

# Resolution learning in deep convolutional networks using scale-space theory

Silvia L. Pintea\*, Nergis Tömen\*, Stanley F. Goes, Marco Loog, and Jan C. van Gemert

**Abstract**—Resolution in deep convolutional neural networks (CNNs) is typically bounded by the receptive field size through filter sizes, and subsampling layers or strided convolutions on feature maps. The optimal resolution may vary significantly depending on the dataset. Modern CNNs hard-code their resolution hyper-parameters in the network architecture which makes tuning such hyper-parameters cumbersome. We propose to do away with hard-coded resolution hyper-parameters and aim to learn the appropriate resolution from data. We use scale-space theory to obtain a self-similar parametrization of filters and make use of the N-Jet: a truncated Taylor series to approximate a filter by a learned combination of Gaussian derivative filters. The parameter  $\sigma$  of the Gaussian basis controls both the amount of detail the filter encodes and the spatial extent of the filter. Since  $\sigma$  is a continuous parameter, we can optimize it with respect to the loss. The proposed N-Jet layer achieves comparable performance when used in state-of-the-art architectures, while learning the correct resolution in each layer automatically. We evaluate our N-Jet layer on both classification and segmentation, and we show that learning  $\sigma$  is especially beneficial when dealing with inputs at multiple sizes.

**Index Terms**—Scale-space theory, Gaussian basis approximation, resolution learning in deep networks.

## I. INTRODUCTION

Resolution defines the *inner scale* at which objects should be observed in an image [1]. To control the resolution in a network, one can change the filter sizes or feature map sizes. Because there is a maximum frequency that can be encoded in a limited spatial extent, the filter sizes and feature map sizes define a lower bound on the resolution encoded in the network. CNNs typically use small filters of  $3 \times 3$  px or  $5 \times 5$  px, where the first layers are forced to look at detailed, local image neighborhoods such as edges, blobs, and corners. As the network deepens, each subsequent convolution increases the receptive field size linearly [2], allowing the network to combine the detailed responses of the previous layer to obtain textures, and object parts. Going even deeper, strategically placed memory-efficient subsampling operations reduce feature maps to half their size which is equivalent to increasing the receptive field multiplicatively. At the deepest layers, the receptive field spans a large portion of the image and objects emerge as combinations of their parts [3]. The resolution, as controlled by the sizes of the receptive field and feature maps, is one of the fundamental aspects of CNNs.

Dr. S.L. Pintea, Dr. N. Tömen, Dr. M. Loog and Dr. J.C. van Gemert are with the Computer Vision Lab, Delft University of Technology, 2628CD Delft, Netherlands (e-mail: s.l.pintea@tudelft.nl; n.tomen@tudelft.nl). MSc. S.F. Goes is with Q.E.F Electronic Innovations.

\* Shared first authorship with equal contributions.

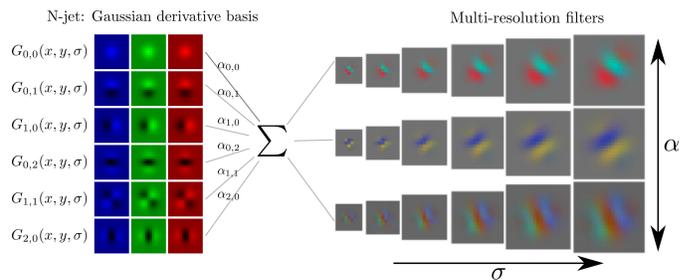


Fig. 1. Illustration of how an N-Jet Gaussian derivative basis parametrizes the shape and size of the filters. A linear combination of Gaussian derivative basis filters (left) weighted by  $\alpha$  parameters span a Taylor series to locally approximate the shape of image filters. The filters are self-similar: the  $\sigma$  parameter can change the size of the filters while keeping its spatial structure intact. Each of the three filters (right) has a different weighted combination of basis filters, while their  $\sigma$  is varied on the horizontal axis. Optimizing for  $\alpha$  learns filters shape, optimizing for  $\sigma$  learns their size.

In modern CNN architectures [4], [5], the resolution is a hyper-parameter which has to be manually tuned using expert knowledge, by changing the filter sizes or the subsampling layers. For example, the popular ResNeXt [5] for the ImageNet dataset starts with a  $7 \times 7$  px filter, followed by  $3 \times 3$  and  $1 \times 1$  px convolutions where the feature maps are subsampled 6 times. The same network on the CIFAR-10 dataset exclusively uses  $3 \times 3$  px convolutions and the feature maps are subsampled 2 times. Hard-coding the resolution hyper-parameters in the network for different datasets affects the extent of the receptive field, and the specific choices made can be restrictive.

In this paper we propose the N-JetNet which can replace CNN network design choices of filter sizes by learning these. We make use of scale-space theory [6], where the resolution is modeled by the  $\sigma$  parameter of the Gaussian function family and its derivatives. Gaussian derivatives allow a truncated Taylor series, called the N-Jet [7], to model a convolutional filter [8] as a linear combination of Gaussian derivative filters, each weighted by an  $\alpha_i$ . We optimize these  $\alpha$  weights instead of individual weights for each pixel in the filter, as done in a standard CNN. The choice of the basis cannot be avoided. In standard CNNs the choice is implicit: an  $N \times N$  pixel-basis, whose size cannot be optimized, because it has no well-defined derivative to the error. In contrast, in the N-Jet model the basis is a linear combination of Gaussian derivatives where the  $\sigma$  parameter controls both the resolution and the filter size, and has a well-defined derivative to the effective filter and therefore to the error. This formulation allows the network to learn  $\sigma$  and thus the resolution of convolutional networks. We exemplify our approach in Fig. 1.

To avoid confusions, we make the following naming conventions: throughout the paper we refer to ‘resolution’ as the *inner scale* as defined in [1]; ‘size’ as the *outer scale* [1] denoting the number of pixels of a filter or a feature map; and ‘scale’ as the parameter controlling the resolution, which is the standard deviation  $\sigma$  parameter of the Gaussian basis [7]. The scale is different from the size of a filter: one can blur a filter and change its scale without necessarily changing its size. However, they are related as increasing the scale of an object (i.e. blurring) increases its size in the image (i.e. the number of pixels it occupies). Here we tie the filter size to the scale parameter by making it a function of  $\sigma$ .

We make the following contributions. (i) We exploit the multi-scale local jet for automatically learning the scale parameter,  $\sigma$ . (ii) We show both for classification and segmentation that our proposed N-Jet model automatically learns the appropriate input resolution from the data. (iii) We demonstrate that our approach generalizes over network architectures and datasets without deteriorating accuracy for both classification and segmentation.

## II. RELATED WORK

**Multiples scales and sizes in the network.** Size plays an important role in CNNs. The highly successful inception architecture [9] uses two filter sizes per layer. Multiple input sizes can be weighted per layer [10], integrated at the feature map level [11], processed at the same time [12], [13], [14], [15] or even made to compete with each other [16]. To process multiple featuremap sizes, spatial pyramids are used [17], [18], [19], [20], alternatively the best input size and network resolution can be selected over a validation set [21]. Scale-equivariant CNNs can be obtained by applying each filter at multiple sizes [22], or by approximating filters with Gaussian basis combinations [23] where the set of scale parameters is not learned, but fixed. Unlike these works, we do not explicitly process our feature maps over a set of predefined fixed sizes. We learn a single scale parameter per layer from the data.

Downsampling and upsampling can be modeled as a bijective function [24], or made adaptive using reinforcement learning [25] and contextual information at the object boundaries [26]. The optimal size for processing an input giving the maximum classification confidence can be selected among multiple sizes [27], [28], or learned by mimicking the human visual focus [29], or minimizing the entropy over multiple input sizes at inference time [30]. Network architecture search can also be used for learning the resolution, at the cost of increased computations [31]. Alternatively, the scale distribution can be adapted per image using dynamic gates [32], or by using self-attentive memory gates [33]. The atrous [34], [35] or dilated convolutions [36], [37] design fixed versions of larger receptive fields without subsampling the image. These are extended to adaptive dilation factors learned through a sub-network [38]. Rather than only learning the filter size, we learn both the filter shape and the size jointly, by relying on scale-space theory.

**Architectures accommodating subsampling.** A pooling operation groups features together before subsampling. Popular

forms of grouping are average pooling [39], and max pooling [40]. Average pooling tends to perform worse than max pooling [41], [42] which is outperformed by their combination [43], [44]. Other forms include pooling based on ranking [45], spatial pyramid pooling [46], spectral pooling [47] and stochastic pooling [48] and stochastic subsampling [49]. The recent BlurPool [50] avoids aliasing effects when sampling, while fractional pooling [51] subsamples with a factor of  $\sqrt{2}$  instead of 2 which allows larger feature maps to be used in more network layers. All these pooling methods use hard-coded feature map subsampling. Our work differs, as we do not use fixed subsampling or strided convolution: we learn the resolution in each layer.

**Fixed basis approximations.** Resolution in images is aptly modeled by scale-space theory [52], [1], [53]. This is achieved by convolving the image with filters of increasing scale, removing finer details at higher scales. Convolution with a Gaussian filter has the property of not introducing any artifacts [54], [55] and the differential structure of images can be probed with Gaussian derivative filters [56], [6] which form the N-Jet [7]: a complete and stable basis to locally approximate any realistic image. Scale-spaces model images at different resolutions by a continuous one-parameter family of smoothed images, parametrized by the value of  $\sigma$  of the Gaussian filter [1]. In this paper we build on scale-space theory and exploit the differential structure of images to optimize  $\sigma$  and thus learn the resolution.

Various mathematical multi-scale image modeling tools have been used in convolutional networks. The classical work of Simoncelli *et al.* [57] proposes the steerable pyramid, defining a set of wavelets for orientation and scale invariance. Similarly, the seminal Scattering transform [58], [59] and its extensions [60], [61] are based on carefully designed complex wavelet basis filters [62] with pre-defined rotations and scales giving excellent results on uniform datasets such as MNIST and textures. Using the Scattering transform as initialization for the first few layers of a CNN has recently [63], [64] been shown to also lead to good results on more varied datasets. Filters can also be approximated as a linear combination over a set of learned low-rank filter basis [65]. Recent work also starts with a filter basis and use a CNN to learn the filter weights. Examples include a PCA basis [11], circular harmonics [66], Gabors [67], and Gaussian derivatives [8]. In this paper we build on the Gaussian derivative basis [8] because it directly offers the tools of Gaussian scale-space to learn CNN resolution.

**Learning kernel shape.** Current methods investigate inherent properties of CNN filters. Filters that go beyond convolution include non-linear Volterra kernels [68], a learned image adaptive bilateral filter [69] and learned image processing operations [70]. For convolutional CNN filters, Sun *et al.* [71] proposes an asymmetric kernel shape, which simulates hexagonal lattices leading to improved results. The active convolution by Jeon and Kim [72] and the deformable CNNs by Dai *et al.* [73] offer an elegant approach to learn a spatial offset for each filter coefficient leading to flexible filters and improved accuracy. [74] learns continuous filters as functions

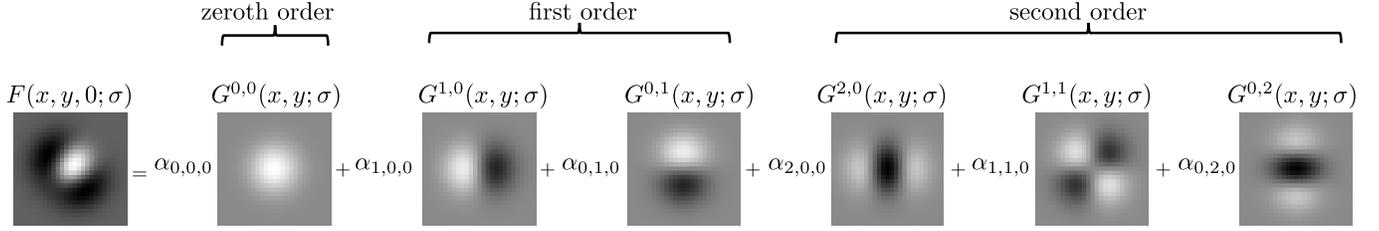


Fig. 2. Representing local image structure with a linear combination of Gaussian basis filters. The patch (left),  $F(x, y, 0; \sigma)$ , is modeled by up to second order Gaussian derivatives, using six  $\alpha$ -coefficients.

over sub-pixel coordinates, allowing learnable resizing of the feature maps. The hierarchical auto-zoom net [75], the scale proposal network [76], and the recurrent scale approximation network [77] explicitly predict the object sizes and adapt the input size accordingly. Our work differs from all these methods because we learn both the filter shape and the size.

Most similar to us, [78], [79] combine free-form filters with learned Gaussian kernels that can adapt the receptive field size. The recent work of Lindeberg *et al.* [80] uses Gaussian derivatives for scale-invariance, however the scales are fixed according to a geometric distribution. Dissimilar to these works we propose to approximate the complete filter using a combination of Gaussian derivatives, while adapting the receptive field size.

### III. LEARNING NETWORK RESOLUTION

#### A. Local image differentials at given scale

Scale-spaces [56], [6], [53] offer a general framework for modeling image structures at various scales. The resolution, or the *inner scale* [1] of an image is modeled by a convolution with a 2D Gaussian. The 1D Gaussian at scale  $\sigma$  is given by  $G(x; \sigma) = \frac{1}{\sigma\sqrt{2\pi}} e^{-\frac{x^2}{2\sigma^2}}$  which is readily extended to 2D as  $G(x, y; \sigma) = G(x; \sigma) G(y; \sigma)$ . The local structure learned in deep networks [1] is linked to the image derivatives. Image pixels are discretely measured, and do not directly offer derivatives. The linearity of the convolution operator allows [81] to take an exact derivative of a slightly smoothed function  $f$  with a Gaussian kernel  $G(\cdot; \sigma)$  with scale  $\sigma$ :

$$\frac{\partial(f(x) * G(x; \sigma))}{\partial x} = \frac{\partial G(x; \sigma)}{\partial x} * f(x), \quad (1)$$

where  $*$  denotes a convolution. This allows taking image derivatives by convolving the image with Gaussian derivatives. Gaussian derivatives in 1D at order  $m$  and scale  $\sigma$  can be defined recursively using the Hermite polynomials [82]:

$$\begin{aligned} G^m(x; \sigma) &= \frac{\partial^m G(x, \sigma)}{\partial x^m} \\ &= \left(\frac{-1}{\sigma\sqrt{2}}\right)^m H_m\left(\frac{x}{\sigma\sqrt{2}}\right) G(x; \sigma), \end{aligned} \quad (2)$$

where  $G(x; \sigma)$  is the Gaussian function and  $H_m(x)$  the  $m$ -th order Hermite polynomial, recursively defined as  $H_i(x) = 2xH_{i-1}(x) - 2(i-1)H_{i-2}(x)$ ;  $H_0(x) = 1$ ;  $H_1(x) = 2x$ . We

define 2D Gaussian derivatives by the product of the partial derivatives on  $x$  and on  $y$ :

$$\begin{aligned} G^{i,j}(x, y; \sigma) &= \frac{\partial^{i+j} G(x, y; \sigma)}{\partial x^i \partial y^j} \\ &= \frac{\partial^i G(x; \sigma)}{\partial x^i} \frac{\partial^j G(y; \sigma)}{\partial y^j}. \end{aligned} \quad (3)$$

#### B. Multi-scale local N-Jet for modeling local image structure

A discrete set of Gaussian derivatives up to  $n^{\text{th}}$  order,  $\{G^{i,j}(x, y; \sigma) \mid 0 \leq i + j \leq n\}$ , can be used in a truncated local Taylor expansion to represent the local scale-space near any given point with increasing accuracy [7]. This allows us to approximate a filter  $F(x)$  around the point  $a$  up to order  $N$  as:

$$F(x) = \sum_{i=0}^N \frac{\frac{\partial^i F(a)}{\partial x^i}}{i!} (x-a)^i + \frac{R(a)}{(N+1)!} (x-a)^{N+1}, \quad (4)$$

where  $R$  is the residual term that corresponds to the approximation error. By absorbing the polynomial coefficients into a value  $\alpha$ , we arrive at a linear combination of Gaussian derivative basis filters which can be used to approximate image filters, as illustrated in Fig. 2. For filter  $F(x, y, c)$  at position  $(x, y)$  and color channel  $c$  the approximation is:

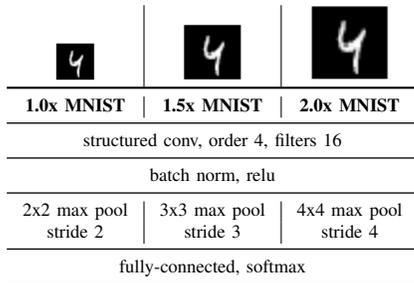
$$\begin{aligned} F(x, y, c; \sigma) &= \sum_{\substack{i+j \leq N \\ 0 \leq i, 0 \leq j}} \alpha_{i,j,c} \frac{\partial^{i+j}}{\partial x^i \partial y^j} G(x, y; \sigma) \\ &+ R(x, y, c; \sigma), \end{aligned} \quad (5)$$

where  $R$  is the residual error, ignored here. Optimizing the  $\alpha$  parameters allows us to switch from learning pixel weights as commonly used in CNNs, to learning the weights of the Gaussian basis filters. We show some examples in Fig. 3 where we optimize the  $\alpha$  parameters of an order-3 RGB Gaussian derivative basis with  $\sigma = 5$  to least squares fit an  $11 \times 11$  px patch. Results show that the fit can approximate well a slightly blurred version of the original patch. Because  $\sigma = 5$  we cannot recover a perfectly sharp faithful copy of the original patch.

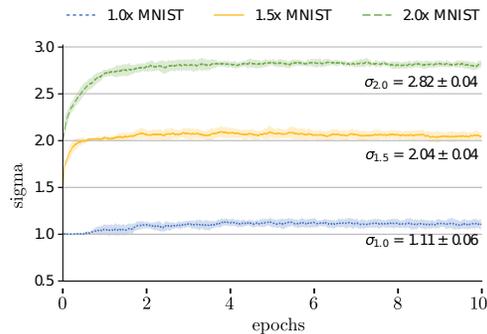
#### C. Learning receptive field size

We have all ingredients to learn the resolution in a convolutional deep neural network (CNN). Resolution is bounded by the size of the CNN filters. We can now dynamically adapt the resolution during training.





(a) Network architecture.



(b) Learned  $\sigma$  on multi-size *MNIST*.

Fig. 6. **Exp 1.2(A):** (a) Toy architecture used for testing whether we can learn the correct data resolution from the inputs. (b) Estimated Gaussian basis scale,  $\sigma$ , on *MNIST* resized 1 $\times$ , 1.5 $\times$  and 2 $\times$ . The estimated  $\sigma$  follows the resizing of the data.

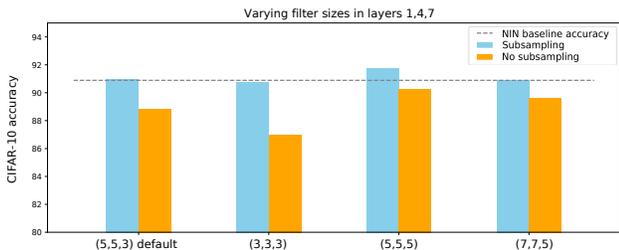


Fig. 7. **Exp 1.1(A):** The impact on accuracy when varying filter sizes and feature map size for the *NIN* baseline on *CIFAR-10*. Smaller filter sizes are more affected by the removal of the subsampling. Setting the resolution hyper-parameters wrong can have a detrimental effect on accuracy.

or strided convolution. For classification models we remove all subsampling operations in the network and add a *safe-subsampling* operation. If the resolution is low (*i.e.*, the  $\sigma$  value is high) then there is no need to keep the feature map at full size, and it can safely be subsampled, to improve memory and speed. For a feature map of size  $s$ , we subsample the feature map to a new size  $\bar{s}$ , where we half its current size as a function of  $\sigma$  as:  $\bar{s} = s (\frac{1}{2})^{\sigma/r}$ , where  $r$  is the *safe-subsampling* hyper-parameter. We apply *safe-subsampling* for all models except for the very deep networks: *Resnet-110* and *EfficientNet*, where it is detrimental by reducing the feature map sizes too much.

**Experimental setup.** In all our N-Jet models we set the order of the Taylor series approximation to 3, unless otherwise specified. When using *safe-subsampling* we remove all pooling layers, and set the stride to 1 in all network layers. For all the models reported, we add a batch normalization layer after the convolutional layers, for robustness, and use the momentum SGD with the momentum set to 0.9. We train for the number of epoch reported in the literature. For the *NIN* baseline model we found the best starting learning rate to be 0.5, while for the *ALLCNN* baseline 0.25. We use the same starting learning rates in our N-Jet models. For our *N-Jet-NIN* model we use an  $L_2$  regularization weight over the Gaussian mixing coefficients,  $\alpha$ , set to 0.01, while for *N-Jet-ALLCNN* we regularize the  $\alpha$ -s with a weight decay of 0.001. When training our *N-Jet-Resnet* models we use an  $L_2$  regularization weight over the Gaussian mixing coefficients,  $\alpha$ , set to 0.0001, and a starting learning rate of 0.1 as indicated in [87]. We evaluate on

relatively small datasets, and therefore we use the lightweight version of *Resnet* where the first block has 16 channels and the last block 64, while for the deeper *Resnet-110* models we use bottleneck blocks with a 4 $\times$  channel expansion. For both the baseline *EfficientNet* and our *N-Jet-EfficientNet* we train the models from scratch, and given the small datasets we use the smallest model B0 [88]. For the *EfficientNet* baseline we use a 0.01 learning rate and a batch size of 32 and we rescale the inputs to 224  $\times$  224 px, since otherwise the model performs poorly, maybe due to the large subsampling. In our *N-Jet-EfficientNet* we keep the input images to their original size. For our deeper models *N-Jet-EfficientNet* and *N-Jet-Resnet-110* we use batches of 16 and a learning rate of 0.001. We release the N-Jet code<sup>1</sup>.

**Experiment 1(A): Validation**

**Experiment 1.1(A): Do resolution hyper-parameters really matter?** We test our assumption that filter sizes and feature map sizes affect accuracy. For this we use the *NIN* baseline trained on *CIFAR-10*. We vary the filter sizes in the layers of the *NIN* which are not 1  $\times$  1 convolutions, and we reset the strides to 1 in all layers, to remove the feature map subsampling. Fig. 7 shows the impact of changing the filter sizes and removing the subsampling. The smaller filter sizes, as in the case when all filter sizes are set to 3, are affected to a greater degree by the removal of the subsampling because they have a smaller receptive field. Selecting the correct filter sizes impacts the overall classification accuracy, and an exhaustive search over all possible filter size combinations is not feasible. This validates the need for learning filter sizes.

**Experiment 1.2(A): Can the image resolution be learned?**

To test resolution learning, we create a toy network architecture depicted in Fig. 6(a). We train the toy architecture on *MNIST* when resizing the images 1 $\times$ , 1.5 $\times$ , and 2 $\times$ . Fig. 6.(b) shows the learned Gaussian basis scale,  $\sigma$ , per setup. The  $\sigma$  values learned for the images resized by 1.5 and 2 do not directly correspond to these values because the operations of sampling and resizing are not commutative: we first discretized the continuous signal into an image and subsequently subsampled it. However, the relative ratio between the learned scales is close:  $(2.0/1.5)\sigma_{1.5} = 2.81 \pm 0.04 \approx \sigma_{2.0} = 2.82 \pm 0.04$ .

<sup>1</sup><https://github.com/SilviaLauraPintea/N-JetNet>

TABLE I

**Exp 2.1(A):** THE EFFECT ON *CIFAR-10* OF VARYING THE FILTER SCALE,  $\sigma$ . HAVING A SUB-OPTIMAL FILTER SCALE  $\sigma$  CAN DECREASE ACCURACY UP TO 3%. THE SAFE-SUBSAMPLING IS AFFECTED SLIGHTLY MORE BY THE CHOICE OF  $\sigma$  THAN THE BASELINE SAMPLING.

Sampling	$\sigma$		
	$\sigma = 0.5$	$\sigma = 1.0$	$\sigma = 2.0$
Baseline	88.76%	90.25%	87.39%
Safe-subsampling	86.29%	89.50%	87.26%

TABLE II

**Exp 2.2(A):** THE IMPORTANCE OF THE HYPER-PARAMETER  $r$  OF THE SAFE-SUBSAMPLING ON THE *CIFAR-10* ACCURACY. THE ACCURACY SLIGHTLY INCREASES AS THE HYPER-PARAMETER  $r$  OF THE SAFE-SUBSAMPLING INCREASES.

	Safe-subsampling hyper-parameter		
	$r = 2.0$	$r = 4.0$	$r = 6.0$
Accuracy	91.59%	91.60%	91.63%
Training time	$\approx 129.91$ min.	$\approx 179.43$ min.	$\approx 193.48$ min.

The learned  $\sigma$  values follow the input resizing, thus the correct filter scales and sizes can be learned from the input.

**Experiment 2(A): Model choices**

**Experiment 2.1(A): Learning sigma.** We test the effect of  $\sigma$  on the performance on the *CIFAR-10* dataset, using the *NIN* backbone. We fix the spatial extent,  $k$ , to 2 and vary sigma in the set  $\{0.5, 1.0, 2.0\}$ . Tab. I shows that a wrong setting of  $\sigma$  can influence the classification accuracy up to 3%. The safe-subsampling setting is affected more by the choice of  $\sigma$  than the baseline subsampling as it relies on the value of  $\sigma$  when deciding how much to subsample the input feature maps. Overall, we note that  $\sigma = 1.0$  achieves the best performance on this setting, therefore we use this value when initializing  $\sigma$  during the learning in our N-Jet models.

**Experiment 2.2(A): Safe-subsampling.** We test the importance of the hyper-parameter  $r$  in the safe-subsampling, with respect to the classification accuracy on *CIFAR-10* using a *NIN* backbone. For this experiment we learn the filter scale  $\sigma$  and set  $k = 2$ . We fix the hyper-parameter  $r$  to one of the values in the set  $\{2.0, 4.0, 6.0\}$ . Tab. II shows the effect on accuracy of different settings of  $r$ . We also show the runtime needed to train the network for different  $r$  settings. As the value of  $r$  increases the accuracy also increases, however also the feature map sizes in the layers of the network increase, which affect the overall computational time. For our subsequent experiments we select  $r = 4.0$  as a trade-off between accuracy and training speed.

**Experiment 3(A): Generalization ability**

**Experiment 3.1(A): Generalization to other datasets.** We compare our *N-Jet-NIN* method with the baseline *NIN* [83]. We test the generalization properties of our method by also reporting scores on two other datasets: *CIFAR-100* and *SVHN*. Tab. III shows the classification results of our *N-Jet-NIN* when compared with the baseline *NIN*. We report mean and standard deviations over 3 runs for our method. We show in Fig. 8 the

TABLE III

**Exp 3.1(A):** DATASET GENERALIZATION. COMPARISON OF OUR *N-Jet-NIN* AND BASELINE *NIN* ON THE *CIFAR-100* AND *SVHN* DATASETS. FOR THE BASELINE MODEL WE ONLY REPORT THE BEST PERFORMANCE WE OBTAIN RERUNNING THE MODELS, WHILE FOR OUR N-JET MODELS WE REPORT MEAN AND STANDARD DEVIATIONS OVER 3 RUNS. WE ACHIEVE COMPARABLE CLASSIFICATION ACCURACY WITH THE BASELINE.

	<i>NIN</i> [83]	<i>N-Jet-NIN</i> (Ours)
<i>SVHN</i>	98.17%	97.55% ( $\pm 0.08$ )
<i>CIFAR-10</i>	90.89%	91.60% ( $\pm 0.08$ )
<i>CIFAR-100</i>	66.14%	68.42% ( $\pm 0.31$ )

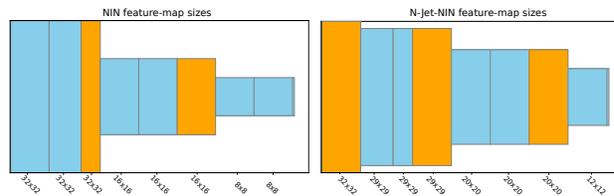


Fig. 8. **Exp 3.1(A):** Dataset generalization. The baseline feature map sizes compared to the learned feature maps sizes by our *N-Jet-NIN* on *CIFAR-10*. We show in orange the subsampling layers. Safe-subsampling dynamically finds the appropriate feature map size.

hard-coded sizes of the baseline feature maps, versus the sizes learned by our *N-Jet-NIN* on *CIFAR-10*. The performance of *N-Jet-NIN* is comparable with the baseline performance, while dynamically learning the appropriate feature map size.

**Experiment 3.2(A): Generalization to other models.** To test the generalization of our N-Jet convolutional layer to different network architectures, we use the *ALLCNN* [86], Resnet [87], and the recent *EfficientNet* [88] backbone network architectures. For *N-Jet-ALLCNN* we use safe-subsampling at every layer, while for *N-Jet-Resnet-32* only at the layers where the original network subsamples. Tab. IV shows the classification accuracy of the baseline models tested by us on *CIFAR-10*, *CIFAR-100*, when compared with our *N-Jet* models. We report mean and standard deviation over 3 repetitions for our models, as well as the number of parameters. Using our proposed N-Jet layers gives similar accuracy to the standard convolutional layers, while avoiding the need to hard-code the filter sizes. For an approximation of order 3 in the N-Jet, there is a small increase in the number of parameters compared to the baseline, except for the *EfficientNet* which uses also kernel sizes larger than  $3 \times 3$  px. When employing larger models – *N-Jet-Resnet-110* and *N-Jet-EfficientNet* – a Gaussian basis combination of order 2 is sufficient to obtain an accuracy comparable to the baseline models, while reducing the number of parameters. In Fig. 9 we show the baseline *ALLCNN* feature map sizes when compared to the *N-Jet-ALLCNN* learned feature map sizes on *CIFAR-10*. The N-Jet model has similar classification accuracy when compared to the *ALLCNN* baseline, while learning at every layer the befitting feature map size. When applied at every layer, the safe-subsampling makes the subsampling continuous and smooth, compared to the baseline.

<sup>1</sup>We use torchinfo (<https://github.com/tyleryep/torchinfo>) to enumerate all the parameters.

TABLE IV

**Exp 3.2(A):** ARCHITECTURE GENERALIZATION. THE CLASSIFICATION ACCURACY ON *CIFAR-10*, *CIFAR-100* WHEN COMPARING THE BASELINE MODELS WITH OUR PROPOSED *N-Jet-ALLCNN*, *N-Jet-Resnet-32*, AND TWO DEEPER MODELS: *N-Jet-Resnet-110* AND *N-Jet-EfficientNet*. FOR THE BASELINE MODELS WE REPORT THE BEST PERFORMANCE, WHILE FOR OUR N-JET MODELS WE REPORT MEAN AND STANDARD DEVIATIONS OVER 3 RUNS. WE ALSO SHOW THE NUMBER OF PARAMETERS FOR EACH MODEL. *N-Jet* NETS OBTAIN COMPARABLE ACCURACY TO THE BASELINES WHILE REDUCING THE NUMBER OF PARAMETERS FOR ORDER 2.

	ALLCNN [86]		N-Jet-ALLCNN (Ours)	
			Order 3	Order 2
# params	0.97 M		1.07 M	0.66 M
<i>CIFAR-10</i>	91.87%		92.48% ( $\pm 0.134$ )	89.91% ( $\pm 0.032$ )
<i>CIFAR-100</i>	67.24%		67.62% ( $\pm 0.863$ )	65.17% ( $\pm 0.228$ )
<hr/>				
	Resnet-32 [87]		N-Jet-Resnet-32 (Ours)	
			Order 3	Order 2
# params	0.47 M		0.52 M	0.31 M
<i>CIFAR-10</i>	92.30%		92.28% ( $\pm 0.260$ )	89.49% ( $\pm 0.304$ )
<i>CIFAR-100</i>	67.89%		67.59% ( $\pm 0.278$ )	65.14% ( $\pm 0.619$ )
<hr/>				
	Resnet-110 [87]		N-Jet-Resnet-110 (Ours)	
			Order 3	Order 2
# params	6.90 M		7.29 M	5.74 M
<i>CIFAR-10</i>	92.83%		93.71% ( $\pm 0.337$ )	93.52% ( $\pm 0.043$ )
<i>CIFAR-100</i>	73.53%		71.73% ( $\pm 0.203$ )	73.66% ( $\pm 0.295$ )
<hr/>				
	EfficientNet [88]		N-Jet-EfficientNet (Ours)	
			Order 3	Order 2
# params	3.60 M		3.51 M	3.48 M
<i>CIFAR-10</i>	92.64%		93.51% ( $\pm 0.110$ )	93.71% ( $\pm 0.029$ )
<i>CIFAR-100</i>	76.19%		75.22% ( $\pm 0.163$ )	76.17% ( $\pm 0.409$ )

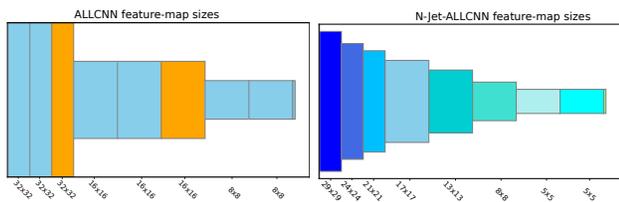


Fig. 9. **Exp 3.2(A):** Architecture generalization. The baseline feature map sizes when compared to the learned N-Jet feature map sizes. We show in different colors the layers at which the size changes. For the standard architecture the subsampling layers are orange. We can find at every layer the appropriate subsampling level.

**Experiment 4(A): Comparison to scale-invariant methods.**

We evaluate on the normal sized  $28 \times 28$  px MNIST and on MNIST resized by a factor of 4 with a size of  $112 \times 112$  px. We compare against a standard CNN with varying filter sizes, and against the Deformable CNN [73], as well as Atrous (dilated) convolutions [34]. We consider 2 and 4-layer toy architectures containing only convolutional layers followed by ReLU activations. Results in table V show that the standard CNN performs well on MNIST, yet results are sensitive to the filter size for  $4 \times$  MNIST. The Deformable CNN [73] is also affected by the change in image size. Our intuition is that the Deformable CNN still relies on the initial  $3 \times 3$  convolutions and optimizing the offsets is difficult under large size changes in the input. For Atrous CNN the dilation factor has to be hard-coded, and we use a dilation factor of 2, as using 4 would imply including prior knowledge. The Atrous performance is also affected by the change in input size. In contrast, our N-Jet model is able to learn the correct

TABLE V

**Exp 4(A):** COMPARISON WITH SCALE-INVARIANT METHODS. WE EVALUATE ON MNIST AND MNIST RESIZED  $4 \times$  OUR N-JET, STANDARD CONVOLUTIONS OF VARYING FILTER SIZES, AS WELL AS USING ATROUS CONVOLUTIONS [34] AND DEFORMABLE CONVOLUTIONS [34]. OUR N-JET PERFORMS WELL ON MNIST  $\times 4$  DESPITE THE INCREASE IN SIZE.

CNN	2-layer architecture	
	MNIST	$4 \times$ MNIST
Standard $3 \times 3$	97.04% ( $\pm 0.22$ )	86.27% ( $\pm 3.44$ )
Standard $5 \times 5$	98.62% ( $\pm 0.08$ )	88.67% ( $\pm 5.97$ )
Standard $9 \times 9$	98.93% ( $\pm 0.08$ )	95.84% ( $\pm 0.76$ )
Standard $11 \times 11$	98.72% ( $\pm 0.20$ )	95.75% ( $\pm 1.50$ )
<hr/>		
Atrous [34]	98.47% ( $\pm 0.25$ )	89.87% ( $\pm 3.67$ )
Deformable [73]	97.54% ( $\pm 0.46$ )	84.09% ( $\pm 0.84$ )
N-Jet (Ours)	99.05% ( $\pm 0.04$ )	98.11% ( $\pm 0.16$ )
<hr/>		
CNN	4-layer architecture	
	MNIST	$4 \times$ MNIST
Standard $3 \times 3$	98.49% ( $\pm 0.20$ )	86.53% ( $\pm 7.15$ )
Standard $5 \times 5$	98.54% ( $\pm 0.37$ )	97.47% ( $\pm 0.41$ )
Standard $9 \times 9$	98.91% ( $\pm 0.18$ )	94.23% ( $\pm 4.01$ )
Standard $11 \times 11$	98.81% ( $\pm 0.20$ )	96.21% ( $\pm 1.40$ )
<hr/>		
Atrous [34]	98.40% ( $\pm 0.38$ )	91.28% ( $\pm 1.88$ )
Deformable [73]	98.91% ( $\pm 0.27$ )	86.45% ( $\pm 1.32$ )
N-Jet (Ours)	99.37% ( $\pm 0.05$ )	98.87% ( $\pm 0.21$ )

resolution and is more accurate.

**B. N-Jet for Image Segmentation**

**Learning the receptive field size for segmentation.** Multi-scale information processing is heavily used in modern segmentation architectures, and seems to be an important performance booster [10]. Here, we focus on two popular mechanisms for multi-scale processing, namely the merging of information at different scales via skip connections in *U-Net* architectures [91], and the pooling of information at different scales via atrous spatial pyramid pooling (ASPP) layers in *DeepLab* architectures [34], [92].

Similar to the classification experiments (Section IV-A), we replace the fixed-size convolutional filters of baseline networks with the N-Jet definition, where we learn the size and scale of the filters in the convolution operations during training.

**Experiment 1(B). Segmentation with U-Net**

**Experiment 1.1(B): Segmentation of multi-scale inputs.** We first evaluate the performance of N-Jet models on a small toy dataset, where each input image is formed by concatenating 4 images (objects) from the *Fashion MNIST* [93] dataset (Fig. 10). Each object is assigned a random scale  $s$ , which determines the factor by which we upsample the original *Fashion MNIST* image, via bilinear interpolation. The scale affects the object sizes — the number of pixels occupied by the object in the image. We construct four different training sets: three where the scale of each object is homogeneous: the discrete variable  $s$  has the probability mass functions  $P(s = 1) = 1$ , or  $P(s = 2) = 1$ , or  $P(s = 4) = 1$ ; and one where the image contains objects on multiple scales:  $s$  has the probability mass function  $P(s) = 0.25$  for values  $s \in \{1, 2, 3, 4\}$ . After rescaling, each object is placed in

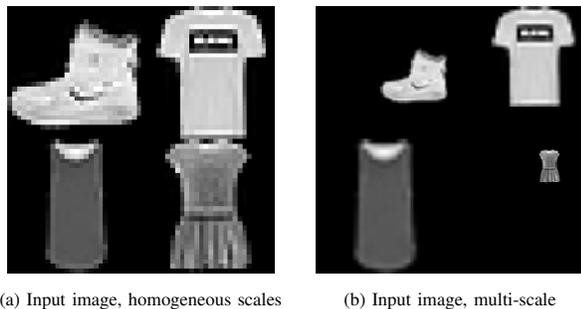


Fig. 10. **Exp 1.1(B)**: Example images from the *multi-scale Fashion MNIST* toy dataset for segmentation. The corresponding segmentation masks are generated by assigning the class label of the corresponding object to pixels whose input grayscale values are above the threshold  $h_{x,y} \geq h_\theta = 0.2$ .

one quadrant of the input image, centered at a uniformly sampled random location. The corresponding ground truth segmentation masks are created by assigning the class label (1...10) of the corresponding object to pixels whose input grayscale values  $h_{x,y}$  are above the threshold  $h_\theta = 0.2$ , by assigning the background label (0) to pixel locations where  $h_{x,y} = 0$ , and by assigning an ignore index to undetermined pixel locations where  $0 < h_{x,y} < h_\theta$ . The ignored pixel locations do not contribute to the loss during training and do not contribute to the accuracy at test time.

Due to the simple nature of the training set, we use a small *U-Net* architecture, where the encoding network has three levels, as opposed to five in the original *U-Net* [91]. This corresponds to two downsampling layers. Each level is composed of two convolutional layers, followed by the ReLU activation layer. The channel dimension is 64 at the first level, and doubles with every downsampling, performed via  $2 \times 2$  max pooling. In the decoding network, we use bilinear upsampling to increase feature map size, and in all convolutional layers we use ‘same’ padding. We train all networks (N-Jet and baseline) for 50 epochs, using the ADAM optimizer [94] and learning rate 0.0001. To accommodate input images of different sizes, and keep with the original *U-Net* implementation, we use a batch size of 1 and no batch normalization, but high momentum  $\beta = (0.9, 0.999)$ . To combat class imbalance, given the especially high frequency of the background class, we weigh the losses with the inverse of class frequencies in the training set. For the N-Jet models, we use filters with basis order 4 and 2. The scale parameter  $\sigma$  is shared between all filters in a convolutional layer.

After training, we evaluate segmentation performance on the validation set using the mean intersection over union (mIoU) over all object classes. Each validation set is constructed in the same way as the corresponding training set, using the *Fashion MNIST* validation images. We find that as we increase the average scale  $s$  of the segmented objects (homogeneous scale case) or the variance of object scales  $s$  (multi-scale case), N-Jet models successfully optimize the scale parameter  $\sigma$  accordingly. This makes N-Jets capable of adapting to different object scales without changing the network architecture, depth or hyper-parameters at all. In contrast, baseline *U-Net* models with fixed filter size cannot adapt their receptive field (RF)

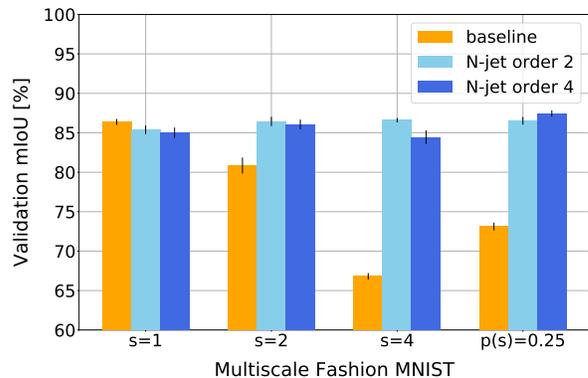


Fig. 11. **Exp 1.1(B)**: mIoU scores averaged over 5 repetitions with different random seeds on *multi-scale Fashion MNIST*. Error bars denote the standard deviation. Validation mIoU decreases dramatically for baseline networks with fixed filter sizes (orange) for larger objects. In comparison, N-Jet models are robust against scale changes, as they can learn the filter size and scale, without changing the architecture or hyper-parameters.

size based on the object scales in the training set, and their segmentation performance decays for larger objects (Fig. 11).

In addition to the robustness of N-Jet networks against changing object scales, we find that N-Jets of only order 2 (where each kernel is defined by only 6 free parameters) is enough to obtain good validation accuracy. In fact, N-Jets of order 4 perform slightly worse for larger  $s$ . This is partly because the reduction of the basis order acts as a regularization via parameter reduction on our simple toy dataset, and partly because *Fashion MNIST* (especially after upscaling) does not contain many high frequency components, which the higher order Gaussian derivatives can capture.

**Experiment 1.2(B): Learning the receptive field size.** While  $\sigma$  optimization is successful for different basis orders, we note that the N-Jet model with basis order 4 has a larger number of free parameters than the baseline *U-Net*. Nevertheless, we observe on our toy dataset that the validation mIoU depends only weakly on the number of parameters, beyond a certain network size. For the multi-scale segmentation task with  $s = 4$ , where the scale of objects are increased by a factor of 4, we find that the receptive field size at the end of the encoding network largely determines the validation mIoU (Fig. 12). To demonstrate this, we vary the number of parameters and the receptive field size at the end of encoding in the baseline *U-Net* models, until we match the N-Jet performance: we increase the kernel size  $k$  from 3 to 4 and 5, and expand the depth of the baseline network by increasing the number of encoding and decoding levels from 3 (10 convolutional layers) to 4 (14 convolutional layers) and 5 (18 convolutional layers). To keep the number of trainable parameters at a reasonable level, for networks with 4 and 5 levels we also decrease the channel width of the layers (by halving or quartering the number of channels in each layer, as given in the legend of Fig. 12).

We find that the N-Jet models can outperform baseline *U-Net* models while using a much smaller number of free parameters, due to  $\sigma$  optimization. In addition, we show that while the receptive field size is a good predictor of performance,

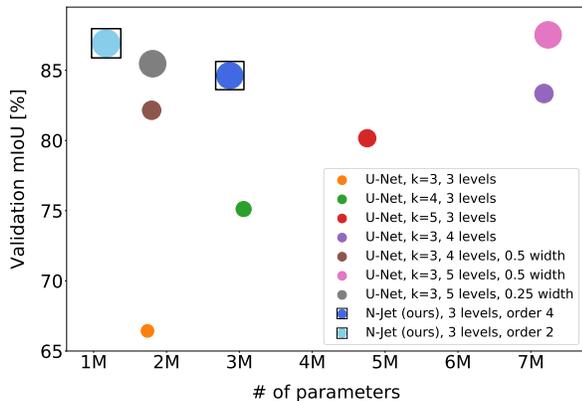


Fig. 12. **Exp 1.2(B):** The effect of the number of parameters (x-axis) and receptive field size (round marker size) on the mIoU scores in the validation set of *multi-scale Fashion MNIST* dataset with  $s = 4$ . We observe that the number of model parameters is a weak predictor of segmentation performance compared to receptive field size which can be learned by N-Jets during training. We note that N-Jet models (in black squares) can perform dramatically better than the baseline model with the same architecture (orange circle), and overall display high mIoU with a low number of parameters.

it cannot be learned during training for the baseline *U-Net*, and would need to be optimized via hyper-parameter scans. This can potentially mean increasing the depth of the network to match the input resolution, which cannot be parallelized. Finally, we observe that slightly better validation mIoU can be obtained by baseline models, with almost 7 times the number of parameters and double the number of layers. We attribute this slight performance boost to the much larger depth, and thus increased number of nonlinearities in the network.

**Experiment 2(B): Segmentation with DeepLabv2.** Next, we consider a more realistic segmentation task on the *Pascal VOC (SBD)* dataset [89], [90] using the *DeepLabv2* architecture [34]. Modern *DeepLab* models take advantage of dilated convolutions to aggregate information from multiple scales in atrous spatial pyramid pooling (ASPP) layers [34], [92]. However, dilated kernels can only be upsampled discretely, based on the dilation rate in units of pixels. In addition, it is typically not possible to determine *a priori* which scales in a dataset contain task-relevant information and the employed dilation rates need to be optimized using excessive hyper-parameter scans. We propose N-Jets as an alternative to optimizing the scales in a continuous way during training, eliminating the need to excessively search for dilation rates for each task.

To that end, we employ the *DeepLabv2* model with a ResNet-101 backbone pretrained on the 20 class subset of the *MS COCO* dataset [95] corresponding to the *Pascal VOC* classes. We retain all the network and training hyper-parameters of the original *DeepLabv2* model and finetune the baseline network with an ASPP output layer on *Pascal VOC* with the batch normalization layers frozen. For the N-Jet network, we replace the 4 convolutional layers of the ASPP layer with different dilation rates with 4 N-Jet layers with independent, learnable scales  $\sigma$  (during finetuning) and we impose weight sharing between the different scales (i.e. same

TABLE VI

**EXP 2(B):** SEGMENTATION MIOU SCORES ON *Pascal VOC* VALIDATION SET ALONG WITH THE NUMBER OF PARAMETERS IN THE ASPP LAYER. N-JET MODELS WITH WEIGHT SHARING, LOWER NUMBER OF PARAMETERS, AND NO HYPER-PARAMETER TUNING ARE AS ACCURATE AS THE BASELINE *DeepLabv2* MODEL WITH WEIGHT SHARING. MOREOVER THE N-JET NETS SUBSTANTIALLY REDUCE THE NUMBER OF PARAMETERS.

Model	mIoU	# parameters
DeepLabv2	76.13	1,548,372
DeepLabv2, weight sharing	74.73	387,093
N-Jet, order 3	75.17	430,101
N-Jet, order 2	74.58	258,069
N-Jet, order 1	74.89	129,045

$\alpha$  values). On top of eliminating the need to manually tune the dilation rates, N-Jet models with weight sharing also have the potential to dramatically reduce the number of parameters in ASPP layers.

We find that *DeepLabv2* with N-Jet output layers indeed allows for parameter reductions (Tab. VI). Using an N-Jet output layer with basis order 3 and weight sharing, we achieve validation mIoU values within 1% of the baseline network, while reducing the number of parameters by nearly a factor of 4. As an additional control, we also train a baseline network with weight sharing within the ASPP layer. Interestingly, we observe that our N-Jet models attain on par or better performance than the weight-tied baseline network, even when we only use a basis order of 1 (each kernel is defined by only 3 free parameters).

It is worth noting that these validation mIoUs are achieved with no hyper-parameter tuning for the N-Jet models, and despite not using N-Jet layers in the pretraining of the *DeepLab* backbone. As it is, we believe N-Jet output layers may be used for multi-scale processing applications with further hyper-parameter tuning of learning rates and regularization parameters, or can be used out of the box to estimate the optimal scale or dilation rates for other architectures.

V. DISCUSSION AND LIMITATIONS

To illustrate the differences between standard convolutional layers and N-Jet convolutional layers, we visualize a set of trained baseline filters compared to the equivalent N-Jet filters (Fig. 13). We find that in many models earlier layers will converge to smaller  $\sigma$  values during training (Fig. 13, top), while deeper layers are prone to learning larger filter sizes (Fig. 13, bottom).

In addition, the strength of the N-Jet representation lies in that it can learn filter sizes, and thus the receptive field size, during training. However, recent work has demonstrated that the effective receptive field (eRF) size of networks can be considerably smaller than what would be expected from the kernel size [96]. We investigate the change in eRF size in our N-Jet models by visualizing the gradients with respect to the input image in our models trained on the multiscale Fashion-MNIST dataset (Fig. 14). We find that, as expected, the eRF size of N-Jet models grows with the size of the training images,

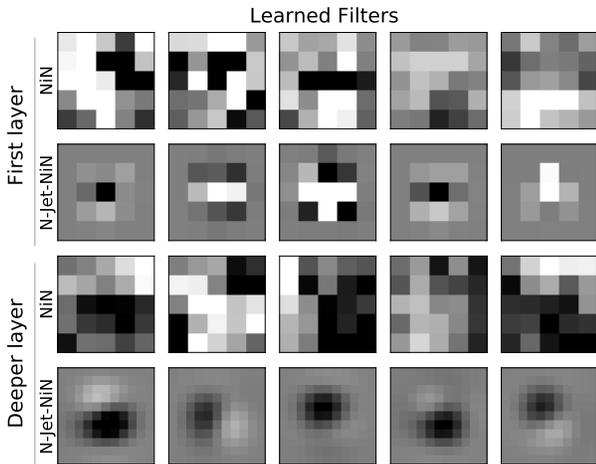


Fig. 13. **Filter visualization.** Learned filters in the *NiN* and *N-Jet-NiN* models. In the first layer, *N-Jet-NiN* converges to a kernel size of  $5 \times 5$ , matching the baseline kernel size (top). In deeper layers, the learned  $\sigma$  value can be much larger, leading to larger kernel sizes: in this example  $11 \times 11$  (bottom). All filter values are centered around 0.

proportionally to the growth of filter sizes. The baseline U-Net model with  $3 \times 3$  kernels cannot learn to adapt its receptive field size during training, its eRF size remains relatively constant as a function of the input image scale.

One of the limitations of our proposed kernels is that they are typically larger than the standard  $3 \times 3$  px, and therefore the convolutions take longer to compute. This comes at no cost in parameters as the size of the N-Jet filters is only affected by the scale parameters,  $\sigma$ . Additionally, computing the Gaussian basis is more expensive because it involves more operations: computing the Hermite polynomials, and obtaining the individual Gaussian basis from these, followed by estimating their linear combinations with the weights  $\alpha$ . For the NiN architecture, our model is  $\approx 2 \times$  slower than the baseline model. As the network depth increases, so do the computations. However, manual architecture search takes a lot longer for finding the appropriate resolution hyper-parameters, because it requires a grid search over all possible filter sizes given a specific network depth and sub-sampling strategy.

## VI. CONCLUSION

We learn the resolution in deep convolutional networks. Learning the resolution frees the network architect from setting resolution related hyper-parameters such as the receptive field size and subsampling layers, which are dataset and network dependent. While we learn the receptive field size and the feature map subsampling for classification, the resolution is also determined by the depth of the network, as each layer increases the resolution linearly. Network depth is not something we learn, and thus we do not learn all resolution hyper-parameters. In addition to hard-coded filter sizes and subsampling layers, current CNN architectures are also designed to share the same filter size in a single layer. Due to computational restraints, our implementation does not make it possible to learn a  $\sigma$  for each filter, rather than per layer. We leave this as potential future work. To conclude, by replacing pixel-weights convolutional

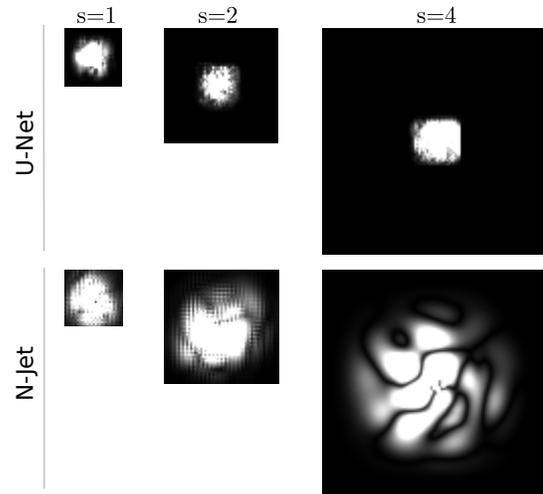


Fig. 14. **Effective receptive field sizes.** Effective receptive field (eRF) size for the N-Jet models trained on the multiscale Fashion-MNIST dataset. Columns ( $s = \{1, 2, 4\}$ ) denote models trained with input images scaled up by a factor of  $s$ . eRF visualizations are obtained by the gradients back-propagated to input pixel space [96]. We find that the eRF size of N-Jet models scale up with the size of the input, while the eRF size of baseline U-Net models stay constant.

layers with our N-Jet convolutional layers we show that we can obtain similar performance as the baseline methods, without tuning the hyper-parameters controlling the resolution.

**Acknowledgements.** This publication is part of the project “Pixel-free deep learning” (with project number 612. 001.805 of the research programme TOP which is financed by the Dutch Research Council (NWO)).

## REFERENCES

- [1] J. J. Koenderink, “The structure of images,” *Biological cybernetics*, vol. 50, no. 5, pp. 363–370, 1984.
- [2] W. Luo, Y. Li, R. Urtasun, and R. Zemel, “Understanding the effective receptive field in deep convolutional neural networks,” in *NIPS*, pp. 4898–4906, 2016.
- [3] M. D. Zeiler and R. Fergus, “Visualizing and understanding convolutional networks,” in *ECCV*, pp. 818–833, Springer, 2014.
- [4] G. Huang, Z. Liu, K. Q. Weinberger, and L. van der Maaten, “Densely connected convolutional networks,” in *CVPR*, 2017.
- [5] S. Xie, R. Girshick, P. Dollár, Z. Tu, and K. He, “Aggregated residual transf o deep neural networks,” in *CVPR*, pp. 5987–5995, IEEE, 2017.
- [6] T. Lindeberg, *Scale-space theory in computer vision*, vol. 256. Springer Science & Business Media, 2013.
- [7] L. Florack, B. T. H. Romeny, M. Viergever, and J. Koenderink, “The gaussian scale-space paradigm and the multiscale local jet,” *IJCV*, vol. 18, no. 1, pp. 61–75, 1996.
- [8] J.-H. Jacobsen, J. C. van Gemert, Z. Lou, and A. W. M. Smeulders, “Structured receptive fields in CNNs,” in *CVPR*, 2016.
- [9] C. Szegedy, W. Liu, Y. Jia, P. Sermanet, S. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, and A. Rabinovich, “Going deeper with convolutions,” in *CVPR*, 2015.
- [10] L.-C. Chen, Y. Yang, J. Wang, W. Xu, and A. L. Yuille, “Attention to scale: Scale-aware semantic image segmentation,” in *CVPR*, 2016.
- [11] G. Ghiasi and C. C. Fowlkes, “Laplacian pyramid reconstruction and refinement for semantic segmentation,” in *ECCV*, 2016.
- [12] S. Bai, V. Koltun, and J. Z. Kolter, “Multiscale deep equilibrium models,” *ICLR workshop*, 2018.
- [13] A. Kanazawa, A. Sharma, and D. Jacobs, “Locally scale-invariant convolutional neural networks,” *NIPS workshop*, 2014.
- [14] T.-W. Ke, M. Maire, and X. Y. Stella, “Multigrid neural architectures,” in *CVPR*, 2017.
- [15] Y. Li, Z. Kuang, Y. Chen, and W. Zhang, “Data-driven neuron allocation for scale aggregation networks,” in *CVPR*, 2019.

- [16] Z. Liao and G. Carneiro, "A deep convolutional neural network module that promotes competition of multiple-size filters," *Pattern Recognition*, 2017.
- [17] E. Agustsson, D. Minnen, N. Johnston, J. Balle, S. J. Hwang, and G. Toderici, "Scale-space flow for end-to-end optimized video compression," in *CVPR*, pp. 8503–8512, 2020.
- [18] T.-Y. Lin, P. Dollár, R. Girshick, K. He, B. Hariharan, and S. Belongie, "Feature pyramid networks for object detection," in *CVPR*, pp. 2117–2125, 2017.
- [19] Z. Liu, G. Gao, L. Sun, and L. Fang, "Ipg-net: Image pyramid guidance network for small object detection," in *CVPR Workshops*, pp. 1026–1027, 2020.
- [20] X. Wang, S. Zhang, Z. Yu, L. Feng, and W. Zhang, "Scale-equalizing pyramid convolution for object detection," in *CVPR*, pp. 13359–13368, 2020.
- [21] T. Yang, S. Zhu, C. Chen, S. Yan, M. Zhang, and A. Willis, "Mutualnet: Adaptive convnet via mutual learning from network width and resolution," in *ECCV*, 2020.
- [22] D. Marcos, B. Kellenberger, S. Lobry, and D. Tuia, "Scale equivariance in cnns with vector fields," *ICML workshop*, 2018.
- [23] I. Sosnovik, M. Szmaja, and A. Smeulders, "Scale-equivariant steerable networks," *ICLR*, 2020.
- [24] M. Xiao, S. Zheng, C. Liu, Y. Wang, D. He, G. Ke, J. Bian, Z. Lin, and T.-Y. Liu, "Invertible image rescaling," *ECCV*, 2020.
- [25] B. Uzkent and S. Ermon, "Learning when and where to zoom with deep reinforcement learning," in *CVPR*, pp. 12345–12354, 2020.
- [26] D. Marin, Z. He, P. Vajda, P. Chatterjee, S. Tsai, F. Yang, and Y. Boykov, "Efficient segmentation: Learning downsampling near semantic boundaries," in *ICCV*, pp. 2131–2141, 2019.
- [27] Y. Li, D. M. Tax, and M. Loog, "Scale selection for supervised image segmentation," *Image and Vision Computing*, vol. 30, no. 12, pp. 991–1003, 2012.
- [28] Y. Li, D. M. Tax, and M. Loog, "Supervised scale-invariant segmentation (and detection)," in *International Conference on Scale Space and Variational Methods in Computer Vision*, pp. 350–361, 2011.
- [29] G. A. Sigurdsson, A. Gupta, C. Schmid, and K. Alahari, "Beyond the camera: Neural networks in world coordinates," *CoRR*, 2020.
- [30] D. Wang, E. Shelhamer, B. Olshausen, and T. Darrell, "Dynamic scale inference by entropy minimization," *CoRR*, 2019.
- [31] C. Liu, L.-C. Chen, F. Schroff, H. Adam, W. Hua, A. L. Yuille, and L. Fei-Fei, "Auto-deeplab: Hierarchical neural architecture search for semantic image segmentation," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 82–92, 2019.
- [32] Y. Li, L. Song, Y. Chen, Z. Li, X. Zhang, X. Wang, and J. Sun, "Learning dynamic routing for semantic segmentation," in *CVPR*, pp. 8553–8562, 2020.
- [33] G.-J. Qi, "Hierarchically gated deep networks for semantic segmentation," in *CVPR*, pp. 2267–2275, 2016.
- [34] L.-C. Chen, G. Papandreou, I. Kokkinos, K. Murphy, and A. L. Yuille, "Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected CRFs," *TPAMI*, 2017.
- [35] G. Papandreou, I. Kokkinos, and P.-A. Savalle, "Modeling local and global deformations in deep learning: Epitomic convolution, multiple instance learning, and sliding window detection," in *CVPR*, 2015.
- [36] F. Yu and V. Koltun, "Multi-scale context aggregation by dilated convolutions," in *ICLR*, 2016.
- [37] F. Yu, V. Koltun, and T. Funkhouser, "Dilated residual networks," in *CVPR*, 2017.
- [38] M. Zhang, J. Zhao, X. Li, L. Zhang, and Q. Li, "Ascnet: Adaptive-scale convolutional neural networks for multi-scale feature learning," in *International Symposium on Biomedical Imaging*, pp. 144–148, 2020.
- [39] K. Fukushima, "A neural network model for selective attention in visual pattern recognition," *Biological Cybernetics*, vol. 55, no. 1, pp. 5–15, 1986.
- [40] T. Serre, L. Wolf, and T. Poggio, "Object recognition with features inspired by visual cortex," in *CVPR*, vol. 2, pp. 994–1000, IEEE, 2005.
- [41] Y.-L. Boureau, F. Bach, Y. LeCun, and J. Ponce, "Learning mid-level features for recognition," in *CVPR*, pp. 2559–2566, IEEE, 2010.
- [42] D. Scherer, A. Müller, and S. Behnke, "Evaluation of pooling operations in convolutional architectures for object recognition," *Artificial Neural Networks-ICANN 2010*, pp. 92–101, 2010.
- [43] Y.-L. Boureau, J. Ponce, and Y. LeCun, "A theoretical analysis of feature pooling in visual recognition," in *ICML*, 2010.
- [44] C.-Y. Lee, P. Gallagher, and Z. Tu, "Generalizing pooling functions in cnns: Mixed, gated, and tree," *TPAMI*, 2017.
- [45] Z. Shi, Y. Ye, and Y. Wu, "Rank-based pooling for deep convolutional neural networks," *Neural Networks*, vol. 83, pp. 21–31, 2016.
- [46] K. He, X. Zhang, S. Ren, and J. Sun, "Spatial pyramid pooling in deep convolutional networks for visual recognition," in *ECCV*, pp. 346–361, Springer, 2014.
- [47] O. Rippel, J. Snoek, and R. P. Adams, "Spectral representations for convolutional neural networks," in *NIPS*, pp. 2449–2457, 2015.
- [48] M. Zeiler and R. Fergus, "Stochastic pooling for regularization of deep convolutional neural networks," in *ICLR*, 2013.
- [49] S. Zhai, H. Wu, A. Kumar, Y. Cheng, Y. Lu, Z. Zhang, and R. Feris, "S3pool: Pooling with stochastic spatial sampling," in *CVPR*, 2017.
- [50] R. Zhang, "Making convolutional networks shift-invariant again," *ICML*, 2019.
- [51] B. Graham, "Fractional max-pooling," *CoRR*, 2014.
- [52] T. Iijima, "Basic theory of pattern observation," *Technical Group on Automata and Automatic Control*, pp. 3–32, 1959.
- [53] A. P. Witkin, "Scale-space filtering," in *International Joint Conference on Artificial Intelligence*, 1983.
- [54] J. Babaud, A. P. Witkin, M. Baudin, and R. O. Duda, "Uniqueness of the gaussian kernel for scale-space filtering," *TPAMI*, vol. 1, pp. 26–33, 1986.
- [55] R. Duits, L. Florack, J. De Graaf, and B. ter Haar Romeny, "On the axioms of scale space theory," *Journal of Mathematical Imaging and Vision*, vol. 20, no. 3, pp. 267–298, 2004.
- [56] L. M. Florack, B. M. ter Haar Romeny, J. J. Koenderink, and M. A. Viergever, "Scale and the differential structure of images," *Image and Vision Computing*, vol. 10, no. 6, pp. 376–388, 1992.
- [57] E. P. Simoncelli, W. T. Freeman, E. H. Adelson, and D. J. Heeger, "Shiftable multiscale transforms," *IEEE transactions on Information Theory*, vol. 38, no. 2, pp. 587–607, 1992.
- [58] J. Bruna and S. Mallat, "Invariant scattering convolution networks," *TPAMI*, vol. 35, no. 8, pp. 1872–1886, 2013.
- [59] S. Mallat, "Group invariant scattering," *Communications on Pure and Applied Mathematics*, vol. 65, no. 10, pp. 1331–1398, 2012.
- [60] F. Cotter and N. Kingsbury, "Visualizing and improving scattering networks," in *MLSP*, 2017.
- [61] L. Sifre and S. Mallat, "Rotation, scaling and deformation invariant scattering for texture discrimination," in *CVPR*, 2013.
- [62] S. Mallat, *A wavelet tour of signal processing*. Academic press, 1999.
- [63] E. Oyallon, E. Belilovsky, and S. Zagoruyko, "Scaling the scattering transform: Deep hybrid networks," in *ICCV*, 2017.
- [64] A. Singh and N. Kingsbury, "Efficient convolutional network learning using parametric log based dual-tree wavelet scatternet," in *CVPR workshop*, 2017.
- [65] Y. Li, S. Gu, L. V. Gool, and R. Timofte, "Learning filter basis for convolutional neural network compression," in *ICCV*, pp. 5623–5632, 2019.
- [66] D. E. Worrall, S. J. Garbin, D. Turmukhambetov, and G. J. Brostow, "Harmonic networks: Deep translation and rotation equivariance," in *CVPR*, July 2017.
- [67] S. Luan, B. Zhang, C. Chen, X. Cao, Q. Ye, J. Han, and J. Liu, "Gabor Convolutional Networks," *CoRR*, 2017.
- [68] G. Zoumpourlis, A. Doumanoglou, N. Vretos, and P. Daras, "Non-linear convolution filters for cnn-based learning," in *ICCV*, 2017.
- [69] V. Jampani, M. Kiefel, and P. V. Gehler, "Learning sparse high dimensional filters: Image filtering, dense crfs and bilateral neural networks," in *CVPR*, 2016.
- [70] Q. Chen, J. Xu, and V. Koltun, "Fast image processing with fully-convolutional networks," in *CVPR*, 2017.
- [71] Z. Sun, M. Ozay, and T. Okatani, "Design of kernels in convolutional neural networks for image classification," in *ECCV*, 2016.
- [72] Y. Jeon and J. Kim, "Active convolution: Learning the shape of convolution for image classification," in *CVPR*, 2017.
- [73] J. Dai, H. Qi, Y. Xiong, Y. Li, G. Zhang, H. Hu, and Y. Wei, "Deformable convolutional networks," in *ICCV*, 2017.
- [74] A. Shocher, B. Feinstein, N. Haim, and M. Irani, "From discrete to continuous convolution layers," *CoRR*, 2020.
- [75] F. Xia, P. Wang, L.-C. Chen, and A. L. Yuille, "Zoom better to see clearer: Human and object parsing with hierarchical auto-zoom net," in *ECCV*, pp. 648–663, Springer, 2016.
- [76] Z. Hao, Y. Liu, H. Qin, J. Yan, X. Li, and X. Hu, "Scale-aware face detection," in *CVPR*, pp. 6186–6195, 2017.
- [77] Y. Liu, H. Li, J. Yan, F. Wei, X. Wang, and X. Tang, "Recurrent scale approximation for object detection in cnn," in *ICCV*, 2017.
- [78] E. Shelhamer, D. Wang, and T. Darrell, "Blurring the line between structure and learning to optimize and adapt receptive fields," *CoRR*, 2019.

- [79] Z. Xiong, Y. Yuan, N. Guo, and Q. Wang, "Variational context-deformable convnets for indoor scene parsing," in *CVPR*, pp. 3992–4002, 2020.
- [80] T. Lindeberg, "Scale-covariant and scale-invariant gaussian derivative networks," 2020.
- [81] R. A. Young, "The gaussian derivative model for spatial vision: I. retinal mechanisms," *Spatial vision*, vol. 2, no. 4, pp. 273–293, 1987.
- [82] J.-B. Martens, "The hermite transform-theory," *IEEE Transactions on Acoustics, Speech, and Signal Processing*, vol. 38, no. 9, pp. 1595–1606, 1990.
- [83] M. Lin, Q. Chen, and S. Yan, "Network in network," *ICLR*, 2014.
- [84] A. Krizhevsky and G. Hinton, "Learning multiple layers of features from tiny images," *Tech Report*, 2009.
- [85] Y. Netzer, T. Wang, A. C. A. Bissacco, B. Wu, and A. Y. Ng, "Reading digits in natural images with unsupervised feature learning," *NIPS workshop*, vol. 2011, no. 2, p. 5, 2011.
- [86] J. Springenberg, A. Dosovitskiy, T. Brox, and M. Riedmiller, "Striving for simplicity: The all convolutional net," in *ICLR workshop*, 2015.
- [87] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *CVPR*, pp. 770–778, 2016.
- [88] M. Tan and Q. Le, "Efficientnet: Rethinking model scaling for convolutional neural networks," in *International Conference on Machine Learning*, pp. 6105–6114, PMLR, 2019.
- [89] B. Hariharan, P. Arbeláez, L. Bourdev, S. Maji, and J. Malik, "Semantic contours from inverse detectors," in *2011 International Conference on Computer Vision*, pp. 991–998, IEEE, 2011.
- [90] M. Everingham, S. A. Eslami, L. Van Gool, C. K. Williams, J. Winn, and A. Zisserman, "The pascal visual object classes challenge: A retrospective," *International journal of computer vision*, vol. 111, no. 1, pp. 98–136, 2015.
- [91] O. Ronneberger, P. Fischer, and T. Brox, "U-net: Convolutional networks for biomedical image segmentation," in *International Conference on Medical image computing and computer-assisted intervention*, pp. 234–241, Springer, 2015.
- [92] L.-C. Chen, G. Papandreou, F. Schroff, and H. Adam, "Rethinking atrous convolution for semantic image segmentation," *arXiv preprint arXiv:1706.05587*, 2017.
- [93] H. Xiao, K. Rasul, and R. Vollgraf, "Fashion-mnist: a novel image dataset for benchmarking machine learning algorithms," *arXiv preprint arXiv:1708.07747*, 2017.
- [94] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," in *3rd International Conference on Learning Representations, ICLR*, 2015.
- [95] T. Lin, M. Maire, S. J. Belongie, J. Hays, P. Perona, D. Ramanan, P. Dollár, and C. L. Zitnick, "Microsoft COCO: common objects in context," in *ECCV*, vol. 8693, pp. 740–755, 2014.
- [96] W. Luo, Y. Li, R. Urtasun, and R. Zemel, "Understanding the effective receptive field in deep convolutional neural networks," in *30th International Conference on Neural Information Processing Systems (NeurIPS)*, pp. 4905–4913, 2016.