Self-Localisation and Autonomous Navigation by a Mobile Robot

Tom Duckett and Ulrich Nehmzow
Department of Computer Science
University of Manchester
Manchester M13 9PL
United Kingdom
`duckettt@cs.man.ac.uk, ulrich@cs.man.ac.uk`

26.3.99

Abstract

This paper provides an overview of a three year project conducted on the subject of self-localisation and autonomous navigation by a mobile robot. The research was carried out in two stages:

Self-Localisation. *Due to the fundamental unreliability of dead reckoning, landmark-based methods were investigated. Several related issues were addressed, including performance evaluation, replication and comparison of existing work, and the development of a novel localisation system based on the results obtained.*

Autonomous Navigation. *This work included the development of novel mechanisms for exploration and map building, as well as validating the previous work on self-localisation, in a complete, navigating mobile robot. Quantitative assessment of the different competences required for navigation was also carried out.*

1 Introduction

In recent years, there has been a great deal of interest in the topic of navigation by mobile robots. Reliable navigation systems are essential if mobile robots are to become widely used in industrial and domestic applications. However, while assembly robots are now ubiquitous in modern production plants, mobile robots are by comparison rarely found outside of scientific research institutions at the present time.

Perhaps the most fundamental competence required for autonomous navigation is that of self-localisation (location recognition). Without the ability to identify locations reliably, a mobile robot will inevitably become lost, and therefore unable to move purposefully between locations. We contend that the large scale use of mobile robots has been prevented so far, at least in part, by a lack of robust mechanisms for self-localisation.

In this paper, we provide an overview of a three year PhD project carried out on the subject of mobile robot self-localisation and its application to the problem of autonomous mobile robot navigation. The work was carried out in two main stages, as follows:

1. *Mobile Robot Self-Localisation.* Due to the fundamental unreliability of self-localisation using dead reckoning, landmark-based methods were investigated. Several important issues were addressed, including quantitative assessment of performance, conducting empirical comparisons of existing approaches – including replication of the work of other researchers, and the development of a novel localisation system based on the results obtained.

2. *Autonomous Mobile Robot Navigation.* A complete navigation system for a Nomad 200 mobile robot was developed, including novel mechanisms for exploration and map building, together with the previous work on self-localisation. Experiments were conducted in which complex, untreated environments of several hundred m^2 were successfully mapped and navigated by the robot, requiring human intervention only to specify a goal location.

Proc. TIMR 99 "Towards Intelligent Mobile Robots", Bristol 1999.
Technical Report Series, Department of Computer Science, Manchester University, ISSN 1361 - 6161.
Report number UMCS-99-3-1. http://www.cs.man.ac.uk/cstechrep/titles99.html

The main achievements of the project were the development of quantitative methods for evaluating navigational competence, the comparative analysis of existing localisation systems, the development of novel mechanisms for exploration, map building and self-localisation by a mobile robot, and the ability to navigate autonomously in complex, untreated environments of middle-scale dimensions[1].

2 Related Work

A wide range of navigation systems have been proposed for mobile robots, and many laboratory prototypes have been developed. Some of these systems have been validated in real world environments, generally consisting of untreated, enclosed areas within public buildings (see e.g., Kortenkamp *et al* [9] for further information). Lack of space prevents a detailed review here; instead, we discuss the relative merits of some representative approaches.

Yamauchi *et al* ([19]) developed a mobile robot which explores and builds its own maps of uncharted territory, using a a global occupancy grid[2] to represent the environment. A disadvantage of this approach is that it depends critically upon accurate laser sensors and precisely corrected odometry, because exact position information is needed to update grid-based maps. Also, the system requires knowledge of the robot's initial location for self-localisation, so the robot cannot recover from becoming lost.

Burgard *et al* ([2]) developed a tour-guide robot which incorporates many state-of-the-art techniques for mobile robot navigation using metric maps, e.g., employing Markov localisation and the use of an "entropy filter" to detect the presence of humans within the robot's sensory range. The work also integrates other areas of AI research, including high-level problem solving and human-robot interaction. This robot has successfully operated in a busy museum environment, though the map required for this task was constructed by the designers using accurate distance measurements taken by hand.

In general, metric maps require large amounts of computation — both of the approaches described above require radio communication with external processors — and may often be over-specified for the navigational task in hand. While effective systems have been built using metric maps, topological maps have, by nature of their compactness, the potential for representing environments which are several orders of magnitude larger than those which can be tractably navigated using metric maps.

One of the most popular approaches for topological navigation is that of hidden Markov models, and their extension to partially-observable Markov decision processes (e.g., Koenig and Simmons [8], Hertzberg and Kirchner [7]). Here, the robot maintains a probability distribution over a set of discrete locations, known as the robot's "belief". Similarly, landmarks and actions are typically defined according to a set of human-defined categories, e.g., possible landmarks might be 'doors', 'junctions', *etc.*, and possible actions might be 'North', 'West', *etc.* A disadvantage of this approach is that it requires pre-installation of a topological map by the system designers. Some important work has been done on extending the Baum-Welch algorithm, which is used for constructing hidden Markov models, in order to allow acquisition of the map by the robot itself (Shatkay [14]). However, this algorithm is computationally expensive, subject to local minima and requires a large amount of data which has to be collected first by the robot under manual control.

3 The Robot

The Nomad 200 robot used in this project was equipped with sixteen ultrasonic range-finding sensors (range up to $6.5m$) and sixteen infrared sensors (range up to $60cm$) spaced at 22.5 degree intervals around its turret (see figure 1). A flux gate compass was used to determine the robot's heading, and also to keep the turret, and therefore the sensors, at a constant orientation during exploration. All computation was carried out in real-time on board the robot's 486 processor.

[1] Middle-scale navigation involves leaving the sensory range of a 'home' or 'goal' location, and navigating within the wider environment [16]. In the case of our Nomad 200 robot, we mean navigation over several hundred m^2.

[2] An occupancy grid is a matrix representation of the robot's environment, where each cell contains a probability of occupancy. The cells are typically labelled as 'occupied', 'empty' or 'unknown' depending on the probability value.

Figure 1: The Nomad 200 Mobile Robot *FortyTwo.*

4 Mobile Robot Self-Localisation

Mobile robot self-localisation is possible using artificial landmarks such as beacons. However, modifying the robot's environment is costly and results in a lack of flexibility in dealing with environments that change over time. The ability to localise autonomously, i.e., using only the robot's own sensory impressions of the environment, is therefore desirable, and a wide variety of different mechanisms have been proposed towards this end (see Borenstein *et al* [1] for a detailed survey).

However, while successful laboratory prototypes have been developed, progress towards autonomous mobile robots which can localise reliably in real world environments has been slower, due to the numerous scientific challenges which must be overcome, e.g., perceptual aliasing[3], sensor noise and other unpredictable variations such as moving people.

There is a second reason, as we see it, why progress here has been slow; objective comparison of the different approaches is largely impossible, because researchers use different robots, apply different quality measures and conduct different experiments to validate their work. Without the ability to compare results, it is perhaps not surprising that overcoming the scientific challenges has been difficult.

This stage of the project therefore investigated quantitative methods which would facilitate the comparative analysis of mobile robot self-localisation systems. An empirical comparison of previous work on self-localisation was conducted, including replication of work done by other researchers. In addition, the results of this analysis were used to guide the development of a novel self-localisation system. Experiments conducted in a variety of middle-scale environments demonstrated that the performance of the new system was significantly better than that of the other approaches investigated.

4.1 Measurement of Performance

The question asked in this piece of work was: *Which landmark recognition mechanisms are best suited for self-localisation by a mobile robot?* The aim was to be able to evaluate different self-localisation systems on the same robot and in the same environments, so that the relative performance of the various approaches could be determined. Two key problems were addressed (see [4, 5] for full details):

- *Measurement of Localisation Quality.*
- *Tracking the Robot's True Location.*

To assess localisation quality, we used the following entropy-based statistic known as the *Uncertainty Coefficient of L given R*, which measures the extent to which the robot's response R (i.e., the location estimates produced by the localisation mechanism under investigation) predicts the robot's true location L. This is defined by

[3]This refers to the situation where several places are perceptually similar enough to be confused by the robot.

$$\begin{aligned} U(L \mid R) &\equiv \frac{H(L) - H(L \mid R)}{H(L)}, \\ H(L) &= -\sum_j p_{\bullet j} \ln p_{\bullet j}, \\ H(L \mid R) &= -\sum_{i,j} p_{ij} ln \frac{p_{ij}}{p_{i\bullet}}, \end{aligned}$$

where $p_{\bullet j} = \sum_i p_{ij}$, $p_{i\bullet} = \sum_j p_{ij}$, and p_{ij} refers to the probability of the robot's response being i when the robot's true location was j. A value of $U(L \mid R) = 1$ implies that the robot's response always predicts its true location, while $U(L \mid R) = 0$ implies that the response never predicts the true location.

In order to calculate the above statistic, some means of tracking the "true" location L of the robot was required. In initial experiments, this was measured by hand, but this process was costly and prone to human error. We therefore developed a novel mechanism for position tracking which was based on retrospectively corrected odometer data.

Our approach has the advantage that is requires no interpretation of the "correct" response by the robot (Cartesian or otherwise), and no optimum standard has to be established by the observer. In addition, only approximate location tracking is required in order to evaluate localisation quality.

4.2 Comparative Study of Landmark Recognition Systems

In the traditional approach to self-localisation, a Cartesian position estimate produced using odometry is continuously recalibrated using sightings on known environmental features or "landmarks" previously encoded in the robot's map (see e.g., Lee [11]). The reliance on *a priori* knowledge of the robot's position means that this approach cannot solve the more general problem of *global self-localisation*, in which the robot must be able to relocalise even after becoming lost. Any autonomous robot operating in the real world is bound to become lost sooner or later, regardless of the accuracy of its sensors, especially in environments which are subject to changes over time.

In this work, we therefore considered systems which use landmarks rather than dead reckoning for localisation. In particular, we concentrated on systems which determine their *own* landmarks, including self-organising neural networks, e.g., Kohonen, ART ([3]) and Restricted Coulomb Energy ([10]) networks, and self-detailing schemes, e.g., occupancy grids ([18]). An alternative would be to use human-defined landmarks (referred to as "natural landmarks" by Borenstein *et al* [1]) such as ceiling lights, doors, *etc.*, but this approach can only be used in environments which contain the appropriate objects, and also suffers from the "correspondence problem" of matching sensor readings to the human-defined categories.

Environment	**RCE Classifier**	**ART2 Network**	**Occupancy Grids**
T-shaped hallway	0.554	0.573	0.732
Conference room	0.552	0.669	0.879
L-shaped corridor	0.460	0.538	0.644
Long straight corridor	0.220	0.265	0.487
Mean $U(L \mid R)$	0.447	0.511	0.686

Table 1: Localisation Quality $U(L \mid R)$. In this example, we compared the performance of three different classifier mechanisms for recognising landmarks, namely an ART network [3], an RCE classifier [10] and an algorithm for matching local occupancy grids [18].

Of the systems considered in the study, we found that the highest level of localisation quality was obtained by the localisation mechanism developed by Yamauchi and Langley ([18]). Some example results are given in table 1. In this system, each place in a topological map is represented using a *local* occupancy grid, i.e., there is no global grid model as in other systems ([2, 15, 19]). During localisation, a recognition grid is constructed from the robot's immediate sensor readings. A hill climbing procedure is then used to compare the recognition grid with the previously stored grids in the robot's map, searching the space of possible translations and rotations to find the best match.

However, the main disadvantage of this localisation system is that it is computationally highly expensive, taking 5 minutes to relocalise on a Decstation 3100 using 43 stored grids, and therefore unsuitable for localisation in real-time on our Nomad 200 robot. In addition, while this approach reduces perceptual aliasing by using as much sensor information as possible, it can never be guaranteed to eliminate the problem completely. We found that this system was still not reliable enough to enable self-localisation in real world environments such as corridors, which contain no guarantee of unique perceptual cues.

4.3 Novel Self-Localisation System

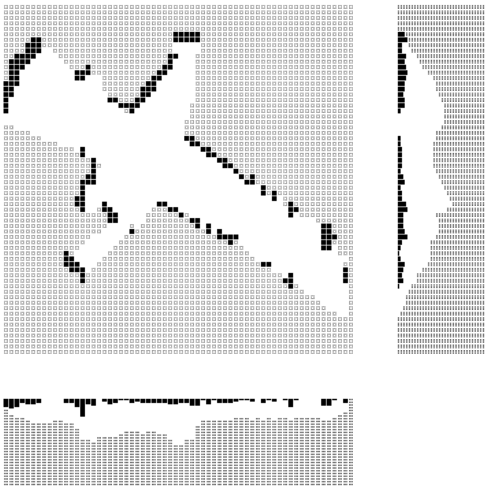

Figure 2: Example Occupancy Grid and Histograms. Occupied cells are shown in black, empty cells in white and unknown cells in grey.

To overcome the high computational requirements of Yamauchi and Langley's approach, we introduced a much faster matching procedure, taking less than 1 second to relocalise on our robot's 486 processor using 43 stored grids. Here, we do not store or compare the actual occupancy grids. Instead, a compass is first used to remove the problem of self-orientation, then each occupancy grid is reduced to a pair of histograms, one in x and one in y (see figure 2). Location recognition then consists of convolving the respective x and y histograms for the current sensor data with those stored previously in the map to find the best matching place (as in Weiss and Puttkamer [17]). Using histograms requires only a fraction of the storage and processing required to compare whole occupancy grids, with little or no loss in overall information content (see [6] for full details).

However, due to problems such as perceptual aliasing and sensor noise, the best matching place is not guaranteed to be the true location of the robot. To overcome these problems, we developed an iterative localisation algorithm (see [5, 6] for details). Here, the robot maintains a multi-modal probability distribution over a set of possible location "hypotheses", reflecting the robot's confidence in its estimate of its position. As the robot moves through its environment, this probability distribution is updated on the basis of the new sensory evidence and the observed motion of the robot since the previous iteration, and multiple Kalman filters are used to adjust the corresponding location estimates. Using this method, the robot can relocalise successfully even after becoming completely lost in environments containing no unique landmarks.

5 Autonomous Mobile Robot Navigation

Current industrial mobile robots are typically only able to navigate in specially treated environments, e.g., by following paths marked by induction loops or beacons, or by using a pre-installed map supplied by the system designer. By contrast, we have developed an adaptable, autonomous navigation system, i.e., the robot has all components necessary for its operation on board, and requires no modifications to its environment to operate successfully.

5.1 Map Building

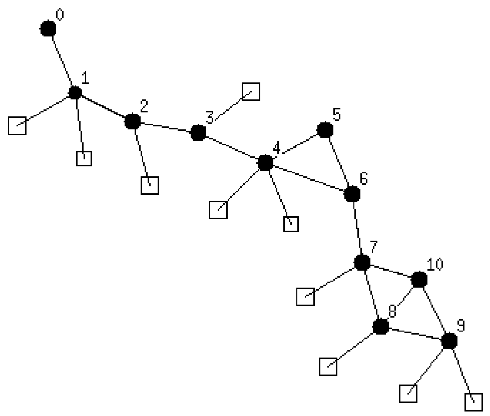

Figure 3: Example of Topological Map Building. Places presumed to exist ("predicted") but not yet visited by the robot are shown by squares. Places actually visited ("confirmed") by the robot are shown by circles.

The robot builds its own topological map through a process of autonomous, map-based exploration. In other words, the robot is able to identify regions of unexplored territory, navigate to the identified areas using its own map, and incrementally update its representation of the environment at the same time. The topological map contains two different types of places:

- *Predicted.* Places presumed to exist but not yet visited by the robot.
- *Confirmed.* Places actually visited by the robot.

The exploration strategy used to acquire the maps consists of continuously trying to expand the territory already charted by the robot, using a neural network to add new 'predicted' places to the map (see section 5.2), and then attempting to verify these predictions by moving into new territory (see figure 3).

From its initial location, the robot adds the first set of 'predicted' places to the map. The robot then tries to navigate to the nearest 'predicted' place, resulting in that node either being replaced by a 'confirmed' place or being deleted from the map, depending on whether the robot can actually travel to a distinct new location in the environment without encountering any obstacles. Whenever another 'confirmed' place is added to the map, the neural network is used again to predict more new places. This process is repeated until all 'predicted' places in the map have either been visited by the robot or deleted.

5.2 Detection of Free Space

To explore an environment in the manner described above, some mechanism is required to add the new 'predicted' places to the map, i.e., to detect areas of unexplored territory. Individual range-finder readings are not well suited for this purpose in real world environments, because of occlusions caused by moving people, and also problems such as sensor noise, cross-talk and specular reflections. Instead, we used an artificial neural network to learn the concept of "free space", combining noisy information obtained from many sensor readings. All of the data required for training the network, including the ability to travel in a particular direction, was collected by the robot itself, thus avoiding the need for manually labelling

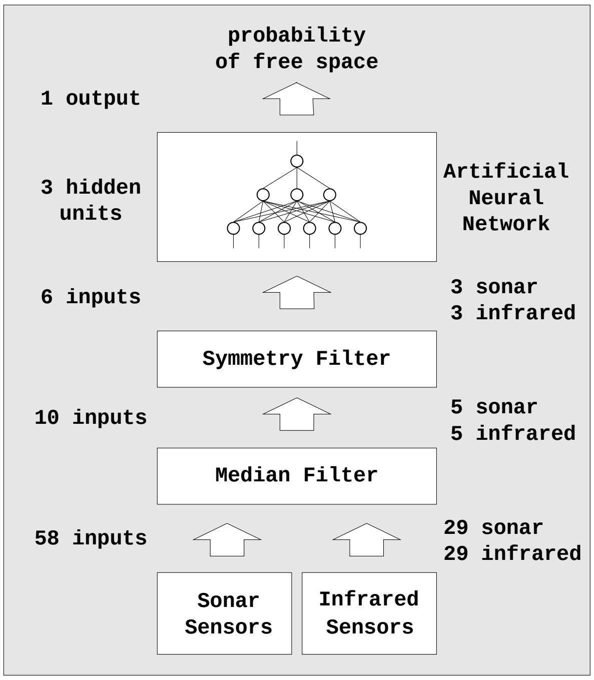

Figure 4: Architecture for Free Space Detection. Two levels of pre-processing are applied to the sensory input, then a neural network is used to predict the presence or absence of free space in a given direction.

the training examples with the desired output categories (e.g., as in the concept learning mechanism described by Mahadevan *et al* [12]).

The sensing strategy used for predicting free space consisted of rotating the robot's turret to obtain a detailed scan of sonar and infrared readings. For data collection, a scan was first taken, then the robot attempted to move as far as possible in an arbitrary direction until an obstacle was detected, recording both the sensor readings from the scan and the distance travelled. A subset of 58 of the sensor readings, centred around the direction of travel, was then used as input to the prediction mechanism (see figure 4). Two pre-processing functions were applied to the sensor readings before input to the neural network:

1. *Median Filter.* Groups of 5 or 6 adjacent sensor readings were combined to produce a single reading by taking their median value, thus reducing the effects of specular reflections (where the sonar readings are "too big") and occlusions caused by moving people (where the sonar readings are "too small").

2. *Symmetry Filter.* This function was used to exploit the bilateral symmetry inherent in the problem of free space prediction; e.g., the left-most sonar reading was combined with the right-most sonar reading by taking the minimum of the two values (i.e., the nearest of the two obstacles detected).

Finally, a fully connected feedforward network with 6 inputs, 3 hidden units and 1 output was trained to associate the sensory input in a given direction with the robot's ability to move by a prespecified distance[4] in that direction, as obtained from the data collected by the robot.

5.3 Way Finding

Dijkstra's algorithm was used to plan paths in the topological map, either to the nearest 'predicted' place during map building, or to a user-specified location after map building has been completed. The localisation mechanism described in section 4.3 was used to determine the robot's current location in the map. Low level motion control was carried out using previously acquired sensor-motor behaviours for obstacle avoidance and contour following ([13]).

[4] a value of $1m$ was used in these experiments

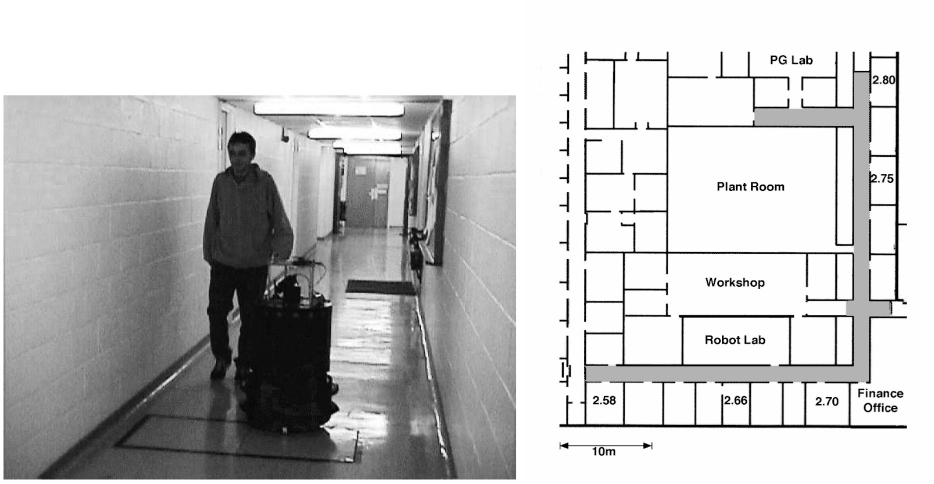

Figure 5: Corridor Environment. This is a busy public area containing few unique landmarks.

6 Results

Quantitative performance measures were used to evaluate the different competences used by the navigating mobile robot. One of the environments used for testing is shown in figure 5. This area is a busy thoroughfare affected by many transient changes such as moving people, doors opening and closing, *etc.*, and is subject to a high level of perceptual aliasing.

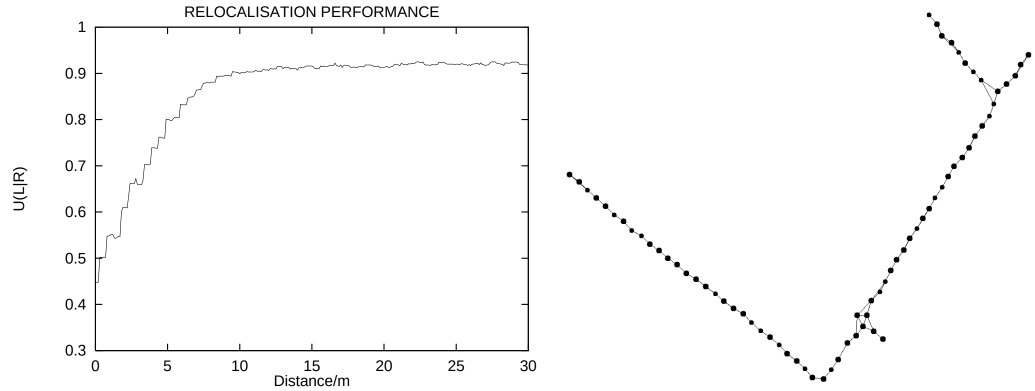

Figure 6: a.) Worst Case Relocalisation Performance. This shows localisation quality against distance travelled by the robot after becoming "lost". b.) Example Map Acquired by the Robot.

To assess self-localisation performance, we used the entropy-based performance measure described in section 4.1. In this experiment, the robot was first moved to a random location in the environment and its localisation mechanism re-initialised, so that the robot was "lost". The robot then used wall-following to accumulate new sensory evidence, and the uncertainty coefficient was measured over the distance travelled. The results shown in figure 6 indicate the ability of the system to relocalise under global uncertainty (i.e., after becoming lost).

To assess map building performance, we compared the performance of our map-based exploration strategy to that of exploration by wall-following, using the distance travelled to make a complete tour of the environment and then return to the starting location. The results showed that the performance of the map-based exploration strategy was comparable to that of wall-following in corridor environments, while having the additional benefit of being able to operate in "open" environments which cannot be explored by wall-following. An example of one of the topological maps acquired by the map-based exploration system is shown in figure 6.

To assess the overall navigational competence of the system, we developed quantitative measures of its *robustness*, that is, its ability to recover from error in navigating to a goal location (see [6]). In this experiment, the robot had to find 100 routes through the corridor environment, navigating from one user-chosen location to another. This was carried out with firstly with the robot knowing its initial location (the "control"), and then with a "lost" robot, where the robot's localisation was re-initialised at the start of each navigation attempt. The results showed a 100% success rate for the control, and a 92% success rate for the lost robot, indicating the effectiveness of our approach.

7 Summary

In summary, the main products of this research project were five-fold:

1. The development of quantitative measures for assessing navigational competence in mobile robots. In particular, quantitative measures of self-localisation, map building and navigational performance were considered.
2. Replication and a comparative study of existing work on mobile robot self-localisation, including work by other researchers, using quantitative performance measures.
3. A case study in "quantitative robotics"; that is, the application of quantitative measures to the design and testing of a novel mobile robot self-localisation system, which relocalises in middle-scale environments without requiring *a priori* knowledge of its position.
4. The development of novel techniques for autonomous exploration and map building of populated, middle-scale environments by a mobile robot, without requiring human intervention.
5. An untethered, self-governing mobile robot which builds its own maps through free exploration, identifies locations, determines routes and navigates reliably between user-specified locations in untreated, middle-scale environments, without requiring off-line processing.

8 Discussion and Future Work

The navigation system as it stands depends critically on a compass sense in order to self-localise and determine the robot's heading. On the Nomad 200, the compass sense was implemented by rotating the robot's turret at small speeds towards North, as indicated by a flux gate compass, thus smoothing out local variations in the magnetic field of the robot's environment. However, while this method is robust in dealing with minor fluctuations, it cannot be guaranteed to work in the presence of severe compass errors, e.g., due to ferrous building materials in the robot's environment. Future work will therefore need to examine the problem of self-orientation in more detail, perhaps integrating perceptual information from the robot's exteroceptive and proprioceptive sensors, as in the work already conducted on self-localisation.

A fundamental problem for any navigating robot is to build consistent maps in very large environments containing loops. Here, the robot's self-localisation mechanism needs to be able to recognise when the starting location has been reached once again after the robot has made a complete circuit of the environment, otherwise the robot will keep on adding new, duplicate copies of previously visited locations to the map *ad infinitum*. This cannot be achieved using exteroception, because of perceptual aliasing, but equally cannot be achieved using proprioception, because of the inevitable drift errors.

So far, navigating mobile robots have only been able to "close the loop" successfully by using accurate range-finding sensors and precisely corrected odometry (e.g., the system described by Thrun [15] has successfully mapped a circular route of $160m$ using laser range-finders and a high resolution metric map). However, this approach will inevitably fail once the size of the environment is increased beyond the limits of the robot's dead reckoning mechanism. We introduced a manual method of "closing the loop" in our previous work on self-localisation ([5]), using retrospectively corrected odometry, but have so far been unsuccessful in our attempts at autonomously mapping a circular environment of $190m$ using *FortyTwo*.

References

[1] J. Borenstein, H. R. Everett and L. Feng, *Navigating Mobile Robots: Systems and Techniques*, A. K. Peters Ltd., Wellesley, MA, 1996.

[2] W. Burgard, A.B. Cremers, D. Fox, D. Hähnel, G. Lakemeyer, D. Schulz, W. Steiner and S. Thrun, *The Interactive Museum Tour-Guide Robot*, Proc. AAAI Thirteenth National Conference on Artificial Intelligence, 1998.

[3] G. Carpenter and S. Grossberg, *ART2: Self-Organization of Stable Category Recognition Codes for Analog Input Patterns*, Applied Optics, Vol. 26, No. 23, pp. 4919-4930, 1987.

[4] T. Duckett and U. Nehmzow, *Quantitative Analysis of Mobile Robot Systems*, Proc. "Towards Intelligent Mobile Robots", British Conference on Mobile Robotics and "Scientific Methods in Mobile Robotics", Manchester UK, 1997. ISSN 1361-6161.

[5] T. Duckett and U. Nehmzow. *Mobile Robot Self-Localization and Measurement of Performance in Middle-Scale Environments*, J. Robotics and Autonomous Systems, No. 24, Vols. 1-2, pp. 57-69, 1998.

[6] T. Duckett and U. Nehmzow, *Knowing Your Place in Real World Environments*, submitted to EUROBOT-99, 1999.

[7] J. Hertzberg and F. Kirchner, *Landmark-Based Autonomous Navigation in Sewerage Pipes*, Proc. First Euromicro Workshop on Advanced Mobile Robots (EUROBOT'96), Los Alamitos, IEEE Computer Society Press, pp. 68-73, 1996.

[8] S. Koenig and R. G. Simmons, *A Robot Navigation Architecture Based on Partially Observable Markov Decision Process Models*, in [9], 1998.

[9] D. Kortenkamp, R. P. Bonasso and R. Murphy, eds. *Artificial Intelligence and Mobile Robots: Case Studies of Successful Robots*, AAAI Press, 1998.

[10] A. Kurz, *Constructing Maps for Mobile Robot Navigation Based on Ultrasonic Range Data*, IEEE Trans. Systems, Man, and Cybernetics - Part B: Cybernetics, Vol. 26, No. 2, April 1996.

[11] D. C. Lee, *The Map-Building and Exploration Strategies of a Simple, Sonar-Equipped Mobile Robot; An Experimental, Qualitative Evaluation*, Ph.D. thesis, University College London, England, 1995.

[12] S. Mahadevan, G. Theocharous and N. Khaleeli, *Rapid Concept Learning for Mobile Robots*, J. Autonomous Robots, Vol. 5, pp. 239-251, 1998.

[13] U. Nehmzow, *Autonomous Acquisition of Sensor-Motor Couplings in Robots*, Dept. Computer Science, University of Manchester, Technical Report Series UMCS-94-11-1, 1994.

[14] H. Shatkay, *Learning Models for Robot Navigation*, Ph.D. Thesis, Department of Computer Science, Brown University, 1998.

[15] S. Thrun, *Learning metric-topological maps for indoor mobile robot navigation*, Artificial Intelligence, Vol. 99, pp. 21-71, 1998.

[16] R. Wehner, M. Lehrer and W. R. Harvey, eds., *J. Experimental Biology*, Vol. 199, Special Issue on Animal Navigation, 1996.

[17] G. Weiss and E. v. Puttkamer, *A Map Based on Laserscans Without Geometric Interpretation*, Intelligent Autonomous Systems-4 (IAS-4), Karlsruhe, Germany, March 1995, pp. 403-407.

[18] B. Yamauchi and P. Langley, *Place Recognition in Dynamic Environments*, Journal of Robotics Systems, Special Issue on Mobile Robots, Vol, 14, No. 2, pp. 107-120, 1997.

[19] B. Yamauchi, A. Schultz and W. Adams, *Mobile Robot Exploration and Map-Building with Continuous Localization*, Proc. IEEE International Conference on Robotics and Automation, Leuven, Belgium, pp. 3715-3720, 1998.