

Spatial Mapping and Map Exploitation: A Bio-inspired Engineering Perspective

Michael Milford¹ and Gordon Wyeth²

¹Queensland Brain Institute

²School of Information Technology and Electrical Engineering
The University of Queensland, Brisbane, Australia
{milford, wyeth}@itee.uq.edu.au

Abstract. Probabilistic robot mapping techniques can produce high resolution, accurate maps of large indoor and outdoor environments. However, much less progress has been made towards robots using these maps to perform useful functions such as efficient navigation. This paper describes a pragmatic approach to mapping system development that considers not only the map but also the navigation functionality that the map must provide. We pursue this approach within a bio-inspired mapping context, and use results from robot experiments in indoor and outdoor environments to demonstrate its validity. The research attempts to stimulate new research directions in the field of robot mapping with a proposal for a new approach that has the potential to lead to more complete mapping and navigation systems.

1 Introduction

The spatial mapping problem has been the subject of much research in the robotics field, resulting in the existence of several well-established mapping algorithms [1]. Sensor and environment uncertainty has caused most robotic mapping and navigation methods to converge to probabilistic techniques. The key strength of probabilistic techniques is their ability to deal with uncertainty and ambiguity in a robot's sensors and environment. Any technique that has some means of handling the uncertainties faced by a mobile robot has an immense advantage over techniques that do not. Probabilistic techniques allow a robot to appropriately use sensory measurements based on their modeled uncertainties. Ambiguous features or landmarks in the environment become useful (albeit less useful than unique features) rather than becoming failure points for the mapping algorithm.

Under appropriate conditions some of these systems can solve the core SLAM (Simultaneous Localization And Mapping) problem in large, real world, static environments. The world representations produced by these systems generally take the form of occupancy grids, or landmark maps. However, while a high resolution (often as small as 10 mm) occupancy grid can faithfully represent the physical structure of an environment, it is not necessarily the most 'usable' representation for general navigation tasks. For environments of any significant size, the amount of data stored in such a map becomes quite large, and requires significant processing before it

can be used by a robot to perform tasks such as goal navigation. The limitations of such representations are perhaps revealed by the imbalance between a relative abundance of competent mapping algorithms [2–8] and a scarcity of robotic systems that use these maps to perform navigation and other useful tasks [9–11].

This problem has been partially addressed by other mapping techniques that produce topological or hybrid metric-topological representations. Techniques such as the Hybrid Spatial Semantic Hierarchy [12,13] produce maps that are perhaps more suited to route planning and navigation than purely grid-based maps. By embedding navigational concepts into the representation, such as the locations of exits from a room, the maps produced are already somewhat ‘navigation-ready’ so to speak. A robot need only pick which exit to go through in order to progress towards a goal, rather than process an occupancy grid map to extract possible exit locations.

The concept of map usage shaping the nature of the map applies not only in spatial mapping, but also conceptual and biological mapping domains. In the conceptual domain the structure and form of conceptual maps created by humans is known to change significantly depending on motivation and context [14]. In rodents, place cell firing varies depending on the behavior of the rat at the time [15,16]. Reward location, movement speed, movement direction, and turning angle can all affect place cell firing and even cause remapping of place fields [15]. Furthermore, the place cell maps can gradually change to provide the rat with efficient trajectories between reward sites [17].

Animals are, in fact, excellent examples of the co-evolution of mapping and map usage processes. Foraging honeybees are known as central-place foragers because they have a common point to start and finish their foraging trips. Their ‘maps’ must provide them with the ability to return to their hive after foraging, a journey typically of two or three kilometers, but stretching as far as 13.5 kilometers [18]. When returning to the nest from a foraging location, the desert ant *Cataglyphis fortis* uses its *home vector*, which is calculated during the outbound trip using weighted angular averaging [19]. Primates possess spatial view cells, that allow them to remember where an object was last seen, even if it is currently obscured, a very useful ability in their natural environment. Each animal or insect represents the environment in a manner which suits the task they must perform in it.

This paper presents a pragmatic bio-inspired approach to developing a robot mapping system that considers both the map and the navigation functionality that the map must provide. The research occurred as part of the RatSLAM project, the aim of which was to use models of rodent hippocampus to produce a robot mapping and navigation system [20–23]. By considering the usability of the maps in navigation tasks, the project developed a robot mapping system with significantly different characteristics to those developed with only the SLAM problem in mind. Furthermore, the bio-inspired approach incorporated concepts from three separate mapping fields; robotic mapping, mapping in nature, and computational models of biological mapping systems.

The hypothesis in this paper is that mapping methods for an autonomous robot must necessarily develop in parallel with mechanisms for map exploitation. Only with awareness of the entire mapping and navigation problem will researchers be able to develop autonomous mobile robot systems that can be deployed across all environments and situations. Since biological systems fulfill both mapping and

navigational needs of animals and insects, this hypothesis is pursued from a bio-inspired perspective. We present results demonstrating the capability of a bio-inspired robotic system developed using the proposed approach.

The paper is organized as follows: We start by performing a comparative review of what is known about the mapping and navigation processes in nature and the state of the art robotic mapping techniques. The characteristics of mapping systems in both areas are discussed with respect to solving the entire mapping and navigation problem, rather than just creating a spatial map, resulting in the proposed approach to developing robotic mapping systems. We then describe an example of the pragmatic use and extension of models of biological systems, and the increase in mapping performance that can be achieved by doing so. The parallel development of mapping and navigation processes is illustrated within the context of the RatSLAM system, and results are presented demonstrating the resultant system's effectiveness in indoor and outdoor environments. The paper concludes with a discussion of the significant issues raised by the study.

2 Mapping and Navigation in Nature and Robots

There are many forces driving the diversity that is apparent when examining biological and robot mapping and navigation systems. In nature, creatures and their navigation systems have evolved to suit their environments and their own physical and sensory characteristics. In the natural world, the range of body types, sensory equipment, environments, and lifestyles has produced a myriad of solutions to the problems facing a creature that needs to move around effectively in its environment. Likewise in the research labs of robotics researchers and the domestic home or industrial workplace, mapping and navigation systems have developed in ways to suit the environments and the sensors available, but less so the purpose of the robot.

Given this diversity it seems a challenging task to identify one research direction that is most likely to yield a complete robot mapping and navigation system. The specific mechanisms, representations, and context of each system differ so greatly that direct comparison is not possible. Fortunately, there are a number of fundamental mapping and navigation issues that are common to all these systems. Through consideration of these issues for all these systems it is possible to define goals for future research in this field, and in the process identify one means by which these goals may be met.

The following sections compare and contrast biological systems and models of biological systems with robotic mapping algorithms in a range of areas, with a focus in the biological area on rodents. The comparison highlights the shortcomings of both types of system as well as their complementary characteristics. The shortcomings are particularly relevant when considering where future research into the mapping and navigation problem should concentrate. After a discussion of possible future research approaches, the final section presents a proposed approach for developing complete robot navigation systems.

2.1 Robustness Versus Accuracy

One of the most significant differences between biological mapping systems and probabilistic methods is the world representations that they use. Many probabilistic systems incorporate high resolution occupancy grid maps, such as those shown in Figure 1. In work by Grisetti, Stachniss *et al.* [7] in the Freiburg campus environment, the occupancy grid contained 625 million squares. By contrast the place cells found in the rodent hippocampus appear quite coarse, and likely do not encode information to such a precise degree. The place cells certainly do not represent occupancy, but instead represent the rodent's location in space. The place cells give localization, but the map itself is widely distributed in other regions.

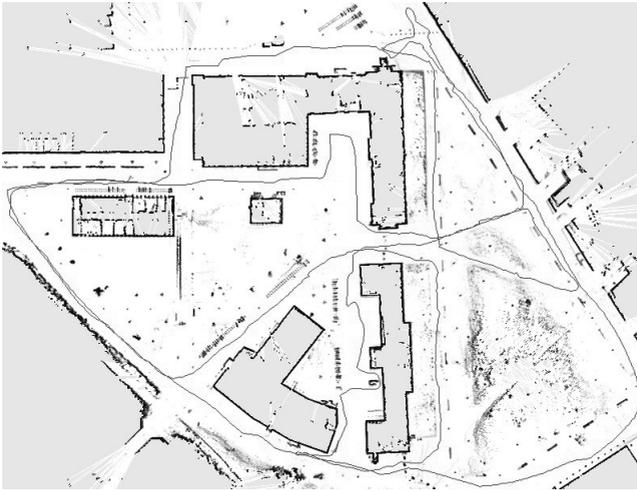


Fig. 1. High resolution grid map of the University of Freiburg campus, measuring 250×250 m, with a grid resolution of 0.01 m. The map was generated using the extended Rao-Blackwellized particle filter method [7]. Adapted from Figure 4 of [7] and reprinted with permission of IEEE. © 2005 IEEE.

Furthermore many probabilistic based mapping methods attempt to produce a Cartesian map with direct spatial correspondence to the physical environment. A 2.5 meter wide, 16.7 meter long corridor should appear as an identically sized region of free space in an occupancy grid map, with occupied cells along its edges. Biological systems relax the requirement that there be direct Cartesian correspondence between a map and the environment. In rodents, place field size and shape can vary significantly depending on the type of environment. For instance, the addition of straight barriers to an environment resulted in the destruction of the place fields located where the barrier was added [24]. Barriers can also cause an apparent remapping from the affected place cells into new place cells. Many animals appear to rely more on robust coping strategies than precise representations of their world. One such example is the desert ant; it navigates back to the general vicinity of its nest and then uses a systematic search procedure to find it [19].

2.2 Map Friendliness Versus Map Usability

Maps that are highly accurate, precise, and Cartesian have a number of advantages and disadvantages. One immediate observation is that they are ‘human-friendly’. Looking at the map in Figure 1, a human can immediately make a number of high level observations about the environment – there are multiple loops in the environment, several buildings, large open spaces, and so on. It is easy for a human to relate parts of the map back to physical places in the environment, and also to compare the map and environment on a global scale.

In contrast, it is hard to present an equivalent map for any animal or insect. Place fields observed from brain readings in the rodent hippocampus convey only the observed locations in the environment to which a single place cell responds. There is no definitive method of combining the environmental features that a rodent perceives with the place fields of each place cell to produce a single ‘conventional’ map. An occupancy grid map could be created by using place cell activity and externally determined ranges to obstacles. It is unlikely however such a map would be representative of the rodent’s actual mapping process. The extent to which rats rely on “range-to-obstacle” concepts is unknown, as is whether rats use a linear geometric scale for ordering the map.

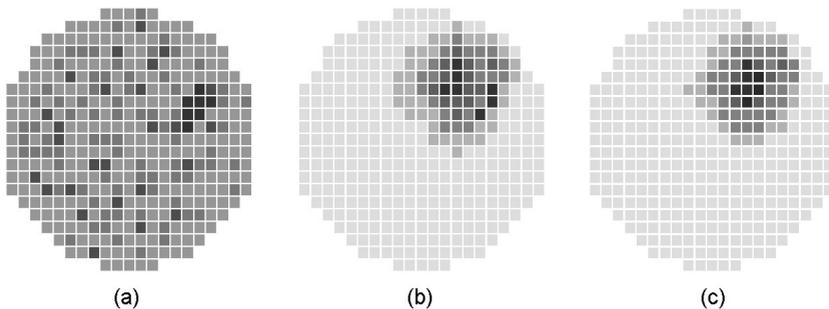


Fig. 2. Sample combined place cell firing rates and area occupancy maps. (a) Time spent by rat at each location in the circular environment, darker squares indicate longer occupancy times. (b) Spike map showing firing rate of one complex-spike cell over the environment, darker squares indicate a higher firing rate. (c) The firing rate array is obtained by dividing the spike map by the time spent in each location, to reveal a single peak of activity.

The maps shown in Figure 1 are human-friendly in that a human can easily understand them. A human could use these maps to navigate efficiently to various locations. However, in doing so a number of high level classification and planning processes would be used. The narrow straight sections of free space in Figure 1 might immediately be classified as pathways, and the large obstacle shapes as buildings. A human might consciously trace along the paths shown in the map, before choosing and following the shortest one. The overall process would be very different to the subconscious navigation that occurs as a person drives to and from work, walks around an office building, or moves around in their house and garden.

This distinction is an important one – the form of most maps created by probabilistic methods is governed by the underlying mathematical algorithms, or

human requirements that the map be in a form suitable for performance analysis. In contrast, biological mapping systems have evolved to produce representations that are suited to the navigation tasks animals must perform every day.

2.3 Sensory Differences

Animals and robot platforms have different sensory equipment and movement capabilities. Probabilistic methods exploit the particular characteristics of the sensors they use. Biological sensors have evolved for important activities such as finding food and mating, with biological navigation systems evolving to suit the sensors and activities. Models of biological systems emulate the biological mapping and navigation mechanisms while using robotic sensors. Some researchers have applied compensatory algorithms to the output of the robotic sensors in order to more closely emulate the biological systems. For instance, the use of certain color spaces minimizes the effect of illumination differences [25], a task that many biological vision sensors achieve effortlessly. Other research approaches have involved plausibly modifying the models so that they take advantage of the differences in sensing equipment. For instance, instead of using a camera, a set of sonar sensors can be used to detect bearings and ranges to environment cues [26].

In the search for a biologically inspired mapping and navigation system, there are two ways to approach this problem of sensory differences. One approach is based on the fact that biological systems manage to perform quite well despite the specific limitations of some of their sensors. Although biological vision sensors can be quite advanced, there is a lack of accurate range sensing equipment (with the exception of such animals as bats). It seems reasonable that given the rapidly increasing computational power of modern computers, it should eventually be possible to create an artificial system equipped only with biologically faithful sensors that can match the navigation performance of animals. This approach dictates the meticulous emulation of the sensory equipment and theorized navigation strategies of animals even if this means purposefully handicapping certain aspects of the systems.

The other approach involves trying to exploit the superior characteristics of some robotic sensors by extending the mapping and navigation models. Given the current inability to match the capabilities of many biological sensors, it also seems reasonable that the superior characteristics of some robotic sensors be exploited to make up for other shortcomings in the models. In this approach the focus is on creating functional robotic systems, rather than faithful replication of proposed biological mechanisms. Biological mapping and navigation mechanisms are modified to accommodate the different characteristics of robot sensors. This approach has received relatively little attention, unlike the fields focusing on biologically inspired mechanical robot design. There has been only a limited amount of research into developing practical biologically inspired robot mapping and navigation systems [27].

2.4 Capability in Real World Environments

Biological navigation systems perform well enough to allow the millions of species of animals using them to function and survive every day. These systems combine advanced sensing, clever mapping and robust navigation mechanisms. There is a reasonable amount of knowledge about the capabilities of their sensors, and

experiments have gathered a significant amount of knowledge about the navigation capabilities of these animals. Many theories have been devised to account for their capabilities, but in some areas research is only starting to scratch the surface. However, there is no question that animals can navigate well in a wide range of complex, dynamic environments.

In most areas, the state of the art in robotic mapping and navigation systems has not yet come close to matching the abilities of animals. In specific subsets of the problem and under certain conditions, these systems may outperform their biological counterparts, but it is with the sacrifice of robustness and flexibility, and is usually accompanied by a host of assumptions and the use of advanced sensing equipment. Nevertheless the best methods can produce impressive maps of large indoor and outdoor environments [6, 7, 10]. Most of the major mapping and navigation problems, such as closing the loop or coping with ambiguity, can be solved by one or another of these methods with appropriate assumptions.

In contrast to the robotic systems, computational models of the rodent hippocampus have only been used in simulation or on robots in small structured environments, and are yet to solve many of the major mapping and navigation problems. The small size and limited complexity of these environments reduces the major mapping and navigation challenges such as closing the loop, dealing with extended sensory ambiguity, and navigating to goals to trivial or non-existent problems. None of these models have been tested in or successfully applied to the large unmodified environments in which one might reasonably expect an autonomous mobile robot to function, such as an office floor or an outdoor set of pathways.

Biologically inspired models are partly limited because the goal for much of the research is to test navigational theories for a particular animal, rather than to produce a fully functional robot system [28]. The uncertainty about biological systems and subsequent speculation has produced models that may only be partially faithful to the biology, with resulting navigation performance that is inferior to that of the animal. In pursuing the development of a biologically plausible model, it is unlikely that a researcher will stumble upon a better performing model by chance – current biological systems are the product of a long sequence of evolution.

2.5 Proposed Approach

Given the current state of robotic and biological mapping and navigation systems, several conclusions can be drawn. It is unlikely that research in the near future will create a perfect model of the higher level mapping and navigation systems, such as those of a rodent, primate, or human. Animal brains are only partially understood; researchers create theories from the readings of a few dozen electrodes, but the theories are far from being comprehensively proven. Even though models may recreate some aspects of animal navigation behavior, there can be no real confidence that the underlying mechanisms driving the artificial system are the same as those in the real animal. Furthermore, even biological systems do not necessarily possess all the capabilities autonomous robots require to function in the challenging environments earmarked for their deployment.

Conventional robot mapping and navigation research is also facing many challenges. The most impressive recent demonstrations of mobile robotics have been

largely made possible by good engineering and incremental improvements in algorithms, sensors, and computational power. The Defence Advanced Research Projects Agency (DARPA) Grand Challenge of 2005 is one prime example of this; while the onboard navigation and mapping systems were state of the art, it was only with an impressive array of expensive, high precision sensing equipment that this specific task could be solved [29, 30]. Some may argue that the continual improvement in sensors and computer power may eventually lead to navigation systems that surpass all the abilities of one of the most adept of navigators – humans. However, it is perhaps more likely that this milestone will be achieved through fundamental methodology changes, rather than steady computational and sensory improvements, although such changes will definitely facilitate the process.

So, where to look for a solution to the mapping and navigation problem? This paper proposes that an eventual solution may be found using a biologically inspired yet completely pragmatic approach, which considers both mapping and navigational requirements. Previous research has investigated bio-inspired robotic mapping methods, although without concurrent consideration of the how the map would be used [31]. However, research on pragmatic models of biological systems to date has still had a heavy emphasis on biological plausibility and has had limited practical success. No research has developed or tested a biologically inspired model under the same conditions and criteria as conventional robot mapping and navigation systems.

For example, the biologically inspired model developed by Arleo [32] can perform localization, mapping and goal recall, but only in a very small artificial arena, with continuous visual input from artificial cues or a distal cue source. Other less biologically faithful models have displayed limited re-localization ability in robot experiments in an artificial arena with visual cues, and have been able to navigate to goals in simple T and two arm mazes in simulation [31]. These approaches have also been fundamentally limited by the rodents on which they are based. The capabilities of rodent mapping and navigation systems do not necessarily fulfill all the desired capabilities of an autonomous mobile robot. If a biologically inspired system is ever to compete with conventional probabilistic methods, it must also contain solutions to the navigational shortcomings of animals. The approach presented in this paper therefore also sought to investigate whether models of the rodent hippocampus can serve as a basis for a complete robot mapping and navigation system.

3 Pragmatic Biological Modeling

To evaluate the potential of current hippocampal models as the basis for a robot navigation system, a model was developed based on the structure shown in Figure 3. Experiments were run on a Pioneer 2DXe mobile robot from Mobile Robots Inc (formerly ActivMedia Robotics). Ideothetic information for path integration is derived from the robot's wheel encoders. Allothetic information for generating the local view is derived from camera images. The testing environment was a flat two by two meter area of a laboratory floor. Colored cylinders were used as visual cues for the robot and were placed just outside the arena in various configurations.

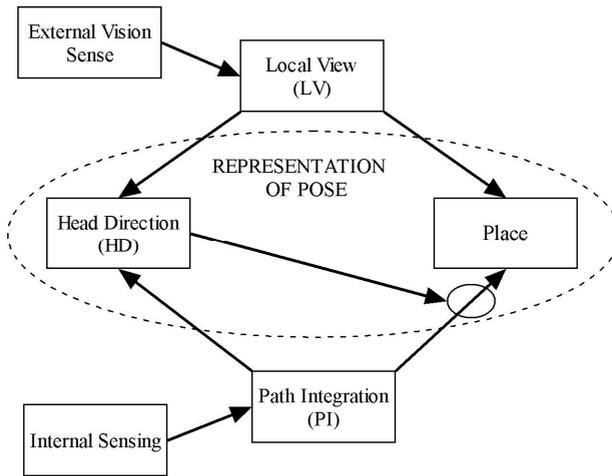


Fig. 3. Overall structure of the hippocampal model

The initial study was based on well known models of the hippocampus [32–34] as shown in Figure 3. The head-direction cells were implemented using a competitive attractor network (CAN) made up of neural units or ‘cells’ roughly corresponding to biological head-direction cells. Each cell is tuned to be maximally active when the robot’s orientation matches the cell’s preferred direction. The cell’s activity level reduces as the robot orientation rotates away from this preferred direction. The cell arrangement reflects their associated robot orientations – nearby cells encode similar robot orientations. When the ensemble activity of the head-direction cells is viewed as a graph, one can see a ‘packet of activity’ that resembles a Gaussian curve (see Figure 4 for examples). The center of this ‘activity packet’ represents the current perceived orientation of the robot.

The place cells were modeled as a two-dimensional CAN, with each cell tuned to be maximally active when the robot is at a specific location. A coarse representation is used, with the path integration system tuned so that each place cell represents a physical area of approximately 250 mm × 250 mm. The cells are arranged in a two-dimensional matrix with full excitatory interconnectivity between all cells. Further details on the implementation can be found in [33], and is similar to the extended implementation described in later sections.

3.1 Experimental Performance

In rotation-only experiments, the robot was able to maintain an accurate estimate of its orientation. However, when the robot moved around the environment, the network’s tracking ability proved to be unstable. Over the period of an hour the robot became lost and its perceived location moved well outside its two by two meter arena.

Closer examination revealed that the place cells and head-direction cells developed multiple firing fields in situations of ambiguous sensory input. While activity in the head-direction cells shifted to represent appropriate changes in the robot’s orientation, activity in the place cells did not shift appropriately to represent robot translation (Figure 4).

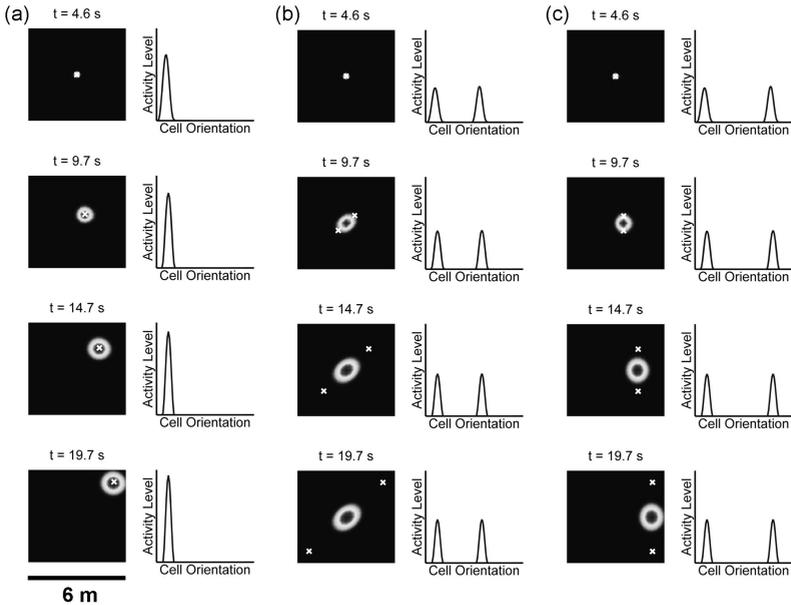


Fig. 4. Place and head-direction cell firing fields under unique and ambiguous conditions. Crosses show the correct location hypotheses (a) Correct update of place cell activity under a unique orientation hypothesis. (b) With opposing orientation hypotheses, place cell firing smears, and does not shift to represent either possible location estimate. (c) Place cell activity shifts in the net (and consequently incorrect) direction of two orientation hypotheses varying by 90 degrees.

The underlying cause of the failures was the separation of pose (the conjunction of location and orientation) into separate location and orientation representations in the place cells and head-direction cells. Without a strong coupling between the two cell types, allothetic input can incompletely correct errors accumulated during exploration leading to corruption in the associative memory. Most significantly, the separation of the state representation removes the ability to correctly update multiple estimates of pose in perceptually ambiguous situations. Because place cells have no directional attributes, any path integration process, regardless of the specific mechanism, must obtain its orientation information from activity in the head-direction cells. However, each location estimate can only be updated by its one-to-many association with all orientation estimates, rather than updated only by the orientation estimate pertinent to the particular location hypothesis. Consequently, location estimates rapidly become incorrect. The correct robot location estimate can only be reinstated through strong unique visual input.

Separate representations of robot orientation and spatial location are inherently unsuitable for mapping and navigation in large, ambiguous environments, as demonstrated by the work presented in this section. The following section describes the implementation of an extended hippocampal model known as RatSLAM, which combines the concept of head-direction and place cells to form a new type of cell known as a *pose cell*.

3.2 A Model of Spatial Pose

To form a population of pose cells, the competitive attractor network structures that were used to model head-direction and place cells were combined to form a three-dimensional structure (Figure 5). The pose cells are arranged in an (x', y', θ') pattern, allowing the robot to simultaneously represent multiple pose estimates in x' , y' , and θ' . Primed co-ordinate variables are used because although the correspondence between cells and physical co-ordinates is initially Cartesian, this relationship can become discontinuous and non-linear as the system learns the environment. For example, in indoor experiments, each pose cell initially fires maximally when the robot is in a 0.25×0.25 meter area and within a 10 degree band of orientation. However as an experiment progresses, the pose volume each pose cell corresponds to can grow, shrink, warp, or even disappear under the influence of visual information.

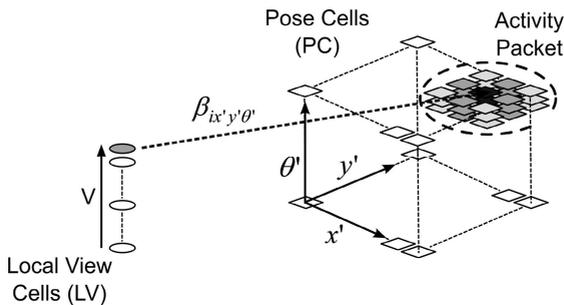


Fig. 5. The local view and pose cell structures. The local view encodes visual information about the environment. The pose cells represent the robot's pose. Co-activated local view and pose cells form associative links. A familiar visual scene activates a local view cell which in turn injects activity into the pose cells associated with it. Re-localization occurs when the activity packet caused by visual input becomes dominant.

The advantage of a unified pose representation is that each cell encodes its own preferred orientation. This allows multiple hypotheses of the robot's pose to propagate in different directions, unlike the more conventional head-direction – place cell model, and enables the mapping of large, ambiguous indoor and outdoor environments (Figure 7c, d) [21, 35].

4 Contextual Mapping

Conventional robot mapping methods typically produce either occupancy grid maps [7] or landmark-based maps [36]. Such maps are appealing from a mathematical perspective because they facilitate the use of probabilistic mapping algorithms. However, such maps are context-less, and store no information regarding the environment that can be used directly by the robot. For example, while an occupancy grid map stores the location of free space and obstacles in the environment, the map does not directly tell the robot about the speeds at which it can traverse different areas

of the environment. A cluttered area might require slow movement with much turning, while a long straight corridor offers the potential for a high speed movement. The occupancy map could perhaps be processed using a ‘clutter’ metric in order to yield this information, but this is a complex and perhaps unnecessary step.

Contextual mapping involves enriching the spatial representation of an environment with additional information relevant to the potential uses of the representation. For example, temporal or speed data provides information about the rate at which different areas of the environment can be traversed, which can be used by a robot to plan the quickest route to a goal. Behavioral data provides information about the appropriate movement behaviors required to move between different places in the environment, which a robot can use to achieve improved movement through the environment. Transition success rate data provides information about how easy it is to cross different areas in the environment, which a robot can use to decide between competing routes to a goal. This section presents a contextual approach to mapping the environment. The map use requirements were that the robot be able to rapidly explore and map its environment, plan and execute the quickest routes to goals, and adapt its map and goal navigation to simple changes in the environment.

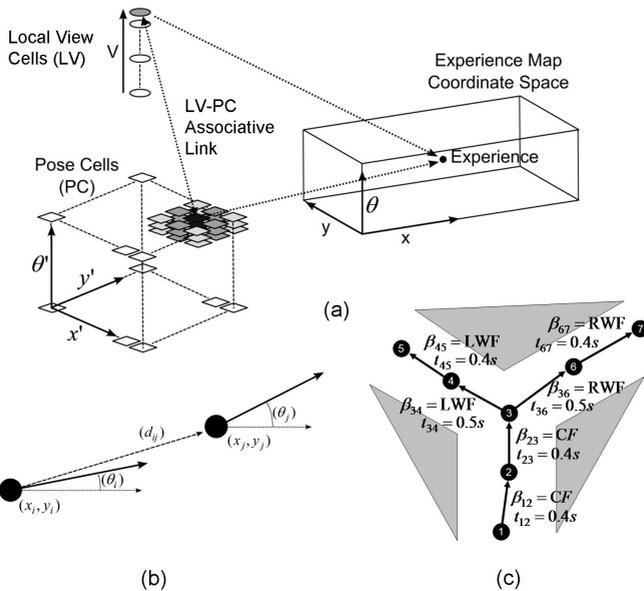


Fig. 6. (a) Experience map co-ordinate space. An experience is associated with certain pose and local view cells, but exists within the experience map's own (x, y, θ) co-ordinate space. (b) A transition between two experiences. Shaded circles and bold arrows show the actual pose of the experiences. d_{ij} is the odometric distance between the two experiences. (c) An example of the behavioral and temporal information. The numbered circles represent experiences, and CF – Centerline Following; LWF – Left Wall Following; RWF – Right Wall Following.

4.1 Experience Mapping

The experience mapping algorithm was developed to create a contextual representation of the environment containing not only a spatial map, but also movement, behavioral, and temporal information relevant to the robot's motion and higher level behaviors such as goal navigation.

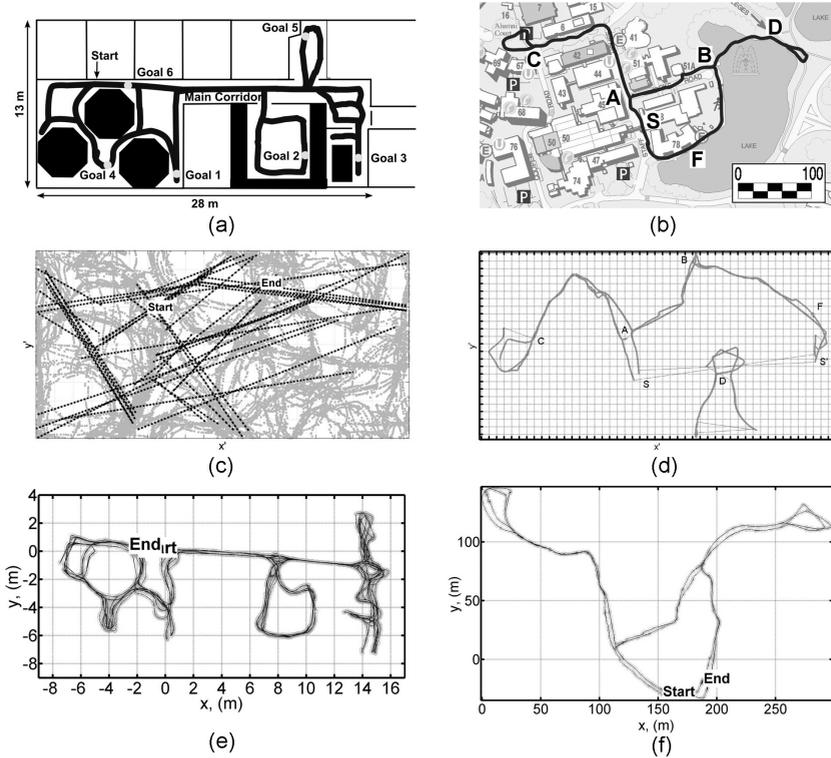


Fig. 7. (a-b) Indoor and outdoor environments. The robot's paths are shown by the thick line. In the outdoor environment the robot started at S and finished at F. An initial loop of the inner loop SABFS was followed by a clockwise traverse of the outer loop SACABDBFS and then a counter-clockwise traverse of the same outer loop. (c-d) The trajectory of the most strongly activated pose cell in (x, y) space, with wrapping in both directions. The thin lines indicate re-localization jumps where the RatSLAM system closed the loop. Each grid square represents 4×4 pose cells in the (x, y) plane. (e-f) Experience maps.

Activity in the pose cells and local view cells drives the creation of experiences. Experiences are represented by nodes in (x, y, θ) space connected by links representing transitions between experiences. Each experience represents a snapshot of the activity within the pose cells and local view cells at a certain time, and in effect represents a specific spatial and visual robot experience (see Figure 6a).

As the robot moves around the environment, the experience mapping algorithm also learns experience *transitions*, which are links between experience nodes.

Transitions store information about the physical movement of the robot between one experience and another, the movement behavior used during the transition, the time duration of the transition, and the traversability of the transition (see Figure 6b, c). This information enables a robot to perform rapid exploration, map correction, route planning and execution, and adaptation to environment change – all with minimal computation [22, 23, 37].

5 Map Exploitation

5.1 Exploration

In a novel environment, an animal or robot must explore in order to acquire information about its surroundings. In classical robotics, exploration is often achieved through greedy frontier-based techniques, where the robot heads towards the nearest unknown area of the environment, or the unknown area which offers the greatest potential information gain [38, 39]. These techniques process occupancy or coverage maps to identify borders of unknown areas in the environment.

By encoding behaviors into the spatial experience map, the RatSLAM system is able to choose movement behaviors that have not yet been used at *behavior intersections* in the environment, where more than one movement is possible, with minimal map processing (Figure 8). Past behavior usage can be detected using the experience map. A recursive tree search function is used to find routes from the currently active experience through the experience map using only transitions of a specific behavior type. If a sufficiently long continuous path is found through the experience map for a certain behavior, that behavior is tagged as having already been followed from the current environment location. During exploration the robot attempts to use previously unused movement behaviors to try to experience new routes through the environment.

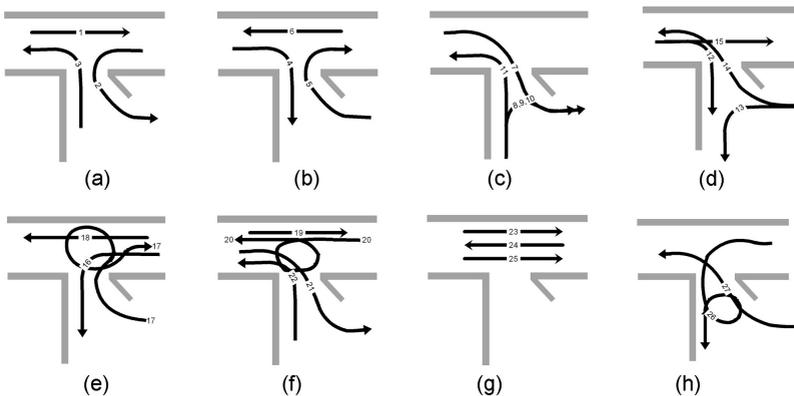


Fig. 8. The robot’s local movement at an intersection during an 80 minute experiment. The labels indicate the order in which the paths were traversed.

5.2 Route Planning and Execution

The experience maps provide the spatial, temporal, and behavioral information used to perform goal directed navigation. Temporal information allows the creation of a temporal map, with gradient climbing techniques used to find the fastest (in the robot's experience) route to the goal. Once a route has been found, the robot uses spatial and behavioral information stored in the experience links to navigate to the goal. If the context was different – for instance, if the robot was required to find the straightest routes to the goal, that information could be added to the map.

By creating maps which are both spatial and navigational, a robot is easily able to use them to navigate to goals. The navigation capabilities of the system were tested over a sequence of 12 goal navigation trials in the environment shown in Figure 7a. Figure 9 shows the temporal maps, planned and executed routes to six goals. The robot successfully navigated to all twelve goals, demonstrating the usability of the spatial and navigational experience maps.

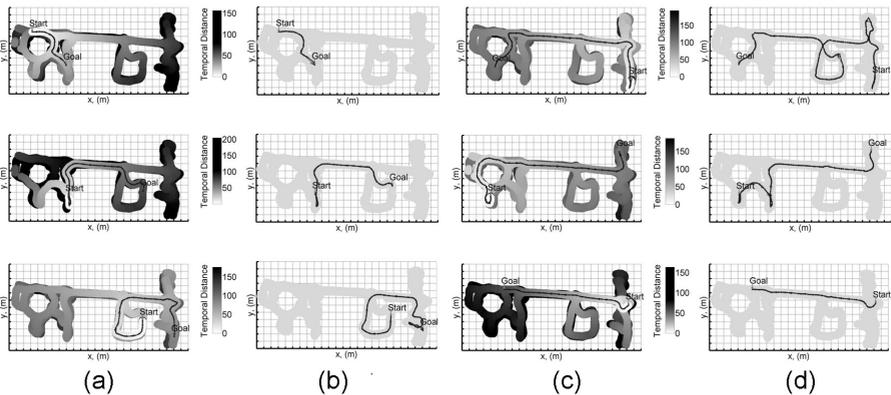


Fig. 9. Temporal maps, planned and actual routes to the goals for six of the twelve trials. Temporal maps and planned routes are shown in the first and third columns (a, c). The actual routes executed by the robot are shown in the second and fourth column (b, d). Five of the executed routes were near optimal, with major errors in one route.

5.3 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 behavior. If enough 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. In this way the robot is able to adapt to simple changes in the environment without any explicit recognition of objects. The system's ability to adapt to changes in the environment was demonstrated in an indoor environment, with the results shown in [22].

6 Discussion

By testing hippocampal models on actual robots in real world environments, it was possible to evaluate their functional navigation performance from a robotics context, rather than just their biological plausibility. The results indicate that functional navigation performance comparable to other more developed navigation systems can be obtained, and demonstrate the potential of a pragmatic bio-inspired approach to the mapping and navigation problem.

However, it might be argued that we have strayed significantly from the original biological models in our research, and that we have, in fact, demonstrated that biological models don't translate well to robot systems. The key finding of the pilot study was that a separate head-direction – place representation system, common to most biological models, could not maintain multiple hypotheses of the robot's pose over time. The separation of the two networks meant that opposing hypotheses of robot orientation were effectively averaged causing the hypotheses of the robot's location to be updated erroneously. To overcome this problem we introduced pose cells, which represent a conjunction of orientation with place. However, even with pose cells, the representation could not be used for meaningful tasks. Discontinuities, ambiguities and collisions in the pose cell representations made path planning impossible.

With consideration of both mapping and navigational requirements, the experience mapping system was developed, which created representations ready-made for navigation and also map adaptation. Surely all of this invention would indicate that the biological models were of little use to translate into a robotic system? We would argue that it is not an error in translation from biology to robotics, but an error in the original biological models themselves. Recent evidence of *grid cells* [40, 41] in the entorhinal cortex (a region of the brain near the hippocampus) indicates that this region of brain integrates ideothetic and allothetic information into a full pose representation – as do the pose cells in the RatSLAM model. Grid cells exhibit many of the same properties of pose cells, such as discontinuity, ambiguity and redundancy. It would seem that the pose cells invented during this research may be closer to biological fact than the previously proposed biological models.

If it is the role of the grid cells to perform the integration of ideothetic and allothetic cues, then the hippocampus may store context-specific episodic sequences combining a number of cues (including the pose information in the grid cells). Perhaps the well defined place representations found in the experience map form a closer analogy to the hippocampal place cells. In future work, we intend to pursue the links from our model to biology. If biological plausibility can be shown, then the RatSLAM model allows predictions to be made regarding neuro-cognitive models of rodent navigation in larger environments than the usual small mazes and arenas.

From a robotics perspective, this research has demonstrated that it is possible to create a bio-inspired robot mapping and navigation system that can function in real-time in large indoor and outdoor environments. Furthermore, the work is one of the few attempts to develop a complete integrated navigation system – one that addresses not only the primary SLAM problem, but also the problems of exploration, navigating to goals, and adaptation. The success of the system can be attributed to the consideration of both mapping and navigation requirements during development, and to the pragmatic adaptation of the biological navigation models on which it is based.

References

1. Thrun, S.: Robotic Mapping: A Survey. In: Exploring Artificial Intelligence in the New Millennium, Morgan Kaufmann, San Francisco (2002)
2. Dissanayake, G., Durrant-Whyte, H., Bailey, T.: A computationally efficient solution to the simultaneous localisation and map building (SLAM) problem. In: 2000 IEEE International Conference on Robotics and Automation, San Francisco, USA, IEEE, Los Alamitos (2000)
3. Newman, P.M., Leonard, J.J.: Consistent, Convergent, and Constant-time SLAM. In: International Joint Conference on Artificial Intelligence, Acapulco, Mexico (2003)
4. Thrun, S.: Probabilistic Algorithms in Robotics. *AI Magazine* 21, 93–109 (2000)
5. Montemerlo, M., Thrun, S., Koller, D., Wegbreit, B.: FastSLAM: A Factored Solution to the Simultaneous Localization and Mapping Problem. In: AAAI National Conference on Artificial Intelligence, Edmonton, Canada (2002)
6. Montemerlo, M., Thrun, S., Koller, D., Wegbreit, B.: FastSLAM 2.0: An improved particle filtering algorithm for simultaneous localization and mapping that provably converges. In: International Joint Conference on Artificial Intelligence, Acapulco, Mexico (2003)
7. Grisetti, G., Stachniss, C., Burgard, W.: Improving Grid Based SLAM with Rao Blackwellized Particle Filters by Adaptive Proposals and Selective Resampling. In: International Conference on Robotics and Automation, Barcelona, Spain (2005)
8. Thrun, S., Koller, D., Ghahramani, Z., Durrant-Whyte, H., Ng, A.Y.: Simultaneous Mapping and Localization With Sparse Extended Information Filters: Theory and Initial Results. In: Fifth International Workshop on Algorithmic Foundations of Robotics, Nice, France (2002)
9. Yamauchi, B., Beer, R.: Spatial Learning for Navigation in Dynamic Environments. *Man and Cybernetics* 26, 496–505 (1996)
10. Thrun, S.: Probabilistic Algorithms and the Interactive Museum Tour-Guide Robot Minerva. *International Journal of Robotics Research* 19, 972–999 (2000)
11. Stentz, A., Hebert, M.: A Complete Navigation System for Goal Acquisition in Unknown Environments. *Autonomous Robots* 2, 127–145 (1995)
12. Kuipers, B., Byun, Y.T.: A Robot Exploration and Mapping Strategy Based on a Semantic Hierarchy of Spatial Representations. *Robotics and Autonomous Systems* 8, 47–63 (1991)
13. Kuipers, B., Modayil, J., Beeson, P., MacMahon, M., Savelli, F.: Local Metrical and Global Topological Maps in the Hybrid Spatial Semantic Hierarchy. In: International Conference on Robotics and Automation, New Orleans, USA (2004)
14. Hmelo-Silver, C.E., Pfeffer, M.G.: Comparing expert and novice understanding of a complex system from the perspective of structures, behaviors, and functions. *Cognitive Science* 28, 127–138 (2003)
15. Kobayashi, T., Nishijo, H., Fukuda, M., Bures, J., Ono, T.: Task-Dependent Representations in Rat Hippocampal Place Neurons. *The Journal of Neurophysiology* 78, 597–613 (1997)
16. Bostock, E., Muller, R., Kubie, J.: Experience-dependent modifications of hippocampal place cell firing. *Hippocampus* 1, 193–205 (1991)
17. Kobayashi, T., Tran, A.H., Nishijo, H., Ono, T., Matsumoto, G.: Contribution of hippocampal place cell activity to learning and formation of goal-directed navigation in rats. *Neuroscience* 117, 1025–1035 (2003)
18. Frisch, K.v.: *The Dance Language and Orientation of Bees*. Harvard University Press, Cambridge (1967)

19. Wehner, R., Gallizzi, K., Frei, C., Vesely, M.: Calibration processes in desert ant navigation: vector courses and systematic search. *Journal of Comparative Physiology A: Sensory, Neural, and Behavioral Physiology* 188, 683–693 (2002)
20. Milford, M.J., Wyeth, G., Prasser, D.: RatSLAM: A Hippocampal Model for Simultaneous Localization and Mapping. In: *International Conference on Robotics and Automation*, New Orleans, USA (2004)
21. Prasser, D., Milford, M., Wyeth, G.: Outdoor Simultaneous Localisation and Mapping using RatSLAM. In: *International Conference on Field and Service Robotics*, Port Douglas, Australia (2005)
22. Milford, M.J., Wyeth, G.F., Prasser, D.P.: RatSLAM on the Edge: Revealing a Coherent Representation from an Overloaded Rat Brain. In: *International Conference on Robots and Intelligent Systems*, Beijing, China (2006)
23. Milford, M., Schulz, R., Prasser, D., Wyeth, G., Wiles, J.: Learning Spatial Concepts from RatSLAM Representations. *Robots and Autonomous Systems* (2006)
24. Muller, R.U., Kubie, J.L.: The effects of changes in the environment on the spatial firing of hippocampal complex-spike cells. *The Journal of Neuroscience* 7, 1951–1968 (1987)
25. Tews, A., Robert, J., Roberts, J., Usher, K.: Is the Sun Too Bright in Queensland? An Approach to Robust Outdoor Colour Beacon Detection. In: *Australasian Conference on Robotics and Automation*, Sydney, Australia (2005)
26. Recce, M., Harris, K.D.: Memory for places: A navigational model in support of Marr's theory of hippocampal function. *Hippocampus* 6, 735–748 (1998)
27. Gaussier, P., Zrehen, S.: Navigating With an Animal Brain: A Neural Network for Landmark Identification and Navigation. In: *Intelligent Vehicles*, 399–404 (1994)
28. Franz, M.O., Mallot, H.A.: Biomimetic robot navigation. *Robotics and Autonomous Systems*, 133–153 (2000)
29. Crane III, C., Armstrong Jr., D., Torrie, M., Gray, S.: Autonomous Ground Vehicle Technologies Applied to the DARPA Grand Challenge. In: *International Conference on Control, Automation, and Systems*, Bangkok, Thailand (2004)
30. Ozguner, U., Redmill, K.A., Broggi, A.: Team TerraMax and the DARPA grand challenge: a general overview. In: *Intelligent Vehicles Symposium*, Parma, Italy (2004)
31. Browning, B.: Biologically Plausible Spatial Navigation for a Mobile Robot. PhD, Computer Science and Electrical Engineering, University of Queensland, Brisbane (2000)
32. Arleo, A.: Spatial Learning and Navigation in Neuro-mimetic Systems: Modeling the Rat Hippocampus. PhD, Department of Computer Science, Swiss Federal Institute of Technology, Lausanne (2000)
33. Milford, M.J., Wyeth, G.: Hippocampal Models for Simultaneous Localisation and Mapping on an Autonomous Robot. In: *Australasian Conference on Robotics and Automation*, Brisbane, Australia (2003)
34. Redish, D.: *Beyond the Cognitive Map*. Massachusetts Institute of Technology, Massachusetts (1999)
35. Milford, M.J., Wyeth, G., Prasser, D.: Simultaneous Localization and Mapping from Natural Landmarks using RatSLAM. In: *Australasian Conference on Robotics and Automation*, Canberra, Australia (2004)
36. Nieto, J., Guivant, J., Nebot, E.: Real Time Data Association for FastSLAM. In: *International Conference on Robotics & Automation*, Taipei, Taiwan (2003)
37. Milford, M.J., Prasser, D., Wyeth, G.: Experience Mapping: Producing Spatially Continuous Environment Representations using RatSLAM. In: *Australasian Conference on Robotics and Automation*, Sydney, Australia (2005)

38. Stachniss, C., Burgard, W.: Mapping and Exploration with Mobile Robots using Coverage Maps. In: International Conference on Intelligent Robots and Systems, Las Vegas, USA (2003)
39. Yamauchi, B.: A Frontier Based Approach for Autonomous Exploration. In: International Symposium on Computational Intelligence in Robotics and Automation, Monterey, USA (1997)
40. Fyhn, M., Molden, S., Witter, M.P., Moser, E.I., Moser, M.-B.: Spatial Representation in the Entorhinal Cortex. *Science* 27, 1258–1264 (2004)
41. Hafting, T., Fyhn, M., Molden, S., Moser, M.-B., Moser, E.I.: Microstructure of a spatial map in the entorhinal cortex. *Nature Neuroscience* 11, 801–806 (2005)