

# Unsupervised Quality Prediction for Improved Single-Frame and Weighted Sequential Visual Place Recognition

Helen Carson, Jason J. Ford, Michael Milford

**Abstract**—While substantial progress has been made in the absolute performance of localization and Visual Place Recognition (VPR) techniques, it is becoming increasingly clear from translating these systems into applications that other capabilities like integrity and predictability are just as important, especially for safety- or operationally-critical autonomous systems. In this research we present a new, training-free approach to predicting the likely quality of localization estimates, and a novel method for using these predictions to bias a sequence-matching process to produce additional performance gains beyond that of a naive sequence matching approach. Our combined system is lightweight, runs in real-time and is agnostic to the underlying VPR technique. On extensive experiments across four datasets and three VPR techniques, we demonstrate our system improves precision performance, especially at the high-precision/low-recall operating point. We also present ablation and analysis identifying the performance contributions of the prediction and weighted sequence matching components in isolation, and the relationship between the quality of the prediction system and the benefits of the weighted sequential matcher.

## I. INTRODUCTION

Visual Place Recognition (VPR) is the task of estimating position using cameras, by matching an image of the current scene to a database of geo-tagged reference images taken along a previous traverse of the route [1]. Substantial progress has been made in this research area in recent years [2]–[4]. While absolute performance has often been the target, when adapting these systems to operate on deployed robots and autonomous systems, other capabilities around localization integrity become just as, if not more, important.

While substantial progress has been made in other domains including aerospace on topics like system integrity [5], this field is relatively underinvestigated in robotics. Initial pilot work in this area has demonstrated some promise, but has had substantial practical limitations including the need for extensive, multi-traverse training data with ground truth correspondences from the target domain. In this paper, we present a new VPR system that provides high performance whilst overcoming some of the practical limitations of previous approaches.

This research is partially supported by an ARC Laureate Fellowship FL210100156 to MM, the QUT Centre for Robotics, the Centre for Advanced Defence Research in Robotics and Autonomous Systems, and received funding from the Australian Government via grant AUS-MURIB000001 associated with ONR MURI grant N00014-19-1-2571. The work of H.Carson was supported in part by an Australian Postgraduate Award.

The authors are with the QUT Centre for Robotics, School of Electrical Engineering and Robotics at the Queensland University of Technology, Brisbane, Australia (e-mail: h.carson@hdr.qut.edu.au, j2.ford@qut.edu.au, michael.milford@qut.edu.au).

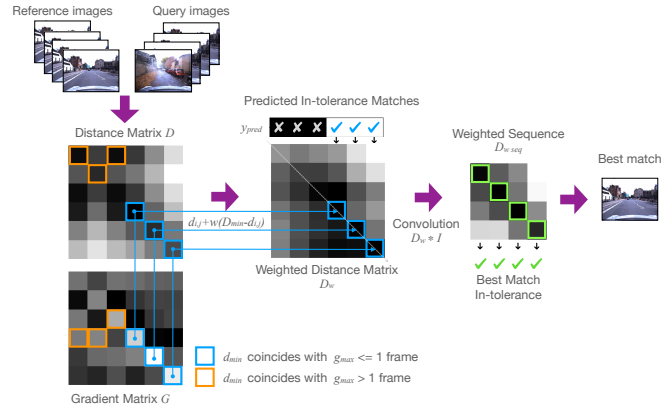


Fig. 1. Visual representation of our combined approach. For each query, matches are predicted to be in-tolerance when the distance matrix minimum and gradient matrix maximum coincide within one frame. These consensus matches are weighted heavily before applying sequence matching, resulting in a higher proportion of sequence matches that are in-tolerance.

We make the following contributions:

- 1) We present a novel, unsupervised method that enables out-of-tolerance matches to be identified and discarded during a post-processing step on various VPR techniques, using a consensus-based approach.
- 2) We use our localization quality prediction technique to anchor sequence-based matching around identifiably good localization points, further improving performance (Fig. 1).

To support our contribution claims, we demonstrate increased average precision over 15 combinations of benchmark datasets and VPR techniques, covering viewpoint variation, seasonal changes, and urban driving scenarios, using both simple and state-of-the-art VPR techniques.

We further provide analysis isolating the performance contributions of both the prediction system and the weighted sequence matching technique, as well as of the relationship between the performance of the prediction system and the subsequent effect on performance of the weighted sequence matcher.

The paper proceeds as follows. We provide background in Section II and describe our approach in Section III. Our experimental setup is outlined in Section IV, and results and discussion in Sections V and VI.

## II. BACKGROUND

In this section, we further define the problem, provide an overview of VPR techniques and related works, and

outline the differences between our methods and similar post-processing techniques.

### A. VPR Techniques

VPR techniques themselves take many forms, including direct image comparison, hand-crafted techniques, machine-learned feature extraction, hierarchical methods, and semantic segmentation [2]. All involve a method of describing images in the form of extracted features prior to comparison, ranging from simple downsampling and patch normalization [6], to complex techniques such as SIFT [7], SURF [8], ORB [9], FAB-MAP [10], Convolutional Neural Network (CNN) architectures [11] such as NetVLAD [12]; and hierarchical methods such as PatchNetVLAD [13]. A localization match is then formed by finding the minimum distance between features in the current scene and all reference images. This match however can be incorrect due to visual aliasing, changes in the environment, dynamic agents, lighting differences and shadows, and many other issues that cause images taken at the same place to look different, and different places to look the same [1].

Applying post-processing techniques to initial VPR estimates can result in large performance improvements [14]. These can include fusion with additional sensor data [15]; use of visual odometry to reject out-of-tolerance estimates [16]; temporal window-based filtering [17]; and Random Sample Consensus (RANSAC) to identify outliers once an error tolerance threshold is tuned [18]. Sequential filtering can also be used to exploit the rich spatio-temporal information provided as a localization system moves along a route, rather than treating each image independently [6], including learned descriptors based on contrastive learning methods such as SeqNet [19] and SeqMatchNet [20]. Extended Kalman Filters (EKFs) [21] and particle filters [22] can also be used to provide continually updated localization estimates, based on motion state-space estimation.

### B. Localization Integrity

We relate our work to the concept of integrity, which aims to provide a real-time determination of whether a current localization estimate is within an acceptable tolerance from ground-truth, or not [23]. As the overall performance of camera-based localization systems becomes higher, and these systems are able to perform reliably in more challenging environments, a system's ability to self-identify when its performance is poor is becoming increasingly important, especially in safety- or operationally-critical applications like autonomous vehicles [24].

In our previous work [25], we proposed a supervised learning method for predicting which matches along a route are in-tolerance, based on extracting features from the distance matrix. However, this supervised technique had substantial practical disadvantages: it required two traverses along a route through a representative environment, along with ground-truth data for each reference and query frame, to enable the prediction method to learn which points were likely out-of-tolerance.

Here we provide two complementary techniques: an unsupervised method for predicting in-tolerance matches, and a method of weighting sequences prior to matching based on these predictions to improve overall precision. In contrast to our previous work, this technique requires no calibration or training data. Our unsupervised method is effectively a consensus-based filter used to "check" the distance-based localization estimate from the VPR technique against a secondary and independent gradient-based estimate. Both of these techniques are suited to applications where high precision and low false positive rates are more important than recall - that is, where it is better to reject good points as bad, rather than the alternative: accepting out-of-tolerance estimates as good.

In the following sections we outline our approach and experimental method, and compare the performance of our proposed techniques to baseline VPR techniques.

## III. APPROACH

In this section we detail our approach to our two contributions, including the formulation of a distance matrix, our new gradient matrix, consensus process, and the exploitation of this process to create an improved weighted sequence matching approach.

### A. Unsupervised prediction

The cornerstone of our unsupervised method is based on the observation that the absolute average gradient of the similarity matrix is typically maximised around the best match. This is most evident when creating a similarity matrix of one set of images against itself: all exact matches produce an instantaneous distance of zero along the diagonal, corresponding with a sharp instantaneous rate of change in the distance vector at each matching point. Within the matrix, similar places may exhibit relatively low distances between each other, but intuitively, the absolute average gradient change should be highest around the closest match when the distance is closest to zero.

Our goal is to identify correct (in-tolerance) points with high precision, such that inaccurate points are discarded and only accurate points remain. Based on our observations, we form the hypothesis that points along a route where the minimum distance and maximum rate-of-change (gradient) of the distance vector coincide are more likely to be exact or very close matches, rather than where these two phenomenon occur at vastly different positions along the route.

1) *Formulation of Distance Matrix:* For completeness sake, here we briefly introduce our distance matrix calculation, formed by computing the euclidean or cosine distance between the features in the current scene (query image), to features in a set of reference images. The candidate best match is then selected as the reference frame with the minimum distance to each query along a route, representing the best estimated location. In this paper we use a distance matrix  $\mathbf{D} \in \mathbb{R}^{n \times m}$  as a representation of feature euclidean distance between all  $m$  queries and  $n$  references, and a distance vector  $\mathbf{d}_j = \mathbf{D}[:, j]$  as the distance between a single

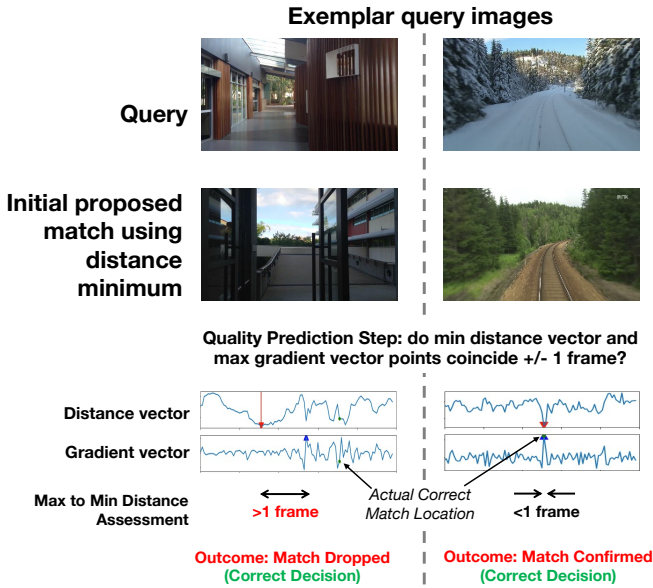


Fig. 2. Examples showing our prediction method applied to proposed “good” matches from a VPR technique. Consensus between the distance vector minimum (used by the VPR technique to generate the match) and gradient vector maximum is shown to predict a good quality (in-tolerance) match.

query image and all reference images, which is the method used in real-time implementation.

Each element  $\mathbf{D}_{ij}$  is the distance between the features in the  $i^{\text{th}}$  database image and the  $j^{\text{th}}$  query image, and the reference with the minimum distance,  $d_0$ , is found as the candidate best match where,  $i_{d0} = \text{argmin}(\mathbf{d}_j)$ . Lower values at the minimum point indicate a stronger match.

2) *Formulation of Gradient Matrix:* We introduce in this work a gradient matrix, which is formulated by computing a modified average gradient around each point along the distance vector, computed as:

$$g = \begin{cases} \frac{1}{2}(d_{i+1} + d_{i-1}) - d_i, & 1 < i < n \\ d_{i+1} - d_i, & i = 1 \\ d_{i-1} - d_i, & i = n \end{cases} \quad (1)$$

We use convolution to enhance the signal using a  $3 \times 3$  kernel, and pad the first two frames with the mean of the gradient vector. A potential negative impact of this in a real-time implementation may be decreased prediction accuracy for the first and second query frames.

A gradient vector,  $\mathbf{g}_j = \mathbf{G}[:, j]$  is formed for the  $j^{\text{th}}$  query, representing the rate of change between the query features, and features in each reference frame along the route. Note that the order of the reference frames must be preserved, but the query images can be encountered in any order in real-time. The maximum gradient for each query can then be hypothesised as an alternative best match candidate:

$$i_{g0} = \text{argmax}(\mathbf{g}_j) \quad (2)$$

3) *Rejection of out-of-tolerance points:* To finalise our technique, we compare the best match generated by the minimum distance vector and maximum gradient vector, and

reject matches where the number of frames between these are greater than one. We form a binary prediction vector  $\mathbf{y}_{pred} \in \{0, 1\}^m$  from this comparison, where  $y_{pred} = 1$  representation predicts the localization estimate is in-tolerance (True Positive) and  $y_{pred} = 0$  represents an out-of-tolerance prediction (False Positive):

$$y_{pred} = \begin{cases} 1, & |i_{g0} - i_{d0}| \leq 1 \\ 0, & \text{otherwise.} \end{cases} \quad (3)$$

The prediction vector can be used as a mask to retain only predicted good candidate matches from the distance matrix. Examples of this consensus based prediction approach are shown in Figure 2.

### B. Anchored/Weighted sequence

The methods described so far have the potential to improve single-frame-based matching, and could be further enhanced by naive application of a sequence-based matching technique. However, a key new opportunity here we pursue is to weight the sequence matching process based on the predictions around the quality of the localization estimates. Here we incorporate our method into SeqSLAM [6] by heavily weighting predicted good points prior to performing sequence matching. This artificially gives each sequence template containing predicted good points a much lower value than the standard implementation, which anchors the sequence around our prediction.

For each query, we weight each minimum value,  $d_0$ , in the distance vector corresponding to a predicted good point by a weighting factor,  $w$ , between 0 and 1 to create a weighted distance matrix,  $\mathbf{D}_w$ :

$$D_w|_{i=i_{d0}, j} = \begin{cases} d_0 - w(d_0 - D_{min}), & y_{pred} = 1 \\ D_{i,j}, & y_{pred} = 0 \end{cases} \quad (4)$$

The weighting factor  $w$  lowers the minimum value in the distance vector towards the overall minimum of the distance matrix. All other values, including predicted bad points ( $y_{pred} = 0$ ), are left unchanged.

To implement sequence matching, we use the computationally efficient method of convolving the weighted distance matrix with an identity matrix with sequence length  $L$ , as described in [20] and [26]:

$$\mathbf{D}_{wseq} = \mathbf{I}_L * \mathbf{D}_w, \quad (5)$$

and identify the new best candidate sequence for each query:

$$i_0 = \text{argmin}(\mathbf{D}_{wseq}[:, j]) \quad (6)$$

Using a weighting factor  $w = 1$  reduces all predicted ‘good’ values to the minimum value in the distance matrix,  $\mathbf{D}_{min}$ . To enable generation of PR-curves, a weighting factor  $w$  less than one (e.g.  $w = 0.99$ ) is recommended.

## IV. EXPERIMENTAL SETUP

To test our hypotheses, we implement our approach using 15 combinations of datasets and VPR techniques, described in the sections below, and compare the performance of our method against those of baseline techniques.

### A. VPR Techniques

We apply our method to Sum-of-Absolute Differences (SAD), using patch-normalized downsampled images [6]; NetVLAD [12] based on a CNN architecture; and recent state-of-the-art approach PatchNetVLAD, which uses fusion of locally-global descriptors on multiple-scales using pre-trained models [13].

### B. Datasets

Four widely used benchmark datasets have been used to implement our technique: Gardens Point Walking [27], Nordland [28], 4Seasons OfficeLoop [29] and Oxford RobotCar [30], [31]. In each case, we use one reference set and one query set taken on the same route for our unsupervised prediction. When benchmarking against our previous supervised prediction technique, we use a second calibration reference/query set over a different route as described in [25], along with ground-truth data for the calibration set.

Gardens Point Walking is a walking dataset obtained on a university campus with viewpoint variation between the reference and query sets. We use the first 100 frames with the left set as reference and right set as query, with 2.5m between each frame. Nordland contains images from a Norwegian railway, and uses a fixed viewpoint, but with large seasonal variation between reference and query sets. We have used two separate sets from this dataset: fall/spring and summer/winter as query/reference sets. The summer/winter set provides a challenging dataset due to significant differences in daylight levels, snow cover, and foliage along the route. In both cases we use 1150 images over 35km from Test Section 2 as defined in [28], with mean distance of 30m between consecutive images, and tunnels and station stops removed.

4Seasons and Oxford RobotCar contain images recorded from passenger vehicles driving in urban environments. Challenges in these sets include presence of dynamic objects, viewpoint variations along the route (e.g. when changing lanes) and environmental differences. In 4Seasons, we use 381 frames for reference and query routes over 500m with little variation apart from dynamic objects on the road, providing an ‘easy’ set that all chosen VPR techniques work well on. For Oxford RobotCar, we use 1800 image frames over approximately 2.25km recorded in sunny conditions as the reference set, in rain for the query set, with 1.25m median distance between frames [32].

Together, the datasets chosen provide a broad set of conditions and challenges to assess the performance of our method.

## V. RESULTS

We present performance results for the overall combination of our two proposed contributions in Section V-A, and the results of our two contributions independently in Sections V-B.1 and V-B.2. We also present an ablation study showing the effect of sequence length and weighting factor on the weighted sequence performance, and experimental results for mean execution time at inference.

### A. Combined system

Precision and recall are commonly used performance measures defined as ratios of in-tolerance localization matches (True Positives), out-of-tolerance matches (False Positives), and False Negatives (FN) (good matches that have been discarded), such that  $\text{precision} = \frac{TP}{TP+FP}$  and  $\text{recall} = \frac{TP}{TP+FN}$  [1]. Area Under the PR-curve (AUC) is a proxy for average precision, with higher values corresponding to higher performance. We impose a tight tolerance of +/-1 image frame as a “good” (in-tolerance) match in all experiments.

Table I includes AUC up to 20% recall for baseline sequences, and sequences weighted with our unsupervised predictions. In 13 out of 15 combinations our weighted sequence has higher or equal AUC than the baseline. A common weighting factor of 0.99 and sequence length of 2 is used in all cases.

PR-curves for all three VPR techniques applied to each dataset are shown in Figure 3 for the baseline sequences using no weighting, weighted sequences using the unsupervised prediction method, and the supervised method from our previous work. In the majority of cases, the PR-curve using weighted sequences has higher precision for the same recall as the unweighted sequences in the higher precision range for each dataset, which is the operating point of interest for safety-critical scenarios. The performance using weighted sequences also matches or exceeds the precision of our supervised technique in several cases, albeit up to a lower recall.

TABLE I  
COMPARISON OF AUC UP TO 20% RECALL FOR BASELINE (BL) AND WEIGHTED UNSUPERVISED SEQUENCES (OURS) WITH SEQUENCE LENGTH 2 (HIGHER IS BETTER).

Dataset	SAD		NetVLAD		PatchNetVLAD	
	BL	Ours	BL	Ours	BL	Ours
GardensPoint	0.54	<b>0.68</b>	<b>0.87</b>	<b>0.87</b>	0.92	<b>0.94</b>
Nordland-fs	0.88	<b>0.98</b>	0.62	<b>0.74</b>	0.95	<b>0.99</b>
Nordland-sw	0.66	<b>0.89</b>	0.27	<b>0.35</b>	<b>0.74</b>	<b>0.74</b>
4seasons	<b>0.99</b>	0.96	<b>0.99</b>	0.98	0.97	<b>0.98</b>
RobotCar	0.50	<b>0.82</b>	0.78	<b>0.81</b>	0.88	<b>0.92</b>
Average	Baseline 0.77			Ours <b>0.84</b>		

### B. Performance of individual contributions

The following sections present an analysis of results for our two contributions independently.

1) *Unsupervised Prediction*: We compare our unsupervised method to the following: (1) baseline VPR technique, sweeping the minimum distance as a threshold, and (2) applying a supervised (rather than unsupervised) prediction to the baseline VPR technique, and (3) generating matches using the maximum value in the gradient matrix only, and sweeping the maximum gradient as a threshold.

Fig. 4 demonstrates the impact of these comparisons as averaged PR-curves for the Nordland fall/spring and 4Seasons datasets, and shows the increased precision over the baseline when supervised and unsupervised predictions are

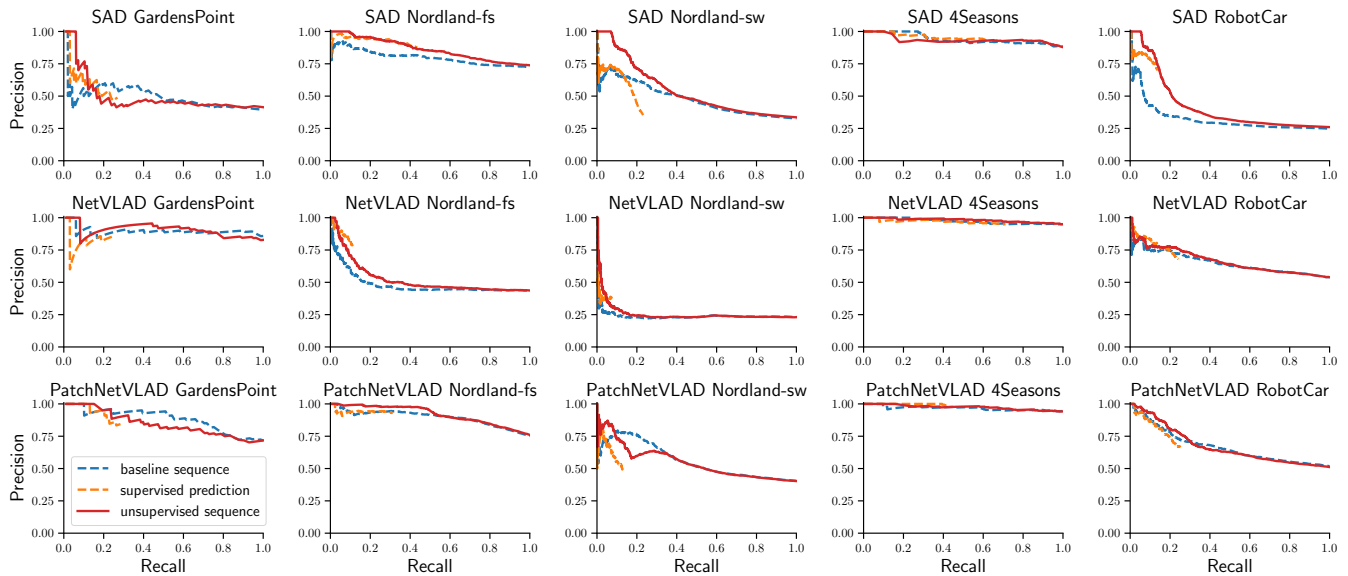


Fig. 3. PR-curves showing that in most cases, the performance impact of our proposed unsupervised approach using weighted sequences (combined system) is higher at high precision/low recall points compared to using a baseline (unweighted) sequence, over 15 combinations of varying datasets and VPR techniques. We also show comparison to the supervised prediction technique (without sequencing), from our previous work.

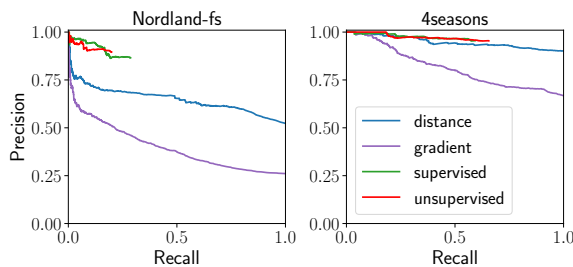


Fig. 4. PR-curves for all three VPR techniques averaged over the Nordland fall/spring and 4Seasons datasets, showing the superior performance using the minimum distance (baseline, blue) versus the maximum gradient only (purple) to determine best matches, and also the improved precision and impact on recall by retaining only ‘good’ matches using both supervised prediction (green) and our unsupervised prediction (red).

used to remove ‘bad’ matches. The curves are truncated to a maximum recall value due to many matches being discarded, such that 100% recall cannot be achieved.

These results also show that for these two datasets, using the maximum gradient alone to determine the best match produces decreased performance compared to the baseline, which uses the minimum distance value. Increased precision is only achieved when the match with the minimum distance and maximum gradient coincide (within  $\pm 1$  one frame), which is the basis for our unsupervised technique.

2) *Weighted Sequences*: The performance of the weighted sequence is directly impacted by the prediction quality, chosen weighting factor, and is also bounded in part by the performance of the underlying VPR technique in the environment. A break-even point will exist for each combination of these characteristics such that below a particular prediction quality threshold, use of the weighted sequence may actually *decrease* the combined system performance.

The weighting factor,  $w$ , tempers the impact of the prediction on performance, and can be chosen based on the estimated performance quality of a given prediction method and dataset / VPR technique combination.

Figure 5 shows the increase in performance by using a weighted sequence, where the unsupervised method (from Section III-A.3) is used to weight the distance matrix at predicted ‘good’ locations, for two datasets using SAD. The potential increase on this specific route and conditions using a perfect prediction is also shown to illustrate how, even with unrealistic perfect prediction, upper bound performance is limited by the underlying system.

To highlight the impact of this precision increase on the number of false positive matches, we can examine a specific example from our results. In Fig. 5(b), at 20% recall, the number of false positives (incorrect matches) in the GardensPoint SAD set is 37 for the unweighted sequence but drops to only 17 for the weighted sequence using unsupervised predictions.

Finally, Figure 6 presents a characterisation of the impact of the prediction performance over a range of weighting factors, for two different scenarios, a good prediction system and a poor prediction system, using the Oxford Robotcar dataset. The thicker line in the middle of the graph shows an unweighted sequence matcher as a baseline. *Below* the unweighted line, we have a badly performing predictor, showing that as the weighting factor goes up, the performance gets progressively worse. Above the unweighted line, we have a good predictor; we can see as the weighting factor goes up in this case, the overall performance goes up.

3) *Ablation Study*: An ablation study of weighted sequences for varying sequence lengths shows that the performance increase persists as sequence length increases, but the relative improvement over the baseline decreases. Examples

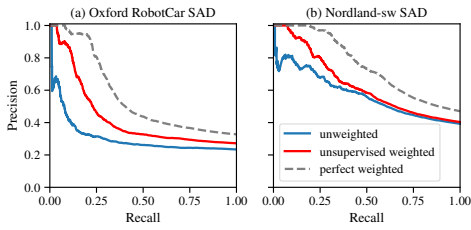


Fig. 5. PR-curves showing the performance improvement of weighted sequences (using unsupervised predictions) over unweighted sequences (baseline), and the potential combined system performance (dashed) that could be obtained with perfect prediction on this specific route (with sequence length = 3) is shown for comparison.

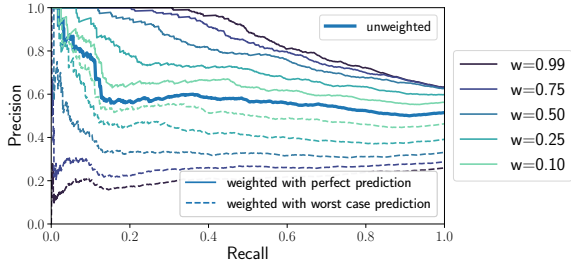


Fig. 6. PR-curves characterizing the relationship between performance prediction and weighting factor for one example route in the Oxford RobotCar dataset using SAD, for both a high quality and poor quality prediction system. Above the unweighted thicker middle line, we have a high quality predictor: as weighted use of the predictions goes up, performance goes up. Below the unweighted line, we have a poor quality predictor: as weighting goes up, performance actually decreases.

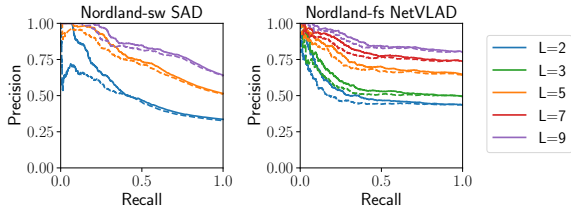


Fig. 7. PR-curves for Nordland-fall/spring dataset with NetVLAD, showing the improved performance for weighted sequences (solid) vs baseline (unweighted) sequences (dashed) for various sequence lengths, with constant weighting factor.

are given in Figure 7 for two combinations of Nordland datasets with varying sequence lengths from  $L=2$  to  $L=9$ .

### C. Computational resources

To evaluate whether our methods are suitable for real-time implementation, we analyse scalability and present experimental results showing mean execution time for our method.

1) *Theoretical explanation of scalability:* At run-time, our prediction method requires generation of a gradient vector for each query point. Both the prediction and weighted sequence filter require convolution to be performed along the gradient and distance vectors. Execution time will therefore increase linearly with reference database size, as the same arithmetic operations are performed for each reference image.

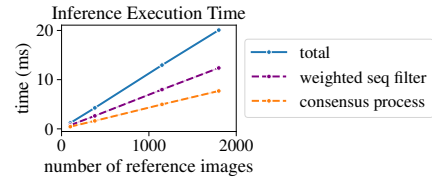


Fig. 8. Mean execution time at inference showing the separate contributions of our two proposed methods, suggesting the method is suitable for real-time applications (under 20ms for up to 1800 reference frames,  $L=3$ ).

2) *Execution time at inference:* Mean execution time for real-time implementation in Python has been measured running on an Apple Silicon M1 chip with 16GB RAM. Figure 8 shows the total mean execution time required for the unsupervised prediction and weighted sequence at inference is less than 20ms per query for datasets presented here up to 1800 reference frames. This includes the time required to implement the sequence filter, which is common with the baseline sequence technique. Note that our implementation is not optimised and can be considered a worst case. Results in Fig 8 provide evidence that our proposed methods are suitable for implementation in real-time applications.

## VI. DISCUSSION

We have presented two complementary methods for improving VPR performance based on: (1) unsupervised prediction of ‘good’ localization estimates along a route, and (2) using sequences weighted with predictions. Both methods improve performance but synergize in particular because the predictions enable the weighted nature of the modified sequence matching technique. Over a comprehensive combination of benchmark datasets and VPR techniques, results show that both techniques lead to improvements in localization performance. A limitation is that these techniques are best suited to applications where high precision, low recall operating points are acceptable.

There are a number of exciting avenues for future work. While the unsupervised nature of the approach presented here has practical advantages, there are a subset of applications where supervised training of a prediction system is feasible. In such situations, it would be interesting to examine how fusing both supervised and unsupervised methods could further improve both performance and versatility. In particular, an unsupervised method could potentially provide an online training signal for a supervised method. In addition, while difference and gradient-based matching and prediction has been shown to be effective here, it may be possible to learn more abstract “factors” that provide further predictive capability around localization performance. Finally, while we have used prediction outputs to weight the sequence matching process, it may be possible to more organically link the prediction processes and weighted sequential matching processes into a single monolithic system, with further potential advantages.

## REFERENCES

- [1] S. Lowry, N. Sunderhauf, P. Newman, J. J. Leonard, D. Cox, P. Corke, and M. Milford, "Visual place recognition: A survey," *IEEE Transactions on Robotics*, vol. 32, pp. 1–19, Feb. 2016.
- [2] P. Yin, S. Zhao, I. Cisneros, A. Abuduweili, G. Huang, M. Milford, C. Liu, H. Choset, and S. Scherer, "General place recognition survey: Towards the real-world autonomy age," 2022. [Online]. Available: <https://arxiv.org/abs/2209.04497>
- [3] Y. D. Yasuda, L. E. G. Martins, and F. A. Cappabianco, "Autonomous visual navigation for mobile robots: A systematic literature review," *ACM Computing Surveys*, vol. 53, pp. 1–34, 2020.
- [4] X. Zhang, L. Wang, and Y. Su, "Visual place recognition: A survey from deep learning perspective," *Pattern Recognition*, May 2021.
- [5] *DO-229 Minimum Operational Performance Standards for GPS/SBAS Airborne Equipment*, RTCA Std., 2020.
- [6] M. J. Milford and G. F. Wyeth, "Seqslam: Visual route-based navigation for sunny summer days and stormy winter nights," in *2012 IEEE International Conference on Robotics and Automation*, May 2012, pp. 1643–1649.
- [7] D. G. Lowe, "Object recognition from local scale-invariant features," *Proceedings of the Seventh IEEE International Conference on Computer Vision*, pp. 1150–1157 vol.2, 1999.
- [8] H. Bay, A. Ess, T. Tuytelaars, and L. Van Gool, "Speeded-up robust features (surf)," *Computer Vision and Image Understanding*, vol. 110, no. 3, pp. 346–359, 2008, similarity Matching in Computer Vision and Multimedia. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S1077314207001555>
- [9] E. Rublee, V. Rabaud, K. Konolige, and G. Bradski, "Orb: An efficient alternative to sift or surf," in *2011 International Conference on Computer Vision*, 2011, pp. 2564–2571.
- [10] M. J. Cummins and P. Newman, "Fab-map: Appearance-based place recognition and mapping using a learned visual vocabulary model," in *ICML*, 2010.
- [11] Z. Chen, A. Jacobson, N. Sunderhauf, B. Upcroft, L. Liu, C. Shen, I. Reid, and M. Milford, "Deep learning features at scale for visual place recognition," in *2017 IEEE International Conference on Robotics and Automation (ICRA)*, 2017, pp. 3223–3230.
- [12] R. Arandjelović, 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.
- [13] S. Hausler, S. Garg, M. Xu, M. Milford, and T. Fischer, "Patch-netvlad: Multi-scale fusion of locally-global descriptors for place recognition," in *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, Jun. 2021, pp. 14 136–14 147.
- [14] S. Schubert and P. Neubert, "What makes visual place recognition easy or hard?" Jun. 2021, arXiv:2106.12671.
- [15] K. A. Tsintotas, L. Bampis, and A. Gasteratos, "Tracking-doseqslam: A dynamic sequence-based visual place recognition paradigm," *IET Computer Vision*, vol. 15, pp. 258–273, 6 2021.
- [16] T. Kanji, "Multi-model hypothesize-and-verify approach for incremental loop closure verification," in *2016 IEEE/SICE International Symposium on System Integration (SII)*. IEEE, 2016, pp. 762–767.
- [17] J. Bruce, A. Jacobson, and M. Milford, "Look no further: Adapting the localization sensory window to the temporal characteristics of the environment," 6 2017. [Online]. Available: <http://arxiv.org/abs/1706.05772><http://dx.doi.org/10.1109/LRA.2017.2724146>
- [18] M. A. Fischler and R. C. Bolles, "Random sample consensus: A paradigm for model fitting with applications to image analysis and automated cartography," *Commun. ACM*, vol. 24, no. 6, p. 381–395, jun 1981. [Online]. Available: <https://doi.org/10.1145/358669.358692>
- [19] S. Garg and M. Milford, "Seqnet: Learning descriptors for sequence-based hierarchical place recognition," *IEEE Robotics and Automation Letters*, vol. 6, pp. 4305–4312, 7 2021.
- [20] S. Garg, M. Vankadari, and M. Milford, "Seqmatchnet: Contrastive learning with sequence matching for place recognition & relocalization," in *5th Annual Conference on Robot Learning*, 2021. [Online]. Available: <https://openreview.net/forum?id=OQMXb0xiCrt>
- [21] H. T. Tran, T. C. Vo, D. L. Tran, Q. N. Nguyen, D. M. Ha, Q. N. Pham, T. Q. Le, T. K. Nguyen, H. T. Do, and M. T. Nguyen, "Extended kalman filter (ekf) based localization algorithms for mobile robots utilizing vision and odometry," in *2022 IEEE 21st Mediterranean Electrotechnical Conference (MELECON)*, 2022, pp. 91–96.
- [22] P. Karkus, D. Hsu, and W. S. Lee, "Particle filter networks with application to visual localization," 2018. [Online]. Available: <https://arxiv.org/abs/1805.08975>
- [23] J. A. Hage, P. Xu, P. Bonnifait, J. Ibañez-Guzmán, and J. Ibanez-Guzman, "Localization integrity for intelligent vehicles through fault detection and position error characterization," *IEEE Transactions on Intelligent Transportation Systems*, pp. 1–13, 2020.
- [24] C. Zhu, M. Meurer, and C. Günther, "Integrity of visual navigation—developments, challenges, and prospects," *Navigation, Journal of the Institute of Navigation*, vol. 69, 6 2022.
- [25] H. Carson, J. J. Ford, and M. Milford, "Predicting to improve: Integrity measures for assessing visual localization performance," *IEEE Robotics and Automation Letters*, vol. 7, no. 4, pp. 9627–9634, 2022.
- [26] T. Fischer and M. Milford, "How many events do you need? event-based visual place recognition using sparse but varying pixels," 2022, arXiv:2206.13673v2.
- [27] Z. Chen, A. Jacobson, U. M. Erdem, M. E. Hasselmo, and M. Milford, "Multi-scale bio-inspired place recognition," in *2014 IEEE International Conference on Robotics and Automation (ICRA)*, 2014, pp. 1895–1901.
- [28] D. Olid, J. M. Fácil, and J. Civera, "Single-view place recognition under seasonal changes," in *PPNIV Workshop at IROS 2018*, 2018, pp. 101–106.
- [29] P. Wenzel, R. Wang, N. Yang, Q. Cheng, Q. Khan, L. von Stumberg, N. Zeller, and D. Cremers, "4Seasons: A cross-season dataset for multi-weather SLAM in autonomous driving," in *Proceedings of the German Conference on Pattern Recognition (GCPR)*, 2020.
- [30] 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. [Online]. Available: <http://dx.doi.org/10.1177/0278364916679498>
- [31] W. Maddern, G. Pascoe, M. Gadd, D. Barnes, B. Yeomans, and P. Newman, "Real-time kinematic ground truth for the oxford robotcar dataset," *arXiv preprint arXiv: 2002.10152*, 2020. [Online]. Available: <https://arxiv.org/pdf/2002.10152>
- [32] 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.