This ICCV paper is the Open Access version, provided by the Computer Vision Foundation. Except for this watermark, it is identical to the accepted version; the final published version of the proceedings is available on IEEE Xplore.

Scalable Place Recognition Under Appearance Change for Autonomous Driving

Anh-Dzung Doan¹, Yasir Latif¹, Tat-Jun Chin¹, Yu Liu¹, Thanh-Toan Do², and Ian Reid¹

¹School of Computer Science, The University of Adelaide ²Department of Computer Science, University of Liverpool

Abstract

A major challenge in place recognition for autonomous driving is to be robust against appearance changes due to short-term (e.g., weather, lighting) and long-term (seasons, vegetation growth, etc.) environmental variations. A promising solution is to continuously accumulate images to maintain an adequate sample of the conditions and incorporate new changes into the place recognition decision. However, this demands a place recognition technique that is scalable on an ever growing dataset. To this end, we propose a novel place recognition technique that can be efficiently retrained and compressed, such that the recognition of new queries can exploit all available data (including recent changes) without suffering from visible growth in computational cost. Underpinning our method is a novel temporal image matching technique based on Hidden Markov Models. Our experiments show that, compared to state-ofthe-art techniques, our method has much greater potential for large-scale place recognition for autonomous driving.

1. Introduction

Place recognition (PR) is the broad problem of recognizing "places" based on visual inputs [26, 6]. Recently, it has been pursued actively in autonomous driving research, where PR forms a core component in localization (i.e., estimating the vehicle pose) [34, 21, 4, 9, 35, 5, 7] and loop closure detection [10, 13]. Many existing methods for PR require to train on a large dataset of sample images, often with ground truth positioning labels, and state-of-the-art results are reported by methods that employ learning [21, 20, 7, 9].

To perform convincingly, a practical PR algorithm must be robust against appearance changes in the operating environment. These can occur due to higher frequency environmental variability such as weather, time of day, and pedestrian density, as well as longer term changes such as seasons and vegetation growth. A realistic PR system must also contend with "less cyclical" changes, such as construction and roadworks, updating of signage, façades and billboards, as well as abrupt changes to traffic rules that affect traffic flow (this can have a huge impact on PR if the database contains images seen from only one particular flow [10, 13]). Such appearance changes invariably occur in real life.

To meet the challenges posed by appearance variations, one paradigm is to develop PR algorithms that are inherently robust against the changes. Methods under this paradigm attempt to extract the "visual essence" of a place that is independent of appearance changes [1]. However, such methods have mostly been demonstrated on more "natural" variations such as time of day and seasons.

Another paradigm is to equip the PR algorithm with a large image dataset that was acquired under different environmental conditions [8]. To accommodate long-term evolution in appearance, however, it is vital to continuously accumulate data and update the PR algorithm. To achieve continuous data collection cost-effectively over a large region, one could opportunistically acquire data using a fleet of service vehicles (e.g., taxis, delivery vehicles) and amateur mappers. Indeed, there are street imagery datasets that grow continuously through crowdsourced videos [30, 14]. Under this approach, it is reasonable to assume that a decent sampling of the appearance variations, including the recent changes, is captured in the ever growing dataset.

Under continuous dataset growth, the key to consistently accurate PR is to "assimilate" new data quickly. This demands a PR algorithm that is *scalable*. Specifically, the computational cost of testing (i.e., performing PR on a query input) should not increase visibly with the increase in dataset size. Equally crucially, updating or retraining the PR algorithm on new data must also be highly efficient.

Arguably, PR algorithms based on deep learning [7, 9] can accommodate new data by simply appending it to the dataset and fine-tuning the network parameters. However, as we will show later, this fine-tuning process is still too costly to be practical, and the lack of accurate labels in the testing sequence can be a major obstacle.

Contributions We propose a novel framework for PR on large-scale datasets that continuously grow due to the incor-

poration of new sequences in the dataset. To ensure scalability, we develop a novel PR technique based on Hidden Markov Models (HMMs) that is lightweight in both training and testing. Importantly, our method includes a topologically sensitive compression procedure that can update the system efficiently, *without* using GNSS positioning information or computing visual odometry. This leads to PR that can not only improve accuracy by continuous adaption to new data, but also maintain computational efficiency. We demonstrate our technique on datasets harvested from Mapillary [30], and also show that it compares favorably against recent PR algorithms on benchmark datasets.

2. Problem setting

We first describe our adopted setting for PR for autonomous driving. Let $\mathcal{D} = \{\mathcal{V}_1, \ldots, \mathcal{V}_M\}$ be a dataset of M videos, where each video

$$\mathcal{V}_i = \{I_{i,1}, I_{i,2}, \dots, I_{i,N_i}\} = \{I_{i,j}\}_{j=1}^{N_i}$$
(1)

is a time-ordered sequence of N_i images. In the proposed PR system, \mathcal{D} is collected in a distributed manner using a fleet of vehicles instrumented with cameras. Since the vehicles could be from amateur mappers, accurately calibrated/synchronized GNSS positioning may not be available. However, we do assume that the camera on all the vehicles face a similar direction, e.g., front facing. The query video is represented as

$$\mathcal{Q} = \{Q_1, Q_2, \dots, Q_T\}\tag{2}$$

which is a temporally-ordered sequence of T query images. The query video could be a new recording from one of the contributing vehicles (recall that our database D is continuously expanded), or it could be the input from a "user" of the PR system, e.g., an autonomous vehicle.

Overall aims For each $Q_t \in \mathcal{Q}$, the goal of PR is to retrieve an image from \mathcal{D} that was taken from a similar location to Q_t , i.e., the FOV of the retrieved image overlaps to a large degree with Q_t . As mentioned above, what makes this challenging is the possible variations in image appearance.

In the envisioned PR system, when we have finished processing Q, it is appended to the dataset

$$\mathcal{D} = \mathcal{D} \cup \{\mathcal{Q}\},\tag{3}$$

thus the image database could grow unboundedly. This imposes great pressure on the PR algorithm to efficiently "internalise" new data and compress the dataset. As an indication of size, a video can have up to 35,000 images.

2.1. Related works

PR has been addressed extensively in literature [26]. Traditionally, it has been posed as an image retrieval problem using local features aggregated via a BoW representation [10, 13, 11]. Feature-based methods fail to match correctly under appearance change. To address appearance change, SeqSLAM [28] proposed to match statistics of the current image sequence to a sequence of images seen in the past, exploiting the temporal relationship. Recent methods have also looked at appearance transfer [31][23] to explicitly deal with appearance change.

The method closest in spirit to ours is [8], who maintain multiple visual "experiences" of a particular location based on localization failures. In their work, successful localization leads to discarding data, and they depend extensively on visual odometry (VO), which can be a failure point. In contrast to [8], our method does not rely on VO; only image sequences are required. Also, we update appearance in both successful and unsuccessful (new place) localization episodes, thus gaining robustness against appearance variations of the same place. Our method also has a novel mechanism for map compression leading to scalable inference.

A related problem is that of visual localization (VL): inferring the 6 DoF pose of the camera, given an image. Given a model of the environment, PnP [24] based solutions compute the pose using 2D-3D correspondences [34], which becomes difficult both at large scale and under appearance change [39]. Some methods address the issue with creating a model locally using SfM against which query images are localized [35]. Given the ground truth poses and the corresponding images, VL can also be formulated as an image to pose regression problem, solving simultaneously the retrieval and pose estimation. Recently, PoseNet [21] used a Convolution Neural Network (CNN) to learn this mapping, with further improvements using LSTMs to address overfitting [41], uncertainty prediction [19] and inclusion of geometric constraints [20]. MapNet [7] showed that a representation of the map can be learned as a network and then used for VL. A downside of deep learning based methods is their high-computational cost to train/update.

Hidden Markov Models (HMMs) [38, 33] have been used extensively for robot localization in indoor spaces [22, 2, 37]. Hansen et al. [15] use HMM for outdoor scene, but they must maintain a similarity matrix between database and query sequences, which is unscalable when data is accumulated continuously. Therefore, we are one of the first to apply HMMs to large urban-scale PR, which requires significant innovation such as a novel efficient-to-evaluate observation model based on fast image retrieval (Sec. 4.2). In addition, our method explicitly deals with temporal reasoning (Sec. 4.1), which could help to combat the confusion from perceptual aliasing problem [36]. Note also that our main contributions are in Sec. 5, which tackles PR on a continuously growing dataset \mathcal{D} .

3. Map representation

When navigating on a road network, the motion of the vehicle is restricted to the roads, and the heading of the ve-



Figure 1: An overview of our idea using HMM for place recognition. Consider dataset $\mathcal{D} = \{\mathcal{V}_1, \mathcal{V}_2\}$ and query \mathcal{Q} . Figure 1a: Because \mathcal{V}_1 and \mathcal{V}_2 are recorded in different environmental conditions, \mathcal{V}_2 cannot be matched against \mathcal{V}_1 , thus there is no connection between \mathcal{V}_1 and \mathcal{V}_2 . Query \mathcal{Q} visits the place covered by \mathcal{V}_1 and \mathcal{V}_2 , and then an unknown place. Figure 1b: Query \mathcal{Q} is firstly localized against only \mathcal{V}_1 . When it comes to the "Overlap region" at time t + 1, it localizes against both \mathcal{V}_1 and \mathcal{V}_2 . The image corresponding to MaxAP at every time step t is returned as the matching result. Figure 1c: A threshold decides if the matching result should be accepted, thus when Q visits an unseen place, the MaxAPs of \mathcal{V}_1 and \mathcal{V}_2 are small, we are uncertain about the matching result. Once \mathcal{Q} is finished, the new place discovered by \mathcal{Q} is added to the map to expand the coverage area. In addition, since Q is matched against both \mathcal{V}_1 and \mathcal{V}_2 , we can connect \mathcal{V}_1 and \mathcal{V}_2 .

hicle is also constrained by the traffic direction. Hence, the variation in pose of the camera is relatively low [35, 32].

The above motivates us to represent a road network as a graph $\mathcal{G} = (\mathcal{N}, \mathcal{E})$, which we also call the "map". The set of nodes \mathcal{N} is simply the set of all images in \mathcal{D} . To reduce clutter, we "unroll" the image indices in \mathcal{D} by converting an (i, j) index to a single number $k = N_1 + N_2 + \cdots + N_{i-1} + j$, hence the set of nodes are

$$\mathcal{N} = \{1, \dots, K\},\tag{4}$$

where $K = \sum_{i=1}^{M} N_i$ is the total number of images. We call an index $k \in \mathcal{N}$ a "place" on the map.

We also maintain a corpus C that stores the images observed at each place. For now, the corpus simply contains

$$\mathcal{C}(k) = \{I_k\}, \quad k = 1, \dots, K, \tag{5}$$

at each cell C(k). Later in Sec. 5, we will incrementally append images to C as the video dataset D grows.

In \mathcal{G} , the set of edges \mathcal{E} connect images that overlap in their FOVs, i.e., $\langle k_1, k_2 \rangle$ is an edge in \mathcal{E} if

$$\exists I \in \mathcal{C}(k_1) \text{ and } \exists I' \in \mathcal{C}(k_2) \text{ such that } I, I' \text{ overlap.}$$
 (6)

Note that two images can overlap even if they derive from different videos and/or conditions. The edges are weighted

by probabilities of transitioning between places, i.e.,

$$w(\langle k_1, k_2 \rangle) = P(k_2 \mid k_1) = P(k_1 \mid k_2), \tag{7}$$

for a vehicle that traverses the road network. Trivially,

$$\langle k_1, k_2 \rangle \notin \mathcal{E} \text{ iff } P(k_2 \mid k_1) = P(k_1 \mid k_2) = 0.$$
 (8)

It is also clear from (7) that \mathcal{G} is undirected. Concrete definition of the transition probability will be given in Sec. 5. First, Sec. 4 discusses PR of \mathcal{Q} given a fixed \mathcal{D} and map.

4. Place recognition using HMM

To perform PR on $Q = \{Q_1, \ldots, Q_T\}$ against a fixed map $\mathcal{G} = (\mathcal{N}, \mathcal{E})$ and corpus \mathcal{C} , we model Q using a HMM [33]. We regard each image Q_t to be a noisy observation (image) of an latent *place state* s_t , where $s_t \in \mathcal{N}$. The main reason for using HMM for PR is to exploit the temporal order of the images in Q, and the high correlation between time and place due to the restricted motion (Sec. 3).

To assign a value to s_t , we estimate the belief

$$P(s_t \mid Q_{1:t}), \quad s_t \in \mathcal{N}, \tag{9}$$

where $Q_{1:t}$ is a shorthand for $\{Q_1, \ldots, Q_t\}$. Note that the

belief is a probability mass function, hence

$$\sum_{s_t \in \mathcal{N}} P(s_t \mid Q_{1:t}) = 1.$$
(10)

Based on the structure of the HMM, the belief (9) can be recursively defined using Bayes' rule as

$$P(s_t|Q_{1:t}) = \eta P(Q_t|s_t) * \sum_{s_{t-1} \in \mathcal{N}} P(s_t|s_{t-1}) P(s_{t-1}|Q_{1:t-1}), \quad (11)$$

where $P(Q_t|s_t)$ is the observation model, $P(s_t|s_{t-1})$ is the state transition model, and $P(s_{t-1}|Q_{1:t-1})$ is the prior (the belief at the previous time step) [33]. The scalar η is a normalizing constant to ensure that the belief sums to 1.

If we have the belief $P(s_t \mid Q_{1:t})$ at time step t, we can perform PR on Q_t by assigning

$$s_t^* = \underset{s_t \in \mathcal{N}}{\operatorname{arg\,max}} P(s_t \mid Q_{1:t})$$
(12)

as the place estimate of Q_t . Deciding the target state in this manner is called maximum a posteriori (MaxAP) estimation. See Fig. 1 for an illustration of PR using HMM.

4.1. State transition model

The state transition model $P(s_t|s_{t-1})$ gives the probability of moving to place s_t , given that the vehicle was at place s_{t-1} in the previous time step. The transition probability is simply given by the edge weights in \mathcal{G} , i.e.,

$$P(s_t = k_2 | s_{t-1} = k_1) = w(\langle k_1, k_2 \rangle).$$
(13)

Again, we defer the concrete definition of the transition probability to Sec. 5. For now, the above is sufficient to continue our description of our HMM method.

4.2. Observation model

Our observation model is based on image retrieval. Specifically, we use SIFT features [25] and VLAD [17] to represent every image. Priority search k-means tree [29] is used to index the database, but it is possible to use other indexing methods [16, 12, 3].

Image representation For every image $I_k \in C$, we seek a nonlinear function $\psi(I_k)$ that maps the image to a single high-dimensional vector. To do that, given a set of SIFT features densely extracted from image I_k : \mathbf{X}_k = $\{x_k^h\} \in \mathbb{R}^{d \times H_k}$, where H_k is the number of SIFT features of image I_k . K-means is used to build a codebook $B = \{b_m \in \mathbb{R}^d \mid m = 1, ..., \mathbb{M}\},$ where M is the size of codebook. The VLAD embedding function is defined as:

$$\phi(x_k) = [..., 0, x_k^h - b_m, 0, ...] \in \mathbb{R}^D$$
(14)

where, b_m is the nearest visual word of feature vector x_k^h . To obtain a single vector, we employ sum aggregation:

$$\psi(I_k) = \sum_{i=1}^{H_k} \phi(x_k) \tag{15}$$

To reduce the impact of background features (e.g., trees, roads, sky) within the vector $\psi(I_k)$, we adopt rotation and normalization (RN) [18], followed by L-2 normalization. In particular, we use PCA to project $\psi(I_k)$ from D to D', where D' < D. In our experiment, we set D' = 4,096. Power-law normalization is then applied on rotated data:

$$\psi(I_k) := |\psi(I_k)|^{\alpha} sign(\psi(I_k))$$
(16)

where, we set $\alpha = 0.5$.

Note that different from DenseVLAD [40] which uses whitening for post-processing, performing power-law normalization on rotated data is more stable.

Computing likelihood We adopt priority search k-means tree [29] to index every image $I_k \in C$. The idea is to partition all data points $\psi(I_k)$ into K clusters by using K-means, then recursively partitioning the points in each cluster. For each query Q_t , we find a set of L-nearest neighbor $L(Q_t)$. Specifically, Q_t is mapped to vector $\psi(Q_t)$. To search, we propagate down the tree at each cluster by comparing $\psi(Q_t)$ to K cluster centers and selecting the nearest one.

The likelihood $P(Q_t|s_t)$ is calculated as follows:

- Initialize $P(Q_t|s_t = k) = e^{\frac{-\beta}{\sigma}}, \forall k \in \mathcal{N}$, where, we set $\beta = 2.5$ and $\sigma = 0.3$ in our experiment.
- For each $I_k \in \mathbf{L}(Q_t)$
 - Find node $\hat{k} = C^{-1}(I_k)$, where C^{-1} is the inverse of corpus C, which finds node \hat{k} storing I_k . - Calculate the probability: $\hat{p} = e^{\frac{-dist(Q_t, I_k)}{\sigma}}$,
 - where dist is the distance between Q_t and I_k .

- If
$$\hat{p} > P(Q_t | s_t = \hat{k})$$
, then $P(Q_t | s_t = \hat{k}) = \hat{p}$.

4.3. Inference using matrix computations

The state transition model can be stored in a $K \times K$ matrix E called the transition matrix, where the element at the k_1 -th row and k_2 -th column of **E** is

$$\mathbf{E}(k_1, k_2) = P(s_t = k_2 \mid s_{t-1} = k_1).$$
(17)

Hence, **E** is also the weighted adjacency matrix of graph \mathcal{G} . Also, each row of E sums to one. The observation model can be encoded in a $K \times K$ diagonal matrix O_t , where

$$\mathbf{O}_t(k,k) = P(Q_t \mid s_t = k). \tag{18}$$

If the belief and prior are represented as vectors $\mathbf{p}_t, \mathbf{p}_{t-1} \in$ \mathbb{R}^{K} respectively, operation (11) can be summarized as

$$\mathbf{p}_t = \eta \mathbf{O}_t \mathbf{E}^T \mathbf{p}_{t-1},\tag{19}$$

where p_0 corresponds to uniform distribution. From this, it can be seen that the cost of PR is $\mathcal{O}(K^2)$.

Computational cost Note that **E** is a very sparse matrix, due to the topology of the graph \mathcal{G} which mirrors the road network; see Fig. 3 for an example **E**. Thus, if we assume that the max number of non-zero values per row in **E** is *r*, the complexity for computing \mathbf{p}_t is O(rK).

Nonetheless, in the targeted scenario (Sec. 2), \mathcal{D} can grow unboundedly. Thus it is vital to avoid a proportional increase in E so that the cost of PR can be maintained.

5. Scalable place recognition based on HMM

In this section, we describe a novel method that incrementally builds and compresses \mathcal{G} for a video dataset \mathcal{D} that grows continuously due to the addition of new query videos.

We emphasize again that the proposed technique functions without using GNSS positioning or visual odometry.

5.1. Map intialization

Given a dataset \mathcal{D} with one video $\mathcal{V}_1 = \{I_{1,j}\}_{j=1}^{N_1} \equiv \{I_k\}_{k=1}^K$, we initialize \mathcal{N} and \mathcal{C} as per (4) and (5). The edges \mathcal{E} (specifically, the edge weights) are initialized as

$$w(\langle k_1, k_2 \rangle) = \begin{cases} 0 & \text{if } |k_1 - k_2| > W, \\ \alpha \exp\left(-\frac{|k_1 - k_2|^2}{\delta^2}\right) & \text{otherwise}, \end{cases}$$

where α is a normalization constant. The edges connect frames that are $\leq W$ time steps apart with weights based on a Gaussian on the step distances. The choice of W can be based on the maximum velocity of a vehicle.

Note that this simple way of creating edges will ignore complex trajectories (e.g., loops). However, the subsequent steps will rectify this issue by connecting similar places.

5.2. Map update and compression

Let $\mathcal{D} = \{\mathcal{V}_i\}_{i=1}^M$ be the current dataset with map $\mathcal{G} = (\mathcal{N}, \mathcal{E})$ and corpus \mathcal{C} . Given a query video $\mathcal{Q} = \{Q_t\}_{t=1}^T$, using our method in Sec. 4 we perform PR on \mathcal{Q} based on \mathcal{G} . This produces a belief vector \mathbf{p}_t (19) for all t.

We now wish to append Q to D, and update G to maintain computational scalability of future PR queries. First, create a subgraph $G' = (N', \mathcal{E}')$ for Q, where

$$\mathcal{N}' = \{K + 1, K + 2, \dots, K + T\},$$
(20)

(recall that there are a total of K places in \mathcal{G}), and \mathcal{E}' simply follows Sec. 5.1 for \mathcal{Q} .

In preparation for map compression, we first concatenate the graphs and extend the corpus

$$\mathcal{N} = \mathcal{N} \cup \mathcal{N}', \ \mathcal{E} = \mathcal{E} \cup \mathcal{E}', \ \text{and} \ \mathcal{C}(K+t) = \{Q_t\}$$
 (21)

for t = 1, ..., T. There are two main subsequent steps: culling new places, and combining old places.

Culling new places For each *t*, construct

$$\mathcal{M}(t) = \{k \in \{1, \dots, K\} \mid \mathbf{p}_t(k) \ge \gamma\}, \qquad (22)$$

where γ with $0 \leq \gamma \leq 1$ is a threshold on the belief. There are two possibilities:

- If M(t) = Ø, then Qt is the image of a new (unseen before) place since the PR did not match a dataset image to Qt with sufficient confidence. No culling is done.
- If M(t) ≠ Ø, then for each k₁ ∈ M(t),
 For each k₂ such that ⟨K + t, k₂⟩ ∈ E:
 * Create new edge ⟨k₁, k₂⟩ with weight w(⟨k₁, k₂⟩) = w(⟨K + t, k₂⟩).
 * Delete edge ⟨K + t, k₂⟩ from E

* Delete edge
$$\langle K + l, k_2 \rangle$$
 fior

- $C(k_1) = C(k_1) \cup C(K + t)$. Once the above is done for all t, for those t where $\mathcal{M}(t) \neq \emptyset$, we delete the node K + t in \mathcal{N} and cell C(K + t) in C, both with the requisite adjustment in the remaining indices. See Figs. 2a and 2b for an illustration of culling.

Combining old places Performing PR on Q also provides a chance to connect places in G that were not previously connected. For example, two dataset videos V_1 and V_2 could have traversed a common subpath under very different conditions. If Q travels through the subpath under a condition that is simultaneously close to the conditions of V_1 and V_2 , this can be exploited for compression.

To this end, for each t where $\mathcal{M}(t)$ is non-empty,

- $k_1 = \min \mathcal{M}(t)$.
- For each k₂ ∈ M(t) where k₂ ≠ k₁ and ⟨k₁, k₂⟩ ∉ E:
 For each k₃ such that ⟨k₂, k₃⟩ ∈ E, ⟨k₁, k₃⟩ ∉ E:
 * Create edge ⟨k₁, k₃⟩ with weight w(⟨k₁, k₃⟩) = w(⟨k₂, k₃⟩).
 * Delete edge ⟨k₂, k₃⟩ from E.

$$- \mathcal{C}(k_1) = \mathcal{C}(k_1) \cup \mathcal{C}(k_2).$$

Again, once the above is done for all t for which $\mathcal{M}(t) \neq \emptyset$, we remove all unconnected nodes from \mathcal{G} and delete the relevant cells in \mathcal{C} , with the corresponding index adjustments. Figs. 2c, 1a and 1c illustrate this combination step.

5.3. Updating the observation model

When Q is appended to the dataset, i.e., $\mathcal{D} = \mathcal{D} \cup Q$, all vector $\psi(Q_t)$ need to be indexed to the k-means tree. In particular, we find the nearest leaf node that $\psi(Q_t)$ belongs to. Assume the tree is balanced, the height of tree is $(\log N/\log \mathbf{K})$, where $N = \sum N_i$, thus each $\psi(Q_t)$ needs to check $(\log N/\log \mathbf{K})$ internal nodes and one leaf node. In each node, it needs to find the closest cluster center by computing distances to all centers, the complexity of which is $O(\mathbf{K}D')$. Therefore, the cost for adding the query video Q is $O(T\mathbf{K}D'(\log N/\log \mathbf{K}))$, where T = |Q|. Assume it is a complete tree, every leaf node contains \mathbf{K} points, thus it has N/\mathbf{K} leaf nodes. For each point $\psi(Q_t)$, instead of exhaustedly scanning N/\mathbf{K} leaf nodes, it only needs to check $\log N/\log \mathbf{K}$ nodes. Hence, it is a scalable operation.



Figure 2: An overview of our idea for scalable place recognition. Graph $\mathcal{G} = \mathcal{G}_1 \cup \mathcal{G}_2$, where $\mathcal{G}_1 = \{1, 2, 3, 4, 5\}$ and $\mathcal{G}_2 = \{6, 7, 8\}$ are disjoint sub-graphs. Query video $\mathcal{Q} = \{Q_1, ..., Q_T\}$ is matched against \mathcal{G} . Figure 2a: Q_t is matched with node k = 3 and 7 (dashed green lines), due to $\mathbf{p}_t(3)$, $\mathbf{p}_t(7) > \gamma$. Figure 2b: Q_t is added to node 3 and 7, new edges are created (blue lines) to maintain the connections between Q_{t-1} , Q_{t+1} and Q_t . Figure 2c: Node 3 and 7 are combined. New edges are generated (blue lines) to maintain the connections within the graph. Note that after matching query \mathcal{Q} against \mathcal{G} , our proposed culling and combining methods connect two disjoint sub-graphs \mathcal{G}_1 and \mathcal{G}_2 together.

5.4. Overall algorithm

Algorithm 1 summarizes the proposed scalable method for PR. A crucial benefit of performing PR with our method is that map \mathcal{G} does not grow unboundedly with the inclusion of new videos. Moreover, the map update technique is simple and efficient, which permits it to be conducted for every new video addition. This enables scalable PR on an ever growing video dataset. In Sec. 6, we will compare our technique with state-of-the-art PR methods.

6. Experiments

We use a dataset sourced from Mapillary [30] which consists of street-level geo-tagged imagery; see supplementary material for examples. Benchmarking was carried out on the Oxford RobotCar [27], from which we use 8 different sequences along the same route; details are provided in supplementary material, and the sequences are abbreviated as Seq-1 to Seq-8. The initial database \mathcal{D} is populated with Seq-1 and Seq-2 from the Oxford RobotCar dataset. Seq-3 to Seq-8 are then sequentially used as the query videos. To report the 6-DoF pose for a query image, we inherit the pose of the image matched using the MaxAP estimation. Following [35], the translation error is computed as the Euclidean distance $||c_{est} - c_{qt}||^2$. Orientation errors $|\theta|$, measured in degree, is the angular difference $2\cos(|\theta|) = trace(R_{at}^{-1}R_{est}) - 1$ between estimated and ground truth camera rotation matrices R_{est} and R_{qt} . Following [21, 20, 7, 42], we compare mean and median errors.

Performance with and without updating the database We investigate the effects of updating database on localization accuracy and inference time. After each query sequence finishes, we consider three strategies: i) No update: \mathcal{D} always contains just the initial 2 sequences, ii) Cull: Update \mathcal{D} with the query and perform culling, Algorithm 1 Scalable algorithm for large-scale PR.

Require: Threshold W for transition probability, threshold γ for PR, initial dataset $\mathcal{D} = \{\mathcal{V}_1\}$ with one video.

- 1: Initialize map $\mathcal{G} = (\mathcal{N}, \mathcal{E})$ and corpus \mathcal{C} (Sec. 5.1).
- 2: Create observation model (Sec. 4.2)
- 3: while there is a new query video \hat{Q} do
- 4: Perform PR on Q using map G, then append Q to D.
- 5: Create subgraph \mathcal{G}' for \mathcal{Q} (Sec. 5.2).
- 6: Concatenate \mathcal{G}' to \mathcal{G} , extend \mathcal{C} with \mathcal{Q} (Sec. 5.2).
- 7: Reduce \mathcal{G} by culling new places (Sec. 5.2).
- 8: Reduce \mathcal{G} by combining old places (Sec. 5.2).
- 9: Update observation model (Sec. 5.3).
- 10: end while
- 11: **return** Dataset \mathcal{D} with map \mathcal{G} and corpus \mathcal{C} .

_	No update		Cull	Cull+combine	
_	Seq-3		6.59m, 3.28°		
	Seq-4	7.42m, 4.64°	2m, 4.64° 5.80m , 3.24°		
	Seq-5	16.21m, 5.97°	15.07m , 5.89 °	15.88m, 5.91°	
	Seq-6	26.02m, 9.02°	18.88m , 6.24 °	19.28m, 6.28°	
	Seq-7	31.83m, 17.99°	30.06m, 17.12°	30.03 m, 17.05 °	
	Seq-8	25.62m, 22.38°	24.28m, 21.99°	24.26 m, 21.54 °	
		No update	Cull	Cull+combine	
	Seq-3	6.06m, 1.65°			
	Seq-4	5.80m, 1.40°	5.54m , 1.39°	5.65m, 1.33 °	
	Seq-5	13.70m, 1.56°	13.12m, 1.52 °	13.05m , 1.55°	
	Seq-6	6.65m, 1.87°	5.76m , 1.75 °	6.60m, 1.85°	
	Seq-7	13.58m, 3.52°	11.80m, 2.81°	10.87m , 2.60°	
Seq-8		13.28m, 4.93°	7.13m , 2.31 °	7.15m, 2.47°	

Table 1: Comparison between 3 different settings of our technique. Mean (top) and median (bottom) errors of 6-DoF pose on Oxford RobotCar are reported.



Figure 3: Illustrating map maintenance w and w/o compression. After each query video Q finishes, we compress the map by culling known places in Q and combining old places on the map which represent the same place. Thus, the size of transition matrix is shrunk gradually. In contrast, if compression is not conducted, the size of transition matrix will continue increasing.

Sequences	No update	Cull	Cull+Combine	
Seq-3	4.03			
Seq-4	4.56	5.05	4.82	
Seq-5	4.24	5.06	4.87	
Seq-6	3.81	4.03	3.72	
Seq-7	3.82	4.18	3.78	
Seq-8	3.77	3.91	3.68	

Table 2: Inference time (ms) on Oxford RobotCar. Cull+Combine has comparable inference time while giving better accuracy (see Table 1) over No update.

and iii) Cull+Combine: Full update with both culling and combining nodes. Mean and median 6-DoF pose errors are reported in Table 1. In general, Cull improves the localization accuracy over No update, since culling adds appearance variation to the map. In fact, there are several cases, in which Cull+Combine produces better results over Cull. This is because we consolidate useful information in the map (combining nodes which represent the same place), and also enrich the map topology (connecting nodes close to each other through culling). Inference times per query with different update strategies are given in Table 2. Without updating, the inference time is stable at (\sim 4ms/query) between sequences, since the size of graph and the database do not change. In contrast, culling operation increases the inference time by about 1ms/query, and Cull+Combine makes it comparable to the No update case. This shows that the proposed method is able to compress the database to an extent that the query time after assimilation of new information remains comparable to the case of not updating the database at all.

Map maintenance and visiting unknown regions Figure 3 shows the results on map maintenance with and with-

Training sequences	VidLoc	MapNet	Our method
Seq-1,2	14.1h	11.6h	98.9s
Seq-3	-	6.2h	256.3s
Seq-4	-	6.3h	232.3s
Seq-5	-	6.8h	155.1s
Seq-6	-	5.7h	176.5s
Seq-7	-	6.0h	195.4s

Table 3: Training/updating time on the Oxford RobotCar.

out compression. Without compression, size of map \mathcal{G} (specifically, adjacency matrix **E**) grows continuously when appending a new query video \mathcal{Q} . In contrast, using our compression scheme, known places in \mathcal{Q} are culled, and redundant nodes in \mathcal{G} (i.e., nodes representing a same place) are combined. As a result, the graph is compressed.

Visiting unexplored area allows us to expand the coverage of our map, as we demonstrate using Mapillary data. We set $\gamma = 0.3$, i.e., we only accept the query frame which has the MaxAP belief ≥ 0.3 . When the vehicle explores unknown roads, the probability of MaxAP is small and no localization results are accepted. Once the query sequence ends, the map coverage is also extended; see Fig. 4.

Comparison against state of the art Our method is compared against state-of-the-art localization methods: MapNet [7] and VidLoc [9]. We use the original authors' implementation of MapNet. VidLoc implementation from MapNet is used by the recommendation of VidLoc authors. All parameters are set according to suggestion of authors.¹

For map updating in our method, Cull+Combine steps

¹Comparisons against [8] are not presented due to the lack of publicly available implementation.



Figure 4: Expanding coverage by updating the map. Locations are plotted using ground-truth GPS for visualization only.

Methods	Seq-3	Seq-4	Seq-5	Seq-6	Seq-7	Seq-8
VidLoc	38.86m, 9.34°	38.29m, 8.47°	36.05m, 6.81°	51.09m, 10.75°	54.70m, 18.74°	47.64m, 23.21°
MapNet	9.31m, 4.37°	8.92m, 4.09°	17.19m, 5.72 °	26.31m, 9.78°	33.68m, 18.04°	26.55m, 21.97°
MapNet (update+ retrain)		8.71m, 3.31°	18.44m, 6.94°	28.69m, 10.02°	36.68m, 19.34°	29.64m, 22.86°
Our method	6.59m, 3.28°	6.01m , 3.11°	15.88m , 5.91°	19.28m , 6.28 °	30.03 m, 17.05 °	24.26m , 21.54 °
Methods	Seq-3	Seq-4	Seq-5	Seq-6	Seq-7	Seq-8
VidLoc	29.63m, 1.59 °	29.86m, 1.57°	31.33m, 1.39°	47.75m, 1.70 °	48.53m, 2.40°	42.26m, 1.94 °
MapNet		4.53m , 1.54°	13.89m, 1.17 °	8.69m, 2.42°	12.49m, 1.71 °	8.08m, 2.02°
MapNet (update+ retrain)	- 4.69m , 1.67°	5.15m, 1.44°	17.39m, 1.87°	11.45m, 3.42°	20.88m, 4.02°	11.01m, 5.21°
Our method	6.06m, 1.65°	5.65m, 1.33 °	13.05m , 1.55°	6.60m , 1.85°	10.87m , 2.60°	7.15m , 2.47°

Table 4: Comparison between our method, MapNet and VidLoc. Mean (top) and median (bottom) 6-DoF pose errors on the Oxford RobotCar dataset are reported.



Figure 5: Qualitative results on the RobotCar dataset.

are used. MapNet is retrained on the new query video with the ground truth from previous predictions. Since VidLoc does not produce sufficiently accurate predictions, we do not retrain the network for subsequent query videos.

Our method outperforms MapNet and VidLoc in terms of the mean errors (see Table 4), and also has a smoother predicted trajectory than MapNet (see Fig. 5). In addition, while our method improves localization accuracy after updating the database (See Table 1), MapNet's results is worse after retraining (See Table 4). This is because MapNet is retrained on a noisy ground truth. However, though our method is qualitatively better than MapNet, differences in median error is not obvious: this shows that median error is not a good criterion for VL, since gross errors are ignored.

Note that our method mainly performs PR; here, comparisons to VL methods are to show that a correct PR paired with simple pose inheritance can outperform VL methods in presence of appearance change. The localization error of our method can likely be improved by performing SfM on a set of images corresponding to the highest belief.

Table 3 reports training/updating time for our method and MapNet and VidLoc. Particularly, for Seq-1 and Seq-2, our method needs around 1.65 minute to construct the k-means tree and build the graph, while MapNet and Vid-Loc respectively require 11.6 and 14.1 hours for training. For updating a new query sequence, MapNet needs about 6 hours of retraining the network, whilst our method culls the database and combine graph nodes in less than 5 minutes. This makes our method more practical in a realistic scenario, in which the training data is acquired continuously.

7. Conclusion

This paper proposes a novel method for scalable place recognition, which is lightweight in both training and testing when the data is continuously accumulated to maintain all of the appearance variation. From the results, our algorithm shows significant potential towards achieving longterm autonomy in terms of scalable inference and database update. Memory scalability is leaved as our future work.

References

- Relja Arandjelovic, Petr Gronat, Akihiko Torii, Tomas Pajdla, and Josef Sivic. Netvlad: Cnn architecture for weakly supervised place recognition. In *CVPR*, 2016.
- [2] Olivier Aycard, François Charpillet, Dominique Fohr, and J-F Mari. Place learning and recognition using hidden markov models. In *IROS*, 1997.
- [3] Artem Babenko and Victor Lempitsky. Tree quantization for large-scale similarity search and classification. In *CVPR*, 2015.
- [4] Eric Brachmann, Alexander Krull, Sebastian Nowozin, Jamie Shotton, Frank Michel, Stefan Gumhold, and Carsten Rother. DSAC-differentiable RANSAC for camera localization. In CVPR, 2017.
- [5] Eric Brachmann and Carsten Rother. Learning less is more-6d camera localization via 3d surface regression. In *CVPR*, 2018.
- [6] Eric Brachmann and Torsten Sattler. Visual Localization: Feature-based vs. Learned Approaches. https://sites.google.com/view/ visual-localization-eccv-2018/home, 2018.
- [7] Samarth Brahmbhatt, Jinwei Gu, Kihwan Kim, James Hays, and Jan Kautz. Geometry-aware learning of maps for camera localization. In *CVPR*, 2018.
- [8] Winston Churchill and Paul Newman. Experience-based navigation for long-term localisation. *The International Journal of Robotics Research*, 2013.
- [9] Ronald Clark, Sen Wang, Andrew Markham, Niki Trigoni, and Hongkai Wen. VidLoc: A deep spatio-temporal model for 6-DoF video-clip relocalization. In *CVPR*, 2017.
- [10] Mark Cummins and Paul Newman. Fab-map: Probabilistic localization and mapping in the space of appearance. *The International Journal of Robotics Research*, 2008.
- [11] Mark Cummins and Paul Newman. Appearance-only slam at large scale with fab-map 2.0. *The International Journal of Robotics Research*, 2011.
- [12] Matthijs Douze, Hervé Jégou, and Florent Perronnin. Polysemous codes. In *ECCV*, 2016.
- [13] Dorian Gálvez-López and Juan D Tardos. Bags of binary words for fast place recognition in image sequences. *IEEE Transactions on Robotics*, 2012.
- [14] Mordechai Haklay and Patrick Weber. OpenStreetMap: User-generated street maps. *IEEE Pervasive Computing*, 2008.
- [15] Peter Hansen and Brett Browning. Visual place recognition using hmm sequence matching. In *IROS*, 2014.
- [16] Herve Jegou, Matthijs Douze, and Cordelia Schmid. Product quantization for nearest neighbor search. *TPAMI*, 2011.
- [17] Hervé Jégou, Matthijs Douze, Cordelia Schmid, and Patrick Pérez. Aggregating local descriptors into a compact image representation. In *CVPR*, 2010.
- [18] Hervé Jégou and Andrew Zisserman. Triangulation embedding and democratic aggregation for image search. In CVPR, 2014.
- [19] Alex Kendall and Roberto Cipolla. Modelling uncertainty in deep learning for camera relocalization. In *ICRA*, 2016.

- [20] Alex Kendall and Roberto Cipolla. Geometric loss functions for camera pose regression with deep learning. In CVPR, 2017.
- [21] Alex Kendall, Matthew Grimes, and Roberto Cipolla. Posenet: A convolutional network for real-time 6-dof camera relocalization. In *CVPR*, 2015.
- [22] Jana Kosecka and Fayin Li. Vision based topological markov localization. In *ICRA*, 2004.
- [23] Yasir Latif, Ravi Garg, Michael Milford, and Ian Reid. Addressing challenging place recognition tasks using generative adversarial networks. In *ICRA*, 2018.
- [24] Vincent Lepetit, Francesc Moreno-Noguer, and Pascal Fua. Epnp: An accurate o (n) solution to the pnp problem. *IJCV*, 2009.
- [25] David G Lowe. Distinctive image features from scaleinvariant keypoints. *IJCV*, 2004.
- [26] Stephanie Lowry, Niko Sünderhauf, Paul Newman, John J Leonard, David Cox, Peter Corke, and Michael J Milford. Visual place recognition: A survey. *IEEE Transactions on Robotics*, 2016.
- [27] Will Maddern, Geoffrey Pascoe, Chris Linegar, and Paul Newman. 1 year, 1000 km: The oxford robotcar dataset. *The International Journal of Robotics Research*, 2017.
- [28] Michael J Milford and Gordon F Wyeth. SeqSLAM: Visual route-based navigation for sunny summer days and stormy winter nights. In *ICRA*, 2012.
- [29] Marius Muja and David G Lowe. Scalable nearest neighbor algorithms for high dimensional data. *TPAMI*, 2014.
- [30] Gerhard Neuhold, Tobias Ollmann, Samuel Rota Bulò, and Peter Kontschieder. The mapillary vistas dataset for semantic understanding of street scenes. In *ICCV*, 2017.
- [31] Horia Porav, Will Maddern, and Paul Newman. Adversarial training for adverse conditions: Robust metric localisation using appearance transfer. In *ICRA*, 2018.
- [32] Cosimo Rubino, Alessio Del Bue, and Tat-Jun Chin. Practical motion segmentation for urban street view scenes. In *ICRA*, 2018.
- [33] Stuart J Russell and Peter Norvig. Artificial intelligence: a modern approach. Malaysia; Pearson Education Limited,, 2016.
- [34] Torsten Sattler, Bastian Leibe, and Leif Kobbelt. Efficient & effective prioritized matching for large-scale image-based localization. *TPAMI*, 2017.
- [35] Torsten Sattler, Will Maddern, Carl Toft, Akihiko Torii, Lars Hammarstrand, Erik Stenborg, Daniel Safari, Masatoshi Okutomi, Marc Pollefeys, Josef Sivic, et al. Benchmarking 6DOF outdoor visual localization in changing conditions. In *CVPR*, 2018.
- [36] Nikolay Savinov, Alexey Dosovitskiy, and Vladlen Koltun. Semi-parametric topological memory for navigation. In *ICLR*, 2018.
- [37] Sebastian Thrun, Wolfram Burgard, and Dieter Fox. A probabilistic approach to concurrent mapping and localization for mobile robots. *Autonomous Robots*, 1998.
- [38] Sebastian Thrun, Wolfram Burgard, and Dieter Fox. Probabilistic robotics. 2005.

- [39] Akihiko Torii, Relja Arandjelovic, Josef Sivic, Masatoshi Okutomi, and Tomas Pajdla. 24/7 place recognition by view synthesis. In *CVPR*, 2015.
- [40] Akihiko Torii, Relja Arandjelovic, Josef Sivic, Masatoshi Okutomi, and Tomas Pajdla. 24/7 place recognition by view synthesis. In CVPR, 2015.
- [41] Florian Walch, Caner Hazirbas, Laura Leal-Taixe, Torsten Sattler, Sebastian Hilsenbeck, and Daniel Cremers. Imagebased localization using lstms for structured feature correlation. In *ICCV*, 2017.
- [42] Peng Wang, Ruigang Yang, Binbin Cao, Wei Xu, and Yuanqing Lin. Dels-3d: deep localization and segmentation with a 3d semantic map. In *CVPR*, 2018.