

Learning Augmented, Multi-Robot Long-Horizon Navigation in Partially Mapped Environments

Abhish Khanal and Gregory J. Stein

Abstract—We present a novel approach for efficient and reliable goal-directed long-horizon navigation for a multi-robot team in a structured, unknown environment by predicting statistics of unknown space. Building on recent work in learning-augmented model based planning under uncertainty, we introduce a high-level state and action abstraction that lets us approximate the challenging Dec-POMDP into a tractable stochastic MDP. Our Multi-Robot Learning over Subgoals Planner (MR-LSP) guides agents towards coordinated exploration of regions more likely to reach the unseen goal. We demonstrate improvement in cost against other multi-robot strategies; in simulated office-like environments, we show that our approach saves 13.29% (2 robot) and 4.6% (3 robot) average cost versus standard non-learned optimistic planning and a learning-informed baseline.

I. INTRODUCTION

We aim to navigate through an unknown environment using multiple robots to find an unseen point goal in minimum expected distance: e.g., for package retrieval. To perform well, the multi-robot team needs to collectively navigate the unexplored region, seeking out promising routes to the goal while avoiding regions that typically lead to dead-ends.

Planning well in an unknown environment requires making inferences about unseen parts of the environment; in theory, robots must envision all possible configurations of the unknown space—including regions unlikely to lead to the goal—to determine how to navigate so that they can most quickly reach the unseen goal. Multi-robot planning under uncertainty can be modeled as a Decentralized Partially Observable Markov Decision Processes (Dec-POMDP) [1], [2]. However, Dec-POMDP planning is computationally intractable in general [3] and it requires access to a distribution over possible environments, difficult to obtain in general.

Learning is often used to help make inferences about unseen space needed to inform good behavior and is an increasingly powerful tool for planning under uncertainty [4], [5] and Dec-POMDP planning [6]. However, despite impressive progress in this domain, particularly for model-free approaches trained via deep reinforcement learning [7], [8], [9], many such strategies struggle to learn effectively in large-scale environments and can be brittle in practice [10].

So as to avoid the computational expense of planning over the space of short-time-horizon primitive actions, many approaches for multi-robot planning introduce a topological action abstraction that simplify planning [11], [12], [13]. Such approaches typically constrain robot motion so that

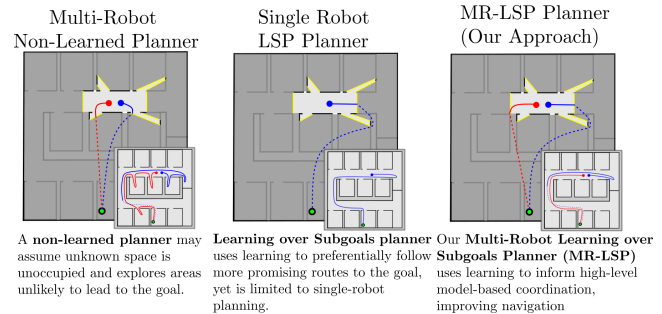


Fig. 1. Our Multi-Robot Learning over Subgoals Planner (MR-LSP) uses learning to guide robots toward coordinated exploration of regions more likely to reach the unseen goal.

each robot intends to leave known space through different *frontiers*, each a boundary between free and unseen space, or via paths constrained to belong to different relative homology [14], [15]. Though these action abstractions simplify planning and coordination between robots, action selection typically relies on simple, greedy heuristics to decide where each robot should navigate next and do not alleviate the challenges of predicting the impact of an action.

The recent Learning over Subgoals Planning abstraction (LSP) [16] overcomes this limitation for single-robot planning. In LSP, learning is used to estimate the goodness of high-level (frontier-associated) actions that enter unseen space, including both the likelihood that such an action will reach the goal and its expected cost. Learning augments a model-based planning abstraction, affording performant and reliable navigation under uncertainty. However, despite LSP’s improved performance in single-robot planning under uncertainty, the LSP state transition model is not straightforwardly extended to support multiple robots, which must have the capacity to concurrently explore different unseen regions.

Leveraging insights from multi-robot topologically constrained planning and the recent Learning over Subgoals Planning (LSP) abstraction, we introduce a multi-robot generalization of the LSP model: Multi-Robot Learning over Subgoal Planning (MR-LSP). Our approach supports multi-robot planning, is reliable, and leverages learning to inform where robots should navigate next. Each robot’s high-level (topologically-constrained) actions correspond to navigation to a *subgoal*—associated with a frontier—and then navigation beyond in an effort to reach the goal. We introduce a new state abstraction and transition model that allows our high-level planner to envision how the robot team will redistribute effort once each robot finishes their respective exploratory action, a key feature of coordinated multi-robot planning. Our abstraction lets us approximate the challenging Dec-

Abhish Khanal and Gregory J. Stein are with the Department of Computer Science, George Mason University, Fairfax, VA, 22030, USA. {akhana17, gjstein}@gmu.edu

POMDP as a stochastic MDP and solve it via sample-based tree search. Learning is used to estimate the goodness of each action and informs planning via a Bellman Equation for our model-based planning abstraction.

We demonstrate the effectiveness of our approach in a simulated office floorplan environment, showing that our approach reduces cost by 13.29% (two robots) and 4.6% (three robots) versus standard non-learned optimistic planning and a competitive learning-informed baseline of our own design.

II. PROBLEM FORMULATION

Our multi-robot team is placed in a partially-mapped environment and tasked to find a point-goal located in unseen space in minimum expected cost (distance). Each robot is equipped with a planar laser scanner, which it uses to localize and build an accurate map of its local surroundings, a problem setting recently coined by Merlin et al. [17] as a *Locally Observable Markov Decision Process* (LOMDP). We assume lossless communication between robots so that all robots have an up-to-date partial map represented as an occupancy grid, as revealed by all members of the team and the poses of each robot—collectively, the team’s belief b_t . Planning is centrally coordinated, and so at each time step, the team’s collective action a_t specifies how each robot should move so as to make progress towards the unseen goal in an effort to minimize the expected cost.

Formally, we represent this problem as a Partially Observable Markov Decision Process [18], [19] (POMDP). The expected cost Q under this model can be written via a belief space variant of the Bellman equation [20]:

$$Q(b_t, a_t) = \sum_{b_{t+1}} P(b_{t+1}|b_t, a_t) \left[R(b_{t+1}, b_t, a_t) + \min_{a_{t+1} \in \mathcal{A}(b_{t+1})} Q(b_{t+1}, a_{t+1}) \right] \quad (1)$$

Where $R(b_{t+1}, b_t, a_t)$ is the cost accumulated by reaching belief state b_{t+1} from b_t by taking action a_t .

III. PRELIMINARIES: SINGLE-ROBOT PLANNING IN A PARTIAL MAP VIA LEARNING OVER SUBGOALS

Even for single robot planning, planning via the POMDP model in Eq. (1) requires both enormous computational effort and also access to a distribution over possible environments, difficult to obtain in general. So as to mitigate the complexities of *single-robot* navigation in a partial map, the *Learning over Subgoals Planner* (LSP) [16] introduces an abstraction in which actions correspond to navigation through *subgoals* placed at boundaries between free and unknown space, simplifying the process of imagining the impact of actions that enter unseen space. Under the LSP abstraction, model-based planning is augmented via predictions from learning, allowing for both reliability and good performance.

For the LSP abstraction, temporally-extended actions correspond to *subgoals*, each associated with a contiguous boundary between free and unseen space. A *high-level action* a_t consists of (1) navigating to the subgoal and then (2) exploring the unknown space beyond in an effort to reach the unseen goal. Planning is done over an abstract belief state:

a tuple $b_t = \langle m_t, \mathcal{S}_u, q_t \rangle$, where m_t is the (partial) map of the environment, \mathcal{S}_u is the set of unexplored subgoals, and q_t is the robot pose. Each high-level action $a_t \in \mathcal{S}_u$ has a binary outcome: with probability $P_S(a_t)$, the robot *succeeds* in reaching the goal or (with probability $1 - P_S(a_t)$) *fails* to reach the goal.

Upon selecting an action a_t , the robot must first move through known space to the boundary, accumulating a cost $D(m_t, q_t, a_t)$. If the robot succeeds in reaching the goal, it accumulates a *success cost* $R_S(a_t)$, the expected cost for robot to reach the goal, and navigation is complete. If it fails to reach the goal, the robot accumulates a cost associated with exploring the region and needing to turn back $R_E(a_t)$ and the state is updated to reflect that the robot has moved (to position $q(a_t)$) and that the subgoal associated with a_t is explored: $b_{t+1} = \langle m_t, \mathcal{S}_u \setminus \{a_t\}, q(a_t) \rangle$. Upon failing to reach the goal, the robot must subsequently select another action from the set of unexplored subgoals: $a_{t+1} \in \mathcal{S}_u \setminus \{a_t\}$.

The expected cost of a high-level action a_t can be written as a Bellman Equation:¹

$$Q(b_t, a_t \in \mathcal{S}_u) = D(m_t, q_t, a_t) + P_S(a_t)R_S(a_t) + (1 - P_S(a_t)) \left[R_E(a_t) + \min_{a_{t+1} \in \mathcal{S}_u \setminus \{a_t\}} Q(b_{t+1}, a_{t+1}) \right] \quad (2)$$

The terms P_S , R_S , and R_E —too difficult to compute exactly—are estimated from images collected on board the robot via learning.

IV. MULTI-ROBOT LEARNING OVER SUBGOALS PLANNING (MR-LSP)

For single robot planning, the Learning over Subgoals planning paradigm (LSP) has demonstrated state-of-the-art performance under uncertainty and reliability-by-design, despite its reliance on learning. Here, we extend the LSP state and action abstraction to support multi-robot planning.

While the LSP action abstraction is designed for single robot planning, generalizing the space of high-level actions to incorporate multi-robot planning is straightforward—the *collective* high-level action assigns each robot a subgoal for it to navigate towards and explore beyond. As such the collective action for an N -robot team can be written as a list of subgoals: $a_t = [\sigma_1, \sigma_2, \dots, \sigma_N]$.

However, generalizing the LSP model to support multiple robots is made complicated by the fact that different robots are executing their respective subgoal-actions *concurrently* and therefore may finish exploration at different times; when one robot has finished exploration beyond one subgoal, the others may not yet be done and may not even have reached the unseen space they seek to explore. Moreover, whenever one robot completes a subgoal-action, the planner has the capacity to reassign the actions of each robot and thus where they should travel next. If planning is to take into

¹The statistics of belief (P_S , R_S , and R_E) are used as input to the planner for planning, instead of the belief b_t , in order to transform the challenging POMDP in Eq. (1) into a simpler stochastic MDP in Eq. (2).

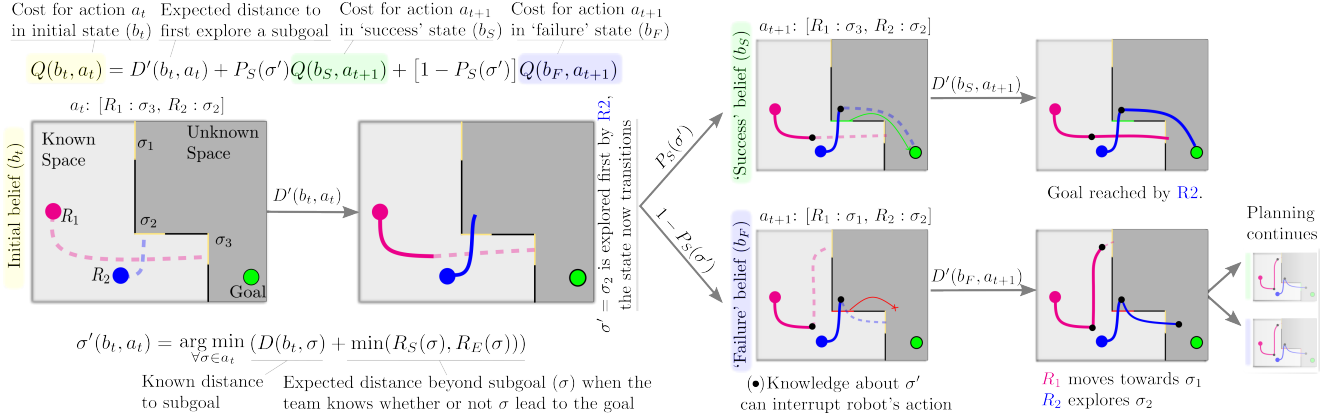


Fig. 2. **Schematic of MR-LSP multi-robot expected cost calculation** A schematic showing how the cost of collective-action under a policy is calculated using the Multi-Robot Learning over Subgoals (MR-LSP) abstraction. For action a_t , the robot-team concurrently explores multiple subgoals until one of the robot discovers whether its assigned subgoal leads to the goal, leading to a transition in the abstract belief state.

account the impact of concurrent action execution, we must augment the LSP state abstraction and state transition model to incorporate these effects.

In this section, we introduce our Multi-Robot Learning over Subgoals Planner (MR-LSP), which introduces a new state and state transition abstraction to model concurrent action execution, an essential component of multi-robot coordination. Key to our approach is the recognition that estimates of the expected time to complete an action (the costs of success R_S and failure R_E) tell us which robot is expected to complete its exploration first, knowledge we can use to simplify the process of imagining the future during planning. We introduce a Bellman Equation for our new multi-robot planning abstraction (Sec. IV-A), which we use in combination with a Monte-Carlo Tree Search-based approach to compute the expected cost of multi-robot collective action (Sec. V).

A. Expected Cost of Multi-Robot High-Level Actions

Under the MR-LSP model, the abstract state of the environment is a tuple $b_t = \langle m_t, \mathcal{S}_u, \mathcal{S}_g, q_t \rangle$, where m_t is the map, \mathcal{S}_u is the set of unexplored subgoals, \mathcal{S}_g is the set of subgoals that are known to lead to the goal and q_t is a list of the robot poses. A high-level *collective action* a_t assigns each robot a subgoal to explore. For an abstract state b_t , the set of high-level actions $\mathcal{A}(b_t)$ specifies all collective-actions for the robot team: $\mathcal{A}(b_t) = \otimes_{i \in I} (\mathcal{S}_u \cup \mathcal{S}_g)$, constrained such that no two robots can explore same subgoal whenever possible. A collective-action a_t is thus a list of subgoals, that specifies to which subgoal each robot will next navigate to and explore beyond in an effort to reveal the goal.

Key to defining our state transition model is the idea that the multi-robot team can select a new collective-action whenever a *single* robot reveals whether or not a subgoal will lead to the goal; the high-level state is updated when the *first* of the robot's subgoal-actions is completed. As in the LSP model, under our multi-robot abstraction we estimate how long it will take a single robot to either reach the goal ($D + R_S$) or to explore ($D + R_E$) beyond a particular subgoal s . Given estimates of the costs of success R_S and

exploration R_E , we can determine which of the robots is expected to complete its high-level action first, yielding both (i) the subgoal σ' that the team will first discover either does or does not lead to the goal and (ii) how long this discovery will take D' . Mathematically,

$$D'(b_t, a_t) = \min_{\forall \sigma \in a_t} (D(b_t, \sigma) + \min(R_S(\sigma), R_E(\sigma)))$$

$$\sigma'(b_t, a_t) = \operatorname{argmin}_{\forall \sigma \in a_t} (D(b_t, \sigma) + \min(R_S(\sigma), R_E(\sigma))) \quad (3)$$

Under our model, the outcome of a collective action a_t reveals only whether σ' leads to the goal or not, and so the outcome of that high-level collective action is binary. After the execution of collective action (a_t), with probability $P_S(\sigma')$ the state transitions to an abstract *success state* (b_S) where the robots can reach goal from subgoal σ' or with probability $1 - P_S(\sigma')$ the state transitions to an abstract *failure state* (b_F) where the robots cannot reach the goal from subgoal σ' . Under both success and failure, the subgoal σ' becomes explored, and so is removed from \mathcal{S}_u , and the robot have traveled a distance D' towards completing their respective high-level subgoal-actions, ending up at new poses $q_t(a_t, D')$. The tuple representation of abstract states associated with success (b_S) and failure (b_F) to find a path to the goal is:

$$b_S = \langle m_t, \mathcal{S}'_u = \mathcal{S}_u \setminus \{\sigma'\}, \mathcal{S}'_g = \mathcal{S}_g \cup \{\sigma'\}, q_t(a_t, D') \rangle$$

$$b_F = \langle m_t, \mathcal{S}'_u = \mathcal{S}_u \setminus \{\sigma'\}, \mathcal{S}'_g = \mathcal{S}_g, q_t(a_t, D') \rangle \quad (4)$$

Given our state and action abstraction and state transition model, we can write a Bellman Equation that defines the expected cost Q for a collective action a_t (see also Fig. 2):

$$Q(b_t, a_t \in \mathcal{A}(b_t)) = D' + P_S(\sigma') \min_{a_{t+1} \in \mathcal{A}(b_S)} Q(b_S, a_{t+1}) + [1 - P_S(\sigma')] \min_{a_{t+1} \in \mathcal{A}(b_F)} Q(b_F, a_{t+1}) \quad (5)$$

Difficult to compute exactly, the terms P_S , R_S , and R_E are estimated via learning from images collected on the robot; see Sec. VI for a discussion of our neural network structure and training process. Owing to the combinatorial explosion of possible actions, we cannot tractably plan while including

all subgoals in the set of candidate actions. Following the example of LSP [16], we limit the number of subgoals under consideration to seven for all MR-LSP experiments.

We note that when a single robot is used, our MR-LSP equation reduces to the single-robot LSP model of Stein et al. [16], and so our approach extends the original LSP model to support multi-robot planning.²

B. Navigation via MR-LSP

The high-level collective action defines the long-horizon robot behavior. Upon selecting a collective action a_t via Eq. (5), each robot makes progress towards each subgoal. We use an A* plan cost computed over the observed grid to select motion primitives for each robot that make progress towards their assigned subgoal. The robots (i) move according to these primitive actions, (ii) observe their surroundings, (iii) update the partial map and the set of subgoals, and (iv) compute a new collective action based on this newly-updated map. This process repeats until the goal is reached.

V. COMPUTING MR-LSP EXPECTED COST VIA SAMPLE-BASED TREE SEARCH

Despite our action abstraction to simplify planning, the space of collective actions can become large in practice, making calculating cost via Eq. (5) computationally intensive. Instead, we rely on sample-based high-level planning and use a Partially Observable UCT (PO-UCT) [21]—a variant of Monte-Carlo tree search. PO-UCT is an anytime planning algorithm and allows us to approximate the expected cost without the need to exhaustively simulate all states.

During expansion of the planning search tree in PO-UCT, we maintain a *rollout history* associated with each node: $\mathcal{H}_{b_t} = [[a_0, n_0, Q_0], \dots, [a_k, n_k, Q_k]]$, that retains each node’s history of (i) executed high-level actions a_t and their outcomes, (ii) the number of times the node has been visited, and (iii) the accumulated expected cost Q_t up to that node. By simulating a collective action a_t from the abstract state b_t , finding the subgoal σ' via Eq. (3), we expand the tree stochastically, where the outcome of a particular action is sampled from a Bernoulli distribution parametrized by the estimated $P_S(\sigma')$. After the action outcome is sampled, the belief transitions to a node either corresponding to a success (b_S) or failure state (b_F), as defined in Eq. (4).

Rollouts proceed similar to other Monte-Carlo Tree Search approaches. The cost of each node corresponding to belief state (b_S or b_F) is the sum of cost accrued to reach the current belief state from the initial belief state (b_t) and a search heuristic corresponding to the lower bound cost for the team to reach the goal if unseen space were assumed to be unoccupied. After each node is visited, its count is incremented, which is used to control the rate of exploration during traversal. In all our experiments, we use 15,000 samples at each planning step.

²Owing to a change in how exploration is treated during concurrent action execution in our MR-LSP model—see Eq. (3)—the single-robot exploration cost R_E of Eq. (2) becomes $R_E \leftarrow \min(R_S, R_E)$, a slight deviation from the original LSP definition. With this (small) change, one-robot MR-LSP via Eq. (5) is equivalent to single-robot LSP planning via Eq. (2).

VI. TRAINING DATA GENERATION AND LEARNING SUBGOAL PROPERTIES

To compute the expected cost during MR-LSP planning, we require the subgoal properties P_S , R_S , and R_E for all subgoals. We train a convolutional neural network, similar to that of [22], to estimate these properties from images collected by the robot.

The convolutional neural network (CNN) takes a 128×512 RGB panoramic image aligned towards subgoal, egocentric position of the subgoal and goal as inputs. The image is passed through 4 convolutional layers, after which the relative-distance features are concatenated, passed through an additional 9 convolutional layers, and finally 5 fully connected layers, which output the subgoal properties. Our CNN estimates the subgoal properties for all subgoals and, with distances D computed from the occupancy grid, are used to compute cost of a collective-action via Eq. (5).

To generate training data, we use a non-learned optimistic planner to navigate previously unseen environments. Observations (images) and are collected from every step and labels of for each subgoal for P_S , R_S , and R_E are computed from the underlying known map. $P_S = 1$ if the subgoal leads to the goal, for which R_S is the distance to reach the goal, and $P_S = 0$ if it does not, where R_E is the distance the robot would travel before reaching a dead end and turning around.

VII. EXPERIMENTAL RESULTS

We conduct simulated experiments in two different environments—our own *guided maze* and *office floorplan* environments—in which the robot navigates from a randomly generated start location to a randomly generated point goal in unseen space. We evaluate the following approaches:

Multi-Robot Learning over Subgoals (MR-LSP) Our approach, in which the expected cost of a collective action is computed via Eq. (5). When only one robot is used, this approach corresponds to LSP planning via Eq. (2).

Non-Learned Optimistic Planner We assume that all unseen space is unoccupied and compute paths through each subgoal to reach the goal. The optimistic plan cost is computed for each subgoal and linear sum assignment [23] is used to ensure that robots pursue different subgoals (a topological constraint on their plans similar to [14]) while minimizing net optimistic plan cost to reach the goal. This planning strategy serves as a *non-learned baseline*.

Linear Sum Assignment using LSP (LSA-LSP) This planning strategy also uses linear sum assignment to enforce that robots select different subgoals to pursue, yet the expected cost associated with each subgoal is computed via the single-robot LSP approach Eq. (2). This planning strategy serves as a *learning-informed baseline*. When one robot is used, this approach corresponds to LSP planning via Eq. (2).

Known-Space Planner Planning in the fully-known map (no uncertainty), providing a lower bound on possible cost. Since the shortest path is known, this planner has the same performance for any number of robots.




Fig. 3. **Navigation in maze environment results** The table shows the average cost (in meters) accrued in 100 experiments for each planner. The scatter plot each shows the performance of our approach versus the baselines for two robots. The result shows planners which use learning (MR-LSP and LSP-LSA) outperform the non-learned planner. The performance of learned baseline (LSP-LSA) is similar to MR-LSP (slightly better for two robots) in a relatively simpler environment.

A. Guided Maze Environment Results

We first perform experiments in our “guided” maze environment, in which a green path on the ground connects the goal to another point in the maze. Initially, the robot team is placed at the center of this path so that two green routes extend from its location, only one leading to the goal. The maze is simply connected—i.e., there exists only a single path to the goal—and so a single robot, even one that understands the significance of the green path, must get lucky to reach the goal quickly by choosing the correct path. However, a two robot team can divide-and-conquer, following both green paths simultaneously and reaching the goal quickly with high reliability.

We evaluate our planners in 100 guided maze environments, and show that our MR-LSP and the LSA-LSP learned baseline perform near-optimally in this environment for 2- and 3-robot experiments, outperforming both the 1-robot trials and the non-learned optimistic baseline. Fig. 3 shows the average cost for each planner in this environment. Both MR-LSP and LSA-LSP planners use learning to evaluate goodness of paths and understand the importance of splitting the team to follow both green paths for reaching the goal. In contrast, the non-learned optimistic planner explores many unlikely paths due to a lack of understanding of the green path’s significance. These behaviors are evident in Fig. 4(b) and (c), highlighting the significance of using learning to prioritize the green path and follow multiple promising routes simultaneously.

B. Office Floorplan Environment Results

We conduct experiments in our simulated *office floorplan* environments. These environments are designed to mimic a typical office building, with multiple hallways intersecting and connecting many offices, including some corner rooms. Clutter in each office simulates furniture and obstructs view from the hallway. We evaluate our planners in 100 proce-




Fig. 4. **Navigation under uncertainty in maze** (a) Images taken on board robot used as inputs to learning. (b) Generally learned planner knows to follow green path to reach goal quickly. (c) Occasionally, MR-LSP with single robot (LSP) explores alternative route, whereas with multiple robots MR-LSP guides coordinated exploration.

durally generated office environments—distinct from those seen during training—with varying numbers of robots and report the effectiveness of our MR-LSP planning approach in Fig. 5.

Our MR-LSP approach seems to both understand the utility of following hallways until the goal can be reached and how to effectively allocate exploratory actions to different team members, improving performance over both LSA-LSP and Non-Learned Optimistic planners. The LSA-LSP learned baseline, which uses the same subgoal estimator as the MR-LSP planner, tends to follow hallways and avoid exploring rooms. However, due to a lack of coordination among the team, it only moderately improves over the non-learned baseline. Fig. 6b shows an example of this behavior. The LSA-LSP planner can only coordinate two robots myopically, resulting in poorer behavior compared to the MR-LSP planner, which quickly explores promising routes to the goal.

In the office floorplan environments, 3-robot MR-LSP experiments show *increased* average cost compared to 2-robot MR-LSP experiments in the same space. Fig. 7 shows a scenario where 2-robot MR-LSP outperforms 3-robot MR-LSP planning. Additional experiments increasing the number of samples for PO-UCT from 15k to 100k showed improved performance, indicating that the slight degradation is owed to a limited computational budget and should not be seen as a limitation of the approach itself. In future work, we will explore this relationship in more depth and reimplement our planner in a faster, compiled language.

VIII. RELATED WORKS

Multi-robot navigation in unknown environment is often modelled as a Decentralized POMDP (Dec-POMDP) [1], [2], [3]. Owing to computational challenges in Dec-POMDP planning, dynamic programming [24] and heuristic search

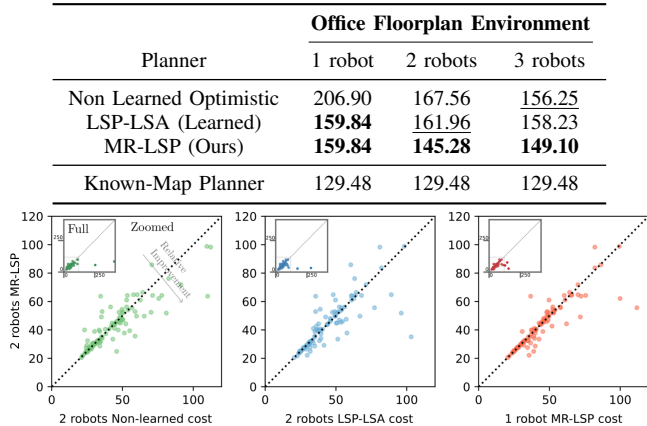


Fig. 5. **Navigation in office floorplan results** The table shows the average cost (in meters) accrued in 100 experiments for each planners. The scatter plot each shows performance of our approach versus the baselines for two robots. In a simulated office floorplan, MR-LSP outperforms both non-learned and learned baselines.

approaches [25] are limited to comparatively short-horizon planning tasks.

Learning for Multi-Robot Planning Learning is frequently used to address the difficulties of planning under uncertainty. Many learning-informed approaches in this domain (e.g., [4], [5]) focus on single-robot planning and do not scale to larger and more complex environments. Multi-agent reinforcement learning has a long history in this domain [26], [27], yet only with recent advances in deep reinforcement learning has it become possible to navigate in somewhat realistic environments [7], [8], [9], [10]. Still, though these approaches work well for small environments, they are often brittle to change [28] and struggle to scale to large-scale environments at the scale of buildings.

Action Abstraction for Multi-Robot Planning To mitigate computational challenges, many approaches in this domain rely on a state or action abstraction to simplify planning. Temporally-extended *macro-actions* are one such action abstraction that help to scale planning under uncertainty [29], [30], [31]. However, these approaches have not proven scalable to building sized environments. Some strategies [32], [33] use variants of Monte Carlo tree search for faster computation in POMDP for long-horizon planning, yet without direct access to a distribution over environments are limited in their ability to reason far into the future. Other approaches to multi-robot planning introduce *topological action abstractions* to simplify planning [11], [12], [13]. Under these approaches, each robot is constrained to leave known space through different *frontiers*, boundaries between known space and unknown space [34], or via paths belonging to different relative homology [15], [14]. Works on navigation using this abstraction generally uses greedy heuristic to select where the robot should navigate next and do not take into account the impact of an action over the longer horizon.

The Learning over Subgoals Planning (LSP) approach [16] uses both a model-based topological abstraction and supervised learning to improve navigation performance, yet is designed only with a single robot planning in mind.

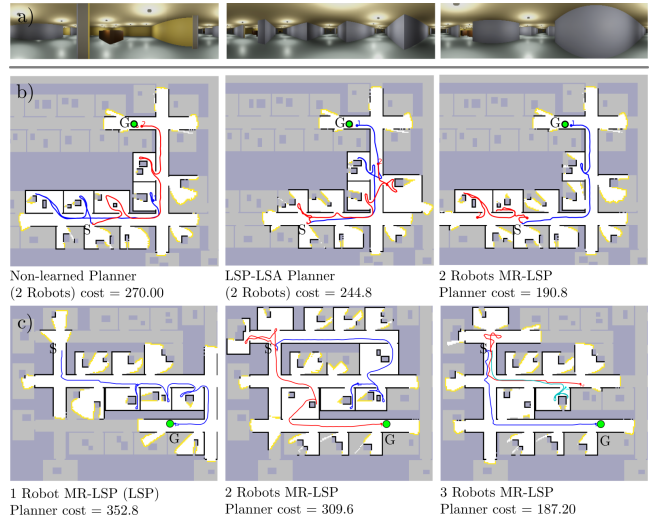


Fig. 6. **Navigation under uncertainty with two robots in the simulated office floorplan** (a) Images from office used as inputs to learning (b) Learned planner improves cost compared to the non-learned planner. MR-LSP performs better over non-learned and learned baselines. (c) MR-LSP generally improves cost as the number of robots increase.




Fig. 7. Increasing the number of PO-UCT samples improves performance for three robots.

IX. CONCLUSION

In this work, we present Multi-Robot Learning over Subgoals (MR-LSP): a novel method for performant and reliable, learning-informed multi-robot navigation through partially-mapped environments. Using learning to estimate the goodness of individual exploratory actions that enter unseen space, our multi-robot team is able to envision the expected long-horizon impact of each robot’s actions and can thus coordinate behavior far into the future to quickly reach the unseen goal, outperforming both learned and non-learned strategies. However, the number of multi-robot actions grows rapidly with the number of robots, imposing a practical limitation on team size. In future work, we would like to extend our model to support more complex and multi-stage tasks, effectively extending the model of Bradley et al. [22] to the multi-robot planning domain.

ACKNOWLEDGEMENT

We thank Kevin Doherty, Chris Bradley, Jana Kosecka, Erion Plaku, and Cameron Nowzari for their thoughtful feedback on this work. We acknowledge funding support from George Mason University.

REFERENCES

- [1] C. Amato, G. Chowdhary, A. Geramifard, N. K. Üre, and M. J. Kochenderfer, "Decentralized control of partially observable markov decision processes," in *52nd IEEE Conference on Decision and Control*, 2013.
- [2] S. Omidshafiei, A.-a. Agha-mohammadi, C. Amato, and J. P. How, "Decentralized control of partially observable markov decision processes using belief space macro-actions," 2015.
- [3] D. S. Bernstein, E. A. Hansen, and S. Zilberstein, "Bounded policy iteration for decentralized pomdps," in *Proceedings of the 19th International Joint Conference on Artificial Intelligence*, 2005.
- [4] S. Gupta, V. Tolani, J. Davidson, S. Levine, R. Sukthankar, and J. Malik, "Cognitive mapping and planning for visual navigation," 2017.
- [5] C. Richter, J. Ware, and N. Roy, "High-speed autonomous navigation of unknown environments using learned probabilities of collision," in *2014 IEEE International Conference on Robotics and Automation (ICRA)*, 2014.
- [6] C. Amato, "Decision-making under uncertainty in multi-agent and multi-robot systems: Planning and learning," in *Proceedings of the 27th International Joint Conference on Artificial Intelligence*, 2018.
- [7] Y. Jin, Y. Zhang, J. Yuan, and X. Zhang, "Efficient multi-agent cooperative navigation in unknown environments with interlaced deep reinforcement learning," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2019.
- [8] J. K. Gupta, M. Egorov, and M. J. Kochenderfer, "Cooperative multi-agent control using deep reinforcement learning," in *AAMAS Workshops*, 2017.
- [9] J. N. Foerster, Y. M. Assael, N. de Freitas, and S. Whiteson, "Learning to communicate with deep multi-agent reinforcement learning," 2016.
- [10] J. Kober, J. A. Bagnell, and J. Peters, "Reinforcement learning in robotics: A survey," *The International Journal of Robotics Research*, 2013.
- [11] B. Yamauchi, "Frontier-based exploration using multiple robots," in *Proceedings of the second international conference on Autonomous agents*, 1998.
- [12] P. Khandelwal and P. Stone, "Multi-robot human guidance using topological graphs," in *AAAI Spring 2014 Symposium on Qualitative Representations for Robots (AAAI-SSS)*, 2014.
- [13] Y. Wang, A. Liang, and H. Guan, "Frontier-based multi-robot map exploration using particle swarm optimization," in *2011 IEEE Symposium on Swarm Intelligence*, 2011.
- [14] S. Bhattacharya, M. Likhachev, and V. Kumar, "Topological constraints in search-based robot path planning," *Autonomous Robots*, 2012.
- [15] S. Bhattacharya, R. Ghrist, and V. Kumar, "Persistent homology for path planning in uncertain environments," *IEEE Transactions on Robotics*, 2015.
- [16] G. J. Stein, C. Bradley, and N. Roy, "Learning over subgoals for efficient navigation of structured, unknown environments," in *Conference on Robot Learning*, 2018.
- [17] M. Merlin, N. Parikh, E. Rosen, and G. Konidaris, "Locally observable markov decision processes," in *ICRA 2020 Workshop on Perception, Action, Learning*, 2020.
- [18] L. P. Kaelbling, M. L. Littman, and A. R. Cassandra, "Planning and acting in partially observable stochastic domains," *Artificial intelligence*, 1998.
- [19] M. L. Littman, A. R. Cassandra, and L. P. Kaelbling, "Learning policies for partially observable environments: Scaling up," in *Machine Learning Proceedings*, 1995.
- [20] J. Pineau and S. Thrun, "An integrated approach to hierarchy and abstraction for pomdps," Carnegie Mellon University, Tech. Rep., 2002.
- [21] D. Silver and J. Veness, "Monte-carlo planning in large pomdps," *Advances in neural information processing systems*, 2010.
- [22] C. Bradley, A. Pacheck, G. J. Stein, S. Castro, H. Kress-Gazit, and N. Roy, "Learning and planning for temporally extended tasks in unknown environments," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2021.
- [23] H. W. Kuhn, "The hungarian method for the assignment problem," in *50 Years of Integer Programming*, 2010.
- [24] E. A. Hansen, D. S. Bernstein, and S. Zilberstein, "Dynamic programming for partially observable stochastic games," in *AAAI*, 2004.
- [25] D. Szer, F. Charpillet, and S. Zilberstein, "MAA*: A heuristic search algorithm for solving decentralized pomdps," *arXiv preprint arXiv:1207.1359*, 2012.
- [26] M. Lauer and M. Riedmiller, "An algorithm for distributed reinforcement learning in cooperative multi-agent systems," in *In Proceedings of the Seventeenth International Conference on Machine Learning*, 2000.
- [27] L. Matignon, G. J. Laurent, and N. Le Fort-Piat, "Independent reinforcement learners in cooperative markov games: a survey regarding coordination problems," *The Knowledge Engineering Review*, 2012.
- [28] P. Henderson, R. Islam, P. Bachman, J. Pineau, D. Precup, and D. Meger, "Deep reinforcement learning that matters," in *AAAI Conference on Artificial Intelligence*, 2018.
- [29] C. Amato, G. D. Konidaris, and L. P. Kaelbling, "Planning with macro-actions in decentralized pomdps," in *Proceedings of the 2014 International Conference on Autonomous Agents and Multi-Agent Systems*, 2014.
- [30] C. Amato, G. Konidaris, G. Cruz, C. A. Maynor, J. P. How, and L. P. Kaelbling, "Planning for decentralized control of multiple robots under uncertainty," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2015.
- [31] C. Amato, G. Konidaris, A. Anders, G. Cruz, J. P. How, and L. P. Kaelbling, "Policy search for multi-robot coordination under uncertainty," *The International Journal of Robotics Research*, 2016.
- [32] M. Hoerger, H. Kurniawati, and A. Elfes, "Multilevel monte-carlo for solving pomdps online," in *The International Symposium of Robotics Research*, 2019.
- [33] C. Amato and F. Oliehoek, "Scalable planning and learning for multi-agent pomdps," in *Proceedings of the AAAI Conference on Artificial Intelligence*, 2015.
- [34] B. Yamauchi, "A frontier-based approach for autonomous exploration," *Proceedings 1997 IEEE International Symposium on Computational Intelligence in Robotics and Automation CIRA'97. 'Towards New Computational Principles for Robotics and Automation'*, 1997.