Provably Scale-Covariant Continuous Hierarchical Networks Based on Scale-Normalized Differential Expressions Coupled in Cascade

Tony Lindeberg¹

Received: 12 March 2019 / Accepted: 9 October 2019 / Published online: 25 October 2019
© The Author(s) 2019

Abstract

This article presents a theory for constructing hierarchical networks in such a way that the networks are guaranteed to be provably scale covariant. We first present a general sufficiency argument for obtaining scale covariance, which holds for a wide class of networks defined from linear and nonlinear differential expressions expressed in terms of scale-normalized scale-space derivatives. Then, we present a more detailed development of one example of such a network constructed from a combination of mathematically derived models of receptive fields and biologically inspired computations. Based on a functional model of complex cells in terms of an oriented quasi quadrature combination of first- and second-order directional Gaussian derivatives, we couple such primitive computations in cascade over combinatorial expansions over image orientations. Scale-space properties of the computational primitives are analysed, and we give explicit proofs of how the resulting representation allows for scale and rotation covariance. A prototype application to texture analysis is developed, and it is demonstrated that a simplified mean-reduced representation of the resulting QuasiQuadNet leads to promising experimental results on three texture datasets.

Keywords Handcrafted · Structured · Deep · Hierarchical · Network · Scale covariance · Scale invariance · Differential expression · Quasi quadrature · Complex cell · Feature detection · Texture analysis · Texture classification · Scale space · Deep learning · Computer vision

1 Introduction

The recent progress with deep learning architectures [1–10] has demonstrated that hierarchical feature representations over multiple layers have higher potential compared to approaches based on single layers of receptive fields.

Although theoretical and empirical advances are being made [11–18], we currently lack a comparable understanding of the nonlinearities in deep networks in the way that scale-space theory provides a deep understanding of early visual receptive fields. Training deep networks is still very much of an art [19]. Moreover, deep nets sometimes perform serious errors. The observed problem with adversarial exam-

ples [20–26] can be taken as an indication that current deep nets may not solve the same type of problem as one might at first expect them to do. For these reasons, it is of interest to develop theoretically principled approaches to capture nonlinear hierarchical relations between image structures at different scales as an extension of the regular scale-space paradigm.

A specific limitation of current deep nets is that they are not truly scale covariant. A deep network constructed by repeated application of compact 3×3 or 5×5 kernels, such as AlexNet [1], VGG-Net [2] or ResNet [5], implies an implicit assumption of a preferred size in the image domain as induced by the discretization in terms of local 3×3 or 5×5 kernels of a fixed size. Spatial max pooling over image neighbourhoods of fixed size, such as over 2×2 neighbourhoods over multiple layers, also implies that nonlinearities are applied relative to a fixed grid spacing. Thereby, due to the nonlinearities in the deep net, the output from the network may be qualitatively different depending on the specific size of the object in the image domain, as varying because of, e.g., different distances between the object and the observer. To

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

handle this lack of scale covariance, approaches have been developed such as spatial transformer networks [27], using sets of subnetworks in a multi-scale fashion [28] or by combining deep nets with image pyramids [29]. Since the size normalization performed by a spatial transformer network is not guaranteed to be truly scale covariant, and since traditional image pyramids imply a loss of image information that can be interpreted as corresponding to undersampling, it is of interest to develop continuous approaches for deep networks that guarantee true scale covariance or better approximations thereof.

An argument that we want to put forward in this article is that truly scale-covariant deep networks with their associated extended notion of truly scale-invariant networks may be conceptually much easier to achieve if we set aside the issues of spatial sampling in the first modelling stage and model the transformations between adjacent layers in the deep network as continuous translation-covariant operators as opposed to discrete filters. Specifically, we will propose to combine concepts from hierarchical families of CNNs with scale-space theory to define continuous families of hierarchical networks, with each member of the family being a rescaled copy of the base network, in a corresponding way as an input image is embedded into a one-parameter family of images, with scale as the parameter, within the regular scale-space framework. Then, a structural advantage of a continuous model as compared to a discrete model is that it can guarantee provable scale covariance in the following way: if the computational primitives that are used for defining a hierarchical network are defined in a multi-scale manner, e.g. from Gaussian derivatives and possibly nonlinear differential expressions constructed from these, and if the scale parameters of the primitives in the higher layers are proportional to the scale parameter in the first layer, then if we define a multi-scale hierarchical network over all the scale parameters in the first layer, the multi-scale network is guaranteed to be truly scale covariant.

This situation is in contrast to the way most deep nets are currently constructed, as a combination of discrete primitives whose scales are instead proportional to the grid spacing. That in turn implies a preferred scale of the computations and which will violate scale covariance unless the image data are resampled to multiple rescaled copies of the input image prior to being used as input to a deep net. If using such spatial resampling to different levels of resolution, then, however, it may be harder to combine information between different multi-scale channels compared to using a continuous model that preserves the same spatial sampling in the input data. Rescaling of the image data prior to later stage processing may also introduce sampling artefacts.

The subject of this article is to first present a general sufficiency argument for constructing provably scale-covariant hierarchical networks based on a spatially continuous model

of the transformations between adjacent layers in the hierarchy. This sufficiency result holds for a very wide class of possible continuous hierarchical networks. Then, we will develop in more detail one example of such a continuous network for capturing nonlinear hierarchical relations between features over multiple scales.

Building upon axiomatic modelling of visual receptive fields in terms of Gaussian derivatives and affine extensions thereof, which can serve as idealized models of simple cells in the primary visual cortex [30–33], we will propose a functional model for complex cells in terms of an oriented quasi quadrature measure, which combines first- and second-order directional affine Gaussian derivatives according to an energy model [34–37]. Compared to earlier approaches of related types [38–42], our quasi quadrature model has the conceptual advantage that it is expressed in terms of scale-space theory in addition to well reproducing properties of complex cells as reported by [34,43–45]. Thereby, this functional model of complex cells allows for a conceptually easy integration with transformation properties, specifically truly provable scale covariance, or a generalization to affine covariance provided that the receptive field responses are computed in terms of affine Gaussian derivatives as opposed to regular Gaussian derivatives.

Then, we will combine such oriented quasi quadrature measures in cascade, building upon the early idea of Fukushima [38] of using Hubel and Wiesel's findings regarding receptive fields in the primary visual cortex [46–48] to build a hierarchical neural network from repeated application of models of simple and complex cells. This will result in a handcrafted network, termed quasi quadrature network, with structural similarities to the scattering network proposed by Bruna and Mallat [41], although expressed in terms of Gaussian derivatives instead of Morlet wavelets.

We will show how the scale-space properties of the quasi quadrature primitive in this representation can be theoretically analysed and how the resulting handcrafted network becomes provably scale covariant and rotation covariant, in such a way that the multi-scale and multi-orientation network commutes with scaling transformations and rotations in the spatial image domain.

As a proof of concept that the proposed methodology can lead to meaningful results, we will experimentally investigate a prototype application to texture classification based on a substantially simplified representation that uses just the average values over image space of the resulting QuasiQuadNet. It will be demonstrated that the resulting approach leads to competitive results compared to classical texture descriptors as well as to other handcrafted networks.

Specifically, we will demonstrate that in the presence of substantial scaling transformations between the training data and the test data, true scale covariance substantially improves

the ability to perform predictions or generalizations beyond the variabilities that are spanned by the training data.

1.1 Structure of this Article

Section 2 begins with an overview of related work, with emphasis on related scale-space approaches, deep learning approaches somehow related to scale, rotation-covariant deep networks, biologically inspired networks, other handcrafted or structured networks including other hybrid approaches between scale space and deep learning.

As a general motivation for studying hierarchical networks that are based on primitives that are continuous over image space, Sect. 3 then presents a general sufficiency argument that guarantees provable scale covariance for a very wide class of networks defined from layers of scale-space operations coupled in cascade.

To provide an additional theoretical basis for a subclass of such networks that we shall study in more detail in this article, based on functional models of complex cells coupled in cascade, Sect. 4 describes a quasi quadrature measure over a purely 1D signal, which measures the energy of first- and second-order Gaussian derivative responses. Theoretical properties of this entity are analysed with regard to scale selectivity and scale selection properties, and we show how free parameters in the quasi quadrature measure can be determined from closed-form calculations.

In Sect. 5, an oriented extension of the 1D quasi quadrature measure is presented over multiple orientations in image space and is proposed as a functional model that mimics some of the known properties of complex cells, while at the same time being based on axiomatically derived affine Gaussian derivatives that well model the functional properties of simple cells in the primary visual cortex.

In Sect. 6, we propose to couple such quasi quadrature measures in cascade, leading to a class of hierarchical networks based on scale-space operations that we term quasi quadrature networks. We give explicit proofs of scale covariance and rotational covariance of such networks and show examples of the type of information that can be captured in different layers in the hierarchies.

Section 7 then outlines a prototype application to texture analysis based on a substantially mean-reduced version of such a quasi quadrature network, with the feature maps in the different layers reduced to just their mean values over image space. By experiments on three datasets for texture classification, we show that this approach leads to promising results that are comparable or better than other handcrafted networks or more dedicated handcrafted texture descriptors. We do also present experiments of scale prediction or scale generalization, which quantify the performance over scaling transformations for which the variabilities in the testing data

are not spanned by corresponding variabilities in the training data.

Finally, Sect. 8 concludes with a summary and discussion.

1.2 Relations to Previous Contribution

This paper constitutes a substantially extended version of a conference paper presented at the SSVM 2019 conference [49] and with substantial additions concerning:

- the motivations underlying the developments of this work and the importance of scale covariance for deep networks (Sect. 1),
- a wider overview of related work (Sect. 2),
- the formulation of a general sufficiency result to guarantee scale covariance of hierarchical networks constructed from computational primitives (linear and nonlinear filters) formulated based on scale-space theory (Sect. 3),
- additional explanations regarding the quasi quadrature measure (Sect. 4) and its oriented affine extension to model functional properties of complex cells (Sect. 5),
- better explanation of the quasi quadrature network constructed by coupling oriented quasi quadrature measures in cascade, including a figure illustration of the network architecture, details of discrete implementation, issues of exact versus approximate covariance or invariance in a practical implementation and experimental results showing examples of the type of information that is computed in different layers of the hierarchy (Sect. 6),
- a more extensive experimental section showing the results of applying a mean-reduced QuasiQuadNet for texture classification, including additional experiments demonstrating the importance of scale covariance and better overall descriptions about the experiments that could not be given in the conference paper because of the space limitations (Sect. 7).

In relation to the SSVM 2019 paper, this paper therefore gives a more general treatment about the notion of scale covariance of more general validity for continuous hierarchical networks, presents more experimental results regarding the prototype application to texture classification and gives overall better descriptions of the subjects treated in the paper, including more extensive references to related literature.

2 Related Work

In the area of scale-space theory, theoretical results have been derived showing that Gaussian kernels and Gaussian derivatives constitute a canonical class of linear receptive fields for an uncommitted vision system [30,31,50–62]. The conditions that specify this uniqueness property are basically

linearity, shift invariance and regularity properties combined with different ways of formalizing the notion that new structures should not be created from finer to coarser scales in a multi-scale representation.

The receptive field responses obtained by convolution with such Gaussian kernels and Gaussian derivatives are truly scale covariant—a property that has been used for designing a large number of scale-covariant and scale-invariant feature detectors and image descriptors [36,63–71]. With the generalization to affine covariance and affine invariance based on the notion of affine scale-space [51,56,66,72,73], these theoretical developments served as a conceptual foundation that opened up for a very successful track of methodology development for image-based matching and recognition in classical computer vision.

In the area of deep learning, approaches to tackle the notion of scale have been developed in different ways. By augmenting the training images with multiple rescaled copies of each training image or by randomly resizing the training images over some scale range (scale jittering), the robustness of a deep net can usually be extended to moderate scaling factors [2,74]. Another basic data-driven approach consists of training a module to estimate spatial scaling factors from the data by a spatial transformer network [27,75]. A more structural approach consists of applying deep networks to multiple layers in an image pyramid [29,76–78], or using some other type of multi-channel approach where the input image is rescaled to different resolutions, possibly combined with interactions or pooling between the layers [79–82]. Variations or extensions of this approach include scale-dependent pooling [83], using sets of subnetworks in a multi-scale fashion [28], using dilated convolutions [84–86], scale-adaptive convolutions [87] or adding additional branches of down-samplings and/or up-samplings in each layer of the network [88,89].

A more specific approach to designing a scale-covariant network is by spatially warping the image data prior to image filtering by a log-polar transformation [90,91]. Then, the scaling and rotation transformations are mapped to mere translations in the transformed domain, although this property only holds provided that the origin of the log-polar transformation can be preserved between the training data and the testing data. Specialized learning approaches for scale-covariant or affine-covariant feature detection have been developed for interest point detection [92,93].

There is a large literature on approaches to achieve rotation-covariant networks [94–103] with applications to different domains including astronomy [104], remote sensing [105], medical image analysis [106,107] and texture classification [108]. There are also approaches to invariant networks based on formalism from group theory [42,109,110].

In the context of more general classes of image transformations, it is worth noting that beyond the classes of

spatial scaling transformations and spatial affine transformations (including rotations), the framework of generalized axiomatic scale-space theory [111,112] does also allow for covariance and/or invariance with regard to temporal scaling transformations [113], Galilean transformations and local multiplicative intensity transformations [32,33].

Concerning biologically inspired neural networks, Fukushima [38] proposed to build upon Hubel and Wiesel's findings regarding receptive fields in the primary visual cortex (see [48]) to construct a hierarchical neural network from repeated application of models of simple and complex cells. Poggio and his co-workers built on this idea and constructed handcrafted networks based on two layers of such models expressed in terms of Gabor functions [39,40,114].

The approach of scattering convolution networks [41, 115,116] is closely related, where directional odd and even wavelet responses are computed and combined with a non-linear modulus (magnitude) operator over a set of different orientations in the image domain and over a hierarchy over a dyadic set of scales.

Other types of handcrafted or structured networks have been constructed by applying principal component analysis in cascade [117] or by using Gabor functions as primitives to be modulated by learned filters [118].

Concerning hybrid approaches between scale space and deep learning, Jacobsen et al. [119] construct a hierarchical network from learned linear combinations of Gaussian derivative responses. Shelhamer et al. [120] compose free-form filters with affine Gaussian filters to adapt the receptive field size and shape to the image data.

Concerning the use of a continuous model of the transformation from the input data to the output data in a hierarchical computation structure, which we will here develop for deep networks from motivations of making it possible for the network to fulfil geometric transformation properties in spatial input data, such a notion of continuous transformations from the input to the output has been proposed as a model for neural networks prior to the deep learning revolution by Le Roux and Bengio [121] from the viewpoint of an uncountable number of hidden units, and suggesting that that makes it possible for the network to represent some smooth functions more compactly.

For an overview of texture classification, which we shall later use as an application domain, we refer to the recent survey by Liu et al. [122] and the references therein.

In this work, we aim towards a conceptual bridge between scale-space theory and deep learning, with specific emphasis on handling the variability in image data caused by scaling transformations. We will show that it is possible to design a wide class of possible scale-covariant networks by coupling linear or nonlinear expressions in terms of Gaussian derivatives in cascade. As a proof of concept that such a construction can lead to meaningful results, we will present a

specific example of such a network, based on a mathematically and biologically motivated model of complex cells and demonstrate that it is possible to get quite promising performance on texture classification, comparable or better than many classical texture descriptors or other handcrafted networks. Specifically, we will demonstrate how the notion of scale covariance improves the ability to perform predictions or generalizations to scaling variabilities in the testing data that are not spanned by the training data.

We propose that this opens up for studying other hybrid approaches between scale-space theory and deep learning to incorporate explicit modelling of image transformations as a prior in hierarchical networks.

3 General Scale Covariance Property for Continuous Hierarchical Networks

For a visual observer that views a dynamic world, the size of objects in the image domain can vary substantially, because of variations in the distance between the objects and the observer and because of objects having physically different sizes in the world. If we rescale an image pattern by a uniform scaling factor, we would in general like the perception of objects in the underlying scene to be preserved.¹ A natural precursor to achieving such a scale-invariant perception of the world is to have a scale-covariant image representation. Specifically, a scale-covariant image representation can often be used as a basis for constructing scale-invariant image descriptors and/or scale-invariant recognition schemes.

In the area of scale-space theory [30,50,52,53,56,57,59,61], theoretically well-founded approaches have been developed to handle the notion of scale in image data and to construct scale-covariant and scale-invariant image representations [36,63–69,71,111]. In this section, we will present a general argument of how these notions can be extended

to construct provably scale-covariant hierarchical networks, based on continuous models of the image operations between adjacent layers.

Given an image $f(x)$, consider a multi-scale representation $L(x; s)$ constructed by Gaussian convolution and then from this scale-space representation defining a family of scale-parameterized possibly nonlinear operators $\mathcal{D}_{1,s_1}$ over a continuum of scale parameters s_1 :

$$F_1(\cdot; s_1) = (\mathcal{D}_{1,s_1} f)(\cdot), \quad (1)$$

where the effect of the Gaussian smoothing operation is incorporated in the operator $\mathcal{D}_{1,s_1}$.

Within the framework of Gaussian scale-space representation [30,50,52,53,56,57,59,61], we could consider these operators as being formed from sufficiently homogeneous possibly nonlinear combinations of Gaussian derivative operators, such that they under a rescaling of the input domain $x' = Sx$ by a factor of S are guaranteed to obey the scale covariance property:

$$F'_1(x'; s'_1) = S^{\alpha_1} F_1(x; s_1) \quad (2)$$

for some constant α_1 and some transformation of the scale parameters $s'_1 = \phi_1(s_1)$. In other words, for any image representation computed over the original image domain x at scale s_1 , it should be possible to find a corresponding representation over the transformed domain $x' = Sx$ at scale s'_1 with a possibly transformed magnitude as determined by the relative amplification factor S^{α_1} .

By concatenating a set of corresponding scale-parameterized differential operators $\mathcal{D}_{k,s_k}$

$$F_k(\cdot; s_1, \dots, s_{k-1}, s_k) = \mathcal{D}_{k,s_k} F_{k-1}(\cdot; s_1, \dots, s_{k-1}) \quad (3)$$

that obey similar scale covariance properties such that

$$F'_k(x'; s'_1, \dots, s'_k) = S^{\alpha_1} \dots S^{\alpha_k} F_k(x; s_1, \dots, s_k), \quad (4)$$

it follows that the combined hierarchical network is guaranteed to be provably scale covariant, see Figs. 1 and 2 for schematic illustrations. Specifically, it is natural to choose the scale parameters s_k in the higher layers proportional to the scale parameter s_1 in the first layer to guarantee scale covariance.

More generally, we could also consider constructing scale-covariant networks from other types of scale-covariant operators that obey similar scaling properties as in Eqs. (2) and (4), for example, expressed in terms of a basis of rescaled Gabor functions or a family of continuously rescaled wavelets. Then, however, the information reducing properties from finer to coarser scales in the representation computed by Gaussian convolution and Gaussian derivatives are, however,

¹ When rescaling an object in the image domain, there are three main scaling effects occurring: (i) how large the object will be in the image domain, (ii) how large the image structures of the object will be relative to the resolution of the image sensor and (iii) how large the object will be relative to the outer dimensions (the size) of the image sensor. In this article, we focus primarily on the first effect, to design mechanisms for achieving scale covariance and scale invariance under variations in the apparent size of objects in the image domain, assuming that the resolution as well as the size of the image sensor is sufficient to sufficiently resolve the interesting image structures over the scale range we are interested in covering. In a practical implementation, the resolution of the image data will additionally imply a lower bound on how fine scale levels can be computed (the inner scale). The size of the image sensor will also impose an additional upper bound on how large objects can be captured (the outer scale). While such effects may also be highly important with regard to a specific application, the topic of this article concerns how to handle the essential geometric effects of the image transformations, leaving more detailed issues of image sampling and handling of image boundaries for future work.

Fig. 1 Commutative diagram for a scale-covariant hierarchical network constructed according to the presented sufficiency result. Provided that the individual differential operators $\mathcal{D}_{k,s_k}$ between adjacent layers are scale covariant, which, for example, holds for the class of homogeneous differential expressions of the form (12) as well as self-similar compositions of such operations that additionally satisfy corresponding homogeneity requirements, it follows that it will be possible to perfectly match the corresponding layers F_k and F'_k under a scaling transformations of the underlying image domain $f'(x') = f(x)$ for $x' = Sx$, provided that the scale parameter s_k in layer k is proportional to the scale parameter s_1 in the first layer, $s_k = r_k^2 s_1$, for some scalar constants r_k . For such a network constructed from scale-space operations based on the Gaussian scale-space theory framework, the scale parameters in the two domains should be related according to $s'_k = S^2 s_k$

not guaranteed to hold. As mentioned above, the Gaussian kernel and the Gaussian derivatives can be uniquely determined from different ways of formalizing the requirement that they should not introduce new image structures from finer to coarser scales in a multi-scale representation [30,31,50–62].

In this overall structure, there is a large flexibility in how to choose the operators $\mathcal{D}_{k,s_k}$. Within the family of operators defined from a scale-space representation, we could consider a large class of differential expressions and differential invariants in terms of scale-normalized Gaussian derivatives [36] that guarantee provable scale covariance.

For example, if we choose to express the first differential operator $\mathcal{D}_{1,s_1}$ in a basis in terms of scale-normalized derivatives [36] (here with the multi-index notation $\partial_x^n = \partial_{x_1^{n_1} \dots x_D^{n_D}}$ for the partial derivatives in D dimensions and $|n| = n_1 + \dots + n_D$)

$$\partial_{\xi^n} = \partial_{x^n, \gamma\text{-norm}} = s^{|n|\gamma/2} \partial_x^n \quad (5)$$

computed from a scale-space representation of the input signal

$$L_1(\cdot; s_1) = g(\cdot; s_1) * f(\cdot) \quad (6)$$

Fig. 2 A hierarchical network defined by coupling scale-covariant differential expressions formulated within the continuous scale-space framework will be guaranteed to be provably scale covariant provided that the scale parameters in higher layers s_k for $k \geq 2$ are proportional to the scale parameter s_1 in the first layer. If the scale normalization parameter γ in the scale-normalized derivative expressions is equal to one, then general differential expressions in terms of such derivatives can be used based on the transformation property (9). If the scale normalization parameter γ is not equal to one, then one can take homogeneous polynomial differential expressions of the form (12) as well as self-similar transformations of such expressions. (In this schematic illustration, the arguments of the layers F_1, F_2, F_{k-1} and F_k , which should be $F_1(\cdot; s_1), F_2(\cdot; s_1, s_2), F_{k-1}(\cdot; s_1, s_2, \dots, s_{k-1})$ and $F_k(\cdot; s_1, s_2, \dots, s_{k-1}, s_k)$, respectively, have been suppressed to simplify the notation. The argument of the input data f should be $f(\cdot)$.)

by convolution with Gaussian kernels

$$g(x; s) = \frac{1}{(2\pi s)^{D/2}} e^{-\frac{|x|^2}{2s}} \quad (7)$$

and with s in (5) determined from s_1 in (6), it then follows that under a rescaling of the image domain $f'(x') = f(x)$ for $x' = Sx$ the scale-normalized derivatives transform according to [36, Eq. (20)]

$$\partial_{\xi^n} L'_1(x'; s'_1) = S^{|n|(\gamma-1)} \partial_{\xi^n} L_1(x; s_1) \quad (8)$$

provided that the scale parameters are matched according to $s'_1 = S^2 s_1$. Specifically, in the special case of choosing $\gamma = 1$, the scale-normalized derivatives will be equal

$$\partial_{\xi^n} L'_1(x'; s'_1) = \partial_{\xi^n} L_1(x; s_1). \quad (9)$$

This implies that any scale-parameterized differential operator $\mathcal{D}_{1,s_k}$ that can be expressed as a sufficiently regular function ψ

$$\mathcal{D}_{1,s_1} f = \psi(\mathcal{J}_{N,s_1} L_1) \quad (10)$$

of the scale-normalized N -jet, $\mathcal{J}_{N,s_1} L_1$, of the scale-space representation L_1 of the input image, which is the union of

all partial derivatives up to order N

$$\mathcal{J}_{N,s_1} L_1 = \cup_{1 \leq |n| \leq N} \partial_{\xi^n} L_1, \quad (11)$$

will satisfy the scale covariance property (2) for $\alpha_1 = 0$. More generally, it is not necessary that all the derivatives are computed at the same scale, although such a choice could possibly be motivated from conceptual simplicity.

In the less specific case of choosing $\gamma \neq 1$, we can consider homogeneous polynomials of scale-normalized derivatives of the form

$$\mathcal{D}_{s_1} f = \sum_{i=1}^I c_i \prod_{j=1}^J L_{1,x^{\beta_{ij}}} \quad (12)$$

for which the sum of the orders of differentiation in a certain term

$$\sum_{j=1}^J |\beta_{ij}| = M \quad (13)$$

does not depend on the index i of that term. The corresponding scale-normalized expression with the regular spatial derivatives replaced by γ -normalized derivatives is

$$\mathcal{D}_{\gamma\text{-norm},s_1} f = s_1^{M\gamma/2} \mathcal{D}_{s_1} f \quad (14)$$

and transforms according to [36, Eq. (25)]

$$\mathcal{D}'_{\gamma\text{-norm},s'_1} f' = S^{M(\gamma-1)} \mathcal{D}_{\gamma\text{-norm},s_1} f \quad (15)$$

under any scaling transformation $f'(x') = f(x)$ for $x' = Sx$ provided that the scale levels are appropriately matched $s'_1 = S^2 s_1$. Such a self-similar form of scaling transformation will also be preserved under self-similar transformations $z \mapsto z^\delta$ of such expressions as well as for a rich family of polynomial combinations as well as rational expressions of such expressions as long as the scale covariance property (2) is preserved.

A natural complementary argument to constrain such self-similar compositions is to preserve the dimensionality of the image data, such that each layer F_k has the same dimensionality [intensity] as the input image f . If a polynomial is used for constructing a composed nonlinear differential expression $\mathcal{D}_{\text{comp},s_1} f$ by combinations of differential expressions of the form (14) and if this composed polynomial is a homogeneous polynomial of order P relative to the underlying partial derivatives ∂_{ξ^n} in the N -jet, in the sense that under a rescaling of the magnitude of the original image data f by a factor of β such that $f'(x') = \beta f(x)$ the differential expression transforms according to

$$\mathcal{D}'_{\text{comp},s'_1} f' = \beta^P \mathcal{D}_{\text{comp},s_1} f, \quad (16)$$

we should then transform that differential expression by the power $1/P$ to preserve the dimensionality of [intensity]. A similar argument applies to differential entities formed from rational expressions of differential expressions of the form (14) as long as the scale covariance property (2) is preserved.

Corresponding reasoning as done here regarding the transformation from the input image f to the first layer F_1 can be performed regarding the transformations $\mathcal{D}_{k,s_k}$ between any pairs of adjacent layers F_{k-1} and F_k , implying that if the differential operators $\mathcal{D}_{k,s_k}$ are chosen from similar families of differential operators as described above regarding the first differential operator $\mathcal{D}_{1,s_1}$, then the entire layered hierarchy will be scale covariant, provided that the scale parameter s_k in layer k is proportional to the scale parameter s_1 in the first layer, $s_k = r_k^2 s_1$, for some scalar constants r_k (see Fig. 1). This opens up for a large class of provably scale-covariant continuous hierarchical networks based on differential operators defined from the scale-space framework, where it remains to be determined which of these possible networks lead to desirable properties in other respects. In the following, we will develop one specific way of defining such a scale-covariant continuous network, by choosing these operators based on functional models of complex cells expressed within the Gaussian scale-space paradigm.

4 The Quasi Quadrature Measure Over a 1D Signal

Consider the scale-space representation [50,52,53,56,57,59, 61]

$$L(\cdot; s) = g(\cdot; s) * f(\cdot) \quad (17)$$

of a 1D signal $f(x)$ defined by convolution with Gaussian kernels

$$g(x; s) = \frac{1}{\sqrt{2\pi}s} e^{-\frac{x^2}{2s}} \quad (18)$$

and with scale-normalized derivatives according to [36]

$$\partial_{\xi^n} = \partial_{x^n, \gamma\text{-norm}} = s^{n\gamma/2} \partial_{x^n}. \quad (19)$$

In this section, we will describe a quasi quadrature entity that measures the local energy in the first- and second-order derivatives in the scale-space representation of a 1D signal and analyse its behaviour to image structures over multiple scales. Later in Sect. 5, an oriented extension of this measure to 2D image space will be used for expressing a functional model of complex cells that reproduces some of the known properties of complex cells.

4.1 Quasi Quadrature Measure in 1D

Motivated by the fact that the first-order derivatives primarily respond to the locally odd component of the signal, whereas the second-order derivatives primarily respond to the locally even component of a signal, it is natural to aim at a differential feature detector that combines locally odd and even components in a complementary manner. By specifically combining the first- and second-order scale-normalized derivative responses in a Euclidean way, we obtain a quasi quadrature measure of the form

$$Q_{x,\text{norm}}L = \sqrt{\frac{sL_x^2 + C s^2 L_{xx}^2}{s^\Gamma}} \quad (20)$$

as a modification² of the quasi quadrature measures previously proposed and studied in [36,37], with the scale normalization parameters γ_1 and γ_2 of the first- and second-order derivatives coupled according to $\gamma_1 = 1 - \Gamma$ and $\gamma_2 = 1 - \Gamma/2$ to enable scale covariance by adding derivative expressions of different orders only for the scale-invariant choice of $\gamma = 1$. This differential entity can be seen as an approximation³ of the notion of a quadrature pair of an odd and even filter [123] as more traditionally formulated based on a Hilbert transform [124, p. 267–272], while confined within the family of differential expressions based on Gaussian derivatives.

Intuitively, this quasi quadrature operator is intended as a measure of the amount of local changes in the signal, not specific to whether the dominant response comes from odd first-order derivatives or even second-order derivatives, and with additional scale selective properties as will be described later in Sect. 4.3.

If complemented by spatial integration, the components of the quasi quadrature measure are specifically related to the following class of energy measures over the frequency

² Compared to the previous work on quasi quadrature measures in [36, 37], we here transform use the previous 1D quasi quadrature measure by a square root function to maintain the same dimensionality as the input signal, which is a useful property when defining hierarchical networks by coupling quasi quadrature measures in cascade.

³ For a true quadrature pair, the Euclidean norm of the two feature responses should be constant for sine waves of all frequencies and thus insensitive to the local phase of the signal. Due to the restriction of filters to first- and second-order Gaussian derivatives only, this property cannot hold for sine waves of all frequencies at all scales simultaneously. Near the scale levels that are determined by applying scale selection to a sine wave of a given frequency, the phase dependency in the response will, however, be moderately low, as described in [36,37]. Since the Euclidean norm of the first- and second-order Gaussian derivative responses tries to mimic these properties of a quadrature pair, although not being able to obey them fully because of the restriction of the filter basis to the square responses of the first- and second-order Gaussian derivatives only, this entity is termed quasi quadrature.

domain (Lindeberg [36, App. A.3]):

$$\begin{aligned} E_{m,\gamma\text{-norm}} &= \int_{x \in \mathbb{R}} s^{m\gamma} L_{x^m}^2 dx \\ &= \frac{s^{m\gamma}}{2\pi} \int_{\omega \in \mathbb{R}} |\omega|^{2m} \hat{g}^2(\omega; s) d\omega. \end{aligned} \quad (21)$$

For the specific choice of $C = 1/2$ and $\Gamma = 0$, the square of the quasi quadrature measure (20) coincides with the proposals by Loog [125] and Griffin [126] to define a metric of the N -jet in scale space, which can specifically been seen as an approximation of the variance of a signal using a Gaussian spatial weighting function.⁴

Figure 3 shows the result of computing this quasi quadrature measure for a Gaussian peak as well as its first- and second-order derivatives. As can be seen, the quasi quadrature measure is much less sensitive to the position of the peak compared to, e.g., the first- or second-order derivative responses. Additionally, the quasi quadrature measure also has some degree of spatial insensitivity for a first-order Gaussian derivative (a local edge model) and a second-order Gaussian derivative.

4.2 Determination of the Parameter C

To determine the weighting parameter C between local second-order and first-order information, let us consider a Gaussian blob $f(x) = g(x; s_0)$ with spatial extent given by s_0 as input model signal.

By using the semi-group property of the Gaussian kernel $g(\cdot; s_1) * g(\cdot; s_2) = g(\cdot; s_1 + s_2)$, it follows that the scale-space representation is given by $L(x; s) = g(x; s_0 + s)$ and that the first- and second-order derivatives of the scale-space representation are

$$L_x = g_x(x; s_0 + s) = -\frac{x}{(s_0 + s)} g(x; s_0 + s), \quad (22)$$

$$L_{xx} = g_{xx}(x; s_0 + s) = \frac{(x^2 - (s_0 + s))}{(s_0 + s)^2} g(x; s_0 + s), \quad (23)$$

⁴ To understand the relationship between the proposed metric of the N -jet with the variance of the signal, which has been previously described by Griffin [126] and Loog [125], consider a 1D signal that is approximated by its second-order Taylor expansion $L(x) = c_0 + c_1 x + c_2 x^2/2$ around $x = 0$ at some scale level in scale space, where $c_0 = L(0)$, $c_1 = L_x(0)$ and $c_2 = L_{xx}(0)$. The variance of this signal with a Gaussian weighting function $g(x) = \exp(-x^2/2s)/\sqrt{2\pi s}$ around $x = 0$ is $V = \int_{x \in \mathbb{R}} (L(x))^2 g(x) dx - (\int_{x \in \mathbb{R}} L(x) g(x) dx)^2 = M_2 - M_1^2$. Solving these integrals gives $M_2 = c_0^2 + c_0 c_2 s + c_1^2 s + 3c_2^2 s^2/4$ and $M_1 = c_0 + c_2 s/2$, from which we obtain $V = s c_1^2 + s^2 c_2^2/2 = s L_x(0)^2 + s^2 (L_{xx}(0))^2/2 = (Q_{x,\text{norm}}L)^2$ at $x = 0$ for $C = 1/2$ and $\Gamma = 0$. A similar result holds if we instead determine a preferred representative of the class of possible signals that has similar 2-jets (the metamer) by its metamery class norm minimizer $L(x) = (c_0 - c_2 s/2) + c_1 x + c_2 x^2/2$ according to Griffin [126, Sect. 2.1].

Fig. 3 1D Gaussian derivatives up to orders 0, 1 and 2 for $s_0 = 1$ with the corresponding 1D quasi quadrature measures (20) computed from them at scale $s = 1$ for $C = 8/11$ (horizontal axis: $x \in [-5, 5]$)

from which the quasi quadrature measure (20) can be computed in closed form

$$Q_{x, \text{norm}} L = \frac{s^{\frac{1-\Gamma}{2}} e^{-\frac{x^2}{2(s+s_0)}} \sqrt{x^2(s+s_0)^2 + Cs(s+s_0-x^2)^2 + 2}}{\sqrt{2\pi} (s+s_0)^{5/2}}. \quad (24)$$

By determining the weighting parameter C such that it minimizes the overall ripple in the squared quasi quadrature measure for a Gaussian input

$$\hat{C} = \operatorname{argmin}_{C \geq 0} \int_{x=-\infty}^{\infty} \left(\partial_x (Q_{x, \text{norm}}^2 L) \right)^2 dx, \quad (25)$$

which is one way of quantifying the desire to have a stable response under small spatial perturbations of the input, we obtain

$$\hat{C} = \frac{4(s+s_0)}{11s}, \quad (26)$$

which in the special case of choosing $s = s_0$ corresponds to $C = 8/11 \approx 0.727$. This value is very close to the value $C = 1/\sqrt{2} \approx 0.707$ derived from an equal contribution condition in [37, Eq. (27)] for the special case of choosing $\Gamma = 0$.

4.3 Scale Selection Properties

To analyse how the quasi quadrature measure selectively responds to image structures of different sizes, which is important when computing the quasi quadrature entity at multiple scales, we will in this section analyse the scale selection properties of this entity.

Let us consider the result of using Gaussian derivatives of orders 0, 1 and 2 as models of different types of local input signals, i.e.

$$f(x) = g_{x^n}(x; s_0) \quad (27)$$

for $n \in \{0, 1, 2\}$. For the zero-order Gaussian kernel, the scale-normalized quasi quadrature measure at the origin is given by:

$$Q_{x, \text{norm}} L|_{x=0, n=0} = \frac{\sqrt{C} s^{1-\Gamma/2}}{2\pi (s+s_0)^2}. \quad (28)$$

For the first-order Gaussian derivative kernel, the scale-normalized quasi quadrature measure at the origin is

$$Q_{x, \text{norm}} L|_{x=0, n=1} = \frac{s_0^{1/2} s^{(1-\Gamma)/2}}{2\pi (s+s_0)^2}, \quad (29)$$

whereas for the second-order Gaussian derivative kernel, the scale-normalized quasi quadrature measure at the origin is

$$Q_{x, \text{norm}} L|_{x=0, n=2} = \frac{3\sqrt{C} s_0 s^{1-\Gamma/2}}{2\pi (s+s_0)^3}. \quad (30)$$

By differentiating these expressions with respect to the scale parameter s , we find that for a zero-order Gaussian kernel the maximum response over scale is assumed at

$$\hat{s}|_{n=0} = \frac{s_0(2 - \Gamma)}{2 + \Gamma}, \quad (31)$$

whereas for the first- and second-order derivatives, respectively, the maximum response over scale is assumed at

$$\hat{s}|_{n=1} = \frac{s_0(1 - \Gamma)}{3 + \Gamma}, \quad (32)$$

$$\hat{s}|_{n=2} = \frac{s_0(2 - \Gamma)}{4 + \Gamma}. \quad (33)$$

In the special case of choosing $\Gamma = 0$, these scale estimates correspond to

$$\hat{s}|_{n=0} = s_0, \quad (34)$$

$$\hat{s}|_{n=1} = \frac{s_0}{3}, \quad (35)$$

$$\hat{s}|_{n=2} = \frac{s_0}{2}. \quad (36)$$

Thus, for a Gaussian input signal, the selected scale level will for the most scale-invariant choice of using $\Gamma = 0$ reflect the spatial extent $\hat{s} = s_0$ of the blob, whereas if we would like the scale estimate to reflect the scale parameter of first- and second-order derivatives, we would have to choose $\Gamma = -1$. An alternative motivation for using finer scale levels for the Gaussian derivative kernels is to regard the positive and negative lobes of the Gaussian derivative kernels as substructures of a more complex signal, which would then warrant the use of finer scale levels to reflect the substructures of the signal ((35) and (36)).

4.4 Spatial Sensitivity of the Quasi Quadrature Measure

Due to the formulation of the quasi quadrature measure in terms of Gaussian derivatives from the N -jet, the spatial sensitivity (phase dependency) of this entity can be estimated from the first-order component in the local Taylor expansion

$$\frac{\sqrt{s} \partial_x (\mathcal{Q}_{x,\text{norm}} L)}{\mathcal{Q}_{x,\text{norm}} L} = \frac{s L_{xx} (s^{1/2} L_x + C s^{3/2} L_{xxx})}{s L_x^2 + C s^2 L_{xx}^2}, \quad (37)$$

where we have expressed this entity in terms of scale-normalized derivatives for $\gamma = 1$ to emphasize the scale-invariant form of the scale-normalized perturbation measure $s^{1/2} \partial_x (\mathcal{Q}_{x,\text{norm}} L)$. Notably, this entity is zero at inflection points where $L_{xx} = 0$.

4.5 Post-Smoothed Quasi Quadrature Measure

To reduce the spatial sensitivity of the quasi quadrature measure, the definition in Eq. (20) can be complemented by spatial post-smoothing

$$(\overline{\mathcal{Q}}_{x,\text{norm}} L)(\cdot; s, r^2 s) = g(\cdot; r^2 s) * (\mathcal{Q}_{x,\text{norm}} L)(\cdot; s), \quad (38)$$

where the parameter r is referred to as the relative post-smoothing scale. When coupling quasi quadrature measures in cascade, this amount of post-smoothing $r^2 s$ will represent the amount of additional Gaussian smoothing before computing derivatives in the next layer in the hierarchical feature representation.

This spatial post-smoothing operation serves as a scale-covariant spatial pooling operation, notably with the support region, as determined by the integration scale $r^2 s$, proportional to the current scale level s , as opposed to the standard application of spatial pooling over neighbourhoods of fixed size in most CNNs, which would then imply violations of scale covariance.

5 Oriented Quasi Quadrature Modelling of Complex Cells

In this section, we will consider an extension of the 1D quasi quadrature measure (20) into an oriented quasi quadrature measure over 2D image space of the form

$$\mathcal{Q}_{\varphi,\text{norm}} L = \sqrt{\frac{\lambda_{\varphi} L_{\varphi}^2 + C \lambda_{\varphi}^2 L_{\varphi\varphi}^2}{s^{\Gamma}}}, \quad (39)$$

where L_{φ} and $L_{\varphi\varphi}$ denote directional derivatives of an affine Gaussian scale-space representation [51] [56, ch. 15]

$$L(\cdot; s, \Sigma) = g(\cdot; s, \Sigma) * f(\cdot) \quad (40)$$

of the form

$$L_{\varphi} = \cos \varphi L_{x_1} + \sin \varphi L_{x_2}, \quad (41)$$

$$L_{\varphi\varphi} = \cos^2 \varphi L_{x_1 x_1} + 2 \cos \varphi \sin \varphi L_{x_1 x_2} + \sin^2 \varphi L_{x_2 x_2}, \quad (42)$$

and with λ_{φ} denoting the variance of the affine Gaussian kernel (with $x = (x_1, x_2)^T$)

$$g(x; s, \Sigma) = \frac{1}{2\pi s \sqrt{\det \Sigma}} e^{-x^T \Sigma^{-1} x / 2s} \quad (43)$$

in direction φ , preferably with the orientation φ aligned with the direction α of either of the eigenvectors of the composed spatial covariance matrix $s \Sigma$, with

$$\Sigma = \frac{1}{\max(\lambda_1, \lambda_2)} \times \begin{pmatrix} \lambda_1 \cos^2 \alpha + \lambda_2 \sin^2 \alpha & (\lambda_1 - \lambda_2) \cos \alpha \sin \alpha \\ (\lambda_1 - \lambda_2) \cos \alpha \sin \alpha & \lambda_1 \sin^2 \alpha + \lambda_2 \cos^2 \alpha \end{pmatrix} \quad (44)$$

normalized such that the main eigenvalue is equal to one.

5.1 Affine Gaussian Derivative Model for Linear Receptive Fields

According to the normative theory for visual receptive fields in Lindeberg [31–33, 112], directional derivatives of affine Gaussian kernels constitute a canonical model for visual receptive fields over a 2D spatial domain. Specifically, it was proposed that simple cells in the primary visual cortex (V1) can be modelled by directional derivatives of affine Gaussian kernels, termed *affine Gaussian derivatives*, of the form

$$T_{\varphi^m}(x_1, x_2; s, \Sigma) = \partial_{\varphi}^m (g(x_1, x_2; s, \Sigma)). \quad (45)$$

Figure 4 shows an example of the spatial dependency of a colour-opponent simple cell that can be well modelled by a first-order affine Gaussian derivative over an R-G colour-opponent channel over image intensities. Corresponding modelling results for non-chromatic receptive fields can be found in [31–33].

5.2 Affine Quasi Quadrature Modelling of Complex Cells

Figure 5 shows functional properties of a complex cell as determined from its response properties to natural images, using a spike-triggered covariance method (STC), which computes the eigenvalues and the eigenvectors of a second-

order Wiener kernel (Touryan et al. [43]). As can be seen from this figure, the shapes of the eigenvectors determined from the nonlinear Wiener kernel model of the complex cell do qualitatively agree very well with the shapes of corresponding affine Gaussian derivative kernels of orders 1 and 2.

Motivated by this property, that mathematical modelling of functional properties of a biological complex cell in terms of a second-order energy model reveals computational primitives similar to affine Gaussian derivatives, combined with theoretical and experimental motivations for modelling receptive field profiles of simple cells by affine Gaussian derivatives, we propose to model complex cells by a possibly post-smoothed oriented quasi quadrature measure of the form (39)

$$(\bar{Q}_{\varphi, \text{norm}} L)(\cdot; s_{\text{loc}}, s_{\text{int}}, \Sigma_{\varphi}) = \sqrt{g(\cdot; s_{\text{int}}, \Sigma_{\varphi}) * (Q_{\varphi, \text{norm}}^2 L)(\cdot; s_{\text{loc}}, \Sigma_{\varphi})} \quad (46)$$

where $s_{\text{loc}} \Sigma_{\varphi}$ represents an affine covariance matrix in direction φ for computing directional derivatives and $s_{\text{int}} \Sigma_{\varphi}$ represents an affine covariance matrix in the same direction for integrating pointwise affine quasi quadrature measures over a region in image space.

The pointwise affine quasi quadrature measure in this expression $(Q_{\varphi, \text{norm}} L)(\cdot; s_{\text{loc}}, \Sigma_{\varphi})$ can be seen as a Gaussian derivative-based analogue of the energy model for complex cells as proposed by Adelson and Bergen [34] and Heeger [35]. It is closely related to a proposal by Koenderink and van Doorn [128] of summing up the squares of first- and second-order derivative responses and nicely compatible with results by De Valois et al. [129], who showed that first- and second-order receptive fields typically occur in pairs that can be modelled as approximate Hilbert pairs.

Fig. 4 Example of a colour-opponent receptive field profile for a double-opponent simple cell in the primary visual cortex (V1) as measured by Johnson et al. [127]. (left) Responses to L-cones corresponding to long wavelength red cones, with positive weights represented by red and negative weights by blue. (middle) Responses to M-cones corresponding to medium wavelength green cones, with positive weights represented by red and negative weights by blue. (right) Idealized model

of the receptive field from a first-order directional derivative of an affine Gaussian kernel $\partial_{\varphi} g(x, y; \Sigma)$ according to (45) for $\sigma_1 = \sqrt{\lambda_1} = 0.6$, $\sigma_2 = \sqrt{\lambda_2} = 0.2$ in units of degrees of visual angle, $\alpha = 157$ degrees and with positive weights for the red-green colour-opponent channel $U = R - G$ with positive values represented by red and negative values by green

Fig. 5 Significant eigenvectors of a complex cell in the cat primary visual cortex, as determined by Touryan et al. [43] from the response properties of the cell to a set of natural image stimuli, using a spike-triggered covariance method (STC) that computes the eigenvalues and the eigenvectors of a second-order Wiener kernel using three different parameter settings (cutoff frequencies) in the system identification method (from left to right). Qualitatively, these kernel shapes agree well with the shapes of first- and second-order affine Gaussian derivatives

Specifically, this pointwise differential entity mimics some of the known properties of complex cells in the primary visual cortex as discovered by Hubel and Wiesel [48] in the sense of: (i) being independent of the polarity of the stimuli, (ii) not obeying the superposition principle and (iii) being rather insensitive to the phase of the visual stimuli. The primitive components of the quasi quadrature measure (the directional derivatives) do in turn mimic some of the known properties of simple cells in the primary visual cortex in terms of: (i) precisely localized “on” and “off” subregions with (ii) spatial summation within each subregion, (iii) spatial antagonism between on- and off-subregions and (iv) whose visual responses to stationary or moving spots can be predicted from the spatial subregions.

The addition of a complementary post-smoothing stage in (46) as determined by the affine Gaussian weighting function $g(\cdot; s_{\text{int}}, \Sigma_{\varphi})$ is closely related to recent results by Westö and May [130], who have shown that complex cells are better modelled as a combination of two spatial integration steps than a single spatial integration. This spatial post-smoothing stage, which serves as a spatial pooling operation, does additionally decrease the spatial sensitivity of the pointwise quasi quadrature measure and makes it more robust to local spatial perturbations.

By choosing these spatial smoothing and weighting functions as affine Gaussian kernels, we ensure an affine-covariant model of the complex cells, to enable the computation of affine invariants at higher levels in the visual hierarchy.

The use of multiple affine receptive fields over different shapes of the affine covariance matrices $\Sigma_{\varphi, \text{loc}}$ and $\Sigma_{\varphi, \text{int}}$ can be motivated by results by Goris et al. [45], who show that there is a large variability in the orientation selectivity of simple and complex cells (see Fig. 6). With respect to this model, this means that we can think of affine covariance matrices of different eccentricities as being present from

isotropic to highly eccentric. By considering the full family of positive definite affine covariance matrices, we obtain a fully affine-covariant image representation able to handle local linearizations of the perspective mapping for all possible views of any smooth local surface patch.

With respect to computational modelling of biological vision, the proposed affine quasi quadrature model constitutes a novel functional model of complex cells as previously studied in biological vision by Hubel and Wiesel [46–48], Movshon et al. [131], Emerson et al. [132], Touryan et al. [43, 133] and Rust et al. [134] and modelled computationally by Adelson and Bergen [34], Heeger [35], Serre and Riesenhuber [135], Einhäuser et al. [136], Kording et al. [137], Merolla and Boahen [138], Berkes and Wiscott [139],

Fig. 6 Statistics of the orientation selectivity of simple cells and complex cells in the primary visual cortex of the Macaque monkey as reported by Goris et al. [45]. With respect to the affine Gaussian derivative model for the receptive fields of simple and complex cells, the large variability in orientation selectivity reported from these biological measurements implies that we should consider derivatives of affine Gaussian kernels for a large variability in the eccentricity of their shapes, as can be parameterized by, e.g., the ratio between the eigenvalues λ_1 and λ_2 of the affine covariance matrix $s \Sigma$ (a highly eccentric affine Gaussian derivative kernel will have more narrow orientation selectivity)

Carandini [140] and Hansard and Horaud [141]. A conceptual novelty of our model, which emulates several of the known properties of complex cells although our understanding of the nonlinearities of complex cells is still limited, is that it is fully expressed based on the mathematically derived affine Gaussian derivative model for visual receptive fields [32] and therefore possible to relate to natural image transformations as modelled by affine transformations over the spatial domain.

In the following, we will use this quasi quadrature model of complex cells for constructing continuous hierarchical networks.

6 Hierarchies of Oriented Quasi Quadrature Measures

Let us in this first study henceforth for simplicity disregard the variability due to different shapes of the affine receptive fields for different eccentricities and assume that $\Sigma = I$.

This restriction enables covariance to scaling transformations and rotations, whereas a full treatment of affine quasi quadrature measures over all positive definite covariance matrices for the underlying affine Gaussian smoothing operation would enable full affine covariance.

An approach that we shall pursue is to build feature hierarchies by coupling oriented quasi quadrature measures (39) or (46) in cascade⁵

$$F_1(x, \varphi_1) = (\mathcal{Q}_{\varphi_1, \text{norm}} L)(x) \quad (47)$$

$$\begin{aligned} F_k(x, \varphi_1, \dots, \varphi_{k-1}, \varphi_k) \\ = (\mathcal{Q}_{\varphi_k, \text{norm}} F_{k-1})(x, \varphi_1, \dots, \varphi_{k-1}), \end{aligned} \quad (48)$$

where we have suppressed the notation for the scale levels assumed to be distributed such that the scale parameter at level k is $s_k = s_0 r^{2(k-1)}$ for some $r > 1$, e.g. $r = 2$. Assuming that the initial scale-space representation L is computed at scale s_0 , such a network can in turn be initiated for different values of s_0 , also distributed according to a geometric distribution.

This construction builds upon an early proposal by Fukushima [38] of building a hierarchical neural network

⁵ If using raw quasi quadrature measures of the form (39) when constructing the hierarchical representation, the Gaussian spatial smoothing operation, underlying the computation of the Gaussian derivatives from which the quasi quadrature measure is computed, implies that a certain amount of spatial integration (spatial pooling) is guaranteed to be performed in the transformation between successive layers. If the post-smoothed quasi quadrature measure (46) is instead used for constructing the feature hierarchy, then the spatial post-smoothing operation in the post-smoothed quasi quadrature measure guarantees that a certain amount of spatial integration (spatial pooling) is also guaranteed in the quasi quadrature measure computed in any layer.

from repeated application of models of simple and complex cells [46–48], which has later been explored in a handcrafted network based on Gabor functions by Riesenhuber and Poggio [39] and Serre et al. [40] and in the scattering convolution networks by Bruna and Mallat [41]. This idea is also consistent with a proposal by Yamins and DiCarlo [142] of using repeated application of a single hierarchical convolution layer for explaining the computations in the mammalian cortex. With this construction, we obtain a way to define continuous networks that express a corresponding hierarchical architecture based on Gaussian derivative-based models of simple and complex cells within the scale-space framework.

Each new layer in this model implies an expansion of combinations of angles over the different layers in the hierarchy. For example, if we in a discrete implementation discretize the angles $\varphi \in [0, \pi[$ into M discrete spatial orientations, we will then obtain M^k different features at level k in the hierarchy. To keep the complexity down at higher levels, we will for $k \geq K$ in a corresponding way as done by Hadji and Wildes [143] introduce a pooling stage over orientations

$$(\mathcal{P}_k F_k)(x, \varphi_1, \dots, \varphi_{K-1}) = \sum_{\varphi_k} F_k(x, \varphi_1, \dots, \varphi_{K-1}, \varphi_k), \quad (49)$$

which sums up the responses for all the orientations in the current layer, before the next successive layer is instead defined by applying oriented quasi quadrature measures to the pooled responses

$$\begin{aligned} F_k(x, \varphi_1, \dots, \varphi_{k-2}, \varphi_{K-1}, \varphi_k) = \\ (\mathcal{Q}_{\varphi_k, \text{norm}} \mathcal{P}_{k-1} F_{k-1})(x, \varphi_1, \dots, \varphi_{K-1}). \end{aligned} \quad (50)$$

In this way, the number of features at any level will be limited to maximally M^{K-1} . The proposed hierarchical feature representation is termed QuasiQuadNet.

Figure 7 gives a schematic illustration of the structure of such a resulting hierarchy using an expansion over $M = 8$ spatial orientations in the image domain over a total number of four layers with the combinatorial expansion over image orientations delimited from layer $K = 3$.

6.1 Scale Covariance

A theoretically attractive property of this family of networks is that the networks are provably scale covariant. Given two images f and f' that are related by a uniform scaling transformation,

$$f'(x') = f(x) \quad \text{with} \quad x' = Sx \quad (51)$$

for some $S > 0$, their corresponding scale-space representations L and L' will be equal

Fig. 7 Schematic illustration of how the quasi quadrature network is constructed from an image, here with a total number of four layers. In the first layer, there is an expansion over all $M = 8$ orientations, leading to a total number of $2M$ independent features L_φ and $L_{\varphi\varphi}$ over all M image orientations from which the dependent feature QL is then computed according to (39). In the second layer, the maps of QL for all the M image orientations are used for another expansion over image orientations, such that a total number of $2M^2$ independent features L_φ and $L_{\varphi\varphi}$ is computed over all pairs of image orientations. To delimit the complexity of the features in higher layers, there is a pooling stage over image orientations by summing up the quasi quadrature responses over all the image orientations before further expansions over image orientations are performed at layer $K = 3$. Thereby, the number of independent features in these layers is delimited by $2M^2$ instead of $2M^3$. By a corresponding pooling stage before layer 4, the number of independent features in this layer is also delimited by $2M^2$ (the grey boxes, which show the independent features L_φ and $L_{\varphi\varphi}$ and the dependent feature QL that are computed in every layer in the hierarchy, are here only shown for one of the several possible paths through the hierarchy. The combinatorial expansion in layer 2 is also only shown for one of the M orientations in layer 1)

$$L'(x'; s') = L(x; s) \quad (52)$$

and so will the scale-normalized derivatives

$$s'^{n/2} L'_{x_i^n}(x'; s') = s^{n/2} L_{x_i^n}(x; s) \quad (53)$$

based on $\gamma = 1$ if the scale levels are matched according to $s' = S^2 s$ [36, Eqs. (16) and (20)].

This implies that if the initial scale levels s_0 and s'_0 underlying the construction in (47) and (48) are related according to $s'_0 = S^2 s_0$, then the first layers of the feature hierarchy will be related according to [37, Eqs. (55) and (63)]

$$F'_1(x', \varphi_1) = S^{-\Gamma} F_1(x, \varphi_1). \quad (54)$$

Higher layers in the feature hierarchy are in turn related according to

$$F'_k(x', \varphi_1, \dots, \varphi_{k-1}, \varphi_k) = S^{-k\Gamma} F_k(x, \varphi_1, \dots, \varphi_{k-1}, \varphi_k) \quad (55)$$

and are specifically equal if $\Gamma = 0$. This means that it will be possible to perfectly match such hierarchical representations under uniform scaling transformations.

6.2 Rotation Covariance

Under a rotation of image space by an angle α ,

$$f'(x') = f(x) \quad \text{with} \quad x' = R_\alpha x, \quad (56)$$

the corresponding feature hierarchies are in turn equal if the orientation angles are related according to $\varphi'_i = \varphi_i + \alpha$ ($i = 1 \dots k$)

$$F'_k(x', \varphi'_1, \dots, \varphi'_{k-1}, \varphi'_k) = F_k(x, \varphi_1, \dots, \varphi_{k-1}, \varphi_k). \quad (57)$$

6.3 Exact Versus Approximate Covariance (or Invariance) in a Practical Implementation

The architecture of the quasi quadrature network has been designed to support scale covariance based on image primitives (receptive fields) that obey the general scale covariance property (4) and to support rotational covariance by an explicit expansion over image rotations of the form (48).

Scale Covariance The statement about true scale covariance in Sect. 6.1 holds in the continuous case, provided that we can represent a continuum of scale parameters.

In a practical implementation, it is natural to sample this space into a set of discrete scale levels with a constant scale ratio between adjacent scale levels. Then, the scale-covariant property will be restricted to spatial scaling factors that can

be perfectly matched between these scale levels. If the scale levels are expressed in units of $\sigma = \sqrt{s}$ and if the scale ratio between adjacent scale levels in these units is r , then exact scale covariance will hold for all scaling factors that are integer powers of r , provided that the image resolution and the image size are sufficient to resolve the relevant image structures. For scaling factors in between these discrete values, there will be an approximation error, which could possibly be reduced by a complementary scale interpolation mechanism.

For a discrete implementation with limited image resolution and limited image size, there will be additional restrictions on how well the discrete implementation approximates the continuous theory. For the implementations underlying this paper, we use a scale-space concept specially designed for discrete signals computed by separable convolution with the discrete analogue of the Gaussian kernel $T(n; s) = e^{-s} I_n(s)$ [145], which is defined in terms of the modified Bessel functions of integer order $I_n(s)$ [146]. This discrete scale-space concept constitutes numerical approximation of the continuous scale-space concept via a spatial discretization of the diffusion equation, which governs the evolution properties over scale of the Gaussian scale-space concept.

Rotational Covariance The statement about true rotational covariance in Sect. 6.2 holds provided that we can represent a continuum of rotation angles. For a continuum of orientation angles, the summation over image orientations in the pooling stage (49) should be replaced by an integral over all the image orientations to guarantee exact covariance to hold for all rotation angles.

In a practical implementation, it is natural to sample the orientation angles on the unit circle into a set of discrete angles with a constant increment between. Then, the rotation-covariant property will be restricted to the set of discrete rotation angles that are spanned by this discretization. For rotation angles in between, there will be an approximation error, which could possibly be reduced by a suitable interpolation mechanism.

With regard to a discrete implementation, there may be additional deviations in how well the discrete approximations of directional derivatives numerically approximate their continuous counterparts. For the implementation underlying this paper, we complement the discrete scale-space concept in [145] with discrete derivative approximations with scale-space properties [147], where small support discrete derivative approximations $\delta_x = (-1/2, 0, 1/2)$ and $\delta_{xx} = (1, -2, 1)$ are applied to the discrete scale-space smoothed image data and directional derivative approximations are then computed from the continuous relationships (41) and (42).

Numerical Approximation of a Truly Covariant Continuous Theory By all steps in the discrete implementation constituting numerical approximations of their corresponding counterparts in the continuous theory, it follows that the dis-

crete implementation will also numerically approximate the desirable covariance properties (or as an extension invariance properties) with respect to scaling transformations and rotations in the image domain. The accuracy of approximation of the combined system will then be a composed effect of the numerical accuracy of the different primitives.

6.4 Experiments

Figures 8, 9 and 10 show examples of computing different layers in such a quasi quadrature network for two texture images and an indoor image, with the combinatorial angular expansion for higher layers delimited at layer $K = 3$.

For the quite regular corduroy image in Fig. 8, we can see that we get clear responses to the stripes in the cloth in layers 1 and 2, with only a minor dominant response in the third layer corresponding to the slight irregularity in the mid-left of the original image.

For the mixed regular/irregular wool image in Fig. 9, we get clear responses to the crochet work in layer 1, with additional clear responses to the different types of repeated crochet structures in different subparts of the image in layer 2, whereas in layer 3 the main strong response is due to the intentional overall irregularity in the pattern.

For the indoor scene in Fig. 10, we can note that the responses are strongest along the edges in the scene for all the layers, with some locally stronger responses in layers 2 and 3 assumed near corners or end-stoppings, especially when the orientations of the oriented quasi quadrature measures at higher levels in the hierarchy are orthogonal to the orientation of the oriented quasi quadrature measure in the first layer ($\varphi_2 \perp \varphi_1$ or $\varphi_3 \perp \varphi_1$). For this image, which is not in any way stationary over image space, we can observe that the spatial structure of the scene can be perceived from the pure magnitude responses of the quasi quadrature measure in layer 3 in the hierarchy.

In these qualitative respects, we can see how the proposed quasi quadrature hierarchy is able to reflect nonlinear hierarchical relations between image structures over different scales.

7 Application to Texture Analysis

In the following, we will use a substantially reduced version of the proposed quasi quadrature network for building an application to texture analysis and evaluate the resulting approach on the KTH-TIPS2b, CuRET and UMD datasets (Figs. 11, 12, 13).

7.1 Mean-Reduced Texture Descriptors

If we make the assumption that a spatial texture should obey certain stationarity properties over image space, we may

Fig. 8 Subset of features in a hierarchy of directional derivatives and quasi quadrature measures computed from a texture image (corduroy/sample_a/42a-scale_4_im_5_col.png) from the KTH-TIPS2 dataset [144] for different combinations of angles in layers 1 and 2 for $s_0 = 2$. (top row) Original image and first-order directional derivatives L_{φ} . (second row) Second-order directional derivatives $L_{\varphi\varphi}$. (third row) Oriented quasi quadrature measures $\mathcal{Q}_{\varphi_1} L$ in layer 1. (fourth row)

First-order directional derivatives computed from $\mathcal{Q}_{\varphi_1} L$ for $\varphi_1 = 0$. (fifth row) Second-order directional derivatives computed from $\mathcal{Q}_{\varphi_1} L$ for $\varphi_1 = 0$. (sixth row) Oriented quasi quadrature measures $\mathcal{Q}_{\varphi_2} \mathcal{Q}_{\varphi_1} L$ in layer 2 for $\varphi_1 = 0$. (bottom row) Oriented quasi quadrature measures $\mathcal{Q}_{\varphi_3} \mathcal{P}_{\varphi_2} \mathcal{Q}_{\varphi_2} \mathcal{Q}_{\varphi_1} L$ in layer 3 for $\varphi_1 = 0$ (the contrast has been reversed for the quasi quadrature measures so that high values are shown as dark and low values as bright; image size: 200×200 pixels)

Fig. 9 Subset of features in a hierarchy of directional derivatives and quasi quadrature measures computed from a texture image (wool/sample_a/22a-scale_7_im_3_col.png) from the KTH-TIPS2 dataset [144] for different combinations of angles in layers 1 and 2 for $s_0 = 2$. (top row) Original image and first-order directional derivatives L_{φ} . (second row) Second-order directional derivatives $L_{\varphi \varphi}$. (third row) Oriented quasi quadrature measures $Q_{\varphi_1} L$ in layer 1. (fourth row)

First-order directional derivatives computed from $Q_{\varphi_1} L$ for $\varphi_1 = 0$. (fifth row) Second-order directional derivatives computed from $Q_{\varphi_1} L$ for $\varphi_1 = 0$. (sixth row) Oriented quasi quadrature measures $Q_{\varphi_2} Q_{\varphi_1} L$ in layer 2 for $\varphi_1 = 0$. (bottom row) Oriented quasi quadrature measures $Q_{\varphi_3} P_{\varphi_2} Q_{\varphi_2} Q_{\varphi_1} L$ in layer 3 for $\varphi_1 = 0$ (the contrast has been reversed for the quasi quadrature measures so that high values are shown as dark and low values as bright; image size: 200×200 pixels)

Fig. 10 Subset of features in a hierarchy of directional derivatives and quasi quadrature measures computed from an indoor image for different combinations of angles in layers 1 and 2 for $s_0 = 2$. (top row) Original image and first-order directional derivatives L_{φ} . (second row) Second-order directional derivatives $L_{\varphi\varphi}$. (third row) Oriented quasi quadrature measures $Q_{\varphi_1} L$ in layer 1. (fourth row) First-order directional derivatives computed from $Q_{\varphi_1} L$ for $\varphi_1 = 0$. (fifth row)

Second-order directional derivatives computed from $Q_{\varphi_1} L$ for $\varphi_1 = 0$. (sixth row) Oriented quasi quadrature measures $Q_{\varphi_2} Q_{\varphi_1} L$ in layer 2 for $\varphi_1 = 0$. (seventh row) Oriented quasi quadrature measures $Q_{\varphi_2} Q_{\varphi_1} L$ in layer 2 for $\varphi_1 = \pi/2$. (bottom row) Oriented quasi quadrature measures $Q_{\varphi_3} P_{\varphi_2} Q_{\varphi_2} Q_{\varphi_1} L$ in layer 3 for $\varphi_1 = 0$ (the contrast has been reversed for the quasi quadrature measures so that high values are shown as dark and low values as bright; image size: 512×350 pixels)

regard it as reasonable to construct texture descriptors by accumulating statistics of feature responses over the image domain, in terms of, e.g., mean values or histograms.

Inspired by the way the SURF descriptor [68] accumulates mean values and mean absolute values of derivative responses and the way Bruna and Mallat [41] and Hadji and Wildes [143] compute mean values of their hierarchical fea-

ture representations, we will initially explore reducing the QuasiQuadNet to just the mean values over the image domain of the following five features:

$$\{\partial_{\varphi} F_k, |\partial_{\varphi} F_k|, \partial_{\varphi\varphi} F_k, |\partial_{\varphi\varphi} F_k|, Q_{\varphi} F_k\}. \quad (58)$$

These types of features are computed for all layers in the feature hierarchy (with $F_0 = L$), which leads to a 4000D

Table 1 Performance results of the mean-reduced QuasiQuadNet in comparison with a selection of among the better methods in the extensive performance evaluation by Liu et al. [157] (our results in slanted font) (the column labelled “feat” states whether the image features are fixed (“F”) or learnt (“L”). The column labelled “class” states whether the classification criterion is fixed (“F”) or learnt (“L”))

	KTH-TIPS2b	CUReT	UMD	Feat	Class
FV-VGGVD [151] (SVM)	88.2	99.0	99.9	L	L
FV-VGGM [151] (SVM)	79.9	98.7	99.9	L	L
MRELBP [152] (SVM)	77.9	99.0	99.4	F	L
FV-AlexNet [151] (SVM)	77.9	98.4	99.7	L	L
<i>Mean-reduced QuasiQuadNet LUV (SVM)</i>	78.3	98.6		F	L
<i>Mean-reduced QuasiQuadNet grey (SVM)</i>	75.3	98.3	97.1	F	L
ScatNet [41] (PCA)	68.9	99.7	98.4	F	L
MRELBP [152]	69.0	97.1	98.7	F	F
BRINT [153]	66.7	97.0	97.4	F	F
MDLBP [154]	66.5	96.9	97.3	F	F
<i>Mean-reduced QuasiQuadNet LUV (NNC)</i>	72.1	94.9		F	F
<i>Mean-reduced QuasiQuadNet grey (NNC)</i>	70.2	93.0	93.3	F	F
LBP [155]	62.7	97.0	96.2	F	F
ScatNet [41] (NNC)	63.7	95.5	93.4	F	F
PCANet [117] (NNC)	59.4	92.0	90.5	L	F
RandNet [117] (NNC)	56.9	90.9	90.9	F	F

Fig. 11 Sample images from the KTH-TIPS2b texture dataset [144]. This dataset consists of images of 11 classes of textures with four samples from each class. Each sample has been photographed from nine distances leading to nine relative scales, with additionally 12 different

pose and illumination conditions for each scale, implying a total number of $11 \times 4 \times 9 \times 12 = 4752$ images. This figure shows one sample from each class, with varying scale, pose and illumination conditions between the samples (most images of size 200×200 pixels)

descriptor⁶ based on $M = 8$ uniformly distributed orientations in $[0, \pi]$, four layers in the hierarchy delimited in complexity by directional pooling for $K = 3$ with four initial scale levels $\sigma_0 = \sqrt{s_0} \in \{1, 2, 4, 8\}$.

⁶ With $M = 8$ orientations in image space and five basic types of features $\{\partial_\varphi F_k, |\partial_\varphi F_k|, \partial_{\varphi\varphi} F_k, |\partial_{\varphi\varphi} F_k|, Q_\varphi F_k\}$, there are $8 \times 5 = 40$ features in layer 1 at a single scale and $8 \times 8 \times 5 = 320$ features in layer 2 due to the additional combinatorial expansion in and similar numbers of 320 features in layers 3 and 4 due to the limitation on combinatorial complexity at layer $K = 3$. For any initial scale level σ_0 , there are therefore a total number of $40 + 3 \times 320 = 1000$ features. Expanded over four initial scale levels $\sigma_0 = \sqrt{s_0} \in \{1, 2, 4, 8\}$, this leads to a total number of 4000 feature dimensions, which we here represent by just their average values over image space.

7.2 Texture Classification on the KTH-TIPS2b Dataset

The second column in Table 1 shows the result of applying this approach to the KTH-TIPS2b dataset [144] for texture classification, see Fig. 11 for sample images from this dataset. The KTH-TIPS2b dataset contains images of 11 texture classes (“aluminium foil”, “cork”, “wool”, “lettuce leaf”, “corduroy”, “linen”, “cotton”, “brown bread”, “white bread”, “wood” and “cracker”) with four physical samples from each class and photographs of each sample taken from nine distances leading to nine relative scales labelled “2”, ..., “10” over a factor of 4 in scaling transformations and additionally 12 different pose and illumination conditions for

Fig. 12 Sample images from the CURET texture dataset [148]. This dataset consists of images of 61 materials, with a single sample for each class, and each sample viewed under different viewing and illumination conditions. Here, we use the selection of 92 viewing and illumination

conditions chosen in [149] leading to a total number of $61 \times 92 = 5612$ images. This figure shows one image of about every third sample, with varying viewing and illumination conditions between the samples (all images of size 200×200 pixels)

Fig. 13 Sample images from the UMD texture dataset [150]. This dataset consists of 25 texture classes, with 40 grey-level images from each class taken from a variety of different distances and viewing direc-

tions; thus, a total number $25 \times 40 = 1000$ images. This figure shows one sample from each of the first 15 classes (all images of size 1280×960 pixels)

each scale, leading to a total number of $11 \times 4 \times 9 \times 12 = 4752$ images. The regular benchmark setup implies that the images from three samples in each class are used for training and the remaining sample in each class for testing over four permutations. Since several of the samples from the same class are quite different from each other in appearance, this implies a non-trivial benchmark which has not yet been saturated.

When using nearest-neighbour classification on the mean-reduced grey-level descriptor, we get 70.2 % accuracy, and

72.1 % accuracy when computing corresponding features from the LUV channels of a colour-opponent representation. When using SVM classification [156], the accuracy becomes 75.3 % and 78.3 %, respectively. Comparing with the results of an extensive set of other methods in Liu et al. [157], out of which a selection of the better results is listed in Table 1, the results of the mean-reduced QuasiQuadNet are better than classical texture classification methods such as local binary patterns (LBP) [155], binary rotation-invariant noise-

Fig. 14 Examples of the scaling variations in the KTH-TIPS2 dataset for one sample each from the classes “wool”, “lettuce” and “brown bread” at a subset of five of the scales in the dataset (the sizes labelled “2”, “4”, “6”, “8” and “10” from top to bottom)

tolerant texture descriptors [153] and multidimensional local binary patterns (MDLBP) [154] and also better than other handcrafted networks, such as ScatNet [41], PCANet [117] and RandNet [117]. The performance of the mean-reduced QuasiQuadNet descriptor does, however, not reach the performance of applying SVM classification to Fischer vectors of the filter output in learned convolutional networks (FV-VGGVD, FV-VGGM [151]).

By instead performing the training on every second scale in the dataset (scales “2”, “4”, “6”, “8”, “10”) and the testing on the other scales (“3”, “5”, “7”, “9”), such that the benchmark does not primarily test the generalization properties between the different very few samples in each class, the classification performance is 98.8 % for the grey-level descriptor and 99.6 % for the LUV descriptor.

Fig. 15 Comparison between scale-covariant matching versus non-covariant matching of texture descriptors on the KTH-TIPS2b dataset [144]. In the experiments underlying this figure, we have used the scale variations in the dataset to perform matching over spatial scaling factors of $S = \sqrt{2}$, 2, $2\sqrt{2}$ and 4, here represented as $\log_2 S$ on the horizontal axis. For non-covariant matching, here represented as red curves, we have used the same scale parameters for the image descriptors in the training data and the test data. For scale-covariant matching, here represented as blue curves, we have adapted the scale levels of the image descriptors to the known scale factor between the training data and the test data. As can be seen from the results, in the presence of substantial scaling variations, the use of scale-covariant matching, as enabled by the provably scale-covariant networks proposed in this article, improves the performance substantially if there are significant scaling variations in the data (all these results have been computed with SVM classification of mean-reduced image descriptors from QuasiQuadNets computed from either pure grey-level images or colour images. The results for the pure grey-level descriptors are indicated by ‘o’, whereas the results for the LUV colour descriptors are indicated by ‘*’)

7.3 Scale-Covariant Matching of Image Descriptors on the KTH-TIPS2b Dataset

An attractive property of the KTH-TIPS2b dataset is that we can use the controlled scaling variations in this dataset (see Fig. 14) to investigate the influence of scale covariance with respect to image descriptors defined from a provably scale-covariant network. To test this property, we constructed partitionings of the dataset into training sets and test sets with known scaling variations between the data.

The scales in the datasets, which we will henceforth refer to as sizes, labelled from “2” to “10”, span a scaling factor of 4, with a relative scaling factor of $\sqrt[4]{2}$ between adjacent sizes. To cover a set of relative scaling factors $S \in \{\sqrt{2}, 2, 2\sqrt{2}, 4\}$, we partitioned the dataset and adapted the scale parameters of the QuasiQuadNet to the relative scaling factors in the following way:

- Relative scaling factor $\sqrt{2}$: Training data at the sizes labelled $\{5, 6, 9, 10\}$ with image descriptors computed at the scales $\sigma_0 \in \{1, 2, 4, 8\}$. Test data at the sizes labelled $\{3, 4, 7, 8\}$ with image descriptors computed at the scales $\sigma_0 \in \{\sqrt{2}, 2\sqrt{2}, 4\sqrt{2}, 8\sqrt{2}\}$.
- Relative scaling factor 2: Training data at the sizes labelled $\{7, 8, 9, 10\}$ with image descriptors computed at the scales $\sigma_0 \in \{1, 2, 4, 8\}$. Test data at the sizes labelled $\{3, 4, 5, 6\}$ with image descriptors computed at the scales $\sigma_0 \in \{2, 4, 8, 16\}$.
- Relative scaling factor $2\sqrt{2}$: Training data at the sizes labelled $\{8, 9, 10\}$ with image descriptors computed at the scales $\sigma_0 \in \{1, 2, 4, 8\}$. Test data at the sizes labelled $\{2, 3, 4\}$ with image descriptors computed at the scales $\sigma_0 \in \{2\sqrt{2}, 4\sqrt{2}, 8\sqrt{2}, 16\sqrt{2}\}$.
- Relative scaling factor 4: Training data at the size labelled $\{10\}$ with image descriptors computed at the scales $\sigma_0 \in \{1, 2, 4, 8\}$. Test data at the size labelled $\{2\}$ with image descriptors computed at scales $\sigma_0 \in \{4, 8, 16, 32\}$.

These partitionings between training sets and test sets have thus been constructed in such a way that there for each image descriptor computed from an image in the test set should exist a corresponding scale-matched image descriptor in the training set.

To measure the influence relative to not adapting the scale levels to scale covariance, we also performed non-covariant classification with all the image descriptors, both in the training data and the test data, computed at the scales $\sigma_0 \in \{1, 2, 4, 8\}$.

The result of this experiment is shown in Fig. 15, which shows graphs of how the accuracy of the texture classification depends on the logarithm of the relative scaling factor $\log_2 S$ between the training data and the test data (see also Table 2). As can be seen from the graphs, the performance is substantially higher for scale-covariant classification compared to non-covariant classification. Although this task is not influenced by the generalization ability of the image descriptors, as measured in the regular experimental setup for the KTH-TIPS2 dataset in the sense that images from all the samples are here included in both the training sets and the test sets, there are nevertheless reasons why the image data cannot be perfectly matched: (i) the support regions for the texture descriptors differ in size due to the scaling transformation, which implies that new image details appear in one of the images relative to the other (see Fig. 14 for an illustration), which in turn challenges the stationarity assumption underlying the image texture descriptor, here represented by mean values only and (ii) the boundary effects at the image boundaries are different between the two image domains, which in particular affect the image features at coarser spatial scales.

Notwithstanding these effects, due to the fact that the addition of new image structures during the scaling transfor-

Table 2 Numerical performance values underlying the graphs in Fig. 15, which quantify the performance of texture classification based on mean-reduced texture descriptors from QuasiQuadNets over scaling transformations with scaling factors of $\sqrt{2}$, 2, $2\sqrt{2}$ and 4 for the KTH-TIPS2b dataset

S	$\sqrt{2}$	2	$2\sqrt{2}$	4
Non-covariant grey	90.6	72.6	46.1	31.4
Scale-covariant grey	97.5	93.8	86.9	79.0
Non-covariant LUV	93.0	79.1	55.2	32.6
Scale-covariant LUV	96.5	93.9	87.8	78.6

mations leads to a violation of full scale covariance because of the a priori delimited image domains in the already given dataset, the primary purpose of this experiment is to conceptually demonstrate how substantial gains in performance can be obtained by having a scale-covariant network, and how such scale-covariant networks are conceptually easier to construct using a continuous model of the filtering operations in the network. Specifically, scale-space theory, which underpins this treatment, has been developed to handle such scaling variations in a theoretically well-founded manner.

7.4 Matching with Scale-Aggregated Covariant Image Descriptors on the KTH-TIPS2b Dataset

In the previous section, we used a priori known information about the structured amounts of scaling transformations in the KTH-TIPS2 dataset for demonstrating the importance of using scale-covariant image descriptors as opposed to non-covariant image descriptors in situations where the scaling transformations are substantial.

A more realistic scenario is that the amount of scaling transformation between the training data and the test data is not a priori known. A useful approach in such a situation is to complement the image descriptors in the training set by scale aggregation, meaning that multiple copies of image descriptors are computed over some set of scale levels, to enable scale-covariant matching of the image descriptors in the sense that for any image descriptor computed from the test set we should as far as possible increase the likelihood for the classification scheme to be able to find a corresponding scale-matched image descriptor in the training set.

To test the scale sensitivity of a composed texture classification scheme to such a scenario, we computed image descriptors for the training data at the following scales $\sigma_0 \in \{1, 2, 4, 8\}$, $\sigma_0 \in \{\sqrt{2}, 2\sqrt{2}, 4\sqrt{2}, 8\sqrt{2}\}$, $\sigma_0 \in \{2, 4, 8, 16\}$, $\sigma_0 \in \{2\sqrt{2}, 4\sqrt{2}, 8\sqrt{2}, 16\sqrt{2}\}$, $\sigma_0 \in \{4, 8, 16, 32\}$ and computed the test data at the single scale $\sigma_0 \in \{1, 2, 4, 8\}$. As training data we used the images at the single size $\{2\}$ and as test data the images from a single one of each of the sizes $\{3, 4, 5, 6, 7, 8, 9, 10\}$, to study the sensitivity to variations

Table 3 Numerical performance values underlying the bottom graphs in Fig. 16, which quantify the performance of texture classification based on mean-reduced texture descriptors from QuasiQuadNets over scaling transformations with different scaling factors S

S	$2^{1/4}$	$2^{1/2}$	$2^{3/4}$	2^1	$2^{5/4}$	$2^{3/2}$	$2^{7/4}$	2^2
Non-aggregated grey	95.5	86.0	77.3	66.7	51.4	39.2	32.8	27.1
Scale-aggregated grey	98.0	97.0	95.3	92.2	89.4	83.5	76.4	71.8
Non-aggregated LUV	97.3	88.6	82.9	73.0	57.6	48.0	39.9	32.2
Scale-aggregated LUV	98.8	97.2	95.6	92.8	89.1	83.9	80.0	76.5

in scaling transformation in steps of $\sqrt[4]{2}$ between adjacent sizes.

The result of this experiment is shown in top figure in Fig. 16, which shows graphs of how the accuracy of the texture classification depends on the logarithm of the relative scaling factor $\log_2 S$ between the training data and the test data. In the top figure, the experiments have been made relative to training data the single size “2” only, and corresponding test data for each one of the sizes “3”, “4”, ..., “10” in the dataset. In the bottom figure, the average result of a set of more extensive experiments is shown, where each one of the sizes “2”, “3”, ..., “9” has been used for defining scale-aggregated training data and the testing data has then been taken from a single size with number label greater than the label for the training data. The graphs in the bottom figure show the average values over all those graphs for equal relative scaling factors between the training data and the test data (see also Table 3). As can be seen from the graphs, the performance is substantially higher for scale-aggregated matching compared to non-aggregated matching. In this way, the experiment demonstrates how the use of a scale-covariant network enables significantly better performance in situations where there are substantial scaling transformations in the test data that are not spanned by corresponding scaling variations in the training data.

A similar way of handling scale variations between training data and test data by computing the image descriptors over a range of scales has also been used for texture classification by Crosier and Griffin [158].⁷ This type of scale matching constitutes an integrated part of the scale-space methodology for relating image descriptors computed from image structures that have been subject to scaling transformations in the image domain. Here, we extend this approach for scale generalization to hierarchical or deep networks, where the scale covariance property of our networks makes such scale matching possible.

⁷ The approach by Crosier and Griffin [158] is also scale covariant, however, not hierarchical or deep. They do furthermore not test for scale prediction or scale generalization over large scaling factors as we target here.

Fig. 16 Comparison between scale-aggregated matching versus non-aggregated matching of texture descriptors on the KTH-TIPS2b dataset [144]. In the experiments underlying this figure, we have used the scale variations in the dataset to perform matching over spatial scaling factors of S between $\sqrt[4]{2}$ and 4 in steps of factors of $\sqrt[4]{2}$, here represented as $\log_2 S$ on the horizontal axis. For non-aggregated matching, here represented as red curves, we have used the image descriptors at the same single scale $\sigma_0 \in \{1, 2, 4, 8\}$ in the training data and the test data. For scale-aggregated matching, here represented as blue curves, we have extended the training data with image descriptors over the set of scale levels $\sigma_0 \in \{1, 2, 4, 8\}$, $\sigma_0 \in \{\sqrt{2}, 2\sqrt{2}, 4\sqrt{2}, 8\sqrt{2}\}$, $\sigma_0 \in \{2, 4, 8, 16\}$, $\sigma_0 \in \{2\sqrt{2}, 4\sqrt{2}, 8\sqrt{2}, 16\sqrt{2}\}$, $\sigma_0 \in \{4, 8, 16, 32\}$ to span scaling variations up to a factor of 4 in steps of $\sqrt{2}$. In the top figure, the experiments have been made relative to training data at size “2” only, and corresponding test data at the sizes “3”, “4”, ..., “10”. In the bottom figure, a set of multiple experiments have been performed with training data at each one of the sizes “2”, “3”, ..., “9”, with testing data for the sizes with greater number labels. The curve in the bottom figure shows the average value of all these experiments as averaged over equal relative scaling factors. As can be seen from the results, in the presence of substantial scaling variations, the use of scale-aggregated matching, as enabled by the provably scale-covariant networks proposed in this article, improves the performance substantially if there are significant scaling variations in the data (all these results have been computed with SVM classification of mean-reduced image descriptors from QuasiQuadNets computed from either pure grey-level images or colour images. The results for the pure grey-level descriptors are indicated by ‘o’, whereas the results for the LUV colour descriptors are indicated by ‘*’)

7.5 Texture Classification on the CURET Dataset

The third column in Table 1 shows the result of applying a similar texture classification approach as is used in Sect. 7.2 to the CURET texture dataset [148], see Fig. 12 for sample images from this dataset. The CURET dataset consists of images of 61 materials, with a single sample for each material, and each sample viewed under 205 different viewing and illumination conditions. For our experiments, we use the selection of 92 cropped images of size 200×200 pixels chosen in [149] from the criterion that a sufficiently large region of texture should be visible for all the materials. This implies a total number of $61 \times 92 = 5612$ images. Following the standard for this dataset, we measure the average value of a set of random partitionings into training and testing data of equal size.

With SVM classification on the mean-reduced Quasi-QuadNet, we get 98.3 % accuracy for the grey-level descriptor and 98.6 % for the colour descriptor. This performance is better than handcrafted PCANet [117] and RandNet [117] and some pure texture descriptors such as local binary patterns [155], multidimensional local binary patterns (MDLBP) [154], binary rotation-invariant noise tolerant texture descriptors [153] and near the learned networks FV-AlexNet and FV-VGGM [151]. For this dataset, the handcrafted ScatNet [41] does, however, perform better and so do the learned networks FV-VGGVD [151] and median robust extended local binary patterns [152].

7.6 Texture Classification on the UMD Dataset

The fourth column in Table 1 shows the result of applying a similar texture classification approach to the UMD texture dataset [150], see Fig. 13 for sample images from this dataset. The UMD dataset consists of 25 texture classes with 40 grey-level images of size 1280×900 pixels from each class, taken from different distances and viewpoints; thus, a total number of $25 \times 40 = 1000$ images. Following the standard for this dataset, we measure the average of random partitions in training and testing data of equal size. When using the same scale levels $\sigma_0 \in \{1, 2, 4, 8\}$ for the training data and the test data, we get 97.1 % accuracy of our mean-reduced grey-level descriptor, which is better than local binary patterns [155], PCANet [117] and RandNet [117].

Noting that this dataset contains significant unstructured scaling variations, which are not taken into account when computing all the image descriptors at the same scale, we also did an experiment with scale-covariant matching, where we expanded the training data to the following scale combinations $\sigma_0 \in \{1, 2, 4, 8\}$, $\sigma_0 \in \{\sqrt{2}, 2\sqrt{2}, 4\sqrt{2}, 8\sqrt{2}\}$, $\sigma_0 \in \{2, 4, 8, 16\}$, $\sigma_0 \in \{2\sqrt{2}, 4\sqrt{2}, 8\sqrt{2}, 16\sqrt{2}\}$, $\sigma_0 \in \{4, 8, 16, 32\}$ and computed the test data at the single scale $\sigma_0 \in \{2, 4, 8, 16\}$. The intention behind this data aggregation

over scales is to make it easier to find a match between the training data and the test data over situations where there are significant scaling transformations between the training data and the test data, with specifically a lack of matching training data at a similar scale as for a given test data. Then, the performance increased from 93.3 to 95.9% using NN classification and from 97.1 to 98.1% using SVM classification on the UMD dataset.

A corresponding expansion of the training data to cyclic permutations over the underlying angles in the image descriptors in the training data, to achieve rotation-covariant matching, did, however, not improve the results.

8 Summary and Discussion

We have presented a theory for defining handcrafted or structured hierarchical networks by combining linear and nonlinear scale-space operations in cascade. After presenting a general sufficiency condition to construct networks based on continuous scale-space operations that guarantee provable scale covariance, we have then in more detail developed one specific example of such a network constructed by applying quasi quadrature responses of first- and second-order directional Gaussian derivatives in cascade.

A main purpose behind this study has been to investigate whether we could start building a bridge between the well-founded theory of scale-space representation and the recent empirical developments in deep learning, while at the same time being inspired by biological vision. The present work is intended as initial work in this direction, where we propose the family of quasi quadrature networks as a new baseline for handcrafted networks with associated provable covariance properties under scaling and rotation transformations.

Specifically, by constructing the network from linear and nonlinear filters defined over a continuous domain, we avoid the restriction to discrete 3×3 or 5×5 filters in most current deep net approaches, which implies an implicit assumption about a preferred scale in the data, as defined by the grid spacing in the deep net. If the input data to the deep net are rescaled by external factors, such as from varying the distance between an observed object and the observer, the lack of true scale covariance as arising from such preferred scales in the network implies that the nonlinearities in the deep net may affect the data in different ways, depending on the size of a projected object in the image domain.

By early experiments with a substantially mean-reduced representation of our provably scale-covariant QuasiQuadNet, we have demonstrated that it is possible to get quite promising performance on texture classification, and comparable or better than other handcrafted networks, although not reaching the performance of applying more refined statistical classification methods on learned CNNs.

By inspection of the full non-reduced feature maps, we have also observed that some representations in higher layers may respond to irregularities in regular textures (defect detection) or corners or end-stoppings in regular scenes.

Concerning extensions of the approach with quasi quadrature networks, we propose to:

- relax the restriction to isotropic covariance matrices with $\Sigma = I$ in Sect. 6 to construct hierarchical networks based on more general affine quasi quadrature measures based on affine Gaussian derivatives that are computed with varying eccentricities of the underlying affine Gaussian kernel to enable affine covariance, which will then also enable affine invariance,
- complement the computation of quasi quadrature responses by a mechanism for divisive normalization [44] to enforce a competition between multiple feature responses and thus increase the selectivity of the image features,
- explore the spatial relationships in the full feature maps that are suppressed in the mean-reduced representation to make it possible for the resulting image descriptors to encode hierarchical relations between image features over multiple positions in the image domain and
- incorporate learning mechanisms into the representation.

Specifically, it would be interesting to formulate learning mechanisms that can learn the parameters of a parameterized model for divisive normalization and to formulate learning mechanisms that can combine quasi quadrature responses over different positions in the image domain to support more general object recognition mechanisms than those that can be supported by a stationarity assumption as explored in the prototype application to texture classification developed in Sect. 7.

For the specific application to texture classification in this work, it does also seem possible that using more advanced statistical classification methods on the QuasiQuadNet, such as Fischer vectors, could lead to gains in performance compared to the mean-reduced representation that we used here, based on just the mean values and the mean absolute values of the filter responses in our hierarchical representation.

Concerning more general developments, the general arguments about scale-covariant continuous networks in Sect. 3 open up for studying wider classes of continuous hierarchical networks that guarantee provable scale covariance. We plan to study such extensions in future work.

Acknowledgements Open access funding provided by Royal Institute of Technology. I would like to thank Ylva Jansson and the anonymous reviewers for valuable comments that improved the presentation of some of the topics in the article.

Open Access This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

References

1. Krizhevsky, A., Sutskever, I., Hinton, G.E.: ImageNet classification with deep convolutional neural networks. In: *Advances in Neural Information Processing Systems*, pp. 1097–1105 (2012)
2. Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition. In: *International Conference on Learning Representations (ICLR 2015)* (2015). [arXiv:1409.1556](https://arxiv.org/abs/1409.1556)
3. LeCun, Y., Bengio, Y., Hinton, G.: Deep learning. *Nature* **521**, 436–444 (2015)
4. Szegedy, C., Liu, W., Jia, Y., Sermanet, P., Reed, S., Anguelov, D., Erhan, D., Vanhoucke, V., Rabinovich, A.: Going deeper with convolutions. In: *Proceedings of Computer Vision and Pattern Recognition (CVPR 2015)*, pp. 1–9 (2015)
5. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: *Proceedings of Computer Vision and Pattern Recognition (CVPR 2016)*, pp. 770–778 (2016)
6. Redmon, J., Divvala, S., Girshick, R., Farhadi, A.: You only look once: unified, real-time object detection. In: *Proceedings of Computer Vision and Pattern Recognition (CVPR 2016)*, pp. 779–788 (2016)
7. Huang, G., Liu, Z., van der Maaten, L., Weinberger, K.Q.: Densely connected convolutional networks. In: *Proceedings of Computer Vision and Pattern Recognition (CVPR 2017)*, pp. 4700–4708 (2017)
8. Xie, S., Girshick, R., Dollár, P., Tu, Z., He, K.: Aggregated residual transformations for deep neural networks. In: *Proceedings of Computer Vision and Pattern Recognition (CVPR 2017)*, pp. 5987–5995 (2017)
9. Ren, S., He, K., Girshick, R., Sun, J.: Faster R-CNN: towards real-time object detection with region proposal networks. *IEEE Trans. Pattern Anal. Mach. Intell.* **39**, 1137–1149 (2017)
10. Sabour, S., Frosst, N., Hinton, G.E.: Dynamic routing between capsules. In: *Advances in Neural Information Processing Systems (NIPS 2017)*, pp. 3856–3866 (2017)
11. Cohen, N., Sharir, O., Shashua, A.: On the expressive power of deep learning: a tensor analysis. In: *Conference on Learning Theory*, pp. 698–728 (2015)
12. Mahendran, A., Vedaldi, A.: Understanding deep image representations by inverting them. In: *Proceedings of Computer Vision and Pattern Recognition (CVPR 2015)*, pp. 5188–5196 (2015)
13. Tishby, N., Zaslavsky, N.: Deep learning and the information bottleneck principle. In: *Information Theory Workshop (ITW 2015)*, pp. 1–5 (2015)
14. Mallat, S.: Understanding deep convolutional networks. *Philos. Trans. R. Soc. A* **374**, 20150203 (2016)
15. Lin, H.W., Tegmark, M., Rolnick, D.: Why does deep and cheap learning work so well? *J. Stat. Phys.* **168**, 1223–1247 (2017)
16. Vidal, R., Bruna, J., Gyries, R., Soatto, S.: Mathematics of deep learning (2017). [arXiv:1712.04741](https://arxiv.org/abs/1712.04741)
17. Wiatowski, T., Bölcskei, H.: A mathematical theory of deep convolutional neural networks for feature extraction. *IEEE Trans. Inf. Theory* **64**, 1845–1866 (2018)
18. Goldfeld, Z., van den Berg, E., Greenewald, K., Melnyk, I., Nguyen, N., Kingsbury, B., Polyanskiy, Y.: Estimating information flow in neural networks (2018). [arXiv:1810.05728](https://arxiv.org/abs/1810.05728)
19. Achille, A., Rovere, M., Soatto, S.: Critical learning periods in deep neural networks (2017). [arXiv:1711.08856](https://arxiv.org/abs/1711.08856)

20. Szegedy, C., Zaremba, W., Sutskever, I., Bruna, J., Erhan, D., Goodfellow, I., Fergus, B.: Intriguing properties of neural networks (2013). [arXiv:1312.6199](#)
21. Nguyen, A., Yosinski, J., Clune, J.: Deep neural networks are easily fooled: high confidence predictions for unrecognizable images. In: Proceedings of Computer Vision and Pattern Recognition (CVPR 2015), pp. 427–436 (2015)
22. Tanay, T., Griffin, L.: A boundary tilting perspective on the phenomenon of adversarial examples (2016). [arXiv:1608.07690](#)
23. Athalye, A., Sutskever, I.: Synthesizing robust adversarial examples (2017). [arXiv:1707.07397](#)
24. Su, J., Vargas, D.V., Kouichi, S.: One pixel attack for fooling deep neural networks (2017). [arXiv:1710.08864](#)
25. Moosavi-Dezfooli, S.M., Fawzi, A., Fawzi, O., Frossard, P.: Universal adversarial perturbations. In: Proceedings of Computer Vision and Pattern Recognition (CVPR 2017) (2017)
26. Baker, N., Lu, H., Erlikhman, G., Kellman, P.J.: Deep convolutional networks do not classify based on global object shape. *PLoS Comput. Biol.* **14**, e1006613 (2018)
27. Jaderberg, M., Simonyan, K., Zisserman, A., Kavukcuoglu, K.: Spatial transformer networks. In: Advances in Neural Information Processing Systems, pp. 2017–2025 (2015)
28. 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), Volume 9908 of Springer LNCS, pp. 354–370 (2016)
29. 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)
30. Koenderink, J.J., van Doorn, A.J.: Generic neighborhood operators. *IEEE Trans. Pattern Anal. Mach. Intell.* **14**, 597–605 (1992)
31. 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)
32. Lindeberg, T.: A computational theory of visual receptive fields. *Biol. Cybern.* **107**, 589–635 (2013)
33. Lindeberg, T.: Invariance of visual operations at the level of receptive fields. *PLoS ONE* **8**, e66990 (2013)
34. Adelson, E., Bergen, J.: Spatiotemporal energy models for the perception of motion. *J. Opt. Soc. Am. A* **2**, 284–299 (1985)
35. Heeger, D.J.: Normalization of cell responses in cat striate cortex. *Vis. Neurosci.* **9**, 181–197 (1992)
36. Lindeberg, T.: Feature detection with automatic scale selection. *Int. J. Comput. Vis.* **30**, 77–116 (1998)
37. Lindeberg, T.: Dense scale selection over space, time and space-time. *SIAM J. Imaging Sci.* **11**, 407–441 (2018)
38. Fukushima, K.: Neocognitron: a self-organizing neural network model for a mechanism of pattern recognition unaffected by shift in position. *Biol. Cybern.* **36**, 193–202 (1980)
39. Riesenhuber, M., Poggio, T.: Hierarchical models of object recognition in cortex. *Nature* **2**, 1019–1025 (1999)
40. Serre, T., Wolf, L., Bileschi, S., Riesenhuber, M., Poggio, T.: Robust object recognition with cortex-like mechanisms. *IEEE Trans. Pattern Anal. Mach. Intell.* **29**, 411–426 (2007)
41. Bruna, J., Mallat, S.: Invariant scattering convolution networks. *IEEE Trans. Pattern Anal. Mach. Intell.* **35**, 1872–1886 (2013)
42. Poggio, T.A., Anselmi, F.: Visual Cortex and Deep Networks: Learning Invariant Representations. MIT Press, Cambridge (2016)
43. Touryan, J., Felsen, G., Dan, Y.: Spatial structure of complex cell receptive fields measured with natural images. *Neuron* **45**, 781–791 (2005)
44. Carandini, M., Heeger, D.J.: Normalization as a canonical neural computation. *Nat. Rev. Neurosci.* **13**, 51–62 (2012)
45. Goris, R.L.T., Simoncelli, E.P., Movshon, J.A.: Origin and function of tuning diversity in Macaque visual cortex. *Neuron* **88**, 819–831 (2015)
46. Hubel, D.H., Wiesel, T.N.: Receptive fields of single neurones in the cat's striate cortex. *J. Physiol.* **147**, 226–238 (1959)
47. Hubel, D.H., Wiesel, T.N.: Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. *J. Physiol.* **160**, 106–154 (1962)
48. Hubel, D.H., Wiesel, T.N.: Brain and Visual Perception: The Story of a 25-Year Collaboration. Oxford University Press, Oxford (2005)
49. 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), Volume 11603 of Springer LNCS, pp. 328–340 (2019)
50. Iijima, T.: Basic theory on normalization of pattern (in case of typical one-dimensional pattern). *Bull. Electrotech. Lab.* **26**, 368–388 (1962). (in Japanese)
51. Iijima, T.: Basic theory on normalization of two-dimensional pattern. *Stud. Inf. Control* **1**, 15–22 (1963). (in Japanese)
52. Witkin, A.P.: Scale-space filtering. In: Proceedings of 8th International Joint Conference of Artificial Intelligence, Karlsruhe, Germany, pp. 1019–1022 (1983)
53. Koenderink, J.J.: The structure of images. *Biol. Cybern.* **50**, 363–370 (1984)
54. 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)
55. Yuille, A.L., Poggio, T.A.: Scaling theorems for zero-crossings. *IEEE Trans. Pattern Anal. Mach. Intell.* **8**, 15–25 (1986)
56. Lindeberg, T.: Scale-Space Theory in Computer Vision. Springer, Berlin (1993)
57. Lindeberg, T.: Scale-space theory: a basic tool for analysing structures at different scales. *J. Appl. Stat.* **21**, 225–270 (1994)
58. 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 of Ph.D. School on Scale-Space Theory, pp. 75–97. Copenhagen, Denmark, Springer (1996)
59. Florack, L.M.J.: Image Structure. Series in Mathematical Imaging and Vision. Springer, Berlin (1997)
60. Weickert, J., Ishikawa, S., Imiya, A.: Linear scale-space has first been proposed in Japan. *J. Math. Imaging Vis.* **10**, 237–252 (1999)
61. ter Haar Romeny, B.: Front-End Vision and Multi-Scale Image Analysis. Springer, Berlin (2003)
62. 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)
63. Lindeberg, T.: Edge detection and ridge detection with automatic scale selection. *Int. J. Comput. Vis.* **30**, 117–154 (1998)
64. Bretzner, L., Lindeberg, T.: Feature tracking with automatic selection of spatial scales. *Comput. Vis. Image Underst.* **71**, 385–392 (1998)
65. 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 I, pp. 117–133 (2000)
66. Mikolajczyk, K., Schmid, C.: Scale and affine invariant interest point detectors. *Int. J. Comput. Vis.* **60**, 63–86 (2004)
67. Lowe, D.G.: Distinctive image features from scale-invariant keypoints. *Int. J. Comput. Vis.* **60**, 91–110 (2004)
68. Bay, H., Ess, A., Tuytelaars, T., van Gool, L.: Speeded up robust features (SURF). *Comput. Vis. Image Underst.* **110**, 346–359 (2008)

69. Tuytelaars, T., Mikolajczyk, K.: A Survey on Local Invariant Features. Volume 3(3) of Foundations and Trends in Computer Graphics and Vision. Now Publishers, Boston (2008)
70. Lindeberg, T.: Scale selection. In: Ikeuchi, K. (ed.) Computer Vision: A Reference Guide, pp. 701–713. Springer, Berlin (2014)
71. Lindeberg, T.: Image matching using generalized scale-space interest points. *J. Math. Imaging Vis.* **52**, 3–36 (2015)
72. Lindeberg, T., Gårding, J.: Shape-adapted smoothing in estimation of 3-D depth cues from affine distortions of local 2-D structure. *Image Vis. Comput.* **15**, 415–434 (1997)
73. Baumberg, A.: Reliable feature matching across widely separated views. In: Proceedings of Computer Vision and Pattern Recognition (CVPR'00), Hilton Head, SC, vol. I, pp. 1774–1781 (2000)
74. Barnard, E., Casasent, D.: Invariance and neural nets. *IEEE Trans. Neural Netw.* **2**, 498–508 (1991)
75. Lin, C.H., Lucey, S.: Inverse compositional spatial transformer networks. In: Proceedings of Computer Vision and Pattern Recognition (CVPR 2017), pp. 2568–2576 (2017)
76. 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)
77. 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)
78. Hu, P., Ramanan, D.: Finding tiny faces. In: Proceedings of Computer Vision and Pattern Recognition (CVPR 2017), pp. 951–959 (2017)
79. 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)
80. 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)
81. 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)
82. 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)
83. 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)
84. Yu, F., Koltun, V.: Multi-scale context aggregation by dilated convolutions (2015). [arXiv:1511.07122](https://arxiv.org/abs/1511.07122)
85. Yu, F., Koltun, V., Funkhouser, T.: Dilated residual networks. In: Proceedings of Computer Vision and Pattern Recognition (CVPR 2017), pp. 472–480 (2017)
86. 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)
87. 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)
88. 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)
89. 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 (2019). [arXiv:1904.05049](https://arxiv.org/abs/1904.05049)
90. Henriques, J.F., Vedaldi, A.: Warped convolutions: efficient invariance to spatial transformations. *Int. Conf. Mach. Learn.* **70**, 1461–1469 (2017)
91. Esteves, C., Allen-Blanchette, C., Zhou, X., Daniilidis, K.: Polar transformer networks. In: International Conference on Learning Representations (ICLR 2018) (2018)
92. Lenc, K., Vedaldi, A.: Learning covariant feature detectors. In: Proceedings of European Conference on Computer Vision (ECCV 2016). Volume 9915 of Springer LNCS, pp. 100–117 (2016)
93. Zhang, X., Yu, F.X., Karaman, S., Chang, S.F.: Learning discriminative and transformation covariant local feature detectors. In: Proceedings of Computer Vision and Pattern Recognition (CVPR 2017), pp. 6818–6826 (2017)
94. Dieleman, S., Fauw, J.D., Kavukcuoglu, K.: Exploiting cyclic symmetry in convolutional neural networks. In: International Conference on Machine Learning (ICML 2016) (2016)
95. Laptev, D., Savinov, N., Buhmann, J.M., Pollefeys, M.: Tl-pooling: transformation-invariant pooling for feature learning in convolutional neural networks. In: Proceedings Computer Vision and Pattern Recognition (CVPR 2016), pp. 289–297 (2016)
96. Worrall, D.E., Garbin, S.J., Turmukhambetov, D., Brostow, G.J.: Harmonic networks: deep translation and rotation equivariance. In: Proceedings of Computer Vision and Pattern Recognition (CVPR 2017), pp. 5028–5037 (2017)
97. Zhou, Y., Ye, Q., Qiu, Q., Jiao, J.: Oriented response networks. In: Proceedings of Computer Vision and Pattern Recognition (CVPR 2017), pp. 519–528 (2017)
98. Marcos, D., Volpi, M., Komodakis, N., Tuia, D.: Rotation equivariant vector field networks. In: Proceedings of International Conference on Computer Vision (ICCV 2017), pp. 5048–5057 (2017)
99. Cohen, T.S., Welling, M.: Steerable CNNs. In: International Conference on Learning Representations (ICLR 2017) (2017)
100. Weiler, M., Geiger, M., Welling, M., Boomsma, W., Cohen, T.: 3D steerable CNNs: learning rotationally equivariant features in volumetric data. In: Advances in Neural Information Processing Systems (NIPS 2018), pp. 10381–10392 (2018)
101. Weiler, M., Hamprecht, F.A., Storath, M.: Learning steerable filters for rotation equivariant CNNs. In: Proceedings of Computer Vision and Pattern Recognition (CVPR 2018), pp. 849–858 (2018)
102. Worrall, D., Brostow, G.: Cubenet: equivariance to 3D rotation and translation. In: Proceeding of European Conference on Computer Vision (ECCV 2018). Volume 11209 of Springer LNCS, pp. 567–584 (2018)
103. Cheng, G., Han, J., Zhou, P., Xu, D.: Learning rotation-invariant and Fisher discriminative convolutional neural networks for object detection. *IEEE Trans. Image Process.* **28**, 265–278 (2018)
104. Dieleman, S., Willett, K.W., Dambre, J.: Rotation-invariant convolutional neural networks for galaxy morphology prediction. *Mon. Not. R. Astron. Soc.* **450**, 1441–1459 (2015)
105. Cheng, G., Zhou, P., Han, J.: Learning rotation-invariant convolutional neural networks for object detection in VHR optical remote sensing images. *IEEE Trans. Geosci. Remote Sens.* **54**, 7405–7415 (2016)
106. Wang, Q., Zheng, Y., Yang, G., Jin, W., Chen, X., Yin, Y.: Multiscale rotation-invariant convolutional neural networks for lung texture classification. *IEEE J. Biomed. Health Inform.* **22**, 184–195 (2017)
107. Bekkers, E.J., Lafarge, M.W., Veta, M., Eppenhof, K.A.J., Pluim, J.P.W., Duits, R.: Roto-translation covariant convolutional networks for medical image analysis. In: International Conference on Medical Image Computing and Computer-Assisted Intervention MICCAI 2018. Volume 11070 of Springer LNCS, pp. 440–448 (2018)
108. Andrearczyk, V., Depeursinge, A.: Rotational 3D texture classification using group equivariant CNNs (2018). [arXiv:1810.06889](https://arxiv.org/abs/1810.06889)

109. Cohen, T., Welling, M.: Group equivariant convolutional networks. In: International Conference on Machine Learning (ICML 2016), pp. 2990–2999 (2016)
110. Kondor, R., Trivedi, S.: On the generalization of equivariance and convolution in neural networks to the action of compact groups (2018). [arXiv:1802.03690](https://arxiv.org/abs/1802.03690)
111. 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)
112. Lindeberg, T.: Normative theory of visual receptive fields (2017). [arXiv:1701.06333](https://arxiv.org/abs/1701.06333)
113. Lindeberg, T.: Time-causal and time-recursive spatio-temporal receptive fields. *J. Math. Imaging Vis.* **55**, 50–88 (2016)
114. Jhuang, H., Serre, T., Wolf, L., Poggio, T.: A biologically inspired system for action recognition. In: International Conference on Computer Vision (ICCV'07), pp. 1–8 (2007)
115. 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)
116. 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)
117. Chan, T.H., Jia, K., Gao, S., Lu, J., Zeng, Z., Ma, Y.: PCANet: a simple deep learning baseline for image classification? *IEEE Trans. Image Process.* **24**, 5017–5032 (2015)
118. Luan, S., Chen, C., Zhang, B., Han, J., Liu, J.: Gabor convolutional networks. *IEEE Trans. Pattern Anal. Mach. Intell.* **27**, 4357–4366 (2018)
119. 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)
120. Shelhamer, E., Wang, D., Darrell, T.: Blurring the line between structure and learning to optimize and adapt receptive fields (2019). [arXiv:1904.11487](https://arxiv.org/abs/1904.11487)
121. Roux, N.L., Bengio, Y.: Continuous neural networks. In: *Artificial Intelligence and Statistics (AISTATS 2007)*. Volume 2 of *Proceedings of Machine Learning Research*, pp. 404–411 (2007)
122. Liu, L., Chen, J., Fieguth, P., Zhao, G., Chellappa, R., Pietikäinen, M.: From BoW to CNN: two decades of texture representation for texture classification. *Int. J. Comput. Vis.* **127**, 74–109 (2019)
123. Gabor, D.: Theory of communication. *IEE J.* **93**, 429–457 (1946)
124. Bracewell, R.N.: *The Fourier Transform and Its Applications*, 3rd edn. McGraw-Hill, New York (1999)
125. Loog, M.: The jet metric. In: *International Conference on Scale Space and Variational Methods in Computer Vision (SSVM 2007)*. Volume 4485 of *Springer LNCS*, pp. 25–31 (2007)
126. Griffin, L.D.: The second order local-image-structure solid. *IEEE Trans. Pattern Anal. Mach. Intell.* **29**, 1355–1366 (2007)
127. Johnson, E.N., Hawken, M.J., Shapley, R.: The orientation selectivity of color-responsive neurons in Macaque V1. *J. Neurosci.* **28**, 8096–8106 (2008)
128. Koenderink, J.J., van Doorn, A.J.: Receptive field families. *Biol. Cybern.* **63**, 291–298 (1990)
129. 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)
130. Westö, J., May, P.J.C.: Describing complex cells in primary visual cortex: a comparison of context and multi-filter LN models. *J. Neurophysiol.* **120**, 703–719 (2018)
131. Movshon, J.A., Thompson, E.D., Tolhurst, D.J.: Receptive field organization of complex cells in the cat's striate cortex. *J. Physiol.* **283**, 79–99 (1978)
132. Emerson, R.C., Citron, M.C., Vaughn, W.J., Klein, S.A.: Nonlinear directionally selective subunits in complex cells of cat striate cortex. *J. Neurophysiol.* **58**, 33–65 (1987)
133. Touryan, J., Lau, B., Dan, Y.: Isolation of relevant visual features from random stimuli for cortical complex cells. *J. Neurosci.* **22**, 10811–10818 (2002)
134. Rust, N.C., Schwartz, O., Movshon, J.A., Simoncelli, E.P.: Spatiotemporal elements of macaque V1 receptive fields. *Neuron* **46**, 945–956 (2005)
135. Serre, T., Riesenhuber, M.: Realistic modeling of simple and complex cell tuning in the HMAX model, and implications for invariant object recognition in cortex. Technical Report AI Memo 2004-017, MIT Computer Science and Artificial Intelligence Laboratory (2004)
136. Einhäuser, W., Kayser, C., König, P., Körding, K.P.: Learning the invariance properties of complex cells from their responses to natural stimuli. *Eur. J. Neurosci.* **15**, 475–486 (2002)
137. Körding, K.P., Kayser, C., Einhäuser, W., König, P.: How are complex cell properties adapted to the statistics of natural stimuli? *J. Neurophysiol.* **91**, 206–212 (2004)
138. Merolla, P., Boahn, K.: A recurrent model of orientation maps with simple and complex cells. In: *Advances in Neural Information Processing Systems (NIPS 2004)*, pp. 995–1002 (2004)
139. Berkes, P., Wiskott, L.: Slow feature analysis yields a rich repertoire of complex cell properties. *J. Vis.* **5**, 579–602 (2005)
140. Carandini, M.: What simple and complex cells compute. *J. Physiol.* **577**, 463–466 (2006)
141. Hansard, M., Horaud, R.: A differential model of the complex cell. *Neural Comput.* **23**, 2324–2357 (2011)
142. Yamins, D.L.K., DiCarlo, J.J.: Using goal-driven deep learning models to understand sensory cortex. *Nat. Neurosci.* **19**, 356–365 (2016)
143. Hadji, I., Wildes, R.P.: A spatiotemporal oriented energy network for dynamic texture recognition. In: *Proceedings of International Conference on Computer Vision (ICCV 2017)*, pp. 3066–3074 (2017)
144. Mallikarjuna, P., Targhi, A.T., Fritz, M., Hayman, E., Caputo, B., Eklundh, J.O.: The KTH-TIPS2 database. KTH Royal Institute of Technology, Stockholm (2006)
145. Lindeberg, T.: Scale-space for discrete signals. *IEEE Trans. Pattern Anal. Mach. Intell.* **12**, 234–254 (1990)
146. Abramowitz, M., Stegun, I.A. (eds.): *Handbook of Mathematical Functions*. Applied Mathematics Series, 55th edn. National Bureau of Standards, Gaithersburg (1964)
147. Lindeberg, T.: Discrete derivative approximations with scale-space properties: a basis for low-level feature extraction. *J. Math. Imaging Vis.* **3**, 349–376 (1993)
148. Dana, K.J., van Ginneken, B., Nayar, S.K., Koenderink, J.J.: Reflectance and texture of real-world surfaces. *ACM Trans. Graph.* **18**, 1–34 (1999)
149. Varma, M., Zisserman, A.: A statistical approach to material classification using image patch exemplars. *IEEE Trans. Pattern Anal. Mach. Intell.* **31**, 2032–2047 (2009)
150. Xu, Y., Yang, X., Ling, H., Ji, H.: A new texture descriptor using multifractal analysis in multi-orientation wavelet pyramid. In: *Proceedings of Computer Vision and Pattern Recognition (CVPR 2010)*, pp. 161–168 (2010)
151. Cimpoi, M., Maji, S., Vedaldi, A.: Deep filter banks for texture recognition and segmentation. In: *Proceedings of Computer Vision and Pattern Recognition (CVPR 2015)*, pp. 3828–3836 (2015)
152. Liu, L., Lao, S., Fieguth, P.W., Guo, Y., Wang, X., Pietikäinen, M.: Median robust extended local binary pattern for texture classification. *IEEE Trans. Image Process.* **25**, 1368–1381 (2016)

153. Liu, L., Long, Y., Fieguth, P.W., Lao, S., Zhao, G.: BRINT: binary rotation invariant and noise tolerant texture classification. *IEEE Trans. Image Process.* **23**, 3071–3084 (2014)
154. Schaefer, G., Doshi, N.P.: Multi-dimensional local binary pattern descriptors for improved texture analysis. In: *Proceedings of International Conference on Pattern Recognition (ICPR 2012)*, pp. 2500–2503 (2012)
155. Ojala, T., Pietikäinen, M., Maenpää, T.: Multiresolution gray-scale and rotation invariant texture classification with local binary patterns. *IEEE Trans. Pattern Anal. Mach. Intell.* **24**, 971–987 (2002)
156. Chang, C.C., Lin, C.J.: LIBSVM: a library for support vector machines. *ACM Trans. Intell. Syst. Technol.* **2**, 27 (2011)
157. Liu, L., Fieguth, P., Guo, Y., Wang, Z., Pietikäinen, M.: Local binary features for texture classification: taxonomy and experimental study. *Pattern Recognit.* **62**, 135–160 (2017)
158. Crosier, M., Griffin, L.D.: Using basic image features for texture classification. *Int. J. Comput. Vis.* **88**, 447–460 (2010)

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, focus of attention and computational modelling of biological vision. He has developed theories and methodologies for continuous and discrete scale-space representation, visual and auditory receptive fields, 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 the author of the book *Scale-Space Theory in Computer Vision*.