

Heuristic Search Planning to Reduce Exploration Uncertainty

David Meger, Ioannis Rekleitis, and Gregory Dudek

Abstract—The path followed by a mobile robot while mapping an environment (i.e. an exploration trajectory) plays a large role in determining the efficiency of the mapping process and the accuracy of any resulting metric map of the environment. This paper examines some important aspects of path planning in this context: the trade-offs between the speed of the exploration process versus the accuracy of resulting maps; and alternating between exploration of new territory and planning through known maps. The resulting motion planning strategy and associated heuristic are targeted to a robot building a map of an environment assisted by a Sensor Network composed of uncalibrated monocular cameras. An adaptive heuristic exploration strategy based on A^* search over a combined distance and uncertainty cost function allows for adaptation to the environment and improvement in mapping accuracy. We assess the technique using an illustrative experiment in a real environment and a set of simulations in a parametric family of idealized environments.

I. INTRODUCTION

Exploration is a pre-requisite behaviour for many essential functions of a mobile robot. During localization and mapping, geometric information is gathered as the robot enters new areas. During visual search, the locations of potential objects are identified from images of new territory. During Sensor Network localization, the robot passes into the sensing or communication range of additional sensors. The common thread is that the system begins with no (or little) information about its environment, and additional information can only be collected when the agent moves into new territory.

For active information gathering tasks such as map building, decision making is an essential component which determines the quality of information collected. The robot's path determines the order and frequency of observation for each feature, which greatly impacts the accuracy of the final map produced, as well as the efficiency of the process. While both speed and accuracy are desired during mapping, these two goals are often in conflict. On one hand, *accurate* mapping is dependent on the robot's position estimate being corrected through repeated measurements of the same landmarks. On the other hand, *efficient* mapping demands minimizing distance traveled; thus, making a return to an already explored landmark undesirable.

D. Meger is with the department of Computer Science, University of British Columbia, Vancouver, British Columbia, Canada, dpmege@cs.ubc.ca

I. Rekleitis is an Adjunct Professor at the School of Computer Science, McGill University, Montreal, Quebec, Canada, yiannis@cim.mcgill.ca

G. Dudek is a faculty member of the School of Computer Science, McGill University, Montreal, Quebec, Canada, dudek@cim.mcgill.ca

All the authors worked at the Centre for Intelligent Machines, McGill University during this research.

At each exploration step, the environment can be partitioned into known and unknown regions. This suggests a natural decomposition of the planning problem into two sub-tasks. First, paths must be planned through the robot's current map, which, while errorful and incomplete, provides at least a rough estimate of the nature of the world. Second, in order to explore new territory, paths must be planned into or through the remainder of the environment which is initially unknown. The current map provides a relatively large amount of information for decision making, so there is some hope for selecting favorable paths in this setting. In contrast, planning through unknown regions is much more challenging, and appears to require heuristic strategies, unless strong prior information, or specific task properties are exploited. Figure 1 illustrates the information available to a robot when planning its motions in a hospital environment instrumented with a camera Sensor Network. At every instant, cameras which have previously observed the robot are candidates for re-visitation, and paths to these cameras can be planned quite accurately. Also, regions of so-far unvisited space give the opportunity for exploration, although the result of moving into these regions is somewhat less predictable.

This paper adapts and extends exploration techniques developed for mapping with a mobile robot to the context of camera Sensor Network self-localization - that is, a network of cameras whose precise positions must be determined by a mobile robot. Illustrative applications are building-security systems and traffic-monitoring networks. Such cameras provide a rich source of visual information for the regions in which they are emplaced and facilitate applications such as automated surveillance [1] and detection of abandoned luggage in airports [2]. These applications commonly assume a map of camera locations, as well as knowledge of the camera imaging properties; or, in other words, that calibration information is known *a priori*. This is rarely true in practice, but mapping and calibration can be completed by a mobile robot operating in the same environment as the camera network, as shown in [3].

This paper presents an exploratory trajectory planning solution for a robot exploring and localizing the cameras within a camera Sensor Network, such as the scenario depicted in Figure 1. Specifically, we propose the use of a planner based on A^* search to optimize local sections of the robot's path with respect to both distance traveled and map uncertainty. This method is derived from, but also extends previous work such as [4] since it provides a parameter which naturally adapts the levels of exploration and re-localization. In addition, we evaluate the effects of our planner when alternated with excursions into unexplored




Fig. 1. A robot's progress through an environment during exploration. Paths can be planned through the known map and to the border of unknown territory (dotted lines). Camera observations (large dots) provide the sensor readings which allow for accurate mapping, particularly when the robot plans to revisit a camera numerous times.

territory, so that our method can be considered a complete exploration algorithm.

The next section will review necessary background material regarding SLAM as well as previous methods for planning to reduce map uncertainty. Section III describes the Localization and Mapping solution for a mobile robot in a camera Sensor Network considered in this paper. An adaptive heuristic search based planner for exploration paths is introduced in Section IV. Experimental results in Section V demonstrate the efficacy of the network localization solution in a large indoor environment and illustrate the performance of the exploration planning methods in simulation.

II. BACKGROUND

The network localization problem is similar to Simultaneous Localization and Mapping (SLAM) since both scenarios involve estimating the pose of the robot and the positions of environment features (landmarks or sensor nodes) from acquired sensor data. Hence, similar estimation approaches are appropriate. In this paper, the extended Kalman filter (EKF) as described in [5] for SLAM is adapted for camera

network localization. The EKF computes the mean μ and covariance P for each map quantity. Many other solutions are possible, but the EKF is used here for computational simplicity and ease of analysis.

Numerous authors have studied the problem of planning paths through the already known map in order to gather additional information and to increase mapping accuracy, e.g. [4], [6], [7], [8]. Many approaches have attempted to reduce the entropy in the map estimates [9], [10], [11], which is the measure of the uncertainty in a distribution and is defined as:

$$H(p(\xi)) \equiv - \int p(\xi) \log(p(\xi)) d\xi \quad (1)$$

For the Gaussian distributions used by an EKF representation of the environment, entropy can be expressed in closed form. Sim and Roy [4] discuss two different measures from information theory for which either the trace or the determinant of the covariance matrix provides the final measure for entropy.

Early work proposed a single-step, greedy choice of the action which maximally minimizes the entropy because optimal planning of multi-step paths requires computational cost exponential in the path length. Recently, Sim and Roy [4] have proposed pruning loops during breadth first search in order to ensure manageable complexity even when planning longer paths under conditions of idealized sensing and a rough initial estimate of landmark locations. In addition, [6] has considered a simulation-based approach which has the potential to generate multi-step paths at the cost of significant computation.

As mentioned earlier, accuracy and efficiency are conflicting goals during exploration. In order to produce paths that compromise between the goals, distance and uncertainty have to be combined into a single cost function. Unfortunately, the two are incommensurable; that is, they lack common units for comparison, so care must be taken in combining their values. Makarenko *et al.* [10] have previously proposed a weighted linear combination of distance and uncertainty for path p :

$$C(p) = \omega_d \text{length}(p) + \omega_u \text{trace}(P(p)) \quad (2)$$

In this cost function, P is the covariance matrix resulting from the EKF and its trace is an approximation of the uncertainty in the map. The choice of weighting factors ω_d and ω_u represents the compromise between distance traveled and mapping uncertainty or *accuracy* versus *efficiency*. We would like to produce a flexible method based on varying only one intrinsic parameter, so normalizing each quantity by its maximum possible value would be desirable. Unfortunately, computation of the maximum possible uncertainty over all paths is intractable, so this value is approximated by simulating the robot along a straight path with no relocalization, which is very likely an overestimate. Once each quantity has been normalized, a single free parameter α in the range $[0, 1]$ is able to specify the contribution of each factor. Based on this formulation, the weights used in our cost function are:

$$\omega_d = \frac{\alpha}{\text{maxdist}} \quad , \quad \omega_u = \frac{1 - \alpha}{\text{maxuncert}}$$

By setting α to the two extremes, zero and one, it is possible to consider only one of the factors at a time: distance only, by setting $\alpha = 1$, and uncertainty only, by setting $\alpha = 0$. Section IV will discuss the effect of varying α on the quality of the resulting paths.

Several authors have considered the collaboration between a Sensor Network and a mobile robot in different sensing scenarios and in some cases with much more capable robotic agents [12], [13]. In addition, our analysis of a single mobile agent within a static network can be viewed as a special case of multi-robot collaboration, which has also been studied by many authors [14], [15], [16].




Fig. 2. The experimental setup used throughout this paper. The robot carries a calibration target which can be easily detected in images taken by the cameras in the network (such as the one mounted on a door here).

III. LOCALIZATION AND MAPPING IN A CAMERA SENSOR NETWORK

An autonomous solution for calibration and mapping of a camera sensor network by a mobile robot has been previously presented in [3], [17]. The crucial details of this method will be reviewed in this section to provide sufficient background to enable subsequent discussion of the exploration planning algorithms.

Given a network of cameras placed inside a building, and a mobile robot, the goal is to autonomously explore the building; locate each camera, by receiving an alert every time the robot enters the field of view (FOV) of the camera; calibrate the internal parameters of the camera; and finally, recover the 3D pose of the camera with respect to the robot. The first step is recognizing the robot when it enters the FOV of a camera. This is accomplished with a specially constructed target mounted on the robot which can be robustly detected in visual imagery; see Figure 2. Our target is comprised of ARTag [18] fiducial markers which have been employed for automated camera calibration and pose estimation previously by [19]. The calibration procedure estimates the pose of the camera with respect to the robot, also known as extrinsic parameters. This information allows the pose of the camera to be added to the map. Finally, an EKF tracks the location of the robot as it moves between cameras and corrects the robot's location as well as the map of camera locations each time a measurement is collected.

As mentioned above, the decision making aspect of exploration is crucial to any map building approach, but particularly to mapping a camera Sensor Network. In this scenario, the robot travels over potentially large distances between camera observations, so odometry error accumulates drastically unless it is corrected actively. That is, camera positions serve as landmarks for planning which must be visited and revisited in order to keep the accumulated uncertainty low. Position measurements obtained when a camera observes the robot repeatedly allow the estimator to reduce

the positional uncertainty of the robot and of the map. To illustrate this point, and to give a basis for comparison, several hand-crafted methods for planning during exploration were considered in this work which explicitly considered the compromise between accurate and efficient mapping:

- **Depth-first exploration:** the robot always moves into unexplored territory, never relocalizing. This strategy provides coverage of the environment with minimal distance traveled, but the uncertainty of the robot position grows quickly. It is worth noting that as the environment is unknown, DFS is the most efficient method possible.
- **Return-to-Origin:** the robot alternates between exploring a new camera position and returning to the first camera it mapped, which has the lowest uncertainty. This strategy allows for accurate relocalization, but it means that the robot must travel increasingly larger distances as it maps cameras further away, thus, introducing measurements from increasingly inaccurate positions.
- **Return-to-Nearest:** in a compromise between the two previous methods, the robot alternates between exploring a new camera position and relocalizing at the nearest previously explored camera. The ability to relocalize accurately depends on the uncertainty of the nearest camera only, which might not be mapped as accurately as cameras which are farther from the robot, however, regressing by only one camera at a time means the extra distance traveled remains small.

IV. ADAPTIVE HEURISTIC PLANNING FOR MAP BUILDING

Exploration of a Sensor Network can be thought of in terms of the graph formed by nodes corresponding to sensor positions connected by edges corresponding to traversable pathways between sensors. This construction will be useful to explain our proposed solution. In graph-based terms, the exploration process consists of two steps: 1) selecting the next node to visit, and 2) planning the best path through the known graph to reach the selected node. The three-hand crafted approaches described earlier undertake these two steps without consideration of the current state of the map and estimator. A planning algorithm which examines the current uncertainty of each camera’s estimated position is able to adaptively select paths which allow better relocalization. An optimal solution could be produced by searching the space of all possible paths and choosing the one with the lowest cost as defined by Eq. 2. Unfortunately, we can show that the number of paths through such a graph-like environment has a worst-case bound of d^k , where d is the maximum node degree in the graph and k is the path length [19]. Thus, exhaustively examining all paths is intractable.

Inspired by the pruned search method of Sim and Roy [4] we consider approximating the minimum cost path in a computationally efficient fashion. To avoid exhaustive search, our planner employs a variant of A^* search. As all informed search techniques, A^* search requires two pieces of information about each node in the graph: $f(n)$, the cost of the

best path found so far from the start to node n ; and $h(n)$, the heuristic function, which is an estimate of the cost from node n to the goal based on some, hopefully cheap, approximation. The search procedure expands nodes in order of increasing expected cost C which is a combination of the two terms:

$$C(n) = f(n) + h(n) \quad (3)$$

We use Eq. 2 for our f function as we seek to optimize paths based on *both* distance and uncertainty. The h function must be “admissible”. That is, it must be an underestimate for the remaining cost that will be required to reach the goal. Computing a reasonably tight lower bound for the uncertainty reduction along a path appears to be an exponential problem in itself. So, noting that A^* with a less informative heuristic function is still more efficient than breadth-first search, we simply leave out the contribution of uncertainty. Instead our h function is the straight-line distance between nodes, as is common in traditional planning. This allows our algorithm to reduce the number of searched nodes, with little overhead computation. Note that uncertainty still guides our search to some extent, as it is a component of the cost function f .

It is important to note that A^* in our case is only approximating optimal paths with respect to the chosen cost function because the situation we consider violates the standard assumptions in two aspects. First, in our context paths which contain loops can often reduce uncertainty, while traditional AI planning techniques often assume no negative weight cycles. Second, our cost function uses a scalar representation of the EKF covariance matrix. Uncertainty in subsequent planning steps cannot be predicted based on such a simple value. In other words, the $trace(P)$ measure used by our planner is not a sufficient statistic, which is required for optimal planning. This lack of optimality is unavoidable since exploration planning is exponentially complex, but is noted here in order to provide a more complete understanding and to provide directions for future work. It is also important to consider the potential for generalization of the method. Since our analysis does not explicitly depend on any properties of the camera sensors used, it could easily be applied to other graph-based exploration problems, such as, Generalized Voronoi Graphs (GVG) [20] and occupancy grids [21].

Figure 3 shows paths generated for four values of the α parameter which weighs the contributions of distance and uncertainty. This example illustrates that A^* search is able to adapt to the environment and compromise between the conflicting goals of efficiency and accuracy. The next section provides results from further experiments which will evaluate the impact of such adaptation on the exploration process.

V. EXPERIMENTAL RESULTS

A. Hallway Mapping

A set of hardware experiments was performed to verify the underlying calibration and mapping framework. A network of seven cameras was deployed in an indoor environment

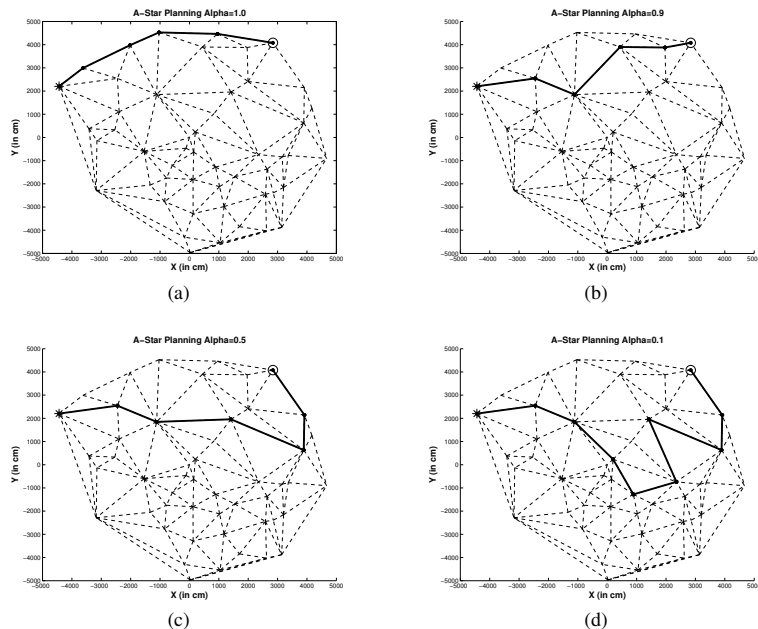


Fig. 3. Paths generated by A^* search using a distance and uncertainty cost function for four values of α . Dark lines indicate the path followed by the robot in each case. At (a) only the distance is considered. The relative contribution of uncertainty increases sequentially for (b)-(d).

with a diameter over 50 m and a mobile robot equipped with the fiducial target described in Section III traversed the environment. The robot passed in front of the cameras and followed a path whose length was over 360 m. Figure 4(a) shows that the estimate of the robot's position based on odometry accumulates error quickly. Figure 4(b) demonstrates that the calibration and mapping system is able to correct this error and produce an accurate map of the environment. For these preliminary experiments, a naive exploration policy was employed. These experiments served to validate the utility of the approach and demonstrate that apparently accurate mapping could be achieved in practice even with very simple path planning. Moreover, these experiments provided verification of the EKF implementation as well as experimentally derived noise statistics for the odometry and the camera estimates. More informed trajectory planning methods will be considered in the remainder of the experimental results.

B. Simulation Environment

We used numerical simulation to evaluate the performance of the approach over a wide range of parametric variations in the environment. This environment was meant to emulate the properties of camera networks such as the one used in the hallway experiments as closely as possible. To accomplish this, nodes were chosen from a uniform distribution over free space with properties similar to the physical hallway setup. The camera heights and distances from the robot location were allowed to vary randomly, with mean values equal to those found in the hallway tests. The same EKF implementation was used for state estimation as was used

to produce Figure 4(b). Various graph sizes were produced by altering the number of nodes so that trends in the exploration results could be examined. Two idealized classes of graphs were used which exemplified near-extremes in terms of connectivity (though the reader is reminded that our solution is independent of particular graph structure). The first class considered was **fully connected graphs** (i.e. cliques) where every pair of nodes is connected by an edge, which represents a scenario where the robot is able to freely traverse the environment without obstacles; see Figure 5(a). The second class examined was **triangulated graphs** where edges are chosen by triangulation of the nodes to produce a planar graph, which again represents obstacle free space, but in this case assumes the robot should move through a sequence of nodes along its path; see Figure 5(b). We also examined graphs generated by stochastically sampling a real environment to produce a roadmap. In this case, we used the floorplan of an actual hospital and triangulated stochastically selected sample locations which were mutually visible. We refer to these as **hospital graphs**; see Figure 5(c). In the remainder of this section, results have been computed over a mixture of the graph types.

C. Single Path Results for the Adaptive Heuristic

The goal of our adaptive heuristic planner is to choose a short-term path through the known graph that allows the robot to arrive at a new node with minimal distance and uncertainty. The simulation environment described earlier was used in order to analyze the performance of our algorithm. For each randomly generated network, the **Return-to-Nearest** strategy was executed for a portion of the network




Fig. 4. (a) Odometry Readings for Hallway Path. (b) EKF Estimate of the Hallway Path. Estimated camera positions with uncertainty ellipses are in red where color is available.

in order to initialize camera estimates in the EKF. A series of planners were then executed to find paths between constant start and goal nodes so that paths returned could be compared. The planners evaluated used A^* search with four different values for α : (0.1, 0.5, 0.9, 1.0). Note that $\alpha = 1.0$ produces shortest distance planning.

Figure 6(a) illustrates the distance traveled for the four choices of α , averaged over twenty instances for each network size. As expected, larger α values produce shorter path lengths since distance is weighted more heavily in the cost function. The relatively graceful increase in distance traveled as α decreases indicates the ease with which the planners are able to find slightly longer paths that perform better with respect to final uncertainty. That is, there is no catastrophic degradation in distance performance as the weighting is changed.

Figure 6(b) shows the final robot uncertainty upon arrival at the chosen goal node for three of the four α values used. Results for $\alpha = 0.5$ are excluded because the results for this method lie extremely close to $\alpha = 0.1$ and $\alpha = 0.9$, which makes visualization difficult. Setting $\alpha < 1$ manages

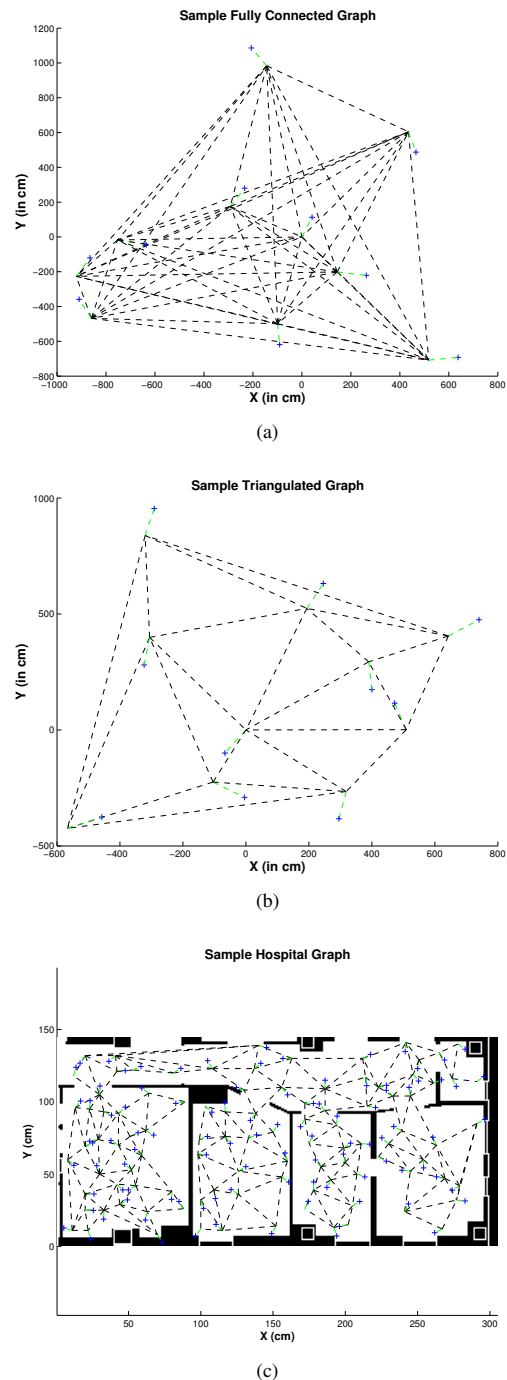


Fig. 5. Example graphs from each of the three types considered. Dotted lines indicate edges between nodes. Camera positions are shown as blue crosses where color is available.

to reduce the uncertainty drastically over the $\alpha = 1$ case, where only distance is considered. In fact, the improvement in uncertainty between $\alpha = 1$ and $\alpha = 0.9$ is much larger than that between $\alpha = 0.9$ and $\alpha = 0.1$. These results are of particular importance since they indicate that by considering uncertainty in the planning process, quality increased dramatically as can be seen in Figure 6(b).

D. Global Exploration Results

The previous section indicates that the adaptive heuristic is able to produce relocalization paths with intuitively favorable and adaptive properties. This method is extended to global exploration which requires iterating two steps: following a previously untraversed arc from a frontier node of the known graph and selecting a relocalization path through nodes which have previously been mapped to arrive at the next frontier node. The first step, which poses a problem similar to that solved by the Frontier-Based Exploration strategy [22] in occupancy-grid SLAM, is challenging since no measurements have yet been made about the destination camera location. We are not able to compute expected distance or uncertainty for this exploratory section of the path. Under this condition, several hand-crafted strategies have been attempted for choosing the next exploration action to take. These include a breadth-first search method where nodes closest to the origin are explored first, and also a spiral search method. In future work, we plan more detailed analysis of the effects of each of these strategies as well as more sophisticated methods such as the use of the A^* search algorithm to evaluate each frontier node. For the purposes of evaluating the relocalization strategies in this section, the same breadth-first exploration ordering was used for each method. The second step in exploration involves selecting the path through the known graph ending at the destination frontier node, which allows for accurate robot relocalization with minimal distance traveled. Each of the hand-crafted trajectories solves this problem by making a generic choice of node to use for relocalization at each stage. The adaptive relocalization strategy can be used instead in order to provide the additional flexibility and relocalization ability that was demonstrated in the previous section. Repeated simulated explorations were conducted to compare the hand-crafted trajectory methods mentioned earlier to the adaptive strategy.

Figure 7(a) illustrates the total distances traveled. These results are not surprising; the **Depth-first** strategy covers the environment with the least robot motion, **Return-to-Nearest** requires slightly more motion, the adaptive heuristic of A^* with $\alpha = 0.9$, slightly more again, and **Return-to-Origin** requires the largest distance traveled. Figure 7(b) presents the final map uncertainty results. The adaptive global strategy is able to produce maps with lower uncertainty than any of the static methods due to the fact that it uses all of the information available in order to choose paths which exploit properties of the current estimate. In terms of computation time, the A^* was slower than the other heuristics without being excessively slow.

VI. CONCLUSIONS

Robot path planning has been considered for reduction of map uncertainty in the context of a mobile robot calibrating and mapping a camera Sensor Network. An adaptive heuristic which produces relocalization trajectories based on the current state of the estimator has been shown to improve performance over a variety of intuitive hand-crafted approaches. The adaptive heuristic planning is able to provide a compromise between efficiency and accuracy in planning relocalization paths. This translates to favorable performance in global mapping when compared to less adaptive strategies; however, there is still room for improvement to the technique by searching more efficiently and explicitly considering overall map uncertainty.

Perhaps the most promising area of future work is the use of the adaptive heuristic trajectory planning in different localization and mapping domains. Numerous previous techniques for planning to reduce map uncertainty employ greedy planners based on entropy reduction techniques. By nature, greedy planning is far from optimal, given that it does not attempt to exploit all of the information available. The A^* search method presented here provides for greater adaptation, while managing to limit computation through the use of a heuristic function to guide the search for solutions. Guiding a stochastic or sampling-based approach such as the method employed by [6] in a similar fashion would allow combination of the benefits of both methods.

The authors believe similar methods can be applied in other mapping domains such as the landmark based EKF, occupancy grid representations such as FastSLAM, and in fact any representation where uncertainty is explicitly modeled in the estimator. The use of such adaptive heuristics will allow robotic mapping to occur with lower error and contribute to the autonomy of robotic agents in general.

REFERENCES

- [1] O. Javed, Z. Rasheed, O. Alatas, and M. Shah, "Knight: a real time surveillance system for multiple and non-overlapping cameras," *The fourth International Conference on Multimedia and Expo (ICME 2003)*, 2003.
- [2] N. Krahnstoeber, P. Tu, T. Sebastian, A. Perera, and R. Collins, "Multi-view detection and tracking of travelers and luggage in mass transit environments," in *Proceedings of the 9th IEEE International Workshop on Performance Evaluation of Tracking and Surveillance*, New York, NY, 2006, pp. 67–74.
- [3] D. Meger, I. Rekleitis, and G. Dudek, "Autonomous mobile robot mapping of a camera sensor network," in *The 8th International Symposium on Distributed Autonomous Robotic Systems (DARS)*, Minneapolis, Minnesota, July 2006, pp. 155–164.
- [4] R. Sim and N. Roy, "Global a-optimal robot exploration in slam," in *International Conference on Robotics and Automation*, 2005, pp. 661 – 666.
- [5] R. Smith, M. Self, and P. Cheeseman, "Estimating uncertain spatial relationships in robotics," *Autonomous Robot Vehicles*, pp. 167 – 193, 1990.
- [6] R. Martinez-Cantin, N. de Freitas, A. Doucet, and J. Castellanos, "Active policy learning for robot planning and exploration under uncertainty," in *In Proceedings of Robotics: Science and Systems (RSS)*, 2007.
- [7] T. Kollar and N. Roy, "Using reinforcement learning to improve exploration trajectories for error minimization," in *In Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, Orlando, 2006.

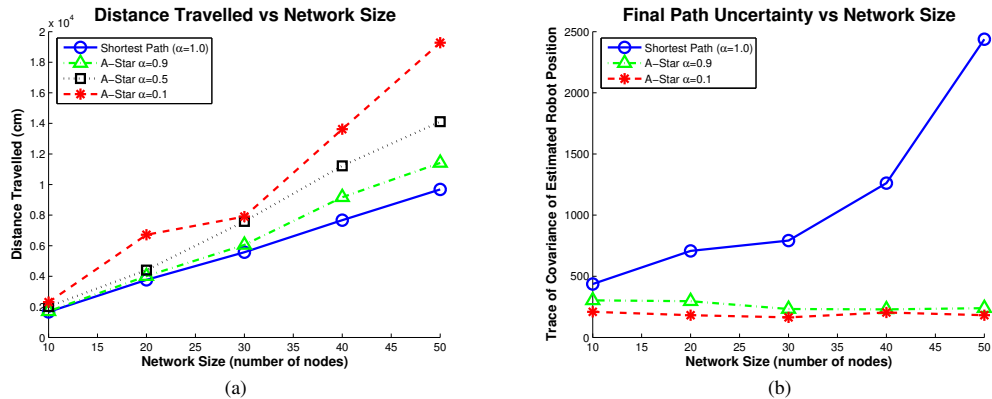


Fig. 6. Results of adaptive relocalization: (a) The distance required to reach a goal node in a partially explored graph is shortest with $\alpha = 1$ representing shortest path planning and increases as α is decreased. (b) The uncertainty with which the robot reaches the goal node in a partially explored graph is largest with $\alpha = 1$ representing shortest path planning and decreases as α is decreased.

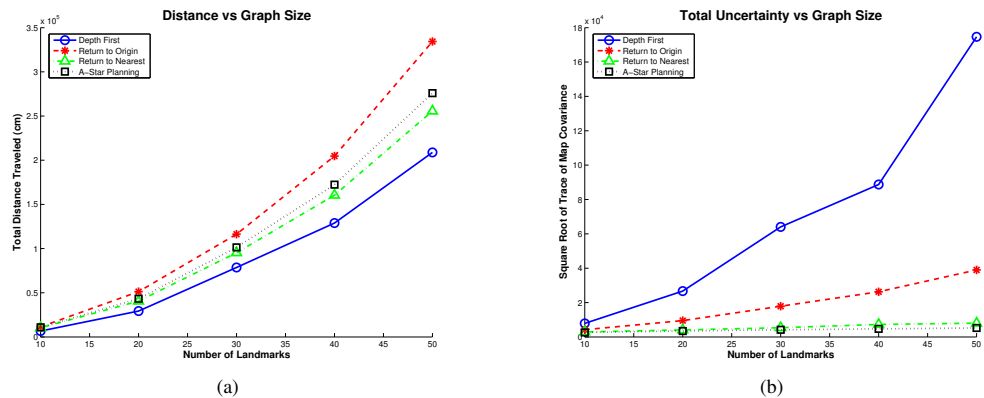


Fig. 7. Results for various exploration strategies: (a) The total distance required to explore the environment for each of the strategies considered. (b) The final map uncertainty for each of the strategies considered.

[8] D. Fox, W. Burgard, and S. Thrun, "Active markov localization for mobile robots," *Robotics and Autonomous Systems*, 1998.

[9] C. Stachniss, D. Haehnel, and W. Burgard, "Exploration with active loop-closing for fastslam," *International Conference on Intelligent Robots and Systems*, 2004.

[10] A. Makarenko, S. Williams, F. Bourgault, and H. Durrant-Whyte, "An experiment in integrated exploration," *International Conference on Intelligent Robots and Systems*, 2002.

[11] S. Hang, N. Kwok, G. Dissanayake, Q. Ha, and G. Fang, "Multi-step look-ahead trajectory planning in slam: Possibility and necessity," *International Conference on Robotics and Automation*, 2005.

[12] M. Batalin and G. S. Sukhatme, "Coverage, exploration and deployment by a mobile robot and communication network," *Telecommunication Systems Journal, Special Issue on Wireless Sensor Networks*, vol. 26, no. 2, pp. 181–196, 2004.

[13] P. Corke, R. Peterson, and D. Rus, "Localization and navigation assisted by cooperating networked sensors and robots," *International Journal of Robotics Research*, vol. 24, no. 9, 2005.

[14] K. O'Hara and T. Balch, "Distributed path planning for robots in dynamic environments using a pervasive embedded network," in *Third International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, 2004.

[15] A. Mourikis and S. Roumeliotis, "Performance analysis of multirobot cooperative localization," *IEEE Transactions on Robotics*, vol. 22, no. 4, to appear - Aug. 2006.

[16] I. M. Rekleitis, G. Dudek, and E. Milios, "Multi-robot collaboration for robust exploration," in *Proceedings of International Conference in Robotics and Automation*, San Francisco, USA, April 2000, pp. 3164–3169.

[17] I. Rekleitis, D. Meger, and G. Dudek, "Simultaneous planning, localization, and mapping in a camera sensor network," *Robotics and Autonomous Systems*, vol. 54, no. 11, pp. 921–932, November 2006.

[18] M. Fiala, "Artag revision 1, a fiducial marker system using digital techniques," in *National Research Council Publication 47419/ERB-1117*, Nov. 2004.

[19] D. Meger, "Planning, localization, and mapping for a mobile robot in a camera network," *Master of Science Thesis - supervisors Ioannis Rekleitis and Gregory Dudek*, 2007.

[20] H. Choset and J. Burdick, "Sensor based planning, part ii: Incremental construction of the generalized voronoi graph," in *Proc. of IEEE Conference on Robotics and Automation*. Nagoya, Japan: IEEE Press, May 1995, pp. 1643 – 1648.

[21] H. Moravec and A. Elfes, "High-resolution maps from wide-angle sonar," in *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, St. Louis, MO, USA, 1985, pp. 116–121.

[22] B. Yamauchi, A. C. Schultz, and W. Adams, "Mobile robot exploration and map-building with continuous localization," in *IEEE Int. Conf. on Robotics and Automation*, Leuven, Belgium, 1998, pp. 2833–2839.