Effect of Representation Size and Visual Ambiguity on RatSLAM System Performance

Michael Milford, David Prasser, Gordon Wyeth
School of Information Technology and Electrical Engineering
The University of Queensland
St Lucia Queensland 4072 Australia
{milford, prasserd, wyeth@} itee.uq.edu.au

Abstract

RatSLAM is a vision-based SLAM system based on extended models of the rodent hippocampus. RatSLAM creates environment representations that can be processed by the experience mapping algorithm to produce maps suitable for goal recall. The experience mapping algorithm also allows RatSLAM to map environments many times larger than could be achieved with a one to one correspondence between the map and environment, by reusing the RatSLAM maps to represent multiple sections of the environment. This paper describes experiments investigating the effects of the environment-representation size ratio and visual ambiguity on mapping and goal navigation performance. The experiments demonstrate that system performance is weakly dependent on either parameter in isolation, but strongly dependent on their joint values.

1 Introduction

RatSLAM is a vision-based SLAM (Simultaneous Localisation And Mapping) system inspired by computational models of the rodent hippocampus. It extends the concept of place cells and head direction cells used in models by Arleo [2000], Stringer *et al.* [2002, 2002] and Browning [2000] to *pose* cells that represent the robot's location and orientation. RatSLAM has been demonstrated performing SLAM in a range of unmodified indoor environments [Milford *et al.*, 2004] and outdoor environments [Prasser *et al.*, 2005].

Work following on from the initial SLAM investigations has addressed the suitability of the RatSLAM representations for use in tasks such as goal navigation [Milford *et al.*, 2005, Milford *et al.*, 2006]. While the RatSLAM representations become stable and coherent over time, phenomena such as multiple representations in the pose cells make them unsuitable for direct use in goal navigation. To create representations that could be used for goal recall, the *experience mapping* algorithm was developed [Milford *et al.*, 2005]. Experience maps preserve the spatial information which is lost from the pose cell representations as they become increasingly discontinuous. The maps share spatial and topological similarities with those produced by the ELDEN system [Yamauchi and Beer, 1996], and store behavioural information at each experience in a manner similar to that used by the Spatial Semantic Hierarchy [Kuipers *et al.*, 2004].

Because the experience maps obtain their spatial information directly from internal robot sensor information, they remove any need to interpret the pose cells' structural arrangement. This allows the pose cell map to be reused multiple times, allowing mapping of environments much larger than the nominal representation size of the pose cell structure. Recent work has demonstrated that the RatSLAM system can successfully perform goal navigation even when the pose cell representation is many times smaller than the actual environment [Milford, et al., 2006]. Referring back to the biological literature from which RatSLAM draws its inspiration, it is interesting also to note recent research suggesting that cells in rats are reused at regular intervals in the environment [Fyhn *et al.*, 2004, Hafting *et al.*, 2005].

The work described in this paper seeks to quantify the relationship between the size of the pose cell structure and system performance, for a fixed environment size. Because the experience mapping algorithm partially relies on visual information to identify experiences, visual ambiguity is also increased to assess its effect. System performance is assessed quantitatively in terms of goal navigation success rates, as well as qualitatively by examining the structure of the pose cell and experience maps.

This paper proceeds as follows. Section 2 describes the core RatSLAM system, with the experience mapping algorithm described in Section 3. The goal recall process is described in Section 4. Section 5 describes the experimental setup and procedure, and Section 6 presents the results. Discussion is provided in Section 7, before the paper concludes in Section 8.

2 RatSLAM

This section briefly describes the RatSLAM system – a more detailed description is given in [Milford *et al.*, 2004]. Figure 1 shows the core structure of the RatSLAM system. The robot's pose is represented by activity in a competitive attractor neural network called the pose cells. Wheel encoder information is used to perform path integration by appropriately shifting the current pose cell activity. Activity can wrap in all three directions in the pose cell matrix. The path integration model results in each pose cell initially representing a 0.25m × 0.25m × 10° space. Vision information is converted into a local view (LV) representation that is associated with the currently active pose cells. If familiar, the current visual scene also causes activity to be injected into the particular pose cells associated with the currently active local view cells.

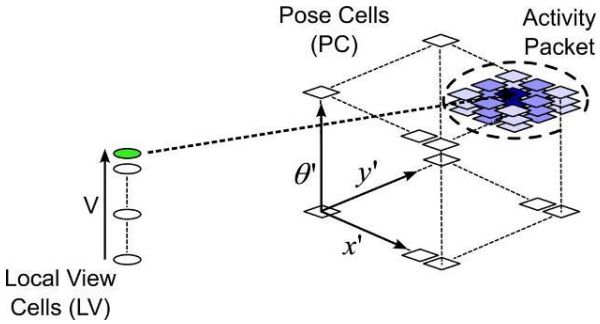

Figure 1 - The core RatSLAM pose cell and local view cell networks.

For small environments, the RatSLAM representations have a high degree of correspondence to the Cartesian layout of the environment. However, as the environments become larger and more complex, a number of phenomena become common. Large accumulated path integration errors are corrected when the robot sees familiar visual scenes. This leads to discontinuities in the pose cell matrix where the dominant packet of activity jumps from one location to another. Because re-localization is not an instantaneous process, the system also learns multiple representations of the same physical areas in the environment. The complementary attribute of collisions in the pose cell matrix also becomes increasingly common – clusters of pose cells become associated with more than one location in the environment. These phenomena render the RatSLAM representations unsuitable for direct use in tasks such as goal recall, and were the reason for the development of the experience mapping algorithm discussed in the following section.

3 Experience Mapping

The premise of the experience mapping algorithm is the creation and maintenance of a collection of experiences and inter-experience links. The algorithm creates experiences to represent certain states of activity in the pose cell and local view networks. The algorithm also learns behavioural, temporal, and spatial information in the form of inter-experience links. Figure 2 shows the relationship between the experience map and the core RatSLAM representations.

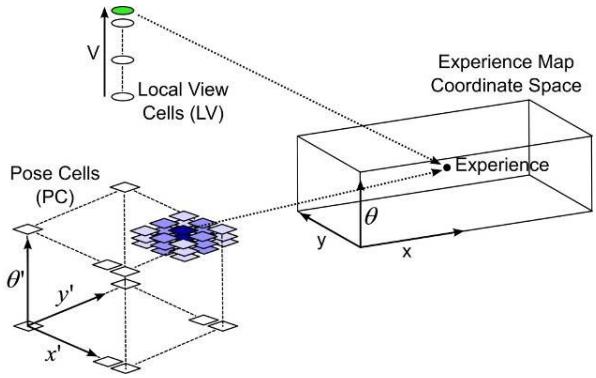

Figure 2 - An experience is associated with certain pose and local view cells, but exists within the experience map's own coordinate space.

3.1 Experiences

Experiences have an activity level that is dependent on how close the activity peaks in the pose cells and local view cells are to the cells associated with the experience. The component of activity determined by the pose cell network activity, $E_{x'y'\theta'}$, is given by:

$$E_{x'y'\theta'} = \begin{cases} 0 & \text{if } r_r' > 1; \\ 0 & \text{if } \theta_r' > 1; \\ 2 - r_r' - \theta_r' & \text{otherwise} \end{cases} \quad (1)$$

$$r_r' = \frac{\sqrt{\left(x_{pc}' - x_i'\right)^2 + \left(y_{pc}' - y_i'\right)^2}}{r_a} \quad (2)$$

$$\theta_r' = \frac{\left|\theta_{pc}' - \theta_i'\right|}{\theta_a} \quad (3)$$

where x'_{pc}, y'_{pc}, and θ'_{pc} are the coordinates in the pose cell matrix of the dominant activity packet, x'_i, y'_i, and θ'_i are the coordinates of the pose cells associated with experience i, r_a is the zone constant for the (x', y') plane, and θ_a is the zone constant for the θ' dimension. The visual scene V_i switches an experience on or off:

$$E_i = \begin{cases} 0 & \text{if } V_{curr} \neq V_i; \\ E_{x'y'\theta'} & \text{if } V_{curr} = V_i \end{cases} \quad (4)$$

where V_{curr} is the current visual scene, and V_i is the visual scene associated with experience i. The most active experience is known as the peak experience. Learning of new experiences is triggered by the peak experience's activity level dropping below a threshold value.

3.2 Experience Transitions

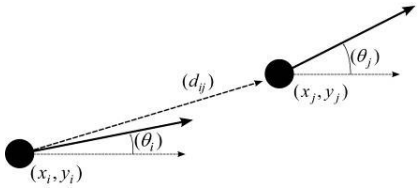

Figure 3 - Links between experiences store several types of information, including odometric information about the robot's movement during the transition.

Inter-experience links store temporal, behavioural, and odometric information about the robot's movement between experiences. Figure 3 shows a transition from experience *i* to experience *j*. The physical movement of the robot during this transition is given by:

$$d\mathbf{p}_{ij} = \mathbf{p}_j - \mathbf{p}_i = \begin{pmatrix} \theta_j \\ x_j \\ y_j \end{pmatrix} - \begin{pmatrix} \theta_i \\ x_i \\ y_i \end{pmatrix} = \begin{pmatrix} d\theta_{ij} \\ dx_{ij} \\ dy_{ij} \end{pmatrix} \quad (5)$$

where $d\mathbf{p}_{ij}$ is a vector describing the position and orientation of experience *j* relative to experience *i*. Repeated transitions between experiences result in an averaging of the odometric information:

$$d\mathbf{p}_{ij}^{new} = A.d\mathbf{p}_{ij}^{old} + B.d\mathbf{p}_{ij}^{curr} \quad (6)$$

where $A = \begin{pmatrix} 1/2 & 0 & 0 \\ 0 & d_s.\cos d\theta & -d_s.\sin d\theta \\ 0 & d_s.\sin d\theta & d_s.\cos d\theta \end{pmatrix}$,

$$B = \begin{pmatrix} 1/2 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix},$$

$$d\theta = \frac{1}{2}\left[\tan^{-1}\left(\frac{dy_{ij}^{curr}}{dx_{ij}^{curr}} \right) - \tan^{-1}\left(\frac{dy_{ij}^{old}}{dx_{ij}^{old}} \right) \right], \text{ and}$$

$$d_s = \left(d_{ij}^{curr} + d_{ij}^{old} \right) \Big/ \left(2d_{ij}^{old} \right).$$

3.3 Map Correction

Discrepancies between a transition's odometric information and the linked experiences' (x, y, θ) coordinates are minimised through a process of map correction:

$$\Delta\mathbf{p}_i = \alpha \left[\sum_{j=1}^{N_f} \left(\mathbf{p}_j - \mathbf{p}_i - d\mathbf{p}_{ij} \right) + \sum_{k=1}^{N_t} \left(\mathbf{p}_k - \mathbf{p}_i - d\mathbf{p}_{ki} \right) \right] \quad (7)$$

where α is a learning rate constant, N_f is the number of links from experience *i* to other experiences, and N_t is the number of links from other experiences to experience *i*. The experience map is subject to the same constraints of any network style learning system – appropriate learning rates must be used to balance rapid convergence with instability. Experimentation has determined that a learning rate of 0.5 results in the map rapidly converging to a stable state.

When the orientation of an experience is changed through the map correction process, the (dx, dy) component of the transitional information must also be updated to account for the rotation:

$$d\mathbf{p}_{ij}^{new} = \begin{pmatrix} 1 & 0 & 0 \\ 0 & \cos\Delta\theta_i & -\sin\Delta\theta_i \\ 0 & \sin\Delta\theta_i & \cos\Delta\theta_i \end{pmatrix} d\mathbf{p}_{ij}^{curr} \quad (8)$$

3.4 Experience Map Adaptation

The experience maps represent changes in the environment through modification of the inter-experience transition information. As well as learning experience transitions, the system also monitors transition 'failures'. Transition failures occur when the robot's current experience switches to an experience other than the one expected given the robot's current movement behaviour. If enough of these failures occur for a particular transition, indicated by the *confidence* level dropping below a certain threshold, then that link is deleted from the experience map. The confidence level is given by:

$$c_{ij} = \frac{n_{ij}}{n_{i\beta_{ij}}} \quad (9)$$

where n_{ij} is the number of times the transition between experience *i* and *j* has occurred, and $n_{i\beta_{ij}}$ is the number of times a transition from experience *i* to any experience using the behaviour β_{ij} has occurred. Current work is addressing the implementation of a scheme more suited to longer term experiments, such as the recency weighted averaging described in [Yamauchi and Beer, 1996].

4 Goal Recall

The temporal link information and local spatial properties of the experience map are used to form a temporal map that can be used for goal recall.

4.1 Temporal Map Creation

To create the temporal map, the peak experience is seeded with a zero time stamp value, and all other experiences are seeded with a 'very large' value. Time stamp values are then assigned to linked experiences based on the peak experience's time stamp and the temporal link information. This process is iterated, with any experience A only updating experience B if they are linked and if experience B's time stamp is larger, shown by (10) and (11):

$$\tau_j = t_i + t_{ij} \quad (10)$$

$$t_j^{k+1} = \min\left(\tau, t_j^k \right) \quad (11)$$

where t_i is the time stamp value of experience i, t_{ij} is the temporal link from experience i to experience j, τ_j is the resultant proposed time stamp value of experience j, $\boldsymbol{\tau}$ is the set of proposed time stamp values for experience j, k is the iteration number, and t_j^{k+1} is the updated time stamp value.

4.2 Route Finding

To find the shortest route to a goal, a gradient climbing procedure is used, starting at the goal location. The route is stored as a sequence of nodes, with each node corresponding to an experience. Each node is time stamped to represent the estimated time it will take from the robot's current position to reach each node along the route to the goal.

4.3 Behaviour Arbitration

RatSLAM uses the goal route as input to its local movement module in order to pick appropriate movement behaviour. The robot chooses the local movement path that is spatially closest to the recalled route, shown in Figure 4.

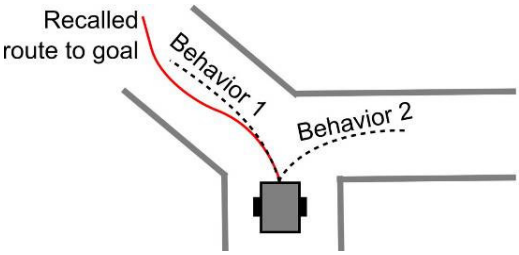

Figure 4 - Behavior arbitration. In this situation the robot would choose to go left, as that local movement behavior more closely matches the proposed route.

4.4 Route Loss Recovery

As the robot navigates toward the goal, the navigation system is constantly finding the shortest route to the goal based on the current peak experience. New visual scenes can trigger the learning of new experiences even when the robot is traversing a previously travelled path. This allows the robot to more completely learn the route but also results in the temporary loss of a functional temporal map, since time stamps cannot be propagated from a brand new experience that is not yet linked to any others.

In this situation the robot uses the last known 'good' route to the goal to navigate. The robot's position along the route is updated based on time elapsed. If an old route is relied on for longer than a certain time period, it is discarded and the robot returns to exploration in order to acquire a fresh route to the goal.

5 Experimental Setup

This section describes the simulator, the environment, and the test procedure for the experiments.

5.1 Simulator

RatSLAM is tested in simulation using a simulator developed specifically for the RatSLAM project. The simulator has been developed in Visual C++ and uses the Open Graphics Library (OpenGL). It has several running modes ranging from full simulation to playback of actual test data. For these experiments the full simulation mode was used to allow greater control over the environment conditions. Figure 5 shows a screenshot of the simulator's graphical user interface.

One of the simulator's key environment variables is N_{vis}, which affects the likelihood of false positive visual re-recognition by the robot. N_{vis} corresponds directly to the likelihood at any one time that the robot will incorrectly perceive the current visual scene as belonging to another, previously experienced part of the environment. This parameter was varied during the experiments to determine the effect that visual ambiguity has on the system performance.

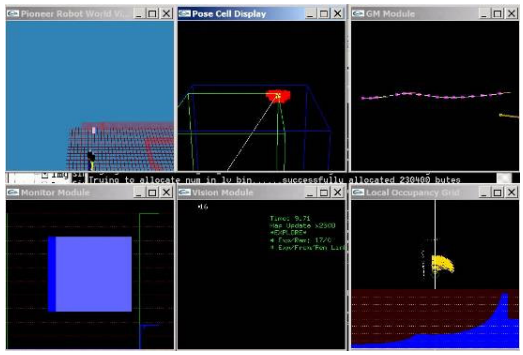

Figure 5 - Simulator graphical user interface. From top left, in a clockwise direction, the windows show (a) the simulated environment (b) the pose cells (c) the experience map (d) system information (e) vision information and (f) laser range information.

5.2 Environment

The testing environment was a model of the open plan office space in our building, measuring ten metres by eight metres. Three of the room's corners were used as goal locations, with the fourth corner used as the starting location for the robot when it commenced navigation to goal number one.

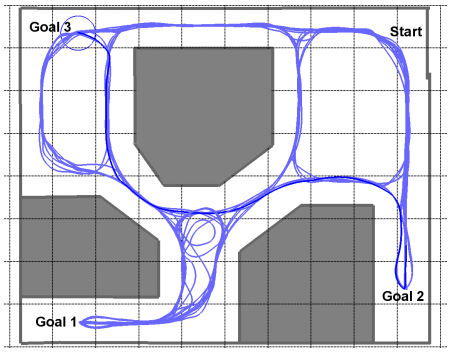

Figure 6 - The environment floorplan, showing the robot trajectory, obstacles and goal locations. Each grid square measures one metre by one metre.

5.3 Procedure

Each trial consisted of 20 minutes of exploration by the robot, followed by navigation in sequence to the three goals. During exploration, the first time the simulated robot approached each of the three goal locations, the currently active experience was tagged. When goal recall commenced, these tagged experiences were used as the desired goal locations within the experience map.

When the goal navigation procedure determined that it had reached a goal, the simulated robot's position was noted and compared to the actual goal location. If the robot was not sufficiently close to the actual goal location, the navigation segment was noted as a failure. If the goal navigation process was unable to find a goal after a set amount of time, the failure was noted and navigation to the next goal was commenced.

Ten trials were carried out for each combination of visual ambiguity, N_{vis}, and pose cell structure dimensions, L_{pc}. The two parameters were varied as given by:

$$N_{vis} : \{0\%, 2\%, 4\%, 6\%, 8\%, 10\%\}$$
$$L_{pc} : \{64\times64, 32\times32, 16\times16, 8\times8\}$$

L_{pc} is defined in terms of the number of cells along each side of the pose cell structure in the (x', y') plane. Using every possible combination of parameters produced 24 separate experiments, for a total of 240 trials, of approximately 100 hours total duration.

Table 1 summarises the ratio between the physical size of the environment (10 × 8 metres) and the nominal representation size of the pose cells, r_s. The nominal size is the maximum environment size that can be represented without wrapping in the pose cell structure, assuming the robot odometry is perfect. The nominal size is calculated by multiplying the (x', y') cell dimensions of the pose cell structure by the nominal (x, y) area represented by a pose cell, as given by the path integration model. For indoor RatSLAM experiments the nominal area represented by each pose cell is 250 mm by 250 mm.

TABLE 1. RATIO BETWEEN ENVIRONMENT SIZE AND NOMINAL POSE CELL REPRESENTATION SIZE

Pose Cell Dimensions (x', y')	$r_s = \frac{\textbf{environment size}}{\textbf{nominal representation size}}$
64 × 64	0.31
32 × 32	1.25
16 × 16	5
8 × 8	20

6 Results

This section presents the results of the experiments described in Section 5.3. Sections 6.1 and 6.2 present the pose cell representations and experience maps generated by the RatSLAM system for a visual ambiguity level of 2%. Section 6.3 presents the system's goal navigation performance over all the trials for each combination of pose cell structure size and visual ambiguity.

6.1 Pose Cell Maps

At a size of 64 × 64, the pose cell structure is large enough to represent the environment without wrapping of across the boundaries of the pose cells, as shown by the top row of Figure 7. For 32 × 32 pose cells, some wrapping occurs. At 16 × 16 pose cells, there is significant wrapping, and many of the pose cells are being reused several times to represent different areas of the environment. At 8 × 8 pose cells, all cells are heavily utilised. It is important to note that the trajectory shown in the figure is only that of the most highly activated cell.

6.2 Experience Maps

The experience maps corresponding to the RatSLAM representations shown in Figure 7 are shown in Figure 8. When the 64 × 64 pose cell structure is used all experience maps are correctly formed. For 32 × 32 and 16 × 16 pose cell structures most experience maps are formed correctly but some have significant errors. Using an 8 × 8 pose cell structure, the experience mapping algorithm is still able in some cases to generate a correct map, although many of the maps have errors and some are highly deformed.

6.3 Goal Recall Performance

With unambiguous vision and a sufficiently large pose cell structure (64 × 64 cells), the robot had a 100% success rate in navigating to all three goals over the ten trials (Figure 9). As the pose cell structure size shrunk, the navigation performance gradually degraded, with just over an 80% success rate for the smallest pose cell structure size.

The goal navigation performance for all pose cell structure sizes also degraded as the visual ambiguity was increased, as can be seen in Figure 10 to Figure 14. Although performance was reduced across all tests, the most significant reduction occurred in the experiments involving the smallest 8 × 8 pose cell structure, as shown in Figure 15. For the 64 × 64 and 32 × 32 cell pose structures, the increase in visual ambiguity had little effect on performance. For the 16 × 16 cell pose structure there was a definite correlation between visual ambiguity and goal navigation performance, a correlation which was even more pronounced for the 8 × 8 cell structure.

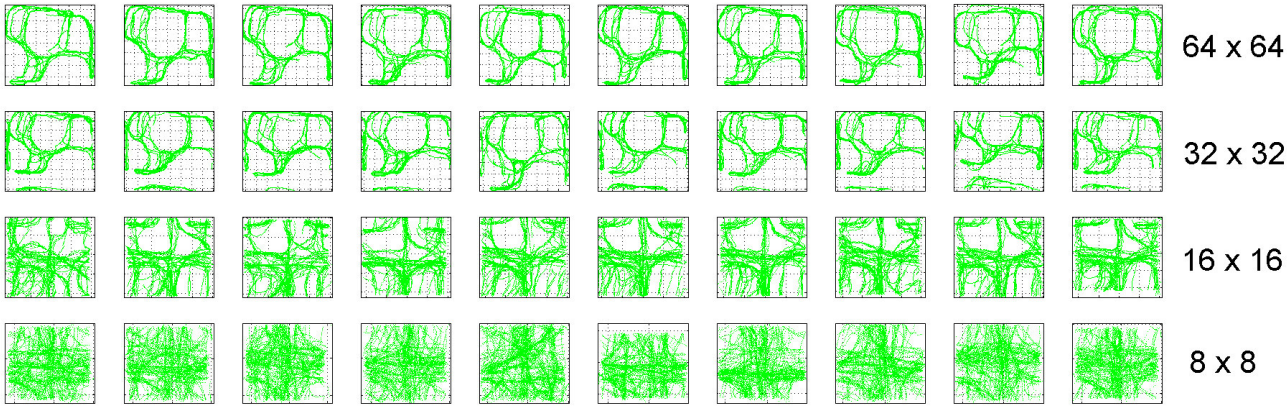

Figure 7 – Pose cell maps under 2% visual ambiguity using each pose cell structure size (shown on the right). Each grid square represents 4×4 pose cells in the $(x' y')$ plane. The trajectory shows the path of the most highly activated pose cell during the experiment.

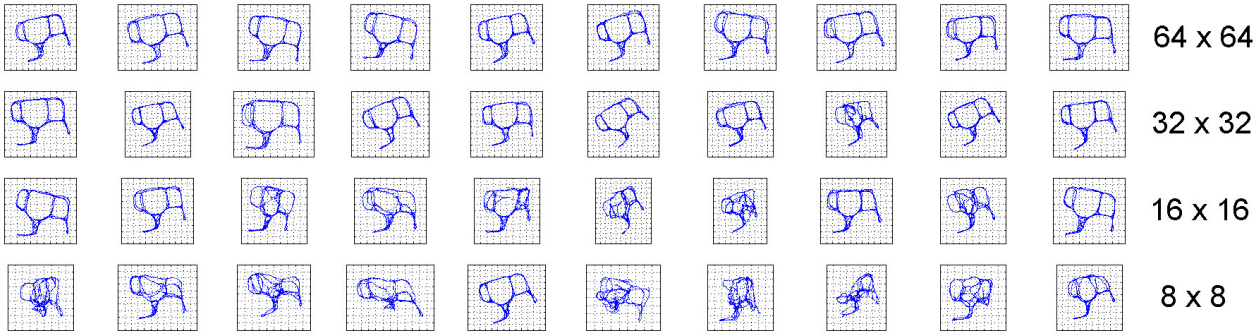

Figure 8 – Experience maps created under 2% visual ambiguity using each pose cell structure size (shown on the right). Each grid square represents one square metre.

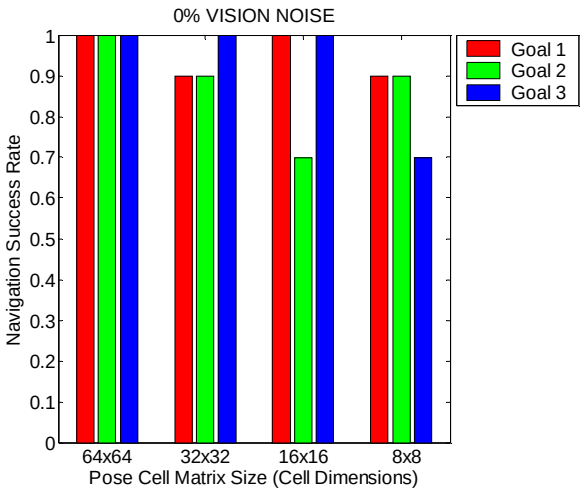

Figure 9 - Goal recall performance with no visual ambiguity.

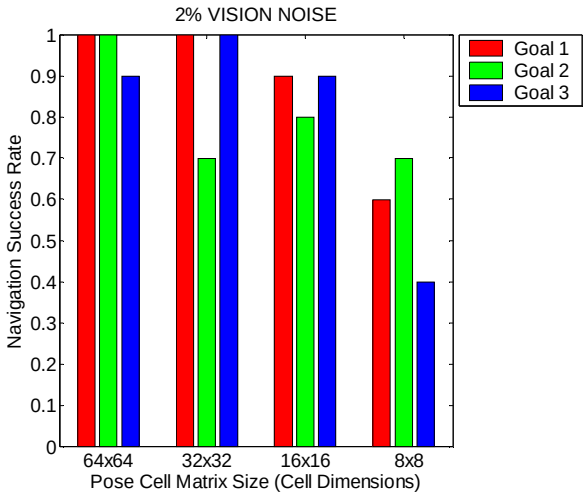

Figure 10 - Goal recall performance with 2% visual ambiguity.

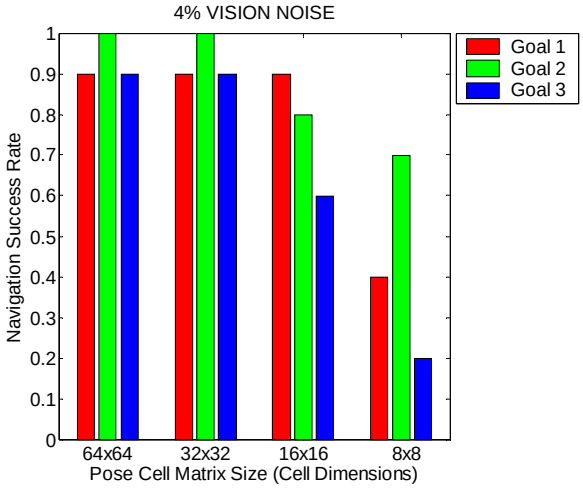

Figure 11 - Goal recall performance with 4% visual ambiguity.

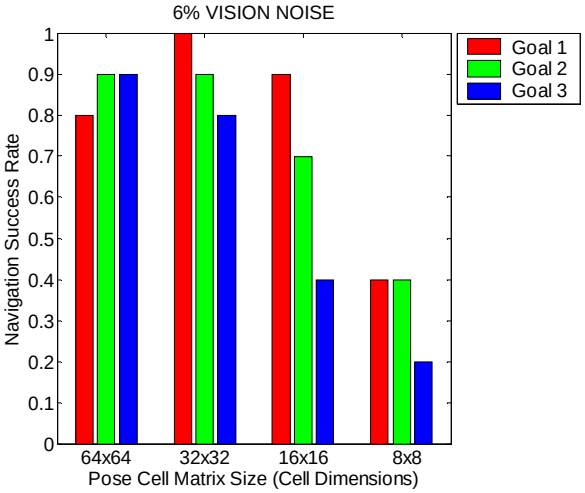

Figure 12 - Goal recall performance with 6% visual ambiguity.

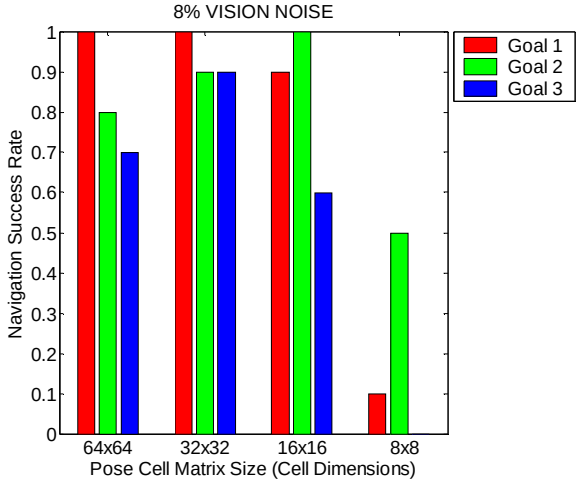

Figure 13 - Goal recall performance with 8% visual ambiguity.

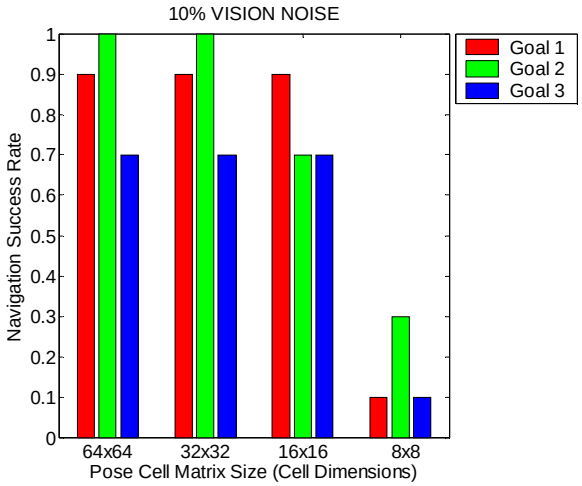

Figure 14 - Goal recall performance with 10% visual ambiguity.

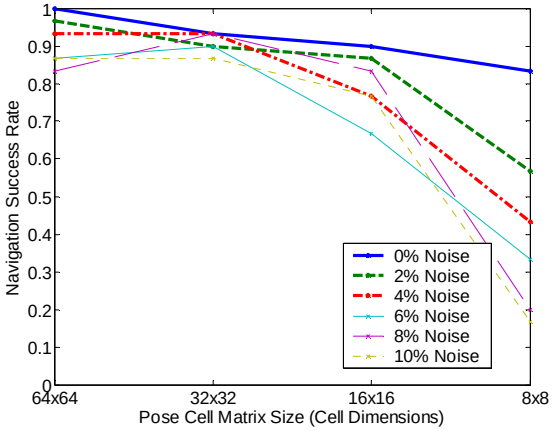

Figure 15 - Goal navigation success rates for varying pose cell structure sizes and visual ambiguity levels.

7 Discussion

The results presented in Section 6 demonstrate that the mapping and goal navigation performance of the RatSLAM system is dependent on both the ratio between the size of the pose cell structure and the environment, as well as on the visual ambiguity of the environment. With low or no visual ambiguity, the system's performance as measured by goal navigation success degrades only gradually as the pose cell structure size is reduced.

In situations of higher visual ambiguity however, there is a limit to how small the pose cell structure can be shrunk before there is a sudden and significant reduction in the robot's ability to navigate successfully to goals. This is most likely due to the experience mapping algorithm's reliance on visual information to distinguish between different experiences that are associated with the same cluster of pose cells. For a pose cell structure size of 8×8 cells, the dominant activity packet is almost the size of the entire pose cell structure, meaning visual information is used almost continuously to differentiate between experiences.

The significance of these results is that to achieve satisfactory RatSLAM system performance, one of two conditions must be met:

- the vision system must have a low false positive recognition rate, or
- the nominal size of the pose cell representation must be at least as large as the environment that is being mapped.

The current RatSLAM implementation achieves a very low false positive recognition rate by matching 12×8 pixel greyscale images using a sum of absolute differences metric using a high sensitivity level. This allows the system to perform goal navigation even when r_s is quite large. However, the sensitivity of the vision module means that places in the environment may be represented by multiple stored scenes, requiring the robot to explore for a longer duration in order to learn all possible representations of each place.

If a vision system with greater generalisation capabilities is used in future work, the value of r_s will need to be reduced. Reducing r_s will increase the quality of the information stored in the pose cell structure and reduce the experience mapping algorithm's reliance on having unique vision information.

8 Conclusion

The work described in this paper has investigated the effect of visual ambiguity and the environment-representation size ratio on RatSLAM mapping and goal navigation performance. The two parameters were found to be related in their effect, with an isolated rise in visual ambiguity or decrease in pose cell structure size causing only a gradual reduction in system performance, for a fixed environment size. However, when both parameters were changed together system performance degraded significantly. The results demonstrated that the experience mapping algorithm increasingly requires unique vision information as the pose cell structure reduces in size. If the vision processing system can be tuned to produce low false positive rates, significant memory and computational reductions can be achieved by shrinking the RatSLAM pose cell structure size. Conversely, if RatSLAM is to be used with a generalising vision system, the pose cell structure must remain large enough to represent the environment with minimal reuse of the cells.

References

[Arleo, 2000] A. Arleo, "Spatial Learning and Navigation in Neuro-mimetic Systems: Modeling the Rat Hippocampus," in *Department of Computer Science*. Lausanne: Swiss Federal Institute of Technology, 2000, pp. 198.

[Browning, 2000] B. Browning, "Biologically Plausible Spatial Navigation for a Mobile Robot," in *Computer Science and Electrical Engineering*. Brisbane: University of Queensland, 2000, pp. 277.

[Fyhn, et al., 2004] M. Fyhn, S. Molden, M. P. Witter, E. I. Moser, and M.-B. Moser, "Spatial Representation in the Entorhinal Cortex," *Science*, vol. 27, pp. 1258-1264, 2004.

[Hafting, et al., 2005] T. Hafting, M. Fyhn, S. Molden, M.-B. Moser, and E. I. Moser, "Microstructure of a spatial map in the entorhinal cortex," *Nature Neuroscience*, vol. 11, pp. 801-806, 2005.

[Kuipers, et al., 2004] B. Kuipers, J. Modayil, P. Beeson, M. MacMahon, and F. Savelli, "Local Metrical and Global Topological Maps in the Hybrid Spatial Semantic Hierarchy," presented at International Conference on Robotics and Automation, New Orleans, United States of America, 2004.

[Milford, et al., 2005] M. J. Milford, D. Prasser, and G. Wyeth, "Experience Mapping: Producing Spatially Continuous Environment Representations using RatSLAM," presented at Australasian Conference on Robotics and Automation, Sydney, Australia, 2005.

[Milford, et al., 2005] M. J. Milford, G. Wyeth, and D. Prasser, "Efficient Goal Directed Navigation using RatSLAM," presented at International Conference on Robotics and Automation, Barcelona, Spain, 2005.

[Milford, et al., 2004] M. J. Milford, G. Wyeth, and D. Prasser, "RatSLAM: A Hippocampal Model for Simultaneous Localization and Mapping," presented at International Conference on Robotics and Automation, New Orleans, USA, 2004.

[Milford, et al., 2004] M. J. Milford, G. Wyeth, and D. Prasser, "Simultaneous Localization and Mapping from Natural Landmarks using RatSLAM," presented at Australasian Conference on Robotics and Automation, Canberra, Australia, 2004.

[Milford, et al., 2006] M. J. Milford, G. F. Wyeth, and D. P. Prasser, "RatSLAM on the Edge: Revealing a Coherent Representation from an Overloaded Rat Brain," presented at International Conference on Robots and Intelligent Systems, Beijing, China, 2006.

[Prasser, et al., 2005] D. Prasser, M. Milford, and G. Wyeth, "Outdoor Simultaneous Localisation and Mapping using RatSLAM," presented at International Conference on Field and Service Robotics, Port Douglas, Australia, 2005.

[Stringer, et al., 2002] S. M. Stringer, E. T. Rolls, T. P. Trappenberg, and I. E. T. de Araujo, "Self-organizing continuous attractor networks and path integration: two-dimensional models of place cells," *Computational Neural Systems*, vol. 13, pp. 429-446, 2002.

[Stringer, et al., 2002] S. M. Stringer, T. P. Trappenberg, E. T. Rolls, and I. E. T. de Araujo, "Self-organizing continuous attractor networks and path integration: one-dimensional models of head direction cells," *Computation in Neural Systems*, vol. 13, pp. 217-242, 2002.

[Yamauchi and Beer, 1996] B. Yamauchi and R. Beer, "Spatial Learning for Navigation in Dynamic Environments," *Systems, Man and Cybernetics*, vol. 26, pp. 496-505, 1996.