

Delta Descriptors: Change-Based Place Representation for Robust Visual Localization

Sourav Garg¹, Ben Harwood², Gaurangi Anand³ and Michael Milford¹

Abstract—Visual place recognition is challenging because there are so many factors that can cause the appearance of a place to change, from day-night cycles to seasonal change to atmospheric conditions. In recent years a large range of approaches have been developed to address this challenge including deep-learnt image descriptors, domain translation, and sequential filtering, all with shortcomings including generality and velocity-sensitivity. In this paper we propose a novel descriptor derived from tracking changes in *any* learned global descriptor over time, dubbed *Delta Descriptors*. Delta Descriptors mitigate the offsets induced in the original descriptor matching space in an unsupervised manner by considering temporal differences across places observed along a route. Like all other approaches, Delta Descriptors have a shortcoming - volatility on a frame to frame basis - which can be overcome by combining them with sequential filtering methods. Using two benchmark datasets, we first demonstrate the high performance of Delta Descriptors in isolation, before showing new state-of-the-art performance when combined with sequence-based matching. We also present results demonstrating the approach working with four different underlying descriptor types, and two other beneficial properties of Delta Descriptors in comparison to existing techniques: their increased inherent robustness to variations in camera motion and a reduced rate of performance degradation as dimensional reduction is applied. Source code is made available at <https://github.com/oravus/DeltaDescriptors>.

I. INTRODUCTION

Visual Place Recognition (VPR) is one of the key enablers for mobile robot localization and navigation. The earlier approaches to VPR predominantly relied on hand-crafted local (SIFT [1], ORB [2]) and global (HoG [3], GIST [4]) image representation methods. The use of local features in BoW [5] and VLAD [6] like encoding techniques based on visual vocabularies has been a popular choice for the task of global place retrieval (kidnapped robot problem). However, the lack of appearance robustness of underlying local features led to the development of appearance-invariant whole-image description techniques, combined with sequence-based matching [7], [8] in order to deal with extreme appearance variations.

Recent advances in deep learning have led to more robust counterparts to hand-crafted local and global feature representations like LIFT [9], DeLF [10], NetVLAD [11],

This work was supported by an AOARD grant: FA2386-19-1-4079.

¹Sourav Garg and Michael Milford are with School of Electrical Engineering and Robotics, Australian Centre for Robotic Vision (ACRV) and the QUT Centre for Robotics (QCR), Queensland University of Technology, Australia. s.garg@qut.edu.au

²Ben Harwood is with Robotics and Autonomous Systems, Data61 CSIRO, Australia and the Department of Electrical and Computer Systems Engineering, Monash University, Australia.

³Gaurangi Anand is with School of Computer Science, Queensland University of Technology, Australia.

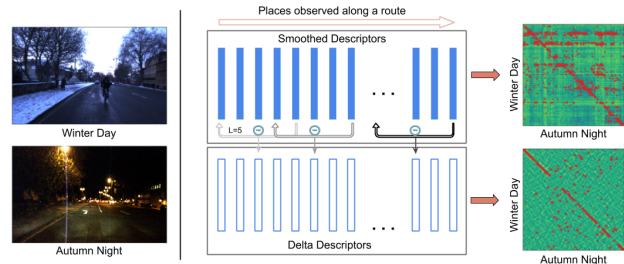


Fig. 1. We propose Delta Descriptor, defined as a high-dimensional signed vector of change measured across the places observed along a route. Using a difference-based description, places can be effectively recognized despite significant appearance variations. The cosine distance matrix between the Winter Day and Autumn Night traverse of the Oxford Robotcar dataset [17] is displayed on the right with red markers indicating the predicted matches and the matrix diagonal representing its ground truth.

and LoST [12]. Furthermore, GANs [13] based night-to-day image translation [14], feature fusion [15], and teacher-student networks [16] have also been explored for VPR. Although achieving state-of-the-art performance, deep learning based image descriptors often suffer from the challenges of data bias that limits their utility for out-of-the-box operations. This is typically resolved by re-training or fine-tuning the CNN. However, this may not always be feasible for all the application scenarios like VPR: supervised learning would require multiple traverses of the new environment such that it captures variations in scene appearance and camera viewpoint.

In this paper, we propose an unsupervised method for transforming existing deep-learnt image descriptors into change-based representations, dubbed *Delta Descriptors*. Due to appearance variations in the environment during a revisit, the original descriptors tend to be offset by a margin, leading to an increased distance between descriptors belonging to the same place. Defined in a difference space [18], Delta Descriptors implicitly deal with such offsets, resulting in a reduced distance between the same place descriptors. This is of particular importance to a mobile robot or autonomous vehicle operating in a new environment that may undergo significant variations in visual appearance during a repeated route traversal. Previously, difference-based approaches have been explored for recognizing objects [18], faces [19] and actions [20]. In this research, the concept of difference-based description is based on measuring the changes across an observed sequence of places which is repeatable across traverses.

Using the proposed Delta Descriptors, we:

- establish that given a fixed sequential span, Delta De-

scriptors perform at par with sequence-based matching, and

- that they achieve state-of-the-art performance when used in conjunction with the sequence-based matching,
- show Delta Descriptors retain high performance when dimensional reduction techniques like PCA are used, in contrast to raw descriptors, especially in the presence of strong appearance variations,
- demonstrate their robustness to variations in camera motion both along and across the repeated traverses, unlike sequence-based matching methods which either require motion information or a sophisticated sequence-search approach, and
- provide insights into selecting sequential span sizes for calculating Delta Descriptors, the role of order in which places are observed, an investigative Multi-Delta Descriptor approach to deal with velocity variability, and image-level visualization of how the proposed descriptors aid in time-series pattern matching.

II. RELATED WORK

A. Hand-crafted and Deep-Learnt Descriptors

The ability to represent images as a compact descriptor remains a key requirement for VPR. Broadly speaking, the purpose of these descriptors is to map images of a particular physical location into a lower dimensional representation. In the context of VPR, the goal of these mappings is to preserve a unique description of each location while removing changing information such as camera orientation, ambient lighting, mobile distractors and seasonal variations.

In the era of predominantly hand-crafted descriptors, the ORB [2] descriptor was designed to utilise a training set of local image patches. These patches were used to generate a set of binary tests that would maximise rotational invariance. While this learning procedure enabled the creation of more computationally efficient features, it also introduced the usual biases that come from training a model on a particular data set.

With the advent of deep learning, hand-crafted descriptors such as SIFT [1] have largely been replaced by learned descriptors such as LIFT [9] and DeLF [10]. Learning these local patch-based representations end-to-end has again yielded improved descriptors, but has also increased the reliance on having training data that accurately models all aspects of the target domain. This data bias is also seen in learned global representations such as NetVLAD [11] and AMOSNet [21].

B. Sequence-based Representations

A vast literature exists for spatio-temporal representation of video data with applications in action classification [22], activity recognition [23], person re-identification [24], dense video captioning [25], 3D semantic labelling [26], 3D shape completion [27], and dynamic scene recognition [28]. This has led to the emergence of 3D CNNs [29], [30] involving 3D convolutions to learn better spatio-temporal representations that are suited to the task at hand. However, most of the aforementioned tasks only require dealing with a limited

number of classes unlike VPR where every other place is a unique observation. Furthermore, methods based on RNNs, LSTM networks [31] and GRUs [32] tend to learn general patterns of how temporal information is ordered, for example, in the case of action or activity recognition. For VPR, such general patterns of order may not exist beyond the overlapping visual regions around a particular visual landmark.

In the context of VPR, there have been some attempts towards developing robust spatio-temporal or sequence-based representations. In [33], authors learnt spatio-temporal landmarks based on SURF features, but these local features needed to be independently tracked on frame-to-frame basis. In [34], authors proposed a bio-inspired place recognition method that used environment-specific discriminative training of different Long-Term Memory (LTM) cells. [35] explored three novel techniques: Descriptor Grouping, Descriptor Fusion, and Recurrent Descriptors, to accrue deep features from multiple views of the scene across the route. [36] proposed a topometric spatio-temporal representation of places using monocular depth estimation, mainly focused on recognizing places from opposing viewpoints. [37] proposed coresets-based visual summarization method for efficient hierarchical place recognition. Doing away with a compact representation, [38] used a variety of image descriptors to represent groups of images and formulated sequence-searching as an optimization problem.

C. Sequence-based Matching

Although sequence-based representations are not that common in VPR literature, use of sequence-based matching has been extensively explored for VPR. Such methods leverage sequential information *after* computing the place matching scores where place representations are typically based on a single image. This leads to enhanced VPR performance [39], particularly in cases where perceptual aliasing is very high, for example, dealing with extreme appearance variations caused by day-night and seasonal cycles, using methods like SeqSLAM [7] and SMART [40]. The follow up work in this direction is comprised of a number of methods that deal with camera velocity sensitivity [8], [41] or velocity estimation [42]. More recent work includes using temporal information and diffusion process within graphs [43], multi-sequence maps based VPR [44], and trajectory attention-based learning for SLAM [45]. In this paper, we use a simplified sequence-based matching technique mainly to analyse performance dynamics in using temporal information in two very different ways: sequential representation and sequential matching.

D. Difference-based Representations

The concept of using difference-based representation has been explored in a few different ways. [18] proposed Generalized Difference Subspace (GDS) as an extension of a difference vector for analyzing shape differences between objects. [19] proposed a novel discriminant analysis based on GDS demonstrating its utility as discriminative feature extractor for face recognition. Recently, [20] extended the concept of




Fig. 2. For five selected descriptor dimensions (across rows), time-series of observed places for two traverses from Oxford Robotcar dataset are displayed for ℓ_2 -normalized Raw (left), Smoothed (middle) and Delta Descriptors (right). The latter two are computed using a 64-frame window. The time-series pairs (blue and orange) should ideally be well aligned with each other.

GDS to tensors for representing and classifying gestures and actions. [46] used difference subspace analysis to maximize inter-class discrimination for effective face recognition as an alternative approach to improving representation ability of samples. [47] proposed a human action recognition method based on difference information between spatial subspace of neighboring frames. [23] used sum of depth difference between consecutive frames to discriminate moving/non-moving objects in order to detect humans. Our proposed method is based on descriptor difference and in essence solves the problem of recognition in a similar way as the GDS-based methods. In particular, Delta descriptors enable dealing with the offset that occurs in the deep-learned place representations when a robot is operating in new environments under significantly different environmental conditions.

III. PROPOSED APPROACH

A vast majority of existing place representation methods use single-image based descriptors for place recognition, typically followed by sequential matching or temporal filtering performed over place-matching scores. In this paper, we follow an alternative approach to place representation and propose Delta Descriptors that leverage the sequential information by measuring the *change* in descriptor as different places are observed over time. We hypothesize that these changes are both unique and consistent across multiple traverses of the environment, despite significant variations in scene appearance. In particular, measuring the change inherently ignores the data bias of the underlying deep-learned image descriptors, elevating the latter’s utility under diverse environment types and appearance conditions.

In this section, we first highlight key observations from the time-series of existing state-of-the-art image descriptors, then define and formulate the Delta Descriptors, and finally, describe an alternative convolutions-based approach to compute the Delta Descriptors more efficiently.

a) Key Observations: In the context of VPR for a mobile robot, images are typically captured as a data stream and converted into high-dimensional descriptors. We consider this stream of image descriptors as a multi-variate time-series.

For some of the descriptor dimensions¹, Figure 2 shows pairs of time-series for first 600 images (with 0.5 meters frame separation) from two different traverses of Oxford Robotcar dataset captured under day and night time conditions respectively. For the Raw descriptors, it can be observed that the consecutive values in the time-series tend to vary significantly even though the adjacent frames have high visual overlap. Moreover, the local variations in the descriptor values are not consistent across the traverses albeit the global time-series patterns appear repeatable. As the underlying deep-learned image descriptors (in most cases) are not trained to be stable against slight perturbations in camera motion or for ignoring the dynamic objects in the scene, such local variations are an expected phenomenon.

b) Defining Delta Descriptors: With these observations, we define delta descriptor, Δ_t , as a high-dimensional signed vector of change measured across a window of length l over a *smoothed* multi-variate time-series, $\mathbf{X}_t \in \mathcal{R}^D$, where t represents the time instant for an observation of a place along a route in the form of an image descriptor. To be more specific, we have

$$\Delta_t = \bar{\mathbf{X}}_{t+l/2} - \bar{\mathbf{X}}_{t-l/2} \quad (1)$$

where $\bar{\mathbf{X}}_t$ represents the smoothed signal obtained from a rolling average of the time series \mathbf{X}_t :

$$\bar{\mathbf{X}}_t = \sum_{t'=t-l/2}^{t+l/2} \mathbf{X}_{t'} / l \quad (2)$$

The middle and the right graphs in Figure 2 show the smoothed time-series and the corresponding Delta Descriptor respectively. It can be observed that the proposed descriptors are much more aligned than the baseline ones, getting rid of the offset in their original values.

c) Simplified Implementation: The formulation for Delta Descriptors presented above is suitable for understanding and visualizing the time-series patterns. However, Equation 1 and 2 can be simplified to a convolution-based calculation of the proposed descriptors:

$$\Delta = \mathbf{X} * \mathbf{W} \quad (3)$$

where convolutions are performed along the time axis of the baseline descriptor, independently per dimension using a 1D convolutional filter $\mathbf{W} = (w_1, w_2, \dots, w_L)$ defined as a vector of length $L = 2l$:

$$w_i = \begin{cases} -1/l, & \text{if } i \leq L/2, \\ 1/l, & \text{otherwise.} \end{cases} \quad (4)$$

For performing visual place recognition, the proposed descriptors are matched using cosine distance². However, for visualization purposes as in Figure 2, individual descriptors

¹The dimension indices were selected using the method described in Section VI-A using NetVLAD descriptors.

²If Euclidean distance is used for this purpose, Delta Descriptors would need to be ℓ_2 -normalized.

are ℓ_2 -normalized, knowing that the Euclidean distance between pairs of normalized descriptors is proportional to the cosine distance between their un-normalized counterparts.

IV. EXPERIMENTAL SETUP

A. Datasets

We used subsets of two different benchmark datasets to conduct experiments: Oxford Robotcar [17] and Nordland [48]. Repeated route traverses from these datasets exhibit significant variations in scene appearance due to changes in environmental conditions caused by time of day and seasonal cycles.

a) *Oxford Robotcar*: This dataset is comprised of 10 km traverses of urban regions of Oxford city captured under a variety of environmental conditions. We used the forward-facing camera imagery from the first 1 km of three different traverses, referred to as Summer Day, Winter Day and Autumn Night in this paper³. For all three traverses, we used a constant frame separation of 0.5 meters, leading to a database of around 2000 images per traverse.

b) *Nordland*: This dataset comprises 728 km train journey across vegetative open environment from Nordland captured under four seasons. We used the first 1750 images (out of 35768) from the Summer and Winter traverse after skipping the first 250 frames where the train was stationary.

B. Parameter: Sequence Length

The concept of Delta descriptors is based on measuring the changes in visual information in the form of places observed during a traverse which are then expected to be preserved across subsequent traverses. When using a very short sequence length, such changes can be erratic due to high visual overlap between adjacent frames. This is partly due to the unstable response output from the underlying CNN as it is not trained to produce smooth variation in descriptor for small variations in camera motion, as shown in Figure 2. In order to choose the sequence length parameter for our experiments, we used the relative distribution of cosine distance between descriptors, obtained by matching a dataset with itself.

Figure 3 shows these distributions for Winter Day traverse from the Oxford Robotcar dataset and Summer traverse from the Nordland dataset, where cosine distance is plotted against frame separation as the median value computed across the whole traverse. We found that using a fixed cosine distance threshold of 0.7 (black horizontal line), a minimum bound on the sequence length parameter can be directly estimated from these distributions (shown with red circles), making sure that the place observation has changed sufficiently enough to robustly measure the changes in descriptor values. Using this method, the sequential span was found to be 38, 36 and 57 for the Oxford datasets: Winter Day, Summer Day and Autumn Night. For the Nordland Summer and Winter traverses, these values were estimated to be 14 and 8. Hence, as a lower bound, we compute Delta Descriptors using a fixed sequence

³corresponding to 2014-07-14-14-49-50, 2014-11-14-16-34-33, 2015-02-03-08-45-10 respectively.




Fig. 3. Median of cosine distance between descriptors of neighboring image frames is plotted against their frame separation for Oxford Winter Day (left) and Nordland Summer (right) traverses to estimate the sequence length parameter for calculating Delta Descriptors.

length of 64 and 16 frames for all the traverses of the Oxford Robotcar and the Nordland dataset respectively.

C. Evaluation

We used Precision-Recall (PR) curves to measure VPR performance. For a given localization radius, precision is defined as the ratio of correct matches to total matches retrieved and recall is defined as the ratio of correct matches to total possible true matches. For datasets used in this paper, a true match exists for every query image. A match for a query is retrieved only when its cosine distance is less than a threshold which is varied to generate the PR curves. We present PR curves for two different values of localization radii: 10 and 40 meters for the Oxford Robotcar dataset and 2 and 10 frames for the Nordland dataset. For some of the experiments, we also report Precision at 100% Recall which is useful for re-ranking based hierarchical localization pipelines [49], [50].

D. Comparative Study

We used the state-of-the-art single image-based descriptor NetVLAD [11] as a baseline in the results, represented as *Raw Descriptors*. *Delta Descriptors* were calculated using Equation 3 with NetVLAD as the underlying descriptor (see Section VI-B for experiments using different underlying descriptor). As the proposed Delta descriptors use sequential information, we also compare them against a naive sequential representation of NetVLAD, achieved by smoothing the baseline descriptors using Equation 2, represented as *Smoothed Descriptors*. Furthermore, we also consider the orthogonal approach to utilizing sequences for VPR that is based on sequential aggregation of match scores, typically obtained by comparing single image descriptors. For this, we use a simplified version of sequence matching which is similar to [7] but only aggregates match scores along a straight line without any velocity searching [51]. We refer to this as *SeqMatch* in the results and use it on top of the Raw, Smoothed and Delta descriptors.

V. RESULTS

In this section, we first present the benchmark comparisons on three pairs of route traverses using two different datasets. Then, we demonstrate the performance effects of PCA transformation, data shuffling within a traverse, variations in camera motion and sequential-span searching using Multi-Delta Descriptors.

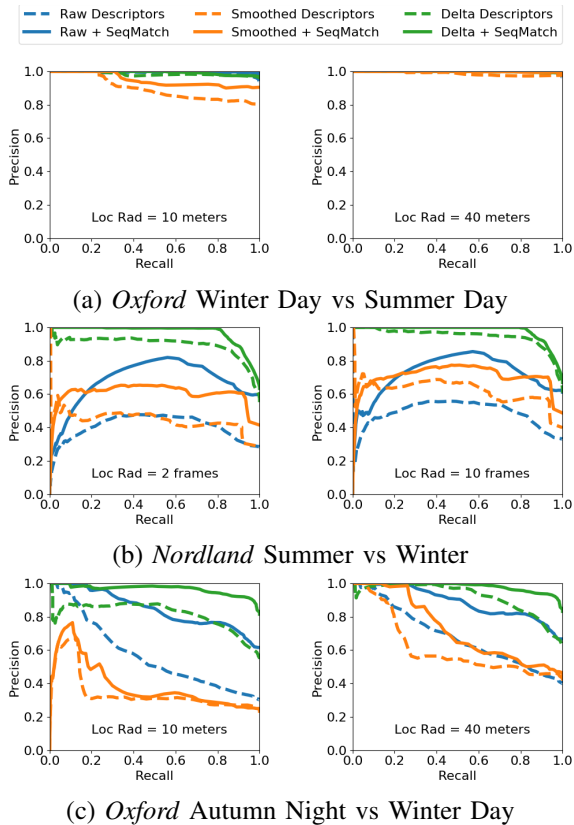


Fig. 4. Precision-Recall performance comparison on three pairs of traverses from the Oxford Robotcar and Nordland datasets.

A. Sequential Representations and/vs Sequential Matching

Figure 4 shows the results for two pairs of traverses from the Oxford Robotcar dataset and one pair from the Nordland dataset. The sequential span L was set to 64 frames (32 meters) for the former and 16 frames for the latter.

As a general trend, it can be observed that Delta Descriptors outperform both Raw and Smoothed descriptors, leading to a much higher recall in the high-precision region. While performance levels are observed to be saturated for the day-day comparison across different seasons in an urban city environment (Figure 4a), the absolute performance of raw descriptors is observed to be quite low when such appearance variations occur in natural open environment (Figure 4b). Such low performance might be due to lack of generalization ability of the NetVLAD descriptors. Thus, in contrast to requiring supervised fine-tuning or re-training, Delta Descriptors tend to mitigate the issue in a completely unsupervised manner.

In Figure 4b, it can also be observed that even using sequence matching on top of raw descriptors (solid blue) cannot achieve performance similar to that attained using Delta Descriptors without sequence matching (dashed green). This particularly highlights the effectiveness of sequence-based place representation as opposed to sequence-based aggregation/filtering of matching scores, given a fixed sequential span.

Figure 4c shows performance trends for a more difficult scenario combining the challenges of both seasonal (au-

turn/winter) and time-of-day (day/night) variations. It can be observed that even without using sequence matching on top, Delta Descriptors perform on par with the Raw+SeqMatch combination, except when considering recall at 100% precision. We observed that the averaging operation in both Smoothed and Delta Descriptors leads to some loss of precision, particularly apparent when the smoothing window size is larger than the considered localization radius. This loss in precision can typically be mitigated by sequential matching where the spurious matching scores are averaged out. In Figure 4, the benefits of sequential matching can be consistently observed in all the results where it not only improves the overall recall performance but simultaneously maintains a high precision level. It is worth noting that a subsequent geometric verification step, commonly employed in hierarchical localization pipelines [39], [49], [50], can further improve the precision performance.

With the use of sequentially-ordered information, Delta descriptors are able to neglect the offsets in the baseline descriptors that occur due to significant variations in scene appearance. This leads to superior performance as compared to the raw descriptors, especially under the challenging scenario of day vs night (see Figure 4c). Furthermore, it can be observed that descriptor smoothing applied naively to the baseline descriptors is of a limited use. This is due to the dilution of discriminative information within the descriptors as they all get closer to the mean of the data. Finally, it can be observed that sequence matching enhances the performance of Delta descriptors more than the raw descriptors, indicating the better representation ability of the former.

B. Dimension Reduction via PCA

Image descriptors obtained through CNNs are typically high-dimensional. For global retrieval tasks, computational complexity is often directly related to the descriptor dimension size. Therefore, dimension reduction techniques like PCA are commonly employed [11], [52], [53], [54]. However, this can lead to significant performance degradation due to extreme variations in scene appearance as the variance distribution in the original descriptor space may not be repeatable. In Figure 5a, we show the effect of PCA-based dimension reduction on the performance of Raw NetVLAD and Delta descriptors using the Oxford Robotcar day-night traverses. It can be observed that the proposed Delta Descriptors are robust to dimension reduction techniques like PCA: even retaining only 50 principal components does not degrade the performance much. On the other hand, the baseline NetVLAD descriptors suffer significant performance drop with PCA even when all 4096 components are retained, highlighting its sensitivity to data centering.

C. Order of Place Observations

The concept of Delta descriptors is based on the sequential order of changes in visual information. In the context of VPR, sequentially observed places typically have some visual overlap which affects the overall dynamics of performance




Fig. 5. (a) Effect of PCA transformation on performance of Raw NetVLAD and Delta Descriptors using Oxford Robotcar Day-Night traverses. * means no PCA transformation. (b) Performance comparisons using Nordland Summer and Winter traverses with random shuffling of data such that the order of places is preserved across the traverses but not within.

when considering either sequential representation or sequential matching. We consider another scenario where both the reference and the query data are shuffled such that the order of images is preserved across traverses but there is no visual overlap between adjacent frames. For this, we used the Nordland dataset and sampled every 100th image out of 35768 images and then performed the shuffling. In Figure 5b, we can observe that even with a sequential span of 2 frames (lacking visual overlap) and localization radius of 1 frame, additional information in the form of sequences can be better utilized with sequence-based descriptors than sequence-based match-score aggregation of single image descriptors, while their combination achieves even higher performance. Furthermore, this experiment also indicates that the concept of Delta Descriptors is applicable in general to describing and matching ordered observations, irrespective of whether or not the adjacent elements are related to each other.

D. Camera Motion Variations & Multi-Delta Descriptors

In our previous experiments, we used a constant frame spacing between the reference and the query traverses (0.5 meters for Oxford Robotcar). In practice, camera velocity may change both within and across the repeated traverses of the environment. In order to observe the effect of such variations on the VPR performance, we conducted another experiment using the first 8600 images⁴ (~ 1 km) from the Winter Day and Autumn Night traverses without any data pre-processing, that is, without motion-based keyframe selection.

For this study, we used a sequential span of 64 frames both for computing Delta Descriptors and sequence matching. Furthermore, in order to deal with variable motion both across and within the traverses, we present a preliminary investigation into sequential-span searching using a Multi-Delta Descriptor approach. To achieve this, for both the reference and the query data, multiple Delta descriptors are computed using a range of sequential spans: $\{30, 40, 50, 60\}$ frames in this case. The match value for any given image pair is calculated as the minimum of the cosine distance over all the possible combinations of sequential spans (16 here) used for computing multiple Delta Descriptors.

⁴only every 5th frame was considered, leading to the dataset size of 1720 images each. Note that this does not affect the camera velocity and was only done to reduce the processing time.

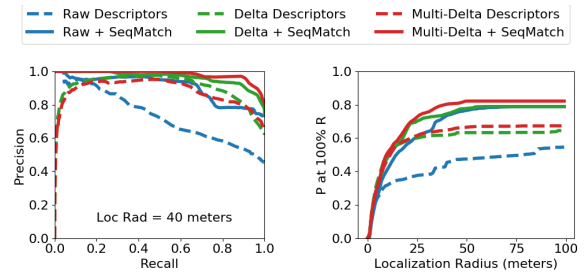


Fig. 6. Performance comparison using Oxford Autumn Night and Winter Day traverses without any pre-processing of camera motion.

In Figure 6, we can observe that even without any data pre-processing (keyframe selection), state-of-the-art performance is achieved by the Delta Descriptors + SeqMatch combination. As we do not use any local velocity-search technique (originally proposed in [7]) in our sequential matching method, performance contribution of SeqMatch is less due to velocity variations as compared against the constant-velocity experiments in Figure 4c. However, the effect of local variations in camera velocity is less detrimental for Delta descriptors, leading to superior absolute performance even without any sequential matching. Finally, it can be observed that the Multi-Delta Descriptor approach further improves the state-of-the-art performance. This also emphasizes the highly discriminative nature of difference-based description that enables accurate match searching within a given range of sequential spans, potentially leading to its applications beyond *exact* repeats of route traversals.

VI. DISCUSSION

A. Visualizing Variations in Activations

In order to visualize how Delta descriptors utilize the difference-based sequential information, we used a Global Max Pooling (GMP) based image descriptor through which image region activations can be directly interpreted which is not trivial for VLAD pooling of NetVLAD. GMP descriptors are extracted from the final *conv* layer of ResNet-50 [55] and Delta Descriptors are calculated using them for Oxford's Winter Day and Summer Day traverses for this experiment.

For visualization purpose, using the ground truth for place matches, dimensions of the GMP descriptor were ranked in order to observe only those which contributed the most to the performance. This was achieved by taking an element-wise product of a known matching pair of descriptors and sorting them in the order of the product value (a higher value indicates that both the descriptors had a similar high activation). This process was repeated for all the pairs of descriptors (~ 2000) and the dimensions that repeatedly ranked higher (higher product value) were selected for visualization.

In Figure 7 (a) and (c), GMP descriptor dimension index 1501 is used. The graph in Figure 7a shows the variation of descriptor values along the route for Raw, Smoothed, and Delta descriptors⁵ (from top to bottom). For image

⁵All the descriptors are ℓ_2 -normalized independently to aid visualization in Euclidean space.

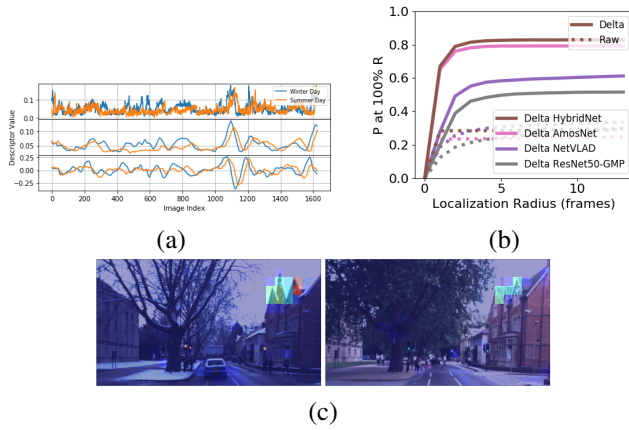


Fig. 7. (a) Time-series of GMP descriptor of images observed along a route for two different traverses. (b) Performance comparison using different underlying descriptors (Raw) to compute Delta Descriptor. (c) 550th image index from the Winter Day and Autumn Day traverse with colored activation masks overlaid (green and red correspond to 20 and 50 respectively).

indices in the range of 400 – 800, both the raw and the smoothed values do not align well across the traverses but are relatively closer in the Delta Descriptor space. Figure 7c displays the 550th image from the Winter (left) and the Summer day traverse (right) where mask color indicates the activation values that increase from blue to green and then to red. It can be observed that the activations for the image from the Summer traverse are lower than that for the Winter traverse due to different lighting conditions around the visual landmark, leading to an increased distance in the original descriptor space. However, as the Delta descriptor only considers changes *within* a traverse, its descriptor values still remain consistent throughout, even though the absolute activation values are lower in one of the traverses.

B. Delta based on Different Underlying Descriptors

For the Nordland (Winter vs Summer) dataset, Figure 7b shows performance comparison between the Raw and Delta descriptors computed using four different underlying descriptors: NetVLAD, GMP (ResNet50), AMOSNet (fc7) [21] and HybridNet (fc7) [21]. It can be observed that irrespective of the underlying descriptor choice, Delta Descriptors lead to consistent performance gains with significant improvements for viewpoint-based descriptors (AMOSNet and HybridNet).

VII. CONCLUSION AND FUTURE WORK

Visual place recognition under large appearance changes is a difficult task. Existing deep-learnt global image description methods do enable effective global image retrieval. However, when operating in new environments where appearance conditions vary drastically due to day-night and seasonal cycles, these methods tend to suffer from an offset in their image description. Our proposed Delta Descriptors are defined in a difference space, which is effective at eliminating description offsets seen in the original space in a completely unsupervised manner. This leads to a significant performance gain, especially for the challenging scenario of day-night VPR. When considering a given sequential span, we have

demonstrated that Delta Descriptors achieve state-of-the-art results when combined with sequential matching. This performance is a strong indicator of the robust representation ability given by Delta Descriptors. Finally, we have presented results for a range of experiments that show the robustness of our method when handling PCA-based dimensional reduction and variations in camera motion both along and across the repeated route traversals.

Our current work can be extended in several ways including estimating the descriptor transformation (offsets) on the fly, learning what visual landmarks are more suited to track changes, and measuring changes independently but simultaneously for different descriptor dimensions. In particular, it would be interesting to see how a framework that learns underlying descriptors would change its behaviour if optimized for place recognition performance using the subsequent Delta Descriptors. The concept of using a difference space [18] is not well-explored in the place recognition literature but is a promising avenue for future research applied to other similar problems where inferring or learning the changes might be more relevant than the representation itself [19], [20], [46]. We believe that our research contributes to the continued understanding of deep-learnt image description techniques and opens up new opportunities for developing and learning robust representations of places that leverage spatio-temporal information.

REFERENCES

- [1] D. G. Lowe, “Distinctive image features from scale-invariant keypoints,” *International journal of computer vision*, vol. 60, no. 2, pp. 91–110, 2004.
- [2] E. Rublee, V. Rabaud, K. Konolige, and G. Bradski, “Orb: An efficient alternative to sift or surf,” in *Computer Vision (ICCV), 2011 IEEE international conference on*. IEEE, 2011, pp. 2564–2571.
- [3] N. Dalal and B. Triggs, “Histograms of oriented gradients for human detection,” in *Computer Vision and Pattern Recognition, 2005. CVPR 2005. IEEE Computer Society Conference on*, vol. 1. IEEE, 2005, pp. 886–893.
- [4] A. Oliva, “Gist of the scene,” *Neurobiology of attention*, vol. 696, no. 64, pp. 251–258, 2005.
- [5] J. Sivic and A. Zisserman, “Video google: A text retrieval approach to object matching in videos,” in *Proceedings of International Conference on Computer Vision (ICCV)*. IEEE, 2003, p. 1470.
- [6] H. Jégou, M. Douze, C. Schmid, and P. Pérez, “Aggregating local descriptors into a compact image representation,” in *Computer Vision and Pattern Recognition (CVPR), 2010 IEEE Conference on*. IEEE, 2010, pp. 3304–3311.
- [7] M. J. Milford and G. F. Wyeth, “Seqslam: Visual route-based navigation for sunny summer days and stormy winter nights,” in *Robotics and Automation (ICRA), 2012 IEEE International Conference on*. IEEE, 2012, pp. 1643–1649.
- [8] T. Naseer, L. Spinello, W. Burgard, and C. Stachniss, “Robust visual robot localization across seasons using network flows,” in *Twenty-Eighth AAAI Conference on Artificial Intelligence*, 2014.
- [9] K. M. Yi, E. Trulls, V. Lepetit, and P. Fua, “Lift: Learned invariant feature transform,” in *European Conference on Computer Vision*. Springer, 2016, pp. 467–483.
- [10] H. Noh, A. Araujo, J. Sim, T. Weyand, and B. Han, “Large-scale image retrieval with attentive deep local features,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2017, pp. 3456–3465.
- [11] R. Arandjelovic, P. Gronat, A. Torii, T. Pajdla, and J. Sivic, “Netvlad: Cnn architecture for weakly supervised place recognition,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2016, pp. 5297–5307.

- [12] S. Garg, N. Suenderhauf, and M. Milford, "Lost? appearance-invariant place recognition for opposite viewpoints using visual semantics," in *Proceedings of Robotics: Science and Systems XIV*, 2018.
- [13] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, "Generative adversarial nets," in *Advances in neural information processing systems*, 2014, pp. 2672–2680.
- [14] A. Anoosheh, T. Sattler, R. Timofte, M. Pollefeys, and L. Van Gool, "Night-to-day image translation for retrieval-based localization," *arXiv preprint arXiv:1809.09767*, 2018.
- [15] 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.
- [16] P.-E. Sarlin, F. Debraine, M. Dymczyk, and R. Siegwart, "Leveraging deep visual descriptors for hierarchical efficient localization," in *Conference on Robot Learning*, 2018, pp. 456–465.
- [17] W. Maddern, G. Pascoe, C. Linegar, and P. Newman, "1 year, 1000 km: The oxford robotcar dataset." *IJ Robotics Res.*, vol. 36, no. 1, pp. 3–15, 2017.
- [18] K. Fukui and A. Maki, "Difference subspace and its generalization for subspace-based methods," *IEEE transactions on pattern analysis and machine intelligence*, vol. 37, no. 11, pp. 2164–2177, 2015.
- [19] K. Fukui, N. Sogi, T. Kobayashi, J.-H. Xue, and A. Maki, "Discriminant analysis based on projection onto generalized difference subspace," *arXiv preprint arXiv:1910.13113*, 2019.
- [20] B. B. Gatto, E. M. d. Santos, A. L. Koerich, K. Fukui, and W. S. Junior, "Tensor analysis with n-mode generalized difference subspace," *arXiv preprint arXiv:1909.01954*, 2019.
- [21] Z. Chen, A. Jacobson, N. Sünderhauf, B. Upcroft, L. Liu, C. Shen, I. Reid, and M. Milford, "Deep learning features at scale for visual place recognition," in *Robotics and Automation (ICRA), 2017 IEEE International Conference on*. IEEE, 2017, pp. 3223–3230.
- [22] R. Girdhar, D. Ramanan, A. Gupta, J. Sivic, and B. Russell, "Actionvlad: Learning spatio-temporal aggregation for action classification," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2017, pp. 971–980.
- [23] A. Jalal, Y.-H. Kim, Y.-J. Kim, S. Kamal, and D. Kim, "Robust human activity recognition from depth video using spatiotemporal multi-fused features," *Pattern recognition*, vol. 61, pp. 295–308, 2017.
- [24] L. Wu, Y. Wang, L. Shao, and M. Wang, "3-d personvlad: Learning deep global representations for video-based person reidentification," *IEEE transactions on neural networks and learning systems*, vol. 30, no. 11, pp. 3347–3359, 2019.
- [25] Y. Li, T. Yao, Y. Pan, H. Chao, and T. Mei, "Jointly localizing and describing events for dense video captioning," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018, pp. 7492–7500.
- [26] S. Song, F. Yu, A. Zeng, A. X. Chang, M. Savva, and T. Funkhouser, "Semantic Scene Completion from a Single Depth Image," pp. 1746–1754, 2016. [Online]. Available: <http://arxiv.org/abs/1611.08974>
- [27] A. Dai, C. R. Qi, and M. Nießner, "Shape completion using 3D-encoder-predictor CNNs and shape synthesis," *Proceedings - 30th IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017*, vol. 2017-Janua, pp. 6545–6554, 2017.
- [28] C. Feichtenhofer, A. Pinz, and R. P. Wildes, "Bags of spacetime energies for dynamic scene recognition," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2014, pp. 2681–2688.
- [29] A. Karpathy, G. Toderici, S. Shetty, T. Leung, R. Sukthankar, and L. Fei-Fei, "Large-scale video classification with convolutional neural networks," in *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, 2014, pp. 1725–1732.
- [30] D. Tran, L. Bourdev, R. Fergus, L. Torresani, and M. Paluri, "Learning spatiotemporal features with 3d convolutional networks," in *Proceedings of the IEEE international conference on computer vision*, 2015, pp. 4489–4497.
- [31] S. Hochreiter and J. Schmidhuber, "Long short-term memory," *Neural computation*, vol. 9, no. 8, pp. 1735–1780, 1997.
- [32] J. Chung, C. Gulcehre, K. Cho, and Y. Bengio, "Empirical evaluation of gated recurrent neural networks on sequence modeling," *arXiv preprint arXiv:1412.3555*, 2014.
- [33] E. Johns and G.-Z. Yang, "Place recognition and online learning in dynamic scenes with spatio-temporal landmarks." in *BMVC*. Citeseer, 2011, pp. 1–12.
- [34] V. A. Nguyen, J. A. Starzyk, and W.-B. Goh, "A spatio-temporal long-term memory approach for visual place recognition in mobile robotic navigation," *Robotics and Autonomous Systems*, vol. 61, no. 12, pp. 1744–1758, 2013.
- [35] J. M. Facil, D. Olid, L. Montesano, and J. Civera, "Condition-invariant multi-view place recognition," *arXiv preprint arXiv:1902.09516*, 2019.
- [36] S. Garg, M. Babu V, T. Dharmasiri, S. Hausler, N. Suenderhauf, S. Kumar, T. Drummond, and M. Milford, "Look no deeper: Recognizing places from opposing viewpoints under varying scene appearance using single-view depth estimation," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2019.
- [37] M. Volkov, G. Rosman, D. Feldman, J. W. Fisher, and D. Rus, "Coresets for visual summarization with applications to loop closure," in *2015 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2015, pp. 3638–3645.
- [38] H. Zhang, F. Han, and H. Wang, "Robust multimodal sequence-based loop closure detection via structured sparsity." in *Robotics: Science and systems*, 2016.
- [39] M. Cummins and P. Newman, "Fab-map: Probabilistic localization and mapping in the space of appearance," *The International Journal of Robotics Research*, vol. 27, no. 6, pp. 647–665, 2008.
- [40] E. Pepperell, P. I. Corke, and M. J. Milford, "All-environment visual place recognition with smart," in *Robotics and Automation (ICRA), 2014 IEEE International Conference on*. IEEE, 2014, pp. 1612–1618.
- [41] O. Vysotska and C. Stachniss, "Lazy data association for image sequences matching under substantial appearance changes," *IEEE Robotics and Automation Letters*, vol. 1, no. 1, pp. 213–220, 2016.
- [42] S. Garg and M. Milford, "Straightening sequence-search for appearance-invariant place recognition using robust motion estimation," *Proceedings of the Australasian conference on robotics and automation (ACRA)*, 2017.
- [43] X. Zhang, L. Wang, Y. Zhao, and Y. Su, "Graph-based place recognition in image sequences with cnn features," *Journal of Intelligent & Robotic Systems*, vol. 95, no. 2, pp. 389–403, 2019.
- [44] O. Vysotska and C. Stachniss, "Effective visual place recognition using multi-sequence maps," *IEEE Robotics and Automation Letters*, 2019.
- [45] E. Parisotto, D. Singh Chaptol, J. Zhang, and R. Salakhutdinov, "Global pose estimation with an attention-based recurrent network," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, 2018, pp. 237–246.
- [46] Q. Zhu, Q. Feng, J. Huang, and D. Zhang, "Sparse representation classification based on difference subspace," in *2016 IEEE Congress on Evolutionary Computation (CEC)*. IEEE, 2016, pp. 4244–4249.
- [47] C.-C. Tseng, J.-C. Chen, C.-H. Fang, and J.-J. J. Lien, "Human action recognition based on graph-embedded spatio-temporal subspace," *Pattern Recognition*, vol. 45, no. 10, pp. 3611–3624, 2012.
- [48] N. Sünderhauf, P. Neubert, and P. Protzel, "Are we there yet? challenging seqslam on a 3000 km journey across all four seasons," in *Proc. of Workshop on Long-Term Autonomy, IEEE International Conference on Robotics and Automation (ICRA)*, 2013, p. 2013.
- [49] S. Garg, N. Sünderhauf, and M. Milford, "Semantic-geometric visual place recognition: A new perspective for reconciling opposing views," *The International Journal of Robotics Research*, 2019.
- [50] R. Mur-Artal and J. D. Tardós, "Orb-slam2: An open-source slam system for monocular, stereo, and rgb-d cameras," *IEEE Transactions on Robotics*, vol. 33, no. 5, pp. 1255–1262, 2017.
- [51] S. Garg and M. Milford, "Fast, compact and highly scalable visual place recognition through sequence-based matching of overloaded representations," in *2020 IEEE International Conference on Robotics and Automation (ICRA)*, 2020.
- [52] S. Garg, N. Suenderhauf, and M. Milford, "Don't look back: Robustifying place categorization for viewpoint- and condition-invariant place recognition," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2018.
- [53] S. Lowry and M. J. Milford, "Supervised and unsupervised linear learning techniques for visual place recognition in changing environments," *IEEE Transactions on Robotics*, vol. 32, no. 3, pp. 600–613, 2016.
- [54] S. Schubert, P. Neubert, and P. Protzel, "Unsupervised learning methods for visual place recognition in discretely and continuously changing environments," *arXiv preprint arXiv:2001.08960*, 2020.
- [55] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 770–778.