

Boosting Performance of a Baseline Visual Place Recognition Technique by Predicting the Maximally Complementary Technique

Connor Malone Stephen Hausler Tobias Fischer Michael Milford

Abstract—One recent promising approach to the Visual Place Recognition (VPR) problem has been to fuse the place recognition estimates of multiple complementary VPR techniques using methods such as shared representative appearance learning (SRAL) and multi-process fusion. These approaches come with a substantial practical limitation: they require all potential VPR methods to be brute-force run before they are selectively fused. The obvious solution to this limitation is to predict the viable subset of methods ahead of time, but this is challenging because it requires a predictive signal within the imagery itself that is indicative of high performance methods. Here we propose an alternative approach that instead starts with a *known* single base VPR technique, and learns to predict the most complementary *additional* VPR technique to fuse with it, that results in the largest improvement in performance. The key innovation here is to use a dimensionally reduced difference vector between the query image and the top-retrieved reference image using this baseline technique as the predictive signal of the most complementary additional technique, both during training and inference. We demonstrate that our approach can train a single network to select performant, complementary technique pairs across datasets which span multiple modes of transportation (train, car, walking) as well as to generalise to unseen datasets, outperforming multiple baseline strategies for manually selecting the best technique pairs based on the same training data.

I. INTRODUCTION

Visual Place Recognition (VPR) is a critical component for robotic and autonomous platforms: providing the ability to recognise previously visited locations solely using image data, which is beneficial for a number of downstream tasks including Simultaneous Localization And Mapping (SLAM) and navigation [1], [2], [3], [4]. VPR is challenging due to the range of appearance and viewpoint changes that can arise when navigating. Appearance changes include diverse seasonal changes, such as summer to winter, or illumination changes like day to night cycles.

Because of the wide variety of appearance and viewpoint changes, it is challenging to design a single approach to VPR that is effective across all possible deployment environments [4]. Two well known examples [5] supporting this observation are that strong techniques such as NetVLAD [6] perform poorly on the Nordland dataset [7] and that the Histogram of Oriented Gradients [8] fail on the Berlin Kudamm dataset [9]. More generally, several works recently highlighted

All authors are with the QUT Centre for Robotics, School of Electrical Engineering and Robotics, QUT, Brisbane, Australia.

This research was partially supported by funding from an Amazon Research Award, grant 81418190, from ARC Laureate Fellowship FL210100156 to MM, and by funding from Intel’s Neuromorphic Computing Lab to TF and MM. The authors acknowledge continued support from the Queensland University of Technology (QUT) through the Centre for Robotics.

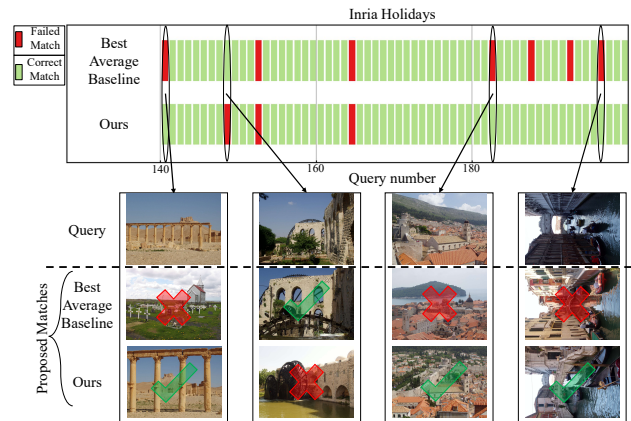


Fig. 1. Here we demonstrate that our approach for dynamically selecting the maximally complementary VPR technique for a given base technique is able to outperform choosing the average best performing technique in the Inria Holidays dataset. Green bars denote successful place matching and red bars denote failure. We include sample images of the queries and top VPR matches to show where our method outperforms the comparison. We also include a failure case in the second column.

that the optimal VPR technique for a given dataset varies significantly [10], [5].

One workaround to this problem is to fuse multiple VPR techniques in tandem, creating a pseudo-multi-sensor fusion with a single sensor [11], [12], building on prior multi-sensor fusion approaches [13], [14], [15]. By using multiple VPR techniques simultaneously, limitations with one technique (in a specific environment) can be offset with the inclusion of additional techniques. This approach yields performance benefits but has the significant downside of requiring substantive computational capability to run all methods in parallel. Recent work has further shown that certain combinations of techniques are more effective than others on specific datasets – a phenomenon that has been termed the *complementarity* of different combinations of techniques [11], [9], [16].

In this work, we propose a new classification approach that learns to predict the most complementary VPR technique for a known base VPR technique on a frame-by-frame basis. Here we define complementary to mean the additional technique that results in the largest performance improvement when fused with the base technique. Importantly, this does not require the brute-force running of all techniques beforehand. Our approach first performs VPR using the known base technique to obtain the top matching reference image. The *difference vector* between the query descriptor and retrieved top reference descriptor is used as the input to a multi-label

classification network. The network then predicts the best additional VPR technique to fuse with the base technique for VPR using multi-process fusion [12], resulting in an updated place match estimate. We demonstrate that our approach can train a single network to predict complementary technique pairs across datasets captured with varying modes of transport (train, car, walking) and outperform manual selection methods given the same training data. In addition, we show that the same network is capable of generalising to unseen datasets. Our proposed solution enables computationally efficient multi-process fusion to be deployed in various settings for VPR with little training and dynamically chosen technique pairs using a given base technique.

We make the following specific contributions:

- 1) We introduce a new ‘anchored’ complementarity search concept for VPR, that assumes a baseline technique and searches for a maximally complementary additional technique without requiring brute force running of these technique candidates
- 2) We present a multi-label classification approach that predicts VPR technique utility for fusion with a baseline technique, using as input the truncated difference vectors between a query image and image match candidate output by this baseline VPR technique.
- 3) We present extensive experimental results demonstrating that the proposed approach outperforms strong baselines across both new parts of familiar environment types but also completely unseen environments.
- 4) We provide quantitative analysis of the characteristics of the chosen technique pairs, including their usage distribution within and across different environmental types.

Source code is publicly available at https://github.com/CMalone-Jupiter/VPR_Pair_Predict.

II. RELATED WORK

A. Fusing Multiple Modalities for Improved VPR

Fusing multiple sources of information is a common technique in both SLAM and VPR [17], [4]. Traditionally this has been achieved through the use of multiple physical sensors, like LiDAR, depth cameras, RADAR and Wi-Fi [18], [19], [20]. Numerous algorithms have been devised to intelligently fuse these different sensors and produce a more robust localization estimate [14], [13].

Recent work has shown that some of the benefits of multi-sensor fusion can still be attained with a single sensor, by fusing different interpretations of the image data that are obtained by applying different image processing techniques on the same image. SRAL [11] showed that a convex optimization problem can be used to optimize the contribution weights of a set of fused image processing modalities, resulting in improved VPR performance. The optimization was computed for a separate training set for each test dataset. Multi-process fusion [12] fuses multiple modalities within a Hidden Markov Model, combining the benefits of both sequences and multiple techniques. Recent works

demonstrated that certain combinations of modalities are more effective than others on specific datasets [9], [16]. However, [9], [16] calibrate dataset specific combinations and do not generalise to more than one dataset at a time.

B. Sensor, Technique and Reference Data Selection

In [21], the sensor selection problem was relaxed from a binary selection to a soft weighted fusion, which enabled the use of convex optimization theorems to approximate the solution to the problem without requiring a brute-force search of all combinations of sensors. In VPR specifically, [22] used a Greedy algorithm to find the optimal set of feature maps to use to improve localization performance, and further work used mutual information to determine the ideal subset of feature maps to use [23].

[24] introduced Bayesian Selective Fusion to intelligently select the optimal reference set of images for VPR. Given a set of reference images taken at different times, their algorithm found an optimal subset of reference times and also weighted the different sets to optimize the localization performance. Another recent work proposes an unsupervised algorithm for automatically selecting optimal parameters for VPR algorithms [25], using only one technique at a time.

C. Place Recognition as a Classification Problem

Classification is the task of evaluating an image or image region and assigning it a label from a predefined set. Multi-label classification is an extension of this task where labels are not mutually exclusive and an image can have more than one label. We note that we differ from recent work that re-framed place recognition from an image retrieval problem to a classification task to overcome issues such as large scale localisation [26] and neuromorphic learning [27]. In our work, we still frame VPR as an image retrieval problem, but use multi-label classification to select the most complementary technique given a baseline technique.

III. PROPOSED APPROACH

Our approach aims to maximise VPR performance across varying scenes by fusing a known base technique with the most performant additional technique for a given query using multi-process fusion. This is achieved by training a multi-label classifier to predict the most suitable VPR technique to fuse with the given base technique. The classifier network uses truncated descriptors from the base technique as input features and produces likelihood values for a set of additional VPR techniques indicative of their additive performance potential on top of the base technique. The prediction of the best pairing is taken as the one with the highest likelihood value. We proceed by first providing a brief summary of the multi-process fusion method adopted from previous work and follow with our proposed multi-label classification approach.

A. Multi-Process Fusion

In VPR, each query image observed by the autonomous platform is compared against a database of \mathcal{D} prior images using a specified feature extraction and matching technique.

This produces a D dimensional similarity vector, containing a list of similarity scores between the query image and each database image. A larger score denotes a stronger VPR match.

In a multi-process fusion (MPF) algorithm [12], a variety of different VPR techniques are used simultaneously and a fusion of similarity scores can improve the overall VPR performance of a system. Formally, given a set of techniques $\mathcal{N} = \{1, \dots, N\}$, each technique $n \in \mathcal{N}$ separately computes a similarity vector \mathbf{s}_n containing the similarity scores between a query image and a set of database images. Since the set of techniques is arbitrary, the distribution of scores within each technique will not be consistent with the distribution in other techniques. Therefore, a normalization process is used prior to fusing the set of \mathcal{N} techniques, so that each similarity vector has a minimum value of zero and a maximum value of one:

$$\hat{\mathbf{s}}_n(i) = \frac{\mathbf{s}_n(i) - \min(\mathbf{s}_n)}{\max(\mathbf{s}_n) - \min(\mathbf{s}_n)} \quad \forall i, \quad n \in \mathcal{N} \quad (1)$$

The collection of similarity vectors $\hat{\mathbf{s}}$ are then summed to produce a combined similarity vector:

$$\mathbf{s}_{\text{MPF}} = \sum_{n=1}^{\mathcal{N}} \hat{\mathbf{s}}_n \quad (2)$$

The matching image X is then the image with the maximum score in \mathbf{s}_{MPF} : $X = \arg \max \mathbf{s}_{\text{MPF}}$ (see Figure 2).

While such MPF algorithms have previously been shown to outperform single technique baselines, typically only a sub-selection of techniques will be suited to a particular dataset. The drastic difference in performance between different datasets and even within different sections of the same datasets of particular methods was shown in [10], [28]. Therefore, the inclusion of poorly performing techniques (which is the case for standard MPF) can potentially offset any advantages of using multiple methods simultaneously. As an alternative, in the following we describe how the best technique to pair with a given base technique can be dynamically chosen using a multi-label classification approach.

B. Multi-Label Classification

Multi-label classification is a problem where there are multiple correct labels which can all describe a query. This is highly relevant to technique selection in place recognition where multiple separate VPR techniques could all accurately localise a query. In this work, we apply this concept to the problem of finding complementary technique pairs for visual place recognition from the set $\mathcal{N} = \{1, \dots, N\}$ given a base technique, $B \in \mathcal{N}$. For a given query, there are typically multiple pairs of VPR techniques which can localise the query within database \mathcal{D} . Therefore we utilise a multi-label classification framework to learn the likelihood that each technique in \mathcal{N} is complementary to the base technique B . This is achieved by using descriptors from B as input features.

1) *Inputs*: Our method summarises the difference between query images and the reference database \mathcal{D} using technique B to compute the difference vector $\mathbf{z}_{B,\Delta}$ between a query descriptor $\mathbf{z}_{B,q}$ and the top VPR match descriptor $\mathbf{z}_{B,r}$ as determined by technique B .

$$\mathbf{z}_{B,\Delta} = \mathbf{z}_{B,q} - \mathbf{z}_{B,r} \quad (3)$$

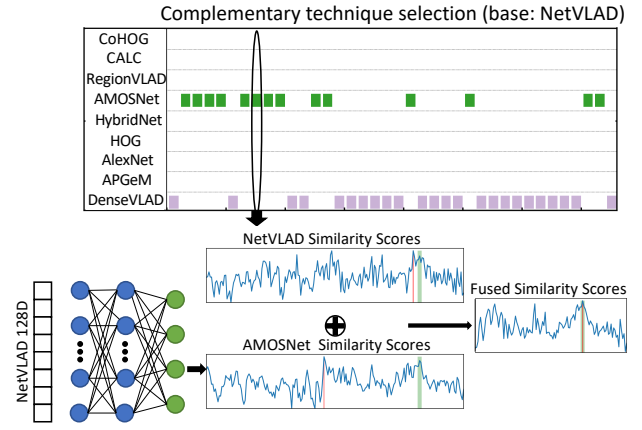


Fig. 2. Given a base technique (NetVLAD), our approach uses a multi-label classification network to determine the most complementary technique pair for VPR using multi-process fusion. *Top*: we demonstrate our approach predicting the most complementary technique pair on a frame-by-frame basis across part of one tested dataset. *Bottom*: for a given query we compute input features using NetVLAD descriptors and use a small classifier network to predict the best technique to pair with NetVLAD. Multi-process fusion is then used to perform VPR using similarity scores for a reference database. In the similarity score plots, the green bar denotes the ground truth and the red bar denotes the top VPR match from the reference database.

To reduce the number of input features used in the network, the difference vector is flattened and truncated. In this work, we use NetVLAD [6] as our base technique B and take the first 128 principal components of the standard PCA reduction.

At training time we make the adjustment that $\mathbf{z}_{B,r}$ will be the descriptor of the ground truth reference image as computed by B , rather than the descriptor of the top VPR match. This removes erroneous training data introduced by incorrect VPR matches from technique B and was found to achieve a higher recall performance.

2) *Network*: Given that the chosen input features are already an abstraction of the relevant image data, the network architecture adopted for the following multi-label classifier is a simple multi-layer perceptron model as in Figure 2. The model has an input layer with neurons corresponding to the length of $\mathbf{z}_{B,\Delta}$; all hidden layers are given a ReLU activation and dropout; and the output layer has $|\mathcal{N}|$ neurons to represent all techniques which could pair with technique B in the eventual multi-process fusion.

3) *Training*: When training the classifier, we utilised the typical methods that are applied in other multi-label classification problems, namely, a sigmoid activation is applied to the networks output to bound the label likelihoods between 0 and 1; and a binary cross entropy (BCE) loss is used to train each label output independent of each other:

$$\text{BCE} = -\frac{1}{\mathcal{N}} \sum_{n=1}^{\mathcal{N}} (y_n \log \hat{y}_n + (1 - y_n) \log(1 - \hat{y}_n)), \quad (4)$$

where y_n is the target for label n and \hat{y}_n is the networks predicted likelihood for label n . In this work, label n represents a VPR technique from the set \mathcal{N} and \hat{y}_n is the networks prediction of the complementarity between technique n and technique B . y_n is 1 if the combined pair n and B retrieve a reference image within distance tolerance of the

corresponding image from \mathcal{D} , else it is 0. A sigmoid activation with *BCE* is used because for any query there could be multiple techniques n from \mathcal{N} where y_n is high. Therefore a standard softmax with cross entropy loss would unduly punish high activations for non-maximum predictions \hat{y}_n where the target y_n is also high.

4) *Prediction*: In the context of visual place recognition, despite training as a multi-label problem, the desired output from the classification network is the single most complementary technique to fuse with technique B for place recognition \hat{Y} . This is simply taken as the output neuron/technique from the network with the highest likelihood value.

$$\hat{Y} = \arg \max(\hat{y}) \quad (5)$$

IV. EXPERIMENTAL SETUP

We demonstrate our proposed approach to selecting complementary VPR techniques for a known base technique through a diverse collection of methods and datasets largely provided by the recent VPR-Bench [10].

A. Datasets

We use the Gardens Point Walking [29], Nordland [7], Pittsburgh [30], and INRIA Holidays [31] benchmark datasets as provided by VPR-Bench [10]. For these datasets, we use the same dataset configurations and ground truth tolerances as used in VPR Bench. In addition, we include traverses from both the Oxford RobotCar [32] and SFU Mountain [33] datasets. Taken together, these datasets are highly varied, covering a wide range of environments with both traversal and non-traversal (images in the dataset are independent to each other) datasets. A variety of illumination, viewpoint, seasonal and structural changes occur between the reference and query sets of these datasets. For completeness, we provide a brief overview of each dataset:

1) *Gardens Point Walking*: The Gardens Point Walking dataset consists of walking traverses through the Queensland University of Technology, Gardens Point campus. The query and reference sets used in this work include a day and night traverse respectively. As well as strong appearance shift, these images also include viewpoint variations.

2) *Nordland*: The Nordland dataset is a common VPR benchmark dataset covering a 728 kilometer train journey through Norway across multiple seasons. We use subsets from the Winter and Summer traverses as query and reference data.

3) *Pittsburgh*: The Pittsburgh dataset is a large-scale driving dataset which includes full 360 degree viewpoint variations for each query image.

4) *Inria Holidays*: Unlike the above datasets, the Inria Holidays dataset is not a continuous traversal dataset. It is comprised of images from a wide range of geographically disconnected indoor, outdoor, coastal, densely vegetated and many more types of scenes.

5) *Oxford RobotCar*: The Oxford RobotCar set is a long-term localisation dataset which was created by driving a route through central Oxford twice a week for over a year. It covers environment changes such as illumination, seasonal

and structural changes. In particular, we adopt a night and day traverse for our query and reference data in this work.

6) *SFU Mountain*: The SFU Mountain dataset consists of more than 8 hours of trail driving captured by a Clearpath Husky robot platform navigating Burnaby Mountain, British Columbia, Canada. It covers a year of changing environmental conditions including sun, rain and snow. This work includes a snow and sunny traverse as query and reference data.

B. Considered Techniques for Fusion

For all experiments, we also adopt the same range of hand-crafted and deep-learned VPR techniques as VPR-Bench. Specifically, we set our base technique $B = \text{NetVLAD}$ [6] and our further set of techniques \mathcal{N} to include RegionVLAD [34], CoHOG [35], HOG [8], AlexNet, AMOSNet [36], HybridNet [36], CALC [37], AP-GeM [38] and DenseVLAD [39] as defined in Section III-B.

C. Multi-Label Datasets

In order to train the multi-label classifier for technique selection, we generated multi-hot labels for each of the datasets above, indicative of which technique pairs would localise respective queries. For a given query, a technique pair was given a positive label if the top VPR match was within tolerance of the ground truth reference place within \mathcal{D} . We limited the technique pairs to those that include the given base technique B , NetVLAD.

Each dataset was then split into geographically separated train, validation and testing sets. Since all of the datasets except Inria Holidays consist of sequential frames captured along a route, we maintained geographical separation between the splits by using the first 60% of queries for training, the next 20% for validation and the final 20% for testing purposes. To demonstrate robustness across multiple datasets, we also provide results on a curated dataset where the respective train, test and validation splits of the Gardens Point, Nordland, Pittsburgh and SFU Mountain datasets were combined.

To prevent the occurrence of queries without positive labels when training the classification network, we remove queries where none of the technique pairs successfully localised. We emphasise that this does not affect performance trends to favour our method; we simply remove redundant data where none of the technique pairs can possibly be successful and therefore place recognition using these pairs is not possible in any case. Table I shows that the Nordland and Oxford RobotCar datasets are most affected by this curation.

D. Baseline Comparisons

For evaluation we compare against two baselines which represent the most likely strategies for manually selecting a single technique pair, inclusive of the base technique B (NetVLAD), to perform visual place recognition across an entire dataset given the same training data. The first baseline we establish for comparison is selecting the NetVLAD technique pair which performs the best on average across the entire training data. For our experiments we found this to be NetVLAD paired with AMOSNet.

TABLE I
NUMBER OF QUERIES WITHIN EACH DATASET

Dataset	Train (all)	Val (all)	Test (all)	Train (curated)	Val (curated)	Test (curated)
SFU Mountain	142	46	50	88	35	47
GP Walking	120	39	41	113	33	38
Nordland	1656	551	553	707	129	82
Pittsburgh	600	199	201	585	197	196
RobotCar Night	-	-	3876	-	-	1891
Inria Holidays	-	-	300	-	-	280
Total	2518	835	5021	1493	394	2534

The second baseline which we use for comparison is selecting the best single technique pair, inclusive of the base technique B (NetVLAD), specifically for each dataset based on average performance across the training sets. These were NetVLAD & Ap-GeM for Gardens Point Walking; NetVLAD & AMOSNet for Nordland; NetVLAD & DenseVLAD for Pittsburgh; and NetVLAD & HybridNet for the SFU Mountain set. Since the RobotCar and Inria Holidays datasets are withheld from training to test as unseen data, we cannot compare them against this second dataset-specific baseline.

E. Metrics

For quantitative results, we use the standard Recall@1 metric which is widely used for evaluating place recognition systems. Recall@1 indicates the percentage of queries which were accurately localised using the top VPR match from the reference database. For this work, the multi-hot labels created for training the classifier network are a query level representation of the Recall@1.

F. Parameter Optimization

During the optimization of network parameters we performed a Bayesian sweep over *batch size*, *learning rate*, *hidden layer size*, *number of hidden layers* and the *dropout chance* at each hidden layer. To maintain a low model complexity and reduce the chance of over-fitting, the search limited the number of hidden layers to 1-3 and the hidden layer size to 32, 64, 128 or 256. The final model parameters were $batch\ size = 8$, $learning\ rate = 4.550325e^{-4}$, $hidden\ layer\ size = 32$, $num\ hidden\ layers = 1$ and $dropout = 0.126450$. The network with these parameters trained for 17 epochs over about 10 seconds and performed inference at about 0.14ms per query using an Intel i7-8665U CPU.

V. RESULTS & DISCUSSION

A. Comparison to Baseline Techniques

In this section, we present the results for using our multi-label classification network to predict complementary VPR techniques to pair with a given base technique on a frame-by-frame basis. We emphasize that all of the following results have been obtained by training a single network on the combined training set outlined in Section IV-C and testing on both a) a geographically separate test set and b) two datasets from completely unseen environmental types.

We start with Table II which compares the recall of our system against multiple baselines. For the purposes of the

TABLE II
RECALL @ 1 FOR ALL DATASETS (%)

Dataset	NetVLAD	Best average train pair	Best dataset-specific train pair	Ours
SFU Mt (Test)	57.5	89.4	87.2	89.4
GP Walk (Test)	44.7	57.9	57.9	60.5
Nordland (Test)	24.4	61.0	61.0	61.0
Pittsburgh (Test)	97.5	87.2	98.5	98.5
Combined Test	57.5	78.5	-	84.8
RobotCar Night	30.9	53.1	-	63.1
Inria Holidays	81.8	90.0	-	96.1
Mean Recall	56.1	73.1	74.6	78.1

‘Mean Recall’ values, results from the ‘Combined Test’ set are excluded and the ‘best dataset specific training pair’ for the unseen RobotCar and Inria datasets are taken as the ‘best average training pair’ as no information about these sets is available at training time in our experiments. We then provide more detailed discussion on the behaviours of our classifier and performance across the tested datasets as well as query level observations for where our approach does/does not perform better in Figure 3.

Table II demonstrates that our multi-label classification approach is able to dynamically select techniques which achieve the same or better recall than selecting a static technique pair based on the observed training data. Importantly this table of results shows that neither the best average or dataset specific training pairs (Columns 2 and 3) guaranteed the best pair observed in training generalised to the test sets. However, our classification network is able to at least equal the performance of the best of these statically chosen pairs. Furthermore, in the case of a dataset with a significantly changing environment, such as the ‘Combined Test’ set where the specific dataset is not known, our method provides a significant 6% improvement in recall by dynamically choosing techniques to pair with NetVLAD.

This advantage is emphasised by our methods ability to perform in unseen datasets, providing an additional 10% and 6% recall in a Day-Night traverse from Oxford RobotCar and the Inria Holidays dataset respectively.

B. Dynamic Technique Selection

Figure 4 illustrates our method’s ability to adapt its prediction of the best technique to pair with NetVLAD across significant environment, lighting and viewpoint changes contained within our combined test dataset. Interestingly, our model learnt to only predict three of the nine possible techniques to pair with NetVLAD (CoHOG, AMOSNet and DenseVLAD). We believe this is because the training data shows that there are only 78 of the 1493 training queries where NetVLAD paired with either of these three methods will not be successful whilst other technique pairs are. Therefore, NetVLAD paired with either CoHOG, AMOSNet or DenseVLAD successfully localises in $\approx 95\%$ of the training data. So whilst we assume the complementarity of different technique pairs to be independent in training, the performance of some technique pairs is likely correlated.

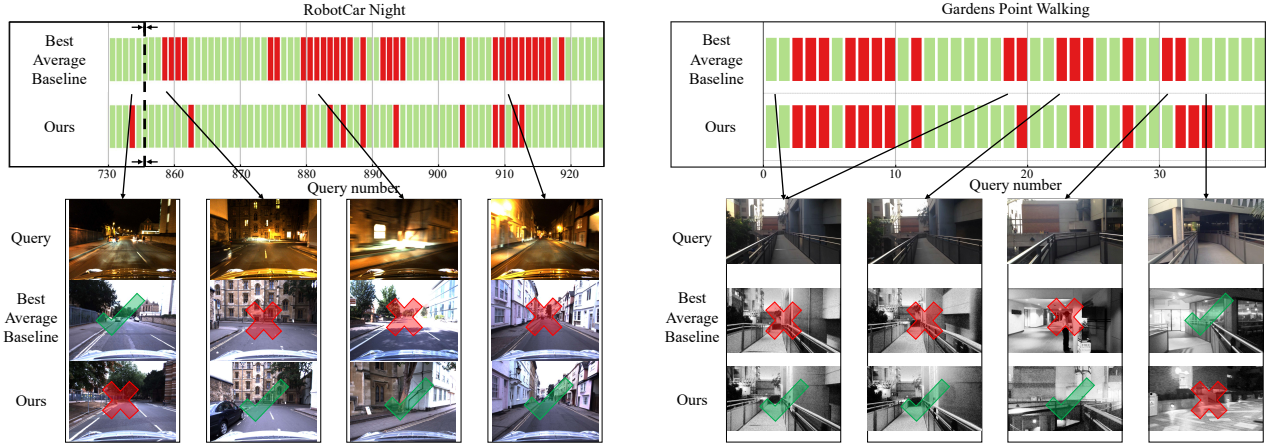


Fig. 3. **Correct/incorrect technique selections per query.** In this figure we show on a query level (each bar denotes an image) where our approach is able to select technique pairs to successfully localise when the baseline did not. We also include some failure examples from our network’s technique selection. Green bars denote successful VPR and red bars denote failure. We include two different datasets to demonstrate the extreme changes where we are able to use a single network for complementary technique selection. The datasets will be **left: Oxford RobotCar, right: Gardens Point Walking.** We include samples from the datasets to show some of the queries and corresponding top VPR matches where these improvements/failures occur.




Fig. 4. **Technique selections per query.** Here we show that in the combined test set (made from SFU Mountain, Gardens Point Walking, Nordland and Pittsburgh), our approach is able to use a single network to perform frame-by-frame predictions for the most complementary technique to pair with NetVLAD without any knowledge of the specific dataset queries are from.

C. Selection Consistency

Figures 4 & 5 show that within each separated dataset there is typically a particular technique which is primarily selected to pair with NetVLAD. This is not unexpected as environment changes are typically consistent within datasets, such as the seasonal changes in Nordland or the viewpoint changes in Pittsburgh. We recognise that our method is less effective in this scenario, however, when there is significant change within a dataset, such as our combined test set, our method is able to effectively predict the most complementary technique for NetVLAD despite these changes. Table II quantitatively shows that our method improves the mean recall by 3.5% over the best comparison (dataset specific baselines, col. 3).

VI. CONCLUSION AND FUTURE WORK

In this work we propose a multi-label classification approach to selecting the most complementary VPR technique on a frame-by-frame basis to add to a baseline technique in order to maximize place recognition performance using multi-process fusion. The results demonstrated that our approach was effective across datasets with significant environment, illumination, viewpoint and structural changes by providing equal or better recall than manually selected baselines, and that it generalized to two unseen environments. Whilst there are some limiting requirements such as the need to store features and models from multiple VPR networks,




Fig. 5. **Technique selections per dataset.** In this figure we show the percentage of each dataset where each of the technique pairs were selected using our classification network. We limit this to only include CoHOG, AMOSNet and DenseVLAD as we found our network learned it only needed to make predictions from these three techniques. As seen, there is typically a dominant technique pair which is predicted for individual datasets.

this approach offers an effective and computationally cheap method for more robust VPR.

This work opens up a number of opportunities for future research. One main observation from the experiments was that within individual datasets there was typically one technique pair primarily selected throughout the dataset. While this did result in better performance than the baselines, it was still below the performance of an “oracle” system (with access to ground truth) which would always choose the best pairing of techniques. This performance gap could potentially be bridged by specifically training input features to aid the multi-label classifier, or through developing loss arguments which encourage a more diverse choice of techniques.

Finally, while maximizing average performance across an entire dataset is a standard research goal, real-world deployments often prioritise other performance properties like optimizing for worst case performance, minimizing “time without a good localization match” and other functionally-related properties. Future work could build on the foundation established here to enable any functionally-driven performance goal to be targeted with this approach.

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 Transactions on Robotics*, vol. 32, no. 1, pp. 1–19, 2015.
- [2] C. Masone and B. Caputo, "A survey on deep visual place recognition," *IEEE Access*, vol. 9, pp. 19 516–19 547, 2021.
- [3] X. Zhang, L. Wang, and Y. Su, "Visual place recognition: A survey from deep learning perspective," *Pattern Recognition*, vol. 113, p. 107760, 2021.
- [4] S. Garg, T. Fischer, and M. Milford, "Where is your place, visual place recognition?" in *International Joint Conference on Artificial Intelligence*, 2021, pp. 4416–4425.
- [5] S. Schubert and P. Neubert, "What makes visual place recognition easy or hard?" *arXiv:2106.12671*, 2021.
- [6] R. Arandjelovic, P. Gronat, A. Torii, T. Pajdla, and J. Sivic, "NetVLAD: CNN Architecture for Weakly Supervised Place Recognition," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 40, no. 6, pp. 1437–1451, 2018.
- [7] N. Sünderhauf, P. Neubert, and P. Protzel, "Are we there yet? challenging seqslam on a 3000 km journey across all four seasons," in *Workshop on Long-Term Autonomy, IEEE International Conference on Robotics and Automation*, 2013.
- [8] N. Dalal and B. Triggs, "Histograms of oriented gradients for human detection," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2005, pp. 886–893.
- [9] S. Hausler and M. Milford, "Hierarchical Multi-Process Fusion for Visual Place Recognition," in *IEEE International Conference on Robotics and Automation*, 2020, pp. 3327–3333.
- [10] M. Zaffar, S. Garg, M. Milford, J. Kooij, D. Flynn, K. McDonald-Maier, and S. Ehsan, "VPR-Bench: An Open-Source Visual Place Recognition Evaluation Framework with Quantifiable Viewpoint and Appearance Change," *International Journal of Computer Vision*, vol. 129, pp. 2136–2174, 2021.
- [11] F. Han, X. Yang, Y. Deng, M. Rentschler, D. Yang, and H. Zhang, "SRAL: Shared Representative Appearance Learning for Long-Term Visual Place Recognition," *IEEE Robotics and Automation Letters*, vol. 2, no. 2, pp. 1172–1179, 2017.
- [12] S. Hausler, A. Jacobson, and M. Milford, "Multi-Process Fusion: Visual Place Recognition Using Multiple Image Processing Methods," *IEEE Robotics and Automation Letters*, vol. 4, no. 2, pp. 1924–1931, 2019.
- [13] M. J. Milford and A. Jacobson, "Brain-inspired sensor fusion for navigating robots," in *IEEE International Conference on Robotics and Automation*, 2013, pp. 2906–2913.
- [14] P. Zhang, J. Gu, E. E. Milios, and P. Huynh, "Navigation with IMU/GPS/digital compass with unscented Kalman filter," in *IEEE International Conference on Mechatronics and Automation*, 2005, pp. 1497–1502.
- [15] J.-w. Hu, B.-y. Zheng, C. Wang, C.-h. Zhao, X.-l. Hou, Q. Pan, and Z. Xu, "A survey on multi-sensor fusion based obstacle detection for intelligent ground vehicles in off-road environments," *Frontiers of Information Technology & Electronic Engineering*, vol. 21, no. 5, pp. 675–692, 2020.
- [16] M. Waheed, M. Milford, K. McDonald-Maier, and S. Ehsan, "Improving Visual Place Recognition Performance by Maximising Complementarity," *IEEE Robotics and Automation Letters*, vol. 6, no. 3, pp. 5976–5983, 2021.
- [17] C. Cadena, L. Carlone, H. Carrillo, Y. Latif, D. Scaramuzza, J. Neira, I. Reid, and J. J. Leonard, "Past, present, and future of simultaneous localization and mapping: Toward the robust-perception age," *IEEE Transactions on robotics*, vol. 32, no. 6, pp. 1309–1332, 2016.
- [18] A. Jacobson, Z. Chen, and M. Milford, "Leveraging variable sensor spatial acuity with a homogeneous, multi-scale place recognition framework," *Biological Cybernetics*, vol. 112, no. 3, pp. 209–225, 2018.
- [19] Z. Wang, Y. Wu, and Q. Niu, "Multi-sensor fusion in automated driving: A survey," *IEEE Access*, vol. 8, pp. 2847–2868, 2019.
- [20] H. Wei, X.-e. Li, Y. Shi, B. You, and Y. Xu, "Multi-sensor fusion glass detection for robot navigation and mapping," in *WRC Symposium on Advanced Robotics and Automation*, 2018, pp. 184–188.
- [21] S. Joshi and S. Boyd, "Sensor selection via convex optimization," *IEEE Transactions on Signal Processing*, vol. 57, no. 2, pp. 451–462, 2009.
- [22] S. Hausler, A. Jacobson, and M. Milford, "Filter Early, Match Late: Improving Network-Based Visual Place Recognition," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2019, pp. 3268–3275.
- [23] J. Maltar, I. Marković, and I. Petrović, "Visual place recognition using directed acyclic graph association measures and mutual information-based feature selection," *Robotics and Autonomous Systems*, vol. 132, p. 103598, 2020.
- [24] T. L. Molloy, T. Fischer, M. Milford, and G. N. Nair, "Intelligent Reference Curation for Visual Place Recognition Via Bayesian Selective Fusion," *IEEE Robotics and Automation Letters*, vol. 6, no. 2, pp. 588–595, 2021.
- [25] J. Mount, M. Xu, L. Dawes, and M. Milford, "Unsupervised Selection of Optimal Operating Parameters for Visual Place Recognition Algorithms Using Gaussian Mixture Models," *IEEE Robotics and Automation Letters*, vol. 6, no. 2, pp. 343–350, 2021.
- [26] G. Berton, C. Masone, and B. Caputo, "Rethinking visual geolocalization for large-scale applications," in *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 4878–4888.
- [27] S. Hussaini, M. Milford, and T. Fischer, "Spiking neural networks for visual place recognition via weighted neuronal assignments," *IEEE Robotics and Automation Letters*, vol. 7, no. 2, pp. 4094–4101, 2022.
- [28] S. Schubert, P. Neubert, and P. Protzel, "Unsupervised learning methods for visual place recognition in discretely and continuously changing environments," in *IEEE International Conference on Robotics and Automation*, 2020, pp. 4372–4378.
- [29] N. Sünderhauf, S. Shirazi, F. Dayoub, B. Uprocroft, and M. Milford, "On the performance of ConvNet features for place recognition," *IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 4297–4304, 2015.
- [30] A. Torii, J. Sivic, T. Pajdla, and M. Okutomi, "Visual place recognition with repetitive structures," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2013, pp. 883–890.
- [31] H. Jegou, M. Douze, and C. Schmid, "Hamming Embedding and Weak Geometric Consistency for Large Scale Image Search," in *European Conference on Computer Vision*, 2008, pp. 304–317.
- [32] W. Maddern, G. Pascoe, C. Linegar, and P. Newman, "1 Year, 1000km: The Oxford RobotCar Dataset," *The International Journal of Robotics Research*, vol. 36, no. 1, pp. 3–15, 2017.
- [33] J. Bruce, J. Wawerla, and R. Vaughan, "The SFU mountain dataset: Semi-structured woodland trails under changing environmental conditions," in *Workshop on Visual Place Recognition in Changing Environments, IEEE International Conference on Robotics and Automation*, 2015.
- [34] A. Khaliq, S. Ehsan, Z. Chen, M. Milford, and K. McDonald-Maier, "A Holistic Visual Place Recognition Approach Using Lightweight CNNs for Significant ViewPoint and Appearance Changes," *IEEE Transactions on Robotics*, vol. 36, no. 2, pp. 561–569, 2020.
- [35] M. Zaffar, S. Ehsan, M. Milford, and K. McDonald-Maier, "COHOG: A light-weight, compute-efficient, and training-free visual place recognition technique for changing environments," *IEEE Robotics and Automation Letters*, vol. 5, no. 2, pp. 1835–1842, 2020.
- [36] Z. Chen, A. Jacobson, N. Sünderhauf, B. Uprocroft, L. Liu, C. Shen, I. Reid, and M. Milford, "Deep learning features at scale for visual place recognition," in *IEEE International Conference on Robotics and Automation*, 2017, pp. 3223–3230.
- [37] N. Merrill and G. Huang, "Lightweight unsupervised deep loop closure," in *Robotics: Science and Systems*, 2018.
- [38] J. Revaud, J. Almazán, R. S. Rezende, and C. R. d. Souza, "Learning with average precision: Training image retrieval with a listwise loss," in *IEEE/CVF International Conference on Computer Vision*, 2019, pp. 5107–5116.
- [39] A. Torii, R. Arandjelovic, J. Sivic, M. Okutomi, and T. Pajdla, "24/7 Place Recognition by View Synthesis," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 40, no. 2, pp. 257–271, 2018.