
Deep Continuous Networks

Nergis Tomen¹ Silvia L. Pintea¹ Jan C. van Gemert¹

Abstract

CNNs and computational models of biological vision share some fundamental principles, which opened new avenues of research. However, fruitful cross-field research is hampered by conventional CNN architectures being based on spatially and depthwise discrete representations, which cannot accommodate certain aspects of biological complexity such as continuously varying receptive field sizes and dynamics of neuronal responses. Here we propose deep continuous networks (DCNs), which combine spatially continuous filters, with the continuous depth framework of neural ODEs. This allows us to learn the spatial support of the filters during training, as well as model the continuous evolution of feature maps, linking DCNs closely to biological models. We show that DCNs are versatile and highly applicable to standard image classification and reconstruction problems, where they improve parameter and data efficiency, and allow for meta-parametrization. We illustrate the biological plausibility of the scale distributions learned by DCNs and explore their performance in a neuroscientifically inspired pattern completion task. Finally, we investigate an efficient implementation of DCNs by changing input contrast.

1. Introduction

Computational neuroscience and computer vision have a long and mutually beneficial history of cross-pollination of ideas (Cox & Dean, 2014; Sejnowski, 2020). The current state-of-the-art in computer vision relies heavily on convolutional neural networks (CNNs), from which multiple analogies can be drawn to biological circuits (Kietzmann et al., 2018). Thus, based on recent developments in deep learning, there has been a growing trend to relate CNNs to biological

circuits (Schrimpf et al., 2018), and employ CNNs as models of biological vision (Zhuang et al., 2020a). Specifically, recent advances in CNNs have enabled researchers to learn more accurate models of the response properties of neurons in the visual cortex (Klindt et al., 2017; Cadena et al., 2019; Ecker et al., 2019), as well as to test decades old hypotheses from neuroscience in the domain of computer vision (Lindsey et al., 2019). Hence, links between CNNs and biological models from neuroscience is fruitful for both research fields.

Contrary to many biological models, feed-forward CNNs typically use spatio-temporally discrete representations: CNNs employ spatially discretized, pixel-based kernels, and input is processed through a depthwise-discrete pipeline made up of successive convolutional layers. To clarify, within our framework we consider CNN depth to be analogous to time, similar to input-processing time-course in biological models. Unlike CNNs, large-scale, neuroscientific neural network models of the visual system often adopt continuous, closed-form expressions to describe spatio-temporal receptive fields, as well as the interaction strength between populations of neurons (Dayan & Abbott, 2001). Among others, such descriptions serve to limit the scope and parameter space of a model, by utilizing prior information regarding receptive field shapes (Jones & Palmer, 1987) and principles of perceptual grouping (Li, 1998). In addition, the choice of continuous—and often analytic—functions help retain some analytical tractability in complex models involving a large number of coupled populations. Our approach draws inspiration from such computational models to propose continuous CNNs.

In this work we aim for a biologically more plausible CNN model: We bring together (a) spatially continuous receptive fields, where both the shape and the scale of the filters are trainable in the continuous domain, and (b) depthwise continuous representations capable of modeling the continuous evolution of neuronal responses in feed-forward CNNs. Continuous receptive fields provide a link between modern CNNs and large-scale rate-based models of the visual system (Ernst et al., 2001). In addition, recent influential work in deep learning has introduced neural ordinary differential equations (ODEs) (Lu et al., 2018; Chen et al., 2018; Ruthotto & Haber, 2019) which propose a continuous depth (or time) interpretation of CNNs, while having spa-

¹Computer Vision Lab, Delft University of Technology, Delft, Netherlands. Correspondence to: Nergis Tomen <n.tomen@tudelft.nl>.

tially discrete filters. Such continuous depth models both offer end-to-end training capabilities with backpropagation which are highly applicable to computer vision problems (e.g. by way of adopting ResNet blocks (He et al., 2016)), as well as help bridge the gap to computational biology where networks are often modelled as dynamical systems which evolve according to differential equations. Building on this, we introduce deep continuous networks (DCNs), which are spatially and depthwise continuous in that the neurons have spatially well-defined receptive fields based on scale-spaces and Gaussian derivatives (Florack et al., 1996) and their activations evolve according to equations of motion comprising convolutional layers. Thus, we combine depthwise and spatial continuity by employing neural ODEs in a network which learns linear weights for a set of analytic basis functions (as opposed to pixel-based weights), which can also intuitively be parametrized as a function of network depth (or time).

Our main contributions are: (i) We provide a theoretical formulation of deep networks with spatially and depthwise continuous representations, building on Gaussian basis functions and neural ODEs; (ii) We demonstrate the applicability of DCN models, namely, that they exhibit a reduction in parameters, improve data efficiency and can be used to parametrize convolutional filters as a function of depth in a straightforward fashion, while achieving performance comparable with or better than ResNet and ODE-Net baselines; (iii) We show that filter scales learned by DCNs are consistent with biological observations and we propose that the combination of our design choices for spatial and depthwise continuity may be helpful in studying the emergence of biological receptive field properties as well as high-level phenomena such as pattern completion; (iv) We explore, for the first time, contrast sensitivity of neural ODEs and suggest that the continuous representations learned by DCNs may be leveraged for computational savings.

We believe DCNs can bring together two communities as they both provide a test bed for hypotheses and predictions pertaining to biological systems, and push the boundaries of biologically inspired computer vision.

2. Deep Continuous Networks

2.1. Neuroscientific motivation

There is little doubt that modern deep learning frameworks will be conducive to effective and insightful collaborations between neuroscience and machine learning (Richards et al., 2019). In particular in vision research, CNNs are becoming increasingly popular for modelling early visual areas (Batty et al., 2017; Ecker et al., 2019; Lindsey et al., 2019). Here we propose the DCN model which can facilitate such investigations by linking the end-to-end trainable but discrete

CNN architectures with the spatio-temporally continuous models of biological vision. Our approach makes it possible to optimize the spatial extent (kernel size) of the filters during training, as well as explicitly model the dynamics of the neuronal responses to input images.

Structured receptive fields. Classical receptive fields (RFs) of cortical neurons display complex response properties with a wide array of selectivity structures already at early visual areas (Van den Bergh et al., 2010). Such response properties may also vary greatly based on multiple factors. For example the RF size (spatial extent) is known to depend on eccentricity (Harvey & Dumoulin, 2011), visual area (Smith et al., 2001) and even cortical layer (Bauer et al., 1999). Similarly, studies have shown that spatial frequency selectivity and receptive field size may co-vary with input contrast (Sceniak et al., 2002).

Based on these observations, we aim to build a model which can accommodate the biological realism better than conventional CNNs, by explicitly modelling the RF size as a trainable parameter. To that end, we adopt a Gaussian scale-space representation for the convolutional filters, which we call structured receptive fields (SRFs) (Jacobsen et al., 2016). Previously, Gaussian scale-spaces have been proposed as a plausible model of biological receptive fields and feature extraction in low-level vision (Florack et al., 1992; Lindeberg, 1993; Lindeberg & Florack, 1994). Here, we are inspired by computational models which investigate the origin of response properties in the visual system, by employing RFs and recurrent interaction functions which scale as a difference of Gaussians (Somers et al., 1995; Ernst et al., 2001).

Neural ODEs. Studies have shown that both the contrast (Albrecht et al., 2002) and spatial frequency (Frazor et al., 2004) response functions of cortical neurons display characteristic temporal profiles. However, temporal dynamics are not incorporated into standard feed-forward CNN models. In addition, it has been suggested that lateral interactions play an important role in the generation of complex and selective neuronal responses (Angelucci & Bressloff, 2006). Such activity dynamics are often computationally modeled using recurrently coupled neuronal populations whose activations evolve according to coupled differential equations (Ben-Yishai et al., 1995; Ernst et al., 2001).

To describe the continuous evolution of feature maps consistent with biological models, we adopt the framework of neural ODEs (Chen et al., 2018). Neural ODE interpretation of ResNet models presents an opportunity to explicitly model the dynamics of feature extractors in feed-forward CNNs. Under certain assumptions, neural ODEs can be interpreted as biologically plausible recurrent interactions (Liao & Poggio, 2016; Rousseau et al., 2019), where the depth dimension represents time. Unlike neural ODEs with pixel-based

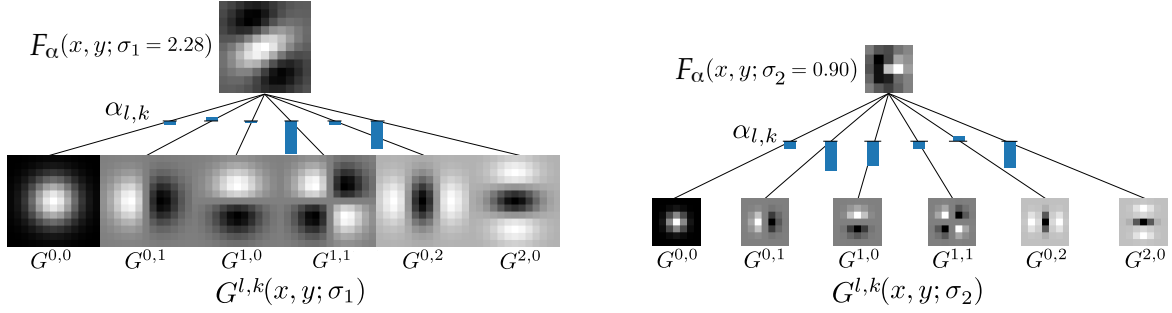


Figure 1. SRF filters based on N-jet filter approximation. Convolutional filters are defined as the weighted sum of Gaussian derivative basis functions up to order 2, with corresponding scales $\sigma_1 = 2.28$ (left) and $\sigma_2 = 0.90$ (right). Our DCN models learn both the coefficients α , and the scale σ end-to-end during training.

convolutional filters, DCNs with structured filters (SRFs) also provide an intuitive way to parametrize the evolution of the kernels as a function of depth.

Deep Continuous Networks. DCNs presented here combine structured receptive fields with neural ODEs. We view spatio-temporally continuous representations in end-to-end trainable networks as a link between modern CNN architectures and computational models of biological vision. Specifically, we are inspired by large-scale models of population activity. In contrast, networks modelling biological phenomena at smaller spatio-temporal scales may require discrete descriptions of biological neurons, such as spatially-discrete photoreceptors or temporally-discrete spiking dynamics. However, continuous rate-based population models provide reasonably good explanations of phenomena observed at the network or systems level (Ben-Yishai et al., 1995; Dayan & Abbott, 2001), which we believe align well with CNNs trained on high-level computer vision tasks.

Taken together, DCNs provide a fully trainable analog to biological models with continuous receptive fields and continuously evolving state variables, while preserving the modularity of the visual hierarchy by stacking spatio-temporally continuous blocks in a feed-forward stream (Fig. 2).

2.2. Structured receptive fields

We use the multiscale local N-jet formulation (Florack et al., 1996) to define the filters in convolutional layers. Structured receptive fields (SRFs) based on the Gaussian N-jet basis functions are highly applicable to CNNs, as they represent a Taylor expansion of the input image or feature maps in a local neighbourhood in space and scale, and can be used to approximate pixel-based filters (Appendix A.1). This means that each filter $F(x, y; \sigma)$ in the network is a weighted sum of N basis functions, which are partial derivatives of the isotropic two-dimensional Gaussian function $G(x, y; \sigma) = \frac{1}{2\pi\sigma^2} e^{-\frac{(x^2+y^2)}{2\sigma^2}}$. The scale, or the spatial extent, of the filter is explicitly modelled in the σ parameter

of the Gaussian, which also indirectly determines the spatial frequency response of the SRF (Fig. 1). Note that the Gaussian SRF formulation allows for learning filters with different aspect ratios, however, in this work we only consider isotropic basis functions with $\sigma = \sigma_x = \sigma_y$.

The N-jet formulation of an SRF filter $F(x, y)$ is given by:

$$\begin{aligned} F_\alpha(x, y; \sigma) &= \sum_{\substack{l+k \leq N \\ 0 \leq l, 0 \leq k}} \alpha_{l,k} G^{l,k}(x, y; \sigma) \\ &= \sum_{\substack{l+k \leq N \\ 0 \leq l, 0 \leq k}} \alpha_{l,k} \frac{\partial^{l+k}}{\partial x^l \partial y^k} G(x, y; \sigma), \quad (1) \end{aligned}$$

where $G^{l,k}(x, y; \sigma)$ are the partial derivatives of the Gaussian $G(x, y; \sigma)$ with respect to x and y , N is the degree of the Taylor polynomial which determines the basis order, and α encodes the expansion coefficients.

N-jet SRFs have favourable properties over pixel-based filters. SRF filters are steerable by the coefficients α and the basis functions are spatially separable. Likewise, due to their spatially continuous description, the filters can be trivially scaled, or rotated, without interpolation. In addition, SRFs can provide parameter efficiency when filters are constructed using a small number of basis functions. In this work we opt for basis order 2 (basis function up to the second order derivative), which yields relatively smooth filters. However, the generalized SRF framework allows for learning more irregular RF shapes by increasing the number of basis functions.

Fig. 1 shows the N-jet approximation of two filters in different scales σ_1 and σ_2 . We note that both the coefficients α and the scale σ are learnable filter parameters. Instead of fixing the scale σ a priori and optimizing for α as in Jacobsen et al. (2016) and Sosnovik et al. (2020), we integrate both these parameters in the network optimization, thus learning not only the shape but also the spatial support of the filters.

2.3. Neural ODEs

We model the continuous evolution of feature maps as a function of depth t within an ‘ODE block’. Formally, an ODE block contains a stack of M convolutional layers, each with its own convolutional filters \mathbf{w}_m with $m = 1 \dots M$, followed by normalization $G_{\text{norm}}(\cdot)$ and non-linear activation $\text{CELU}(\cdot)$ functions. Following the notations of [Chen et al. \(2018\)](#) and [Ruthotto & Haber \(2019\)](#), we define the equations of motion for the feature states $\mathbf{h} \in \mathbb{R}^n$ as:

$$\begin{aligned} \frac{d\mathbf{h}(t)}{dt} &= f(\mathbf{h}(t), t, \mathbf{w}_m, \mathbf{d}_m) \\ &= G_{\text{norm}} [\mathbf{K}_2(\mathbf{w}_2)g(\mathbf{K}_1(\mathbf{w}_1)g(\mathbf{h}) + \mathbf{d}_1t) + \mathbf{d}_2t] \end{aligned} \quad (2)$$

where $g(\mathbf{x}) = \text{CELU}(G_{\text{norm}}(\mathbf{x}))$, the linear operators $\mathbf{K}_m(\cdot) \in \mathbb{R}^{n \times n}$ denote the convolution operators parametrized by \mathbf{w}_m . The filters $\mathbf{w}_m(\theta)$, $\mathbf{d}_m(\theta)$ are functions of learnable parameters θ . In conventional CNNs \mathbf{w}_m are typically 3×3 kernels where the learnable parameters correspond to pixel weights. In the DCN model we define the filters \mathbf{w}_m using the Gaussian SRF, thus, the learnable parameters are basis coefficients α and scale σ , and the kernel size is not fixed but scales with σ and is learned (see Section 2.4).

The CNN convolution operator, $\mathbf{K}_m(\mathbf{w}_m)$ with 2 input and 2 output channels can be written as

$$\mathbf{K}_m(\mathbf{w}_m) = \begin{pmatrix} \mathbf{K}_m^{1,1}(\mathbf{w}_m^{1,1}) & \mathbf{K}_m^{1,2}(\mathbf{w}_m^{1,2}) \\ \mathbf{K}_m^{2,1}(\mathbf{w}_m^{2,1}) & \mathbf{K}_m^{2,2}(\mathbf{w}_m^{2,2}) \end{pmatrix} \quad (3)$$

with \mathbf{w}_m^{ji} the convolutional kernels for input channel i and output channel j of the m -th convolution. The time-offset terms $\mathbf{d}_m t$ in Eq. 2 makes the ODE an explicit function of t , which separates the ODE block implementation from a simple convolutional block with weight sharing over depth.

In accordance with conventional ResNet blocks, we pick $M = 2$. Based on the implementation by [Chen et al. \(2018\)](#), G_{norm} is defined as group normalization ([Wu & He, 2018](#)). For generalized compatibility with ODEs and the adjoint method, we choose a non-linear activation function with a theoretically unique and bounded adjoint, namely continuously differentiable exponential linear units, or CELU ([Barron, 2017](#)). Similarly, we keep the linear dependence of the equations of motion on continuous network depth t . Finally, we adapt the GPU implementation of ODE solvers¹ to solve the equations of motion for a predefined time interval $t \in [0, T]$ using the adaptive step size DOPRI method.

Time vs. depth In the neural ODE definition ([Chen et al., 2018](#)), the discrete depth of feed-forward networks such as ResNets is reimaged as a continuous dimension denoted by time t , where the input image defines the initial

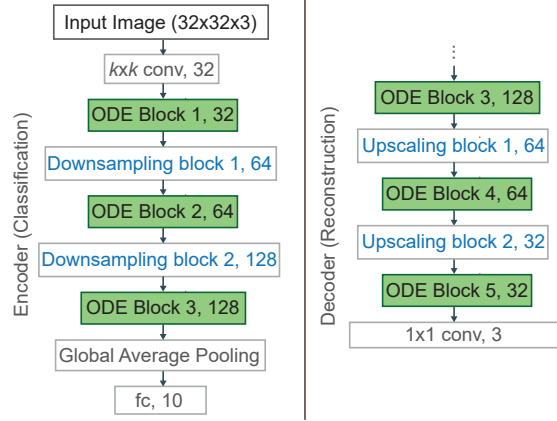


Figure 2. DCN model architecture with CIFAR-10 input images. Convolutional kernel size k is learned during training. The equations of motion (Eq. 2) are solved within ODE blocks.

conditions $\mathbf{h}(0)$. For the rest of this paper, we use the interpretation that the number of function evaluations performed by the numerical ODE solver is analogous to network depth. In this sense, continuous ‘depth’ or ‘time’ refers to the continuous variable t within the ODE blocks, while the full architecture is still modular and composed of multiple ODE blocks. It is also important to note that when we talk about the spatio-temporal dynamics of DCNs, we refer to the temporal evolution of the feature maps in the ODE blocks and not to input dynamics, as in a video. While DCNs are primarily feed-forward networks, the ODE definition makes it possible for DCNs to model time-varying neuronal activations via the ‘continuous depth’, even in response to static input images (see Section 4 for more detailed comparisons with recurrent neural networks).

2.4. Deep Continuous Networks with SRFs and ODEs

We formulate deep continuous networks (DCNs) by employing learnable, continuous SRF filter descriptions to define the weights \mathbf{w} in the evolution of a neural ODE. This means that for DCNs, each \mathbf{w}_m^{ji} in Eq. 2 is a discretization of the continuous SRF filter $F_{\alpha_m^{ji}}(x, y; \sigma_m)$ given in Eq. 1, sampled in $[-2\sigma_m, 2\sigma_m]$. α_m^{ji} and σ_m are trainable filter parameters, where $\sigma_m^{ji} = \sigma_m$ is shared between the filters in a convolutional layer m unless stated otherwise. All our code is available at².

Network architecture and training. We construct DCNs by stacking ODE blocks separated by downsampling blocks (Fig. 2). Each downsampling block is a sequence of normalization, nonlinear activation and strided convolution. We use a convolutional layer for increasing the channel dimensionality at the input level and employ global average pooling and a fully connected layer at the output level. We train

¹<https://github.com/rtqichen/torchdiffeq/>

²<https://github.com/ntomen/Deep-Continuous-Networks>

Table 1. CIFAR-10 validation accuracies of DCN models, averaged over 3 runs, compared to baseline models. ODE-Net and ResNet-blocks baselines are as introduced in *Chen et al. (2018). DCNs perform on par with spatially and/or temporally discrete baselines, despite having a lower number of trainable parameters.

Model	Continuity		Accuracy (%)	Parameters
	Spatial	Temporal		
ODE-Net *	x	✓	89.6 ± 0.3	560K
ResNet-blocks *	x	x	89.0 ± 0.2	555K
ResNet-SRF-blocks	✓	x	88.3 ± 0.03	426K
ResNet-SRF-full	✓	x	89.3 ± 0.4	323K
DCN-ODE	✓	✓	89.5 ± 0.2	429K
DCN-full	✓	✓	89.2 ± 0.3	326K
DCN σ^{ji}	✓	✓	89.7 ± 0.3	472K

our networks using cross-entropy loss and the CIFAR-10 dataset (Krizhevsky, 2009). (See Appendices A.2-A.3 for further details regarding training parameters.)

As a baseline without spatial continuity, we compare DCN performance to the ODE-Net introduced in Chen et al. (2018), where the convolutions within the ODE blocks are performed using discrete, pixel-based kernels, with 3×3 parameters. As a baseline without (depthwise) temporal continuity, we define the ‘ResNet-blocks’ model where the ODE blocks are replaced by generic, discrete ResNet blocks, comprising two convolutional layers and a skip connection, with comparable number of parameters to the ODE-Net. This is also a baseline model used in Chen et al. (2018). In the ResNet-SRF-blocks model, we provide the discrete-depth and continuous-space baseline by replacing the 3×3 filter definition of ResNet-blocks with SRF definitions.

We test two versions of DCNs and ResNet-SRF-blocks to quantify the viability of SRF filters outside of the ODE blocks. In DCN-ODE and ResNet-SRF-blocks we use the SRF filters only within the ODE (ResNet) blocks, and for the remaining layers we use discrete kernels with the same hyperparameters as the baselines. In the second version, DCN-full (ResNet-SRF-full), we use spatially continuous kernels everywhere, including the downsampling layers.

As an additional demonstration of the versatility of DCNs, we conduct an image reconstruction experiment on CIFAR-10. We use the feature maps generated by encoder networks (output of ODE Block 3 in Fig. 2), as input to a decoder network. The decoder networks are composed of 2 DCN-ODE, ODE-Net or ResNet blocks, separated by bilinear upscaling layers and 1×1 convolutions to reduce dimensionality.

Finally, we investigate the case where we drop scale sharing within a layer, and optimize the scale parameter σ_m^{ji} independently for each input channel i and output channel j , which we call DCN σ^{ji} .

Meta-parametrization. DCNs enable us to parametrize the trainable filter parameters α and σ as a function of depth t . This both enables the kernels to vary smoothly over depth, and lets us define temporal dynamics for the neuronal responses in our network. We test the viability of such models by introducing DCN variants where σ and/or α are defined using linear or quadratic functions of t and learnable parameters $a, b, c, a_s, b_s, a_\alpha$ and b_α (Table 4).

3. Experimental Analysis

3.1. Parameter reduction and data efficiency

Similar to biological models, where analytical receptive fields limit the scope of the model using prior information, we find that DCNs are more parameter efficient compared to baseline networks. Evaluated on CIFAR-10, DCNs perform on par with baselines, despite using SRFs of a small basis order 2, which means each filter shape is defined by only 6 free parameters as opposed to 9 for the conventional 3×3 kernels (Table 1). In addition, we find that parameter reduction via the use of SRFs with a small basis order also leads to data efficiency. When trained on a subset of CIFAR-10 images (small-data regime), DCNs outperform the discrete baseline networks (Table 2). We also find accuracy increase over the small-data performance reported by Arora et al. (2020) for the convolutional neural tangent kernel (CNTK) model (Table 2).

Moreover, we train encoder-decoder networks to reconstruct CIFAR-10 images using mean squared error (MSE) loss. We find that the DCN models outperform discrete baseline models on the validation set (Table 3), despite having a lower number of parameters as before. Additional details and example images are shown in Appendix A.4.

We find that meta-parametrized DCN variants match the classification performance of baselines and may outperform DCNs with static weights (Table 5). This is an interesting finding as we test only a few models, with little hyperparameter optimization, indicating that DCNs can potentially be used to parametrize the dependence of convolutional kernel weights on network depth, for further parameter reduction.

3.2. Link with biological models

Scale fitting. As an advantage over conventional CNNs, it is possible to directly investigate the optimal receptive field (RF) size in each DCN block after training, since DCNs fit the kernel scale σ explicitly. We observe an upward trend in the SRF scale σ with the depth of the convolutional layer within the network (Fig. 3a). While the RF size grows with depth also in conventional CNNs, it typically grows in a predetermined manner: for a cascade of convolutional layers the RF size is a linear function of depth given constant kernel size. Thus, the receptive field size at every CNN layer

Table 2. Validation accuracies for the DCN-ODE model and baselines trained on a subset of CIFAR-10 (small-data regime). First two rows show small-data baseline accuracies taken from [†]Arora et al. (2020). ResNet-blocks and ODE-Net models are implemented by us, as in Table 1. The DCN model outperforms spatially and/or temporally discrete baselines for medium training set size as parameter efficiency leads to data efficiency. All results are averaged over 3 runs.

Model	# images per class										
	2	4	8	16	32	52	64	103	128	512	1024
ResNet34 [†]	17.5±2.5	19.5±1.4	23.3±1.6	28.3±1.4	33.2±1.2	–	41.7±1.1	–	49.1±1.3	–	–
CNTK [†]	18.8±2.1	21.3±1.9	25.5±1.9	30.5±1.2	36.6±0.9	–	42.6±0.7	–	48.9±0.7	–	–
ResNet-blocks	16.7±0.8	19.6±1.0	22.0±1.3	28.1±1.7	35.4±0.9	39.8±0.6	41.6±1.5	49.0±0.2	50.9±0.6	70.4±1.2	76.8±0.7
ODE-Net	16.8±2.8	20.5±0.8	23.1±2.5	29.8±0.8	36.4±1.0	41.7±1.2	42.3±0.2	48.6±0.5	50.7±0.7	71.7±1.5	77.4±0.5
DCN-ODE	16.4±1.6	19.8±0.7	26.5±0.9	31.2±0.6	37.7±0.6	44.5±0.8	48.0±1.3	54.2±0.8	58.2±0.7	75.5±0.8	79.7±0.3

Table 3. DCNs achieve lower MSE loss in the reconstruction task than discrete baselines on the CIFAR-10 validation set, despite using a smaller number of parameters. See also Appendix A.4 for reconstructed image examples. All results are averaged over 3 runs.

Model	Reconstruction Loss (%)
ResNet-blocks	21.0 ± 0.4
ODE-Net	20.2 ± 1.3
DCN-ODE	17.1 ± 0.3

is fixed depending on the architecture and hyperparameters. This is a limitation of CNNs which the visual system does not necessarily have. DCNs, on the other hand, can *learn* RF sizes which grow non-linearly as a function of depth, which seems to be in line with the behaviour in downstream visual areas (Smith et al., 2001).

In addition, we plot the distribution of learned σ^{j_i} in different ODE blocks of the model DCN σ^{j_i} (Fig. 3b). Note that the scale parameter σ controls the bandwidth of the SRF filters and is thus related to their spatial frequency response. We find that the σ^{j_i} distributions after training are approximately log-normal and display a positive skew, which is consistent with the scale and spatial frequency tuning distributions in the primate visual system (Yu et al., 2010). We believe these results are promising for bridging the gap between deep learning and traditional models of biological systems.

Pattern completion. Established models from computational neuroscience, with continuous temporal dynamics and well-defined recurrent interaction structures, such as the Ermentrout-Cowan model (Bressloff et al., 2001), or neural field models (Amari, 1977), display interesting high-level phenomena such as spontaneous pattern formation and travelling waves (Coombes, 2005). Such models employ local, distance-dependent interactions, similar to the SRF-based ODE blocks in the DCN formulation. Based on

Table 4. Meta-parametrization of filter parameters σ and α as a function of depth t in different DCN variants.

Model	Parametrization
DCN $\sigma(t)$	$\sigma = 2^{at+b}$
DCN $\sigma(t^2)$	$\sigma = 2^{at^2+bt+c}$
DCN $\sigma(t), \alpha(t)$	$\sigma = 2^{a_s t + b_s}, \alpha = a_\alpha t + b_\alpha$

Table 5. CIFAR-10 validation accuracies averaged over 3 runs for DCN models with meta-parametrization.

Model	Accuracy (%)
DCN $\sigma(t)$	89.97 ± 0.30
DCN $\sigma(t^2)$	89.93 ± 0.28
DCN $\sigma(t)$ and $\alpha(t)$	89.88 ± 0.25

this resemblance, we explore whether DCNs may display similar *emergent* properties. Specifically, we hypothesize that DCNs can perform well in the case of locally missing information in images, through pattern completion at the feature map level.

We test this hypothesis on the DCN models trained on CIFAR-10 classification by masking $n \times n$ pixels of the validation images at test time. The masks have zero pixel values, and are placed at the center of the image. We find that when confronted with a small patch of missing information at test time, DCNs can generate feature maps similar to those obtained from intact images. Specifically, we observe that the total difference:

$$D(t) = \frac{1}{A} \sum |\mathbf{h}^{\text{im}}(t) - \mathbf{h}^{\text{im,masked}}(t)| \quad (4)$$

between the feature maps generated by an intact image $\mathbf{h}^{\text{im}}(t)$ and a masked image $\mathbf{h}^{\text{im,masked}}(t)$, normalized by the amplitude of the intact image A , is reduced within an ODE block (Fig. 4). In terms of the overall classification

performance with images masked at test time, we find that DCNs are marginally more robust against zero-masking than baselines (Fig. 3c).

3.3. Contrast robustness and computational efficiency

The selectivity of neuronal responses is invariant to contrast in mammalian vision (Sclar & Freeman, 1982; Skottun et al., 1987). However, we observe that DCN and ODE-Net models are sensitive to changes in input contrast. This is not unexpected since ODE blocks compute the solution to the initial value problem posed by the equations of motion and the input $\mathbf{h}(0)$. To quantify this sensitivity we vary the contrast c of the input images at test time, where for each image H in the CIFAR-10 validation set we define the network input as $\widehat{H} = cH$. When naively changing the input contrast c this way, we find that the validation accuracy decays rapidly for both models (solid lines in Fig. 5, top).

Empirically, we notice that, with the appropriate choice of normalization functions, the input contrast c has a direct effect on the time scales of the solution $\mathbf{h}(t)$. This means that under different contrast values c , the feature map trajectories within an ODE block may converge faster, and a more efficient DCN implementation might be possible. Based on this observation, we heuristically test whether scaling the integration time interval T (used during training) of ODE block 1 by the input contrast at test time, as $\widehat{T} = cT$, can improve contrast robustness at test time. We find that with the scaled integration interval, DCN validation accuracy is relatively robust against changes in contrast c , compared to naive baselines and ODE-Net, until $c \ll 1$ when time scales become too fast and the ODE solver becomes unstable for all models (dashed lines in Fig. 5, top).

Interestingly, we observe a reduction in the number of function evaluations (NFEs) in ODE block 1 for $c < 1$ (Fig. 5, middle). Furthermore, we show that as long as the error tolerance of the ODE blocks are not decreased, this effect can be exploited by scaling the input feature maps of all ODE blocks by c for significant computational savings. We find that decreasing c leads to considerable efficiency improvements, where total NFEs can be reduced from 102 to 60, (for $c = 1$ and 0.06), with less than 0.5% loss in accuracy (Fig. 5, bottom).

4. Related Work

Our proposed DCN networks extend prior work on continuous filters and continuous depth neural ODEs.

Spatially continuous filter representations. Structured filters have been traditionally used in computer vision for extracting image structure at multiple scales. N-jet filter basis is first introduced by Florack et al. (1996) based on previous work on Gaussian scale-spaces (Florack et al., 1992;

Lindeberg, 2013). We use the N-jet basis, which enables a spatially continuous representation, with a learnable scale parameter σ , to approximate convolutional filters.

Similar to the N-jet basis, a set of oriented multi-scale wavelets, called a steerable pyramid, is proposed by Simoncelli et al. (1992) and complex wavelets have been used by Mallat (2012) and Bruna & Mallat (2013) as part of scattering transforms. CNN filters based on linear combinations of Gabor wavelets are adopted by Luan et al. (2018), while Worrall et al. (2017) propose circular harmonics, as spatially continuous filter representations.

Similar to our approach, Shelhamer et al. (2019) combine free-form filters with Gaussian kernels, thus learning the filter resolution. Likewise, Xiong et al. (2020) learn filter sizes using Gaussian kernels optimized using variational inference. Finally, Loog & Lauze (2017) integrate continuous scale-selectivity through a regularization hyper-parameter. Here, we use the N-jet framework based on Gaussian derivatives as in Jacobsen et al. (2016) and Pintea et al. (2021), however our main motivation is retaining compatibility with biological models. Also, unlike Jacobsen et al. (2016) we learn the scale parameter σ during training.

Continuous depth representations in deep networks.

Along with work by Lu et al. (2018) and Ruthotto & Haber (2019), networks continuous in the depth (or time) dimension have been proposed by Chen et al. (2018) under the name neural ordinary differential equations (ODEs). They propose ODE-Nets based on the ResNet formulation (He et al., 2016) for classification tasks, which we use as a baseline. In this work we focus mainly on image classification, however, there is extensive ongoing work on generative models and normalizing flows using the neural ODE continuous depth interpretation (Salman et al., 2018; Grathwohl et al., 2019). We note that DCNs can be readily incorporated into continuous flow models, as well as other spatio-temporally continuous CNN interpretations based on partial differential equations (Ruthotto & Haber, 2019).

Even though the adjoint method described in Chen et al. (2018) offers considerable computational savings, especially in terms of memory, recent work has improved upon it both in terms of stability, computational efficiency and performance (Dupont et al., 2019; Finlay et al., 2020; Zhuang et al., 2020b). Likewise, the contrast robust formulation of DCNs, as well as the synergy between the $\mathcal{O}(1)$ memory complexity of the adjoint method and spatially separable SRF filters (the implementations of which may otherwise inflate the memory cost) provide potential computational benefits over conventional CNNs where the number of function evaluations is fixed.

Other studies have suggested that, similar to our DCN variants where the filter definitions are independent of depth,

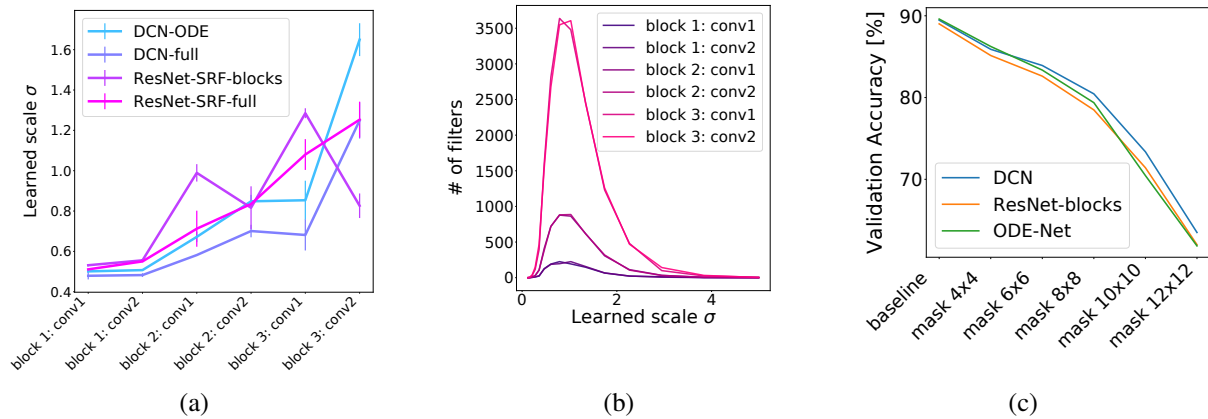


Figure 3. (a) Learned σ values increase with depth within the network. (b) σ^{ji} distributions within the ODE blocks display a positive skew in line with biological observations. (c) CIFAR-10 validation accuracies on the pattern completion task with increasing mask size.

neural ODEs based on ResNet architectures with weight sharing can be interpreted as recurrent neural networks (Kim et al., 2016; Rousseau et al., 2019), which bridges the gap between deep learning, dynamical systems and the primate visual cortex (Liao & Poggio, 2016; Massaroli et al., 2020). Similar to these works, we illustrate the parallels between neural ODEs and the dynamical systems approach of the computational models of biological circuits. As a novel contribution, we extend neural ODEs to DCNs, where not only the depth of the network is continuous but also the shape and spatial resolution of the filters are end-to-end trainable.

CNNs and RNNs as models of biological networks.

There is extensive prior work on CNNs and recurrent neural networks (RNNs) for modeling biological computation. The visual cortex is highly recurrent (Dayan & Abbott, 2001; Liao & Poggio, 2016) which is thought to be responsible for complex neuronal dynamics (Ben-Yishai et al., 1995; Angelucci & Bressloff, 2006). Accordingly, computational models with lateral connections (Sompolinsky et al., 1988; Ernst et al., 2001) and more recently RNNs (Laje & Buonomano, 2013; Mante et al., 2013; Mastrogiuseppe & Ostojic, 2018) have been extensively used as models of biological neural computation. For example the first-order reduced and controlled error (FORCE) algorithm, have been used to reproduce the dynamics of different biological circuits (Sussillo & Abbott, 2009; Laje & Buonomano, 2013; Carnevale et al., 2015; Rajan et al., 2016; Enel et al., 2016). Similarly, optimization via gradient-based algorithms such as the Hessian-free method (HF) or stochastic gradient-descent (SGD) have been adopted to replicate experimental observations (Mante et al., 2013; Barak et al., 2013; Song et al., 2016). It has also been suggested to use spiking recurrent networks (Kim & Chow, 2018; Kim et al., 2019) and incorporate synaptic dynamics (Ba et al., 2016; Miconi et al., 2018) for improved physiological realism.

Bringing together the power of CNNs and neuroscience,

recurrent convolutional networks (RCNNs) have been proposed (Liang & Hu, 2015; Spoerer et al., 2017; Hu & Mihaïlas, 2018), which can emulate biological lateral connectivity structures and extra-classical receptive field effects. Similar to our work where depthwise-continuity mimics recurrent networks, it has been shown that adding recurrent layers to convolutional deep networks can facilitate pattern completion in a manner consistent with psychophysical and electrophysiological experiments (Tang et al., 2018). Furthermore, our DCN models have the potential to compress the depth of the network, by replacing multiple sequential layers with meta-parametrized ODE blocks, which are analogous to recurrent networks with continuously evolving filter parameters. In a similar line of work, it has previously been shown that shallow networks with recurrently connected layers can achieve high object recognition performance while retaining brain-like representations, and specifically reproducing the population dynamics in area IT of the visual system much more closely than feed-forward deep CNNs (Kar et al., 2019; Kubilius et al., 2019).

In contrast to standard RNNs, our model is based on the ResNet inspired model of neural ODEs, and in its current form (Eq. 2), does not accept time-variant input. In that sense, the spatio-temporal dynamics of DCNs refer to the dynamics of the feature maps, or neuronal responses, and not the input. Nevertheless, this gives DCNs the ability to model time-varying responses, even to static input images.

In addition, DCNs with weight sharing can be thought of as recurrent networks (Rousseau et al., 2019) and can be easily modified to process time-variant input (such as videos). However, in this paper we consider DCN models as an extension of conventional feed-forward CNNs, with extended temporal dynamics and continuous spatial representations, which are applicable to feed-forward models of the visual system similar to works by Schrimpf et al. (2018); Lindsey et al. (2019); Ecker et al. (2019); Zhuang et al. (2020a).

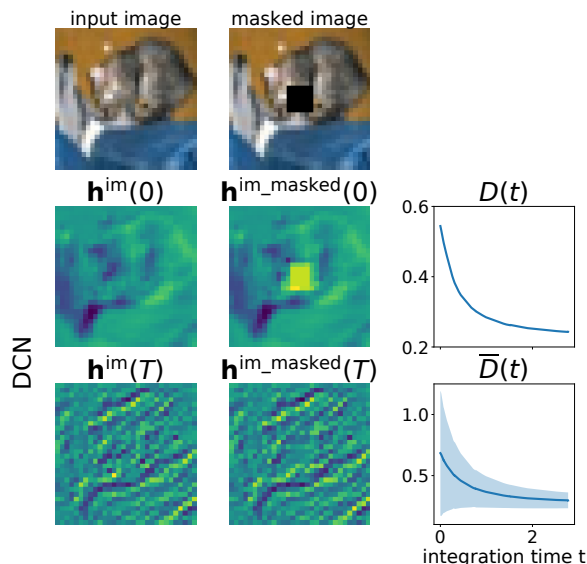


Figure 4. Pattern completion in the DCN feature maps during classification of masked images. Feature maps in a single channel of ODE block 1 are shown for an example image. We find that the difference $D(t)$ between the feature maps $\mathbf{h}^{\text{im}}(t)$ of an intact image and $\mathbf{h}^{\text{im_masked}}(t)$ of a masked image is reduced as $t \rightarrow T$. We also show the mean $\bar{D}(t)$ for 1000 validation images (bottom right), where the shaded area is the standard deviation over different images. Example feature maps from baseline models are provided in Appendix A.5.

5. Discussion

We introduce DCNs, CNN models which learn spatio-temporally continuous representations, consistent with biological models. We show that DCNs can match baseline performance in an image classification task and outperform baselines in the small-data regime and in a reconstruction task, while using a smaller number of parameters. Similarly, we propose different methods of meta-parametrization of the convolutional filter as a function of depth, which may not only be applicable to network compression, but also for modelling the temporal profiles of biological responses. As a further link with biological models, we have demonstrated that the learned filter scale distributions in DCNs are compatible with experimental observations. This makes the DCN models viable for future neuroscientific investigations regarding the emergence of RF sizes. In addition, we have presented the capability of DCNs to reduce errors in feature maps caused by masking. Finally, we have empirically shown an interesting interplay between the input contrast to ODE blocks and the time scales of the solutions, which can be capitalized on for computational savings.

However, one of the biggest limitations of DCN models is that they may become unstable during training. Combining neural ODEs with scale fitting may lead to explod-

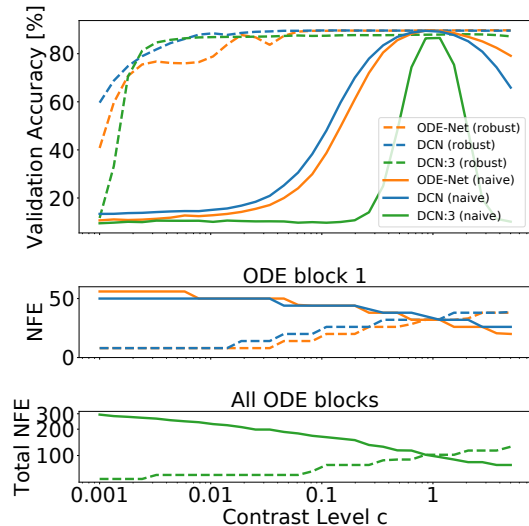


Figure 5. On the CIFAR-10 validation set, DCNs are more robust than baseline ODE-Nets to changes in input contrast c at test time (top). Interestingly, the number of function evaluations (NFEs) in the first ODE block (middle) or the whole DCN network (bottom) can be reduced considerably by modulating c .

ing filter sizes at large learning rates. Especially for meta-parametrization, it would be advisable to clip the integration time and filter parameters within a reasonable range.

Nevertheless, we believe there are exciting future research opportunities involving DCNs. Neural ODE formulations provide an interesting opportunity for establishing a theoretical understanding of deep networks based on dynamical systems. The interplay of input contrast and integration time is one such observation which requires further investigation. Similarly, our choice of filters based on well-behaved Gaussian derivatives allow for further analytical studies, unlike conventional CNNs.

Similarly, DCNs offer interesting possibilities for biological modelling. The inbuilt smooth evolution of filters in DCNs can be used, for example, to incorporate response dynamics such as synaptic depression or short-term potentiation (Ba et al., 2016; Micconi et al., 2018). Likewise, the equations of motion can be modified to reflect axonal delays or generate oscillations. Taken together, we believe by offering a link between dynamical systems, biological models and CNNs, DCNs display an interesting potential to bring together ideas from both fields.

Acknowledgements

Authors thank the reviewers for insightful comments, and Prof. Dr. Marco Loog for fruitful discussions. This publication is part of the project “Pixel-free deep learning” (TOP grant with project number 612.001.805) which is financed by the Dutch Research Council (NWO).

References

- Albrecht, D. G., Geisler, W. S., Frazor, R. A., and Crane, A. M. Visual cortex neurons of monkeys and cats: temporal dynamics of the contrast response function. *Journal of Neurophysiology*, 88(2):888–913, 2002.
- Amari, S. Dynamics of pattern formation in lateral-inhibition type neural fields. *Biological Cybernetics*, 27(2):77–87, 1977.
- Angelucci, A. and Bressloff, P. C. Contribution of feedforward, lateral and feedback connections to the classical receptive field center and extra-classical receptive field surround of primate v1 neurons. *Progress in Brain Research*, 154:93–120, 2006.
- Arora, S., Du, S. S., Li, Z., Salakhutdinov, R., Wang, R., and Yu, D. Harnessing the power of infinitely wide deep nets on small-data tasks. In *International Conference on Learning Representations (ICLR)*, 2020.
- Ba, J., Hinton, G. E., Mnih, V., Leibo, J. Z., and Ionescu, C. Using fast weights to attend to the recent past. In *Advances in Neural Information Processing Systems (NeurIPS)*, pp. 4331–4339, 2016.
- Barak, O., Sussillo, D., Romo, R., Tsodyks, M., and Abbott, L. From fixed points to chaos: three models of delayed discrimination. *Progress in Neurobiology*, 103:214–222, 2013.
- Barron, J. T. Continuously differentiable exponential linear units. *arXiv preprint arXiv:1704.07483*, 2017.
- Batty, E., Merel, J., Brackbill, N., Heitman, A., Sher, A., Litke, A., Chichilnisky, E., and Paninski, L. Multilayer recurrent network models of primate retinal ganglion cell responses. *ICLR*, 2017.
- Bauer, U., Scholz, M., Levitt, J. B., Obermayer, K., and Lund, J. S. A model for the depth-dependence of receptive field size and contrast sensitivity of cells in layer 4c of macaque striate cortex. *Vision research*, 39(3):613–629, 1999.
- Ben-Yishai, R., Bar-Or, R. L., and Sompolinsky, H. Theory of orientation tuning in visual cortex. *Proceedings of the National Academy of Sciences*, 92(9):3844–3848, 1995.
- Bressloff, P. C., Cowan, J. D., Golubitsky, M., Thomas, P. J., and Wiener, M. C. Geometric visual hallucinations, euclidean symmetry and the functional architecture of striate cortex. *Philosophical Transactions of the Royal Society of London. Series B: Biological Sciences*, 356(1407):299–330, 2001.
- Bruna, J. and Mallat, S. Invariant scattering convolution networks. *TPAMI*, 35(8):1872–1886, 2013.
- Cadena, S. A., Denfield, G. H., Walker, E. Y., Gatys, L. A., Tolias, A. S., Bethge, M., and Ecker, A. S. Deep convolutional models improve predictions of macaque V1 responses to natural images. *PLoS Computational Biology*, 15(4):e1006897, 2019.
- Carnevale, F., de Lafuente, V., Romo, R., Barak, O., and Parga, N. Dynamic control of response criterion in premotor cortex during perceptual detection under temporal uncertainty. *Neuron*, 86(4):1067–1077, 2015.
- Chen, R. T. Q., Rubanova, Y., Bettencourt, J., and Duvenaud, D. Neural ordinary differential equations. *NeurIPS*, 2018.
- Coombes, S. Waves, bumps, and patterns in neural field theories. *Biological Cybernetics*, 93(2):91–108, 2005.
- Cox, D. D. and Dean, T. Neural networks and neuroscience-inspired computer vision. *Current Biology*, 24(18):R921–R929, 2014.
- Dayan, P. and Abbott, L. F. *Theoretical neuroscience: computational and mathematical modeling of neural systems*. Massachusetts Institute of Technology Press, 1st edition, dec 2001. ISBN 0262041995.
- Dupont, E., Doucet, A., and Teh, Y. W. Augmented neural ODEs. In *NeurIPS*, pp. 3140–3150, 2019.
- Ecker, A. S., Sinz, F. H., Froudarakis, E., Fahey, P. G., Cadena, S. A., Walker, E. Y., Cobos, E., Reimer, J., Tolias, A. S., and Bethge, M. A rotation-equivariant convolutional neural network model of primary visual cortex. In *International Conference on Learning Representations (ICLR)*, 2019.
- Enel, P., Procyk, E., Quilodran, R., and Dominey, P. F. Reservoir computing properties of neural dynamics in prefrontal cortex. *PLoS Computational Biology*, 12(6):e1004967, 2016.
- Ernst, U., Pawelzik, K., Sahar-Pikielny, C., and Tsodyks, M. Intracortical origin of visual maps. *Nature Neuroscience*, 4(4):431–436, 2001.
- Finlay, C., Jacobsen, J.-H., Nurbekyan, L., and Oberman, A. M. How to train your neural ODE. *arXiv preprint arXiv:2002.02798*, 2020.
- Florack, L., Romeny, B. T. H., Viergever, M., and Koenderink, J. The gaussian scale-space paradigm and the multiscale local jet. *IJCV*, 18(1):61–75, 1996.
- Florack, L. M., ter Haar Romeny, B. M., Koenderink, J. J., and Viergever, M. A. Scale and the differential structure of images. *Image and Vision Computing*, 10(6):376–388, 1992.
- Frazor, R. A., Albrecht, D. G., Geisler, W. S., and Crane, A. M. Visual cortex neurons of monkeys and cats: temporal dynamics of the spatial frequency response function. *Journal of Neurophysiology*, 91(6):2607–2627, 2004.
- Grathwohl, W., Chen, R. T. Q., Bettencourt, J., Sutskever, I., and Duvenaud, D. Ffjord: Free-form continuous dynamics for scalable reversible generative models. *ICLR*, 2019.
- Harvey, B. M. and Dumoulin, S. O. The relationship between cortical magnification factor and population receptive field size in human visual cortex: constancies in cortical architecture. *Journal of Neuroscience*, 31(38):13604–13612, 2011.
- He, K., Zhang, X., Ren, S., and Sun, J. Deep residual learning for image recognition. In *CVPR*, pp. 770–778, 2016.
- Hu, B. and Mihalas, S. Convolutional neural networks with extra-classical receptive fields. *CoRR*, abs/1810.11594, 2018. URL <http://arxiv.org/abs/1810.11594>.
- Jacobsen, J.-H., van Gemert, J., Lou, Z., and Smeulders, A. W. Structured receptive fields in CNNs. In *CVPR*, pp. 2610–2619, 2016.
- Jones, J. P. and Palmer, L. A. An evaluation of the two-dimensional gabor filter model of simple receptive fields in cat striate cortex. *Journal of Neurophysiology*, 58(6):1233–1258, 1987.

- Kar, K., Kubilius, J., Schmidt, K., Issa, E. B., and DiCarlo, J. J. Evidence that recurrent circuits are critical to the ventral stream’s execution of core object recognition behavior. *Nature Neuroscience*, 22(6):974–983, 2019.
- Kietzmann, T. C., McClure, P., and Kriegeskorte, N. Deep neural networks in computational neuroscience. *BioRxiv*, pp. 133504, 2018.
- Kim, C. M. and Chow, C. C. Learning recurrent dynamics in spiking networks. *Elife*, 7:e37124, 2018.
- Kim, J., Kwon Lee, J., and Mu Lee, K. Deeply-recursive convolutional network for image super-resolution. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 1637–1645, 2016.
- Kim, R., Li, Y., and Sejnowski, T. J. Simple framework for constructing functional spiking recurrent neural networks. *Proceedings of the National Academy of Sciences*, 116(45):22811–22820, 2019.
- Klindt, D., Ecker, A. S., Euler, T., and Bethge, M. Neural system identification for large populations separating “what” and “where”. In *Advances in Neural Information Processing Systems (NeurIPS)*, pp. 3506–3516, 2017.
- Krizhevsky, A. Learning multiple layers of features from tiny images. Technical report, University of Toronto, Department of Computer Science, 04 2009.
- Kubilius, J., Schrimpf, M., Hong, H., Majaj, N. J., Rajalingham, R., Issa, E. B., Kar, K., Bashivan, P., Prescott-Roy, J., Schmidt, K., Nayebi, A., Bear, D., Yamins, D. L., and DiCarlo, J. J. Brain-like object recognition with high-performing shallow recurrent ANNs. In *Advances in Neural Information Processing Systems (NeurIPS)*, pp. 12785–12796, 2019.
- Laje, R. and Buonomano, D. V. Robust timing and motor patterns by taming chaos in recurrent neural networks. *Nature Neuroscience*, 16(7):925–933, 2013.
- Li, Z. A neural model of contour integration in the primary visual cortex. *Neural Computation*, 10(4):903–940, 1998.
- Liang, M. and Hu, X. Recurrent convolutional neural network for object recognition. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 3367–3375, 2015. doi: 10.1109/CVPR.2015.7298958. URL <https://doi.org/10.1109/CVPR.2015.7298958>.
- Liao, Q. and Poggio, T. Bridging the gaps between residual learning, recurrent neural networks and visual cortex. *arXiv preprint arXiv:1604.03640*, 2016.
- Lindeberg, T. Discrete derivative approximations with scale-space properties: A basis for low-level feature extraction. *Journal of Mathematical Imaging and Vision*, 3(4):349–376, 1993.
- Lindeberg, T. *Scale-space theory in computer vision*, volume 256. Springer Science & Business Media, 2013.
- Lindeberg, T. and Florack, L. *Foveal scale-space and the linear increase of receptive field size as a function of eccentricity*. KTH Royal Institute of Technology, 1994.
- Lindsey, J., Ocko, S. A., Ganguli, S., and Deny, S. A unified theory of early visual representations from retina to cortex through anatomically constrained deep CNNs. In *International Conference on Learning Representations (ICLR)*, 2019.
- Loog, M. and Lauze, F. Supervised scale-regularized linear convolutional filters. In *BMVC*, 2017.
- Lu, Y., Zhong, A., Li, Q., and Dong, B. Beyond finite layer neural networks: Bridging deep architectures and numerical differential equations. In *International Conference on Machine Learning (ICML)*, pp. 3276–3285. PMLR, 2018.
- Luan, S., Chen, C., Zhang, B., Han, J., and Liu, J. Gabor Convolutional Networks. *IEEE Transactions on Image Processing*, 27(9):4357–4366, 2018.
- Mallat, S. Group invariant scattering. *Communications on Pure and Applied Mathematics*, 65(10):1331–1398, 2012.
- Mante, V., Sussillo, D., Shenoy, K. V., and Newsome, W. T. Context-dependent computation by recurrent dynamics in prefrontal cortex. *nature*, 503(7474):78–84, 2013.
- Massaroli, S., Poli, M., Park, J., Yamashita, A., and Asama, H. Dissecting neural ODEs. *arXiv preprint arXiv:2002.08071*, 2020.
- Mastrogiuseppe, F. and Ostojic, S. Linking connectivity, dynamics, and computations in low-rank recurrent neural networks. *Neuron*, 99(3):609–623, 2018.
- Miconi, T., Stanley, K. O., and Clune, J. Differentiable plasticity: training plastic neural networks with backpropagation. In *Proceedings of the 35th International Conference on Machine Learning (ICML)*, volume 80, pp. 3556–3565, 2018.
- Pintea, S. L., Tomen, N., Goes, S. F., Loog, M., and van Gemert, J. C. Resolution learning in deep convolutional networks using scale-space theory. *arXiv preprint arXiv:2106.03412*, 2021.
- Rajan, K., Harvey, C. D., and Tank, D. W. Recurrent network models of sequence generation and memory. *Neuron*, 90(1):128–142, 2016.
- Richards, B. A., Lillicrap, T. P., Beaudoin, P., Bengio, Y., Bogacz, R., Christensen, A., Clopath, C., Costa, R. P., de Berker, A., Ganguli, S., Gillon, C. J., Hafner, D., Kepecs, A., Kriegeskorte, N., Latham, P., Lindsay, G. W., Miller, K. D., Naud, R., Pack, C. C., Poirazi, P., Roelfsema, P., Sacramento, J., Saxe, A., Scellier, B., Schapiro, A. C., Senn, W., Wayne, G., Yamins, D., Zenke, F., Zylberberg, J., Therien, D., and Kording, K. P. A deep learning framework for neuroscience. *Nature Neuroscience*, 22(11):1761–1770, 2019.
- Rousseau, F., Drumetz, L., and Fablet, R. Residual networks as flows of diffeomorphisms. *Journal of Mathematical Imaging and Vision*, pp. 1–11, 2019.
- Ruthotto, L. and Haber, E. Deep neural networks motivated by partial differential equations. *Journal of Mathematical Imaging and Vision*, pp. 1–13, 2019.
- Salman, H., Yadollahpour, P., Fletcher, T., and Batmanghelich, K. Deep diffeomorphic normalizing flows. *arXiv preprint arXiv:1810.03256*, 2018.

- Sceniak, M. P., Hawken, M. J., and Shapley, R. Contrast-dependent changes in spatial frequency tuning of macaque v1 neurons: effects of a changing receptive field size. *Journal of Neurophysiology*, 88(3):1363–1373, 2002.
- Schrimpf, M., Kubilius, J., Hong, H., Majaj, N. J., Rajalingham, R., Issa, E. B., Kar, K., Bashivan, P., Prescott-Roy, J., Geiger, F., Schmidt, K., Yamins, D. L. K., and DiCarlo, J. J. Brain-score: Which artificial neural network for object recognition is most brain-like? *bioRxiv preprint*, 2018.
- Sclar, G. and Freeman, R. Orientation selectivity in the cat’s striate cortex is invariant with stimulus contrast. *Experimental Brain Research*, 46(3):457–461, 1982.
- Sejnowski, T. J. The unreasonable effectiveness of deep learning in artificial intelligence. *Proceedings of the National Academy of Sciences*, 2020.
- Shelhamer, E., Wang, D., and Darrell, T. Blurring the line between structure and learning to optimize and adapt receptive fields. *arXiv preprint arXiv:1904.11487*, 2019.
- Simoncelli, E. P., Freeman, W. T., Adelson, E. H., and Heeger, D. J. Shiftable multiscale transforms. *IEEE transactions on Information Theory*, 38(2):587–607, 1992.
- Skottun, B. C., Bradley, A., Sclar, G., Ohzawa, I., and Freeman, R. D. The effects of contrast on visual orientation and spatial frequency discrimination: a comparison of single cells and behavior. *Journal of Neurophysiology*, 57(3):773–786, 1987.
- Smith, A. T., Singh, K. D., Williams, A., and Greenlee, M. W. Estimating receptive field size from fmri data in human striate and extrastriate visual cortex. *Cerebral Cortex*, 11(12):1182–1190, 2001.
- Somers, D. C., Nelson, S. B., and Sur, M. An emergent model of orientation selectivity in cat visual cortical simple cells. *Journal of Neuroscience*, 15(8):5448–5465, 1995.
- Sompolinsky, H., Crisanti, A., and Sommers, H.-J. Chaos in random neural networks. *Physical Review Letters*, 61(3):259, 1988.
- Song, H. F., Yang, G. R., and Wang, X.-J. Training excitatory-inhibitory recurrent neural networks for cognitive tasks: a simple and flexible framework. *PLoS Computational Biology*, 12(2):e1004792, 2016.
- Sosnovik, I., Szmaja, M., and Smeulders, A. Scale-equivariant steerable networks. *ICLR*, 2020.
- Spoerer, C. J., McClure, P., and Kriegeskorte, N. Recurrent convolutional neural networks: a better model of biological object recognition. *Frontiers in psychology*, 8:1551, 2017.
- Sussillo, D. and Abbott, L. F. Generating coherent patterns of activity from chaotic neural networks. *Neuron*, 63(4):544–557, 2009.
- Tang, H., Schrimpf, M., Lotter, W., Moerman, C., Paredes, A., Caro, J. O., Hardesty, W., Cox, D., and Kreiman, G. Recurrent computations for visual pattern completion. *Proceedings of the National Academy of Sciences*, 115(35):8835–8840, 2018.
- Van den Bergh, G., Zhang, B., Arekens, L., and Chino, Y. M. Receptive-field properties of v1 and v2 neurons in mice and macaque monkeys. *Journal of Comparative Neurology*, 518(11):2051–2070, 2010.
- Worrall, D. E., Garbin, S. J., Turmukhambetov, D., and Brostow, G. J. Harmonic networks: Deep translation and rotation equivariance. In *CVPR*, July 2017.
- Wu, Y. and He, K. Group normalization. In *ECCV*, pp. 3–19, 2018.
- Xiong, Z., Yuan, Y., Guo, N., and Wang, Q. Variational context-deformable convnets for indoor scene parsing. In *CVPR*, pp. 3992–4002, 2020.
- Yu, H.-H., Verma, R., Yang, Y., Tibballs, H. A., Lui, L. L., Reser, D. H., and Rosa, M. G. Spatial and temporal frequency tuning in striate cortex: functional uniformity and specializations related to receptive field eccentricity. *European Journal of Neuroscience*, 31(6):1043–1062, 2010.
- Zhuang, C., Yan, S., Nayebi, A., Schrimpf, M., Frank, M., DiCarlo, J., and Yamins, D. Unsupervised neural network models of the ventral visual stream. *bioRxiv*, 2020a.
- Zhuang, J., Dvornek, N., Li, X., Tatikonda, S., Papademetris, X., and Duncan, J. Adaptive checkpoint adjoint method for gradient estimation in neural ODE. *arXiv preprint arXiv:2006.02493*, 2020b.