Provably Correct Learning

Anayo K. Akametalu

Advisor: Dr. Claire J. Tomlin

Hybrid Systems Lab, UC Berkeley

Biped Walking

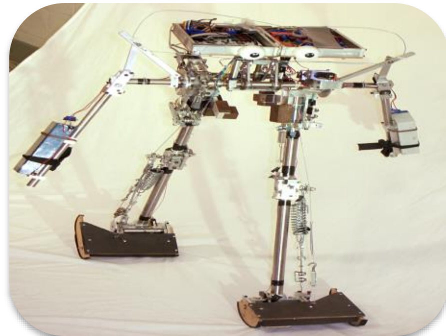

Adaptive Control

Autonomous Driving

Learning in Robotic Systems

Helicopter Flight

Domestic Robots

[Images: BMW, Cornell, MIT UC Berkeley, Stanford, and Tohuku Univers

What about Safety (Correctness)?

Learning

for Safety-Critical Systems

Can we learn to fly from *scratch*?

Asymptotically, Yes!

…but must avoid crashing first

We need a framework for safe learning

Model Assumptions

with quadrotor example

Deterministic dynamical system with bounded additive uncertainty

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

States $x \in \mathcal{X}$
Vertical position and velocity

Control $u \in \mathcal{U}$
Thrust

Disturbance $d(x) \in \mathcal{D}$
Uncertainty in payload
Ground effect
Nonlinearities

Hamilton-Jacobi-Isaacs (HJI) [Mitchell 2005]

Safety as a Dynamical Game

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

$$u \in \mathcal{U}, d \in \mathcal{D}$$

Avoid $\mathcal{K}$:
control u vs. disturbance d

$l(x)$: negative in keep-out and positive otherwise

Propagate dynamics backwards in time:

$$\frac{\partial J(x,t)}{\partial t} = -\min\left\{0, \boxed{\max_{u\in\mathcal{U}}} \min_{d\in\mathcal{D}} \frac{\partial J(x,t)}{\partial x}^T f(x,u,d)\right\}$$

$$J(x,0) = l(x)$$

Unsafe Set: $Pre(\mathcal{K}) = \{x : J(x) \leq 0\}$

Set of initial states for which disturbance can win!

Least Restrictive Control Law: $u \in \begin{cases} \mathcal{U}, & \text{if } J(x) > 0 \\ u^*(x), & \text{otherwise} \end{cases}$

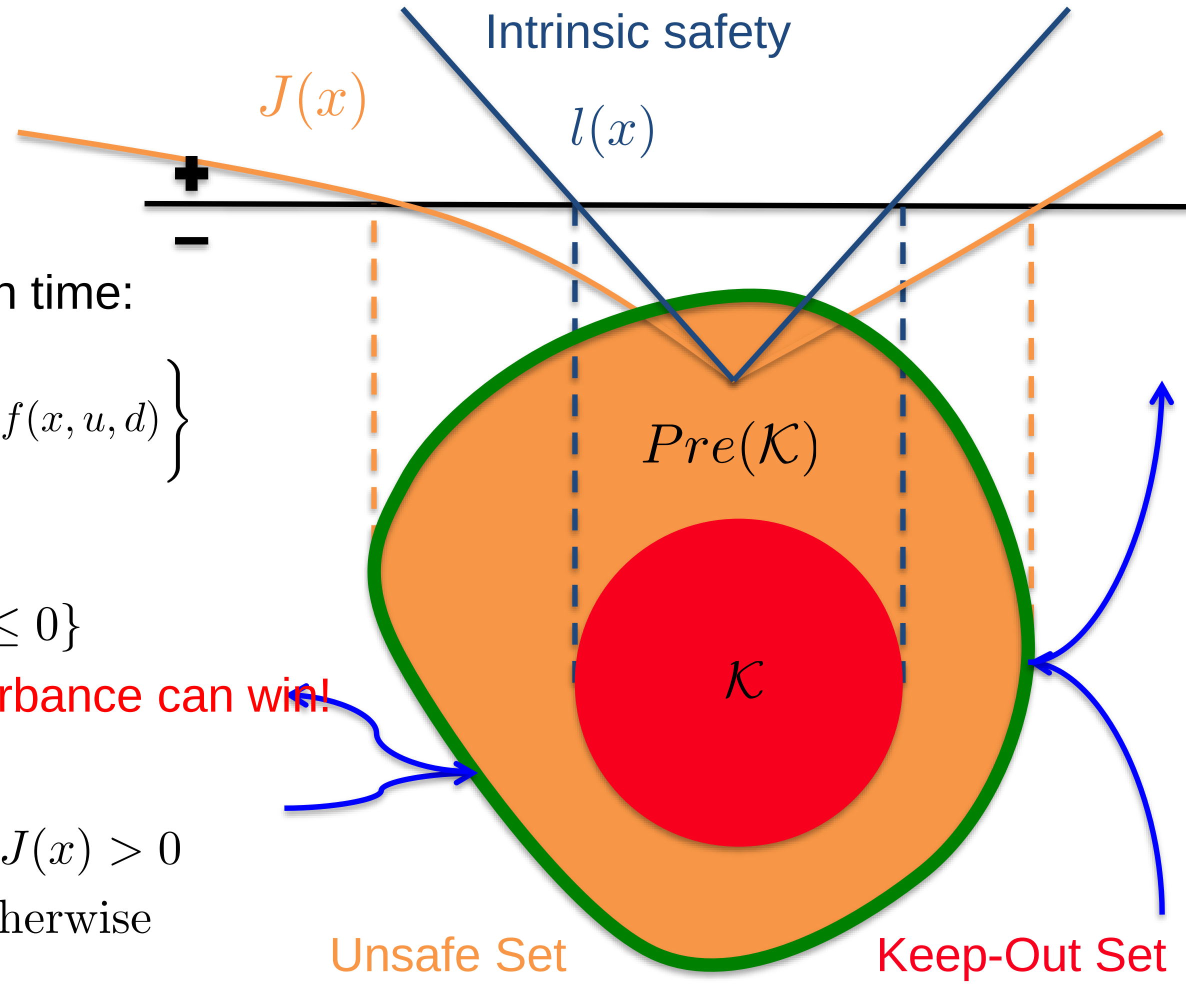

Safe Learning

with Initial Zero-Controller

First thing the quadrotor does is drop

Can we learn to fly from *scratch*?

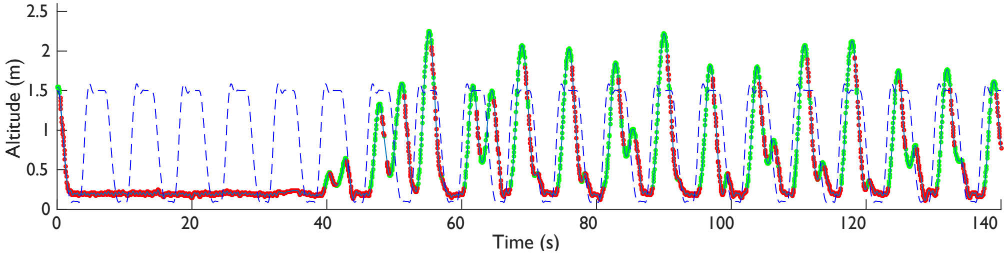

After about 1 minute,
it can roughly track the trajectory

Soon, it starts experimenting

…but the safe controller steps in

Akametalu, Fisac, Zeilinger, Kaynama, Gillula, and Tomlin, CDC 2014

Guaranteed Safety

…but too restrictive

What if we knew the dynamics exactly?

Learn disturbance function then recompute value function

[Gillula and Tomlin, 2012]

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

$$\frac{\partial J(x,t)}{\partial t} = -\min\left\{0, \max_{u \in \mathcal{U}} \min_{d \in \mathcal{D}} \frac{\partial J(x,t)}{\partial x}^T f(x,u,d)\right\}$$

Unsafe Set: $Pre(\mathcal{K}) = \{x : J(x) \leq 0\}$

Least Restrictive Control Law: $u \in \begin{cases} \mathcal{U}, & \text{if } J(x) > 0 \\ u^*(x), & \text{otherwise} \end{cases}$

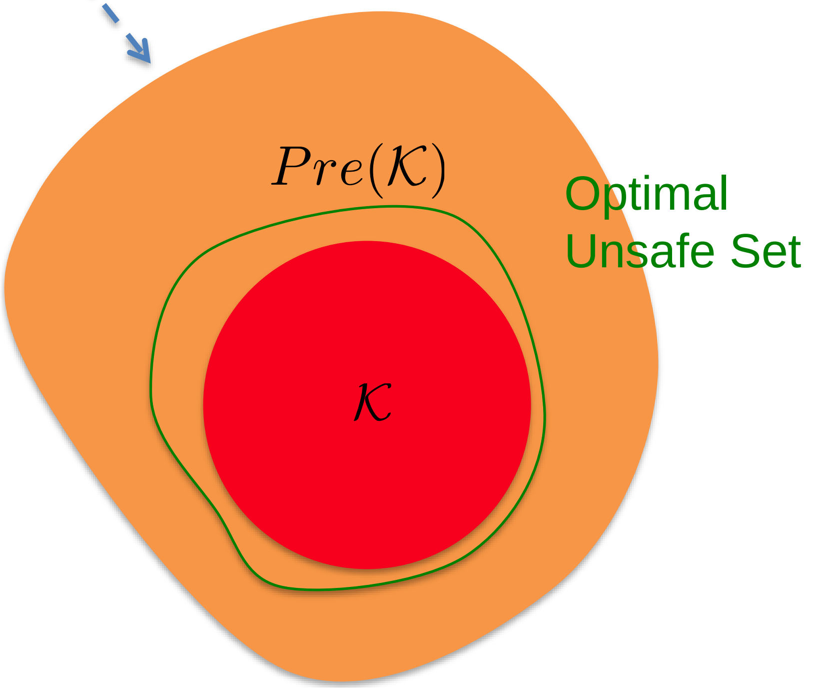

What if we learned a poor model?

Model Validation

Validate model online and adjust controller

Infinite number of candidate unsafe sets for $\mathcal{D}(x)$

$u^*(x)$

Initialize:
Active unsafe set = smallest candidate set

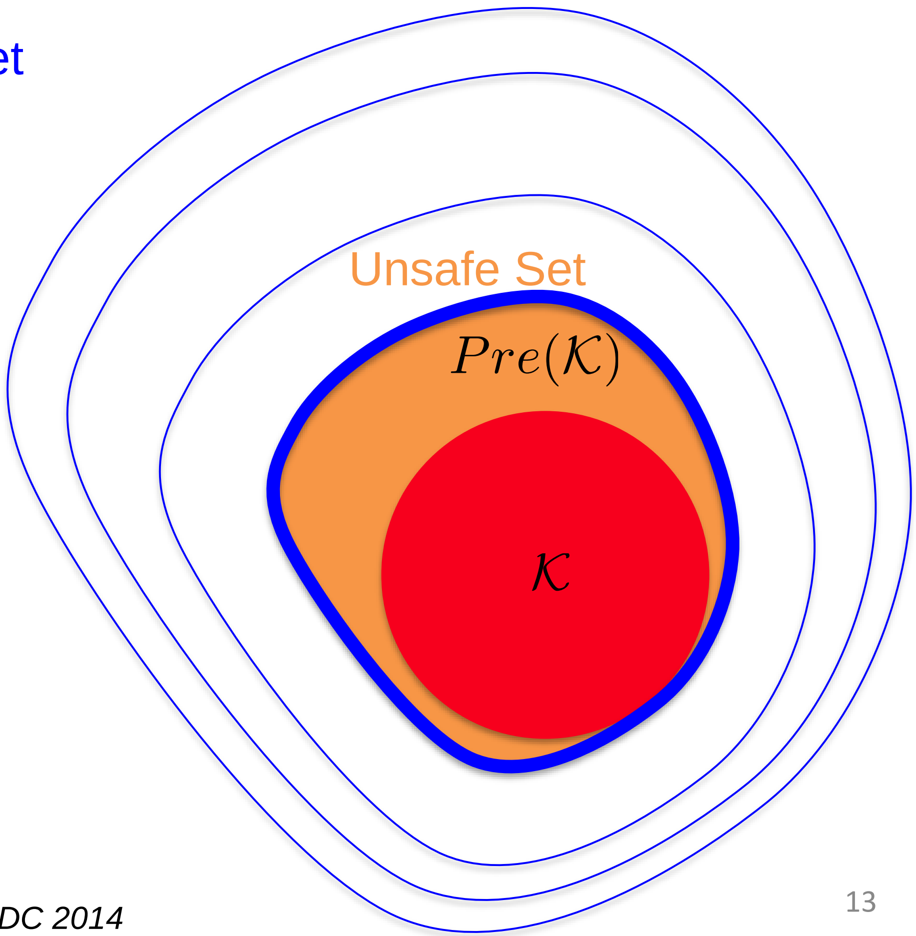

Akametalu, Fisac, Zeilinger, Kaynama, Gillula, and Tomlin, CDC 2014

Model Validation

Validate model online and adjust controller

Infinite number of candidate unsafe sets for $\mathcal{D}(x)$

Initialize:
Active unsafe set = smallest candidate set

Validate measured disturbance at
visited states against $\mathcal{D}(x)$

Detected model inaccuracy

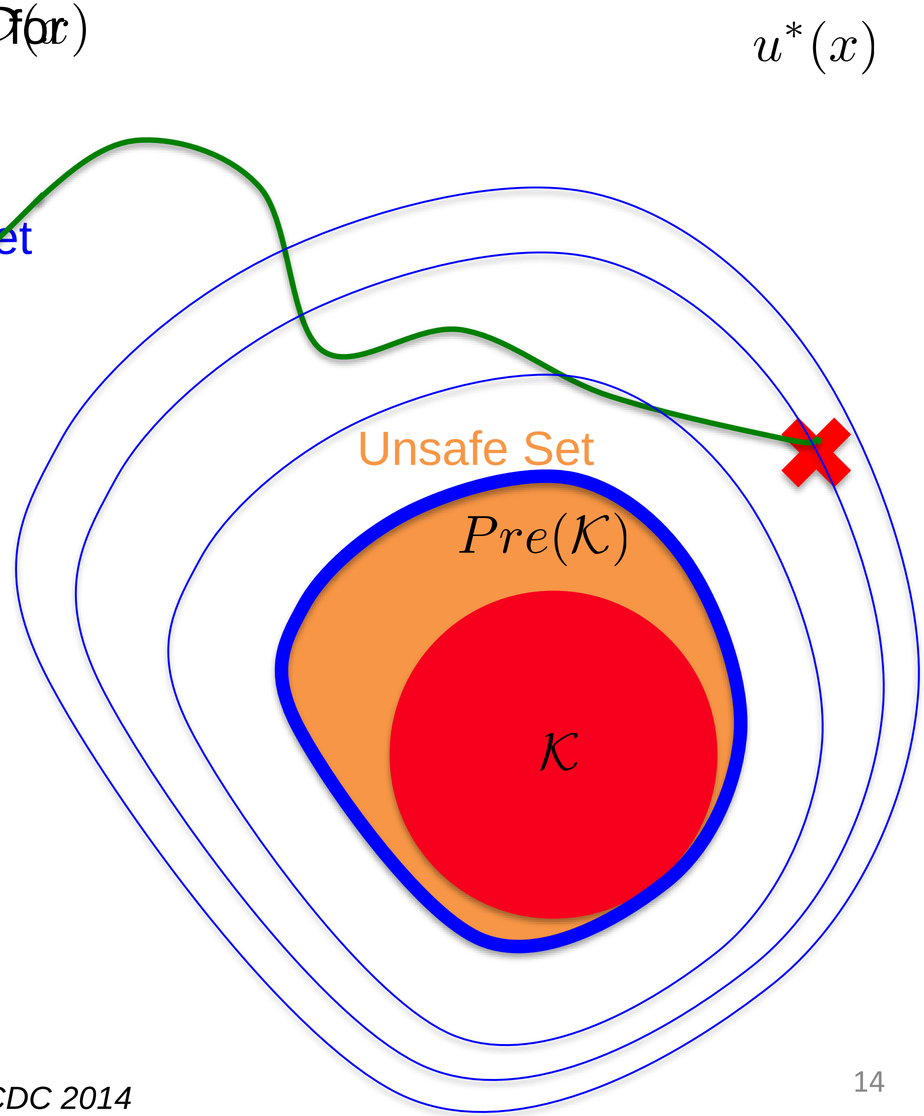

Akametalu, Fisac, Zeilinger, Kaynama, Gillula, and Tomlin, CDC 2014

Model Validation

Validate model online and adjust controller

Infinite number of candidate unsafe sets for $\mathcal{D}(x)$

Initialize:
Active unsafe set = smallest candidate set

Validate measured disturbance at visited states against $\mathcal{D}(x)$

Detected model inaccuracy

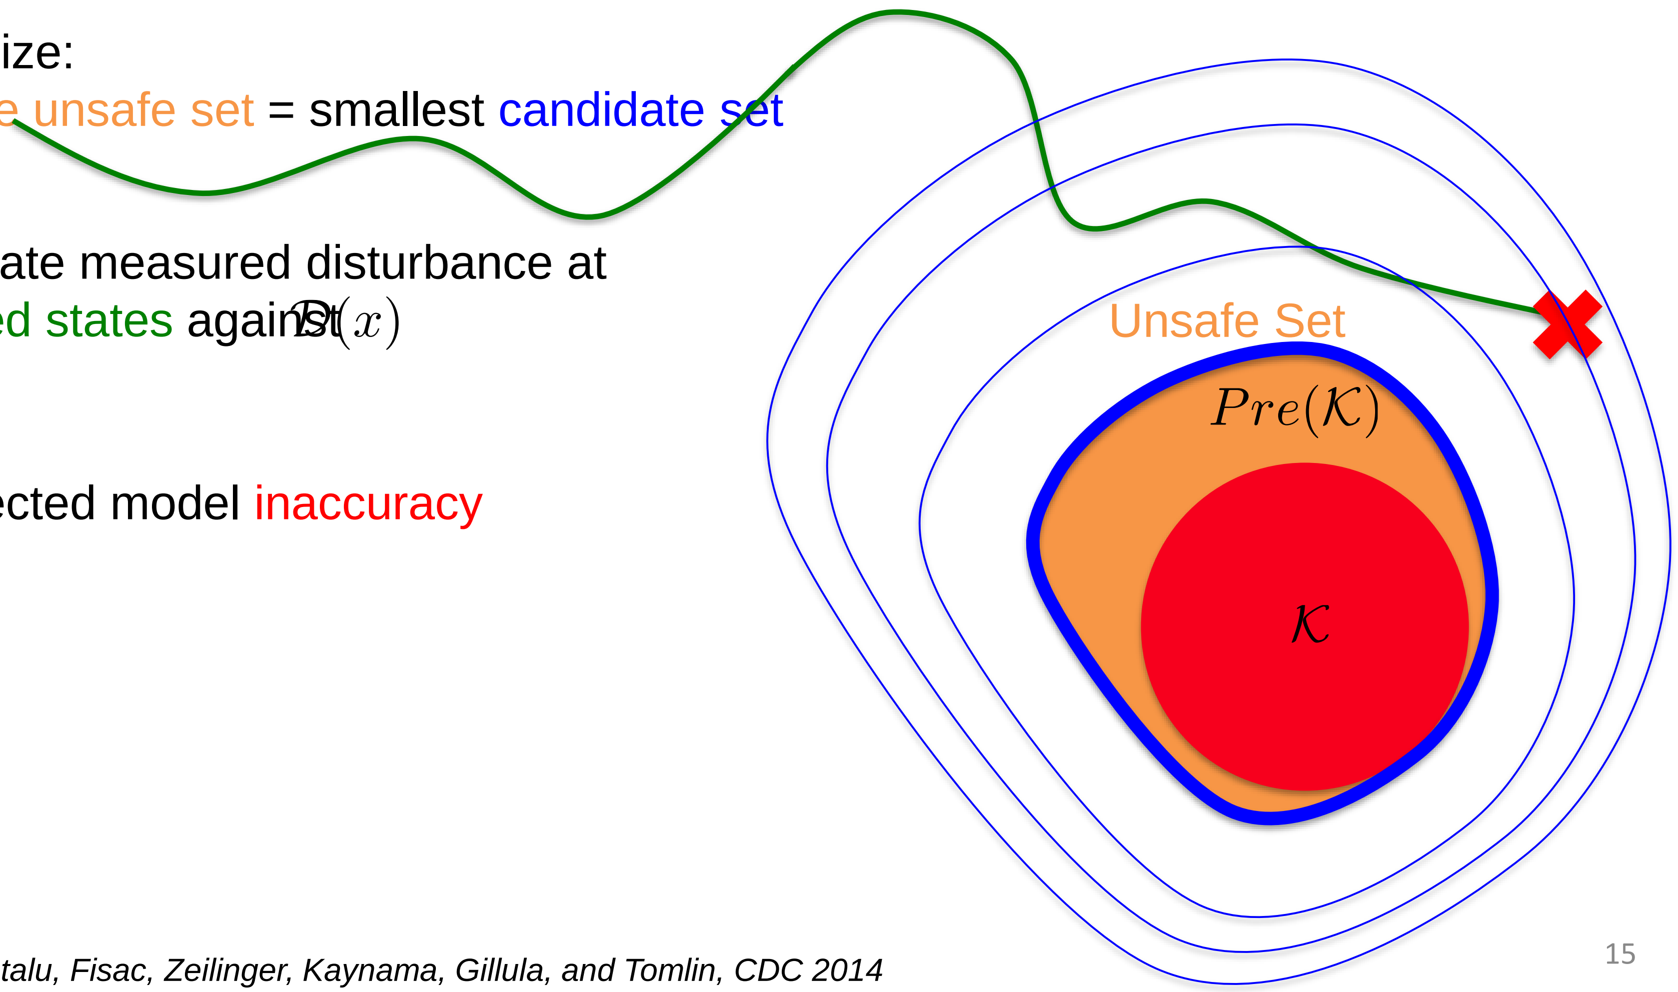

Akametalu, Fisac, Zeilinger, Kaynama, Gillula, and Tomlin, CDC 2014

Model Validation

Validate model online and adjust controller

Infinite number of candidate unsafe sets for $\mathcal{D}(x)$

Initialize:
Active unsafe set = smallest candidate set

Validate measured disturbance at visited states against $\mathcal{D}(x)$

Detected model inaccuracy

Expand unsafe set to include current sate

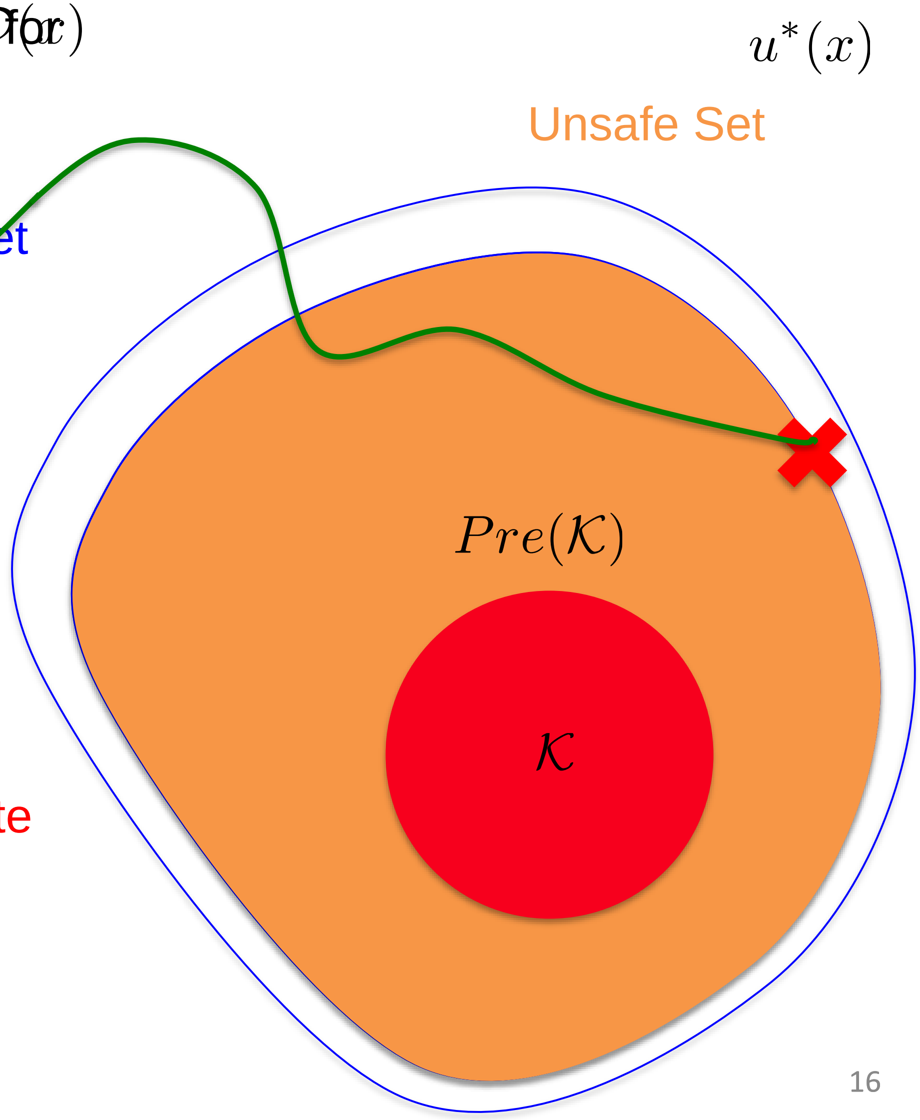

Model Validation

Validate model online and adjust controller

Infinite number of candidate unsafe sets for $\mathcal{D}(x)$

Initialize:
Active unsafe set = smallest candidate set

Validate measured disturbance at
visited states against $\mathcal{D}(x)$

Detected model inaccuracy

Expand unsafe set to include current sate

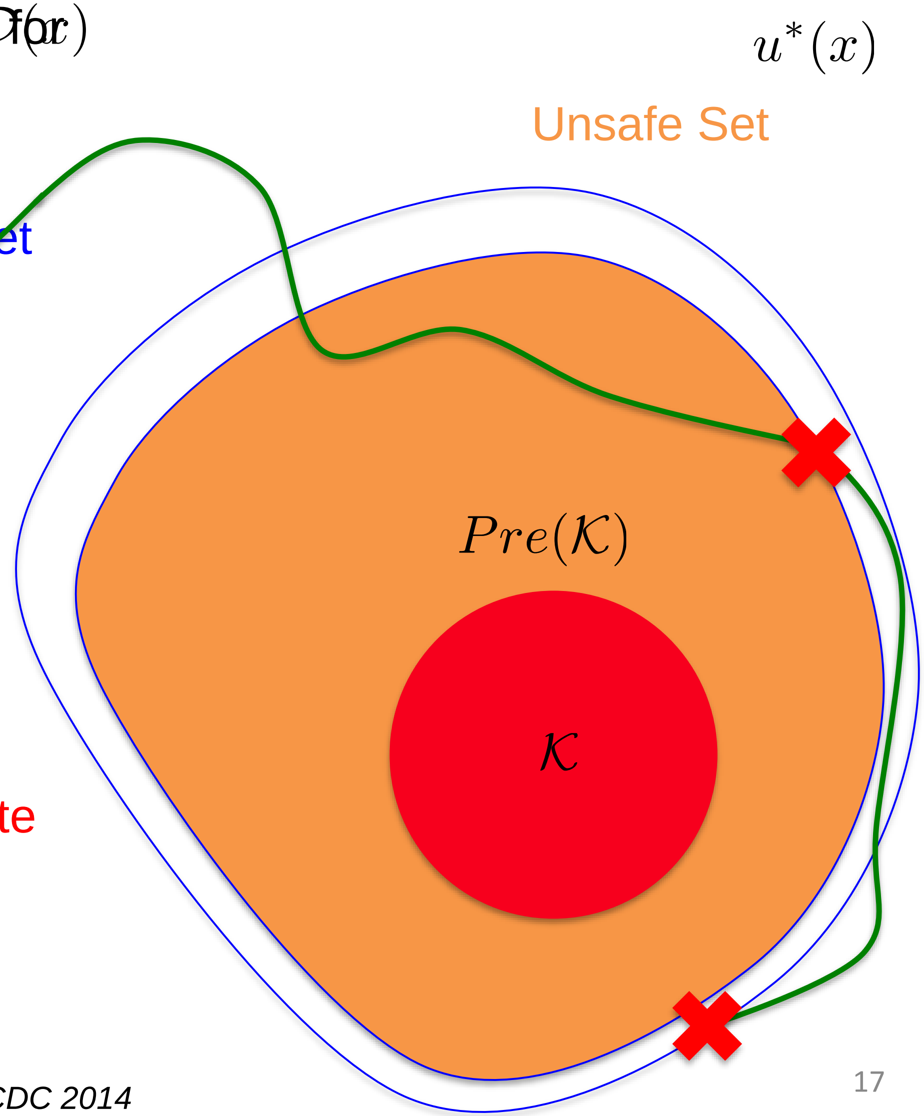

Akametalu, Fisac, Zeilinger, Kaynama, Gillula, and Tomlin, CDC 2014

Model Validation

Validate model online and adjust controller

Infinite number of candidate unsafe sets for $\mathcal{D}(x)$

Initialize:
Active unsafe set = smallest candidate set

Validate measured disturbance at visited states against $\mathcal{D}(x)$

Detected model inaccuracy

Expand unsafe set to include current sate

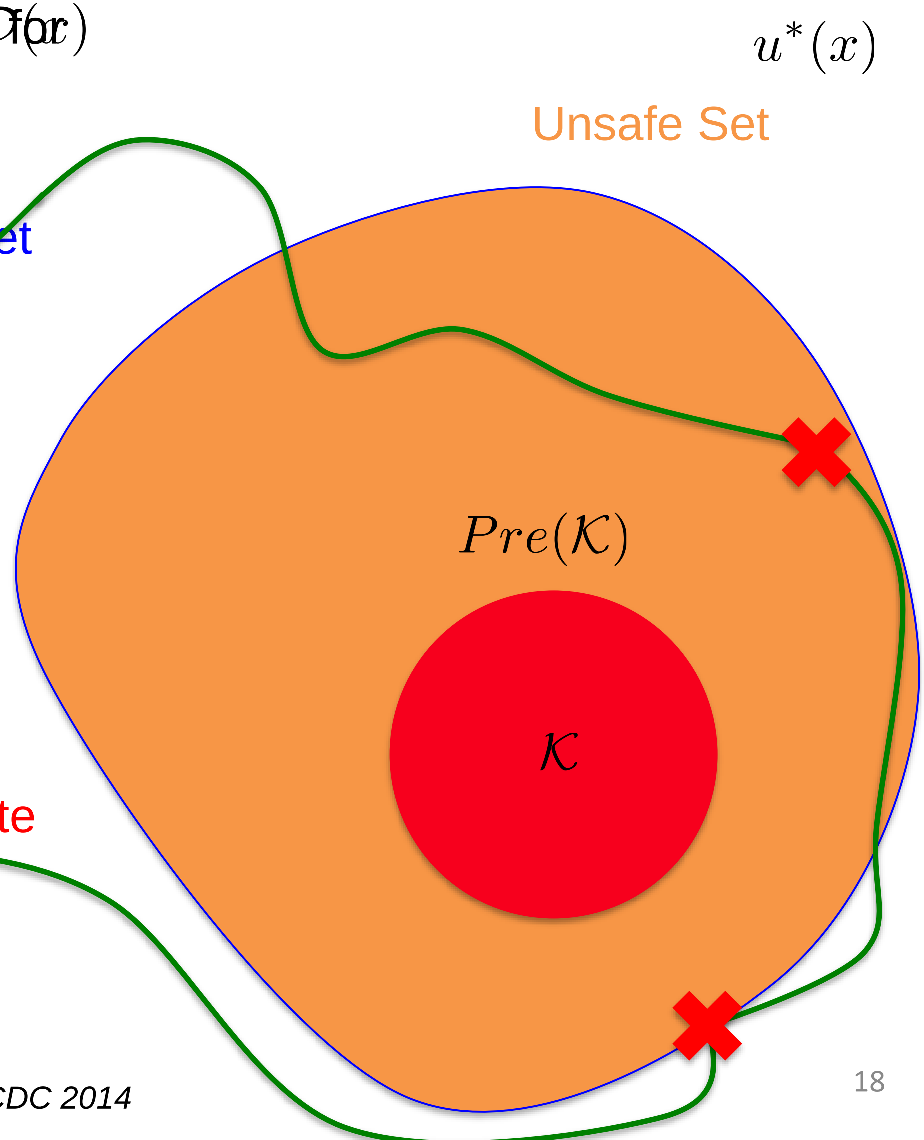

Akametalu, Fisac, Zeilinger, Kaynama, Gillula, and Tomlin, CDC 2014

Model Validation

Validate model online and adjust controller

Infinite number of candidate unsafe sets for $\mathcal{D}(x)$

Initialize:
Active unsafe set = smallest candidate set

Validate measured disturbance at
visited states against $\mathcal{D}(x)$

Detected model inaccuracy

Expand unsafe set to include current sate

Recompute $\mathcal{D}(x)$ and repeat

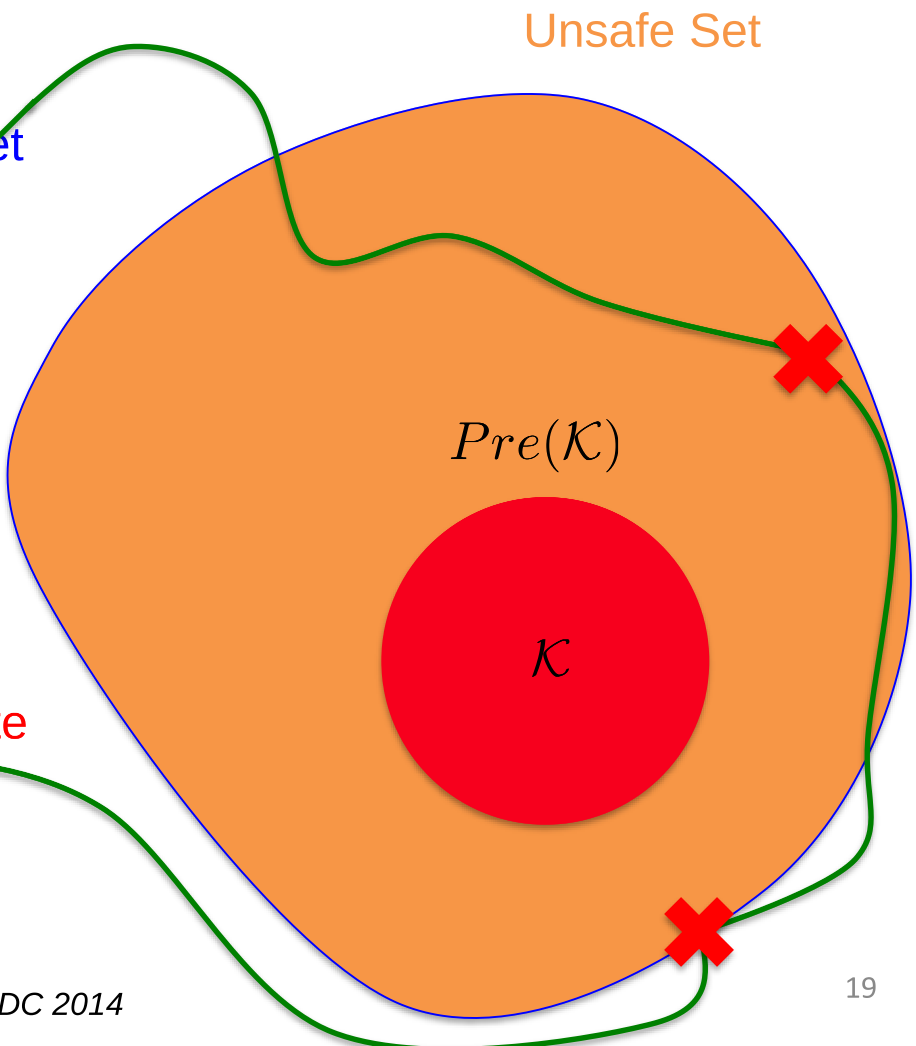

Akametalu, Fisac, Zeilinger, Kaynama, Gillula, and Tomlin, CDC 2014

Model Validation

Validate model online and adjust controller

Infinite number of candidate unsafe sets for $\mathcal{D}(x)$

$u^*(x)$

Initialize:
Active unsafe set = smallest candidate set

Validate measured disturbance at
visited states against $\mathcal{D}(x)$

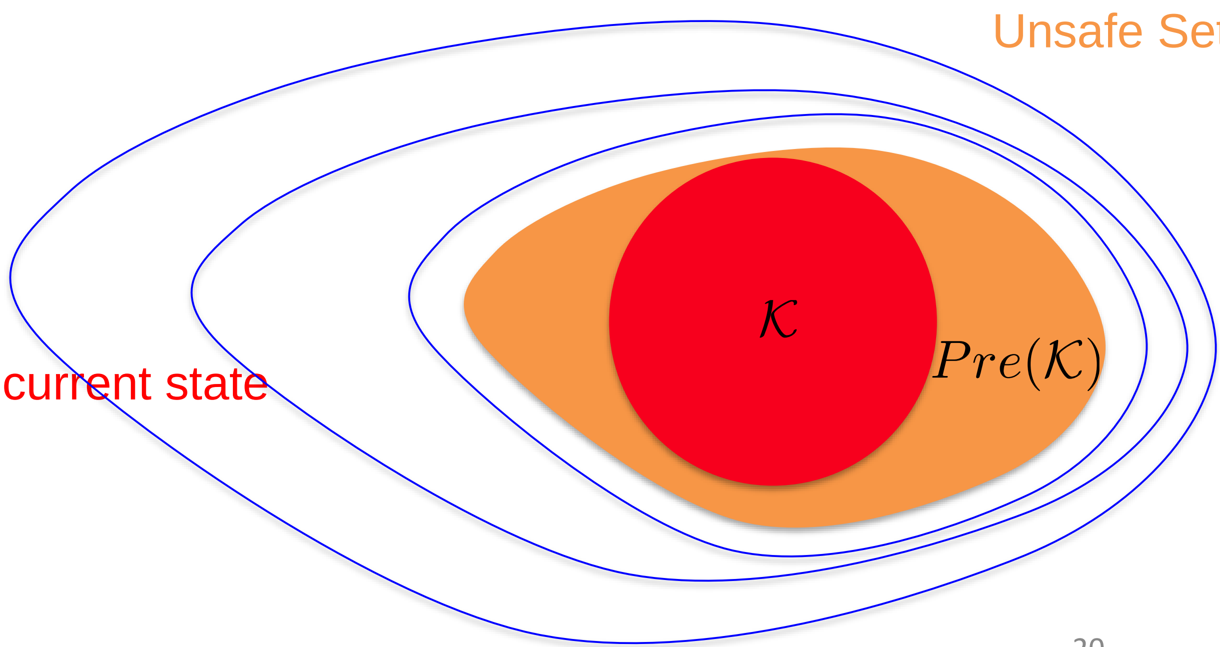

Detected model inaccuracy

Expand unsafe set to include current state

Recompute $\mathcal{D}(x)$ and repeat

Akametalu, Fisac, Zeilinger, Kaynama, Gillula, and Tomlin, CDC 2014

Model Validation (Demo)

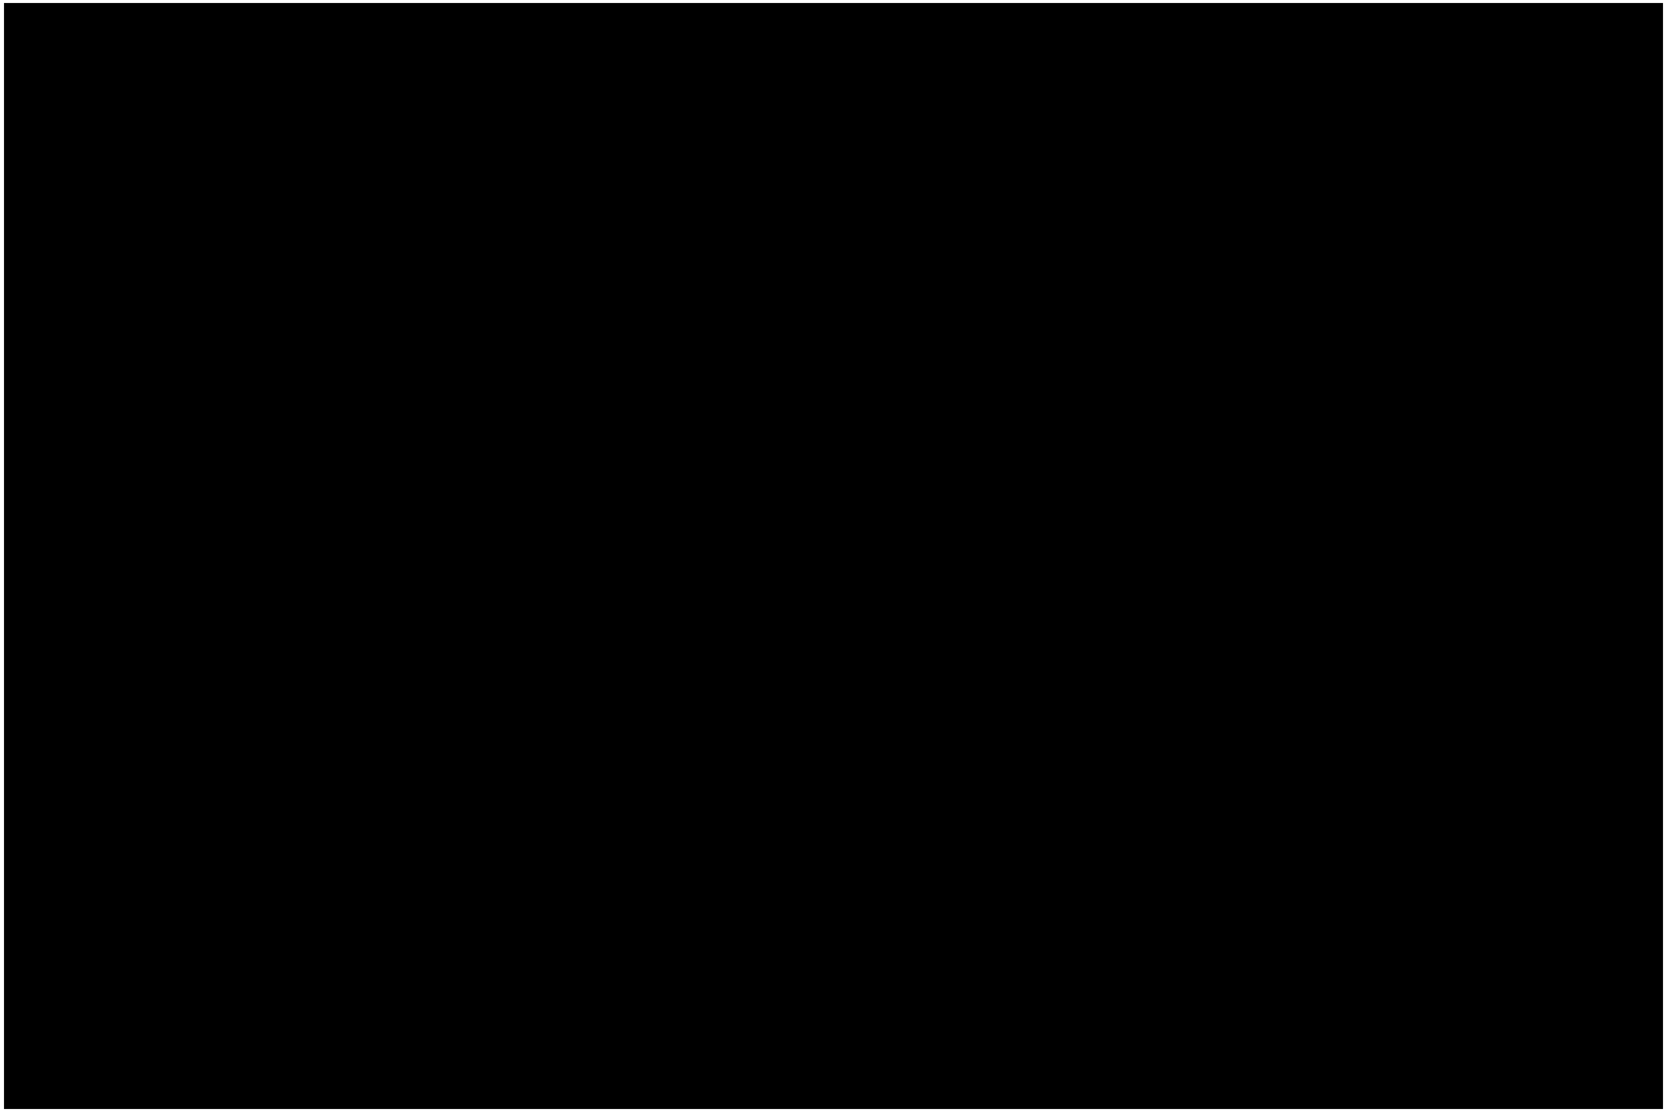

Fisac, Akametalu , Zeilinger, Kaynama, Gillula, and Tomlin, in progress

Summary

Framework for Safe Learning

General

Constraints satisfied given accurate model

Model Inference

Reduced conservatiness

Allows flexibility in learning new models

Model Validation

Computationally cheap online model validation

Robust to modeling error

Allows flexibility in learning new models

Thank you. Questions?