2013 IEEE/RSJ International Conference on
Intelligent Robots and Systems (IROS)
November 3-7, 2013. Tokyo, Japan

Autonomous Movement-Driven Place Recognition Calibration for Generic Multi-Sensor Robot Platforms

Adam Jacobson, *Student Member, IEEE*, Zetao Chen and Michael Milford, *Member, IEEE*

Abstract— **In this paper we present a method for autonomously tuning the threshold between learning and recognizing a place in the world, based on both how the rodent brain is thought to process and calibrate multisensory data and the pivoting movement behaviour that rodents perform in doing so. The approach makes no assumptions about the number and type of sensors, the robot platform, or the environment, relying only on the ability of a robot to perform two revolutions on the spot. In addition, it self-assesses the quality of the tuning process in order to identify situations in which tuning may have failed. We demonstrate the autonomous movement-driven threshold tuning on a Pioneer 3DX robot in eight locations spread over an office environment and a building car park, and then evaluate the mapping capability of the system on journeys through these environments. The system is able to pick a place recognition threshold that enables successful environment mapping in six of the eight locations while also autonomously flagging the tuning failure in the remaining two locations. We discuss how the method, in combination with parallel work on autonomous weighting of individual sensors, moves the parameter dependent RatSLAM system significantly closer to sensor, platform and environment agnostic operation.**

I. Introduction

Current robotic systems are capable of performing Simultaneous Localization And Mapping (SLAM) [1] operations in a large array of different environments using a multitude of different techniques [2-6]. However, most of these methods require explicit knowledge of the sensor type, characteristics (such as camera calibration parameters or a laser's field of view) and the robot-relative pose of the sensor. Typical deployment requires a precise calibration step to configure system parameters for each environment and often involves a human "in the loop". No currently available mapping system to date can be deployed "out of the box" onto an unknown robot platform with unknown sensors in an unknown environment.

Attempts to automate sensor calibration [7, 8] have generally required explicit knowledge of the sensors and the interrelations between the system and sensor dynamics. These calibration techniques are also performed in specific locations reducing their ability to be performed as required in the field. Calibration utilising explicit sensor knowledge requires intimate knowledge of the sensor and models of its behaviour. Utilisation of these models is computationally intensive and they do not generally account for potential damage and wear to the robotic platforms over their lifetime.

A. Jacobson, Z. Chen and M. Milford are with the School of Electrical Engineering and Computer Science at the Queensland University of Technology, Brisbane, Australia, michael.milford@qut.edu.au. This work was supported by an Australian Research Council Discovery Project DP1212775 to MM, and an Australian Research Council scholarship to AJ.

In contrast to artificial systems, rodents are capable of rapidly performing sensor calibration from birth within a large range of diverse environments. Rat pups have been seen to demonstrate particular behaviours theorized to help them calibrate their sensory stream from just four days after birth [9, 10]. Adult rats have also been seen to rapidly adapt to changes in the environment or changes in their own sensing abilities during their adult life [11]. Recently the neuroscience community has also shown that it is possible to integrate novel sensory devices into a rat brain and have the rats subsequently learn to utilise this novel input [12].

In this paper, we describe an autonomous place recognition calibration technique that is inspired by how rodents are thought to process and calibrate multisensory data and integrate it with the OpenRatSLAM [13-17] mapping system. Using pivoting movement behaviours, seen in Fig. 1, the calibration technique generates a threshold and confidence score indicating the quality of the threshold. The calibration behaviour compares the minimum difference scores between previously stored sensory snapshots to evaluate a sensory recognition threshold. The technique is fully autonomous and agnostic of sensor type or number and robot-relative sensor pose.

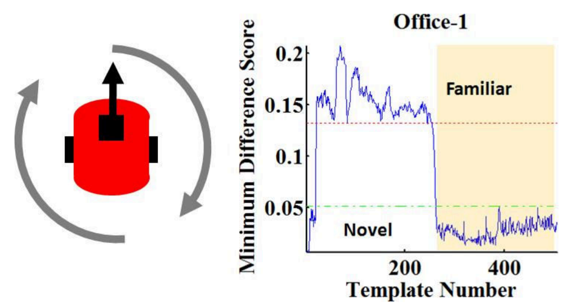

Fig. 1: The calibration technique involves two rotations of a robot, the first providing novel sensory input and the second familiar input. Comparison of the two types of input yields both an estimate of an appropriate place recognition threshold and a quality metric.

We test the calibration system in multiple locations spread over two different environments on a Pioneer 3DX robot, and evaluate its effectiveness in mapping experiments in these environments. Although the robot is a differential drive robot, ongoing work suggests that the calibration behaviours will readily extend to Ackermann steering vehicles performing small loops. The work presented here is intended to complement [18], in which a brain-inspired sensor fusion algorithm was developed to autonomously weight the trustworthiness of different sensor modalities online during navigation. That work and indeed many other research papers [13, 19, 20] have revealed time and time again the critical role of picking an appropriate initial recognition threshold, a problem we address here.

The paper proceeds as follows. In Section 2 we review the literature on robot calibration techniques, along with robot and rat approaches to mapping changing environments. Section 3 presents our approach, briefly reviewing the OpenRatSLAM system, then describing in detail the sensor pre-processing, super template algorithms and the generation of movement behaviours, autonomous threshold calibration and the development of a threshold confidence score. In Section 4 we present the experimental setup, with calibration and resultant mapping results from multiple locations over two different environments presented in Section 5. Finally Section 6 discusses the outcome of the research.

II. Background

In this section, we review robot and rodent research with a focus on autonomous calibration techniques. Robotics research has developed a number of different sensing modalities, such as cameras, inertial measurement units (IMU) and ranging sensors. Utilisation of these sensing modalities within SLAM systems generally requires calibration for particular environment types to function correctly.

Many algorithms have been proposed to calibrate cameras for robot perception. In [21] and [22], 3-D calibration objects are used for the camera calibration which requires that the geometry of calibration objects in 3-D space is known with high precision. An easier method proposed by Zhang [23] calibrates the camera by moving it around a 2-D calibration object to get multiple views of that object. This calibration technique requires no knowledge of the object motion and can estimate the intrinsic parameters of the camera and calculate the relative position between the camera and the calibration objects. Zhang [24] also proposed to used a 1-D calibration object to calibrate the camera [24, 25].

IMU calibration can also be performed, using techniques such as the Factorization method proposed by [26], which requires no prior knowledge of the sensor configuration. Ranging sensors too are calibrated using methods such as that by Fuchs and Hirzinger [27], which simultaneously estimates the intrinsic parameters and the depth distance distortion for a ToF camera.

A number of techniques for multi-sensor calibration have been proposed [28, 29]. Lobo and Dias [28] proposed an algorithm for camera-IMU calibration which independently estimated the relative orientation and translation between the sensors. Their method required a turntable and a pendulum unit. Mirzaei and Roumeliotis [29] also presented a method for camera-IMU calibration in which they used a planar calibration object and tracked its corner points. Those measurements were then fused with IMU data by an extended Kalman filter to estimate the relative pose between sensors and the IMU biases.

The calibration techniques presented here either require explicit knowledge of sensor types or the assistance of calibration objects. These calibration techniques rely on standard and static models of systems in which calibration is being performed.

In contrast, it has been shown that from birth rat pups, all starting with the same sensory equipment, manage to rapidly become adept navigators in an incredible range of environments. This ability to navigate in a diverse range of environments is gained from initial calibration of sensory equipment from birth. As soon as four days after birth, rat pups perform movement behaviours such as pivoting to calibrate their sensory equipment [9]. Adult rats have also been seen to actively perform movement behaviours after their sensory equipment have been altered or sensory cues have been removed [11]. In [12], after the integration of infrared sensors into the touch region of a rat's brain, the rat actively performed exploratory strategies to learn to interpret and use the novel sensory input without the loss of the original touch sensation.

This work aims to replicate the sensory calibration rodents perform through the combination of a pivoting movement behaviour and raw sensory data analysis techniques.

III. Approach

In this section we describe the implementation of our sensor calibration technique through the use of movement behaviours and threshold tuning. We will also provide details of the SLAM backend for this paper along with the sensor pre-processing and the super template learning process which is used for validation of the calibration technique.

A. *OpenRatSLAM*

OpenRatSLAM [16] is an open source version of the RatSLAM system [13-15]. Like RatSLAM, OpenRatSLAM has many tuneable parameters, one of which is the recognition threshold that we set out to autonomously calibrate in this work. OpenRatSLAM has a limited false positive filtering capability, which removes the "no false positives" requirement that is commonly used throughout the literature [2].

OpenRatSLAM consists of three modules – the local view cells, the pose cells and the experience map. Each local view cell encodes a distinct sensory snapshot of the environment, while the pose cell neural network performs filtering of odometry and place recognition information. Finally the experience map provides a graphical, semi-metric map of the environment's layout. More information on RatSLAM and OpenRatSLAM can be found in [13] and [16].

B. *Sensor Pre-Processing*

This system is designed to integrate multiple different sensing modalities, all with different sensing capability, ranges and dimensions. To facilitate the integration of all these sensing modalities a standard method for data storage was developed. The Robot Operating System (ROS) was used as the base framework for integrating these sensing modalities.

Each sensor reading is normalised to a value between 0 and 1 by dividing by the maximum possible sensor value and stored in single line vectors. To enable the integration of forward facing and panoramic imaging systems, the forward facing sensors are resized to a resolution of 12×9 and the panoramic images are resized to 32×9 – the specific resolution is arbitrary and decreasing resolution has been shown to have no significant detrimental effect [30]. Sensors which output images also have their output converted to greyscale and are then separated into single line vectors, seen in Fig. 2.

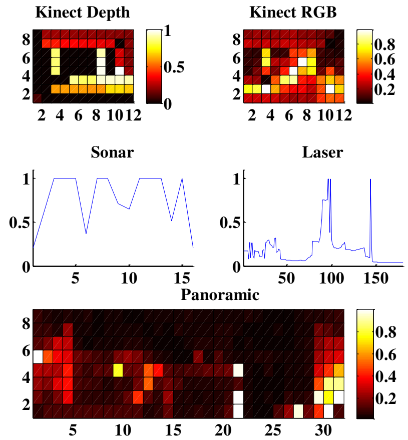

Fig. 2: Reconstructed visualisation of the pre-processed sensory data.

C. Super Template Creation and Evaluation

To allow for robust processing of multisensory inputs, sensory information from single sensors is concatenated into a single line vector called a super template, seen in Fig. 3.

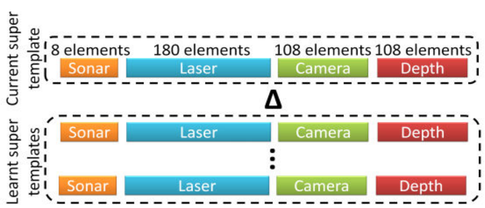

Fig. 3: Super Templates are constructed from concatenation of individual sensor vectors and evaluated using Sum of Absolute differences.

Super template creation occurs when a novel sensory input is detected from analysis of previous stored super templates. Super templates are evaluated utilising a Sum of Absolute Differences (SAD) to determine whether a particular sensory snapshot is novel or familiar. The difference score d between the current proposed super template and a previous super template j is given by:

$$d(j)=\frac{1}{s}\sum_{x=0}^{s-1}\left|p_x-p_x^j\right| \tag{1}$$

where p_x is the value of element x in the current super template, p_x^j is the value of element x in a previous super template j, n is the number of templates learned thus far and s is the length of the template. The best matching template b is the one with the lowest difference score, given by:

$$b=\arg\min_{0\le j<n} d(j) \tag{2}$$

A threshold s_{thresh} determines whether the new template is considered novel ($m=0$) or familiar ($m=1$):

$$m=\begin{cases}1, & d(b)\le s_{thresh}\\ 0, & d(b)>s_{thresh}\end{cases} \tag{3}$$

It is this threshold, s_{thresh}, which is central to the RatSLAM system and which normally requires careful calibration to ensure a minimisation of false template matches. The super template m is then used as the local view cell module in the RatSLAM model.

D. Threshold Tuning

Calibration of the template threshold is important within the RatSLAM system, typically being tuned offline on representative data. The template threshold allows the RatSLAM system to discriminate between novel and familiar sensory environments; a threshold which is too high will create false template matches and potentially distort the map generation, while a threshold which is too low will prevent adequate loop closure. In this work, the primary goal is to minimise false template matches whilst maximising correct template matches, where fewer correct template matches is preferable than many incorrect template matches. This goal is different to that of a system such as FAB-MAP [2], which has lower recognition latencies but is more sensitive to false positives.

To automate the threshold generation, evaluation of novel and familiar templates during the calibration movements was required. Movement behaviours were developed to enable comparison of novel and familiar templates by travelling over the same path twice, the first pass being novel and the second pass being familiar. The algorithm to calculate the recognition threshold is:

$$s_{thresh}=\max_{F_1\le j<F_2} D(j) \tag{4}$$

where D is an array of minimum difference scores for all templates stored during the execution of the movement behaviour. F_1 indicates the start of the familiar region and F_2 indicates the end of the familiar region.

Calculating the maximum difference score for the familiar traverse identifies the maximum difference between two familiar templates. Any difference score below this value can be deemed to be a template match.

E. Confidence Score

The confidence score was created to enable comparison of the quality of different, autonomously generated thresholds for a particular robot sensor configuration. The score, C, is calculated by:

$$C=\frac{MinN-MaxF}{MaxF-MinF} \tag{5}$$

where

$$MaxF=\max_{F_1\le j<F_2} D(j) \tag{6}$$

$$MinF=\min_{F_1\le j<F_2} D(j) \tag{7}$$

$$MinN=\min_{N_1\le j<N_2} D(j) \tag{8}$$

- *MaxF* is the maximum difference score of the familiar region, *MinF* is the minimum difference score of the familiar region and *MinN* is the minimum difference score of the novel region. N_1 and N_2 indicate the start and the end of the novel region respectively.
- *(MinN-MaxF)* is the difference between novel and familiar regions. This number will become negative when the difference scores for the familiar region is greater than the novel difference scores, identifying that

there is no distinct difference between the novel and familiar regions. This also gives an indication of the potential for aliasing.

- *(MaxF-MinF)* identifies the magnitude of potential variance between matching template difference scores. This number also gives evidence on the variance and potential noise sensed by the sensory suite.

The confidence score, *C*, provides a measure of the system's confidence in the generated threshold, with larger values indicating a more trustworthy threshold. When the confidence value drops to 1, it indicates that the variance in the familiar region minimum difference scores is as large as the minimum difference between the familiar and novel region minimum difference scores, suggesting that the calibration procedure has encountered unexpected problems and should probably be repeated.

F. Movement Behaviours

The movement behaviours were developed under the assumption that the robotic platform can safely move within its immediate environment. With this assumption, a rotational movement behaviour was developed. To ensure maximum applicability, the behaviour was developed without using wheel encoder feedback (or any other self-motion information), as it cannot be guaranteed that all robotic platforms have that particular feedback mechanism. Instead, this work utilises raw sensor data to determine whether a particular action is complete.

The rotation behaviour executes two 360° turns. The method utilised to detect a full rotation involves the robotic platform performing sum of absolute differences, rotating until a familiar snapshot is detected utilising a filtered familiarity score. To enable the detection of a familiar sensory snapshot, the difference scores are filtered using:

$$r = k_1 \times r' + (1 - k_1)d(b) \qquad (9)$$

where *r'* is the previous value of the filtered difference score, k_1 is the filter coefficient with a value between 0 and 1 and *d(b)* is the difference score of the minimum difference score described in equation 1 and 2.

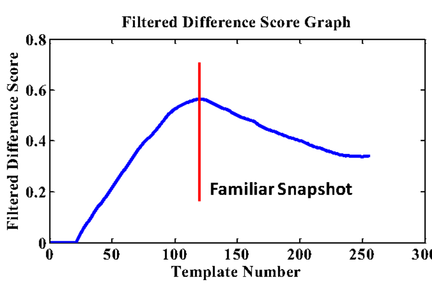

Fig. 4: The rotation behaviour utilises filtered difference scores to detect the completion of a full rotation. Once a negative gradient is detected, a familiar scene is detected and a full rotation is completed.

A familiar template is detected by a negative gradient in the filtered difference scores as seen in Fig. 4. Once the familiar template is detected, this indicates the completion of a full rotation.

This behaviour allows autonomous identification of a full rotation without explicit knowledge of odometry information. The behaviour is performed twice to complete two rotations, the first lap is denoted as novel and the second lap is familiar. Whilst this behaviour is being completed, data is stored for later threshold value computation.

IV. Experimental Setup

In this section we describe the testing environments, robot platform and system parameters. ROS bag-files for both testing environments are available for readers to download and process at the following link: https://wiki.qut.edu.au/display/cyphy/Michael+Milford+Datasets+and+Downloads

A. Robot Platform

The testing platform was an *Adept Mobilerobots* Pioneer 3DX robot (Fig. 5). The standard sensor suite included a SICK laser, a Firewire PointGrey Camera utilising a Catadioptric mirror, and 16 ultrasonic range sensors. To further test the systems utility, sensors were added to simulate deployment on different robotic platforms. For the Office environment, a Microsoft Kinect utilising both RGB and Depth images was incorporated. For the Car Park environment, two sideways facing USB cameras were added.

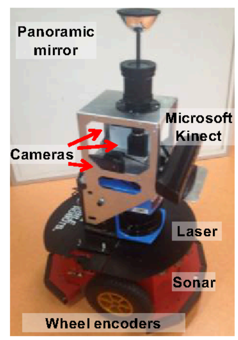

Fig. 5: Pioneer 3DX robot with SICK laser, Microsoft Kinect, sonar ring, panoramic mirror, FireWire camera, two webcameras and wheel encoders.

B. Testing Environments

Testing was performed in two varied settings, an Office environment and a Car Park environment. These locations were chosen to test the system's flexibility because they are two distinctly different environments in which a robot could operate.

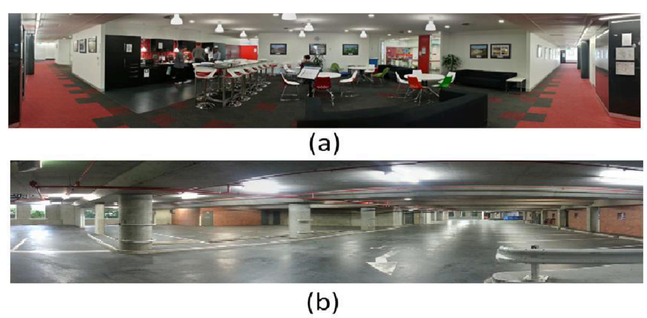

Fig. 6: (a) Office and (b) Car Park testing environments.

The Office environment consists of narrow hallways and cluttered seating areas, shown in Fig. 6a. The Car Park environment is a visually bland, open environment, shown in Fig. 6b. Both test sites are located at the Queensland University of Technology's Gardens Point Campus. The testing procedure for both environments was the same; the robot was placed in a number of locations within the environment, seen in Fig. 7, and the autonomous sensor calibration behaviour was initiated. After calibration, the robot was then tele-operated around the environment in the paths seen in Fig. 8 and Fig. 9, and map generation was performed offline to evaluate the effectiveness of the calculated thresholds. The experimentation locations were not isolated and the public was free to walk through the environments.

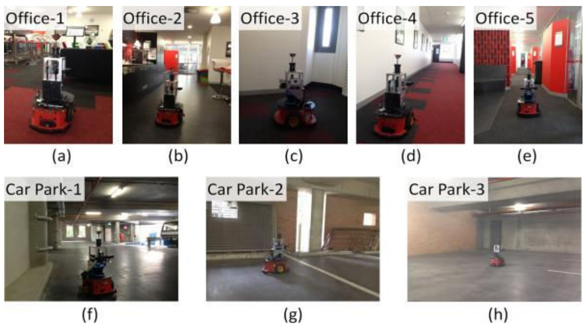

Fig. 7: Visualization of the individual calibration locations. Diverse locations were utilised to test the flexibility of the system.

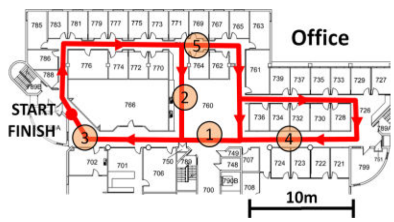

Fig. 8: Calibration locations and robot paths for the Office environment.

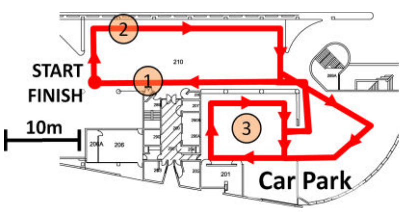

Fig. 9: Calibration locations and robot paths for the Car Park environment.

V. Results

In this section we present results for all calibration locations in both environments. We show four key sets of results – difference scores for each calibration location, generated thresholds, confidence scores for each threshold and resultant OpenRatSLAM experience maps. We also present outcomes of manual threshold tuning to provide "before and after" comparison. The accompanying video illustrates the method, experiments and results.

A. Manually Tuned Experimentation

To provide a baseline performance indicator, we manually selected a "small" and "large" threshold and ran the mapping experiment. Fig. 10a demonstrates the template graph and experience maps which have been generated with a manually selected threshold of 0.08. A (too) high recall rate results in false positives and a corrupted experience map. Fig. 10b show the template graph and experience maps with a threshold value of 0.01, resulting in insufficient loop closures and a badly distorted map due to poor odometry data.

Thresholds were then manually tuned to determine the range within which OpenRatSLAM is able to generate topologically correct maps of the Office and Car Park environments. For the Office environment the thresholds needed to be between 0.04 and 0.07 before OpenRatSLAM would operate correctly. For the Car Park environment appropriate threshold values ranged between 0.027 and 0.07. The wide range of acceptable thresholds for these two environments was promising in that it suggested that the autonomous tuning procedure presented here will in general have a large target to hit.

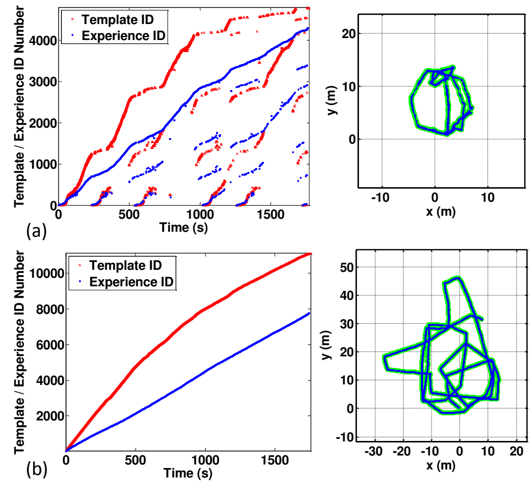

Fig. 10: Template graph and experience maps for the Office environment with threshold of (a) 0.08 and (b) 0.01. Threshold of 0.08 results in false positive matches and consequent distortion in the experience map. The 0.01 theshold results in a disjointed map with no loop closures.

B. Autonomous Threshold Calibration

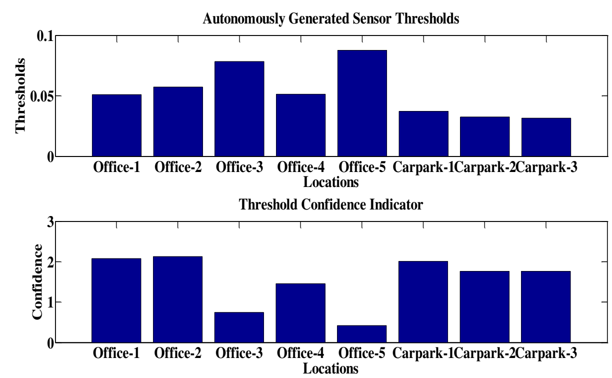

Fig. 11: Graph of autonomously generated threshold and confidence scores. Confidence scores outline the reliabilty of the autonomously generated threshold. Confidence scores below one indicate a likely failed calibration attempt.

Fig. 11 shows the autonomously generated thresholds and confidences scores. It can be seen that for the Car Park environment, the generated thresholds and their respective confidence scores are relatively constant for all three calibration locations. The Office environment displays a larger amount of fluctuation between threshold values.

The thresholds for Office-3 and Office-5 are significantly larger than the other Office thresholds. It can also be seen that the confidence scores for Office-3 and Office-5 are

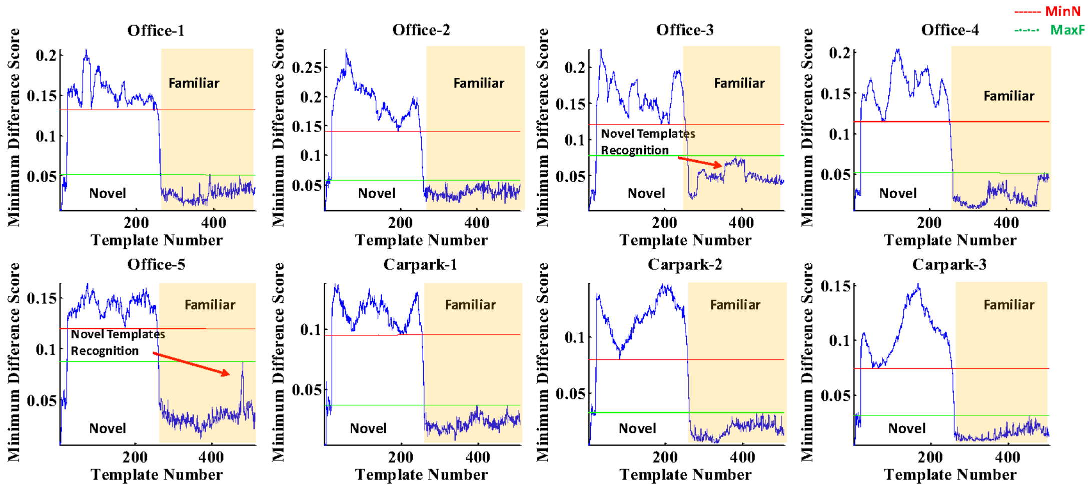

Fig. 12: Difference scores for each calibration location showing the distinction between novel and familiar sensory snapshots. The novel template recognition seen in difference graphs for Office-3 and Office-5 caused the generation of a threshold deemed unreliable by the confidence score. MinN and MaxF are utilised by the confidence score calculation.

considerably smaller than the other Office samples indicating the thresholds generated from these samples are likely unreliable. Analysis of the difference scores for each calibration location, shown in Fig. 12, identifies clear and noticeable spikes in the difference score readings during the second, familiar traverse of the environment. These spikes indicate that there were novel sensory snapshots being detected during the familiar region of the data acquisition.

Analysis of the data has shown that during the familiar (second) rotation, novel information was learnt. In these two cases this novel information source was identified as people walking past the robot, which created novel sensory snapshots and caused the large variance in difference scores. In a normal usage scenario, the operator could be notified of the poor confidence score and instructed to repeat the calibration, possibly at a different location.

C. *Experience Maps*

The generation of experience maps was performed to evaluate the effectiveness of the threshold calibration techniques for mapping. Fig. 13 shows the OpenRatSLAM maps generated with the autonomously generated thresholds. It can be seen that in all cases, except Office-3 and Office 5, the autonomously generated thresholds produce maps which are topologically correct representations of the environments.

The experience maps for Office-3 and Office-5 are corrupted due to a number of incorrect loop closures caused by large thresholds. This was expected due to the low confidence scores. It can also be seen that the Car Park experience maps in Fig. 13 are slightly distorted when compared to the original floor plan seen in Fig. 9. This distortion is due to the wheel slip and encoder errors, not false positive or negative template matches. The Car Park experience maps in Fig. 13, although not metrically correct, are topologically correct and are likely to still be useful for accurate robot navigation. Work presented in [14] shows that topological maps with similar distortion were quite effective as the basis for long term robot navigation in a similar type of environment.

D. *Compute*

All experiments were run at real-time speed on a dual-core Intel i7 Ivy-Bridge processor. Compute for this system scales with time due analysis of super templates and the CPU usage on all four hyper-threads was never above 40% per hyper-thread for all experiment.

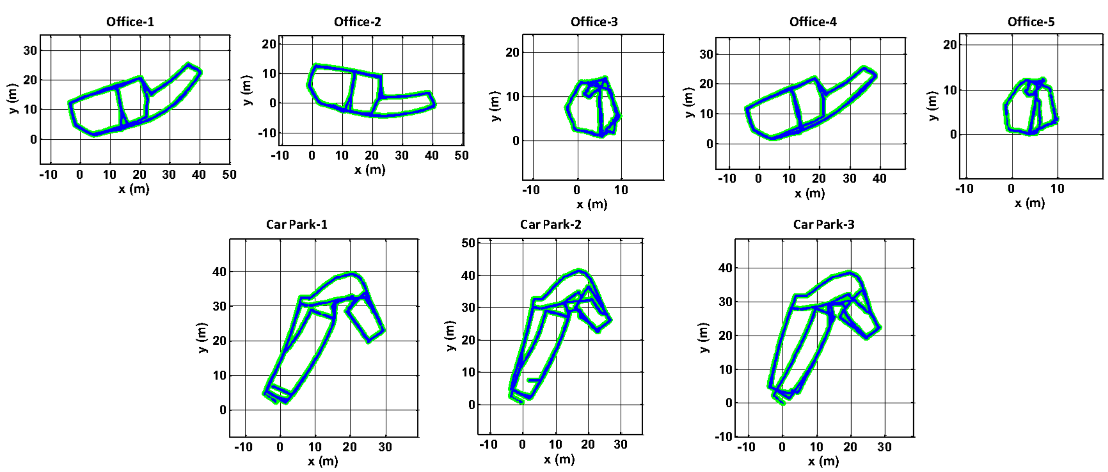

Fig. 13: OpenRatSLAM experience maps generated witht the autonmous threhsolds. All maps are topologically correct, except for maps generated from the Office-3 and Office-5 location thresholds.

VI. Discussion and Future Work

The results presented in this paper are a promising step towards our overall goal of an autonomous, movement driven calibration scheme enabling robot mapping in unknown environments with unknown sensor configurations. This is of course a challenging aim and we have attempted to address possible failure modes through the development of a confidence score that can be used to forecast the reliability of autonomously calibrated thresholds. The confidence score metric correctly identified the decreased reliability of two calibration thresholds resulting from dynamic environments which were unsuitable for use in mapping.

The calibration method does not currently provide any ability to conduct further recognition threshold tuning after the initial calibration step. Future work will look to filter out errors caused by dynamic environments and will look to combine these initial calibrations with an ongoing threshold tuning scheme to account for situations where the robot may travel from a feature rich environment to a visually bland environment.

Future work will also integrate the autonomous threshold calibration techniques described here with the brain-inspired sensor fusion algorithm developed to autonomously and dynamically weight the utility of different sensor modalities in map creation, seen in [18]. The method provided here ensures that an appropriate recognition threshold is calculated shortly after a robot is turned on, while the work in [18] builds on (and relies upon) a correct initial threshold to then dynamically weight the trustworthiness of sensors during online operation. Integrating the two systems, we will expand our experiments to a wider range of platforms (including Ackermann steering vehicles), sensor suites and environments. The development of an alternative, "donut" movement behaviour is currently under development for use by Ackermann steering vehicles, which will utilise the same method to detect the completion of a 360° turn as described above.

References

[1] G. Dissanayake, H. Durrant-Whyte, and T. Bailey, "A computationally efficient solution to the simultaneous localisation and map building (SLAM) problem," in *IEEE International Conference on Robotics and Automation*, San Francisco, USA, 2000, pp. 1009-1014.

[2] M. Cummins and P. Newman, "Highly scalable appearance-only SLAM - FAB-MAP 2.0," in *Robotics: Science and Systems*, Seattle, United States, 2009.

[3] 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.

[4] H. Andreasson, T. Duckett, and A. Lilienthal, "A Minimalistic Approach to Appearance-Based Visual SLAM," *IEEE Transactions on Robotics,* vol. 24, pp. 1-11, 2008.

[5] L. M. Paz, P. Pinies, J. D. Tardos, and J. Neira, "Large-Scale 6-DOF SLAM With Stereo-in-Hand," *IEEE Transactions on Robotics,* vol. 24, pp. 946-957, 2008.

[6] K. Konolige and M. Agrawal, "FrameSLAM: From Bundle Adjustment to Real-Time Visual Mapping," *IEEE Transactions on Robotics,* vol. 24, pp. 1066-1077, 2008.

[7] D. Caltabiano, G. Muscato, and F. Russo, "Localization and self-calibration of a robot for volcano exploration," in *Robotics and Automation, 2004. Proceedings. ICRA '04. 2004 IEEE International Conference on*, 2004, pp. 586-591 Vol.1.

[8] N. Roy and S. Thrun, "Online Self-Calibration For Mobile Robots," in *IEEE International Conference on Robotics and Automation*, Detroit, MI, 1999, pp. pp. 2292-2297.

[9] I. Golani, G. Bronchti, D. Moualem, and P. Teitelbaum, ""Warm-up" along dimensions of movement in the ontogeny of exploration in rats and other infant mammals," *Proceedings of the National Academy of Sciences of the United States of America,* vol. 78, pp. 7226-7229, 1981.

[10] P. Stratton, M. Milford, G. Wyeth, and J. Wiles, "Using Strategic Movement to Calibrate a Neural Compass: A Spiking Network for Tracking Head Direction in Rats and Robots," *PLoS One,* vol. 6, 2011.

[11] A. Cheung, D. Ball, M. Milford, G. Wyeth, and J. Wiles, "Maintaining a Cognitive Map in Darkness: The Need to Fuse Boundary Knowledge with Path Integration," *PLoS Computational Biology,* vol. 8, 2012.

[12] E. E. Thomson, R. Carra, and M. A. Nicolelis, "Perceiving invisible light through a somatosensory cortical prosthesis," *Nat Commun,* vol. 4, p. 1482, Feb 12 2013.

[13] 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.

[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] D. Ball, S. Heath, M. Milford, and G. Wyeth. (2012). *OpenRatSLAM*. Available: http://code.google.com/p/ratslam/

[17] 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.

[18] M. Milford and A. Jacobson, "Brain-inspired Sensor Fusion for Navigating Robots, *International Conference on Robotics and Automation*, Karlsruhe, Germany, 2013.

[19] M. Milford, F. Schill, P. Corke, R. Mahony, and G. Wyeth, "Aerial SLAM with a Single Camera Using Visual Expectation," in *International Conference on Robotics and Automation*, Shanghai, China, 2011.

[20] M. Milford and G. Wyeth, "Single Camera Vision-Only SLAM on a Suburban Road Network," in *International Conference on Robotics and Automation*, Pasadena, United States, 2008.

[21] R. Tsai, "An efficient and accurate camera calibration technique for 3D machine vision," in *Proceedings of the CVPR*, Miami Beach, USA, 1986, pp. 364–374.

[22] J.Y.Weng, P.Cohen, and M.Herniou, "Camera calibration with distortion model and accuracy evaluation," *IEEE Trans. Pattern Anal. Mach. Intell,* vol. 14, pp. 965–980, 1992.

[23] Z.Y.Zhang, "A flexible new technique for camera calibration," *IEEE Trans. on Pattern Analysis and Machine Intelligence,* vol. 22, pp. 1330–1334, 2000.

[24] Z. Y. Zhang, "Camera calibration with one-dimensional objects," *IEEE Trans. Pattern Anal. Mach. Intell,* vol. 26, pp. 892–899, 2004.

[25] F. C. Wu, Z. Y. Hu, and H. J. Zhu, "Camera calibration with moving onedimensional objects," *Pattern Recognition,* vol. 38, pp. 755–765, 2005.

[26] M. Hwangbo and T. Kanade, "Factorization-based calibration method for MEMS inertial measurement unit," in *In IEEE Int.Conf. on Robotics and Automation*, 2008.

[27] S. Fuchs and G. Hirzinger, "Extrinsic and depth calibration of ToF-cameras," in *IEEE Conf. on Computer Vision and Pattern Recognition*, 2008.

[28] J. Lobo and J. Dias, "Relative pose calibration between visual and inertial sensors," *Int. Journal of Robotics Research,* vol. 26, pp. 561–575, 2007.

[29] F. Mirzaei and S. Roumeliotis, "A Kalman filter-based algorithm for IMU-camera calibration: Observability analysis and performance evaluation," *IEEE Trans. on Robotics,* vol. 24, pp. 1143–1156, 2008.

[30] M. Milford, "Visual Route Recognition with a Handful of Bits," in *Robotics: Science and Systems VIII*, Sydney, Australia, 2012.