



Queensland University of Technology
Brisbane Australia

This may be the author's version of a work that was submitted/accepted for publication in the following source:

[Chen, Zetao, Jacobson, Adam, Erdem, Ugur, Hasselmo, Michael, & Milford, Michael](#)

(2014)

Multi-scale bio-inspired place recognition.

In Tan, J & Hamel, W R (Eds.) *Proceedings of the 2014 IEEE International Conference on Robotics and Automation (ICRA 2014)*.

Institute of Electrical and Electronics Engineers Inc., United States of America, pp. 1895-1901.

This file was downloaded from: <https://eprints.qut.edu.au/73412/>

© Consult author(s) regarding copyright matters

This work is covered by copyright. Unless the document is being made available under a Creative Commons Licence, you must assume that re-use is limited to personal use and that permission from the copyright owner must be obtained for all other uses. If the document is available under a Creative Commons License (or other specified license) then refer to the Licence for details of permitted re-use. It is a condition of access that users recognise and abide by the legal requirements associated with these rights. If you believe that this work infringes copyright please provide details by email to qut.copyright@qut.edu.au

Notice: *Please note that this document may not be the Version of Record (i.e. published version) of the work. Author manuscript versions (as Submitted for peer review or as Accepted for publication after peer review) can be identified by an absence of publisher branding and/or typeset appearance. If there is any doubt, please refer to the published source.*

<https://doi.org/10.1109/ICRA.2014.6907109>

Multi-scale Bio-inspired Place Recognition

Zetao Chen, Adam Jacobson, Uğur M. Erdem, Michael E. Hasselmo and Michael Milford

Abstract— This paper presents a novel place recognition algorithm inspired by the recent discovery of overlapping and multi-scale spatial maps in the rodent brain. We mimic this hierarchical framework by training arrays of Support Vector Machines to recognize places at multiple spatial scales. Place match hypotheses are then cross-validated across all spatial scales, a process which combines the spatial specificity of the finest spatial map with the consensus provided by broader mapping scales. Experiments on three real-world datasets including a large robotics benchmark demonstrate that mapping over multiple scales uniformly improves place recognition performance over a single scale approach without sacrificing localization accuracy. We present analysis that illustrates how matching over multiple scales leads to better place recognition performance and discuss several promising areas for future investigation.

I. INTRODUCTION

Robotic mapping and localization systems typically operate at either one fixed spatial scale, or over two, combining a local and a global scale [1-3]. In contrast, recent high profile discoveries in neuroscience have indicated that animals, such as rodents, navigate the world using multiple parallel maps, with each map encoding the world at a specific spatial scale [4, 5]. The multi-scale rodent mapping system consists of neurons that encode areas ranging from several square centimetres to several square meters, with many intermediate scales represented in-between. Unlike hybrid metric-topological multi-scale robot mapping systems, rodent maps are homogeneous, distinguishable only by scale. Although theoretical studies have highlighted computational benefits of a multi-scale mapping system [6, 7], no real world experiments have been done to investigate these principles.

In this paper, we present a biologically-inspired multi-scale mapping system mimicking the rodent multi-scale map. Our approach utilizes multiple arrays of Support Vector Machines, with each array trained to perform place recognition at a specific spatial scale, and a process for combining place recognition hypotheses from these different

spatial scales (Figure 1). Unlike traditional probabilistic robotics methods, where spatial specificity is passively determined by sensor observation models, our approach intentionally creates parallel training systems to map the sensor input to the environment at different spatial scales.




Figure 1. Illustration of our multi-scale place recognition system. Multiple parallel SVMs are trained to recognize places at different spatial scales and filter out hypotheses not supported by all scales. In this example, the number one ranked match at the highest spatial precision (yellow) is not supported by matches at lower spatial precisions. In contrast, the second ranked match is supported at all scales and is consequently correctly chosen as the place match.

We conduct experiments on three real world datasets and compare single- and multi-scale place recognition performance. We extend an initial pilot study [8] by presenting a new, adaptive method for combining multi-scale spatial hypotheses based on SVM firing score, rather than the manual approach described in [8]. For the first time the method is able to achieve a uniform improvement across all presented studies, improving the recall rate at 100% precision by an average factor of 77%. We present a new visualization method that illustrates how place hypotheses at different scales are combined and present for the first time results on the benchmark 70 km Eynsham dataset.

The paper is organized as follows. Section II discusses related place recognition and mapping techniques. In Section III we describe the components of the multi-scale place learning system. Experiments are presented in Section IV, with results shown in Section V. Finally we conclude the paper in Section VI by discussing ongoing and future work.

II. BACKGROUND

In this section, we summarize single- and two-scale robotic mapping methods and describe recent evidence regarding

ZC, AJ and MM are with the School of Electrical Engineering and Computer Science at the Queensland University of Technology, zetao.chen@student.qut.edu.au. UE and MH are with Center for Memory and Brain and Graduate Program for Neuroscience at Boston University. This work was supported by an Australian Research Council Discovery Project DP1212775 awarded to MM, and an Office of Naval Research ONR MURI N00014-10-1-0936 and Silvio O. Conte Center Grant P50 NIMH MH094263 to UME and MEH.

multiple parallel mapping systems in the mammalian brain.

A. *Robotic Mapping Methods*

There has been extensive research in robotics mapping and localization [9] over the past two decades. Since cameras have become a common sensor modality in many robot platforms, a large number of vision-based mapping and place recognition algorithms have been proposed. Here we briefly discuss some of the key systems, although the list is by no means exhaustive.

Some of the most significant vision-based mapping and localization algorithms proposed over the past decade include FAB-MAP [10, 11], MonoSLAM [12], RatSLAM [13-16], FrameSLAM [17] and SeqSLAM [18, 19]. Those methods primarily focus on performing localization at one fixed scale.

Hybrid approaches combine local and global mapping techniques. Atlas [1] is a hybrid SLAM algorithm which combines existing small-scale mapping algorithms with a global topology to achieve real-time large-scale navigation. Similarly, a hybrid extension to the Spatial Semantic Hierarchy [2, 3] builds local maps using metric SLAM methods but represents the structure of large-scale space using a topological map. The mapping frameworks in these and other approaches have generally been limited to two distinct scales and heterogeneous, in that different types of maps are used at different scales. Interestingly, the concept of multiple scale maps has perhaps been most thoroughly investigated in the temporal rather than spatial domain [20].

B. *A Multi-Scale Neuronal Map*

In contrast to mobile robots, which are typically only capable of operating within one specific type of environment, rats are the second most widespread mammal after humans. Wild rats found in Alaska are genetically similar to “urban” rats that spend their time living in the walls of a house or sewer, as are burrow rats that never come above ground and rats living in an open air Asian market or Australian garbage dump. Rats across all these environments map and perform place recognition using the same neuronal machinery. Underpinning these impressive rodent capabilities is a sophisticated sensory processing and mapping system.

The rodent entorhinal-hippocampal (EC) formation, and more specifically the *medial entorhinal cortex* (MEC) stores spatial information (a map) in the form of a highly regular, grid-like representation of space [5]. The primary spatially-responsive cell in MEC is called a “grid” cell [5] – grid cells fire whenever the rodent (and likely other animals such as bats [21]), is located at a vertex of a regular grid of locations over the environment. Recent evidence has shown that grid cells encode *multiple, discrete* scales of grid locations, in steps of $\sqrt{2}$ [4]. The area encoded by a cell at each grid vertex can vary from a few square centimetres to tens of square metres. The upper limit, if there is one, is unknown. This integrated, multi-scale representation has

been shown to have a number of theoretical advantages including efficient mapping of arbitrarily large environments [6, 7].

C. *RatSLAM*

The approach described here is perhaps most closely related to the RatSLAM system. RatSLAM is a computational model of the part of the rodent brain thought to be responsible for mapping called the hippocampus. The system creates a simplified single-scale neural spatial map similar to that stored in the MEC [5]. RatSLAM has been demonstrated in a number of experiments mapping large environments [16] or over long periods of time [14]. Much of the representational power of RatSLAM is drawn from the temporal dynamics of the network over time.

III. APPROACH AND METHODOLOGY

Our overall approach comprises three stages: image feature extraction, place learning and place recall at each spatial scale, and combination of place match hypotheses from each scale to produce an overall place match hypothesis.

A. *Feature Extraction*

Dimensional reduction was performed before the images were input to the SVMs. We implemented two commonly used feature extraction methods – Principal Component Analysis (PCA) and Gist.

1) *Principal Component Analysis*

PCA [22] is an efficient dimensional reduction method. We applied PCA to down sampled camera images (32×10) and selected the top 38 principal eigenvectors, representing 90% of the dataset variance.

2) *Gist*

A variety of experimental studies have demonstrated that humans perform rapid categorization of a scene by integrating only the coarse global information or “gist” [23, 24]. Using the model proposed by Oliva [25], we extracted Gist features from down sampled (128×40) images resulting in a 512-dimensional feature. We then extracted the top 32 principle eigenvectors, which captured approximately 90% of the total variance.

B. *Learning Algorithm*

We chose Support Vector Machines (SVM) [26] to learn and recall places based on their ease of use (in contrast to more biologically plausible models such as continuous attractor networks [15]) and applicability to the place recognition problem, where both positive and negative place match training exemplars are available.

We denote the data and its label as (x_i, y_i) where $x_i \in R^N$ is an N-dimensional feature and $y_i \in \{+1, -1\}$ is its label. With the assumption that the data can be separated by a hyperplane in some Hilbert space H, we search for the optimal separating hyperplane that maximizes its distance to

the closest points in the training data, resulting in a discriminant function:

$$f(x) = \sum_{i=1}^m a_i y_i K(x_i, x) + b \quad (1)$$

The sign of $f(x)$ indicates the classification result and a_i and b can be found by solving a constrained minimization problem, efficiently calculated using the SMO algorithm [27]. Because visual place recognition data is typically non-linearly separable, a slack variable is introduced to enable violations of the classification margin while minimizing the classification error. A parameter C controls this trade-off.

The kernel function $K(x_i, x)$ maps the input data to a high dimensional feature space H where the non-linearly separable data may be separated linearly. Here we use the popular RBF kernel.

C. Place Recognition using an Array of SVMs

Each array of SVMs produces a firing matrix representing the matching scores of the testing segments on the trained SVMs. The left matrix in Figure 2 shows a summary of SVM firing scores over a 20 segment test dataset when compared to a 20 segment training dataset. Element $M(i, j)$ indicates the response of the i^{th} SVM from a training dataset to the j^{th} segment in a test dataset. Similarly, the j^{th} column provides a place recognition distribution over all training segments. Firing scores are normalized to sum to one for each place recognition distribution:

$$M(i, j) = \frac{M(i, j)}{\sum_i M(i, j)} \quad (2)$$

D. Multi-Scale Place Matching Verification

Place recognition hypotheses produced by each array of SVMs are only as accurate as the average size of a segment in that array. For example, in Figure 2, if a 1000 metre long dataset traversed at approximately constant speed is divided into 20 segments (left, Figure 2), place matches are only accurate to a 50 meters, compared to 25 meters for 40 segments and so on (right, Figure 2). Here we present a two-step method for combining place recognition hypotheses at these different spatial scales.




Figure 2. An example confusion matrix showing the distribution of SVM scores over an input dataset comprising 20 segments (left) and 40 segments (right).

1) Finding Overlapping Hypotheses

To identify overlapping place recognition hypotheses (consistent place matches reported at every mapping scale, see Fig. 3), we first normalize the reported place recognition matches to the scale space of the smallest segment size. For K arrays of SVMs, after normalization the matching scores of each array are:

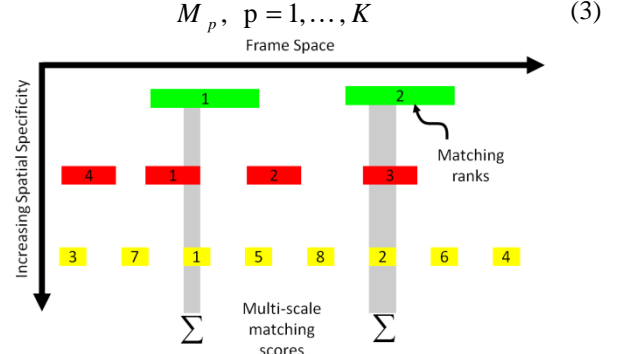


Figure 3. Overlapping SVM matching scores are combined at the smallest spatial scale in order to accept or reject place match hypotheses. In this case, $K=3$.

Suppose there are L_p training segments for the matching score M_p . For a segment j in a test data set, its coherence measurement on each training segment $c(i, j), i=1, \dots, L_p$ depends on whether there are spatially overlapping hypotheses over all SVMs scales. Without support at all spatial scales, the system reports “no coherent” match ($c=0$):

$$c(i, j) = \begin{cases} 1, & M_p(i, j) > 0, \forall p \\ 0, & \text{else} \end{cases} \quad (4)$$

2) Finding the Best Overlapping Match Candidate

It is possible to have multiple competing place recognition hypotheses supported by all other spatial scales. To determine the most likely hypothesis, we sum the firing scores of the overlapping SVMs at each spatial scale and classify segment j to the class $C(j)$ with largest accumulative firing scores:

$$C(j) = \arg \max_i \sum_p M_p(i, j), \forall c(i, j) = 1 \quad (5)$$

IV. EXPERIMENTAL SETUP

In this section, we describe the data sets used and the SVM training procedure.

A. Datasets

Three datasets are tested in this experiment with details summarized in Table I. Each dataset consists of two traverses along the same route with the first traverse used for training and the second traverse for testing.

The Eynsham dataset (Figure 4a) is a large 70 km road-based dataset (2×35 km traverses) used in the [10] FAB-MAP and SeqSLAM studies. Panoramic images were captured at 7 meter intervals using a Ladybug 2 camera.

We also performed experiments on two additional datasets

first in [8], in order to demonstrate the improvement from using the new multi-scale combination process. The Rowrah dataset was collected from a forward-facing camera mounted on a motorbike. The Campus dataset was collected using a GoPro Hero 1 camera mounted on a bicycle pushed by an experimenter. Dataset descriptions are provided in Table I.

TABLE I
DATASET DESCRIPTIONS

Dataset Name	Total Distance	Number of Frames per Traverse	Distance between frames
Eynsham	70 km	9575	6.7 m (median)
Campus	1.6 km	1000	1.6 m
Rowrah	2k m	1570	1.26




Figure 4. The (a) 70 km Eynsham, (b) 2 km Rowrah and (c) 1.6 km Campus environment, each of which consists of two traverses along the same route

B. Ground Truth

We used the 40 metre tolerance GPS-derived ground truth provided with the Eynsham dataset, consistent with the tolerance used in the original FAB-MAP study [10]. Ground truth for the Rowrah and Campus datasets was obtained by interpolating manually selected frame correspondences.

C. Training and Testing Procedure

Images from the first traverse of the environment were used for training while images from the second traverse were used to evaluate performance. The overall training procedure consisted of the following three steps: dataset segmentation, feature extraction and SVM training.

1) Dataset Segmentation

The images in each dataset were grouped into a total of S subsequent segments ($S/2$ segments per traverse). For the sake of intuition, in the later experiment results, we refer to different SVM arrays by the spatial size encoded by each segment. For example, the Eynsham dataset contains 9575 frames and therefore splitting it into 320 segments (160 segments per traverse) resulted in an average segment size of 30 frames, representing approximately 200 meters.

Place recognition using a segment size of 200 meters is only accurate to a distance of 200 meters. To achieve more precise place match reporting, we segmented the smallest spatial scale for the Eynsham dataset using overlapping segments offset by 6 frames corresponding to approximately 40 meters (Figure 5). For example, segment 1 covers frames 1 to 30, segment 2 overlaps segment 1 and covers frames 7 to 36. Larger segment sizes were not overlapped. We trialed a

variety of segment sizes, listed in Table II.




Figure 5. The Eynsham dataset was divided into overlapping segments in order to increase the spatial precision of the place matching process.

TABLE II
DATASET SEGMENTATION AND SEGMENT SIZES

Number of Segments	Eynsham (m)	Campus (m)	Rowrah (m)
320	200		
160	400		
80	800		
30		53	67
43		40	47
90			22
170		9	

2) Feature Extraction

Two feature types (as discussed in Section IIIA) were extracted from each dataset. The feature vectors from all frames in a segment were consolidated into a single vector and input into each of the SVMs.

3) SVM Training

To train a SVM model for each segment, we labeled the images in that segment as positive examples and other ($S-I$) groups of images as negative examples. Cross validation was used to choose optimal values for the kernel width θ and the slack variable C .

V. RESULTS

We present three sets of results on the Eynsham dataset – performance comparison between single and multi-scale place recognition, ground truth plots and illustrative multi-scale place recognition combination plots. We also summarize performance on the Rowrah and Campus datasets.

A. Single- and Multi-scale Place Recognition

This section presents precision recall (PR) curves resulting for the single- and multi-scale place recognition experiments using Gist and PCA features on the Eynsham dataset (Figures 6 and 9). Two set of combinations are shown – “combined 200 400 meter system” and “combined 200 400 and 600 meter system”, as well as results from using a single 200 meter scale. Figures 7 and 10 display the relative error rate (number of false positives divided by the total number of recalled frames) for one, two and three spatial scales at varying recall rates.

Multi-scale matching consistently improves the

performance. Using Gist features, the recall rate at 100% precision improves by a factor of nearly 3 from 8% using a single-scale to 28% when using three scales. The corresponding improvement using PCA features is from 14% to 23%. Using two scales results in an intermediate performance improvement. Matching coverage is relatively evenly distributed throughout the dataset, with the largest gap at 100% precision measuring approximately 2 km in length.

Although our focus is on the improvement potential offered by adopting a multi-scale approach, we provide absolute comparison metrics here. The maximum recall rate of 28% at 100% precision is superior to the baseline FAB-MAP performance, comparable to the motion-model FAB-MAP performance and less than the approximately 50% recall achieved with FAB-MAP using both a motion model and epipolar geometric verification, and comparable to the 20 frame SeqSLAM implementation [18].




Figure 6. Precision recall curves demonstrating the single- (“200 meter system”) and multi-scale (“combined 200, 400 meter system” and “combined 200, 400 and 800 meter system”) place recognition performance using gist feature on the Eynsham dataset




Figure 7. Error rates using (from left to right in each bar cluster) one, two or three spatial mapping scales at various recall rates using Gist features.

B. Ground Truth Plots

Figure 11a-b presents ground truth plots showing the true positives (green circles), false positives (blue squares) and false negatives (red stars) output by the single and three-scale systems for the Eynsham dataset at an identical recall rate. Straight lines connect the matching segments. The introduction of multiple matching scales removes most of the false positive matches from the single scale results. Further analysis of the multi-scale false positive match is provided in the following subsection.




Figure 8. Place recognition coverage at 100% precision and 28% recall on the Eynsham dataset. Coverage is generally distributed with a worst case gap approximately 2 km long.




Figure 9. Precision recall curves demonstrating the single- (“200 meter system”) and multi-scale (“combined 200, 400 meter system” and “combined 200, 400 and 800 meter system”) place recognition performance using PCA feature on the Eynsham dataset.




Figure 10. Error rates using (from left to right in each bar cluster) one, two or three spatial mapping scales at various recall rates using PCA features.

C. Multi-hypothesis Combination Plots

Figures 11c-f show examples of how place match hypotheses at varying scales are combined together. In general, a large number of false positives at the smallest spatial scale (bottom yellow row) are eliminated due to lack of support from larger spatial scales. The examples in (c-d) show how secondary ranked spatially specific matches are correctly chosen as the overall place match due to support from other spatial scales. In (e) the best ranked spatially specific match is correctly supported by the other spatial scales, while (f) shows a failure case where the incorrect 4th ranked spatially specific match is more strongly supported by the other spatial scales than the 1st ranked and correct spatially specific match.

D. Campus and Rowrah Dataset

In these two datasets, the proposed method improves the recall rate by an average factor of 74.79% across all experiments at 100% precision. This performance represents a 34% improvement in the recall rate at 100% precision over that presented in the original study [8]. A larger improvement is achieved using Gist than PCA; at 100% precision, the recall rate for Gist was improved by an average of 81.7% over all experiments, versus 67.9% for PCA.

VI. DISCUSSION AND FUTURE WORK

We have demonstrated that implementing a multi-scale place recognition system improves place recognition performance by combining the output from parallel mapping frameworks, each trained to recognize places at a specific spatial scale. Although this paper presents a specific visual pre-processing techniques and learning mechanisms, we believe that the novel multi-scale combination concept should generalize to other sensor types, sensor processing schemes and learning methods. In this section we discuss several areas of current and future work.

The current system assumes that the camera is moving at a constant speed during the training and testing stages.

Incorporating an odometry source will allow the system to allocate segments directly based on spatial distances travelled rather than (in effect) time. Moreover, incorporating odometry information will enable us to expand our current system to two-dimensional unconstrained movement in large open environments. Testing the system in open field environments will be more analogous to many current rodent experiments and may increase the likelihood of generating neuroscience insights.

The next step beyond odometry-driven segmentation is data-driven segmentation, where an environment is segmented based on local self-similarity. Such an approach would avoid inefficient representations of large bland spaces with small spatial scale maps. Furthermore, in large open spaces, precise localization is often not possible; in such a situation it may be possible to fall back to a less spatially specific place recognition estimate that uses broader visual cues. It may also be possible to improve the algorithm's efficiency by performing selective hypothesis validation using a "top-down" approach; only searching for finer scale place matches in areas of a map that are matched at the broadest level.

Recent work using RatSLAM has shown that biologically

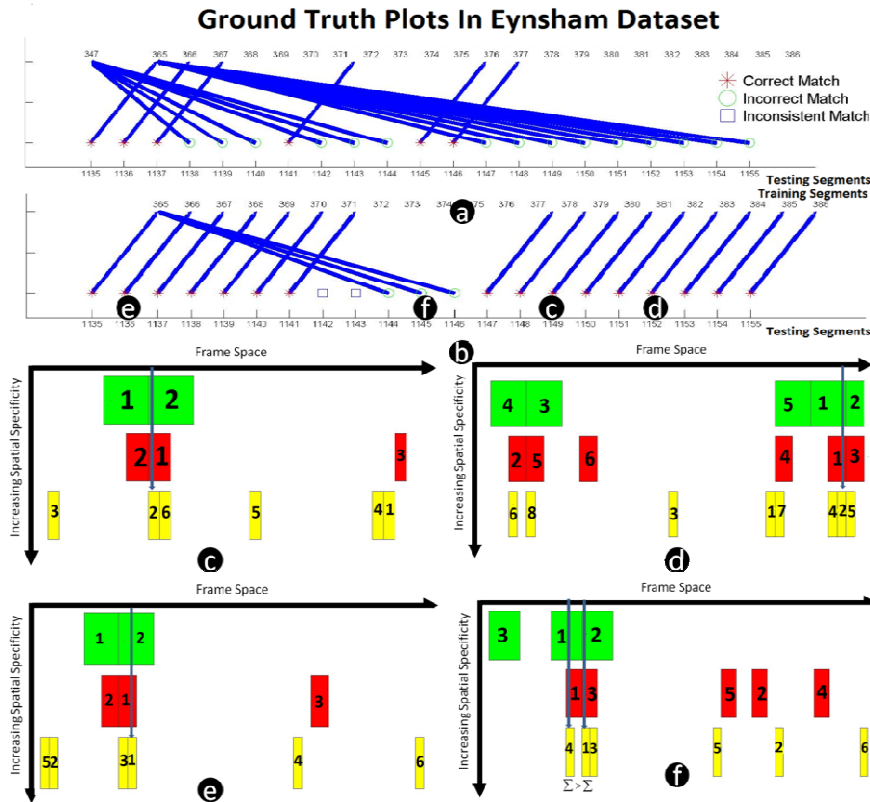


Figure 11. Ground truth plots for the (a) single and (b) multi-scale Eynsham dataset. (c-d) show examples of secondary-ranked spatially specific place matches (yellow) that became the primary overall place match hypothesis due to support from other spatial scales. In (e) the first ranked spatially specific match is supported, while (f) shows a failure case where a secondary ranked spatially specific match is incorrectly chosen as the overall match due to more significant support from the other spatial scales than the correct, first ranked spatially specific match.

inspired algorithms can perform online sensor fusion to enable place recognition in changing environmental conditions, such as over day-night cycles [28, 29]. An obvious extension to this research would be to use a multi-scale mapping framework to exploit the variable spatial specificity of different sensor modalities, such as cameras, range finders and WiFi. By integrating these multi-sensor fusion systems with a biologically-inspired, multi-scale mapping framework, it may be possible to combine their functional capabilities to produce a highly capable, general purpose robot mapping and navigation system.

REFERENCES

- [1] M. Bosse, P. Newman, J. Leonard, M. Soika, W. Feiten, and S. Teller, "An atlas framework for scalable mapping," in *International Conference on Robotics and Automation*, Taipei, Taiwan, 2003, pp. 1899-1906.
- [2] B. Kuipers, J. Modayil, P. Beeson, M. MacMahon, and F. Savelli, "Local Metrical and Global Topological Maps in the Hybrid Spatial Semantic Hierarchy," in *International Conference on Robotics and Automation*, New Orleans, USA, 2004.
- [3] B. Kuipers and Y. T. Byun, "A Robot Exploration and Mapping Strategy Based on a Semantic Hierarchy of Spatial Representations," *Robotics and Autonomous Systems*, vol. 8, pp. 47-63, 1991.
- [4] H. Stensola, T. Stensola, T. Solstad, K. Froland, M. Moser, and E. Moser, "The entorhinal grid map is discretized," *Nature*, vol. 492, pp. 72-78, 2012.
- [5] T. Hafting, M. Fyhn, S. Molden, M.-B. Moser, and E. I. Moser, "Microstructure of a spatial map in the entorhinal cortex," *Nature*, vol. 11, pp. 801-806, 2005.
- [6] Y. Burak and I. R. Fiete, "Accurate path integration in continuous attractor network models of grid cells," *PLoS Computational Biology*, vol. 5, 2009.
- [7] P. E. Welinder, Y. Burak, and I. R. Fiete, "Grid cells: the position code, neural network models of activity, and the problem of learning," *Hippocampus*, vol. 18, pp. 1283-1300, 2008.
- [8] Z. Chen, A. Jacobson, and M. Milford, "Bio-inspired Place Recognition over Multiple Spatial Scales," presented at the Australasian Conference on Robotics and Automation, Sydney, Australia, 2013.
- [9] G. Dissanayake, P. M. Newman, S. Clark, H. Durrant-Whyte, and M. Csorba, "A solution to the simultaneous localisation and map building (SLAM) problem," *IEEE Transactions on Robotics and Automation*, vol. 17, pp. 229-241, June, 2001 2001.
- [10] M. Cummins and P. Newman, "Highly scalable appearance-only SLAM - FAB-MAP 2.0," in *Robotics: Science and Systems*, Seattle, United States, 2009.
- [11] M. Cummins and P. Newman, "FAB-MAP: Probabilistic Localization and Mapping in the Space of Appearance," *International Journal of Robotics Research*, vol. 27, pp. 647-665, 2008.
- [12] A. J. Davison, I. D. Reid, N. D. Molton, and O. Stasse, "MonoSLAM: Real-Time Single Camera SLAM," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 29, pp. 1052-1067, 2007.
- [13] D. Ball, S. Heath, J. Wiles, G. Wyeth, P. Corke, and M. Milford, "OpenRatSLAM: an open source brain-based SLAM system," *Autonomous Robots*, pp. 1-28, 2013/02/21 2013.
- [14] M. Milford and G. Wyeth, "Persistent Navigation and Mapping using a Biologically Inspired SLAM System," *International Journal of Robotics Research*, vol. 29, pp. 1131-1153, 2010.
- [15] M. J. Milford, *Robot Navigation from Nature: Simultaneous Localisation, Mapping, and Path Planning Based on Hippocampal Models* vol. 41. Berlin-Heidelberg: Springer-Verlag, 2008.
- [16] M. Milford and G. Wyeth, "Mapping a Suburb with a Single Camera using a Biologically Inspired SLAM System," *IEEE Transactions on Robotics*, vol. 24, pp. 1038-1053, 2008.
- [17] K. Konolige and M. Agrawal, "FrameSLAM: From Bundle Adjustment to Real-Time Visual Mapping," *IEEE Transactions on Robotics*, vol. 24, pp. 1066-1077, 2008.
- [18] M. Milford, "Vision-based place recognition: how low can you go?," *International Journal of Robotics Research*, vol. 32, pp. 766-789, 2013.
- [19] M. Milford and G. Wyeth, "SeqSLAM: Visual Route-Based Navigation for Sunny Summer Days and Stormy Winter Nights," in *IEEE International Conference on Robotics and Automation*, St Paul, United States, 2012.
- [20] Peter Biber and T. Duckett, "Experimental analysis of sample-based maps for long-term SLAM," *The International Journal of Robotics Research* vol. 28, pp. 20-33, 2009.
- [21] J. G. Heys, K. M. MacLeod, C. F. Moss, and M. E. Hasselmo, "Bat and rat neurons differ in theta frequency resonance despite similar coding of space," *Science*, vol. 340, pp. 363-367, 2013.
- [22] I. T. Jolliffe, *Principal Component Analysis*, 2 ed.: Springer, 2002.
- [23] I. Biederman, "Aspects and extension of a theory of human image understanding.," *Computational processes in human vision: An interdisciplinary perspective*, 1988.
- [24] M. C. Potter, "Meaning in visual search," *Science* vol. 187, pp. 965-966, 1975.
- [25] A. Oliva and A. Torralba, "Modeling the shape of the scene: A holistic representation of the spatial envelope," *International journal of computer vision*, vol. 42, pp. 145-175, 2001.
- [26] V. Vapnik, "The support vector method of function estimation," *Nonlinear Modeling*, pp. 55-85, 1998.
- [27] N. Cristianini and S. T. John, *An introduction to support vector machines and other kernel-based learning methods*: Cambridge university press, 2000.
- [28] A. Jacobson, Z. Chen, and M. Milford, "Brain-based Sensor Fusion for Navigating Robots," in *IEEE International Conference on Robotics and Automation*, Karlsruhe, Germany, 2013.
- [29] A. Jacobson and M. Milford, "Towards Brain-based Sensor Fusion for Navigating Robots," in *Proceedings of the 2012 Australasian Conference on Robotics & Automation*, 2012.