

**Master Thesis**



**Czech  
Technical  
University  
in Prague**

**F3**

**Faculty of Electrical Engineering  
Department of Cybernetics**

## **Improving the HardNet Descriptor**

**Bc. Milan Pultar**

**Supervisor: M.Sc. Dmytro Mishkin**

**Field of study: Open Informatics**

**Subfield: Computer Vision and Image Processing**

**May 2020**



## I. Personal and study details

Student's name: **Pultar Milan** Personal ID number: **476840**  
Faculty / Institute: **Faculty of Electrical Engineering**  
Department / Institute: **Department of Cybernetics**  
Study program: **Open Informatics**  
Specialisation: **Computer Vision and Image Processing**

## II. Master's thesis details

Master's thesis title in English:

**Improving the HardNet Descriptor**

Master's thesis title in Czech:

**Vylepšení HardNet deskriptoru**

Guidelines:

Two view matching and registration is the problem of finding geometric relations between pairs of images. It has many applications 3D reconstruction, SLAM, medical image processing, in autonomous driving, etc. An integral part of most image matching algorithms is a local feature descriptor, which transforms an image patch to a compact vector representation. The current state-of-the-art descriptor in many standard benchmarks is based on the HardNet [1] convolutional neural network (CNN).

The objective of the thesis is to improve the HardNet and/or the image matching pipeline using one or more of the following:

- use of multiple training datasets
- exploiting different sampling procedures of (hard) examples during training
- modification of the HardNet CNN architecture, e.g. by the method [2]
- considering alternative loss functions, e.g. [3]
- adapting HardNet to specific detectors for selection of keypoints
- combination of global or local image pre-preprocessing

The proposed improvements will be evaluated on a standard benchmarks for image matching in various domains, e.g. [4].

Bibliography / sources:

[1] Mishchuk et al, Working hard to know your neighbor's margins: Local descriptor learning loss, NIPS 2017

[2] Liu et al. DARTS: Differentiable Architecture Search. ICLR 2019

[3] Revaud et al. Learning with Average Precision: Training Image Retrieval with a Listwise Loss. ICCV 2019

[4] Balntas et.al. HPatches: A benchmark and evaluation of handcrafted and learned local descriptor. CVPR 2017

Name and workplace of master's thesis supervisor:

**Mgr. Dmytro Mishkin, Visual Recognition Group, FEE**

Name and workplace of second master's thesis supervisor or consultant:

Date of master's thesis assignment: **19.12.2019** Deadline for master's thesis submission: **22.05.2020**

Assignment valid until: **30.09.2021**

\_\_\_\_\_  
Mgr. Dmytro Mishkin  
Supervisor's signature

\_\_\_\_\_  
doc. Ing. Tomáš Svoboda, Ph.D.  
Head of department's signature

\_\_\_\_\_  
prof. Mgr. Petr Páta, Ph.D.  
Dean's signature

### III. Assignment receipt

The student acknowledges that the master's thesis is an individual work. The student must produce his thesis without the assistance of others, with the exception of provided consultations. Within the master's thesis, the author must state the names of consultants and include a list of references.

\_\_\_\_\_  
Date of assignment receipt

\_\_\_\_\_  
Student's signature



## Acknowledgements

I wish to express my sincere appreciation to my supervisor Dmytro Mishkin for his support and patient guidance. He educated me in the fields of computer vision and helped me pursue my goals. I would also like to thank Jiří Matas for his enthusiastic encouragement and for giving me valuable advice in academia as well as in day-to-day life. I would like to extend my thanks to all members of CMP for creating a friendly and inspiring environment. Finally, I would like to thank Nicole for her love and kind support during my studies.

## Declaration

I declare that the presented work was developed independently and that I have listed all sources of information used within it in accordance with the methodical instructions for observing the ethical principles in the preparation of university theses.

Prague

22 May 2020

.....  
signature

## Abstract

We consider the problem of local feature descriptor learning for wide baseline stereo focusing on the HardNet descriptor, which is close to state-of-the-art. AMOS Patches dataset is introduced, which improves robustness to illumination and appearance changes. It is based on registered images from selected cameras from the AMOS dataset. We provide recommendations on the patch dataset creation process and evaluate HardNet trained on data of different modalities. We also introduce a dataset combination and reduction methods, that allow comparable performance on a significantly smaller dataset.

HardNet8, consistently outperforming the original HardNet, benefits from the architectural choices made: connectivity pattern, final pooling, receptive field, CNN building blocks found by manual or automatic search algorithms – DARTS. We show impact of overlooked hyperparameters such as batch size and length of training on the descriptor quality. PCA dimensionality reduction further boosts performance and also reduces memory footprint.

Finally, the insights gained lead to two HardNet8 descriptors: one performing well on a variety of benchmarks – HPatches, AMOS Patches and IMW Phototourism, the other is optimized for IMW Phototourism.

**Keywords:** local feature descriptor, HardNet, image matching, convolutional neural network, illumination robust, dataset creation, dataset reduction, combining datasets, architecture search, loss function, compression of embeddings

**Supervisor:** M.Sc. Dmytro Mishkin  
Czech Technical University in Prague,  
Faculty of Electrical Engineering,  
Department of Cybernetics,  
Karlovo namesti 13,  
121 35 Prague 2,  
Czech Republic

## Abstrakt

Uvažujeme problém učení lokálního deskriptoru pro wide baseline stereo, zaměřujeme se na deskriptor HardNet, který je blízko state-of-the-art. Představujeme datový soubor AMOS Patches, který zlepšuje odolnost vůči změnám osvětlení a vzhledu. Je založen na registrovaných obrázcích z vybraných kamer z datového souboru AMOS. Dáváme doporučení k procesu vytváření datových sad obrázkových výřezů a testujeme HardNet naučený na datech různého charakteru. Představujeme také metody pro kombinaci datových sad a jejich redukci, díky kterým dává učení na výrazně méně datech srovnatelné výsledky.

HardNet8, který překonává původní HardNet, těží z provedených rozhodnutí o architektuře: schéma propojení, final polling, receptivní pole, stavební bloky CNN nalezené pomocí manuálních nebo automatických vyhledávacích algoritmů - DARTS. Ukazujeme vliv často přehlížených hyperparametrů, jako je velikost batche a délka tréninku, na kvalitu deskriptoru. Komprese výstupů sítě pomocí PCA dále zlepšuje výsledky a také snižuje paměťovou náročnost.

Na základě získaných poznatků představujeme dvě varianty HardNet8 deskriptoru: jeden má dobré výsledky na HPatches, AMOS Patches a IMW Phototourism, druhý je optimalizován pro IMW Phototourism.

**Klíčová slova:** deskriptor bodu zájmu, HardNet, registrace obrázků, konvoluční neuronová síť, robustnost ke změně osvětlení, tvorba datasetu, redukce datasetu, kombinování datasetů, hledání architektury, ztrátová funkce, komprese vektorové reprezentace

**Překlad názvu:** Vylepšení HardNet deskriptoru

# Contents

<b>1 Introduction</b>	<b>1</b>
<b>2 Related work</b>	<b>5</b>
<b>3 HardNet Descriptor</b>	<b>7</b>
3.1 Local Descriptor Evaluation Methods	8
<b>4 Improving HardNet</b>	<b>11</b>
4.1 Datasets	11
4.1.1 Existing Datasets	11
4.1.2 AMOS Patches	12
4.1.3 Experiments with AMOS Patches	22
4.1.4 Experiments with IMW Phototourism	24
4.1.5 Combining Datasets	26
4.2 Architecture	28
4.2.1 Automated Search	28
4.2.2 Manual Search	30
4.2.3 Final Pooling	32
4.3 Compression of Embeddings	33
4.4 Loss Function	34
4.5 Hyperparameters	35
<b>5 Evaluation on Standard Benchmarks</b>	<b>39</b>
<b>6 Conclusion</b>	<b>41</b>
<b>Bibliography</b>	<b>43</b>

## Figures

1.1 The image matching pipeline. . . . .	1	4.16 Impact of the embeddings compression. . . . .	34
1.2 A subset of the AMOS Patches dataset.	3	4.17 HardNet performance depending on the margin applied during training. . . . .	35
3.1 HardNet descriptor architecture. . . . .	7	4.18 HardNet performance depending on the number of positives sampled in a batch during training. . . . .	36
3.2 Pipeline of the IMW PT benchmark.	8	4.19 HardNet performance depending on the training batch size. . . . .	37
4.1 The AMOS dataset. . . . .	13	4.20 Evaluation of HardNet trained for different number of epochs. . . . .	37
4.2 The AMOS Patches extraction pipeline. . . . .	14		
4.3 MODS registration failure example.	16		
4.4 AMOS dataset manually pruned views. . . . .	16		
4.5 HardNet mAP score in HPatches matching task evaluated for different sizes of AMOS patches training dataset	19		
4.6 Deterministic selection of feature points for AMOS Patches. . . . .	19		
4.7 HardNet mAP score in HPatches matching task as a function of the number of source views for a batch. . . . .	21		
4.8 Comparison of two approaches to sampling of patches. . . . .	22		
4.9 Visualization of a strategy of sampling of pairs of patches enforcing patch overlap to be small. . . . .	23		
4.10 The mAP score of HardNet trained on AMOS patches displaced by different shifts. . . . .	24		
4.11 Evaluation of HardNet trained on subsets of AMOS Views. . . . .	25		
4.12 Visualization of the loss function values on AMOS Views images at the end of training. . . . .	26		
4.13 Evaluation of the dataset reduction method. . . . .	27		
4.14 Modules found by the DARTS algorithm. . . . .	29		
4.15 ResNet modules. . . . .	32		


## Tables

4.1 Patch sampling: Influence of the response function on HPatches matching score (mAP). . . . .	17
4.2 Patch sampling: Influence of the image averaging on HPatches matching score (mAP). . . . .	18
4.3 Patch sampling: Influence of the response averaging on HPatches matching score (mAP). . . . .	18
4.4 Hyperparameter search for a deterministic selection of patches. . . . .	20
4.5 AMOS Patches mAP scores of HardNet trained on data of different modality . . . . .	24
4.6 Evaluation of dataset reduction approaches. . . . .	25
4.7 Training datasets comparison. . . . .	27
4.8 Combining datasets: Influence of the training dataset(s) and combining methods on the measured metrics. . . . .	28
4.9 Evaluation of models found by the DARTS algorithm. . . . .	29
4.10 Variants of the HardNet architecture	30
4.11 Evaluation of the variants of the HardNet architecture . . . . .	30
4.12 ResNet architectures . . . . .	31
4.13 Evaluation of the ResNet architectures. . . . .	31
4.14 Evaluation of HardNet variants with different input size. . . . .	32
4.15 Evaluation of final pooling variants.	33
5.1 Final evaluation of the proposed descriptors. . . . .	39

# Chapter 1

## Introduction

We consider the wide baseline stereo, where the goal is to find correspondences between pixels and regions and estimate the relative camera pose from one image to the other. The list of applications is vast, it can be applied in the tasks of structure from motion (SfM) [53, 47], image retrieval [55], relocalization [69, 51], simultaneous localization and mapping, self-localization and in applications like autonomous driving [12] and panorama stitching [6]. The pipeline consists of several steps. First, feature points are detected in the two input images, then small image patches are cropped around these points. Patches are then normalized with respect to rotation, scale and possibly other factors like contrast. A descriptor maps to a vector representation for each of the patches. Subsequently tentative correspondences are created between the two sets of descriptors via nearest neighbor matching in the descriptor space and verified by RANSAC [10] and the geometry relation is then obtained from the set of inliers [17]. Two popular choices of an image relation model are homography – for the planar surfaces – and epipolar geometry, which are represented by essential or fundamental matrices [18].



**Figure 1.1:** The wide baseline stereo pipeline. The stage, which is studied in current work is highlighted. Image credit: CTU Computer Vision Methods course <sup>1</sup>.

A local feature descriptor is a mapping from a (small) image to a metric space, which is robust to changes in acquisition conditions so that descriptors related to the same 3D points are similar and dissimilar otherwise. A descriptor is typically used to establish tentative correspondences between detected local features like blobs, corners or more complex structures. Performance of the whole WBS pipeline relies heavily on the qualities of the descriptor, namely its performance, speed, universality and robustness to changes in acquisition conditions.

<sup>1</sup><https://cw.fel.cvut.cz/wiki/courses/mpv/start>

The earliest descriptors were handcrafted. One of the most famous is SIFT [33], which computes a histogram of weighted gradient orientations to form a 128-dimensional descriptor. Other descriptors [65, 15] compute the relative ordering of pixel values. Recent years have witnessed a noticeable effort to move to descriptors obtained by deep learning [38, 62]. Existing work explore possible architectures [62], loss functions [5, 38, 20, 25] and improvements of robustness to viewpoint changes by introducing large scale datasets from 3D reconstruction [43, 34].

In some works there has been an effort to train detector and descriptor as a single model [45, 48, 14]. Although these methods were able to perform well in the most recent benchmark [24], they are still worse than the combination of handcrafted DoG detector (used in SIFT) and HardNet-like descriptor. We focus on improving the descriptor part, namely using the HardNet architecture [38] with the triplet margin loss function.

Regarding acquisition conditions, we refer to [40] for an introduction to wide-baseline matching and an overview of 4 main factors - sensor type, geometry, illumination, and appearance - which influence the performance of the descriptor. In this work we focus on the last three, that is we want to improve robustness to changes in the intensity, location and direction of light sources; to change in camera pose, its scale and resolution; to season-related changes and changes in visibility in the scene (fog, snow).

Robustness to illumination and appearance changes has received comparatively little attention, yet it is a bigger challenge for modern descriptors [41, 4]. Their performance degrades significantly if the images are taken at different times of day, different seasons or weather conditions [64]. We tackle this problem by leveraging information from 24/7 webcams located worldwide [23, 22].

We make the following contributions. In Section 4.1.2 we present a method for selecting suitable static views for the training of a descriptor together with a fixed number of the most representative images from a (too) large collection of camera feeds. We also introduce AMOS Patches<sup>2</sup> - extracted image crop-outs forming sets of veridical patch correspondences, see Figure 1.2. This dataset improves robustness to changes in illumination and appearance of the trained descriptor.

In the second part of Section 4.1.2 we focus on further improving the HardNet [38] descriptor by making adjustments to data preparation and training process. In Sections 4.1.3 and 4.1.4 we present experiments with AMOS Patches and IMW Phototourism [24]. They are related to data modality and the number of source views. We also introduce a method for dataset reduction and compare a few ways of combining datasets.

In Section 4.2 we conduct a systematic study on the influence of different architectural and choices on the descriptor performance: final pooling, architecture style, receptive field size.

In Section 4.3 we explain our approach to compressing network outputs that

---

<sup>2</sup>The dataset and contributing images are available at [https://github.com/pultarmi/AMOS\\_patches](https://github.com/pultarmi/AMOS_patches)




increases the descriptor performance.

Section 4.4 presents experiments and choices related to the loss function.

Hyperparameters are set in section 4.5.

We summarize our findings in Section 6.

Part of the following work is based on the published paper [46]. The main author is also author of this thesis.



**Figure 1.2:** A subset of the AMOS Patches dataset originating from 24/7 webcams located world-wide. Each column contains 20 corresponding patches.



## Chapter 2

### Related work

The literature on local feature descriptors is vast [11]. The earliest descriptors were handcrafted. One of the most widely used is SIFT [33], which uses a weighted histogram of gradient orientations and is moderately robust to local deformations and intensity changes[37]. DSP-SIFT [13] is a modification, which uses also pooling of gradient orientations across different scales.

The family of order-based descriptors like LIOP [65] or MROGH [15] operates on the relative order of pixel intensities in the patch instead of on the intensities themselves. Relative order (sorting) is invariant to any monotonically increasing intensity transformation. Descriptors like SymFeat [19], SSIM [54] and learned DASC [27] encode local symmetries and self-similarities.

There are two main ways towards achieving robustness to illumination change: by descriptor construction and by learning on the appropriate dataset. Normalization of the patch mean and variance is a simple but powerful method, which is implemented in both SIFT [33] and modern learned descriptors [62, 38]. It makes the descriptor invariant to affine changes in pixel intensities in the patch. Photometric normalization might be more complex than mean-var normalization, e.g. as done by the learned RGB2NIR [67] or hand-crafted LAT [50]. HalfSIFT [26] treats opposite intensity gradient directions as equal, trading off half of the SIFT dimensionality for being contrast reversal invariant. It works well in medical imaging and infrared-vs-visible matching.

Binary descriptors were developed to fill the demand for compact and fast representations, although they have usually worse performance than the real-valued ones. BRIEF [7] and BRISK [29] construct a descriptor based on comparison of pairs of randomly selected pixels. ORB [49] is an improvement of BRIEF with an estimation of patch orientation using FAST local feature point detector. The similarity between descriptors is computed efficiently with Hamming distance. L2 distance is commonly used in the case of real-valued descriptors.

Current research on local feature descriptors is shifting towards those obtained by deep learning. Significant advancement has been made in recent years [4, 24]. In the following text we focus on real-valued ones. A simple three-layer network was used in [57] together with a stochastic sampling of hard examples in a metric learning ap-

proach. L2Net [62] introduced relatively simple convolutional network architecture, which outputs a vector of length 128. HardNet [38] adopted the L2Net architecture and introduced the triplet loss function with hard-in-batch sampling. GeoDesc [34] uses geometry constraints from multi-view reconstructions. SOSnet [63] introduced second order similarity regularization, which is added to the loss function.


Learned descriptors for multimodal data, e.g. infrared-vs-visible mostly include Siamese convolution networks with modality-specific branches, like the Quadruplet Network [3]. The decision which branch to use for a specific patch comes from an external source or a domain classifier. HNet [32] uses an auto-encoder network and style transfer methods like CycleGAN [70] for emulating different modalities.

There are many datasets potentially suitable for training and testing of local feature descriptors. Most of them comprise images only, i.e. without patch-level correspondences. DTU Robot [2] contains photos of objects in a lab environment, camera positioning is operated by a robotic arm, light sources are artificial. RobotCar dataset [35] contains camera feed from a vehicle, which was repeatedly traversing a predefined route in central Oxford. TILDE [64] dataset comprises images from 6 scenes, it was made for training and evaluation of illumination-robust detector. Other datasets are OxfordAffine [37], Aachen Day-Night [52], GDB [66], SymBench [19], etc. However, none of these datasets contain also keypoint correspondences, which are needed for the extraction of patches.

Despite the importance of the topic, the number of patch-level datasets is small, especially for training and testing illumination-robust descriptors. UBC Phototour dataset [6] contains patches from 3 scenes, it covers more viewpoint related changes. HPatches dataset [4] contains 116 scenes, each represented by 6 images. It is the only patch-based dataset we know of which contains sequences with illumination changes. 57 scenes contain 3D objects under varying photometric conditions, while 59 scenes are planar with changes in camera pose. The PhotoSynth (PS) dataset [43] comprises 30 scenes reconstructed by the COLMAP [53] SfM software. Although the scenes contain photometric changes, the trained descriptor achieves better results only on the viewpoint split of HPatches.

## Chapter 3

### HardNet Descriptor



**Figure 3.1:** HardNet descriptor architecture. Image credit: [38].

We have selected the HardNet descriptor [38] as an object to study and improve given its good performance in various benchmarks [24]. HardNet descriptor is a real-valued convolutional network with the output vector of length 128 and the input patch of size  $32 \times 32$  pixels. The architecture adopts the VGG [58] style, it is composed of convolutional layers, each followed by batch normalization and ReLU activation. Dropout with  $p = 0.3$  is used. The inputs to the network are mean and variance normalized.

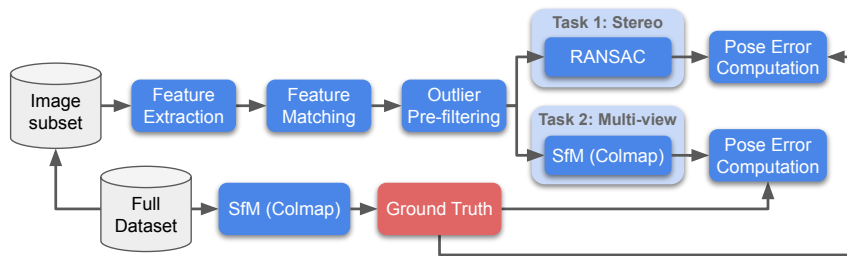
The task of learning a local feature descriptor is formulated as a metric learning problem, where the goal is to learn an embedding for the input patches such that corresponding (positive) patches are closer than the noncorresponding (negative) ones. For training we use the hard-in-batch triplet margin loss [38]. Let  $(p_{i,1}, p_{i,2}), i = 1, \dots, n$  be pairs of corresponding patches in a batch of size  $n$ . Let  $d$  be the Euclidean norm. The loss function is the following:

$$L = \frac{1}{n} \sum_{i=1, n} \max(0, M + d_{\text{pos}}^i - \min(d_{\text{neg1}}^i, d_{\text{neg2}}^i)), \quad (3.1)$$
$$d_{\text{pos}}^i = d(p_{i,1}, p_{i,2}),$$
$$d_{\text{neg1}}^i = \min_{j=1, \dots, n, j \neq i} d(p_{i,1}, p_{j,2}),$$
$$d_{\text{neg2}}^i = \min_{j=1, \dots, n, j \neq i} d(p_{j,1}, p_{i,2}),$$

where  $M=1$  is called the margin. By minimizing the loss we aim to minimize

the distance between the positives while maximizing the distance to the hardest negatives from both positive patches. The loss is zero if the difference between the latter and the former distance is at least  $M$ .

### 3.1 Local Descriptor Evaluation Methods



**Figure 3.2:** Pipeline of the IMW PT benchmark. The input is a set of images. It detects keypoints, extracts features and computes matches. Camera pose is then computed after running outlier filtering and Random sample consensus (RANSAC). Image credit: [24].

There are two types of evaluation. One way is to evaluate either solely the descriptor using pre-extracted patches, the second way is to evaluate some end-goal-metric of the whole image matching pipeline by plugging in the descriptor and computing the precision of the pipeline as a whole. An example of end-goal-metric is camera pose accuracy. We describe and use both approaches.

The important property of the first approach is zero dependence on other parts of the image matching pipeline. The procedure is generally simpler, faster and we can make clear judgments on the descriptor. On the other hand, it provides us limited information on how the image matching pipeline will perform with the descriptor plugged in. One of the used metrics is a false positive rate (FPR), which measures the percentage of false positives at a given recall level. Specifically, FPR95 - FPR at 0.95 recall – is usually used in UBC Phototour. It is defined as

$$\text{FPR95} = \frac{\text{FP}}{\text{FP} + \text{TN}}, \text{ s.t. } \frac{\text{TP}}{\text{TP} + \text{FN}} = 0.95, \quad (3.2)$$

considering ascending ordering of samples according to assigned confidence score. We observed this metric has little correlation with the actual performance of the descriptor and so its usability is limited to a sanity check of convergence during the training of a model.

Another metric is mean average precision (mAP), formulated as

$$\begin{aligned} \text{mAP} &= \frac{1}{Q} \sum_{i=1}^Q \text{AP}_i, \\ \text{AP}_i &= \sum_{j=1}^{N_i} P(s_{i,j}) \nabla r(s_{i,j}), \end{aligned} \quad (3.3)$$

where  $s_{i,j}$  is the  $j$ -th largest assigned confidence score in query  $i$ ,  $\nabla$  denotes the difference in recall between samples  $s_{i,j}$  and  $s_{i,j-1}$ ,  $P$  denotes precision, considering samples  $s_{i,1}, \dots, s_{i,j}$ .  $\text{AP}_i$  values are averaged over  $Q$  queries, where each query contains  $N_i$  samples. This metric is measured in Hatches [4] in three tasks: verification, retrieval and matching. In verification the descriptor is used to assign a score for a pair of patches indicating whether they correspond. In retrieval there is a query patch and the descriptor is used to retrieve the corresponding patch from a given set of patches which contains several positives as well as a big number of negatives. In matching, which is most relevant for the image matching pipeline, the task is to find the corresponding patch given a query patch and a set of patches which contains a single positive.

The second approach to evaluate a descriptor is more connected to its typical application. In IMW Phototourism (IMW PT) [24] the image matching pipeline, see Figure 3.2, is used to estimate the relative pose between two cameras. Then the benchmark computes the angular error between the estimated and ground-truth rotation and translation vectors. A threshold is used to a binary result, indicating whether the estimate is good enough within the tolerance limit, and the average accuracy is calculated over all camera pairs. By integrating over a set of thresholds from  $0^\circ$  to  $X^\circ$ , we obtain mean average accuracy  $\text{mAA}(5^\circ)$  for  $X=5$  and  $\text{mAA}(10^\circ)$  for  $X=10$ .

In IMW PT [24] we use 8k Difference of Gaussians (DoG) keypoints, DEGEN-SAC [9] for model estimation and "both" matching strategy, i.e. with the requirement that matches are mutually nearest. FLANN library [44] is used for approximate nearest neighbour search.





## Chapter 4

# Improving HardNet

### 4.1 Datasets

The quality of the learned local descriptor depends on the data it is trained on. There are a variety of datasets [2, 35] of image sequences suitable for descriptor training, but only a few of them contain also keypoint correspondences, which are needed for the extraction of patches. In the current section we will describe existing datasets and the one we present.

#### 4.1.1 Existing Datasets

**UBC Phototour [61, 16].** UBC Phototour dataset<sup>1</sup> (also called Brown) comprises 3 scenes: Liberty, Notredame and Yosemite. Point cloud 3D reconstructions and dense depth maps were constructed from original images with the help of Bundler [60] structure-from-motion library to obtain sets of corresponding patches. Each scene was reconstructed using Difference of Gaussian and Harris feature point detectors. Liberty sequence has been often been used as a single dataset for training[38]. It was used for testing as well, but with the rise of learned descriptors, the usually used FPR95 metric is already under 1%, making this benchmark saturated.

**HPatches [6].** This dataset<sup>2</sup> is a part of a benchmark for local feature descriptors. It contains 116 scenes, 6 images each, which are split to illumination (57 scenes) and viewpoint subset, based on the main nuisance factor. In the illumination split the scene is captured by a static webcam during a different time of the day and the viewpoint split contains significant camera pose changes. Images are related by a homography to the reference image. DoG, Hessian-Hessian and Harris-Laplace detectors were used to extract keypoints. The smallest patch size in the reference image is 16x16 pixels, while the extracted patch size is 65x65. Patches are further divided into difficulty levels (easy, hard and tough) according to the amount of

<sup>1</sup>available at <http://matthewalunbrown.com/patchdata/patchdata.html>

<sup>2</sup>available at <https://github.com/hpatches/hpatches-dataset>

synthetic noise added before projection. The applied transformation is composed of random rotation, translation, isotropic and anisotropic scaling.

**The PhotoSynth Dataset [43].** The PhotoSynth (PS) dataset<sup>3</sup> contains patches from 30 scenes with illumination and viewpoint changes. COLMAP [53] SfM software was used to create 3D reconstructions with verified patch correspondences. Each patch is scale and rotation normalized. Minimum crop size is  $20 \times 20$  pixels, maximum is  $128 \times 128$  pixels. The cropped patches are resized to  $48 \times 48$ px. Around 400 000 patches were extracted from each scene. The trained descriptor achieved state-of-the-art performance at the time on the HPatches benchmark in the viewpoint split. It however sacrifices robustness to photometric changes.

**IMW Phototourism 2020 [24].** This dataset<sup>4</sup> was recently created for the upcoming CVPR Image Matching Workshop and is a part of extensive evaluation software. Currently it contains 13 scenes for training, 3 for validation and 10 for testing. COLMAP [53] SfM software was used to create 3D reconstructions with verified patch correspondences. Each folder with a 3D reconstructed scene contains 2D to 3D point correspondences. It does not provide information about scale and rotation and the extracted patches cannot be easily normalized. However, we have observed that it is not a big issue in terms of the performance of the trained descriptor. Because the images are obtained from Phototourism data, it covers some changes in weather and lighting. The number of extracted patches is proportional to the total number of keypoints in all images and varies highly across the scenes. E.g. Brandenburg Gate contains over 2 million of patches, while Trevi Fountain contains over 15 million.

#### 4.1.2 AMOS Patches

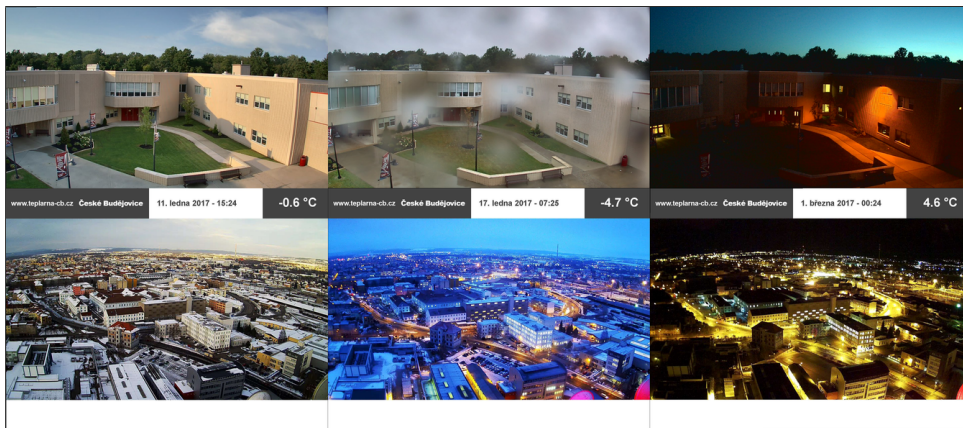
There was little training data available for improving local descriptor robustness to photometric changes. We propose to fill this gap by introducing AMOS Patches dataset. It is created from the AMOS collection of world-wide static webcams and covers various illumination and season-related acquisition conditions.

AMOS [22, 23] is a continuously growing publicly available dataset collected from outdoor webcams, currently containing over one billion (or 20 TB) images. It is organized into individual camera directories, which are split into folders according to the year and month of the acquisition. The size of the images varies, and so does their quality and the number of images in each camera directory. A typical AMOS camera is static and has approximately size of  $300 \times 300$  pixels. Many cameras store images in all seasons and during the whole day.

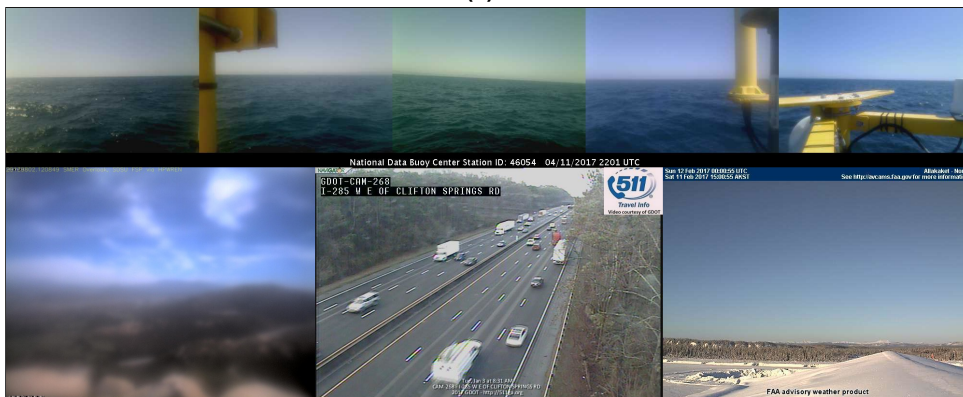
The advantage of static cameras lies in the fact that they show the same structure under different weather and lighting conditions. Therefore, if observing a static scene, they *seem to be* highly suitable for training of local feature descriptor robust to illumination and appearance changes.

<sup>3</sup>available at <https://github.com/rmitra/PS-Dataset>

<sup>4</sup>available at <https://github.com/vcg-uvic/image-matching-benchmark>



(a):




(b):

**Figure 4.1:** The AMOS dataset [22, 23] - example images from (a) cameras contributing to the AMOS patches set and (b) cameras unsuitable for descriptor training because of blur, dynamic content or dominant sky.

We learned the hard way that using this type of data is not trivial. Firstly, due to the dataset size, it is not feasible with moderate computing power to fit such data into memory. Moreover, preprocessing would take a prohibitive amount of time. Secondly, the training procedure is sensitive to the misregistration of the images and the presence and size of moving objects. Many cameras experience technical issues such as: being out of focus, rotating over time, displaying highly dynamic scene (e.g. sky, sea waves), which all significantly hurt the performance of the trained descriptor, as discussed later.

Therefore, we developed a pipeline for the creation of AMOS Patches, shown in Figure 4.2, which entails several steps to create a clean dataset with veridical patch correspondences. These methods focus on the selection of cameras and images, detection of view switching in a camera and the registration of images. Because of several not easily detectable problems, it was still necessary to perform a final manual check of the selected image sets.



**Figure 4.2:** The AMOS Patches extraction pipeline: camera selection to filter out dynamic or empty scenes, appearance clustering to remove redundant images, viewpoint reclustering to tackle switching cameras, precise registration for further filtering, manual pruning for the final selection of views and patch sampling. The numbers above denote the number of cameras in each stage.

**Camera selection.** The first step — camera selection — aims at choosing a subset of cameras that are suitable for training, i.e. do not produce very dark images, are sharp and do not display moving objects like cars or boats.

The procedure uses two neural networks, a sky detector [36] and an object detector [1], and computes simple statistics for each of 20 randomly chosen images in each camera directory.

The camera selection took approximately one week on a single PC (Intel Xeon CPU E5-2620) with one GPU GTX Titan X. Processing more images by the neural network detectors increases both the precision of the method and the running time. Our choice is therefore based on the (limited) available computation power.

Each image is then checked whether it satisfies the following conditions:

- $f_1$  : sky area  $< 50\%$  *not empty*
- $f_2$  : no detected cars or boats *not dynamic*
- $f_3$  :  $\text{Var}(\nabla^2 \text{ image}) \geq 180$  *sharp*
- $f_4$  : mean pixel intensity  $> 30$  *not black*
- $f_5$  : image size  $> (700, 700)$  *large*

A camera is kept if at least 14 out of the 20 images pass the check.

The filter  $f_5$  is the most restrictive, it removes 91% of the cameras – AMOS contains mostly low-resolution images. The reasoning behind using  $f_5$  is that images of smaller size often observe a motorway or are blurred. Also, such cameras would not generate many patches. We want to select only a relatively small subset of the cameras with the predefined characteristics and therefore an incorrect removal of a camera is not a problem.

Several cameras were removed because of corrupted image files. The resulting set contains 474 camera folders which were subject to subsequent preprocessing.

**Appearance clustering by K-means.** The resulting data is of sufficient quality, but it is highly redundant: images shot in 10 minute intervals are often indistinguishable and very common. To select sufficiently diverse image sets, we run the K-means clustering algorithm with  $K=120$  to keep the most representative images. We use the fc6 layer of the ImageNet-pretrained AlexNet [28] network as the global image descriptor. While not being the state-of-the-art, AlexNet is still the most effective architecture in terms of speed [8], with acceptable quality.

At this stage of the pipeline, there are  $K=120$  images for each of the  $C=474$  cameras selected, a feasible number for training with the computational resources available.

Feature descriptor training with patches selected from this image set was not successful. We were unable to achieve accuracy higher than 49.1 mean average precision (mAP) in the HPatches matching task; the state-of-the-art is 59.1 mAP – GeoDesc [34].

**Viewpoint clustering with MODS.** After examining the data closely, we found that many of the cameras switch between a few views, which breaks our assumption for the generation of ground truth correspondences via identity transformation. In order to filter out the non-matching views, we run MODS [39], a fast method for two-view matching, and split each camera folder into clusters, called views, by applying a threshold on the number of inliers and the difference between the homography matrix and the identity transform.

Let  $(x_1, x_2, \dots, x_K)$  be a set of images in a camera folder in arbitrary order. MODS matching is first run on pairs  $(x_1, x_2), (x_1, x_3), \dots, (x_1, x_K)$ . Image  $x_1$  becomes the reference image in a newly created view, which contains  $x_i$  for which the registration yields more than 50 inliers and  $\text{SAD}(H(x_1, x_i), I_3) < 50$ . SAD denotes the sum of absolute differences,  $H$  denotes a homography matrix normalized by the element in position (3, 3),  $I_3$  is 3x3 identity matrix. All images in the created view are then removed from the processed image set. The step is repeated until no images remain.

We observed that the number of the resulting views in one camera folder depends on phenomena other than camera movement. For example, in cases where there is a fog or very rainy weather, MODS fails to match most of the image pairs and many of them form a single element cluster, which is excluded from further processing. For each camera, we keep only the view with the largest number of images, if it has more than 50. Each remaining view is reduced to 50 images by random selection.

**Registration with GDB-ICP.** While the MODS method is effective in matching and subsequent reclustering of camera sequences, in most cases the estimate of the global homography is not sufficiently precise. MODS often outputs a homography valid for only a small area in the image, see the example is shown in Figure 4.3. Therefore, the views contain also images that are not correctly aligned. To alleviate the problem, we run Generalized Dual Bootstrap-ICP [66] to prune the set of views, keeping those where this second registration is successful.


The registration proceeds as follows. Each view folder contains images  $(x_1, x_2, \dots, x_{50})$ ,



#### 4. Improving HardNet

where image  $x_1$  is the MODS reference. The GDB-ICP registration is run on pairs  $(x_1, x_2), (x_1, x_3), \dots, (x_1, x_{50})$  and warped images  $x'_2, x'_3, \dots, x'_{50}$  are obtained. If registration fails on any pair, the whole view is removed.

After the precise registration with GDB-ICP, 151 views remained. It is feasible to manually inspect such a set.



**Figure 4.3:** MODS registration failure, most of the correspondences are on moving structures. Top: an image pair with marked inliers. Bottom: misregistered image (left) and the reference.

**Manual Pruning.** A few problems remain, see Figure 4.4, such as dynamic scenes, undetected sky (the sky detector fires mostly on the clear blue sky). As a precaution, we also removed views with very similar content and views from different cameras observing the same place from a different viewpoint. We tried to use the scene segmentation network [68] to detect moving objects, but the result was not satisfactory. The final selection is therefore done by hand, resulting in a set of 27 folders with 50 images each, which we call AMOS Views. In this step we also select another 7 views which form a test set in our AMOS Patches matching benchmark.



**Figure 4.4:** AMOS dataset manually pruned views. Examples of dynamic scenes (left, center) and a cloud-dominated scene not removed during camera selection (right).

#### ■ Extracting AMOS Patches

The last phase of the AMOS Patches pipeline consists of sampling images to obtain patch centers, scales and angles, and subsequent cropping from source images.

Here we explore two approaches. In the first one we select patch centers and scales randomly based on a probability map. The probability map construction is discussed below. In the second approach we deterministically pick local maxima in the Hessian detector response map. In both cases we use each selected point to extract a set of corresponding patches of size  $96 \times 96$  pixels. Each patch set contains 50 corresponding patches, i.e. one per each image in a view. Here we assume the images in each view are well registered, meaning that a mapping between image coordinates is (almost) identity. The assumption holds due to previous parts of the pipeline. The number of patch sets is manually set in the first approach, and semi-automatically determined in the second approach.

In the following experiments we train for 10 "epochs", 5 million pairs of samples each, with batch size = 1024, learning rate = 20, SGD optimizer with momentum = 0.9 and a linear learning rate decay schedule unless stated otherwise. We use "epoch" not in a conventional sense as "going through all examples in the dataset," but in the sense of "going through all generated pairs."

### ■ Sampling of Feature Points

Here we explore a sampling of feature points based on a probability map. There are many ways how to construct the probability distribution. First, it relies on the choice of a response function which is evaluated for each position in an input image. Secondly, it depends on the choice of input images for the response function and how we average the obtained results. In the following text we explain these choices in more detail and find the best setting. Once the random distribution is computed, scales and angles are chosen randomly and the patches are extracted.

**Response Function.** First we decide on the response function. In Table 4.1 we compare several possibilities. The differences in mAP scores are not large. The best mAP is achieved by the square root of the determinant of the Hessian matrix. If we do not extract the root, the result is slightly worse. The worst results are achieved by the uniform distribution together with the square root of the determinant of the Hessian matrix and nonmax suppression (NMS). We choose the  $\sqrt{\text{Hessian}}$  as the optimal response function. In the next step we explore two ways how we can perform the sampling procedure: by averaging before or after computing the response.

Weighting	mAP
Uniform	56.20
Hessian	56.39
$\sqrt{\text{Hessian}}$	<b>56.49</b>
NMS( $\sqrt{\text{Hessian}}$ )	56.18

**Table 4.1:** Patch sampling: Influence of the response function on HPatches matching score (mAP). By "Hessian" we denote shortly the determinant of the Hessian matrix.

**Representative Image.** The first way of sampling of the patches comprises acquiring a representative image, on which we subsequently evaluate the chosen response function - the output is then normalized and is used as the random distribution. We compare three possibilities: random, mean and median, see Figure 4.2. "Random" refers to choosing a single image from the dataset. In "mean" and "median" we apply the corresponding averaging over all images in a view. One can see that all three choices are comparable. "mean" averaging gives the best result by a small margin.

Image	mAP
random	56.49
median	56.44
mean	<b>56.58</b>

**Table 4.2:** Patch sampling: Influence of the image averaging on HPatches matching score (mAP). Weighting function is  $\sqrt{\text{Hessian}}$ .

**Averaging the Responses.** In this second way of sampling, we first evaluate the response function for all images in a view. The random distribution is then acquired by averaging all the responses or by choosing a single response (denoted as none), see Figure 4.3. The latter variant achieves the highest mAP score. It does not beat the best choice in the previous paragraph.

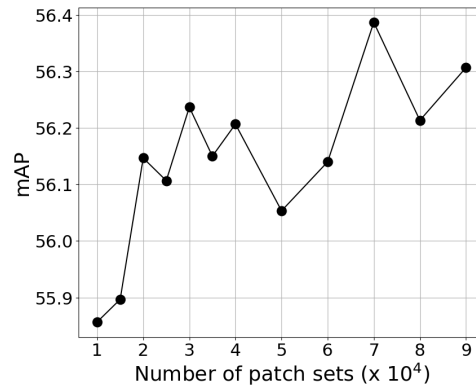
Averaging	mAP
none	<b>56.49</b>
mean	56.10
median	56.45

**Table 4.3:** Patch sampling: Influence of the response averaging on HPatches matching score (mAP). Weighting function is  $\sqrt{\text{Hessian}}$ .

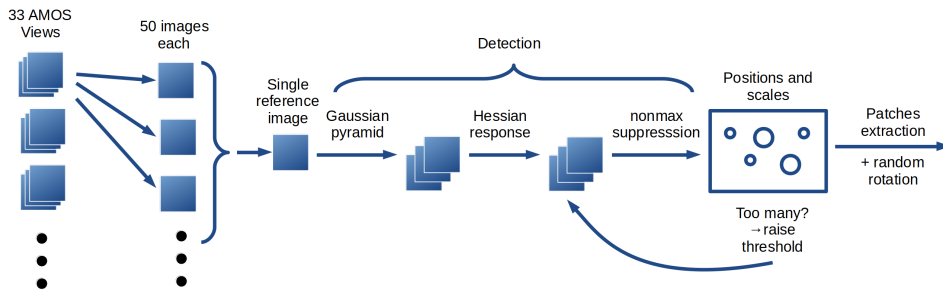
Overall, we conclude that the process of construction of a probability map is not as important as one may think. Although the images contain a lot of noise, choosing a random representative image is just as good as some of the more robust approaches.

**Number of Patch Sets.** We have determined the optimal way of sampling patches using a probability distribution, but we do not know what should be the size of the dataset. For this reason we created several datasets containing a different number of patch sets, from 10 000 to 90 000. Figure 4.5 shows that the performance of the trained descriptor roughly increases with the dataset size. We select extracting around 30 000 patch sets as a reasonable trade-off between performance and dataset size.





**Figure 4.5:** HardNet mAP score in HPatches matching task evaluated for different sizes of AMOS patches training dataset. Each value is an average over 3 different randomly generated datasets of the same size.



**Figure 4.6:** Part of the AMOS Patches extraction pipeline [46], which selects feature points in a deterministic way. Here we consider we already have filtered and registered set of camera views as input. The reference image is per-pixel median over the images. Detection part is fully described in section 4.1.2. In the end of the pipeline, patches are extracted from all 50 images using the same obtained set of positions and scales.

## ■ Deterministic Selection of Feature Points

Here we describe the second approach to extracting AMOS Patches. We noticed the previous method is somewhat distant from the common implementation in image matching, where feature points typically represent local maxima in the output of a response function. This approach is therefore very similar to the Hessian detector. One difference is that we do not use nonmax suppression in a 3D scalespace, but instead apply it for each scale in the image pyramid separately.

The input is a representative (e.g. median) image of a single AMOS View. From this input, a sequence of blurred images is created using incrementally larger Gaussian filter. Let  $\sigma^{k_1}, \sigma^{k_1+1}, \dots, \sigma^{k_2}, k_i \in \mathbb{N}_0, \sigma = 1.8$  empirically set, be chosen sequence of standard deviations of the normal distribution and the levels be indexed  $k_1, k_1 + 1, \dots, k_2$  according to the assigned sigma. The Gaussian kernel has width  $6\sigma^{k_i}$  in each level  $k_i$  and is convolved over the image. Output of this step is a sequence of images.

In the next step, Hessian kernel is convolved over all images, the response output

is multiplied by  $1/\sigma^{4k_i}$  in each level and then maxpooling is applied with kernel size equal to  $2\sigma^{k_i} \cdot bs$ , where  $bs$  denotes base scale and is usually around 16 (half of the patch size). Every local maxima under empirically chosen threshold = 0.00016 is discarded so that abundance of feature points is still generated per image. If too many feature points are generated, the threshold is multiplied by 2 and the algorithm starts over. We decided to keep at most 2000 feature points per input image (i.e. per single view).

The output of the described procedure is a set of locations in the representative image together with scales. Angles are chosen randomly for each patch set. As in the previous approach, the last step is the extraction of patches.

$s_{min}$	$s_{max}$	HPatches subset			AMOS Patches
		illum	view	full	
HardNet					
24	140	54.58	59.10	56.88	45.94
24	252	54.28	58.68	56.52	46.20
43	140	55.00	60.04	57.56	45.38
43	252	54.97	60.09	57.57	45.32
30	175	<b>55.32</b>	<b>60.16</b>	<b>57.78</b>	<b>46.43</b>
30	315	55.01	60.02	57.56	<b>46.43</b>
54	175	53.98	58.81	56.43	44.24
54	315	53.51	58.60	56.10	44.93
HardNet8					
24	140	57.69	62.61	60.19	47.37
24	252	57.73	62.62	60.22	47.47
43	140	58.02	63.76	60.94	46.33
43	252	57.99	63.48	60.78	46.26
30	175	<b>58.42</b>	<b>63.93</b>	<b>61.22</b>	47.60
30	315	58.27	63.69	61.02	<b>47.75</b>
54	175	55.79	61.06	58.47	44.87
54	315	55.84	61.06	58.49	45.11

**Table 4.4:** HPatches and AMOS Patches matching scores (mAP).  $s_{min}, s_{max}$  denote the smallest and largest patch scale, which is computed as  $s_{min} = bs \cdot \sigma^{k_1}, s_{max} = bs \cdot \sigma^{k_2}, \sigma = 1.8$ .  $bs$  denotes the base scale, which is also proportional to the kernel size of maxpool operator in detection of points of each scale.

**Hyperparameters.** In our implementation of the Hessian detector in AMOS Patches pipeline,  $k_1, k_2, bs$ , are the parameters which we search for in the ablation study, see Table 4.4. We pick  $k_1 = 0, k_2 = 4, bs = 30$ , corresponding to min and max patch scales  $s_{min} = 30, s_{max} = 175$ , as the optimal setting for which the mean of HPatches full mAP and AMOS test mAP scores are the highest. With this setting we outperform the previous, probability-based, approach to sampling.

## ■ AMOS Patches Test Set


In addition to the training data, we also create AMOS Patches test set, which is extracted from 7 pruned views, 10 images each. Keypoints are obtained using the SIFT detector. We consider the matching task and follow the HPatches protocol.

## ■ Sampling of Patches for a Batch

In this section we test different approaches to sampling of patches for a batch. The composition of a batch usually has an impact on the efficiency of the training procedure, because of the averaging of the gradients over the samples. We use the hard-in-batch structured loss function, where the composition furthermore has an influence on the distribution of the negative patches. Therefore, one may expect that the way of sampling to be important in our training. In this section we use probability map-based extraction of patches to create a dataset with 30 000 patch sets.

**Uniform Sampling.** In this way of sampling we select  $B$  patch sets for each batch. From each set we pick randomly two positive patches. Both selections are done using the uniform distribution. We train for 20 epochs with batch size = 256. The training dataset consists of AMOS Patches from 27 views in total. In Figure 4.7 we show the performance (mAP) as a function of the number of views we select from for each batch. Notice the monotone trend, which we interpret as follows. Reducing the number of views increases the number of negative patches from the same scene, which are often the most difficult to distinguish.

If we increase the batch size to 1024 and choose the number of source views = 6, then the model achieves 56.57 mAP on HPatches full split. Note that we did not use this batch size in Figure 4.7, because it would be higher than the minimal number of patch sets there is in a view, i.e. the trend would be obscured for # views = 1.



**Figure 4.7:** HardNet mAP score in HPatches matching task as a function of the number of source views for a batch. Views are selected randomly in each iteration. Dataset consists of 27 views in total.



**Figure 4.8:** Two approaches to sampling of patches. Left: using a probability map, right: using multiscale Hessian detector with scale-wise nonmaxima suppression. Both variants contain 186 patches. Patch centers and scales marked on single AMOS Views image.


**Sampling from Pairs of Images.** The previous approach is somewhat counter-intuitive since patches from many images from a single view may co-appear in a batch, creating a situation that does not occur in stereo matching. In the new strategy of sampling we run sequentially over all camera views until the batch is filled, selecting patches always only from a pair of images - and the selection of the pair is random so that all patches have equal chance to contribute. HardNet trained with such sampling approach achieves 56.42 mAP on HPatches. Therefore, contrary to our intuition, this approach is comparable to uniform sampling.

**Avoiding Collisions.** Another, in a sense orthogonal, approach to sample patches for a batch uses information about overlaps between patches generated from a camera view. During the sampling step, after selecting a sample from the pool of available patches, the pool is iterated and overlapping patches are discarded, see Figure 4.9. The idea of this method comes from the "first geometrically inconsistent (GI)" [39] criterion, which is used in image matching, where it applies a threshold on the ratio between the distance to the first and second keypoints, which considers a ratio between descriptors of patches, which are geometrically inconsistent, i.e. not overlapping. Our idea was to provide only such patches as negatives in the triplet loss.

The trained model achieves 55.70 mAP on HPatches and so this approach is inferior to the previous ones. Our hypothesis is that it may be due to preference of smaller – therefore containing less information – patches during sampling. Also, in real stereo matching patches do overlap (see Figure 4.8, right).

### ■ 4.1.3 Experiments with AMOS Patches

**The Importance of View Registration.** In this experiment we show the importance of the precise alignment of images. We displace each patch by different shifts and observe the influence on the HPatches matching score, see Figure 4.10. Note how the performance of the descriptor improves with a small shift, but then quickly deteriorates. We use the patch selection approach using probability map, #source views = 27 (all), 30000 patch sets and Hessian weighting without averaging.

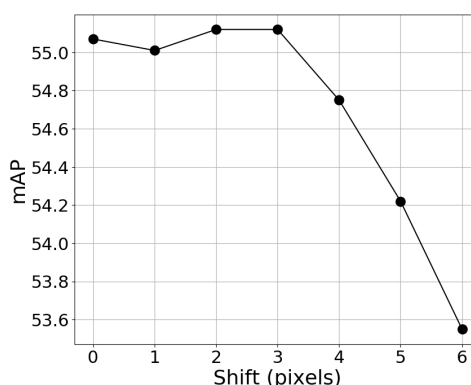


**Figure 4.9:** A strategy of sampling of pairs for a batch, in which we enforce patch overlap across the same scale to be small.

**Data Modality.** To the best of our knowledge, learned descriptors have usually been trained on grayscale data – Brown, PS and HPatches dataset contain only such. Here we experiment with other modalities. We extract RGB patches from AMOS Views and compare two architectures: HardNet8 and HardNet8 with learned color conversion. The latter model has prepended two more convolutional layers with kernel size = 1 which have 10 filters (first layer) and 1 filter (second layer) similar to [42]. We can see in Table 4.5 that this model improves the mAP score on the AMOS Patches test set. The difference is however not large and using RGB data has its drawbacks, for example bigger size of the dataset or incompatibility with other benchmarks (HPatches, UBC Phototour). We believe the improvement is not significant enough to compensate the demerits.

We also tried to make use of MegaDepth [30], a tool for single-view depth prediction, to enrich the modality our dataset. MegaDepth pretrained model was applied to each image in our AMOS Views dataset. The estimated depth maps were then during training concatenated with source images to form an input with two channels (denoted G+D). The result is significantly worse than the baseline.

**Number of AMOS Views.** During manual pruning in the AMOS Patches extraction pipeline we have the possibility to select up to 146 views. To check if more views are beneficial or not, we prepared 3 versions which are ordered with respect to subset relation. In this experiment we train for 20 epochs with batch size = 3072 and compare three architectures. One can see in Figure 4.11 that using a



**Figure 4.10:** Hpatches matching. The mAP score of HardNet trained on AMOS patches displaced by different shifts.

Modality	Model	mAP
G	HardNet8	49.09
G+D	HardNet8, input channels = 2	45.74
RGB	HardNet8, input channels = 3	47.56
RGB	HardNet8, learned color conversion	<b>49.31</b>

**Table 4.5:** AMOS Patches mAP scores of HardNet trained on data of different modality. G denotes grayscale, D denotes (estimated mono-)depth.

bigger training dataset does not necessarily improve performance. We make a similar observation also later in the case of IMW PhotoTourism data. Based on this result, we recommend the smaller subset for training of HardNet and the medium dataset for training of HardNet8.

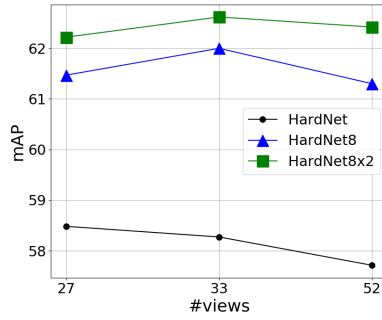
**Distribution of Loss Values.** We noticed that during training the average values of the loss function per patch set are distributed according to where the patch was extracted, see Figure 4.12. It holds approximately that patch sets with higher average loss are in parts of the image which view dynamic content such as trees, grass or sky. And vice versa, patch sets with low average loss are positioned on static content such as buildings or walkways.

#### 4.1.4 Experiments with IMW Phototourism

**Dataset Reduction.** As we mentioned earlier, some of the IMW PT datasets are very large. The extracted patches from Colosseum Exterior scene occupy 32.6 GB, while patches from Trevi Fountain scene – even 65.4 GB. The goal of this experiment is to explore a data reduction technique and compare the smaller versions to the full dataset. Note that we made a related observation with AMOS Patches, where we found that using smaller number of cameras might be beneficial.

The description of the dataset reduction procedure follows. In the first step we sample batches randomly in the same manner as during training, obtain the





**Figure 4.11:** Evaluation of HardNet trained on subsets of AMOS Views. We measure mean average precision (mAP) on the HPatches benchmark.

descriptors for each batch by using HardNet pretrained on AMOS Patches, and save the observed differences

$$e(p_1) = d(p_1, p_2) - \min(d_{\text{neg1}}, d_{\text{neg2}})$$

for each patch  $p_1$ . Let  $\bar{e}_i$  be the mean value of such differences for patch set  $i$  over all patches it contains and

$$\bar{e} = \frac{1}{N} \sum_{j=1}^N \bar{e}_j$$

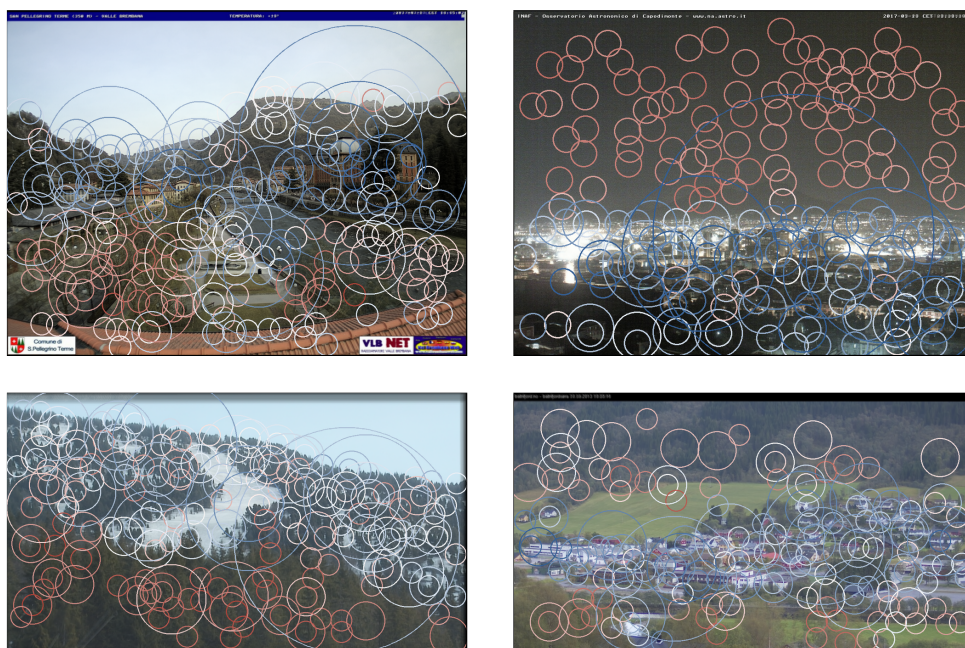
be the mean value over all  $N$  patch sets in the dataset. Then we sort the patch sets  $i = 1, \dots, N$  in ascending order according to the value of  $l_i = |e_i - \bar{e}|$ . To reduce the dataset to  $X$  number of patches, we add only those patch sets, which have the lowest distance  $l_i$ , so that the overall number of patches is less or equal to  $X$ . Such patch sets are of "medium" hardness relative to the given dataset. The remaining patches are then added randomly from the first unused patch set with the lowest  $l_i$ .

We also tried other sorting approaches to reducing a dataset, which we show in Table 4.6. "Medium" refers to the above described procedure. In "Low" and "High" we select patch sets with the lowest (easy patches) and highest (hard patches)  $\bar{e}_i$  values respectively. Note that taking the hardest patches leads to the worst performance on all datasets.

$\bar{e}_i$	HPatches	AMOS Patches	IMW PT
Low	<b>43.20</b>	<b>47.05</b>	68.78
Medium	42.38	45.91	<b>69.77</b>
High	40.64	42.99	66.77

**Table 4.6:** Evaluation of dataset reduction approaches. The first column denotes which patch sets are selected first based on the  $\bar{e}_i$  values. Each reduced version contains 1.6 million patches. Mean average precision (mAP) is measured on HPatches and AMOS Patches. In the IMW PT benchmark we measure mean average accuracy (mAA) with the threshold of  $10^\circ$ .

Figure 4.13 shows that on the IMW PT benchmark some of the reduced versions lead to better performance. Bigger architecture achieves better score with any



**Figure 4.12:** Visualization of the loss function values on AMOS Views images at the end of training. Values range from 0.8 (blue), over 1 (white) to 1.2 (red). Note that value above 1 means that the negative example is closer than positive, i.e. wrong match. Here margin in the loss function [38] is equal to 1.

number of patches. On the contrary, on HPatches the result is consistently worse on all reduced versions than if we train on the full dataset. For smaller number of patches the smaller architecture gives better mAP measure, while for more patches it becomes the opposite.

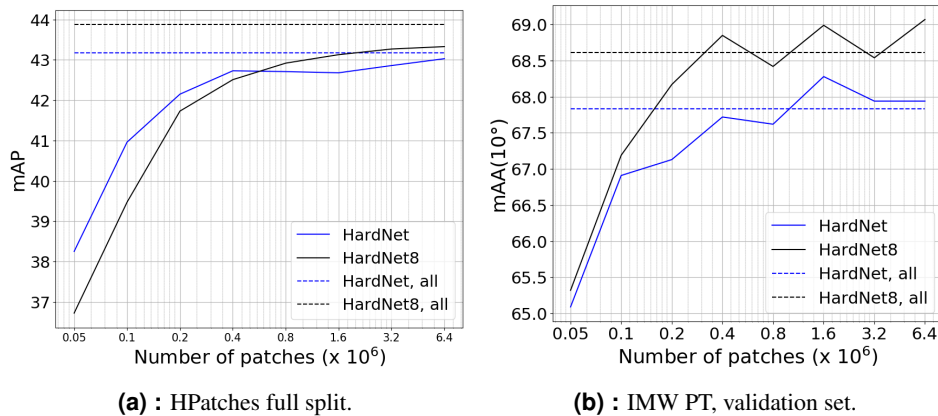
**Comparison of Scenes.** Here we make use of the described procedure of dataset reduction and of the observation that training on a smaller version leads to comparable, and sometimes better, performance. We consider several datasets from the collection of IMW PT benchmark, for each make a smaller version which contains exactly 1.6 million patches and train the HardNet architecture for 10 epochs, 5 millions samples each. The results are listed in tab 4.7. One can notice that good performance on one benchmark does not imply similar result on the other two.

#### 4.1.5 Combining Datasets

When multiple diverse datasets are available, a logical next step is to try to combine them for training. For this purpose we developed and used own software<sup>5</sup> and summarize our findings in Table 4.8 for Liberty, Notredame and AMOS Patches datasets. We compare three methods: In "Per Epoch" we split each epoch into  $D$  contiguous parts –  $D$  is number of datasets – and in each part we sample always from a single dataset. In "Per Batch" each batch is composed of patches from a

<sup>5</sup>Public version available at [https://github.com/pultarmi/HardNet\\_MultiDataset](https://github.com/pultarmi/HardNet_MultiDataset)





**Figure 4.13:** Evaluation of the dataset reduction method. The Trevi dataset is used. Dashed lines represent performance of the model trained on full dataset. The measure is mean average precision (mAP) on HPatches and mean average accuracy (mAA) with the threshold of  $10^\circ$  on IMW PT.

Dataset	HPatches	IMW PT	AMOS Patches
Buckingham Palace	35.99	67.87	43.42
Colosseum Exterior	42.38	<b>69.77</b>	45.91
Brandenburg Gate	40.58	68.54	44.12
Hagia Sophia Interior	<b>49.27</b>	67.46	39.86
Trevi Fountain	42.68	68.28	<b>46.78</b>
Prague Old Town Square	42.63	67.42	46.46

**Table 4.7:** Training datasets comparison. HardNet is trained on a dataset specified in the left column. Each dataset contains exactly 1.6 million patches. Mean average precision (mAP) is measured on HPatches and AMOS Patches. On the IMW PT benchmark we measure mean average accuracy (mAA) with the threshold of  $10^\circ$ .

single random dataset. In "In Batch" we split each batch into  $D$  blocks, where each block is composed of patches from a different dataset. "Per Batch" method most often consistently achieves better results on all benchmarks, while "Per Epoch" tends to forget the dataset, which comes first in the last epoch, as can be seen in Table 4.8 in the case of Lib+AMOS.

Combining the datasets leads to some improvement. However, one has to be careful about several things. First, in some situations we often experienced the opposite behaviour – e.g. HardNet8 trained on Colosseum achieves  $70.75$  mAA( $10^\circ$ ), while adding Liberty leads to  $70.21$  mAA( $10^\circ$ ). This negative effect can, in a sense, be reverted using compression of descriptors (explained in section 4.3), as we observed that dimensionality reduction works better if the model was trained on more datasets. Also, if AMOS Patches is present in the ensemble of datasets, it tends to influence the result significantly by improving result in HPatches and AMOS Patches, while worsening it in IMW PT. For this reason, we advise to tune the probability of sampling from AMOS using e.g. settings available in our software.

Train on	Combine	HPatches	AMOS Patches	IMW PT
Liberty	–	52.96	44.21	68.17
Notredame	–	52.44	44.65	68.56
AMOS	–	58.27	46.96	61.50
Lib+Notre	Per Epoch	53.14	45.63	68.82
	Per Batch	<b>53.92</b>	<b>45.84</b>	<b>68.89</b>
	In Batch	53.02	45.02	68.14
Lib+AMOS	Per Epoch	59.23	48.62	65.44
	Per Batch	56.63	46.66	<b>68.25</b>
AMOS+Lib	Per Epoch	56.63	46.66	<b>68.25</b>
	Per Batch	<b>59.28</b>	<b>48.43</b>	66.47
	In Batch	57.77	46.77	66.04

**Table 4.8:** Combining datasets: Influence of the training dataset(s) and combining methods on the measured metrics. The HardNet architecture is used. Lib+AMOS denotes that Liberty is sampled first in each epoch, for AMOS+Lib it is the opposite. Mean average precision (mAP) is measured on HPatches and AMOS Patches. In the IMW PT benchmark we measure mean average accuracy (mAA) with the threshold of  $10^\circ$ .

## 4.2 Architecture

The architecture of a model is one of the paramount factors influencing the performance of the descriptor. A lot of work in the field of computer vision is focused on discovering architectures that achieve higher performance, are faster or more compact. However, there is a paucity of literature available describing new architectures for a local feature descriptor. A convolutional architecture was introduced in [62] and was adopted by others [38, 34]. There has been little effort to find a better architecture since. Better architecture can generally be obtained either by manual search or by using an algorithmic solution. Here we explore both approaches.

### 4.2.1 Automated Search

Differentiable Architecture Search (DARTS) [31] is an algorithm which searches a space of architectures formulated in a continuous manner by using gradient descent. The bilevel optimization problem it approximates can be expressed as follows. Credit: [31]:

$$\begin{aligned} \min_{\alpha} \mathcal{L}_{val}(w^*(\alpha), \alpha), \\ \text{s.t. } w^*(\alpha) = \arg \min_w \mathcal{L}_{train}(w, \alpha), \end{aligned} \quad (4.1)$$

where  $\mathcal{L}_{train}$  and  $\mathcal{L}_{val}$  is the loss function on the training and validation set,  $\alpha$  is the architecture and  $w$  are the associated weights. The search space is represented

by a sequence of cells, which are composed of  $k$  ordered nodes. The input to each node (called a connection or edge) is equal to the sum of its two predecessors. The output of a cell is defined as a concatenation of all its nodes. The edges between the nodes correspond to operations that are to be learned. During architecture search each of these are assigned parameters that are weighted using the softmax operator.


After the search procedure is finished, the operation corresponding to the highest weight is assigned to each edge. It is recommended to use a larger number of cells at this point due to smaller resource requirements.

In our experiments we make use of the publicly available implementation<sup>6</sup> and plug in our data loaders, so that Liberty and Notredame from UBC Phototour [61, 16] is the training and validation dataset respectively. We run architecture search for 40 epochs, 100 000 samples each, learning rate = 2.5, batch size = 64. Other settings are kept default.

Subsequent training is performed for 40 epochs, 500 000 tuples each, with batch size = 128 and learning rate = 0.025. In Table 4.9 we list the results for several runs and compare with HardNet. In Figure 4.14 we present the found cells constituting the architecture DARTS 1. Reduction module is used in two positions in the architecture (1/3 and 2/3 of the length), all other cells are normal. The results of the DARTS models are significantly worse than those of HardNet. It might be for example due to small batch size - higher is not feasible due to GPU memory. Also, the method could be improved by changing the search space. We leave such modifications for future work.

Architecture	HPatches	AMOS Patches	IMW PT
DARTS 1	36.56	26.42	57.38
DARTS 2	28.48	17.41	46.53
DARTS 3	24.54	13.36	40.76
HardNet	<b>52.96</b>	<b>44.21</b>	<b>68.17</b>

**Table 4.9:** Evaluation of models found by the DARTS algorithm, one model per row. Several trainings were run due to the recommendation in [31]. We measure mean average precision (mAP) for HPatches and AMOS Patches and mean average accuracy mAA( $10^\circ$ ) for the CVPR IMW benchmark.



**Figure 4.14:** Modules found by the DARTS algorithm visualized using the provided tool [31].

<sup>6</sup>Available at <https://github.com/quark0/darts>

## 4.2.2 Manual Search

**VGG [58] Style.** HardNet is a VGG style network with blocks of convolutional layers of decreasing spatial size and an increasing number of channels. Here we introduce and evaluate several modifications to the architecture. With respect to the original version, HardNet7x2 contains two times more channels in each layer, HardNet8 has one convolutional layer added in the third block and HardNet9 contains two more layers, one in the second and one the third block. HardNet8x2 has two times more channels in each layer than HardNet8. See tab 4.10 for a detailed overview.

In Table 4.11 we list the results for each architecture

HardNet	HardNet7x2	HardNet8	HardNet8x2	HardNet9
1.3M param	5.3M param	4.7M param	19M param	5.3 param
Block 1, spatial size 32x32				
Conv 3x3x32	Conv 3x3x64	Conv 3x3x32	Conv 3x3x64	Conv 3x3x32
Conv 3x3x32	Conv 3x3x64	Conv 3x3x32	Conv 3x3x64	Conv 3x3x32
Block 2, spatial size 16x16				
Conv 3x3x64/2	Conv 3x3x128/2	Conv 3x3x64/2	Conv 3x3x128/2	Conv 3x3x64/2
Conv 3x3x64	Conv 3x3x128	Conv 3x3x64	Conv 3x3x128	Conv 3x3x64
-	-	-	-	Conv 3x3x128
Block 3, spatial size 8x8				
Conv 3x3x128/2	Conv 3x3x256/2	Conv 3x3x128/2	Conv 3x3x256/2	Conv 3x3x128/2
Conv 3x3x128	Conv 3x3x256	Conv 3x3x128	Conv 3x3x256	Conv 3x3x256
-	-	Conv 3x3x256	Conv 3x3x512	Conv 3x3x256
Global pooling block				
Dropout(0.3)				
Conv 8x8x128	Conv 8x8x256	Conv 8x8x256	Conv 8x8x256	Conv 8x8x256
Flatten(), L2Norm()				

**Table 4.10:** Variants of the HardNet architecture. Each convolutional layer is followed by batch normalization and ReLU activation.

Architecture	HPatches	AMOS Patches	IMW PT
HardNet	52.96	44.21	68.17
HardNet7x2	<b>54.56</b>	<b>44.39</b>	68.67
HardNet8	53.67	43.77	<b>69.43</b>
HardNet8x2	54.55	42.96	69.18
HardNet9	54.21	42.85	69.34

**Table 4.11:** Evaluation of the variants of the HardNet architecture. Mean average precision (mAP) is measured on HPatches and AMOS Patches. In the CVPR IMW benchmark we measure mean average accuracy (mAA) with the threshold of  $10^\circ$

**ResNet [21] Style.** Another popular architecture type is ResNet. It consists of blocks of convolutional layers called ResBlock. The output from each block is summed with its input, so the network learns an incremental change instead of a direct transformation. In Table 4.12 we list the architectures we tested. Evaluation of these models is shown in Table 4.13. Each model was trained on the Liberty dataset for 20 epochs, 5 million samples each. The results are almost on par with those of HardNet models. However, in none of our experiments it was superior and it has a higher GPU memory requirement, so we can not advise to use this architecture type for the learning of a local feature descriptor.

ResNet2	ResNet3	ResNet3x
1.7 param	2M params	2.1M params
Block 1, spatial size 32x32		
Conv 3x3x64	Conv 3x3x32	Conv 3x3x32
Conv 3x3x64	Conv 3x3x32	Conv 3x3x64
Block 2, spatial size 16x16		
ResBlock2/2 32	ResBlock3/2 16	ResBlock3/2 32
ResBlock2 64	ResBlock3 32	ResBlock3 64
Block 3, spatial size 8x8		
ResBlock2/2 64	ResBlock3/2 64	ResBlock3/2 64
ResBlock2 128	ResBlock3 128	ResBlock3 128
Global pooling block		
Dropout(0.3)		
Conv 8x8x256	Conv 8x8x256	Conv 8x8x128
Flatten(), L2Norm()		

**Table 4.12:** ResNet architectures. Each convolutional layer is followed by batch normalization and ReLU activation.

Architecture	HPatches	AMOS Patches	IMW PT
ResNet2	52.63	40.57	68.31
ResNet3	52.96	40.72	69.01
ResNet3x	52.25	38.50	68.27
HardNet	<b>52.96</b>	<b>44.21</b>	<b>68.17</b>

**Table 4.13:** Evaluation of the ResNet architectures together with the baseline HardNet. All models are trained on Liberty. Mean average precision (mAP) is measured on HPatches and AMOS Patches. On the CVPR IMW benchmark we measure mean average accuracy (mAA) with the threshold =  $10^\circ$ .

**Input Size.** The input to the HardNet architecture is of size  $32 \times 32$  pixels. Due to the observation that enlarging the model architecture further from HardNet8 does not lead to better performance, we perform the following experiment. We make minor adjustments to the HardNet8 architecture so that it expects inputs of a bigger size.

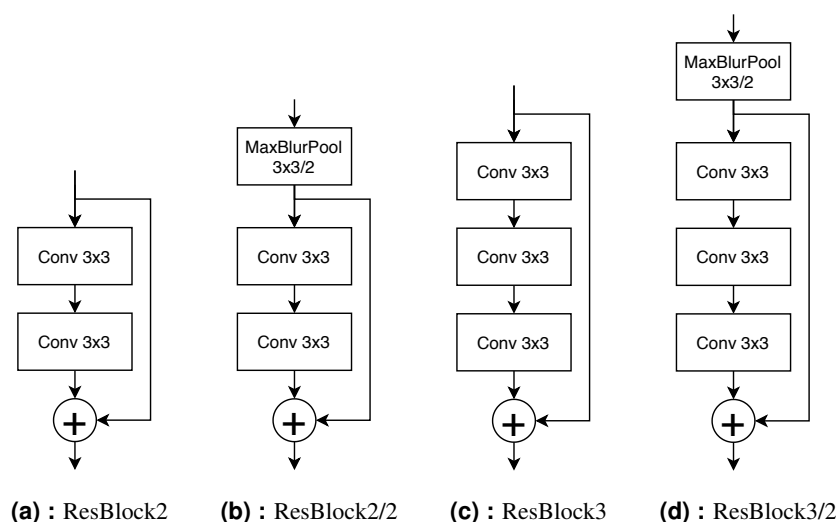


Figure 4.15: ResNet modules.

First, we change the stride to 2 in the penultimate layer, then input size is  $64 \times 64$  pixels. If we also set the kernel size of the last convolutional layer to  $6 \times 6$  pixels, the input size is then  $48 \times 48$  pixels. We compare these three models in Table 4.14. Notice the two opposite trends: bigger input size leads to better performance on HPatches and AMOS Patches, but such a model gives inferior results on the IMW PT benchmark.

Model	HPatches	AMOS Patches	IMW PT
HardNet8	52.37	42.42	<b>69.06</b>
HardNet8, input size = 48	53.34	43.12	68.41
HardNet8, input size = 64	<b>54.19</b>	<b>43.98</b>	68.25

**Table 4.14:** Evaluation of HardNet variants with different input size. Mean average precision (mAP) is measured on HPatches and AMOS Patches. In the IMW PT benchmark we measure mean average accuracy (mAA) with the threshold of  $10^\circ$ .

### 4.2.3 Final Pooling

The HardNet architecture uses global pooling implemented by a  $8 \times 8 \times 128$  convolution as the last layer in the model – i.e. in fact a fully connected linear layer – unlike SIFT, which uses local pooling, where the output vector is composed of 16 blocks based on the relative position in the input patch. In this section we compare different architectures and try several pooling approaches. All models are trained on the Liberty dataset for 20 epochs, 5 million samples each.

**Local pooling.** In this experiment we try to mimic the mechanism of SIFT and replace the last convolutional layer in order to increase its stride so that the

convolution is performed over separate blocks. In Table 4.15 we can observe that the performance of these models is worse than of the baseline.

**Increased Receptive Field.** We also test modifications to the HardNet8 architecture which increase its receptive field. In Table 4.15 we can observe that such architectures do not outperform the baseline. Enlarging the architecture by 3 layers even fails to learn.

Final Layer(s)	HPatches	AMOS Patches	IMW PT
Local Pooling			
Conv 4x4x16/4	53.28	44.50	68.50
Conv 3x3x16/2, p=1	52.76	<b>44.78</b>	68.36
Increased Receptive Field			
(Conv 3x3/2, p=1) <sup>1</sup> → Conv 4x4, p=0	52.31	42.00	68.29
(Conv 3x3/2, p=1) <sup>2</sup> → Conv 2x2, p=0	50.25	39.14	67.62
(Conv 3x3/1, p=0) <sup>3</sup> → Conv 2x2, p=0	n/c	n/c	n/c
Non-learned			
MaxPool 8x8	44.23	38.58	67.11
AvgPool 8x8	n/c	n/c	n/c
Global Pooling (original)			
Conv 8x8x256	<b>53.67</b>	43.77	<b>69.43</b>

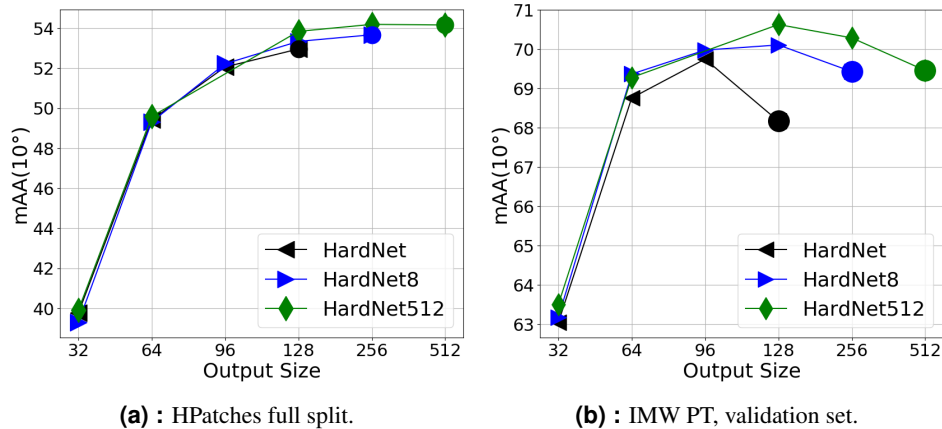
**Table 4.15:** Final pooling variants. We list modifications to the HardNet8 architecture by substituting the last convolutional layer. The last row represents no modification.  $X^y$  means the layer X is stacked  $y$  times. We write n/c for models which failed to learn. Mean average precision (mAP) is measured on HPatches and AMOS Patches. On the IMW PT benchmark we measure mean average accuracy (mAA) with the threshold of  $10^\circ$ .

**Non-learned Pooling.** Another possibility is to compute the maximum or average per input channel. MaxPool achieves worse result and AvgPool fails to learn, see Table 4.15.

## 4.3 Compression of Embeddings

During manual architecture search we found that some of the larger models, e.g. HardNet8, achieve better performance. These architectures however also output longer output vector, which brings disadvantages such as higher memory usage and slower nearest-neighbour search. If we reduce the output size of HardNet8 to 128 to match the vector length of SIFT, we get inferior results - 68.75 mAA( $10^\circ$ ) vs 69.43 mAA( $10^\circ$ ) on the IMW PT benchmark. Another way to decrease the output size is to use a dimensionality reduction technique. Here we make use of principal

component analysis (PCA) to compress the feature embeddings. See in Figure 4.16 that dimensionality reduction is beneficial. We can observe that the best combination is HardNet8 with output size 512 followed by compression to size 128. In this experiment we train on the Liberty dataset from UBC Phototour and we fit PCA on the model outputs from the training data.



**Figure 4.16:** Impact of the embeddings compression. PCA dimensionality reduction method is used. HardNet8-512 denotes HardNet8 with the number of channels in the last layer changed to 512. The measure is mean average precision (mAP) on HPatches and mean average accuracy (mAA) with the threshold of  $10^\circ$  on IMW PT. Original, non-compressed dimension is marked with a circle.


## 4.4 Loss Function

In this section we explore several modifications to the hard-in-batch triplet margin loss function. First, we find the optimal value for the margin in the original version, viz eq. 3.1. Then we introduce a generalization of the loss by sampling  $n$ -tuples of positives instead of pairs. Finally, we test an existing implementation of a variant with a dynamic margin.

**Deciding on the Margin.** The purpose of this experiment is to determine the best value for the margin in the loss function. In Figure 4.17 we can see that the performance roughly increases with the margin on both the CVPR and HPatches benchmarks. We can observe that the optimal value for margin is around 0.5 in case of the CVPR IMW benchmark. In HPatches the difference in evaluation between margins above 0.4 is almost negligible.

**Sampling More Positives.** The hard-in-batch triplet margin loss can be generalized in the following way. Let  $p_1, \dots, p_B$  be a set of patches in a batch with labels  $l_1, \dots, l_B$ . Corresponding patches have the same label. The generalized loss is then





**Figure 4.17:** HardNet performance depending on the margin applied during training. Mean average precision (mAP) is measured on HPatches. In the IMW PT benchmark we measure mean average accuracy (mAA) with the threshold of  $10^\circ$ .

formulated as

$$\mathcal{L} = \sum_{l \in \text{unique(labels)}} \max(0, M + d_{\text{pos}}^l - d_{\text{neg}}^l),$$

$$d_{\text{pos}}^l = \max_{\substack{p_i, l_i=l \\ p_j, l_j=l \\ p_i \neq p_j}} d(p_i, p_j), d_{\text{neg}}^l = \min_{\substack{p_i, l_i=l \\ p_j, l_j \neq l}} d(p_i, p_j).$$

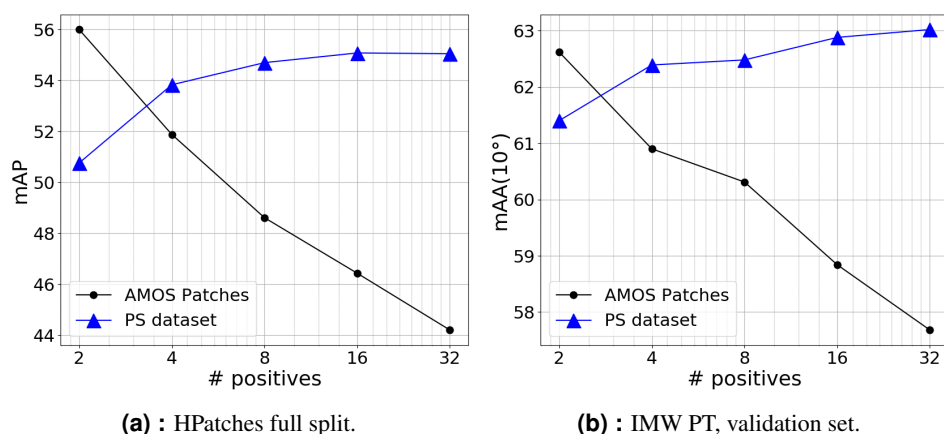
where  $d_{\text{pos}}^l$  denotes the biggest distance within a set of positives and  $d_{\text{neg}}^l$  denotes the smallest distance to the hardest negative outside of the set. Furthermore, we control the maximum number of positive patches in a batch, i.e. the batch size is variable and we only set its upper limit.

Evaluation is shown in Figure 4.18 for two training datasets. Notice there is an opposite trend for the datasets - only for the PS dataset it is beneficial to sample more positives. We think the negative effect in AMOS Patches might be caused by an amount of noise in the dataset due to dynamic content such as moving cars, so that from a batch with more positives we actually sample false positives more often.

## 4.5 Hyperparameters

We train our final models with learning rate = 3 found by the LR Range test [59]. Also, we use 1-Cycle style training [59]. It comprises of two phases. First, the learning rate increases linearly to the initial value, then it decreases to 0 by following a cosine annealing schedule. In the following text we describe also other decisions we make regarding the training setup.

**Augmentation.** During training, the patches are first horizontally and vertically flipped, both with probability 0.5. Then they are transformed to the input size.




**Figure 4.18:** HardNet performance depending on the number of positives sampled in a batch during training. Different curves represent different training datasets. The measure is mean average precision (mAP) on HPatches and mean average accuracy (mAA) with the threshold of  $10^\circ$  on IMW PT.

For AMOS Patches we apply random affine transformation and cropping to get patches of smaller size. First, random rotation from range  $(-25^\circ, 25^\circ)$ , scaling from range  $(0.8, 1.4)$  and shear are applied. Second, we crop a  $64 \times 64$  center of a patch. Then we crop a random area with scale in  $(0.7, 1.0)$  and aspect ratio in  $(0.9, 1.1)$  w.r.t. to the input patch and the patch is resized to  $32 \times 32$  pixels. These transformed patches are the input for training.


For Liberty we apply only resize transformation from  $64 \times 64$  to  $32 \times 32$  pixels. We tested other variants such as random affine transformation with or without Gaussian blur before downscaling, but none achieved better results.

**Batch Size.** Batch size is an important factor in the setup of the training procedure. See in Figure 4.19 the influence of the batch size on the performance of the trained descriptor. We use HardNet architecture and train on Liberty for 10 epochs, 5 million samples each. Notice that there is an increase in performance, measured by mAP score on AMOS Patches and HPatches benchmarks, if we increase the batch size. It is likely that a higher batch size would be even more beneficial, but we are limited by the GPU memory. On a 32 GB machine the maximum batch size is approximately 8192. The best performance on the IMW PT benchmark is achieved for batch size around 4096. Interestingly, a further increase worsens the results on both HPatches and IMW PT.

**Number of Epochs.** To decide on the number of epochs for which we train HardNet and HardNet8 models on the Liberty dataset for a different number of epochs and observe its performance. See in Figure 4.20 that it steadily increases on HPatches with the number of epochs, while it peaks at 1 epoch in the IMW PT benchmark and then decreases.



**Figure 4.19:** HardNet performance depending on the training batch size. Mean average precision (mAP) is measured on HPatches. In the IMW PT benchmark we measure mean average accuracy (mAA) with the threshold of 5° and 10°.



**Figure 4.20:** Evaluation of HardNet trained for different number of epochs. Mean average precision (mAP) is measured on HPatches. In the IMW PT benchmark we measure mean average accuracy (mAA) with the threshold of 10°.



## Chapter 5

### Evaluation on Standard Benchmarks

We evaluate two HardNet8 variants which use the improvements validated in the previous chapter. The first, HardNet8-Univ, was trained on Liberty and AMOS Patches datasets for 20 epochs, 5 million samples each, with the generalized loss function with sampling of 2 positive patches, batch size = 3072.

The second model, HardNet8-PT, has 512 output channels and was trained on Liberty, Notredame and Colosseum Exterior datasets for 40 epochs, 5 million samples each, batch size = 9000. The embeddings from the model are compressed by PCA – which is fitted on Liberty – to dimensionality of 128.

Evaluation is performed on the test set of IMW PT – the previous experiments were evaluated on the validation set. HPatches and AMOS Patches comprise of a single test split which is used both in the experiments and final evaluation.

Results are shown in Table 5.1. HardNet8-Univ works reasonably well on a wide range of conditions, be it viewpoint or illumination change, and improves state-of-the-art on the standard benchmarks. It is only outperformed by GeoDesc in the viewpoint split of HPatches. HardNet8-PT further improves the performance on the IMW PT benchmark in the Stereo 8k task.

Model	Trained on	HPatches subset			AMOS Patches	IMW PT Stereo 8k
		illum	view	full		
SIFT	–	23.33	28.88	26.15	37.08	45.84
HardNet	Liberty	49.86	55.63	52.79	43.32	55.43
SOSNet	Liberty	50.66	56.73	53.75	43.26	55.87
HardNetPS	PS Dataset	48.55	67.43	58.16	31.83	50.51
GeoDesc	GL3d [56]	50.52	<b>67.48</b>	59.15	25.50	51.11
HardNet8-Univ	AMOS+Liberty	<b>58.20</b>	63.60	<b>60.95</b>	<b>47.20</b>	56.22
HardNet8-PT	Lib+Coloss+Notre	51.04	55.81	53.46	45.01	<b>57.58</b>

**Table 5.1:** Final evaluation of the proposed HardNet8 descriptors. Mean average precision (mAP) is measured on HPatches and AMOS Patches. On the IMW PT benchmark we measure mean average accuracy (mAA) with the threshold of  $10^\circ$ .

Apart from benchmarking, we also conjecture that significantly better results can be achieved for a real-world task if HardNet8 is retrained on a specific combination

## 5. Evaluation on Standard Benchmarks

---

of datasets, suitable for the use case – e.g. AMOS Patches for illumination changes with no viewpoint change, while Liberty, Notredame, Colosseum Exterior and others for purely viewpoint changes.

## Chapter 6

### Conclusion

We have studied various factors, which influence the performance of a local feature descriptor, and made improvements to the HardNet descriptor.

AMOS Patches dataset is presented, which improves robustness of trained local feature descriptors to illumination and appearance changes. It is based on registered images from selected cameras from the AMOS dataset. Based on our experience we give recommendations related to using webcams for the local descriptor learning: (i) When picking cameras for training manually, a small and diverse subset is better than a bigger one with similar views or imprecise alignment of images; (ii) When sampling patches for a batch during training, using less source cameras is beneficial.

We present a study on combining and filtering heterogeneous datasets into a single compact dataset for the learning of a local descriptor. A technique for dataset reduction was described, we showed that training on a reduced dataset of by an order of magnitude smaller size yields comparable performance. Enriching data modality by switching to RGB or concatenating patches with estimated mono-depth maps leads to limited or no improvement.

HardNet8 architecture is presented, achieving consistently better results than the original HardNet. Other architectural and training choices were explored. We show ResNet-like architectures and models found by the DARTS algorithm are inferior to those of the VGG-like models. Similar observation was made about switching to local pooling or increasing the receptive field. Conversely, bigger training batch size is beneficial. Also, using PCA to compress the network outputs leads to better results on the IMW PT benchmark.

We also present a set of the findings, which might be useful for the practitioners, such as that the hardest negative often goes from the same image, training on more datasets might deteriorate the descriptor quality and others.

By combining all the improvements we found, we have trained two models<sup>1</sup>. One improves state-of-the-art on the HPatches, AMOS Patches and IMW PT. The other, specifically tuned to the IMW PT, further improves results on that dataset at the cost of being slightly worse on HPatches and AMOS patches.

---

<sup>1</sup>Trained models available at [https://github.com/pultarmi/HardNet\\_MultiDataset](https://github.com/pultarmi/HardNet_MultiDataset)







## Bibliography

- [1] <https://github.com/kuangliu/torchcv>. reviewed on December 2018.
- [2] H. Aanæs, A. Dahl, and K. Steenstrup Pedersen. Interesting interest points. *IJCV*, 97:18–35, 2012.
- [3] C. A. Aguilera, A. D. Sappa, C. Aguilera, and R. Toledo. Cross-spectral local descriptors via quadruplet network. *Sensors*, 17(4), 2017.
- [4] V. Balntas, K. Lenc, A. Vedaldi, and K. Mikolajczyk. HPatches: A benchmark and evaluation of handcrafted and learned local descriptors. In *CVPR*, volume 4, page 6, 2017.
- [5] V. Balntas, E. Riba, D. Ponsa, and K. Mikolajczyk. Learning local feature descriptors with triplets and shallow convolutional neural networks. In *BMVC*, 2016.
- [6] M. Brown and D. G. Lowe. Automatic panoramic image stitching using invariant features. *IJCV*, 74(1):59–73, 2007.
- [7] M. Calonder, V. Lepetit, C. Strecha, and P. Fua. Brief: Binary robust independent elementary features. In *ECCV*, pages 778–792. Springer, 2010.
- [8] A. Canziani, A. Paszke, and E. Culurciello. An analysis of deep neural network models for practical applications. *arXiv preprint arXiv:1605.07678*, 2016.
- [9] O. Chum, T. Werner, and J. Matas. Two-view geometry estimation unaffected by a dominant plane. In *CVPR*, volume 1, pages 772–779. IEEE, 2005.
- [10] R. S. Consensus. A paradigm for model fitting with applications to image analysis and automated cartography. *MA Fischler, RC Bolles*, 6:381–395, 1981.
- [11] G. Csurka and M. Humenberger. From handcrafted to deep local invariant features. *arXiv preprint arXiv:1807.10254*, 2018.
- [12] A. Dewan, T. Caselitz, and W. Burgard. Learning a local feature descriptor for 3d lidar scans. *arXiv preprint arXiv:1809.07494*, 2018.

- [13] J. Dong and S. Soatto. Domain-size pooling in local descriptors: Dsp-sift. In *CVPR*, pages 5097–5106, 2015.
- [14] M. Dusmanu, I. Rocco, T. Pajdla, M. Pollefeys, J. Sivic, A. Torii, and T. Sattler. D2-net: A trainable cnn for joint detection and description of local features. *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.
- [15] B. Fan, F. Wu, and Z. Hu. Aggregating gradient distributions into intensity orders: A novel local image descriptor. In *CVPR*, June 2011.
- [16] M. Goesele, N. Snavely, B. Curless, H. Hoppe, and S. M. Seitz. Multi-view stereo for community photo collections. In *ICCV*, pages 1–8. IEEE, 2007.
- [17] R. I. Hartley. Projective reconstruction and invariants from multiple images. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 16(10):1036–1041, 1994.
- [18] R. I. Hartley and A. Zisserman. *Multiple View Geometry in Computer Vision*. Cambridge University Press, ISBN: 0521540518, second edition, 2004.
- [19] D. Hauage and N. Snavely. Image matching using local symmetry features. In *CVPR*, pages 206–213, 2012.
- [20] K. He, Y. Lu, and S. Sclaroff. Local descriptors optimized for average precision. In *CVPR*, June 2018.
- [21] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *CVPR*, pages 770–778, 2016.
- [22] N. Jacobs, W. Burgin, N. Fridrich, A. Abrams, K. Miskell, B. H. Braswell, A. D. Richardson, and R. Pless. The global network of outdoor webcams: properties and applications. In *Proceedings of the 17th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems*, pages 111–120. ACM, 2009.
- [23] N. Jacobs, N. Roman, and R. Pless. Consistent temporal variations in many outdoor scenes. In *CVPR*, pages 1–6. IEEE, 2007.
- [24] Y. Jin, D. Mishkin, A. Mishchuk, J. Matas, P. Fua, K. M. Yi, and E. Trulls. Image Matching across Wide Baselines: From Paper to Practice. *arXiv*, 2020.
- [25] M. Keller, Z. Chen, F. Maffra, P. Schmuck, and M. Chli. Learning deep descriptors with scale-aware triplet networks. In *CVPR*, June 2018.
- [26] A. Kelman, M. Sofka, and C. V. Stewart. Keypoint descriptors for matching across multiple image modalities and non-linear intensity variations. In *CVPR*, 2007.
- [27] S. Kim, D. Min, B. Ham, M. N. Do, and K. Sohn. Dasc: Robust dense descriptor for multi-modal and multi-spectral correspondence estimation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 39(9), Sept 2017.

- [28] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pages 1097–1105, 2012.
- [29] S. Leutenegger, M. Chli, and R. Y. Siegwart. Brisk: Binary robust invariant scalable keypoints. In *2011 International conference on computer vision*, pages 2548–2555. Ieee, 2011.
- [30] Z. Li and N. Snavely. Megadepth: Learning single-view depth prediction from internet photos. In *CVPR*, 2018.
- [31] H. Liu, K. Simonyan, and Y. Yang. Darts: Differentiable architecture search. *ICLR*, 2019.
- [32] W. Liu, X. Shen, C. Wang, Z. Zhang, C. Wen, and J. Li. H-net: Neural network for cross-domain image patch matching. In *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI*, July 2018.
- [33] D. G. Lowe. Distinctive image features from scale-invariant keypoints. *IJCV*, 60(2):91–110, 2004.
- [34] Z. Luo, T. Shen, L. Zhou, S. Zhu, R. Zhang, Y. Yao, T. Fang, and L. Quan. Geodesc: Learning local descriptors by integrating geometry constraints. In *ECCV*, September 2018.
- [35] W. Maddern, G. Pascoe, C. Linegar, and P. Newman. 1 Year, 1000km: The Oxford RobotCar Dataset. *IJRR*, 36(1):3–15, 2017.
- [36] R. P. Mihail, S. Workman, Z. Bessinger, and N. Jacobs. Sky segmentation in the wild: An empirical study. In *IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 1–6. IEEE, 2016.
- [37] K. Mikolajczyk, T. Tuytelaars, C. Schmid, A. Zisserman, J. Matas, F. Schafalitzky, T. Kadir, and L. Van Gool. A comparison of affine region detectors. *IJCV*, 65(1):43–72, 2005.
- [38] A. Mishchuk, D. Mishkin, F. Radenovic, and J. Matas. Working hard to know your neighbor’s margins: Local descriptor learning loss. In *Advances in Neural Information Processing Systems*, pages 4826–4837, 2017.
- [39] D. Mishkin, J. Matas, and M. Perdoch. Mods: Fast and robust method for two-view matching. *Computer Vision and Image Understanding*, 141:81–93, 2015.
- [40] D. Mishkin, J. Matas, M. Perdoch, and K. Lenc. WxBS: Wide Baseline Stereo Generalizations. In *BMVC*. BMVA, Sept. 2015.
- [41] D. Mishkin, J. Matas, M. Perdoch, and K. Lenc. Wxbs: Wide baseline stereo generalizations. In *BMVC*, September 2015.

- [42] D. Mishkin, N. Sergievskiy, and J. Matas. Systematic evaluation of convolution neural network advances on the imagenet. *Computer Vision and Image Understanding*, 161:11 – 19, 2017.
- [43] R. Mitra, N. Doiphode, U. Gautam, S. Narayan, S. Ahmed, S. Chandran, and A. Jain. A large dataset for improving patch matching. *arXiv preprint arXiv:1801.01466*, 2018.
- [44] M. Muja and D. G. Lowe. Fast approximate nearest neighbors with automatic algorithm configuration. *VISAPP (1)*, 2(331-340):2, 2009.
- [45] H. Noh, A. Araujo, J. Sim, T. Weyand, and B. Han. Large-scale image retrieval with attentive deep local features. In *ICCV*, pages 3456–3465, 2017.
- [46] M. Pultar, M. Dmytro, and J. Matas. Leveraging outdoor webcams for local descriptor learning. *Computer Vision Winter Workshop (CVWW)*, feb 2019.
- [47] A. Resindra, A. Torii, and M. Okutomi. Structure from motion using dense cnn features with keypoint relocalization. *IPSN Transactions on Computer Vision and Applications*, 10, Dec 2018.
- [48] J. Revaud, P. Weinzaepfel, C. De Souza, N. Pion, G. Csurka, Y. Cabon, and M. Humenberger. R2d2: Repeatable and reliable detector and descriptor. *NeurIPS*, 2019.
- [49] E. Rublee, V. Rabaud, K. Konolige, and G. Bradski. Orb: An efficient alternative to sift or surf. In *ICCV*, pages 2564–2571. Ieee, 2011.
- [50] S. Ryu, S. Kim, and K. Sohn. Lat: Local area transform for cross modal correspondence matching. *Pattern Recognition*, 63, 2017.
- [51] P.-E. Sarlin, C. Cadena, R. Siegwart, and M. Dymczyk. From coarse to fine: Robust hierarchical localization at large scale. In *CVPR*, pages 12716–12725, 2019.
- [52] T. Sattler, W. Maddern, C. Toft, A. Torii, L. Hammarstrand, E. Stenborg, D. Safari, M. Okutomi, M. Pollefeys, J. Sivic, F. Kahl, and T. Pajdla. Benchmarking 6DOF outdoor visual localization in changing conditions. In *CVPR*, pages 8601–8610. IEEE Computer Society, 2018.
- [53] J. L. Schonberger and J.-M. Frahm. Structure-from-motion revisited. In *CVPR*, pages 4104–4113, 2016.
- [54] E. Shechtman and M. Irani. Matching local self-similarities across images and videos. In *CVPR*, 2007.
- [55] T. Shen, Z. Luo, L. Zhou, R. Zhang, S. Zhu, T. Fang, and L. Quan. Matchable image retrieval by learning from surface reconstruction. *arXiv preprint arXiv:1811.10343*, 2018.
- [56] T. Shen, Z. Luo, L. Zhou, R. Zhang, S. Zhu, T. Fang, and L. Quan. Matchable image retrieval by learning from surface reconstruction. In *ACCV*, 2018.

- [57] E. Simo-Serra, E. Trulls, L. Ferraz, I. Kokkinos, P. Fua, and F. Moreno-Noguer. Discriminative learning of deep convolutional feature point descriptors. In *ICCV*, pages 118–126, 2015.
- [58] K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. *ICLR*, 2014.
- [59] L. N. Smith. A disciplined approach to neural network hyper-parameters: Part 1—learning rate, batch size, momentum, and weight decay. *arXiv preprint arXiv:1803.09820*, 2018.
- [60] N. Snavely, S. M. Seitz, and R. Szeliski. Photo tourism: exploring photo collections in 3d. In *ACM Siggraph 2006 Papers*, pages 835–846. 2006.
- [61] N. Snavely, S. M. Seitz, and R. Szeliski. Modeling the world from internet photo collections. *IJCV*, 80(2):189–210, 2008.
- [62] Y. Tian, B. Fan, and F. Wu. L2-net: Deep learning of discriminative patch descriptor in euclidean space. In *CVPR*, 2017.
- [63] Y. Tian, X. Yu, B. Fan, F. Wu, H. Heijnen, and V. Balntas. Sosnet: Second order similarity regularization for local descriptor learning. In *CVPR*, pages 11016–11025, 2019.
- [64] Y. Verdie, K. M. Yi, P. Fua, and V. Lepetit. TILDE: A temporally invariant learned detector. In *CVPR*, 2015.
- [65] Z. Wang, B. Fan, G. Wang, and F. Wu. Exploring local and overall ordinal information for robust feature description. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 38(11), Nov 2016.
- [66] G. Yang, C. V. Stewart, M. Sofka, and C.-L. Tsai. Registration of challenging image pairs: Initialization, estimation, and decision. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 29(11), 2007.
- [67] T. Zhi, B. R. Pires, M. Hebert, and S. G. Narasimhan. Deep material-aware cross-spectral stereo matching. In *CVPR*, June 2018.
- [68] B. Zhou, H. Zhao, X. Puig, T. Xiao, S. Fidler, A. Barriuso, and A. Torralba. Semantic understanding of scenes through the ADE20k dataset. *IJCV*, 2018.
- [69] Q. Zhou, T. Sattler, M. Pollefeys, and L. Leal-Taixe. To learn or not to learn: Visual localization from essential matrices. *arXiv preprint arXiv:1908.01293*, 2019.
- [70] J.-Y. Zhu, T. Park, P. Isola, and A. A. Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *IEEE International Conference on Computer Vision (ICCV)*, 2017.