RobotiXX

Learning Robot Navigation in Challenging Environments

Xuesu Xiao, Ph.D.

Assistant Professor

Computer Science, George Mason University

Disaster Robotics

Modular Snake (Mexico City Earthquake)

[***XX*** *et al., ICRA15*]

- Overhead Cameras
- Locomotive Reduction

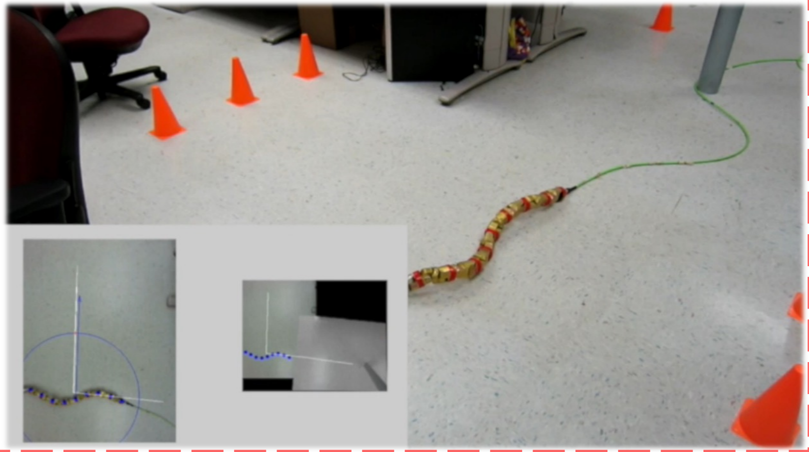

EMILY (Greece Refugee Crisis)

[***XX*** *et al., IROS17, D,* ***XX****, M, SSRR17*]

- UAV-USV Team
- Visual Pose Stabilization
- Visual Navigation

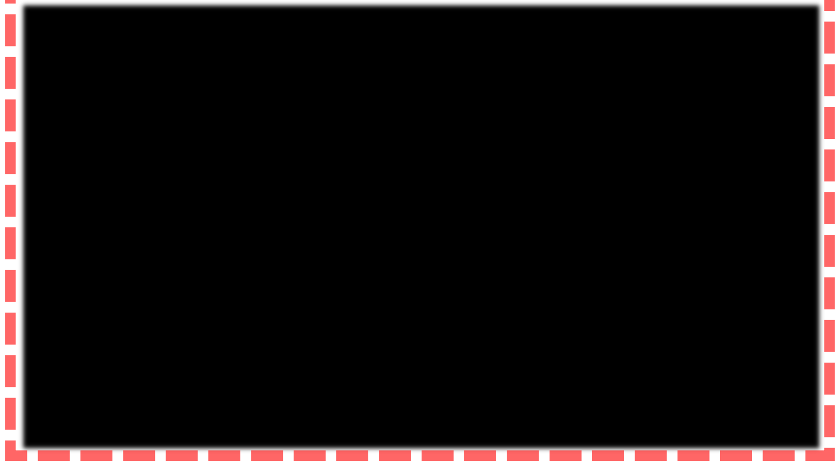

PackBot (Fukushima Daiichi)

[*F,* ***XX****, M, THMS21,*
XX *et al., RA-L20,*
XX *et al., FSR19*
XX *et al., SSRR19a,*
XX *et al., SSRR19b,*
XX *et al., IROS18,*
XX *et al., SSRR18*
(Best Paper Finalist),
XX *et al., SSRR17*]

- Viewpoint Theory
- Risk-Awareness
- Tethered Flight

Learning Navigation in Challenging Environments

Highly-Constrained Environments

[***XX*** *et al., RA-L21a,* ***XX*** *et al., ICRA21,* W, ***XX*** *et al., IROS21*]

Offroad Environments

[***XX*** *et al., RA-L21b, K, S, A, R,* ***XX*** *et al., IROS22, D, P, N,* ***XX****, ICRA24, D, P,* ***XX****, under review*]

Social Environments

[***XX*** *et al., CoRL22, K, N,* ***XX*** *et al., RA-L22, N, N, P, D,* ***XX****, IROS23*]

The BARN Challenge and Datasets

ICRA 2022
Philadelphia
[***XX*** *et al. RAM22*]

ICRA 2023
London
[***XX*** *et al. RAM23*]

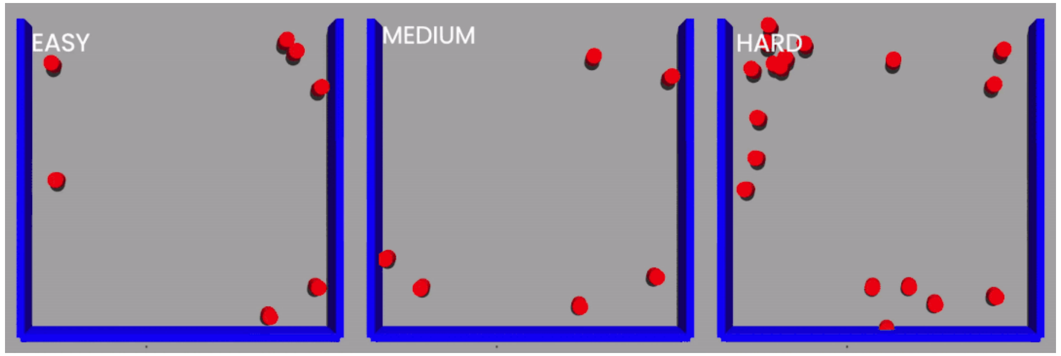

DynaBARN [N, J, H, X, L, ***XX***, *S, SSRR22*]

Benchmark Autonomous Robot Navigation (BARN)
[P, T, ***XX***, *S. SSRR20*]

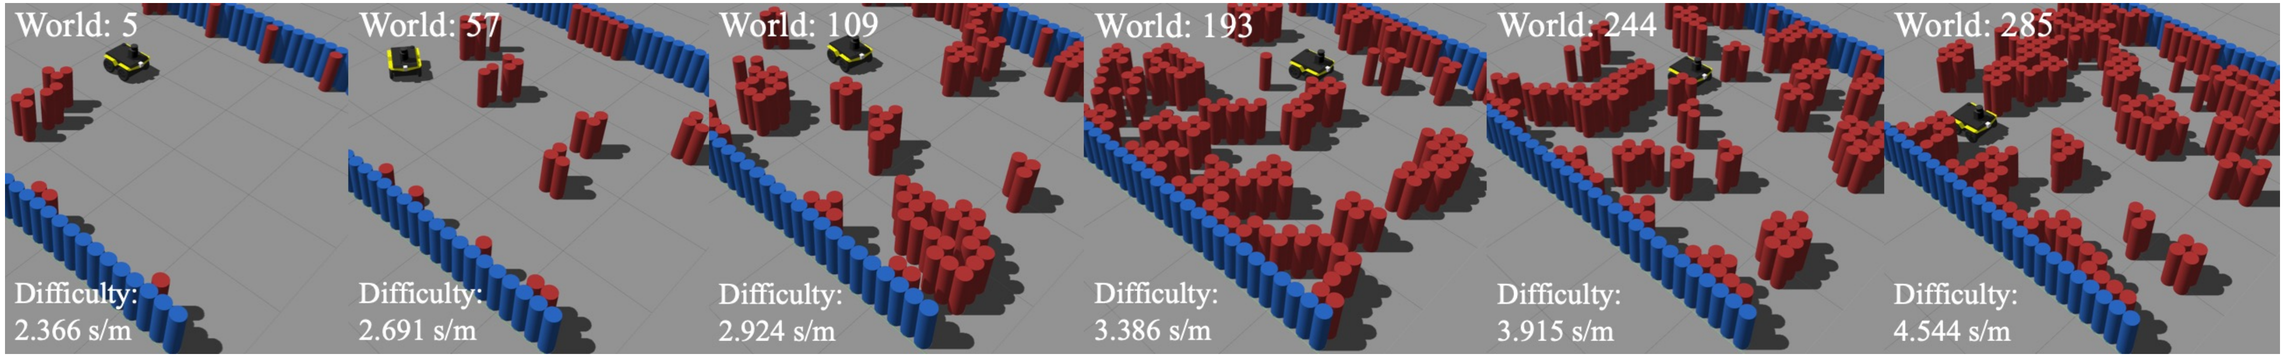

Learning from Hallucination (LfH)

Motivation: Highly-constrained conditions require **more computation** for classical sampling- or optimization-based methods or high-quality, but **more expensive training data** for learning methods.

Inspiration: Agile maneuvers in relatively (or completely) open spaces can be optimal for certain highly constrained environments.

Solution: Hallucinate many highly constrained environments to generate training data to learn an agile motion planner.

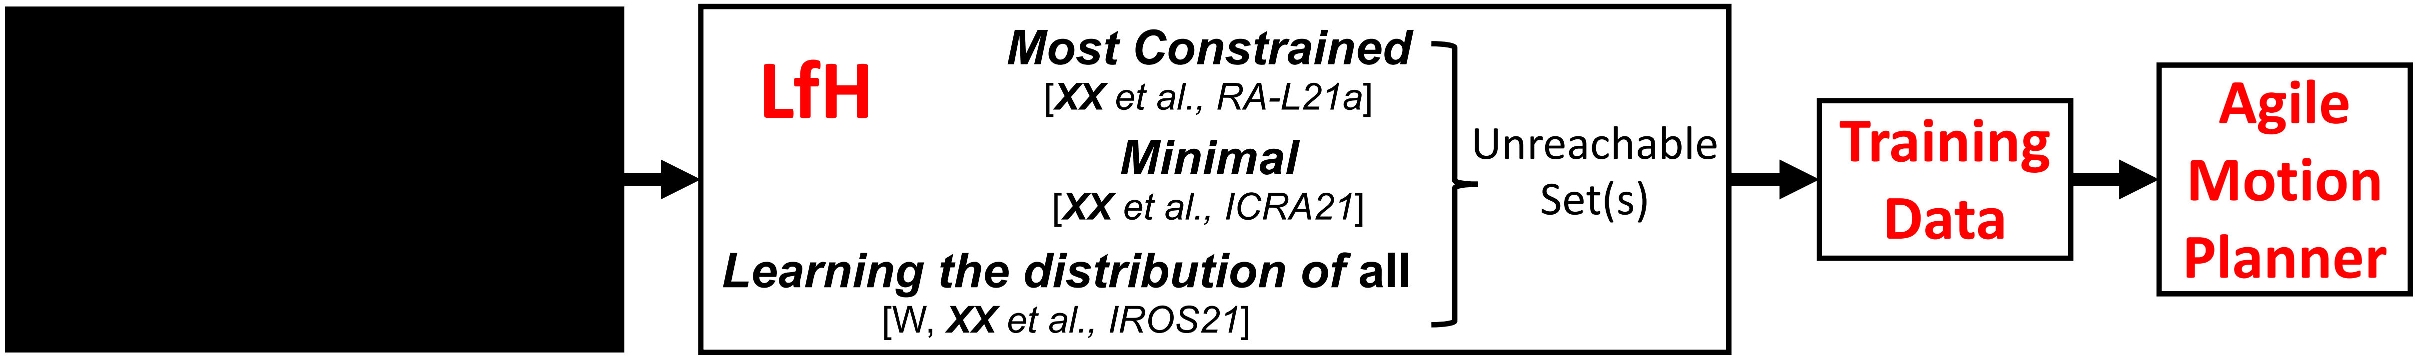

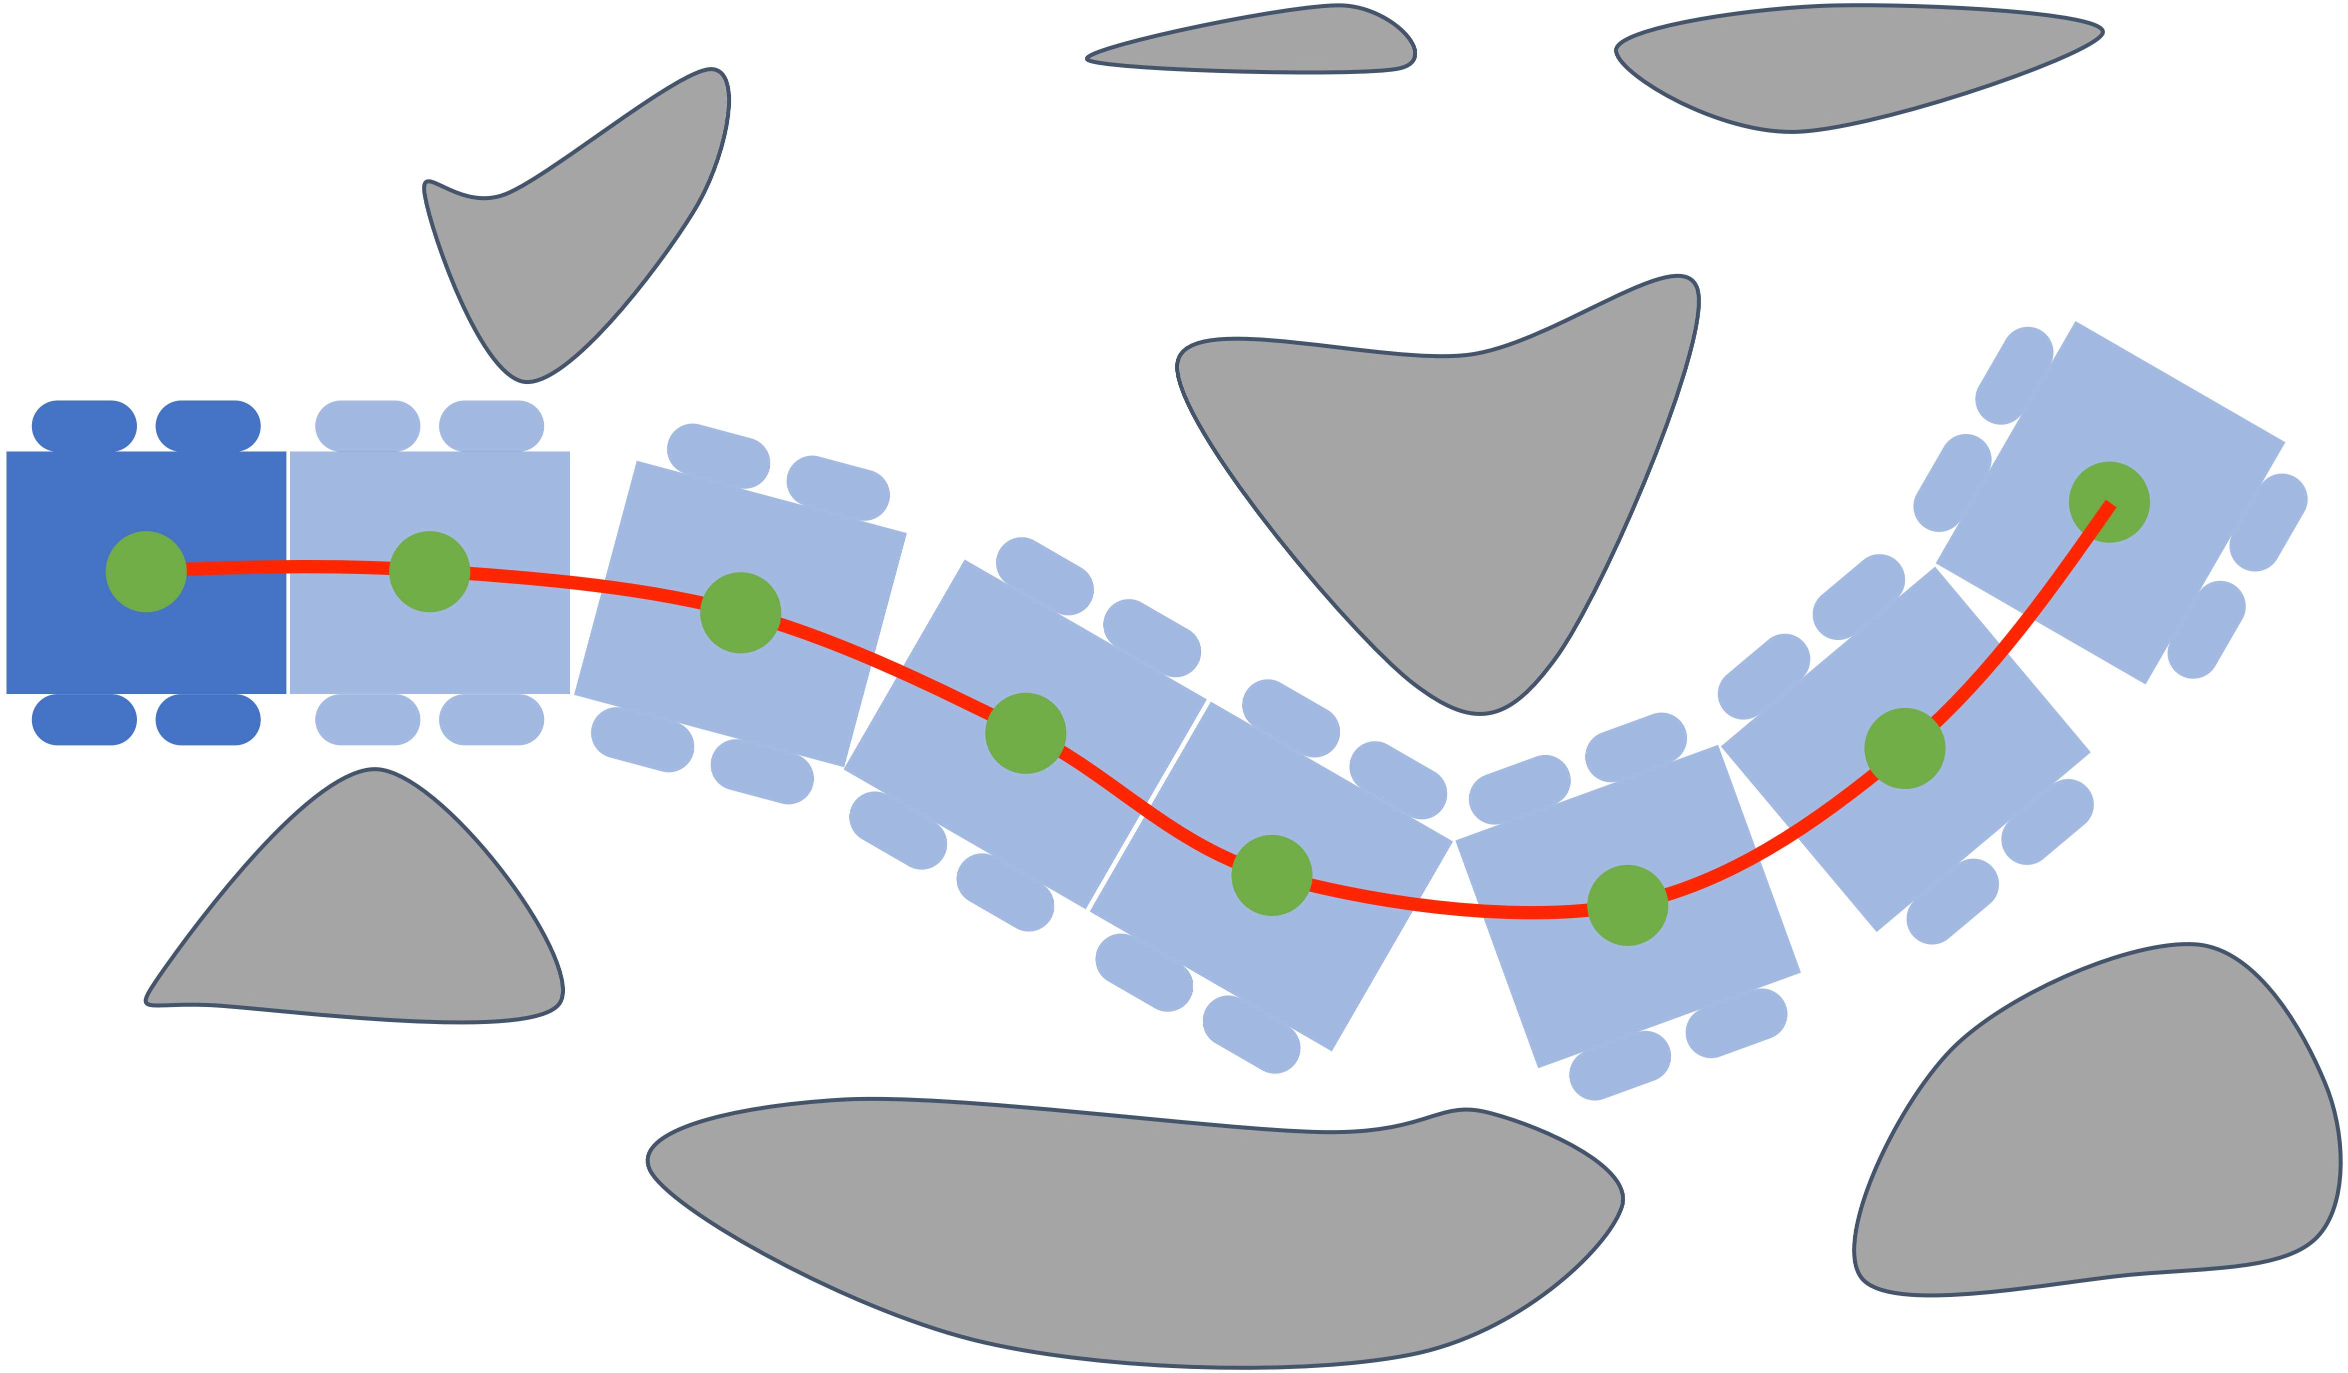

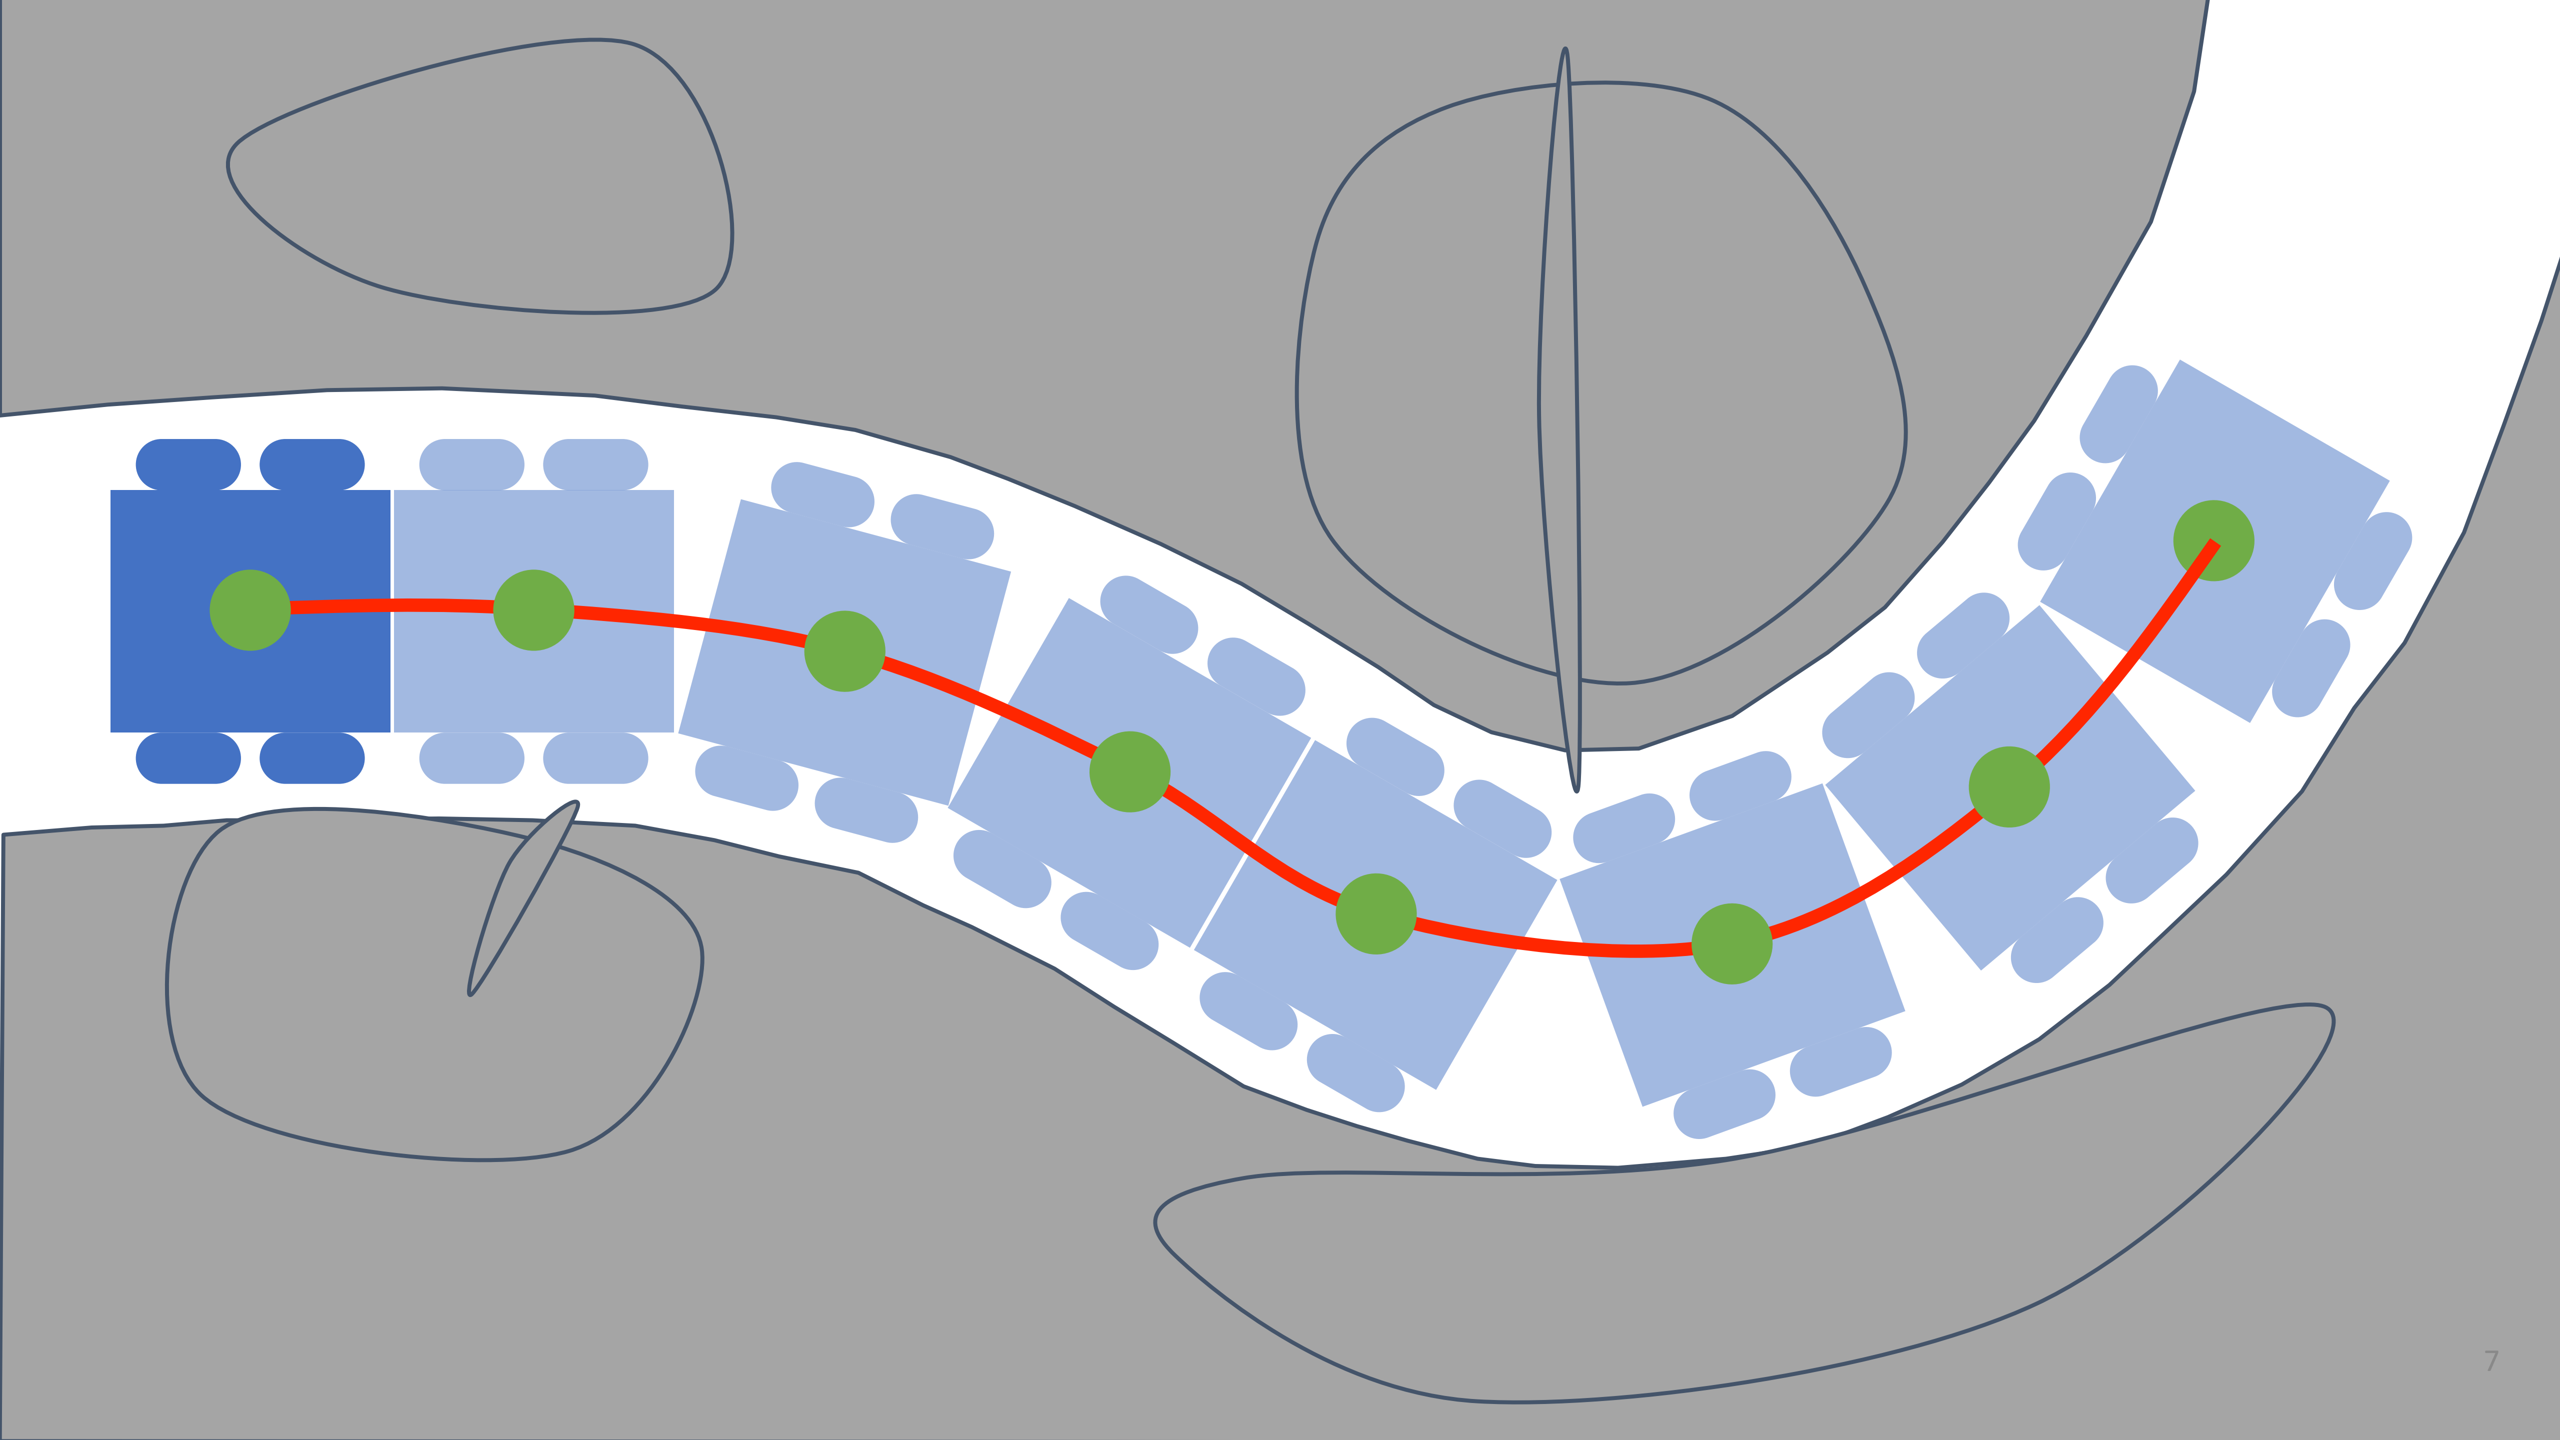

Three Learning from Hallucination (LfH) methods

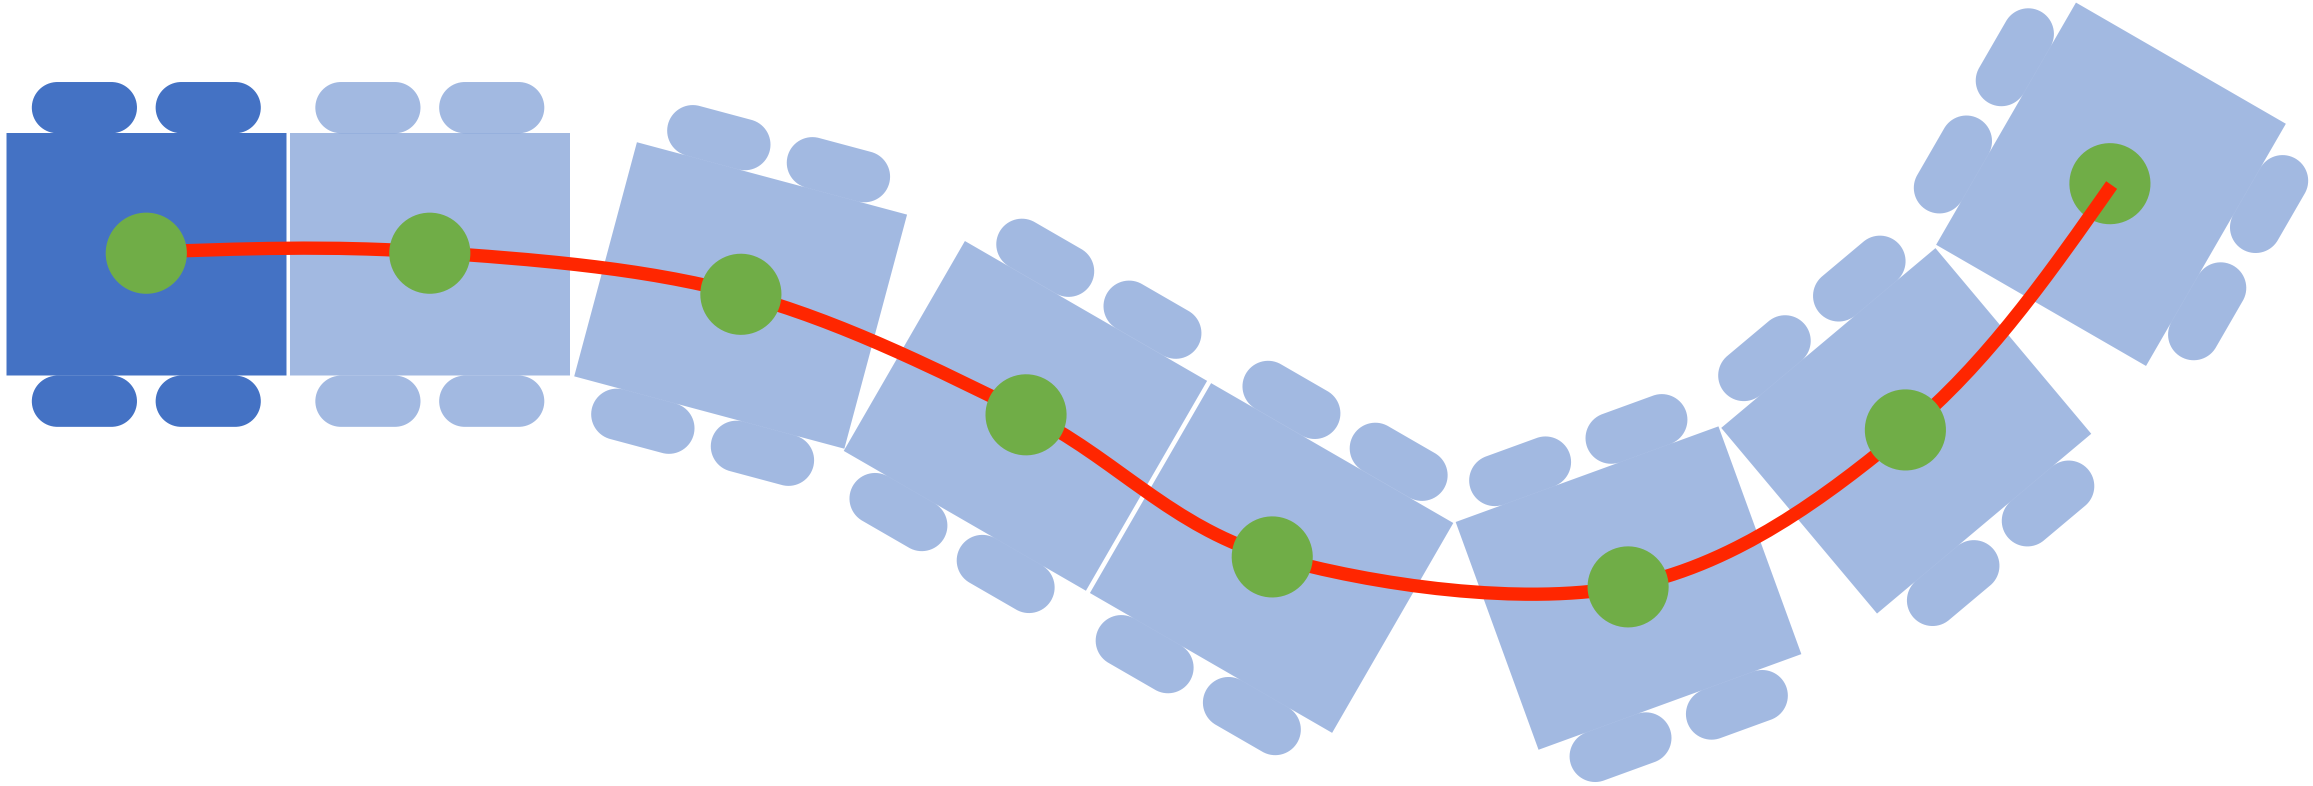

(1) The most constrained (maximal) obstacle hallucination

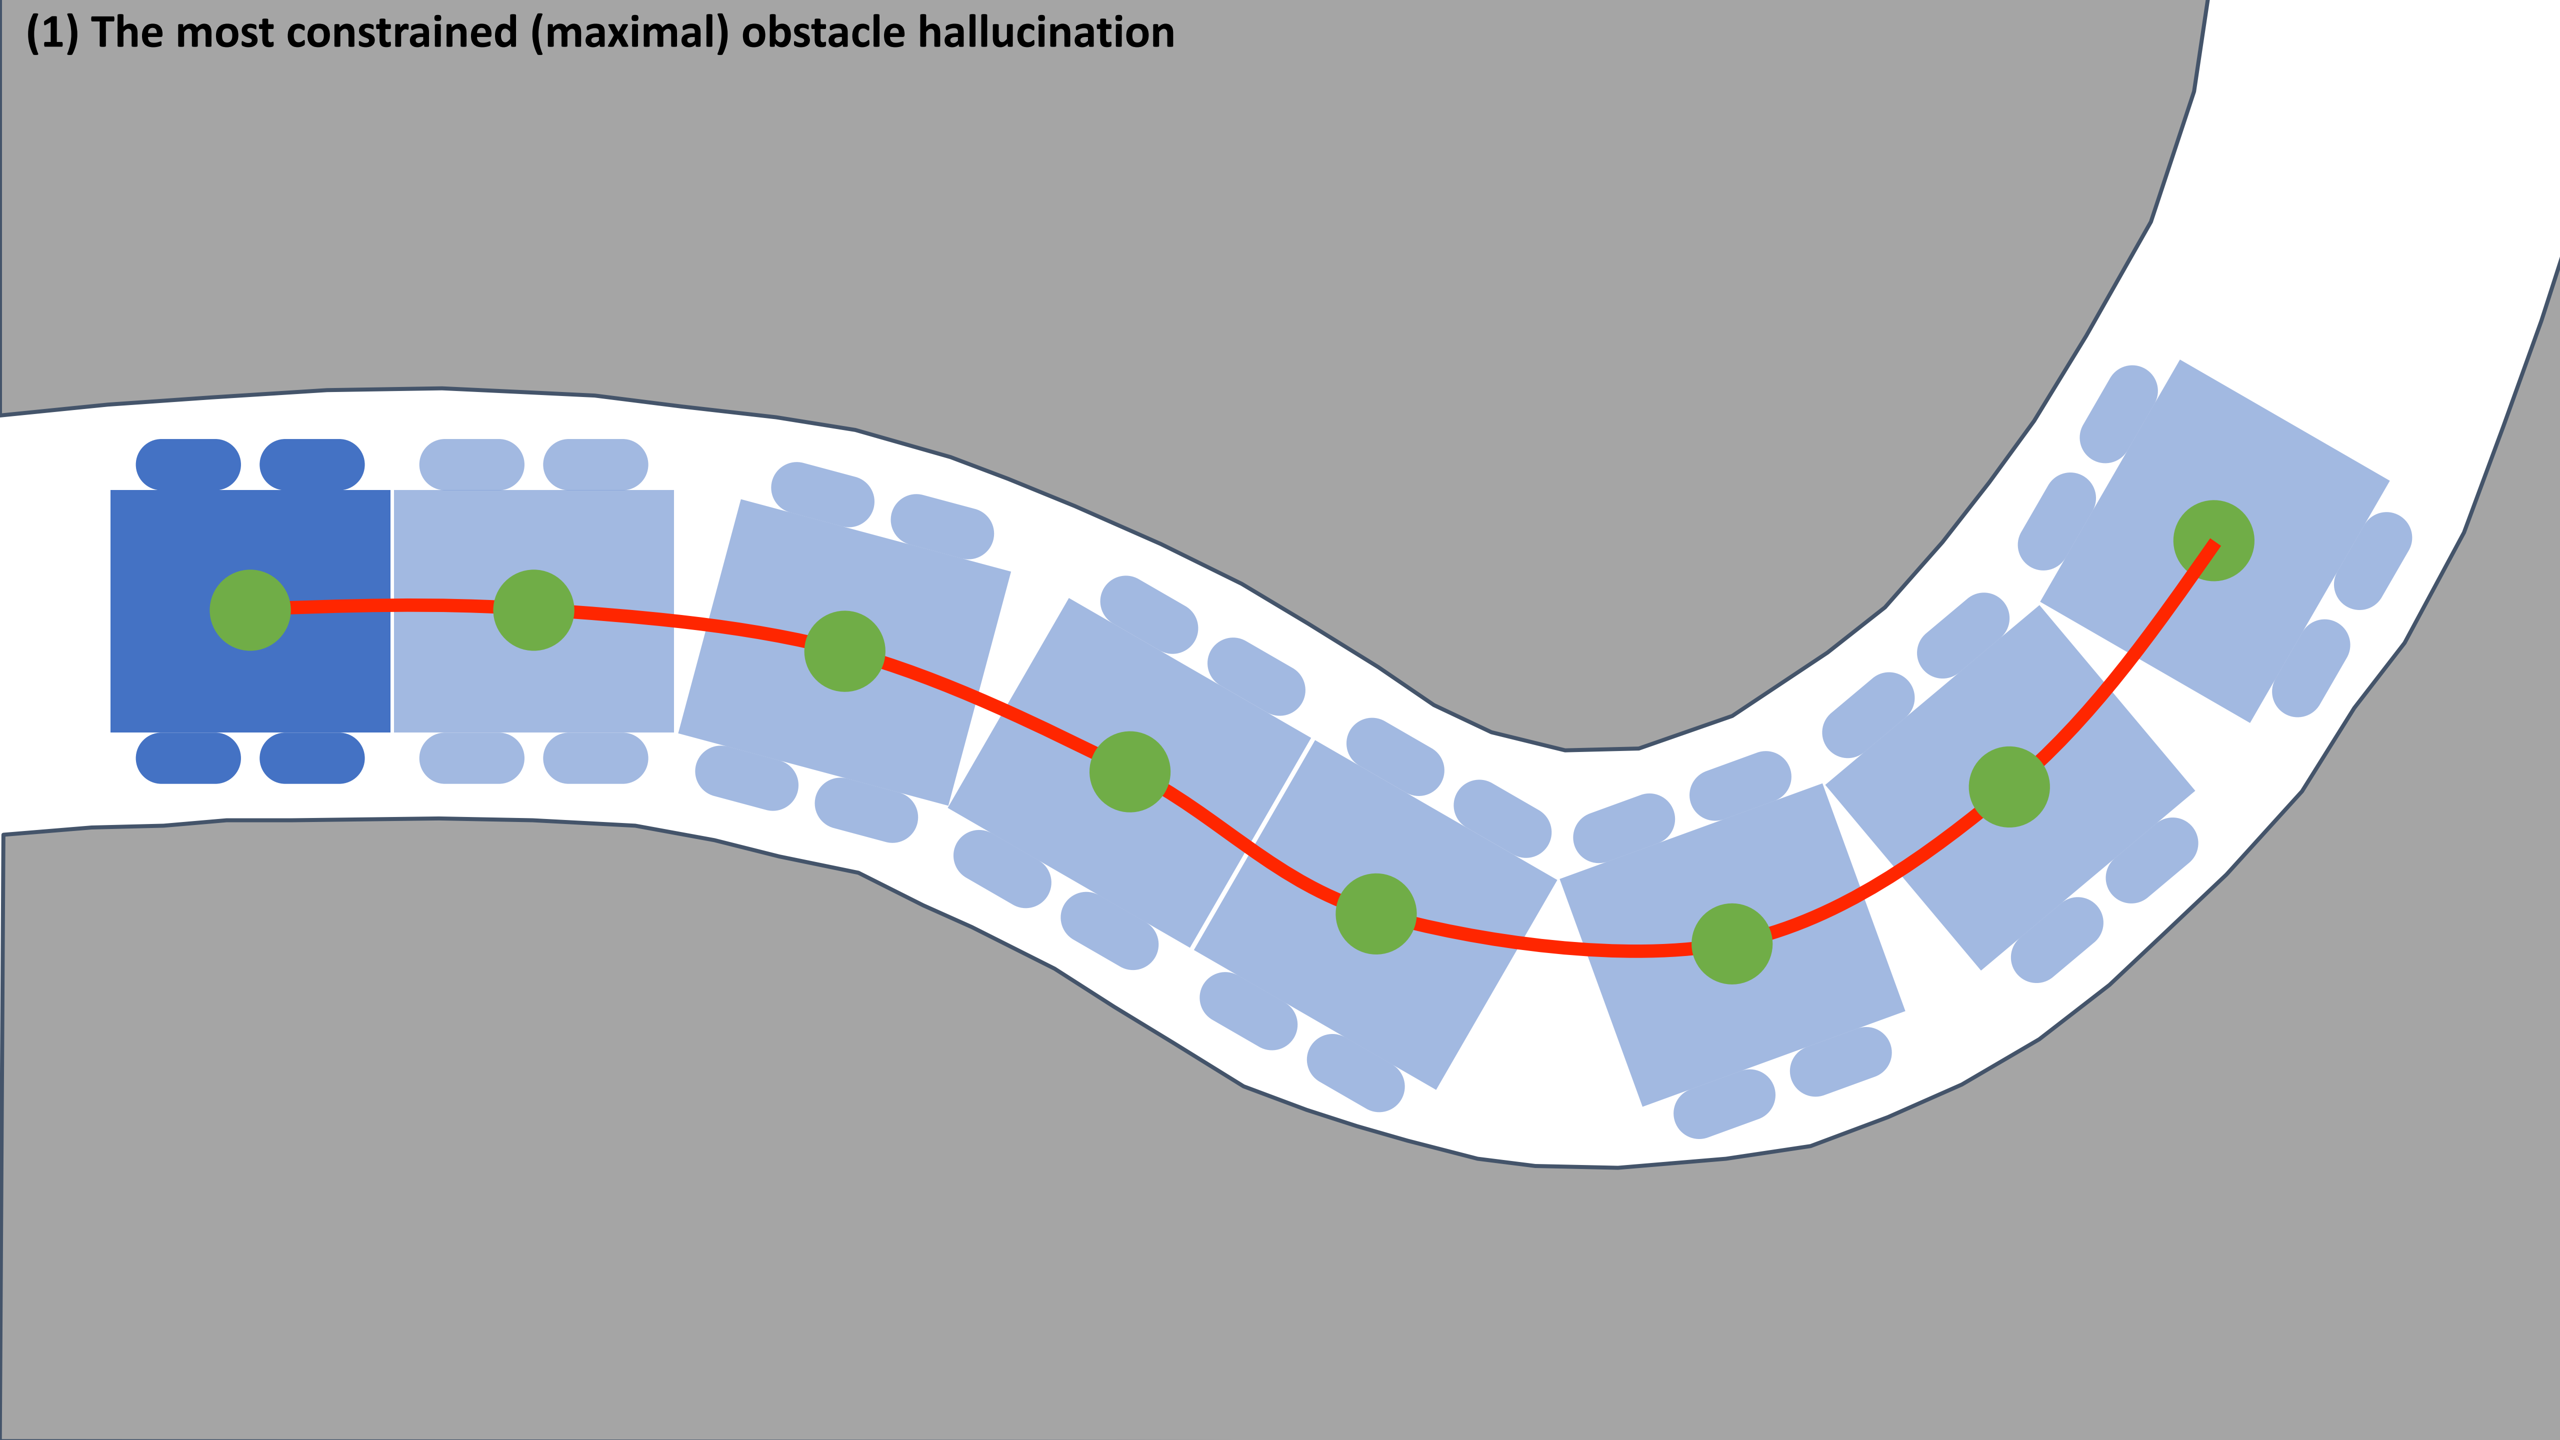

(2) A minimal obstacle hallucination

Then randomly sample additional obstacles

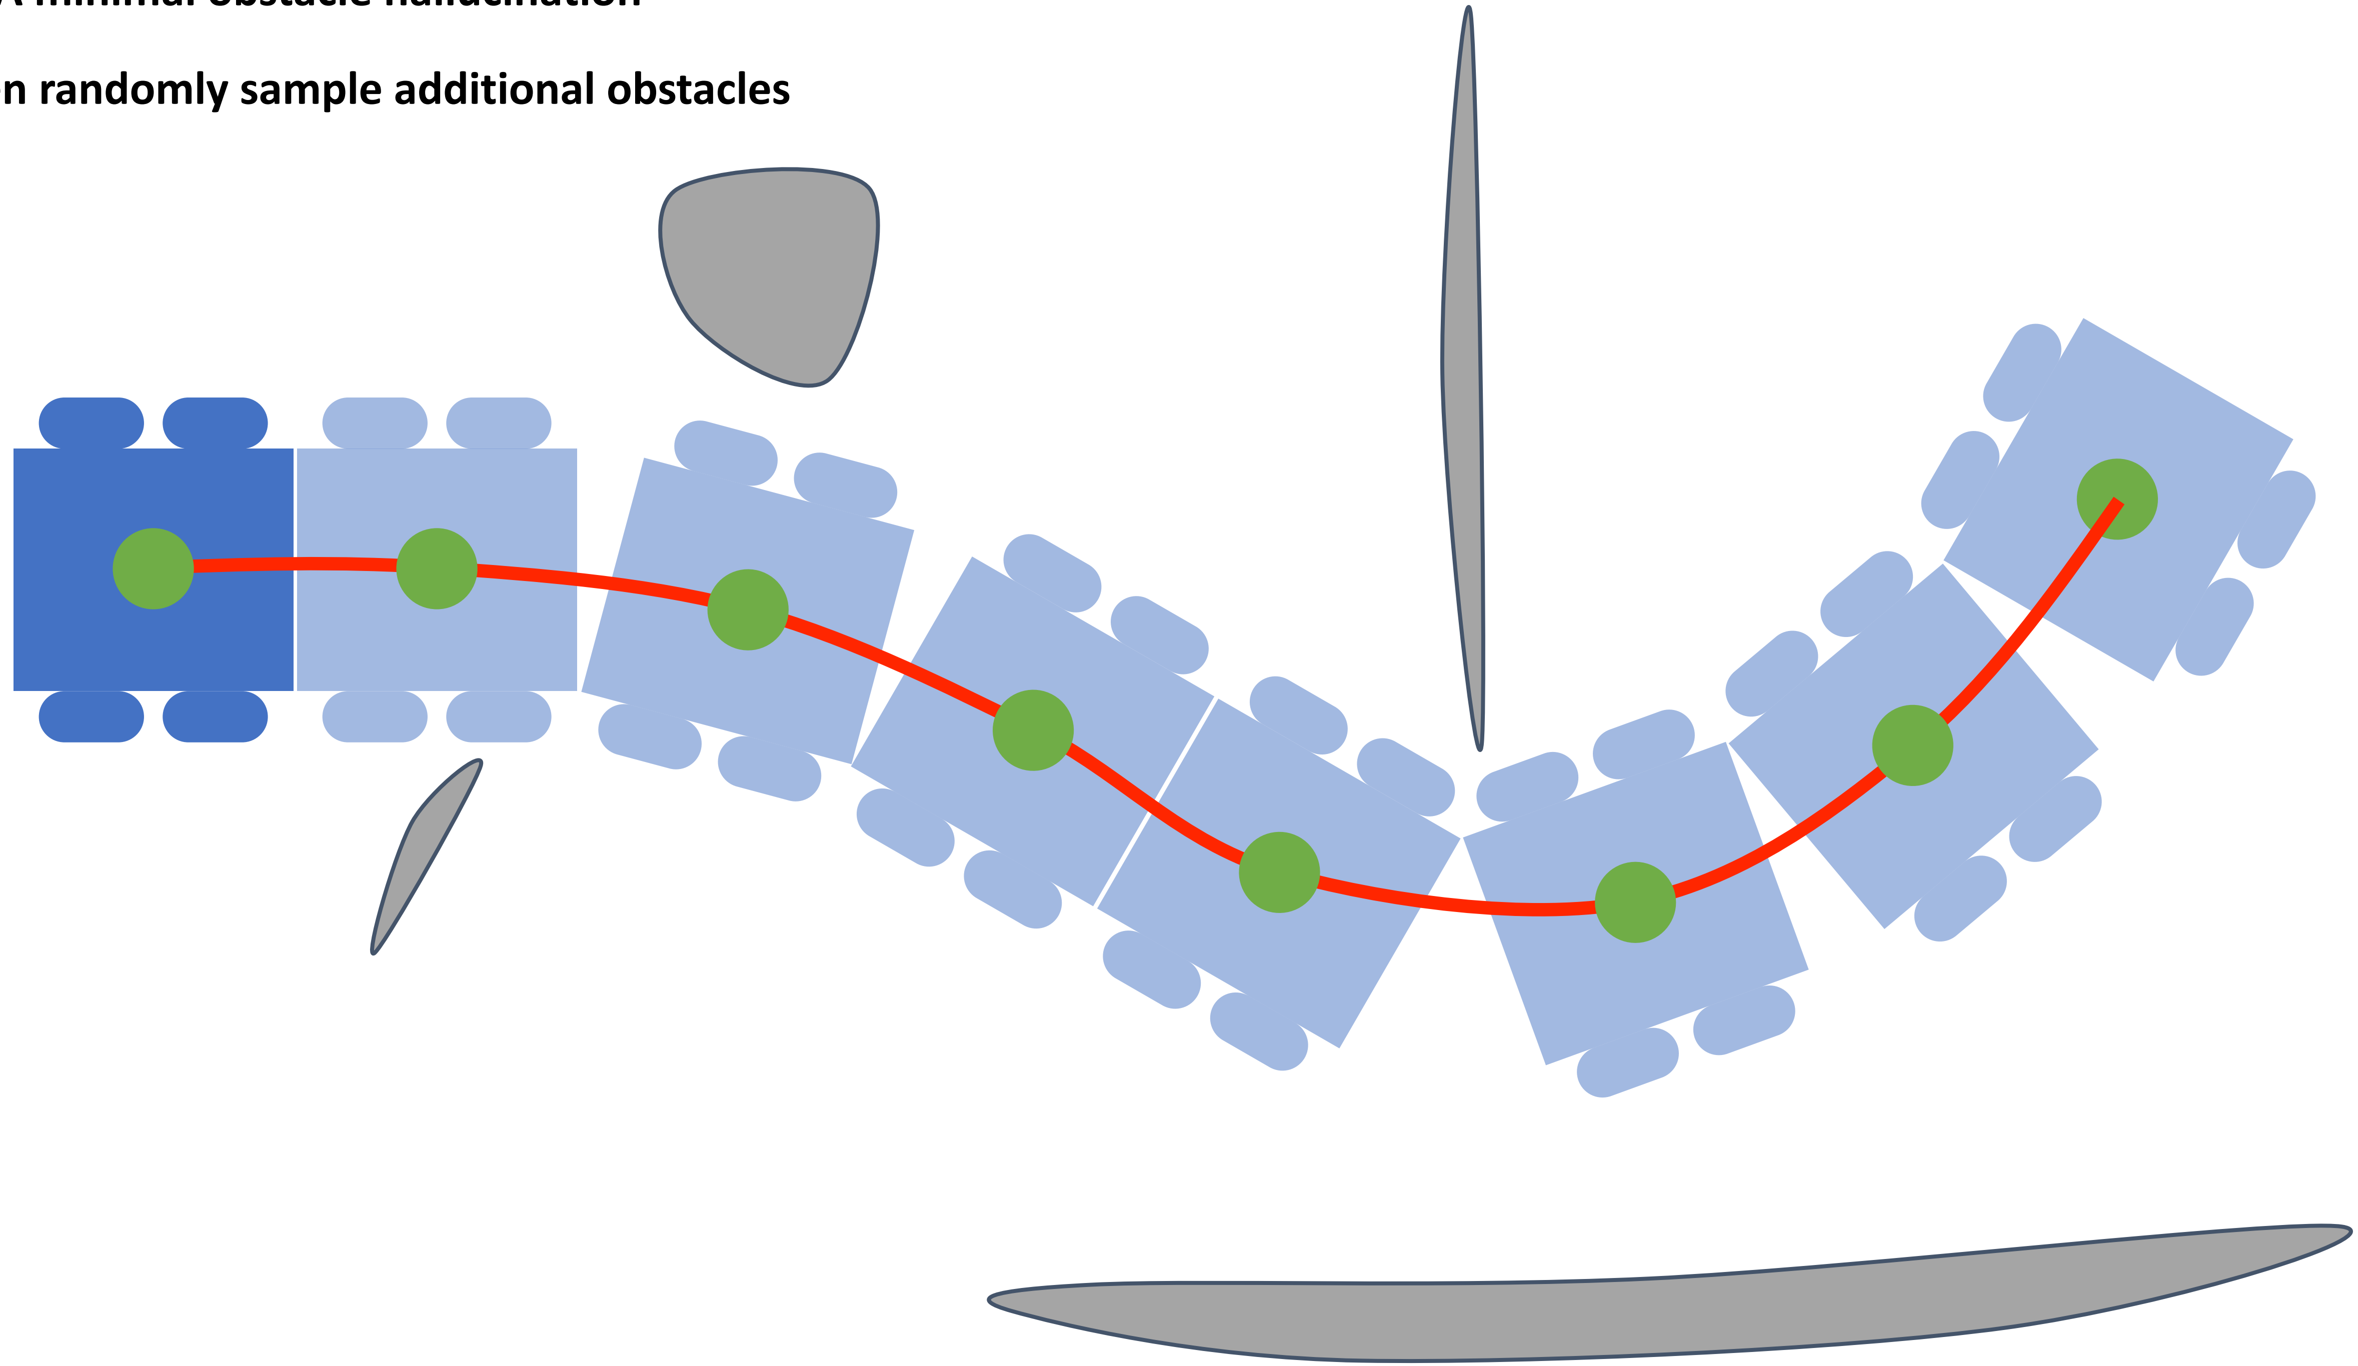

(3) Learn obstacle distributions

Then sample from learned distributions

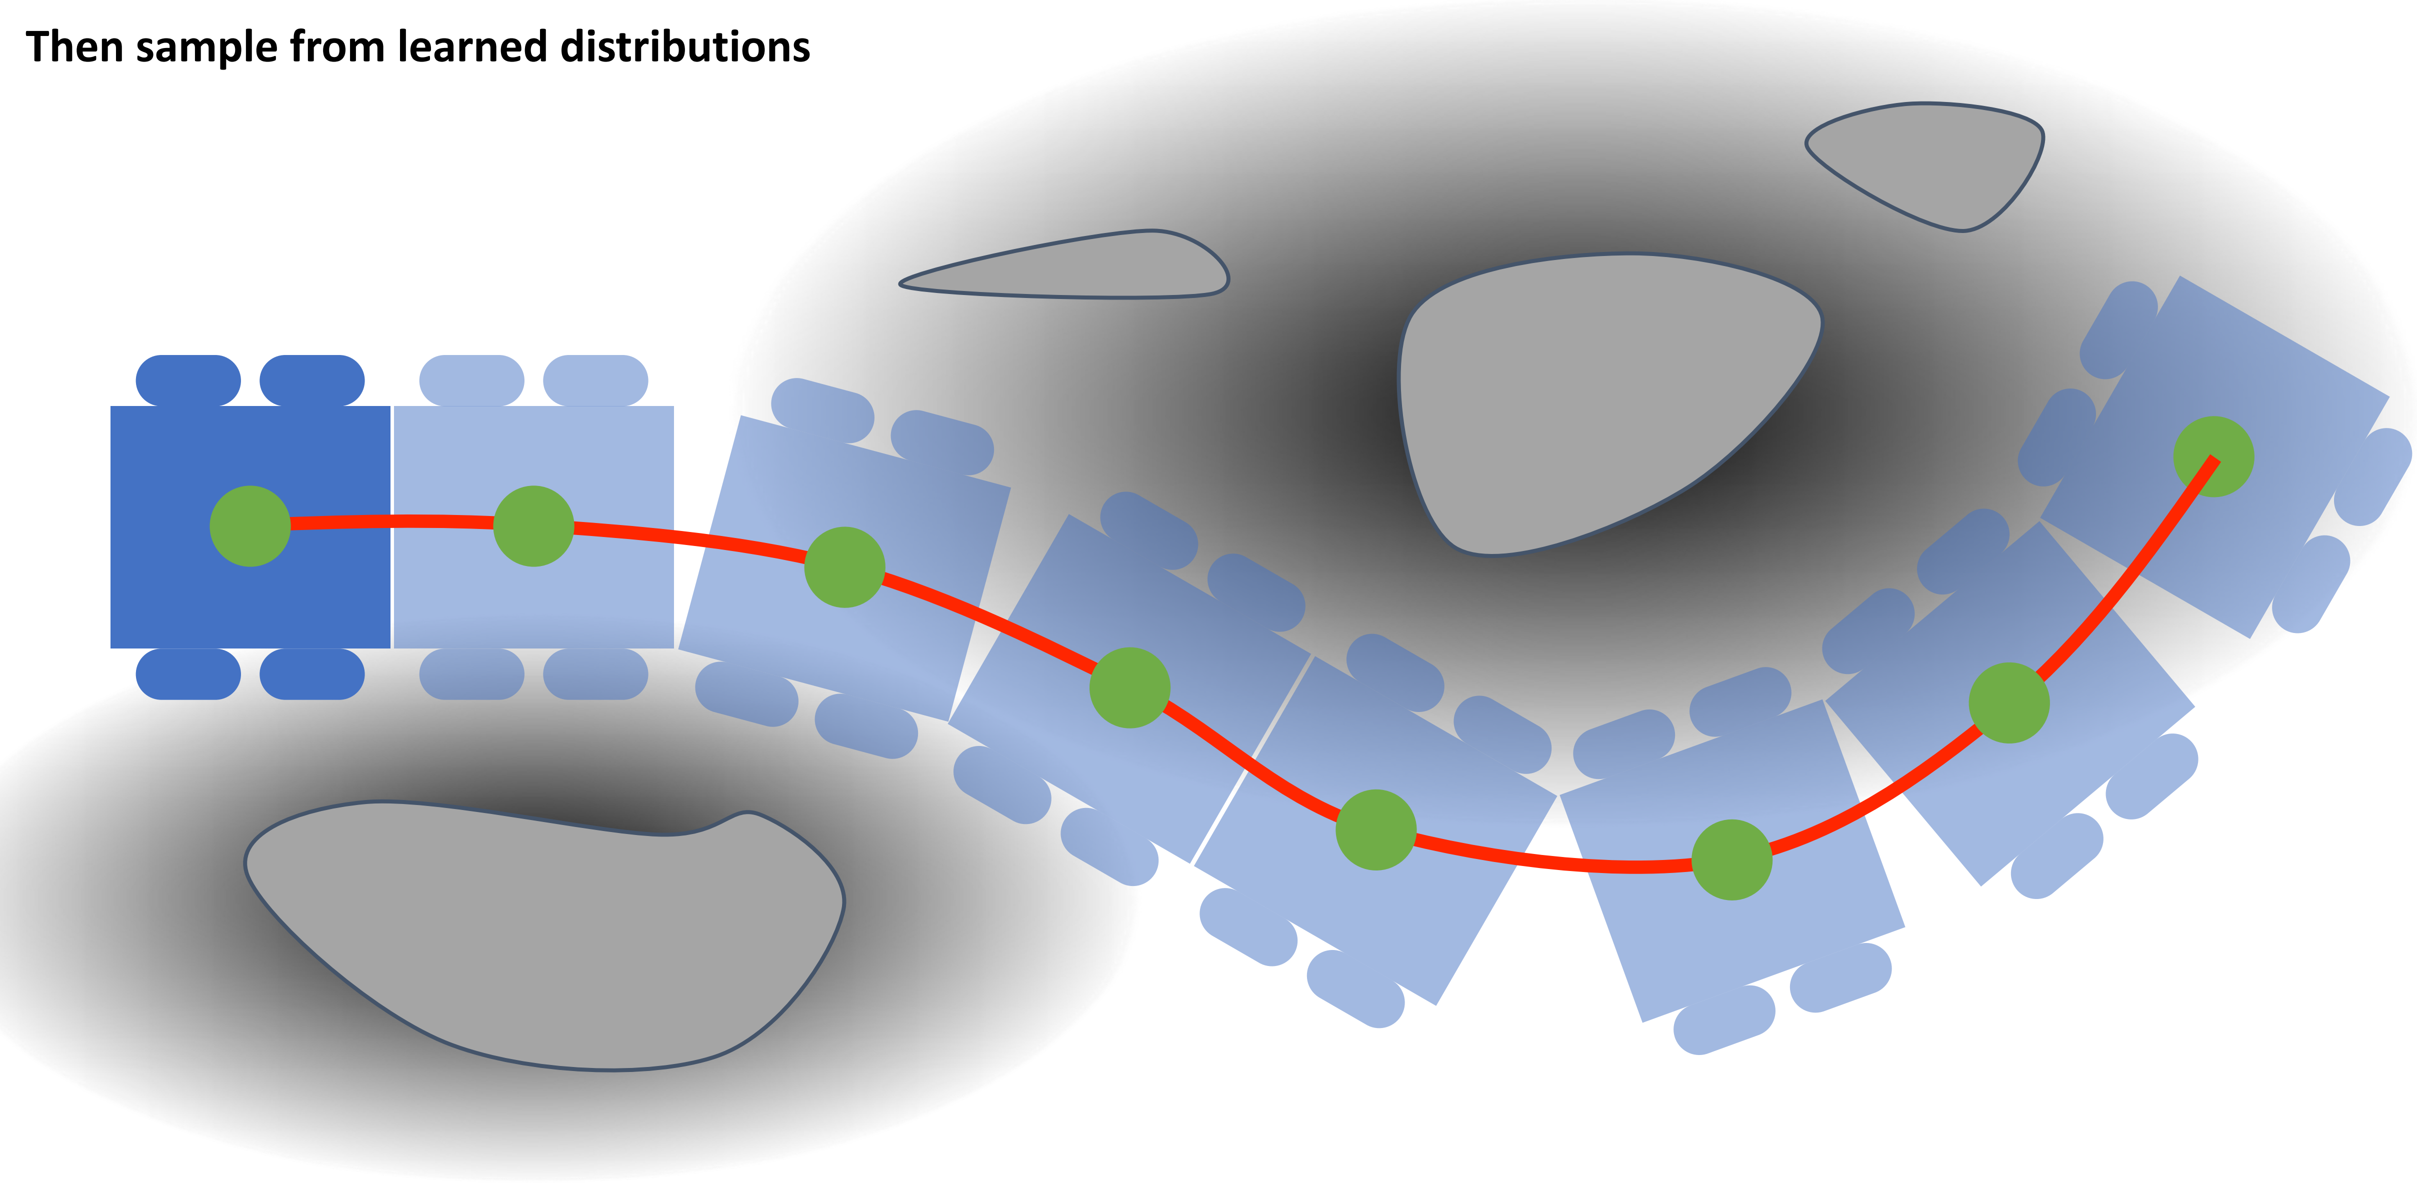

Motion Planning:

Find a motion planner as a function

$$f(\cdot)$$

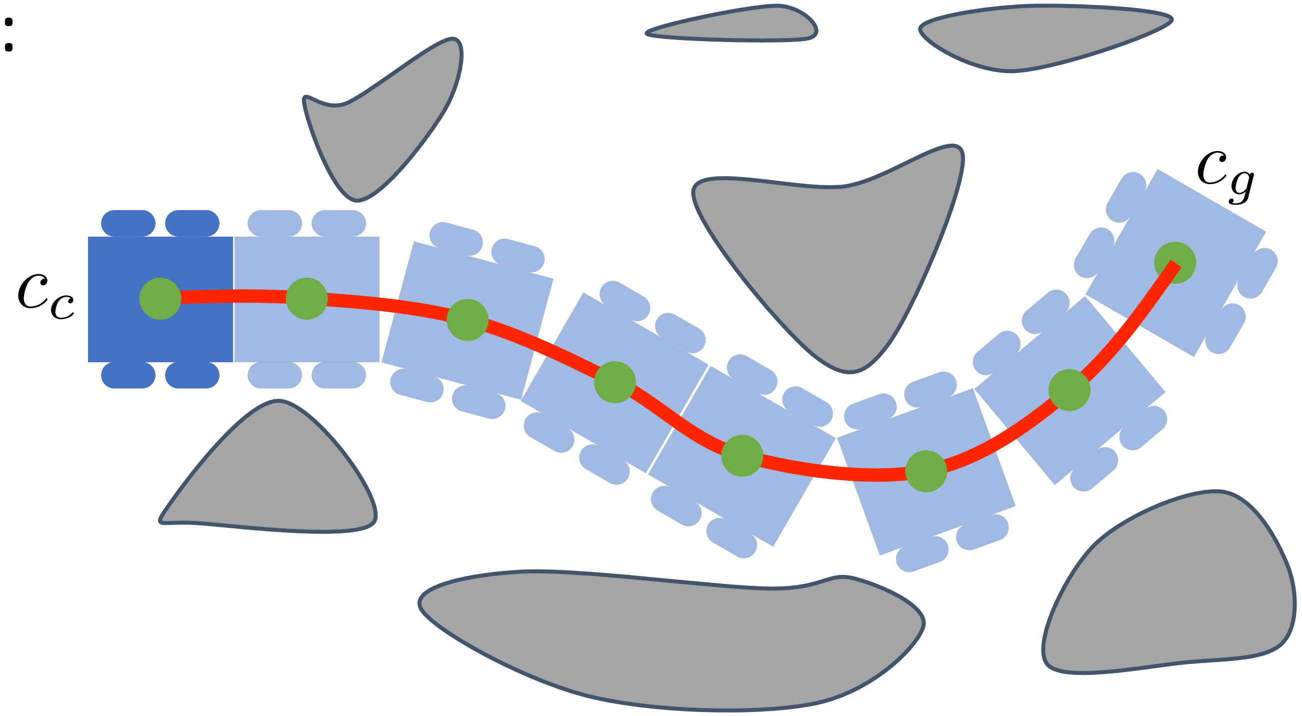

Hallucination:

Find a Hallucinator as a function

$$f^{-1}(\cdot)$$

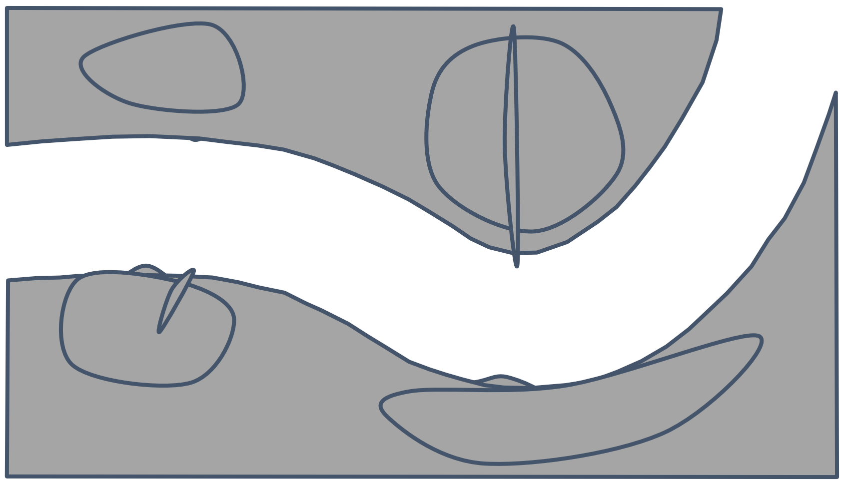

$$C_{obst} \overset{\longrightarrow}{\cancel{\longleftarrow}} p$$

Motion Planning:

Find a motion planner as a function

$$f(\cdot)$$

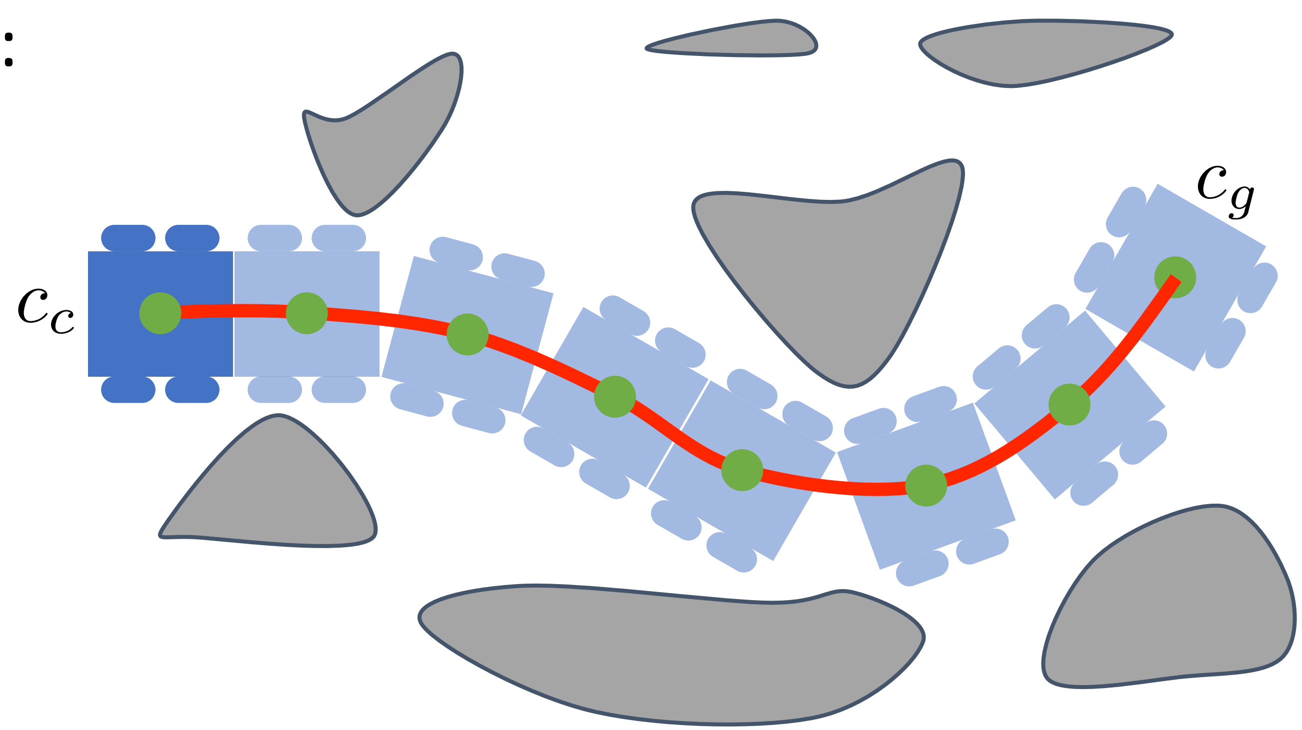

Hallucination:

Find a Hallucinator as a function

$$f^{-1}(\cdot)$$

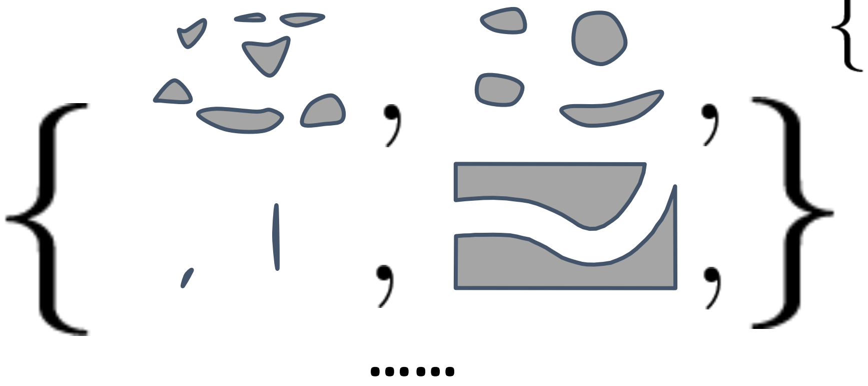

$$\{C^i_{obst}\}_{i=0}^{\infty}$$

p

Learning from *Most Constrained* Hallucination

[*XX et al., RA-L21a*]

- Hallucinating **the most constrained unreachable set**

$$C^*_{obst} = g(p \mid c_c, c_g) \quad \text{iff} \quad \forall C_{obst} \in \mathscr{C}_{obst},$$

$$f^*(C_{obst} \mid c_c, c_g) = p \implies C_{obst} \subseteq C^*_{obst},$$

- Train function approximator

$$g_\theta^{-1}(\cdot) : C^*_{obst} \to p$$

- During agile deployment
  - Seek help from coarse global plan
  - Create runtime hallucination
  - Query learned g_θ^{-1}

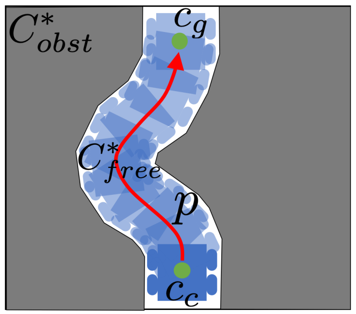

$p = f(C^*_{obst} \mid c_c, c_g)$

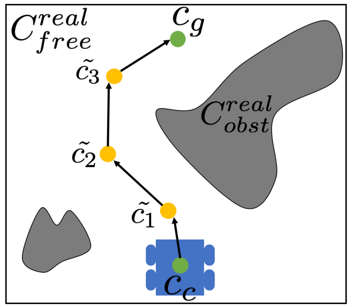

$\{\tilde{c}_j \mid 1 \leq j \leq m\} = gp(C^{real}_{obst} \mid c_c, c_g)$

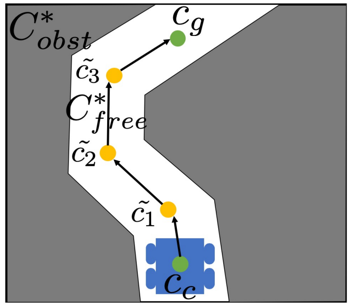

$C^*_{obst} \approx h(\{\tilde{c}_j \mid 1 \leq j \leq m\})$

Learning from *Minimal* Hallucination

[*XX et al., ICRA21*]

- **Motivation:** What if a global path is not available for runtime hallucination?
- **Inspiration:** Not every obstacle is required to make a plan optimal.
- **Solution:** Identify a (not unique) minimal unreachable set, randomly sample obstacles in addition to this set for training, and deploy without runtime hallucination (Hallucinated Learning and Sober Deployment, HLSD).

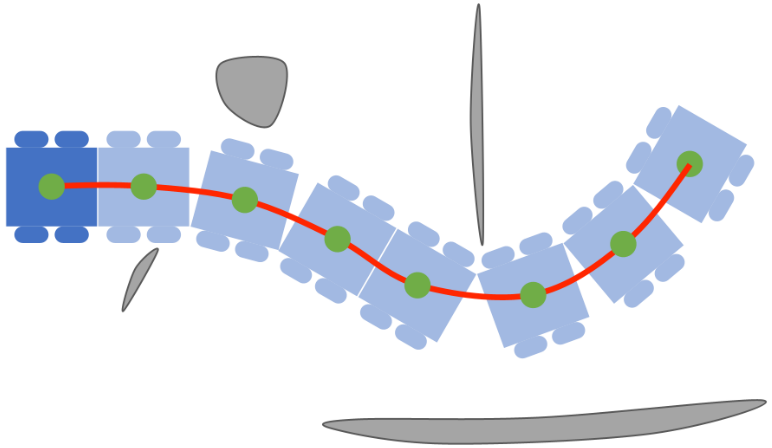

Learning from *Minimal* Hallucination

[*XX et al., ICRA21*]

Hallucinating **a minimal unreachable set**

$$\cancel{C^*_{obst} = g(p \mid c_c, c_g)} \; C^{min}_{obst} = o(p \mid c_c, c_g)$$

$$\mathscr{C}^{min}_{obst} \doteq \{C^{min}_{obst} \mid \forall c \in C^{min}_{obst}, f(C^{min}_{obst} \backslash \{c\} \mid c_c, c_g) \neq f(C^{min}_{obst} \mid c_c, c_g)\}$$

Point Mass
Holonomic
Shortest-Path

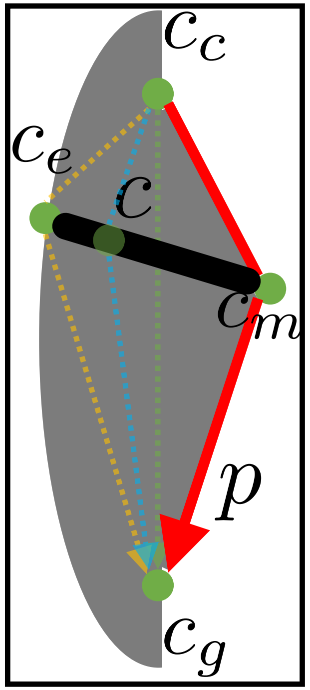

Realistic
Robot

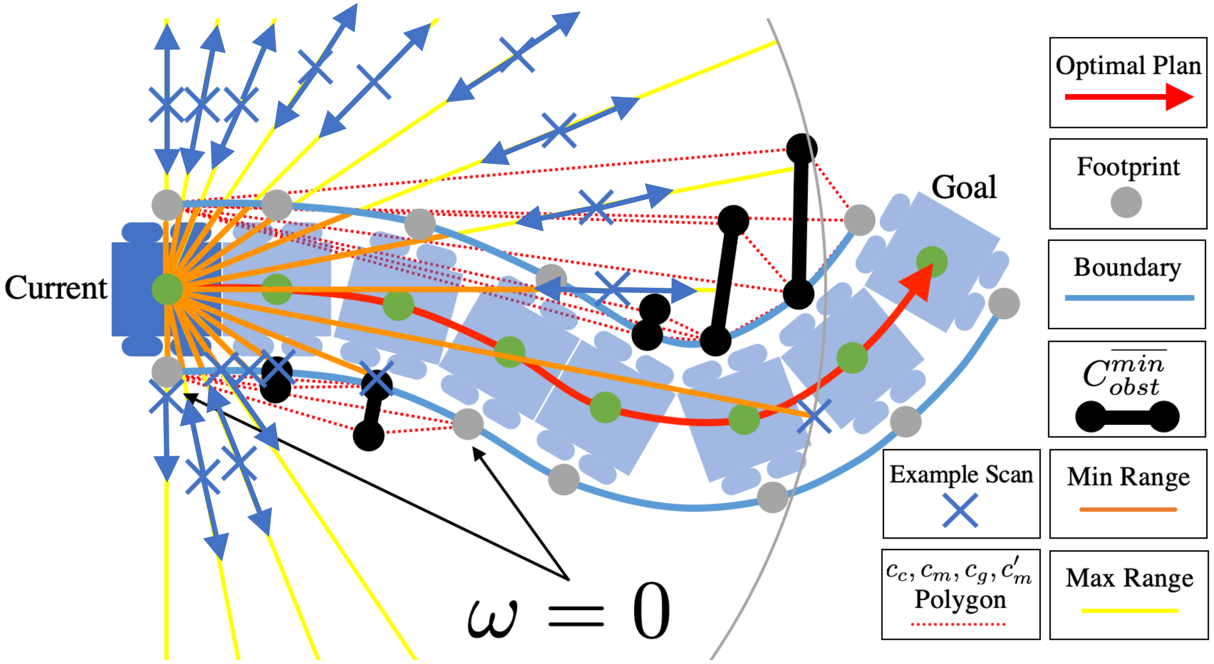

Learning from Learned Hallucination (LfLH)

[*W, **XX** et al., IROS21*]

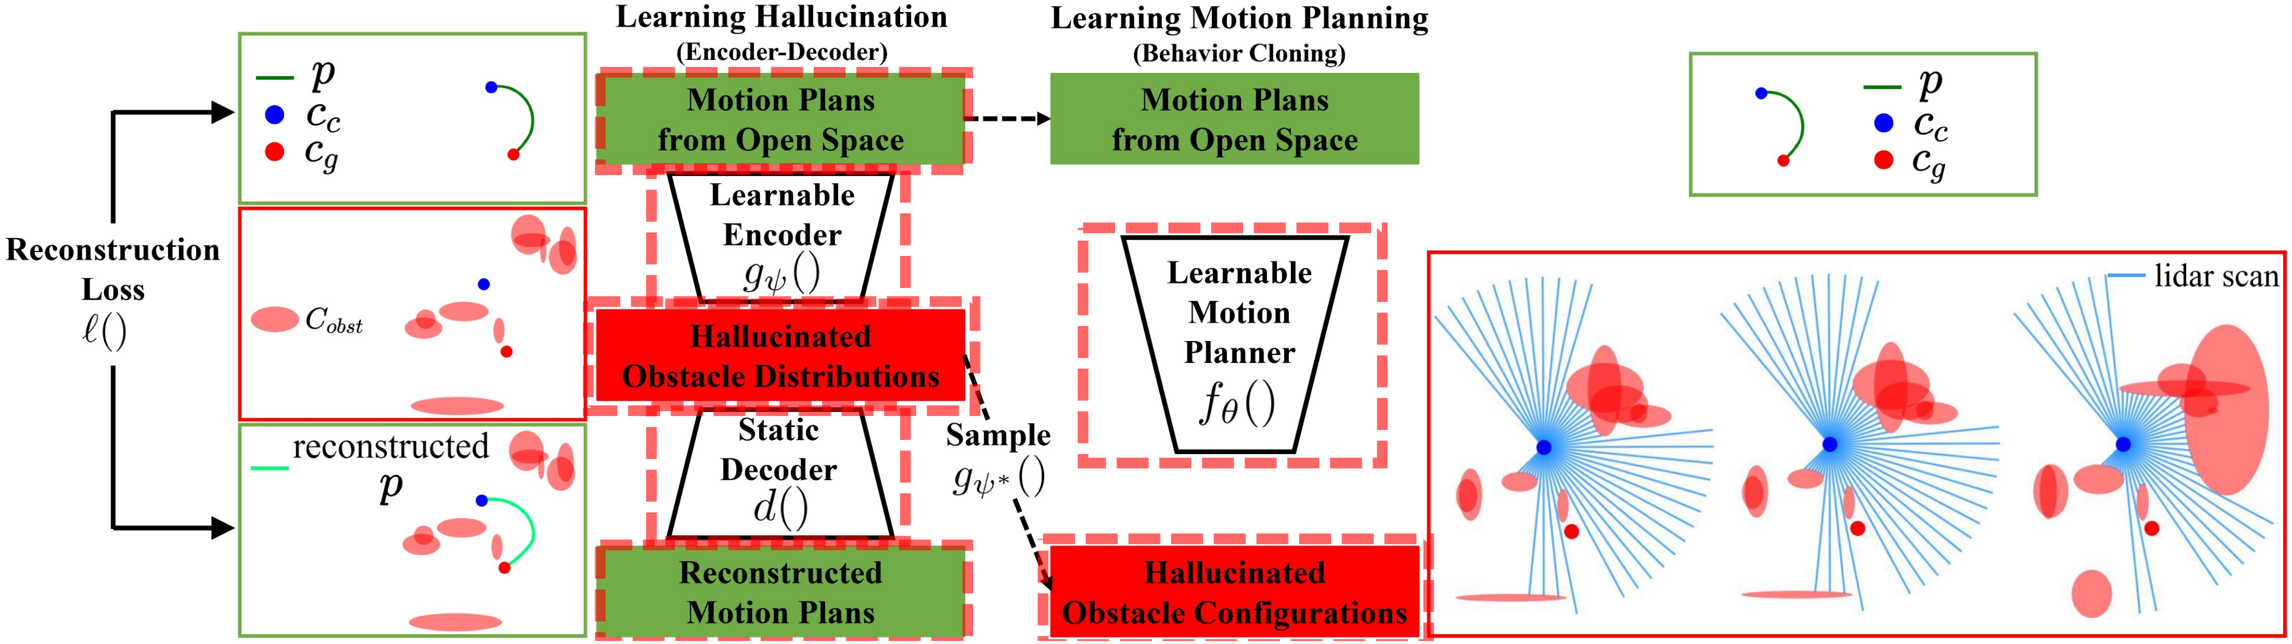

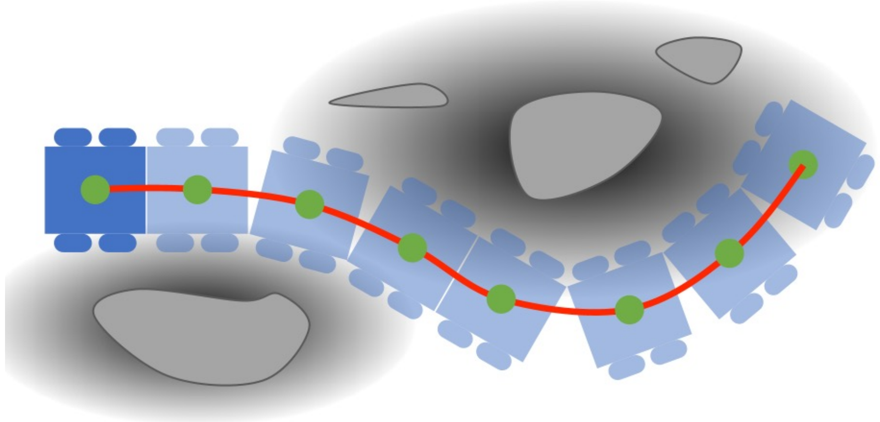

Dyna-LfLH: Learning Agile Navigation in Dynamic Environments from Learned Hallucination

Saad Abdul Ghani, Zizhao Wang, Peter Stone, and Xuesu Xiao

George Mason University and The University of Texas at Austin

The University of Texas at Austin

Learning Navigation in Challenging Environments

Highly-Constrained Environments

[***XX*** *et al., RA-L21a,* ***XX*** *et al., ICRA21,* W, ***XX*** *et al., IROS21*]

Offroad Environments

[***XX*** *et al., RA-L21b, K, S, A, R,* ***XX*** *et al., IROS22, D, P, N,* ***XX****, ICRA24, D, P,* ***XX****, under review*]

Social Environments

[***XX*** *et al., CoRL22, K, N,* ***XX*** *et al., RA-L22, N, N, P, D,* ***XX****, IROS23*]

Learning Inverse Kinodynamics [*XX et al., RA-L21b*]

- **Objective:** Navigate a mobile robot to track a reference trajectory during deployment as quickly and as accurately as possible.

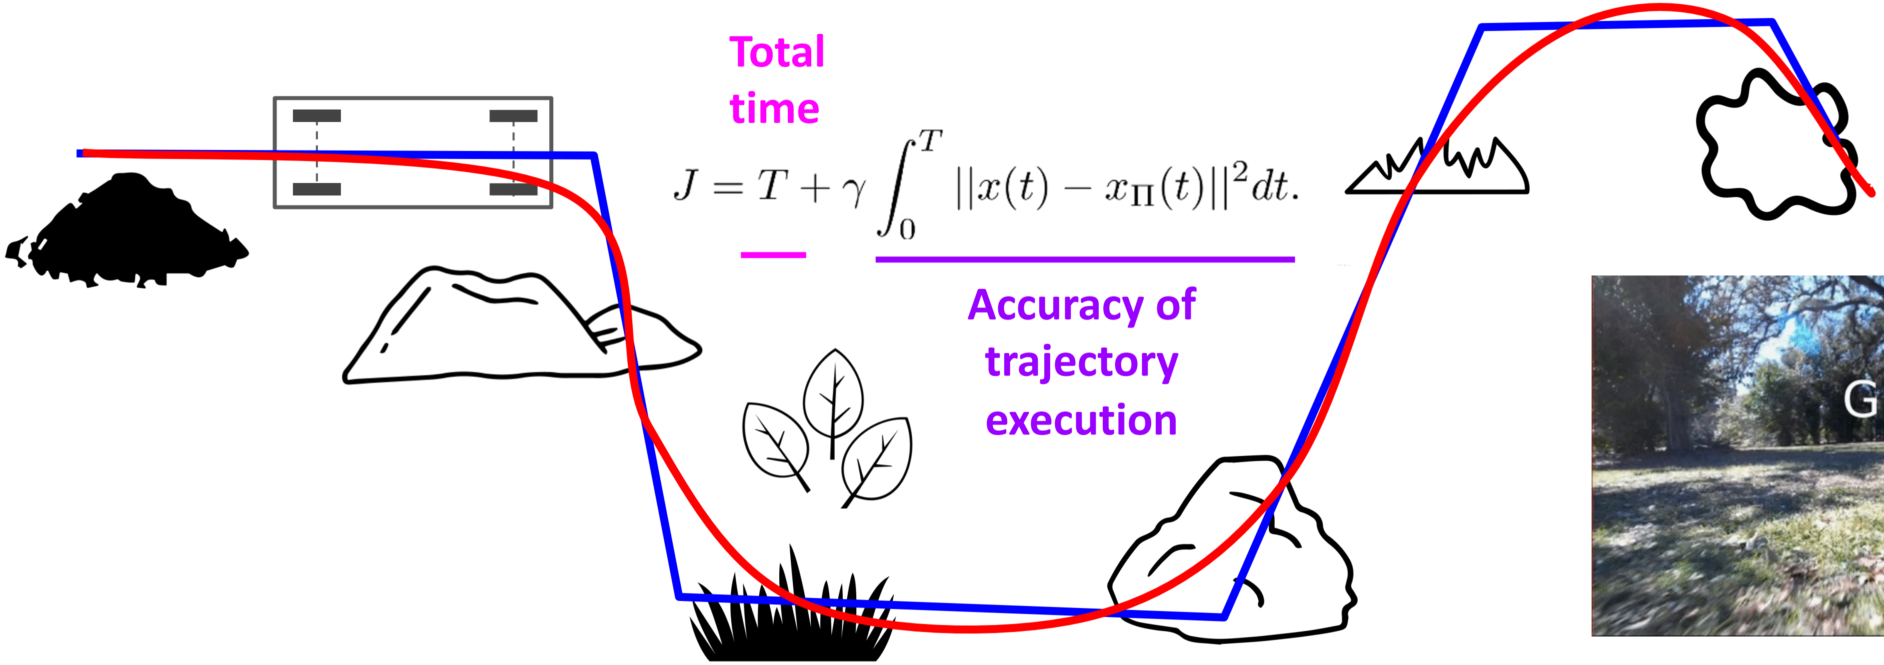

$$J = T + \gamma \int_0^T ||x(t) - x_\Pi(t)||^2 dt.$$

Learning Inverse Kinodynamics [*XX et al., RA-L21b*]

Problem: kinodynamic function depends on unknown world states!

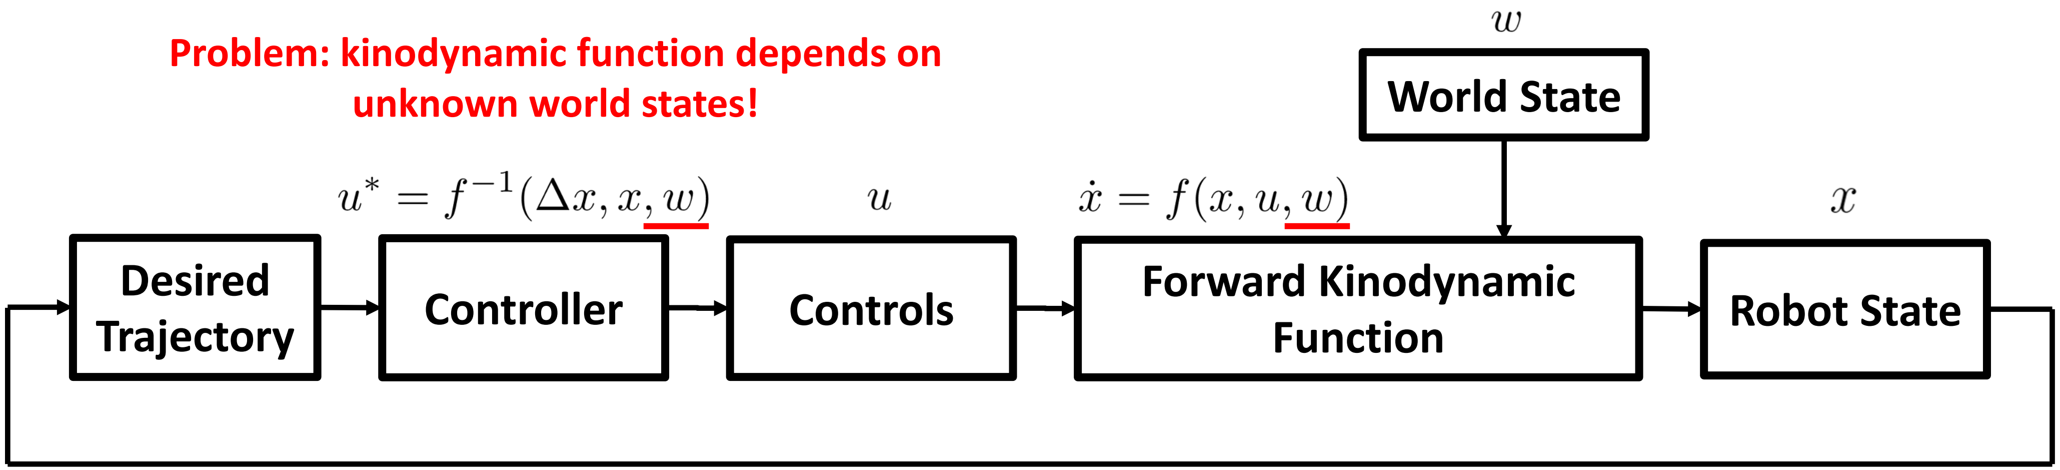

Learning Inverse Kinodynamics *[XX et al., RA-L21b]*

Our Approach

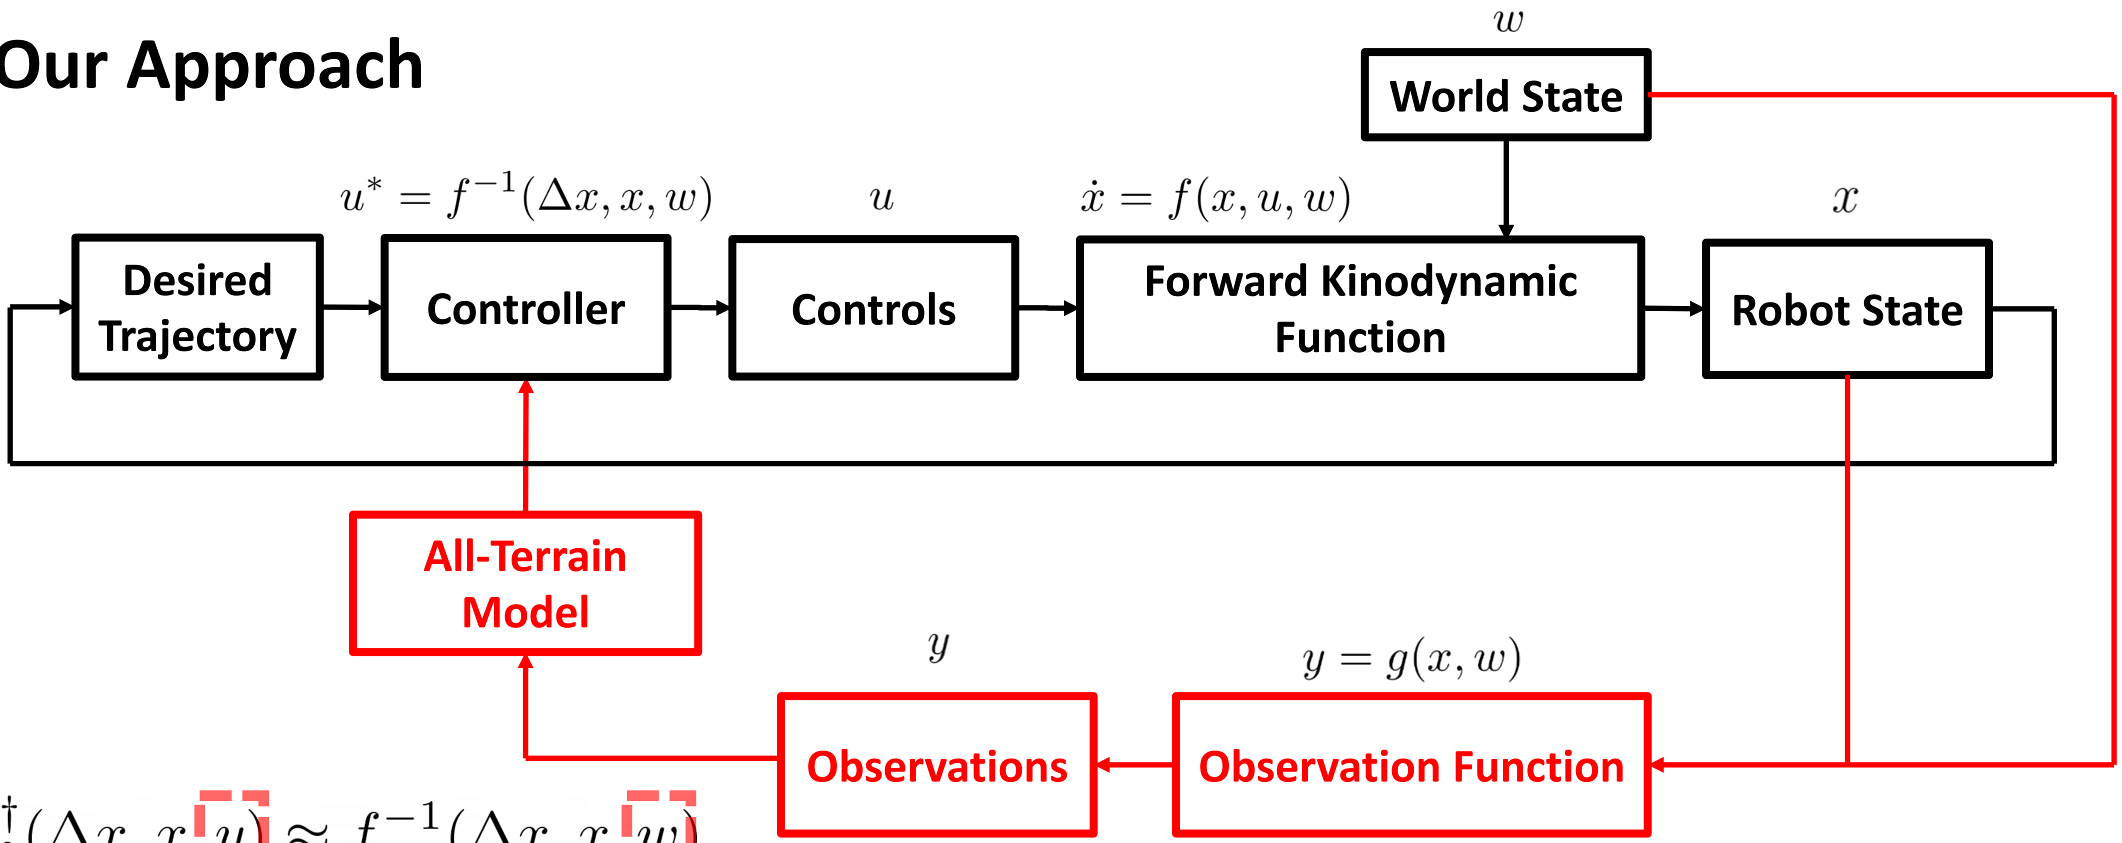

$$f_\theta^\dagger(\Delta x, x, \boxed{y}) \approx f^{-1}(\Delta x, x, \boxed{w})$$

Learning IMU-IKD [*XX et al., RA-L21b*]

Implementation

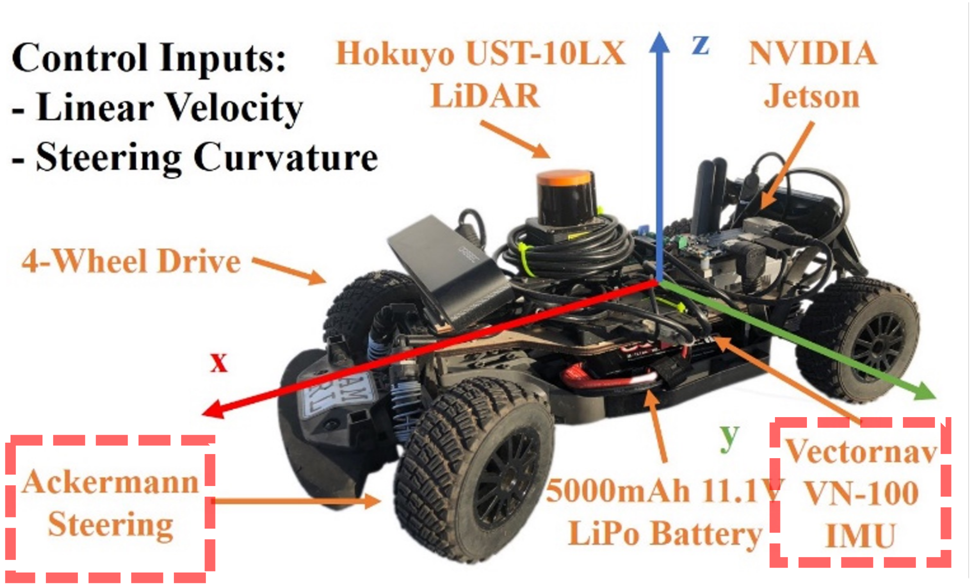

UT Automata

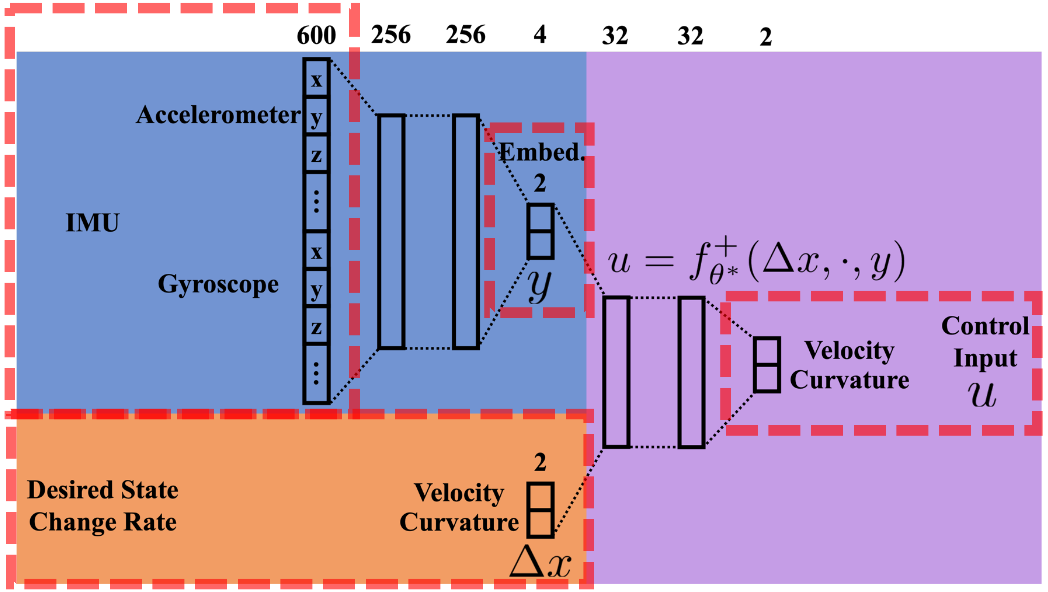

Neural Network Architecture

Learning IMU-IKD [*XX et al., RA-L21b*]

Baseline Ablation IMU-IKD Reference

Seen Terrain

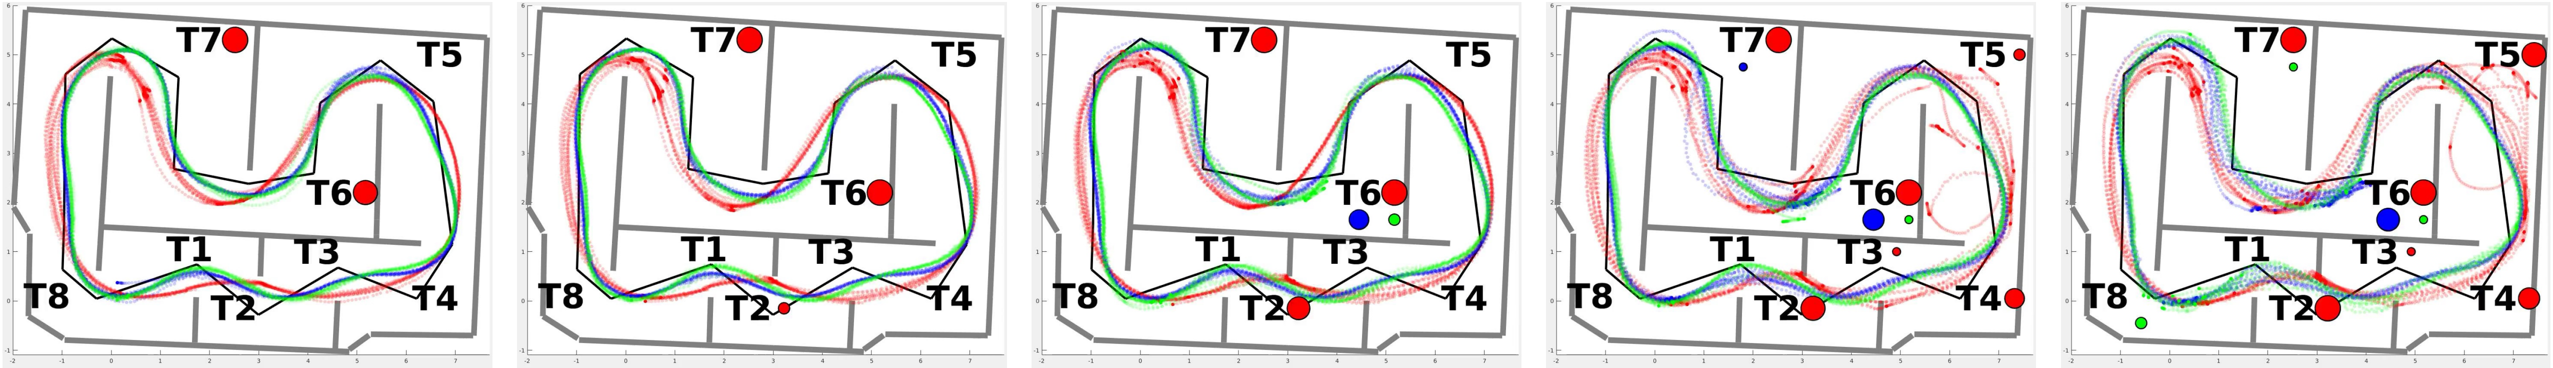

1.6m/s 1.7m/s 1.8m/s 1.9m/s 2.0m/s

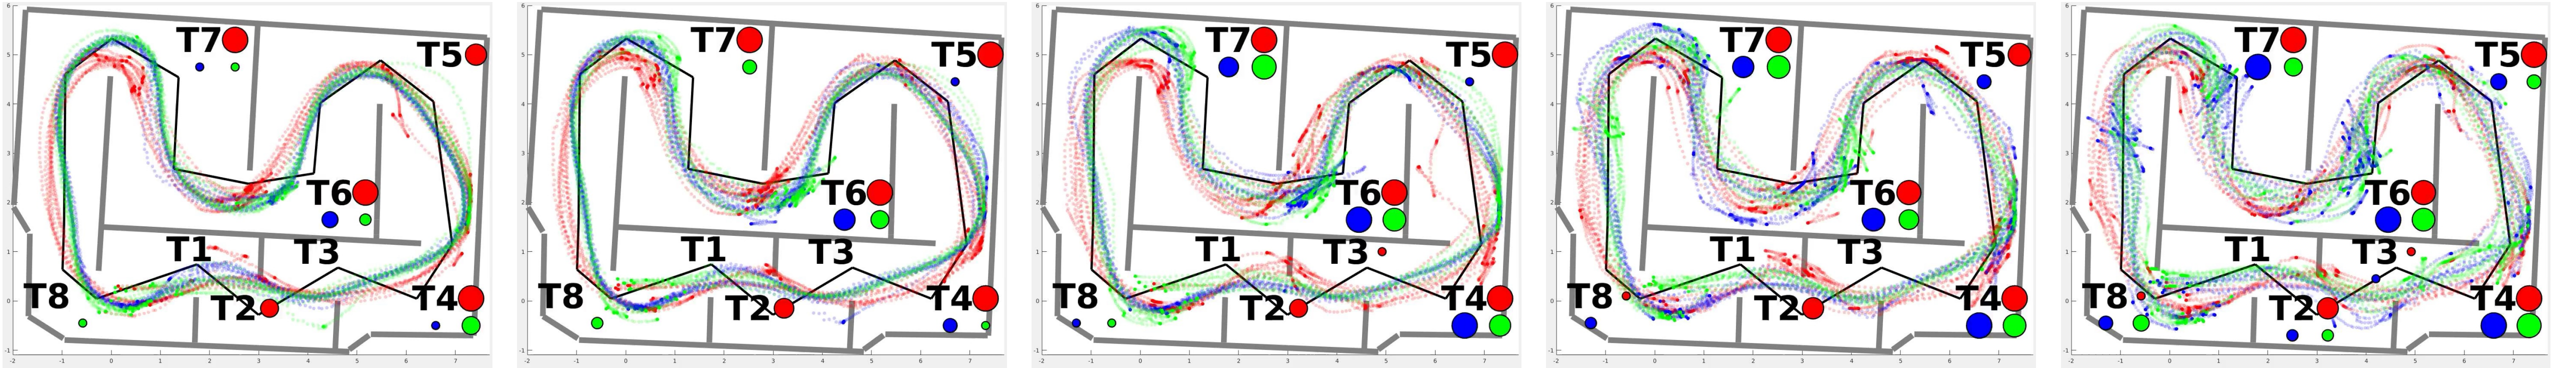

2.1m/s 2.2m/s 2.3m/s 2.4m/s 2.5m/s

Learning IMU-IKD [*XX et al., RA-L21b*]

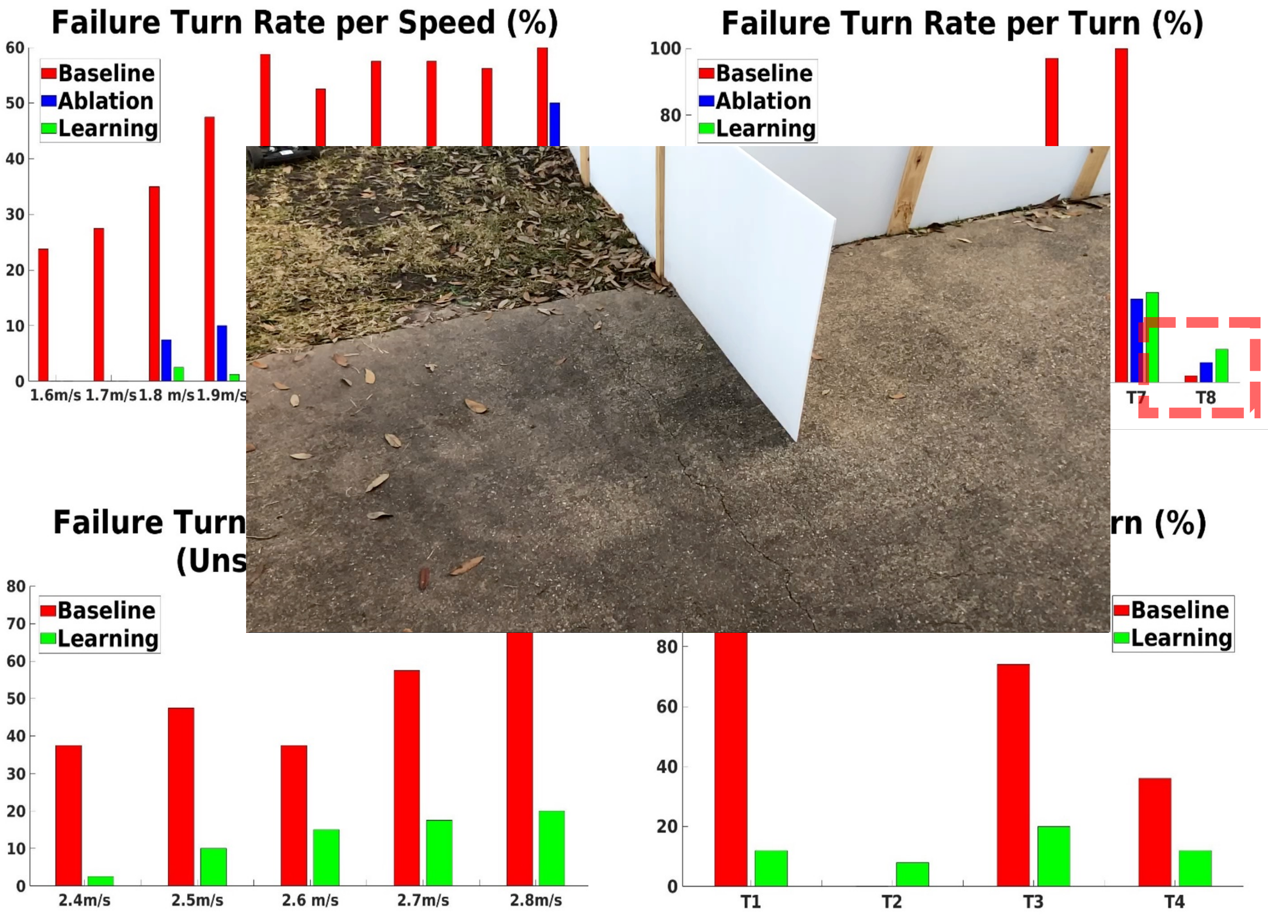

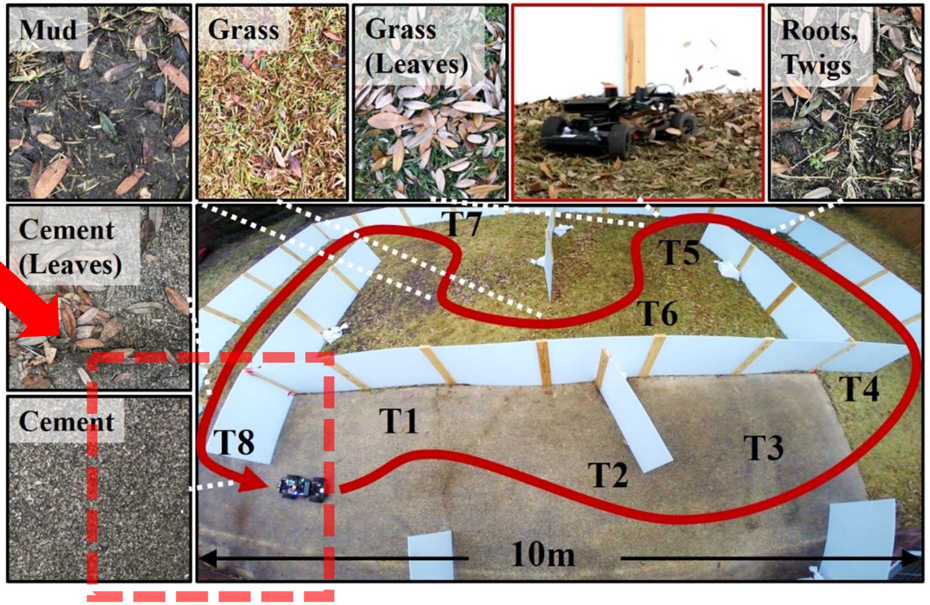

Learning Visual-Inertial IKD (VI-IKD) *[K, S, A, R, XX et al., IROS22]*

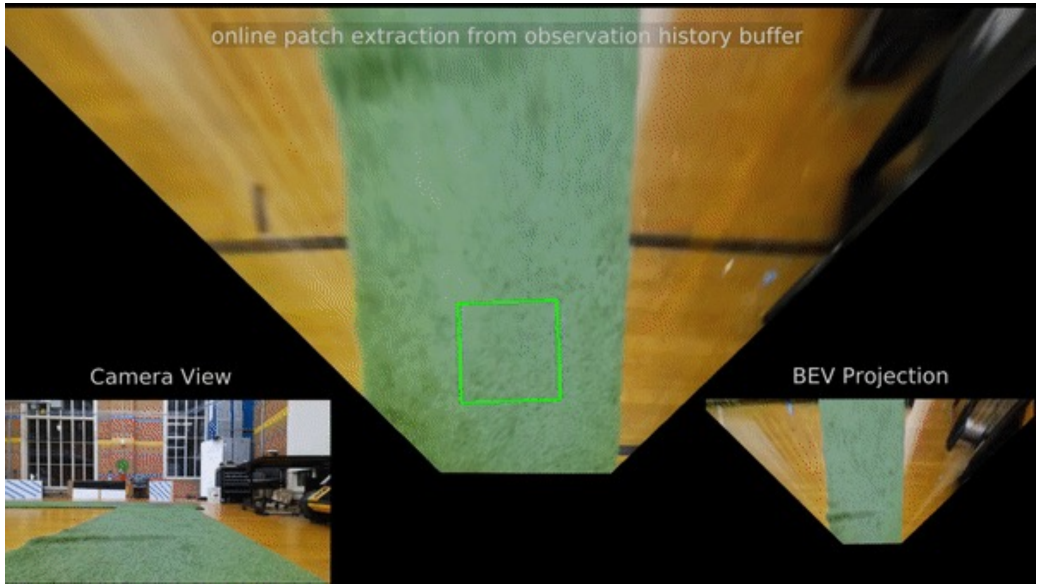

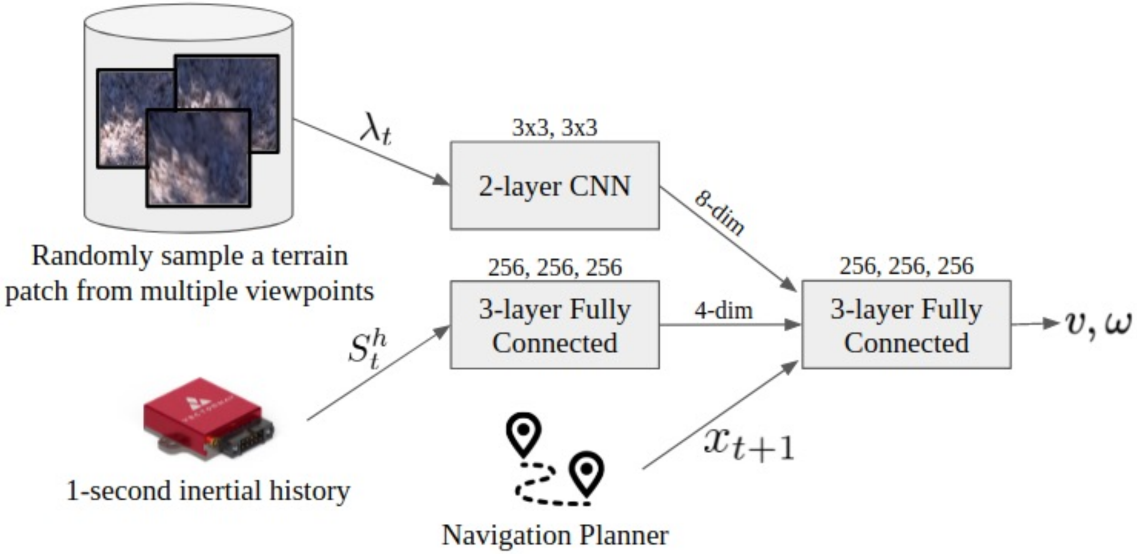

Image Patch Extraction to Anticipate Upcoming Kinodynamic Changes

Adding Vision in Addition to Inertia in the Observation y for Better Representation of Unknown World State w


IMU-IKD model: uses only inertial information (no anticipation)

Off-Road Challenges

From High-Speed to Vertically Challenging Terrain

The Verti-Wheelers Project

[*D, P, N, **XX**, ICRA24*]

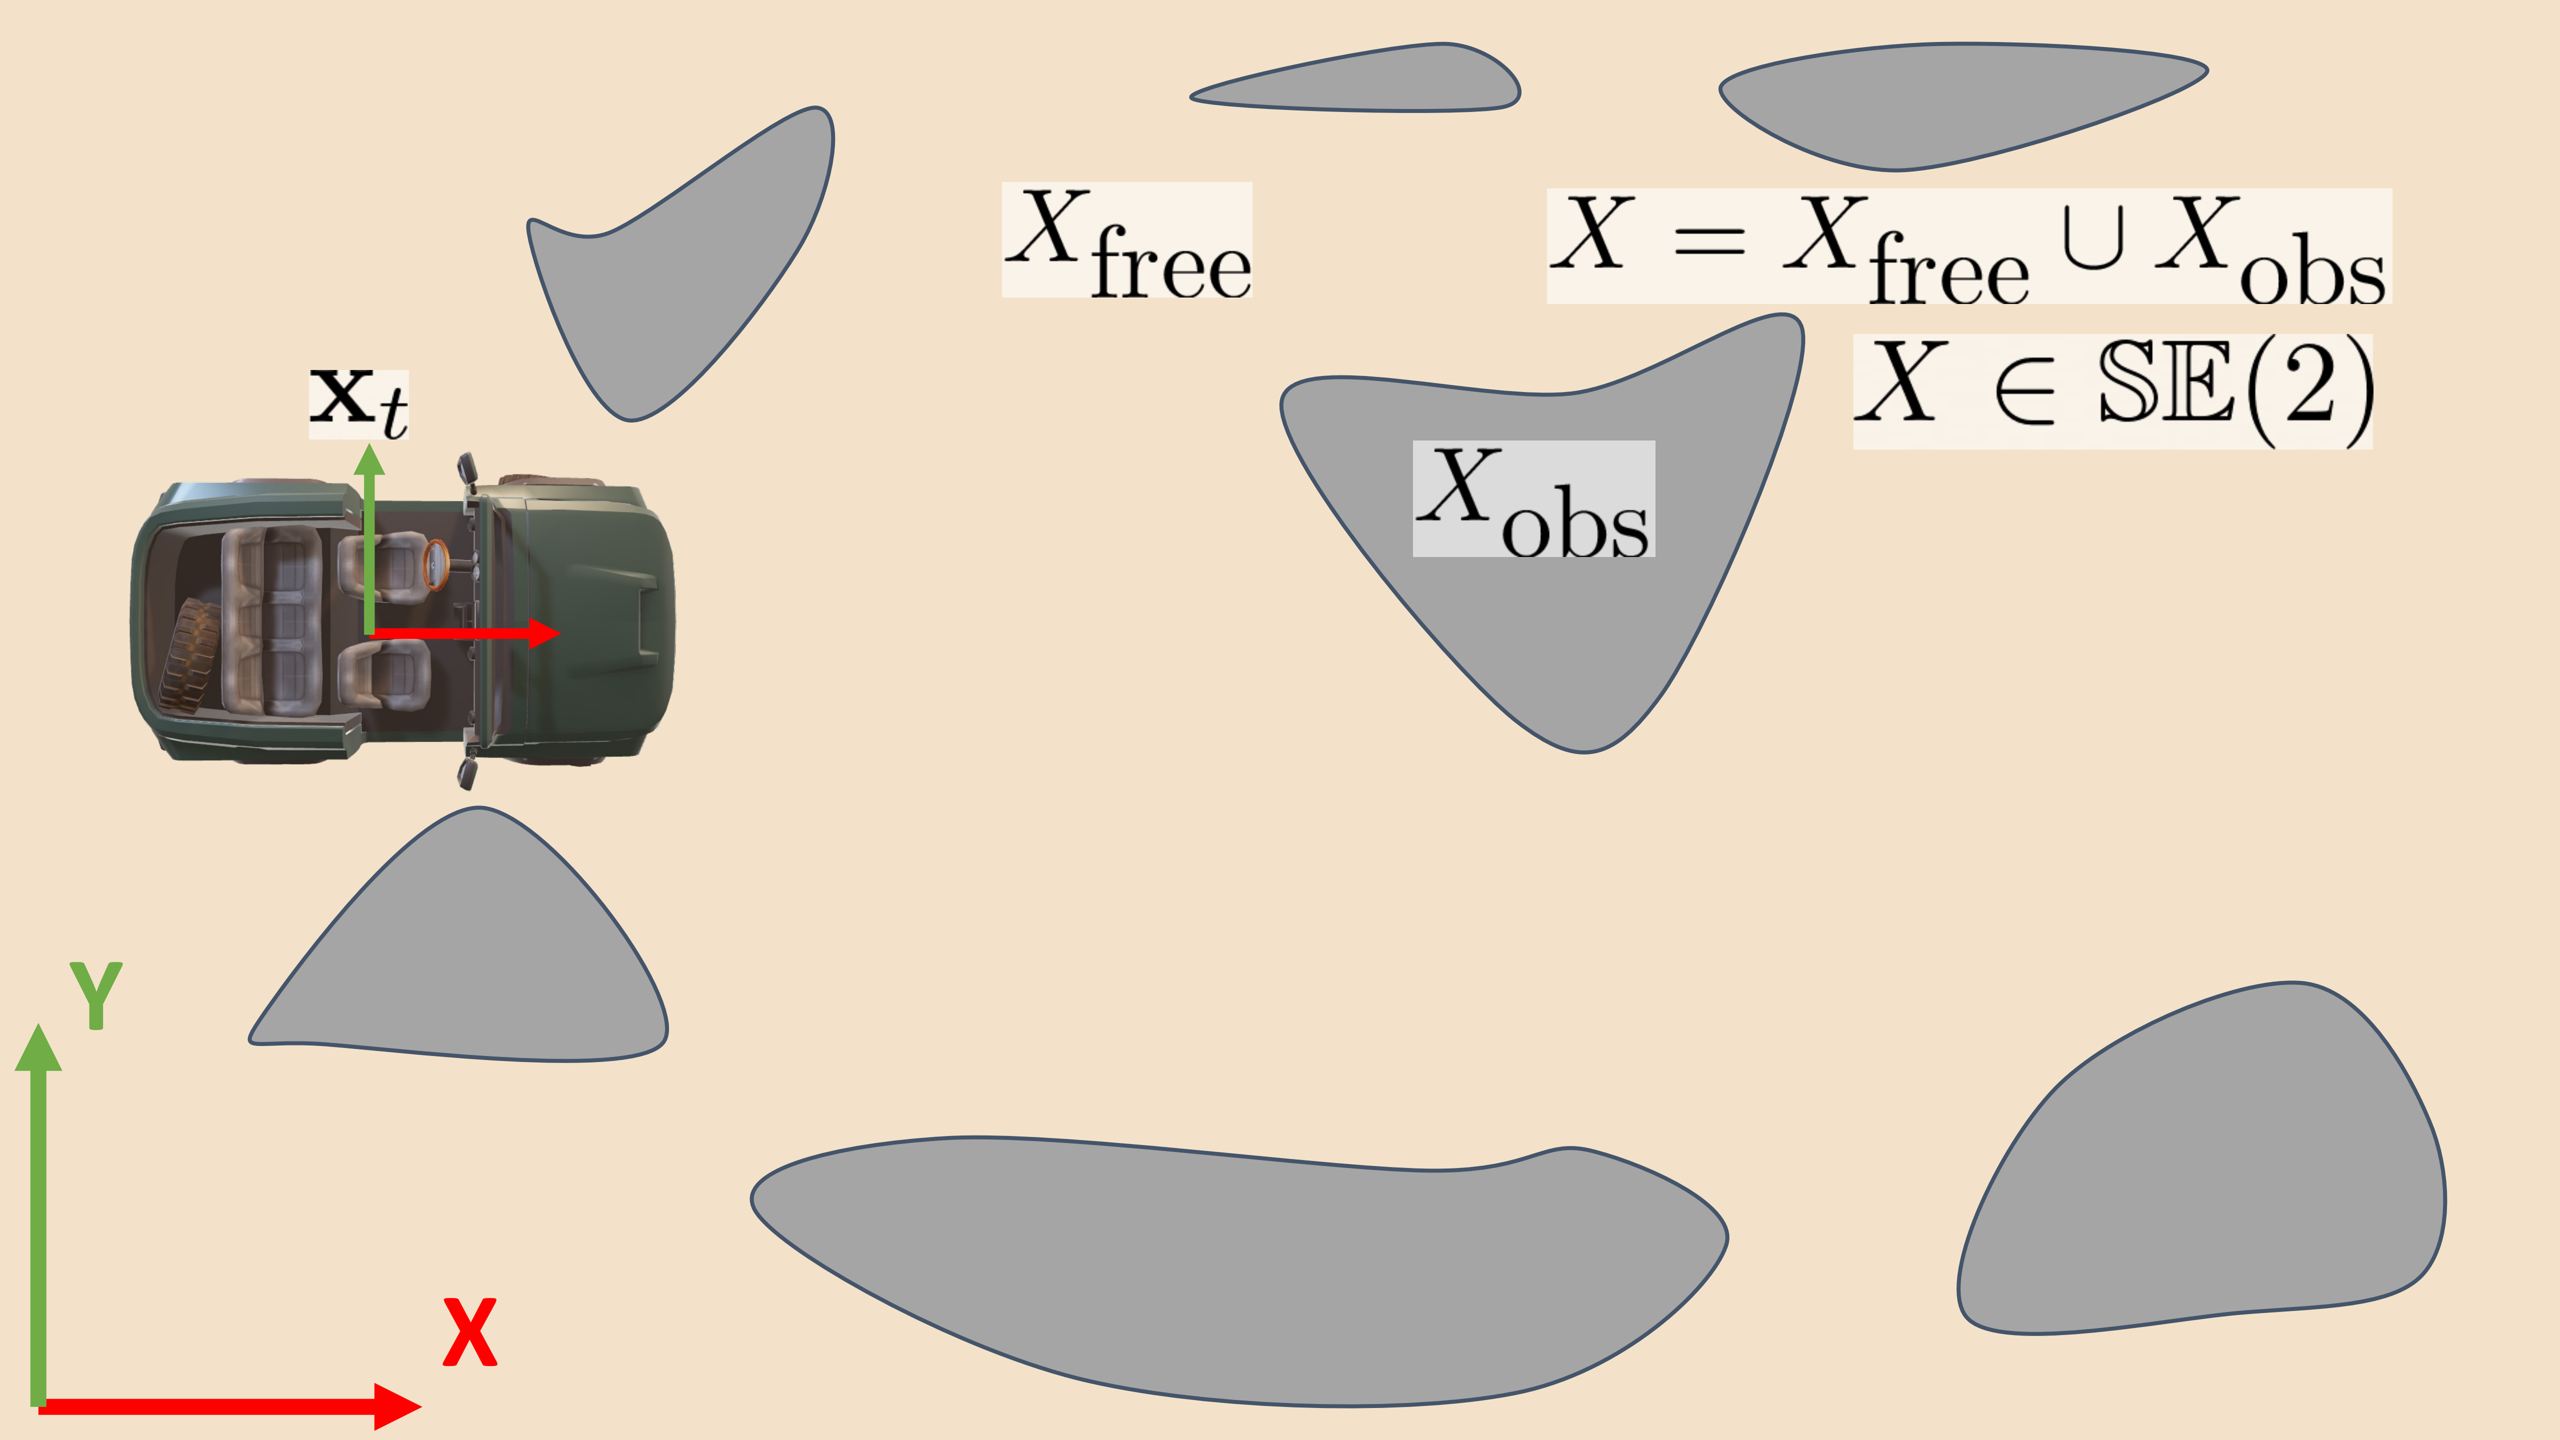

X_{free}
$X = X_{\text{free}} \cup X_{\text{obs}}$
$X \in \mathbb{SE}(2)$
$\mathbf{x}_t$
X_{obs}
Y
X

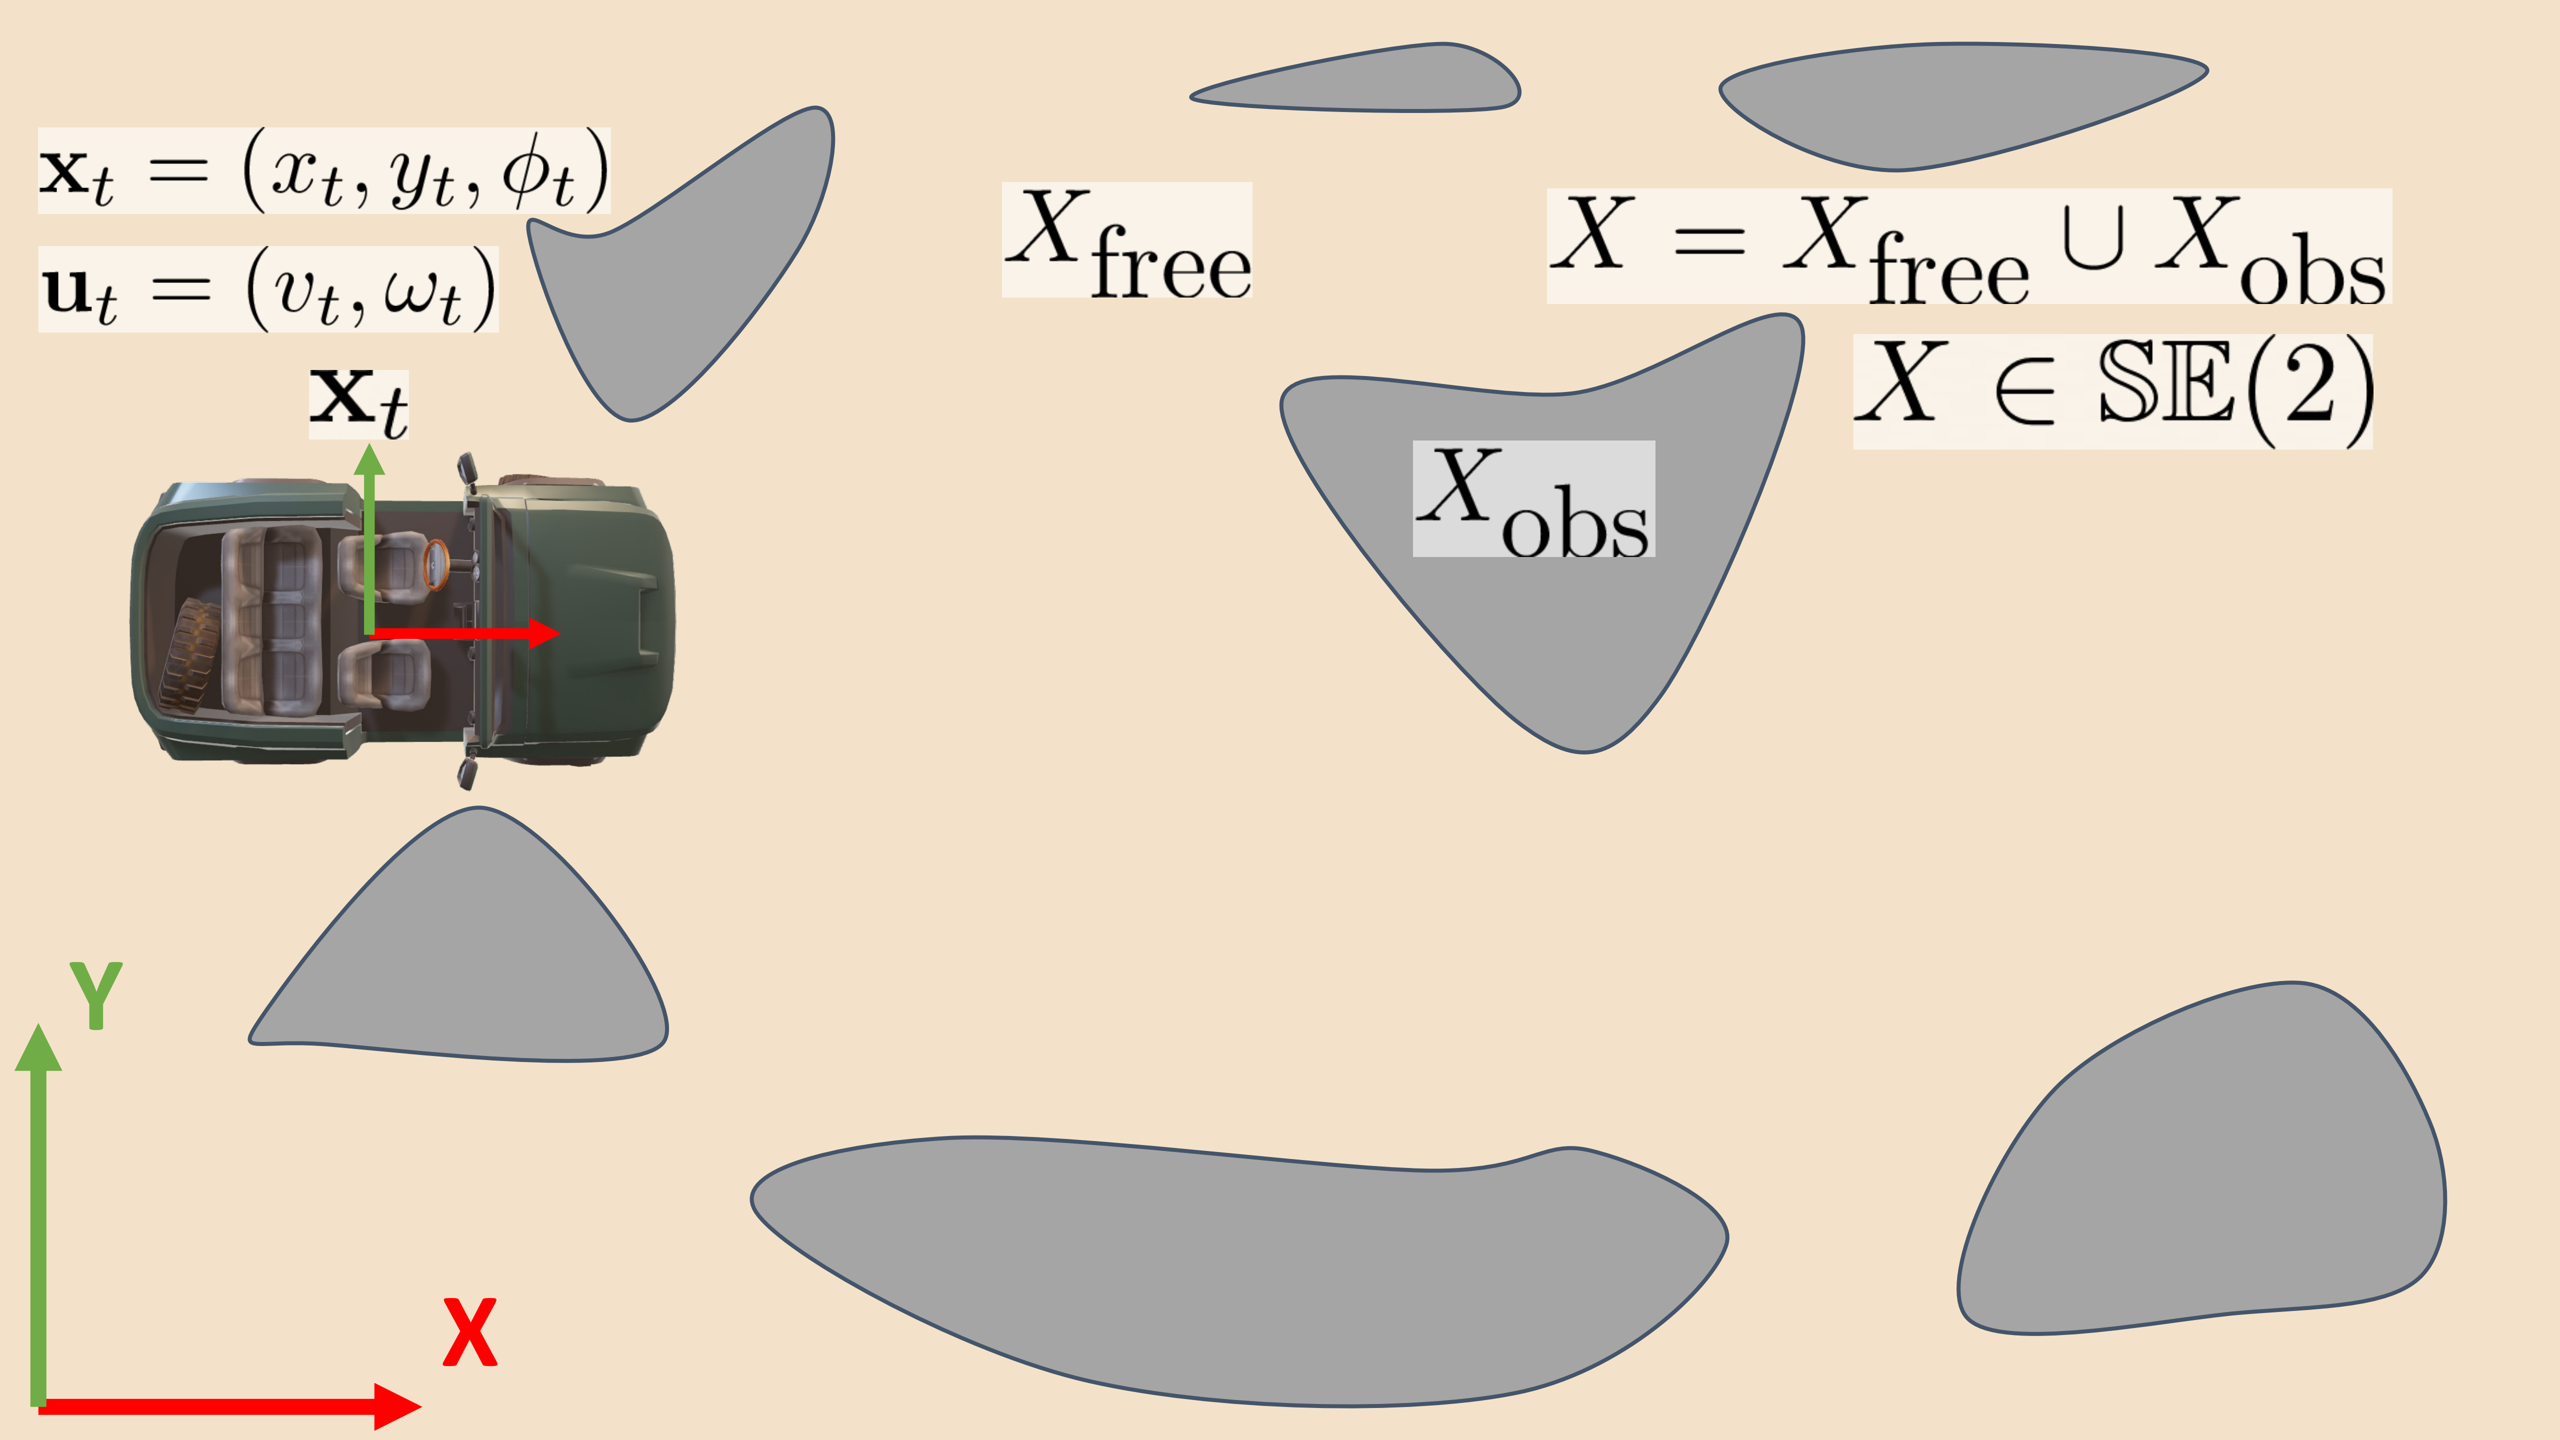

$\mathbf{x}_t = (x_t, y_t, \phi_t)$
$\mathbf{u}_t = (v_t, \omega_t)$
$\mathbf{x}_t$
X_{free}
$X = X_{\text{free}} \cup X_{\text{obs}}$
$X \in \mathbb{SE}(2)$
X_{obs}
Y
X

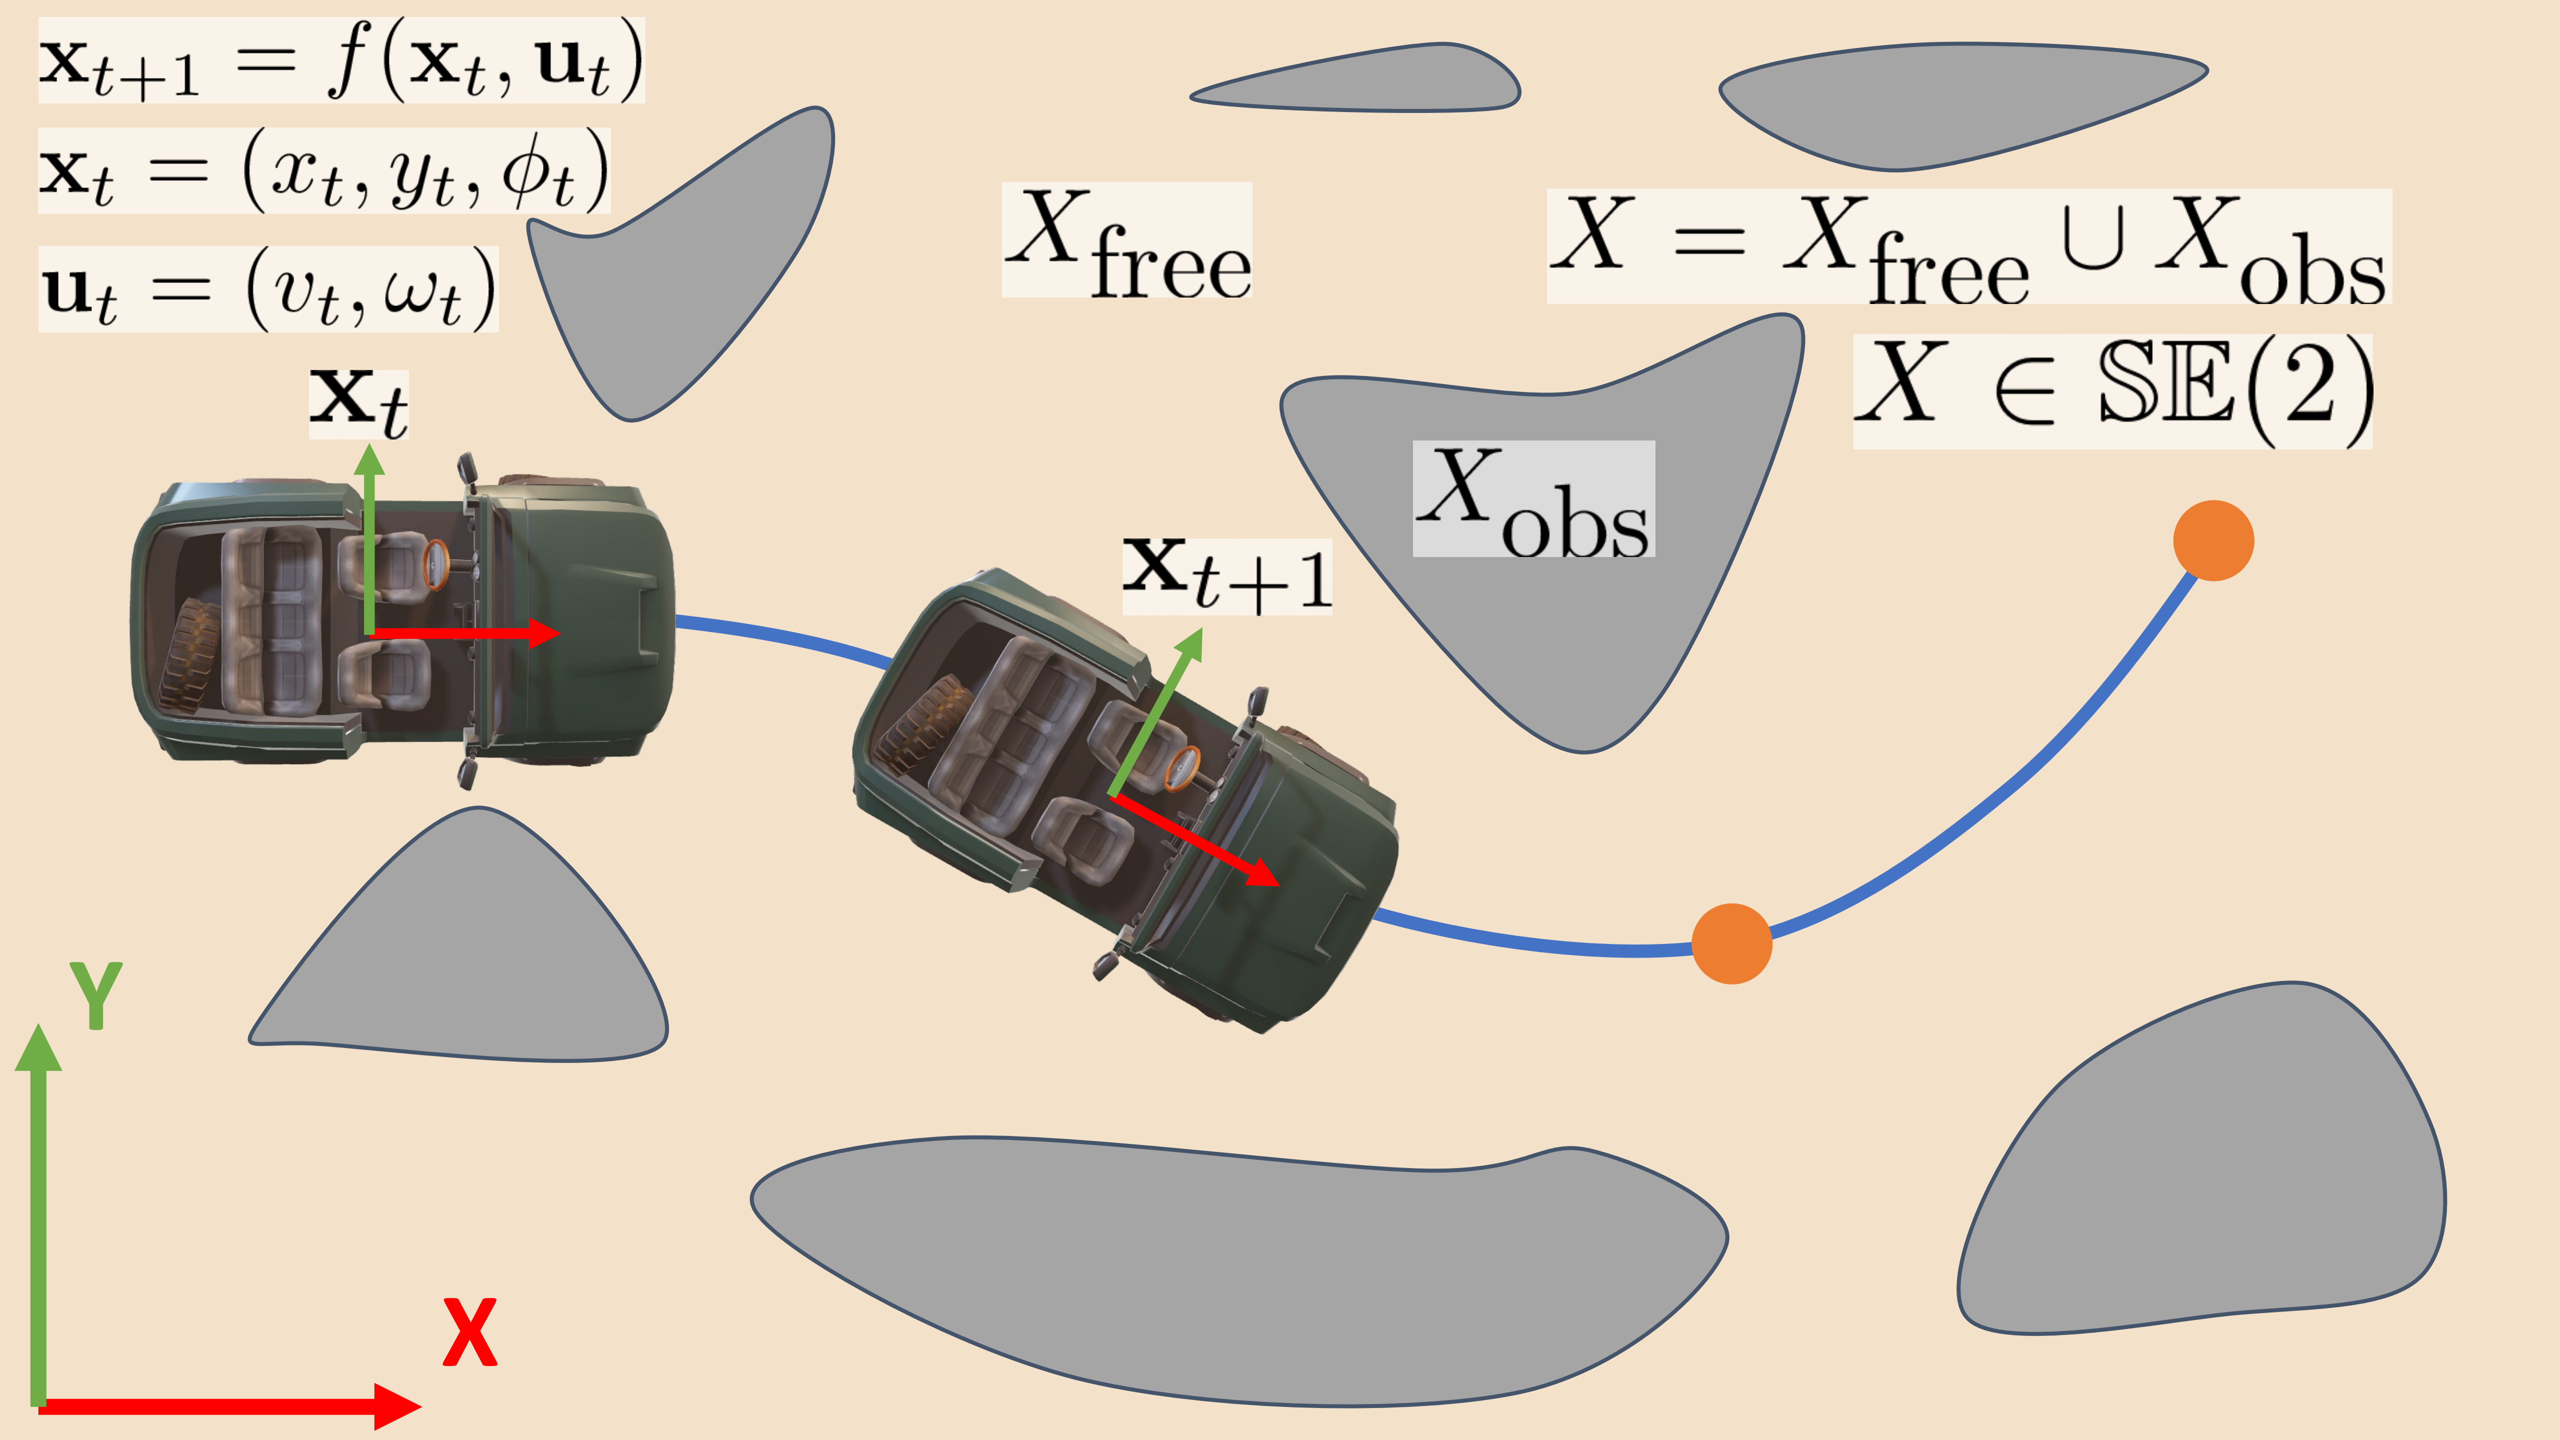

$\mathbf{x}_{t+1} = f(\mathbf{x}_t, \mathbf{u}_t)$
$\mathbf{x}_t = (x_t, y_t, \phi_t)$
$\mathbf{u}_t = (v_t, \omega_t)$
X_{free}
$X = X_{\text{free}} \cup X_{\text{obs}}$
$X \in \mathbb{SE}(2)$
X_{obs}
$\mathbf{x}_t$
$\mathbf{x}_{t+1}$
Y
X

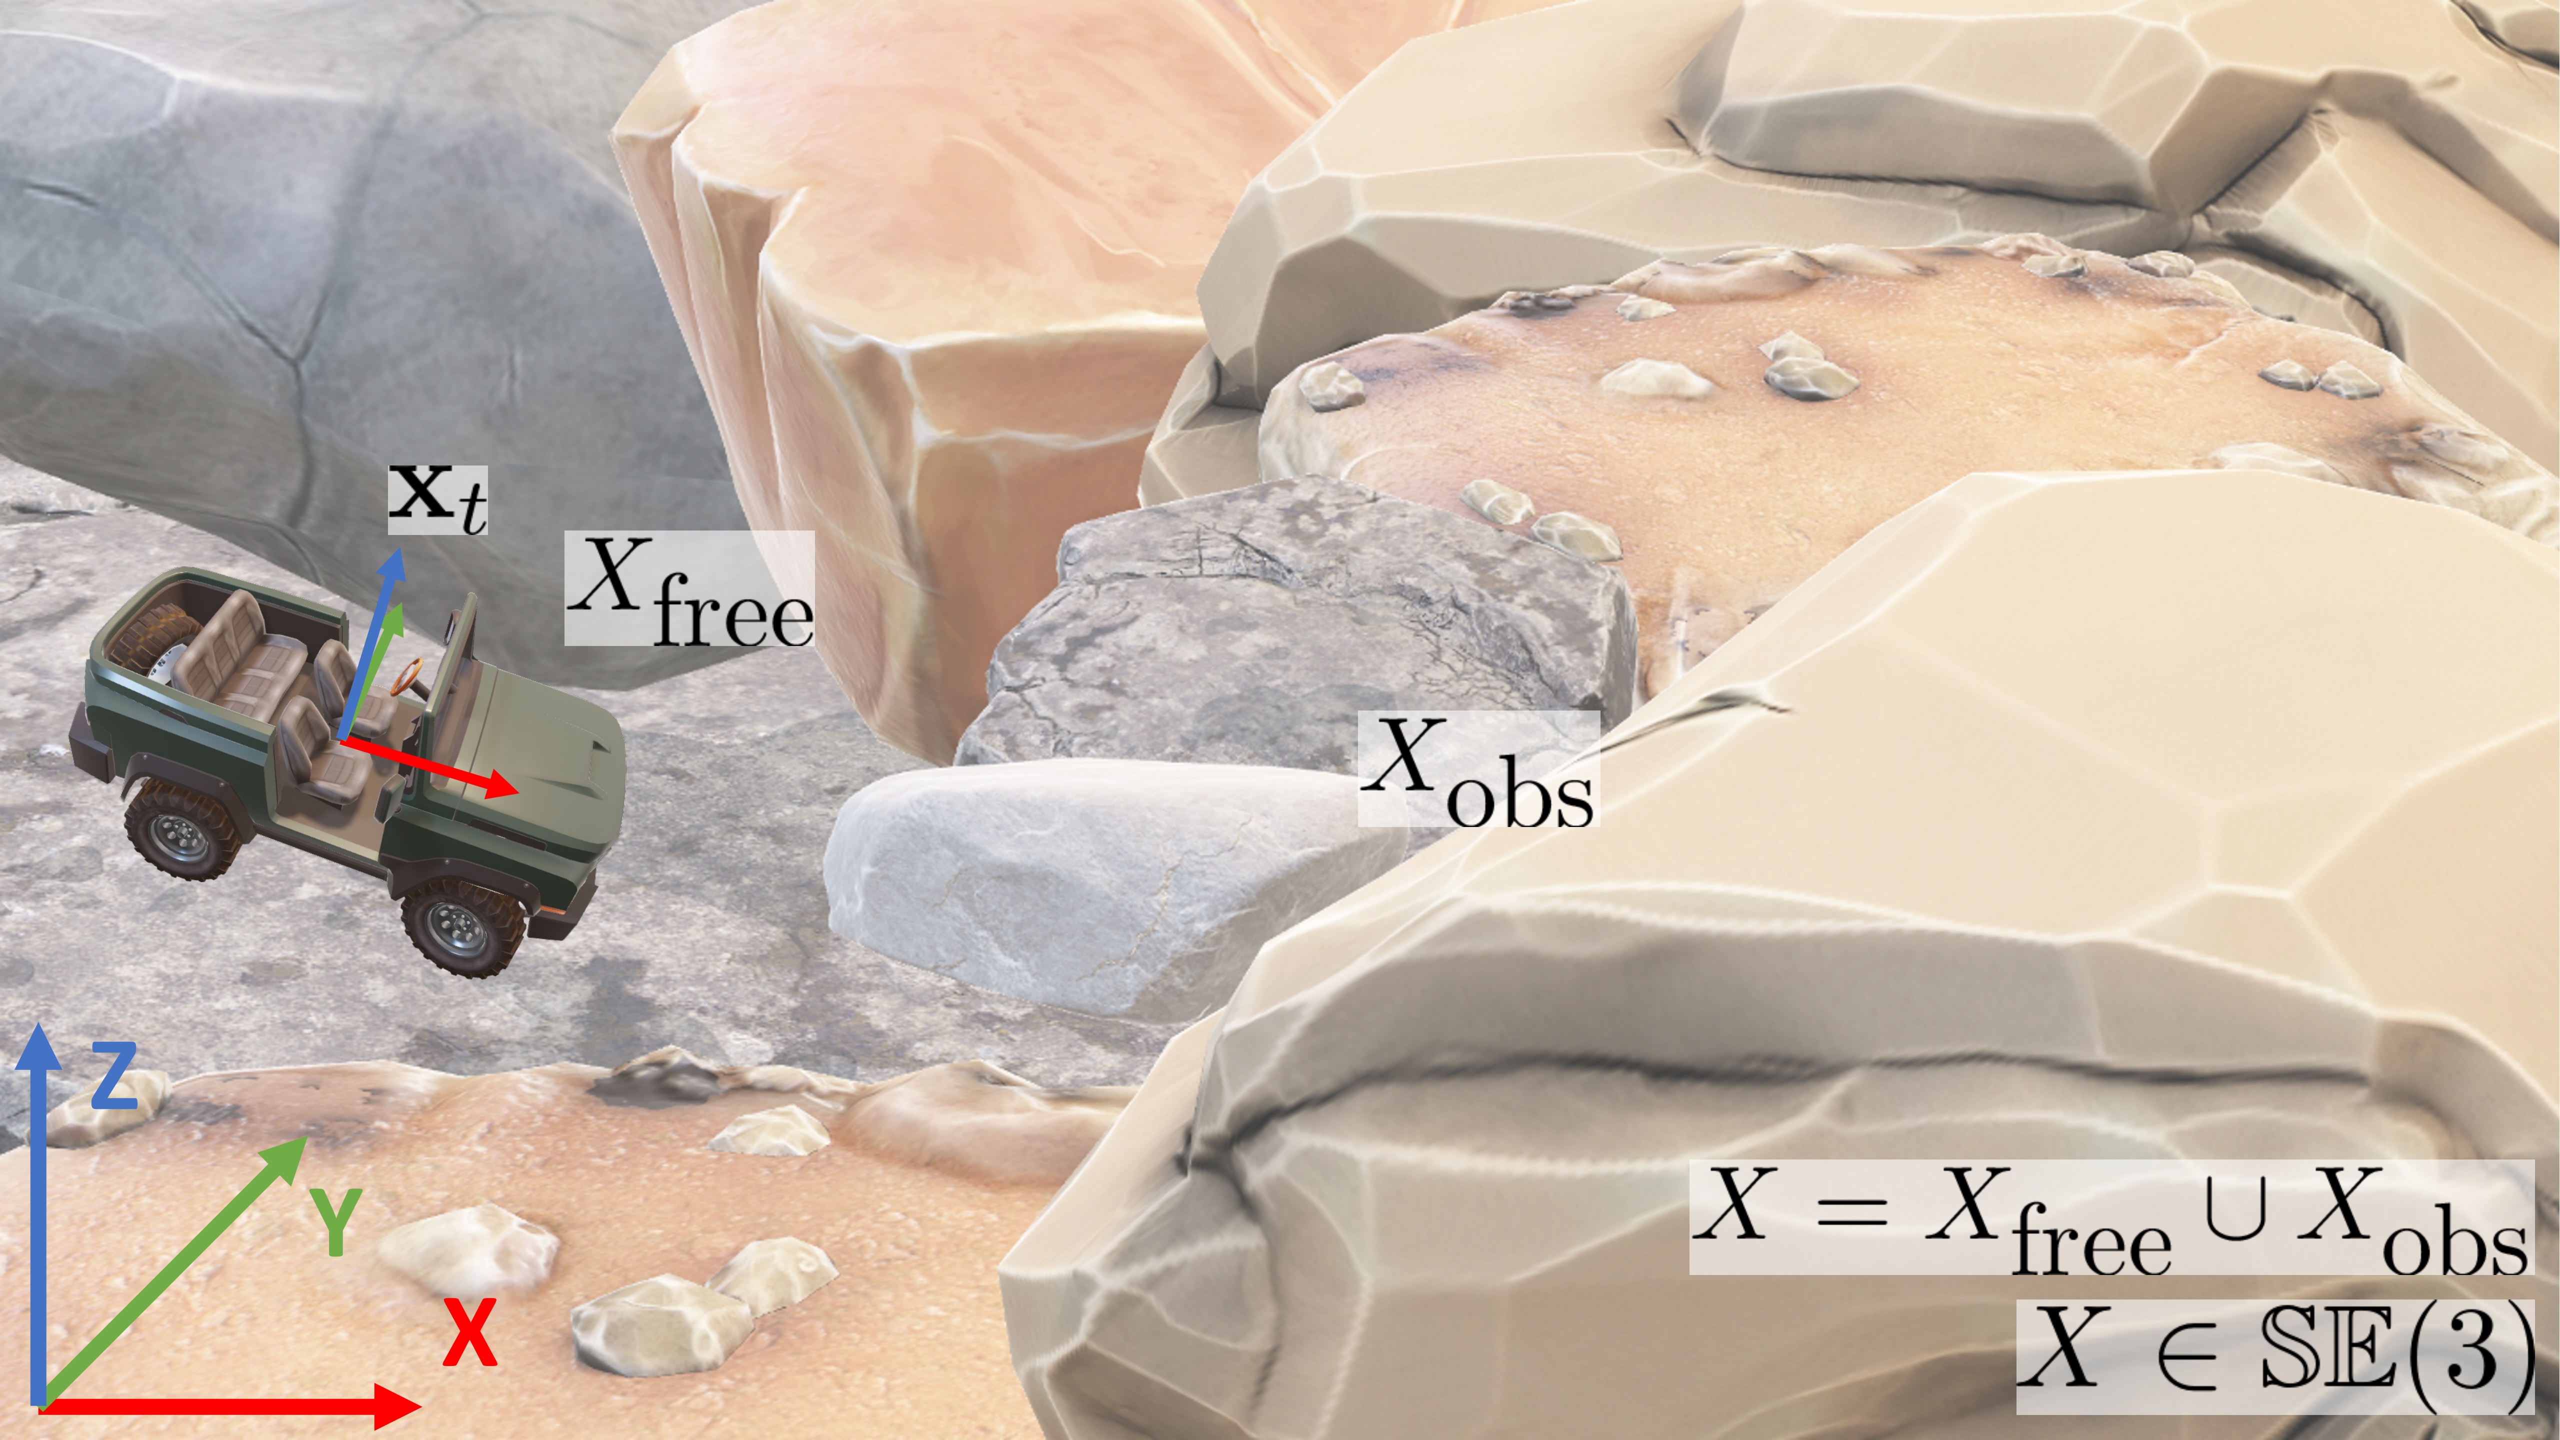

$\mathbf{x}_t$
X_{free}
X_{obs}
Z
Y
X
$X = X_{\text{free}} \cup X_{\text{obs}}$
$X \in \mathbb{SE}(3)$

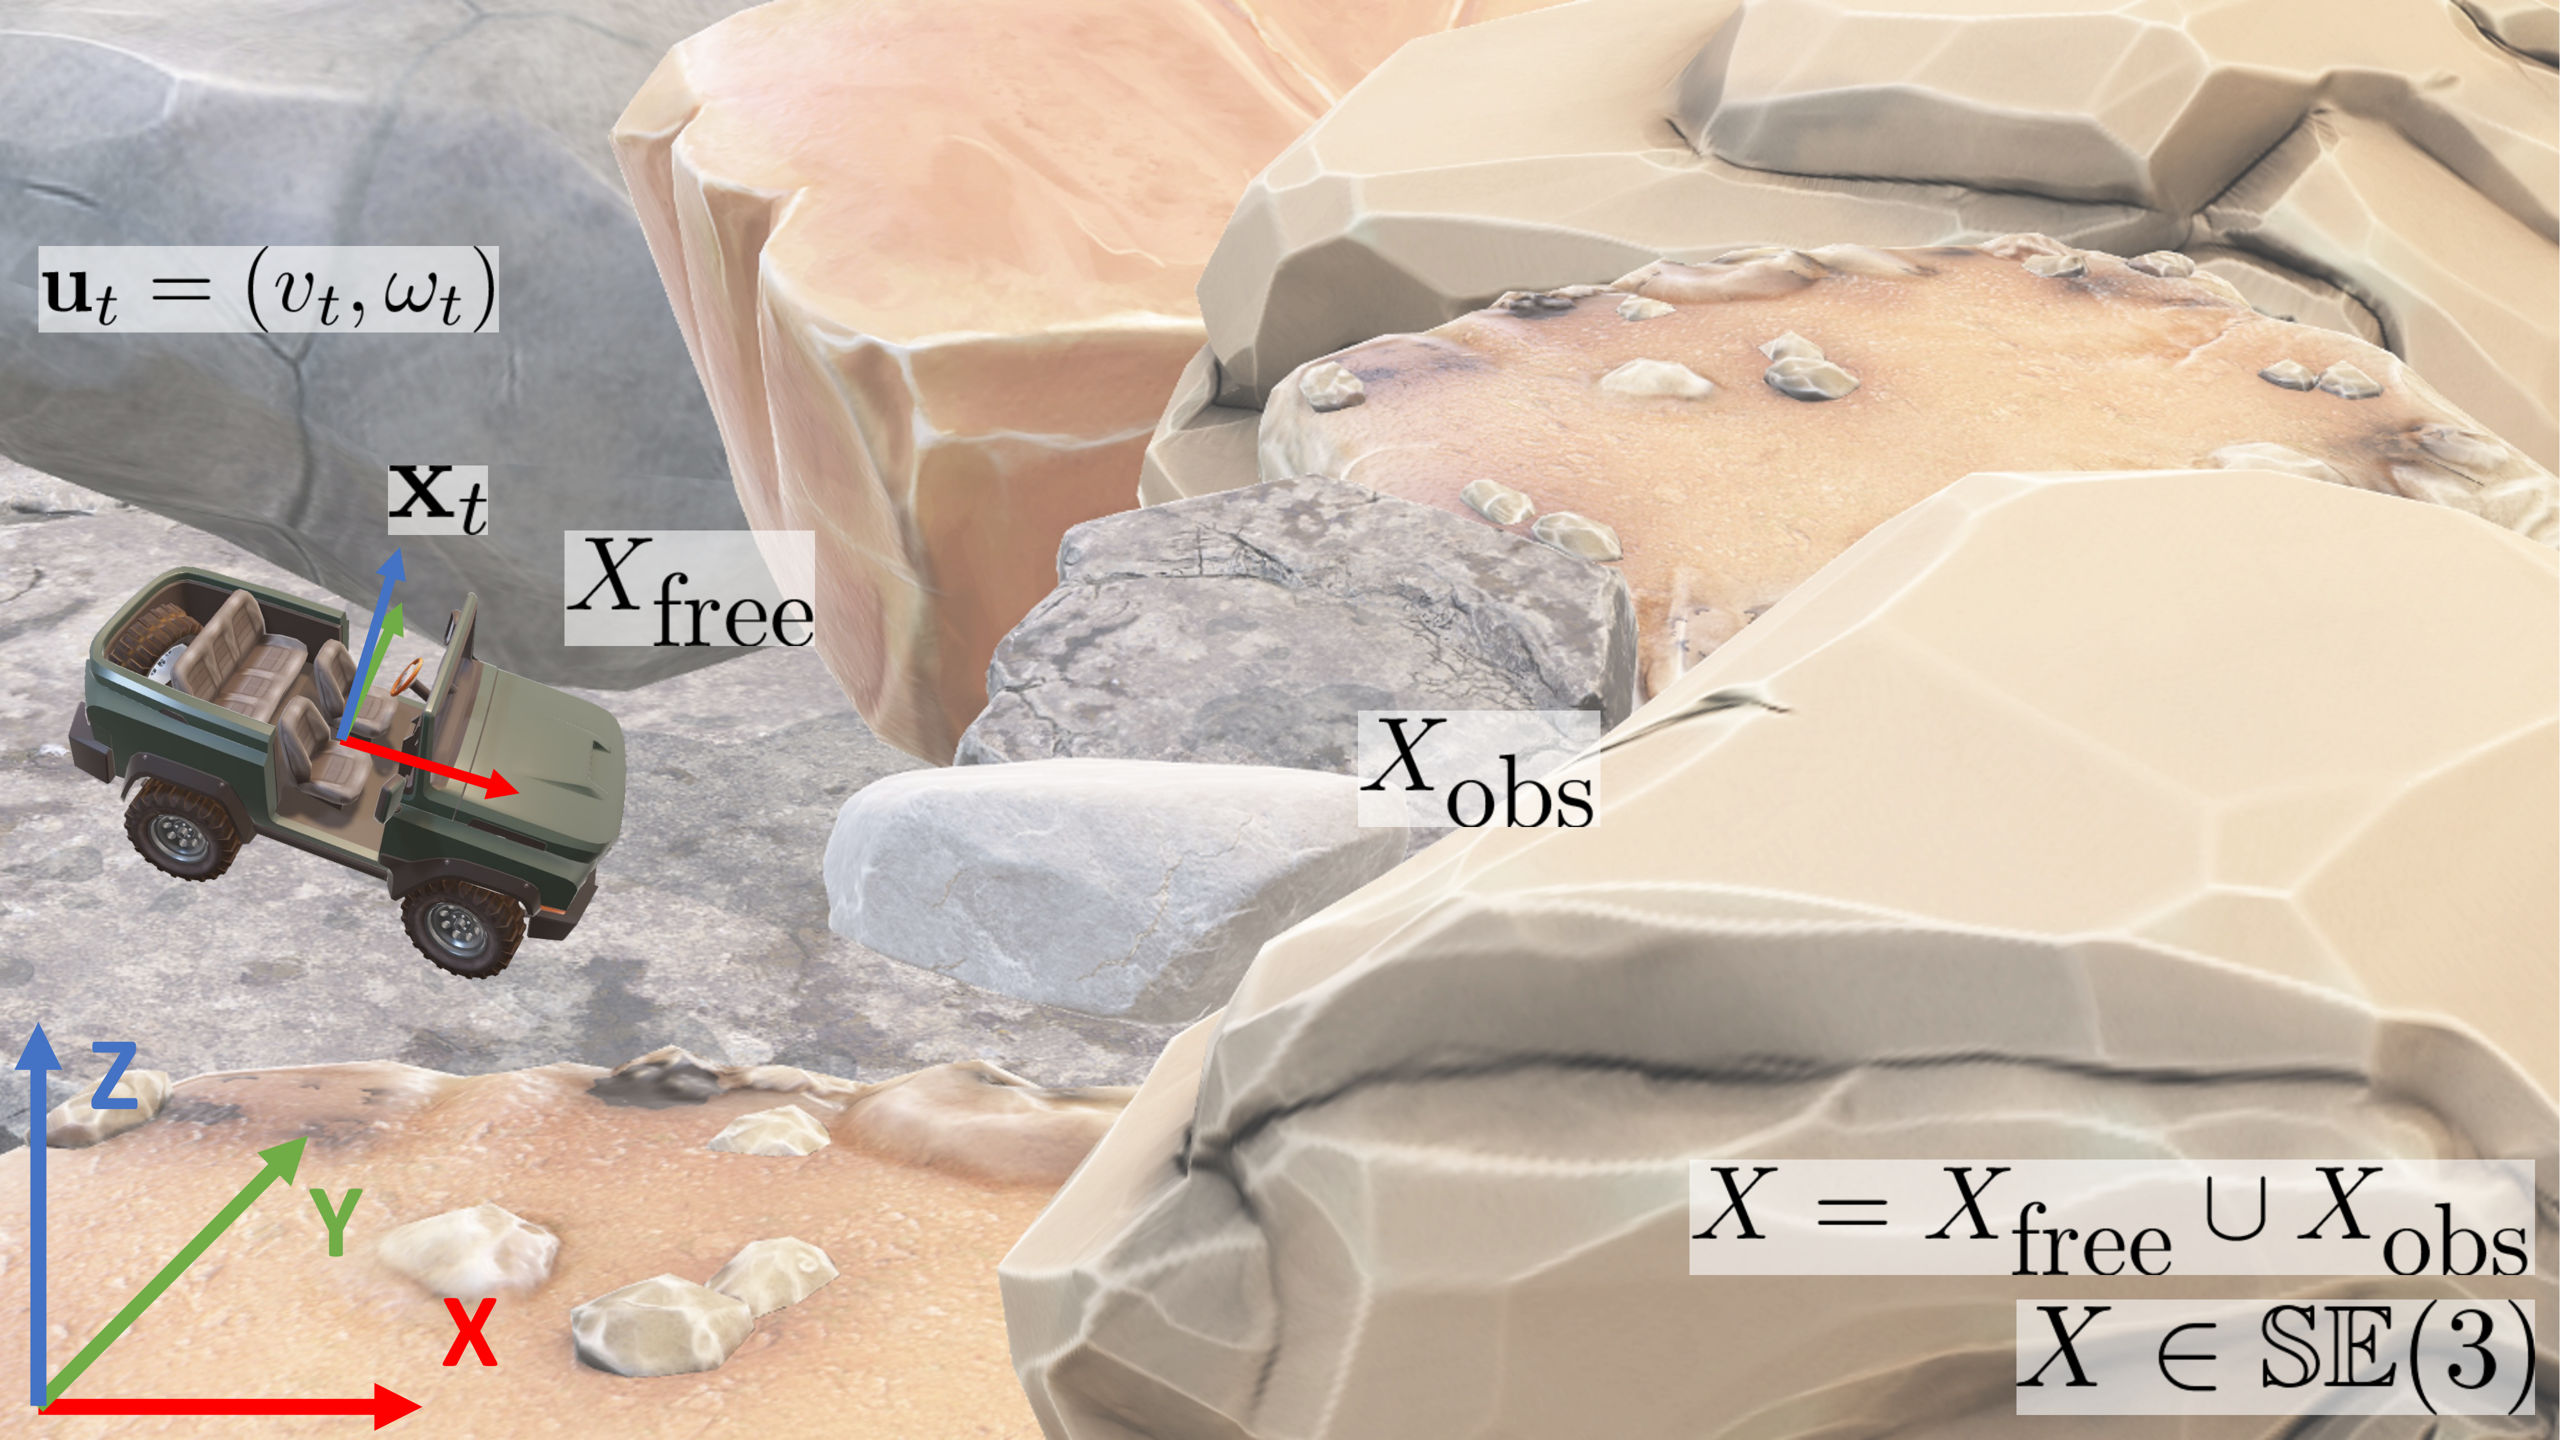

$\mathbf{u}_t = (v_t, \omega_t)$
$\mathbf{x}_t$
X_{free}
X_{obs}
Z
Y
X
$X = X_{\text{free}} \cup X_{\text{obs}}$
$X \in \mathbb{SE}(3)$

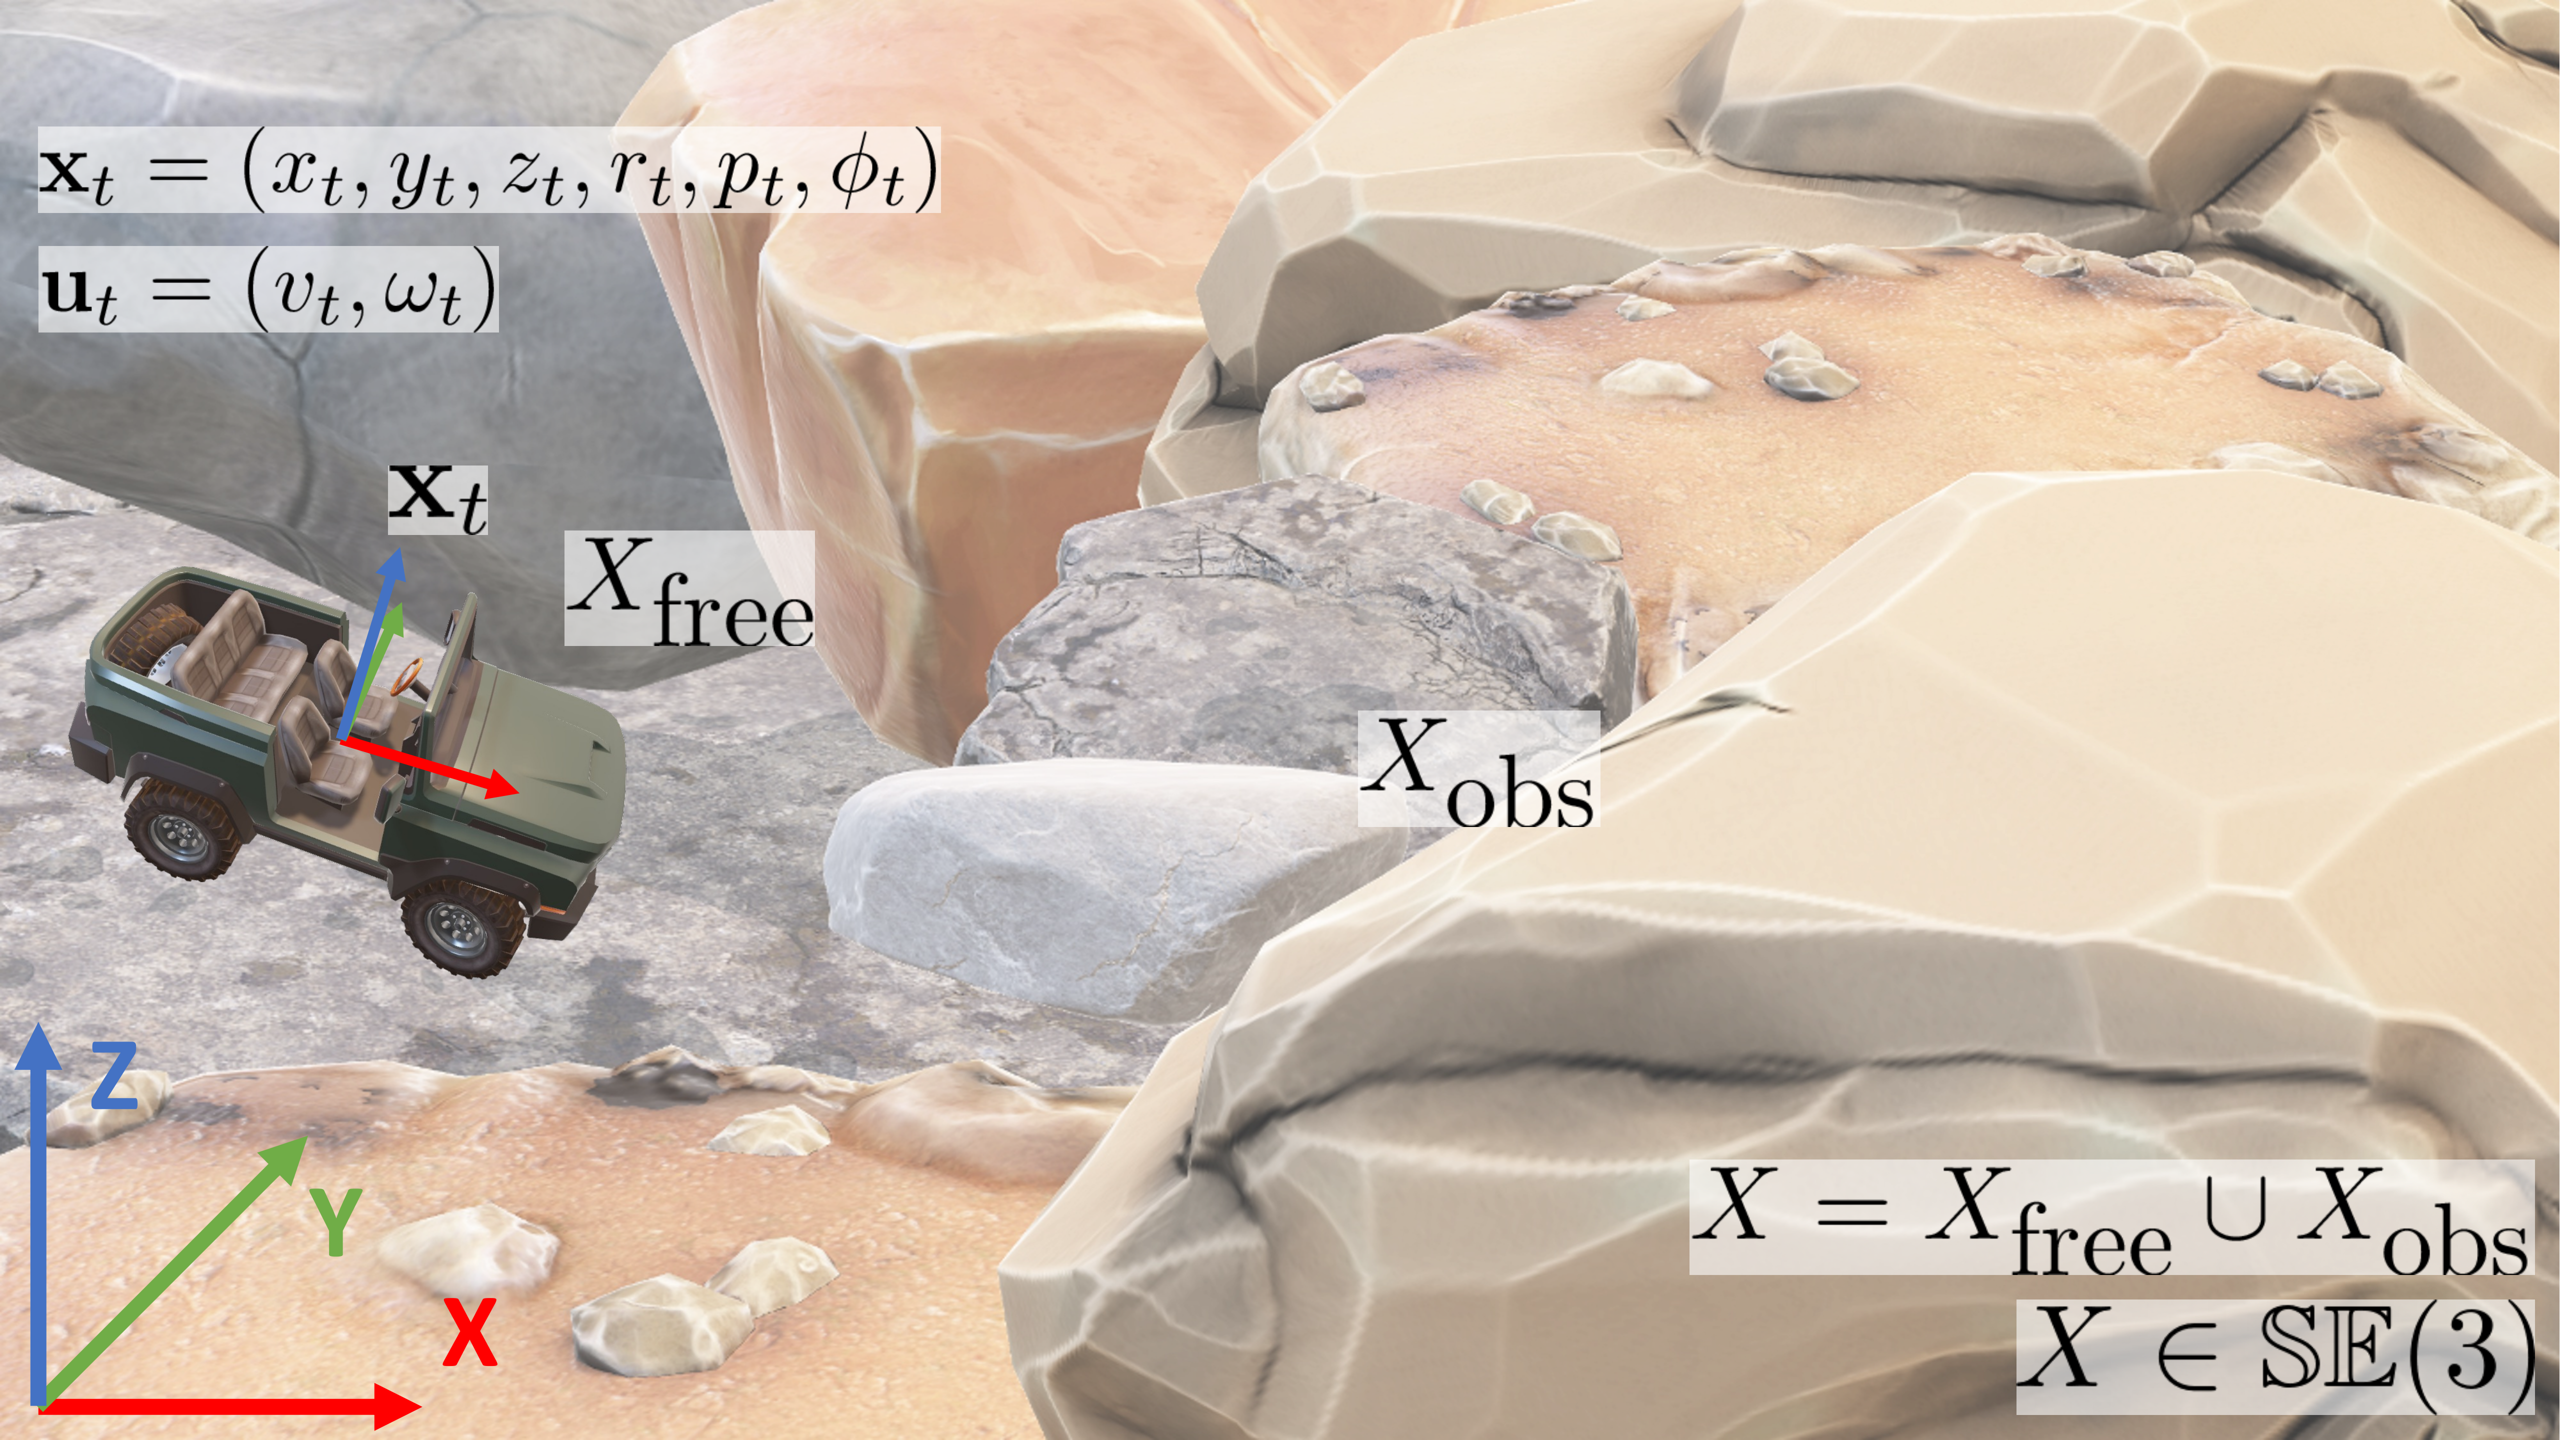

$\mathbf{x}_t = (x_t, y_t, z_t, r_t, p_t, \phi_t)$
$\mathbf{u}_t = (v_t, \omega_t)$
$\mathbf{x}_t$
X_{free}
X_{obs}
Z
Y
X
$X = X_{\text{free}} \cup X_{\text{obs}}$
$X \in \mathbb{SE}(3)$

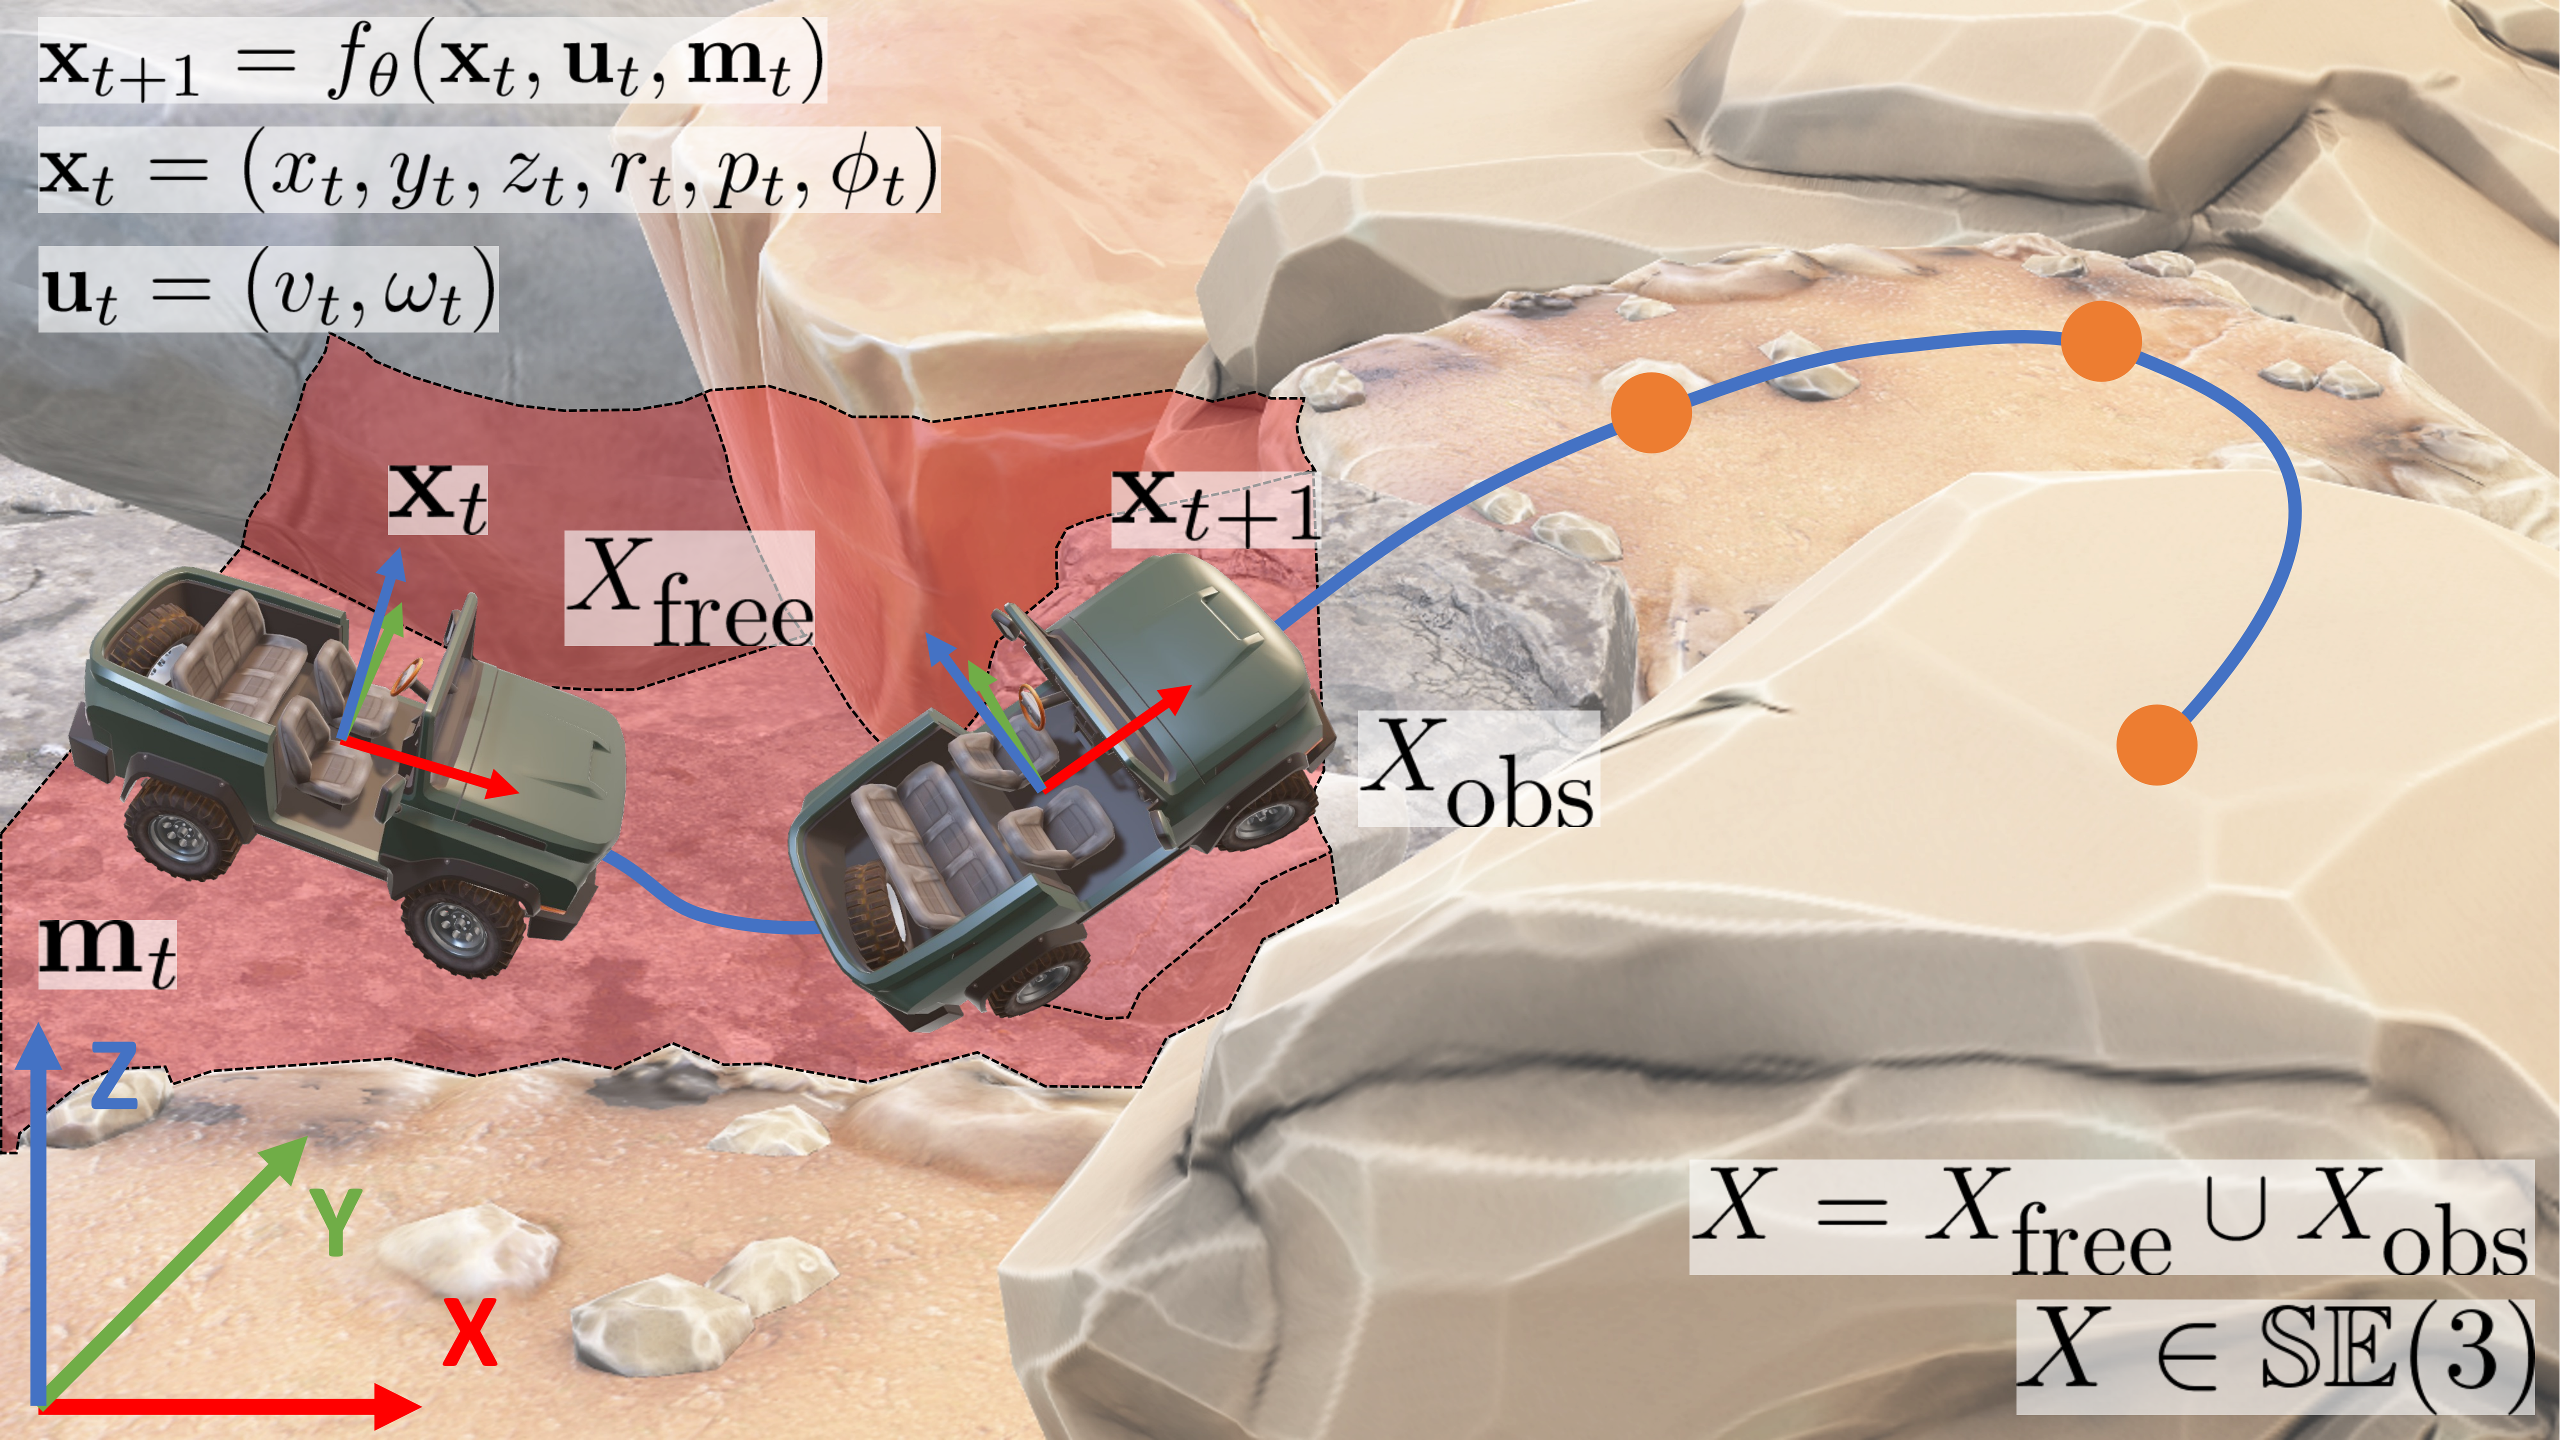
$\mathbf{x}_{t+1} = f_\theta(\mathbf{x}_t, \mathbf{u}_t, \mathbf{m}_t)$
$\mathbf{x}_t = (x_t, y_t, z_t, r_t, p_t, \phi_t)$
$\mathbf{u}_t = (v_t, \omega_t)$
$\mathbf{x}_t$
X_{free}
$\mathbf{x}_{t+1}$
X_{obs}
$\mathbf{m}_t$
Z
Y
X
$X = X_{\text{free}} \cup X_{\text{obs}}$
$X \in \mathbb{SE}(3)$

Planning with Learned 6-DoF Forward Model

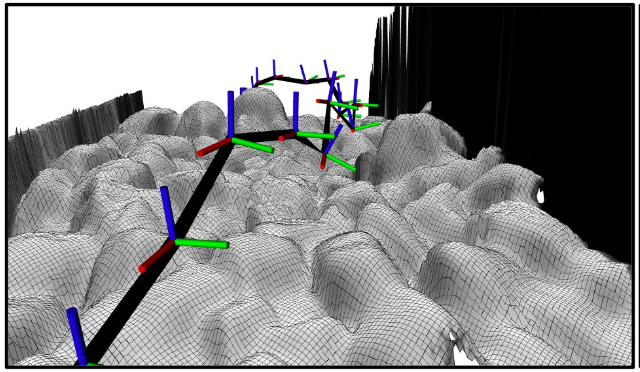

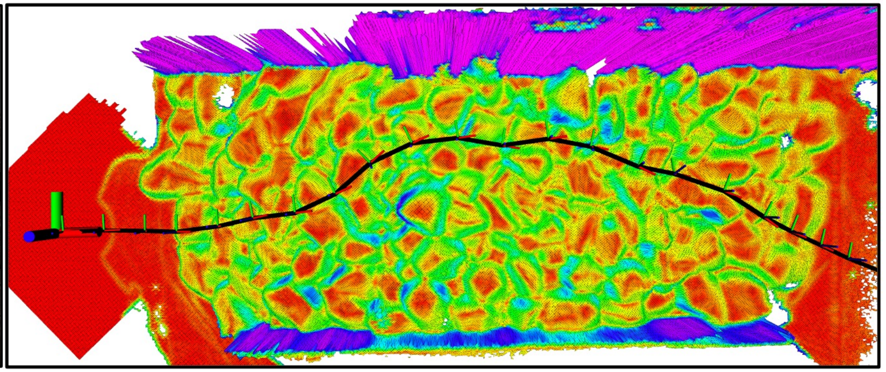

*[D, P, **XX**, under review]*

Multi-Robot Mobility on Vertically Challenging Terrain?

Learning Navigation in Challenging Environments

Highly-Constrained Environments

[***XX** et al., RA-L21a, **XX** et al., ICRA21,* W, ***XX** et al., IROS21*]

Offroad Environments

[***XX** et al., RA-L21b, K, S, A, R, **XX** et al., IROS22, D, P, N, **XX**, ICRA24, D, P, **XX**, under review*]

Social Environments

[***XX** et al., CoRL22, K, N, **XX** et al., RA-L22, N, N, P, D, **XX**, IROS23*]

SCAND: A Large-Scale Dataset of Socially Compliant Navigation Demonstration [K, N, **XX** *et al.*, RA-L22]

- 25 miles (8.7 hours) of real-world data (~0.5TB)
- 138 trajectories, 15 days
- Data collected on two robots: Jackal and Spot
- Indoor and outdoor environments @ UT Austin
- Four different human demonstrators
- Coarse labels of social interactions

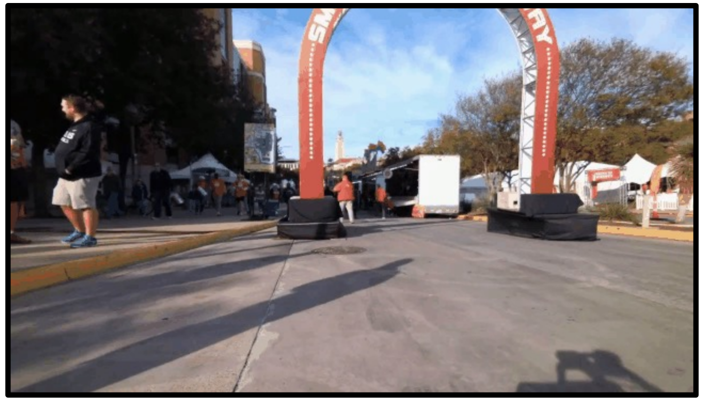

MuSoHu: Multi-Modal Social Human Navigation Dataset *[N, N, P, D, **XX**, IROS23]*

- 50km, 10 hours, 150 trials, 7 humans, **and counting**!

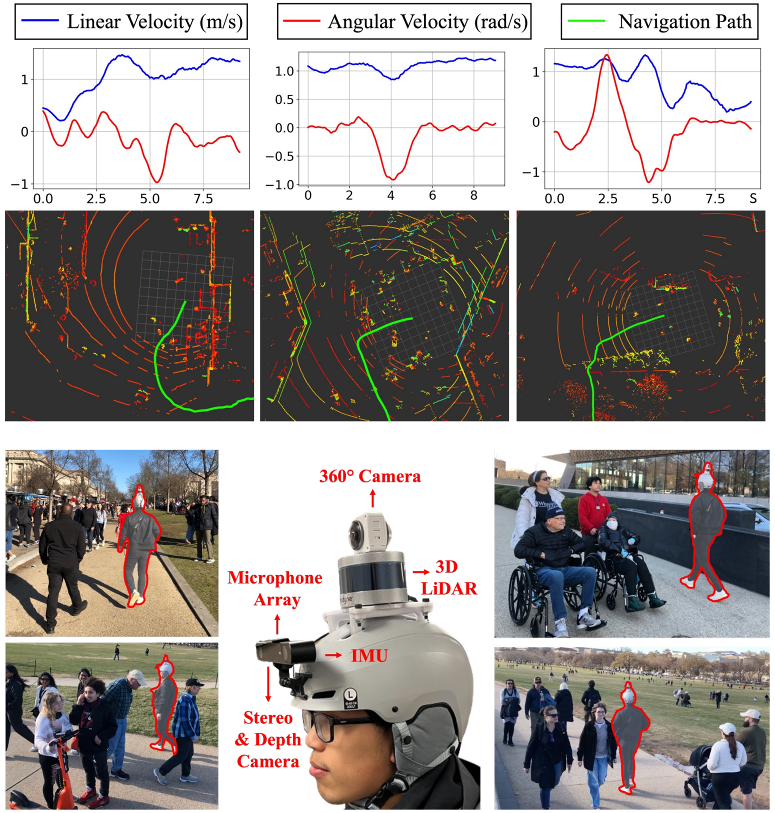

Performer-MPC: Socially Compliant Navigation Behavior by Real-Time Transformers [*XX et al., CoRL22*]

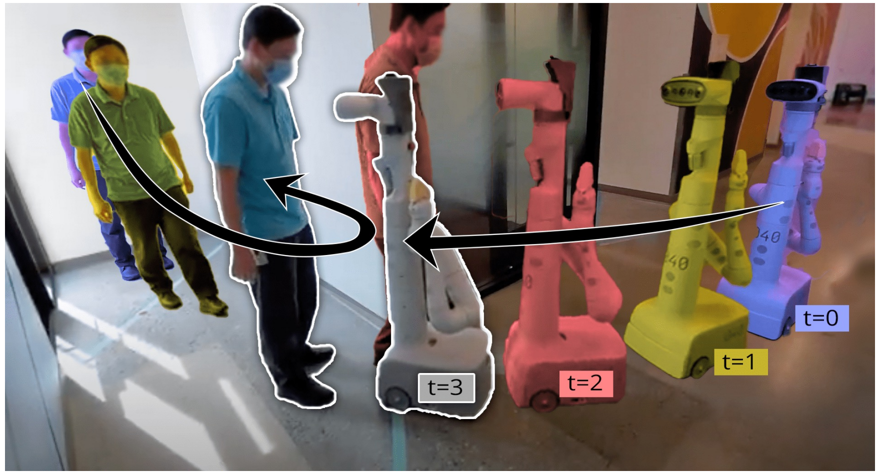

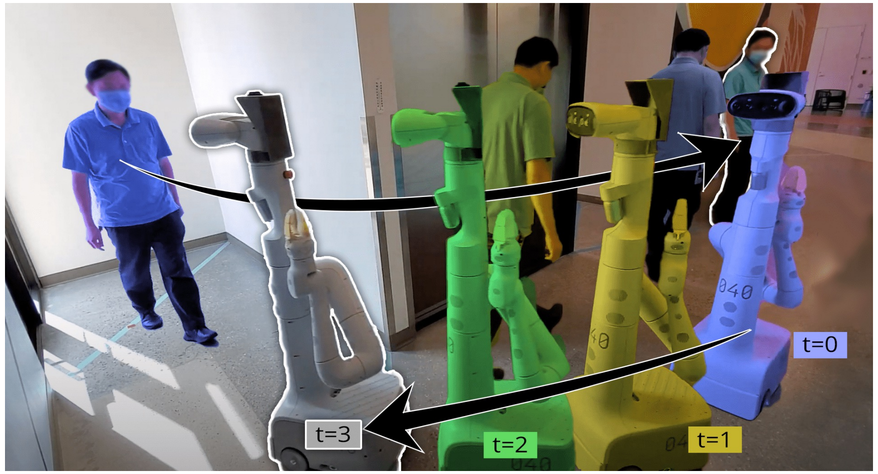

Planning the most efficient, shortest length, minimal time plan?

Social compliance improves motion planning performance!

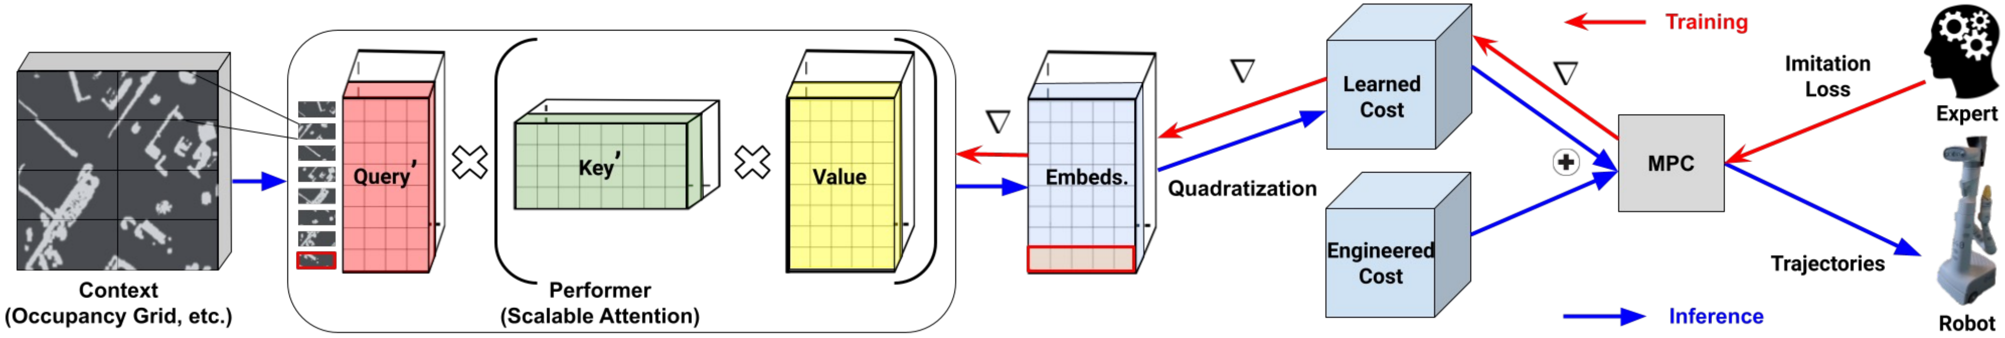

Performer-MPC: Socially Compliant Navigation Behavior by Real-Time Transformers [***XX*** *et al., CoRL22*]

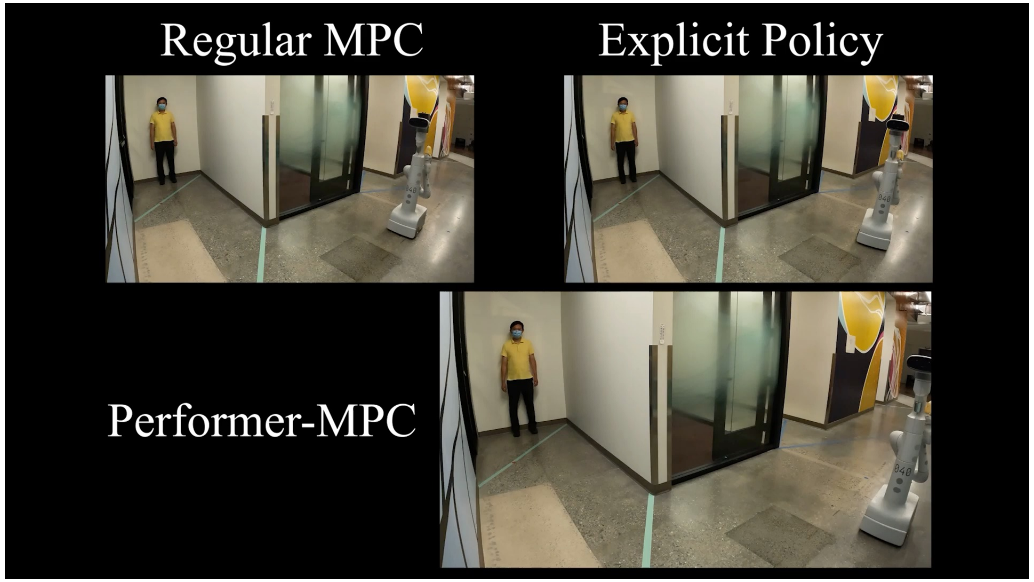

Blind Corner
Learning to anticipate Pedestrians

Pedestrian Obstruction
Learning to respect comfort distance

Team Coordination on Graphs *[O, L, H, W,* ***XX****, S, IROS23]*

Joint State Graph (JSG)

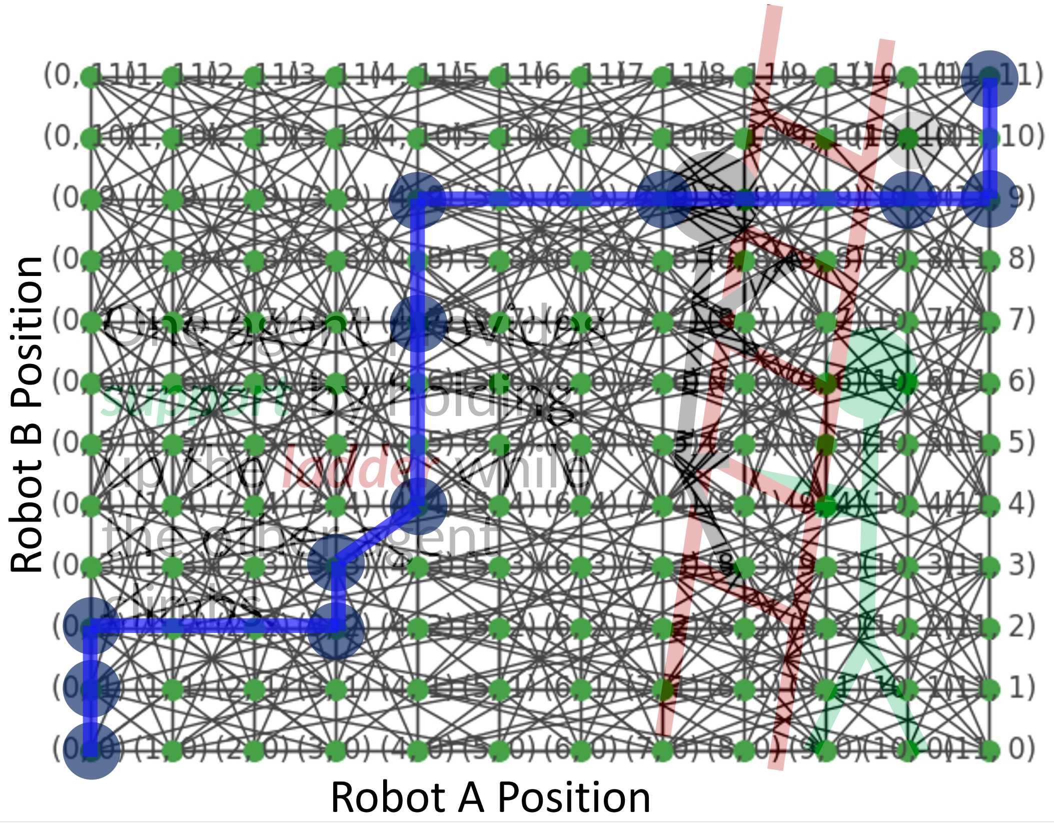

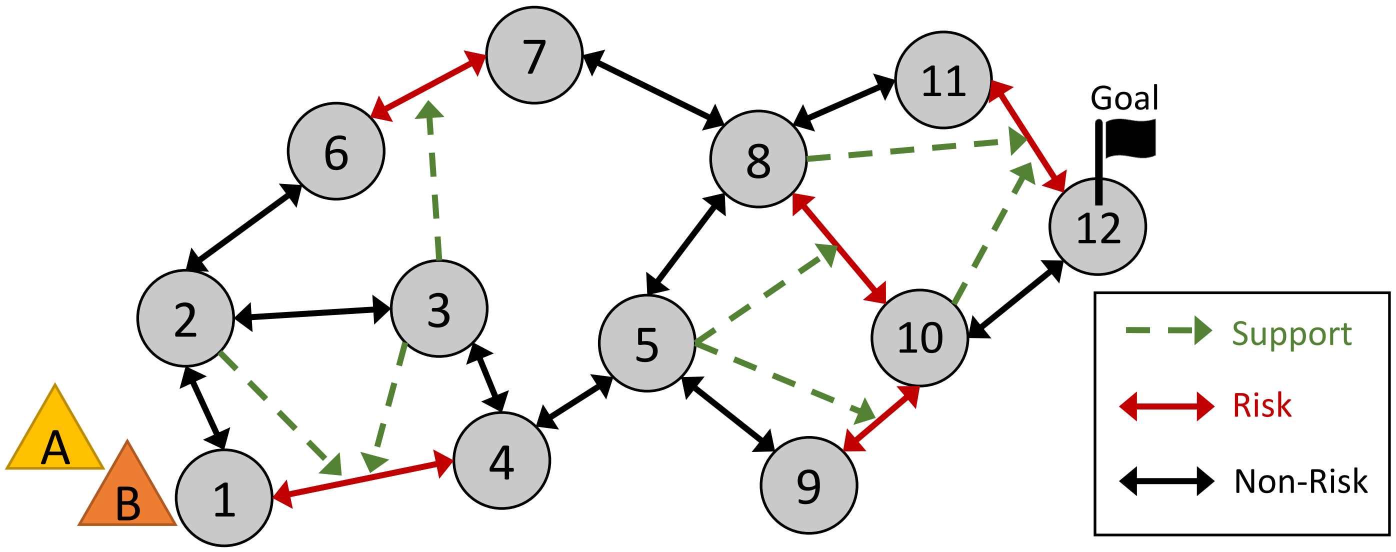

Team Coordination on Graphs *[O, L, H, W, XX, S, IROS23]*

Critical Joint State Graph (CJSG)

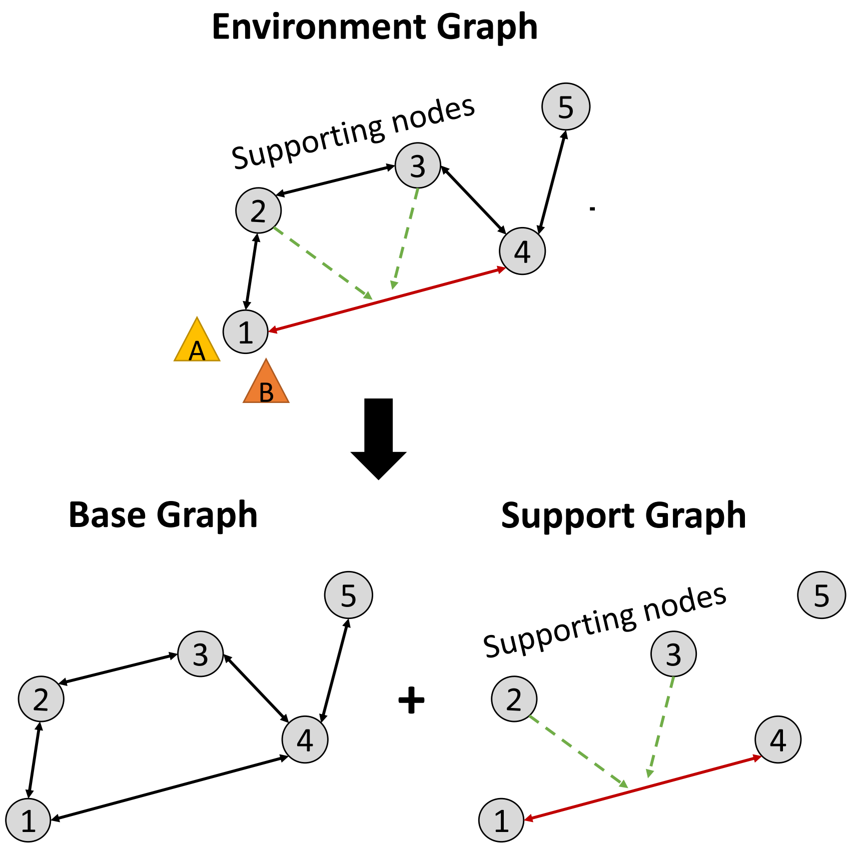

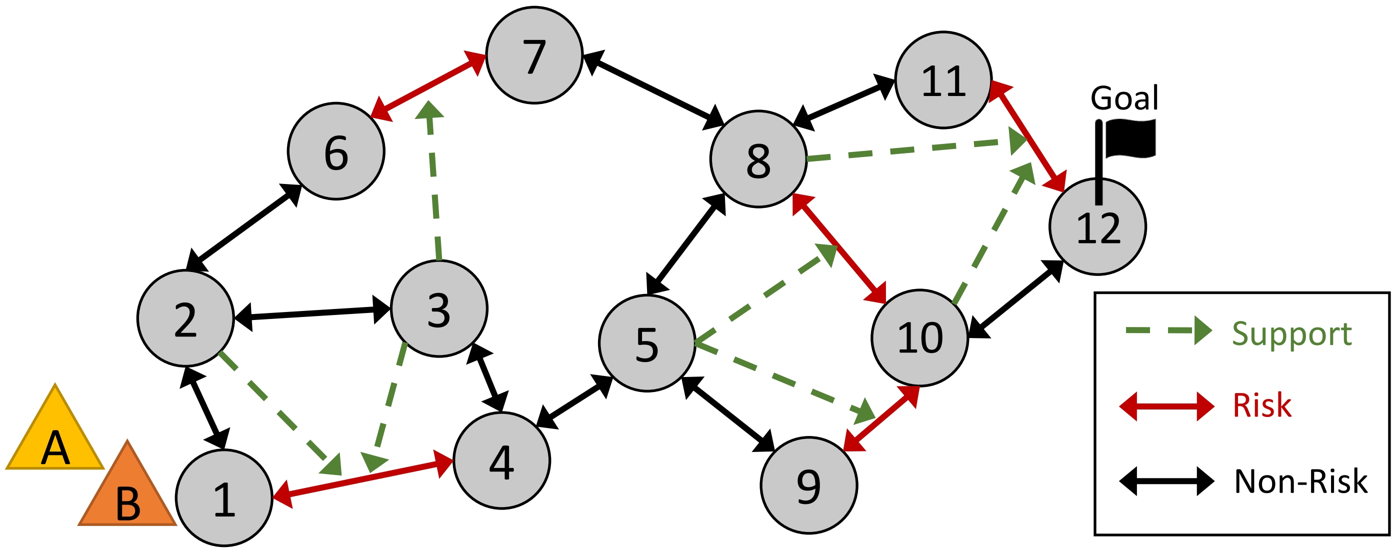

Adaptive Planner Parameter Learning (APPL): Leveraging Non-Expert Interactions in Social Environments

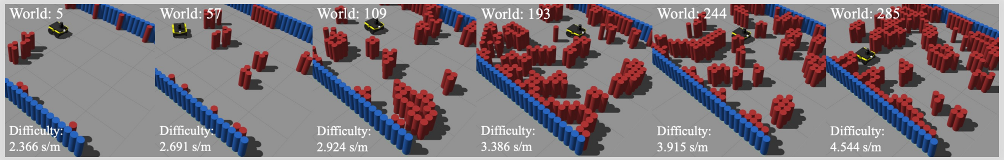

APPLR
[X, D, N, ***XX*** *et al., ICRA21*]

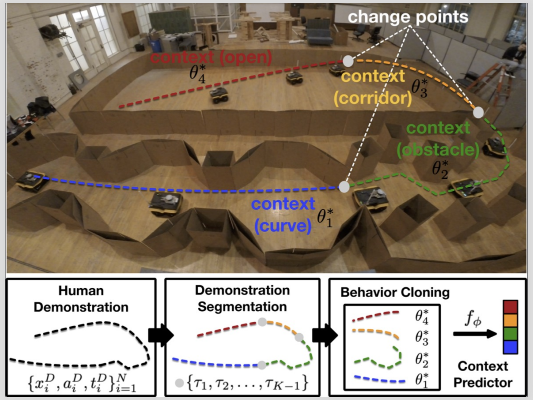

APPLD
[***XX*** *et al., RA-L20*]

Demonstration

Interventions

Evaluative Feedback

Reinforcement Learning

Classical Autonomy System

APPLE
[*W,* ***XX*** *et al., RA-L21*]

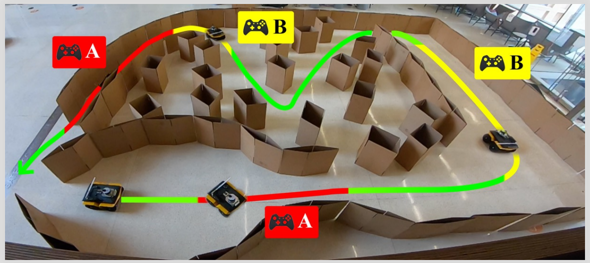

APPLI
[*W,* ***XX*** *et al., ICRA21*]

Learning Navigation in Challenging Environments

Highly-Constrained Environments

[***XX*** *et al., RA-L21a,* ***XX*** *et al., ICRA21,* W, ***XX*** *et al., IROS21*]

Offroad Environments

[***XX*** *et al., RA-L21b, K, S, A, R,* ***XX*** *et al., IROS22, D, P, N,* ***XX****, ICRA24, D, P,* ***XX****, under review*]

Social Environments

[***XX*** *et al., CoRL22, K, N,* ***XX*** *et al., RA-L22, N, N, P, D,* ***XX****, IROS23*]

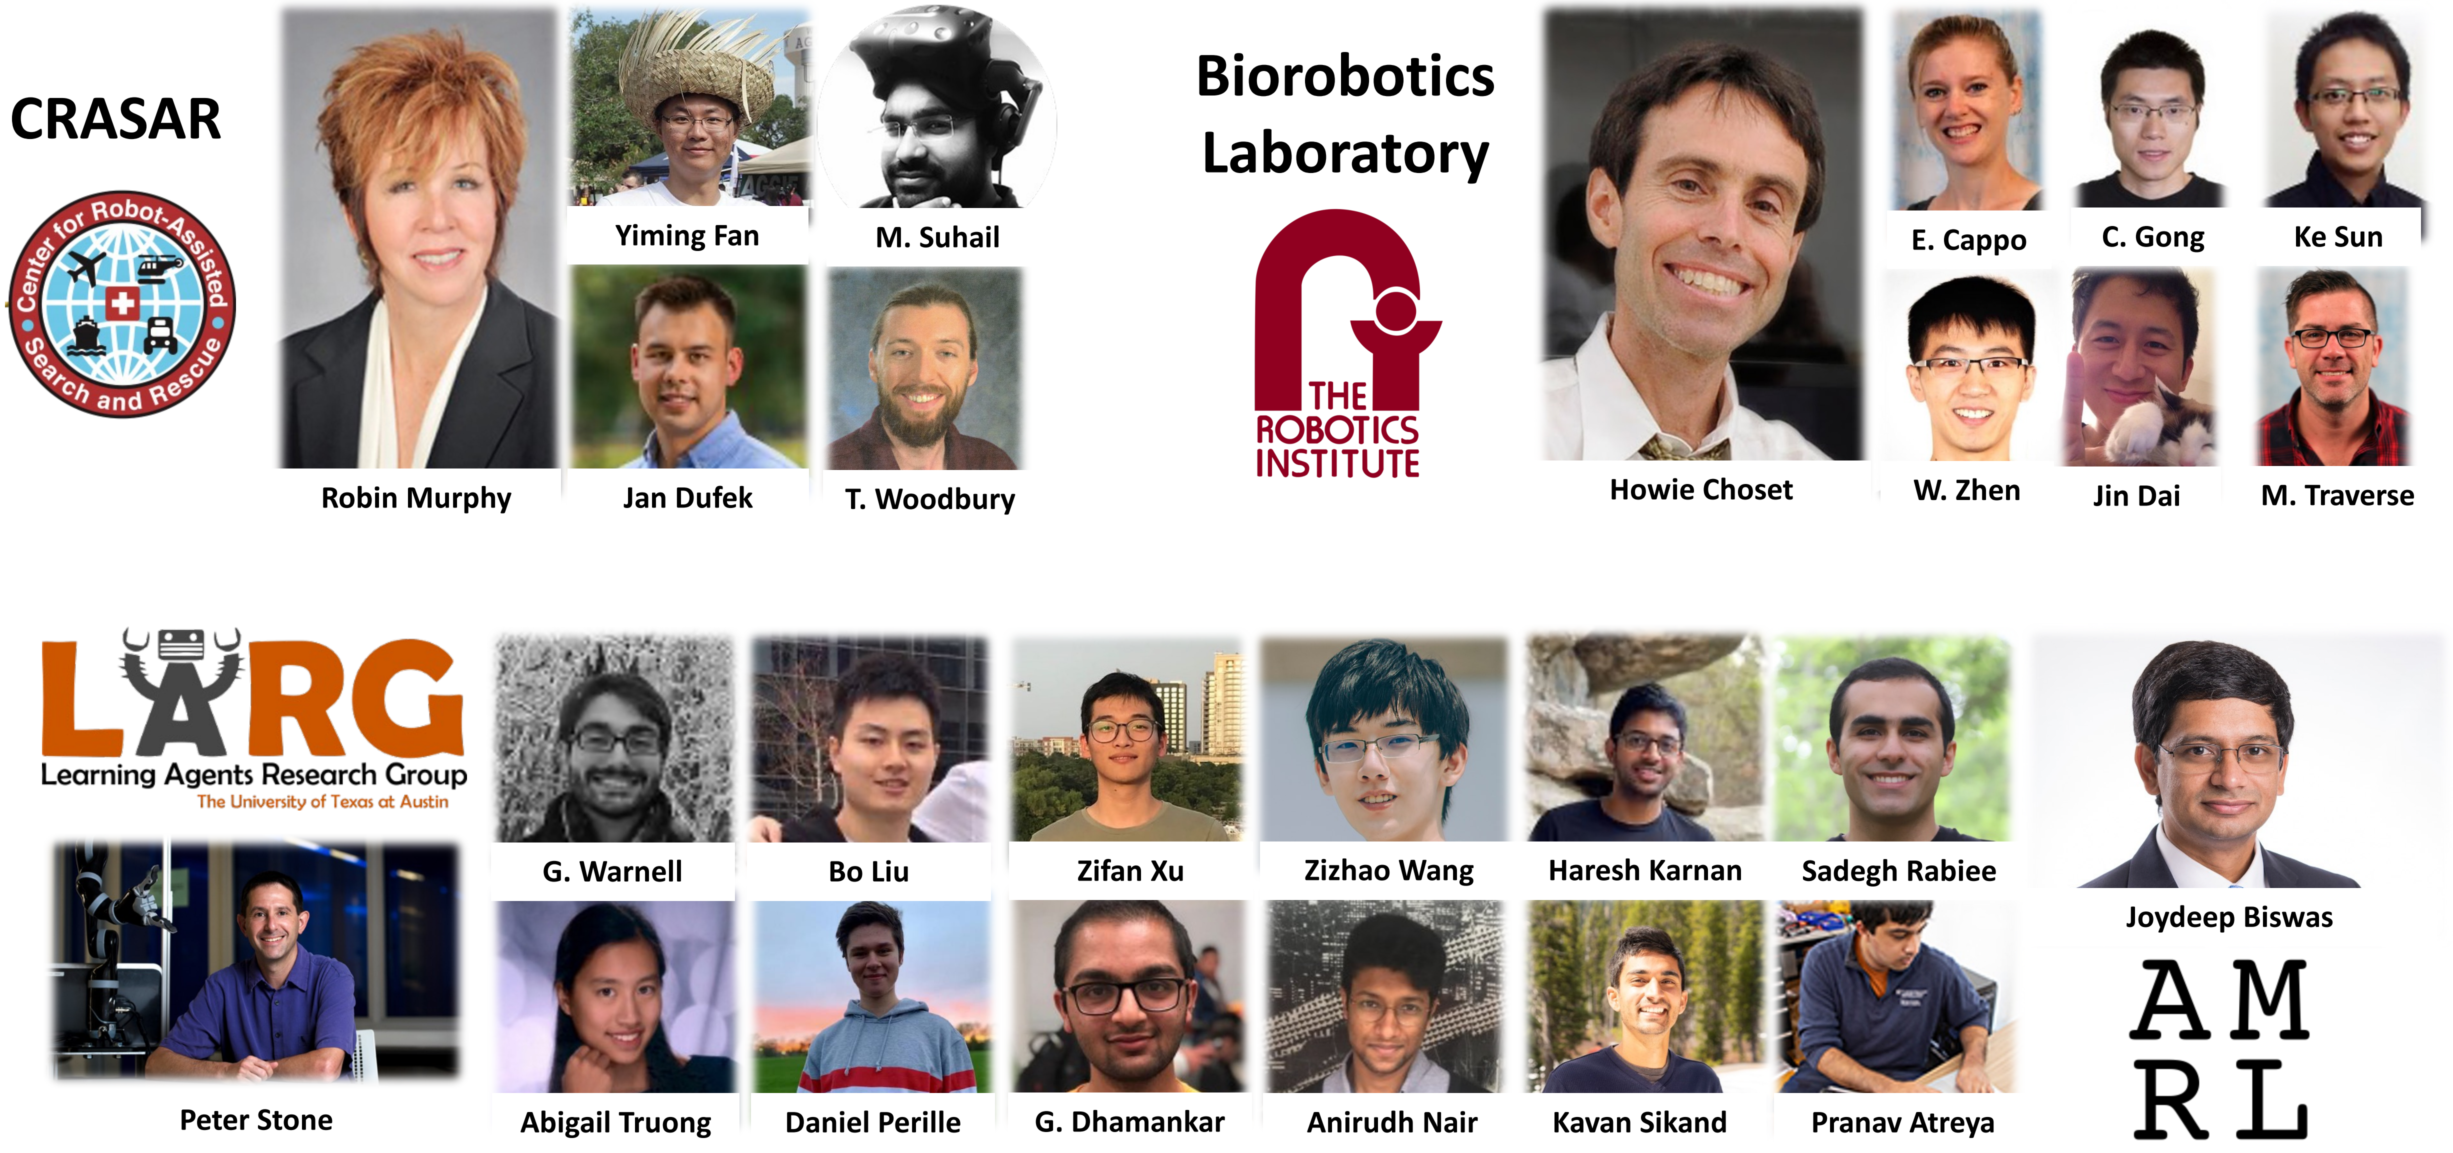
CRASAR
Center for Robot-Assisted Search and Rescue
Robin Murphy
Yiming Fan
M. Suhail
Jan Dufek
T. Woodbury
Biorobotics Laboratory
THE ROBOTICS INSTITUTE
Howie Choset
E. Cappo
C. Gong
Ke Sun
W. Zhen
Jin Dai
M. Traverse
LARG
Learning Agents Research Group
The University of Texas at Austin
Peter Stone
G. Warnell
Bo Liu
Zifan Xu
Zizhao Wang
Haresh Karnan
Sadegh Rabiee
Abigail Truong
Daniel Perille
G. Dhamankar
Anirudh Nair
Kavan Sikand
Pranav Atreya
Joydeep Biswas
AMRL

Everyday Robots

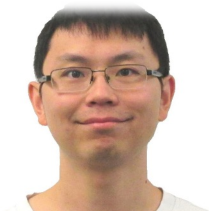
Tingnan Zhang


K. Choromanski


Edward Lee


Anthony Francis


Jake Varley


Stepehn Tu


Sumeet Singh

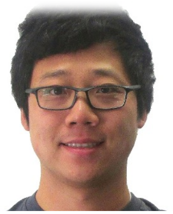
Peng Xu


Fei Xia


Mikael Persson


D. Kalashnikov


Leila Takayama


Roy Frostig


Jie Tan


Carolina Parada


Vikas Sindhwani

Google Research


Daigo Shishika


Sara Oughourli

GEORGE MASON UNIVERSITY | College of Engineering and Computing


Zechen Hu

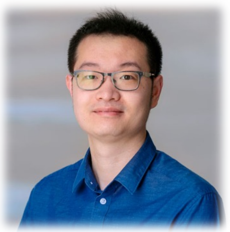
Xuan Wang


Aniket Datar


Chenhui Pan


Aaron Nguyen


Mohammad Nazeri


Amir Payandeh


Manshi Limbu

RobotiX

Learning Navigation in Challenging Environments

Highly-Constrained Environments

[***XX*** *et al., RA-L21a,* ***XX*** *et al., ICRA21,* W, ***XX*** *et al., IROS21*]

Offroad Environments

[***XX*** *et al., RA-L21b, K, S, A, R,* ***XX*** *et al., IROS22, D, P, N,* ***XX****, ICRA24, D, P,* ***XX****, under review*]

Social Environments

[***XX*** *et al., CoRL22, K, N,* ***XX*** *et al., RA-L22, N, N, P, D,* ***XX****, IROS23*]