Research Article

Convolutional Neural Networks Do Work with Pre-Defined Filters

Christoph Linse¹, Erhardt Barth¹, Thomas Martinetz¹

1. Institute for Neuro- and Bioinformatics, University of Lübeck, Lübeck, Germany

We present a novel class of Convolutional Neural Networks called Pre-defined Filter Convolutional Neural Networks (PFCNNs), where all $n \times n$ convolution kernels with n > 1 are pre-defined and constant during training. It involves a special form of depthwise convolution operation called a Predefined Filter Module (PFM). In the channel-wise convolution part, the $1 \times n \times n$ kernels are drawn from a fixed pool of only a few (16) different pre-defined kernels. In the 1×1 convolution part linear combinations of the pre-defined filter outputs are learned. Despite this harsh restriction, complex and discriminative features are learned. These findings provide a novel perspective on the way how information is processed within deep CNNs. We discuss various properties of PFCNNs and prove their effectiveness using the popular datasets Caltech101, CIFAR10, CUB-200-2011, FGVC-Aircraft, Flowers102, and Stanford Cars. Our implementation of PFCNNs is provided on Github <u>https://github.com/Criscraft/PredefinedFilterNetworks</u>.

1. Introduction

Over the years the computer vision community has been shifting its focus from using pre-defined features towards training end-to-end systems such as Convolutional Neural Networks (CNNs), which have become state-of-the-art in many visual applications for several reasons. CNNs have empirically proven their good generalization abilities^[1], state-of-the-art performance in image recognition^[2], especially in domains of unconstrained image data (taken in the wild)^[3]. The features learned by CNNs can be generic, such that they can be used for a number of different applications^[4].

However, CNNs typically have millions of weights and usually operate in the over-parameterized regime, which is also called the modern interpolating regime^[5]. Pruning experiments show that many of these weights serve no particular function and could be omitted without any loss in performance. It

seems that the superiority of CNNs over traditional image recognition techniques comes with the cost of having an enormous pool of superfluous weights that makes the networks inefficient and less transparent.

In this work, we try to significantly reduce the number of trainable weights in a CNN by applying certain restrictions on the convolutional weights. We apply only a few different pre-defined kernels and do not adjust them during training to save time, computational resources, and energy.

We mold this idea into a Pre-defined Filter Module (PFM), essentially a depthwise convolution layer consisting of a channel-wise $1 \times 3 \times 3$ convolution and a subsequent $C \times 1 \times 1$ convolution that convolves over all input channels C. In the following, we will abbreviate the latter layer with 1×1 convolution. The former layer utilizes a pool of 16 different $1 \times 3 \times 3$ edge filter kernels, each of which is applied to a single input channel individually. As the pre-defined filters are frozen we only adjust the 1×1 convolution weights according to the training data. Learning in this context means finding linear combinations of the pre-defined filter outputs. Training is done end-to-end by finding linear combinations of the outputs of these pre-defined filters using gradient descent.

We construct the novel architecture PFNet18 (Pre-defined Filter Network 18) by replacing all convolution layers with n > 1 of ResNet18 by our PFMs. This eliminates many adjustable weights and PFNet18 requires only 13% of the training parameters of a ResNet18. Please note, that we use the PFMs in the entire network, not only in the first layer. We give empirical evidence that PFNet18 with edge filters can outperform ResNet18 on some fine-grained image datasets. Furthermore, we show that the choice of filters matters, i.e. that edge filters lead to better recognition rates than random filters.

The concept of PFCNNs is fundamentally different from fine-tuning, where network weights are initialized with pre-trained weights. Our pre-defined Filter Convolutional Neural Networks (PFCNNs) apply pre-defined filters for all $n \times n$ convolution kernels with n > 1, which are not changed during training. In our approach, we keep the pre-defined $1 \times n \times n$ convolution kernels and do not alter them during training. There are only a few different kernels (16 in our experiments) that are reused within each layer.

The paper is structured as follows: Section 2 gives a short overview of the related work. Section 3 provides all details about the pre-defined filters and our architecture PFNet18. Section 4 presents the training details and the image classification datasets. The results are presented in Section 5 and subsequently discussed in Section 6. The paper completes with our conclusions in Section 7.

2. Related work

The concept of pre-defined filters has a long history in vision. However, the limitation to fixed, predefined spatial filters seems to be novel in the intermediate layers of neural networks. In 1980 Fukushima^[6] used pre-defined filters in self-organizing neural networks. In today's era of deep networks, pre-defined filters are sometimes used as a pre-processing step to boost recognition performance. Ma et al.^[7] used pre-defined filters to incorporate domain knowledge into their training. They, however, replaced only the first convolutional layer of CNNs with trainable traditional image filters (Gabor filters).

Gavrikov and Keuper^[8] showed that learning linear combinations of fixed, random filters is a successful strategy to solve image classification problems, especially with wide CNNs (many channels). They also showed that the random filters can be shared across layers to further reduce the number of weights. In our study, we also learn linear combinations of filter outputs. However, we reduce the number of random filters to only 16 hand-picked kernels, that are applied in each depth-wise convolution operation. We show that the choice of these kernels matters, i.e. that edge filters outperform random kernels.

3. Pre-defined Filter Networks

3.1. Pre-defined Filter Module

The Pre-defined Filter Module (PFM) is illustrated in Figure 1. It is structured like a depthwise convolutional layer. In the first part of the module, each input channel is convolved independently with exactly one $1 \times 3 \times 3$ filter kernel from a pool of k pre-defined kernels. In our experiments, we use k = 16 different edge kernels. The order in which the kernels are distributed over the input channels is fixed. The input channel with index i is convolved with the kernel with index $i \mod k$.

The width of the intermediate part of the module can be chosen with the parameter f, as can be seen in Figure 1. Conceptually, f determines the number of copies of the n_{in} input channels prior to convolution. The number of intermediate channels after applying the pre-defined filter kernels is $n_{int} = n_{in} \cdot f$. Our implementation requires $n_{int} \mod k = 0$ with the number of pre-defined kernels k. Another requirement is $n_{int} \ge n_{in}$.

In the second part of the module, an $n_{int} \times 1 \times 1$ convolution creates an arbitrary number of linear combinations of the channels. By default, the first part of the PFM is frozen during training and only the weights in the 1×1 convolution are adjusted to the training set.



Figure 1. Pre-defined Filter Module (PFM). $1 \times n \times n$ kernels are taken from a small pool of kernels and are applied channel-wise as known from depthwise convolution. The order in which the kernels are distributed over the input channels is fixed.

3.2. Choice of the parameter f

The order in how the pre-defined filters are applied to the input channels should have no effect on the set of functions F that can be learned by the network. The same should apply to a permutation of the channels in the input image, i.e., the order of kernels should not be part of the architecture. This has implications for the choice of the parameter f.

First, we discuss the first PFM that is applied to the RGB input image. For f = k, each kernel is applied once to each RGB input channel. In this case, the order of the input channels or the pre-defined filters is irrelevant. If we permute the input channels or the pre-defined filters we can always permute the weights in the consecutive 1×1 layer to get the same result as before.

Second, we discuss the second and following PFMs. Here, we choose f = 1. Again, the order of the pre-defined filters has no effect on the set of functions F that can be learned by the network. If two input kernels at positions A and B are swapped, this operation can be undone in two steps. First, the preceding 1×1 layer swaps its output channels A and B by permuting its weights. If skip connections exist, all preceding PFMs have to adjust their 1×1 layer. Second, one needs to swap the inputs of the consecutive 1×1 layer, which can also be done by permuting the corresponding weights.

The order of kernels would matter for $f \in [2, k - 1]$, which we avoid. This can easily be seen with some combinatorics. Please remember, that the input channel with index *i* is convolved with the kernel $i \mod k$. As shown in Figure 1, the input channels are copied prior to convolution. Thus, each input channel receives a set of f different pre-defined kernels. However, if the order would be arbitrary we had $\binom{k}{f}$ filter combinations that could occur for a single input channel. This is larger than k for $f \in [2, k - 1]$. Therefore, if two pre-defined filters are swapped this could introduce new filter combinations, and the set of functions F could change.

3.3. PFNet18 architecture

The PFNet18 architecture is described in Table 1. For comparison, ResNet18 architecture is presented in Table 2. Starting from ResNet18, we replace the convolution layers with Pre-defined Filter Modules. The first module has f = 16, which is also the number of pre-defined kernels. Thus, the first layer has $n_{int} = 3 \cdot 16 = 48$ intermediate channels. To be consistent with ResNet18, we choose $n_{out} = 64$ output channels for the 1×1 convolution. After the first PFM, we stack 8 residually connected Basic Blocks and replace the convolution operations with PFMs with f = 1. The residual connections are kept. After the last block, there is an adaptive average pooling layer and a fully connected layer just like in ResNet18. PFNet18 has 1.46 million parameters, which is only 13% of the 11.23 million parameters of the ResNet18.

Layers	Output size	Kernel	
Pre-defined Filter Module	$\begin{bmatrix} 3 \times 224 \times 224 \\ 48 \times 112 \times 112 \\ 64 \times 112 \times 112 \end{bmatrix}$	$egin{bmatrix} ext{input,} \ 1 imes3 imes3, ext{depthwise} \ 48 imes1 imes1, ext{pixelwise} \end{bmatrix}$	
MaxPooling	64 imes 56 imes 56	3 imes 3, 64, stride 2	
Double HF Residual Module	64 imes56 imes56	$egin{bmatrix} 1 imes3 imes3, ext{depthwise}\ 64 imes1 imes1, ext{pixelwise} \end{bmatrix} imes2 \end{cases}$	
Double HF Residual Module	128 imes28 imes28	$egin{bmatrix} 1 imes3 imes3, ext{depthwise}\ 64 imes1 imes1, ext{pixelwise} \end{bmatrix}$	
Double HF Residual Module	128 imes28 imes28	$egin{bmatrix} 1 imes 3 imes 3, ext{depthwise} \ 128 imes 1 imes 1, ext{pixelwise} \end{bmatrix}$	
Double HF Residual Module	256 imes14 imes14	$egin{bmatrix} 1 imes3 imes3, ext{depthwise}\ 128 imes1 imes1, ext{pixelwise} \end{bmatrix}$	
Double HF Residual Module	256 imes14 imes14	$egin{bmatrix} 1 imes 3 imes 3, ext{depthwise} \ 256 imes 1 imes 1, ext{pixelwise} \end{bmatrix}$	
Double HF Residual Module	512 imes7 imes7	$egin{bmatrix} 1 imes 3 imes 3, ext{depthwise} \ 256 imes 1 imes 1, ext{pixelwise} \end{bmatrix}$	
Double HF Residual Module	512 imes7 imes7	$egin{bmatrix} 1 imes 3 imes 3, ext{depthwise}\ 512 imes 1 imes 1, ext{pixelwise} \end{bmatrix}$	
Classification layer	512 imes1 imes1	Adaptive average pool	
		fully connected, softmax	

 Table 1. Details of the PFNet18 architecture.

Layers	Output size	Kernel	
Convolution	64 imes112 imes112	7 imes7, 64, stride 2	
MaxPooling	64 imes 56 imes 56	3 imes 3, 64, stride 2	
Basic Block	64 imes 56 imes 56	$\begin{bmatrix} 64\times3\times3\\ 64\times3\times3 \end{bmatrix}\times2$	
Basic Block	128 imes28 imes28	$\begin{bmatrix} 64\times 3\times 3\\ 128\times 3\times 3\end{bmatrix}$	
Basic Block	128 imes28 imes28	$egin{bmatrix} 128 imes3 imes3\ 128 imes3 imes3\end{bmatrix}$	
Basic Block	256 imes14 imes14	$egin{bmatrix} 128 imes 3 imes 3 \ 256 imes 3 imes 3 \end{bmatrix}$	
Basic Block	256 imes14 imes14	$\begin{bmatrix} 256\times3\times3\\ 256\times3\times3 \end{bmatrix}$	
Basic Block	512 imes7 imes7	$\begin{bmatrix} 256\times3\times3\\512\times3\times3 \end{bmatrix}$	
Basic Block	512 imes7 imes7	$egin{bmatrix} 512 imes3 imes3\ 512 imes3 imes3 \end{bmatrix}$	
Classification layer	512 imes1 imes1	Adaptive average pool	
		fully connected, softmax	

Table 2. Details of the ResNet18 architecture.

3.4. Choice of pre-defined filters

In our experiments, we employ only 16 different pre-defined $1 \times 3 \times 3$ filters. We choose edge filters because they provide gradient information and are often found in rudimentary forms in trained CNNs^{[9][10]}. Figure 2 presents 8 uneven and 8 even edge filters in different orientations including horizontal, vertical, and diagonal. We suppose that employing edge filters in PFCNNs will lead to an appropriate bias that will help the network to learn robust features. The kernel elements w_i are normalized such that $\sum_i w_i = 0$ and $\sum_i |w_i| = 1$. The constant component is missing in the convolution kernels, which should bias the network toward the processing of edges and shapes. As there are only 8 linearly independent kernels (some of the kernels have a flipped sign) the kernels span an 8-dimensional space.



4. Experimental setup

We train and test PFNet18 on several image classification datasets and compare its performance with ResNet18. The following sections present the benchmark datasets and all training details.

4.1. Datasets

Example images for each dataset are presented in Table 3. The Caltech101 dataset^[11] has a total of 8677 images, which were collected from the internet. They are scaled to be about 300 pixels wide. The dataset contains 101 categories, as well as a background class, which is ignored in the experiments. The categories are quite diverse and reach from animals and plants to electronic products and vehicles. The objects are shown in cluttered environments or in front of realistic or white backgrounds. As no official split is available, we randomly pick 20 training images and 10 test images per class.

The Caltech-UCSD Birds-200-2011 dataset (CUB-200-2011)^[12] originated from the Caltech-UCSD Birds 200 dataset^[13] from 2010. We use the official split with 5994 training and 5794 test images. The average width of the images is about 470 pixels. There are 200 challenging bird categories (classes), which have plenty of inter-class variation due to deformations, plumage color, lighting, perspective, and pose.

The Fine-Grained Visual Classification of Aircraft dataset (FGVC-Aircraft)^[14,] provides 6667 images for training and 3333 images for testing. It contains large images with an average width of 1100 pixels. We use the 100 categories of airplane models. The airplanes are rigid and do not deform. However, the same airplane model can appear quite different depending on the advertisement, airlines, perspective, and cropping.

The images from the 102 Category Flower dataset (Flowers102)^[15] show one or several blossoms of 102 flower types. As the other datasets in our study have training and test sets only, we merge the official training and validation set of Flowers102 to our training set and we test on the official test set. Accordingly, the training set contains 2040 images, and the test set 6149 images. The images have a mean width of 630 pixels. While some flower types can have subtle differences in their appearance, there are large intra-class variations due to scaling, perspective, lighting, and color variants. We only use categorical information within the dataset.

The StanfordCars dataset^[16] from Stanford University provides 8144 training and 8041 test images. The average width of an image is 700 pixels. The dataset covers 196 different car models. Usually, one car is shown per image. The cars can be on the road or indoors, shown from the front, side, or back, introducing many variations to the dataset. CIFAR10^[17] contains 50000 training images and 10000 test images of the size 32×32 pixels. The dataset provides the 10 classes plane, car, bird, cat, deer, dog, frog, horse, ship, and truck. For our experiments on the CIFAR10 dataset, we adjusted the architectures to be compatible with small image shapes by removing the max pool layer and having a stride of 1 in the first convolutional layer and a kernel size of 3. These adjustments apply for both, PFNet18 and ResNet18.



Table 3. Images from the benchmark datasets showing each two instances of four classes.

4.2. Training details

The network weights are initialized using Kaiming normal initialization^[18]. We apply the Lamb optimizer^[19] to minimize the cross-entropy loss of our models. The batch size is 64. We train for 300 epochs. The initial learning rate of $lr_{init} = 0.003$ is step-wise reduced to $lr_{init} \cdot 10^{-2}$. The weight decay is 1. We use the Pytorch framework^[20] to implement the models and to perform the training on an NVidia GTX 1080 Ti. The experiments are repeated 5 times with random seeds. Each seed affects the weight initialization of the networks, the mini-batch aggregation, and random effects during data augmentation. For image augmentation, we apply random cropping and random horizontal flipping.

5. Results

5.1. Benchmarks

The average test performance of the PFNet18 and ResNet18 models is presented in Table 4. PFNet18 outperforms ResNet18 on the Caltech101, FGVC-Aircraft, and Flowers102 datasets. On Caltech101, PFNet18 has an almost 10% higher test accuracy than ResNet18 and on Flowers102 the improvement is 7%. On the Stanford Cars dataset, both architectures perform very similarly. ResNet18 achieves higher accuracies on the CIFAR10 and the CUB-200-2011 datasets. The experiments show that the restriction to employ only 16 different pre-defined $1 \times 3 \times 3$ filters is sufficient to learn complex relationships in image data. The PFNet18 models reached high accuracy by simply finding linear combinations of pre-defined filter outputs. This is intriguing because it demonstrates that the spatial kernels do not have to be learned at all in many cases.

Dataset	PFNet18	ResNet18	
Caltech101	65.60±0.66	57.19±0.78	
CIFAR10	92.15±0.18	94.33±0.11	
CUB-200-2011	53.00±0.55	58.51±0.53	
FGVC-Aircraft	75.49±0.29	73.32±1.06	
Flowers102	80.66±0.35	73.40±0.34	
Stanford Cars	77.06±0.42	77.90±0.37	

Table 4. Test accuracy on benchmark datasets for training from scratch and 5 different seeds.

5.2. Feature visualization

The promising results indicate that our choice of pre-defined filters introduces a good bias for at least half of the datasets. In an ablation study in Section 5.5 we study the choice of filters quantitatively. In this section, we study the learned filters of PFNet18 qualitatively using the feature visualization technique^[21]. Feature visualization reveals specific input patterns that maximize the activation of specific network units. The outputs of this technique are input images that show characteristics of the features processed within the network. The idea is to initialize an input image with Gaussian noise and to modify the image such that it maximizes the activation of a specific neuron using gradient ascent. Hence, this procedure is also called activation maximization. Let θ be the network weights and let $f_{ij}(\theta, x)$ be the jth feature map in layer i given θ and some input image x: for activation maximization, the input image x^* with

$$x^* = rg \max_{x \in \mathbb{R}^{C imes H imes W}} f_{ij}(heta, x)$$

is determined^[21]. We perform the optimization using gradient ascent on a fixed θ . Such generated images often look unnatural and regularization approaches have been used in the literature to improve the visual quality^[22]. Between the update steps $x \leftarrow x + \mu \nabla_x f_{ij}(\theta, x)$ with learning rate μ we apply random image transformations to support the gradient ascent to find robust maxima. Our transformations include random rotation, scaling, blurring, cropping, pixel rolling, and shifting/scaling the image tensor distribution toward the normal distribution. Similar transformations have been applied in the works^{[23][24][25]}.

Table 5 shows feature visualizations of PFNet18 and ResNet18 trained on the Caltech101 and Flowers102 datasets. Table 5 is best viewed digitally with zoom. The images were picked by hand to illustrate the variation and the characteristics of the features in different layers.

The results indicate that both PFNet18 and ResNet18 are able to learn complex visual features. One observes that the features processed at the end of the first block consist of simple, repetitive textures. At the end of the second block, more details are added to these textures and it is possible to see a difference between the datasets. The models that were trained on the Flowers102 dataset produce leaf and flower prototypes. The models that were trained on Caltech101 also include rectangular shapes and more variety in general. Similar observations apply for the third and fourth blocks, where the first object instances appear. The feature visualization of the classification layer reveals what the networks expect to look like cougar face, face, motorbike, and saxophone (Caltech101) and Mexican aster, primula, tree mallow, and water lily (Flowers102).

Feature visualization reveals that both PFNet18 and ResNet18 learned object-specific features. On the Caltech101 and Flowers102 datasets, the PFNets18's feature visualizations look more convincing. On Caltech101, exactly one instance of the target class is generated with a plausible shape. The cougar's

face consists of a head with two eyes and ears. The human face has a nose, two eyes (the second one is difficult to see), hair, and a flat chin. The motorbike has two wheels, a saddle, and a handle. The saxophone is a long stick with buttons and a shimmering surface. ResNet18 does not generate such clear shapes and seems to focus more on textures. Among many feature visualizations and also the other classes, no examples could be found that look as convincing as those of PFNet18. The results show that PFCNNs are able to learn complex, object-specific features from only 20 training images per class.

Layer	Caltech101, PFNet18	Caltech101, ResNet18	Flowers102, PFNet18	Flowers102, ResNet18
Block 1				
Block 2				
Block 3				
Block 4				
Classes				
Class reference				

Table 5. Feature visualization for models trained on the Caltech101 and Flowers102 dataset. Each image is an input that maximizes the activation of a specific channel in a specific layer. The visualized classes are cougar face (top left), face (top right), motorbike (bottom left), and saxophone (bottom right) for

the Caltech101 dataset and Mexican aster (top left), primula (top right), tree mallow (bottom left) and water lily (bottom right) for the Flowers102 dataset. Best viewed with zoom.

5.3. Computational efficiency

Computational efficiency and model size are important factors for training and deploying deep networks, especially on limited hardware with harsh energy and memory requirements. Table 6 summarizes relevant computational aspects of the models. PFNet18 has only 1.46 million parameters and requires only 6 MB space on the disk, whereas ResNet18 needs 45 MB and has 11.23 million parameters. The reduction of parameters does not lead to a significant change in training time, as the duration of a backward pass is around 65-70 ms for a batch size of $64 \times 3 \times 224 \times 224$ on an NVidia GTX 1080 Ti. PFNet18 needs an additional time of 10 ms per batch. This is interesting because the number of mult-adds of PFNet18 is only $0.26 \cdot 10^9$ while ResNet18 has $1.81 \cdot 10^9$ mult-adds. Counting Mult-Add operations means counting the FLOPs of a model and to divide the result by 2. Although PFNet18 requires much fewer computations in total, its implementation requires more nodes in the computational graph and more distinct GPU calls. Our implementation seems to be not very efficient on current GPU hardware, which is optimized for a few but large tensor operations. Therefore, we assume that there is still much room for optimizations in the implementation of PFMs.

Model	Parameters (10 ⁶)	Size (MB)	FP (ms)	BP (ms)	GPU Memory FP (GB)	Mult-Adds (10 ⁹)
PFNet18	1.46	6	37	70	3.9	0.26
ResNet18	11.23	45	27	65	3.0	1.81

Table 6. Computational efficiency of PFNet18 and ResNet18 considering model size and speed. FP denotes forward pass and BP denotes backward pass. During BP the hyperparameters described above are used. The input tensors have the shape $64 \times 3 \times 224 \times 224$.

5.4. Aliasing effects

Aliasing may occur in CNNs due to spatial sub-sampling. In PFNet18 and ResNet18 there are 3 skip connections with a 1×1 convolution and a stride of 2 where information is lost. The aliasing effects do not always seem to affect the classification performance. To test this, we blur the input data of these 3 convolution operations with a 3×3 Gaussian filter and train the models again. For ResNet18, there is no significant change in the test accuracy on the Flowers102 dataset. For PFNet18, however, there was a great improvement in test accuracy from 72.40 ± 0.59 to 80.66 ± 0.35 .

The aliasing issue is studied in Figure 3 where different convolution operations with stride 2 are applied to an input image. The input image contains two white lines on a black background. The 1×1 convolution does only capture one of the lines because of spatial sampling. With stride 2, our $1 \times 3 \times 3$ pre-defined filters do not always capture both lines, either. This is a problem for PFNet18 because each input channel is convolved with exactly one $1 \times 3 \times 3$ pre-defined filter. The architecture lacks other convolution steps on the same input channel that could add redundancy. This means that information can be irrevocably lost. Thus, PFNet18 relies on the skip connections to compensate for aliasing effects and additional aliasing within the skip connections must be avoided.

Our results suggest that ResNet18 is more robust against aliasing in the skip connections. In Figure 3 we see that random 3×3 filters often capture both lines, which illustrates that ResNet18 has a lower risk of losing image information by aliasing. In addition, each convolution layer from ResNet18 convolves each input channel as many times as there are output channels. This gives more opportunities to propagate information from each input channel to the next layer.



Figure 3. Aliasing effects. Various stride 2 convolutional operations are applied on the input image showing two horizontal lines. Orange color denotes positive kernel weights, black denotes zero, and blue negative weights.

5.5. Ablation study

An ablation study is conducted to study the relevance of single elements in the PFM as well as the importance of the choice of pre-defined filters. Table 7 presents the results of the ablation study. The experiments are conducted on the Flowers102 dataset with the same hyperparameters as in the previous experiments. Similar to Table 4 the results show average values for five different seeds.

First, we show that the choice of pre-defined kernels matters. We pick 16 random pre-defined kernels from a uniform distribution. These kernels are frozen during training. Only the 1×1 convolutions are adjusted to the Flowers102 dataset. In this setting, the performance drops to $73.98 \pm 1.24\%$. However, interestingly, it is still as good as ResNet18.

In another experiment, the random pre-defined kernels are unfrozen, such that the model can determine the 16 kernels during training. Note that we give each PFM its own 16 kernels to optimize. This model gave a test accuracy of $74.54 \pm 0.88\%$, which is a small improvement. When the pre-defined filters are initialized with the edge filters and allowed to be optimized during training, the test accuracy is 80.11 ± 0.62 , which is similar to the default experiment. Our findings indicate that the proposed set of filter kernels provides a beneficial bias for the Flowers102 dataset.

Second, we study the importance of the first ReLU in the PFMs. The first ReLU is removed from all PFMs. This means that the convolution with the pre-defined kernels and the subsequent 1×1 convolution form one linear operation. The resulting filter will have a kernel that is differently organized than the convolution kernels in ResNet18. The performance drops to 75.21 ± 0.23 . The ReLU seems to be important, possibly due to the increased amount of non-linearity.

Model	Description	Accuracy
PFNet18	No aliasing, default	80.66±0.35
PFNet18	Aliasing	72.40±0.59
PFNet18	First ReLU removed, no aliasing	75.21±0.23
PFNet18	16 Trainable filters, edge init., no aliasing	80.11±0.62
PFNet18	16 Trainable filters, random init., no aliasing	74.54±0.88
PFNet18	16 frozen filters, random init., no aliasing	73.98±1.24
ResNet18	Default	73.40±0.34
ResNet18	No aliasing	74.36 ±0.57

Table 7. Test accuracy of variants of PFNet18 and ResNet18 on the Flowers102 dataset for training fromscratch.

6. Discussion

Pre-defined filter kernels lead to significant performance improvements on half of the datasets that we evaluated. This improvement is achieved with only 16 different spatial $1 \times 3 \times 3$ filter kernels and only 13% of the trainable weights of ResNet18. Despite the restriction of PFCNNs to only learn combinations of pre-defined filter outputs, discriminative features emerge during training. For PFCNNs, the ability to perform visual recognition is based on the appropriate combination of filter outputs. These findings provide a new perspective on the information processing within deep CNNs and they show once again, that many weights in conventional CNNs are redundant.

We found that aliasing can significantly reduce the performance of PFCNNs. Since this effect is much lower in CNNs, one can assume that the CNNs learn how to deal with aliasing; which, however, implies that resources need to be spent to deal with the problem.

Feature visualization shows that the combinations of the pre-defined filter outputs yield complex, object-specific features. Note that these features emerged from only 20 training images per class. Compared to ResNet18, our PFCNNs seem to also consider the shape of the recognized objects, while ResNet18 seems to focus on textures.

We discovered that the choice of pre-defined filters matters. When using random kernels instead of edge filters, the test accuracy on the Flowers102 dataset drops to $73.98 \pm 1.24\%$, which, however, is still as good as ResNet18, which is remarkable. We conclude that edge filters add a suitable bias to the image recognition problem.

7. Conclusion

We introduced a novel class of CNNs called Pre-defined Filter Convolutional Neural Networks (PFCNNs), which utilizes fixed pre-defined filters in all $n \times n$ convolution kernels with n > 1 while keeping the end-to-end training paradigm. In our implementation, the PFNet18 architecture is a ResNet18 where we replaced the convolution operations with so-called Pre-defined Filter Modules (PFMs). These modules consist of a depthwise convolution with pre-defined weights that are not changed during training; only the weights of the subsequent 1×1 convolution are adjusted. PFNet18 has only 13% of the weights of ResNet18 but outperforms ResNet18 on the Caltech101, FGVC-Aircraft, and Flowers102 datasets with an absolute increase of 10%. On the CIFAR10, CUB-200-2011, and Stanford Cars datasets, PFNet18 does not reach higher performances than ResNet18 but still performs

well with much fewer parameters. The results demonstrate that our choice of taking edge filters as pre-defined filter kernels is a useful bias for image data. Interestingly, the pre-defined edge filters are useful biases not only in the first but also in the higher layers of the CNN. Our results imply that it is unnecessary to train the spatial kernels of a CNN to reach reasonable test accuracies on image data, which saves most trainable weights. In contrast to pruning, where weights are eliminated, we save weights by excessive weight sharing using a small pool of 16 pre-defined kernels. Our approach uses only 13% of the parameters of ResNet18, which may be useful for mobile devices and other applications where model size is critical.

Many questions regarding PFCNNs arise, e.g., the reduction of weights by the PFMs could be combined with more recent, efficient architectures to get very light models with much fewer parameters. These models are in stark contrast to the usual over-parameterized approaches but still achieve reasonable results. We hope that PFCNNs will lead to interesting comparative studies and a better understanding of the way information is processed internally by CNNs. It is left to future research to explore how PFCNNs perform on large-scale datasets such as ImageNet^[261], and how PFCNNs perform in transfer-learning scenarios. Another question involves how the choice (and number) of filters affect the performance of PFCNNs as well as the robustness against input perturbations and adversarial attacks. In addition, the width of the networks could be increased to allow for more filters to be linearly combined in each layer. This might boost the performance of the PFCNNs according to Gavrikov and Keuper^[8] who found that increasing the width of networks with fixed, random spatial kernels leads to an improvement of the test accuracy on CIFAR10. Also, the question remains to what extent the results can be transferred to other domains, for instance, audio processing. Overall it seems that the power of deep networks mainly lies in their ability to learn how to combine filters, rather than in the learning of spatial convolution kernels.

Acknowledgments

The work of Christoph Linse was supported by the Bundesministerium für Wirtschaft und Klimaschutz through the Mittelstand-Digital Zentrum Schleswig-Holstein Project. The Version of Record of this contribution was presented at the 2023 International Joint Conference on Neural Networks (IJCNN), and is available online at <u>https://doi.org/10.1109/IJCNN54540.2023.10191449</u>

References

- [^]Linse C, Martinetz T (2022). "Large Neural Networks Learning from Scratch with Very Few Data and without Regularization." arXiv. doi:<u>10.48550/ARXIV.2205.08836</u>.
- [^]He K, Zhang X, Ren S, Sun J. "Deep Residual Learning for Image Recognition." In: 2016 IEEE Conferenc e on Computer Vision and Pattern Recognition (CVPR). Las Vegas, NV, USA: IEEE; 2016. p. 770-778. doi: <u>10.1109/CVPR.2016.90</u>.
- 3. [^]Alshazly H, Linse C, Barth E, Martinetz T (2019). "Handcrafted versus CNN Features for Ear Recognitio n". Symmetry. 11 (12): 1493. doi:<u>10.3390/sym1121493</u>.
- 4. [^]Hertel L, Barth E, Käster T, Martinetz T (2017). "Deep Convolutional Neural Networks as Generic Featu re Extractors." arXiv:1710.02286 [cs]. arXiv: <u>1710.02286</u>.
- 5. [△]Belkin M, Hsu D, Ma S, Mandal S (2019). "Reconciling modern machine-learning practice and the clas sical bias-variance trade-off". Proceedings of the National Academy of Sciences. 116 (32): 15849–1585
 4. doi:10.1073/pnas.1903070116.
- ^AFukushima K. "Neocognitron: A self-organizing neural network model for a mechanism of pattern rec ognition unaffected by shift in position." Biological Cybernetics. 36(4): 193--202. doi:<u>10.1007/BF00344</u>.
 <u>251</u>. Apr. 1980.
- 7. [^]Ma Y, Luo Y, Yang Z (2020). "PCFNet: Deep neural network with predefined convolutional filters". Neu rocomputing. 382: 32--39. doi:<u>10.1016/j.neucom.2019.11.075</u>.
- 8. ^a, ^bGavrikov P, Keuper J. "Rethinking 1x1 Convolutions: Can we train CNNs with Frozen Random Filter s?" arXiv. Jan 2023. Available from: <u>arXiv:2301.11360</u>.
- 9. [△]Krizhevsky A, Sutskever I, Hinton GE. "ImageNet Classification with Deep Convolutional Neural Netwo rks." In: Pereira F, Burges CJ, Bottou L, Weinberger KQ, editors. Advances in Neural Information Processi ng Systems. Vol. 25. Curran Associates, Inc.; 2012.
- [^]Gavrikov P, Keuper J. "CNN Filter DB: An Empirical Investigation of Trained Convolutional Filters." In:
 2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). New Orleans, LA, USA:
 IEEE; 2022. p. 19044-19054. doi:<u>10.1109/CVPR52688.2022.01848</u>.
- 11. ^ALi F-F, Fergus R, Perona P. "Learning Generative Visual Models from Few Training Examples: An Incre mental Bayesian Approach Tested on 101 Object Categories." In: 2004 Conference on Computer Vision a nd Pattern Recognition Workshop. Washington, DC, USA: IEEE; 2004. p. 178–178. doi:10.1109/CVPR.200 <u>4.383</u>.

- 12. [^]Wah C, Branson S, Welinder P, Perona P, Belongie S. "The Caltech-UCSD Birds-200-2011 Dataset." Cal ifornia Institute of Technology; 2011. Report No.: CNS-TR-2011-001.
- 13. [^]Welinder P, Branson S, Mita T, Wah C, Schroff F, Belongie S, Perona P. Caltech-UCSD Birds 200. Califor nia Institute of Technology; 2010. Report No.: CNS-TR-2010-001.
- 14. [^]Maji S, Rahtu E, Kannala J, Blaschko M, Vedaldi A. "Fine-Grained Visual Classification of Aircraft." arX iv:1306.5151 [cs]. 2013 Jun. arXiv: <u>1306.5151</u>.
- ^ANilsback M-E, Zisserman A. "Automated Flower Classification over a Large Number of Classes." In: 20 o8 Sixth Indian Conference on Computer Vision, Graphics & Image Processing. Bhubaneswar, India: IEE E; 2008. p. 722-729. doi:<u>10.1109/ICVGIP.2008.47</u>.
- 16. [△]Krause J, Stark M, Deng J, Fei-Fei L. "3D Object Representations for Fine-Grained Categorization." In:
 2013 IEEE International Conference on Computer Vision Workshops. Sydney, Australia: IEEE; 2013. p. 55
 4-561. doi:10.1109/ICCVW.2013.77.
- 17. ^AKrizhevsky A. (2009). "Learning multiple layers of features from tiny images."
- [^]He K, Zhang X, Ren S, Sun J. Delving Deep into Rectifiers: Surpassing Human-Level Performance on Im ageNet Classification. In: 2015 IEEE International Conference on Computer Vision (ICCV). Santiago, Chil e: IEEE; 2015. p. 1026-1034. doi:<u>10.1109/ICCV.2015.123</u>.
- 19. [△]You Y, Li J, Reddi S, Hseu J, Kumar S, Bhojanapalli S, Song X, Demmel J, Hsieh C-J (2020). "Large batc h optimization for deep learning: Training BERT in 76 minutes." In: International Conference on Learni ng Representations.
- 20. [△]Paszke A, Gross S, Massa F, Lerer A, Bradbury J, Chanan G, Killeen T, Lin Z, Gimelshein N, Antiga L, Des maison A, Kopf A, Yang E, DeVito Z, Raison M, Tejani A, Chilamkurthy S, Steiner B, Fang L, Bai J, Chintal a S. "PyTorch: An Imperative Style, High-Performance Deep Learning Library." In: Wallach H, Larochell e H, Beygelzimer A, d'Alché-Buc F, Fox E, Garnett R, editors. Advances in Neural Information Processing Systems 32. Curran Associates, Inc.; 2019. p. 8024–8035.
- 21. ^a. ^bErhan D, Bengio Y, Courville A, Vincent P (2009). "Visualizing higher-layer features of a deep netwo rk". University of Montreal. **1341** (3): 1–13.
- 22. [△]Mahendran A, Vedaldi A. "Understanding deep image representations by inverting them." In: 2015 IEE
 E Conference on Computer Vision and Pattern Recognition (CVPR). Boston, MA, USA: IEEE; 2015. p. 5188
 –5196. doi:10.1109/CVPR.2015.7299155.
- 23. [^]Yosinski J, Clune J, Nguyen A, Fuchs T, Lipson H. "Understanding Neural Networks Through Deep Visua lization." arXiv. 2015. Number: arXiv:1506.06579. Available from: <u>arXiv:1506.06579</u>.

- 24. ^AMordvintsev A, Olah C, Tyka M (2015). "Inceptionism: Going deeper into neural networks."
- 25. [^]Nguyen A, Yosinski J, Clune J. "Deep neural networks are easily fooled: High confidence predictions for unrecognizable images." In: 2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR). Boston, MA, USA: IEEE; 2015. p. 427–436. doi:<u>10.1109/CVPR.2015.7298640</u>.
- 26. [△]Russakovsky O, Deng J, Su H, Krause J, Satheesh S, Ma S, Huang Z, Karpathy A, Khosla A, Bernstein M, Berg AC, Fei-Fei L (2015). "ImageNet Large Scale Visual Recognition Challenge". International Journal of Computer Vision. 115 (3): 211–252. doi:<u>10.1007/s11263-015-0816-y</u>.

Declarations

Funding: The work of Christoph Linse was supported by the Bundesministerium für Wirtschaft und Klimaschutz through the Mittelstand-Digital Zentrum Schleswig-Holstein Project.
Potential competing interests: No potential competing interests to declare.