

# Patch Reordering: A Novel Way to Achieve Rotation and Translation Invariance in Convolutional Neural Networks

Xu Shen,<sup>†</sup> Xinmei Tian,<sup>†</sup> Shaoyan Sun,<sup>†</sup> Dacheng Tao<sup>‡</sup>

<sup>†</sup> CAS Key Laboratory of Technology in Geo-spatial Information Processing and Application System, University of Science and Technology of China Hefei, Anhui, China 230027

<sup>‡</sup> Centre for Artificial Intelligence and the Faculty of Engineering and Information Technology, University of Technology Sydney, 81 Broadway Street, Ultimo, NSW 2007, Australia

## Abstract

Convolutional Neural Networks (CNNs) have demonstrated state-of-the-art performance on many visual recognition tasks. However, the combination of convolution and pooling operations only shows invariance to small local location changes in meaningful objects in input. Sometimes, such networks are trained using data augmentation to encode this invariance into the parameters, which restricts the capacity of the model to learn the content of these objects. A more efficient use of the parameter budget is to encode rotation or translation invariance into the model architecture, which relieves the model from the need to learn them. To enable the model to focus on learning the content of objects other than their locations, we propose to conduct patch ranking of the feature maps before feeding them into the next layer. When patch ranking is combined with convolution and pooling operations, we obtain consistent representations despite the location of meaningful objects in input. We show that the patch ranking module improves the performance of the CNN on many benchmark tasks, including MNIST digit recognition, large-scale image recognition, and image retrieval.

## Introduction

In recent years, convolutional neural networks (CNNs) have achieved state-of-the-art performance on many computer vision tasks, including image recognition (Szegedy et al. 2015; Simonyan and Zisserman 2014a; He et al. 2016), semantic segmentation (Long, Shelhamer, and Darrell 2015), image captioning (Karpathy and Li 2015; Donahue et al. 2015; Fang et al. 2015), action recognition (Gkioxari, Girshick, and Malik 2015; Simonyan and Zisserman 2014b), and video captioning (Yao et al. 2015; Pan et al. 2016). The success of CNNs comes from their ability to learn the two-dimensional structures of images for which objects and patterns may appear at different locations. To detect and learn patterns despite their locations, the weights of local filters are shared when applied to different positions in the image.

Since distortions or shifts of the input can cause the positions of salient features to vary, weight sharing is very important for CNNs to detect invariant elementary features regardless of location changes of these features (Lecun et al. 1998). In addition, pooling also reduces the sensitivity of the

output to small local shifts and distortions by reducing the resolutions of input feature maps. However, another important property of weight sharing or pooling is that the location of detected features in the output feature maps is identical to that of the corresponding local patch in the input feature maps. As a result, the location change of the input visual patterns in lower layers will propagate to higher convolutional layers. Due to the typically small local spatial support for pooling (e.g.,  $2 \times 2$ ) and convolution (e.g.,  $9 \times 9$  kernel size), large global location changes of patterns in input (e.g., global rotation or translation of objects) will even propagate to the feature maps of the final convolutional layer (as shown in Fig. 1). Consequently, the following fully connected layers have to learn the location invariance to produce consistent predictions or representations, which restricts the use of the parameter budget for achieving more powerful outputs.

In this paper, we introduce a *Patch Reordering (PR)* module that can be embedded into a standard CNN architecture to improve the rotation and translation invariance capabilities. Output feature maps of the convolutional layer are first divided into multiple tiers of non-overlapped local patches at different spatial pyramid levels. We reorder these local patches at each level based on their energy (e.g., L1 or L2 norm of activations of the patch). To retain the spatial consistency of local patterns, we only reorder the patches of a given level locally (i.e., within each single patch of its upper level). In convolutional layers, a location change of the patterns in input feature maps will result in a corresponding location change of the output feature maps, while the local patterns (activations) in the output are equivalent. As a result, ranking these local patterns in a specific order leads to a consistent representation despite the locations of local patterns in input, that is, rotation or translation invariance. The proposed architecture can be inserted into any convolutional layers and allows for end-to-end training of the models for which they are applied. In addition, we do not need any extra training supervision or modification to the training process or any preprocessing of input images.

## Related Work

The equivalence and invariance of CNN representations to input image transformations were investigated in (Lenc and Vedaldi 2015; Cohen and Welling 2015; Gens and Domingos 2014). Specifically, Cohen and Welling (Cohen and

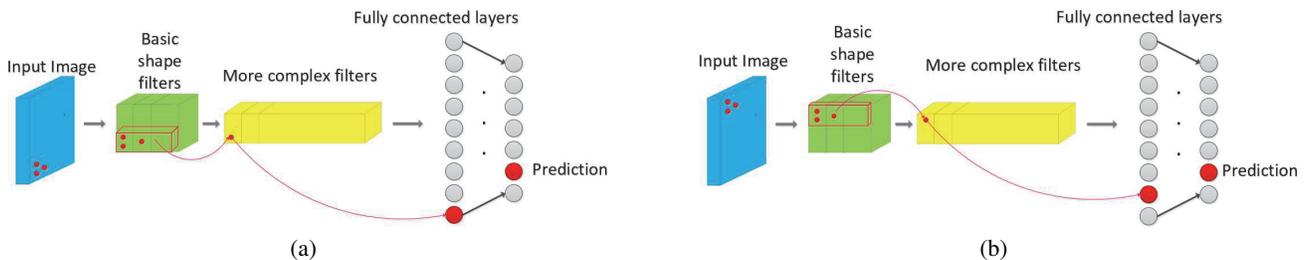


Figure 1: Large global location changes of patterns in input (e.g., global rotation or translation of objects) will propagate to the feature maps of the final convolutional layer. As a result, the fully connected layers have to encode the invariance of location changes into its parameters, which restricts the capacity of the model to learn the content of these objects.

Welling 2015) showed that a linear transform of a good visual representation was equivalent to a combination of the elementary irreducible representations using the theory of group representations. Lenc and Vedaldi (Lenc and Vedaldi 2015) estimated the linear relationships between representations of the original and transformed images. Gens and Domingos (Gens and Domingos 2014) proposed a generalization of CNNs that formed feature maps over arbitrary symmetry groups based on the theory of symmetry groups in (Lenc and Vedaldi 2015), resulting in feature maps that were more invariant to symmetry groups. Bruna and Mallat (Bruna and Mallat 2013) proposed a wavelet scattering network to compute a translation invariant image representation. Local linear transformations were adopted in the feature learning algorithms in (Sohn and Lee 2012) for the purpose of transformation-invariant feature learning.

Numerous recent works have focused on introducing spatial invariance in deep learning architectures explicitly. For unsupervised feature learning, Sohn and Lee (Sohn and Lee 2012) presented a transform-invariant restricted Boltzmann machine that compactly represented data by its weights and their transformations, which achieved invariance of the feature representation via probabilistic max pooling. Each hidden unit was augmented with a latent transformation assignment variable that described the selection of the transformed view of the weights associated with the unit in (Kivinen and Williams 2011). In both works, the transformed filters were only applied at the center of the largest receptive field size. In tied convolutional neural networks (Le et al. 2010), invariance was learned explicitly by square-root pooling hidden units computed by partially un-tied weights. Here, additional learned parameters were needed when untying weights.

The latest two works on incorporating spatial invariance in CNNs are described in (Kanazawa and Sharma 2014; Jaderberg et al. 2015). In (Kanazawa and Sharma 2014), feature maps in CNNs were scaled or rotated to multiple levels, and the same kernel was convolved across the input at each scale. Then, the responses of the convolution at each scale were normalized and pooled at each spatial location to obtain a locally scale-invariant representation. In this model, only limited scales were considered, and extra modules were needed in the feature extraction process. To address different transformation types in input images, Jaderberg *et al.*

(Jaderberg et al. 2015) proposed inserting a spatial transformer module between CNN layers, which explicitly transformed an input image into a proper appearance and fed the transformed input into the CNN model.

In conclusion, all aforementioned related works improve the transform invariance of deep learning models by adding *extra feature extraction modules, more learnable parameters, or extra transformations on input images*, which makes the trained CNN model problem-dependent and not generalizable to other datasets. In contrast, in this paper, we propose a very simple reordering on feature maps during the training of CNN models. No extra feature extraction modules or more learnable parameters are needed. Therefore, it is very easy to apply the trained model to other vision tasks.

## Patch Reordering in Convolutional Neural Networks

Weight sharing in CNNs allows feature detectors to detect features regardless of their spatial locations in the image; however, the corresponding location of output patterns varies when subject to location changes of the local patterns in the input. Learning invariant representations causes parameter redundancy problems in current CNN models. In this section, we will reveal this phenomenon and propose the formulation of our *Patch Reordering* module.

### Parameter Redundancy in Convolutional Neural Networks

Let  $X^{l-1} = \{\vec{x}_{ij}, i = 1, \dots, h, j = 1, \dots, w\}$  denote the output feature maps of a convolutional layer with  $m$  elements ( $m = h \times w \times c$  for  $c$  feature maps with height  $h$  and width  $w$ ). Each  $\vec{x}_{ij}$  is a  $c$ -dimensional input feature vector corresponding to location  $(i, j)$ . If it is followed by a fully connected layer,  $X^l$  can be computed by

$$X^l = f(WX^{l-1}) = f\left(\sum_{i=1}^h \sum_{j=1}^w \vec{w}_{ij}^T \vec{x}_{ij}\right), \quad (1)$$

where  $f$  is a non-linear activation function and  $\vec{w}_{ij}$  are the weights for location  $(i, j)$ . If there is some location change (such as a rotation or translation) of the input features, the resulting new input becomes  $X'^{l-1} = \{\vec{x}'_{i'j'}, i' = 1, \dots, h, j' = 1, \dots, w\}$ . Since there are no value changes

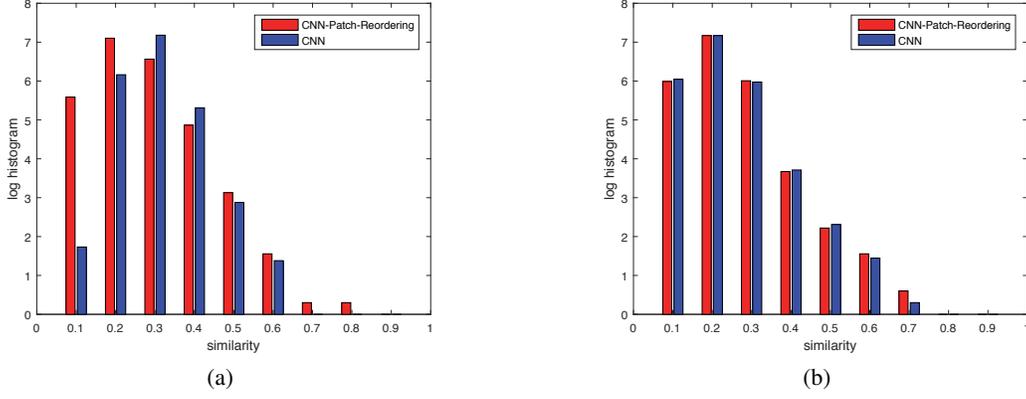


Figure 2: Log histogram of similarity between weights in fc6 (a) and fc7 (b) layers in a CNN with (CNN-Patch-Reordering) and without the (CNN) patch reordering module. In conventional CNNs, the correlation of parameters in fc6 is much higher than that in fc7, while that for PR-CNN is quite consistent. This phenomenon implies that the location change of input visual patterns leads to a higher parameter redundancy in the subsequent layers.

(except cropping or padding), for  $\vec{x}_{ij}$  in any position  $(i, j)$ , we can always find its correspondence  $\vec{x}'_{i'j'}$  in the transformed input, i.e.  $\vec{x}_{ij} = \vec{x}'_{i'j'}$ . If the network learns to be invariant under this type of location change, the output (or representation) should remain the same. Specifically,

$$X^l \approx X^l \Rightarrow f\left(\sum_{i'=1}^h \sum_{j'=1}^w \vec{w}_{i'j'}^T \vec{x}'_{i'j'}\right) \approx f\left(\sum_{i=1}^h \sum_{j=1}^w \vec{w}_{ij}^T \vec{x}_{ij}\right). \quad (2)$$

Then, in the monotonous section of  $f$ , we have

$$\sum_{i'=1}^h \sum_{j'=1}^w \vec{w}_{i'j'}^T \vec{x}'_{i'j'} \approx \sum_{i=1}^h \sum_{j=1}^w \vec{w}_{ij}^T \vec{x}_{ij}. \quad (3)$$

Since  $\vec{x}_{ij} = \vec{x}'_{i'j'}$ , the aforementioned equation can be simplified as:

$$\sum_{i'=1}^h \sum_{j'=1}^w (\vec{w}_{i'j'} - \vec{w}_{ij})^T \vec{x}_{ij} \approx \vec{0} \quad (4)$$

Because  $\vec{x}_{ij}$  varies as the input image changes, we have  $\vec{w}_{i'j'} - \vec{w}_{ij} \approx \vec{0}$ . That is to say, encoding rotation or translation invariance into CNNs leads to highly correlated parameters in higher layers. Therefore, the capacity of the model decreases. To validate this redundancy in CNN models, we compare the log histogram of cosine similarities between weights in fc6 and fc7 in AlexNet (Krizhevsky, Sutskever, and Hinton 2012). Fig. 2 shows that parameter redundancy of the model is significantly reduced because of a more consistent feature map after patch reordering.

### Patch Reordering

If one object is located at different positions in two images, the same visual features of the object will locate at different positions in their corresponding convolution feature maps.

The feature maps generated by deep convolutional layers are analogous to the feature maps in traditional methods (Chatfield et al. 2011; Coates and Ng 2011). In those methods, image patches or SIFT vectors are densely extracted and then encoded. These encoded features compose the feature maps and are pooled into a histogram of bins. Reordering of the pooled histogram achieves translation and rotation invariance. Likewise, since the deep convolutional feature maps are the encoded representations of images, reordering can be applied in a similar way.

Since convolutional kernels function as feature detectors, each activation in the output feature maps corresponds to a match of a specific visual pattern. Therefore, when the feature detectors slide through the whole input feature maps, the locations with matched patterns generate very high responses and *vice versa*. Consequently, the “energy” distribution (L1 norm or L2 norm) of the local patches in the output feature maps presents some heterogeneity. Furthermore, patches with different energies correspond to different parts of the input object. Naturally, if we rank the patches by their energies in a descending or ascending order, regardless of how we change the location of visual patterns by rotation or translation in the input, the output order will be quite consistent. Finally, and rotation- and translation-invariant representation is generated.

**Forward Propagation** The details of the patch reordering module are illustrated in Fig. 3. The feature maps are divided into  $(n \times n)^l$  non-overlapped patches at level- $l$ . Here,  $n$  is a predefined parameter (e.g., 2 or 3). Then, we rank the patches by energy (L1 or L2 norm) within each patch of level  $l - 1$ :

$$En_{i \in (w_1^{l-1}, w_2^{l-1}); j \in (h_1^{l-1}, h_2^{l-1})}^{l-1} = \sum_{i=w_1^{l-1}}^{w_2^{l-1}} \sum_{j=h_1^{l-1}}^{h_2^{l-1}} |X_{i,j}| \text{ or } X_{i,j}^2. \quad (5)$$

The patches are located from the upper left to the lower right in descending order of energy. The offset of each pixel in

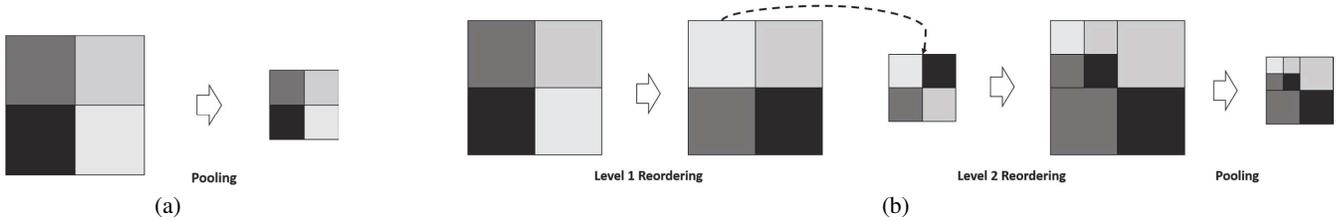


Figure 3: Side-by-side comparison of the structure of (a) a conventional convolutional layer and (b) the proposed convolutional layer with patch reordering module. Feature maps are divided into  $(n \times n)^l$  non-overlapped patches at level- $l$ . Here, we take  $n = 2$  and  $l = 2$  as an example. The four patches of level 1 are first reordered so that patches with higher energy precedes other patches. Then, we repeat this process within each single patch in the previous level (here, we only show the reordering of the first patch in level 2). One visual example of the output feature maps between original and rotated/translated input images with/without patch reordering is illustrate in Supplemental Material.

the patch  $(o_i^l, o_j^l)$  can be obtained from the gap between the target patch location and the source patch location. Finally, the output feature map can be computed by

$$Z_{i,j} = X_{i+\sum_{k=1}^l o_i^k, j+\sum_{k=1}^l o_j^k}. \quad (6)$$

**Backward Propagation** During the back-propagation process, we simply pass the error from the output pixel to its corresponding input pixel:

$$E_{i,j}^X = E_{i-\sum_{k=1}^l o_i^k, j-\sum_{k=1}^l o_j^k}^Z. \quad (7)$$

## Experiments

In this section, we evaluate our proposed CNN with patch reordering module on several supervised learning tasks, and compare our model with state-of-the-art methods, including traditional CNNs, SI-CNN (Kanazawa and Sharma 2014), and ST-CNN (Jaderberg et al. 2015). First, we conduct experiments on the distorted versions of the MNIST handwriting dataset as in (Jaderberg et al. 2015; Kanazawa and Sharma 2014). The experimental results show that patch reordering is capable of achieving comparable or better classification performance. Second, to test the effectiveness of patch reordering on CNNs for large-scale real-world image recognition tasks, we compare our model with AlexNet (Krizhevsky, Sutskever, and Hinton 2012) on ImageNet-2012 dataset. The results demonstrate that patch reordering improves the learning capacity of the model and encodes translation and rotation invariance into the architecture even when trained on raw images only. Finally, to evaluate the generalization ability of the proposed model on other vision tasks with real-world transformations of images, we apply our model to solve the image retrieval task on UK-Bench (Nister and Stewenius 2006) dataset. The improvement in the retrieval performance reveals that the proposed model has a good generalization ability and is better at solving real-world transformation variations.

We implement our method using the open-source Caffe framework (Jia et al. 2014). For patch energy, we have tested both L1 and L2 norm and found that they did not show much difference. Our code and model will be available online. For SI-CNN and ST-CNN, we directly report their results from the original papers on MNIST. For ImageNet-2012, since

Method	R	T
FCN	2.1	2.9
CNN	1.2	1.3
SI-CNN	0.9	-
ST-CNN	<b>0.8</b>	0.8
PR-CNN(ours)	<b>0.8</b>	<b>0.7</b>

Table 1: Classification error on the transformed MNIST dataset. The different distorted MNIST datasets are R (rotation) and T (translation). All models have the same number of parameters and use the same basic architectures.

these two methods did not report their results on this dataset, we forked from the github for re-implementation.

## MNIST

In this section, we use the MNIST handwriting dataset to evaluate all deep models. In particular, different neural networks are trained to classify MNIST data that have been transformed via rotation (R) and translation (T). The rotated dataset was generated from rotating digits with a random angle sampled from a uniform distribution  $U[-90^\circ, 90^\circ]$ . The translated dataset was generated by randomly locating the  $28 \times 28$  digit in a  $42 \times 42$  canvas.

Following (Jaderberg et al. 2015), all networks use ReLU activation function and softmax classifiers. All CNN networks have a  $9 \times 9$  convolutional layer (stride 1, no padding), a  $2 \times 2$  max-pooling layer with stride 2, a subsequent  $7 \times 7$  convolutional layer (stride 1, no padding), and another  $2 \times 2$  max-pooling layer with stride 2 before the final classification layer. All CNN networks have 64 filters per layer. For SI-CNN, convolutional layers are replaced by rotation-invariant layers using six angles from  $-90^\circ$  to  $90^\circ$ . For ST-CNN, the spatial transformer module is placed at the beginning of the network. In our patch reordering CNN, the patch reorder module is applied to the second convolutional layer. The feature maps are divided into  $4^l$  blocks at level  $l$ . Here, we set  $l = 1$ . All networks are trained with SGD for 150000 iterations, with a 256 batch size, 0.01 base learning rate, and no weight decay or dropout. The learning rate was reduced by a factor of 0.1 every 50000 iterations. Weights were initialized randomly, and all networks shared the same random

Method	Ori	R	T
CNN	57.1/80.2	36.6/57.7	46.5/70.8
CNN-Data-Aug	56.6/79.8	36.5/58.3	50.0/73.9
SI-CNN	57.2/80.2	36.8/58.9	-
ST-CNN	59.1/81.7	37.3/59.3	51.4/75.3
PR-CNN(ours)	<b>60.4/82.4</b>	<b>40.7/63.3</b>	<b>54.9/78.0</b>

Table 2: Classification accuracy on the ImageNet-2012 validation dataset. The evaluation datasets are Ori (original), R (rotation), and T (translation).

seed.

The experimental results are summarized in Table 1. It shows that our model achieves better performance under translation and comparable performance under rotation. Because our model does not need any extra learnable parameters, feature extraction modules, or transformations on training images, the comparable performance still reflects the superiority of the patch reordering CNN. For ST-CNN, the best results reported in (Jaderberg et al. 2015) is obtained by training with a more narrow class of transformations selected manually (affine transformations). In our method, we did not optimize with respected to transformation classes. Therefore the comparison is unfair for our PR-CNN. We should compare with the most general ST-CNN defined for a class of projection transformations: 0.8(R) and 0.8(T).

## ImageNet-2012

The ImageNet-2012 dataset consists of images from 1000 classes and is split into three subsets: training (1.3M), validation (50K), and testing (100K images with held-out class labels). The classification performance is evaluated using the top-1 and top-5 accuracy. The former is a multi-class classification accuracy. The latter is the main evaluation criterion used in ILSVRC and is defined as the proportion of images whose ground-truth category is not in the top-5 predicted categories. We use this dataset to test the performance of our model on a large-scale image recognition task.

CNN models are trained on raw images and tested on both raw and transformed images. For all transform types, specific transformations are applied to the original images. Then, the transformed images are rescaled to have a smallest image side of 256 pixels. Finally, the center  $224 \times 224$  crop is used for test. The rotated (R) dataset is generated by randomly rotating original images from  $-45^\circ$  to  $45^\circ$  with a uniform distribution. The translated dataset (T) is generated by randomly shifting an image by a proportion of  $U[-0.2, 0.2]$ .

All the models follow the architecture of AlexNet. For SI-CNN, the first, second and fifth convolutional layers are replaced by rotation-invariant layers using six angles from  $-90^\circ$  to  $90^\circ$ . For ST-CNN, the input is fed into a spatial transformer network before the AlexNet. The spatial transformer network uses bilinear sampling with an affine transformation function. As in (Jaderberg et al. 2015), the size of the spatial transformer network is about half the size of AlexNet. For our PR-CNN, the feature maps are divided into  $4^l$  blocks at level  $l = 2$ , and the patch reorder module is

Method	FC6	FC7
CNN	3.381	3.438
CNN-Data-Aug	3.340	3.441
SI-CNN	3.431	3.452
ST-CNN	3.430	3.446
PR-CNN(ours)	<b>3.574</b>	<b>3.539</b>

Table 3: Performance of CNN models on the UK-Bench retrieval dataset. Here, we use the 4096-dimensional feature of fc6 and fc7 for evaluation.

applied to the fifth convolutional layer.

To train SI-CNN and PR-CNN, we use a base learning rate of 0.01 and decay it by 0.1 every 200,000 iterations. Both networks are trained for 700,000 iterations. We use a momentum of 0.9, a weight decay of 0.0005, and a weight clip of 35. The convolutional kernel weights and bias are initialized by  $\mathcal{N}(0, 0.01^2)$  and 0.1, respectively. The weights and bias of fully connected layers are initialized by  $\mathcal{N}(0, 0.005^2)$  and 0.1. The bias learning rate is set to be  $2 \times$  the learning rate for the weights. For ST-CNN, since it does not converge under the aforementioned setting, we fine-tune the network with the classification network initialized by the pre-trained AlexNet. The spatial transformer module consists of 2 convolutional layers, 2 pooling layers, and 3 fully connected layers. The first convolutional layer filters the input with 48 kernels of size  $11 \times 11 \times 3$  with a stride of 4 pixels, then is connected by a  $3 \times 3$  pooling layer with stride 2. The second convolutional layer has 48 kernels of size  $5 \times 5 \times 48$  with a stride of 2 pixels, followed by a  $3 \times 3$  max pooling layer with stride 2. The output of the pooling layer is fed into two fully connected layers with 48 neurons. Finally, the third fully connected layer maps the output into 6 affine parameters. Then the 6-dimensional output is fed into the spatial transformer layer to get the transformed input image. During the fine-tuning, the learning rate of the spatial transformer is set to be  $0.01 \times$  that of the classification network. We use a base learning rate of 0.001 and decay it by 0.1 every 50,000 iterations, the training process converges after approximately 200,000 iterations.

The results are presented in Table 2. It shows that data augmentation, feature map augmentation, transform preprocessing and patch reordering are all effective ways to improve the rotation or translation invariance of CNNs. Our PR-CNN not only achieves more consistent representation faced with location changes in input but also relieves the models from encoding invariance. It improves the classification accuracy of the model even for the original test images.

## UK-Bench

We also evaluate our PR-CNN model on the popular image retrieval benchmark dataset UK-Bench (Nister and Szwed 2006). This dataset includes 2550 groups of images, each containing 4 relevant samples concerning a certain object or scene from different viewpoints. Each of the in total 10200 images is used as one query to perform image retrieval, targeting at finding each image’s 3 counterparts. We choose UK-Bench since the viewpoint variation in the

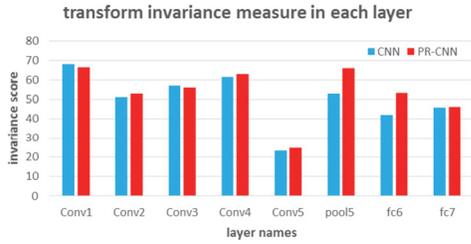


Figure 4: Transform invariance measure (the larger, the better). By applying patch reordering on feature maps during training, the invariance of the following layers is significantly improved.

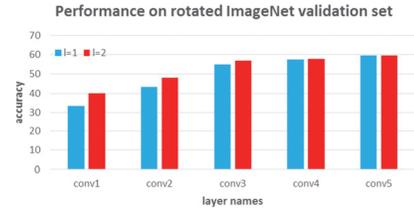
dataset is very common. Although many of the variation types are beyond the three types of geometry transformations that we attempt to address with, we demonstrate the effectiveness of PR-CNN for solving many severe rotation, translation and scale variance cases in image retrieval task.

We directly apply the models trained on ImageNet-2012 for evaluation. The outputs of the fc6 and fc7 layers are used as the feature for each image. Then, we compute the root value of each dimension and perform L2 normalization. To perform image retrieval on UK-Bench, the Euclidean distances of the query image with respect to all 10200 database images are computed and sorted. Images with the smallest distances are returned as top ranked images. NS-Score (average top four accuracy) is used to evaluate the performance, and a score of 4.0 indicates that all the relevant images are successfully retrieved in the top-four results.

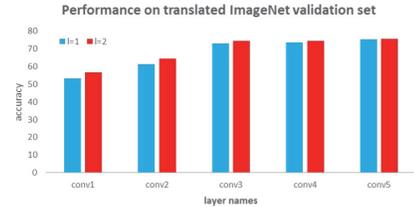
As shown in Table 3, data augmentation, feature map augmentation or spatial transformer network does not present considerable capacity of transform invariance when applied to an unrelated new task. Maybe these models need to be well fine-tuned when transferred to a new dataset and the Spatial Transformer block is content and task dependent. Patch reordering is better for transferring by encoding invariance only into architecture, which is irrelevant to the content of input. It demonstrates that our PR-CNN model can be seamlessly transferred to other image recognition based applications (e.g. image retrieval) without any re-training/fine-tuning. Meanwhile, for other models, fc7 presents better invariance than fc6. However, for our PR-CNN, fc6 is better. We can find some clues from Fig. 2, that is, fc6 presents less parameter redundancy than fc7 in PR-CNN.

### Measuring Invariance

We evaluate the transform invariance achieved by our model using the invariance measure proposed in (Goodfellow et al. 2009). In this approach, a neuron is considered to be firing when its response is above a certain threshold  $t_i$ . Each  $t_i$  is chosen to satisfy the condition that  $G(i) = \sum |h_i(x) > t_i|/N$  is greater than 0.01, where  $N$  is the number of inputs. Then, the *local firing rate*  $L(i)$  is computed as the proportion of transformed inputs to which a neuron fires. To ensure that a neuron is selective and with a high local firing rate (invariance to the set of the transformed inputs), the invariance score of a neuron is computed based on the ratio of its



(a)



(b)

Figure 5: Performance of PR-CNN when the PR module is applied to different layers with different levels. The performance drops significantly when we perform patch reordering in low layers. This is because patch reordering breaks local spatial correlations among patches, which is vital for recognizing meaningful visual patterns in lower layers.

invariance to selectivity, i.e.,  $L(i)/G(i)$ . We report the average of the top 20% highest scoring neurons ( $p = 0.2$ ), as in (Kanazawa and Sharma 2014). Please refer to (Goodfellow et al. 2009) for more details.

Here, we build the transformed dataset by applying rotation ( $[-45^\circ, 45^\circ]$  with a step size of  $9^\circ$ ) and translation ( $[-0.2, 0.2]$  with a step size of 0.04) on the 50,000 validation images of ImageNet-2012. Fig. 4 shows the invariance score of CNN and PR-CNN measured at the end of each layer. We can see that by applying patch reordering to feature maps during training, the invariance of the subsequent layers is significantly improved.

### Effect of Patch Reordering on Different Layers

To investigate the effect of applying patch reordering to different convolutional layers and the effect of pyramid levels, we train different PR-CNN models with patch reordering applied to 1 ~ 5 convolutional layers with level 1 or 2. When  $l = 1$ , we divide the feature maps into  $4 \times 4$  blocks. For  $l = 2$ , the feature maps are first divided into  $2 \times 2$  blocks, and each block is further divided into  $2 \times 2$  sub-blocks. The experimental results are presented in Fig. 5. We can see that the performance drops significantly when we perform patch reordering in low layers. Meanwhile, multi-level reordering does not result in a significant difference to single-level reordering in regard to higher convolutional layers. Low-level features, such as edges and corners, are detected in low layers, and they must be combined in a local spatial range to conduct further recognition. Because patch reordering breaks this local spatial correlation and treats each block as an independent feature, the generated representation becomes less meaningful. This explanation can also clarify the phenomenon that the multi-level division of feature maps

significantly improves model performance in lower layers because a hierarchical reordering will preserve more local spatial relationships than will a single one.

## Conclusion

In this paper, we introduce a very simple and effective way to improve the rotation and translation invariance of CNN models. By reordering the feature maps of CNN layers, the model is relieved from encoding location invariance into its parameters. Meanwhile, CNN models are able to generate more consistent representations when faced with location changes of local patterns in input. Our architecture does not need any extra parameters or pre-processing on input images. Experiments show that our model outperforms CNN models in both image recognition and image retrieval tasks.

**Acknowledgments** This work is supported by NSFC under the contracts No.61572451 and No.61390514, the 973 project under the contract No.2015CB351803, the Youth Innovation Promotion Association CAS CX2100060016, Fok Ying Tung Education Foundation WF2100060004, the Fundamental Research Funds for the Central Universities WK2100060011, Australian Research Council Projects: FT-130101457, DP-140102164, and LE140100061.

## References

- Bruna, J., and Mallat, S. 2013. Invariant scattering convolution networks. *IEEE Trans. Pattern Anal. Mach. Intell.* 35:1872–1886.
- Chatfield, K.; Lempitsky, V.; Vedaldi, A.; and Zisserman, A. 2011. The devil is in the details: an evaluation of recent feature encoding methods. In *Proceedings of the British Machine Vision Conference*, 76.1–76.12.
- Coates, A., and Ng, A. Y. 2011. The importance of encoding versus training with sparse coding and vector quantization. In *ICML*, 921–928.
- Cohen, T. S., and Welling, M. 2015. Transformation properties of learned visual representations. *ICLR*.
- Donahue, J.; Hendricks, L. A.; Guadarrama, S.; Rohrbach, M.; Venugopalan, S.; Saenko, K.; and Darrell, T. 2015. Long-term recurrent convolutional networks for visual recognition and description. *CVPR*.
- Fang, H.; Gupta, S.; Iandola, F.; Srivastava, R.; Deng, L.; Dollar, P.; Gao, J.; He, X.; Mitchell, M.; Platt, J. C.; Zitnick, L.; and Zweig, G. 2015. From captions to visual concepts and back. *CVPR*.
- Gens, R., and Domingos, P. M. 2014. Deep symmetry networks. *NIPS*.
- Gkioxari, G.; Girshick, R.; and Malik, J. 2015. Contextual action recognition with r\*cnn. *ICCV*.
- Goodfellow, I. J.; Le, Q. V.; Saxe, A. M.; Lee, H.; and Ng, A. Y. 2009. Measuring invariances in deep networks. *NIPS*.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. *CVPR*.
- Jaderberg, M.; Simonyan, K.; Zisserman, A.; and Kavukcuoglu, K. 2015. Spatial transformer networks. *NIPS*.
- Jia, Y.; Shelhamer, E.; Donahue, J.; Karayev, S.; Long, J.; Girshick, R.; Guadarrama, S.; and Darrell, T. 2014. Caffe: Convolutional architecture for fast feature embedding. *arXiv preprint arXiv:1408.5093*.
- Kanazawa, A., and Sharma, A. 2014. Locally scale-invariant convolutional neural networks. *NIPS*.
- Karpathy, A., and Li, F.-F. 2015. Deep visual-semantic alignments for generating image description. *CVPR*.
- Kivinen, J. J., and Williams, C. K. I. 2011. Transformation equivariant boltzmann machines. In *Artificial Neural Networks and Machine Learning*, 1–9.
- Krizhevsky, A.; Sutskever, I.; and Hinton, G. E. 2012. Imagenet classification with deep convolutional neural networks. *NIPS* 1097–1105.
- Le, Q. V.; Ngiam, J.; Chen, Z.; hao Chia, D. J.; Koh, P. W.; and Ng, A. Y. 2010. Tiled convolutional neural networks. *NIPS*.
- Lecun, Y.; Bottou, L.; Bengio, Y.; and Haffner, P. 1998. Gradient-based learning applied to document recognition. In *Proceedings of the IEEE*, 2278–2324.
- Lenc, K., and Vedaldi, A. 2015. Understanding image representations by measuring their equivariance and equivalence. *CVPR*.
- Long, J.; Shelhamer, E.; and Darrell, T. 2015. Fully convolutional networks for semantic segmentation. *CVPR*.
- Nister, D., and Stewenius, H. 2006. Scalable recognition with a vocabulary tree. In *CVPR*, 2161–2168.
- Pan, Y.; Mei, T.; Yao, T.; Li, H.; and Rui, Y. 2016. Joint modeling embedding and translation to bridge video and language. *CVPR*.
- Simonyan, K., and Zisserman, A. 2014a. Very deep convolutional networks for large-scale image recognition. *CoRR* abs/1409.1556.
- Simonyan, K., and Zisserman, A. 2014b. Two-stream convolutional networks for action recognition in videos. *NIPS*.
- Sohn, K., and Lee, H. 2012. Learning invariant representations with local transformations. In *ICML*.
- Szegedy, C.; Liu, W.; Jia, Y.; Sermanet, P.; Reed, S.; Anguelov, D.; Erhan, D.; Vanhoucke, V.; and Rabinovich, A. 2015. Going deeper with convolution. *CVPR*.
- Yao, L.; Torabi, A.; Cho, K.; Ballas, N.; Pal, C.; Larochelle, H.; and Courville, A. 2015. Describing videos by exploiting temporal structure. *ICCV*.