

Aerial Robotic Autonomy

> *Uncertainty-aware Navigation using Deep Neural Networks*

Huan Nguyen, Kostas Alexis
Autonomous Robots Lab

Introduction

Robotic operations take place in a diverse set of challenging environments.

Key challenges:

- Limited communication
- Perception degradation
- Noisy state estimate/sensor observations
- Hard-to-detect obstacles




Fly through video to allow understanding of the overall environment




6x

Conventional Autonomy Stack



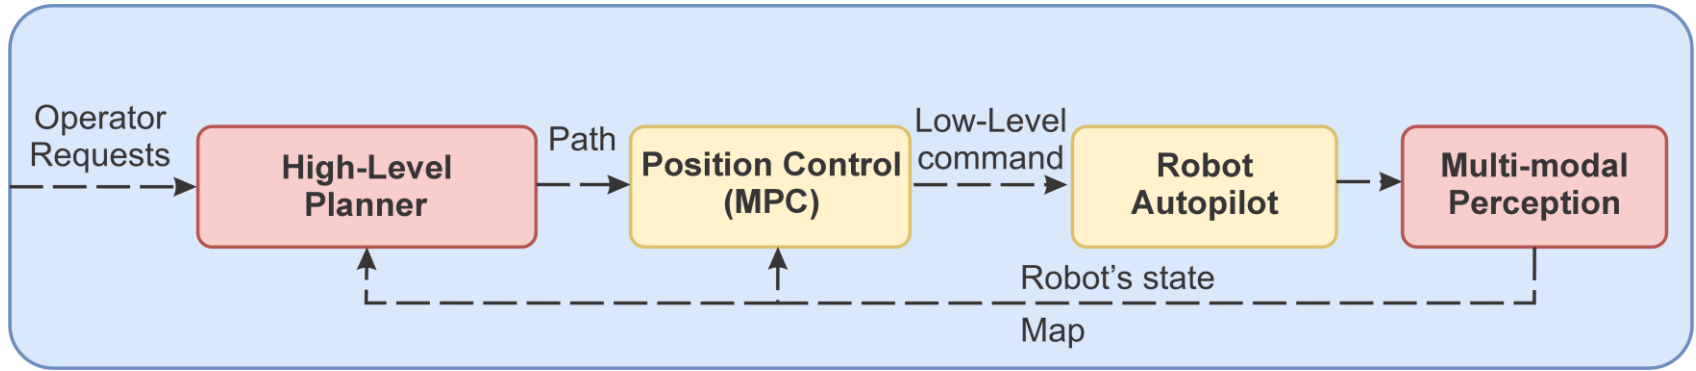
Requires position estimates and map to be accurate

Conventional Autonomy Stack



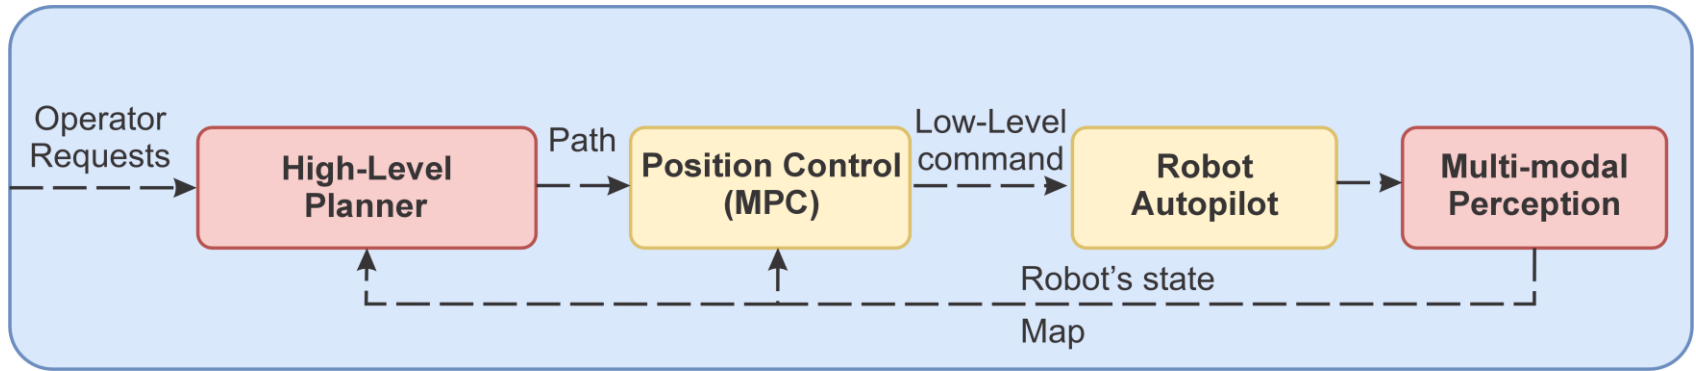
- SLAM can **fail/drift** significantly
- Mapping can require high **computational** cost

Conventional Autonomy Stack



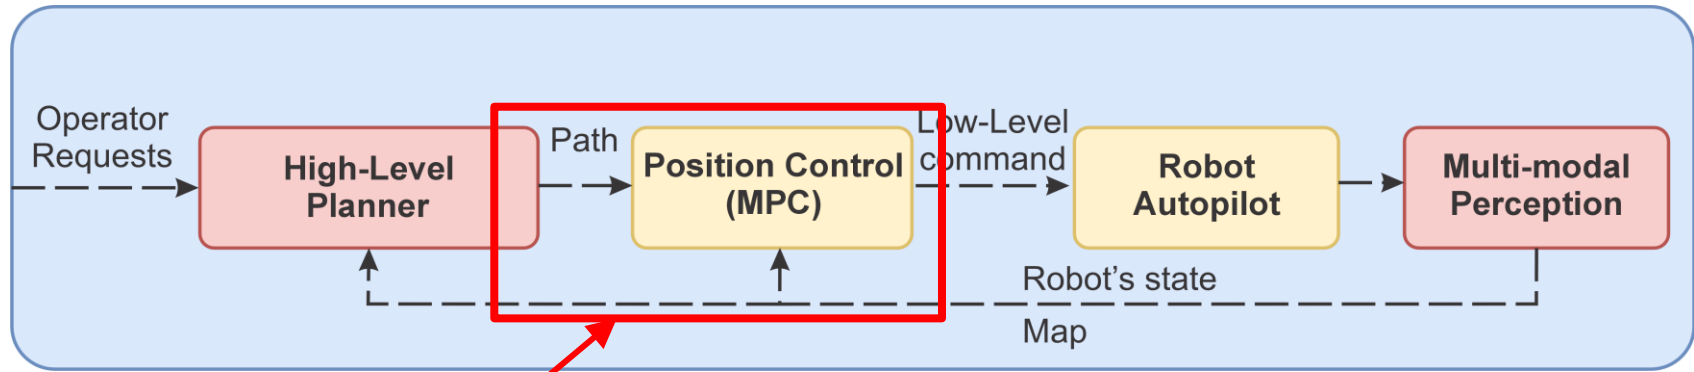
- SLAM can **fail/drift** significantly
- Mapping can require high **computational** cost → **speed limit**

Conventional Autonomy Stack



- SLAM can **fail/drift** significantly
- Mapping can require high **computational** cost → **speed limit**
- **Noisy + high-dimensional** exteroceptive data
- **Uncertain** robot's state estimate

Conventional Autonomy Stack



Resilient learning-based navigation methods

- “Last-resort” policy when SLAM fails/drifts (**Redundancy**)
- A safety layer – utilizing same inputs as mapping-planning modules (**Resourcefulness**)
- Resilient to noise (**Robustness**)

Related Work



Agile [2]

Dagger imitation learning

Require position tracking controller

Robustness to noise (up to some levels)



BADGR [3]

Score a Motion Primitive Libraries

No position feedback


No uncertainty consideration

ORACLE (ours) [1]

Score a Motion Primitive Libraries

No position feedback

Uncertainty-aware



PAC-Bayes [4]

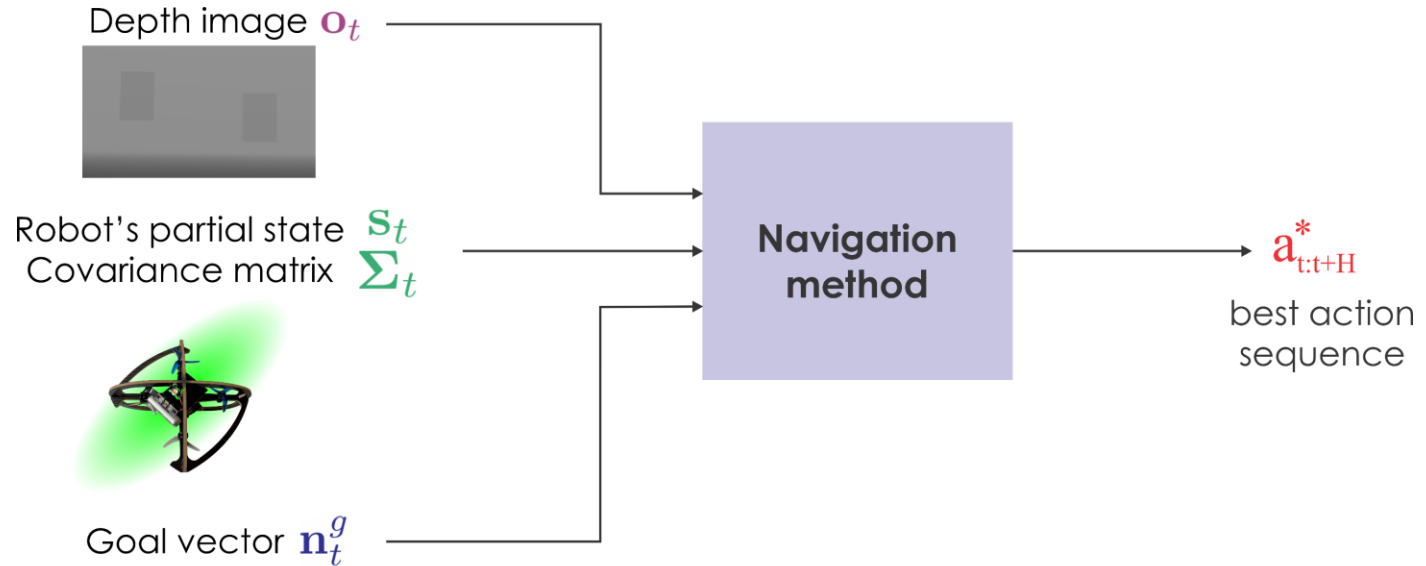
Score a Motion Primitive Libraries

Require position tracking controller

No uncertainty consideration

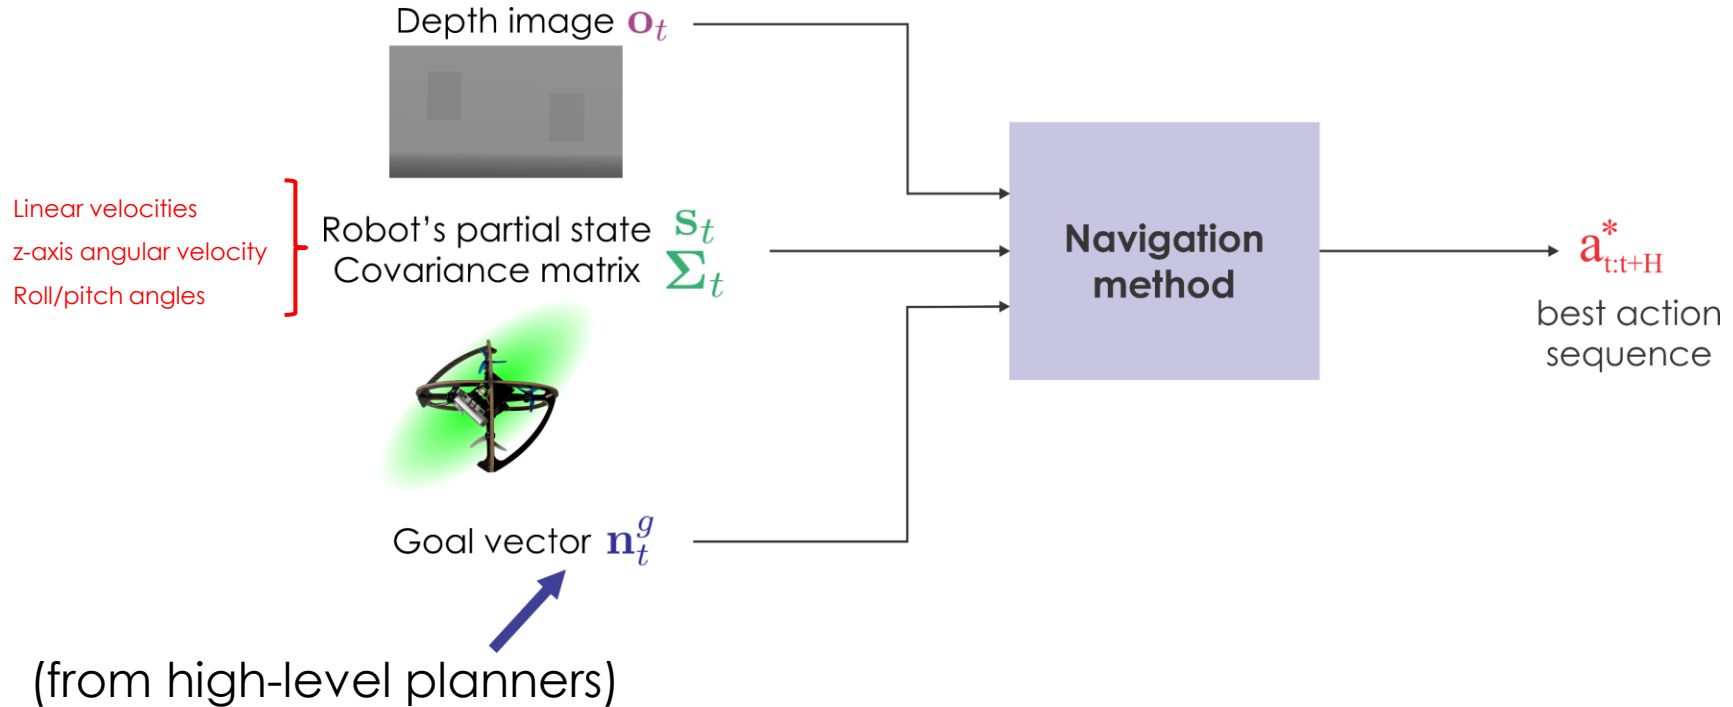
Only in simulation

Problem Formulation

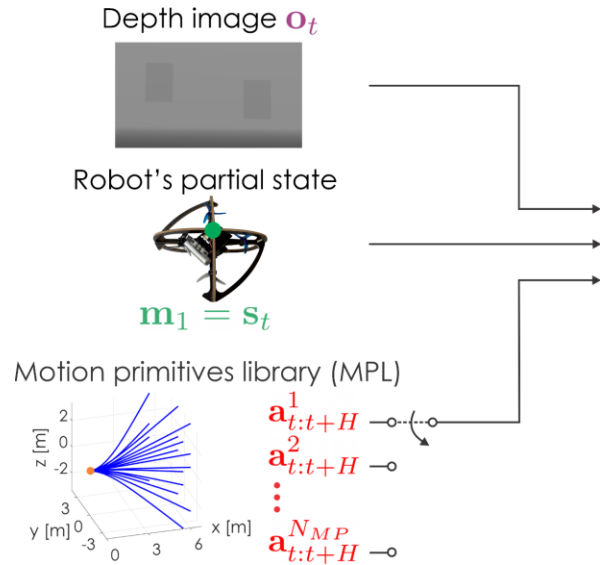


(from high-level planners)


Problem Formulation



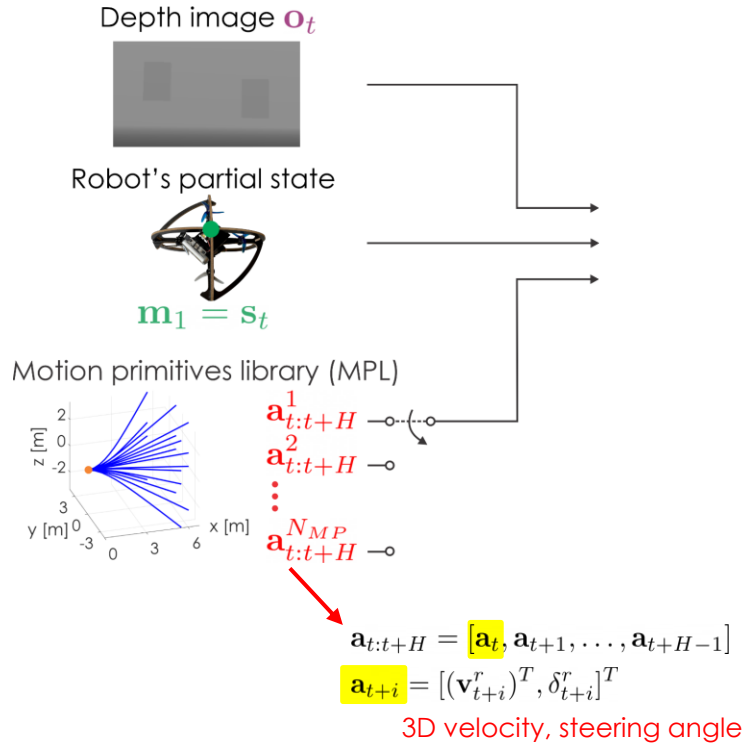
Approach




ORACLE



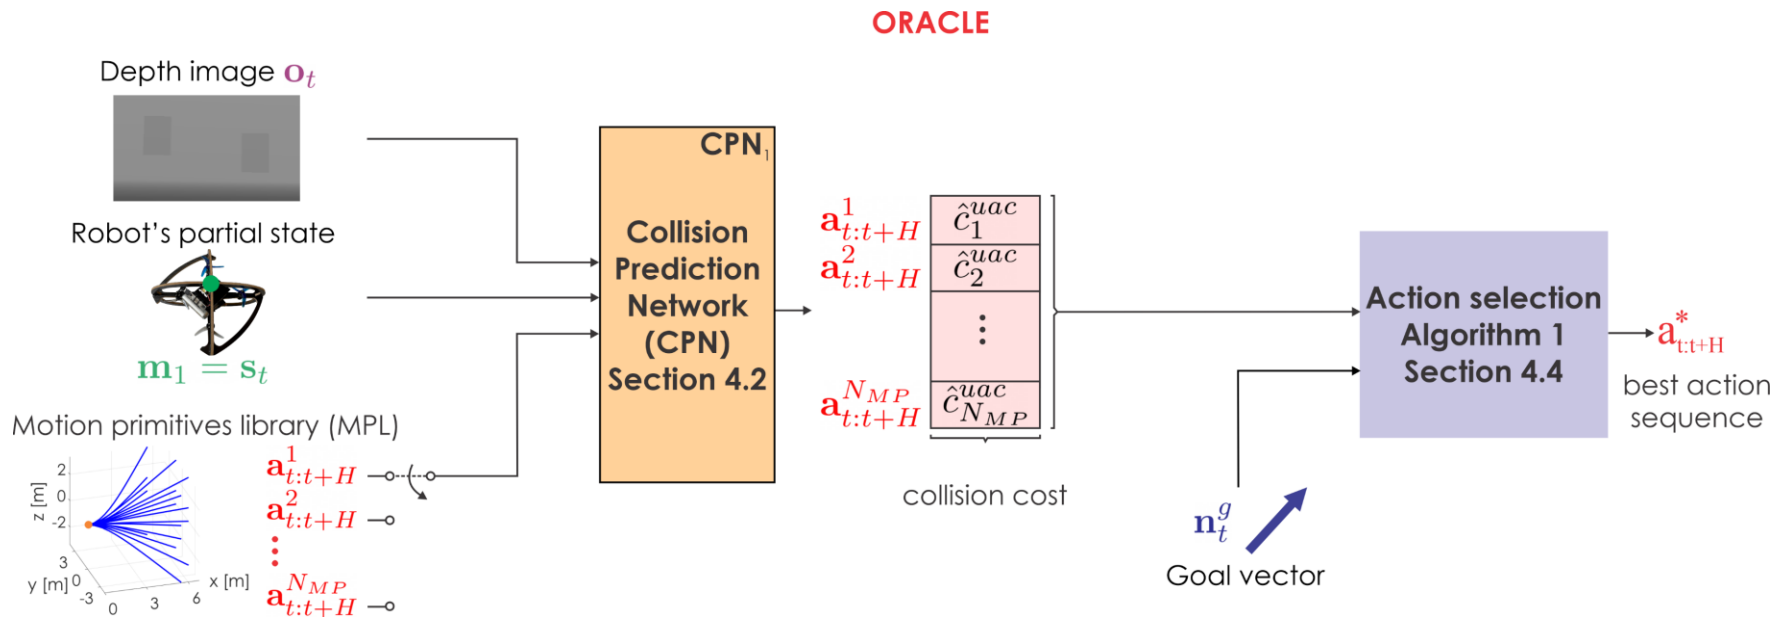
Approach



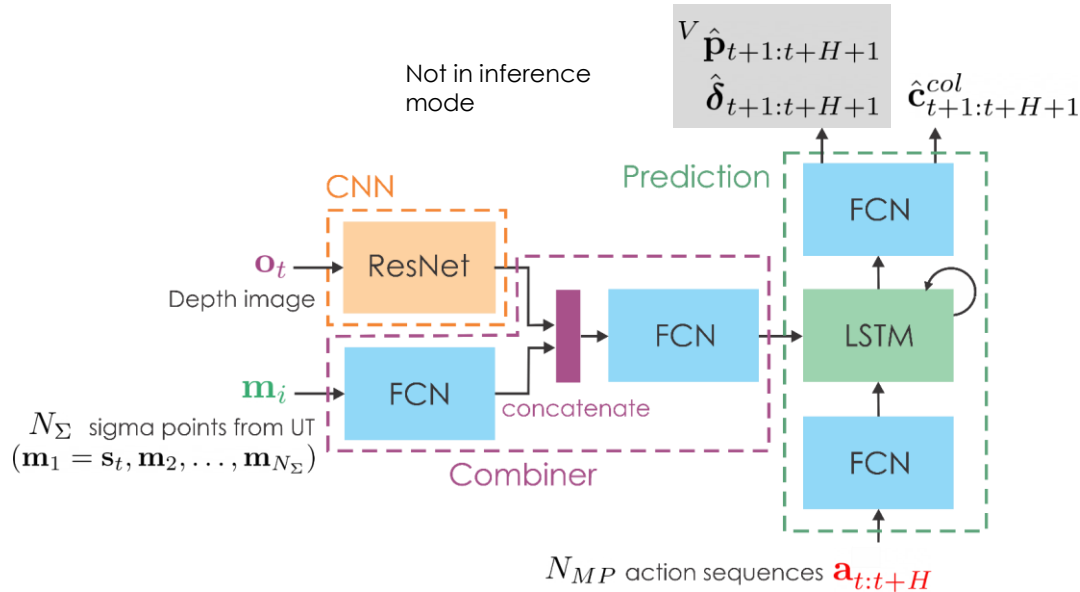
ORACLE



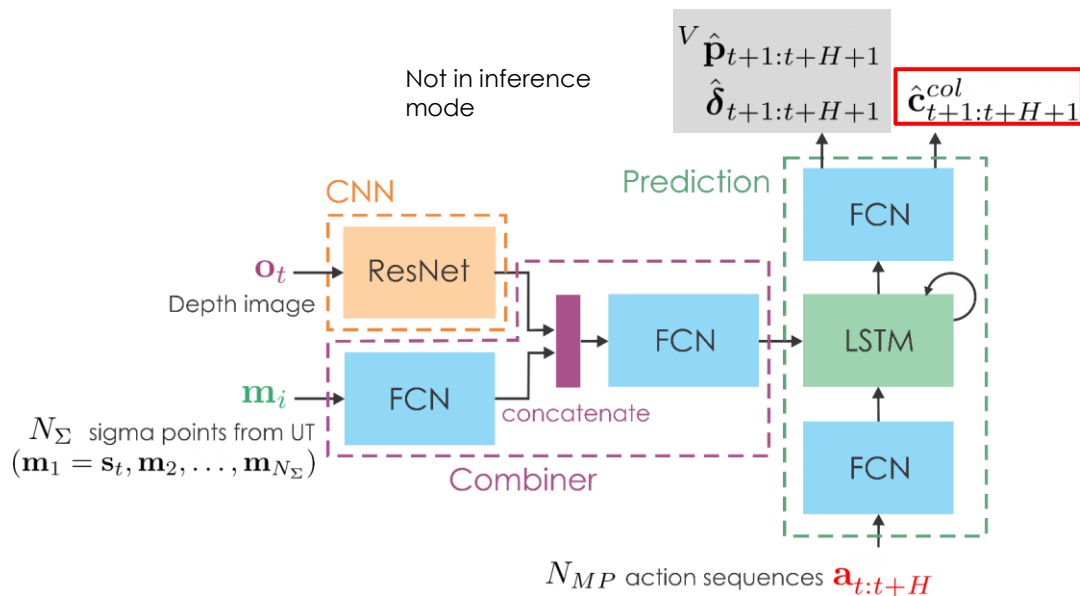
Approach



CPN Architecture & Training Losses




CPN Architecture & Training Losses



Binary Cross Entropy for collision prediction

$$L_{CPN} = \alpha_1 L_{BCE}(\mathbf{c}_{t+1:t+H+1}^{col}, \hat{\mathbf{c}}_{t+1:t+H+1}^{col}) + \alpha_2 L_{MSE}(\mathcal{V} \mathbf{p}_{t+1:t+H+1}, \mathcal{V} \hat{\mathbf{p}}_{t+1:t+H+1}) + \alpha_3 L_{MSE}(\delta_{t+1:t+H+1}, \hat{\delta}_{t+1:t+H+1})$$


CPN Architecture & Training Losses



$$L_{CPN} = \alpha_1 L_{BCE}(\mathbf{c}_{t+1:t+H+1}^{col}, \hat{\mathbf{c}}_{t+1:t+H+1}^{col}) + \alpha_2 L_{MSE}(V \mathbf{p}_{t+1:t+H+1}, V \hat{\mathbf{p}}_{t+1:t+H+1}) + \alpha_3 L_{MSE}(\delta_{t+1:t+H+1}, \hat{\delta}_{t+1:t+H+1})$$

Mean Squared Error for position and yaw predictions

Data Collection

- Roll out action sequences in RotorS simulator 
- Match dynamics with the real robot
- Format of one datapoint


$$\mathbf{d}_{in} = (\mathbf{o}_t, \mathbf{s}_t, \mathbf{a}_{t:t+H})$$

$$\mathbf{d}_{out} = (\mathbf{c}_{t+1:t+H+1}^{col}, \mathbf{p}_{t+1:t+H+1}^V, \delta_{t+1:t+H+1})$$


collision
status
0: non-
collision
1: collision

relative
position

relative
yaw

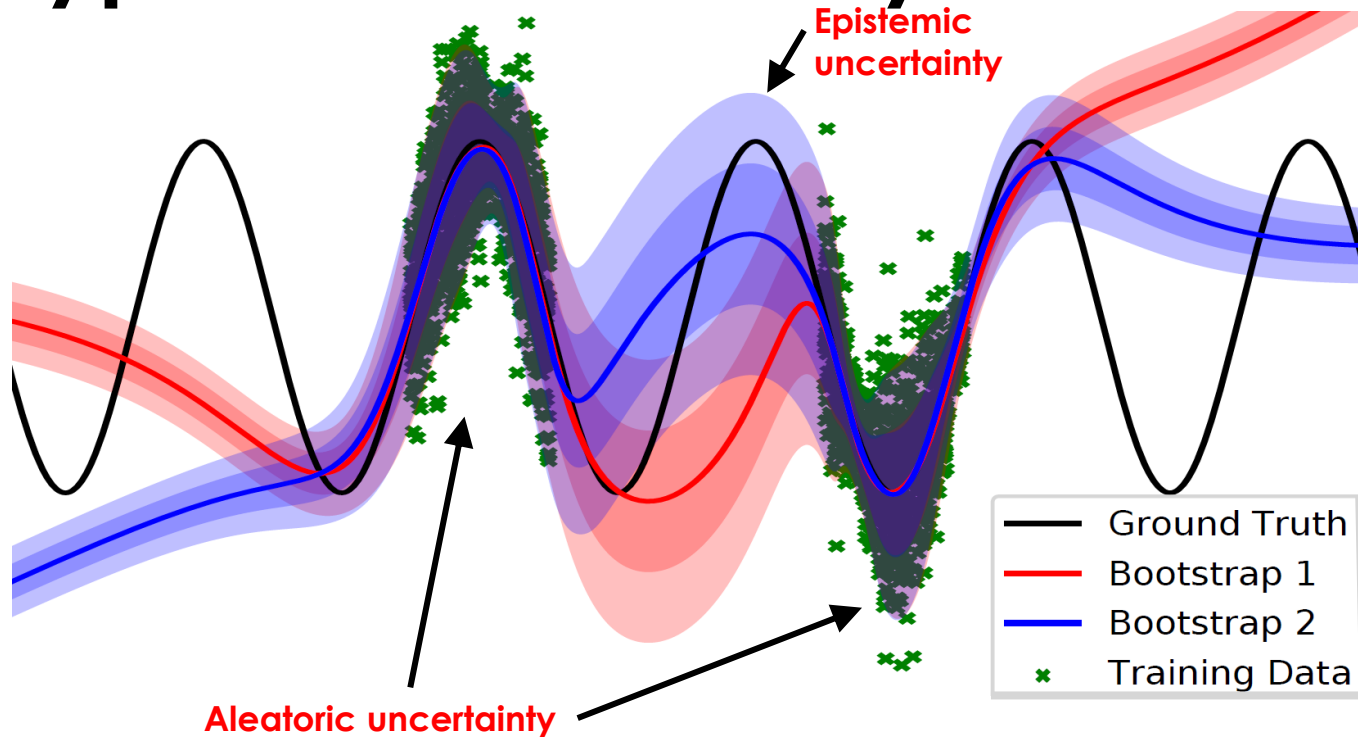


Randomized environments



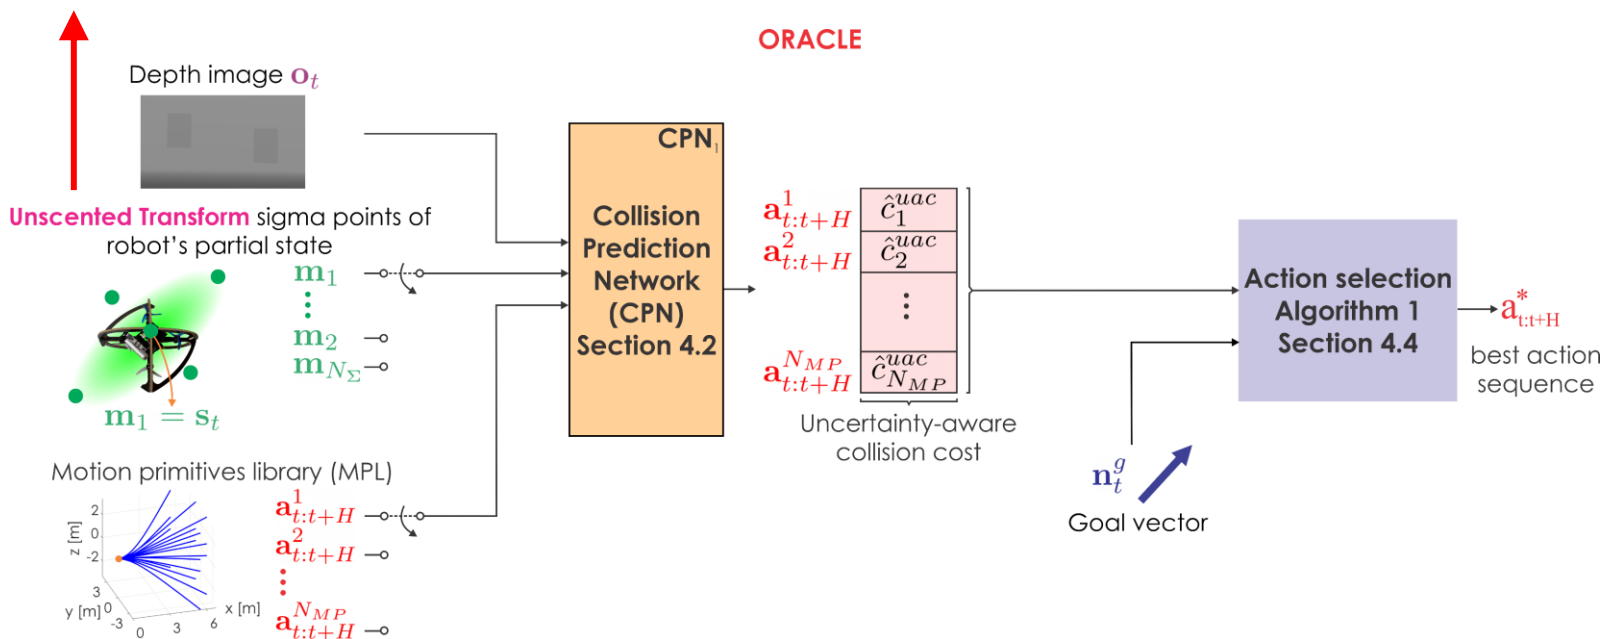
Onboard camera when random
action sequences are rolled out

Two types of uncertainty in DL



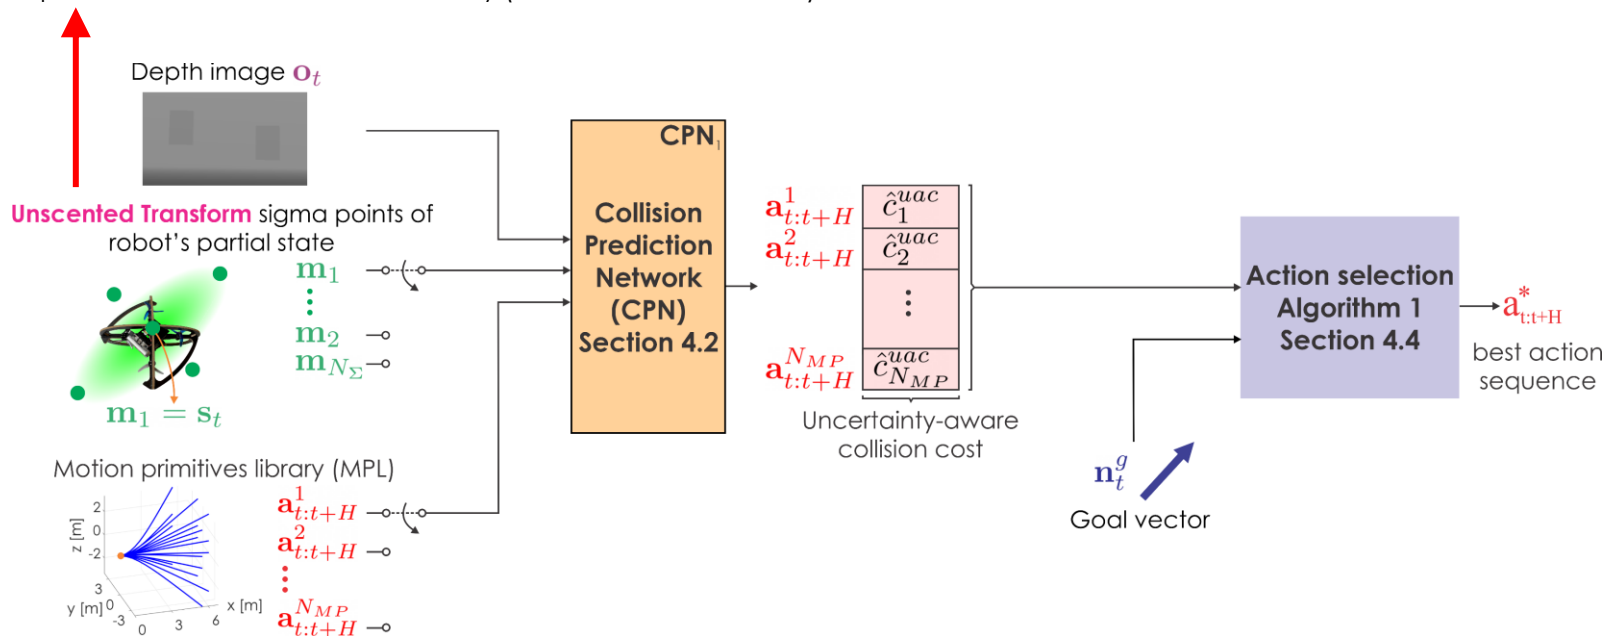
Uncertainty-Aware Safe Navigation

- **Aleatoric uncertainty:** prediction uncertainty due to the noise in input data
- Cannot be reduced with more training data



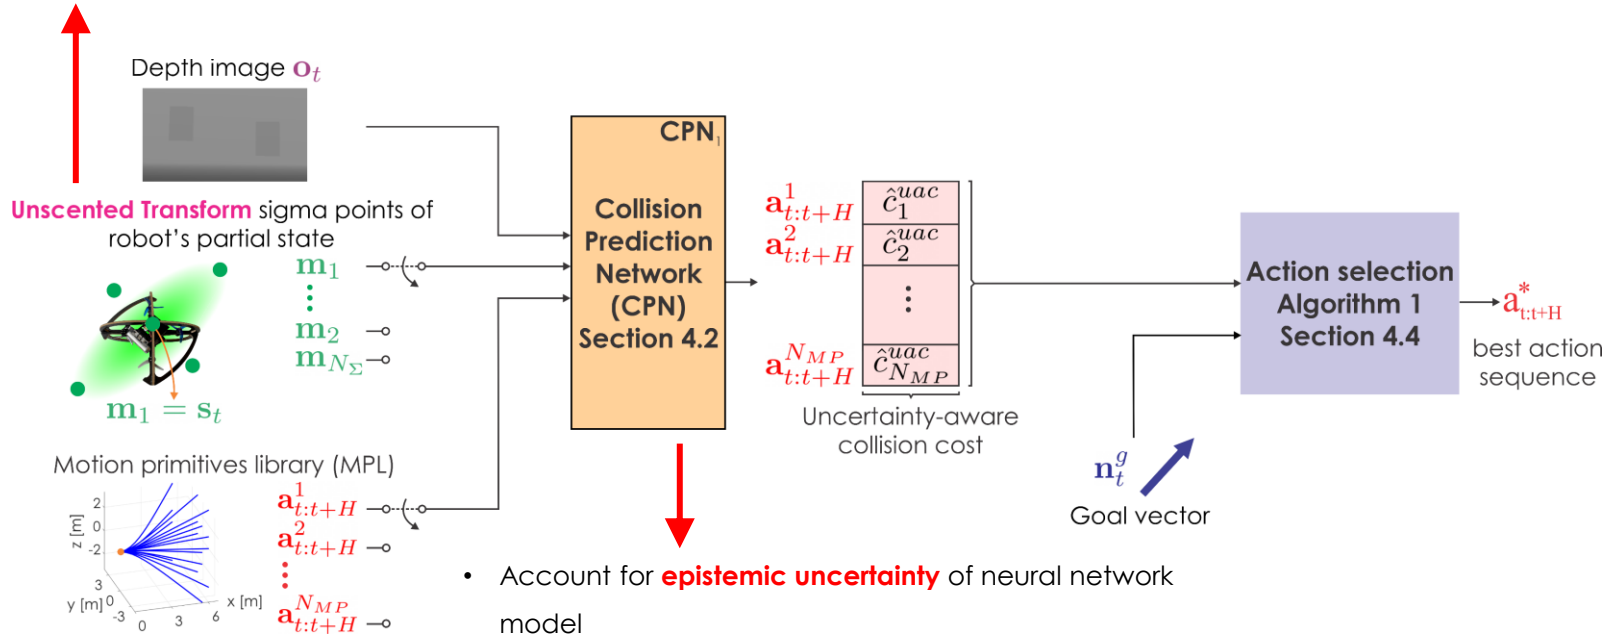
Uncertainty-Aware Safe Navigation

- **Aleatoric uncertainty:** prediction uncertainty due to the noise in input data
- Cannot be reduced with more training data
- Adapt to current estimation uncertainty (vs train with fixed noise)



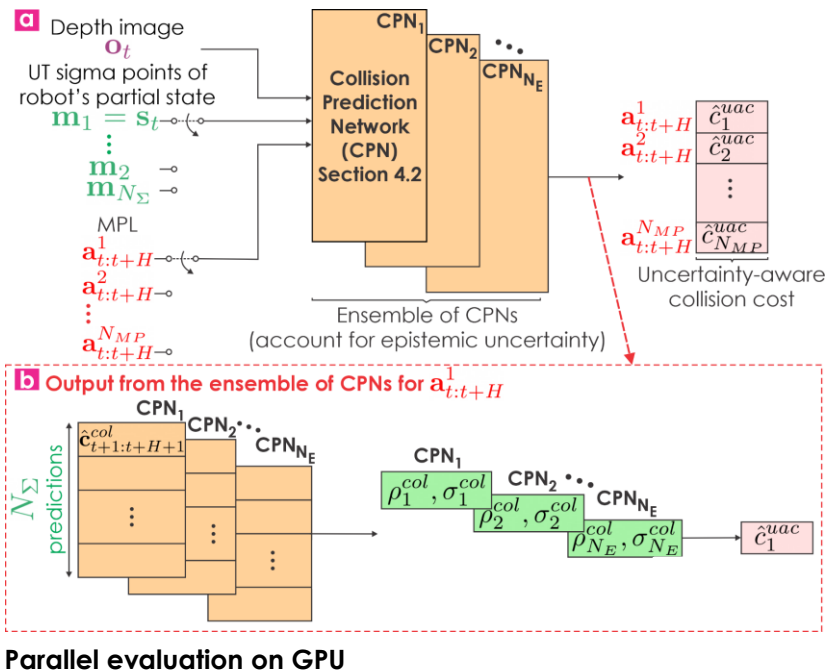
Uncertainty-Aware Safe Navigation

- **Aleatoric uncertainty**: prediction uncertainty due to the noise in input data
- Cannot be reduced with more training data
- Adapt to current estimation uncertainty (vs train with fixed noise)

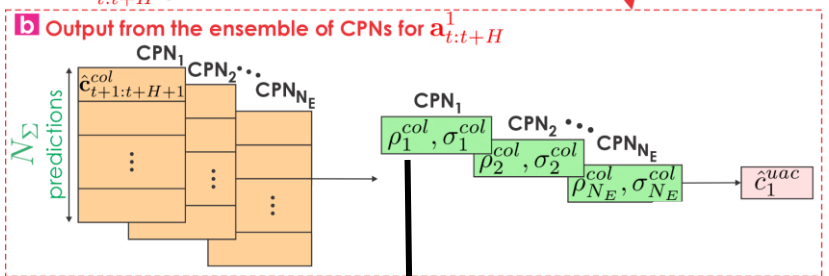



- Account for **epistemic uncertainty** of neural network model
- Can not be ignored for unseen noisy input data

Get Uncertainty-aware Collision Costs





Get Uncertainty-aware Collision Costs



Parallel evaluation on GPU

$$\hat{c}^{col} = \sum_{i=1}^H \hat{c}_{t+i}^{col} e^{-\lambda(i-1)}, \lambda > 0 \quad \text{Unscented Transform}$$



Get Uncertainty-aware Collision Costs



Parallel evaluation on GPU

$$\left[\begin{aligned} \sigma_{tot}^{col} &= \frac{1}{N_E} \sum_{n=1}^{N_E} [\sigma_n^{col} + (\rho_n^{col} - \bar{\rho}^{col})^2] \\ \hat{c}^{uac} &= \bar{\rho}^{col} + \alpha \sqrt{\sigma_{tot}^{col}}, \alpha > 0 \end{aligned} \right.$$

Get Uncertainty-aware Collision Costs

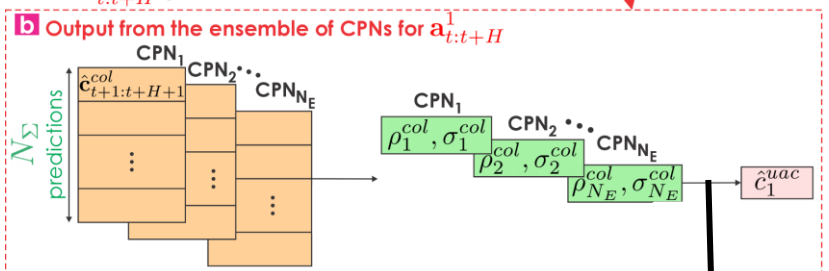
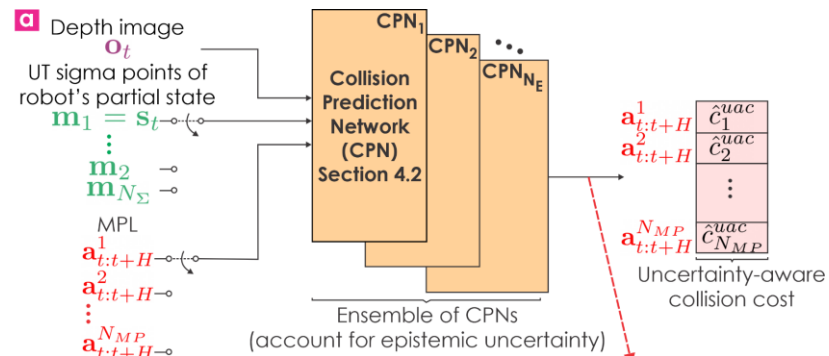


Mean of variance

$$\sigma_{tot}^{col} = \frac{1}{N_E} \sum_{n=1}^{N_E} [\sigma_n^{col} + (\rho_n^{col} - \bar{\rho}^{col})^2]$$

$$\hat{c}^{uac} = \bar{\rho}^{col} + \alpha \sqrt{\sigma_{tot}^{col}}, \alpha > 0$$

Get Uncertainty-aware Collision Costs



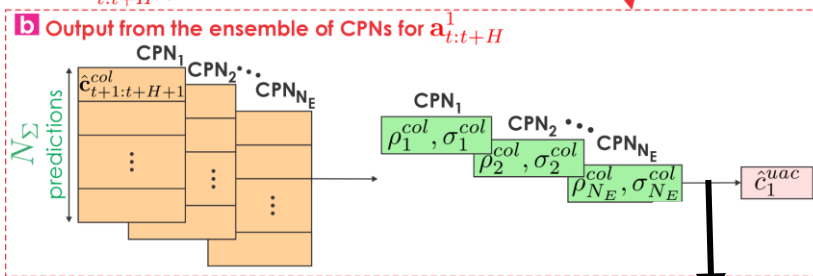
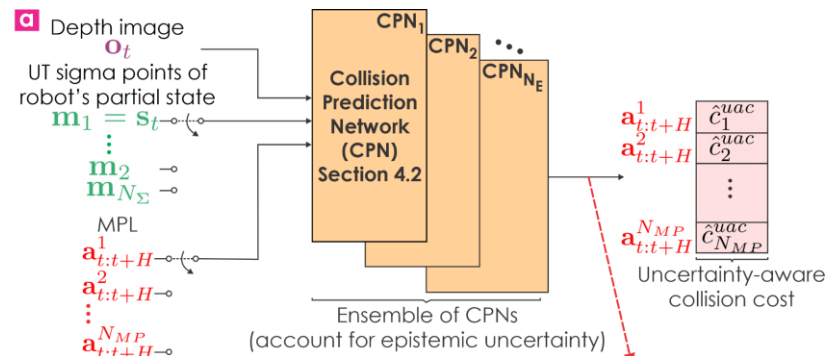
Parallel evaluation on GPU

Variance of mean

$$\sigma_{tot}^{col} = \frac{1}{N_E} \sum_{n=1}^{N_E} [\sigma_n^{col} + (\rho_n^{col} - \bar{\rho}^{col})^2]$$

$$\hat{c}^{uac} = \bar{\rho}^{col} + \alpha \sqrt{\sigma_{tot}^{col}}, \alpha > 0$$

Get Uncertainty-aware Collision Costs




Parallel evaluation on GPU

$$\left[\begin{aligned} \sigma_{tot}^{col} &= \frac{1}{N_E} \sum_{n=1}^{N_E} [\sigma_n^{col} + (\rho_n^{col} - \bar{\rho}^{col})^2] \\ \hat{c}^{uac} &= \bar{\rho}^{col} + \alpha \sqrt{\sigma_{tot}^{col}}, \alpha > 0 \end{aligned} \right.$$


Upper Confidence Bound

Simulation Results (Gazebo)



Simulated environments:

- Noisy velocity and/or depth image
- Poisson disc sampling

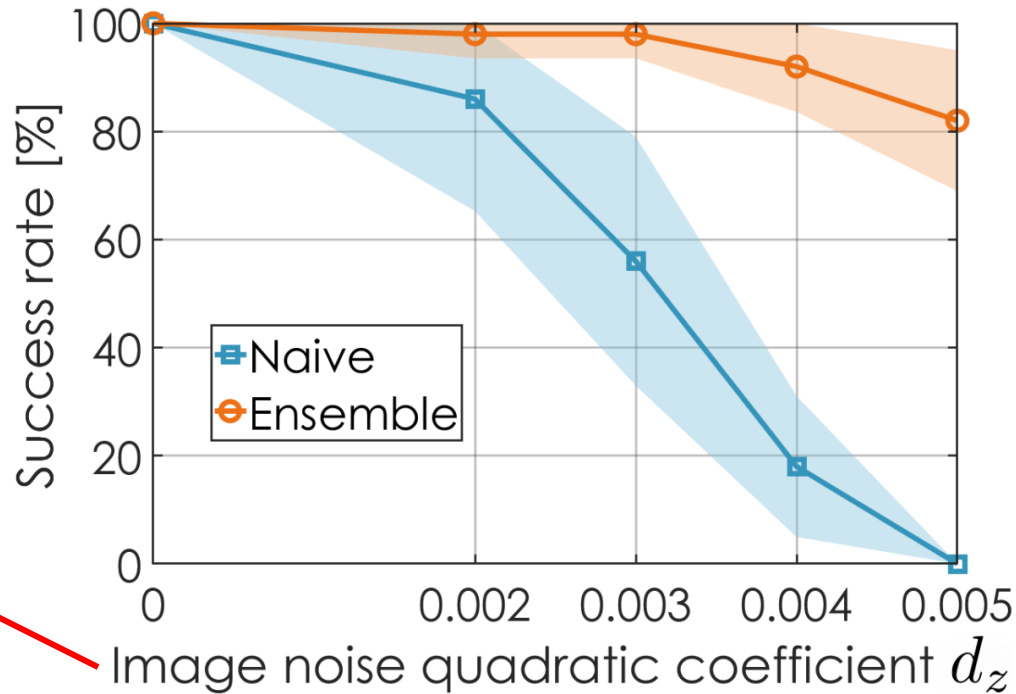


Green: safe action sequences


Blue: chosen action sequence

Methods	Naive	Ensemble	ORACLE
UT	No	No	Yes
Ensemble	No	Yes	Yes

Simulation Results (Gazebo)



Simulation Results (Gazebo)



Intel Realsense D435

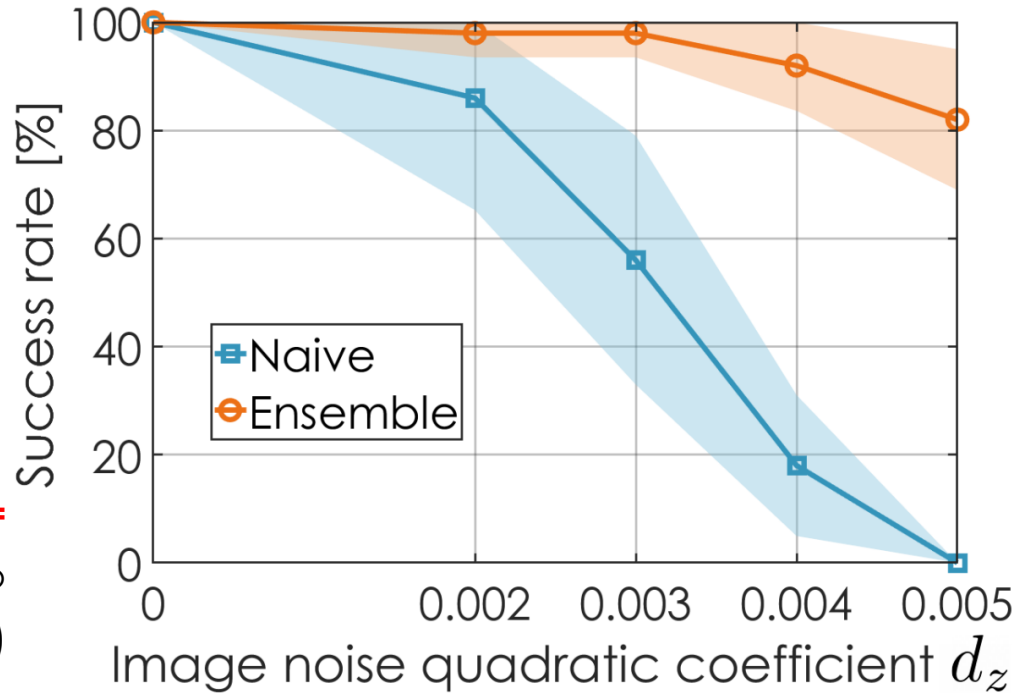
$$z_{noisy} = z + \mathcal{N}(0, \sigma_z(z))$$

$$\sigma_z(z) = d_z z^2$$




Image noise quadratic coefficient d_z

Simulation Results (Gazebo)

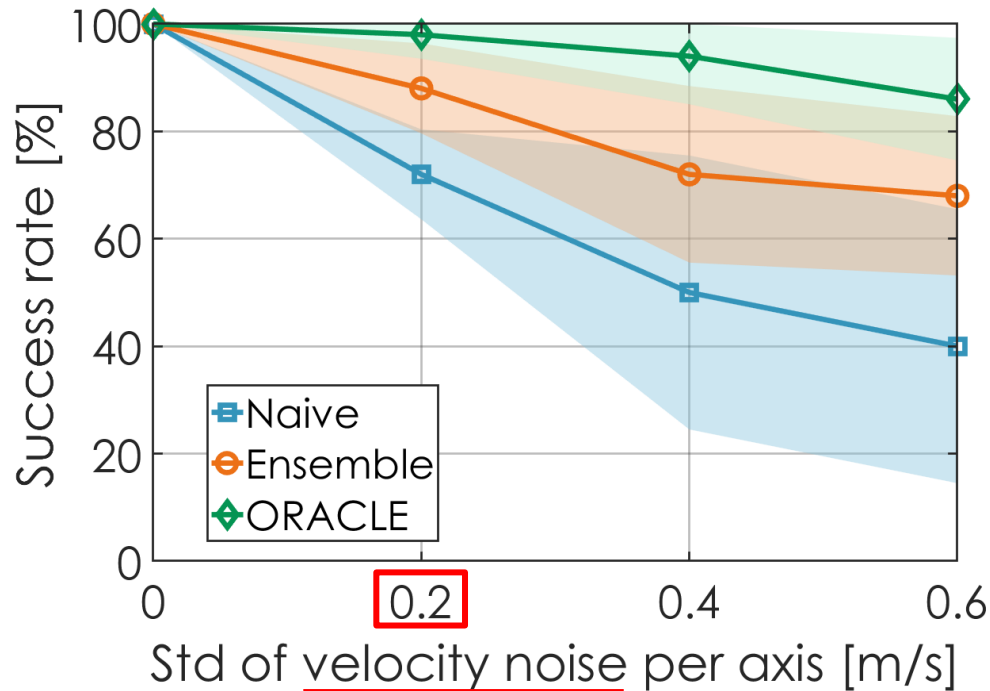


- **ORACLE** = (no velocity noise)
Ensemble outperforms **Naive**

Simulation Results (Gazebo)




Simulation Results (Gazebo)



Velocity noise in y-axis (main reactive axis) simulated in Agile


Simulation Results (Gazebo)



- **ORACLE** outperforms **Naïve** and **Ensemble** baselines
- **ORACLE** outperforms **Naïve** and **Ensemble** baselines
- **Ensemble** performs reasonably well without UT. **Why?**

Velocity noise in y-axis (main reactive axis) simulated in Agile

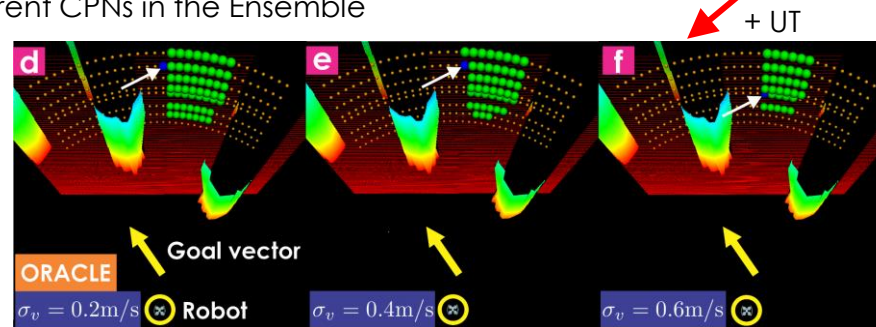
Simulation Results (Gazebo)



Green: safe action sequences

Blue with arrow: chosen action sequence


Predictions from 2 different CPNs in the Ensemble



- Using the Ensemble leads to **more conservative** set of **safe** action sequences
- Using the Ensemble + UT **enhances conservativeness**

Simulation Results (Gazebo)

- **ORACLE** outperforms **Naive** and **Ensemble** baselines



Simulation Results (Flightmare)

- **ORACLE** vs **Naïve**, **Ensemble**, **Agile** (pretrained weights)
- ORACLE and its variants haven't seen this type of environment in training



Onboard RGB images in an indicative mission

Simulation Results (Flightmare)

Noise levels	Metrics	Naive	Ensemble	ORACLE	Agile
No noise	Success rate [%]	65 (31)	80 (35)	90 (10)	55 (29)
	Traveled distance of every run [m]	42.3 (14.8)	47.2 (10.5)	50.2 (5.4)	38.4 (16.3)
	Traveled distance of successful runs [m]	51.2 (1.5)	51.4 (1.7)	51.4 (2.4)	48.4 (0.4)
	Average acceleration [m/s^2]	0.7 (0.1)	0.6 (0.1)	0.6 (0.1)	1.5 (0.4)
	Average jerk [m/s^3]	5.5 (1.1)	4.6 (1.0)	4.4 (0.8)	8.6 (2.4)
Velocity noise	Success rate [%]	55 (19)	60 (18)	82 (15)	35 (25)
	Traveled distance of every run [m]	40.5 (15.3)	39.9 (16.5)	49.0 (13.1)	33.4 (16.6)
	Traveled distance of successful runs [m]	51.6 (1.5)	51.7 (1.6)	53.9 (5.8)	50.3 (0.6)
	Average acceleration [m/s^2]	1.0 (0.2)	0.8 (0.2)	0.9 (0.2)	1.4 (0.2)
	Average jerk [m/s^3]	13.1 (2.2)	8.9 (1.6)	9.7 (1.3)	8.0 (1.7)
Image noise	Success rate [%]	50 (25)	66 (32)	73 (34)	32 (28)
	Traveled distance of every run [m]	46.3 (14.6)	45.6 (15.1)	57.0 (18.1)	31.8 (15.0)
	Traveled distance of successful runs [m]	53.4 (4.2)	55.3 (4.0)	60.6 (16.6)	48.4 (0.3)
	Average acceleration [m/s^2]	1.0 (0.1)	0.8 (0.1)	0.8 (0.1)	1.4 (0.3)
	Average jerk [m/s^3]	8.3 (1.0)	6.3 (1.0)	5.4 (0.8)	8.5 (1.6)
Both noise	Success rate [%]	32 (16)	53 (16)	64 (27)	24 (25)
	Traveled distance of every run [m]	42.9 (18.1)	44.1 (15.5)	52.8 (22.1)	30.2 (15.5)
	Traveled distance of successful runs [m]	54.0 (5.4)	58.8 (9.7)	62.1 (10.8)	50.2 (0.6)
	Average acceleration [m/s^2]	1.2 (0.1)	0.9 (0.1)	1.0 (0.1)	1.5 (0.3)
	Average jerk [m/s^3]	15.3 (1.8)	9.9 (1.1)	10.4 (1.3)	8.5 (1.8)

Simulation results in Flightmare with 10 forests x 10 runs/forest,
the metrics are given in mean (std) values

Simulation Results (Flightmare)

Noise levels	Metrics	Naive	Ensemble	ORACLE	Agile
No noise	Success rate [%]	65 (31)	80 (35)	90 (10)	55 (29)
	Traveled distance of every run [m]	42.3 (14.8)	47.2 (10.5)	50.2 (5.4)	38.4 (16.3)
	Traveled distance of successful runs [m]	51.2 (1.5)	51.4 (1.7)	51.4 (2.4)	48.4 (0.4)
	Average acceleration [m/s ²]	0.7 (0.1)	0.6 (0.1)	0.6 (0.1)	1.5 (0.4)
	Average jerk [m/s ³]	5.5 (1.1)	4.6 (1.0)	4.4 (0.8)	8.6 (2.4)
Velocity noise	Success rate [%]	55 (19)	60 (18)	82 (15)	35 (25)
	Traveled distance of every run [m]	40.5 (15.3)	39.9 (16.5)	49.0 (13.1)	33.4 (16.6)
	Traveled distance of successful runs [m]	51.6 (1.5)	51.7 (1.6)	53.9 (5.8)	50.3 (0.6)
	Average acceleration [m/s ²]	1.0 (0.2)	0.8 (0.2)	0.9 (0.2)	1.4 (0.2)
	Average jerk [m/s ³]	13.1 (2.2)	8.9 (1.6)	9.7 (1.3)	8.0 (1.7)
Image noise	Success rate [%]	50 (25)	66 (32)	73 (34)	32 (28)
	Traveled distance of every run [m]	46.3 (14.6)	45.6 (15.1)	57.0 (18.1)	31.8 (15.0)
	Traveled distance of successful runs [m]	53.4 (4.2)	55.3 (4.0)	60.6 (16.6)	48.4 (0.3)
	Average acceleration [m/s ²]	1.0 (0.1)	0.8 (0.1)	0.8 (0.1)	1.4 (0.3)
	Average jerk [m/s ³]	8.3 (1.0)	6.3 (1.0)	5.4 (0.8)	8.5 (1.6)
Both noise	Success rate [%]	32 (16)	53 (16)	64 (27)	24 (25)
	Traveled distance of every run [m]	42.9 (18.1)	44.1 (15.5)	52.8 (22.1)	30.2 (15.5)
	Traveled distance of successful runs [m]	54.0 (5.4)	58.8 (9.7)	62.1 (10.8)	50.2 (0.6)
	Average acceleration [m/s ²]	1.2 (0.1)	0.9 (0.1)	1.0 (0.1)	1.5 (0.3)
	Average jerk [m/s ³]	15.3 (1.8)	9.9 (1.1)	10.4 (1.3)	8.5 (1.8)

- Naïve-Agile similar performance

Simulation results in Flightmare with 10 forests x 10 runs/forest,
the metrics are given in mean (std) values

Simulation Results (Flightmare)

Noise levels	Metrics	Naive	Ensemble	ORACLE	Agile
No noise	Success rate [%]	65 (31)	80 (35)	90 (10)	55 (29)
	Traveled distance of every run [m]	42.3 (14.8)	47.2 (10.5)	50.2 (5.4)	38.4 (16.3)
	Traveled distance of successful runs [m]	51.2 (1.5)	51.4 (1.7)	51.4 (2.4)	48.4 (0.4)
	Average acceleration [m/s ²]	0.7 (0.1)	0.6 (0.1)	0.6 (0.1)	1.5 (0.4)
	Average jerk [m/s ³]	5.5 (1.1)	4.6 (1.0)	4.4 (0.8)	8.6 (2.4)
Velocity noise	Success rate [%]	55 (19)	60 (18)	82 (15)	35 (25)
	Traveled distance of every run [m]	40.5 (15.3)	39.9 (16.5)	49.0 (13.1)	33.4 (16.6)
	Traveled distance of successful runs [m]	51.6 (1.5)	51.7 (1.6)	53.9 (5.8)	50.3 (0.6)
	Average acceleration [m/s ²]	1.0 (0.2)	0.8 (0.2)	0.9 (0.2)	1.4 (0.2)
	Average jerk [m/s ³]	13.1 (2.2)	8.9 (1.6)	9.7 (1.3)	8.0 (1.7)
Image noise	Success rate [%]	50 (25)	66 (32)	73 (34)	32 (28)
	Traveled distance of every run [m]	46.3 (14.6)	45.6 (15.1)	57.0 (18.1)	31.8 (15.0)
	Traveled distance of successful runs [m]	53.4 (4.2)	55.3 (4.0)	60.6 (16.6)	48.4 (0.3)
	Average acceleration [m/s ²]	1.0 (0.1)	0.8 (0.1)	0.8 (0.1)	1.4 (0.3)
	Average jerk [m/s ³]	8.3 (1.0)	6.3 (1.0)	5.4 (0.8)	8.5 (1.6)
Both noise	Success rate [%]	32 (16)	53 (16)	64 (27)	24 (25)
	Traveled distance of every run [m]	42.9 (18.1)	44.1 (15.5)	52.8 (22.1)	30.2 (15.5)
	Traveled distance of successful runs [m]	54.0 (5.4)	58.8 (9.7)	62.1 (10.8)	50.2 (0.6)
	Average acceleration [m/s ²]	1.2 (0.1)	0.9 (0.1)	1.0 (0.1)	1.5 (0.3)
	Average jerk [m/s ³]	15.3 (1.8)	9.9 (1.1)	10.4 (1.3)	8.5 (1.8)

- **Naive-Agile** similar performance
- Simulation with noise: **ORACLE, Ensemble** outperform **Agile**

Simulation results in Flightmare with 10 forests x 10 runs/forest, the metrics are given in mean (std) values

Experiment 1: ORACLE in cluttered corridor

Waypoint: 35 m ahead of the robot



- Reference forward speed: 2.5 m/s
- Realsense D455 + Realsense T265 + PixRacer IMU
- Position estimate and map - not provided to CPN

Code Examples and Tasks

- https://github.com/ntnu-arl/aerial_gym_simulator
- https://ntnu-arl.github.io/aerial_gym_simulator/
- <https://github.com/ntnu-arl/ORACLE>

Find out more

- Nguyen, H., Andersen, R., Boukas, E. and Alexis, K., 2024. Uncertainty-aware visually-attentive navigation using deep neural networks. *The International Journal of Robotics Research*, 43(6), pp.840-872.
- A. Loquercio et al., "Learning high-speed flight in the wild," *Science Robotics*, 2021
- G. Kahn et al., "Badgr: An autonomous self-supervised learning-based navigation system," *RAL* 2021
- S. Veer et al., "Probably approximately correct vision-based planning using motion primitives," *CoRL* 2020