

Towards Safe Aviation Autonomy

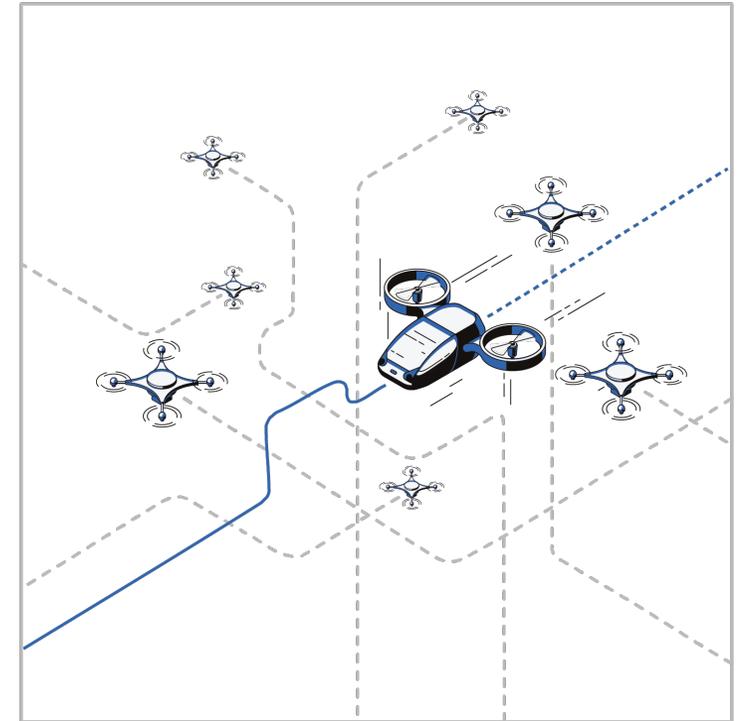
Marco Pavone

Autonomous Systems Laboratory
Stanford University

Autonomous Vehicle Research Group
NVIDIA

NSF-PIRE Workshop
June 9th, 2021

pavone@stanford.edu



Stanford
University



Advanced air mobility is on the rise

DRONE INDUSTRY INSIGHTS
The Drone Market Environment 2019

The infographic is a large grid of logos categorized into several main sections:

- Hardware:**
 - Drone Platforms:** Includes sub-sections for Agriculture (Aerodrone, AgriDrone, etc.), Delivery Systems (Dronamics, Zipline, etc.), Safety & Security (Arcturus, etc.), Fixed-Wing (Aeromoo, etc.), VTOL Fixed-Wing (Aerovell, etc.), and Lighter-Than-Air (Aense, etc.).
 - Counter-Drone Solutions:** Includes Aeronaught, AeroDefense, etc.
 - Components & Systems:** Includes navigation and guidance systems (AdaPilot, etc.), propulsion and power (Propeller, etc.), and launch and recovery systems (AeroLaunch, etc.).
- Software:**
 - Flight, Fleet & Operation Management:** Includes Aeriosense, AirHub, etc.
 - Open Source Infrastructure, SDK:** Includes DroneKit, etc.
 - Navigation, CV and AI:** Includes Shield AI, etc.
 - UTM, LAANC Suppliers:** Includes AirMap, etc.
 - Data Analytics, Workflow, CV and AI:** Includes SDR, etc.
- Services:**
 - Drone-as-a-Service Providers:** Includes Aerium, etc.
 - Education, Simulation, Training:** Includes AeroCameras, etc.
 - System Integration, Engineering, Advisory:** Includes AlphaLink, etc.
 - Maintenance:** Includes Drone Doctor, etc.
 - Market Research & Consulting:** Includes Drone Analyst, etc.
 - User Groups, Networks:** Includes Drone Community, etc.
 - Media, News, Blogs & Magazines:** Includes Drone News, etc.
 - Marketplaces:** Includes Drone Marketplace, etc.
 - Shows, Conferences, Events:** Includes Drone Expo, etc.
 - Coalitions, Organizations & Initiatives:** Includes Drone Alliance, etc.
 - Insurance:** Includes Drone Insurance, etc.

Key technological enabler:
vehicle autonomy

AI has been fueling self-driving car progress

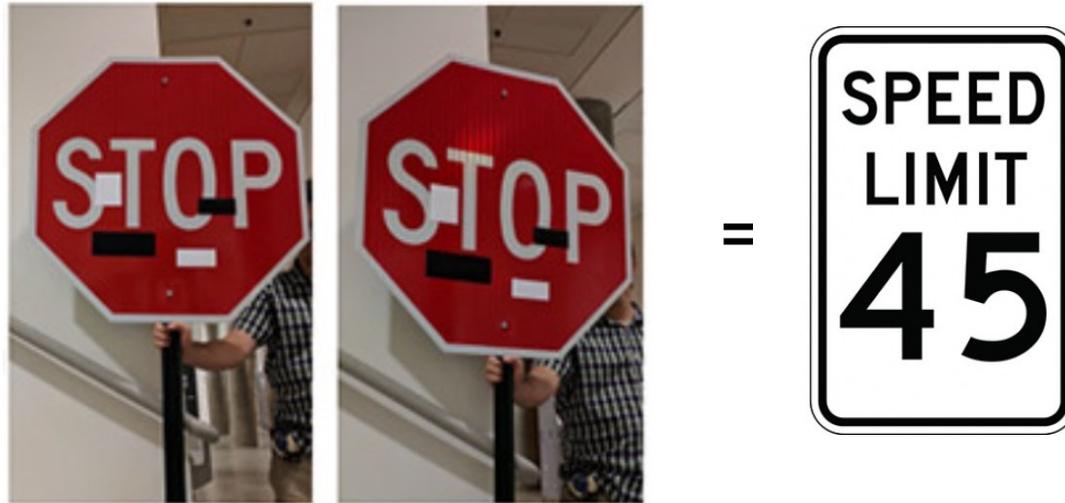


2005



2017

Yet, AI-based algorithms still suffer from unpredictability and erratic behavior

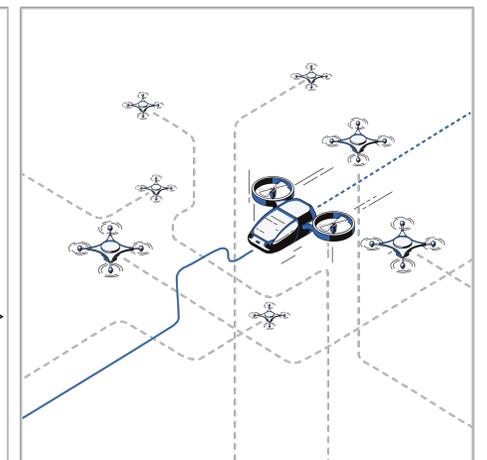
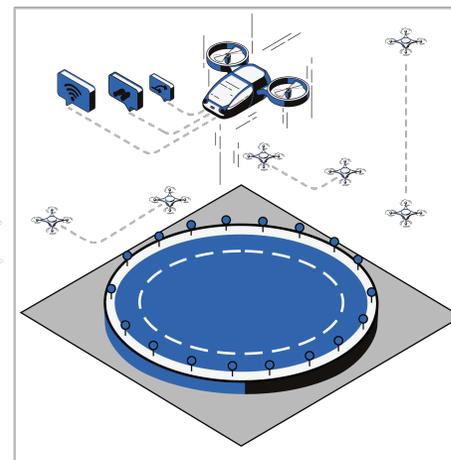
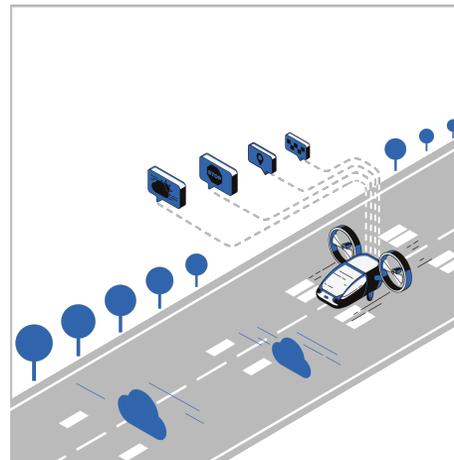


NASA ULI center– motivation and project goal

Challenge

Can we trust AI for *safety-critical* applications, and in particular for future aviation systems?

Project goal: develop mathematical and algorithmic techniques to enable *trusted* AI-based aviation systems



NASA ULI center– approach

Technical approach:

1. Verified AI-based software components
2. Fault detection, isolation, and recovery (FDIR) for AI-based aviation systems
3. Tools for airspace management with AI-driven agents



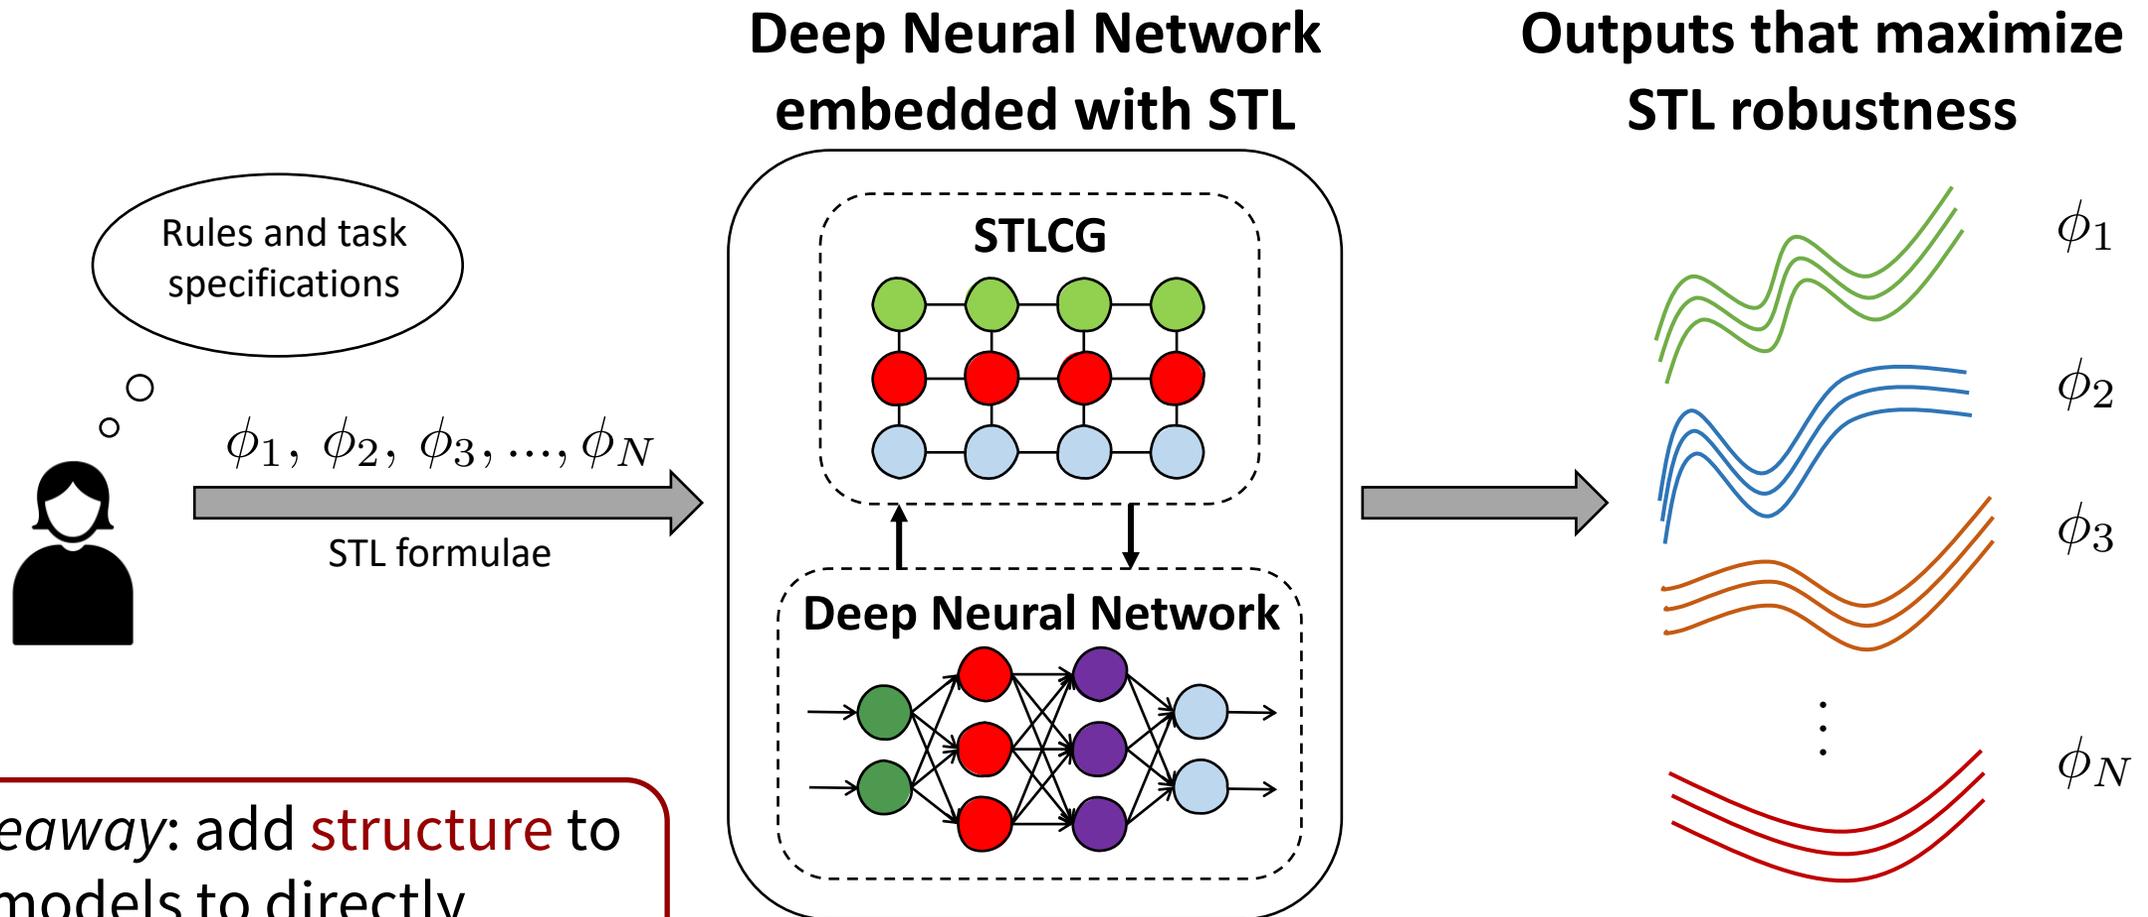
System integration and validation:

1. Testing at the Boeing Flight & Autonomy Lab at Stanford
2. Graduation exercises with hardware-in-the-loop and real-world flight testing



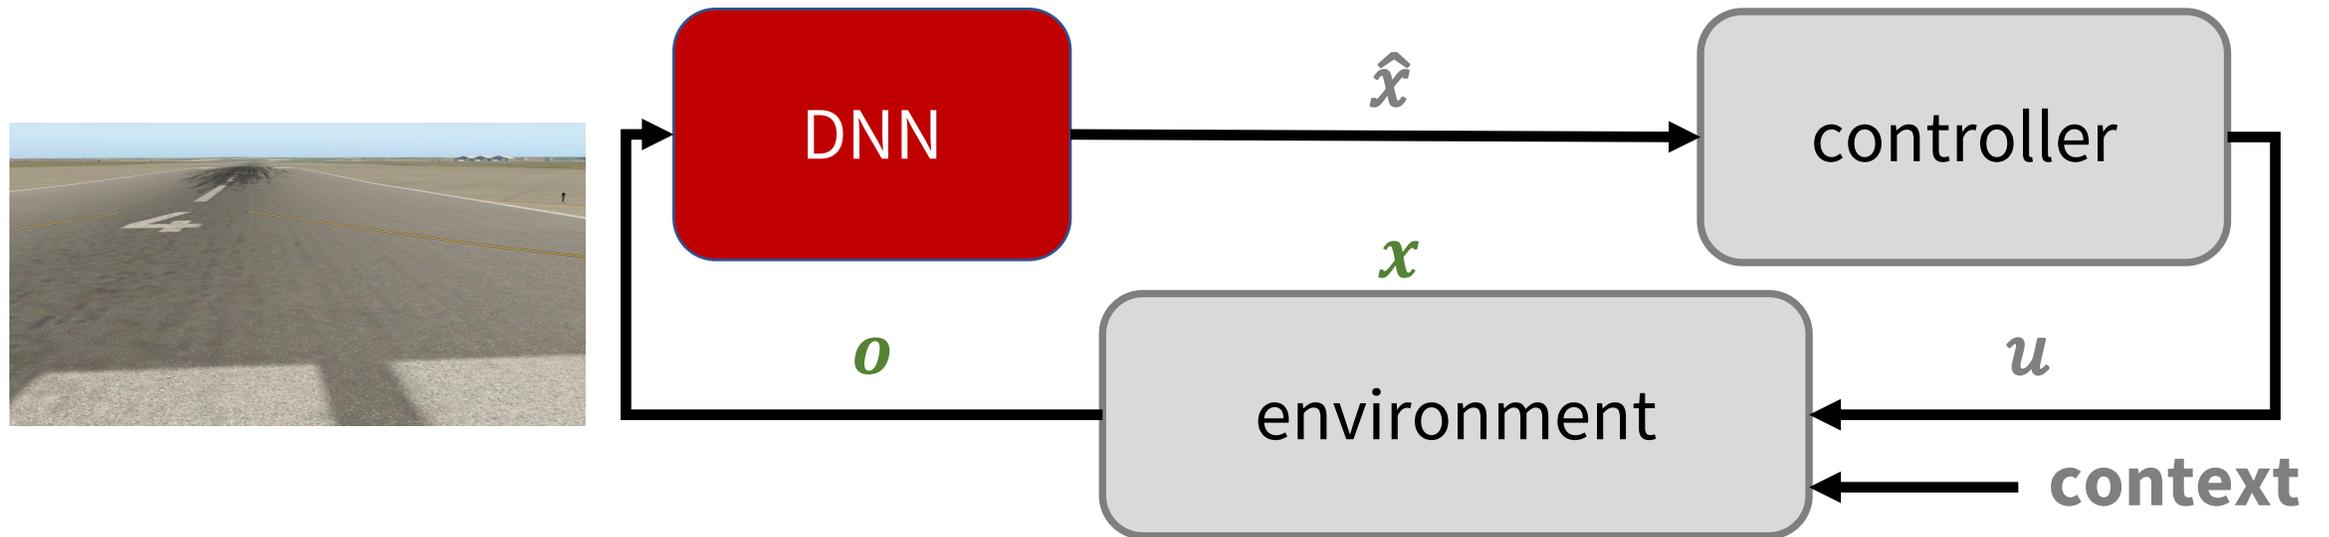
Adding structure to models

[Leung, Arechiga, Pavone, WAFR '20, Leung, Pavone, CDC '21]



Takeaway: add **structure** to ML models to directly embed desirable properties

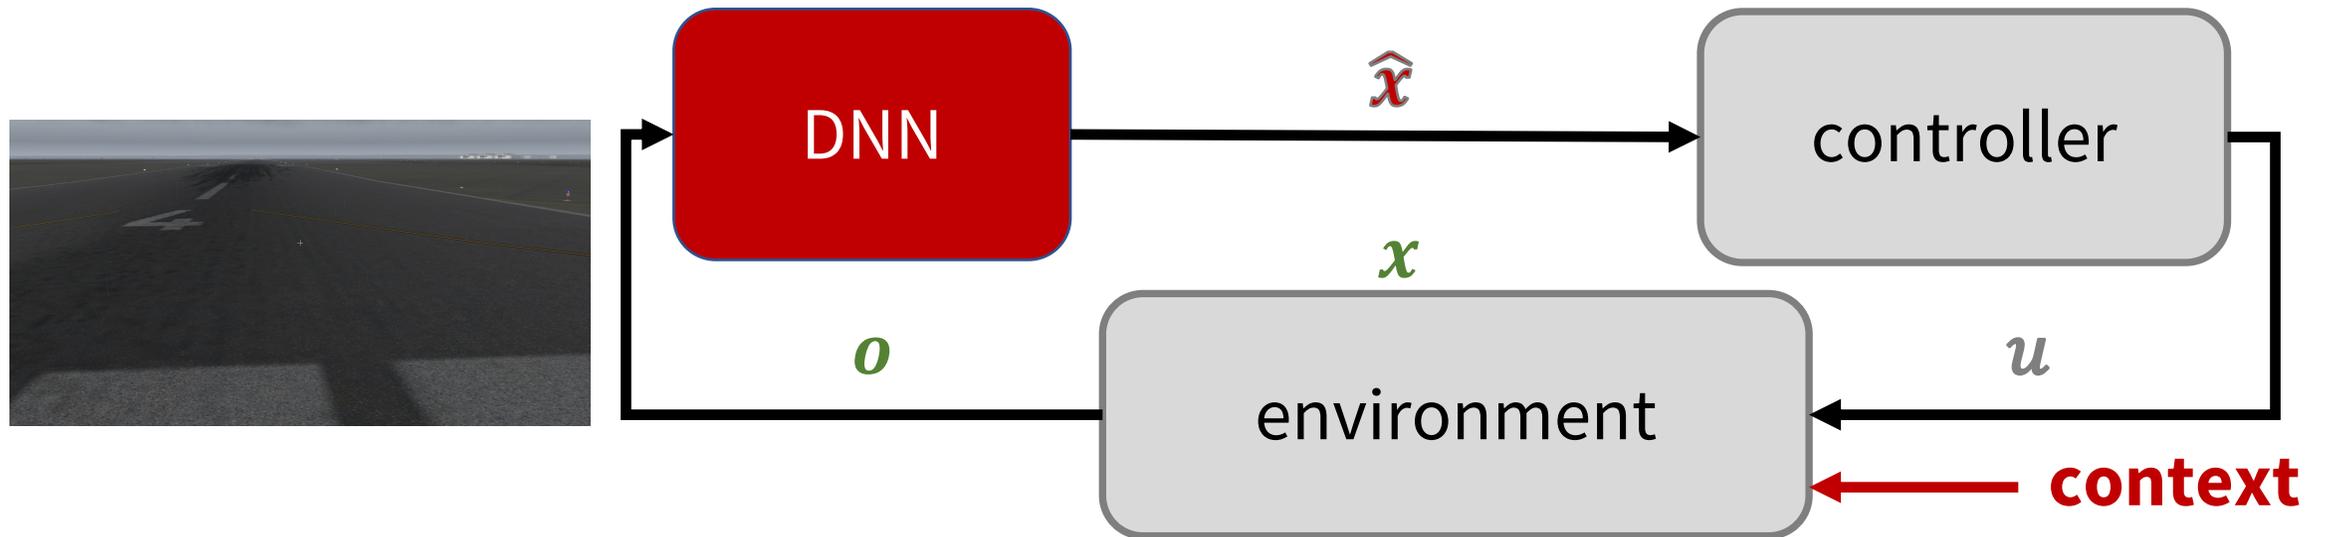
DNNs are only as good as their datasets



DNN trained on $(o, x) \sim p_{\text{train}}(o, x)$

These examples come from specific contexts (e.g., weather, time of day)

DNNs are only as good as their datasets

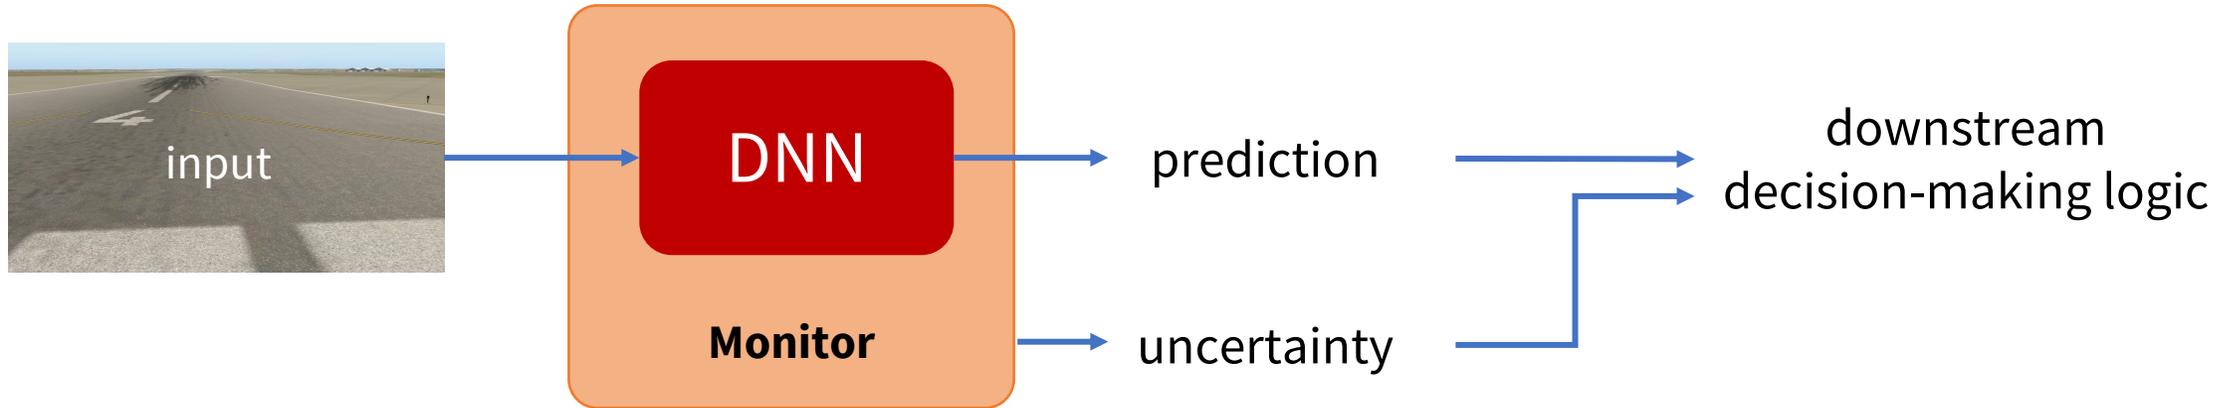


DNN trained on $(o, x) \sim p_{\text{train}}(o, x)$

Out-of-distribution observations can lead to **faults in perception**, and impact **closed-loop performance**

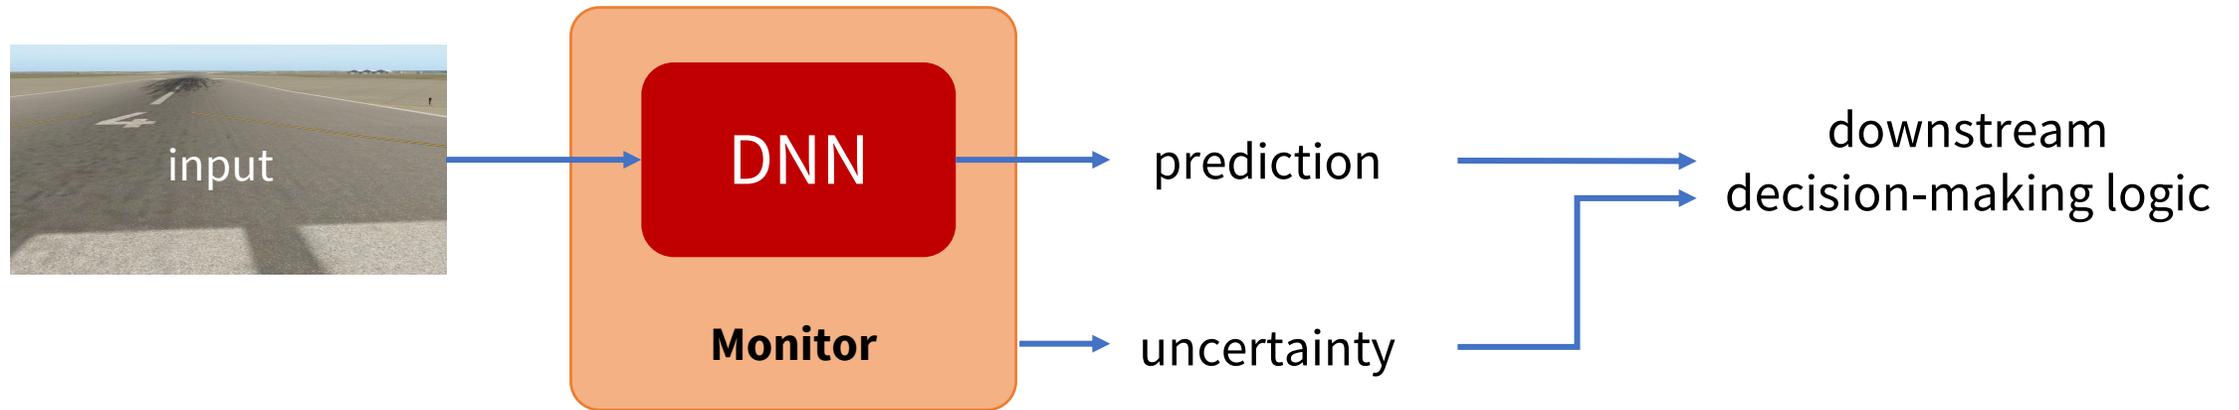
Run-time monitors for anomaly detection

[Sharma, Azizan, Pavone, UAI '21]



Run-time monitors for anomaly detection

[Sharma, Azizan, Pavone, UAI '21]



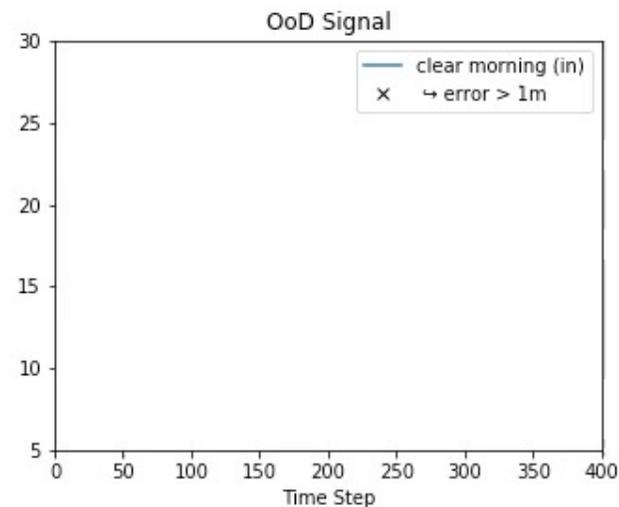
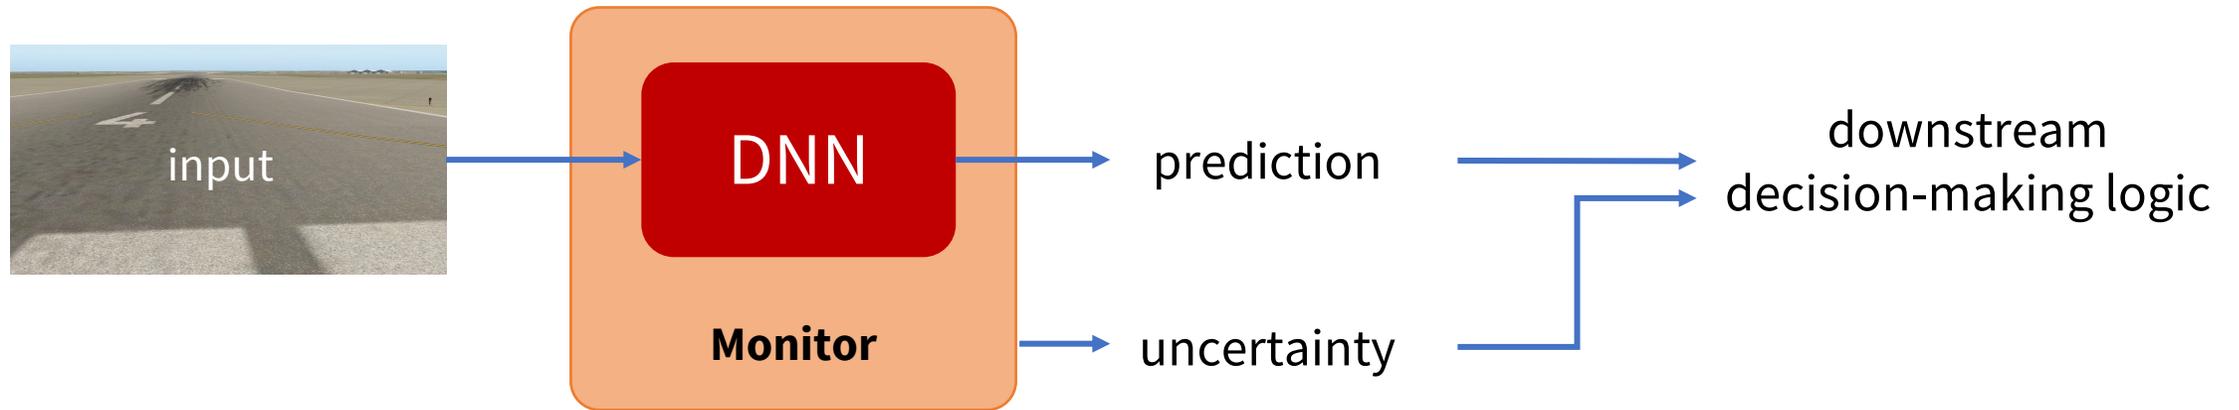
SCOD: Sketching Curvature for OoD Detection

A **general purpose**, **model-agnostic** approach to adding an **efficient** (minimal computational burden) and **informative** (capable of detecting OoD inputs) uncertainty measure to **pre-trained** DNN models



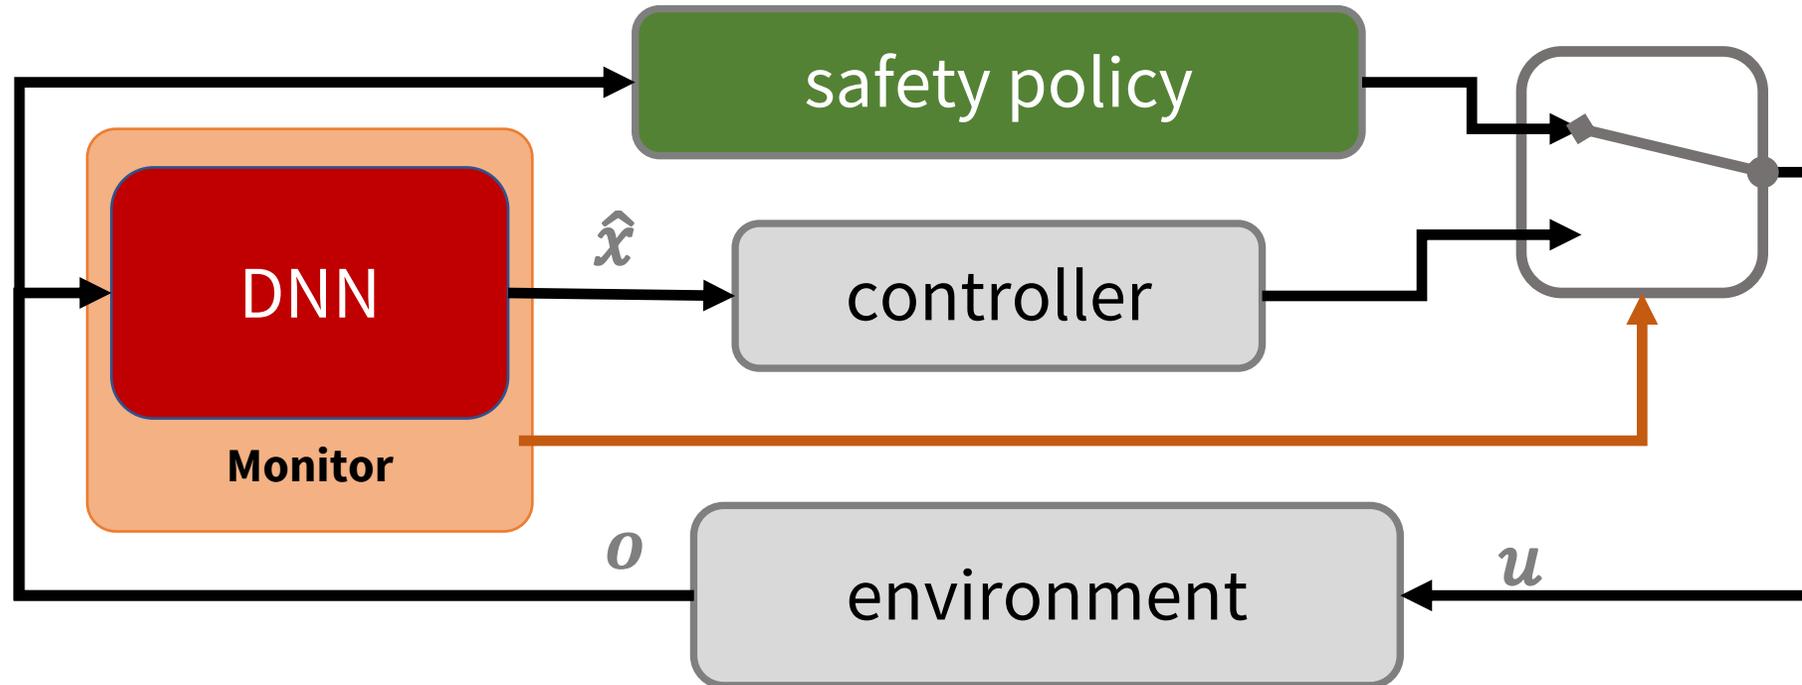
Run-time monitors for anomaly detection

[Sharma, Azizan, Pavone, UAI '21]



Takeaway: run-time monitors are key to build high-confidence AI systems

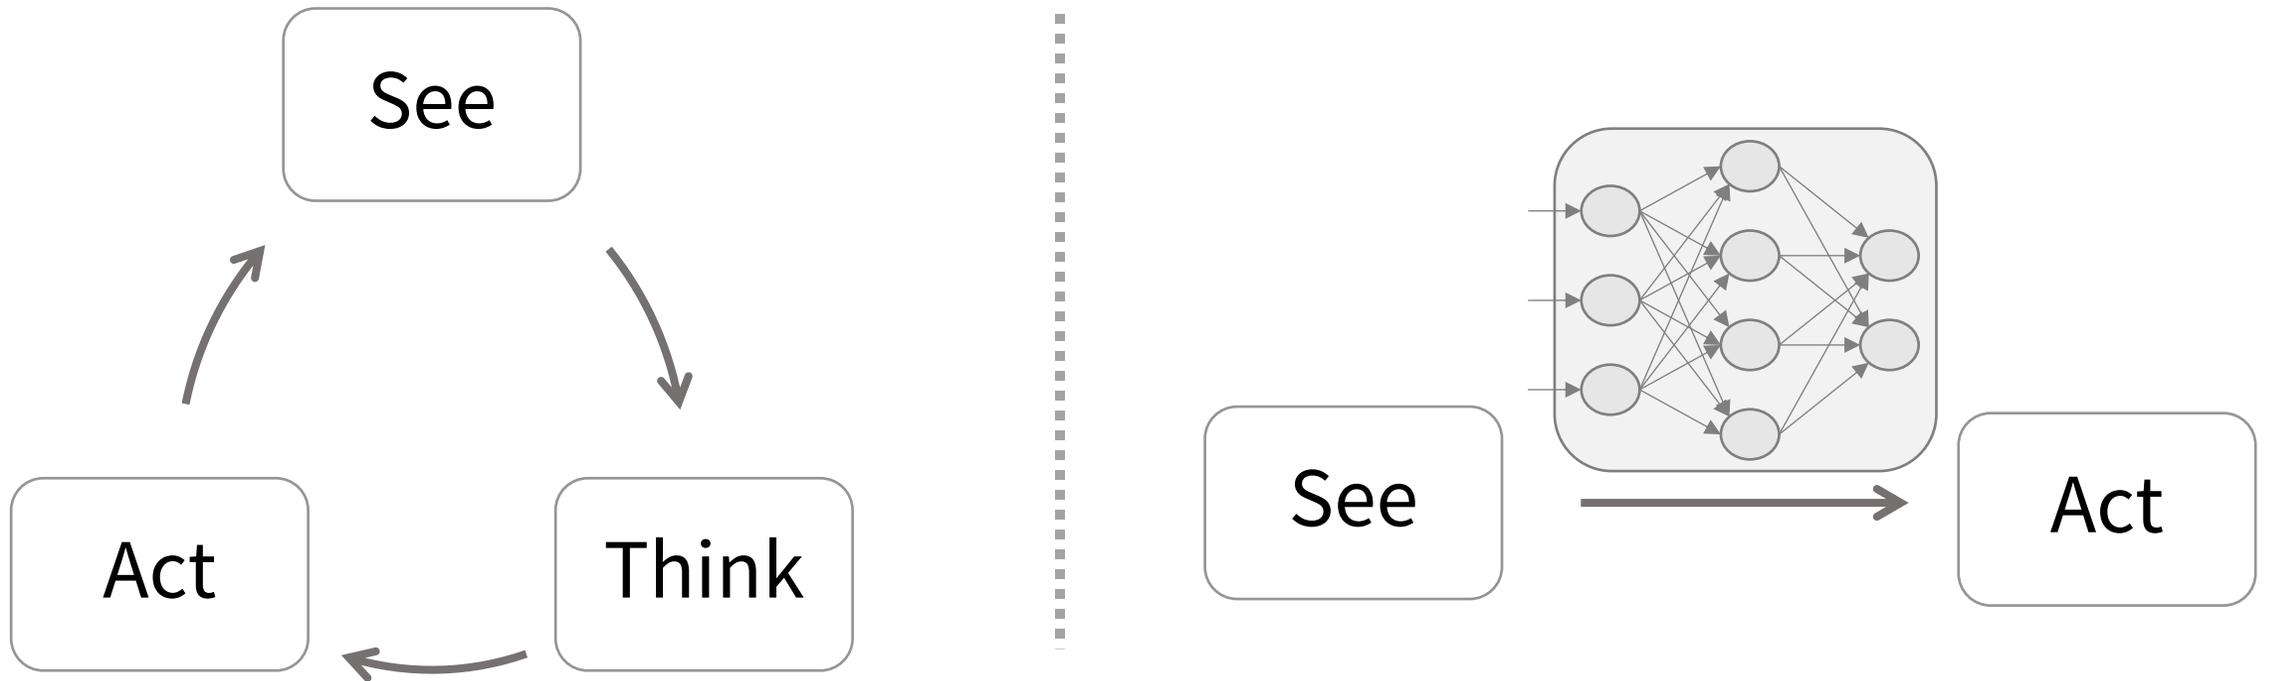
Reacting to anomalies



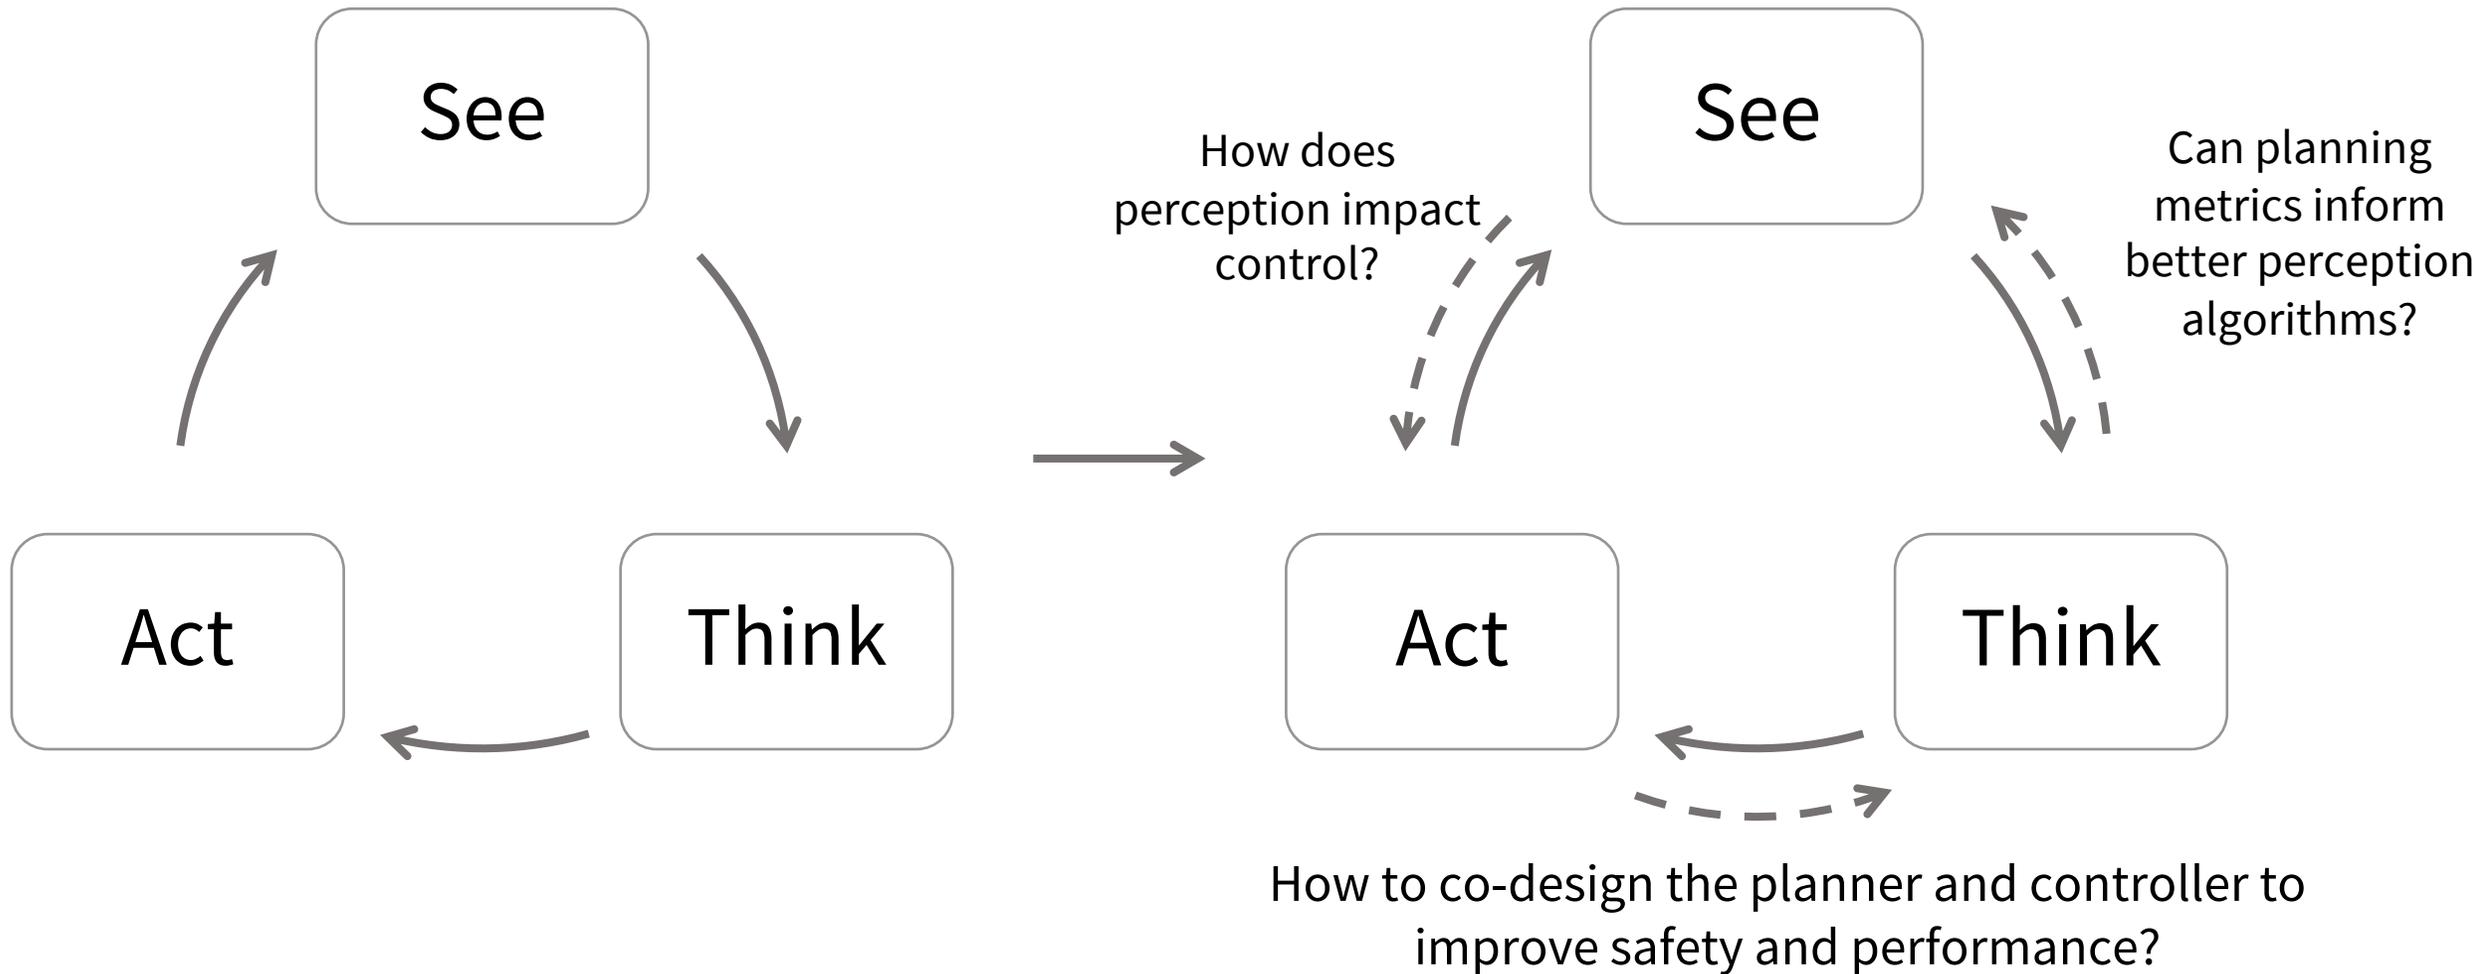
Takeaway: safe operations require seamless transition from AI-based policies to **recovery policies** (and back)

Can we architect things better?

On architecting the autonomy stack



Highly integrated yet modular autonomy stacks



Conclusions

