
# Virtual Research Presentation Conference

## Hierarchical Information-based Planning for Coverage in Large Unknown Environments

**Principal Investigator: Sung Kim (347)**
**Co-Is: Kyohei Otsu (347), Amanda Bouman (347), Aliakbar Aghamohammadi (347)**
**Program: Spontaneous Concept**

Assigned Presentation #: RPC-126

Jet Propulsion Laboratory
California Institute of Technology

# Tutorial Introduction


**Abstract**

In this project, we tackle the coverage planning optimization problem for large unknown environments. In ***large*** planetary surface/subsurface with ***no/little map*** information, autonomous exploration for ***coverage***, i.e., mapping of the environment and detection of targets of interest, is one of the key technologies for NASA's future missions. For example, Lunar Surface Innovation Initiative (LSII) seeks extreme access capability for efficient exploration in permanently shaded surfaces or unknown subsurface voids at a large scale. Such autonomous exploration capability enables detection of scientific investigation targets and/or human habitat areas, without human operators in the loop. The challenge, however, is that the optimization of coverage planning for unknown environments is hard to scale up—its *complexity grows exponentially* as the environment size gets larger. Thus, the state-of-the-art coverage planners solve for myopic, suboptimal solutions for the local area around the robot. In this work, we developed an **efficient information-based coverage** planner for large unknown environments. The two key ideas are 1) to represent the probabilistic robot and world state as a **compact graph representation** with embedded semantics, and 2) to employ **hierarchical planning-under-uncertainty framework** to pursue local optimality and global completeness at the same time.


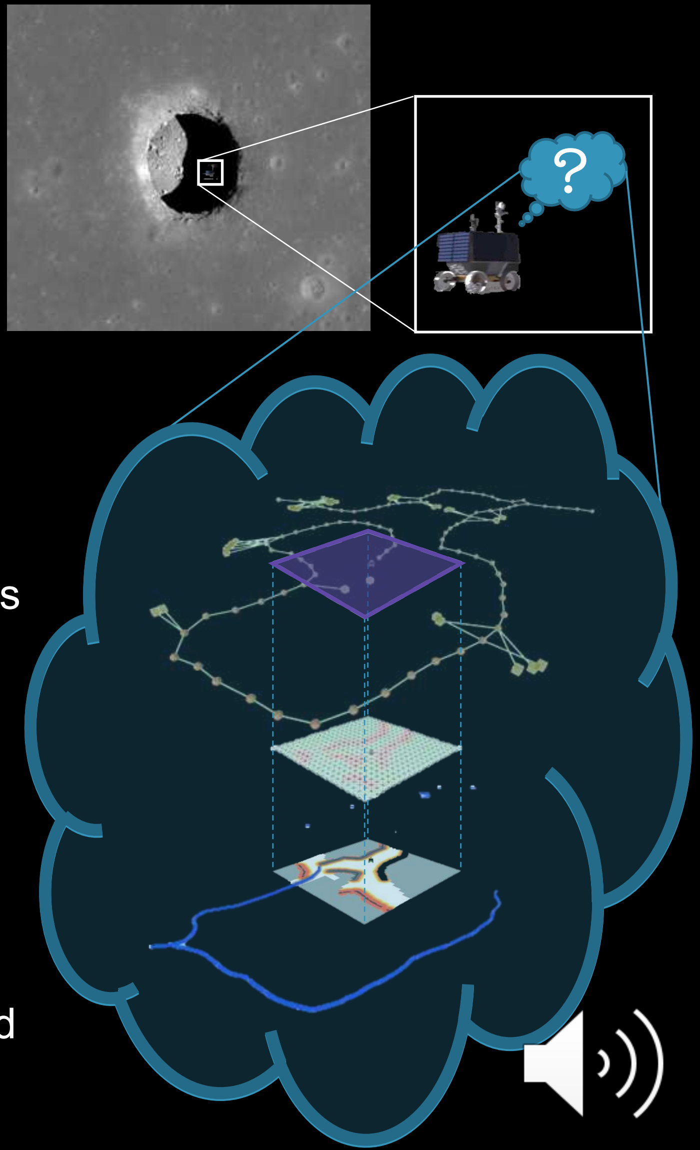

# Problem Description

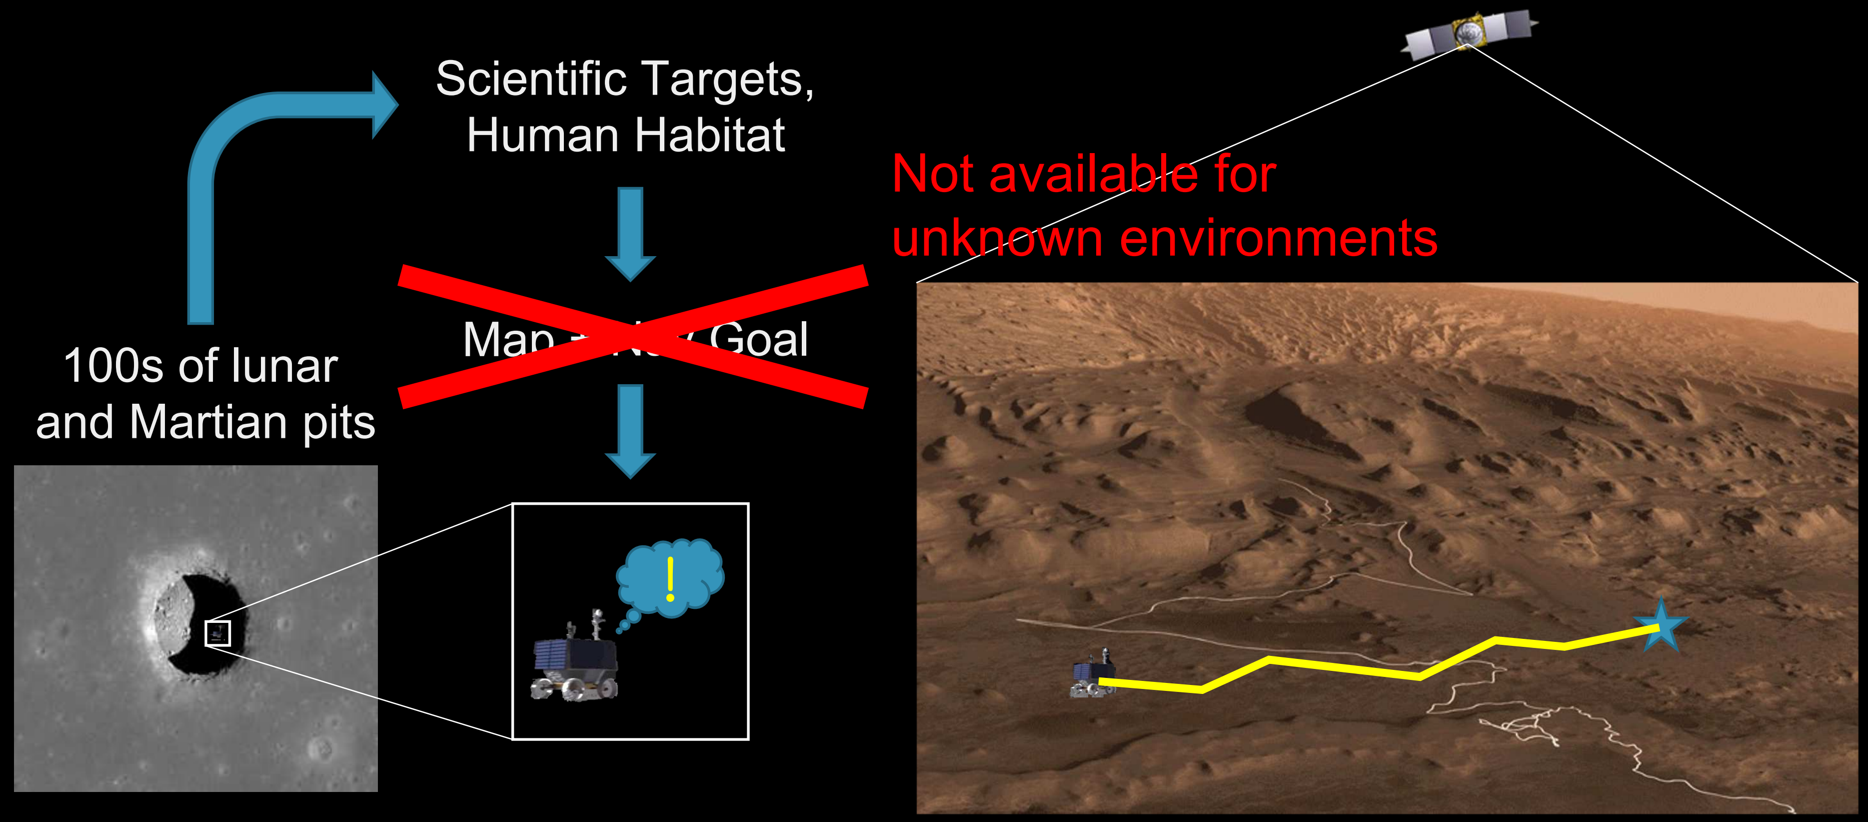

# Problem Description

Scientific Targets,
Human Habitat

No Map + No Nav Goal

100s of lunar
and Martian pits

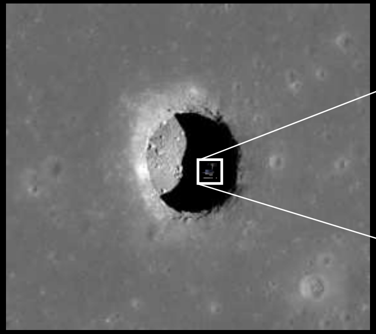

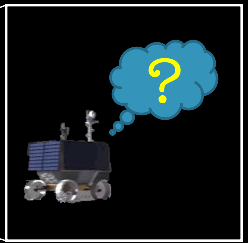


**State-of-the-art Coverage Planners**

- Next-Best-View: Myopic, suboptimal solution
- Belief Space Planner: Not scalable to large problems

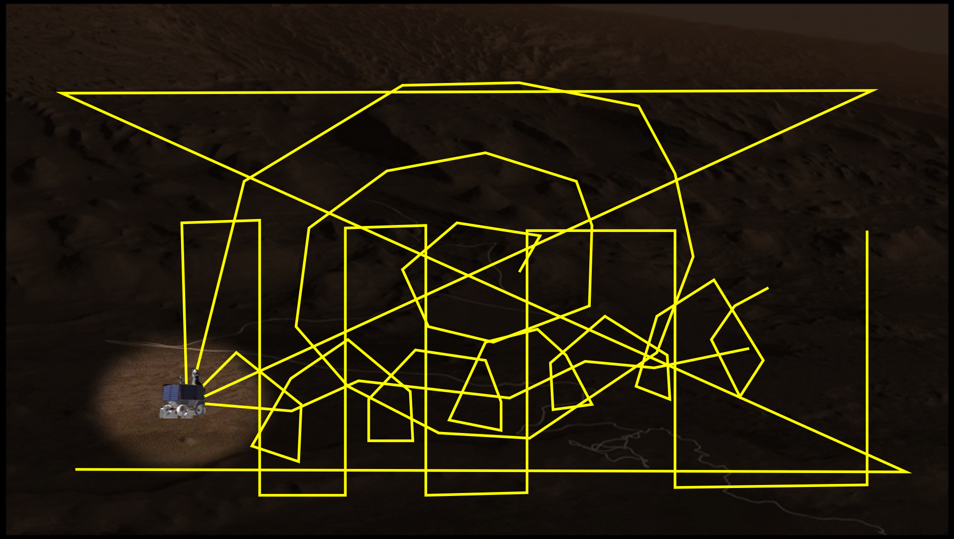

Computationally intractable
to solve for exact solutions!

**Auto-Exploration**

1) Build the map &
2) Find targets of interest
3) Given limited time/resources

# Methodology

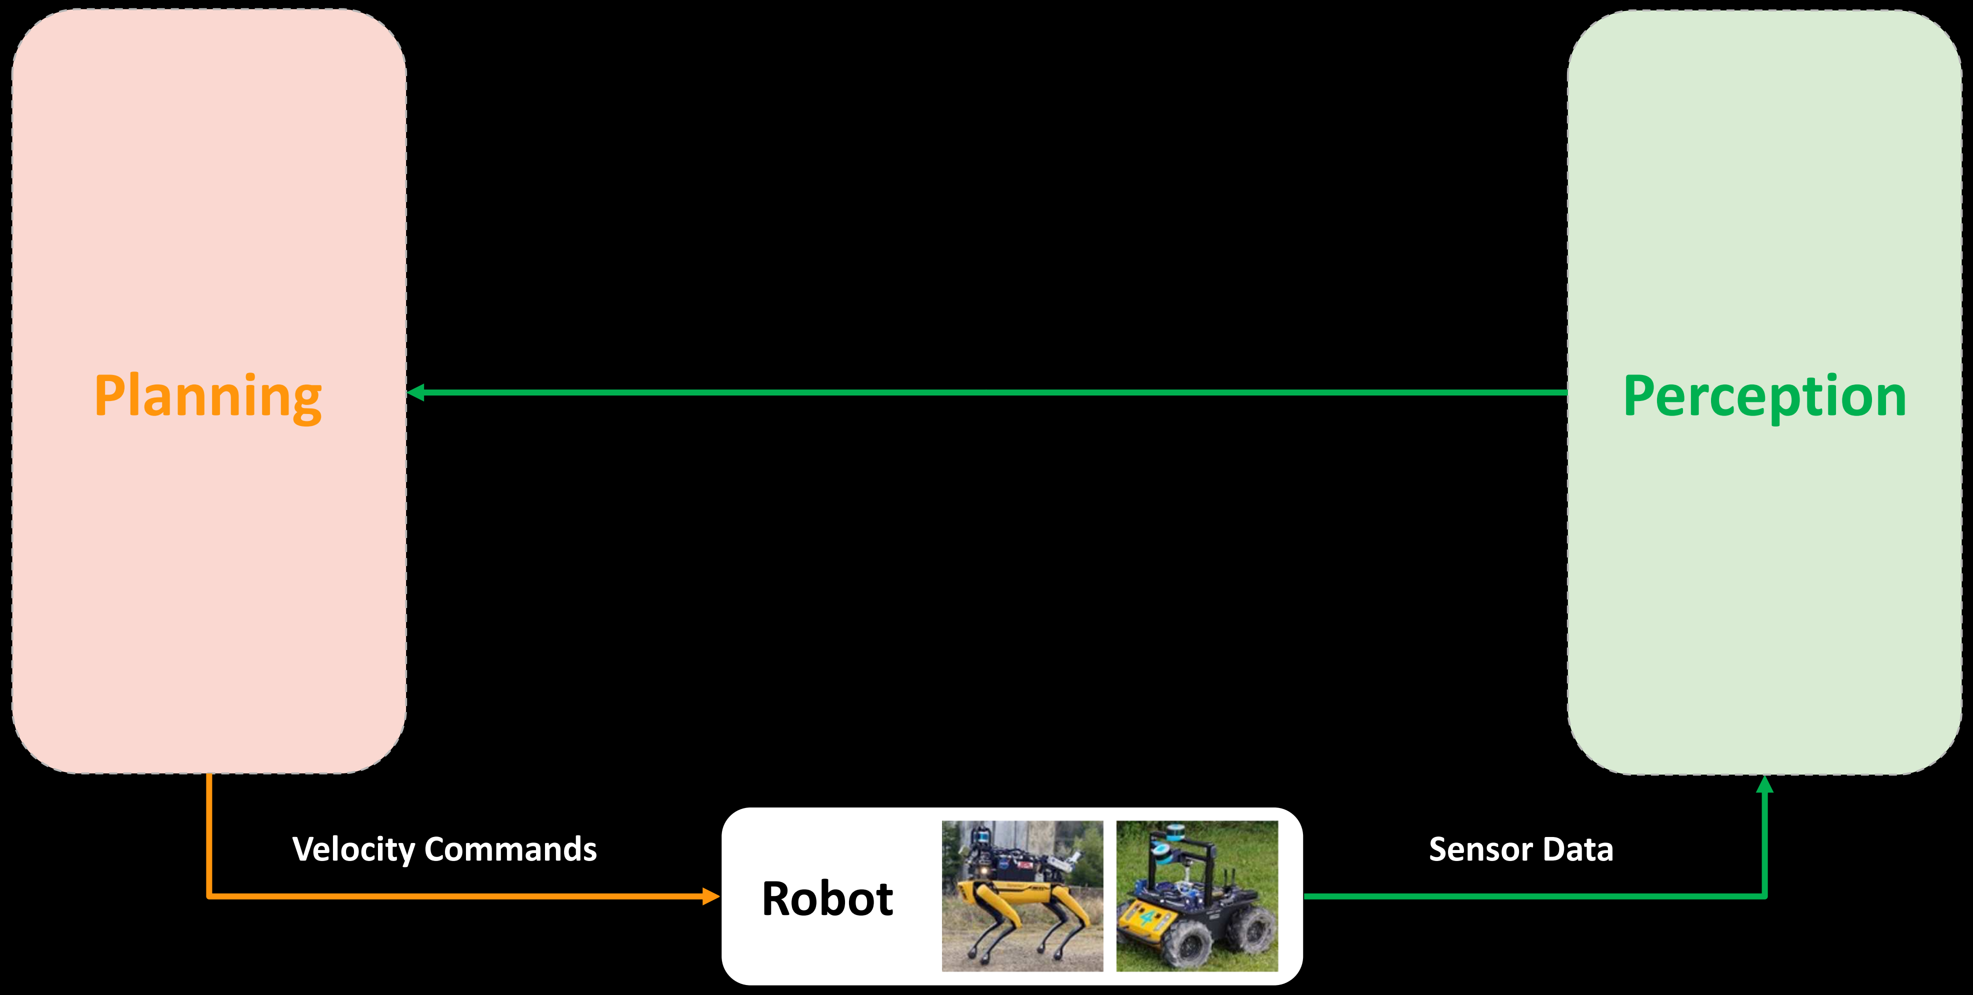

# Methodology

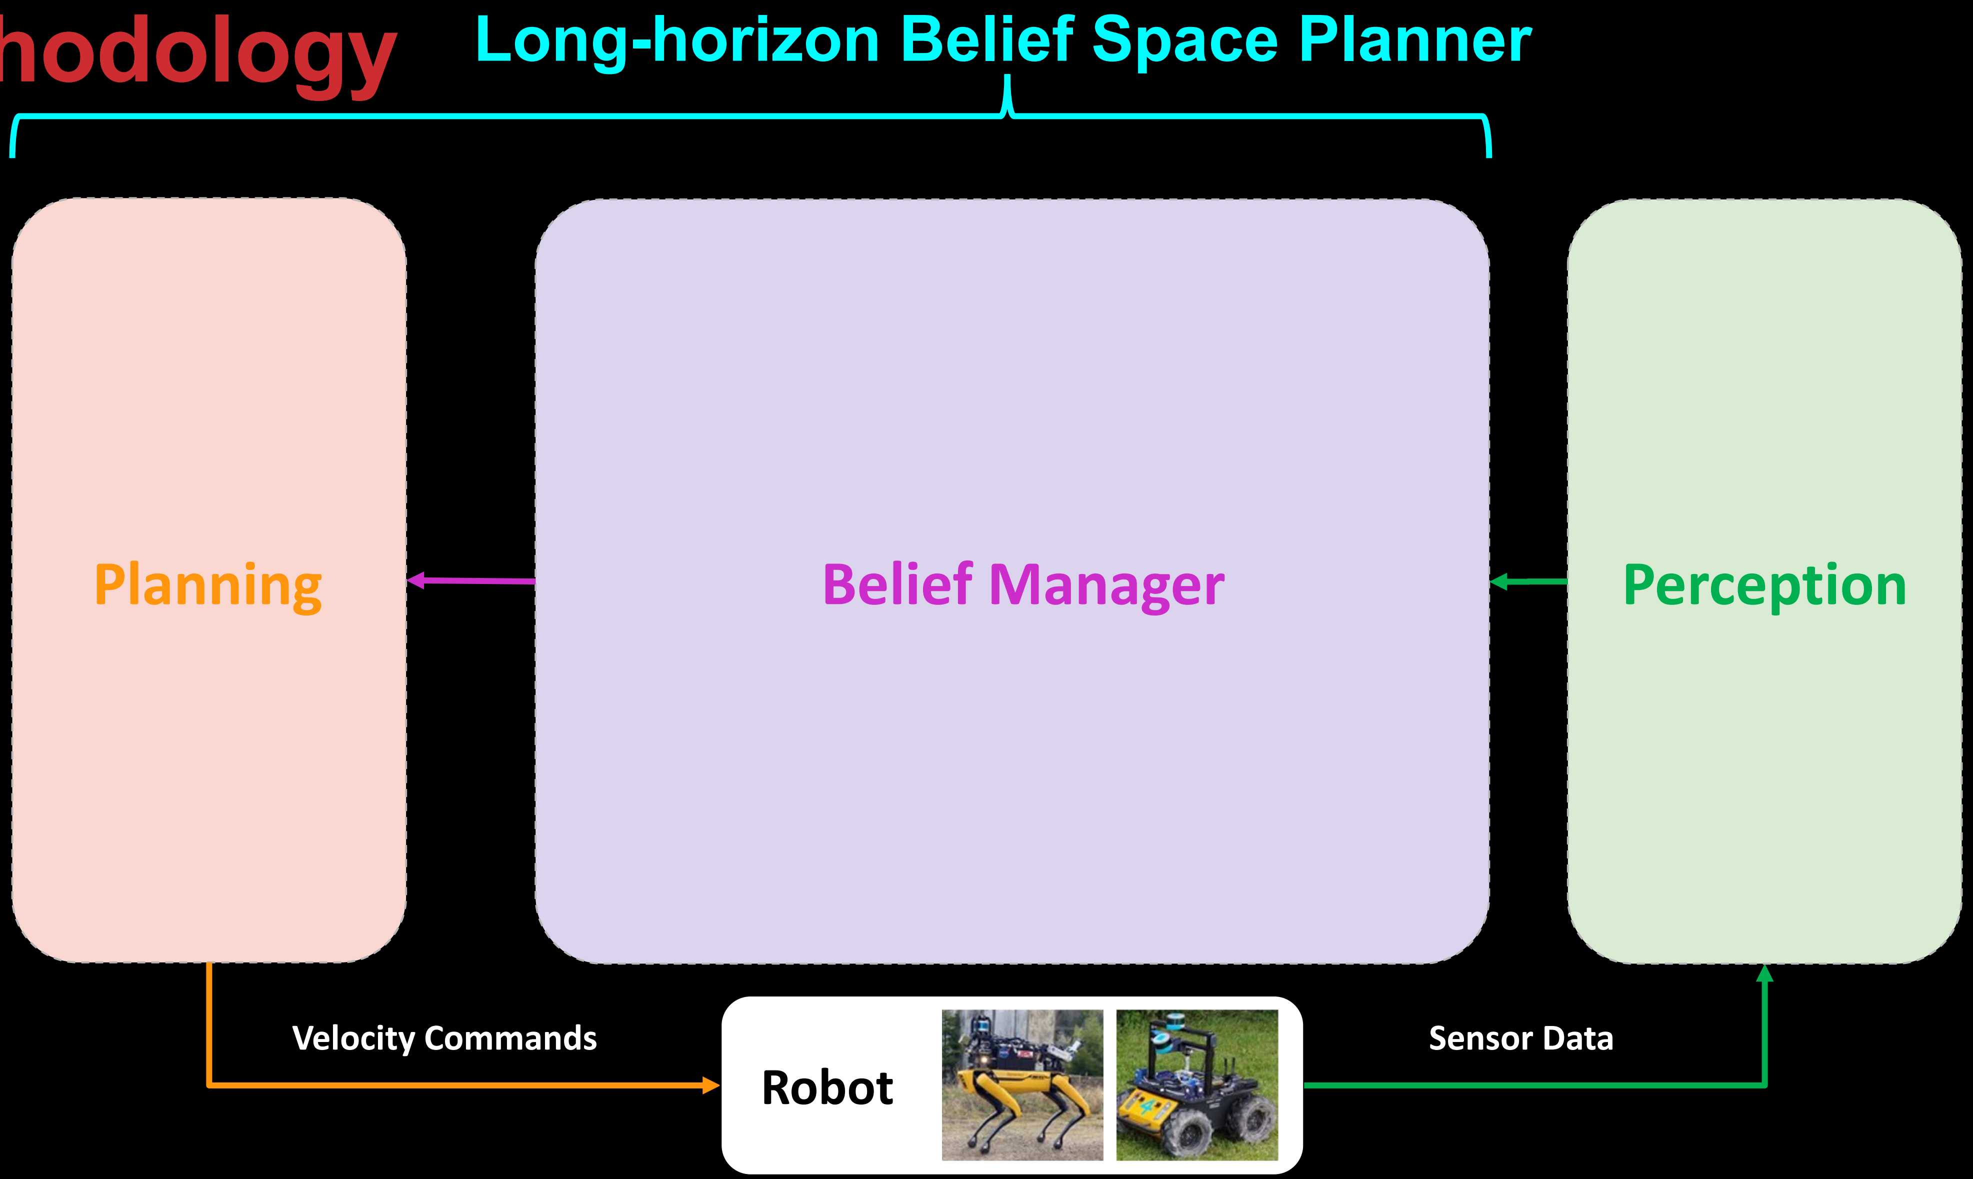

# Methodology

## Long-horizon Belief Space Planner

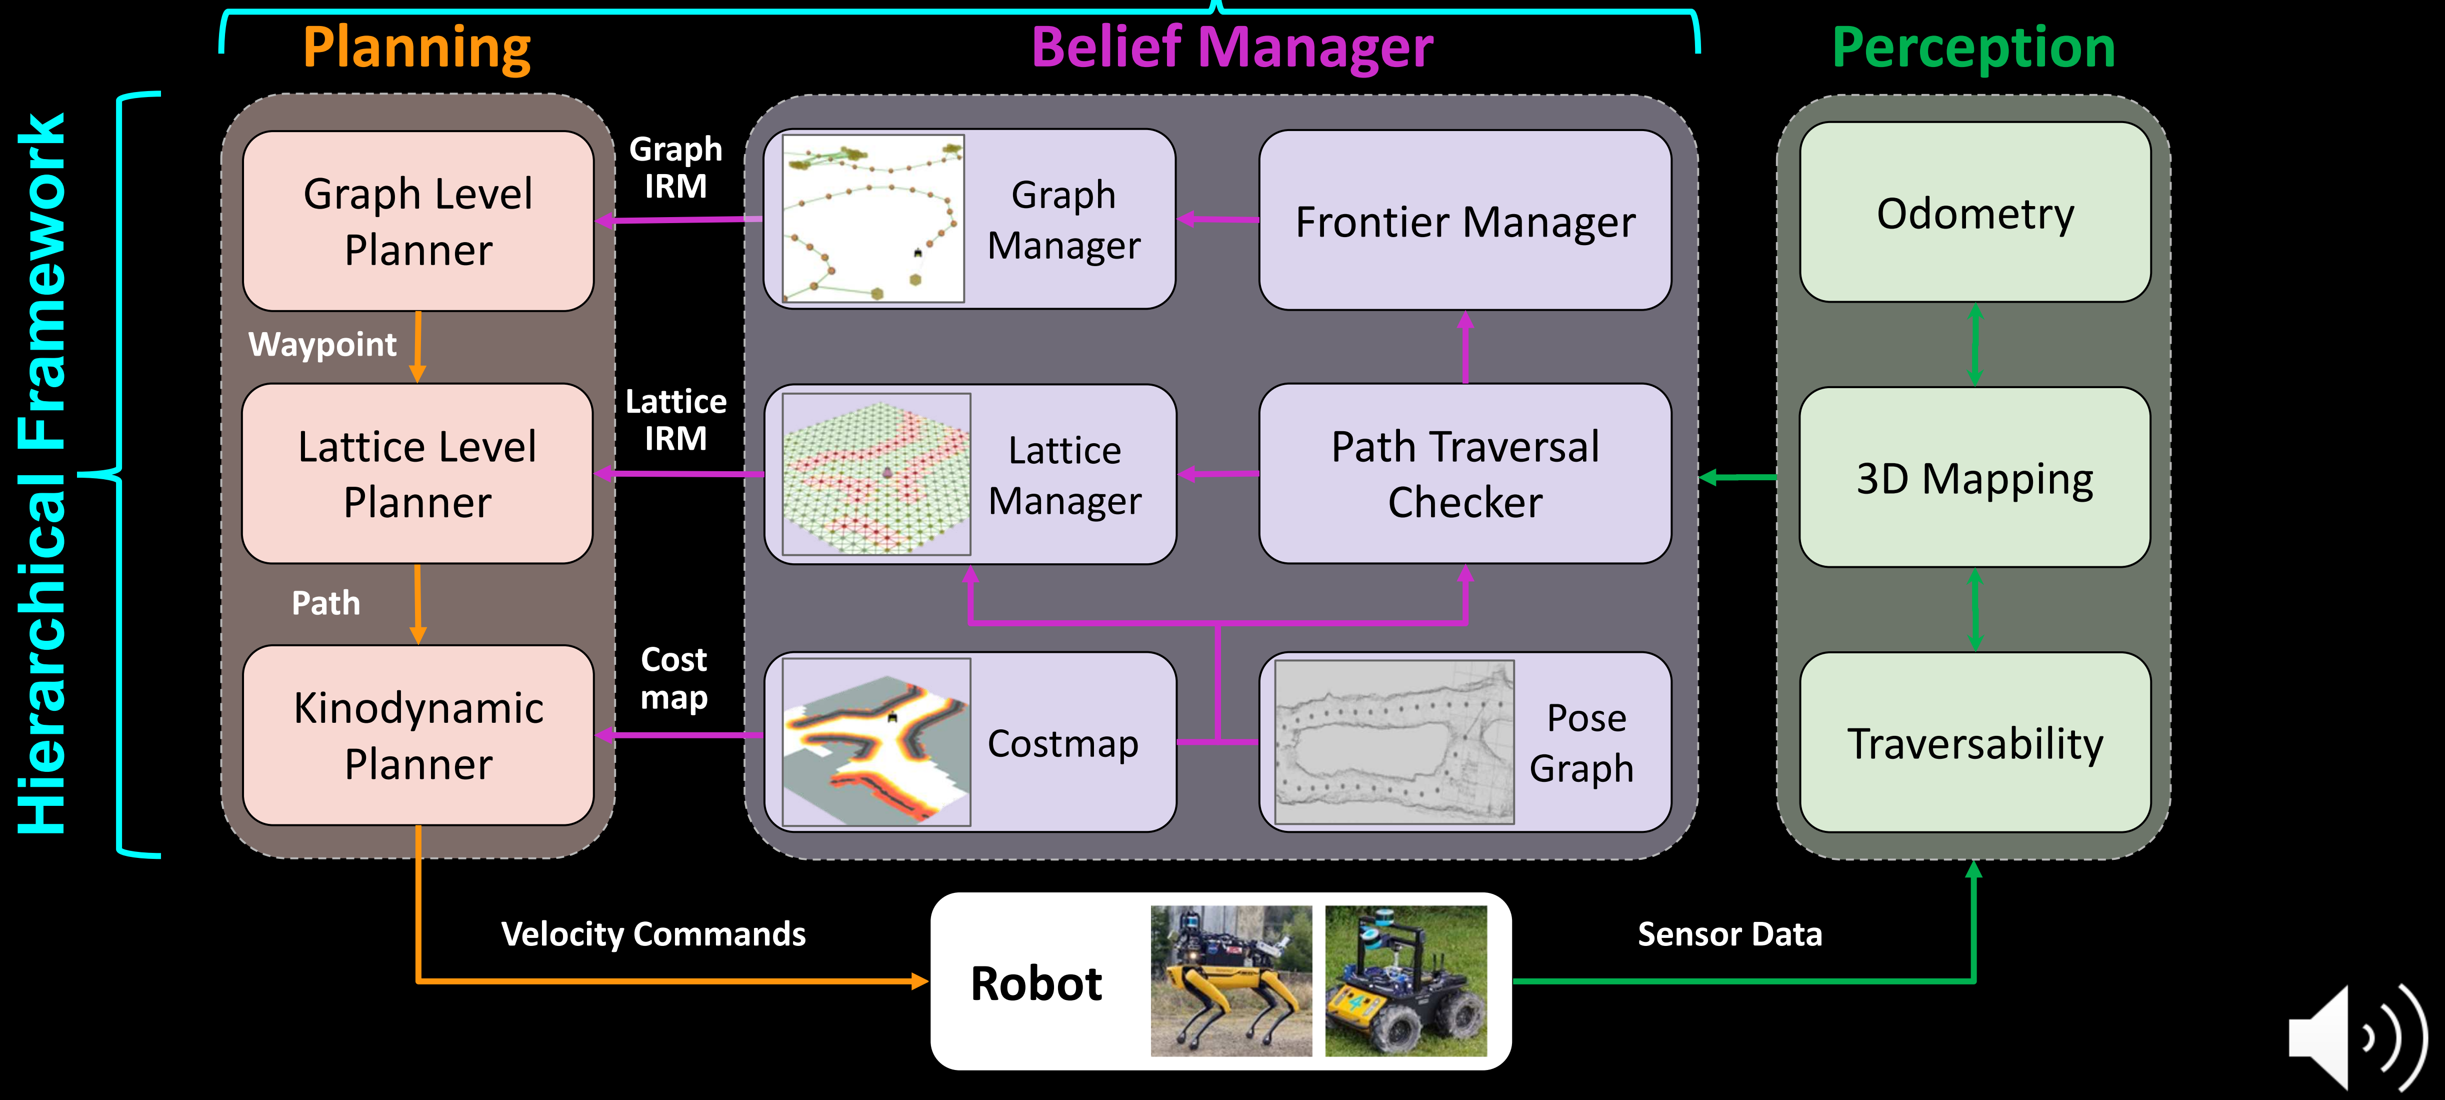

# Methodology

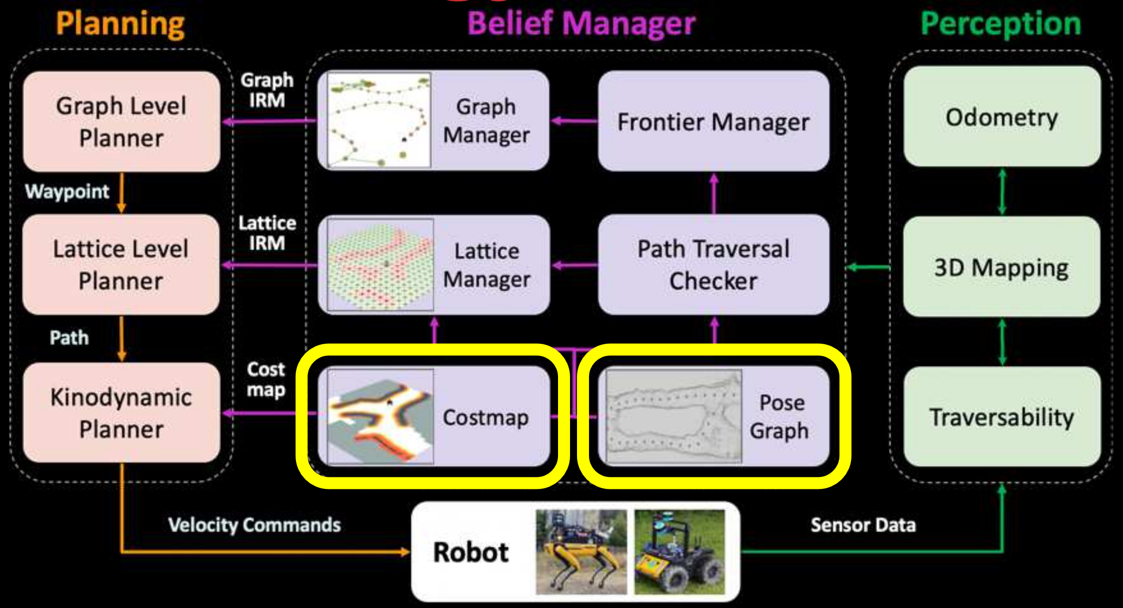


**Belief State:**

$$b = p(s) = p(W, Q)$$

World State

$$W = (W_{occ}, W_{cov})$$

Robot Pose

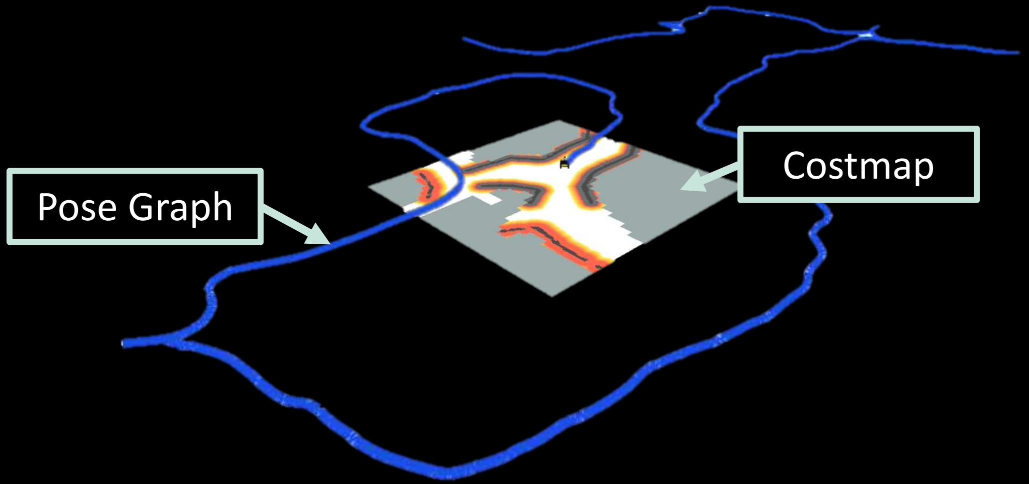

# Methodology

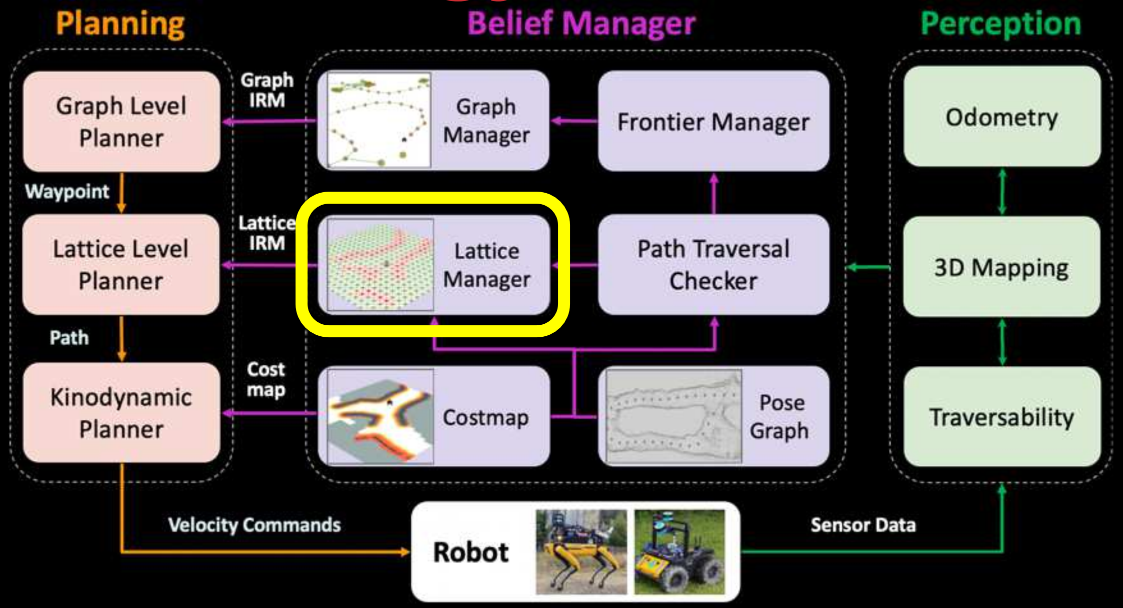


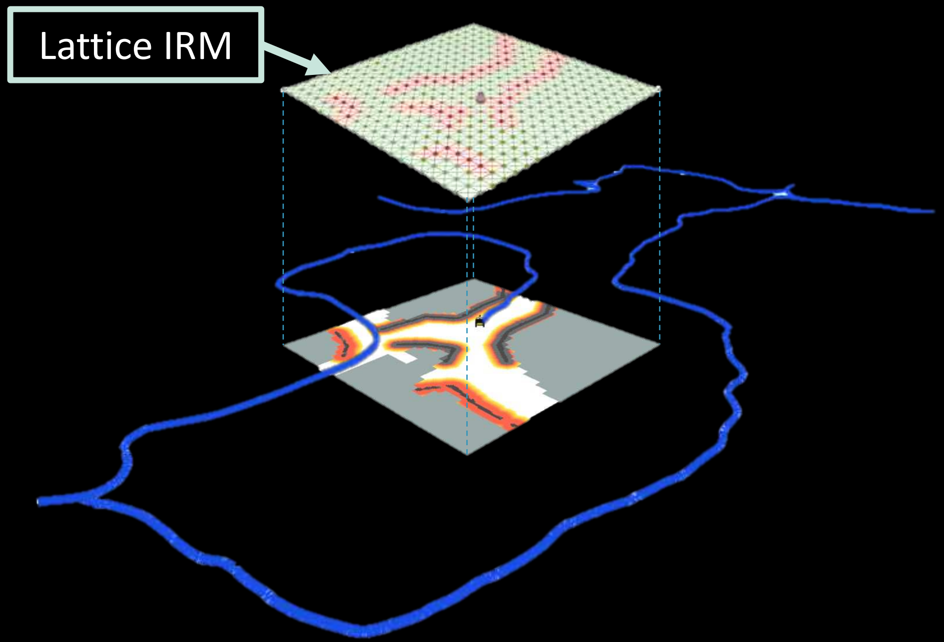


**Local Belief State:**

$$\boxed{b^{\ell} = p(s^{\ell}) = p(W^{\ell}, Q^{\ell})}$$

# Methodology

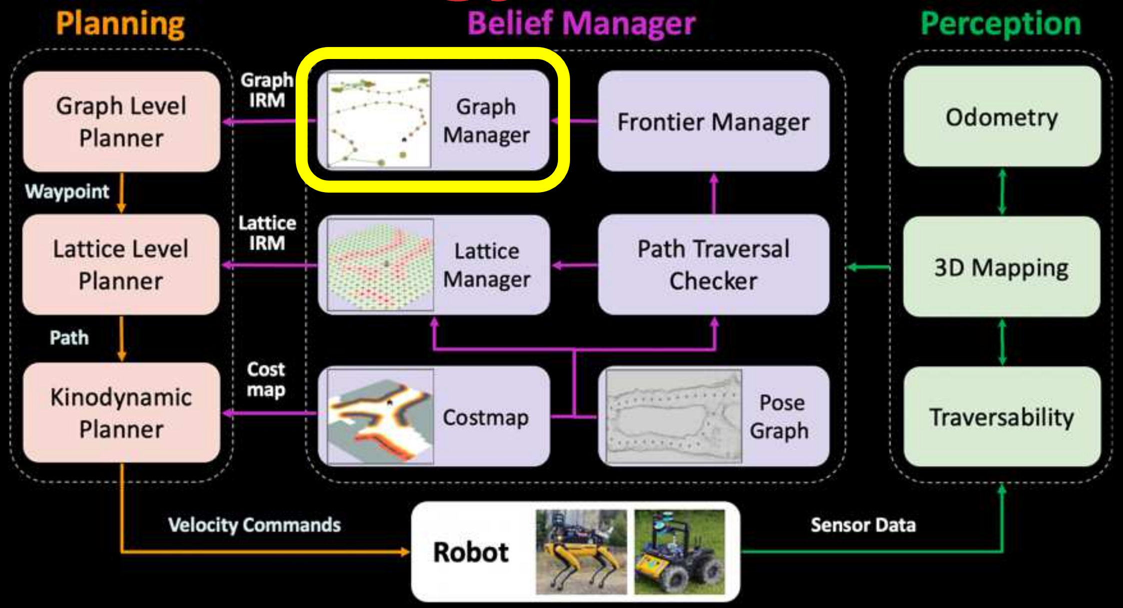


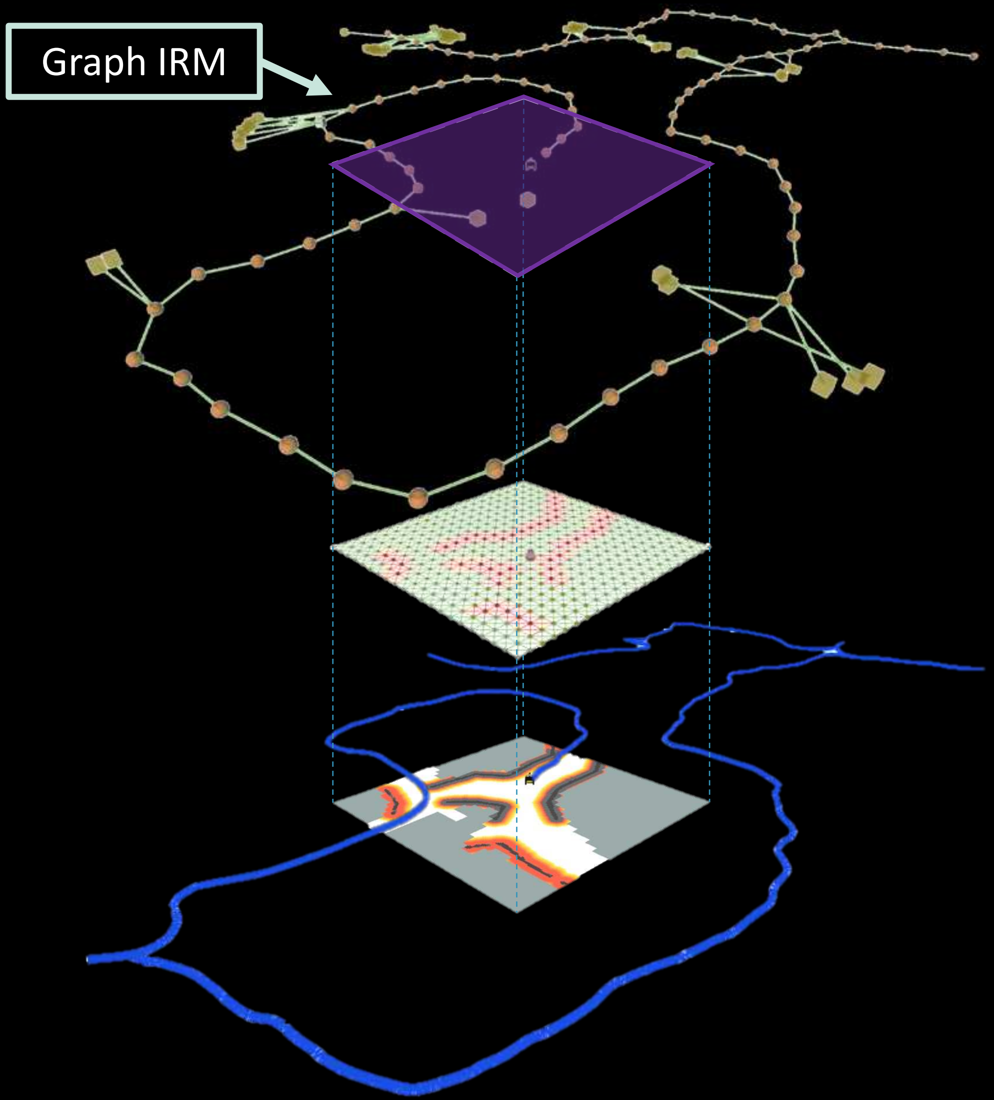


**Global Belief State:**

$$b^g = p(s^g) = p(W^g, Q^g)$$

**Local Belief State:**

$$b^\ell = p(s^\ell) = p(W^\ell, Q^\ell)$$

# Methodology

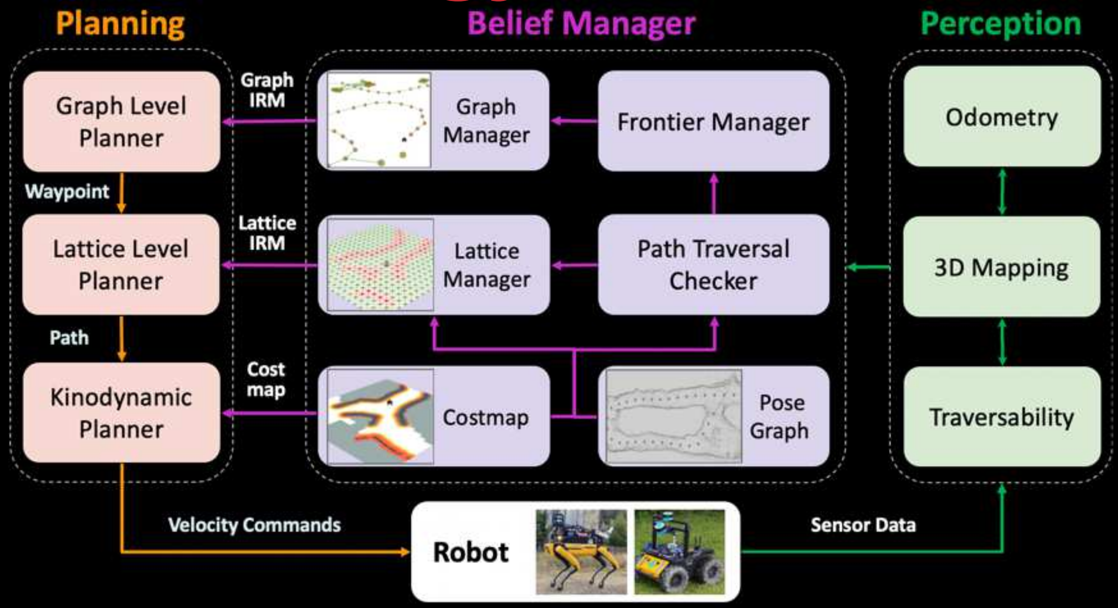


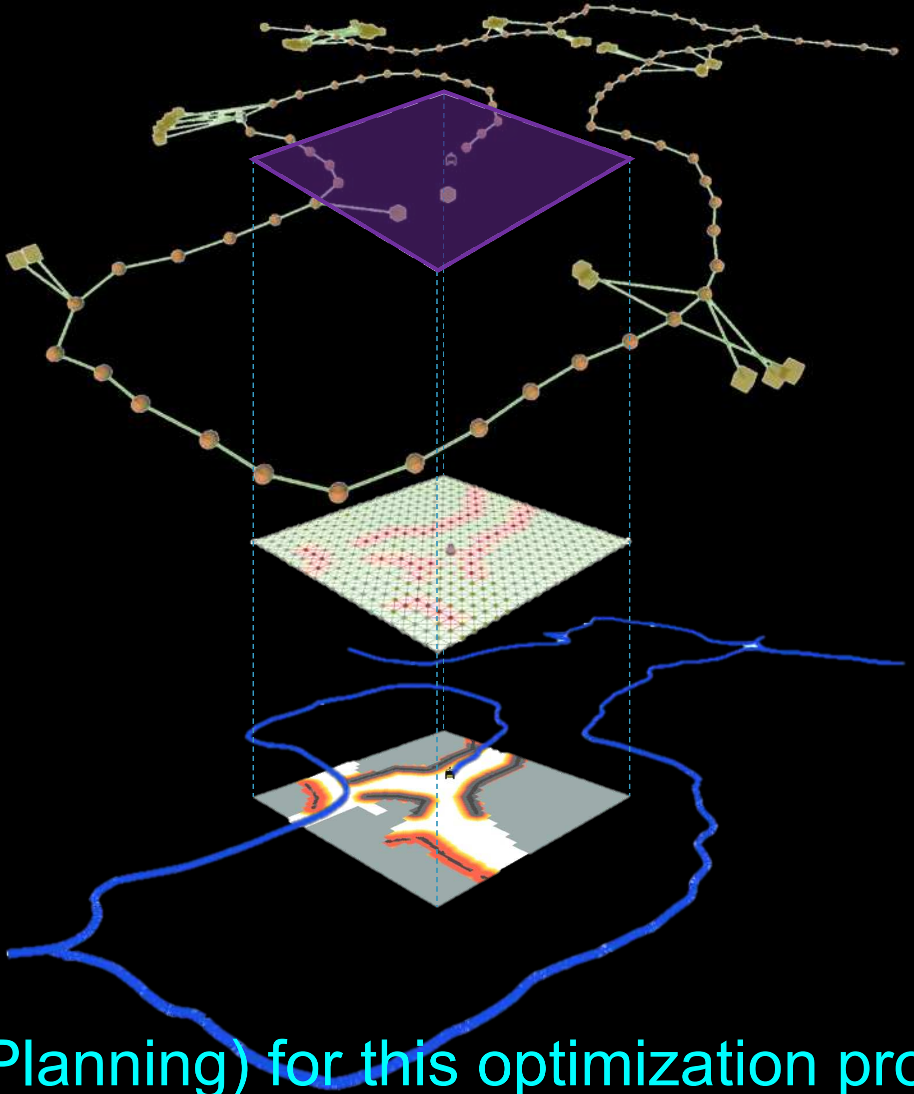

**Reward Function:**

$$R(s,a) = f\big(\underbrace{IG(W_{cov}, z)}_{\text{Information Gain}}, \underbrace{C(W_{occ}, a)}_{\text{Action Cost}}\big)$$

➔ **POMCP** (Partially Observable Monte Carlo Planning) for this optimization problem!

# Methodology

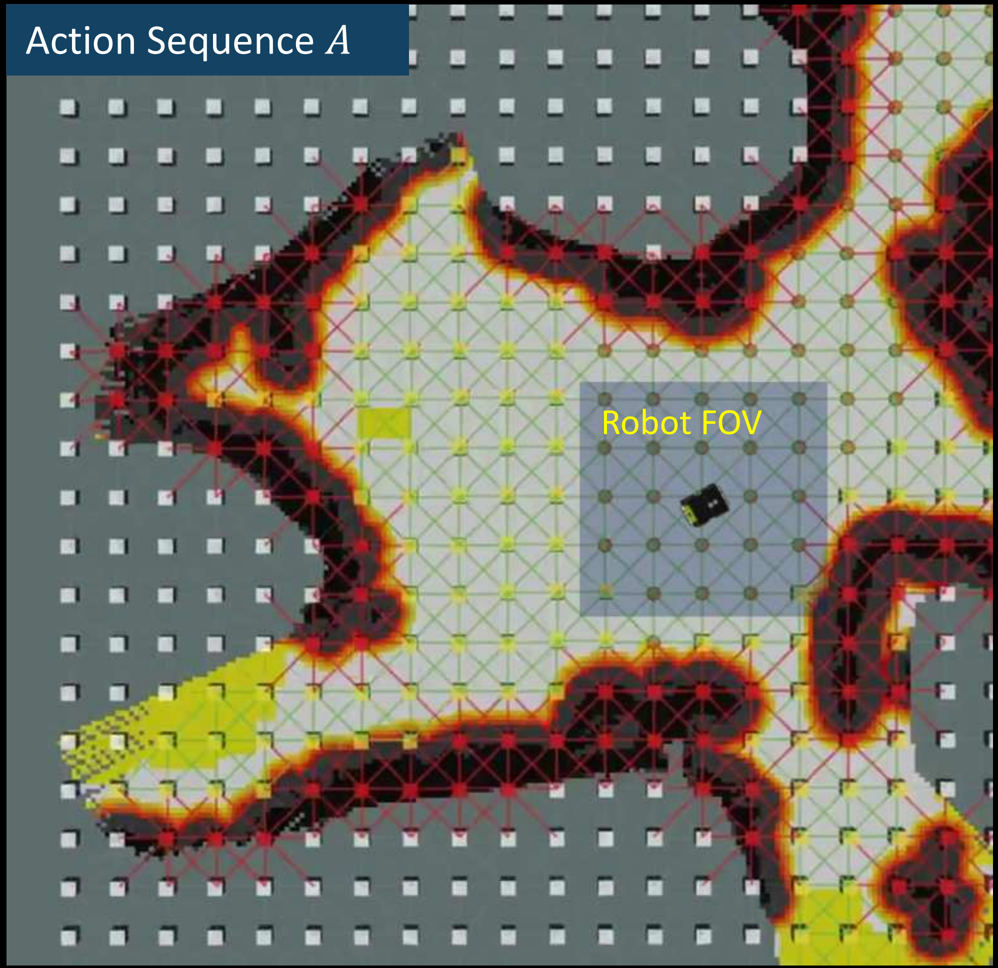


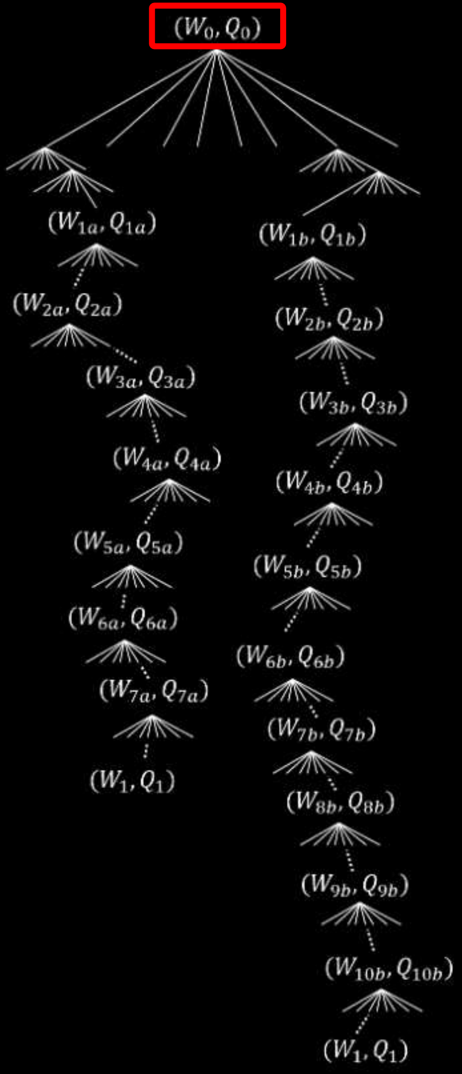


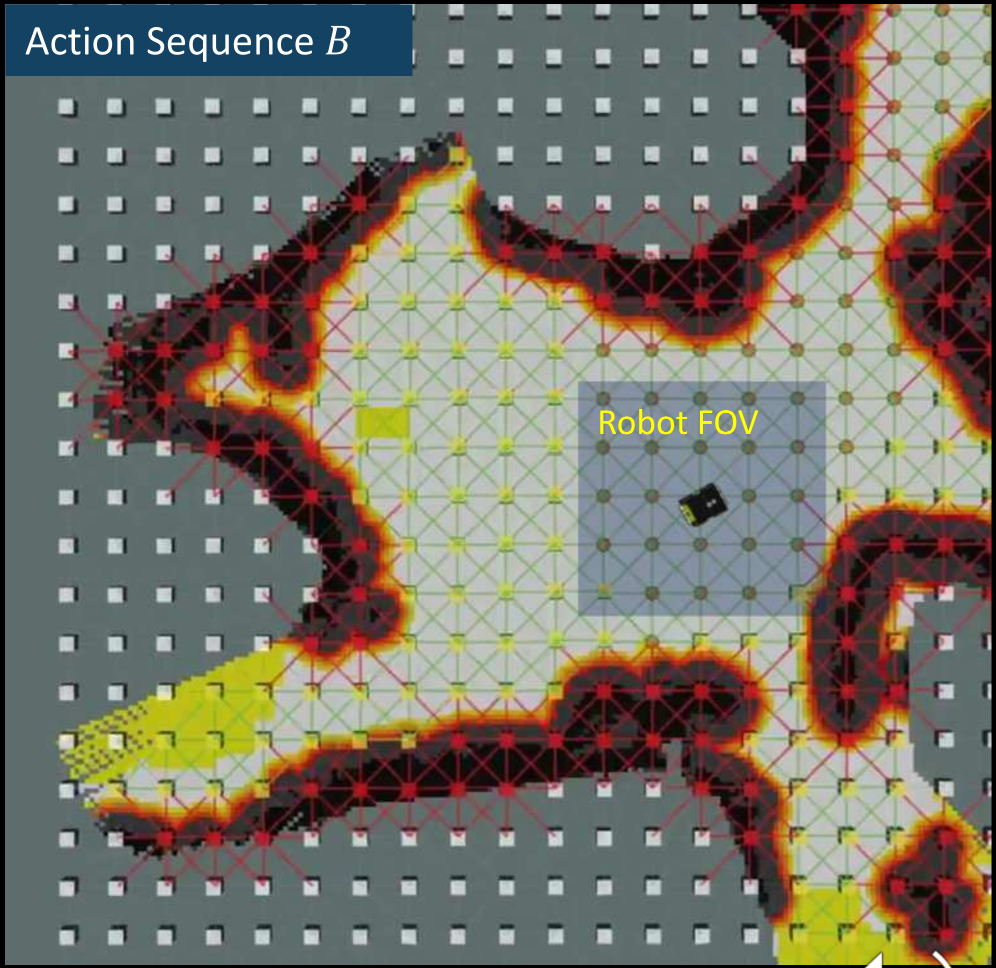

# Methodology

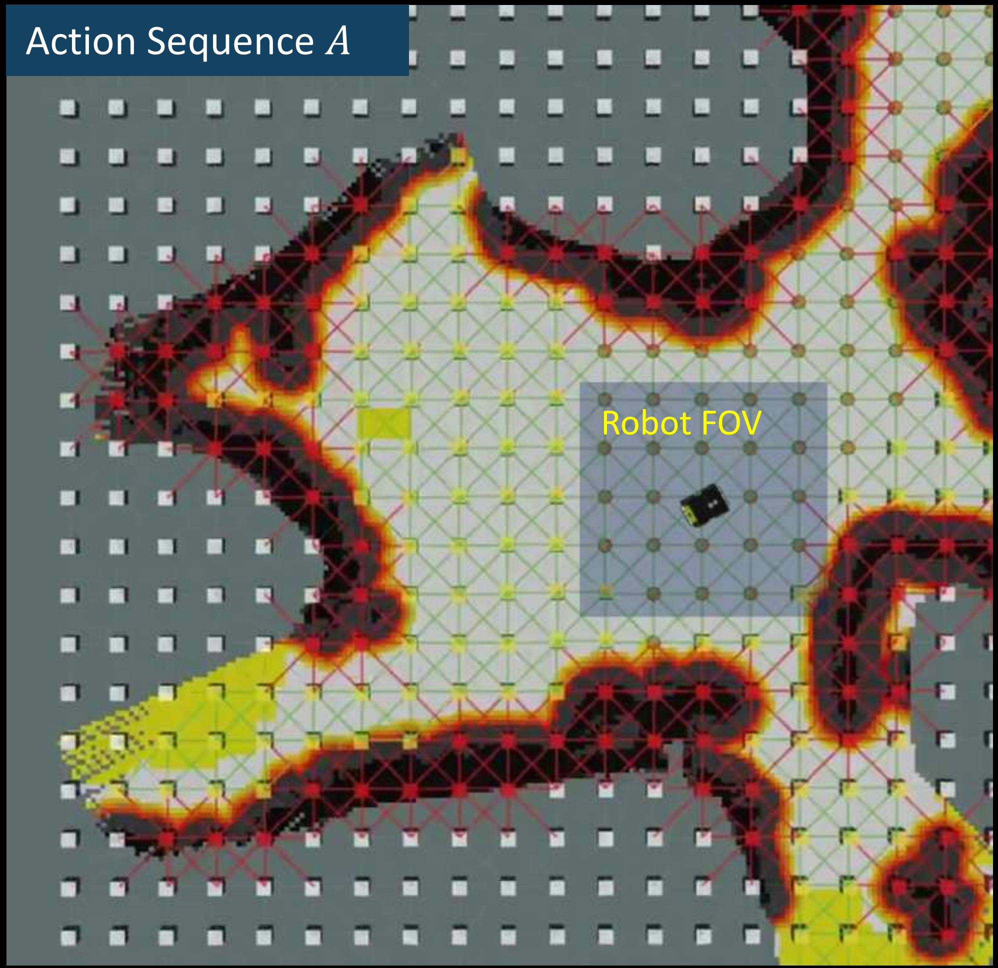


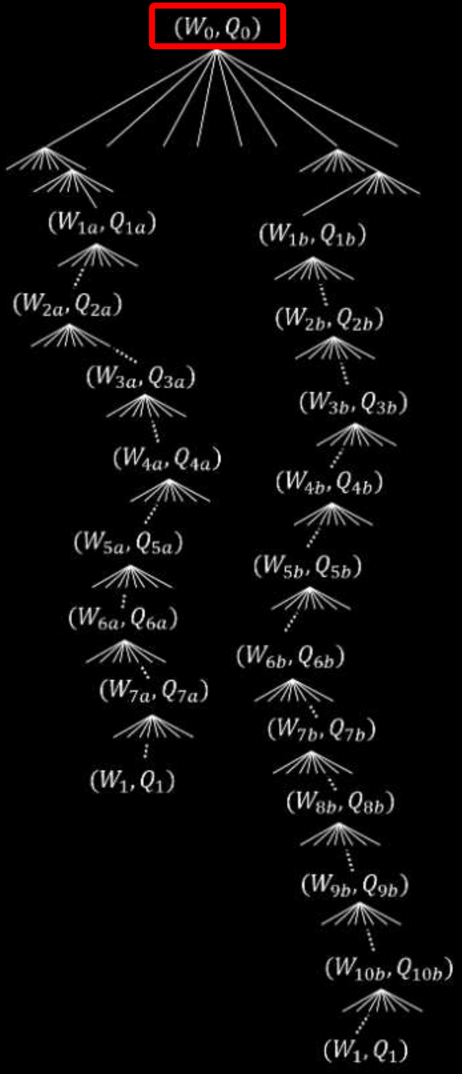


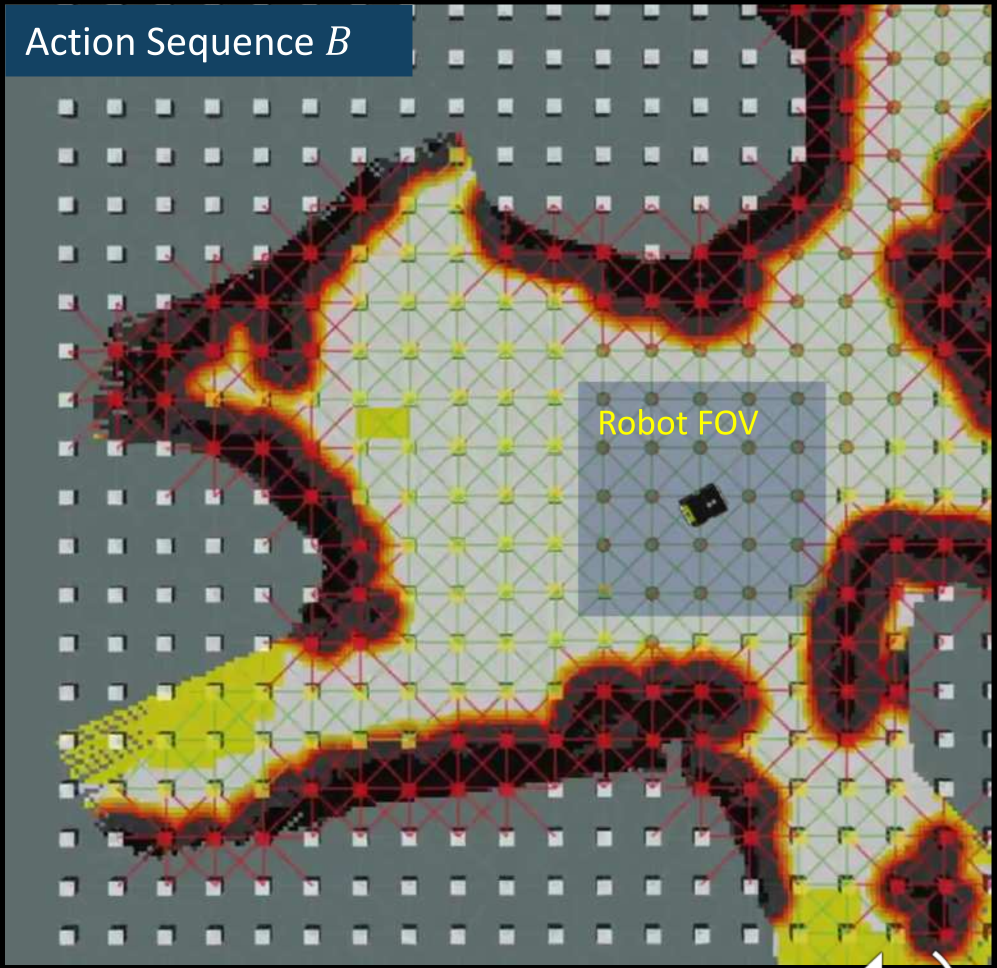

# Methodology

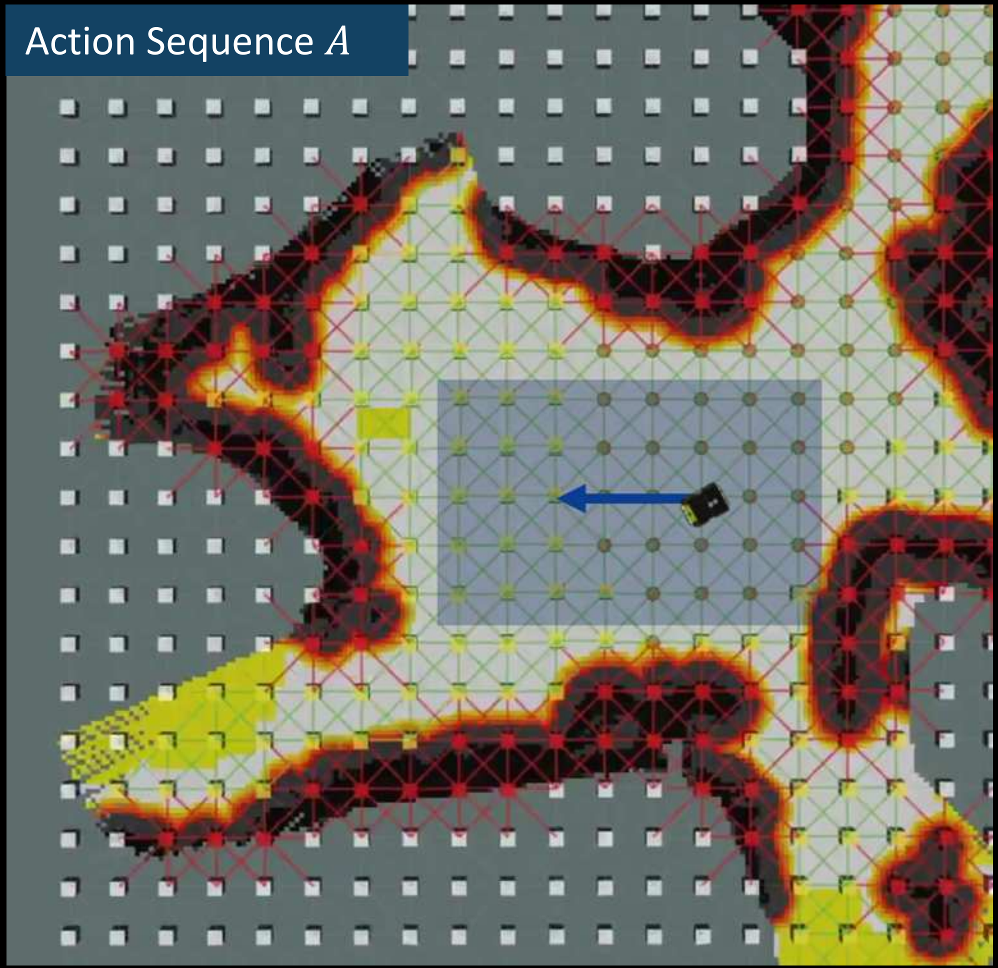


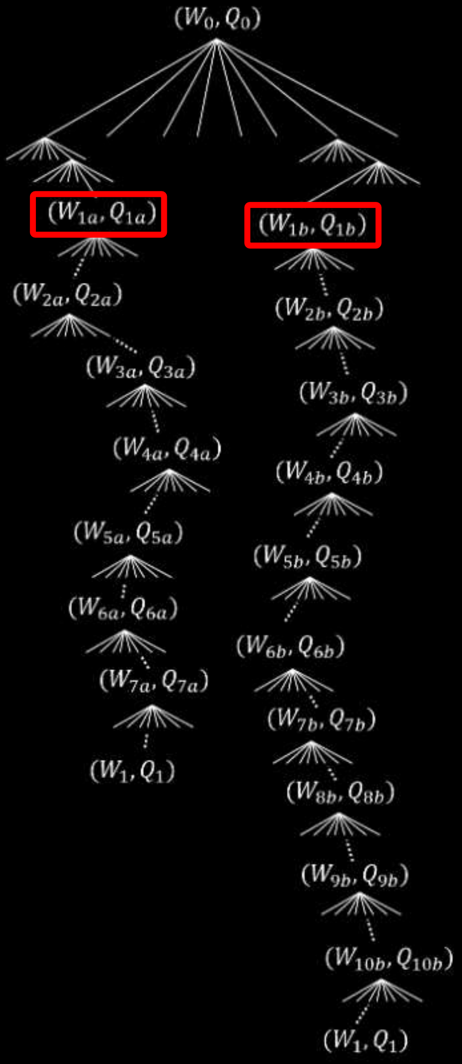


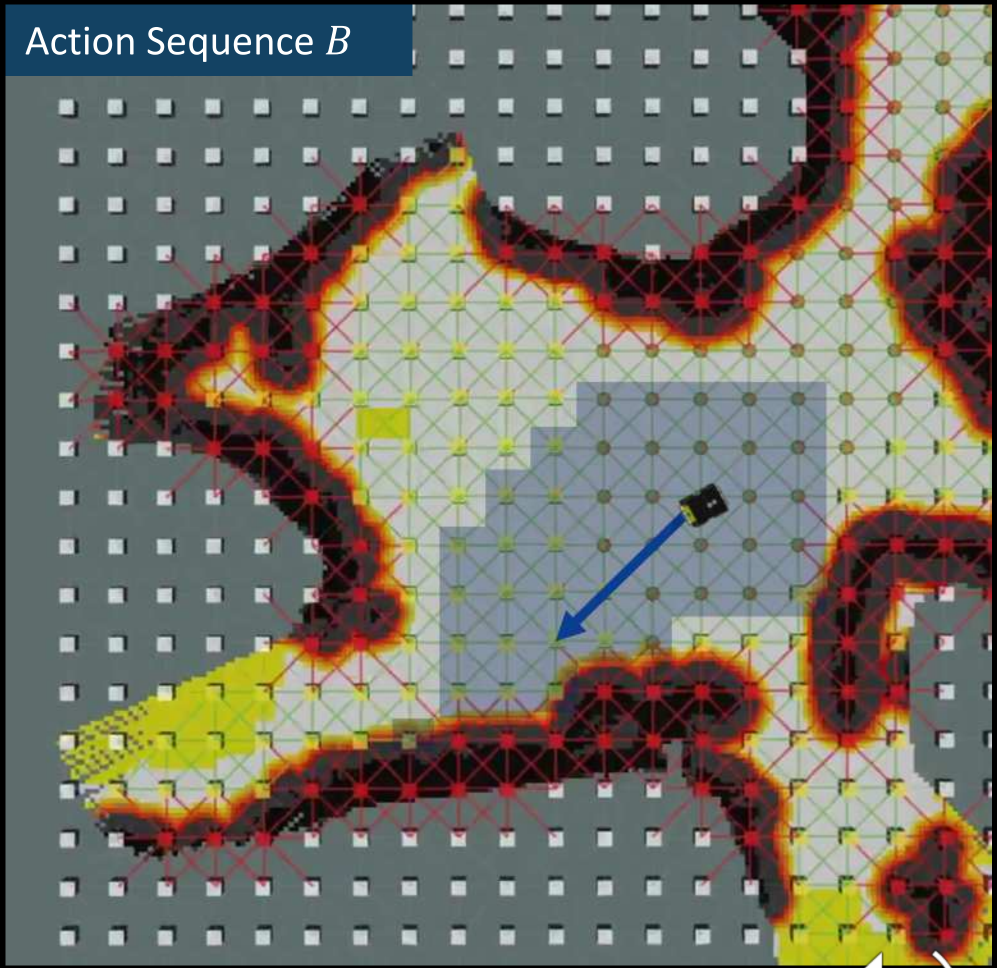

# Methodology

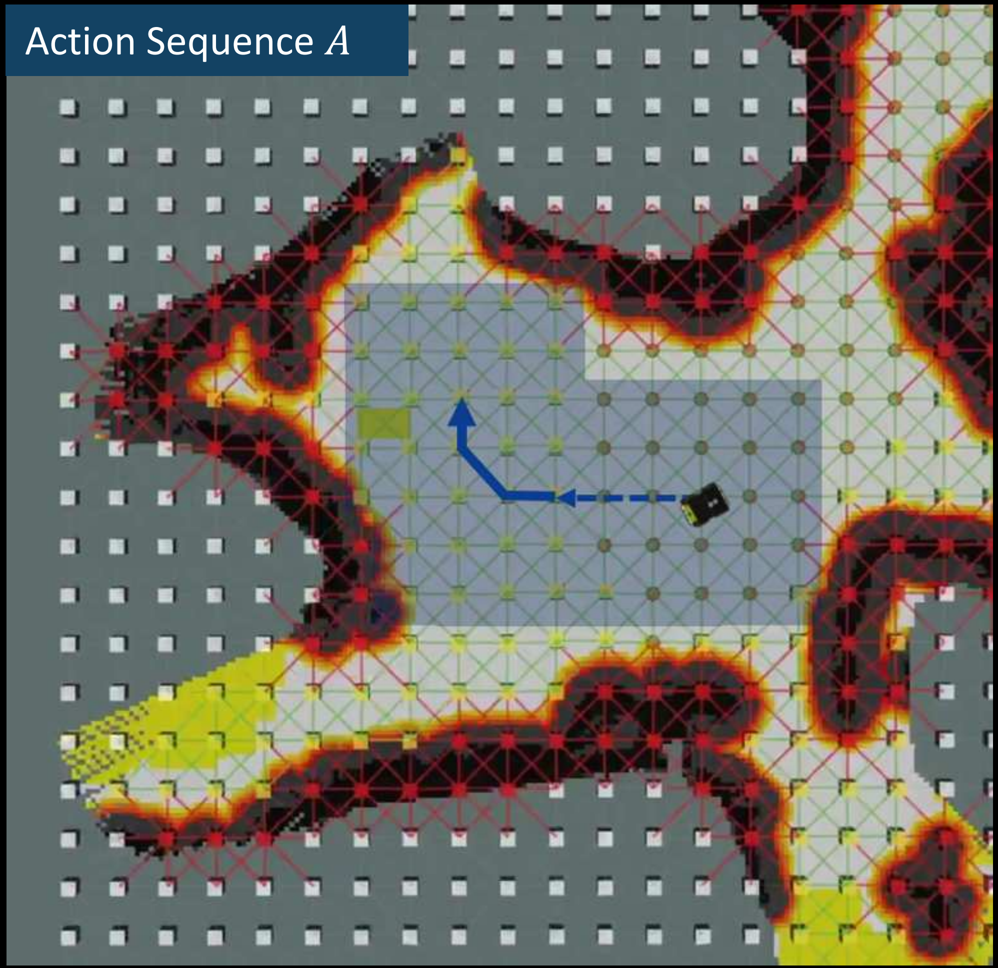


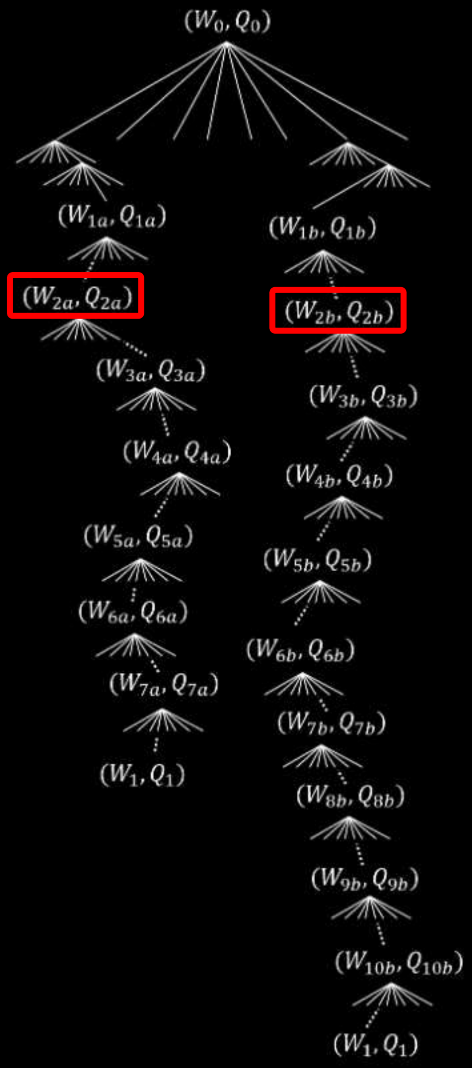


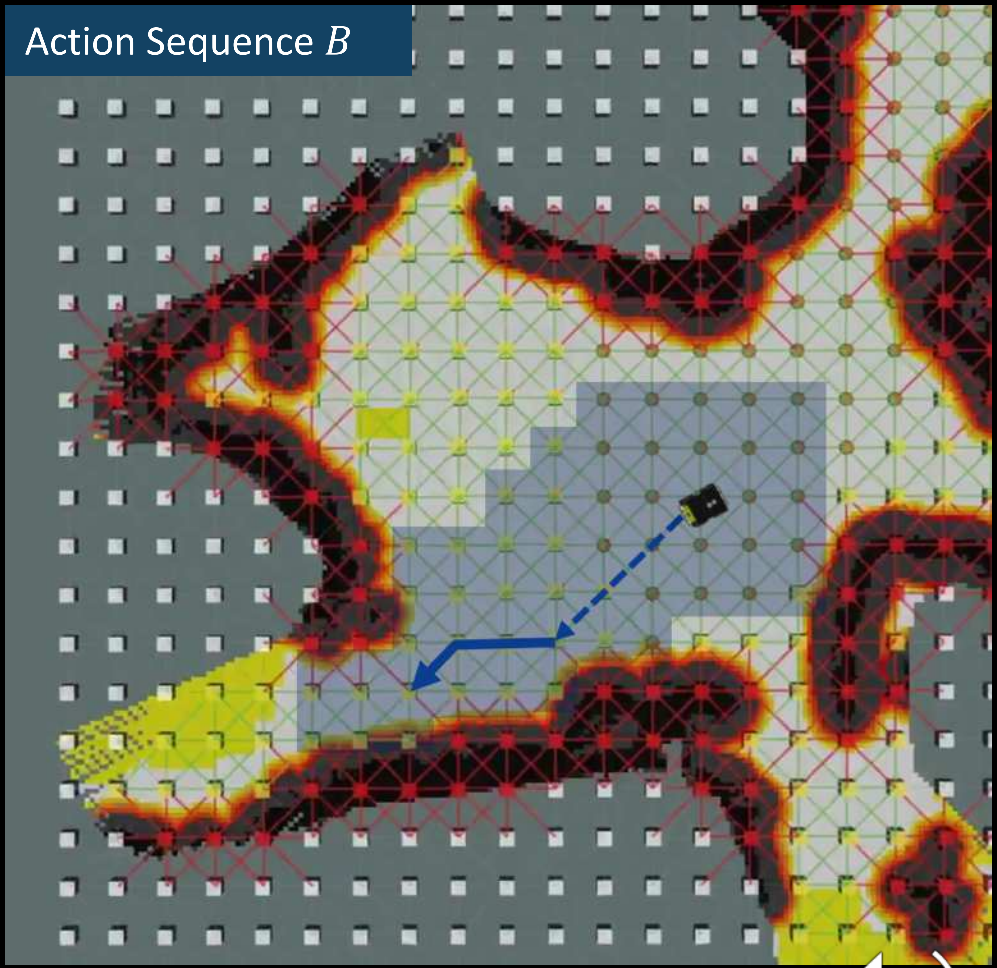

# Methodology

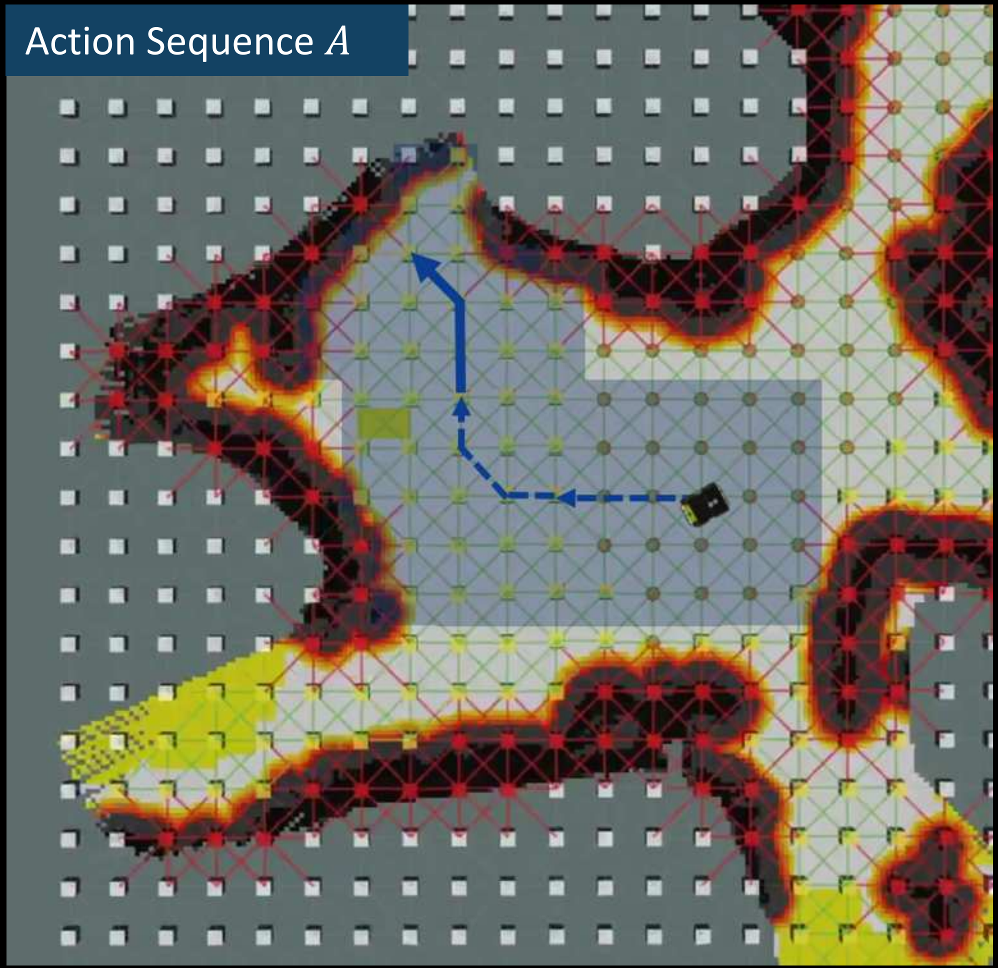


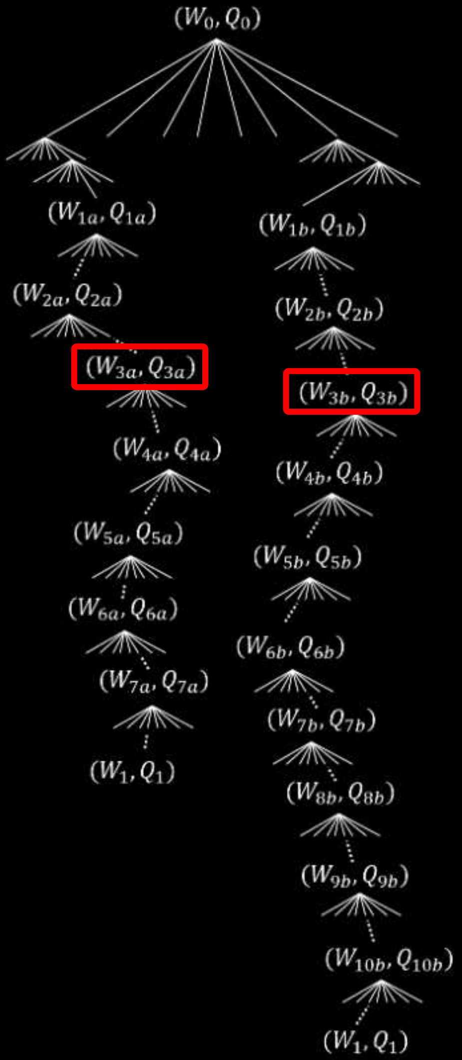


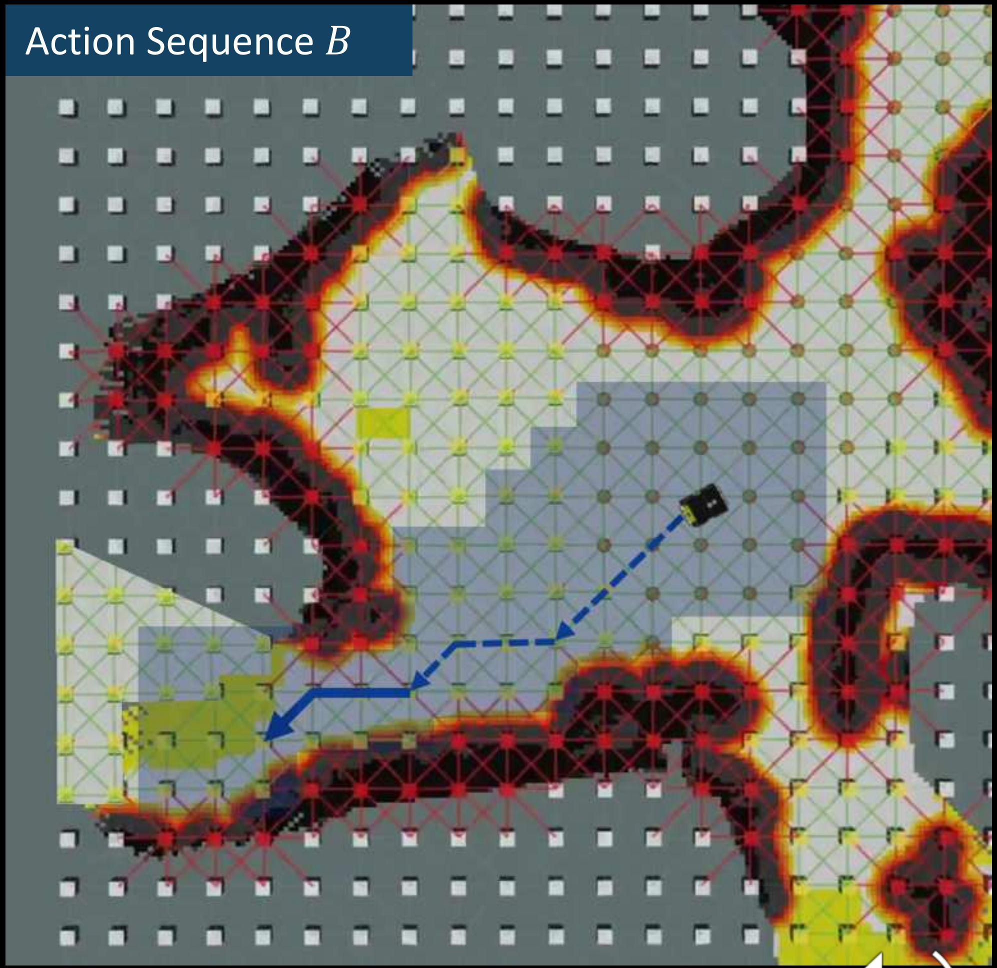

# Methodology

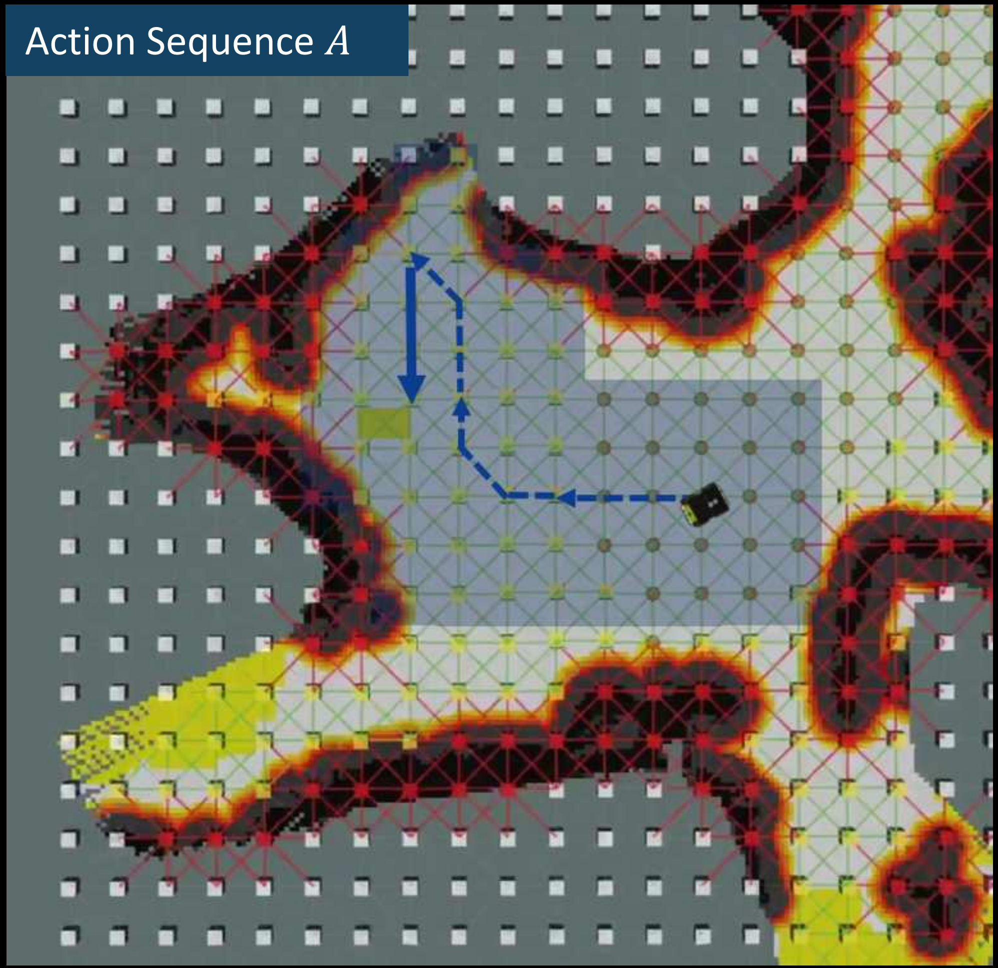


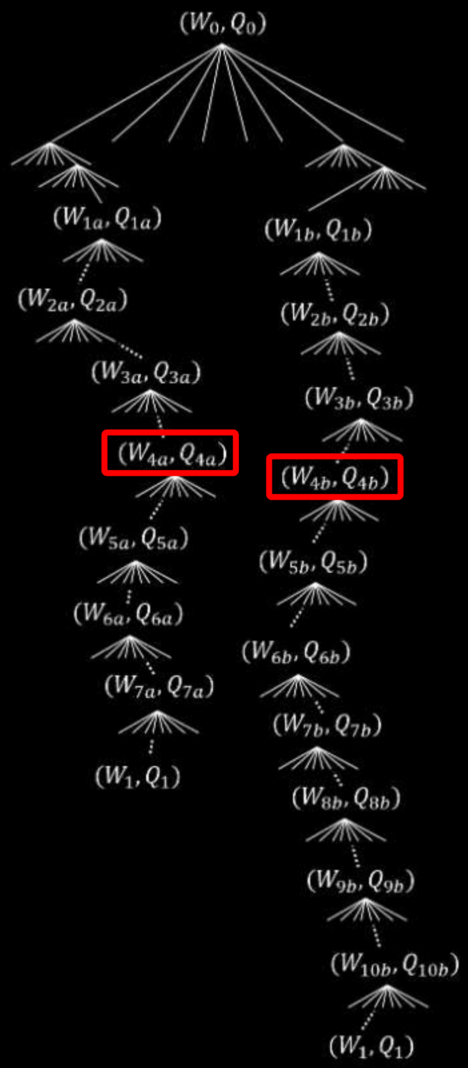


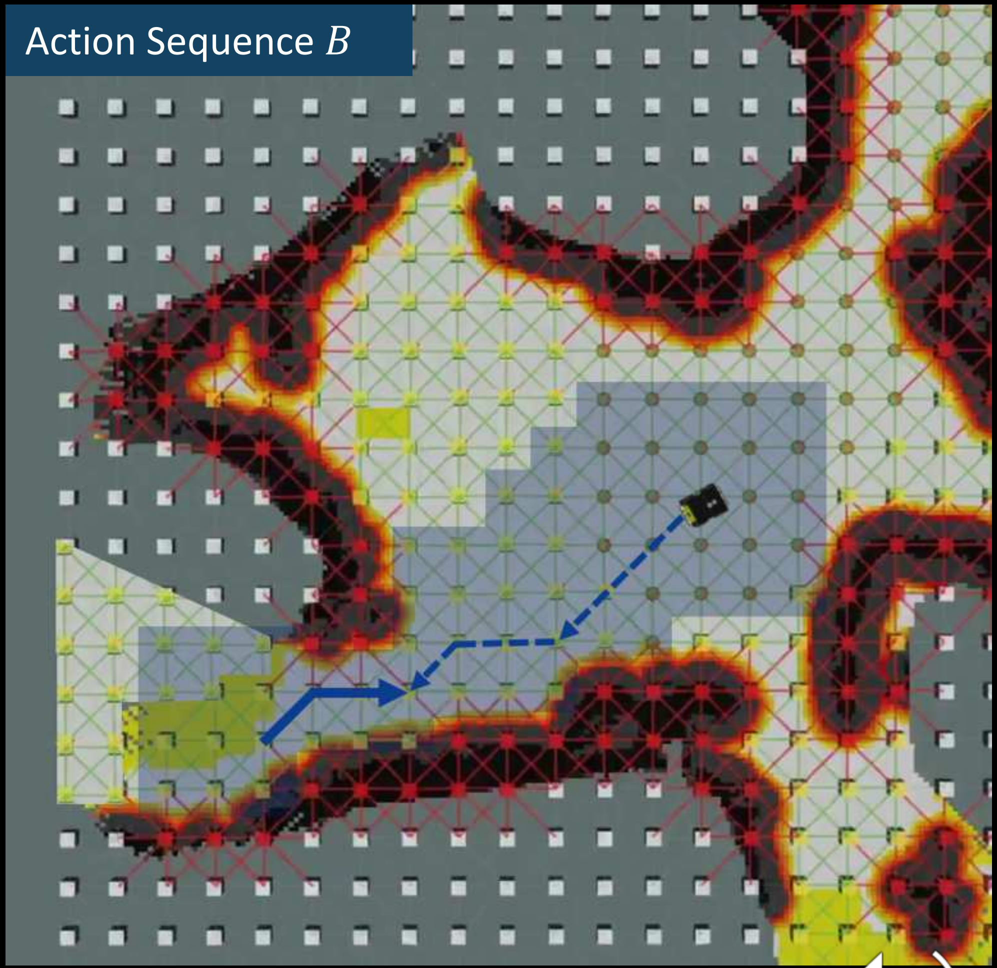

# Methodology

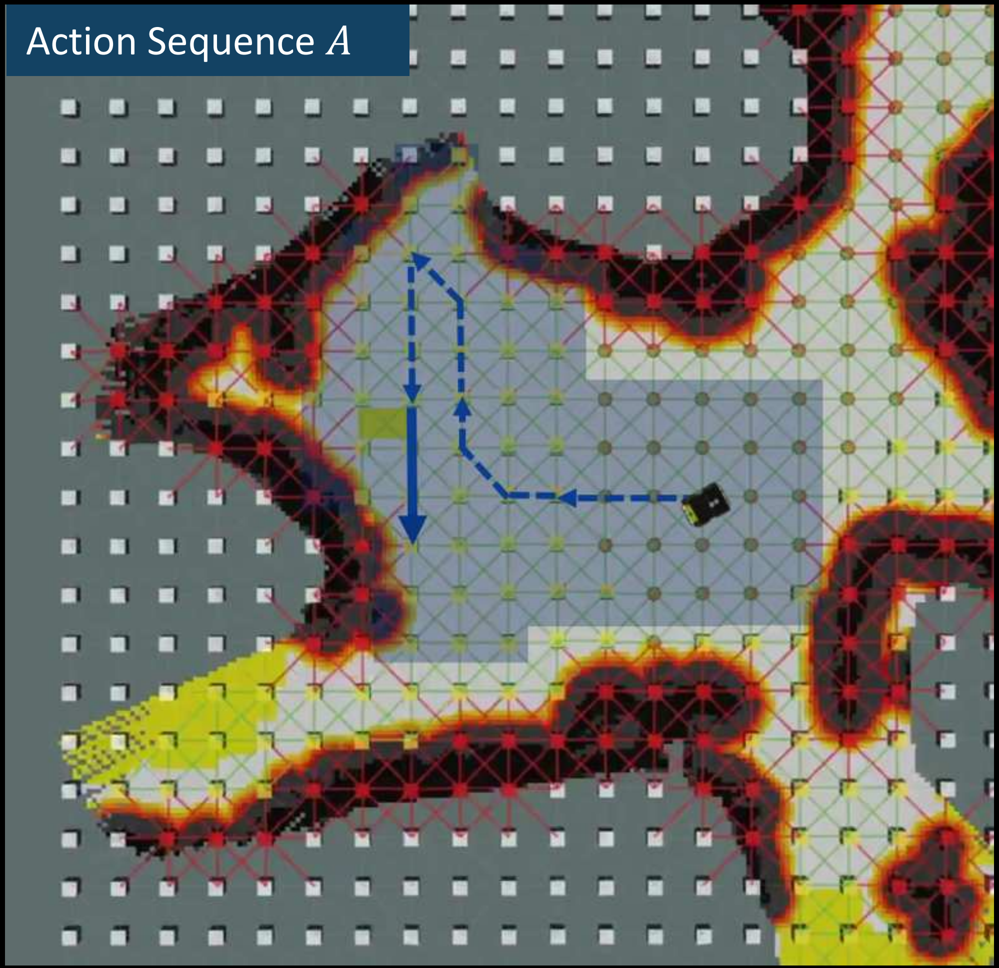


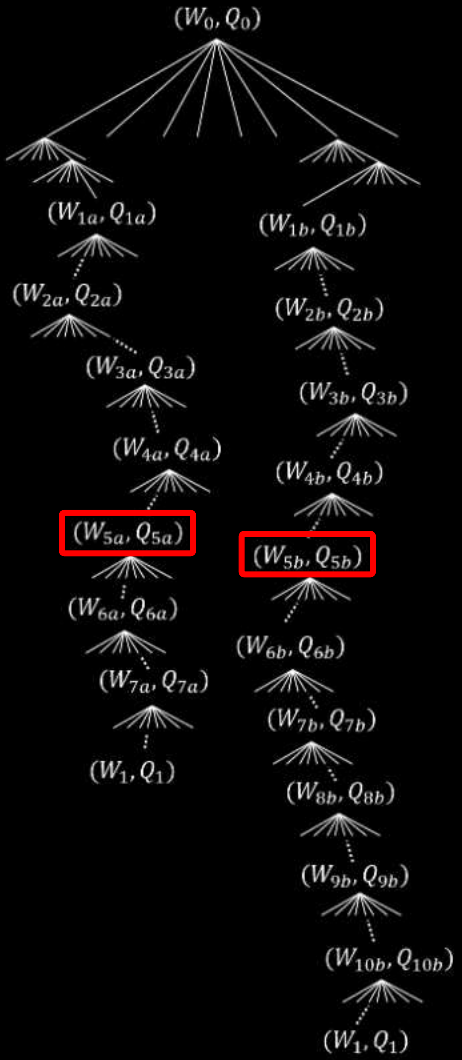


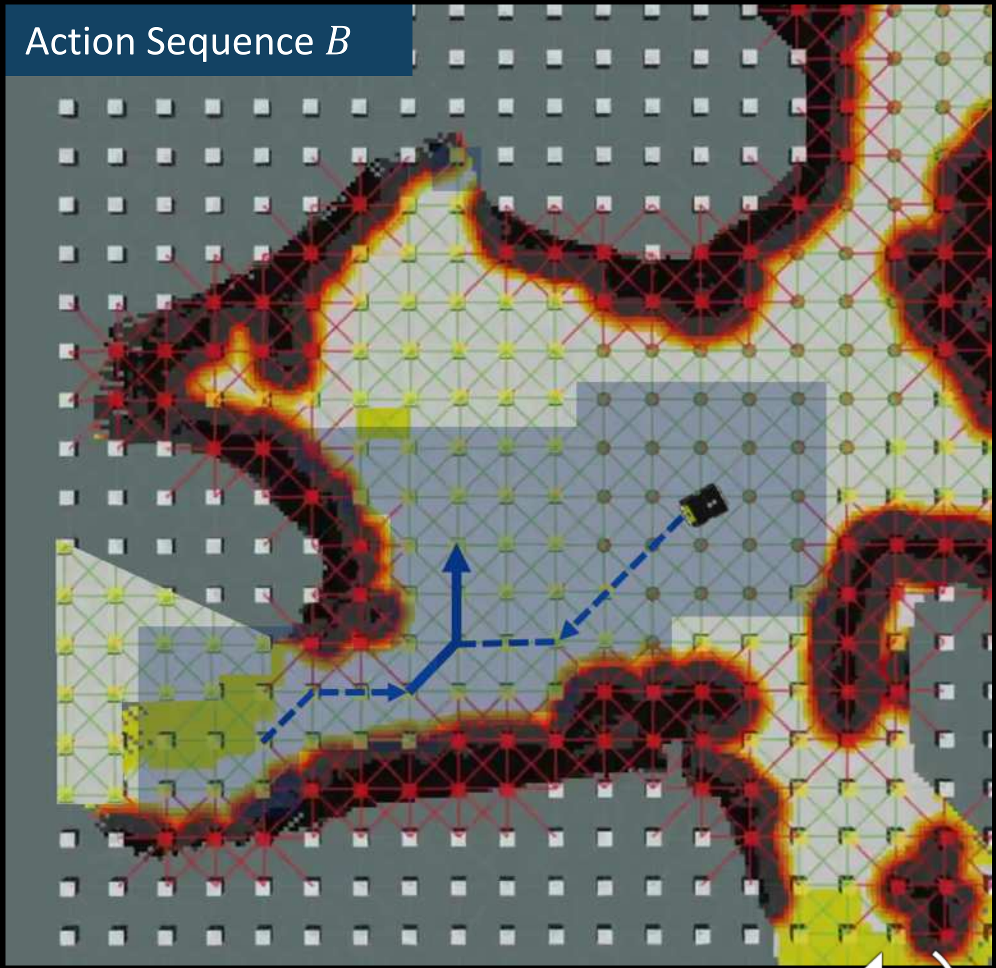

# Methodology

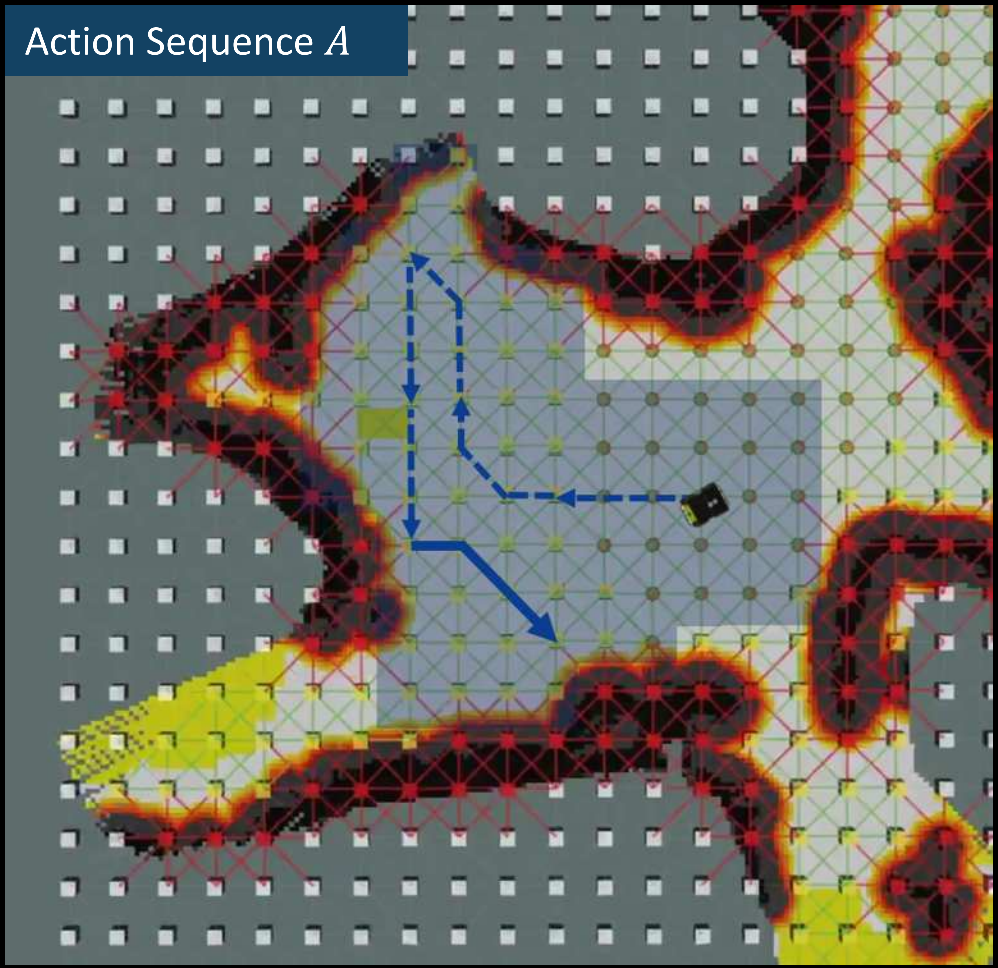


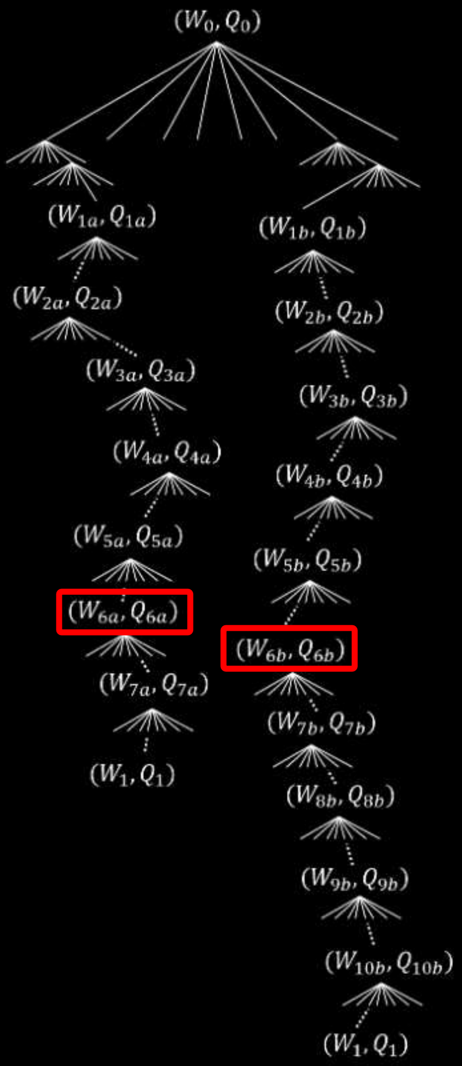


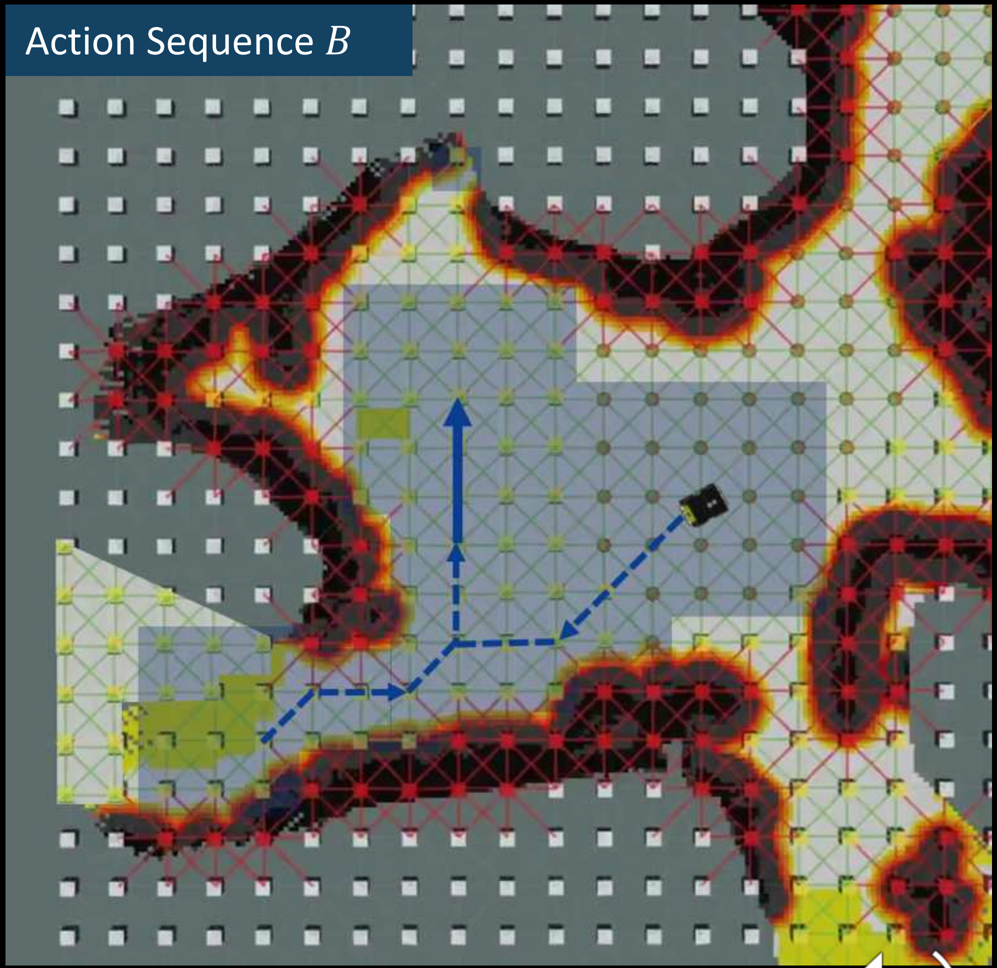

# Methodology

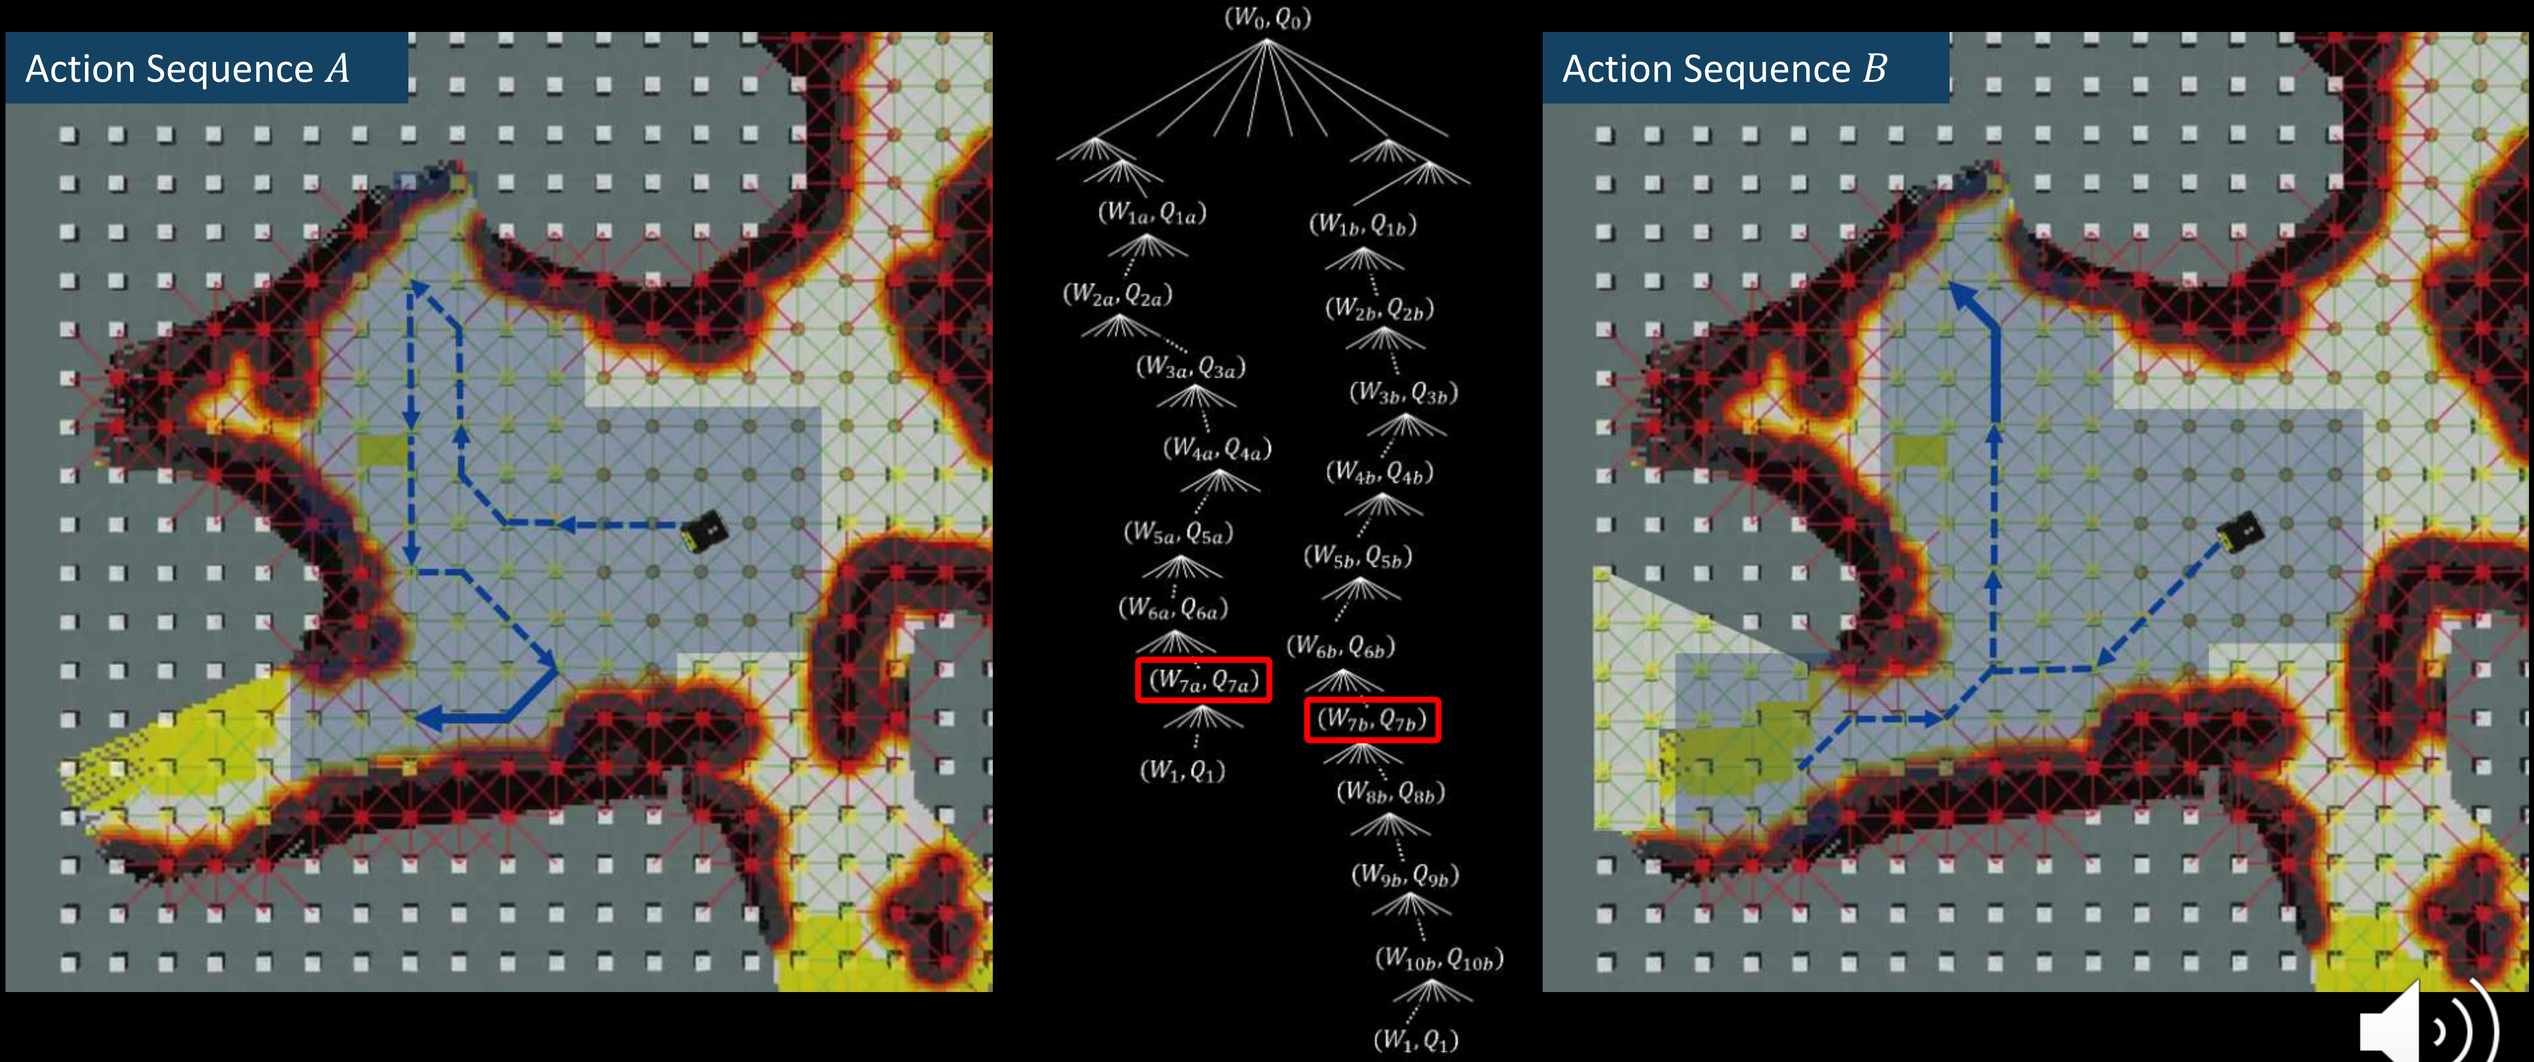

# Methodology

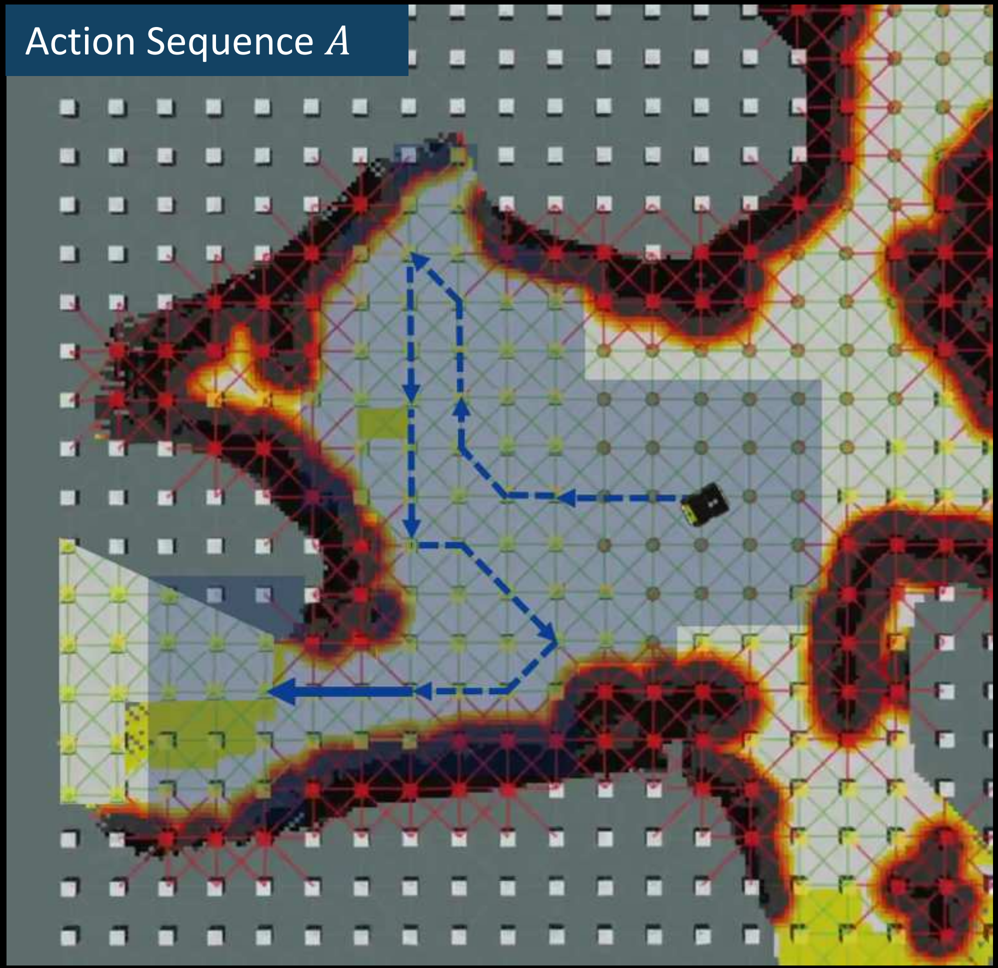


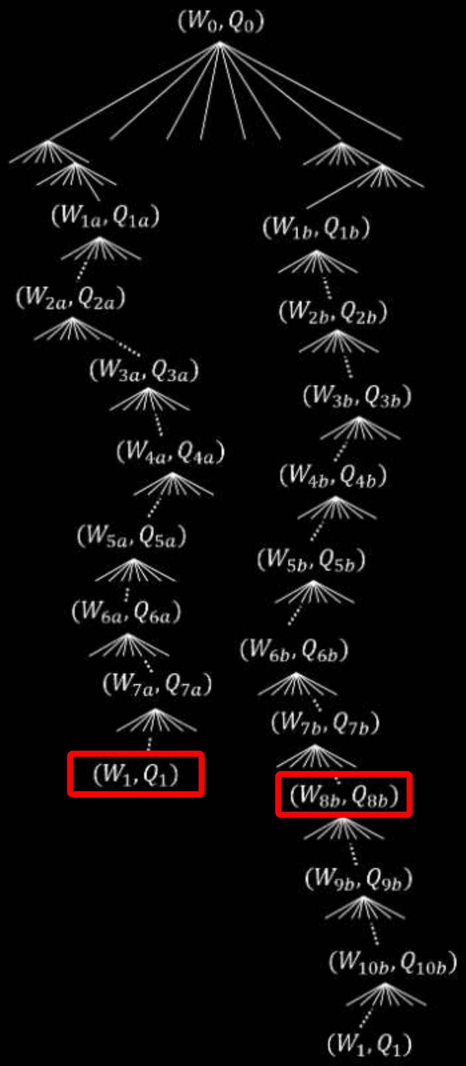


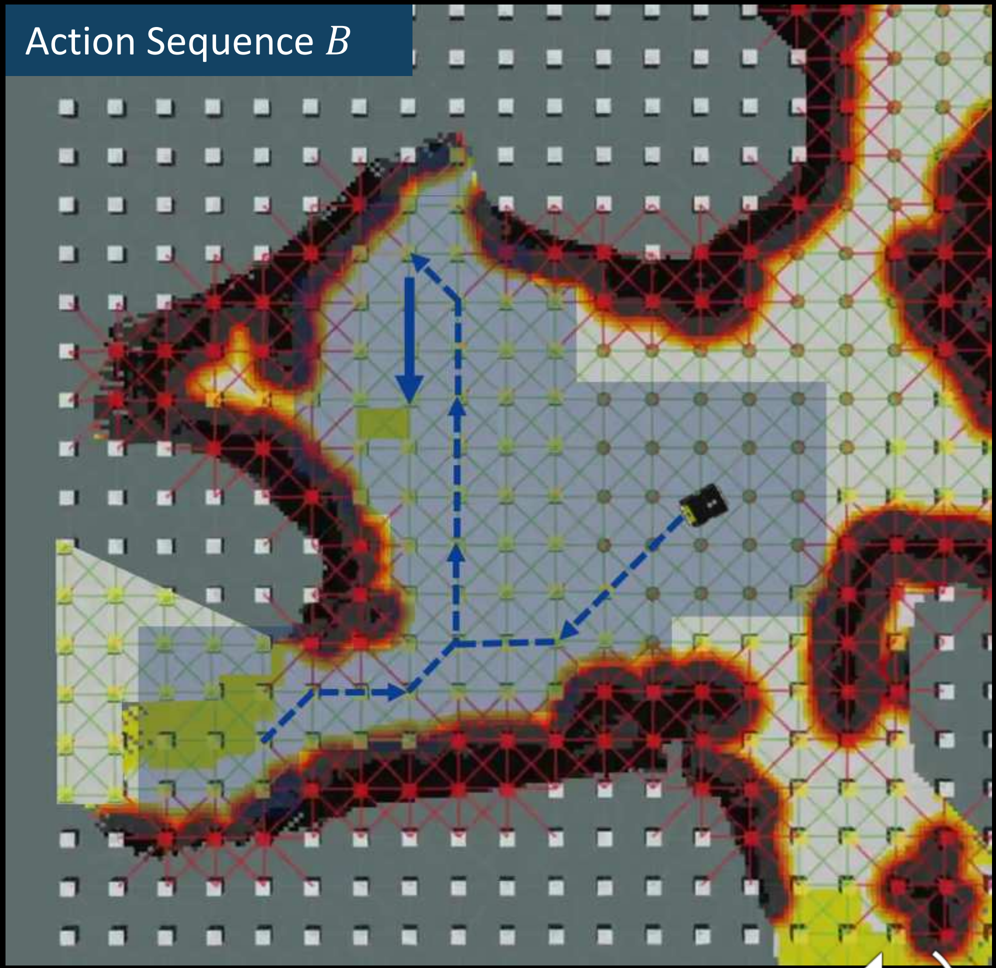

# Methodology

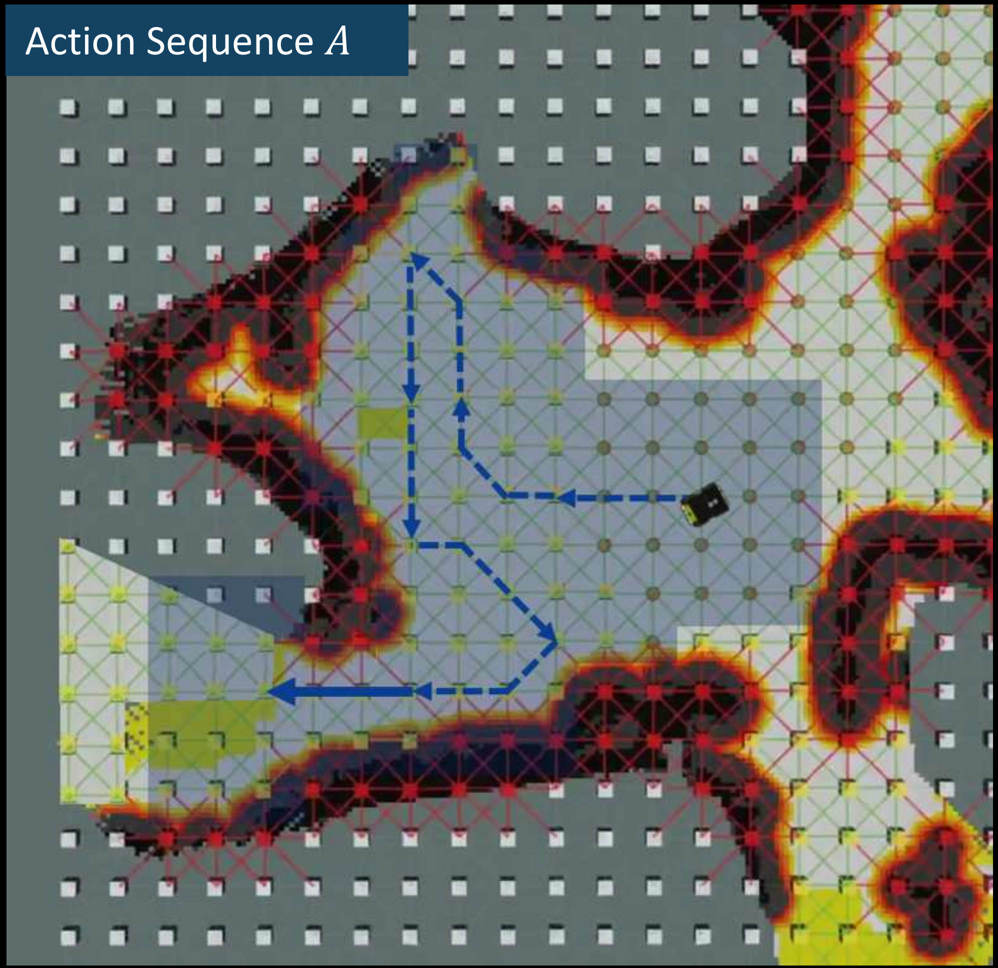


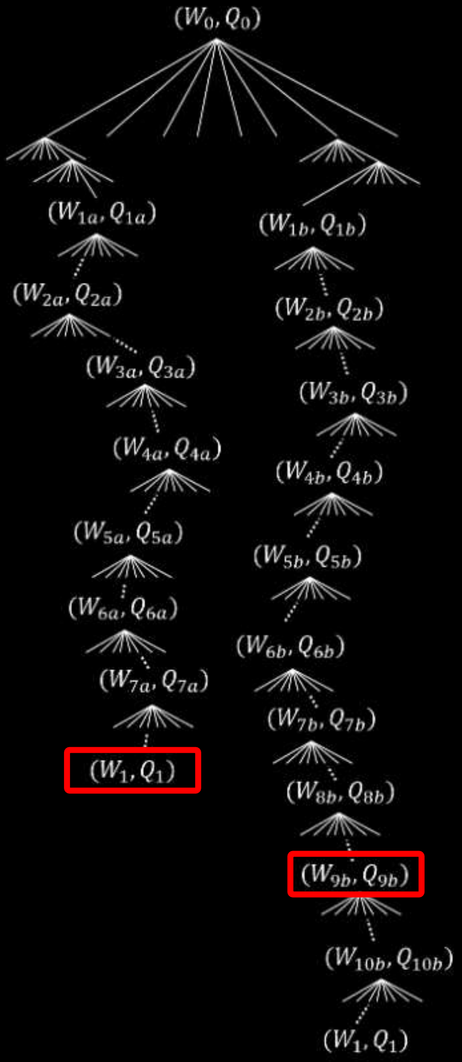


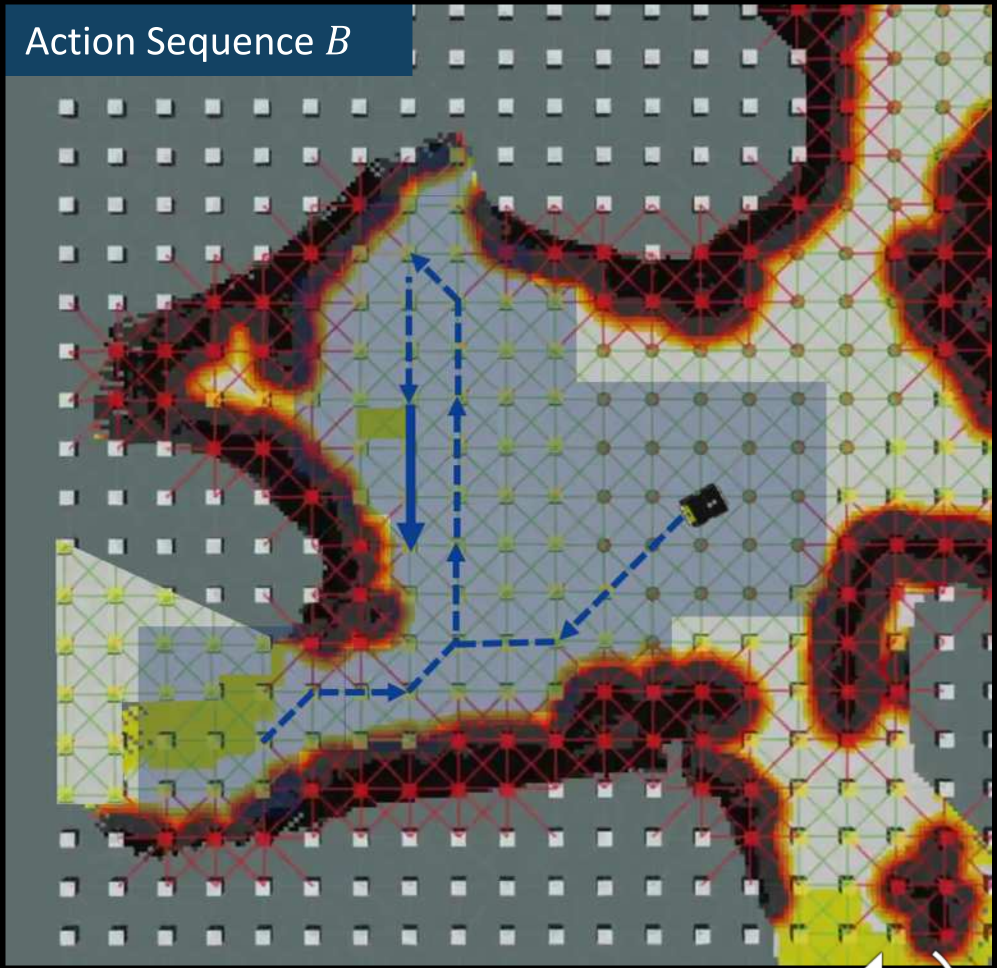

# Methodology

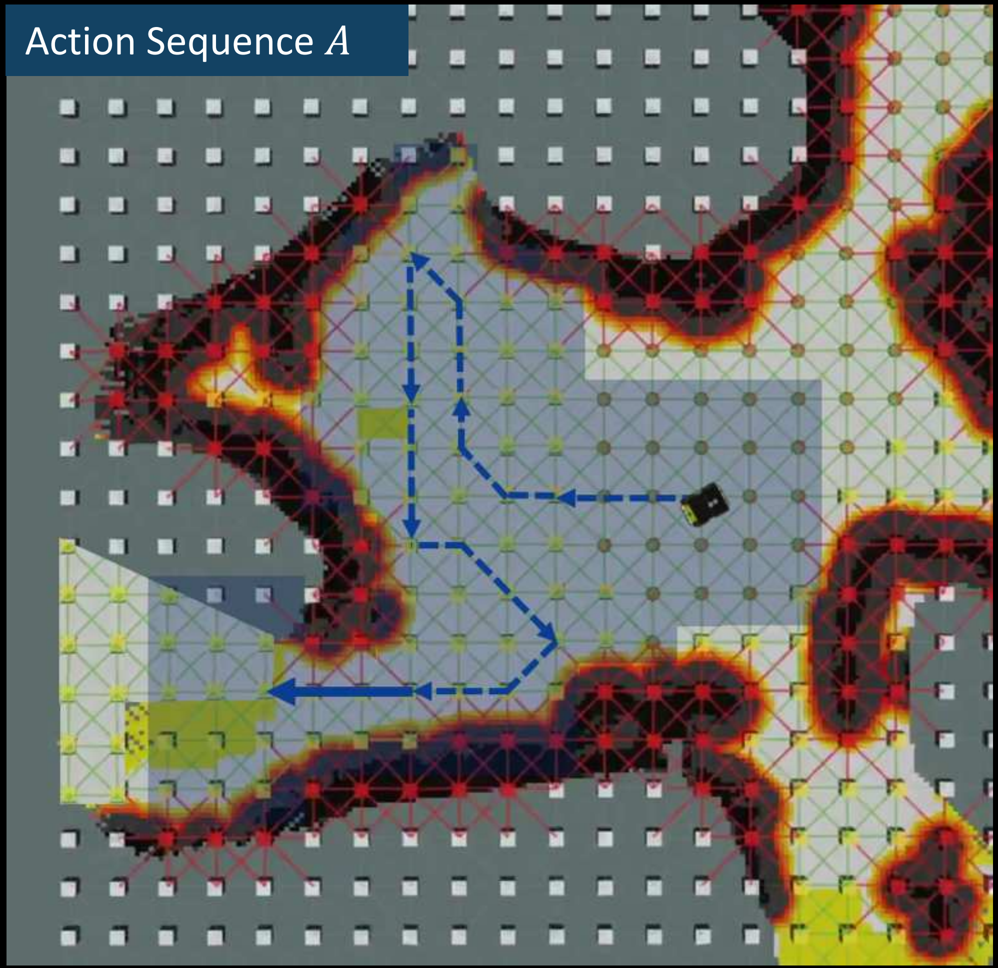


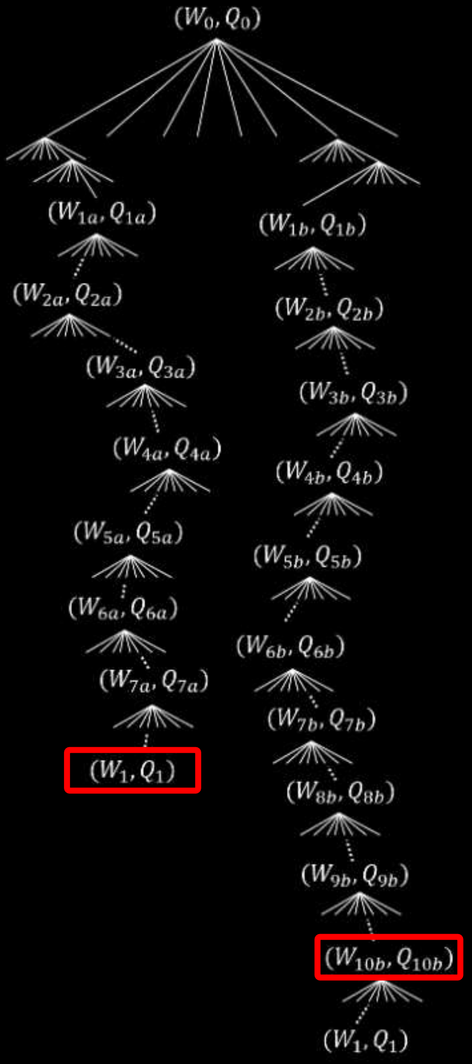


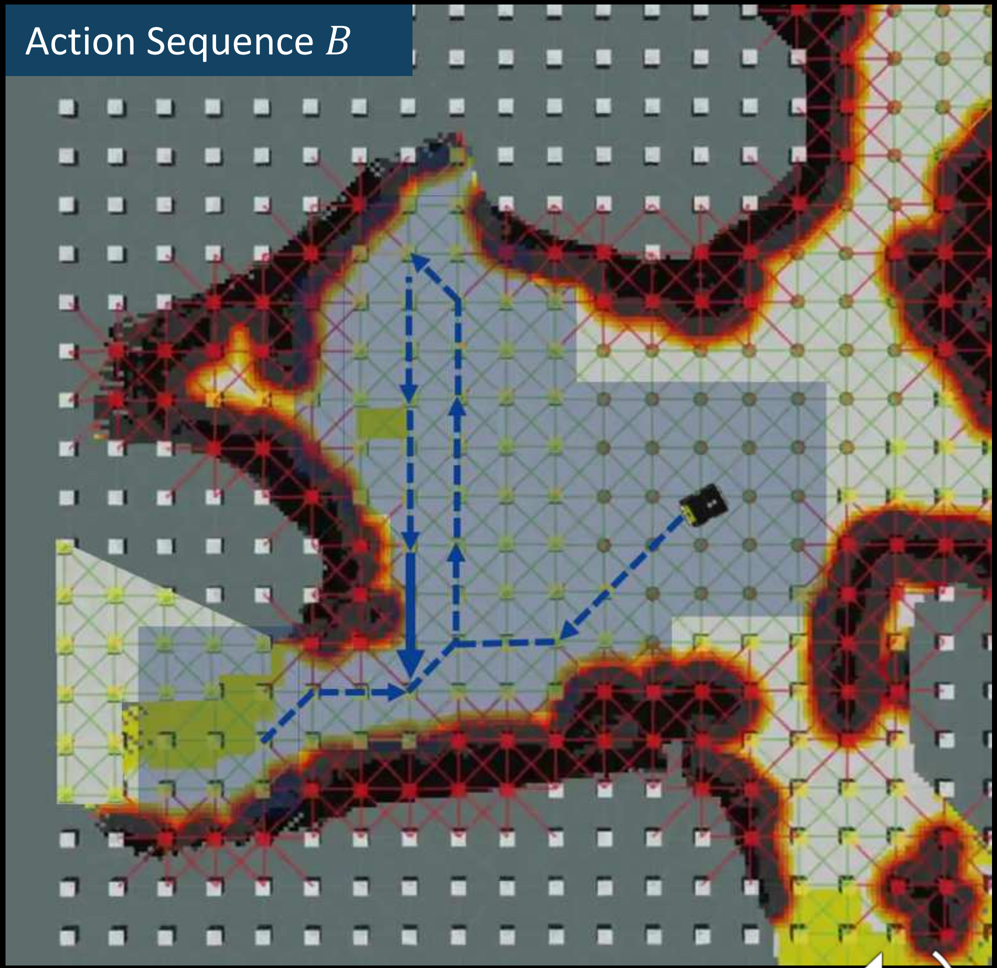

# Methodology

Action Sequence $A$

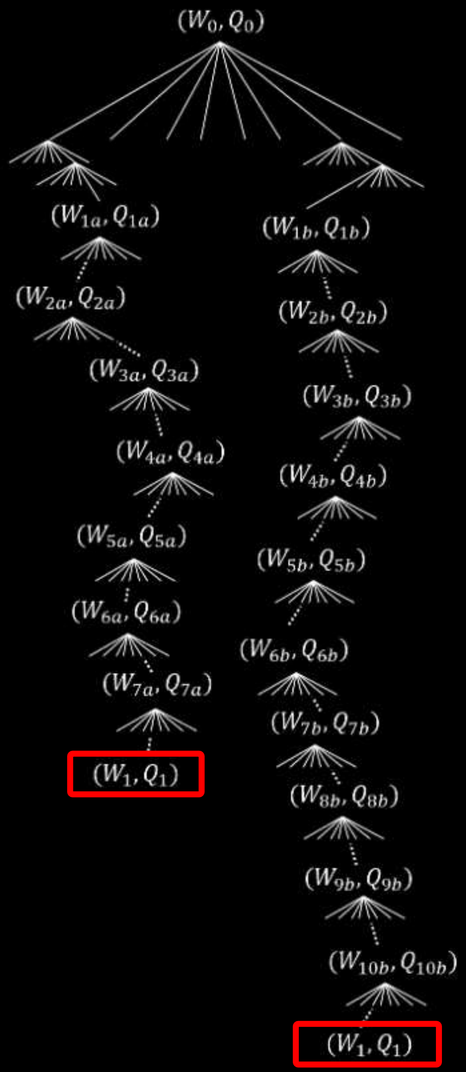


Action Sequence $B$

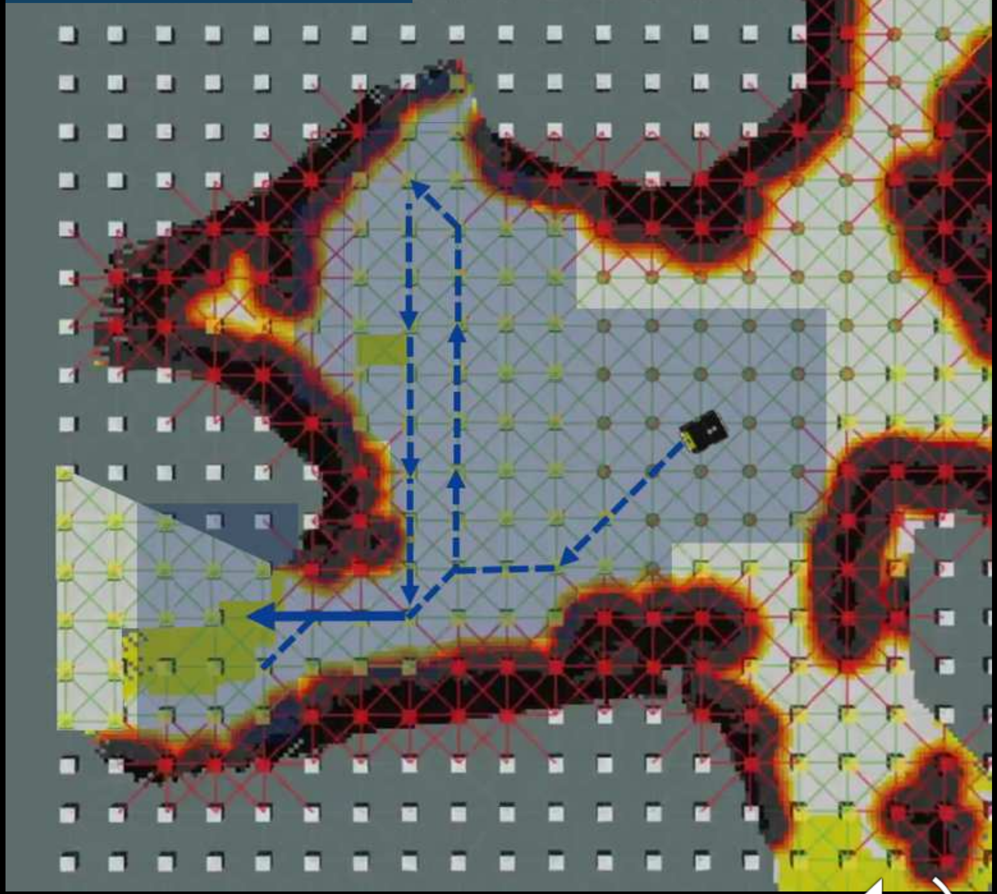

# Methodology

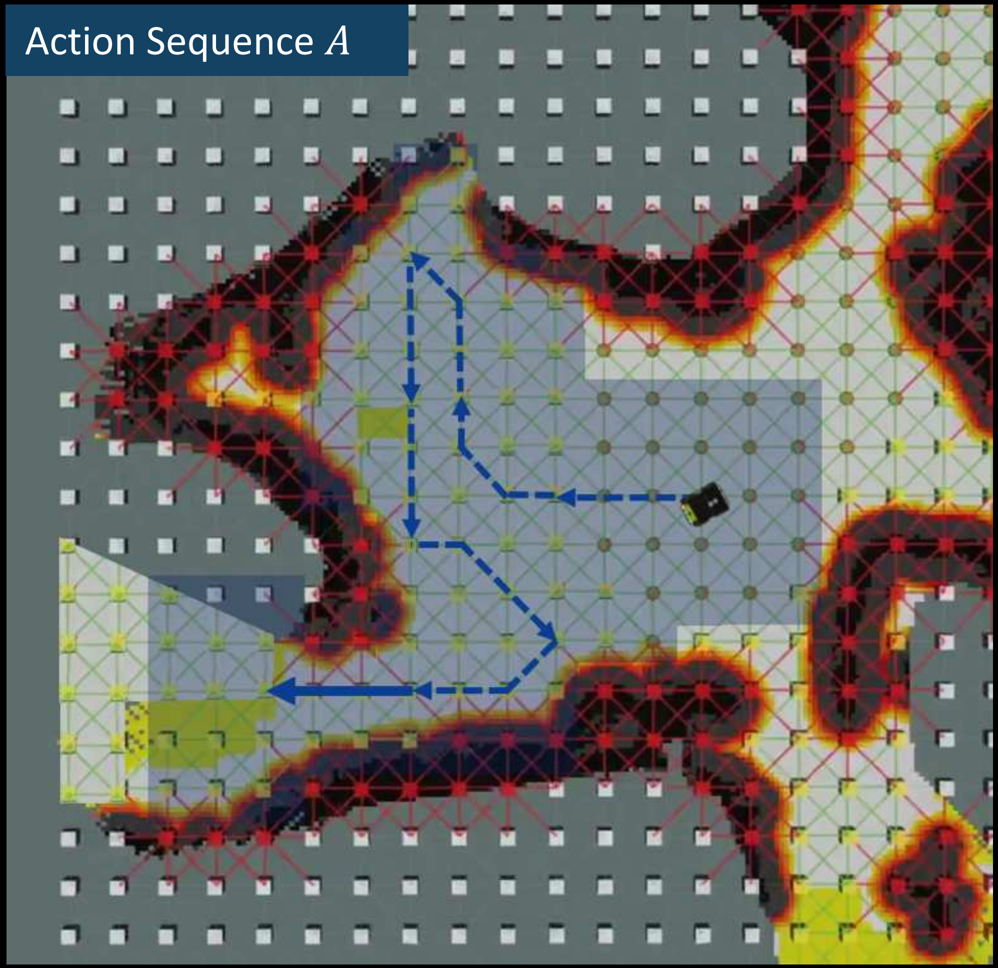


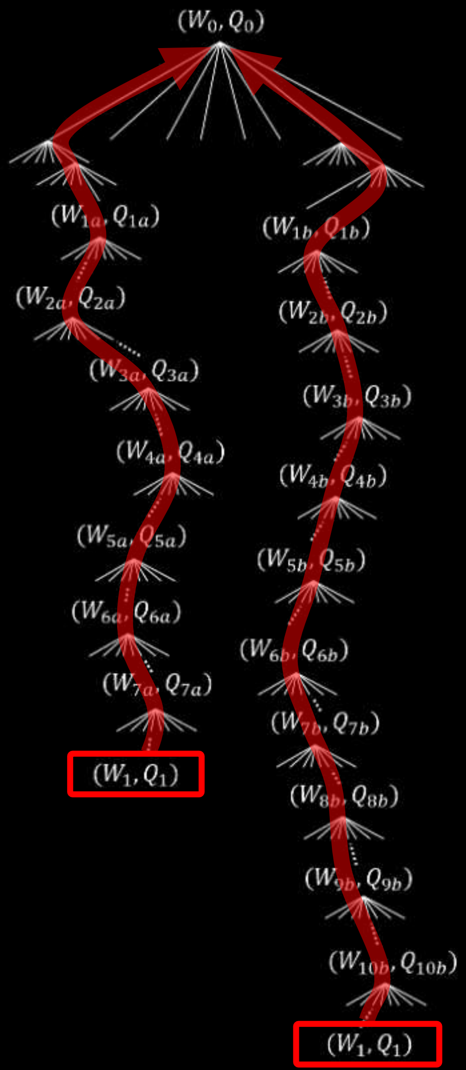


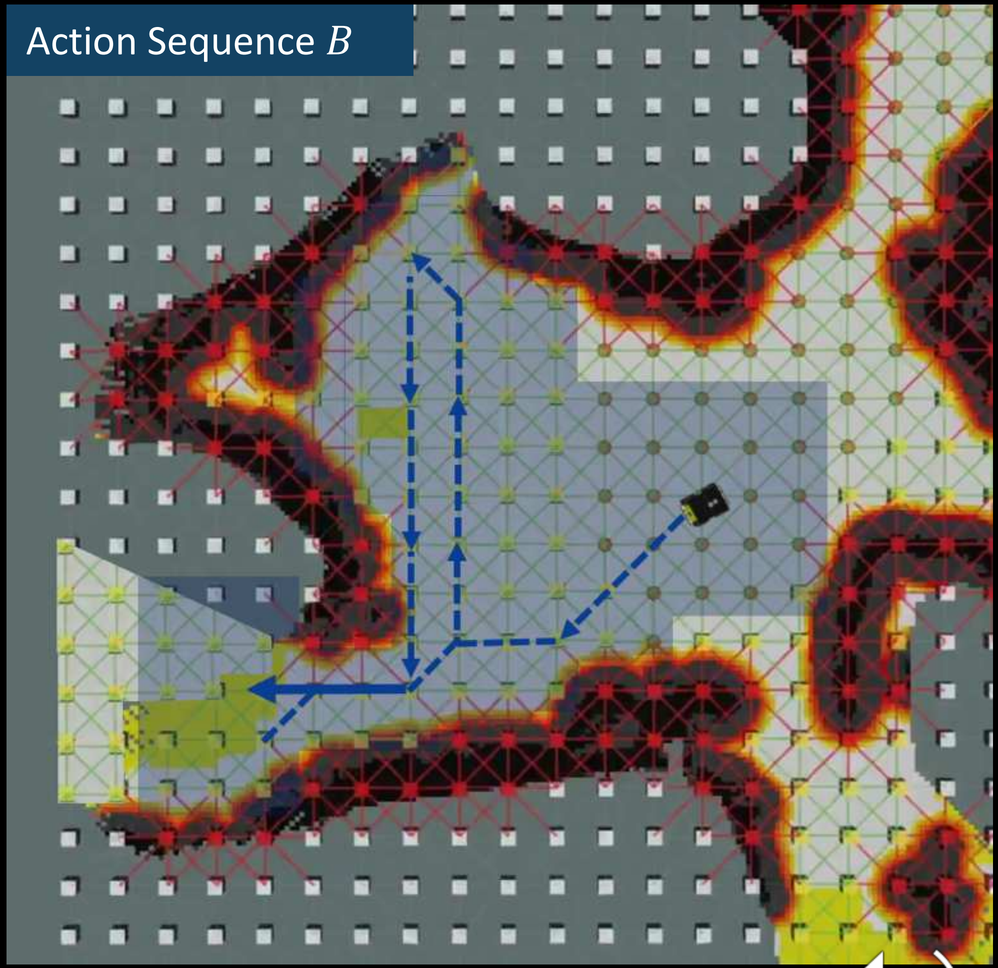

# Methodology

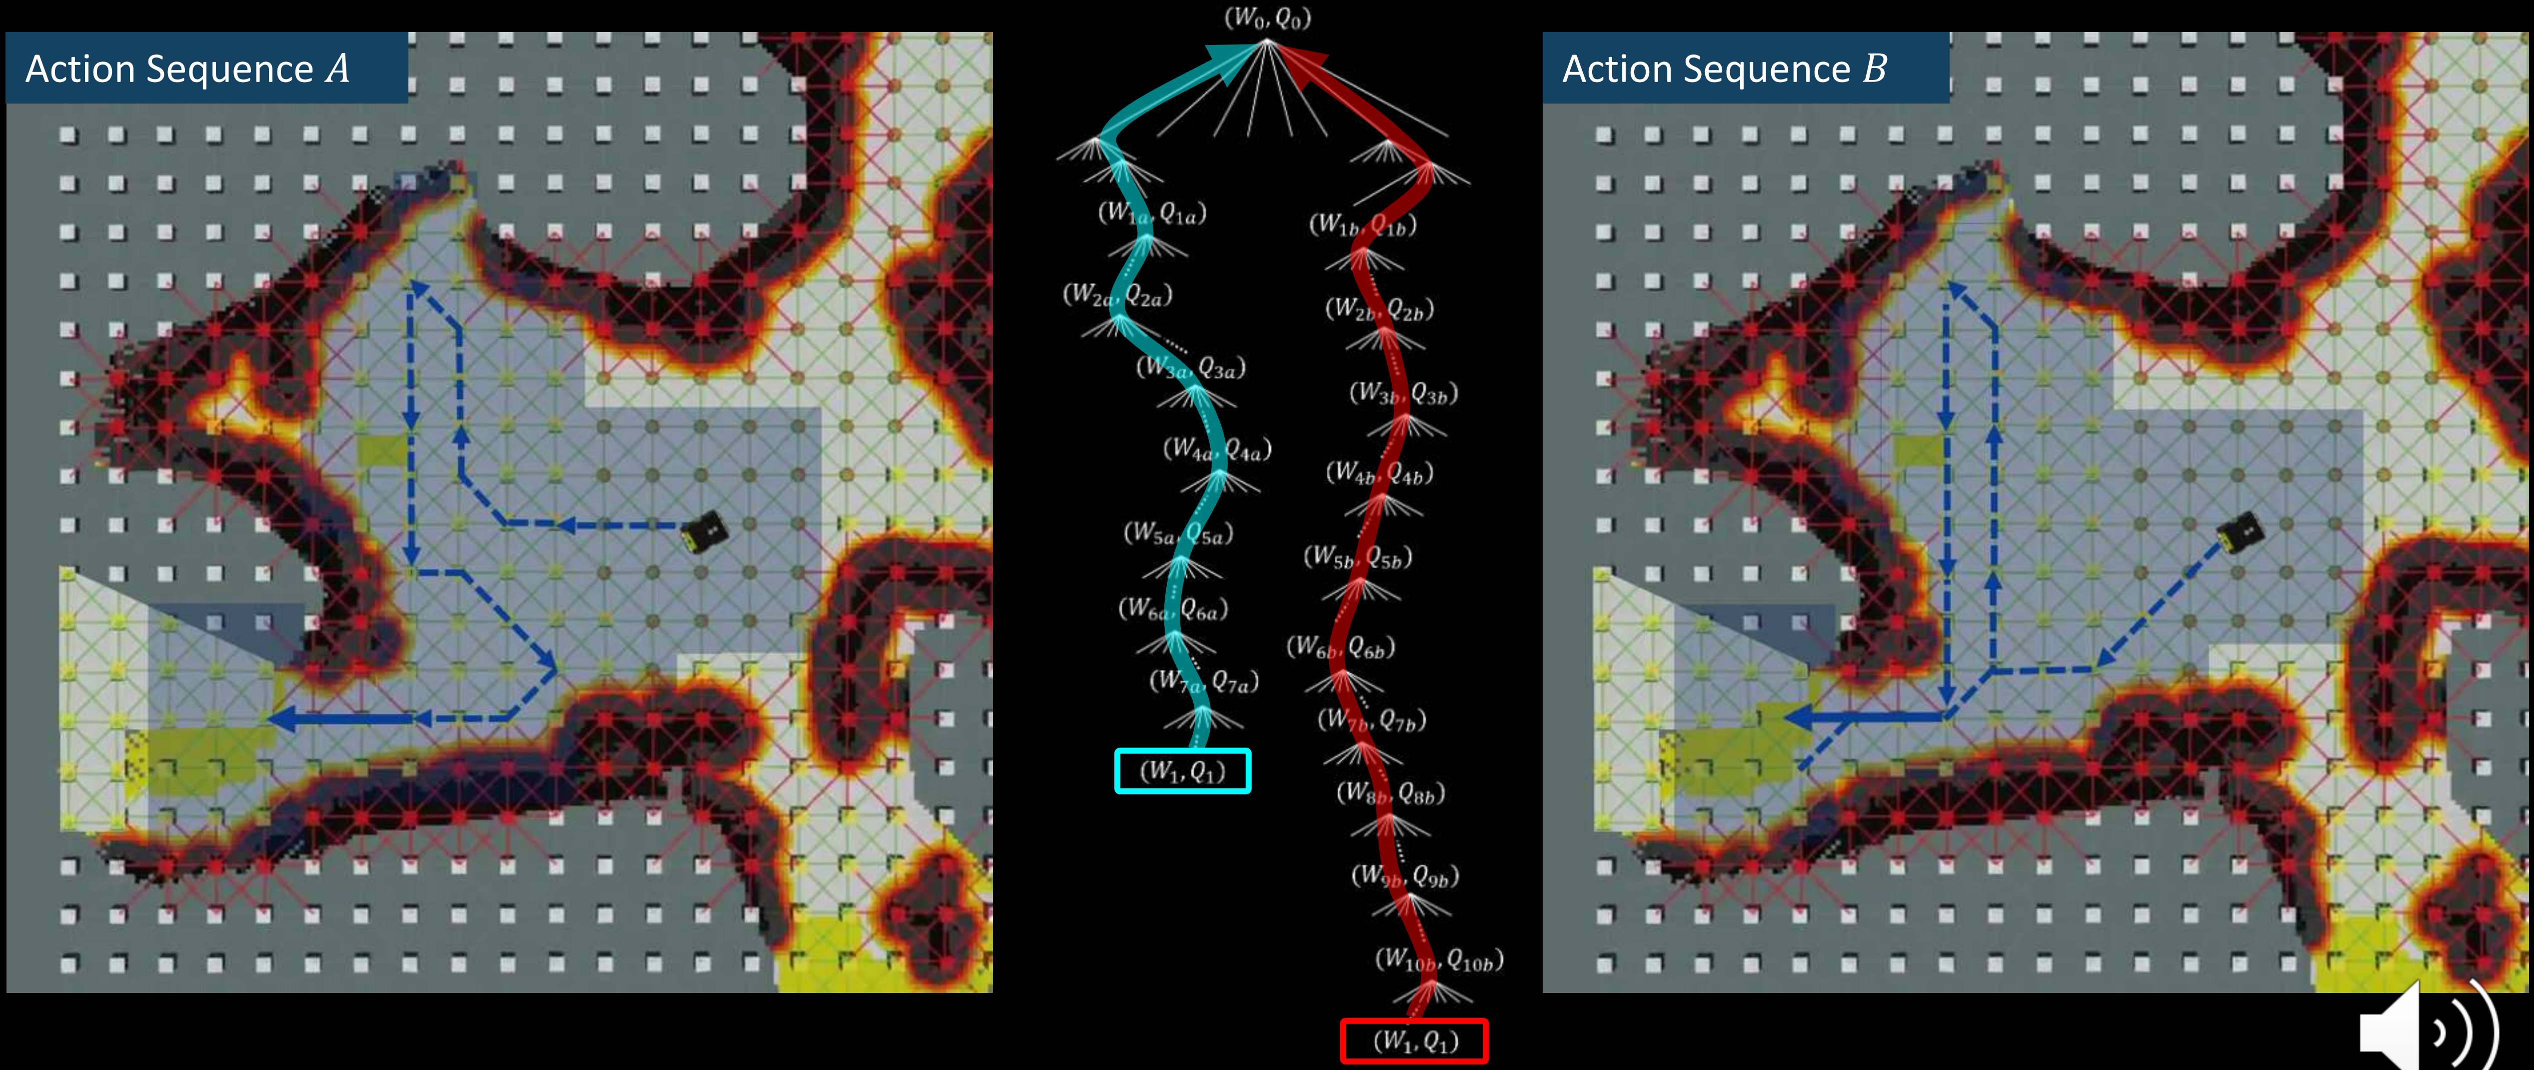

# Results

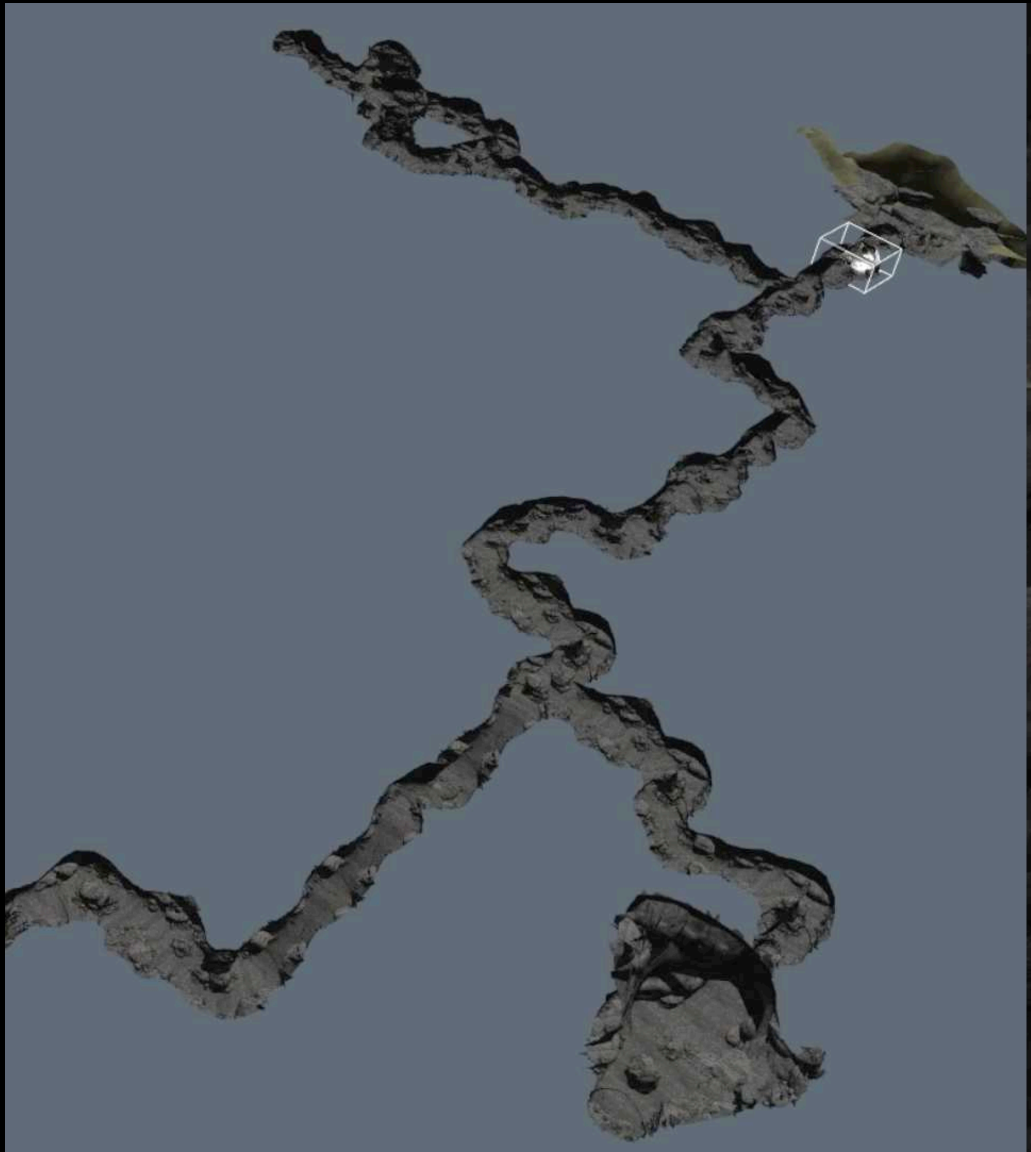

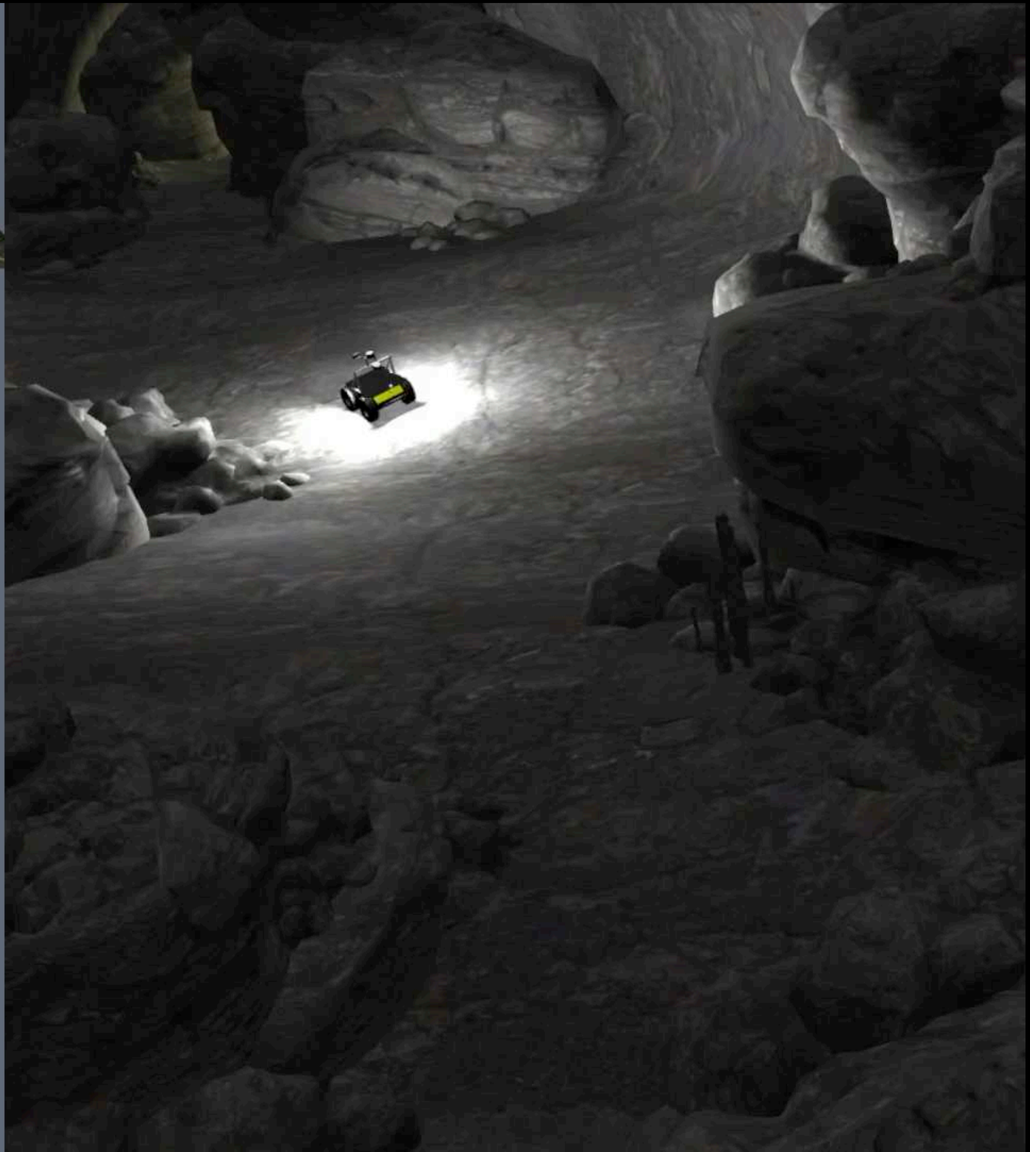

# Results


Onboard View
Sim World
Planning
Belief Manager
Perception
Graph Level Planner
Graph IRM
Graph Manager
Frontier Manager
Odometry
Waypoint
Lattice Level Planner
Lattice IRM
Lattice Manager
Path Traversal Checker
3D Mapping
Path
Kinodynamic Planner
Cost map
Costmap
Pose Graph
Traversability
Velocity Commands
Robot
Sensor Data
Costmap

# Results

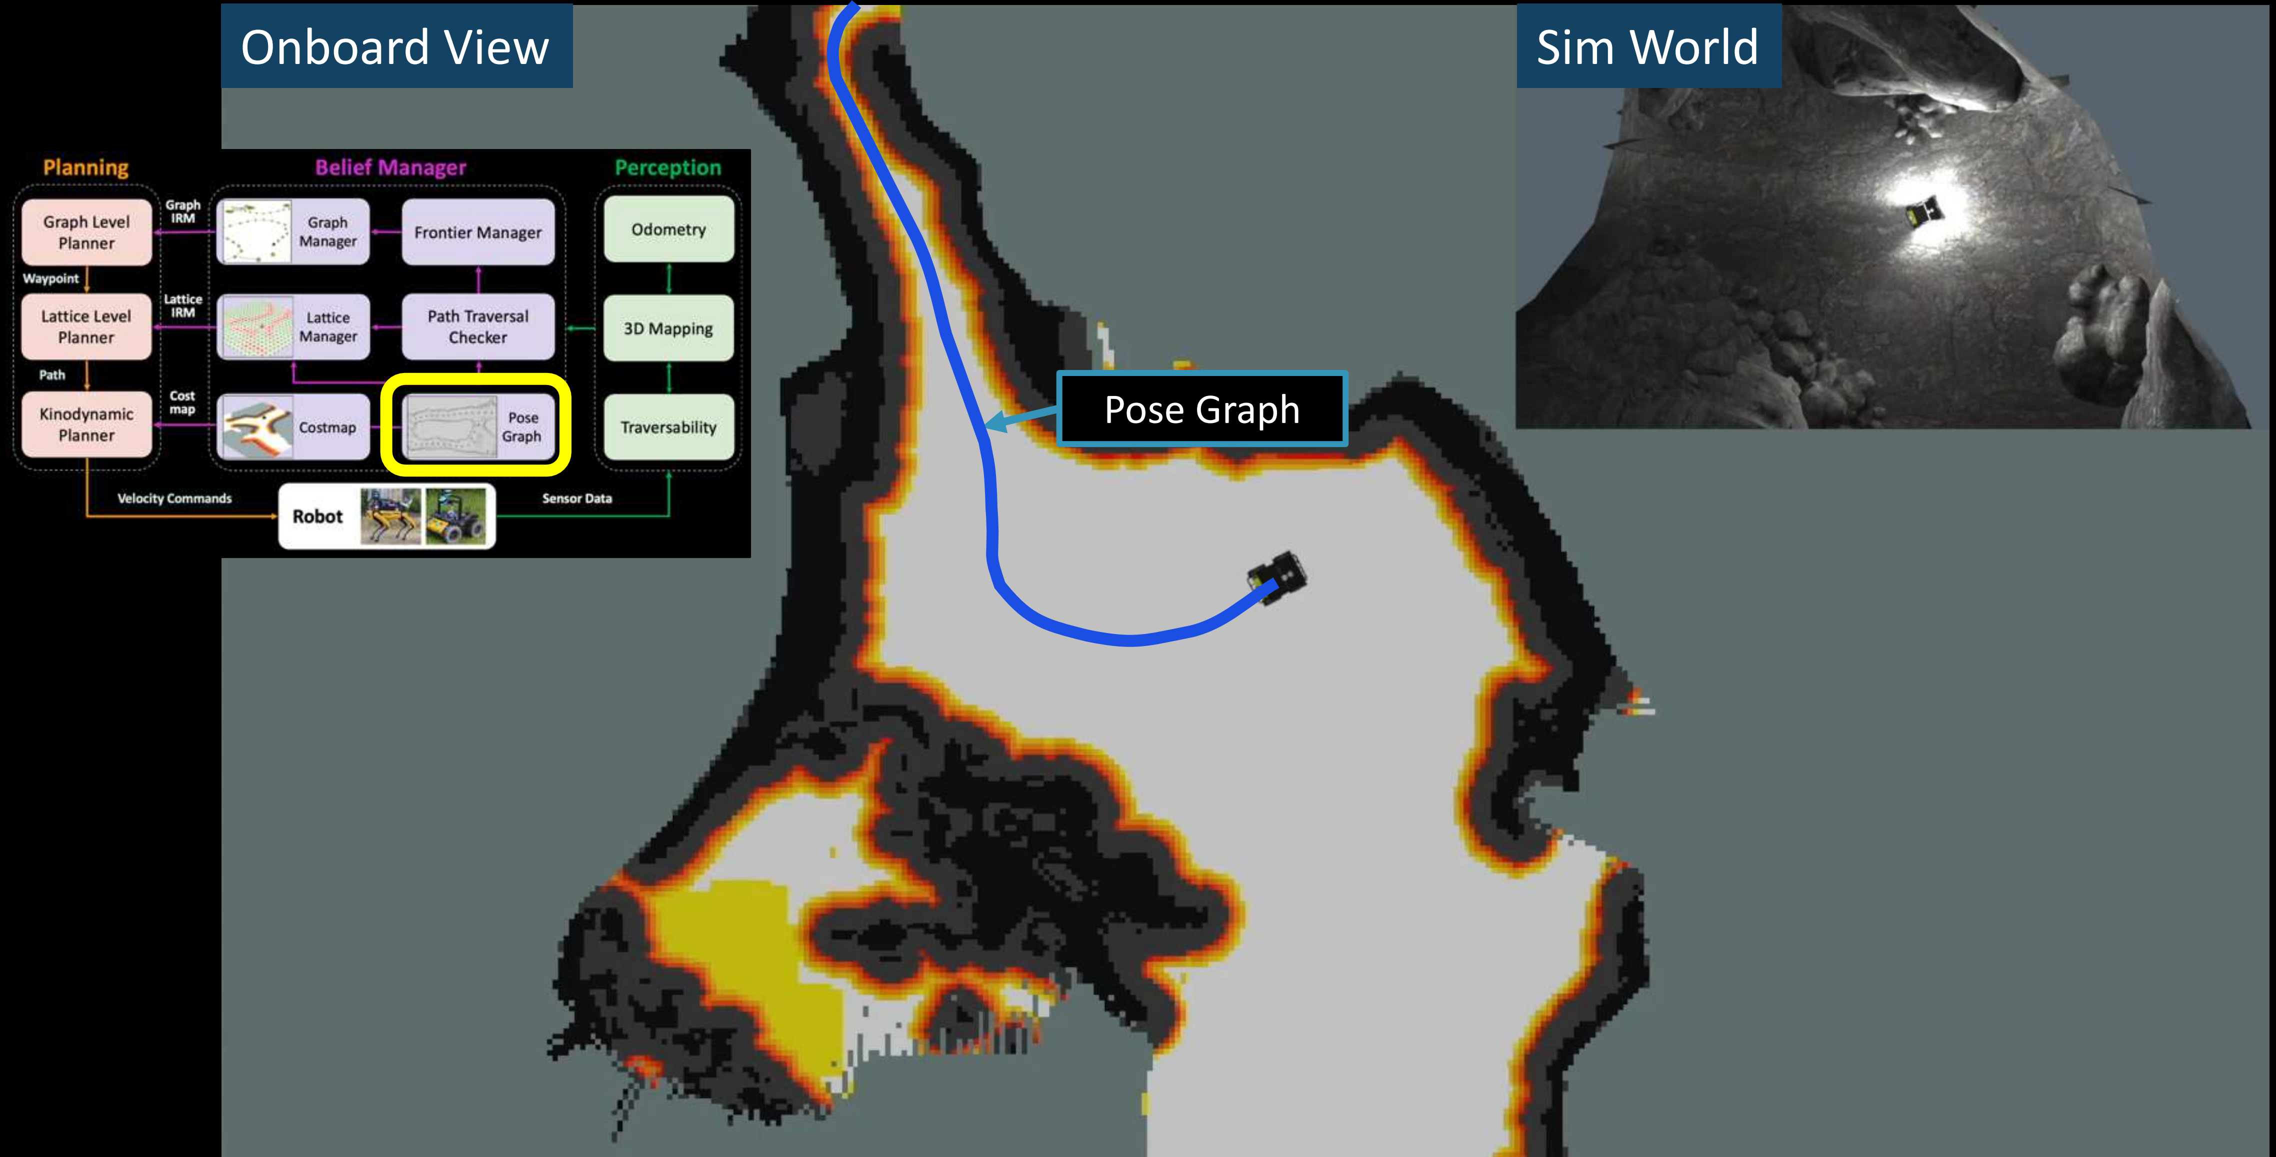
Onboard View
Sim World
Pose Graph
Planning
Belief Manager
Perception
Graph Level Planner
Graph IRM
Graph Manager
Frontier Manager
Odometry
Waypoint
Lattice Level Planner
Lattice IRM
Lattice Manager
Path Traversal Checker
3D Mapping
Path
Kinodynamic Planner
Cost map
Costmap
Pose Graph
Traversability
Velocity Commands
Robot
Sensor Data

# Results

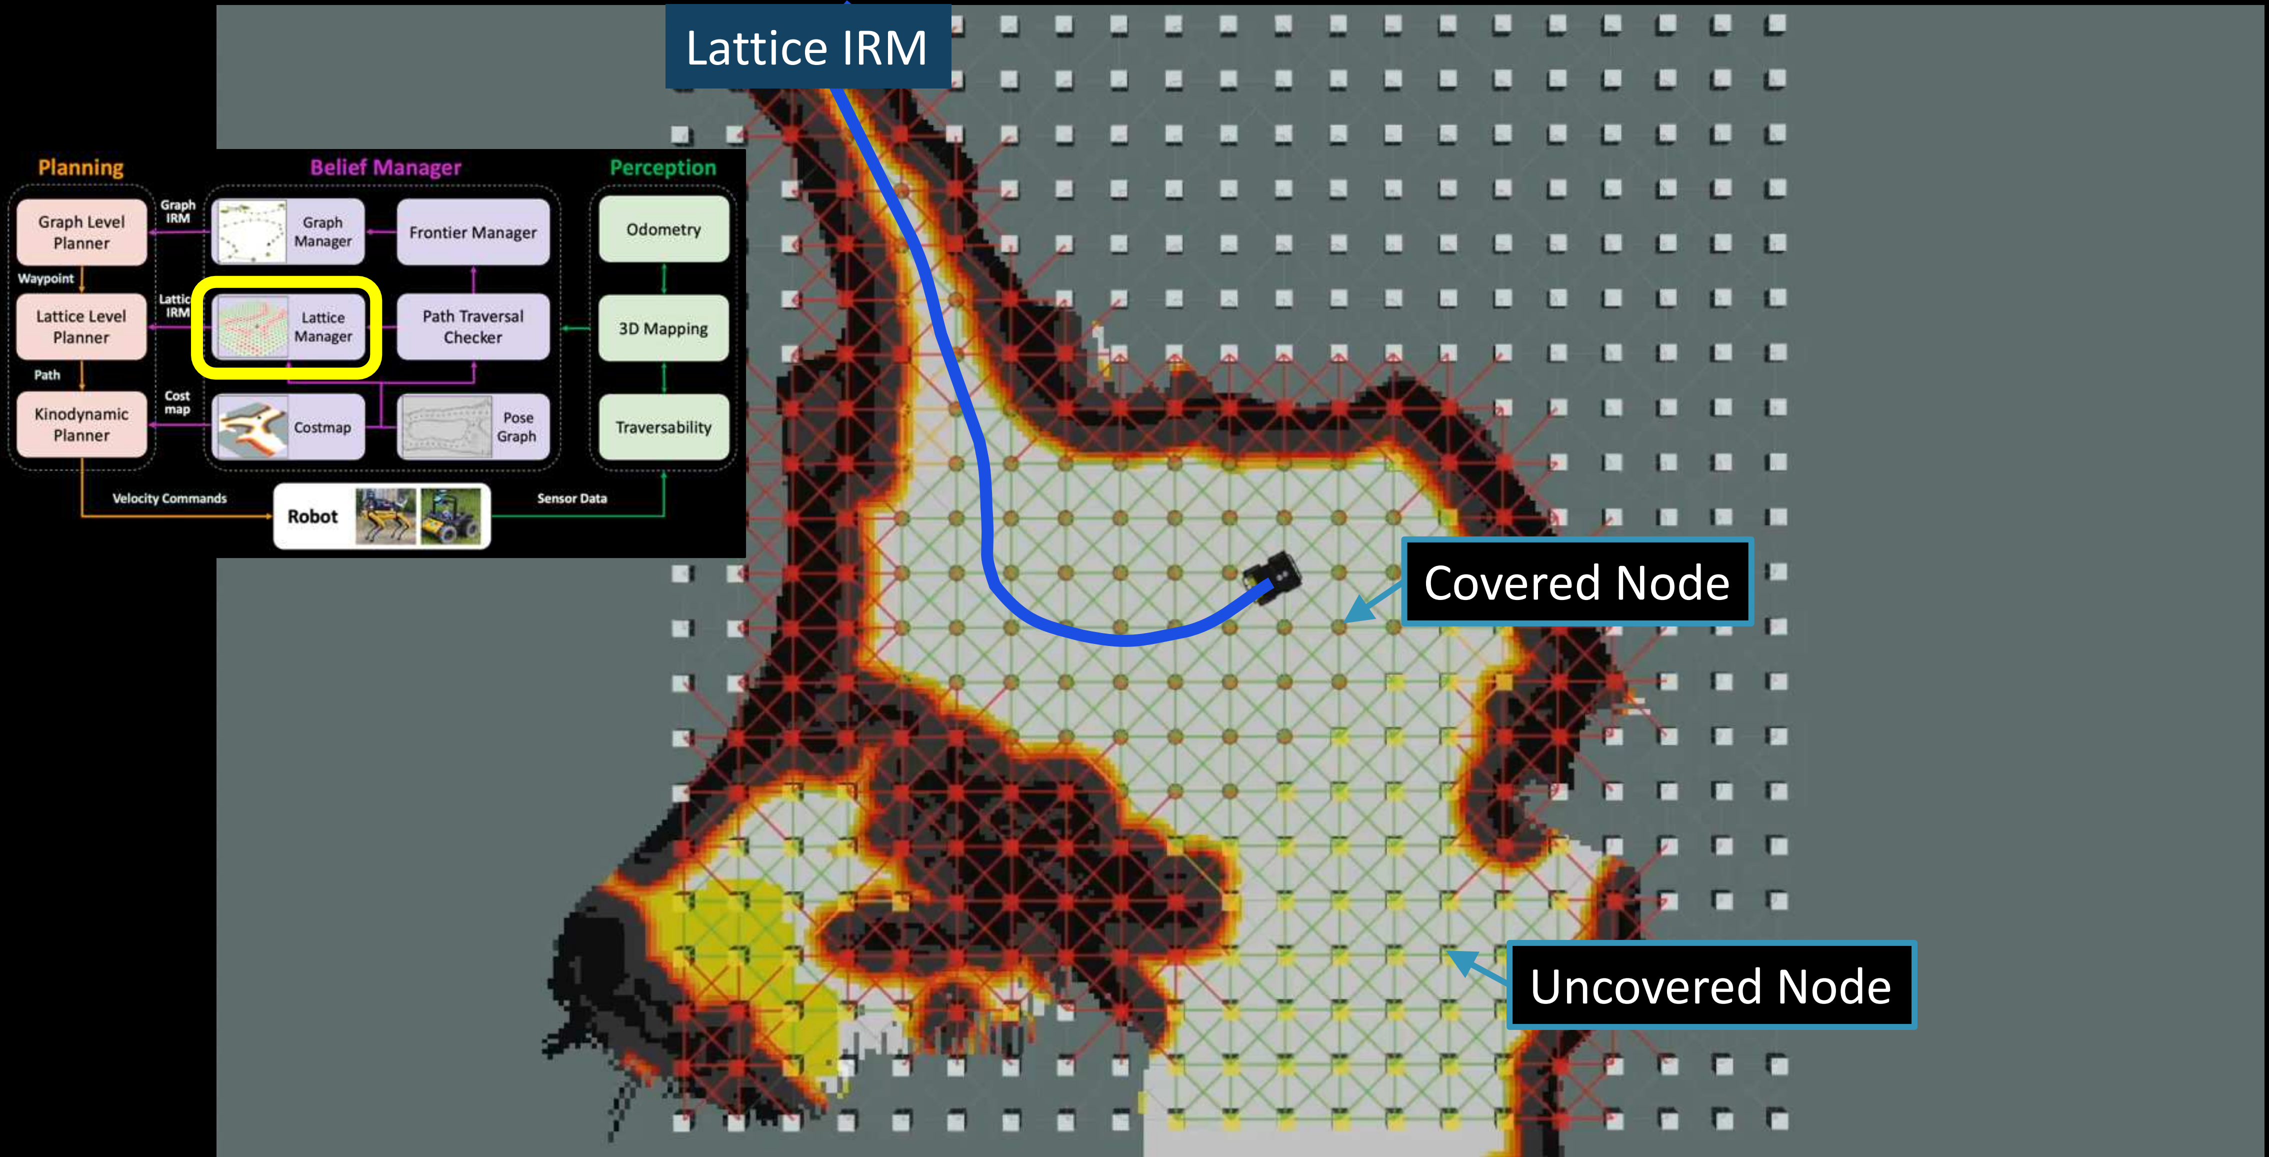
Lattice IRM
Planning
Belief Manager
Perception
Graph Level Planner
Graph IRM
Graph Manager
Frontier Manager
Odometry
Waypoint
Lattice Level Planner
Lattice IRM
Lattice Manager
Path Traversal Checker
3D Mapping
Path
Kinodynamic Planner
Cost map
Costmap
Pose Graph
Traversability
Velocity Commands
Robot
Sensor Data
Covered Node
Uncovered Node

# Results

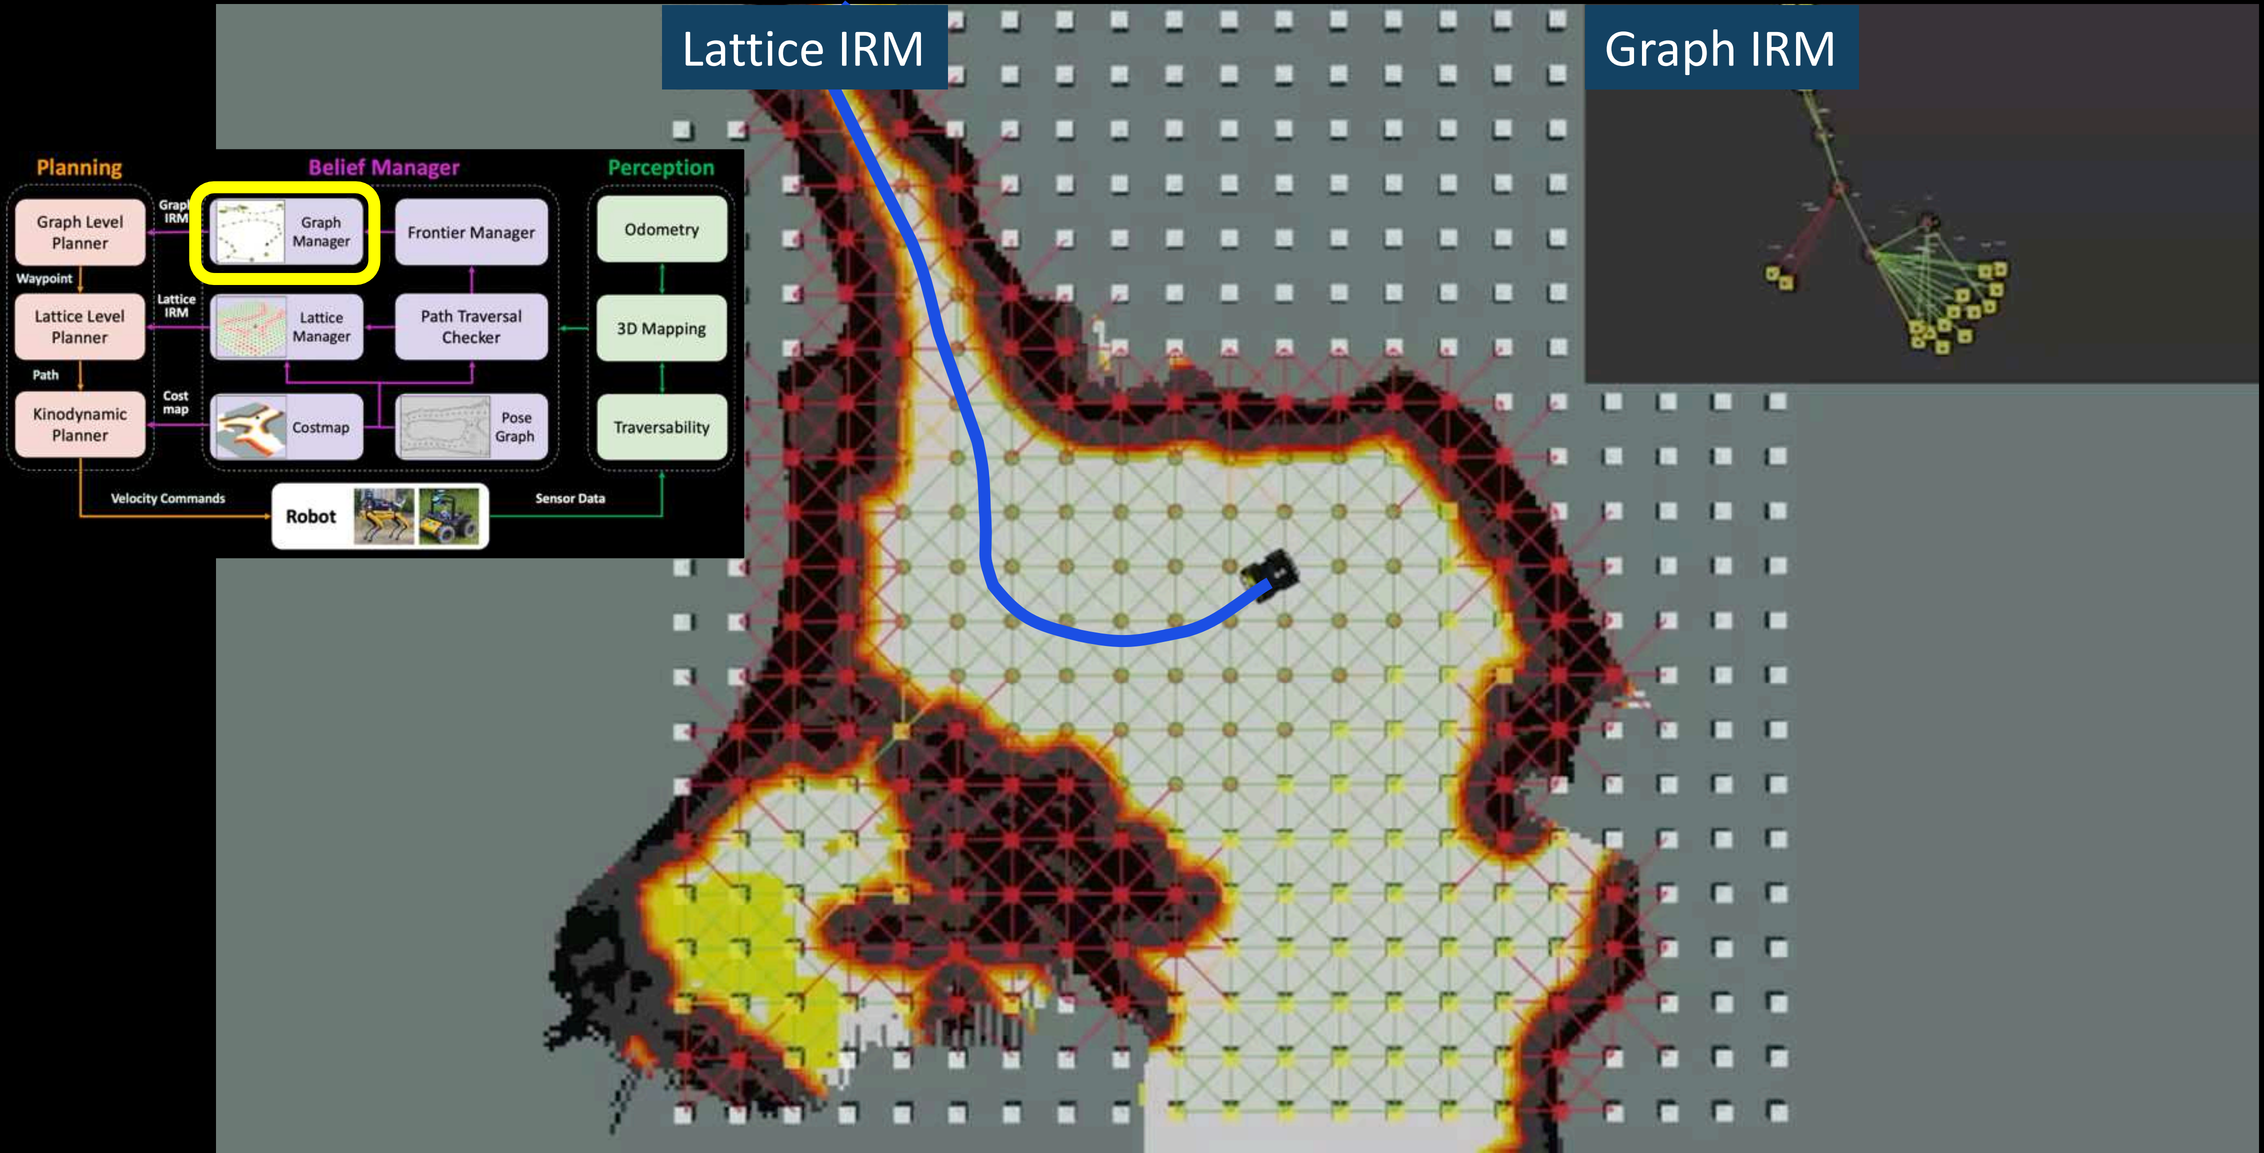

Lattice IRM
Graph IRM
Planning
Belief Manager
Perception
Graph Level Planner
Graph IRM
Graph Manager
Frontier Manager
Odometry
Waypoint
Lattice Level Planner
Lattice IRM
Lattice Manager
Path Traversal Checker
3D Mapping
Path
Kinodynamic Planner
Cost map
Costmap
Pose Graph
Traversability
Velocity Commands
Robot
Sensor Data

# Results

# Results

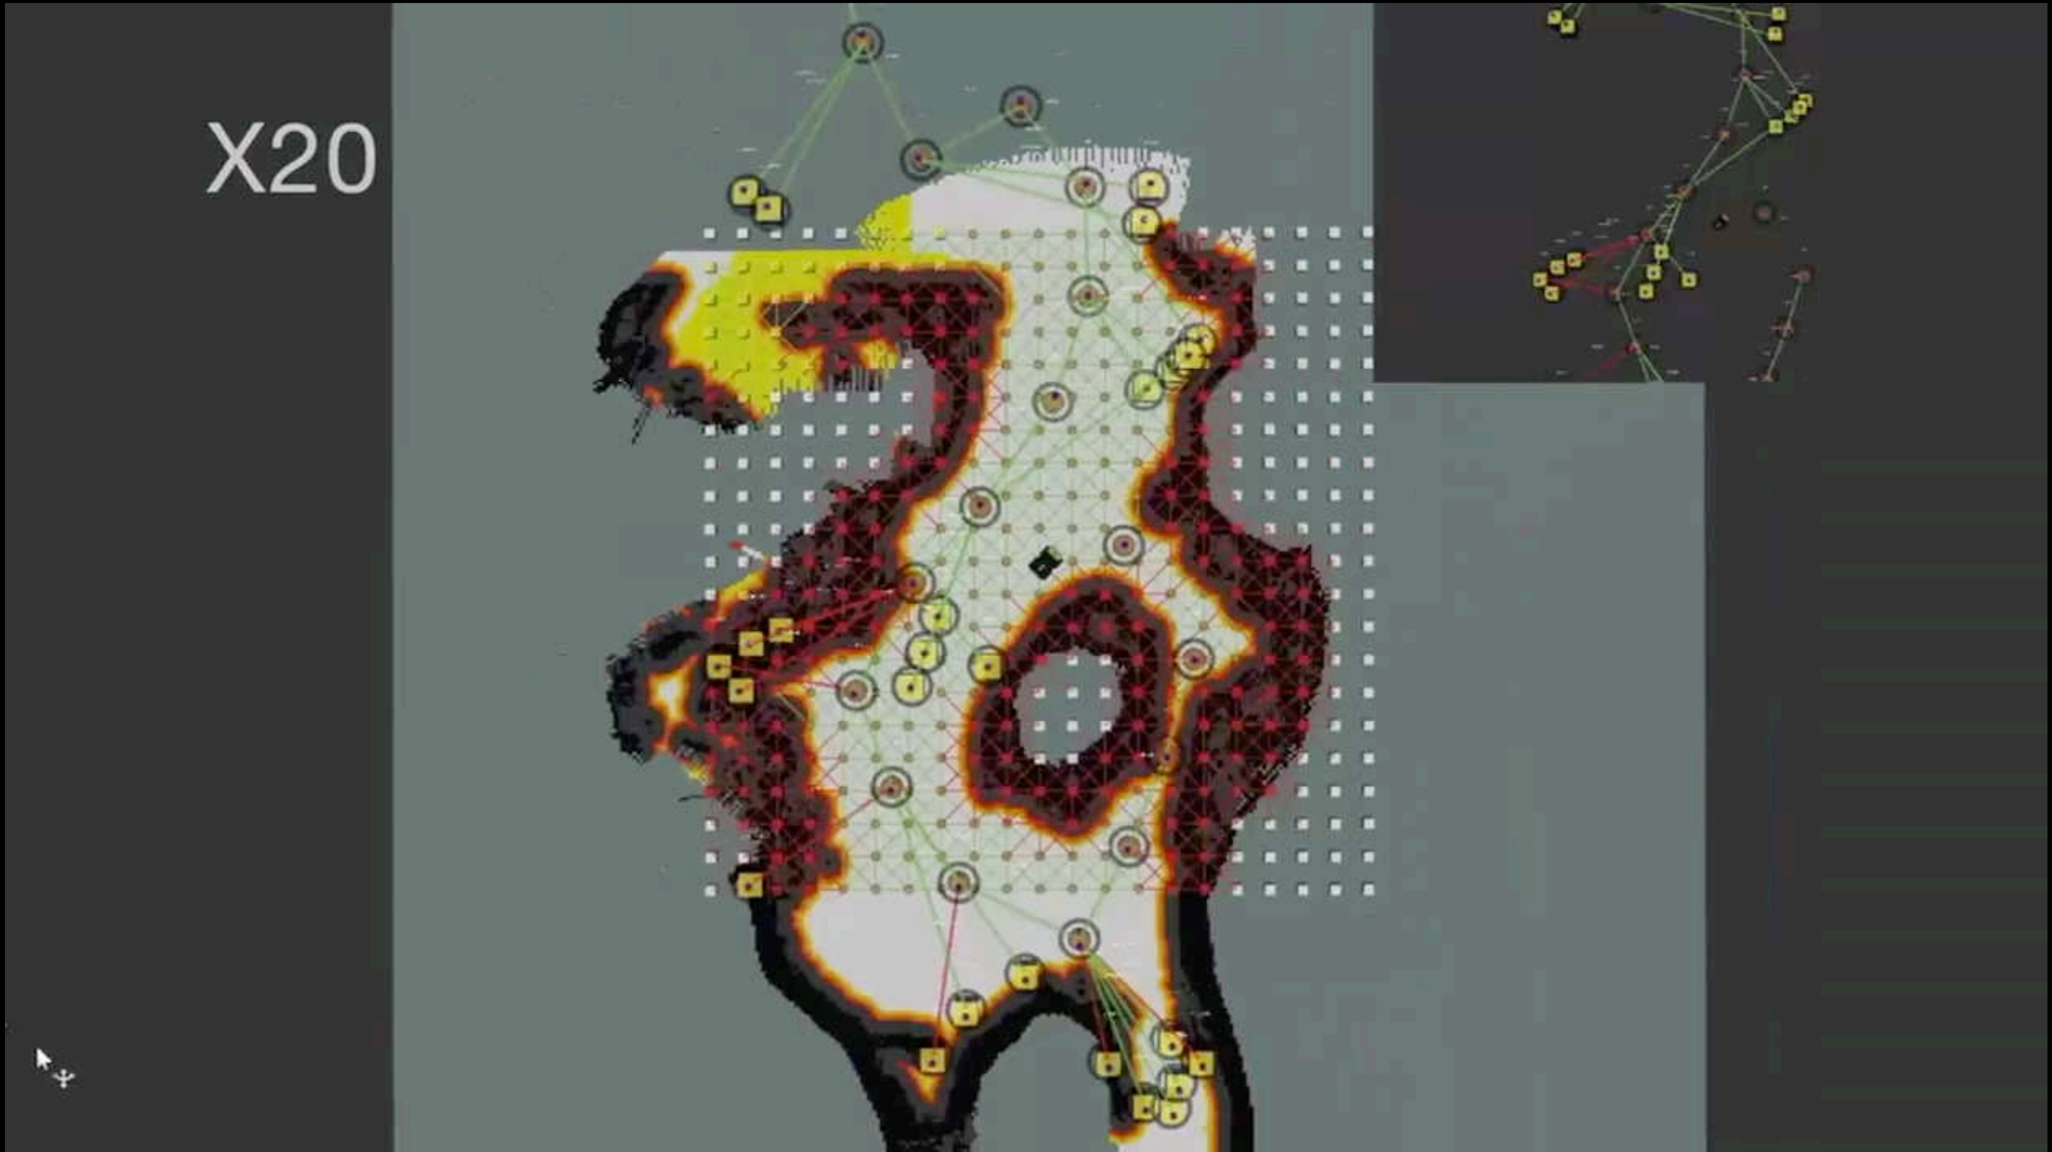

# Results

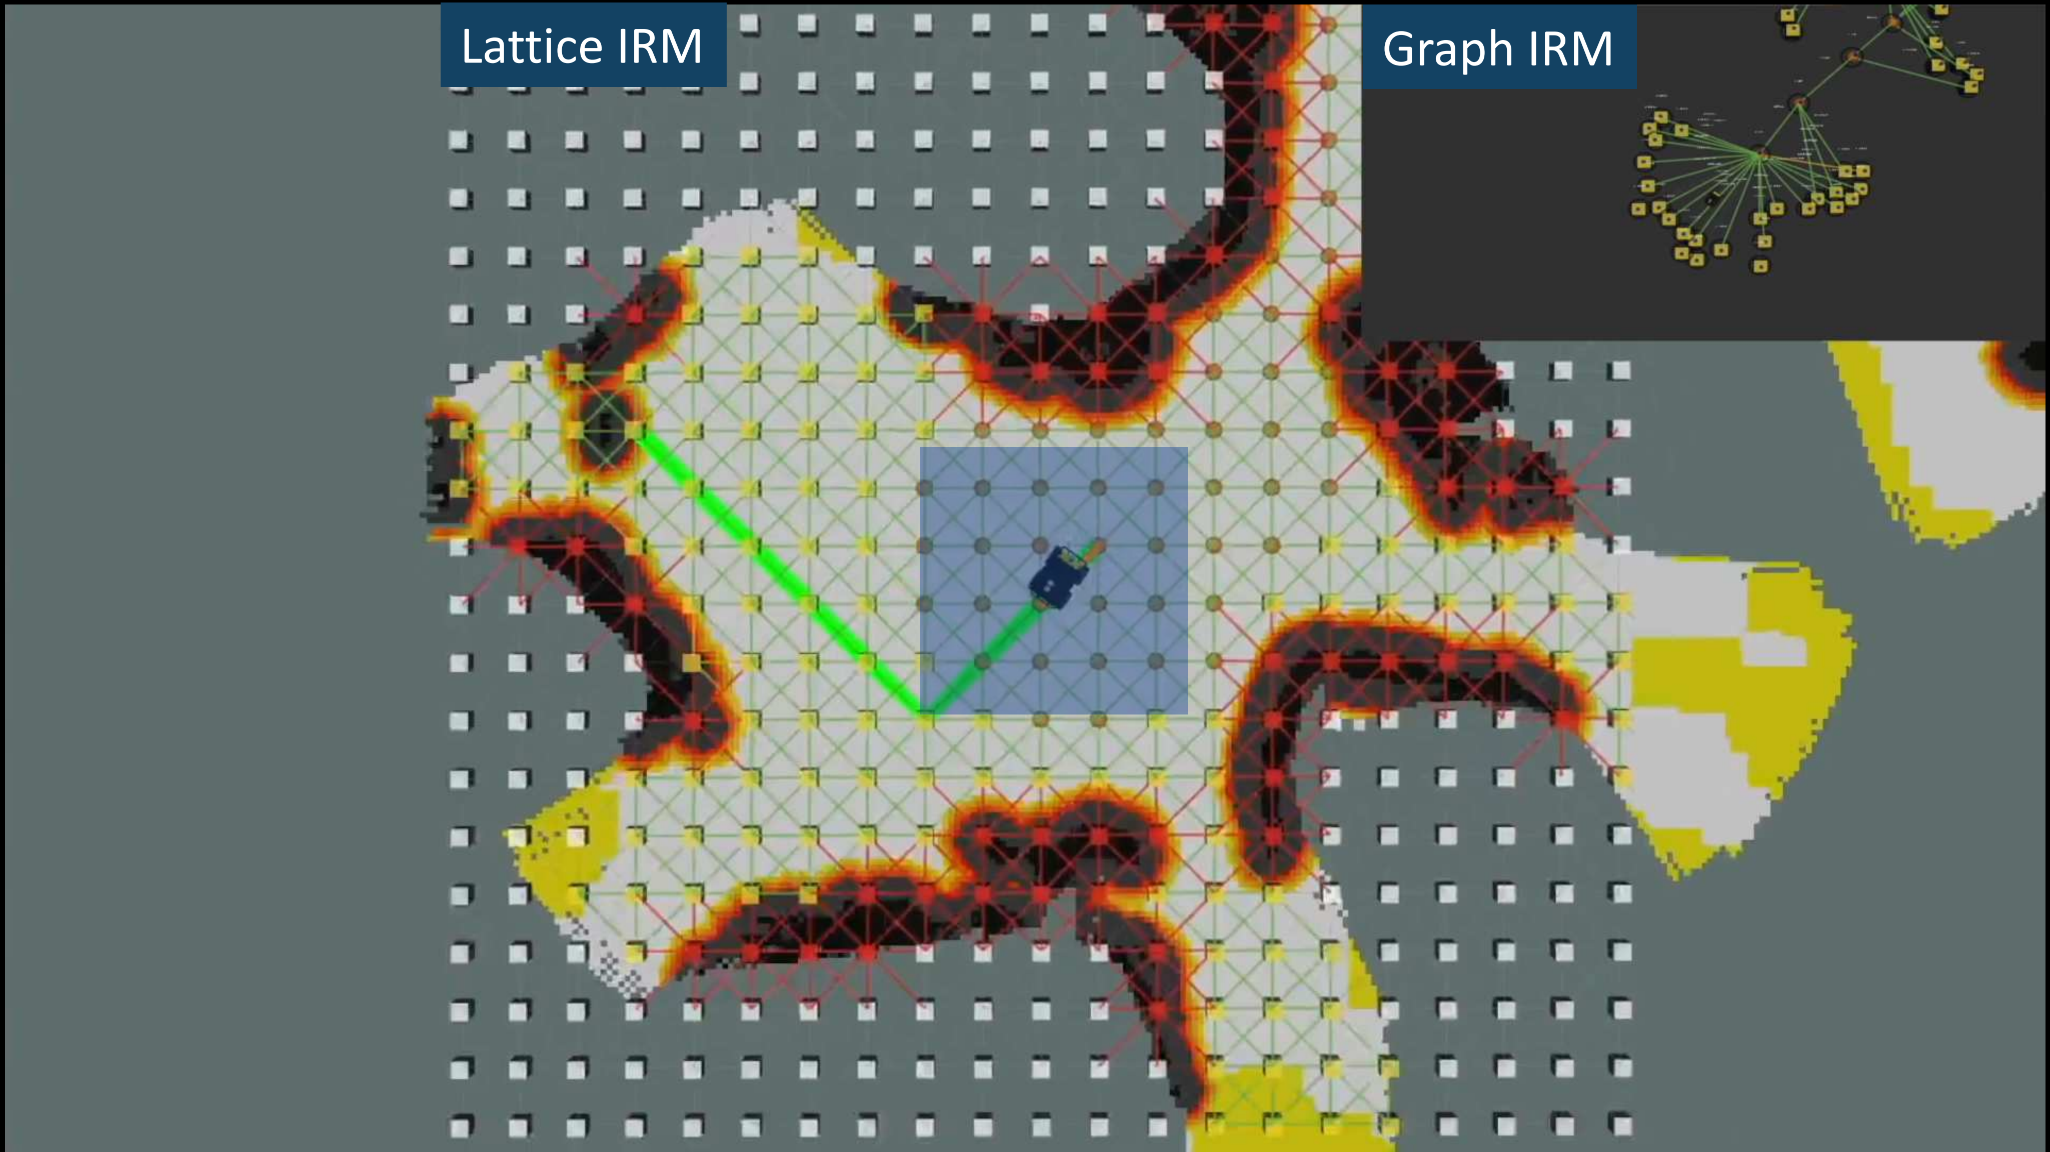

# Results

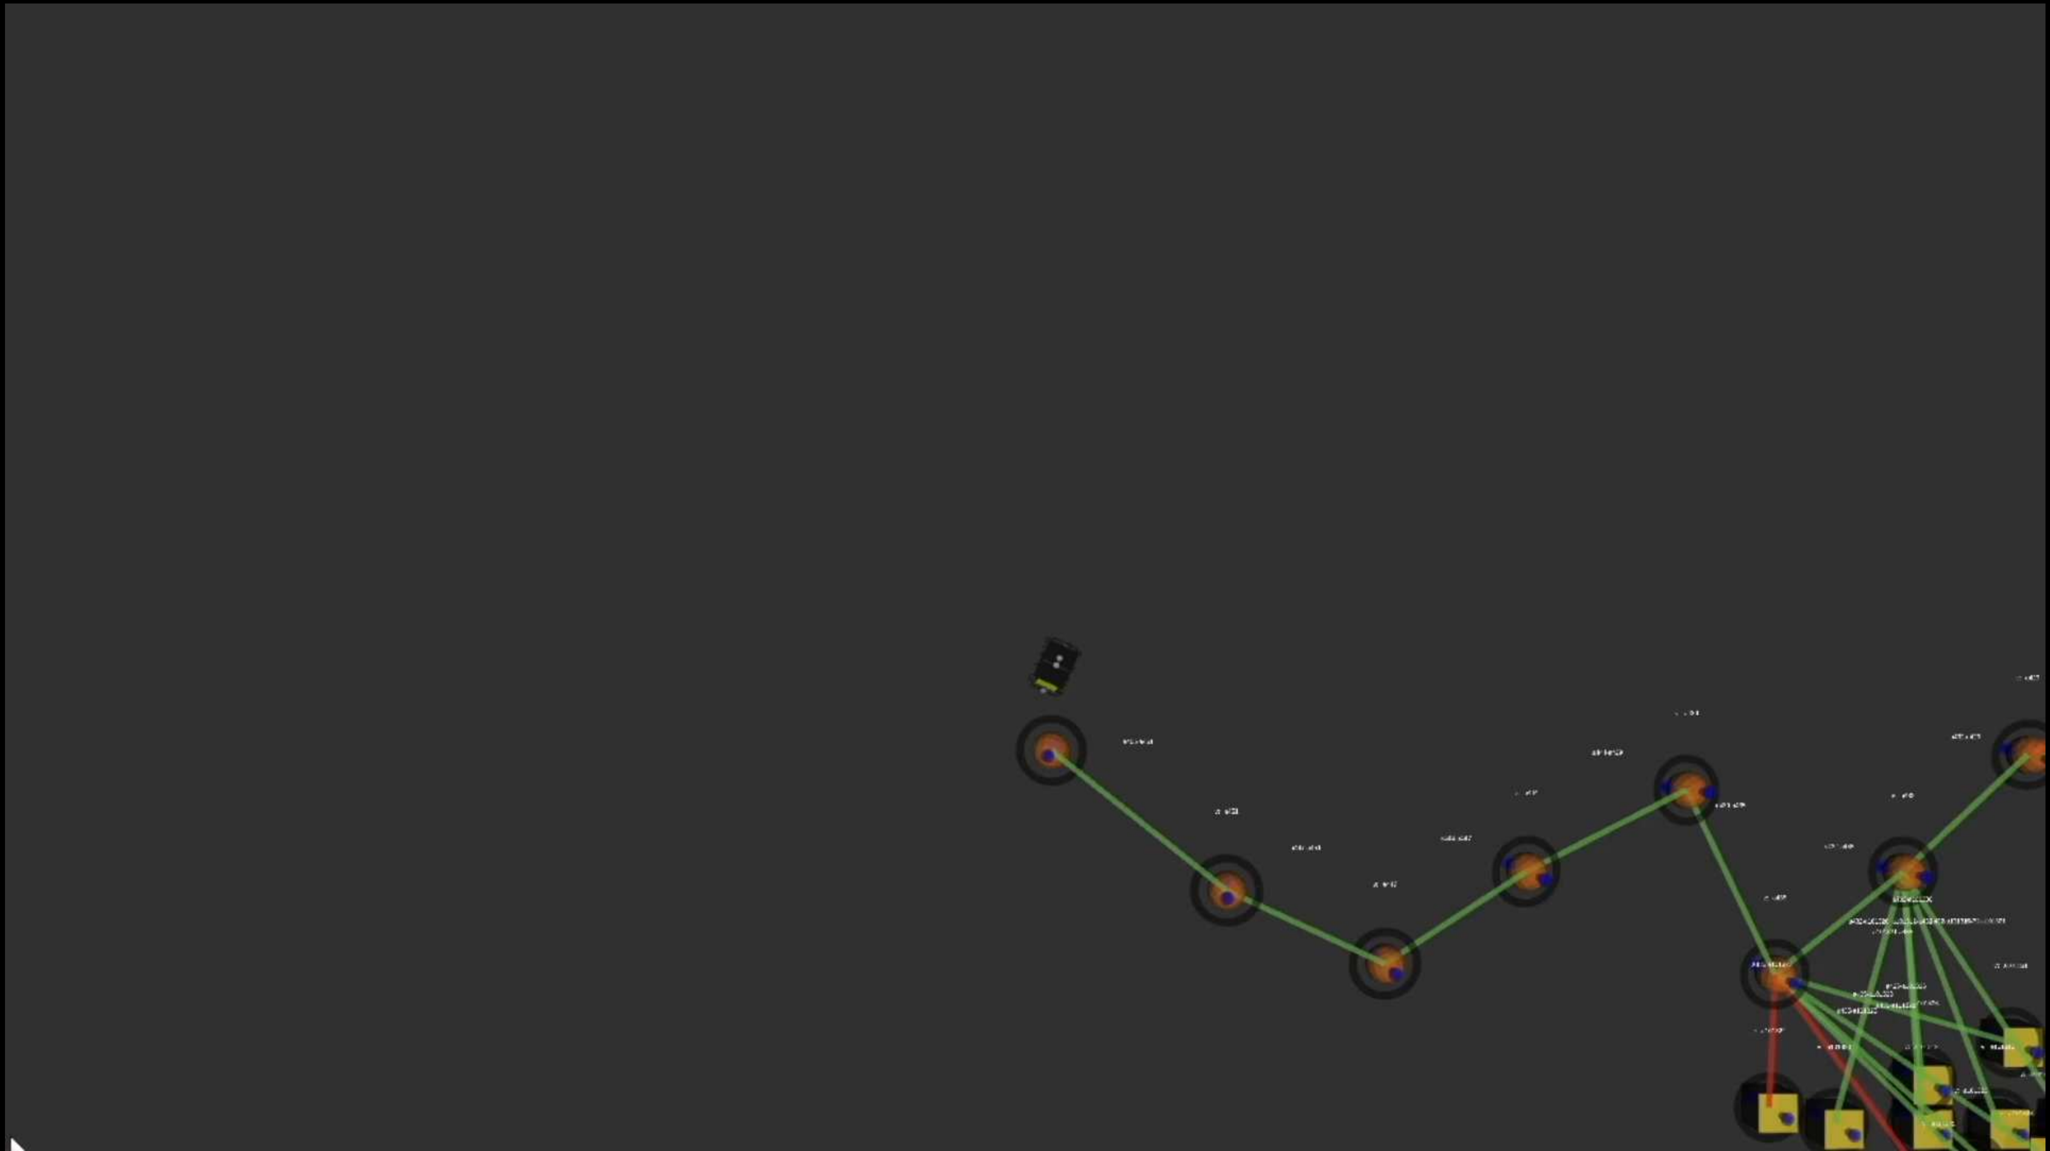

# Results

- Enabled large-scale autonomous exploration in unknown environments
- Theoretically grounded approach
- Demonstrated simulation and (initial) hardware tests

- Adaptive replanning in dynamic environments
- Multi-robot coverage planning
- Learning-based approaches

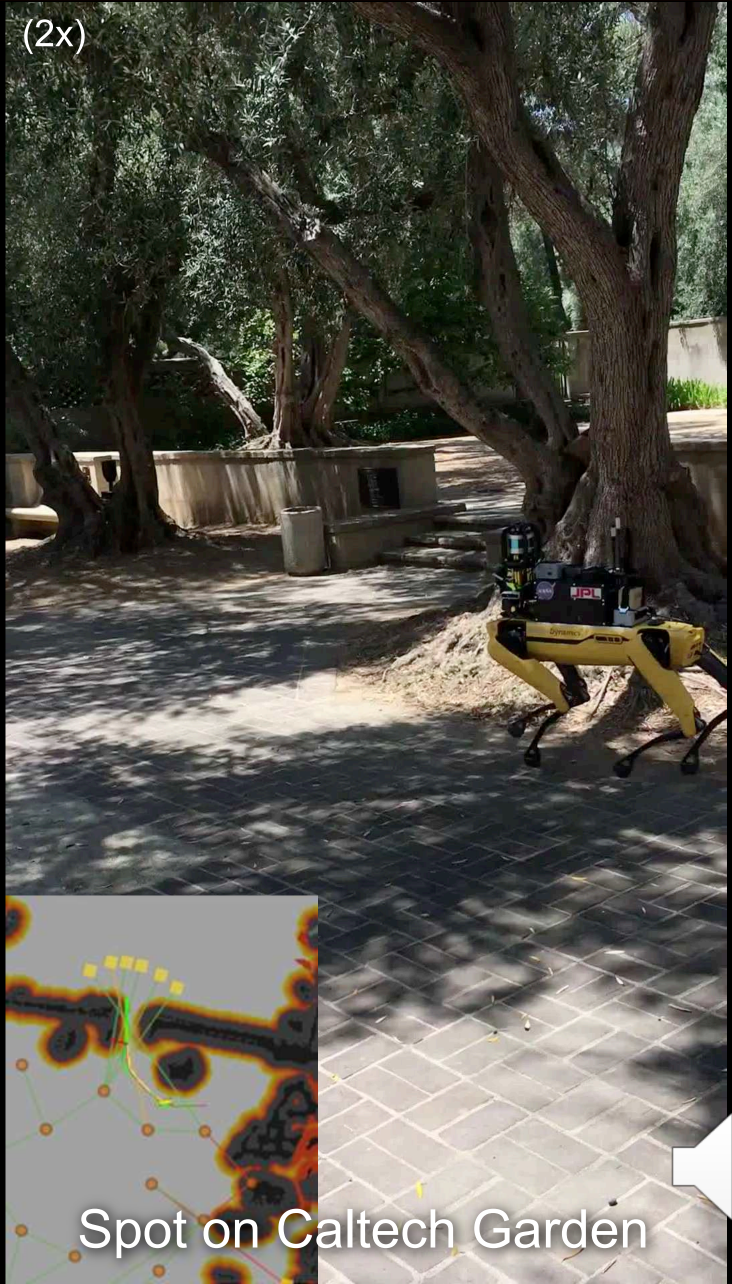

# Publications and References

**Sung-Kyun Kim**, Amanda Bouman, Gautam Salhotra, David D. Fan, Kyohei Otsu, Joel Burdick, Ali-akbar Agha-mohammadi, "PLGRIM: Hierarchical Value Learning for Large-scale Coverage in Unknown Environments," IEEE Robotics and Automation Letters (RA-L), 2020. *In preparation*.

**Belief Space Planning**

[1] **Sung-Kyun Kim**, Rohan Thakker, and Ali-akbar Agha-mohammadi, "Bi-directional Value Learning for Risk-aware Planning Under Uncertainty," IEEE Robotics and Automation Letters (RA-L), vol. 4, no. 3, pp. 2493-2500, 2019.

[2] **Sung-Kyun Kim**, Oren Salzman, and Maxim Likhachev, "POMHDP: Search-based Belief Space Planning using Multiple Heuristics," International Conference on Automated Planning and Scheduling (ICAPS), pp. 734-744, 2019.

[3] Ali-akbar Agha-mohammadi, Saurav Agarwal, **Sung-Kyun Kim**, Suman Chakravorty, and Nancy M. Amato, "SLAP: Simultaneous Localization and Planning for Physical Mobile Robots via Enabling Dynamic Replanning in Belief Space," IEEE Transactions on Robotics (TRO), vol. 34, no. 5, pp. 1195-1214, 2018.

**Belief Space Representation**

[4] Ali-akbar Agha-mohammadi, Eric Heiden, Karol Hausman, and Gaurav S. Sukhatme, "Confidence-rich 3D Grid Mapping: Toward High-speed Vision-based UAV Navigation," International Journal of Robotics Research (IJRR), vol. 38, pp. 1352-1374, 2019.

[5] Ali-akbar Agha-mohammadi, Suman Chakravorty, and Nancy Amato, "FIRM: Sampling-based Feedback Motion Planning Under Motion Uncertainty and Imperfect Measurements," International Journal of Robotics Research (IJRR), v 33, no. 2, pp. 268-304, 2014.

# Backup Slides

- Objective
  - Maximize the coverage of the environment within a given time
- Approach
  - **Longer-horizon** planning under uncertainty
    - Compared to frontier-based exploration with one-step look-ahead greedy policy
  - **Compact** environment representation
    - To reduce the problem complexity (dimensionality)
  - **Hierarchical** framework to scale to large environments
    - To reduce the problem complexity (planning horizon, i.e., history)

- Objective
  - Maximize the **information gain** per **action cost** for a finite planning horizon, in a receding horizon control fashion (online replanning)
- Approach
  - **Longer-horizon** planning under uncertainty
    - Compared to frontier-based exploration with one-step look-ahead greedy policy
  - **Compact** environment representation
    - To reduce the problem complexity (dimensionality)
  - **Hierarchical** framework to scale to large environments
    - To reduce the problem complexity (planning horizon, i.e., history)

- Objective
  - Maximize the **information gain** per **action cost** for a finite planning horizon, in a receding horizon control fashion (online replanning)
- Approach
  - **Longer-horizon** planning under uncertainty
    - Based on POMCP (Partially Observable Monte Carlo Planning) algorithm
  - **Compact** environment representation
    - To reduce the problem complexity (dimensionality)
  - **Hierarchical** framework to scale to large environments
    - To reduce the problem complexity (planning horizon, i.e., history)

- Objective
  - Maximize the **information gain** per **action cost** for a finite planning horizon, in a receding horizon control fashion (online replanning)
- Approach
  - **Longer-horizon** planning under uncertainty
    - Based on POMCP (Partially Observable Monte Carlo Planning) algorithm
  - **Compact** environment representation
    - Information Roadmap (IRM) that compactly encodes high-fidelity information
  - **Hierarchical** framework to scale to large environments
    - To reduce the problem complexity (planning horizon, i.e., history)

- Objective
  - Maximize the **information gain** per **action cost** for a finite planning horizon, in a receding horizon control fashion (online replanning)
- Approach
  - **Longer-horizon** planning under uncertainty
    - Based on POMCP (Partially Observable Monte Carlo Planning) algorithm
  - **Compact** environment representation
    - Information Roadmap (IRM) that compactly encodes high-fidelity information
  - **Hierarchical** framework to scale to large environments
    - Cascaded global-local planners for local optimality and global completeness

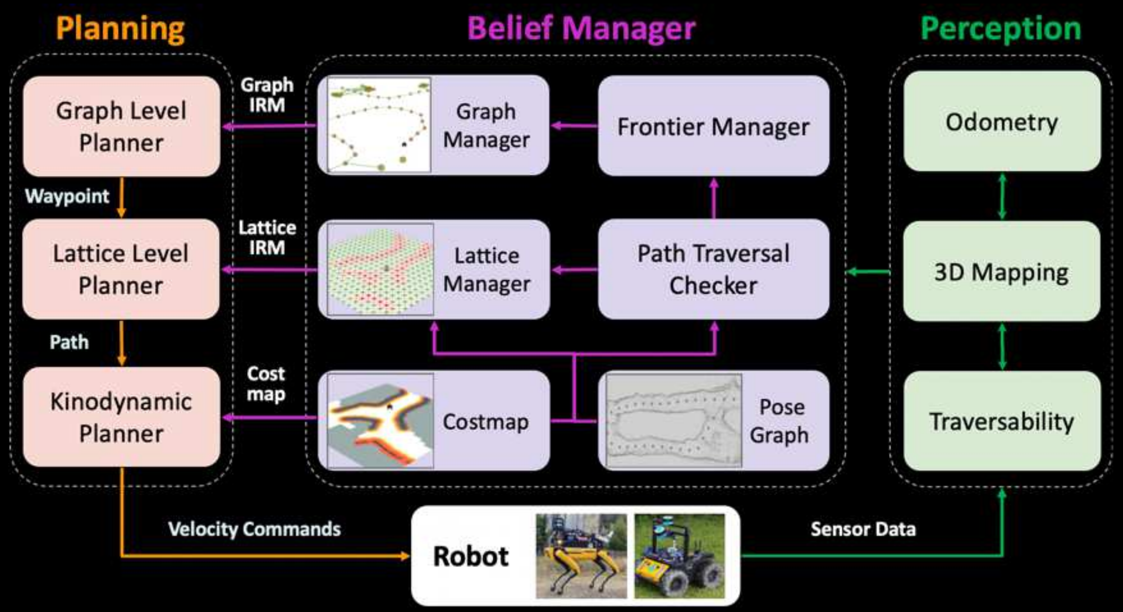


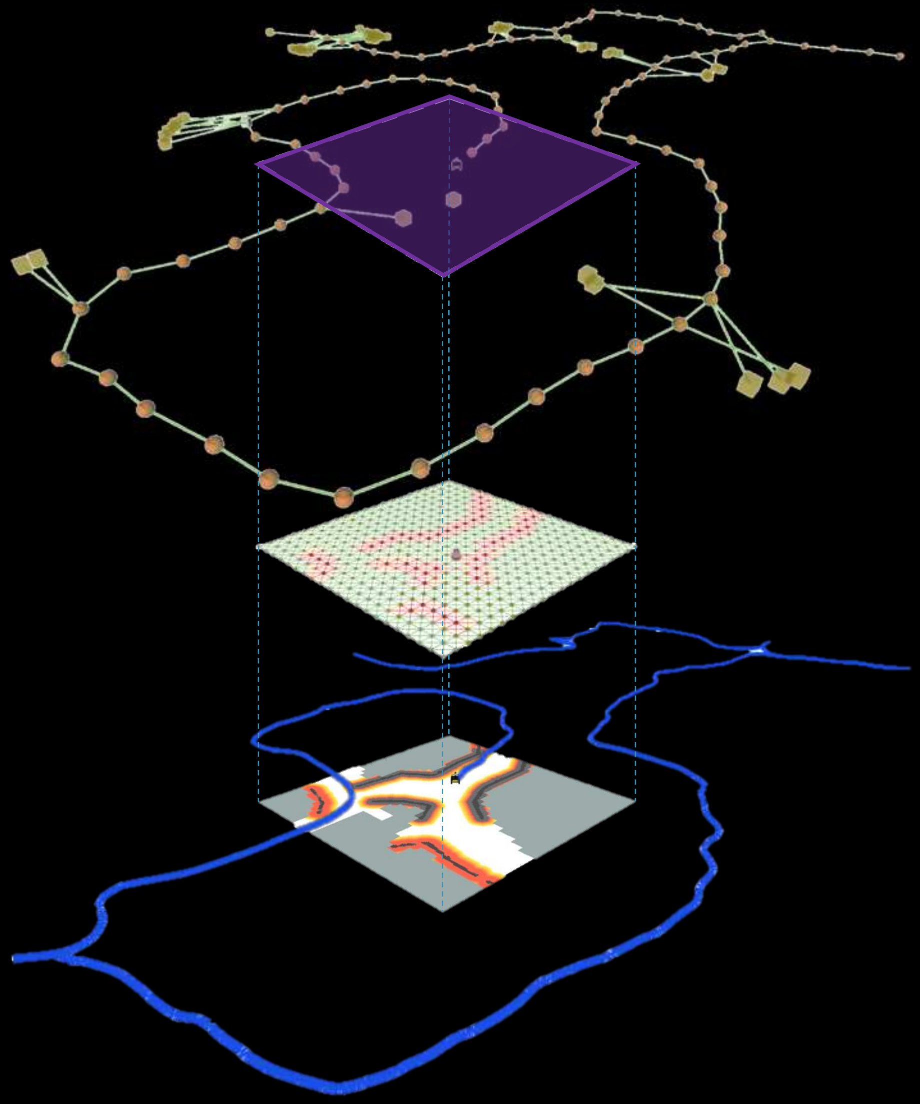

**Optimal Policy:**

$$\pi^*(b) = \arg\max_{\pi} \mathbb{E} \sum_{t=0}^{T} \gamma^t r(b_t, \pi(b_t))$$

**Belief Reward**

$$= \int_s R(s,a)\, b(s)\, ds$$

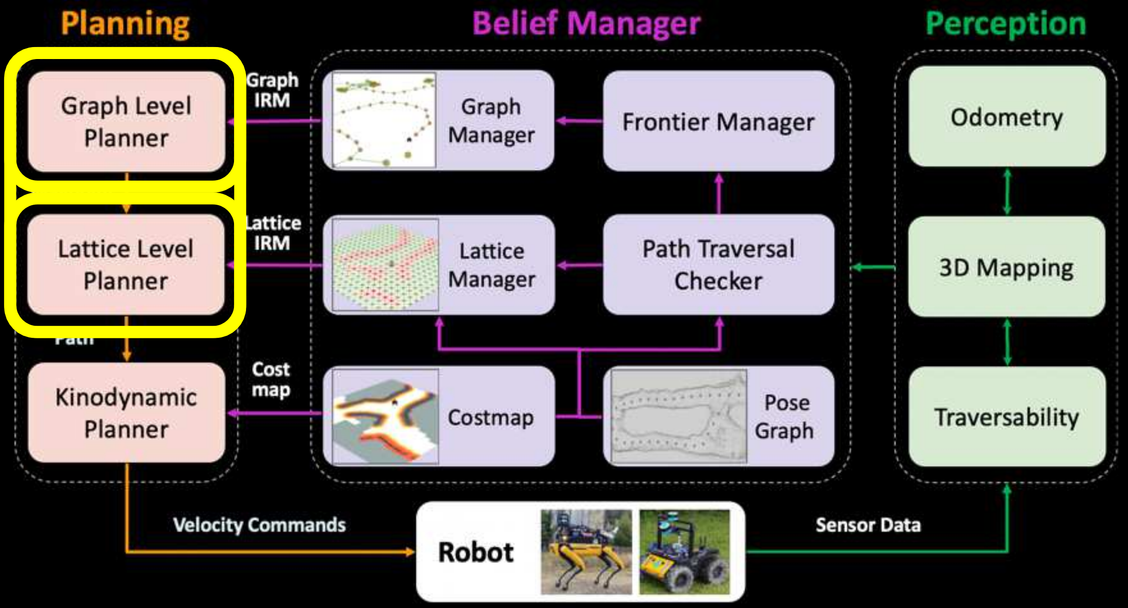


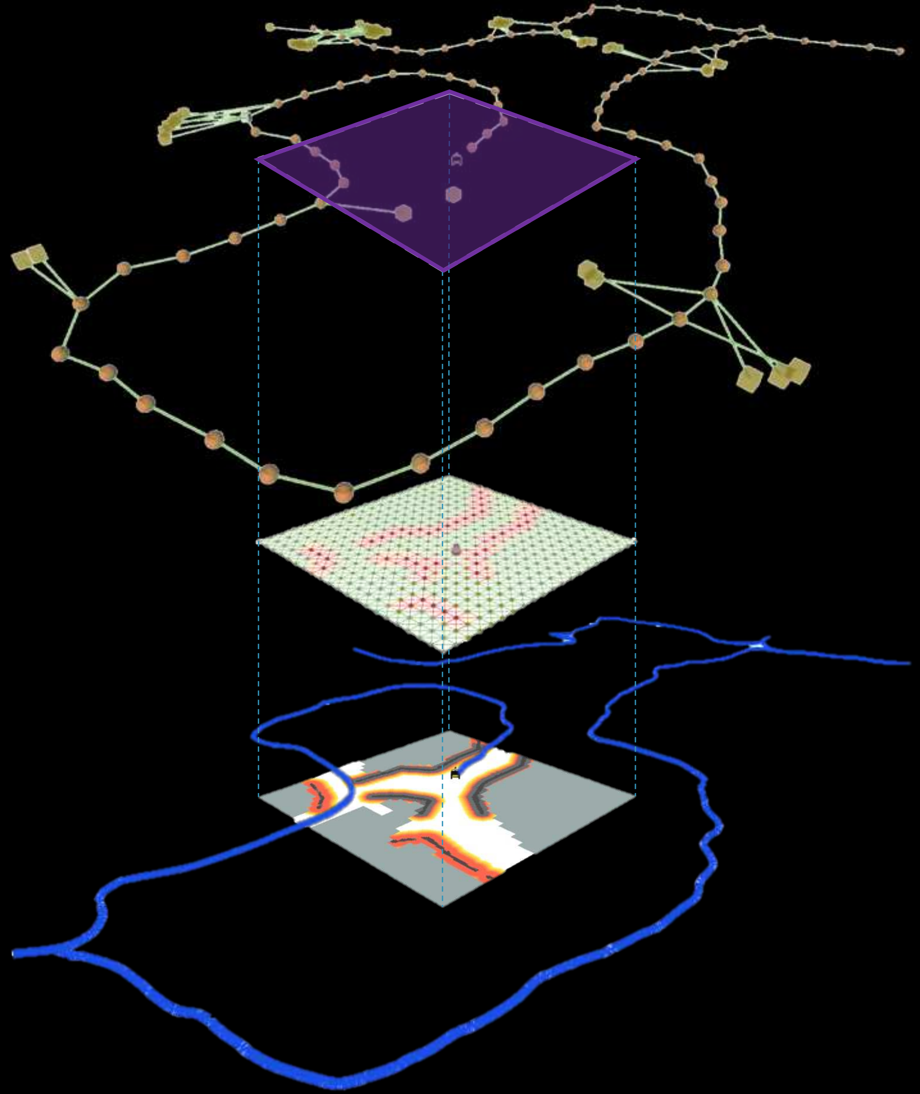

**Graph-Level Policy:** $\pi^g : \mathbb{B}^g \rightarrow \mathbb{A}^g, \ \mathbb{A}^g = \Theta^\ell$

**Lattice-Level Policy:** $\pi^\ell : \mathbb{B}^\ell \times \Theta^\ell \rightarrow \mathbb{A}^\ell$

**Hierarchical Policy:** $\pi(b) = \pi^\ell\left(b^\ell; \pi^g(b^g)\right)$

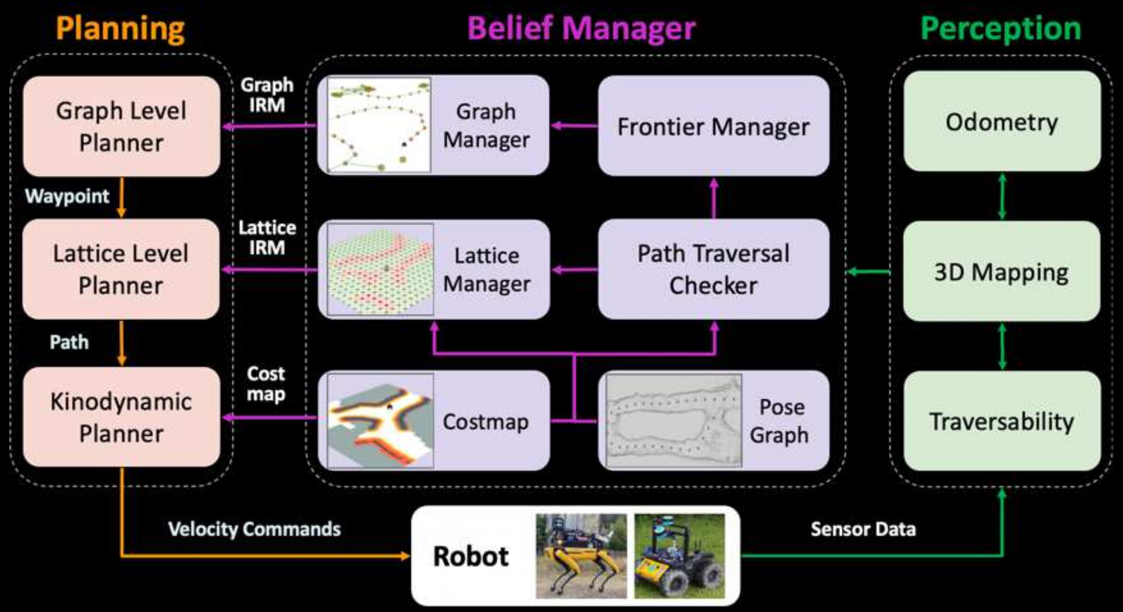


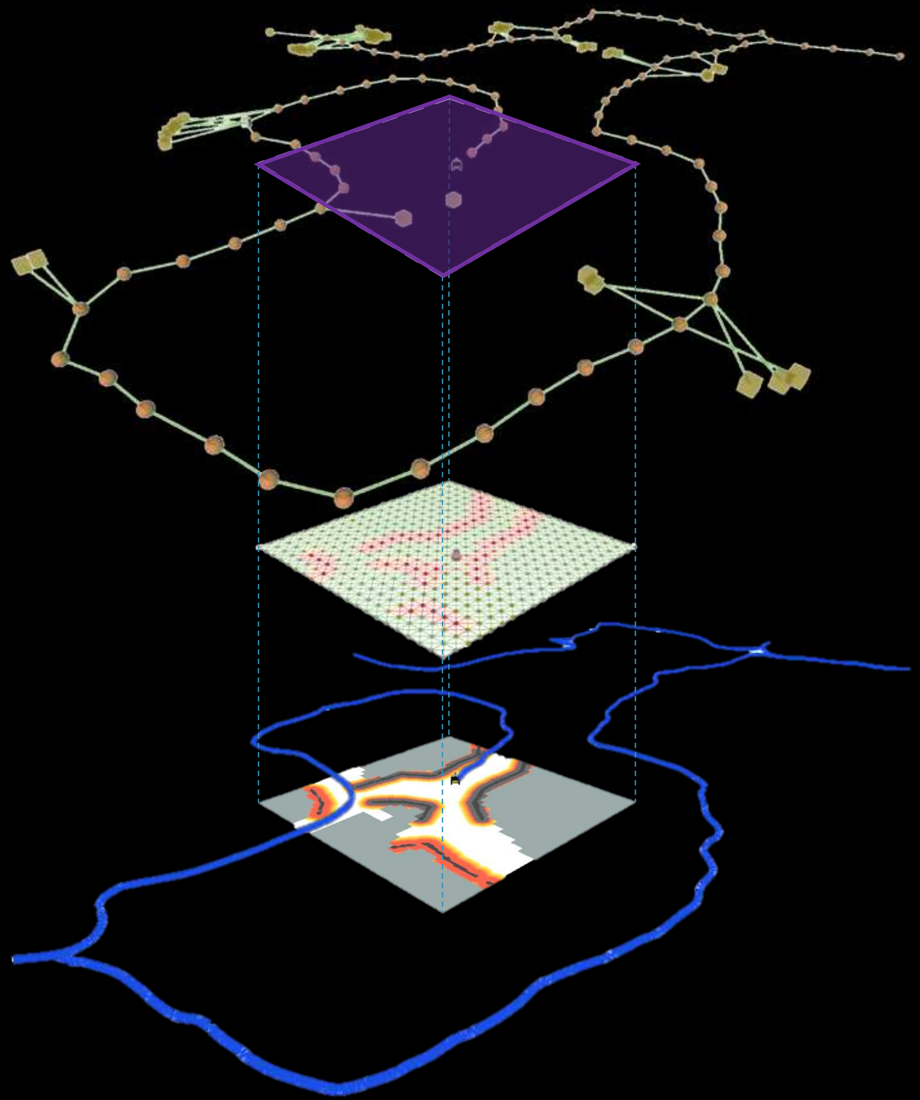

**Optimal Hierarchical Policy:**

$$\pi^*(b) = \arg\max_{\pi} \mathbb{E}\left[\sum_{t=0}^{T} \gamma^t \underbrace{r^\ell\left(b_t^\ell, \pi^\ell\left(b_t^\ell;\ \pi^g(b_t^g)\right)\right)}_{\text{Hierarchical Reward}}\right]$$