Experiments in Evidence Based Localisation for a Mobile Robot

Tom Duckett and Ulrich Nehmzow
Department of Computer Science,
University of Manchester
Oxford Road,
Manchester M13 9PL,
England.
t.duckett@cs.man.ac.uk, u.nehmzow@cs.man.ac.uk

7/4/97

Abstract

This paper addresses the problem of localisation in autonomous mobile robot navigation, i.e., the task of identifying places after prior exploration and mapbuilding by the robot. In particular, the work is concerned with the more general problem of relocalisation without using past experience (i.e., knowing roughly where you are to start with), referred to here as the lost robot problem. *In the experiments presented here, the robot had to relocalise after being moved to a randomly chosen location, its sensors being disabled during that move. The robot therefore had no* a priori *knowledge of its position, and had to use current sensory perceptions and map knowledge alone to relocalise.*

A perception-based localisation method is presented which is resilient to the problem of perceptual aliasing (i.e., perceptual identity of distinct locations), and is capable of relocalising even in environments where no single place has a unique perceptual signature. During an exploration phase, the robot builds a map of its environment, using a self–organising neural network to cluster its perceptual space. The robot is then moved to an arbitrary location, where it will attempt to relocalise. By actively exploring, and accumulating evidence through the use of relative odometry between competing place memories, the robot is able to establish its location with respect to perceptual landmarks very quickly.

1 The Lost Robot Problem

In the field of mobile robot navigation, (self-) localisation refers to the task of establishing one's position with respect to the known world. In particular, this paper is concerned with the worst case of relocalisation where the robot has no idea of its approximate location, having become lost through some arbitrary circumstance. In this situation, the robot is unable to localise using past experience, i.e., no *a priori* estimate of position is available.

As we are interested in the use of autonomous mobile robots in unmodified environments, we avoid the use of pre-installed maps or external devices such as markers or beacons for position evaluation. To be fully autonomous, the robot must therefore rely on its own perceptions for dealing with the interrelated sub-problems of exploration, mapbuilding and relocalisation. The space of possible perceptions available to the robot for carrying out these tasks may be divided into two broad categories;

1. *Exteroception.* The robot's current perceptions of the outside world. A robot's exteroceptors may include range-finding sensors, tactile sensors, video cameras, etc.

2. *Proprioception.* The robot's perceptions of its own interactions with the world. In particular, odometry refers to the proprioceptor mechanism used for dead reckoning in mobile robots.

Proc. AISB workshop on "Spatial Reasoning in Mobile Robots and Animals", Manchester 1997.
Technical Report Series, Department of Computer Science, Manchester University, ISSN 1361 - 6161.
Report number UMCS-97-4-1. http://www.cs.man.ac.uk/csonly/cstechrep/titles97.html

A localisation system based solely on proprioception will be unsuitable for two reasons. Firstly, as no *a priori* information will be available to the robot when trying to relocalise, it will not be possible to initialise dead reckoning. Secondly, any proprioceptive sensor system will be subject to drift errors, which cannot be compensated through proprioception.

Exteroception offers potential solutions to these problems, allowing places to be identified and drift errors to be corrected on the basis of perceived environmental features. However, this also introduces another problem; an exteroception-based localisation system will be affected by *perceptual aliasing* sooner or later in reasonably complex environments. This refers to the situation where two or more places may be perceptually similar enough to be confused by the robot. In other words, the robot is unable to distinguish distinct places based on their perceptual signatures alone.

This research therefore addresses the question of how the different forms of perception available to the robot might be combined to provide a solution to the lost robot problem. In the experiments presented here, the robot had to relocalise after being moved to a randomly chosen location, its sensors being disabled during that move. Thus, the robot had no idea of its approximate location based on past experience. In addition, the problem of perceptual aliasing was deliberately made worse by restricting the robot to a subset of the available sensory information, e.g., using only the short-range infrared sensors of the Nomad 200 robot used in the experiments.

2 Related Work

There are various approaches in the literature to achieving localisation. Some authors combine proprioception and exteroception, relying on global odometry corrected by sightings on known environmental features. For example, Kurz [3] uses a self-organising neural network to cluster the robot's sensor readings, associating each cluster with an (x, y) coordinate obtained by dead reckoning. This system will fail however if two or more places share the same perceptual category, as a lost robot would be unable to distinguish between such places without prior knowledge of its approximate position.

In a different approach, Zimmer [9] combines position information from odometry with the corresponding sensor readings in the input vector to a continuously adaptable mapbuilding and localisation system. However, localisation may fail if reasonably accurate odometer information becomes unavailable and several stored sensor situations match the current sensor data. In general, approaches which require global odometry will not solve the lost robot problem, because the prior knowledge of the robot's position will be unavailable, and the correct identification of observed features can never be guaranteed with certainty in reasonably complex environments.

Another approach attempts to uniquely identify distinct places without proprioception, using more and more sensory information from the robot's exteroceptors to disambiguate similar looking places. For example, Yamauchi and Langley [8] rotate the turret of a stationary robot equipped with range sensors to construct a detailed grid model of the robot's immediate environment, effectively increasing the robot's perceptual resolution by a factor of s, where s is the number of rotational steps used to build the grid model. A separate evidence grid is constructed for each different "place" in the robot's environment, each grid cell containing a probability of occupancy by an object in the corresponding region of physical space.

Unlike previous grid-based approaches, this approach avoids the need for a globally consistent metric world model, instead using comparison of current and stored perceptions to evaluate both the current place occupied by the robot and the relative metric displacement of the robot within that particular place. Also, over-dependency on accurate sensing is avoided, because the compared perceptions will be subject to the same distortions, e.g., specular reflection. However, this system still does not overcome the problem of perceptual aliasing, and fails completely if the wrong place has been identified. Thus, it can be seen that increasing perceptual resolution alone will not solve the lost robot problem, as perceptual aliasing can never be eliminated with certainty in complex environments.

Another exteroception-based approach attempts to deal with perceptual ambiguity by incorporating history of local sensor sequences into the recognition of locations. The basic idea is to uniquely identify places by adding context information to their perceptual signatures. Nehmzow *et al* [4] combine past and present sensorimotor experience of the robot in the input to a self–organising feature map. Similarly, Tani and Fukumara [7] add previous sensor patterns to the input field of a neural network controller.

A disadvantage of this approach is that the identical set of previously stored locations must be revisited

Figure 1: *The Manchester 'FortyTwo'.* The robot's sonar and infrared sensors are mounted on the turret. In these experiments, the turret was kept at a constant orientation to provide a "compass sense". Exploration was carried out using the separate translational and rotational motors located in the base of the robot.

for matching of history information to be possible. This means that relocalisation can only be achieved by restricting the navigational behaviour of the robot, e.g., to following fixed paths. In general, this approach is limited by the fixed length of history stored, and does not generalise to cover arbitrarily large areas of perceptual ambiguity (e.g., long featureless corridors) in the robot's environment. Also, these methods are not robust to sensor noise; one freak perception will affect the next n inputs to the system, where n is the number of previous perceptions taken into account for localisation.

3 System Architecture

The mapbuilding and localisation system presented here was is composed of a hierarchy of modules (see figure 2), executed concurrently on a Nomad 200 robot (see figure 1). During both mapbuilding and relocalisation, the robot explores its environment using a reactive behaviour previously acquired using instinct rules [5]. In the experiments presented here, either wall-following or random wandering behaviours were used for exploration.

The robot begins by building a map of its environment, using a self-organising classifier to cluster its exteroceptive sensor readings. The robot is then moved to a random position in its environment, its sensors being disabled during that move. The robot then begins to actively explore its environment, and attempts to relocalise on the basis of evidence accumulated from both relative odometry and the robot's self-organising classifier mechanism. Using this sensory evidence, the robot is able to choose between the competing place memories contained in the map. A winning location "hypothesis" emerges once a perceptually unique path has been traced through the map.

3.1 Perceptual Clustering

The first stage of the problem is for the robot to recognise locally distinctive places. In order to deal with sensor noise, the robot should be capable of generalising to extract the most important features in a given situation without being distracted by the finer detail of individual sensor images. In these experiments, the self-organising neural network architecture ART2 [1] was used to cluster the robot's sonar and infrared sensor readings. This also helps to avoid the "correspondence problem" of matching specific environmental features against a detailed world model. The ART network is effectively used as a "black box" for classifying the robot's sensory input, and various other classification paradigms could also be used, e.g., Kohonen self-organising feature map [4]. (See [2] for a full discussion of the motivations for using ART.)

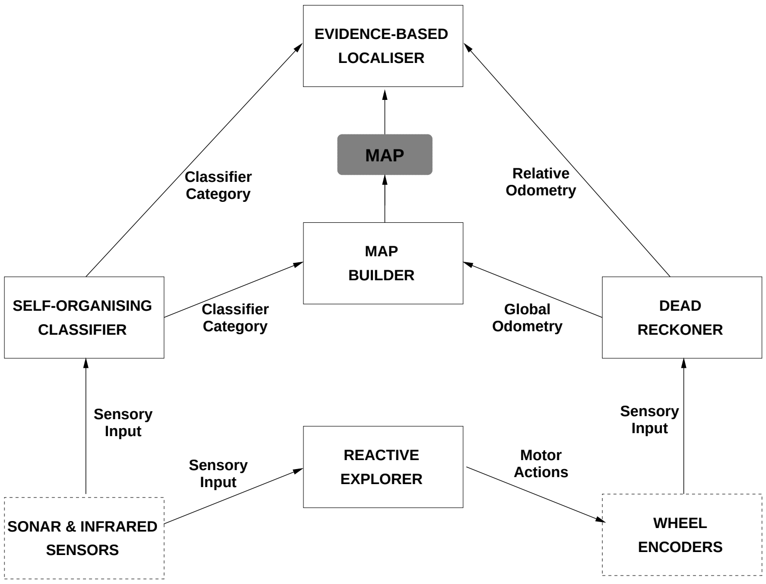

Figure 2: *System Architecture.* The solid-edged boxes denote behavioural modules, and arrows the dependencies between different parts of the system. The shaded box denotes the representation used for the map, and the dashed boxes show hardware components.

In previous work (e.g., [5]), the rotational and turret motors of the Nomad 200 robot have been controlled in unison, so that the sensors always point in the same direction relative to the heading of the robot. Here, however, the rotational motor was controlled independently, and the turret left stationary, in order to provide a consistent "compass sense". In other words, the robot explores its environment with its sensors always facing in the same global direction rather than the direction of travel. Thus, the appearance of locations to the robot depends on the robot's position alone, not its orientation, providing the robot with a "canonical view" of places. In the experiments conducted so far, no problem has been found with angular drift of the turret against the rotational movement of the robot's wheel motors, though this will need to be investigated thoroughly in subsequent large scale testing.

3.2 Map Building

The map constructed by the robot consists of a discrete set of place memories, each consisting of an ART category associated with an (x, y) coordinate obtained by dead-reckoning (see figure 3). This is similar to the mapbuilding mechanism described by Kurz [3], where the robot's (x, y) coordinates are averaged as the robot moves through a particular "situation area", i.e., an area of physical space associated with a particular classifier category. A new place is entered into the map whenever a new ART category is perceived, or when the robot has travelled by more than a prespecified distance threshold D from the nearest stored point in the map[1].

3.3 Relocalisation

The relocalisation system works by comparing old and new information at each iteration, or in other words, by considering the *changes* in the moving robot's perceptions over time. In this system, the robot maintains a set of competing place hypotheses (based on the place memories stored in the map), each being associated with a confidence level which reflects the robot's certainty in that particular place being the robot's true location. The hypothesis which emerges with the single highest confidence level (see below) is regarded as the winner, and taken to be the robot's actual location.

The algorithm used for performing relocalisation in these experiments is provided in figure 4. The first stage of the localisation process is to determine the place memory cell that most closely resembles

[1] In our experiments, $D = 0.25m$.

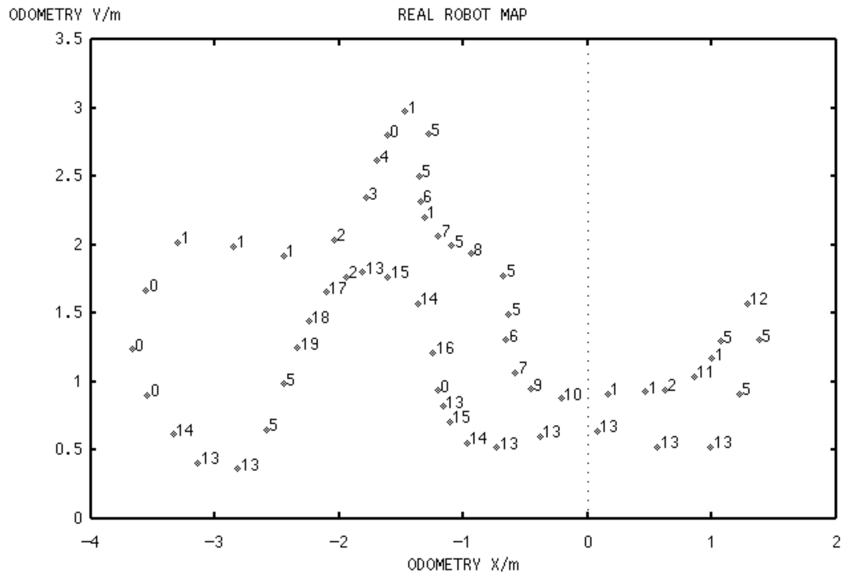

Figure 3: *Map Created By Real Robot.* The stored place memories follow the path of the wall-following robot around an enclosed environment (approximate size $6m \times 4m$). The numbers shown correspond to the perceptual categories assigned by the self-organising classifier. Perceptual aliasing occurs whenever two or more points share the same classifier category.

0. *Initialisation.* Formulate set of hypotheses, $H = \{h_0, ..., h_n\}$, consisting of all stored place memories which match the currently perceived ART category. Initialise confidence levels: $\forall h_i \in H$, assign $conf(h_i) = 1$
1. Wait until ART category changes or robot has moved by distance T.
2. For each hypothesis h_i add relative change in odometry $(\triangle x, \triangle y)$ to the coordinates of h_i, (x_{h_i}, y_{h_i}).
3. Generate second set of candidate hypotheses, $H' = \{h'_0, ..., h'_N\}$.
4. For each h_i, find the nearest h'_j, recording distance apart, d_{h_i}.
5. Accumulate evidence, given distance threshold T, minimum confidence level MIN, gain factor $GAIN > 1$ and decay factor $0 < DECAY < 1$:
   $\forall h_i \in H$

   if $d_{h_i} < T$ then

   replace (x_{h_i}, y_{h_i}) with $(x_{h'_j}, y_{h'_j})$ from the matching h'_j

   let $conf(h_i) = conf(h_i) \times GAIN$

   else

   let $conf(h_i) = conf(h_i) \times DECAY$

   if $conf(h_i) < MIN$ then delete h_i.
6. Delete any duplicates in H, preserving the one with the highest confidence level.
7. Add all remaining h'_j from H' which are not already contained in H to H, assigning $conf(h'_j) = 1$.
8. Repeat from step 1.

Figure 4: *Localisation Algorithm.*

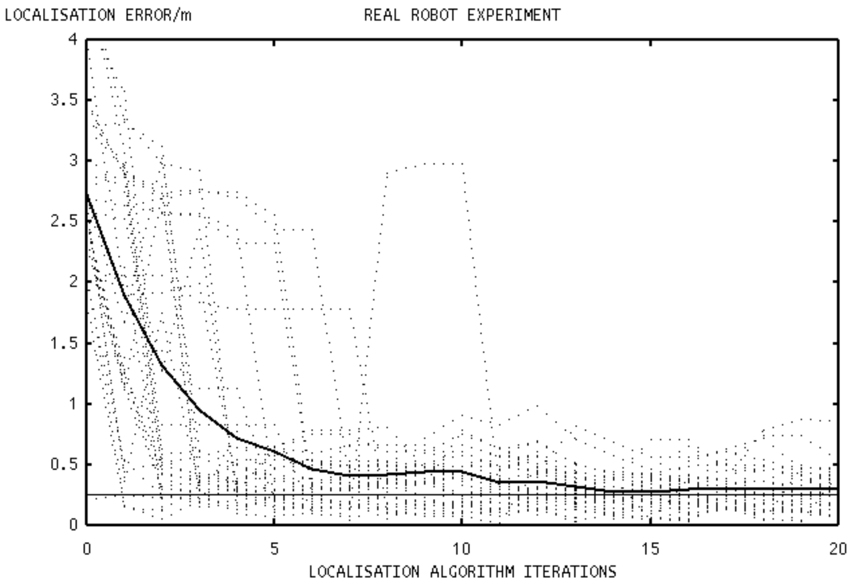

Figure 5: *Localisation Error on Real Robot.* The thick line shows the mean error taken over 30 trials. The dotted lines show the individual trials, and the horizontal line at $y = 0.25m$ denotes the distance threshold D used during mapbuilding, i.e., the upper limit of the best theoretically possible error range.

the current sensory perception. Because of perceptual aliasing, typically a number of place memory cells emerge as equally likely candidates for the robot's current position. The robot then performs an exploratory movement of an arbitrary nature, whose *relative* displacement in Cartesian space is observed by the localisation system. The current perception, previous perception and displacement are then used to narrow down the list of candidate cells for the robot's current position, providing positive or negative reinforcement of the confidence level associated with each of these cells[2].

4 Experiments

The system was tested on a Nomad 200 robot in an enclosed environment corresponding to the map shown in figure 3. Only the robot's infrared sensors were used here to deliberately exacerbate the problem of perceptual aliasing. Figure 5 shows the results when the robot was "lost" at 30 randomly chosen locations, and then told to relocalise using the same clockwise wall-following behaviour for exploration as used in the mapbuilding phase. The mean number of steps taken for the localisation error to fall below the distance threshold D was 7.0 in a mean time of 27.3 seconds, with the robot travelling at 0.10ms^{-1}.

In the results shown, the localisation error is calculated as the difference between the "actual" position of the robot obtained from the wheel encoders[3] and the winning position estimate created by the localisation algorithm. Where two or more hypotheses share the same confidence level, the error is taken from the hypothesis which lies furthest from the actual location. As the system localises to the stored place in the map which lies closest to the actual position of the robot, the best theoretically possible result should be a localisation error of between zero and D, the distance threshold used in mapbuilding.

Further inspection of the results showed that it is not only the sequence of ART categories perceived which determines the robot's estimated position, but also the *size* of the corresponding regions in Cartesian space (i.e., the number of neighbouring place memories sharing the same perceptual category). Thus, the robot continues to accumulate useful information even when no change is detected by the ART net-

[2]For these experiments, $D = 0.25m$, $GAIN = 3.0$, $DECAY = 0.7$, $MIN = 0.5$, $T = 0.50m$. A value of $T = 2D$ was chosen, because this corresponds to the maximum distance theoretically possible between stored place memories after mapbuilding has been completed.

[3]A problem here was that global odometry was used to measure the "true" position of the robot, as we currently have no independent means of measuring the robot's exact position. The results shown in figure 5 will therefore only be as accurate as the odometer drift error will allow.

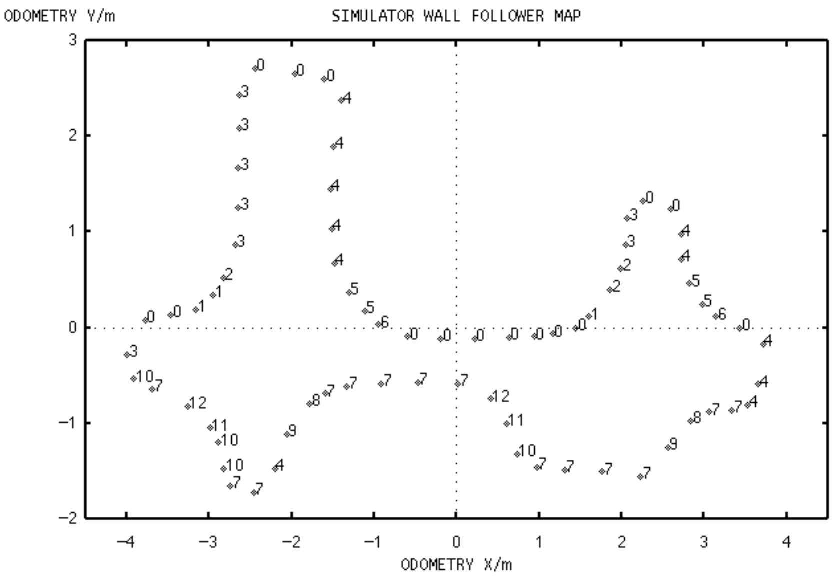

Figure 6: *Map Created By Wall Following Simulator Robot.* No single place memory has a unique ART category, an extreme case of perceptual aliasing.

work. This aspect was never deliberately "programmed" into the system, rather it emerges from the use of relative odometry in the localisation algorithm.

The experiment was also repeated on the robot using an anticlockwise wall-following behaviour for relocalisation rather than the standard clockwise wall-follower used for mapbuilding. This showed the same level of performance as before, demonstrating the ability to relocalise independently of a fixed temporal sequence of landmarks as in previous work, e.g., [4].

Work currently in progress includes conducting a set of long-range experiments on the Nomad 200, involving distances of the order of hundreds of meters, to determine the limits of the current system.

5 Simulations

While it was possible to make useful qualitative judgements about the behaviour of the localisation system on the real robot, the use of the wheel encoders for calculating the localisation error makes it difficult to make a really accurate assessment of the performance of the localisation algorithm. It is also very difficult to isolate the effects of different possible sources of error using the real robot, e.g., sensor noise, wheel slippage, chaotic variation in the exploration strategy, etc. For this reason, and for this reason only, the following tests were conducted on the Nomadic Robot Simulator, where different types of error could be carefully controlled.

5.1 Wall Following

In each of the following simulations, the same sensor readings were taken from the same simulated robot over the same time period. A simulated environment was deliberately constructed so that no single place in the map had a unique perceptual signature (see figure 6). Figure 7 shows the results of four simulations under various different error conditions.

In the first simulation (the "control"), no errors were introduced. In the second, a one in ten chance of misclassification by the ART network was introduced at each iteration of the localisation algorithm, when an incorrect category would be randomly selected. In the third, an artificial odometer drift error was introduced by continuously integrating the distance travelled by the robot, and adjusting the robot's current odometer readings by 25% of this distance in a direction chosen randomly at the start of each trial. Both the 10% misclassification error and 25% drift error were included in the fourth simulation, providing a test of the system's ability to relocalise under highly adverse conditions.

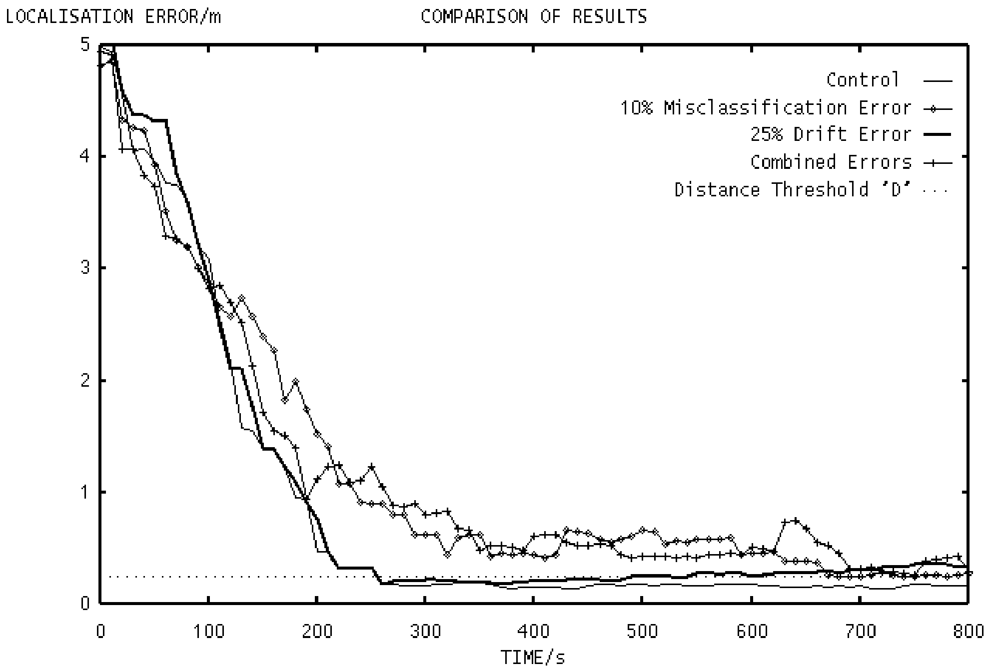

Figure 7: *Comparison of the Different Wall Following Simulations.* Mean errors taken over 30 trials each. An identical wall-following behaviour was used throughout. NB. The times taken from the simulator are the result of the speed of the operating system used, and should not be taken as "realtime", these are provided for relative comparison of the results only.

The results show a graceful degradation in performance with respect to the errors introduced, drift error in particular having little effect. The mean number of steps taken for the localisation error to fall below D was 3.6 in the control, and 4.0 in the third simulation (25% drift error added), both taking a similar time to localise[4].

Misclassification errors also made little difference to the time taken to localise, although the localisation algorithm had to work harder to achieve this. For example, it took a mean of 12.4 steps for the localisation error to fall below D in the second simulation (10% misclassification error) compared to the 3.6 steps taken in the control. This is because each misclassification produced by the ART network will trigger an extra iteration through the localisation algorithm, plus another one when the correct category is perceived once more.

5.2 Random Wandering

A further set of simulations involved exploring another enclosed environment (see figure 8) using a random wandering behaviour based on smooth, continuous obstacle avoidance (see [5] for details) instead of wall-following. Sonar was used for input to the ART network rather than infrared, as much of the space away from the walls of the enclosure contained no features discernible by infrared. The aim of this simulation was to test the independence of the mapbuilding and localisation systems from any particular exploration strategy. The wandering behaviour also meant that the robot could approach places from arbitrary directions, testing the ability of the system to provide "canonical views" of places, as discussed previously.

While this simulation showed that the mapbuilding and localisation systems do not rely on following a fixed route as in wall-following, it could take a long time both to build a complete map even in the small environment used. This is hardly surprising, as the exploration strategy chosen basically relies on blind chance and "the law of averages" to eventually steer the robot towards uncharted areas of the map. This was also reflected in the results obtained for relocalisation. While the robot never failed to relocalise

[4]This compares with a mean value of 7.0 steps in the experiments with the robot.

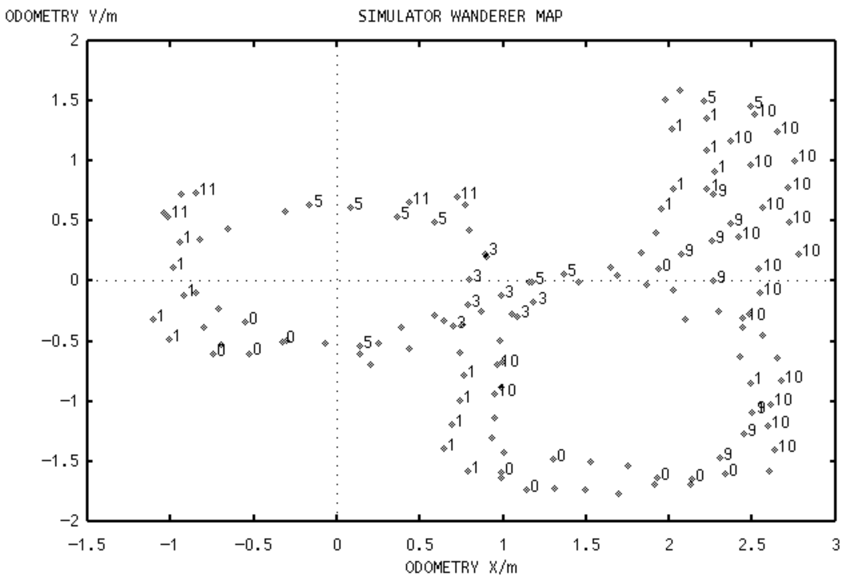

Figure 8: *Map Created by Wandering Simulator Robot.* Not all of the ART categories are shown here for clarity.

eventually, it could take a long time before it was able to find a perceptually unique path through the enclosure. The mean number of steps taken was 10.2, but there was a lot of variation around this level (standard deviation 7.5 steps), depending on the luck of the localiser in being party to a good piece of exploration by the wandering behaviour.

These results show a clear advantage in principled exploration over random exploration. In particular, mapbuilding by following a canonical path (i.e., wall-following) was shown to be more efficient, requiring far fewer place memories to cover a similar sized area than the random strategy. Furthermore, a map constructed from a single tour of the environment can contain all of the information needed by the lost robot to relocalise using the same exploration strategy. (See [6] on the use of canonical paths in animal and robot navigation.)

6 Summary

A complete robot mapbuilding and localisation system has been presented which can localise even in environments where no single place has a unique perceptual signature, both on a robot and using a simulator. No prior estimate of position is required for relocalisation, so the robot can recover its position even after becoming completely lost. During an active exploration phase, evidence is accumulated to decide between competing place memories. Perceptual changes are correlated with a self-acquired map of the environment to provide positive or negative reinforcement of the competing hypotheses. The use of relative odometry only during relocalisation means that the localisation algorithm operates independently of any global reference frame. In addition, wall-following was shown to be a superior exploration strategy over random wandering.

7 Conclusions

The localisation algorithm presented successfully combines exteroception and proprioception without reliance on global odometry. During relocalisation, the robot can continue to collect useful evidence even when moving between stored places which share the same perceptual signature. More generally, it can be seen that a winning position estimate emerges once a perceptually unique path has been traced through the map. This applies regardless of the length of the path, or in other words, the amount of sensory experience required to disambiguate similar looking locations. The localisation method therefore generalises in principle to cover arbitrary levels of perceptual aliasing in the environment, and suffers a graceful degradation in performance with respect to the possible errors introduced in the simulations presented here.

Drift errors in the robot's odometry have little effect on localisation performance, proprioception only being used to determine the approximate relative displacement of the robot between places. Misclassification errors in the self-organising classifier have a greater effect on performance, requiring more processing to achieve successful relocalisation. Freak perceptions during mapbuilding would also lead to "ghost" place memories being stored in the map, i.e., cells in the map with no corresponding place in reality. There are (at least) two ways of dealing with potential misclassification errors in future work.

Firstly, the effects of misclassification might be reduced by allowing a fuzzy match between current and stored sensor images, rather than insisting on absolute boundaries between distinct perceptual "categories". Future work will therefore involve investigating alternatives to the self-organising classifier mechanism used here. Secondly, any false place memories contained in the map could be removed over time by incorporating forgetting as well as aquisition of map knowledge during mapbuilding. We are therefore currently investigating the use of lifelong learning, where the map would be continuously updated on the basis of new perceptions, as in Zimmer [9].

The map*building* scheme, as it currently stands, is dependent on a global reference frame, and would be negatively affected by odometry errors. We are therefore also investigating the use of continuous relocalisation during mapbuilding to compensate for drift errors. A more robust compass sense is also needed, as the current system would fail if the robot's turret failed to face in the same global direction throughout. We are also investigating "pro-active" exploration, where exploration might be driven by the state of the internal world model rather than being purely reactive, and scaling the whole system up to work in large, dynamic environments.

To conclude, the key achievements of the work are the handling of perceptual aliasing and the ability to localise when no predetermined estimate of position is available. The underlying principles are the maintenance of competing place memories and the incremental accumulation of evidence over time to choose between them.

References

[1] 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.

[2] T. Duckett and U. Nehmzow, "A Robust, Perception-Based Localisation Method for a Mobile Robot", Dept. Computer Science, University of Manchester, Technical Report Series UMCS-96-11-1, 1996.

[3] 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.

[4] U. Nehmzow, T. Smithers and J. Hallam, "Location Recognition in a Mobile Robot Using Self-Organising Feature Maps", in G. Schmidt (ed.), "Information Processing in Autonomous Mobile Robots", Springer Verlag, 1991.

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

[6] U. Nehmzow, "Animal and Robot Navigation", Robotics and Autonomous Systems, Vol. 15, No. 1 - 2, pp. 71-81, 1995.

[7] J. Tani and N. Fukumara, "Learning Goal–Directed Sensory–Based Navigation of a Mobile Robot", Neural Networks, Vol. 7, No. 3, pp. 553-563, 1994.

[8] B. Yamauchi and P. Langley, "Place Recognition in Dynamic Environments", to appear in Journal of Robotics Systems, Special Issue on Mobile Robots, 1996.

[9] U. R. Zimmer, "Self–Localization in Dynamic Environments", IEEE/SOFT International Workshop BIES'95, Tokyo, Japan, May 30-31 1995.