Distributed Inference-Based Multi-Robot Exploration

Andrew J. Smith, and Geoffrey A. Hollinger

Abstract— **This work proposes a technique for distributed multi-robot exploration that leverages novel map inference to increase the team's cumulative exploration efficiency. The multi-robot team uses a distributed algorithm to coordinate the exploration using both the inferred and observed portions of the map. Individual robots select exploration poses by accounting for expected information gain and travel costs. Robots resolve conflicts between exploration goals with local auctions of expected travel costs. The benefits of inference-informed exploration are demonstrated in both simulated explorations and hardware trials. The proposed technique is compared against frontier-based and information-based exploration approaches. These comparisons evaluate the performance of the three exploration methods with decaying communication and a varied number of agents. Including inference in the coordination leads to a 13.15% reduction in the cumulative exploration path length in the trials conducted.**

I. INTRODUCTION

This paper presents a method of coordinated exploration for decentralized multi-robot teams that leverages distributed map-inference techniques. This coordination method is applicable to multi-robot teams exploring environments for civilian, search and rescue, military, and research applications. Prior methods of exploration have been largely frontier based [1] and have either greedily selected frontiers or used a market or auction to distribute frontiers between agents [2]. Recent work has used information-theorietic based approaches to explore and optimize the path of exploring robots [3]. This paper builds upon these ideas by combining potential information gain, distributed markets, and novel methods of map inferences to create a robust and efficient distributed exploration technique.

The developed inference techniques uses observed portions of the map to, first, estimate the outer boundaries of the explorable area then, second, to infer the structure inside the boundary. The outer boundary is inferred using a heuristic method to extend the observed map boundaries. The internal structure is inferred by comparing observed map structure against a library of map structure. The agents then coordinate the exploration by conducting local markets of potential observation poses. Results are presented to demonstrate the value and accuracy of the inference as well as the benefit inference provides to multi-robot exploration. Experiments were conducted with teams of robots exploring an indoor space with limited ability communicate. Our strategy reduces the time required to explore an environment by 13.15% and 12.34% for varied communication strength and numbers of agents, respectively, against a market frontier based approach.

A. J. Smith and G. A. Hollinger are with the School of MIME at Oregon State University. {smithan7, geoffrey.hollinger}@oregonstate.edu

II. METHODS

The inference developed in this work is comprised of two components: perimeter inference used to estimate the outer boundary of the area being explored, and structural inference used to infer internal map structure. The primary objective of the perimeter inference is to estimate the outermost boundary of the explorable space. The process is a heuristic method that begins by identifying the convex hull of the observed map and extending the observed map structure on the hull into unobserved portions of the map. Then, structural inference uses the observed portions of the map and the inferred perimeter to infer unobserved internal map structure and potential breaches in the inferred perimeter. The inference is performed by simulating a sparse 360^o laser range scan at a sampled pose in the environment. If the sampled scan encounters only observed portions of the map then the scan is added to the structural library of priors. If the simulated scans encounter an unobserved portion of the environment, then the simulated scan is compared, using maximum likelihood, against the library of prior scans. If the likelihood of a match surpasses a threshold then the matching library entry is merged into the agent's reward map. To increase the strength of the structural inference beyond the observations of the current exploration, $86,246$ simulated laser scans from 124 different environments are seeded into the structural inference library.

By inferring the unobserved free space, the exploring robots have additional information to plan how to explore the

Fig. 1. Example of a partially explored map with inference. Green indicates areas with low reward, corresponding to places that have been observed, and red indicates area of high reward, corresponding with unobserved inferred areas.

remaining space as shown in Figure 1. Exploring robots use the inference in their individual planners by searching over both the inferred and observed space for potential exploration poses. Then the robots will use the inference to calculate the reward of exploring those poses. The reward of each pose is the predicted information gain of observing from that pose and each cost is a weighted travel cost to reach each pose. The value of each pose is the reward minus the cost of each pose. Robots explore by selecting the pose with the maximum value. This approach prioritizes exploration by providing reward for observing cells that are inferred to be free or cells that were inferred to leave the current boundaries of the inferred perimeter; i.e. a doorway to a new wing of a building.

After each agent has selected exploration goal poses from their individual markets, the next step is to coordinate the exploration with local agents. Each agent broadcasts their goal pose and expected travel cost to all agents in communication range. All local agents sharing the same goal pose form a single bid auction. When each robot broadcasts their goal pose they are allowing other robots the opportunity to underbid them or to underbid the other robot. When this occurs the robot who was underbid resolves the conflict by selecting another goal pose. This auction-based approach allows each robot to explore their optimal exploration pose if is not in conflict with other robots. However, as the number of robots participating in the exploration increases the chances of conflicting goal poses increase. Two poses are considered conflicting when they have overlapping observations of the same area. If this occurs, the reward for observing the overlapped area by the robot arriving second is decreased. The purpose of devaluing, instead of removing, conflicted reward from the reward map is that in the absence of another place to explore, robot a_0 will still move in the direction of the conflicted goal. Although it will be explored when the agent arrives, it is possible that it may branch into new areas that have not been explored, moving the robot closer to unexplored areas.

The result of the inference based coordination is a goal pose for each robot to explore. Each robot uses the observed map structure to make their own inferences about the unobserved portions of the map. Then, robots sample poses from the observed and inferred portions of the map to create an internal market of poses that fully observe the inferred free space. Robots then select the exploration pose with the highest value from their internal market and set it as their exploration goal. Robots settle local conflicts by broadcasting their selected goal pose and travel cost. This approach provides a reasonable method of incorporating the benefits of map inferences into a distributed exploration by building upon many of the strengths of market-based coordination and information-theoretic approaches to exploration. This allows for a fully distributed approach to coordination and informed goal selection.

III. EXPERIMENTS AND RESULTS

To evaluate the performance of the inference, two series of simulated explorations were conducted. The first simulation consisted of a single robot exploring a previously unexplored building using the described market-based frontier exploration method. It is assumed that the agent has a 360^o laser scanner and the ability to perform SLAM. While the robot explored the building it recorded both its observations and the output of the inference. To evaluate the performance of the inference, the commonly used classification metrics of precision and recall are used. Precision describes how accurate the inferred information is while recall describes how complete the inferred information is. These two metrics combine to allow for an accurate description of how much additional information is gained by the inference (recall) and how useful that information is (precision). The naive (without inference) exploration of the environment is used as a baseline for comparison. Increases in the inferred recall indicates that the inference provides information describing a larger portion of the map than is currently observed while the inference precisions indicates the accuracy of the inferred information.

For the simulations, test environments were randomly selected from the set of maps and then starting locations were randomly selected from unoccupied cells in the chosen map. The robot then explored the map using a greedy frontier based approach on the naive map. At each time step, the exploring robot records the current map information gathered through direct observations and provided by the map inference. A total of 300 trials were completed. To account for the wide range of exploration times required to explore maps of different sizes, the results have been normalized with respect to mission time duration. Results for recall and precision are presented below in Figures 2 and 3, respectively.

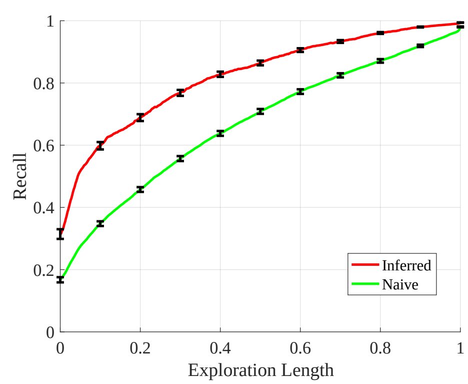

Fig. 2. The mean recall of the naive and inference costmaps for the exploration trials completed. The error bar indicates the standard error of the mean. Recall provides a measure of how much of the explored space is observed or inferred.

As can be seen in Figure 2, a significant amount of information is gained by the perimeter and structural inference, especially early in the exploration. The peak mean

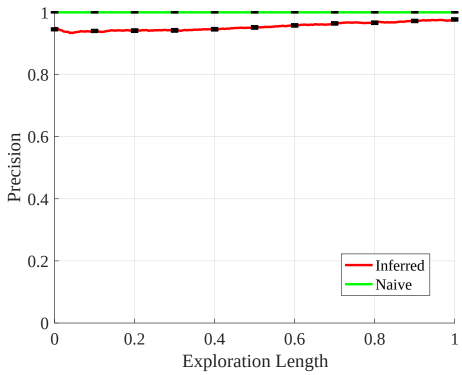

Fig. 3. The mean precision of naive and inference costmaps for the exploration trials. The error bar indicates the standard error of the mean. Precision provides a measure of how correct the observations or inference are. Notice that the naive observations have perfect precision; this is because no measurement errors are provided and all observations are assumed to be correct.

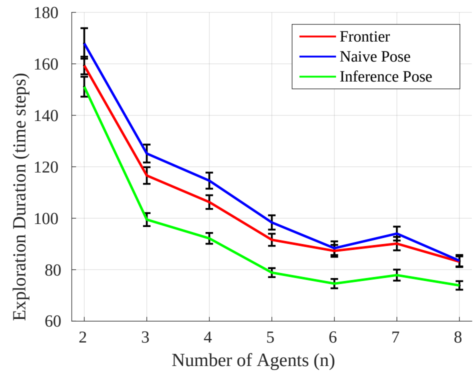

Fig. 4. Exploration efficiency for varying number of agents in the simulated trials. As can be seen across the range of agents tested the inference informed pose selection method results in more efficient exploration of the environment. The error bar indicates the standard error of the mean for the 50 iterations of each of the 7 testing scenarios.

gain in recall occurs at 5.6% of the exploration duration with an increase in recall of 108.84% by the inference over direct observation. The mean gain in recall throughout the exploration is 34.47%. This gain in recall provides the agent with an additional 34.47% information with which to plan the exploration. By looking at Figure 3 it can be seen that the information provided by the inference is also accurate and generally infers the map structure correctly, with a mean precision across the exploration of 0.9539 for the inference.

To evaluate inference's contribution to exploration efficiency multiple tests were conducted with different exploration methods, levels of communication, and number of agents. For the purpose of these tests, exploration efficiency is measured by the number of time steps required to complete an exploration. The goal selection methods tested were the proposed inference-informed pose market, a naive (without inference) pose market, and a frontier market. All three methods used an internal market for exploration goal selection and then broadcast their bids in the open local auctions. The primary difference between the three methods was the method of sampling goal poses and calculating each goal poses reward. The communication ability of each agent was varied from unrestricted global communication between all agents to range restricted line-of-sight communication. Tests were conducted to identify the performance of the system with two to eight agents. This range was chosen as the exploration gain of additional agents appeared to be negligible. For each test scenario, either number of agents or communication strength, a random starting location was chosen on the map for each test iteration for a total of 50 iterations.

The developed exploration method that leverages map inference to inform pose selection outperforms the naive pose and frontier based explorers by ($\mu = 17.70\%$, $\sigma = 5.37$) and ($\mu = 12.34\%$, $\sigma = 2.97$), respectively, across the range of team sizes tested as shown in Figure 4. This shows that the inference informed exploring team is able to identify better exploration goals for the members of the team resulting in reduced time to fully explore the same space. As the number of team-members is increased the inference informed exploring team continues to outperform the two baselines suggesting the inference leads to improved coordination.

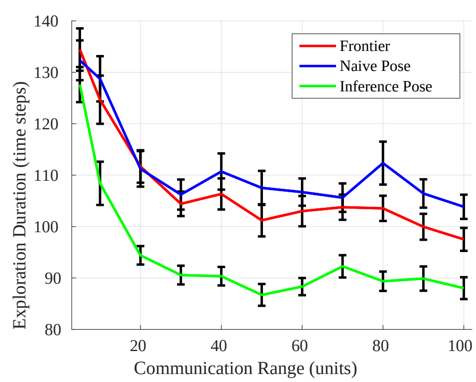

Fig. 5. Exploration efficiency for varying communication ability in the simulated trials with three robots. As can be seen across the range of communication tested the inference informed pose selection method results in more efficient exploration of the environment. The error bar indicates the standard error of the mean for the 50 iterations of each of the 11 testing scenarios.

The developed exploration method that leverages map inference to inform pose selection outperforms the naive pose and frontier based explorers by ($\mu = 16.89\%$, $\sigma = 4.92$) and ($\mu = 13.15\%$, $\sigma = 3.26$), respectively, across the communication ranges tested as shown in Figure 5. This shows that as the communication between agents degrades the developed pose base inference degrades gracefully. This is because, similar to the two baselines, the inference informed robot uses the auction based approach to resolve local conflicts. When two robots come into contact they resolve exploration conflicts and retain the other robots goal locations. So, even after communication between them has been severed they continue to account for the other robot's broadcast exploration goal. This behavior results in the robots separating from one another and effectively exploring the space even with unreliable communication. The inference assists the exploration in two ways. First, as discussed before the inference improves the estimate of pose rewards. Second, the inference allows for the marketed poses to be placed in additional locations. The inference informed pose market allows for poses to be placed in both the observed and inferred free portions of the map.

Hardware trials were conducted on Pioneer P3-DX robots to verify the functionality of the combined inference and coordination on hardware. For this trial global communication was assumed and a fixed team size of three agents was conducted with each of the coordination methods in one environment. During the hardware trials it was demonstrated that the pioneers were capable of using the inference based coordination to explore an indoor space. Figure 6 shows the map constructed by the exploration team and their paths through the environment during the mapping.

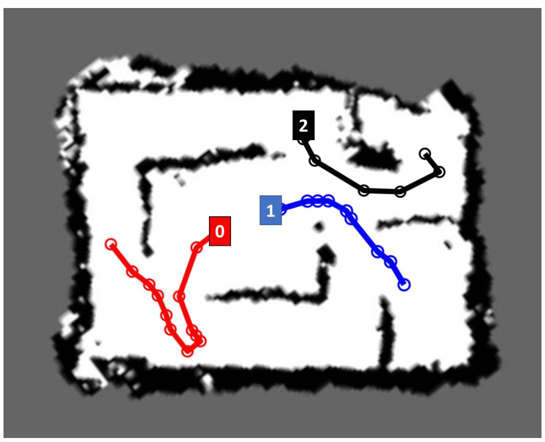

Fig. 6. The resulting merged map of the test environment and the paths of the three pioneers during the exploration. Numbered rectangles indicate starting positions of each pioneer.

IV. CONCLUSIONS AND DISCUSSION

This work presents a novel map inference based coordination method for distributed multi-robot exploration. Exploring robots sample potential poses from the observed and inferred portions of the map and use an internal market to select their goal pose for exploration. To resolve conflicting goal poses between agents, each agent broadcasts their goal pose and travel cost in an open auction. Robots use the developed map inference techniques to sample and evaluate potential goal poses. The map inference was tested across 126 different maps and provided an average gain in information of 34.47% with a mean precision of 95.1% in the simulated trials. The inference based exploration method increased the team's cumulative exploration efficiency against a naive (without inference) information pose and frontier based exploring teams in the simulated trials. The distributed coordination method was demonstrated to be robust to varying numbers of agents, outperforming the naive and frontier based exploration methods by 13.15% and 16.89%, and communication ability, outperforming the naive and frontier based exploration methods by 12.34% and 17.70%. The developed system was then demonstrated through hardware trials on a team of robots.

REFERENCES

[1] B. Yamauchi, "A frontier-based approach for autonomous exploration," in *Computational Intelligence in Robotics and Automation*. IEEE, 1997, pp. 146–151.

[2] R. Zlot, A. Stentz, M. B. Dias, and S. Thayer, "Multi-robot exploration controlled by a market economy," in *Robotics and Automation, 2002. Proc. IEEE International Conference on*, vol. 3, 2002, pp. 3016–3023.

[3] B. Charrow, G. Kahn, S. Patil, and S. Liu, "Information-theoretic planning with trajectory optimization for dense 3d mapping," *Proc. Robotics Science and Systems Conf. 2015*.