Fast and Reliable Online Retraining and Adaptation for Robot Planning Despite Missing World Knowledge

Gregory J. Stein, George Mason University
https://cs.gmu.edu/~gjstein

Overview

We aim to develop a principled approach for improving robot behavior during deployment for long-horizon planning in partially-mapped environments, emphasizing *reliability*, *data efficiency*, and *performance*.

Prior Work: learning-informed long-horizon planning under uncertainty

Under our model-based planning abstraction, subgoals are places to explore

The classroom is unlikely to reach the faraway goal.

The hallway connects long distances and is likely to reach the goal.

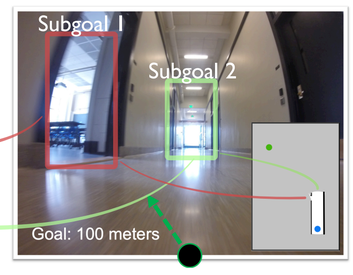

When executing subgoal-action 1...

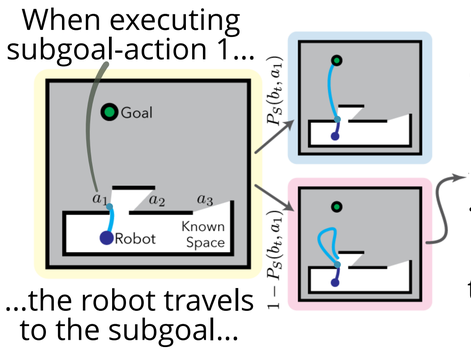

...the robot travels to the subgoal...

... there is a likelihood the action will succeed in reaching the goal,

... and a likelihood the goal cannot be reached, requiring that the agent then select another action.

Central Challenge: improving despite uncertainty

Planning well requires reasoning about uncertainty, often done using learning. When learning informs bad behavior, determining when and how to improve is difficult.

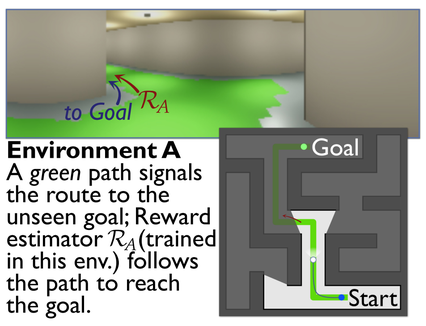

Environment A
A *green* path signals the route to the unseen goal; Reward estimator $\mathcal{R}_A$ (trained in this env.) follows the path to reach the goal.

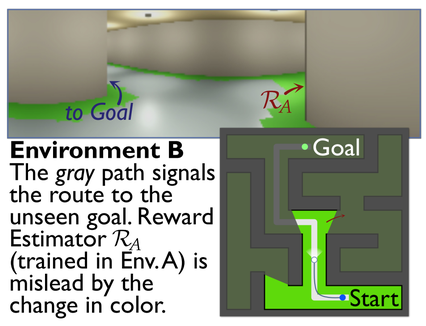

Environment B
The *gray* path signals the route to the unseen goal. Reward Estimator $\mathcal{R}_A$ (trained in Env. A) is mislead by the change in color.

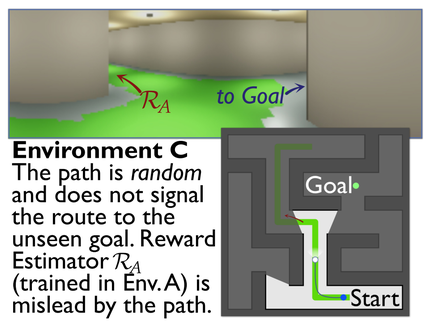

Environment C
The path is *random* and does not signal the route to the unseen goal. Reward Estimator $\mathcal{R}_A$ (trained in Env. A) is mislead by the path.

How can the robot use its own experience to quickly and reliably improve during deployment?

How can it update learned behavior to match expert guidance?

Model-based learning-informed planning will afford fast and reliable improvement

Develop a principled approach for fast and reliable *planner selection* despite uncertainty.

Our approach will help make *domain adaptation* more reliable, even without prior knowledge.

Input: from Env. B (no translation) **Output: Env. A-like** (translated from B)

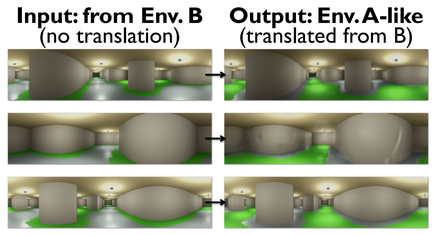

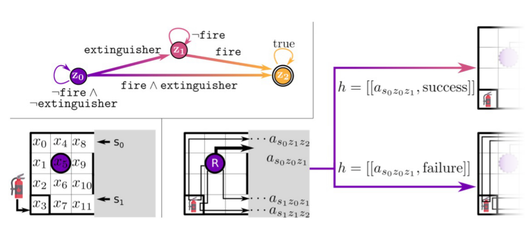

Leverage *expert guidance* to improve learning-informed behavior for multi-stage tasks.

Fast and Reliable Planner Selection

Deployment-time planner selection via learning-informed model-based planning.

Guarantees on sub-linear regret bounds on planner selection

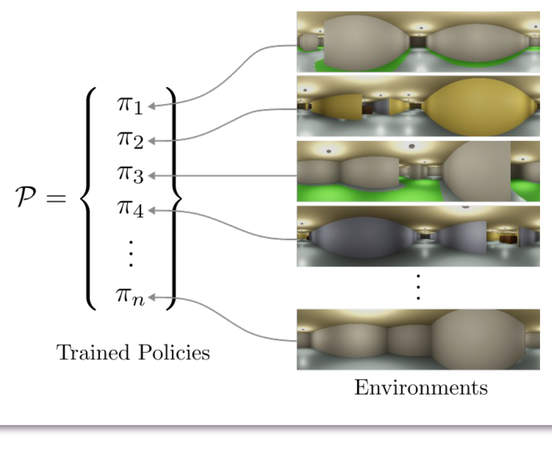

At the intersection (left), estimators $\mathcal{R}_A$ and $\mathcal{R}_{A \to B}$ disagree on where to go.

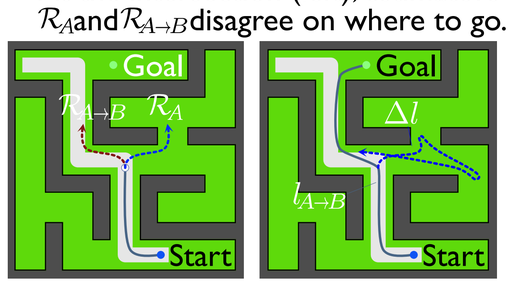

After reaching the goal with $\mathcal{R}_{A \to B}$, we ask how well $\mathcal{R}_A$ **could** have done, a cost that informs model selection.

Interventions via Expert Guidance

No Intervention **With Intervention**

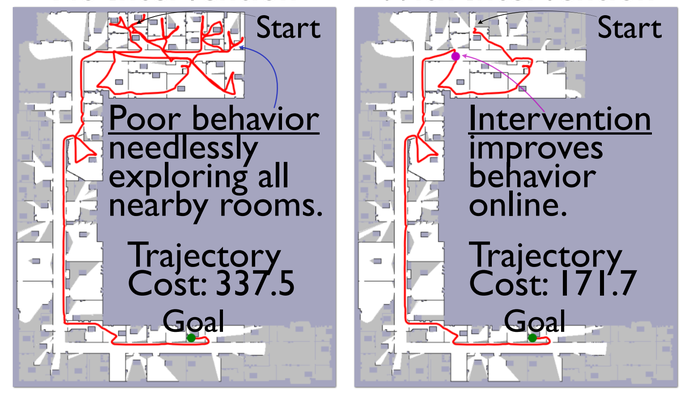

Societal Impacts

Our contributions will be a key enabler of home and service robots.

We will further develop tools for explainable AI under uncertainty.

Our work will help democratize the training of autonomous decision-making systems.

Image Credit: https://www.pexels.com/photo/bed-empty-equipments-floor-236380/

Integrating Research, Education, and Outreach

Mentorship of two undergraduate researchers during academic terms

Summer mentorship of 1–2 high-school summer students through the GMU "ASSIP" program.

Tools will help lower the barrier to entry for robot learning; integration of contributed tools into robotics curriculum at GMU.

Publications

[1] Gregory J. Stein, Christopher Bradley, and Nicholas Roy. “Learning over Subgoals for Efficient Navigation of Structured, Unknown Environments”. In: Conference on Robot Learning (CoRL). 2018.

[2] Christopher Bradley, Adam Pacheck, Gregory J. Stein, Sebastian Castro, Hadas Kress-Gazit, and Nicholas Roy. “Learning and Planning for Temporally Extended Tasks in Unknown Environments”. In: International Conference on Robotics and Automation (ICRA). 2021.

[3] Gregory J. Stein. “Generating High-Quality Explanations for Navigation in Partially-Revealed Environments”. In: Advances in Neural Information Processing Systems (NeurIPS). 2021.

[4] Abhishek Paudel and Gregory J. Stein. “Data-Efficient Policy Selection for Navigation in Partial Maps via Subgoal-Based Abstraction”. ArXiv (under review). 2023.

2023 FRR & NRI Principal Investigators' Meeting
May 2-3, 2023

Award ID#: 2232733