

---

This is an electronic reprint of the original article.  
This reprint may differ from the original in pagination and typographic detail.

Melekhov, Iaroslav; Laskar, Zakaria; Li, Xiaotian; Wang, Shuzhe; Kannala, Juho  
**Digging Into Self-Supervised Learning of Feature Descriptors**

*Published in:*  
Proceedings - 2021 International Conference on 3D Vision, 3DV 2021

*DOI:*  
[10.1109/3DV53792.2021.00122](https://doi.org/10.1109/3DV53792.2021.00122)

Published: 01/01/2022

*Document Version*  
Peer-reviewed accepted author manuscript, also known as Final accepted manuscript or Post-print

*Please cite the original version:*  
Melekhov, I., Laskar, Z., Li, X., Wang, S., & Kannala, J. (2022). Digging Into Self-Supervised Learning of Feature Descriptors. In *Proceedings - 2021 International Conference on 3D Vision, 3DV 2021* (pp. 1144-1155). (International Conference on 3D Vision proceedings). IEEE. <https://doi.org/10.1109/3DV53792.2021.00122>

# Digging Into Self-Supervised Learning of Feature Descriptors

Iaroslav Melekhov\*

Zakaria Laskar\*

Xiaotian Li

Shuzhe Wang

Juho Kannala

Aalto University

## Abstract

*Fully-supervised CNN-based approaches for learning local image descriptors have shown remarkable results in a wide range of geometric tasks. However, most of them require per-pixel ground-truth keypoint correspondence data which is difficult to acquire at scale. To address this challenge, recent weakly- and self-supervised methods can learn feature descriptors from relative camera poses or using only synthetic rigid transformations such as homographies. In this work, we focus on understanding the limitations of existing self-supervised approaches and propose a set of improvements that combined lead to powerful feature descriptors. We show that increasing the search space from in-pair to in-batch for hard negative mining brings consistent improvement. To enhance the discriminativeness of feature descriptors, we propose a coarse-to-fine method for mining local hard negatives from a wider search space by using global visual image descriptors. We demonstrate that a combination of synthetic homography transformation, color augmentation, and photorealistic image stylization produces useful representations that are viewpoint and illumination invariant. The feature descriptors learned by the proposed approach perform competitively and surpass their fully- and weakly-supervised counterparts on various geometric benchmarks such as image-based localization, sparse feature matching, and image retrieval.*

## 1. Introduction

Many geometric computer vision tasks require robust estimation of local descriptors such as image alignment [5], structure-from-motion [28, 64], image retrieval [24, 56]. With the advent of deep learning models, robust local descriptors can be learnt with highly over-parameterized networks such as Convolutional Neural Networks (CNN). Consequently, much of the latest works [6, 17, 20, 44, 57] have since focused on designing appropriate datasets, labels and loss functions to learn useful local representations.

It is well-known that CNNs are data demanding models that require large amount of data with labelled supervision. While in many cases large amount of data can be obtained, structuring it in meaningful partitions or labels is a costly process. To counter this problem, existing approaches [24, 55] use SfM pipelines to generate labels. Despite the effectiveness, the process incurs significant computational burden. Furthermore, integrating new stream of data to the existing dataset requires additional iterations of the costly spatial descriptor matching process. The paradigm of self-supervised learning attempts to provide a solution by enabling label-free training of CNNs. Most self-supervised methods focus on auxiliary pretext tasks [18, 49, 83] such as image reconstruction – generative modelling, or discriminative optimization of the latent space based on contrastive learning [3, 19]. The key requirement is to learn transferable representations that can adapt to downstream tasks with limited supervision. Recent progress [12, 14, 25, 31, 58] in the image classification domain shows the crucial role of positives (similar) and negatives (dissimilar) in the contrastive loss function to learn robust representations. In this paper, we delve deeper into these factors for the task of local image descriptors learning.

Existing self-supervised methods for learning local descriptors [13, 17] have achieved some success by mining positives and negatives from homography related image pairs (in-pair sampling). On the other hand, *supervised* local descriptor learning methods such as HardNet [44] show consistent improvement with in-batch negative sampling. In addition, image retrieval methods [24, 55] have demonstrated remarkable results by mining global negative descriptors from a large database of thousands of images. In this paper, we propose a scalable method for sampling hard-negative local features from a wide search space that operates in a coarse-to-fine fashion by leveraging image descriptors.

Learning keypoint descriptors that are invariant to illumination changes is a challenging task since it is not easy to acquire suitable data. Synthetic color augmentation (CA) which is widely used to improve model robustness and performance in many computer vision applications [21, 33, 46, 77] can enhance illumination invariance

\*indicates equal contribution. Correspondence to: iaroslav.melekhov@aalto.fi

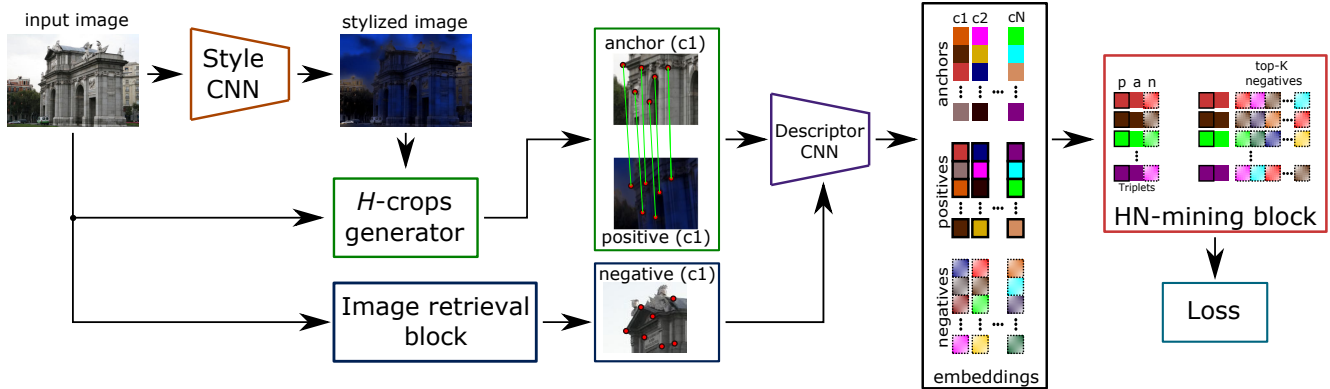


Figure 1: **Overview of the proposed pipeline for learning local image descriptors.** Given an image and a set of interest points (top-left), we create a stylized copy of the image with a forward-pass through Style CNN. A pair of homography related crops ( $c_i$ ) is randomly sampled from the original and stylized image in H-crops generator. We pass the image crops through Descriptor CNN to obtain representation of corresponding interest points in the feature space (embeddings). For each  $a_i$  and  $p_i$  descriptor from positive pair (anchors-positives), we create a list of hard-negative descriptors (shown with color gradient) using the proposed coarse-to-fine sampling procedure by leveraging global visual descriptors (Image retrieval block). Finally, the ranking loss is utilized to optimize Descriptor CNN. See Sec. 3 for a detailed description.

of local descriptors only to a certain extent. Following [57], we propose to utilize photorealistic image stylization [37] at training time. In contrast to [57], the more diverse style images are considered which help to increase appearance variations of our training data. We combine these findings to train existing state-of-the-art local descriptor models without any labels. During training and evaluation, only a subset of local descriptors from dense descriptors produced by the network is used. The selection is based on local keypoint detectors such as SuperPoint [17] providing efficient training iterations – by limiting hard-negative mining only to a subset of local descriptors. The proposed method is compared with strong supervised and weakly-supervised baselines [57, 76] on a wide range of benchmarks such as visual localization, sparse feature matching, and image retrieval. The key findings are analyzed and complemented by in-depth ablation study.

In summary, this work makes the following contributions: (1) we demonstrate that improving the quality of hard negative samples during training leads to more powerful and discriminative local keypoint descriptors; (2) by leveraging visual image descriptors, we propose a coarse-to-fine method for mining hard negative local descriptors at scale; (3) we investigate several CA techniques to improve the robustness of learned feature descriptors to illumination changes and show that photorealistic image stylizations yield best performance. We demonstrate that, together, these contributions largely close the gap between self-supervised and fully-supervised approaches for feature descriptor learning.

## 2. Related work

**Local image descriptors.** Due to their generalization ability, traditional local features [8, 11, 39, 60] are still popular and widely used by many existing systems where establishing pixel correspondences is essential [34, 48, 64, 65]. The recent success of deep learning in various computer vision applications has motivated researchers to work on learned CNN-based approaches for interest point detection [7, 36, 45, 63, 74, 84], description [6, 26, 40, 44, 69, 70] or both detection and description [17, 20, 22, 41, 50, 57, 75, 82]. Although, those approaches have demonstrated improved performance over classical hand-crafted methods on challenging benchmarks such as vision localization and image matching, most of them require ground-truth pixel correspondences between two views [20, 30, 57, 66]. Such accurate ground-truth labels are quite tedious to obtain since it requires information about the 3D representation of the scene (e.g. 3D reconstruction, dense depth maps or scene flow). To address this issue, recent research directions have shifted towards weakly- and self-supervised methods. These methods rely on the supervision from epipolar geometry [76, 78, 81, 86], principles of Reinforcement Learning [9, 72, 73], relative camera poses [9, 76], camera intrinsics [79], and rigid synthetic 2D transformations [13, 17, 57, 68]. In this work, we show that with careful choices regarding hard-negative mining, color augmentation and photorealistic image stylization, we can achieve superior performance to supervised methods using only synthetic homographies. In principle, the proposed ideas can be also applied to supervised CNN-based approaches.

**Illumination invariance in feature descriptors.** Learn-

ing features that are invariant to illumination changes is a challenging task. To increase robustness to illumination changes, methods like [20, 76, 79, 80] utilize a large SfM database of images with different lighting conditions. Domain adaptation-based methods [1, 47, 52] utilize GANs [23, 87], view synthesis [47] and few-shot learning [4] to reduce the domain gap between the training and the test (target) image distributions. Leveraging photorealistic image stylization to improve robustness to day-night variations was introduced in [57]. In the follow-up work, Melekhov *et al.* [42] extend the original method [57] by considering more diverse style categories. Pautrat *et al.* [51] propose a model predicting several dense descriptors with different levels of invariance. For a particular input image, the most relevant invariance is automatically selected using the concept of meta-descriptors based on a NetVLAD [2] layer. Inspired by [42, 57], we utilize the idea of image stylization to improve the robustness of local descriptors to appearance variations.

**Self-supervised learning.** Self-supervised learning paves the way for training models on large amount of unlabelled data. Recent methods [12, 25] in the classification domain are based on contrastive loss where representations from augmented copies of the original image – positives, are contrasted against representations from other images in the batch or memory bank – negatives. The quality of negatives was analyzed by several methods [14, 31, 58] in improving self-supervised image representation learning. For local descriptors, [35, 43, 59] consider only the augmented input image pair to sample positive and negative local descriptors (in-pair sampling). In this work, we propose methods to improve the quality of both positives and negatives.

### 3. Method

In this section, we describe our method for local image descriptors learning that takes a single RGB image and produces a set of local feature descriptors in a self-supervised manner. Supervised methods leverage known scene geometry and utilize image- and keypoint level correspondences to increase robustness and invariance of local features to viewpoint and appearance variations. However, obtaining such ground-truth correspondence might be difficult especially for large-scale scenes. We analyze and carefully tailor improvements to existing self-supervised methods for learning invariant local descriptors.

We first review the main components of the proposed approach and the key ideas behind it, then we describe the hard-negative sampling procedure and training loss. An overview of our method is schematically illustrated in Figure 1. Given an input image  $I_i \in \mathbb{R}^{H \times W \times 3}$  and a set of interest points  $\mathbf{k}_i$ , our goal is to predict local descriptors  $\mathbf{d}_i$  corresponding to index  $i$  which would be robust to appear-

ance changes and have good generalization performance. To obtain interest points, we use SuperPoint [17] detector.

**Illumination Invariance** Illumination invariance can be seen as learning context independent representations. In classification domain, various forms of CA methods such as RandAugment [15] aim to remove the contextual bias in image representations. However, these methods apply various illumination transformations uniformly over both salient foreground object and background context which might be troublesome for local descriptor learning. Following [42, 57] we rely on image stylization [37] to improve the robustness of local features under drastic illumination variations. Li *et al.* [37] propose a closed-form solution to image stylization based on a deep learning model trained in a self-supervised fashion by minimizing the sum of the reconstruction loss and perceptual loss. The method photorealistically transfers style of a reference photo to a content photo preserving local scene geometry. We use this pre-trained model as a part of our pipeline (**Style CNN** in Figure 1) and manually select reference images representing two style categories, *i.e.* night and dusk, from the contributed views of the Amos Patches dataset [53]. In contrast to [57], for each category we consider *multiple* style images which are then randomly applied to each image in the training set. The style categories and image stylization results are provided in the supplementary material.

**Viewpoint Invariance.** To model viewpoint invariance, recent self-supervised methods [13, 17] apply a known homography transformation  $\mathbf{H}$ , *i.e.*  $\hat{I} = \mathbf{H}(I)$  followed by random crop. As in [17], we use similar types of the homographic adaptation operations: translation, rotation, perspective transform, and scale whose parameters are sampled from a uniform distribution during training (**H-crops generator**). In addition to stylization, we apply synthetic CA independently on the anchor and positive crops by adding blur, random per-pixel Gaussian noise, adaptive histogram equalization (CLAHE) along with augmentation in brightness, contrast, hue, and saturation. The full list of augmentations is also presented in supplementary.

**Descriptor CNN.** We use a fully-convolutional neural network architecture which takes an input crop  $C_i \in \mathbb{R}^{H_c \times W_c \times 3}$  and outputs a dense map of L2-normalized fixed length descriptors, *i.e.*  $\mathbb{R}^{H_c \times W_c \times d}$ . In this work, we consider two strong baseline methods (supervised and weakly-supervised), namely R2D2 [57] and CAPS [76]. Sec. 4.1 provides an overview of these baselines in more detail. In order to perform a fair comparison, we directly use R2D2 and CAPS models as a structure of Descriptor CNN.

**Hard-negative (HN) mining block.** Supervised metric learning requires both positive and negative samples to learn meaningful representations. Without both positives and negatives, the model will collapse by converging to 0-valued

representations. Existing unsupervised approaches [17, 57] for learning local descriptors sample positives and negatives from a homography related image pair  $(I, \hat{I})$  with local descriptors  $A \in I, P \in \hat{I}$  respectively. The images in the pair are dubbed as the anchor ( $I$ ) and to the positive ( $\hat{I}$ ). Each local descriptor  $a_i \in A$  is associated with one  $p_i \in P$  due to known homography matrix  $\mathbf{H}$ . Next, for each matching pair  $(a_i, p_i)$  the index  $n$  of the non-matching descriptor  $p_n$  is obtained as follows:

$$n = \underset{n=1..M, n \neq i}{\operatorname{argmax}} s(a_i, p_n), \quad (1)$$

where  $M$  is the number of interest points in  $I$ ;  $s(\cdot)$  is a similarity function such as a dot-product. We refer to this baseline of selecting negatives as in-pair sampling. Several works [24, 44, 55] have shown that the quality of negatives has a strong impact in supervised learning of robust representations. Inspired by [44], we propose in-batch hard negatives sampling strategy that searches non-matching descriptors over all interest point embeddings  $p \in B \setminus I$ , where  $B \sim D$  is a mini-batch sampled from the dataset  $D$ .

Finding the local hard negatives is an exhaustive process. Sampling negatives over a search space  $D_n \sim D$  which is much larger than  $B$ ,  $|D_n| \gg |B|$  ideally should improve the quality of negatives but incurs significant computational cost. Rather than mining hard-negatives at the local descriptor level, we propose a coarse-to-fine search method (**Image retrieval block** in Figure 1). For each image  $I \in B$  we first find its nearest neighbor  $I^n \in D_n$  in the feature space using visual descriptor  $g_n$ . This results in a mini-batch  $B_n$  consisting of hard-negative image samples (the images look similar to the anchors  $I \in B$  but *not* homography related as the positives), *i.e.*,  $B_n = \{I_i^n\}_{i=1..|B|}$ , which are then used to mine non-matching feature descriptors at the local level. Given the batch,  $B_n$  the local descriptor pool now consists of  $p_n \in B' \setminus I$  where  $B' = (B \cup B_n)$ . Such coarse-to-fine strategy leveraging visual image descriptors makes the process of mining hard-negative local descriptors tractable. Computing visual descriptors can be formulated in two ways. On the one hand, one can use state of the art image retrieval models, such as [55, 56], trained on supervised data generated by SfM systems to extract global descriptor from the entire image. On the other hand we propose a completely self-supervised approach whereby the global descriptors are constructed from the set of local descriptors extracted using the current network state. These local descriptors are then aggregated by some mean operation, such as GeM [55]. Specifically, given a dense feature map  $X_f \in \mathbb{R}^{H_c \times W_c \times N_d}$ , visual descriptor  $g_f \in \mathbb{R}^{N_d}$  is obtained by summing  $l2$  normalized  $X_f$  over the spatial resolution. The final representation is again  $l2$  normalized.

**Loss functions.** Similar to [26, 57], we formulate our problem as the optimization of a differentiable approximation

( $AP'$ ) of the Average Precision (AP) ranking metric at training time:

$$\mathcal{L}_{AP} = \frac{1}{N_a} \sum_i \left( 1 - AP' \left( a_i, p_i, \{p_{n_k}^i\}_{k=1..K} \right) \right) \quad (2)$$

where  $N_a$  is the number of anchor descriptors in a mini-batch  $B'$ . In contrast to image matching [26], the number of positive and negative samples is highly imbalanced in our setting. Indeed, for each anchor descriptor in  $B'$  there is only one positive sample and several hundred negatives. To address this issue, we propose to select *top-K* hard-negative local descriptors  $\{p_{n_k}^i\}_{k=1..K}$  for  $a_i$ . Results presented in Sec. 4.4 demonstrate consistent improvement over the original AP loss [26, 57]. For Triplet loss, the proposed idea of intra-batch hard-negative mining also demonstrates better performance compared to in-pair sampling (see supplementary for more details).

## 4. Experiments

Here, we describe the experimental setting and validate that (1) increasing the search space from within-pair to intra-batch for hard negative mining improves the discriminativeness of local image descriptors and leads to better performance, (2) photorealistic image stylization in addition to color augmentations demonstrates good generalization performance and improves results, especially when query images are taken during challenging illumination conditions, (3) self-supervised methods for local image descriptors learning achieve competitive results with their fully- and weakly-supervised counterparts. In order to highlight the effectiveness of the proposed ideas, we further present an extensive ablation study.

### 4.1. Baselines

In this work, we consider two recent approaches, R2D2 [57] and CAPS [76], as the baseline methods. Since we focus exclusively on local image descriptors learning, SuperPoint [17] keypoint detector is used in all our experiments and the baselines.

**R2D2** [57] can directly extract dense descriptors and a heatmap of keypoints from an input image. It consists of L2-Net backbone network with two heads for jointly learning repeatable and reliable matches. The network was trained on pairs of image crops with known pixel correspondences based on either synthetic homography or 3D reprojection. We consider the model trained with 3D supervision as our baseline.

**CAPS** [76] is a CNN-based local image descriptor trained in a weakly-supervised manner. Specifically, it can learn feature descriptors using only scene labels and relative camera poses between views. It utilizes SuperPoint keypoints






Figure 2: **Mean matching accuracy (MMA) on the HPatches [5] dataset.** For each method, the MMA with different pixel error thresholds is reported. The proposed self-supervised models (R2D2-U and CAPS-U) achieve better overall performance for stricter matching thresholds (up to 3 pixels) compared to their supervised (R2D2-S and CAPS-S) counterparts. To avoid image clutter, for R2D2-U and CAPS-U, we only provide the results obtained by the (▣, ▤) models.

and an epipolar constraint on pixel locations of matched points as a supervision signal. The resulting descriptor consists of two concatenated 128-dimensional feature vectors, *i.e.* coarse- and fine-level descriptors, extracted at different parts of the model encoder (an ImageNet-pretrained ResNet-50 architecture [16, 27]). In our experiments, we utilize *only* fine-level descriptors to optimize the proposed criterion (Eq. 2).

## 4.2. Notations

We use the following symbolic notation throughout the paper: (▣) denotes models trained with synthetic CA; (▤) indicates photorealistic stylized images; (▣, ▤) if both CA and style transfer are used during training. Finally, we refer the models utilizing image retrieval pre-trained networks for extracting visual descriptors with a suffix  $\text{*-gd}$  and  $\text{*-selfgd}$  denotes the self-supervised way of computing visual descriptors on the fly using the current state of the proposed Descriptor CNN (*c.f.* Figure 1).

## 4.3. Benchmarks and metrics

We verify the proposed ideas on the following benchmarks.

**Sparse feature matching.** For this experiment, we evaluate our descriptors on the full images provided by the HPatches dataset [5]. The dataset consists of 116 image sequences with varying photometric and viewpoint changes. Each sequence contains a reference image and 5 source images taken under a different viewpoint. For all image pairs, the estimated ground-truth homography matrix with respect to the reference image is provided. We follow the standard evaluation protocol [20] and use the mean matching accuracy (MMA) as the metric, *i.e.* the average percentage of

correct matches per image under a certain pixel threshold. In addition to MMA, we also report a homography estimation score [17, 20], precision and recall computed for a 3-pixel error threshold.

**Image-based localization.** To verify the generalization performance of the proposed approach, we evaluate our method on indoor and outdoor visual localization benchmarks. Specifically, we consider the large-scale *outdoor* Aachen Day-Night v1.1 [85] and Tokyo24/7 [71] datasets for localization under severe illumination conditions. Image-based localization in complex *indoor* environments is challenging due to the large viewpoint changes and weakly textured scenes. We use InLoc [67], a large-scale indoor dataset with strong appearance changes between query images and the reference 3D map. The localization performance on the Aachen Day-Night v1.1 and InLoc benchmarks is reported as the percentage of correctly localized queries. For InLoc, we follow the localization pipeline<sup>1</sup> proposed in [61, 62]. The query is successfully localized if its camera position and orientation are both within a certain threshold.

The performance on the Tokyo24/7 dataset is evaluated using Recall@N, which is the number of queries that are correctly localized given  $N$  nearest neighbor database images. Specifically, for a given query we first obtain a ranked list of database images,  $L$  based on Euclidean distance between their global NetVLAD representations. The top-100 ranked database images,  $L' \in L$  are re-ranked according to their similarity score based on the number of geometrically verified inliers between the query and each database image.

**Image retrieval.** We evaluate the proposed local image descriptors on the image retrieval benchmark. Specifi-

| Method             | Illumination |              |              | Viewpoint    |              |              |
|--------------------|--------------|--------------|--------------|--------------|--------------|--------------|
|                    | H            | Precision    | Recall       | H            | Precision    | Recall       |
| Root SIFT [2]      | 0.933        | <b>0.782</b> | 0.799        | 0.566        | 0.651        | 0.527        |
| HardNet [44]       | 0.940        | 0.702        | 0.731        | 0.664        | 0.701        | 0.734        |
| SOSNet [70]        | 0.933        | 0.748        | 0.821        | <b>0.698</b> | 0.727        | 0.760        |
| SuperPoint [17]    | 0.912        | 0.710        | 0.811        | 0.671        | 0.685        | 0.750        |
| D2-Net [20]        | 0.905        | 0.725        | 0.775        | 0.617        | 0.666        | 0.664        |
| LISRD [51]         | <b>0.947</b> | 0.766        | 0.920        | 0.688        | 0.731        | 0.757        |
| R2D2 [57]          | 0.940        | 0.762        | 0.837        | 0.692        | 0.720        | 0.732        |
| CAPS [76]          | 0.888        | 0.757        | <u>0.938</u> | <u>0.692</u> | 0.723        | 0.699        |
| R2D2-(▣, ▤)        | 0.944        | 0.764        | <u>0.838</u> | 0.678        | <u>0.732</u> | <u>0.739</u> |
| R2D2-(▣, ▤)-selfgd | 0.933        | 0.761        | 0.817        | 0.678        | 0.715        | 0.705        |
| R2D2-(▣, ▤)-gd     | <b>0.947</b> | <u>0.766</u> | 0.826        | <b>0.698</b> | 0.726        | 0.720        |
| CAPS-(▣, ▤)        | 0.933        | 0.750        | 0.884        | 0.671        | 0.742        | 0.728        |
| CAPS-(▣, ▤)-selfgd | <u>0.937</u> | 0.756        | 0.893        | 0.661        | 0.740        | 0.752        |
| CAPS-(▣, ▤)-gd     | 0.919        | <u>0.757</u> | 0.890        | 0.681        | <u>0.747</u> | <u>0.762</u> |

Table 1: **Evaluation results on HPatches.**

We follow [51] and report the accuracy of homography estimation, precision and recall for error thresholds of 3 pixels. The best score between supervised baselines, **R2D2** and **CAPS**, and the proposed methods for each category is underlined. The overall best score is in **bold**.

<sup>1</sup><https://github.com/cvg/Hierarchical-Localization>















Figure 3: **Qualitative matching results.** We evaluate the proposed unsupervised pipeline with two different backbone models, R2D2 (top row) and CAPS (bottom row) on test samples of the MegaDepth dataset [38]. Both models have been trained on stylized images with color augmentations using in-batch hard negative sampling. Green lines indicate correspondences estimated by RANSAC (best viewed digitally). The proposed method predicts reliable matches even without training on annotated correspondences.

|                 | Method  | Supervision | Training data | Aachen v1.1<br>% localized queries |          |         |                    |          |         | InLoc<br>% localized queries |          |         |           |          |         |
|-----------------|---|-------------|---------------|------------------------------------|----------|---------|--------------------|----------|---------|------------------------------|----------|---------|-----------|----------|---------|
|                 |   |             |               | Day (824 images)                   |          |         | Night (191 images) |          |         | DUC1                         |          |         | DUC2      |          |         |
|                 |   |             |               | 0.25m, 2°                          | 0.5m, 5° | 5m, 10° | 0.25m, 2°          | 0.5m, 5° | 5m, 10° | 0.25m, 2°                    | 0.5m, 5° | 5m, 10° | 0.25m, 2° | 0.5m, 5° | 5m, 10° |
| Super           | R2D2 [57]   | OF          | A+R           | 88.6                               | 95.4     | 98.9    | 72.8               | 89.0     | 97.4    | 27.8                         | 42.4     | 54.5    | 22.1      | 34.4     | 42.7    |
|                 | R2D2 <sup>+</sup>   | OF          | A             | 87.7                               | 94.7     | 98.7    | 69.6               | 86.4     | 95.3    | 29.8                         | 43.4     | 55.1    | 21.4      | 34.4     | 43.5    |
|                 | CAPS [76]   | SL+RP       | M             | 85.3                               | 93.8     | 97.9    | 75.9               | 88.5     | 97.9    | 38.4                         | 59.1     | 74.7    | 34.4      | 52.7     | 61.8    |
|                 |   |             |               |                                    |          |         |                    |          |         |                              |          |         |           |          |         |
| Self-supervised | R2D2-  | -           | A             | 87.4                               | 94.9     | 98.3    | 63.9               | 80.1     | 92.1    | 29.3                         | 46.5     | 57.6    | 29.8      | 42.0     | 51.1    |
|                 | R2D2-  | -           | A             | 88.0                               | 94.8     | 98.2    | 70.2               | 86.4     | 95.8    | 27.8                         | 46.5     | 57.6    | 29.8      | 42.0     | 51.1    |
|                 | R2D2-  | -           | M             | 87.4                               | 94.7     | 98.3    | 72.3               | 88.5     | 97.4    | 26.8                         | 46.0     | 57.1    | 26.7      | 42.7     | 48.9    |
|                 | R2D2-  | -           | M             | 87.5                               | 94.9     | 98.3    | 71.7               | 86.4     | 96.9    | 27.8                         | 43.4     | 54.5    | 28.2      | 41.2     | 51.9    |
|                 | R2D2-  | -           | M             | 88.1                               | 94.8     | 98.1    | 71.2               | 88.0     | 95.8    | 30.8                         | 43.9     | 54.0    | 28.2      | 37.4     | 46.6    |
|                 | R2D2-  | -           | M+P           | 88.2                               | 95.1     | 98.5    | 73.3               | 90.1     | 97.4    | 29.3                         | 47.0     | 58.1    | 25.2      | 41.2     | 51.9    |
|                 | CAPS-  | -           | M             | 85.8                               | 93.8     | 98.2    | 67.0               | 82.2     | 96.9    | 35.9                         | 54.0     | 69.2    | 38.9      | 50.4     | 64.9    |
|                 | CAPS-  | -           | M             | 85.1                               | 93.2     | 97.8    | 71.7               | 87.4     | 97.9    | 39.9                         | 61.6     | 70.7    | 38.2      | 51.1     | 66.4    |
|                 | CAPS-  | -           | M             | 87.0                               | 93.8     | 98.3    | 73.8               | 89.0     | 97.4    | 35.9                         | 53.0     | 65.7    | 32.8      | 47.3     | 61.1    |
|                 | CAPS-  | -           | M             | 86.9                               | 93.8     | 98.1    | 71.7               | 89.0     | 97.4    | 35.9                         | 53.5     | 64.6    | 30.5      | 45.0     | 58.8    |
|                 | CAPS-  | -           | M+P           | 85.4                               | 93.2     | 97.9    | 72.3               | 88.5     | 97.9    | 38.9                         | 56.1     | 69.2    | 40.5      | 54.2     | 66.4    |

(a) Evaluations results for Aachen v1.1 [85] and InLoc [67]

| Method                    | Recall              |       |       |       |
|---------------------------|---------------------|-------|-------|-------|
|                           | r@1                 | r@5   | r@10  |       |
| DenseVLAD [71]            | 67.10               | 74.20 | 76.10 |       |
| NetVLAD-TokyoTM [2]       | 71.10               | 83.10 | 86.20 |       |
| SIFT [39]                 | 73.33               | 80.00 | 84.40 |       |
| InLoc [67]                | 62.54               | 67.62 | 70.48 |       |
| Dense Pixel Matching [35] | 77.14               | 84.44 | 86.67 |       |
| R2D2-[A+R] [57]           | 78.10               | 82.54 | 84.13 |       |
| CAPS-[M] [76]             | 83.49               | 86.67 | 87.94 |       |
| Self-supervised           | R2D2-[A]-[A]        | 79.05 | 83.17 | 85.71 |
|                           | R2D2-[A]-[I]        | 79.37 | 84.76 | 85.71 |
|                           | R2D2-[A]-selfgd-[M] | 78.41 | 82.54 | 83.81 |
|                           | R2D2-[A]-gd-[M]     | 77.78 | 82.86 | 85.40 |
|                           | R2D2-[A]-[M+P]      | 79.68 | 84.44 | 86.03 |
|                           | CAPS-[A]-[M]        | 82.54 | 86.03 | 87.94 |
|                           | CAPS-[A]-selfgd-[M] | 81.27 | 85.71 | 86.98 |
|                           | CAPS-[A]-gd-[M]     | 82.54 | 86.35 | 87.94 |
|                           | CAPS-[A]-[M+P]      | 82.86 | 87.30 | 88.25 |

(b) Performance on Tokyo24/7 [71]

Table 2: **Indoor and outdoor localization performance.** The supervised baseline methods, **R2D2** and **CAPS**, are color-coded. Best results in each category for each localization benchmark are in **bold**. **Legend:** *Supervision:* OF – ground-truth pixel correspondences based on optical flow [57]; SL – ground-truth scene labels; RP – relative camera poses. *Training data:* A – Aachen database images; R – Random images from the Internet [57]; M – MegaDepth dataset; P – Phototourism dataset. For Tokyo24/7, the training datasets for our supervised baselines and the proposed approach are given in [.] parentheses.

cally, the revisited Oxford (ROxford5k) and revisited Paris (RParis6k) datasets proposed by Radenović *et al.* [54] have been used. Similarly to the Tokyo24/7 evaluation protocol, for each query image we first create a ranked list of database images based on cosine similarity of their global representations in the feature space and then compute 2D-2D correspondences between query and top-100 database images which are then re-ranked based on the number of inliers verified by RANSAC. To compute global descriptors, the model "retrievalSfM120k-resnet101-gem" proposed in [55] is used. We evaluate performance by mean average precision (mAP) and mean precision at 1, 5, and 10 (mP@k), as defined by the protocol [54].

#### 4.4. Results

First, we experimentally compare the proposed models with their supervised oracles [57, 76] on the **sparse feature matching** benchmark by calculating MMA for different pixel thresholds illustrated in Figure 2. Following D2-Net [20], we separately report results for scenes cor-

responding to illumination changes (Illumination), view-point changes (Viewpoint), and for the whole dataset (Overall). The proposed unsupervised version of R2D2 (R2D2-U) shows a clear improvement over its supervised counterpart (R2D2-S) achieving better results for all evaluation thresholds. At the same time, although CAPS-U demonstrates competitive overall performance, it falls slightly behind its supervised baseline (CAPS-S [76]) on the illumination benchmark at larger pixel thresholds (>3px). Our hypothesis is that the original CAPS [76] model that leverages coarse-to-fine architecture (*c.f.* Sec. 4.3) allows to learn more discriminative descriptors leading to better matching performance. We also report homography estimation accuracy, precision (the percentage of correct matches over all the predicted matches) and recall (the ratio of correct matches over the total number of ground truth correspondences) using the same HPatches dataset. Following [51], we use mutual nearest neighbor matcher and a threshold of 3 pixels to consider a match to be correct. The results are summarized in Table 1. As can be seen, the proposed unsu-


























|                 | Method  | Training data | Hard-negative mining | ROxford5k    |              |              |                 |              |              |              |              |              |              |              |              | RParis6k        |              |              |            |              |              |            |            |              |              |              |              |
|-----------------|---|---------------|----------------------|--------------|--------------|--------------|-----------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|-----------------|--------------|--------------|------------|--------------|--------------|------------|------------|--------------|--------------|--------------|--------------|
|                 |   |               |                      | mAP          |              |              | mP@k [1, 5, 10] |              |              |              |              |              | mAP          |              |              | mP@k [1, 5, 10] |              |              |            |              |              |            |            |              |              |              |              |
| Supervised      | Baseline [55]   | -             | -                    | 72.60        | 54.89        | 27.37        | 91.43           | 80.74        | 74.12        | 91.43        | 80.57        | 73.86        | 61.43        | 47.50        | 39.79        | 86.15           | 68.98        | 43.86        | 98.57      | 95.14        | 93.43        | 98.57      | 98.57      | 96.14        | 95.71        | 87.71        | 82.86        |
|                 | R2D2 [57]   | A+R           | in-pair              | 76.39        | <b>59.09</b> | <b>33.20</b> | 97.06           | 88.38        | 81.55        | 95.71        | 89.71        | 83.57        | <b>85.71</b> | <b>61.14</b> | <b>47.14</b> | 85.73           | 68.51        | 42.86        | <b>100</b> | 97.24        | 95.81        | <b>100</b> | <b>100</b> | <b>99.14</b> | <b>100</b>   | <b>92.57</b> | 81.57        |
|                 | R2D2*   | A             | in-pair              | 75.36        | 58.61        | 32.69        | 97.06           | 87.35        | 79.56        | 95.71        | 89.43        | 82.57        | 85.71        | 59.14        | 45.86        | 85.39           | 68.36        | 42.93        | <b>100</b> | 97.43        | 96.14        | <b>100</b> | 99.71      | 98.86        | 98.57        | 91.14        | 82.43        |
|                 | CAPS [76]   | M             | -                    | 75.27        | <b>59.67</b> | <b>34.77</b> | 94.12           | 86.76        | 80.44        | 94.29        | 89.14        | 84.14        | <b>85.71</b> | <b>66.00</b> | <b>50.57</b> | <b>88.92</b>    | <b>69.87</b> | <b>44.47</b> | <b>100</b> | <b>99.43</b> | <b>97.43</b> | <b>100</b> | <b>100</b> | <b>100</b>   | <b>98.57</b> | <b>94.86</b> | <b>87.14</b> |
|                 | R2D2-(  )        | M             | in-pair              | 70.99        | 55.38        | 29.67        | 95.59           | 84.12        | 75.44        | 94.29        | 86.57        | 78.43        | 82.86        | 54.57        | 42.29        | 84.57           | 67.95        | 42.57        | <b>100</b> | 97.14        | 96.00        | <b>100</b> | 99.14      | 96.00        | 100          | 99.29        | 81.86        |
| Self-supervised | R2D2-(  )        | A             | in-batch             | 73.16        | 56.47        | 30.05        | 95.59           | 85.51        | 77.87        | 94.29        | 87.14        | 79.29        | 82.86        | 55.43        | 42.57        | 86.09           | <b>68.60</b> | <b>43.04</b> | <b>100</b> | 98.00        | 96.43        | <b>100</b> | 99.71      | 99.43        | 97.14        | 91.71        | <b>82.71</b> |
|                 | R2D2-(  )        | M             | in-batch             | 75.09        | 58.69        | 33.17        | 97.06           | 87.06        | 80.15        | 95.71        | 89.71        | 83.29        | 85.71        | 60.29        | 45.71        | 86.38           | 68.60        | 42.83        | <b>100</b> | 98.00        | <b>96.86</b> | <b>100</b> | <b>100</b> | <b>99.57</b> | 98.57        | 91.43        | 82.57        |
|                 | R2D2-(  )        | M             | in-batch             | 75.70        | 58.79        | 33.03        | 97.06           | 86.57        | 80.88        | 95.71        | 89.14        | 82.57        | 85.57        | 60.00        | 46.29        | 86.49           | 68.67        | 42.86        | <b>100</b> | 98.29        | 96.43        | <b>100</b> | 99.29      | <b>100</b>   | 92.00        | 82.14        |              |
|                 | R2D2-(  )-gd     | M             | in-batch             | 75.97        | 58.70        | 32.85        | 97.06           | 88.24        | 81.03        | 95.71        | <b>90.00</b> | 83.14        | 85.71        | 60.29        | 46.86        | 86.31           | 68.63        | 42.85        | <b>100</b> | 98.00        | 96.29        | <b>100</b> | 99.43      | <b>100</b>   | 91.14        | <b>82.71</b> |              |
|                 | R2D2-(  )-selfgd | M             | in-batch             | 73.47        | 57.57        | 32.16        | 97.06           | 85.59        | 76.91        | 95.71        | 88.86        | 81.29        | 85.71        | 58.57        | 46.43        | 86.04           | 68.48        | 42.76        | <b>100</b> | 97.71        | 96.14        | <b>100</b> | 99.71      | 99.29        | 98.57        | 89.71        | 82.29        |
|                 | R2D2-(  )        | M+P           | in-batch             | <b>76.65</b> | 59.07        | 33.15        | <b>97.06</b>    | <b>88.82</b> | <b>81.62</b> | <b>95.71</b> | 89.43        | <b>83.71</b> | <b>85.71</b> | 60.57        | <b>47.14</b> | <b>87.02</b>    | <b>68.80</b> | 42.85        | <b>100</b> | <b>98.57</b> | <b>96.86</b> | <b>100</b> | <b>100</b> | <b>99.57</b> | 98.57        | 91.43        | 81.86        |
|                 | CAPS-(  )        | M             | in-batch             | 74.68        | 58.64        | 32.54        | 95.59           | 87.06        | 78.38        | 94.29        | 89.71        | 82.71        | 82.86        | 59.14        | 46.29        | 87.88           | 69.25        | 43.46        | <b>100</b> | 99.14        | 96.71        | <b>100</b> | 99.71      | 100          | 92.86        | 84.14        |              |
|                 | CAPS-(  )-gd     | M             | in-batch             | 77.06        | 59.67        | 32.89        | 95.59           | 88.82        | <b>82.04</b> | 94.29        | 90.57        | <b>84.71</b> | 82.86        | 60.00        | 47.57        | 87.36           | 69.19        | 43.54        | <b>100</b> | 98.86        | 96.71        | <b>100</b> | 99.86      | 100          | 94.00        | 84.71        |              |
|                 | CAPS-(  )-selfgd | M             | in-batch             | <b>77.17</b> | 59.54        | 32.71        | <b>97.06</b>    | <b>89.85</b> | 81.47        | <b>95.71</b> | <b>90.86</b> | 84.43        | <b>85.71</b> | 60.57        | 47.29        | 87.61           | 69.14        | 43.33        | <b>100</b> | 98.86        | 96.71        | <b>100</b> | <b>100</b> | <b>99.57</b> | <b>100</b>   | 93.43        | 83.71        |
|                 | CAPS-(  )        | M+P           | in-batch             | 76.49        | <b>59.72</b> | 33.60        | 95.59           | 88.16        | 81.54        | 94.29        | 89.71        | 84.14        | 84.29        | 62.00        | 48.14        | 87.66           | 69.12        | 43.34        | <b>100</b> | 98.86        | 97.00        | <b>100</b> | <b>100</b> | <b>100</b>   | 93.14        | 84.71        |              |

Table 3: **Image retrieval performance.** We report results for the  $\mathcal{R}\text{Oxford5k}$  and  $\mathcal{R}\text{Paris6k}$  datasets following the evaluation protocol proposed in [54]. We use the same color-map for the supervised baseline approaches as in Table 2a. The proposed unsupervised models demonstrate competitive performance

pervised models demonstrate competitive performance. Interestingly, the model utilizing global image descriptor for hard-negative mining (CAPS-(, )-gd) achieves the best Recall and Precision performance over all other methods.

**Visual localization.** We evaluate the performance of the proposed pipeline on the task of image-based localization. Note, that we train our approaches only on outdoor datasets, *i.e.* Megadepth and Phototourism but evaluate on both outdoor (Aachen v1.1) and indoor (InLoc) to test the generalization of the proposed models. We use R2D2, CAPS and R2D2\* as baseline models where R2D2\* represents R2D2 model trained with SuperPoint detector, *i.e.*, the setup which is closer to our pipeline (see Sec. 3). For unsupervised methods, we report results obtained by the models utilize in-batch hard-negative sampling during training. The percentage of correctly localized query images is presented in Table 2a. First, if we compare R2D2-() and R2D2-(, ) trained on Aachen database images (A), we can notice that using only synthetic color augmentation is not sufficient to handle drastic illumination changes. Increasing the number of training samples leads to better localization performance. The model R2D2-(, ) trained on (M+P) is comparable to the supervised counterpart on day-time and surpasses it on night-time queries by a noticeable margin: up to **+0.5** for (0.5m, 5°) and **+1.1** for (5m, 10°). For indoor pose estimation, the same model can outperform the supervised baseline in all metrics: up to **+1.5** and **+2.9** for the finest threshold (0.25m, 2°) for scene DUC1 and DUC2 respectively. Leveraging visual descriptors (\*-gd and \*-selfgd) provides better outdoor localization results over the models trained with only stylization and color augmentations (R2D2-(, ) and CAPS-(, )). However, in indoors, visual descriptors *cannot* provide strong generalization performance leading to a marginal improvement (for the R2D2 backbone) or even poor results (with the CAPS backbone). We hypothesize that such behaviour is caused by larger gap in data distribution during training and evaluation. Interestingly, the transferring accuracy depends on the Descriptor CNN (*c.f.* Figure 1) structure. The unsuper-

vised model trained on (M+P) and based on the CAPS backbone *cannot* improve performance of outdoor localization compared to its supervised baseline (CAPS [76]). However, large gains are seen for indoor setting: up to **+6.1** (34.4 vs. 40.5) for the most rigorous threshold of scene DUC2. Table 2b shows the localization performance on Tokyo24/7. Similarly to the results on Aachen, R2D2-(, ) is *better* than its supervised counterpart (R2D2-[A+R]) in all metrics: **+1.58**, **+1.9**, and **+1.9**, respectively. Qualitative results of the proposed models on test images of MegaDepth are illustrated in Figure 3.

**Image retrieval.** As shown in Table 3, the proposed unsupervised models show competitive image retrieval performance and surpass their supervised counterparts in mP@k and mAP on  $\mathcal{R}\text{Oxford5k}$  and  $\mathcal{R}\text{Paris6k}$ . Most notably, utilizing visual image descriptors with the model based on CAPS backbone leads to better retrieval accuracy which is consistent with the results obtained for the camera localization benchmark. However, for the R2D2 backbone model, mining harder negative samples with the help of visual descriptors leads to a marginal improvement (R2D2-(, ) vs. R2D2-(, )-gd) which can be explained by the network capacity. In  $\mathcal{R}\text{Paris6k}$  both the supervised and unsupervised methods attain comparable performance. Notably, the biggest performance improvement of 5% is observed across mP@10 under Hard setting.


We observe that increasing the size of dataset ( $A \rightarrow M \rightarrow M + P$ ) performs favourably on the different benchmarks. Stylization provides substantial improvements on Aachen Night time dataset, while performs comparably on other benchmarks. This indicates that the proposed network is not biased to the chosen styles. Furthermore, the proposed coarse-to-fine method with global descriptors also performs favourably compared to in-batch sampling especially on Aachen and  $\mathcal{R}\text{Oxford5k}$  datasets. Finally, the self-supervised global mining approach \*-selfgd performs comparably to \*-gd which utilizes image retrieval models pre-trained on SfM data. Improving the on the fly strategy of mining hard negative local samples conditioned on visual




image descriptors is a possible future direction.

## 5. Ablation Study

The goal of this section is to investigate the benefits of using intra-batch hard negative mining compared to in-pair sampling used in recent CNN-based interest point detectors and descriptors [20, 57]. For in-batch negative mining, we experiment with different sampling strategies: consider *all* negative samples in a mini-batch, take *k* negative samples *randomly*, and consider only *top-k* samples.

**In-pair vs. in-batch negative mining.** To evaluate the effectiveness of the sampling procedure we conduct controlled experiments summarized in Table 4a. The results are obtained by the R2D2-() model trained on the MegaDepth dataset using the AP loss function. We use the following evaluation benchmarks: sparse feature matching (HPatches), image retrieval (ROxford5k), and visual localization (Aachen v1.1) (see Sec. 4.3). As shown in Table 4a, in-batch hard-negative sampling significantly outperforms in-pair mining in all metrics. We observe similar behaviour with metric learning losses, *e.g.*, Triplet loss. The results obtained by models trained with Triplet loss are provided in supplementary.

**In-batch sampling: all vs. random vs. topK.** In contrast to *pairwise* ranking losses, the AP loss is defined on a ranked list of samples *L*. Let us assume we have a mini-batch consisting of *M* homography related image crop pairs where each crop has *p* interest points extracted by SuperPoint [17] detector. Each keypoint is associated with a feature vector and has one positive and  $M \times (p-1)$  negative samples within a mini-batch. Therefore, we consider 3 ways of constructing the list *L*: take all possible negatives, randomly select *k* samples from  $M \times (p-1)$ , and select *top-k* negatives (the proposed strategy). For each setting, we train our unsupervised R2D2-() model on (A) and evaluate on two benchmarks: image retrieval and image-based localization. The results are presented in Table 4b. The proposed "top-k negative" sampling approach surpasses its counterparts with a noticeable margin. The number of hard negative samples *k* is set to 30 after empirical evaluation for all our experiments. Interestingly, unlike Triplet loss [44], the AP loss with random sampling will *not* lead to overfit, but evaluation performance is low.

## 6. Conclusion

In this work, we have presented a self-supervised local descriptor learning framework. We introduced three main contributions: (1) a coarse-to-fine method to efficiently sample hard-negative local descriptors by leveraging visual image descriptors, (2) the modified version of AP loss which can carefully handle imbalance between the number of positive and negative keypoint descriptors and (3) photoreal-

| Metric    |                 |       | Negative sampling type |  |
|-----------|-----------------|-------|------------------------|--|
|           |                 |       | in-pair                | in-batch                                       |
| ROxford5k | mAP             | M     | 55.38                  | <b>58.69</b>                                   |
|           |                 | H     | 29.67                  | <b>33.17</b>                                   |
|           | mP@k [1, 5, 10] | M     | [94.29, 86.57, 78.43]  | [ <b>95.71</b> , <b>89.71</b> , <b>83.29</b> ] |
|           |                 | H     | [82.86, 54.57, 42.29]  | [ <b>85.71</b> , <b>60.29</b> , <b>45.71</b> ] |
| HPatches  | MMA             | 1px   | 0.239 / 0.425 / 0.332  | <b>0.254 / 0.439 / 0.346</b>                   |
|           |                 | 3px   | 0.585 / 0.677 / 0.631  | <b>0.630 / 0.707 / 0.669</b>                   |
|           |                 | 5px   | 0.648 / 0.742 / 0.695  | <b>0.706 / 0.784 / 0.745</b>                   |
| Aachen    |                 | day   | 87.9 / 94.2 / 97.9     | <b>88.2 / 95.5 / 98.7</b>                      |
|           |                 | night | 66.5 / 79.1 / 91.6     | <b>68.1 / 83.8 / 94.8</b>                      |

(a) **In-batch vs. in-pair** hard negative sampling. Intra-batch negative mining performs consistently better over in-pair counterpart in all metrics.

| Metric      |                                  |       | Baseline    | Negative samples in mini-batch |               |              |
|-------------|----------------------------------|-------|-------------|--------------------------------|---------------|--------------|
|             |                                  |       |             | <i>all</i>                     | <i>random</i> | <i>top-k</i> |
| ROxford5k   | mAP                              | E     | 72.60       | <u>72.85</u>                   | 72.06         | <b>75.27</b> |
|             |                                  | M     | 54.89       | <u>57.06</u>                   | 56.91         | <b>57.83</b> |
|             |                                  | H     | 27.37       | <u>30.53</u>                   | 31.28         | <b>31.56</b> |
|             | mP@k [1, 5, 10]                  | E     | 94.12       | <b>97.06</b>                   | 95.59         | <b>97.06</b> |
|             |                                  |       | 80.74       | 84.41                          | <u>84.63</u>  | <b>88.24</b> |
|             |                                  |       | 74.12       | 77.40                          | <u>77.72</u>  | <b>80.44</b> |
|             |                                  | M     | 91.43       | <b>95.71</b>                   | 94.29         | <b>95.71</b> |
|             |                                  |       | 80.57       | <u>88.10</u>                   | <b>88.86</b>  | <b>88.86</b> |
|             |                                  |       | 73.86       | <u>81.24</u>                   | 81.00         | <b>81.41</b> |
|             |                                  | H     | 61.43       | 81.43                          | 81.43         | <b>82.86</b> |
|             |                                  |       | 47.50       | 56.43                          | <u>57.14</u>  | <b>59.43</b> |
|             |                                  |       | 39.79       | 43.57                          | <u>44.57</u>  | <b>45.43</b> |
| Tokyo24/7   | recall                           | r@1   | 67.10       | <u>76.15</u>                   | 74.92         | <b>79.05</b> |
|             |                                  | r@5   | 74.20       | <u>81.27</u>                   | 80.63         | <b>84.13</b> |
|             |                                  | r@10  | 76.10       | <u>83.49</u>                   | 81.90         | <b>85.71</b> |
| Aachen v1.1 | 0.25m, 2°<br>0.5m, 5°<br>5m, 10° | day   | 87.7        | 87.5                           | 88.0          | <b>88.0</b>  |
|             |                                  |       | <u>94.7</u> | 94.4                           | 94.2          | <b>94.8</b>  |
|             |                                  |       | <b>98.7</b> | <u>98.3</u>                    | 98.1          | 98.2         |
|             |                                  | night | 69.6        | <b>70.2</b>                    | 67.0          | <b>70.2</b>  |
|             |                                  |       | <b>86.4</b> | 85.9                           | 84.3          | <b>86.4</b>  |
|             |                                  |       | <u>95.3</u> | <b>95.8</b>                    | 94.8          | <b>95.8</b>  |
| InLoc       | 0.25m, 2°<br>0.5m, 5°<br>5m, 10° | DUC1  | 27.8        | <b>28.3</b>                    | 27.3          | 27.8         |
|             |                                  |       | 42.4        | <b>46.0</b>                    | 43.4          | <u>43.9</u>  |
|             |                                  |       | <u>54.4</u> | <b>55.6</b>                    | 53.0          | 54.0         |
|             |                                  | DUC2  | 22.1        | 23.7                           | <u>26.0</u>   | <b>27.5</b>  |
|             |                                  |       | 34.4        | 35.1                           | <u>36.6</u>   | <b>38.2</b>  |
|             |                                  |       | 42.7        | 40.5                           | <u>44.3</u>   | <b>49.6</b>  |

(b) **Intra-batch negative sampling.** For a listwise loss, the *top-k* sampling strategy leads to a better performance over its counterparts in almost all metrics. The best score is in bold and the second best score is underlined.

Table 4: **Ablation study.** We analyze the influence of different design choices of the proposed approach. See Sec. 5 for more details.

istic image stylization that combined with synthetic color augmentation can significantly enhance the robustness of learned local descriptors to illumination changes. The proposed improvements together result in a self-supervised approach demonstrates favourable performance on a wide range of geometric tasks compared to strong fully- and weakly-supervised baseline models.

## A. Overview

In this supplementary material, we provide additional results and visualization of the proposed self-supervised framework for local image descriptors learning. In Sec. B, we provide the details of the training data and training procedure. Sec. C shows the evaluation on the recently proposed Rotated Day-Night Image Matching dataset (RDNIM) [51] representing strong illumination and viewpoint changes. To verify generalization performance of our method, we present ablation study on different detectors in Sec. D. In Sec. E, we demonstrate more qualitative results of matches produced by our method on outdoor and indoor benchmarks.

## B. Training details

In this work, we use three large-scale outdoor datasets: MegaDepth (M) [38], Phototourism (P), and database images of Aachen v1.1 (A). For M and P, the official training and validation splits have been utilized to optimize the proposed approach presented in the main part in Figure 1. All models are trained in an end-to-end fashion with the Adam optimizer [32] and an initial learning rate of  $10^{-3}$  and  $10^{-4}$  for unsupervised versions of R2D2 and CAPS models respectively. We use the  $\left(\begin{smallmatrix} \blacksquare & \blacksquare \\ \blacksquare & \blacksquare \end{smallmatrix}\right)$  notation for the models trained on synthetic homographies with color augmentations while the  $\left(\begin{smallmatrix} \blacksquare & \blacksquare \\ \blacksquare & \blacksquare \end{smallmatrix}\right)$  encodes the models trained on images with color augmentations and stylization. The training was performed on a single GeForce RTX 2080Ti GPU. The style categories and image stylization utilized during training are illustrated in Figure 5.

**Synthetic color augmentation.** To improve generalization performance of the proposed approach, we apply synthetic color augmentation separately on the source and target image crops. Specifically, we utilize histogram equalization (CLAHE) and add per-pixel Gaussian noise and Gaussian blur along with color augmentation in contrast, brightness, hue, and saturation. The full set of color augmentations is presented in Figure 4.

## C. Evaluation on RDNIM

In order to verify the robustness of local image descriptors to severe appearance and viewpoint changes, Pautrat *et al.* [51] propose the RDNIM dataset. The dataset originates from the DNIM dataset – a subset of the AMOS database [29] represents a large number of images taken at regular time intervals by outdoor webcams with fixed positions and orientations. The DNIM dataset consists of sequences, with a total of 1722 day-time and night-time images. Pautrat *et al.* [51] provide two benchmarks, where the images of each sequence are paired with either the day-time or the night-time image. The images are then augmented with homographic warps including rotations simi-

| Method              | Day reference   |              |              | Night reference |              |              |
|---------------------|-----------------|--------------|--------------|-----------------|--------------|--------------|
|                     | H               | Precision    | Recall       | H               | Precision    | Recall       |
| LISRD [51]          | 0.198           | 0.291        | 0.317        | 0.262           | 0.371        | 0.384        |
| LISRD*              | <b>0.358</b>    | <b>0.433</b> | <b>0.526</b> | <b>0.442</b>    | <b>0.541</b> | <b>0.621</b> |
| HardNet [44]        | 0.249           | 0.225        | 0.224        | 0.325           | 0.359        | 0.365        |
| SOSNet [70]         | 0.226           | 0.218        | 0.226        | 0.252           | 0.288        | 0.296        |
| SuperPoint [17]     | 0.178           | 0.191        | 0.214        | 0.235           | 0.259        | 0.296        |
| R2D2 [57]           | 0.215           | 0.219        | 0.216        | 0.254           | 0.289        | 0.280        |
| CAPS [76]           | 0.159           | 0.258        | 0.278        | 0.212           | 0.308        | 0.310        |
| Unsupervised (ours) | R2D2-[M]        | 0.233        | 0.238        | 0.244           | <u>0.281</u> | 0.312        |
|                     | R2D2-[M+P]      | 0.222        | <u>0.245</u> | <u>0.253</u>    | 0.274        | <u>0.317</u> |
|                     | R2D2-[M]-selfgd | <u>0.237</u> | 0.229        | 0.228           | 0.278        | 0.304        |
|                     | R2D2-[M]-gd     | 0.229        | 0.230        | 0.239           | 0.274        | 0.307        |
|                     | CAPS-[M]        | 0.180        | <u>0.267</u> | 0.256           | 0.250        | <u>0.337</u> |
|                     | CAPS-[M+P]      | 0.193        | <u>0.267</u> | 0.254           | 0.250        | 0.336        |
|                     | CAPS-[M]-selfgd | <u>0.215</u> | <u>0.263</u> | <u>0.282</u>    | <u>0.271</u> | 0.331        |
|                     | CAPS-[M]-gd     | 0.209        | <u>0.267</u> | <u>0.282</u>    | 0.262        | <u>0.350</u> |
|                     |                 |              |              |                 |              |              |
|                     |                 |              |              |                 |              |              |

Table 5: **Evaluation results on RDNIM [51]** We follow [51] and report the accuracy of homography estimation, precision and recall for error thresholds of 3 pixels. The best score between supervised (weakly-) baselines, *i.e.*, R2D2 and CAPS, and the proposed unsupervised methods for each category is underlined. The overall best score is **bold**.

| Method          | Supervision | Hard-negative sampling | Aachen v1.1         |          |         |                    |          |         |
|-----------------|-------------|------------------------|---------------------|----------|---------|--------------------|----------|---------|
|                 |             |                        | % localized queries |          |         |                    |          |         |
|                 |             |                        | Day (824 images)    |          |         | Night (191 images) |          |         |
|                 |             |                        | 0.25m, 2°           | 0.5m, 5° | 5m, 10° | 0.25m, 2°          | 0.5m, 5° | 5m, 10° |
| R2D2 [57]-[A+R] | OF          | in-pair                | 86.5                | 95.0     | 98.2    | 69.1               | 84.5     | 93.2    |
| R2D2 [57]-[A+R] | OF          | in-batch               | 88.6                | 95.5     | 98.7    | 71.2               | 86.4     | 94.8    |
| R2D2-[A]        | -           | in-batch               | 87.9                | 94.5     | 98.4    | 63.9               | 79.1     | 90.6    |
| R2D2-[A]        | -           | in-batch               | 87.1                | 94.3     | 98.4    | 68.1               | 84.3     | 92.7    |
| R2D2-[M+P]      | -           | in-batch               | 88.1                | 95.1     | 98.6    | 72.1               | 88.2     | 95.4    |

Table 6: **Localization with Triplet loss.** The percentage of correctly localized queries of Aachen Day-Night v1.1 [85] under specific thresholds. **Legend:** *Supervision*: OF – ground-truth pixel correspondences based on optical flow [57]. The training datasets are given in [.] parentheses: A – Aachen database images; R – Random images from the Internet [57]; M – MegaDepth dataset; P – Phototourism dataset.

larly to [17].

The evaluation of the state-of-the-art local descriptors with SuperPoint detector is summarized in Table 5. For LISRD, in addition to the results presented in [51], we also evaluate one of the models publicly available at <https://github.com/rpautrat/LISRD>. We refer to this model as LISRD\*. The proposed models perform better than the fully- and weakly-supervised counterparts (R2D2 [57] and CAPS [76], respectively) in all metrics. The models have been trained on stylized images with synthetic color augmentation, *i.e.*,  $\left(\begin{smallmatrix} \blacksquare & \blacksquare \\ \blacksquare & \blacksquare \end{smallmatrix}\right)$ . Notably, using larger model (CAPS) and visual image descriptors (-gd) leads to better performance which is consistent with the results obtained for image retrieval and camera relocation benchmarks presented in the main part.

|       | Error threshold | SIFT detector |             |             | SuperPoint detector |             |             | D2-Net detector           |             |             |
|-------|-----------------|---------------|-------------|-------------|---------------------|-------------|-------------|---------------------------|-------------|-------------|
|       |                 | SIFT          | Ours (R2D2) | Ours (CAPS) | SuperPoint          | Ours (R2D2) | Ours (CAPS) | D2-Net (SS <sup>2</sup> ) | Ours (R2D2) | Ours (CAPS) |
| day   | 0.25m, 2°       | 76.2          | <b>80.0</b> | 77.3        | 86.8                | <b>87.4</b> | 85.1        | <b>84.0</b>               | 83.2        | 79.2        |
|       | 0.5m, 5°        | 84.3          | <b>88.6</b> | 88.5        | 94.1                | <b>94.7</b> | 93.2        | 92.1                      | <b>92.2</b> | 89.0        |
|       | 5m, 10°         | 91.4          | 94.9        | <b>95.4</b> | <b>98.3</b>         | <b>98.3</b> | 97.8        | <b>97.8</b>               | 97.3        | 97.3        |
| night | 0.25m, 2°       | 52.1          | 56.5        | <b>61.4</b> | 68.1                | <b>72.3</b> | 71.7        | <b>73.8</b>               | 68.1        | 67.0        |
|       | 0.5m, 5°        | 63.3          | 71.2        | <b>77.1</b> | 85.9                | <b>88.5</b> | 87.4        | <b>92.1</b>               | 87.3        | 85.3        |
|       | 5m, 10°         | 74.5          | 84.8        | <b>93.3</b> | 96.9                | 97.4        | <b>97.9</b> | <b>98.4</b>               | 97.2        | 97.9        |

(a) Localization performance of the proposed keypoint descriptor with different detectors

|       |           | SIFT detector |             |                |                    |                | SuperPoint detector |             |                |                    |                | D2-Net detector |             |                |                    |                |
|-------|-----------|---------------|-------------|----------------|--------------------|----------------|---------------------|-------------|----------------|--------------------|----------------|-----------------|-------------|----------------|--------------------|----------------|
|       |           | SIFT          | Ours selfgd | Ours (R2D2) gd | Ours (CAPS) selfgd | Ours (CAPS) gd | SuperPoint          | Ours selfgd | Ours (R2D2) gd | Ours (CAPS) selfgd | Ours (CAPS) gd | D2-Net (SS)     | Ours selfgd | Ours (R2D2) gd | Ours (CAPS) selfgd | Ours (CAPS) gd |
| day   | 0.25m, 2° | 76.2          | 79.2        | <b>79.9</b>    | 79.1               | 78.5           | 86.8                | <b>88.1</b> | 87.5           | 86.9               | 87.0           | 84.0            | 83.4        | <b>85.0</b>    | 82.6               | 82.9           |
|       | 0.5m, 5°  | 84.3          | 88.1        | 88.2           | <b>89.1</b>        | 88.8           | 94.1                | 94.8        | <b>94.9</b>    | 93.8               | 93.8           | 92.1            | 92.5        | <b>92.6</b>    | 90.5               | 91.1           |
|       | 5m, 10°   | 91.4          | 93.6        | 94.4           | 95.4               | <b>95.6</b>    | <b>98.3</b>         | 98.1        | <b>98.3</b>    | 98.1               | <b>98.3</b>    | 97.8            | 97.5        | 97.8           | 97.5               | <b>98.1</b>    |
| night | 0.25m, 2° | 52.1          | 49.7        | 52.4           | 64.9               | <b>66.0</b>    | 68.1                | 71.2        | 71.7           | 71.7               | <b>73.8</b>    | <b>73.8</b>     | 69.6        | 70.7           | 71.2               | 69.1           |
|       | 0.5m, 5°  | 63.3          | 67.5        | 71.2           | <b>83.8</b>        | 81.7           | 85.9                | 88.0        | 86.4           | <b>89.0</b>        | <b>89.0</b>    | <b>92.1</b>     | 87.4        | 88.0           | 84.8               | 87.4           |
|       | 5m, 10°   | 74.5          | 78.5        | 84.8           | <b>95.8</b>        | 94.8           | 96.9                | 95.8        | 96.9           | <b>97.4</b>        | <b>97.4</b>    | <b>98.4</b>     | 97.4        | 97.4           | 97.9               | <b>98.4</b>    |

(b) Localization performance of the proposed method utilizing *visual image descriptors* during training.

Table 7: **Visual localization performance on Aachen Day-Night v1.1 [85]**. To verify generalization performance, we evaluate the proposed learned local descriptor on Aachen v1.1 using different keypoint detectors. The overall best score for each detector is in **bold**.

## D. Ablation study

**Triplet Loss.** In addition to the results presented in our main work, we verify the proposed ideas of hard-negative sampling and photo-realistic image stylization using Triplet loss. For in-pair sampling, the closest non-matching descriptor is selected for each positive-anchor pair from a set of descriptors extracted from an input image pair. To create our supervised baseline, we adapt the model by Revaud *et al.* [57] and train it with Triplet loss using ground-truth pixel correspondences produced by optical flow. We evaluate the baselines and our models on the Aachen Day-Night v1.1 [85] dataset and report the percentage of correctly localized queries under specific error thresholds. The results are presented in Table 6. Interestingly, the triplet margin loss and the list-wise loss discussed in the main part behave very similarly. The in-batch hard negative sampling leads to better results over in-pair mining. Utilizing both synthetic color augmentation and image stylization is very powerful and can significantly improve localization performance.

**Detectors.** In order to analyze generalization performance of the proposed method, we evaluate our local descriptors on different keypoint detectors. Specifically, we train our models with SuperPoint keypoints but utilize SIFT [39] and D2-Net [20] detectors during evaluation. We consider the problem of visual localization as our benchmark and report the number of correctly localized query images of Aachen

Day-Night [85] under different settings in Table 7. As shown in Table 7, our models are not limited to SuperPoint detector but perform favorably with other detectors. The models leveraging visual image descriptors for hard-negative mining during training demonstrate better generalization performance and localization results (Table 7b, D2-Net detector).

## E. Qualitative results

We visualize the correspondences estimated by our pipeline (with CAPS backbone network) and its weakly-supervised strong counterpart on the image pairs of Aachen Day-Night [85] in Figure 6 and of InLoc [67] in Figure 7. The matches are verified by using the `findFundamentalMatrix` function with a RANSAC threshold of 1 [86]. For a randomly selected query image, we choose the database image with the most inlier correspondences verified by the camera pose solver. Figure 8 shows qualitative results on InLoc queries obtained by the model based on R2D2 architecture. The correspondences are illustrated in green.

## References

- [1] Asha Anooosheh, Torsten Sattler, Radu Timofte, Marc Pollefeys, and Luc Van Gool. Night-to-day image translation for retrieval-based localization. In *International Conference on Robotics and Automation (ICRA)*, pages 5958–5964, 2019. 3
- [2] Relja Arandjelović, Petr Gronat, Akihiko Torii, Tomas Pajdla, and Josef Sivic. NetVLAD: CNN architecture for weakly supervised place recognition. In *Proceedings of*

<sup>1</sup>We evaluate the single scale model only due to technical restrictions of our hardware

```

1 import albumentations as A
2 from albumentations.pytorch import ToTensorV2
3
4
5 def get_imagenet_mean_std():
6     mean, std = [0.485, 0.456, 0.406], [0.229, 0.224, 0.225]
7     return mean, std
8
9
10 def train_img_augmentations():
11     mean, std = get_imagenet_mean_std()
12     augs = A.Compose(
13         [
14             A.OneOf(
15                 [
16                     A.RandomBrightnessContrast(brightness_limit=0.3,
17                                                 contrast_limit=0.2,
18                                                 p=0.7),
19                     A.HueSaturationValue(hue_shift_limit=10,
20                                         sat_shift_limit=30,
21                                         val_shift_limit=0,
22                                         p=0.7),
23                     A.CLAHE(p=0.5),
24                     A.ToGray(p=0.5),
25                     A.ChannelShuffle(p=0.1),
26                 ],
27                 p=0.6,
28             ),
29             A.OneOf(
30                 [
31                     A.GaussianBlur(p=0.5),
32                     A.Blur(p=0.5)
33                 ],
34                 p=0.3,
35             ),
36             A.OneOf([A.GaussNoise(p=0.5),
37                     A.FAAAdditiveGaussianNoise(p=0.5)],
38                     p=0.1),
39             A.Normalize(mean, std),
40             ToTensorV2(),
41         ],
42         p=1,
43     )
44     return augs

```

Figure 4: **Color augmentation** utilized in the proposed pipeline. We used Albumentations [10] library to generate random color transforms of training images

the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), pages 5297–5307, 2016. 3, 6

- [3] Philip Bachman, R Devon Hjelm, and William Buchwalter. Learning representations by maximizing mutual information across views. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 32, pages 15535–15545. Curran Associates, Inc., 2019. 1
- [4] Sungyong Baik, Hyo Jin Kim, Tianwei Shen, Eddy Ilg, Kyoung Mu Lee, and Chris Sweeney. Domain adaptation of learned features for visual localization. In *Proceedings of the British Machine Vision Conference (BMVC)*, 2020. 3
- [5] 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)*, pages 3852–3861, 2017. 1, 5
- [6] Vassileios Balntas, Edgar Riba, Daniel Ponsa, and Krystian Mikolajczyk. Learning local feature descriptors with triplets and shallow convolutional neural networks. In *Proceedings of the British Machine Vision Conference (BMVC)*, 2016. 1, 2
- [7] Axel Barroso-Laguna, Edgar Riba, Daniel Ponsa, and Krystian Mikolajczyk. KeyNet: Keypoint detection by handcrafted and learned CNN filters. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 5835–5843, 2019. 2
- [8] 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. Springer Berlin Heidelberg, 2006. 2
- [9] Aritra Bhowmik, Stefan Gumfold, 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 4947–4956, 2020. 2
- [10] Alexander Buslaev, Vladimir I. Iglovikov, Eugene Khvedchenya, Alex Parinov, Mikhail Druzhinin, and Alexandr A. Kalinin. Albumentations: Fast and flexible image augmentations. *Information*, 11(2), 2020. 11
- [11] Michael Calonder, Vincent Lepetit, Christoph Strecha, and Pascal Fua. BRIEF: Binary robust independent elementary features. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 778–792. Springer Berlin Heidelberg, 2010. 2
- [12] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *Proceedings of the International Conference on Machine Learning (ICML)*, 2020. 1, 3
- [13] 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. 1, 2, 3
- [14] Ching-Yao Chuang, Joshua Robinson, Yen-Chen Lin, Antonio Torralba, and Stefanie Jegelka. Debaised contrastive learning. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 33, pages 8765–8775. Curran Associates, Inc., 2020. 1, 3
- [15] Ekin D. Cubuk, Barret Zoph, Jonathon Shlens, and Quoc V. Le. RandAugment: Practical automated data augmentation with a reduced search space. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, pages 3008–3017, 2019. 3
- [16] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. ImageNet: A large-scale hierarchical image database. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 248–255, 2009. 5
- [17] 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 (CVPR) Workshops*, pages 337–349, 2018. 1, 2, 3, 4, 5, 8, 9
- [18] Carl Doersch, Abhinav Gupta, and Alexei A. Efros. Unsupervised visual representation learning by context prediction. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 1422–1430, 2015. 1
- [19] Alexey Dosovitskiy, Jost Tobias Springenberg, Martin Riedmiller, and Thomas Brox. Discriminative unsupervised feature learning with convolutional neural networks. In *Advances in Neural Information Processing Systems (NIPS)*, volume 27, pages 766–774. Curran Associates, Inc., 2014. 1






a) Dusk styles



b) Night styles



c) Stylization

Figure 5: **The set of images used for photo-realistic stylization.** In contrast to [57], we use two style categories, *i.e.* Dusk and Night, and 7 style examples for each category. We randomly sample a style image from these two categories and apply it to each particular image from the training dataset.

- [20] 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)*, pages 8084–8093, 2019. [1](#), [2](#), [3](#), [5](#), [6](#), [8](#), [10](#)

- [21] Hao-Shu Fang, Jianhua Sun, Runzhong Wang, Minghao Gou, Yong-Lu Li, and Cewu Lu. InstaBoost: Boosting instance segmentation via probability map guided copy-pasting. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 682–691, 2019. [1](#)

- [22] Hugo Germain, Guillaume Bourmaud, and Vincent Lepetit.





Figure 6: **Example correspondences** produced by our method (left) and its weakly-supervised counterpart, *i.e.* CAPS [76] (right), on night-time queries of Aachen Day-Night (v1.1) [85]






Figure 7: **Example correspondences** produced by our method (left) and its weakly-supervised counterpart, *i.e.* CAPS [76] (right), on queries of InLoc [67]






Figure 8: **Example correspondences** produced by our method (left) and its supervised counterpart, *i.e.* R2D2 [57] (right), on queries of InLoc [67]



- S2DNet: Learning image features for accurate sparse-to-dense matching. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 626–643. Springer International Publishing, 2020. [2](#)
- [23] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *Advances in Neural Information Processing Systems (NIPS)*, volume 27, pages 2672–2680. Curran Associates, Inc., 2014. [3](#)
- [24] Albert Gordo, Jon Almazán, Jerome Revaud, and Diane Larlus. End-to-end learning of deep visual representations for image retrieval. *International Journal on Computer Vision (IJCV)*, 124(2):237–254, 2017. [1](#), [4](#)
- [25] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 9726–9735, 2020. [1](#), [3](#)
- [26] Kun He, Yan Lu, and Stan Sclaroff. Local descriptors optimized for average precision. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 596–605, 2018. [2](#), [4](#)
- [27] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 770–778, 2016. [5](#)
- [28] Jared Heinly, Johannes Lutz Schönberger, Enrique Dunn, and Jan-Michael Frahm. Reconstructing the world\* in six days. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3287–3295, 2015. [1](#)
- [29] Nathan Jacobs, Nathaniel Roman, and Robert Pless. Consistent temporal variations in many outdoor scenes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1–6, 2007. [9](#)
- [30] Wei Jiang, Eduard Trulls, Jan Hosang, Andrea Tagliasacchi, and Kwang Moo Yi. COTR: Correspondence transformer for matching across images. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 6207–6217, 2021. [2](#)
- [31] Yannis Kalantidis, Mert Bulent Sariyildiz, Noe Pion, Philippe Weinzaepfel, and Diane Larlus. Hard negative mixing for contrastive learning. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 33, pages 21798–21809. Curran Associates, Inc., 2020. [1](#), [3](#)
- [32] Diederik P. Kingma and Jimmy Ba. Adam: A Method for Stochastic Optimization. In *International Conference on Learning Representations (ICLR)*, 2015. [9](#)
- [33] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. ImageNet classification with deep convolutional neural networks. In *Advances in Neural Information Processing Systems (NIPS)*, volume 25, pages 1097–1105. Curran Associates, Inc., 2012. [1](#)
- [34] Mathieu Labbé and François Michaud. RTAB-Map as an open-source lidar and visual simultaneous localization and mapping library for large-scale and long-term online operation. *Journal of Field Robotics*, 36(2):416–446, 2019. [2](#)
- [35] Zakaria Laskar, Iaroslav Melekhov, Hamed Rezazadegan Tavakoli, Juha Ylioinas, and Juho Kannala. Geometric image correspondence verification by dense pixel matching. In *Proceedings of the IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 2510–2519, 2020. [3](#), [6](#)
- [36] Karel Lenc and Andrea Vedaldi. Learning covariant feature detectors. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 100–117. Springer International Publishing, 2016. [2](#)
- [37] Yijun Li, Ming-Yu Liu, Xueting Li, Ming-Hsuan Yang, and Jan Kautz. A closed-form solution to photorealistic image stylization. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 468–483. Springer International Publishing, 2018. [2](#), [3](#)
- [38] 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. [6](#), [9](#)
- [39] David G. Lowe. Distinctive image features from scale-invariant keypoints. *International Journal on Computer Vision (IJCV)*, 60(2):91–110, 2004. [2](#), [6](#), [10](#)
- [40] 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)*, pages 2522–2531, 2019. [2](#)
- [41] 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)*, pages 6588–6597, 2020. [2](#)
- [42] Iaroslav Melekhov, Gabriel J. Brostow, Juho Kannala, and Daniyar Turmukhambetov. Image stylization for robust features. *ArXiv preprint arXiv:2008.06959*, 2020. [3](#)
- [43] Iaroslav Melekhov, Aleksei Tiulpin, Torsten Sattler, Marc Pollefeys, Esa Rahtu, and Juho Kannala. DGC-Net: Dense geometric correspondence network. In *Proceedings of the IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 1034–1042, 2019. [3](#)
- [44] Anastasiia Mishchuk, Dmytro Mishkin, Filip Radenovic, and Jiri Matas. Working hard to know your neighbor's margins: Local descriptor learning loss. In *Advances in Neural Information Processing Systems (NIPS)*, volume 30, pages 4826–4837. Curran Associates, Inc., 2017. [1](#), [2](#), [4](#), [5](#), [8](#), [9](#)
- [45] Dmytro Mishkin, Filip Radenović, and Jiri Matas. Repeatability is not enough: Learning affine regions via discriminability. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 287–304. Springer International Publishing, 2018. [2](#)
- [46] Rohit Mohan and Abhinav Valada. EfficientPS: Efficient panoptic segmentation. *International Journal on Computer Vision (IJCV)*, 129(5):1551–1579, 2020. [1](#)
- [47] Markus S. Mueller, Thorsten Sattler, Marc Pollefeys, and Boris Jutzi. Image-to-image translation for enhanced feature matching, image retrieval and visual localization. *ISPRS annals*, IV-2/W7:111–119, 2019. [3](#)

- [48] 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*, 33(5):1255–1262, 2017. [2](#)
- [49] Mehdi Noroozi and Paolo Favaro. Unsupervised learning of visual representations by solving jigsaw puzzles. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 69–84. Springer International Publishing, 2016. [1](#)
- [50] 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)*, volume 31, pages 6234–6244. Curran Associates, Inc., 2018. [2](#)
- [51] 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. Springer International Publishing, 2020. [3](#), [5](#), [6](#), [9](#)
- [52] Horia Porav, Maddern Will, and Paul Newman. Adversarial training for adverse conditions: Robust metric localisation using appearance transfer. In *International Conference on Robotics and Automation (ICRA)*, pages 1011–1018, 2018. [3](#)
- [53] Milan Pultar, Dmytro Mishkin, and Jiri Matas. Leveraging outdoor webcams for local descriptor learning. *ArXiv preprint arXiv:1901.09780*, 2020. [3](#)
- [54] Filip Radenović, Ahmet Iscen, Giorgos Tolias, Yannis Avrithis, and Ondřej Chum. Revisiting oxford and paris: Large-scale image retrieval benchmarking. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5706–5715, 2018. [6](#), [7](#)
- [55] Filip Radenović, Giorgos Tolias, and Ondřej Chum. CNN image retrieval learns from BoW: Unsupervised fine-tuning with hard examples. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 3–20. Springer International Publishing, 2016. [1](#), [4](#), [6](#), [7](#)
- [56] Filip Radenović, Giorgos Tolias, and Ondřej Chum. Fine-tuning CNN image retrieval with no human annotation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 41(7):1655–1668, 2019. [1](#), [4](#)
- [57] 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. Curran Associates, Inc., 2019. [1](#), [2](#), [3](#), [4](#), [5](#), [6](#), [7](#), [8](#), [9](#), [10](#), [12](#), [15](#)
- [58] Joshua David Robinson, Ching-Yao Chuang, Suvrit Sra, and Stefanie Jegelka. Contrastive learning with hard negative samples. In *International Conference on Learning Representations (ICLR)*, 2021. [1](#), [3](#)
- [59] Ignacio Rocco, Relja Arandjelović, and Josef Sivic. Convolutional neural network architecture for geometric matching. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2553–2567, 2017. [3](#)
- [60] 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, 2011. [2](#)
- [61] Paul-Edouard Sarlin, Cesar Cadena, Roland Siegwart, and Marcin Dymczyk. From coarse to fine: Robust hierarchical localization at large scale. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 12708–12717, 2019. [5](#)
- [62] 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 4937–4946, 2020. [5](#)
- [63] 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 3929–3937, 2017. [2](#)
- [64] Johannes Lutz Schönberger 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](#), [2](#)
- [65] Johannes Lutz Schönberger, Enliang Zheng, Marc Pollefeys, and Jan-Michael Frahm. Pixelwise view selection for unstructured multi-view stereo. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 501–518. Springer International Publishing, 2016. [2](#)
- [66] Jiaming Sun, Zehong Shen, Yuang Wang, Hujun Bao, and Zhou Xiaowei. 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. [2](#)
- [67] 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 *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1293–1307, 2018. [5](#), [6](#), [10](#), [14](#), [15](#)
- [68] Jiexiong Tang, Rares Ambrus, Vitor Guizilini, and Hanme Kim. Neural outlier rejection for self-supervised keypoint learning. In *Proceedings of the International Conference on Machine Learning (ICML)*, 2020. [2](#)
- [69] Yurun Tian, Bin Fan, and Fuchao Wu. L2-Net: Deep learning of discriminative patch descriptor in euclidean space. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6128–6136, 2017. [2](#)
- [70] 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 11008–11017, 2019. [2](#), [5](#), [9](#)
- [71] Akihiko Torii, Relja Arandjelović, Josef Sivic, Masatoshi Okutomi, and Tomas Pajdla. 24/7 place recognition by view synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 257–271, 2015. [5](#), [6](#)
- [72] Prune Truong, Stefanos Apostolopoulos, Agata Mosinska, Samuel Stucky, Carlos Ciller, and Sandro De Zanet. GLAM-points: Greedily learned accurate match points. In *Proceed-*

- ings of the *IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 10731–10740, 2019. [2](#)
- [73] Michał Tyszkiewicz, Pascal Fua, and Eduard Trulls. DISK: Learning local features with policy gradient. In *Advances in Neural Information Processing Systems*, volume 33, pages 14254–14265. Curran Associates, Inc., 2020. [2](#)
- [74] Yannick Verdie, Kwang Moo Yi, Vincent Lepetit, and Pascal Fua. TILDE: A temporally invariant learned detector. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5279–5288, 2015. [2](#)
- [75] Bing Wang, Changhao Chen, Zhaopeng Cui, Jie Qin, Chris Xiaoxuan Lu, Zhengdi Yu, Peijun Zhao, Zhen Dong, Fan Zhu, Niki Trigoni, and Andrew Markham. P2-net: Joint description and detection of local features for pixel and point matching. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 16004–16013, 2021. [2](#)
- [76] 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)*, pages 757–774. Springer International Publishing, 2020. [2](#), [3](#), [4](#), [5](#), [6](#), [7](#), [9](#), [13](#), [14](#)
- [77] Yuxin Wu, Alexander Kirillov, Francisco Massa, Wan-Yen Lo, and Ross Girshick. Detectron2. <https://github.com/facebookresearch/detectron2>, 2019. [1](#)
- [78] Guandao Yang, Tomasz Malisiewicz, and Serge Belongie. Learning data-adaptive interest points through epipolar adaptation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, pages 1–7, 2019. [2](#)
- [79] Heng Yang, Wei Dong, Luca Carlone, and Vladlen Koltun. Self-supervised geometric perception. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 14350–14361, 2021. [2](#), [3](#)
- [80] Tsun-Yi Yang, Duy-Kien Nguyen, Huub Heijnen, and Vasilios Balntas. UR2KiD: Unifying retrieval, keypoint detection, and keypoint description without local correspondence supervision. *ArXiv preprint arXiv:2001.07252*, 2020. [3](#)
- [81] Yuan Yao, Yasamin Jafarian, and Hyun Soo Park. MONET: Multiview semi-supervised keypoint detection via epipolar divergence. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 753–762, 2019. [2](#)
- [82] 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. Springer International Publishing, 2016. [2](#)
- [83] Richard Zhang, Phillip Isola, and Alexei A. Efros. Colorful image colorization. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 649–666. Springer International Publishing, 2016. [1](#)
- [84] Xu Zhang, Felix X. Yu, Svebor Karaman, and Shih-Fu Chang. Learning discriminative and transformation covariant local feature detectors. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4923–4931, 2017. [2](#)
- [85] Zichao Zhang, Torsten Sattler, and Davide Scaramuzza. Reference pose generation for long-term visual localization via learned features and view synthesis. *International Journal on Computer Vision (IJCV)*, 129(4):821–844, 2021. [5](#), [6](#), [9](#), [10](#), [13](#)
- [86] Qunjie Zhou, Torsten Sattler, and Laura Leal-Taixé. Patch2Pix: Epipolar-guided pixel-level correspondences. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4669–4678, 2021. [2](#), [10](#)
- [87] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 2242–2251, 2017. [3](#)