Evaluation of feature-based room maps

Kevin Nickels and Michael Sutton
{knickels,msutton}@engr.trinity.edu
Department of Engineering Science
Trinity University
San Antonio TX 78212
Ph: 210-999-7543 Fax: 210-999-8037

Abstract

In this paper, we describe a method for evaluating the consistency and sufficiency of a proposed feature-based room map of unspecified origin with regard to a single sensor measurement. This method analyzes the possibility of the proposed map giving rise to the given dense sonar scan. The features from which the map is derived can be from any of a multitude of sensors.

These quality ratings allow hypothesized room maps to be ruled out if they are inconsistent with observed data, and allow an autonomous robotic system to arrive at a "most plausible explanation" for a room map based on sensor measurements.

Keywords: Mobile Robotics, Map Evaluation, Feature-Based Maps

1 Introduction

Many artificial intelligence researchers use the abilities of humans and animals as motivation, if not inspiration, for new capabilities of artificial systems. The ability to navigate in a partially unknown environment has led to much work in the area of mobile robot navigation and map-building.

In navigation tasks, a map of the area is often provided as input to the system, and the task given is to move around in the area. Localization, a subset of navigation, is the process of determining the current position of the robot in a specified environment. In map-building tasks, the goal is to create a map of the area based on information that has been gathered by the mobile robot. The tasks are often intertwined, as a knowledge of the current position of the robot in the environment is required to integrate new sensor readings from the robot into the map [12].

There are several different ways for a robot to gather information about its environment. Due to the low cost and ease of use, sonar transceivers are a popular sensor. The Polaroid 6500, one of the most popular transceiver kits, provides circuitry to detect and return the TOF (time of flight) for a sonar pulse [14]. If reflections of sonar pulses are not considered, this information translates directly into the distance between the sensor and the nearest obstacle in the view of the sensor.

Visual information can also be used to gather information about the surroundings of a mobile robot. Structured lighting or stereo vision, for example, can provide range information about the environment [8].

There are two broad categories of environmental representation: grid-based and feature-based. In a grid-based representation, the environment is segmented into squares[1]. Based on measurements of the environment and a-priori information, each square is assigned a value of *empty* or *full*. It is also possible to assign a numeric value to each grid, where one end of the numeric range represents *empty* and the other end represents *full*. A perfect map would have *full* cells at each location containing obstacles to the mobile robot, and *empty* cells elsewhere. Elfes [5] uses such a probabilistic occupancy grid representation to fuse range data from sonar and stereo range measurements, and to arrive at a grid-based map suitable for navigation. Pagac et al. [13] utilize a probabilistic model for the sonar sensor to fuse new measurements into an occupancy grid for an autonomous vehicle.

However, most localization and navigation techniques that do not include explicit map-building use a feature-based representation of the environment. For example, Crowley [3] uses a global world-map that consists of line segments. Drumheller [4] assumes a complete outline of the room is available. Leonard et al. [11] assumes a model which is complete enough to estimate sonar returns. For these reasons, we will consider only feature-based representations in environmental models.

In the past few years, several methods have been developed that allow autonomous robots to utilize various forms of *intelligent sensors* to observe their environment. Barshan and Kuc [2], and later, Kleeman and Kuc [7] use a linear array of sound navigation and ranging, or sonar, transponders to determine if an echo corresponds to a plane,

[1] in the case of a three dimensional environmental model, a cube is the basic element.

a corner, an edge, or unknown. Han and Hahn [6] describe a similar system utilizing two pairs of ultrasonic sensors and time-of-flight information. Kuc [9] has also analyzed the shape of the return sonar signal to perform recognition on the object reflecting the sonar pulse. This system can distinguish between ball bearings, machine washers, and O-rings of different sizes by searching a database of sonar return signatures. Leonard and Durrant-Whyte [11] extract regions of constant depth (RCDs) from dense sonar scans, and characterizes the behavior of RCDs that correspond to corners, walls, or cylinders in the environment. Leonard et al. [12] use this for simultaneous mobile robot localization and feature tracking.

These intelligent sensors, as well as others utilizing techniques such as computer vision and laser range finding can provide detailed information about portions of a mobile robot's environment. Often these feature locations are sufficient for localization and navigation in a known environment. However, a truly flexible autonomous system needs the ability to evaluate complete room maps in the light of observations, and to test hypotheses about its environment.

2 System Architecture

This section describes our procedure for evaluating feature-based maps. First, we cover some notation and define our data structures for the problem. Then we describe the models used to simulate the behavior of sonar returns in structured indoor environments. Finally, we describe our method for arriving at particular metrics that reflect the quality of a proposed map.

2.1 Definitions and Representations

We define a map $\mathbf{M}$ to be a collection of line segments $\mathbf{L}$ and corners $\mathbf{C}$. Each line segment $\mathbf{l} \in \mathbf{L}$ is represented by a parameter vector: $\mathbf{l} = [R, \theta, V, p_1, p_2]$ that contains the parameters (R, θ) for the line in Hessian normal form [1] (that is, the perpendicular distance and angle representation), a binary variable V indicating which half-plane of the line is visible, and two optional endpoints $p_1 = (x_1, y_1)$ and $p_2 = (x_2, y_2)$ to the line segment. Each corner $\mathbf{c} \in \mathbf{C}$ is also represented by a parameter vector: $\mathbf{c} = [x_c, y_c, \theta]$ that contains the location and orientation of the corner. An unoccluded corner is assumed to be visible if the observer is on the inner half plane of each wall that comprises the corner.

We make a distinction between a map $\mathbf{M}$ and a collection of features such as lines, corners, edges, and cylinders. A dense set of features approximates a map as we have defined it, but may require much more storage. A map may or may not have a large amount of support for the lines which comprise it: it is the purpose of the algorithm presented in this section to evaluate the quality of a hypothesized map.

We next define several quality metrics for maps. A map may be generated by any of a number of sensors, or may be hypothesized from other incomplete information such as previous experiments. This map will be compared with a single range measurement (we will use a dense sonar scan, although other range measurement could also be used) to determine the quality of this map as an explanation for the given measurement.

An observation that is used to evaluate a proposed map is called the **validation observation.** An observation that is simulated from a proposed map is called the **prediction observation.** An observation is comprised of some set of individual measurements taken from the same location in the room. A measurement is computed from a dense sonar scan. A map is **consistent** if it does not contradict any measurement in the validation observation. An **explanation** is comprised of a set of measurements from the validation observation that match well with the prediction observation.

2.2 Verification Observation

One of the more popular sonar transducers used in mobile robotics is the Polaroid 6500 module [14]. This module, which includes a transmitter, transducer, receiver, and thresholding circuit, is standard on many mobile robots. This along with its low cost has lead to the widespread use and modeling of this module. Kuc and Siegel [10] present an in-depth analysis of the physical behavior of sonar.

We follow Leonard et al. [11] in utilizing the Regions of Constant Distance (RCDs) in sonar scans as our base environmental measurement. As scans are taken from different, but similar, locations in the environment the behavior of RCDs arising from corners and walls differentiates itself. Leonard et al. use a tracking algorithm to identify clusters of RCDs and track them as the mobile robot moves through the environment. We shall simply consider them as an indication of a strong sonar return.

Our verification observation is computed by extracting the RCDs with width greater than $\beta = 10^\circ$ from a dense sonar scan. Each RCD is assigned a distance and bearing. The next section describes the prediction of measurements and observations from a proposed map.

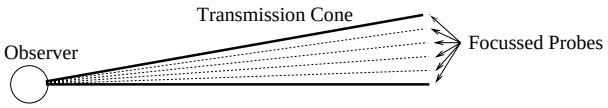

Figure 1: Computing a range measurement from sonar probes.

2.3 Prediction Observation

Given a map $\mathbf{M}$, we wish to predict the range R measured at a given orientation θ from a known origin (s_x, s_y). This is a problem in the simulation of sonar returns. Kuc and Siegel [10] have presented a thorough and accurate simulation method for air-based sonar. Leonard and Durrant-Whyte [11] give a simplified version of this method that is sufficient for tracking and navigation purposes. We present below an extension to Leonard and Durrant-Whyte's method that accounts for reflection and deals with incompletely specified line segments.

We assume that the map is given, and the observer location and orientation are known. An observation is made by concatenating measurements taken at some small fixed rotation interval (we used 0.5° in our experiments) from the observer orientation. We now discuss the computation of an observation.

While sonar does not behave as a ray tracer, the theoretical behavior of small focused pulses of sound are useful in estimating the behavior of a practical sonar system. A theoretical return is calculated for a large number of these *probes* within the sonar transmission cone (which has been estimated to have a half-angle of 5.7° [11]), and the minimum measurement for these is taken as the range measurement. Figure 1 illustrates the relationship of probes to range measurements. This approximation yields expected behavior while decreasing computational burden over a continuous transmission cone.

The computation of a sonar probe is largely a problem in geometry. The proximity of the probe to any corners in the map is computed. If a probe intersects a corner, the sonar probe returns the TOF from the observer to the corner. We encode this knowledge in an *observation equation* for corners,

$$\mathbf{h}_c(\mathbf{x}(k), \mathbf{C}) = \sqrt{(x_c - x(k))^2 + (y_c - y(k))^2},$$

where the vector $\mathbf{x}(k) = \begin{bmatrix} x(k) & y(k) \end{bmatrix}$ gives the location of the observer.

For the purposes of simulation, an exact intersection is not required. If the probe passes within a threshold region g_c of the corner, it detects the corner. Note that this is an approximation, as corners may be detected by a double reflection off the adjacent walls as well as a direct sonar return.

The intersection of the probe with any walls in the map is also computed. Recall that a wall is represented by a line segment that may or may not have endpoints specified. The normal wall observation function, the orthogonal distance from the wall to the observer, only applies in certain conditions. The angle between the probe and the normal to the wall must be less than the reflectance angle $\beta_t/2$ of the wall, so that the sonar return is detected by the transducer. The ray from the observer to the wall that intersects the wall at a right angle must intersect within the endpoints of the line segments. If the reading would fall outside the specified endpoints, no return is generated. If the angle between the probe and the normal to the wall is between $\beta_t/2$ and $\pi/2$, the sonar pulse is assumed to specularly reflect off the wall.

Previous researchers vary in their approaches to specular reflections. Kuc and Siegel [10] simulate the physical behavior of sonar faithfully, and deal with specular reflections. In Leonard and Durrant-Whyte [11], one simplification that is made is to ignore reflections of this type (in their terminology, RCDs of order greater than two). The very high visibility of corners in sonar scans, and the reason that corner RCDs extend behind the adjacent walls, is due to the specular reflections off adjacent walls.

In this work, we also account for sonar returns that specularly reflect off walls, and return the image of a corner to the observer. This case is common, and can be seen in the upper right return of Figure 3. While many researchers look upon this case as an error, we believe that it contains useful information, and should be exploited to help validate proposed maps. If a probe is determined to specularly reflect off a wall, the probe return should be the TOF from the observer to the reflection point added to a new probe taken at the reflection point in the reflected probe direction. This new probe is computed in an identical fashion to those described above. In the current implementation, we do not allow more than one specular reflection (not counting the double reflections that arise from corners). We define $\theta(k)$ to be the angle of intersection of the sonar pulse with respect to the normal of the wall. If $\theta(k) < \beta_t/2$,

$$\mathbf{h}_w(\mathbf{x}(k), \mathbf{l}) = V(R - x(k)cos(\theta) - y(k)sin(\theta))$$

otherwise,

$$\mathbf{h}_w(\mathbf{x}(k), \mathbf{l}) = \sqrt{(W_x - x(k))^2 + (W_y - y(k))^2} + \mathbf{h}([W_x, W_y, \phi]),$$

where ϕ is the reflection angle of the sonar pulse off the wall, and (W_x, W_y) is the location along the wall of the pulse reflection.

The minimum probe within a range measurement is returned as the value of the range measurement. The range measurements are concatenated to form a dense sonar scan, and RCDs with width greater than $\beta = 10^\circ$ are extracted from this scan. This list of RCDs forms the predicted observation.

2.4 Evaluation of Maps

The predicted observation given the maps is computed as described above. Validation gates are placed around each measurement in the observation, as seen in Figures 4 and 5. The validation observation is compared to this predication, and the number of measurements in the validation observation that are explained and not explained is recorded. A measurement is explained if it falls within the validation gate of one or more predicted measurements. A measurement is not explained if it does not fall within the validation gate of any predicted measurement. A validation measurement is contradicted if it falls within the angular portion of a validation gate, but is of shorter range. A good explanation of the use of validation gates in this fashion can be found in [11].

Two metrics are used to quantify the quality of a proposed room map. As defined in Section 2.1, map *consistency* is the minimum sufficient condition for a proposed map to be possible. This condition states that none of the measurements in the validation observation are not explained by the prediction observation.

Definition 1 *The consistency metric $Con(\mathbf{M})$ is defined to be zero if a map is not consistent, and equal to the ratio of explained measurements to measurements in the validation observation if the map is consistent:*

$$Con(\mathbf{M}) = \begin{cases} 0 & \exists \ \textit{contradiction} \\ \frac{m_{exp}}{m_{vo}} & \textit{otherwise} \end{cases}$$

where m_{exp} is the number of explained measurements, and m_{vo} is the number of measurements in the validation observation.

As illustrated in Figure 2, a highly consistent map will have a validation observation that is explained well by the map. A map with a low but nonzero consistency measure is likely very sparse. Note that the consistency metric does not evaluate how extensively a validation observation tests the map, but how well the map matches the given observation.

The second metric evaluates the extent that a given explanation has tested the map. For example, a prediction observation that only has a few RCDs in it will turn out to be consistent with many possible room maps. The *sufficiency metric* $Suf(\mathbf{M})$ attempts to determine how well the prediction observation is explained.

Definition 2 *The metric $Suf(\mathbf{M})$ is defined to be the ratio of explained measurements to measurements in the prediction observation:*

$$Suf(\mathbf{M}) = \frac{m_{exp}}{m_{po}}$$

where m_{exp} is the again the number of explained measurements, and m_{po} is the total number of measurements in the prediction observation.

A map with a high sufficiency rating indicates that the predictions generated from the map are matching the validation measurements well. This rating does not convey how well the map exercises the validation observation, just how well the prediction measurements are matching.

In order to illustrate these definitions, we show in Figure 2 examples of several maps. The map is shown as solid lines. The verification observation is shown as a series of dashed rays. Explained verification measurements are shown as circles.

Using these two metrics, we can not only evaluate how well a proposed map models the (single) validation observation, we can also tell how well the validation observation exercises this match. A map with high scores in both consistency and sufficiency is likely to be useful for localization and navigation, and can be propagated throughout the system.

3 Results

In this section, we present some results of the application of the metrics described in Section 2. We present validation observations, prediction observations, and the evaluation of two proposed room maps for a simple room.

Figure 3 shows the validation observation, overlaid on a hand-measured room map. This observation was computed by finding the regions of constant depth in a single dense sonar scan. The range measurements which contributed to the RCDs are shown in the figure.

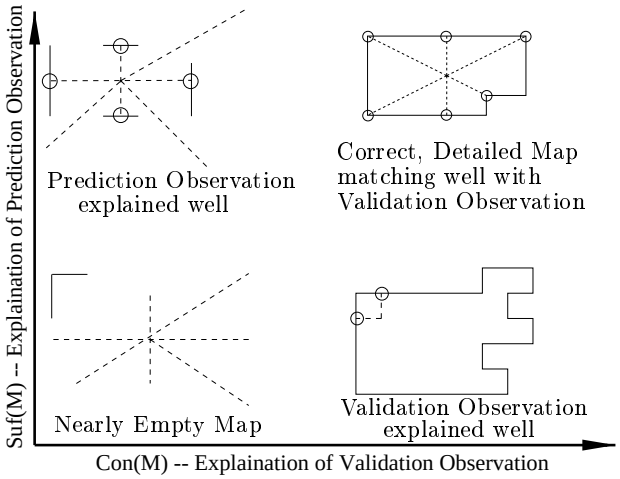

Figure 2: Consistency and Sufficiency Metrics

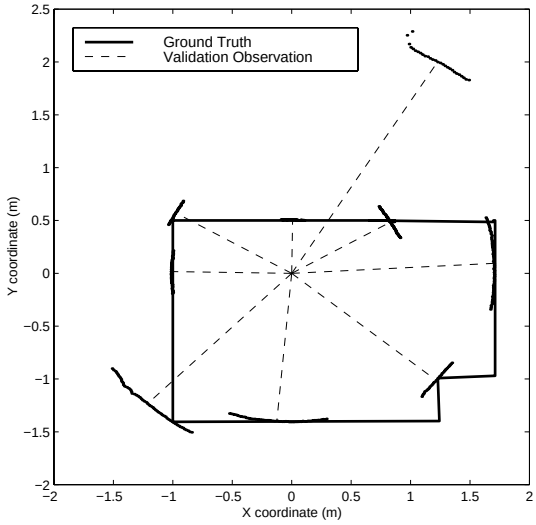

Figure 3: Actual room map and validation observation.

Figures 4 and 5 illustrate the observations predicted for two representative maps, with the associated validation gates. In the first case, the bump-out at the lower right of the room is missing. This leads to two missing measurements, directly and from the reflection in the top wall, and leads to an inconsistent map. The second map is consistent and explains many of the measurements.

4 Discussion

The maps shown in Figures 4 and 5 illustrate the use of the consistency and sufficiency metrics. In Figure 4, map $\mathbf{M}_1$ shows a likely proposed map for the data. However, $\mathbf{M}_1$ does not contain the edge in the lower right corner of the room. Since the validation measurement is within the angular validation gate, but shorter in range than the associated prediction measurement, this leads to a contradiction, which might invalidate this map from consideration. A separate contradiction is detected in the image of this corner in the north wall.

The hand measured room contains a door moulding along the upper wall, leading to a corner return along this wall. This feature is not present in either $\mathbf{M}_1$ or $\mathbf{M}_2$. However, since there is not a predicted observation in this direction, this missing feature will not lead to a contradiction. In general, contradictions will not arise from missing features, but may arise due to misplaced features.

Another point in $\mathbf{M}_2$ worth discussion is the upper right corner. The range measurements from the right wall are approximately the same range as the measurements from the corner. The RCD extracted from the validation scan combines the range measurements from these features into a single measurement. The prediction observation anticipates this, as can be seen from the overlapping validation gates. The validation measurement matches with the prediction measurement from the wall, which reduces the size of the explanation but does not cause serious

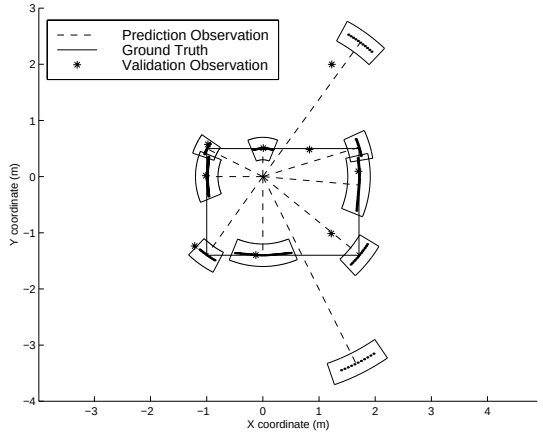

Figure 4: Map $\mathbf{M}_1$ and associated observations. $Con(\mathbf{M}_1) = 0$, $Suf(\mathbf{M}_1) = 5/10$

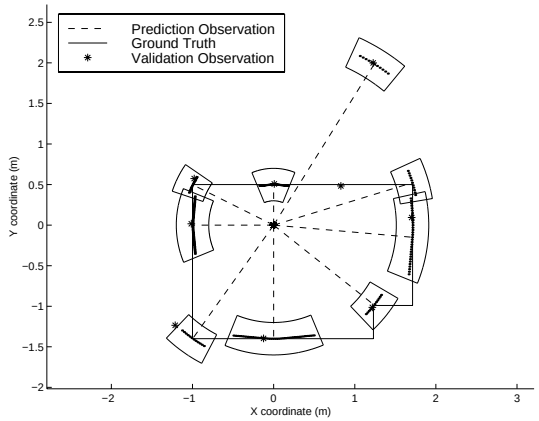

Figure 5: Map M_2 and associated observations. $Con(\mathbf{M}_2) = 7/9$, $Suf(\mathbf{M}_2) = 7/9$

problems.

Figure 5 shows a good match with the observed features. The two points where validation measurements are not matched are the door moulding on the upper wall and the lower left corner. This map would very likely survive the proposal stage and be propagated throughout the system.

5 Conclusions

We have illustrated a method for evaluating the quality of a proposed room map by comparing predicted sonar features from that map with observed features on a validation observation. This method will allow a system to check several possible maps for quality before attempting to use one for motion planning, navigation or further investigation of the environment.

Advantages to the use of a feature-based room map instead of simply recording all the features that are observed include storage size, possibilities of extrapolation from observed features, and ease of use in geometric algorithms.

Increasing the allowable complexity of the room map is high on our research agenda. The addition of doorways and allowances for more topologically complex rooms need to be investigated. In addition, the postulation of missing measurements that would increase the overall belief of the map would be helpful both for environmental map building and for learning guidance of mobile robots.

6 Acknowledgements

Thanks to J. Leonard for providing the sonar data.

References

[1] D. Ballard and C. Brown. *Computer Vision*. Prentice-Hall, 1982.

[2] B. Barshan and R. Kuc. Differentiating sonar reflections from corners and planes by employing an intelligent sensor. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 12(6):560–569, June 1990.

[3] J. L. Crowley. Navigation for an intelligent mobile robot. *IEEE Journal of Robotics and Automation*, RA-1(1):31–41, March 1985.

[4] M. Drumheller. Mobile robot localization using sonar. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, PAMI-9(2):325–332, 1987.

[5] A. Elfes. Sonar-based real-world mapping and navigation. *IEEE Journal of Robotics and Automation*, RA-3(3):249–265, 1987.

[6] Y. Han and H. Hahn. Localization and classification of target surfaces using two pairs of ultrasonic sensors. In *Proceedings IEEE International Conference Robotics and Automation*, pages 637–643, Detroit, MI, May 1999.

[7] L. Kleeman and R. Kuc. An optimal sonar array for target localalization and classification. In *Proceedings IEEE International Conference Robotics and Automation*, pages 3130–3135, San Diego, CA, May 1994.

[8] D. Kriegman, E. Triendl, and T. Binford. Stereo vision and navigation in buildings for mobile robots. *IEEE Transactions on Robotics and Automation*, 5(6), December 1989.

[9] R. Kuc. Fusing binaural sonar information for object recognition. In *Proceedings International Conference Multisensor Fusion and Integration for Intelligent Systems*, Washington, DC, USA, December 1996.

[10] R. Kuc and M. W. Siegel. Physically based simulation model for acoustic sensor robot navigation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 9(6):766–778, 1987.

[11] J. J. Leonard and H. F. Durrant-Whyte. *Directed sonar sensing for mobile robot navigation.* Kluwer Academic Publishers, Norwell, MA USA, 1992.

[12] J. J. Leonard, H. F. Durrant-Whyte, and I. J. Cox. Dynamic map building for an autonomous mobile robot. *International Journal of Robotics Research*, 11(4):286–298, August 1992.

[13] D. Pagac, E. M. Nebot, and H. Durrent-Whyte. An evidential approach to map-building for autonomous vehicles. *IEEE Transactions on Robotics and Automation*, 14(4), August 1998.

[14] Polaroid Corporation, Commercial Battery Division. *Ultrasonic ranging system*, 1984.