Generating Sonar Maps in Highly Specular Environments

Andrew Howard, Les Kitchen

Computer Vision and Machine Intelligence Laboratory
Department of Computer Science
University of Melbourne, Victoria 3053
Australia
andrbh@cs.mu.OZ.AU, ljk@cs.mu.OZ.AU

Abstract

We address the problem of building environment maps from ultrasonic range data obtained from multiple viewpoints. We present a novel environment modelling technique called the 'response grid' that allows us to build occupancy maps in highly specular environments. We present two different approaches that utilise this technique: a Bayesian probabilisitic approach and a Dempster-Shafer evidential reasoning approach. Both approaches can be implemented in real-time with modest computational resources, and as such are suitable for use in mobile robot navigation tasks. We present and compare the experimental results obtained by these methods in a highly specular indoor environment.

1 Introduction

Building environment maps from sensory data is an important aspect of mobile robot navigation, particularly for those applications in which robots must function in unstructured environments. Ultrasonic range sensors are, superficially, an attractive sensor modality to use in building such maps, due mainly to their low cost, high speed and simple output. Unfortunately, these sensors have a number of properties that make map building a non-trivial process. In particular, standard sensors have very poor angular resolution and can generate misleading range values in specular environments. The first of these problems can be largely overcome by combining range measurements from multiple viewpoints. Elfes [1] and Moravec [2] describe an approach in which range measurements from multiple viewpoints are combined in a two-dimensional 'occupancy grid'. Each cell in the grid is assigned a value indicating the probability that the cell is occupied. Unfortunately, the occupancy grid approach does not work well in specular environments. Specular reflection may occur whenever an ultrasonic pulse encounters a smooth extended surface [3]. In such cases the pulse may not be reflected back to the ultrasonic sensor; in effect, the surface may appear to be invisible. In ordinary office environments which contain smooth walls and glass doors specular reflection is common. In this paper, we improve on earlier grid-based approaches by introducing the concept of a 'response grid'. The intent of the response grid framework is to produce an approach which has the advantages of the occupancy grid framework, but also performs well in specular environments.

The response grid framework attempts to model the behaviour of ultrasonic range sensors in a more physically realistic fashion. A number of other authors have considered the physical behaviour of such sensors in some detail [4, 5, 6]; in this paper, however, we are only conerned with physical behaviour insofar as it allows us to generate two dimensional occupancy maps in specular environments. The basic notion encapsulated by the response grid is that a cell may generate a response (i.e. appear to be occupied) when viewed from one direction, but will not generate a response when viewed from another. In the original occupancy map framework, this would present a contradiction, since this approach assumes that an occupied cell should generate responses in *every* direction.

In this paper, we present two different methods for generating occupancy maps from ultrasonic range data within the response grid framework. The difference between the two methods lies in the different techniques they use to combine data obtained from multiple viewpoints. We present and compare a Bayesian probabilistic reasoning approach with a Dempster-Shafer evidential reasoning approach. The Bayesian approach can be viewed, to some extent, as a generalisation of the approaches described by Elfes [1] and by Moravec [2]. In

Section 3 we compare the experimental results for each method and consider their relative advantages. Note that both methods can be implemented in real-time with modest computational resources, which makes them well suited to mobile robot navigation tasks.

2 Framework

2.1 The Response Grid

We model the environment as a set of cells arranged in a regular two-dimensional grid. A given cell in the grid may either be entirely empty, or else it may contain one or more surfaces which reflect ultrasonic pulses. A pulse entering a cell will do one of three things:

- If the cell is entirely empty, the pulse will pass through the cell unaffected.
- If the cell contains one or more reflective surfaces, the pulse may be reflected back to the ultrasonic detector.
- If the cell contains one or more reflective surfaces, the pulse may be reflected away from the ultrasonic detector.

The behaviour of a pulse reflecting off a surface contained in the cell will depend upon the orientation of the surface and the direction of the incoming pulse. If a pulse which is propagating in some given direction is reflected back to the detector, the cell is said to have a *response* in that direction. In this paper, we determine the *occupancy* of each cell by assuming that any cell that generates a response in one or more directions must contain at least one surface and is therefore *occupied*.

The model described above can be expressed mathematically as follows. The occupancy of a cell at location (x, y) is measured by the state variable Occ which can have one of two values:

$$\mathrm{Occ}(x, y) = [\mathit{occupied}, \mathit{unoccupied}]. \tag{1}$$

The response of a cell in some direction ϕ is measured by the state variable Res:

$$\mathrm{Res}(x, y, \phi) = [\mathit{response}, \mathit{no\ response}]. \tag{2}$$

The direction ϕ is allowed to take discrete values between 1 and n. The two state variables are bound together by a logical implication. In order to express this clearly, we define a set of propositions:

$$\begin{aligned} O &: \mathrm{Occ}(x, y) = \mathit{occupied} \\ R_\phi &: \mathrm{Res}(x, y, \phi) = \mathit{response}. \end{aligned} \tag{3}$$

That is, the proposition O states that the cell at (x, y) is occupied and the proposition R_ϕ states that the cell at (x, y) generates a response in direction ϕ. We can therefore write the implication:

$$O \Leftarrow R_1 \vee R_2 \cdots \vee R_{n-1} \vee R_n. \tag{4}$$

In this paper we present two methods for using ultrasonic range data to determine the response properties of each cell, and thence to determine occupancy.

2.2 Method 1: Bayesian

The objective of the Bayesian method is to determine, for each cell, the *probability* that the cell is occupied. That is, for a cell at (x, y) we wish to determine the probability that the proposition O is true. Rather than attempting to calculate this probability directly from range data, we compute it indirectly. Since O is related to the cell responses by the logical implication expressed in Equation 4, the probability that the cell is occupied must be given by:

$$p(O) = p(R_1 \vee R_2 \cdots \vee R_{n-1} \vee R_n). \tag{5}$$

To expand the right hand side of this equation, we make use of the observation that for two independent propositions A and B, the probability of either A or B being true is given by [7]

$$p(A \vee B) = p(A) + p(B) - p(A)p(B). \tag{6}$$

Applying this to the above equation, one obtains after a little algebra:

$$p(O) = 1 - \Pi_\phi(1 - p(R_\phi)). \tag{7}$$

This expression can be used to compute the probability that a cell is occupied once we have determined the cell response probabilities. Note that this involves the determination of n *separate* probabilities, since $\{R_1 \cdots R_n\}$ is a set of *independent* propositions.

Consider now just one proposition R_ϕ corresponding to the response in direction ϕ. We can apply Bayes' rule to determine the probability that R_ϕ is true, given a range measurement r:

$$p(R_\phi \mid r) = \frac{p(r \mid R_\phi)p(R_\phi)}{p(r)}. \tag{8}$$

In this expression $p(R_\phi)$ is the *prior probability* of obtaining a response. In Bayesian approaches it is very common to set the prior probability to 0.5 to indicate no opinion. In this case we demand that $p(O) = 0.5$, which implies that the prior probabilities for $p(R_\phi)$ are given by:

$$p(R_\phi) = 1 - (0.5)^{\frac{1}{n}}. \tag{9}$$

The other important term in Equation 8 is the *sensor model* $p(r \mid R_\phi)$. The sensor model indicates the probability of obtaining the measurement r, given that the proposition R_ϕ is true. Let s be the distance between the cell and the sensor; the sensor model we use is:

$$p(r \mid R_\phi) = \begin{cases} 0.05 & \text{if } s < r \\ 0.5 & \text{if } s > r \\ \alpha/r & \text{if } s = r \end{cases} . \qquad (10)$$

The first two parts of this model are self-explanatory, but the third requires some explanation. When the sensor records a range r we know that *something* at this range has generated a response. Due to the finite angular resolution of the detector, however, there may be a number of cells that could have generated this response. In general, the number of such cells will be proportional to the measured range, so the probability that any individual cell generated the response is *inversely proportional* to the measured range. In effect, we give more weight to short range measurements than to long range measurements. This rule is both physically plausible and intuitively appealing. The α that appears in the above rule is a normalisation constant: summing over the probabilities assigned to each cell at range r should yield a total probability of 1. The value of α is determined by the spatial dimensions of the cells making up the map and by the angular resolution of the ultrasonic sensor.

The model we use ignores statistical errors associated with the range value itself; that is, our model is *ideal*. We justify this simplification by noting that for common types of ultrasonic sensors such errors are of the order of one or two centimetres. In practical applications, such as mobile robotics, these errors are insignificant compared to errors arising from other sources, such as robot (mis)localisation.

2.3 Method 2: Dempster-Shafer

The Dempster-Shafer Theory of Evidence [8] can be applied to our problem in a straight-forward fashion. The objective is to determine the *support* for the proposition O. From the logical implication expressed in Equation 4, we can write the following:

$$\mathrm{Sup}(O) = \mathrm{Sup}(R_1 \vee R_2 \cdots \vee R_{n-1} \vee R_n). \qquad (11)$$

To expand the right hand side of this equation, we use the Dempster-Shafer analogue of Equation 7 . Given two independent propositions A and B, the support for the combined proposition $A \vee B$ is given by [9]:

$$\mathrm{Sup}(A \vee B) = \mathrm{Sup}(A) + \mathrm{Sup}(B) - \mathrm{Sup}(A)\mathrm{Sup}(B). \quad (12)$$

Applying this to the expression above, one obtains:

$$\mathrm{Sup}(O) = 1 - \Pi_\phi(1 - \mathrm{Sup}(R_\phi)), \qquad (13)$$

which allows us to calculate the support for the proposition that a cell is occupied, once we have determined the support for each of the propositions $\{R_1 \cdots R_n\}$.

Consider now just one possible response direction. In order to determine the support for the proposition that the cell responds in this direction, the Dempster-Shafer approach requires that we construct a relevant *frame of discernment*, which is a set which contains all the propositions of interest. In our case, the frame of discernment contains just two propositions: that the cell responds, or that it does not:

$$\Theta = \{R_\phi, \neg R_\phi\}. \qquad (14)$$

In general, support for these propositions cannot be determined directly from the available evidence. Instead, support is computed indirectly via the *mass distribution*. The mass distribution allows us to allocate a 'weight' to any element in the frame of discernment, or to any proper subset of the frame of discernment. Thus in our application, the mass distribution can assign weight to any element in the set:

$$2^\Theta = \{R_\phi, \neg R_\phi, R_\phi \vee \neg R_\phi\}. \qquad (15)$$

The weight assigned to an element in this set indicates the amount of evidence that supports that element directly. Support for propositions in the frame of discernment can be calculated according to the rule:

$$\mathrm{Sup}(A) = \sum_{A \subseteq B} m(B). \qquad (16)$$

Thus the support for the proposition R_ϕ is trivially equal to $m(R_\phi)$. The mass distribution is somewhat similar to a probability distribution in that it must be non-negative and must sum to one. Consider the following mass distribution:

$$\begin{aligned} m(R_\phi) &= 0 \\ m(\neg R_\phi) &= 0 \\ m(R_\phi \vee \neg R_\phi) &= 1. \end{aligned} \qquad (17)$$

For brevity we write this as $(0, 0, 1)$. This distribution corresponds to complete ignorance, since it will yield no support for either R_ϕ or its negation. Contrast this with the distribution $(1, 0, 0)$, which indicates complete support for R_ϕ and no support for its negation; and with $(0.5, 0.5, 0)$, which indicates a contradiction - there is equal support for R_ϕ and its negation. One of the attractive features of the Dempster-Shafer approach is the way in which the ignorance and contradiction are clearly distinguished; this is not true of the Bayesian approach.

In our problem, every new measurement arriving from the ultrasonic sensor is treated as a new piece of evidence for which we must generate a mass distribution. These separate distributions must then be combined to generate a collective opinion. If $m(\cdot \mid r_1)$ and $m(\cdot \mid r_2)$ represent mass distributions arising from independent measurements, the combined mass distribution $m(\cdot \mid r_1, r_2)$ is given by Dempster's rule [8]:

$$m(C \mid r_1, r_2) = \frac{1}{1-\kappa} \sum_{A \cap B = C} m(A \mid r_1) m(B \mid r_2), \tag{18}$$

where κ is defined as

$$\kappa = \sum_{A \cap B = \emptyset} m(A \mid r_1) m(B \mid r_2), \tag{19}$$

and A, B and C are elements of the set $\{R_\phi, \neg R_\phi, R_\phi \vee \neg R_\phi\}$. Inspecting the above equation, one can see that when a contradiction exists (for example when there is support for both a proposition and its negation), κ will be non-zero. As a result κ is usually interpreted as indicating the degree of contradiction between the two mass distributions. Dempster's rule can be thought of as Dempster-Shafer equivalent of Bayes' rule.

The rule we use for generating mass distributions from range measurements is analogous to the sensor model used in the Bayesian case. Let $m(\cdot \mid r)$ be the mass distribution resulting from the measurement r; let s be the distance between the cell and the sensor; the sensor model we use is:

$$m(\cdot \mid r) = \begin{cases} (0, 0.95, 0.05) & \text{if } s < r \\ (0, 0, 1) & \text{if } s > r \\ (\alpha/r, 0, 1-\alpha/r) & \text{if } s = r \end{cases} . \tag{20}$$

The reasoning used to arrive at this model is identical to that used in the Bayesian case. We note only that the initial, or prior, mass distribution we use corresponds to complete ignorance – $(0, 0, 1)$.

3 Implementation and Experimental Results

The implementation of the two methods described in the preceeding sections is remarkably similar. The map is represented by a two-dimensional array, with each cell corresponding to a region of the environment. Each cell has associated with it a value which represents the cell's occupancy, and an array of values representing the cell's responses. The meaning of these values will depend upon the method being used. When a measurement is taken, the first step is to determine which of the cells in the array should be updated to reflect this measurement. If we imagine that the pulse emitted by the ultrasonic sensor propagates outward through a conical region of space, then only the cells corresponding to this region of space should be updated. For the Polaroid sensors used in our experiments, this will be all cells within about $\pm 10°$ of the center line of the sensor. Furthermore, for each of the candidate cells, only one of the responses needs to be updated – that corresponding to the direction of the emitted pulse. Once this determination has been made, the appropriate cell response can be updated and the overall cell occupancy recomputed.

The results presented in this section where obtained using a small mobile robot with a single Texas Instraments/Polaroid ultrasonic range-finder attached to a pivoting head. The sensor has an unobstructed view and can rotate through $360°$ in $7.5°$ increments. Consequently, 48 range readings are generated by each 'sweep' of the pivoting sensor head. The experiments where conducted in a relatively complex environment containing a number of boxes, a hatstand and a chair. The robot travels in a more-or-less straight line between the obstacles, taking range readings as it goes. The results shown in this section include readings from about 30 complete sweeps of the sensor head, about 1440 individual readings. The robot's location is determined by simple odometry. The robot has an onboard 486 processor which is fast enought to generate maps in real time (i.e. it can incorporate new range measurements at the rate at which they are acquired).

Figure 1 shows the maps produced for varying values of n (i.e. varying numbers of response directions). Each cell represents a region 4 centimetres square. In these maps, cells which are probably occupied (or for which we have strong support) appear darker than cells which are probably unoccupied. The dotted line shows the path of the robot.

Consider first the $n = 8$ result. The Bayesian and Dempster-Shafer maps are quite similar (notwithstanding the fact that in the Bayesian map, 'unknown' cells appear as gray, whereas in the Dempster-Shafer map they appear white). These maps clearly show the various boxes, the hatstand and the chair in the test environment. A surprising feature of these maps is that multiple reflections have *not* manifested themselves as spurious features. Looking next at the $n = 4$ and $n = 1$ results, it can be seen that there is a sharp decrease in the quality of the maps as n becomes small. Note particularly the $n = 1$ case, where we are effectively ignoring the response behaviour and attempting to compute cell occupancy directly (the Bayesian $n = 1$ case corresponds to the method described by Elfes [1]). The reason for this fall off is simple – the evidence obtained from multiple measurements may be contradictory when n is small. In the extreme case, $n = 1$, the entire approach collapses. At the other extreme, increasing the value of n much beyond 8 leads to a decrease in the quality of the results. When n becomes large, the chances that two measurements will fall into the same response 'bin' be-

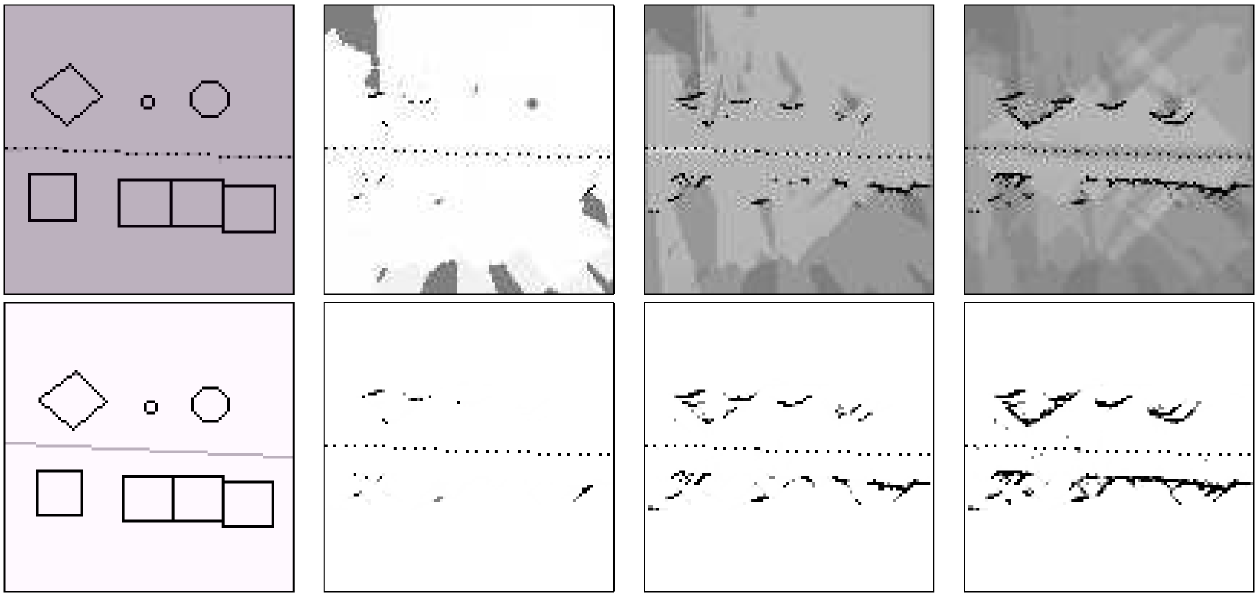

Figure 1: Occupancy maps. Top row: Bayesian ideal, $n = 1, 4, 8$. Bottom row: Dempster-Shafer ideal, $n = 1, 4, 8$.

comes very small. Consequently, there is no combination of evidence occuring at this level. For our particular experimental configuration, the optimal value of n is about 8.

4 Conclusion

The remarkable thing about the two methods presented in this paper is the similarity of the results they produce, with the Bayesian method appearing somewhat faster and slightly more efficient in its use of memory. The experimental results clearly demonstrate the advantage of the response grid framework in highly specular environments. We are also exploring some dynamic extensions which will make the response grid approach more suitable for use in changing environments.

References

[1] Alberto Elfes, "Occupancy grids: A stochastic spatial representation for active robot perception", in *Proceedings of the Sixth Conference on Uncertainty in AI*, 2929 Campus Drive, San Mateo, CA 94403, July 1990, Morgan Kaufmann Publishers, Inc.

[2] Hans Moravec, "Sensor fusion in certainty grids for mobile robots", *AI Magazine*, pp. 61–74, Summer 1988.

[3] R A Jarvis, "A perspective on range finding techniques for computer vision", *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 5, no. 2, 1983.

[4] Ömür Bozma and Roman Kuc, "Building a sonar map in a specular environment using a single mobile sensor", *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 23, no. 12, 1991.

[5] M K Brown, "Feature extraction techniques for recognising solid objects with an ultrasonic range sensor", *IEEE Journal of Robotics and Automation*, vol. RA-1, no. 4, 1985.

[6] Ruman Kuc and Victor Brian Viard, "A physically based navigation strategy for sonar-guided vehicles", *International Journal of Robotics Research*, vol. 10, no. 2, 1991.

[7] James O. Berger, *Statistical Decision Theory and Bayesian Analysis, Second Edition*, Springer-Verlag, New-York, 1985.

[8] J D Lowrance, T M Strat, L P Wesley, T D Garvey, E H Ruspini, and D E Wilkins, "The theory, implementation, and practice of evidential reasoning", Tech. Rep., SRI Project 5071, 1991, Final report.

[9] J Baldwin, "Support logic programming", in *Fuzzy sets - Theory and Applications, Proceedings of NATO Advanced Study Institute*, A Jones et al, Ed. 1986, Reidel Pub. Co., Norwell, MA.