

Learning Steerable Filters for Rotation Equivariant CNNs

Maurice Weiler^{1,2}

Fred A. Hamprecht²

Martin Storath²

¹AMLab / QUVA Lab, University of Amsterdam

²HCI/IWR, University of Heidelberg

m.weiler@uva.nl

{fred.hamprecht, martin.storath}@iwr.uni-heidelberg.de

Abstract

In many machine learning tasks it is desirable that a model's prediction transforms in an equivariant way under transformations of its input. Convolutional neural networks (CNNs) implement translational equivariance by construction; for other transformations, however, they are compelled to learn the proper mapping. In this work, we develop Steerable Filter CNNs (SFCNNs) which achieve joint equivariance under translations and rotations by design. The proposed architecture employs steerable filters to efficiently compute orientation dependent responses for many orientations without suffering interpolation artifacts from filter rotation. We utilize group convolutions which guarantee an equivariant mapping. In addition, we generalize He's weight initialization scheme to filters which are defined as a linear combination of a system of atomic filters. Numerical experiments show a substantial enhancement of the sample complexity with a growing number of sampled filter orientations and confirm that the network generalizes learned patterns over orientations. The proposed approach achieves state-of-the-art on the rotated MNIST benchmark and on the ISBI 2012 2D EM segmentation challenge.

1. Introduction

Convolutional neural networks are extremely successful predictive models when the input data has spatial structure. One principal reason is that the convolution operation exhibits translational equivariance so that feature extraction is independent of the spatial position. For many types of images it is desirable to make feature extraction orientation independent as well. Typical examples are biomedical microscopy images or astronomical data which do not show a prevailing global orientation. Consequently, the output of a network processing such data should be equivariant w.r.t. the orientation of its input – if the input is rotated, the output should transform accordingly. Even when there is a predominant direction in an image as a whole, the low level features in the first layers such as edges usually appear in all orientations; see e.g. the filterbanks visualized in [1].

In both cases, conventional CNNs are compelled to learn rotated versions of the same filter, introducing redundant degrees of freedom and increasing the risk of overfitting.

1.1. Contribution

We propose a rotation-equivariant CNN architecture which shares weights over filter orientations to improve generalization and to reduce sample complexity. A key property of our network is that its filters are learned such that they are steerable. This approach avoids interpolation artifacts which can be severe at the small length scale of typical filter kernels. We accomplish the steerability of the learned filters by representing them as linear combinations of a fixed system of atomic steerable filters.

In all intermediate layers of the network, we utilize group convolutions to ensure an equivariant mapping of feature maps. Group-convolutional networks were proposed by Cohen and Welling [2] who considered four filter orientations. An advantage of our construction is that we can achieve an arbitrary angular resolution w.r.t. the sampled filter orientations. Indeed, our experiments show that results improve significantly when using more than four orientations.

An important practical aspect of CNNs is a proper weight initialization. Since the weights to be learned serve as expansion coefficients for the steerable function space, common weight initialization schemes need to be adapted. Here, we generalize the results found by Glorot and Bengio [3] and He et al. [4] to networks which learn filters as a composition of (not necessarily steerable) atomic filters.

Our network achieves state-of-the-art results on two important rotation-equivariant/invariant recognition tasks: (i) The proposed approach is the first to obtain an accuracy higher than 99% on the rotated MNIST dataset, which is the standard benchmark for rotation-invariant classification. (ii) A processing pipeline based on the proposed SFCNN layers ranks among the top three entries in the ISBI 2012 electron microscopy segmentation challenge [5].

Figure 1 gives an overview over the key concepts utilized in Steerable Filter CNNs.

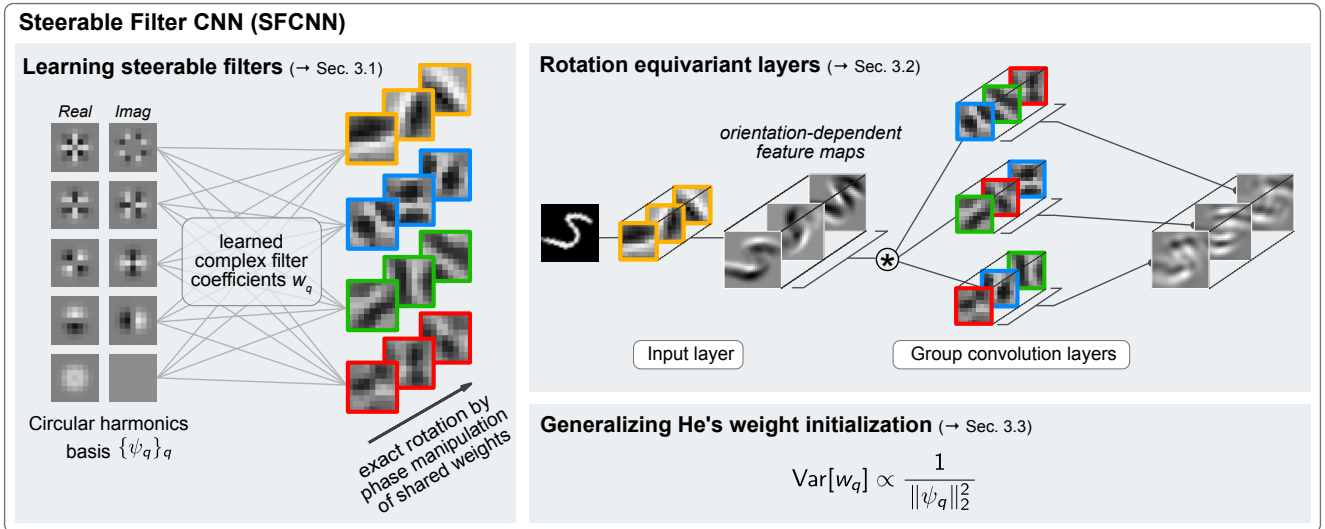


Figure 1: Key concepts of the proposed Steerable Filter CNN: The filters are parameterized in a steerable function space with shared weights over filter orientations. Exact filter rotations are achieved by a phase manipulation of the expansion coefficients w_q . All layers are designed to be jointly translation and rotation equivariant. The weights w_q serve as expansion coefficients of a fixed filter basis $\{\psi_q\}_q$ rather than pixel values. Therefore, we adapt He’s weight initialization scheme to this more general case which implies to normalize the basis filter energies.

2. Equivariance properties of CNNs

Equivariance is the property of a function to commute with the actions of a symmetry group acting on its domain and codomain. Formally, given a transformation group G , a function $f : X \rightarrow Y$ is said to be equivariant if

$$f(\varphi_g^X(x)) = \varphi_g^Y(f(x)) \quad \forall g \in G, x \in X,$$

where $\varphi_g^{(\cdot)}$ denotes a *group action* in the corresponding space. A special case of equivariance is invariance for which $\varphi_g^Y = \text{id}$.

In many machine learning tasks a set of transformations is known a-priori under which the prediction should transform in an equivariant way. Including such knowledge directly into the model can greatly facilitate learning by freeing up model capacity for other factors of variation. As an example consider a segmentation problem where the goal is to learn a mapping from an image space \mathcal{I} to label images in \mathcal{L} , which we formalize by a ground truth segmentation map $\mathcal{S} : \mathcal{I} \rightarrow \mathcal{L}$. The learning process involves fitting a model $\mathcal{M} : \mathcal{I} \rightarrow \mathcal{L}$ to approximate the ground truth. For segmentation tasks, however, translations of the input image $I \in \mathcal{I}$ should typically lead to a translated segmentation map. Specifically, one has

$$\mathcal{S}(\mathcal{T}_d I) = \mathcal{T}_d \mathcal{S}(I) \quad \forall d \in \mathbb{R}^2, I \in \mathcal{I}, \quad (1)$$

where \mathcal{T}_d is an action of the translation group $T = (\mathbb{R}^2, +)$ which shifts the image or segmentation by $d \in \mathbb{R}^2$. The group action partitions the image space in equivalence classes $T.I = \{\mathcal{T}_d I \mid d \in \mathbb{R}^2\}$ which are known as *group orbits* and comprise all images that are related by the action. Note that the translation equivariance (1) of the ground truth

segmentation function implies a mapping of whole orbits in \mathcal{I} to orbits in \mathcal{L} . It is therefore possible to reformulate the ground truth as $\tilde{\mathcal{S}} : \mathcal{I}/T \rightarrow \mathcal{L}/T$, where $(\cdot)/T$ denotes the *quotient space* resulting from collapsing equivalent images in an orbit to a single element. Instead of fitting an unrestricted model \mathcal{M} to \mathcal{S} it is advantageous to incorporate the transformation behavior into the model by construction. The crucial consequence is that this reduces the hypothesis space to models $\tilde{\mathcal{M}} : \mathcal{I}/T \rightarrow \mathcal{L}/T$.

CNN layers, which transform feature maps ζ by convolving them with filters Ψ , are by construction equivariant under translations, that is, $(\mathcal{T}_d \zeta) * \Psi = \mathcal{T}_d (\zeta * \Psi)$. Therefore, their hypothesis space is restricted to $\tilde{\mathcal{M}}$.¹ As consequence, patterns learned at one specific location evoke the same response at each other location which leads to reduced sample complexity and enhanced generalization.

Besides translations, there are often further transformations like rotations, mirroring or dilations under which the model should be equivariant. Enforcing equivariance under an extended transformation group G leads to an enhanced generalization over larger orbits $G.I$ and reduces the hypothesis space further to $\tilde{\mathcal{M}} : \mathcal{I}/G \rightarrow \mathcal{L}/G$.

3. Steerable Filter CNNs

Here, we develop Steerable Filter CNNs (SFCNNs) which achieve equivariance under joint translations and discrete rotations. The key concept leading to translation equivariance of CNNs is *translational weight sharing*. We

¹In practice one often uses strided pooling layers which make the prediction more robust to local deformations but reduce the equivariance to a subgroup determined by the stride.

extend the transformation group under which our networks' layers are equivariant by additionally *sharing weights over filter orientations*. This implies to perform convolutions with several rotated versions of each filter. The rotational weight sharing leads to an improved sample complexity and to an enhanced generalization over orbits consisting of images connected by translations and discrete rotations.

In the following sections we introduce our parametrization of steerable filters, propose the network design in terms of these filters and derive a weight initialization scheme adapted to the filter parametrization. For the formal derivations we assume the images, feature maps and filters to be defined on the continuous domain \mathbb{R}^2 . The effects resulting from a discretized implementation are investigated in the experimental section.

3.1. Parametrization of the steerable filters

At the heart of convolutional neural networks lies the concept of learning filter kernels. Our construction demands for filters whose responses can be computed accurately and economically for several filter orientations. Simultaneously the filters should not be restricted in their expressive power, i.e. in the patterns to be learned. All of these requirements are met by learning linear combinations of a system of steerable filters. Here we describe a suitable construction of steerable filters for learning in CNNs.

A filter Ψ is rotationally steerable in the sense of Hel-Or and Teo [6], when its rotation by an arbitrary angle θ can be expressed in a function space spanned by a fixed set of atomic basis functions $\{\psi_q\}_{q=1}^Q$. This definition includes the classical formulation of steerability by Freeman and Adelson [7] as a specific choice of basis. Formally, a steerable filter $\Psi : \mathbb{R}^2 \rightarrow \mathbb{R}$ satisfies

$$\rho_\theta \Psi(x) = \sum_{q=1}^Q \kappa_q(\theta) \psi_q(x), \quad (2)$$

for all angles $\theta \in (-\pi, \pi]$ and for angular expansion coefficient functions κ_q . Here ρ_θ denotes both the rotation operator defined by $\rho_\theta \Psi(x) = \Psi(\rho_{-\theta} x)$ when acting on a function as well as a counterclockwise rotation by the angle θ when acting on a coordinate vector. As pointed out by Freeman and Adelson [7], the rotation by steerability is analytic and exact even for signals sampled on a grid. In contrast to rotations by interpolation the approach does not suffer from interpolation artifacts. An important practical consequence of steerability is that the response of each orientation can be synthesized from the atomic responses $f * \psi_q$; that is, $(f * \rho_\theta \Psi)(x) = \sum_{q=1}^Q \kappa_q(\theta) (f * \psi_q)(x)$.

A basis of a steerable function space which is particularly easy to handle is given by circular harmonics; see e.g. [8, 9]. They are defined by a sinusoidal angular part multiplied with a radial function $\tau : \mathbb{R}^+ \rightarrow \mathbb{R}$, i.e.

$$\psi_k(r, \phi) = \tau(r) e^{ik\phi}, \quad (3)$$




Figure 2: Illustration of the circular harmonics $\psi_{jk}(r, \phi) = \tau_j(r) e^{ik\phi}$ sampled on a 9×9 grid. Each row shows a different radial part j , the angular frequencies are arranged in the columns. For larger scales there are higher frequency filters not shown here.

where (r, ϕ) denote polar coordinates of $x = (x_1, x_2)$ and $k \in \mathbb{Z}$ is the angular frequency. By construction, ψ_k can be rotated by multiplication with a complex exponential,

$$\rho_\theta \psi_k(x) = e^{-ik\theta} \psi_k(x). \quad (4)$$

In our network, we utilize a system of circular harmonics ψ_{jk} with $j = 1, \dots, J$, and $k = 0, \dots, K_j$ where the additional index j controls the radial part of $\psi_{jk} = \tau_j(r) e^{ik\phi}$. Figure 2 shows the real and imaginary parts of the atoms used in the experiments where we chose Gaussian radial parts $\tau_j(r) = \exp(-(r - \mu_j)^2 / 2\sigma^2)$ with $\mu_j = j$. The maximum angular frequencies K_j are limited to the point where aliasing effects occur. We found this system to be convenient for learning as the filters are approximately orthogonal and radially localized.

The learned filters are then defined as linear combinations of the elementary filters, that is,

$$\tilde{\Psi}(x) = \sum_{j=1}^J \sum_{k=0}^{K_j} w_{jk} \psi_{jk}(x), \quad (5)$$

with weights $w_{jk} \in \mathbb{C}$. The complex phase of the weights allows rotating the atomic filters *with respect to each other*. Such a composed filter can subsequently be steered *as a whole* by phase manipulation of the atoms via

$$\rho_\theta \tilde{\Psi}(x) = \sum_{j=1}^J \sum_{k=0}^{K_j} w_{jk} e^{-ik\theta} \psi_{jk}(x). \quad (6)$$

We select a single orientation by taking their real part

$$\Psi(x) = \text{Re} \tilde{\Psi}(x) \quad (7)$$

and let $\rho_\theta \Psi = \text{Re} \rho_\theta \tilde{\Psi}$.

3.2. Equivariant network architecture

The basic building blocks of the proposed SFCNN are three equivariant layer types which we introduce in this section.

Input layer: The first layer $l = 1$ of our network ingests an image with C channels $I_c : \mathbb{R}^2 \rightarrow \mathbb{R}$, $c = 1, \dots, C$ and convolves these with \hat{C} rotated filters $\rho_\theta \Psi_{\hat{c}c}^{(1)}$, where $\Psi_{\hat{c}c}^{(1)} : \mathbb{R}^2 \rightarrow \mathbb{R}$, $\hat{c} = 1, \dots, \hat{C}$, are filter channels of the

form (7). This results in pre-nonlinearity features

$$\begin{aligned} y_{\hat{c}}^{(1)}(x, \theta) &= \sum_{c=1}^C \left(I_c * \rho_{\theta} \Psi_{\hat{c}c}^{(1)} \right) (x) \\ &= \sum_{c=1}^C \left(I_c * \operatorname{Re} \sum_{j=1}^J \sum_{k=0}^{K_j} w_{\hat{c}cj k} e^{-ik\theta} \psi_{jk} \right) (x) \\ &= \operatorname{Re} \sum_{c=1}^C \sum_{j=1}^J \sum_{k=0}^{K_j} w_{\hat{c}cj k} e^{-ik\theta} (I_c * \psi_{jk}) (x), \end{aligned} \quad (8)$$

where the filters are rotated by in total Λ equidistant orientations $\theta \in \Theta = \{0, \dots, 2\pi \frac{\Lambda-1}{\Lambda}\}$. In this setting the rotational weight sharing is reflected by the phase manipulation of the weights $w_{\hat{c}cj k}$ which themselves are independent of the angle θ . A higher resolution in orientations can be achieved by simply expanding the tensor containing the phase-factors.

As usual, after the convolution step a bias $\beta_{\hat{c}}^{(1)}$ is added and a nonlinearity σ is applied, so that we end up with the first layer's feature map given by

$$\zeta_{\hat{c}}^{(1)}(x, \theta) = \sigma \left(y_{\hat{c}}^{(1)}(x, \theta) + \beta_{\hat{c}}^{(1)} \right).$$

Note that the resulting representation $\zeta_{\hat{c}}^{(1)}$ depends on a spatial location x and an orientation angle θ , i.e. on the transformation group applied to the filters.

Group-convolutional layers: To process the resulting feature maps further we utilize group convolutions which naturally generalize spatial convolutions from translations to more general transformation groups. Given a feature map $\zeta : G \rightarrow \mathbb{R}$ and a filter $\Psi : G \rightarrow \mathbb{R}$ living on a group G , their group convolution is defined by $(\zeta \circledast \Psi)(g) = \int_G \zeta(h) \Psi(h^{-1}g) d\lambda(h)$, where we use the symbol \circledast to distinguish group convolutions from the spatial convolution operator $*$, and λ denotes a Haar measure. The resulting feature map is again a function on the group. In analogy to spatial convolutions, group convolutions are equivariant under the group operation, i.e. $(\varphi_h(\zeta) \circledast \Psi)(g) = \varphi_h(\zeta \circledast \Psi)(g)$, $\forall h, g \in G$, where φ_h is given by $\varphi_h \zeta(g) = \zeta(h^{-1}g)$. For a deeper discussion of group convolutions in neural networks we refer to [2].

The feature maps calculated by the input layer are functions on the semidirect product group $\mathbb{R}^2 \rtimes \Theta \leq \operatorname{SE}(2)$. Keeping the parameterization by (x, θ) , the group convolutions with summation over input channels can be explicitly instantiated as

$$\begin{aligned} y_{\hat{c}}^{(l)}(x, \theta) &= \sum_{c=1}^C \left(\zeta_c^{(l-1)} \circledast \Psi_{\hat{c}c}^{(l)} \right) (x, \theta) \\ &= \sum_{c=1}^C \sum_{\phi \in \Theta} \int_{\mathbb{R}^2} \zeta_c^{(l-1)}(u, \phi) \Psi_{\hat{c}c}^{(l)}((u, \phi)^{-1}(x, \theta)) du \\ &= \sum_{c=1}^C \sum_{\phi \in \Theta} \left(\zeta_c^{(l-1)}(\cdot, \phi) * \rho_{\phi} \Psi_{\hat{c}c}^{(l)}(\cdot, \theta - \phi) \right) (x) \\ &= \sum_{c=1}^C \sum_{\phi \in \Theta} \left(\zeta_c^{(l-1)}(\cdot, \phi) * \mathcal{R}_{\phi} \Psi_{\hat{c}c}^{(l)}(\cdot, \theta) \right) (x). \end{aligned} \quad (9)$$

Here the multiplication with the inverse group element, $(u, \phi)^{-1}(x, \theta) = (\rho_{-\phi}(x - u), \theta - \phi)$, was evaluated by switching to a representation of the group. We further introduced the action \mathcal{R}_{ϕ} defined by

$$\mathcal{R}_{\phi} \Psi(x, \theta) := \rho_{\phi} \Psi(x, \theta - \phi)$$

which transforms functions on the group by rotating them spatially and shifting their orientation components cyclically. The above equation reveals that the group convolution can be decomposed into a spatial convolution, rotation and linear combination. In analogy to the first layer we make use of the steerable filters which on the group are defined by $\Psi_{\hat{c}c}^{(l)}(x, \theta) = \operatorname{Re} \sum_{j=1}^J \sum_{k=0}^{K_j} w_{\hat{c}cj k} \theta \psi_{jk}(x)$. Note that the additional orientation dimension is reflected by an additional index of the weight tensor. Inserting the steerable filters in (9) we obtain the pre-nonlinearity feature maps of the group-convolutional layers

$$\begin{aligned} y_{\hat{c}}^{(l)}(x, \theta) &= \sum_{c=1}^C \sum_{\phi \in \Theta} \left(\zeta_c^{(l-1)}(\cdot, \phi) * \operatorname{Re} \sum_{j,k} w_{\hat{c}cj k, \theta - \phi} e^{-ik\phi} \psi_{jk} \right) (x) \\ &= \operatorname{Re} \sum_{c=1}^C \sum_{\phi \in \Theta} \sum_{j,k} w_{\hat{c}cj k, \theta - \phi} e^{-ik\phi} \left(\zeta_c^{(l-1)}(\cdot, \phi) * \psi_{jk} \right) (x). \end{aligned} \quad (10)$$

As before, a bias $\beta_{\hat{c}}^{(l)}$ is added and the activation function σ is applied, $\zeta_{\hat{c}}^{(l)}(x, \theta) = \sigma(y_{\hat{c}}^{(l)}(x, \theta) + \beta_{\hat{c}}^{(l)})$.

By the linearity of the steerability and the convolution, one can implement the layers either by a direct convolution with linearly combined filters, or by linearly combining the responses of the atomic filters. We implemented both approaches and found that in typical operation regimes the first option is faster since the kernels to be linearly combined have a smaller spatial extent than the atomic responses of the second option.

Output layer: After the last group-convolutional layer we extract the information of interest for the specific task. For rotation-invariant classification we pool globally over both the orientation dimension and the remaining spatial resolution. A pooling over orientations is also done for rotation-equivariant segmentation where spatial dimensions remain and the output rotates according to the rotation of the network's input. If the orientation itself is of interest it could be kept as extra feature.

Equivariance: Each individual layer $L_{(\cdot)}$ of the network is equivariant under joint translations and rotations in the group $\mathbb{R}^2 \rtimes \Theta$: Rotating the input image leads to a transformation $L_{\text{in}}(\rho_{\phi} I) = \mathcal{R}_{\phi} L_{\text{in}}(I)$ of the first layer's feature maps. The subsequent group-convolutional layers then transform like $L_{\text{gconv}}(\mathcal{R}_{\phi} \zeta) = \mathcal{R}_{\phi} L_{\text{gconv}}(\zeta)$. When using orientation pooling in the output layer the resulting feature

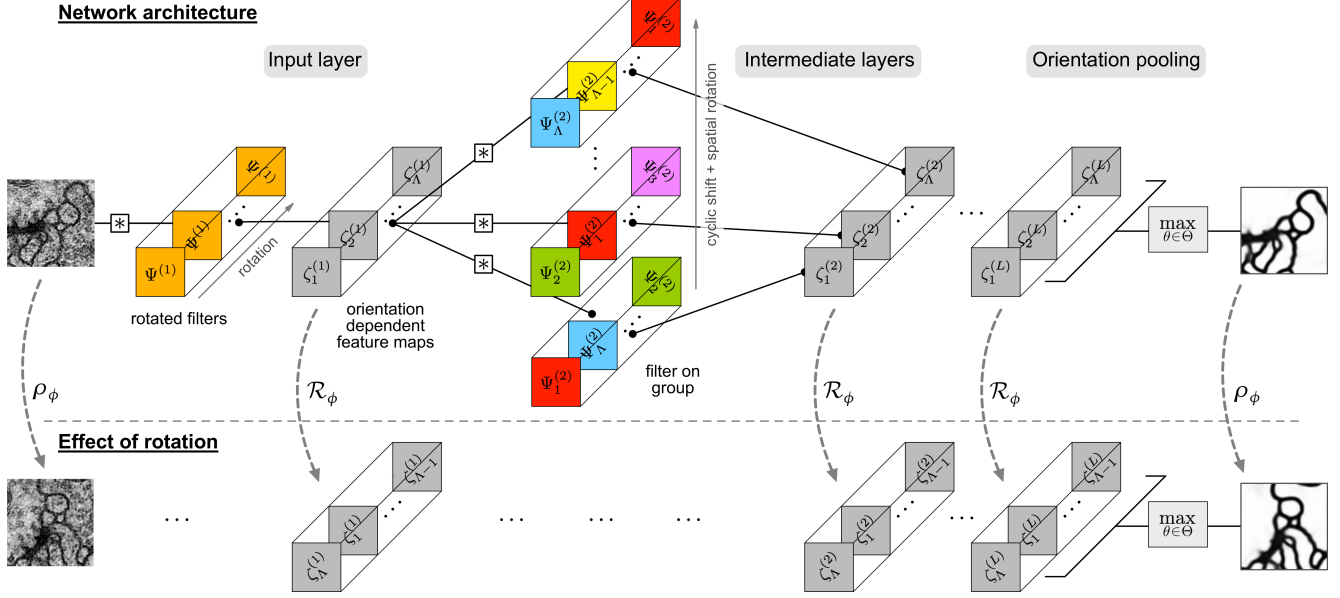


Figure 3: *Top*: Basic structure of a typical SFCNN for rotation-equivariant segmentation. For clarity, we display only a single group-convolutional layer and a single feature channel and omit pooling and normalization layers. Rotated Greek letters represent the spatial orientations of the filters and the feature maps. Orientation components are abbreviated as subscript, i.e. $\Psi_\lambda = \Psi(\cdot, \theta_\lambda)$. Filters in the same color share their weights as they are connected by rotations. The weight sharing of the filters on the group is prescribed by the group convolution (9). After the last group-convolutional layer we pool over orientations to obtain predictions which are invariant under rotations of local patches in the field of view. *Bottom*: Visualization of the layerwise rotation-equivariance. Applying a rotation ρ_ϕ to the input image results in a joint spatial rotation operation and cyclic shift over orientation indices \mathcal{R}_ϕ of the feature maps $\zeta^{(l)}$. This transformation behavior can be understood intuitively when paying attention to the relative orientation of each layer’s input and filters.

maps are rotated: $L_{\text{out}}(\mathcal{R}_\phi \zeta) = \rho_\phi L_{\text{out}}(\zeta)$. Overall, this implies the equivariance of a whole network,

$$(L_{\text{out}} \circ L_{\text{gconv}}^d \circ L_{\text{in}})(\rho_\phi I) = \rho_\phi (L_{\text{out}} \circ L_{\text{gconv}}^d \circ L_{\text{in}}(I)),$$

where d is the number of group-convolutional layers. The layers’ equivariance is proven in the supplementary.

The top part of Figure 3 visualizes the building blocks of a typical SFCNN for rotation-equivariant segmentation. An overview over the transformation behavior of the feature maps under rotation of the input is given in the bottom part. The spatial rotation and cyclic shift over orientation channels \mathcal{R}_ϕ of the feature maps on the group can be understood intuitively when paying attention to the relative orientation of each layer’s input and filters.

Compared to a conventional CNN which independently learns filters in Λ orientations in a rotation-invariant recognition task, a corresponding SFCNN consumes Λ times less parameters to extract the same representation.

SFCNN incur a small computational overhead for building the filter kernels from the circular harmonics basis which we found to be negligible. The computational cost of SFCNNs is therefore equivalent to that of a conventional CNN when the effective number of channels coincide, i.e. when $I_{\text{CNN}} = \Lambda I_{\text{SFCNN}}$.

3.3. Generalizing He’s weight initialization scheme

An important practical aspect of training deep networks is an appropriate initialization of their weights. When the

weights’ variances are chosen too high or low, the signals propagating through the network are amplified or suppressed exponentially with depth. Glorot and Bengio [3] and He et al. [4] investigated this issue and came up with initialization schemes which are accepted as a standard for random weight initialization. In contrast to [3] and [4] our filters are not parameterized in a pixel basis but as a linear combination of a system of atomic filters with weights serving as expansion coefficients. To be specific, we consider filters $\Psi_{\hat{c}cx} = \sum_{q=1}^Q w_{\hat{c}cq} \psi_{qx}$ which are built from Q , not necessarily steerable, real valued atomic filters which map C input channels to \hat{C} output channels. This assumption is more general than that of the aforementioned works since they only consider the pixel basis $\psi_{qx}^{\text{Dirac}} = \delta_{q,x}$, i.e. atomic filters which are zero everywhere but at one pixel.

Most of the further assumptions are identical to those in [4]: We assume the activations and gradients to be i.i.d. and to be independent from the weights. Further, the weights themselves are initialized to be mutually independent and have zero mean. An important difference is that we do *not* restrict the weights to be identically distributed because of the inherent asymmetry of the different atomic filters. All biases are initialized to be zero and the nonlinearities are chosen to be ReLUs. These assumptions lead to the initialization conditions

$$\text{Var}[w_q] = \frac{2}{CQ \|\psi_q\|_2^2} \quad \text{or} \quad \text{Var}[w_q] = \frac{2}{\hat{C}Q \|\psi_q\|_2^2}$$

for the forward or backward pass, respectively. A detailed derivation is given in the supplementary material.

As discussed in [4], the difference between both initializations cancels out for intermediate layers. Note that our results include those of He et al. [4], that is, $\text{Var}[w_q] = \frac{2}{n_{\text{in}}}$ or $\text{Var}[w_q] = \frac{2}{n_{\text{out}}}$, for $\psi_{qx}^{\text{Dirac}} = \delta_{q,x}$ with $\|\psi_{qx}^{\text{Dirac}}\|_2^2 = 1$. We further want to point out that the learned filters are combined of products $w_q \psi_q$ which implies that the factors $\|\psi_q\|_2^2$ counterbalance different energies of the basis filters. A convenient way to initialize the network is hence to normalize all filters to unit norm and subsequently initialize the weights uniformly by $\text{Var}[w_q] = \frac{2}{CQ}$ or $\text{Var}[w_q] = \frac{2}{\tilde{C}Q}$. In our group-convolutional layers the filters additionally comprise orientation channels. From the perspective of weight initialization these have the same effect as conventional channels, therefore we propose to normalize their weights variance with an additional factor of Λ . We emphasize that using normalization layers like batch normalization does not obviate the need for a proper weight initialization. This is because such layers scale activations as a whole while our initialization conditions indicates that the relative scale of the summands contributing to each activation needs to be adapted. Further details, in particular on initializing weights of complex-valued filters, are given in the supplementary material.

4. Prior and related work

A priori knowledge about transformation-invariance of images can be exploited in manifold ways. A commonly utilized technique is data augmentation, see e.g. [10]. The basic idea is to enrich the training set by transformed samples. Augmenting datasets allows to train larger models and is easily applicable without modifying the network architectures. When the augmenting transformations form a group G the additional images $I \subseteq \mathcal{I}$ lie on the orbit $G.I$. In contrast to equivariant models the hypothesis space is not restricted to the quotient space \mathcal{I}/G under the utilized symmetry group but the equivariance needs to be learned explicitly by the network. This demands for a high learning capacity which makes the network prone to overfitting.

Recent work focuses on incorporating equivariance to various transformations directly into the network’s architecture. Invariance to specific transformations can be achieved by applying them to the input and subsequently pooling their responses [11, 12, 13]. In [14] the regions in symmetry space to pool over are learned to become invariant only to nuisance deformations. Another approach is to resample the input and apply standard convolutions. Henriques and Vedaldi [15] achieve equivariance w.r.t. Abelian symmetry groups by fixing a sampling grid according to the symmetry while in [16, 17] the network itself estimates the grid. In [18] transformations are dealt with by convolving with

filters which are steered by a subnetwork.

In particular, there has recently been a considerable interest in rotation-equivariant CNNs. The work [19] introduces four operations which are easily included into existing networks and enrich both the batch- and feature dimension with transformed versions of their content. In [2], the feature maps resulting from transformed filters are treated as functions of the corresponding symmetry-group which allows to use group-convolutional layers. As their computational cost is coupled to the size of the group, Cohen and Welling [20] propose to alternatively use steerable representations as composition of elementary feature types. Besides translations and rotations, the aforementioned works also incorporate reflections, i.e. they operate on the dihedral group. Their current limitation is the restriction to rotations by the angle $\frac{\pi}{2}$, thus to four orientations. In [21], several rotated versions of the same image are sent through a conventional CNN. The resulting features are subsequently pooled over the orientation dimension. The approach can be easily extended to other transformations. On the downside, the equivariance is only w.r.t. global transformations. Marcos et al. [22] perform convolutions with rotated versions of a each filter in a shallow network followed by a global pooling over orientations. These ideas were extended to networks which additionally propagate the orientation of the maximum response [23]. In both approaches the filter rotation is based on bicubic interpolation, allowing for fine resolutions with respect to the orientation but causing interpolation artifacts. Worrall et al. [24] achieve continuous resolution in orientations by working with complex valued steerable filters and feature maps. However, this requires the angular frequencies of the feature maps to be kept disentangled. Rotation-equivariant feature extraction can also be achieved by using group-convolutional scattering transforms [25]. A fundamental difference to our work is that the filter banks are fixed rather than learned.

5. Experimental results

We evaluate the proposed SFCNNs on two datasets exhibiting rotational symmetries. On the rotated MNIST dataset we first investigate specific network properties like the accuracy’s dependence on the number of sampled orientations and the generalization of learned patterns over orientations. With the insights gained in these experiments we benchmark the model and the proposed initialization scheme. To evaluate the segmentation capabilities of SFCNNs on real world data we run a further experiment on the ISBI 2012 EM segmentation challenge.

5.1. Rotated MNIST

In our first experiments we investigate the equivariance properties of the proposed network architecture on the **rotated MNIST dataset** (mnist-rot) which is the standard




Figure 4: *Left*: Test error versus number of sampled filter orientations for different training subsets from mnist-rot. Shaded regions highlight the standard deviations over several runs. The accuracy improves significantly with increasing angular resolution until it saturates at around 12 to 16 orientations. *Right*: Rotational generalization capabilities of a conventional CNN and a SFCNN with $\Lambda = 16$ using different data augmentation strategies. In this experiment the training set consists of unrotated MNIST digits while the test set for each angle contains the remaining digits, rotated to the corresponding angle.

benchmark for rotation-equivariant models. The dataset contains the handwritten digits of the classical MNIST dataset, rotated to random orientations in $[0, 2\pi)$. It is split in 12000 training and 50000 test images; model selection is done by training on 10000 images and validating on the 2000 remaining samples in the training set.

For our initial experiments we utilize the classification-SFCNN given in Table 1 in the supplementary material as baseline. It consists of one steerable input layer which maps the input images to the group, five following group-convolutional layers and three fully connected layers. After every two steerable filter layers we perform a spatial 2×2 max-pooling. The orientation dimension and the remaining spatial dimensions are pooled out globally after the last convolutional layer. Details on the further training setup are given in the supplementary material.

Sampled orientations: The number of sampled orientations Λ is a parameter specific to our network, so we first explore its influence on the test accuracy. We are further interested in the network’s sample complexity, i.e. the dependence on the size of the training set. The accuracies resulting when varying these parameters are reported in Figure 4 (left). As expected, the test error and its standard deviation decrease with the size of the training data set. We observe that the accuracy improves significantly when increasing the number of orientations until it saturates at around 12 to 16 angles. Up to this point, the gain of adding more sampled orientations is considerable. For example, in almost all cases, increasing the angular resolution from 2 to 4 sampled orientations provides a higher gain in accuracy than sticking with 2 orientations and doubling the number of training samples. We want to emphasize that the possibility of SFCNNs to go beyond the four sampled orientations of [19, 2, 20] leads to a significant gain in accuracy. Note that the case $\Lambda = 1$ corresponds to conventional CNNs.

Rotational generalization: In order to test how well the networks generalize learned patterns over orientations we

conduct an experiment where we train them on unrotated digits and record their accuracy over the orientation of rotated digits. Specifically, we take the first 12000 samples of the conventional MNIST dataset to train a SFCNN with $\Lambda = 16$ as well as a conventional CNN of comparable size using either no augmentation, augmentation by rotations which are multiples of either $\frac{\pi}{4}$ or $\frac{\pi}{2}$ or augmentation by rotations which are densely sampled from $[0, 2\pi)$. As test set we take the remaining 58000 samples and record the test errors’ dependence on the orientation of this dataset. To obtain a fair comparison between the networks we experiment with conventional CNNs with the same number of parameters or the same number of channels like the SFCNN. Since both show the same behavior we only report the accuracies of the network with the same number of channels which performs slightly better. The results are plotted in Figure 4 (right). One can see that, lacking rotational equivariance, the conventional CNN does not generalize well over orientations. When using rotational augmentation the error reduces considerably on average, it grows, however, for small angles in a neighborhood of zero. This is the case because the network needs to learn to detect the augmented samples additionally which demands an increased learning capacity. The SFCNN on the other hand generalizes quite well over orientations even without augmentation. In continuous space we would expect the test error curve to be $\frac{2\pi}{\Lambda}$ -periodic because of the rotational equivariance. The deviations from this behavior can be attributed to the sampling effects of using digitized images. As to be expected for $\Lambda = 16$ orientations, the accuracy is not influenced by augmentation with $\frac{\pi}{2}$ -rotations since the additional samples lie on the group orbit on which the network is invariant. In contrast to conventional CNNs, SFCNNs do not show an increased error for small angles in a neighborhood of zero when using augmentation. This indicates that the cost of learning rotated versions of each digit is negligible thanks to the approximate rotation equivariance. An augmentation by rotations which are multiples of $\frac{\pi}{4}$ or by continuous rotations give




Figure 5: Experimental results on the ISBI 2012 challenge. The shown patches are cropped from slice 30 of the training data set which we used for validation. *Left*: Raw EM image. *Mid-left*: Binary membrane ground truth segmentation. *Mid-right*: Probability map predicted by the proposed network. *Right*: Top 6 of more than 100 entries of the [leaderboard](#), accessed on November 13, 2017. Higher values mean better accuracy.

Method	Test Error (%)
Ours – CoeffInit, train time augmentation	0.714 ± 0.022
Ours – CoeffInit	0.880 ± 0.029
Ours – HeInit	0.957 ± 0.025
Marcos et al. [23] – test time augmentation	1.01
Marcos et al. [23]	1.09
Laptev et al. [21]	1.2
Worrall et al. [24]	1.69
Cohen and Welling [2] - G-CNN	2.28 ± 0.0004
Schmidt and Roth [29]	4.0
Sohn and Lee [11]	4.2
Cohen and Welling [2] - conventional CNN	5.03 ± 0.0020
Larochelle et al. [30]	10.4 ± 0.27

Table 1: Test errors on the rotated MNIST dataset. We distinguish He initialization (HeInit) from the proposed initialization scheme (CoeffInit).

very similar results. Both seem to act as a regularization preventing the filters to overfit on the pixel grid.

We conclude that SFCNNs outperform the rotational generalization of CNNs for all levels of augmentation.

Benchmarking: Based on the insights from the above experiments we fix the number of sampled orientations to $\Lambda = 16$ and tune the network further to the slightly larger architecture given in Table 2 in the supplementary material. The results are reported in Table 1. Using the SFCNN with He’s weight initialization and no data augmentation, we obtain a test error of 0.957% which already exceeds the previous state-of-the-art. The proposed initialization scheme, adapted to filter coefficients, significantly improves the test error to 0.880%. When additionally augmenting the dataset with continuous rotations during training time the error decreases further to 0.714%. To summarize, our approach reduces the best previously published error by a factor of 29%.

5.2. ISBI 2012 2D EM segmentation challenge

In a second experiment we evaluate the performance of our model on the ISBI 2012 electron microscopy segmentation challenge [5]. The goal of the challenge is to predict the locations of the cell boundaries in the *Drosophila* ventral nerve cord from EM images which is a key step for investigating the connectome of the brain. The dataset consists of 30 train and test slices of size 512×512 px with a binary segmentation ground truth provided for the training set. Figure 5 shows an exemplary raw EM image with the corresponding ground truth segmentation mask and our network’s prediction. An important property of the dataset is

that the images have no preferred orientation which makes it suitable for evaluating rotation-equivariant networks.

We build on an established pipeline introduced in [27] where a crucial step is the boundary prediction via a conventional CNN. In the present experiment, we replaced their network by a SFCNN with a U-net design [31]. The network architecture is visualized in Figure 1 in the supplementary material. As loss function we chose a pixel wise binary cross entropy loss. The dataset was augmented by random elastic deformations, flips and rotations by multiples of $\frac{\pi}{2}$ during train time. In the experiment on rotational generalization we found that augmenting samples by transformations in a subgroup under which the network is equivariant does not have any effect. We therefore sampled $\Lambda = 17$ orientations which is mutually prime with the 4 augmented orientations. This way the augmented images do not fall into a subgroup w.r.t. which the network is invariant.

Segmentation predictions are evaluated by the challenge hosters and ranked w.r.t. the foreground-restricted Rand score V^{Rand} and the information score V^{Info} ; for an explanation of these metrics see [5]. The current [leaderboard](#) in Figure 5 (right) shows that our approach yields top-tier results. In particular, it improves upon the results of [27].

6. Conclusion

We have developed a rotation-equivariant CNN whose filters are learned such that they are steerable. Layer-wise equivariance is obtained by using group convolutions. He’s weight initialization scheme is extended to general filter bases which empirically leads to an increased accuracy. Our network allows sampling an arbitrary number of filter orientations which improves the performance until a saturation is reached. We confirmed experimentally that SFCNNs generalize learned patterns over orientations and therefore achieve a lower sampling complexity than CNNs in rotation-equivariant recognition tasks. The proposed SFCNNs achieve state-of-the-art results on rotated MNIST and the ISBI 2012 2D EM segmentation challenge.

Acknowledgement. We would like to thank T. Beier, C. Pape, N. Rahaman and I. Arganda-Carreras for their technical support and U. Köthe and T. Cohen for valuable discussions. This work was partially supported by the German Research Foundation (DFG grant STO1126/2-1).

References

- [1] M. Zeiler and R. Fergus, “Visualizing and understanding convolutional networks,” in *European Conference on Computer Vision (ECCV)*. Springer, 2014, pp. 818–833. 1
- [2] T. Cohen and M. Welling, “Group equivariant convolutional networks,” in *International Conference on Machine Learning (ICML)*, 2016. 1, 4, 6, 7, 8
- [3] X. Glorot and Y. Bengio, “Understanding the difficulty of training deep feedforward neural networks,” in *International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2010. 1, 5
- [4] K. He, X. Zhang, S. Ren, and J. Sun, “Delving deep into rectifiers: Surpassing human-level performance on imagenet classification,” in *IEEE International Conference on Computer Vision (ICCV)*, 2015, pp. 1026–1034. 1, 5, 6
- [5] I. Arganda-Carreras, S. C. Turaga, D. R. Berger, D. Cireşan, A. Giusti, L. M. Gambardella, J. Schmidhuber, D. Laptev, S. Dwivedi, J. M. Buhmann *et al.*, “Crowdsourcing the creation of image segmentation algorithms for connectomics,” *Frontiers in neuroanatomy*, vol. 9, 2015. 1, 8
- [6] Y. Hel-Or and P. C. Teo, “Canonical decomposition of steerable functions,” *Journal of Mathematical Imaging and Vision*, vol. 9, no. 1, pp. 83–95, 1998. 3
- [7] W. Freeman and E. Adelson, “The design and use of steerable filters,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 13, no. 9, pp. 891–906, 1991. 3
- [8] Y.-N. Hsu and H. Arsenault, “Optical pattern recognition using circular harmonic expansion,” *Applied Optics*, vol. 21, no. 22, pp. 4016–4019, 1982. 3
- [9] J. Rosen and J. Shamir, “Circular harmonic phase filters for efficient rotation-invariant pattern recognition,” *Applied Optics*, vol. 27, no. 14, pp. 2895–2899, 1988. 3
- [10] A. Krizhevsky, I. Sutskever, and G. Hinton, “Imagenet classification with deep convolutional neural networks,” in *Advances in Neural Information Processing Systems (NIPS)*, 2012, pp. 1097–1105. 6
- [11] K. Sohn and H. Lee, “Learning invariant representations with local transformations,” in *International Conference on Machine Learning (ICML)*, 2012, pp. 1311–1318. 6, 8
- [12] A. Kanazawa, A. Sharma, and D. Jacobs, “Locally scale-invariant convolutional neural networks,” *Preprint arXiv:1412.5104*, 2014. 6
- [13] C. Zhang, S. Voinea, G. Evangelopoulos, L. Rosasco, and T. Poggio, “Discriminative template learning in group-convolutional networks for invariant speech representations,” in *Annual Conference of the International Speech Communication Association*, 2015. 6
- [14] R. Gens and P. Domingos, “Deep symmetry networks,” in *Advances in Neural Information Processing Systems (NIPS)*. Curran Associates, Inc., 2014, pp. 2537–2545. 6
- [15] J. F. Henriques and A. Vedaldi, “Warped convolutions: Efficient invariance to spatial transformations,” in *Proceedings of the 34th International Conference on Machine Learning*, vol. 70, 2017, pp. 1461–1469. 6
- [16] M. Jaderberg, K. Simonyan, A. Zisserman, and K. Kavukcuoglu, “Spatial transformer networks,” in *Advances in Neural Information Processing Systems (NIPS)*, 2015, pp. 2017–2025. 6
- [17] C.-H. Lin and S. Lucey, “Inverse compositional spatial transformer networks,” *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. 6
- [18] J.-H. Jacobsen, B. de Brabandere, and A. W. Smeulders, “Dynamic steerable blocks in deep residual networks,” *arXiv preprint arXiv:1706.00598*, 2017. 6
- [19] S. Dieleman, J. De Fauw, and K. Kavukcuoglu, “Exploiting cyclic symmetry in convolutional neural networks,” in *International Conference on Machine Learning (ICML)*, 2016. 6, 7
- [20] T. Cohen and M. Welling, “Steerable CNNs,” in *International Conference on Learning Representations (ICLR)*, 2017. 6, 7
- [21] D. Laptev, N. Savinov, J. Buhmann, and M. Pollefeys, “TI-POOLING: transformation-invariant pooling for feature learning in Convolutional Neural Networks,” in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016, pp. 289–297. 6, 8
- [22] D. Marcos, M. Volpi, and D. Tuia, “Learning rotation invariant convolutional filters for texture classification,” in *International Conference on Pattern Recognition (ICPR)*, 2016. 6
- [23] D. Marcos, M. Volpi, N. Komodakis, and D. Tuia, “Rotation equivariant vector field networks,” *arXiv:1612.09346*, 2016. 6, 8
- [24] D. Worrall, S. Garbin, D. Turmukhambetov, and G. Brostow, “Harmonic networks: Deep translation and rotation equivariance,” *Preprint arXiv:1612.04642*, 2016. 6, 8
- [25] L. Sifre and S. Mallat, “Rotation, scaling and deformation invariant scattering for texture discrimination,” in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2013, pp. 1233–1240. 6
- [26] T. M. Quan, D. G. Hilderbrand, and W.-K. Jeong, “Fusionnet: A deep fully residual convolutional neural network for image segmentation in connectomics,” *arXiv preprint arXiv:1612.05360*, 2016. 8

- [27] T. Beier, C. Pape, N. Rahaman, T. Prange, S. Berg, D. D. Bock, A. Cardona, G. W. Knott, S. M. Plaza, L. K. Scheffer, U. Koethe, A. Kreshuk, and F. A. Hamprecht, “Multi-cut brings automated neurite segmentation closer to human performance,” *Nature Methods*, vol. 14, no. 2, pp. 101–102, 2017. [8](#)
- [28] M. Drozdal, E. Vorontsov, G. Chartrand, S. Kadoury, and C. Pal, “The importance of skip connections in biomedical image segmentation,” in *International Workshop on Large-Scale Annotation of Biomedical Data and Expert Label Synthesis*. Springer, 2016, pp. 179–187. [8](#)
- [29] U. Schmidt and S. Roth, “Learning rotation-aware features: From invariant priors to equivariant descriptors,” in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2012, pp. 2050–2057. [8](#)
- [30] H. Larochelle, D. Erhan, A. Courville, J. Bergstra, and Y. Bengio, “An empirical evaluation of deep architectures on problems with many factors of variation,” in *International Conference on Machine Learning (ICML)*, 2007, pp. 473–480. [8](#)
- [31] O. Ronneberger, P. Fischer, and T. Brox, “U-net: Convolutional networks for biomedical image segmentation,” in *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer, 2015, pp. 234–241. [8](#)
- [32] S. Ioffe and C. Szegedy, “Batch normalization: Accelerating deep network training by reducing internal covariate shift,” in *International Conference on Machine Learning (ICML)*, 2015, pp. 448–456.
- [33] D. Kingma and J. Ba, “Adam: A method for stochastic optimization,” in *International Conference on Learning Representations (ICLR)*, 2015.
- [34] N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov, “Dropout: a simple way to prevent neural networks from overfitting,” *Journal of Machine Learning Research*, vol. 15, no. 1, pp. 1929–1958, 2014.