Decoupling Makes Weakly Supervised Local Feature Better

Kunhong Li[1,2]   Longguang Wang[3]   Li Liu[3,4]   Qing Ran[5]   Kai Xu[3]   Yulan Guo[1,2,3*]
[1]Sun Yat-Sen University   [2]The Shenzhen Campus of Sun Yat-Sen University
[3]National University of Defense Technology   [4]University of Oulu   [5]Alibaba Group

Abstract

*Weakly supervised learning can help local feature methods to overcome the obstacle of acquiring a large-scale dataset with densely labeled correspondences. However, since weak supervision cannot distinguish the losses caused by the detection and description steps, directly conducting weakly supervised learning within a **joint** training describe-then-detect pipeline suffers limited performance. In this paper, we propose a **decoupled** training describe-then-detect pipeline tailored for weakly supervised local feature learning. Within our pipeline, the detection step is decoupled from the description step and postponed until discriminative and robust descriptors are learned. In addition, we introduce a line-to-window search strategy to explicitly use the camera pose information for better descriptor learning. Extensive experiments show that our method, namely PoSFeat (Camera **Po**se **S**upervised **Feat**ure), outperforms previous fully and weakly supervised methods and achieves state-of-the-art performance on a wide range of downstream tasks.*
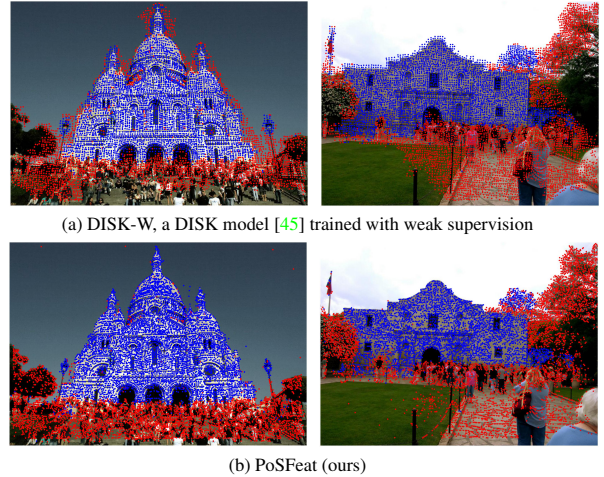
(a) DISK-W, a DISK model [45] trained with weak supervision

(b) PoSFeat (ours)

Figure 1. An illustration of the influence of ambiguity for weakly supervised local feature methods. Keypoints that succeed and fail to create landmarks are shown. (a) With **joint** training describe-then-detect pipeline, DISK-W [45] produces inaccurate keypoints that are out of the objects. (b) With our **decoupled** training describe-then-detect pipeline, PoSFeat can produce more reasonable keypoints. Best viewed in color.

1. Introduction

Finding pixel correspondences is a fundamental problem in computer vision. Sparse local feature [5,11,21,35,47], as one of the mainstream methods to find correspondences, has been widely applied in many areas, such as simultaneous localization and mapping (SLAM) [28, 51], structure from motion (SfM) [1, 39], and visual localization [15, 37, 50].

Traditional sparse local feature methods [5, 21, 35] follow a *detect-then-describe* pipeline. Specifically, keypoints are first detected and then patches centered at these keypoints are used to generate descriptors. Early methods [14, 16, 18, 21] focus on the detection step and are proposed to distinguish distinctive areas to detect good keypoints. Later works pay more attention to the description step and make attempts to design powerful descriptors using advanced representations [5, 7, 35].

Motivated by the success of deep learning, many efforts [22, 26, 30, 43, 49] have been made to replace the detection or description step in the *detect-then-describe* pipeline with CNNs. Recent works [12, 23, 33, 45] find that keypoints and descriptors are interdependent and propose a **joint** training describe-then-detect pipeline. Specifically, the description network and detection network are combined into a single CNN and optimized jointly. The **joint** training describe-then-detect pipeline achieves better performance than the *detect-then-describe* pipeline, especially under challenging conditions [17,44]. However, these methods are fully supervised and rely on dense ground-truth correspondence labels for training.

Because collecting a large dataset with pixel-level ground-truth correspondences is expensive, self-supervised and weakly supervised learning are investigated for training. Specifically, DeTone *et al.* [11] used a single im-

---

age and a virtual homography to generate image pairs to conduct self-supervised learning. However, homography transformation cannot cover complicated geometry transformations in real-world settings, resulting in limited performance. Noh *et al.* [29] used landmark labels to train the local feature network, which suffers extremely poor performance on viewpoint changes. Owing to the convenience of collecting camera poses, Wang *et al.* [46] introduced camera poses as weak supervision for descriptor learning. Although weakly supervised learning achieves promising results within the *detect-then-describe* pipeline, directly applying it to the *joint training describe-then-detect* pipeline is hard to produce satisfying results [45].

When a detection network and a description network are jointly optimized within a *joint training describe-then-detect* pipeline with only weak supervision (*e.g.*, camera pose), the loss produced by these two components cannot be distinguished. Specifically, when only one component is failed (Fig. 2), both the detection network and the description network cannot be correctly updated within a *joint training describe-then-detect* pipeline. As a result, the description network is hard to produce highly discriminative descriptors, and the detection network may produce false detected keypoints that are out of object boundaries, as shown in Fig. 1.

In this paper, we propose a *decoupled training describe-then-detect* pipeline tailored for weakly supervised local feature learning. Our main insight is that, with only weak supervision, the detection network relies heavily on a good descriptor for accurate keypoint detection (Fig. 1). Consequently, we decouple the detection network from the description network to postpone it until a discriminative and robust descriptor is learned. Different from the *detect-then-describe* pipeline that relies on low-level structures for early detection, our keypoints detection depends on the higher-level structures encoded in the descriptors. As a result, better robustness is achieved. In contrast to the *joint training describe-then-detect* pipeline that simultaneously perform detection and description optimization, the two networks are trained separately and thus the loss function for these two components are decoupled to address the ambiguity. It is demonstrated that our *decoupled training describe-then-detect* pipeline facilitates local feature methods to achieve much better performance with only weak supervision. Our contributions can be summarized as:

(1) We introduce a *decoupled training describe-then-detect* pipeline for weakly supervised local feature learning. This simple yet efficient pipeline significantly improves the performance of weakly supervised local features.

(2) We propose a line-to-window search strategy to exploit the weak supervision of camera poses for descriptor learning. This strategy can make full use of the geometric information of camera poses to reduce the search space and


Figure 2. Motivation of decoupling. Two reasonable keypoints can be matched incorrectly due to the low discriminativeness of descriptors (e.g., caused by repetitive textures). Meanwhile, two false detected keypoints can also be matched with a high descriptor similarity. Best viewed in color.

learn highly discriminative descriptors.

(3) Our method achieves state-of-the-art performance on three datasets and largely closes the gap between fully and weakly supervised methods.

2. Related Works

2.1. Fully Supervised Local Feature Methods

Fully supervised methods conduct local feature learning using pixel-level ground-truth correspondences to provide supervision. Following the *detect-then-describe* pipeline, early learning-based methods [4,13,22,26,38,43] use CNNs to perform the detection or description steps. Specifically, QuadNet [38] and Key.Net [4] were proposed to use CNNs for keypoint detection. HardNet [26] and SOSNet [43] were developed to leverage CNNs to extract descriptors. Later, LIFT [49] and LFNet [30] were introduced to integrate both detection and description steps into an end-to-end architecture to achieve better performance. Note that, LIFT [49] also introduced a decoupled training to address unstable training issue with full supervision in a *detect-then-describe* pipeline.

Recent works [12, 23, 33, 45] follow a *joint training describe-then-detect* pipeline in which detection and description are combined into a single CNN and optimized jointly. Specifically, Dusmanu *et al.* [12] first used a CNN to extract dense features and then selected local maxima of the dense feature map as keypoints. Revaud *et al.* [33] further took both the repeatablity and reliability of the descriptors into consideration for better keypoint detection. Tyszkiewicz *et al.* [45] used policy gradient to address the discreteness during the selection of sparse keypoints (namely, DISK). Luo *et al.* [23] adopted deformable convolution to model the geometry information and detected keypoints at multiple scales. By jointly optimizing the detection network and the description network, *joint training describe-then-detect* pipeline achieves better performance than previous *detect-then-describe* pipeline.
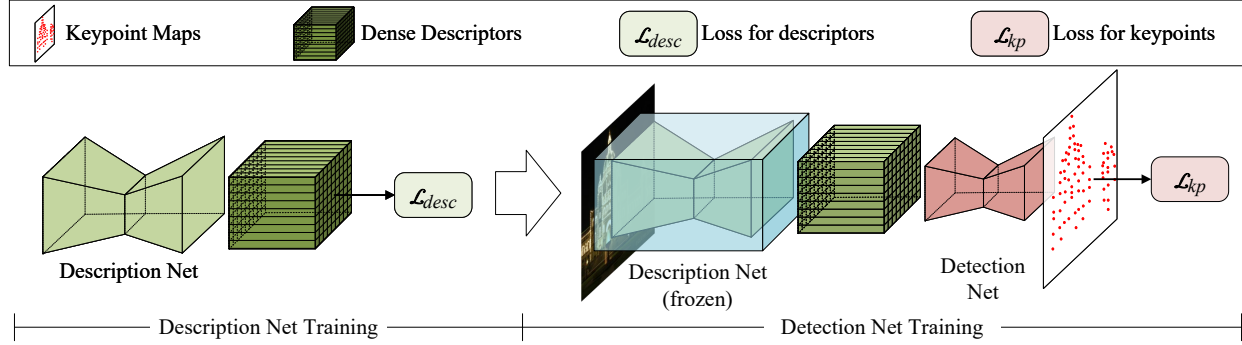
Figure 3. The proposed ***decoupled*** *training describe-then-detect* pipeline. The detection network is decoupled from the description network and postponed until good descriptors are obtained.

2.2. Self-Supervised Local Feature Methods

As a large dataset with densely labeled correspondences is difficult to collect, self-supervised learning has been studied for local feature learning. Specifically, DeTone *et al.* [11] used a virtual homography to generate an image pair from a single image to conduct self-supervised learning. This method uses a CNN pretrained on synthetic data as a teacher of the detection network. Differently, Christiansen *et al.* [10] proposed an end-to-end framework to train both the detection network and the description network using virtual homography in a self-supervised manner. Later, Parihar *et al.* [31] leveraged the homography to enhance the robustness of descriptors to rotation. Nevertheless, simple homography transformations used in these self-supervised methods may not hold in real cases.

2.3. Weakly Supervised Local Feature Methods

Noh *et al.* introduced DELF [29], which is trained with an image retrieval task, to achieve local feature extraction. However, the keypoints detected by DELF are sensitive to viewpoint changew and thus cannot be applied in real world settings. For camera poses are easy to collect, Wang *et al.* [46] used them as weak supervision and introduced an epipolar loss for descriptor learning. This method follows a *detect-then-describe* pipeline and relies on an off-the-shelf detection method (*e.g.*, SIFT) to detect keypoints. Recently, Tyszkiewicz *et al.* [45] developed DISK-W to integrate weakly supervised learning in a ***joint** training describe-then-detect* pipeline by adopting policy gradient. Nevertheless, when DISK-W is directly trained with a weakly supervised loss (rather than a fully-supervised loss), it suffers a notable performance drop on pixel-wise metrics. As weakly supervised loss cannot distinguish between errors introduced by false keypoints and inaccurate descriptors, this ambiguity hinders the ***joint** training describe-then-detect* pipeline to learn good local features.

2.4. Learning-based Matcher Methods

Since a Brute Force Matcher (also named NN matcher) usually produces low quality raw matches, learning-based matchers are proposed to achieve better matching results. Sarlin *et al.* [36] proposed SuperGlue to achieve robust matching with a graph neural network (GNN) and an optimal transport algorithm. Chen *et al.* [9] improved the architecture of GNN to increase the efficiency of descriptor enhancement. Zhou *et al.* [52] proposed a weakly supervised network to refine raw matches using patch matches as prior. Sun *et al.* [41] introduced a detector-free matcher to achieve pixel correspondence in a coarse-to-fine manner. Note that, most matcher methods are not the direct competitors of local feature methods. Instead, they can be considered as a post processing step and combined with local features to achieve improved performance.

3. Decoupled Training Describe-then-Detect Pipeline

3.1. Overview

The ***decoupled*** *training describe-then-detect* pipeline is shown in Fig. 3. We train the description net and detection net individually to suppress the loss ambiguity caused by weak supervision. During training, we first leave out the detection network and optimize the description network to learn good descriptors with a line-to-window strategy. The description network is then frozen to train a detection network for keypoint detection. We follow CAPS [46] to use ResUNet as the description net, which produces a feature map with 1/4 resolution and 128 dimensions as dense descriptors. Additionally, we design a shallow detection net to detect keypoints at the original resolution. For more details about the network architecture, please refer to the supplementary material.
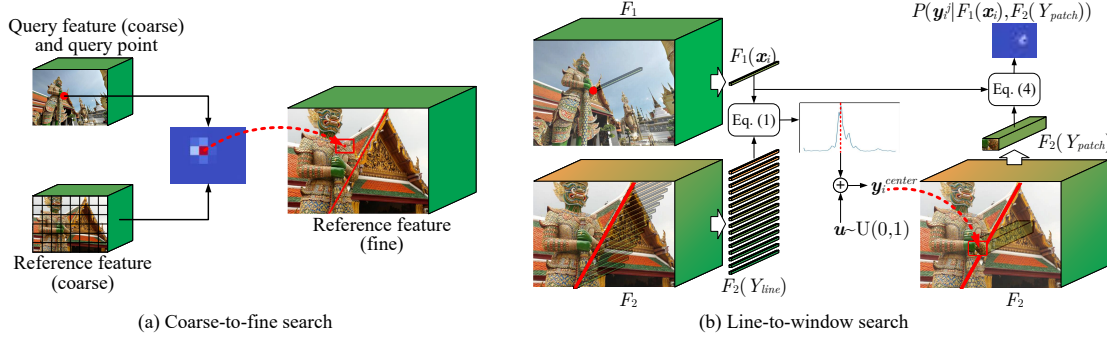
Figure 4. An illustration of the coarse-to-fine search strategy (a) and our line-to-window search strategy (b). The red line in F_2 denotes the epipolar line corresponding to the query point in F_1.

3.2. Feature Description

Following the widely used paradigm [46], we impose supervision only on sparse query points sampled from paired images to conduct training of the description network. We first split an image into small grids of size $g_d \times g_d$, and randomly sample one point per grid as a query point. Then, we translate relative camera pose into an epipolar constraint and introduce a line-to-window search strategy to reduce search space (Sec. 3.2.1). Moreover, we formulate a loss function by encouraging the predicted matches to obey the epipolar constraint (Sec. 3.2.2).

3.2.1 Line-to-Window Search

Given a query point x in the query image I_1, our goal is to find its correspondence in the reference image I_2. Since repetitive structures widely exist in a natural image, the commonly used coarse-to-fine strategy [41, 46] usually selects a mismatched patch such that inferior performance is produced (Fig. 4(a)). Intuitively, the correspondence of the query point x is constrained in an epipolar line in the reference image. Therefore, we introduce a line-to-window search strategy to reduce search space for better performance. Our line-to-window search strategy consists of two search steps, as illustrated in Fig. 4(b).

Search along An Epipolar Line. For a query point $x_i \in I_1$, we first calculate its corresponding epipolar line L_{x_i} in the reference image I_2 based on the relative camera pose. Then, we uniformly sample N_{line} points along this epipolar line to formulate the search space $Y_{line} = \{y_i^j\}(j = 1, ..., N_{line})$. Next, we calculate the matching probability of x_i over Y_{line}:

$$P(y_i^j|F_1(x_i), F_2(Y_{line})) = \frac{\exp(F_1(x_i)^{\mathrm{T}} F_2(y_i^j))}{\sum_{Y_{line}} \exp(F_1(x_i)^{\mathrm{T}} F_2(y_i^k))}, \tag{1}$$

where F_1 and F_2 are the feature maps for I_1 and I_2, respectively. Afterwards, we select $\overline{y}_i$ with the maximum prob-

ability from Y_{line} to determine the coarse location of the correspondence of x_i:

$$\overline{y}_i = \arg\max_{y_i^j} P(y_i^j|F_1(x_i), F_2(Y_{line})). \tag{2}$$

Search in A Local Window. Due to the discreteness of the candidates in Y_{line}, the resultant corresponding point $\overline{y}_i$ can be far from the groundtruth. To remedy this, a subsequent search is conducted in a local window. First, we calculate the center of the local window:

$$y_i^{center} = \overline{y}_i + 0.5 \cdot w_{patch} \cdot u, \tag{3}$$

where w_{patch} is the window size of a local patch, $u \in \mathbb{R}^2$ is a noise vector drawn from a uniform distribution $\mathrm{U}(0, 1)$ to avoid the convergence to trivial solution $F(x) \equiv 0$. Then, a local patch $Y_{patch} \subset I_2$ centered at y_i^{center} is cropped from F_2 as the search space. Next, we calculate the matching probability of x_i over Y_{patch}:

$$P(y_i^j|F_1(x_i), F_2(Y_{patch})) = \frac{\exp(F_1(x_i)^{\mathrm{T}} F_2(y_i^j))}{\sum_{Y_{patch}} \exp(F_1(x_i)^{\mathrm{T}} F_2(y_i^k))}. \tag{4}$$

Because directly selecting the point with the maximum probability in the local patch is non-differentiable, we calculate the correspondence $\hat{y}_i$ in a differentiable manner:

$$\hat{y}_i = E(y_i^j) = \sum_{y_i^j \in Y_{patch}} y_i^j \cdot P(y_i^j|F_1(x_i), F_2(Y_{patch})). \tag{5}$$

Compared to the previous coarse-to-fine search strategy [46], our line-to-window search strategy can make better use of the camera pose information to reduce search space and further improve the discriminativeness of descriptors (as demonstrated in Sec. 4.3).

3.2.2 Loss Function

With only weak supervision of camera pose, we calculate the distance of the correspondence $\hat{\boldsymbol{y}}_i$ to the epipolar line $L_{\boldsymbol{x}_i}$ as the loss of query point $\boldsymbol{x}_i$ [46]:

$$\mathcal{L}_{epi}(\hat{\boldsymbol{y}}_i, \boldsymbol{x}_i) = \text{distance}(\hat{\boldsymbol{y}}_i, \ L_{\boldsymbol{x}_i}). \tag{6}$$

Then, we use the weighted sum of the losses over all query points as the final loss:

$$\mathcal{L}_{desc} = \frac{\sum_i \frac{M_i}{\sigma(\boldsymbol{x}_i)} \cdot \mathcal{L}_{epi}(\hat{\boldsymbol{y}}_i, \boldsymbol{x}_i)}{\sum_i \frac{M_i}{\sigma(\boldsymbol{x}_i)}}. \tag{7}$$

Here, M_i is a binary mask (which is used to exclude query points whose epipolar lines are not in the reference image) and $\sigma(\boldsymbol{x}_i)$ is the variance of the probability distribution over Y_{patch},

$$\sigma(\boldsymbol{x}_i) = \|\hat{\boldsymbol{y}}_i^2 - E(\boldsymbol{y}_i^{j2})\| \tag{8}$$

3.3. Feature Detection

After feature description learning, the description network is frozen to produce dense descriptors for keypoint detection, as shown in Fig. 3. Since selecting discrete sparse keypoints is non-differentiable, we adopt policy gradient introduced in DISK [45] to achieve network training.

First, dense descriptors F_1 and F_2 are respectively extracted from I_1 and I_2, and fed to a detection network to produce keypoint heatmaps. Then, we divide these heatmaps into grids of size $g_k \times g_k$ and select at most one keypoint from each grid cell. Specifically, we establish a probability distribution P_{kp} over each grid cell based on the heatmap scores in this cell. Afterwards, P_{kp} is used to probabilistically select candidate keypoints $Q_1 = \{\boldsymbol{x}_1, \boldsymbol{x}_2, \cdots\}$ and $Q_2 = \{\boldsymbol{y}_1, \boldsymbol{y}_2, \cdots\}$ from I_1 and I_2, respectively. Next, a matching probability P_m is calculated based on the feature similarity $S_{i,j}$ between each pair of candidate keypoints $(\boldsymbol{x}_i, \boldsymbol{y}_j)$. With only camera pose supervision, we adopt an epipolar reward similar to Eq. 6 to encourage $\boldsymbol{y}_j$ to be close to the epipolar line of $\boldsymbol{x}_i$ (*i.e.*, $L_{\boldsymbol{x}_i}$):

$$R(\boldsymbol{x}_i, \boldsymbol{y}_j) = \begin{cases} \lambda_p, & \text{if distance}(\boldsymbol{y}_j, L_{\boldsymbol{x}_i}) \le \epsilon \\ \lambda_n, & \text{if distance}(\boldsymbol{y}_j, L_{\boldsymbol{x}_i}) > \epsilon \end{cases}, \tag{9}$$

where the reward threshold ϵ is empirically set to 2. The overall loss function is defined as:

$$\mathcal{L}_{kp} = - \frac{1}{|Q_1| + |Q_2|} \Big(\sum_{\boldsymbol{x}_i, \boldsymbol{y}_j} \mathcal{L}_{rew}(\boldsymbol{x}_i, \boldsymbol{y}_j)$$
$$+ \lambda_{reg} \big(\sum_{\boldsymbol{x}_i} \log P_{kp}(\boldsymbol{x}_i) + \sum_{\boldsymbol{y}_j} \log P_{kp}(\boldsymbol{y}_j) \big) \Big), \tag{10}$$

where λ_{reg} is a regularization penalty and the reward loss $\mathcal{L}_{rew}(\boldsymbol{x}_i, \boldsymbol{y}_j)$ is defined as:

$$\mathcal{L}_{rew}(\boldsymbol{x}_i, \boldsymbol{y}_j) = P_m(\boldsymbol{x}_i, \boldsymbol{y}_j) \cdot R(\boldsymbol{x}_i, \boldsymbol{y}_j) \cdot \log(P_{kp}(\boldsymbol{x}_i) P_{kp}(\boldsymbol{y}_i)). \tag{11}$$

Since our descriptors are well optimized, P_m can suppress spurious points with low scores. In contrast, in a joint pipeline, descriptors are under-optimized such that spurious points cannot be well distinguished. Please refer to the supplementary material for more details.

4. Experiments

4.1. Experimental Settings

Datasets The MegaDepth dataset [20] was used for training. We used a subset of the training split of CAPS [46]. Totally, 127 out of 196 scenes were used as the training set. **Implementation Details** During the training phase, images were resized to 640×480 with breaking the aspect ratio. All networks were trained using a SGD optimizer with nesterov momentum [42]. The learning rate is set to 1×10^{-3} and the batch size was set to 6. The description network was trained for 100000 iterations, and the detection network was trained for 5000 iterations. All experiments were conducted using Pytorch on a single NVIDIA RTX3090 GPU. In our experiments, the number of sampled points N_{line} was set to 100, the window size w_{patch} was set to 0.1 (normalized height and width), and the grid size g_d and g_k were set to 16 and 8, respectively. Following [45], λ_p, λ_n, and λ_{reg} were set to 1, -0.25, and -0.001, respectively. For more details, please refer to the supplementary material.

4.2. Comparison with Previous Methods

4.2.1 Feature Matching

Settings. We first evaluate our method on the widely used HPatches dataset [3]. Following D2-Net [12], 8 high-resolution scenes are removed and the remaining 52 scenes with illumination changes and 56 scenes with viewpoint changes are included for evaluation. Mean matching accuracy (MMA) [12] with thresholds ranging from 1 to 10 is used for evaluation. We also use a weighted sum of MMA at different thresholds for overall evaluation:

$$\text{MMAscore} = \frac{\sum_{\text{thr} \in [1,10]} (2 - 0.1 \cdot \text{thr}) \cdot \text{MMA@thr}}{\sum_{\text{thr} \in [1,10]} (2 - 0.1 \cdot \text{thr})}. \tag{12}$$

Three families of methods are included for comparison:

- **Patch-based methods:** Hessian-Affine keypoints [25] with Root-SIFT [2] (Hes. Aff. + Root-SIFT), affine region detector HesAffNet [27] with HardNet++ [26] (HAN + HN++), and SIFT [21] with ContextDesc [22] (SIFT + ContextDesc).

Figure 5. Results achieved on the HPatches dataset [3]. Mean match accuracy (MMA) achieved at different thresholds are illustrated. Learning based methods with weak supervision are shown in solid lines while other methods are shown in dashed lines. The numbers of keypoints and matches for each method are also reported.

Methods	Num. Feature	Num. match
Hes. Aff. + Root-SIFT	6710	2851
HAN + HN++	3860	1960
SIFT + ContextDesc	4066	1744
D2-Net	2994	1182
R2D2	4996	1850
ASLFeat	4013	2009
DISK	7705	3851
DELF	4590	1940
SuperPoint	1562	883
SIFT + CAPS	4386	1450
DISK-W	6760	3976
PoSFeat (Ours)	8192	4275

(a) Image pairs (b) SIFT + CAPS (c) DISK-W (d) PoSFeat (Ours)

Figure 6. Visualization results achieved on HPatches. For simplicity, only successfully matched keypoints are shown and colored according to their match errors. The colorbar is shown on the right. Best viewed in color.

Methods	MMAscore Overall	MMAscore Illumination	MMAscore Viewpoint
Hes. Aff. + Root-SIFT [2]	0.584	0.544	0.624
HAN [27] + HN++ [26]	0.633	0.634	0.633
SIFT [21] + ContextDesc [22]	0.636	0.613	0.657
D2Net [12]	0.519	0.605	0.440
R2D2 [33]	0.695	0.727	0.665
ASLFeat [23]	0.739	0.795	0.687
DISK [45]	0.763	0.813	0.716
DELF [29]	0.571	**0.903**	0.262
SuperPoint [11]	0.658	0.715	0.606
SIFT [21] + CAPS [46]	0.699	0.764	0.639
DISK-W [45]	0.719	0.803	0.649
PoSFeat (Ours)	**0.775**	0.826	**0.728**

Table 1. MMAscore results achieved by different methods on the HPatches dataset [3]. The MMAscores are calculated from Fig. 5.

- **Fully supervised dense feature methods:** D2-Net [12], R2D2 [33], ASLFeat [23], and DISK [45].

- **Weakly supervised dense feature methods:** DELF [29], SuperPoint [11], DISK-W [45], and SIFT with CAPS [46] (SIFT + CAPS).

Results. As shown in Fig. 5 and Table 1, the proposed PoSFeat outperforms all previous works, with the highest MMAscore being achieved. Compared to existing weakly supervised methods, our method produces significant performance improvements. Specifically, our method outper-

forms DISK-W by notable margins under both illumination (0.826 vs. 0.803) and viewpoint (0.728 vs. 0.649) changes, and therefore achieves higher overall MMAscore (0.775 vs. 0.719). We also visualize the matching results in Fig. 6. It can be seen that our PoSFeat produces more reasonable keypoints and less wrong matches. Compared to fully supervised methods, our method still performs favorably with higher MMA scores. This clearly demonstrates the superiority of our method. Note that, because DELF detects keypoints in a low resolution feature map with a fixed grid, it produces the best results under illumination change. However, our method significantly surpasses DELF under viewpoint change (0.728 vs. 0.262) and achieves much better overall performance (0.775 vs. 0.571).

4.2.2 Visual Localization

Settings. We then evaluate our method on the visual localization task with the Aachen Day-Night dataset [50]. We adopt the official visual localization pipeline[1] used in the local feature challenge of workshop on long-term visual localization under changing conditions. This challenge only evaluates the pose of night-time query images. Accuracy with different thresholds are used as metrics, including

---

[1] https://github.com/tsattler/visuallocalizationbenchmark/tree/master/local_feature_evaluation

Method	Aachen Day-Night v1			Aachen Day-Night v1.1		
	(0.5m,2°)	(1m,5°)	(5m, 10°)	(0.5m,2°)	(1m,5°)	(5m, 10°)
SP [11]	74.5	78.6	89.8	-	-	-
D2-Net [12]	74.5	86.7	**100**	-	-	-
R2D2 [33]	76.5	**90.8**	**100**	71.2	86.9	97.9
ASLFeat [23]	**81.6**	87.8	**100**	-	-	-
ISRF [24]	-	-	-	69.1	**87.4**	**98.4**
LISRD [32]	-	-	-	<u>73.3</u>	86.9	97.9
PoSFeat (Ours)	**81.6**	**90.8**	**100**	**73.8**	**87.4**	**98.4**
DualRC-Net [19]	-	-	-	71.2	86.9	97.9
SP+SuperGlue [36]	79.6	90.8	100	73.3	88.0	98.4
Sparse-NCNet [34]	76.5	84.7	98.0	-	-	-
LoFTR [41]	-	-	-	72.8	88.5	99.0
Patch2Pix [52]	79.6	87.8	100	-	-	-
SP+SGMNet [9]	77.6	88.8	99.0	72.3	85.3	97.9

Table 2. Results achieved by different methods on the Aachen Day-Night dataset [50]. 'LISRD' represents LISRD with Super-Point keypoints and AdaLAM [8]. Two categories of methods are presented, including feature methods (top) and matchers (bottom).

(0.5m, 2°), (1m, 5°), and (5m, 10°).

We compare our method with two families of methods:

- **Local feature methods:** D2-Net [12], SuperPoint [11], R2D2 [33], ASLFeat [23], ISRF [24], and LISRD [32].

- **Matcher methods:** DualRC-Net [19], SuperGlue [36] + SuperPoint, SparseNCNet [34], LoFTR [41], Patch2Pix [52], and SGMNet [9] + SuperPoint. As mentioned in Sec 2.4, matchers are the cooperators instead of the direct competitors of local features. Therefore, we group them separately.

Results. As shown in Table 2, our PoSFeat achieves the state-of-the-art performance among the feature methods. Specifically, on Aachen Day-Night v1, our method achieves the best accuracy in terms of all metrics. Note that, although ASLFeat is a fully supervised method, our PoSFeat still outperforms it on (1m, 5°). On Aachen Day-Night v1.1, our method also produces the best performance in all metrics. Note that, although R2D2 [33], ISRF [24], and LISRD [32] are fully-supervised and trained on the Aachen Day-Night dataset, our PoSFeat still achieves better results. We additionally include matcher methods for further comparison. Although these methods take pairs of images as inputs, our PoSFeat achieves comparable or even better performance.

4.2.3 3D Reconstruction

Settings. We finally evaluate our method on the 3D reconstruction task. We conduct experiments on the ETH local feature benchmark [40]. Four metrics are used for evaluation, including the number of registered images (# Imgs), the number of sparse points (# Pts), track length, and the mean reprojection error (Reproj. Err.).

Four families of methods were included for comparison:

- **Patch-based method:** Root-SIFT [2, 21].

Subset	Method	# Imgs	# Pts	Track Length	Reproj. Err. (px)
South Building (128 imgs)	Root-SIFT [2, 21]	128	108k	6.32	**0.55**
	SuperPoint [11]	128	**160k**	7.83	0.92
	RFP [6]	128	102k	7.86	0.88
	DISK [45]	128	115k	**9.91**	0.59
	DISK-W [45]	128	<u>154k</u>	<u>9.63</u>	0.63
	PoSFeat (Ours)	128	148k	9.47	<u>0.58</u>
Madrid Metropolis (1344 imgs)	Root-SIFT [2, 21]	500	116k	6.32	**0.60**
	SuperPoint [11]	438	29k	<u>9.03</u>	1.02
	D2-Net [12]	501	84k	6.33	1.28
	ASLFeat [23]	613	96k	8.76	0.90
	CAPS [46]	**851**	<u>242k</u>	6.16	1.03
	CoAM [48]	<u>702</u>	**256k**	6.09	1.30
	PoSFeat (Ours)	419	72k	**9.18**	<u>0.86</u>
Gendar-menmarkt (1463 imgs)	Root-SIFT [2, 21]	1035	339k	5.52	**0.70**
	SuperPoint [11]	967	93k	7.22	1.03
	D2-Net [12]	1053	250k	5.08	1.19
	ASLFeat [23]	1040	221k	**8.72**	1.00
	CAPS [46]	**1179**	**627k**	5.31	1.00
	CoAM [48]	<u>1072</u>	<u>570k</u>	6.60	1.34
	PoSFeat (Ours)	956	240k	<u>8.40</u>	<u>0.92</u>
Tower of London (1576 imgs)	Root-SIFT [2, 21]	806	239k	7.76	**0.61**
	SuperPoint [11]	681	52k	8.67	0.96
	D2-Net [12]	785	180k	5.32	1.24
	ASLFeat [23]	<u>821</u>	222k	**12.52**	0.92
	CAPS [46]	**1104**	**452k**	5.81	0.98
	CoAM [48]	804	239k	5.82	1.32
	PoSFeat (Ours)	778	<u>262k</u>	<u>11.64</u>	<u>0.90</u>

Table 3. Results achieved by different methods on the ETH local feature benchmark.

- **Fully supervised dense feature methods:** Reinforced Feature Points [6] (RFP), DISK [45], DISK-W [45], D2-Net [12], and ASLFeat [23].

- **Weakly supervised dense feature methods:** SuperPoint [11] and CAPS [46].

- **Fully supervised matcher method:** CoAM [48].

Results. As shown in Table 3, our method performs favorably against previous methods on the 3D reconstruction task. Specifically, our method produces the lowest reprojection error among all learning-based methods. Moreover, our method achieves the best or second best performance in terms of track length, which demonstrates that our keypoints are robust and thus can be tracked across a large amount of images.

4.3. Ablation Study

In this section, we first conduct ablation experiments on the HPatches dataset [3] to demonstrate the effectiveness of our *decoupled training describe-then-detect* pipeline and line-to-window search strategy. Then, we conduct experiments to study the effectiveness of hyper-parameters in our method, *i.e.*, the number of points sampled from the epipolar line N_{line} and the window size w_{patch}. Results and model settings are shown in Fig. 7 and Table 4.

Decoupled Training Describe-then-Detect Pipeline. We first constructed a network variant (Model 2) following the *joint training describe-then-detect* pipeline. That is, the description network and the detection network are jointly op-
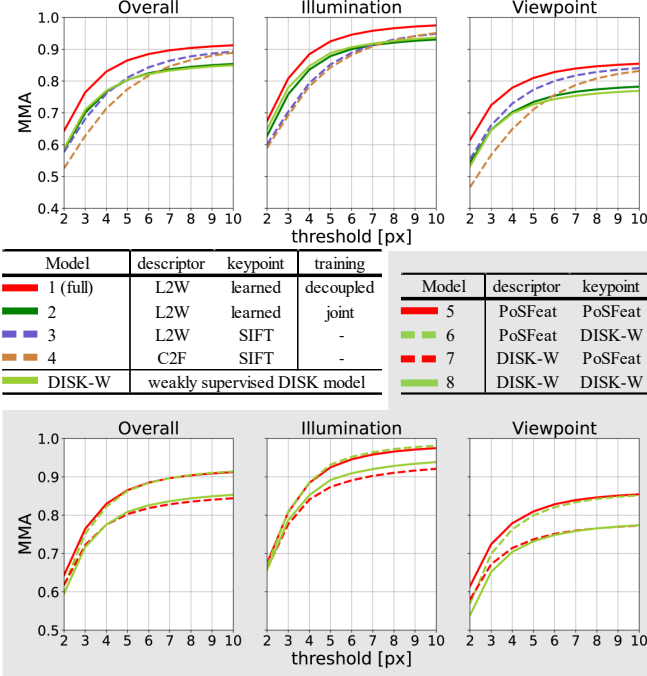
Figure 7. Ablation results on HPatches. "L2W" denotes our line-to-window search strategy (illustrated in Fig 4(b)) and "C2F" denotes the coarse-to-fine search strategy [46] (illustrated in Fig 4(a)). "learned" means that the keypoints are generated by a detection network and "SIFT" mean that SIFT keypoints (OpenCV default settings) are used. "decoupled" means the proposed decoupled training pipeline is adopted and 'joint' means the description network and the detection network are jointly optimized.

timized. Then, we developed Model 3 based on the *detect-then-describe* pipeline. Specifically, the description network is combined with SIFT keypoints in Model 3.

As shown in Fig. 7, with only weak supervision, the ambiguity during optimization limits the performance of **joint** *training describe-then-detect* approaches (Model 2 and DISK-W). Moreover, Model 2 is even inferior to Model 3 under viewpoint change. Compared to Models 2 and 3, Model 1 with our **decoupled** *training describe-then-detect* pipeline produces much higher accuracy. This clearly demonstrates that our **decoupled** *training describe-then-detect* pipeline is well suitable to weakly supervised learning to achieve superior performance.

We further test different combinations of keypoints and descriptors (Models 5-8). It can be observed that the improvement mainly comes from the descriptor, and the keypoints are slightly improved on the viewpoint change. Besides, we also illustrate the keypoints produced by our method and DISK-W in Fig. 1. DISK-W generates considerable inaccurate keypoints out of objects (*e.g.*, in the sky). In contrast, our model detects more reasonable keypoints. That is because those mismatched descriptors and

N_{line} \ w_{patch}	0.075	0.100	0.125
75	0.7703	0.7705	0.7666
100	0.7726	0.7748	0.7732
125	0.7732	0.7745	0.7744

Table 4. MMAscore achieved by our description network with different values of N_{line} and w_{patch} on the HPatches dataset.

erroneous keypoints produced from two different components do not influence each other within our **decoupled** *training describe-then-detect* pipeline.

Line-to-Window Search Strategy. To validate the effectiveness of our line-to-window search strategy, we developed a network variant (Model 4) by replacing our search strategy with a coarse-to-fine one (as proposed in [46], illustrated in Fig. 4(a)). For fair comparison with Model 3, SIFT keypoints are employed in this network variant. It can be observed that Model 3 outperforms Model 4 by significant margins. That is because, our line-to-window search strategy can make full use of the geometry information of camera poses to reduce the search space for accurate localization of correspondences. Consequently, higher accuracy can be achieved.

Number of Sampled Points N_{line} **and Window Size** w_{patch}**.** We conduct experiments to study the effects of N_{line} and w_{patch} during our line-to-window search. More sampled points and a large window size are beneficial to the performance at the expense of higher computational cost. To achieve a trade-off between performance and computational complexity, $w_{patch} = 0.100$ and $N_{line} = 100$ are used as the default setting.

5. Conclusion

In this paper, we introduce a **decoupled** *training describe-then-detect* pipeline tailored for weakly supervised local feature learning. Within our pipeline, the detection network is decoupled from the description network and postponed until discriminative and robust descriptors are obtained. In addition, we propose a line-to-window search strategy to explicitly use the camera pose information to reduce search space for better descriptor learning. Extensive experiments show that our method achieves the state-of-the-art performance on three different evaluation frameworks and significantly closes the gap between fully-supervised and weakly supervised methods.

References

[1] 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*, 54(10):105–112, 2011. 1

[2] Relja Arandjelović and Andrew Zisserman. Three Things Everyone Should Know to Improve Object Retrieval. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2911–2918, 2012. 5, 6, 7

[3] Vassileios Balntas, Karel Lenc, Andrea Vedaldi, and Krystian Mikolajczyk. HPatches: A Benchmark and Evaluation of Handcrafted and Learned Local Descriptors. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. 5, 6, 7

[4] Axel Barroso-Laguna, Edgar Riba, Daniel Ponsa, and Krystian Mikolajczyk. Key.Net: Keypoint Detection by Handcrafted and Learned CNN Filters. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 5836–5844, 2019. 2

[5] Herbert Bay, Tinne Tuytelaars, and Luc Van Gool. SURF: Speeded Up Robust Features. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 404–417, 2006. 1

[6] Aritra Bhowmik, Stefan Gumhold, Carsten Rother, and Eric Brachmann. Reinforced Feature Points: Optimizing Feature Detection and Description for a High-Level Task. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4948–4957, 2020. 7

[7] Michael Calonder, Vincent Lepetit, Mustafa Ozuysal, Tomasz Trzcinski, Christoph Strecha, and Pascal Fua. BRIEF: Computing a Local Binary Descriptor Very Fast. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 34(7):1281–1298, 2011. 1

[8] Luca Cavalli, Viktor Larsson, Martin Ralf Oswald, Torsten Sattler, and Marc Pollefeys. Handcrafted Outlier Detection Revisited. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 770–787, 2020. 7

[9] Hongkai Chen, Zixin Luo, Jiahui Zhang, Lei Zhou, Xuyang Bai, Zeyu Hu, Chiew-Lan Tai, and Long Quan. Learning to Match Features with Seeded Graph Matching Network. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 6301–6310, 2021. 3, 7

[10] Peter Hviid Christiansen, Mikkel Fly Kragh, Yury Brodskiy, and Henrik Karstoft. UnsuperPoint: End-to-end Unsupervised Interest Point Detector and Descriptor. *arXiv preprint arXiv:1907.04011*, 2019. 3

[11] Daniel DeTone, Tomasz Malisiewicz, and Andrew Rabinovich. SuperPoint: Self-Supervised Interest Point Detection and Description. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, pages 224–236, 2018. 1, 3, 6, 7

[12] 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 *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 1, 2, 5, 6, 7

[13] Patrick Ebel, Anastasiia Mishchuk, Kwang Moo Yi, Pascal Fua, and Eduard Trulls. Beyond Cartesian Representations for Local Descriptors. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 253–262, 2019. 2

[14] Steve R Gunn. Edge Detection Error in the Discrete Laplacian of Gaussian. In *Proceedings of the International Conference on Image Processing (ICIP)*, volume 2, pages 515–519. IEEE, 1998. 1

[15] Yulan Guo, Michael Choi, Kunhong Li, Farid Boussaid, and Mohammed Bennamoun. Soft exemplar highlighting for cross-view image-based geo-localization. *IEEE Transactions on Image Processing*, 2022. 1

[16] Chris Harris, Mike Stephens, et al. A Combined Corner and Edge Detector. In *Alvey vision conference*, volume 15, pages 147–151. Citeseer, 1988. 1

[17] 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. *International Journal of Computer Vision (IJCV)*, 129(2):517–547, 2021. 1

[18] Stefan Leutenegger, Margarita Chli, and Roland Y Siegwart. BRISK: Binary Robust Invariant Scalable Keypoints. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 2548–2555. IEEE, 2011. 1

[19] Xinghui Li, Kai Han, Shuda Li, and Victor Prisacariu. Dual-Resolution Correspondence Networks. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 33, 2020. 7

[20] Zhengqi Li and Noah Snavely. MegaDepth: Learning Single-View Depth Prediction From Internet Photos. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2041–2050, 2018. 5

[21] David G Lowe. Distinctive Image Features from Scale-Invariant Keypoints. *International Journal of Computer Vision (IJCV)*, 60(2):91–110, 2004. 1, 5, 6, 7

[22] Zixin Luo, Tianwei Shen, Lei Zhou, Jiahui Zhang, Yao Yao, Shiwei Li, Tian Fang, and Long Quan. ContextDesc: Local Descriptor Augmentation with Cross-Modality Context. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 1, 2, 5, 6

[23] 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 *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 1, 2, 6, 7

[24] Iaroslav Melekhov, Gabriel J Brostow, Juho Kannala, and Daniyar Turmukhambetov. Image Stylization for Robust Features. *arXiv preprint arXiv:2008.06959*, 2020. 7

[25] Krystian Mikolajczyk and Cordelia Schmid. Scale & Affine Invariant Interest Point Detectors. *International Journal of Computer Vision (IJCV)*, 60(1):63–86, 2004. 5

[26] Anastasiia Mishchuk, Dmytro Mishkin, Filip Radenovic, and Jiri Matas. Working Hard to Know Your Neighbor's Mar-

gins: Local Descriptor Learning Loss. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 30, 2017. 1, 2, 5, 6

[27] Dmytro Mishkin, Filip Radenovic, and Jiri Matas. Repeatability Is Not Enough: Learning Affine Regions via Discriminability. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 284–300, 2018. 5, 6

[28] Raul Mur-Artal and Juan D Tardós. ORB-SLAM2: An Open-Source SLAM System for Monocular, Stereo, and RGB-D Cameras. *IEEE Transactions on Robotics (TR)*, 33(5):1255–1262, 2017. 1

[29] Hyeonwoo Noh, Andre Araujo, Jack Sim, Tobias Weyand, and Bohyung Han. Large-Scale Image Retrieval With Attentive Deep Local Features. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 3456–3465, 2017. 2, 3, 6

[30] Yuki Ono, Eduard Trulls, Pascal Fua, and Kwang Moo Yi. LF-Net: Learning Local Features from Images. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2018. 1, 2

[31] Udit Singh Parihar, Aniket Gujarathi, Kinal Mehta, Satyajit Tourani, Sourav Garg, Michael Milford, and K Madhava Krishna. RoRD: Rotation-Robust Descriptors and Orthographic Views for Local Feature Matching. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2021. 3

[32] Rémi Pautrat, Viktor Larsson, Martin R Oswald, and Marc Pollefeys. Online Invariance Selection for Local Feature Descriptors. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 707–724, 2020. 7

[33] Jerome Revaud, Cesar De Souza, Martin Humenberger, and Philippe Weinzaepfel. R2D2: Reliable and Repeatable Detector and Descriptor. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 32, pages 12405–12415, 2019. 1, 2, 6, 7

[34] Ignacio Rocco, Relja Arandjelović, and Josef Sivic. Efficient neighbourhood consensus networks via submanifold sparse convolutions. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 605–621. Springer, 2020. 7

[35] Ethan Rublee, Vincent Rabaud, Kurt Konolige, and Gary Bradski. ORB: An Efficient Alternative to SIFT or SURF. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 2564–2571. IEEE, 2011. 1

[36] Paul-Edouard Sarlin, Daniel DeTone, Tomasz Malisiewicz, and Andrew Rabinovich. SuperGlue: Learning Feature Matching With Graph Neural Networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4938–4947, 2020. 3, 7

[37] Torsten Sattler, Bastian Leibe, and Leif Kobbelt. Improving Image-Based Localization by Active Correspondence Search. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 752–765, 2012. 1

[38] Nikolay Savinov, Akihito Seki, Lubor Ladicky, Torsten Sattler, and Marc Pollefeys. Quad-Networks: Unsupervised Learning to Rank for Interest Point Detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1822–1830, 2017. 2

[39] Johannes L Schonberger and Jan-Michael Frahm. Structure-from-Motion Revisited. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4104–4113, 2016. 1

[40] Johannes L Schonberger, Hans Hardmeier, Torsten Sattler, and Marc Pollefeys. Comparative Evaluation of Hand-Crafted and Learned Local Features. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1482–1491, 2017. 7

[41] Jiaming Sun, Zehong Shen, and Yuang Wang. LoFTR: Detector-Free Local Feature Matching With Transformers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 8922–8931, 2021. 3, 4, 7

[42] Ilya Sutskever, James Martens, George Dahl, and Geoffrey Hinton. On the Importance of Initialization And Momentum in Deep Learning. In *Proceedings of the International Conference on Machine Learning (ICML)*, pages 1139–1147, 2013. 5

[43] Yurun Tian, Xin Yu, Bin Fan, Fuchao Wu, Huub Heijnen, and Vassileios Balntas. SOSNet: Second Order Similarity Regularization for Local Descriptor Learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 11016–11025, 2019. 1, 2

[44] Carl Toft, Will Maddern, Akihiko Torii, Lars Hammarstrand, Erik Stenborg, Daniel Safari, Masatoshi Okutomi, Marc Pollefeys, Josef Sivic, Tomas Pajdla, et al. Long-Term Visual Localization Revisited. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 2020. 1

[45] Michał J. Tyszkiewicz, Pascal Fua, and Eduard Trulls. DISK: Learning Local Features with Policy Gradient. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 33, 2020. 1, 2, 3, 5, 6, 7

[46] Qianqian Wang, Xiaowei Zhou, Bharath Hariharan, and Noah Snavely. Learning Feature Descriptors Using Camera Pose Supervision. In *Proceedings of the European Conference on Computer Vision (ECCV)*, volume 12346, pages 757–774, 2020. 2, 3, 4, 5, 6, 7, 8

[47] Xupeng Wang, Ferdous Ahmed Sohel, Mohammed Bennamoun, Yulan Guo, and Hang Lei. Persistence-based interest point detection for 3D deformable surface. In *International Conference on Computer Graphics Theory and Applications (GRAPP)*, pages 58–69, 2017. 1

[48] Olivia Wiles, Sebastien Ehrhardt, and Andrew Zisserman. Co-Attention for Conditioned Image Matching. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 15920–15929, 2021. 7

[49] Kwang Moo Yi, Eduard Trulls, Vincent Lepetit, and Pascal Fua. LIFT: Learned Invariant Feature Transform. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 467–483, 2016. 1, 2

[50] Zichao Zhang, Torsten Sattler, and Davide Scaramuzza. Reference Pose Generation for Long-term Visual Localization via Learned Features and View Synthesis. *International Journal of Computer Vision (IJCV)*, 129(4):821–844, 2021. 1, 6, 7

[51] Yong Zhao, Shibiao Xu, Shuhui Bu, Hongkai Jiang, and Pengcheng Han. GSLAM: A General SLAM Framework and Benchmark. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 1110–1120, 2019. 1

[52] Qunjie Zhou, Torsten Sattler, and Laura Leal-Taixe. Patch2pix: Epipolar-guided pixel-level correspondences. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4669–4678, 2021. 3, 7