

# A case for using rotation invariant features in state of the art feature matchers

Georg Bökman

Fredrik Kahl

{bokman, fredrik.kahl}@chalmers.se

Chalmers University of Technology

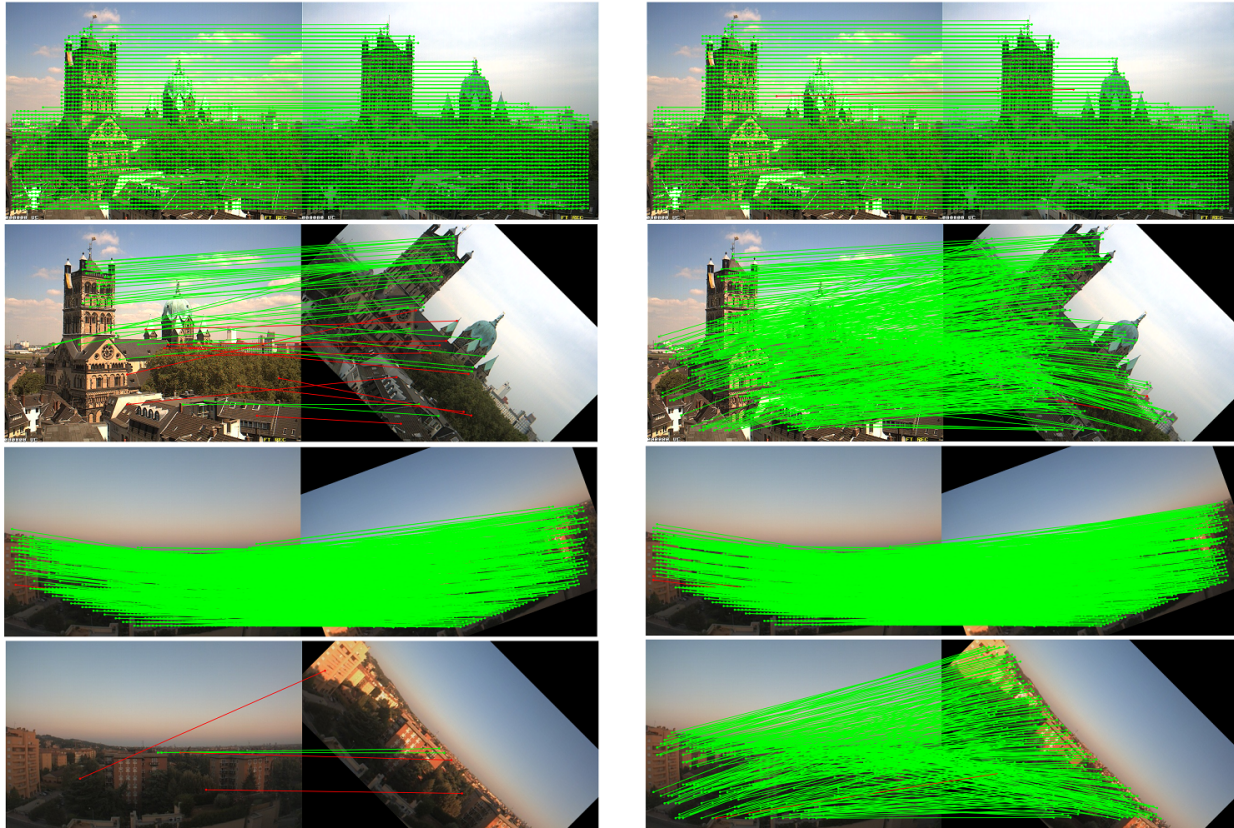


Figure 1. *Left:* LoFTR [39] finds good matches between images under illumination and small viewpoint changes, but the performance deteriorates completely under large rotation changes. *Right:* Our model SE2-LoFTR-4\*, which is obtained by making the LoFTR backbone CNN rotation equivariant, performs well both in the situations where the original LoFTR does, and under large rotation changes. The image pairs are from two HPatches illumination sequences, the bottom three pairs are augmented by inplane rotations of 45°, 20° and 45° respectively. Matches with reprojection error under 10px are green, others red.

## Abstract

*The aim of this paper is to demonstrate that a state of the art feature matcher (LoFTR) can be made more robust to rotations by simply replacing the backbone CNN with a steerable CNN which is equivariant to translations and image rotations. It is experimentally shown that this boost is obtained without reducing performance on ordinary illumination and viewpoint matching sequences.*

## 1. Introduction

Finding corresponding points across images is a basic building block in many computer vision tasks including structure-from-motion, motion estimation and visual localization [20, 41]. Hand-crafted methods have dominated the field for a long time, but now we see major improvements with learning-based methods, making image matching work despite strong illumination changes and wide baselines.

Classical descriptors have typically been rotation invari-

ant by design. In the deep learning era, rotation invariant descriptors have been researched, but state of the art feature extractors usually rely on ordinary CNNs for keypoint (or dense) description, which means that the obtained descriptors are not rotation invariant. In this paper we show that rotation invariant features can easily be obtained by using steerable CNNs for feature extraction and our experimental results indicate that while this leads to better matching of rotated images it does not deteriorate matching of non-rotated images.

We use LoFTR as a base architecture and perform extensive matching on HPatches and geometrically transformed versions of HPatches. Three different equivariant alternatives are developed and experimentally compared with the baseline.

## 2. Related work and relevant theory

This work stands on the shoulders of two giants – the image matching literature and the literature on group equivariant neural networks. Next, we mention related work from each of these two areas.

### 2.1. Matching

The goal of matching (sub-)pixels in two images is typically to use the obtained correspondences to estimate some form of transformation between the two images. Examples include estimating an essential matrix, a fundamental matrix or a homography [20].

The classical matching pipeline can be broken up into four steps.

1. *Detection*. Finding locations of interesting keypoints in both images.
2. *Description*. Encoding the image content around these keypoints. The encodings are often called descriptors or features.
3. *Matching*. Matching the obtained descriptors to determine which locations in the two images correspond to each other.
4. *Filtering*. Processing the matches to filter out outliers, i.e. incorrect matches.

The two first steps are usually coupled, while the last steps have been more independent. Many recent neural network based approaches however incorporate several of the steps end-to-end. The most famous method for detection and description is SIFT [32], where the detection is based on finding extrema in scale-space and the description of each keypoint is based on image gradients of close by image patches. SIFT descriptors are (approximately) rotation invariant by design. Other classical rotation invariant descriptors include [4, 27, 37]. The archetypal matching method is

nearest neighbour matching and the most common outlier filtering method is RANSAC [17].

CNN-based approaches for detection and description of keypoints include for instance D2-Net, R2D2-Net and DISK [15, 36, 44], but most related to our approach are the ones that use rotation invariant features such as LIFT, GIFT and others [26, 31, 49, 51]. While we believe that all of these approaches could be used to good effect in combination with more recent CNN based matching pipelines to obtain the benefit of rotation invariant features, none of the required modifications would be as straightforward or natural as just directly replicating a CNN architecture and replacing the layers with their steerable analogues, which is what we suggest.

Another line of research skips detection completely and instead hinges on densely describing every pixel in both images (or downsampled versions). This approach has recently been used with great success in [16, 39, 43]. We will base our experiments on LoFTR [39], which uses a CNN for dense feature description and transformer layers for further feature processing and matching. The approach is described in more detail in Section 3.1.

Learning-based methods have also been applied to the matching step. In [38], a graph neural network is trained to perform matching, similar to our base model LoFTR.

Once a set of matches (with outliers) has been obtained, deep learning can also be used for the filtering step or to directly regress the transformation between the two images (e.g. a homography or an essential matrix). Approaches include [40, 50, 52, 53] and the rotation invariant [6].

### 2.2. Group equivariant neural networks

We will use the framework of  $E(2)$  equivariant steerable CNNs [46] to construct rotation invariant dense descriptors. Let us first briefly describe the broader context of group equivariant neural networks before going into a bit more detail on the  $E(2)$ -steerable ones.

A key property of CNN layers is that they are translation equivariant, i.e. if  $T$  denotes a translation,  $L$  a CNN layer and  $I$  an image, then  $L(T(I)) = T(L(I))$ . This property makes CNNs process image features equivalently regardless of where in an image they appear.

Extensive research has been done on describing neural networks that are equivariant under more general groups than translations [2, 8, 9, 11, 24, 25, 47]. Under mild assumptions, linear layers that are group equivariant have been shown to be required to have a convolutional form. The theory of linear equivariant layers is formulated in terms of *representation theory* [18].

More broadly, group equivariant neural networks are part of the growing framework of geometric deep learning [7, 19].

### 2.2.1 E(2) equivariant steerable CNNs

Steerable CNNs [10,46,48] are examples of group equivariant neural networks. Let us sketch the most relevant parts of representation theory to be able to explain steerable CNNs. In the following we will think of  $G$  as being a subgroup of  $SO(2)$ , but the theory applies in more general settings.

A tuple  $(\mathbb{R}^m, \rho)$  is said to be a representation of  $G$  if  $\rho$  is a function that associates to each  $g \in G$  an  $m \times m$  matrix  $\rho(g)$ , such that  $\rho(g_1 g_2) = \rho(g_1) \rho(g_2)$  for all  $g_1, g_2 \in G$  (succinctly,  $\rho$  is a group homomorphism)<sup>1</sup>. This means that the group product in  $G$  is encoded in  $m \times m$  matrix multiplication, however this encoding is not required to be *faithful*, i.e. one to one. A simple valid representation for any group and any vector space  $\mathbb{R}^m$  is the trivial representation where for all  $g$ ,  $\rho(g)$  is the identity matrix<sup>2</sup>.

In an ordinary CNN, the input image is processed into feature maps, which associate to every pixel a vector in  $\mathbb{R}^m$  with  $m$  channels. The analogues in  $G$ -steerable CNNs of feature maps are called feature fields. Like a feature map, a feature field associates each pixel to a feature vector, but whereas a CNN feature vector has no specific structure, a vector in a feature field is interpreted to lie in a representation  $(\mathbb{R}^m, \rho)$  of the group  $G$ . Concretely, this means that the feature at each pixel is a vector  $v \in \mathbb{R}^m$  and when the pixel grid is acted on by  $g \in G$  (i.e. rotated),  $v$  is changed to  $\rho(g)v$ . An input image corresponds to a feature field in the trivial representation – when the image is rotated this does not change the actual RGB-values, it only changes their positions. Features in the trivial representation are said to be *G-invariant*. In contrast, consider rotating an image gradient field. The gradient vectors change position and they also rotate the same amount as the underlying pixel grid. A gradient field is a representation  $(\mathbb{R}^2, \rho)$  where  $\rho(g)$  is the 2D rotation matrix associated with the rotation  $g$ .

In our experiments we use *regular representations* of finite subgroups  $G$  of  $SO(2)$ . We refer to [18,46] for a rigorous definition of regular representations, but note that they are a way to associate a feature to each group element of  $G$ . In other words, regular representations have the same number of dimensions as  $G$  has elements, e.g. if  $G$  has four elements then the regular representation has the vector space  $\mathbb{R}^4$ . The matrices  $\rho(g)$  in this case are permutation matrices which faithfully encode the group multiplication of  $G$ . Taking multiple copies of the regular representation we can construct feature fields with any multiple of  $|G|$  number of channels.

Designing a steerable CNN architecture involves not only specifying the number of channels  $m$  in each feature field, but also what representations these channels should be

<sup>1</sup>More generally, one can instead of  $\mathbb{R}^m$  consider an arbitrary vector space  $V$  and let  $\rho$  be a group homomorphism  $G \rightarrow GL(V)$ .

<sup>2</sup>Technically this amounts to  $m$  copies of the trivial representation, which is one dimensional.

interpreted as belonging to. In [10,46] it is shown how given a specification of input and output feature fields one can use the theory of irreducible representations [18] to solve for the linear layers/convolution kernels that are  $G$ -equivariant between the inputs and outputs. This leads to even more weight sharing than using ordinary convolutions. We refer to [46] for details. They also provide the Pytorch package `e2cnn` which we use for our experiments.

## 3. Models

The base model for our experiments is LoFTR [39]. To investigate the usefulness of rotation invariant features, we evaluate the original LoFTR as well as three different versions with steerable CNNs as feature extracting backbone.

### 3.1. LoFTR

LoFTR (Local Feature TRansformer) is a four stage deep neural network for finding point correspondences between two images  $I^A, I^B$  of the same scene. Briefly, the four stages are as follows. For more details we refer to [39].

1. A feature pyramid CNN is used to extract features in the two images. This CNN is referred to as the *backbone*. Features are obtained both at a coarse level (standard: 1/8 of the image size) and a fine level (standard: 1/2 of the image size).
2. The coarse features are concatenated with positional encodings and fed into a type of transformer called a *LoFTR module*.
3. The transformed coarse features are matched between the images using a differentiable matching module. As part of this matching, a confidence value is computed for each possible correspondence between the coarse feature maps. The output is a set of matching positions  $\hat{i}^A, \hat{i}^B$  in the two images which are mutual most confident matches and whose matching confidences are above a threshold hyperparameter  $\theta_c$ .
4. Around the predicted locations of the coarse matches, patches of the fine level feature maps are cropped out and processed in another LoFTR module. The output corresponding patches are used to compute matches from each  $\hat{i}^A$  to a subpixel position  $\hat{i}^B$ .

LoFTR is trained using a loss on the coarse match confidences as well as a loss on the fine matches. At validation/test time, after LoFTR has been used to find good matches between two images, RANSAC is used to estimate for instance an essential matrix or a homography.

The backbone CNN is a ResNet [21] style feature pyramid network [30]. It has an initial stride 2 convolution halving the image dimensions, followed by three ResNet blocks



	Number of learnable parameters in backbone	Relative number of intermediate backbone features	Time for last training epoch
LoFTR	5.9M	100%	1h 53m
SE2-LoFTR-4*	1.1M	100%	1h 56m
SE2-LoFTR-4	4.1M	200%	2h 28m
SE2-LoFTR-8*	540k	100%	1h 58m

Table 1. Brief comparison of the different models considered.

with strides 1, 2, and 2, i.e. with outputs of 1/2, 1/4 and 1/8 of the original image dimensions respectively. This is followed by layers that iteratively upsample to 1/2 of the original image dimensions. Coarse features are extracted at the end of the downsampling blocks and fine features at the end of the upsampling.

### 3.2. SE2-LoFTR

We create three LoFTR variants by replacing the CNN backbone with steerable CNNs. This modification is easily implemented with the help of the `e2cnn` Pytorch package [46]. We call the modified models SE2-LoFTR as we will consider convolutions that are steerable under certain subgroups  $G$  of  $SO(2)$ , which combined with translation equivariance means that the layers will be equivariant under subgroups of  $SE(2)$ , the group of rotations and translations.

As explained in Section 2.2.1, implementing a  $G$ -equivariant steerable CNN involves selecting for each layer, which representation of  $G$  the features should be interpreted as lying in. In [46] the authors found that using regular representations of finite subgroups of  $SO(2)$  works very well compared to other choices and so we will use this approach.

We will consider  $G$  as being either the group  $C_4$  of quarter rotations in the plane or the group  $C_8$  of eighth rotations. The features that are extracted by the backbone will lie in the trivial representation of  $G$  and hence be invariant to quarter or eighth rotations respectively. Having this type of discrete approximation of rotation invariance has been seen to work well for image classification [5, 9, 46].

Input and output features are defined to be in the trivial representation of  $G$  and all intermediate features are chosen to be in the regular representation of  $G$ . However, the feature pyramid structure of the LoFTR backbone means that there are features output at the middle of the network. As we want these coarse features to be rotation invariant rather than living in the regular representation of  $G$ , we add a readout layer going from the middle of the network (where the features are in the regular representation) to features in the trivial representation of  $G$ . The outputs of this readout layer are used as coarse features.

The reason for not using trivial representations through-

out the network is expressivity. Using regular representations allows for much more expressive layers [46].

An advantage of using the regular representation rather than particular other for the features is that we can apply pointwise nonlinearities while preserving equivariance. Thus we use an ordinary ReLU nonlinearity as in the original LoFTR backbone. `BatchNorm2d` layers are replaced by their  $G$ -equivariant analogues `InnerBatchNorm` layers.

The following is a list of the three SE2-LoFTR variants that we experiment with.

- *SE2-LoFTR-4\**. Here the LoFTR backbone is changed to be  $C_4$ -steerable.  $C_4$  is the group of quarter rotations. For every four channels in an intermediate feature map of the LoFTR backbone we use one regular representation in the steerable feature field. Since the regular representation has  $|C_4| = 4$  dimensions, this makes the total number of channels equal between SE2-LoFTR-4\* and LoFTR. However, SE2-LoFTR-4\* has far fewer learnable parameters than LoFTR (see Table 1) because of the increased weight sharing in the equivariant layers.
- *SE2-LoFTR-4*. Like SE2-LoFTR-4\* but for every *two* channels in an intermediate feature map of the LoFTR backbone we use one regular representation channel. This brings the number of parameters closer to that of LoFTR but increases the number of features and hence both the model size and compute time.
- *SE2-LoFTR-8\**. Like SE2-LoFTR-4\* but the considered group is  $C_8$  – the group of eighth rotations. For every eight channels in an intermediate feature map of the LoFTR backbone we use one regular representation copy in the feature field.

Indicating the models preserving the number of channels of the original model with an asterisk is consistent with [46]. In practice, these are the most similar to the original model in terms of compute and memory consumption but not in terms of number of trainable parameters.

In Table 1 we list information about the four models considered in this paper. The slight overhead in training time for SE2-LoFTR- $\{4,8\}^*$  over LoFTR is a consequence of the steerable layers, but this overhead can be alleviated at test time, by converting the steerable layers to ordinary CNN layers (see [46, Section 2.8]).

#### 3.2.1 A comment on the positional encoding

Before being passed into the LoFTR module, the coarse features are concatenated with positional encodings to enable the network to differentiate between absolute positions in the images. This means that the obtained features are not

actually rotation invariant, but the fact that a part of them is invariant can still facilitate matching of image patches related by a rotation.

### 3.2.2 Implementation details

Our code is a modification of the official LoFTR implementation and will be made available at [github.com/georg-bn/se2-loftr](https://github.com/georg-bn/se2-loftr). In particular, all hyperparameters are kept fixed from LoFTR, but we train on eight GPUs rather than sixteen, which reduces the batch size by a factor of 2. This should however not affect performance too drastically as the LoFTR implementation scales relevant hyperparameters with the batch size. We use the *dual-softmax* version of the LoFTR matching in all experiments as this version was best performing in [39]. During testing, we fix the matching confidence threshold  $\theta_c$  to 0.2, the same value as is used during training. Tuning this parameter for a specific test set can improve performance. We train all models on identical cluster compute nodes containing eight NVIDIA A100-SXM4-40GB GPUs.

## 4. Datasets/benchmarks

We train on MegaDepth [29], and test on HPatches sequences [3] as well as geometrically modified versions of HPatches sequences.

### 4.1. MegaDepth

We train the models on MegaDepth [29], using the exact same setup as LoFTR. MegaDepth consists of 196 sets of images collected from the internet, each set corresponding to a particular tourist attraction. The task is to determine the relative pose given two images in the same set. Following LoFTR [39] and DISK [44], validation is done on the *Sacre Coeur* and *St Peter's Square* sets. There is no separate test set and early stopping is performed on the validation set. We hence argue that the results on MegaDepth should be interpreted as validation results only.

#### 4.1.1 Metrics

The performance on MegaDepth is measured as in [38, 39] in terms of the area under curve (AUC) of the pose accuracy up to a specific threshold. The considered thresholds are  $5^\circ$ ,  $10^\circ$  and  $20^\circ$ . The pose accuracy at angle  $a$  is defined as the proportion of relative poses whose rotation and translation are both within  $a$  from the ground truth. Each AUC is normalized so that the max score is 100%. Early validation stopping is done on the AUC@ $10^\circ$ .

### 4.2. HPatches

Testing of the models is done on HPatches sequences [3] where the task is to find a homography relating two im-

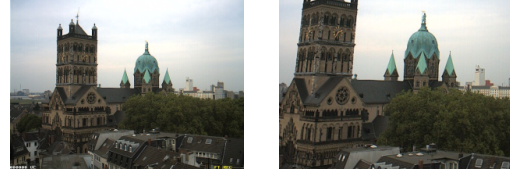


Figure 2. Left: An image from the HPatches sequences. Right: The modified image in HPatches-h0.3.

Pose error AUC (%)	@ $5^\circ$	@ $10^\circ$	@ $20^\circ$
LoFTR (paper)	52.8	69.2	81.3
LoFTR (re-trained)	52.2	68.9	81.0
SE2-LoFTR-4*	51.3	68.0	80.3
SE2-LoFTR-4	52.6	69.2	81.4
SE2-LoFTR-8*	50.1	66.9	79.4

Table 2. Validation results on MegaDepth.

ages. Again we follow the LoFTR setup. We use the original HPatches sequences and the later added ones from [1, 12, 22, 33, 34, 45]. Following D2-Net [15] eight image sequences are removed<sup>3</sup> which leaves 108 sequences. Out of these, 52 sequences contain *illumination* changes and 56 contain *viewpoint* changes. Each sequence contains six images and ground truth homographies relating the first to each of the following five. We refer to the first image in each sequence as the *A*-image and to the five images in each sequence that are to be related to *A* as the *B*-images.

We follow [13] in resizing all images to size  $640 \times 480$  (or  $480 \times 640$  for upright images). This might slightly deviate from the approach used by the LoFTR authors who state “All images are resized with shorter dimensions equal to 480.”. As the aim of the resizing is to make the AUC of corner errors comparable across images, we believe it to be reasonable to use a uniform size for all images.

#### 4.2.1 Modifications of HPatches sequences

As we want to test the robustness of the models under rotations, we introduce modifications of the HPatches sequences as follows. For an angle  $a \in (0^\circ, 90^\circ]$ , we rotate each *B*-image by  $a$  either clockwise or anticlockwise (this is chosen randomly). We call the obtained modified HPatches sequence HPatches- $ra$ , and will consider HPatches-r20 and HPatches-r45. Examples of these sequences can be seen in Figure 1.

To test if enforcing increased robustness under rotations influences robustness under other types of geometric transformations, we introduce a further modified HPatches variant. Each *B*-image is warped as follows. Let the image

<sup>3</sup>High resolution images are removed. This removal can be considered a historical artefact stemming from the fact that some methods were not able to handle the high resolutions.

	MMA	All @3px	@5px	@10px	Illumination			Viewpoint		
					@3px	@5px	@10px	@3px	@5px	@10px
<b>HPatches</b>										
LoFTR (re-evaluated)		93.4	95.2	96.1	98.4	98.7	99.3	88.9	91.9	93.1
LoFTR (re-trained)		93.4	95.1	96.0	98.6	98.9	99.5	88.7	91.5	92.7
SE2-LoFTR-4*		94.3	96.2	97.3	98.3	98.7	99.4	90.5	93.9	95.3
SE2-LoFTR-4		94.5	96.5	97.5	98.5	98.9	99.6	90.8	94.2	95.5
SE2-LoFTR-8*		94.1	96.1	97.3	98.2	98.7	99.4	90.3	93.8	95.3
<b>HPatches-r20</b>										
LoFTR (re-evaluated)		86.0	90.5	92.6	89.5	94.6	97.1	82.7	86.7	88.5
LoFTR (re-trained)		86.6	91.0	92.9	90.2	95.2	97.4	83.2	87.1	88.7
SE2-LoFTR-4*		89.4	93.8	95.8	90.7	95.2	97.4	88.1	92.5	94.3
SE2-LoFTR-4		89.7	94.0	95.9	91.1	95.5	97.6	88.4	92.5	94.3
SE2-LoFTR-8*		89.5	93.9	95.8	90.5	95.2	97.5	88.5	92.7	94.3
<b>HPatches-r45</b>										
LoFTR (re-evaluated)		35.0	41.5	46.2	37.7	45.2	50.3	32.5	38.1	42.4
LoFTR (re-trained)		30.7	38.1	44.3	30.3	39.3	46.9	31.1	37.1	41.9
SE2-LoFTR-4*		72.6	83.9	89.2	69.6	81.9	88.3	75.3	85.7	90.0
SE2-LoFTR-4		70.9	84.1	89.9	68.4	82.7	89.4	73.3	85.4	90.4
SE2-LoFTR-8*		80.4	89.3	93.0	78.8	88.7	93.1	81.9	89.9	92.8
<b>HPatches-h0.3</b>										
LoFTR (re-evaluated)		87.7	92.7	95.0	88.9	94.6	97.3	86.5	90.9	92.9
LoFTR (re-trained)		88.0	93.0	95.2	89.8	95.3	97.8	86.4	90.8	92.7
SE2-LoFTR-4*		88.9	94.2	96.6	89.3	95.0	97.6	88.6	93.4	95.6
SE2-LoFTR-4		89.3	94.4	96.6	89.5	95.0	97.5	89.1	93.8	95.8
SE2-LoFTR-8*		88.2	93.6	96.2	87.7	93.8	96.8	88.6	93.4	95.7

Table 3. Mean match accuracy in percent on HPatches and modified versions. HPatches- $ra$  contains images that are rotated by  $a$  degrees from the original HPatches sequences and HPatches-h0.3 contains images warped using the scheme described in Section 4.2.1.

dimensions be  $h \times w$ , and choose a value  $s > 0$  to control the amount of warping.

1. For each of the four image corners, sample an offset of at most  $(s \cdot h, s \cdot w)$ .
2. Warp the corners outwards from the image center by the respective offsets. I.e. so that the upper left corner  $(0, 0)$  is moved to somewhere in the rectangle  $[-s \cdot h, 0] \times [-s \cdot w, 0]$  etc.

This warp is determined by four points and hence corresponds to a homography, which is used to alter the ground truth homography from  $A$  to  $B$ . The resulting  $B$ -images are slightly skewed zoom-ins of the originals. We denote the obtained modified sequence by HPatches- $hs$ , and will consider HPatches-h0.3 in the experiments. An example image of HPatches-h0.3 is shown in Figure 2.

The modifications in HPatches- $ra$  and HPatches- $hs$  are carried out on top of the illumination/viewpoint transformations that are already present in HPatches. Hence, these datasets are quite a bit more challenging than the original HPatches.

#### 4.2.2 Metrics

The AUC@ $x$  metric is used as in [39]. For an image correspondence  $A, B$  with ground truth homography  $H$ , the *corner error* is defined as the mean of the distances between the corners of  $A$  warped by  $H$  and the corners of  $A$  warped by an estimated homography  $\tilde{H}$ . The metric used is then the area under curve (AUC) of the corner error accuracy up to thresholds 3px, 5px and 10px. Each AUC score is normalized so that the max is 100%.

We furthermore report the mean match accuracy (MMA) at the same thresholds. The MMA is defined as the mean proportion of matches that are within a certain threshold in reprojection error. It was used in e.g. [14, 16, 55].

To compute these metrics, we use the `immatch` package [54]. The estimated homography  $\tilde{H}$  is obtained by running the OpenCV function `findHomography` on the top 1000 most confident matches with RANSAC as the outlier filtering method. We use the default OpenCV hyperparameters for RANSAC.

The metrics are sensitive to for instance image size, RANSAC hyperparameters as well as the matching threshold  $\theta_c$ . As we do not know the exact settings used for the results in the LoFTR paper, the results are not directly comparable.

The metrics are sensitive to for instance image size, RANSAC hyperparameters as well as the matching threshold  $\theta_c$ . As we do not know the exact settings used for the results in the LoFTR paper, the results are not directly comparable.

## 5. Results

For comparison, we retrain a LoFTR version using our training setup (in particular lower batch size) and evaluate it along with the pretrained version supplied by the authors.

Corner error AUC	All @3px	@5px	@10px	Illumination			Viewpoint		
	@3px	@5px	@10px	@3px	@5px	@10px	@3px	@5px	@10px
<b>HPatches</b>									
LoFTR (paper)	65.9	75.6	84.6						
LoFTR (re-evaluated)	65.3	74.8	84.4	81.8	88.7	94.2	50.2	62.1	75.5
LoFTR (re-trained)	65.3	75.5	84.4	81.5	88.4	94.0	50.5	63.6	75.6
SE2-LoFTR-4*	65.2	75.1	84.7	80.8	87.9	93.8	51.0	63.3	76.3
SE2-LoFTR-4	66.2	76.6	86.0	82.2	89.2	94.4	51.5	65.0	78.3
SE2-LoFTR-8*	65.8	76.0	85.5	80.4	87.8	93.8	52.3	65.2	77.9
<b>HPatches-r20</b>									
LoFTR (re-evaluated)	52.5	65.3	78.0	61.0	74.7	86.7	44.7	56.7	70.0
LoFTR (re-trained)	51.3	64.7	77.4	62.8	76.4	87.6	40.9	53.9	68.1
SE2-LoFTR-4*	55.3	68.4	80.8	64.5	76.9	87.7	46.9	60.6	74.5
SE2-LoFTR-4	54.9	68.3	80.9	64.0	76.4	87.7	46.6	61.1	74.7
SE2-LoFTR-8*	56.1	68.7	81.1	65.6	77.8	88.3	47.5	60.3	74.6
<b>HPatches-r45</b>									
LoFTR (re-evaluated)	15.1	22.9	32.0	16.2	26.5	37.8	14.2	19.7	26.8
LoFTR (re-trained)	10.3	16.9	27.3	9.64	18.0	31.0	11.1	16.1	24.1
SE2-LoFTR-4*	31.0	47.1	66.3	31.6	50.3	70.3	30.5	44.4	62.8
SE2-LoFTR-4	25.4	42.3	64.0	24.2	42.9	66.5	26.7	41.8	61.8
SE2-LoFTR-8*	38.8	55.2	73.3	42.0	59.6	77.8	36.0	51.3	69.2
<b>HPatches-h0.3</b>									
LoFTR (re-evaluated)	44.9	59.0	74.9	54.4	69.4	83.7	36.3	49.6	66.9
LoFTR (re-trained)	44.7	58.7	74.2	54.2	68.8	83.0	36.1	49.5	66.3
SE2-LoFTR-4*	45.4	60.0	75.8	55.5	70.6	84.0	36.2	50.3	68.3
SE2-LoFTR-4	45.9	60.5	76.2	55.7	70.5	84.2	37.0	51.4	69.0
SE2-LoFTR-8*	45.7	60.0	75.6	54.9	69.4	83.4	37.4	51.4	68.4

Table 4. AUC of corner error in percent on HPatches and modified versions. HPatches- $r_a$  contains images that are rotated by  $a$  degrees from the original HPatches sequences and HPatches-h0.3 contains images warped using the scheme described in Section 4.2.1.

The retrained version uses updated code which fixes a bug in the positional encoding. We furthermore report the numbers from the LoFTR paper where available.

### 5.1. MegaDepth

Validation results on MegaDepth are shown in Table 2. We see that performance correlates quite well with the number of parameters in each model, but the performance is similar across all models. All models perform better than DRC-Net [28] and SuperPoint+SuperGlue [13, 38] which are the comparisons in the LoFTR paper [39, Table 3].

### 5.2. HPatches

Qualitative results are shown in Figure 1. Tables 3 and 4 show the test results on the ordinary HPatches sequences as well as modified versions described in Section 4.2.1. The most important score is the corner error AUC in Table 4 as the MMA (Table 3) can be unreliable when the amount of extracted matches fluctuates. The difference between the LoFTR (paper) row and the LoFTR (re-evaluated) row is due to differences in how the evaluation is performed as described in Section 4.2.2.

Let us point out a couple of interesting results.

- The SE2 versions clearly outperform the baseline on the rotated HPatches-r datasets.

- The SE2 versions perform very well on the unmodified HPatches. This means that using rotation invariant features is not detrimental to general performance. In particular, the SE2 versions quite consistently outperform the baseline on the sequences of HPatches with viewpoint changes.
- The SE2 versions outperform the baseline on the warped HPatches-h0.3, showing that rotation invariant features can improve robustness under more general geometric transformations than rotations.
- The SE2-LoFTR-8\* model outperforms the other SE2 versions on the rotated HPatches-r datasets. Its features are by design invariant to 45° rotations so the great performance on HPatches-r45 is unsurprising.

### 5.3. Limitations

We have only trained each model once due to computational limitations. This is similar to most papers in the field. However, we have saved checkpoints from 5 epochs of the training run, corresponding to the top 5 validation scores on AUC@10° on MegaDepth. We can evaluate each of these on HPatches to get a sense of the variability of the test scores. Table 5 shows that the variability is not too high.

	Corner error AUC			MMA		
	@3px	@5px	@10px	@3px	@5px	@10px
LoFTR (paper)	(65.9, , )	(75.6, , )	(84.6, , )			
LoFTR (re-evaluated)	(65.3, , )	(74.8, , )	(84.4, , )	(93.4, , )	(95.2, , )	(96.1, , )
LoFTR (re-trained)	(65.3, 65.8, 64.2)	(75.5, 75.5, 74.1)	(84.4, 84.6, 83.5)	(93.4, 93.8, 93.3)	(95.1, 95.4, 95.0)	(96.0, 96.3, 95.9)
SE2-LoFTR-4*	(65.2, 65.6, 65.1)	(75.1, 75.7, 75.0)	(84.7, 85.5, 84.7)	(94.3, 94.5, 94.1)	(96.2, 96.4, 96.0)	(97.3, 97.5, 97.0)
SE2-LoFTR-4	(66.2, 66.3, 65.5)	(76.6, 76.6, 76.0)	(86.0, 86.0, 85.6)	(94.5, 94.7, 94.2)	(96.5, 96.7, 96.2)	(97.5, 97.6, 97.2)
SE2-LoFTR-8*	(65.8, 65.8, 64.6)	(76.0, 76.0, 74.6)	(85.5, 85.5, 84.6)	(94.1, 94.1, 94.0)	(96.1, 96.2, 96.0)	(97.3, 97.4, 97.2)

Table 5. Variability among test scores of the top five validation checkpoints of a training run. The scores presented are over all ordinary HPatches sequences in the format (Score of checkpoint with best MegaDepth validation score, Highest score on HPatches among the five checkpoints, Lowest score on HPatches among the five checkpoints). I.e. the first value of each triple corresponds to the value presented in all other tables.

We tested implementing a naive data augmentation scheme where images were rotated by some multiple of quarter rotations, but could not get reasonable results as the coarse matching too often failed to produce any matches at all during training. Any form of data augmentation technique would however likely require much longer training times as the training data set implicitly is enlarged.

The `e2cnn` package enables many more general representations for the feature fields of the backbone CNN. We have limited the scope of this paper to regular representations but cannot claim to know that this is the best choice.

## 6. Conclusion and future work

We argue that steerable CNNs should be a more common tool in feature matching pipelines than they currently are and have showed promising results with a very simple modification of the LoFTR model. Robustness to rotations is achieved while keeping the matching performance for non-rotated image pairs and with little computational overhead.

We note that the whole pipeline is not equivariant due to the transformer layers. It would be interesting to investigate whether further boosts can be obtained by designing the transformer layers in an equivariant manner as well.

Finally, we mention that rotation invariant features can be of great value in domains such as matching aerial images where there is no canonical image orientation or non-rigid matching where the orientations of the features might change between the images. Steerable CNN feature extraction could also be combined with the approach proposed in [42], where features are obtained on rectified planar surfaces of the input images – these rectified planar regions can have arbitrary relative planar rotations between the two images.

## Acknowledgements

We thank the anonymous reviewers whose comments helped improve the paper. This work was supported by the Swedish Foundation for Strategic Research and the

Wallenberg AI, Autonomous Systems and Software Program (WASP) funded by the Knut and Alice Wallenberg Foundation. The computations were enabled by resources provided by the Swedish National Infrastructure for Computing (SNIC), partially funded by the Swedish Research Council through grant agreement no. 2018-05973.

## References

- [1] Henrik Aanæs, Anders Lindbjerg Dahl, and Kim Steenstrup Pedersen. Interesting Interest Points. *Int J Comput Vis*, 97(1):18–35, Mar. 2012. 5
- [2] Jimmy Aronsson. Homogeneous vector bundles and  $G$ -equivariant convolutional neural networks. *arXiv:2105.05400 [cs, math, stat]*, May 2021. 2
- [3] Vassileios Balntas, Karel Lenc, Andrea Vedaldi, and Krystian Mikołajczyk. HPatches: A Benchmark and Evaluation of Handcrafted and Learned Local Descriptors. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3852–3861, July 2017. 5
- [4] Herbert Bay, Tinne Tuytelaars, and Luc Van Gool. Surf: Speeded up robust features. In *European Conference on Computer Vision*, pages 404–417. Springer, 2006. 2
- [5] Erik J. Bekkers, Maxime W. Lafarge, Mitko Veta, Koen AJ Eppenhof, Josien PW Pluim, and Remco Duits. Roto-Translation Covariant Convolutional Networks for Medical Image Analysis. *arXiv:1804.03393 [cs, math]*, June 2018. 4
- [6] Georg Bökman, Fredrik Kahl, and Axel Flinth. ZZ-Net: A Universal Rotation Equivariant Architecture for 2D Point Clouds. *CVPR*, 2022. 2
- [7] Michael M. Bronstein, Joan Bruna, Taco Cohen, and Petar Veličković. Geometric Deep Learning: Grids, Groups, Graphs, Geodesics, and Gauges. *arXiv:2104.13478 [cs, stat]*, May 2021. 2
- [8] Taco Cohen, Maurice Weiler, Berkay Kicanaoglu, and Max Welling. Gauge equivariant convolutional networks and the icosahedral CNN. In *Int. Conf. Machine Learning*, 2019. 2
- [9] Taco Cohen and Max Welling. Group equivariant convolutional networks. In *Int. Conf. Machine Learning*, 2016. 2, 4
- [10] T.S. Cohen and M. Welling. Steerable cnns. In *ICLR*, 2017. 3



- [11] Taco S. Cohen. *Equivariant Convolutional Networks (PhD Thesis)*. PhD thesis, University of Amsterdam, June 2021. 2
- [12] Kai Cordes, Bodo Rosenhahn, and Jorn Ostermann. Increasing the accuracy of feature evaluation benchmarks using differential evolution. In *2011 IEEE Symposium on Differential Evolution (SDE)*, pages 1–8, Paris, France, Apr. 2011. IEEE. 5
- [13] Daniel DeTone, Tomasz Malisiewicz, and Andrew Rabinovich. SuperPoint: Self-Supervised Interest Point Detection and Description. *arXiv:1712.07629 [cs]*, Apr. 2018. 5, 7
- [14] Mihai Dusmanu, Ignacio Rocco, Tomas Pajdla, Marc Pollefeys, Josef Sivic, Akihiko Torii, and Torsten Sattler. D2-Net: A Trainable CNN for Joint Description and Detection of Local Features. In *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 8084–8093, Long Beach, CA, USA, June 2019. IEEE. 6
- [15] 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. In *CVPR*, 2019. 2, 5
- [16] Johan Edstedt, Mårten Wadenbäck, and Michael Felsberg. Deep Kernelized Dense Geometric Matching. *arXiv:2202.00667 [cs]*, Feb. 2022. 2, 6
- [17] M. A. Fischler and R. C. Bolles. Random sample consensus: A paradigm for model fitting with application to image analysis and automated cartography. *Commun. Assoc. Comp. Mach.*, 24:381–395, 1981. 2
- [18] William Fulton and Joe Harris. *Representation Theory*, volume 129 of *Graduate Texts in Mathematics*. Springer New York, New York, NY, 2004. 2, 3
- [19] Jan E. Gerken, Jimmy Aronsson, Oscar Carlsson, Hampus Linander, Fredrik Ohlsson, Christoffer Petersson, and Daniel Persson. Geometric Deep Learning and Equivariant Neural Networks. *arXiv:2105.13926 [hep-th]*, May 2021. 2
- [20] Richard Hartley and Andrew Zisserman. *Multiple View Geometry in Computer Vision*. Cambridge University Press, Cambridge, UK ; New York, 2nd ed edition, 2003. 1, 2
- [21] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep Residual Learning for Image Recognition. *arXiv:1512.03385 [cs]*, Dec. 2015. 3
- [22] Nathan Jacobs, Nathaniel Roman, and Robert Pless. Consistent Temporal Variations in Many Outdoor Scenes. In *2007 IEEE Conference on Computer Vision and Pattern Recognition*, pages 1–6, June 2007. 5
- [23] F. Kahl, R. Hartley, and K. Åström. Critical configurations for N-view projective reconstruction. In *CVPR*, 2001. 8
- [24] Risi Kondor and Shubhendu Trivedi. On the Generalization of Equivariance and Convolution in Neural Networks to the Action of Compact Groups. *arXiv:1802.03690 [cs, stat]*, Nov. 2018. 2
- [25] Leon Lang and Maurice Weiler. A Wigner-Eckart Theorem for Group Equivariant Convolution Kernels. *arXiv:2010.10952 [cs]*, Jan. 2021. 2
- [26] Karel Lenc and Andrea Vedaldi. Learning Covariant Feature Detectors. *arXiv:1605.01224 [cs]*, Sept. 2016. 2
- [27] Stefan Leutenegger, Margarita Chli, and Roland Y Siegwart. BRISK: Binary robust invariant scalable keypoints. In *2011 International Conference on Computer Vision*, pages 2548–2555. Ieee, 2011. 2
- [28] Xinghui Li, Kai Han, Shuda Li, and Victor Prisacariu. Dual-Resolution Correspondence Networks. *NeurIPS*, page 12, 2020. 7
- [29] Zhengqi Li and Noah Snavely. MegaDepth: Learning Single-View Depth Prediction from Internet Photos. In *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2041–2050, Salt Lake City, UT, USA, June 2018. IEEE. 5
- [30] Tsung-Yi Lin, Piotr Dollár, Ross Girshick, Kaiming He, Bharath Hariharan, and Serge Belongie. Feature Pyramid Networks for Object Detection. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 936–944, Honolulu, HI, July 2017. IEEE. 3
- [31] Yuan Liu, Zehong Shen, Zhixuan Lin, Sida Peng, Hujun Bao, and Xiaowei Zhou. GIFT: Learning Transformation-Invariant Dense Visual Descriptors via Group CNNs. *arXiv:1911.05932 [cs]*, Nov. 2019. 2
- [32] David G. Lowe. Distinctive Image Features from Scale-Invariant Keypoints. *International Journal of Computer Vision*, 60(2):91–110, Nov. 2004. 2
- [33] Krystian Mikolajczyk and Cordelia Schmid. A Performance Evaluation of Local Descriptors. *IEEE Trans. Pattern Anal. Mach. Intell.*, 27(10):1615–1630, Oct. 2005. 5
- [34] Jean-Michel Morel and Guoshen Yu. ASIFT: A New Framework for Fully Affine Invariant Image Comparison. *SIAM J. Imaging Sci.*, 2(2):438–469, Jan. 2009. 5
- [35] C. Olsson, F. Kahl, and M. Oskarsson. The registration problem revisited: Optimal solutions from points, lines and planes. In *CVPR*, 2006. 8
- [36] Jerome Revaud, Philippe Weinzaepfel, César Roberto de Souza, and Martin Humenberger. R2D2: repeatable and reliable detector and descriptor. In *NeurIPS*, 2019. 2
- [37] Ethan Rublee, Vincent Rabaud, Kurt Konolige, and Gary Bradski. ORB: An efficient alternative to SIFT or SURF. In *2011 International Conference on Computer Vision*, pages 2564–2571. Ieee, 2011. 2
- [38] Paul-Edouard Sarlin, Daniel DeTone, Tomasz Malisiewicz, and Andrew Rabinovich. SuperGlue: Learning Feature Matching with Graph Neural Networks. *arXiv:1911.11763 [cs]*, Mar. 2020. 2, 5, 7
- [39] Jiaming Sun, Zehong Shen, Yuang Wang, Hujun Bao, and Xiaowei Zhou. LoFTR: Detector-free local feature matching with transformers. *CVPR*, 2021. 1, 2, 3, 5, 6, 7
- [40] Weiwei Sun, Wei Jiang, Eduard Trulls, Andrea Tagliasacchi, and Kwang Moo Yi. ACNe: Attentive Context Normalization for Robust Permutation-Equivariant Learning. *arXiv:1907.02545 [cs]*, Jan. 2021. 2
- [41] Carl Toft, Will Maddern, Akihiko Torii, Lars Hammarstrand, Erik Stenborg, Daniel Safari, Masatoshi Okutomi, Marc Pollefeys, Josef Sivic, Tomas Pajdla, Fredrik Kahl, and Torsten Sattler. Long-term visual localization revisited. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(4):2074–2088, 2022. 1

- [42] Carl Toft, Daniyar Turmukhambetov, Torsten Sattler, Fredrik Kahl, and Gabriel J. Brostow. Single-Image Depth Prediction Makes Feature Matching Easier. In *European Conference on Computer Vision (ECCV)*, 2020. 8
- [43] P. Truong, M. Danelljan, F. Yu, and L. Van Gool. Warp consistency for unsupervised learning of dense correspondences. In *ICCV*, 2021. 2
- [44] Michał Tyszkiewicz, Pascal Fua, and Eduard Trulls. DISK: Learning local features with policy gradient. *NeurIPS*, 2020. 2, 5
- [45] Vassilios Vonikakis, Dimitris Chrysostomou, Rigas Kouskouridas, and Antonios Gasteratos. Improving the robustness in feature detection by local contrast enhancement. In *2012 IEEE International Conference on Imaging Systems and Techniques Proceedings*, pages 158–163, July 2012. 5
- [46] Maurice Weiler and Gabriele Cesa. General e(2)-equivariant steerable cnns. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019. 2, 3, 4
- [47] Maurice Weiler, Patrick Forré, Erik Verlinde, and Max Welling. Coordinate Independent Convolutional Networks – Isometry and Gauge Equivariant Convolutions on Riemannian Manifolds. *arXiv:2106.06020 [cs, stat]*, June 2021. 2
- [48] Maurice Weiler, Mario Geiger, Max Welling, Wouter Boomsma, and Taco Cohen. 3D Steerable CNNs: Learning Rotationally Equivariant Features in Volumetric Data. *arXiv:1807.02547 [cs, stat]*, Oct. 2018. 3
- [49] Kwang Moo Yi, Eduard Trulls, Vincent Lepetit, and Pascal Fua. LIFT: Learned Invariant Feature Transform. *arXiv:1603.09114 [cs]*, July 2016. 2
- [50] Kwang Moo Yi, Eduard Trulls, Yuki Ono, Vincent Lepetit, Mathieu Salzmann, and Pascal Fua. Learning to Find Good Correspondences. *arXiv:1711.05971 [cs]*, May 2018. 2
- [51] Kwang Moo Yi, Yannick Verdie, Pascal Fua, and Vincent Lepetit. Learning to Assign Orientations to Feature Points. *arXiv:1511.04273 [cs]*, Apr. 2016. 2
- [52] Jiahui Zhang, Dawei Sun, Zixin Luo, Anbang Yao, Lei Zhou, Tianwei Shen, Yurong Chen, Long Quan, and Hongen Liao. Learning two-view correspondences and geometry using order-aware network. *International Conference on Computer Vision (ICCV)*, 2019. 2
- [53] Zhen Zhong, Guobao Xiao, Linxin Zheng, Yan Lu, and Jiayi Ma. T-Net: Effective permutation-equivariant network for two-view correspondence learning. In *ICCV*, 2021. 2
- [54] Qunjie Zhou. Image-matching-toolbox. <https://github.com/GrumpyZhou/image-matching-toolbox>. (accessed 2022-04-04). 6
- [55] Qunjie Zhou, Torsten Sattler, and Laura Leal-Taixe. Patch2Pix: Epipolar-Guided Pixel-Level Correspondences. In *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4667–4676, Nashville, TN, USA, June 2021. IEEE. 6