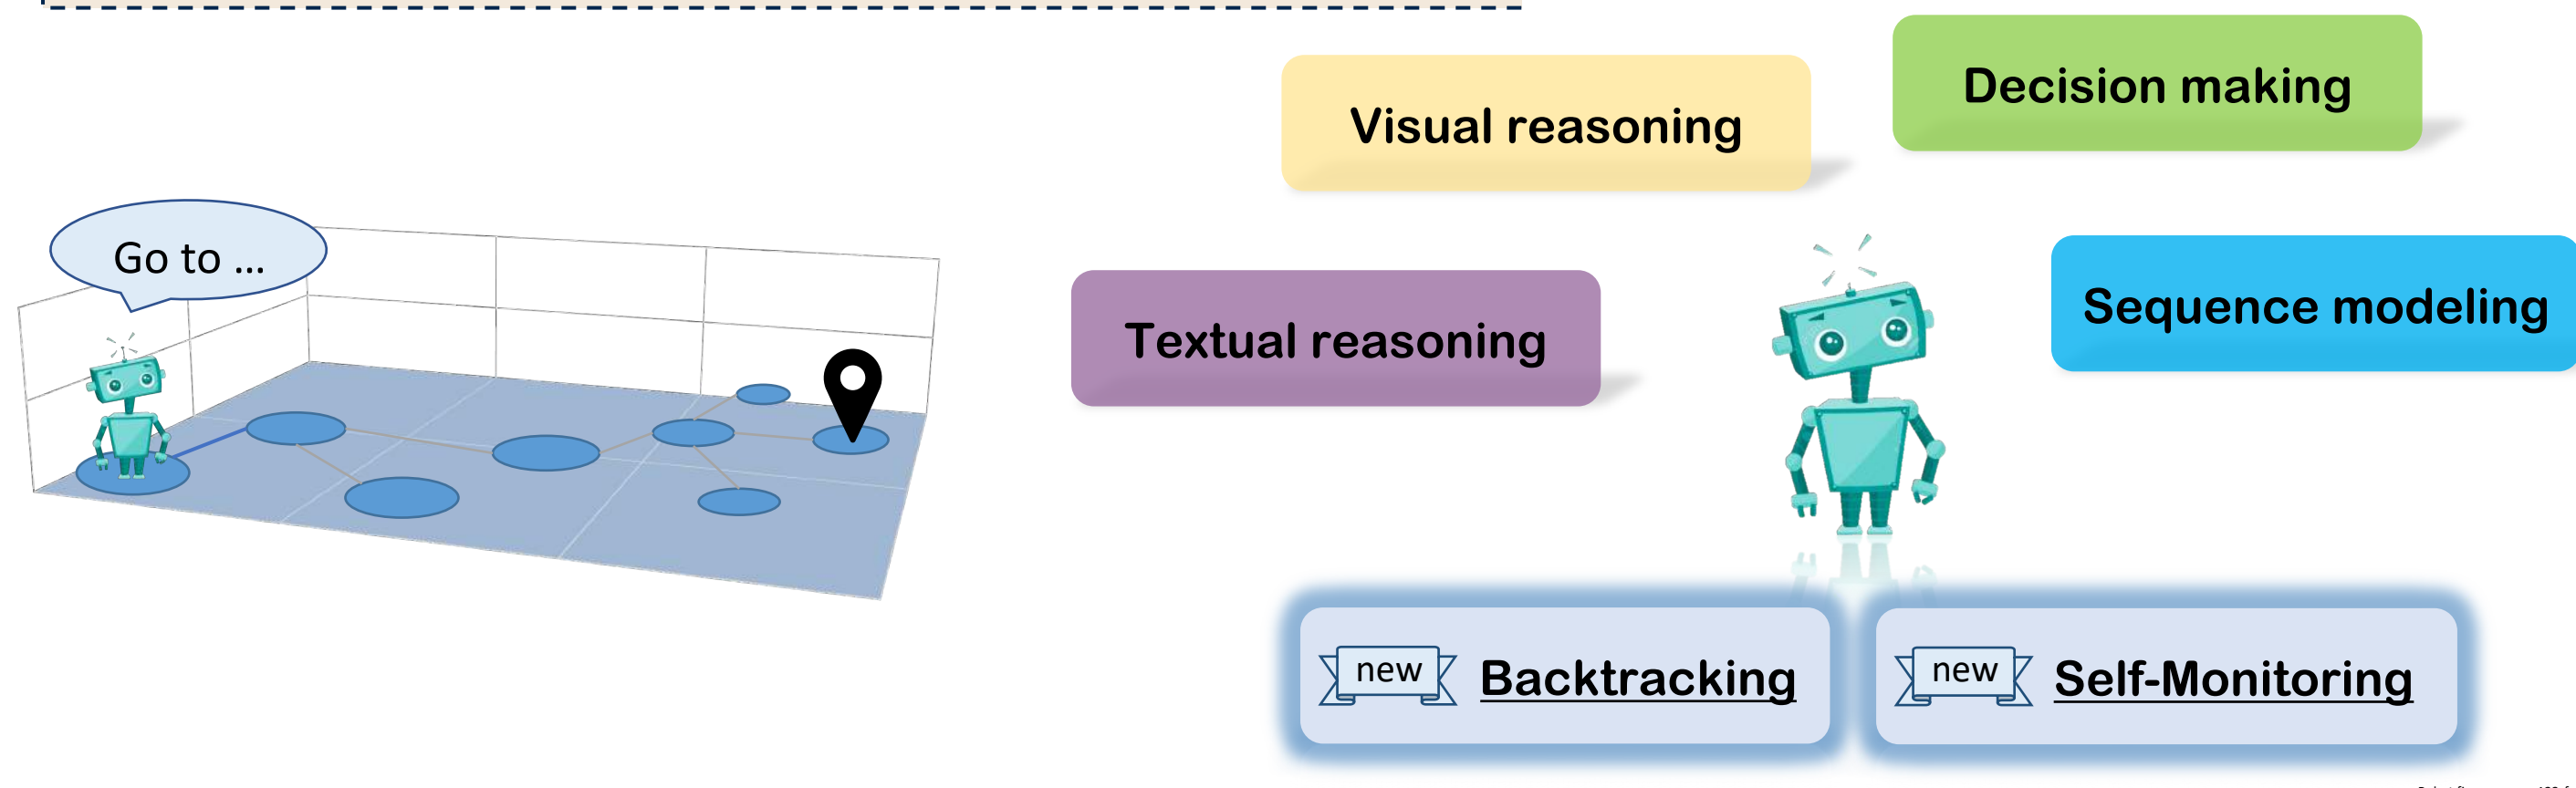



1 GOAL-DRIVEN REASONING

What do we need for goal-driven navigation?



2 VISION-AND-LANGUAGE NAVIGATION (VLN)

Exit the bedroom and go towards the table. Go to the stairs on the left of the couch. Wait on the third step.

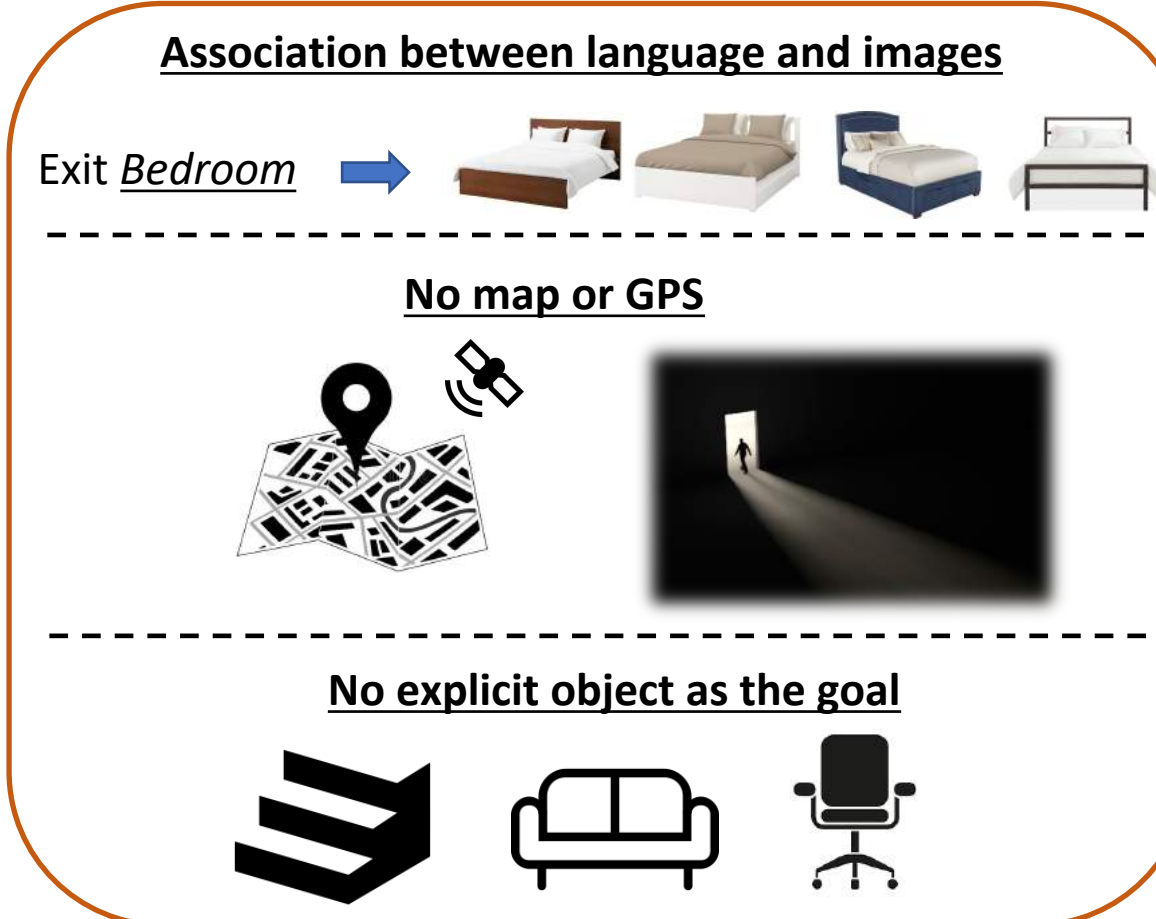


Matterport3D



Panoramic camera

Challenges:



3 INSTRUCTION GROUNDING


Instruction following:

We expect textual attention to match with agent's actions — ground instruction to actions.

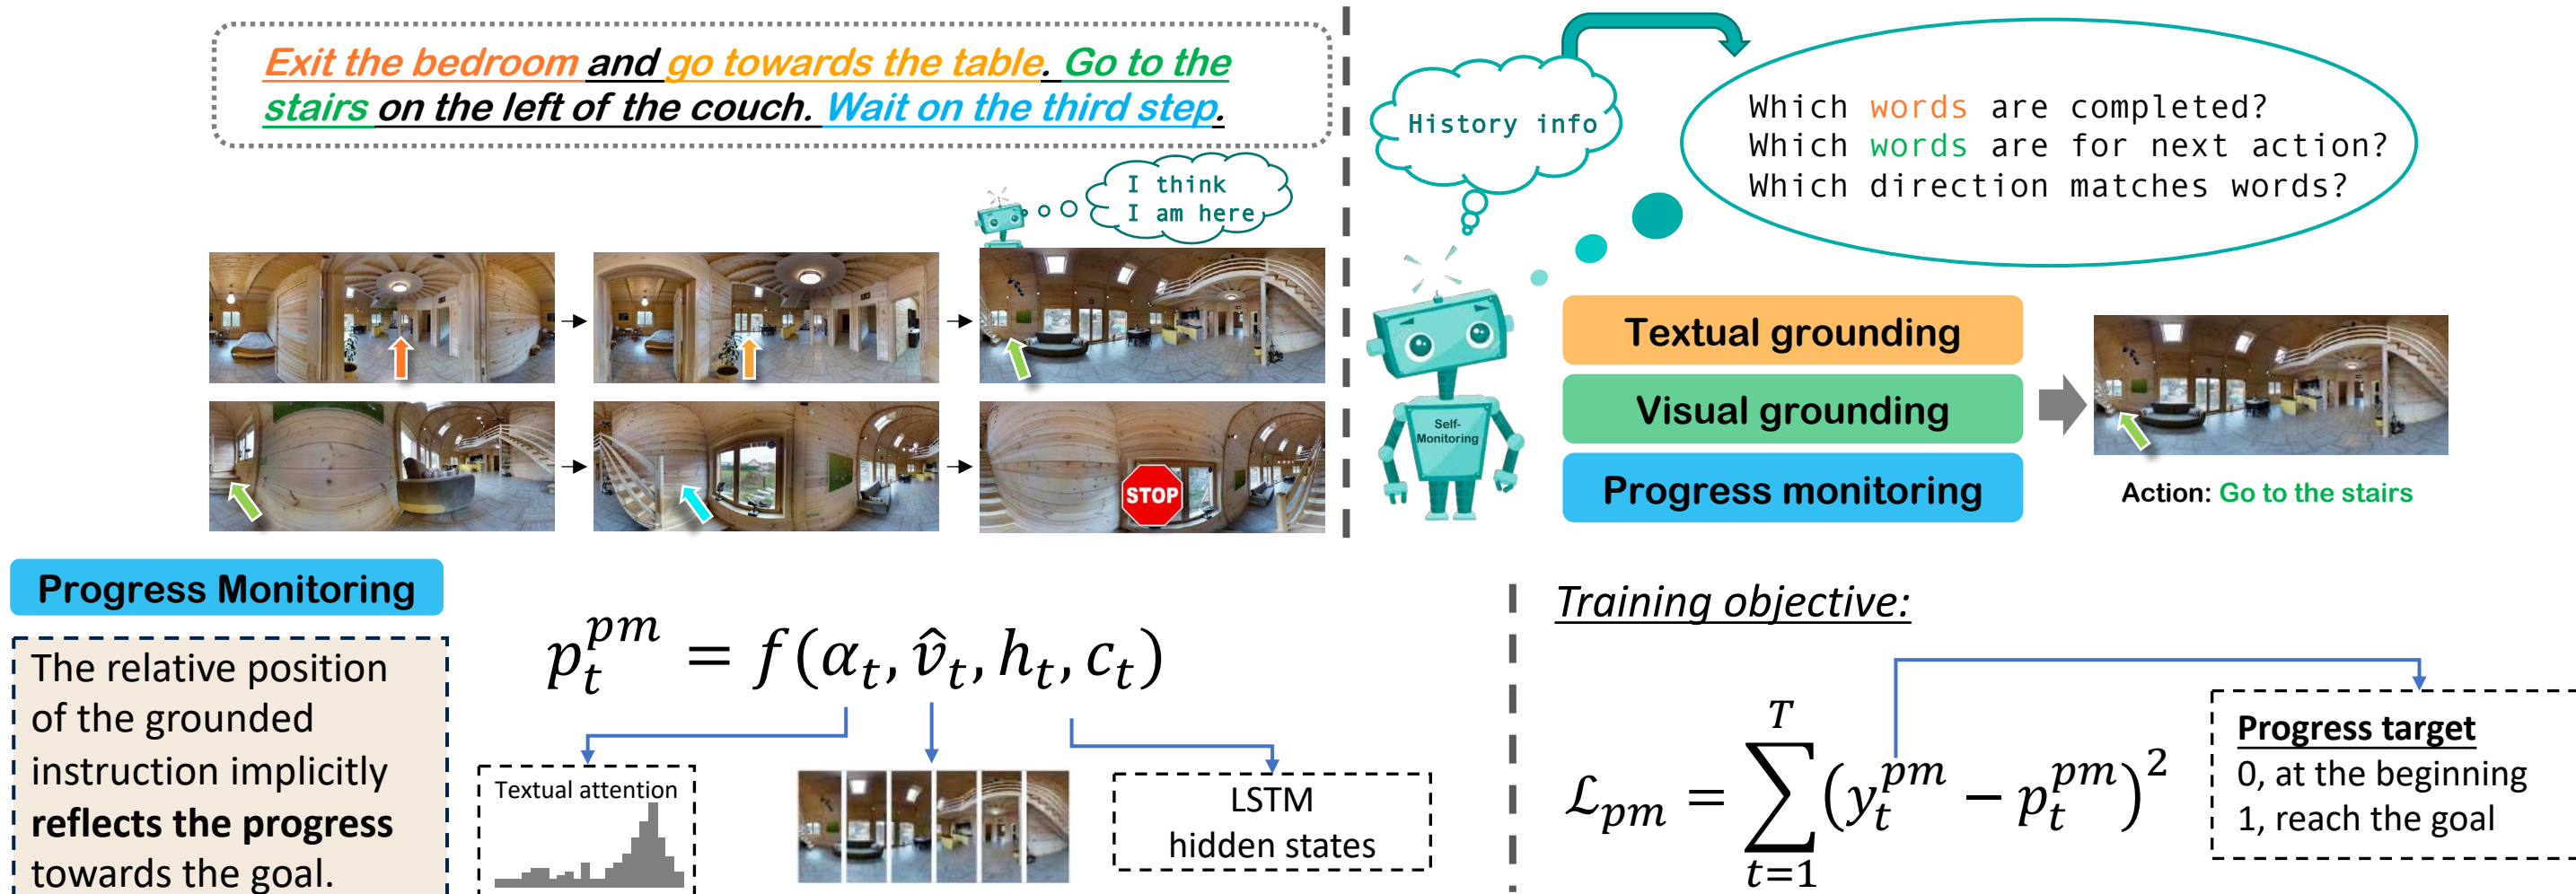
Exit the bedroom and go towards the table. Go to the stairs on the left of the couch. Wait on the third step.



Incorrect grounding from existing work




4 SELF-MONITORING AGENT



5 GROUNDING QUALITATIVE ANALYSIS



3 EXPERIMENTS & ANALYSIS

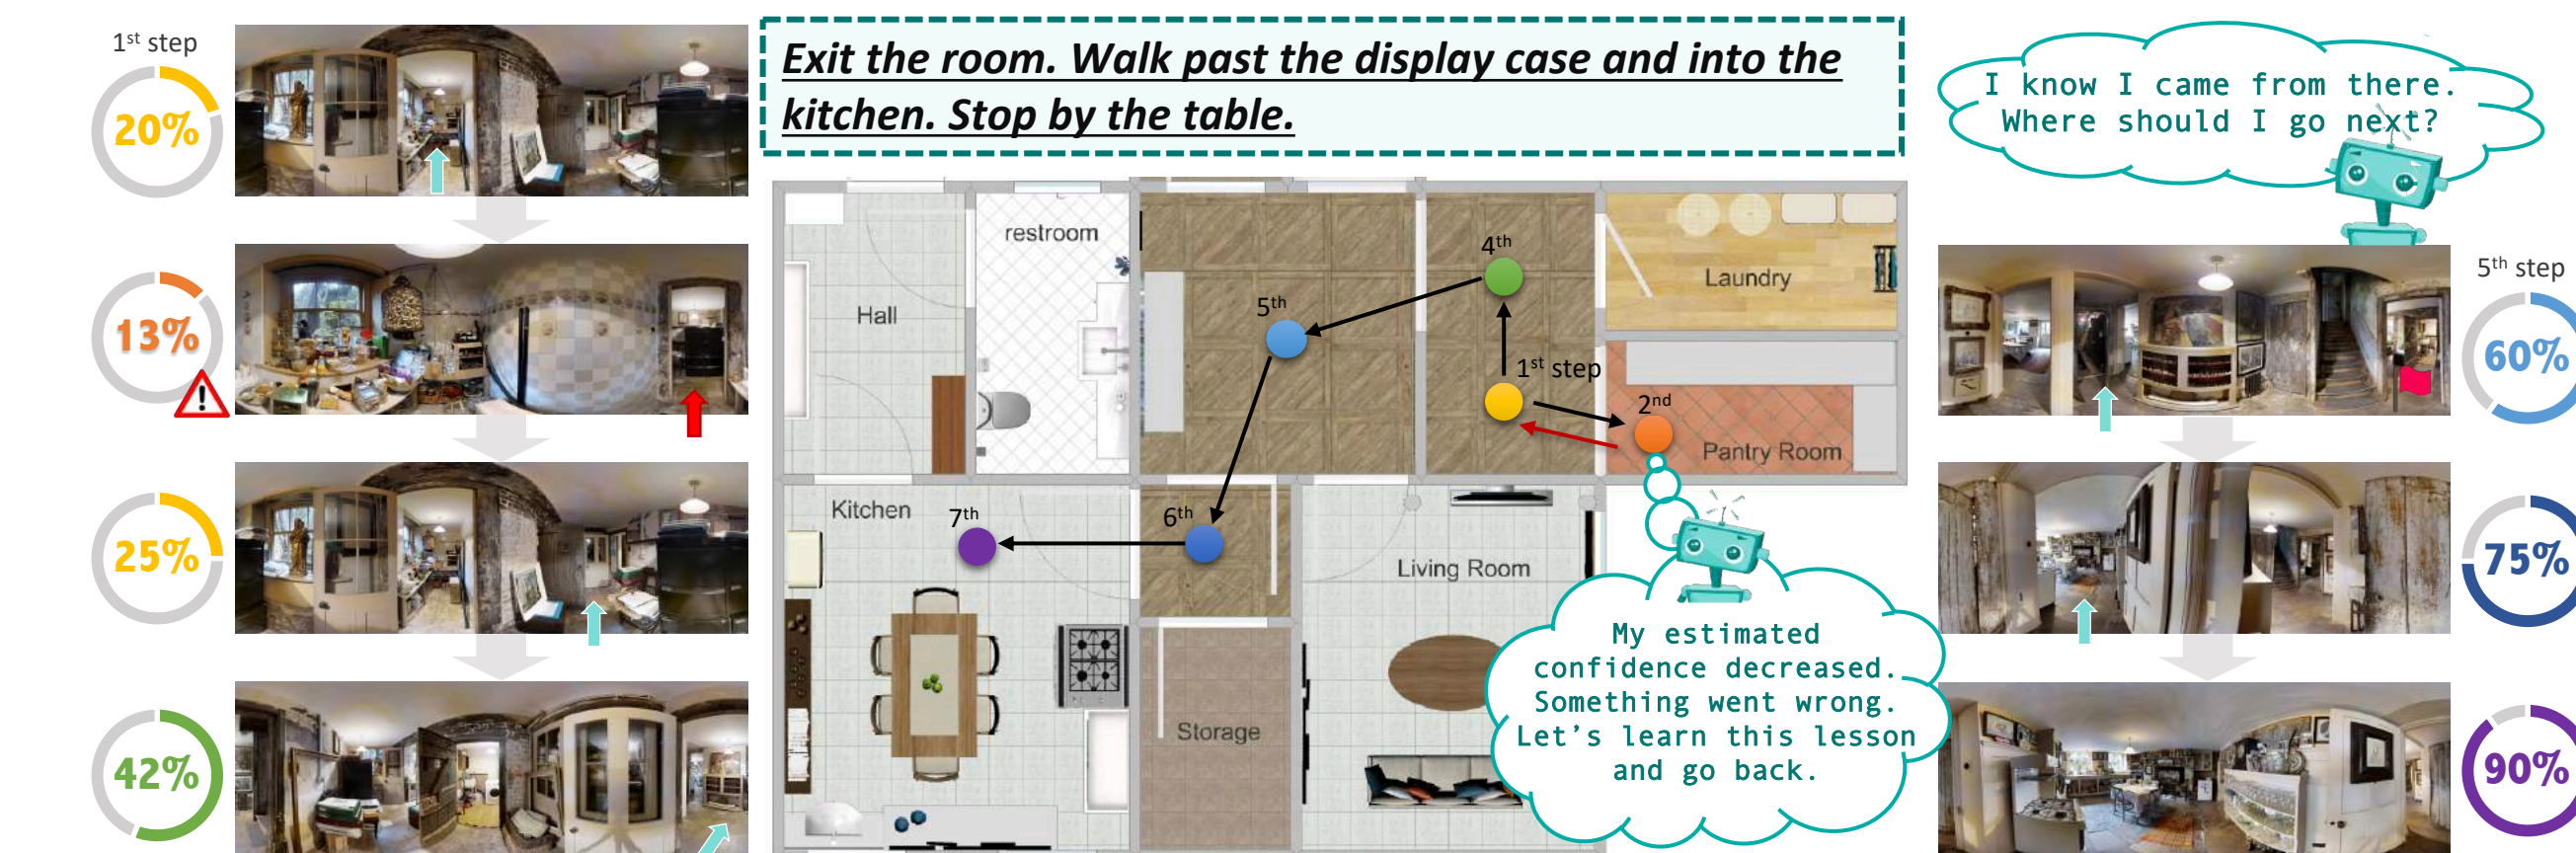


Evaluation metrics:

- 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.

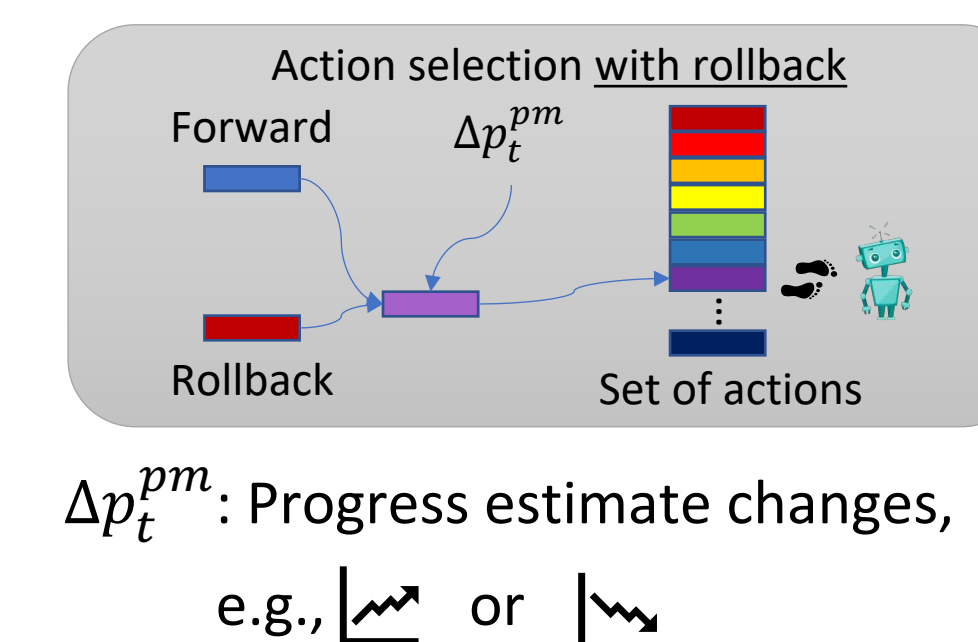
1 END-TO-END LEARNED BACKTRACKING AGENT

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?

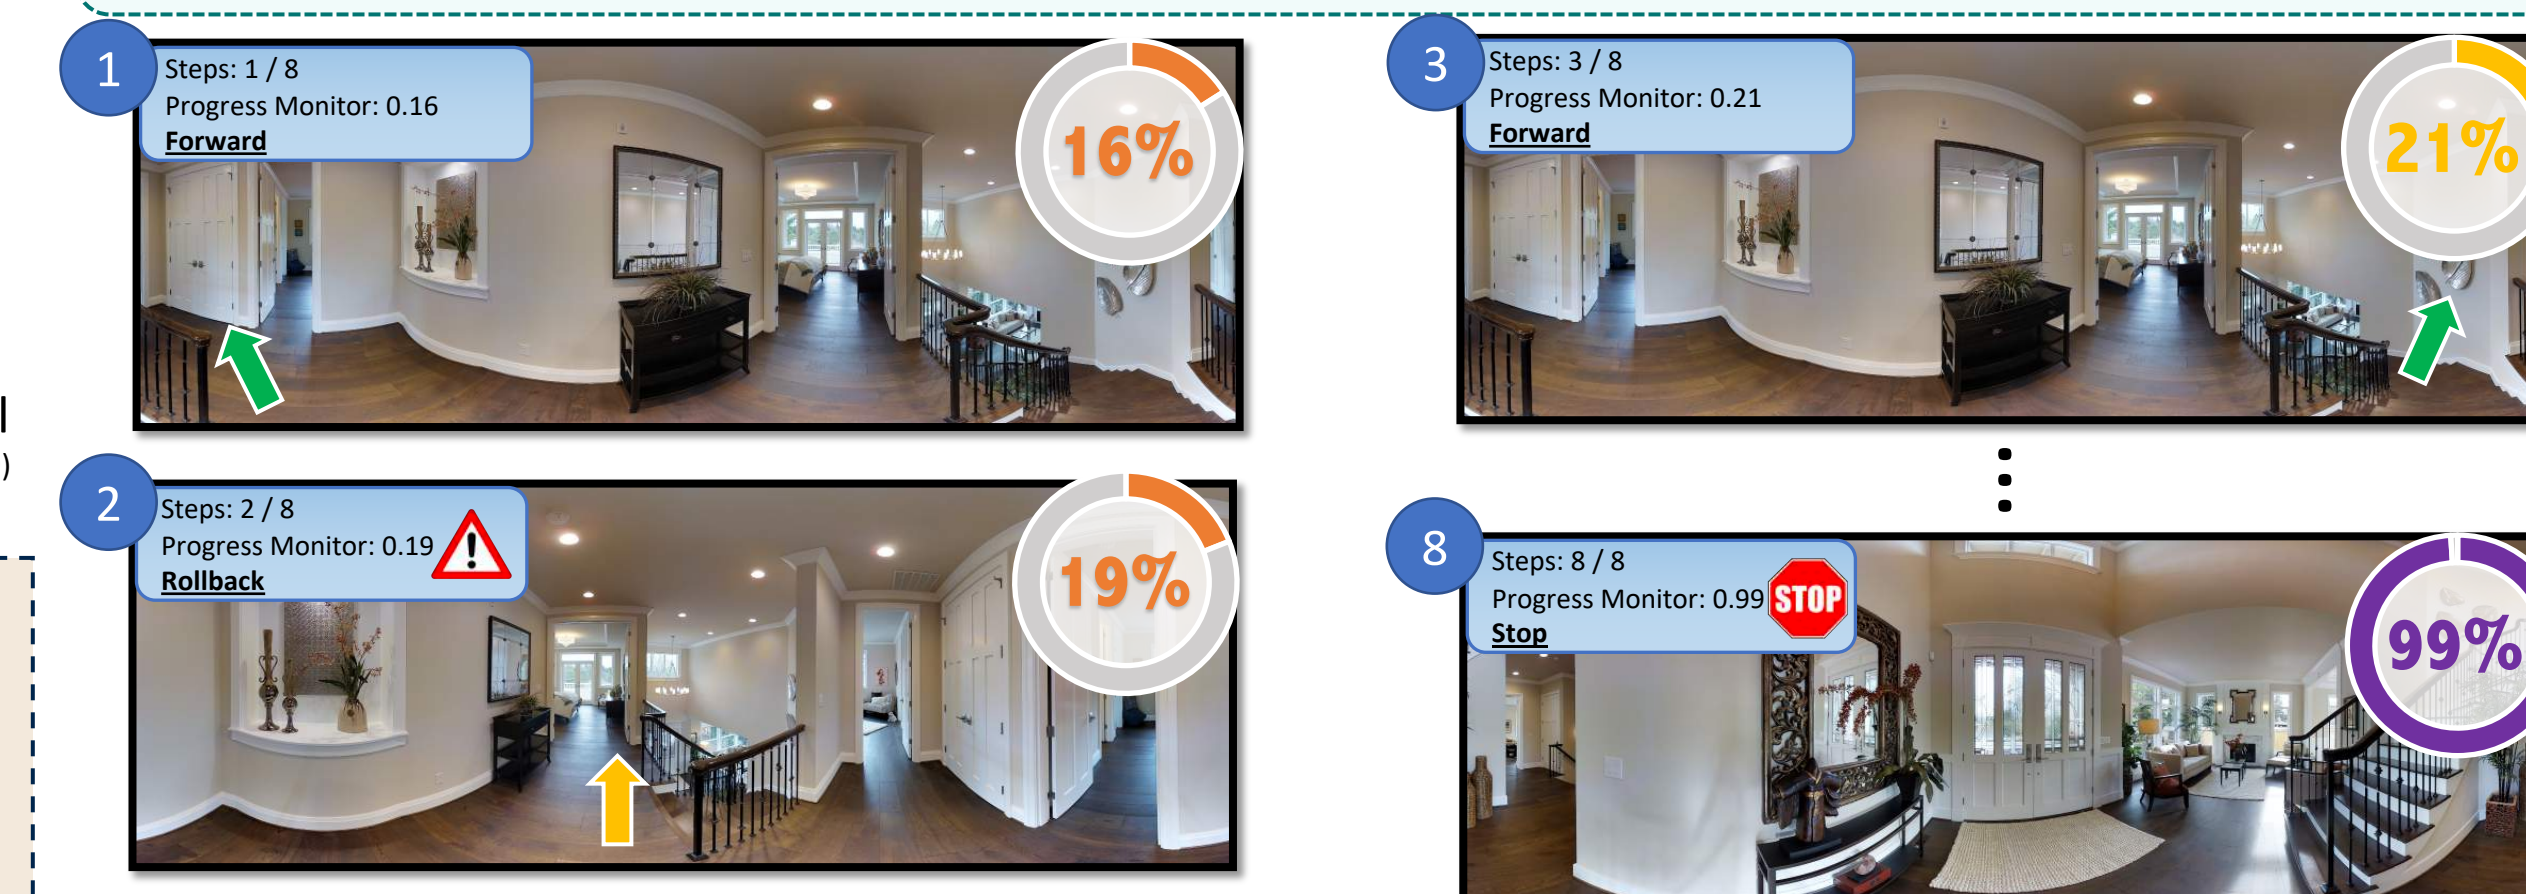


Progress Marker: which way to go?

- Local Graph Search: know which directions have been visited, estimate which one is likely lead to the goal.



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.



Sanity check: manually blocking rollback

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

Reference:

- [1] Anderson et al., "Vision-and-Language Navigation: Interpreting visually-grounded navigation instructions in real environments", CVPR 2018.
- [2] Wang et al., "Look Before You Leap: Bridging Model-Free and Model-Based Reinforcement Learning for Planned-Ahead Vision-and-Language Navigation", ECCV 2018.
- [3] Fried et al., "Speaker-Follower Models for Vision-and-Language Navigation", Neural 2018.
- [4] Wang et al., "Reinforced Cross-Modal Matching and Self-Supervised Imitation Learning for Vision-Language Navigation", CVPR 2019.
- [5] Ma et al., "The Regretful Agent: Heuristic-Aided Navigation through Progress Estimation", CVPR 2019.

Acknowledgement:

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