arXiv:2009.06385v1 [cs.CV] 14 Sep 2020

Adaptive Convolution Kernel for Artificial Neural Networks

F. Boray Tek, İlker Çam, Deniz Karlı

September 15, 2020

Abstract

Many deep neural networks are built by using stacked convolutional layers of fixed and single size (often 3×3) kernels. This paper describes a method for training the size of convolutional kernels to provide varying size kernels in a single layer. The method utilizes a differentiable, and therefore backpropagation-trainable Gaussian envelope which can grow or shrink in a base grid. Our experiments compared the proposed adaptive layers to ordinary convolution layers in a simple two-layer network, a deeper residual network, and a U-Net architecture. The results in the popular image classification datasets such as MNIST, MNIST-CLUTTERED, CIFAR-10, Fashion, and "Faces in the Wild" showed that the adaptive kernels can provide statistically significant improvements on ordinary convolution kernels. A segmentation experiment in the Oxford-Pets dataset demonstrated that replacing a single ordinary convolution layer in a U-shaped network with a single 7×7 adaptive layer can improve its learning performance and ability to generalize.

1 Introduction

Neural network-based pattern recognition is the state-of-the-art approach to solving many visual problems. The most successful solutions are based on stacked convolutional layers Krizhevsky et al. (2017); Iandola et al. (2016); He et al. (2015). The stacked deep hierarchy allows increasingly complex and discriminative representations (features) which also become easier to classify. Though biological neurons are functionally different, there is firm evidence that biological neurons in the visual cortex perform in a similar way to neurons in convolutional layers Poggio and Serre (2013). In the late 1960s, Hubel and Wiesel Hubel (1962) discovered three types of cells in the visual cortex: simple, complex, and hyper-complex (i.e. end-stopped cells). The simple cells are sensitive to the orientation of the excitatory input, whereas the hyper-complex cells are activated by particular types of orientation, motion, and *size* of the stimuli.

The common convolutional layer in a neural network is composed of several fixed-size convolution kernels with trainable/learnable weights (coefficients) LeCun et al. (1998); GoodFellow et al. (2016). There are two important properties

of a convolutional neuron which differentiates it from a fully connected neuron: 1) it has a local receptive field. 2) it shares its weights with all other neurons at the same layer (assuming a single kernel). Therefore, the same local (non)linear transformation is applied to all regions of the input. Thus, it calculates the same transformation for an input window regardless of its position in the image. However, it is neither scale- nor rotation-invariant, and the size, shape, or orientation of the kernel also affect the output. Though many practitioners often employ basic 3×3 kernels for all tasks, others have tried varied size and shape kernels and different input samplings to improve robustness Li et al. (2017); Dai et al. (2017); Ronneberger et al. (2015); Jeon and Kim (2017). These works are reviewed in Section 2.

In this study, we describe a new and adaptive model of the convolution layer where the kernel sizes are learned during training. In this unique setting, a single convolution layer can tune and accommodate several kernel sizes at the same time. Such a layer can compute a multi-scale representation from the same input. This is achieved by an additional function which limits and controls the size of the kernel (illustrated in Figure 1). Therefore, the first important question of this paper is: can a differentiable and trainable functional form effectively control the receptive field of a kernel? We tested this ability on an auto-encoder network to learn ordinary image processing operators (e.g., Sobel filter, Gaussian blur). The second question is whether the new adaptively sized convolution kernels can provide any advantage over ordinary fixed-size kernels. In two different network structures, (simple CNN and residual) we substitute the ordinary convolution layers with the adaptive layers to compare their learning and generalization performances. We used the popular MNIST, MNIST-Cluttered, CIFAR-10, Fashion, LFW-Faces ("Labelled Faces in the Wild") datasets for the comparisons. Finally, we replaced a single convolution layer in a U-net architecture with an adaptive layer and tested it in segmentation.

The main contributions of the current paper are as follows: 1) a formal description of the two-dimensional adaptive kernel model based on a Gaussian envelope function, 2) a demonstration that the adaptive envelope makes the kernels less prone to overfitting than ordinary large kernels, 3) a demonstration that their performance is comparable to or better than the ordinary 3×3 kernels which are commonly used in vision applications.

2 Related Works

The concept of receptive field has attracted attention since the earliest studies of artificial neural networks. The stacked topology of networks enables a neuron in the deeper layers of the network to have an enlarged effective field of view on the input. Recently Luo Luo et al. (2016), found that the effective receptive field grows with the square root of the depth, and in contrary to general belief, the receptive fields of the top-layer neurons may not extend to cover the whole input domain.

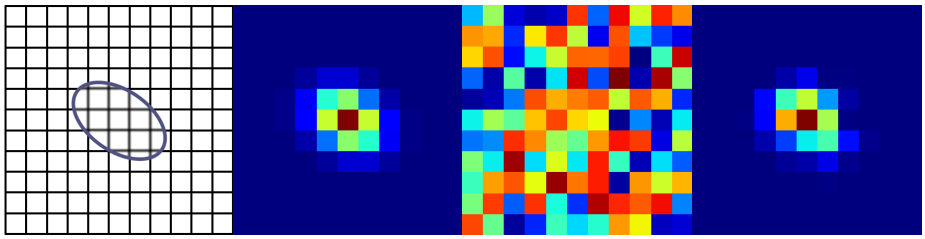

(a) Grid (b) U (c) W (d) $U \circ W$

Figure 1: Illustration of the proposed weight envelope: (a) An arbitrary differentiable envelope function on a base grid. (b) An example Gaussian envelope (U). (c) Randomly generated weights (W). (d) Weights are masked by envelope (b) through element-wise multiplication ($U \circ W$).

The importance of the receptive field in convolution was spotted by signal processing researchers, long before the deep learning community focused on it. The most relevant work is on *atrous* (dilated) convolution Holschneider et al. (1989) which used up-sampled convolution kernels by inserting zeros between the coefficients. The atrous or dilated convolution is widely used in a range of deep learning applications where multi-scale processing is crucial. The applications include image classification Yu et al. (2017), semantic segmentation Chen et al. (2018), speech synthesis&recognition Oord et al. (2016), and image denoising Tian et al. (2020).

Dilated convolution enlarges the effective receptive fields. It has been used to provide multi-scale representation in various network configurations Yu and Koltun (2016); Yu et al. (2017); Guo et al. (2020); Li et al. (2019). However, it does not solve the problem of scale completely by itself. The network architectures such as U-net aim to increase the scale tolerance of the network by creating multi-scale feature maps Ronneberger et al. (2015), whereas others such as inception included parallel convolution paths containing different fixed-size convolutions to extract multi-scale information Szegedy et al. (2017).

The shape or orientation of the kernels was also a concern. For example, Li et al. Li et al. (2017) studied optimizing kernel shapes using Lasso to create arbitrary shape kernels for audio inputs, as an alternative to commonly used square kernels optimal for natural images. Weiler et al. Weiler et al. (2018) employed steerable kernels trained from a harmonic functional basis to create orientation-sensitive kernels.

The receptive field of a convolution operation can be changed by varying the locations where input is sampled while maintaining the size of the kernel. The ordinary convolution uses fixed sampling locations with respect to the current position (i) of the kernel, (e.g., $\{i-2,\ i-1,\ i,\ i+1,\ i+2\}$), whereas dilated convolution would sample sparsely (e.g., $\{i-4,\ i-2,\ i,\ i+2,\ i+4\}$) using a fixed sampling parameter. The active convolution model Jeon and Kim (2017) attempted to learn input sampling offsets (pn) from training data, (e.g., $\{i+p0,$

$i+p1$, $i+p2$, $i+p3$, $i+p4\}$). Similarly, Dai's deformable convolution model Dai et al. (2017) dynamically computed sampling offsets per input and per location by performing additional convolutions on the input.

The adaptive model proposed here differs from these approaches in three ways: 1) it does not change the way that input is sampled, 2) it does not use secondary convolutions to compute parameters, 3) it learns the kernel size from training data. The parameters are static and not computed per input, meaning that after training the kernels are fixed.

The proposed model can be seen as an aperture-only, case-specific exemplar of a generalized form, the adaptive locally connected neuron Tek (2019, 2020) which can learn its receptive field location and aperture using a Gaussian focus attachment.

3 Method

The proposed kernel model learns the receptive field size of the kernel by training a smooth envelope function that can grow or shrink in a base kernel grid. The following sections explain the role of the envelope function, provide an appropriate functional form to construct the envelope, and discuss its parameters. Before starting, note that although commonly referred to as convolution, the operation that is studied and used in neural networks is more appropriately termed cross-correlation. Therefore, for mathematical consistency, we continue with the term cross-correlation instead of convolution, although we use the terms interchangeably for the sake of consistency with the literature.

A 2-D-matrix cross-correlation computes its output $O = [o_{i,j}]$ by calculating the weighted sum of the ($n \times n$ shaped) kernel coefficients $W = [w_{k,l}]$ times the input $X = [x_{i,j}]$ across all possible locations i, j. Therefore, the output matrix of valid size $(M - n) \times (N - n)$ can be expressed in the following form:

$$O = X \star W = \left[\sum_{k=-\lfloor n/2 \rfloor}^{\lfloor n/2 \rfloor} \sum_{l=-\lfloor n/2 \rfloor}^{\lfloor n/2 \rfloor} x_{(i+k,j+l)} w_{(k+\lfloor n/2 \rfloor, l+\lfloor n/2 \rfloor)} \right]_{i=\lfloor n/2 \rfloor, j=\lfloor n/2 \rfloor}^{M-\lfloor n/2 \rfloor, N-\lfloor n/2 \rfloor} \tag{1}$$

where, for simplicity, we can ignore the precise offsets, subscripts (e.g., $w_{k+\lfloor n/2 \rfloor, l+\lfloor n/2 \rfloor}$) and index limits to use the following form (2) which is sufficient for our discussions:

$$O = X \star W = \left[\sum_{k}^{n} \sum_{l}^{n} x_{(i+k,j+l)}\, w_{k,l} \right]_{i,j}^{M,N} \tag{2}$$

3.1 The envelope function

In the adaptive model, the kernel coefficient matrix W is paired with an envelope $U = [u_{k,l}]$ which controls kernel growth through an element-wise multiplication

(i.e. the Hadamard product):

$$O = X \star (W \circ U) = \left[\sum_{k}^{n} \sum_{l}^{n} x_{(i+k,j+l)}\, w_{k,l}\, u_{k,l} \right]_{i,j}^{M,N} \tag{3}$$

It may seem as if we are adding just another weight; however, the envelope coefficients are not independent of each other. Here, we define the envelope on a two-dimensional Euclidean space since it is the most common case which can be generalized to further dimensions. We omitted input channels in our notation; if the input contains channels, the weight matrix is three or more dimensional, so the envelope U must be repeated on that dimension. As illustrated in Figure 1, the envelope resides in a base grid which is also the kernel domain. Let us assume an $n \times n$ base grid for an odd-sized square kernel; and let U_f be a smooth and differentiable function defined in this domain by a parameter set $\boldsymbol{\theta} \in \mathbb{R}^p$ (4):

$$\begin{gathered} U_f \colon ((k,l), \boldsymbol{\theta}) \mapsto u_{k,l} \in \mathbb{R} \qquad \text{where} \\ \{(k,l) \mid k,l \in \{1,2,..,n\}\} \quad \text{and} \quad \boldsymbol{\theta} = \{\theta_1, \theta_2, ..\theta_p\} \end{gathered} \tag{4}$$

Thus, a functional form can be chosen or designed for U_f to control the envelope shape represented by the coefficients $u_{k,l}$ which mask the weights. When U_f is differentiable with respect to the parameters $\boldsymbol{\theta}$, the error derivatives can be calculated using the chain rule; and the updates can be performed using (5):

$$w_{k,l}^{'} := w_{k,l} - \eta \frac{\partial E}{\partial O} \frac{\partial O}{\partial w_{k,l}} \qquad \theta_p^{'} := \theta_p - \eta \frac{\partial E}{\partial O} \frac{\partial O}{\partial \theta_p} \tag{5}$$

where η denotes the learning rate, $w'_{k,l}$ and $\theta_p^{'}$ denote the updated kernel weight coefficient and envelope parameter respectively, and E denotes an error term. Though they seem disconnected, the updates of the envelope coefficients and weights are related. We elaborate this point by inspecting the partial derivatives of E with respect to $w_{k,l}$ and θ_p. The expression for the derivative $\partial E / \partial w_{k,l}$ (7) includes the focus coefficient $u_{k,l}$ as a scaler coefficient:

$$\frac{\partial E}{\partial w_{k,l}} = \sum_{i}^{M} \sum_{j}^{N} \frac{\partial E}{\partial o_{i,j}} \frac{\partial o_{i,j}}{\partial w_{k,l}} = \sum_{i}^{M} \sum_{j}^{N} \frac{\partial E}{\partial o_{i,j}}\, x_{(i+k,j+l)}\, u_{k,l} \tag{6}$$

$$= u_{k,l} \sum_{i}^{M} \sum_{j}^{N} \frac{\partial E}{\partial o_{i,j}}\, x_{(i+k,j+l)} \tag{7}$$

Thus, the envelope not only controls the forward signal but also affects the weight updates. Likewise, we calculate the derivative with respect to the envelope parameter θ_p:

$$\frac{\partial E}{\partial \theta_p} = \sum_{i}^{M} \sum_{j}^{N} \frac{\partial E}{\partial o_{i,j}} \frac{\partial o_{i,j}}{\partial \theta_p} \tag{8}$$

where

$$\frac{\partial o_{i,j}}{\partial \theta_p} = \sum_k^n \sum_l^n x_{i+k,j+l}\, w_{k,l} \frac{\partial u_{k,l}}{\partial \theta_p}. \tag{9}$$

Thus, we can write the following expression:

$$\frac{\partial E}{\partial \theta_p} = \sum_i^M \sum_j^N \frac{\partial E}{\partial o_{i,j}} \left(\sum_k^n \sum_l^n x_{i+k,j+l}\, w_{k,l} \frac{\partial u_{k,l}}{\partial \theta_p} \right) \tag{10}$$

We see that the derivative with respect to the envelope parameter is accumulated over both the input image and kernel, unlike the weight derivative (7) which is only accumulated over the whole image. This is because $u_{k,l}$ values are not independent of each other.

3.2 Choosing an Envelope Function

A Gaussian form is the primary candidate for the envelope function because it is continuous and differentiable, and it neither creates nor enhances extrema (Lindeberg, 2011). Its center parameter ($\boldsymbol{\mu}$) controls the position, the covariance parameter (Σ) smoothly controls the orientation and spread of the form, and s performs the normalization:

$$U_f(\boldsymbol{g}, (\mu, \boldsymbol{\Sigma})) = s\, e^{-\frac{1}{2}(\mathbf{g}-\mu)'\Sigma^{-1}(\mathbf{g}-\mu)} \tag{11}$$

In two-dimensional Euclidean space $\boldsymbol{g} \in \mathbb{R}^2$, the center is two-dimensional, $\boldsymbol{\mu} = \langle \mu_x, \mu_y \rangle$ and covariance is a 2×2 matrix $\Sigma = \left[\begin{smallmatrix} \sigma^2_{xx} & \sigma^2_{xy} \\ \sigma^2_{yx} & \sigma^2_{yy} \end{smallmatrix}\right]$. However, we exclude the rotation (and the ellipsoid kernels) from the current discussion (for an extended introduction see İlker Çam (2019); Cam and Tek (2018)). Although it is possible to train the kernel position ($\boldsymbol{\mu}$) as in Jeon and Kim (2017), we did not observe any benefit from doing so in our preliminary studies. Hence, here μ is initialized to the center of the grid and not trained. There remains only one trainable parameter σ_u which controls the size of the circular envelope shape in the set of parameters: $\boldsymbol{\theta} = \{\mu_x = 0.5n,\ \mu_y = 0.5n,\ \sigma_u\}$.

During the feed-forward execution the envelope function U_f is computed on the normalized grid coordinates $\boldsymbol{g} = \langle k/n, l/n \rangle$ with the current aperture σ_u to produce envelope coefficients $u_{k,l}$ which are multiplied element-wise with the weights $w_{k,l}$ prior to the convolution. Figure 2 depicts the weight kernels, envelope matrices and product kernels ($W \circ U$) for two example cases with relatively smaller and larger aperture (σ_u) values.

Let us denote l_2-norm of a vector as $\|\boldsymbol{x}\|_2 = \sqrt{x_1^2 + x_2^2}$, for any given $\boldsymbol{x} = \langle x_1, x_2 \rangle$. Then the partial derivative with respect to σ_u, which can replace $\partial u_{k,l}/\partial \theta_p$ in (11), can be expressed as below (12):

$$\frac{\partial U_f}{\partial \sigma_u} = s \frac{\|\boldsymbol{g} - \boldsymbol{\mu}\|_2^2}{2\sigma_u^3} e^{(-\|\boldsymbol{g}-\boldsymbol{\mu}\|_2^2/(2\sigma_u^2))} \tag{12}$$

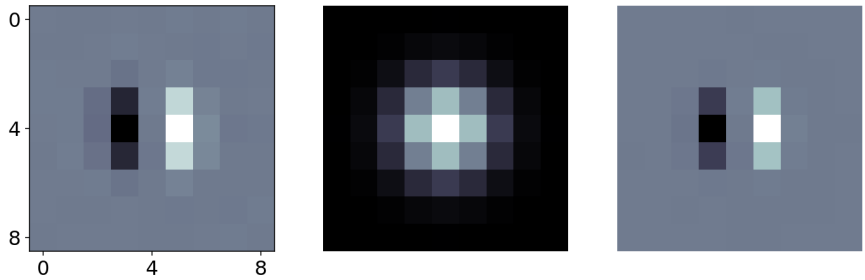

(a) Weight, smaller envelope, effective kernel on a 9×9 grid

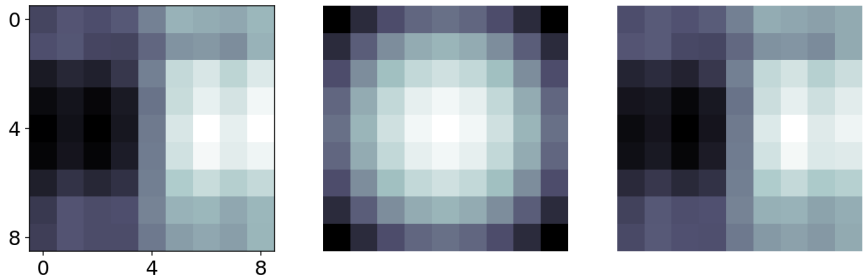

(b) Weight, larger envelope, effective kernel on a 9×9 grid

Figure 2: Examples of envelopes and effective (product) kernels.

3.3 Initialization of Envelope Parameters

Recent studies demonstrated that the initialization of weights in a neural network is crucial to improve its training and generalization capacity (He et al., 2015; Glorot and Bengio, 2010). These studies usually inspect the forward signal and backward gradient flows to suggest an optimal weight-initialization strategy. A common approach is to adjust the variance of the weights so that layer inputs and outputs have equal variance. However, in the adaptive kernel, the envelope coefficients scale the weights and change the variance of the propagated signals. Moreover, since the total fan-in of an adaptive kernel is larger than its effective fan-in it is not clear what value should be used for calculating the weight variance as recommended by common initialization schemes (He et al., 2015; Glorot and Bengio, 2010).

Nevertheless, we could derive an appropriate initialization variance for the weights in the envelope's presence. However, during the training, the updates to σ_u would change the envelope, the product kernel, and the output variance. Therefore, we approach this problem from an alternative perspective where we normalize and scale the envelope U to keep the variance of the weights unchanged by the element-wise multiplication operation $W \circ U$. Although it is not possible to keep the variance of the individual weights $w_{k,l}$ unchanged, it is possible to maintain the mean of the variances. Let us write single summa-

tion $\sum_{k,l}^{nxn}$ instead of $\sum_k^n \sum_l^n$ to simplify the notation and define the mean of variances along an $n \times n$ matrix $A = [a_{k,l}]_{n\times n}$ by:

$$\mathrm{MVar}[A] = \frac{1}{n^2} \sum_{k,l}^{n\times n} \mathrm{Var}(a_{k,l}) \tag{13}$$

Then, Theorem 3.1 states that the mean of the variances of the weights will be unchanged by the Hadamard multiplication of the envelope matrix if the mean of the expected value of the squared envelope coefficients is 1.

Theorem 3.1. $\mathrm{MVar}[W \circ U] = \mathrm{MVar}[W] = \sigma_w^2$ *when* $\frac{1}{n^2} \sum_{k,l}^{n\times n} \mathbb{E}(u_{k,l}^2) = 1.$

Proof. Since $u_{k,l}$ and $w_{k,l}$ are independent for any k, l and weights are i.i.d with zero mean (i.e. $\mathbb{E}(w_{k,l} = 0)$), we have

$$\mathrm{Var}(u_{k,l}w_{k,l}) = \mathbb{E}(u_{k,l}^2 w_{k,l}^2) - [\mathbb{E}(u_{k,l})\mathbb{E}(w_{k,l})]^2 \tag{14}$$

$$= \mathbb{E}(u_{k,l}^2)\mathbb{E}(w_{k,l}^2) = \mathbb{E}(u_{k,l}^2)\sigma_w^2 \tag{15}$$

Next, we consider the mean of variances.

$$\mathrm{MVar}(U \circ W) = \frac{1}{n^2} \sum_{k,l}^{n\times n} \mathrm{Var}(u_{k,l}w_{k,l}) \tag{16}$$

$$= \sigma_w^2 \frac{1}{n^2} \sum_{k,l}^{n\times n} \mathbb{E}(u_{k,l}^2) \tag{17}$$

$$= \mathrm{MVar}(W) \frac{1}{n^2} \sum_{k,l}^{n\times n} \mathbb{E}(u_{k,l}^2). \tag{18}$$

Hence the result follows. □

Likewise, we may consider the backward propagation of the error variance. Using (7) and the assumptions $\mathbb{E}(x_{(i+k,j+l)}) = 0$ and $\mathbb{E}(w_{k_l}) = 0$, we formulate the mean of the gradient variances of the weights as follows (20)(see the appendix for the derivation):

$$\mathrm{MVar}\left(\frac{\partial E}{\partial W}\right) = \frac{1}{n^2} \sum_{k,l}^{n\times n} \mathrm{Var}\left(\frac{\partial E}{\partial w_{k_l}}\right) \tag{19}$$

$$= \sigma_x^2 \left[\frac{1}{n^2} \sum_{k,l}^{n\times n} \mathbb{E}(u_{k,l}^2)\right] \sum_{i,j}^{M,N} \mathbb{E}\left[\left(\frac{\partial E}{\partial o_{i,j}}\right)^2\right] \tag{20}$$

Thus, Eq.(20) states that the envelope coefficients affect the mean of the gradient variances of the weights minimally if the mean expected value of the

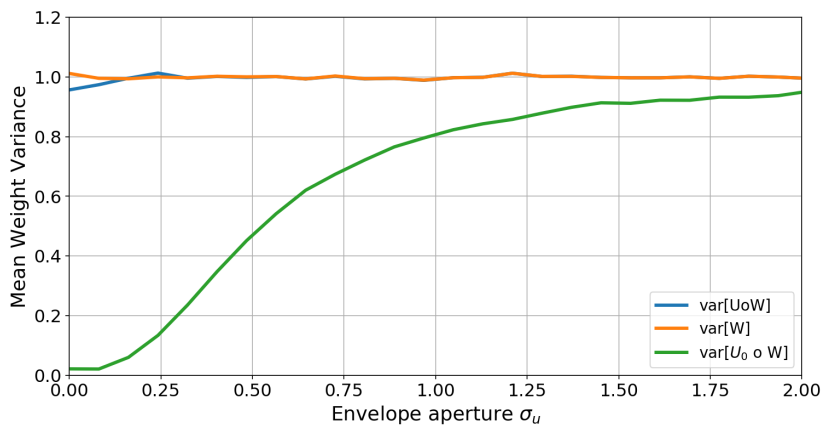

Figure 3: Mean of the kernel variances against increasing aperture σ_u. U: normalized, U_0: unnormalized envelope.

squared envelope coefficients is 1.0. However, the envelope function is a deterministic function of the random parameter σ_u which has an unknown probability density, because it will be learned by the network. However, we see that the function U_f can be scaled so that $u_{k,l}$ sum to a constant value irrespective of the σ_u value. Therefore, satisfying the condition on Theorem 3.1 translates to a condition on the norm of the $\|U\|_2^2 = \sum_{k,l}^{n\times n} u_{k,l}^2 = n^2$. In practice, a convolution layer would have more than one kernel to calculate multiple outputs. Hence, in the forward run, each kernel q calculates its envelope using its own σ_u^q, then normalizes itself using s_u^q:

$$s_u^q = \frac{n}{\sqrt{\sum_{\mathbf{g}\in A\times A} \left(e^{-\left(\frac{(\mathbf{g}-\mu)^T(\mathbf{g}-\mu))}{2(\sigma_u^q)^2}\right)}\right)^2}} \tag{21}$$

To test this proposition empirically, we set up a simple experiment. In a loop of increasing aperture (σ_u) values, we calculated the corresponding envelope matrix and also randomly sampled weight matrices of size $n\times n\times channels\times filters$ from a uniform distribution (normal distribution was also tested). The weight sample variance was calculated along the channel and filter dimensions. Figure 3 shows the mean of the variances of the weight matrices (MVar[W]) against the mean of the variances of the product kernels (MVar[$W\circ U$]) for increasing envelope (U) aperture σ_u value. It can be seen that when we normalized U_f using (21), negligibly small deviations occurred in the product kernel variance at very low aperture values, whereas the larger aperture envelopes maintained the mean of the weight variances perfectly.

On the other hand, we do not have formal guidance on the initialization of σ_u, except that it must be positive and non-zero. In practice, we have noticed that initializing σ_u in the range $[1/\mathrm{n}, \mathrm{n}]$ works well (n=kernel size). However,

during training the σ_u value must be monitored and clipped to stay above a value (e.g., 1/n) to prevent over-shrinking by high gradient values and fluctuations that may occur.

4 Experiments

We divided the experiments into four sections. First, we investigated whether the proposed adaptive kernel can learn common image processing filters. Second, we set up a simple convolutional network and compared the adaptive kernels against the ordinary convolutional kernels in their learning and generalization performance. Next, we repeated the same comparison in a popular deep architecture, ResNet He et al. (2015). Finally, we tested the adaptive kernels in a U-net Ronneberger et al. (2015) architecture for segmentation. We implemented the proposed model in Python 3 using Keras & Tensorflow Chollet et al. (2015). All code and a demo are available in Tek (2020).

4.1 Learning Basic Image Processing Kernels

We set up a simple auto-encoder network to test whether the new adaptive kernels are able to synthesize some basic image processing kernels. The network configuration can be found in the supplementary materials. The network took a single image (e.g., Figure 4a) as input to learn the outputs of nine different image processing kernels of size 9×9 pixels. The targets included the output of 3×3 Laplace, horizontal and vertical Sobel kernels, Gauss smoothing kernels of different variance, and applications of Laplace and Sobel to the Gauss-smoothed outputs, as shown in Figure 4d. Figure 4e shows the outputs of the network after 500 training iterations using stochastic gradient descent optimizer with a learning rate and momentum of 0.1 and 0.9, respectively. The training converged after 150-200 iterations, as seen in Figure 4b. In addition, we observed that the minimum square error of the adaptive convolution network was slightly lower than an ordinary convolution network which contain kernels of equal size (9×9). However, both networks were able to learn the kernels. Figure 4c shows that most of the aperture parameters (σ_u) converged at around 150-200 epochs in the adaptive network. The initial weight (W) and envelope (U) kernels are shown in Figures 4f and 4g with the learned envelope and final product kernels ($U \circ W$) are given in Figures 4h and 4i, respectively. It can be seen that the envelopes were successfully learned in the presence of the weights and vice versa. In conclusion, the adaptive kernels were able to learn basic image processing kernels of different size and character.

4.2 Comparisons in Simple Convolutional Network

Next, we compared the proposed adaptive kernels (ACONV) to the ordinary convolution (CONV) kernels in a simple convolutional classification network.

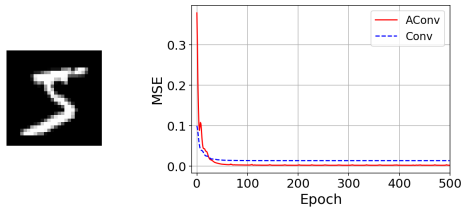

(a) Input (b) Train Loss

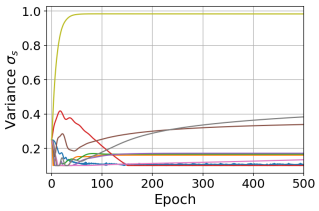

(c) Apertures

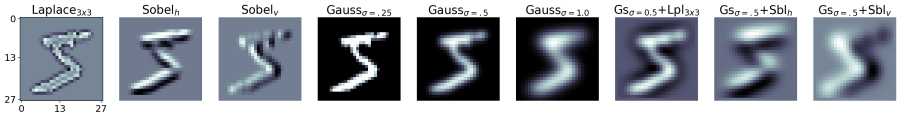

(d) Targets

(e) Predictions

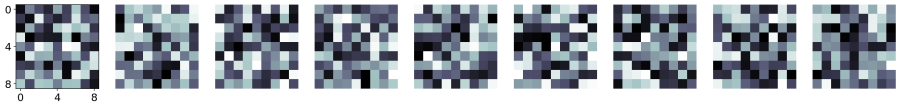

(f) Weights before training

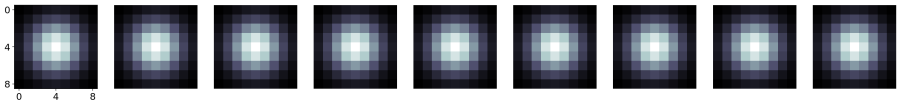

(g) Envelopes before training

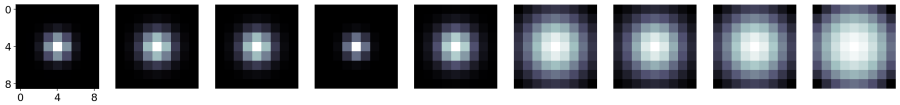

(h) Envelopes after training

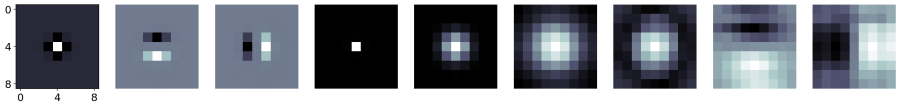

(i) Product kernels after training

Figure 4: The auto-encoder with the adaptive convolution kernels learning basic image kernels. a) Input image. b) Mean Squared Loss. c) Change of σ_u during training. d) Output images (targets) created by image processing kernels (Laplace, Sobel, Gauss and combinations). e) Network predicted outputs after 2500 training updates. f) Initial weights. g) Initial envelopes. h) Learned envelopes. i) Effective (product) kernels).

The basic network configuration is summarized in Table 1 (see the supplementary figures for a plot of the network graph). The network was built using two consecutive convolutional layers (CONV or ACONV) of 32 kernels followed by a single max-pool layer of size 2×2 and a dense classification layer of 256 units surrounded by two drop-out layers. All neuronal layers were followed by batch normalization (BN) and rectified linear activation units (RELU). The network output was formed of c Softmax units, where c was equal to the number of dataset categories.

The popular gray-scale MNIST character recognition dataset (LeCun et al., 1998) (MNIST) was the first place to start. More challenging datasets were also used: a cluttered version of MNIST data (CLT), comprised of randomly transformed MNIST samples superimposed on cluttered 60×60 backgrounds (Jaderberg et al., 2015); the CIFAR-10 general object classification dataset which is composed of 32×32×3 RGB images of ten concrete categories such as car, plane, bird, horse (Krizhevsky, 2009); and the FASHION (clothes) dataset which is arranged similarly to MNIST (Xiao et al., 2017) to include 10 categories such as t-shirt, pullover, and coat. These almost-standard datasets had already been separated into training (60000) and test (validation) instances (10000). The tests also included the "Faces in the Wild" dataset (LFW-Faces) (Huang et al., 2007) as a benchmark for face verification. The LFW-Faces set contains 13233 images of 5749 people; in order to reduce the number of output classes, individuals with less than 20 images were excluded from the experiments, resulting in a dataset of 3023 (2267 training, 756 validation) images of 62 people.

All comparisons were repeated with different kernel sizes n (3×3, 5×5, 7×7, 9×9) and with 5 different random initializations. To ensure a fair comparison, we fine-tuned the ordinary kernel network to get the best validation accuracy, then replaced the ordinary kernel with the adaptive kernel. We then tuned the model-specific parameters such as the initial σ_u's before comparing the two cases. In other words, for the given configuration, we compared the maximum performance we could achieve using the ordinary kernel with the adaptive kernel. We used stochastic gradient descent optimizer with an initial learning rate of $0.1 \times \eta_{dset}$ (dataset specific multiplier) and a momentum of 0.9 and gradient clip value of 1.0. We further employed an adaptive learning rate schedule which monitors the validation loss and drops the learning rate by a factor of 0.9 when no improvement is seen in the past 10 epochs. Table 1 lists other important parameters used in training. We initialized the aperture values σ_u^q with linearly spaced values in the range [0.1,0.5]. In addition, we attached a clip function to the optimizer to clip the σ_u values within range $[1/n, n]$ after each update.

Figure 5 shows a comparison of the mean validation accuracies with respect to training epochs. In all five datasets, the adaptive layers performed better than their fixed-size counterparts. As anticipated, the 3×3 ACONV kernels performed the least effectively, since there is limited room to operate the adaptive aperture. In contrast, the 7×7 and 9×9 kernels often performed the best.

Table 2 summarizes the results that were calculated using Algorithm 1 across five repeats. Comparing the mean peak validation accuracies, the adaptive filter reached higher validation accuracies in all five datasets. The t-tests compared

Algorithm 1: Training, validation, and optimization procedure

Input: network, dataset, $N_{repeats}$, N_{epochs}, n: kernel size, η_{dset}: learning rate multiplier, $m = 0.9$: momentum rate. *OPT: SGD (stochastic gradient descent with momentum) or Adam or SGD with Cyclic Schedule*

Output: BestTestResults: list of best test accuracies.

```
begin
    BestTestResults = []
    for r ← 0 to N_repeats − 1 do
        EpochAccuracyList = []
        trainX,trainy,testX,testy = split(dataset)
        for e ← 0 to N_epochs − 1 do
            for each batch (Xinputs, targets) in (trainX, trainy) do
                params ← network.trainableparams
                pred ← network.output(Xinputs)
                loss ← categoricalcrossentropy(targets, pred)
                updates ← OPT(loss, params, 0.1* η_dset, m, clipvalue=1.0)
                if type(network) == focused then
                    updates.append(clip(params.sigma, 1/n, n))
                network.update(updates)
            score ← accuracyscore(network, testX,testy)
            EpochAccuracyList.append(score)
        maxscore ← max(EpochAccuracyList)
        BestTestResults.append(maxscore)
```

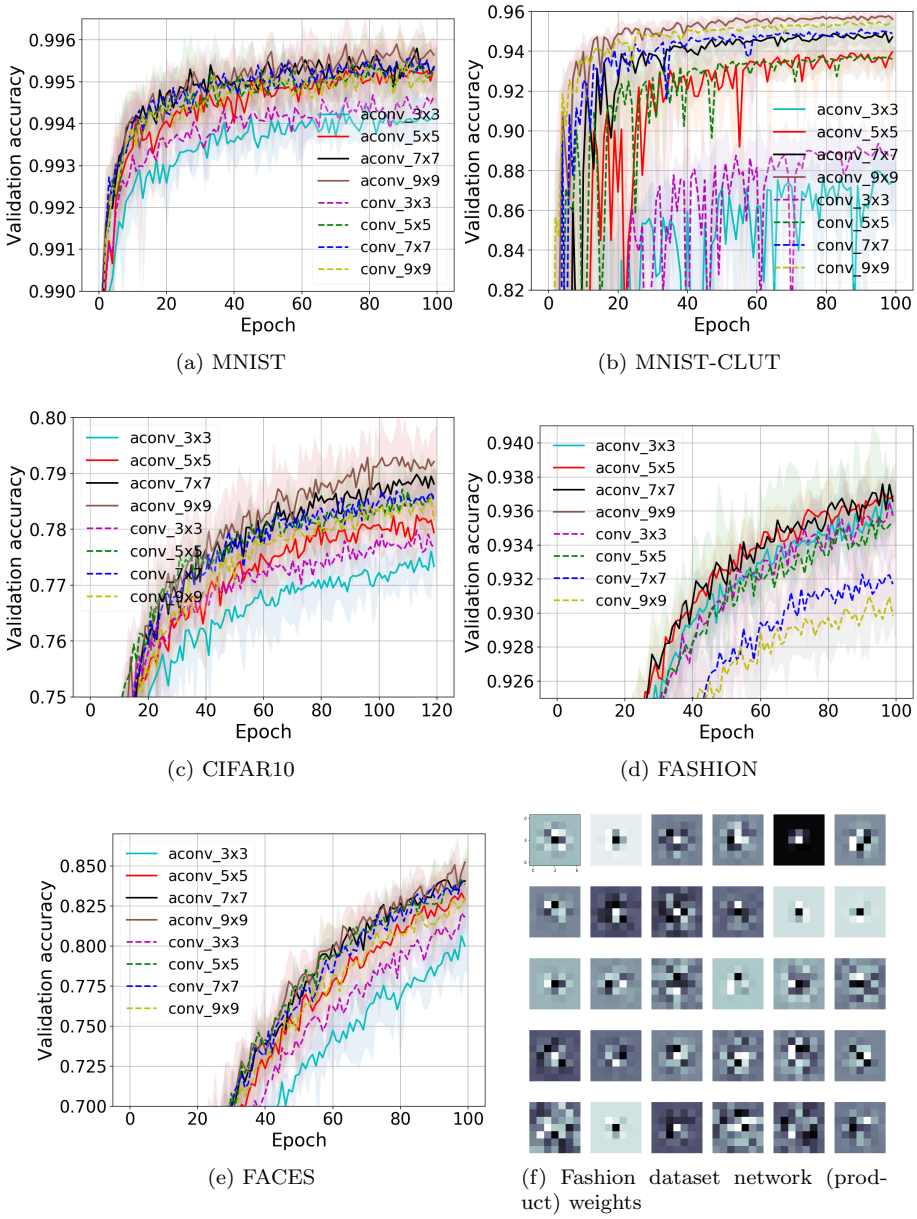

(a) MNIST

(b) MNIST-CLUT

(c) CIFAR10

(d) FASHION

(e) FACES

(f) Fashion dataset network (product) weights

Figure 5: Simple Convolutional Network Comparisons: a-e) Validation Accuracy Plots. f) Learned effective (product) kernels (ACONV-1) on FASHION dataset.)

Table 1: Left: simple convolutional network. Right: ResNet network and training parameters.

Simple Network	MNIST	CLUT	CFR10	Fashion	Faces
Num Params	1.6M	75K	1.4M	0.9M	0.9M
Batch	128	128	64	256	8
Augment	False	False	False	False	True
Epoch	100	100	120	100	100
Dropouts	0.5	0.5	0.5	0.5	0.5
η_{dset} (x 0.1)	0.1	1.0	0.1	0.1	0.01

ResNet Params	MNIST	CLUT	CFR10	Fashion	Faces
Filters	16	16	16	16	16
Num Blocks	1	1	3	2	2
Num Layers	11	11	20	14	14
Num Params	75K	75K	1.4M	0.9M	0.9M
Batch	128	128	128	128	8
Augment	False	False	True	True	True
Epoch	50	50	200	100	150
Dropout	0.5	-	-	-	-
η_{dset} ($\times$ [1e-3$\rightarrow$0.5])	1.0	1.0	0.1	0.1	0.1

Table 2: Validation performances of simple ordinary (CONV) and adaptive (ACONV) convolution networks on popular image classification sets. The two-tailed t-tests are included for each case. N (repeats)=5, p: p-value, highlights indicate *: p-value<0.05. Best size indicates the best performing (and t-test comparison) kernel size.

	MNIST		CLT		CIFAR-10		Fashion		LFW-Faces	
	Mn±std	Max	Mn±std	Max	Mn±std	Max	Mn±std	Max	Mn±std	Max
CONV	99.58±2e-4	99.61	95.59±2e-3	95.9	78.74±2e-3	78.95	93.31±6e-4	93.38	83.87±7e-3	85.66
ACONV	**99.63±2e-4**	**99.66**	**95.9±9e-4**	**96.06**	**79.63±5e-3**	**80.1**	**93.84±2e-3**	**94.12**	**85.79±6e-3**	**86.93**
T-Test (t,p)	2.88	0.02*	2.46	0.039*	3.28	0.011*	6.16	2.7e-4*	3.89	4.5e-3*
Best size	9×9		9×9		9×9		7×7		9×9	

the means of peak validation accuracies when using the kernel size of the maximum peak performance and results confirmed that the accuracy improvements were all statistically significant.

Another observation was about the performance of ordinary kernels in different sizes. Although 3×3 sized kernels are preferred in most applications, we observed that the larger kernels produced significantly better results in our setting. Finally, Figure 5e depicts the learned kernels during the Fashion dataset training which demonstrates the varying kernel sizes.

Table 3: Validation performances of Residual ordinary (CONV) and adaptive (ACONV) convolution networks (ResNet) in popular image classification sets. The two-tailed t-tests are included for each case. N (repeats)=5, p: p-value, highlights indicate *: p-value<0.05.

	MNIST		CLT		CIFAR-10		Fashion		LFW-Faces	
	Mn±std	Max	Mn±std	Max	Mn±std	Max	Mn±std	Max	Mn±std	Max
CONV	99.69±2e-4	99.71	98.93±4e-4	99.01	91.3±2e-3	91.71	93.95±1e-4	94.12	**96.15±7e-3**	**97.48**
ACONV	**99.70±1e-4**	**99.73**	**99.06±8e-4**	**99.17**	**92.21±3e-3**	**92.68**	**94.72±2e-3**	**95.01**	94.83±9e-3	96.06
T-Test (t,p)	0.67	0.51	2.72	**0.02***	5.12	**9e-4***	7.01	**1e-4***	-2.19	0.059
Best size	5×5		7×7		5×5		7×7		3×3	

4.3 Comparisons in Deep Residual Network (ResNet)

Next, we compared the proposed adaptive kernel (ACONV) to ordinary convolution (CONV) kernels in a modern successful network architecture setup for classification, ResNet He et al. (2015). For the comparison, we redefined the basic convolutional block to employ either an adaptive or ordinary convolutional layer selectively (see the supplementary figures). We used the same datasets as the previous experiments. To train the networks faster (and avoid local minima) we employed a one-cycle learning rate schedule function that starts the learning rate from 0.001 before rising $0.5\times\eta_{dset}$ (η_{dset}: dataset learning rate multiplier) in the first half of the training session and then dropping down to 0.001 again towards the end of the training. We used data augmentation in the Fashion, CIFAR-10 and Faces datasets. Table 1 shows the remaining parameters.

The comparison plots of the validation performances are shown in Figure 6. By inspecting the plots, we observed clear performance gains from adaptive 5×5 and 7×7 kernels in the MNIST-CLUT, CIFAR-10, and Fashion datasets. It was difficult to identify the best performers in the MNIST and Faces tests from the accuracy plots. Table 3 demonstrates that the peak performances were those of the adaptive convolution (ACONV), with the exception of the Faces dataset. The mean peak validation accuracy differences in MNIST-CLUT, CIFAR-10, and Fashion were statistically significant. Further inspection of the Faces dataset results (see also Figure 6e), revealed that the highest mean peak accuracy was achieved by the ACONV 5×5 network (96.28%); however, all comparisons were made at the kernel size which achieved the maximum peak accuracy, which was (CONV) 3×3 in this case.

The active convolution model by Jeon Jeon and Kim (2017) was also tested on the CIFAR-10 dataset, and test accuracy was reported as 92.46% (single value) for an active convolution ResNet of 5 blocks and 32 layers. In contrast, we used 3 blocks and 20 layers which produced a maximum accuracy of 92.68% and mean of 92.21%.

To explain the differences between the adaptive and ordinary convolutions, we computed the deep Taylor Montavon et al. (2017); Alber et al. (2018) decompositions for three examples selected from the MNIST, Fashion and CIFAR-10 validation sets. The decompositions depicted in Figure 6f represent the relevancy of individual input pixels back-traced from the network predictions. The heatmaps computed for CONV 3×3 and CONV 7×7 networks show that the larger kernel size network caused smoother and fuzzier relevancy regions. On the flip side, ACONV 7×7 input heatmaps were smoother than CONV 3×3 but sharper than CONV 7×7. Moreover, the CIFAR-10 heatmaps revealed that ACONV 7×7 input representation and focus was even more precise than CONV 3×3.

4.4 Comparisons in U-Net for Segmentation

The next experiment investigated the performance of the adaptive convolution in a U-shaped network architecture Ronneberger et al. (2015) which allows end-to-

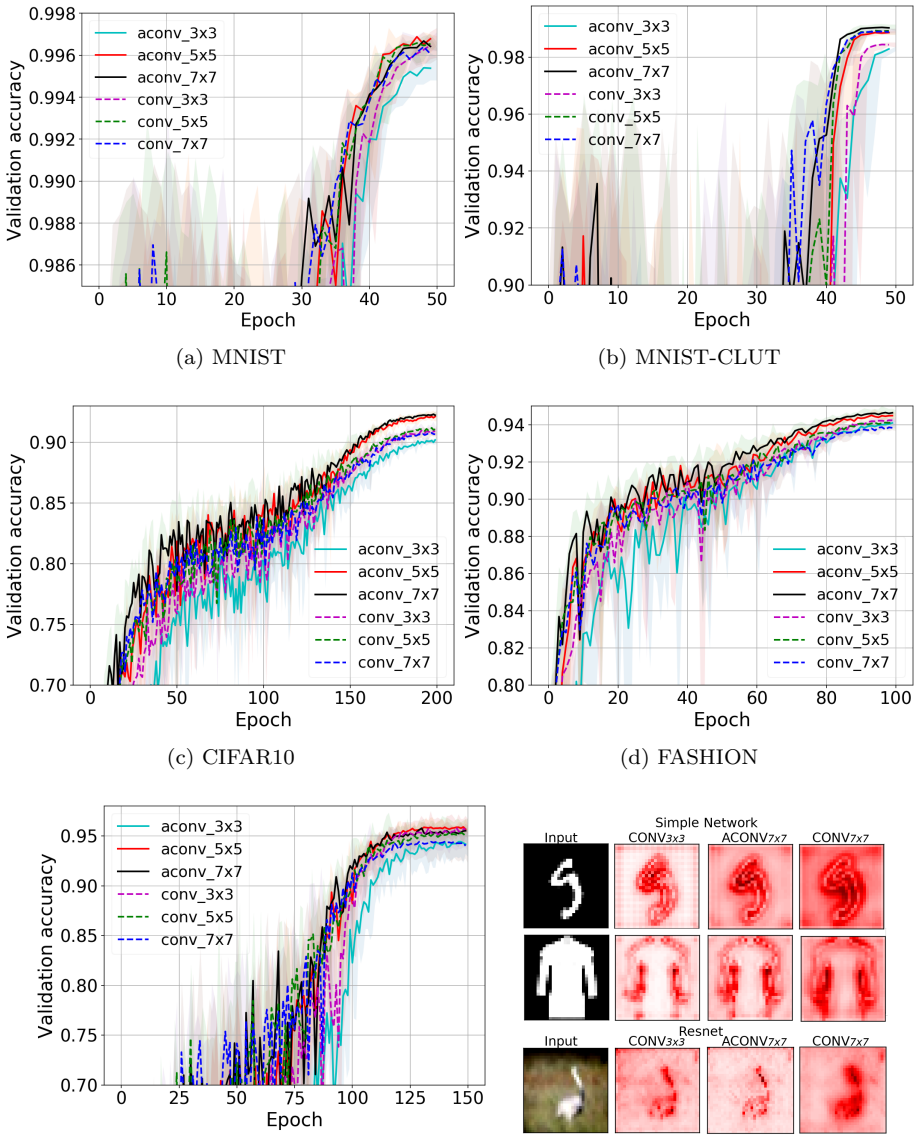

(a) MNIST

(b) MNIST-CLUT

(c) CIFAR10

(d) FASHION

(e) FACES

(f) Deep Taylor Decompositions

Figure 6: Experiment 3: Resnet Comparisons. a-e) Validation Accuracy Plots. f) Deep Taylor Decompositions.

end image segmentation. The baseline code collected from Keras library Chollet et al. (2015) implements an efficient U-net architecture by using convolution, separable convolution, and deconvolution layers (see the supplementary for the configuration used). We took the first, ordinary convolution layer and replaced it with an adaptive kernel layer, then compared it against the original. The separable convolutions or deconvolution layers in the network were not replaced because they were not implemented in our adaptive kernel framework. The network was setup to learn tri-map segmented outputs of the images in the Oxford Pets-III dataset Parkhi et al. (2012), where the output classes are pet, border, and background. The dataset includes 7349 pictures of different dog and cat breeds together with their tri-map segmentation annotations.

We resized all images to 128×128 and then used the splits provided in the dataset to create training and validation sets of 3680 and 3669 instances. We used sparse categorical entropy loss and Adam optimizer using a batch size of 64 and a fixed learning rate of 0.01. We employed and compared 7×7 adaptive kernels with 7×7 and 3×3 (baseline) ordinary kernels. All three networks were trained with five random initializations. Figure 7a and 7b plot the mean loss and validation accuracies over 75 training epochs. We observed that the networks started overfitting at shifted iterations (>≈40 epochs). However, the mean training loss value reached by ACONV 7×7 was lower than that of the ordinary kernels. Moreover, the mean peak validation accuracy achieved by ACONV 7×7 (86.21%±3e-3) was significantly higher than that of CONV 7×7 (85.54%±3e-3) with (p=0.0062), but only marginally higher than the accuracy reached by CONV 3×3 (85.89%±3e-3) (p=0.12).

Figures 7c through 7q compare the segmentation outputs qualitatively. While the output maps look very similar, the adaptive convolution layer network produced slightly more accurate border regions.

4.5 Time Complexity

The additional complexity of the adaptive kernel is due to the calculation of the envelope function which depends on the kernel size only; it is independent of the input width and height. During training, the envelope function must be calculated for each batch using the current aperture values of individual filters σ_u^q while during back-propagation, an extra gradient is calculated for σ_u^q. In the MNIST training, we recorded the following forward+backward mean batch (128) step-times for different kernel sizes of ACONV, where the time for the ordinary CONV kernel of the same size is given in parentheses: {3 × 3: 215us (166us), 5×5: 238us (180us), 7×7: 321us (282us), 9×9: 340us (285us)}, on a laptop equipped with i7-8565U and NVIDIA 1650 GPU. Therefore, the adaptive network was ≈1.2 to ≈1.3 times slower than an ordinary kernel of the same size in training. However, we must note that the current implementation was not optimized for speed at all. In addition, in run-time the overhead of the envelope can be removed by using the learned product kernels ($U \circ W$).

5 Discussions

The experiments demonstrated the feasibility of the proposed adaptive kernel model. First, the adaptive kernel was able to learn different image processing filters without encountering any difficulty. The learned kernel shapes demonstrated that the envelope and weights were able to co-adapt successfully during the training.

Second, the comparative tests in a simple convolutional network configuration demonstrated the learning and generalization performances on popular image classification datasets. In all datasets, the adaptive kernels provided significant but slight improvements in generalization performance, more than the potential gains that would be achieved by using the ordinary kernels of larger size.

In the ResNet architecture, the use of adaptive kernels resulted in better generalization performance compared to the ordinary kernels in all datasets except Faces, in which the maximum peak validation accuracy was in favor of the ordinary kernel of size 3×3. However, we noted that the mean accuracy of the 5×5 adaptive kernel was higher. Therefore, we recommend employing 5×5 or 7×7 adaptive kernels for potential performance gains in ResNet architectures. An additional insight gained from our experiments was that, in contrast to the widely accepted usage of 3×3 ordinary convolution kernels, the larger kernels may work better in ResNet for some datasets.

The kernels learned in the Fashion dataset (Figure 5f) verified that the adaptive layers were able to create varying-sized kernels. Furthermore, the deep Taylor decomposition analysis of the compared Resnets (Figure 6f) displayed evidence for multi-scale representation computed by the adaptive convolution networks.

In a brief segmentation experiment, we tested the adaptive kernels in an efficient U-net architecture by replacing a single convolution layer, which resulted in an improved segmentation performance against the same (larger) size kernels; however, there was no significant gain compared to smaller 3×3 kernel layer. However, the current state-of-the-art image segmentation methods use more complex architectures and architecture-search algorithms Zhang et al. (2020); Howard et al. (2019). Therefore, it would be appropriate to study adaptive kernels for segmentation in a dedicated study.

In summary, the experiments demonstrated that the adaptive kernel model is an effective alternative to the ordinary convolution kernel. It can create varying-sized kernels in a single layer. It is less prone to overfitting than an ordinary large convolution kernel (5×5, 7×7, 9×9) while providing better or comparable performance to the widely employed 3×3 ordinary kernel.

6 Conclusion

In conclusion, we here propose an adaptive convolution kernel which is able to learn its size by training with backpropagation. The new model is standalone,

modular and compatible with existing Keras and Tensorflow backends. Hence, one can easily import and attach the proposed adaptive layer into a network and train it with any stride or dilation factor. The single additional requirement is to apply a clip (callback) to the aperture parameter (σ_u) to keep it above a minimum positive value during training iterations.

There were some limitations to our study, which may be addressed by future work. The current state-of-the-art networks require large resources to set up, tune and optimize on larger datasets. It will be interesting to observe the learning performance of the adaptive kernels in a state-of-the-art network on one of the large datasets. Next, it will be necessary to set up a dedicated segmentation study to compare the adaptive kernel model against the ordinary kernels and other adaptive methods such as deformable or active convolution models.

Funding

This work was supported by TUBITAK 1001 programme (no: 118E722), Isik University BAP programme (no: 16A202), and NVIDIA hardware donation of a Tesla K40 GPU unit.

Acknowledgments

Earlier implementation and experiments were conducted by İ. Çam; a draft was prepared by İ. Çam, F. B. Tek coded the kernels again in Keras/Tensorflow performed the current experiments, and wrote the current paper. D. Karlı contributed to the mathematical model and proofs. Thanks to Mert Mısırlıoğlu for helping with the ResNet experiment setup.

References

A. Krizhevsky, I. Sutskever, G. E. Hinton, Imagenet classification with deep convolutional neural networks, in: Communications of the ACM, 2017.

F. N. Iandola, M. W. Moskewicz, K. Ashraf, S. Han, W. J. Dally, K. Keutzer, Squeezenet: Alexnet-level accuracy with 50x fewer parameters and <1mb model size, arXiv abs/1602.07360 (2016).

K. He, X. Zhang, S. Ren, J. Sun, Deep residual learning for image recognition, in: IEEE CVPR, 2015.

T. Poggio, T. Serre, Models of visual cortex, Scholarpedia 8 (2013) 3516.

D. . T. W. Hubel, Receptive fields, binocular interaction and functional architecture in the cat's visual cortex., J Physiology 160 (1962) 106–54.

Y. LeCun, L. Bottou, Y. Bengio, P. Haffner, Gradient-based learning applied to document recognition, in: Proc. of the IEEE, volume 86, 1998, pp. 2278–2324.

I. GoodFellow, Y. Bengio, A. Courville, Deep Learning, The MIT Press, 2016.

X. Li, F. Li, X. Fern, R. Raich, Filter shaping for convolutional neural networks, in: ICLR, 2017.

J. Dai, H. Qi, Y. Xiong, Y. Li, G. Zhang, H. Hu, Y. Wei, Deformable convolutional networks, in: ICLR, 2017.

O. Ronneberger, P. Fischer, T. Brox, U-net: Convolutional networks for biomedical image segmentation, in: MICCAI, 2015, pp. 18–29.

Y. Jeon, J. Kim, Active convolution: Learning the shape, in: IEEE CVPR, 2017.

W. Luo, Y. Li, R. Urtasun, R. Zemel, Understanding the effective receptive field in deep convolutional neural networks, in: Adv. in Neural Information Processing Systems, 2016, pp. 4898–4906.

M. Holschneider, R. Kronland-Martinet, J. Morlet, P. Tchamitchian, A real-time algorithm for signal analysis with the help of the wavelet transform, Wavelets: Time-Frequency Methods and Phase Space (1989) 289–297.

F. Yu, V. Koltun, T. Funkhouser, Dilated residual networks, in: IEEE CVPR, 2017.

L. Chen, G. Papandreou, I. Kokkinos, K. Murphy, A. L. Yuille, Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs, IEEE Transactions on Pattern Analysis and Machine Intelligence 40 (2018) 834–848.

A. Oord, S. Dieleman, H. Zen, K. Simonyan, O. Vinyals, A. Graves, N. Kalchbrenner, A. Senior, K. Kavukcuoglu, Wavenet: a generative model for raw audio, ArXiv, Corr abs/1609.03499 (2016).

C. Tian, Y. Xu, W. Zuo, Image denoising using deep cnn with batch renormalization, Neural Networks 121 (2020) 461 – 473.

F. Yu, V. Koltun, Multi-scale context aggregation by dilated convolutions, in: ICLR, 2016.

B. Guo, K. Song, H. Dong, Y. Yan, Z. Tu, L. Zhu, Nernet: Noise estimation and removal network for image denoising, J. Vis. Commun. Image R. 71 (2020) 102851.

H. Li, F. Qi, G. Shi, C. Lin, A multiscale dilated dense convolutional network for saliency prediction with instance-level attention competition, J. Vis. Commun. Image R. 64 (2019) 102611.

C. Szegedy, S. Ioffe, V. Vanhoucke, A. Alemi, Inception-v4, inception-resnet and the impact of residual connections on learning, in: Conf. on Artificial Intelligence, 2017.

M. Weiler, F. A. Hamprecht, M. Storath, Learning steerable filters for rotation equivariant cnns, in: IEEE CVPR, 2018.

F. B. Tek, Uyarlanır yerel bağlı nöron modelinin İncelemesi, Bilişim Teknolojileri Dergisi 12 (2019) 307–317.

F. B. Tek, Adaptive locally connected neural network, Neurocomputing, Online, In Press (2020).

T. Lindeberg, Generalized Gaussian scale-space axiomatics comprising linear scale-space, affine scale-space and spatio-temporal scale-space, Journal of Mathematical Imaging and Vision 40 (2011) 36–81.

İlker Çam, Learning Filter Scale and Orientation In Convolution Neural Networks, M. sc., Isik University, 2019.

I. Cam, F. B. Tek, Learning filter scale and orientation in cnns, arXiv preprint arXiv:1803.00388 (2018).

K. He, X. Zhang, S. Ren, J. Sun, Delving deep into rectifiers: Surpassing human-level performance on imagenet classification, in: ICCV, 2015, pp. 1026–1034.

X. Glorot, Y. Bengio, Understanding the difficulty of training deep feedforward neural networks, in: Proc. of Machine Learning Research, volume 9, 2010, pp. 249–256.

F. Chollet, et al., Keras, `https://keras.io`, 2015.

F. B. Tek, 2020. URL: `https://github.com/btekgit/AdaptiveCNN`.

M. Jaderberg, K. Simonyan, A. Zisserman, K. Kavukcuoglu, Spatial transformer networks, in: NeurIPS, volume 28, 2015.

A. Krizhevsky, Learning Multiple Layers of Features from Tiny Images, Technical Report, Canadian Institute For Advanced Research, 2009.

H. Xiao, K. Rasul, R. Vollgraf, Fashion-mnist: a novel image dataset for benchmarking machine learning algorithms, arXiv cs.LG/1708.07747 (2017).

G. B. Huang, M. Ramesh, T. Berg, E. Learned-Miller, Labeled Faces in the Wild: A Database for Studying Face Recognition in Unconstrained Environments, Technical Report 07-49, University of Massachusetts, Amherst, 2007.

G. Montavon, S. Lapuschkin, A. Binder, W. Samek, K.-R. Müller, Explaining nonlinear classification decisions with deep taylor decomposition, Pattern Recognition 65 (2017) 211 – 222.

M. Alber, S. Lapuschkin, P. Seegerer, M. Hägele, K. T. Schütt, G. Montavon, W. Samek, K.-R. Müller, S. Dähne, P.-J. Kindermans, innvestigate neural networks!, 2018.

O. M. Parkhi, A. Vedaldi, A. Zisserman, C. V. Jawahar, Cats and dogs, in: IEEE CVPR, 2012.

H. Zhang, C. Wu, Z. Zhang, Y. Zhu, Z. Zhang, H. Lin, Y. Sun, T. He, J. Mueller, R. Manmatha, M. Li, A. Smola, Resnest: Split-attention networks, arXiv, Corr 2004.08955 (2020).

A. Howard, M. Sandler, G. Chu, L.-C. Chen, B. Chen, M. Tan, W. Wang, Y. Zhu, R. Pang, V. Vasudevan, Q. V. Le, H. Adam, Searching for mobilenetv3, arXiv, Corr 1905.02244 (2019).

A Derivation of mean of the variances of the weight derivatives

We assume that the inputs and weights are i.i.d (independent and identically distributed) and the expected values are zero, i.e. $\mathbb{E}(x_{(i,j)}) = 0$ and $\mathbb{E}(w_{k_l}) = 0$. Let us recall the expression for the mean of variances,

$$\mathrm{MVar}\left(\frac{\partial E}{\partial W}\right) = \frac{1}{n^2}\sum_{k,l}^{n\times n}\mathrm{Var}\left(\frac{\partial E}{\partial w_{k,l}}\right) \tag{22}$$

Here, the expected value of the derivative is also zero $\mathbb{E}\left(\frac{\partial E}{\partial w_{k,l}}\right) = 0$ by the independence of variables and $\mathbb{E}(x_{i,j}) = 0$. Hence:

$$\mathrm{Var}\big(\frac{\partial E}{\partial w_{k,l}}\big) = \mathbb{E}\left(\big[\frac{\partial E}{\partial w_{k,l}}\big]^2\right) = \mathbb{E}\left((u_{k,l}^2)\left[\sum_{i,j}^{M,N}\frac{\partial E}{\partial o_{i,j}}x_{i+j,k+l}\right]^2\right) \tag{23}$$

By independence, the last line equals:

$$= \mathbb{E}(u_{k,l}^2)\left[\sum_{i,j}^{M,N}\mathbb{E}\left(\left[\frac{\partial E}{\partial o_{i,j}}x_{i+j,k+l}\right]^2\right) + 2\sum_{i,j}^{M,N}\sum_{p,r}^{i-1,j-1}\mathbb{E}\left(\frac{\partial E}{\partial o_{i,j}}\frac{\partial E}{\partial o_{p,r}}x_{i+k,j+l}\ x_{i+p,k+r}\right)\right] \tag{24}$$

The second term is zero since $\mathbb{E}(x_{i,j}) = 0$ and independence of input $x_{i,j}$ from the other variables. Then,

$$\mathrm{Var}\big(\frac{\partial E}{\partial w_{k,l}}\big) = \mathbb{E}(u_{k,l}^2)\sum_{i,j}^{M,N}\mathbb{E}(x_{i+k,j+l}^2)\,\mathbb{E}\left(\big[\frac{\partial E}{\partial o_{i,j}}\big]^2\right) \tag{25}$$

Since $x_{i,j}$ are i.i.d with variance σ_x^2 and expectation zero, we can write

$$\mathrm{Var}\big(\frac{\partial E}{\partial w_{k,l}}\big) = \sigma_x^2 \; \mathbb{E}(u_{k,l}^2) \; \sum_{i,j}^{M,N} \mathbb{E}\Big(\big[\frac{\partial E}{\partial o_{i,j}}\big]^2\Big) \tag{26}$$

Then the mean of variances is as follows:

$$\mathrm{MVar}\big(\frac{\partial E}{\partial W}\big) = \sigma_x^2 \Big[\frac{1}{n^2} \sum_{k,l}^{n\times n} \mathbb{E}(u_{k,l}^2)\Big] \sum_{i,j}^{M,N} \mathbb{E}\Big[\big(\frac{\partial E}{\partial o_{i,j}}\big)^2\Big] \tag{27}$$

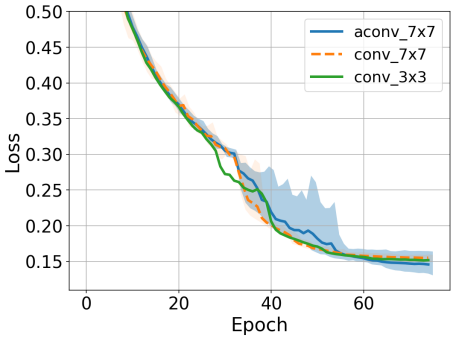

(a) Train Categorical Cross Entropy loss

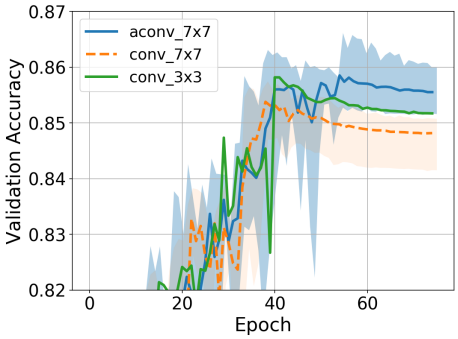

(b) Validation Accuracy

(c) Input (d) True mask (e) ACONV 7×7 (f) CONV 7×7 (g) CONV 3×3

(h) Input (i) True mask (j) ACONV 7×7 (k) CONV 7×7 (l) CONV 3×3

(m) Input (n) True mask (o) ACONV 7×7 (p) CONV 7×7 (q) CONV 3×3

Figure 7: U-Net segmentation comparisons. a-b) Training and validation accuracy plots of U-net with adaptive convolution (ACONV 7×7) and ordinary convolutions (CONV 7×7 and CONV 3×3). c-q) Input, true mask, and predictions of the networks for three different input images from the validation set.