

Cross-Descriptor Visual Localization and Mapping

Mihai Dusmanu¹ Ondrej Miksik² Johannes L. Schönberger² Marc Pollefeys^{1,2}
¹ Department of Computer Science, ETH Zürich ² Microsoft MR & AI Lab, Zürich

Abstract

Visual localization and mapping is the key technology underlying the majority of mixed reality and robotics systems. Most state-of-the-art approaches rely on local features to establish correspondences between images. In this paper, we present three novel scenarios for localization and mapping which require the continuous update of feature representations and the ability to match across different feature types. While localization and mapping is a fundamental computer vision problem, the traditional setup supposes the same local features are used throughout the evolution of a map. Thus, whenever the underlying features are changed, the whole process is repeated from scratch. However, this is typically impossible in practice, because raw images are often not stored and re-building the maps could lead to loss of the attached digital content. To overcome the limitations of current approaches, we present the first principled solution to cross-descriptor localization and mapping. Our data-driven approach is agnostic to the feature descriptor type, has low computational requirements, and scales linearly with the number of description algorithms. Extensive experiments demonstrate the effectiveness of our approach on state-of-the-art benchmarks for a variety of handcrafted and learned features.

1. Introduction

Mixed reality and robotics blend the physical and digital worlds to unlock the next evolution in human, machine, and environment interaction. This promise largely relies on the technical capability to build and localize against maps of the environment. For example, collaborative experiences require co-localization of multiple devices into the same coordinate space. Similarly, re-localization against existing maps enables persistence and retrieval of digital content attached to the real world. Most systems build upon localization and mapping pipelines predominantly based on vision sensors using local features for image matching.

Traditionally, these systems run in real-time on the device, but the industry has been increasingly moving localization and mapping capabilities to the cloud (e.g., Facebook LiveMaps [2], Google VPS [42], Magic Leap’s Mag-




Figure 1: Mixed reality and robotics systems typically do not store raw images. This complicates deployment of new feature representations for visual localization and mapping. Our approach enables the continuous update of feature representations and the ability to match across heterogeneous devices using different features.

icverse [3], or Microsoft Azure Spatial Anchors [23]). This is for a variety of reasons, such as reducing on-board compute or enabling collaborative experiences and crowd-sourced mapping. In most settings, images are typically not shared across devices or uploaded to the cloud primarily due to privacy reasons [25], but also to reduce bandwidth and storage requirements. Instead, image features are computed locally on the device and only the extracted keypoints and descriptors are shared. In other words, it is impossible to re-extract features whenever a different representation is required, as the images are no longer available.

There are two fundamental limitations of existing systems. First, they cannot adopt new feature algorithms because incompatibilities of feature representations make re-localization against maps built using older features impossible. In a world of continuous progress on local features in the research community, this severely limits progress on localization and mapping. One might argue for sim-

ply re-building the maps from scratch whenever a significantly improved feature algorithm is available. However, content attached to the old maps would be lost and mapping is an expensive process, where it could take weeks or even months until the whole area is re-visited. Second, co-localization and collaborative mapping scenarios with devices using different features is impossible. The situation is further complicated by the fact that many devices implement specific algorithms in hardware for efficiency reasons, making a client-side upgrade of algorithms impossible. This also means that many of the existing commercial solutions might be significantly behind the state-of-the-art in the research community, because they cannot easily upgrade their algorithms and representations.

In this paper, we first define three novel scenarios addressing the challenges in a world of changing local feature representations (*c.f.* Figures 1 and 2):

- *Continuous deployment* of feature representations without requiring an explicit re-mapping phase.
- *Cross-device localization* when the localization and mapping devices use different features.
- *Collaborative mapping* with multiple heterogeneous devices and features.

Note that these scenarios are completely different from existing industrial or academic setups, where generally a *fixed* local feature extraction algorithm is used for co-localization and throughout the evolution of a map. Section 3 introduces the scenarios in more detail.

As a first step towards enabling these scenarios, we focus on local feature descriptors and propose what to our knowledge is the first principled and scalable approach to the underlying challenges. Our learned approach translates descriptors from one representation to another without any assumptions about the structure of the feature vectors and enables matching of features with incompatible dimensionality as well as distance metrics. For instance, we can match 512 dimensional binary BRIEF [10] against 128 dimensional floating point SIFT [28] or even deep learning models such as HardNet [33] or SOSNet [51] and vice versa. Our method has linear scalability in the number of algorithms and is specifically designed to have a small computational footprint to enable deployment on low-compute devices. The training data is generated automatically by computing different descriptors from the same image patches.

We evaluate our approach on relevant geometric tasks in the context of the newly proposed scenarios. We first consider localization (pose estimation) on the Aachen Day & Night benchmark [46]. Next, we assess the performance for 3D mapping from crowd-sourced images using the benchmark of Schönberger *et al.* [48]. In the supplementary material, we also show results on the HPatches descriptor benchmark [7], the Image Matching Workshop challenge [22], and the InLoc Indoor Visual Localization dataset [50]. Our

experiments demonstrate the effectiveness and high practical relevance of our method for real-world localization and mapping systems.

To summarize our contributions, we i) introduce three novel scenarios to localization and mapping in a world of changing feature representations, ii) propose the first principled and scalable approach tackling these newly introduced scenarios, and iii) demonstrate the effectiveness of our method on challenging real-world datasets.

2. Related work

Local image descriptors. Local descriptors are typically extracted from normalized images patches defined by local feature frames [28, 31, 44]. The most notable handcrafted descriptors include binary BRIEF [10], gradient-based SIFT [28], and their variants [44, 32, 52, 6]. Recently, the community has moved to data-driven models leveraging large-scale datasets, various triplet losses, and hard-negative mining techniques [8, 33, 29, 51, 16]. A related line of research aims to reformulate the entire pipeline in an end-to-end trainable fashion [53, 39]. Finally, the describe-then-detect approach extracts local features from dense feature maps [38, 13, 15, 43, 30]. While it is possible to handcraft relations between some of these descriptors (*e.g.* SIFT [28] \leftrightarrow Hue-SIFT [52]), they are typically not designed for inter-compatibility. In contrast, our approach aims to enable matching of local descriptors with incompatible feature representations, dimensionalities or metrics.

Visual localization and mapping. Given the efficiency of local features and the well-established theory of (sparse) multi-view geometry, a majority of large-scale mapping [4, 19] and localization [27, 45] is based on local image features. Furthermore, there is a growing number of open-source tools for Simultaneous Localization and Mapping (*e.g.*, ORB-SLAM [36, 37]), visual odometry (*e.g.*, LIBVISO [26, 18]) or Structure-from-Motion (*e.g.*, COLMAP [47], OpenMVG [35], Theia [49]) rendering this research domain more accessible. Recently, these pipelines were adopted for evaluation of visual localization and mapping – Schönberger *et al.* [48] and Jin *et al.* [22] considered mapping and localization performance using phototourism collections, while Sattler *et al.* [46] introduced several datasets for localization under challenging scenarios (*e.g.*, seasonal / temporal changes). Also related to our work, the problem of collaborative mapping between multiple agents (*e.g.*, robots, drones) has been thoroughly studied [17, 11, 34] in the field of robotics. However, all these systems and associated benchmarks use the same local image features throughout the whole process and are thus unable to tackle any of the scenarios discussed in Section 3.

Domain translation. Several local feature inversion techniques [14, 40] reconstructing the original image from its




Figure 2: **Continuous deployment and cross-device localization.** (a) Translating the local descriptors associated with the map enables continuous deployment of new feature representations. (b) Translating the query descriptors enables cross-device localization when queries and map use different feature representations.

features were proposed. One way to enable descriptor translation would be to reconstruct the image and then follow the traditional pipeline to extract a different descriptor. In our early experiments with feature inversion networks, we noticed significant limitations in generalizing to new scenes and poor quality of recovered low-level gradient information. As such, it was not possible to successfully match images using this feature inversion based approach.

Image-to-image translation [54, 21, 55] attempts to convert one representation of a scene to another (e.g., RGB image to semantic labels). One could imagine adapting these models to translate between different local feature representations. However, similarly to feature inversion techniques, these approaches typically have very large computational and memory footprint, as they employ deep convolutional neural networks to produce full resolution images. In contrast, we propose to use shallow multi-layer perceptrons, taking a single descriptor as input and predicting its translation – a solution more suitable for low-compute devices.

Finally, research on domain adaptation (c.f. surveys of Csurka [12] and Zhuang *et al.* [56]) tackles the problem of adapting algorithms learned on a source data distribution to a dataset with a related but different distribution. In the case of deep learning, this is particularly interesting when the source domain has enough annotations to allow training while the target domain has little to no annotations available. In our case, instead of adapting to a different input data distribution, we try to adapt the output distributions to make them compatible for matching.

3. Scenarios

As our first contribution, we identify, introduce, and formalize three novel scenarios for localization and mapping. In all scenarios, we assume there are no images stored in the resulting maps. Therefore, it is impossible to re-extract or replace the underlying feature representation.

3.1. Continuous deployment

Given a sparse 3D map, associated local feature descriptors and no access to the original image data, the goal is to develop a mechanism enabling the continuous deployment of new feature representations (Figure 2 (a)) without the need for re-mapping. In other words, we aim to *translate* the features of the map from one representation to another. For instance, it should be possible to switch from handcrafted to learned descriptors (e.g., SIFT [28] → SOSNet [51]), change their dimensionality, or update the model weights (e.g., different training data, loss, or architecture).

Matching the new feature representation against a database of translated descriptors might lead to worse performance when compared to matching using a single feature representation. This is, however, only a transient issue and not of serious concern in practice, since the translated descriptors can be used in a boot-strapping fashion – as users re-visit the environment, the translated descriptors are gradually replaced by the newly extracted ones. Thus, the map will eventually contain the latest descriptors to close the performance gap. The key points are that i) the devices run only a single algorithm (due to real-time constraints) or cannot be easily updated (due to hardware implementation), ii) the scenario avoids an explicit re-mapping phase, which is generally very expensive and time-consuming.

3.2. Cross-device localization

The previous scenario assumes the same feature extraction algorithm is running on all client devices. However, this is generally not the case, as legacy devices often cannot be updated due to hardware limitations. In this case, a mechanism for backwards-compatibility is required to match across different device versions (see Figure 2 (b)). Similarly, specialized devices (e.g., headsets, autonomous cars, mapping platforms) could take advantage of additional on-board compute for a better localization perfor-

mance compared to devices such as mobile phones or light-weight drones. It would be beneficial to co-localize these devices inside the same maps to avoid a fragmented scene representation. The same functionality could be extended to enable localization between devices of different vendors, but, in this situation, new challenges arise due to differences in previous steps of the pipeline – notably feature detection, as discussed in detail in Section 6.

3.3. Collaborative mapping

Finally, we consider multiple heterogeneous devices *collaboratively* mapping the same environment. For the aforementioned reasons, it is fairly unlikely for all devices to use the same features as in a standard academic setup, even when produced by a single vendor. Consequently, we end up with each device providing different types of features from which we need to build a single coherent map. Therefore, we need a mechanism for translating all descriptors into a common representation that can be used for establishing correspondences (see Figure 1).

4. Method

In this section, we present the first principled solution to cross-descriptor localization and mapping. We start by formalizing the descriptor translation problem. We then propose to use a separate multi-layer perceptron (MLP) for each pair of description algorithms, trained using a translation loss (c.f. Figure 3 (a)). Finally, we extend this formulation by using an auto-encoder inspired model that embeds all descriptors from different algorithms into a joint embedding space (c.f. Figure 3 (b)). To make sure that the common embedding is suitable for establishing correspondences between different algorithms, we leverage additional supervision through a matching loss.

4.1. Descriptor translation

A feature description algorithm A defines a handcrafted or learned function mapping images to vectors as $A : \mathcal{I}_{h \times w} \rightarrow \mathbb{R}^n$, where $\mathcal{I}_{h \times w}$ is the set of images of size $h \times w$ and n is the embedding dimension. In the case of local image descriptors, the domain is the set of patches \mathcal{I} normalized based on the estimated feature geometry (position, scale, orientation, affine shape).

Let $A_1 : \mathcal{I} \rightarrow \mathbb{R}^{n_1}$ and $A_2 : \mathcal{I} \rightarrow \mathbb{R}^{n_2}$ be two description algorithms. A translation function $t_{1 \rightarrow 2} : \mathbb{R}^{n_1} \rightarrow \mathbb{R}^{n_2}$ maps a feature vector of A_1 onto the manifold of A_2 , satisfying $t_{1 \rightarrow 2}(A_1(p)) = A_2(p)$ for all patches $p \in \mathcal{I}$. Our experiments (Section 5) empirically demonstrate that using MLPs to approximate the translation functions $t_{1 \rightarrow 2}$ is highly effective in enabling matching between different descriptors which, in turn, facilitates the scenarios introduced in Section 3. Note that we do not propose a matching algorithm but instead directly map descriptors from one space

to another. Our approach thus relies on traditional descriptor comparison strategies, lending itself to efficient matching implementations (e.g., approximate nearest neighbor). Apart from the corresponding descriptors being extracted from the same patches, we make no additional assumptions regarding the descriptor dimension or the underlying algorithm (handcrafted or learned, binary or floating point).

In the following, we consider a set of description algorithms $\mathcal{A} = \{A_1, A_2, \dots\}$. Our training batches consist of normalized images patches $\mathcal{P} = \{p_1, p_2, \dots\}$. Each description algorithm can be applied independently on a patch $p \in \mathcal{P}$ to obtain the associated descriptor, i.e., $A_i(p) = a_i$.

4.2. Pair network

Given a pair of different description algorithms (A_i, A_j) , one could approximate the translation function $t_{i \rightarrow j}$ by a multi-layer perceptron $N_{i \rightarrow j}$ (c.f. Figure 3 (a)). This model translates features extracted using algorithm A_i into the feature representation of algorithm A_j . To this end, we define the translation loss as the ℓ_2 prediction error:

$$\mathcal{L}_{i \rightarrow j}^T = \frac{1}{|\mathcal{P}|} \sum_{p \in \mathcal{P}} \|N_{i \rightarrow j}(A_i(p)) - A_j(p)\| \quad (1)$$

If the target algorithm A_j produces binary descriptors, the regression loss is replaced by a classification loss:

$$\mathcal{L}_{i \rightarrow j}^T = \frac{1}{|\mathcal{P}|} \sum_{p \in \mathcal{P}} \text{BCE}(N_{i \rightarrow j}(A_i(p)), A_j(p)) \quad , \quad (2)$$

where BCE is the binary cross entropy loss function defined as $\text{BCE}(x, y) = -y \log(x) - (1 - y) \log(1 - x)$.

To compute the distance d between two descriptors a_i, a_j extracted using algorithms A_i, A_j , respectively, one can translate from A_i to A_j by using the associated MLP ($d = \|N_{i \rightarrow j}(a_i) - a_j\|$) or conversely ($d = \|N_{j \rightarrow i}(a_j) - a_i\|$). In most practical applications, the direction of translation is known. In the case of continuous software updates, the outdated map descriptors are “migrated” to the new ones, while for cross-device localization the query descriptors are always translated to map descriptors. However, it is not obvious how to use this approach in the collaborative mapping scenario as there are $\mathcal{O}(|\mathcal{A}|^2)$ possible translation directions. One could use a “progressive” strategy that translates the weaker to the better descriptor for each image pair, but deciding which descriptor is more suitable is not straightforward as it generally depends on the specific scene.¹

4.3. Encoder-decoder network

To address this issue, we turn our attention towards auto-encoders. Each description algorithm A_i now has an associated encoder E_i and decoder D_i that ideally satisfy

¹Please refer to the supplementary material for a scene-by-scene breakdown on the HPatches Sequences dataset [7] that supports this argument.

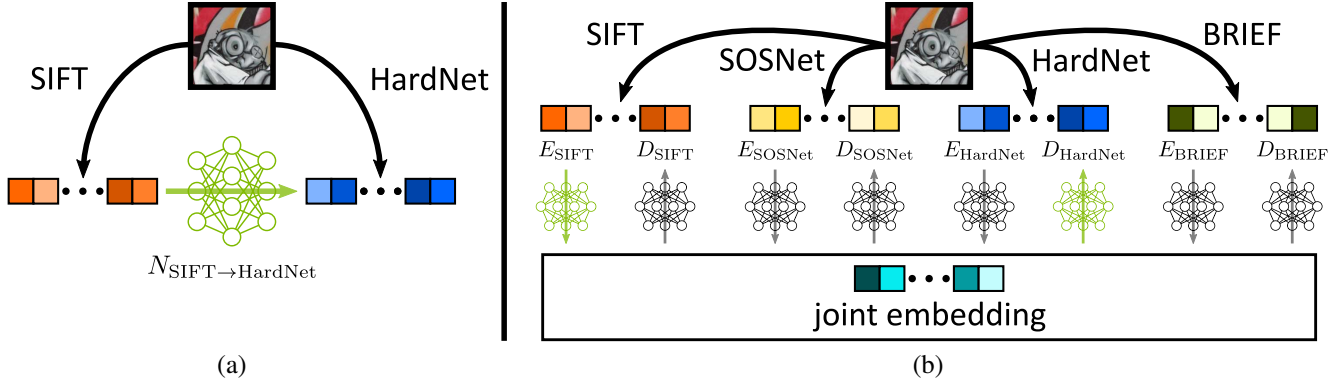


Figure 3: **Method overview.** (a) The pair network is trained for each description algorithm pair independently. (b) The encoder-decoder network is trained for all description algorithms at once. All descriptors are mapped to a joint embedding space. In green, we highlight the networks that need to be used in order to translate from SIFT to HardNet.

$D_i(E_i(a_i)) = a_i$ for all descriptors a_i extracted using A_i . We propose an extension to the identity loss traditionally used for training auto-encoders [9], such that the codes (*i.e.*, the output of encoders / the input of decoders) live in a joint space for all algorithms. Thus, for a pair of descriptor extractors A_i, A_j , one can chain the encoder of A_i and the decoder of A_j to obtain the mapping from A_i to A_j and vice versa, *i.e.*, $N_{i \rightarrow j} = D_j \circ E_i$ and $N_{j \rightarrow i} = D_i \circ E_j$ (see Fig. 3 (b)). In this way, we reduce the number of required networks to $\mathcal{O}(|\mathcal{A}|)$ and achieve linear scalability.

Furthermore, we use a triplet ranking loss to facilitate direct matching in the joint embedding space. Therefore, one can compare two descriptors a_i, a_j extracted using different algorithms A_i, A_j by mapping both descriptors to the joint space and computing the distance as $d = \|E_i(a_i) - E_j(a_j)\|$, instead of using the directional translation mentioned above. This is of particular interest in the collaborative mapping scenario, as it offers an elegant way to match descriptors coming from different algorithms by translating everything to the joint embedding space. Moreover, it allows compatibility with existing pipelines that suppose all images use the same feature.

Translation loss. As above, we consider the translation loss $\mathcal{L}_{i \rightarrow j}^T$ for each algorithm pair $A_i, A_j \in \mathcal{A}^2$. Note that, for $i = j$, this loss is equivalent to the traditional auto-encoder loss [9]. The overall translation loss is defined as the average over all algorithm pairs as:

$$\mathcal{L}^T = \frac{1}{|\mathcal{A}|^2} \sum_{A_i, A_j \in \mathcal{A}^2} \mathcal{L}_{i \rightarrow j}^T . \quad (3)$$

Matching loss. We use a triplet margin loss for each description algorithm pair $A_i, A_j \in \mathcal{A}^2$ to allow matching in the joint embedding space:

$$\mathcal{L}_{i \rightarrow j}^M = \frac{1}{|\mathcal{P}|} \sum_{p \in \mathcal{P}} \max(m + \text{pos}(p) - \text{neg}(p), 0) , \quad (4)$$

where $\text{pos}(p)$ and $\text{neg}(p)$ are the distances to the positive and negative samples, respectively. Given two corresponding descriptors $A_i(p) = a_i, A_j(p) = a_j$, the positive sample for the embedding of a_i is simply the embedding of a_j :

$$\text{pos}(p) = \|E_i(A_i(p)) - E_j(A_j(p))\| . \quad (5)$$

Following HardNet [33], we use in-batch hardest negative mining by selecting the closest non-matching embedding of a descriptor extracted using algorithm A_j as negative sample for the embedding of the current descriptor a_i :

$$\text{neg}(p) = \arg \min_{p' \in \mathcal{P}, p' \neq p} \|E_i(A_i(p)) - E_j(A_j(p'))\| . \quad (6)$$

The overall matching loss is then defined as the sum over all description algorithm pairs as:

$$\mathcal{L}^M = \frac{1}{|\mathcal{A}|^2} \sum_{A_i, A_j \in \mathcal{A}^2} \mathcal{L}_{i \rightarrow j}^M . \quad (7)$$

Final loss. The final loss is a weighted sum of the translation and matching losses as: $\mathcal{L} = \mathcal{L}^T + \alpha \mathcal{L}^M$.

4.4. Implementation details

Training dataset. Our training dataset contains 7.4 million patches from 3190 random internet images of the Oxford-Paris revisited retrieval dataset distractors [41]. The normalized image patches are extracted around Difference-of-Gaussians (DoG) detections according to the estimated scale and orientation. Training descriptors are computed by applying all description algorithms on each patch to obtain a set of corresponding descriptors. For DoG keypoint extraction we use COLMAP [47] with default parameters.

Training methodology. We train all networks for 5 epochs using Adam [24] as optimizer with a learning rate of 10^{-3} and a batch size of 1024. Similar to Mishchuk *et al.* [33], we use a margin $m = 1$. According to validation results, we fix

the weighting of the matching loss $\alpha = 0.1$. The encoders and decoders are MLPs with 2 hidden layers of sizes 1024 and 256 for hand-crafted and learned features, respectively. We use ReLU followed by batch normalization [20] after each linear layer apart from the last. For consistency with state-of-the-art floating point local descriptors [28, 33, 51], the joint embedding is 128-dimensional and ℓ_2 normalized. For binary descriptors [10], we use sigmoid activation after the last linear layer and, at test time, we use a threshold of 0.5 for binarization to allow efficient matching using bitwise operations. Similarly, if needed [28, 33, 51], the output of the network is ℓ_2 normalized. Please refer to the supplementary material for additional details.

5. Experimental evaluation

We evaluate our approach on the tasks of visual localization and mapping in the context of the scenarios introduced in Section 3. Please refer to the supplementary material for additional results on the HPatches descriptor benchmark [7], the Image Matching Workshop challenge [22], and the InLoc Indoor Visual Localization dataset [50]. We consider four well-established widely used descriptors – two handcrafted descriptors, with BRIEF [10] as a binary and SIFT [28] as a gradient-based approach as well as two state-of-the-art learned descriptors, HardNet [33] and SOSNet [51]. In our initial experiments, we found the pair network to perform on par with the encoder-decoder approach in terms of accuracy and runtime. However, the encoder-decoder architecture provides us with a joint embedding and covers all three discussed scenarios. Thus, all results presented in this section were obtained using the encoder-decoder approach. Please refer to the supplementary material for a comparison of the two approaches.

5.1. Visual localization

We consider the task of visual localization against pre-built maps on the challenging Aachen Day & Night localization benchmark [46]. We start by matching each database image with its 20 nearest spatial neighbors. COLMAP [47] is used for triangulation using the provided camera poses and intrinsics. Next, each query image is matched against its top 50 retrieval results according to NetVLAD [5]. Finally, we use COLMAP’s image registrator for localization with known intrinsics. According to standard procedure, the poses are submitted to the evaluation system [1] and we report the percentage of localized images at different real-world thresholds in Table 1.

Localization under continuous deployment. To evaluate this scenario, we build the reference map from a single description algorithm (*e.g.*, SIFT). Then, all the feature descriptors of the map are translated to a target feature representation (*e.g.*, SIFT \rightarrow HardNet). For the queries, we

Scenario	Database descriptor	Query descriptor	% localized queries			
			Day (824 images)		Night (98 images)	
			0.25m, 2°	0.5m, 5°	0.25m, 2°	0.5m, 5°
Standard	BRIEF	BRIEF	76.1	81.4	32.7	36.7
	SIFT	SIFT	82.5	88.7	52.0	61.2
	HardNet	HardNet	86.2	92.2	64.3	72.4
	SOSNet	SOSNet	86.4	92.7	65.3	75.5
Continuous deployment	BRIEF \rightarrow	SIFT	74.9 ^{-1.2}	80.5 ^{-0.9}	31.6 ^{-1.1}	36.7 ^{0.0}
		HardNet	81.4 ^{+5.3}	86.7 ^{+5.3}	44.9 ^{+12.2}	49.0 ^{+12.3}
		SOSNet	81.6 ^{+5.5}	86.9 ^{+5.5}	42.9 ^{+10.2}	46.9 ^{+10.2}
	SIFT \rightarrow	BRIEF	66.6 ^{-15.9}	73.1 ^{-15.6}	19.4 ^{-32.6}	23.5 ^{-37.7}
		HardNet	83.4 ^{+0.9}	90.9 ^{+2.2}	59.2 ^{+7.2}	66.3 ^{+5.1}
		SOSNet	84.2 ^{+1.7}	91.4 ^{+2.7}	55.1 ^{+3.1}	62.2 ^{+1.0}
	HardNet \rightarrow	BRIEF	70.5 ^{-15.7}	76.7 ^{-15.5}	22.4 ^{-41.9}	26.5 ^{-45.9}
		SIFT	81.2 ^{-5.0}	88.0 ^{-4.2}	41.8 ^{-22.5}	51.0 ^{-21.4}
		SOSNet	85.8 ^{-0.4}	92.4 ^{+0.2}	61.2 ^{-3.1}	68.4 ^{-4.0}
	SOSNet \rightarrow	BRIEF	68.8 ^{-17.6}	74.8 ^{-17.9}	18.4 ^{-46.9}	20.4 ^{-55.1}
		SIFT	81.7 ^{-4.7}	87.5 ^{-5.2}	42.9 ^{-22.4}	49.0 ^{-26.5}
		HardNet	85.9 ^{-0.5}	92.4 ^{-0.3}	63.3 ^{-2.0}	69.4 ^{-6.1}
Cross-device	BRIEF \leftarrow	SIFT	65.8 ^{-10.3}	71.6 ^{-9.8}	22.4 ^{-10.3}	24.5 ^{-12.2}
		HardNet	68.8 ^{-7.3}	74.0 ^{-7.4}	20.4 ^{-12.3}	26.5 ^{-10.2}
		SOSNet	66.6 ^{-9.5}	71.4 ^{-10.0}	22.4 ^{+10.3}	24.5 ^{-12.2}
	SIFT \leftarrow	BRIEF	75.6 ^{-6.9}	80.8 ^{-7.9}	28.6 ^{-23.4}	36.7 ^{-24.5}
		HardNet	81.2 ^{-1.3}	87.3 ^{-1.4}	46.9 ^{-5.1}	55.1 ^{-6.1}
		SOSNet	80.1 ^{-2.4}	87.3 ^{-1.4}	42.9 ^{-9.1}	46.9 ^{-14.3}
	HardNet \leftarrow	BRIEF	82.8 ^{-3.4}	88.8 ^{-3.4}	43.9 ^{-20.4}	49.0 ^{-23.4}
		SIFT	84.7 ^{-1.5}	91.0 ^{-1.2}	58.2 ^{-6.1}	67.3 ^{-5.1}
		SOSNet	86.2 ^{0.0}	92.7 ^{+0.5}	64.3 ^{0.0}	69.4 ^{-3.0}
	SOSNet \leftarrow	BRIEF	82.5 ^{-3.9}	88.2 ^{-4.5}	44.9 ^{-20.4}	49.0 ^{-26.5}
		SIFT	84.0 ^{-2.4}	91.1 ^{-1.6}	51.0 ^{-14.3}	57.1 ^{-18.0}
		HardNet	85.3 ^{-1.1}	91.9 ^{-0.8}	66.3 ^{+1.0}	72.4 ^{-3.1}
Collab.	Embed \leftarrow	BRIEF	80.6 ^{+4.5}	86.7 ^{+5.3}	48.0 ^{+15.3}	50.0 ^{+13.3}
		SIFT	82.8 ^{+0.3}	89.0 ^{+0.3}	50.0 ^{-2.0}	57.1 ^{-4.1}
		HardNet	85.1 ^{-1.1}	91.7 ^{-0.5}	55.1 ^{-9.2}	61.2 ^{-11.2}
		SOSNet	84.8 ^{-1.6}	90.9 ^{-1.8}	57.1 ^{-8.2}	60.2 ^{-15.3}

Table 1: **Localization under continuous deployment.** A reference map is built using the database description algorithm. The descriptors of this map are translated to a target query descriptor. **Cross-device localization.** A reference map is built using the database description algorithm. The descriptors of query images are translated to be compatible with the map. **Localization to collaborative maps.** The database images are partitioned in 4 balanced subsets, one for each description algorithm. Both database and query descriptors are mapped to the common embedding space. The absolute performance difference (^{superscript}) is color-coded according to the baseline used as reference.

use descriptors extracted directly by the target feature description algorithm (*e.g.*, HardNet) and match them against the translated features. The resulting matches are used for localization. Our approach not only enables updating descriptors without re-mapping the environment, but in many cases also unexpectedly improves the re-localization per-

formance. In particular, hand-crafted descriptors [10, 28] are significantly better after translating them to learned ones [33, 51]. For SIFT, the increase is of more than 2% and 5% in absolute performance at the largest threshold for day and night queries, respectively, and for BRIEF the difference is even more accentuated with 5% and 10%. While there is still a gap w.r.t. re-building the map from scratch using the learned descriptors, the performance can be further improved using a bootstrapping strategy (c.f. Section 3.1).

Cross-device localization. Inverse to the previous experiment, we now build maps using a description algorithm (e.g., SIFT) and instead translate the query descriptors (e.g., HardNet \rightarrow SIFT). Similarly to the previous case, compared to standard localization, we notice increased performance when translating hand-crafted descriptors to learned ones. Maps built from state-of-the-art descriptors (HardNet, SOSNet) preserve their advantages no matter the description algorithm used for queries. Not surprisingly, translating other descriptors to BRIEF [10] yields worse performance due to the more limited discriminative power of binary descriptors. Overall, our approach enables localization with a good performance in this previously infeasible scenario.

Localization to collaborative maps. For the final scenario, we randomly split the database images of the map into 4 balanced subsets (one for each description algorithm) and translate all descriptors to the joint embedding space (i.e., SIFT, HardNet, $\dots \rightarrow$ embedding). Similarly, all query descriptors are translated to the joint embedding space for matching. Although the map is built from heterogeneous descriptors, the localization performance at day is comparable to the best results in the previous experiments. Moreover, note that the performance of HardNet and SOSNet in this scenario is just slightly below the current state-of-the-art patch-based descriptor [51]. In addition, BRIEF queries achieve a substantial improvement of more than 5% for both day and night folds at the coarsest threshold. This is a very encouraging result for cloud-based solutions, as we enable multiple devices implementing different algorithms to not only collaborate on mapping a location, but also remain competitive with the state-of-the-art in localization.

5.2. Collaborative mapping

Next, we consider the large-scale Structure-from-Motion benchmark of Schönberger *et al.* [48]. Similar to other datasets [46, 22], to evaluate the camera poses, we generate pseudo-ground-truth intrinsics and extrinsics via an initial SfM process. Only the images registered in this step are considered.² For each method, we exhaustively match all remaining images using a mutual nearest-neighbors matcher with symmetric second nearest neighbor ratio test [28] (the threshold is set to 0.9 for all methods). Given

²Please refer to the supplementary material for additional details.

the feature matches, COLMAP [47] is used for geometric verification and sparse reconstruction with known intrinsics. In the context of collaborative mapping (c.f. Section 3.3), we randomly split the database into balanced sets corresponding to different description algorithms.

The resulting point-cloud statistics are reported in Table 2. We evaluate two variants of our method: i) **Embed** taking advantage of the joint embedding space and ii) **Progressive** using the hierarchy BRIEF \rightarrow SIFT \rightarrow HardNet \rightarrow SOSNet based on the rankings reported in [7, 51]. In this variant, the translation direction is chosen *online* towards the strongest in each pair. Given two images with descriptors X and Y respectively, we translate the lower one to the higher one in the order and then run matching on the obtained descriptors. For instance, in an image pair with BRIEF and HardNet, BRIEF would be translated to HardNet, while in an image pair with only SIFT descriptors, there would be no translation.

First, we compare the joint embedding approach with progressive translation. Both models achieve comparable results, but the jointly embedded map consistently registers more images and reconstructs more 3D points, albeit at the cost of shorter track lengths. Furthermore, the joint embedding also significantly simplifies the translation and matching in the reconstruction pipeline.

Next, we compare our approach with the traditional setup, where one has access to the same feature descriptor for all dataset images. While these baselines are not able to tackle collaborative mapping with heterogeneous descriptors, it serves as an upper-bound for our method. Both our approaches outperform maps built from only BRIEF or SIFT, but achieve slightly inferior results than state-of-the-art HardNet and SOSNet. However, it is important to keep in mind that, in our case, the dataset is split into 4 random subsets. Therefore, only half of the images have HardNet or SOSNet features. Finally, we compare against baselines that only have access to their associated subset of dataset images to demonstrate the actual merit of our work. This corresponds to the realistic setup in which 4 types of devices with incompatible descriptors map the same environment but the mapping system cannot jointly leverage them. In this case, both our approaches enable collaboration across devices and significantly outperform all baselines – we consistently register four times more images.

Figure 4 provides co-visibility statistics of the 3D points for the **Embed** model. An overwhelming majority of the 3D points have at least 2 different description algorithms in their tracks and more than 50% have 3 or 4. Furthermore, the co-occurrence matrix demonstrates that all descriptors contributed (almost) equally to the map. We refer to the supplementary material for statistics on the other datasets.

Dataset		Madrid Metropolis – 453 images					Gendarmenmarkt – 985 images					Tower of London – 730 images							
Method	% localized images			Num. 3D pts.	Track length	Reproj. error	% localized images			Num. 3D pts.	Track length	Reproj. error	% localized images			Num. 3D pts.	Track length	Reproj. error	
	0.25m 2°	0.5m 5°	∞				0.25m 2°	0.5m 5°	∞				0.25m 2°	0.5m 5°	∞				
Standard	BRIEF	57.0	64.2	72.4	18.3K	6.84	0.63	59.9	68.7	80.7	52.3K	6.31	0.82	64.9	68.5	74.2	48.1K	7.70	0.66
	SIFT	78.1	83.7	95.1	39.4K	6.71	0.83	68.8	77.1	95.6	121.4K	5.53	0.95	74.2	76.7	97.1	90.0K	7.14	0.81
	HardNet	89.2	95.4	100	47.4K	7.15	0.92	84.2	90.7	99.8	135.4K	6.40	1.01	83.0	87.7	100	104.5K	7.56	0.87
	SOSNet	92.7	96.3	100	46.0K	7.22	0.92	85.5	89.2	99.9	128.4K	6.56	1.02	85.2	89.2	100	101.3K	7.67	0.86
Real-world	BRIEF	3.1	4.9	6.0	1.9K	4.79	0.53	2.9	3.7	10.4	4.5K	4.36	0.72	10.4	11.5	11.8	5.8K	4.65	0.57
	SIFT	14.1	16.6	21.6	6.1K	4.70	0.77	5.4	6.8	20.9	15.7K	4.29	0.89	13.0	16.0	17.1	15.2K	4.55	0.74
	HardNet	10.8	13.9	21.4	8.8K	5.34	0.87	13.8	15.9	23.2	25.7K	4.67	0.95	16.0	17.8	18.6	22.1K	5.23	0.80
	SOSNet	16.8	19.4	21.2	7.5K	4.70	0.85	13.6	14.8	23.5	26.7K	4.95	0.96	17.3	18.8	23.0	22.9K	5.32	0.80
Ours	Embed	80.6	84.3	92.1	36.3K	7.26	0.86	74.1	82.3	95.6	103.4K	6.45	0.98	77.3	81.2	97.9	88.5K	7.57	0.81
	Progressive	77.0	82.8	88.7	31.2K	7.59	0.86	76.2	82.7	94.8	92.2K	6.58	0.98	79.2	83.2	96.6	76.0K	7.81	0.82

Table 2: **Collaborative mapping.** We report different reconstruction statistics from internet collected images. The first four rows represent the standard evaluation protocol where each description algorithm has access to all images. The next four rows present more realistic scenarios where each algorithm only has access to a quarter of the data. Finally, the last two rows represent variants of our method enabling cross-descriptor reconstruction using the same splits as above.




Figure 4: **Co-visibility statistics.** For the **Embed** approach on the *Tower of London* dataset, we report the % of 3D points containing 1 – 4 distinct algorithms in their tracks on the left. On the right, we visualize the co-occurrence, *i.e.*, the percentage of 3D points containing descriptors originating from 2 given description algorithms in their tracks.

6. Discussion

The main limitation comes from the **same keypoint** assumption, as our work only considers the compatibility of different description algorithms. While our assumption is likely applicable for devices produced by a single manufacturer that has control over the entire pipeline, different manufacturers likely use different keypoint detection algorithms. Possible directions for future research include studying cross-detector repeatability or improving feature inversion techniques for better generalization with different local feature detection pipelines. Despite our proposed solution being only a first step, it still has significant real-world applications: i) migration of existing maps from, *e.g.*, handcrafted descriptors to learned ones and, similarly, continuous updates of learned descriptors in case of novel architectures / losses / datasets, ii) support of legacy devices inside maps built with newer descriptors, and iii) collabora-

tive mapping between devices of the same manufacturer with different compute capabilities. Furthermore, while we have made a significant step towards enabling cross-descriptor localization and mapping, several challenges still remain. **Information loss.** Each description algorithm only encodes a subset of the information available in the input patches and these subsets might be distinct for different methods. This is further exacerbated in efficient real-world applications taking advantage of quantization and dimensionality reduction. **One-to-many associations.** A one-to-one mapping may not even exist as two given patches might lead to a similar descriptor for one algorithm but completely different feature vectors for another one. Similarly to other data-driven approaches, translation performance has to be verified experimentally as there are no formal guarantees. To address these limitations, future research could try to take advantage of local or global context to better disambiguate the visual information.

7. Conclusion

We have identified, introduced, and formalized three novel scenarios for localization and mapping in the presence of heterogeneous feature representations. Towards addressing the open challenges underlying these scenarios, we presented the first principled solution to cross-descriptor localization and mapping. We demonstrated the effectiveness of our approach on representative and well-established benchmarks. We believe our work will not only spark new research on the topic of translating local image features but will also make an immediate impact on commercial applications in the area of cloud-based localization and mapping.

Acknowledgments. This work was supported by the Microsoft MR & AI Lab Zürich PhD scholarship.

References

- [1] Long-Term Visual Localization Benchmark. <https://www.visuallocalization.net>. 6
- [2] Inside Facebook Reality Labs: Research updates and the future of social connection. <https://tech.fb.com/inside-facebook-reality-labs-research-updates-and-the-future-of-social-connection/>, 2019. 1
- [3] Rony Abovitz. What Is The Magicverse (And Why)? <https://www.magicleap.com/en-us/news/oped/magicverse>, 2019. 1
- [4] Sameer Agarwal, Yasutaka Furukawa, Noah Snavely, Ian Simon, Brian Curless, Steven M. Seitz, and Richard Szeliski. Building Rome in a Day. *Communications of the ACM*, 2011. 2
- [5] Relja Arandjelovic, Petr Gronat, Akihiko Torii, Tomas Pajdla, and Josef Sivic. NetVLAD: CNN architecture for weakly supervised place recognition. In *Proc. CVPR*, 2016. 6
- [6] Relja Arandjelovic and Andrew Zisserman. Three things everyone should know to improve object retrieval. In *Proc. CVPR*, 2012. 2
- [7] Vassileios Balntas, Karel Lenc, Andrea Vedaldi, and Krystian Mikolajczyk. HPatches: A benchmark and evaluation of handcrafted and learned local descriptors. In *Proc. CVPR*, 2017. 2, 4, 6, 7
- [8] Vassileios Balntas, Edgar Riba, Daniel Ponsa, and Krystian Mikolajczyk. Learning local feature descriptors with triplets and shallow convolutional neural networks. In *Proc. BMVC.*, 2016. 2
- [9] Hervé Bourlard and Yves Kamp. Auto-Association by Multilayer Perceptrons and Singular Value Decomposition. *Biological Cybernetics*, 1988. 5
- [10] Michael Calonder, Vincent Lepetit, Christoph Strecha, and Pascal Fua. BRIEF: Binary Robust Independent Elementary Features. In *Proc. ECCV*, 2010. 2, 6, 7
- [11] Titus Cieslewski, Simon Lynen, Marcin Dymczyk, Stéphane Magnenat, and Roland Siegwart. Map API - Scalable Decentralized Map Building for Robots. In *Proc. ICRA*, 2015. 2
- [12] Gabriela Csurka. A Comprehensive Survey on Domain Adaptation for Visual Applications. *Advances in Computer Vision and Pattern Recognition*, 2017. 3
- [13] Daniel DeTone, Tomasz Malisiewicz, and Andrew Rabinovich. SuperPoint: Self-Supervised Interest Point Detection and Description. In *CVPR Workshops*, 2018. 2
- [14] Alexey Dosovitskiy and Thomas Brox. Inverting visual representations with convolutional networks. In *Proc. CVPR*, 2016. 2
- [15] Mihai Dusmanu, Ignacio Rocco, Tomas Pajdla, Marc Pollefeys, Josef Sivic, Akihiko Torii, and Torsten Sattler. D2-Net: A Trainable CNN for Joint Detection and Description of Local Features. In *Proc. CVPR*, 2019. 2
- [16] Ebel, Patrick and Mishchuk, Anastasiia and Yi, Kwang Moo and Fua, Pascal and Trulls, Eduard. Beyond Cartesian Representations for Local Descriptors. In *Proc. ICCV*, 2019. 2
- [17] Dieter Fox, Wolfram Burgard, Hannes Kruppa, and Sebastian Thrun. A Probabilistic Approach to Collaborative Multi-Robot Localization. *Autonomous robots*, 2000. 2
- [18] Andreas Geiger, Julius Ziegler, and Christoph Stiller. StereoScan: Dense 3D Reconstruction in Real-time. In *Intelligent Vehicles Symposium (IV)*, 2011. 2
- [19] Jared Heinly, Johannes L. Schönberger, Enrique Dunn, and Jan-Michael Frahm. Reconstructing the World* in Six Days *(As Captured by the Yahoo 100 Million Image Dataset). In *Proc. CVPR*, 2015. 2
- [20] Sergey Ioffe and Christian Szegedy. Batch Normalization: Accelerating Deep Network Training by Reducing Internal Covariate Shift. *arXiv preprint arXiv:1502.03167*, 2015. 6
- [21] Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A. Efros. Image-to-Image Translation with Conditional Adversarial Networks. In *Proc. CVPR*, 2017. 3
- [22] Yuhe Jin, Dmytro Mishkin, Anastasiia Mishchuk, Jiri Matas, Pascal Fua, Kwang Moo Yi, and Eduard Trulls. Image Matching across Wide Baselines: From Paper to Practice. *arXiv preprint arXiv:2003.01587*, 2020. 2, 6, 7
- [23] Neena Kamath. Announcing Azure Spatial Anchors for collaborative, cross-platform mixed reality apps. <https://azure.microsoft.com/en-us/blog/announcing-azure-spatial-anchors-for-collaborative-cross-platform-mixed-reality-apps/>, 2019. 1
- [24] Diederik P. Kingma and Jimmy Lei Ba. Adam: A Method for Stochastic Optimization. *arXiv preprint arXiv:1412.6980*, 2014. 5
- [25] Alex Kipman. Azure Spatial Anchors approach to privacy and ethical design. <https://www.linkedin.com/pulse/azure-spatial-anchors-approach-privacy-ethical-design-alex-kipman>, 2019. 1
- [26] Bernd Kitt, Andreas Geiger, and Henning Lategahn. Visual Odometry based on Stereo Image Sequences with RANSAC-based Outlier Rejection Scheme. In *Intelligent Vehicles Symposium (IV)*, 2010. 2
- [27] Yunpeng Li, Noah Snavely, Dan Huttenlocher, and Pascal Fua. Worldwide Pose Estimation using 3D Point Clouds. In *Proc. ECCV*, 2012. 2
- [28] David G. Lowe. Distinctive image features from scale-invariant keypoints. *IJCV*, 2004. 2, 3, 6, 7
- [29] Zixin Luo, Tianwei Shen, Lei Zhou, Siyu Zhu, Runze Zhang, Yao Yao, Tian Fang, and Long Quan. GeoDesc: Learning local descriptors by integrating geometry constraints. In *Proc. ECCV*, 2018. 2
- [30] Zixin Luo, Lei Zhou, Xuyang Bai, Hongkai Chen, Jiahui Zhang, Yao Yao, Shiwei Li, Tian Fang, and Long Quan. ASLFeat: Learning Local Features of Accurate Shape and Localization. In *Proc. CVPR*, 2020. 2
- [31] Krystian Mikolajczyk and Cordelia Schmid. Scale & Affine Invariant Interest Points Detectors. *IJCV*, 2003. 2
- [32] Krystian Mikolajczyk and Cordelia Schmid. A Performance Evaluation of Local Descriptors. *IEEE PAMI*, 2005. 2
- [33] Anastasiya Mishchuk, Dmytro Mishkin, Filip Radenovic, and Jiri Matas. Working hard to know your neighbor's mar-

- gins: Local descriptor learning loss. In *Advances in NeurIPS*, 2017. 2, 5, 6, 7
- [34] John G. Morrison, Dorian Gálvez-López, and Gabe Sibley. MOARSLAM: Multiple Operator Augmented RSLAM. In *Distributed Autonomous Robotic Systems*. Springer, 2016. 2
- [35] Pierre Moulon, Pascal Monasse, Renaud Marlet, and Others. OpenMVG. <https://github.com/openMVG/openMVG>. 2
- [36] Raúl Mur-Artal, J. M. M. Montiel, and Juan D. Tardós. ORB-SLAM: a Versatile and Accurate Monocular SLAM System. *IEEE Transactions on Robotics*, 2015. 2
- [37] Raúl Mur-Artal and Juan D. Tardós. ORB-SLAM2: an Open-Source SLAM System for Monocular, Stereo and RGB-D Cameras. *IEEE Transactions on Robotics*, 2017. 2
- [38] Hyeonwoo Noh, Andre Araujo, Jack Sim, Tobias Weyand, and Bohyung Han. Large-Scale Image Retrieval with Attentive Deep Local Features. In *Proc. ICCV*, 2017. 2
- [39] Yuki Ono, Eduard Trulls, Pascal Fua, and Kwang Moo Yi. LF-Net: Learning local features from images. In *Advances in NeurIPS*, 2019. 2
- [40] Francesco Pittaluga, Sanjeev J Koppal, Sing Bing Kang, and Sudipta N Sinha. Revealing Scenes by Inverting Structure From Motion Reconstructions. In *Proc. CVPR*, 2019. 2
- [41] Filip Radenović, Ahmet Iscen, Giorgos Tolias, Yannis Avrithis, and Ondřej Chum. Revisiting Oxford and Paris: Large-Scale Image Retrieval Benchmarking. In *Proc. CVPR*, 2018. 5
- [42] Tilman Reinhardt. Google Visual Positioning Service. <https://ai.googleblog.com/2019/02/using-global-localization-to-improve.html>, 2019. 1
- [43] Jerome Revaud, Philippe Weinzaepfel, César Roberto de Souza, and Martin Humenberger. R2D2: Repeatable and Reliable Detector and Descriptor. In *Advances in NeurIPS*, 2019. 2
- [44] Ethan Rublee, Vincent Rabaud, Kurt Konolige, and Gary Bradski. ORB: an efficient alternative to SIFT or SURF. In *Proc. ICCV*, 2011. 2
- [45] Torsten Sattler, Bastian Leibe, and Leif Kobbelt. Improving Image-Based Localization by Active Correspondence Search. In *Proc. ECCV*, 2012. 2
- [46] Torsten Sattler, Will Maddern, Carl Toft, Akihiko Torii, Lars Hammarstrand, Erik Stenborg, Daniel Safari, Masatoshi Okutomi, Marc Pollefeys, Josef Sivic, Fredrik Kahl, and Tomas Pajdla. Benchmarking 6DoF Outdoor Visual Localization in Changing Conditions. In *Proc. CVPR*, 2018. 2, 6, 7
- [47] Johannes Lutz Schönberger and Jan-Michael Frahm. Structure-from-Motion Revisited. In *Proc. CVPR*, 2016. 2, 5, 6, 7
- [48] Johannes L. Schönberger, Hans Hardmeier, Torsten Sattler, and Marc Pollefeys. Comparative Evaluation of Hand-Crafted and Learned Local Features. In *Proc. CVPR*, 2017. 2, 7
- [49] Chris Sweeney. Theia Multiview Geometry Library: Tutorial & Reference. <http://theia-sfm.org>. 2
- [50] Hajime Taira, Masatoshi Okutomi, Torsten Sattler, Mircea Cimpoi, Marc Pollefeys, Josef Sivic, Tomas Pajdla, and Akihiko Torii. InLoc: Indoor Visual Localization with Dense Matching and View Synthesis. In *Proc. CVPR*, 2018. 2, 6
- [51] Yurun Tian, Xin Yu, Bin Fan, Fuchao Wu, Huub Heijnen, and Vassileios Balntas. SOSNet: Second Order Similarity Regularization for Local Descriptor Learning. In *Proc. CVPR*, 2019. 2, 3, 6, 7
- [52] Koen Van De Sande, Theo Gevers, and Cees Snoek. Evaluating Color Descriptors for Object and Scene Recognition. *IEEE PAMI*, 2009. 2
- [53] Kwang Moo Yi, Eduard Trulls, Vincent Lepetit, and Pascal Fua. LIFT: Learned invariant feature transform. In *Proc. ECCV*, 2016. 2
- [54] Jun-Yan Zhu, Philipp Krähenbühl, Eli Shechtman, and Alexei A. Efros. Generative visual manipulation on the natural image manifold. In *Proc. ECCV*, 2016. 3
- [55] Zhu, Jun-Yan and Park, Taesung and Isola, Phillip and Efros, Alexei A. Unpaired Image-To-Image Translation Using Cycle-Consistent Adversarial Networks. In *Proc. ICCV*, 2017. 3
- [56] Fuzhen Zhuang, Zhiyuan Qi, Keyu Duan, Dongbo Xi, Yongchun Zhu, Hengshu Zhu, Hui Xiong, and Qing He. A Comprehensive Survey on Transfer Learning. *arXiv preprint arXiv:1911.02685*, 2019. 3