

Multi-robot task allocation for collaborative adaptive informative sampling in structured environments

Nicholas Fung^{1†}, John Rogers^{1†}, Carlos Nieto², Stephanie Kemna³, Gaurav Sukhatme³, Henrik I. Christensen²

I. INTRODUCTION

We look to improve the efficiency of teams of autonomous robots for information gathering approaches. These approaches are useful for environmental modeling, persistent monitoring, or mapping chemical signals in the air. An effective way of creating spatial models is by using Gaussian Process (GP) regression. Information-theoretic metrics, such as entropy or mutual information, can be used on top of the GP model’s predictions to decide where to place sensors or robots. In such case we speak of informative sampling. When the robot is creating such a model while sampling data this is called on-line or adaptive informative sampling.

In this work we present our late breaking results extending our prior work [1] on multi-robot adaptive information gathering in structured indoor environments. This previous work implemented an information gathering rate adaptive algorithm to task robots with sensing an environment and building a GP model of RSSI signal readings. Multiple robots were coordinated by partitioning the environment into a number of segments equal to the number of robots and assigning each robot to its own region. In this paper we present a method for oversegmenting the sampling space into more regions than robots. A task assignment algorithm allows a central controller to dynamically task robots to different partition regions.

Related work

Informative sampling was pioneered in robotics by Singh et al. and Low et al. [2], [3]. Atanasov et al. [4] developed an efficient decentralized and optimal solution for a multi-sensor active information acquisition problem that is non-myopic and memory-efficient. They show they can reduce the complexity in the planning horizon and in the number of sensors. Their work was applied to multi-robot active SLAM. Kemna et al. [5] describe an on-line decentralized multi-robot coordination approach, where Voronoi partitions are used to coordinate robots’ actions. In contrast to most former methods, their approach recalculates the Voronoi partitions during the sampling while the model of the environment is being built. Furthermore the GP’s estimates are used to

balance informativeness across vehicle partitions. This work assumes unstructured environments.

II. APPROACH

A. Signal Modeling

This work addresses the use of a task-allocation algorithm to coordinate work between multiple robots to cover a spatially segmented, structured indoor environment. We compare a new method of oversegmenting the space and using task allocation, with our former method that segments the space into as many regions as there are robots. Previous work initially assigns specific robots to their own region and does not revisit this decision. The goal for the robot team is to collect received signal strength indicator (RSSI) measurements to model the signal strength over a known indoor environment.

For this experiment, a ground truth model was trained using real world data collected in a controlled test environment using robotic sensor platforms that were manually driven through a building environment that contained three static radio signal generators, each with unique identification. Each RSSI measurement by any robot is used to train the GP model using libGP [6]. This is the ground truth model is only used to generate sensor readings in a simulation environment for the experiment.

During experimentation, the simulated robots build a GP model based on simulated sensor readings. When considering potential points to task a robot to measure, the central controller uses the predictive variance according to its current GP model. The controller computes the differential entropy $h(\mathbf{x})$ at each of these points according to Equation 1.

$$h(\mathbf{x}) = \ln(\sigma\sqrt{2\pi\epsilon}) \quad (1)$$

Robots are tasked to points of interest according to the predicted differential entropy and the predicted travel cost. We refer to the *rate* of information gain as the differential entropy between two points divided by the travel cost to that location. We dynamically maximize the *rate* of information gain, which we refer to as *information rate adaptive* or *dl-adaptive* [1]. As the robots move and sense the environment, the GP model is updated and new predictive variances are used to influence future navigation goals.

B. Multi-robot coordination and task allocation

We augment the *dl-adaptive* approach with a partitioning routine which analyzes an all-pairs distance matrix for

[†] equal contributions

²University of California San Diego cnietogr@ucsd.edu

³University of Southern California kemna@usc.edu

¹US Army Research Laboratory (ARL) nicholas.c.fung.civ, john.g.rogers59.civ@mail.mil

structural separations corresponding to rooms via a recursive normalized graph cuts procedure [7]. This partitioning divides the work of exploring the complete space irrespective of the potential information gain. The baseline partitioning approach, proposed in [1] chooses as many partitions as there are robots. Each robot is assigned one partition to sample for the duration of the test run. Our new approach implements oversegmentation and a task allocation algorithm. We hypothesize that oversegmenting the space will result in more efficient tasking of the robots.

For the multi-robot coordination problem, we partition the sampling space and assign a robot to each partition. An optimal partitioning would divide the region into separate partitions such that each partition takes an equal and minimal amount of time to explore. Within an individual partition, the assigned robot evaluates points of interest based on the entropy and travel cost.

There are several cases related to the number of partitions in relation to the number of robots. If the number of robots is equal to the number of partitions, i.e. *equal partitioning*, each robot is assigned its closest partition to explore. If there are more partitions than robots, i.e. *over partitioning*, a task assignment procedure is used to instruct each robot to sample an individual partition. This assignment can be revised over time as information is incorporated into the model. As the information gathering operation progresses, the entropy in the sampled partitions will be reduced, and robots will be reallocated to other partitions to improve the model.

The task assignment routine begins by adding up the entropy of each point of interest within a partition. These partitions are then sorted on the entropy measure and the highest entropy partitions are each assigned to a robot. In the case of *oversegmenting*, some partitions will not initially be assigned to a robot for sampling. To ensure that robots continue to collect information, once the total entropy of their assigned partition falls to a ratio of 70% of the original value, the partition assignment procedure is repeated for that robot. At this point, the robot will be assigned to the region with the largest cumulative entropy in the model that does not currently have a robot assigned to it. It is possible for this new assignment to the same partition if it has the largest entropy sum. Figure 1 shows partitions created for our test environments, showing fixed cuts for two to four robots (left three columns of images) and oversegmentation using the normalized cut (rightmost column of images).

The hysteresis parameter of reducing to 70% of the original cumulative entropy is used to encourage robots to work locally on improving their model before moving on to other regions. If this parameter were set so that the robot would be reassigned after any reduction in entropy, then the resulting behavior would mimic an alternative where no partitions were formed. We have shown in previous work [1] under review that using partitions helps to improve models more quickly than the ad-hoc alternative without partitions.

III. EXPERIMENTATION

We performed the following experiments:




Fig. 1. Partitions selected by normalized graph cuts on inverse path cost affinity matrix. The leftmost three columns show fixed cuts made for between two and four robots. The rightmost column shows the partitions selected when segments up to the normalized cut cost of 0.7 are made. Open cells are shaded according to partition assignment in unique colors.

- The floor plan is segmented with an equal number of robots and partitions.
- The floor plan is segmented with at least as many partitions as robots, and then oversegmented until the normalized cut cost metric exceeds 0.7.

Task assignment of the robots is carried out by a centralized controller according to the GP model built by the robots. We assume full communication between all robots and this controller.

The experiments are carried out on three different floor plans labeled Bank, Hospital Floor 1, and Hospital Floor 2. The environments were seeded with 3 radio signal generators labeled 2, 3, and 4 (radio 1 is located on the exploration vehicle) placed within the environment under two different configurations: A and B. Including the “Max” radio configuration, in which all the radio signals are active, this results in 8 different radio signal configurations per floor plan.

IV. RESULTS

We evaluate the simulation results using the root-mean squared error (RMSE) between the ground truth (GP model of real world data) and the GP model created through sensor measurements by the simulated robot. We measure the area under the RMSE curve to evaluate how quickly the RMSE decreases as the robots build their GP models. A small area under curve shows that the error fell quickly, demonstrating an improvement on performance. The results show that the area under the curve drops with addition of more robots, and it also drops significantly when oversegmenting and applying the task assignment algorithm.

Figure 2 shows an overview of the final area under the RMSE curve for all experiments, showing the results for two to four robots for every configuration. Red bars show the results for runs with an equal number of segments and robots,




Fig. 2. Montage of the area under the RMSE curve results, across all configurations. Red bars are for equal number of partitions and robots, and blue bars are for the oversegmented partitioning and task assignment. In each subfigure, the x axis shows results for 2, 3, and 4 robots.

and blue bars show the results for the oversegmentation with task assignment. The majority of the graphs show an improvement in RMSE reduction when using more robots, for either algorithm. In addition, in all but very few cases the oversegmentation with task allocation method outperforms the method that uses an equal number of segments to robots. These effects are more pronounced with larger and more complex environments: Hospital 2F is larger than Hospital 1F, both of which are significantly larger than Bank.

For a clearer visual analysis, Figure 3 shows a more detailed graph with the area under the curve for one specific scenario: Hospital Floor 2 configuration 2A. The gains from utilizing oversegmentation are very pronounced for this scenario.

Our results show that modeling performance improves when we divide up work using oversegmentation and task allocation. This approach keeps robots working locally until that portion of the model is improved, at which point the

robots can be reallocated to continue their sampling tasks. This improvement increases with the number of robots and the size and complexity of the environment, showing the biggest improvements for the simulations with four robots.

V. CONCLUSIONS

In this work, we describe a task allocation algorithm for adaptive information sampling in structured environments. This task allocation algorithm operates on an “oversegmentation” of spatial regions by making use of the normalized graph cut algorithm. By configuring robots to collect information in nearby regions before moving on to other areas, via a hysteresis parameter, multi-robot teams can be effectively coordinated to collect information efficiently.

REFERENCES

- [1] N. Fung, J. Rogers, S. Kemna, and C. Nieto, “Multi-robot coordination for adaptive informative sampling in structured environments,” in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)* (under review). IEEE, 2018.
- [2] A. Singh, A. Krause, C. Guestrin, W. Kaiser, and M. Batalin, “Efficient Planning of Informative Paths for Multiple Robots,” in *International Joint Conference on Artificial Intelligence (IJCAI)*, 2007.
- [3] K. H. Low, J. M. Dolan, and P. Khosla, “Adaptive Multi-Robot Wide-Area Exploration and Mapping,” in *Proceedings of the 7th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 2008, pp. 23–30.
- [4] N. Atanasov, J. Le Ny, K. Daniilidis, and G. J. Pappas, “Decentralized active information acquisition: Theory and application to multi-robot SLAM,” *ICRA*, pp. 4775–4782, 2015.
- [5] S. Kemna, J. G. Rogers III, C. Nieto-Granda, S. Young, and G. S. Sukhatme, “Multi-Robot Coordination through Dynamic Voronoi Partitioning for Informative Adaptive Sampling in Communication-Constrained Environments,” in *International Conference on Robotics and Automation (ICRA)*, May 2017.
- [6] M. Blum, “libgp: a C++ Library for Gaussian Process Regression,” GitHub, <https://github.com/mblum/libgp>, retrieved August, 2017.
- [7] J. Shi and J. Malik, “Normalized cuts and image segmentation,” *IEEE Transactions on pattern analysis and machine intelligence*, vol. 22, no. 8, pp. 888–905, 2000.




Fig. 3. Comparison of the area under the curve between equal number of partitions as robots vs “oversegmented” partitioning and task assignment on Hospital Floor 2, configuration 2A, for 2 to 4 robots.