Sonar-based Feature Recognition and Robot Navigation Using a Neural Network

Puneet Goel* **and Gaurav S. Sukhatme**†
Department of Computer Science
Institute for Robotics and Intelligent Systems
University of Southern California
Los Angeles, CA 90089-0781

Abstract

An efficient strategy based on an artificial neural network is presented to enable a mobile robot to move autonomously and build a rough map of its environment. The paper describes a technique based on a neural network which uses data from sonars to recognize features online. Our method is systematic, simple and yields very good results. Moreover, it is generalizable in the sense that one can train the robot to recognize particular features of interest. We have implemented and tested this technique successfully on board a mobile robot which uses a 68332 processor and does not have support for floating point arithmetic and file structure. Once trained, the robot can recognize the features on the fly and react to them appropriately.

1 Introduction

In order to autonomously navigate and perform useful tasks, a mobile robot needs to recognize features (e.g. corners, intersections) in its environment and react to them appropriately. Autonomous navigation is a difficult problem for traditional robotics techniques, primarily because of the noise and variability associated with real world scenes.

A popular technique using sonar is to detect obstacles via reflection of the sound from them. If the distance between the robot and obstacle is small, evasive action is taken. In order to implement obstacle avoidance the designer has to manually model the sonar. In typical indoor environments, the problems posed by specular reflection from active time-of-flight devices such as sonars are severe. Sonar based wall following and obstacle avoidance are commonly implemented in mobile robots. An example of wall following is [11] in which the robot has to go through successive runs in the environment before learning the important features for wall following. In our work we describe a system that associates multiple key features in an interior environment with basic navigation primitives. Once trained, the system can freely explore novel interior spaces without bumping into obstacles.

Various groups have implemented mapping using sonar. While [7] [4] use occupancy grid models, [10] uses a target tracking model. In the target tracking paradigm, the main problem is to associate the sensor readings with features while the occupancy-grid model suffers from enumerating and updating the probabilities for all possible environmental situations. Moreover both strategies need to explicitly model the sensor (sonar) used.

The learning community has also attempted to identify navigable features automatically. The philosophy behind this approach is that the robot does not require a detailed, geometric map of the environment to function properly. Galles [5] used tree representations to recognize landmarks such as corners, intersections and then navigate between them. Kuipers et. al. [8] also focussed on identifying distinctive positions and developed control methods for navigating between them. Kirchner and Hertzberg [6] used modular control software (finite state automata) which reflected a hierarchical decomposition of the global task to enable a robot to navigate autonomously through a network of sewage pipes recognizing and responding to corners and crossings. Thrun [13] trained an artificial neural network to extract landmarks/features from sensor data. The mobile robot first learned which features were best suited for localization using Bayesian techniques. In [9] the authors describe an architecture which is an instance of feature extraction ap-

*This work was performed during a summer internship at IS Robotics by the first author
†Contact author for correspondence

proach to mapping and navigation in office environments. Simmons and Koenig [12] used partially observable Markov Models to robustly track a robot's location and integrated it with a planning and execution monitoring approach to control the robot's actions. All the approaches explicitly derive mathematical models or complicated structures for recognizing features. We propose a technique which is relatively simple and easy to implement.

In this work we train a neural network to detect features. Using a neural network for navigation is not new in robotics. Several groups have developed sensor fusion for robot navigation using neural networks. Cao et. al. [3] used sonars for determining the relative position of robot from obstacles and pattern recognition (using neural network) on a vision-based system for navigation. [15] also used vision-based navigation complemented with neural network. Since specular reflection and wide beam pattern from sonars contributes to potentially large errors, it is common to use sonars complemented with a vision system for navigation. We have been able to train a neural network on sonar data and robot relies solely on it for navigation. We do not use a vision system for feature recognition unlike most of the work cited above. The neural network is implemented on a 68332 micro processor which does not support floating point arithmetic. The network takes approximately two minutes for training on this processor and is able to recognize the features in real time. All the processing takes place on board and once trained the robot can make a rough map of the environment online while moving autonomously.

The rest of the paper is organized as follows. Section 2 presents the robot model. Section 3 describes the neural network used in our experiments. Section 4 describes the features selected and the collection of data. Section 5 presents the experimental results. Section 6 concludes with a summary and a discussion of ongoing and future research.

2 Robot Model

The Fetch II (manufactured by IS Robotics[1]) robot used for experiments is a four wheeled robot shown in Figure 1. The wheels on the same side are tightly coupled. The encoders return two distinct speeds; one for the right pair of wheels and other for the left pair of wheels. The robot is provided with narrow beam, page electric sonars with a beam angle of ten degrees. The robot has two rings of eight sonars each. The front has one full ring of eight sonars while the other eight are distributed, with four on back and two each on both sides. These sonars have a blind zone of 30 cm (that is, they cannot see anything that is at less than or equal to 30 cm from them) and from 30 cm to 45 cm they detect a second echo instead of the first echo. Hence if an obstacle is 40 cm away the sonars detect it at a distance of 80 cm.

[1] www.isr.com

Figure 1: The Fetch Robot from IS Robotics.

The Fetch robot is controlled using a behavior based controller. Behavioral control is a decentralized approach to the architecture of robotic control systems [2]. The navigation and control code that communicates with ports is written in L. C routines can be called from the L code but the C code has to reside in a library. The robot does not support floating point arithmetic. This complicates the implementation of the normalization step (that keeps all the weights of the links and of the hidden and output units in the neural network between 0 and 1) which is an important step in using neural networks. This is not a simple scaling problem and we discuss it in the next section.

3 Neural Network

Neural Networks are best suited for classification problems where the input is high dimensional and noisy [1]. Since we had sixteen sonars which are noisy sources of information we decided to use a neural network to detect features. An artificial neural network, trained with the well known error backpropagation algorithm [14] was used to process the raw data given by the sonars as a pattern. The network associated the data with the appropriate feature in the environment.

The neural network is trained using the sonar data collected for features shown in Figure 2. The robot was placed at the sites A, B, C and D (see Figure 3) in different positions to get the appropriate data.

A total of eighty five patterns were used to train the network. The network had sixteen input units, one for each sonar, nine hidden units and nine output units. The backpropagation equations used are:

$$(\triangle \mathbf{w}_{ji})_n = \eta \delta_j \mathbf{o}_i \; + \; \gamma (\triangle \mathbf{w}_{ji})_{n-1} \tag{1}$$

$$\delta_j = \mathbf{o}_j(\mathbf{1} - \mathbf{o}_j)(\mathbf{t}_j - \mathbf{o}_j) \;\; for\ output\ unit \tag{2}$$

$$\delta_j = \mathbf{o}_j(\mathbf{1} - \mathbf{o}_j)\Sigma_k \delta_k \mathbf{w}_{kj} \;\; otherwise \tag{3}$$

where $\triangle \mathbf{w}_{ji}$ represents the change in the weight of the link joining j^{th} and i^{th} units, δ represents the change in weights of output and hidden units, and η and γ denote learning rate and momentum respectively. $\mathbf{o}_j$ represents j^{th} output unit while $\mathbf{t}_j$ is the desired output for that unit.

On studying the equations 1-3 carefully one notices that the number **1** occurs prominently. The general trend is to normalize the input between 0 and 1 before feeding it to the neural network. Also the most popular choice for a non-linear output function that allows computation of arbitrary functions is the sigmoid function.

$$\mathbf{o}_j = \frac{1}{1 + e^{-(\mathbf{net}_j - \mathbf{T}_j)}} \tag{4}$$

$$\mathbf{net}_j = \Sigma_k \mathbf{hidden}_k \mathbf{w}_{kj} \tag{5}$$

where $\mathbf{net}_j$ represents the sum of the products of weights of hidden units and links joining them to j^{th} output unit while $\mathbf{T}_j$ is the threshold for the j^{th} output unit.

The problem in using this convention is that it needs support for floating point numbers since all quantities are scaled between 0 and 1. As stated in the previous section we did not have floating point support on the robot. It was observed that the maximum output from the sonars was 4504. We decided to use 45 (divide the sonar data by 100 before feeding it to neural net) as the maximum input. At first sight it seems that scaling the quantities between 0 and 45 solves the problem. However if we scale everything from 0 to 45 then the maximum weight of a link joining the input and hidden layer can be 45 and so the update in the hidden layer is of the order 45^2 (since a unit in the input layer can have a maximum weight of 45) and therefore the update in the output layer will be of order 45^3 which is not desirable. So we need something more than just scaling. Our solution is to scale the weight of the units by 45 and that of the links by 1000 (so as to be precise up to third decimal place) but when updating the weights we divide the change by the scaling factor to preserve the right effect. The learning rate and momentum are also scaled by 1000. So now the equations become:

$$(\triangle \mathbf{w}_{ji})_n = \frac{\eta \delta_j \mathbf{o}_i}{45 * 1000} + \frac{\gamma (\triangle \mathbf{w}_{ji})_{n-1}}{1000} \tag{6}$$

$$\delta_j = \frac{\mathbf{o}_j(\mathbf{45} - \mathbf{o}_j)(\mathbf{t}_j - \mathbf{o}_j)}{45 * 45} \;\; for\ output\ unit \tag{7}$$

$$\delta_j = \frac{\mathbf{o}_j(\mathbf{45} - \mathbf{o}_j)\Sigma_k \delta_k \mathbf{w}_{kj}}{45 * 45 * 1000} \;\; otherwise \tag{8}$$

$$\mathbf{o}_j = \frac{45}{1 + e^{-\frac{(\mathbf{net}_j - \mathbf{T}_j)}{45}}} \tag{9}$$

We used a bias unit with weight 45. The sigmoid function is approximated using a look up table indexed on the value $\frac{(\mathbf{net}_j - \mathbf{T}_j)}{45}$. The code for the neural network is written in C and resides in a library. The L routine for navigation control interacts with the sonars directly, passes the sonar values to the C procedure, gets the neural network output, processes it and issues the appropriate navigation command based on the output.

4 Features

The commonly occurring features in an indoor office-like environment were chosen for recognition so that the robot would be able to move comfortably in various similar environments. A multilayered network with a single hidden layer was used to recognize nine features (Figure 2) namely (0) Taking a turn (1) Front Closed (2) Corridor (3) Right Closed (4) Back Closed (5) Front and Left Closed (6) Back and Right Closed (7) Back and Left Closed (8) Front and Right Closed. Feature number zero (taking a turn) was used to avoid detection of a feature while robot was turning. This prevented the robot from responding to a feature while turning hence avoiding the possibility of getting lost by getting stuck in an infinite loop for turning.

The training data was collected in a corridor environment. Sites A, B, C and D in Figure 3 were used for collection of data. The corridors used were of varying width ranging from 1.2 m to 1.6 m. The collection of data for features 1 to 8 was done with the robot stationary. There are 9 features and we used 4 sites. The robot was placed appropriately to collect the data for the required features using only the above mentioned sites. The response to each feature was also defined and hard-coded. For example, if the robot sensed that the front and left directions are closed the appropriate action is a right turn. Similarly if it senses front, left and right closed it will return along the path it came from. Similar responses were defined for the other corners and intersections.

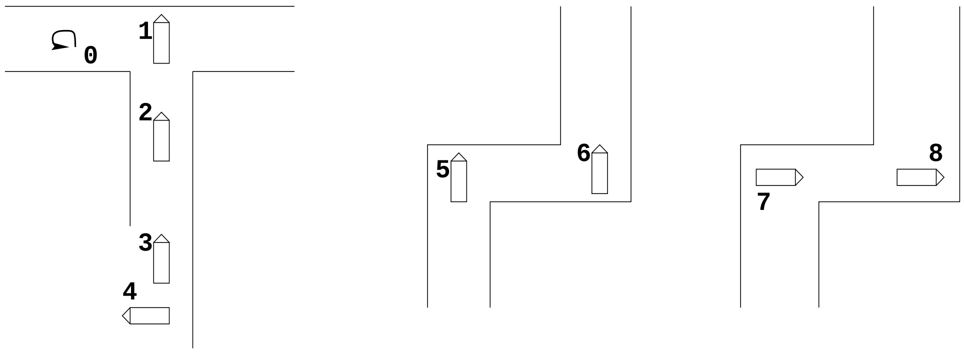

Figure 2: Features 0-8 on which the neural network was trained. (0) Taking turn (1)Front Closed (2)Corridor (3)Right Closed (4)Back Closed (5)Front & Left Closed (6)Back & Right Closed (7)Back & Left Closed (8)Front & Right Closed.

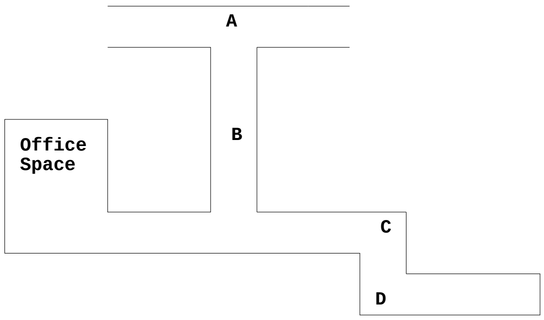

Figure 3: The training environment. A, B, C, D denote the training sites.

To reduce the number of false alarms (incorrect recognition) we experimented and calculated the confidence level given by neural network above which the robot would be sure about the identification of the feature. The neural network produces weights (which we identify with confidence levels) between 0 and 45 for various output units. We chose 25 as the threshold. If an output unit had weight greater than 25, only then was the feature corresponding to that output unit assumed to be correctly identified. Further, to make the implementation robust, the robot responded to a feature only if the feature was detected for at least three consecutive time steps. This improvement prevented the robot from reacting to false alarms. The data rate was 20 Hz so it would take no more than a fraction of second for the robot to react to the feature. Using these two heuristics the false alarm rate dropped dramatically.

5 Experimental Results

The aim of our study is to show that a neural network is capable of recognizing features even in environments which are unknown to it. Hence very few test sites were used for our training patterns. To show that this work is generalizable (i.e. the robot will be able to recognize features in environments different from the training environment) we conducted various experiments (Figure 4 and Figure 5). The Figure 4 shows the movement of robot in an environment similar to the one in which it was trained. The office space was totally new to it. There were a lot of cupboards and various machines (copiers etc) in that space but irrespective of this new environment the robot turned appropriately at the corners and toured the space. There were also people walking in and out of the office space.

All the training patterns were collected in the corridors of width 1.2 m - 1.6 m. Table 1 shows instances of sonar data obtained for four different features. We wanted to observe the response of the robot when it is moving in corridors with different widths. For this we chose a test site that had corridors of width ranging from 2.25 m to 2.75 m. The first plot in Figure 5 shows this test run. The robot performed well and there were no false alarms. There was a door (marked in the figure) on the right hand side at the end. If that door happened to be open the robot went through it otherwise the robot retraced its path. No human intervention was made during this run.

We also tested our implementation in an environment which had different kind of walls (so that the reflections were not similar to the ones in the training patterns) and also had wider corridors. For this we chose a garage-like environment which had brick walls with no paint and the corridor was approximately 3.5

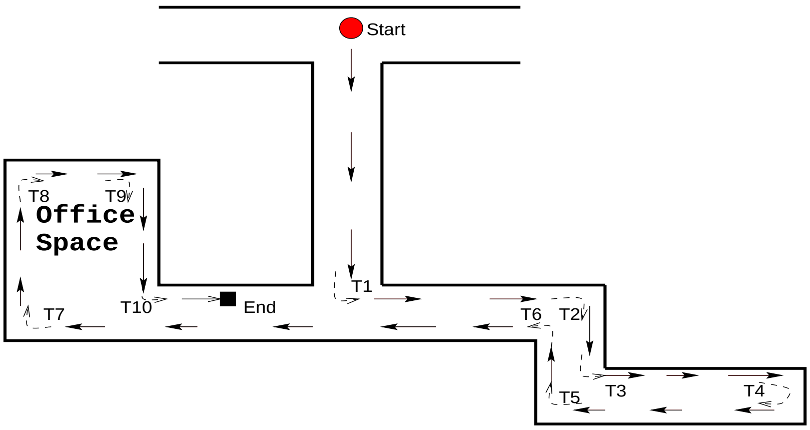

Figure 4: A test environment same as the training environment except the office space. Solid arrows show the path traversed by robot. The turns are marked by dotted arrows and Ti denotes the i^{th} turn.

Feature	2 Lt. Side Sensors.	2 Rt. Side. Sensors	8 Front Sensors	4 Back Sensors
Back Closed	4508 4508	4508 4508	4508 4508 4508 4508 4508 4508 4508 4508	537 583 615 598
Corridor	599 614	516 525	4508 4508 4508 4508 4508 4508 4508 4508	4508 1861 1870 4508
Front Closed	4508 1962	4508 4508	4508 4508 4508 1238 1227 1328 4508 4508	4508 4508 4508 4508
Left & Front closed	823 763	1386 4508	4508 4508 4508 1338 1326 4508 4508 1294	4508 4508 4508 4508

Table 1: Typical sample of sonar data for 4 different Features.

m wide. The test run is shown in the second plot of Figure 5. Again the robot performed very well and no false alarms were raised. One notices that unlike the other two plots where the robot takes a 180^o turn when it reaches the dead end, here the robot takes a 90^o turn (because it observes the end as front and left closed), gets to the other dead end, makes another 90^o turn and hence retraces its path. This happened because the corridor was too wide as compared to the one used for training and hence the robot thinks that it still has one more direction to explore and goes in that direction.

The experiments show that robot faithfully follows the wall and therefore provides a very good sketch of the corners, openings etc in the environment. It takes a 90^o turn or a 180^o turn depending on whether there is anything more to explore on either side of the corner or not. The first test in Figure 5 show that it goes through the door if it is open and explores the area inside. This results in an obvious exploration strategy. In the test environment of Figure 4 the robot moved confidently through the 'office space' and avoided hitting objects. It was also observed that if the robot gets surrounded by a bunch of objects (people, machines etc) and does not have an open space for continuing the motion, it starts rotating and tries to look for an opening. As soon as it finds one, it escapes through it and continues the navigation task. This behavior is not hard-coded and is the result of the robot detecting confined spaces as various kinds of corners and its reaction to them explains this typical behavior. The important thing here is avoiding hitting objects in a congested environment.

6 Conclusions and Future Work

In this paper we presented a neural network based technique to enable a mobile robot to comfortably move in an environment via efficient feature recognition and navigation primitives. The main thrust of the work is the generalizability and simplicity of the approach. There are no sophisticated modules or cumbersome mathematical equations to estimate the distance of the obstacles from the sonar data. Instead, a pattern recognition approach based on a neural network is used to detect features. The responses to the features are hard-coded. We have shown that sonars complemented with a neural network can result in an efficient, simple exploration strategy. We realize that our system cannot recognize complex objects by only

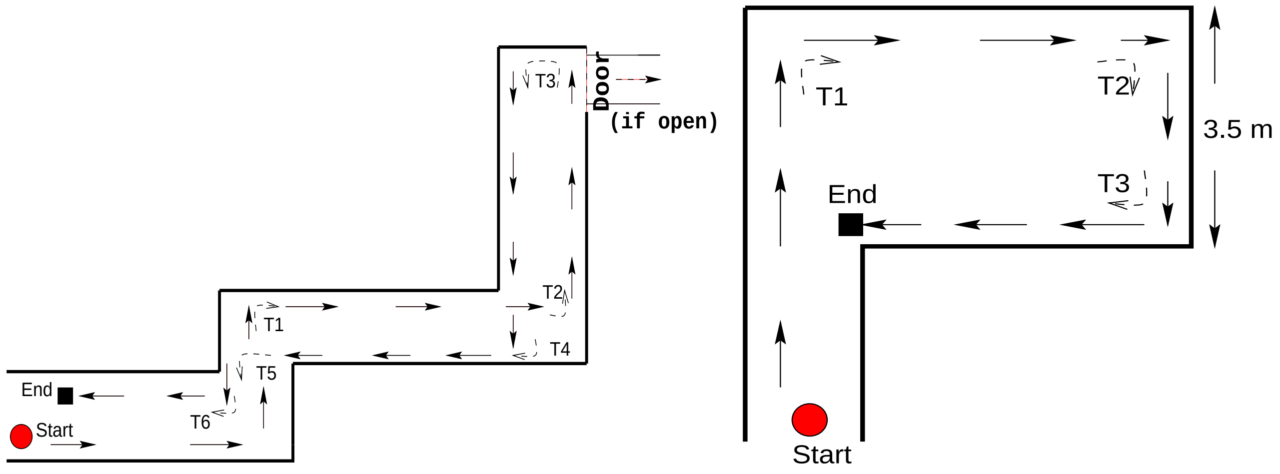

Figure 5: The figures show the path that the robot traversed in unknown environments. Solid arrows show the path traversed by robot. The turns are marked by dotted arrows and Ti denotes the i^{th} turn.

using sonars. However, we have shown that it is possible to obtain a rough map (openings, corners etc) of the environment without using expensive vision systems. Moreover our method is quite fast. The training takes only two minutes and after training the robot is free to move around autonomously and is capable of recognizing features and responding to them in real time.

As stated in section 2, the sonars have a blind spot and human intervention becomes necessary when the robot is going to hit a wall due to the blind spot. In the future we would like to use short range infrared rangers (which are available on the Fetch robot) to react to features nearby and then blend the data from sonars and Infrared rangers to make the robot completely autonomous. We also plan to investigate other exploration strategies which provide the robot with some memory so that it avoids going to sites which it has already visited.

Acknowledgments

The authors thank S. Finney for help with sonars and K. Barrett for help in migrating the code on-board. The authors also thank J. Aspinall and P. Pook for support and creative ideas. This research was sponsored by IS Robotics, Inc. under performance of DARPA ITO contract F30602-96-C-0280 during the first author's summer internship in 1999. The Fetch II robot development was sponsored by ONR and the Naval EOD Technology Division.

References

[1] George Bekey and Kenneth Y. Goldberg. *Neural Networks in robotics.* Kluwer Academic, Boston, 1993.

[2] R.A. Brooks. A robust layered control system for a mobile robot. *IEEE Journal of Robotics and Automation*, 2(1):14–23, March 1986.

[3] J. Cao and E. Hall. Sensor fusion for the navigation of autonomous guided vehicle using neural networks. In *SPIE International Conference*, November 1998.

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

[5] D. Galles. Map building and following using teleo-reactive trees. Technical report, Stanford University Robotics Laboratory, 1993.

[6] Frank Kirchner and Joachim Hertzberg. A prototype study of an auonomous robot platform for sewerage system maintenance. *Autonomous Robots*, 1997.

[7] Kurt Konolige. Improved occupancy grids for map building. *Autonomous Robots*, 4:351–367, 1997.

[8] B. Kuipers, R. Froom, W.K. Lee, and D. Pierce. The semantic hierarchy in robot learning. *Robot Learning*, 1993.

[9] C. Kunz, T. Willeke, and I. R. Nourbakhsh. Automatic mapping of dynamic office environments. *Autonomous Robots*, 7(2), 1999.

[10] J. J. Leonard and H. F. Durrant-Whyte. Directed sonar sensing for mobile robot navigation. Kluwer Academic Publishers, 1992.

[11] Carlos Puchol. Achieving high performance sonar-based wall following. In *International Symposium on Artificial Intelligence in Real-Time Control*, 1994.

[12] R. Simmons and S. Koenig. Probabilistic robot navigation in partially observable environments. In *Proceedings of the Fourteenth International Joint Conference on Artificial Intelligence*, pages 1080–1087, 1995.

[13] S. Thrun. Bayesian landmark learning for mobile robot localization. *Machine Learning*, 33(1), 1998.

[14] K. Warwick, G.W. Irwin, and K.J. Hunt. *Neural Networks for Control and Systems.* P. Peregrinus on behalf of IEE, London, 1992.

[15] J. Weng and S. Chen. Incremental learning for vision-based navigation. In *International Conference on Pattern Recognition*, pages 44–49, August 1996.