

Autonomous Exploration via Regions of Interest

Robert Grabowski, Pradeep Khosla and Howie Choset

Carnegie Mellon University
Pittsburgh, Pennsylvania 15213
{grabowski, pkk, choset}@cs.cmu.edu

Abstract

*We describe a new paradigm for exploration of unknown spaces based on maximizing the understanding of obstacles rather than the exposure of free space. We look at the interaction between multiple sensor readings and how they combine to resolve obstacles. Taking a **next best view** approach, we generate an inverse sensor model that identifies regions in space where a new sensor reading has maximal utility with respect to increasing the resolution of that reading. Fusion of multiple models is exploited to generate regions of interest that direct exploration in such a way as to maximize the robots understanding of its space. These techniques are applied to a team of small robots called Millibots.*


1 Introduction

In order to operate independently and effectively, a robot must be able to autonomously explore its own space. Autonomous exploration is a recursive process that utilizes the relationship between sensing and movement of the robots through a map. Features of the map such as polylines, frontiers, openings and edges are extracted and processed to direct the movements of the robot to further generate new information. Ultimately, the success of exploration depends on the quality of the coupling between sensing, map building and navigation.

2 Free-Space Approach

Researchers have posed many variants to the exploration process [1][3][8][9]. Most are based on developing a global map and assessing the free-space of that map to generate plans for moving the robots. That is, robots are directed to new places in the map that will clear the most open area. A commonly employed technique is frontier exploration where the robot utilizes an occupancy map generated by fusing a collection of range readings obtained throughout motion. A threshold function is then applied to the occupancy map to separate the map into two classes – obstacles and frontiers. Points between open, explored space (low occupancy) and closed, explored space (high occupancy) represent the edges of obstacle. Regions between open explored space (low occupancy) and unexplored space represent frontier points. The collection of adjacent frontier points is dubbed a *frontier*.

Yamauchi achieves exploration by selecting one of the existing frontiers and directs the robot towards it. In doing so, the frontier is expanded [9]. Simmons extends this notion to robot teams by developing an arbitration method for splitting frontiers among multiple members of a cooperating team to minimize search overlap [8]. Similarly, Burgard applies a cost metric for maximizing free space exposure for a team of robots [3].



A team of small robots dubbed Millibots

Banos provides the first hint as to the importance of obstacles in the exploration process. He utilizes map information to define the contours of both free space and existing obstacles. He then guides his exploration in such a way as to maximize free space exposure while maintaining a minimum degree of contact with these obstacles for the purpose of localization [1]. However, this approach does little to ensure the quality of the definition of the obstacles being utilized and fails to improve that quality during exploration.

The fundamental shortcoming of the free-space approach is that it fails to account for the importance of obstacle definition in the exploration process. First, obstacle definition is important in the navigation sense. In autonomous exploration, local features such as polylines, walls, doors and edges are derived from obstacles to generate the next actions of the robot. Consequently, if the representation of local obstacles is poor, any actions derived from that representation are equally flawed. Critical targets (with respect to a given mission) and even viable search areas may be missed during the exploration process due to choices based on poor map information. Second, robots are increasingly utilizing map features to correct for inherent errors in odometry. A robot can exploit the differences in position and orientation with respect to common references before and after local moves to correct for these errors. Such an approach depends on the comparison between the extracted features before and after each move. If the obstacles are not defined well, the features are less discernable and the correlation is weaker. In the free-space approach, there is no emphasis on gaining a clear understanding of the obstacles that make up the world. Obstacle definition is opportunistic and relies on the assumption that the robot

will gain the necessary information as it moves.

3 Occupancy Maps

A critical step to obstacle-based exploration is the ability to assess the current quality of the map representation and subsequently the ability to extract relevant features from it. That is, how much confidence can we place in the definition of the features extracted from the map. The most common map representation for robots utilizing range sensors is the occupancy map [5]. The environment is divided into a grid of homogeneous cells where each cell represents the probability that the corresponding region in the world is occupied. Occupancy values range from zero to one. An occupancy value near zero corresponds to an open cell and indicates with highest probability that the corresponding region is free of any obstruction. An occupancy value near one indicates the opposite or that the region is occupied. When no previous knowledge is available the occupancy value of each grid is initialized at 0.5 (equally likely to be occupied or free).

Given a range reading and its position in space, a sensor model is used to infer the occupancy of that region. In the case of sonar sensors, the sensor model takes on the shape of an arc with the distance to the end of the arc equal to the range value for that reading (Figure 1b). The area within the arc represents low probabilities of occupancy while readings near the outer edge of the arc represent high probabilities of occupancy. Areas outside the arc are not changed. Each time a range reading comes in, the model is applied and the occupancy map is updated. The occupancy map is updated by fusing the local sensor's notion of occupancy with the existing global notion of occupancy via a Bayesian update rule [5].

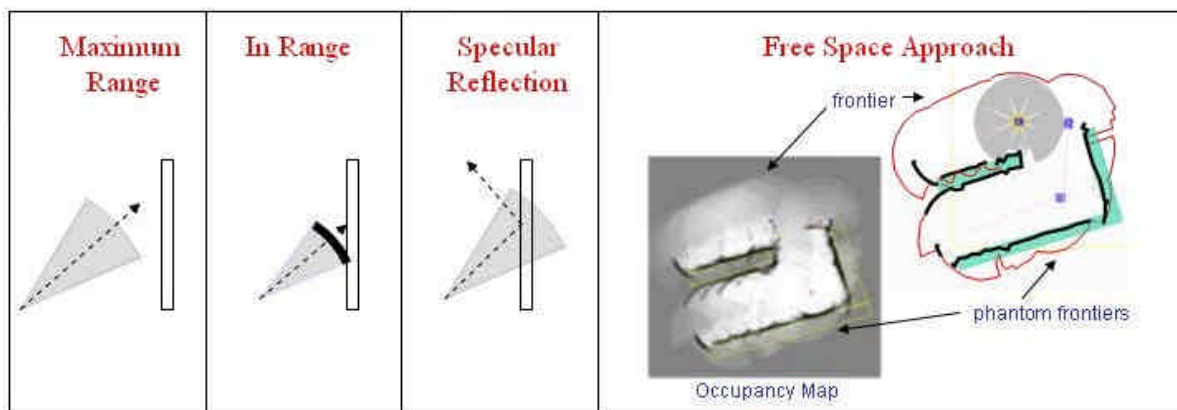


Figure 1: Specular Reflection

(a) A sonar reading taken of an obstacle beyond range. It updates only the free space hypothesis. (b) A sensor reading from an adjacent obstacle. It updates both free space and the likely position of the obstacle (c) A sonar reading experiencing specular reflection. The map is updated as if no obstacle is present (d) Multiple readings build a global occupancy map (left). Dark areas represent obstacles. Light areas indicate free space. Specular reflection produced phantom frontiers that project beyond existing obstacles.

4 Occupancy Map and Specular Reflection

One major shortcoming of employing the occupancy map for exploration is specular reflection and its effect on the boundary definitions of the map. Specular reflection is a phenomenon where the energy being emitted by the sensor strikes an adjacent object but the incidence angle is sufficiently shallow that it causes the return echo to be reflected away from the detector. Even though the obstacle is close, the sensor returns a maximum range reading erroneously indicating the space is open (Figure 1c). Specular reflection has the adverse effect of sweeping out large areas of the map beyond adjacent obstacles.

However, a similar case arises when there is no obstacle within range of the sensor. In this case, no obstacle is present so no echo is returned. Consequently, the sensor times out and returns a value indicating maximum range (Figure 1a). The sensor model is adjusted to update only the areas inside the arc and ignores its edges. Given the limited sensor range of most sensors and the potentially vast spaces of the environment, maximum range readings account for a large portion of the makeup of the map.

Unfortunately, without additional information, it is impossible to disambiguate between range readings that derive from no reflection and those that are the result of specular reflection. While additional local sensor readings, taken from different poses, may resolve the closed space, it does nothing to resolve the erroneously reported open space. This can have significant implications when utilized in the free-space approach. Without some means for discriminating between the two, specular reflection can result in the generation of phantom

frontiers. Phantom frontiers are regions that extend beyond the sensed side of an obstacle and erroneously indicate viable frontiers of exploration (Figure 1d). If a robot bases its exploration strategy on a map with phantom frontiers, it will generate ill-posed search solutions such as directing the robot around obstacles to regions that are ultimately inaccessible. Conversely, in an obstacle-based approach, exploration is conducted with respect to local obstacle definition. Therefore the robot is not guided by the generation of free space and does not suffer from the effects of specular reflection on the occupancy map.

5 Next Best View

We seek a method that not only generates new information about the unexplored space of the map but also ensures an adequate level of quality of the representation of existing obstacles. To accomplish this, we take a *next best view* approach to sensing. That is, where can we take the next sensor readings as to improve the resolution of existing obstacles? The next best view (NBV) philosophy is not new to robotics. However, the majority of NBV approaches resides in the vision domain and is primarily based on assessing the geometric properties of occlusion [2][7]. Banos is one of the few to have applied the next best view approach to navigation. He chooses positions that seek to maximize exposure to free space while maintaining a minimum degree of contact with existing obstacles [1]. However, none of these described techniques seek to improve the resolution of existing obstacles or takes into account for the




Figure 2: Developing an Inverse Sensor Model

(a) Geometric representation of sensor occupancy model impinging on an obstacle at the origin. (b) Inverse sensor model developed by determining regions where a second measurement would be sufficiently separate in bearing from the first. Regions close to original bearing would be more coupled so are suppressed. Regions away from original bearing add greater independent information and are desirable. (c) The next best view with respect to the current point can be found by choosing a reading inside the inverse sensor model.

interaction between sensor and obstacle.

6 Inverse Sensor Model

Our main contribution to this paper is the way we develop the *next best view* approach. We can be more efficient in our movement and sensing if we take advantage of the interaction between obstacles and sensors. In conventional map building, we are given a range reading and its position and utilize a sensor model to derive a notion of occupancy (Figure 2a). We apply this model for each cell in the occupancy map that the sonar model impinges. Then via an update rule, the derived occupancy from the sensor model is fused with the existing notion of occupancy in the map. Geometrically the occupancy model for a sonar sensor takes on the appearance of a wedge (Figure 2a). Points that lie inside the wedge are less in range than the reading and consequently low in probability of occupancy. Points that lie on the outer edge of the wedge indicate a high probability to occupancy. However, due to the poor angular resolution of sonar, there is an equally likely chance of an obstacle being present anywhere along the arc of the segment. Consequently, the ability for that occupancy cell to represent an underlying obstacle is reduced.

However, with respect to resolving ability, we can turn this around to develop an inverse sensor model. We start with an existing obstacle point in the occupancy map and derive respective positions where a new sensor reading would better resolve the underlying obstacle. We develop the parameters of the inverse model based on the effect of fusing together two (or more) sensor readings impinging on the same point. First, we define a metric for assessing the resolution ability of multiple sensor readings in an occupancy map as a function of separation in pose (Figure 3). If we consider the readings to be conditionally

independent, then the combined result is the product of the individual probabilities. We assert that the relationship between the geometrical overlap (common area) of two or more sensor models is directly related to the dependence in their readings with respect to defining that point. In terms of resolution ability, the tighter the combined distribution (or overlap in distribution), the better the definition of that point in space with respect to the underlying obstacle. For example, Figure 3b shows the definition of a point developed from fusing two sensor readings. When the overlap is large (small separation angle), the distribution of the fusion is widely distributed. As a result, the resolution of the underlying point is poor. Conversely, when the overlap is small (large separation angle), the distribution is tight. The result of a tight distribution is that the resolution of the point is higher. Figure 3c shows a plot of the overlap of two sonar readings as a function of separation angle. In this test, one sensor reading is fixed while the other is rotated about the point under test (Figure 3a). The combined distribution of two sensor readings is measured via a Monte Carlo method and counts only points above the obstacle threshold used in feature extraction. For sonar sensors, the overlap is greatest at 0 degrees and quickly falls off after 15 degrees. Consequently, we assign a higher utility to readings that come from bearings that exceed 15 degrees.

Choset reaches a similar conclusion in developing the Arc-Transversal Median (ATM) algorithm [4]. In this work, Choset describes a robust method for indicating the location of an obstacle by looking at the intersecting points of adjacent sonar arcs. He makes this method robust by only considering readings that “stably” intersect. A stable intersection is a pair of readings whose intersection does not significantly change with small perturbations in either of the readings. The ATM work derives a critical intersection angle of 30 degrees. The

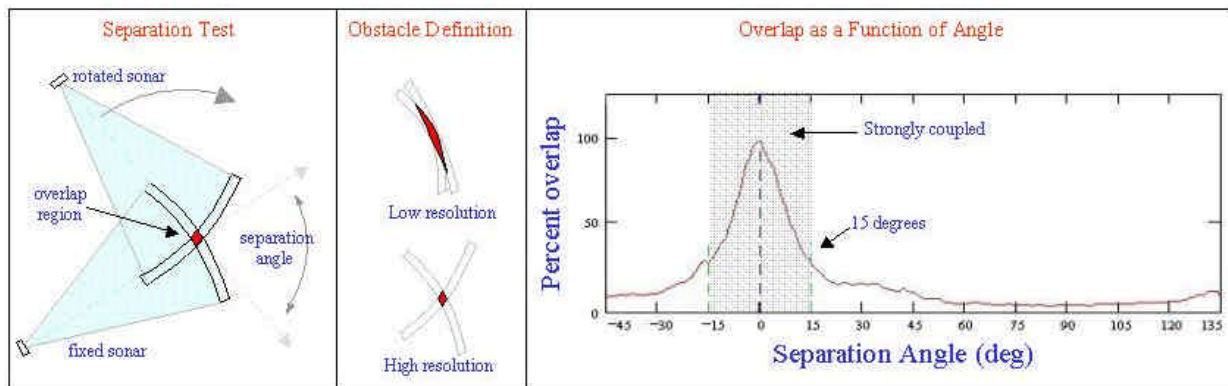


Figure 3: Assessing Pose Separation

(a) A separation test is conducted by looking at the fusion of two sonar readings impinging on the same point. One reading is fixed; the other is rotated about the point under test. (b) The resulting obstacle definition of a point from two overlapping sensors, one at 10 degrees, the other at 90. (c) A plot of the overlap of distributions between the two sensor arcs as a function of angle. The overlap is greatest at 0 degrees and quickly falls of after 15 deg.

derivation is based on evaluating the increase in effective azimuth resolution for a sonar sensor. These results are consistent with our work.

However, our work goes beyond developing a critical angle. We expose the sensitivity of dependence as a function of angle. For example, our tests show the utility of pose separation for readings that are separated by 30 degrees is on the same order as poses separated by 90 degrees. Furthermore, whereas we currently only exploit the critical angle, work is underway to develop a continuous utility function that better captures the dependence on angle.

Utilizing this relationship, we develop an inverse sensor model that is composed of three regions (Figure 2b). The region along the same bearing as the original reading represents measurements that would be strongly coupled to each other. As such, this region has negative utility and make up the suppression part of the model. Given a choice, we would not take new readings in this area. Regions to either side that differ in bearing from the original are more independent in terms of uniquely defining an obstacle. These regions have a positive utility and make up the separability part of the model. New sensor readings from these areas are desirable. Together, these two readings define the inverse sensor model that represents the next best view with respect to increasing the resolution of that obstacle point (Figure 2c).

In practice, we make a few adjustments to the model. First, we limit the range of the model to equal the range of the sensor. Second, we limit the angular width of the sensor model to be on the order of 90 degrees on either side. Limiting the angular range of the model makes the

update of the map tractable. Our future work will examine the optimum coverage of the model. Furthermore, it recognizes the limitations in pose separation imposed by the incident angle of specular reflection. Third, since the same inverse model applies equally to all the obstacle points for each sensor reading, we can reduce the calculations by defining a composite, inverse sensor model that batch processes all the corresponding cells for a given sensor reading (Figure 4b).

7 Regions of Interest

The inverse sensor model defines a *next best view* methodology for defining optimum places for new sensor readings with respect to individual readings. However, we can extend this methodology to generate a composite map that represents the overall utility for moving the robot to new positions with respect to map resolution. We take multiple inverse sensor models and fuse them together in much the same fashion as the occupancy map (Figure 4b). The result of this fusion is the development of a **Regions of Interest** map (Figure 4c). Areas where multiple separability regions of individual sensor models overlap have a higher utility than areas derived from a single occurrence. The result of fusion produces regions of interest that are not only on the frontiers of exploration but are also designed to improve local obstacle definition in the process.

We can see the utility of this approach if we compare the exploration potential against an instance of the free space approach. The standard free-space approach is based on extracting frontiers from an occupancy map by evaluating

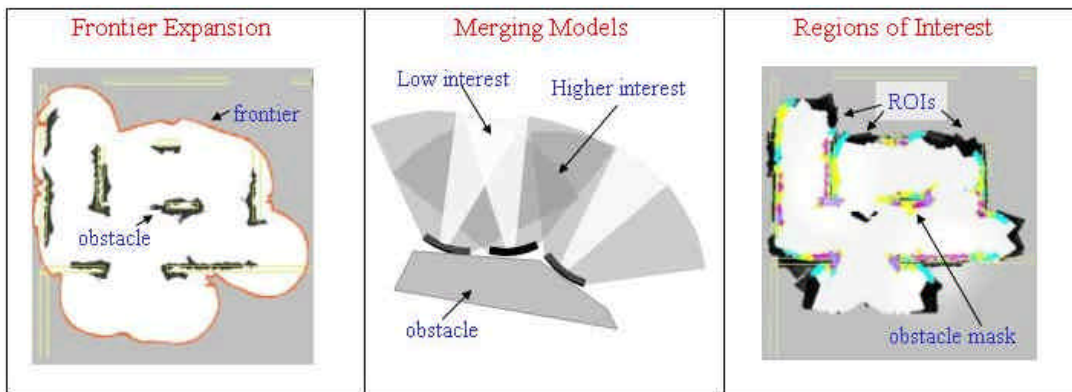


Figure 4: Developing an Inverse Sensor Model

(a) The standard method for exploration is based on extracting frontiers from an occupancy map (boundary between open and unexplored). Specular reflection erroneously exposes regions beyond local obstacles. In this case, there is essentially one frontier that extends around the entire perimeter of the map. (b) Multiple inverse sensor models can be fused together in much the same fashion as an occupancy map to build higher utility for measurement. (c) Merging multiple regions of interest into a common map provide an alternate means for generating exploration plans. Unexplored regions are gray. Light areas are open space. Dark regions are regions of interest. The darker the value, the more utility it has for resolving local obstacles.

boundaries between open and unexplored space. However, the effects of specular reflection generate phantom frontiers that can erroneously expose large regions beyond local obstacles. The result of this experiment was severe enough to produce a single frontier that extends around the entire perimeter of the map as several phantom frontiers overlap (Figure 4a). From a navigation perspective, this bulk frontier does little to guide the robot. Certainly given such a large frontier, the robot could simply select the closest part of the frontier and move towards it. However, since the frontier is actually based on false readings, this is not an effective means for navigation.

Conversely, the same sensor readings were utilized to generate the *Regions of Interest* map in Figure 4c. Again, the map was produced by merging each of the individual inverse sensor models into a global map. Regions in white indicate low utility for new sensing. Regions in black have a higher utility for new sensing and indicate the best regions to explore. In this case, guiding the robot towards the regions of interest produces a more efficient plan for exploration that is not susceptible to the effects of specular reflection.

8 Exploration and Future work

Currently we adopt a simple method for generating the *next best view* for the robot based on the Regions of Interest map. We attempt to find the closest region of interest with the highest utility for sensing. Utilizing the pose map, we generate a number of random exploration positions about the robot and keep only those that fall in a region of interest. From there, we select the closest region of interest and plan a path through the occupancy map to guide the robot.

Given the metrics developed for the inverse sensor model that indicate the utility of sensing from a new position, we are working on developing a recursive, maximum likelihood estimator for exploration. We hope to develop a system that optimizes on both sensing and movement during exploration. Our goal is to execute the most efficient plan that minimizes the movement and sensing of the robots while maximizing their impact. This is essential for our work. We are targeting our research towards a team of small, heterogeneous, resource-limited robots called Millibots [6]. Small robots by their very nature are more limited in sensing, processing and mobility. As such, efficient and effective exploration is a key to their success.

9 Conclusion

In this paper we have presented a new approach to exploration based on the perspective of the obstacles that

make up the world. To support this approach, we have developed a means for measuring the quality of a robot's map with respect to obstacle resolution by assessing the degree of pose separation of the readings used to define that cell. We have applied this information to build an inverse sensor model that identifies regions in space where a new sensor reading has maximal utility with respect to increasing the resolution of that reading. We exploited the fusion of multiple models to generate a *Regions of Interest* map that direct exploration in such a way as to maximize the robots understanding of the entire space.

10 References

- [1] Banos, H., Mao, E., Latombe, J.C., Murali, T.M., and Efrat A., "Planning Robot Motion Strategies for Efficient Model Construction." Robotics Research -- The 9th Int. Symposium, J.M. Hollerbach and D.E. Koditschek (eds.), Springer, pp. 345-352, 2000.
- [2] Banta, J., "A 'Best-Next-View' Algorithm for Three Dimensional Scene Reconstruction Using Range Images,"
- [3] Burgard, W., Fox, D., Moors, M., Simmons, R., and Thrun, S., "Collaborative Multi-Robot Exploration," In Proceedings of the IEEE International Conference on Robotics and Automation, San Francisco CA, April 2000..
- [4] Choset, H., Nagatani, K., Lazar, N. "The Arc-Traversal Median Algorithm: an Approach to Increasing Ultrasonic Sensor Accuracy," IEEE/ICRA Detroit, 1999.
- [5] Elfes, A., "Using Occupancy Grids for Mobile Robot Perception and Navigation," IEEE Computer, pages 46--57, June 1989.
- [6] Grabowski, R., Navarro-Serment, L. E., Paredis, C.J.J., and Khosla, P. "Heterogeneous Teams of Modular Robots for Mapping and Exploration," Autonomous Robots - Special Issue on Heterogeneous Multirobot Systems.
- [7] Pito, R. "A Sensor-based Solution to the "Next Best View" Problem." Proceedings of 13th International Conference on Pattern Recognition. Held: Vienna, Austria, 25-29 Aug. 1996. (USA: IEEE Comput. Soc. Press, 1996. p. 941-5 vol.1)
- [8] Simmons, R., Apfelbaum, D., Burgard, W., Fox, D., Moors, M., Thrun S., and Younes H., "Coordination for Multi-Robot Exploration and Mapping." In *Proceedings of the National Conference on Artificial Intelligence*, Austin TX, August 2000.
- [9] Yamauchi, B., "Frontier-Based Exploration Using Multiple Robots," In Proc. of the Second International Conference on Autonomous Agents (Agents'98), Minneapolis, MN., 1998.