AMR LAB

UVA ENGINEERING
LINK LAB

IROS KYOTO 2022
IEEE/RSJ International Conference on Intelligent Robots and Systems

Meta-Learning-Based Proactive Online Planning for UAVs Under Degraded Conditions

Esen Yel[1]*, Shijie Gao[2]*, Nicola Bezzo[2]

*Equal Contribution

[1]Now at Stanford University, [1]University of Virginia

IEEE/RSJ International Conference on Intelligent Robots and Systems

IEEE Robotics and Automation Letters (RA-L)

Introduction

Motivations:

- Autonomous robotic systems are subject to many challenges in real world applications such as:
  - Component failures
  - System aging
  - Model changes
- These situations cause the system to operate under degraded conditions and potentially become unsafe
- The onboard controller or control inputs may not be always accessible

Introduction

Objectives:

- To predict the future states and state uncertainties of a system with an unforeseen fault at runtime
- To monitor if the system will violate safety constraints
- To replan the desired trajectories to improve safety
- To continuously monitor and update the prediction models using runtime data

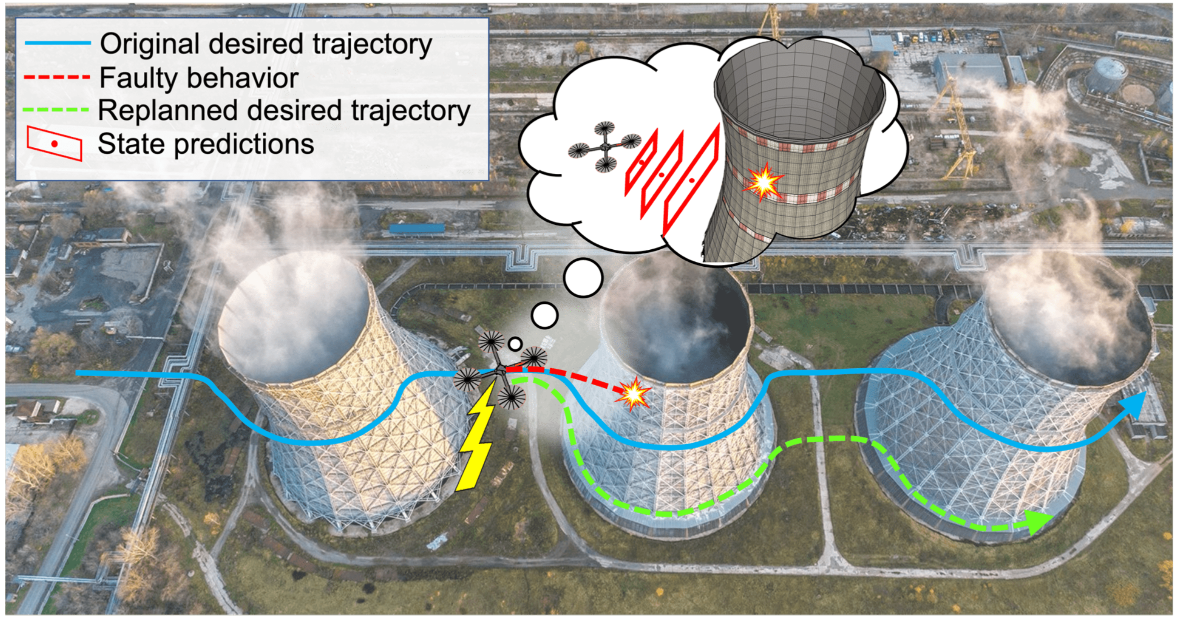

Approach Overview

- Meta-learning-based framework for future state and state uncertainty prediction
- **Offline stage:** Meta-training using data from various actuator faults
- **Online stage:** Predictions and safety monitoring for the new faulty system and replanning to preserve safety

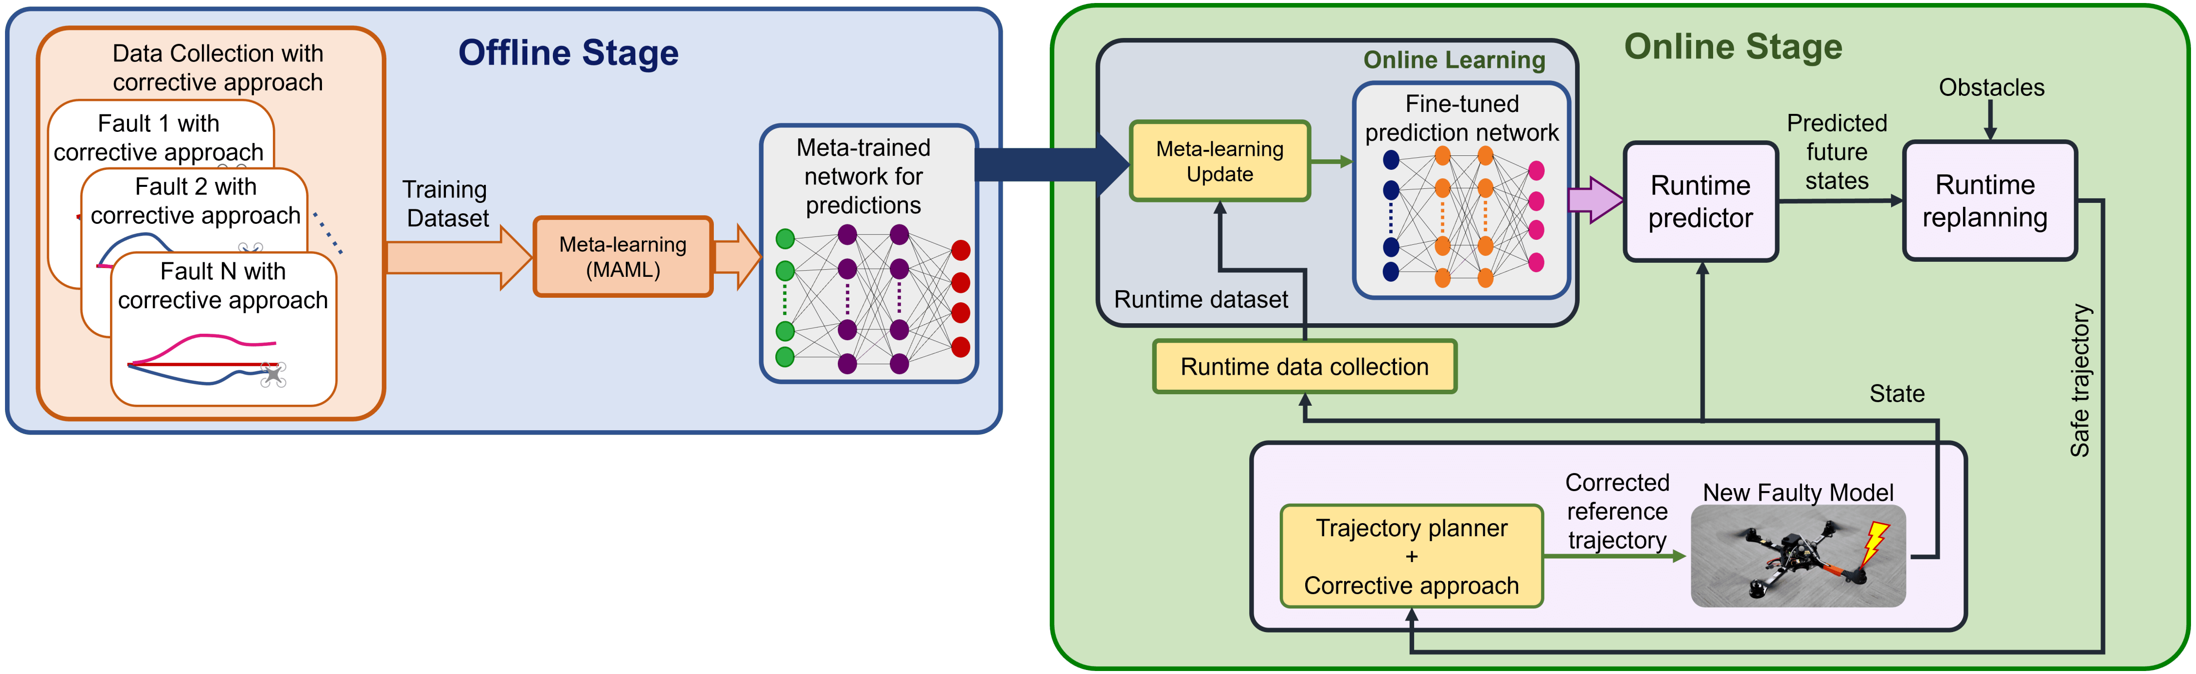

Approach: Offline Meta-Training

- During the offline stage, state and uncertainty data from a discrete fault set are collected.
- The element-wise mean and variance is computed for each fault:

$$\bar{\boldsymbol{x}}_i(k) = \frac{\sum_{j=1}^{N} \boldsymbol{x}_i^j(k)}{N}, \boldsymbol{\sigma}_i(k) = \sqrt{\frac{\sum_{j=1}^{N} |\boldsymbol{x}_i^j(k) - \bar{\boldsymbol{x}}_i(k)|^2}{N-1}}$$

$$\forall k \in [0, T_\tau], \forall i \in \{1, \dots, |\mathcal{F}|\}$$

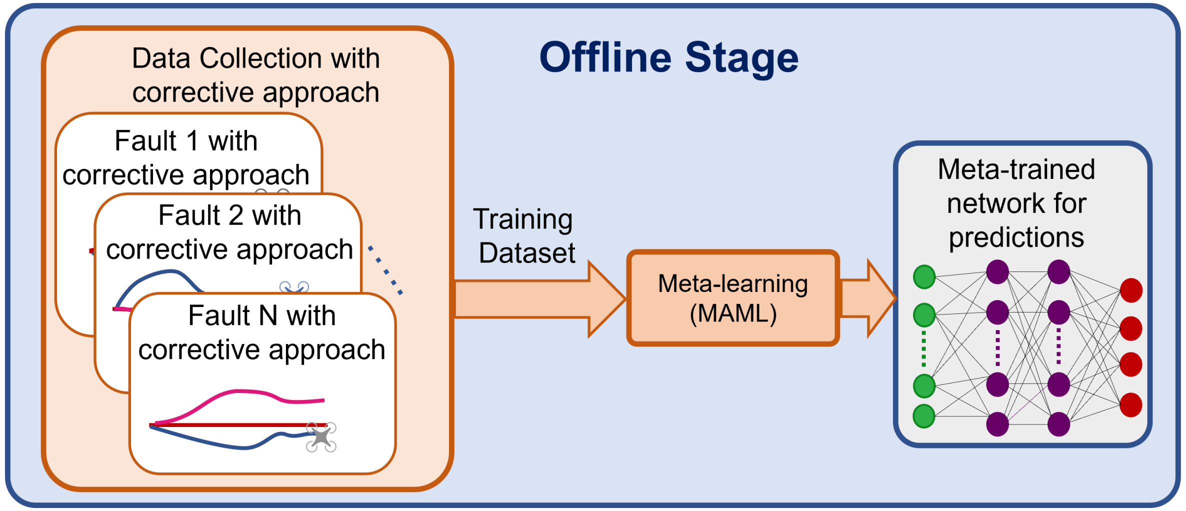

- A meta-network is trained to predict the future states and state uncertainties:

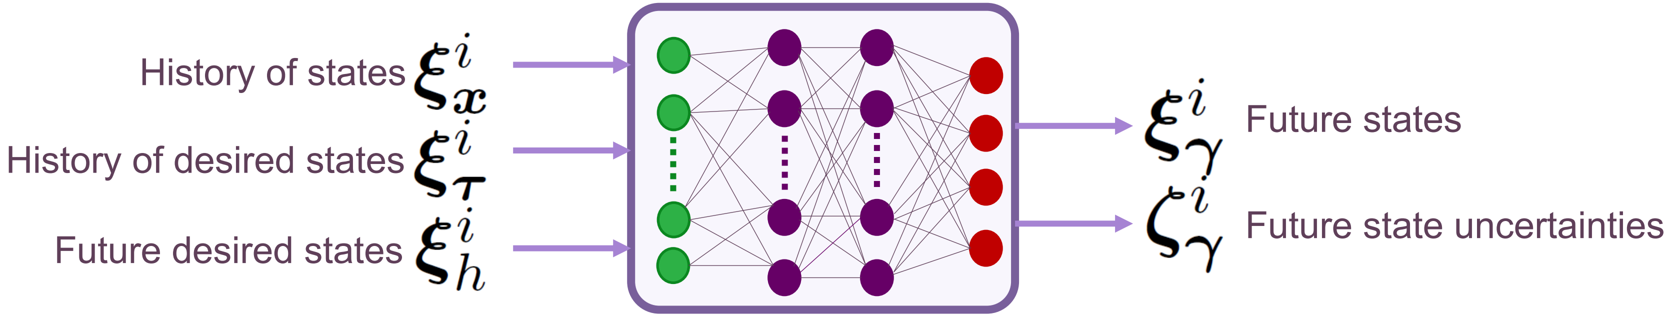

Approach: Online Update

- During the online stage, a meta-learned model is fine-tuned with the online data collected with a new fault:

$$\boldsymbol{\chi}_h^*(k) = \begin{bmatrix} \boldsymbol{\xi}_{\boldsymbol{x}}^*(k) \\ \boldsymbol{\xi}_{\boldsymbol{\tau}}^*(k) \\ \boldsymbol{\xi}_h^*(k) \end{bmatrix} \qquad \gamma_h^*(k) = \begin{bmatrix} \xi_\gamma^*(k) \\ \zeta_\gamma^*(k) \end{bmatrix}$$

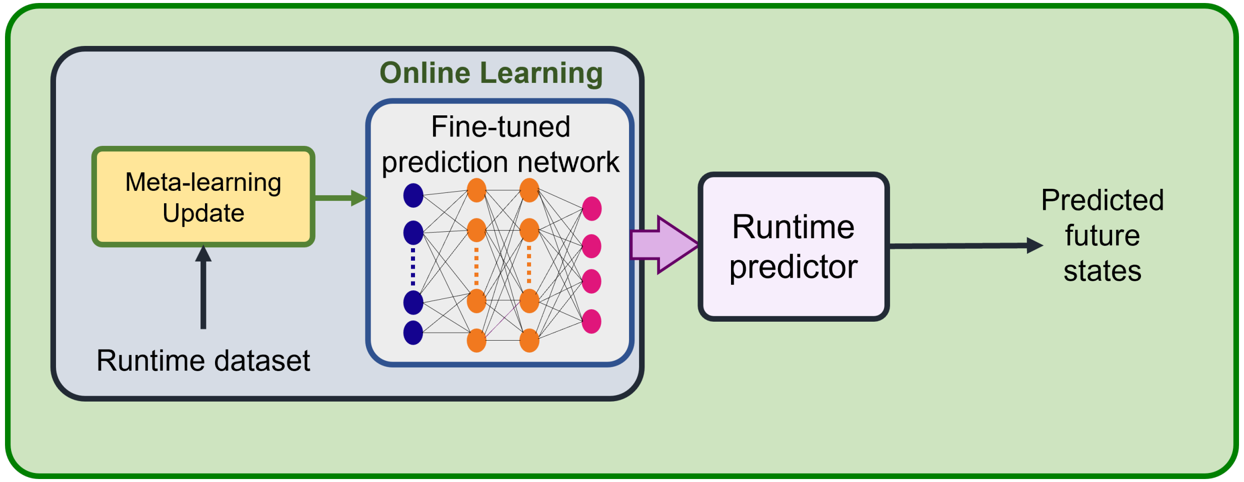

- **Runtime validation:** We constantly monitor the model and re-update it if:
  - The observed states leave the predicted region
  - The distance between the runtime input and training set becomes larger than a given threshold

Approach: Online Predictions and Replanning

- The network is used to predict future states and state uncertainties at runtime:

$$\begin{bmatrix} \tilde{x}(k+\delta_H : \delta_H : k+H) \\ \tilde{y}(k+\delta_H : \delta_H : k+H) \\ \tilde{\sigma}_x(k+\delta_H : \delta_H : k+H) \\ \tilde{\sigma}_y(k+\delta_H : \delta_H : k+H) \end{bmatrix} = \boldsymbol{h}^*_\phi(\boldsymbol{\chi}^*_h(k)) + \begin{bmatrix} x^*(k-T)\vec{\mathbf{1}} \\ y^*(k-T)\vec{\mathbf{1}} \\ 0 \\ 0 \end{bmatrix}$$

$$\forall k \geq T + K_p$$

- If the predicted region collides with obstacles, the trajectory is replanned by sampling and testing waypoints:

$$s^*(k+t) = \begin{cases} 0 & \text{if } \tilde{R}_{\boldsymbol{p}}(k+t) \cap \boldsymbol{\mathcal{O}} \neq \emptyset \\ 1 & \text{otherwise} \end{cases}$$

$$\forall t \in \{\delta_H, 2\delta_H, \dots, H\}$$

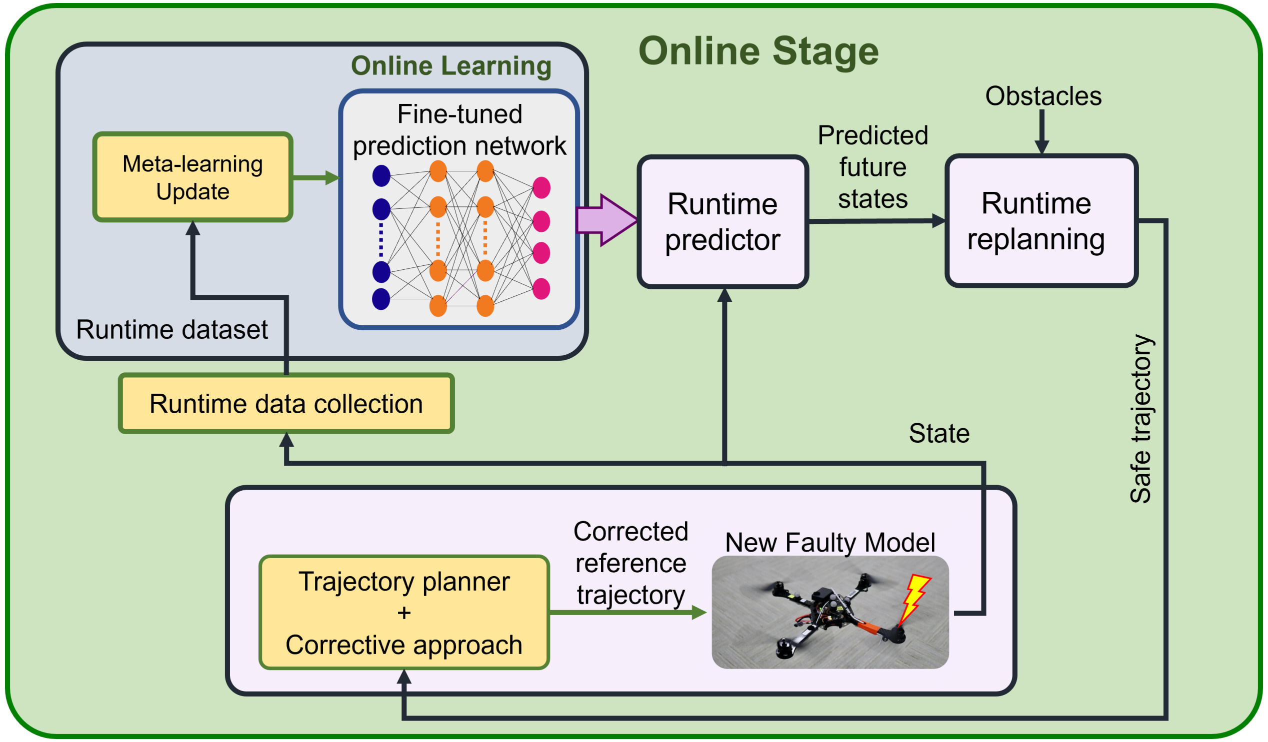

Results: Simulations

The UAV collides with the object without the meta-learning prediction and replanning approach

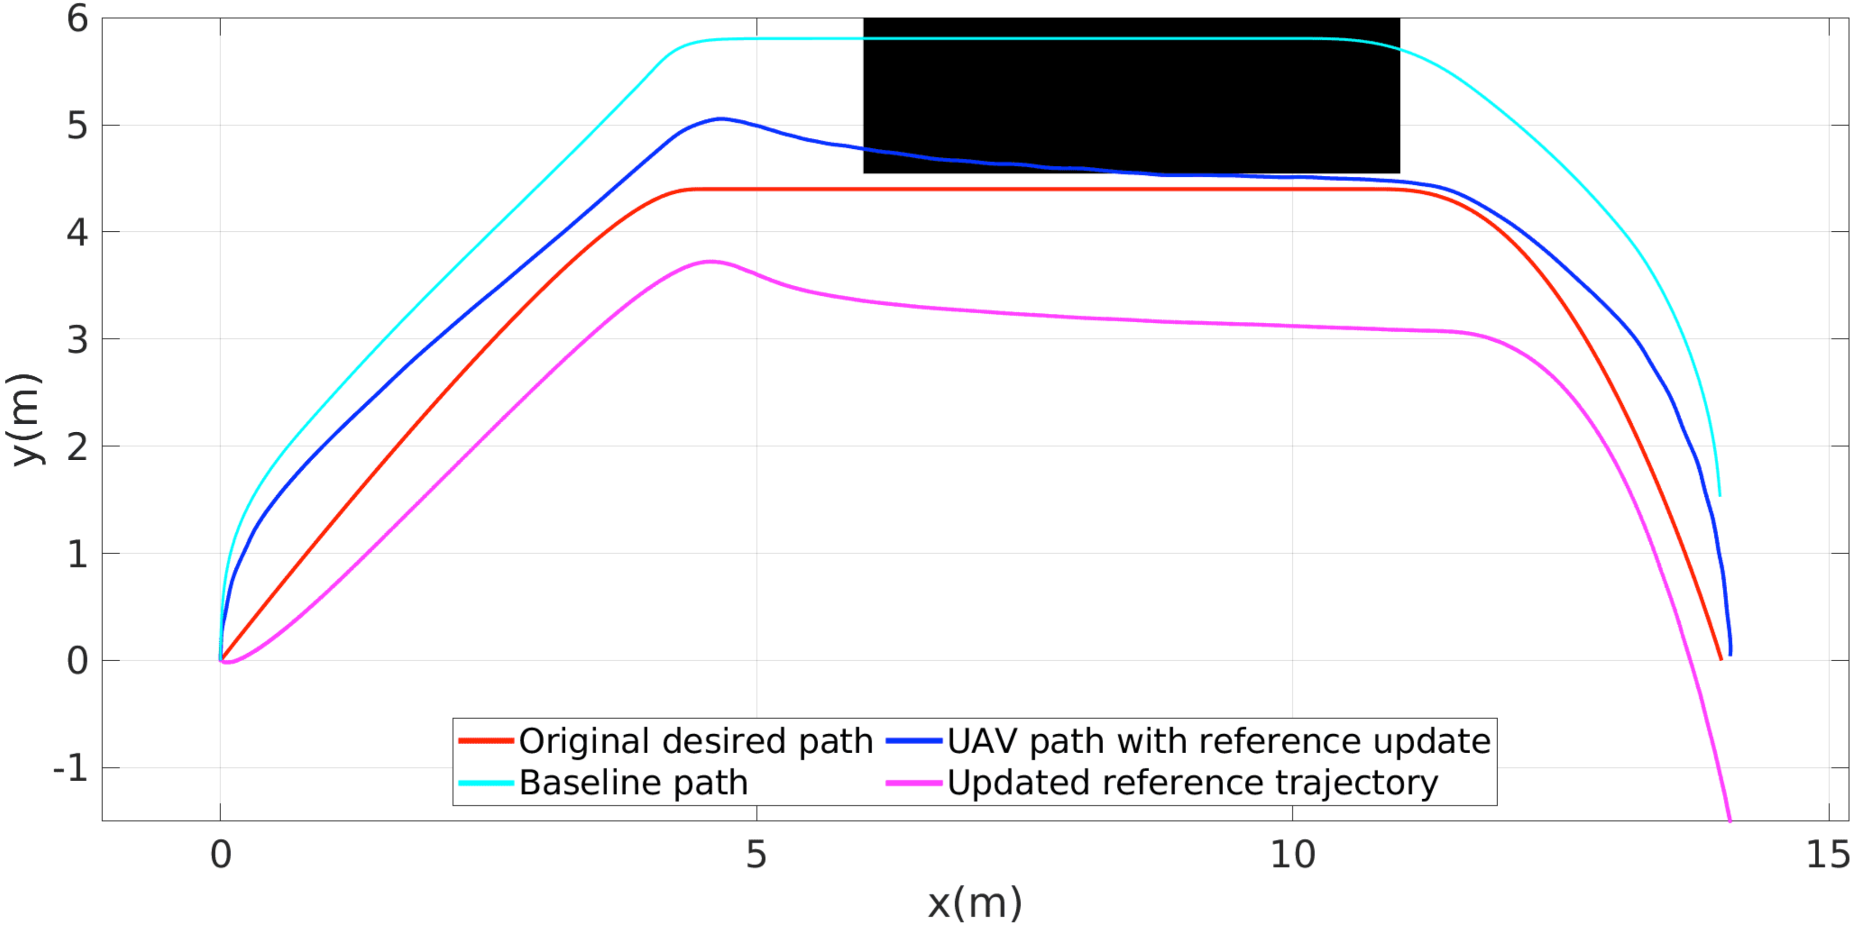

Results: Simulations

The UAV detects the possible collision and replanned for safety

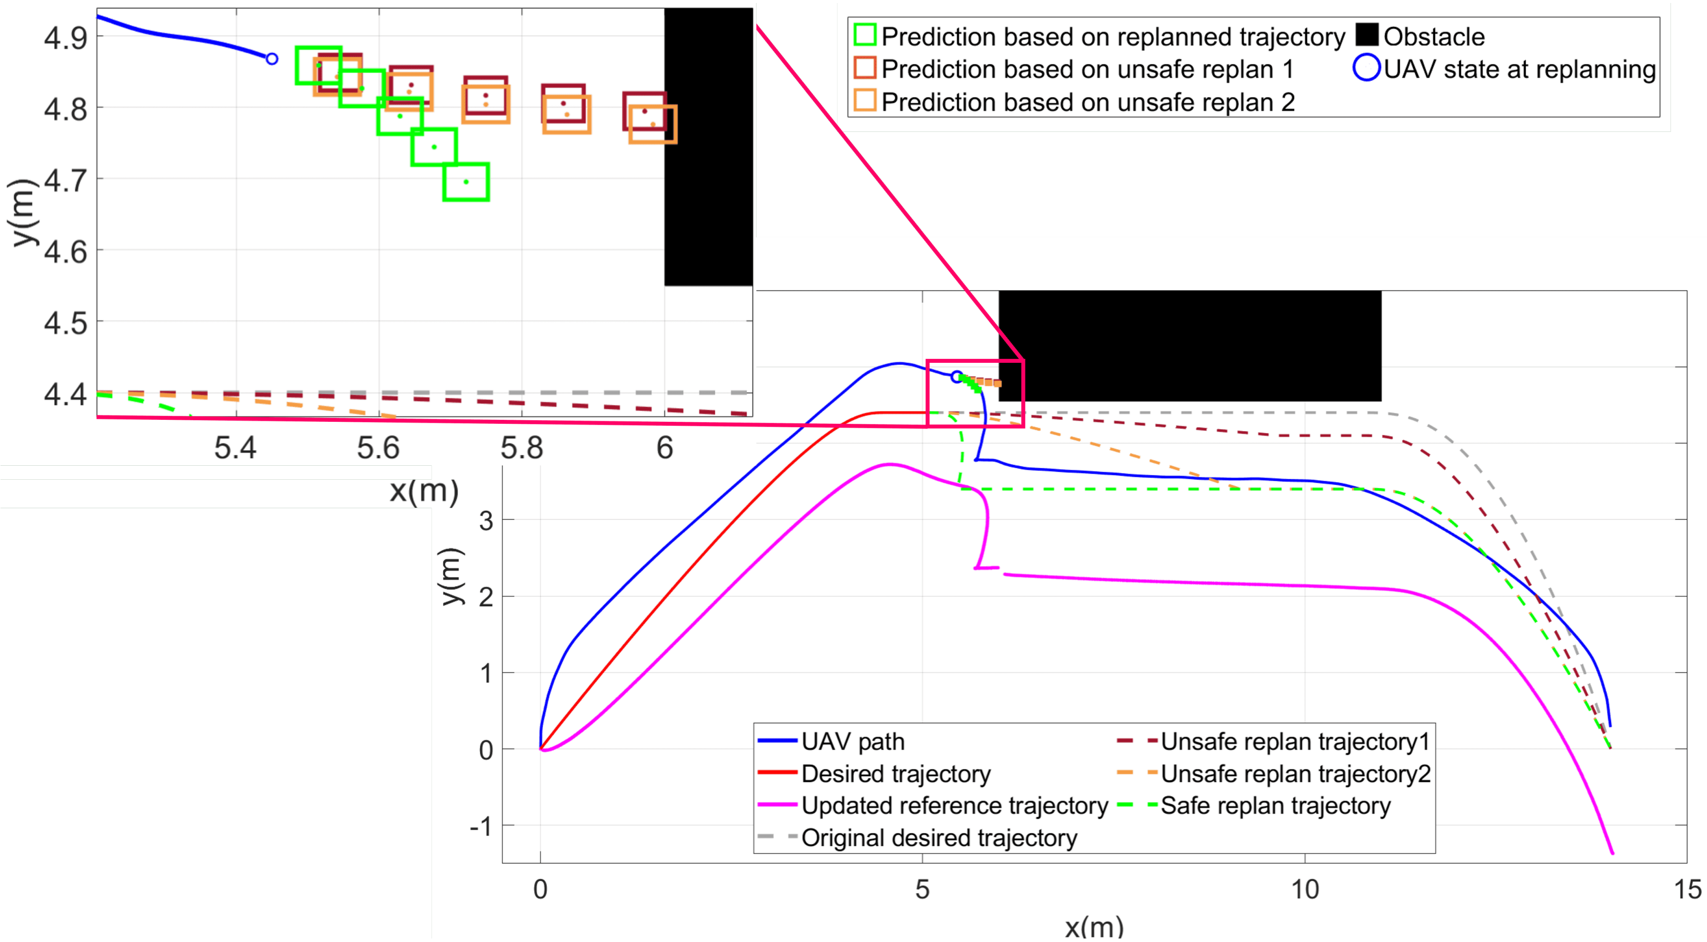

Results: Experiments

The UAV would collide with the obstacle without proactive replanning.

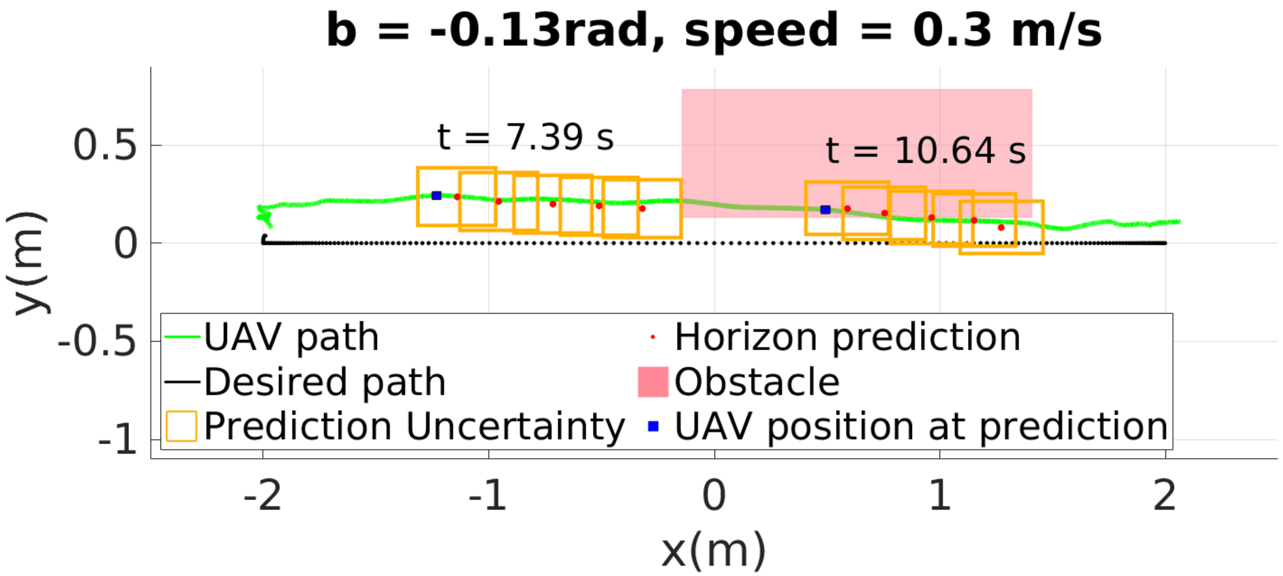

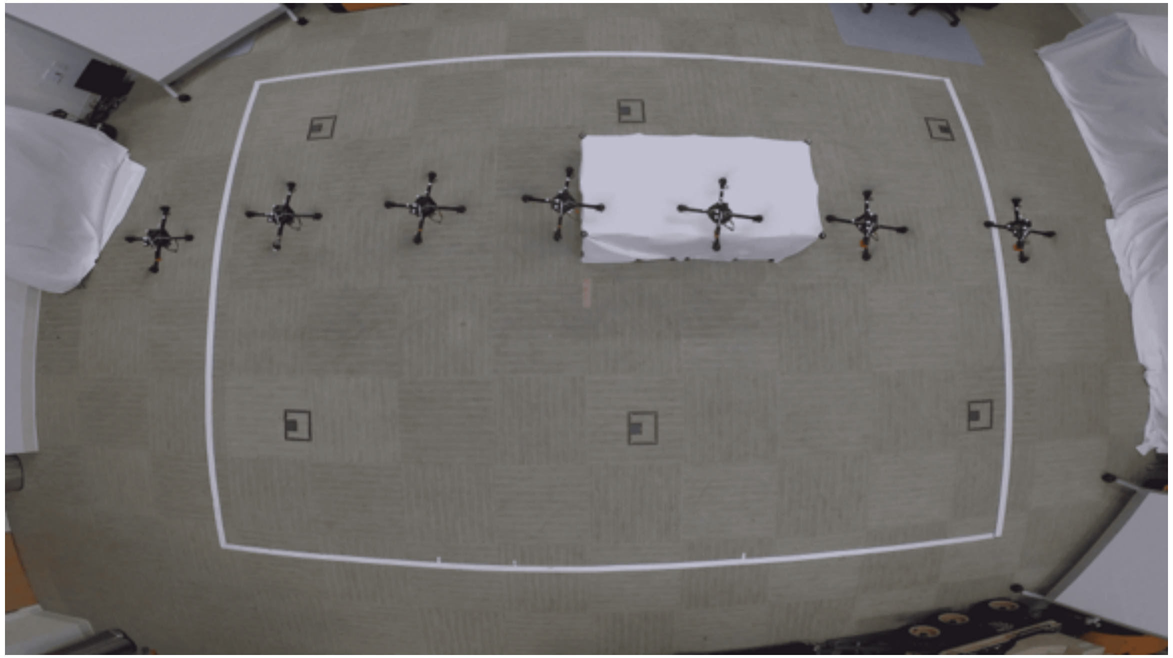

Results: Experiments

The UAV predicts the possible collision and replanned the path for safety

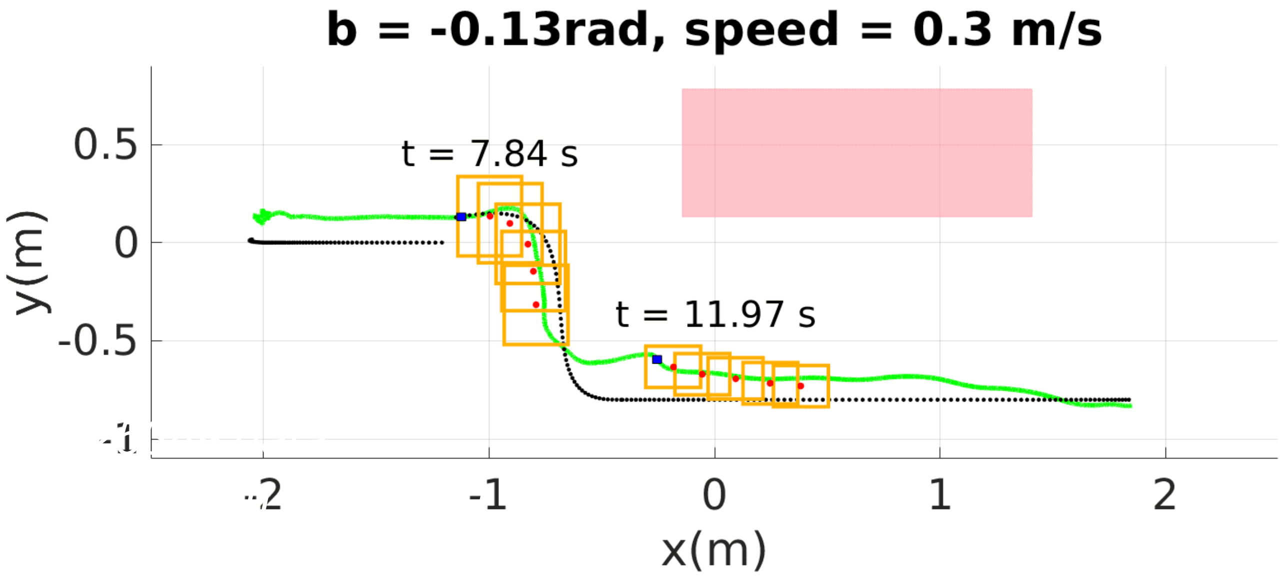

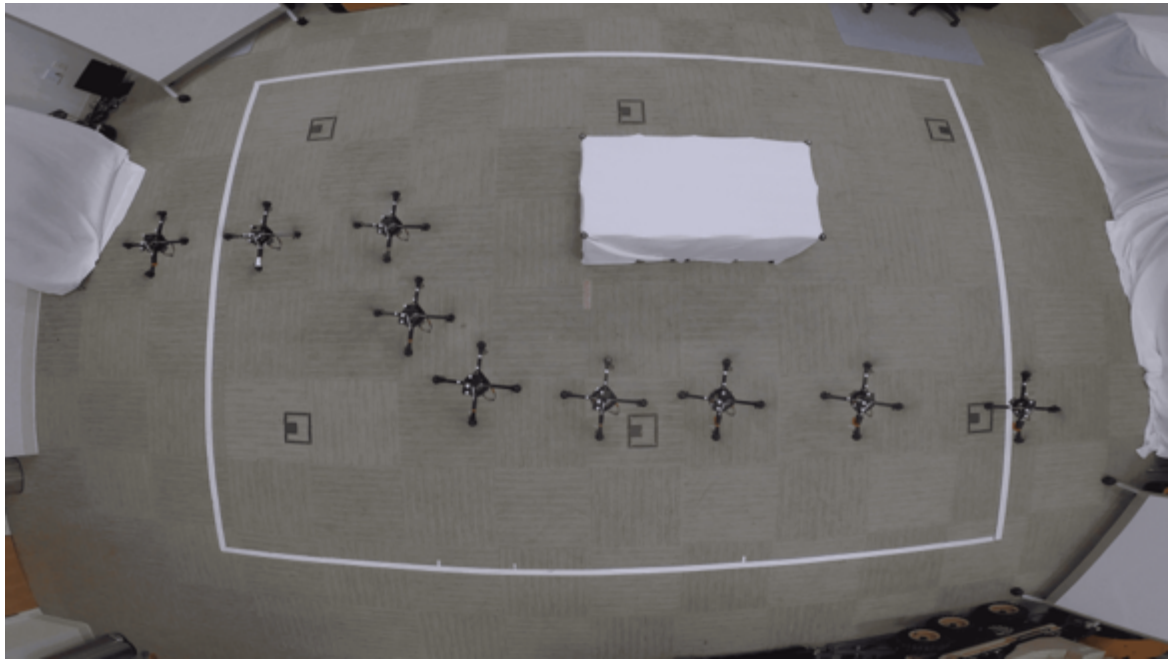

Results: Experiments

When there is no possible collision, replanning is not triggered.

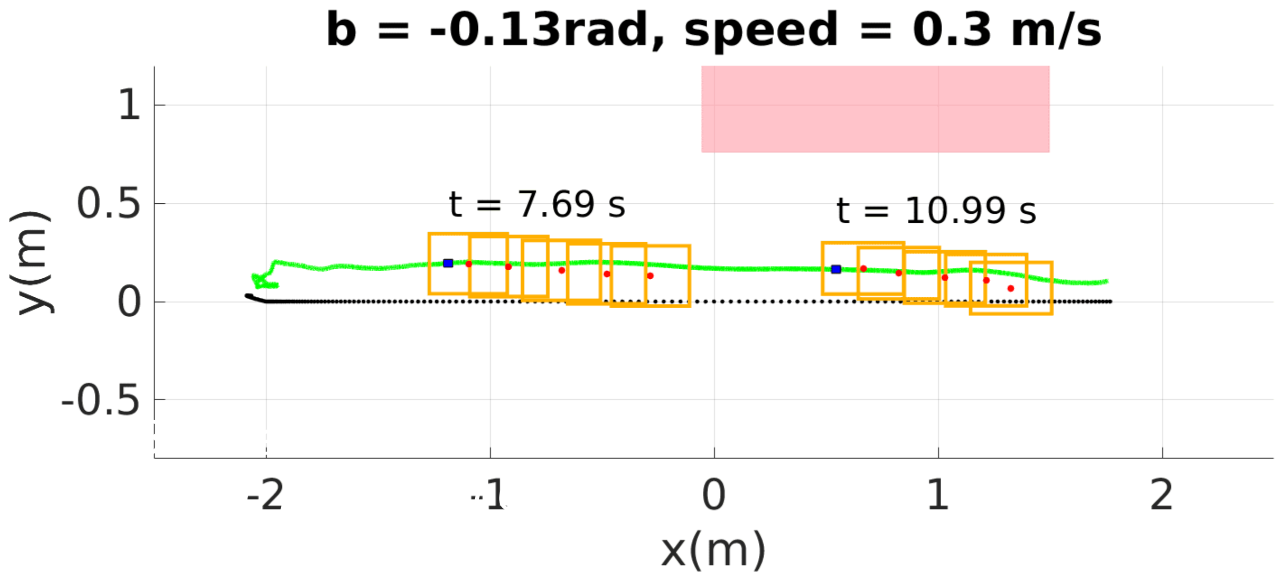

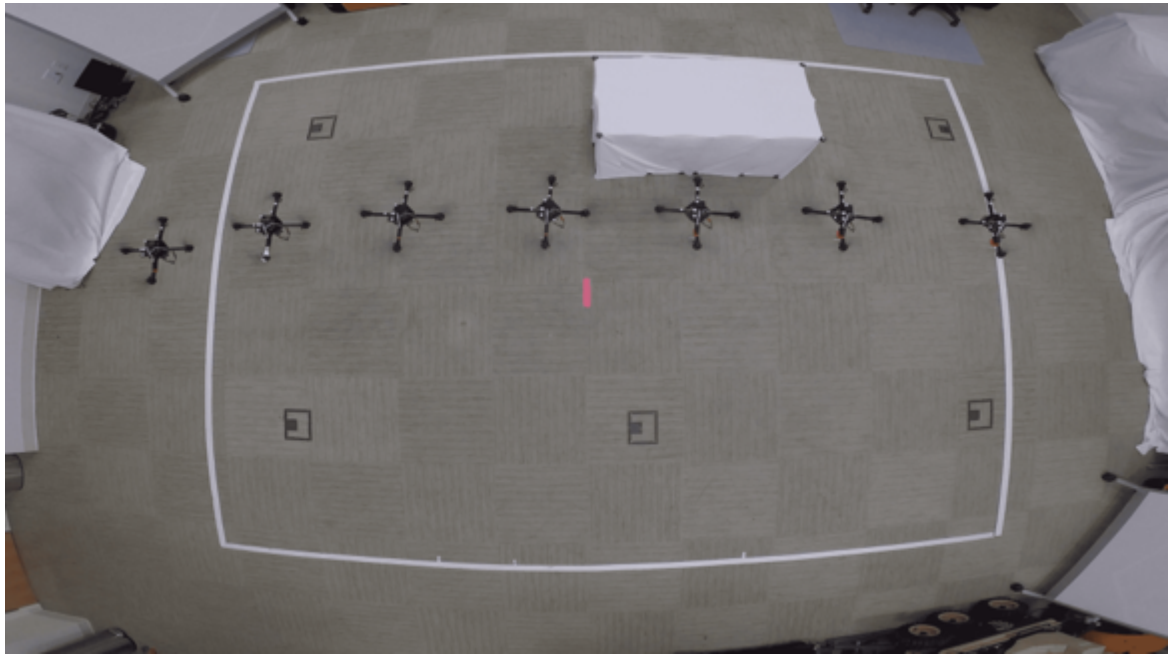

Results: Experiments

The approach is validated with a more cluttered environment to demonstrates the robustness.

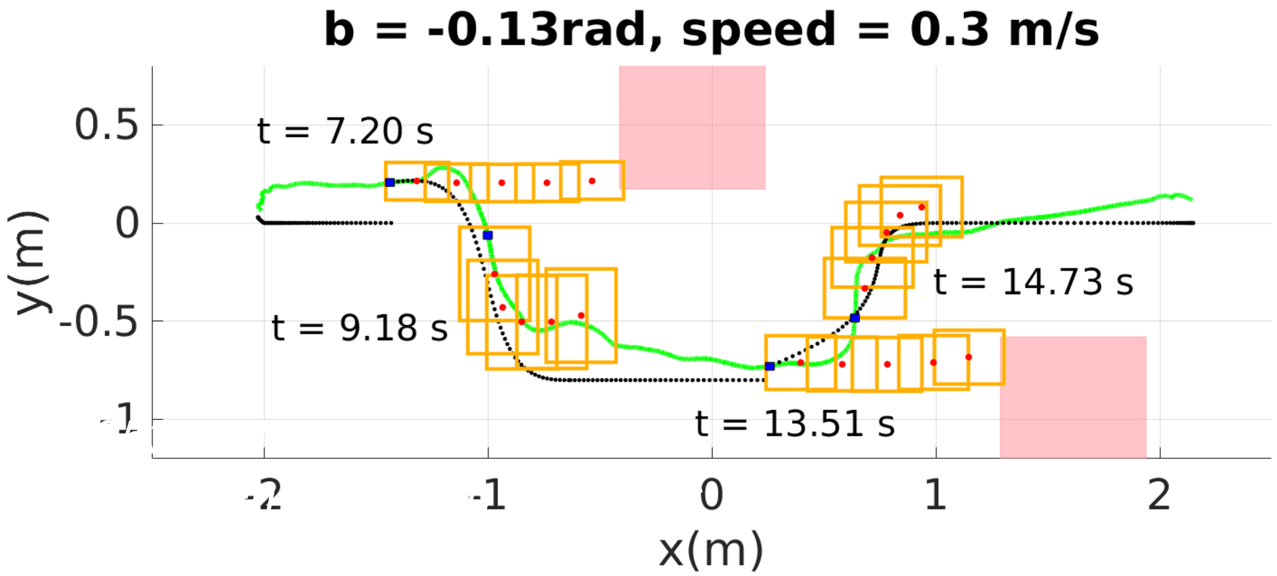

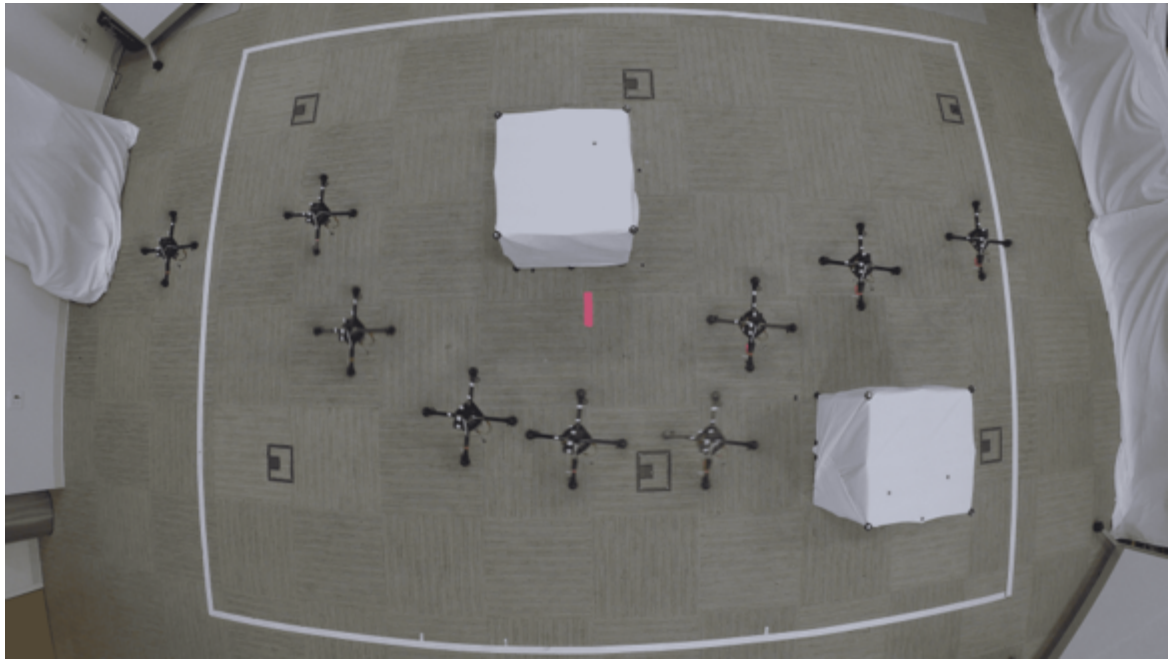

Conclusions and Future Work

Summary:

- Meta-learning for future state and state uncertainty prediction under unknown faults
- Proactive replanning to avoid collisions
- Runtime model monitoring

Current and Future Work:

- Extension to dynamic environments
- Extension to time-varying faults
- Incorporating various planning techniques to provide guarantees about finding a safe trajectory solution

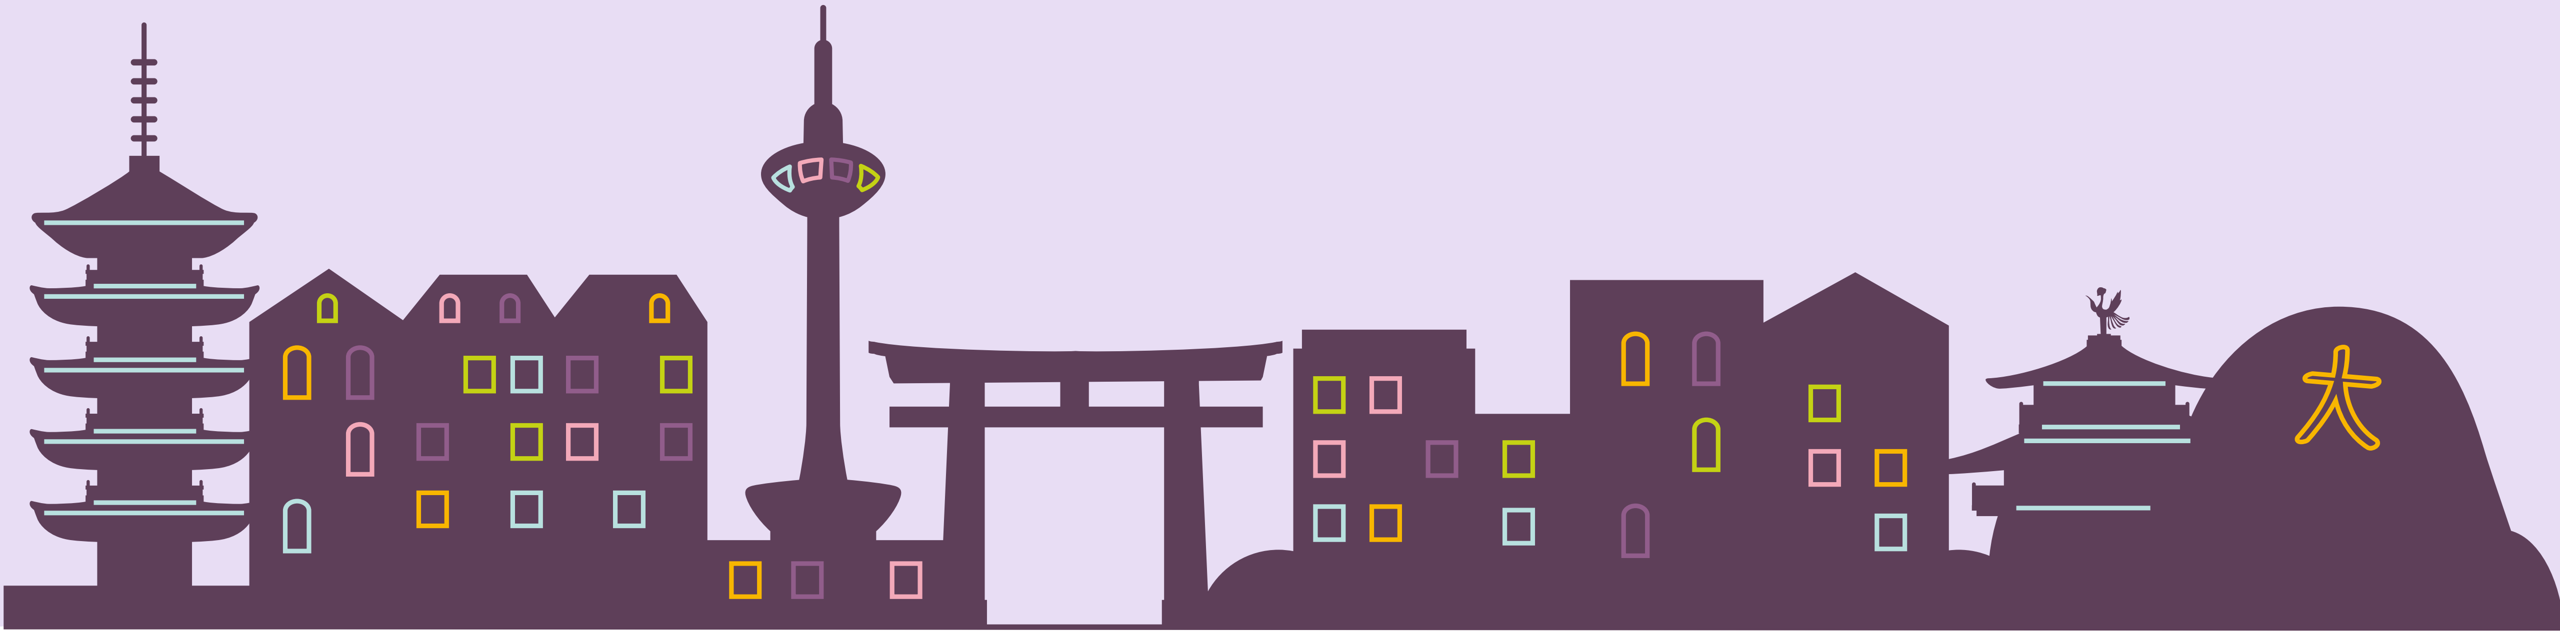

Thank you!

Esen Yel: esenyel@virginia.edu

Shijie Gao: sjgao@virginia.edu

Nicola Bezzo: nbezzo@virginia.edu

https://www.bezzorobotics.com/

UVA ENGINEERING

LINK LAB

AMR LAB

DARPA