

MISSION-RELEVANT COLLABORATIVE OBSERVATION AND LOCALIZATION

Ashley W. Stroupe

Carnegie Mellon University, Pittsburgh, Pennsylvania, USA

ashley@ri.cmu.edu

Tucker Balch

Georgia Institute of Technology, Atlanta, Georgia, USA

tucker@cc.gatech.edu

Abstract We propose an approach for the integration of collaborative observation and localization within a separate multi-robot mission. The proposed approach will allow individual robots within a team to evaluate and choose a task based on the needs of the mission and the needs, capabilities, and task choices of teammates. Tasks will include mission-related tasks as well as task for self-localization, teammate-localization, and collaborative mapping. To date, work has been done in collaborative observation and mapping and.

1. INTRODUCTION

Near-term goals for multi-robot systems include applications in complex environments with large robot teams, such as planetary science, construction, surveillance/reconnaissance, and search and rescue. Successful completion of such complex tasks with limited resources and time will require taking full advantage of multi-robot systems' potential to improve efficiency, capability, and reliability, as well as addressing complex dynamic environments.

Suppose a robot team could create and maintain a consistent and accurate understanding of the environment while performing a high-level mission. We hypothesize that such a capability can substantially improve the system's effectiveness. Accuracy is required to perform individual tasks precisely and efficiently; consistency is required for effective cooperation among agents in

tasks that cannot be accomplished individually. To achieve a consistent and accurate understanding of the environment, robots must make and combine multiple observations of objects and teammates and share results.

Integration of this collaborative mapping and localization with a complex mission increases the complexity of the task selection process. More tasks, capabilities and joint effects of teammates, and priorities of different tasks (which change with time and state) must be considered.

The proposed approach will deal with task selection, mapping, and localization, each in a multi-robot context. To date, work has been done in cooperative observation of static and dynamic targets and localization. Future work will include the design of explicitly collaborative localization and observation and design and implementation of a behavioral framework for incorporating localization and mapping into a multi-robot mission.

2. RELATED WORK

Localization maps sensor data to robot pose using a representation of the environment. Approaches range from simple geometry to probabilistic modeling of pose estimates. Triangulation uses range and bearing to landmarks (sometimes with previous position and motion model) to find a pose most consistent with sensor information. More complex geometric approaches process data to create “images.” Closest match and interpolation with stored templates provides pose. Most current approaches represent pose by a probability distribution or density. Kalman-Bucy filters and Bayes’ Rule filters use Gaussian distributions, allowing simple mathematics for updates. Markov Localization and Monte Carlo Localization approaches allow any type of distribution at the cost of complexity and/or resolution. Many localization approaches have been applied to small multi-robot systems. Using triangulation, teammates take turns exploring the environment and serving as stationary landmarks (leap-frog). In probabilistic approaches, leap-frog is used, probabilistic estimates of observed teammate poses are shared and used in updates, or individual maps are registered for temporary cooperation.

Mapping provides an environment representation for localization and task selection. This representation can be location of objects or landmarks or terrain elevations. Approaches are similar to those for localization, and mapping and localization are often done simultaneously. Simple approaches relate relative position of objects to world position using robot pose, typically in a discrete world. More complex approaches can provide a probability of grid cells being occupied or traversable, templates, or a probabilistic position estimate of objects using Kalman-Bucy filters, Bayesian updates, or Markov/Monte Carlo. Such probabilistic approaches have also been used to track dynamic objects. Multi-robot approaches typically divide the problem into single-robot problems by dividing area or using leap-frog approaches. Multiple robots can also jointly update a single grid cell occupancy map.

Some work has been done in determining how to deploy a team to maximize coverage without consideration of overlapping observations.

Members of a team must allocate mission tasks among agents. The simplest approach, applicable when a single task is required, is division of the task space. Some approaches are dynamic, changing results as new information is discovered and faults occur. Task selection is typically based on the value of different tasks to a robot agent. The evaluation of value can be done in a reactive or behavior-based fashion, using predefined functions, learned functions, or traditional planning approaches with predefined models.

The primary limitations of these approaches are the lack of integration and the inability to fully utilize the multi-robot system. Implementations of task selection algorithms focus on few choices, typically a single high-level goal, and do not incorporate localization and map uncertainty in the decision process. Systems exhibit no explicit movement for obtaining good team observations, and knowledge of teammate decision processes (direct and indirect) are not fully addressed.

3. APPROACH

3.1. Task Selection

Robot tasks are specified by the mission directly (such as scientific analysis) or indirectly (such as need for localization and a map). The selected task is that which is determined most valuable. Task value depends on current priorities of different tasks, current task status, current state and pose estimate (including uncertainty), and resource availability. For example, mapping will have higher value when map certainty is low, when other mission-related tasks have low priority or value (due to low map quality), or when mapping sensors are more accurate. Value may also include an estimate of cost to complete the task. Unlike previous approaches, value may also be affected by teammate availability to participate in the task and the current needs of teammates. For example, localizing a teammate with a high-priority task will have higher value than localizing one with a low-priority task, and some mission tasks may have low value when other necessary resources are unavailable.

Given that task selection partially depends on the choices of others, agreement must be reached among participating agents. This may lead to an iterative process of evaluating value. For efficiently selecting a task, the number of agents considered in evaluation and included in agreement may be reduced by considering subsets of agents defined by proximity, availability, or team assignment.

Values for behaviors can be defined by functions; the highest-valued function at any time selects the most appropriate behavior. Functions may be partially generic and partially specific to task and robot, but may also change as aspects of the mission become more or less important and the state

changes. An example of this type of evaluation is shown. Mission tasks requiring accuracy become less valuable as pose uncertainty increases; at some level of uncertainty, explicitly improving pose estimate becomes the highest valued task and is selected.

3.2. Collaborative Localization

Collaborative localization includes both cooperative localization (using observations of and by teammates to update pose) and collaborative localization (explicitly moving to aid a lost teammate). When teammates encounter and observe each other, each makes a relative and absolute position estimate of the other. This absolute estimate can be used by the other to update. Should a teammate call for help, one or more robots must answer the call and attempt to localize the lost robot. The value of localizing the teammate will depend on the value of that robot's current task, the estimated value of other robots to localize the teammate, sensor model, and the travel distance and time required to sufficiently localize the teammate. Robots with high value on helping decide which robot(s) will proceed via negotiation. This decision will depend on how uncertain the lost robot is (larger uncertainties will require more measurements to correct and more robots to cover the possible space in which the lost robot resides) and how far away the robot is from landmarks (if robots must drive farther out of sight of landmarks, their uncertainty will grow making observations less valuable and travel more risky).

For reliability, localization must also include the ability to explicitly self-localize in the event that assistance is unavailable or very costly. The robot must maximize its chance of localization by maximizing its chance of locating landmarks. The probability of locating landmarks will depend on the uncertainty in position and heading, the chosen direction of travel, and the distribution of landmarks within the area potentially traversed by choosing that direction. The resulting value of explicit self-localization depends directly on this probability.

3.3. Collaborative Mapping and Tracking

Collaborative mapping and observations attempts to take full advantage of the multi-robot team. Improved accuracy can be achieved by taking multiple observations on objects. Using individual sensor models, the value of observing at a point can be determined using the expected improvement from the observation and the cost to achieve that observation. By looking at sets of robot-position assignments, team deployment can be optimized for improving estimate quality. This optimization can be done for static or dynamic targets.

Analytical optimization can be done for simple value functions over small numbers of robots and objects. For larger teams and more complex

functions, acceptable solutions may be found using approximation techniques, such as hill-climbing, simulated annealing, space discretization, coarse-to-fine solutions, and search space reduction based on proximity to position and targets. Coarse solutions in a discrete space can provide a good starting position for further iterations or hill-climbing/simulated annealing.

For static targets, value may be directly evaluated as described. Efficiency will mostly be achieved through search space reduction and, if necessary, dividing larger teams into smaller teams. To leverage multiple observations by a single robot, selecting multiple points along or near a provided trajectory (defined by mission tasks, for example) reduces search.

For dynamic objects, when single robots cannot be assigned full-time tracking duty, iterative solutions become more important. The first solution estimate will depend on a preliminary position and velocity estimate, providing a range of area that may be traversed by the object. Initially, robots will take up vantage points around this area. As more information is collected, uncertainty is reduced and the area collapses, allowing agents to obtain more useful positions. The driving force will be time; good observation points lose value if they cannot be achieved before the target arrives. Also, the value of an observation point may need to include many potential target points.

4. PROGRESS

All of the implementation and experimentation to date have been conducted using *TeamBots* multi-robot architecture and Minnow robots.

4.1. Cooperative Observation

We have implemented and tested several aspects of cooperative observation, including the combining simultaneous observations of single static and dynamic objects from multiple robots and optimized selection of observation points of multiple static targets. This approach combines observations with Bayes' Rule, using a Gaussian distribution to represent position and simple mathematical updates. By combining observations, the resulting quality of position estimates is improved.

4.1.1. Combining Observations

The approach first transforms local parameters directly obtained from sensor data into global parameters, using geometrical relations and a rotation transformation, so that observations may be combined using a simple matrix approach. The result is a covariance and mean of the combined distribution:

$$C' = C_1 - C_1 [C_1 + C_2]^{-1} C_1$$

$$\hat{X}' = \hat{X}_1 + C_1 [C_1 + C_2]^{-1} (\hat{X}_2 - \hat{X}_1)$$

This approach assumes that positional uncertainty is absent or previously incorporated. It also assumes that measurement distributions are Gaussian and that observations are simultaneous. These assumptions make processing very efficient and real-time even for large teams.

4.1.2. Distributed Sensing Experiments

Several experiments were conducted in distributed sensing, including locating static targets, visually tracking dynamic objects, and increasing effective field of view of agents. In the first, an object was placed at a series of points and observed by three robots. Observations were combined in pairs and in all. The resulting position estimates provided more accurate results, with fewer outliers and lower maximum error. In the second, a robot is given the task of pushing a box to a goal position; the box is unseen at an unknown location. The robot receives a position estimate from a teammate (that cannot itself reach the object) allowing it to immediately proceed to the object without searching and manipulate it with a more accurate combined estimate. In the third experiment, three robots visually track a target by using the combined position estimate from all three robots' observations. When a robot is blindfolded, it can still track the target using the observations of teammates. This approach to distributed sensing was also employed in RoboCup in 2000, allowing increased coverage of the field.

4.1.3. Optimized Mapping Simulations

For optimized mapping, a team of robots is deployed to best improve a map of landmarks by optimizing the expected improvement and expected cost (resulting value). Expected improvement uses sensor models, landmark positions, and the method of combining observations to predict the distributions resulting from all observations. Expected cost uses driving time and distance. Robot position uncertainty is ignored, as it is hard to predict. Simulation results show that the resulting position estimates improve upon single-robot position estimates and on multi-robot position estimates with individually chosen observation points. Simulations also verify that, as expected, results are highly dependent on choice of value function and the relative weights of reward and cost; such values need to be tuned to the specific team and application desired.

4.2. Localization

Driven by the need for a localization approach requiring minimal processing and applicable with minimal (and noisy) landmark information and availability, we have developed and implemented a constraint-based approach to landmark localization. Localization has been implemented for individual robots but can easily include teammates as landmarks.

4.2.1. Localization Updates

Different types of measurements on landmarks lead to different types of geometric constraints on robot position and heading. A range, for example, provides a circular constraint around a point landmark while a bearing to a point landmark provides a single heading value and a linear constraint on position. These geometric constraints are easily computed algebraically. Each independent estimate can be combined, along with previous estimates and motion models, to efficiently determine a new position estimate. In this way, all available information is used without requiring the same quantity or quality of information from all landmarks or measurements. The current implementation does not yet provide any estimate of result quality. Additionally, the implementation does not yet use teammates as landmarks (they are not equipped to sense teammate position and orientation accurately), though this would be an easy expansion of the approach.

4.2.2. Localization Experiments

The localization approach was tested by driving a robot along a circuit, defined by waypoints, among landmarks. As the robot reached each landmark, it would record its position as recorded by on-board odometry and by localization; ground truth was also measured. Several experiments were conducted in which the number of landmarks was varied and type of landmark was varied for position estimation.

The localization approach eliminated the general odometry drift, even with very few landmarks that could only provide range for position estimation. As more information was added (such as bearing and more landmarks), performance further improved. Position and heading errors were non-zero, resulting from the noisy sensing, but were consistent through time and did not grow like odometry.

The approach was also implemented in RoboCup middle-size soccer, 2001. In this environment only four bearing-only landmarks are available, and never more than two are visible at any time. Qualitative results indicate that odometry drift is eliminated; robots returning to home position did not demonstrate the characteristic drift toward the edge of the field and the tendency to lose ability to orient down-field. Observed position errors in going home and practice remained on the order of 50 cm and 10 degrees, unlike the 2 m, 45 degree (or more) error observed in odometry in the same time.

5. CONCLUSION

We propose an approach to integrate collaborative localization and observation needs with the needs of a separate, complex mission. Robots will select an applicable task based on the current value of available tasks, including mapping and localization of self and of teammates. Mapping will be done by optimizing multiple simultaneous observation points for available

teammates, providing a team-optimized result. Similarly, observation of dynamic objects will optimize multiple, time-distributed observation points. Collaborative localization will not only take advantage of fortuitous or opportunistic encounters, but will allow explicit movement to self-localize or localize teammates.

Preliminary results in combining multiple observations, collaborative observation point selection, and localization indicate that these separate aspects are quite promising. Further work will be done in developing an integrated collaborative multi-robot system. This will include developing the framework for task selection as well as value functions for some mission tasks, mapping, and localization. This will also include reducing the search space for team deployment in observation of static and dynamic targets in order to provide acceptable results in an efficient manner. The system will be implemented and tested in simulation and on a real-robot team.

6. REFERENCES

- G. Dissanayake, H. Durant-Whyte, and T. Bailey. "A Computationally Efficient Solution to the Simultaneous Localisation and Map Building (SLAM) Problem." Proceedings ICRA, 2000.
- D. Fox, W. Burgard, H. Kruppa, and S. Thrun. "A Probabilistic Approach to Collaborative Multi-Robot Localization." Autonomous Robots on Heterogeneous Multi-Robot Systems, Special Issue, vol. 8, no. 3, 2000.
- D. Fox, W. Burgard, and S. Thrun. "Markov Localization for Mobile Robots in Dynamic Environments." Journal of Artificial Intelligence Research, vol. 11, 1999.
- S. Lenser and M. Veloso. "Sensor Resetting Localization for Poorly Modeled Mobile Robots." Proceedings ICRA, 2000.
- J. J. Leonard and H. F. Durrant-Whyte. "Dynamic Map Building for an Autonomous Mobile Robot." International Journal of Robotics Research, vol. 11, no. 4, 1992.
- I. M. Rekleitis, G. Dudek, and E. E. Milios. "Multi-Robot Collaboration for Robust Exploration." Proceedings ICRA, 2000.
- S. I. Roumeliotis and G. A. Bekey. "Synergetic Localization for Groups of Mobile Robots." *Proceedings of the IEEE Conference on Decision and Control, 2000*, 2000.
- A. W. Stroupe, M. C. Martin, and T. Balch. "Distributed Sensor Fusion for Object Position Estimation by Multi-Robot Systems." Proceedings ICRA, 2001.
- A. Stroupe, K. Sikorski, and T. Balch. "Multi-Modal Landmark Localization." Submitted to ICRA 2002.