Natural Landmark Recognition using Neural Networks for Autonomous Vacuuming Robots

Sylvia C. Wong, George G. Coghill and Bruce A. MacDonald
Department of Electrical and Electronic Engineering
University of Auckland, Private Bag 92019,
Auckland, New Zealand
{*s.wong, g.coghill, b.macdonald*}*@auckland.ac.nz*

Abstract

Two types of neural networks were trained and tested on a real robot for a natural landmark recognition task. The neural networks investigated were the multilayer perceptron (MLP) and learning vector quantisation (LVQ). The intended application is for autonomous vacuuming robots in completely unknown indoor environments, using a novel topological world model and region filling algorithm. A topological world model based on natural landmarks is built incrementally while the robot systematically cleans the environment. The implementation of this world model depends on robust and accurate recognition of natural landmarks. Both types of neural network were found to be able to successfully recognise the natural landmarks selected.

Keywords: neural networks, vacuum cleaning, mobile robots, landmark-based representation

1 Introduction

Autonomous domestic vacuum cleaners will be welcomed in most homes as vacuuming is a chore that provides little intrinsic satisfaction. For a proper clean, the robot has to completely cover all the floor area in a room. A map of the environment ensures that no area is missed by the robot. Cao [2] first recognised the different requirements of map construction for point-to-point navigation and region filling navigation. In point-to-point navigation, a map of the environment is built in an exploration phase to expose all path segments. This map is then used for planning optimal paths between points during navigation. In region filling navigation, all exposed floor area is to be covered. As a vacuum cleaning robot has to visit all exposed floor area each time it vacuums, it is more convenient and efficient to build the map while vacuuming. Even though a separate exploration phase is superfluous, vacuuming navigation may be improved after the first vacuuming episode by using the previously built world model. The world model also needs to contain information specific to vacuuming, for example, the area cleaned.

Since Cao various people have worked on the problem of world modelling for autonomous vacuuming robots. Hofner [4] has developed a path planner for cleaning robots working in public areas. However, the system relies on *a priori* knowledge of the environment. For domestic vacuum cleaners, it is preferable that users do not have to input full floor plans of their houses.

González [3] uses an occupancy grid to describe the environment. The floor area, represented in a grid map, is subdivided into rectangular regions. As a result, it can only handle rectangular obstacles. Lang [8] does not use any standard world modelling methods. The robot first follows the outermost walls of the room. At the same time, coordinates of points along the walls are remembered and marked as ends of cleaning tracks. Afterwards these tracks are followed to completely cover the area within the outermost walls. The problem with this approach is that the points are remembered as absolute coordinates only and dead reckoning is used to locate and follow these tracks. As a result, the system is very susceptible to odometry error.

Recognising the need for localisation to reduce the odometry error, Ulrich [11] has proposed a navigation strategy that allows the robot to recalibrate its odometry while cleaning. First the robot maps the border of the room and records the length of walls and compass readings along them. This is stored in an occupancy grid. Once the border is mapped, the robot chooses the direction with the most uncleaned pixels and moves in a straight line in that direction until reaching a wall or an obstacle. If the trajectory ends with a wall, the robot can recalibrate its estimation of its orientation and one x-y component of its odometry. To keep a good estimation of its odometry, the robot chooses paths that end successively with perpendicular walls. With this navigation strategy, odometry estimation can be kept accurate without the use of a complicated localisation module. The only problem is that the cleaning paths generated are highly redundant as already cleaned surfaces are covered repeatedly to reach the opposite wall for recalibration.

Landmark-based representation methods provide an alternative to using only metrical information to construct world models of environments [6]. A topological model is created using natural landmarks as nodes, and connectivity between landmarks as edges. This way of representing the environment is very similar to that used by humans and other animals [10]. For example, we would say that the supermarket is to the right of the bookstore and across the road from the restaurant. In contrast to metrical methods, these landmark-based methods are robust against sensor and actuator noise. For example Zimmer [13] successfully implemented a world modelling and localisation system using neural networks on a very low budget mobile robot platform with only touch and light sensors. For domestic autonomous vacuum cleaners, the cost and size of mobile robot platforms should be kept at a minimum. Therefore, landmark-based methods are suitable for world modelling in this application.

One way to recognise natural landmarks is to create explicit rules for the types of landmark selected [6, 9]. For example, when a robot is in a concave corner, range sensors should return short range readings for two sides of the robot. In these representations, a description of the environment only exists around the landmarks. All other areas are described only as paths between landmarks. An alternative to predefined landmarks is to use a clustering algorithm, for example self-organising maps, to partition the environment into regions [7, 13]. A node is assigned to each region, and neighbouring regions are connected together with edges.

Section 2 of this paper presents a landmark-based world model for region filling navigation. Section 3 then describes and compares two types of neural networks that have been implemented for recognising natural landmarks.

2 Topological World Model

In the proposed world model, the environment to be vacuumed is described topologically using concave and convex corners as landmarks. These landmarks are connected with travelling paths. Figure 1 shows how a typical environment is represented in the proposed world model. The symbols used in all the figures in this section are shown in figure 2. This world model can be stored as a bidirectional graph, where landmarks are the nodes and travel paths are the edges. The location of the robot at any time is given by the node it is at or the edge it is on.

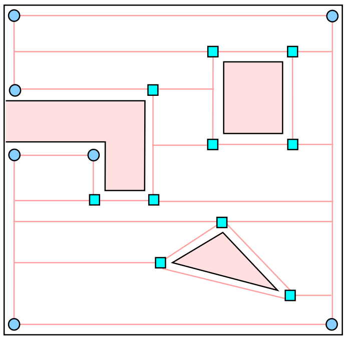

Figure 1: Topological world model for a typical environment.

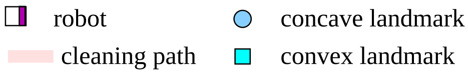

Figure 2: Symbols used in the figures.

The topological world model of an environment is constructed incrementally while the robot is cleaning. Starting at a corner of the room, the robot moves in a zigzag pattern down the room until it is completely covered (figure 3). The distance between tracks shown in figure 3 is for visualisation purposes only. In a real implementation, there will be overlap between tracks to ensure a more thorough clean.

The algorithm assumes cleaning starts at a corner of a room (figure 4(a)). As shown in the figure, the initial node has two edges. These two edges represent the two unexplored directions of the initial landmark. Nodes, and corresponding edges, are added to the world model whenever the robot discovers a new landmark. Figure 4(b) depicts the robot discovering its second landmark while cleaning the first track after leaving the starting corner. A new node rep-

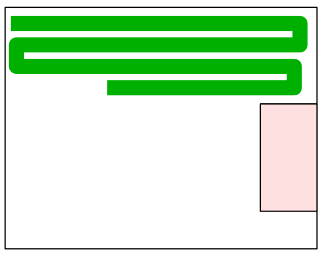

Figure 3: Vacuuming a room systematically using a zigzag pattern.

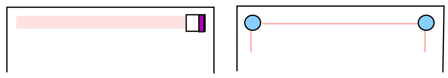

Figure 4: (a) Initial world model. (b) Addition of a new node at the end of the first track.

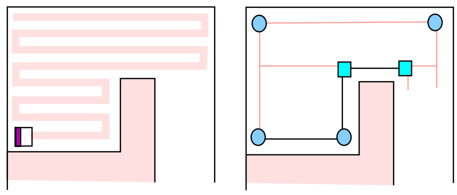

Figure 5: Search for an uncleaned region when there are no more uncleaned tracks left at current position. The path returned by breath-first search is highlighted.

resenting the new landmark is added to the world model. This new node is connected to the initial node. Also added is a new edge representing the new unexplored direction. Because of the way the world model is constructed, information about whether an environment is completely covered is stored in the topology of the world model. Any unexplored, thus uncleaned, direction is represented by an edge that is not connected at both ends. Only when all edges are properly connected is an environment completely cleaned. In this case, the environment is also fully mapped.

If the robot is at a concave corner and there are no more free tracks in the current position (figure 5), a breath-first search is carried out on the world model to find the closest incompletely connected edge. The search is done to find any uncleaned directions in the environment. If the search fails to find an edge that is not connected at both ends, the room is completely covered. If the search finds such an edge, such as the case shown in figure 5, the path returned by the search is followed using a wall following strategy.

The zigzag cleaning pattern tries to cover the room from top to bottom. Sometimes, a new region on the 'top side' is discovered while the robot is at a convex corner (figure 6(a)). In this case, the robot cleans this newly found region first before moving on to clean the rest of the room (figure 6(b)).

Free-standing obstacles are handled by first covering the area on one side of the obstacle and then the other side. The robot resumes its normal top-to-bottom zigzag cleaning pattern only after both sides of the free-standing obstacles are covered. Figure 7 depicts such a situation. In figure 7(a), the robot has nearly finished cleaning one side of an obstacle. In figure 7(b), the robot has moved from the first to the

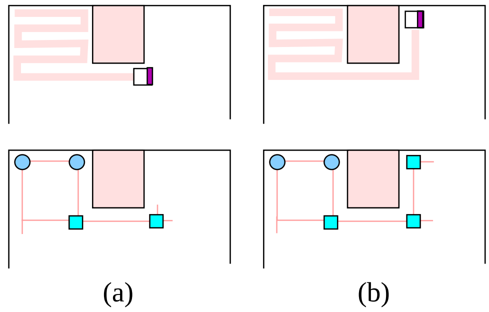

Figure 6: (a) A new 'top side' region is discovered. (b) The newly discovered region is cleaned first.

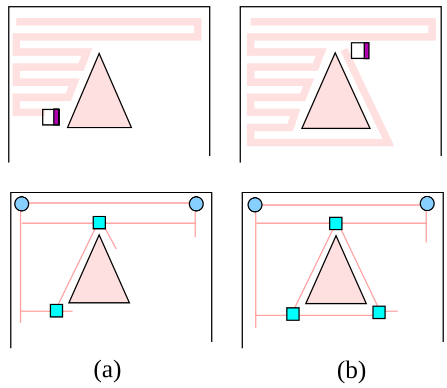

Figure 7: (a) Cleaning one side of a free standing obstacle. (b) Moved to the other side of the obstacle to start cleaning the other side from the top.

second side of the obstacle, and is starting to clean, in a zigzag pattern, the second side.

The topological world model and region filling algorithm proposed in this section have been successfully implemented in simulation. A more in depth description of the two can be found in [12].

3 Recognition of natural landmarks

The implementation of the world model described in section 2 depends on accurate and robust recognition of concave and convex corners. Two different types of neural networks, multilayer perceptron (MLP) [1] and learning vector quantisation (LVQ) [5], were implemented for the task of recognising when the robot is at either of the two chosen types of landmarks. The recognition system was trained and tested on a mobile robot in various laboratories at the university. A picture of the mobile robot used is shown in figure 8. It has a single ultrasonic transducer on the top. A stepper motor is used to rotate the transducer around to detect obstacles from all directions of the mobile robot. A vector of 48 readings is returned

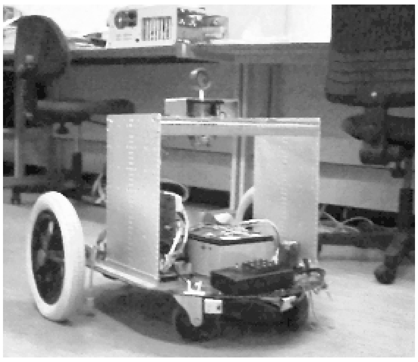

Figure 8: Maxifander in a university laboratory.

from a single 360-degree scan.

Sonar range data collected were categorised into three groups — concave corners, convex corners and everything else. The tasks of the two neural networks were thus to learn this classification and to predict which group a new sonar range vector belonged to.

3.1 Preprocessing

Both concave and convex corners are local features. To reduce the influence of far away objects on the recognition process, the measured range data was cut if it was over a certain threshold. In the case of MLP, the range vector was also normalised so its value was within the range of the output activation function of the neural network.

To make the classification independent of the orientation of the robot, each vector of 48 range readings was virtually rotated into the orientation most occupied by obstacles [7]. After this virtual rotation, index 0 of the vector would always be pointing towards the direction where the sonar range sensor measured the shortest distances. This most occupied orientation was calculated using the following equation:

$$\vec{d}_{\mathrm{MOO}} = \frac{1}{n}\sum_{i=1}^{n} \vec{d}_i \qquad (1)$$

where $n = 48$ is the number of readings in each vector, $\vec{d}_i$ is a vector originating from the centre of the robot denoting sonar sensor reading for direction i, and $\vec{d}_{\mathrm{MOO}}$ is the vector for the most occupied orientation.

Using this equation, all 48 points in the vector were used to calculate the most occupied orientation. This made the process of finding the most occupied orientation more robust to noise than if only the shortest range in the vector was used.

	MLP	LVQ
concave	90%	90%
convex	60%	100%
others	95%	88%

Table 1: Accuracy achieved on the test set for MLP and LVQ.

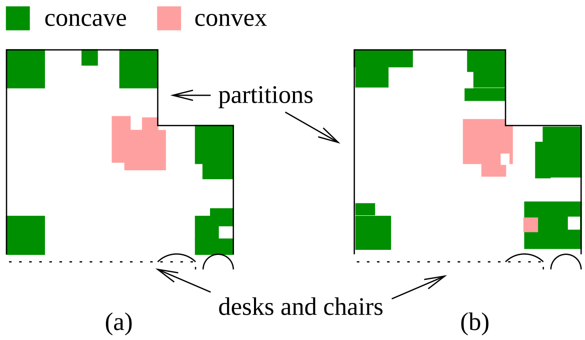

Figure 9: Classification with (a) MLP (b) LVQ.

3.2 Multilayer Perceptron

Backpropagation with a momentum term and flat spot elimination was used to train a three layer feed-forward neural network. Networks with various configurations were trained multiple times on this recognition problem using the training set to find the network that achieved the lowest mean square error on the test set. The parameters that were varied in the networks were number of hidden layer neurons, learning rate and momentum term.

The lowest mean square error achieved was 0.0955 with 8 hidden neurons, a learning rate of 0.4 and a momentum term of 0.25. Classification accuracy achieved in the test set is shown in table 1, where accuracy is defined as the number of accurate predictions divided by the total number of samples in the test set. The result of classification using MLP on a test environment is shown in figure 9(a).

3.3 Learning Vector Quantisation

According to Kohonen, all the different LVQ variants should yield similar accuracy [5]. OLVQ1 was picked for this problem because of its fast training time. Networks with different numbers of neurons, or codebook vectors, distributed among the three classes were trained for 40 epochs. A network with 30 neurons yielded good results. The accuracy achieved is shown in table 1. The result of classification using LVQ on the test environment is shown in figure 9(b).

3.4 Discussion

Supervised neural networks were chosen for this application because the landmark types to be recognised were predefined. The neural networks were thus trained to generate rules statistically to classify sonar range data. This is different from existing use of neural networks in landmark-based world models, where unsupervised neural networks partition environments into separate regions according to similarity of input sensory data [7, 13].

It can be seen from table 1 that the two neural networks gave different accuracy rates for the three categories to be classified. Despite the difference in accuracies, the resultant classification was quite similar (see figure 9). This is due to the fact that misclassification has mostly occurred at the boundaries between different zones. Misclassification at boundaries is insignificant because it does not affect the implementation of the world model proposed in section 2. This shows that accuracy alone is not a good indication on how well a neural network performs for the target application.

Also note that desks and chairs were not present in the training set. They were added to the test environment to see how well the two networks generalised on unseen objects. It can be seen from figure 9 that both networks generalised well to handle the desks and chairs. It can be concluded that both types of neural network are suitable for this application.

4 Conclusions

In this paper, a novel topological world model and region filling algorithm for autonomous vacuuming robots is presented. Humans, and other animals, model their environment using topologies of landmarks. This type of representation does not rely solely on an absolute coordinate system. Therefore, a coherent world model can be constructed with noisy sensor data as long as the landmarks are properly recognised. This is especially useful in completely unknown environment where no *a priori* world model is available for localisation.

Landmark recognition is central to the implementation of the proposed world model on a real robot. This paper shows that the natural landmarks selected can be easily recognised by a neural network. Even though the work presented here is in its early stages, it still shows that it is feasible to carry out region filling navigation using a landmark-based world model with obstacles of various shapes.

References

[1] Christopher M. Bishop. *Neural Networks for Pattern Recognition*. Oxford University Press, 1995.

[2] Zuo Liang Cao, Yuyu Huang, and Ernest L. Hall. Region filling operations with random obstacle avoidance for mobile robots. *Journal of Robotic Systems*, 5(2):87–102, 1988.

[3] E. González, A. Suárez, C. Moreno, and F. Artigue. Complementary regions: a surface filling algorithm. In *Proceedings IEEE International Conference on Robotics and Automation*, pages 909–914, Minneapolis, Minnesota, April 1996.

[4] Christian Hofner and Günther Schmidt. Path planning and guidance techniques for an autonomous mobile cleaning robot. *Robotics and Autonomous Systems*, 14:199–212, 1995.

[5] Teuvo Kohonen. *Self-Organizing Maps*. Springer, second edition, 1997.

[6] Benjamin J. Kuipers and Yung-Tai Byun. A robust, qualitative method for robot spatial learning. In *Proceedings Seventh National Conference on Artificial Intelligence AAAI-88*, pages 774–779, St Paul, Minnesota, August 1988.

[7] Andreas Kurz. Constructing maps for mobile robot navigation based on ultrasonic range data. *IEEE Transactions on Systems, Man and Cybernetics - Part B: Cybernetics*, 26(2):233–242, April 1996.

[8] Sherman Y. T. Lang and Bing-Yung Chee. Coordination of behaviours for mobile robot floor cleaning. In *Proceedings Intl. Conference on Intelligent Robots and Systems*, pages 1236–1241, Victoria, B.C., Canada, October 1998.

[9] Maja J. Mataric. Integration of representation into goal-driven behavior-based robots. *IEEE Transactions on Robotics and Automation*, 8(3):304–312, June 1992.

[10] J. O'Keefe and J. Dostrovsky. The hippocampus as a spatial map: Preliminary evidence from unit activity in the freely moving rat. *Brain Research*, 34:171–175, 1971.

[11] Iwan R. Ulrich, Francesco Mondada, and J. D. Nicoud. Autonomous vacuum cleaner. *Robotics and Autonomous Systems*, 19(3–4):233–245, March 1997.

[12] Sylvia Wong, George Coghill, and Bruce MacDonald. Landmark-based world model for autonomous vacuuming robots. In *Proceedings International ICSC Congress on Intelligent Sys-*

tems and Applications (ISA), Wollongong, Australia, 2000.

[13] Uwe R. Zimmer. Robust world-modelling and navigation in a real world. *Neurocomputing*, 13(2–4):247–260, 1996.