

A hierarchical model of goal directed navigation selects trajectories in a visual environment



Uğur M. Erdem^{a,*}, Michael J. Milford^b, Michael E. Hasselmo^a

^a Center for Memory and Brain and Graduate Program for Neuroscience, Boston University, 2 Cummington Mall, Boston, MA 02215, USA

^b School of Electrical Engineering and Computer Science, Queensland University of Technology, 2 George St, Brisbane, QLD 4000, Australia

ARTICLE INFO

Article history:

Received 3 September 2013

Revised 17 June 2014

Accepted 9 July 2014

Available online 29 July 2014

Keywords:

Navigation
Path planning
Grid cell
Place cell
Hippocampus
SLAM
RatSLAM

ABSTRACT

We have developed a Hierarchical Look-Ahead Trajectory Model (HiLAM) that incorporates the firing pattern of medial entorhinal grid cells in a planning circuit that includes interactions with hippocampus and prefrontal cortex. We show the model's flexibility in representing large real world environments using odometry information obtained from challenging video sequences. We acquire the visual data from a camera mounted on a small tele-operated vehicle. The camera has a panoramic field of view with its focal point approximately 5 cm above the ground level, similar to what would be expected from a rat's point of view. Using established algorithms for calculating perceptual speed from the apparent rate of visual change over time, we generate raw dead reckoning information which loses spatial fidelity over time due to error accumulation. We rectify the loss of fidelity by exploiting the loop-closure detection ability of a biologically inspired, robot navigation model termed RatSLAM. The rectified motion information serves as a velocity input to the HiLAM to encode the environment in the form of grid cell and place cell maps. Finally, we show goal directed path planning results of HiLAM in two different environments, an indoor square maze used in rodent experiments and an outdoor arena more than two orders of magnitude larger than the indoor maze. Together these results bridge for the first time the gap between higher fidelity bio-inspired navigation models (HiLAM) and more abstracted but highly functional bio-inspired robotic mapping systems (RatSLAM), and move from simulated environments into real-world studies in rodent-sized arenas and beyond.

© 2014 Elsevier Inc. All rights reserved.

1. Introduction

The ability to successfully navigate to a predefined location is often a life crucial task for many higher order organisms. The goal location might be a food source, a temporary shelter, a nest, or some other desired location. Squirrels are effective at rediscovering their previously stashed food sources (Jacobs & Liman, 1991). Rats can learn to revisit or to avoid known food locations (Brown, 2011; Olton & Schlosberg, 1978). Mice learn to avoid an unpleasant environment, such as a water-maze, by finding an out-of-sight escape platform after only a handful of learning trials (Morris, Garrud, Rawlins, & O'Keefe, 1982; Redish & Touretzky, 1998; Steele & Morris, 1999). If a visible goal location is in the field-of-view of the agent, the navigation task becomes trivial: The agent proceeds towards the visible goal location avoiding potential obstacles on the way. However, if the goal location is out of visual range or hidden (as in the water-maze) then navigation mechanisms based

on cognitive capabilities that can exploit the previously encoded and currently out of view goal location become important to guide the agent to the goal. Such a navigation mechanism would not necessarily need to pinpoint the goal location. It would be sufficient to guide the agent to the general goal location neighborhood such that the goal is in the visual range of the agent. Consequently, the visually driven navigation system can take over to home the agent into the goal location, an approach that has been used successfully by the robotic mapping system used in this research (Milford & Wyeth, 2009).

There is compelling evidence gathered from physiological and behavioral data suggesting the existence of spatial cognitive mechanisms in the brain representing the agent's spatial environment and aiding it during goal-directed navigation experiments. The entorhinal cortex and hippocampus play a role in goal-directed behavior towards recently learned spatial locations in an environment. Rats show impairments in finding the spatial location of a hidden platform in the Morris water-maze after lesions of the hippocampus, postsubiculum, or entorhinal cortex (Morris, Garrud, Rawlins, & O'Keefe, 1982; Steele & Morris, 1999; Steffenach,

* Corresponding author.

Witter, Moser, & Moser, 2005; Taube, Kesslak, & Cotman, 1992). Recordings from several brain areas in behaving rats show neural spiking activity relevant to goal-directed spatial behavior, including grid cells in the entorhinal cortex that fire when the rat is in a repeating regular array of locations in the environment falling on the vertices of tightly packed equilateral triangles (Hafting, Fyhn, Molden, Moser, & Moser, 2005), place cells in the hippocampus that respond to mostly unique spatial locations (O'Keefe & Nadel, 1978), head direction cells in the postsubiculum that respond to narrow ranges of allocentric head direction (Taube, 2007), and cells that respond to translational speed of running (O'Keefe, Burgess, Donnett, Jeffery, & Maguire, 1998).

Some of the evidence related to the goal-directed navigation planning include forward sweeping events of spiking activity in rat place cell ensembles that have been observed during vicarious trial and error experiments (Johnson & Redish, 2007; Pfeiffer & Foster, 2013) and sharp wave ripple events during goal-directed spatial tasks (Davidson, Kloosterman, & Wilson, 2009; Foster & Wilson, 2006; Jadhav, Kemere, German, & Frank, 2012; Louie & Wilson, 2001). Furthermore, brief sequences of place cell ensemble activity encoding trajectories from an agent's current location have been observed to be strongly biased towards the agent's predicted goal location (Pfeiffer & Foster, 2013).

In this work we combine two biologically inspired models that generate and maintain representations of their environment as collections of simulated spatially tuned neurons such as grid cells and place cells.

The first one of these models is the RatSLAM model (Milford, Wyeth, & Prasser, 2004) which has been implemented on real robotic agents and has been shown to match or outperform the state of the art probabilistic robotic systems in encoding and navigating large environments over long periods of time (Milford & Wyeth, 2009; Prasser, Milford, & Wyeth, 2006). However, the current RatSLAM model is not easily scalable and its goal directed navigation module is less biologically plausible than its Simultaneous Localization and Mapping (SLAM) component.

The second model we use in our work is the HiLAM (Erdem & Hasselmo, 2013), a biologically inspired goal-directed navigation model based on look-ahead trajectories in a hierarchical collection of simulated grid cells and place cells. While HiLAM is highly capable in simulating behavioral goal-directed navigation experiments, it is prone to failure in the presence of noisy and degraded input, since it does not have mechanisms in place to detect and to correct for the stochastic loss of fidelity in its state representation. Consequently, like many other high fidelity computational models, the HiLAM has not been previously tested on real life data.

In this work we combine the RatSLAM model and the HiLAM such that their individual fortes complement each other in generating and maintaining stable spatial maps using real life visual data (RatSLAM) and in using the generated maps for goal-directed path planning in a biologically plausible manner (HiLAM).

2. Material and methods

The framework presented in this work shows collaboration between two previously developed computational models for spatial mapping and navigation. While the RatSLAM model generates rectified odometry data, the Hierarchical Look-Ahead Trajectory Model (HiLAM) provides a mechanism for goal directed navigation. We also show the scalability of the HiLAM using odometry data extracted from noisy real-life visual information collected from a small remote controlled vehicle referred to as the "agent". Using ground truth extracted from external cameras, we show the goal directed navigation accuracy in two environments, a small open-field square indoor maze and an outdoor area that is larger

than the indoor maze by two orders of magnitude. We first extract the unrectified odometry data from the optic flow information implicit in the camera's field of view. Then, we rectify the raw odometry data by detecting loop-closure points in time and space using the RatSLAM model. Finally, we form spatial representations using grid cells and place cells in the HiLAM and select trajectories to goal locations using hierarchical linear look-ahead probes in this model.

2.1. Hierarchical Look-Ahead Trajectory Model (HiLAM)

In the HiLAM, head direction cells modulated by proprioceptive velocity data provide inputs to downstream grid cells driven by a *phase interference model* (Blair, Gupta, & Zhang, 2008; Burgess, 2008; Burgess, Barry, & O'Keefe, 2007; Hasselmo, 2008). Several grid cells with different scales and field spacings converge to form a single place cell. Each place cell also provides downstream spiking input to a single reward cell proposed to represent prefrontal cortex mechanisms, i.e., place cells and reward cells have a bijective topology (Fig. 1).

A head direction cell is a neuron that significantly increases its firing rate when the rat's allocentric head orientation in the world horizontal plane, i.e., the head azimuth, approaches a specific angle which is referred to as its *preferred direction* (Sargolini et al., 2006; Taube, Muller, & Ranck, 1990). The head direction cell's preferred direction depends on the environmental cues and proprioceptive inputs. The head direction cells simulated in the HiLAM are cosine tuned and velocity modulated. Given the agent's instantaneous velocity vector $\mathbf{v}(t)$ and the preferred direction θ_i of a simulated head direction cell i its output d_i can be given as:

$$d_i(t) = \mathbf{v}(t) \cdot \begin{bmatrix} \cos(\theta_i) \\ \sin(\theta_i) \end{bmatrix} \quad (1)$$

A grid cell is a neuron type which increases its firing rate significantly when the animal traverses a regular array of periodic locations in the environment. The collection of locations where an individual grid cell fires, i.e., the grid cell's firing fields, forms a two dimensional periodic pattern with regular inter-field intervals and similar field areas. More specifically, the firing fields of a single grid cell tile the infinite two dimensional plane as the vertices of equilateral triangles. Extensive experimental data show the existence of grid cells with different inter-field spacing and field areas along the dorsal to ventral axis of the medial entorhinal cortex (Barry & Burgess, 2007; Hafting et al., 2005; Stensola et al., 2012). In a single rat, grid cells in the medial entorhinal cortex are organized in anatomically overlapping modules with distinct firing field orientation and discrete scales (Stensola et al., 2012). The simulated grid cells found in the HiLAM use a variant of the persistent spiking model (Hasselmo, 2008) which belongs to the class of phase interference models (Burgess, Barry, & O'Keefe, 2007) for grid cells. The spiking output of the j th grid cell g_j can be defined as:

$$\begin{aligned} \phi_{(ij)}(t) &= 2\pi \left(ft + b_j \int_0^t d_i(\tau) d\tau \right) \\ s_{(ij)}(t) &= H \left(\cos(\phi_{(ij)}(t) + \psi_{(ij)}) - s_{thr} \right) \\ g_j(t) &= \prod_{s \in S_j} s(t) \end{aligned} \quad (2)$$

where $\phi_{(ij)}$ is the phase of the persistent spiking cell modulated by the i th head direction cell, f is the frequency, b_j is the scaling factor for all persistent spiking cells projecting to the j th grid cell, $s_{(ij)}$ is the persistent spiking cell signal, ψ is the phase offset, s_{thr} is the action potential threshold, H is the Heaviside function satisfying $H(0) = 0$, and S_j is the set of persistent spiking

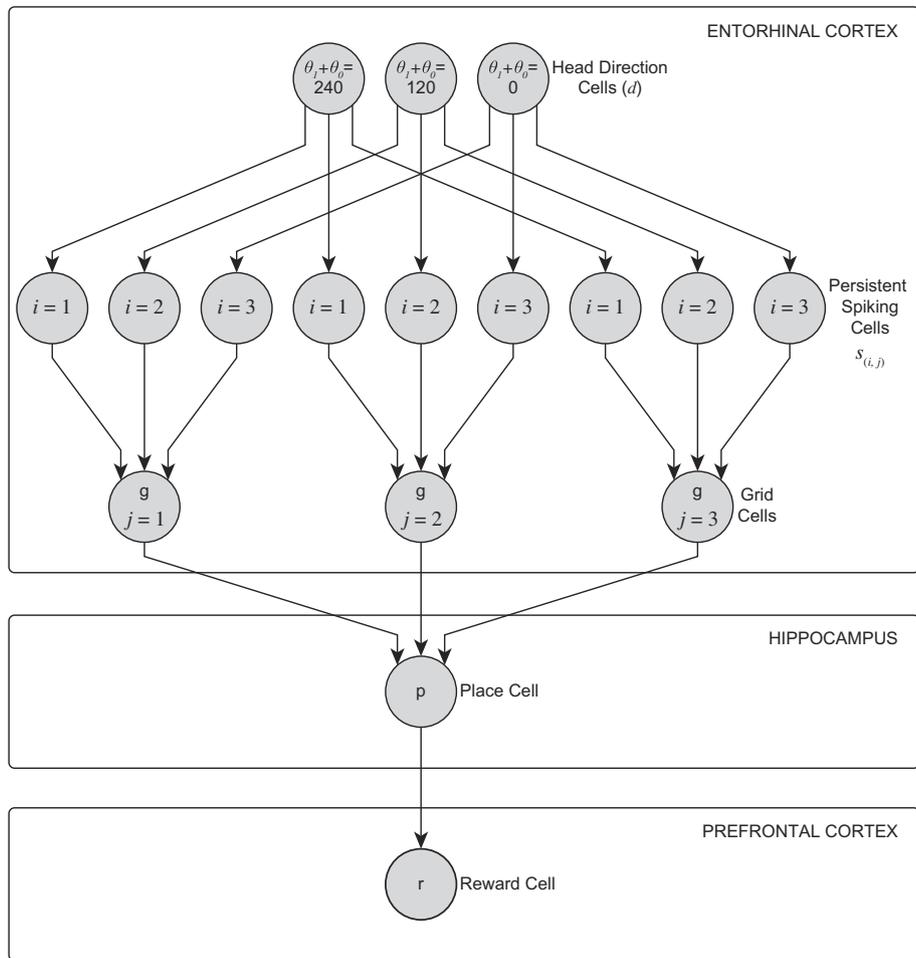


Fig. 1. The network topology showing how head direction cells in entorhinal cortex are proposed to drive persistent spiking cells that generate grid cells in entorhinal cortex. In this model, the grid cells drive the activity of place cells in the hippocampus. Input from place cells converges with reward representations to drive the activity of reward cells proposed to occur in the prefrontal cortex.

cells projecting to grid cell j . In summary, a grid cell is the conjunction of its immediate predecessors consisting of persistent spiking cells.

A place cell increases its firing rate when the animal crosses a compact region of the environment (O’Keefe and Nadel, 1978).

The firing field of a place cell is called its *place field*. Each place cell has mostly unique place fields making them good candidates to encode a rat’s spatial environment. In HiLAM a simulated place cell is the conjunction of all its inputs which are provided by a set of predecessor grid cells:

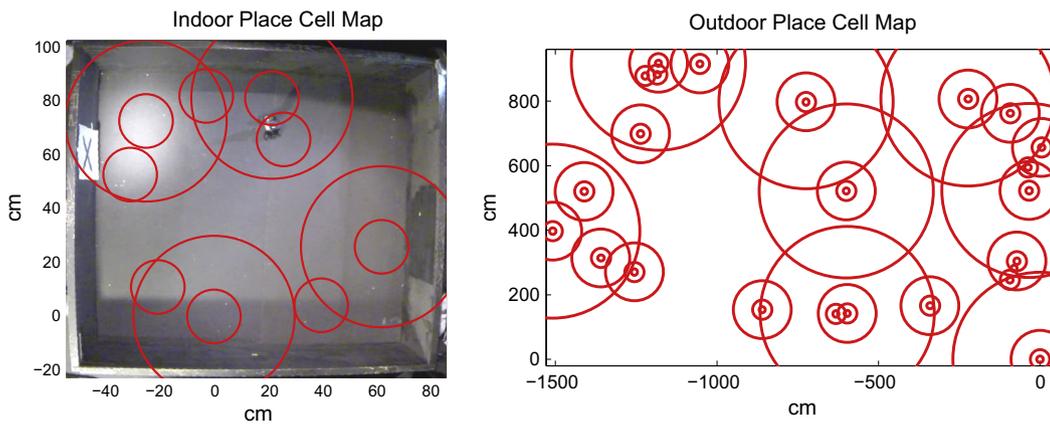


Fig. 2. Two examples of hierarchical place cell maps after the exploration phase. The red circles represent place cell firing fields. The coordinate system origin denotes the starting location of the exploration phase. (Left) Indoor 125 cm by 125 cm square maze overlaid by a hierarchical place cell map with 2 levels. The first level’s place field radii are $\rho_0 = 10$ cm and the second level’s relative scale factor is $a = 3$, i.e., place cell firing field radii of the second level are $3 \times 10 = 30$ cm. The object visible at the coordinates (20 cm, 70 cm) is the remote controlled vehicle. (Right) Outdoor arena one order of magnitude larger than the indoor maze overlaid by a hierarchical place cell map with 4 levels. This map is obtained by adding 2 more levels to the map of the indoor maze. Relative scaling factor between layers is $a = 3$. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

$$p_k(t) = \prod_{g \in G_k} g(t) \quad (3)$$

where G_k is the set of grid cells projecting to place cell k and p_k is the k th place cell.

A reward cell is a theoretical prefrontal cortex (PFC) cell representing whether a place cell is associated with a reward (or goal) or not. Each reward cell receives its input from a single unique place cell. The necessary and sufficient condition for a reward cell to fire is satisfied when (i) its associated place cell is a goal place cell and (ii) its associated place cell generates action potentials.

In the original look-ahead trajectory model (Erdem & Hasselmo, 2012) the agent initially generates a place cell map of its environment by exploring and recruiting place cells to represent salient locations, e.g., food sources, escape platforms, etc. In HiLAM the place cell map does not necessarily need to be dense, i.e., place fields might be non-overlapping and distant from each other. During the goal directed navigation phase the agent picks a previously visited location as its goal and marks all place cells with place fields containing the goal location as *goal place cells*. Successively, the agent generates look-ahead linear trajectory probes starting from its current location towards several samples of candidate orientations while stationary. If a look-ahead probe crosses a place field containing the chosen goal location, then it causes the respective reward cell to generate action potentials and is considered as the *winning probe*. The agent then moves towards the direction of the winning probe.

In the HiLAM (Erdem & Hasselmo, 2013) the agent represents its environment at different scales creating a *scale space* (Lindeberg, 1993; Sporring, Nielsen, Florack, & Johansen, 1997). The spatial resolution of each level decreases going from lower levels to higher ones. Equivalently, place field radii of place cells belonging to different levels decrease going from higher to lower levels.

During the exploration phase, the agent recruits place cells from each level to encode salient locations as long as no other previously recruited place cell's place field already contains that location (see Fig. 2). During the navigation phase, the agent generates multiple look-ahead linear trajectory probes with different bearings starting from its current location. This time, however, probes propagate at all levels of the hierarchy *simultaneously* but at different speeds proportional with that level's relative scale. For instance, if the scale of a level l is three times larger relative to the previous level $l - 1$ then the probe at level l will propagate *simultaneously* three times faster than the probe at level $l - 1$. The necessary condition guaranteeing that some probe at level l will always be able to cross a goal place field at level $l - 1$ can be given as follows (Erdem & Hasselmo, 2013):

$$\gamma_0 > 2\rho_0 a^{-1} \sqrt{1+a} \quad (4)$$

In Eq. (4), γ_0 is the probe range at the lowest level of the hierarchy, ρ_0 is the place field radius at the lowest level of the hierarchy, and a is the relative scale factor between consecutive levels of the hierarchy. This condition also guarantees that the agent will reach the goal location after a finite amount of probe scans (Erdem & Hasselmo, 2013).

The faster probe propagation at levels with larger scales allows probes to cover longer ranges. More importantly, since the probe propagation happens *simultaneously* at all levels of the hierarchy, we can extend the maximum probe length of the hierarchical model by simply adding levels on top of the hierarchy while keeping the total time allocated for a single probe constant. Assuming that the network noise accumulates faster during look-ahead linear trajectory probes due to the absence of sensory cues, and that the noise accumulation is directly proportional to the duration of a single probe, the place cell hierarchy allows the coverage of longer probe ranges guaranteeing noise accumulation levels limited from above. Theoretical details of the hierarchical

look-ahead linear trajectory model can be found in (Erdem & Hasselmo, 2013).

2.2. RatSLAM

RatSLAM is a state of the art robotic mapping and navigation system inspired at a high level by the neural processes underlying navigation in the rodent hippocampus and entorhinal cortex (Ball et al., 2013; Milford, Wiles, & Wyeth, 2010; Milford & Wyeth, 2009; Milford et al., 2004). Here we use three key RatSLAM components: visual self-motion estimation, visual place recognition, and map relaxation to form a stable map of the environment which is then used as the input into the HiLAM. In this implementation the sole sensory input is low resolution visual imagery sampled from the small remote-controlled vehicle (referred to here as the agent) as it moves around the environments. No other sensory modalities such as vehicle wheel encoders, inertial measurement units or a compass are used. In this implementation we omit the use of the *pose cell* (Milford, 2008) component of RatSLAM for two reasons: firstly, the spatial cells used in the HiLAM model provide a higher level of biological fidelity, and secondly because the filtering capability provided by the pose cells is not required in the experimental environments required here.

2.2.1. Visual self-motion estimation

To calculate rotational changes (rotational movement by the agent), mean absolute image intensity differences \mathbf{D} between the two most recent consecutive images are calculated over all horizontal (rotation) offsets:

$$\mathbf{D} = \underset{\Delta x \in \sigma}{\text{ming}}(\Delta x) \quad (5)$$

where σ is the complete range of relative horizontal image offsets (0–100 pixels), and $g(\cdot)$ is given by:

$$g(\Delta x, j) = \frac{1}{s} \sum_{x=0}^w \sum_{y=0}^w |p_{\text{mod}(x+\Delta x, w), y} - p_{x, y}^{\text{last}}| \quad (6)$$

where s is the area in pixels of the image, w is the image width in pixels and p is the pixel intensity value. The resultant difference profile \mathbf{D} is shown in Fig. 3c.

The horizontal pixel shift Δx_m corresponding to the minimum difference score is multiplied by a gain constant, ζ , to obtain a rotational velocity estimate, ω :

$$\omega = \zeta \Delta x_m \quad (7)$$

ζ is calculated by dividing the image horizontal field of view (360°) by the down sampled horizontal pixel resolution (100 pixels).

The minimum mean absolute image intensity difference D_m is multiplied by a gain constant, ν , to obtain an estimate of the agent's translational speed, s :

$$s = D_m \quad (8)$$

The gain constant is determined empirically for each environment using a short traverse of the agent over a known distance. While the simplicity of the method means there is no guarantee of consistent scale between different areas of an environment, extensive studies in robotics have shown that such an approach produces maps of an environment that are sufficiently metric to enable robot navigation (Ball et al., 2013; Milford, Schill, Corke, Mahony, & Wyeth, 2011; Milford & Wyeth, 2008).

2.2.2. Visual place recognition

Place recognition is performed by comparing the current camera image to all images (also referred to as *templates*) that have previously been learnt by the recognition system during prior exploration by the agent (Fig. 4). In a similar manner to the visual

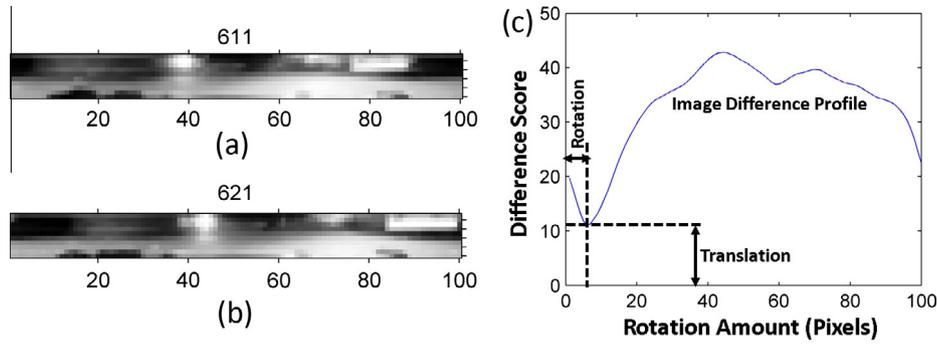


Fig. 3. (a) Visual self-motion estimation based on samples from the camera mounted on the agent. (a and b) Consecutive panoramic images are compared using a Sum of Absolute Differences over a complete range of relative rotation offsets to generate a (c) difference profile. The change in agent orientation is calculated using pixel shift corresponding to the minimum difference score, while agent translation is calculated using the minimum difference score.

self-motion estimation process, the mean absolute image intensity differences D_j between the current image and each learnt image j is calculated:

$$D_j = \frac{1}{s} \sum_{x=0}^s \sum_{y=0}^s |p_{x,y} - p_{x,y}^j| \quad (9)$$

where s is the area in pixels of the image and p is the pixel intensity value. If the minimum difference score over all previously learnt images is below a recognition threshold, the current image is matched to the corresponding learnt image, and a local view (LV) cell associated with the scene is activated. Otherwise the current camera image is learnt as a novel visual scene. More sophisticated computer vision methods for image matching can also be used for the visual place recognition phase of RatSLAM. However, the performance of the template matching algorithm we use based on the mean absolute difference between the new and the previously experienced view proved to be good enough for the experiments presented.

2.2.3. RatSLAM experience mapping

The experience mapping algorithm provides a mechanism for using vision-based place recognition to correct for the accumulation of self-motion errors over time in order to produce a stable and

locally metric map of space. An experience map contains representations of distinct places, called experiences, e , and links between experiences describing the transitions, t , between these places (Fig. 5). In this simplified RatSLAM implementation, each experience is defined by an active local view cell V^i . However, each experience is positioned at a location \mathbf{p}^i in *experience space*, which is similar to real world Cartesian space but with connectivity constraints. The complete state of an experience can be defined as the 2-tuple:

$$e_i = \{V^i, \mathbf{p}^i\} \quad (10)$$

The creation of a new experience is triggered by the visual place recognition algorithm learning a novel visual scene, while the re-activation of an existing experience is triggered by the recognition of a familiar visual scene. In either case, a transition link l_{ij} is learnt from the previously active experience e_i to the currently active experience e_j . These links encode the value of the change in position, $\Delta\mathbf{p}^{ij}$, computed directly from visual self-motion estimates:

$$l_{ij} = \{\Delta\mathbf{p}^{ij}\} \quad (11)$$

The visual self-motion information defines the *initial* location in the experience map space of a newly created experience:

$$e_j = \{V^j, \mathbf{p}^i + \Delta\mathbf{p}^{ij}\} \quad (12)$$

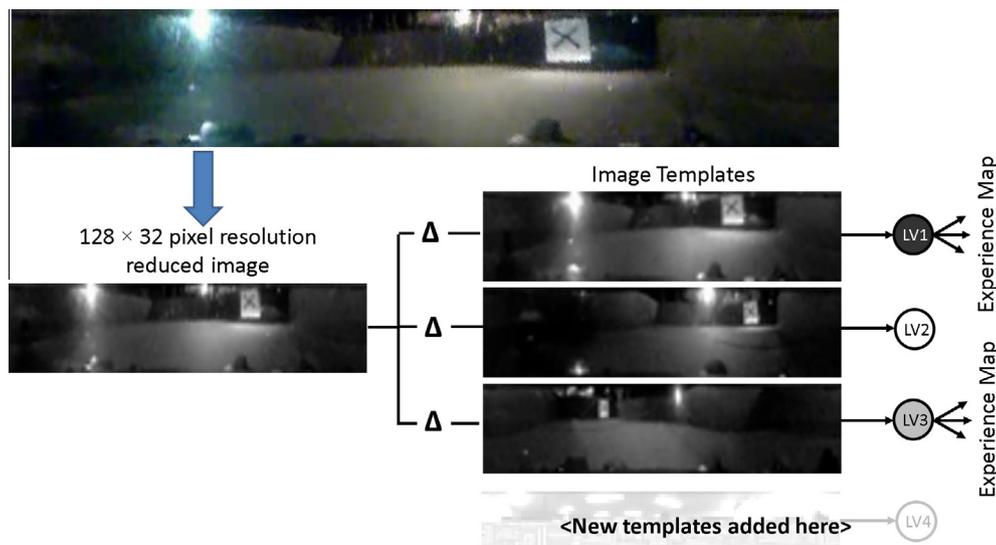


Fig. 4. Visual place recognition system. The current camera image from the agent is downsampled and compared to all existing image templates stored in an image database. If a strong image match is found, that image is used to perform a loop closure in the experience map. If no strong image matches are found, the current image is added to the image database.

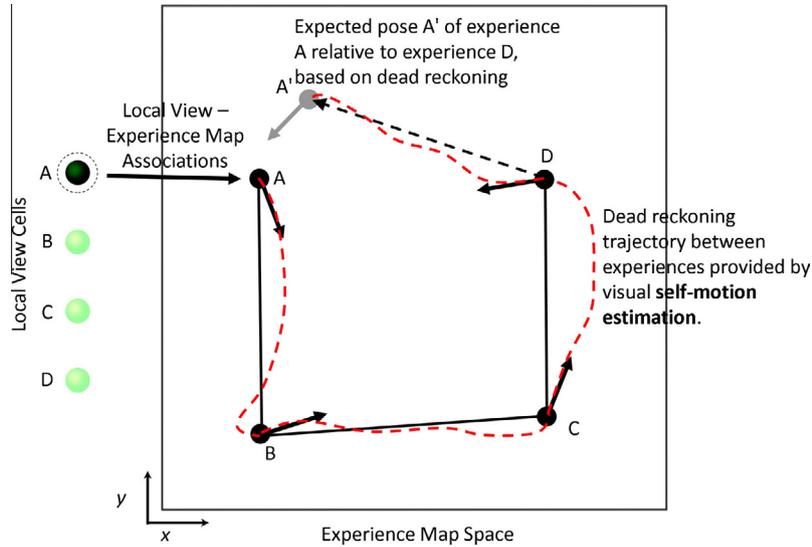


Fig. 5. Each local view cell represents a distinct visual scene in the environment, and becomes active when the camera sees that scene. Active local view cells and visual odometry drive the creation of experience nodes (filled black circles) in the experience map, a semi-metric graphical representation of places in the environment and their interconnectivity. Recognized visual scenes can drive re-activation of previously learnt nodes in the map in a process known as “loop closure”. Black lines code the relative motion information between experiences and are a net sum of motion (red dashed lines) between experiences. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

Initially, the spatial relationships between linked experiences exactly match the spatial information provided by the self-motion estimates. However, when the first familiar visual scene is recognized, a process of loop closure occurs, where the experience node associated with that scene is reactivated, rather than learning a new experience. Unless the self-motion estimates are perfect, a discrepancy between the relative locations of the two most recently activated experiences is introduced. To distribute this error throughout the map, a process of graph relaxation is performed, which minimizes the discrepancy between inter-experience self-motion estimates and relative location in experience map space. The process involves changing the location of each experience by $\Delta \mathbf{p}^i$:

$$\Delta \mathbf{p}^i = \alpha \left[\sum_{j=1}^{N_f} (\mathbf{p}^j - \mathbf{p}^i - \Delta \mathbf{p}^{ij}) + \sum_{k=1}^{N_r} (\mathbf{p}^k - \mathbf{p}^i - \Delta \mathbf{p}^{ki}) \right] \quad (13)$$

where α is a correction rate constant (0.5), N_f is the number of links from experience e_i to other experiences, and N_r is the number of links from other experiences to experience e_i . Eq. (13) is applied to all experiences a set number of times per second of system operation (15 Hz for the experiments described here) as the map is formed, and results in the experience map arranging itself so as to average out self-motion errors throughout the map, maximizing the local metric topology of any area of the map.

3. Experimental procedure

Experiments were performed in two distinct environments, a small square rat arena and an outdoor area more than two orders of magnitude larger. The larger area enabled us to test the scalability of the HiLAM.

3.1. Agent

To gather data in the two testing environments, we used a high speed, miniature remote control vehicle (referred to as the agent) equipped with a panoramic imaging setup, as shown in Fig. 6. The agent was equipped with a Kogeto dot panoramic combined

lens-mirror paired with a 720p miniature camera (808 #16 Micro Key Chain Camera). Raw camera images were cropped and unwrapped to create 480×80 pixel panoramic images with a field of view corresponding to approximately 360° horizontally by 60° vertically. These unwrapped images were gathered at 30 Hz as the agent moved through the environment but down sampled to an effective rate of 3 Hz before being input to the visual odometry and visual place recognition algorithms. Downsampling achieved two main benefits: perhaps counterintuitively, small amounts of motion that are not detectable in consecutive low resolution images at 30 Hz can be more reliably detected at 3 Hz, since the apparent visual change is greater; and computation time is reduced by an order of magnitude. The original RatSLAM system has been demonstrated running at real-time speed in very large environments (6 km²) (Milford & Wyeth, 2008); we would hope to replicate this scalability in future work combining the RatSLAM and HiLAM models.

3.2. Indoor square arena

The first experiment was conducted in a 125 cm \times 125 cm indoor square arena normally used in rodent experiments in the Hasselmo Lab (Fig. 7). A downwards facing GoPro camera mounted directly over the center of the arena captured video which was processed to yield a ground truth trajectory for the agent. The arena wall had one high contrast black cue and several external lights to provide sufficient illumination for the small camera sensor. The agent was tele-operated around the arena for 3 min 51 s over a distance of approximately 2500 cm at an average speed of 11 cm/s. The agent’s path during the experiment was determined by the human operator to cover the arena with slight emphasis given to laps following the perimeter.

The Hierarchical Look-Ahead Trajectory Model used for the square arena consisted of two levels. The first level, providing the highest resolution, contained place cells with place fields having 10 cm radii. The second level of the hierarchy had a scaling factor of 3 relative to the first level. Consequently, the second level’s place field radii were $10 \times 3 = 30$ cm. The look-ahead probe range of the first level was 70 cm. The second level’s probe range was

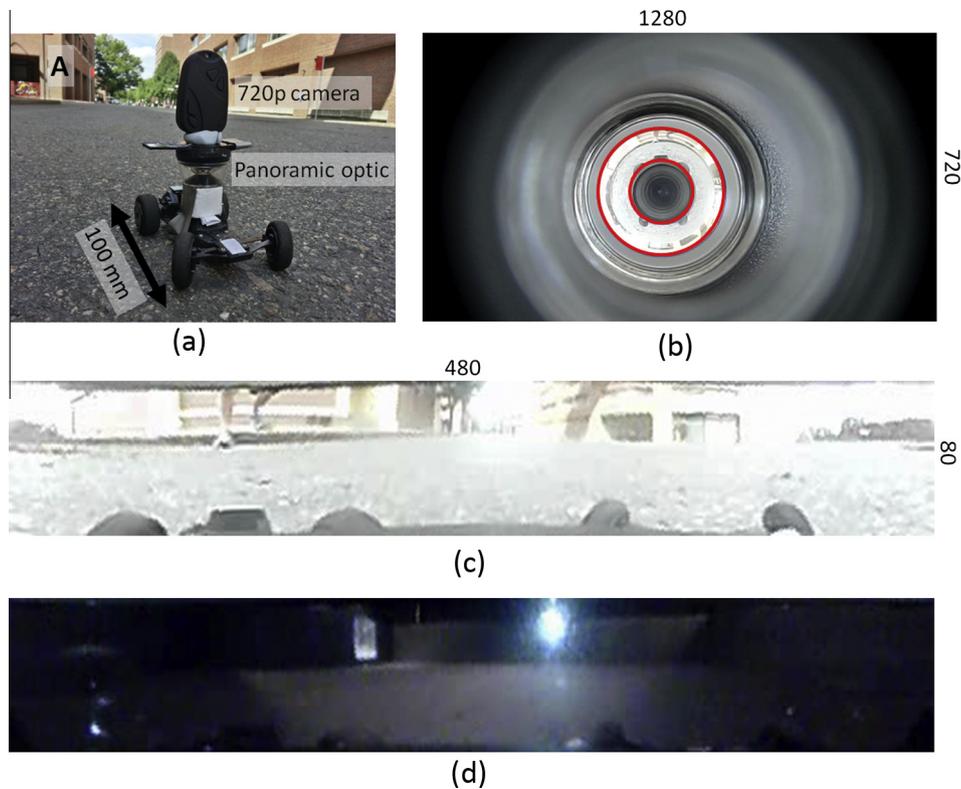


Fig. 6. (a) Experimental agent with miniature panoramic imaging system shown in the outdoor environment. (b) 1280×720 pixel raw camera images were cropped and unwrapped (red circles show boundaries of unwrapped area) to form (c and d) 480×80 pixel panoramic images corresponding to an approximate field of view of $360^\circ \times 60^\circ$. (c) Unwrapped road image. (d) Unwrapped square arena image. The “A” in part (a) indicates the vertical corner of a distal building in the outdoor environment which we use as an orientating landmark (for the reader) throughout the paper. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

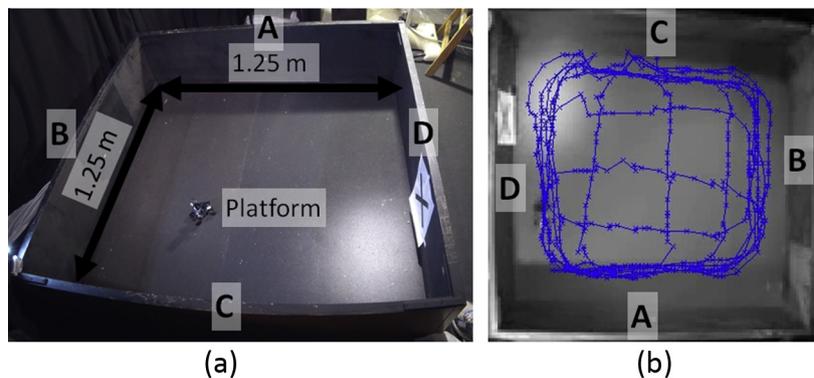


Fig. 7. (a) Square $125 \text{ cm} \times 125 \text{ cm}$ experimental arena. Video from a downward facing GoPro camera mounted directly over the center of the arena was processed to extract a ground truth trajectory for the agent. The four arena walls are labeled A to D for viewer orientation (see Fig. 14).

$70 \times 3 = 210 \text{ cm}$. Each full look-ahead scan consisted of probes spanning the egocentric bearing range between -90° and 90° with 10° increments in a clockwise frame. The 0° bearing is the agent's forward facing heading angle. The speed and duration of a single look-ahead probe at the lowest level of the hierarchy were set to 70 cm/s and 1 s respectively. The speed of probes at a higher level is scaled by that level's relative scale. For instance, the speed of a probe propagating at the third level is $70 \times 3^2 = 630 \text{ cm/s}$ given that the relative scaling factor of that level is three. The full set of parameter values is given in Table 1.

3.3. Outdoor road arena

The second experiment was conducted in a $1700 \times 1000 \text{ cm}$ road area. A GoPro camera mounted to overlook the area at an

angle (not directly above due to practical considerations) captured video which was processed to yield a ground truth trajectory for the agent, which was then converted into standard (x, y) co-ordinates using a homography transformation. The environment was somewhat dynamic with lighting changes and pedestrian traffic through the area. The agent was tele-operated around the arena for $7 \text{ min } 58 \text{ s}$ over a distance of approximately $46,800 \text{ cm}$ at an average speed of 98 cm/s . The agent's path during the experiment was determined by the human operator to cover the arena with slight emphasis given to laps following the perimeter.

The Hierarchical Look-Ahead Trajectory Model used for the outdoor arena consisted of four levels. The model was obtained by adding two more levels on top of the model used to encode the small indoor arena. Hence, relative scaling between two consecutive levels remained at 3. The addition of two levels allowed the

Table 1
HiLAM parameters for the experiments.

	Probe bearing range	Probe bearing increment (°)	Probe speed (cm/s)	Probe duration (s)	Level count	ρ_0 (cm)	α
Indoor	$[-90^\circ, 90^\circ]$	0	70	1	2	10	3
Outdoor	$[-90^\circ, 90^\circ]$	10	70	1	4	10	3

model to encode the outdoor arena at coarser resolutions equivalently extending the maximum probe range to $70 \times 3^3 = 1890$ cm while keeping the single probe time at 1 s as in the indoor maze. The full set of parameter values is given in Table 1.

4. Results

In this section we present results from the vision-based self-motion estimation and place recognition processes, map formation and navigation probes in the two environments. The main interaction between the two processes of the hybrid model is as follows. The RatSLAM process computes agent's self-motion estimates in terms of odometry data based on visual cues. The odometry data then are input to the HiLAM process in the form of the corrected internal representation of the velocity vectors at each sampled position on the trajectory. The HiLAM process then uses this corrected trajectory to generate the hierarchical place field map encoding the explored area. Consecutively, HiLAM uses the hierarchical place field map to compute navigation paths towards goal locations.

In the hybrid model's implementation presented in this work the interaction between the two models, i.e., RatSLAM and HiLAM, is based on a serialized batch processing approach. In other words, the collection of the odometry data by the RatSLAM model using visual cues and the generation of the hierarchical place field map using the odometry by the HiLAM model is almost mutually exclusive. First the odometry information is collected and computed then the hierarchical place field map is generated and used for goal directed navigation. We talk about this approach more in the discussion section.

4.1. Self-motion estimation, place recognition and stable map formation

The maps of the agent's trajectory through the environment for the indoor and outdoor environments are shown in Figs. 9 and 10. Using just visual self-motion estimates (visual odometry), the estimate of the agent's location drifts rapidly over time. However, when using visual place recognition to perform loop closures, the system is able to form a stable and representative map of the agent's trajectory through the environment (compare with the

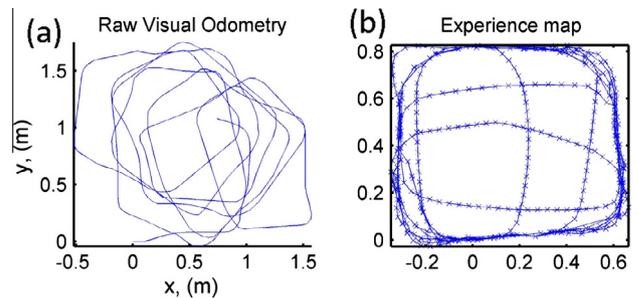


Fig. 9. Loop closure and map stabilization for the agent's trajectory in the indoor square arena. (a) Tracking agent location using raw visual self-motion estimates alone leads to significant dead reckoning drift, but with the addition of vision-induced loop closures, a stable map is achieved in (b).

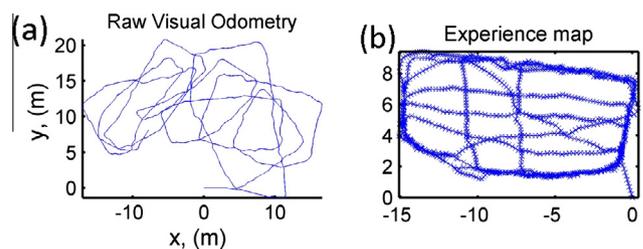


Fig. 10. Loop closure and map stabilization for the agent's trajectory in the larger outdoor road arena. (a) The drift is somewhat worse than for the indoor environment, but the addition of vision-induced loop closures results in a stable map in (b).

ground truth plots in Figs. 7b and 8b). This stable mapping of space is used as the input to the HiLAM.

4.2. Navigation probes

In this section we present the results of goal-directed navigation computed using the HiLAM model. In order to show the importance of the loop closure used by the RatSLAM process we show the results of goal-directed navigation using the same start and goal locations on the place cell map generated using the raw

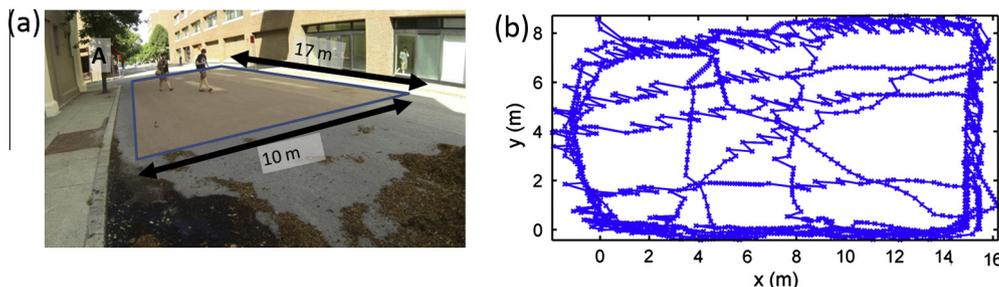


Fig. 8. (a) Outdoor 1700 cm \times 1000 cm road arena. Video from a GoPro camera mounted at the side of the arena (sample frame shown) was processed to extract a (b) ground truth trajectory for the vehicle, which was then converted into (x, y) co-ordinates using a homogenous transform calculated using the measured arena dimensions. Due to the size, illumination changes and dynamic nature of the environment the ground truth is less accurate than for the indoor environment, but is still useful for evaluating the topological correctness of the generated maps. The "A" in part (a) indicates the vertical corner of a distal building which we use as an orientating landmark (for the reader) throughout the paper.

(no loop closure) odometry estimation and the place cell map generated using the corrected (with loop closure) odometry estimation. We compare both raw and corrected navigation results to emphasize the crucial role of loop closure correction of raw odometry estimations. For all experiments the start and goal locations were chosen in a pseudo-random fashion to reasonably span the arenas.

In these figures, we show both the scan locations and probes and highlight the winning probes from each scan location. To aid in visualization, we also plot the nearest views to each scan location previously learnt by the system. In Section 5 we discuss multiple strategies a robot could use to follow the successful probe scans to a goal location, as in past RatSLAM work (Milford & Wyeth, 2009).

4.2.1. Without a stable map

To show the importance of achieving a stable, approximately metric map, we present navigation probes in the indoor arena calculated using a map produced without place recognition enabled, and hence no loop closure. For a start and goal location seemingly at opposite sides of the arena (Fig. 11a), two sets of scans are required to plan a path to the goal, including a first scan with probe length 210 cm (contrast with the scan length of 70 cm required to plan a path across the arena in Fig. 13). However, examination of the start and goal locations using the ground truth plot reveal that they are actually only about 20 cm apart, demonstrating that probe-based navigation is critically dependent upon the quality of the underlying spatial map. The camera views for the two scan locations are shown in Fig. 12 and confirm that the locations are close to each other and that lengthy probes are unnecessary.

4.2.2. With a stable map

A sample navigation probe for the indoor arena is shown in Fig. 13. Two sets of 70 cm long scans are required to scan to the goal location. The corresponding views from the agent at each successful scan location and orientation are shown in Fig. 14. The relative orientation of the goal location is indicated by a vertical red line (relative to the successful scan orientation), with the bottom of that line approximately corresponding to the goal location in the environment. The first successful scan heads across the center of the arena from the starting location near wall A to a location near wall C. The second successful scan turns right by

approximately 45° and intersects with the place field of the goal location. The ground truth plot (Fig. 13b) shows that both the start/goal locations and the scan locations are similarly located in both the RatSLAM map and in the ground truth plot.

A sample navigation probe for the outdoor environment is shown in Figs. 15 and 16. The probe consists of four sets of scans with the following successful scan lengths at each step: 1890 cm, 630 cm, 70 cm, 70 cm. The probe steps across the arena from right to left, homing in on the goal location. Fig. 17 shows the camera views corresponding to the four probe locations and orientations. The successful scans at each scan step shift in orientation as the overall probe homes in on the goal location. Once again the close correspondence between navigation probe and start/goal locations in the RatSLAM map and the ground truth plot can be seen in Fig. 15b.

5. Discussion

State-of-the-art goal-directed robotic navigation systems perform extremely well for limited durations and within relatively static environments. Higher level living organisms however appear not to suffer from the degrading effects of persistent navigation for extended periods of time and in dynamic environments. The technical challenge is bridging the spatial representation that autonomous systems use and the spatial representation created by grid cells in the entorhinal cortex and place cells in the hippocampus. Grid cells show stable firing over long time periods (10 min) even in darkness (Hafting et al., 2005), indicating robust path integration despite the noise inherent in neural systems, achieving an outcome that is challenging for state-of-the-art robotic navigation systems. If the robust biological mechanisms of grid cells could be implemented in robots they would provide a dramatic advance over current robot capabilities.

In this work we have demonstrated how two seemingly different biologically inspired navigation and mapping models can be put to work together, complementing each other in areas that are not their strongest suits. The RatSLAM (Milford et al., 2004) system has been shown to perform well generating encoded representations of space via visual information for extended periods of time and for both small and large environments (Milford &

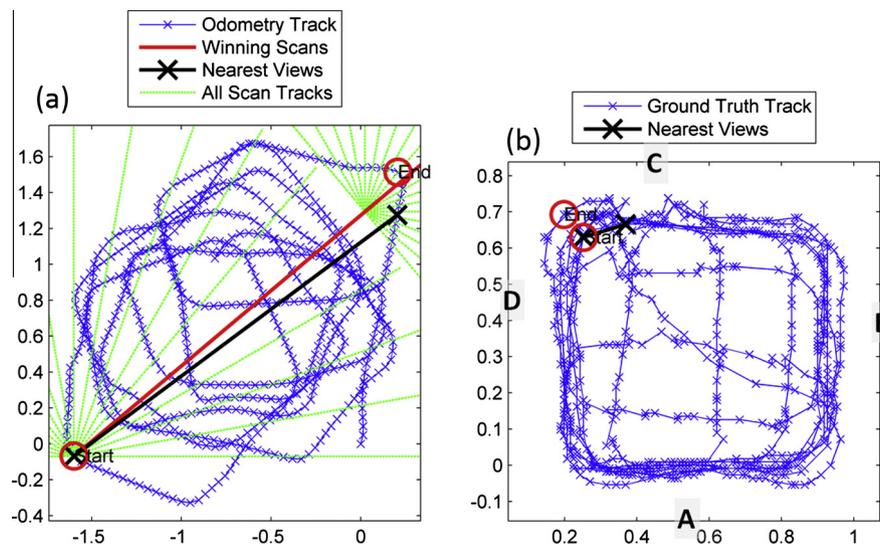


Fig. 11. (a) Navigation probes across the indoor arena using an unstable, uncorrected map representation. For the labeled start and goal locations, two scans are required to plan a seemingly lengthy path across the environment. However, the corresponding locations in the ground truth plot (b) reveal that the start and goal locations are actually almost on top of each other. Even with reasonably accurate self-motion information, the inevitable accumulation of errors over time mean that place recognition and loop closure is required to create a map representation that can be used for navigation.

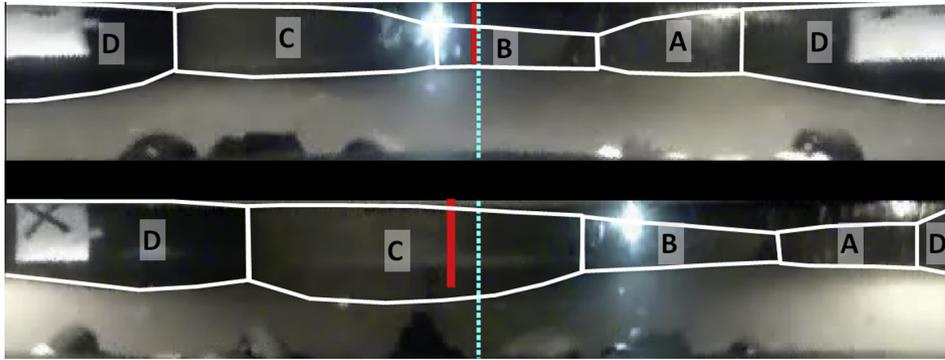


Fig. 12. Panoramic camera views at each of the scan locations using an uncorrected map (black crosses in Fig. 11a). From examining the position and size of the cross landmark in the first person views, it can be seen that the actual physical distance between the two probe locations is quite small, unlike the large distance encoded in the probes shown in Fig. 11a.

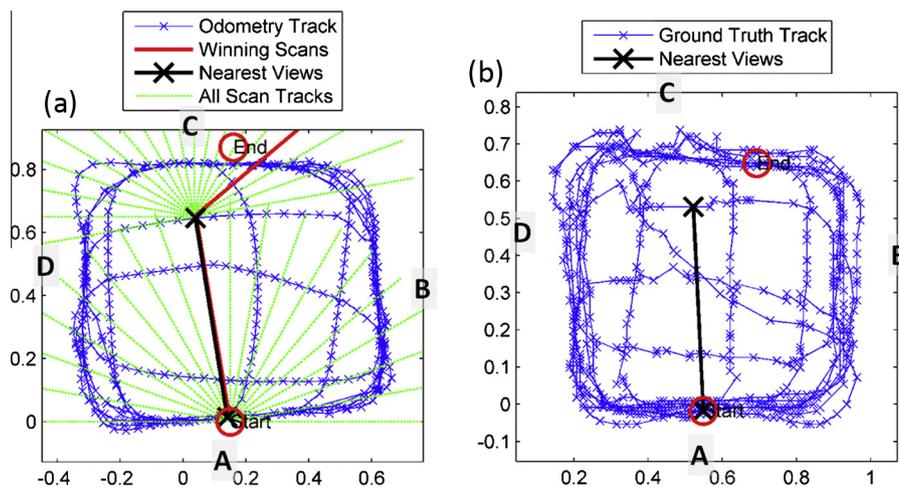


Fig. 13. (a) Navigation probe across the indoor arena. Green dashed lines indicate the scans at each step, with the red line overlaying the winning scan. Note that the first scan wins based on the second level place field with larger size surrounding the goal (End), and the second scan wins based on the first level place field with smaller size surrounding the goal (End). The black line shows the corresponding nearest recallable view stored in the map representation of the environment, with each black cross corresponding to the first person views shown in Fig. 14. Red circles indicate start and goal (end) locations. (b) Ground truth equivalent extracted using overhead tracking. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

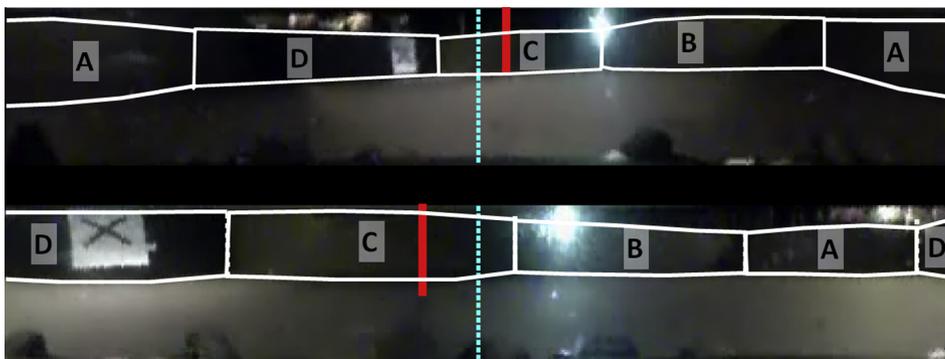


Fig. 14. Panoramic camera views at each of the probe locations (black crosses in Fig. 13a). Each view is the closest learnt view to each scan location. The center of each image corresponds to the planned forward direction of movement of the agent along the winning scan line. The projected goal location is shown by a thick red vertical line, with the bottom of that line approximately corresponding to the goal location in the environment. Wall outlines are indicated by white lines. The probe starts when the agent is adjacent to wall A but facing directly away from it (the wall is visible on either side due to the panoramic view). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

Wyeth, 2008). However, RatSLAM's goal navigation system is less biologically plausible than its mapping system and has no current mechanism for efficient scaling to larger environments. On the other hand, HiLAM (Erdem & Hasselmo, 2013) provides a

framework to encode space in a hierarchical grid cell and place cell topology and to provide guidance towards a preselected goal in the environment. HiLAM is highly scalable while theoretically guaranteeing success in finite steps and providing upper limits for noise

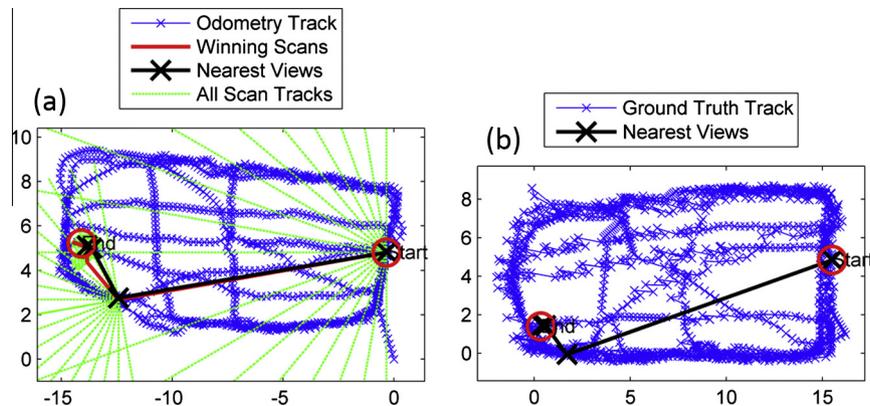


Fig. 15. (a) Sample probes shown in green are projected across the agent trajectory shown in blue that was generated by both visual odometry and vision-based loop closure. Red circles indicate start and goal locations. Red lines indicate the winning scans, and black lines indicate the nearest view path. The distal building cue marked by “A” in Figs. 6, 8 and 17 is located approximately 2000 cm s to the left of the bottom edge of the graph shown in (a). (b) Ground truth equivalent extracted using overhead tracking. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

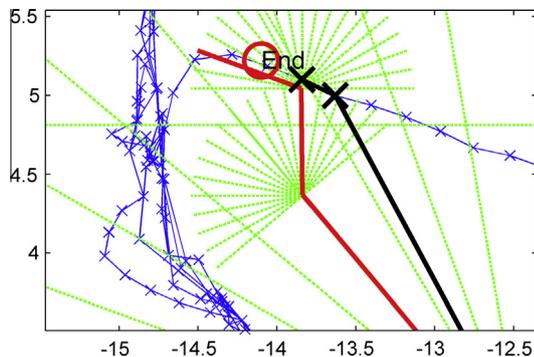


Fig. 16. Zoom of the final two scans in the outdoor probe, showing the significant change of absolute direction for the second last scan, which is mirrored in the camera views from the agent shown in Fig. 17. Black crosses show the nearest learnt views to each of the scan locations.

accumulation levels provided certain assumptions hold. However, HiLAM is a path integration system; it strongly relies on velocity data as its only input and is therefore highly susceptible to inaccuracies in the velocity parameter. In this work, we have demonstrated that a unified system combining RatSLAM and HiLAM can perform better than its individual parts alone even in the case where the visual input is from a noisy, real life system.

The current unified system does not contain a feedback loop between the two models, i.e., the generation of visual odometry and velocity data (RatSLAM) is independent of the goal directed navigation component (HiLAM). This architecture does seem consistent with the training and test paradigm that most physiological and behavioral tasks rely on. For instance, in Morris water-maze experiments (Morris et al., 1982; Steele & Morris, 1999) the rats presumably learn their task environment during the training trials, which might correspond to the exploration phase in our system involving visual odometry generation. Consecutively, in Morris water-maze experiments control rats perform the task almost flawlessly during test trials which might correspond to the goal directed navigation phase in our experiments. However, the absence of a feed-back loop between the two models in our framework might also be preventing potential improvements in performance. We are developing methods to improve the cooperation between RatSLAM and HiLAM by adding a feedback loop.

The hybrid model implementation presented in this work involving two previous models, i.e., RatSLAM and HiLAM, is based on a serialized batch processing approach. More specifically, the interplay between the two models is mutually exclusive. HiLAM

generates the hierarchical place field map only after RatSLAM collects and processes the odometry data. This approach is definitely not the only possible one. However, it is a reasonable method to show the feasibility of interaction between two biologically inspired goal directed navigation models published previously. Note that if the RatSLAM did not use the loop closure approach to correct for stochastic drift in the odometry estimation, using a real-time interaction between the two models, where the RatSLAM's estimated odometry data is fed as soon as it is calculated to the HiLAM, would be the preferred way. However, use of loop closure requires periodic processing of past odometry data to correct the location representations in the previously coded odometry data. We are currently looking for new improved methods that enable the corrective propagation triggered by loop closure detection in the space of place field maps. Such a method would allow us to switch from the current batch processing implementation to a real-time continuous interaction between the RatSLAM and the HiLAM.

The current version of HiLAM (Erdem & Hasselmo, 2013) is capable of delivering “in a beeline” global directions towards the selected goal location and does not explicitly take into account potential obstacles in the environment, though an earlier version did model obstacles (Erdem & Hasselmo, 2012). In more complex environments, the introduction of a local motion planner (Milford & Wyeth, 2009) would enable the system to balance global navigation instructions provided by HiLAM with local considerations such as navigating around static or dynamic obstacles. Ultimately, the introduction of concepts such as barriers into the HiLAM model would enable the system to appropriately reward or penalize navigation probes based on known accessible or no-go areas of the environment. Performing active robot navigation using both the paths planned by HiLAM's probes and the continuous localization capability provided by RatSLAM should be feasible based on previous successful active robot navigation experiments using RatSLAM (Milford & Wyeth, 2009). The simplest and most direct method would be to instruct the robot to follow the winning probes provided by HiLAM (the red lines in Fig. 13), using RatSLAM to provide both dead-reckoning and also place-recognition when crossing previously visited locations.

An open question in the biological representation of space is the trigger to associate hippocampal cells to certain spatial locations. There is compelling evidence that the association trigger might not only depend on spatial cues but on context as well (Komorowski, Manns, & Eichenbaum, 2009). Further understanding of how brain prioritizes contextual and spatial associations could have significant impact on selection of sensory cues to encode locations and their organization in a persistent database

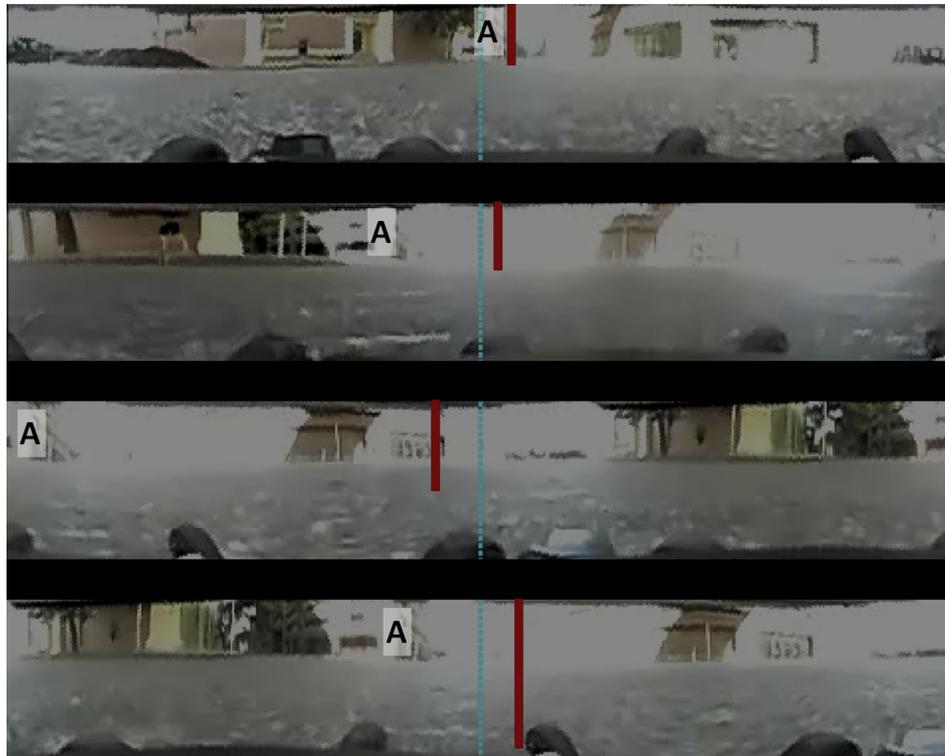


Fig. 17. Panoramic camera views for the four scan locations shown in Fig. 15a (black crosses). The center of each image corresponds to the forward direction of movement along the probe. Images have been artificially darkened by 40% only for presentation purposes. The projected goal location is shown by a thick red vertical line, with the bottom of that line approximately corresponding to the goal location in the environment. The “A” indicates the vertical corner of a distal building which we use as an orientating landmark (for the reader) throughout the paper (see Figs. 6 and 8). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

in robotic SLAM systems. Another interesting biological phenomenon not very well understood so far is the *remapping* of the place cells (Bostock, Muller, & Kubie, 1991; Jeffery, 2011; Markus et al., 1995; Muller & Kubie, 1987). It is not yet very clear why or how the remapping happens. More insight into this phenomenon might result in more efficient encoding of space in robotic navigation.

The HiLAM tries to find the best direction towards the goal location from the agent’s current location by generating hypotheses about possible future trajectories in the spatial coordinate system and picking the one that signals high probability of arrival to the goal location depending on previous experience. A variation of HiLAM might perform the hypothesis search in visual experience space instead of (or in collaboration with) the spatial space. RatSLAM’s spatio-visual experience map would be an excellent candidate search space. In this case the query for the goal directed navigation would be a view of the goal location instead of an abstraction of the goal. Furthermore, once the goal location enters the visual range of the robot during goal directed navigation, the local visual navigation may easily take over to guide the robot towards its intended destination (Milford & Wyeth, 2009). We are currently working on expanding our unified model to accommodate goal queries in visual experience space. Together, we think the combined HiLAM and RatSLAM models provide a unique method for exploring, in a biologically relevant but functionally grounded manner, how animals and robots might best make navigation decisions based on their sensory-spatial representations of the world.

Acknowledgments

This work was supported by the ONR MURI Grant N00014-10-1-0936, the ONR Grant N00014-09-1-0641 and an Australian Research Council Discovery Project DP1212775.

References

- Ball, D., Heath, S., Wiles, J., Wyeth, G., Corke, P., & Milford, M. (2013). OpenRatSLAM: An open source brain-based SLAM system. *Autonomous Robots*, 1–28.
- Barry, C., & Burgess, N. (2007). Learning in a geometric model of place cell firing. *Hippocampus*, 17(9), 786–800. <http://dx.doi.org/10.1002/hipo.20324>.
- Blair, H. T., Gupta, K., & Zhang, K. (2008). Conversion of a phase- to a rate-coded position signal by a three-stage model of theta cells, grid cells, and place cells. *Hippocampus*, 18(12), 1239–1255. <http://dx.doi.org/10.1002/hipo.20509>.
- Bostock, E., Muller, R. U., & Kubie, J. L. (1991). Experience-dependent modifications of hippocampal place cell firing. *Hippocampus*, 1(2), 193–205. <http://dx.doi.org/10.1002/hipo.450010207>.
- Brown, M. (2011). Social influences on rat spatial choice. *Comparative Cognition and Behavior Reviews*, 6, 5–23. <http://dx.doi.org/10.3819/ccbr.2011.60003>.
- Burgess, N. (2008). Grid cells and theta as oscillatory interference: Theory and predictions. *Hippocampus*, 18(12), 1157–1174. <http://dx.doi.org/10.1002/hipo.20518>.
- Burgess, N., Barry, C., & O’Keefe, J. (2007). An oscillatory interference model of grid cell firing. *Hippocampus*, 17(9), 801–812. <http://dx.doi.org/10.1002/hipo.20327>.
- Davidson, T. J., Kloosterman, F., & Wilson, M. A. (2009). Hippocampal replay of extended experience. *Neuron*, 63(4), 497–507. <http://dx.doi.org/10.1016/j.neuron.2009.07.027>.
- Erdem, U. M., & Hasselmo, M. E. (2012). A goal-directed spatial navigation model using forward trajectory planning based on grid cells. *The European Journal of Neuroscience*, 35(6), 916–931. <http://dx.doi.org/10.1111/j.1460-9568.2012.08015.x>.
- Erdem, U. M., & Hasselmo, M. E. (2013). A biologically inspired hierarchical goal directed navigation model. *Journal of Physiology, Paris*. <http://dx.doi.org/10.1016/j.jphysparis.2013.07.002>.
- Foster, D. J., & Wilson, M. A. (2006). Reverse replay of behavioural sequences in hippocampal place cells during the awake state. *Nature*, 440(7084), 680–683. <http://dx.doi.org/10.1038/nature04587>.
- Hafting, T., Fyhn, M., Molden, S., Moser, M.-B., & Moser, E. I. (2005). Microstructure of a spatial map in the entorhinal cortex. *Nature*, 436(7052), 801–806. <http://dx.doi.org/10.1038/nature03721>.
- Hasselmo, M. E. (2008). Grid cell mechanisms and function: Contributions of entorhinal persistent spiking and phase resetting. *Hippocampus*, 18(12), 1213–1229. <http://dx.doi.org/10.1002/hipo.20512>.
- Jacobs, L., & Liman, E. (1991). Grey squirrels remember the locations of buried nuts. *Animal Behaviour*, 41, 103–110. <http://www.sciencedirect.com/science/article/pii/S0003347205805068>.

- Jadhav, S. P., Kemere, C., German, P. W., & Frank, L. M. (2012). Awake hippocampal sharp-wave ripples support spatial memory. *Science (New York, N.Y.)*, 336(6087), 1454–1458. <http://dx.doi.org/10.1126/science.1217230>.
- Jeffery, K. J. (2011). Place cells, grid cells, attractors, and remapping. *Neural Plasticity*, 2011, 182602. <http://dx.doi.org/10.1155/2011/182602>.
- Johnson, A., & Redish, A. D. (2007). Neural ensembles in CA3 transiently encode paths forward of the animal at a decision point. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, 27(45), 12176–12189. <http://dx.doi.org/10.1523/JNEUROSCI.3761-07.2007>.
- Komorowski, R. W., Manns, J. R., & Eichenbaum, H. (2009). Robust conjunctive item-place coding by hippocampal neurons parallels learning what happens where. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, 29(31), 9918–9929. <http://dx.doi.org/10.1523/JNEUROSCI.1378-09.2009>.
- Lindeberg, T. (1993). *Scale-space theory in computer vision* (pp. 444). Springer. <http://www.amazon.com/Scale-Space-Computer-Springer-International-Engineering/dp/0792394186>.
- Louie, K., & Wilson, M. A. (2001). Temporally structured replay of awake hippocampal ensemble activity during rapid eye movement sleep. *Neuron*, 29(1), 145–156. [http://dx.doi.org/10.1016/S0896-6273\(01\)00186-6](http://dx.doi.org/10.1016/S0896-6273(01)00186-6).
- Markus, E., Qin, Y., Leonard, B., Skaggs, W., McNaughton, B., & Barnes, C. (1995). Interactions between location and task affect the spatial and directional firing of hippocampal neurons. *The Journal of Neuroscience*, 15(11), 7079–7094. <http://www.jneurosci.org/content/15/11/7079.abstract>.
- Milford, M. J. (2008). *Robot navigation from nature: Simultaneous localisation, mapping, and path planning based on hippocampal models* (1st ed., Vol. 41, pp. 196). Berlin, Heidelberg: Springer Verlag. <http://www.amazon.com/dp/3540775196>.
- Milford, M. J., Wyeth, G. F., & Prasser, D. (2004). RatSLAM: A hippocampal model for simultaneous localization and mapping. In *Proceedings. ICRA'04. 2004 IEEE international conference on robotics and automation, 2004* (Vol. 1, pp. 403–408). IEEE. <http://dx.doi.org/10.1109/ROBOT.2004.1307183>.
- Milford, M. J., Schill, F., Corke, P., Mahony, R., & Wyeth, G. (2011). Aerial SLAM with a single camera using visual expectation. *IEEE International Conference on Robotics and Automation*, 2011, 2506–2512. <http://dx.doi.org/10.1109/ICRA.2011.5980329>.
- Milford, M., Wiles, J., & Wyeth, G. (2010). Solving navigational uncertainty using grid cells on robots. *PLoS Computational Biology*, 6(11).
- Milford, M., & Wyeth, G. (2008). Mapping a suburb with a single camera using a biologically inspired SLAM system. *IEEE Transactions on Robotics*, 24(5), 1038–1053.
- Milford, M., & Wyeth, G. (2009). Persistent navigation and mapping using a biologically inspired SLAM system. *The International Journal of Robotics Research*, 29(9), 1131–1153. <http://dx.doi.org/10.1177/0278364909340592>.
- Morris, R. G. M., Garrud, P., Rawlins, J. N. P., & O'Keefe, J. (1982). Place navigation impaired in rats with hippocampal lesions. *Nature*, 297(5868), 681–683. <http://dx.doi.org/10.1038/297681a0>.
- Muller, R., & Kubie, J. (1987). The effects of changes in the environment on the spatial firing of hippocampal complex-spike cells. *The Journal of Neuroscience*, 7(7), 1951–1968. <http://www.jneurosci.org/content/7/7/1951.abstract>.
- O'Keefe, J., Burgess, N., Donnett, J. G., Jeffery, K. J., & Maguire, E. A. (1998). Place cells, navigational accuracy, and the human hippocampus. *Philosophical Transactions of the Royal Society of London, Series B: Biological Sciences*, 353(1373), 1333–1340. <http://dx.doi.org/10.1098/rstb.1998.0287>.
- O'Keefe, J., & Nadel, L. (1978). The hippocampus as a cognitive map. *Philosophical Studies*, 2(04), 487–533. <http://dx.doi.org/10.1017/S0140525X00063949>.
- Olton, D. S., & Schlosberg, P. (1978). Food-searching strategies in young rats: Win-shift predominates over win-stay. *Journal of Comparative and Physiological Psychology*, 92(4), 609–618. <http://dx.doi.org/10.1037/h0077492>.
- Pfeiffer, B. E., & Foster, D. J. (2013). Hippocampal place-cell sequences depict future paths to remembered goals. *Nature*. <http://dx.doi.org/10.1038/nature12112>.
- Prasser, D., Milford, M., & Wyeth, G. (2006). Outdoor simultaneous localisation and mapping using RatSLAM. *Field and Service Robotics*, 25, 143–154. http://dx.doi.org/10.1007/978-3-540-33453-8_13.
- Redish, A. D., & Touretzky, D. S. (1998). The role of the hippocampus in solving the Morris water maze. *Neural Computation*, 10(1), 73–111. <http://dx.doi.org/10.1162/089976698300017908>.
- Sargolini, F., Fyhn, M., Hafting, T., McNaughton, B. L., Witter, M. P., Moser, M.-B., et al. (2006). Conjunctive representation of position, direction, and velocity in entorhinal cortex. *Science (New York, N.Y.)*, 312(5774), 758–762. <http://dx.doi.org/10.1126/science.1125572>.
- Sporring, J., Nielsen, M., Florack, L., & Johansen, P. (1997). *Gaussian scale-space theory* (pp. 266). Springer.
- Steele, R. J., & Morris, R. G. M. (1999). Delay-dependent impairment of a matching-to-place task with chronic and intrahippocampal infusion of the NMDA-antagonist D-AP5. *Hippocampus*, 9(2), 118–136. [http://dx.doi.org/10.1002/\(SICI\)1098-1063\(1999\)9:2<118::AID-HIPO4>3.0.CO;2-8](http://dx.doi.org/10.1002/(SICI)1098-1063(1999)9:2<118::AID-HIPO4>3.0.CO;2-8).
- Steffenach, H.-A., Witter, M., Moser, M.-B., & Moser, E. I. (2005). Spatial memory in the rat requires the dorsolateral band of the entorhinal cortex. *Neuron*, 45(2), 301–313. <http://dx.doi.org/10.1016/j.neuron.2004.12.044>.
- Stensola, H., Stensola, T., Solstad, T., Frøland, K., Moser, M.-B., & Moser, E. I. (2012). The entorhinal grid map is discretized. *Nature*, 492(7427), 72–78. <http://dx.doi.org/10.1038/nature11649>.
- Taube, J. S. (2007). The head direction signal: Origins and sensory-motor integration. *Annual Review of Neuroscience*, 30(1), 181–207. <http://dx.doi.org/10.1146/annurev.neuro.29.051605.112854>.
- Taube, J. S., Kesslak, J. P., & Cotman, C. W. (1992). Lesions of the rat postsubiculum impair performance on spatial tasks. *Behavioral and Neural Biology*, 57(2), 131–143. <http://www.ncbi.nlm.nih.gov/pubmed/1586352>.
- Taube, J., Muller, R., & Ranck, J. (1990). Head-direction cells recorded from the postsubiculum in freely moving rats. I. Description and quantitative analysis. *The Journal of Neuroscience*, 10(2), 420–435. <http://www.jneurosci.org/content/10/2/420.abstract>.