

Dynamically Modulating Visual Place Recognition Sequence Length For Minimum Acceptable Performance Scenarios

Connor Malone¹, Ankit Vora², Thierry Peynot¹ and Michael Milford¹

Abstract—Mobile robots and autonomous vehicles are often required to function in environments where critical position estimates from sensors such as GPS become uncertain or unreliable. Single image visual place recognition (VPR) provides an alternative for localization but often requires techniques such as sequence matching to improve robustness, which incurs additional computation and latency costs. Even then, the sequence length required to localize at an acceptable performance level varies widely; and simply setting overly long fixed sequence lengths creates unnecessary latency, computational overhead, and can even degrade performance. In these scenarios it is often more desirable to meet or exceed a set target performance at minimal expense. In this paper we present an approach which uses a calibration set of data to fit a model that modulates sequence length for VPR as needed to exceed a target localization performance. We make use of a coarse position prior, which could be provided by any other localization system, and capture the variation in appearance across this region. We use the correlation between appearance variation and sequence length to curate VPR features and fit a Multi-Layer Perceptron (MLP) for selecting the optimal length. We demonstrate that this method is effective at modulating sequence length to maximize the number of sections in a dataset which meet or exceed a target performance whilst minimizing the median length used. We show applicability across several datasets and reveal key phenomena like generalization capabilities, the benefits of curating features and the utility of non-state-of-the-art feature extractors with nuanced properties.

I. INTRODUCTION

Localization is a critical system capability for most mobile robot and autonomous vehicle applications. Correctly determining the location of a platform within an environment is important for key behaviours and considerations such as navigation, safety, mission completion and platform retrieval. One approach to localization which has been the subject of extensive research in recent years is visual place recognition (VPR) [1]. Visual place recognition provides a localization system which uses purely visual information from a camera and therefore can be implemented relatively cheaply compared to alternatives such as LiDAR-based systems [2]. Localization is achieved in VPR by comparing an image from the platform’s current position to a stored database of images captured when previously traversing the environment. Typically VPR is used either independently for localization,

This research was partially supported by funding from Ford Motor Corporation and NVIDIA, ARC Laureate Fellowship FL210100156 to MM, the Australian Government via grant AUSMURIB000001 associated with ONR MURI grant N00014-19-1-2571, and by the QUT Centre for Robotics.

¹The authors are with the QUT Centre for Robotics, School of Electrical Engineering and Robotics, Queensland University of Technology.

²The authors are with the Ford Motor Company.

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

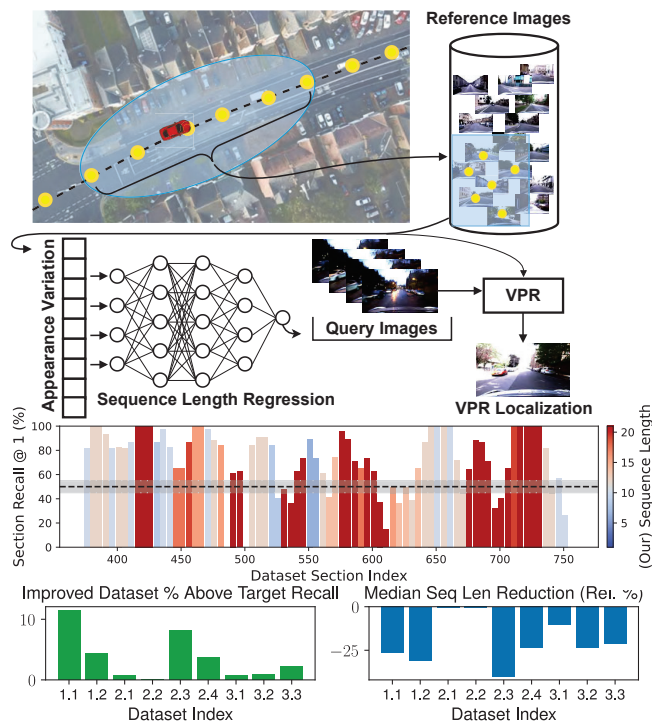


Fig. 1: **Overview of our method.** Our approach modulates VPR sequence length to exceed a target localization performance using coarse position priors and a Multi-Layer Perceptron (MLP). The multi-coloured bar plot shows how the sequence length changes over the dataset (see colour scale) to exceed the target localization performance (black dotted line). The green and blue bar plots show the improvements made over fixed sequence lengths across several datasets by increasing the proportion of the dataset that exceeds the target recall or reducing the median sequence length for the same performance (Explained in Fig. 5 and Sect. V-A).

or as a supporting system in simultaneous localization and mapping (SLAM) applications for loop closure detection [1].

Research in the VPR field often focuses on improving robustness to challenges such as viewpoint and appearance changes due to perspective and the environment [3], [4]. As a result, the metrics that receive the most attention revolve around achieving the highest possible absolute performance *on average* across a single or series of datasets. However, experiences of companies deploying autonomous systems such as self-driving vehicles demonstrate clearly that other aspects of performance are equally important, such as worst case performance [5] or if the system introspectively detects failures [6]. One of the important but

under-investigated functionally-relevant metrics is ensuring that a system will consistently exceed a target localization performance throughout an environment whilst minimizing unnecessary computational overhead or localization latency; this is the subject of the research presented in this paper.

Localization on deployed autonomous platforms is typically aided by additional position sensors such as GPS or IMU, however, these sensors are not always reliable [7]. Even so, they can be effective at supporting VPR localization [8]. Localization with VPR can be further improved by incorporating temporal information through the use of sequence matching [9]. However, the sequence length required to localize at a certain target performance level fluctuates across a dataset (Fig. 1). The obvious fix - setting a fixed sequence length well in excess of what is needed - is not practicable for two reasons. The suitable ‘overkill’ sequence length is not obvious ahead of time, it leads to unnecessarily long and computationally intensive localization processing, and in practice sometimes even degrades performance. Our work addresses this challenge with the following contributions:

- 1) We present an approach which dynamically modulates sequence length for a chosen VPR method to exceed a target localization performance by combining coarse position priors and a Multi-Layer Perceptron (MLP).
- 2) We introduce both a new practical localization metric that measures the consistency of exceeding a target performance and the use of a Leaky ReLU function inside the MSE loss to train networks towards this objective whilst minimizing sequence length.
- 3) We explore the curation of correlated features within VPR feature vectors for this task and demonstrate a clear range of feature types from helpful to counter-productive, in the context of predicting sequence length required to achieve a certain performance level.
- 4) We present experimental results across several datasets, demonstrating the proposed approach effectively modulates sequence length using limited calibration data to outperform fixed sequence lengths.
- 5) We provide analysis on how the approach generalizes across datasets and environmental conditions.

II. RELATED WORK

Here we briefly touch on relevant research relating to visual place recognition and hybrid localization systems.

A. Visual Place Recognition

In recent years there has been substantial research published on VPR for localization tasks using handcrafted, convolutional neural network and now transformer-based feature extractors [10]–[13]. Despite impressive performance, appearance variations due to viewpoint shifts, lighting and environmental conditions and even temporal changes remain as major challenges for long term localization [4].

Outside of improving feature extractors, some works create more robust localization systems by combining or aggregating the outputs of multiple state of the art (SOTA) networks and exploit the strengths of all [14]–[17]. Whilst

proven to be effective, these approaches can have limitations such as requiring multiple VPR networks, being dataset specific or simply requiring additional training using off-the-shelf feature extractors. An alternative to avoid some of these limitations is the addition of sequence-to-sequence VPR.

Sequence matching has been proven in numerous works to improve the localization performance of VPR methods [18]–[20]. This improved performance is typically at the cost of additional latency. Critically, most current VPR research does not consider that the sequence length required to localize a query can vary throughout a dataset [21]. For practical applications this is an important consideration as a sequence length may under perform for some dataset sections whilst introducing unnecessary latency in others.

Some works have shown localization accuracy can be maintained while reducing latency by modulating sequence length based on the strength of localization hypotheses over a sweep of sequence lengths [21]. Approaches in the SLAM and particle filter literature implicitly modulate sequence length through weighting sensory input, motion information and state estimation particles [22], [23]. We differ from these by aiming to address sequence length directly based on visual queues with no inherent assumptions around the importance of recent information.

B. Hybrid Localization Systems

Localization systems for autonomous platforms can also often access position information from many other sensors including GPS, IMUs and LiDAR [24]. However, challenges such as IMU drift [25], textureless surfaces for LiDAR [24], minimal satellite connections or horizontal dilution of position (HDOP) for GPS [7], and multi-path signal reflectance for both LiDAR and GPS can result in uncertain/inaccurate position information from these sensors. To compensate for this degraded position information, many works develop systems which fuse VPR with these sensors to create a more robust localization system [7], [24], [26]–[28]. Most of the works addressing this problem propose some form of behaviour tree, voting system or sensor filter to select or fuse the most appropriate sensors. Our work in this paper fundamentally differs from these approaches as it uses the uncertain position information to alter operational parameters of the VPR system to maintain consistent localization.

III. APPROACH

Our proposed approach addresses the challenge of sequence length selection for a target localization performance using two main assumptions. The first is that the sequence length required for successful localization using a given VPR method is partly related to the variation in appearance across consecutive images [9]. Secondly, that coarse/uncertain position estimates are often available to autonomous platforms.

In this work, we leverage coarse position estimates both as a position prior to reduce the visual place recognition search space and to capture the level of appearance variation in the current region of the environment (Section III-A). We then curate correlated features from the VPR feature

vectors (Section III-B) to train a small MLP regression model (or, fully-connected network) to select a sequence length which enables VPR to exceed a target localization performance (Section III-B.1). We demonstrate this approach is particularly suited for deployment applications through a training pipeline which minimizes sequence length while maximizing the proportion of an environment where the target localization performance is exceeded (Section III-C).

A. Measuring Appearance Variation

We begin by providing some of the typical VPR conventions utilised in this work for capturing the appearance variation across the coarse/uncertain position prior. In single image VPR, the process for performing localization generally follows that a query image, \mathcal{Q} , captured from the current location is compared to a database of reference images, \mathcal{R} , captured when previously traversing the environment. The reference image found to have the highest similarity with the query image is considered the current localization estimate. This process can be augmented by using sequence-to-sequence comparisons, where the similarity becomes the sum of similarities between a sequence of query and reference database images. For the purpose of similarity calculations, VPR methods often represent the query and reference images as n dimensional feature vectors. Therefore creating the query and reference database feature vectors, q_n and R_n .

Two factors which can contribute to the sequence length required to successfully localize a query image are; the difference in appearance between query and reference images \mathcal{Q} and \mathcal{R} ; and the inherent appearance variation between consecutive images in a traverse. In this work we chose to capture and exploit the appearance variation in regions of the reference database. To achieve this we assume that the provided coarse/uncertain position estimate allows the VPR search space within \mathcal{R} to be reduced to a set of m images. We then reduce the feature vectors from the corresponding images, R_{mn} , to a single vector, v_n , which represents the appearance variation across each individual feature within the position prior.

For each feature, j , in the feature vectors of R_{mn} we start by taking the average difference between the feature value for image i and every other image across the search space:

$$\bar{\Delta}_{ij} = \frac{1}{m} \sum_{k=1}^m (R_{ij} - R_{kj}) . \quad (1)$$

This is performed for every image in m and every feature in n , leaving the average difference array $\bar{\Delta}_{mn}$. We then simply take the standard deviation of average distance values for each feature, j , in $\bar{\Delta}_{mn}$:

$$v_j = \sqrt{\frac{1}{m-1} \sum_{i=1}^m (\bar{\Delta}_{ij} - \bar{\bar{\Delta}}_j)^2} . \quad (2)$$

Leaving the vector v_n to represent the appearance variation across the coarse/uncertain position prior. We take this representation a step further by analysing which features in this vector are most correlated with the sequence length needed to exceed a target localization performance.

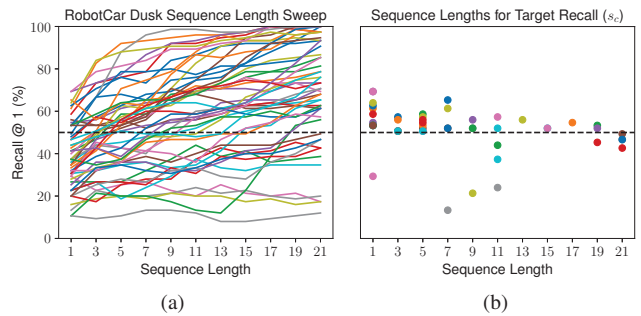


Fig. 2: **a)** Each line represents the change in Recall for increasing sequence lengths in a particular dataset ‘chunk’. The dotted line represents the chosen target recall value for this dataset. **b)** Each dot represents the minimum sequence length on the corresponding colored line in *a)* which exceeds the target recall; or the sequence length closest to the target performance in the case it is not exceeded.

B. Curating Correlated Features

To reduce noise in the MLP input signal, we determine the features from v_n that are most correlated with the sequence length required for a chosen target localization performance.

1) *Choosing Target Performance:* To choose a target localization performance, we separate the dataset into subsets of sequential images with uniform length, referred to as ‘chunks’ from here on, which represent possible coarse/uncertain position priors, m . We then calculate the performance of a sweep of sequence lengths across these ‘chunks’. We evaluate performance using the Recall@1 metric typical to localization tasks. This metric represents the percentage of queries where the top VPR match in the reference database corresponds to the true location [1].

Using a visualization of the sweep results, as in Figure 2a, the target localization performance is selected based on what can be reasonably maintained across the majority of dataset ‘chunks’. After selecting the dataset specific target performance, the required sequence length, s , for each chunk is either set to the smallest sequence length that exceeds the target performance, or the length which comes closest in the case where performance never reaches the target (Fig 2b).

2) *Assessing Feature Correlations:* It is not a new concept that different VPR networks, features and layers have varying robustness and utility in different contexts [29]–[32]. Modern VPR networks integrate this idea into training pipelines and reject features unhelpful to localization tasks [13]. Our approach incorporates it to curate features from pretrained VPR networks for the task of modulating sequence length.

To evaluate the correlation between features and sequence lengths we utilise the adjusted mutual information (AMI) score [33]. This score is easily interpretable, with a score of 1 indicating perfectly correlated data; 0 indicating random/independent data; and negative values indicating data less correlated than would be expected by random chance.

After extracting the appearance variation descriptors, v_n , for all, c , ‘chunks’; we calculate the AMI scores across v_{cn} between each feature, j , and the sequence lengths s_c . This yields a set of n AMI scores which quantifies the correlation

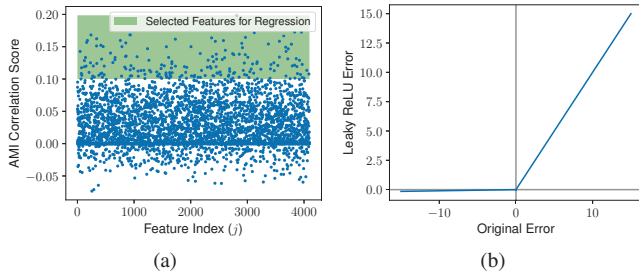


Fig. 3: **a)** The full feature vector v_n is reduced to select features, v_p , using a threshold for the correlation with target sequence length, s . **b)** The Leaky ReLU activation used to control the penalty for over and under predictions separately.

between each feature and the required sequence length for the target performance (Figure 3a):

$$AMI\ Score_j = AMI(s_c, v_{c_j}), \text{ for } j = 1, \dots, n. \quad (3)$$

Importantly, Figure 3a demonstrates that there are both a large number of, features which show some correlation with the required sequence length (> 0), and features which have a correlation equal to or less than random chance (≤ 0). This indicates that there are features which would be harmful to the proposed regression approach. Therefore we set a threshold parameter, $\alpha > 0$, which can be tuned to optimise the appearance variation features to the set p , which are passed on as input, v_{cp} , to train the MLP regression model.

C. Training for Exceeding Target Performance

To perform sequence length selection we establish the problem as a multivariate regression using v_{cp} and s_c as the input and output training signals respectively. Accordingly, we separate the data in these sets into training, validation and test splits and define a small MLP as follows. We set the number of input neurons to p , the number of hidden layers to L , the number of neurons per hidden layer to N and the output of the network to a single neuron. However, unlike other regression problems, in the proposed VPR setting it is more acceptable to over-predict a sequence length and exceed the target performance rather than under-predict and under-perform. This ensures that we maximize the proportion of the environment/dataset where localization performance at least maintains the target value.

1) *Leaky ReLU MSE Loss Function:* To account for this nuance in our training objective, we choose to not use the standard mean squared error (MSE) loss which is typical for regression problems. Instead, our approach applies a leaky ReLU activation function before taking the mean of squared errors. This allows us to separately control the penalty for over-predicting and under-predicting sequence lengths. With a batch size b , sequence length predictions \hat{s}_b , and sequence length targets s_b , this would be achieved using the following:

$$\text{Leaky ReLU MSE} = \frac{1}{b} \sum_{i=1}^b [\text{Leaky ReLU}(s_i - \hat{s}_i)]^2, \quad (4)$$

$$\text{Leaky ReLU} = \begin{cases} \beta x & , x \geq 0 \\ \gamma x & , x < 0 \end{cases}. \quad (5)$$

Where β and γ can be set to control the penalty for under and over predicting respectively. Typically these parameters would be set to values such as $\beta = 1$ and $\gamma = 0.01$, resulting in a shape similar to Figure 3b. Using this loss we are then able to train the regression model for our application focused objective using a standard neural network training pipeline.

IV. EXPERIMENTAL SETUP

In this work, we show that our proposed method for modulating sequence length to exceed a target localization performance is effective across a range of datasets and conditions (Section IV-B). We evaluate using the experiment setup in Section IV-C, metrics in Section IV-D, and compare against other logical alternative approaches (Section IV-E).

A. VPR Method

As mentioned in Section III-B.2, it is known that different VPR networks possess unique properties which are useful in varying contexts. Whilst we have formulated our approach so that it is agnostic to the specific VPR method used for feature extraction, we acknowledge the appearance variation measurement assumes features within VPR feature vectors are spatially related across consecutive images. This is not necessarily a property in VPR methods that use clustering algorithms to aggregate features such as NetVLAD [11] and SALAD [13]. Therefore, in this work we demonstrate our approach using HybridNet [34] and exploit its viewpoint dependency to produce the desired spatially related features.

B. Datasets

To validate our method is robust to different conditions, we perform testing on three place recognition datasets which offer a range of challenges such as illumination, environmental and temporal changes. Examples are in Figure 4.

1) *Oxford RobotCar:* The Oxford RobotCar dataset [35] captures over 100 repetitions of a route through Oxford across a 1 year period. It includes many challenging conditions found in long term localization including illumination, seasonal and temporal changes due to construction. We use a sunny day traverse for the reference database and include common adverse driving conditions using a dusk evening and an overcast day traverse as query datasets. We refer to these as *RobotCar Dusk* and *RobotCar Overcast*.

2) *Nordland:* Nordland [36] is an established benchmark dataset used for visual place recognition and localization. The dataset includes images from a 728km train journey in all four seasons of the year. The changing conditions and perceptual aliasing throughout makes this a tough and unique VPR dataset. We make use of this dataset by using the summer, spring and fall sets in a combination of query and reference pairs. We refer to these as *Nordland Fall-Spring*, *Nordland Fall-Summer* and *Nordland Spring-Summer*.

3) *Ford AV:* The Ford AV dataset [37] presents a challenging driving dataset largely captured from 2017-2018. It includes images captured from vehicles driven throughout Michigan in a range of times, conditions and environments including freeways, city-centers, university campus and suburban areas. In this work we test on query and corresponding



Fig. 4: Example images from the datasets used in our work. **Top Row:** RobotCar Dusk, Overcast and Sunny. **2nd Row:** Nordland Fall, Spring and Summer. **3rd and 4th Rows:** query and reference images respectively from the Ford 1 2017, Ford 3, Ford 4 and Ford 1 2022 datasets.

reference datasets from each of these environment types. We refer to these by their ‘Log’ numbers (and year where applicable), *Ford 1 2017*, *Ford 3*, *Ford 4* and *Ford 1 2022*.

Target Localization Performance Values: In order to train and test our approach on these datasets, we needed to determine what level of VPR performance could reasonably be maintained across each of them. Using HybridNet [34] as our VPR feature extractor, we analysed the recall values from a sweep of sequence lengths, as in Section III-B.1, to set the target localization/recall performance for each dataset. These values can be found in Table I.

C. Experiments

To evaluate the effectiveness of our approach in maximizing the proportion of a dataset/environment that exceeds a target recall, we establish the following experiment protocols. In general, we separate each dataset into overlapping ‘chunks’ and perform VPR within these ‘chunks’. At inference, the appearance variation for a chunk is computed, using the curated feature indices from training, and passed through the MLP to obtain a sequence length to use in that ‘chunk’.

As datasets with uncertain GPS/coarse positions which are also functional as VPR datasets are difficult to find, we manually set the coarse position prior size, m , as in Section III-B.1. We create ‘chunks’ through the dataset with a step size smaller than m to better represent possible position prior locations and generate more training and test data. All datasets except Nordland have been sampled at roughly 1m intervals. We therefore set the chunk size, m , to be 75 images to represent a coarse position estimate with a 75m length for all datasets except in the case for Nordland. We sample these ‘chunks’ from the datasets with a step size of 15 images.

For the purpose of calibrating the proposed MLP regression model, we separate the datasets into training, validation and testing splits of 30%, 20% and 50% respectively. Combined with the selected chunk and step size, the characteristics of the datasets utilised in this work are summarised in Table I. During training we assess regression performance

TABLE I: Summary of Tested Datasets’ Characteristics

Dataset (Target Recall)	Images	Chunks			
		Train	Valid	Test	Total
RobotCar Dusk (50%)	3876	229	153	383	765
RobotCar Overcast (50%)	3876	229	153	383	765
Ford 1 2017 (75%)	5924	351	234	585	1170
Ford 3 (75%)	4602	274	183	458	915
Ford 4 (75%)	3087	184	123	308	615
Ford 1 2022 (40%)	3987	238	159	398	795
Nord Fall-Spring (85%)	11923	711	474	1185	2370
Nord Fall-Summer (85%)	11923	711	474	1185	2370
Nord Spring-Summer (85%)	11923	711	474	1185	2370

purely based on the sequence length training signal, s , determined using Section III-B.1. For localization performance across the test splits we perform VPR for every image within each chunk and calculate the performance for each chunk independently. Testing VPR localization performance in this way could be considered a representation of the chance for localization success at each chunk given that the true position could be any point within the uncertain/coarse position prior.

D. Metrics

As mentioned, throughout this work we use the recall@1 metric widely adopted for assessing the localization performance of place recognition systems. To address the ultimate objective for this work of maintaining consistent VPR performance throughout an environment, we simply use the percentage of dataset ‘chunks’ that meet or exceed the target recall@1. Further to this point, we use the median sequence length across all ‘chunks’ to quantify how efficiently the system is able to achieve this localization consistency.

E. Comparison Methods

For evaluation we compare our approach to four alternatives which demonstrate the possible performance from using a fixed sequence length throughout the entire dataset. To provide fair comparisons we perform VPR within the reduced search space from the coarse/uncertain position prior for *all* methods. The first comparison uses single image VPR (no sequences) and establishes the baseline for performance (Table II, col. 1). Secondly we compare to the performance of using the smallest fixed sequence length from the training set which on average achieves the target localization performance (Table II, col. 3). Finally we compare to fixed sequence length approaches that either match the consistency metric (% of ‘chunks’; Fig. 5, bottom) or median sequence length (Fig. 5, top) of our dynamic sequence length approach.

F. Network Parameters and Optimisation

We optimised hyperparameters to the following values. We set the coarse position prior, $m = 75$. HybridNet feature vectors have dimensions, $n = 4096$. We set hidden layers $L = 3$, hidden layer neurons $N = 128$, AMI threshold $\alpha = 0.125$, loss values $\beta = 1$ and $\gamma = 0.01$, and the number of input p is unique to each dataset depending on the number of features above the AMI threshold in the training sets.

TABLE II: **Sects %**: The percentage of ‘chunks’ which meet or exceed the target performance. **Len**: Median sequence length across ‘chunks’.

Dataset (Target Recall)	No Sequence		Ours		Fixed (Train Len)	
	Sects %	Len	Sects %	Len	Sects %	Len
RC Dusk (50%)	22.5	1	72.3	11	43.6	5
RC Overcast (50%)	55.6	1	77.8	9	55.6	1
Ford 1 2017 (75%)	36.8	1	66.7	7	59.3	5
Ford 3 (75%)	16.2	1	51.3	13	19.0	3
Ford 4 (75%)	28.6	1	44.2	3	43.2	5
Ford 1 2022 (40%)	24.4	1	38.4	13	29.6	3
Nord Fall Spring (85%)	22.3	1	78.8	17	71.1	5
Nord Fall Summer (85%)	52.0	1	82.5	13	76.5	5
Nord Spring Summer (85%)	13.0	1	81.9	15	69.3	5




Fig. 6: ‘Chunks’ from Ford 1 2017 showing longer sequence lengths can sometimes degrade performance.

V. RESULTS

In this section, we present results that demonstrate the advantage of using our sequence length modulation approach for consistently exceeding a target localization performance. Table II and Fig. 5-7 provide quantitative evidence to support the benefit of our approach towards this objective, whilst Tables III, IV and V provide further analysis of its properties.

A. Exceeding Target Localization Performance

We start by analysing the results in Table II and Fig. 5 comparing our approach to the alternative sequence length strategies described in Section IV-E. Importantly, it can be seen in the table that using any sequence-based VPR method typically results in higher percentage of ‘chunks’ exceeding the target localization performance over single image VPR (col. 1). Given all VPR was performed using the coarse position priors, this demonstrates that the majority of improvement from our approach is coming from modulating sequence length rather than the position prior itself. Fig. 1 illustrates how much our approach can modulate sequence length over a dataset, using RobotCar Dusk, while Fig. 6 highlights the scenario where simply setting long fixed sequence lengths can sometimes degrade performance.

First we compare to the smallest fixed sequence length that exceeds the target localization performance across the training split used to fit our MLP regression model (col. 3). Across all datasets, our approach (col. 2) is able to exceed the target performance for a higher percentage of ‘chunks’ than this comparison. This result clearly demonstrates the challenges this work attempts to address; that traditional localization metrics do not guarantee consistent performance throughout a dataset; and that fixed sequence lengths are not guaranteed to be suitable/adequate across entire datasets. Our approach is able to make more effective use of the




Fig. 5: Delta improvements of our approach over fixed sequence lengths when matching its median sequence length (**Top**) or consistency (**Bottom**) metrics. Dataset indices follow order in Table II.

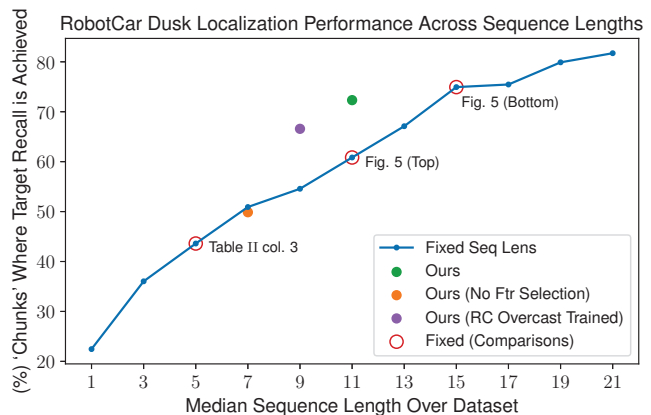


Fig. 7: Localization performance of varying sequence lengths and ablations of our approach for the RobotCar Dusk dataset.

training data by dynamically modulating the sequence length to exceed performance throughout a greater proportion of the dataset. Notably the median sequence length our approach uses is generally higher, likely to accommodate sections of the dataset where a longer sequence length is necessary.

Next we compare to the fixed sequence lengths that either match the consistency or median sequence length metrics of our approach (Fig. 5). These comparisons are necessary for the performance analysis of our approach, however it is important to note that they are not practical methods for choosing fixed sequence lengths as they use ‘oracle’ type information not available at test time. When comparing to using the fixed sequence length which most closely matches the percentage of ‘chunks’ (*Sects %*) from our approach, Fig. 5 shows our approach consistently achieves a lower median sequence length. This demonstrates that our approach is able to achieve this localization consistency with lower latency and therefore more efficiency. Furthermore, using the fixed sequence length matching the median length of our approach (*Len*) clearly shows that our approach exceeds the target performance for a higher percentage of dataset ‘chunks’.

We provide Fig. 7 as a visualisation of how the performance of our approach compares to the entire sweep of fixed sequence lengths for the RobotCar Dusk dataset. This figure can be interpreted as an operating curve for the sweep of

TABLE III: Removing Feature Selection

Dataset	Ours (no feat. sel)		Ours		Fixed (Match Len)	
	Sects %	Len	Sects %	Len	Sects %	Len
RC Dusk	49.9	7	72.3	11	50.9	7
RC Overcast	69.2	5	77.8	9	68.1	5
Ford 1 2017	71.3	13	66.7	7	78.1	13
Ford 3	22.0	5	51.3	13	25.3	5
Ford 4	40.3	3	44.2	3	36.0	3
Ford 1 2022	32.2	7	38.4	13	30.7	7
Nord Fall Spr.	76.2	15	78.8	17	76.5	15
Nord Fall Sum.	82.6	13	82.5	13	81.6	13
Nord Spr. Sum.	79.0	13	81.9	15	79.3	13

TABLE IV: Generalization to Separate Datasets

Test Data →	Ford 1 2017		Ford 3		Ford 4	
Train Data ↓	Sects %	Len	Sects %	Len	Sects %	Len
Ford 1 2017	66.7	7	43.0	9	62.3	15
Ford 3	80	17	51.3	13	68.2	15
Ford 4	57.6	7	50	21	44.2	3
All	80.3	19	53.3	17	65.3	17

sequence lengths. Therefore to prove that our approach is useful towards the objective of this work and more effective than fixed sequence lengths, we show that the operating point of our approach lies above this operating curve.

B. Contribution of Correlated Feature Selection

In our approach we provide a method for curating the most correlated features to use as the MLP regression model input (Section III-B). Fig. 3a shows that there are particular features which are harmful to sequence length regression according to AMI scores. However, we also provide Table III and Fig. 7 to demonstrate this experimentally. Table III clearly shows that removing optimal feature selection from our approach typically lowers the percentage of dataset ‘chunks’ that exceed the target localization performance. Furthermore, using the fixed comparison which matches the median sequence length shows there would be minimal or no improvement over using a fixed sequence length. This importance of feature selection is reiterated in Fig. 7 where the performance of our system changes from above to below the operating curve when removing this process.

C. Generalization to Other Datasets

For most experiments we trained an MLP regression model for each dataset using a small calibration set and tested on a separate unseen section from the same dataset. Table IV provides results which removes this limitation and explores the generalization capabilities of our approach. It shows the performance of each Ford dataset trained model on the other Ford datasets. We include results for the Ford 1 2017, Ford 3 and Ford 4 datasets which share the same target localization performance and train a single model combining all calibration sets.

Interestingly, Table IV shows that when our approach is tested on completely different datasets to its training, it tends to select higher sequence lengths overall. This appears to generally result in a higher percentage of dataset ‘chunks’

TABLE V: Using SOTA VPR Feature Extractors

Dataset (Target)	NetVLAD [11]					
	No Seq		Ours		Fixed (Match Len)	
	Sects %	Len	Sects %	Len	Sects %	Len
Ford 1 2017 (75%)	52.1	1	69.6	7	73.8	7
Ford 3 (75%)	30.6	1	54.1	7	53.3	7
Ford 4 (75%)	35.0	1	39.0	3	43.2	3
Ford 1 2022 (40%)	23.1	1	39.4	13	42.5	13
Dataset (Target)	SALAD [13]					
	No Seq		Ours		Fixed (Match Len)	
	Sects %	Len	Sects %	Len	Sects %	Len
Ford 1 2017 (85%)	65.8	1	87.0	3	77.3	3
Ford 3 (85%)	34.3	1	42.4	5	45.6	5
Ford 4 (85%)	41.9	1	43.8	3	47.7	3
Ford 1 2022 (50%)	36.4	1	48.0	5	50.5	5

that exceed the target localization performance. However, it is likely that the sequence length selections are not as appropriate as on the trained dataset. This is evidenced in column 1 where the model trained on Ford 4 has the same median sequence length but a lower percentage of ‘chunks’ than the model trained specifically for Ford 1 2017.

In addition, we include a data point in Fig. 7 which demonstrates the generalizability of our approach when train and test queries are captured under different conditions. In the figure we show the performance of our approach trained for the RobotCar Overcast dataset when tested on RobotCar Dusk. This model still performs above the operating curve and therefore is beneficial over a fixed sequence length, despite exceeding the target performance in a lower percentage of ‘chunks’ than one specifically trained for RobotCar Dusk.

D. Performance With SOTA Feature Extractors

As described earlier, our proposed approach is particularly suited to VPR feature extractors that maintain some spatial relationship within their features across consecutive images. For completeness and to demonstrate this point we include results in Table V using NetVLAD [11] and SALAD [13] as the VPR feature extractors across the Ford datasets. Table V clearly shows that using our approach with these feature extractors typically results in a lower percentage of ‘chunks’ exceeding the target localization performance compared to the equivalent fixed sequence length. To further support that this is a result of the particular features, it was observed that the typical AMI score for feature correlation was two orders of magnitude lower than the AMI scores observed for the HybridNet features (≈ 0.001 vs ≈ 0.1). This is likely because both NetVLAD and SALAD cluster features and therefore remove spatial relationships. Using the fine-tuned DINOv2 transformer without the clustering from SALAD may present a feature extractor more compatible with our approach.

VI. CONCLUSION

In this work we have addressed the challenge of exceeding a target localization performance by modulating VPR sequence length using an MLP regression model. The results have demonstrated that the approach is able to dynamically

select sequence lengths to exceed a target performance level over a greater proportion of each dataset than fixed sequence lengths. Thereby mitigating some of the latency and computation costs, and potential performance degradation from simply setting overly long fixed sequence lengths. In addition, they demonstrated the effectiveness of curating features from pretrained networks, the utility of non-SOTA feature extractors with nuanced properties, and our approach offers some generalizability to unseen query conditions.

This work aims to emphasise and support the importance of more practical performance metrics in the field of robotics. Showing that, while important, consistent and reliable performance is not always represented by typical metrics. Using visual place recognition, the proposed approach was shown to more effectively maintain localization performance which is critical when other positioning sensors become unreliable.

ACKNOWLEDGEMENT

The authors acknowledge support and funding from the Ford Motor Company through the Ford-QUT Alliance as well as continued support from the Queensland University of Technology (QUT) through the Centre for Robotics.

REFERENCES

- [1] S. Schubert, P. Neubert, S. Garg, M. Milford, and T. Fischer, "Visual place recognition: A tutorial," *IEEE Robotics & Automation Magazine*, pp. 2–16, 2023.
- [2] T. Barros, R. Pereira, L. Garrote, C. Premebida, and U. J. Nunes, "Place recognition survey: An update on deep learning approaches," *arXiv preprint arXiv:2106.10458*, 2021.
- [3] S. Lowry *et al.*, "Visual place recognition: A survey," *IEEE transactions on robotics*, vol. 32, no. 1, pp. 1–19, 2015.
- [4] S. Schubert and P. Neubert, "What makes visual place recognition easy or hard?" *arXiv preprint arXiv:2106.12671*, 2021.
- [5] S. Hausler *et al.*, "Improving worst case visual localization coverage via place-specific sub-selection in multi-camera systems," *IEEE Robotics and Automation Letters*, vol. 7, no. 4, 2022.
- [6] 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.
- [7] T. Saito and Y. Kuroda, "Mobile robot localization using multiple observations based on place recognition and gps," in *IEEE International Conference on Robotics and Automation*, 2013, pp. 1548–1553.
- [8] O. Vysotska, T. Naseer, L. Spinello, W. Burgard, and C. Stachniss, "Efficient and effective matching of image sequences under substantial appearance changes exploiting gps priors," in *IEEE international conference on robotics and automation (ICRA)*, 2015, pp. 2774–2779.
- [9] M.-A. Tomiță, M. Zaffar, B. Ferrarini, M. J. Milford, K. D. McDonald-Maier, and S. Ehsan, "Sequence-based filtering for visual route-based navigation: Analyzing the benefits, trade-offs and design choices," *IEEE Access*, vol. 10, pp. 81 974–81 987, 2022.
- [10] C. Valgren and A. J. Lilienthal, "Sift, surf & seasons: Appearance-based long-term localization in outdoor environments," *Robotics and Autonomous Systems*, vol. 58, no. 2, pp. 149–156, 2010.
- [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] 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.
- [13] S. Izquierdo and J. Civera, "Optimal transport aggregation for visual place recognition," *arXiv preprint arXiv:2311.15937*, 2023.
- [14] N. Keetha *et al.*, "Anyloc: Towards universal visual place recognition," *IEEE Robotics and Automation Letters*, vol. 9, no. 2, 2024.
- [15] A. Ali-bey, B. Chaib-draa, and P. Giguère, "Mixvpr: Feature mixing for visual place recognition," in *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 2023, pp. 2998–3007.
- [16] 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.
- [17] C. Malone, S. Hausler, T. Fischer, and M. Milford, "Boosting performance of a baseline visual place recognition technique by predicting the maximally complementary technique," in *2023 IEEE International Conference on Robotics and Automation*, 2023, pp. 1919–1925.
- [18] E. Pepperell, P. I. Corke, and M. J. Milford, "All-environment visual place recognition with smart," in *2014 IEEE international conference on robotics and automation (ICRA)*. IEEE, 2014, pp. 1612–1618.
- [19] D. Bai, C. Wang, B. Zhang, X. Yi, and X. Yang, "Sequence searching with cnn features for robust and fast visual place recognition," *Computers & Graphics*, vol. 70, pp. 270–280, 2018.
- [20] S. Garg and M. Milford, "Seqnet: Learning descriptors for sequence-based hierarchical place recognition," *IEEE Robotics and Automation Letters*, vol. 6, no. 3, pp. 4305–4312, 2021.
- [21] J. Bruce, A. Jacobson, and M. Milford, "Look no further: Adapting the localization sensory window to the temporal characteristics of the environment," *IEEE Robotics and Automation Letters*, vol. 2, no. 4, pp. 2209–2216, 2017.
- [22] I. A. Kazerouni, L. Fitzgerald, G. Dooly, and D. Toal, "A survey of state-of-the-art on visual slam," *Expert Systems with Applications*, vol. 205, p. 117734, 2022.
- [23] W. Maddern, M. Milford, and G. Wyeth, "Cat-slam: probabilistic localisation and mapping using a continuous appearance-based trajectory," *The International Journal of Robotics Research*, vol. 31, no. 4, 2012.
- [24] S. Yi, S. Worrall, and E. Nebot, "Integrating vision, lidar and gps localization in a behavior tree framework for urban autonomous driving," in *2021 IEEE International Intelligent Transportation Systems Conference (ITSC)*. IEEE, 2021, pp. 3774–3780.
- [25] Z. Zhang, H. Wang, and W. Chen, "A real-time visual-inertial mapping and localization method by fusing unstable gps," in *13th World Congress on Intelligent Control and Automation*, 2018, pp. 1397–1402.
- [26] J. G. Rogers, J. R. Fink, and E. A. Stump, "Mapping with a ground robot in gps denied and degraded environments," in *2014 American Nuclear Conference*. IEEE, 2014, pp. 1880–1885.
- [27] D. Liu, Y. Cui, X. Guo, W. Ding, B. Yang, and Y. Chen, "Visual localization for autonomous driving: Mapping the accurate location in the city maze," in *International Conference on Pattern Recognition (ICPR)*. IEEE, 2021.
- [28] P.-E. Sarlin *et al.*, "Orienternet: Visual localization in 2d public maps with neural matching," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023, pp. 21 632–21 642.
- [29] Z. Chen, O. Lam, A. Jacobson, and M. Milford, "Convolutional neural network-based place recognition," in *Australasian Conference on Robotics and Automation*, vol. 2, 2014, p. 4.
- [30] N. Sünderhauf, S. Shirazi, F. Dayoub, B. Uprocft, and M. Milford, "On the performance of convnet features for place recognition," in *Intelligent Robots and Systems (IROS), 2015 IEEE/RSJ International Conference on*. IEEE, 2015, pp. 4297–4304.
- [31] S. Garg, N. Sünderhauf, 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.
- [32] C. Malone, S. Garg, M. Xu, T. Peynot, and M. Milford, "Improving road segmentation in challenging domains using similar place priors," *IEEE Robotics and Automation Letters*, vol. 7, no. 2, 2022.
- [33] N. X. Vinh, J. Epps, and J. Bailey, "Information theoretic measures for clusterings comparison: is a correction for chance necessary?" in *Proceedings of the 26th annual international conference on machine learning*, 2009, pp. 1073–1080.
- [34] Z. Chen *et al.*, "Deep learning features at scale for visual place recognition," in *Robotics and Automation (ICRA), 2017 IEEE International Conference on*. IEEE, 2017, pp. 3223–3230.
- [35] 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.
- [36] 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.
- [37] S. Agarwal, A. Vora, G. Pandey, W. Williams, H. Kourous, and J. McBride, "Ford multi-av seasonal dataset," *The International Journal of Robotics Research*, vol. 39, no. 12, pp. 1367–1376, 2020.