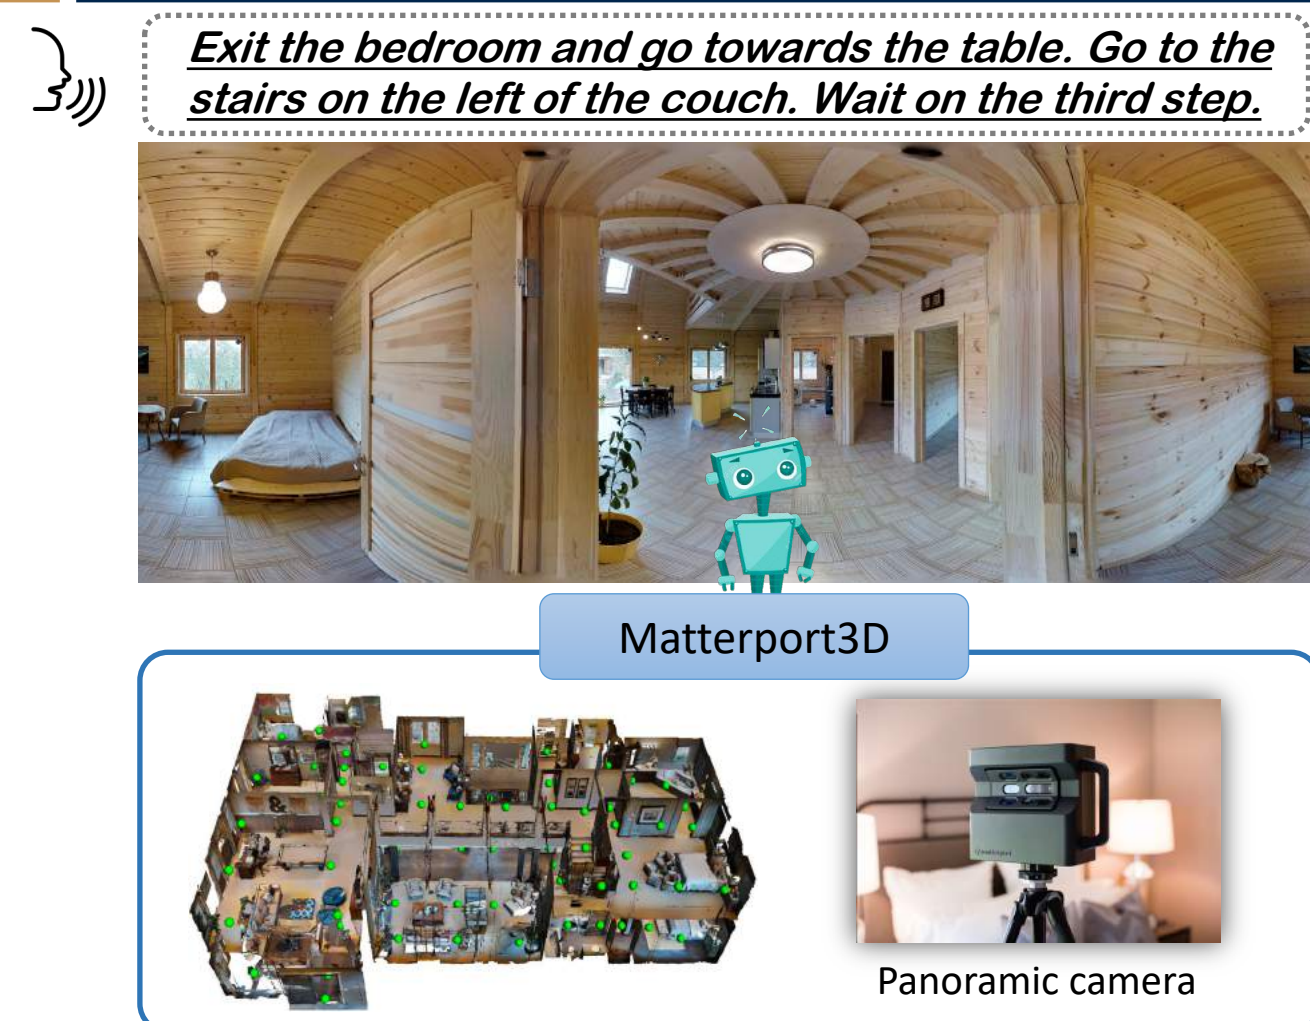
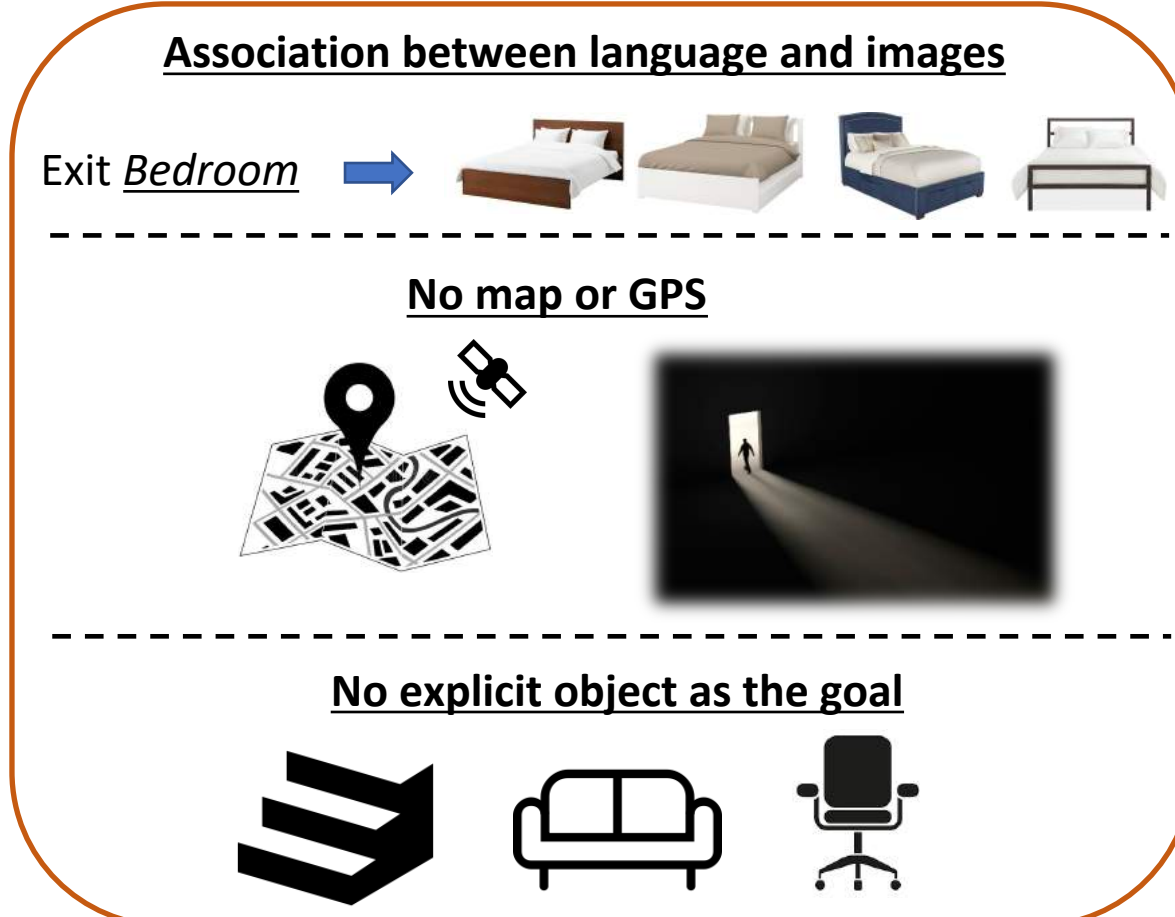


1 PROBLEM: VISION-AND-LANGUAGE NAVIGATION (VLN)

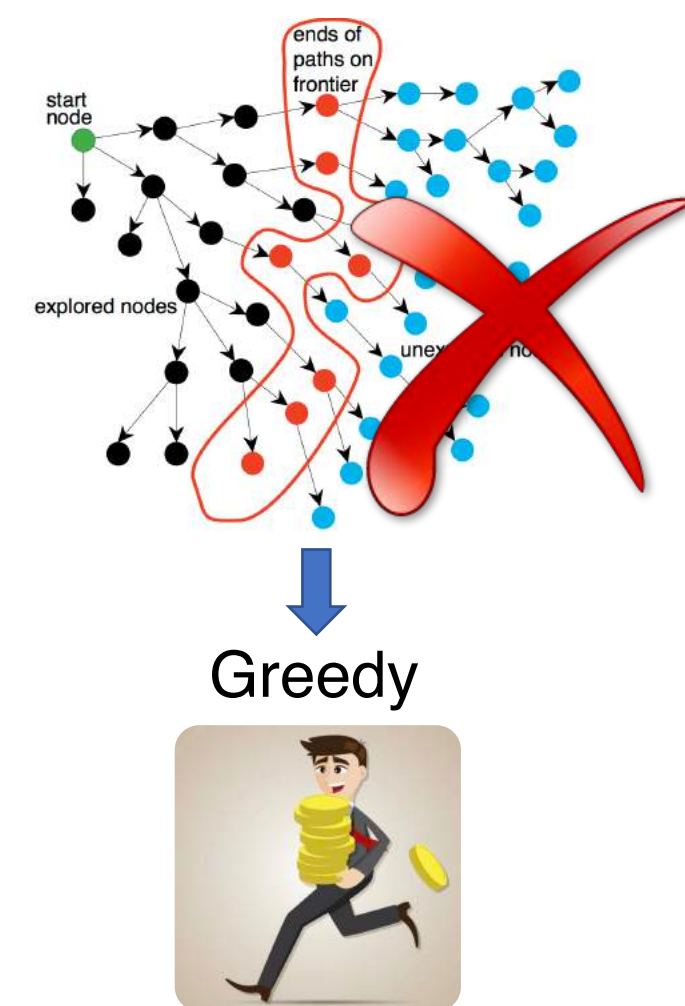


Challenges:



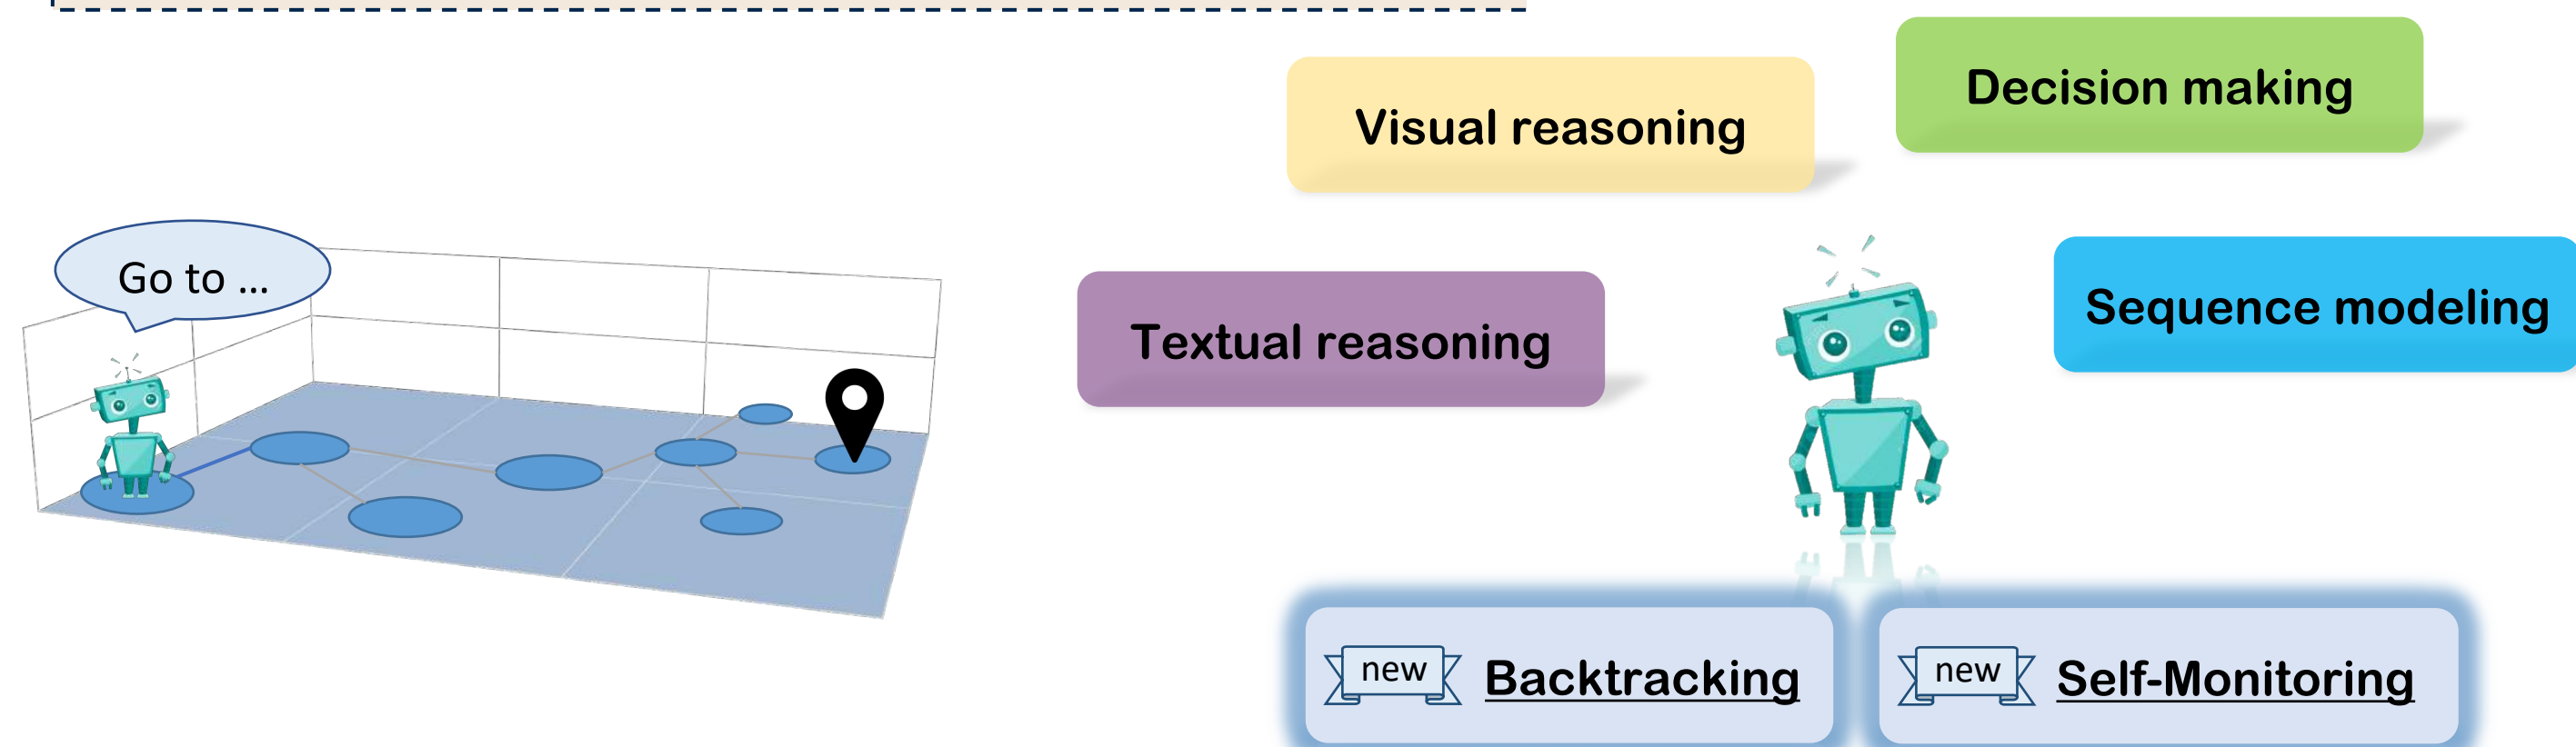
2 INTUITION: NAVIGATION AS GRAPH SEARCH

- Prior methods rely heavily on **beam search**, which requires a thorough exploration of trajectories. We instead focus on **greedy** navigation without frontiers.
- We view the problem as **graph search**, leveraging our prior work on a progress monitor – *a learned heuristic*.
- We propose to **learn a backtracking** mechanism in an end-to-end manner as well as **progress markers** to mark states.

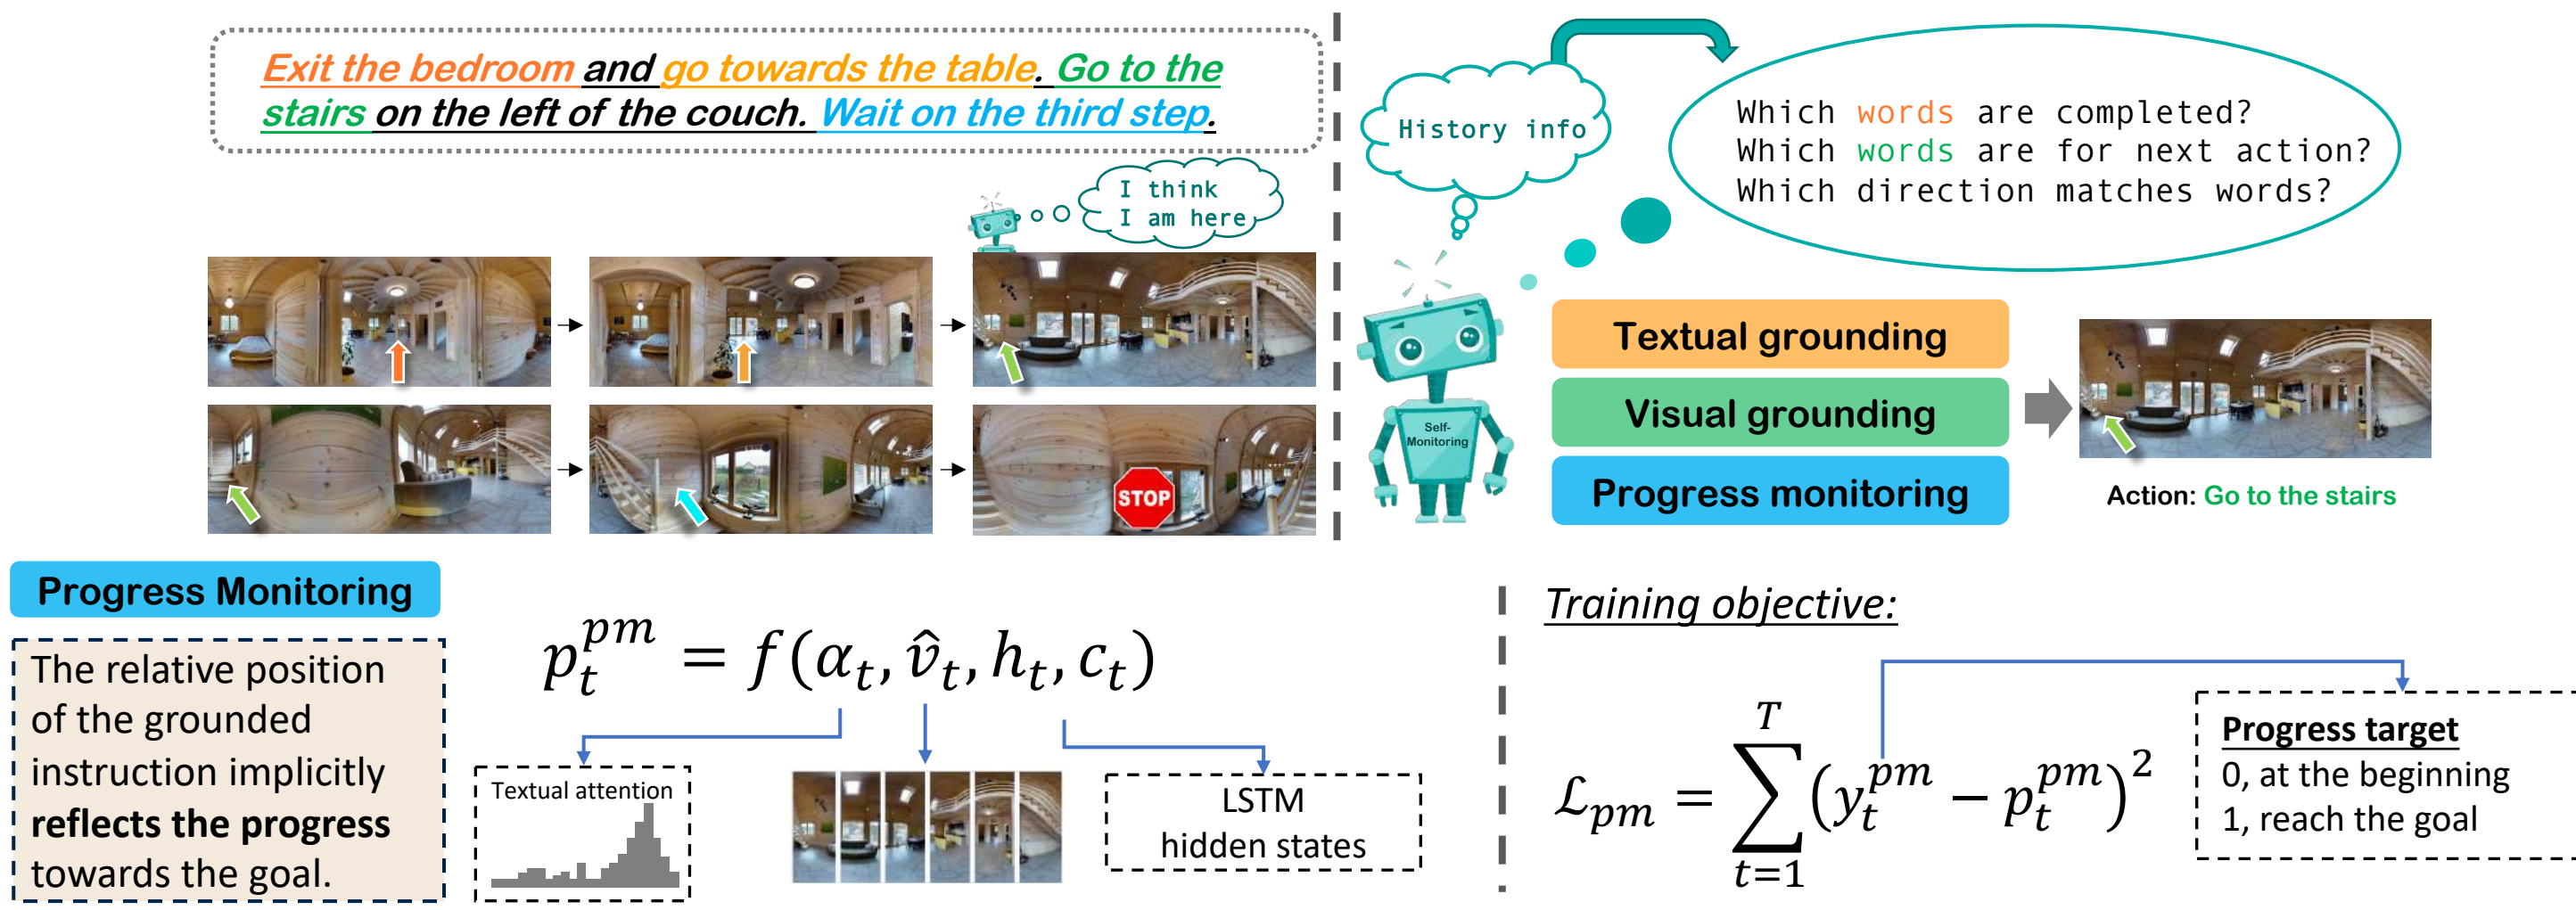


3 GOAL-DRIVEN REASONING WITH LEARNED BACKTRACKING

What do we need for goal-driven navigation?



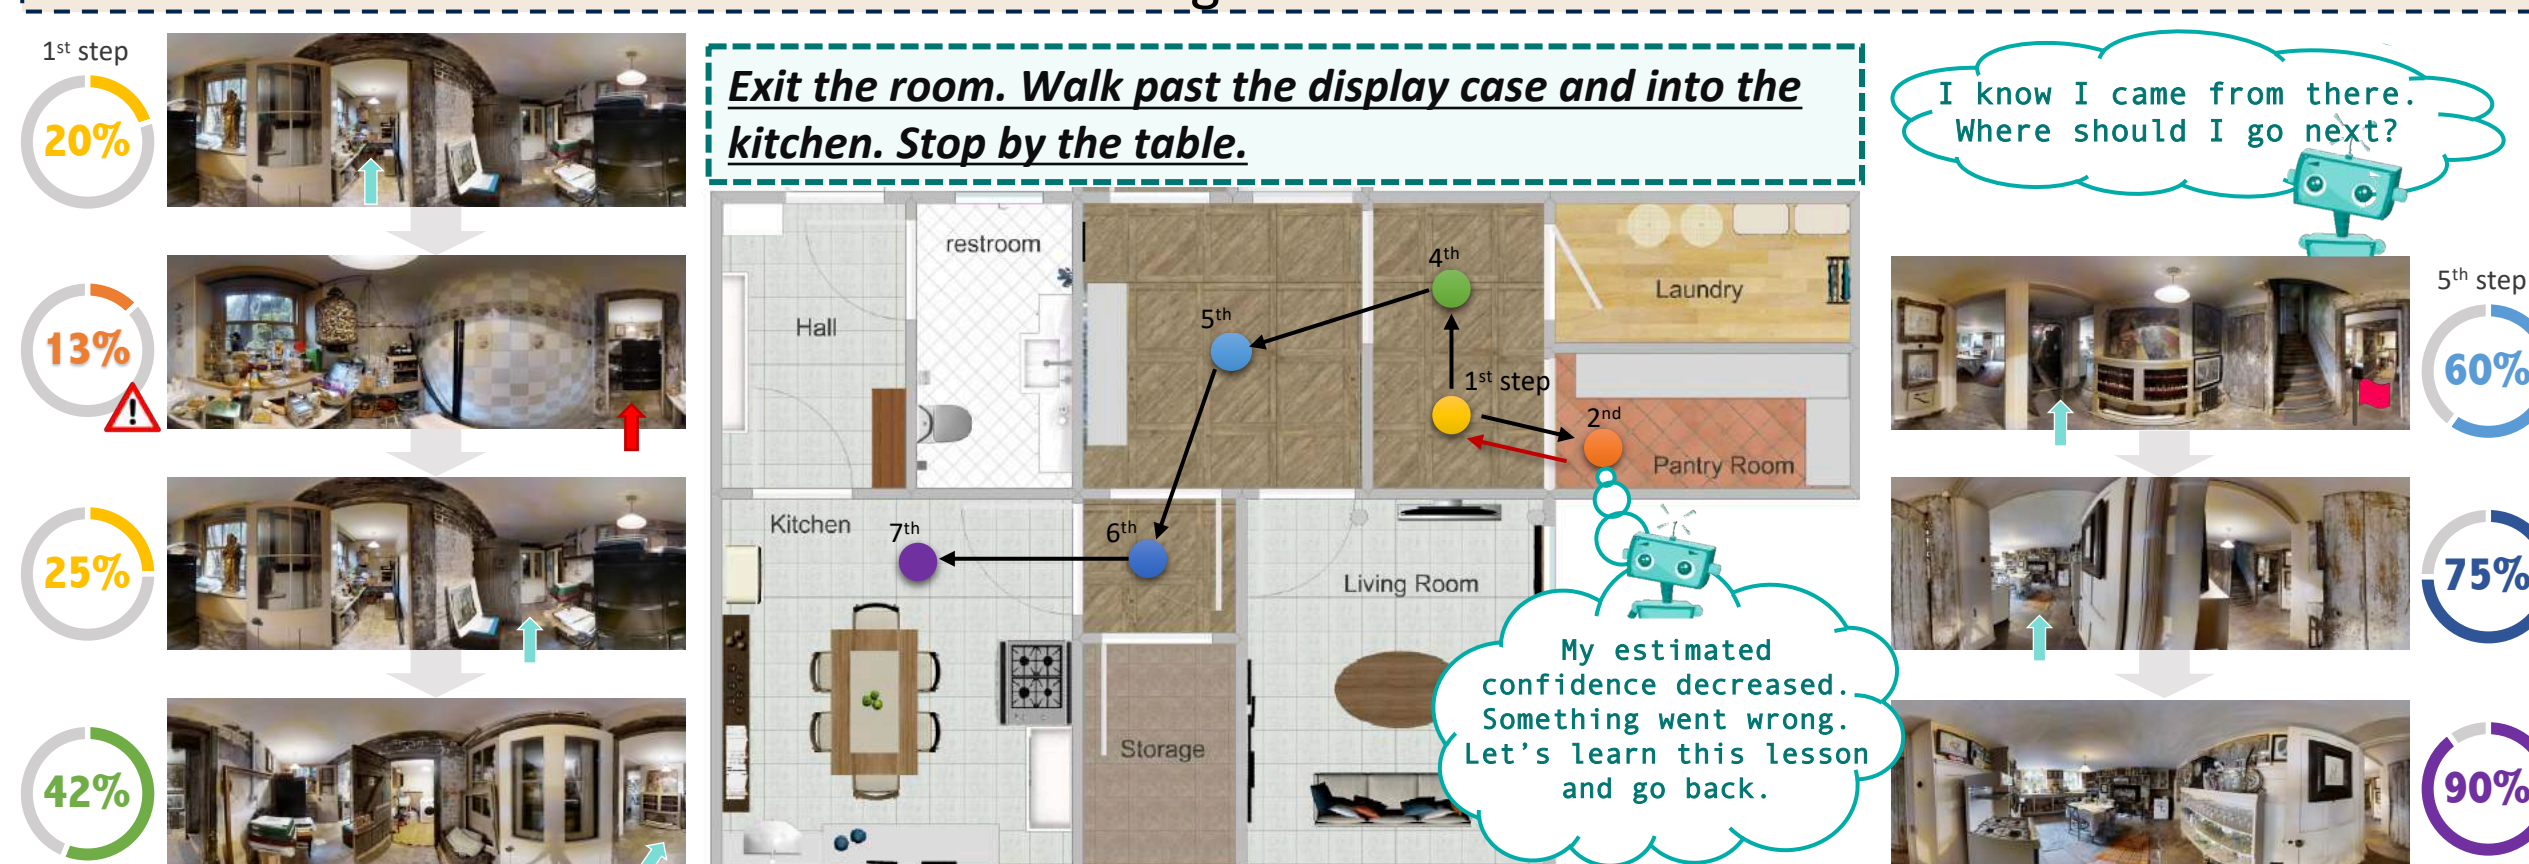
4 SELF-MONITORING AGENT



1 END-TO-END LEARNED BACKTRACKING AGENT

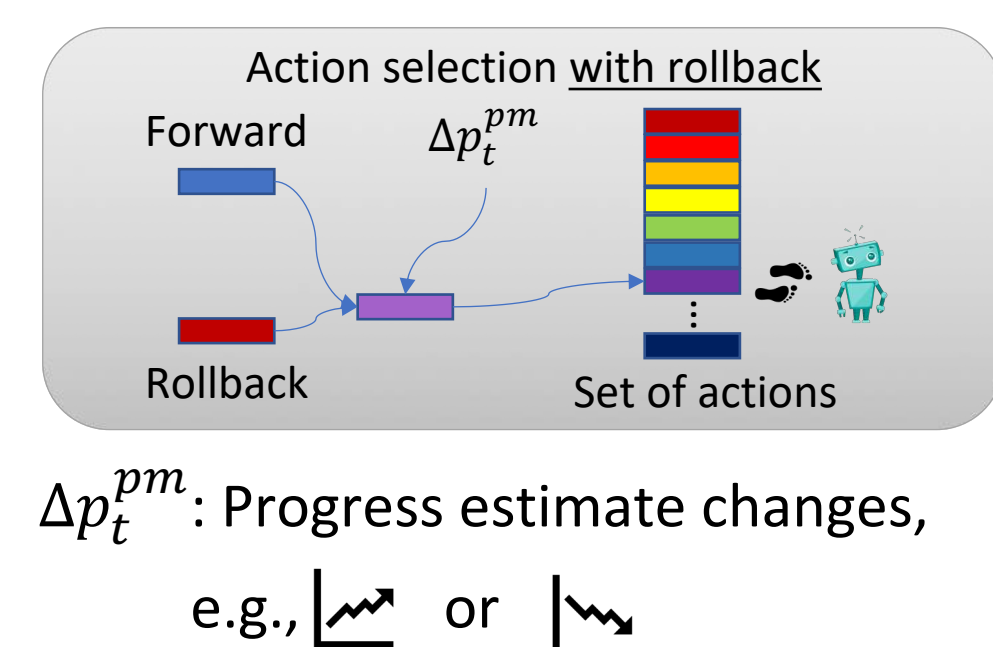


Leverages the self-monitoring mechanism through time to decide when to **roll back** to a previous location and resume the instruction following task.




2 BACKTRACKING FRAMEWORK

Regret Module: forward or rollback?




Progress Marker: which way to go?
Local Graph Search:

- know which directions have been visited
- estimate which one is likely lead to the goal.



3 QUANTITATIVE ANALYSIS



Evaluation metrics:


- Navigation Error (NE):** mean of the shortest path distance between the agent's final position and the goal location.
- Success Rate (SR):** the percentage of final positions less than 3m away from the goal location.
- Success rate weighted by Path Length (SPL):** SR weighted with trajectory lengths.

4 QUALITATIVE RESULTS & OTHER ANALYSIS

Walk down the hall way and make a right at the stairs and walk down the stairs. Make a hard left at the bottom of the stairs and wait by the Bamboo plant.



Number of unsuccessful examples involving **rollback** action reduced



Sanity check: manually blocking rollback

Val-Unseen	Rollback	NE	SR
Regretful	⊖	5.80	0.46
Regretful	✓	5.36	0.48

Reference:

- Anderson et al., "Vision-and-Language Navigation: Interpreting visually-grounded navigation instructions in real environments", CVPR 2018.
- Wang et al., "Look Before You Leap: Bridging Model-Free and Model-Based Reinforcement Learning for Planned-Ahead Vision-and-Language Navigation", ECCV 2018.
- Fried et al., "Speaker-Follower Models for Vision-and-Language Navigation", Neural 2018.
- Wang et al., "Reinforced Cross-Modal Matching and Self-Supervised Imitation Learning for Vision-Language Navigation", CVPR 2019.
- Ma et al., "Self-Monitoring Navigation Agent via Auxiliary Progress Estimation", ICLR 2019.

Acknowledgement:

This research was partially supported by DARPA's Lifelong Learning Machines (L2M) program, under Cooperative Agreement HR0011-18-2-001.