

Visual Place Recognition: A Tutorial

Stefan Schubert[†], Peer Neubert[‡], Sourav Garg^{††}, Michael Milford^{‡‡} and Tobias Fischer^{‡‡}

I. INTRODUCTION

Localization is an essential capability for mobile robots, enabling them to build a comprehensive representation of their environment and interact with the environment effectively towards a goal. A rapidly growing field of research in this area is Visual Place Recognition (VPR), which is the ability to recognize previously seen places in the world based solely on images. The volume of published research on VPR has shown a significant and continuous growth over the years, from two papers with “visual place recognition” and seven papers with “place recognition” in the title in 2006, to 65 and 163 papers, respectively, in 2022¹. A number of survey and benchmarking papers have discussed the challenges, open questions, and achievements in the field of VPR [1], [2], [3], [4].

This present work is the first tutorial paper on visual place recognition. It unifies the terminology of VPR and complements prior research in two important directions:

- 1) It provides a systematic introduction for newcomers to the field, covering topics such as the formulation of the VPR problem, a generic algorithmic pipeline, an evaluation methodology for VPR approaches, and the major challenges for VPR and how they may be addressed.
- 2) As a contribution for researchers acquainted with the VPR problem, it examines the intricacies of different VPR problem types regarding input (database or query set), data processing (online or offline) and output (one or multiple matches per query image). The tutorial also discusses the subtleties behind the evaluation of VPR algorithms, e.g., the evaluation of a VPR system that has to find *all* matching database images per query, as opposed to just a *single* match.

Practical code examples in Python illustrate to prospective practitioners and researchers how VPR is implemented and evaluated. The corresponding source code is avail-

[†] S. S. is with the Chemnitz University of Technology, Germany. Email: stefan.schubert@etit.tu-chemnitz.de

[‡] P. N. is with the University of Koblenz, Germany.

^{††} S. G. is currently with the University of Adelaide, Australia (work done while at the Queensland University of Technology, Australia).

^{‡‡} M. M. and T. F. are with the QUT Centre for Robotics, Queensland University of Technology, Australia.

S. S. acknowledges support by the German Federal Ministry for Economic Affairs and Climate Action. S. G., M. M. and T. F. acknowledge support by the QUT Centre for Robotics, funding from ARC Laureate Fellowship FL210100156 to M. M., and a grant from Intel Labs to T. F. and M. M.

The authors would like to thank Dr Mark Zolotas, Dr Alejandro Fontan and Somayeh Hussaini for valuable insights on drafts of the paper.

¹Source: Google Scholar with query *allintitle: "title"*

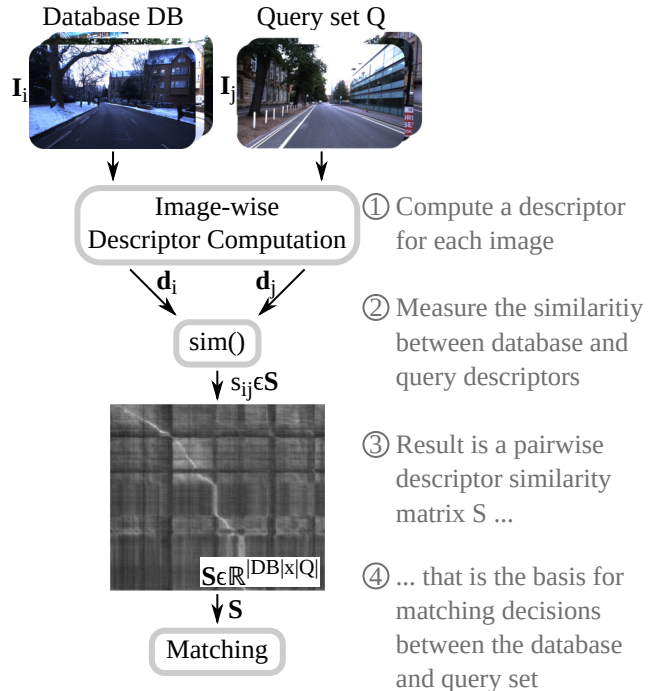


Fig. 1: This figure illustrates the key steps and components of VPR as outlined in Section I. Historically, the matching decisions in step 4 were used for loop closure in Simultaneous Localization and Mapping (SLAM); Section II provides an overview of the history and relevance of the VPR problem. While this figure illustrates a common use case where incoming imagery in the query set is compared to a database, Section III distinguishes different VPR problem categories based on this pipeline and also relates them to VPR use cases. The details of each shown computational step will be discussed in Section IV, followed by details on the evaluation of VPR pipelines in Section V. (Photos from [5])

able online, along with a list of other open-source implementations from the literature: https://github.com/stschubert/VPR_Tutorial. The link also provides a Jupyter notebook written in Python that guides through a basic VPR pipeline. It allows users to experiment with additional image descriptors, benchmark datasets, and evaluation metrics.

The following Section I-A provides a basic introduction to the VPR problem. Section I-B then outlines the structure of the remainder of this tutorial.

A. The basics of Visual Place Recognition (VPR)

VPR involves matching one or multiple image sets in order to determine which images show the same places in the world.

These image sets are typically recorded with a mobile device such as a cell phone or an AR/VR headset, or with a camera mounted on a variety of platforms, such as a robot, uncrewed aerial vehicle, car, bus, train, bicycle, or boat.

Essentially, VPR is an image retrieval problem where the context is to recognize previously seen places. This context provides additional information and structure beyond a general image retrieval setup. Many VPR methods exploit the context to match images of the same places in a wide range of environments, including those with significant appearance and viewpoint differences. For example, one piece of additional information that is often exploited is that *consecutive* images taken by a camera mounted on a car will depict spatially close places in the world.

Figure 1 provides an overview of the typical steps and components of a basic VPR pipeline. Given a reference set composed of database images $\mathbf{I}_i \in DB$ of known places, and one or multiple query images $\mathbf{I}_j \in Q$, the goal is to find matches between these two sets, i.e., those instances where image j from the query set shows the same place as image i from the database. To find these matches, it is essential to compute one or multiple descriptors \mathbf{d} for each image – these descriptors should be similar for images showing the same place and dissimilar for different places. A descriptor is typically represented as a numerical vector (e.g., 128-D or 4,096-D). Conceptually, we can think of a matrix \mathbf{S} of all pairwise descriptor similarities s_{ij} between the database and query images as the basis for deciding which images should be matched. In practice, we must carefully choose the algorithms used to compute and compare the image descriptors \mathbf{d} , taking into account the specific challenges and context of the VPR problem at hand. The remainder of this tutorial will provide more detail on these aspects and discuss the VPR problem from a broader theoretical and practical perspective.

B. The VPR problem and its details as reflected in this tutorial

Section II of this tutorial will outline the relevance and history of the VPR problem, as well as its relation to other areas, particularly its importance for topological Simultaneous Localization And Mapping (SLAM), where the database DB corresponds to the set of previously visited places in the map. In fact, one of the original drivers for VPR research was the generation of loop closures for SLAM systems, that is, recognizing a place when revisiting it (e.g., in a loop) and tying the current observation with that already in the map (i.e., closure) [6]. One of the earliest examples of such a topological SLAM system is FAB-MAP [7], also referred to as ‘appearance-only SLAM’, where loop closure generation is based on appearance only (i.e., images), thus different from 3D/metric SLAM systems such as ORB-SLAM [8] where the map and the visual landmarks are expressed in 3D.

The definition of a “*place*” is an integral aspect of VPR. In this tutorial, we follow the definition that two images must have some visual overlap, i.e., shared image content like same buildings, to be considered as “taken at the same place” [2]. This definition allows to subsequently estimate the camera transformation between matched images for tasks

like visual localization, mapping, or SLAM – indeed, the required amount of visual overlap depends on the specific application. We note that an alternative definition used in particular by some researchers [9] is that two places are matching purely based on their position, without taking the orientation, and in turn visual overlap, into account. Section III will present different applications for VPR and discuss the various subtypes of VPR problems that arise from variations in the available input data, the required data processing, and the requested output.

VPR algorithms are often tailored to the particular properties of an application. Section IV will provide details on a *generic VPR pipeline* that serves as a common basis for diverse practical settings and their unique characteristics. From this section onward, this tutorial includes practical code examples in Python.

It is important to note that not all VPR algorithms address the same VPR problem, e.g., regarding the requested number of image matches per query. This is particularly critical when it comes to evaluating and comparing the performance of different VPR algorithms. Section V explains and discusses the evaluation pipelines that consider various datasets, ground truth subtleties, and different performance metrics.

The properties of the underlying data have a significant impact on the difficulty of the resulting VPR problem and the suitability of a particular algorithm. Section VI will discuss challenges such as severe appearance changes due to varying illumination or weather conditions, large viewpoint changes between images of the same place, and perceptual aliasing, i.e., the challenge that images taken at two distinct places can appear remarkably similar. This section will also present common ways of addressing these challenges to improve robustness, performance, runtime and memory efficiency. These approaches include methodological extensions of the general purpose pipeline that partially build upon a robotic context (e.g., with image sets recorded as videos along trajectories) where VPR differs from pure image retrieval. This often allows the exploitation of additional knowledge and information such as spatio-temporal sequences (i.e., consecutive images in the database DB and query Q are also neighboring in the world) or intra-set similarities (i.e., similarities *within* DB or Q).

II. HISTORY, RELEVANCE AND RELATED AREAS

Visual Place Recognition (VPR) research can be traced back to advances in visual SLAM, visual geo-localization, and image retrieval applied to images of places [10]. In the robotics literature, VPR has historically been called loop closure detection and was mainly used for this purpose for visual SLAM [10]. VPR gained more prominence in the field as the earlier metric SLAM methods based on global and local bundle adjustment techniques could only handle limited-size environments, thus paving way for topological SLAM techniques based on bag-of-words approaches, such as FAB-MAP [7]. In addition to its relevance within SLAM pipelines, VPR also remains a crucial component of localization-only pipelines where the map is available a priori.

Early VPR research primarily focused on place recognition under constant or slightly varying environmental conditions. Addressing appearance changes due to more severe condition changes, such as day-night cycles or seasonal shifts, emerged in the late 2000s. These methods relied for example on local feature matching [11] or on continuously updating appearance-based maps [12]. Since then, research on VPR under challenging conditions has steadily increased, for example tackling the challenging day-night shift [13]. Recent works make heavy use of datasets with condition changes that have appeared since 2012 [1], [3]. In 2014, the use of deep learning for VPR [14] emerged as a way to handle challenging data and has since proven effective in changing environments [15]. In addition to images and image descriptors, VPR research has also explored the use of additional information, such as sequences, intra-set similarities, weak GPS signals, or odometry, to improve performance [16].

In terms of the relationship between VPR and other fields, we recommend the following tutorials: [10] provides an overview of probabilistic SLAM and includes a section on loop closure detection, although a lot of progress has been made in VPR as this tutorial was published more than 15 years ago. [6] specifically investigates the loop closure problem in SLAM. [17] discusses visual odometry, which involves estimating the ego-motion of an agent based on visual input. Visual odometry is thus complementary to VPR, and can be combined with VPR to detect loop closures when building a SLAM system. It is important to note, however, that the scope of this tutorial is limited to providing an accessible introduction to VPR and its core concepts. Aspects such as the integration of VPR methods into a complete SLAM or re-localization system are beyond the scope of this tutorial and would require discussing many additional aspects, such as batch optimization, which are not directly related to VPR.

Beyond loop-closure detection, VPR is necessary if global position sensors such as global navigation satellite systems (GNSS) like GPS, Galileo, or BeiDou are not available or are inaccurate. In urban environments, buildings or other structures can lead to “urban canyons” that block line-of-sight satellite signals, causing occlusions that prevent a GNSS receiver from obtaining accurate position information. In addition to occlusions, reflections of GNSS signals off buildings and other structures, so called non-line-of-sight signals, can further hinder the GNSS accuracy. This issue is not limited to urban environments, as similar occlusions and reflections can occur in natural environments, such as in valleys or canyons. Similarly, indoor environments and caves also hinder GNSS due to the absorption or reflection of satellite signals by walls.

Alternatively, VPR can serve as a redundant component in autonomous systems for fault tolerance and general GNSS outages, such as satellite service disruptions, degradation, or position/time anomalies. It is worth noting that all GNSS systems can potentially be hacked or blocked for non-military use by a central authority. Other systems may not be equipped with a GNSS receiver due to cost or security concerns. In the

case of robotic extraterrestrial missions, installing a GNSS system may be too expensive or time-consuming.

III. VPR PROBLEM CATEGORIES AND USE CASES

In the localization and mapping literature, VPR has been used in different ways depending on three key attributes of its formulation: the *input*, which deals with how the reference and query images are made available (i.e., single-session vs. multi-session); *data processing*, that defines the mode of operation (i.e., online vs. batch); and *output*, that determines the kind of expected output (i.e., single-best-match vs. multi-match). The following Section III-A explains these problem categories in more detail. Section III-B then presents different VPR use cases using these categories. Table I summarizes these use cases, along with their required input and data processing. Note that there might be exceptions and deviations from these categories, such as [18] that uses multiple disjoint sequences as reference. However, we believe that the proposed taxonomy serves as a good starting point for future research to organize the various VPR use cases.

A. VPR problem categories

We distinguish three main dimensions along which VPR problems can vary, creating different VPR problem categories or subtypes:

- 1) **Input – single-session VPR vs. multi-session VPR:** *Are there two separate input sets, one for the database DB and one for the query Q , or is it a single set that is compared to itself?* Single-session VPR is the matching of images within a single set of images, so that the query set Q equals the database DB (i.e., $Q = DB$). A practical consideration in this case is the suppression of matches with recently acquired images – while full SLAM systems typically rely on a motion model for such suppression, standalone VPR systems often use heuristics. In contrast, multi-session VPR is the matching of the two disjoint image sets (i.e., $DB \cap Q = \emptyset$), which were recorded at different times (e.g., summer and winter) or by different platforms (e.g., mobile robot and cell phone).
- 2) **Data processing – online VPR vs. batch VPR:** *Are the images available and processed individually, one after the other, or are they all available in a single batch from the beginning?* Online VPR has to deal with a growing set Q (i.e., $Q \neq \text{const}$) and a set DB that is either given (i.e., $DB = \text{const}$) or also growing (i.e., $DB \neq \text{const}$). In contrast, batch VPR can build upon the full sets Q (i.e., $Q = \text{const}$) and DB (i.e., $DB = \text{const}$). Growing image sets in the case of online VPR limit the number of viable methods. For example, approaches like descriptor standardization [19] based on the statistics of *all* image descriptors or similarity matrix decomposition [20] cannot be used without modifications. This is further discussed in Section VI. Note that for single-session VPR, the difference between a growing or constant query set pertains to whether VPR is being performed online or in batch mode.

3) **Output – single-best-match VPR vs. multi-match VPR:**

Is the intended output for a query image a single image from the database that shows the same place, or do we request all images of this place? Single-best-match VPR only returns the best matching database image $I_i^* \in DB$ per query image $I_j \in Q$. In contrast, the aim in multi-match VPR is finding *all* matching database images for each query image. In practice, the difference between single-best-match VPR and multi-match VPR often boils down to finding either the maximum similarity between a query and all database images or all similarities above a certain threshold, as shown in Section IV-E. Identifying all matching images is often more challenging than finding only one correct match, as it requires an explicit decision for each database image whether it shows the same place as the query image or not [1].

Let us illustrate these problem categories with the example of determining the rough pose [x, y, heading, floor] of a cell phone in a building, e.g., to guide persons to desired places. To achieve this, we first need to map the building before the person can use their cell phone to localize in the building. For this first step of mapping the building, we could use a manually controlled mobile robot equipped with a camera to collect a query set Q^{mapping} of images together with some additional sensor data like odometry. Given these images of all places, we can run a mapping algorithm that processes all images and other data to obtain a metric map of the building, which associates all images in Q^{mapping} with metric poses. Part of this mapping is a *single-session batch multi-match VPR* for loop closure detection that compares the whole image set Q^{mapping} (batch VPR) to itself (single-session VPR) to find *all* loop closures for each image (multi-match VPR). Here, batch processing the whole set Q^{mapping} allows the application of computationally expensive but accurate algorithms.

After mapping (potentially years later), the second step is the actual localization of a cell phone using its camera stream. When localizing, we treat the robot’s mapping query set Q^{mapping} as database DB^{loc} and compare it to query images Q^{loc} from a cell phone’s camera. To determine the location of a cell phone, a *multi-session online single-best-match VPR* can be used that compares the stream of query images Q^{loc} to the fixed database DB^{loc} (multi-session VPR) online (online VPR) to find the best matching database image (single-best-match VPR) with its corresponding pose information.

In summary, VPR can be used for a variety of different use cases, as discussed in more detail in the following Section III-B and shown in Table I. Here, each use case typically requires a certain combination of the *input* (single-session/multi-session VPR) and *data processing* (online/batch VPR) VPR categories. The choice of the *output* category (single-best-match/multi-match VPR) also depends on the use case, and in particular on the algorithm that is used after VPR. For example, for pure (re-)localization [21], one may only need a single best match, so that the choice of the output category depends mainly on the use case, while for graph-based SLAM, the required output category also depends on

TABLE I: Combinations of the VPR *input* and *data processing* categories (cf. Section III-A) with corresponding use cases (cf. Section III-B). In single-session VPR, there is a single input set that is compared to itself. This is for example the case in online SLAM, where this set grows while the robot is exploring its environment (online VPR), and in mapping, where pre-recorded data is processed in a batch manner. In the case that the input consists of two sets, the database DB and the query set Q (multi-session VPR), one can again distinguish the case that data needs to be processed online as it is being collected (online VPR) or in a batch as in multi-session mapping. In multi-session online VPR, DB can be either growing (as in multi-robot mapping) or fixed (as in visual (re-)localization).

		Data processing		
		Online VPR	Batch VPR	
Input	Single-session VPR	Online SLAM		Mapping
	Multi-session VPR	DB grows Multi-Robot Mapping	DB const. Visual (Re-)Localization	Multi-Session Mapping

the post-processing after VPR, as explained in the following. In graph-based SLAM [10], each node encodes the pose of an image in Q . The corresponding edges of connected nodes represent the transformation between them. An edge can be established either between temporally consecutive nodes (using the odometry) or between nodes that were identified as loop closure by VPR. Here, single-best-match VPR could be used to match and fuse two nodes which correspond to the same place to represent each place always by only one node. Alternatively, multi-match VPR could be used to create multiple edges between all existing nodes of the same place. This is particularly helpful if we cannot guarantee that there is a single node for each place in the graph, or if we perform a batch optimization of the poses using a robust optimization approach that can benefit from the additional information provided by multiple matches while handling potential outlier matchings.

B. VPR use cases

VPR is a key component in a variety of robotic applications, including autonomous driving, agricultural robotics, and robotic parcel delivery, as well as in the creation of a metaverse. Some common tasks that VPR is used for include:

- 1) **Candidate selection for 6 DoF visual localization** [21]: 6 Degree Of Freedom (DoF) visual localization (also termed city-scale/natural geo-localization) involves estimating the 6D pose (position and orientation) of a camera in a particular environment. *Multi-session online VPR with fixed DB* is used to select candidates $I_i \in DB$ that have the highest similarity to the current query images $I_j \in Q$ (cf. Figure 2, top). These candidates can then be used for a

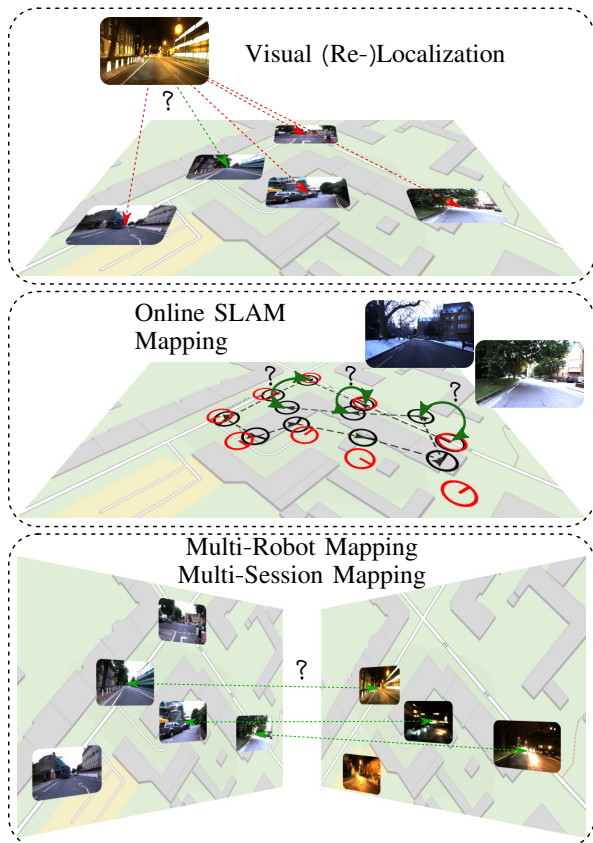


Fig. 2: Overview of Visual Place Recognition use cases. Top: VPR is used to select a small set of candidate images from the database for visual (re-)localization, which are processed with computationally expensive 6-DoF pose estimation methods. Middle: VPR can be used to detect loop closures in a SLAM pipeline (green arrows) or to re-localize after mislocalization or if the robot was moved (kidnapped robot). Bottom: In multi-session mapping, VPR can again be used for loop closure detection, but this time in a batch manner where all images are known in advance. In multi-robot mapping, VPR is used to merge maps by detecting places that have been visited by multiple robots. (Map data from OpenStreetMap)

computationally intensive 6D pose estimation using local image descriptors and more complex algorithms, which would be infeasible for the complete DB set.

- 2) **Loop closure detection and re-localization for online SLAM** [10]: Online SLAM is used to estimate the current pose of a camera while creating a map of the environment at the same time. *Single-session online VPR* is used for loop closure detection (i.e., the recognition of previously visited places), as shown in Figure 2 (mid) to compensate for accumulated errors in odometry data and create a globally consistent map. It is also used for re-localization in the event of mislocalization or if the camera/robot was moved by hand (known as the kidnapped robot problem).
- 3) **Loop closure detection for mapping** [10]: Mapping (also full SLAM or offline SLAM) involves estimating the entire path at once to generate a map. This allows for the use of *single-session batch VPR* for loop closure detection

(cf. Figure 2, mid), which is based on slower but more robust algorithms that run on powerful hardware.

- 4) **Loop closure detection for multi-session mapping** [22]: Multi-session mapping combines the results of multiple SLAM missions performed repeatedly over time in the same environment. *Multi-session batch VPR* is used to find shared places between the individual maps of all missions for map merging (cf. Figure 2, bottom). Alternatively, *multi-session online VPR with a given DB* can be used to detect previously mapped areas (potentially for loop closure) and include unseen areas of the map in real-time.
- 5) **Detection of shared places for multi-robot mapping** [23]: Multi-robot mapping (also termed decentralized SLAM) involves the distributed mapping of an environment using multiple robots. Here, *multi-session online VPR with a growing DB* is used to find shared places between the individual maps of each robot for subsequent map merging, as shown in Figure 2 (bottom).

In summary, this section provided an overview of the different problem categories and corresponding subtypes of VPR and discussed common use cases where VPR is applied.

IV. A GENERIC PIPELINE FOR VPR

This section outlines a generic pipeline for Visual Place Recognition (VPR). The steps involved in this pipeline are shown in Figure 1. The inputs to the pipeline are two sets of images DB and Q (these may be the same for single-session VPR, as explained in Section III). The pipeline produces matching decisions, meaning that for each query image $I_j \in Q$, one or more database images $I_i \in DB$ can be associated. The pipeline includes these intermediate steps and components: 1) computing image-wise descriptors, 2) pairwise comparing of descriptors to 3) create a descriptor similarity matrix S , and 4) making matching decisions. In the following subsections, we will discuss each of these elements in more detail. Extensions to this generic pipeline that can be used to improve performance and robustness against various challenges are presented in Section VI.

A. Inputs: The database and query image sets


To recap, two sets of images serve as the input in a VPR pipeline: the database set DB and a set of current images in the query set Q . The DB set, which is also called the reference set, represents a map of known places, and is often recorded under ideal conditions (e.g., sunny), or by a different platform than Q (e.g., a second robot). The query set Q , on the other hand, is the “live view” recorded by a different platform than DB or after DB – potentially days, months, or even years later. Both sets will have a geographical overlap and share some or all seen places.

There are different VPR problem categories: using just a query set Q (single-session VPR) or using both the DB and Q sets (multi-session VPR). Also, the image sets can either be specified before processing (batch VPR) or grow during an online run (online VPR).

Code Snippet 1 provides example code for loading a dataset with both image sets Q and DB , as well as the ground truth

```
# load dataset GardensPoint Walking [24] with two
# image sets DB and Q and ground truth
from datasets.load_dataset \
    import GardensPointDataset
dataset = GardensPointDataset()
DB, Q, GT, GTsoft = dataset.load()
```

Output:



Code Snippet 1: In this example, the input to a VPR algorithm are two disjoint image sets: the database DB and query Q . We load the GardensPoint Walking dataset [24] and ground-truth information about correspondences. This ground truth only serves for later evaluation and will neither be available nor required when deploying the algorithm.

matrices \mathbf{GT} and \mathbf{GT}^{soft} . Briefly, \mathbf{GT} is a logical matrix that indicates whether corresponding images show the same or different places, while \mathbf{GT}^{soft} is a dilated version of \mathbf{GT} that accounts for image pairs with small visual overlap, avoiding penalization for matches in such cases. We detail these matrices in Section V-B.

B. Image-wise descriptor computation

This section describes the process of computing image descriptors, which are abstractions of images that extract features from raw pixels in order to be more robust against changes in appearance and viewpoint (step 1 in Figure 1, see also Code Snippet 2). The tutorial covers two primary types of image descriptors:

- 1) **Holistic descriptors** (also called global descriptors) represent an image $\mathbf{I}_i \in DB, Q$ with a single vector $\mathbf{d}_i \in \mathbb{R}^d$ (cf. Code Snippet 2). This allows for efficient pairwise descriptor comparisons with low runtimes. Note that when exhaustive k-nearest neighbor search (kNN) is used to obtain the nearest neighbors for a candidate selection of similar database descriptors, the execution time scales linearly with both the descriptor dimension and the number of images contained in the database.
- 2) **Local descriptors** encode an image \mathbf{I}_i with a set $D_i = \{\mathbf{d}_k \mid k = 1, \dots, K\}$ of vectors $\mathbf{d}_k \in \mathbb{R}^d$ at K regions of interest. They often provide better performance than holistic descriptors, but require computationally expensive methods for local feature matching like a left-right check

```
# compute holistic HDC-DELf descriptors [28]
from feature_extraction. \
    feature_extractor_holistic import HDCDELf
feature_extractor = HDCDELf()
DDB = feature_extractor.compute_features(DB)
DQ = feature_extractor.compute_features(Q)
```

Code Snippet 2: The main source of information about image correspondences are image descriptors. Since holistic image descriptors allow for efficient pairwise descriptor comparisons, we compute a holistic HDC-DELf [28] descriptor for each image (step 1 in Figure 1).

(also termed mutual matching), a homography estimation, a computation of the epipolar constraint, or deep-learning matching techniques. Therefore, local descriptors are typically used in a hierarchical pipeline, where first the holistic descriptors are used to retrieve the top- K matches, which are then re-ranked using local descriptor matching.

The abstraction of the VPR pipeline in terms of holistic and local descriptors serves as the foundation for many localization, mapping, and SLAM solutions. Alternative approaches include place classification [25], regional descriptors [26], and incremental bags of binary words [27]. Furthermore, in Section VI we list common shortcomings of this VPR pipeline and ways to address them.

To convert a set of local descriptors from a single image into a holistic descriptor, one can use **local feature aggregation** methods like Bag of Visual Words (BoVW), Vector of Locally Aggregated Descriptors (VLAD) or Hyper-Dimensional Computing (HDC) [28]. In a hierarchical pipeline, this allows a local descriptor to be used for both candidate selection (after aggregation) and verification (with the raw local descriptors).

As the descriptor computation is one of the first steps in a pipeline for VPR, it has a significant impact on the performance of subsequent steps and the overall performance of the VPR system. The algorithm used to obtain the descriptors determine how well the descriptors are suited for a specific environment, the degree of viewpoint change, or the type of environmental condition change. For example, CNN-based holistic descriptors like AlexNet-conv3 [15] perform well in situations with low or negligible viewpoint changes, but perform poorly with large viewpoint changes. On the other hand, VLAD-based descriptors like NetVLAD [29] tend to perform better in settings with large viewpoint changes.

Additionally, the specific training data of deep-learned descriptors affect the performance in different environments. For example, some descriptors may perform better in urban environments, while others may be more effective in natural environments or in specific geographic regions such as Western cities [30].

C. Descriptor similarity between two images

To compare the image descriptors of two images, a measure of similarity or distance must be calculated (see step 2 of Figure 1 and Code Snippet 3). This process compares the descriptors \mathbf{d}_i and \mathbf{d}_j (holistic) or D_i and D_j (local) of

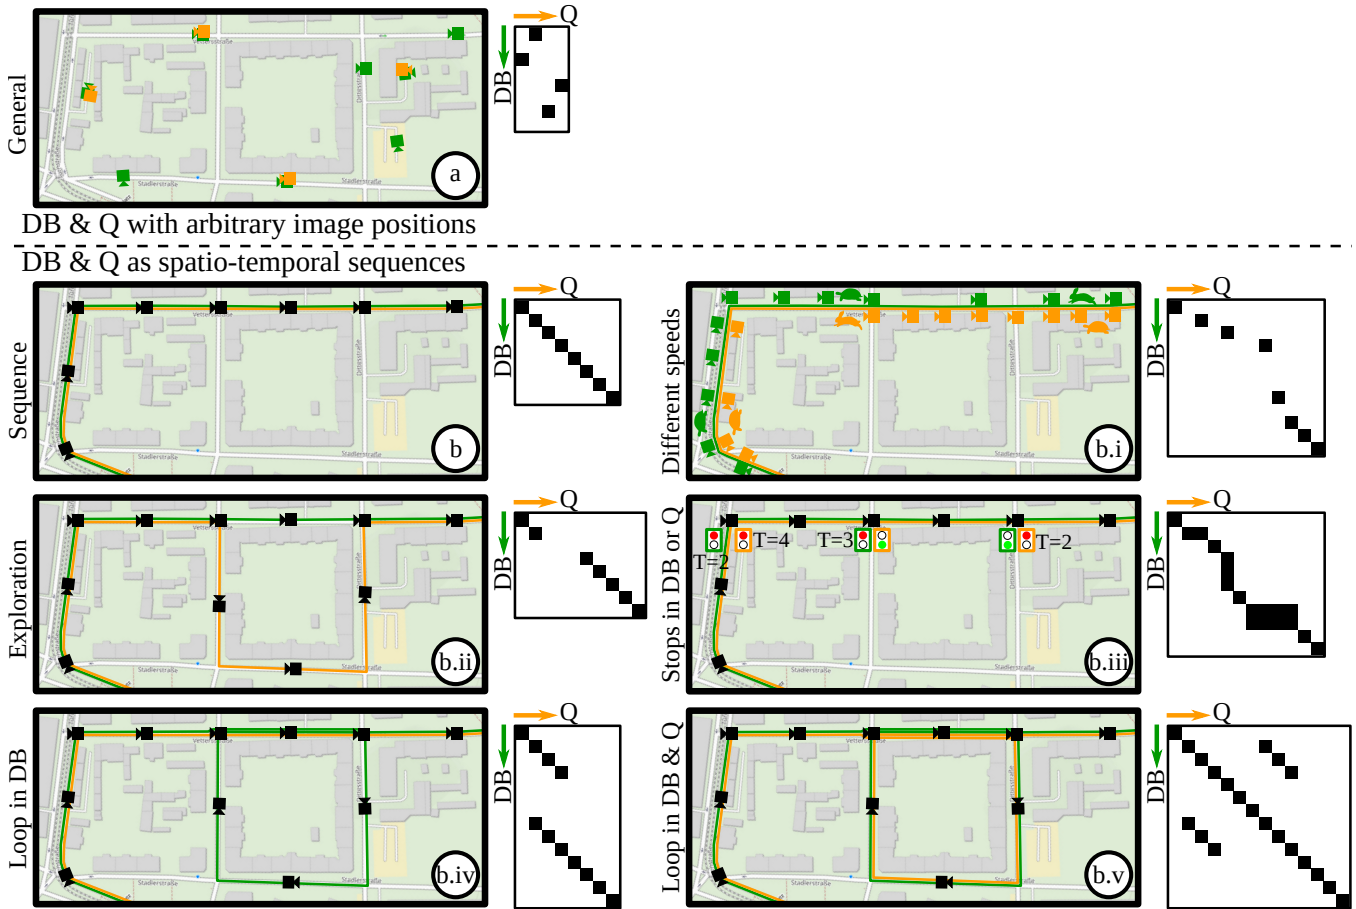


Fig. 3: Relation between the similarity matrix \mathbf{S} and the trajectory during the database and query run. Green and orange cameras depict images in DB and Q , respectively. Green and orange lines indicate that images were recorded as video along a trajectory (also called spatio-temporal sequence). In b.i), a rabbit/turtle indicate fast/slow speeds when traversing the route, and similarly in b.iii) traffic lights indicate stops in Q ($T=2$), DB ($T=3$) or both ($T=\{2,4\}$) for T time steps.

images i and j . Note that similarity s_{ij} and distance $dist_{ij}$ can be related through inversely proportional functions such as

$$s_{ij} = -dist_{ij}, \quad (1)$$

or the reciprocal

$$s_{ij} = \frac{1}{dist_{ij}}. \quad (2)$$

Holistic descriptors can be compared more efficiently than local descriptors, as they only require simple and computationally efficient metrics like the cosine similarity

$$s_{ij} = \frac{\mathbf{d}_i^T \cdot \mathbf{d}_j}{\|\mathbf{d}_i\| \cdot \|\mathbf{d}_j\|}, \quad (3)$$

or the negative Euclidean distance

$$s_{ij} = -\|\mathbf{d}_i - \mathbf{d}_j\|. \quad (4)$$

In contrast, comparing local descriptors requires more complex and computationally expensive algorithmic approaches, as previously mentioned in Section IV-B.

D. The pairwise similarity matrix \mathbf{S}

The pairwise descriptor similarity matrix \mathbf{S} is a key component of VPR. As shown in step 3 of Figure 1, \mathbf{S} contains all calculated similarities s_{ij} between the descriptors of images in the database and query sets. In single-session VPR, \mathbf{S} has dimensions $|Q| \times |Q|$, while in multi-session VPR, \mathbf{S} has dimensions $|DB| \times |Q|$. Depending on the approach used, \mathbf{S} may be dense (if all descriptors are compared) or sparse (if only a subset of descriptors is compared using approximate nearest neighbor search or sequence-based comparison strategies).


The overall appearance of \mathbf{S} is influenced by the camera's trajectories during acquisition of Q and DB , as illustrated in Figure 3. The pattern of high similarities within \mathbf{S} can have a significant impact on the performance of the VPR pipeline, and may enable or hinder the use of certain algorithmic steps for performance improvements. The following relations between camera trajectories and the appearance of \mathbf{S} can be observed (cf. Figure 3 for corresponding examples in a map):

- a) **General:** If images in DB and Q are recorded at arbitrary positions without a specific order, there are no discernible

```
# compare all descriptors using cosine similarity
# normalize descriptors
DDB = DDB /
    np.linalg.norm(DDB, axis=1, keepdims=True)

DQ = DQ /
    np.linalg.norm(DQ, axis=1, keepdims=True)
# dot product between descriptors using matrix
# multiplication
S = np.matmul(DDB, DQ.transpose())
```

Output:



Code Snippet 3: To compare database and query descriptors to obtain the descriptor similarities \mathbf{S} (steps 2 and 3 in Figure 1), we use their cosine similarity (computed by the inner product of the normalized descriptor vectors). Although we might not want to compute the full similarity matrix \mathbf{S} of all possible pairs in a large-scale practical application, it can be useful for visual inspection purposes.

patterns in \mathbf{S} . This is typical for general visual localization and global geo-localization.

b) **Sequence:** If images in DB and Q are recorded along trajectories as spatio-temporal sequences (i.e., consecutive images are also neighbors in the world), continuous lines of high similarities may be observed in \mathbf{S} . This setup is typical for many robotic tasks, including online SLAM, mapping, and multi-robot/multi-session mapping (Section III). In this setup, sequence-based methods can be used for performance improvements (cf. Section VI). The camera’s trajectories can affect \mathbf{S} in the following ways:

- i) **Speed:** If the camera moves at the same speed in the same locations in DB and Q , lines of high similarities with 45° slope will be observed. Otherwise, the slope will vary.
- ii) **Exploration:** If a place shown in a query image Q is not present in DB , the line of high similarities will be discontinuous.
- iii) **Stops:** If the camera stops temporarily (zero velocity) during either the database run or the query run, it will result in multiple consecutive matches in the other set.

- **Stops in DB :** Stops in the database run will result in a vertical line (within the same column) of high similarities in \mathbf{S} .
- **Stops in Q :** Stops in the query run will result in a horizontal line (within the same row) of high similarities in \mathbf{S} .
- **Stops in DB & Q :** If the camera stops in both the database run and query run at the same place, a block of high similarities will be observed in \mathbf{S} .

- iv) **Loops in DB :** Loops in DB can result in multiple matching database images for a single query image in Q . Unlike stops, the multiple matching images due to a loop are not consecutive in their image set.
- v) **Loops in DB & Q :** Loops in DB and Q can result in additional matching query images for a single database image in DB . This results in a more complex structure of high similarities in \mathbf{S} .

E. Output: Matching decisions

The output of a VPR system is a set of matching decisions $m_{ij} \in \mathbf{M}$ (step 4 in Figure 1 and Code Snippet 4) with $\mathbf{M} \in \mathbb{B}^{|Q| \times |Q|}$ (single-session VPR) or $\mathbf{M} \in \mathbb{B}^{|DB| \times |Q|}$ (multi-session VPR) that indicate whether the i -th database/query image and the j -th query image show the same place ($m_{ij} = true$) or different places ($m_{ij} = false$). Existing techniques for matching range from choosing the best match per query or a simple thresholding of the pairwise descriptor similarities $s_{ij} \in \mathbf{S}$ to a geometric verification with a comparison of the spatial (using e.g., the epipolar constraint) or semantic constellation of the scene. For example in Code Snippet 4, M_1 is computed by selecting the best matching database image per query image, i.e., the maximum similarity s_{ij} per column in $\mathbf{S} \in \mathbb{R}^{|DB| \times |Q|}$ (single-best-match VPR). Another example is the computation of M_2 in Code Snippet 4, where a similarity threshold θ is applied to \mathbf{S} : If $s_{ij} \geq \theta$, the i -th and j -th images are assumed to show the same place (multi-match VPR). The next section is concerned with the performance evaluation of these outputs.

V. EVALUATION OF THE PERFORMANCE

This section is concerned with the evaluation of the matching decisions \mathbf{M} , or the pairwise similarities \mathbf{S} , which allows the comparison of different VPR methods. This requires datasets, corresponding ground truth, and performance metrics. In the following, we outline these components for evaluation and discuss their properties and potential pitfalls.

A. Datasets

For VPR, a dataset is composed of one or multiple image sets that have to be matched in order to find shared places. For example, the popular Nordland dataset [31] provides four image sets, one for each season, i.e., spring, summer, fall, and winter. These can be arbitrarily combined for VPR, but a typical choice might be to use *summer* as DB and *spring*, *fall* or *winter* as Q .


Existing datasets vary in the type of environment as well as in the type and degree of appearance and viewpoint change:

```
# match images based on S
from matching import matching

# best matching per query in S for
# single-best-match VPR
M1 = matching.best_match_per_query(S)


# find matches with S>=thresh using an auto-tuned
# threshold for multi-match VPR
M2 = matching.thresholding(S, thresh='auto')
```

Output:



M₁: M₂:

Examples for correct and wrong matches from M₂:



Code Snippet 4: The output of a VPR pipeline is typically a set of discrete matchings, i.e. pairs of query and database images. To obtain matchings for a query image from the similarity matrix (step 4 in Figure 1), we can either find the single best matching database image (M_1) or try to find all images in the database that show the same place as the query image (M_2).

The **type of environment** includes indoor environments, urban environments, suburban environments and natural environments like countryside, forests or lakes. **Appearance changes** occur due to dynamic objects like pedestrians, time of day with lighting changes and moving shadows or day versus night, weather that is sunny, cloudy, overcast, rainy, foggy, or snowy, seasons like spring, summer, fall and winter or dry and wet season, elapsed time with roadworks, construction sites or new and demolished buildings up to modern vs. historical imagery, or catastrophic scenarios, e.g., after an earthquake. **Viewpoint changes** between images of the same place range from nearly pixel-aligned to left-to-right side of the walkway, left-to-right side of the street, panoramic-aligned images to single image, panoramic-aligned images to panoramic-aligned images, bikeway-to-street, aerial-to-ground, or inside-

to-outdoor. For a comprehensive overview of existing datasets, please refer to [1], [3].

B. The ground truth

Ground truth data tells us which image pairs in a dataset show the same places, and which show different places. This data is necessary for evaluating the results of a place recognition method. The ground truth is either directly given as a set of tuples indicating which images in the database DB and the query set Q belong to same places, or it is provided via GNSS coordinates or poses using their maximum allowed distances. Alternatively, some datasets are sampled so that images with the same index in each image set show the same place.

Definition of the ground truth: To evaluate a VPR result, the definition of a logical ground truth matrix \mathbf{GT} is required. This matrix has the same dimensions as \mathbf{S} and \mathbf{M} , i.e., $\mathbf{GT} \in \mathbb{B}^{|Q| \times |Q|}$ or $\mathbf{GT} \in \mathbb{B}^{|DB| \times |Q|}$. The elements $gt_{ij} \in \mathbf{GT}$ define whether the i -th image in Q or DB and the j -th image in Q show the same place ($gt_{ij} = true$) or different places ($gt_{ij} = false$). Their values are set using the ground truth matches from the dataset.

An additional way of evaluating VPR performance that is used by some researchers is the soft ground truth matrix \mathbf{GT}^{soft} . The soft ground truth matrix addresses the problem that we do not expect a VPR method to match images with a very small visual overlap, i.e., $gt_{ij} = false$, as illustrated in Figure 4. However, if a method indeed matches these images with small overlap, we avoid penalization by setting $gt_{ij}^{soft} = true$. Image pairs without any visual overlap are also labeled $gt_{ij}^{soft} = false$. Therefore, \mathbf{GT}^{soft} is a dilated version of \mathbf{GT} , i.e., it contains all true values contained in \mathbf{GT} , as well as additional true values for image pairs with small visual overlap. Image pairs must be matched if

$$\mathbf{GT} = true . \quad (5)$$

Image pairs can, but do not need to be necessarily matched, if

$$\neg \mathbf{GT} \wedge \mathbf{GT}^{soft} = true . \quad (6)$$

Note that we use \neg to denote the logical negation operator. These are usually ignored during evaluation. Image pairs must not be matched if

$$\mathbf{GT}^{soft} = false . \quad (7)$$

How \mathbf{GT} and \mathbf{GT}^{soft} are actually used for evaluation is presented in the following.

C. Metrics

This section presents established metrics to evaluate a VPR method, including precision and recall, the precision-recall curve, area under the precision-recall curve, maximum recall at 100% precision, and recall@ K [32]. All metrics are implemented in the associated code repository (see Section I). These metrics are based on

- the pairwise descriptor similarities $s_{ij} \in \mathbf{S}$ (cf. Section IV-D) or the image matches $m_{ij} \in \mathbf{M}$ (cf. Section IV-E)

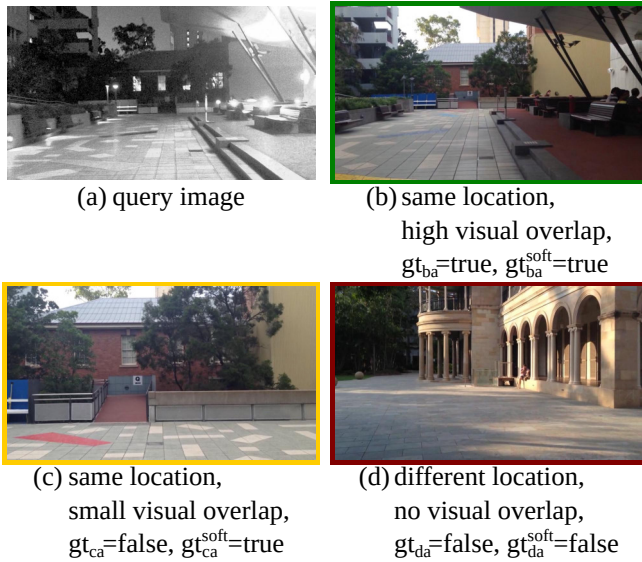


Fig. 4: Relation between the visual overlap of a query image (a) and reference images (b-d), and their corresponding ground truth values gt_{ij} and gt_{ij}^{soft} . Since (c) and (a) only have a small visual overlap, we do not expect a VPR method to match both images and set $gt_{ca} = false$. However, we also avoid penalization in case the VPR method indeed matches both images by setting $gt_{ca}^{soft} = true$.

- with corresponding ground truth $gt_{ij} \in \mathbf{GT}$ (and $gt_{ij}^{soft} \in \mathbf{GT}^{soft}$ in case the soft ground truth is used).

For single-best-match VPR, the evaluation only considers the best matching image pair per query with the highest similarity s_{i^*j} :

$$i^* = \arg \max_i s_{ij} . \quad (8)$$

Precision and Recall: Precision P and recall R are important metrics in the information retrieval domain. In the context of VPR, **precision** P represents the ratio of correctly matched images of same places to the total number of matched images with

$$P = \frac{\#TP}{\#TP + \#FP} . \quad (9)$$

Recall R expresses the ratio of correctly matched images of same places to the total number of ground-truth positives (GTP):

$$R = \frac{\#TP}{\#GTP} . \quad (10)$$

In the case of single-best-match VPR, the number of ground-truth positives refers to the total number of query images for which a ground-truth match exists, i.e.,

$$\#GTP = \sum_{\forall j} \begin{cases} 1, & \text{if } \exists i : gt_{ij} \\ 0, & \text{otherwise} \end{cases} , \quad (11)$$

whereas in the case of multi-match VPR, the number of ground-truth positives refers to the number of actually match-

ing image pairs, i.e.,

$$\#GTP = \sum_{\forall i,j} \begin{cases} 1, & \text{if } gt_{ij} \\ 0, & \text{otherwise} \end{cases} . \quad (12)$$

$\#TP$ and $\#FP$ are the number of correctly matched and wrongly matched image pairs. More specifically, **true positives** TP are actual matching image pairs that were classified as matches:

$$\#TP = \sum_{\forall i,j} \begin{cases} 1, & \text{if } gt_{ij} \wedge m_{ij} \\ 0, & \text{otherwise} \end{cases} . \quad (13)$$

For single-best-match VPR, only i^* from Eq. (8) is evaluated in Eq. (13) for each query image. The same is true for the following Eq. (14).

False positives FP are non-matching image pairs that were incorrectly classified as matches:

$$\#FP = \sum_{\forall i,j} \begin{cases} 1, & \text{if } \neg gt_{ij} \wedge m_{ij} \\ 0, & \text{otherwise} \end{cases} . \quad (14)$$

Note that when using the soft ground truth, image pairs with $\neg gt_{ij} \wedge gt_{ij}^{soft} = true$ (cf. Eq. (6)) are ignored during the computation of $\#TP$ and $\#FP$. While **false negatives** FN are indirectly involved in the calculation of recall R , **true negatives** TN are usually not evaluated due to the typically imbalanced classification problem of VPR with $\#TN \gg \#TP, \#FP, \#FN$.

Precision-recall curve: Precision-recall curves can be used to avoid actual matching decisions, which are often made after VPR using a computationally expensive verification algorithm. The idea is to make matching decisions with $\mathbf{M} = \mathbf{S} \geq \theta_k$ over a range of thresholds $\theta = \{\min(\mathbf{S}), \dots, \max(\mathbf{S})\}$. For instance, the number of true positives $\#TP_k$ for one specific $\theta_k \in \theta$ is then computed with

$$\#TP_k = \sum_{\forall i,j} \begin{cases} 1, & \text{if } gt_{ij} \wedge (s_{ij} \geq \theta_k) \\ 0, & \text{otherwise} \end{cases} . \quad (15)$$

Following Eqs. (9) and (10), this leads to two vectors of precision and recall values $P(\theta_k)$ and $R(\theta_k)$, which in combination formulate the **precision-recall curve**. The full pipeline for the computation of the precision-recall curve is depicted in Figure 5 and code is provided in Code Snippet 5.

Area under the precision-recall curve (AUPRC): The **area under the precision-recall curve** (also termed average precision, avgP) can be used to compress a precision-recall curve into a single number, as shown in Code Snippet 6. In Figure 5, the AUPRC is visualized as the green area under the precision-recall curve.

Maximum recall at 100% precision: The maximum recall at 100% precision (short $R@100P$) represents the maximum recall where $P = 1$ (100%), i.e., the maximum recall without false positives FP (cf. Eq. (14)). In the past, this metric was important to evaluate VPR methods for loop closure detection in SLAM. Keeping the precision at $P = 1$ avoids wrong loop closures and, consequently, mapping errors. However, since the advent of robust graph optimization techniques for


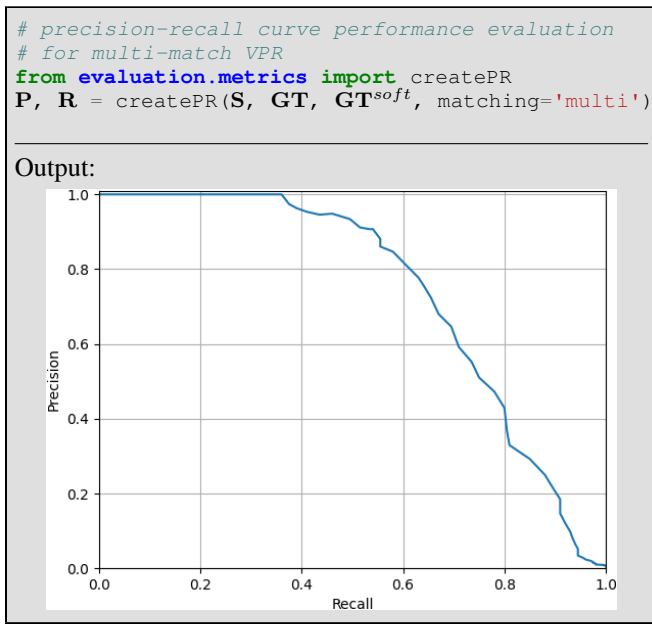



Fig. 5: The evaluation pipeline for multi-match VPR, including the precision-recall curve and the area under the precision-recall curve (AUPRC). Given the similarity matrix S and ground truth GT and GT^{soft} (cf. Section V-B), a range of thresholds $\theta_k \in \{\min(S), \dots, \max(S)\}$ is applied with $S \geq \theta_k$ to obtain binary matching decisions $m_{ij} \in M_k$ for each θ_k . In combination with the ground truth, these can be labeled as either True Positives (TP), False Positives (FP), False Negatives (FN) or True Negatives (TN), and converted into a precision-recall curve and the area under the precision-recall curve (AUPRC).



Code Snippet 5: To evaluate the quality of a similarity matrix S , we can apply a series of decreasing thresholds θ to match more and more image pairs. Combined with ground-truth information, each threshold leads to a different set of true positives, false positives, true negatives and false negatives, which then provides one point on the precision-recall curve. In this example, we create the precision-recall curve for multi-match VPR.

SLAM [33], the avoidance of wrong loop closures became less relevant. With robust graph optimization, it is more important to find enough correct loop closures (TP) than to avoid wrong loop closures (FP). Therefore, using multi-match VPR to identify all loop closures should be preferred over tuning the $R@100P$ for such applications.

If the precision never reaches $P = 1$, the maximum recall at 100% precision is undefined. Therefore, the maximum recall at 99% or 95% precision has been used alternatively.



Code Snippet 6: Finally, to summarize the place recognition quality in a single number, we can use the area under the precision-recall curve (AUPRC).

Recall@K: The $\text{recall@}K$ (also termed *top-K success rate*) is an often used metric for the evaluation of image classifiers. For place recognition, it is defined as follows: For each query image, given the K database images with the K highest similarities s_{ij} , the $\text{recall@}K$ measures the rate of query images with at least one actually matching database image. That means this metric requires at least one matching image in the database for each query image, which corresponds to a typical localization scenario without exploration. For mapping with newly visited places, the metric is not defined. In such a scenario, an implementation of $\text{recall@}K$ could simply ignore all query images without a matching database image – however, this workaround would not evaluate the (in)ability of a method to handle exploration during the query run, i.e., new places which are not part of the database set.

The $\text{recall@}K$ is particularly suited for visual localization tasks, where the K most similar database or query images are retrieved for a subsequent geometric verification. Note that for VPR in the context of localization without exploration (i.e., all query images have at least one matching reference image), the $\text{recall@}1$ and the precision at 100% recall are identical.

Mean, best case and worst case performance: To get a comprehensive understanding of how well a VPR method performs in different environments, with different types of appearance and viewpoint changes, it is best practice to

evaluate it using multiple datasets. The aforementioned metrics measure the performance on each single dataset. One can get a more condensed view of the overall performance by considering the mean, best case and worst case performance. The mean performance allows for a quick comparison with other evaluated methods. The best case performance shows the maximum achievable performance and reveals potential strengths of an approach: if the best case performance of a method is higher than that of the compared methods, this method is well suited for the conditions under which the best case performance was achieved. The worst case performance reveals the weaknesses of a method and its sensitivity to certain conditions or trajectories (cf. Figure 3). For example, if the worst case performance of a method is lower than the worst case performance of the compared methods, it indicates that this method is less robust and struggles with the specific property of at least one of the evaluated datasets.

We would like to note that there are various other metrics for evaluating VPR methods, including those that take computational time into account. We refer interested readers to [32] for a comprehensive overview, which also includes examples of performance evaluations for the same algorithm across multiple metrics.

VI. CHALLENGES AND COMMON WAYS OF ADDRESSING THEM

The previous sections introduced a generic pipeline for VPR, and how to evaluate such a pipeline. In this section, we go beyond this basic pipeline and enlist typical challenges that researchers face in the field of VPR and the ways that prior work has addressed them.

A. Scalability

A major challenge in VPR is how to scale up the system to handle large numbers of images in the database or query set. As discussed in Section IV-B, holistic image descriptors allow for fast retrieval. To reduce the computational effort for descriptor comparison, dimensionality reduction techniques like random projections, or Principal Component Analysis (PCA) have been proposed.

However, the computation time for recognizing places is typically still proportional to the number of images in the database DB or query set Q . To further improve efficiency, approximate nearest neighbor search (e.g., a combination of KD-tree and product quantization as with DELF [34]) can be employed instead of a linear search of all database descriptors, which leads to a sublinear time complexity. Additionally, incorporating coarse position data from weak GPS signals can increase efficiency as it reduces the search space [35].

Finally, to compensate for the reduced accuracy of holistic descriptors, hierarchical place recognition can be employed. This approach re-ranks the top- K retrieved matches from holistic descriptors through geometric verification with local image descriptors [7].

B. Appearance variations

When a robot revisits a place, its current image observation often experiences significant variations in appearance (as

discussed in Section V-A), which can negatively affect the performance. To reduce the discrepancy between the query observation and the observation stored in the database, techniques such as illumination invariant images [36], shadow removal [37], appearance change prediction [38], linear regression [39], and deep learning based methods using generative adversarial networks [40] can be used to convert all images into a reference condition [19]. Such techniques require that the correspondence between each image and its actual condition is provided by human supervision or a condition classifier (e.g., database: summer, query: winter).

To avoid such condition-specific approaches that are trained or designed only for specific conditions (e.g., Night-To-Day-GAN: night and day, shadow removal: different times of day), a condition-wise descriptor standardization can be used to significantly improve performance over a wide range of conditions [19]: This standardization normalizes each dimension of the descriptors from one condition to zero mean and unit standard deviation (e.g., once for the database in summer, once for the query set in winter). Furthermore, if appearance variations occur not only *across* the query and database traverses (e.g., database: sunny, query: rainy) but also *within* a traverse (e.g., database: sunny→cloudy→overcast→rainy, query: sunny), descriptors can be clustered and then standardized per cluster. Besides addressing individual appearance challenges as above, a common trend in recent research has been to train deep architectures on large-scale, diverse datasets [30] to achieve global [41] and local [42] descriptors that are robust to appearance variations. Alternatively, one can combine the strengths of multiple descriptors by simply concatenating them (which sums up their dimensionalities) or combining them using techniques such as hyperdimensional computing (which limits the dimensionality) [28].

C. Viewpoint variations

A robot may revisit a place from a different viewpoint. For drones, this change could be due to a varying 6-DoF pose, and for an on-road vehicle, it could be due to changes in lanes and direction of travel. In addition to recognizing a local feature or region from different viewpoints, one also needs to deal with often limited visual overlap between an image pair captured from different viewpoints. The problem of viewpoint variations becomes even more challenging when simultaneously affected by appearance variations that widen the scope for perceptual aliasing (the problem of distinct places looking very similar, as discussed in Section I and detailed in, e.g., [2]). A popular solution to deal with viewpoint variations is to learn holistic descriptors by aggregating local features in a permutation-invariant manner, that is, independent of their pixel locations, as in NetVLAD [29].

D. Improving performance

In addition to the approaches mentioned earlier, there are several ways to improve VPR performance by using task-specific knowledge.

Sequence-based methods leverage sequences in the database and query set, which lead to continuous lines of

high similarities in the similarity matrix \mathbf{S} (cf. Section IV-D and Figure 3). We can divide these methods into two categories: **Similarity-based sequence methods** use the similarities $s_{ij} \in \mathbf{S}$ to find linear segments of high similarities, e.g., SeqSLAM [13], or continuous lines of high similarities with potentially varying slope, e.g., based on a flow network. One can also use available odometry information to find sequences with varying slope. On the other hand, a sequence of holistic image descriptors can be combined into a single vector. A **sequence descriptor** defined for a place thus accounts for the visual information observed in the preceding image frames [43], [44]. These sequence descriptors can be compared between the database and the query sets to obtain place match hypotheses.

Besides leveraging descriptor similarities \mathbf{S} between the database and query sets, **intra-database and intra-query similarities** \mathbf{S}^{DB} and \mathbf{S}^Q , i.e., descriptor similarities *within* the database and query sets, can be used to improve performance. For example, in [16], the intra-set similarities \mathbf{S}^{DB} and \mathbf{S}^Q are used in combination with \mathbf{S} and sequence information to formulate a factor graph that can be optimized to refine the similarities in \mathbf{S} . In this graph, the intra-set similarities are used to connect images within the database or query sets that are likely to show the same or different places due to a high or low intra-set similarity. For example, let us suppose that the l -th query image has high similarities s_{il} and s_{jl} to the i -th and j -th database images. Let us further suppose that the similarity s_{kl} to the k -th database image is low, although the i -th, j -th and k -th database images have high intra-database similarities s_{ij}^{DB} , s_{ik}^{DB} and s_{jk}^{DB} . The graph optimization then detects that the similarity s_{kl} between the k -th database image and the l -th query image is also likely to be high.

Methods such as experience maps [45] and co-occurrence maps [46] can be used in cases where the robot **frequently revisits the same places**. These “memory-based” methods continually observe each place and create a descriptor every time the appearance changes. During a comparison of a new query descriptor with this aggregated “map”, only one descriptor of a similar condition needs to be matched to recognize the place, reducing the need for condition-invariant descriptors. Several approaches go beyond these memory-based techniques by modeling spatiotemporal dynamics to forecast feature persistence [47], expected outdoor conditions [48], or map occupancy [49].

In the case of robot localization with known places (i.e., each visited query place is guaranteed to be in the database and no exploration beyond this mapped area is performed), VPR can benefit from **place-specific classifiers**, which can improve accuracy with reduced map storage or retrieval time [50]. A similar approach is to train a deep learning-based **place classifier** that directly outputs a place label for a given image, or to create environment-specific descriptors. Another direction is to exploit known place types for place type matching to limit the number of potential matches between the database and query set [15]. For example, instead of searching through all database images, if the query image was taken in a forest, such semantic categorization constrains

the database images to only those that were also taken in a forest.

VII. CONCLUSIONS

Visual Place Recognition (VPR) is a well established problem that has found widespread interest and use in both computer vision and robotics. In this tutorial, we have described the visual place recognition task, including its various problem categories and subtypes, their typical use cases, and how it is typically implemented and evaluated. Additionally, we discussed a number of methods that can be used to address common challenges in VPR.

There are a number of open challenges such as system integration, enriched reference maps, view synthesis, and the design of a “one-fits-all” solution that still need to be tackled by the community. While we do not discuss these challenges in this tutorial, we refer the interested reader to [1], [2].

REFERENCES

- [1] S. Schubert and P. Neubert, “What makes visual place recognition easy or hard?” *arXiv 2106.12671*, 2021.
- [2] S. Garg, T. Fischer, and M. Milford, “Where is your place, visual place recognition?” in *International Joint Conference on Artificial Intelligence (IJCAI)*, 2021, pp. 4416–4425.
- [3] C. Masone and B. Caputo, “A survey on deep visual place recognition,” *IEEE Access*, vol. 9, pp. 19 516–19 547, 2021.
- [4] S. Lowry, N. Sünderhauf, P. Newman, J. J. Leonard, D. Cox, P. Corke, and M. J. Milford, “Visual place recognition: A survey,” *IEEE Transactions on Robotics (TRO)*, vol. 32, pp. 1–19, 2016.
- [5] W. Maddern, G. Pascoe, C. Linegar, and P. Newman, “1 year, 1000 km: The Oxford RobotCar dataset,” *The International Journal of Robotics Research (IJRR)*, vol. 36, no. 1, pp. 3–15, 2017.
- [6] K. A. Tsintotas, L. Bampis, and A. Gasteratos, “The revisiting problem in simultaneous localization and mapping: A survey on visual loop closure detection,” *IEEE Transactions on Intelligent Transportation Systems (T-ITS)*, vol. 23, no. 11, pp. 19 929–19 953, 2022.
- [7] M. Cummins and P. Newman, “FAB-MAP: Probabilistic localization and mapping in the space of appearance,” *The International Journal of Robotics Research (IJRR)*, vol. 27, no. 6, pp. 647–665, 2008.
- [8] R. Mur-Artal, J. M. M. Montiel, and J. D. Tardós, “ORB-SLAM: A Versatile and Accurate Monocular SLAM System,” *IEEE Transactions on Robotics (TRO)*, vol. 31, no. 5, pp. 1147–1163, 2015.
- [9] P. Yin, S. Zhao, I. Cisneros, A. Abuduweili, G. Huang, M. Milford, C. Liu, H. Choset, and S. Scherer, “General place recognition survey: Towards the real-world autonomy age,” *arXiv 2209.04497*, 2022.
- [10] H. Durrant-Whyte and T. Bailey, “Simultaneous localization and mapping: Part I,” *IEEE Robotics & Automation Magazine (RA-M)*, vol. 13, no. 2, pp. 99–110, 2006.
- [11] 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.
- [12] F. Dayoub and T. Duckett, “An adaptive appearance-based map for long-term topological localization of mobile robots,” in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2008, pp. 3364–3369.
- [13] M. J. Milford and G. F. Wyeth, “SeqSLAM: Visual route-based navigation for sunny summer days and stormy winter nights,” in *IEEE International Conference on Robotics and Automation (ICRA)*, 2012, pp. 1643–1649.
- [14] Z. Chen, O. Lam, A. Jacobson, and M. Milford, “Convolutional neural network-based place recognition,” *Australasian Conference on Robotics and Automation (ACRA)*, 2014.
- [15] N. Sünderhauf, S. Shirazi, F. Dayoub, B. Ucpocroft, and M. Milford, “On the performance of convnet features for place recognition,” in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2015, pp. 4297–4304.
- [16] S. Schubert, P. Neubert, and P. Protzel, “Fast and memory efficient graph optimization via ICM for visual place recognition,” in *Robotics: Science and Systems (RSS)*, 2021.

- [17] D. Scaramuzza and F. Fraundorfer, "Visual odometry [tutorial]," *IEEE Robotics & Automation Magazine (RA-M)*, vol. 18, no. 4, pp. 80–92, 2011.
- [18] O. Vysotska and C. Stachniss, "Effective visual place recognition using multi-sequence maps," *IEEE Robotics and Automation Letters (RA-L)*, vol. 4, no. 2, pp. 1730–1736, 2019.
- [19] S. Schubert, P. Neubert, and P. Protzel, "Unsupervised learning methods for visual place recognition in discretely and continuously changing environments," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2020, pp. 4372–4378.
- [20] K. Ho and P. Newman, "Detecting loop closure with scene sequences," *International Journal of Computer Vision (IJCV)*, vol. 74, no. 3, pp. 261–286, 2007.
- [21] T. Sattler, W. Maddern, C. Toft, A. Torii, L. Hammarstrand, E. Stenborg, D. Safari, M. Okutomi, M. Pollefeys, J. Sivic, F. Kahl, and T. Pajdla, "Benchmarking 6DOF outdoor visual localization in changing conditions," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018, pp. 8601–8610.
- [22] J. McDonald, M. Kaess, C. Cadena, J. Neira, and J. J. Leonard, "6-DoF multi-session visual SLAM using anchor nodes," in *European Conference on Mobile Robots (ECMR)*, 2011, pp. 69–76.
- [23] T. Cieslewski, S. Choudhary, and D. Scaramuzza, "Data-efficient decentralized visual SLAM," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2018, pp. 2466–2473.
- [24] A. Glover, "Day and night, left and right," 2014.
- [25] T. Weyand, I. Kostrikov, and J. Philbin, "PlaNet - photo geolocation with convolutional neural networks," in *European Conference on Computer Vision (ECCV)*, 2016, p. 37–55.
- [26] X. Zhang, L. Wang, and Y. Su, "Visual place recognition: A survey from deep learning perspective," *Pattern Recognition*, vol. 113, 2021.
- [27] E. Garcia-Fidalgo and A. Ortiz, "iBoW-LCD: An appearance-based loop-closure detection approach using incremental bags of binary words," *IEEE Robotics and Automation Letters*, vol. 3, no. 4, pp. 3051–3057, 2018.
- [28] P. Neubert and S. Schubert, "Hyperdimensional computing as a framework for systematic aggregation of image descriptors," in *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021, pp. 16 938–16 947.
- [29] R. Arandjelovic, P. Gronat, A. Torii, T. Pajdla, and J. Sivic, "NetVLAD: CNN architecture for weakly supervised place recognition," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016, pp. 15 297–15 307.
- [30] F. Warburg, S. Hauberg, M. Lopez-Antequera, P. Gargallo, Y. Kuang, and J. Civera, "Mapillary street-level sequences: A dataset for lifelong place recognition," in *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020, pp. 2626–2635.
- [31] N. Sünderhauf, P. Neubert, and P. Protzel, "Are we there yet? Challenging SeqSLAM on a 3000 km journey across all four seasons," in *IEEE International Conference on Robotics and Automation (ICRA) Workshop on Long-Term Autonomy*, 2013.
- [32] M. Zaffar, S. Garg, M. Milford, J. Kooij, D. Flynn, K. McDonald-Maier, and S. Ehsan, "VPR-Bench: An open-source visual place recognition evaluation framework with quantifiable viewpoint and appearance change," *International Journal of Computer Vision (IJCV)*, pp. 1–39, 2021.
- [33] N. Sünderhauf and P. Protzel, "Switchable constraints for robust pose graph SLAM," in *IEEE International Conference on Intelligent Robots and Systems (IROS)*, 2012, pp. 1879–1884.
- [34] H. Noh, A. Araujo, J. Sim, T. Weyand, and B. Han, "Large-scale image retrieval with attentive deep local features," in *IEEE International Conference on Computer Vision (ICCV)*, 2017, pp. 3456–3465.
- [35] 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.
- [36] W. Maddern, A. Stewart, C. McManus, B. Upcroft, W. Churchill, and P. Newman, "Illumination invariant imaging: Applications in robust vision-based localisation, mapping and classification for autonomous vehicles," in *International Conference on Robotics and Automation (ICRA) Workshop on Visual Place Recognition in Changing Environments*, 2014.
- [37] C. McManus, W. Churchill, W. Maddern, A. D. Stewart, and P. Newman, "Shady dealings: Robust, long-term visual localisation using illumination invariance," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2014, pp. 901–906.
- [38] P. Neubert, N. Sünderhauf, and P. Protzel, "Superpixel-based appearance change prediction for long-term navigation across seasons," *Robotics and Autonomous Systems*, vol. 69, pp. 15–27, 2015.
- [39] S. Lowry and M. J. Milford, "Supervised and unsupervised linear learning techniques for visual place recognition in changing environments," *IEEE Transactions on Robotics (TRO)*, vol. 32, no. 3, pp. 600–613, 2016.
- [40] H. Porav, W. Maddern, and P. Newman, "Adversarial training for adverse conditions: Robust metric localisation using appearance transfer," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2018, pp. 1011–1018.
- [41] G. Berton, C. Masone, and B. Caputo, "Rethinking visual geolocalization for large-scale applications," in *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022, pp. 4878–4888.
- [42] B. Cao, A. Araujo, and J. Sim, "Unifying deep local and global features for image search," in *European Conference on Computer Vision (ECCV)*, 2020, pp. 726–743.
- [43] S. Garg and M. Milford, "SeqNet: Learning Descriptors for Sequence-Based Hierarchical Place Recognition," *IEEE Robotics and Automation Letters (RA-L)*, vol. 6, no. 3, pp. 4305–4312, 2021.
- [44] R. Mereu, G. Trivigno, G. Berton, C. Masone, and B. Caputo, "Learning sequential descriptors for sequence-based visual place recognition," *IEEE Robotics and Automation Letters (RA-L)*, vol. 7, no. 4, pp. 10 383–10 390, 2022.
- [45] W. Churchill and P. Newman, "Experience-based navigation for long-term localisation," *The International Journal of Robotics Research (IJRR)*, vol. 32, no. 14, pp. 1645–1661, 2013.
- [46] E. Johns and G. Yang, "Feature co-occurrence maps: Appearance-based localisation throughout the day," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2013, pp. 3212–3218.
- [47] D. M. Rosen, J. Mason, and J. J. Leonard, "Towards lifelong feature-based mapping in semi-static environments," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2016, pp. 1063–1070.
- [48] C. Linegar, W. Churchill, and P. Newman, "Work smart, not hard: Recalling relevant experiences for vast-scale but time-constrained localisation," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2015, pp. 90–97.
- [49] T. Krajník, J. P. Fentanes, J. M. Santos, and T. Duckett, "FreMEn: Frequency map enhancement for long-term mobile robot autonomy in changing environments," *IEEE Transactions on Robotics (TRO)*, vol. 33, no. 4, pp. 964–977, 2017.
- [50] C. McManus, B. Upcroft, and P. Newman, "Learning place-dependant features for long-term vision-based localisation," *Autonomous Robots*, vol. 39, pp. 363–387, 2015.