Learning to Explore using Active Neural SLAM

ICLR-20

Webpage: https://devendrachaplot.github.io/projects/Neural-SLAM
Code: https://github.com/devendrachaplot/Neural-SLAM

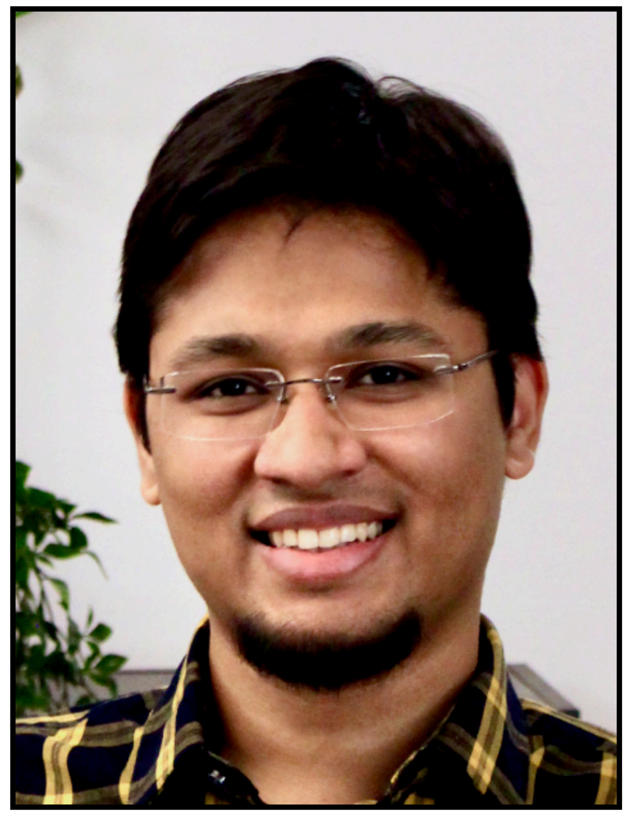

Devendra Singh Chaplot

Dhiraj Gandhi

Saurabh Gupta

Abhinav Gupta

Ruslan Salakhutdinov

Exploration

Exploration

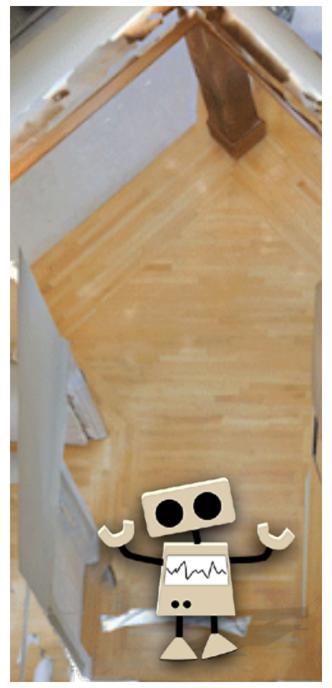

Exploration

Exploration

Exploration

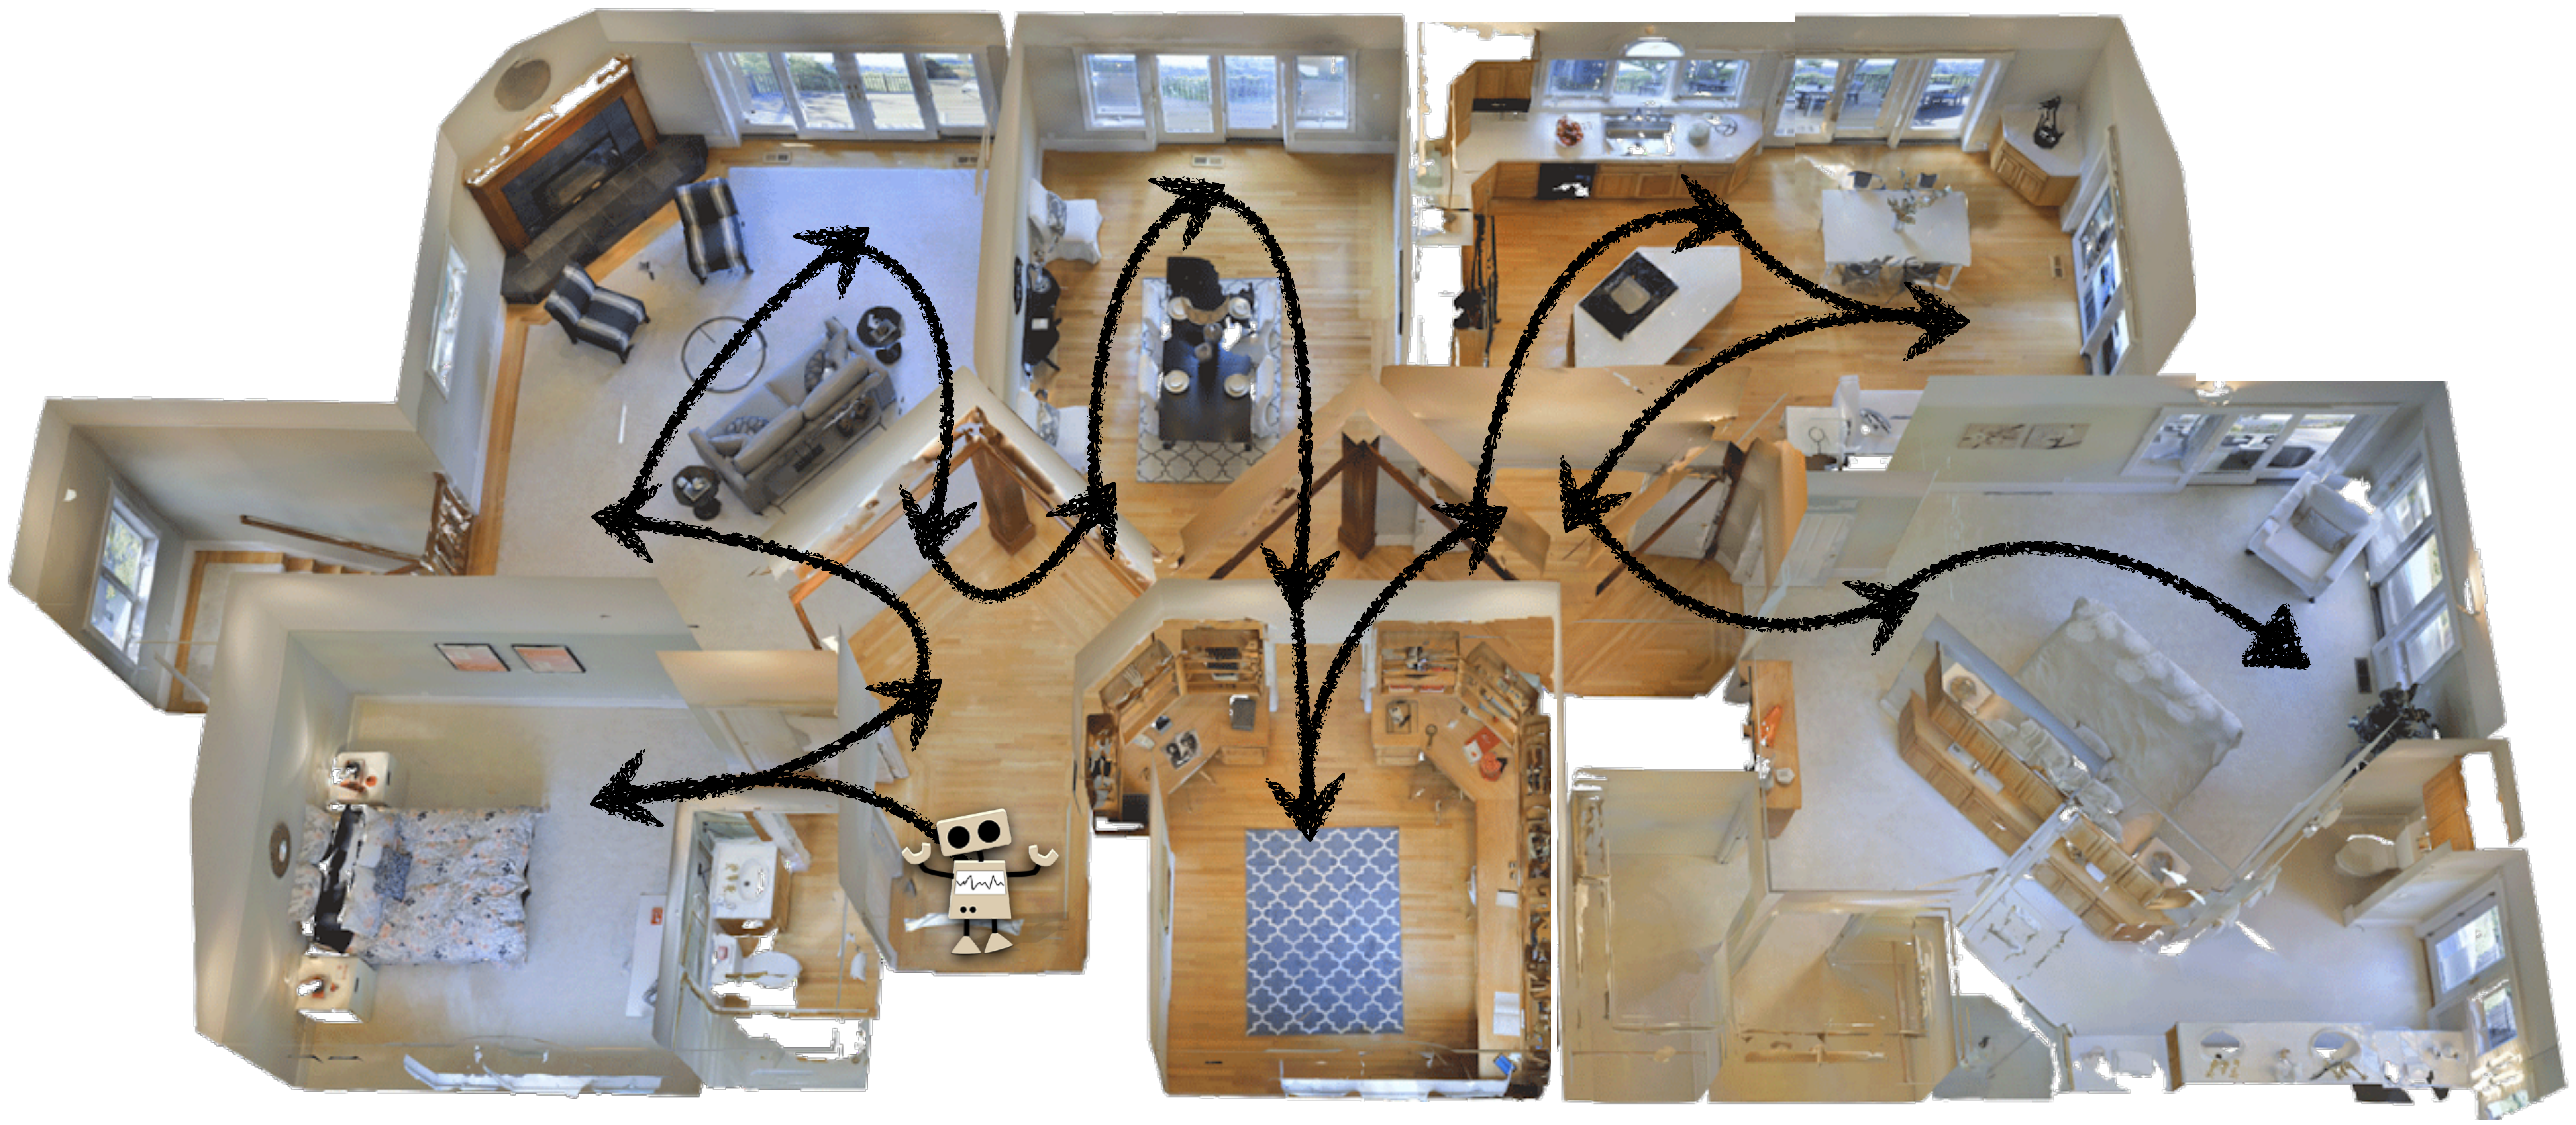

Exploration

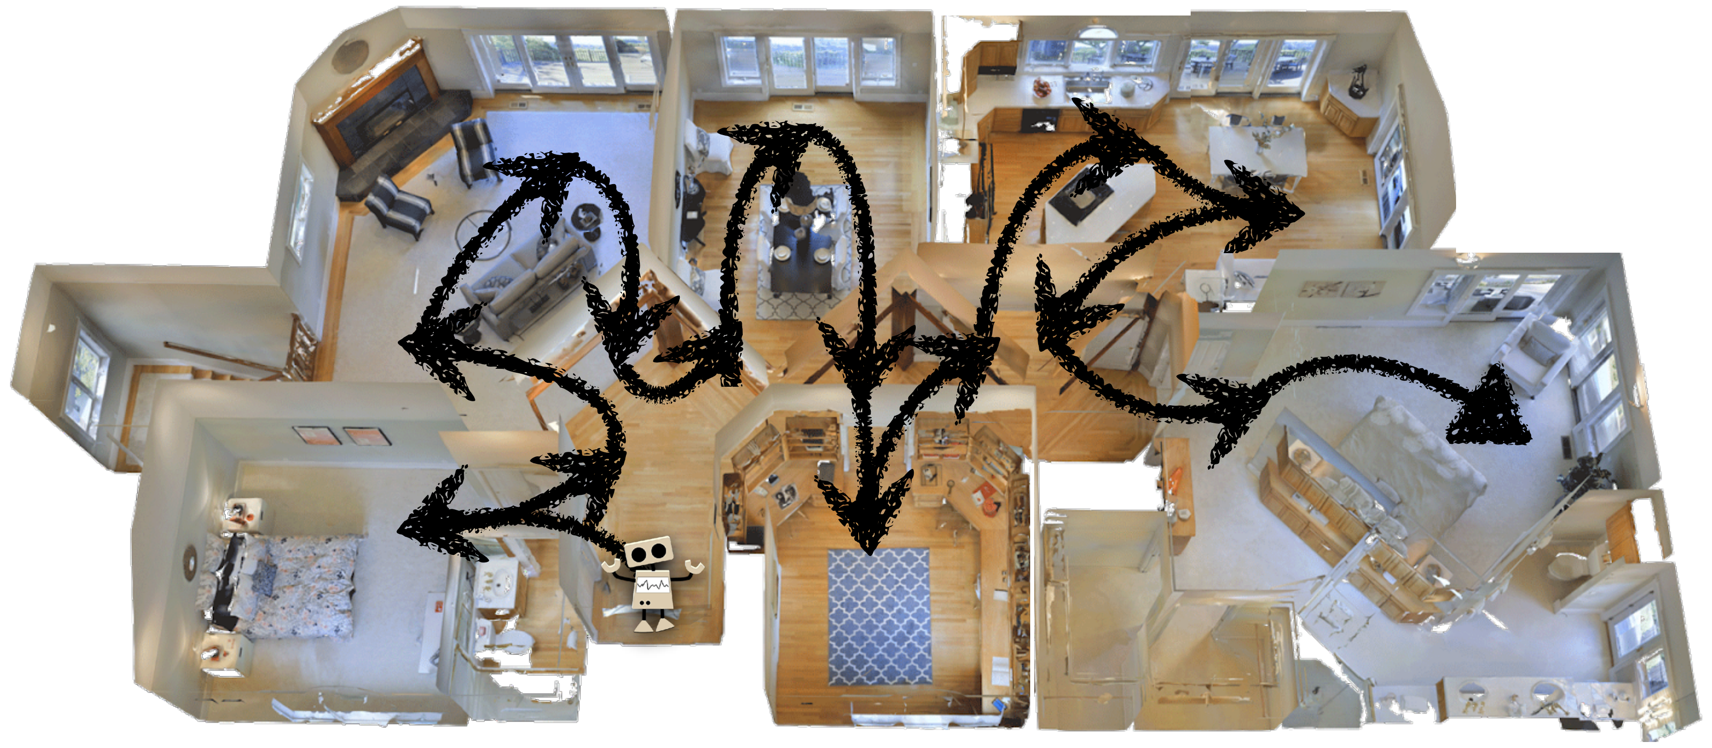

Exploration

- How to efficiently explore an unseen environment?
  - Memory/Mapping: Where have you been?
  - State/Pose Estimation: Where are you now?
  - Planning: Where do you need to go?

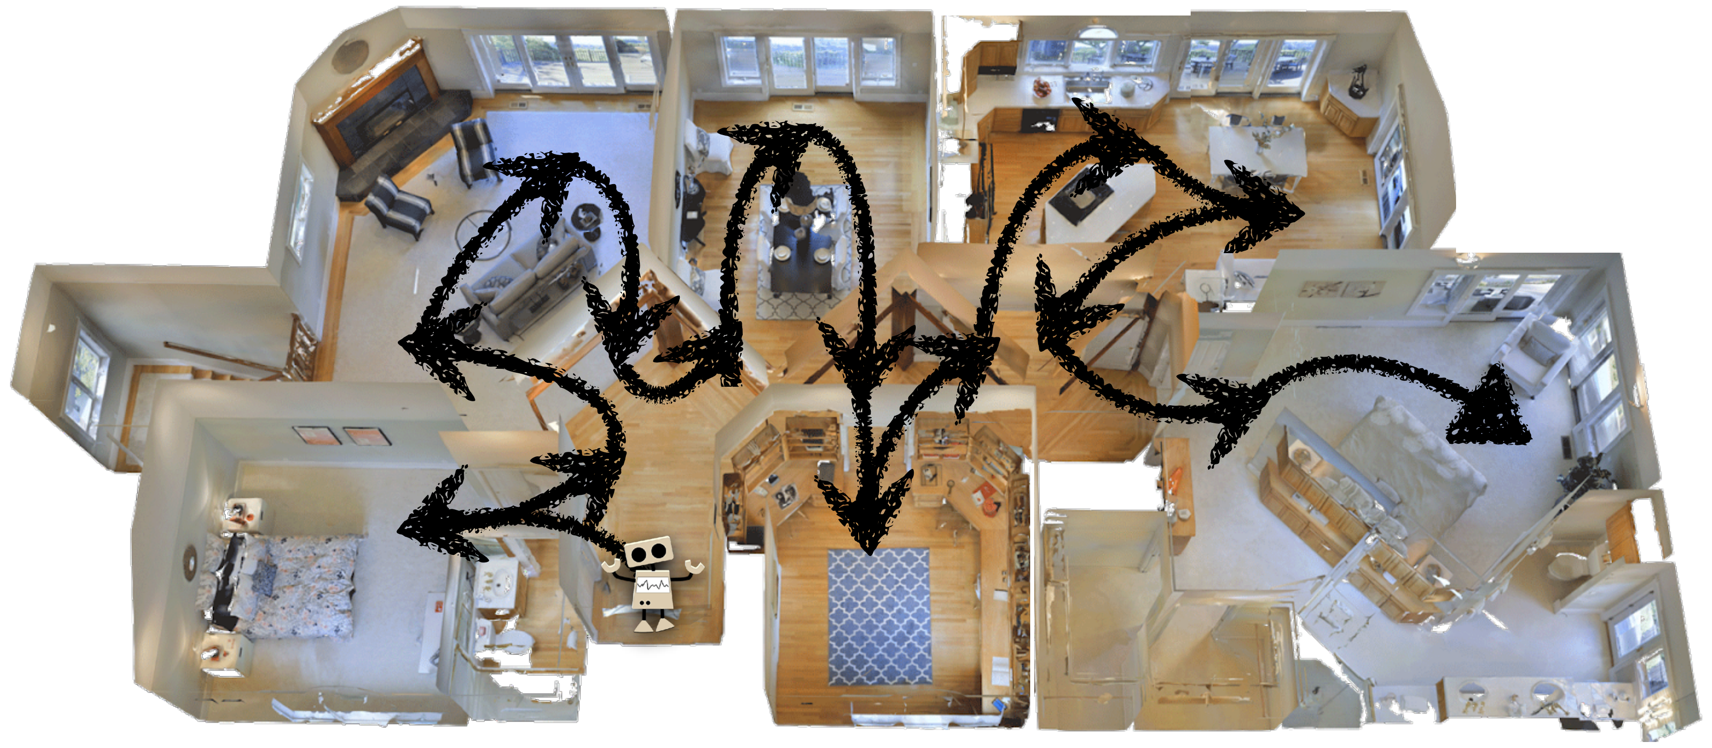

Exploration

- How to efficiently explore an unseen environment?
  - Memory/Mapping: Where have you been?
  - State/Pose Estimation: Where are you now?
  - Planning: Where do you need to go?

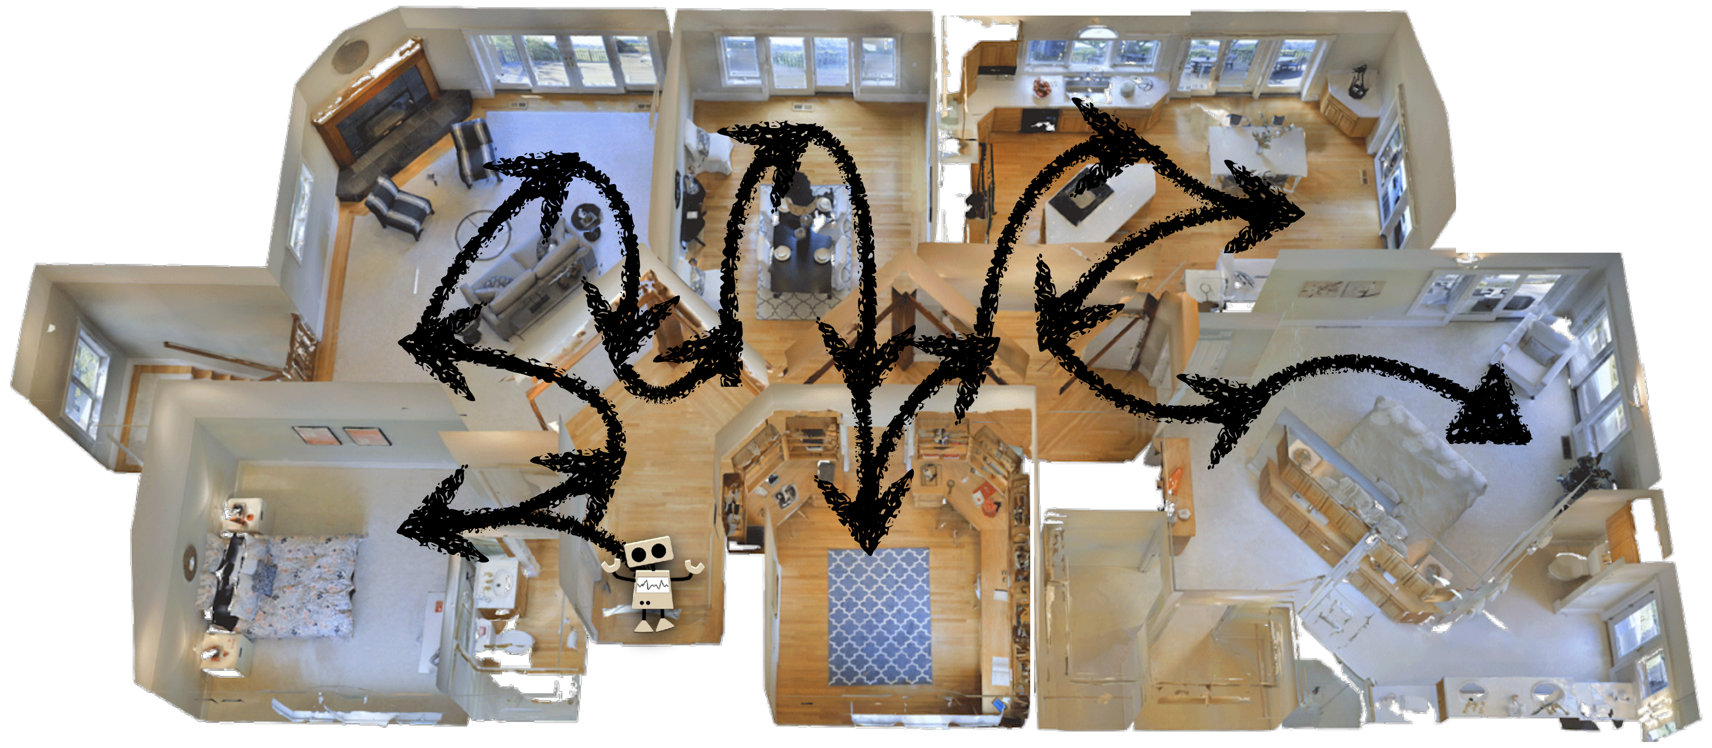

- Limitations of end-to-end RL:
  - High sample complexity
  - Ineffective in large environments

Exploration

- How to efficiently explore an unseen environment?
  - Memory/Mapping: Where have you been?
  - State/Pose Estimation: Where are you now?
  - Planning: Where do you need to go?

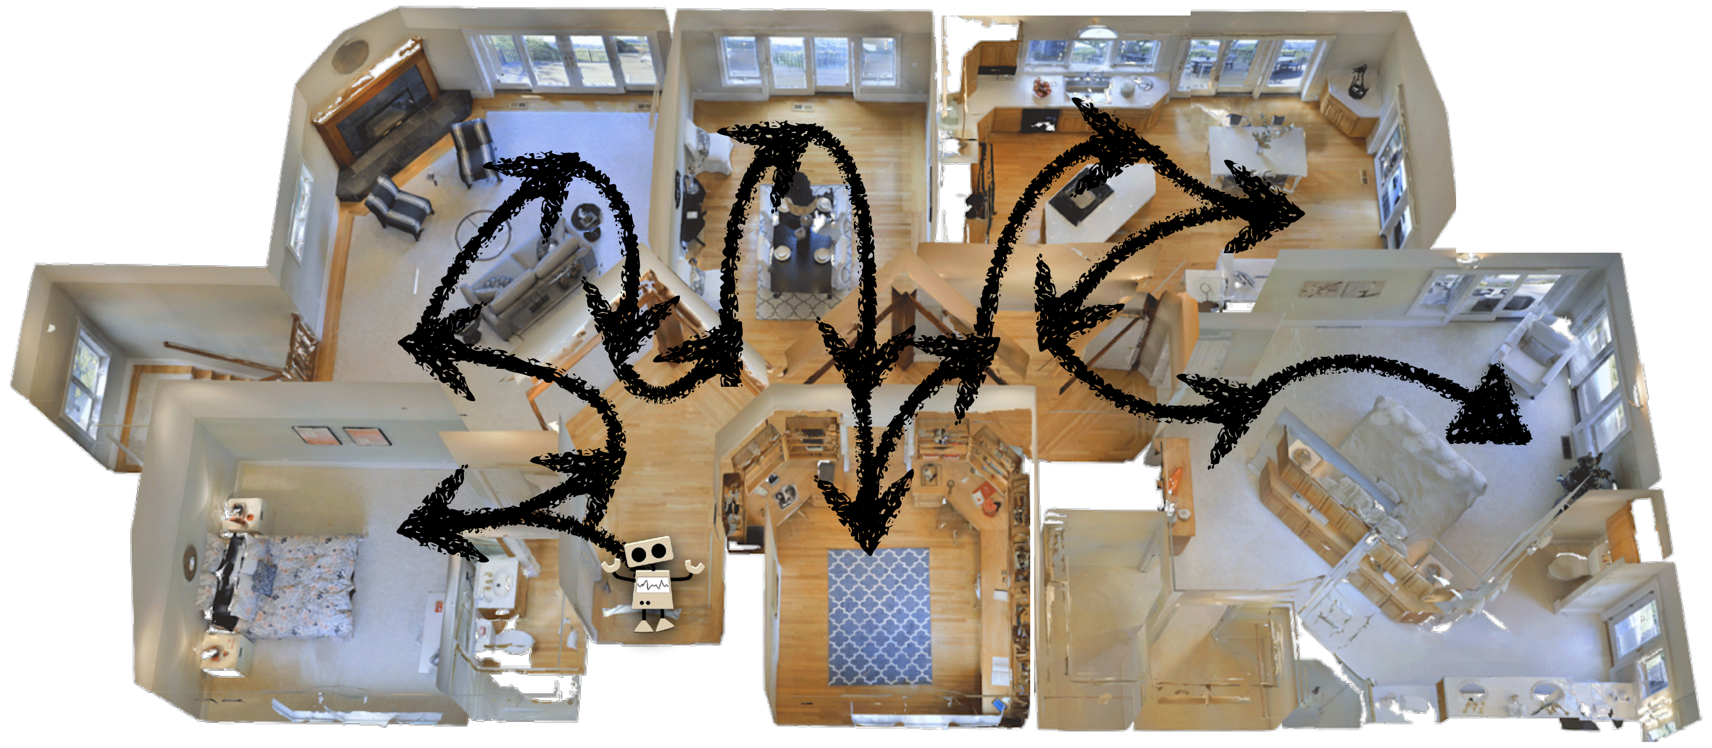

- Limitations of end-to-end RL:
  - High sample complexity
  - Ineffective in large environments

- Our solution: Active Neural SLAM
  - Structured spatial representations
  - Hierarchical policies
  - Analytical planners

Active Neural SLAM: Overview

Active Neural SLAM: Overview

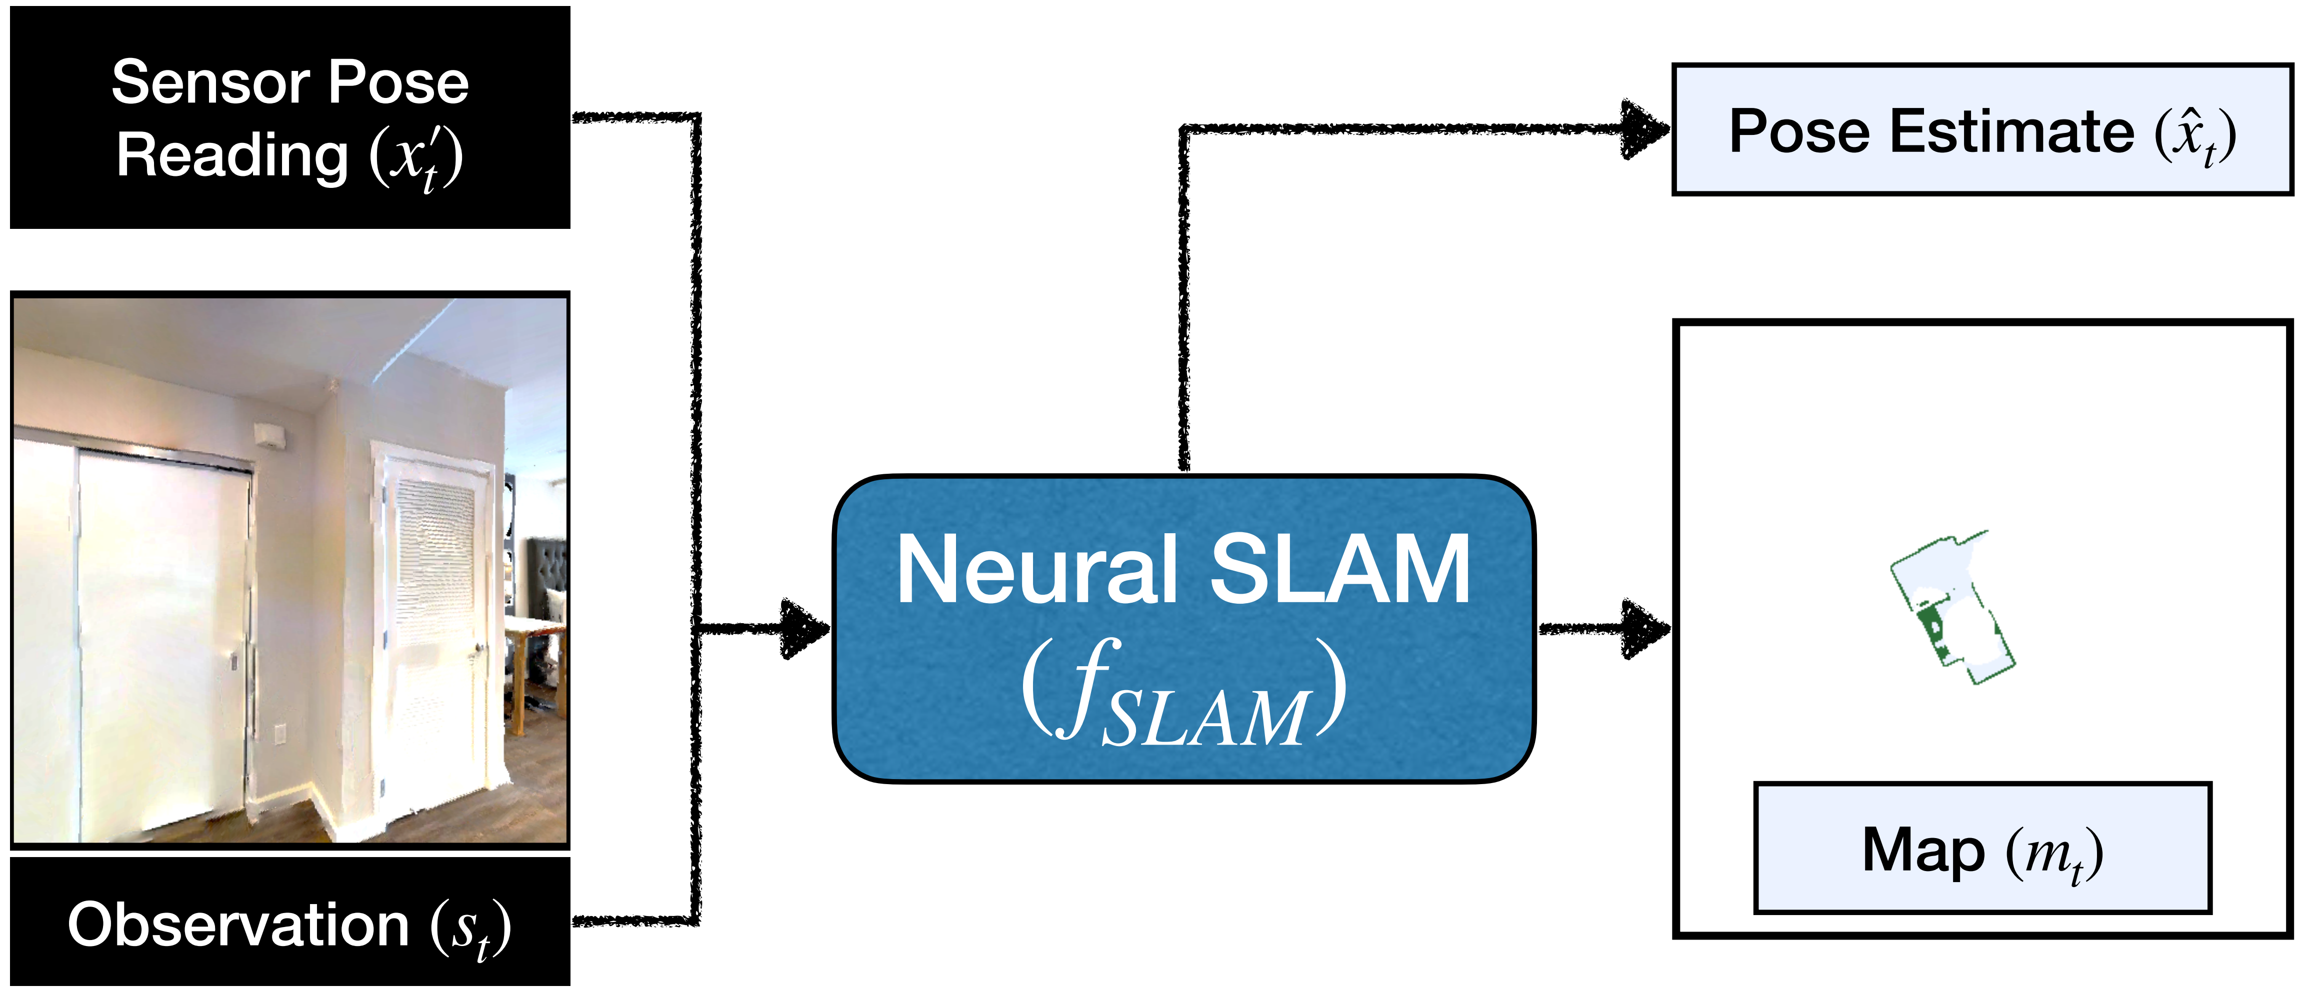

Active Neural SLAM: Overview

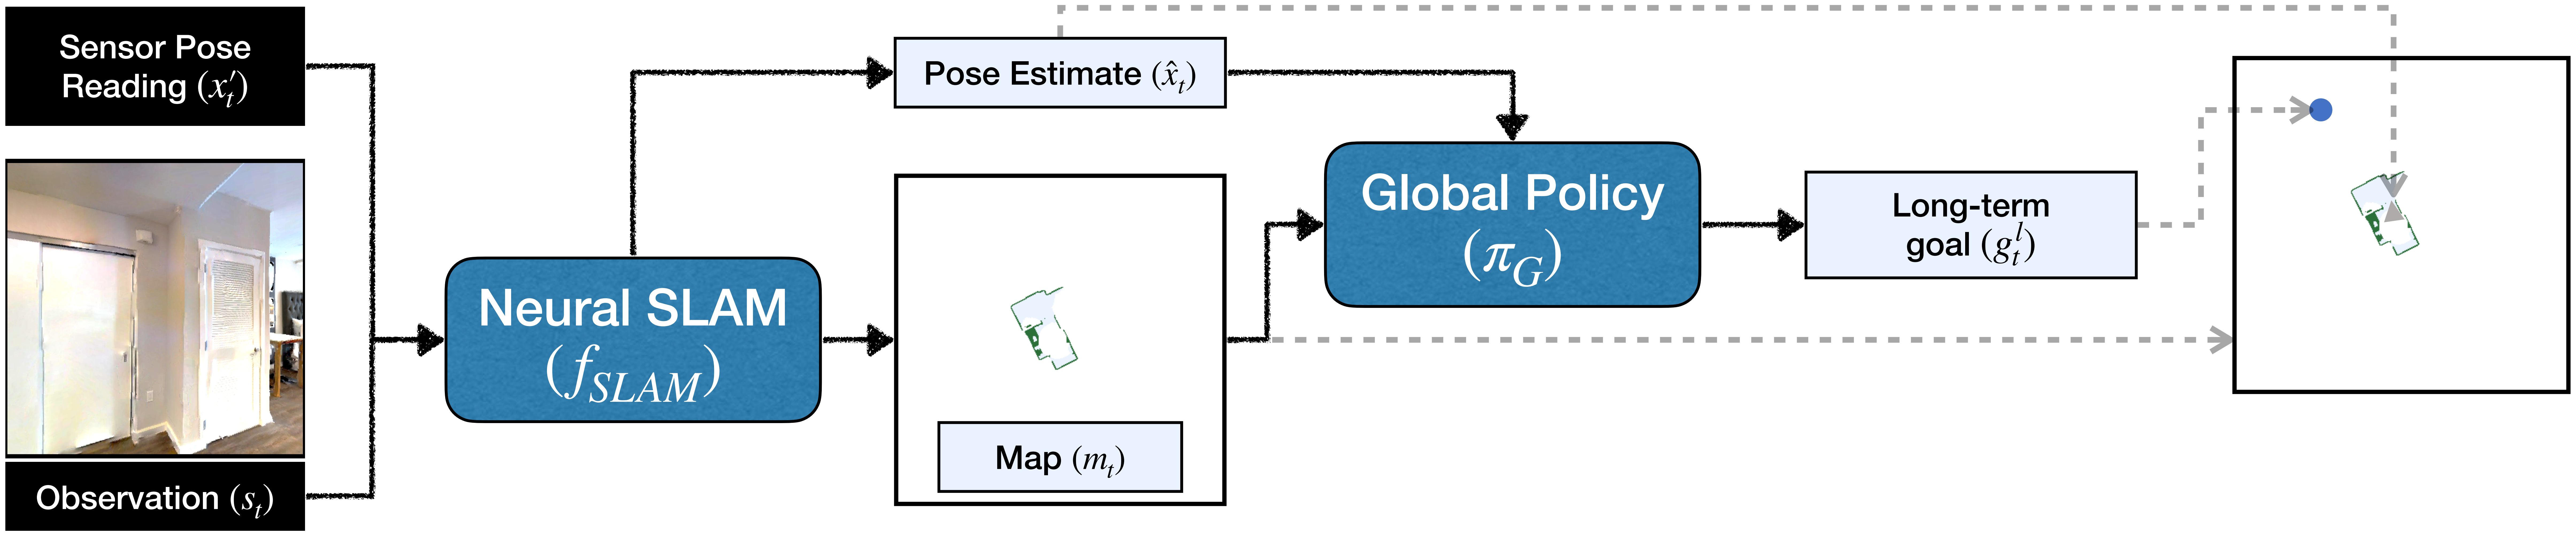

Active Neural SLAM: Overview

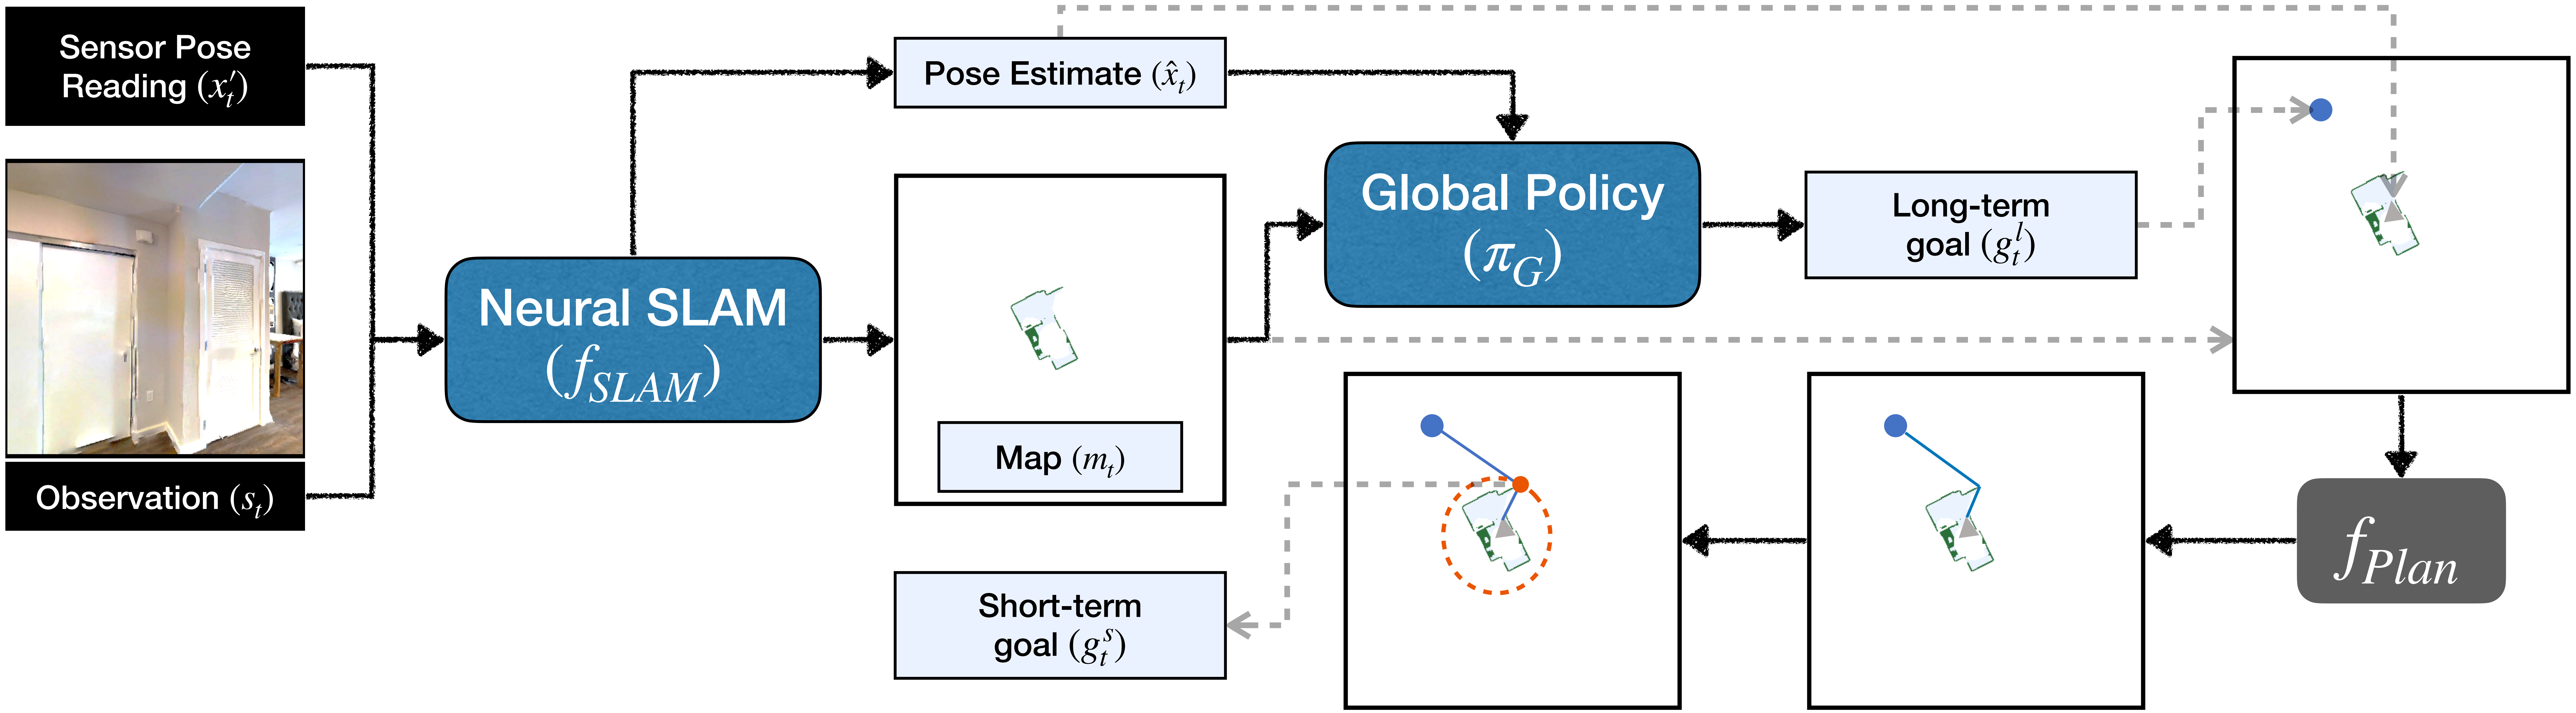

Active Neural SLAM: Overview

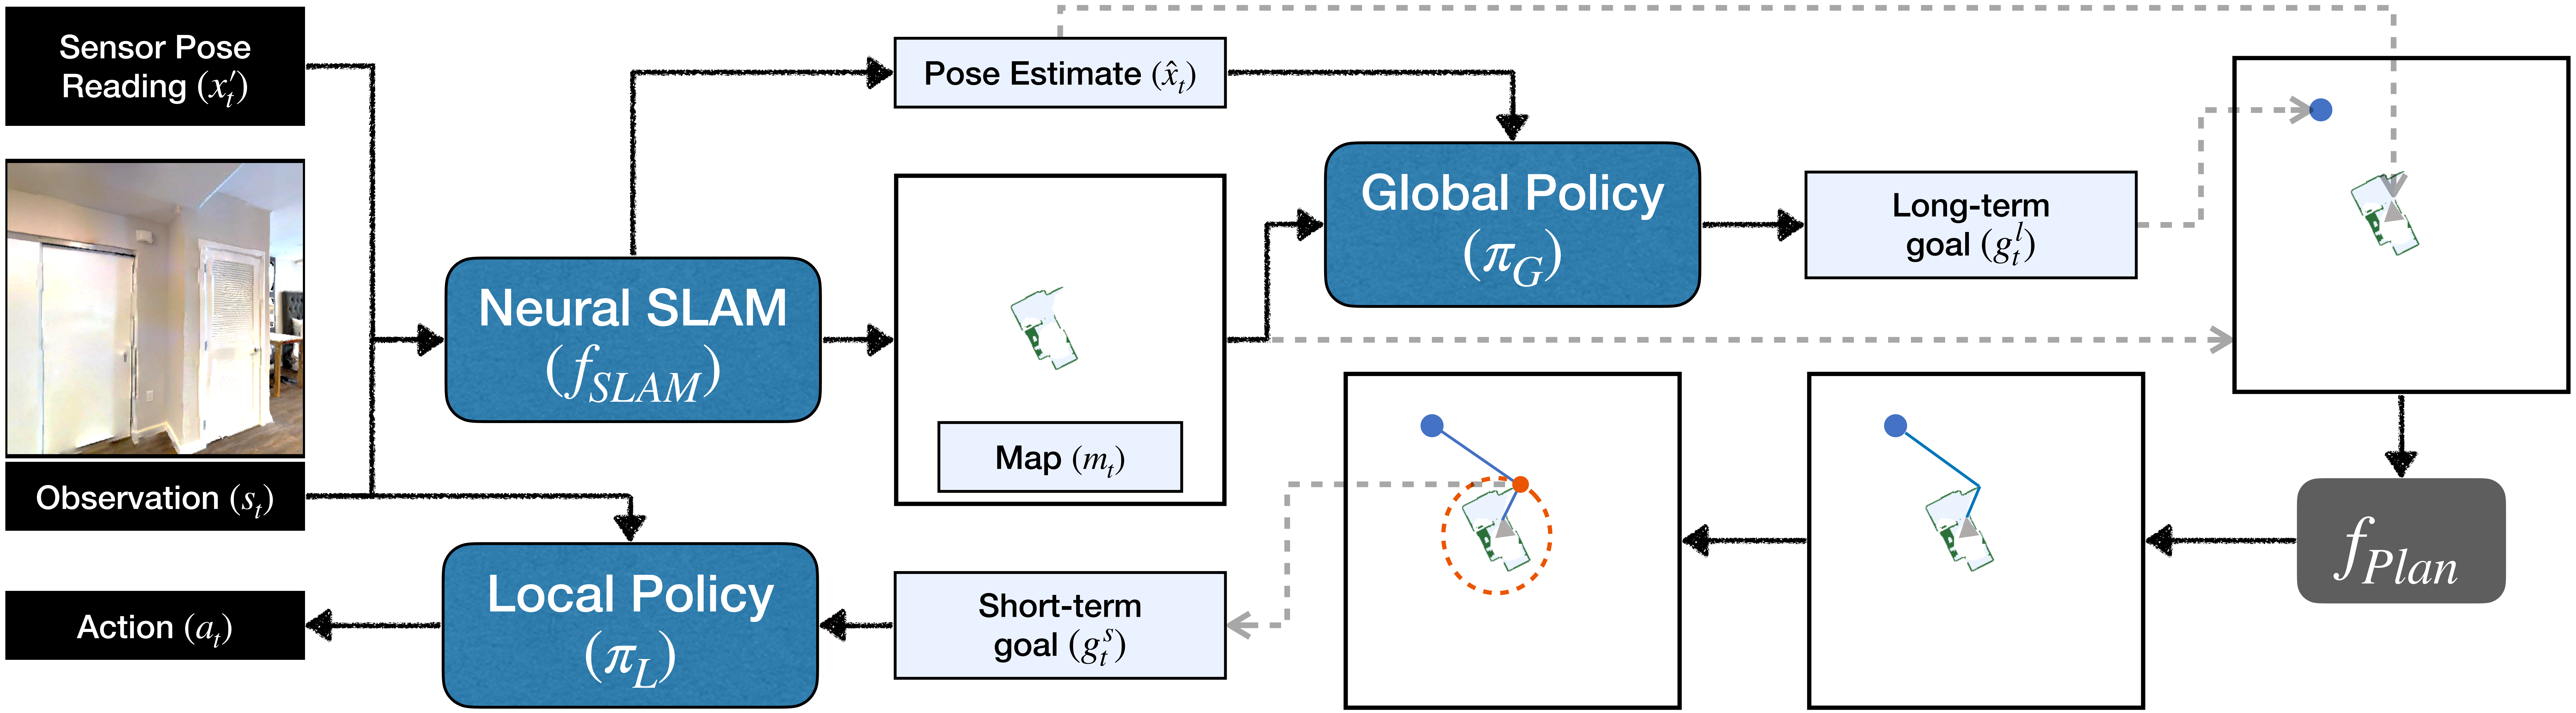

Neural SLAM Module

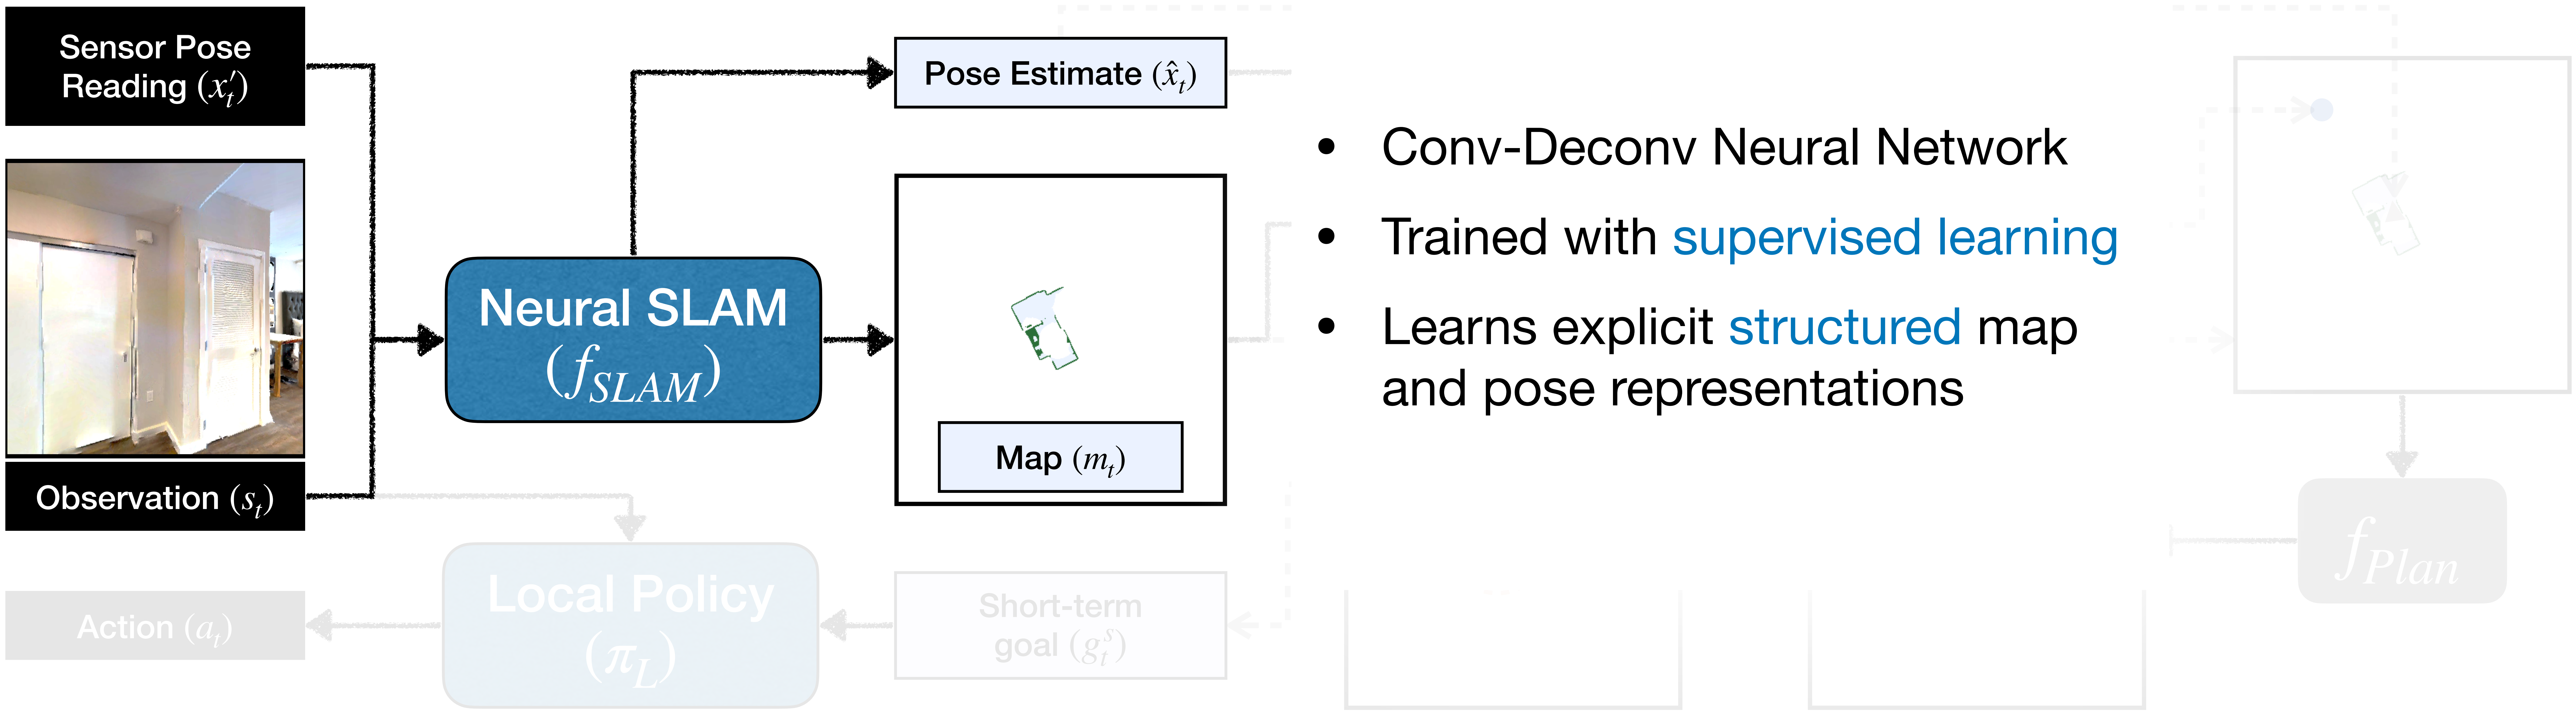

- Conv-Deconv Neural Network
- Trained with supervised learning
- Learns explicit structured map and pose representations

Global Policy

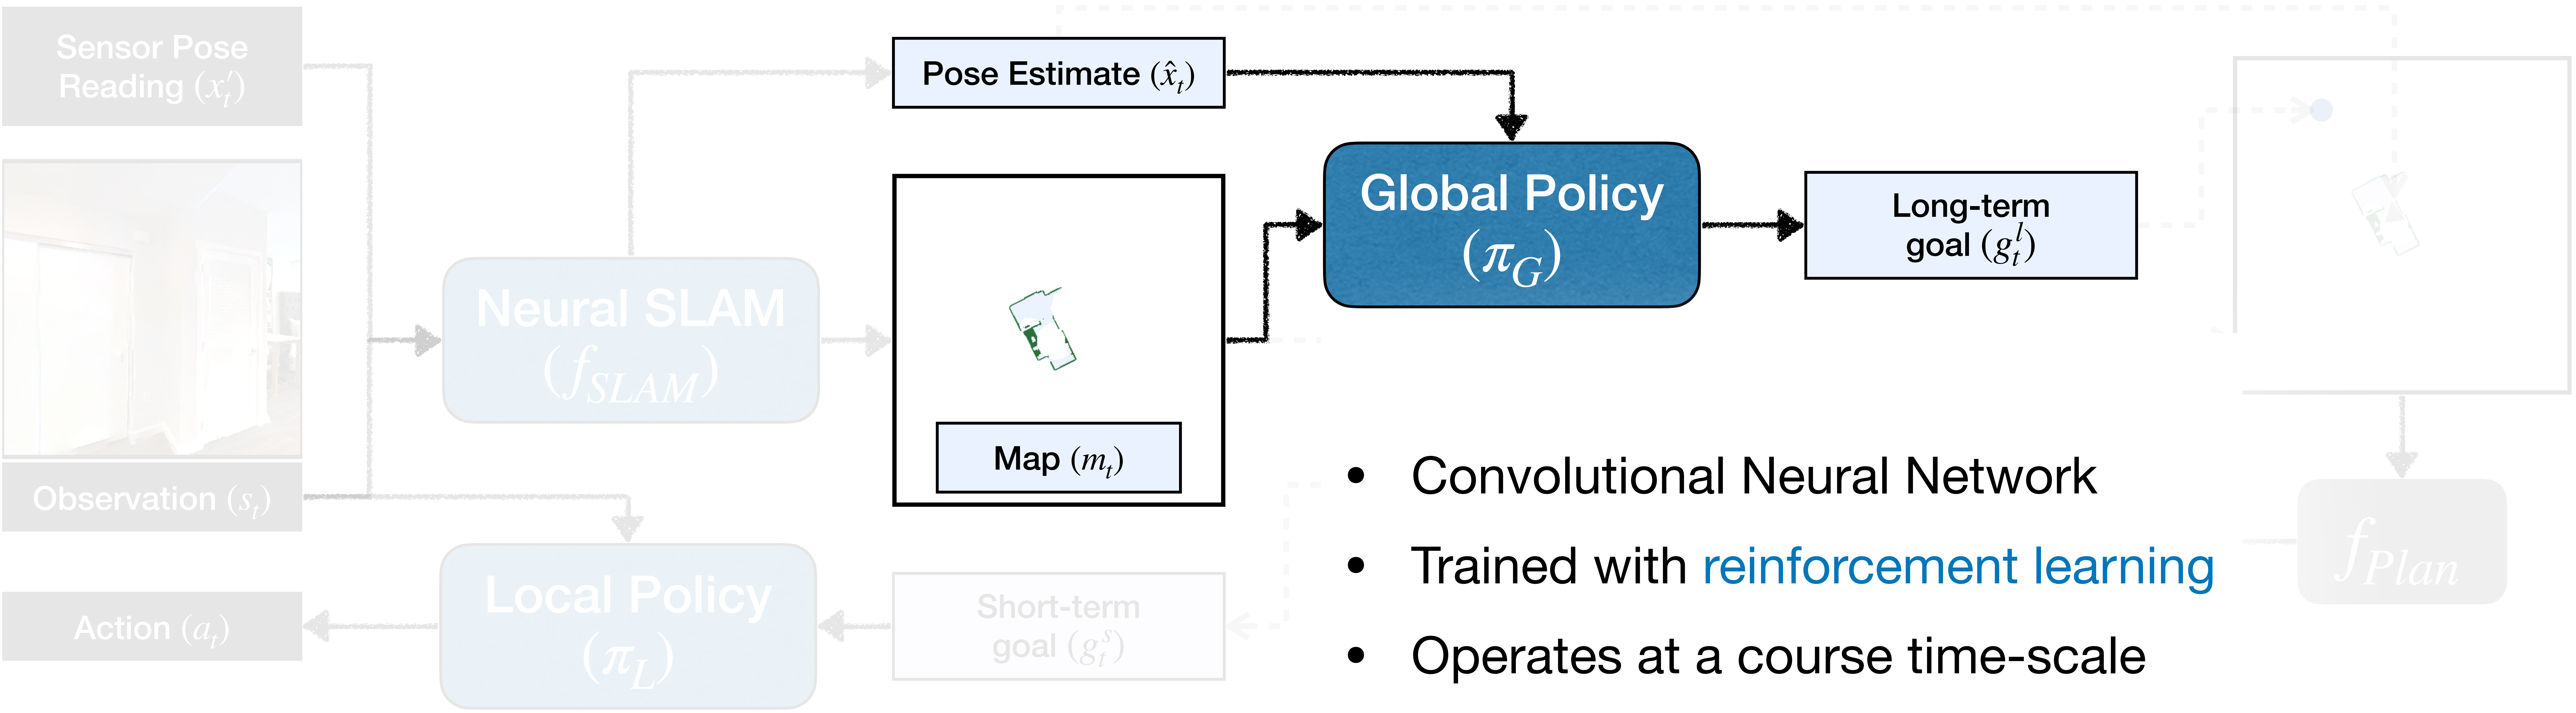

- Convolutional Neural Network
- Trained with reinforcement learning
- Operates at a course time-scale

Local Policy

- Convolutional Neural Network
- Trained with imitation learning
- Operates at a fine time-scale

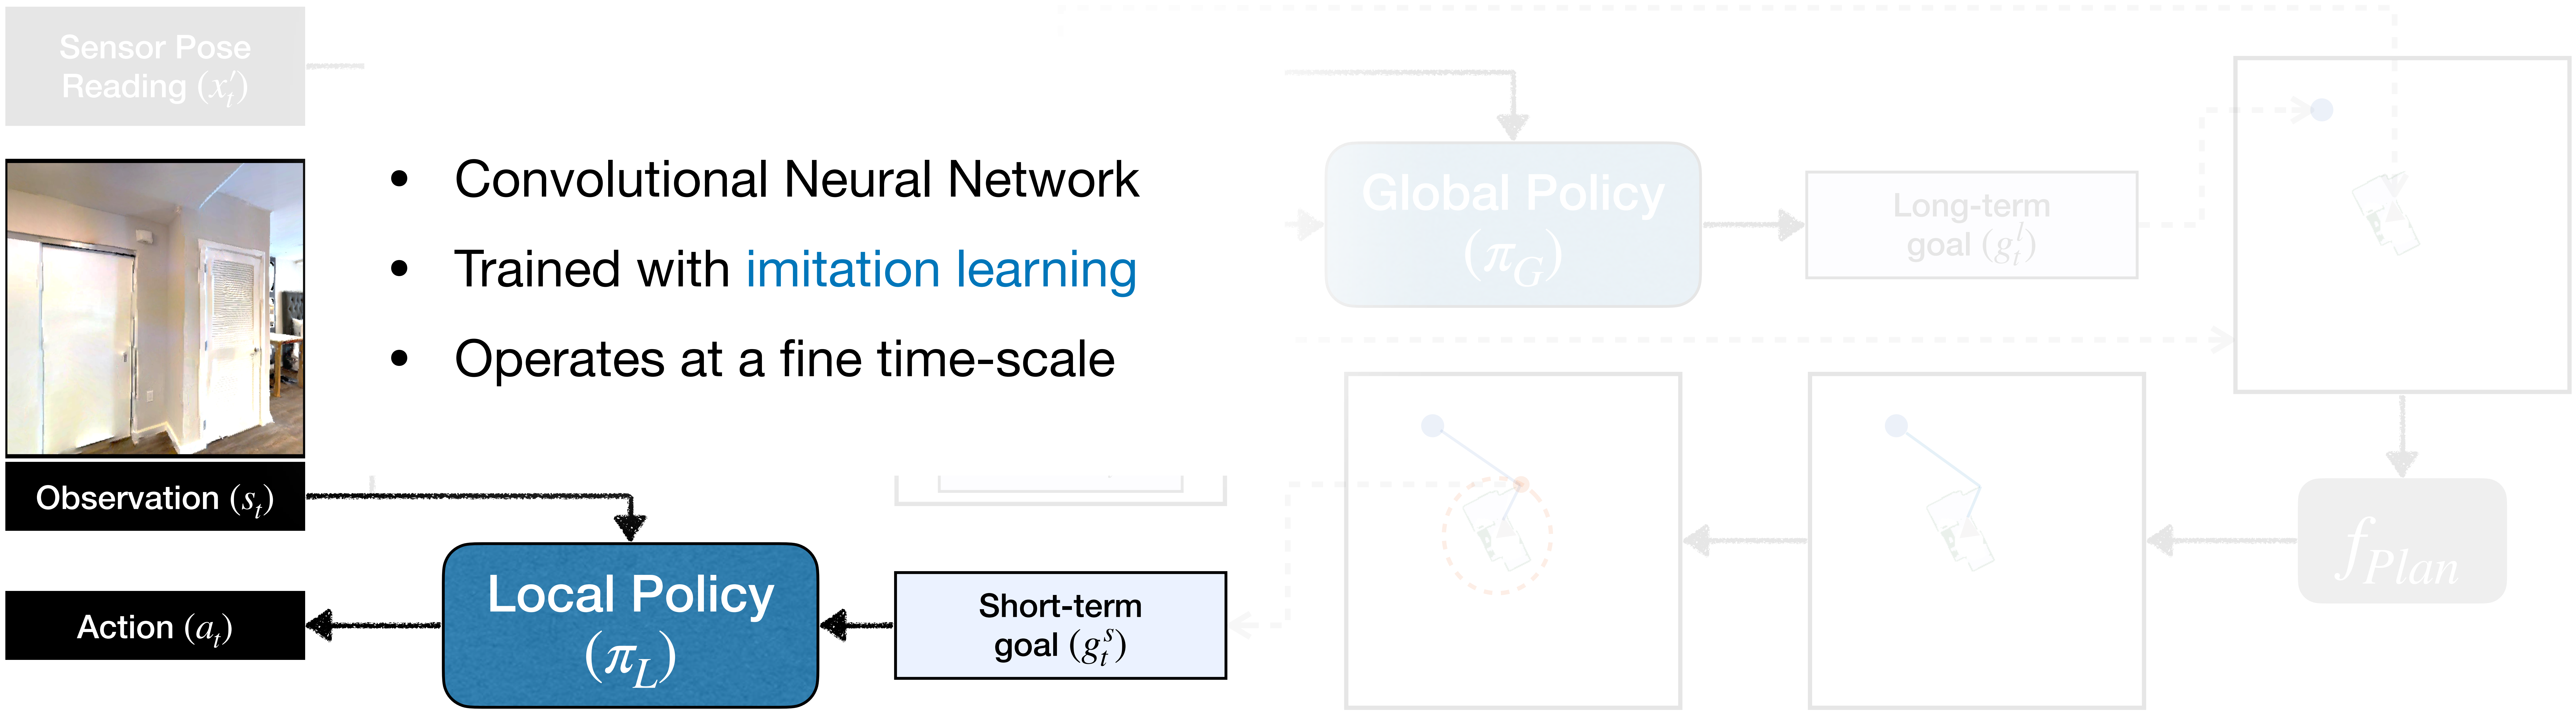

Neural SLAM Module

Neural SLAM Module

Sensor Pose Reading (x'_{t-1})

Observation (s_{t-1})

Sensor Pose Reading (x'_t)

Observation (s_t)

Neural SLAM Module

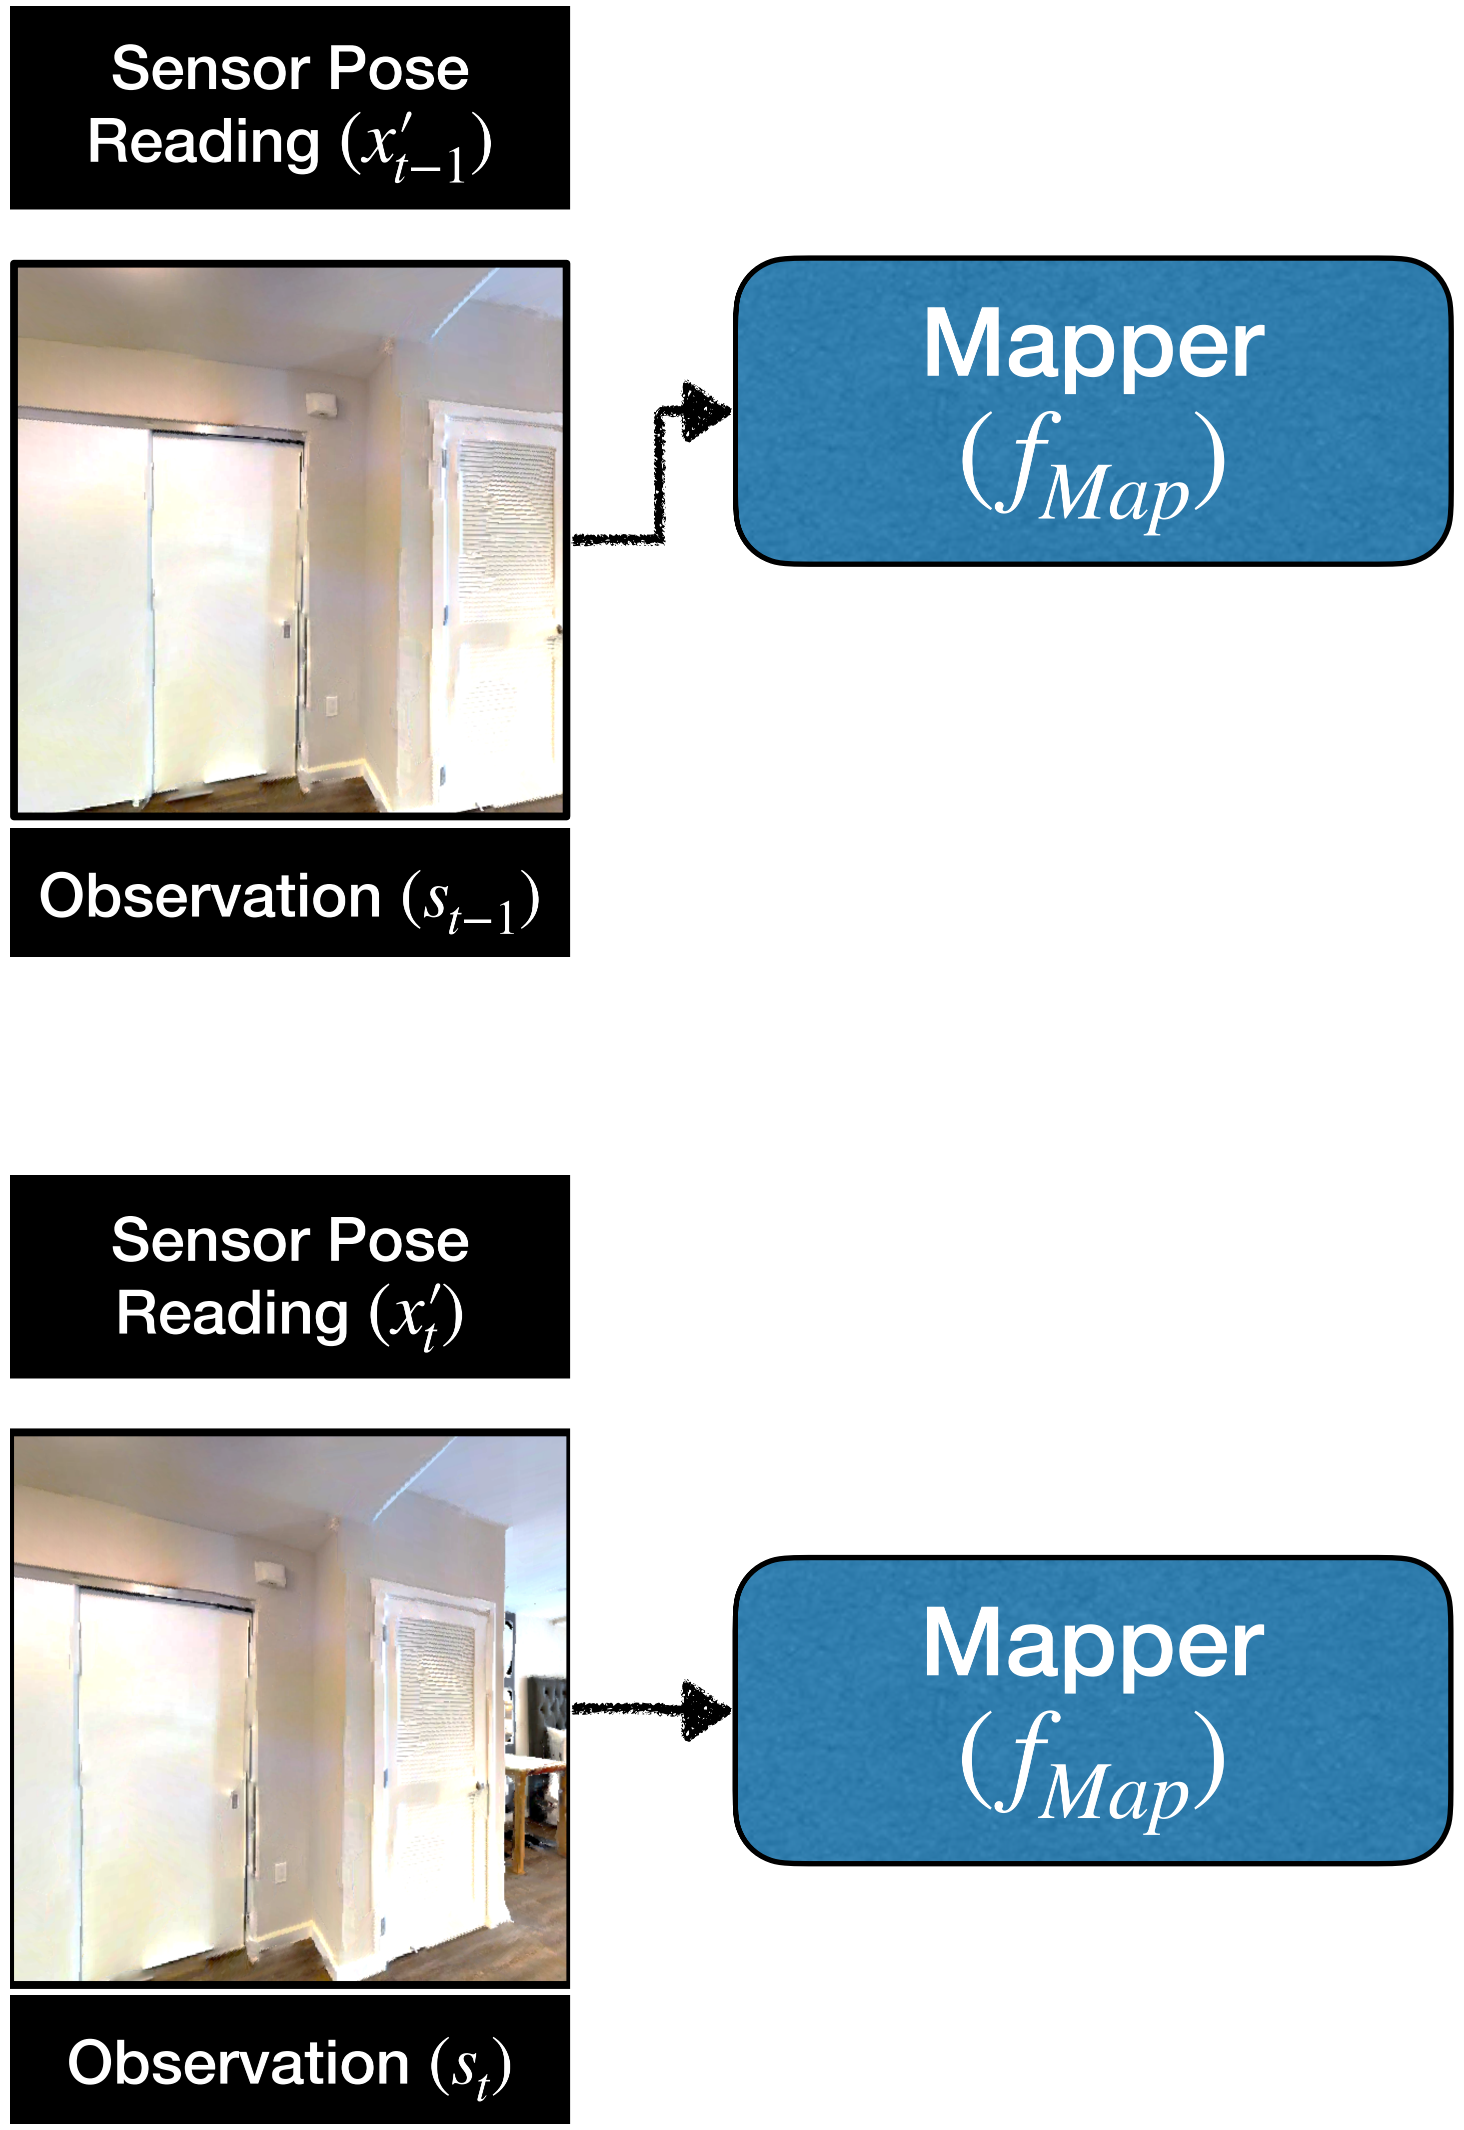

Neural SLAM Module

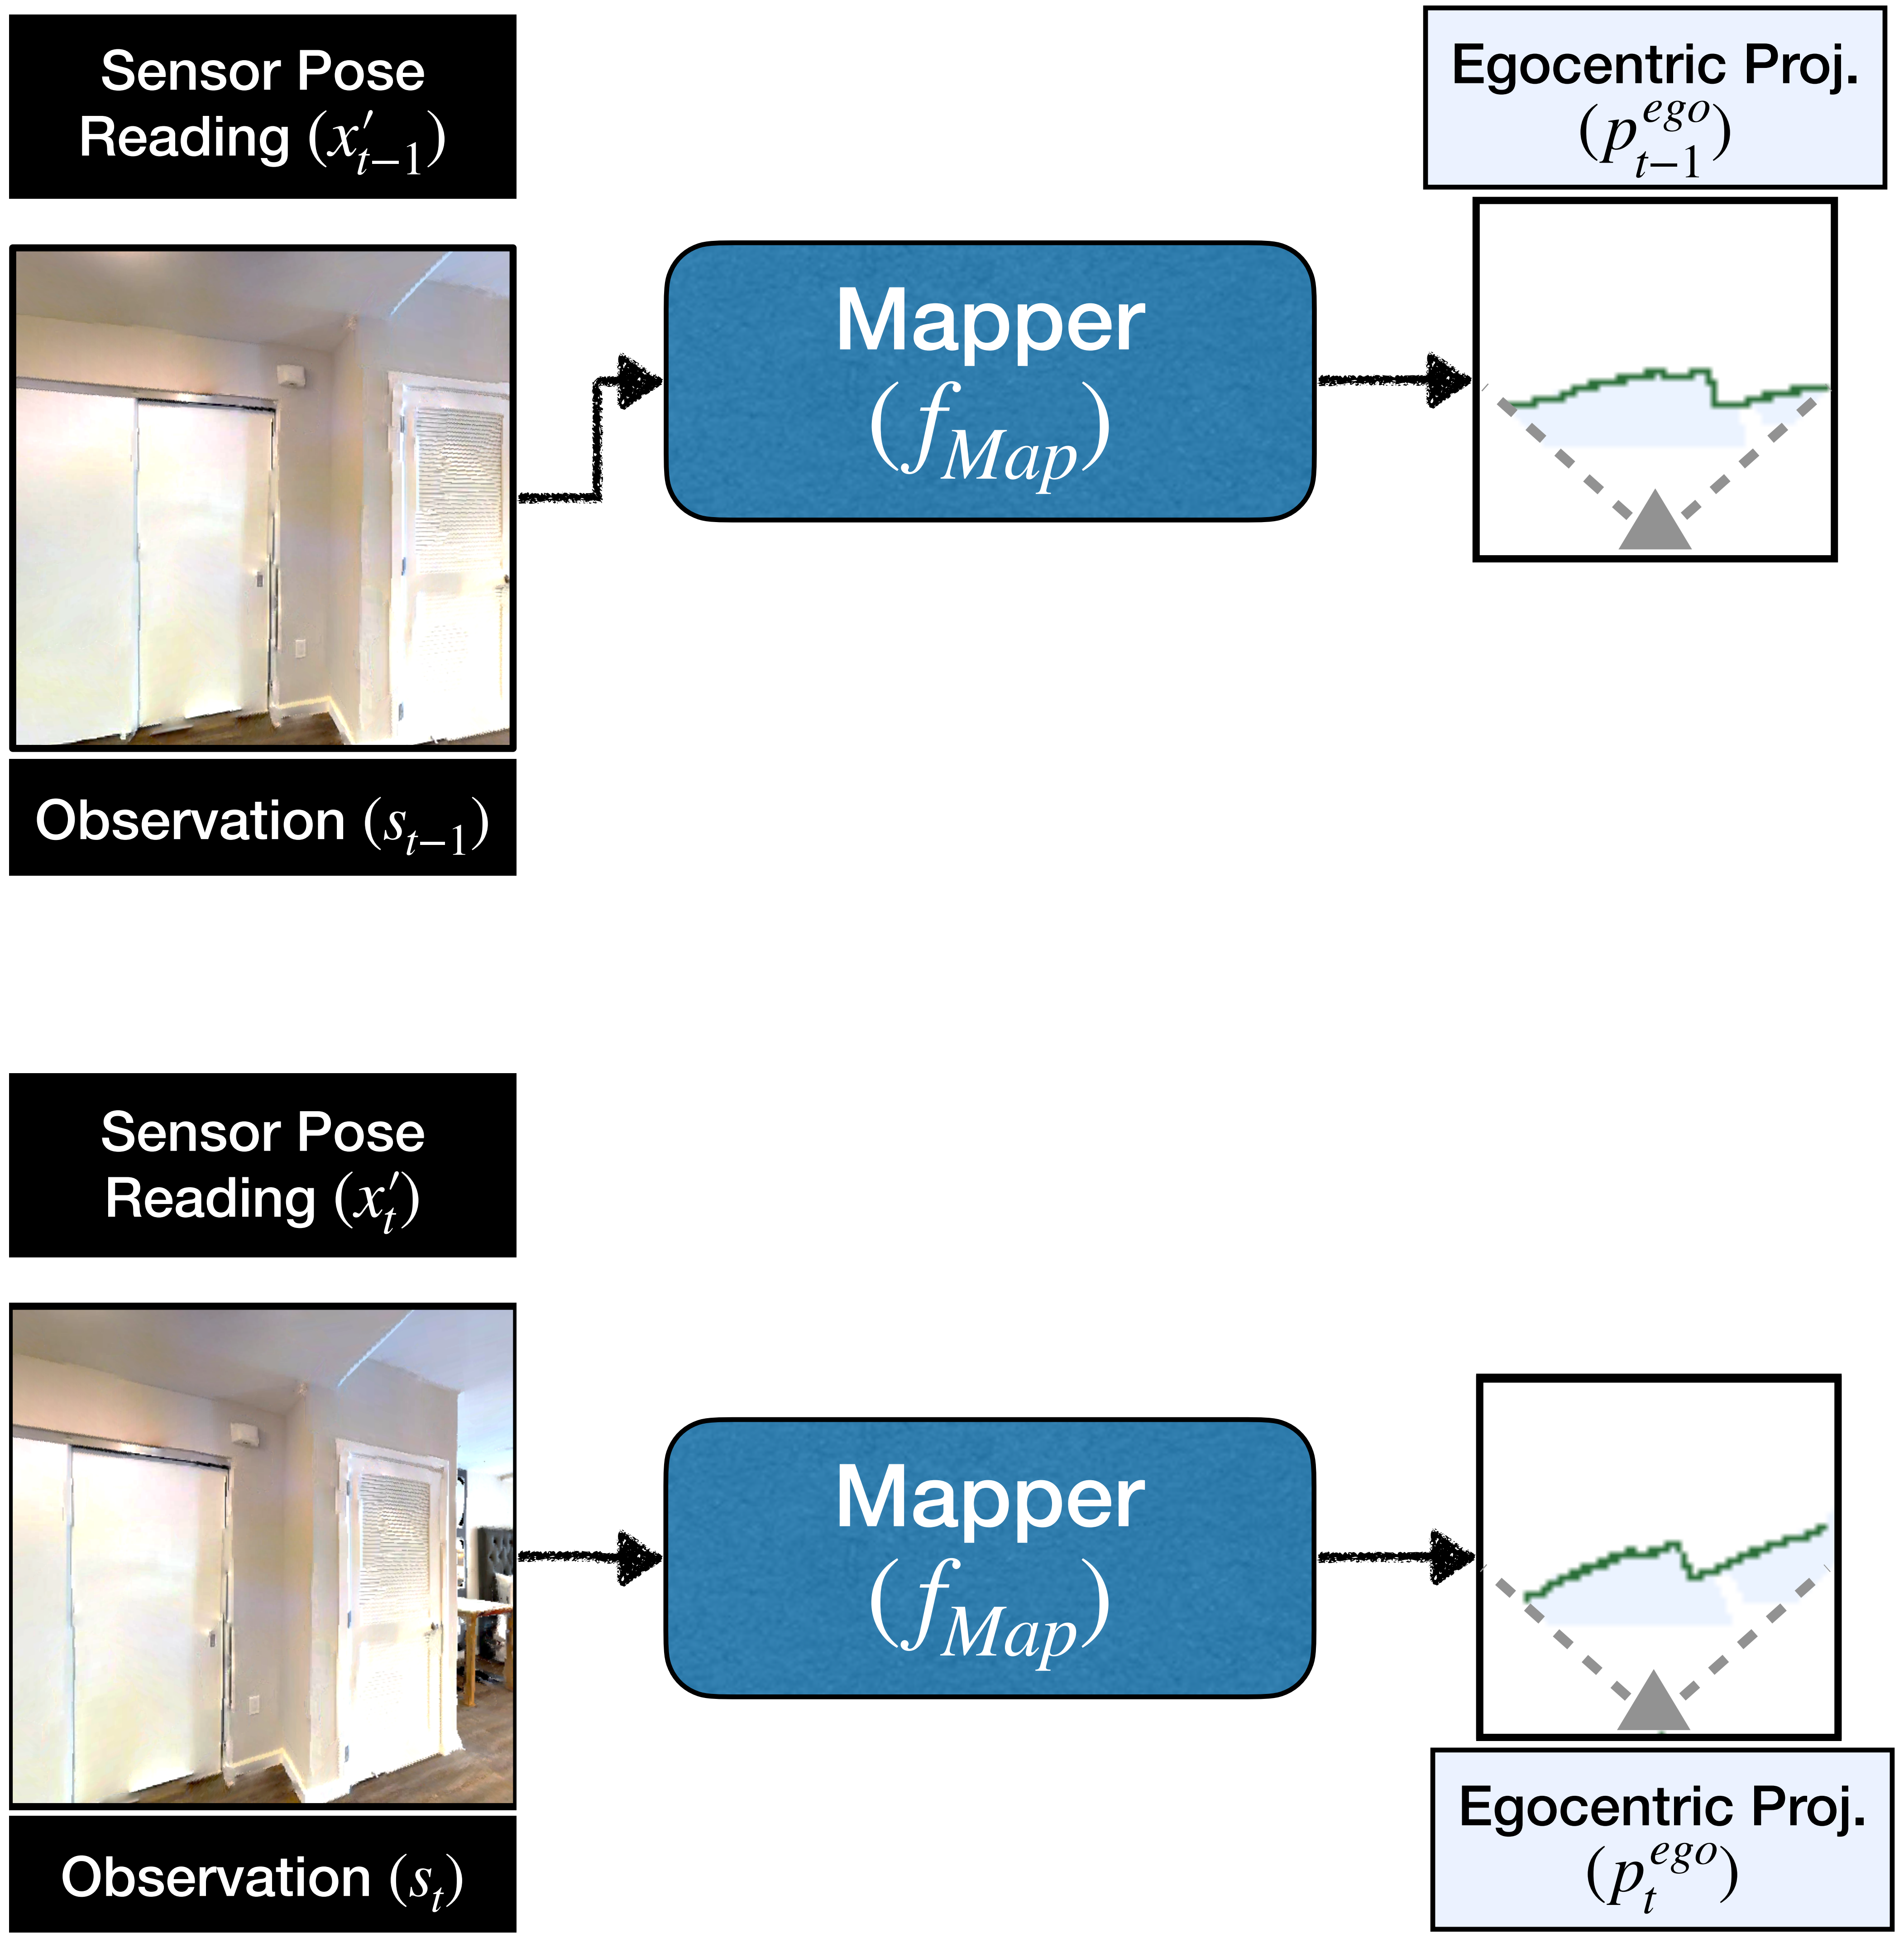

Neural SLAM Module

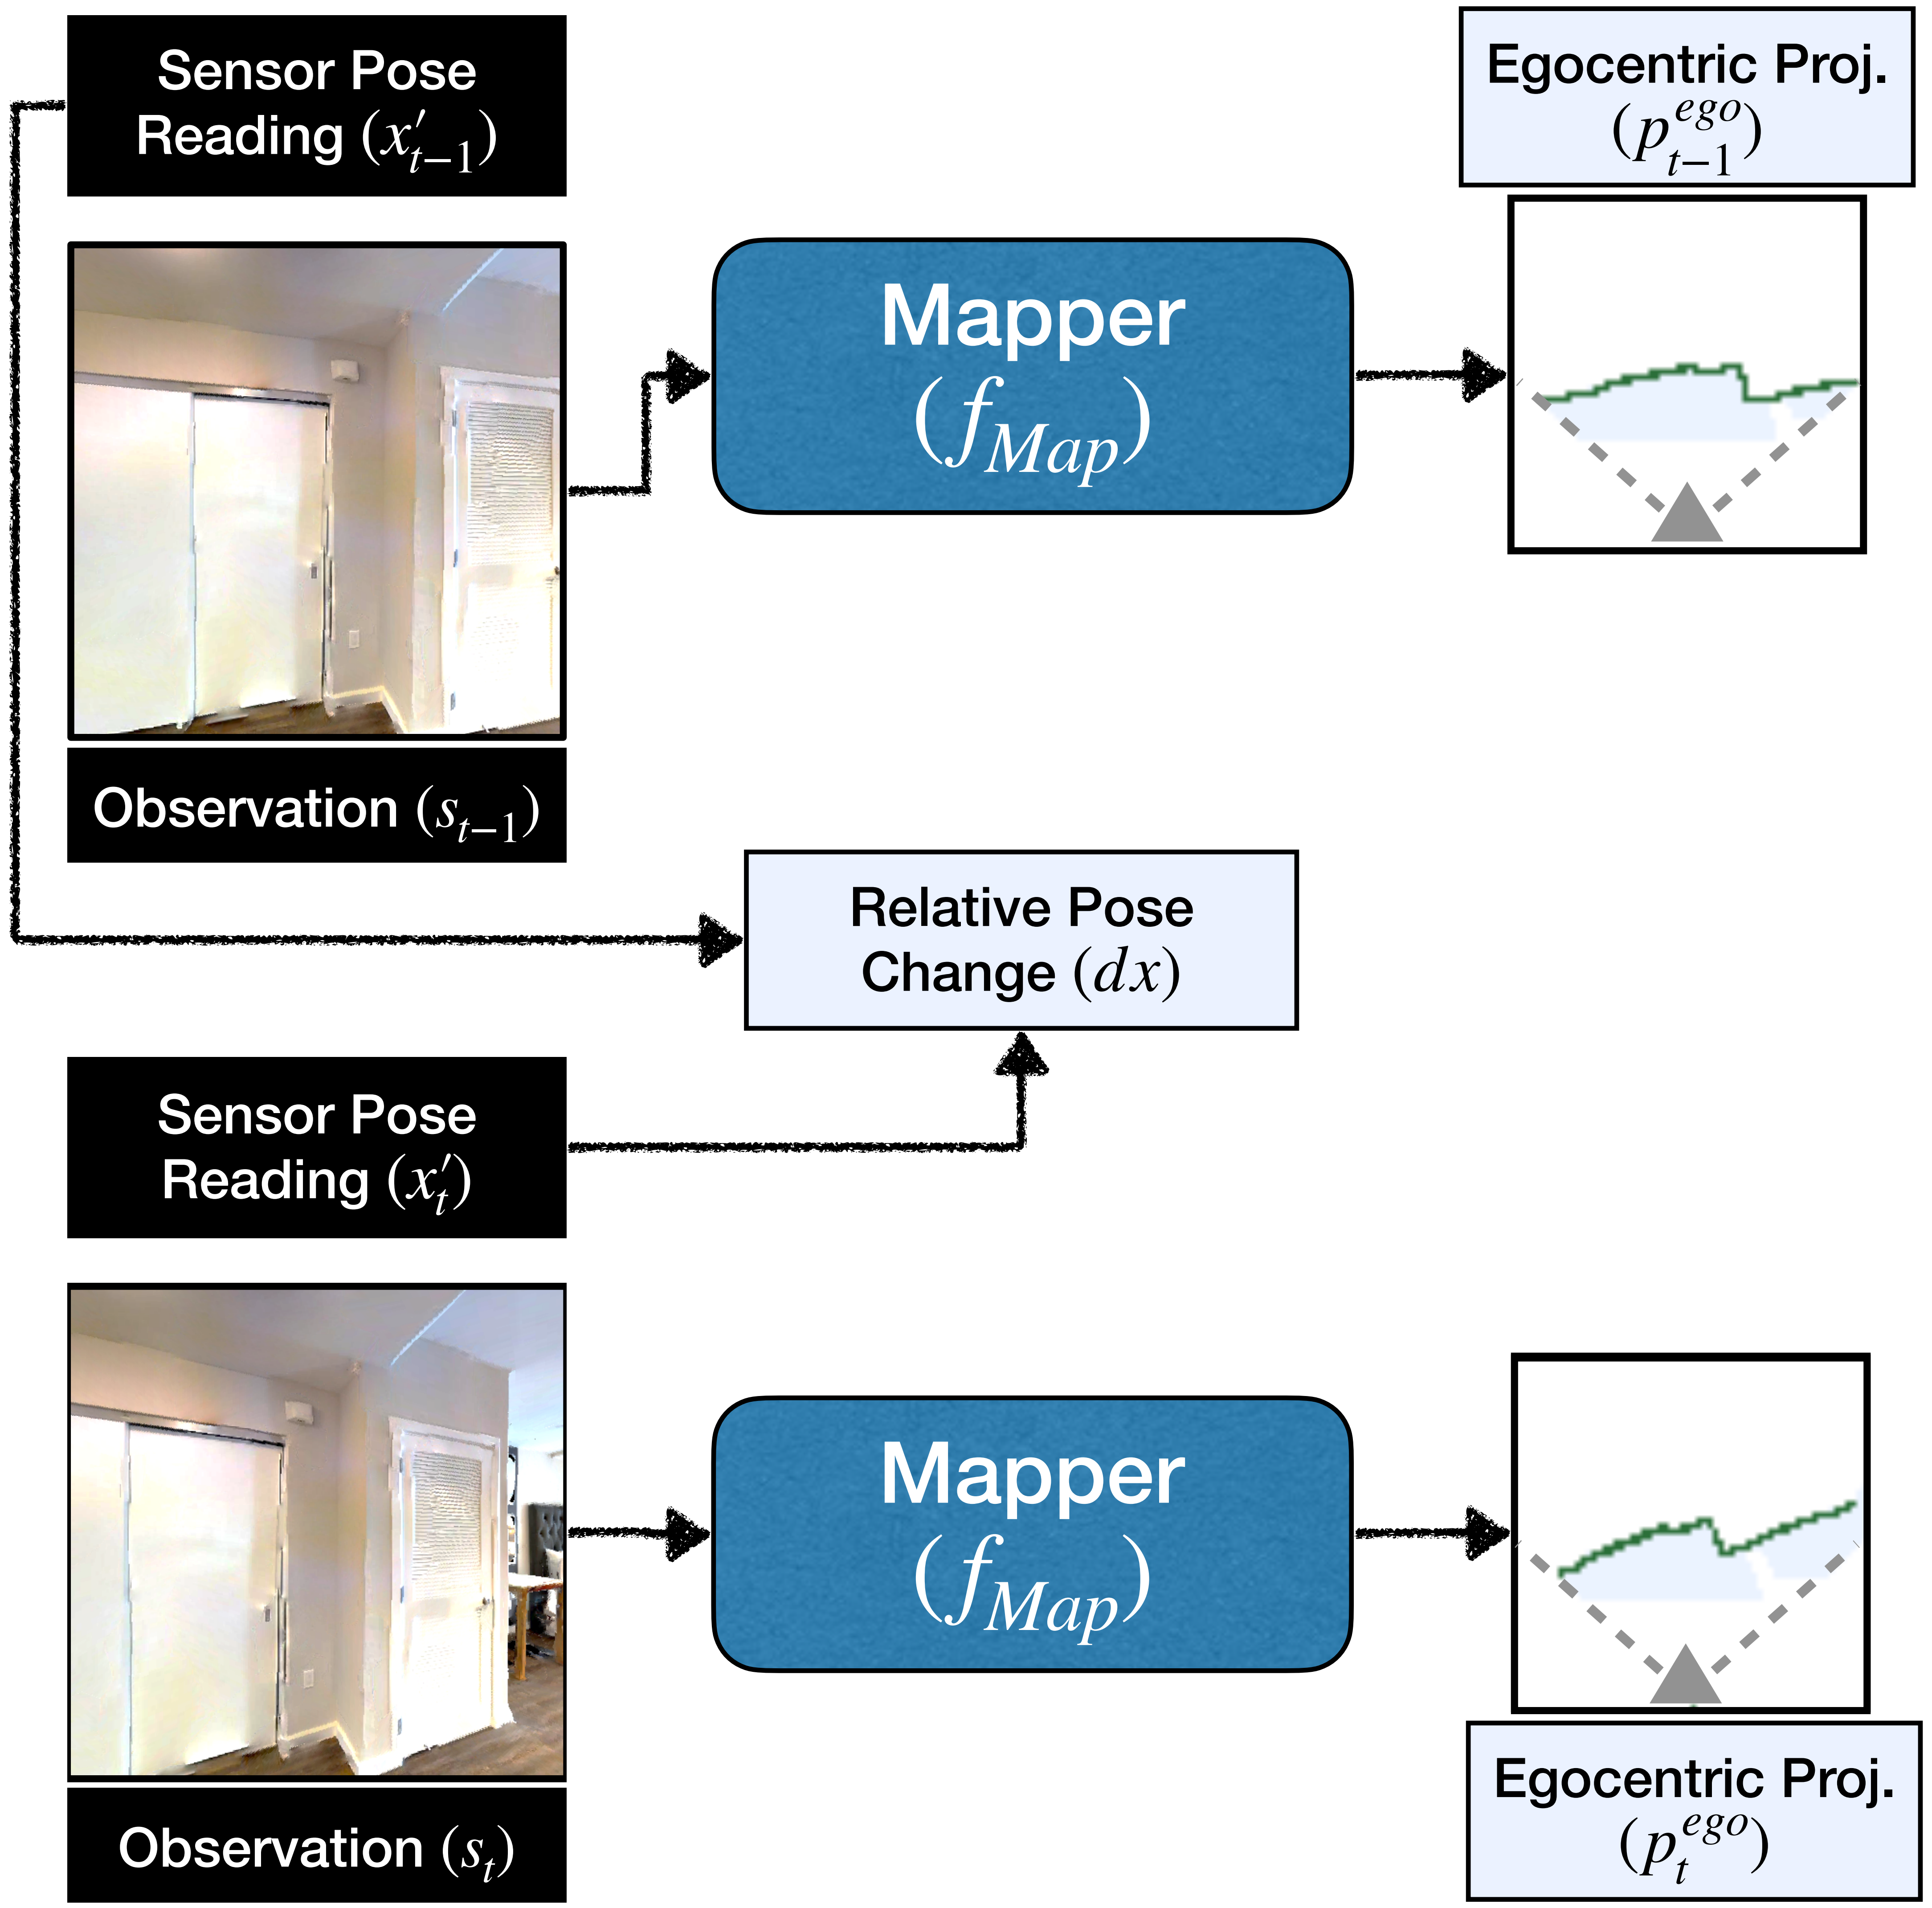

Neural SLAM Module

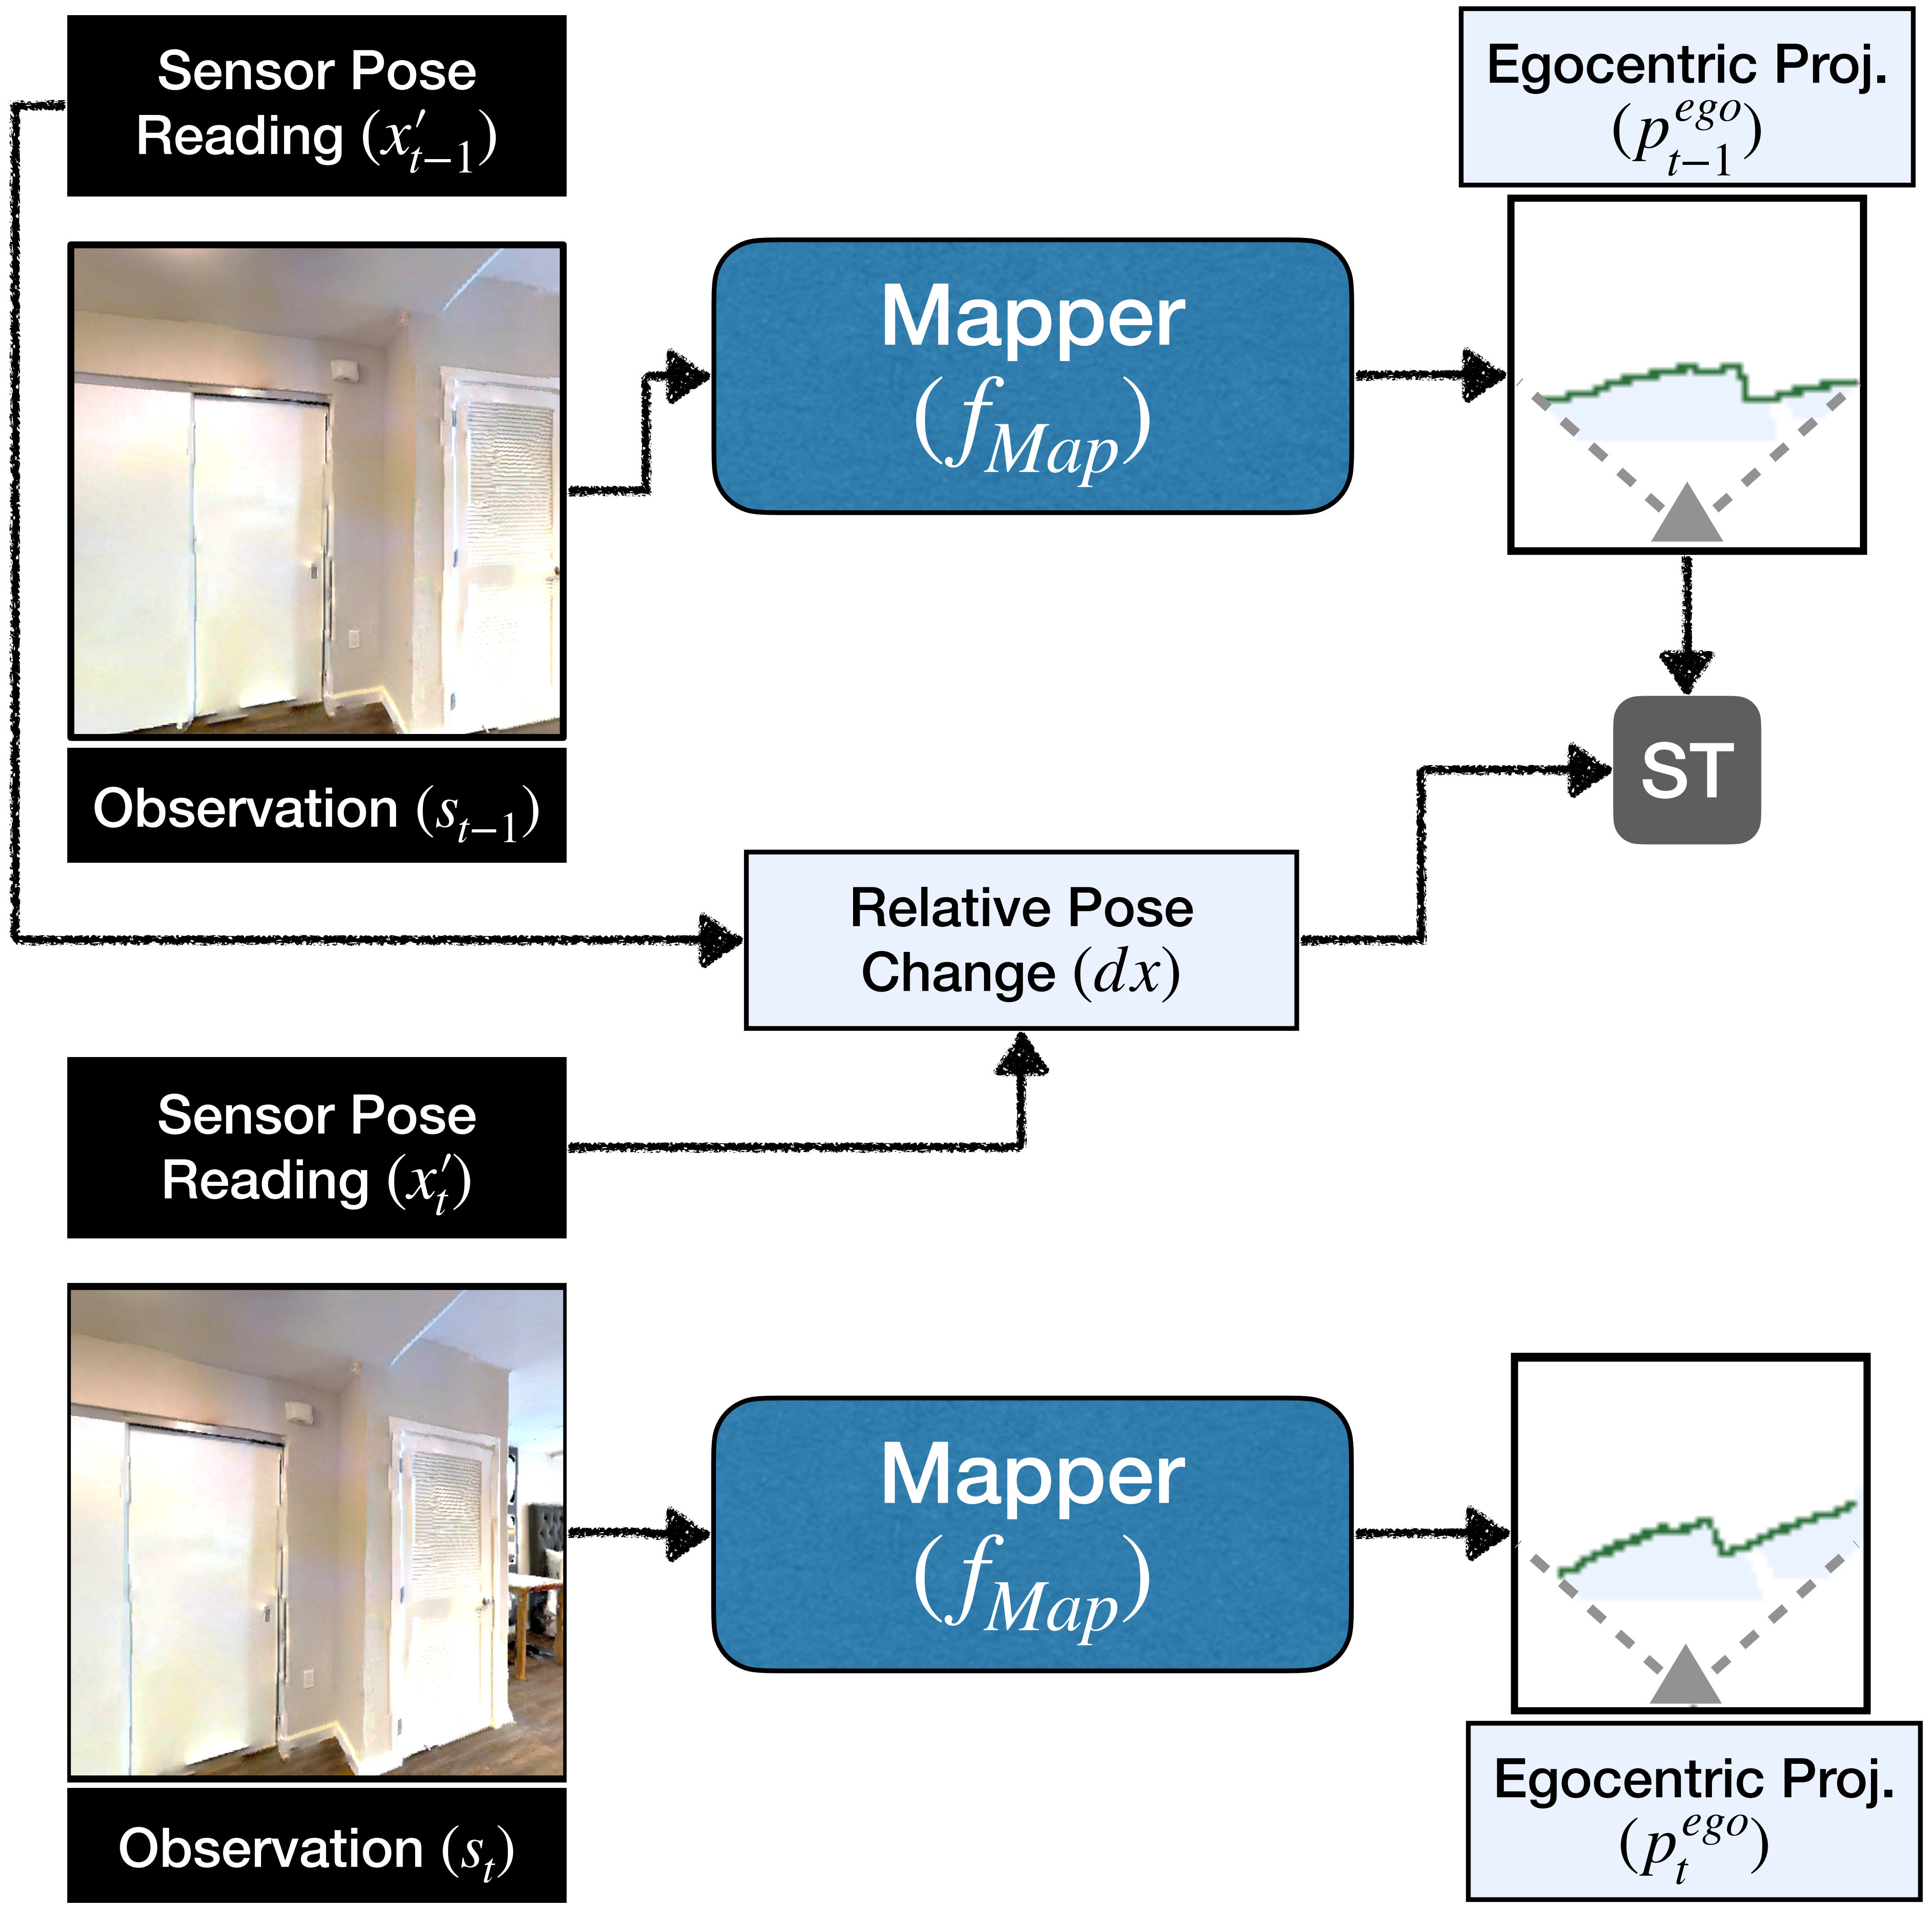

Neural SLAM Module

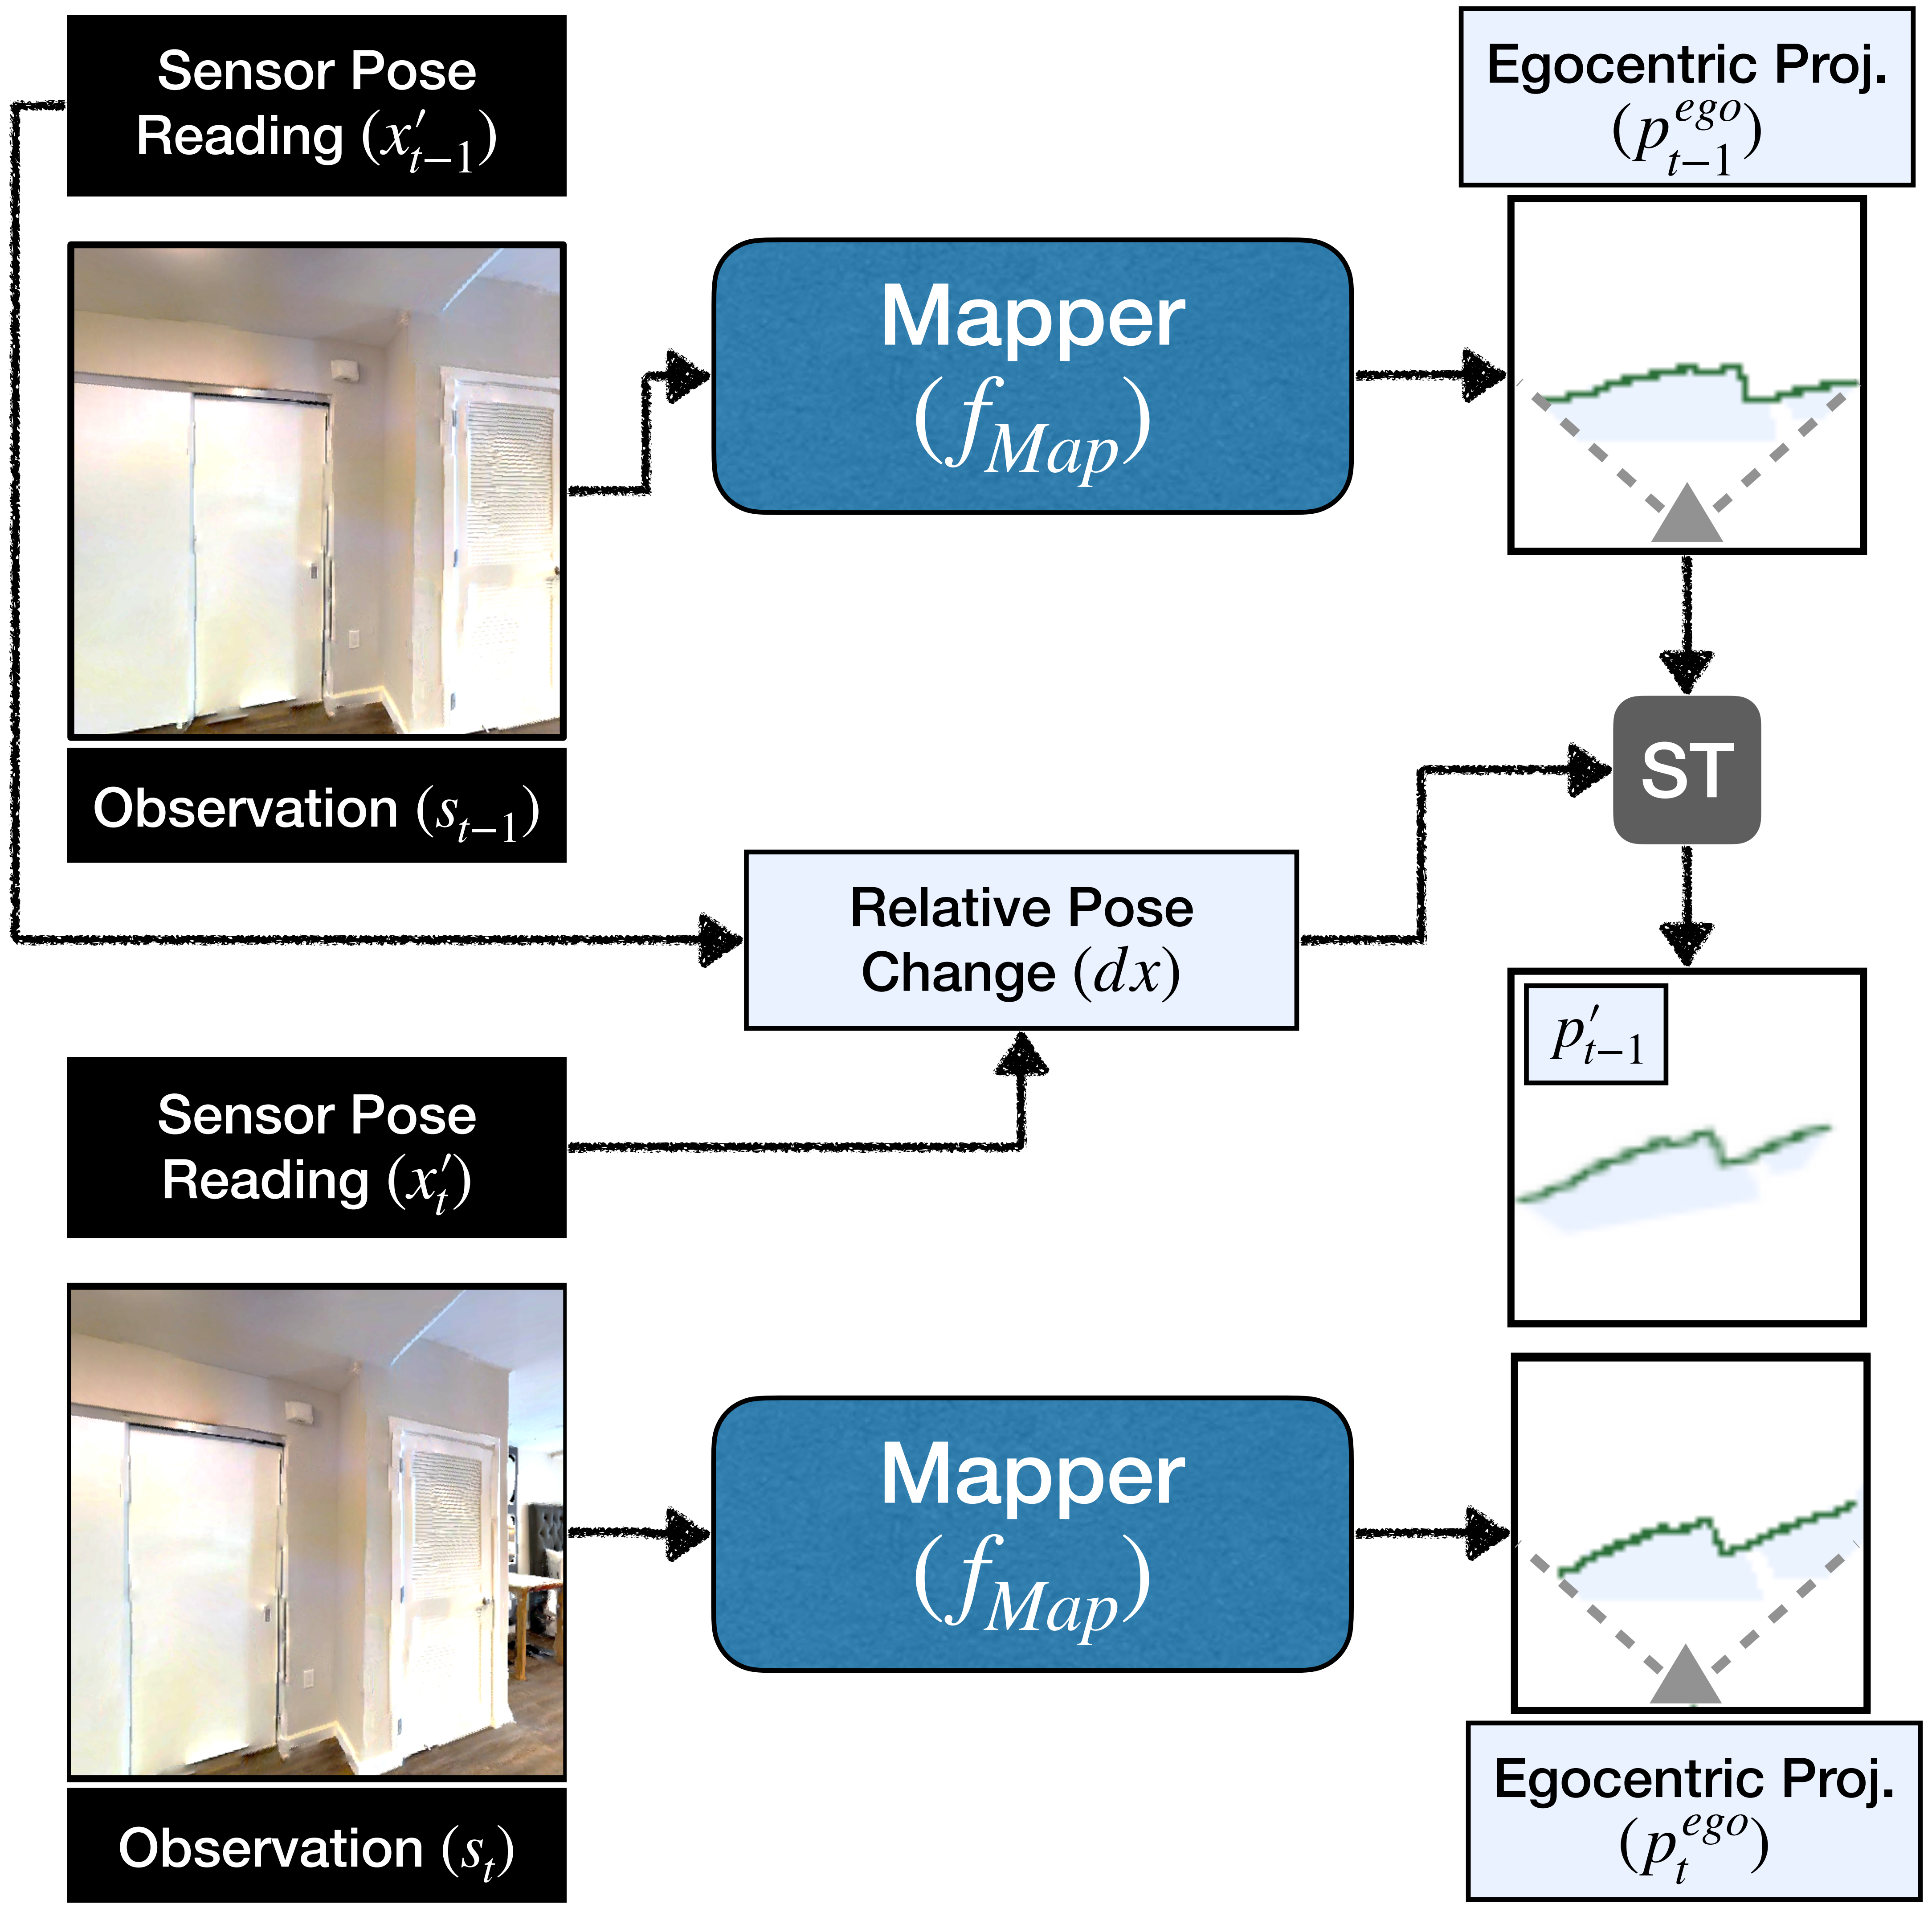

Neural SLAM Module

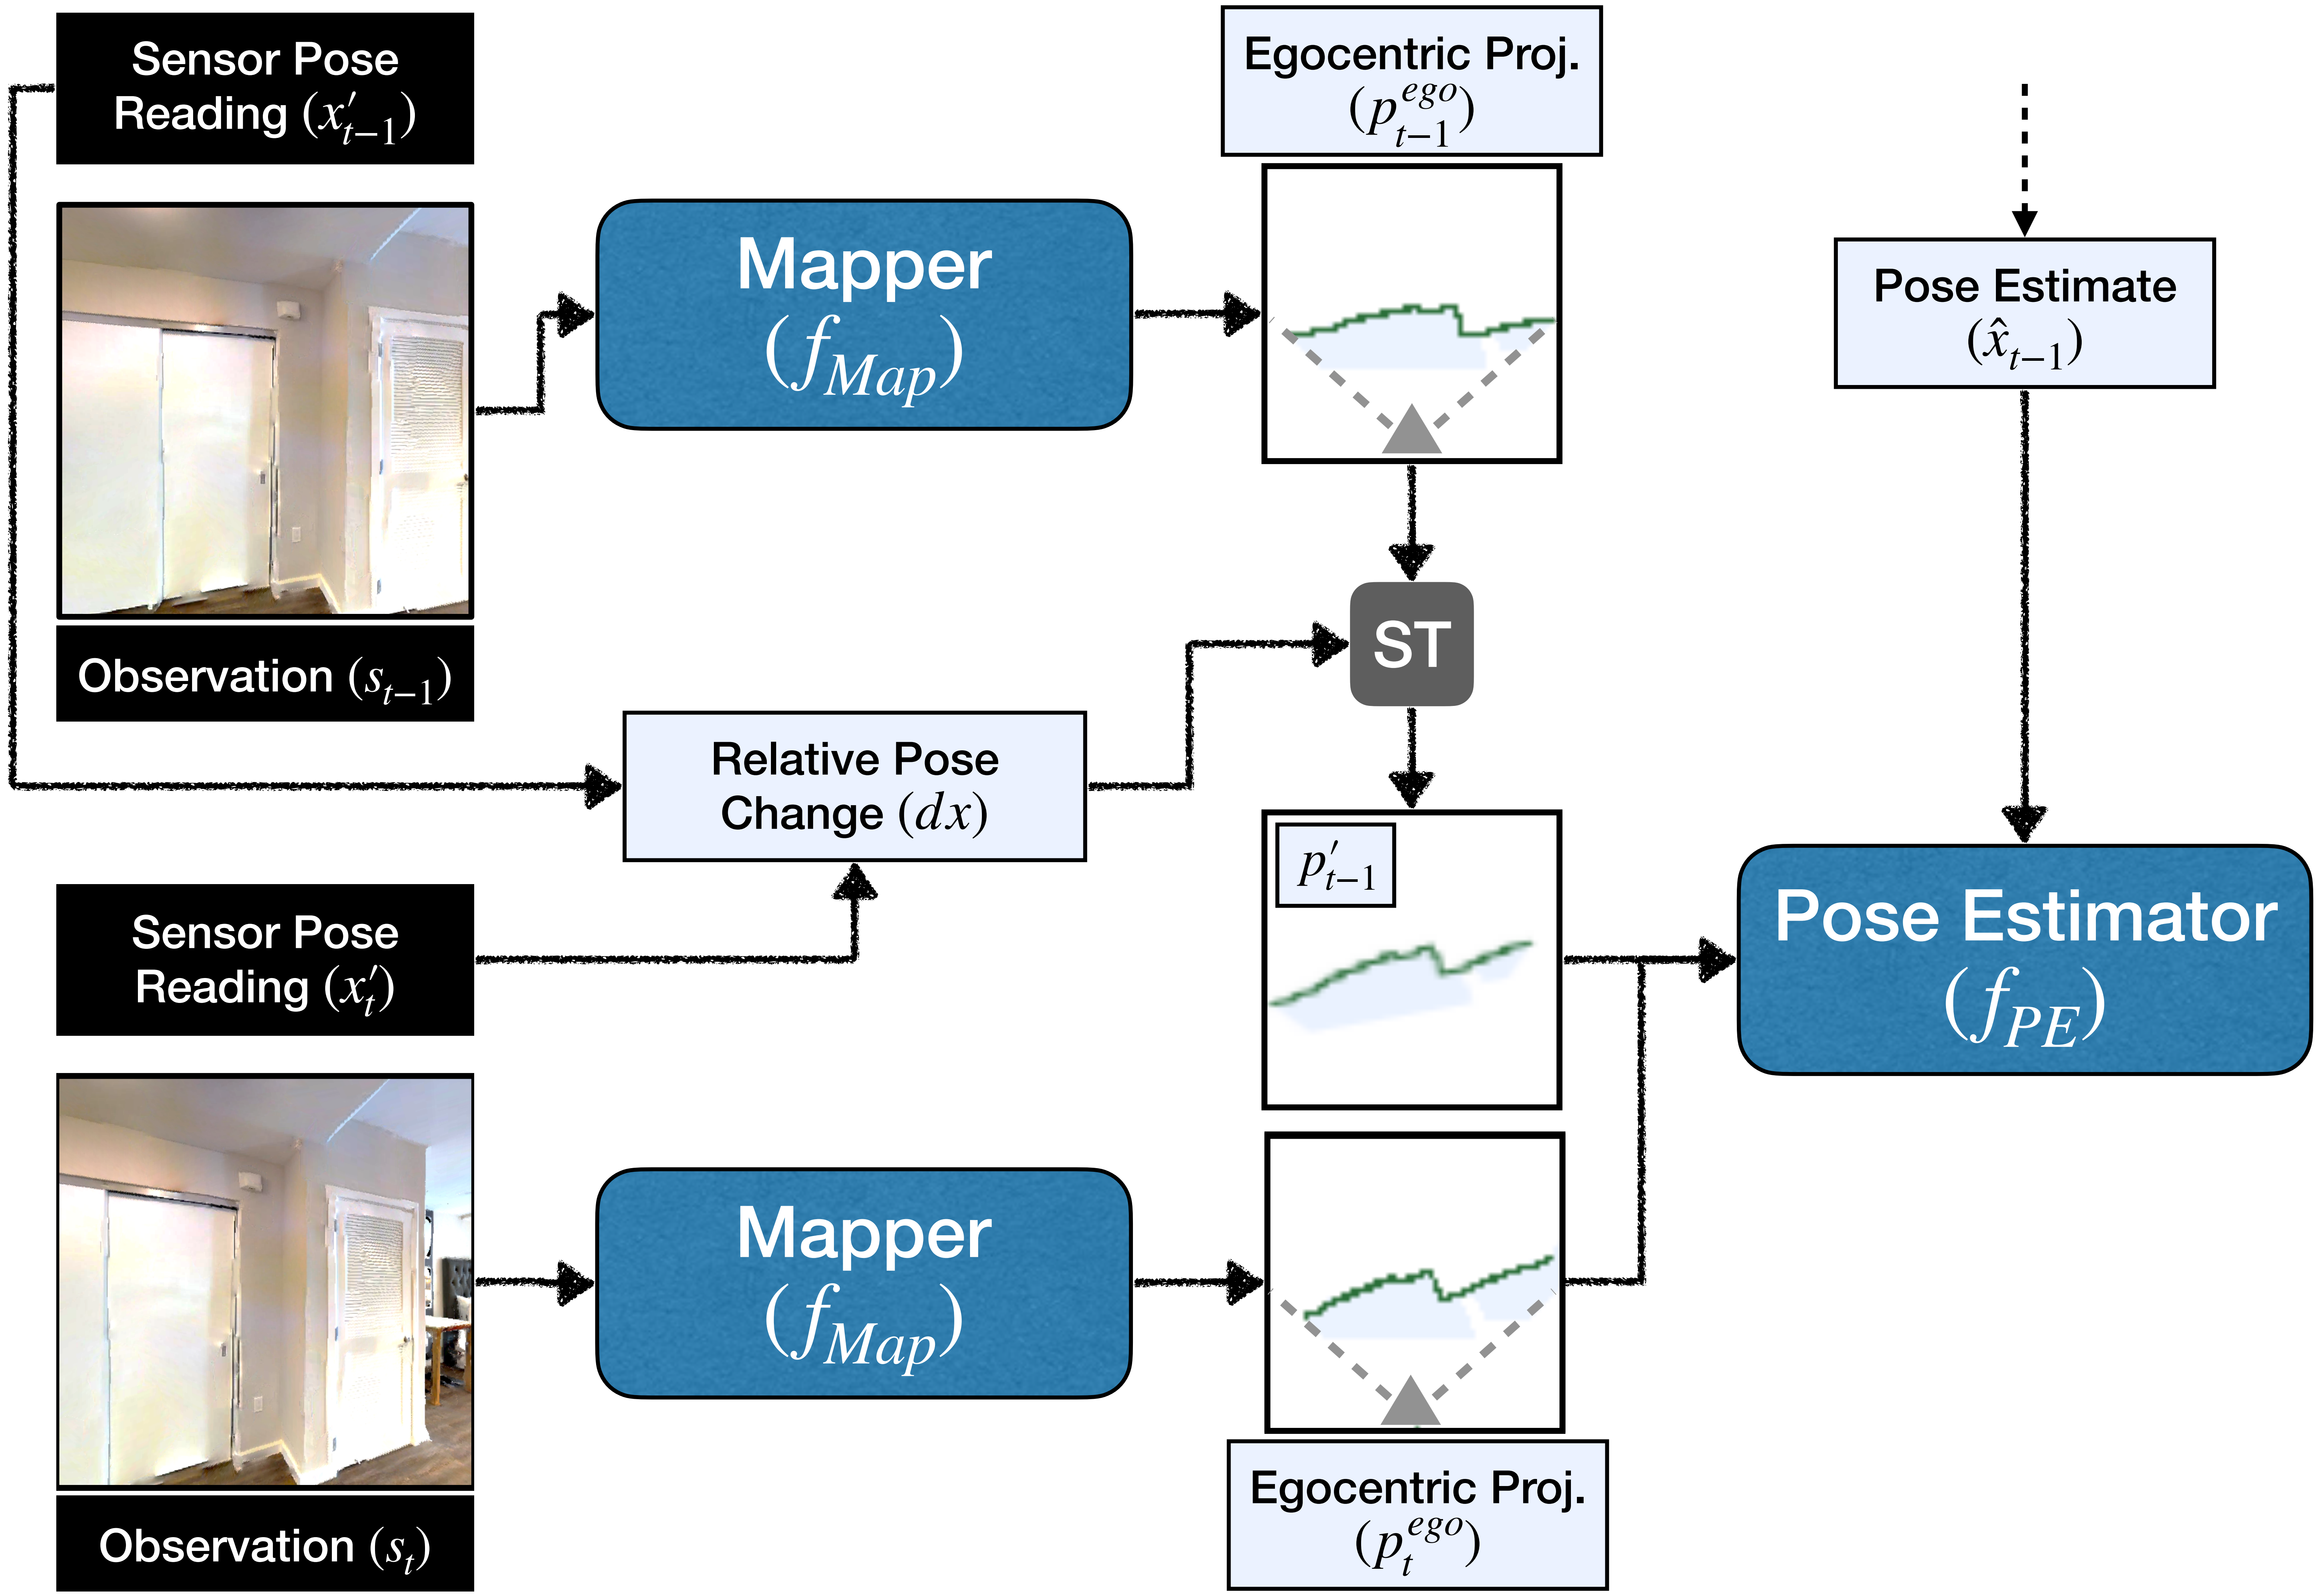

Neural SLAM Module

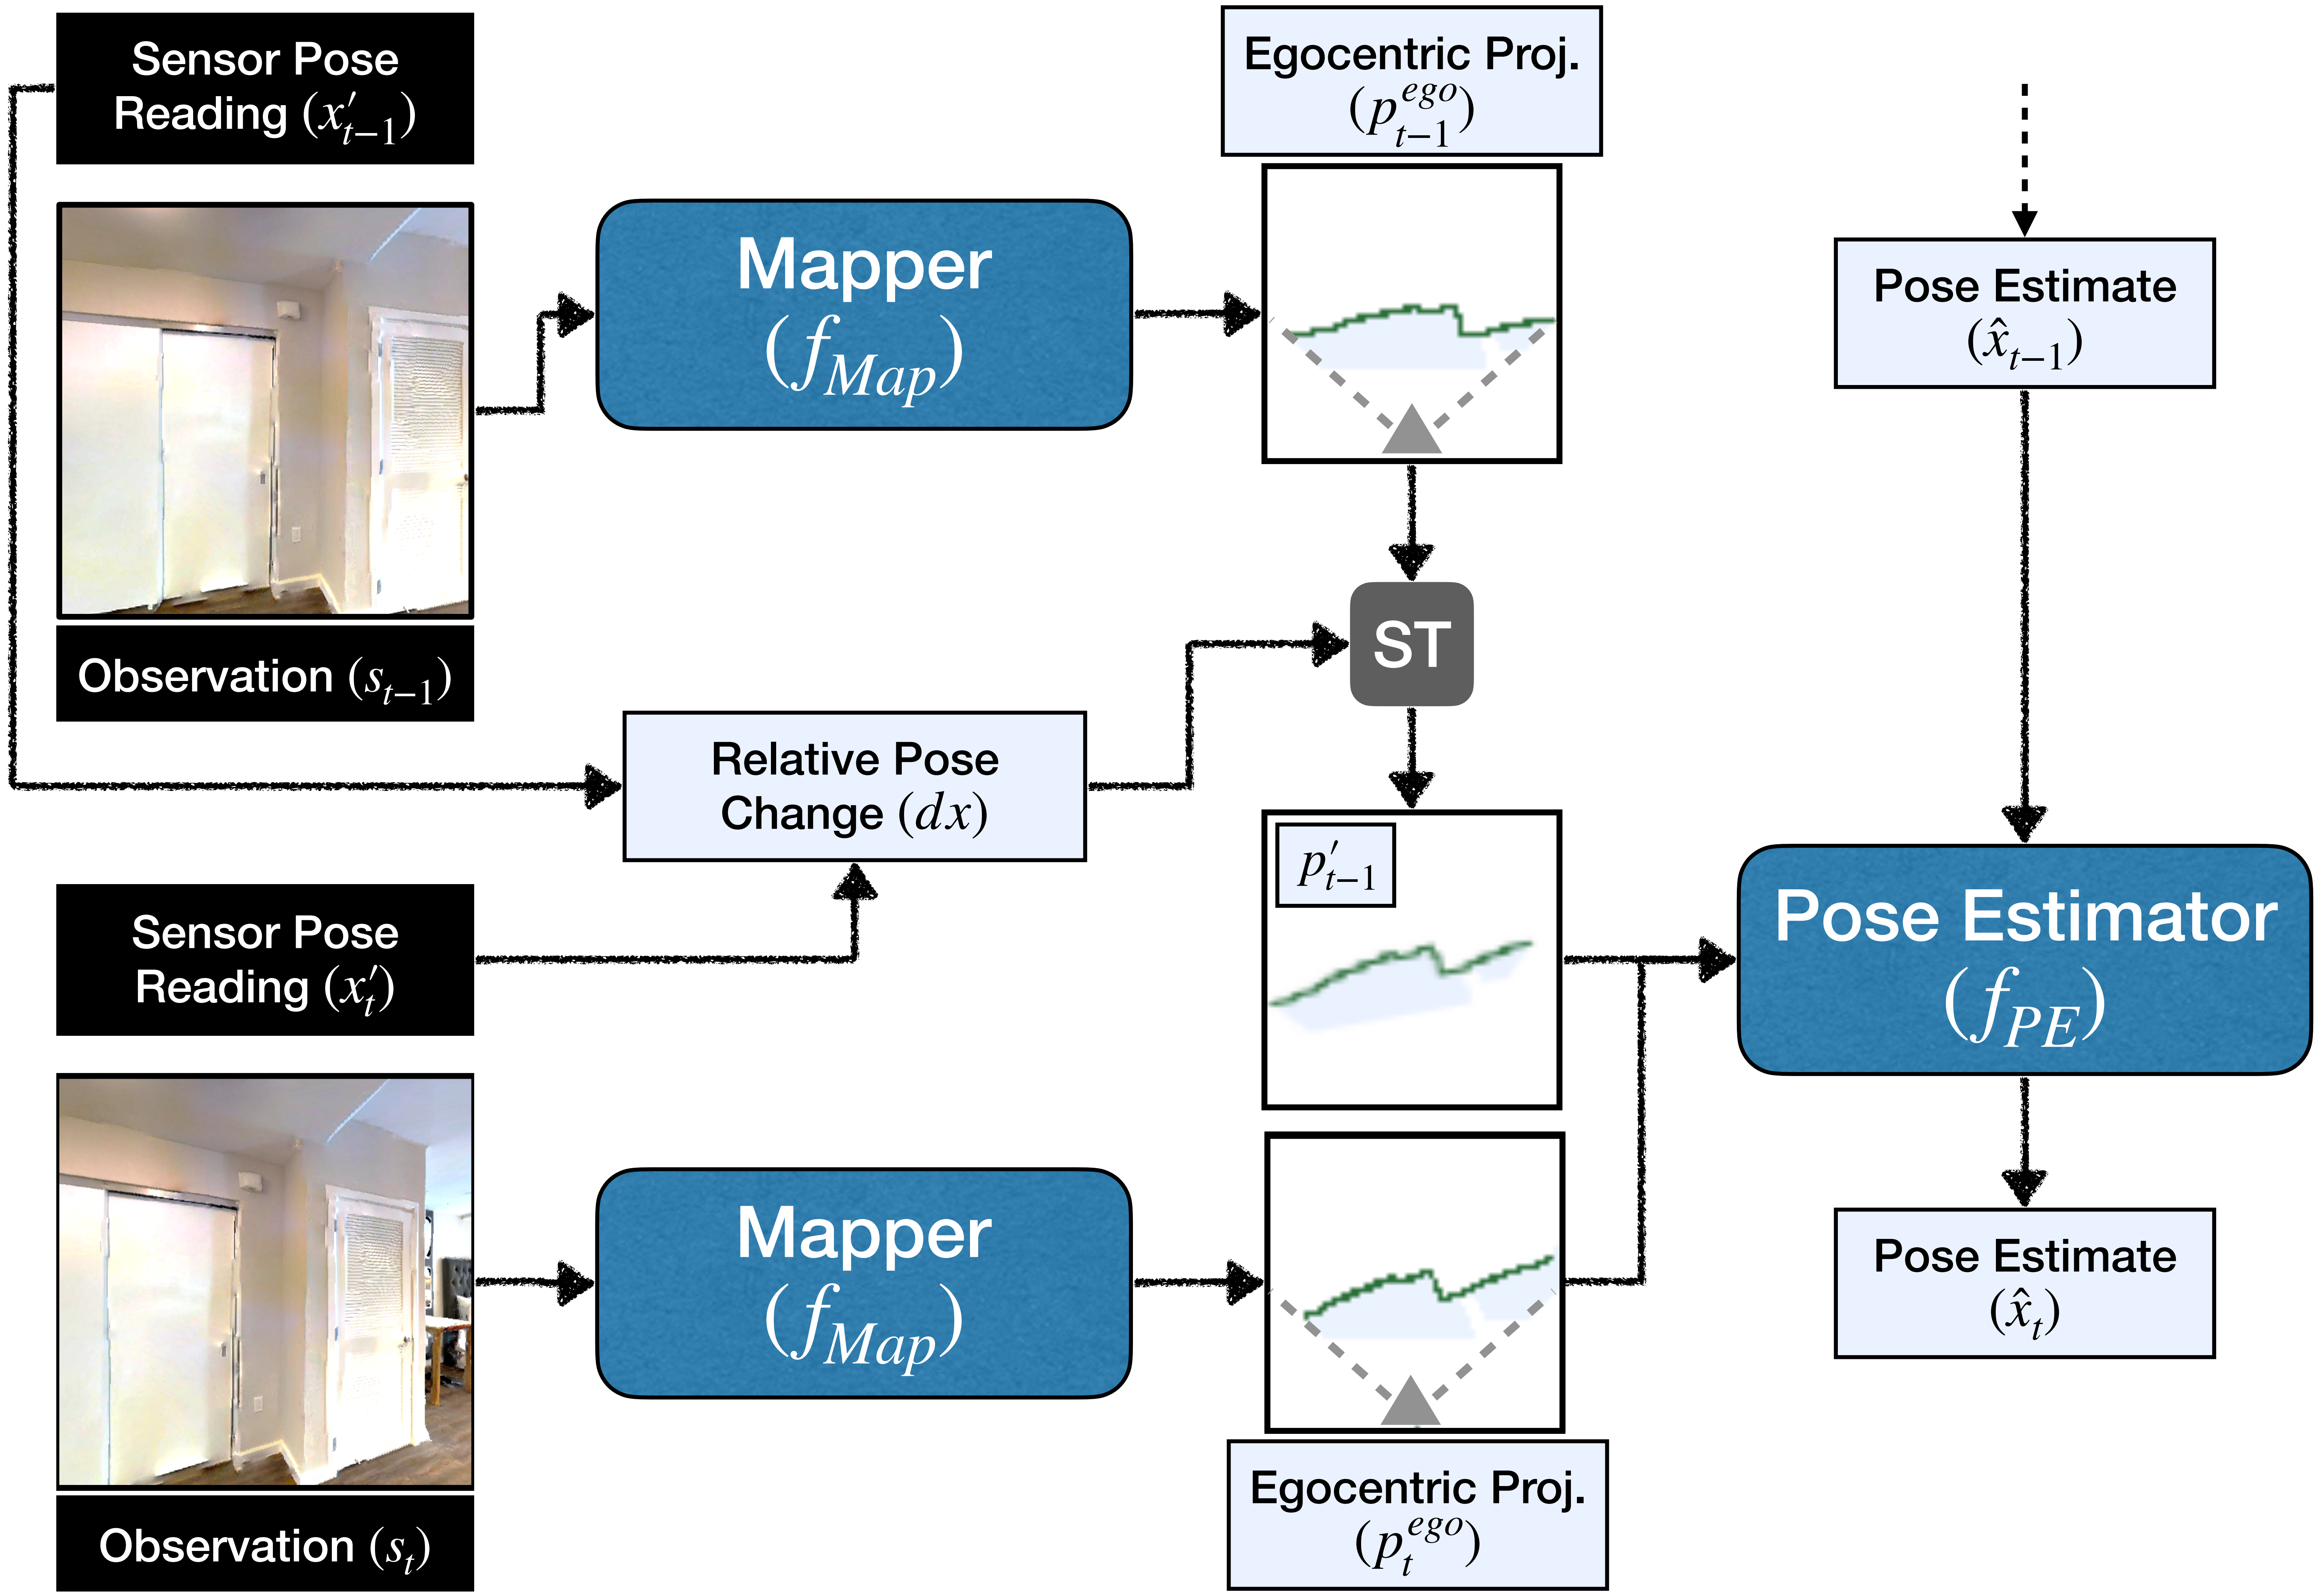

Neural SLAM Module

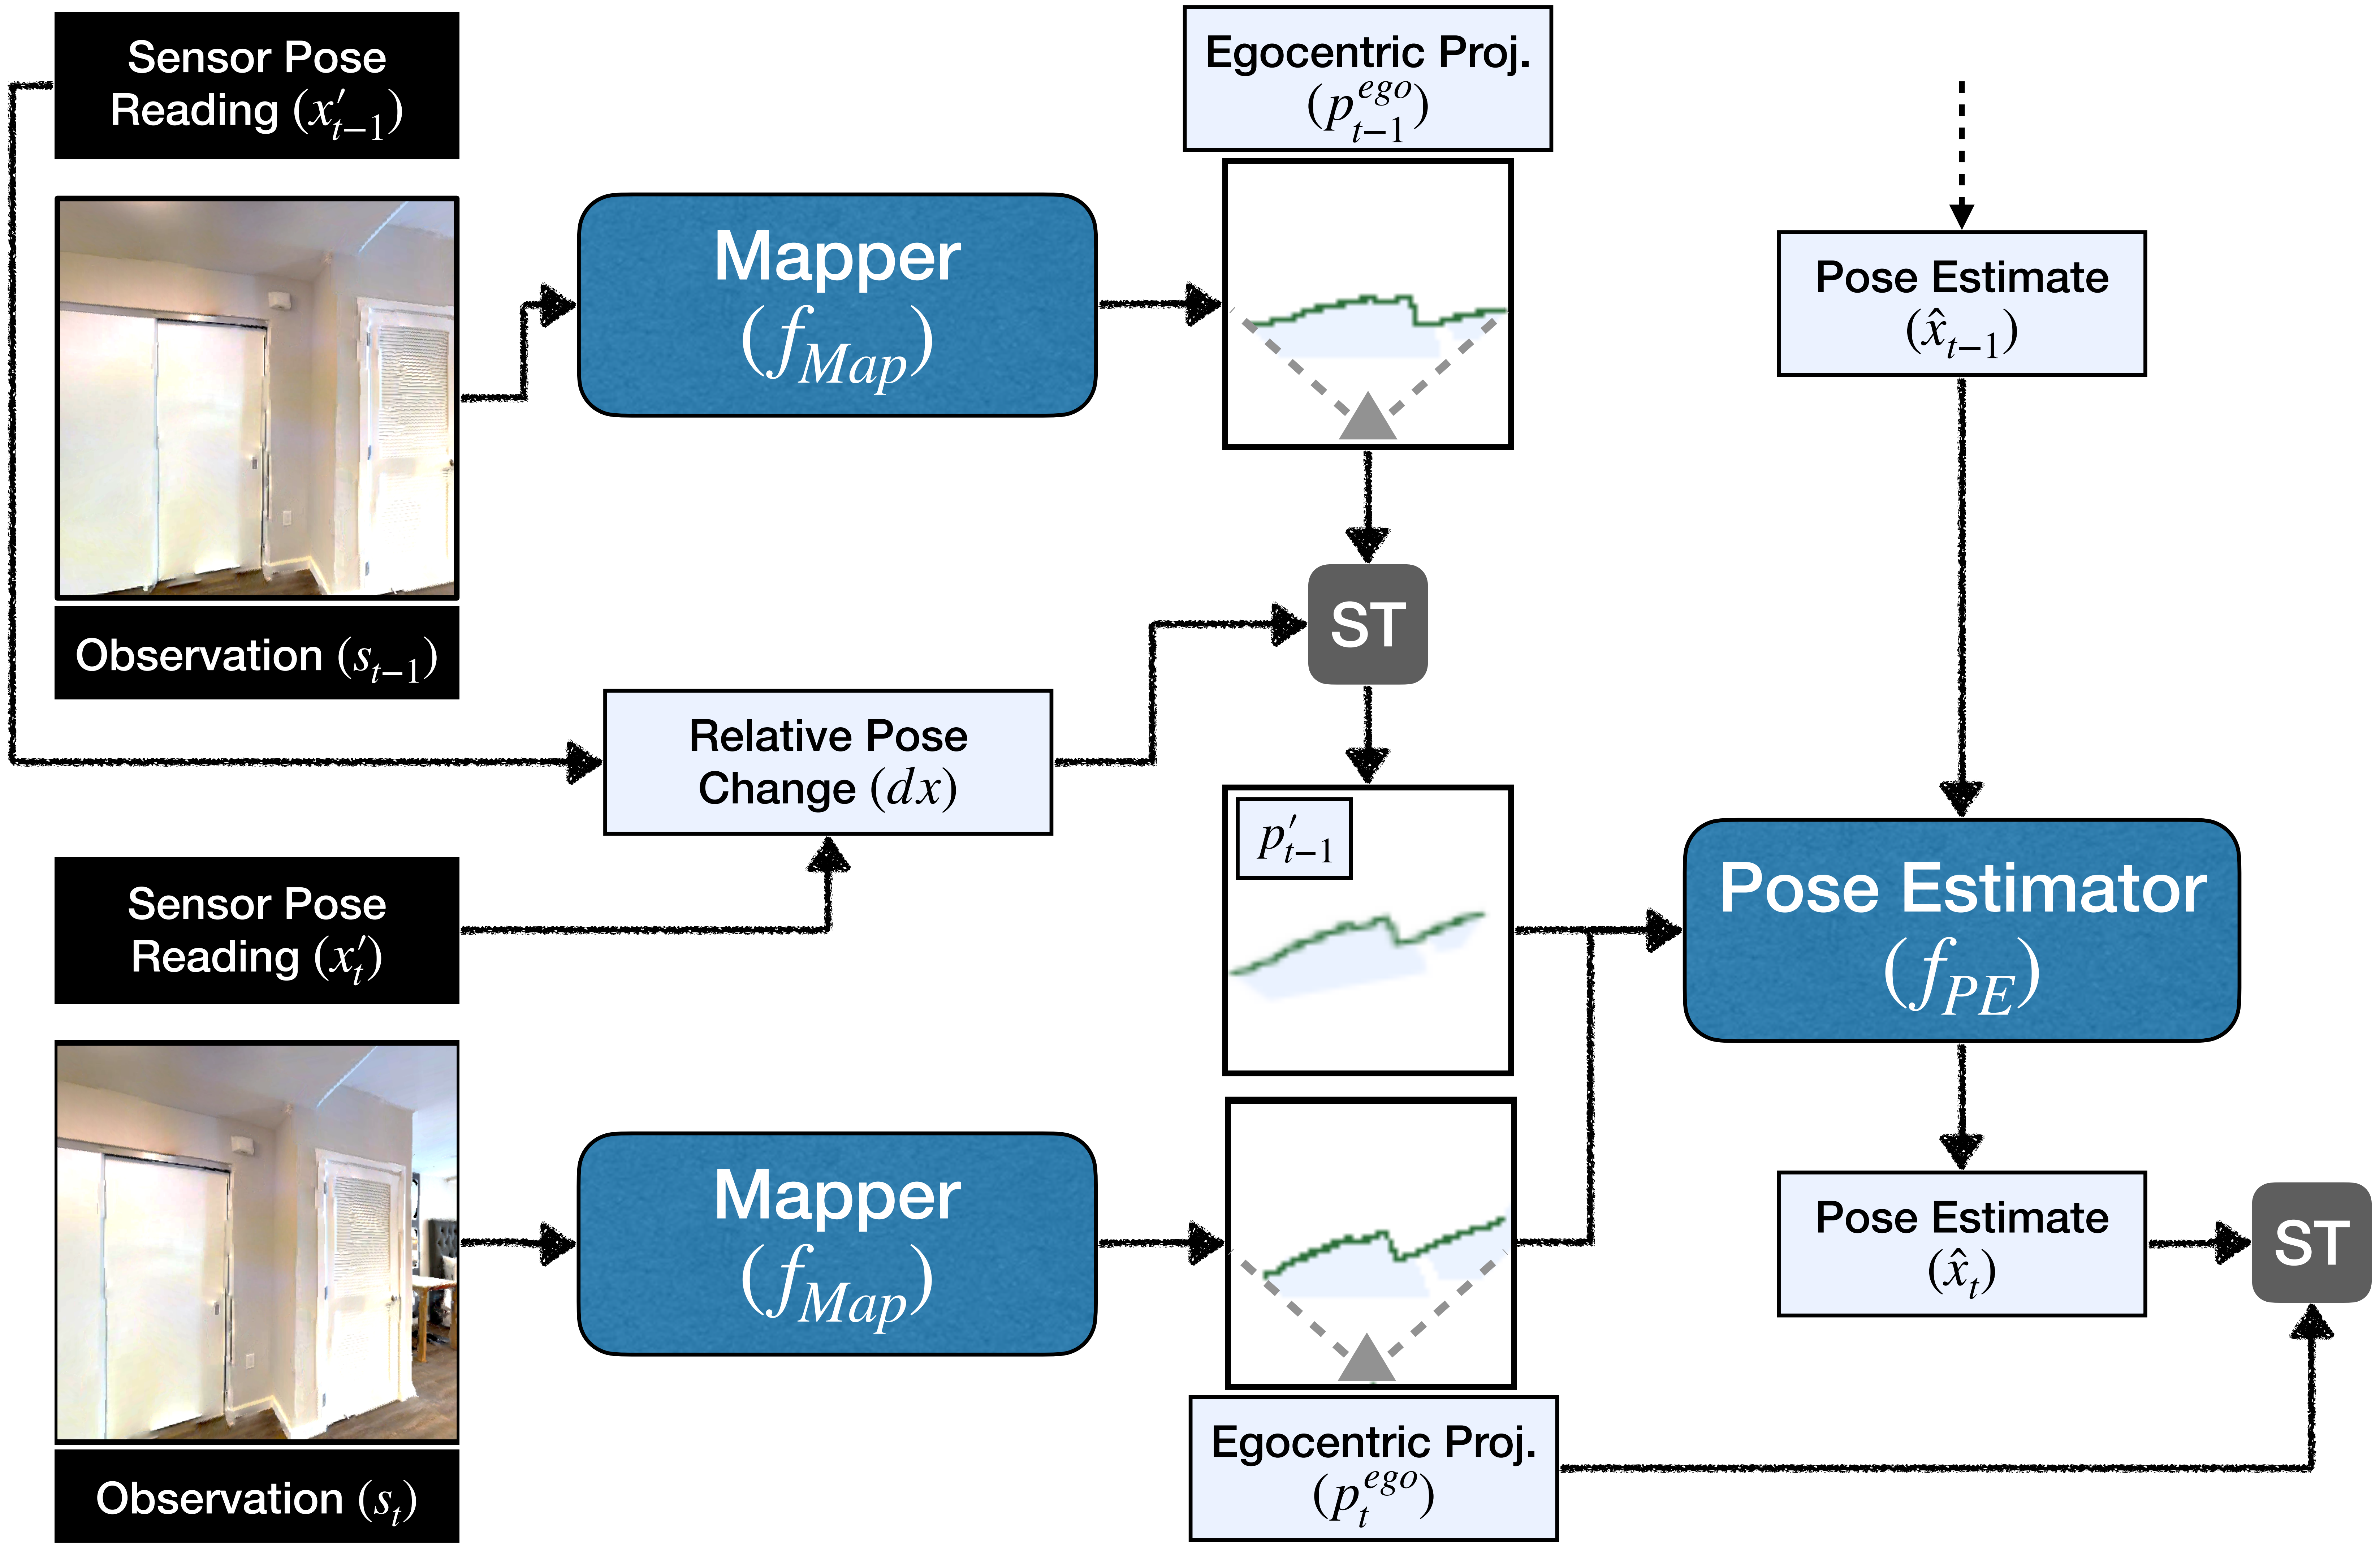

Neural SLAM Module

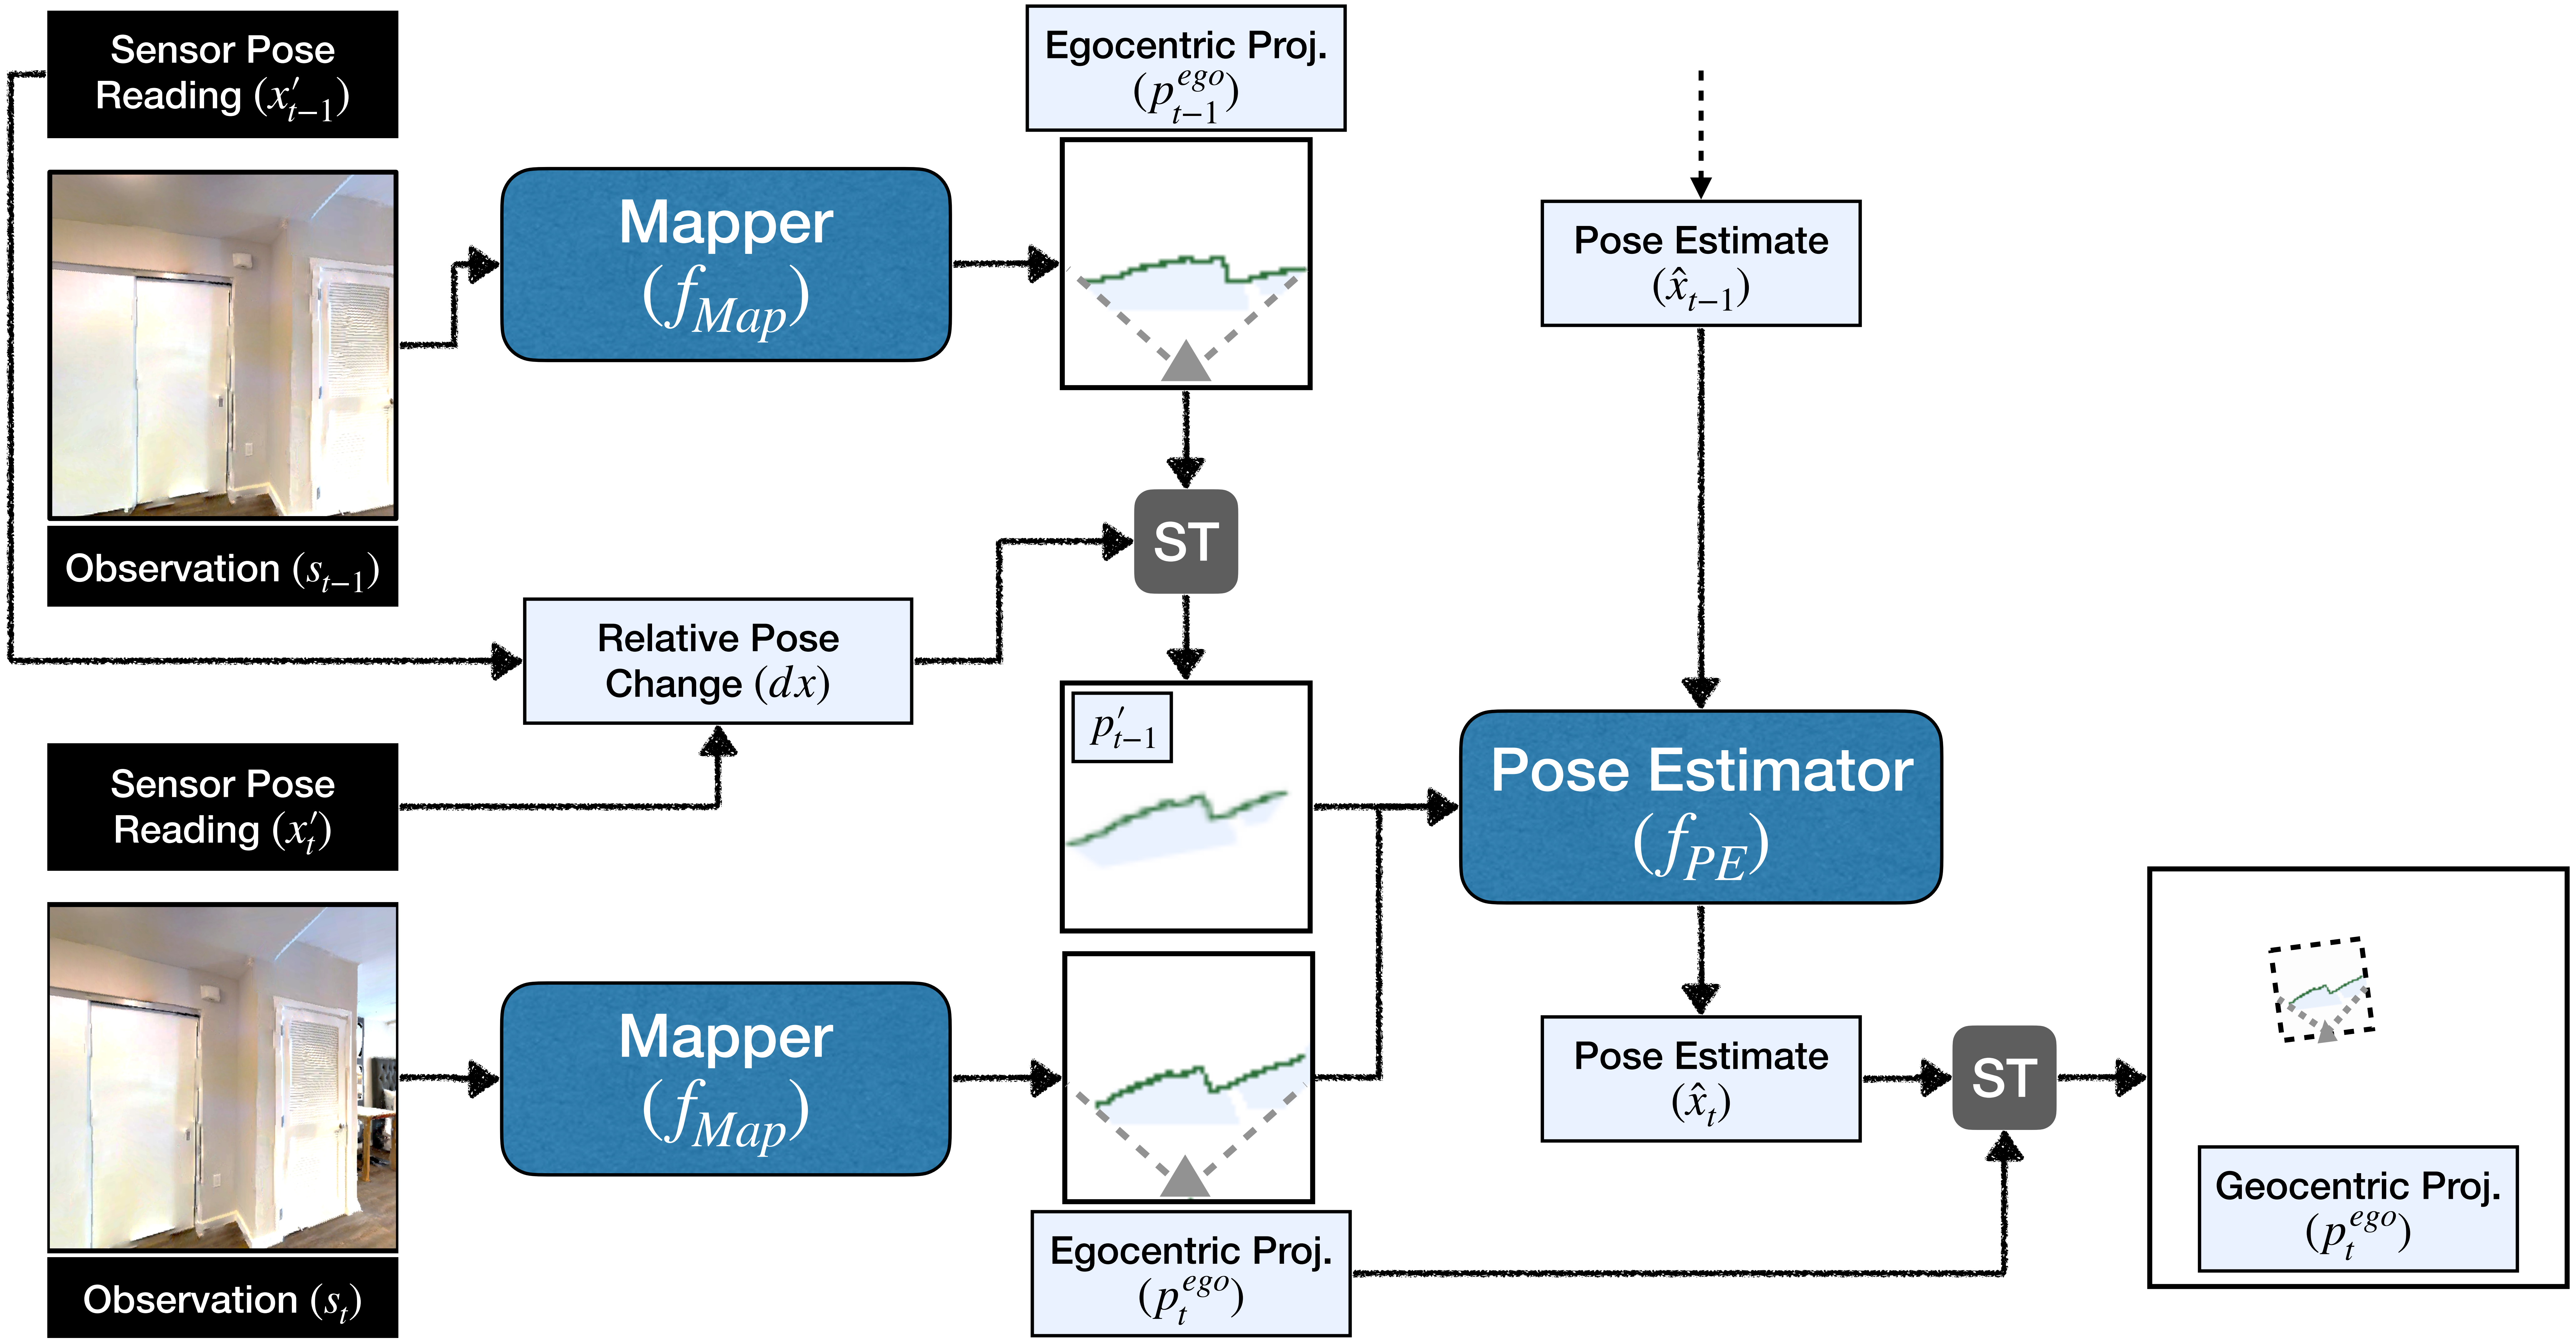

Neural SLAM Module

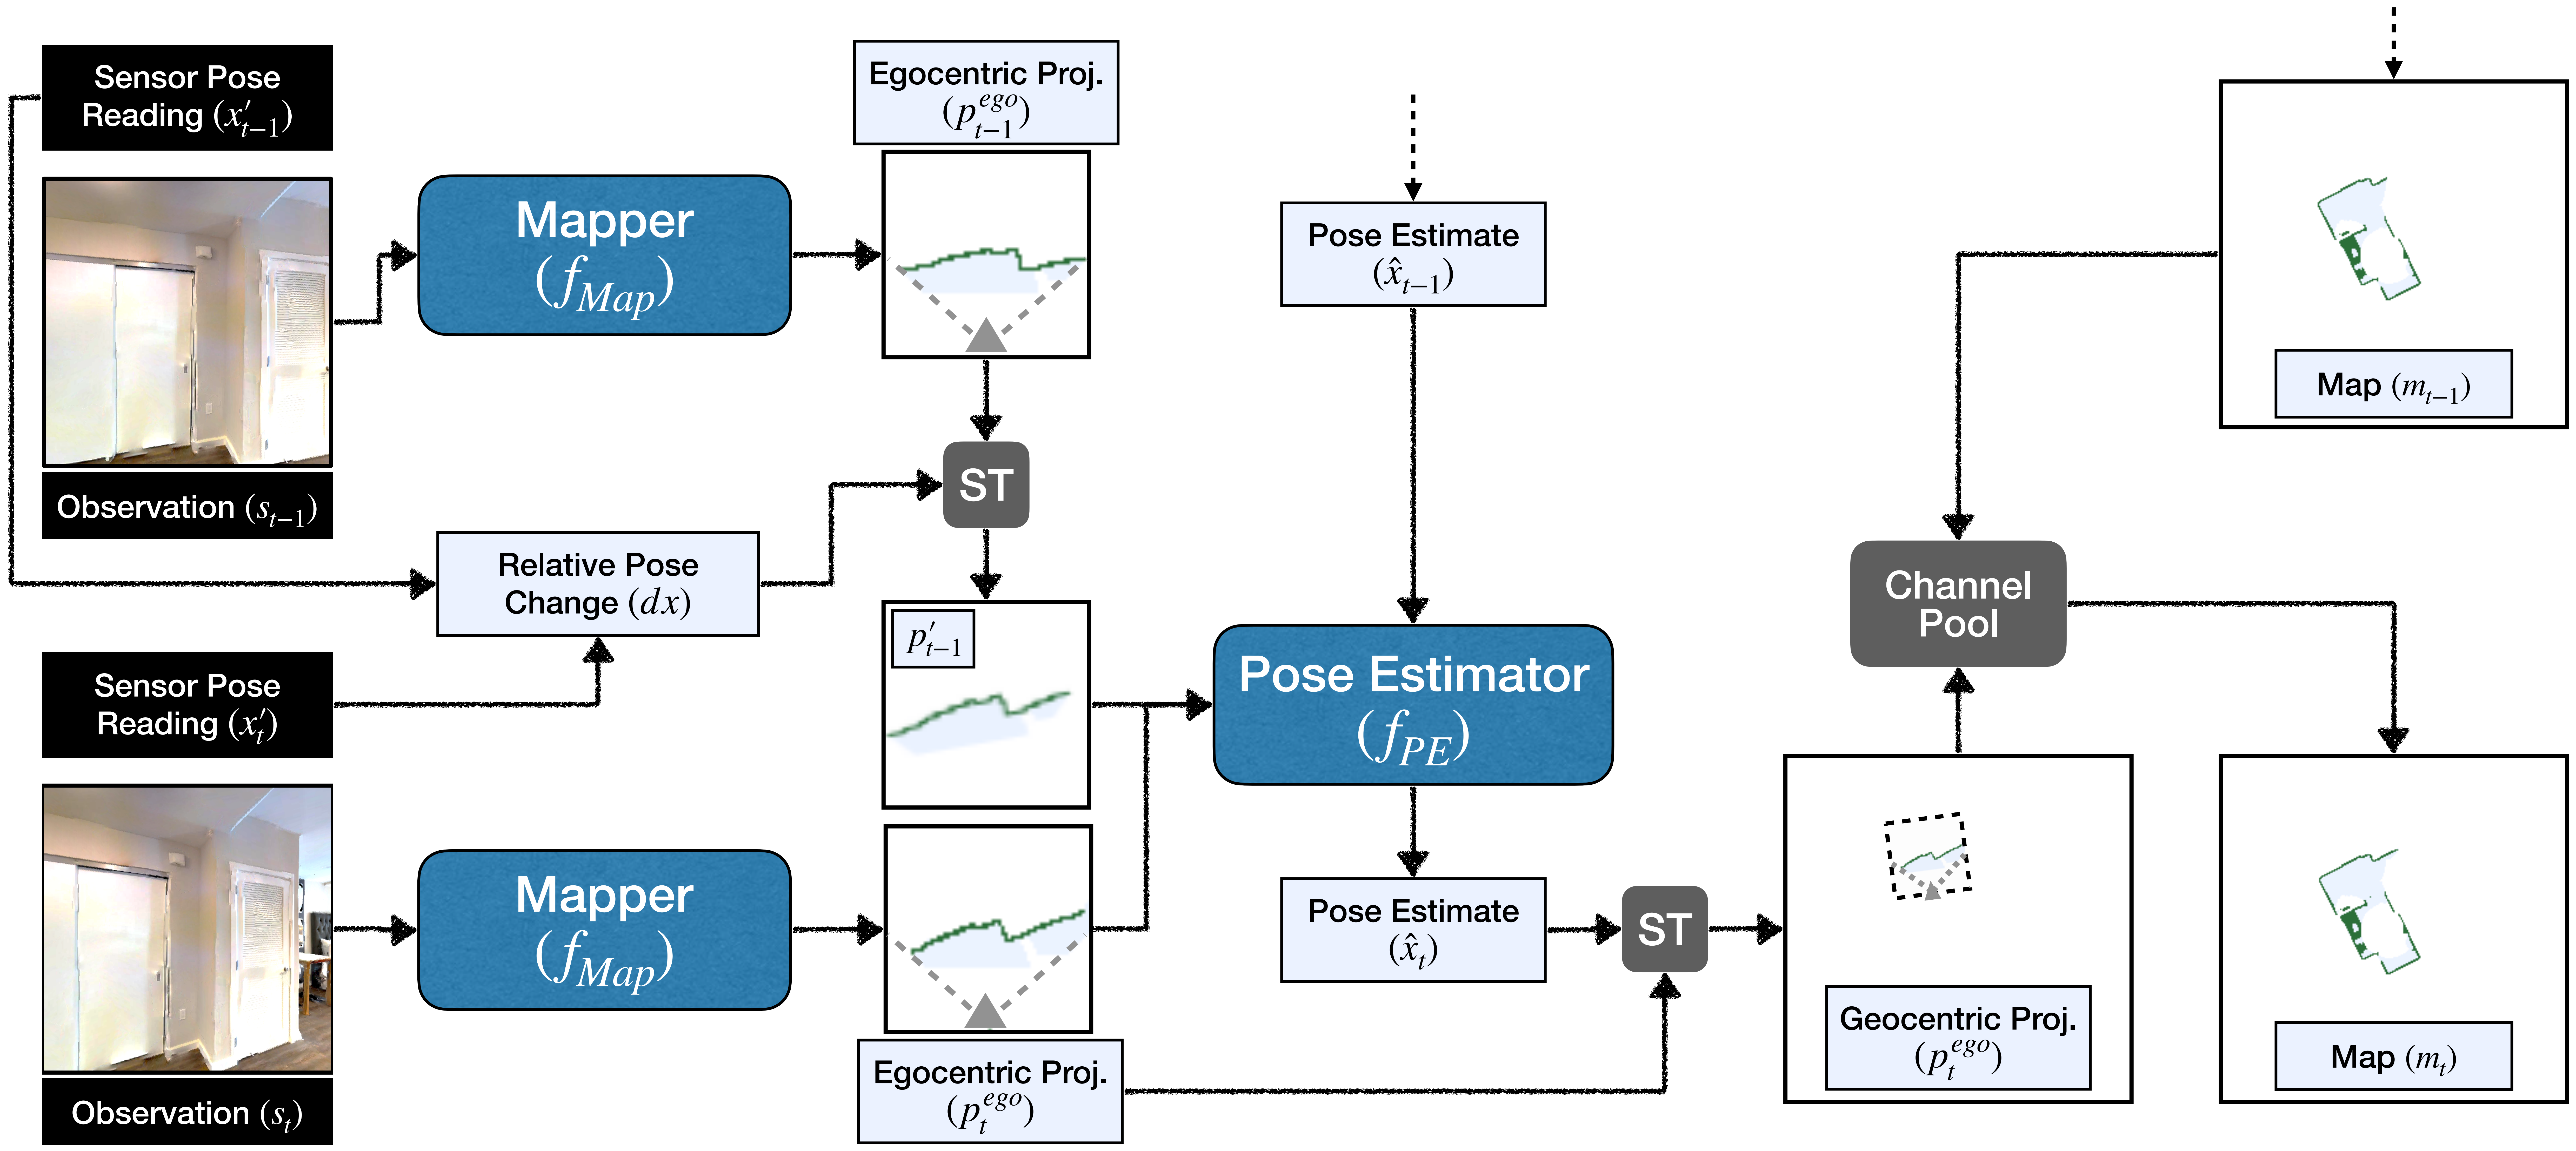

Neural SLAM Module

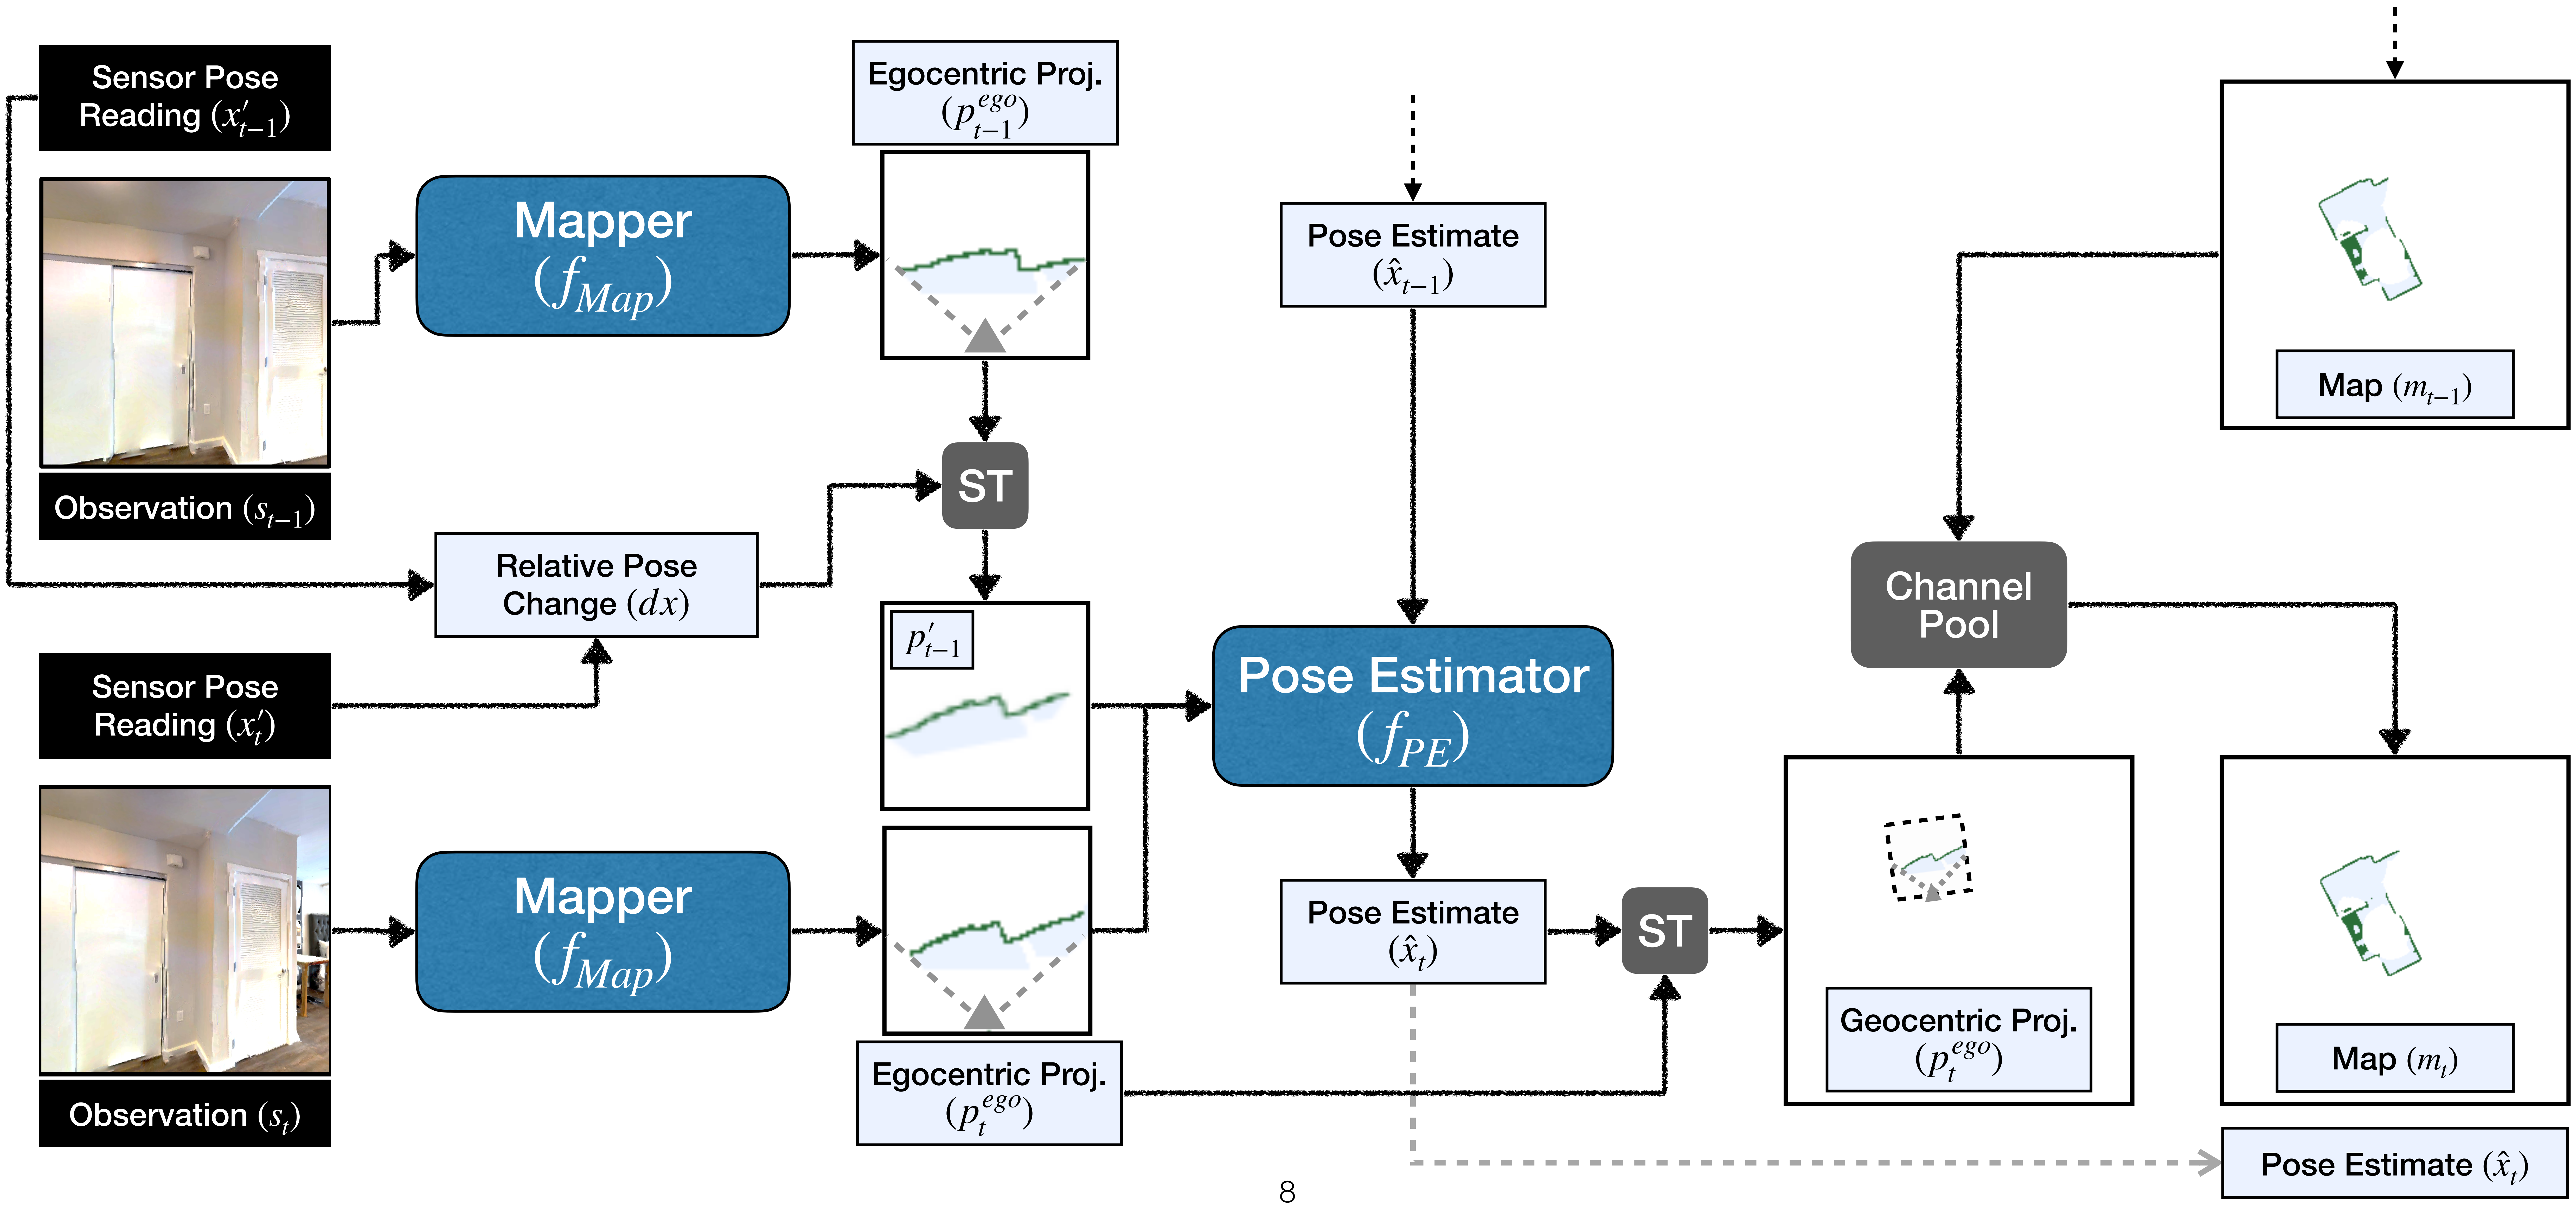

Neural SLAM Module

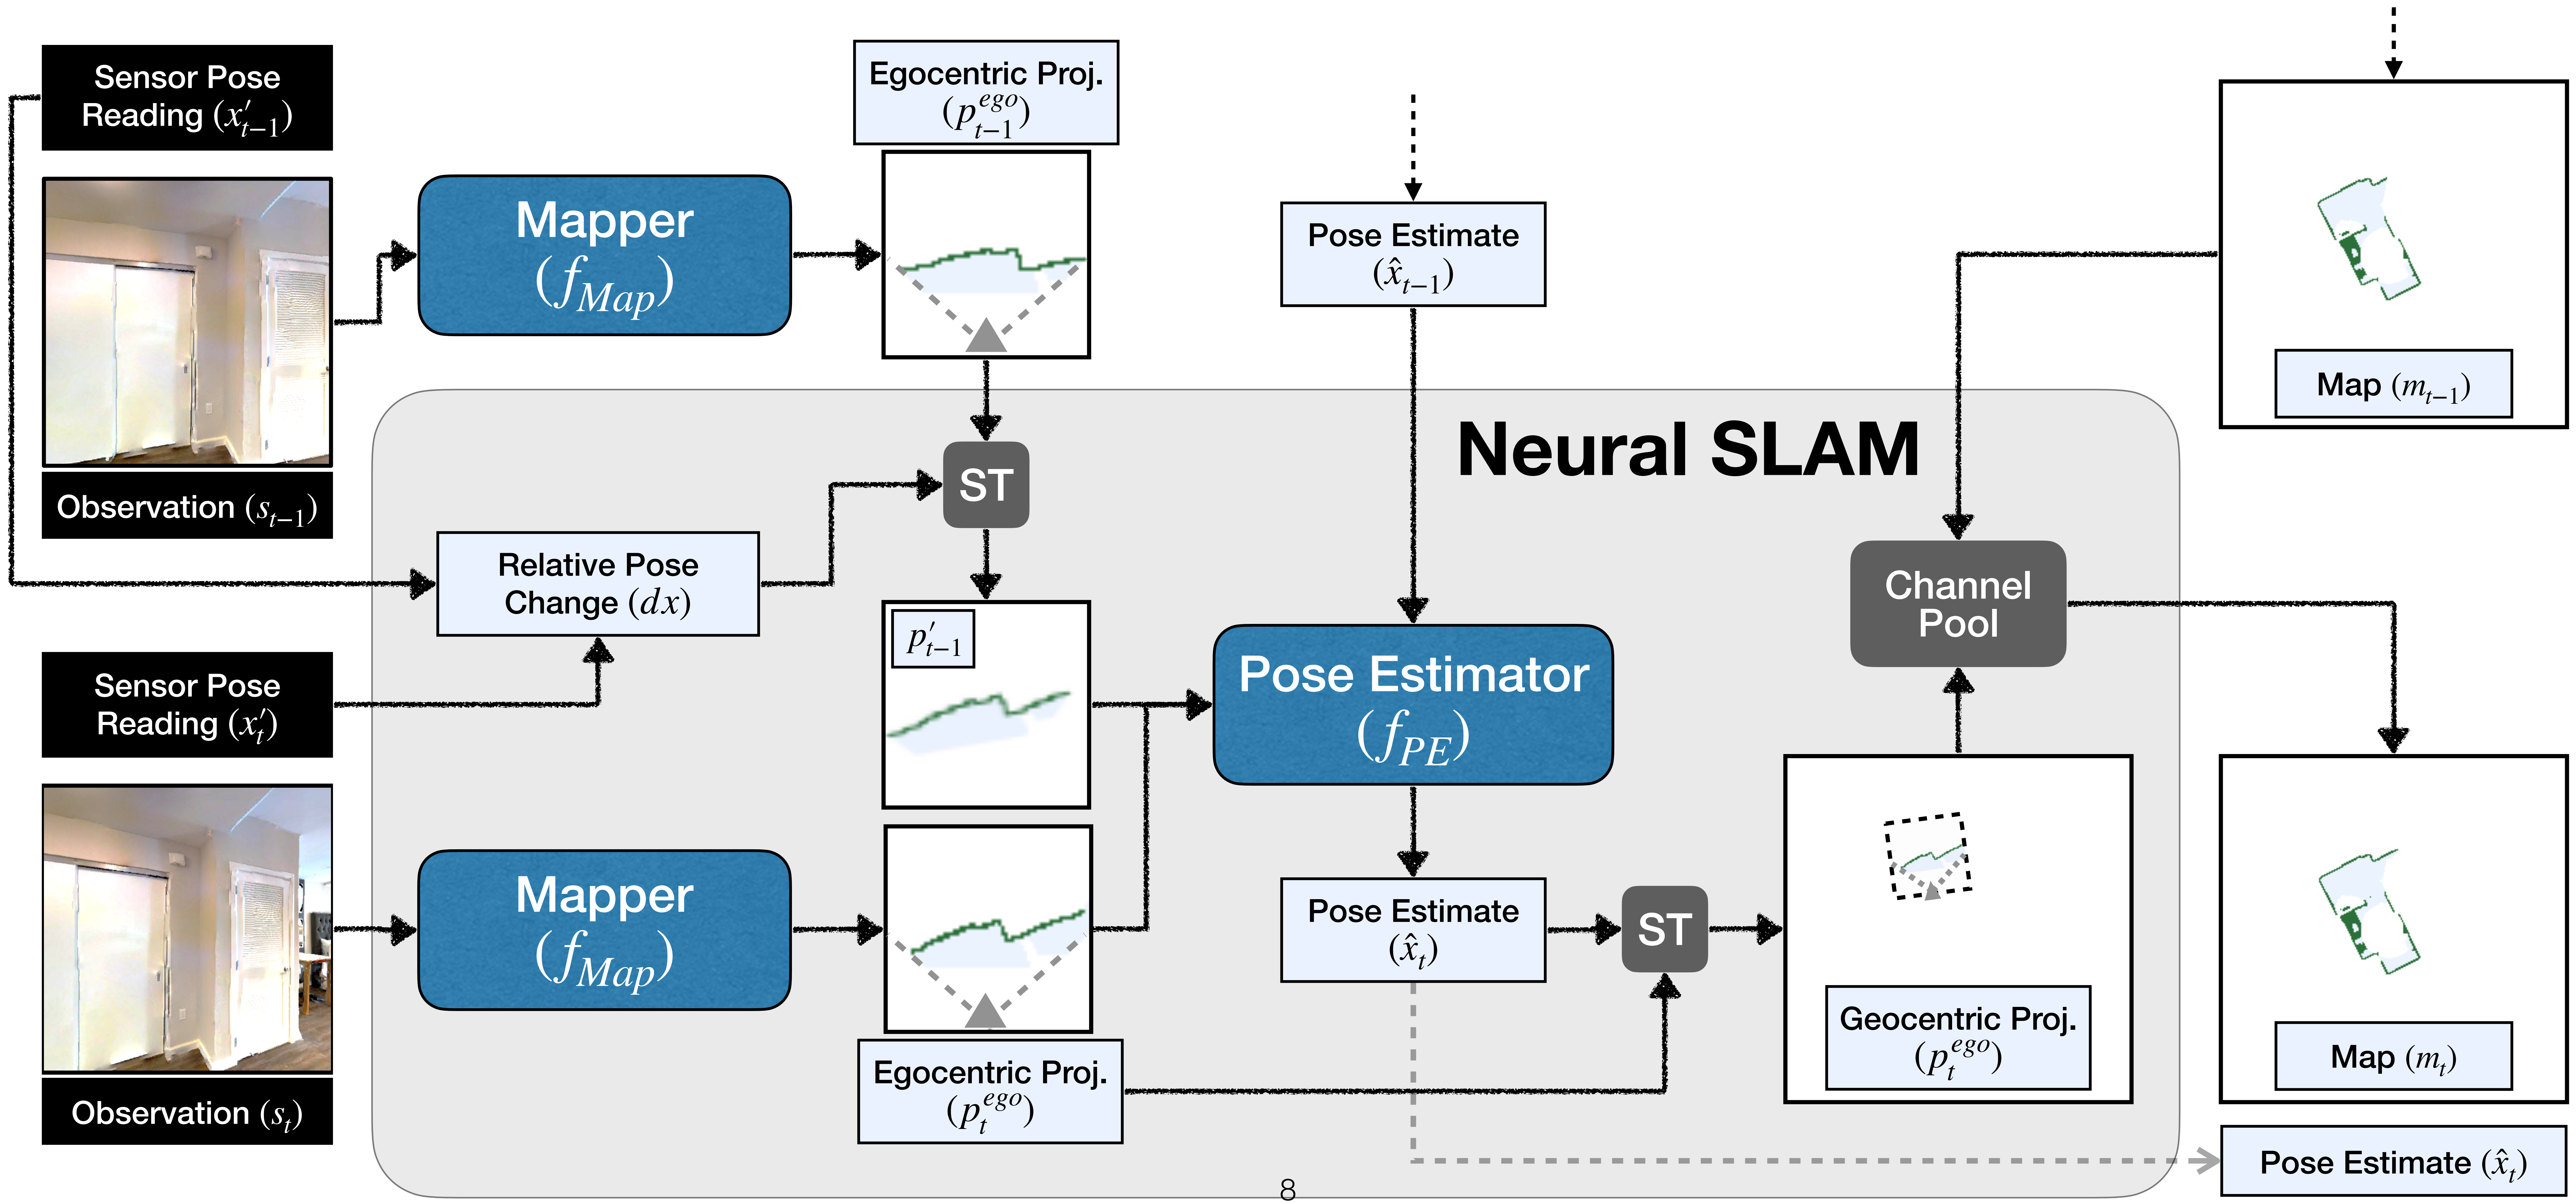

Exploration: Task Setup

Exploration: Task Setup

- **Visual Realism:** Gibson [2] and Matterport3D [3] datasets in Habitat simulator [1]

[1] Savva et al. 2019, [2] Xia et al. 2018, [3] Chang et al. 2017

Exploration: Task Setup

- **Visual Realism:** Gibson [2] and Matterport3D [3] datasets in Habitat simulator [1]
- **Physical Realism:** Actuation and motion sensor noise models based on real-data

[1] Savva et al. 2019, [2] Xia et al. 2018, [3] Chang et al. 2017

Exploration: Task Setup

- **Visual Realism:** Gibson [2] and Matterport3D [3] datasets in Habitat simulator [1]
- **Physical Realism:** Actuation and motion sensor noise models based on real-data
- Objective: Maximize the explored area
    - A cell is explored when it is known to be traversable
    - Metrics:
        - **Coverage (m^2)** - absolute explored area or coverage
        - **% Coverage** - percentage of the environment explored

[1] Savva et al. 2019, [2] Xia et al. 2018, [3] Chang et al. 2017

Exploration: Task Setup

- **Visual Realism:** Gibson [2] and Matterport3D [3] datasets in Habitat simulator [1]
- **Physical Realism:** Actuation and motion sensor noise models based on real-data
- Objective: Maximize the explored area
    - A cell is explored when it is known to be traversable
    - Metrics:
        - **Coverage (m^2)** - absolute explored area or coverage
        - **% Coverage** - percentage of the environment explored
- Fixed episode length of 1000 steps

[1] Savva et al. 2019, [2] Xia et al. 2018, [3] Chang et al. 2017

Exploration: Task Setup

- **Visual Realism:** Gibson [2] and Matterport3D [3] datasets in Habitat simulator [1]
- **Physical Realism:** Actuation and motion sensor noise models based on real-data
- Objective: Maximize the explored area
  - A cell is explored when it is known to be traversable
  - Metrics:
    - **Coverage (m^2)** - absolute explored area or coverage
    - **% Coverage** - percentage of the environment explored
- Fixed episode length of 1000 steps
- All methods trained for 10 million frames

[1] Savva et al. 2019, [2] Xia et al. 2018, [3] Chang et al. 2017

Demo Video: Exploration

https://youtu.be/tlyz68j_jvE

Exploration Results

Exploration Results

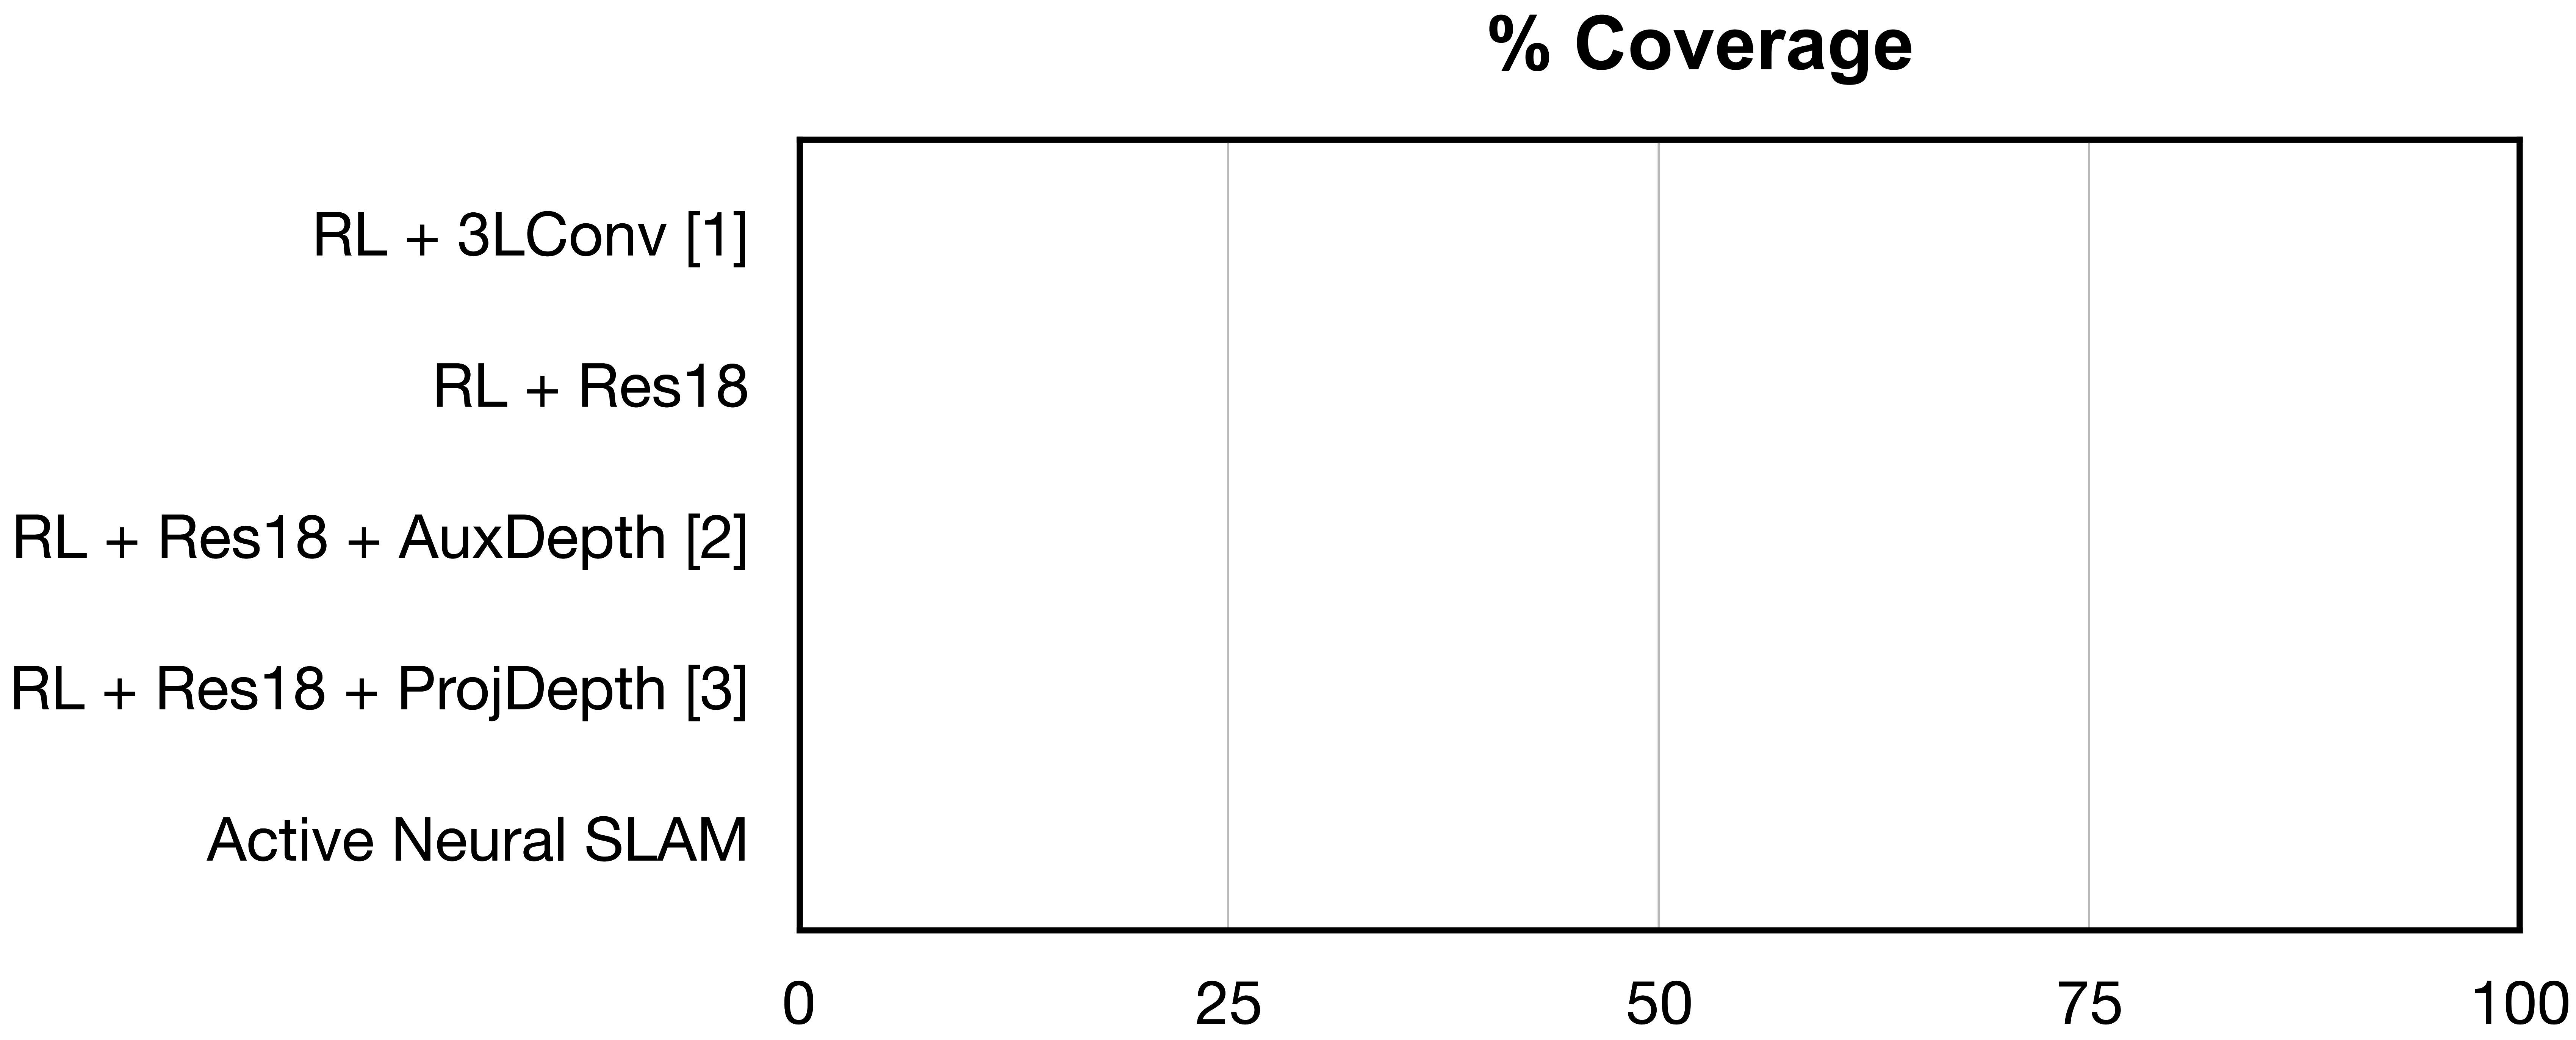

**Adapted from [1] Lample & Chaplot. AAAI-17, [2] Mirowski et al. ICLR-17, [3] Chen el al. ICLR-19*

Exploration Results

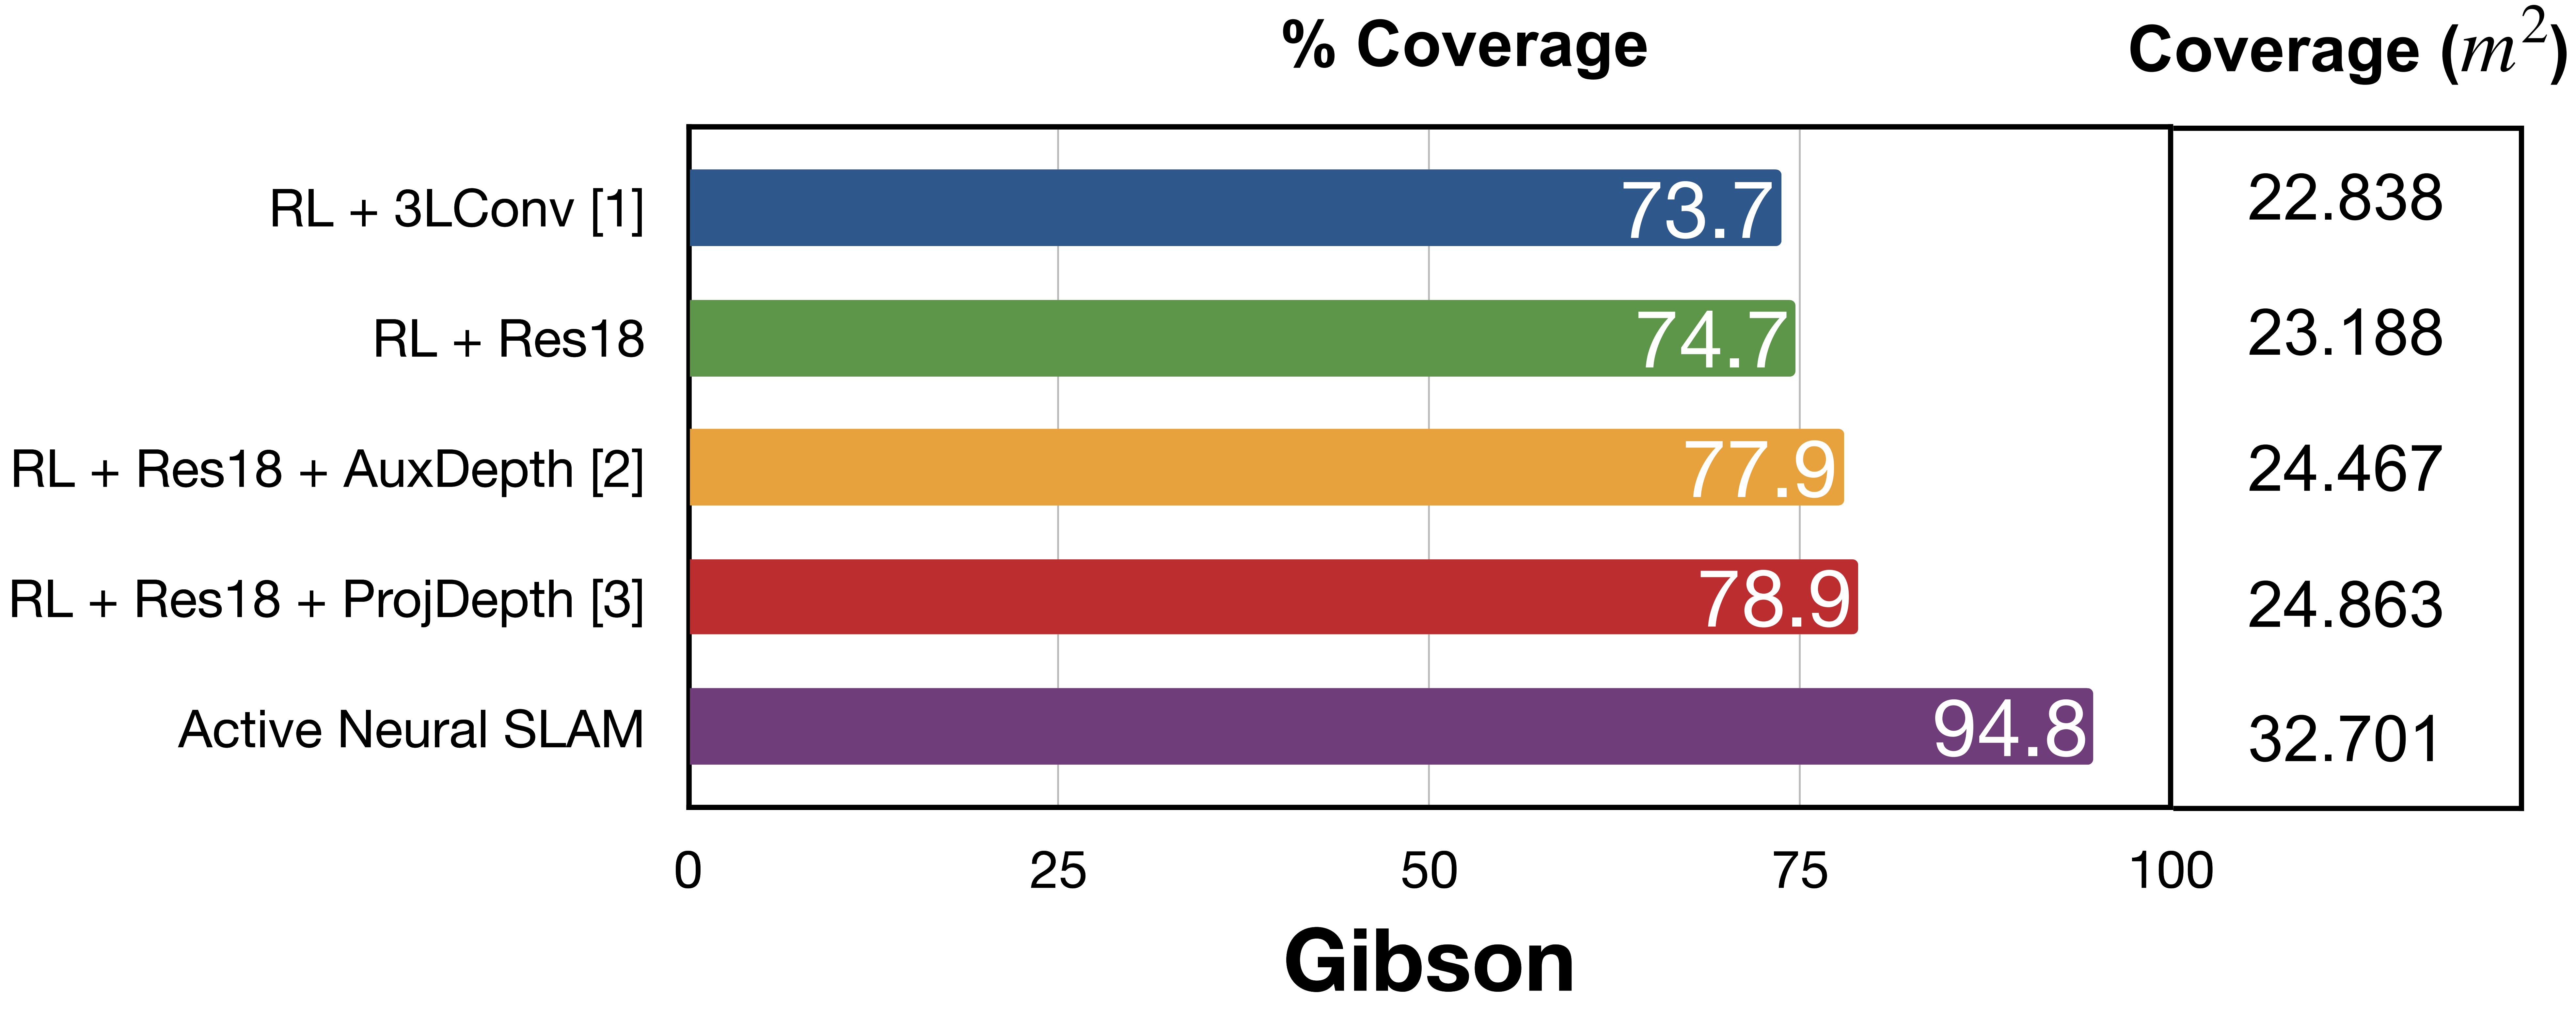

**Adapted from [1] Lample & Chaplot. AAAI-17, [2] Mirowski et al. ICLR-17, [3] Chen el al. ICLR-19*

Exploration Results

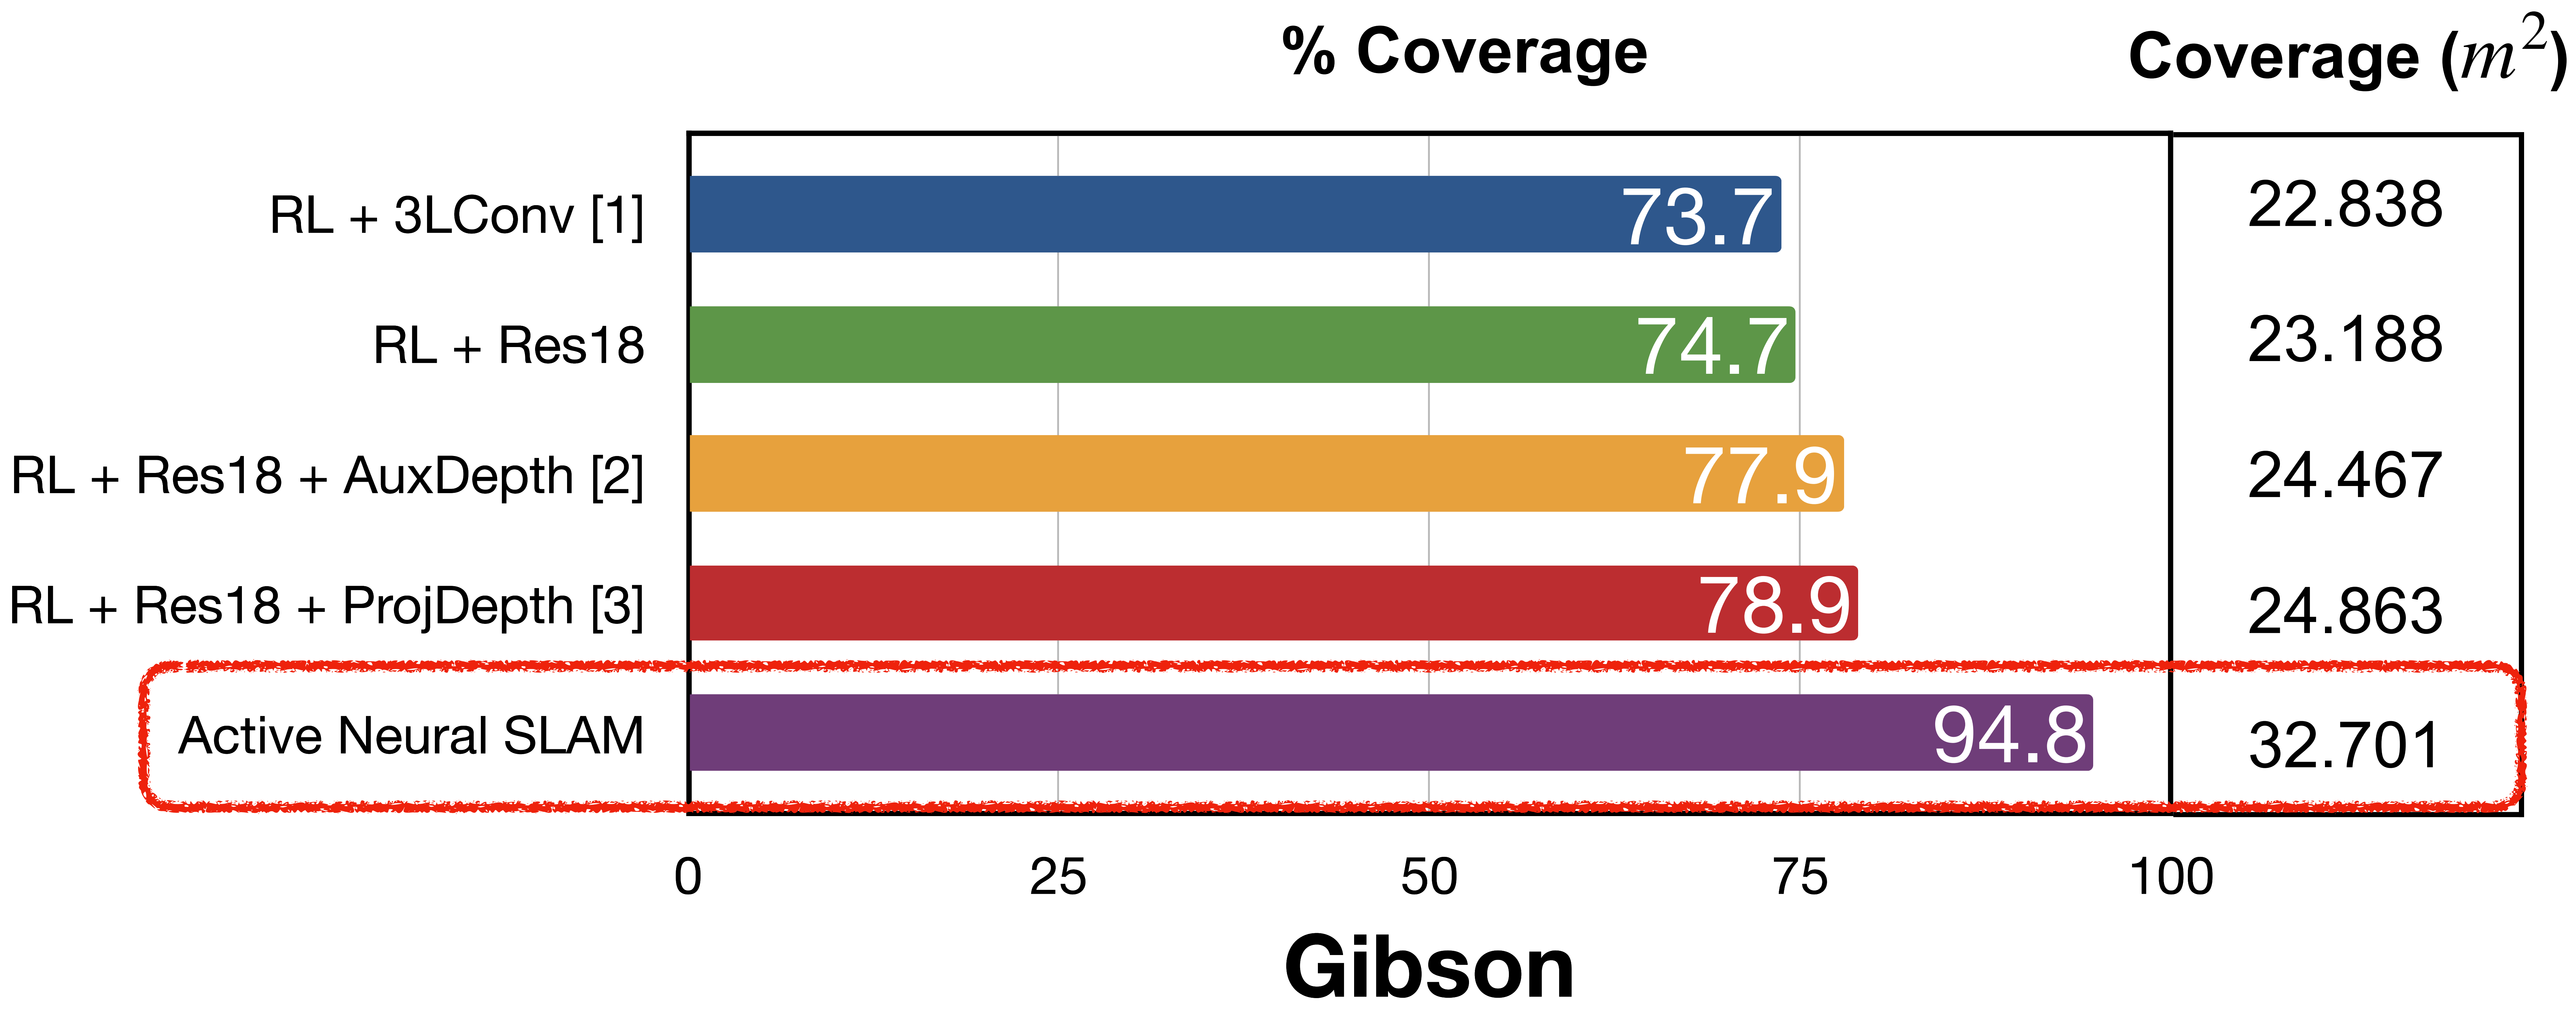

*Adapted from [1] Lample & Chaplot. AAAI-17, [2] Mirowski et al. ICLR-17, [3] Chen el al. ICLR-19

Exploration Results

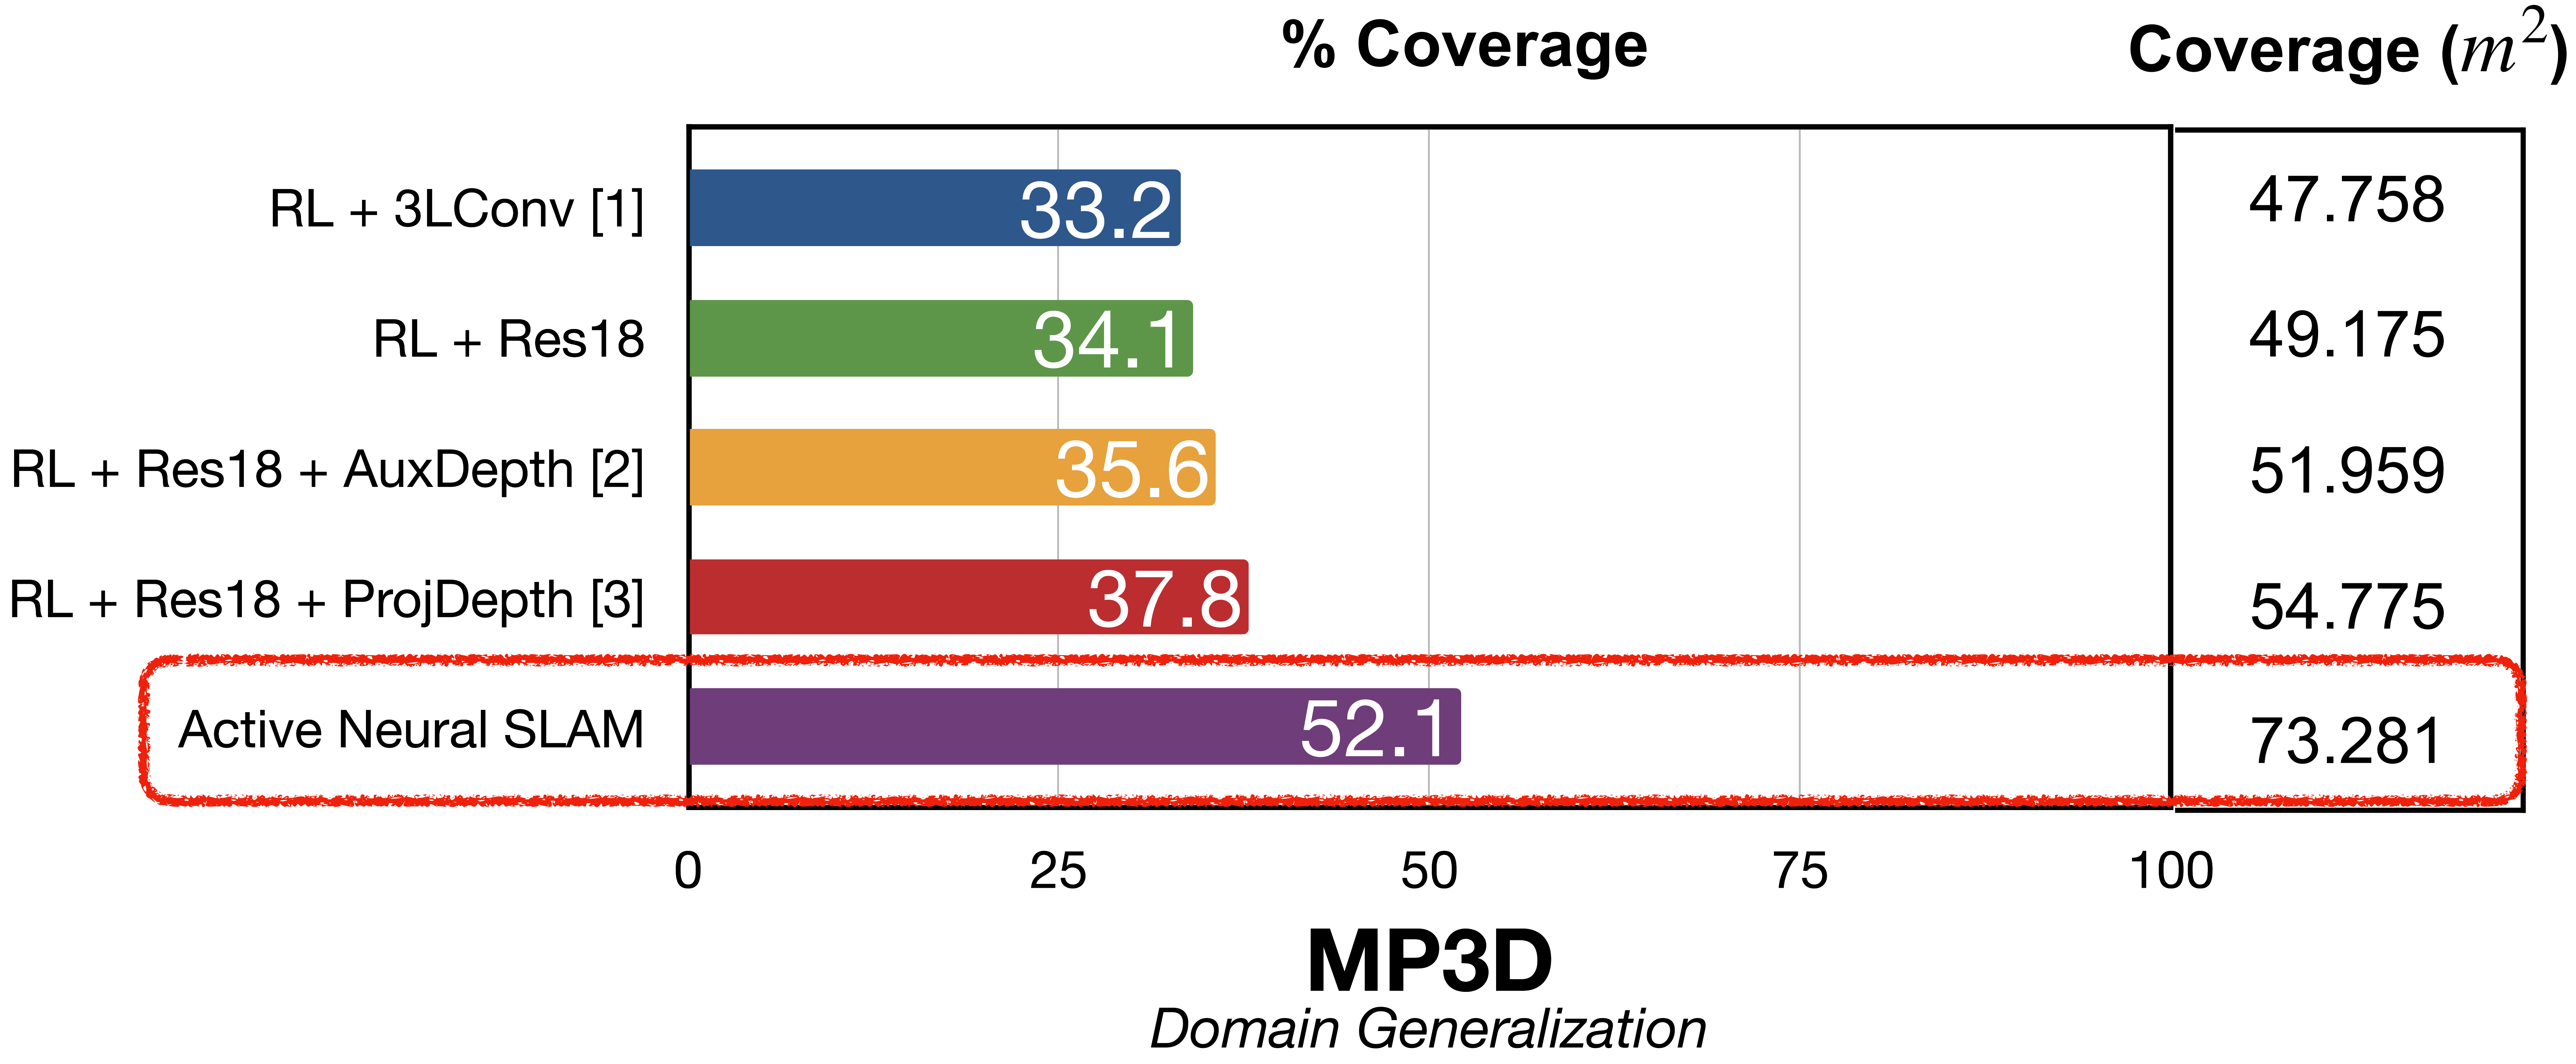

*Adapted from [1] Lample & Chaplot. AAAI-17, [2] Mirowski et al. ICLR-17, [3] Chen el al. ICLR-19

Ablation

Method	Gibson Val Overall % Cov.	Gibson Val Overall Cov. (m2)	Gibson Val Large % Cov.	Gibson Val Large Cov. (m2)	Gibson Val Small % Cov.	Gibson Val Small Cov. (m2)
ANS w/o Local Policy + Det. Planner	0.941	32.188	0.845	53.999	0.980	23.464
ANS w/o Global Policy + FBE	0.925	30.981	0.782	49.731	0.982	23.481
ANS w/o Pose Estimation	0.916	30.746	0.771	49.518	0.973	23.237
ANS	**0.948**	**32.701**	**0.862**	**55.608**	**0.983**	**23.538**

Ablation

Method	Gibson Val Overall % Cov.	Gibson Val Overall Cov. (m2)	Gibson Val Large % Cov.	Gibson Val Large Cov. (m2)	Gibson Val Small % Cov.	Gibson Val Small Cov. (m2)
ANS w/o Local Policy + Det. Planner	0.941	32.188	0.845	53.999	0.980	23.464
ANS w/o Global Policy + FBE	0.925	30.981	0.782	49.731	0.982	23.481
ANS w/o Pose Estimation	0.916	30.746	0.771	49.518	0.973	23.237
ANS	**0.948**	**32.701**	**0.862**	**55.608**	**0.983**	**23.538**

Replace Local Policy by
Analytical Deterministic Policy

Ablation

Method	Gibson Val Overall % Cov.	Gibson Val Overall Cov. (m2)	Gibson Val Large % Cov.	Gibson Val Large Cov. (m2)	Gibson Val Small % Cov.	Gibson Val Small Cov. (m2)
ANS w/o Local Policy + Det. Planner	0.941	32.188	0.845	53.999	0.980	23.464
ANS w/o Global Policy + FBE	0.925	30.981	0.782	49.731	0.982	23.481
ANS w/o Pose Estimation	0.916	30.746	0.771	49.518	0.973	23.237
ANS	**0.948**	**32.701**	**0.862**	**55.608**	**0.983**	**23.538**

Replace Local Policy by Analytical Deterministic Policy

Replace Global Policy by Frontier-based Exploration

Ablation

Local Policy does not improve much over deterministic policy

Method	Gibson Val Overall % Cov.	Gibson Val Overall Cov. (m2)	Gibson Val Large % Cov.	Gibson Val Large Cov. (m2)	Gibson Val Small % Cov.	Gibson Val Small Cov. (m2)
ANS w/o Local Policy + Det. Planner	0.941	32.188	0.845	53.999	0.980	23.464
ANS w/o Global Policy + FBE	0.925	30.981	0.782	49.731	0.982	23.481
ANS w/o Pose Estimation	0.916	30.746	0.771	49.518	0.973	23.237
ANS	**0.948**	**32.701**	**0.862**	**55.608**	**0.983**	**23.538**

Replace Local Policy by Analytical Deterministic Policy

Replace Global Policy by Frontier-based Exploration

Ablation

Local Policy does not improve much over deterministic policy

Method	Gibson Val Overall		Gibson Val Large		Gibson Val Small	
	% Cov.	Cov. (m2)	% Cov.	Cov. (m2)	% Cov.	Cov. (m2)
ANS w/o Local Policy + Det. Planner	0.941	32.188	0.845	53.999	0.980	23.464
ANS w/o Global Policy + FBE	0.925	30.981	0.782	49.731	0.982	23.481
ANS w/o Pose Estimation	0.916	30.746	0.771	49.518	0.973	23.237
ANS	**0.948**	**32.701**	**0.862**	**55.608**	**0.983**	**23.538**

Replace Local Policy by Analytical Deterministic Policy

Replace Global Policy by Frontier-based Exploration

Global Policy and Pose Estimation mostly help in Large maps

Pointgoal: Task Transfer

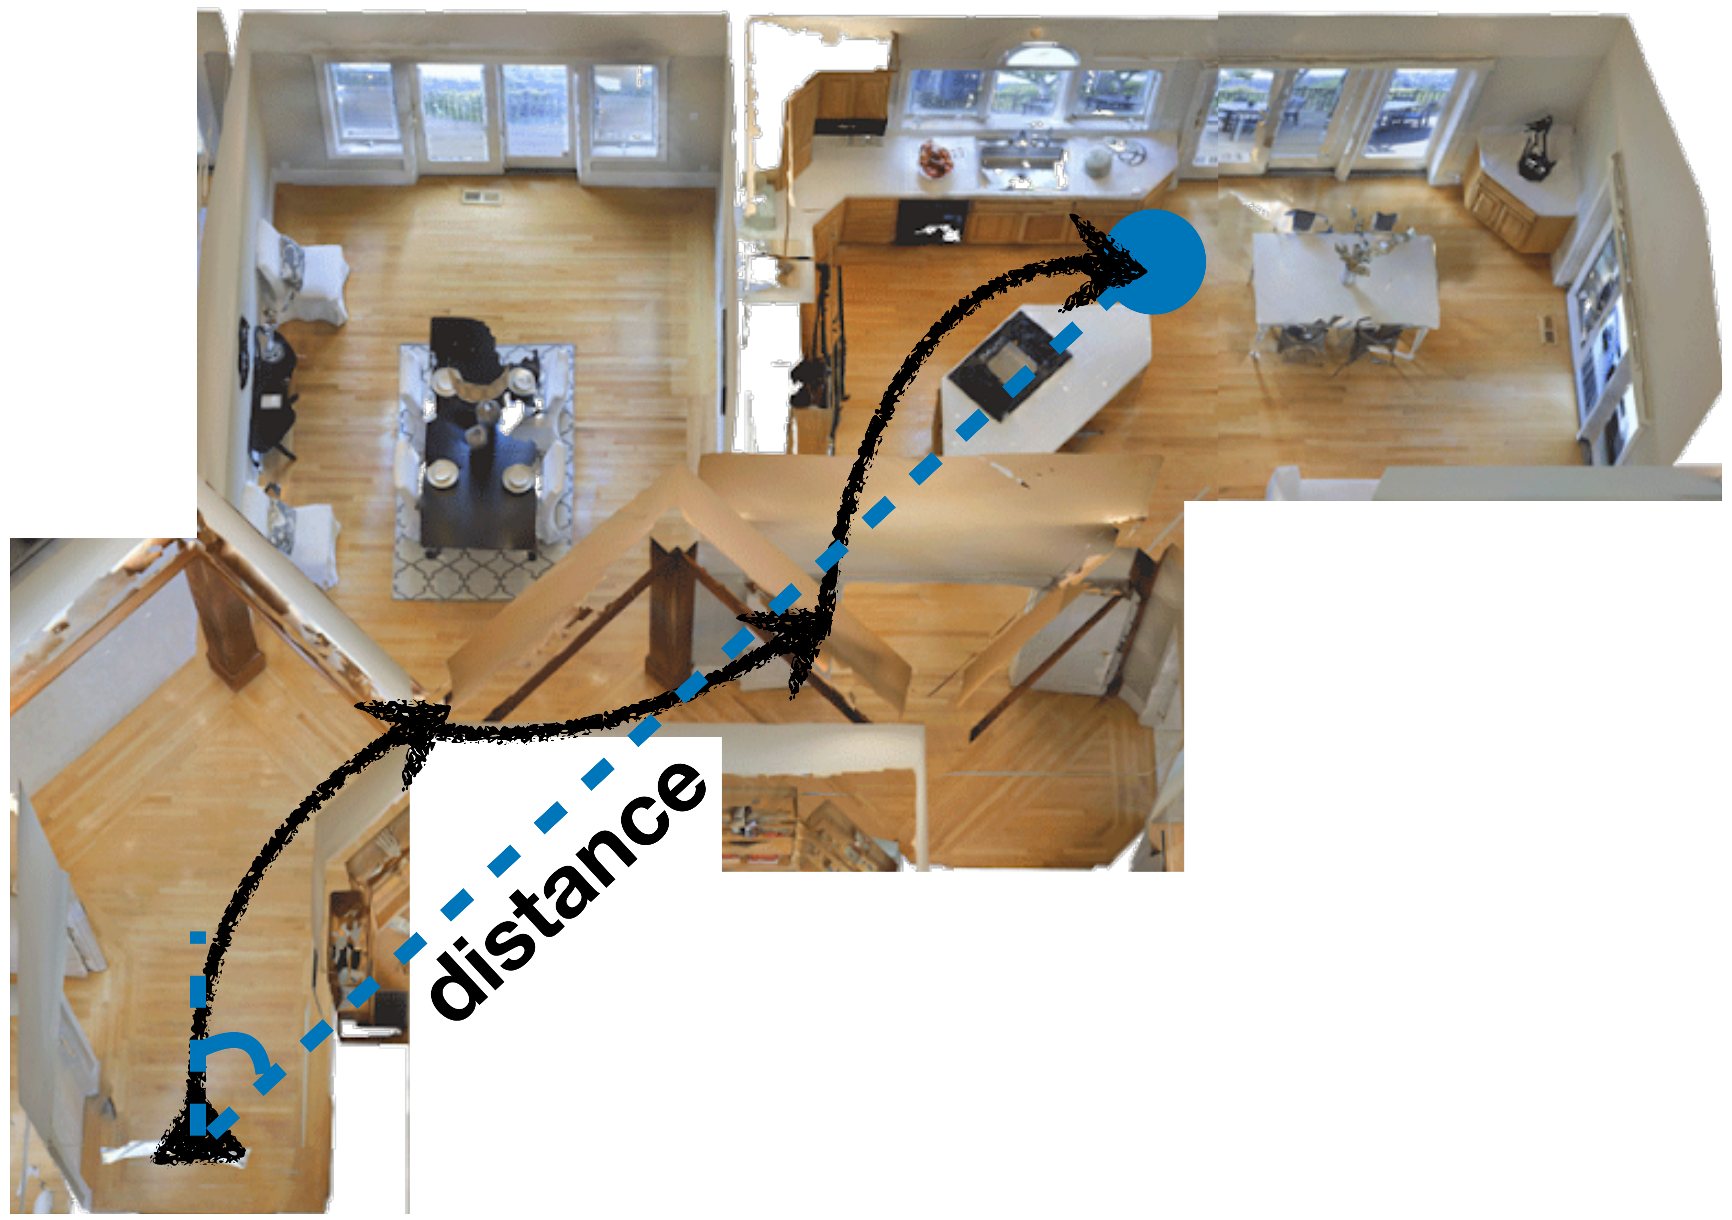

Pointgoal: Task Transfer

- Objective: Navigate to goal coordinates

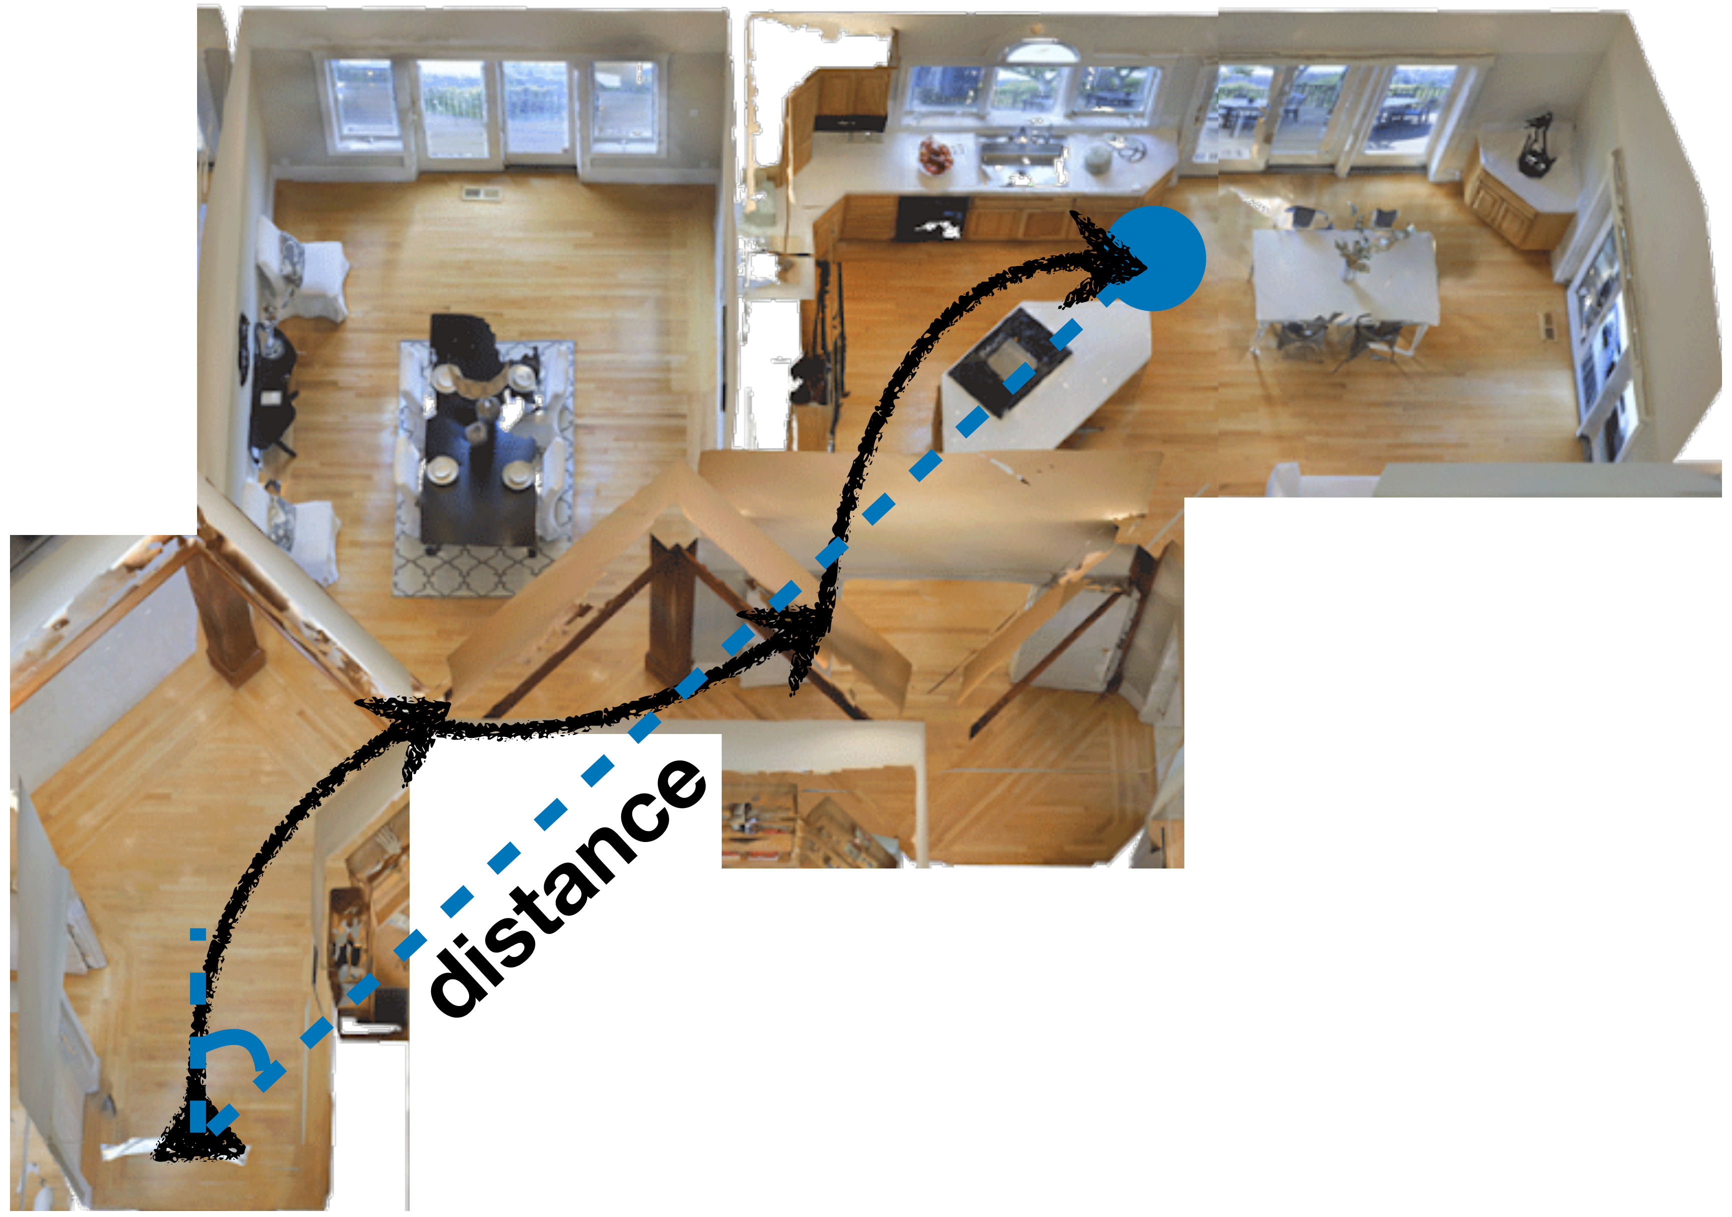

Pointgoal: Task Transfer

- Objective: Navigate to goal coordinates
- Metric: Success weighted by inverse Path Length (SPL)

$$\frac{1}{N}\sum_{i=1}^{N} Success * \frac{ShortestPathLength}{PathLength}$$

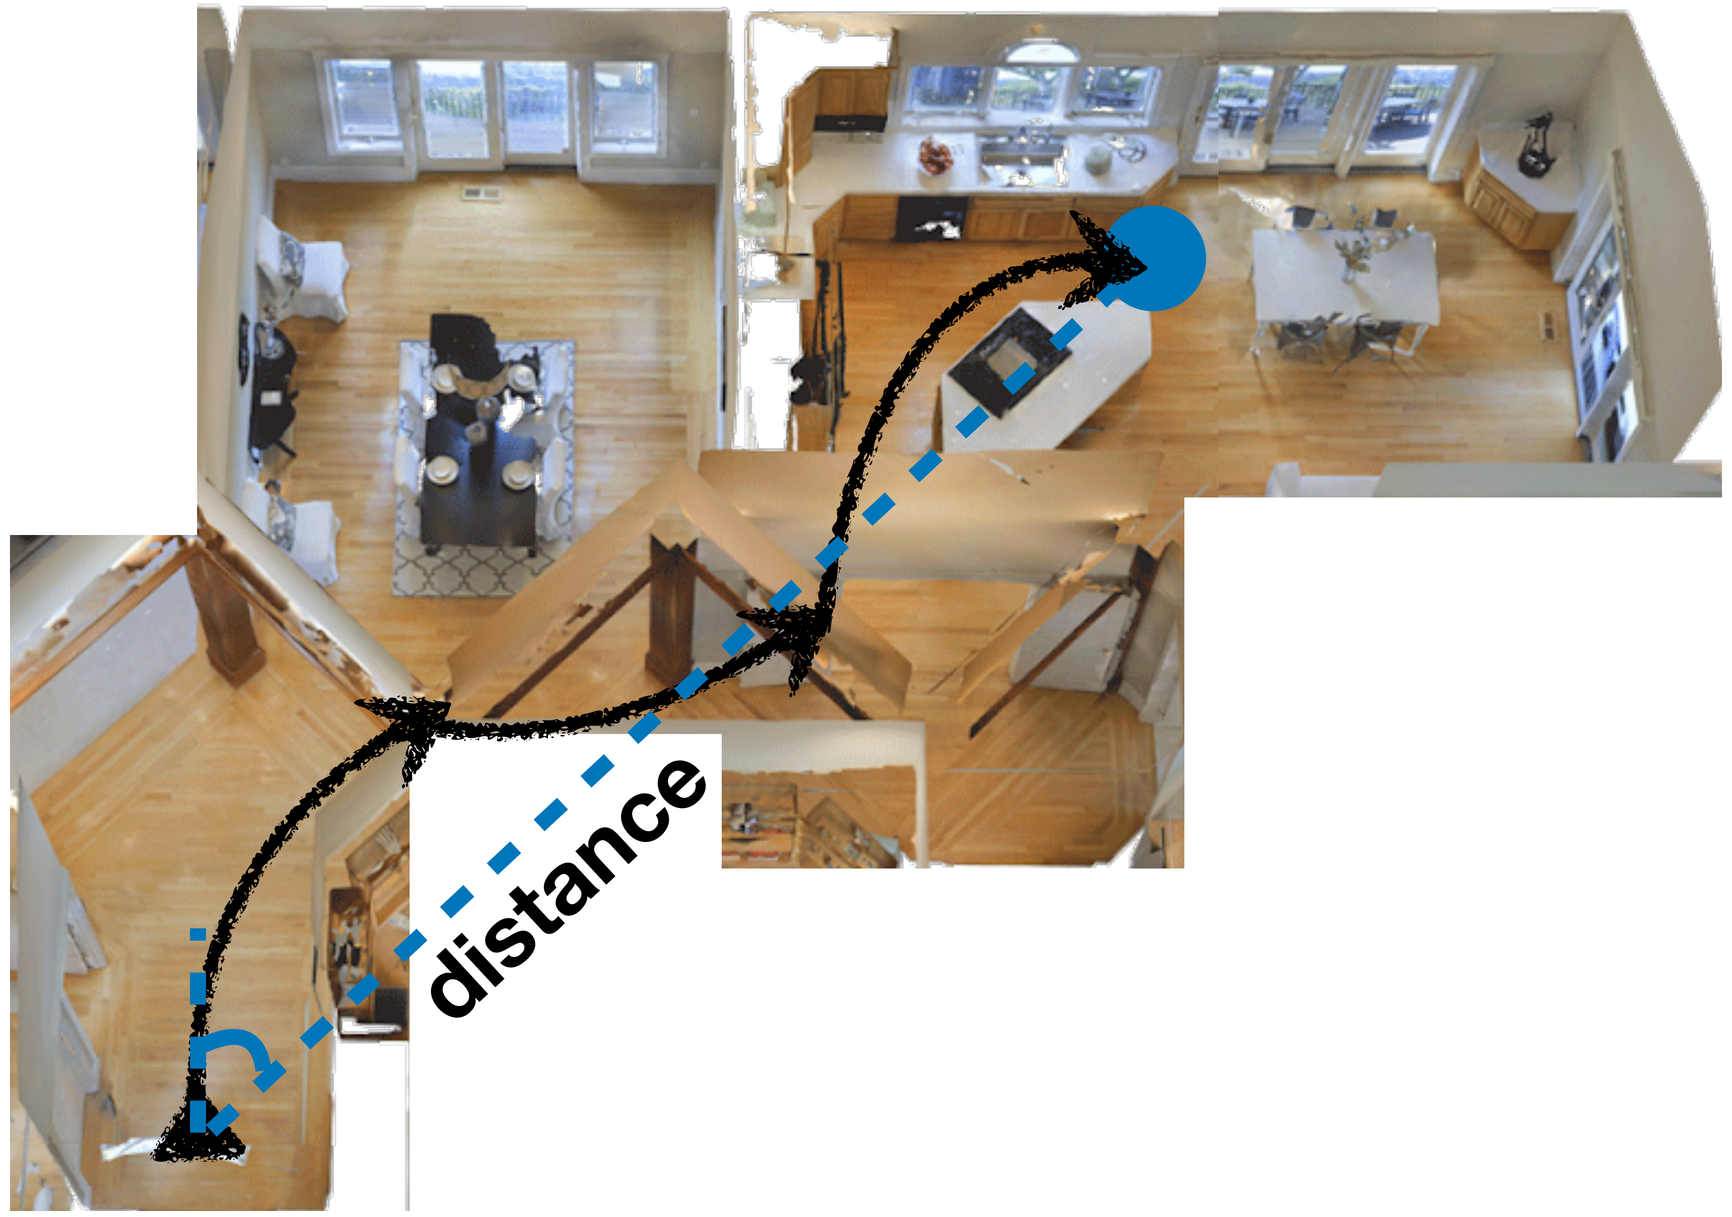

Pointgoal: Task Transfer

- Objective: Navigate to goal coordinates
- Metric: Success weighted by inverse Path Length (SPL)

$$\frac{1}{N}\sum_{i=1}^{N} Success * \frac{ShortestPathLength}{PathLength}$$

- Global Policy -> always gives the pointgoal as the long-term goal

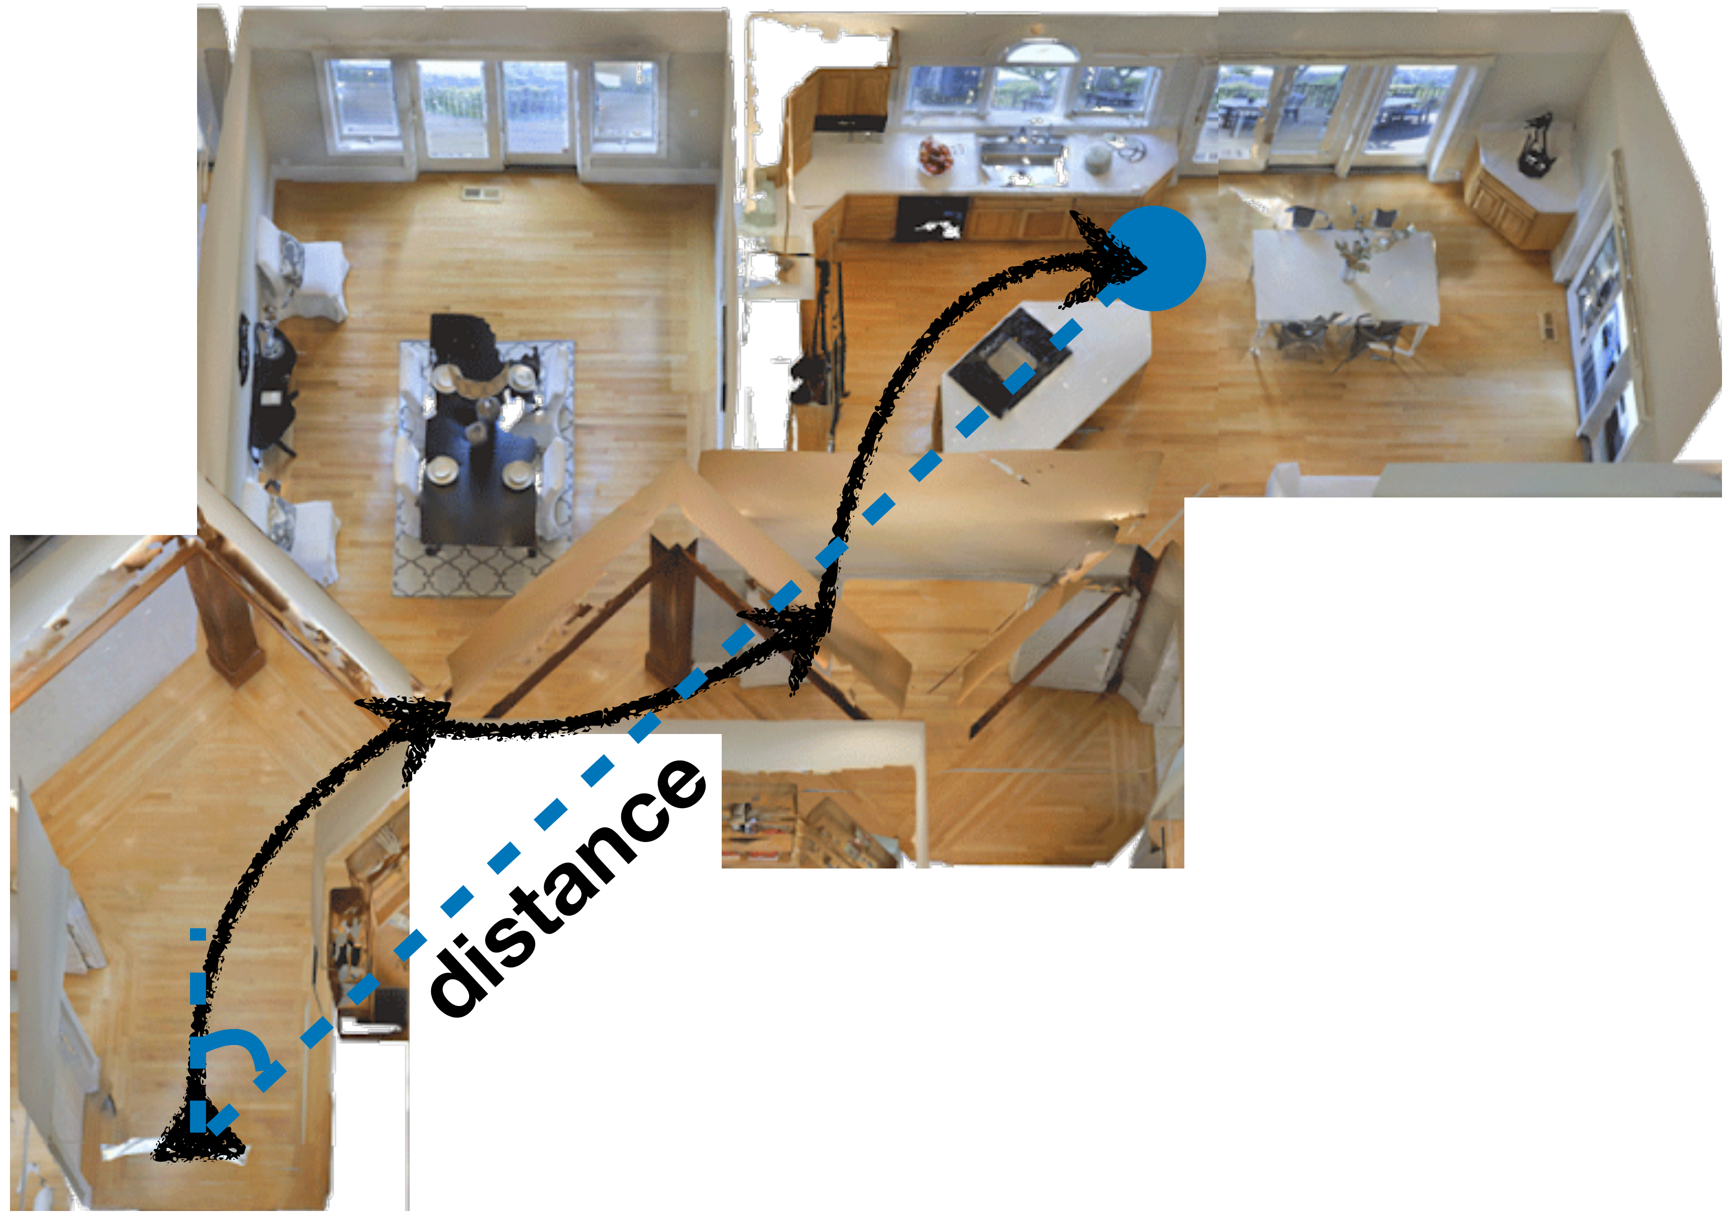

Harder Datasets

- **Hard-GEDR**
  - Higher Geodesic to Euclidean distance ratio (GEDR)
  - Avg GEDR 2.5 vs 1.37, minimum GEDR is 2

- **Hard-Dist**
  - Higher Geodesic distance
  - Avg Dist 13.5m vs 7.0m, minimum Dist is 10m

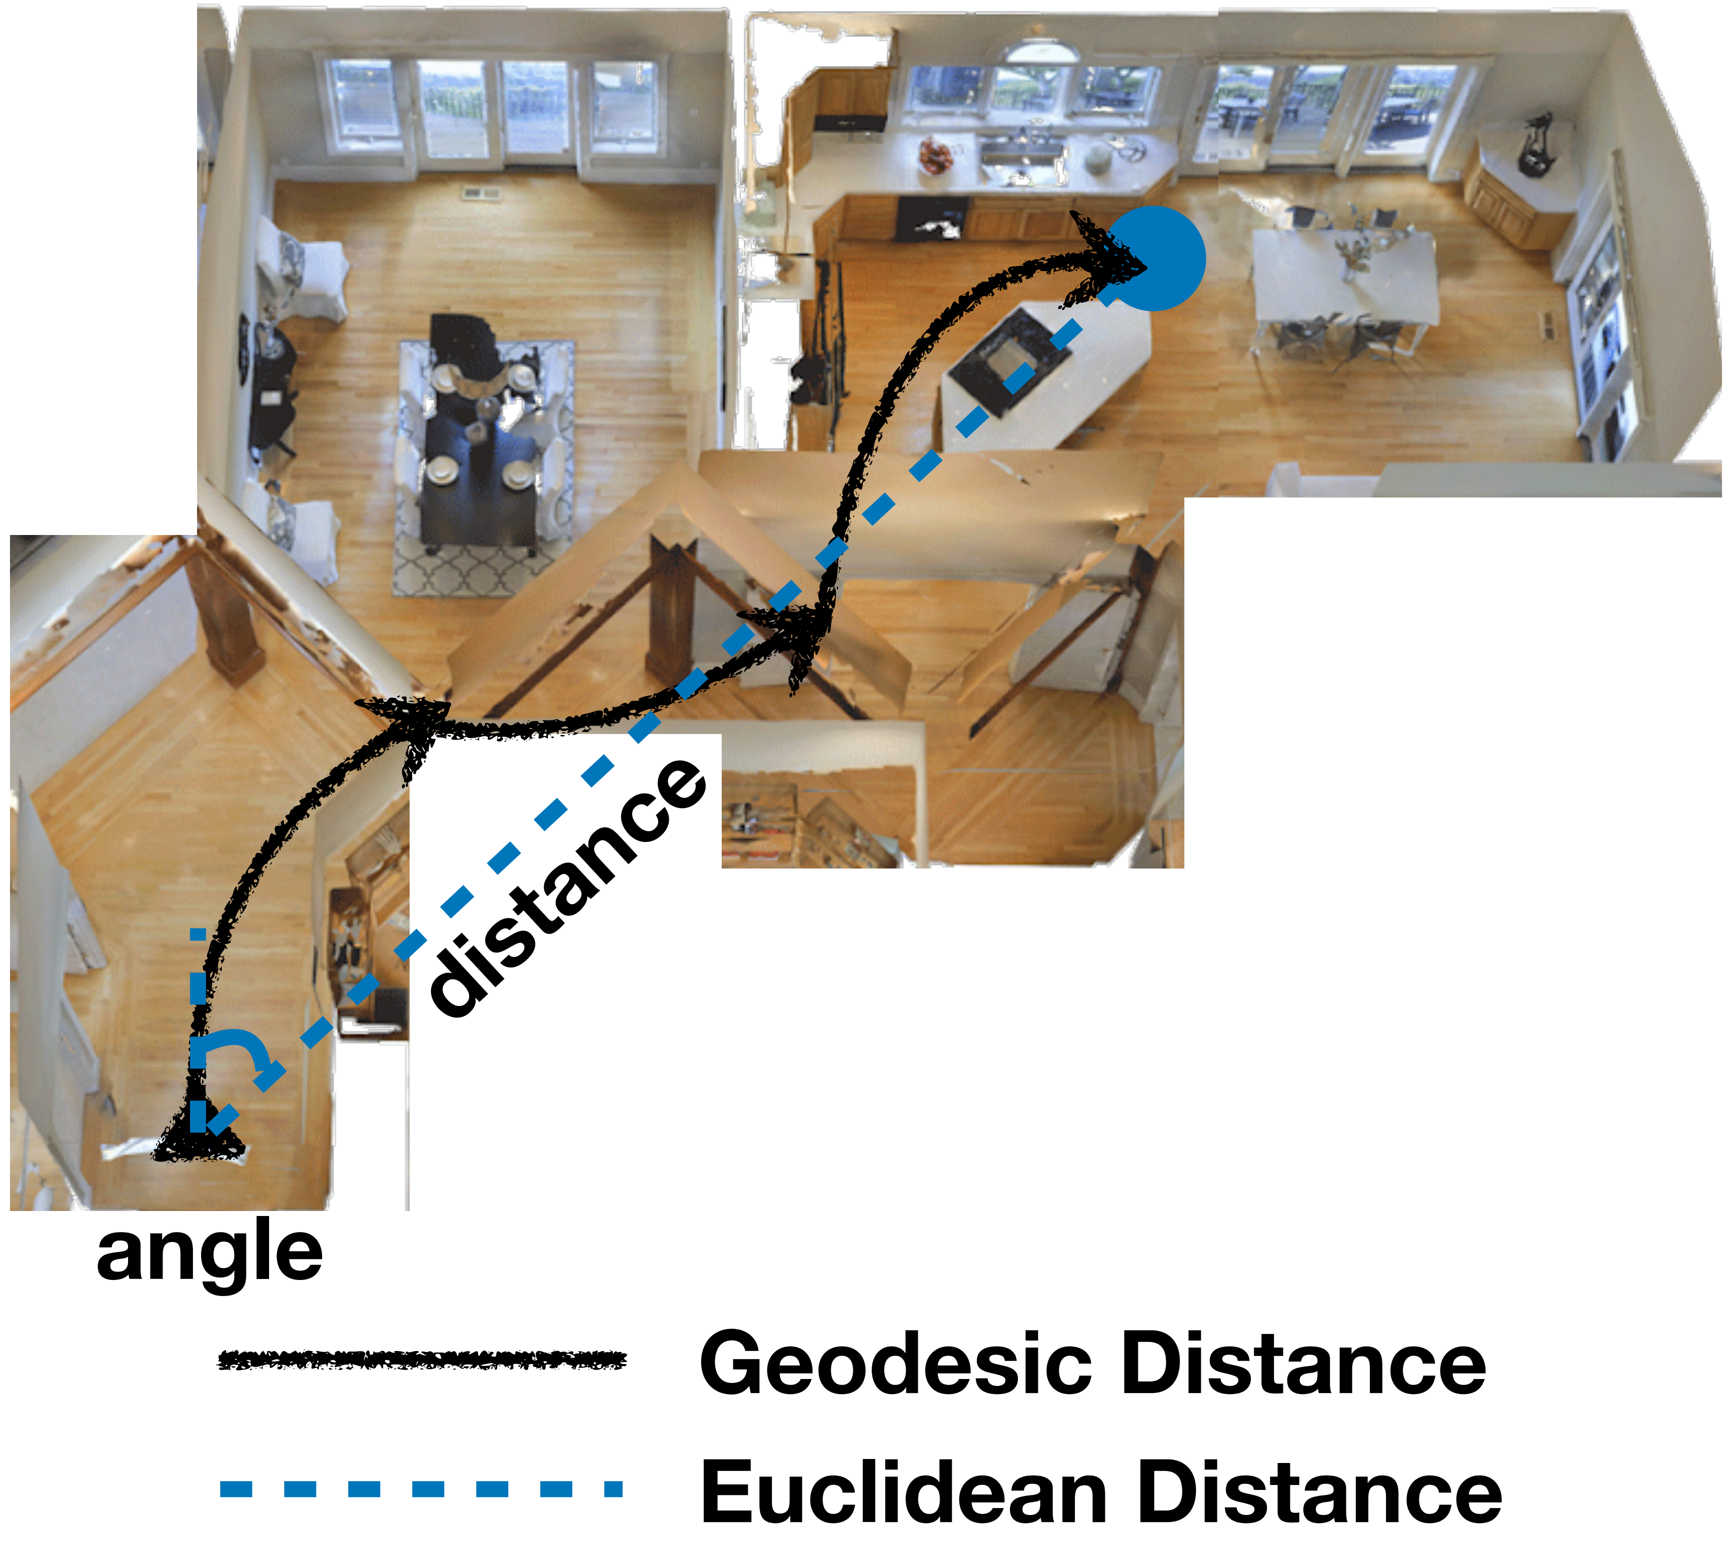

PointGoal Results

PointGoal Results

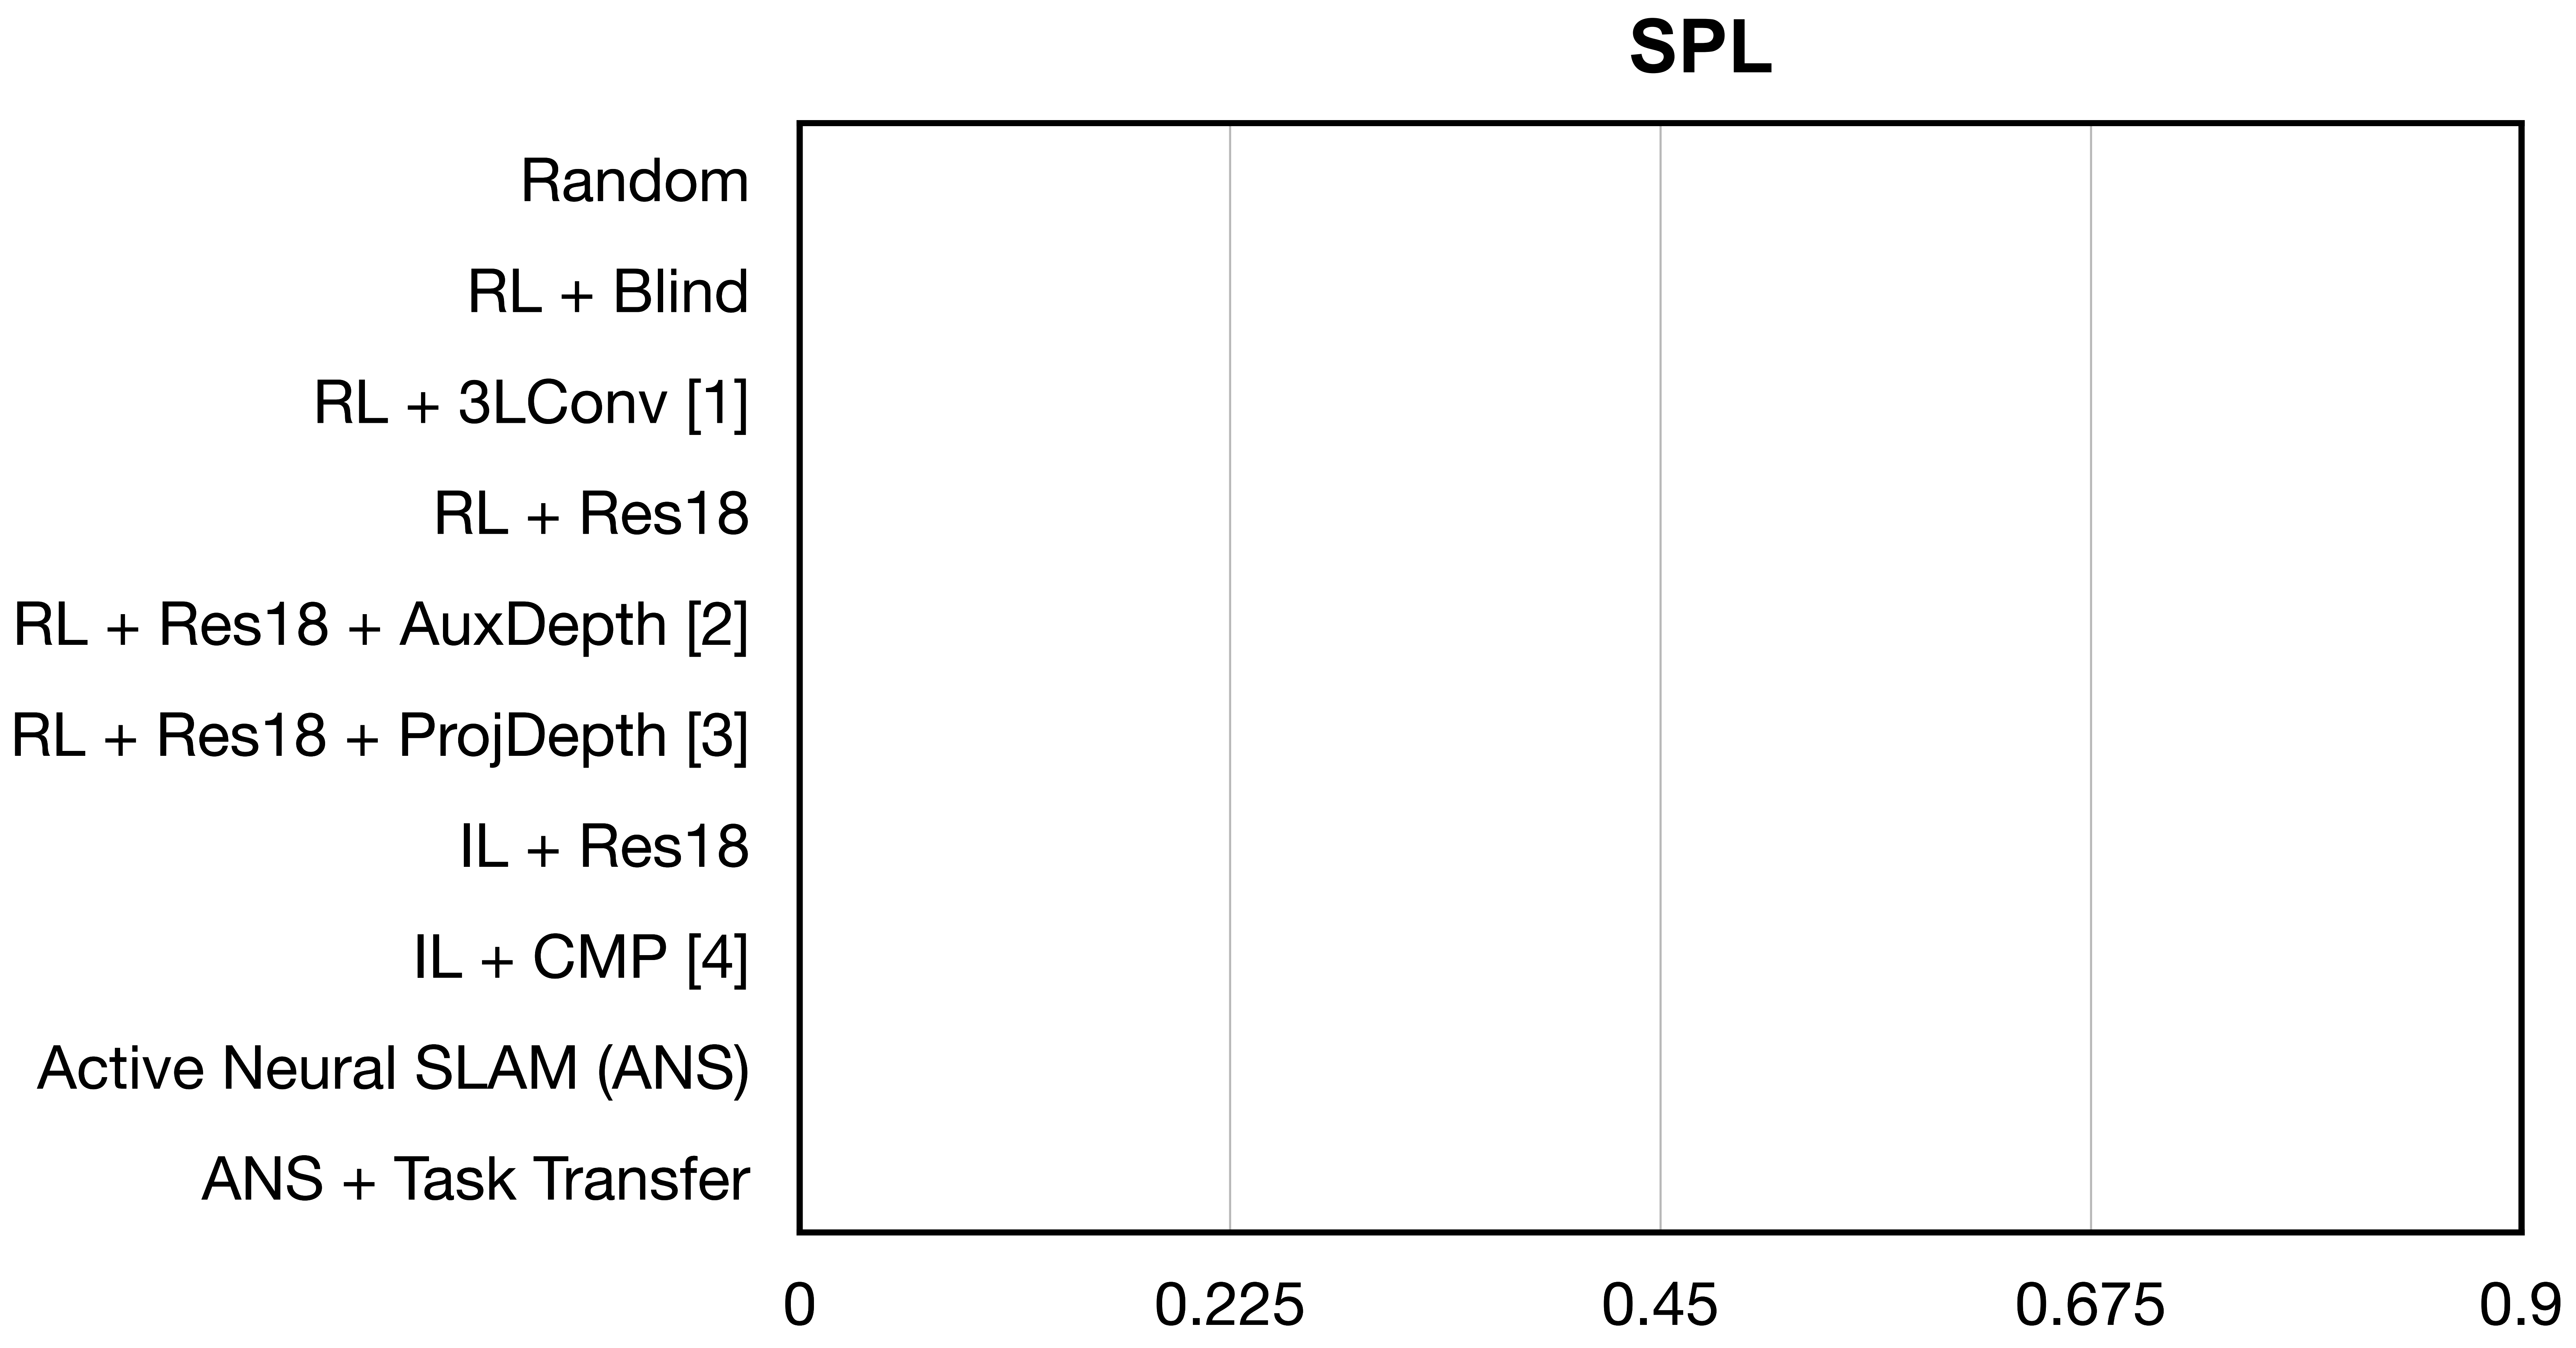

**Adapted from [1] Lample & Chaplot. AAAI-17, [2] Mirowski et al. ICLR-17, [3] Chen el al. ICLR-19, [4] Gupta et al. CVPR-17*

PointGoal Results

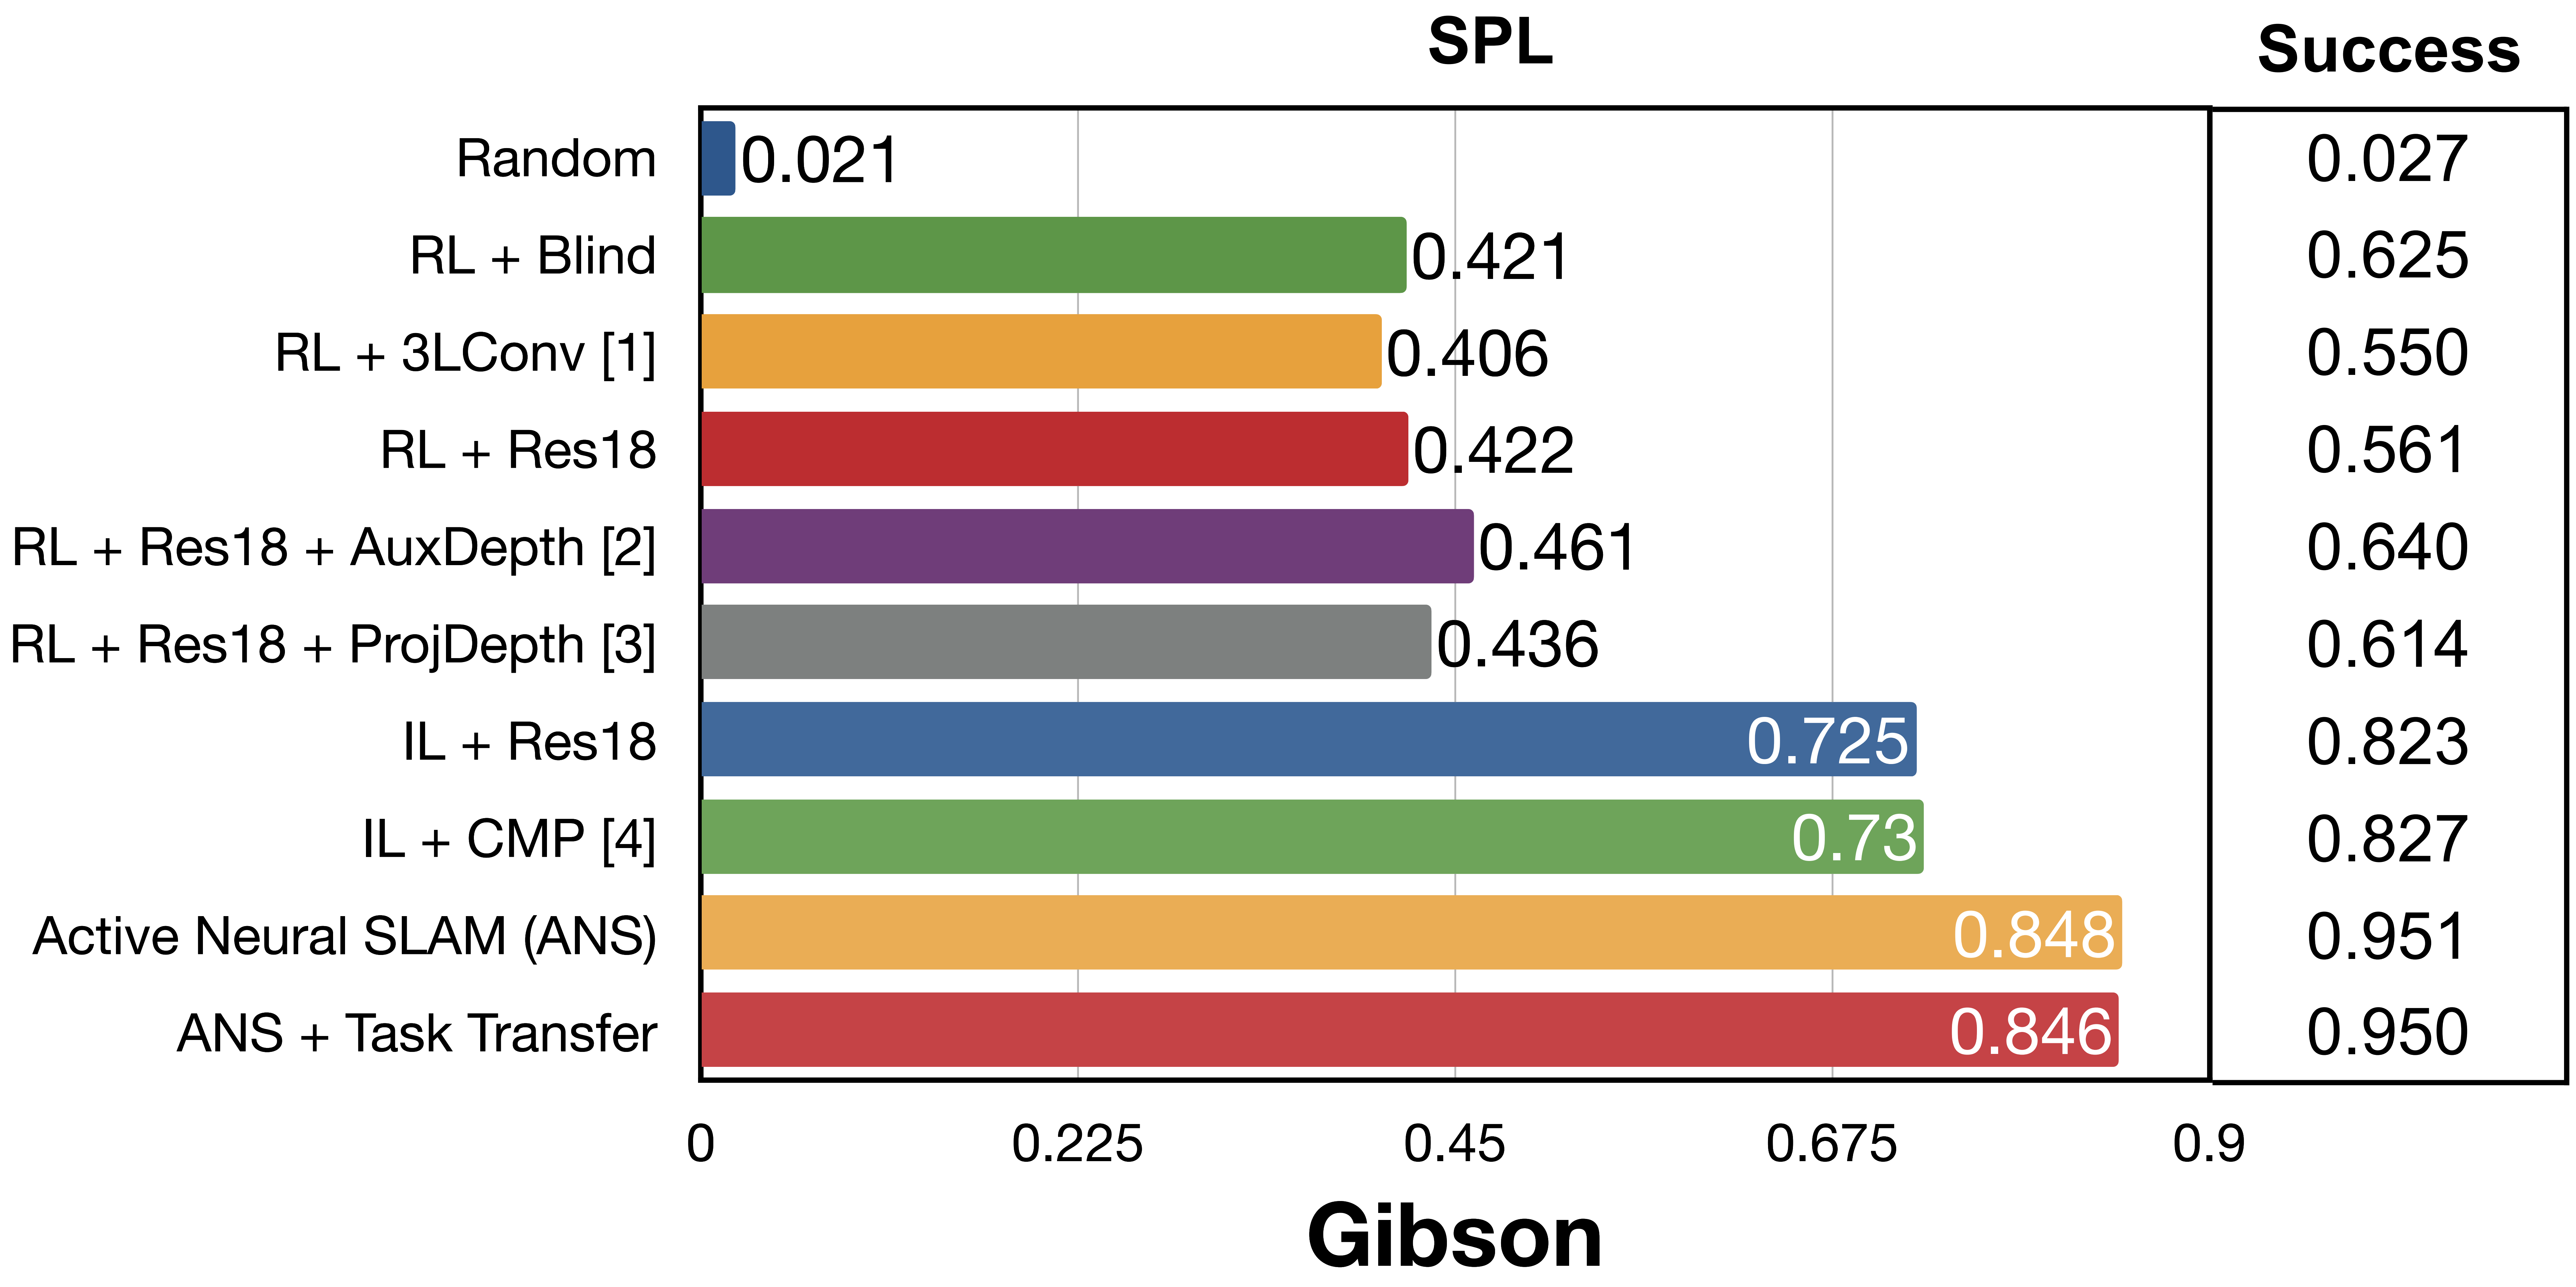

**Adapted from [1] Lample & Chaplot. AAAI-17, [2] Mirowski et al. ICLR-17, [3] Chen el al. ICLR-19, [4] Gupta et al. CVPR-17*

PointGoal Results

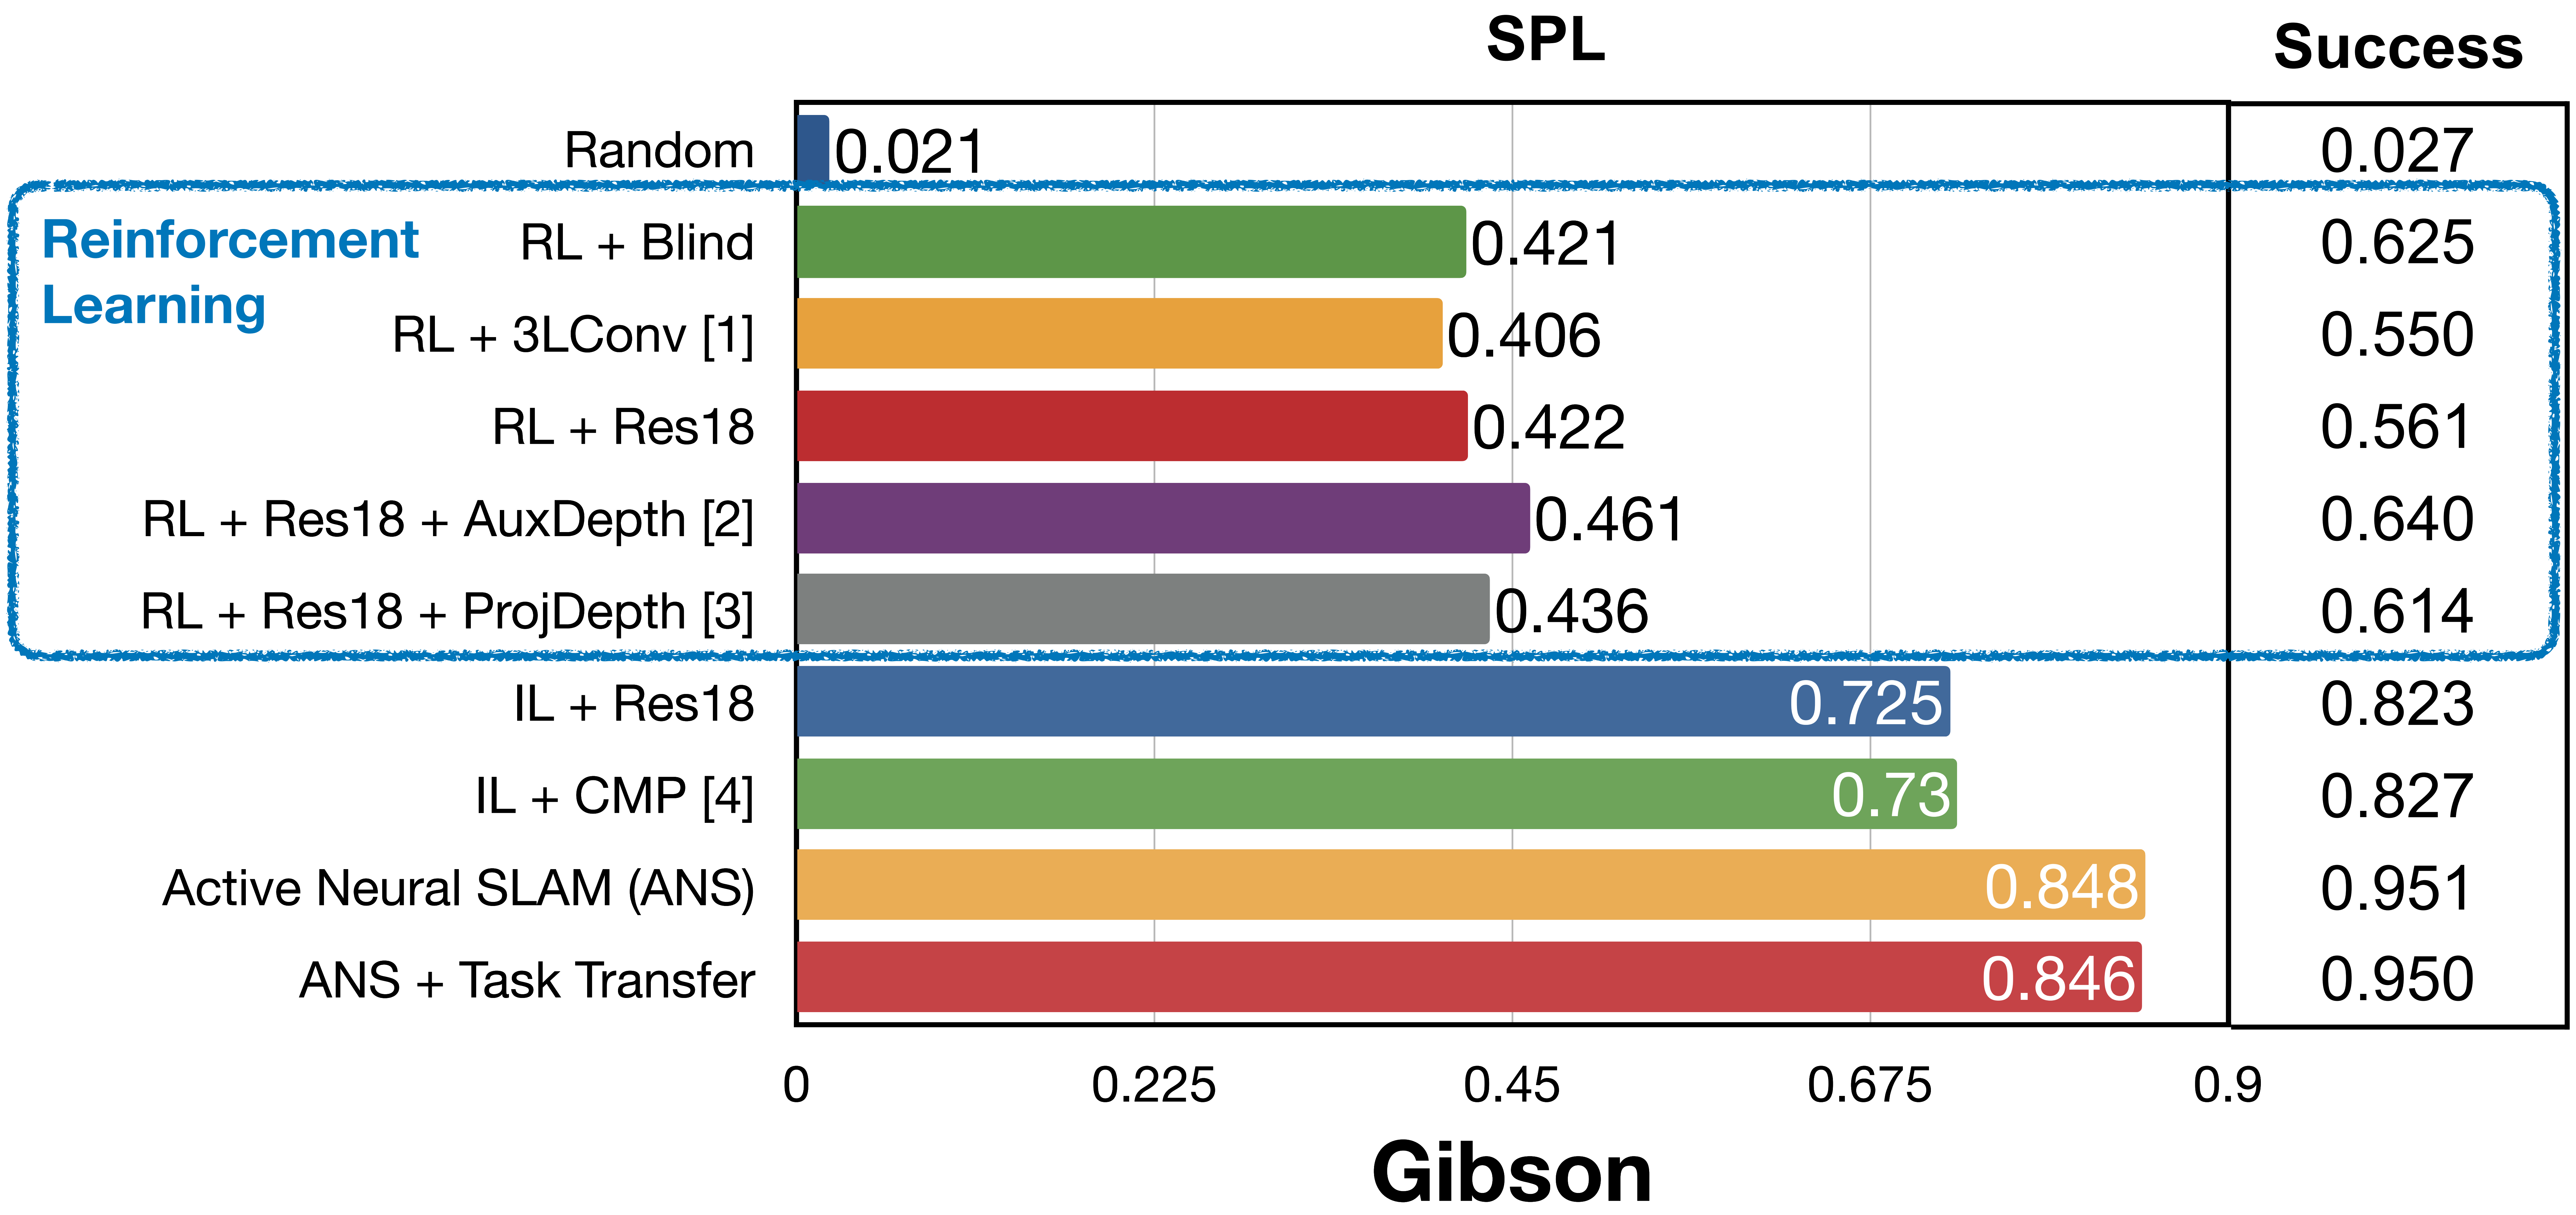

Adapted from [1] Lample & Chaplot. AAAI-17, [2] Mirowski et al. ICLR-17, [3] Chen el al. ICLR-19, [4] Gupta et al. CVPR-17

PointGoal Results

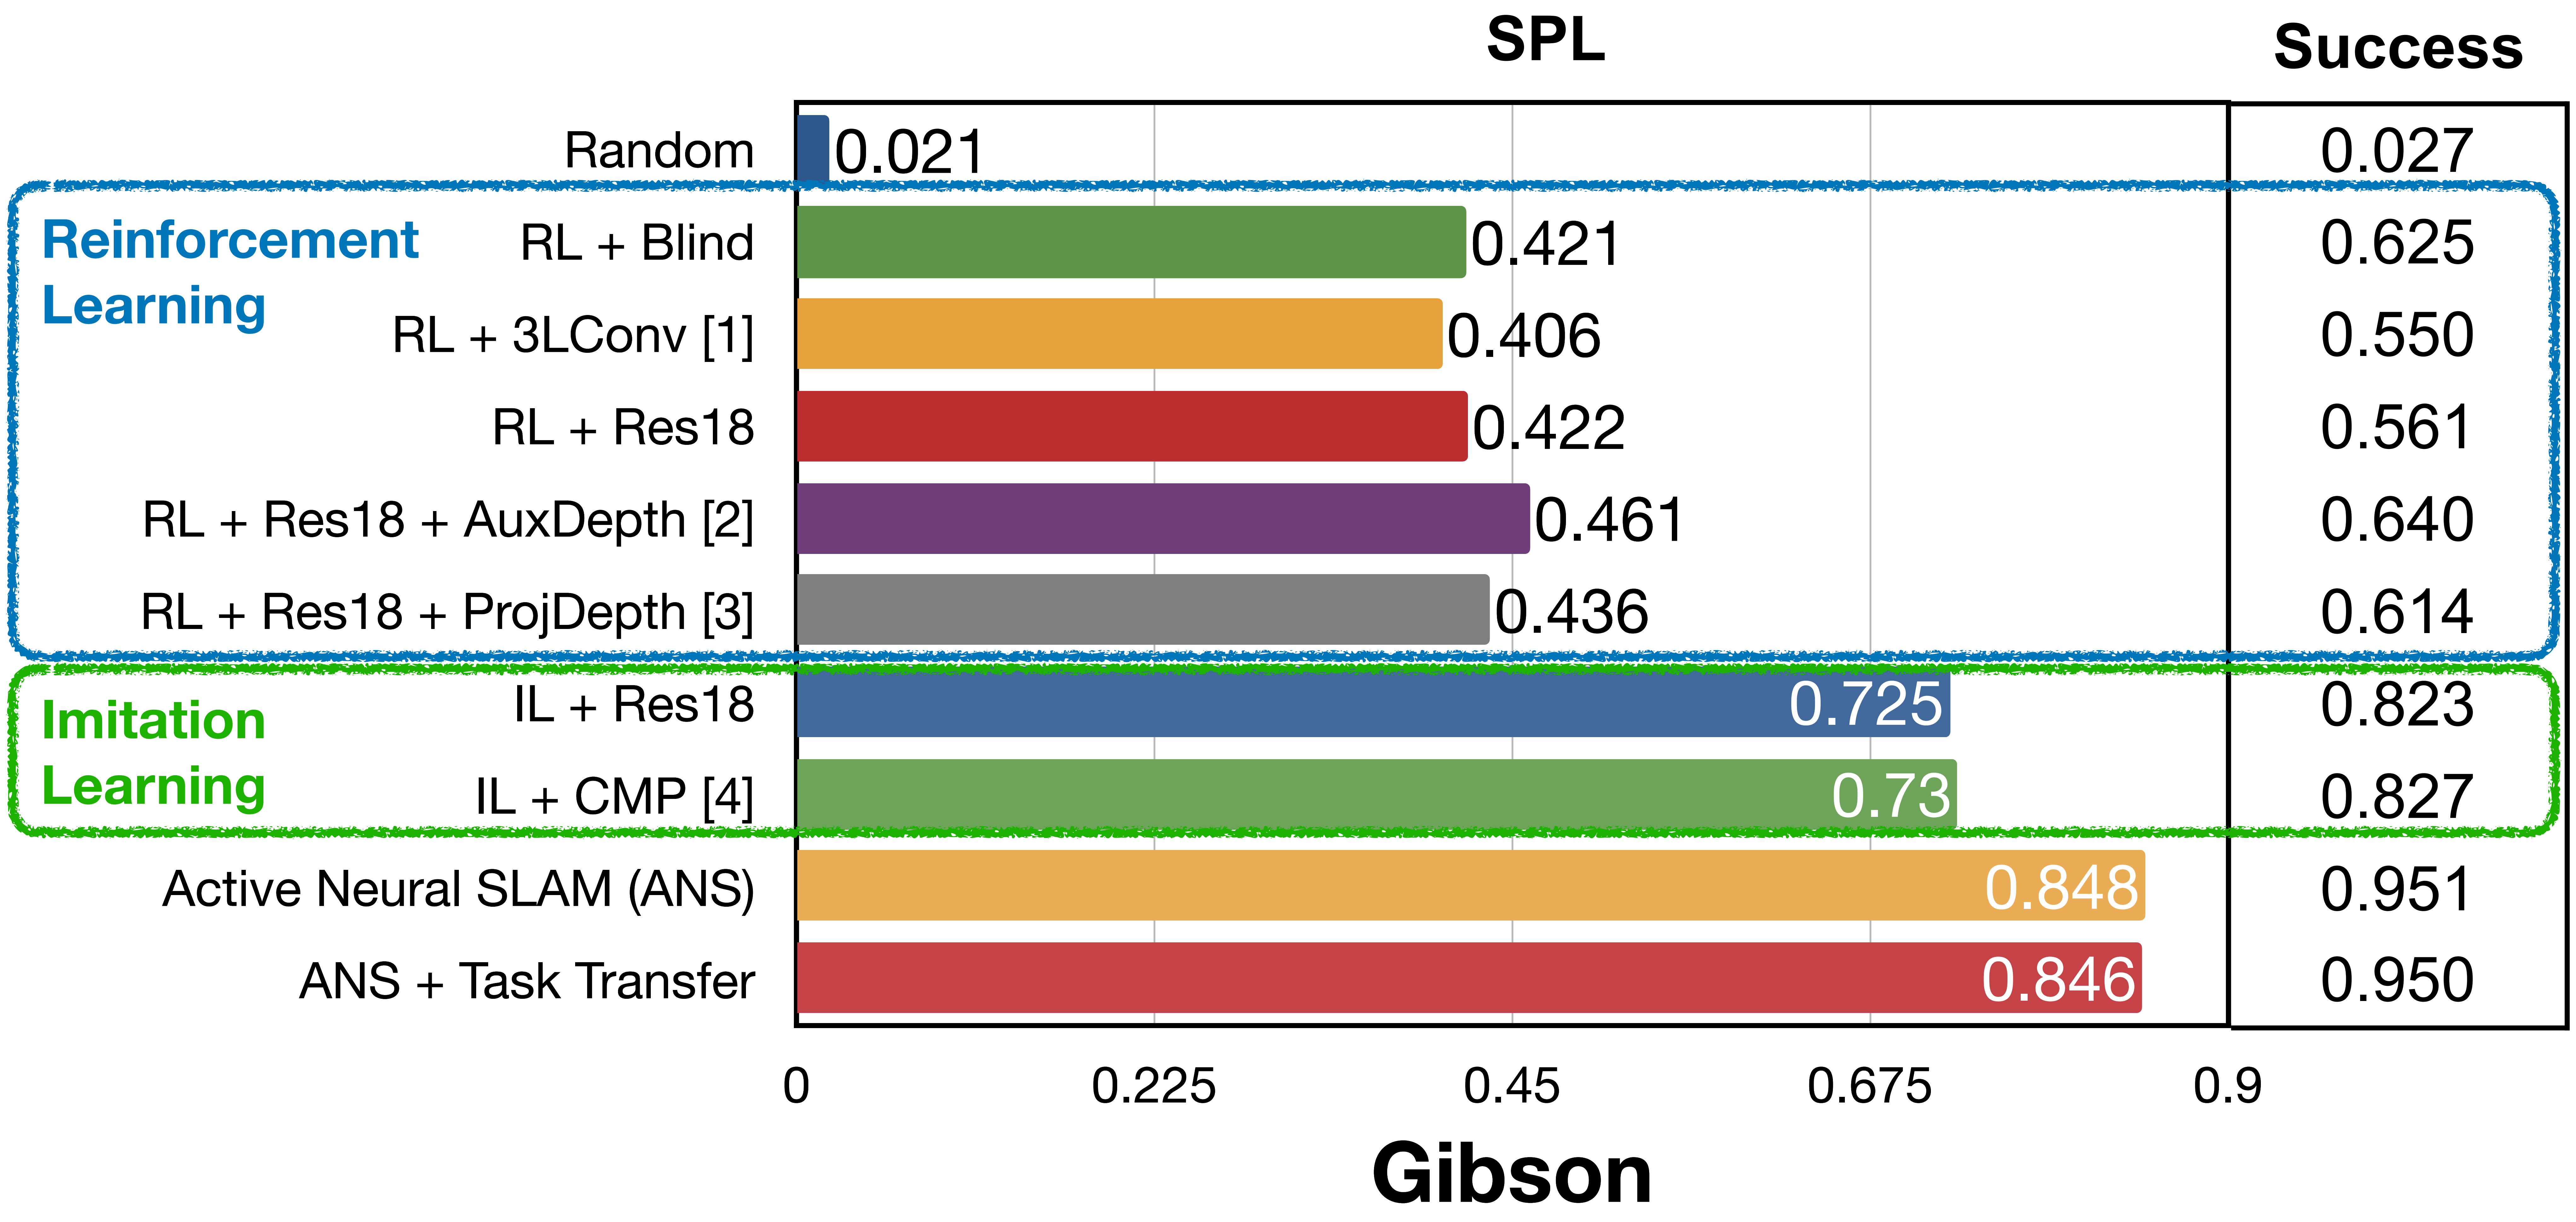

*Adapted from [1] Lample & Chaplot. AAAI-17, [2] Mirowski et al. ICLR-17, [3] Chen el al. ICLR-19, [4] Gupta et al. CVPR-17

PointGoal Results

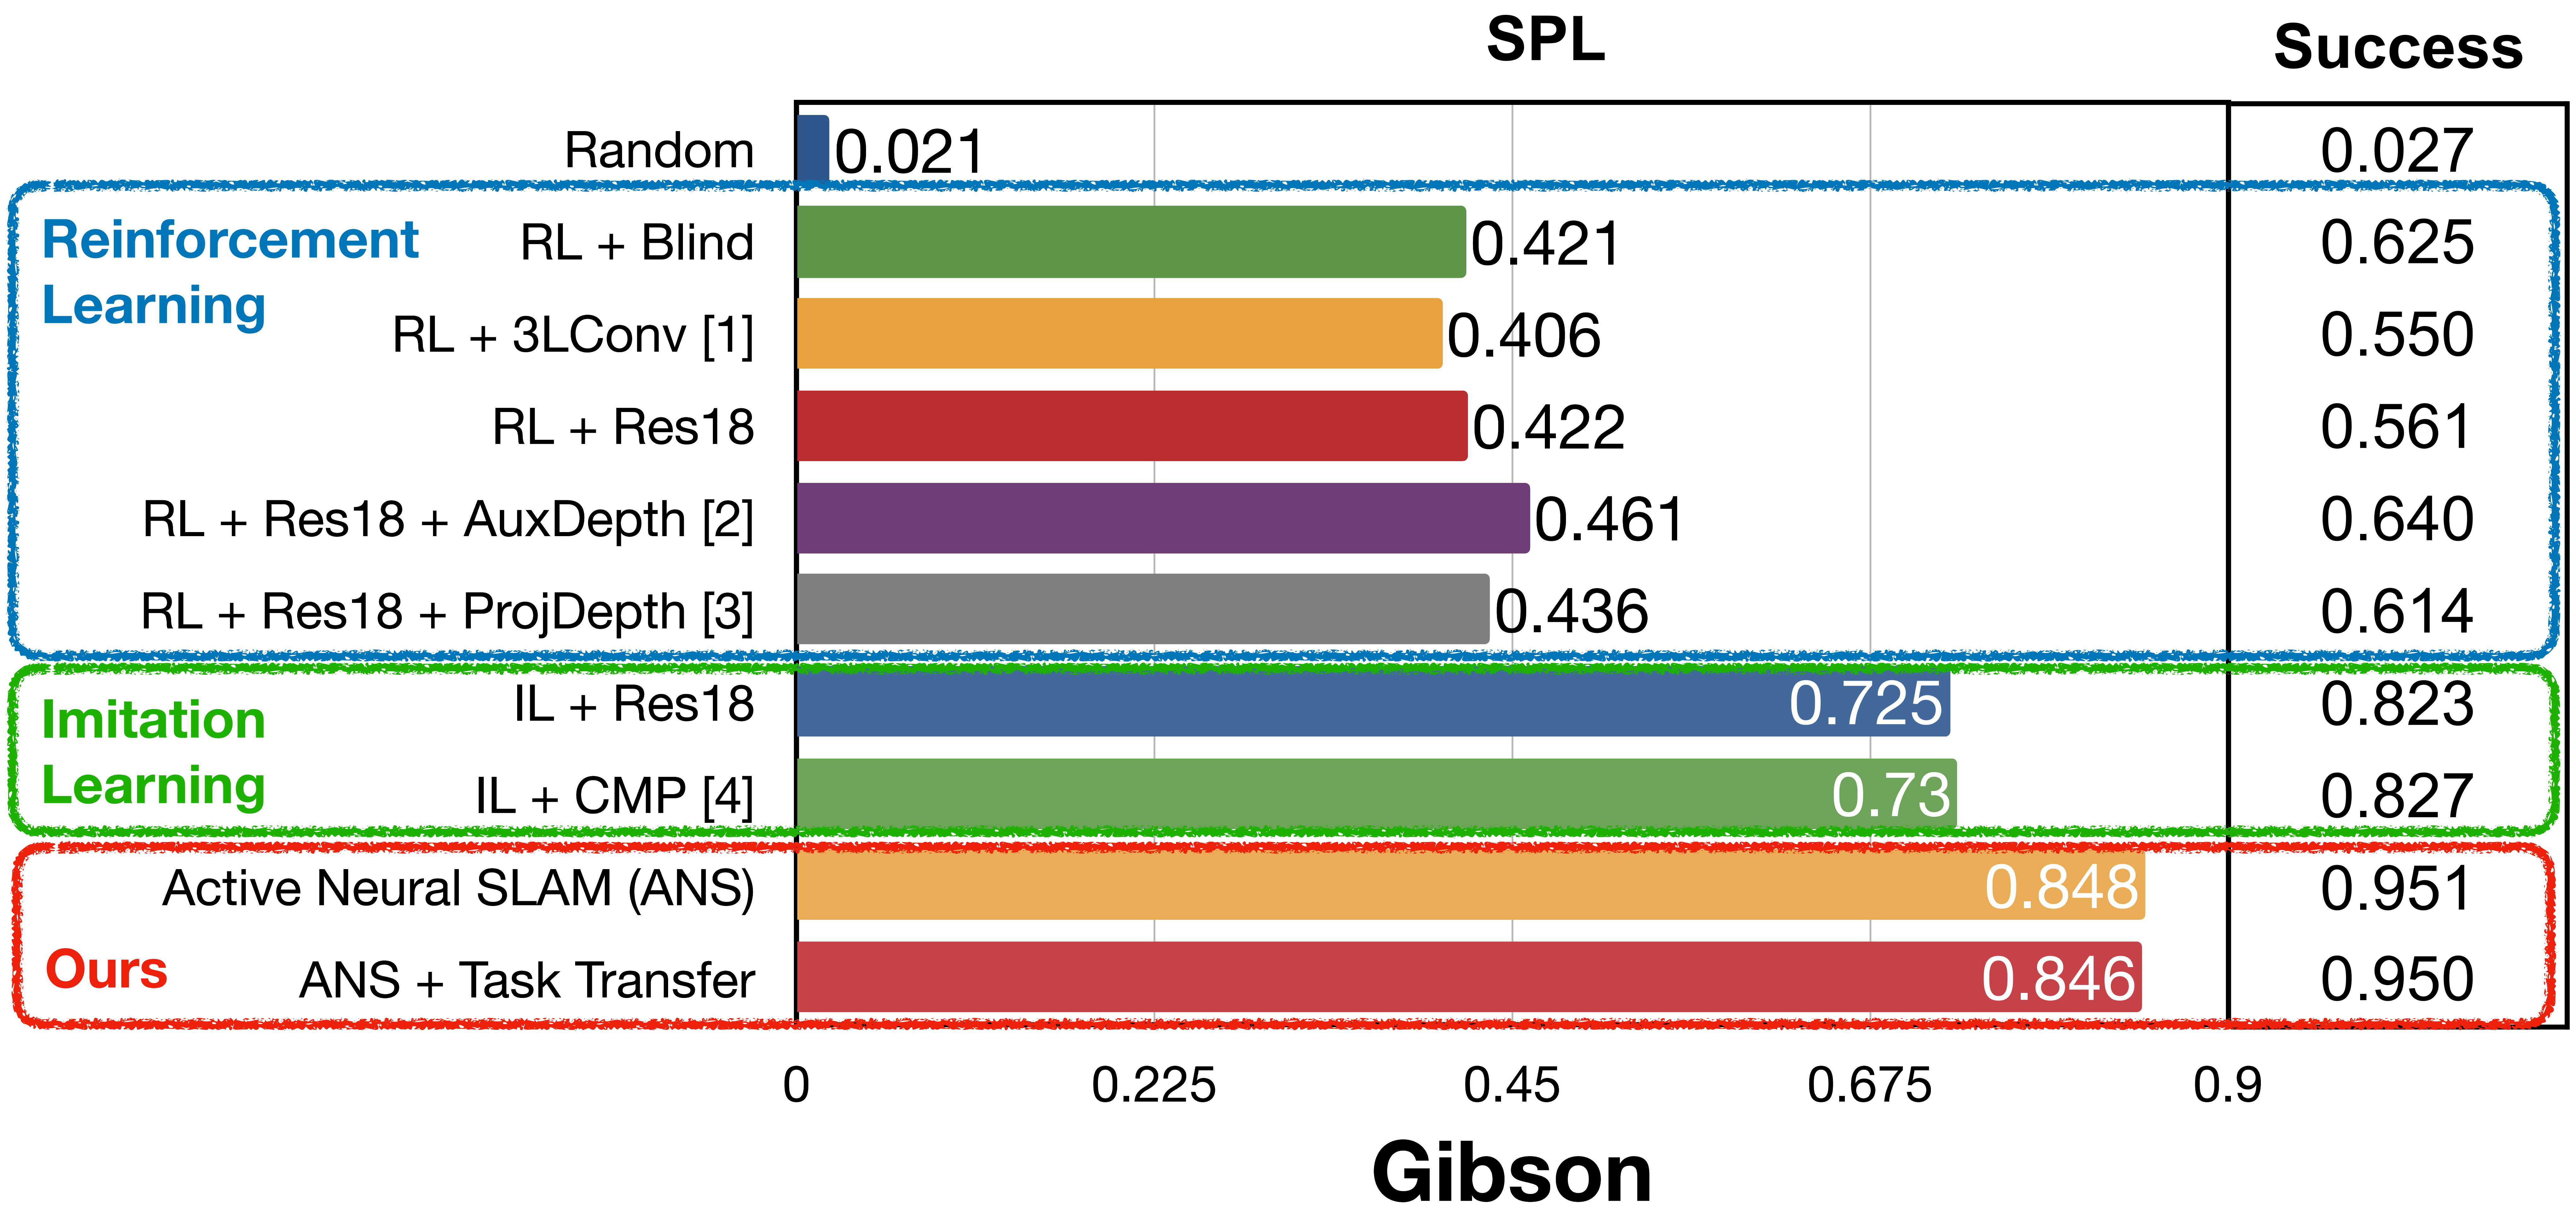

**Adapted from [1] Lample & Chaplot. AAAI-17, [2] Mirowski et al. ICLR-17, [3] Chen el al. ICLR-19, [4] Gupta et al. CVPR-17*

PointGoal Results

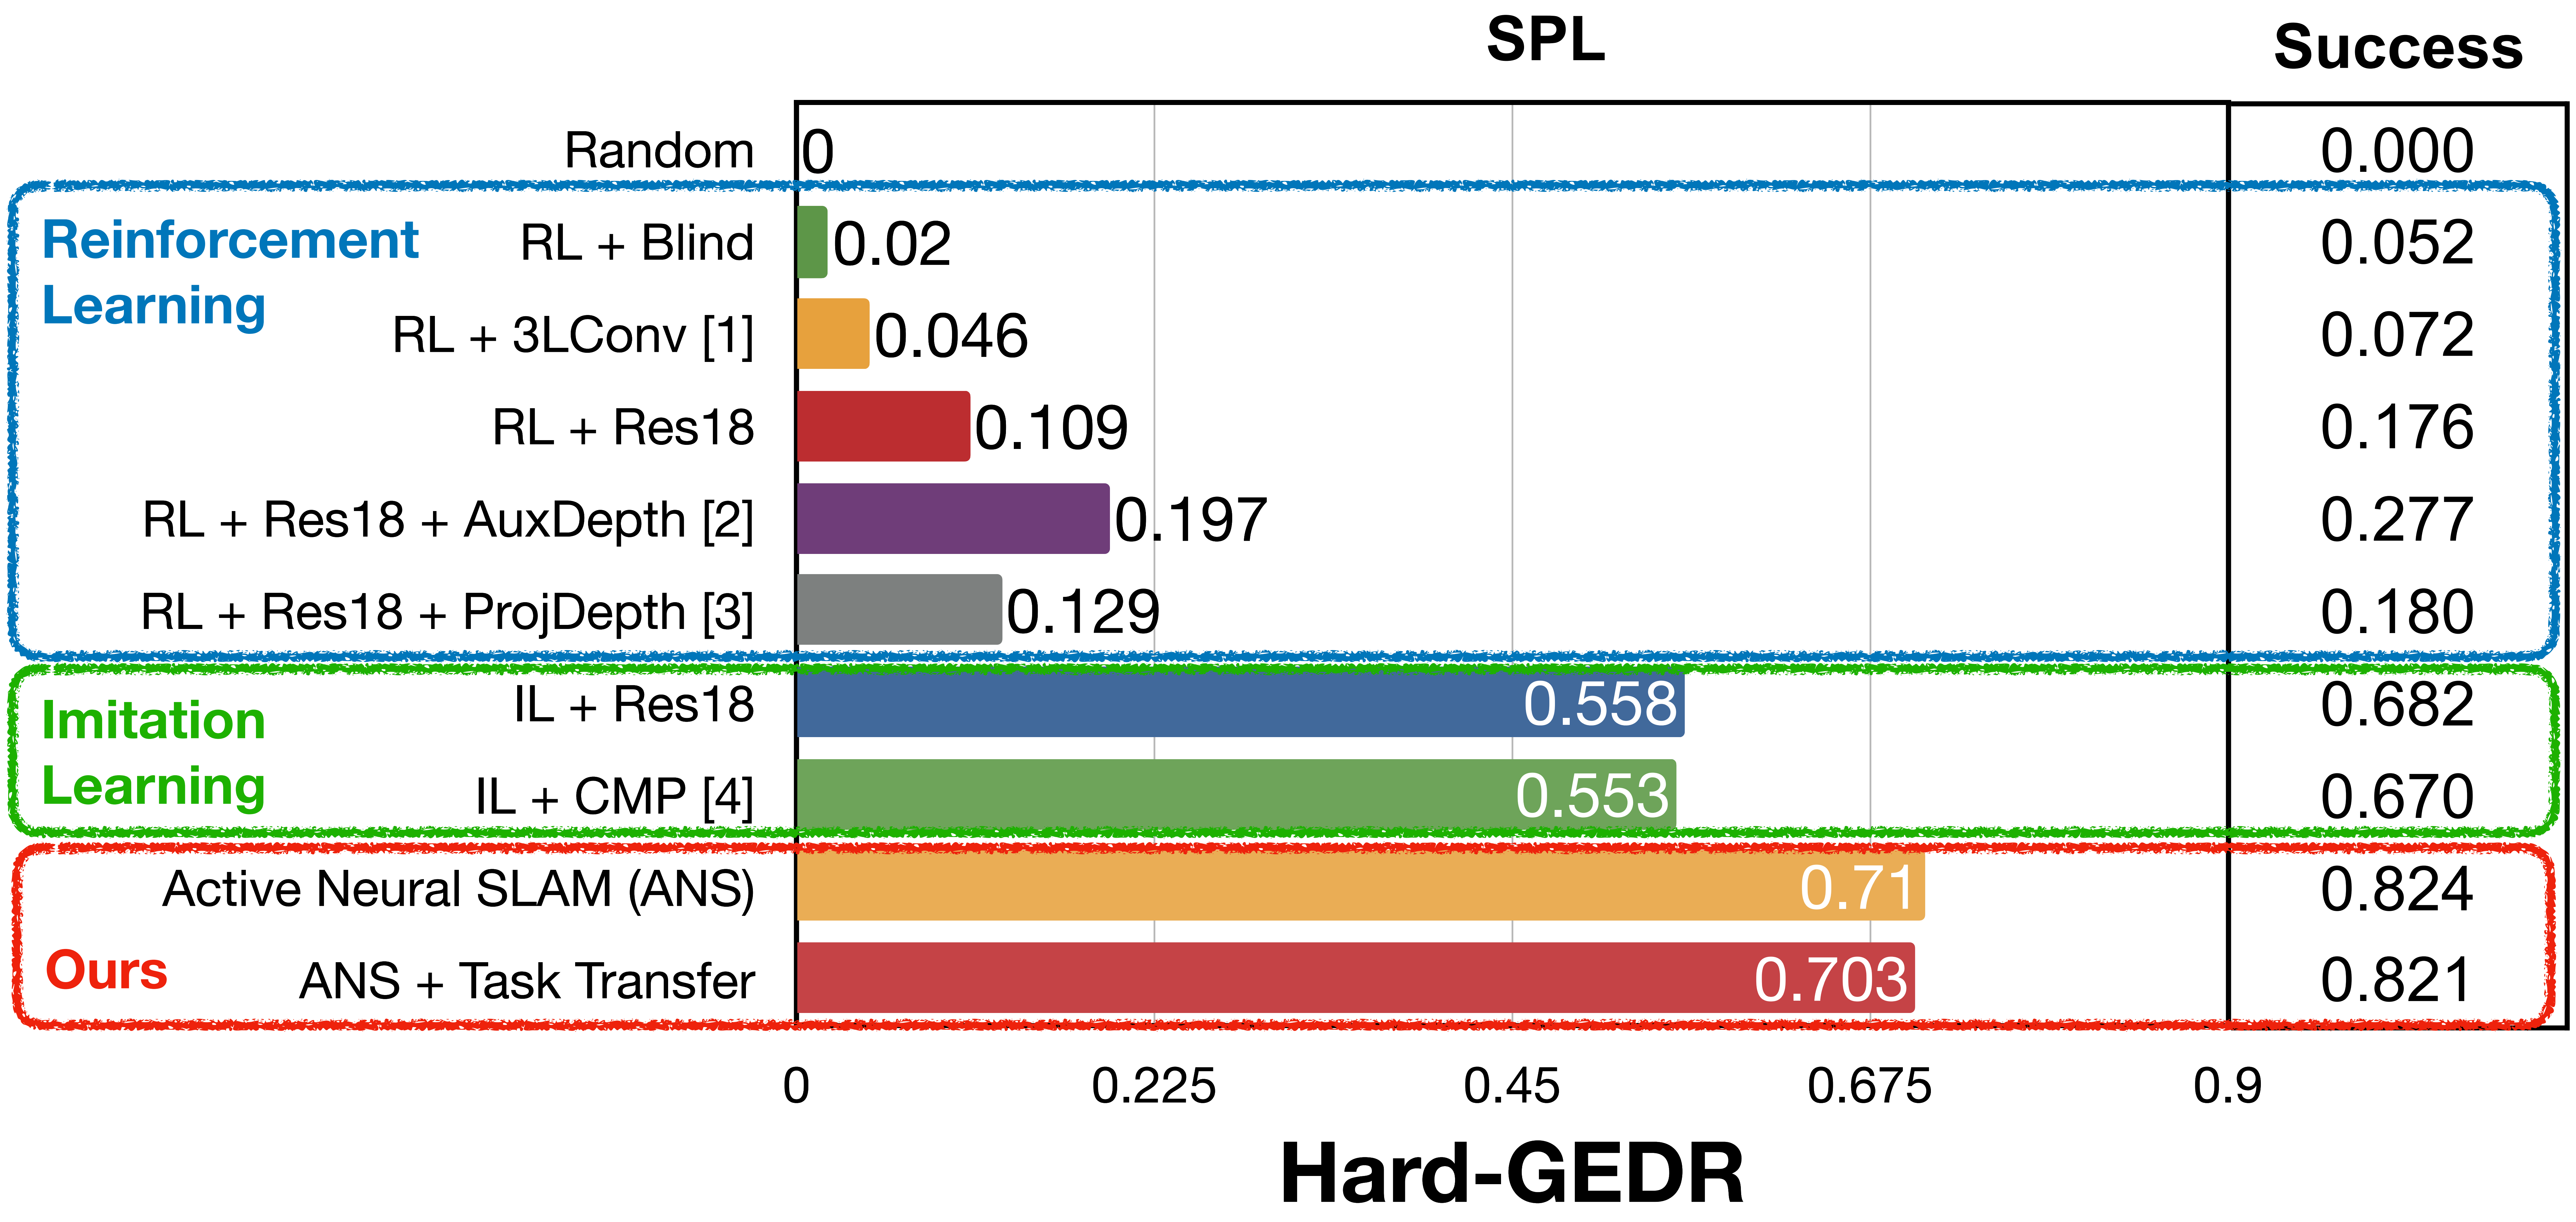

**Adapted from [1] Lample & Chaplot. AAAI-17, [2] Mirowski et al. ICLR-17, [3] Chen el al. ICLR-19, [4] Gupta et al. CVPR-17*

PointGoal Results

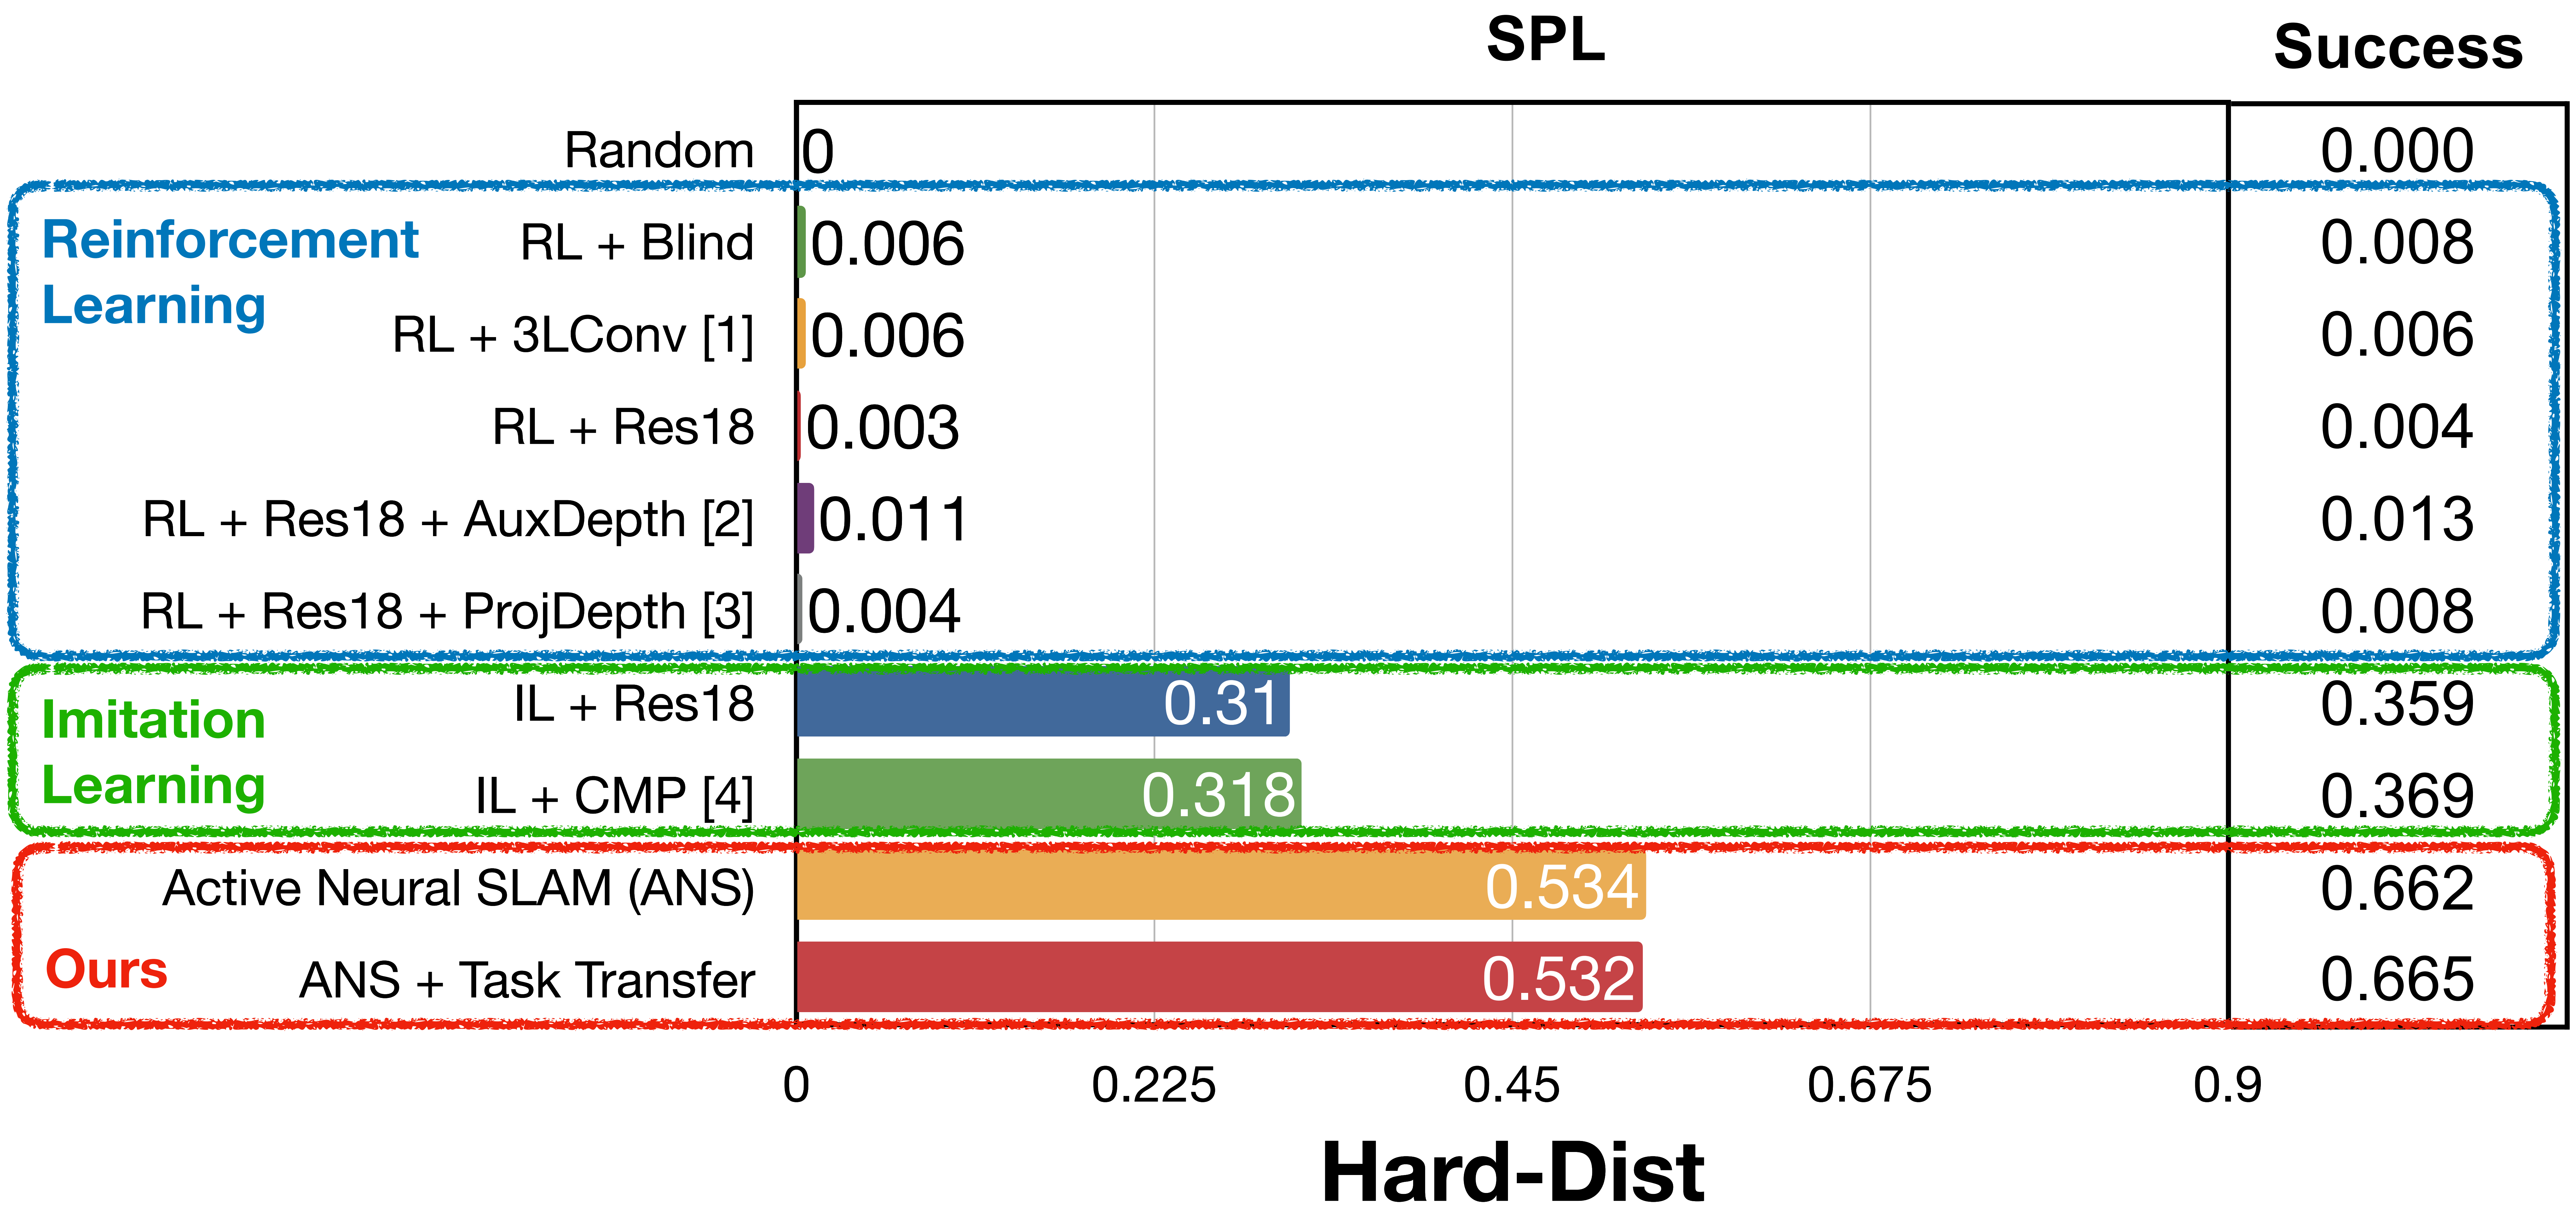

**Adapted from [1] Lample & Chaplot. AAAI-17, [2] Mirowski et al. ICLR-17, [3] Chen el al. ICLR-19, [4] Gupta et al. CVPR-17*

PointGoal Results

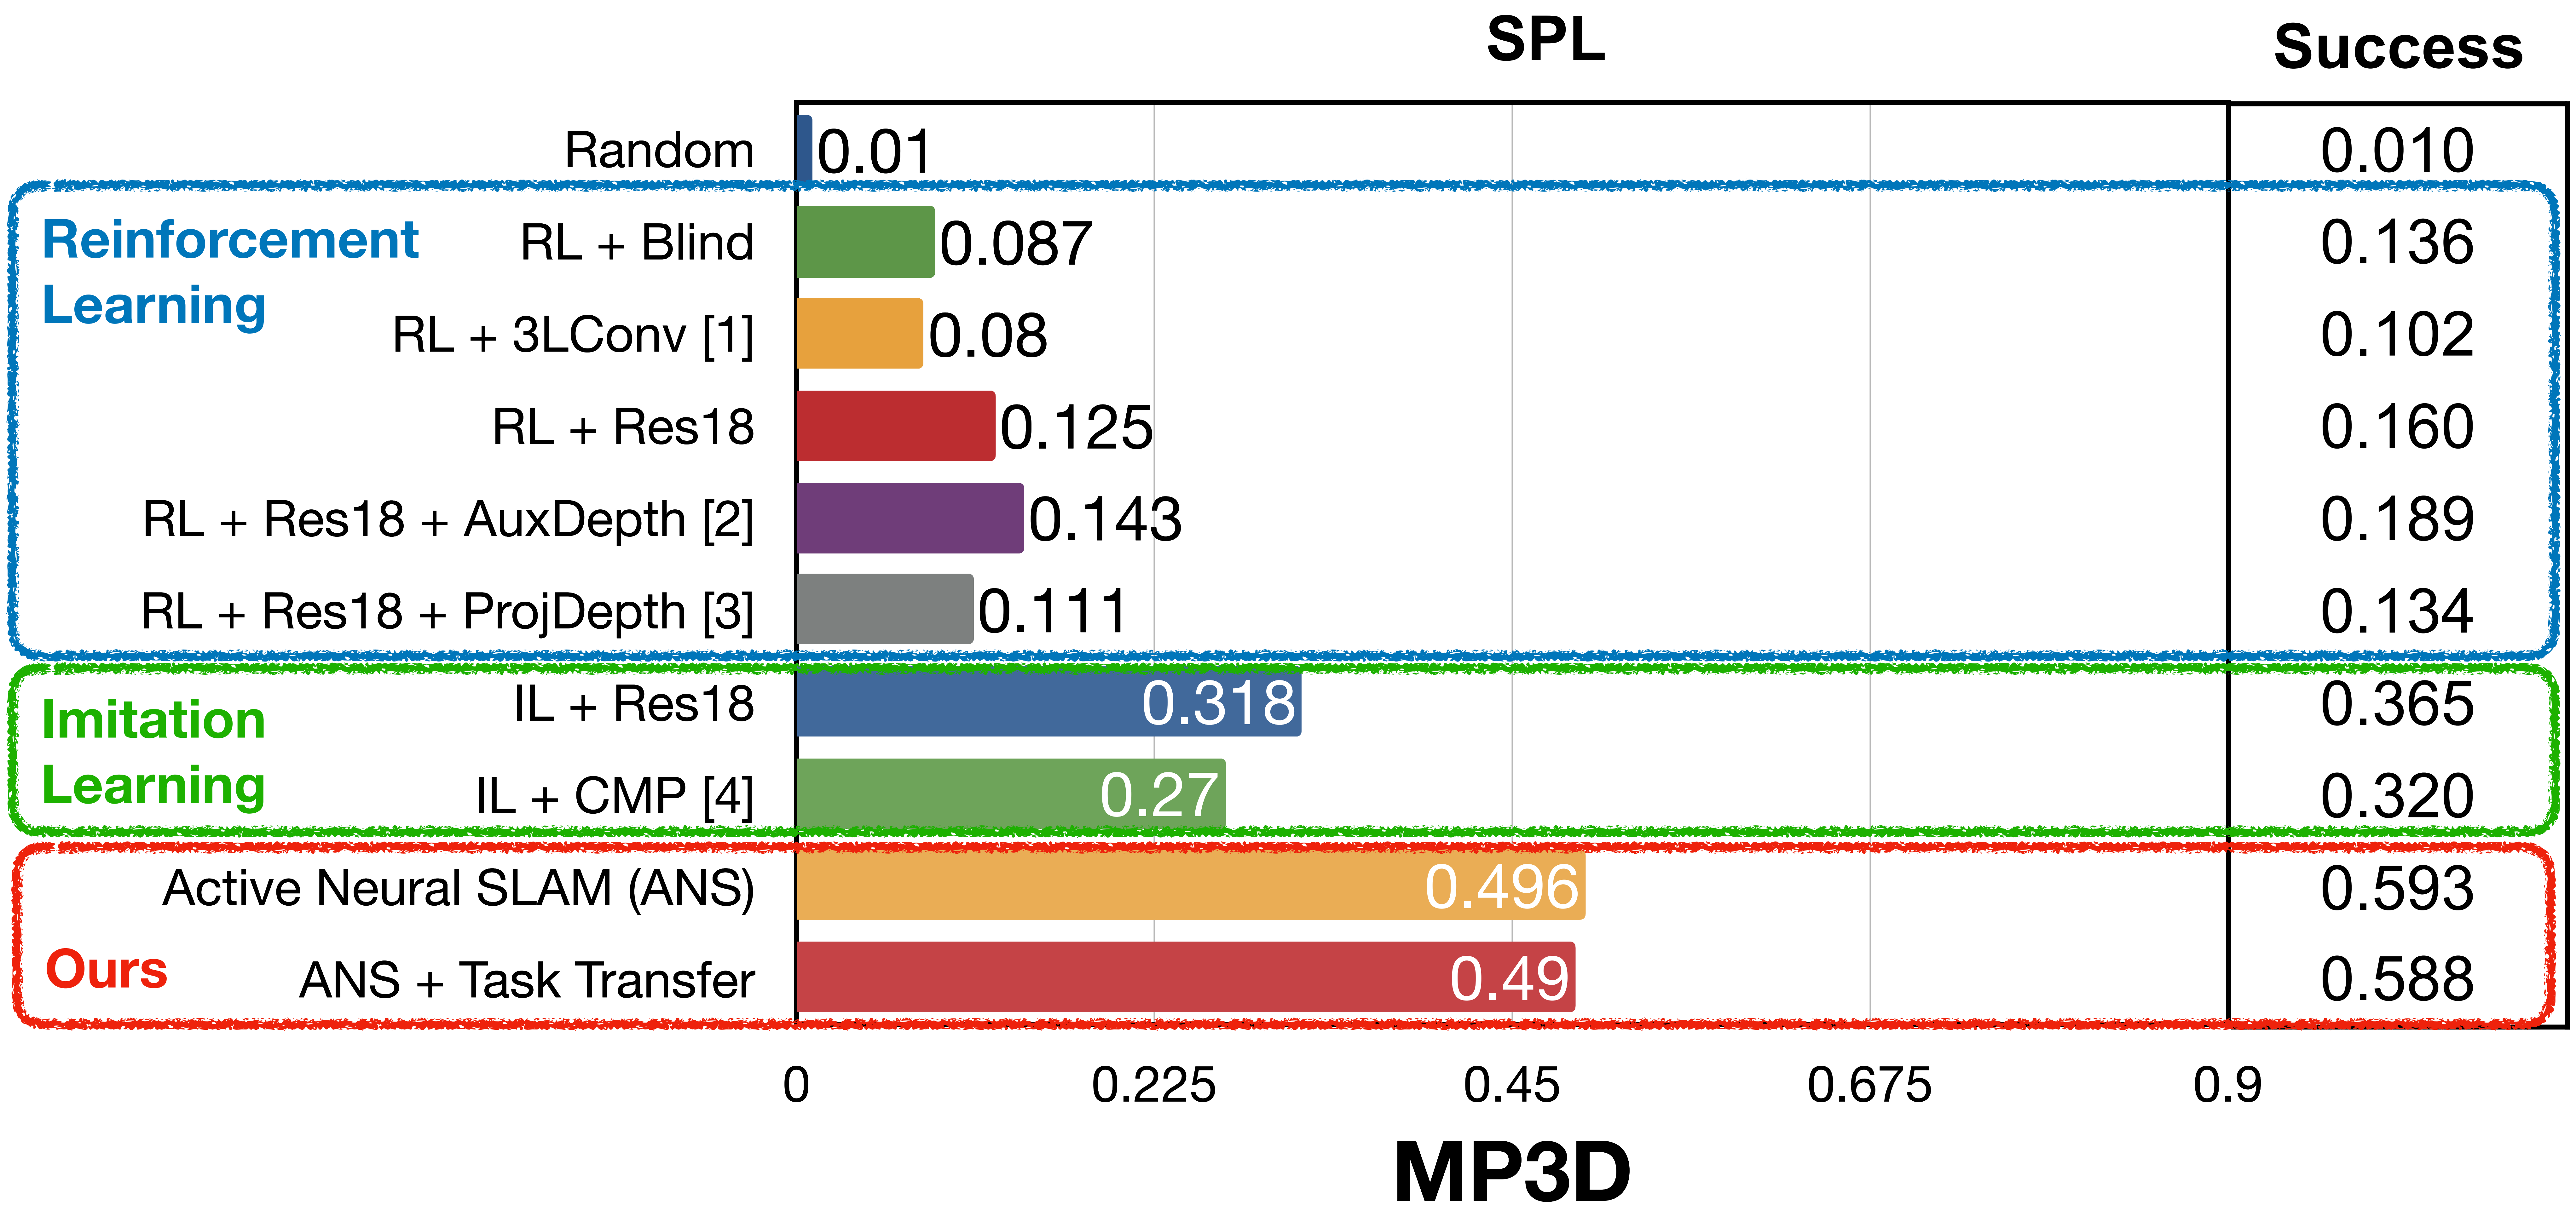

**Adapted from [1] Lample & Chaplot. AAAI-17, [2] Mirowski et al. ICLR-17, [3] Chen el al. ICLR-19, [4] Gupta et al. CVPR-17*

Results

Domain Generalization **Goal Generalization**

Gibson (Matterport 3D) (Harder goals)

Exploration

Task Generalization Pointgoal

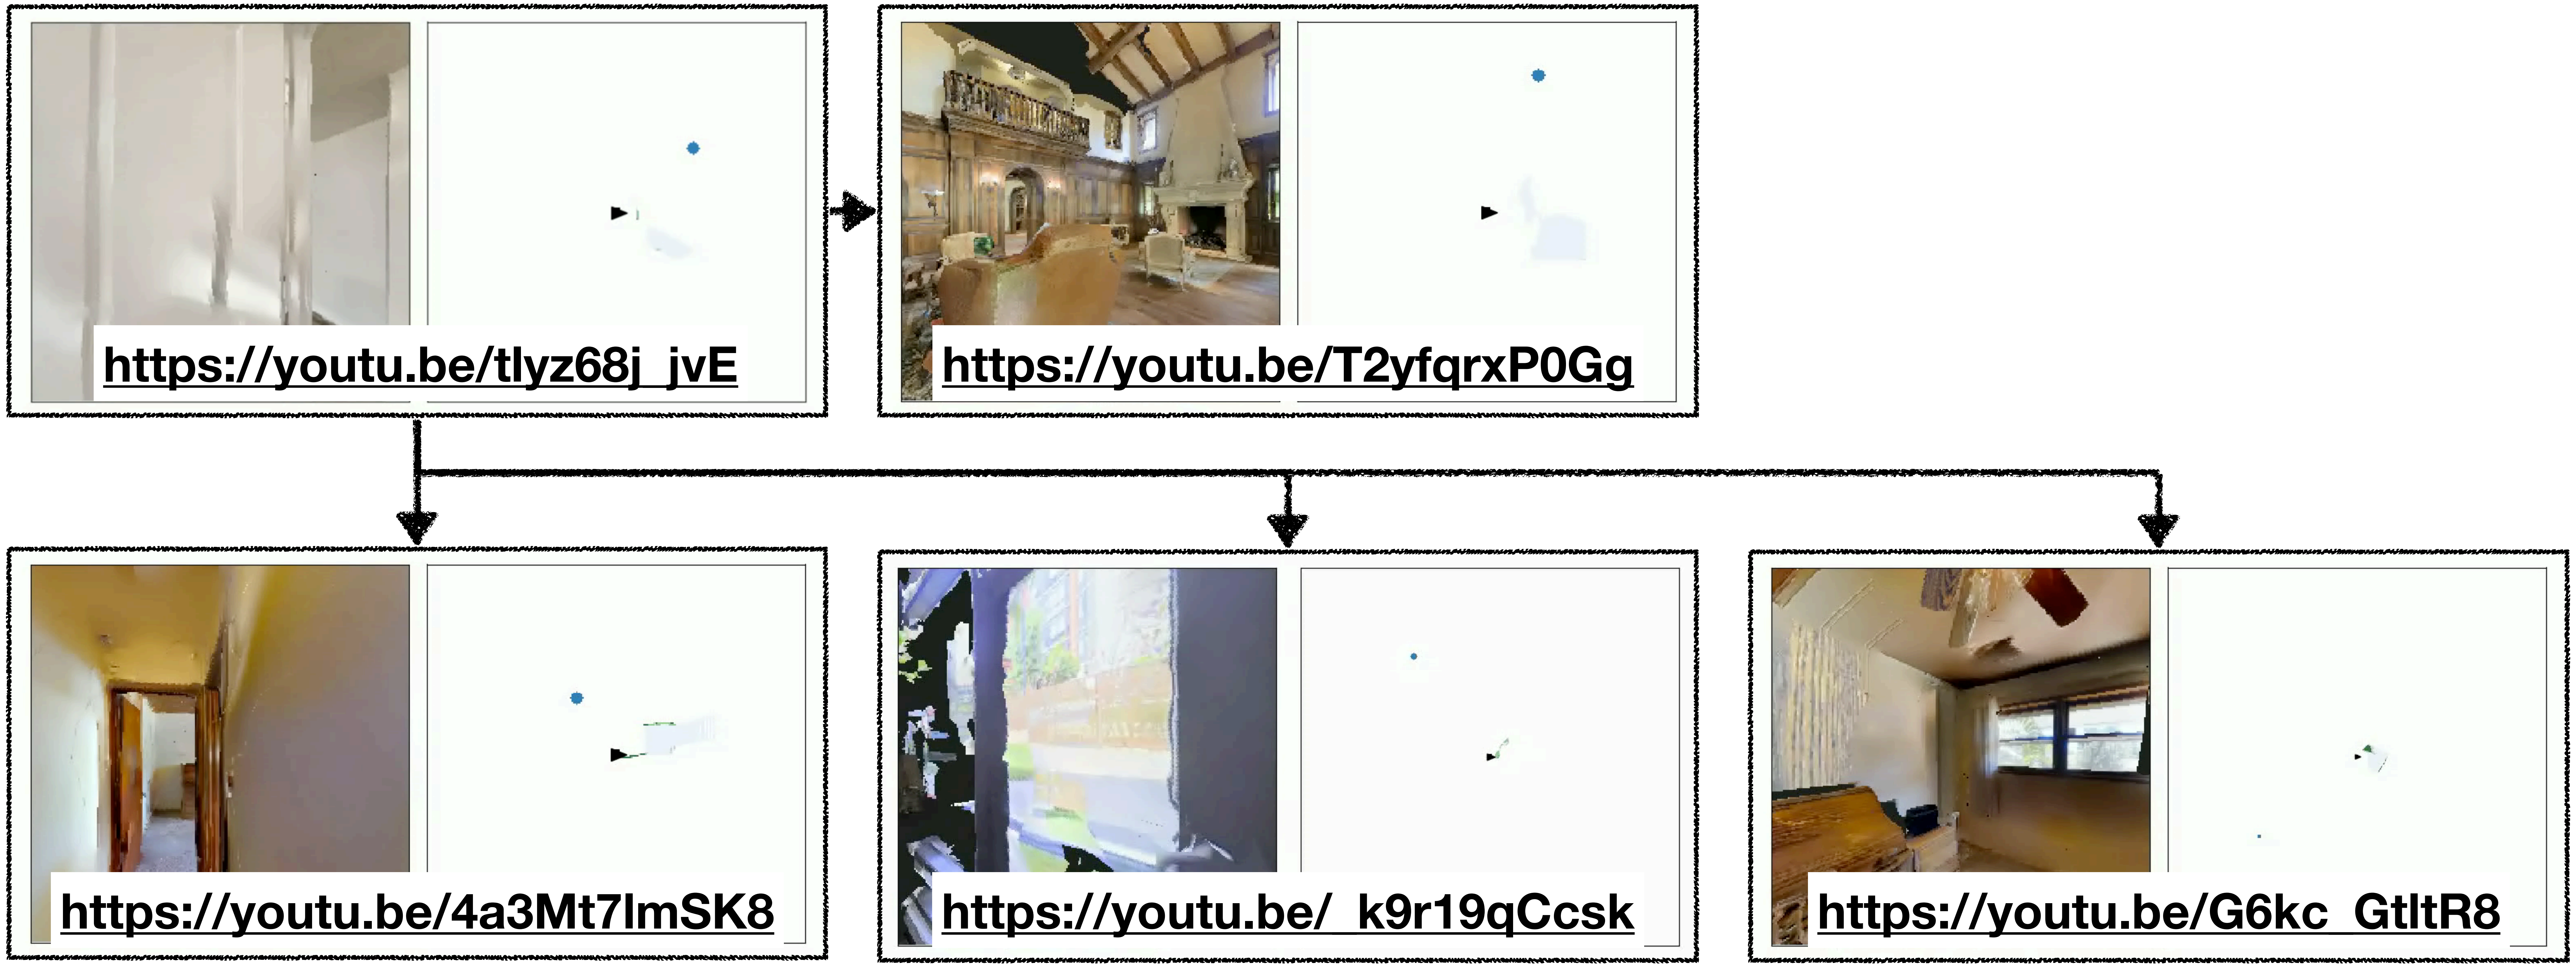

Results

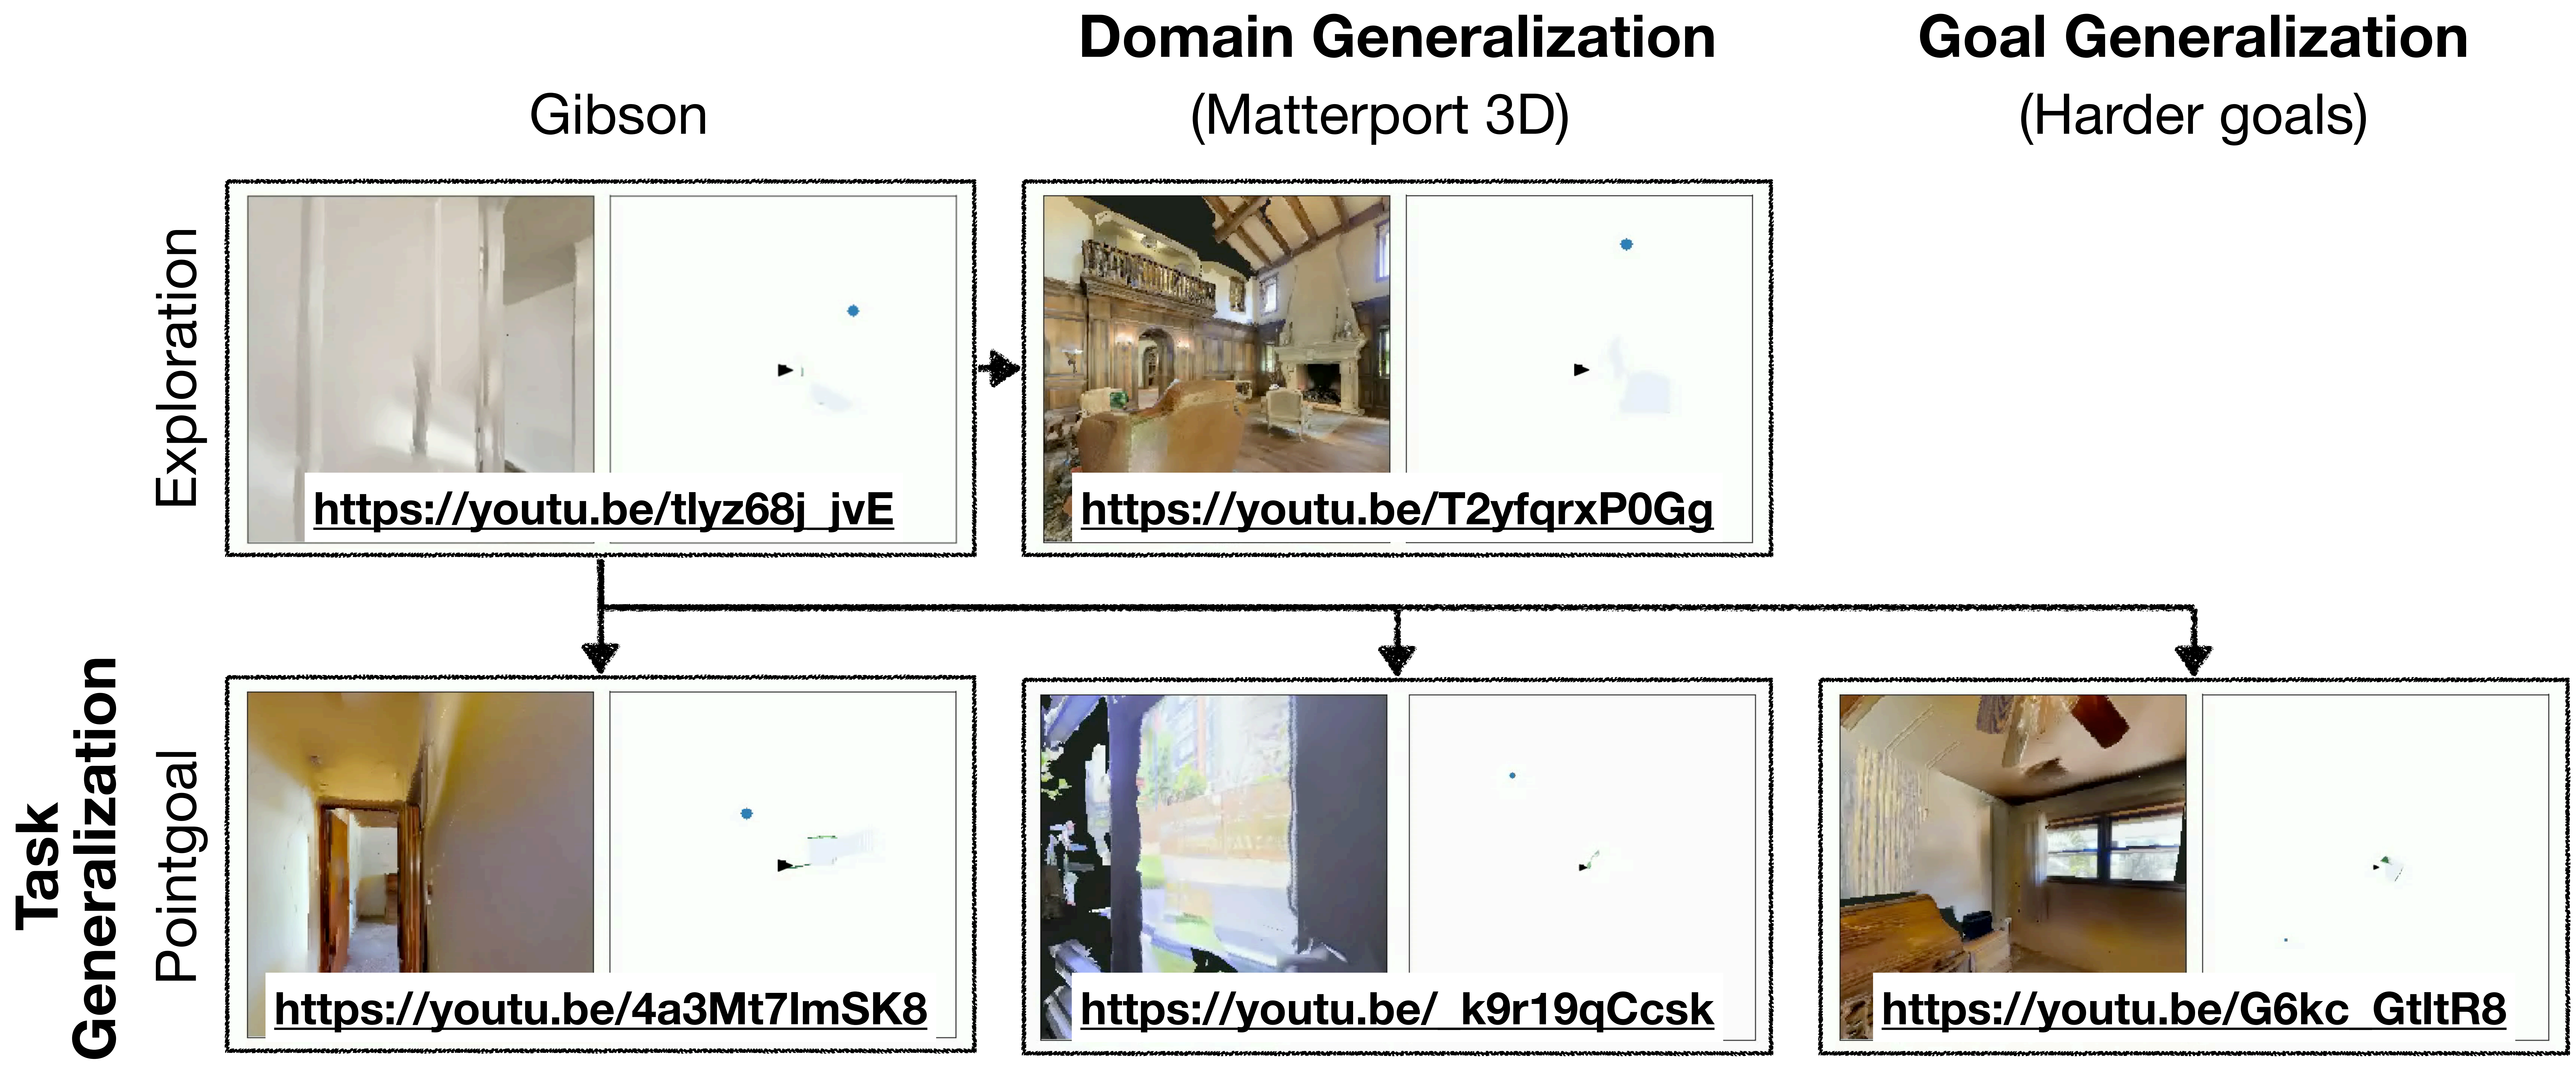
Domain Generalization
Goal Generalization
Gibson
(Matterport 3D)
(Harder goals)
Exploration
https://youtu.be/tlyz68j_jvE
https://youtu.be/T2yfqrxP0Gg
Task Generalization
Pointgoal
https://youtu.be/4a3Mt7lmSK8
https://youtu.be/_k9r19qCcsk
https://youtu.be/G6kc_GtltR8

Winner of CVPR 2019 Habitat Challenge

RGB Leaderboard

Rank	Team	SPL
1	**Active Neural SLAM (Arnold)**	0.805
1	Mid-level-Features	0.800
3	CHROMA	0.712
4	ARF-RL	0.699
5	MTank	0.260

RGBD Leaderboard

Rank	Team	SPL
1	**Active Neural SLAM (Arnold)**	0.948
2	Pansy	0.927
3	Titardrew	0.868
4	Hiccup	0.846
5	CHROMA	0.843

Winner of CVPR 2019 Habitat Challenge

RGB Leaderboard

Rank	Team	SPL
1	**Active Neural SLAM (Arnold)**	0.805
1	Mid-level-Features	0.800
3	CHROMA	0.712
4	ARF-RL	0.699
5	MTank	0.260

RGBD Leaderboard

Rank	Team	SPL
1	**Active Neural SLAM (Arnold)**	0.948
2	Pansy	0.927
3	Titardrew	0.868
4	Hiccup	0.846
5	CHROMA	0.843

Sim-to-Real Transfer

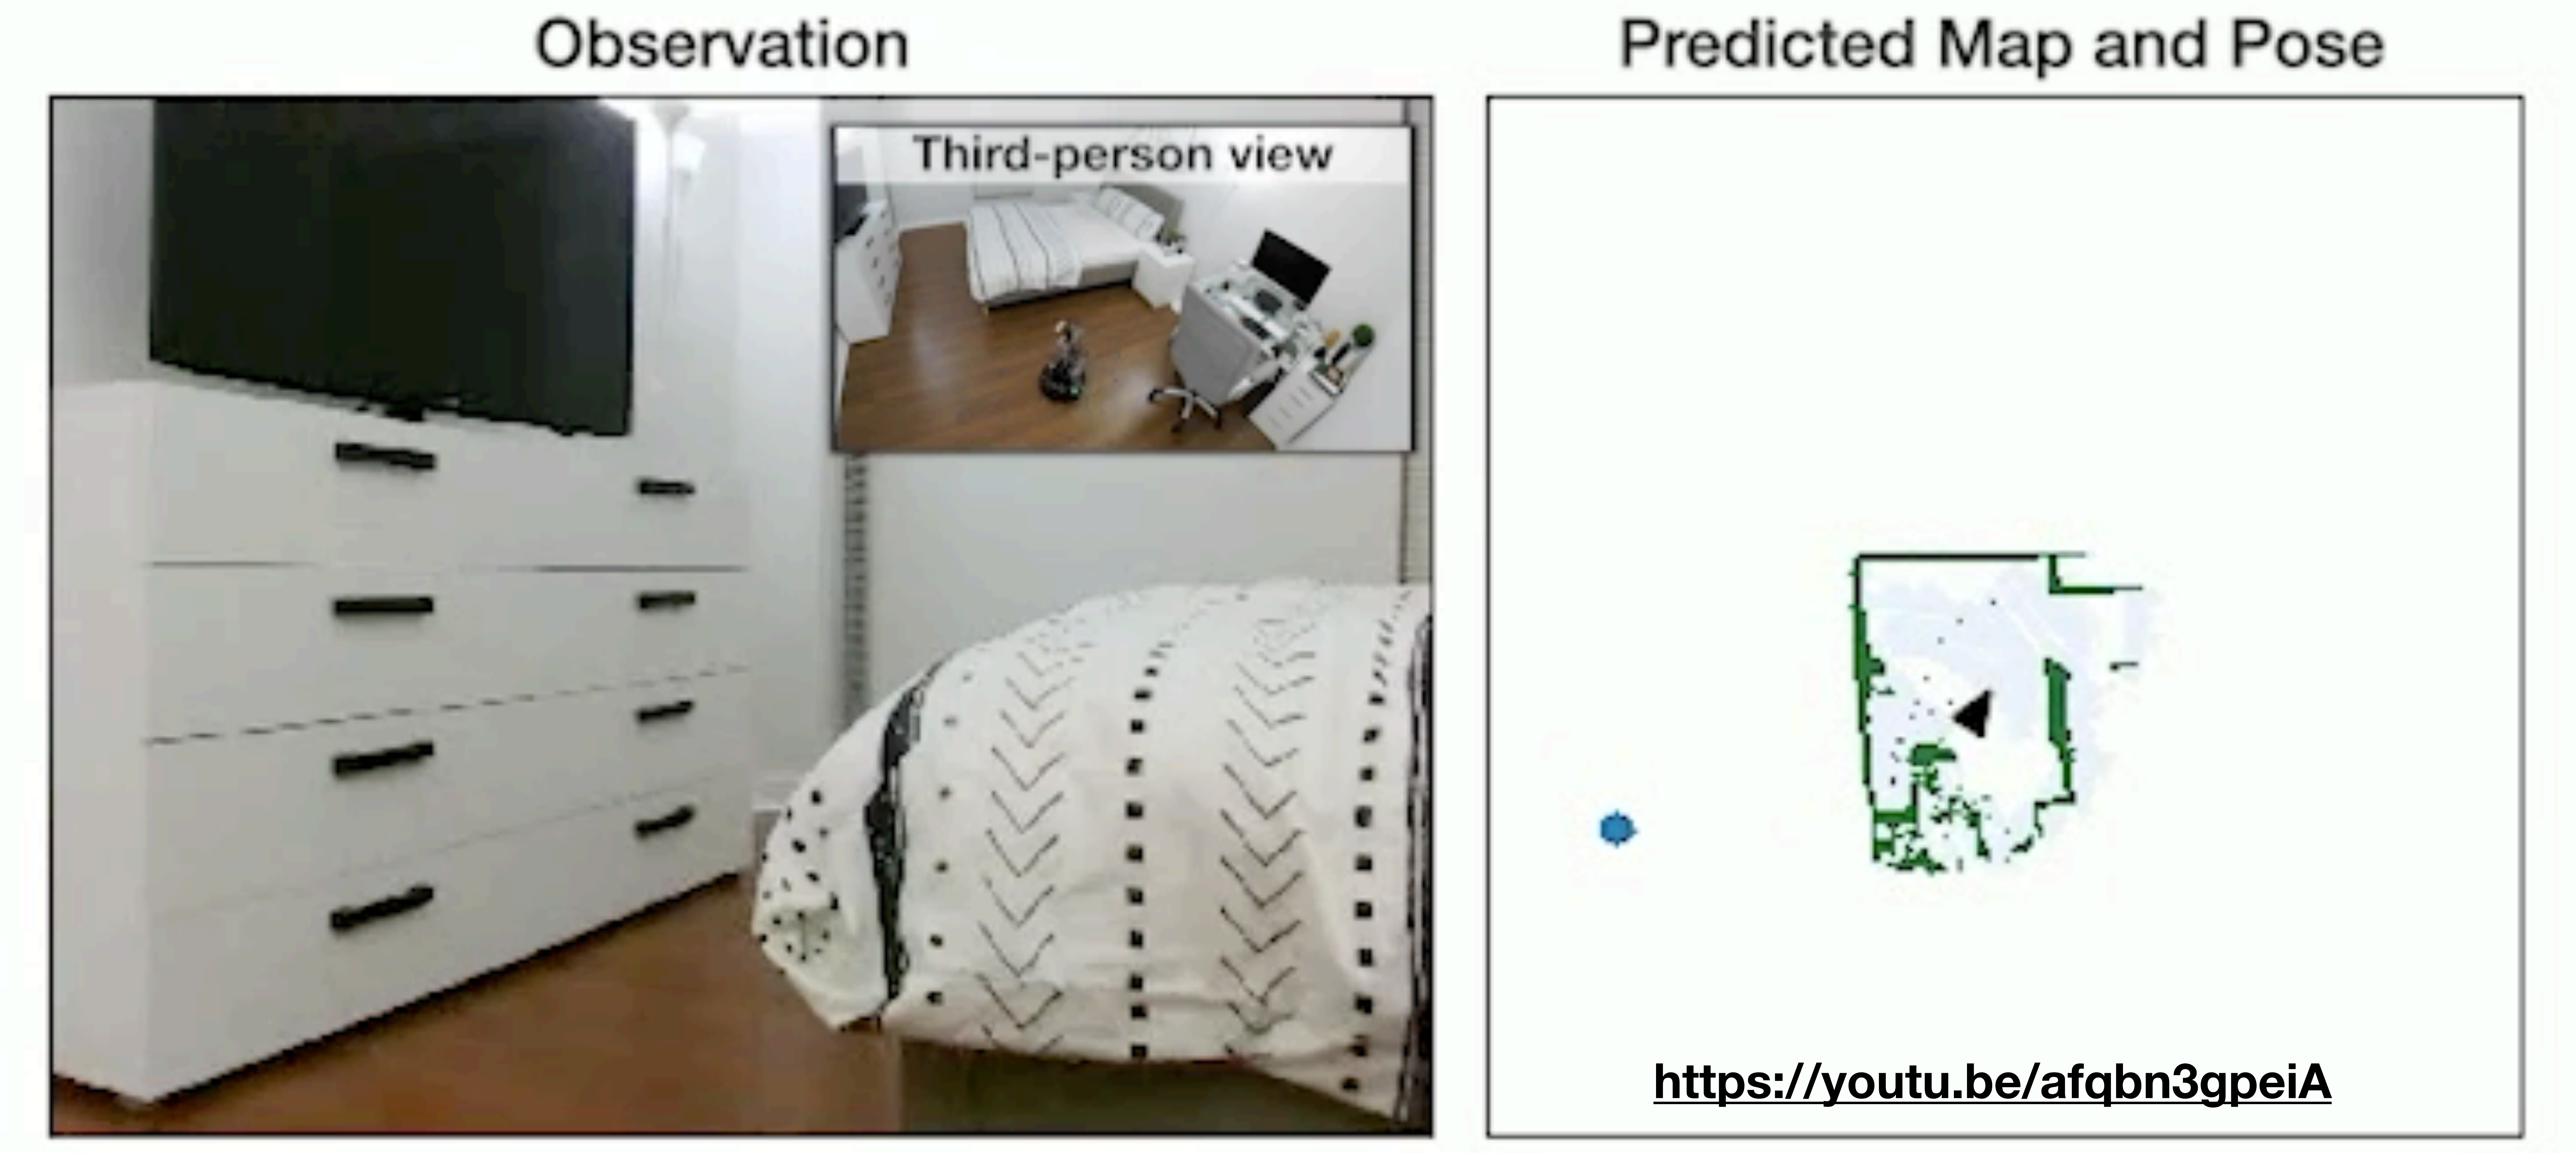

https://youtu.be/afqbn3gpeiA

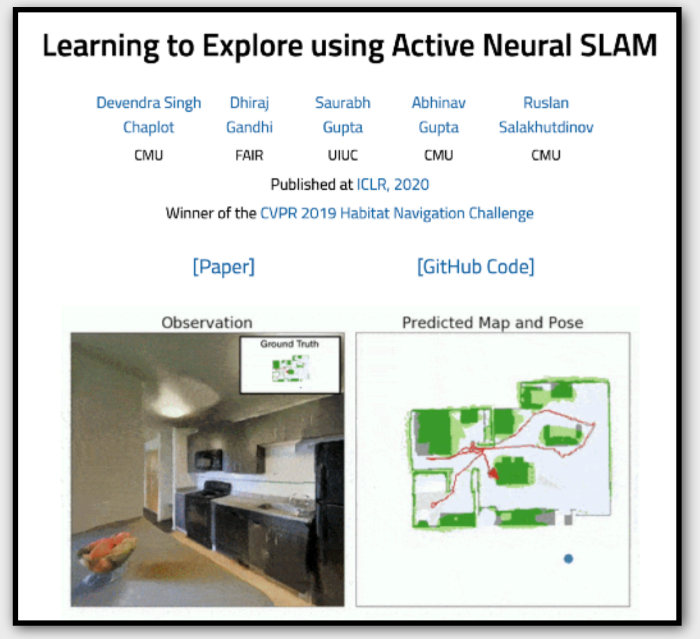

Learning to Explore using Active Neural SLAM

Webpage: https://devendrachaplot.github.io/projects/Neural-SLAM
Code: https://github.com/devendrachaplot/Neural-SLAM

Thank you

Devendra Singh Chaplot

Webpage: http://devendrachaplot.github.io/
Email: chaplot@cs.cmu.edu
Twitter: @dchaplot