Learning-Enhanced Autonomous Navigation for GPS-Denied Vehicles

ART-017

Brendan Englot, Associate Professor, Stevens Institute of Technology

U.S. Army DEVCOM Armaments Center

ANNUAL RESEARCH REVIEW 2022

Project Goals

- Develop a **multi-robot localization and mapping framework** that allows a team of unmanned ground vehicles (UGVs) to build more accurate, comprehensive LIDAR-based maps in GPS-denied environments by **regularly exchanging information.**
- Implement **autonomous multi-robot exploration**, which allows the UGVs to leverage their cooperative localization and mapping capability to efficiently build a complete map of an unknown, GPS-denied environment within a set of specified boundaries.
- **Integrate machine learning tools into the above capabilities** to facilitate both efficient planning and decision-making, using graph networks, and the construction of predictive perceptual data products, such as learning-enhanced grid maps and terrain traversability maps, that will enable high-performance motion planning.

Toward Comms-Efficient Multi-robot Localization and Mapping

Mapping the Stevens campus - Hoboken, NJ

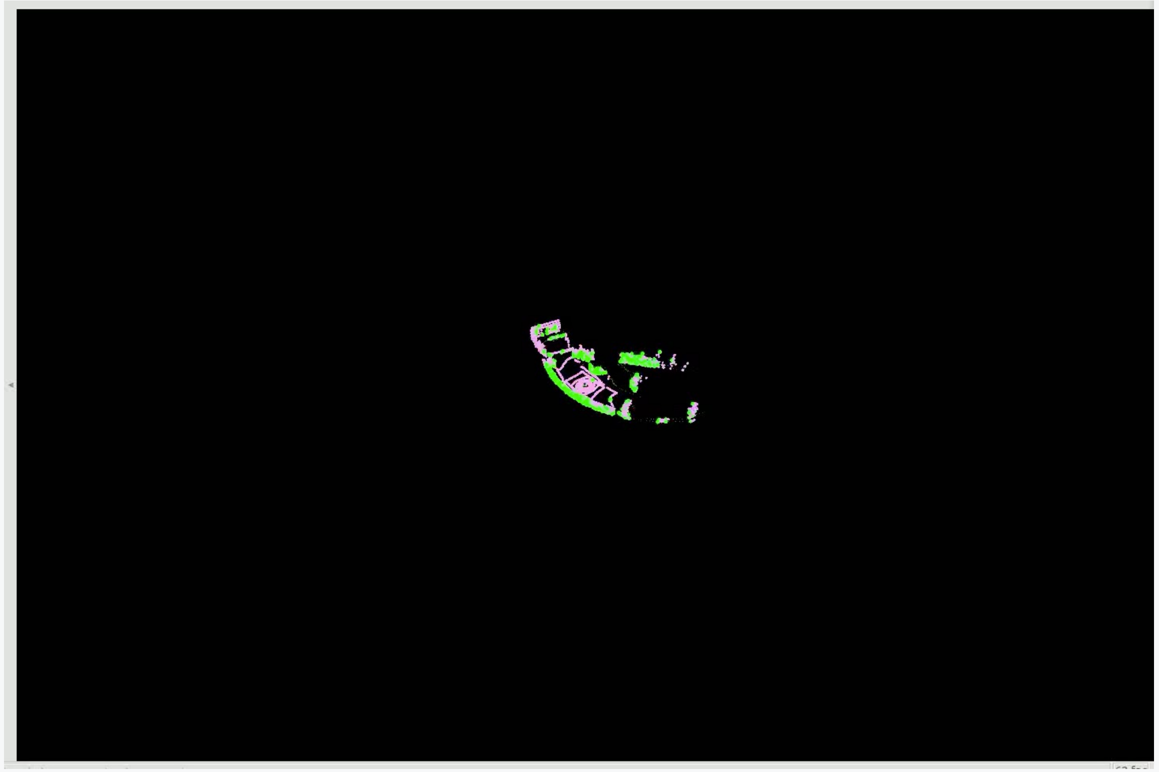

Mapping a mine underground – Louisville, KY

- Our Lightweight, Ground-Optimized Lidar Odometry and Mapping (LeGO-LOAM) algorithm supports high-performance mobile robot localization and mapping in GPS-denied environments over variable terrain

Our Proposed Multi-robot SLAM Architecture

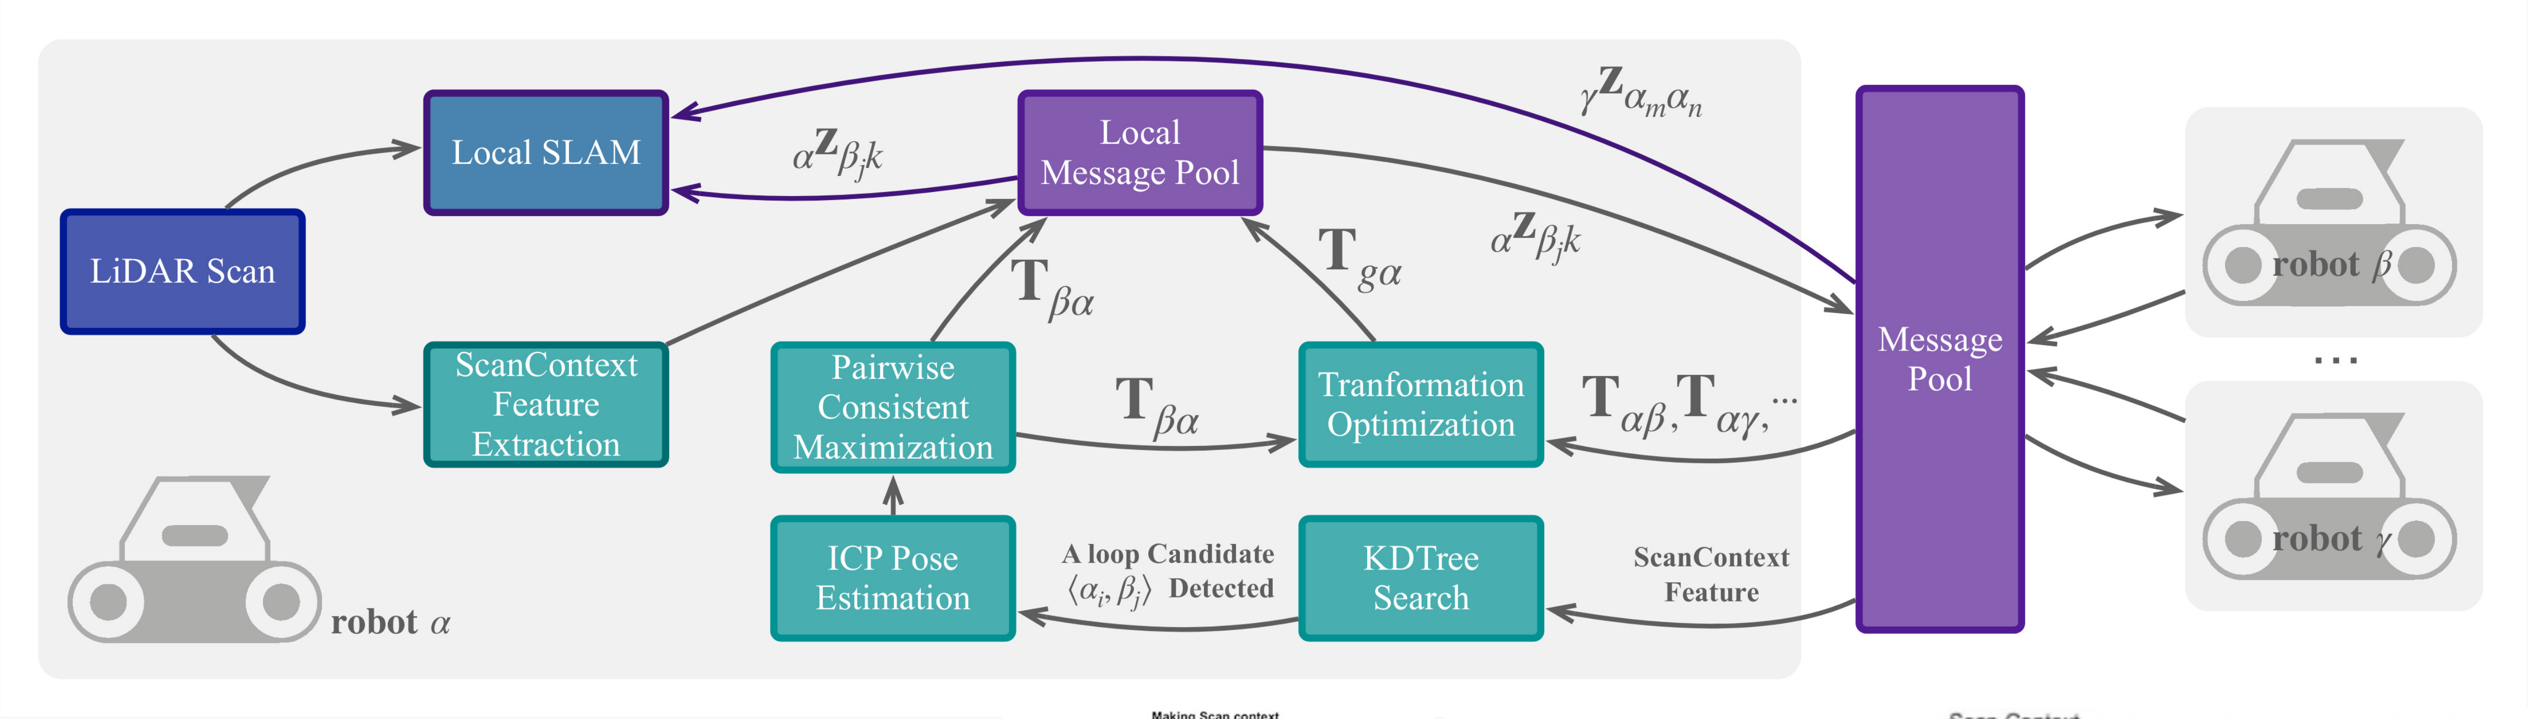

- **Data-Efficient Lidar Scan Descriptors:** Scan Context is used to permit efficient message-passing to identify potential inter-robot loop closures, before exchanging the full-density lidar point clouds

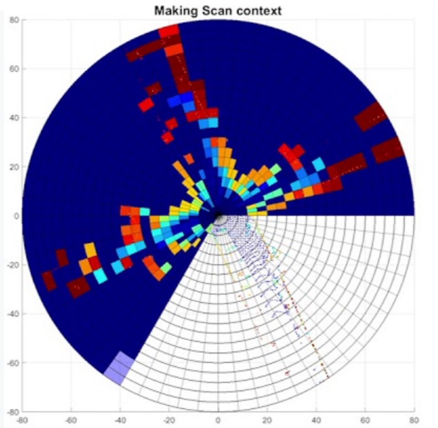

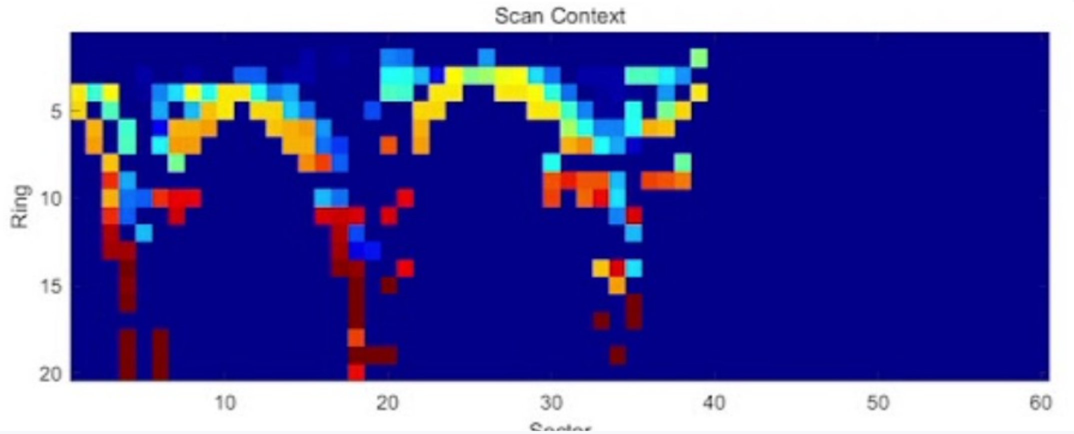

Multi-robot SLAM Experiments (3-robot test)

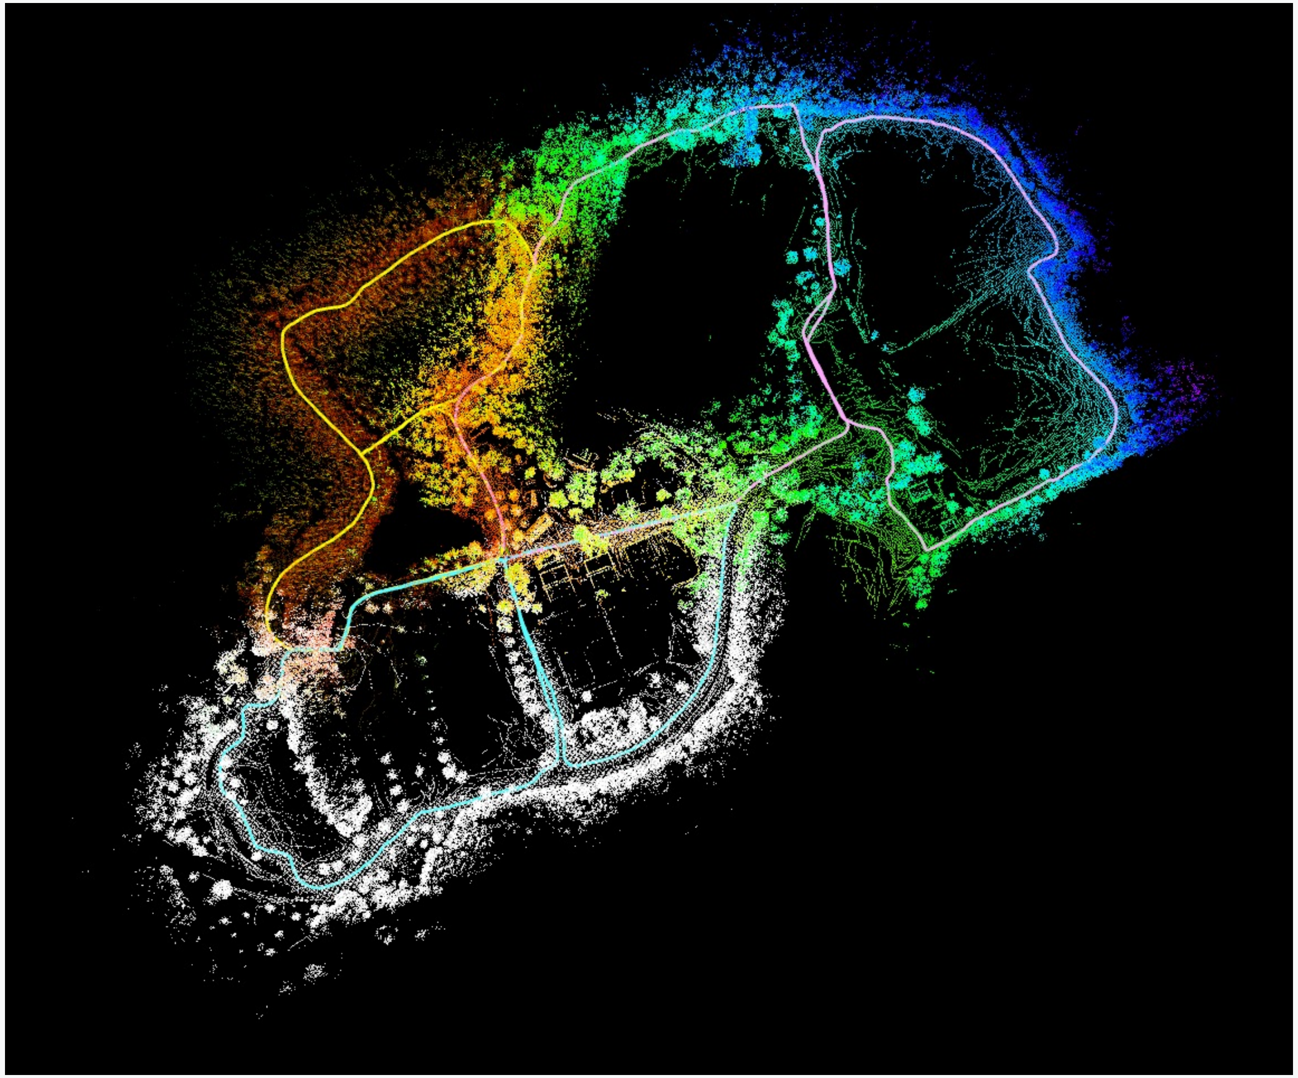

“Park” Dataset

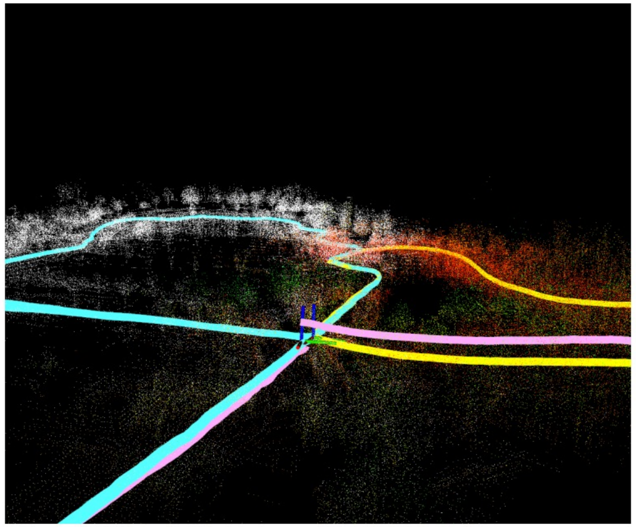

Without inter-robot constraints

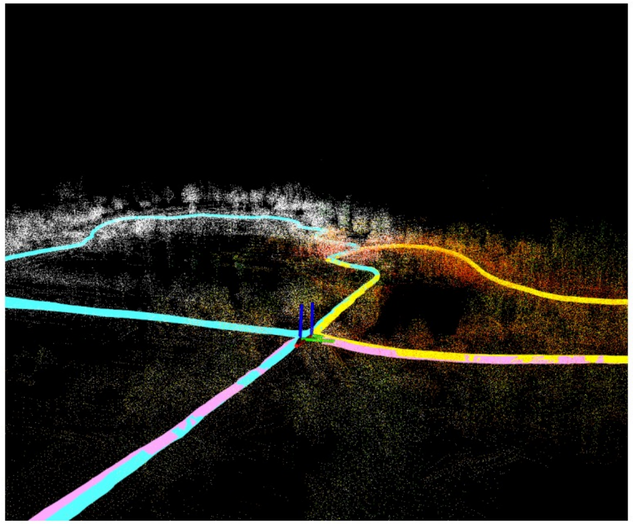

With inter-robot constraints

Multi-robot SLAM Experiments (2-robot test)

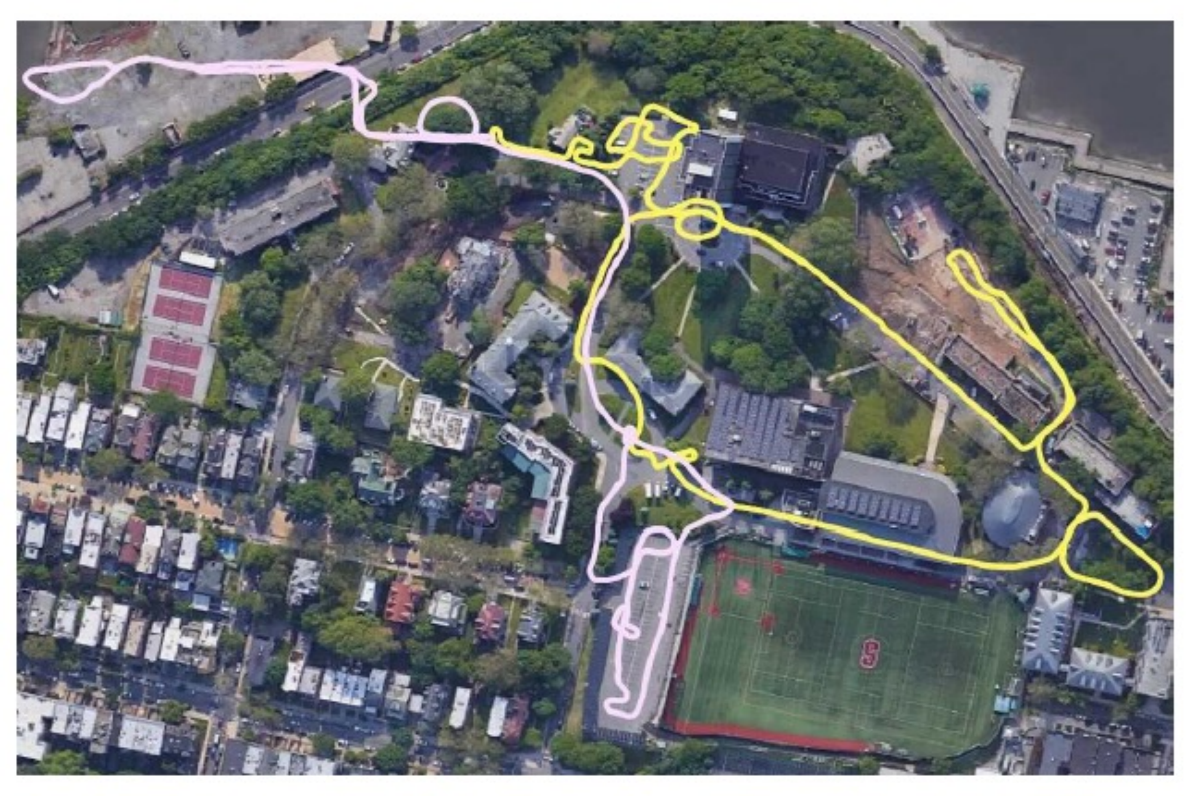

DOOR-SLAM, a competing framework (distributed Gauss-Seidel with pairwise consistent measurement set maximization)

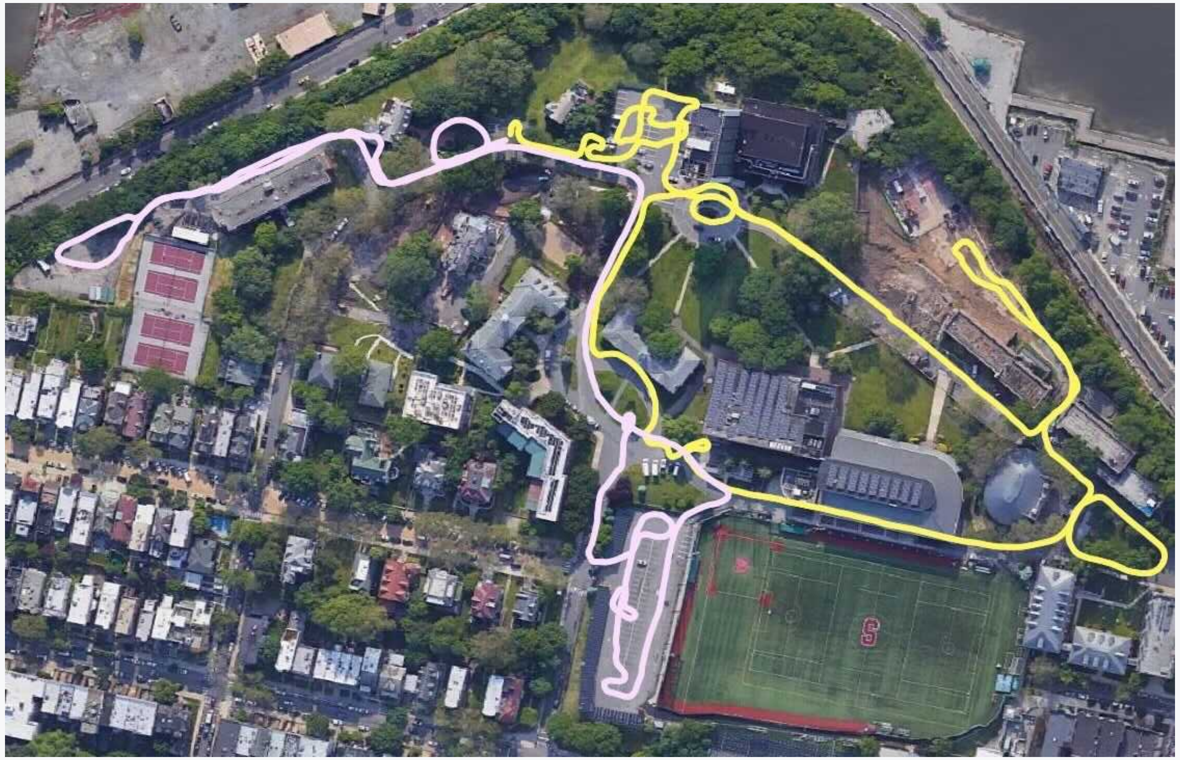

Our proposed framework

"Stevens" Dataset

Communications Requirements of the Proposed Framework

Message Info	Mean (kB)	Min (kB)	Max (kB)	No. Total Msgs.	
				Stevens	Park
SC Feature & Local Pose	4.08	4.04	4.12	3936	4222
Feature Cloud (Edge)	9.65	5.08	15.90	37	837
Feature Cloud (Planar)	71.31	54.74	85.98	37	837
Feature Cloud (Other)	70.91	50.58	83.99	37	837
Coordinate Transformation	0.70	0.70	0.70	3	711
Inter-Robot Loop Closure	0.12	0.12	0.12	3	72

- **Scan Context (SC)** serves as an important screening mechanism preventing the need to send larger messages more frequently

Exploring and Mapping Unknown Environments

Multi-Objective Path Planning to Frontiers

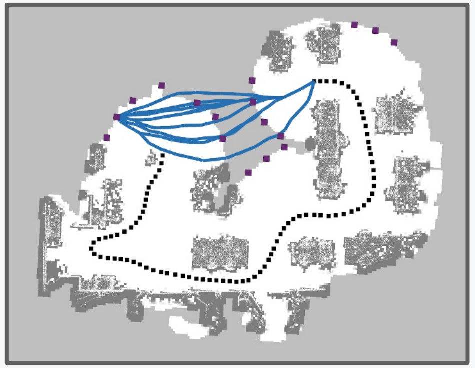

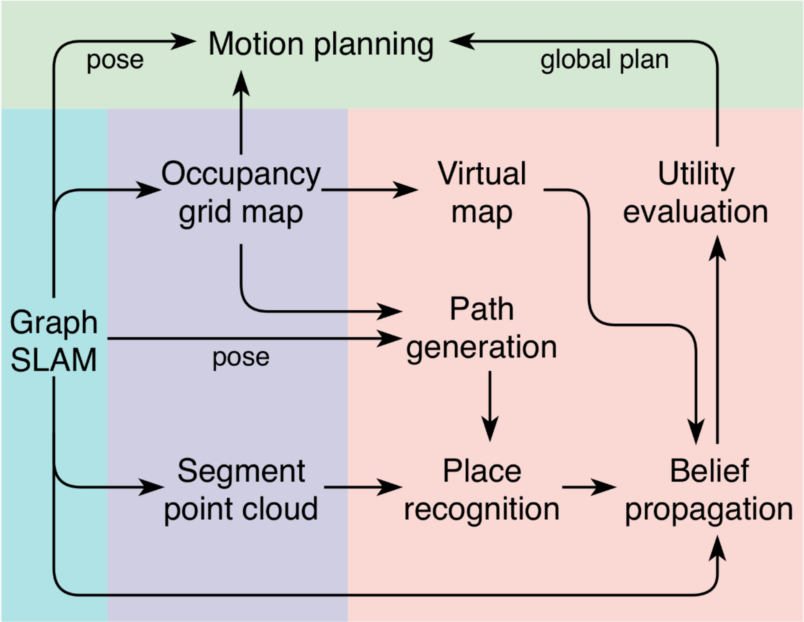

Cost Maps used to Generate Paths:

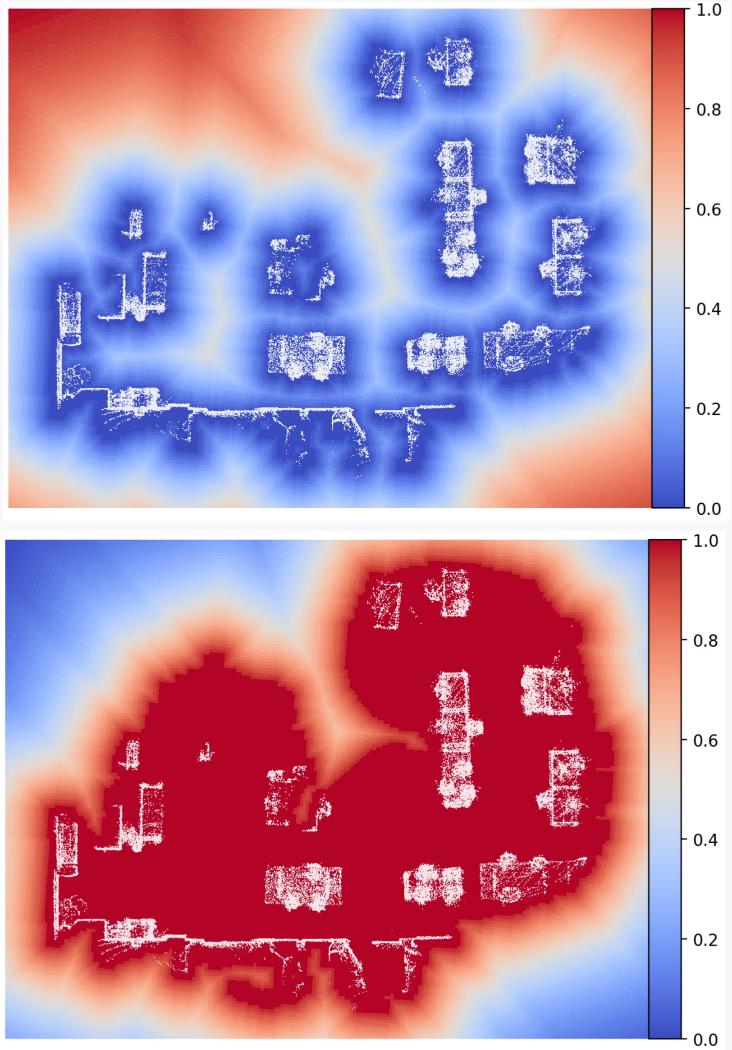

Localization Costmap

Exploration Costmap

- **Expectation-Maximization (EM) Inspired Autonomous Exploration of Unknown Environments:** A novel planning and decision-making architecture for exploring unknown environments with UGVs, which emphasizes the construction of the most accurate map possible of the surrounding structures in the environment

Exploring and Mapping Unknown Environments

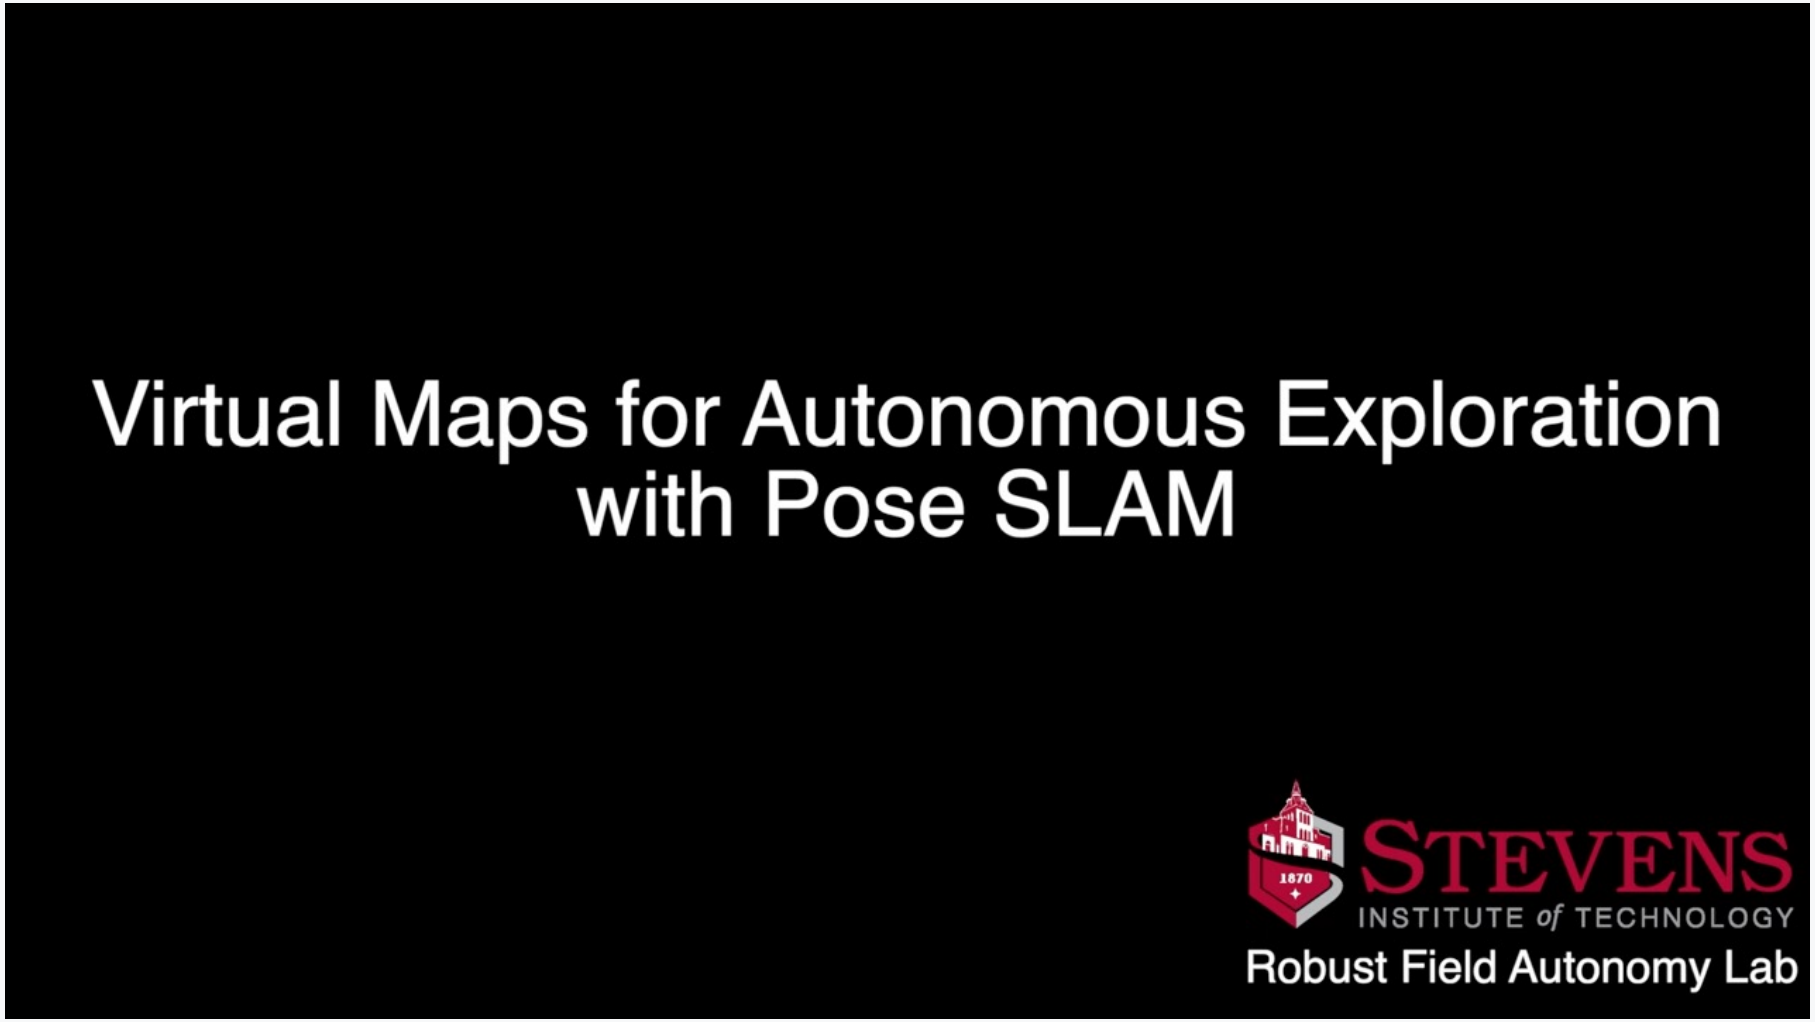

- Here our UGV explores a series of indoor environments without any prior information, and maps the environments as accurately and efficiently as possible by making predictions about its future lidar observations

Graph-based Exploration

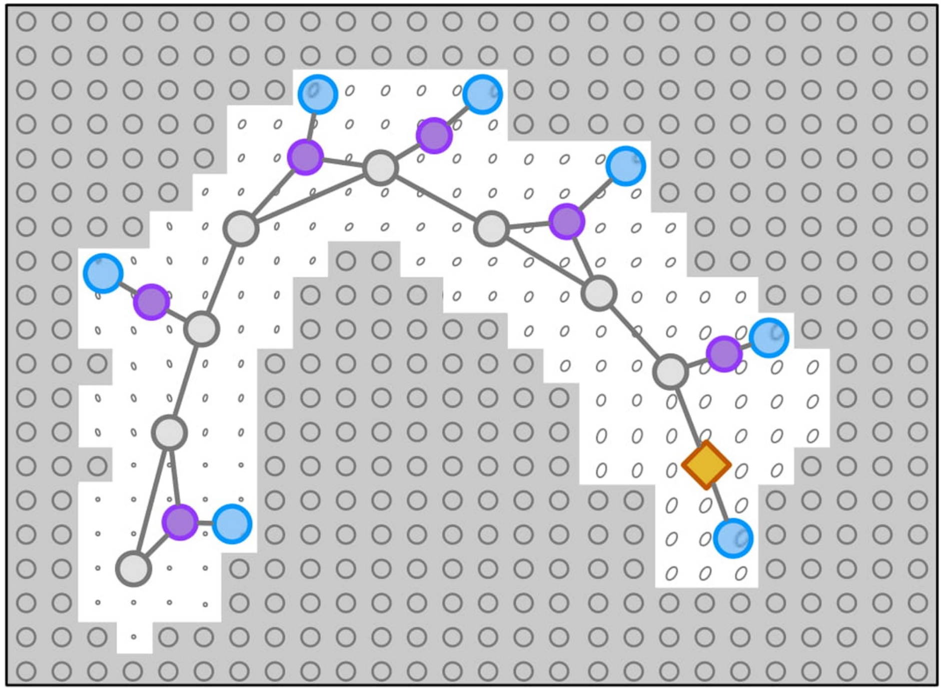

The grayscale color represents the probability of occupancy of a given map cell. The ellipse of each cell is the estimation error covariance of each *virtual landmark*.

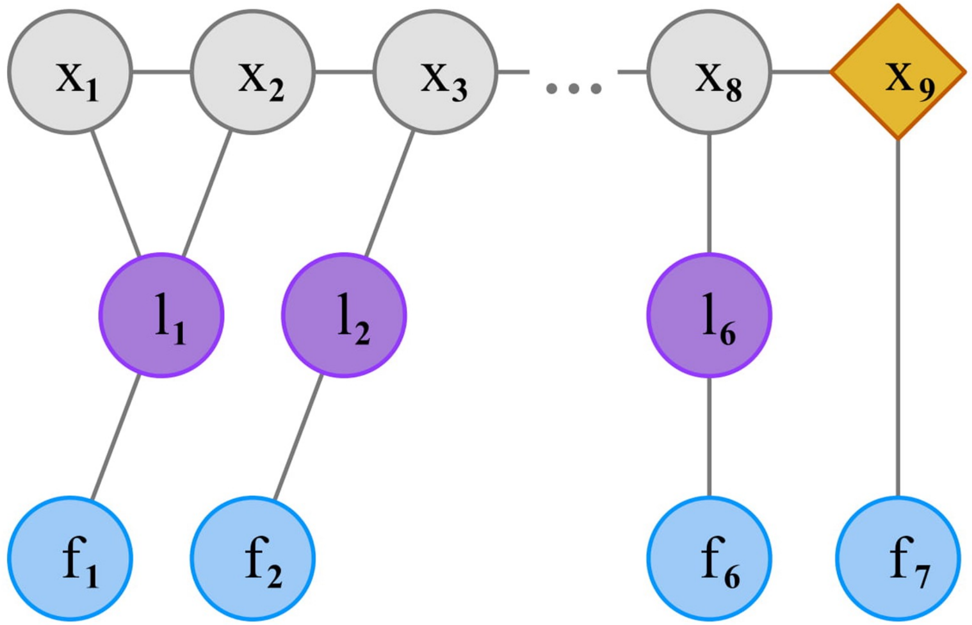

An input *exploration graph* is extracted from the current exploration state. Each edge in this graph is weighted with the Euclidean distance between the two vertices connected

Exploration with Deep Reinforcement Learning (DRL) on Graphs

Advantages of DRL:

- The policy can be trained without human intervention.
- The policy is non-myopic and robust across a variety of test environments.
- The policy can be scalable to testing in larger, higher-dimensional spaces.

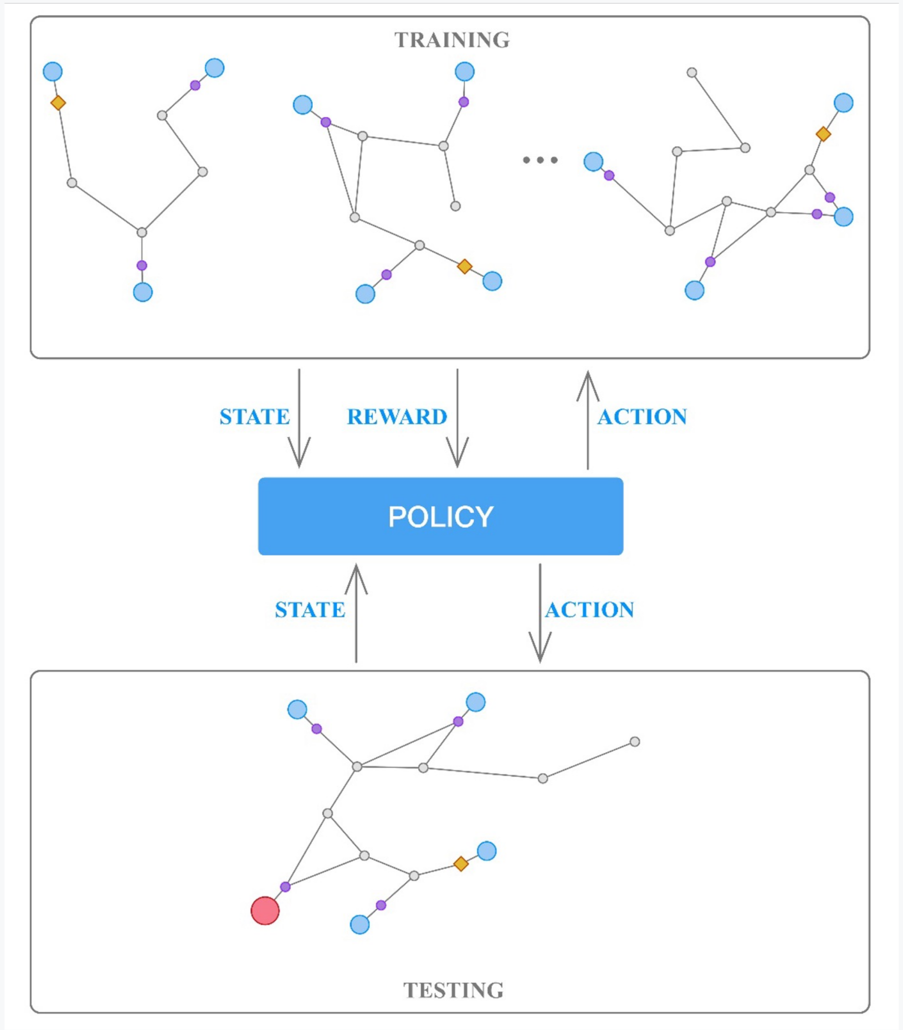

State Space and Action Space

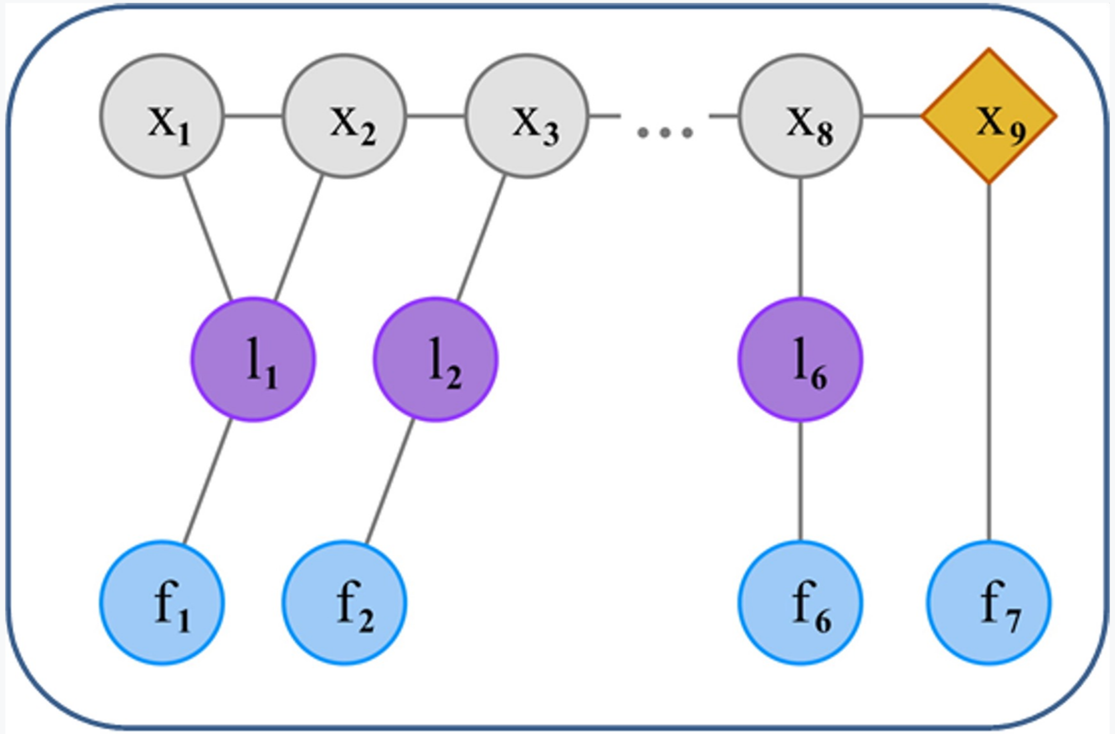

State

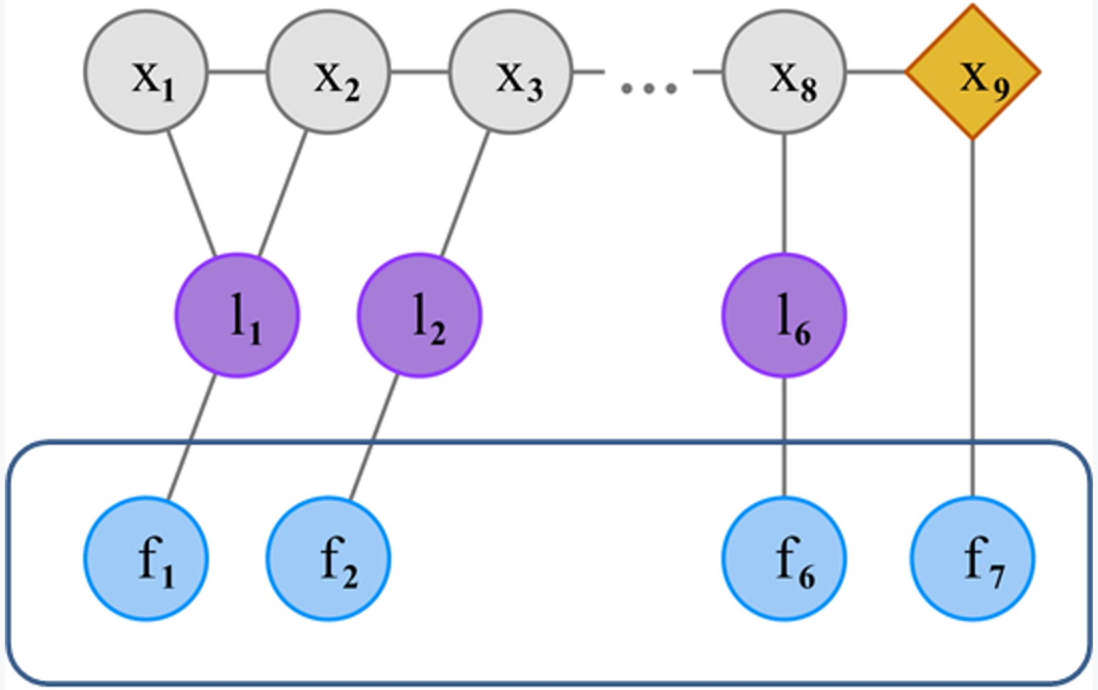

Actions

Computational Study of DRL-GNN Exploration

Exploration Test on 40m x 40m Maps

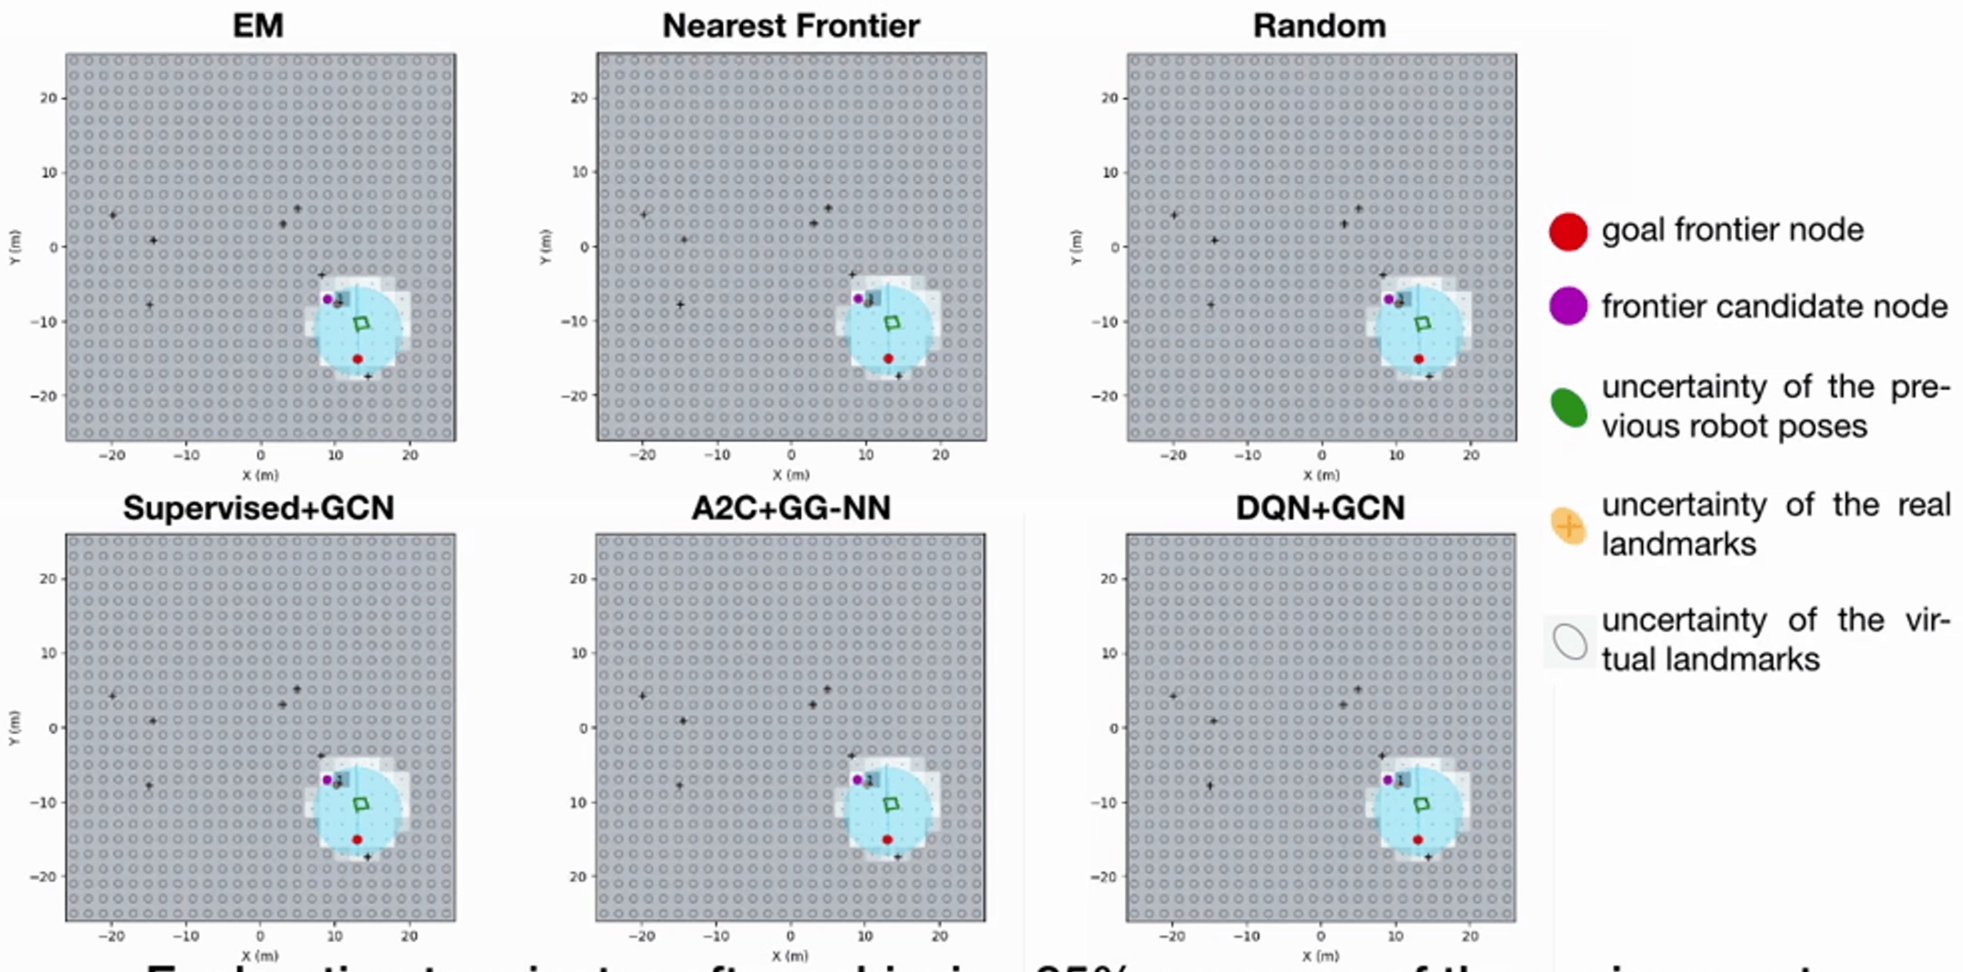

Exploration terminates after achieving 85% coverage of the environment.

Computational Study of DRL-GNN Exploration

Exploration Trained on
40m x 40m Map
Test on
60m x 60m
80m x 80m
100m x 100m Maps

- goal frontier node
- frontier candidate node
- uncertainty of the previous robot poses
- uncertainty of the real landmarks
- uncertainty of the virtual landmarks

A2C+GG-NN

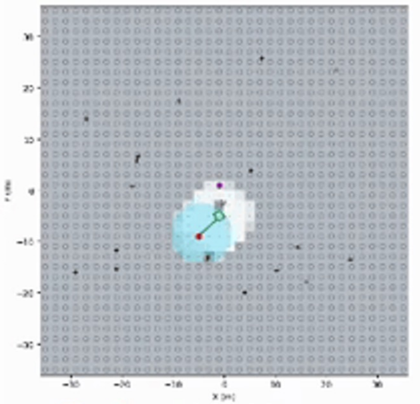

Test on 60x60m map

A2C+GG-NN

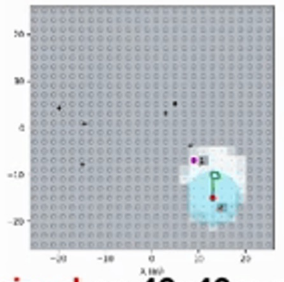

Trained on 40x40m map

A2C+GG-NN

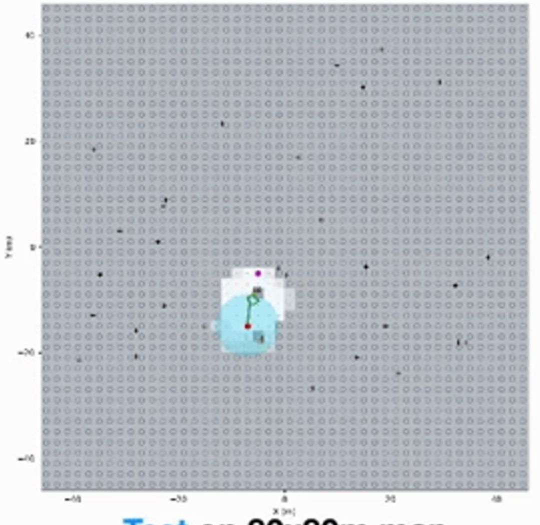

Test on 80x80m map

A2C+GG-NN

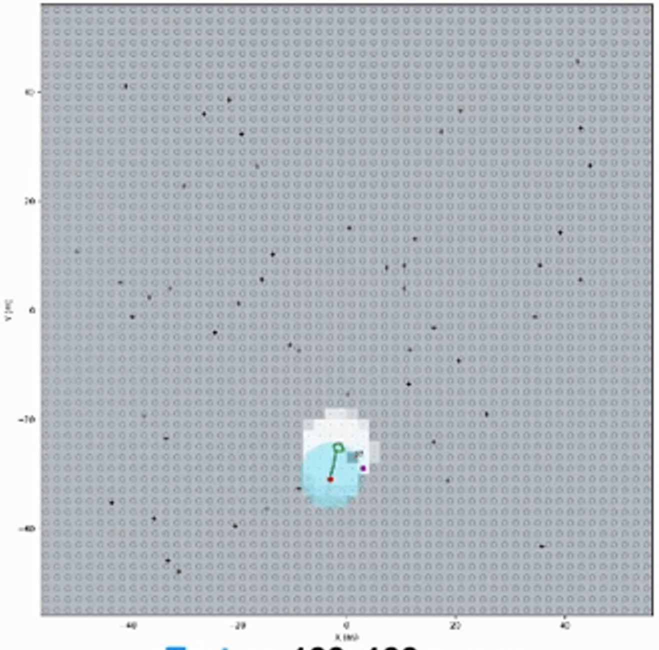

Test on 100x100m map

Scaling up to 3D with Zero-Shot Transfer

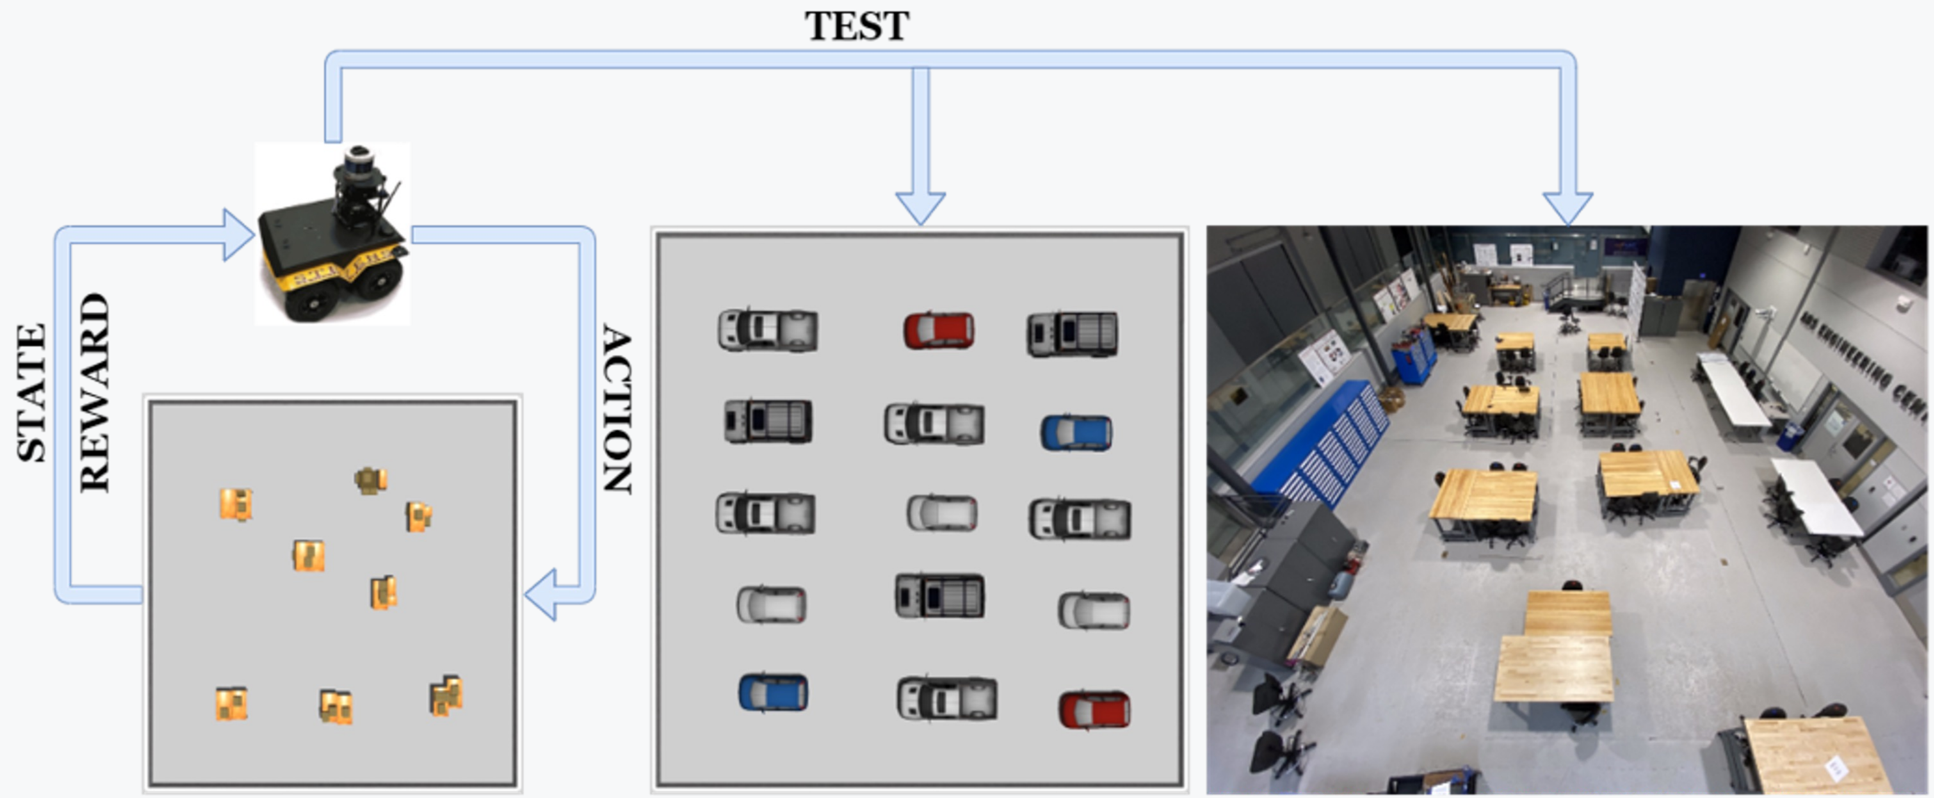

- Can we use the lowest number of training environments (e.g., a single environment)?
- Can the policy be used in a testing environment of a different size, containing new objects?
- Can the policy from the simulation guide the robot in a real environment?

Exploration Graph for Transfer Learning

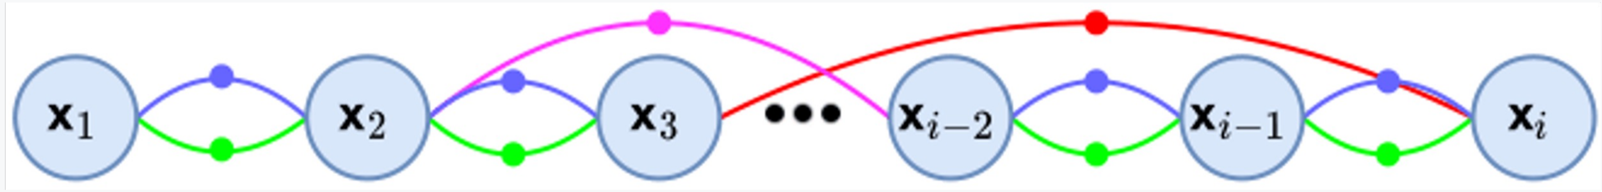

blue factors are provided by odometry measurements;
green factors are obtained from sequential scan matching of two consecutive poses;
red factors represent the loop closures provided by point cloud segment matching;
magenta factors are loop closures generated by pose matching.

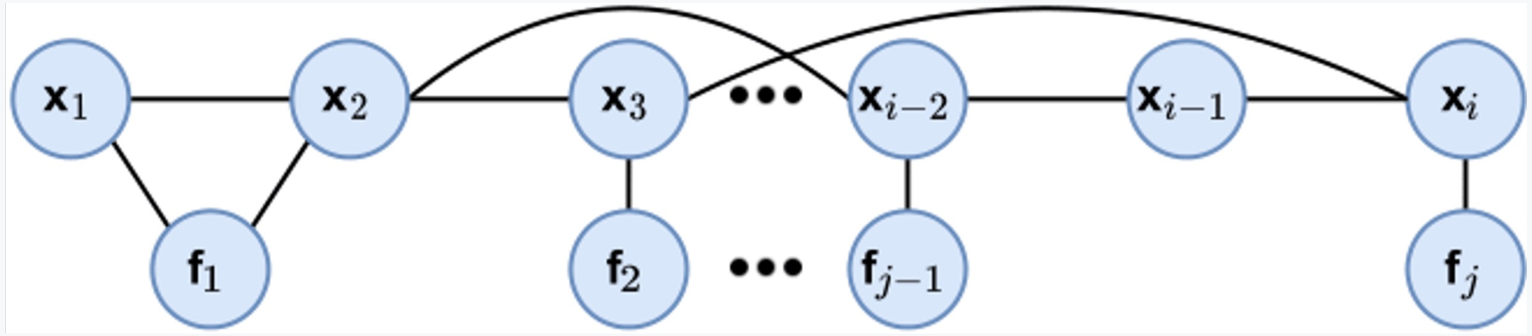

The current pose x is connected to the nearest frontier, and any frontiers whose paths achieve place revisiting are connected to the prior poses they revisit.

Exploration Graph for Transfer Learning

The cyan path is the path to the selected frontier

Pink lines are predicted loop closure constraints

The yellow ellipsoid is the uncertainty of the current pose

Gray paths are for unselected frontiers

The cost map represents the uncertainty of the current virtual map

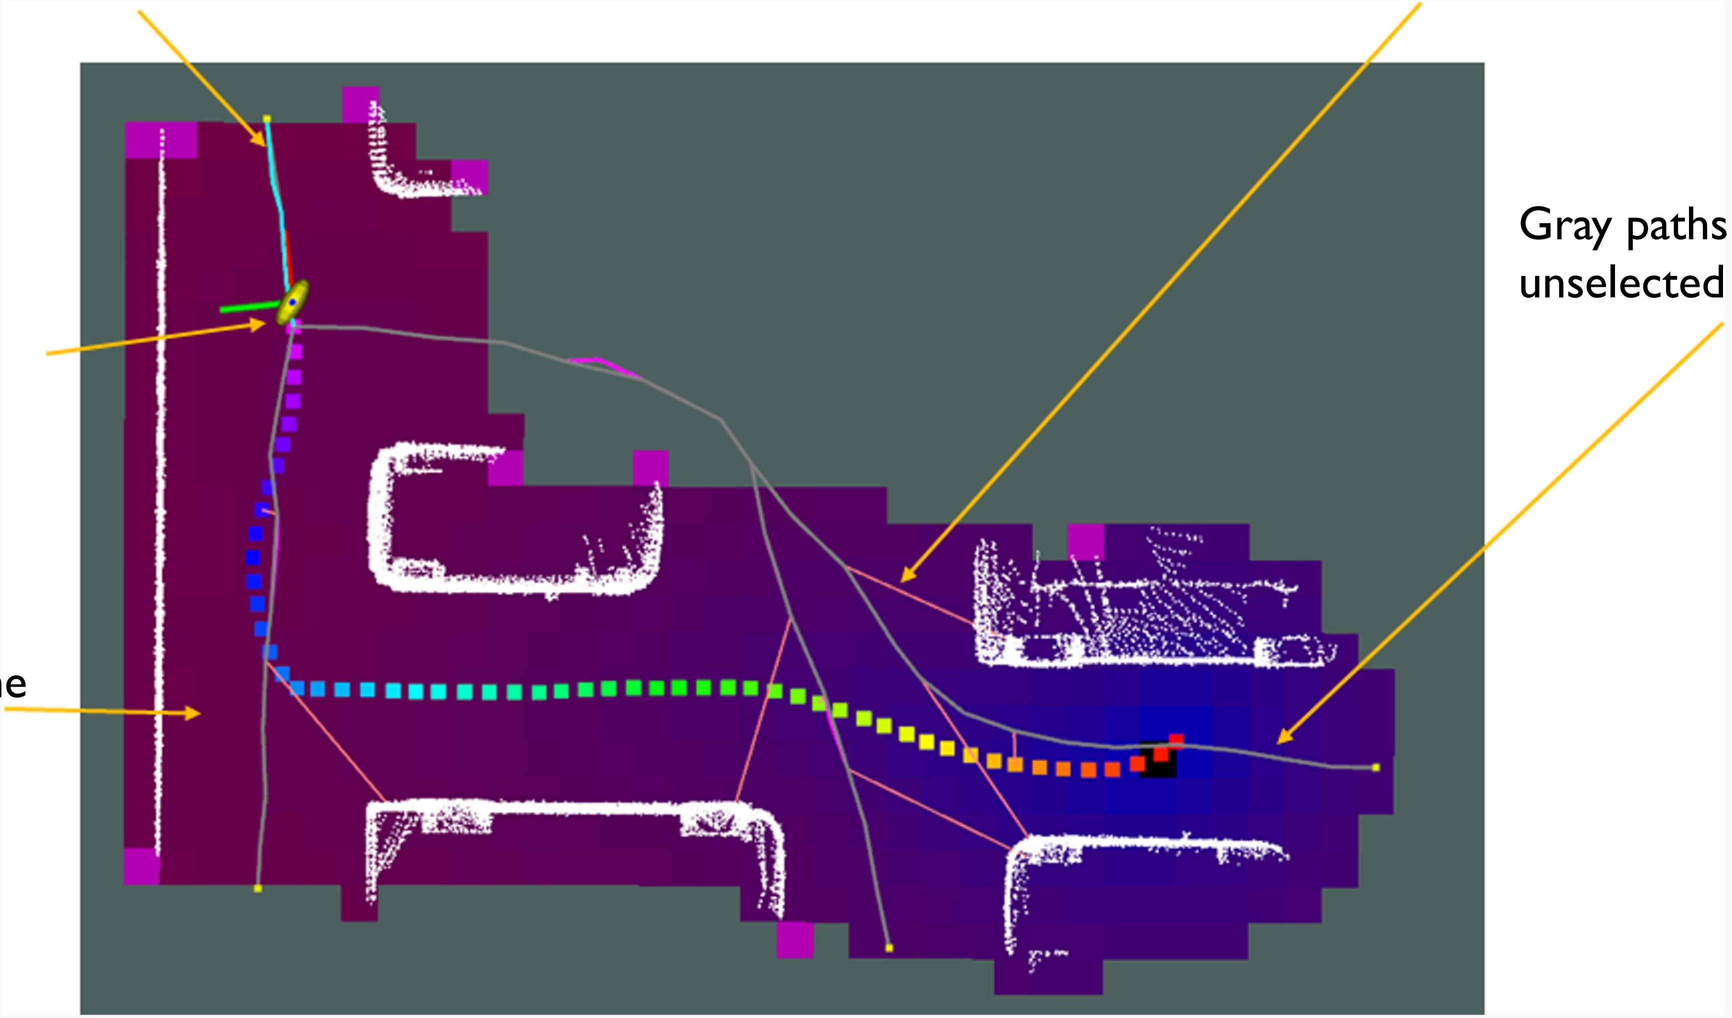

Exploration Graph for Transfer Learning

pose node

frontier node

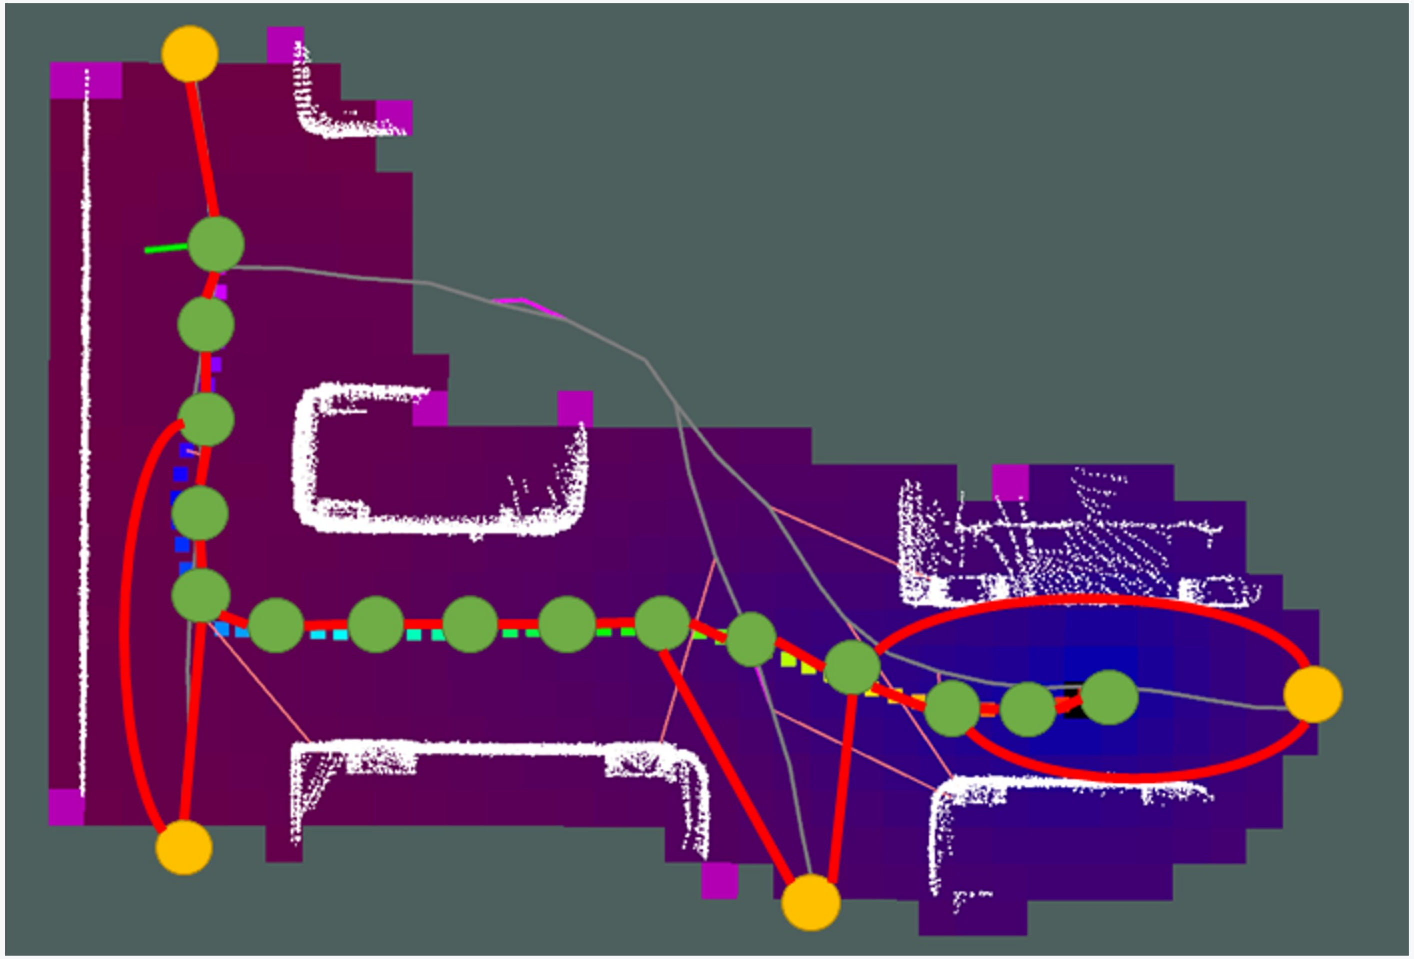

Exploration Experiment

Real-world Testing

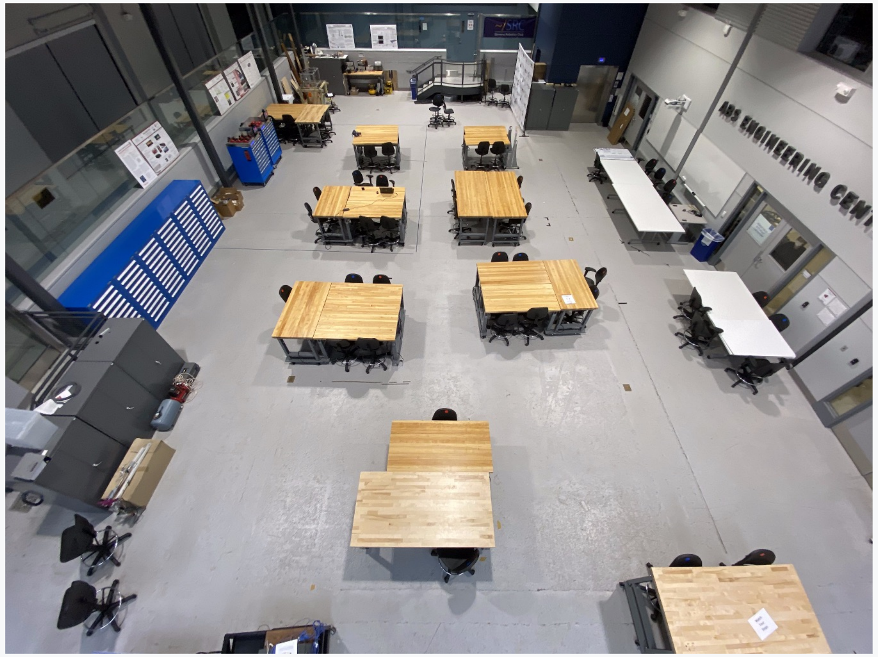

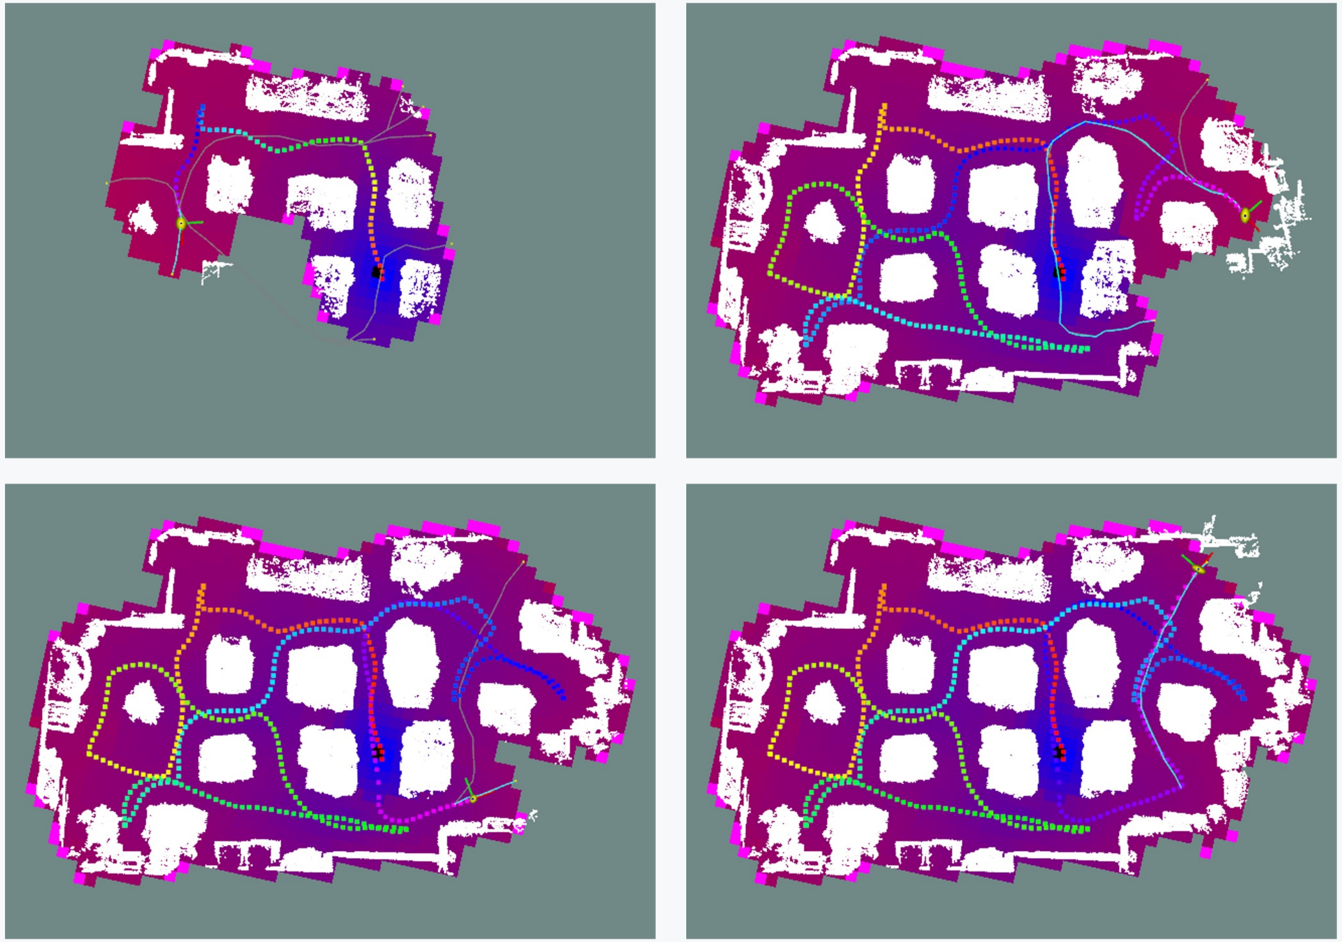

Next Steps: Learning from A Portfolio of Algorithms

Uncertainty-aware and Heuristic Exploration Algorithms have different advantages/disadvantages in different environments

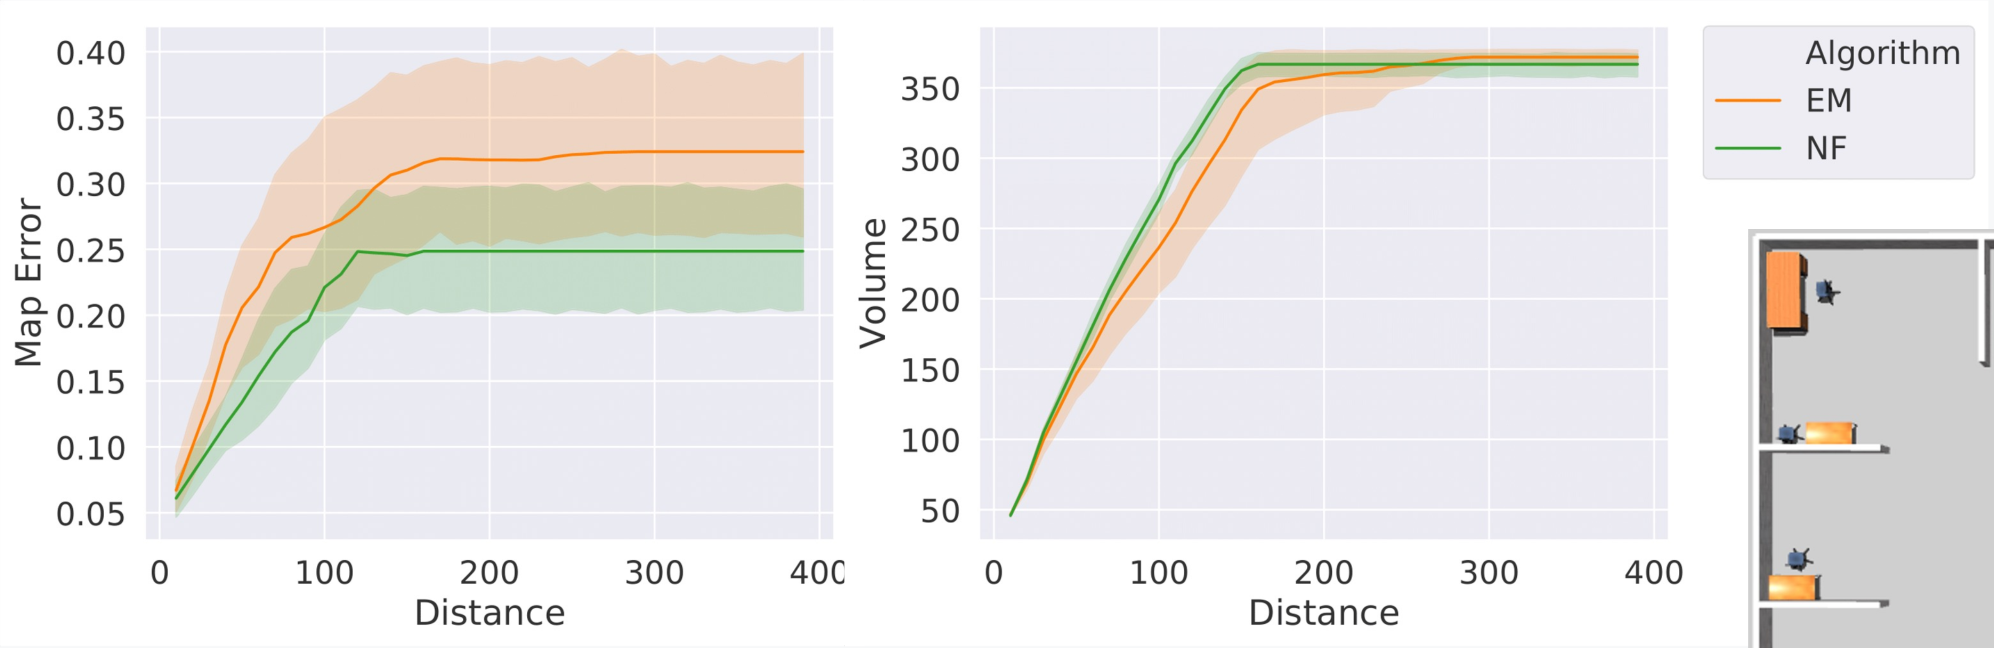

Next Steps: Learning from A Portfolio of Algorithms

Uncertainty-aware and Heuristic Exploration Algorithms have different advantages/disadvantages in different environments

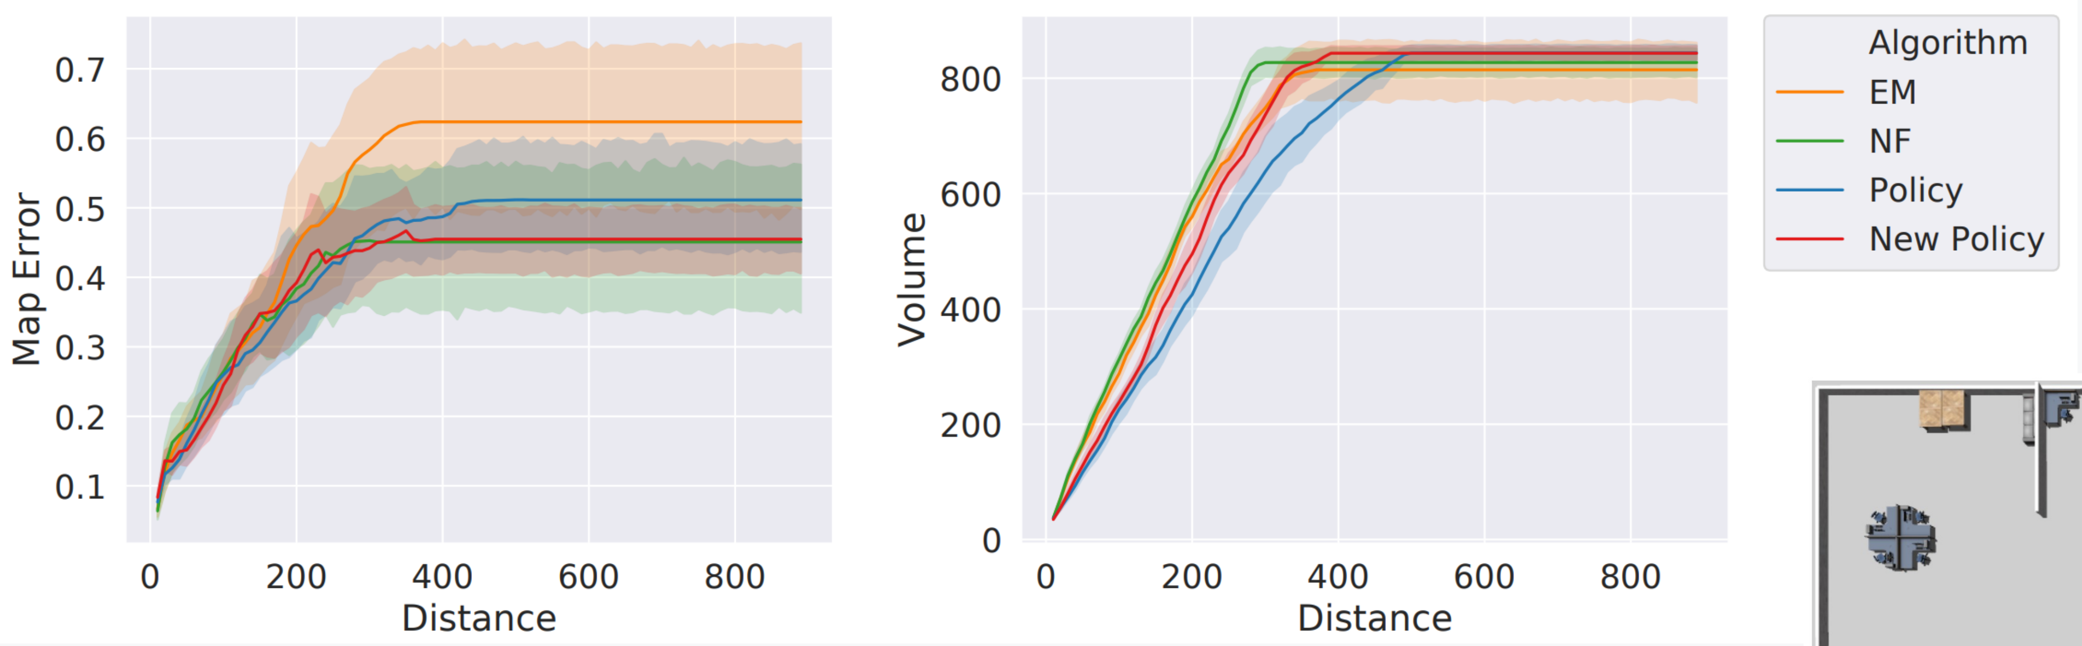

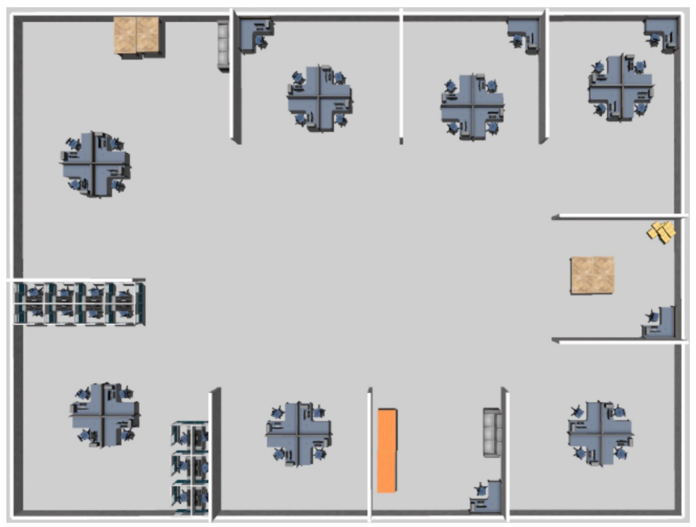

Next Steps: Autonomous Navigation with Multi-Robot Teams

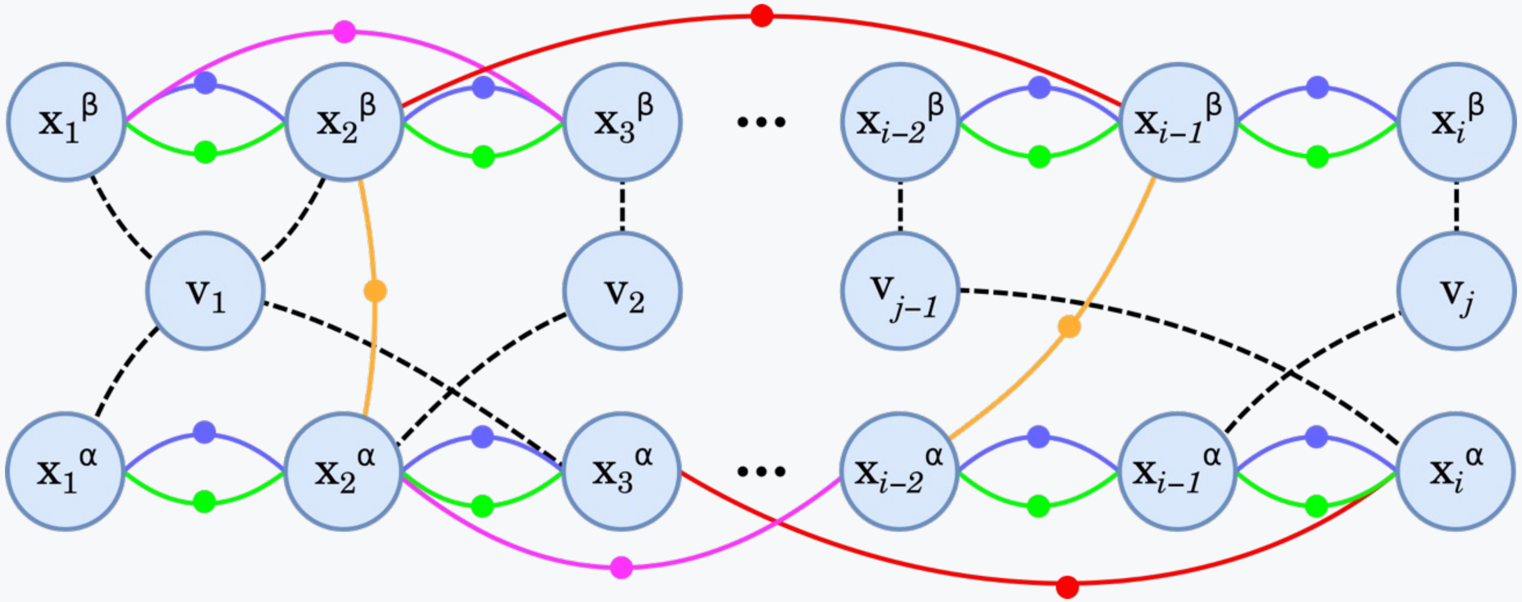

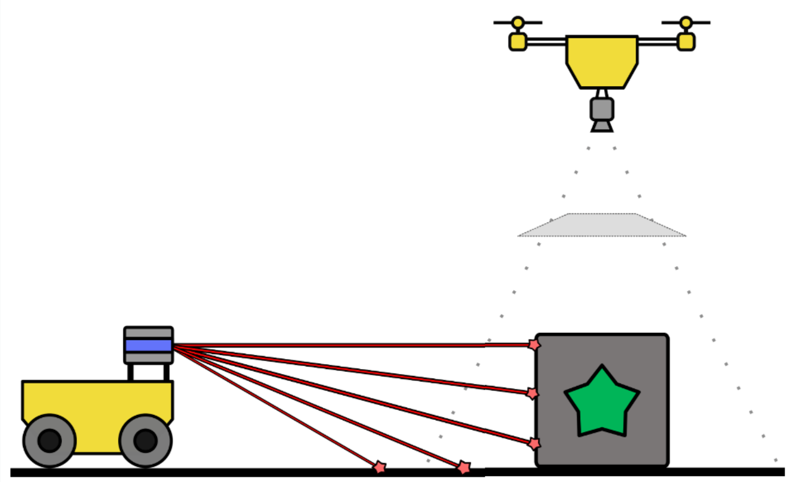

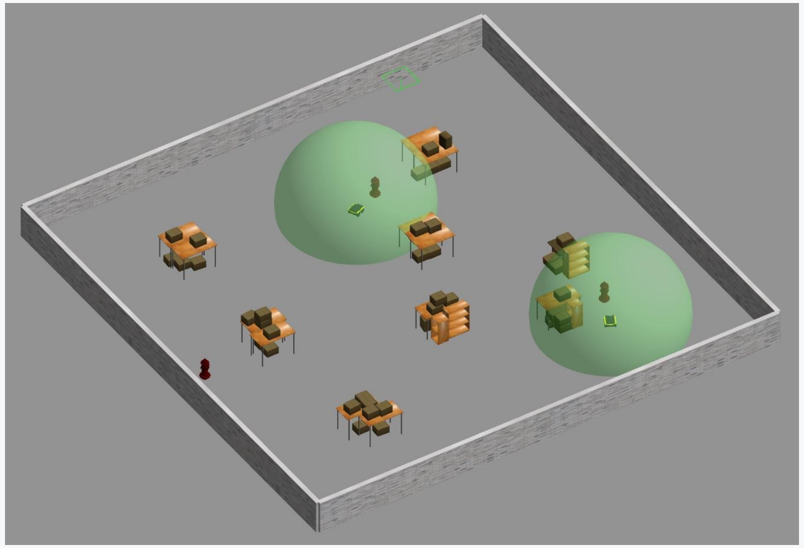

- Currently building a simulation to support heterogeneous multi-robot systems exploring an unknown environment
- Implementing the more generalized problem of "persistent monitoring" – the robots must repeatedly observe a collection of targets greater in number than themselves

THANK YOU

Stay connected with us online.