Proceedings of the 2003 IEEE/RSJ
Intl. Conference on Intelligent Robots and Systems
Las Vegas, Nevada · October 2003

# An Enhanced Occupancy Map for Exploration via Pose Separation

Robert Grabowski, Pradeep Khosla and Howie Choset

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

***Abstract***

*We develop a new occupancy map that respects the role of the sensor measurement bearing and how it relates to the resolution of the existing occupancy map. We borrow an idea from Konolige for recording and tracking, in an occupancy-like map, the bearing at which sensor readings originate with respect to a given cell. Our specific contribution is in the way we process the sensor pose information, which is the bearing of the sensor readings when it indicates the presence of an obstacle in a particular cell. For each cell in the occupancy map, we calculate the greatest separation of incident poses, and then store that information in a new two-dimensional array called a pose map. A cell in the pose map measures the quality of information contained in the corresponding cell of the occupancy map. We merge the new pose map with the existing map to generate an enhanced occupancy map. Exploration plans derived from the enhanced occupancy map are more efficient and complete in that they do not guide the robot around phantom obstacles nor incorrectly classify narrow openings as closed commonly found in conventional occupancy maps.*

## 1 Introduction

Robots are increasingly showing potential for support of human operations such as surveillance, reconnaissance and rescue. Robots can provide remote, distributed, real-time sensing and actuation while reducing the risks inherent to the human operators. However, while robots may be tasked to work in close conjunction with an operator, a degree of autonomy is necessary to free the operator from direct supervision and coordination. Robots must be able to operate in an unstructured environment with little or no previous information. As such, one of the core competencies will be the ability to execute autonomous exploration. Autonomous exploration is a recursive process that utilizes the relationship between sensing and movement of the robots through a map. Robots collect sensor information about their general surroundings, and coupled with the knowledge of robot position, fuse this information to build a global map.

In this paper, we develop a framework for assessing the quality of map information with respect to its ability to represent underlying obstacles. This is essential because planners determine their motion from map information. However, the occupancy map is susceptible to the effects of conditionally dependent sensor readings - especially at the fringes of exploration.

**Figure 1: A team of Millibots exploring with sonar sensors**

To consider this dependency, we define a new metric that quantifies the role of both range and the bearing of the sensor and use this metric when fusing multiple readings in an occupancy map. For each cell in the occupancy map we record the angle of the sensors which indicated that the cell has an obstacle. We then compute a scalar value that is based on the greatest separation of angle. We define that scalar value in a new map called the pose map. A cell in the pose map measures the quality of information of the corresponding cell in the occupancy map. We merge the pose information with the existing map to generate an enhanced occupancy map that provides better discriminative ability with respect to exploration. Ultimately, the success of exploration depends on the quality of the coupling between sensing, map building and navigation.

## 2 Occupancy Maps

### 2.1 Definition

The most common mapping technique is based on the occupancy map. One of the strengths of the occupancy map is its ability to fuse multiple range readings of many types into a globally consistent notion of occupancy [4]. In an occupancy map, 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 and ranges in value between zero and 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 is initialized at 0.5 (equally likely to be occupied or free).

Any sensor that produces a range reading and defines a sensor model can be utilized in the map. Given a range reading and its position in space, a sensor model is used to infer the occupancy of that region of space. 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. 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 [4].

### 2.2 Exploration

Exploration is accomplished by extracting relevant features from the occupancy map that indicate viable areas to search. The most informative feature is the frontier. Frontiers are commonly defined as regions between open, explored space and unexplored space and represent the boundaries of known space. Once a frontier has been selected, he utilizes the occupancy map to find the shortest, obstacle-free path to the target frontier. Yamauchi conducts exploration by extracting frontiers directly from the occupancy map to guide the motion of the robot [10]. His robot selects the centroid of the closest frontier and then relies on local obstacle avoidance to navigate towards the frontier. Simmons extends this form of exploration to select an appropriate frontier that will both maximize the exposure of new space and minimize robot travel [8]. This added utility allows for multiple robots to work in the same space while minimizing overlap.

Banos describes a method of extracting polylines from the occupancy map to be utilized in motion planning [1]. One set of polylines represents the profiles of obstacles. The other set of polylines represent frontiers. He utilizes the polylines to geometrically assess a *next best view* that will both maximize exposure of unexplored space while maintaining a minimum degree of contact with existing obstacles for the purposes of localization. Moreover, he utilizes polylines extracted from the map to infer "safe regions" of the map that take into account for the specular properties of his sensors.

In each case, information is extracted directly from features (such as lines, walls, doors and edges) extracted from the occupancy map but does not take into account for the quality of that information. Specifically, none of these methods address the confidence in the ability of the occupancy map to represent an underlying obstacle.

### 2.3 Conditional Independence

A major shortcoming of the occupancy map with respect to exploration arises from the commonly applied assumption of conditional independence between the sensor readings and the state of the cells that make up the map. This conditional independence argument is critical in the application of real-time, sensor-based mapping. To completely account for the dependence between multiple readings would require the calculation of joint sensor models and normalization of readings over large regions of space. On the other hand, if the assumption of independence is allowed, the calculations are reduced from a complex mixture of joint probabilities to the product of individual probabilities. This assumption has proven acceptable in most mapping applications where the robot is expected to collect large amounts of data as it moves around the space. However, the utility of the assumption quickly breaks down at the frontiers of autonomous exploration where there are fewer sensor readings and the majority of these readings are taken from the same relative bearing.

We can see the adverse effects from a simple example. Assume a series of sensor readings are taken about a point obstacle from the same general distance and bearing. Clearly little information is gained beyond the first few readings. However, if we assume the readings are conditionally independent, then the update rule fuses all the readings without discrimination. The result is the underlying point obstacle is represented in the occupancy map by a large cluster of cells distributed about that point. Clearly if the robot where to then formulate a plan for exploration based on the current map, it may direct its movements around a large obstacle that simply does not exist. Even worse, if the effect occurs around a narrow opening (such as a door or window), it may erroneously indicate the opening as closed. Figure 2 shows an example of a team of robots mapping a maze by navigating from local sensors. Since the conditional independence argument was assumed without correction, the resulting map closed off several openings. Consequently, the inability to recognize and discriminate the quality of map information leads to ill-posed search solutions.

## 3 Pose Separation as Metric for Map Quality

We propose an enhancement to the occupancy map that provides a measure of map quality to each point in the occupancy map by evaluating the separation of the poses that derive that cell. The greater the separation in the poses that make up a cell, the more confidence we give to that cell with respect to feature extraction in exploration. Separation of pose is a strong indicator in defining the resolution of an obstacle point (a cell with higher occupancy values).

The role of pose information is not new to mapping. Konolidge exploits pose information to improve the map building process in his MURIEL method [7]. For each cell in the occupancy map, he defines a structure called "pose buckets" that discretizes the possible poses incident on that point into an array of 64 angles and 3 ranges. When a range reading is taken, its pose is tested against the corresponding pose bucket in the array. If the reading is unique, it is added to the occupancy map and the pose bucket is filled. If a reading is already present in that bucket, the new reading is discarded. This method of reading rejection addresses the conditional dependence argument of multiple sensor readings taken in close proximity to one another.

Rejection of redundant readings allows the fast calculation of occupancy while reducing the complication of strongly dependent readings. However, pose rejection via discrete buckets has its drawbacks. First the discriminative abilities of this approach are a function of the resolution of the array. A coarsely discritized array provides good discriminative qualities (less coupling between readings) but reduces the resolution of the map. A finely discritized array improves map resolution but reduces the ability to discriminate between two closely coupled readings.

More importantly this method does not fully leverage the power of pose to discriminate points in the occupancy map. That is, we look at the degree of separation between readings and its impact on the resolution ability of obstacles. We look at assessment of the readings from a different perspective. We develop a metric for assessing the resolution ability of merging two or more sensor readings in an occupancy map as a function of separation in pose (Figure 3). To make this assessment, we look at the fusion of two sensor readings impinging on the same point in space. If we consider the readings to be conditionally independent, then the combined result is the

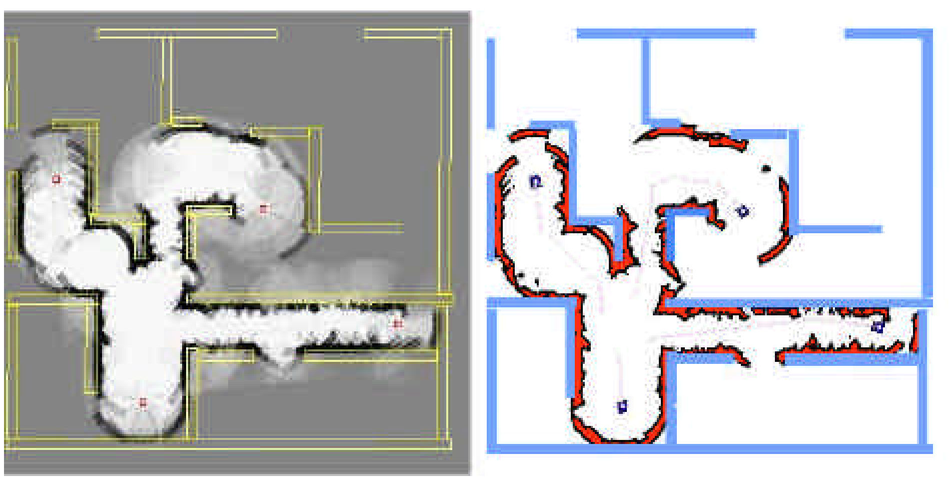

**Figure 2: Conditional Dependence and Exploration.**

**a) Raw occupancy map generated by a team of Millibots using only local navigation strategies to move throughout a maze. b) Exploration map generated by thresholding occupancy map.**

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. Figure 3b shows the definition of a point developed from fusing two sensor readings. When the overlap is large (small separation angle), the resolution of the underlying point is poor. Conversely, when the overlap is small (large separation angle), 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 the two sensors 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 [3]. 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 derivation is based on evaluating the increase in effective azimuth resolution for a sonar sensor. These results are consistent with our work.

## 4 Enhanced Occupancy Map

We exploit the concept of pose separation to assess the quality of each occupancy point with respect to its ability to represent an underlying obstacle. We add a bearing array to each occupancy cell that discretizes the possible poses incident on that cell (Figure 4a). We use the structure as a placeholder array that tracks and regulates the number of unique sensor readings used to make up an occupancy point. Moreover, the array acts as a filter for rejecting duplicate measurements - that is, readings that would be stored in the same array bin. However, rather than simply rejecting similar pose readings, we exploit the pose information to generate a separate pose map that represents the distinctedness of its associated occupancy point (Figure 4c). The value of each pose cell is derived by measuring the greatest degree of separation of the existing readings in the pose array. For example if a cell is built up from two readings taken at 10 and 45 degrees, then the separation for that cell is 35 degree. Conversely, if a cell is built up from 3 readings of 10, 30 and 50 degrees, then the separation for that cell is 40 degrees. Currently we only account for the maximum separation between cells. However, we are already seeing good results. Our future work will explore a weighted separation to account for the utility of readings clustered between the maximums as well.

By applying the metric of pose separation to each occupancy cell, we build a composite pose map (Figure 4c). Each cell of the pose map represents the diversity of the information used to build the corresponding occupancy point and hence the quality of that point with respect to feature extraction. Moreover, the quality of the

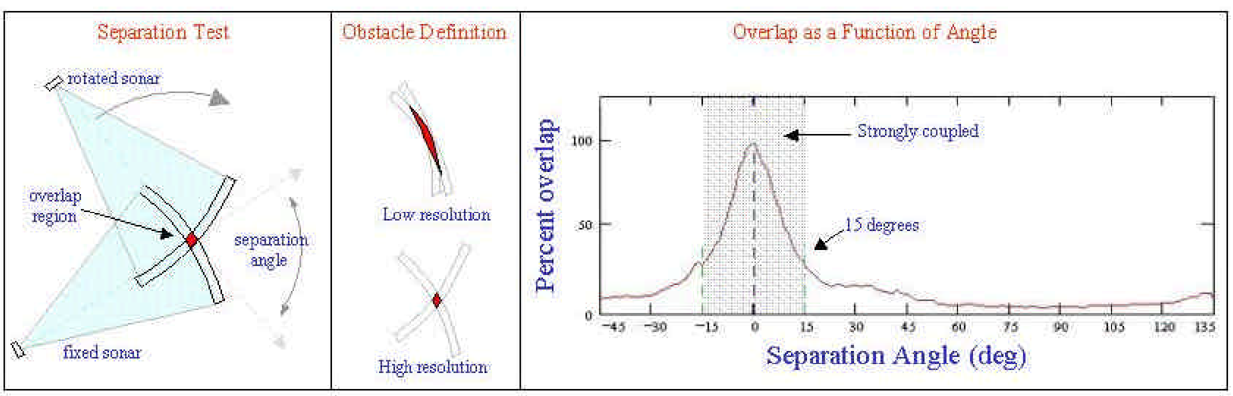


**Figure 3: Assessing Pose Separation**

**(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 source 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.**

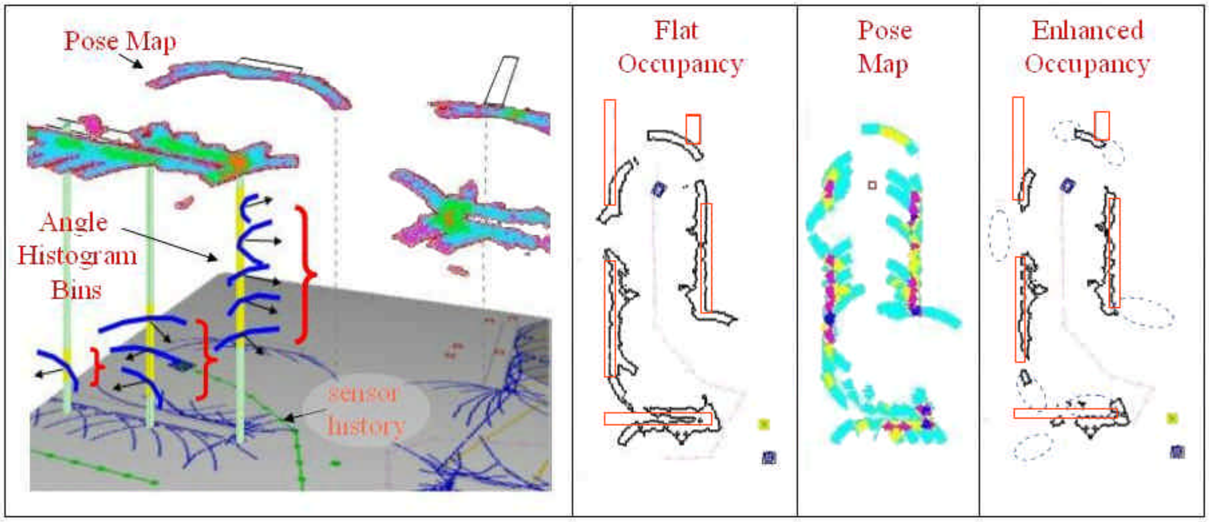


**Figure 4: The Pose Map for Assessing Quality of an Occupancy Map**

**(a) An angle array is added to each cell of an occupancy map that tracks the number of readings taken from a particular angle and is used to measure the greatest separation in pose. (b) Obstacles extracted from an occupancy map generated from a robot passing through a hallway. (c) The pose map generated from the same operation - the darker the value, the greater number of distinct poses. (d) Obstacles extracted from an enhanced occupancy map generated by combining the flat notion of occupancy with the pose map. Note the openings in the hallway are less obstructed. The areas rejected by the combination also give a hint as to good areas to explore.**

occupancy point is now a function of the degree of separation of poses and not the discretization of the space.

Now that we have a measure of occupancy quality, we can generate an enhanced occupancy map that provides better feature extraction with respect to exploration. The enhanced occupancy map is constructed by merging the existing occupancy map with the new pose map in the following fashion. First we define a filter threshold based on the discriminative qualities of pose separation (on the order of 15 degrees for sonar sensors). If the pose value for that cell exceeds the threshold, the original occupancy point is assigned. Consequently, it is utilized in feature extraction for exploration planning. If the pose map value is below the threshold, that point is devalued and essentially ignored in the extraction process. The result is an enhanced occupancy map is more discriminative at the fringes of exploration.

Figure 4 shows the result of a typical mapping experiment. In Figure 4b, sonar readings have been fused without any regard to pose quality. For the most part, the closed space represents the underlying obstacles to a reasonable degree. However, in the corners of the map and at the fringes of exploration, obstacle definition is not as clear. These are places where a majority of the readings are close to each other in proximity. As a result they are more strongly coupled and generate poor map features. In some cases, the map indicates openings in the environment as closed. As such, the robot may never go back and further explore that area. This feature is essential during exploration so that the robot does not inadvertently reject possible exploration paths based on poor obstacle definition. Figure 4c gives an indication of the quality of these obstacle points in the form of a pose map. It clearly shows that the edges of obstacles, where readings are more widely distributed, have a higher resolution capability. Conversely, it also shows the low utility to readings at the fringes of exploration. Figure 4d shows the combination of the pose map and the occupancy map. This map places less confidence in the readings at the fringes. As a result, the areas once closed off are now open. Moreover, the obstacles at the forefront of exploration are weighted less. Plans based on these features will not unnecessarily guide the robot around obstacles that do not exist.

## 5 Precursor to Exploration

The enhanced occupancy map has a utility beyond producing better features with respect to exploration. If we take a closer look at the regions masked out when the occupancy map is merged with the pose map, we get a good indicator as to potentially viable areas to explore. For example, in Figure 4d, the dotted circles indicate the regions filtered out in the enhanced occupancy map. These correspond to regions in the occupancy map that did not have sufficient pose separation and were deemed unfit for feature extraction. However, they still retain some useful information in that they indicate the potential presence of an obstacle. In this sense, they are precursor to areas that should be further explored. This fits in well with the notion of obstacle-based exploration. In obstacle-based exploration, we direct the robot in such a way as to improve the understanding of the obstacles that surround it. In another IROS paper, we develop a *next best view* methodology that generates new positions for sensing that will not only expose new space but concurrently increase the resolution of existing obstacles [6]. Together with the feature resolution utility of the enhanced occupancy map, this obstacle-based, next best view approach increases the effectiveness of autonomous exploration.

## 6 Conclusion

In this paper we have presented a new means for evaluating range data to generate maps for the purpose of autonomous exploration. To support this approach, we have developed a means for measuring the quality of a team'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 a pose map that represents the distinctedness for each occupancy cell. In turn we have utilized the pose map to generate an enhanced occupancy map that is better suited for use in feature extraction in the exploration process. While pertinent to all robot platforms that utilize range sensors, this methodology is especially important when working with small, heterogeneous, resource-limited robot teams [5]. Small robots by their very nature are limited in processing, sensing and mobility. It is essential to be efficient as possible in every aspect of operation including sensing and movement. Small robots cannot afford to be redirected around obstacles that do not exist. Moreover, the effects of phenomenon, such as specular reflection and conditional dependence are magnified making it essential to get as much as possible out of existing information.

## 7 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] 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.

[3] Choset, H., Nagatani, K., Lazar, N. "The Arc-Traversal Median Algorithm: an Approach to Increasing Ultrasonic Sensor Accuracy," IEEE/ICRA Detroit, 1999.

[4] Elfes, A., "Using Occupancy Grids for Mobile Robot Perception and Navigation," IEEE Computer, pages 46--57, June 1989.

[5] 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.

[6] Grabowski, R., Khosla, P., Choset, H., "Autonomous Exploration via Regions of Interest," to be published in IROS 2003.

[7] Konolige, K., "Improved Occupancy Grids for Map Building," 1997, *Autonomous Robot*s, vol. 4, 351–367.

[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 Intelligenc*e, Austin TX, August 2000.

[9] Thrun, S., "Learning Metric-topological Maps for Indoor Mobile Robot Navigation," Artificial Intelligence, 99(1):21-71, 1998.

[10] Yamauchi, B., "Frontier-Based Exploration Using Multiple Robots," In Proc. of the Second International Conference on Autonomous Agents (Agents`98), Minneapolis, MN., 1998.