



Scale-Covariant and Scale-Invariant Gaussian Derivative Networks

Tony Lindeberg¹

Received: 19 March 2021 / Accepted: 20 October 2021 / Published online: 23 December 2021
© The Author(s) 2021

Abstract

This paper presents a hybrid approach between scale-space theory and deep learning, where a deep learning architecture is constructed by coupling parameterized scale-space operations in cascade. By sharing the learnt parameters between multiple scale channels, and by using the transformation properties of the scale-space primitives under scaling transformations, the resulting network becomes provably scale covariant. By in addition performing max pooling over the multiple scale channels, or other permutation-invariant pooling over scales, a resulting network architecture for image classification also becomes provably scale invariant. We investigate the performance of such networks on the MNIST Large Scale dataset, which contains rescaled images from the original MNIST dataset over a factor of 4 concerning training data and over a factor of 16 concerning testing data. It is demonstrated that the resulting approach allows for scale generalization, enabling good performance for classifying patterns at scales not spanned by the training data.

Keywords Scale covariance · Scale invariance · Scale generalisation · Scale selection · Gaussian derivative · Scale space · Deep learning

1 Introduction

Variations in scale constitute a substantial source of variability in real-world images, because of objects having different size in the world and being at different distances to the camera.

A problem with traditional deep networks, however, is that they are not covariant with respect to scaling transformations in the image domain. In deep networks, nonlinearities are performed relative to the current grid spacing, which implies that the deep network does not commute with scaling transformations. Because of this lack of ability to handle scaling variations in the image domain, the performance of deep networks may be very poor when subject to testing data at scales that are not spanned by the training data.

One way of achieving scale covariance in a brute force manner is by applying the same deep net to multiple rescaled

copies of the input image. Such an approach is developed and investigated in [1]. When working with such a scale-channel network it may, however, be harder to combine information between the different scale channels, unless the multi-resolution representations at different levels of resolution are also resampled to a common reference frame when information from different scale levels is to be combined.

Another approach to achieve scale covariance is by applying multiple rescaled nonlinear filters to the same image. For such an architecture, it will specifically be easier to combine information from multiple scale levels, since the image data at all scales have the same resolution.

For the primitive discrete filters in a regular deep network, it is, however, not obvious how to rescale the primitive components, in terms of *e.g.*, local 3×3 or 5×5 filters or max pooling over 2×2 neighbourhoods in a sufficiently accurate manner over continuous variations of spatial scaling factors. For this reason, it would be preferable to have a continuous model of the image filters, which are then combined together into suitable deep architectures, since the support regions of the continuous filters could then be rescaled in a continuous manner. Specifically, if we choose these filters as scale-space filters, which are designed to handle scaling transformations in the image domain, we have the potential of constructing

The support from the Swedish Research Council (contract 2018-03586) is gratefully acknowledged.

✉ Tony Lindeberg
tony@kth.se

¹ Computational Brain Science Lab, Division of Computational Science and Technology, KTH Royal Institute of Technology, 100 44 Stockholm, Sweden

a rich family of hierarchical networks based on scale-space operations that are provably scale covariant [2].

The subject of this article is to develop and experimentally investigate one such hybrid approach between scale-space theory and deep learning. The idea that we shall follow is to define the layers in a deep architecture from scale-space operations, and then use the closed-form transformation properties of the scale-space primitives under scaling transformations to achieve provable scale covariance and scale invariance of the resulting continuous deep network. Specifically, we will demonstrate that this will give the deep network the ability to generalize to previously unseen scales that are not spanned by the training data. This generalization ability implies that training can be performed at some scale(s) and testing at other scales, within some predefined scale range, with maintained high performance over substantial scaling variations.

Technically, we will experimentally explore this idea for one specific type of architecture, where the layers are parameterized linear combinations of Gaussian derivatives up to order two. With such a parameterization of the filters in the deep network, we also obtain a compact parameterization of the degrees of freedom in the network, with of the order of 16k or 38k parameters for the networks used in the experiments in this paper, which may be advantageous in situations when only smaller sets of training data are available. The overall principle for obtaining scale covariance and scale invariance is, however, much more general and applies to much wider classes of possible ways of defining layers from scale-space operations.

1.1 Structure of this Article

This paper is structured as follows: Section 2 begins with an overview of related work, with emphasis on approaches for handling scale variations in classical computer vision and deep networks. Section 3 introduces the technical material with a conceptual overview description of how the notions of scale covariance and scale invariance enable scale generalization, *i.e.*, the ability to perform testing at scales not spanned by the training data. Section 4 defines the notion of Gaussian derivative networks, gives their conceptual motivation and proves their scale covariance and scale invariance properties. Section 5 describes the result of applying a single-scale-channel Gaussian derivative network to the regular MNIST dataset. Section 6 describes the result of applying multi-scale-channel Gaussian derivative networks to the MNIST Large Scale dataset, with emphasis on scale generalization properties and scale selection properties. Finally, Sect. 7 concludes with a summary and discussion.

1.2 Relations to Previous Contribution

This paper is an extended version of a paper presented at the SSVM 2021 conference [3] and with substantial additions concerning:

- a wider overview of related work (Sect. 2),
- a conceptual explanation about the potential advantages of scale covariance and scale invariance for deep networks, specifically with regard to how these notions enable scale generalization (Sect. 3),
- more detailed mathematical definitions regarding the foundations of Gaussian derivative networks (Sect. 4.2) as well as more detailed proofs regarding their scale covariance (Sect. 4.3) and scale invariance (Sect. 4.4) properties,
- a more detailed treatment of the scale selection properties of the resulting scale channel networks (Sect. 6.1) as well as a discussion about issues to consider when training multi-scale-channel networks (Sect. 6.2).

In relation to the SSVM 2021 paper, this paper therefore (1) gives a more general treatment about the importance of scale generalization, (2) gives a more detailed treatment about the theory of the presented Gaussian derivative networks that could not be included in the conference paper because of the space limitations, (3) describes scale selection properties of the resulting scale channel networks and (4) gives overall better descriptions of the subjects treated in the paper, including (5) more extensive references to related literature.

2 Relations to Previous Work

In classical computer vision, it has been demonstrated that scale-space theory constitutes a powerful paradigm for constructing scale-covariant and scale-invariant feature detectors and making visual operations robust to scaling transformations [4–13]. In the area of deep learning, a corresponding framework for handling general scaling transformations has so far not been as well established.

Concerning the relationship between deep networks and scale, several researchers have observed robustness problems of deep networks under scaling variations [14,15]. There have been some approaches developed to aim at scale-invariant convolutional neural networks (CNNs) [16–19]. These approaches have, however, not been experimentally evaluated on the task of generalizing to scales not present in the training data [17,18], or only over a very narrow scale range [16,19].

For studying scale invariance in a more general setting, we argue that it is essential to experimentally verify the scale-invariant properties over sufficiently wide scale ranges. For scaling factors moderately near one, a network could in principle learn to handle scaling transformations by mere training, *e.g.*, by data augmentation of the training data. For wider ranges of scaling factors, that will, however, either not be possible or at least very inefficient, if the same network, with a uniform internal architecture, is to represent both a large number of very fine-scale and a large number of very coarse-scale image structures. Currently, however, there is a lack of datasets that both cover sufficiently wide scale ranges and contain sufficient amounts of data for training deep networks. For this reason, we have in a companion work created the MNIST Large Scale dataset [20,21], which covers scaling factors over a range of 8, and over which we will evaluate the proposed methodology, compared to previously reported experimental work that cover a scale range of the order of a factor of 3 [16,19,22].

Provably scale-covariant or scale-equivariant¹ networks that incorporate the transformation properties of the network under scaling transformations, or approximations thereof, have been recently developed in [1–3,22–25]. In principle, there are two types of approaches to achieve scale covariance by expanding the image data over the scale dimension: (1) either by applying multiple rescaled filters to each input image or (2) by rescaling each underlying image over multiple scaling factors and applying the same deep network to all these rescaled images. In the continuous case, these two dual approaches are computationally equivalent, whereas they may differ in practice depending upon how the discretization is done and depending upon how the computational components are integrated into a composed network architecture.

Still, all the components of a scale-covariant or scale-invariant network, also the training stage and the handling of boundary effects in the scale direction, need to support true scale invariance or a sufficiently good approximation thereof, and need to be experimentally verified on scale generalization tasks, in order to support true scale generalization.²

A main purpose of this article is to demonstrate how provably scale-covariant and scale-invariant deep networks can be constructed using straightforward extensions of estab-

lished concepts in classical scale-space theory, and how the resulting networks enable very good scale generalization. Extensions of ideas and computational mechanisms in classical scale-space theory to handling scaling variations in deep networks for object recognition have also been recently explored in [26]. Using a set of parallel scale channels, with shared weights between the scale channels, as used in other scale-covariant networks as well as in this work, has also been recently explored for object recognition in [27].

In contrast to some other work, such as [23], the scale-covariant and scale-invariant properties of our networks are not restricted to scaling transformations that correspond to integer scaling factors. Instead, true scale covariance and scale invariance hold for the scaling factors that correspond to ratios of the scale values of the associated scale channels in the network. For values in between, the results will instead be approximations, whose accuracy will also be determined by the network architecture and the training method.

In the implementation underlying this work, we sample the space of scale channels by multiples of $\sqrt{2}$, which has also been previously demonstrated to be a good choice for many classical scale-space algorithms, as a trade-off between computational accuracy, as improved by decreasing this ratio, and computational efficiency, as improved by increasing this ratio. In complementary companion work for the dual approach of constructing scale-covariant deep networks by expansion over rescaled input images for multiple scaling factors [1], this choice of sampling density has also been shown to be a good trade-off in the sense that the accuracy of the network is improved substantially by decreasing the scale sampling factor from a factor of 2 to $\sqrt{2}$, however, very marginally by decreasing the scale sampling factor further again to $\sqrt[4]{2}$.

A conceptual simplification that we shall make, for simplicity of implementation and experimentation only, is that all the filters between the layers in the network operate only on information within the same scale channel. A natural generalization to consider would be to also allow the filters to also access information from neighbouring scale channels,³ as used in several classical scale-space methods [28–32] and also used within the notion of group convolution in the deep learning literature [33], which would then enable interactions between image information at different scales. A possible problem with allowing the filters to extend to neighbouring scale channels, however, is that the boundary effects in the scale direction, caused by a limited number of scale channels, may then become more problematic, which may affect

¹ In the deep learning literature, the terminology “scale equivariance” has become common for what is called “scale covariance” in scale-space theory. In this paper, we use the terminology “scale covariance” to keep consistency with the earlier scale-space literature [12].

² For example, in a companion work [1], we noticed that for an alternative sliding window approach studied in that work, full scale generalization is not achieved, because the support regions of the receptive fields in the training stage do not handle all the sizes of the input in a uniform manner, thereby negatively affecting the scale generalization properties, although the overall network architecture would otherwise support true scale invariance.

³ With regard to the Gaussian derivative networks described in this article, the transformation between adjacent layers in Equation (7), should then be expanded with a sum of corresponding contributions from a set of neighbouring scale channels, in order to allow for interactions between image information at different scales.

the scale generalization properties. For this reason, we restrict ourselves to within-scale-channel processing only, in our first implementation, although the conceptual formulation would with straightforward extensions also apply to scale interactions and group convolutions.

Spatial transformer networks have been proposed as a general approach for handling image transformations in CNNs [34,35]. Plain transformation of the feature maps in a CNN by a spatial transformer will, however, in general, not deliver a correct transformation and will therefore not support truly invariant recognition, although spatial transformer networks that instead operate by transforming the input will allow for true transformation-invariant properties, at the cost of more elaborate network architectures with associated possible training problems [36,37].

Concerning other deep learning approaches that somehow handle or are related to the notion of scale, deep networks have been applied to the multiple layers in an image pyramid [38–43], or using other multi-channel approaches where the input image is rescaled to different resolutions, possibly combined with interactions or pooling between the layers [44–46]. Variations or extensions of this approach include scale-dependent pooling [47], using sets of subnetworks in a multi-scale fashion [48], dilated convolutions [49–51], scale-adaptive convolutions [52] or adding additional branches of down-samplings and/or up-samplings in each layer of the network [53,54]. The aims of these types of networks have, however, not been to achieve scale invariance in a strict mathematical sense, and they could rather be coarsely seen as different types of approaches to enable multi-scale processing of the image data in deep networks.

In classical computer vision, receptive field models in terms of Gaussian derivatives have been demonstrated to constitute a powerful approach to handle image structures at multiple scales, where the choice of Gaussian derivatives as primitives in the first layer of visual processing can be motivated by mathematical necessity results [55–66].

In deep networks, mathematically defined models of receptive fields have been used in different ways. Scattering networks have been proposed based on Morlet wavelets [67–69]. Gaussian derivative kernels have been used as structured receptive fields in CNNs [70]. Gabor functions have been proposed as primitives for modulating learned filters [71]. Affine Gaussian kernels have been used to compose free-form filters to adapt the receptive field size and shape to the image data [72].

As an alternative approach to handle scaling transformations in deep networks, the image data has been spatially warped by a log-polar transformation prior to the image filtering steps [73,74], implying that the scaling transformation is mapped to a mere translation in the log-polar domain. Such a log-polar transformation does, however, violate translational

covariance over the original spatial domain, as otherwise obeyed by a regular CNN applied to the original input image.


Approaches to handling image transformations in deep networks have also been developed based on formalism from group theory [33,75–77]. A general framework for handling basic types of natural image transformations in terms of spatial scaling transformations, spatial affine transformations, Galilean transformations and temporal scaling transformations in the first layer of visual processing based on linear receptive fields and with relations to biological vision has been presented in [78], based on generalized axiomatic scale-space theory [12].

The idea of modelling layers in neural networks as continuous functions instead of discrete filters has also been advocated in [79–82]. The idea of reducing the number of parameters of deep networks by a compact parameterization of continuous filter shapes has also been recently explored in [83], where PDE layers are defined as parameterized combinations of diffusion, morphological and transport processes and are demonstrated to lead to a very compact parameterization. Conceptually, there are structural similarities between such PDE-based networks and the Gaussian derivative networks that we study in this work in the sense that: (1) the Gaussian smoothing underlying the Gaussian derivatives correspond to a diffusion process and (2) the ReLU operations between adjacent layers correspond to a special case of the morphological operations, whereas the approaches differ in the sense that (3) the Gaussian derivative networks do not contain an explicit transport mechanism, while the primitive kernels in the Gaussian derivative layers instead comprise explicit spatial oscillations or local ripples. Other types of continuous models for deep networks in terms of PDE layers have been studied in [84,85].

Concerning the notion of scale covariance and its relation to scale generalization, a general sufficiency result was presented in [2] that guarantees provable scale covariance for hierarchical networks that are constructed from continuous layers defined from partial derivatives or differential invariants expressed in terms of scale-normalized derivatives. This idea was developed in more detail for a hand-crafted quasi-quadrature network, with the layers representing idealized models of complex cells, and experimentally applied to the task of texture classification. It was demonstrated that the resulting approach allowed for scale generalization on the KTH-TIPS2 dataset, enabling classification of texture samples at scales not present in the training data.

Concerning scale generalization for CNNs, [20] presented a multi-scale-channel approach, where the same discrete CNN was applied to multiple rescaled copies of each input image. It was demonstrated that the resulting scale-channel architectures had much better ability to handle scaling transformations in the input data compared to a regular vanilla CNN, and also that the resulting approach lead to good scale

Fig. 1 Illustration of the importance of having matching support regions of receptive fields when handling scaling transformations in the image domain. For this figure, we have simulated the effect of varying the distance between the object and the camera by varying the amount of zoom for a zoom lens. The left column illustrates the effect of having a fixed receptive field size in the image domain, and how that fixed receptive field size affects the backprojected receptive fields in the world, if there are significant scale variations. In the right column, the receptive field sizes are matched under the scaling transformation, as enabled by scale covariant receptive field families and scale channel networks, which makes it possible to define deep networks that are invariant to scaling transformations, which in turn enables scale generalization



generalization, for classifying image patterns at scales not spanned by the training data.

The subject of this article is to complement the latter works, and specifically combine a specific instance of the general class of scale-covariant networks in [2] with deep learning and scale-channel networks [20], where we will choose the continuous layers as linear combinations of Gaussian derivatives and demonstrate how such an architecture allows for scale generalization.

3 Scale Generalization Based on Scale Covariant and Scale Invariant Receptive Fields

A conceptual problem that we are interested in addressing is to be able to perform *scale generalization*, *i.e.*, being able to train a deep network for image structures at some scale(s), and then being able to test the network at other scales, with maintained high recognition accuracy, and without complementary use of data augmentation.

The motivation for addressing this problem is that scale variations are very common in real-world image data,

because of objects being of different sizes in the world, and because of variations in the viewing distance to the camera. Regular CNNs are, however, known to perform very poorly when exposed to testing data at scales that are not spanned by the training data.

Thus, a main goal of this work is to equip deep networks with prior knowledge to handle scaling transformations in image data, and specifically the ability to generalize to *new scales that are not spanned by the training data*.

The methodology that we will follow to achieve such scale generalization is to require that the image descriptors or the receptive fields in the deep network are to be truly scale covariant and scale invariant. The underlying idea is that the output from the image processing operations in the deep network should thus remain sufficiently similar under scaling transformations, as will be later formalized into the commutative diagram in Fig. 2. Operationalized, we will require that the receptive field responses can be matched under scaling transformations. Specifically, we may require that the support regions of the receptive fields can be matched under scaling transformations.

Figure 1 illustrates these effects for scaling transformations in an indoor scene, caused by varying the amount of

zoom for a zoom lens, which leads to similar scaling transformations as when varying the distance between the object and the camera. In the left column, we have marked receptive fields of constant size. Then, because of the scaling transformations, the backprojections of these receptive fields will be different in the world because of the scaling transformation, which may cause problems for a deep network, depending on how it is constructed. In the right column, the support regions follow the scaling transformations in a matched and scale-covariant way, making it much easier for a deep network to recognize the object of interest under scaling transformations.

3.1 Scale Covariance and Scale Invariance

With the scaling operator⁴ \mathcal{S}_s defined by

$$(\mathcal{S}_s f)(x) = f(\mathcal{S}_s^{-1}x) = f\left(\frac{x}{s}\right). \tag{1}$$

and with Γ denoting a family of feature map operators, the notion of *scale covariance* means that the feature maps should commute with scaling transformations and transform according to

$$\Gamma'(\mathcal{S}_s f) = \mathcal{S}_s(\Gamma(f)), \tag{2}$$

where Γ' represents some possibly transformed feature map operator within the same family Γ , and adapted to the rescaled image domain. If we parameterize the feature map operator Γ by a scale parameter s , we can make a commutative diagram as shown in Fig. 2.

Implicit in Fig. 2 is the notion that the deep network performs processing at multiple scales. The motivation for such a multi-scale approach is that for a vision system that observes an *a priori* unknown scene, there is no way to know in advance what scale is appropriate for processing the image data. In the absence of further information about what scales are appropriate, the only reasonable approach is to consider image representations at multiple scales. This motivation is similar to a corresponding motivation underlying scale-space theory [12,55–65,86], which has developed a systematic approach for handling scaling transformations and image structures at multiple scales for hand-crafted image operations. Regarding deep networks that are to learn their image




Fig. 2 Commutative diagram for a scale-parameterized feature map operator Γ_s that is applied to image data under scaling transformations. The commutative diagram, which should be read from the lower left corner to the upper right corner, means that irrespective of whether the input image is first subject to a scaling transformation and then the computation of a feature map, or whether the feature map is computed first and then transformed by a scaling transformation, we should get the same result. Note, however, that this definition of scale covariance assumes a multi-scale representation of the image data, and that direct availability to the image representations at the matching scale levels $s' = Ss$ is necessary to complete the commutative diagram

representations from image data, we proceed in a conceptually similar manner, by considering deep networks with *multiple scale channels*, that perform similar types of processing operations in all the scale channels, although over multiple scales.

Scale invariance does in turn mean that the final output from the deep network Λ , for example the result of max pooling or average pooling over the multiple scale channels, should not in any way be affected by scaling transformations in the input

$$\Lambda(\mathcal{S}_s f) = \Lambda(f), \tag{3}$$

at least over some predefined range of scale that defines the capacity of the system, where scale covariance of the receptive fields in the lower layers in the network makes scale invariance possible in the final layer(s).

3.2 Approach to Scale Generalization

By basing the deep network on image operations that are provably scale covariant and scale invariant, we can train on some scale(s) and test at other scale(s).

The scale covariant and scale invariant properties of the image operations will then allow for *transfer* between image information at different scales.

In the following, we will develop this approach for one specific class of deep networks. Conceptually similar approaches to scale generalization of deep networks are applied in [2, Figures 15–16] and [20]. Conceptually similar approaches to scale invariance and scale generalization for hand-crafted computer vision operations are applied in [4–11,13,87].

⁴ From the view-point of group-equivariant CNNs [33] as applied to constructing scale-covariant or scale-equivariant CNNs based on formalism from group theory [22,23,25], this scaling operator can be seen as corresponding to a representation of the scaling group. In the more technical presentation that follows later in Sect. 4, the influence of this scaling group will be in terms of an analysis of the transformation properties under scaling transformations of the receptive field responses in the studied class of deep networks.

3.3 Influence of the Inner and the Outer Scales of the Image

Before starting with the technical treatment, let us, however, remark that when designing scale-covariant and scale-invariant image processing operations, for a given discrete image there will be only a finite range of scales over which sufficiently good approximations to continuous scale covariance and scale invariant properties can be expected hold.

For a discrete image of finite resolution, there may at finer scales be interference with the inner scale of the image, given by the resolution of the sampling grid. For example, for the set of gradually zoomed images in Fig. 1, the surface texture is clearly visible to the observer in the most zoomed-in image, whereas the same surface structure is far less visible to the observer in the most zoomed-out image. Correspondingly, if we would zoom in further to the telephone in the center of the image, peripheral parts of it may fall outside the image domain, corresponding to interference with the outer scale of the image, given by the image size. Sufficiently good numerical approximations to scale covariance and scale invariance can therefore only be expected to hold over a subrange of the scale interval for which the relevant scales of the interesting image structures are sufficiently within the range of scales defined by the inner and the outer scales of the image.

In the experiments to be presented in this work, we will handle these issues by a manual choice of the scale range for scale-space analysis as determined by properties of the dataset. When designing a computer vision system to analyze complex scenes with a priori unknown scales in the image data, there may, however, in addition be useful to add complementary mechanisms to handle the fact that an object seen from a far distance may contain far less fine-scale image structures than the same object seen from a nearby distance by the same camera. This implies that the recognition mechanisms in the system may need the ability to handle both the presence and the absence of image information over different subranges of scale, which goes beyond the continuous notions of scale covariance and scale invariance studied in this work. In a similar way, some mechanism may be needed to handle problems caused by the full object not being visible within the image domain or with a sufficient margin around its boundaries to reduce boundary effects of the spatially extended scale-space filters. For the purpose of the presentation in this paper, we will, however, leave automated handling of these issues for future work.

4 Gaussian Derivative Networks

In a traditional deep network, the filter weights are usually free variables with few additional constraints. In scale-space theory, on the other hand, theoretical results have

been presented showing that Gaussian kernels and their corresponding Gaussian derivatives constitute a canonical class of image operations⁵ [55–65]. In classical computer vision based on hand-crafted image features, it has been demonstrated that a large number of visual tasks can be successfully addressed by computing image features and image descriptors based on Gaussian derivatives, or approximations thereof, as the first layer of image features [4,5,8–10,13,28]. One could therefore raise the question if such Gaussian derivatives could also be used as computational primitives for constructing deep networks.

4.1 Gaussian Derivative Layers

Motivated by the fact that a large number of visual tasks have been successfully addressed by first- and second-order Gaussian derivatives, which are the primitive filters in the Gaussian 2-jet⁶, let us explore the consequences of using linear combinations of first- and second-order Gaussian derivatives as the class of possible filter weight primitives in a deep network.⁷ Thus, given an image f , which could either be the input image to the deep net, or some higher layer F_k in the deep network, we first compute its scale-space representation by smoothing with the Gaussian kernel

$$L(x, y; \sigma) = (g(\cdot, \cdot; \sigma) * f(\cdot, \cdot))(x, y), \quad (4)$$

where

$$g(x, y; \sigma) = \frac{1}{2\pi\sigma^2} e^{-(x^2+y^2)/2\sigma^2}. \quad (5)$$

Then, for simplicity with the notation for the spatial coordinates (x, y) and the scale parameter $\sigma \in \mathbb{R}_+$ suppressed, we consider arbitrary linear combinations of first- and second-order Gaussian derivatives as the class of possible linear

⁵ In this paper, we follow the school of scale-space axiomatics based on causality [57] or non-enhancement of local extrema [65,88], by which smoothing with the Gaussian kernel is the unique operation for scale-space smoothing, and Gaussian derivatives constitute a unique family of derived receptive fields [59,60]. For the alternative school of scale-space axiomatics based on scale invariance [55], a wider class of primitive smoothing operations is permitted [86,89,90], corresponding to the α scale spaces.

⁶ The N -jet of a signal f is the set of partial derivatives of f up to order N . The Gaussian N -jet of a signal f , is the set of Gaussian derivatives up to order N , computed by preceding the differentiation operator by Gaussian smoothing at some scale (or scales).

⁷ A complementary motivation for using first- and second-order Gaussian derivative operators as the basis for linear receptive fields, is that this is the lowest order of combining odd and even Gaussian derivative filters. First- and second-order derivatives naturally belong together, and serve as providing mutual complementary information, that approximate quadrature filter pairs [4,91,92]. In biological vision, approximations of first- and second-order derivatives also occur together in the primary visual cortex [93].




Fig. 3 The 2-D Gaussian kernel with its Cartesian partial derivatives up to order two for $\sigma = 4$




Fig. 4 First- and second-order directional derivatives of the 2-D Gaussian kernel computed from linear combinations of Cartesian partial derivatives according to equations (7) and (8) for $\sigma = 4$ and $\varphi = \pi/6$

filtering operations on f (or correspondingly for F_k):

$$J_{2,\sigma}(f) = C_0 + C_x L_\xi + C_y L_\eta + \frac{1}{2}(C_{xx} L_{\xi\xi} + 2C_{xy} L_{\xi\eta} + C_{yy} L_{\eta\eta}), \quad (6)$$

where $L_\xi, L_\eta, L_{\xi\xi}, L_{\xi\eta}$ and $L_{\eta\eta}$ are first- and second-order scale-normalized derivatives [4] according to $L_\xi = \sigma L_x, L_\eta = \sigma L_y, L_{\xi\xi} = \sigma^2 L_{xx}, L_{\xi\eta} = \sigma^2 L_{xy}$ and $L_{\eta\eta} = \sigma^2 L_{yy}$ (see Fig. 3 for an illustration of Gaussian derivative kernels) and we have also added an offset term C_0 for later use in connection with non-linearities between adjacent layers.

Based on the notion that the set of Gaussian derivatives up to order N is referred to as the Gaussian N -jet, we will refer to the expression (6), which is to be used as computational primitives for Gaussian derivative layers, as the *linearly combined Gaussian 2-jet*.

Since directional derivatives can be computed as linear combinations of partial derivatives, for first- and second-order derivatives we have (see Fig. 4 for an illustration)

$$L_\varphi = \cos \varphi L_x + \sin \varphi L_y, \quad (7)$$

$$L_{\varphi\varphi} = \cos^2 \varphi L_{xx} + 2 \cos \varphi \sin \varphi L_{xy} + \sin^2 \varphi L_{yy}, \quad (8)$$

it follows that the parameterized second-order kernels of the form (6) span all possible linear combinations of first- and second-order directional derivatives.

The corresponding affine extension of such receptive fields, by replacing the rotationally symmetric Gaussian kernel (5) for scale-space smoothing by a corresponding affine Gaussian kernel, does also constitute a good idealized model for the receptive fields of simple cells in the primary visual cortex [66,78]. In this treatment, we will, however, for simplicity restrict ourselves to regular Gaussian derivatives, based on partial derivatives and directional derivatives of rotationally symmetric Gaussian kernels.

In contrast to previous work in computer vision or functional modelling of biological visual receptive fields, where Gaussian derivatives are used as a first layer of linear receptive fields, we will, however, here investigate the consequences of coupling such receptive fields in cascade to form deep hierarchical image representations.

4.2 Definition of a Gaussian Derivative Network

To model Gaussian derivative networks with multiple feature channels in each layer, let us assume that the input image f consists of N_0 image channels, for example with $N_0 = 1$ for a grey-level image or $N_0 = 3$ for a colour image. Assuming that each layer in the network $k \in [1, K]$ should consist of $N_k \in \mathbb{Z}_+$ parallel feature channels, the feature channel $F_1^{c_{out}}$ with feature channel index $c_{out} \in [1, N_1]$ in the first layer is given by⁸

$$F_1^{c_{out}}(x, y; \sigma_1) = \sum_{c_{in} \in [1, N_0]} J_{2,\sigma_1}^{1,c_{out},c_{in}}(f^{c_{in}}(\cdot, \cdot))(x, y; \sigma_1), \quad (9)$$

where $J_{2,\sigma_1}^{1,c_{out},c_{in}}$ is the linearly combined Gaussian 2-jet according to (6) at scale⁹ σ_1 that represents the contribution

⁸ Concerning the notation in this expression, please note that the notation for the linearly combined Gaussian 2-jet seen in isolation $J_{2,\sigma_1}^{1,c_{out},c_{in}}(f^{c_{in}})$ should be understood as referring the linearly combined Gaussian 2-jet at a single scale. When we in addition lift the scale parameter as a parameter argument of the function to the form $J_{2,\sigma_1}^{1,c_{out},c_{in}}(f^{c_{in}}(\cdot, \cdot))(x, y; \sigma_1)$, the resulting expression should instead then be understood as σ_1 being allowed to vary, as a parameter argument of a function.

⁹ For the initial theoretical treatment, we can first consider this representation as being defined for all scales $\sigma_1 \in \mathbb{R}_+$. For a practical

from the image channel with index c_{in} in the input image f to the feature channel with feature channel index c_{out} in the first layer and σ_1 is the scale parameter in the first layer. The transformation between adjacent layers k and $k + 1$ is then for $k \geq 1$ given by

$$F_{k+1}^{c_{out}}(x, y; \sigma_{k+1}) = \sum_{c_{in} \in [1, N_k]} J_{2, \sigma_{k+1}}^{k+1, c_{out}, c_{in}}(\theta_k^{c_{in}}(F_k^{c_{in}}(\cdot, \cdot; \sigma_k)))(x, y; \sigma_{k+1}), \tag{10}$$

where $J_{2, \sigma_{k+1}}^{k+1, c_{out}, c_{in}}$ is the linearly combined Gaussian 2-jet according to (6) that represents the contribution to the output feature channel $F_{k+1}^{c_{out}}$ with feature channel index $c_{out} \in [1, N_{k+1}]$ in layer $k + 1$ from the input feature channel $F_k^{c_{in}}$ with feature channel index $c_{in} \in [1, N_k]$ in layer k and $\theta_k^{c_{in}}$ represents some nonlinearity, such as a ReLU stage. The parameter σ_{k+1} is the scale parameter for computing layer $k + 1$ and the parameter σ_k is the scale parameter for computing layer k .

Of course, the parameters $C_0, C_x, C_y, C_{xx}, C_{xy}$ and C_{yy} in the linearly combined Gaussian 2-jet according to (6) should be allowed to be learnt differently for each layer $k \in [1, K]$ and for each combination of input channel c_{in} and output channel c_{out} . Written out with explicit index notation for the layers and the input and output feature channels for these parameters, the explicit expression for $J_{2, \sigma_k}^{k, c_{out}, c_{in}}$ in (9) and (10) is given by

$$J_{2, \sigma_k}^{k, c_{out}, c_{in}}(h^{c_{in}}) = C_0^{k, c_{out}, c_{in}} + C_x^{k, c_{out}, c_{in}} \sigma_k L_x + C_y^{k, c_{out}, c_{in}} \sigma_k L_y + \frac{1}{2}(C_{xx}^{k, c_{out}, c_{in}} \sigma_k^2 L_{xx} + 2C_{xy}^{k, c_{out}, c_{in}} \sigma_k^2 L_{xy} + C_{yy}^{k, c_{out}, c_{in}} \sigma_k^2 L_{yy}), \tag{11}$$

with L_x, L_y, L_{xx}, L_{xy} and L_{yy} here denoting the partial derivatives of the scale-space representation obtained by convolving the argument $h^{c_{in}}$ with a Gaussian kernel (5) with standard deviation σ_k according to (4)

$$L_{x^\alpha y^\beta} = \partial_{x^\alpha y^\beta} (g(\cdot, \cdot; \sigma_k) * h^{c_{in}}(\cdot, \cdot)), \tag{12}$$

and with $h^{c_{in}}$ representing the image or feature channel with index c_{in} in either the input image f or some higher feature layer F_{k-1} .

When coupling several combined smoothing and differentiation stages in cascade in this way, with pointwise

implementation, however, the values of σ_1 as well as for the higher scale levels σ_k will be restricted to a discrete grid, as defined from Equations (13) and (14).

nonlinearities in between, it is natural to let the scale parameter for layer k be proportional to an initial scale level σ_0 , such that the scale parameter σ_k in layer k is $\sigma_k = \beta_k \sigma_0$ for some set of $\beta_k \geq \beta_{k-1} \geq 1$ and some minimum scale level $\sigma_0 > 0$. Specifically, it is natural to choose the relative factors for the scale parameters according to a geometric distribution

$$\sigma_k = r^{k-1} \sigma_0 \tag{13}$$

for some $r \geq 1$. By gradually increasing the size of the receptive fields in this way, the transition from lower to higher layers in the hierarchy will correspond to gradual transitions from local to regional information. For continuous network models, this mechanism thus provides a way to gradually increase the receptive field size in deeper layers without explicit need for a discrete subsampling operation that would imply larger steps in the receptive fields size of deeper layers.

By additionally varying the parameter σ_0 in the above relationship, either continuously with $\sigma_0 \in \mathbb{R}$ or according to some self-similar distribution

$$\sigma_0 = \gamma^i \tag{14}$$

for some set of integers $i \in \mathbb{Z}$ and some $\gamma > 1$, the Gaussian derivative network defined from (9) and (10) will constitute a multi-scale representation that will also be provably scale covariant, as will be formally shown below in Sect. 4.3. Specifically, the network obtained for each value of σ_0 , with the derived values of σ_1 in (9) as well as of σ_k and σ_{k+1} in (10), will be referred to as a *scale channel*. Letting the initial scale levels be given by a self-similar distribution in this way reflects the desire that in the absence of further information the deep network should be agnostic with regard to preferred scales in the input.

A similar idea of using Gaussian derivatives as structured receptive fields in convolutional networks has also been explored in [70], although not in the relation to scale covariance, multiple scale channels or using a self-similar sampling of the scale levels at consecutive depths.

4.3 Provable Scale Covariance

To prove that a deep network constructed by coupling linear filtering operations of the form (6) with pointwise nonlinearities in between is scale covariant, let us consider two images f and f' that are related by a scaling transformation $f'(x', y') = f(x, y)$ for $x' = Sx, y' = Sy$ and some spatial scaling factor $S > 0$.

A basic scale covariance property of the Gaussian scale-space representation is that if the scale parameters σ and σ' in the two image domains are related according to $\sigma' = S\sigma$,

then the Gaussian scale-space representations are equal at matching image points and scales [4, Eq. (16)]

$$L'(x', y'; \sigma') = L(x, y; \sigma). \tag{15}$$

Additionally regarding spatial derivatives, it follows from a general result in [4, Eq. (20)] that the scale-normalized spatial derivatives will also be equal (when using scale normalization power $\gamma = 1$ in the scale-normalized derivative concept):

$$L'_{\xi\alpha\eta\beta}(x', y'; \sigma') = L_{\xi\alpha\eta\beta}(x, y; \sigma). \tag{16}$$

Applied to the linearly combined Gaussian 2-jet in (6), it follows that also the linearly combined Gaussian 2-jets computed from two mutually rescaled scalar image patterns f and f' will be related by a similar scaling transformation

$$J_{2,\sigma'}(f'(\cdot, \cdot))(x', y'; \sigma') = J_{2,\sigma}(f(\cdot, \cdot))(x, y; \sigma) \tag{17}$$

provided that the image positions (x, y) and (x', y') and the scale parameters σ and σ' are appropriately matched, and that the parameters $C_0, C_x, C_y, C_{xx}, C_{xy}$ and C_{yy} in (6) are shared between the two 2-jets at the different scales.

Thus, provided that the image positions and the scale levels are appropriately matched according to $x' = Sx, y' = Sy$ and $\sigma' = S\sigma$, it holds that the corresponding feature channels F_1^{out} and $F_1'^{\text{out}}$ in the first layer according to (9) are equal up to a scaling transformation:

$$\begin{aligned} F_1'^{\text{out}}(x', y'; \sigma'_1) &= \sum_{c_{\text{in}} \in [1, N_0]} J_{2,\sigma'_1}^{1,c_{\text{out}},c_{\text{in}}}(f'^{c_{\text{in}}}(\cdot, \cdot))(x', y'; \sigma'_1) \\ &= \sum_{c_{\text{in}} \in [1, N_0]} J_{2,\sigma_1}^{1,c_{\text{out}},c_{\text{in}}}(f^{c_{\text{in}}}(\cdot, \cdot))(x, y; \sigma_1) \\ &= F_1^{\text{out}}(x, y; \sigma_1), \end{aligned} \tag{18}$$

provided that the parameters $C_0^{1,c_{\text{out}},c_{\text{in}}}, C_x^{1,c_{\text{out}},c_{\text{in}}}, C_y^{1,c_{\text{out}},c_{\text{in}}}, C_{xx}^{1,c_{\text{out}},c_{\text{in}}}, C_{xy}^{1,c_{\text{out}},c_{\text{in}}}$ and $C_{yy}^{1,c_{\text{out}},c_{\text{in}}}$ in the 2-jets according to (11) for $k = 1$ are shared between the scale channels for corresponding combinations of feature channels, as indexed by c_{in} and c_{out} .

By continuing a corresponding construction of higher layers, by applying similar operations in cascade, with pointwise non-linearities such as ReLU stages in between, according to (10) with the initial scale levels σ_0 and σ'_0 in (13) related according to $\sigma'_0 = S\sigma_0$, it follows that also the corresponding feature channels $F_{k+1}'^{\text{out}}$ and F_{k+1}^{out} in the higher layers in the hierarchy will be equal up to a scaling transformation, since the input data from the adjacent finer layers $F_k'^{\text{cin}}$ and F_k^{cin} are




Fig. 5 Commutative diagram for a scale-covariant Gaussian derivative network constructed by coupling linear combinations of scale-normalized Gaussian derivatives in cascade, with nonlinear ReLU stages in between. Because of the transformation properties of the individual layers under scaling transformations, it will be possible to perfectly match the corresponding layers F_k and F_k' under a scaling transformation of the underlying image domain $f'(x') = f(x)$ for $x' = Sx$ and $y' = Sy$, provided that the scale parameter σ_k in layer k is proportional to the scale parameter σ_1 in the first layer, $\sigma_k = r_k^2 \sigma_1$, for some scalar constant $r_k > 1$. For such a network, the scale parameters in the two domains should be related according to $\sigma'_k = S\sigma_k$. Note, however, that for a scale-discretized implementation, this commutative property holds exactly over a continuous image domain only if the scale levels σ and σ' are part of the scale grid, thus specifically only for discrete scaling factors S that can be exactly represented on the discrete scale grid. For other scaling factors, the results will instead be numerical approximations, with the accuracy of the approximation determined by the combination of the network architecture with the learning algorithm. (In this schematic illustration, we have for simplicity suppressed the notation for multiple feature channels in the different layers, and also suppressed the notation for the pointwise non-linearities between adjacent layers)

already known to be related by plain scaling transformation (see Fig. 5 for an illustration):

$$\begin{aligned} F_{k+1}'^{\text{out}}(x', y'; \sigma'_{k+1}) &= \sum_{c_{\text{in}} \in [1, N_k]} J_{2,\sigma'_{k+1}}^{k+1,c_{\text{out}},c_{\text{in}}}(\theta_k^{c_{\text{in}}}(F_k'^{c_{\text{in}}}(\cdot, \cdot; \sigma'_k)))(x', y'; \sigma'_{k+1}) \\ &= \sum_{c_{\text{in}} \in [1, N_k]} J_{2,\sigma_{k+1}}^{k+1,c_{\text{out}},c_{\text{in}}}(\theta_k^{c_{\text{in}}}(F_k^{\text{cin}}(\cdot, \cdot; \sigma_k)))(x, y; \sigma_{k+1}) \\ &= F_{k+1}^{\text{out}}(x, y; \sigma_{k+1}), \end{aligned} \tag{19}$$

again assuming that the parameters $C_0^{k,c_{\text{out}},c_{\text{in}}}, C_x^{k,c_{\text{out}},c_{\text{in}}}, C_y^{k,c_{\text{out}},c_{\text{in}}}, C_{xx}^{k,c_{\text{out}},c_{\text{in}}}, C_{xy}^{k,c_{\text{out}},c_{\text{in}}}$ and $C_{yy}^{k,c_{\text{out}},c_{\text{in}}}$ in the 2-jets according to (11) are shared between the scale channels for corresponding layers and combinations of feature channels, as indexed by c_{in} and c_{out} .

A pointwise nonlinearity, such as a ReLU stage, trivially commutes with scaling transformations and does therefore not affect the scale covariance properties.

In a recursive manner, we thereby prove that scale covariance in the lower layers imply scale covariance in any higher layer.

4.4 Provable Scale Invariance

To prove scale invariance after max pooling over the scale channels, let us assume that we have an *infinite* set of scale channels \mathcal{S} , with the initial scale values σ_0 in (14) either continuously distributed with

$$\mathcal{S} = \{\sigma \in \mathbb{R}_+\} \tag{20}$$

or discrete in the set

$$\mathcal{S} = \{\sigma_i = \gamma^i, \forall i \in \mathbb{Z}\}. \tag{21}$$

for some $\gamma > 1$.

The max pooling operation over the scale channels for feature channel c in layer k will then at every image position (x, y) and with the scale parameter in this layer σ_k and the relative scale factor r according to (13) return the value

$$F_{k,sup}^c(x, y) = \sup_{\sigma_k} F_k^c(x, y; \sigma_k) = \sup_{\sigma_0 \in \mathcal{S}} F_k^c(x, y; r^{k-1} \sigma_0). \tag{22}$$

Let us without essential loss of generality assume that the scaling transformation is performed around the image point (x, y) , implying that we can without essential loss of generality assume that the origin is located at this point. The set of feature values before the scaling transformation is then given by

$$M = \{F_k^c(0, 0; r^{k-1} \sigma_0) \forall \sigma_0\} \tag{23}$$

and the set of feature values after the scaling transformation

$$M' = \{F_k^c(0, 0; r^{k-1} \sigma'_0) \forall \sigma'_0\}. \tag{24}$$

With the relationship $\sigma'_0 = S \sigma_0$, these sets are clearly equal provided that $S > 0$ in the continuous case or $S = \gamma^j$ for some $j \in \mathbb{Z}$ in the discrete case, implying that the supremum of the set is preserved (or the result of any other permutation invariant pooling operation, such as the average). Because of the closedness property under scaling transformations, the scaling transformation just shifts the set of feature values over scales along the scale axis.

In this way, the result after max pooling over the scale channels is essentially scale invariant, in the sense that the

result of the max pooling operation at any image point follows the geometric transformation of the image point under scaling transformations.

When using a *finite* number of scale channels, the result of max pooling over the scale channels is, however, not guaranteed to be truly scale invariant, since there could be boundary effects, implying that the maximum over scales moves in to or out from a finite scale interval, because of the scaling transformation that shifts the position of the scale maxima on the scale axis.

To reduce the likelihood of such effects occurring, we propose as design criterion to ensure that there should be a sufficient number of additional scale channels below and above the effective training scales¹⁰ in the training data. The intention behind this is that the learning algorithm could then learn that the image structures that occur below and above the effective training scales are less relevant, and thereby associate lower values of the feature maps to such image structures. In this way, the risk should be reduced that erroneous types of image structures are being picked up by the max pooling operation over the multiple scale channels, by the scale channels near the scale boundaries thereby having lower magnitude values in their feature maps than the central ones.

A similar scale boundary handling strategy is used in the scale channel networks based on applying a fixed CNN to a set of rescalings of the original image in [20], as opposed to the approach here based on applying a set of rescaled CNNs to a fixed size input image.

5 Experiments with a Single-Scale-Channel Network

To investigate the ability of these types of deep hierarchical Gaussian derivative networks to capture image structures with different image shapes, we first did initial experiments with the regular MNIST dataset [94]. We constructed a 6-layer network in PyTorch [95] with 12-14-16-20-64 channels in the intermediate layers and 10 output channels, intended to learn each type of digit, see Fig. 6 for an illustration.

¹⁰ To formally define the notion of “effective training scale”, we can consider the set of scale values of the scale channels that lead to the maximum value over scales for a max pooling network (or the range of scales that contain the dominant mass over scales for an average pooling network), formed as the union of all the samples in the training data of a specific size, and with some additional cutoff function to select the majority of the responses, specifically with some suppression of spurious outliers. Since these resulting scale levels could be expected to vary between different training samples of roughly the same size, the notion of “effective training scales” should be a scale interval rather than a single scale, and may vary depending upon the properties of the training data.

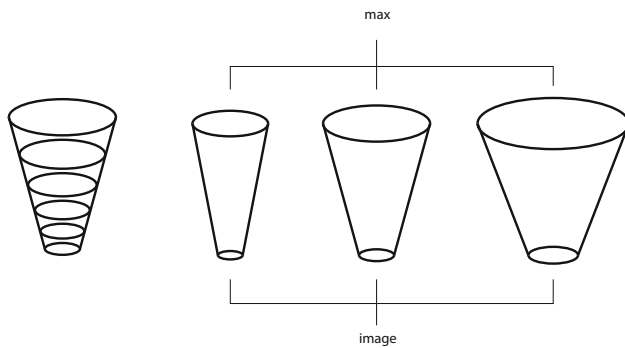


Fig. 6 (left) Schematic illustration of the architecture of the single-scale-channel network, with 6 layers of receptive fields at successively coarser levels of scale. (right) Schematic illustration of the architecture of a multi-scale-channel network, with multiple parallel scale channels over a self-similar distribution of the initial scale level σ_0 in the hierarchy of Gaussian derivative layers coupled in cascade

We chose the initial scale level $\sigma_0 = 0.9$ pixels and the relative scale ratio $r = 1.25$ in (13), implying that the maximum value of σ is $0.9 \times 1.25^5 \approx 2.7$ pixels relative to the image size of 28×28 pixels. The individual receptive fields do then have larger spatial extent because of the spatial extent of the Gaussian kernels used for image smoothing and the larger positive and negative side lobes of the first- and second-order derivatives.

We used regular ReLU stages between the filtering steps, but no spatial max pooling or spatial stride, and no fully connected layer, since such operations would destroy the scale covariance. Instead, the receptive fields are solely determined from linear combinations of Gaussian derivatives, with successively larger receptive fields of size $\sigma_0 r^{k-1}$, which enable a gradual integration from local to regional image structures. In the final layer, only the value at the central pixel is extracted, or here for an even image size, the average over the central 2×2 neighbourhood, which, however, destroys full scale covariance. To ensure full scale covariance, the input images should instead have odd image size.

The network was trained on 50 000 of the training images in the dataset, with the offset term $C_0^{k, c_{out}, c_{in}}$ that serves as the bias for the nonlinearity and the weights $C_x^{k, c_{out}, c_{in}}$, $C_y^{k, c_{out}, c_{in}}$, $C_{xx}^{k, c_{out}, c_{in}}$, $C_{xy}^{k, c_{out}, c_{in}}$ and $C_{yy}^{k, c_{out}, c_{in}}$ of the Gaussian derivatives in (11) initiated to random values and trained individually for each layer and feature channel by stochastic gradient descent over 40 epochs using the Adam optimizer [96] set to minimize the binary cross-entropy loss. We used a cosine learning curve with maximum learning rate of 0.01 and minimum learning rate 0.00005 and using batch normalization over batches with 50 images. The experiment lead to 99.93% training accuracy and 99.43% test accuracy on the test dataset containing 10 000 images.

Notably, the training accuracy does not reach 100.00%, probably because of the restricted shapes of the filter weights,

as determined by the *a priori* shapes of the receptive fields in terms of linear combinations of first- and second-order Gaussian derivatives. Nevertheless, the test accuracy is quite good given the moderate number of parameters in the network ($6 \times (12 + 12 \times 14 + 14 \times 16 + 16 \times 20 + 20 \times 64 + 64 \times 10) = 15\,864$).

5.1 Discrete Implementation

In the numerical implementation of scale-space smoothing, we used separable smoothing with the discrete analogue of the Gaussian kernel $T(n; s) = e^{-s} I_n(s)$ for $s = \sigma^2$ in terms of modified Bessel functions I_n of integer order [97].¹¹ The discrete derivative approximations were computed by central differences, $\delta_x = (-1/2, 0, 1/2)$, $\delta_{xx} = (1, -2, 1)$, $\delta_{xy} = \delta_x \delta_y$, etc., implying that the spatial smoothing operation can be shared between derivatives of different order, and implying that scale-space properties are preserved in the discrete implementation of the Gaussian derivatives [98]. This is a general methodology for computing Gaussian derivatives for a large number of visual tasks.

By computing the Gaussian derivative responses in this way, the scale-space smoothing is only performed once for each scale level, and there is no need for repeating the scale-space smoothing for each order of the Gaussian derivatives.

6 Experiments with a Multi-scale-Channel Network


To investigate the ability of a multi-scale-channel network to handle spatial scaling transformations, we made experiments on the MNIST Large Scale dataset [20,21]. This dataset contains rescaled digits from the original MNIST dataset [94] embedded in images of size 112×112 , see Fig. 7 for an illustration.

For training, we used either of the datasets containing 50 000 rescaled digits with relative scale factors 1, 2 or 4, respectively, henceforth referred to as training sizes 1, 2 and 4. For testing, the dataset contains 10 000 rescaled digits with relative scale factors between 1/2 and 8, respectively, with a relative scale ratio of $\sqrt[4]{2}$ between adjacent testing sizes.

To investigate the properties of a multi-scale-channel architecture experimentally, we created a multi-scale-channel network with 8 scale channels with their initial scale values σ_0 between $1/\sqrt{2}$ and 8 and a scale ratio of $\sqrt{2}$ between

¹¹ This way of implementing Gaussian convolution on discrete spatial domain corresponds to the solution of a purely spatial discretization of the diffusion equation $\partial_s L = \frac{1}{2} \nabla^2 L$ that describes the effect of Gaussian convolution (4), with the continuous Laplacian operator ∇^2 replaced by the five-point operator ∇_5^2 defined by $(\nabla_5^2 L)(x, y) = L(x+1, y) + L(x-1, y) + L(x, y+1) + L(x, y-1) - 4L(x, y)$ [97].

Fig. 7 Sample images from the MNIST Large Scale dataset [20,21]. This figure shows digits for sizes in the range [1, 4], for which there are training data. In addition, the MNIST Large Scale dataset contains testing data over the wider size range [1/2, 8]



adjacent scale channels. For each scale channel, we used a Gaussian derivative network of similar architecture as the single-scale-channel network, with 12-14-16-20-64 channels in the intermediate layers and 10 output channels, and with a relative scale ratio $r = 1.25$ (13) between adjacent layers, implying that the maximum value of σ in each channel is $\sigma_0 \times 1.25^5 \approx 3.1 \sigma_0$ pixels.

Importantly, the parameters $C_0^{k,c_{out},c_{in}}$, $C_x^{k,c_{out},c_{in}}$, $C_y^{k,c_{out},c_{in}}$, $C_{xx}^{k,c_{out},c_{in}}$, $C_{xy}^{k,c_{out},c_{in}}$ and $C_{yy}^{k,c_{out},c_{in}}$ in (11) are shared between the scale channels, implying that the scale channels together are truly scale covariant, because of the parameterization of the receptive fields in terms of scale-normalized Gaussian derivatives. The batch normalization stage is also shared between the scale channels. The output from max pooling over the scale channels is furthermore truly scale invariant, if we assume an infinite number of scale channels, so that scale boundary effects can be disregarded.

Figure 8 shows the result of an experiment to investigate the ability of such a multi-scale-channel network to generalize to testing at scales not present in the training data.

For the experiment shown in the top figure, we have used 50 000 training images for training size 1, and 10 000 testing images for each one of the 19 testing sizes between 1/2 and 8 with a relative size ratio of $\sqrt[4]{2}$ between adjacent testing sizes. The red curve shows the generalization performance for a single-scale-channel network with $\sigma_0 = 1$, whereas the blue curve shows the generalization performance for the multi-scale-channel network with 8 scale channels between $1/\sqrt{2}$ and 8.

As can be seen from the graphs, the generalization performance is very good for the multi-scale-channel network, for sizes in roughly in the range $1/\sqrt{2}$ and $4\sqrt{2}$. For smaller testing sizes near 1/2, there are discretization problems due to sampling artefacts and too fine scale levels relative to the grid spacing in the image, implying that the transformation properties under scaling transformations of the underlying Gaussian derivatives are not well approximated in the dis-

crete implementation. For larger testing sizes near 8, there are problems due to boundary effects and that the entire digit is not visible in the testing stage, implying a mismatch between the training data and the testing data. Otherwise, the generalization performance is very good over the size range between 1 and 4.

For the single-scale-channel network, the generalization performance to scales far from the scales in the training data is on the other hand very poor.

In the middle figure, we show the result of a similar experiment for training images of size 2 and with the initial scale level $\sigma_0 = 2$ for the single-scale-channel network. The bottom figure shows the result of a similar experiment performed with training images of size 4 and with the initial scale level $\sigma_0 = 4$ for the single-scale-channel network.

Figure 9 shows a joint visualization of the generalization performance for all these experiments, where we have also zoomed in on the top performance values in the range 98–99%. In addition to the results from the default network with 12-16-24-32-64-10 feature channels, we do also show results obtained from a larger network with 16-24-32-48-64-10 feature channels, which has more degrees of freedom in the training stage (a total number of $6 \times (16 + 16 \times 24 + 24 \times 32 + 32 \times 48 + 48 \times 64 + 64 \times 10) = 38\,496$ parameters) and leads to higher top performance and also somewhat better generalization performance. As can be seen from the graphs, the performance values for training sizes 1, 2 and 4, respectively, are quite similar for testing data with sizes in the range between 1 and 4, a size range for which the discretization errors in the discrete implementation can be expected to be low (a problem at too fine scales) and the influence of boundary effects causing a mismatch between what parts of the digits are visible in the testing data compared to the training data (a problem at too coarse scales).

To conclude, the experiment demonstrates that it is possible to use the combination of (1) scale-space features as computational primitives for a deep learning method with

Fig. 8 Experiments showing the ability of a multi-scale-channel network to generalize to new scale levels not present in the training data. In the top row, all training data are for size 1, whereas we evaluate on multiple testing sets for each one of the sizes between 1/2 and 8. The red curve shows the generalization performance for a single-scale-channel network for $\sigma_0 = 1$, whereas the blue curve shows the result for a multi-scale-channel network covering the range of σ_0 -values between $1/\sqrt{2}$ and 8. As can be seen from the result, the generalization ability is much better for the multi-scale-channel network compared to the single-scale-channel network. In the middle row, a similar type of experiment is repeated for training size 2 and with $\sigma_0 = 2$ for the single-scale-channel network. In the bottom row, a similar experiment is performed for training size 4 and with $\sigma_0 = 4$ for the single-scale-channel network. (Horizontal axis: Scale of testing data)


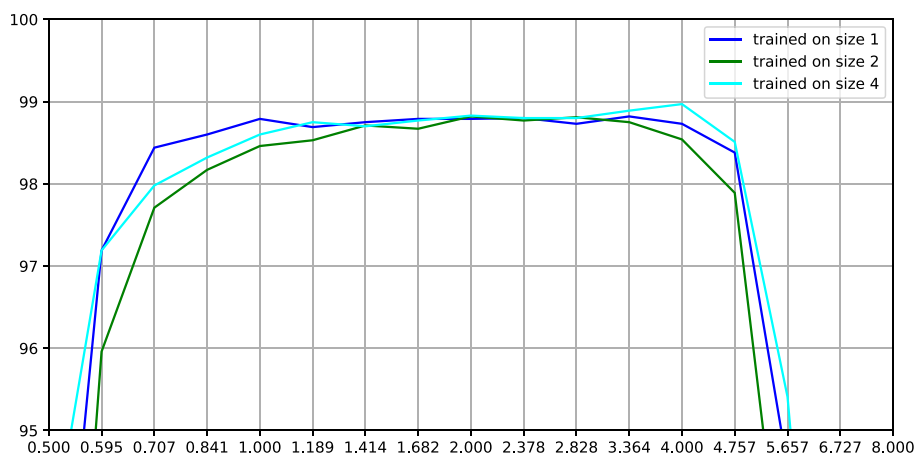



Fig. 9 Joint visualization of the generalization performance when training a multi-scale-channel Gaussian derivative network on training data with sizes 1, 2 and 4, respectively. As can be seen from the graphs, the performance is rather similar for all these networks over the size range between 1 and 4, a range for which the discretization errors in the discrete implementation can be expected to be low (a problem at too fine scales) and the influence of boundary effects causing a mismatch between what parts of the digits are visible in the testing data compared to a training data (a problem at too coarse scales). The top figure shows the results for the default network with 12-16-24-32-64-10 feature channels. The bottom figure shows the results for a larger network with 16-24-32-48-64-10 feature channels. (Horizontal axis: Scale of testing data)

Scale generalization performance when training on different sizes: Default network



Scale generalization performance when training on different sizes: Larger network



(2) the closed-form transformation properties of the scale-space primitives under scaling transformations to (3) make a deep network generalize to new scale levels not spanned by the training data.

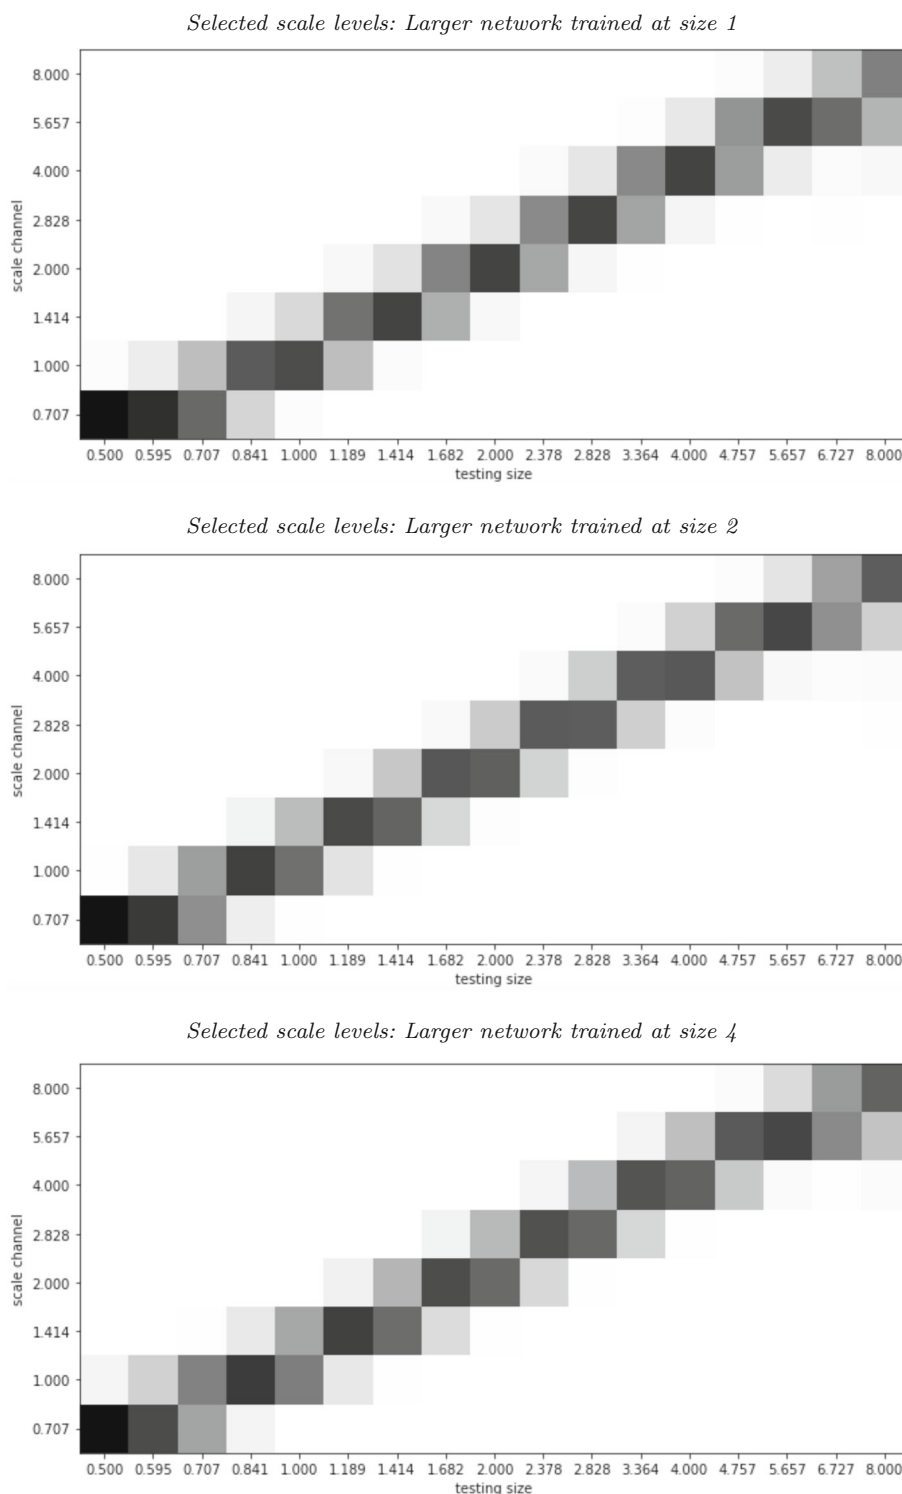
6.1 Scale Selection Properties

Since the Gaussian derivative network is expressed in terms of scale-normalized derivatives over multiple scales, and the max-pooling operation over the scale channels implies detecting maxima over scale, the resulting approach shares similarities to classical methods for scale selection based on local extrema over scales of scale-normalized derivatives [4,5,99]. The approach is also closely related to the scale selection approach in [100,101] based on choosing the scales at which a supervised classifier delivers class labels with the highest posterior.

A limitation of choosing only a single maximum over scales compared to processing multiple local extrema over scales as in [4,5,99], however, is that the approach may be sensitive to boundary effects at the scale boundaries, implying that the scale generalization properties may be affected depending on how many coarser-scale and/or finer-scale channels are being processed relative to the scale of the image data.

In Fig. 10, we have visualized the scale selection properties when applying these multi-scale-channel Gaussian derivative networks to the MNIST Large Scale dataset. For each one of the training sizes 1, 2 and 4, we show what scales are selected as function of the testing size. Specifically, for each testing size, shown on the horizontal axis, we display a histogram of the scale channels at which the maximum over scales is assumed over all the samples in the testing set, with the finest scale at the bottom and the coarsest scale at the top.

Fig. 10 Visualization of the scale channels that are selected in the max pooling stage, when training the larger network for each one of the sizes 1, 2 and 4. For each testing size, shown on the horizontal axis, the vertical axis displays a histogram of the scale levels at which the maximum over the scale channels is assumed for the different samples in the testing set, with the lowest scale at the bottom and the highest scale at the top. As can be seen from the figure, there is a general tendency of the composed classification scheme to select coarser scale levels with increasing size of the image structures, in agreement with the conceptual similarity to classical methods for scale selection based on detecting local extrema over scales of scale-normalized derivatives, the difference being that here only the global maximum over scale is used, as opposed to the detection of multiple local extrema over scale in classical scale selection methods



As can be seen from the figure, the network has an overall tendency of selecting coarser scale levels with increasing size of the image structures it is applied to. Specifically, the overall tendency is that the selected scale is proportional to the size of the testing data, in agreement with the theory for scale selection based on local extrema over scales of scale-normalized derivatives.

Except for minor quantization variations due to the discrete bins, these scale selection histograms are very similar for the different training sizes 1, 2 and 4. In combination with the previously presented experiments, the results in this paper thus demonstrate that it is possible to use scale-space operations as computational primitives in deep networks, and to use the transformation properties of such computational

primitives under scaling transformations to perform scale generalization.

6.2 Properties of Multi-scale-Channel Network Training

While the scale generalization graphs in Fig. 9 and the scale selection histograms in Fig. 10 are very similar for the different training scales 1, 2 or 4, we can note that the results are, however, not fully identical.

This can partly be explained from the fact that the training error does not approach zero (the training accuracy for the larger multi-scale channel networks trained at a single scale reaches the order of 99.6–99.7%), implying that the net effects of the training error on the properties of the network may be somewhat different when training the network on different datasets, here for different training sizes.

Additionally, even if the training error would have approached zero, different types of scale boundary effects and image boundary effects could occur for the different networks, depending on what single size training data the multi-scale channel network is trained on, implying that the learning algorithm could lead to sets of filter weights in the networks with somewhat different properties, because of a lack of full scale covariance in the training stage, although the architecture of the network is otherwise fully scale covariant in the continuous case.

The training process is initiated randomly, and during the gradient descent optimization, the network has to learn to associate large values of the feature map for the appropriate class for each training sample in some scale channel, while also having to learn to not associate large values for erroneous classes for any other training samples in any other scale channel.

Training of a multi-scale-channel network could therefore be considered as a harder training problem than the training of a single-scale-channel network. Experimentally, we have also observed that the training error is significantly larger for a multi-scale-channel network than for a single-scale-channel network, and that the training procedure had not converged fully after the 20 epochs that we used for training the multi-scale-channel networks.

A possible extension of this work, is therefore to perform a deeper study regarding the task of training multi-scale-channel networks, and investigate if this training task calls for other training strategies than used here, specifically considering that the computational primitives in the layers of the Gaussian derivative networks are also different from the computational primitives in regular CNNs, while we have here used a training strategy suitable for regular CNNs.

Notwithstanding these possibilities for improvements in the training scheme, the experiments in the paper demonstrate that it is possible to use scale-space operations as

computational primitives in deep networks, and to use the transformation properties of such computational primitives under scaling transformations to perform scale generalization, which is the main objective of this study.

7 Summary and Discussion

We have presented a hybrid approach between scale-space theory and deep learning, where the layers in a hierarchical network architecture are modelled as continuous functions instead of discrete filters, and are specifically chosen as scale-space operations, here in terms of linear combinations of scale-normalized Gaussian derivatives. Experimentally, we have demonstrated that the resulting approach allows for scale generalization and enables good performance for classifying image patterns at scales not spanned by the training data.

The work is intended as a proof-of-concept of the idea of using scale-space features as computational primitives in a deep learning method, and of using their closed-form transformation properties under scaling transformations to perform extrapolation or generalization to new scales not present in the training data.

Concerning the choice of Gaussian derivatives as computational primitives in the method, it should, however, be emphasized that the necessity results that specify the uniqueness of these kernels only state that the first layer of receptive fields should be constructed in terms of Gaussian derivatives. Concerning higher layers, further studies should be performed concerning the possibilities of using other scale-space features in the higher layers, that may represent the variability of natural image structures more efficiently, within the generality of the sufficiency result for scale-covariant continuous hierarchical networks in [2].

Acknowledgements I would like to thank Ylva Jansson for sharing her code for training and testing networks in PyTorch and for valuable comments on an earlier version of this manuscript.

I would also like to thank the anonymous reviewers for valuable comments, which helped to improve the presentation.

Funding Open access funding provided by Royal Institute of Technology.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copy-

right holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Jansson, Y., Lindeberg, T.: Exploring the ability of CNNs to generalise to previously unseen scales over wide scale ranges. In: International Conference on Pattern Recognition (ICPR 2020), pp. 1181–1188. (2021). Extended version in [arXiv:2004.01536](https://arxiv.org/abs/2004.01536)
- Lindeberg, T.: Provably scale-covariant continuous hierarchical networks based on scale-normalized differential expressions coupled in cascade. *J. Math. Imaging Vis.* **62**, 120–148 (2020)
- Lindeberg, T.: Scale-covariant and scale-invariant Gaussian derivative networks. In: Proceedings of Scale Space and Variational Methods in Computer Vision (SSVM 2021), volume 12679 of Springer LNCS, pp. 3–14 (2021)
- Lindeberg, T.: Feature detection with automatic scale selection. *Int. J. Comput. Vis.* **30**, 77–116 (1998)
- Lindeberg, T.: Edge detection and ridge detection with automatic scale selection. *Int. J. Comput. Vis.* **30**, 117–154 (1998)
- Bretzner, L., Lindeberg, T.: Feature tracking with automatic selection of spatial scales. *Comput. Vis. Image Underst.* **71**, 385–392 (1998)
- Chomat, O., de Verdiere, V., Hall, D., Crowley, J.: Local scale selection for Gaussian based description techniques. In: Proceedings of European Conference on Computer Vision (ECCV 2000), volume 1842 of Springer LNCS, Dublin, Ireland, pp. 117–133 (2000)
- Mikolajczyk, K., Schmid, C.: Scale and affine invariant interest point detectors. *Int. J. Comput. Vis.* **60**, 63–86 (2004)
- Lowe, D.G.: Distinctive image features from scale-invariant keypoints. *Int. J. Comput. Vis.* **60**, 91–110 (2004)
- Bay, H., Ess, A., Tuytelaars, T., van Gool, L.: Speeded up robust features (SURF). *Comput. Vis. Image Underst.* **110**, 346–359 (2008)
- Tuytelaars, T., Mikolajczyk, K.: A Survey on Local Invariant Features. Foundations and Trends in Computer Graphics and Vision, vol. 3. Now Publishers, New York (2008)
- Lindeberg, T.: Generalized axiomatic scale-space theory. In: Hawkes, P. (ed.) Advances in Imaging and Electron Physics, vol. 178, pp. 1–96. Elsevier, Amsterdam (2013)
- Lindeberg, T.: Image matching using generalized scale-space interest points. *J. Math. Imaging Vis.* **52**, 3–36 (2015)
- Fawzi, A., Frossard, P.: Manitest: are classifiers really invariant? In: British Machine Vision Conference (BMVC 2015) (2015)
- Singh, B., Davis, L.S.: An analysis of scale invariance in object detection—SNIP. In: Proceedings of Computer Vision and Pattern Recognition (CVPR 2018), pp. 3578–3587 (2018)
- Xu, Y., Xiao, T., Zhang, J., Yang, K., Zhang, Z.: Scale-invariant convolutional neural networks. *arXiv preprint arXiv:1411.6369* (2014)
- Kanazawa, A., Sharma, A., Jacobs, D.W.: Locally scale-invariant convolutional neural networks. In: NIPS 2014 Deep Learning and Representation Learning Workshop (2014). *arXiv preprint arXiv:1412.5104*
- Marcos, D., Kellenberger, B., Lobry, S., Tuia, D.: Scale equivariance in CNNs with vector fields. In: ICML/FAIM 2018 Workshop on Towards Learning with Limited Labels: Equivariance, Invariance, and Beyond (2018). *arXiv preprint arXiv:1807.11783*
- Ghosh, R., Gupta, A.K.: Scale steerable filters for locally scale-invariant convolutional neural networks. In: ICML Workshop on Theoretical Physics for Deep Learning (2019). *arXiv preprint arXiv:1906.03861*
- Jansson, Y., Lindeberg, T.: Exploring the ability of CNNs to generalise to previously unseen scales over wide scale ranges. In: International Conference on Pattern Recognition (ICPR 2020), pp. 1181–1188 (2021)
- Jansson, Y., Lindeberg, T.: MNIST Large Scale dataset. Zenodo (2020). <https://www.zenodo.org/record/3820247>
- Sosnovik, I., Szmaja, M., Smeulders, A.: Scale-equivariant steerable networks. In: International Conference on Learning Representations (ICLR 2020) (2020)
- Worrall, D., Welling, M.: Deep scale-spaces: Equivariance over scale. In: Advances in Neural Information Processing Systems (NeurIPS 2019), pp. 7366–7378 (2019)
- Lindeberg, T.: Provably scale-covariant networks from oriented quasi quadrature measures in cascade. In: Proceedings of Scale Space and Variational Methods in Computer Vision (SSVM 2019), vol. 11603 of Springer LNCS, pp. 328–340 (2019)
- Bekkers, E.J.: B-spline CNNs on Lie groups. In: International Conference on Learning Representations (ICLR 2020) (2020)
- Singh, B., Najibi, M., Sharma, A., Davis, L.S.: Scale normalized image pyramids with AutoFocus for object detection. *IEEE Trans. Pattern Anal. Mach. Intell.* (2021)
- Li, Y., Chen, Y., Wang, N., Zhang, Z.: Scale-aware trident networks for object detection. In: Proceedings of International Conference on Computer Vision (ICCV 2019), pp. 6054–6063 (2019)
- Schiele, B., Crowley, J.: Recognition without correspondence using multidimensional receptive field histograms. *Int. J. Comput. Vis.* **36**, 31–50 (2000)
- Linde, O., Lindeberg, T.: Object recognition using composed receptive field histograms of higher dimensionality. *Int. Conf. Pattern Recognit.* **2**, 1–6 (2004)
- Laptev, I., Lindeberg, T.: Local descriptors for spatio-temporal recognition. In: Proceedings of ECCV'04 Workshop on Spatial Coherence for Visual Motion Analysis, vol. 3667 of Springer LNCS, Prague, Czech Republic, pp. 91–103 (2004)
- Linde, O., Lindeberg, T.: Composed complex-cue histograms: an investigation of the information content in receptive field based image descriptors for object recognition. *Comput. Vis. Image Underst.* **116**, 538–560 (2012)
- Larsen, A.B.L., Darkner, S., Dahl, A.L., Pedersen, K.S.: Jet-based local image descriptors. In: Proceedings of European Conference on Computer Vision (ECCV 2012), vol. 7574 of Springer LNCS, pp. 638–650. Springer (2012)
- Cohen, T., Welling, M.: Group equivariant convolutional networks. In: International Conference on Machine Learning (ICML 2016), pp. 2990–2999 (2016)
- Jaderberg, M., Simonyan, K., Zisserman, A., Kavukcuoglu, K.: Spatial transformer networks. In: Proceedings of Neural Information Processing Systems (NIPS 2015), pp. 2017–2025 (2015)
- Lin, C.H., Lucey, S.: Inverse compositional spatial transformer networks. In: Proceedings of Computer Vision and Pattern Recognition (CVPR 2017), pp. 2568–2576 (2017)
- Finnveden, L., Jansson, Y., Lindeberg, T.: Understanding when spatial transformer networks do not support invariance, and what to do about it. In: International Conference on Pattern Recognition (ICPR 2020), pp. 3427–3434 (2021). Extended version in [arXiv:2004.11678](https://arxiv.org/abs/2004.11678)
- Jansson, Y., Maydanskiy, M., Finnveden, L., Lindeberg, T.: Inability of spatial transformations of CNN feature maps to support invariant recognition. *arXiv preprint arXiv:2004.14716* (2020)
- Sermanet, P., Eigen, D., Zhang, X., Mathieu, M., Fergus, R., LeCun, Y.: OverFeat: integrated recognition, localization and detection using convolutional networks. *arXiv preprint arXiv:1312.6229* (2013)
- Girshick, R.: Fast R-CNN. In: Proceedings of International Conference on Computer Vision (ICCV 2015), pp. 1440–1448 (2015)

40. Lin, T.Y., Dollár, P., Girshick, R., He, K., Hariharan, B., Belongie, S.: Feature pyramid networks for object detection. In: Proceedings of Computer Vision and Pattern Recognition (CVPR 2017) (2017)
41. Lin, T.Y., Goyal, P., Girshick, R., He, K., Dollár, P.: Focal loss for dense object detection. In: Proceedings of International Conference on Computer Vision (ICCV 2017), pp. 2980–2988 (2017)
42. He, K., Gkioxari, G., Dollár, P., Girshick, R.: Mask R-CNN. In: Proceedings of International Conference on Computer Vision (ICCV 2017), pp. 2961–2969 (2017)
43. Hu, P., Ramanan, D.: Finding tiny faces. In: Proceedings of Computer Vision and Pattern Recognition (CVPR 2017), pp. 951–959 (2017)
44. Ren, S., He, K., Girshick, R., Zhang, X., Sun, J.: Object detection networks on convolutional feature maps. *IEEE Trans. Pattern Anal. Mach. Intell.* **39**, 1476–1481 (2016)
45. Nah, S., Kim, T.H., Lee, K.M.: Deep multi-scale convolutional neural network for dynamic scene deblurring. In: Proceedings of Computer Vision and Pattern Recognition (CVPR 2017), pp. 3883–3891 (2017)
46. Chen, L.C., Papandreou, G., Kokkinos, I., Murphy, K., Yuille, A.L.: DeepLab: semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected CRFs. *IEEE Trans. Pattern Anal. Mach. Intell.* **40**, 834–848 (2017)
47. Yang, F., Choi, W., Lin, Y.: Exploit all the layers: fast and accurate CNN object detector with scale dependent pooling and cascaded rejection classifiers. In: Proceedings of Computer Vision and Pattern Recognition (CVPR 2016), pp. 2129–2137 (2016)
48. Cai, Z., Fan, Q., Feris, R.S., Vasconcelos, N.: A unified multi-scale deep convolutional neural network for fast object detection. In: Proceedings of European Conference on Computer Vision (ECCV 2016), vol. 9908 of Springer LNCS, pp. 354–370 (2016)
49. Yu, F., Koltun, V.: Multi-scale context aggregation by dilated convolutions. In: International Conference on Learning Representations (ICLR 2016) (2016)
50. Yu, F., Koltun, V., Funkhouser, T.: Dilated residual networks. In: Proceedings of Computer Vision and Pattern Recognition (CVPR 2017), pp. 472–480 (2017)
51. Mehta, S., Rastegari, M., Caspi, A., Shapiro, L., Hajishirzi, H.: ESPNet: efficient spatial pyramid of dilated convolutions for semantic segmentation. In: Proceedings of European Conference on Computer Vision (ECCV 2018), pp. 552–568 (2018)
52. Zhang, R., Tang, S., Zhang, Y., Li, J., Yan, S.: Scale-adaptive convolutions for scene parsing. In: Proceedings of International Conference on Computer Vision (ICCV 2017), pp. 2031–2039 (2017)
53. Wang, H., Kembhavi, A., Farhadi, A., Yuille, A.L., Rastegari, M.: ELASTIC: improving CNNs with dynamic scaling policies. In: Proceedings of Computer Vision and Pattern Recognition (CVPR 2019), pp. 2258–2267 (2019)
54. Chen, Y., Fang, H., Xu, B., Yan, Z., Kalantidis, Y., Rohrbach, M., Yan, S., Feng, J.: Drop an octave: Reducing spatial redundancy in convolutional neural networks with octave convolution. In: Proceedings of International Conference on Computer Vision (ICCV 2019) (2019)
55. Iijima, T.: Basic theory on normalization of pattern (in case of typical one-dimensional pattern). *Bull. Electrotech. Lab.* **26**, 368–388 (1962). (in Japanese)
56. Witkin, A.P.: Scale-space filtering. In: Proceedings of 8th International Joint Conference on Artificial Intelligence, Karlsruhe, Germany, pp. 1019–1022 (1983)
57. Koenderink, J.J.: The structure of images. *Biol. Cybern.* **50**, 363–370 (1984)
58. Babaud, J., Witkin, A.P., Baudin, M., Duda, R.O.: Uniqueness of the Gaussian kernel for scale-space filtering. *IEEE Trans. Pattern Anal. Mach. Intell.* **8**, 26–33 (1986)
59. Koenderink, J.J., van Doorn, A.J.: Generic neighborhood operators. *IEEE Trans. Pattern Anal. Mach. Intell.* **14**, 597–605 (1992)
60. Lindeberg, T.: Scale-Space Theory in Computer Vision. Springer, New York (1993)
61. Lindeberg, T.: Scale-space theory: a basic tool for analysing structures at different scales. *J. Appl. Stat.* **21**, 225–270 (1994)
62. Florack, L.M.J.: Image Structure. Series in Mathematical Imaging and Vision, Springer, New York (1997)
63. Weickert, J., Ishikawa, S., Imiya, A.: Linear scale-space has first been proposed in Japan. *J. Math. Imaging Vis.* **10**, 237–252 (1999)
64. ter Haar Romeny, B.: Front-End Vision and Multi-Scale Image Analysis. Springer, New York (2003)
65. Lindeberg, T.: Generalized Gaussian scale-space axiomatics comprising linear scale-space, affine scale-space and spatio-temporal scale-space. *J. Math. Imaging Vis.* **40**, 36–81 (2011)
66. Lindeberg, T.: A computational theory of visual receptive fields. *Biol. Cybern.* **107**, 589–635 (2013)
67. Bruna, J., Mallat, S.: Invariant scattering convolution networks. *IEEE Trans. Pattern Anal. Mach. Intell.* **35**, 1872–1886 (2013)
68. Sifre, L., Mallat, S.: Rotation, scaling and deformation invariant scattering for texture discrimination. In: Proceedings of Computer Vision and Pattern Recognition (CVPR 2013), pp. 1233–1240 (2013)
69. Oyallon, E., Mallat, S.: Deep roto-translation scattering for object classification. In: Proceedings of Computer Vision and Pattern Recognition (CVPR 2015), pp. 2865–2873 (2015)
70. Jacobsen, J.J., van Gemert, J., Lou, Z., Smeulders, A.W.M.: Structured receptive fields in CNNs. In: Proceedings of Computer Vision and Pattern Recognition (CVPR 2016), pp. 2610–2619 (2016)
71. Luan, S., Chen, C., Zhang, B., Han, J., Liu, J.: Gabor convolutional networks. *IEEE Trans. Pattern Anal. Mach. Intell.* **27**, 4357–4366 (2018)
72. Shelhamer, E., Wang, D., Darrell, T.: Blurring the line between structure and learning to optimize and adapt receptive fields. arXiv preprint [arXiv:1904.11487](https://arxiv.org/abs/1904.11487) (2019)
73. Henriques, J.F., Vedaldi, A.: Warped convolutions: efficient invariance to spatial transformations. *Int. Conf. Mach. Learn.* **70**, 1461–1469 (2017)
74. Esteves, C., Allen-Blanchette, C., Zhou, X., Daniilidis, K.: Polar transformer networks. In: International Conference on Learning Representations (ICLR 2018) (2018)
75. Poggio, T.A., Anselmi, F.: Visual Cortex and Deep Networks: Learning Invariant Representations. MIT Press, Cambridge (2016)
76. Lapev, D., Savinov, N., Buhmann, J.M., Pollefeys, M.: TI-pooling: transformation-invariant pooling for feature learning in convolutional neural networks. In: Proceedings of Computer Vision and Pattern Recognition (CVPR 2016), pp. 289–297 (2016)
77. Kondor, R., Trivedi, S.: On the generalization of equivariance and convolution in neural networks to the action of compact groups. In: International Conference on Machine Learning (ICML 2018) (2018)
78. Lindeberg, T.: Normative theory of visual receptive fields. *Heliyon* **7**(e05897), 1–20 (2021)
79. Roux, N.L., Bengio, Y.: Continuous neural networks. In: Artificial Intelligence and Statistics (AISTATS 2007), vol. 2 of Proceedings of Machine Learning Research, pp. 404–411 (2007)
80. Wang, S., Suo, S., Ma, W.C., Pokrovsky, A., Urtasun, R.: Deep parametric continuous convolutional neural networks. In: Proceedings of Computer Vision and Pattern Recognition (CVPR 2018), pp. 2589–2597 (2018)
81. Wu, W., Qi, Z., Fuxin, L.: PointConv: deep convolutional networks on 3D point clouds. In: Proceedings of Computer Vision and Pattern Recognition (CVPR 2019), pp. 9621–9630 (2019)

82. Shocher, A., Feinstein, B., Haim, N., Irani, M.: From discrete to continuous convolution layers. arXiv preprint [arXiv:2006.11120](https://arxiv.org/abs/2006.11120) (2020)
83. Duits, R., Smets, B., Bekkers, E., Portegies, J.: Equivariant deep learning via morphological and linear scale space PDEs on the space of positions and orientations. In: Proceedings of Scale Space and Variational Methods in Computer Vision (SSVM 2021), vol. 12679 of Springer LNCS, pp. 27–39 (2021)
84. L. Ruthotto and E. Haber Deep neural networks motivated by partial differential equations. *J. Math. Imaging Vis.* **62**, 352–364 (2020)
85. Z. Shen and L. He and Z. Lin and J. Ma Partial differential operator based equivariant convolutions. In: International Conference on Machine Learning (ICML 2020). 8697–8706 (2020)
86. Duits, R., Florack, L., de Graaf, J., ter Haar Romeny, B.: On the axioms of scale space theory. *J. Math. Imaging Vis.* **22**, 267–298 (2004)
87. Lindeberg, T.: Invariance of visual operations at the level of receptive fields. *PLoS ONE* **8**, e66990 (2013)
88. Lindeberg, T.: On the axiomatic foundations of linear scale-space. In: Sporring, J., Nielsen, M., Florack, L., Johansen, P. (eds.) *Gaussian Scale-Space Theory: Proceedings*, pp. 75–97. PhD School on Scale-Space Theory, Copenhagen, Denmark, Springer (1996)
89. Pauwels, E.J., Fiddelaers, P., Moons, T., van Gool, L.J.: An extended class of scale-invariant and recursive scale-space filters. *IEEE Trans. Pattern Anal. Mach. Intell.* **17**, 691–701 (1995)
90. Felsberg, M., Sommer, G.: The monogenic scale-space: a unifying approach to phase-based image processing in scale-space. *J. Math. Imaging Vis.* **21**, 5–26 (2004)
91. Koenderink, J.J., van Doorn, A.J.: Representation of local geometry in the visual system. *Biol. Cybern.* **55**, 367–375 (1987)
92. Lindeberg, T.: Dense scale selection over space, time and space-time. *SIAM J. Imag. Sci.* **11**, 407–441 (2018)
93. Valois, R.L.D., Cottaris, N.P., Mahon, L.E., Elfer, S.D., Wilson, J.A.: Spatial and temporal receptive fields of geniculate and cortical cells and directional selectivity. *Vis. Res.* **40**, 3685–3702 (2000)
94. LeCun, Y., Bottou, L., Bengio, Y., Haffner, P.: Gradient-based learning applied to document recognition. *Proc. IEEE* **86**, 2278–2324 (1998)
95. Paszke, A., Gross, S., Chintala, S., Chanan, G., Yang, E., De Vito, Z., Lin, Z., Desmaison, A., Antiga, L., Lerer, A.: Automatic differentiation in PyTorch. In: Proceedings of Neural Information Processing Systems (NIPS 2017) (2017)
96. Kingma, P.D., Ba, J.: Adam: a method for stochastic optimization. In: International Conference for Learning Representations (ICLR 2015) (2015)
97. Lindeberg, T.: Scale-space for discrete signals. *IEEE Trans. Pattern Anal. Mach. Intell.* **12**, 234–254 (1990)
98. Lindeberg, T.: Discrete derivative approximations with scale-space properties: a basis for low-level feature extraction. *J. Math. Imaging Vis.* **3**, 349–376 (1993)
99. Lindeberg, T.: Scale selection. In: Ikeuchi, K. (ed.) *Computer Vision*. Springer, Berlin (2021). https://doi.org/10.1007/978-3-030-03243-2_242-1
100. Loog, M., Li, Y., Tax, D.M.J.: Maximum membership scale selection. In: *Multiple Classifier Systems*, vol. 5519 of Springer LNCS, pp. 468–477 (2009)
101. Li, Y., Tax, D.M.J., Loog, M.: Scale selection for supervised image segmentation. *Image Vis. Comput.* **30**, 991–1003 (2012)

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Tony Lindeberg is a Professor of Computer Science at KTH Royal Institute of Technology in Stockholm, Sweden. He was born in Stockholm in 1964, received his M.Sc. degree in 1987, his Ph.D. degree in 1991, became docent in 1996, and was appointed professor in 2000. He was a Research Fellow at the Royal Swedish Academy of Sciences between 2000 and 2010. His research interests in computer vision relate to scale-space representation, image features, object recognition, spatio-temporal recognition, video analysis, deep networks and computational modelling of biological vision. He has developed theories and methodologies for continuous and discrete scale-space representation, visual and auditory receptive fields, hierarchical and deep networks, detection of salient image structures, automatic scale selection, scale-invariant image features, affine invariant features, affine and Galilean normalization, temporal, spatio-temporal and spectro-temporal scale-space concepts as well as spatial and spatio-temporal image descriptors for image-based recognition. He has also worked on topics in medical image analysis and gesture recognition. He is author of the book *Scale-Space Theory in Computer Vision*.