



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:

Yu, Litao, Jacobson, Adam, & Milford, Michael
(2018)

Rhythmic representations: Learning periodic patterns for scalable place recognition at a sublinear storage cost.

IEEE Robotics and Automation Letters, 3(2), pp. 811-818.

This file was downloaded from: <https://eprints.qut.edu.au/150933/>

© 2018 IEEE

© 2018 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.

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/LRA.2018.2792144>

Rhythmic Representations: Learning Periodic Patterns for Scalable Place Recognition at a Sub-Linear Storage Cost

Litao Yu, Adam Jacobson and Michael Milford

Abstract—Robotic and animal mapping systems share many challenges and characteristics: they must function in a wide variety of environmental conditions, enable the robot or animal to navigate effectively to find food or shelter, and be computationally tractable from both a speed and storage perspective. With regards to map storage, the mammalian brain appears to take a diametrically opposed approach to all current robotic mapping systems. Where robotic mapping systems attempt to solve the data association problem to minimise representational aliasing, neurons in the brain intentionally break data association by encoding large (potentially unlimited) numbers of places with a single neuron. In this paper, we propose a novel method based on supervised learning techniques that seeks out regularly repeating visual patterns in the environment with mutually complementary co-prime frequencies, and an encoding scheme that enables storage requirements to grow sub-linearly with the size of the environment being mapped. To improve robustness in challenging real-world environments while maintaining storage growth sub-linearity, we incorporate both multi-exemplar learning and data augmentation techniques. Using large benchmark robotic mapping datasets, we demonstrate the combined system achieving high-performance place recognition with sub-linear storage requirements, and characterize the performance-storage growth trade-off curve. The work serves as the first robotic mapping system with sub-linear storage scaling properties, as well as the first large-scale demonstration in real-world environments of one of the proposed memory benefits of these neurons.

I. INTRODUCTION

Visual place recognition - recognising whether a current camera image matches to those stored in a map or database - is a fundamental component of most robotic mapping and navigation systems[1]. These mapping systems are typically developed and evaluated based on the quality of the map they can produce, the robustness of representations and their associated computational requirements. Much emphasis has been placed on solving the “data association” problem - making sure that there are no incorrectly or aliased map-landmark associations.

Navigation neurons found in the brain of many mammals such as rodents, known as “grid cells” [2] (see Fig. 1), have highly aliased data associations with locations in the environment - each cell encodes an arbitrary number of physical locations laid out in a triangular tessellating grid [3], [4]. There has been much interest in the theoretical advantages of such a neural representation including implications for memory storage, error correction [5] and scalability that could revolutionize how artificial systems including robots are developed.

In this paper, we propose a novel approach to discover regularly repeating visual patterns in an environment, and to




Fig. 1: Neurons in the mammalian brain known as grid cells intentionally alias their representation of the environment; each neuron represents an arbitrary number of places in a regularly repeating pattern.

encode these regularly repeated patterns in frame sequences (see Fig. 1). We adopt a supervised learning approach and take advantage of statistical properties to identify periodicity in the world being mapped. To perform place recognition, a global location estimate is reconstructed from identifying the phase of these learned patterns in a current camera image. In this way, the storage requirements scale up sub-linearly with the number of encoded places in the environment (or better). Extensive experiments on three real-world datasets demonstrate successful place recognition while retaining sub-linear storage growth.

We present new research that significantly extends a pilot study [6] by developing a number of new contributions that address scalability to large, challenging real world environments, including:

- A method for actively finding and learning periodic visual patterns from frame sequences using supervised learning techniques and best practice for maximizing their utility in sub-linear mapping,
- Techniques for optimizing for minimum storage requirements that balance the number of periodic patterns, their period lengths and their separability,
- Development of a multi-exemplar training scheme that improves place recognition performance in perceptually challenging environments where multiple training examples are available, while maintaining sub-linear storage growth,
- Visual data augmentation techniques for improving performance when multiple-examples are not available, and
- Comprehensive performance evaluation on several large benchmark datasets including characterizations of the tradeoff between storage scalability and place recognition performance, and analysis of the benefits of multi-exemplar and augmentation-based training.

Together these contributions represent a significant step

towards enabling a sub-linear, highly scalable map encoding scheme for autonomous systems, and provide for the first time a real-world data-based test of one of the primary postulated memory benefits of this universal spatial encoding scheme found in the mammalian brain.

The paper proceeds as follows. Section II provides an overview of data compression in signal processing with relevance to the approach presented here. Section III describes the components of our proposed approach in detail. Experimental results and analysis are presented in section IV, with Section V discussing the findings and future areas of research.

II. BACKGROUND

Data compression has a broad range of applications in signal processing, which is to encode data into compact representations by taking advantage of perceptual and statistical properties of data to provide superior analytical results. In image processing, we can use cosine transform to compress a BitMap (BMP) image as a JPEG format with a tolerable information loss but a much smaller data size. In computer vision, images and videos are usually represented as high-dimensional visual feature vectors. The goal of encoding images into compact codes is to simultaneously reduce the storage cost and accelerate the computation. To achieve this, the most discriminant information contained in high-dimensional data is usually embedded into a lower-dimensional space for further analysis. Usually, the embeddings are in a discrete format such as hashing [7]. However, the discrete data representations suffer from data collisions when data size is large, so it is not the best option for unique mapping in visual place recognition. To avoid data collision, visual information can also be embedded in continuous, rather than discrete lower-dimensional spaces [8].

For multimedia data compression, there are two encoding families based on machine learning techniques: binary embedding and vector quantization, both of which are designed to compress continuous sensor data into discrete feature spaces. The idea of binary embedding is to represent feature vectors as compact binary codes, so the Euclidean distance between two vectors could be approximated by Hamming distance in the binary space [7]. The advantage of binary embedding is due to the efficient Hamming distance computation, which can be implemented by the XOR and POPCOUNT operations. Different from binary embedding, vector quantization (VQ) adopts a codebook as a dictionary to quantize the feature vectors into a set of codewords, and the distances between any two codewords are pre-computed and stored in a lookup table [9]. When the original feature space is decomposed into the Cartesian product of several low-dimensional subspaces, vector quantization becomes product quantization (PQ) [10], [11], [12]. Compared with binary embedding, PQ based encoding methods have a lower information loss and thus can achieve a better accuracy, at the cost of a slightly lower computational speed. Both binary embedding and vector quantization are effective encoding techniques with regards to calculation and storage costs.




Fig. 2: An illustrative scenario with only two visual patterns available: buildings and trees, and both of them appear periodically. Each column represents a frame, so the frame sequence simulates a virtual camera moving forward. The combination of the two landmarks can represent at most $3 \times 4 = 12$ distinct locations.

Although binary embedding and vector quantization have different encoding strategies, the underlying mechanism is clustering, i.e., similar statistical patterns have the same codes. Both of the two techniques suffer from code collisions as the data size increases, and they have linear storage growth with the number of data instances.

Computational and storage requirements are of particular importance for mobile robotic and autonomous systems. It is important to differentiate between at least three different goals - achieving highly compact but ultimately linear storage growth, achieving sub-linear computational requirements (but with linear or worse storage growth), and achieving sub-linear storage growth. For the first goal, various techniques have been applied in robotic applications. For example, Locality Sensitive Hashing (LSH) is used to deal with the problem of stereo correspondence estimation [13], multi-model fusion techniques are adopted in humanoid robots to process large-scale language data [14], and Local Difference Binary (LDB) descriptor is applied to obtain a robust global image description for place recognition and loop closure detection [15]. In [16] the authors proposed to compress sensory data from tactile skins. Similarly, the distributed sensor data with high-frequencies can be compressed as coresets for streaming motion [17]. To achieve the sub-linear computation, [18] builds an index on the original database to reduce the computation cost. While significant effort has been devoted towards efficient computation with low absolute storage costs, there has been little work examining how sub-linear storage growth might be achieved for SLAM systems, or examining how natural systems achieve this sub-linear storage growth with a unique one-to-many neuronal mapping system.

III. APPROACH

In this section, we describe our proposed encoding model for scalable place recognition, based on supervised learning techniques. The system comprises a periodic template learning phase, a database encoding phase, and a global place reconstruction phase.

A. Learning Periodic Patterns from Frame Sequences

For clarity, we start with a toy example, since in the actual system we are looking for the underlying visual patterns that are often not intuitive. In Fig. 2, we show an illustrative scenario with only two visual patterns: different buildings and different trees. In this example, each column represents a frame in a sequence. By observing the frame sequence, we can see the style of building cycles in every 3 frames, and the type of the tree regularly changes in every 4 frames, respectively. Consequently, the combination of the two ideal periodic patterns can uniquely represent at most $3 \times 4 = 12$ different locations. In place recognition systems, if there are more than two periodic patterns with different lengths, only storing the pattern information is sufficient for global place estimation. For example, we can use periodic landmarks and their positions to describe a frame, just like the visual templates used to detect interesting points as SIFT descriptors in image processing, and these descriptors can be then aggregated into visual feature vectors for further analysis. However, we will show that it is possible to learn latent periodic patterns from a wide variety of data. To enable the SLAM system to automatically analyse the feature vectors, in [6] the authors proposed to apply spectrogram to find the regularly repeated patterns in frame sequences. The spectral methods assume the signal is composed of phase-shifted sine and cosine curves with scaled factors and offsets, but such a strong assumption is not valid in most real applications. Also, the thresholds need to be neatly set for signal discretization and template matching, because a high threshold may lead to the loss of matched templates, while a low threshold usually results in mismatches or redundancies.

In our proposed data compression model, temporally periodic patterns are learned from the frame sequence in a database. Given an (integer) period τ , our system will look for a linear separation for each possible (integer) phase of that period that separates frames in that phase from frames in all other phases. This allows us to train completely distinct classifiers for each phase of a period, which is much less restrictive than training a single template for all phases of that period, as is implicit in spectral methods.

Let a location database be represented as a frame sequence as $\mathbf{X} = \{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_N\}$ where N is the total number of data instances, and \mathbf{x}_i is the i -th frame in the database ($1 \leq i \leq N$). If each frame is represented as a d -dimensional visual vector, i.e., $\mathbf{x}_i \in \mathbb{R}^d$, the size of the database is $\mathbb{R}^{N \times d}$. When there is a cyclic visual pattern with the length τ , τ templates are generated in each period. For the j -th template within a period ($1 \leq j \leq \tau$), we assign a binary label $y_i^{(j)} \in \{1, -1\}$ for each frame \mathbf{x}_i to indicate if it can match the j -th template:

$$y_i^{(j)} = \begin{cases} 1 & i \bmod \tau = j - 1, \\ -1 & \text{otherwise.} \end{cases} \quad (1)$$

Consequently, the task of determining whether a frame \mathbf{x}_i can match the j -th template in a period becomes a binary classification problem. The weight vector \mathbf{w}_j and bias b_j can

be simply obtained by solving a linear SVM as follows:

$$\begin{aligned} \min_{\mathbf{w}_j, b_j, \xi_{ij}} \quad & \frac{1}{2} \|\mathbf{w}_j\|^2 + C \sum_{i=1}^N \xi_{ij}, \\ \text{s.t.} \quad & y_i^{(j)} (\mathbf{w}_j^\top \mathbf{x}_i + b_j) \geq 1 - \xi_{ij}, \\ & \xi_{ij} \geq 0, 1 \leq i \leq N, \end{aligned} \quad (2)$$

where C is the penalty parameter to balance the hinge loss and functional margin. Here we mainly focus on the loss function to make the templates as linearly separable as possible, and we can simply set $C = \log N$.

Simultaneously considering τ templates within a period, these binary classifiers can be integrated into a multi-class SVM model:

$$\begin{aligned} \min_{\mathbf{w}_j, b_j, \xi_{ij}} \quad & \frac{1}{2} \sum_{j=1}^{\tau} \|\mathbf{w}_j\|^2 + C \sum_{j=1}^{\tau} \sum_{i=1}^N \xi_{ij}, \\ \text{s.t.} \quad & y_i^{(j)} (\mathbf{w}_j^\top \mathbf{x}_i + b_j) \geq 1 - \xi_{ij}, \\ & \xi_{ij} \geq 0, 1 \leq i \leq N, 1 \leq j \leq \tau. \end{aligned} \quad (3)$$

This multi-class SVM can be efficiently computed by some toolboxes such as scikit-learn¹. The statistical property of the weight vector \mathbf{w}_j is straightforward: when all sequenced frames within the database are represented by $\lfloor N/\tau \rfloor$ periods and can be perfectly segmented by τ classifiers, each classifier has the minimum covariance with its positively classified data instances. Thus, these classifiers can be just considered as the templates within a period.

The optimised weight vector \mathbf{w}_j and bias b_j ($1 \leq j \leq \tau$) are able to determine if a frame is at the j -th position within a period with length τ , i.e., the template that can be matched by a frame \mathbf{x} is calculated by:

$$f(\mathbf{x}|\tau) = \arg \max_j (\mathbf{w}_j^\top \mathbf{x} + b_j), \quad 1 \leq j \leq \tau. \quad (4)$$

Note that in order to keep the sub-linearity for data compression, we do not use kernel SVMs. In the kernel case, the weight vector in the decision function $f(\cdot)$ is represented as a linear combination of support vectors. Although the kernel decision function is more discriminative than the linear one, it cannot achieve the sub-linear data compression, because when either the dimension or the size of database increases, the number of support vectors also increases accordingly.

B. Database Encoding

As we have seen, we can use linear SVMs to learn τ periodic templates given a period τ and the N frames of a dataset in which we wish to localise. However, unless $\tau > N$, this is obviously not sufficient to uniquely identify a frame. The core idea of our method is to learn two or more such cyclic patterns $\{\tau_1, \tau_2, \dots\}$, such that frames can be uniquely identified. In this subsection, we show how the periods τ can be chosen to allow for unique identification while minimising storage requirements.

Assuming there are several candidates for τ available, we simply select r cyclic patterns with periods $\tau_1, \tau_2, \dots, \tau_r$

¹<http://scikit-learn.org/>

to estimate a frame position within a database when given an arbitrary frame \mathbf{x} from it. The position of \mathbf{x} could be determined by the *phase matches*, which are represented as a candidate set $\{j_1, j_2, \dots, j_r\}$, $1 \leq j_k \leq \tau_k$. The possible index i is calculated by:

$$i = a_k \cdot \tau_k + j_k, \quad (5)$$

where $k \in \{1, \dots, r\}$, and a_k is a natural number. To identify \mathbf{x}_i with $\{j_1, j_2, \dots, j_r\}$, its index i needs to be the unique solution of Eq. (5). Thus the selections of $\tau_1, \tau_2, \dots, \tau_r$ should satisfy:

$$lcm(\tau_1, \tau_2, \dots, \tau_r) \geq N, \quad (6)$$

where $lcm(\cdot)$ is the least common multiple operator. This condition guarantees the index mapping is unique, and there are sufficient “slots” to store all frames in the database.

If we manage to make $\tau_1, \tau_2, \dots, \tau_r$ co-prime, then Eq. (6) will be equivalent to $\prod_{k=1}^r \tau_k \geq N$. Given this constraint and if $r = 2$, when given a τ_1 , τ_2 needs to be at least N/τ_1 . In Fig. 3, we illustrate how the selection of the period lengths affects the total number of templates if $N = 100$ and $r = 2$. For each τ_1 on the x-axis, we found the smallest τ_2 that is both co-prime with τ_1 and that satisfies $\tau_1 \times \tau_2 \geq N$. Then on the y-axis we report $\tau_1 + \tau_2$, which is proportional to the storage cost. We can see the minimum storage for 100 place estimations is 21 when $r = 2$, so $\tau_1 = 10$ and $\tau_2 = 11$ would be the best period pair.




Fig. 3: The minimum number of templates required when $N = 100$ and $r = 2$.

Assuming that our place recognition algorithm could correctly identify the phase matches j_k (see Eq. (5)) corresponding to a query \mathbf{x} , the memory requirements for our system are to store the weight vectors $\mathbf{w}_1^{(k)}, \mathbf{w}_2^{(k)}, \dots, \mathbf{w}_j^{(k)}$ and biases $b_1^{(k)}, b_2^{(k)}, \dots, b_j^{(k)}$ for each possible k . In other words, we need to allocate memory for $\sum_{k=1}^r \tau_k$ vectors of size $d + 1$. Thus, the minimal storage requirements are achieved when $\prod_{k=1}^r \tau_k \geq N$. In unconstrained cases, the solution is $\sqrt[r]{N}$, but $\tau_1, \tau_2, \dots, \tau_r$ also need to satisfy they are coprime integers. Since we have not found a closed-form solution to this constrained problem, we instead propose to sample candidates from $\lceil \sqrt[r]{N} \rceil, \lceil \sqrt[r]{N} \rceil + 1, \dots, \lceil \sqrt[r]{N} \rceil + m$ with the least training errors. The training error e is the fraction

of misclassified training samples of τ linear models among the whole training set. Therefore, only storing r groups of periodic templates can reduce the space complexity from $O(N)$ to $O(r\sqrt[r]{N})$.

C. Reconstructing a Global Place Estimate

The localisation of an arbitrary frame from the database \mathbf{x} is implemented by the intersection operation. We first calculate the phase matches of r periodic patterns: $j_k = f(\mathbf{x}|\tau_k)$ for $k \in \{1, 2, \dots, r\}$ by applying Eq. (4), and then generate r candidate sets P_1, P_2, \dots, P_r , where $P_k = \{f(\mathbf{x}|\tau_k), f(\mathbf{x}|\tau_k) + \tau_k, \dots, f(\mathbf{x}|\tau_k) + \lfloor N/\tau_k \rfloor\}$. The index of \mathbf{x} in the original frame database is calculated by $f(\mathbf{x}|\tau_1, \tau_2, \dots, \tau_r) = \bigcap_{l=1}^r P_l$.

Before the online searching phase, given a query image, the system should first determine if the location that the image represents can be found in the database. Some appearance-based SLAM systems such as [19], proposed to set a lower bound of likelihood, which is calculated in the training procedure and set by users. The low likelihood that falls below the lower bound means the query image cannot match any places in the database. Our proposal is a deterministic approach, and there would be at least one phase match when given a query even if it is actually an outlier. So a lower bound of decision value can also be set to determine if a query can match a template in a periodic pattern. Alternatively, an auxiliary classifier could be trained when there are negative exemplars available, which are not descriptive to any locations in the database.

The retrieval of our method consists of $r - 1$ 1d intersections of r sets with size N/τ , which is lower than $\sqrt[r]{N}$. Since the sets are already sorted, the 1d intersection can be achieved in $O(r\sqrt[r]{N})$: the intersection of a pair of sorted sets is linear in the sum of the sizes of both sets, so its time complexity is $O(\sqrt[r]{N})$, and we only need to perform $r - 1$ such intersections on all sets.

D. Improving Robustness Through Multi-Exemplar Training and Augmentation

A common method for improving the robustness of recognition techniques is to use, where available, multiple different examples of an object or place in training. For example, ImageNet has over ten million images with one thousand categories [20]. With large-scale and well-labelled training images, the classifiers trained on such datasets have near-human performance in very challenging recognition tasks. For the mobile place localisation systems, the robot should be able to memorise the scenes by revisiting the same places multiple times from different perspectives, or under distinct appearance conditions, to improve their discriminative power. In this case, the periodic patterns are essentially learned in a shared space rather than the original feature space.

Since our proposed data compression model for scalable place recognition is also based on supervised learning techniques, using frames taken under different appearance conditions has the potential to improve recognition accuracy

and robustness. However, there is not always multi-exemplar training data available. In this case, we can apply image augmentation methods such as Gaussian blur, flipping, random cropping and elastic transformation to simulate the multi-exemplar data environment.

IV. EXPERIMENT AND ANALYSIS

In this section, we describe the datasets used, the image preprocessing methods, the evaluation metrics and the experimental results. We also provide analysis of our proposed model, breaking down the performance contributions of the core system, enhancements including multi-exemplar training and augmentation, and provide an analysis of the trade-off between performance and storage scaling.

A. Datasets and Experiment Settings

To evaluate our proposed data compression model for scalable visual place recognition, we experimented with three different datasets: Nordland Train dataset, Aerial Brisbane dataset and Oxford RobotCar dataset.

1) *Nordland Train dataset*: The Nordland Line² is a 729-kilometre railway between Trondheim and Bodø, Norway. This dataset contains four long videos captured by placing a camera at the front of a train facing forward along the railway track. The four videos describe the front views in four seasons, and each video is about ten hours long. This dataset was first used for visual navigation across seasons[21], but we only use it to test the compressibility of our proposed model. In the first part of our evaluation, the queries are all from the reference data, and the model aims to find the exact positions of them in the database. To pre-process the video data, we first extracted the keyframes, then used the optical flow of the ground directly in front of the train to estimate the velocity then normalised it. Finally, the four subsets contain 10,713, 7,403, 9,267 and 7,276 frames, respectively.

2) *Aerial Brisbane dataset*: The Aerial Brisbane dataset is generated by taking a snapshot from NearMap³, which describes the Brisbane region in Queensland, Australia. The total size of the image is 7526×6562 , and each pixel is an actual geographic area of 4.777×4.777 square metres. The image was then segmented to 224×224 -pixel frames with 112-pixel strides, so in our setting the dataset can represent 3,705 different places.

We use this dataset to test if our model can recognise locations in visual changing environment. To simulate the environment, the query images and reference data are from different sources. We collected several snapshots of the aerial map taken at different times ranging from 19/05/2013 to 24/06/2017. One image was selected as the query to search the absolute locations of the patches on the map, and the rest reference frames are used to train the model.

²<https://nrkbeta.no/2013/01/15/nordlandsbanen-minute-by-minute-season-by-season/>

³<http://maps.au.nearmap.com/>

3) *Oxford RobotCar dataset*: The Oxford RobotCar dataset [22] contains over 100 repetitions of a consistent route through Oxford, UK. The dataset captures many different combinations of weather, traffic and pedestrians. For our testing, we used 5 subsets of the Oxford RobotCar dataset generated from a fixed route captured at different times of day using images captured by the Point Grey Bumblebee XB3. This dataset describes a small area with a limited number of locations. Since the geographic positions are described in consecutive northing and easting values as GPS data, we applied KMeans on the normalised coordinates to generate 100 clusters, so the GPS coordinates falling into the same cluster are considered as a unique location on the map. We ran our place recognition model to test if the periodic encoding can successfully capture the cyclic properties of the frame sequences to enable accurate reconstruct the location estimation.

For all of the three datasets, we utilised the deep visual features extracted from popular ConvNet architecture to describe the frames. Specifically, we used the second output of the fully-connected layer of the VGG16 model [23] and then applied L2 normalisation. Thus each frame is represented by a 4,096d visual feature vector. On the Oxford RobotCar dataset, each location is visually represented by multiple frames. To reduce the noise and fit class conditional densities to the data with multiple exemplars, we further applied the Linear Discriminant Analysis (LDA) to reduce the dimensionality to 64.

In our experiment, we compared our model with KD-Tree[24], Iterative Quantization (ITQ)[25] and Optimised Product Quantization (OPQ)[11] on Nordland Train and Aerial Brisbane datasets. KD-Tree is a well-known method for approximate nearest neighbour search, which builds a binary search tree as the index for a fixed-sized database. Although a KD-Tree can effectively accelerate the computation, this technique does not compress the data. ITQ and OPQ are discrete embedding approaches, which encode the high-dimensional data into compact codes for fast computing. However, they cannot achieve unique mapping for place recognition because code collision is inevitable. Furthermore, these techniques compress data in an absolute way, i.e., the compressed data size is proportional to the actual data size, which is not sub-linear.

Let T be the level of the period length, which is the minimum value for τ in data encoding. Assume we set $T = 2\sqrt{N}$ and $r = 2$ for our model, and use a 256-bit binary vector and a 256d integer vector to encode a 4,096d visual instance for ITQ and OPQ, respectively. When data size is comparably small, our proposed model has a higher memory cost, but it increases sub-linearly when the database becomes extremely large.

We used the compression ratio to demonstrate how our model can achieve sub-linear storage and used the accuracy metric to evaluate place recognition performance. We also compared the computational speed of our model with the baseline models.

Our experiment was conducted on a desktop with Intel(R)




Fig. 4: The training errors on Aerial Brisbane dataset when setting different period lengths.

i7-7700K CPU 4.20GHz with 4 processors, 32GB RAM, and Windows 10 operating system with a Python 3.6 computational environment.

B. Place Recognition Results When All Queries Are from Reference Data

We first investigated how the period length τ affects the training error. We tested different lengths of periods and trained the linear SVMs on the Aerial Brisbane dataset, with the training errors illustrated in Fig. 4. We can see longer periods lead to a lower training error rate e and a higher compression ratio when r is fixed. In the extreme case, when a whole frame database has only one period, i.e., $r = 1$ and $\tau = N$, no data compression is implemented, and the model reverts to brute-force search.

We then investigate the data compression results for Nordland Train and Aerial Brisbane datasets when only two periodic patterns are available, i.e., $r = 2$. We also tested the training errors at two periodic levels: $T = \sqrt{N}$ and $T = 2\sqrt{N}$, respectively. For each subset of Nordland Train dataset, as well as the Aerial Brisbane dataset, we tried seven different values of τ in training the linear SVMs and recorded the error rates, and then the system automatically selected the best period pair. Based on the selection of periods, we analyse the storage cost when applying our proposed data compression approach and selecting the length of period T at the level of \sqrt{N} on the two datasets. In a 32-bit operating system, the memory cost for a float number is 4 bytes. The storage comparison of the datasets is summarised in Table I. From the table, we can see that our proposed model is able to encode very large frame databases with high compression ratios. If the data size increases linearly, applying two-period values and several templates can make the storage increase in a sub-linear manner. When the number of frames is more than 10,000, our model only takes about 1/50 memory to store all data instances.

We used the frames from the reference data as queries and applied our models for location estimation. We show the place recognition results of Nordland Train and Aerial Brisbane datasets in Table II. In all of the four subsets of the Nordland Train dataset, none of the accuracies falls below 98% even when we apply the extreme compression method. For example, in the spring subset, 10,713 frames record different views of places along the Norland railway. Applying our proposed data compression model can still achieve 99.46% accuracy. If we set longer periods, i.e., $T = 2\sqrt{N}$,




Fig. 5: The compression ratios and accuracy comparisons when r changes. Note the queries are all from reference data. (1. Nordland spring; 2. Nordland summer; 3. Nordland fall; 4. Nordland winter; 5. Aerial Brisbane)

the compression ratio doubles, but the recognition accuracy is higher, which is very close to 100%. On the Aerial Brisbane dataset, our system achieved the very near-perfect accuracy of 99.92% (only 3 mismatches) when $T = \sqrt{N}$. When the length of period doubles, the recognition accuracy is 100%.

We conducted an experiment using both Nordland Train and Aerial Brisbane datasets evaluating the performance of the KD-Tree, ITQ, OPQ, and the brute-force search techniques, recording the average search time. The comparison is displayed in Table III. By using a few learned periodic templates and the matching approach introduced in section III-C, the computational efficiency is significantly increased compared to the exhaustive search, although it is a lower than KD-Tree, ITQ and OPQ. However, KD-Tree only builds an index on the database but does not implement the data compression at all. Both ITQ and OPQ could compress the data in an absolute manner, but they cannot achieve the unique mapping required for place recognition, even when the distance between a query and its matched frame is zero. Considering the compression ratio, search accuracy and speed, it is worth applying our proposed model to large-scale place recognition systems.

Then we used more than two periods by setting $r = 3$ and $r = 5$ respectively and re-ran our model on Nordland Train and Aerial Brisbane datasets. As is discussed in section III-B, we can choose different available periods for data compression. A larger value of r means the system can achieve a lower compression ratio. Fig. 5 summarises the compression ratios and the accuracy comparisons. From the two figures, it can be seen that although applying more periodic patterns can achieve an even higher compression ratio, the recognition accuracy falls significantly, even though the queries are all from reference data. The reason is that the value of τ is in direct proportion to the number of negative data instances in Eq. (1), i.e. when the number of templates increases within a period, there are fewer positive-labelled instances in the learning process, which makes the periodic templates more linearly separable. By contrast, the training error rate is higher on a more “balanced” dataset.

We tested our system using data from different times of day for the Aerial Brisbane and Oxford Robotcar datasets. For Aerial Brisbane dataset, we used one image from this

TABLE I: The data compression results ($r = 2$ and $T = \sqrt{N}$). Note the 4097 in the Compressed size column comes from 4096 plus the bias.

Dataset	Original size	Original storage	Compressed size	Compressed storage	Compression ratio
Norland (spring)	$\mathbb{R}^{10713 \times 4096}$	175,521,792 bytes	$\mathbb{R}^{211 \times 4097}$	3,457,868 bytes	0.0207
Norland (summer)	$\mathbb{R}^{7403 \times 4096}$	121,290,752 bytes	$\mathbb{R}^{181 \times 4097}$	2,966,228 bytes	0.0246
Norland (fall)	$\mathbb{R}^{9267 \times 4096}$	151,830,528 bytes	$\mathbb{R}^{201 \times 4097}$	3,293,988 bytes	0.0217
Norland (winter)	$\mathbb{R}^{7276 \times 4096}$	119,209,984 bytes	$\mathbb{R}^{179 \times 4097}$	2,933,452 bytes	0.0246
Aerial Brisbane	$\mathbb{R}^{3705 \times 4096}$	60,702,720 bytes	$\mathbb{R}^{125 \times 4097}$	2,048,500 bytes	0.0337

TABLE II: Best learned periods and place recognition results on Aerial Brisbane dataset and Nordland train dataset ($r = 2$ and $T = \sqrt{N}$). Note all queries are from the reference data.

Dataset	$T = \sqrt{N}$					$T = 2\sqrt{N}$				
	τ_1	$e (\times 10^{-3})$	τ_2	$e (\times 10^{-3})$	Accuracy	τ_1	$e (\times 10^{-3})$	τ_2	$e (\times 10^{-3})$	Accuracy
Norland (spring)	105	2.99	106	2.89	0.9946	211	0.84	212	0.65	0.9990
Norland (summer)	88	8.91	93	8.65	0.9846	179	2.70	180	2.30	0.9960
Norland (fall)	98	4.64	103	4.64	0.9912	197	0.54	200	0.22	0.9992
Norland (winter)	87	1.51	92	1.52	0.9971	174	0.27	175	0.00	0.9996
Aerial Brisbane	62	0.54	63	0.27	0.9992	122	0.00	123	0.00	1.0000

TABLE III: The comparisons of different search methods (the search time is recorded on Aerial Brisbane dataset).

Method	Compressed scale	Search time	Unique mapping
KD-Tree	No compression	0.000215	Yes
ITQ	Linear	0.000116	No
OPQ	Linear	0.000241	No
Brute-force	No compression	0.125948	Yes
Ours	Sub-linear	0.000503	Yes

dataset as queries, and used another image taken at a different time as the reference. For Oxford RobotCar dataset, we used the frames taken on a distinct date as queries to search their locations. When applying the brute-force search, the accuracy is 0.9582 on Aerial Brisbane dataset, and 0.236 on Oxford RobotCar dataset, respectively. Using our model to compress the reference data, the accuracy dropped to 0.6835 and 0.183 on the two datasets. The result signifies solely applying the data compression model cannot achieve a satisfactory recognition accuracy under different appearance conditions. As is introduced in Section III-D, we adopted different data augmentation approaches, including Gaussian blur, random cropping, flipping, elastic transformation, contrast normalisation, etc. We separately experimented with these augmented data sources, then merged them as a whole training set to train a unified place recognition model. Note that for Oxford RobotCar dataset, we did not use channel invert and grey-scale augmentations because the raw image data is grey-scale. The accuracies are displayed in Fig. 6 for the two datasets. It can be seen that although each single augmented data source has limited power to help obtain a discriminative model, their combination can effectively boost the accuracy by 2% and 1% on the two datasets, respectively.

C. Results in Visually Changing Environments

Next, we trained the models with the multi-reference set, where these frames are taken under different appearance conditions. We used different combinations of reference sources and evaluated their performance on Aerial Brisbane and Oxford RobotCar datasets, and the accuracy curves are

plotted in Fig. 7. The accuracy improves steadily as the number of data sources increases and learning with the multi-source data and image augmentation enables the data compression model to better deal with the various appearance conditions while keeping the storage sub-linear.

As is discussed in Section III-B, we set longer periods to further reduce the training error e , but with a lower data compression ratio. We tested different lengths of periods by setting $T = \sqrt{N}$, $T = 2\sqrt{N}$, $T = 3\sqrt{N}$, and $T = 4\sqrt{N}$ respectively, and the accuracy curve is plotted in Fig. 8. We can conclude that setting a longer period can improve the recognition accuracy, with the sacrifice of the compression rate. In this real application scenario, if we set $T = 4\sqrt{N}$, the data compression ratio is 0.131 on Aerial Brisbane dataset and 0.81 on Oxford RobotCar dataset, respectively, but the recognition performance is boosted.

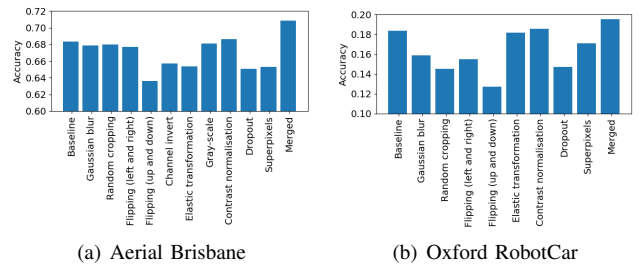


Fig. 6: The accuracies of different data augmentations.

V. DISCUSSIONS AND CONCLUSIONS

We have presented a novel image-based map encoding scheme that deliberately seeks out and learns mutually supportive visual pattern frequencies in the environment to enable place recognition with sub-linear storage growth as the environment size increases. The system is based on the nature of neural mapping systems in the mammalian brain that does not appear to approach the data association problem central to most robotic mapping systems the same




Fig. 7: The accuracy curves when there are multiple reference data sources ($r = 2$).

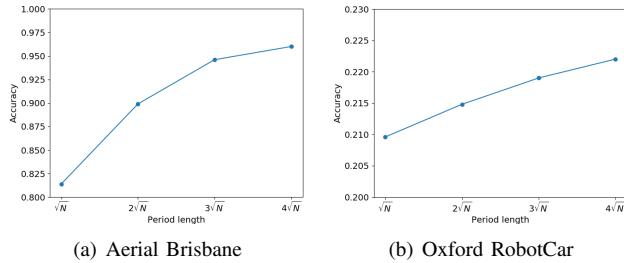


Fig. 8: The accuracy curve when setting different lengths of periods ($r = 2, 4$ data sources with augmentation).

way; instead each neural map “unit” is associated with an arbitrarily large number of places in the environment distributed at regular intervals.

Results on large real-world datasets show that the fundamental premise is valid and that high-performance place recognition can be achieved with a mapping system whose map storage scales sub-linearly with environment size. The system is agnostic of any particular types of features or feature frequencies and its performance across a range of environments shows that, perhaps surprisingly, repetitive visual patterns can usually be found.

We applied the data augmentation and multi-source training data, which are generic methods for visual recognition tasks, to make our model more applicable under different appearance conditions. In future work, we could also design a more sophisticated system by integrating some advanced machine learning techniques to better capture the spatial properties of the periodic patterns and improve the recognition performance. Alternatively, we could apply some existing matching schemes such as SeqSLAM [26], to further improve the stability by taking consideration of multi-frame integration.

ACKNOWLEDGEMENTS

This work was supported by an Asian Office of Aerospace Research and Development Grant FA2386-16-1-4027 and an ARC Future Fellowship FT140101229 to MM.

REFERENCES

[1] S. Lowry, N. Sünderhauf, P. Newman, J. J. Leonard, D. Cox, P. Corke, and M. J. Milford, “Visual place recognition: A survey,” *IEEE Trans. on Robotics*, vol. 32, no. 1, pp. 1–19, 2016.

[2] L. M. Giacomo, E. A. Zilli, E. Fransén, and M. E. Hasselmo, “Temporal frequency of subthreshold oscillations scales with entorhinal grid cell field spacing,” *Science*, vol. 315, no. 5819, pp. 1719–1722, 2007.

[3] T. Hafting, M. Fyhn, S. Molden, M.-B. Moser, and E. I. Moser, “Microstructure of a spatial map in the entorhinal cortex,” *Nature*, vol. 436, no. 7052, pp. 801–806, 2005.

[4] E. I. Moser, E. Kropff, and M.-B. Moser, “Place cells, grid cells, and the brain’s spatial representation system,” *Annual review of neuroscience*, vol. 31, 2008.

[5] S. Sreenivasan and I. Fiete, “Grid cells generate an analog error-correcting code for singularly precise neural computation,” *Nature neuroscience*, vol. 14, no. 10, pp. 1330–1337, 2011.

[6] A. Jacobson, W. Scheirer, and M. Milford, “Déjà vu: Scalable place recognition using mutually supportive feature frequencies,” in *IEEE IROS*, 2017.

[7] J. Wang, T. Zhang, j. song, N. Sebe, and H. T. Shen, “A survey on learning to hash,” *IEEE Trans. on Pattern Analysis and Machine Intelligence*, vol. PP, no. 99, pp. 1–1, 2017.

[8] M. Bosse and R. Zlot, “Keypoint design and evaluation for place recognition in 2d lidar maps,” *Robotics and Autonomous Systems*, vol. 57, no. 12, pp. 1211 – 1224, 2009.

[9] A. Gersho and R. M. Gray, *Vector quantization and signal compression*. Springer Science & Business Media, 2012, vol. 159.

[10] H. Jegou, M. Douze, and C. Schmid, “Product quantization for nearest neighbor search,” *IEEE Trans. on Pattern Analysis and Machine Intelligence*, vol. 33, no. 1, pp. 117–128, 2011.

[11] T. Ge, K. He, Q. Ke, and J. Sun, “Optimized product quantization,” *IEEE Trans. on Pattern Analysis and Machine Intelligence*, vol. 36, no. 4, pp. 744–755, 2014.

[12] L. Yu, Z. Huang, F. Shen, J. Song, H. T. Shen, and X. Zhou, “Bilinear optimized product quantization for scalable visual content analysis,” *IEEE Trans. on Image Processing*, vol. 26, no. 10, pp. 5057 – 5069, 2017.

[13] P. Heise, B. Jensen, S. Klose, and A. Knoll, “Fast dense stereo correspondences by binary locality sensitive hashing,” in *IEEE ICRA*, 2015, pp. 105–110.

[14] W. Takano and Y. Nakamura, “Bigram-based natural language model and statistical motion symbol model for scalable language of humanoid robots,” in *IEEE ICRA*, 2012, pp. 1232–1237.

[15] R. Arroyo, P. F. Alcantarilla, L. M. Bergasa, J. J. Yebe, and S. Bronte, “Fast and effective visual place recognition using binary codes and disparity information,” in *IEEE IROS*, 2014, pp. 3089–3094.

[16] B. Hollis, S. Patterson, and J. Trinkle, “Compressed sensing for tactile skins,” in *IEEE ICRA*, 2016, pp. 150–157.

[17] D. Feldman, A. Sugaya, and D. Rus, “An effective coresets compression algorithm for large scale sensor networks,” in *ACM/IEEE IPSN*, 2012, pp. 257–268.

[18] S. M. Siam and H. Zhang, “Fast-seqslam: A fast appearance based place recognition algorithm,” in *IEEE ICRA*, 2017, pp. 5702–5708.

[19] M. Cummins and P. Newman, “Accelerated appearance-only SLAM,” in *IEEE ICRA*, 2008.

[20] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, A. C. Berg, and L. Fei-Fei, “ImageNet Large Scale Visual Recognition Challenge,” *International Journal of Computer Vision (IJCV)*, vol. 115, no. 3, pp. 211–252, 2015.

[21] P. Neubert, N. Sünderhauf, and P. Protzel, “Appearance change prediction for long-term navigation across seasons,” in *IEEE EECR*, 2013, pp. 198–203.

[22] W. Maddern, G. Pascoe, C. Linegar, and P. Newman, “1 Year, 1000km: The Oxford RobotCar Dataset,” *The International Journal of Robotics Research (IJRR)*, vol. 36, no. 1, pp. 3–15, 2017.

[23] K. Simonyan and A. Zisserman, “Very deep convolutional networks for large-scale image recognition,” *arXiv preprint arXiv:1409.1556*, 2014.

[24] J. L. Bentley, “Multidimensional binary search trees used for associative searching,” *ACM Commun.*, vol. 18, no. 9, pp. 509–517, 1975.

[25] Y. Gong, S. Lazebnik, A. Gordo, and F. Perronnin, “Iterative quantization: A procrustean approach to learning binary codes for large-scale image retrieval,” *IEEE Trans. on Pattern Analysis and Machine Intelligence*, vol. 35, no. 12, pp. 2916–2929, 2013.

[26] M. J. Milford and G. F. Wyeth, “Seqslam: Visual route-based navigation for sunny summer days and stormy winter nights,” in *IEEE ICRA*, 2012, pp. 1643–1649.