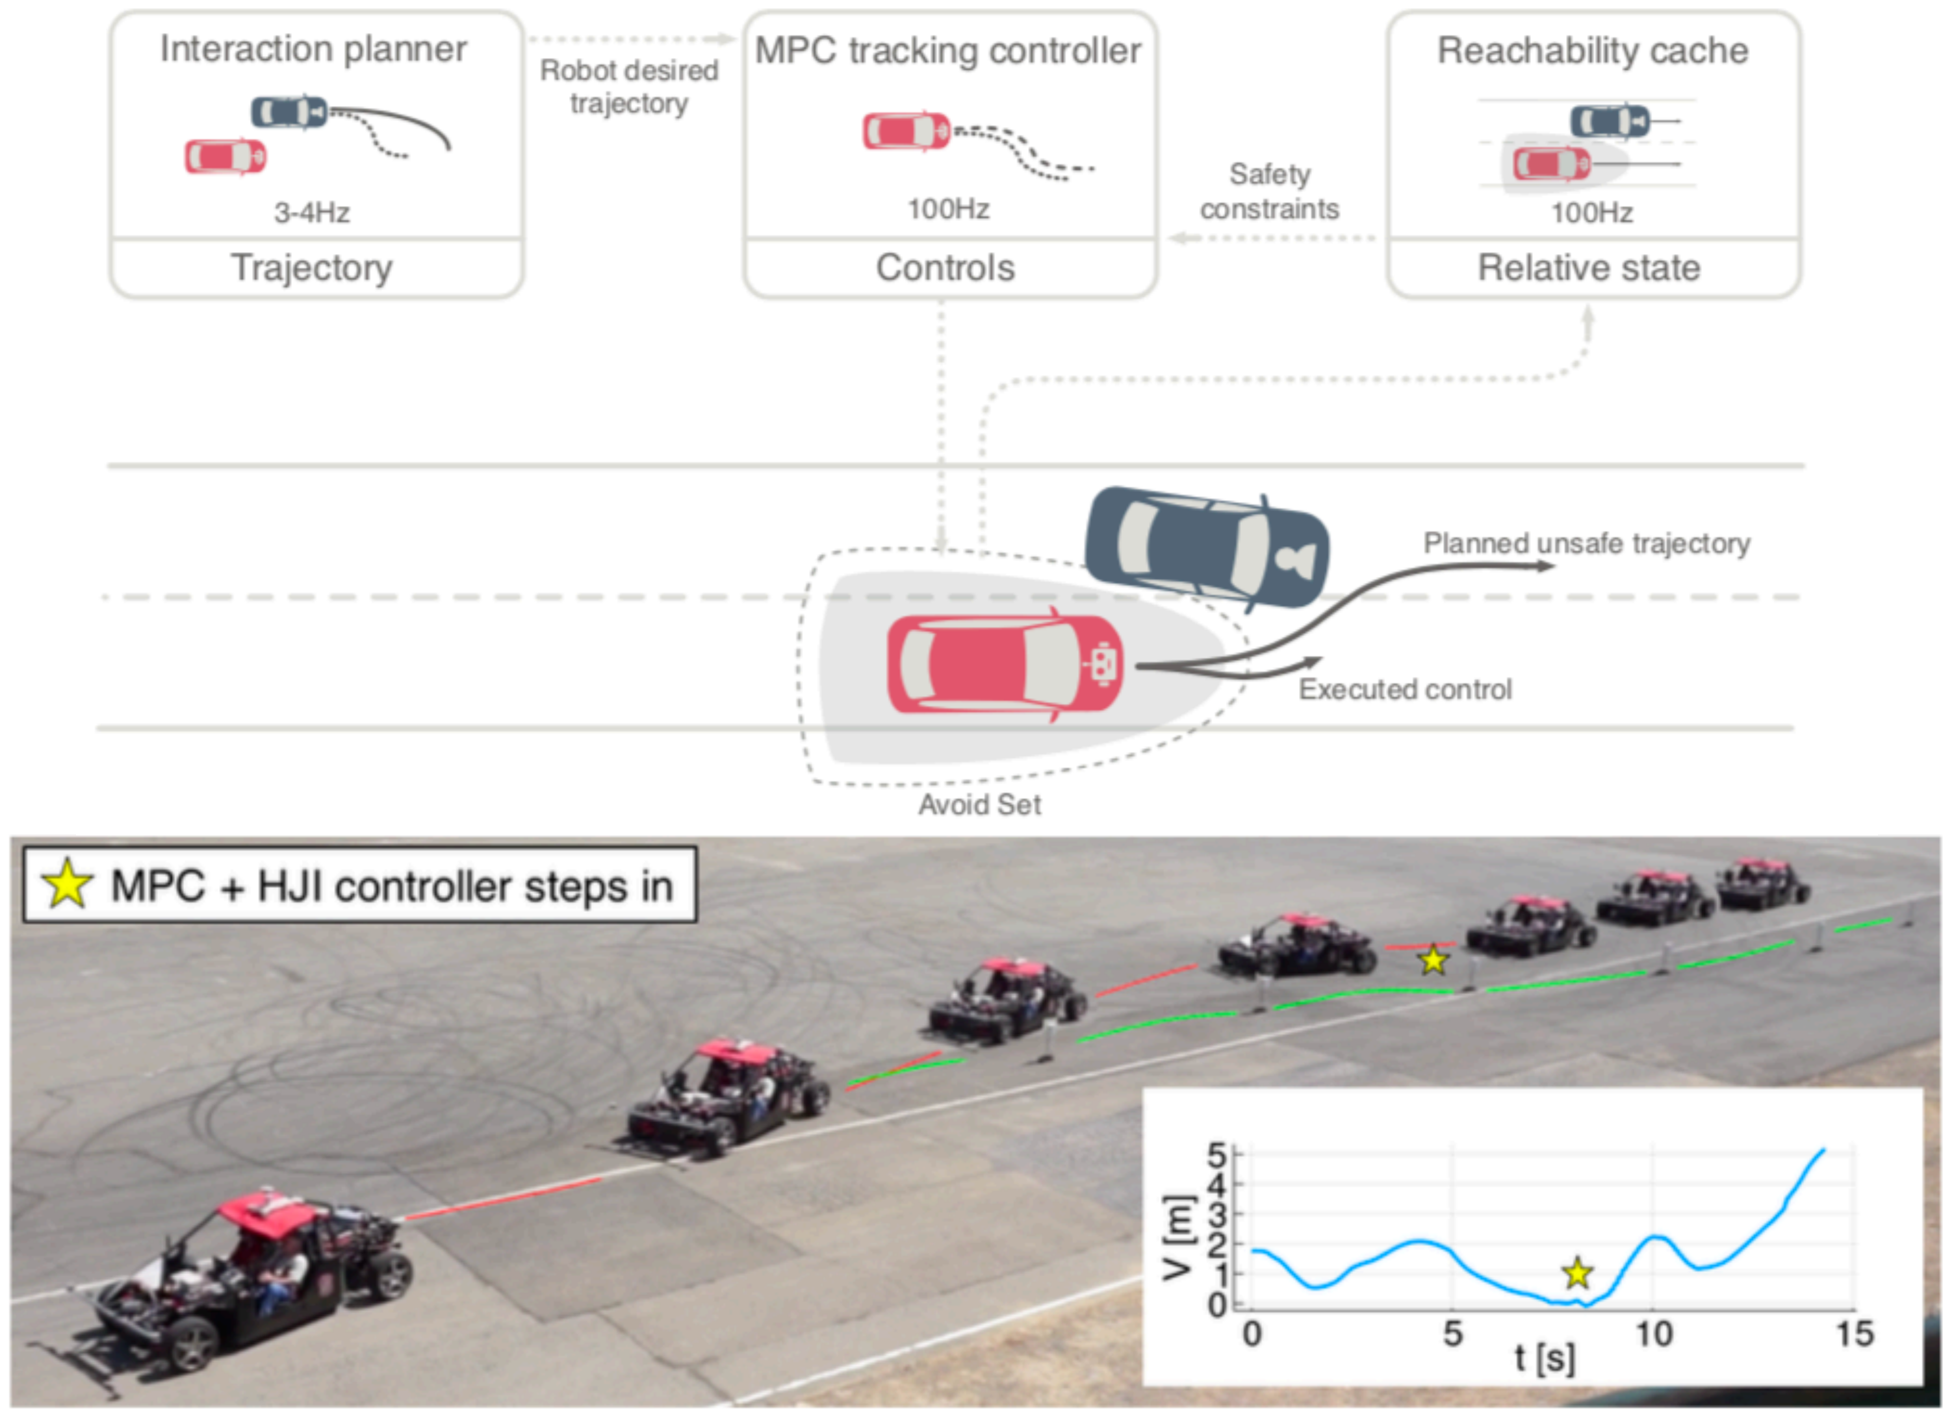

Leung et al. On Infusing Reachability-Based Safety Assurance within Probabilistic Planning Frameworks for Human-Robot Vehicle Interactions. 2018.
https://stanfordasl.github.io/wp-content/papercite-data/pdf/Leung.Schmerling.Chen.ea.ISER18.pdf

- How to guarantee safety with other sentient agent?
- HJI Reachability in MPC as constraint
- **Only tested one maneuver in real life**
- **No sensor, actuator noise. Simulated human agent.**
- Uses saved reachability analysis of pairwise interaction
- Not sure how this can realistically translate (many interactions modeled beforehand)

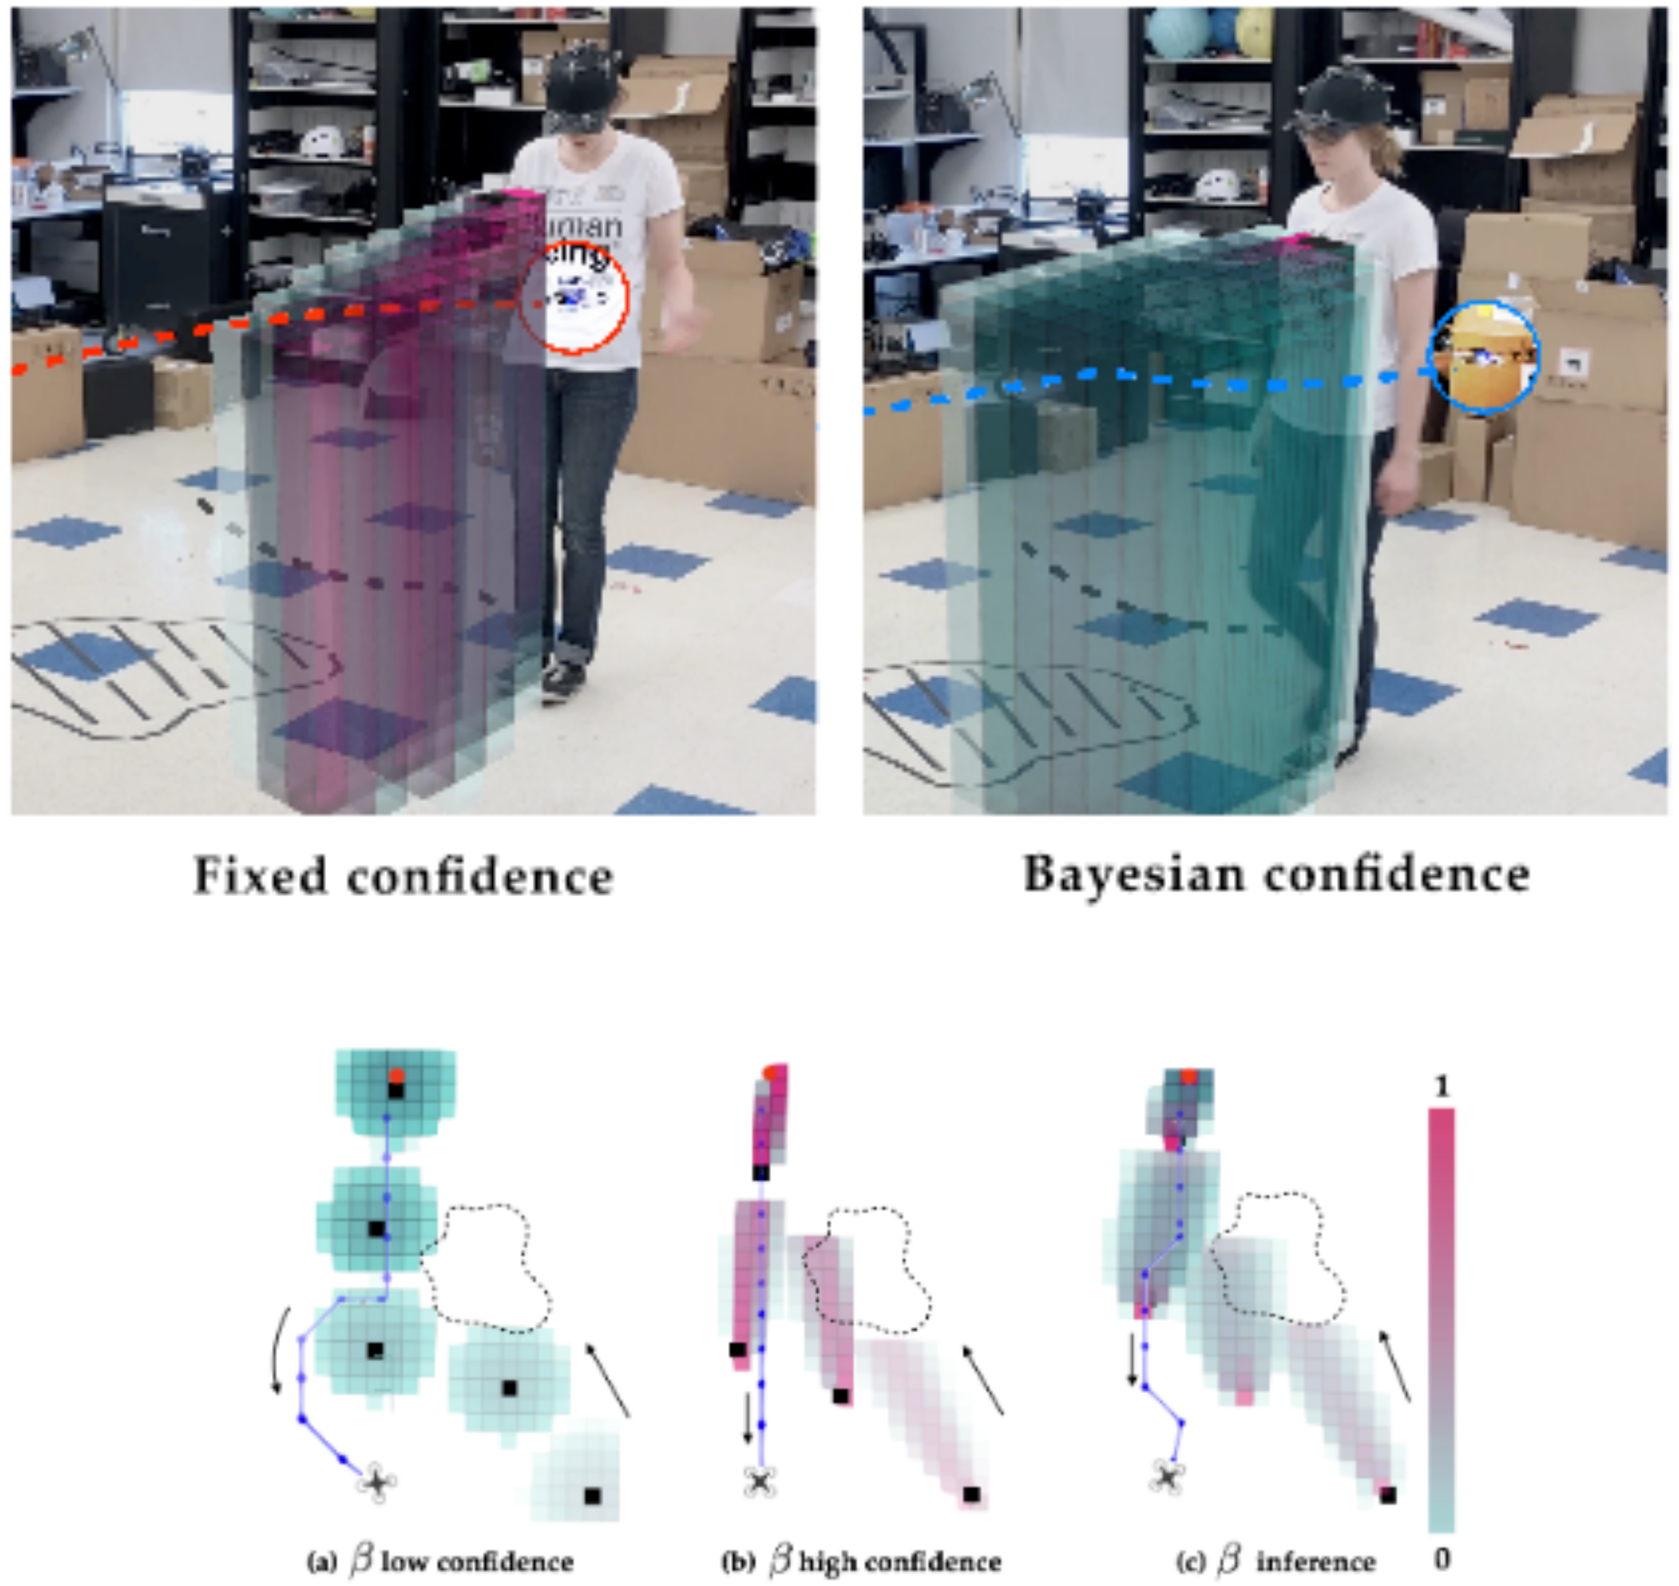

Fisac et al. Probabilistically Safe Robot Planning with Confidence-Based Human Predictions. 2018.
https://arxiv.org/pdf/1806.00109.pdf

- **Vicon for state estimation**
- Estimate human model confidence to adjust planning risk

air lab

Control Lyapunov to certify stability,

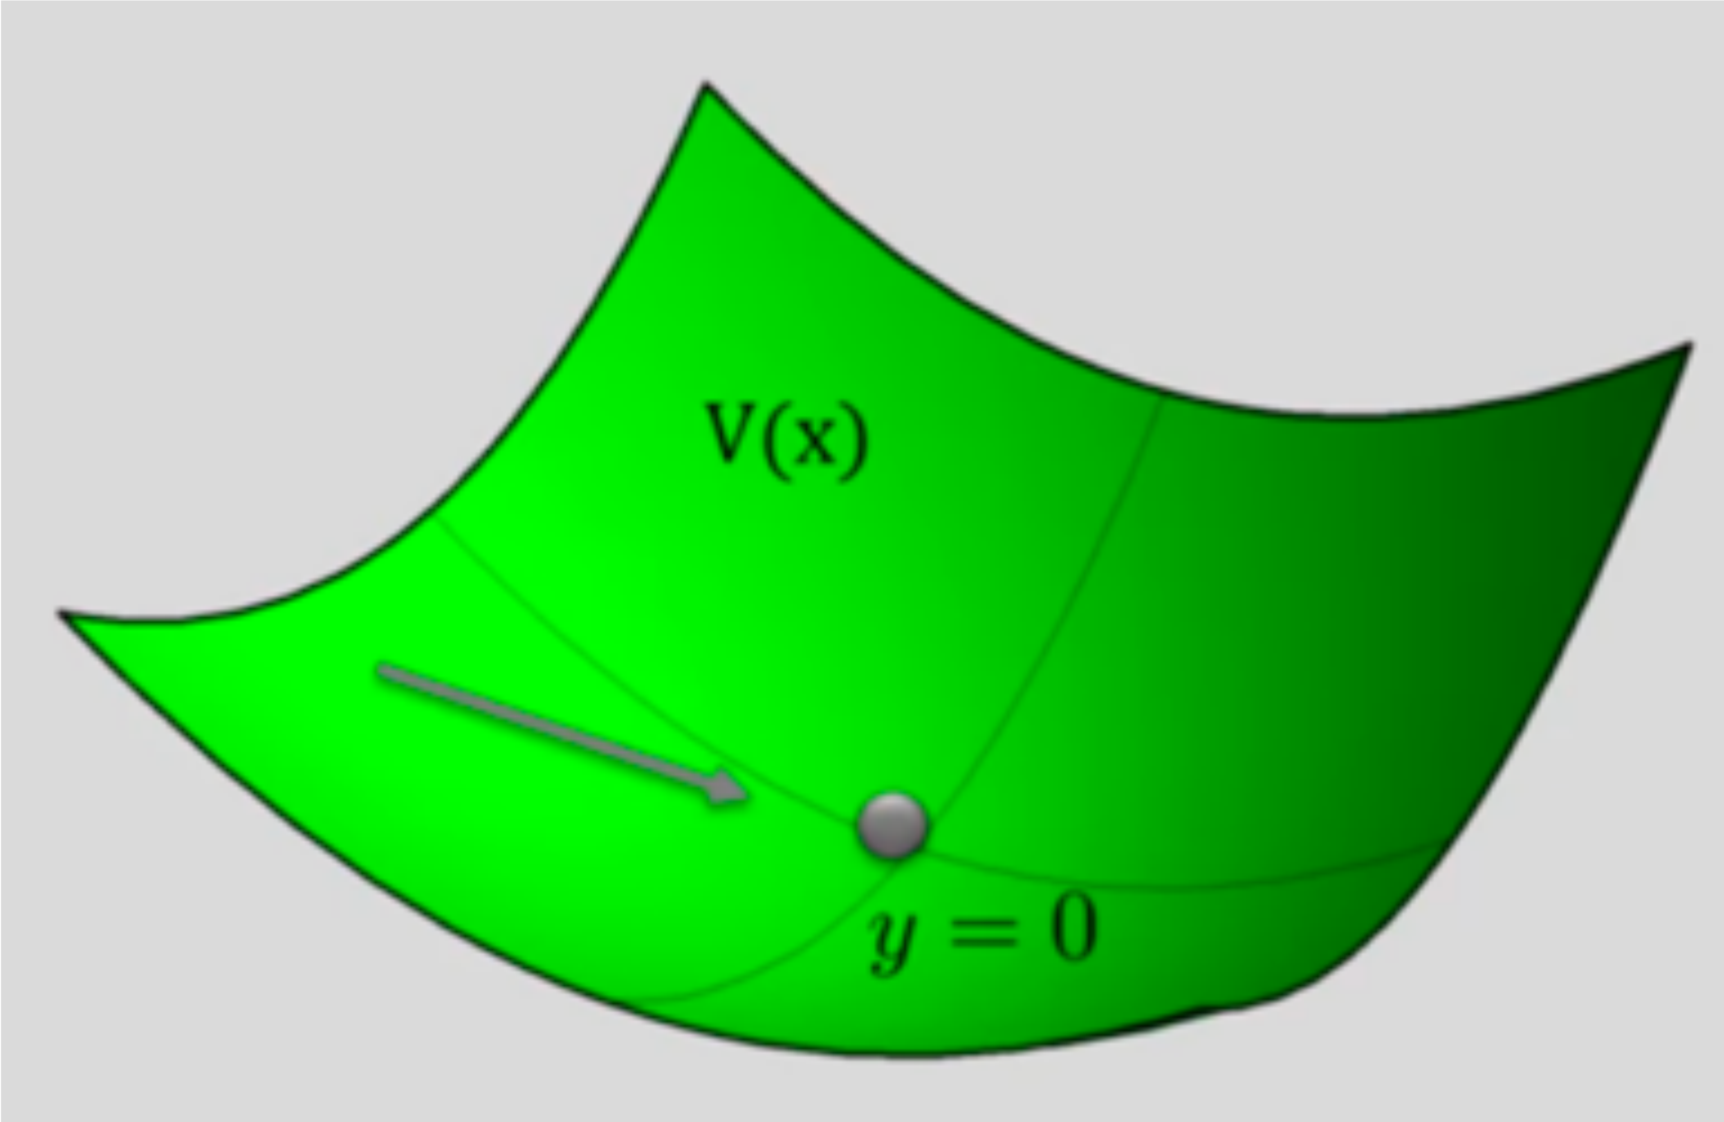

For dynamical system ensure,

$\dot{V}(x, u) \leq -\alpha V(x)$

Control Barrier Function to certify safety

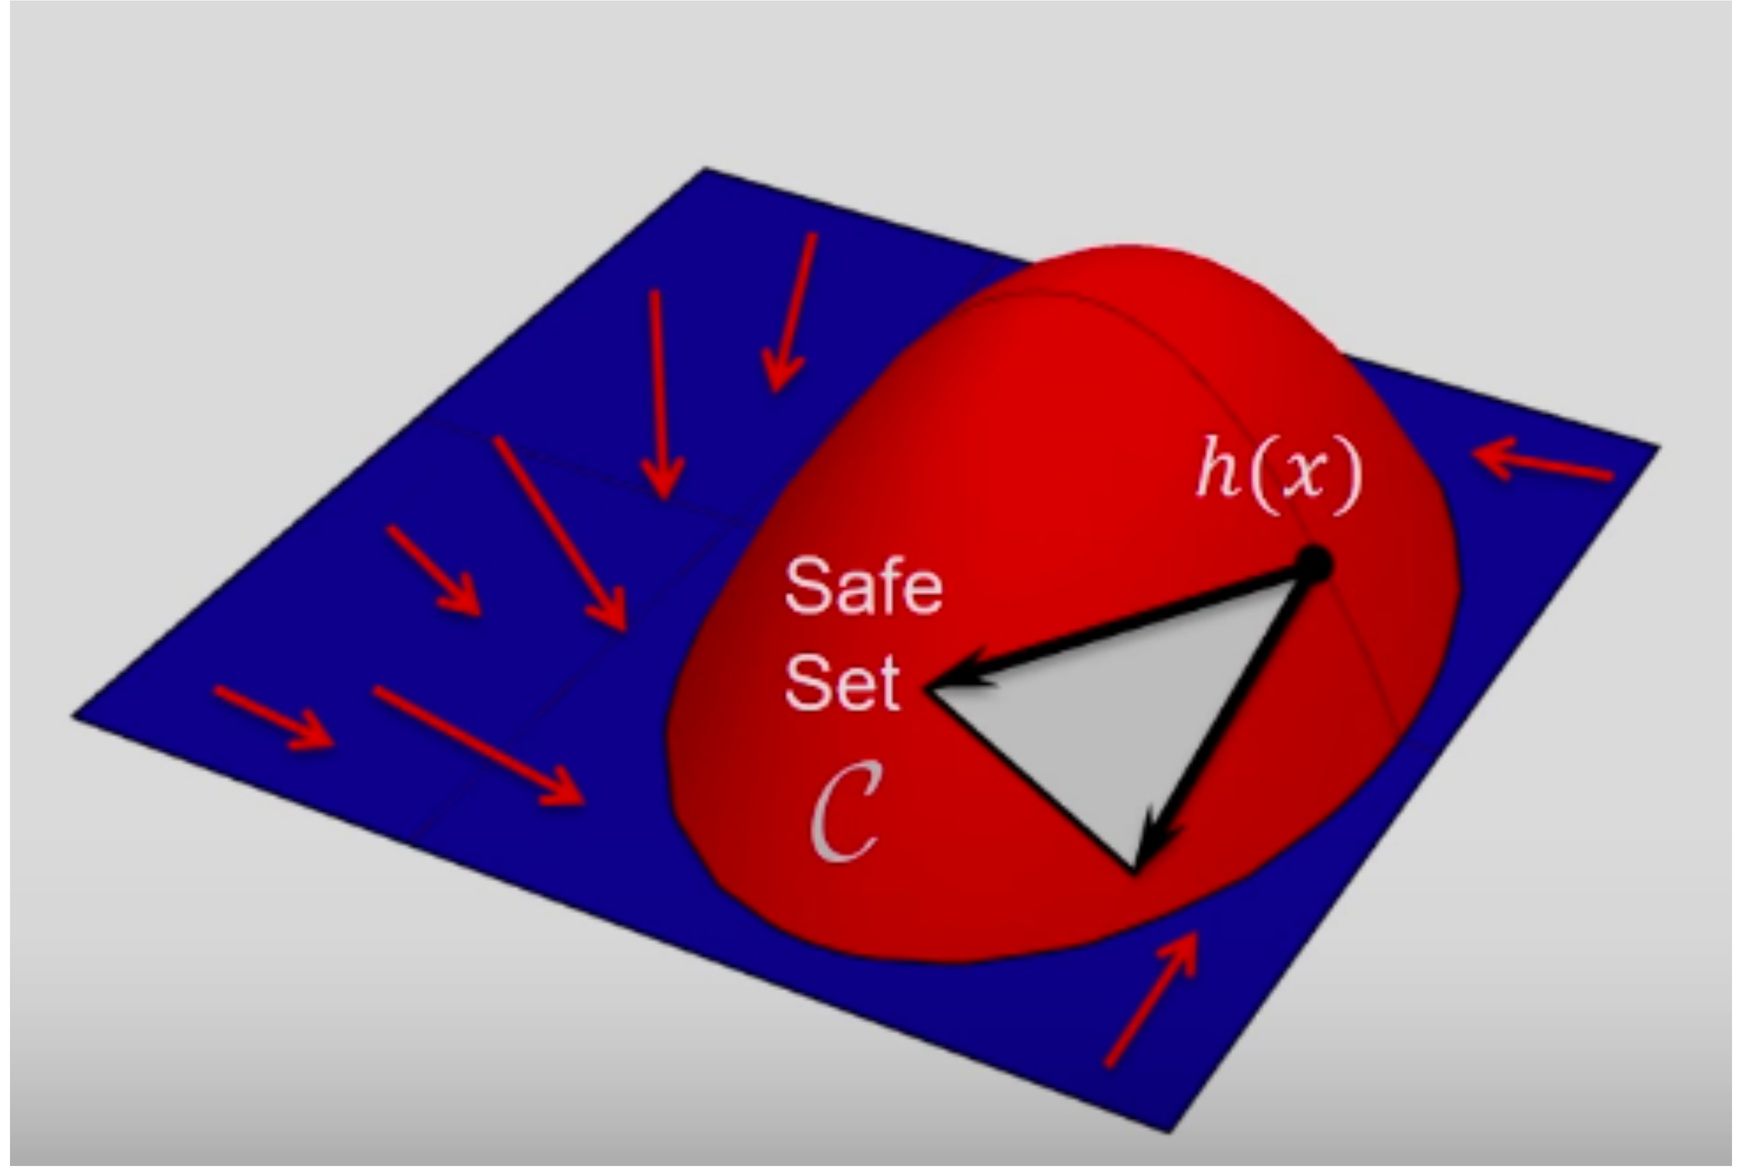

Given a control system and a safe set C

$\exists\ u$ s.t. $\dot{h}(x, u) \geq 0 \Rightarrow C$ is invariant

Compute control action that guarantees vehicle never escape safe set

A cohesive QP framework to achieve control objective while guaranteeing safety

Adapted from Aaron Ames's Slides on Control Barrier Functions

Control Barrier Functions - Applications

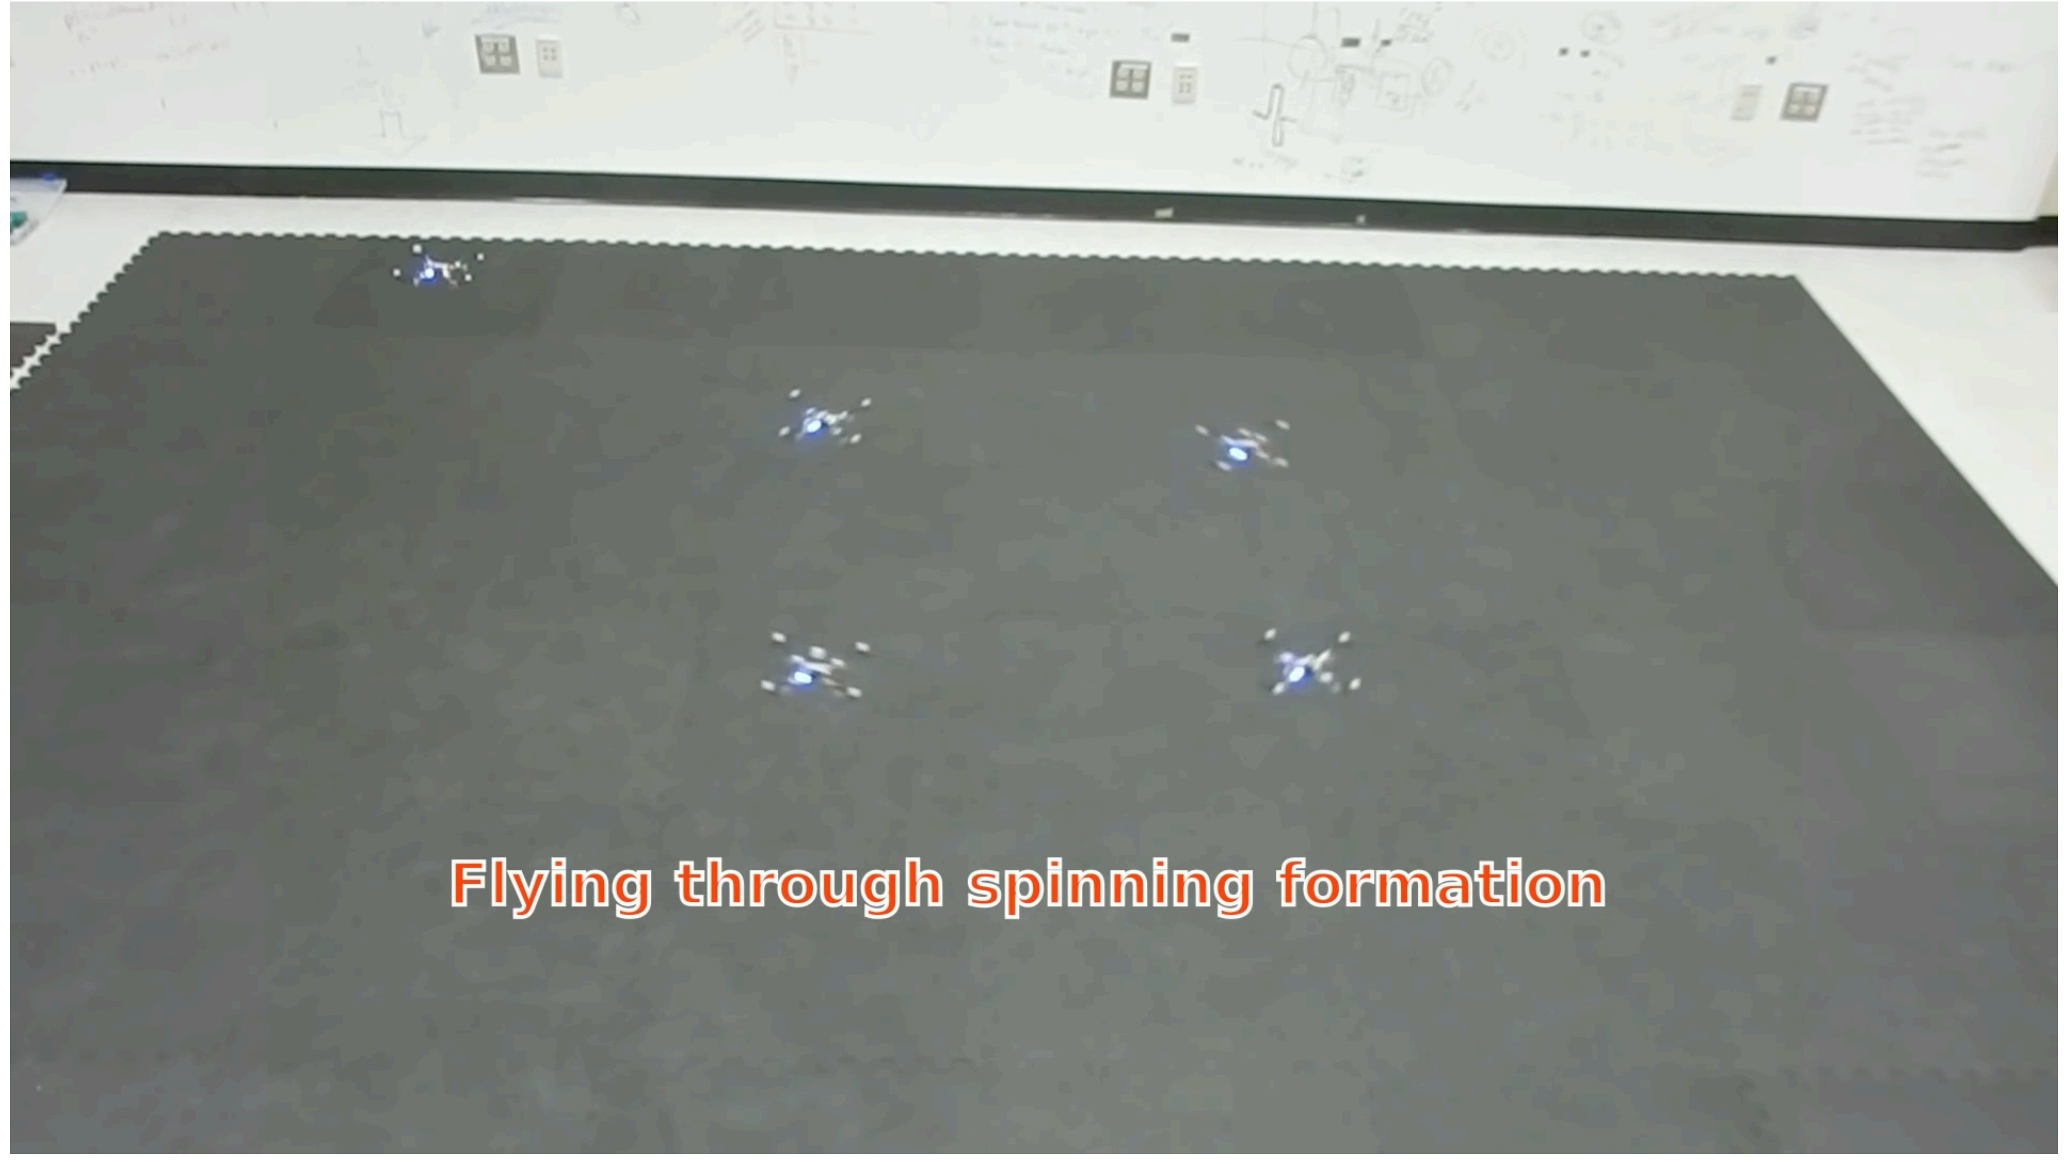

https://www.youtube.com/watch?v=rK9oyqccMJw
Wang et al. Safe Certificate-Based Maneuvers for Teams of Quadrotors Using Differential Flatness. 2017.

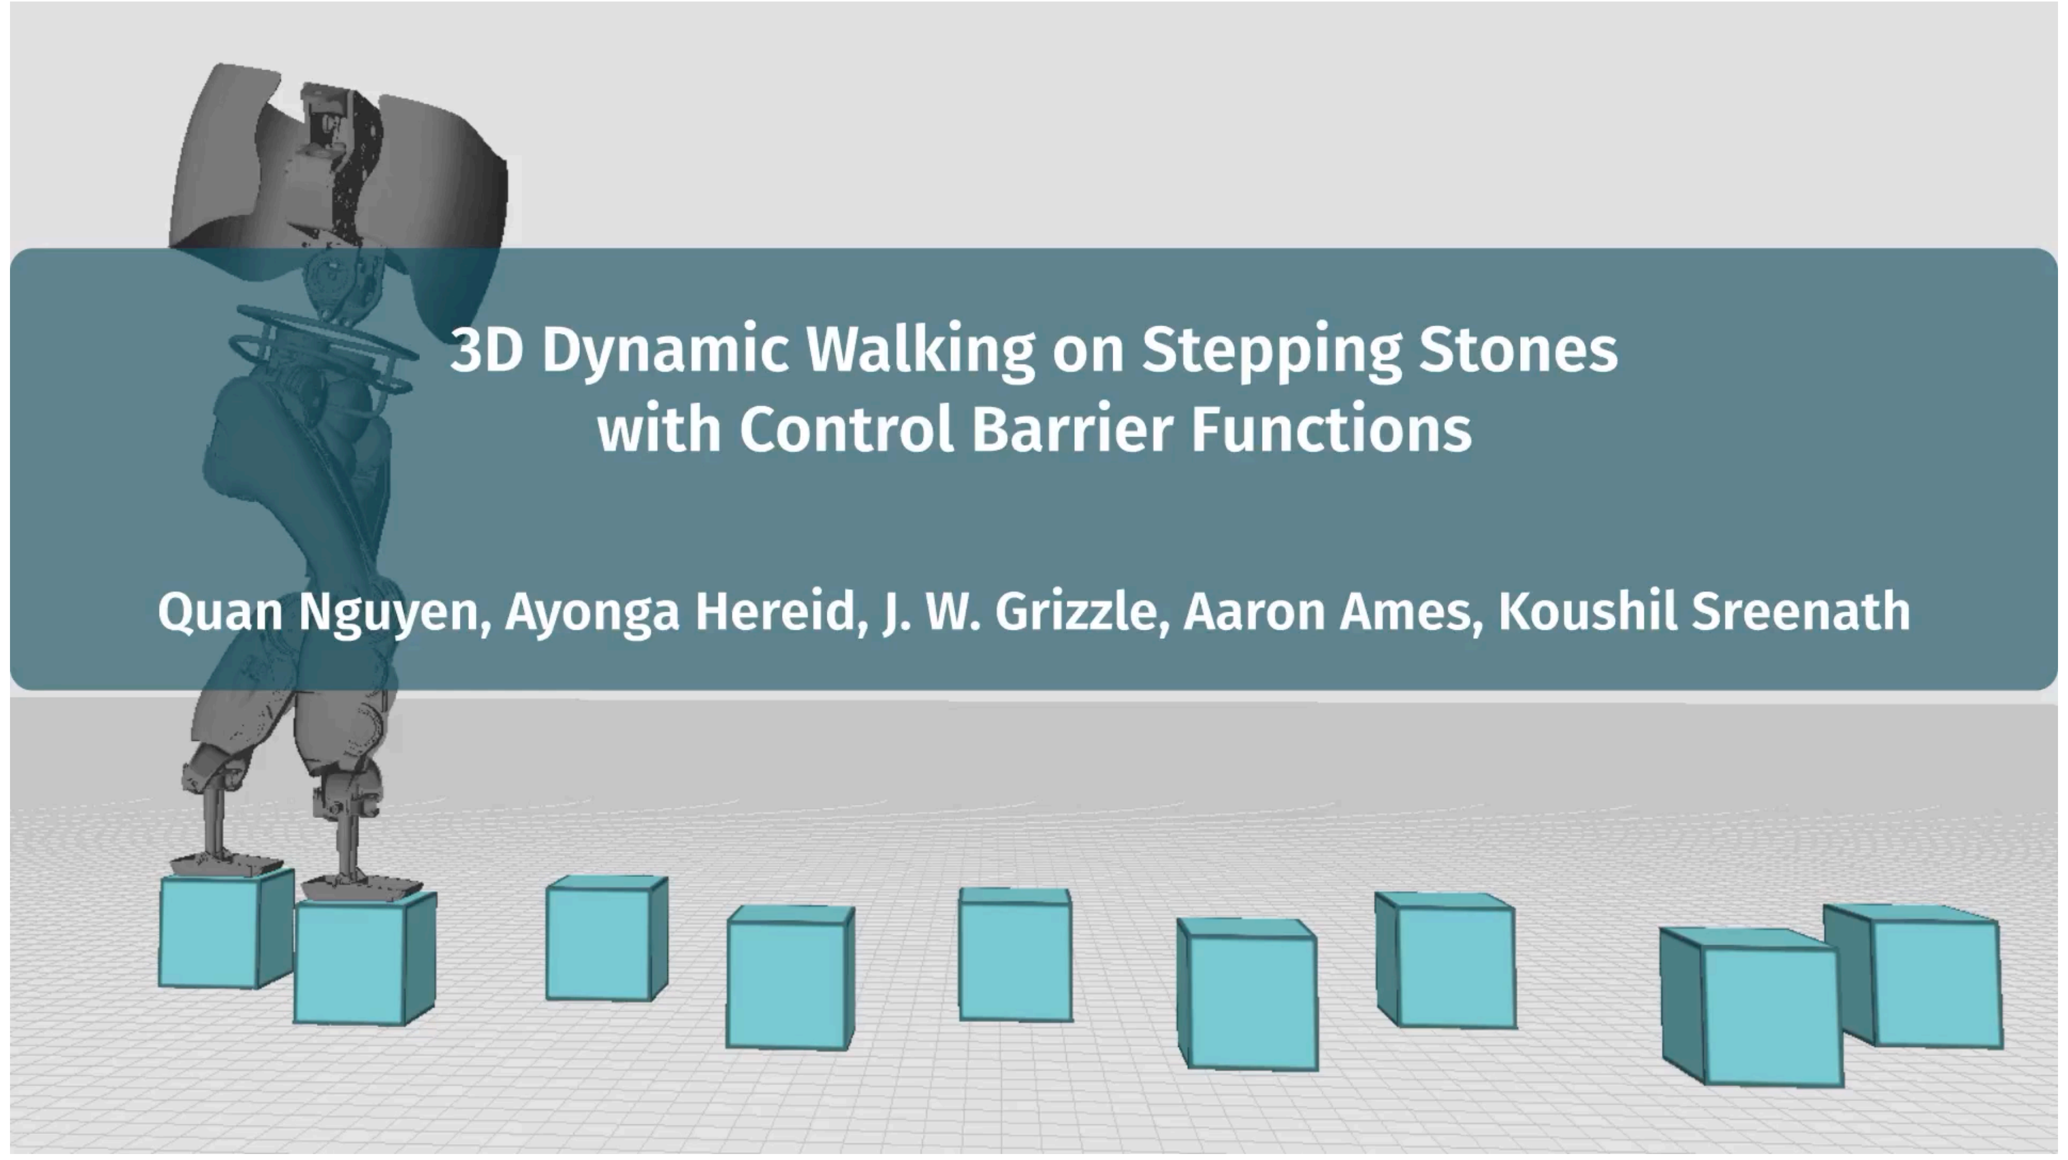

Nguyen et al. 3D Dynamic Walking on Stepping Stones with Control Barrier Functions.

- Modify nominal trajectories to avoid collisions
- Models safety region with super ellipsoids
- **During actual flights, quad rotors often crash due to strong wind**
- **Unclear on real time capability**

air lab

How to bring formal safety guarantees to the real world?

- Real-time
- Account for sensor, state estimation uncertainties
- Theoretical relaxations of safety guarantees

Looking Forward - Safe and Efficient Learning

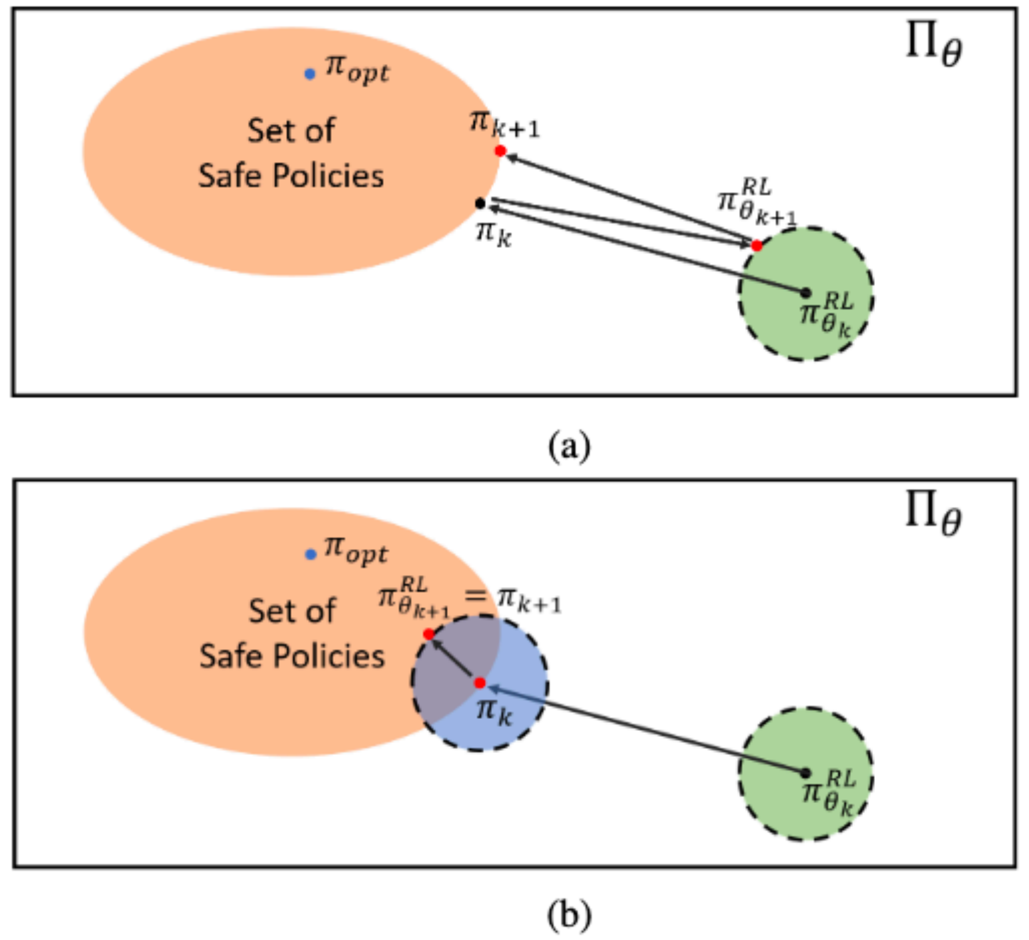

Figure 2: Illustration of policy iteration process, where we try to learn the optimal safe policy, π_{opt}. (a) Policy optimization with barrier-compensating controller. Next policy is updated around the previous RL controller, $\pi^{RL}_{\theta_k}$; (b) Policy optimization with barrier-guided controller. Next policy is updated around previous deployed controller, π_k.

https://arxiv.org/abs/1903.08792
Cheng et al. End-to-End Safe Reinforcement Learning through Barrier Functions for Safety-Critical Continuous Control Tasks. 2019.

Safe RL w/ Barrier Functions to guide learning process

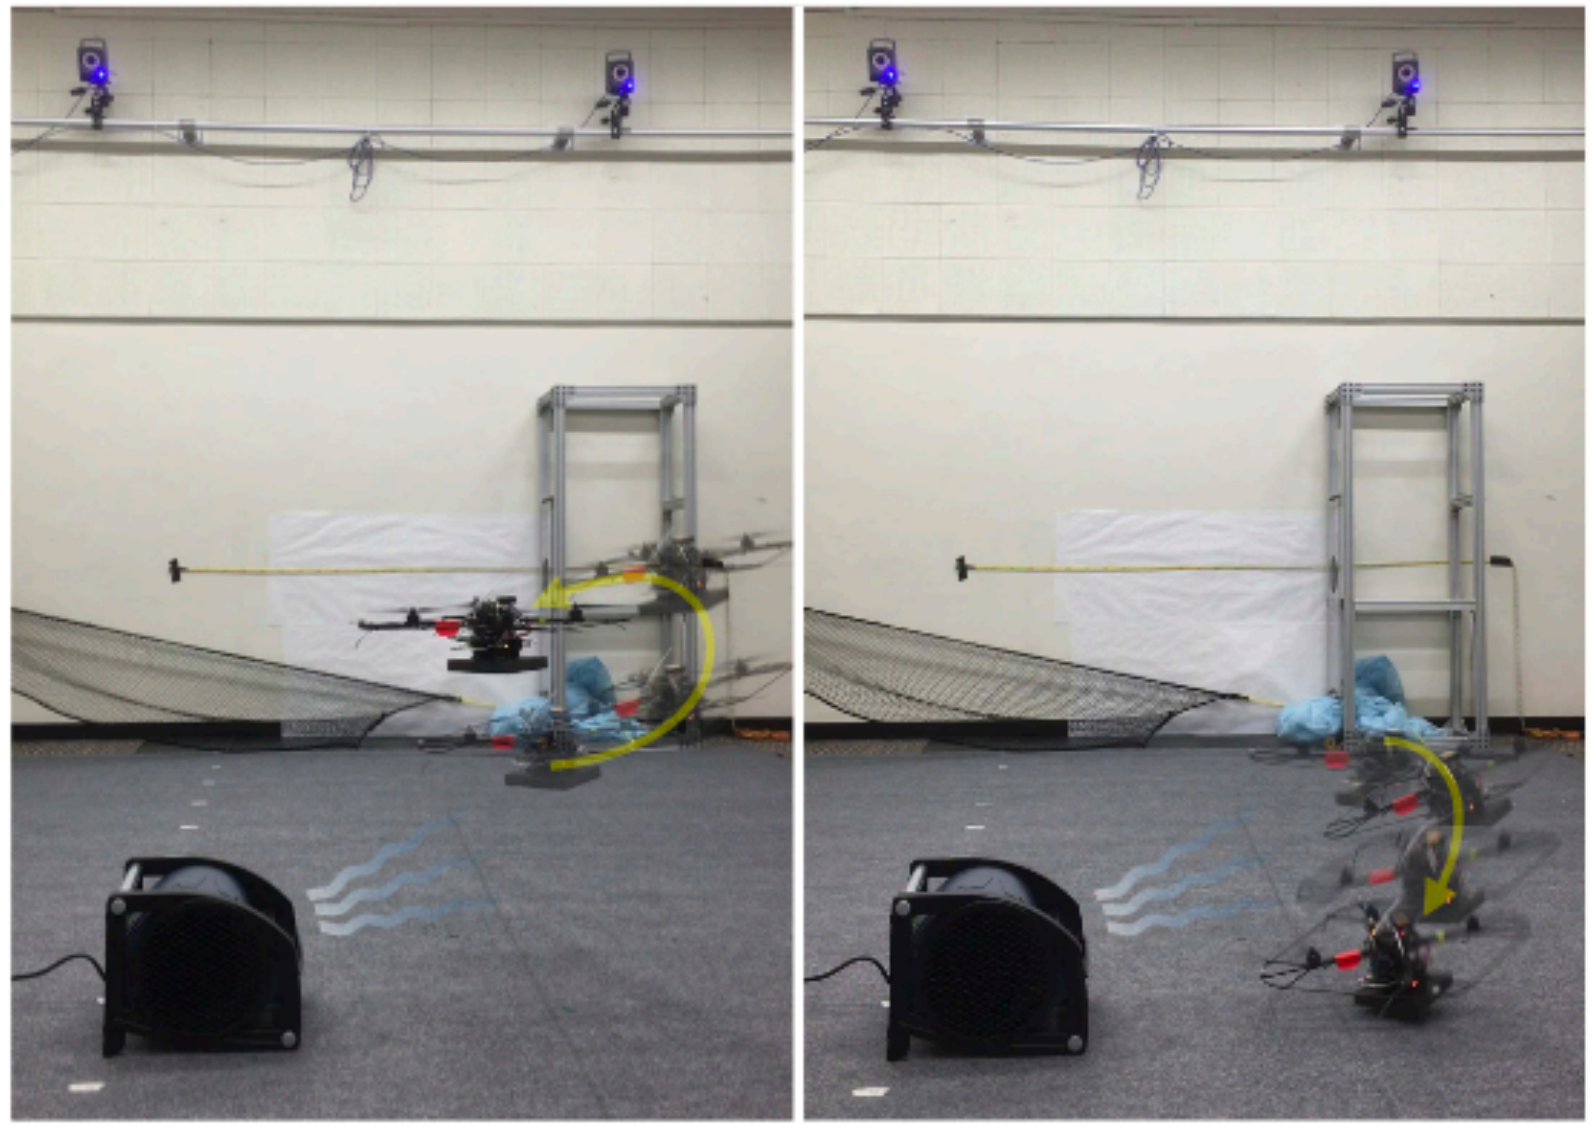

(a) With online guarantee validation (b) Without online guarantee validation

https://arxiv.org/abs/1705.01292
Fisac et al. https://arxiv.org/abs/1705.01292. 2018.

HJ reachability w/ approximate knowledge
*** allows safe learning in the real world**

air lab

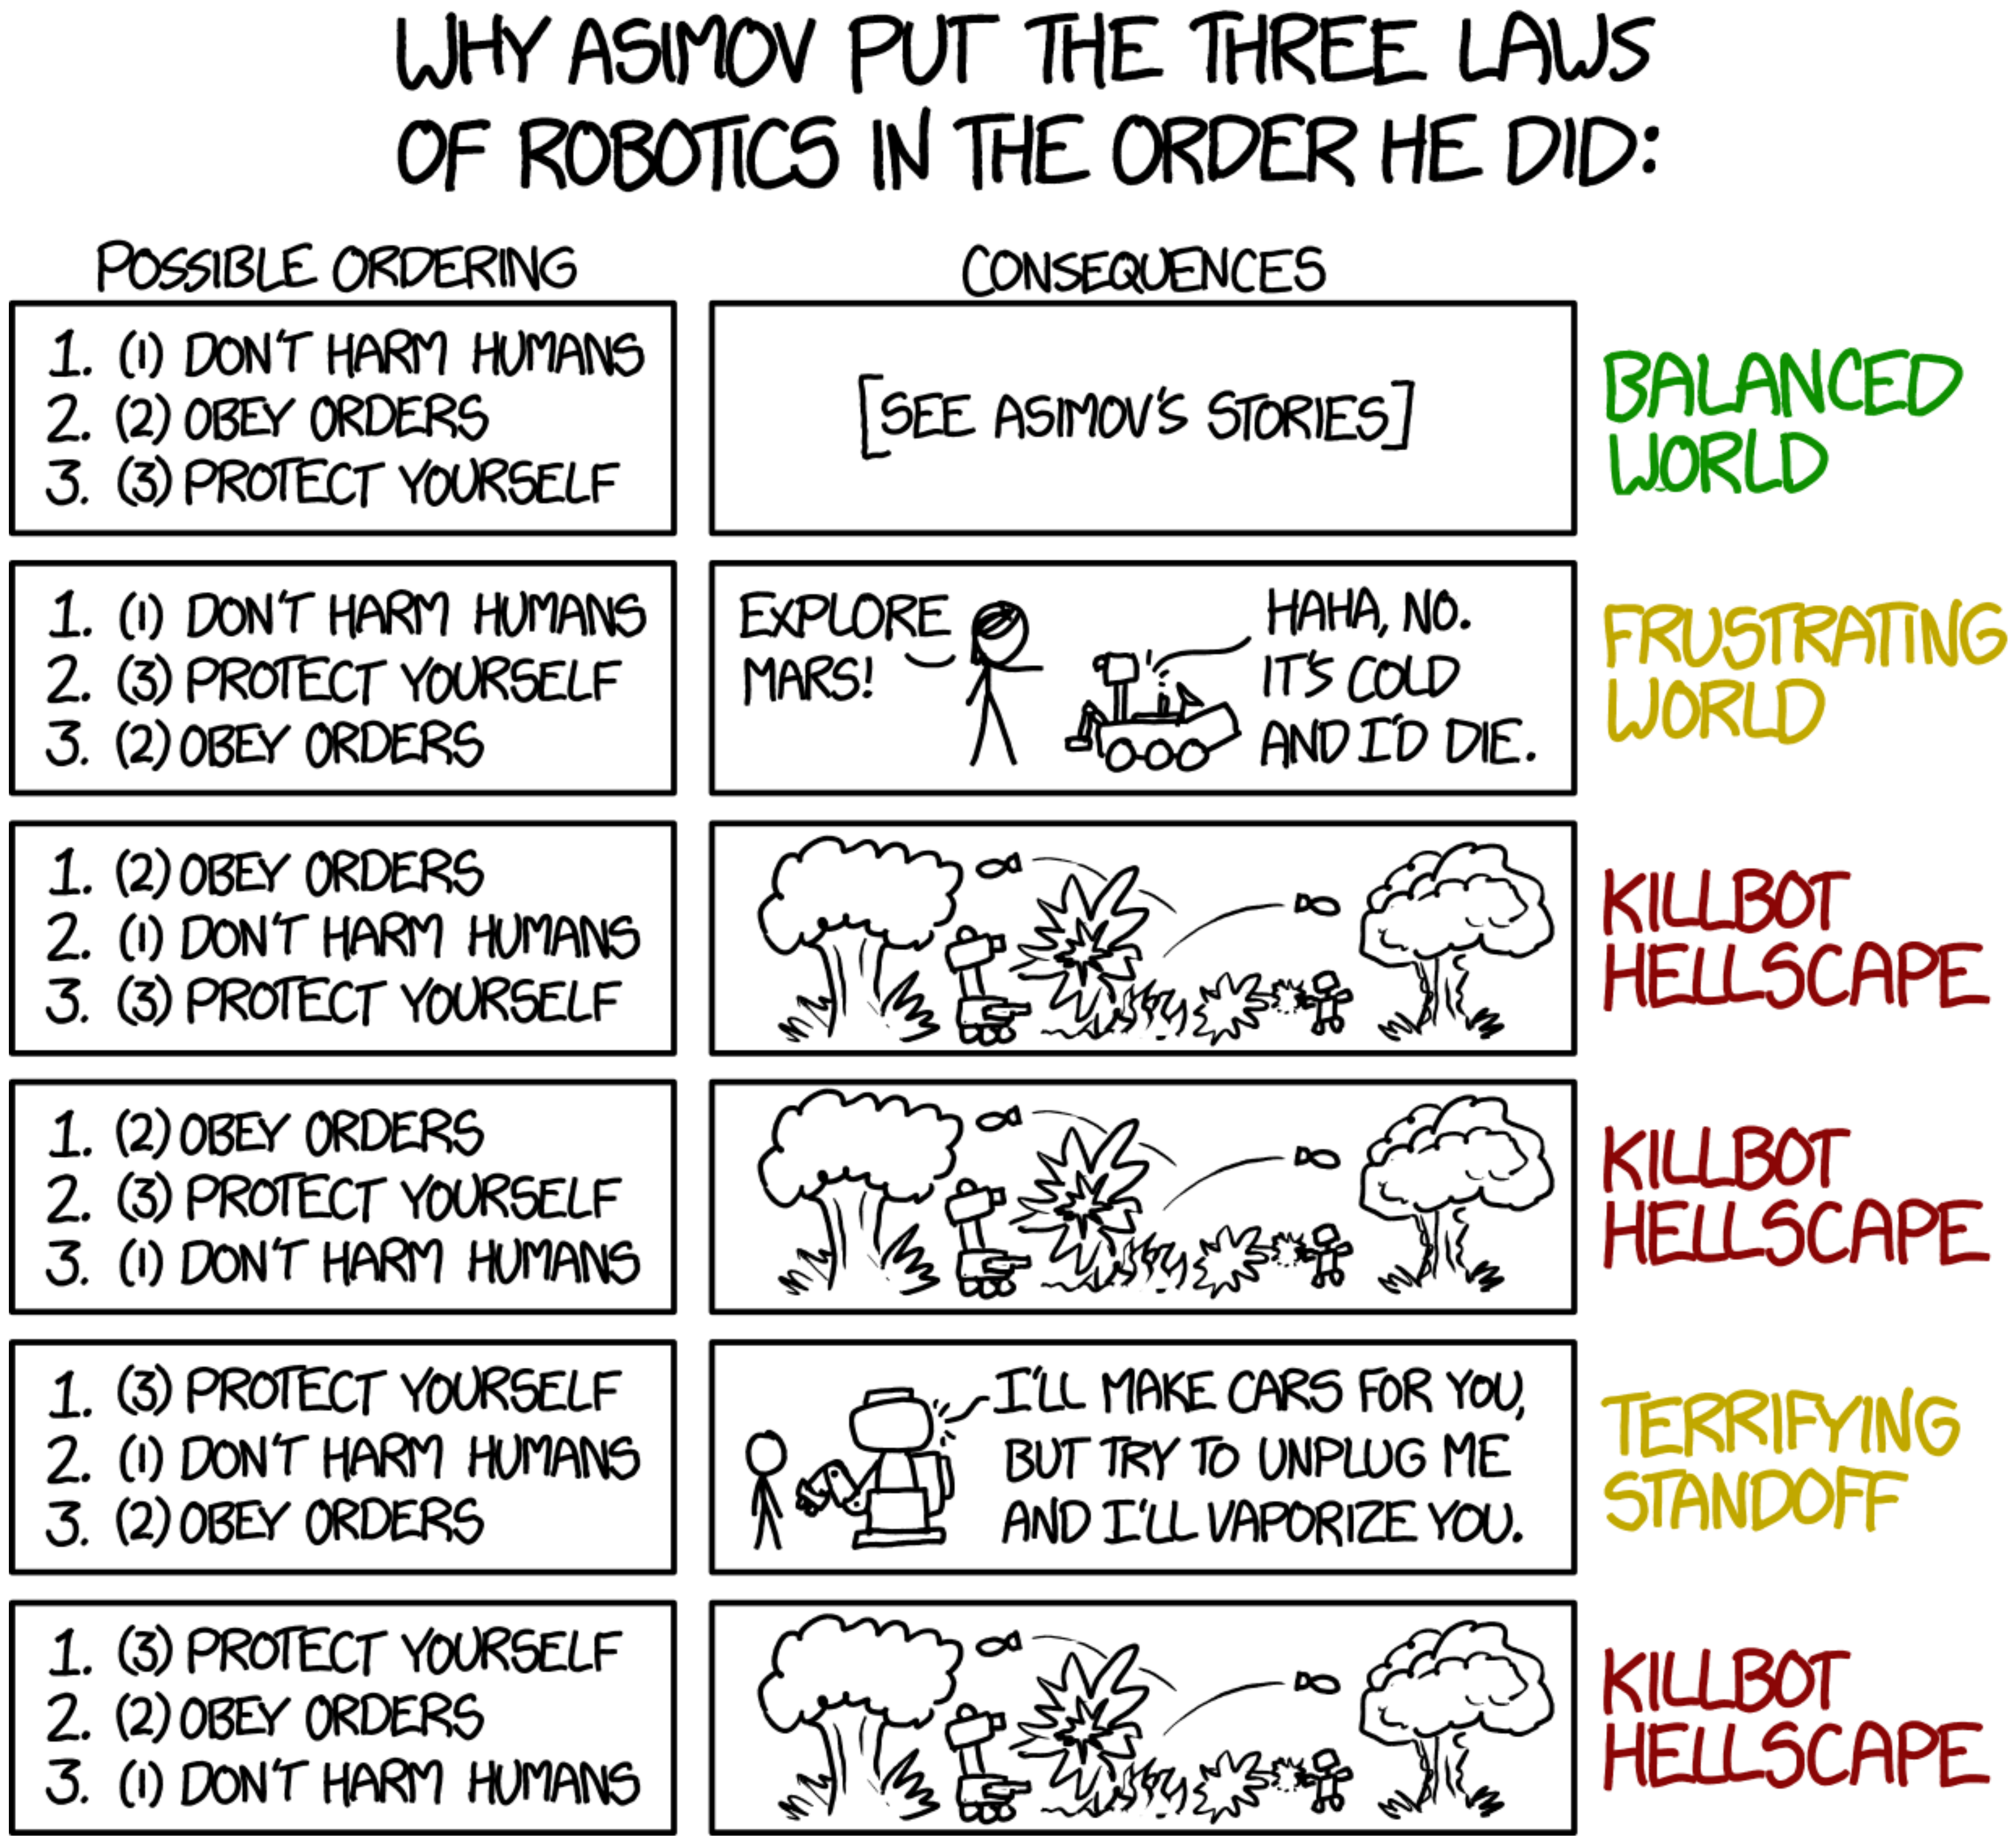

xkcd: The Three Laws of Robotics

Thanks!