



This is an author produced version of *Extending a Hippocampal Model for Navigation Around a Maze Generated from Real-World Data.*

White Rose Research Online URL for this paper:
<http://eprints.whiterose.ac.uk/108448/>

Proceedings Paper:

Boorman, L.W. orcid.org/0000-0001-5189-0232, Damianou, A.C., Martinez-Hernandez, U. et al. (1 more author) (2015) Extending a Hippocampal Model for Navigation Around a Maze Generated from Real-World Data. In: Biomimetic and Biohybrid Systems. 4th International Conference, Living Machines 2015, July 28 - 31, 2015, Barcelona, Spain. Lecture Notes in Computer Science, 9222 . Springer International Publishing , pp. 441-452. ISBN 978-3-319-22978-2

https://doi.org/10.1007/978-3-319-22979-9_44



*promoting access to
White Rose research papers*

eprints@whiterose.ac.uk
<http://eprints.whiterose.ac.uk/>

Extending a hippocampal model for navigation around a maze generated from real-world data

Luke W. Boorman, Andreas C. Damianou, Uriel Martinez-Hernandez, Tony J. Prescott

Sheffield Robotics, University of Sheffield, Sheffield, S10 2TN, U.K.
{l.boorman}@shef.ac.uk

Keywords: Autobiographical memory, Episodic memory, Hippocampus, Robotics, Temporally restricted Boltzmann machine, Navigation

Abstract. An essential component in the formation of understanding is the ability to use past experience in order to comprehend the here and now, and to aid selection of future action. This past experience is stored as memories which are then available for recall at very short notice, allowing for understanding of short and long term action. Autobiographical memory (ABM) is a form of temporally organised memory and is the organisation of episodes and contextual information from an individual's experience into a coherent narrative, which is key to a sense of self. Formation and recall of memories is essential for effective, adaptive behaviour in the world, providing contextual information necessary for planning actions and memory functions, such as event reconstruction. Here we have tested and developed a previously defined computational memory model, based on understanding of hippocampal structure and function, as the base for later developing a synthetic model of human ABM (SAM). The hippocampal model chosen has functions analogous to that of human ABM. We trained the model on real-world sensory data and demonstrate successful, biologically plausible memory formation and recall, in a navigational task. The hippocampal model will be further extended for application in a biologically inspired system for human-robot interaction, which will go beyond current approaches offered by SAM systems.

1 Introduction

For robots to interact with humans in a social manner they need to function in a flexible way. Greater flexibility can arise from robots making inference on future behaviour by recalling relevant past experience. This approach has parallels with human autobiographical memory (ABM), which is defined as the recollection of events from one's life. Human ABM is nonetheless very complex and its exact function is not completely understood. In contrast, navigational memory, especially in the rodent hippocampus has seen much research and is better characterised [1]. Navigational memory does not give insight into the higher level functionality required by human memory, such as comprehension, memory storage and memory recall through language, but instead our current understanding of rodent navigational memory offers a

starting point for the development of a SAM system. A model for spatial navigation, based on a mapping between biological hippocampal function and a temporally restricted Boltzmann machine has been described previously [2, 3]. Following an extensive training phase the model was able to accurately predict the location of a virtual agent within a simple maze from exposure to location specific artificial features and images [4]. The hippocampal model has applicability for robot navigation, but also mimics the key functions of human autobiographical memory, such as memory compression, pattern separation and pattern completion [5]. However, the ability of the current model as a tool for robot navigation and its potential for memory storage and recall has only seen limited testing, as for example, the original navigational task used to test the model, had an overly simplistic maze structure, with carefully placed sensory inputs. Thus, to prepare the model for future use in a more general model of memory, we set out to expand the scope of the model. The first step involved increasing the ‘realness’ of the original maze, by using real-world images from Google Street-view and removing the more ‘artificial’ sensory inputs. The second step involved the construction of a larger and more complex maze. Tests showed the model was able to accurately locate a virtual agent using the real-world sensory image data with the original maze structure. However the model was not able to accurately navigate the larger maze. The software implementation of the original model was adapted to increase flexibility by allowing the automatic adjustment of the number of hippocampal ‘cells’ used for encoding the agents position and the number of ‘cells’ used to store the memories which relate to each spatial location. The model will be tested on the larger more complex maze once the modifications have been completed.

2 The hippocampus, a unitary coherent particle filter model

The rodent hippocampal system has been the basis for much invasive neuroscience research into memory function, especially following the discovery of the function of specific cells in the hippocampus, such as place [6] and grid cells [7]. The hippocampus is thought by many to play a central role in navigation and spatial reasoning [1]. However, it is not only thought to be involved in spatially related tasks, but has also been linked to a wide range of other non-spatial memory functions [8] and even beyond memory, with involvement, for example in decision making and emotion [9]. The hippocampus has been modelled extensively, with the majority of studies using tasks based on navigation to verify their models. The well characterised navigational abilities of the hippocampus, offer direct potential for localisation and mapping for autonomous robots, similar to current simultaneous location and mapping (SLAM) methods [4]. However, the models based on hippocampal function can also offer parallels with features of human autobiographical memory, such as compression, pattern separation and pattern completion [5]. An established unitary coherent particle filter hippocampus (UCPF-HC) model [2] was previously tested using a navigational task [4], for applicability as a tool for SLAM.

Spatial navigation tasks generally require two phases, the first to learn an environment and the second a test to demonstrate the quality of the learning of the environ-

ment, e.g. by traversing the environment to locate a food supply. Previous work has suggested that animals learning can be modelled as a particle filter [10], with sequential learning described using a machine learning approach such as a Temporally Restricted Boltzmann Machine (TRBM). A previous model [2] mapped hippocampal circuitry with a TRBM to produce a navigational learning system (fig. 1). The system was extended across a succession of papers, such as to include learning using a biomimetic sub-theta cycling [3] and to accept visual information using feature extraction [4]. The model is extended within this paper to increase the ‘realness’ of the task, firstly by changing the sensory information within the virtual maze and secondly by building a maze with flexible topology. These extensions are a first step towards a transition towards non-spatial memory, namely to produce a surrogate of human ABM.

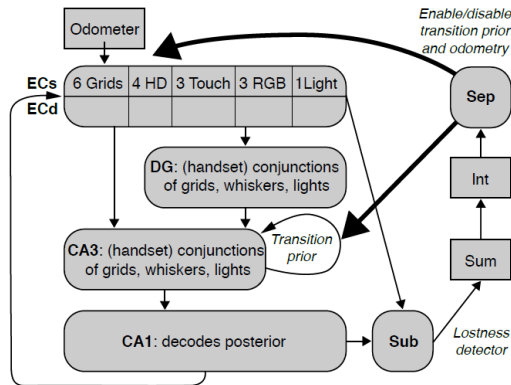


Fig. 1. The hippocampal model, adapted from Fox and Prescott (2010a).

3 The hippocampal model implementation and navigational task

The ability of the existing hippocampal model to store and recall memories was previously tested using a navigational task, based on a simplistic maze shaped as a plus, which had 13 unique locations (Fig. 2). The details of the model implementation are summarised here, for a more complete description see Saul et al. [4].

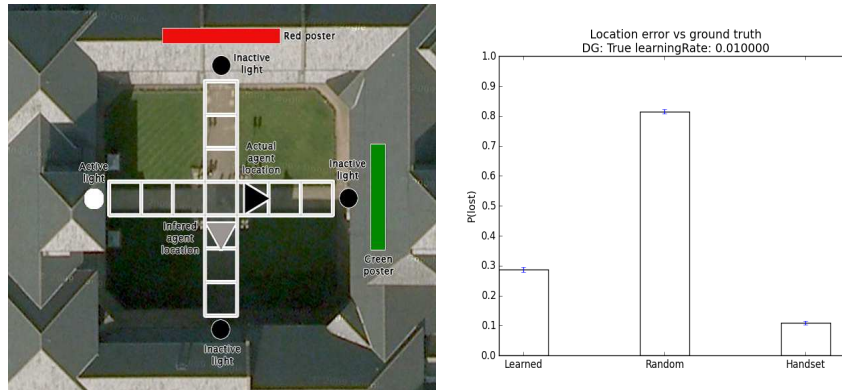


Fig. 2. Left: Plus-shaped maze of the courtyard of the Department of Computer Science, University of Sheffield, taken from Saul et al. [4]. The 13 tiles are marked as white boxes, with light sources at the end of each arm and coloured posters at the ends of two arms. The actual and inferred (by the hippocampal model) locations of the virtual agent are shown. Right: Assessment of the navigational accuracy of the original hippocampal model [4]. The bars show the proportion of steps in which the agent is lost over 3000 randomised steps, as a decimal percentage. ‘Learned’ used the hippocampal model with the weights set by training using 30,000 randomised steps over the maze. ‘Random’ used randomised weights in the model and ‘Handset’ used fixed GPS to include an exaction location of the agent during model training again over 30,000 steps.

The navigational task required a virtual agent to make inference about its current location within the maze, e.g., it produces (x,y) location coordinates. The virtual agent first learns the environment by following a randomly generated path around the maze, while being exposed to the various sensory inputs at each location. The agent is then sent around the maze, using a section of the original path. The agent then reports its inferred location, using sensory information available at each location and information on its previous location. This is made possible within the hippocampal model, by having two underlying systems. The first is a coordinate system which allows the agent to keep track of its location by the inclusion of a series of biologically equivalent cells (Fig. 1), including: 13 place cells, which encode for each maze location: a 3×2 grid of grid cells, which have a unique encoding for of (x,y) ; and 4 head direction cells, which encode direction (e.g., North, South, East, West). The second system is the sensory input and processing modality, which is used by the model to infer its current position within the maze. The ‘plus’ maze has been built to have sensory information available at each location, which, when combined, is unique to that location, thus aiding inference (Fig 2. Left). The sensory inputs made available to the model included: photographs relating to specific directions at each location; coloured markers at the ends of two arms of the maze; touch/whisker sensors marking the location of walls; and an active light source on the west arm. The sensory inputs were processed as an equivalent to compression in ABM. For example, each of the images was converted to a 100-digit binary vector, by extracting Speeded-Up Robust Fea-

tures (SURF) features [11] for each images and merging similar feature vectors using a k-means cluster approach.

The initial learning phase involved training the model (updating the weights within the TRBM), by having the virtual agent take 30,000 random steps around the maze. At each location a number of inputs were made available to the model, including the processed sensory inputs, grid cells marking the previous location, and head directions marking the current direction the agent was facing. Following the training phase, the virtual agent was subject to 3000 movements around the same 'plus' maze. The TRBM weights were fixed. The agent at each step took the sensory inputs, head direction cell firings and previous grid cell encoding and made predictions about its current location by firing the grid cells. The location of the virtual agent was accurately inferred for around 72% of the steps (Fig. 2, right). A secondary loop was included within the model to allow for the detection and reaction to the agent being lost. This lostness system used the differences between the model's expected sensory inputs and its inferred position and the received sensory inputs; if the difference was too great then the priors would be reset (e.g., grid cells were cleared). The model's ability to store and recall memories in an unsupervised manner was compared to a control condition, named 'handset'. The same methods and inputs were used to train and test the model. However, during the training phase the model weights were set by using grid cells' firings based on the exact coordinates of each location (equivalent to using global position system; GPS), rather than using the previous location. The 'handset' control condition offers greater accuracy in predicting the location of the virtual agent for around 90% of the 3000 randomised steps within the maze. In addition to the 'Handset' condition an additional 'Random' condition used randomised weights in the model with the virtual agent's location again predicted for the 3000 randomised steps. Unsurprisingly the ability to detect the location of virtual agent was much lower than for the 'Learned' and 'Handset' conditions, at around 20%.

The hippocampal model demonstrated reliable memory storage and recall in a biologically inspired fashion. The model was able to predict the current location of an agent using inference of past location and processed sensory data specific to that location. There are, however, a number of limitations in both the experimental design used to test the model and the implementation of the model. These not only limit the ability of the model to spatially navigate more 'real' environments, but also limit the applicability of the model for non-spatially specific tasks. We therefore extend the model to use more realistic sensory information and extend the model to navigate mazes with non-fixed topologies.

4 Extending the hippocampal model and navigational task

4.1 Making the task more realistic

To address the limitations imposed by the navigational task used in the original testing of the hippocampal model, a new experimental design has been developed and tested. These modifications are listed below with the ability of the model to navigate the maze assessed for each set of changes.

A novel path and a randomised start. The original navigation task used to test the hippocampal model used both a fixed start location and a section of the same path used in the learning phase. Thus, here the navigational ability of the trained model was tested using a novel path of 100 steps with a randomised start location (Fig. 3, left). The model was able to successfully infer the location of the agent approximately 65% of the time, which is only slightly less accurate than using a 3000 step segment of the original path, where the agent was accurately placed around 72% of the time.

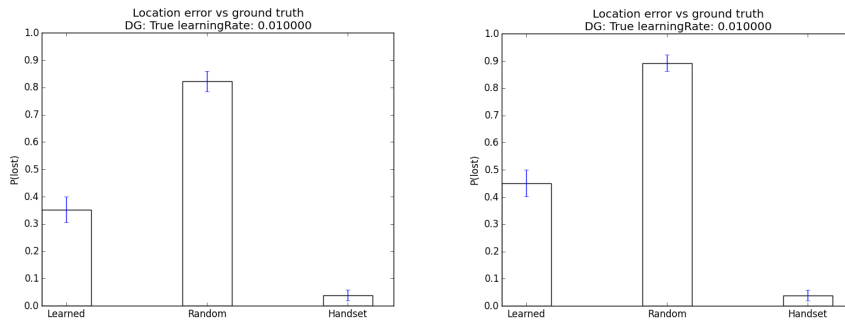


Fig. 3. Left: Assessment of the navigational accuracy of the original hippocampal model, using a novel test path (not a section taken from the training path), with a randomised start. The bars show the proportion of steps that the agent is lost over 100 randomised steps, as a decimal percentage. ‘Learned’ used the hippocampal model with the weights set by training using 30,000 randomised steps over the maze. ‘Random’ used randomised weights in the model and ‘Handset’ used fixed location coordinates (equivalent to using the global positioning system, GPS) to include an exact location of the agent during model training, again over 30,000 steps. Right: Assessment of the navigational accuracy of the original hippocampal model, without the lights or colours placed at the end of each arm of the maze, again using a novel 100 step test path and randomised start.

Non-realistic sensory information removed from the maze. The visual sensory inputs to the original maze were deemed to be unrealistic, for example light sources and coloured panels were ‘carefully’ placed at the extremities of the arms of the plus maze. Thus, the coloured panels and light sources were disabled, as these were not considered to be authentic in a real-world maze environment. The colour and light inputs would also be difficult to place in a complex, non-uniform maze, especially when compared to the original ‘plus’ maze. The hippocampal model when trained (30,000 steps) and tested (100 steps) was able to reliably predict the location of the virtual agent for around 55% of the steps (Fig. 3, right), without the lights or colour inputs included.

‘Real-world’ image data. The model was tested using real-world sensory data from Google Street-view, while retaining the original ‘plus’ maze. The original images used by Saul et al. [4] for the sensory input within the ‘plus’ maze were taken in the Department of Computer Science building courtyard at the university of Sheffield, specifically for use within the model. The images were considered to be relatively

limited both in their framing and content (constrained features available). Google Street-view offers 360° degree panoramic images of the roads around the world, taken with a car which has 360° cameras. Images were extracted manually from Google Street-view on-line. Images were collect at the intersection of Division Street and Carver Street (Fig. 4, left) in the centre of Sheffield, UK. The images were taken from 7 locations along both Division and Carver streets, which produced the identical structure to that of the original ‘plus’ maze. The hippocampal model was trained using the Google Street-view images with 30,000 randomised steps around the maze. The model was then tested using 100 novel randomised steps, with a randomised start and was able to predict the virtual agent’s location within the maze for around 72% of the steps (Fig. 4, right). This suggests the model is able to reliably store and recall memories based on real-world data.

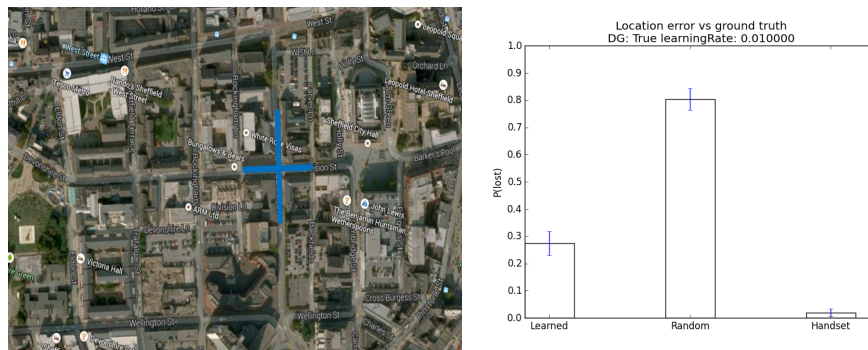


Fig. 4. Left: The locations from where images were manually extracted from Google Street-view and used to generate a ‘plus’ maze. Right: Assessing the ability of the hippocampal model to use real-world sensory data for navigation. The model was trained and tested using the Google Street-view images.

4.2 A flexible, reconfigurable maze

A major limitation of the original experimental design is the fixed size and shape of the ‘plus’ maze, which also carries through into the implementation of the original model. The model has fixed numbers of encoded location tiles, as well as a fixed number of place and grid cells. This lack of flexibility severely limits the applicability of the model to a range of both navigational and non-spatial memory problems. To incrementally update the model, the navigational task was retained but a new larger and more complex maze was built using Google Street-view images for testing the hippocampal model.

Generating the maze. The previous development to the maze, taking ‘real-world’ images, relied on a manual online extraction of Google Street-view images. This approach was considered to be too repetitive, open to experimenter bias and would likely produce errors for the generation of a larger more complex maze. Google offers a

free-to-use uniform resource locator (URL) based application interface (API), which allows the user to request and download Google Street-view images, for example:

<https://maps.googleapis.com/maps/api/streetview?size=480x480&location=53.3794166,-1.4774962&heading=0&pitch=0&fov=90>

The API allows the user to specify the location of the required image as longitude and latitude (Fig. 5, left). Google will not return the image of the exact location requested, but the image nearest to the requested location taken by the Google car when it drove past. The approximate image location makes the generation of the maze more difficult. For example, identical images are often returned for nearby locations and the identical images need to be detected and merged. In addition the non-exact location of the returned images means the physical spacing between images can vary, which adds complexity when attaching images to a fixed grid location. The maze generation module was built to standardise the spacing between images by using the minimum spacing between unique returned images. The Google Street-view API also allows the user to specify a series of parameters with each image request, and includes:

1. Heading (direction in degrees). Four directions were requested for each location e.g. North 0° , East 90° , South 180° and West 270° (Fig. 5, left).
2. Pitch (vertical angle of camera). All images were taken as straight forward (0°).
3. Field of view (width of the view of the image; 1 to 120°). All images taken with width 90° (image of each of the four directions makes up the complete 360°).
4. Resolution (image resolution as x and y). All images were taken as 640×640 pixels.




Fig. 5. Left: Exemplar Google Street-view images extracted using the API, for a single location. The four images shown represent the four directions (left = North, middle = East, right = South and below = West). The images have 0° pitch, 90° field-of-view and a resolution of 640×640 . Right: An aerial view of the centre of Sheffield, UK, from Google Maps is overlaid with the points marked in yellow that were used to predefine the maze. The blue line shows the included sections of road.

The automated generation of the maze uses pre-defined latitude and longitude location coordinates, in decimal degrees, of the start and end of each road and the inter-

sections of one road with another (see yellow circles overlaid onto the map in Fig. 5, right), to be included. A minimum step spacing of 0.000125 decimal degrees was used as a stepping between requested image locations along each road. Returned blocks of four images (North, South, East, and West) were checked to make sure they were not identical to previous images. The image block was then assigned a grid reference, with x (East positive, west negative) and y (South positive, north negative). The image files were stored as jpegs and named (for example, as '012-000-N-001.jpg'), using the nomenclature: x (zero padded to the 3 digits), y (zero padded to the 3 digits), direction as N, S, E or W and the image identifier (zero padded to the 3 digits). The image identifier was included to allow storage of multiple sets of images (e.g., to allow for inclusion of more images from the Google Street-view history feature at a later date). A total of 708 images were requested and were used to build the maze (blue lines in Fig. 5, right).

A flexible hippocampal model. The original model [4] was pre-set to use the 'plus' maze. In order for the model to accept the new more complex maze, a substantial number of modifications were required. The complexity of the model made this a slow process as the 'plus' maze was hard coded at multiple points. The first step involved a completely new maze generation module, which was provided the complete folder of images extracted from Google Street-view and generated the maze from the image files available within the folder. This allows for flexible maze generation, as different sets of images will produce different maze topologies in an automated manner. The increased complexity of the maze called for an easier approach to check the automatically generated maze. The checks need to confirm the appropriate images were present at each location and were correct in terms of their location, direction and whether images at adjacent locations were consecutive. A graphical user interface (GUI) was developed to load the images and allowed a human user to navigate the maze using key presses, showing which direction steps were possible for each grid location. A map was automatically generated from the available images and included the location of the agent and the previous path of the agent (Fig. 6).




Fig. 6. Left: Google maps aerial image of Sheffield. Right: An autonomously generated top-down map of the maze. The map was generated from the locations of the images available in the given folder. The white boxes represent the locations where images are available, the red box represents the current location of the virtual agent and the green boxes represent the previous path of the agent.

The sensory inputs at each location within the original plus maze were hard coded into the model. Therefore each sensory input had to be adapted for the autonomously generated maze and ‘dictionaries’ of sensory inputs built. Examples include:

1. Touch input – This gives details of the walls that are present [left, forward, right] e.g., dead end [1,1,1], intersection [0,0,0].
2. Available moves – Can the agent go; forwards, left, right or do a u-turn at each position.
3. Available grid locations – This gives the resulting location of the virtual agent following each type of move.
4. Grid cell encoding – The encoding of the x and y locations. The grid cell encoding had to be adapted by adjusting the length of the binary encoded vectors of x and y to have enough resolution to accept the full range of x (original ‘plus’ maze 0 to 6, new maze 0 to 42) and y (‘plus’ maze 0 to 6, new maze 0 to 26) values.
5. Head direction – The direction the agent is facing.

The GUI was found to be useful for testing the adjustments made to the model and was used throughout the different phases of training and testing. The modified hippocampal model was trained on the data, but was unable to accurately infer the location of the agent within the maze. Further modifications are required to the implementation of the model to allow it to use the more complex and larger maze.

5 Conclusions and future directions

This paper developed an existing hippocampal model [4] for initial application in a spatial navigation task and with the future aim of further developing the model for application to non-spatial memory. We firstly demonstrated the ability of the model to predict locations within a ‘realistic’ simple maze, which was built using data similar to that used by humans for spatial navigation (images from Google Street-view) and the pre-existing ‘un-realistic’ sensory streams were disabled, such as the lights and colours placed at each branch of the maze. Secondly to assess whether the model could scale to a more complex, less structured task, the spatial size and complexity of the maze used within the navigational task was extended. This ability to scale is essential in producing a system to mirror the functions of human ABM, where multiple streams of sensory information will need to be compressed and processed ready for storage and multiple predictions will need to be made through recall of multiple previous memories.

An automatic maze generation system was produced to generate the larger more complex maze. Initial testing with the larger maze revealed the model implementation

was hardcoded for the original ‘plus-maze’. The hippocampal model implementation was updated to allow for learning of mazes with different topologies, using scalable numbers of grid and place cells. An interactive visualisation tool was used to verify location specific compression of sensory information and to observe the agent moving around the maze using the paths generated for learning and testing. Testing showed the model was unable to predict the location of the virtual agent within the larger maze. A number of avenues are currently being explored to develop the model to learn and accurately navigate the larger maze, these include replacing the TRBM with deep Gaussian process techniques [12] and using active learning to improve the learning phase, for example by exploring the maze with the agent having a preference for novelty, rather than using a random exploration approach.

5.1 Towards a model of human autobiographical memory

The model and navigational task developed here are limited in that they replicate only spatial memory, however, they offer useful insight into practical implementations of memory systems, which go beyond spatial memory, such as human ABM. The testing of the model has demonstrated the model has some of the functionality associated with human ABM [5]:

- **Compression.** The model uses extrinsic compression with the extraction of SURF features from the Street-view images and the nearest neighbour clustering of features. The model also demonstrates intrinsic compression with the encoding of the sensory data within the TRBM during learning.
- **Pattern separation.** The model was able to infer the location of a virtual agent within the Street-view plus-maze, using the provided sensory inputs available. The ability to recall unique locations demonstrates the ability of the model to store and recall information specific to each location, despite the need for compression. Thus, information is not merged where it is deemed to be essential.
- **Pattern completion.** The initial model implementation used a series of well-placed sensory features, we removed these features, replacing them with real-world imagery. The model demonstrated pattern completion, as it was able to infer the location of the virtual agent within the plus-maze with only the limited sensory information available e.g. inferring location using only a set of image features for that location.

5.2 Future objectives

The original model and experimental design have been developed to overcome the immediate limitations. However, our interest is in developing a synthetic autobiographical memory system (SAM) for application in robotics. We aim to build a system that will use both spatial and non-spatial memory formation and recall. This will be used with the iCub humanoid platform and will endow the robot with the ability to interact socially with humans, through recognition of humans, remembering actions and generation of specific actions and language. We will work towards the conver-

gence of SAM systems based on human memory function [13, 14] with our models which use well characterised biology.

6 Acknowledgements

The authors would like to thank the EU for funding this research (EU grant no. 612139 WYSIWYD - “What You Say Is What You Did”), our colleagues at Sheffield Robotics and Google for providing the Street-view images and API.

1. Moser, E.I., Kropff, E., Moser, M.-B.: Place Cells, Grid Cells, and the Brain's Spatial Representation System. *Annual Review of Neuroscience* 31, 69-89 (2008)
2. Fox, C., Prescott, T.: Hippocampus as unitary coherent particle filter. In: *Neural Networks (IJCNN), The 2010 International Joint Conference on*, pp. 1-8. IEEE, (Year)
3. Fox, C., Prescott, T.: Learning in a unitary coherent hippocampus. *Artificial Neural Networks-ICANN 2010*, pp. 388-394. Springer (2010)
4. Saul, A., Prescott, T., Fox, C.: Scaling up a Boltzmann machine model of Hippocampus with visual features for mobile robots. In: *Robotics and Biomimetics (ROBIO), 2011 IEEE International Conference on*, pp. 835-840. IEEE, (Year)
5. Evans, M., Fox, C., Prescott, T.: Machines Learning - Towards a New Synthetic Autobiographical Memory. In: Duff, A., Lepora, N., Mura, A., Prescott, T., Verschure, P.M.J. (eds.) *Biomimetic and Biohybrid Systems*, vol. 8608, pp. 84-96. Springer International Publishing (2014)
6. O'Keefe, J.: Place units in the hippocampus of the freely moving rat. *Experimental Neurology* 51, 78-109 (1976)
7. Hafting, T., Fyhn, M., Molden, S., Moser, M.-B., Moser, E.I.: Microstructure of a spatial map in the entorhinal cortex. *Nature* 436, 801-806 (2005)
8. Eichenbaum, H.: Is the rodent hippocampus just for ‘place’? *Current Opinion in Neurobiology* 6, 187-195 (1996)
9. Cameron, H.A., Glover, L.R.: Adult Neurogenesis: Beyond Learning and Memory. *Annual Review of Psychology* 66, 53-81 (2015)
10. Courville, N.D.D.A.C.: The pigeon as particle filter. *Advances in Neural Information Processing Systems* 20, 369-376 (2008)
11. Bay, H., Tuytelaars, T., Van Gool, L.: Surf: Speeded up robust features. *Computer Vision-ECCV 2006*, pp. 404-417. Springer (2006)
12. Damianou, A.C., Lawrence, N.D.: Deep Gaussian Processes. *arXiv preprint arXiv:1211.0358* (2012)
13. Pointeau, G., Petit, M., Dominey, P.F.: Embodied simulation based on autobiographical memory. *Biomimetic and Biohybrid Systems*, pp. 240-250. Springer (2013)
14. Pointeau, G., Petit, M., Dominey, P.F.: Successive Developmental Levels of Autobiographical Memory for Learning Through Social Interaction. *Autonomous Mental Development, IEEE Transactions on* 6, 200-212 (2014)