

Scale-Equivariant Neural Networks with Decomposed Convolutional Filters

Wei Zhu^{*1}, Qiang Qiu², Robert Calderbank^{1,2}, Guillermo Sapiro², and Xiuyuan Cheng¹

¹Department of Mathematics, Duke University

²Department of Electrical and Computer Engineering, Duke University

Abstract

Encoding the input scale information explicitly into the representation learned by a convolutional neural network (CNN) is beneficial for many vision tasks especially when dealing with multiscale input signals. We study, in this paper, a scale-equivariant CNN architecture with joint convolutions across the space and the scaling group, which is shown to be both sufficient and necessary to achieve scale-equivariant representations. To reduce the model complexity and computational burden, we decompose the convolutional filters under two pre-fixed separable bases and truncate the expansion to low-frequency components. A further benefit of the truncated filter expansion is the improved deformation robustness of the equivariant representation. Numerical experiments demonstrate that the proposed scale-equivariant neural network with decomposed convolutional filters (ScDCFNet) achieves significantly improved performance in multiscale image classification and better interpretability than regular CNNs at a reduced model size.

1 Introduction

Convolutional neural networks (CNNs) have achieved great success in machine learning problems such as image classification [8], object detection [17], and semantic segmentation [10, 18]. Compared to fully-connected networks, CNNs through spatial weight sharing have the benefit of being translation-equivariant, i.e., translating the input leads to a translated version of the output. This property is crucial for many vision tasks such as image recognition and segmentation. However, regular CNNs are not equivariant to other important group transformations such as rescaling and rotation, and it is beneficial in some applications to also encode such group information explicitly into the network representation.

Several network architectures have been designed to achieve (2D) rotation-equivariance [3, 14, 25, 26, 31], and the feature maps of such networks typically include an extra index for the rotation group $SO(2)$. Building on the idea of group convolutions proposed in [5] for discrete symmetry groups, the authors in [3] and [25] constructed rotation-equivariant CNNs by conducting group

^{*}Email: zhu@math.duke.edu.

convolutions jointly across the space and $SO(2)$ using steerable filters [6]. Scale-equivariant CNNs, on the other hand, have only been studied in a less general setting in the existing literature [7, 13, 28]. In particular, to the best of our knowledge, a joint convolution across the space and the scaling group \mathcal{S} has yet been proposed to achieve scale-equivariance in the most general form. This is possibly because of two difficulties one encounters when dealing with the scaling group: First, unlike $SO(2)$, it is an acyclic and unbounded group; second, an extra index in \mathcal{S} incurs a significant increase in model parameters and computational burden, which is further exacerbated by a lack of the counterpart of “steerable filters” for the scaling group. Moreover, since the scaling transformation is rarely perfect in practice (due to changing view angle or numerical discretization), one needs to quantify and promote the deformation robustness of the equivariant representation (i.e., is the model still “approximately” equivariant if the scaling transformation is “contaminated” by a nuisance input deformation), which, to the best of our knowledge, has yet been studied in prior works.

The purpose of this paper is to address the aforementioned theoretical and practical issues in the construction of scale-equivariant CNN models. Specifically, our contribution is three-fold:

1. We propose a general scale-equivariant CNN architecture with a joint convolution over \mathbb{R}^2 and \mathcal{S} , which is proved in Section 4 to be both sufficient and necessary to achieve scale-equivariance.
2. A truncated decomposition of the convolutional filters under a pre-fixed separable basis on the two geometric domains (\mathbb{R}^2 and \mathcal{S}) is used to reduce the model size and computational cost.
3. We prove the representation stability of the proposed architecture up to equivariant scaling action of the input signal.

Our contribution to the family of group-equivariant CNNs is non-trivial; in particular, the scaling group unlike the rotation group is acyclic and non-compact. This poses challenges both in theory and in practice, so that many previous works on group-equivariant CNNs cannot be directly extended. We introduce new algorithm design and mathematical techniques to obtain the first general scale-equivariant CNN in literature with both computational efficiency and proved representation stability.

2 Related Work

Mixed-scale and scale-equivariant CNNs. Incorporating multiscale information into a CNN representation has been studied in many existing works. The Inception net [21] and its variants [20, 22] stack filters of different sizes in a single layer to address the multiscale salient features. Building on this idea, the selective kernel network [9] utilizes a nonlinear learnable mechanism to aggregate information from multiple scales. Dilated convolutions [15, 23, 29, 30] have also been used to combine multiscale information without increasing the model complexity. Although the effectiveness of such models have been empirically demonstrated in various vision tasks, there is still a lack of interpretability of their ability to encode the input scale information. Group-equivariant CNNs, on the other hand, explicitly encode the group information into the

network representation. Cohen and Welling [5] proposed CNNs with group convolutions that are equivariant to several finite discrete symmetry groups. This idea is later generalized in [4] and applied mainly to the rotation groups $SO(2)$ and $SO(3)$ [3, 24, 25]. Although scale-equivariant CNNs have also been proposed in the literature [7, 13, 28], they are typically studied in a less general setting. In particular, none of the previous works proposed to conduct joint convolutions over $\mathbb{R}^2 \times \mathcal{S}$ as a necessary and sufficient condition to impose scale-equivariance, for which reason they are thus variants of a special case of our proposed architecture where the convolutional filters in \mathcal{S} are Dirac delta functions (c.f. Remark 1.)

Representation stability to input deformations. Input deformations typically induce noticeable variabilities within object classes, some of which are uninformative for the vision tasks. Models that are stable to input deformations are thus favorable in many applications. The scattering transform [2, 11, 12] computes translation-invariant representations that are Lipschitz continuous to deformations by cascading predefined wavelet transforms and modulus poolings. A joint convolution over $\mathbb{R}^2 \times SO(2)$ is later adopted in [19] to build roto-translation scattering with stable rotation/translation-invariant representations. These models, however, use pre-fixed wavelet transforms in the networks, and are thus nonadaptive to the data. DCFNet [16] combines a pre-fixed filter basis and learnable expansion coefficients in a CNN architecture, achieving both data adaptivity and representation stability inherited from the filter regularity. This idea is later extended in [3] to produce rotation-equivariant representations that are Lipschitz continuous in L^2 norm to input deformations modulo a global rotation, i.e., the model stays approximately equivariant even if the input rotation is imperfect. To the best of our knowledge, a theoretical analysis of the deformation robustness of a scale-equivariant CNN has yet been studied, and a direct generalization of the result in [3] is futile because the feature maps of a scale-equivariant CNN is typically not in L^2 (c.f. Remark 2.)

3 Scale-Equivariant CNN and Filter Decomposition

Group-equivariance is the property of a mapping $f : X \rightarrow Y$ to commute with the group actions on the domain X and codomain Y . More specifically, let G be a group, and D_g, T_g , respectively, be group actions on X and Y . A function $f : X \rightarrow Y$ is said to be G -equivariant if

$$f(D_g x) = T_g(f(x)), \quad \forall g \in G, x \in X.$$

G -invariance is thus a special case of G -equivariance where $T_g = \text{Id}_Y$. For learning tasks where the feature $y \in Y$ is known a priori to change equivariantly to a group action $g \in G$ on the input $x \in X$, e.g. image segmentation should be equivariant to translation, it would be beneficial to reduce the hypothesis space to include only G -equivariant models. In this paper, we consider mainly the scaling-translation group $\mathcal{ST} \cong \mathcal{S} \times \mathbb{R}^2 \cong \mathbb{R} \times \mathbb{R}^2$. Given $g = (\beta, v) \in \mathcal{ST}$ and an input image $x^{(0)}(u, \lambda)$ ($u \in \mathbb{R}^2$ is the spatial position, and λ is the unstructured channel index, e.g. RGB channels of a color image), the scaling-translation group action $D_g = D_{\beta, v}$ on $x^{(0)}$ is defined as

$$D_{\beta, v} x^{(0)}(u, \lambda) := x^{(0)}\left(2^{-\beta}(u - v), \lambda\right). \quad (1)$$




Figure 1: (a) A special case of scale-equivariant CNN with only (multiscale) spatial convolutions. The previous works on scale-equivariant CNNs [7, 13, 28] are all variants of this architecture. (b) The general case of scale-equivariant CNN with joint convolutions (Theorem 1) where information transfers among different scales. See Remark 1 for more explanation.

Constructing scale-equivariant CNNs thus amounts to finding an architecture \mathcal{A} such that each trained network $f \in \mathcal{A}$ commutes with the group action $D_{\beta, v}$ on the input and a similarly defined group action $T_{\beta, v}$ (to be explained in Section 3.1) on the output.

3.1 Scale-Equivariant CNNs

Inspired by [3] and [25], we consider scale-equivariant CNNs with an extra index $\alpha \in \mathcal{S}$ for the scaling group $\mathcal{S} \cong \mathbb{R}$: for each $l \geq 1$, the l -th layer output is denoted as $x^{(l)}(u, \alpha, \lambda)$, where $u \in \mathbb{R}^2$ is the spatial position, $\alpha \in \mathcal{S}$ is the scale index, and $\lambda \in [M_l] := \{1, \dots, M_l\}$ corresponds to the unstructured channels. We use the continuous model for formal derivation, i.e., the images and feature maps have continuous spatial and scale indices. In practice, the images are discretized on a Cartesian grid, and the scales are computed only on a discretized finite interval. Similar to [3], the group action $T_{\beta, v}$ on the l -th layer output is defined as a scaling-translation in space as well as a shift in the scale channel:

$$T_{\beta, v} x^{(l)}(u, \alpha, \lambda) := x^{(l)}(2^{-\beta}(u - v), \alpha - \beta, \lambda), \quad \forall l \geq 1. \quad (2)$$

A feedforward neural network is said to be scale-equivariant, i.e., equivariant to \mathcal{ST} , if

$$x^{(l)}[D_{\beta, v} x^{(0)}] = T_{\beta, v} x^{(l)}[x^{(0)}], \quad \forall l \geq 1, \quad (3)$$

where we slightly abuse the notation $x^{(l)}[x^{(0)}]$ to denote the l -th layer output given the input $x^{(0)}$. The following Theorem shows that scale-equivariance is achieved if and only if joint convolutions are conducted over $\mathcal{S} \times \mathbb{R}^2$ as in (4) and (5).

Theorem 1. *A feedforward neural network with an extra index $\alpha \in \mathcal{S}$ for layerwise output is scale-equivariant if and only if the layerwise operations are defined as (4) and (5):*

$$x^{(1)}[x^{(0)}](u, \alpha, \lambda) = \sigma \left(\sum_{\lambda'} \int_{\mathbb{R}^2} 2^{-2\alpha} x^{(0)}(u + u', \lambda') W_{\lambda', \lambda}^{(1)}(2^{-\alpha} u') du' + b^{(1)}(\lambda) \right), \quad (4)$$

$$x^{(l)}[x^{(l-1)}](u, \alpha, \lambda) = \sigma \left(\sum_{\lambda'} \int_{\mathbb{R}^2} \int_{\mathbb{R}} 2^{-2\alpha} x^{(l-1)}(u + u', \alpha + \alpha', \lambda') \cdot W_{\lambda', \lambda}^{(l)}(2^{-\alpha} u', \alpha') d\alpha' du' + b^{(l)}(\lambda) \right), \quad \forall l > 1, \quad (5)$$

where $\sigma : \mathbb{R} \rightarrow \mathbb{R}$ is a pointwise nonlinear function.

We defer the proof of Theorem 1, as well as those of other theorems, to the appendix.

Remark 1. When the (joint) convolutional filter $W_{\lambda', \lambda}^{(l)}(u, \alpha)$ takes the special form $W_{\lambda', \lambda}^{(l)}(u, \alpha) = V_{\lambda', \lambda}^{(l)}(u) \delta(\alpha)$, the joint convolution (5) over $\mathbb{R}^2 \times \mathcal{S}$ reduces to only a (multiscale) spatial convolution

$$x^{(l)}[x^{(l-1)}](u, \alpha, \lambda) = \sigma \left(\sum_{\lambda'} \int_{\mathbb{R}^2} x^{(l-1)}(u + u', \alpha, \lambda') V_{\lambda', \lambda}^{(l)}(2^{-\alpha} u') 2^{-2\alpha} du' + b^{(l)}(\lambda) \right),$$

i.e., the feature maps at different scales do not transfer information among each other (see Figure 1a). The previous works [7, 13, 28] on scale-equivariant CNNs are all based on this special case of Theorem 1.

Although the joint convolutions (5) on $\mathbb{R}^2 \times \mathcal{S}$ provide the most general way of imposing scale-equivariance, they unfortunately also incur a significant increase in the model size and computational burden. Following the idea of [3] and [16], we address this issue by taking a truncated decomposition of the convolutional filters under a pre-fixed separable basis, which will be discussed in detail in the next section.

3.2 Separable Basis Decomposition

We consider decomposing the convolutional filters $W_{\lambda', \lambda}^{(l)}(u, \alpha)$ under the product of two function bases, $\{\psi_k(u)\}_k$ and $\{\varphi_m(\alpha)\}_m$, which are the eigenfunctions of the Dirichlet Laplacian on, respectively, the unit disk $D \subset \mathbb{R}^2$ and $[-1, 1]$, i.e.,

$$\begin{cases} \Delta \psi_k = -\mu_k \psi_k & \text{in } D, \\ \psi_k = 0 & \text{on } \partial D, \end{cases} \quad \text{and} \quad \begin{cases} \varphi_m'' = -\nu_m \varphi_m & \text{in } [-1, 1] \\ \varphi_m(-1) = \varphi_m(1) = 0. \end{cases} \quad (6)$$

In particular, the spatial basis $\{\psi_k\}_k$ satisfying (6) is the Fourier-Bessel (FB) basis [1]. In the continuous formulation, the spatial “pooling” operation is equivalent to rescaling the convolutional filters in space. We thus assume, without loss of generality, that the convolutional filters are compactly supported as follows

$$W_{\lambda', \lambda}^{(1)} \in C_c(2^{j_1} D), \quad \text{and} \quad W_{\lambda', \lambda}^{(l)} \in C_c(2^{j_l} D \times [-1, 1]), \quad \forall l > 1.$$

Let $\psi_{j,k}(u) := 2^{-2j} \psi_k(2^{-j} u)$, then $W_{\lambda', \lambda}^{(l)}$ can be decomposed under $\{\psi_{j_l, k}\}_k$ and $\{\varphi_m\}_m$ as

$$W_{\lambda', \lambda}^{(1)}(u) = \sum_k a_{\lambda', \lambda}^{(1)}(k) \psi_{j_1, k}(u), \quad W_{\lambda', \lambda}^{(l)}(u, \alpha) = \sum_m \sum_k a_{\lambda', \lambda}^{(l)}(k, m) \psi_{j_l, k}(u) \varphi_m(\alpha), \quad l > 1 \quad (7)$$

where $a_{\lambda',\lambda}^{(1)}(k)$ and $a_{\lambda',\lambda}^{(l)}(k, m)$ are the expansion coefficients of the filters under the joint bases. During training, the basis functions are fixed, and only the expansion coefficients are updated. In practice, we truncate the expansion to only low-frequency coefficients (i.e., $a_{\lambda',\lambda}^{(l)}(k, m)$ are non-zero only for $k \in [K]$, $m \in [K_\alpha]$), which are kept as the trainable parameters. This directly leads to a reduction of network parameters and computational burden. More specifically, let us compare the l -th convolutional layer (5) of a scale-equivariant CNN with and without truncated basis decomposition:

Number of trainable parameters: Suppose the filters $W_{\lambda',\lambda}^{(l)}(u, \alpha)$ are discretized on a Cartesian grid of size $L \times L \times L_\alpha$. The number of trainable parameters at the l -th layer of a scale-equivariant CNN without basis decomposition is $L^2 L_\alpha M_{l-1} M_l$. On the other hand, in an ScDCFNet with truncated basis expansion up to K leading coefficients for u and K_α coefficients for α , the number of parameters is instead $KK_\alpha M_{l-1} M_l$. Hence a reduction to a factor of $KK_\alpha/L^2 L_\alpha$ in trainable parameters is achieved for ScDCFNet via truncated basis decomposition. In particular, if $L = 5$, $L_\alpha = 5$, $K = 8$, and $K_\alpha = 3$, then the number of parameters is reduced to $(8 \times 3)/(5^2 \times 5) = 19.2\%$.

Computational cost: Suppose the size of the input $x^{(l-1)}(u, \alpha, \lambda)$ and output $x^{(l)}(u, \alpha, \lambda)$ at the l -th layer are, respectively, $W \times W \times N_\alpha \times M_{l-1}$ and $W \times W \times N_\alpha \times M_l$, where $W \times W$ is the spatial dimension, N_α is the number of scale channels, and $M_{l-1}(M_l)$ is the number of the unstructured input (output) channels. Let the filters $W_{\lambda',\lambda}^{(l)}(u, \alpha)$ be discretized on a Cartesian grid of size $L \times L \times L_\alpha$. The following theorem shows that, compared to a regular scale-equivariant CNN, the computational cost in a forward pass of ScDCFNet is reduced again to a factor of $KK_\alpha/L^2 L_\alpha$.

Theorem 2. *Assume $M_l \gg L^2, L_\alpha$, i.e., the number of the output channels is much larger than the size of the convolutional filters in u and α , then the computational cost of an ScDCFNet is reduced to a factor of $KK_\alpha/L^2 L_\alpha$ when compared to a scale-equivariant CNN without basis decomposition.*

Apart from reducing the model size and computational burden, similar to [3], truncating the filter decomposition has the further benefit of improving the deformation robustness of the equivariant representation, i.e., the equivariance relation (3) still approximately holds true if the spatial scaling of the input $D_{\beta,v}x^{(0)}$ is contaminated by a local deformation (e.g., due to changing view angle or numerical discretization.) This will be addressed in detail in the next section.

4 Representation Stability of ScDCFNet to Input Deformation

We study, in this section, the representation stability of ScDCFNet to input deformations modulo a global scale change, i.e., the input undergoes not only a scale change but also a small spatial distortion. To quantify the distance between different feature maps at each layer, we define the norm of $x^{(l)}$ as

$$\|x^{(0)}\|^2 = \frac{1}{M_0} \sum_{\lambda=1}^{M_0} \int \left| x^{(0)}(u, \lambda) \right|^2 du, \quad \|x^{(l)}\|^2 = \sup_{\alpha} \frac{1}{M_l} \sum_{\lambda=1}^{M_l} \int \left| x^{(l)}(u, \alpha, \lambda) \right|^2 du, \quad l \geq 1. \quad (8)$$

Remark 2. The definition of $\|x^{(l)}\|$ is different from that of RotDCFNet [3], where an L^2 norm is taken for the α index as well. The reason why we adopt the L^∞ norm for α in (8) is that $x^{(l)}$ is typically not L^2 in α , since the scaling group \mathcal{S} , unlike $SO(2)$, has infinite Haar measure.

We next quantify the representation stability of ScDCFNet under three mild assumptions on the convolutional layers and input deformations. First,

(A1) The pointwise nonlinear activation $\sigma : \mathbb{R} \rightarrow \mathbb{R}$ is non-expansive.

Next, we need a bound on the convolutional filters under certain norms. For each $l \geq 1$, define A_l as

$$\begin{cases} A_1 := \pi \max \left\{ \sup_{\lambda} \sum_{\lambda'=1}^{M_0} \|a_{\lambda',\lambda}^{(1)}\|_{\text{FB}}, \frac{M_0}{M_1} \sup_{\lambda'} \sum_{\lambda=1}^{M_1} \|a_{\lambda',\lambda}^{(1)}\|_{\text{FB}} \right\}, \\ A_l := \pi \max \left\{ \sup_{\lambda} \sum_{\lambda'=1}^{M_{l-1}} \sum_m \|a_{\lambda',\lambda}^{(l)}(\cdot, m)\|_{\text{FB}}, \frac{2M_{l-1}}{M_l} \sum_m \sup_{\lambda'} \sum_{\lambda=1}^{M_l} \|a_{\lambda',\lambda}^{(l)}(\cdot, m)\|_{\text{FB}} \right\}, \end{cases} \quad (9)$$

where the Fourier-Bessel (FB) norm $\|a\|_{\text{FB}}$ of a sequence $\{a(k)\}_{k \geq 0}$ is a weighted l^2 norm defined as $\|a\|_{\text{FB}}^2 := \sum_k \mu_k a(k)^2$, where μ_k is the k -th eigenvalue of the Dirichlet Laplacian on the unit disk defined in (6). We next assume that each A_l is bounded:

(A2) For all $l \geq 1$, $A_l \leq 1$.

The boundedness of A_l is facilitated by truncating the basis decomposition to only low-frequency components (small μ_k), which is one of the key idea of ScDCFNet explained in Section 3.2. After a proper initialization of the trainable coefficients, (A2) can generally be satisfied. The assumption (A2) implies several bounds on the convolutional filters at each layer (c.f. Lemma 2 in the appendix), which, combined with (A1), guarantees that an ScDCFNet is layerwise non-expansive:

Proposition 1. Under the assumption (A1) and (A2), an ScDCFNet satisfies the following.

(a) For any $l \geq 1$, the mapping of the l -th layer, $x^{(l)}[\cdot]$ defined in (4) and (5), is non-expansive, i.e.,

$$\|x^{(l)}[x_1] - x^{(l)}[x_2]\| \leq \|x_1 - x_2\|, \quad \forall x_1, x_2.$$

(b) Let $x_0^{(l)}$ be the l -th layer output given a zero bottom-layer input, then $x_0^{(l)}(\lambda)$ depends only on λ .

(c) Let $x_c^{(l)}$ be the centered version of $x^{(l)}$ after removing $x_0^{(l)}$, i.e.,

$$x_c^{(0)}(u, \lambda) := x^{(0)}(u, \lambda) - x_0^{(0)}(\lambda) = x^{(0)}(u, \lambda), \quad x_c^{(l)}(u, \alpha, \lambda) := x^{(l)}(u, \alpha, \lambda) - x_0^{(l)}(\lambda), \quad l \geq 1,$$

then $\|x_c^{(l)}\| \leq \|x_c^{(l-1)}\|$, $\forall l \geq 1$. As a result, $\|x_c^{(l)}\| \leq \|x_c^{(0)}\| = \|x^{(0)}\|$.

Finally, we make an assumption on the input deformation modulo a global scale change. Given a C^2 function $\tau : \mathbb{R}^2 \rightarrow \mathbb{R}^2$, the spatial deformation D_τ on the feature maps $x^{(l)}$ is defined as

$$D_\tau x^{(0)}(u, \lambda) = x^{(0)}(\rho(u), \lambda), \quad \text{and} \quad D_\tau x^{(l)}(u, \alpha, \lambda) = x^{(l)}(\rho(u), \alpha, \lambda), \quad l \geq 1, \quad (10)$$

where $\rho(u) = u - \tau(u)$. We assume a small local deformation on the input:

(A3) $|\nabla\tau|_\infty := \sup_u \|\nabla\tau(u)\| < 1/5$, where $\|\cdot\|$ is the operator norm.

The following theorem demonstrates the representation stability of an ScDCFNet to input deformation modulo a global scale change.

Theorem 3. *Let D_τ be a small spatial deformation defined in (10), and let $D_{\beta,v}, T_{\beta,v}$ be the group actions corresponding to an arbitrary scaling $2^{-\beta} \in \mathbb{R}_+$ centered at $v \in \mathbb{R}^2$ defined in (1) and (2). In an ScDCFNet satisfying (A1), (A2), and (A3), we have, for any L ,*

$$\left\| x^{(L)}[D_{\beta,v} \circ D_\tau x^{(0)}] - T_{\beta,v} x^{(L)}[x^{(0)}] \right\| \leq 2^{\beta+1} (4L|\nabla\tau|_\infty + 2^{-jL}|\tau|_\infty) \|x^{(0)}\|. \quad (11)$$

Theorem 3 gauges how approximately equivariant is ScDCFNet if the input undergoes not only a scale change $D_{\beta,v}$ but also a nonlinear spatial deformation D_τ , which is important both in theory and in practice because the scaling of an object is rarely perfect in reality.

5 Numerical Experiments

In this section, we conduct several numerical experiments for the following three purposes.

1. To verify that ScDCFNet indeed achieves scale equivariance (3).
2. To illustrate that ScDCFNet significantly outperforms regular CNNs at a much reduced model size in multiscale image classification.
3. To show that a trained ScDCFNet auto-encoder is able to reconstruct rescaled versions of the input by simply applying group actions on the image codes, demonstrating that ScDCFNet indeed explicitly encodes the input scale information into the representation.

The experiments are tested on the Scaled MNIST (SMNIST) and Scaled Fashion-MNIST (SFashion) datasets, which are built by rescaling the original MNIST and Fashion-MNIST [27] images by a factor randomly sampled from a uniform distribution on $[0.3, 1]$. A zero-padding to a size of 28×28 is conducted after the rescaling. If mentioned explicitly, for some experiments, the images are resized to 64×64 for better visualization.

Before going into the details of the numerical results, we need to clarify the implementation of the spatial pooling module of ScDCFNet. Given a feature $x^{(l)}(u, \alpha, \lambda)$, the traditional average-pooling in u with the same spatial kernel size across α destroys scale equivariance (3). To remedy this, we first convolve $x^{(l)}$ with a scale-specific low-pass filter before downsampling the convolved signal on a coarser spatial grid. Specifically, we have $\tilde{x}^{(l)}(u, \alpha, \lambda) = \int_{\mathbb{R}^2} x^{(l)}(Nu + u', \alpha, \lambda) \eta(2^{-\alpha}u') 2^{-2\alpha} du'$, where $\tilde{x}^{(l)}$ is the feature after pooling, η is a low-pass filter, e.g., a Gaussian kernel, and $N \in \mathbb{Z}_+$ is the pooling factor. We will refer to this as scale-equivariant average-pooling in what follows.

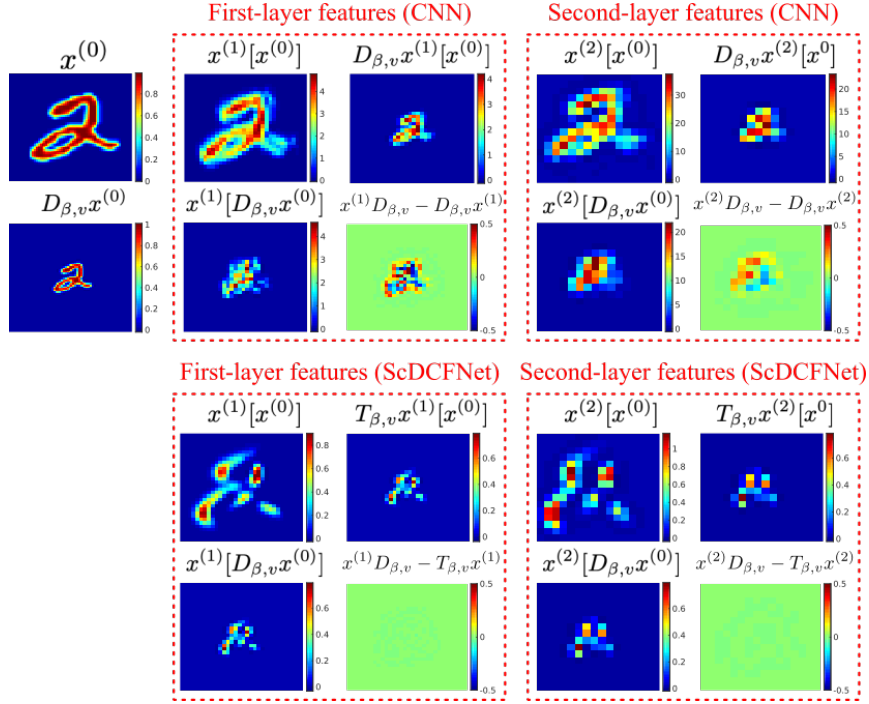


Figure 2: Verification of scale-equivariance in Section 5.1. Given the original input $x^{(0)}$ and its rescaled version $D_{\beta,v}x^{(0)}$, the four figures in each dashed rectangle are: $x^{(l)}[x^{(0)}]$ (l -th layer feature of the original input), $x^{(l)}[D_{\beta,v}x^{(0)}]$ (l -th layer feature of the rescaled input), $T_{\beta,v}x^{(l)}[x^{(0)}]$ (rescaled l -th layer feature of the original input), and the difference $(x^{(l)}[D_{\beta,v}x^{(0)}] - T_{\beta,v}x^{(l)}[x^{(0)}])$ displayed in a (signal intensity) scale relative to the maximum value of $x^{(l)}[D_{\beta,v}x^{(0)}]$. It is clear that even after numerical discretization, scale-equivariance still approximately holds for ScDCFNet, i.e., $x^{(l)}[D_{\beta,v}x^{(0)}] - T_{\beta,v}x^{(l)}[x^{(0)}] \approx 0$, but not for a regular CNN.

5.1 Verification of Scale Equivariance

We first verify that ScDCFNet indeed achieves scale-equivariance (3). Specifically, we compare the feature maps of a two-layer ScDCFNet with randomly generated truncated filter expansion coefficients and those of a regular CNN. The exact architectures are detailed in Appendix B.1. Figure 2 displays the first- and second-layer feature maps of an original image $x^{(0)}$ and its rescaled version $D_{\beta,v}x^{(0)}$ using the two comparing architectures. Feature maps at different layers are rescaled to the same spatial dimension for visualization. The four images enclosed in each of the dashed rectangle correspond to: $x^{(l)}[x^{(0)}]$ (l -th layer feature of the original input), $x^{(l)}[D_{\beta,v}x^{(0)}]$ (l -th layer feature of the rescaled input), $T_{\beta,v}x^{(l)}[x^{(0)}]$ (rescaled l -th layer feature of the original input, where $T_{\beta,v}$ is understood as $D_{\beta,v}$ for a regular CNN due to the lack of a scale index α), and the difference $x^{(l)}[D_{\beta,v}x^{(0)}] - T_{\beta,v}x^{(l)}[x^{(0)}]$. It is clear that even with numerical discretization, which can be modeled as a form of input deformation, ScDCFNet is still approximately scale-equivariant, i.e., $x^{(l)}[D_{\beta,v}x^{(0)}] \approx T_{\beta,v}x^{(l)}[x^{(0)}]$, whereas a regular CNN does not have such a property.

5.2 Multiscale Image Classification

Without batch-normalization		SMNIST test accuracy (%)		SFashion test accuracy (%)	
Architectures	Ratio	$N_{\text{tr}} = 2000$	$N_{\text{tr}} = 5000$	$N_{\text{tr}} = 2000$	$N_{\text{tr}} = 5000$
CNN, $M = 32$	1.00	92.60 ± 0.17	94.86 ± 0.25	77.74 ± 0.28	82.57 ± 0.38
ScDCFNet, $M = 16$					
$K = 10, K_\alpha = 3$	0.84	93.75 ± 0.02	95.70 ± 0.09	78.95 ± 0.31	83.51 ± 0.71
$K = 8, K_\alpha = 3$	0.67	93.91 ± 0.30	95.71 ± 0.10	79.22 ± 0.50	83.06 ± 0.32
$K = 5, K_\alpha = 3$	0.42	93.52 ± 0.29	95.19 ± 0.13	79.74 ± 0.44	83.46 ± 0.69
$K = 10, K_\alpha = 2$	0.56	93.68 ± 0.23	95.54 ± 0.21	79.01 ± 0.61	83.43 ± 0.60
$K = 8, K_\alpha = 2$	0.45	93.67 ± 0.19	95.51 ± 0.20	79.15 ± 0.59	83.44 ± 0.37
$K = 5, K_\alpha = 2$	0.28	93.51 ± 0.30	95.35 ± 0.21	78.57 ± 0.53	82.95 ± 0.46
ScDCFNet, $M = 8$					
$K = 10, K_\alpha = 2$	0.14	93.68 ± 0.17	95.21 ± 0.12	79.11 ± 0.76	82.92 ± 0.68
$K = 8, K_\alpha = 2$	0.11	93.39 ± 0.25	95.25 ± 0.47	78.43 ± 0.76	82.53 ± 0.58
$K = 5, K_\alpha = 2$	0.09	93.21 ± 0.20	94.99 ± 0.12	77.97 ± 0.37	82.21 ± 0.67
$K = 10, K_\alpha = 1$	0.09	93.26 ± 0.52	95.22 ± 0.31	78.71 ± 0.46	83.05 ± 0.51
$K = 8, K_\alpha = 1$	0.06	93.61 ± 0.11	95.03 ± 0.31	78.69 ± 0.72	82.96 ± 0.27
$K = 5, K_\alpha = 1$	0.04	93.23 ± 0.13	94.84 ± 0.38	78.38 ± 0.81	82.14 ± 0.25
With batch-normalization		SMNIST test accuracy (%)		SFashion test accuracy (%)	
Architectures	Ratio	$N_{\text{tr}} = 2000$	$N_{\text{tr}} = 5000$	$N_{\text{tr}} = 2000$	$N_{\text{tr}} = 5000$
CNN, $M = 32$	1.00	94.78 ± 0.17	96.58 ± 0.17	79.79 ± 0.40	84.38 ± 0.17
ScDCFNet, $M = 16$					
$K = 10, K_\alpha = 3$	0.84	95.69 ± 0.15	96.99 ± 0.23	81.12 ± 0.42	84.97 ± 0.16
$K = 8, K_\alpha = 3$	0.67	95.72 ± 0.29	97.15 ± 0.24	81.41 ± 0.35	85.11 ± 0.26
$K = 5, K_\alpha = 3$	0.42	95.31 ± 0.21	96.75 ± 0.13	81.73 ± 0.14	84.70 ± 0.19
$K = 10, K_\alpha = 2$	0.56	95.58 ± 0.16	96.92 ± 0.07	81.21 ± 0.20	84.77 ± 0.15
$K = 8, K_\alpha = 2$	0.45	95.76 ± 0.17	96.74 ± 0.22	80.84 ± 0.28	84.70 ± 0.18
$K = 5, K_\alpha = 2$	0.28	95.43 ± 0.09	96.54 ± 0.20	81.23 ± 0.23	84.65 ± 0.36

Table 1: Classification accuracy on the SMNIST and SFashion dataset with and without batch-normalization. The architectures are detailed in Table 2. In particular, M stands for the number of the first-layer (unstructured) output channels, which is doubled after each layer, and K/K_α is the number of basis function in u/α used for filter decomposition. The networks are tested with different training data size, $N_{\text{tr}} = 2000$ and 5000 , and the means and standard deviations after three independent trials are reported. The column “ratio” stands for the ratio between the number of trainable parameters of the current architecture and that of the baseline CNN.

We next demonstrate the improved performance of ScDCFNet in multiscale image classification. The experiments are conducted on the SMNIST and SFashion datasets, and a regular CNN is used as a performance benchmark. Both networks are comprised of three convolutional layers with the exact architectures (Table 2) detailed in Appendix B.2. Unlike $SO(2)$, the scaling group $\mathcal{S} \cong \mathbb{R}$ is unbounded, and we thus compute only the feature maps $x^{(l)}(u, \alpha, \lambda)$ with the α index restricted to the scale interval $[-1.6, 0]$ ($2^{-1.6} \approx 0.3$), which is discretized uniformly into $N_\alpha = 9$ channels. The performance of the comparing architectures with and without batch-normalization is shown in Table 1. It is clear that, by limiting the hypothesis space to

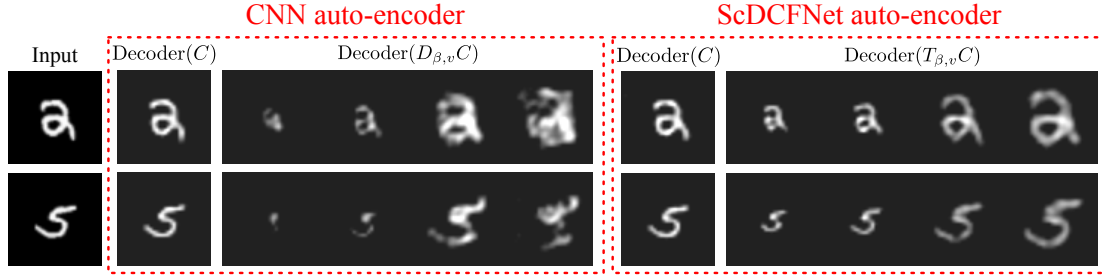


Figure 3: Reconstructing rescaled versions of the original test image by manipulating its image code C according to the group action (2). The first two images on the left are the original inputs; $\text{Decoder}(C)$ denotes the reconstruction using the (unchanged) image code C ; $\text{Decoder}(D_{\beta,v}C)$ and $\text{Decoder}(T_{\beta,v}C)$ denote the reconstructions using the “rescaled” image codes $D_{\beta,v}C$ and $T_{\beta,v}C$ respectively according to (1) and (2). Unlike the regular CNN auto-encoder, the ScDCFNet auto-encoder manages to generate rescaled versions of the original input, suggesting that it successfully encodes the scale information directly into the representation.

scale-equivariant models and taking truncated basis decomposition to reduce the model size, ScDCFNet achieves a significant improvement in classification accuracy with a reduced number of trainable parameters. The advantage of ScDCFNet is more pronounced when the number of training samples is small ($N_{\text{tr}} = 2000$), suggesting that, by hardwiring the input scale information directly into its representation, ScDCFNet is less susceptible to overfitting the limited multiscale training data.

We also observe that even when a regular CNN is trained with data augmentation (random cropping and rescaling), its performance is still inferior to that of an ScDCFNet without manipulation of the training data. In particular, although the accuracies of the regular CNNs trained on 2000 SMNIST and SFashion images after data augmentation are improved to, respectively, 93.85% and 79.41%, they still underperform the ScDCFNets without data augmentation (93.91% and 79.94%) using only a fraction of trainable parameters. Moreover, if ScDCFNet is trained with data augmentation, the accuracies can be further improved to 94.30% and 80.62% respectively. This suggests that ScDCFNet can be combined with data augmentation for optimal performance in multiscale image classification.

5.3 Image Reconstruction

In the last experiment, we illustrate the ability of ScDCFNet to explicitly encode the input scale information into the representation. To achieve this, we train an ScDCFNet auto-encoder on the SMNIST dataset with images resized to 64×64 for better visualization. The encoder stacks two scale-equivariant convolutional blocks with 2×2 average-pooling, and the decoder contains a succession of two transposed convolutional blocks with 2×2 upsampling. A regular CNN auto-encoder is also trained for comparison (see Table 3 in Appendix B.3 for the detailed architecture.)

Our goal is to demonstrate that the image code produced by the ScDCFNet auto-encoder contains

the scale information of the input, i.e., by applying the group action $T_{\beta,v}$ (2) to the code C of a test image before feeding it to the decoder, we can reconstruct rescaled versions of original input. This property can be visually verified in Figure 3. In contrast, a regular CNN auto-encoder fails to do so.

6 Conclusion

We propose, in this paper, a scale-equivariant CNN with joint convolutions across the space \mathbb{R}^2 and the scaling group \mathcal{S} , which we show to be both sufficient and necessary to impose scale-equivariant network representation. To reduce the computational cost and model complexity incurred by the joint convolutions, the convolutional filters supported on $\mathbb{R}^2 \times \mathcal{S}$ are decomposed under a separable basis across the two domains and truncated to only low-frequency components. Moreover, the truncated filter expansion leads also to improved deformation robustness of the equivariant representation, i.e., the model is still approximately equivariant even if the scaling transformation is imperfect. Experimental results suggest that ScDCFNet achieves improved performance in multiscale image classification with greater interpretability and reduced model size compared to regular CNN models.

For future work, we will study the application of ScDCFNet in other more complicated vision tasks, such as object detection/localization and pose estimation, where it is beneficial to directly encode the input scale information into the deep representation. Moreover, the memory usage of our current implementation of ScDCFNet scales linearly to the number of the truncated basis functions in order to realize the reduced computational burden explained in Theorem 2. We will explore other efficient implementation of the model, e.g., using filter-bank type of techniques to compute convolutions with multiscale spatial filters, to significantly reduce both the computational cost and memory usage.

References

- [1] M. Abramowitz and I. A. Stegun. *Handbook of mathematical functions: with formulas, graphs, and mathematical tables*, volume 55. Courier Corporation, 1965.
- [2] J. Bruna and S. Mallat. Invariant scattering convolution networks. *IEEE transactions on pattern analysis and machine intelligence*, 35(8):1872–1886, 2013.
- [3] X. Cheng, Q. Qiu, R. Calderbank, and G. Sapiro. RotDCF: Decomposition of convolutional filters for rotation-equivariant deep networks. In *International Conference on Learning Representations*, 2019.
- [4] T. Cohen, M. Geiger, and M. Weiler. A general theory of equivariant cnns on homogeneous spaces. *arXiv preprint arXiv:1811.02017*, 2018.
- [5] T. Cohen and M. Welling. Group equivariant convolutional networks. In *International conference on machine learning*, pages 2990–2999, 2016.

- [6] W. T. Freeman and E. H. Adelson. The design and use of steerable filters. *IEEE Transactions on Pattern Analysis & Machine Intelligence*, (9):891–906, 1991.
- [7] A. Kim, A. Sharma, and D. Jacobs. Locally scale-invariant convolutional neural networks. *arXiv preprint arXiv:1412.5104*, 2014.
- [8] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pages 1097–1105, 2012.
- [9] X. Li, W. Wang, X. Hu, and J. Yang. Selective kernel networks. 2019.
- [10] J. Long, E. Shelhamer, and T. Darrell. Fully convolutional networks for semantic segmentation. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2015.
- [11] S. Mallat. Recursive interferometric representation. In *Proc. of EUSICO conference, Denmark*, 2010.
- [12] S. Mallat. Group invariant scattering. *Communications on Pure and Applied Mathematics*, 65(10):1331–1398, 2012.
- [13] D. Marcos, B. Kellenberger, S. Lobry, and D. Tuia. Scale equivariance in cnns with vector fields. *arXiv preprint arXiv:1807.11783*, 2018.
- [14] D. Marcos, M. Volpi, N. Komodakis, and D. Tuia. Rotation equivariant vector field networks. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 5048–5057, 2017.
- [15] D. M. Pelt and J. A. Sethian. A mixed-scale dense convolutional neural network for image analysis. *Proceedings of the National Academy of Sciences*, 115(2):254–259, 2018.
- [16] Q. Qiu, X. Cheng, R. Calderbank, and G. Sapiro. DCFNet: Deep neural network with decomposed convolutional filters. In *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pages 4198–4207, Stockholmsmässan, Stockholm Sweden, 10–15 Jul 2018. PMLR.
- [17] S. Ren, K. He, R. Girshick, and J. Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. In C. Cortes, N. D. Lawrence, D. D. Lee, M. Sugiyama, and R. Garnett, editors, *Advances in Neural Information Processing Systems 28*, pages 91–99. Curran Associates, Inc., 2015.
- [18] O. Ronneberger, P. Fischer, and T. Brox. U-net: Convolutional networks for biomedical image segmentation. In N. Navab, J. Hornegger, W. M. Wells, and A. F. Frangi, editors, *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015*, pages 234–241, Cham, 2015. Springer International Publishing.
- [19] L. Sifre and S. Mallat. Rotation, scaling and deformation invariant scattering for texture discrimination. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1233–1240, 2013.

- [20] C. Szegedy, S. Ioffe, V. Vanhoucke, and A. A. Alemi. Inception-v4, inception-resnet and the impact of residual connections on learning. In *Thirty-First AAAI Conference on Artificial Intelligence*, 2017.
- [21] C. Szegedy, W. Liu, Y. Jia, P. Sermanet, S. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, and A. Rabinovich. Going deeper with convolutions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1–9, 2015.
- [22] C. Szegedy, V. Vanhoucke, S. Ioffe, J. Shlens, and Z. Wojna. Rethinking the inception architecture for computer vision. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2818–2826, 2016.
- [23] P. Wang, P. Chen, Y. Yuan, D. Liu, Z. Huang, X. Hou, and G. Cottrell. Understanding convolution for semantic segmentation. In *2018 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 1451–1460. IEEE, 2018.
- [24] M. Weiler, M. Geiger, M. Welling, W. Boomsma, and T. Cohen. 3d steerable cnns: Learning rotationally equivariant features in volumetric data. In *Advances in Neural Information Processing Systems*, pages 10381–10392, 2018.
- [25] M. Weiler, F. A. Hamprecht, and M. Storath. Learning steerable filters for rotation equivariant cnns. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 849–858, 2018.
- [26] D. E. Worrall, S. J. Garbin, D. Turmukhambetov, and G. J. Brostow. Harmonic networks: Deep translation and rotation equivariance. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5028–5037, 2017.
- [27] H. Xiao, K. Rasul, and R. Vollgraf. Fashion-mnist: a novel image dataset for benchmarking machine learning algorithms, 2017.
- [28] Y. Xu, T. Xiao, J. Zhang, K. Yang, and Z. Zhang. Scale-invariant convolutional neural networks. *arXiv preprint arXiv:1411.6369*, 2014.
- [29] F. Yu and V. Koltun. Multi-scale context aggregation by dilated convolutions. In *International Conference on Learning Representations*, 2016.
- [30] F. Yu, V. Koltun, and T. Funkhouser. Dilated residual networks. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, July 2017.
- [31] Y. Zhou, Q. Ye, Q. Qiu, and J. Jiao. Oriented response networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 519–528, 2017.

A Proofs

A.1 Proof of Theorem 1

Proof of Theorem 1. We note first that (3) holds true if and only if the following being valid for all $l \geq 1$,

$$T_{\beta,v}x^{(l)}[x^{(l-1)}] = x^{(l)}[T_{\beta,v}x^{(l-1)}], \quad (12)$$

where $T_{\beta,v}x^{(0)}$ is understood as $D_{\beta,v}x^{(0)}$. We also note that the layer-wise operations of a general feedforward neural network with an extra index $\alpha \in \mathcal{S}$ can be written as

$$x^{(1)}[x^{(0)}](u, \alpha, \lambda) = \sigma \left(\sum_{\lambda'} \int_{\mathbb{R}^2} x^{(0)}(u + u', \lambda') W^{(1)}(u', \lambda', u, \alpha, \lambda) du' + b^{(1)}(\lambda) \right), \quad (13)$$

and, for $l > 1$,

$$x^{(l)}[x^{(l-1)}](u, \alpha, \lambda) = \sigma \left(\sum_{\lambda'} \int_{\mathbb{R}^2} \int_{\mathbb{R}} x^{(l-1)}(u + u', \alpha + \alpha', \lambda') W^{(l)}(u', \alpha', \lambda', u, \alpha, \lambda) d\alpha' du' + b^{(l)}(\lambda) \right). \quad (14)$$

To prove the sufficient part: when $l = 1$, (1), (2), and (4) lead to

$$\begin{aligned} T_{\beta,v}x^{(1)}[x^{(0)}](u, \alpha, \lambda) &= x^{(1)}[x^{(0)}] \left(2^{-\beta}(u - v), \alpha - \beta, \lambda \right) \\ &= \sigma \left(\sum_{\lambda'} \int x^{(0)} \left(2^{-\beta}(u - v) + u', \lambda' \right) W_{\lambda',\lambda}^{(1)} \left(2^{-(\alpha-\beta)} u' \right) 2^{-2(\alpha-\beta)} du' + b^{(1)}(\lambda) \right) \\ &= \sigma \left(\sum_{\lambda'} \int x^{(0)} \left(2^{-\beta}(u - v + \tilde{u}), \lambda' \right) W_{\lambda',\lambda}^{(1)} \left(2^{-\alpha} \tilde{u} \right) 2^{-2\alpha} d\tilde{u} + b^{(1)}(\lambda) \right), \end{aligned}$$

and

$$\begin{aligned} &x^{(1)}[D_{\beta,v}x^{(0)}](u, \alpha, \lambda) \\ &= \sigma \left(\sum_{\lambda'} \int_{\mathbb{R}^2} D_{\beta,v}x^{(0)}(u + u', \lambda') W_{\lambda',\lambda}^{(1)} \left(2^{-\alpha} u' \right) 2^{-2\alpha} du' + b^{(1)}(\lambda) \right) \\ &= \sigma \left(\sum_{\lambda'} \int_{\mathbb{R}^2} x^{(0)}(2^{-\beta}(u + u' - v), \lambda') W_{\lambda',\lambda}^{(1)} \left(2^{-\alpha} u' \right) 2^{-2\alpha} du' + b^{(1)}(\lambda) \right). \end{aligned}$$

Hence $T_{\beta,v}x^{(1)}[x^{(0)}] = x^{(1)}[D_{\beta,v}x^{(0)}]$.

When $l > 1$, we have

$$T_{\beta,v}x^{(l)}[x^{(l-1)}](u, \alpha, \lambda) = x^{(l)}[x^{(l-1)}] \left(2^{-\beta}(u - v), \alpha - \beta, \lambda \right)$$

$$\begin{aligned}
&= \sigma \left(\sum_{\lambda'} \int \int x^{(l-1)} \left(2^{-\beta}(u-v) + u', \alpha - \beta + \alpha', \lambda' \right) \right. \\
&\quad \left. W_{\lambda', \lambda}^{(l)} \left(2^{-(\alpha-\beta)} u', \alpha' \right) 2^{-2(\alpha-\beta)} du' d\alpha' + b^{(l)}(\lambda) \right) \\
&= \sigma \left(\sum_{\lambda'} \int \int x^{(l-1)} \left(2^{-\beta}(u-v + \tilde{u}), \alpha - \beta + \alpha', \lambda' \right) W_{\lambda', \lambda}^{(l)} \left(2^{-\alpha} \tilde{u}, \alpha' \right) 2^{-2\alpha} d\tilde{u} d\alpha' + b^{(l)}(\lambda) \right)
\end{aligned}$$

and

$$\begin{aligned}
&x^{(l)} [T_{\beta, v} x^{(l-1)}](u, \alpha, \lambda) \\
&= \sigma \left(\sum_{\lambda'} \int \int T_{\beta, v} x^{(l-1)}(u + u', \alpha + \alpha', \lambda') W_{\lambda', \lambda}^{(l)} \left(2^{-\alpha} u', \alpha' \right) 2^{-2\alpha} d\alpha' du' + b^{(l)}(\lambda) \right) \\
&= \sigma \left(\sum_{\lambda'} \int \int x^{(l-1)}(2^{-\beta}(u + u' - v), \alpha + \alpha' - \beta, \lambda') W_{\lambda', \lambda}^{(l)} \left(2^{-\alpha} u', \alpha' \right) 2^{-2\alpha} d\alpha' du' + b^{(l)}(\lambda) \right)
\end{aligned}$$

Therefore $T_{\beta, v} x^{(l)}[x^{(l-1)}] = x^{(l)}[T_{\beta, v} x^{(l-1)}]$, $\forall l > 1$.

To prove the necessary part: when $l = 1$, we have

$$\begin{aligned}
&T_{\beta, v} x^{(1)}[x^{(0)}](u, \alpha, \lambda) = x^{(1)}[x^{(0)}] \left(2^{-\beta}(u-v), \alpha - \beta, \lambda \right) \\
&= \sigma \left(\sum_{\lambda'} \int x^{(0)} \left(2^{-\beta}(u-v) + u', \lambda' \right) W^{(1)} \left(u', \lambda', 2^{-\beta}(u-v), \alpha - \beta, \lambda \right) du' + b^{(1)}(\lambda) \right) \\
&= \sigma \left(\sum_{\lambda'} \int x^{(0)} \left(2^{-\beta}(u + u' - v), \lambda' \right) \right. \\
&\quad \left. W^{(1)} \left(2^{-\beta} u', \lambda', 2^{-\beta}(u-v), \alpha - \beta, \lambda \right) 2^{-2\beta} du' + b^{(1)}(\lambda) \right),
\end{aligned}$$

and

$$\begin{aligned}
&x^{(1)}[D_{\beta, v} x^{(0)}](u, \alpha, \lambda) = \sigma \left(\sum_{\lambda'} \int D_{\beta, v} x^{(0)}(u + u', \lambda') W^{(1)}(u', \lambda', u, \alpha, \lambda) du' + b^{(1)}(\lambda) \right) \\
&= \sigma \left(\sum_{\lambda'} \int x^{(0)} \left(2^{-\beta}(u + u' - v), \lambda' \right) W^{(1)}(u', \lambda', u, \alpha, \lambda) du' + b^{(1)}(\lambda) \right)
\end{aligned}$$

Hence for (12) to hold when $l = 1$, we need

$$W^{(1)}(u', \lambda', u, \alpha, \lambda) = W^{(1)} \left(2^{-\beta} u', \lambda', 2^{-\beta}(u-v), \alpha - \beta, \lambda \right) 2^{-2\beta}, \forall u, \alpha, \lambda, u', \lambda', v, \beta. \quad (15)$$

Keeping $u, \alpha, \lambda, u', \lambda', \beta$ fixed while changing v in (15), we obtain that $W^{(1)}(u', \lambda', u, \alpha, \lambda)$ does not depend on the third variable u . Thus $W^{(1)}(u', \lambda', u, \alpha, \lambda) = W^{(1)}(u', \lambda', 0, \alpha, \lambda)$, $\forall u$. Define $W_{\lambda', \lambda}^{(1)}(u')$ as

$$W_{\lambda', \lambda}^{(1)}(u') := W^{(1)}(u', \lambda', 0, 0, \lambda).$$

Then, for any given $u', \lambda', u, \alpha, \lambda$, setting $\beta = \alpha$ in (15) leads to

$$\begin{aligned} W^{(1)}(u', \lambda', u, \alpha, \lambda) &= W^{(1)}(2^{-\alpha}u', \lambda', 2^{-\alpha}(u - v), 0, \lambda) 2^{-2\alpha} \\ &= W^{(1)}(2^{-\alpha}u', \lambda', 0, 0, \lambda) 2^{-2\alpha} = W_{\lambda', \lambda}^{(1)}(2^{-\alpha}u') 2^{-2\alpha}. \end{aligned}$$

Hence (13) can be written as (4).

For $l > 1$, a similar argument leads to

$$W^{(l)}(u', \alpha', \lambda', u, \alpha, \lambda) = W^{(l)}(2^{-\beta}u', \alpha', \lambda', 2^{-\beta}(u - v), \alpha - \beta, \lambda) 2^{-2\beta}, \forall u, \alpha, \lambda, u', \alpha', \lambda', v, \beta. \quad (16)$$

Again, keeping $u, \alpha, \lambda, u', \alpha', \lambda', \beta$ fixed while changing v in (16) leads us to the conclusion that $W^{(l)}(u', \alpha', \lambda', u, \alpha, \lambda)$ does not depend on the fourth variable u . Define

$$W_{\lambda', \lambda}^{(l)}(u', \alpha') := W^{(l)}(u', \alpha', \lambda', 0, 0, \lambda).$$

After setting $\beta = \alpha$ in (16), for any given $u', \alpha', \lambda', u, \alpha, \lambda$, we have

$$\begin{aligned} W^{(l)}(u', \alpha', \lambda', u, \alpha, \lambda) &= W^{(l)}(2^{-\alpha}u', \alpha', \lambda', 2^{-\alpha}(u - v), 0, \lambda) 2^{-2\alpha} \\ &= W^{(l)}(2^{-\alpha}u', \alpha', \lambda', 0, 0, \lambda) 2^{-2\alpha} = W_{\lambda', \lambda}^{(l)}(2^{-\alpha}u') 2^{-2\alpha}. \end{aligned}$$

This concludes the proof of the Theorem. \square

A.2 Proof of Theorem 2

Proof of Theorem 2. In a regular scale-equivariant CNN, the l -th convolutional layer (5) is computed as follows:

$$y(u, \alpha, \alpha', \lambda, \lambda') = \int_{\mathbb{R}^2} x^{(l-1)}(u + u', \alpha + \alpha', \lambda') W_{\lambda', \lambda}^{(l)}(2^{-\alpha}u', \alpha') 2^{-2\alpha} du', \quad (17)$$

$$z(u, \alpha, \lambda, \lambda') = \int_{\mathbb{R}} y(u, \alpha, \alpha', \lambda, \lambda') d\alpha', \quad (18)$$

$$x^{(l)}(u, \alpha, \lambda) = \sigma \left(\sum_{\lambda'=1}^{M_{l-1}} z(u, \alpha, \lambda, \lambda') + b^{(l)}(\lambda) \right). \quad (19)$$

The spatial convolutions in (17) take $2W^2L^2N_\alpha L_\alpha M_l M_{l-1}$ flops (there are $N_\alpha L_\alpha M_l M_{l-1}$ convolutions in u , each taking $2W^2L^2$ flops.) The summation over α' in (18) requires $L_\alpha N_\alpha W^2 M_l M_{l-1}$ flops. The summation over λ' , adding the bias, and applying the nonlinear activation in (19) requires an additional $W^2 N_\alpha M_l (2 + M_{l-1})$ flops. Thus the total number of floating point computations in a forward pass through the l -th layer of a regular scale-equivariant CNN is

$$W^2 N_\alpha M_l (2L^2 L_\alpha M_{l-1} + L_\alpha M_{l-1} + M_{l-1} + 2). \quad (20)$$

On the other hand, in an ScDCFNet with separable basis truncation up to KK_α leading coefficients, (5) can be computed via the following steps:

$$y(u, \alpha, \lambda', m) = \int_{\mathbb{R}} x^{(l-1)}(u, \alpha + \alpha', \lambda') \varphi_m(\alpha') d\alpha' \quad (21)$$

$$z(u, \alpha, \lambda', k, m) = \int_{\mathbb{R}^2} y(u + u', \alpha, \lambda', m) \psi_{j_l, k}(2^{-\alpha} u') 2^{-2\alpha} du' \quad (22)$$

$$x^{(l)}(u, \alpha, \lambda) = \sigma \left(\sum_{m=1}^{K_\alpha} \sum_{k=1}^K \sum_{\lambda'=1}^{M_{l-1}} z(u, \alpha, \lambda', k, m) a_{\lambda', \lambda}^{(l)}(k, m) + b^{(l)}(\lambda) \right). \quad (23)$$

The convolutions in α (21) require $2L_\alpha N_\alpha W^2 K_\alpha M_{l-1}$ flops (there are $W^2 K_\alpha M_{l-1}$ convolutions in α , each taking $2L_\alpha N_\alpha$ flops.) The spatial convolutions in (22) take $2W^2 L^2 N_\alpha M_{l-1} K_\alpha K$ flops ($N_\alpha M_{l-1} K_\alpha K$ convolutions in u , each taking $2W^2 L^2$ flops.) The last step (23) requires an additional $2W^2 N_\alpha M_l (1 + K K_\alpha M_{l-1})$ flops. Hence the total number of floating point computation for an ScDCFNet is

$$2W^2 N_\alpha (K K_\alpha M_{l-1} M_l + M_l + L^2 M_{l-1} K_\alpha K + L_\alpha K_\alpha M_{l-1}). \quad (24)$$

In particular, when $M_l \gg L^2, L_\alpha$, the dominating terms in (20) and (24) are, respectively, $2W^2 N_\alpha M_l M_{l-1} L^2 L_\alpha$ and $2W^2 N_\alpha M_{l-1} M_l K K_\alpha$. Thus the computational cost in an ScDCFNet has been reduced to a factor of $\frac{KK_\alpha}{L^2 L_\alpha}$. \square

A.3 Proof of Proposition 1

Before proving Proposition 1, we need the following two lemmas.

Lemma 1. *Suppose that $\{\psi_k\}_k$ are the FB bases, and $F(u) = \sum_k a(k) \psi_{j, k}(u) = \sum_k a(k) 2^{-2j} \psi_k(2^{-j} u)$ is a smooth function on $2^j \overline{B(0, 1)}$, then*

$$\int |F(u)| du, \quad \int |u| |\nabla F(u)| du, \quad 2^j \int |\nabla F(u)| du \leq \pi \|a\|_{FB} = \pi \left(\sum_k \mu_k \cdot a(k)^2 \right)^{1/2}. \quad (25)$$

This is Lemma 3.5 and Proposition 3.6 in [16] after rescaling u . Lemma 1 easily leads to the following lemma.

Lemma 2. *Let $a_{\lambda', \lambda}^{(l)}(k, m)$ be the coefficients of the filter $W_{\lambda', \lambda}^{(l)}(u, \alpha)$ under the joint bases $\{\psi_k\}_k$ and $\{\varphi_m\}_m$ defined in (7), and define $W_{\lambda', \lambda, m}^{(l)}(u)$ as*

$$W_{\lambda', \lambda, m}^{(l)}(u) := \sum_k a_{\lambda', \lambda}^{(l)}(k, m) \psi_{j_l, k}(u). \quad (26)$$

We have

$$B_{\lambda', \lambda}^{(1)}, C_{\lambda', \lambda}^{(1)}, 2^{j_1} D_{\lambda', \lambda}^{(1)} \leq \pi \|a_{\lambda', \lambda}^{(1)}\|_{FB}, \quad B_{\lambda', \lambda, m}^{(l)}, C_{\lambda', \lambda, m}^{(l)}, 2^{j_l} D_{\lambda', \lambda, m}^{(l)} \leq \pi \|a_{\lambda', \lambda}^{(l)}(\cdot, m)\|_{FB}, \quad \forall l > 1,$$

where

$$\begin{cases} B_{\lambda',\lambda}^{(1)} := \int |W_{\lambda',\lambda}^{(1)}(u)| du, & B_{\lambda',\lambda,m}^{(l)} := \int |W_{\lambda',\lambda,m}^{(l)}(u)| du, \quad l > 1, \\ C_{\lambda',\lambda}^{(1)} := \int |u| |\nabla_u W_{\lambda',\lambda}^{(1)}(u)| du, & C_{\lambda',\lambda,m}^{(l)} := \int |u| |\nabla_u W_{\lambda',\lambda,m}^{(l)}(u)| du, \quad l > 1, \\ D_{\lambda',\lambda}^{(1)} := \int |\nabla_u W_{\lambda',\lambda}^{(1)}(u)| du, & D_{\lambda',\lambda,m}^{(l)} := \int |\nabla_u W_{\lambda',\lambda,m}^{(l)}(u)| du, \quad l > 1. \end{cases} \quad (27)$$

We thus have

$$B_l, C_l, 2^{jl} D_l \leq A_l,$$

where

$$\begin{aligned} B_1 &:= \max \left\{ \sup_{\lambda} \sum_{\lambda'=1}^{M_0} B_{\lambda',\lambda}^{(1)}, \frac{M_0}{M_1} \sup_{\lambda'} \sum_{\lambda=1}^{M_1} B_{\lambda',\lambda}^{(1)} \right\}, \\ C_1 &:= \max \left\{ \sup_{\lambda} \sum_{\lambda'=1}^{M_0} C_{\lambda',\lambda}^{(1)}, \frac{M_0}{M_1} \sup_{\lambda'} \sum_{\lambda=1}^{M_1} C_{\lambda',\lambda}^{(1)} \right\}, \\ D_1 &:= \max \left\{ \sup_{\lambda} \sum_{\lambda'=1}^{M_0} D_{\lambda',\lambda}^{(1)}, \frac{M_0}{M_1} \sup_{\lambda'} \sum_{\lambda=1}^{M_1} D_{\lambda',\lambda}^{(1)} \right\}, \end{aligned} \quad (28)$$

and, for $l > 1$,

$$\begin{aligned} B_l &:= \max \left\{ \sup_{\lambda} \sum_{\lambda'=1}^{M_{l-1}} \sum_m B_{\lambda',\lambda,m}^{(l)}, \frac{2M_{l-1}}{M_l} \sum_m B_{l,m} \right\}, \quad B_{l,m} := \sup_{\lambda'} \sum_{\lambda=1}^{M_l} B_{\lambda',\lambda,m}^{(l)}, \\ C_l &:= \max \left\{ \sup_{\lambda} \sum_{\lambda'=1}^{M_{l-1}} \sum_m C_{\lambda',\lambda,m}^{(l)}, \frac{2M_{l-1}}{M_l} \sum_m C_{l,m} \right\}, \quad C_{l,m} := \sup_{\lambda'} \sum_{\lambda=1}^{M_l} C_{\lambda',\lambda,m}^{(l)}, \\ D_l &:= \max \left\{ \sup_{\lambda} \sum_{\lambda'=1}^{M_{l-1}} \sum_m D_{\lambda',\lambda,m}^{(l)}, \frac{2M_{l-1}}{M_l} \sum_m D_{l,m} \right\}, \quad D_{l,m} := \sup_{\lambda'} \sum_{\lambda=1}^{M_l} D_{\lambda',\lambda,m}^{(l)}. \end{aligned} \quad (29)$$

In particular, (A2) implies that $B_l, C_l, 2^{jl} D_l \leq 1, \forall l$.

Proof of Proposition 1. To simplify the notation, we omit (l) in $W_{\lambda',\lambda}^{(l)}$, $W_{\lambda',\lambda,m}^{(l)}$, and $b^{(l)}$, and let $M = M_l$, $M' = M_{l-1}$. The proof of (a) for the case $l = 1$ is similar to Proposition 3.1(a) of [16] after noticing the fact that

$$\int_{\mathbb{R}^2} |W(2^{-\alpha}u)| 2^{-2\alpha} du = \int_{\mathbb{R}^2} |W(u)| du, \quad (30)$$

and we include it here for completeness. From the definition of B_1 in (28), we have

$$\sup_{\lambda} \sum_{\lambda'} B_{\lambda',\lambda}^{(1)} \leq B_1, \quad \text{and} \quad \sup_{\lambda'} \sum_{\lambda} B_{\lambda',\lambda}^{(1)} \leq B_1 \frac{M}{M'}.$$

Thus, given two arbitrary functions x_1 and x_2 , we have

$$\begin{aligned}
& \left| \left(x^{(1)}[x_1] - x^{(1)}[x_2] \right) (u, \alpha, \lambda) \right|^2 \\
&= \left| \sigma \left(\sum_{\lambda'} \int x_1(u + u', \lambda') W_{\lambda', \lambda} (2^{-\alpha} u') 2^{-2\alpha} du' + b(\lambda) \right) \right. \\
&\quad \left. - \sigma \left(\sum_{\lambda'} \int x_2(u + u', \lambda') W_{\lambda', \lambda} (2^{-\alpha} u') 2^{-2\alpha} du' + b(\lambda) \right) \right|^2 \\
&\leq \left| \sum_{\lambda'} \int x_1(u + u', \lambda') W_{\lambda', \lambda} (2^{-\alpha} u') 2^{-2\alpha} du' - \sum_{\lambda'} \int x_2(u + u', \lambda') W_{\lambda', \lambda} (2^{-\alpha} u') 2^{-2\alpha} du' \right|^2 \\
&= \left| \sum_{\lambda'} \int (x_1 - x_2)(u + u', \lambda') W_{\lambda', \lambda} (2^{-\alpha} u') 2^{-2\alpha} du' \right|^2 \\
&\leq \left(\sum_{\lambda'} \int |(x_1 - x_2)(u + u', \lambda')|^2 |W_{\lambda', \lambda}(2^{-\alpha} u')| 2^{-2\alpha} du' \right) \left(\sum_{\lambda'} \int |W_{\lambda', \lambda}(2^{-\alpha} u')| 2^{-2\alpha} du' \right) \\
&= \left(\sum_{\lambda'} \int |(x_1 - x_2)(u + u', \lambda')|^2 |W_{\lambda', \lambda}(2^{-\alpha} u')| 2^{-2\alpha} du' \right) \left(\sum_{\lambda'} B_{\lambda', \lambda}^{(1)} \right) \\
&\leq B_1 \sum_{\lambda'} \int |(x_1 - x_2)(\tilde{u}, \lambda')|^2 |W_{\lambda', \lambda}(2^{-\alpha} (\tilde{u} - u))| 2^{-2\alpha} d\tilde{u}
\end{aligned}$$

Therefore, for any α ,

$$\begin{aligned}
& \sum_{\lambda} \int \left| \left(x^{(1)}[x_1] - x^{(1)}[x_2] \right) (u, \alpha, \lambda) \right|^2 du \\
&\leq \sum_{\lambda} \int B_1 \sum_{\lambda'} \int |(x_1 - x_2)(\tilde{u}, \lambda')|^2 |W_{\lambda', \lambda}(2^{-\alpha} (\tilde{u} - u))| 2^{-2\alpha} d\tilde{u} du \\
&= B_1 \sum_{\lambda'} \int |(x_1 - x_2)(\tilde{u}, \lambda')|^2 \left(\sum_{\lambda} \int |W_{\lambda', \lambda}(2^{-\alpha} (\tilde{u} - u))| 2^{-2\alpha} du \right) d\tilde{u} \\
&= B_1 \sum_{\lambda'} \int |(x_1 - x_2)(\tilde{u}, \lambda')|^2 \left(\sum_{\lambda} B_{\lambda', \lambda}^{(1)} \right) d\tilde{u} \\
&\leq B_1^2 \frac{M}{M'} \sum_{\lambda'} \int |(x_1 - x_2)(\tilde{u}, \lambda')|^2 d\tilde{u} \\
&= B_1^2 M \|x_1 - x_2\|^2 \\
&\leq M \|x_1 - x_2\|^2,
\end{aligned}$$

where the last inequality makes use of the fact that $B_1 \leq A_1 \leq 1$ under (A2) (Lemma 2.) Therefore

$$\|x^{(1)}[x_1] - x^{(1)}[x_2]\|^2 = \sup_{\alpha} \frac{1}{M} \sum_{\lambda} \int \left| \left(x^{(1)}[x_1] - x^{(1)}[x_2] \right) (u, \alpha, \lambda) \right|^2 du \leq \|x_1 - x_2\|^2.$$

This concludes the proof of (a) for the case $l = 1$. To prove the case for any $l > 1$, we first recall from (29) that

$$\sup_{\lambda} \sum_{\lambda'} \sum_m B_{\lambda', \lambda, m}^{(l)} \leq B_l, \quad \text{and} \quad \sum_m B_{l, m} \leq B_l \frac{M}{2M'}, \quad \text{where} \quad B_{l, m} = \sup_{\lambda'} \sum_{\lambda} B_{\lambda', \lambda, m}^{(l)}.$$

Thus, for two arbitrary functions x_1 and x_2 , we have

$$\begin{aligned} & \left| \left(x^{(l)}[x_1] - x^{(l)}[x_2] \right) (u, \alpha, \lambda) \right|^2 \\ &= \left| \sigma \left(\sum_{\lambda'} \int_{\mathbb{R}^2} \int_{\mathbb{R}} x_1(u + u', \alpha + \alpha', \lambda') W_{\lambda', \lambda} (2^{-\alpha} u', \alpha') 2^{-2\alpha} d\alpha' du' + b(\lambda) \right) \right. \\ & \quad \left. - \sigma \left(\sum_{\lambda'} \int_{\mathbb{R}^2} \int_{\mathbb{R}} x_2(u + u', \alpha + \alpha', \lambda') W_{\lambda', \lambda} (2^{-\alpha} u', \alpha') 2^{-2\alpha} d\alpha' du' + b(\lambda) \right) \right|^2 \\ &\leq \left| \sum_{\lambda'} \int_{\mathbb{R}^2} \int_{\mathbb{R}} (x_1 - x_2)(u + u', \alpha + \alpha', \lambda') W_{\lambda', \lambda} (2^{-\alpha} u', \alpha') 2^{-2\alpha} d\alpha' du' \right|^2 \\ &= \left| \sum_{\lambda'} \int_{\mathbb{R}^2} \int_{\mathbb{R}} (x_1 - x_2)(u + u', \alpha + \alpha', \lambda') 2^{-2\alpha} \sum_m W_{\lambda', \lambda, m} (2^{-\alpha} u') \varphi_m(\alpha') d\alpha' du' \right|^2 \\ &= \left| \sum_{\lambda'} \sum_m \int_{\mathbb{R}^2} G_m(u + u', \alpha, \lambda') 2^{-2\alpha} W_{\lambda', \lambda, m} (2^{-\alpha} u') du' \right|^2 \\ &\leq \left(\sum_{\lambda'} \sum_m \int_{\mathbb{R}^2} |G_m(u + u', \alpha, \lambda')|^2 |W_{\lambda', \lambda, m} (2^{-\alpha} u')| 2^{-2\alpha} du' \right) \\ & \quad \left(\sum_{\lambda'} \sum_m \int_{\mathbb{R}^2} |W_{\lambda', \lambda, m} (2^{-\alpha} u')| 2^{-2\alpha} du' \right) \\ &= \left(\sum_{\lambda'} \sum_m \int_{\mathbb{R}^2} |G_m(\tilde{u}, \alpha, \lambda')|^2 |W_{\lambda', \lambda, m} (2^{-\alpha} (\tilde{u} - u))| 2^{-2\alpha} d\tilde{u} \right) \left(\sum_{\lambda'} \sum_m B_{\lambda', \lambda, m}^{(l)} \right) \\ &\leq B_l \sum_{\lambda'} \sum_m \int_{\mathbb{R}^2} |G_m(\tilde{u}, \alpha, \lambda')|^2 |W_{\lambda', \lambda, m} (2^{-\alpha} (\tilde{u} - u))| 2^{-2\alpha} d\tilde{u}, \end{aligned}$$

where

$$G_m(u, \alpha, \lambda') := \int_{\mathbb{R}} (x_1 - x_2)(u, \alpha + \alpha', \lambda') \varphi_m(\alpha') d\alpha'. \quad (31)$$

We claim (to be proved later in Lemma 3) that

$$M' \|G_m\|^2 = \sup_{\alpha} \sum_{\lambda'} \int_{\mathbb{R}^2} |G_m(u, \alpha, \lambda')|^2 du \leq 2M' \|x_1 - x_2\|^2, \quad \forall m. \quad (32)$$

Thus, for any α ,

$$\begin{aligned}
& \sum_{\lambda} \int_{\mathbb{R}^2} \left| \left(x^{(l)}[x_1] - x^{(l)}[x_2] \right) (u, \alpha, \lambda) \right|^2 du \\
& \leq \sum_{\lambda} \int_{\mathbb{R}^2} B_l \sum_{\lambda'} \sum_m \int_{\mathbb{R}^2} |G_m(\tilde{u}, \alpha, \lambda')|^2 |W_{\lambda', \lambda, m}(2^{-\alpha}(\tilde{u} - u))| 2^{-2\alpha} d\tilde{u} du \\
& = B_l \sum_{\lambda'} \sum_m \int_{\mathbb{R}^2} |G_m(\tilde{u}, \alpha, \lambda')|^2 \left(\sum_{\lambda} \int_{\mathbb{R}^2} |W_{\lambda', \lambda, m}(2^{-\alpha}(\tilde{u} - u))| 2^{-2\alpha} du \right) d\tilde{u} \\
& = B_l \sum_{\lambda'} \sum_m \int_{\mathbb{R}^2} |G_m(\tilde{u}, \alpha, \lambda')|^2 \left(\sum_{\lambda} B_{\lambda', \lambda, m}^{(l)} \right) d\tilde{u} \\
& \leq B_l \sum_{\lambda'} \sum_m \int_{\mathbb{R}^2} |G_m(\tilde{u}, \alpha, \lambda')|^2 B_{l, m} d\tilde{u} \\
& = B_l \sum_m \left(\sum_{\lambda'} \int_{\mathbb{R}^2} |G_m(\tilde{u}, \alpha, \lambda')|^2 d\tilde{u} \right) B_{l, m} \\
& \leq B_l \cdot 2M' \|x_1 - x_2\|^2 \sum_m B_{l, m} \\
& \leq B_l^2 \cdot 2M' \|x_1 - x_2\|^2 \frac{M}{2M'} \leq M \|x_1 - x_2\|^2.
\end{aligned}$$

Therefore

$$\|x^{(l)}[x_1] - x^{(l)}[x_2]\|^2 = \sup_{\alpha} \frac{1}{M} \sum_{\lambda} \int_{\mathbb{R}^2} \left| \left(x^{(l)}[x_1] - x^{(l)}[x_2] \right) (u, \alpha, \lambda) \right|^2 du \leq \|x_1 - x_2\|^2.$$

To prove (b), we use the method of induction. When $l = 0$, $x_0^{(0)}(u, \lambda) = 0$ by definition. When $l = 1$, $x_0^{(1)}(u, \alpha, \lambda) = \sigma(b^{(1)}(\lambda))$. Suppose $x_0^{(l-1)}(u, \alpha, \lambda) = x_0^{(l-1)}(\lambda)$ for some $l > 1$, then

$$\begin{aligned}
x_0^{(l)}(u, \alpha, \lambda) &= \sigma \left(\sum_{\lambda'} \int_{\mathbb{R}^2} \int_{\mathbb{R}} x_0^{(l-1)}(u + u', \alpha + \alpha', \lambda') W_{\lambda', \lambda}^{(l)}(2^{-\alpha} u', \alpha') 2^{-2\alpha} d\alpha' du' + b^{(l)}(\lambda) \right) \\
&= \sigma \left(\sum_{\lambda'} x_0^{(l-1)}(\lambda') \int_{\mathbb{R}^2} \int_{\mathbb{R}} W_{\lambda', \lambda}^{(l)}(2^{-\alpha} u', \alpha') 2^{-2\alpha} d\alpha' du' + b^{(l)}(\lambda) \right) \\
&= \sigma \left(\sum_{\lambda'} x_0^{(l-1)}(\lambda') \int_{\mathbb{R}^2} \int_{\mathbb{R}} W_{\lambda', \lambda}^{(l)}(u', \alpha') d\alpha' du' + b^{(l)}(\lambda) \right) \\
&= x_0^{(l)}(\lambda).
\end{aligned}$$

Part (c) is an easy corollary of part (a). More specifically, for any $l > 1$,

$$\|x_c^{(l)}\| = \|x^{(l)} - x_0^{(l)}\| = \|x^{(l)}[x^{(l-1)}] - x_0^{(l)}[x_0^{(l-1)}]\| \leq \|x^{(l-1)} - x_0^{(l-1)}\| = \|x_c^{(l-1)}\|.$$

□

Lemma 3. Suppose $\varphi \in L^2(\mathbb{R})$ with $\text{supp}(\varphi_m) \subset [-1, 1]$ and $\|\varphi\|_{L^2} = 1$, and x is a function of three variables

$$\begin{aligned} x : \mathbb{R}^2 \times \mathbb{R} \times [M] &\rightarrow \mathbb{R} \\ (u, \alpha, \lambda) &\mapsto x(u, \alpha, \lambda) \end{aligned}$$

with $\|x\|^2 := \sup_{\alpha} \frac{1}{M} \sum_{\lambda} \int_{\mathbb{R}^2} |x(u, \alpha, \lambda)|^2 du$. Define $G(u, \alpha, \lambda)$ as

$$G(u, \alpha, \lambda) := \int_{\mathbb{R}} x(u, \alpha + \alpha', \lambda) \varphi(\alpha') d\alpha'.$$

Then we have

$$M\|G\|^2 = \sup_{\alpha} \sum_{\lambda} \int_{\mathbb{R}^2} |G(u, \alpha, \lambda)|^2 du \leq 2M\|x\|^2. \quad (33)$$

Proof of Lemma 3. Notice that, for any α , we have

$$\begin{aligned} \sum_{\lambda} \int_{\mathbb{R}^2} |G(u, \alpha, \lambda)|^2 du &= \sum_{\lambda} \int_{\mathbb{R}^2} \left| \int_{-1}^1 x(u, \alpha + \alpha', \lambda) \varphi(\alpha') d\alpha' \right|^2 du \\ &\leq \sum_{\lambda} \int_{\mathbb{R}^2} \left(\int_{-1}^1 |x(u, \alpha + \alpha', \lambda)|^2 d\alpha' \right) \|\varphi\|_{L^2}^2 du \\ &= \int_{-1}^1 \left(\sum_{\lambda} \int_{\mathbb{R}^2} |x(u, \alpha + \alpha', \lambda)|^2 du \right) d\alpha' \\ &\leq \int_{-1}^1 M\|x\|^2 d\alpha' = 2M\|x\|^2. \end{aligned}$$

Thus

$$\sup_{\alpha} \sum_{\lambda} \int_{\mathbb{R}^2} |G(u, \alpha, \lambda)|^2 du \leq 2M\|x\|^2.$$

□

A.4 Proof of Theorem 3

To prove Theorem 3, we need the following two Propositions.

Proposition 2. In an ScDCFNet satisfying (A1) and (A3), we have

1. For any $l \geq 1$,

$$\left\| x^{(l)}[D_{\tau}x^{(l-1)}] - D_{\tau}x^{(l)}[x^{(l-1)}] \right\| \leq 4(B_l + C_l)|\nabla\tau|_{\infty}\|x_c^{(l-1)}\|. \quad (34)$$

2. For any $l \geq 1$, we have

$$\|T_{\beta,v}x^{(l)}\| = 2^\beta \|x^{(l)}\|, \quad (35)$$

and

$$\left\| x^{(l)}[T_{\beta,v} \circ D_\tau x^{(l-1)}] - T_{\beta,v} D_\tau x^{(l)}[x^{(l-1)}] \right\| \leq 2^{\beta+2} (B_l + C_l) |\nabla \tau|_\infty \|x_c^{(l-1)}\|, \quad (36)$$

where the first $T_{\beta,v}$ in (36) is replaced by $D_{\beta,v}$ when $l = 1$.

3. If (A2) also holds true, then

$$\left\| x^{(l)}[D_{\beta,v} \circ D_\tau x^{(0)}] - T_{\beta,v} D_\tau x^{(l)}[x^{(0)}] \right\| \leq 2^{\beta+3} l |\nabla \tau|_\infty \|x^{(0)}\|, \quad \forall l \geq 1. \quad (37)$$

Proposition 3. In an ScDCFNet satisfying (A1) and (A3), we have, for any $l \geq 1$,

$$\left\| T_{\beta,v} D_\tau x^{(l)} - T_{\beta,v} x^{(l)} \right\| \leq 2^{\beta+1} |\tau|_\infty D_l \|x_c^{(l-1)}\| \leq 2^{\beta+1} |\tau|_\infty D_l \|x^{(0)}\|. \quad (38)$$

If (A2) also holds true, then

$$\left\| T_{\beta,v} D_\tau x^{(l)} - T_{\beta,v} x^{(l)} \right\| \leq 2^{\beta+1-j_l} |\tau|_\infty \|x^{(0)}\|. \quad (39)$$

Proof of Theorem 3. Putting together (37) and (39), we have

$$\begin{aligned} & \left\| x^{(L)}[D_{\beta,v} \circ D_\tau x^{(0)}] - T_{\beta,v} x^{(L)}[x^{(0)}] \right\| \\ & \leq \left\| x^{(L)}[D_{\beta,v} \circ D_\tau x^{(0)}] - T_{\beta,v} D_\tau x^{(L)}[x^{(0)}] \right\| + \left\| T_{\beta,v} D_\tau x^{(L)}[x^{(0)}] - T_{\beta,v} x^{(L)}[x^{(0)}] \right\| \\ & \leq 2^{\beta+3} L |\nabla \tau|_\infty \|x^{(0)}\| + 2^{\beta+1-j_L} |\tau|_\infty \|x^{(0)}\| \\ & = 2^{\beta+1} (4L |\nabla \tau|_\infty + 2^{-j_L} |\tau|_\infty) \|x^{(0)}\| \end{aligned}$$

This concludes the proof of Theorem 3. \square

Finally, we need to prove Proposition 2 and Proposition 3, where the following lemma from [16] is useful.

Lemma 4 (Lemma A.1 of [16]). Suppose that $|\nabla \tau|_\infty < 1/5$, $\rho(u) = u - \tau(u)$, then at every point $u \in \mathbb{R}^2$,

$$||J\rho| - 1| \leq |\nabla \tau|_\infty (2 + |\nabla \tau|_\infty), \quad (40)$$

where $J\rho$ is the Jacobian of ρ , and $|J\rho|$ is the Jacobian determinant. As a result,

$$||J\rho| - 1|, ||J\rho^{-1}| - 1| \leq 4|\nabla \tau|_\infty, \quad (41)$$

and,

$$|J\rho|, |J\rho^{-1}| \leq 2. \quad (42)$$

Proof of Proposition 2. Just like Proposition 1(a), the proof of Proposition 2(a) for the case $l = 1$ is similar to Lemma 3.2 of [16] after the change of variable (30). We thus focus only on the proof for the case $l > 1$. To simplify the notation, we denote $x^{(l)}[x^{(l-1)}]$ as $y[x]$, and replace $x_c^{(l-1)}$, $W^{(l)}$, $b^{(l)}$, M_{l-1} , and M_l , respectively, by x_c , W , b , M' , and M . By the definition of the deformation D_τ (10), we have

$$\begin{aligned} D_\tau y[x](u, \alpha, \lambda) &= \sigma \left(\sum_{\lambda'} \int_{\mathbb{R}^2} \int_{\mathbb{R}} x(\rho(u) + u', \alpha + \alpha', \lambda') W_{\lambda', \lambda} (2^{-\alpha} u', \alpha') 2^{-2\alpha} d\alpha' du' + b(\lambda) \right), \\ y[D_\tau x](u, \alpha, \lambda) &= \sigma \left(\sum_{\lambda'} \int_{\mathbb{R}^2} \int_{\mathbb{R}} x(\rho(u + u'), \alpha + \alpha', \lambda') W_{\lambda', \lambda} (2^{-\alpha} u', \alpha') 2^{-2\alpha} d\alpha' du' + b(\lambda) \right). \end{aligned}$$

Thus

$$\begin{aligned} & |(D_\tau y[x] - y[D_\tau x])(u, \alpha, \lambda)|^2 \\ &= \left| \sigma \left(\sum_{\lambda'} \int_{\mathbb{R}^2} \int_{\mathbb{R}} x(\rho(u) + u', \alpha + \alpha', \lambda') W_{\lambda', \lambda} (2^{-\alpha} u', \alpha') 2^{-2\alpha} d\alpha' du' + b(\lambda) \right) \right. \\ &\quad \left. - \sigma \left(\sum_{\lambda'} \int_{\mathbb{R}^2} \int_{\mathbb{R}} x(\rho(u + u'), \alpha + \alpha', \lambda') W_{\lambda', \lambda} (2^{-\alpha} u', \alpha') 2^{-2\alpha} d\alpha' du' + b(\lambda) \right) \right|^2 \\ &\leq \left| \sum_{\lambda'} \int_{\mathbb{R}^2} \int_{\mathbb{R}} (x(\rho(u) + u', \alpha + \alpha', \lambda') - x(\rho(u + u'), \alpha + \alpha', \lambda')) W_{\lambda', \lambda} (2^{-\alpha} u', \alpha') 2^{-2\alpha} d\alpha' du' \right|^2 \\ &= \left| \sum_{\lambda'} \int_{\mathbb{R}^2} \int_{\mathbb{R}} (x_c(\rho(u) + u', \alpha + \alpha', \lambda') - x_c(\rho(u + u'), \alpha + \alpha', \lambda')) W_{\lambda', \lambda} (2^{-\alpha} u', \alpha') 2^{-2\alpha} d\alpha' du' \right|^2 \\ &= \left| \sum_{\lambda'} \sum_m \int_{\mathbb{R}^2} \int_{\mathbb{R}} (x_c(\rho(u) + u', \alpha + \alpha', \lambda') - x_c(\rho(u + u'), \alpha + \alpha', \lambda')) \cdot \right. \\ &\quad \left. W_{\lambda', \lambda, m} (2^{-\alpha} u') \varphi_m(\alpha') 2^{-2\alpha} d\alpha' du' \right|^2, \end{aligned}$$

where the second equality results from the fact that $x(u, \alpha, \lambda) - x_c(u, \alpha, \lambda) = x_0(\lambda)$ depends only on λ (Proposition 1(b).) Just like the proof of Proposition 1(a), we take the integral of α' first, and define

$$G_m(u, \alpha, \lambda') := \int_{\mathbb{R}} x_c(u, \alpha + \alpha', \lambda') \varphi_m(\alpha') d\alpha'. \quad (43)$$

Thus

$$\begin{aligned} & |(D_\tau y[x] - y[D_\tau x])(u, \alpha, \lambda)|^2 \\ &\leq \left| \sum_{\lambda'} \sum_m \int_{\mathbb{R}^2} (G_m(\rho(u) + u', \alpha, \lambda') - G_m(\rho(u + u'), \alpha, \lambda')) W_{\lambda', \lambda, m} (2^{-\alpha} u') 2^{-2\alpha} du' \right|^2 \\ &= \left| \sum_{\lambda'} \sum_m \int_{\mathbb{R}^2} G_m(v, \alpha, \lambda') W_{\lambda', \lambda, m} (2^{-\alpha} (v - \rho(u))) 2^{-2\alpha} dv \right|^2 \end{aligned}$$

$$\begin{aligned}
& - \sum_{\lambda'} \sum_m \int_{\mathbb{R}^2} G_m(v, \alpha, \lambda') W_{\lambda', \lambda, m} (2^{-\alpha}(\rho^{-1}(v) - u)) 2^{-2\alpha} |J\rho^{-1}(v)| dv \Big|^2 \\
& = |E_1(u, \alpha, \lambda) + E_2(u, \alpha, \lambda)|^2,
\end{aligned}$$

where

$$\begin{aligned}
E_1(u, \alpha, \lambda) &= \sum_{\lambda'} \sum_m \int_{\mathbb{R}^2} [W_{\lambda', \lambda, m} (2^{-\alpha}(v - \rho(u))) - W_{\lambda', \lambda, m} (2^{-\alpha}(\rho^{-1}(v) - u))] \cdot \\
& \quad 2^{-2\alpha} G_m(v, \alpha, \lambda') dv, \\
E_2(u, \alpha, \lambda) &= \sum_{\lambda'} \sum_m \int_{\mathbb{R}^2} W_{\lambda', \lambda, m} (2^{-\alpha}(\rho^{-1}(v) - u)) [1 - |J\rho^{-1}(v)|] \cdot \\
& \quad 2^{-2\alpha} G_m(v, \alpha, \lambda') dv.
\end{aligned}$$

Therefore

$$\begin{aligned}
M \|D_\tau y[x] - y[D_\tau x]\|^2 &= \sup_\alpha \sum_\lambda \int_{\mathbb{R}^2} |(D_\tau y[x] - y[D_\tau x])(u, \alpha, \lambda)|^2 du \\
&\leq \sup_\alpha \sum_\lambda \int_{\mathbb{R}^2} |E_1(u, \alpha, \lambda) + E_2(u, \alpha, \lambda)|^2 du \\
&= M \|E_1 + E_2\|^2
\end{aligned}$$

Hence

$$\|D_\tau y[x] - y[D_\tau x]\| \leq \|E_1 + E_2\|. \quad (44)$$

We thus seek to estimate $\|E_1\|$ and $\|E_2\|$ individually.

To bound $\|E_2\|$, we let

$$k_{\lambda', \lambda, m}^{(2)}(v, u, \alpha) := W_{\lambda', \lambda, m} (2^{-\alpha}(\rho^{-1}(v) - u)) [1 - |J\rho^{-1}(v)|] 2^{-2\alpha}.$$

Then

$$E_2(u, \alpha, \lambda) = \sum_{\lambda'} \sum_m \int_{\mathbb{R}^2} G_m(v, \alpha, \lambda') k_{\lambda', \lambda, m}^{(2)}(v, u, \alpha) dv,$$

and, for any given v and α

$$\begin{aligned}
\int_{\mathbb{R}^2} |k_{\lambda', \lambda, m}^{(2)}(v, u, \alpha)| du &= \int_{\mathbb{R}^2} |W_{\lambda', \lambda, m} (2^{-\alpha}(\rho^{-1}(v) - u))| [1 - |J\rho^{-1}(v)|] 2^{-2\alpha} du \\
&= [1 - |J\rho^{-1}(v)|] \int_{\mathbb{R}^2} |W_{\lambda', \lambda, m}(\tilde{u})| d\tilde{u} \\
&\leq 4|\nabla \tau|_\infty B_{\lambda', \lambda, m}^{(l)},
\end{aligned}$$

where the last inequality comes from (41). Moreover, for any given u and α ,

$$\begin{aligned} \int_{\mathbb{R}^2} \left| k_{\lambda', \lambda, m}^{(2)}(v, u, \alpha) \right| dv &= \int_{\mathbb{R}^2} |W_{\lambda', \lambda, m}(2^{-\alpha}(\rho^{-1}(v) - u))| |1 - |J\rho^{-1}(v)|| 2^{-2\alpha} dv \\ &= \int_{\mathbb{R}^2} |W_{\lambda', \lambda, m}(\tilde{v} - 2^{-\alpha}u)| \cdot ||J\rho(2^\alpha \tilde{v})| - 1| d\tilde{v} \\ &\leq 4|\nabla \tau|_\infty B_{\lambda', \lambda, m}^{(l)}, \end{aligned}$$

where the last inequality is again because of (41). Thus, for any given α ,

$$\begin{aligned} \sum_{\lambda} \int_{\mathbb{R}^2} |E_2(u, \alpha, \lambda)|^2 du &= \sum_{\lambda} \int_{\mathbb{R}^2} \left| \sum_{\lambda'} \sum_m \int_{\mathbb{R}^2} G_m(v, \alpha, \lambda') k_{\lambda', \lambda, m}^{(2)}(v, u, \alpha) dv \right|^2 du \\ &\leq \sum_{\lambda} \int_{\mathbb{R}^2} \left(\sum_{\lambda'} \sum_m \int_{\mathbb{R}^2} |G_m(v, \alpha, \lambda)|^2 \left| k_{\lambda', \lambda, m}^{(2)}(v, u, \alpha) \right| dv \right) \cdot \\ &\quad \left(\sum_{\lambda'} \sum_m \int_{\mathbb{R}^2} \left| k_{\lambda', \lambda, m}^{(2)}(v, u, \alpha) \right| dv \right) du \\ &\leq \sum_{\lambda} \int_{\mathbb{R}^2} \left(\sum_{\lambda'} \sum_m \int_{\mathbb{R}^2} |G_m(v, \alpha, \lambda)|^2 \left| k_{\lambda', \lambda, m}^{(2)}(v, u, \alpha) \right| dv \right) \left(\sum_{\lambda'} \sum_m 4|\nabla \tau|_\infty B_{\lambda', \lambda, m}^{(l)} \right) du \\ &\leq 4|\nabla \tau|_\infty B_l \sum_m \sum_{\lambda'} \int_{\mathbb{R}^2} |G_m(v, \alpha, \lambda)|^2 \left(\sum_{\lambda} \int_{\mathbb{R}^2} \left| k_{\lambda', \lambda, m}^{(2)}(v, u, \alpha) \right| du \right) dv \\ &\leq 4|\nabla \tau|_\infty B_l \sum_m \sum_{\lambda'} \int_{\mathbb{R}^2} |G_m(v, \alpha, \lambda)|^2 \left(\sum_{\lambda} 4|\nabla \tau|_\infty B_{\lambda', \lambda, m}^{(l)} \right) dv \\ &\leq 16|\nabla \tau|_\infty^2 B_l \sum_m \left(\sum_{\lambda'} \int_{\mathbb{R}^2} |G_m(v, \alpha, \lambda)|^2 dv \right) B_{l, m} \\ &\leq 16|\nabla \tau|_\infty^2 B_l \sum_m M' \|G_m\|^2 B_{l, m} \end{aligned}$$

Since $\|G_m\|^2 \leq 2\|x_c\|^2$ (by Lemma 3), and $\sum_m B_{l, m} \leq \frac{M}{2M'} B_l$ by definition (29), we thus have

$$\sum_{\lambda} \int_{\mathbb{R}^2} |E_2(u, \alpha, \lambda)|^2 du \leq 16|\nabla \tau|_\infty^2 B_l \frac{M}{2M'} B_l \cdot 2M' \|x_c\|^2 = M(4|\nabla \tau|_\infty B_l \|x_c\|)^2, \quad \forall \alpha. \quad (45)$$

Taking \sup_{α} on both sides gives us

$$\|E_2\| \leq 4|\nabla \tau|_\infty B_l \|x_c\|. \quad (46)$$

Similarly, to bound $\|E_1\|$, we introduce

$$k_{\lambda', \lambda, m}^{(1)}(v, u, \alpha) := [W_{\lambda', \lambda, m}(2^{-\alpha}(v - \rho(u))) - W_{\lambda', \lambda, m}(2^{-\alpha}(\rho^{-1}(v) - u))] 2^{-2\alpha}.$$

Then

$$E_1(u, \alpha, \lambda) = \sum_{\lambda'} \sum_m \int_{\mathbb{R}^2} G_m(v, \alpha, \lambda') k_{\lambda', \lambda, m}^{(1)}(v, u, \alpha) dv,$$

and, for any given v and α , we have

$$\int_{\mathbb{R}^2} \left| k_{\lambda', \lambda, m}^{(1)}(v, u, \alpha) \right| du, \quad \int_{\mathbb{R}^2} \left| k_{\lambda', \lambda, m}^{(1)}(v, u, \alpha) \right| dv \leq 4|\nabla \tau|_{\infty} C_{\lambda', \lambda, m}^{(l)}. \quad (47)$$

The proof of (47) is exactly the same as that of Lemma 3.2 in [16] after a change of variable, and we thus omit the detail. Similar to the procedure we take to bound $\|E_2\|$, (47) leads to

$$\|E_1\| \leq 4|\nabla \tau|_{\infty} C_l \|x_c\|. \quad (48)$$

Putting together (44), (46), and (48), we thus have

$$\|D_{\tau} y[x] - y[D_{\tau} x]\| \leq \|E_1 + E_2\| \leq \|E_1\| + \|E_2\| \leq 4(B_l + C_l) |\nabla \tau|_{\infty} \|x_c\|.$$

This concludes the proof of (a).

To prove (b), given any $\beta \in \mathbb{R}$, and $v \in \mathbb{R}^2$, we have

$$\begin{aligned} \|T_{\beta, v} x^{(l)}\|^2 &= \sup_{\alpha} \frac{1}{M_l} \sum_{\lambda} \int_{\mathbb{R}^2} \left| T_{\beta, v} x^{(l)}(u, \alpha, \lambda) \right|^2 du \\ &= \sup_{\alpha} \frac{1}{M_l} \sum_{\lambda} \int_{\mathbb{R}^2} \left| x^{(l)}(2^{-\beta}(u - v), \alpha - \beta, \lambda) \right|^2 du \\ &= \sup_{\alpha} \frac{1}{M_l} \sum_{\lambda} \int_{\mathbb{R}^2} \left| x^{(l)}(\tilde{u}, \alpha - \beta, \lambda) \right|^2 2^{2\beta} d\tilde{u} \\ &= 2^{2\beta} \|x^{(l)}\|^2 \end{aligned}$$

Thus (35) holds true. As for (36), we have

$$\begin{aligned} &\left\| x^{(l)} [T_{\beta, v} \circ D_{\tau} x^{(l-1)}] - T_{\beta, v} D_{\tau} x^{(l)} [x^{(l-1)}] \right\| \\ &= \left\| T_{\beta, v} x^{(l)} [D_{\tau} x^{(l-1)}] - T_{\beta, v} D_{\tau} x^{(l)} [x^{(l-1)}] \right\| \\ &= 2^{\beta} \left\| x^{(l)} [D_{\tau} x^{(l-1)}] - D_{\tau} x^{(l)} [x^{(l-1)}] \right\| \\ &\leq 2^{\beta+2} (B_l + C_l) |\nabla \tau|_{\infty} \|x_c^{(l-1)}\|, \end{aligned}$$

where the first equality holds valid because of Theorem 1, and the second equality comes from (35).

To prove (c), for any $0 \leq j \leq l$, define y_j as

$$y_j = x^{(l)} \circ x^{(l-1)} \circ \dots \circ T_{\beta, v} \circ D_{\tau} x^{(j)} \circ \dots \circ x^{(0)}.$$

We thus have

$$\begin{aligned}
& \left\| x^{(l)} [D_{\beta,v} \circ D_\tau x^{(0)}] - T_{\beta,v} D_\tau x^{(l)} [x^{(0)}] \right\| = \|y_l - y_0\| \leq \sum_{j=1}^l \|y_j - y_{j-1}\| \\
& = \sum_{j=1}^l \left\| x^{(l)} \circ \dots \circ T_{\beta,v} \circ D_\tau x^{(j)} \circ \dots \circ x^{(0)} - x^{(l)} \circ \dots \circ x^{(j)} \circ T_{\beta,v} \circ D_\tau x^{(j-1)} \circ \dots \circ x^{(0)} \right\| \\
& \leq \sum_{j=1}^l \left\| T_{\beta,v} \circ D_\tau x^{(j)} [x^{(j-1)}] - x^{(j)} [T_{\beta,v} \circ D_\tau x^{(j-1)}] \right\| \\
& \leq \sum_{j=1}^l 2^{\beta+2} (B_j + C_j) |\nabla \tau|_\infty \|x_c^{(j-1)}\| \\
& \leq \sum_{k=1}^l 2^{\beta+2} \cdot 2 |\nabla \tau|_\infty \|x^{(0)}\| = 2^{\beta+3} l |\nabla \tau|_\infty \|x^{(0)}\|,
\end{aligned}$$

where the second inequality is because of Proposition 1(a), the third inequality is due to (36), and the last inequality holds true because $B_l, C_l \leq A_l \leq 1$ under (A2) (Lemma 2.) This concludes the proof of Proposition 2. \square

Proof of Proposition 3. The second inequality in (38) is due to Proposition 1(c). Because of (35), the first inequality in (38) is equivalent to

$$\left\| D_\tau x^{(l)} - x^{(l)} \right\| \leq 2 |\tau|_\infty D_l \|x_c^{(l-1)}\| \quad (49)$$

Just like Proposition 2(a), the proof of (49) for the case $l = 1$ is similar to Proposition 3.4 of [16] after the change of variable (30). A similar strategy as that of Proposition 2(a) can be used to extend the proof to the case $l > 1$. More specifically, denote $x^{(l-1)}, x_c^{(l-1)}, W^{(l)}, b^{(l)}$, respectively, as x, x_c, W , and b to simplify the notation. We have

$$\begin{aligned}
& \left| \left(D_\tau x^{(l)} [x] - x^{(l)} [x] \right) (u, \alpha, \lambda) \right|^2 \\
& = \left| \sigma \left(\sum_{\lambda'} \int_{\mathbb{R}^2} \int_{\mathbb{R}} x(\rho(u) + u', \alpha + \alpha', \lambda') W_{\lambda', \lambda} (2^{-\alpha} u', \alpha') 2^{-2\alpha} du' d\alpha' + b(\lambda) \right) \right. \\
& \quad \left. - \sigma \left(\sum_{\lambda'} \int_{\mathbb{R}^2} \int_{\mathbb{R}} x(u + u', \alpha + \alpha', \lambda') W_{\lambda', \lambda} (2^{-\alpha} u', \alpha') 2^{-2\alpha} du' d\alpha' + b(\lambda) \right) \right| \\
& \leq \left| \sum_{\lambda'} \int_{\mathbb{R}^2} \int_{\mathbb{R}} x(\rho(u) + u', \alpha + \alpha', \lambda') W_{\lambda', \lambda} (2^{-\alpha} u', \alpha') 2^{-2\alpha} du' d\alpha' \right. \\
& \quad \left. - \sum_{\lambda'} \int_{\mathbb{R}^2} \int_{\mathbb{R}} x(u + u', \alpha + \alpha', \lambda') W_{\lambda', \lambda} (2^{-\alpha} u', \alpha') 2^{-2\alpha} du' d\alpha' \right|
\end{aligned}$$

$$\begin{aligned}
&= \left| \sum_{\lambda'} \int_{\mathbb{R}^2} \int_{\mathbb{R}} x_c(\rho(u) + u', \alpha + \alpha', \lambda') \sum_m W_{\lambda', \lambda, m}(2^{-\alpha} u') \varphi_m(\alpha') 2^{-2\alpha} du' d\alpha' \right. \\
&\quad \left. - \sum_{\lambda'} \int_{\mathbb{R}^2} \int_{\mathbb{R}} x_c(u + u', \alpha + \alpha', \lambda') \sum_m W_{\lambda', \lambda, m}(2^{-\alpha} u') \varphi_m(\alpha') 2^{-2\alpha} du' d\alpha' \right| \\
&= \left| \sum_{\lambda'} \sum_m \int_{\mathbb{R}^2} G_m(v, \alpha, \lambda') k_{\lambda', \lambda, m}(v, u, \alpha) du' \right|,
\end{aligned}$$

where

$$\begin{aligned}
G_m(u, \alpha, \lambda') &:= \int_{\mathbb{R}} x_c(u, \alpha + \alpha', \lambda') \varphi_m(\alpha') d\alpha', \\
k_{\lambda', \lambda, m}(v, u, \alpha) &:= 2^{-2\alpha} [W_{\lambda', \lambda, m}(2^{-\alpha}(v - \rho(u))) - W_{\lambda', \lambda, m}(2^{-\alpha}(v - u))].
\end{aligned}$$

Similar to (47), we have the following bound

$$\int_{\mathbb{R}^2} |k_{\lambda', \lambda, m}(v, u, \alpha)| du, \quad \int_{\mathbb{R}^2} |k_{\lambda', \lambda, m}(v, u, \alpha)| dv \leq 2|\nabla \tau|_{\infty} D_{\lambda', \lambda, m}^{(l)}. \quad (50)$$

Again, the proof of (50) is the same as that of Proposition 3.4 in [16] after a change of variable. The rest of the proof follows from a similar argument as in (45) and (46). \square

B Experimental Details in Section 5

B.1 Verification of scale-equivariance

The ScDCFNet used in this experiment has two convolutional layers, each of which is composed of a scale-equivariant convolution (4) or (5), a batch-normalization, and a 2×2 scale-equivariant average-pooling. The expansion coefficients $a_{\lambda', \lambda}^{(1)}(k)$ and $a_{\lambda', \lambda}^{(2)}(k, m)$ are sampled from a Gaussian distribution and truncated to $K = 8$ and $K_{\alpha} = 3$ leading coefficients for u and α respectively. Similarly, a regular CNN with two convolutional layers and randomly generated 5×5 convolutional kernels is used as a baseline for comparison.

B.2 Multiscale Image Classification

In the experiments on multiscale image classification on the SMNIST and SFashion dataset, the network architectures for the ScDCFNet and the regular CNN are shown in Table 2. Stochastic gradient descent (SGD) with decreasing learning rate from $10^{-2}(10^{-1})$ to $10^{-4}(10^{-3})$ is used to train all networks without (with) batch-normalization for 160 epochs.

Layer	(Regular) CNN	ScDCFNet
1	c3x3x1xM ReLU ap2x2	sc(9)9x9x1xM ReLU sap2x2
2	c3x3xMx2M ReLU ap2x2	sc(9)9x9x L_α xMx2M ReLU sap2x2
3	c3x3x2Mx4M ReLU ap2x2	sc(9)9x9x L_α x2Mx4M ReLU sap2x2
4	fc64 ReLU fc10 softmax-loss	fc64 ReLU fc10 softmax-loss

Table 2: Network architectures used for the experiments in Section 5.2. **cLxLxM’xM**: a regular convolutional layer with M’ input channels, M output channels, and LxL spatial kernels. **sc(N_α)LxLxM’xM**: the first-layer convolution operation (4) in ScDCFNet, where N_α is the number of the uniform grid points to discretize the scale interval $[-1.6, 0]$, and LxL is the spatial kernel size on the largest scale $\alpha = 0$. **sc(N_α)LxLx L_α xM’xM**: the l -th layer ($l > 1$) convolution operation (5) in ScDCFNet, where the extra symbol L_α stands for the filter size in α . **apLxL(sapLxL)**: the regular(scale-equivariant) LxL average-pooling. **fcM**: a fully connected layer with M output channels. Batch-normalization layers are added to each convolutional layer if adopted during training.

B.3 Image Reconstruction

The network architectures for the SDCFNet and regular CNN auto-encoders are shown in Table 3. The filter expansion in the SDCFNet auto-encoder is truncated to $K = 8$ and $K_\alpha = 3$. SGD with decreasing learning rate from 10^{-2} to 10^{-4} is used to train both networks for 20 epochs.

Layer	Regular auto-encoder	ScDCF auto-encoder
1	c7x7x1x8 ReLU ap2x2	sc(15)13x13x1x4 ReLU sap2x2
2	c7x7x8x16 ReLU ap2x2	sc(15)13x13x3x4x8 ReLU sap2x2
3	fc128 ReLU fc4096 ReLU	fc128 ReLU fc4096 ReLU
4	ct7x7x16x8 ReLU us2x2	ct7x7x16x8 ReLU us2x2
5	ct7x7x8x1 ReLU us2x2	ct7x7x8x1 ReLU us2x2

Table 3: Architectures of the auto-encoders used for the experiment in Section 5.3. The encoded representation is the output of the second layer. **ctLxLxM’xM**: transposed-convolutional layers with M’ input channels, M output channels, and LxL spatial kernels. **us2x2**: 2x2 spatial upsampling. See the caption of Table 2 for the definitions of other symbols. Batch-normalization (not shown in the table) is used after each convolutional layer.