

Rotated Ring, Radial and Depth Wise Separable Radial Convolutions

Wolfgang Fuhl
University of Tübingen
Sand 14, 72076 Tübingen, Germany
wolfgang.fuhl@uni-tuebingen.de

Enkelejda Kasneci
University of Tübingen
Sand 14, 72076 Tübingen, Germany
enkelejda.kasneci@uni-tuebingen.de

Abstract

Simple image rotations significantly reduce the accuracy of deep neural networks. In addition, training with all possible rotations increases the data set, which also increases the training duration. In this work we deal with trainable rotation invariant convolutions as well as with the construction of nets, since the fully connected layers can also only be rotation invariant with a one-dimensional input. On the one hand we show that our approach is rotationally invariant for different models and on different public data sets. We also discuss the influence of purely rotational invariant features on the accuracy. The rotationally adaptive convolution models presented here are more computationally intensive than normal convolution models, so we also present a depth wise separable approach with radial convolution.

1. Introduction

Deep neural networks have become the state of the art in image processing. Many applications requiring image classification [24], landmark regression [38], image segmentation [26] or 3D reconstruction [25] are realized by deep neural networks. This requires a large annotated data set [22] as well as additional data manipulation [28] to obtain robust and efficient networks. This data manipulation requires additional computational effort during training and extends the data set many times over. One of the most frequently used data manipulations is the mirroring and rotation of images. However, mirroring can only be applied to two axes in 2D images and is therefore cheap. In contrast to this, the rotation offers an infinite number of possibilities. Modern training methods use optimizers like SGD [3] or ADAM [20] which in turn use one or more gradient moments [33]. In addition, techniques such as Batch Normalization [17] have been developed to stabilize the training and improve generalization. The training of really deep nets was only possible with the Residual Layers [14] and there

are many techniques to include data manipulation according to optimization strategies in the training [28]. While all these techniques further stabilize the training and promote generalization, object rotation is still a major challenge today. Some examples of applications are carried eye trackers with adjustable cameras like the Dikablis Professional, food [31], objects in a kitchen [8], galaxies [1], the underwater world from the perspective of a diver, plankton [27], grabbing robot arms with a moving camera [32], and many more. In these tasks it is desirable to use neural networks, which by definition are rotation invariant to reduce the training time and to ensure that the objects can be detected under all possible rotations. This variety of applications has already produced some scientific work in the field of rotational invariance.

The first group of approaches deals with the integration of rotation into the training data. Here, different transformations are applied to the data and a separate mesh is trained for each transformation. In the end, either the mean value of all nets [4, 9, 23] or the maximum result of a rotation map of all nets is used [36, 34]. The second group is about learning the transformations themselves [7, 15, 18]. Here nets are trained, which transform the data to a uniform representation. As in the first group it is necessary to rotate the data during the training. The last group is the complete integration of trainable rotated filters [11, 30, 2, 5, 6, 37, 29, 39], to which this work also belongs. This means that the built-in layers are learned holistically with the net and can be trained with the back propagation algorithms. This group is also called steerable approaches. For this last group it is not necessary to transform the data for the training.

In this work we propose two novel approaches for integrated rotation invariant training of deep neural networks. Our work can be seen as an extension to the rotated weights approaches [11, 30] which reduces the performance of neural networks drastically. This is done by rotating the filter per filter level outgoing of the center and selecting the maximum response per rotation ring. Since our approach increases the computational costs of deep neural networks

we also propose a radial convolution approach which only consists of a radial weight vector. In addition, we further reduce the computational complexity of our approaches with the technique of depth wise separable convolutions which were proposed in MobileNet [16]. We tested our layers with different Models and on different publicly available data sets. Our main contributions are as follows:

- 1 Rotated ring convolutions.
- 2 Separable depth wise rotated convolutions.
- 3 Radial Convolutions.
- 4 Publicly available CUDA implementations for both.

2. Related Work

Since the state of the art is divided into three main groups, we have also divided this section into these three sub sections. The first group deals with the separate learning of transformations by individual neural networks and the subsequent combination for classifications. The second group learns how to transform input data to a unified representation. Thereby any net can be used subsequently. In the last group the filters are rotated and fully integrated into the back propagation for training. In the following the individual approaches are described in detail.

2.1. Data Augmentation and Ensemble building

In this group, the transformations are applied to the data during training and a variety of nets are trained with them. This ensemble of nets increases the accuracy on the transformed data. But also the calculation time, because many nets have to be trained. In [4] for example such an ensemble was trained where each net has learned its own transformation. For the application the results of the nets were averaged and in case of a classification the maximum was selected. In [9] galaxies were transformed and merged. This representation was used for the training. By this the mesh learned to see the galaxy from different points of view in parallel, which made the mesh more robust and accurate. A very similar approach was presented in [23]. Here, Siamese networks with shared weights were used. A max pooling was applied to the feature vectors of these networks instead of merging these feature vectors. The output of the max pooling was used to train a subsequent classifier. Another approach is [36] which internally uses an ensemble of oriented edge detectors. The output of these edge filters is projected on a two-dimensional rotation map using different meshes. This map is then used by a classifier to determine the class and rotation. [34] is a similar approach, because it uses a set of meshes to create a two-dimensional rotational map. However, compared to [36], wavlets are used instead of the oriented edge filters.

2.2. The learning of Transformations

The first approach to learning about transformations comes from unsupervised learning. For this purpose [15] used an autoencoder and let it learn the reverse transformation. The autoencoder was trained with as many transformations as possible on the data and the classifier was trained on the output, which received a uniform input from the autoencoder. A similar approach was presented in [18]. Here a small network with weakly monitored learning was trained to transform the input data into a uniform output. This small net can then be integrated at any point in the net. In [7] we went one step further and learned deformable convolutions and a deformable region of interest pooling. However, these have to be trained for the different transformations. As in the first group, the data with the applied transformations must be available in the training for learning transformations. This is a big disadvantage, because it prolongs the training and you cannot guarantee to have considered all transformations in the training.

2.3. Steerable approaches

This section deals with the methods that transform the filters and therefore do not require rotation of the training data. The goal here is that the filters in the network adapt themselves so that a rotation on the input data does not matter. These approaches are also called controllable approaches and the theoretical framework was presented in [6]. Based on the rotation of filters, a number of approaches have been proposed [12, 30, 5, 6, 37, 29, 39] which differ only minimally. Based on the work in [6], group-equivariant convolution neural networks (GCNNs) were proposed in [5]. These consist of groups of rotated convolutions which are then used to perform a pooling operation along the output layers. These groups of convolutions are limited to 90 degrees. An extension without the restriction to 90 degrees are the harmonic nets (N-Nets) [37]. They limit the filters to circular harmonics and require summing over several convolutions which makes these nets, like many others, very computationally demanding compared to standard convolutions. Another approach are the vector field nets [29, 30]. Here the filter is rotated by predefined rotations and the maximum is selected from the output. Additionally, not only the maximum is passed, but also the orientation, which led to the naming of the nets. Oriented Response Networks (ORN) [39] are similar networks where the number of orientations is arbitrary. For this purpose the filter is rotated based on the orientations and the maximum is selected from the convolution. The difference to vector field networks is that the filter is actively rotated but still the amount of orientation is preliminary fixed to four or eight and the response as well as the orientation is passed. Since this is not rotation invariant, as with vector fields, the orientation is also passed to the next layer, the

108
109
110
111
112
113
114
115
116
117
118
119
120
121
122
123
124
125
126
127
128
129
130
131
132
133
134
135
136
137
138
139
140
141
142
143
144
145
146
147
148
149
150
151
152
153
154
155
156
157
158
159
160
161162
163
164
165
166
167
168
169
170
171
172
173
174
175
176
177
178
179
180
181
182
183
184
185
186
187
188
189
190
191
192
193
194
195
196
197
198
199
200
201
202
203
204
205
206
207
208
209
210
211
212
213
214
215




Figure 1. Visual description of the depth wise convolutions with the radial filter construction (RSDW). The left and right side are the depth wise convolutions and in the center are the radial filters. In comparison to RAD the filters are applied as 2D Convolutions per channel and not as tensors.

authors proposed two ways to make it rotation invariant. First is a SIFT feature like alignment and the second one is maximum selection. A further variant of this approach of rotating filters and a novel rotated pooling based module using the filter orientations was presented in Rotationally-Invariant Convolution Module (RIC) [12]. Here the authors have limited themselves to 3×3 convolutions and rotated only the weights between the corners. This turns one filter into exactly four for all rotations. The other weights (corners and center) are identical. Like in [12] we do not use any rotation in the training phase and evaluate different rotations in the validation phase.

3. Method




Figure 2. Visual description of the filter construction of our radial convolutions (RAD). The same color stands for the same value. On the left side of the arrow are the parameters without bias term and on the right side the applied convolution. Each input convolution tensor has other radial parameters per input channel.

Figure 2 shows the radial convolutions. As you can see, the weights to be learned (Always on the left side of an arrow) correspond to an integer distance from the center of the filter (Always on the right side of an arrow). The filters are then constructed so that the distance of the index in the filter to the center is the index of the parameters.

$$F_{c_{in},c_{out}}(i,j) = W_{c_{in},c_{out}}(\text{round}(\text{sqrt}(i^2 + j^2))) \quad (1)$$

This indexing is described in Equation 1 where i, j is the 2D Filter Index and W is the weight field. Since the convolutions are tensors and not only 2D filters, each channel in the tensor has its own weight field c_{in} . This is of course also valid for the output channels c_{out} where each channel

has a separate tensor. In the following this method will be called RAD.

We have also extended this method with the technique of depth wise convolutions. In the first step, a tensor based depth wise convolution is performed, which determines the number of channels in the middle layer. This can be seen in Figure 1. In the middle layer the radial convolutions are applied not as tensor but as individual 2D convolutions for each input channel. Thus, the first depth wise convolution determines the number of output and input channels in the middle layer in Figure 1. At the end, a depth wise convolution is performed again to create a desired output depth. This method is called RSDW in the following.

$$F_{c_{in}/out}(i,j) = W_{c_{in}/out}(\text{round}(\text{sqrt}(i^2 + j^2))) \quad (2)$$

Equation 2 shows the change for the middle section of Figure 1. Since we only use one 2D convolution for each input channel and the input and output layers have the same size, we only need one additional filter index for the channels c_{in}/out .

Both methods (RAD, RSDW) are by definition rotation invariant, because each feature in e.g. a 3×3 field can be rotated around the central pixel without changing the multiplication equation with the filter. In addition, all pixels with the same distance from the center can even be permuted arbitrarily without changing the multiplication equation. Of course, this is also a disadvantage, because the learned features are less meaningful. However, since this is general already the case for purely rotation-invariant features [11, 30, 2, 5, 6, 37, 29, 39], this is not a major disadvantage.

With our third approach (RING) we address the reduced meaningfulness of the features from the first approach. For this purpose we use a similar approach of the rotated filters [39]. However, in our method, each individual ring of the filter rotates. This can be seen in Figure 3. In the case of a 3×3 convolution, this corresponds to the rotated filters




Figure 3. Visual description of the rotated ring convolutions (RING). In the top row the rotations for a 3×3 are shown. For a 5×5 convolution the outer ring is rotated independently. This can be seen in the bottom row.

of [39]. In the case of a 5×5 convolution (lower part of figure 3), another ring is added. This second or third ring, if the central weight is considered a separate ring, will also rotate independently. This independent rotation of the rings makes it possible to learn more than just the features with rotated filters. An example is shown in Figure 4 where each pattern represents the same feature.




Figure 4. Visual example of features that rotate within themselves. All three patches represent the same feature and the two outer rings are always independently rotated around the central pixel.

As you can see, the outer and inner ring is always rotated around the central pixel and both rotate independently. Another advantage of looking at them independently is that not all combinations of both rotations result in their own filters, but each ring gets its own convolution filters, reducing the number of convolutions to $8+16$ instead of $8*16$. At the end the maximum of each ring is selected and added together with the bias term. Compared to the rotated filters [39], our method does not require a fixed number F of filters which has to be set in advance, because each ring is always completely rotated. This gives for our method $\sum_{i=2}^f 2*(i+i-2)$ convolutions for symmetric filters with width and height f .

To train these rotating filters using the backpropagation algorithm, the error as well as the gradient must be rotated. To make this easier and to use the fast cuDNN convolutions we have simply added several storage stages, which allow to assign the error to each rotated filter and thus to each weight. This follows the idea of the rotated filters [39] with the difference that we add separate stages for multiple rings.

Algorithm 1 describes the modifications which have to be made to the back propagation algorithm in order to use the normal convolution together with the cuDNN im-

Data: Data,Weights,WRot

Result: Output,ORot

Function Forward is

```
WRot=RotateRings(Weights);
ORot=cuDNNFWD(WRot,Data);
Output=Comp&MaxPerRing(ORot);
```

end

Data: ErrorIn,ORot,WRot

Result: ErrorOut,ERot

Function Backward is

```
ERot=MaxPerRingBWD(ErrorIn,ORot);
ErrorOut=cuDNNBWD(WRot,ERot);
```

end

Data: ERot,Data

Result: Grad

Function CompGradis

```
GRot=cuDNNCompGrad(Data,ERot);
Grad=RotateRingsGrad(GRot);
```

end

Algorithm 1: Algorithmic description of the modifications to the convolution procedure for the forward, backward, and gradient computation work flow.

plementations for convolution. In the forward algorithm the filters have to be created based on the ring rotations (RotateRings). Afterwards the cuDNN convolution can be used. To combine the results of the individual rings and to get the correct output depth the combination and maximum selection must be executed after the convolution (Comp&MaxPerRing). For the back propagation of the error, the input gradient must be assigned to the maxima selected in the forward step (MaxPerRingBWD). Afterwards the cuDNN back propagation can be used. To calculate the gradients, the assigned error and the input data can be used directly with the cuDNN gradient calculation. Finally, only the calculated gradients have to be assigned to the individual weights in the not rotated filter (RotateRingsGrad). Afterwards the weights can be updated with any optimizer.

Table 1 shows the required parameters for each ap-

324
325
326
327
328
329
330
331
332
333

378
379
380
381
382
383
384
385
386
387

334
335
336

388
389
390

337
338
339
340
341
342
343
344

391
392
393
394
395
396
397
398

345
346

399
400

347
348
349
350
351

401
402
403
404
405

352
353
354

406
407
408

355
356

409
410

357
358
359
360
361

411
412
413
414
415

362
363
364
365
366
367

416
417
418
419
420
421

368
369
370
371
372
373
374

422
423
424
425
426
427
428

375
376
377

429
430
431

Table 1. The necessary weights without bias term. Conv represents a normal tensor convolution. In case of RSDW, out1 stands for the output channels of the first depth wise convolution and out2 for the output channels of the last depth wise convolution.

| Method | Weights |
|--------------------------|--------------------------------------|
| Conv 3×3 | $3*3*in*out$ |
| Conv 5×5 | $5*5*in*out$ |
| RIC 3×3 [11] | $5*in*out$ |
| ORN 3×3 [39] | $3*3*in*out$ |
| ORN 5×5 [39] | $5*5*in*out$ |
| RAD 2 (ours) | $2*in*out$ |
| RAD 3 (ours) | $3*in*out$ |
| RSDW 2 (ours) | $(in*out1) + (2*out1) + (out1*out2)$ |
| RSDW 3 (ours) | $(in*out1) + (3*out1) + (out1*out2)$ |
| RING 3×3 (ours) | $3*3*in*out$ |
| RING 5×5 (ours) | $5*5*in*out$ |

proach. As you can see, no approach needs more parameters than the usual convolution. However, the number of parameters for RIC [11] as well as for RAD is much lower. In the case of RSDW it depends on the number of depth wise convolution where the depth of the first is $out1$ and the depth of the second is $out2$.

Table 2. The complexity of one layer execution without biastern. Conv represents a normal tensor convolution. In case of RSDW, out1 stands for the output channels of the first depth wise convolution and out2 for the output channels of the last depth wise convolution.

| Method | Complexity |
|--------------------------|--|
| Conv 3×3 | $w*h*3*3*in*out$ |
| Conv 5×5 | $w*h*5*5*in*out$ |
| RIC 3×3 [11] | $4*w*h*3*3*in*out$ |
| ORN-F 3×3 [39] | $F*w*h*3*3*in*out$ |
| ORN-F 5×5 [39] | $F*w*h*5*5*in*out$ |
| RAD 2 (ours) | $w*h*3*3*in*out$ |
| RAD 3 (ours) | $w*h*5*5*in*out$ |
| RSDW 2 (ours) | $(w*h*out1) * ((in) + (3*3) + (out2))$ |
| RSDW 3 (ours) | $(w*h*out1) * ((in) + (5*5) + (out2))$ |
| RING 3×3 (ours) | $w*h*3*3*in*out$ |
| RING 5×5 (ours) | $(8*w*h*in*out)*(3*3 + 2*5*5)$ |

In Table 2 the complexity of a layer is given based on the input depth in and surface $w \times h$. As you can see, only the depth wise convolutions for certain depths can reduce the complexity. All other methods are one constant factor more expensive than ordinary convolutions. In case of ORN [39] this is the predefined number of rotated filters F . For RIC [11] the four integrated rotations and for RING the sum of the rotations of the individual rings ($\sum_{i=2}^f 2 * (i + i - 2)$).

4. Neural Network Models

Figure 5 shows the architectures we used in our experiments. To showcase the applicability of the rotation invariant layers we used diverging architectures. As an example the first model (Figure 5 a) is a small model with batch normalization and the second model (Figure 5 b) a ResNet-34. The first model is to evaluate the applicability of our layers to batch normalization and the second model to show that the layers can also be used with modern residual networks. We also used a small classical neural network (Figure 5 c)). This was used to show the impact of the different rotation invariant layers to models without batch normalization. The last model (Figure 5 d) is a fully convolutional neural network [26] with additional connections between the resolution stages. Those models are called U-Nets [35] and the interconnections improve the semantic segmentation result. This network was only used with the VOC2012 [10] data set and the semantic segmentation task. For training and evaluation we used the DLIB [19] library for deep neural networks. In this library we have also integrated our rotation invariant layers and the state of the art approaches against which we compare the proposed approach.

As can be seen in all models the last layer is always a pooling stage which reduces the output tensor dimension to one. This is to achieve the rotation invariance of the model since the tensor itself contains also spatial information for the following fully connected stages. We will show the difference empirically in our first evaluation.

5. Data sets

In this section all used data sets are described as well as the training parameters and algorithms for weight initialization we have used. In addition, we provide the parameters and optimization procedures as well as the used batch size per data set. We use as little data augmentation as possible to ensure an easy reproduction of our results and described this in detail too.

CIFAR10 [21] is a publicly available data set with a total of 60,000 images. It has ten classes which have to be estimated given an image. Each image in this data set has a resolution of 32×32 and three color channels (red, green, and blue). The training set consists of 50,000 images with 5,000 images per class. For validation 10,000 images are provided with 1,000 images per class.

Training: As optimizer, we used ADAM [20] with weight decay of $5 * 10^{-5}$, momentum one with 0.9 and momentum two with 0.999. The batch size was fixed to fifty during training. The initial learning rate was set to 10^{-3} as well as the training was conducted for 300 epochs. After each 50 epochs the learning rate was reduced by 10^{-1} . The weights of our models were initialized using formula 16 from [13] and all bias terms are set to 0 initially. We did not use

486
487
488
489
490
491
492
493
494
495
496
497
498
499
500
501
502
503
504
505
506
507
508
509
510
511
512
513
514
515
516
517
518
519
520
521
522
523
524
525
526
527
528
529
530
531
532
533
534
535
536
537
538
539

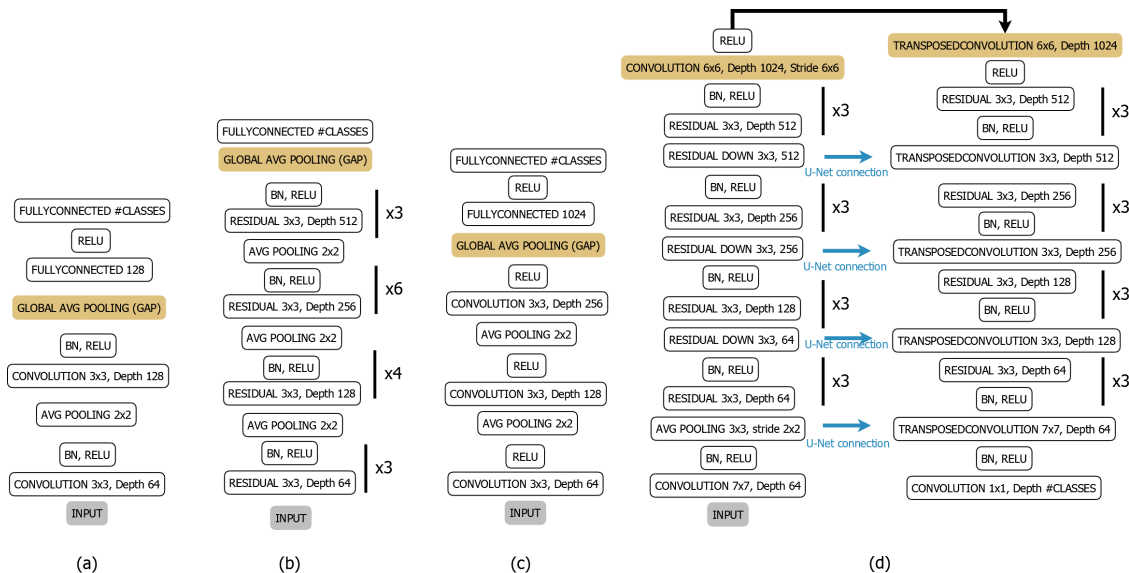


Figure 5. All used architectures in our evaluation. (a) is a small neural network model with batch normalization. (b) represents a ResNet-34 with batch normalization and residual blocks [14]. (c) is a small classical neural network model without batch normalization. (d) is a so called U-Net [35] with residual blocks [14] and fully convolutional [26].

any data augmentation technique during training. For pre-processing we used constant mean subtraction (mean-red 122.782, mean-green 117.001, mean-blue 104.298) and division by 256.0 for the input image. For the evaluation we evaluated the rotations 0° , 90° , 180° , and 270° separately to show the rotation invariance of the proposed approaches. It has to be noted that no rotation was used during training.

CIFAR100 [21] is also a public data set but with one hundred instead of ten classes. It has the same task as in CIFAR10 which is selecting a class given an image. The image resolution for CIFAR100 is 32×32 with three color channels (red, green, and blue). The data set contains 60,000 images and is split into a training and a validation set. The training set contains 50,000 images with 500 examples per class and the validation set contains 10,000 images with 100 examples per class. Therefore, CIFAR100 is a balanced data set with the same amount of images as in CIFAR10.

Training: AS optimizer we used SGD [3] with first momentum 0.9. Weight decay was set to $5 * 10^{-4}$ and we used a fixed batch size of 50 during the training. The initial learning rate was set to 10^{-1} and we trained for 300 epochs in total. After each 50 epochs we reduced the learning rate by 10^{-1} . We did not use any data augmentation but used a image preprocessing. This preprocessing was constant mean subtraction (mean-red 122.782, mean-green 117.001, mean-blue 104.298) and division by 256.0 for the input image. For weight initialization we used formula 16 from [13] and all bias terms are set to 0. For the evaluation we evaluated the rotations 0° , 90° , 180° , and 270° separately to show the rotation invariance of the proposed approaches. It has to be noted that no rotation was used

during training.

VOC2012 [10] is a public available data set. It contains annotations for object detection, classification and semantic segmentation. We used the annotations for semantic segmentation only. In semantic segmentation the task is to classify each pixel in an image regarding his class affiliation. Overall the VOC2012 data set has twenty classes and the background class. Each image can contain multiple objects but not all twenty objects have to be present. It is also possible that the same object class is present multiple times. The training set consists of 3,507 segmented objects on 1,464 images and the validation set has 3,422 segmented objects on 1,449 images. Additionally there is a third image set without any annotations. This set can be used for initializing the weights using unsupervised training like in an auto encoder. In our training and evaluation we did not use the third set. In addition, the data set is unbalanced which makes it even more difficult.

Training: We used a fixed batch size of ten during training. The initial learning rate was set to 10^{-1} . The entire training procedure contained 800 epochs whereby after each 200 epochs the learning rate was reduced by 10^{-1} . As optimizer we used SGD with momentum [33] set to 0.9 and weight decay of $1 * 10^{-4}$. For weight initialization we used formula 16 from [13] and all bias terms are set to 0. For data augmentation we used cropping of 227×227 regions out of each input image. Additionally we used random color offsets during training. As preprocessing of the images we used constant mean subtraction (mean-red 122.782, mean-green 117.001, mean-blue 104.298) and divided each value by 256.0. For the evaluation we evaluated the rotations 0° ,

594
595
596
597
598
599
600
601
602
603
604
605
606
607
608
609
610
611
612
613
614
615
616
617
618
619
620
621
622
623
624
625
626
627
628
629
630
631
632
633
634
635
636
637
638
639
640
641
642
643
644
645
646
647

90°, 180°, and 270° separately to show the rotation invariance of the proposed approaches. It has to be noted that no rotation was used during training.

6. Evaluation

In general, it should be noted here that rotationally invariant features always perform worse compared to rotationally sensitive features if evaluated only on one rotation. This is due to the fact that for some tasks rotation sensitive features contain more general information [12, 39, 29, 30]. However, on average across all rotations, rotationally invariant traits perform better. Since our models were trained without additional data manipulation and rotation, we have limited the evaluation to four simple rotations (0°, 90°, 180°, and 270°), which can be calculated without interpolation. In all tables we give the results for each rotation to show that the trained features are rotation invariant compared to conventional convolutions.

For comparison, we have always replaced all convolution layers with the rotationally invariant layers except for rotational pooling (RP) [12]. Here the authors have shown that it works best if you use it only in the first layer. Therefore, in our evaluation, we also used it only in the first layer together with the rotational invariant layer when &RP is specified. For the RIC [12] approach, we only used filters with a size 3×3 , since no construction rule was specified for larger filters. For ORN [39] we always used all rotations, where the number is specified with $-XY$ each. To make the ORN [39] approach rotationally invariant, we used the ORPooling [39] by the authors.

Table 3 shows the results on CIFAR10 using the model c) from Figure 5. Global pooling has been removed for this evaluation. As you can see, when comparing the entries without rotational pooling [12], the invariant features are not sufficient to make the entire network invariant to a rotation on the input data. One approach to making the entire network invariant is to use rotational pooling [12]. Here, for each different rotation of the filter, the output is rotated as well. In the case of back propagation, the error is also rotated. This makes the nets rotational invariant as seen in Table 3 for the entries with RP. We have used RP together with the invariant features only in the first layer, because the authors in [12] have shown that it works best there. If you compare the results with RP from Table 3 and the results with global pooling from Table 4 you can see that all nets together with the rotation invariant features are rotation invariant. Also, global pooling allows higher accuracy, so we will not use RP for the following evaluations and will use global pooling. Another thing that stands out in the tables 3, 4, and 5 is that RING 3×3 and ORN-8 3×3 [39] always gives the same results. This is because the RING method only rotates one ring here and these are the eight filter rotations of ORN-8 [39]. Thus, for the input 3×3

Table 3. Shows the accuracy on CIFAR10 with the global pooling removed. RP refers to the rotational pooling of [12]. Only the first convolution is replaced by the rotational invariant filters and the rotational pooling is applied subsequently. For all evaluations we used model c) from Figure 5.

| Method | 0° | 90° | 180° | 270° |
|-------------------------------|---------------|---------------|---------------|---------------|
| CNN 3×3 | 79.71% | 28.62% | 33.65% | 29.04% |
| CNN 5×5 [11] | 81.36% | 30.25% | 36.08% | 29.42% |
| RIC 3×3 [11] | 69.74% | 29.98% | 31.11% | 28.59% |
| RIC & RP 3×3 [11] | 53.21% | 53.21% | 53.21% | 53.21% |
| ORN-8 3×3 [39] | 72.89% | 34.66% | 34.36% | 33.65% |
| ORN-16 5×5 [39] | 76.31% | 36.21% | 37.59% | 36.41% |
| ORN-8 & RP 3×3 [39] | 58.17% | 58.17% | 58.17% | 58.17% |
| ORN-16 & RP 5×5 [39] | 61.48% | 61.48% | 61.48% | 61.48% |
| RAD 2 (ours) | 65.05% | 29.14% | 30.99% | 29.85% |
| RAD 3 (ours) | 68.92% | 25.61% | 34.68% | 32.76% |
| RAD 2 & RP (ours) | 50.26% | 50.26% | 50.26% | 50.26% |
| RAD 3 & RP (ours) | 57.35% | 57.35% | 57.35% | 57.35% |
| RSDW 2 (ours) | 62.21% | 26.43% | 29.87% | 27.39% |
| RSDW 3 (ours) | 66.71% | 26.41% | 31.35% | 29.18% |
| RSDW 2 & RP (ours) | 48.01% | 48.01% | 48.01% | 48.01% |
| RSDW 3 & RP (ours) | 55.76% | 55.76% | 55.76% | 55.76% |
| RING 3×3 (ours) | 72.89% | 34.66% | 34.36% | 33.65% |
| RING 5×5 (ours) | 79.58% | 37.93% | 39.90% | 38.59% |
| RING & RP 3×3 (ours) | 58.17% | 58.17% | 58.17% | 58.17% |
| RING & RP 5×5 (ours) | 62.90% | 62.90% | 62.90% | 62.90% |

both approaches are identical in our evaluation, since for ORN [39] we used all rotations of the outermost filter ring, whereas for ORN [39] the always the complete filter was rotated, which makes a difference for the filter size of 5×5 .

In the Tables 4 and 5 the results can be seen on CIFAR10 and CIFAR100 for the models a), b), and c) with global pooling. All rotationally invariant layers and approaches, show the rotational invariance clearly in comparison to ordinary convolutions, as they give the same results for each rotation. The least expensive and least parameterized method in our evaluation is RSDW, which uses the approach of depth wise convolutions [] together with radial convolutions (RAD). RSDW also has the worst results of all rotation invariants but is still better than conventional convolutions on average over all rotations. The RAD approach significantly improves the results and has both lower calculation costs and parameters compared to RIC [12]. RIC, on the other hand, is always significantly better than RAD with a radius of 2 and for model b) even better than RAD with radius 3. If RIC [12] is compared directly with RAD, the calculation costs for RIC [12] are four times as high and the number of parameters in the case of RAD with radius 2 is almost three times as high.

Likewise, it can be seen in the Tables 4 and 5 that for the orientation 0° the CNN with a filter size of 5×5 always performs best. This is also the orientation of the training

648
649
650
651
652
653
654
655
656
657
658
659
660
661
662
663
664
665
666
667
668
669
670
671
672
673
674
675
676
677
678
679
680
681
682
683
684
685
686
687
688
689
690
691
692
693
694
695
696
697
698
699
700
701

702
703
704
705
706
707
708
709
710
711
712
713
714
715
716
717
718
719
720
721
722
723
724
725
726
727
728
729
730
731
732
733
734
735
736
737
738
739
740
741
742
743
744
745
746
747
748
749
750
751
752
753
754
755

756 Table 4. Shows the accuracy on CIFAR10 where the letters a),
757 b), and c) denote the models in Figure 5. Where possible, each
758 method was evaluated with different filter sizes $X \times X$.

| 759 Method | 0° | 90° | 180° | 270° |
|---------------------------|---------------|---------------|---------------|---------------|
| 760 CNN 3 × 3(a) | 77.85% | 35.84% | 47.38% | 35.61% |
| 761 CNN 5 × 5(a) | 78.20% | 35.20% | 43.80% | 35.96% |
| 762 RIC 3 × 3(a) [11] | 66.73% | 66.73% | 66.73% | 66.73% |
| 763 ORN-8 3 × 3(a) [39] | 74.88% | 74.88% | 74.88% | 74.88% |
| 764 ORN-16 5 × 5(ab) [39] | 75.15% | 75.15% | 75.15% | 75.15% |
| 765 RAD 2(ours) (a) | 64.41% | 64.41% | 64.41% | 64.41% |
| 766 RAD 3(ours) (a) | 67.17% | 67.17% | 67.17% | 67.17% |
| 767 RSDW 2(ours) (a) | 60.56% | 60.56% | 60.56% | 60.56% |
| 768 RSDW 3(ours) (a) | 62.75% | 62.75% | 62.75% | 62.75% |
| 769 RING 3 × 3(ours) (a) | 74.88% | 74.88% | 74.88% | 74.88% |
| 770 RING 5 × 5(ours) (a) | 76.10% | 76.10% | 76.10% | 76.10% |
| 771 CNN 3 × 3(b) | 81.93% | 30.61% | 34.82% | 31.79% |
| 772 CNN 5 × 5(b) | 87.78% | 33.56% | 37.34% | 30.39% |
| 773 RIC 3 × 3(b) [11] | 70.74% | 70.74% | 70.74% | 70.74% |
| 774 ORN-8 3 × 3(b) [39] | 77.63% | 77.63% | 77.63% | 77.63% |
| 775 ORN-16 5 × 5(b) [39] | 79.03% | 79.03% | 79.03% | 79.03% |
| 776 RAD 2(ours) (b) | 67.06% | 67.06% | 67.06% | 67.06% |
| 777 RAD 3(ours) (b) | 69.83% | 69.83% | 69.83% | 69.83% |
| 778 RSDW 2(ours) (b) | 64.45% | 64.45% | 64.45% | 64.45% |
| 779 RSDW 3(ours) (b) | 65.82% | 65.82% | 65.82% | 65.82% |
| 780 RING 3 × 3(ours) (b) | 77.63% | 77.63% | 77.63% | 77.63% |
| 781 RING 5 × 5(ours) (b) | 79.96% | 79.96% | 79.96% | 79.96% |
| 782 CNN 3 × 3(c) | 81.01% | 31.61% | 39.12% | 32.06% |
| 783 CNN 5 × 5(c) | 84.36% | 32.25% | 38.08% | 26.42% |
| 784 RIC 3 × 3(c) [11] | 66.31% | 66.31% | 66.31% | 66.31% |
| 785 ORN-8 3 × 3(c) [39] | 74.27% | 74.27% | 74.27% | 74.27% |
| 786 ORN-16 5 × 5(c) [39] | 75.68% | 75.68% | 75.68% | 75.68% |
| 787 RAD 2(ours) (c) | 65.70% | 65.70% | 65.70% | 65.70% |
| 788 RAD 3(ours) (c) | 66.84% | 66.84% | 66.84% | 66.84% |
| 789 RSDW 2(ours) (c) | 61.40% | 61.40% | 61.40% | 61.40% |
| 790 RSDW 3(ours) (c) | 64.25% | 64.25% | 64.25% | 64.25% |
| 791 RING 3 × 3(ours) (c) | 74.27% | 74.27% | 74.27% | 74.27% |
| 792 RING 5 × 5(ours) (c) | 76.83% | 76.83% | 76.83% | 76.83% |

793 data. For all other orientations RING with a filter size of
794 5×5 is best. The second best is ORN-16 with a filter size
795 of 5×5 and needs only $\frac{2}{3}$ of run time compared to RING.

796 As you can see in Table 6, the rotation invariant fil-
797 ters also work for semantic segmentation. Note that we
798 also scaled the validation data to 227×227 to get a one-
799 dimensional vector in the central part of the network. As
800 you can see, the resource-saving method RAD is already
801 better than the conventional convolution and the more ex-
802 pensive method RING improves the results significantly.

803 7. Conclusion

804 In this work we have shown several new approaches for
805 the rotationally invariant feature extraction in deep neural
806

810 Table 5. Shows the accuracy on CIFAR100 where the letters a),
811 b), and c) denote the models in Figure 5. Where possible, each
812 method was evaluated with different filter sizes $X \times X$.

| 813 Method | 0° | 90° | 180° | 270° |
|--------------------------|---------------|---------------|---------------|---------------|
| 814 CNN 3 × 3(a) | 48.97% | 23.30% | 30.75% | 24.05% |
| 815 CNN 5 × 5(a) | 49.83% | 24.41% | 32.99% | 25.50% |
| 816 RIC 3 × 3(a) [11] | 39.10% | 39.10% | 39.10% | 39.10% |
| 817 ORN-8 3 × 3(a) [39] | 47.41% | 47.41% | 47.41% | 47.41% |
| 818 ORN-16 5 × 5(a) [39] | 48.50% | 48.50% | 48.50% | 48.50% |
| 819 RAD 2(ours) (a) | 37.66% | 37.66% | 37.66% | 37.66% |
| 820 RAD 3(ours) (a) | 39.51% | 39.51% | 39.51% | 39.51% |
| 821 RSDW 2(ours) (a) | 33.56% | 33.56% | 33.56% | 33.56% |
| 822 RSDW 3(ours) (a) | 34.20% | 34.20% | 34.20% | 34.20% |
| 823 RING 3 × 3(ours) (a) | 47.41% | 47.41% | 47.41% | 47.41% |
| 824 RING 5 × 5(ours) (a) | 49.03% | 49.03% | 49.03% | 49.03% |
| 825 CNN 3 × 3(b) | 67.31% | 20.29% | 21.16% | 19.92% |
| 826 CNN 5 × 5(b) | 71.20% | 19.39% | 22.73% | 18.90% |
| 827 RIC 3 × 3(b) [11] | 49.12% | 49.12% | 49.12% | 49.12% |
| 828 ORN-8 3 × 3(b) [39] | 57.43% | 57.43% | 57.43% | 57.43% |
| 829 ORN-16 5 × 5(b) [39] | 60.39% | 60.39% | 60.39% | 60.39% |
| 830 RAD 2(ours) (b) | 45.20% | 45.20% | 45.20% | 45.20% |
| 831 RAD 3(ours) (b) | 48.33% | 48.33% | 48.33% | 48.33% |
| 832 RSDW 2(ours) (b) | 34.82% | 34.82% | 34.82% | 34.82% |
| 833 RSDW 3(ours) (b) | 36.91% | 36.91% | 36.91% | 36.91% |
| 834 RING 3 × 3(ours) (b) | 57.43% | 57.43% | 57.43% | 57.43% |
| 835 RING 5 × 5(ours) (b) | 62.40% | 62.40% | 62.40% | 62.40% |
| 836 CNN 3 × 3(c) | 48.53% | 18.37% | 22.29% | 18.14% |
| 837 CNN 5 × 5(c) | 50.30% | 19.81% | 25.01% | 18.92% |
| 838 RIC 3 × 3(c) [11] | 36.01% | 36.01% | 36.01% | 36.01% |
| 839 ORN-8 3 × 3(c) [39] | 43.86% | 43.86% | 43.86% | 43.86% |
| 840 ORN-16 5 × 5(c) [39] | 45.72% | 45.72% | 45.72% | 45.72% |
| 841 RAD 2(ours) (c) | 35.06% | 35.06% | 35.06% | 35.06% |
| 842 RAD 3(ours) (c) | 37.95% | 37.95% | 37.95% | 37.95% |
| 843 RSDW 2(ours) (c) | 30.19% | 30.19% | 30.19% | 30.19% |
| 844 RSDW 3(ours) (c) | 32.73% | 32.73% | 32.73% | 32.73% |
| 845 RING 3 × 3(ours) (c) | 43.86% | 43.86% | 43.86% | 43.86% |
| 846 RING 5 × 5(ours) (c) | 46.89% | 46.89% | 46.89% | 46.89% |

847 Table 6. Shows the pixel wise accuracy on VOC2012. The model
848 for all evaluations is d) from Figure 5.

| 849 Method | 0° | 90° | 180° | 270° |
|--------------------------|---------------|---------------|---------------|---------------|
| 850 CNN 3 × 3(d) | 84.65% | 62.85% | 73.66% | 53.78% |
| 851 RAD 3(ours) (d) | 71.25% | 71.28% | 71.81% | 72.02% |
| 852 RING 3 × 3(ours) (d) | 79.93% | 79.16% | 79.85% | 80.12% |

853 networks. We compared our approaches with the state of
854 the art on CIFAR10 and CIFAR100. Different models were
855 trained and no rotation was applied to the training data.
856 As the evaluation with four rotations has shown, our ap-
857 proaches are also rotation invariant. Additionally, we tested
858 our approaches in fully convolutional neural networks for
859 semantic segmentation on the VOC2012 data set. In addi-
860 tion to the description and evaluation we provide the CUDA
861
862
863

implementations of our approaches. As well as a description of how to integrate them into existing frameworks. Future work will go into the direction of head mounted eye trackers, since there the eye cameras can be shifted and rotated. There we want to show the advantages of the rotation invariant approaches even further. Especially for the transfer between training on synthetic images and the usage with real eyes.

References

- [1] Mohamed Abd El Aziz, IM Selim, and Shengwu Xiong. Automatic detection of galaxy type from datasets of galaxies image based on image retrieval approach. *Scientific Reports*, 7(1):1–9, 2017. 1
- [2] Vincent Andrearczyk, Julien Fageot, Valentin Oreiller, Xavier Montet, and Adrien Depeursinge. Local rotation invariance in 3d cnns. *arXiv preprint arXiv:2003.08890*, 2020. 1, 3
- [3] Léon Bottou. Stochastic gradient learning in neural networks. *Proceedings of Neuro-Nimes*, 91(8):12, 1991. 1, 6
- [4] Dan Ciregan, Ueli Meier, and Jürgen Schmidhuber. Multi-column deep neural networks for image classification. In *2012 IEEE conference on computer vision and pattern recognition*, pages 3642–3649. IEEE, 2012. 1, 2
- [5] Taco Cohen and Max Welling. Group equivariant convolutional networks. In *International conference on machine learning*, pages 2990–2999, 2016. 1, 2, 3
- [6] Taco S Cohen and Max Welling. Steerable cnns. *arXiv preprint arXiv:1612.08498*, 2016. 1, 2, 3
- [7] Jifeng Dai, Haozhi Qi, Yuwen Xiong, Yi Li, Guodong Zhang, Han Hu, and Yichen Wei. Deformable convolutional networks. In *Proceedings of the IEEE international conference on computer vision*, pages 764–773, 2017. 1, 2
- [8] Dima Damen, Hazel Doughty, Giovanni Maria Farinella, Sanja Fidler, Antonino Furnari, Evangelos Kazakos, Davide Moltisanti, Jonathan Munro, Toby Perrett, Will Price, and Michael Wray. Scaling egocentric vision: The epic-kitchens dataset. In *European Conference on Computer Vision (ECCV)*, 2018. 1
- [9] Sander Dieleman, Kyle W Willett, and Joni Dambre. Rotation-invariant convolutional neural networks for galaxy morphology prediction. *Monthly notices of the royal astronomical society*, 450(2):1441–1459, 2015. 1, 2
- [10] M. Everingham, L. Van Gool, C. K. I. Williams, J. Winn, and A. Zisserman. The PASCAL Visual Object Classes Challenge 2012 (VOC2012) Results. <http://www.pascal-network.org/challenges/VOC/voc2012/workshop/index.html>. 5, 6
- [11] Patrick Follmann and Tobias Bottger. A rotationally-invariant convolution module by feature map back-rotation. In *2018 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 784–792. IEEE, 2018. 1, 3, 5, 7, 8
- [12] Patrick Follmann and Tobias Bottger. A rotationally-invariant convolution module by feature map back-rotation. In *2018 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 784–792. IEEE, 2018. 2, 3, 7
- [13] Xavier Glorot and Yoshua Bengio. Understanding the difficulty of training deep feedforward neural networks. In *Proceedings of the thirteenth international conference on artificial intelligence and statistics*, pages 249–256, 2010. 5, 6
- [14] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 1, 6
- [15] Geoffrey E Hinton, Alex Krizhevsky, and Sida D Wang. Transforming auto-encoders. In *International conference on artificial neural networks*, pages 44–51. Springer, 2011. 1, 2
- [16] Andrew G Howard, Menglong Zhu, Bo Chen, Dmitry Kalenichenko, Weijun Wang, Tobias Weyand, Marco Andreetto, and Hartwig Adam. Mobilenets: Efficient convolutional neural networks for mobile vision applications. *arXiv preprint arXiv:1704.04861*, 2017. 2
- [17] Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. *arXiv preprint arXiv:1502.03167*, 2015. 1
- [18] Max Jaderberg, Karen Simonyan, Andrew Zisserman, et al. Spatial transformer networks. In *Advances in neural information processing systems*, pages 2017–2025, 2015. 1, 2
- [19] Davis E King. Dlib-ml: A machine learning toolkit. *Journal of Machine Learning Research*, 10(Jul):1755–1758, 2009. 5
- [20] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. 1, 5
- [21] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009. 5, 6
- [22] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pages 1097–1105, 2012. 1
- [23] Dmitry Laptev, Nikolay Savinov, Joachim M Buhmann, and Marc Pollefeys. Ti-pooling: transformation-invariant pooling for feature learning in convolutional neural networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 289–297, 2016. 1, 2
- [24] Yann LeCun, Yoshua Bengio, et al. Convolutional networks for images, speech, and time series. *The handbook of brain theory and neural networks*, 3361(10):1995, 1995. 1
- [25] Rongjian Li, Tao Zeng, Hanchuan Peng, and Shuiwang Ji. Deep learning segmentation of optical microscopy images improves 3-d neuron reconstruction. *IEEE transactions on medical imaging*, 36(7):1533–1541, 2017. 1
- [26] Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully convolutional networks for semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3431–3440, 2015. 1, 5, 6
- [27] Alessandra Lumini and Loris Nanni. Deep learning and transfer learning features for plankton classification. *Ecological informatics*, 51:33–43, 2019. 1
- [28] Aleksander Madry, Aleksandar Makelov, Ludwig Schmidt, Dimitris Tsipras, and Adrian Vladu. Towards deep learning models resistant to adversarial attacks. *arXiv preprint arXiv:1706.06083*, 2017. 1

918
919
920
921
922
923
924
925
926
927
928
929
930
931
932
933
934
935
936
937
938
939
940
941
942
943
944
945
946
947
948
949
950
951
952
953
954
955
956
957
958
959
960
961
962
963
964
965
966
967
968
969
970
971

| | | |
|------|--|------|
| 972 | [29] Diego Marcos, Michele Volpi, Nikos Komodakis, and Devis Tuia. Rotation equivariant vector field networks. In <i>Proceedings of the IEEE International Conference on Computer Vision</i> , pages 5048–5057, 2017. 1, 2, 3, 7 | 1026 |
| 973 | | 1027 |
| 974 | | 1028 |
| 975 | | 1029 |
| 976 | [30] Diego Marcos, Michele Volpi, and Devis Tuia. Learning rotation invariant convolutional filters for texture classification. In <i>2016 23rd International Conference on Pattern Recognition (ICPR)</i> , pages 2012–2017. IEEE, 2016. 1, 2, 3, 7 | 1030 |
| 977 | | 1031 |
| 978 | | 1032 |
| 979 | | 1033 |
| 980 | [31] Javier Marin, Aritro Biswas, Ferda Ofli, Nicholas Hynes, Amaia Salvador, Yusuf Aytar, Ingmar Weber, and Antonio Torralba. Recipe1m+: A dataset for learning cross-modal embeddings for cooking recipes and food images. <i>IEEE transactions on pattern analysis and machine intelligence</i> , 2019. 1 | 1034 |
| 981 | | 1035 |
| 982 | | 1036 |
| 983 | | 1037 |
| 984 | | 1038 |
| 985 | | 1039 |
| 986 | [32] Lerrel Pinto, Dhiraj Gandhi, Yuanfeng Han, Yong-Lae Park, and Abhinav Gupta. The curious robot: Learning visual representations via physical interactions. In <i>European Conference on Computer Vision</i> , pages 3–18. Springer, 2016. 1 | 1040 |
| 987 | | 1041 |
| 988 | | 1042 |
| 989 | [33] Ning Qian. On the momentum term in gradient descent learning algorithms. <i>Neural networks</i> , 12(1):145–151, 1999. 1, 6 | 1043 |
| 990 | | 1044 |
| 991 | | 1045 |
| 992 | [34] Rosemberg Rodriguez, Eva Dokladalova, and Petr Dokládál. Rotation invariant cnn using scattering transform for image classification. In <i>2019 IEEE International Conference on Image Processing (ICIP)</i> , pages 654–658. IEEE, 2019. 1, 2 | 1046 |
| 993 | | 1047 |
| 994 | | 1048 |
| 995 | | 1049 |
| 996 | [35] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In <i>International Conference on Medical image computing and computer-assisted intervention</i> , pages 234–241. Springer, 2015. 5, 6 | 1050 |
| 997 | | 1051 |
| 998 | | 1052 |
| 999 | | 1053 |
| 1000 | | 1054 |
| 1001 | [36] Rosemberg Rodriguez Salas, Petr Dokládál, and Eva Dokladalova. Red-nn: Rotation-equivariant deep neural network for classification and prediction of rotation. 2019. 1, 2 | 1055 |
| 1002 | | 1056 |
| 1003 | | 1057 |
| 1004 | [37] Daniel E Worrall, Stephan J Garbin, Daniyar Turmukhambetov, and Gabriel J Brostow. Harmonic networks: Deep translation and rotation equivariance. In <i>Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition</i> , pages 5028–5037, 2017. 1, 2, 3 | 1058 |
| 1005 | | 1059 |
| 1006 | | 1060 |
| 1007 | | 1061 |
| 1008 | [38] Yuxin Wu and Kaiming He. Group normalization. In <i>Proceedings of the European conference on computer vision (ECCV)</i> , pages 3–19, 2018. 1 | 1062 |
| 1009 | | 1063 |
| 1010 | | 1064 |
| 1011 | [39] Yanzhao Zhou, Qixiang Ye, Qiang Qiu, and Jianbin Jiao. Oriented response networks. In <i>Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition</i> , pages 519–528, 2017. 1, 2, 3, 4, 5, 7, 8 | 1065 |
| 1012 | | 1066 |
| 1013 | | 1067 |
| 1014 | | 1068 |
| 1015 | | 1069 |
| 1016 | | 1070 |
| 1017 | | 1071 |
| 1018 | | 1072 |
| 1019 | | 1073 |
| 1020 | | 1074 |
| 1021 | | 1075 |
| 1022 | | 1076 |
| 1023 | | 1077 |
| 1024 | | 1078 |
| 1025 | | 1079 |