



HAL
open science

Scale Equivariant U-Net

Mateus Sangalli, Samy Blusseau, Santiago Velasco-Forero, Jesus Angulo

► **To cite this version:**

Mateus Sangalli, Samy Blusseau, Santiago Velasco-Forero, Jesus Angulo. Scale Equivariant U-Net. 33rd British Machine Vision Conference, Nov 2022, Londres, United Kingdom. <hal-03805743>

HAL Id: hal-03805743

<https://hal.science/hal-03805743v1>

Submitted on 7 Oct 2022

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire HAL, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



HAL Authorization

SCALE EQUIVARIANT U-NET

Mateus Sangalli, Samy Blusseau, Santiago Velasco-Forero, Jesús Angulo

Center for Mathematical Morphology

Mines ParisTech, PSL University

Fontainebleau, France

{mateus.sangalli, sammy.blusseau, santiago.velasco, jesus.angulo}@minesparis.psl.eu

ABSTRACT

In neural networks, the property of being equivariant to transformations improves generalization when the corresponding symmetry is present in the data. In particular, scale-equivariant networks are suited to computer vision tasks where the same classes of objects appear at different scales, like in most semantic segmentation tasks. Recently, convolutional layers equivariant to a semigroup of scalings and translations have been proposed. However, the equivariance of subsampling and upsampling has never been explicitly studied even though they are necessary building blocks in some segmentation architectures. The U-Net is a representative example of such architectures, which includes the basic elements used for state-of-the-art semantic segmentation. Therefore, this paper introduces the Scale Equivariant U-Net (SEU-Net), a U-Net that is made approximately equivariant to a semigroup of scales and translations through careful application of subsampling and upsampling layers and the use of aforementioned scale-equivariant layers. Moreover, a scale-dropout is proposed in order to improve generalization to different scales in approximately scale-equivariant architectures. The proposed SEU-Net is trained for semantic segmentation of the Oxford Pet IIIT and the DIC-C2DH-HeLa dataset for cell segmentation. The generalization metric to unseen scales is dramatically improved in comparison to the U-Net, even when the U-Net is trained with scale jittering, and to a scale-equivariant architecture that does not perform upsampling operators inside the equivariant pipeline. The scale-dropout induces better generalization on the scale-equivariant models in the Pet experiment, but not on the cell segmentation experiment.

1 Introduction

Convolutional Neural Networks (CNN) are based on convolutional layers and achieve state-of-the-art performance in many image analysis tasks. A translation applied to the inputs of a CNN is equivalent to a translation applied to its features maps, a property illustrated by Figure 1(a). This property is a particular case of group equivariance [2] and helps improve the generalization of the network to new data if the data has translation symmetry. An operator $\phi : \mathcal{X} \rightarrow \mathcal{Y}$ is equivariant w.r.t. a group if applying a group action in the input and then ϕ , amounts applying a group action to the output of ϕ given the original inputs. This is illustrated in Figure 1. In addition to translations, group actions can model many interesting classes of spatial transformations such as rotations, scalings, and affine transformations. Group equivariant CNNs [2] are a generalization of CNNs that are equivariant to some transformation group. Many approaches focus on equivariance to rotations, in different kinds of data [2, 18, 17, 14] and to scalings [20, 3, 7].

Deep scale-spaces [19] introduce neural networks equivariant to the action of semigroups, instead of groups. Semigroup actions are considered as they can model non-invertible transformations, and the authors focus on equivariance to downsampling in discrete domains as a way to address equivariance to scalings without creating spurious information through interpolation. This seminal work laid the basis to define scale-equivariant CNNs, although it only focused on convolutional layers and did not address the equivariance of pooling and upsampling layers, which are key elements in many neural architectures, such as U-Net.

The U-Net [10] has become famous for its great performance in semantic segmentation. It is a fully convolutional neural network, *i.e.* a CNN without any dense layer, and therefore it is equivariant to a certain subgroup of translations. However, architectures like U-Net are not scale equivariant *a priori*, and experiments show they are not in practice [11]



Figure 1: Example of equivariance in the cases of translation and scaling. In this case, ϕ is an ideal operator that computes the semantic segmentation of images. The operators T_v and R_s are, respectively, a translation and a re-scaling.




Figure 2: Example where a U-Net trained on one scale and is applied to predict an output on the training(a) and an unseen(b) scale. The image with the unseen scale represents the same object but the U-Net no longer segments it correctly.

as illustrated by Figure 2. A scale-equivariant counterpart of such an architecture is desirable as scale symmetry is frequently present in semantic segmentation data. For example, in urban scenes, objects of the same class appear at different scales depending on their distances to the camera.

In this work we introduce the Scale-Equivariant U-Net (SEU-Net) based on semigroup cross-correlations [19] and an adapted use of pooling and upsampling. The rest of the paper is organized as follows. In Section 2 we discuss some of the related work in the literature. In Section 3 we review the semigroup equivariant neural networks. The main contribution of this paper, the SEU-Net, is introduced in Section 4 along with its fundamental building blocks. The whole architecture is tested empirically for its equivariance in Section 5. More precisely, we test the SEU-Net¹ in segmentation tasks where the test images are in scales unseen during training, on the Oxford-IIIT Pet [9] and the DIC-HeLa cell [15] datasets. The SEU-Net is shown to overperform the U-Net even when the latter is training with large values of scale jittering. The paper ends in Section 6 with some conclusions and perspectives for future work.

2 Related Work

Scale-equivariance and scale-invariance are topics already discussed in the deep learning literature [20, 3, 5, 7, 13]. The experimental benchmarks found in those papers are interesting as a first way to measure equivariance, but tend to be based on very simple tasks, such as the classification of re-scaled digits from the MNIST dataset or low resolution images of clothes from the Fashion-MNIST dataset. In [12], combinations of base filters are optimized to minimize the equivariance error of discrete scale convolutions. This is applied to classification, tracking and geometry estimation, but not segmentation.

In [19], instead of treating the scaling as an invertible operation, such as it would behave in a continuous domain, it is considered the action of downsampling the input image in a discrete domain. Therefore a *semigroup*-equivariant generalization of the convolution is introduced. Specifically, the focus is put on a semigroup of scalings and translations. These operators can be efficiently applied even on large images, since applying it at larger scales has the same computational cost. In [19] the semigroup equivariant models were applied to classification and semantic segmentation of datasets of large images, achieving better results compared to matched non-equivariant architectures. Yet, the role of

¹Code available at <https://github.com/mateussangalli/ScaleEquivariantUNet>

scale-equivariance was not isolated, as the performance of the models was not measured for inputs on scales unseen in the training set. Later on, this approach was revisited by [11], where the Gaussian scale-space originally used was generalized to other scale-spaces and the models were tested in experiments where the networks are trained in one fixed scale and tested on unseen scales, albeit on synthetic or simple datasets. In all these approaches, the authors either avoided pooling and upsampling in their architectures, or used them but did not discuss their impact on scale equivariance.

While scale-equivariance has been a topic in the literature for some time, as far as we know a scale-equivariant U-Net has not yet been proposed, contrary to the rotation-equivariance case [1]. Moreover, the current benchmarks for scale-equivariance were either based on simple datasets like MNIST or did not explicitly measure the equivariance in their segmentation or classification experiments, by training the networks on one fixed scale and testing on unseen scales. Here we propose semantic segmentation experiments based on natural data which measure the equivariance of the predictions.

3 Semigroup Equivariant Convolutional Networks

In this work and following [19], image scalings are restricted to image downscalings, which can be viewed as actions of a *semigroup* on images. As illustrated by Figure 1, we seek equivariance with respect to both downscalings and translations. Hence, the network layers are designed to be equivariant with respect to a semigroup combining both transformations.

3.1 Semigroup Equivariance

A semigroup, contrary to a group, can model non-invertible transformations, e.g. the downsampling operation in a discrete domain. In the following, (G, \cdot) denotes a discrete semigroup.

Let \mathcal{X} be a set, a family of mappings $(\varphi_g)_{g \in G}$ from \mathcal{X} to itself, is a semigroup action on \mathcal{X} if it is homomorphic to the semigroup, that is, if either $\forall g, h \in G, \varphi_g \circ \varphi_h = \varphi_{g \cdot h}$ (left action), or $\forall g, h \in G, \varphi_g \circ \varphi_h = \varphi_{h \cdot g}$ (right action). In this paper we will consider in particular the following right action, acting on \mathcal{F} the set of functions from G to \mathbb{R}^n .

$$\forall u, g \in G, \forall f \in \mathcal{F}, \quad R_u(f)(g) = f(u \cdot g). \quad (1)$$

Given two sets \mathcal{X} and \mathcal{Y} , a mapping $H : \mathcal{X} \rightarrow \mathcal{Y}$ is said equivariant with respect to G if there are semigroup actions $(\varphi_g)_{g \in G}$ and $(\psi_g)_{g \in G}$ on \mathcal{X} and \mathcal{Y} respectively, such that $\forall g \in G, \quad H \circ \varphi_g = \psi_g \circ H$. This definition gets more intuitive when $\mathcal{X} = \mathcal{Y}$ is a set of images and $(\varphi_g)_{g \in G} = (\psi_g)_{g \in G}$ are scalings or translations, as illustrated in Figure 1.

3.2 Scale-cross-correlation

For $\gamma > 1$ an integer, let $\mathcal{S}_\gamma = \{\gamma^n | n \in \mathbb{N}\}$, endowed with the multiplication, the semigroup representing discrete scalings of base γ . Then we consider the semigroup $G = \mathcal{S}_\gamma \times \mathbb{Z}^2$ of discrete scalings and translations, endowed with the internal operator “ \cdot ”, defined by

$$\forall k, l \in \mathbb{N}, z, y \in \mathbb{Z}^2 \quad (\gamma^k, z) \cdot (\gamma^l, y) = (\gamma^{k+l}, \gamma^k y + z). \quad (2)$$

Following (1), the action of this semigroup on functions mapping $\mathcal{S}_\gamma \times \mathbb{Z}^2$ to \mathbb{R} is $R_{\gamma^k, z}[f](\gamma^l, y) = f(\gamma^{k+l}, \gamma^k y + z)$. In analogy to convolutions, which are linear and equivariant to translations, a key step in equivariant CNNs is defining linear operators which are equivariant to some class of operators. The *semigroup cross-correlation*, defined for an image $f : G \rightarrow \mathbb{R}$ and a filter $h : G \rightarrow \mathbb{R}$ is a generalization of the convolution which is linear and equivariant to the action R_g of a semigroup. When applied to the semigroup of scales and translations, we obtain the *scale-cross-correlation*. Both were introduced in [19]. The scale-cross-correlation is written²

$$(f \star_G h)(\gamma^k, z) = \sum_{(\gamma^l, y) \in G} R_{\gamma^k, z}[f](\gamma^l, y) h(\gamma^l, y) = \sum_{l \geq 0} \sum_{y \in \mathbb{Z}^2} f(\gamma^{k+l}, \gamma^k y + z) h(\gamma^l, y). \quad (3)$$

This operator is suited for single channel images on G , but it can be easily extended to multichannel images. Let the input $f = (f_1, \dots, f_n) \in (\mathbb{R}^n)^G$ be a signal with n channels. Assuming the output has m channels, the filter is of the form $h : G \rightarrow \mathbb{R}^{n \times m}$. We compute the operator $f \star_G h$ at channel $o \in \{1, \dots, m\}$ as $(f \star_G h)_o := \sum_{c=1}^n (f_c \star_G h_{c,o})$.

The resulting map is equivariant to scalings and translations: $(R_g f \star_G h)_o = R_g((f \star_G h)_o)$. Note that the composition of operators which commute with R_g still commutes with R_g , for which concatenating scale-cross-correlation layers followed by pointwise activation functions and batch normalization yields equivariant architectures.

²The equations in the case of a general semigroup can be found in Appendix A.

3.3 Lifting and Projection

The operators of the previous section are defined on the set of functions with the semigroup as a domain, $\mathcal{F} = (\mathbb{R}^n)^G$, but images input to networks are functions $f : \mathbb{Z}^2 \rightarrow \mathbb{R}^n$. In this section we review the lifting and projections layers - operators which map images to functions on the semigroup and vice-versa.

Lifting. A *lifting* operator Λ is used to map an input function $f : \mathbb{Z}^2 \rightarrow \mathbb{R}^n$ into a function $\Lambda f : G \rightarrow \mathbb{R}^n$. Once lifted to the semigroup space, linear equivariant operators can be applied according to Eq. (3). As pointed out in [11], a sufficient condition to have equivariance of the composition of the lifting followed by the semigroup cross-correlations, is that $\Lambda \circ R'_{\gamma^k, z} = R_{\gamma^k, z} \circ \Lambda$, where $R'_{\gamma^k, z}$ is the re-scaling action for images on \mathbb{Z}^2 : $R'_{\gamma^k, z}[f](t) = f(\gamma^k t + z)$.

Whereas in [11] several liftings are explored, in this paper we set the lifting to the Gaussian scale-space Λ_G , like in [19]. For an image $f : \mathbb{Z}^2 \rightarrow \mathbb{R}$ a point $z \in \mathbb{Z}^2$ and a scale level $k \in \mathbb{N}$,

$$\Lambda_G(f)(\gamma^k, z) = (f * \mathcal{G}_{\gamma^k})(z) \quad (4)$$

where $*$ is the classic discrete convolution and \mathcal{G}_{γ^k} the discrete Gaussian kernel with scale γ^k .

Projection. To project back into the image space, we apply a *max-projection* along the scale dimension, defined by $\forall z \in \mathbb{Z}^2 \Pi[f](z) = \sup_{k \in \mathbb{N}} \{f(\gamma^k, z)\}$. To be consistent with the lifting, we would like to have $R'_{\gamma^k, z} \circ \Pi = \Pi \circ R_{\gamma^k, z}$. Instead, we have $R'_{\gamma^k, z} \Pi f(y) = \sup_{l \in \mathbb{N}} f(\gamma^l, \gamma^k y + z)$ and $\Pi R_{\gamma^k, z} f(y) = \sup_{l \in \mathbb{N}} f(\gamma^{l+k}, \gamma^k y + z) = \sup_{l \geq k} f(\gamma^l, \gamma^k y + z)$ so that $R'_{\gamma^k, z} \Pi f(y) = \max\{\Pi R_{\gamma^k, z} f(y), \max_{0 \leq l < k} f(\gamma^l, \gamma^k y + z)\}$. The previous expression will be equivariant if the scale where the maximum is attained is smaller than k , but in general we can only hope for approximate equivariance for small enough k . The approximate equivariance will be empirically verified in experiments in Section 5. Note that other projections (e.g. sum or average) have the same flaw, as this is intrinsic to the semigroup-equivariant approach, even though it was omitted previously in the literature.

4 Scale-Equivariant U-Net

Recall that the U-Net [10], illustrated in Figure 3, is a CNN architecture for semantic segmentation based on an auto-encoder structure with skip connections linking the encoder and decoder. As such, it has four main components: convolution blocks, pooling, upsampling and skip connections. In this section we aim to propose the Scale-Equivariant U-Net (SEU-Net), in order to have a U-Net with increased generalization capacity.

In the framework of the previous section, a network can be written as $\Gamma = \Pi \circ \Sigma \circ \Lambda$, where Λ and Π are the lifting and projection respectively, and Σ is the core part of the network mapping the lifted space to itself. We already saw that $\Lambda \circ R'_{\gamma^k, z} = R_{\gamma^k, z} \circ \Lambda$ and we assume $R'_{\gamma^k, z} \circ \Pi \approx \Pi \circ R_{\gamma^k, z}$. Hence, to build a (approximately) scale-equivariant network, it is sufficient to have $\Sigma \circ R_{\gamma^k, z} = R_{\gamma^k, z} \circ \Sigma$. In particular, a way to render the U-Net scale-equivariant is to design scale-equivariant versions of its components in Σ . Convolutions are already rendered equivariant by scale-cross-correlations, and pointwise-non-linearities, batch-normalization and skip connections are equivariant as is. The rest of this section is dedicated to examining the remaining components: subsampling and upsampling.

4.1 Subsampling

Classical pooling operators naively applied scale by scale do not result in scale-equivariant poolings in the lifting space. For example, the max-pooling $\text{MP}[f](x) = \max_{y \in N} f(rx + y)$ with strides $r \in \mathbb{N}$ and neighborhood $N \subseteq \mathbb{Z}^2$ (usually a $r \times r$ square). Its naive extension to the lifted space $\text{MP}'[f](\gamma^k, x) = \text{MP}[f(\gamma^k, \cdot)](x) \forall k \in \mathbb{N}$ does not commute with $R_{\gamma^k, x}$.

Strided convolutions however, generalize well to this scenario, written as the subsampling operator $D_t[f](\gamma^k, x) = f(\gamma^k, tx)$ following a scale-cross-correlation. We can verify that it is scale-equivariant: $D_t[R_{\gamma^k, x} f](\gamma^l, y) = (R_{\gamma^k, x} f)(\gamma^l, ty) = f(\gamma^{l+k}, \gamma^k ty) = D_t[f](\gamma^{l+k}, \gamma^k y) = R_{\gamma^k, x}[D_t f](\gamma^l, y)$. We use strides as the subsampling in our networks, with a stride of $t = 2$.

4.2 Upsampling

Upsampling blocks are a well established part of modern neural network architectures for segmentation and other tasks. In order to extend upsampling to a scale-equivariant setting, we look at the case where f is defined on a continuous domain. In that case, the downsampling D_{γ^l} has an inverse U_{γ^l} which is the natural upsampling.




Figure 3: Illustration of the SEU-Net architecture. We parametrize it by the height (i.e. number of subsamplings/upsamplings set to four in this example) and number of filters C in the first layer (after pooling we double the number of filters and after subsampling we halve it). The values H, W represent the height, width of the inputs images, S is the number of scales in the lifting and C_{in} , are the number of channels in the input image or convolutional filters and C_{out} is the number of classes. Two convolutions are performed between subsamplings and between upsamplings. Feature maps connected by skip connections have the same spatial dimensions.

In the discrete case the problem becomes more complicated as downscaling is not invertible, but for $k, l \in \mathbb{N}$ we can define an upsampling U_{γ^l} as an operator satisfying $\forall x \in \mathbb{Z}^2$

$$U_{\gamma^k}[f](\gamma^l, \gamma^k x) = f(\gamma^l, x) \quad \text{and} \quad U_{\gamma^{lk}} = U_{\gamma^l} \circ U_{\gamma^k}. \quad (5)$$

With this, we have $D_{\gamma^k} \circ U_{\gamma^k} = \text{id}$. For all k , $U_{\gamma^k}(f)$ values are only restricted in the points $y \in k\mathbb{Z}^2 = \{kx | x \in \mathbb{Z}^2\}$, and the values on the other pixels can be defined in several ways (e.g. copies, interpolation) as long as it satisfies (5). Now, if $U_{\gamma^l} R_{\gamma^k, x} f = R_{\gamma^k, \gamma^l x} U_{\gamma^l} f$ for any f then $\Sigma \circ R_{\gamma^k, x} = R_{\gamma^k, x} \circ \Sigma$. Indeed let $\psi_i = L_i \cdots D_{\gamma^l} L_1, i = 1, \dots, m$ denote the part of a SEU-Net of height m before the i -th downsampling block, where $L_j, j = 1, \dots, m$, are blocks that commute with $R_{\gamma^k, x}$ (constructed by scale-cross-correlations, pointwise activations and batch normalization). Denote $\phi_m = L_{m+1} \psi_m$ and $\phi_i = L_i C(U_{\gamma^l} \phi_{i+1}, \psi_i), i = m, \dots, 1$ where C denotes concatenation. With the above hypothesis, we have $\phi_i R_{\gamma^k, x} f = R_{\gamma^k, \gamma^l x} \phi_i f$. In particular, $R_{\gamma^k, x} \phi_0 f = \phi_0 R_{\gamma^k, x} f$, and we notice that ϕ_0 is precisely Σ .

The sufficient condition $U_{\gamma^l} R_{\gamma^k, x} f = R_{\gamma^k, \gamma^l x} U_{\gamma^l} f$ is not verified in general (see Appendix B), but Proposition 1 introduces a setting where it does.

Proposition 1. For $N \in \mathbb{N}^*$ and $i \in \{1, \dots, N\}$, let $\mathcal{U}_i = \{U_{\gamma^{n_i+t}} f_i | t \in \mathbb{N}\}$, where each $f_i : G \rightarrow \mathbb{R}^n$ is a function on G and each n_i an integer. Let $n_0 \leq \min\{n_i | i = 1, \dots, N\}$ and $\mathcal{U} = \bigcup_{i=1}^N \mathcal{U}_i$. Then for all $f \in \mathcal{U}$, and $k, l \in \mathbb{N}$ such that $k - l \leq n_0$, we have³ $U_{\gamma^l} R_{\gamma^k, x} f = R_{\gamma^k, \gamma^l x} U_{\gamma^l} f$.

This property states that upsampling behaves as an equivariant operator as long as the input image is an upsampling of some image in a base scale. It can be interpreted as saying that the downscaling should not destroy information of the images in \mathcal{U} . We model this by constraining the scaling factors of the downscaling actions and assuming that the objects of interest in an image are sufficiently big. We would like to point out that this hypothesis is never verified but reasonable for most of the datasets for semantic segmentation.

Before moving on to the experimental part, let us sum up the theoretical properties of a SEU-Net $\Gamma = \Pi \circ \Sigma \circ \Lambda$. By our construction we can hope for an approximated scale-equivariance $\Gamma \circ R'_{\gamma^k, z} \approx R'_{\gamma^k, z} \circ \Gamma$. Two approximations prevent from exact equivariance: The approximated equivariance of the projection operator Π , which is intrinsic to the lifting approach, and the assumption to guarantee an equivariant upsampling, which is never verified in practice. We will see in our experiments that the SEU-Net shows a high degree of scale-equivariance despite these approximations. Each of these approximations is intrinsic to the problem. If the problem was formulated in a continuous domain $\mathcal{S}_\gamma \times \mathbb{R}^2$ upsampling would be theoretically equivariant, but its implementation would have the same problems.

³For a proof, see Appendix C.

5 Experiments

In this section we test the proposed SEU-Net in two segmentation tasks where we evaluate its generalization to unseen scales. We train the SEU-net on a set where objects have roughly the same scale and test it on a wide range of scales. For these experiments we use a scale base of $\gamma = 2$, downsampling D_2 and upsampling U_2 computed by bilinear interpolation. Quantitative results will be measured using Intersection over Union (IoU) and consistency. We define consistency as follows: given a segmentation neural network ϕ , the consistency is the probability of assigning the same label to a pixel after it has been transformed, formally $\text{Cons}(\phi, s) = P(\phi(R_{s,0}[f])(x) = R_{s,0}[\phi(f)](x))$.

We compare the SEU-Net to U-Net and to the SResNet [19], a scale-equivariant residual architecture which applies subsampling but no upsampling inside the equivariant pipeline, i.e. it only applies an upsampling *after* the projection layer. Hence, it does not benefit from the features that made U-Net more suitable for segmentation, namely the skip connections at several upsampling stages.

Scale Dropout. In order to produce more robust results with respect to scale changes we propose the use of *Scale Dropout* before the projection layers. Given a feature map $f : S_\gamma \times \mathbb{Z}^2 \rightarrow \mathbb{R}^n$, we compute its scale dropout of rate $p \in [0, 1]$ as $\text{ScaleDropout}_p(f)(s, x) = X(s)f(s, x)$ where $X(s)$ is a Bernoulli variable of parameter p , i.e. $P(X(s) = 1) = 1 - p$ and $P(X(s) = 0) = p$. In our experiments we use values $p = 0$ (no dropout) and $p = 0.25$.

5.1 Oxford-IIIT Pet Dataset

The Oxford-IIIT Pet ⁴ dataset [9] consists of pictures containing cats and dogs. The relevant labeling for this paper, the trimaps, is the segmentation of the images into three classes: the animal, the background and the boundaries of the animal. In Figure 5(a) and (b) we see an example of an image and its corresponding ground truth. The dataset was loaded from the TensorFlow package⁵, where it is divided into 3680 training samples and 3669 test samples. To make the validation set we removed 200 test samples. During training and testing images are resized to 224×224 pixels. We define multiple test sets by re-scaling the original test set by $s \in \{2^{\frac{1}{2}} | i \in \{-4, -3, \dots, 4\}\}$. We used bilinear interpolation to up-scale images.

Both the U-Net and SEU-Net have height four and contain sixteen filters in the first layer and use 3×3 filters. The SEU-Net truncates at four scales, and filters have depth one in the scales dimension (their values is different from zero in one scale value). The networks are trained using the Adam [6] optimizer with categorical cross-entropy loss. Training the U-Net, SResNet and SEU-Net takes approximately 24, 73, and 97 seconds per epoch respectively, on a Tesla P100-SXM2-16Gb GPU.

Comparison with data augmentation. We also performed scale jittering in the U-Net to compare the effect of the equivariant network with the effect of data augmentation. Scale jittering is performed by rescaling the image by a randomly chosen scale α and either random cropping or padding to the original image. We trained a U-Net with scale jittering in the interval $[\frac{1}{4}, 4]$, equal to the interval of test scales.

Results. The overall results in terms of the IoU are shown in Figure 4. Firstly we notice that the SEU-Net increases performance compared to both SResNet and U-Net. SResNet, however, does not consistently generalize better than the U-Net. Dropout improves the quality of SEU-Net, particularly for more extreme scales, indeed, for larger scales the augmented U-Net has a better IoU than the SEU-Net without scale dropout, but not than the one with scale dropout. The augmented U-Net loses performance scale 1, it would probably need to be larger to retain the same performance. We show examples of the predictions of the U-Net and SEU-Net in Figure 5.

5.2 Cell Segmentation

We also evaluate the models in a medical image segmentation dataset, namely the DIC-C2DH-HeLa dataset [15] of HeLa cells on a flat glass recorded by differential interference contrast (DIC). We used 83 images for train/validation and 83 for testing. Figure 7 (a) and (b) shows an example from the test set with its labels at different scales. Models are trained with the AdamW optimizer [8]. Like in the previous experiment, we first train the models in the training set and test in the test set re-scaled by different values. We also perform scale jittering, but now for both U-Net and SEU-Net. For U-Net we trained models with scale jittering with ranges 4 (α is chosen each step from the interval $[\frac{1}{4}, 4]$) and 1.5 (α is chosen from the interval $[\frac{2}{3}, \frac{3}{2}]$) and for SEU-Net we used only the range 1.5 jittering.

⁴<https://www.robots.ox.ac.uk/~vgg/data/pets/>, CC BY-SA 4.0 license

⁵https://www.tensorflow.org/datasets/catalog/oxford_iiit_pet




Figure 4: Overall results in terms of IoU and Consistency for each scale of the Pet dataset.

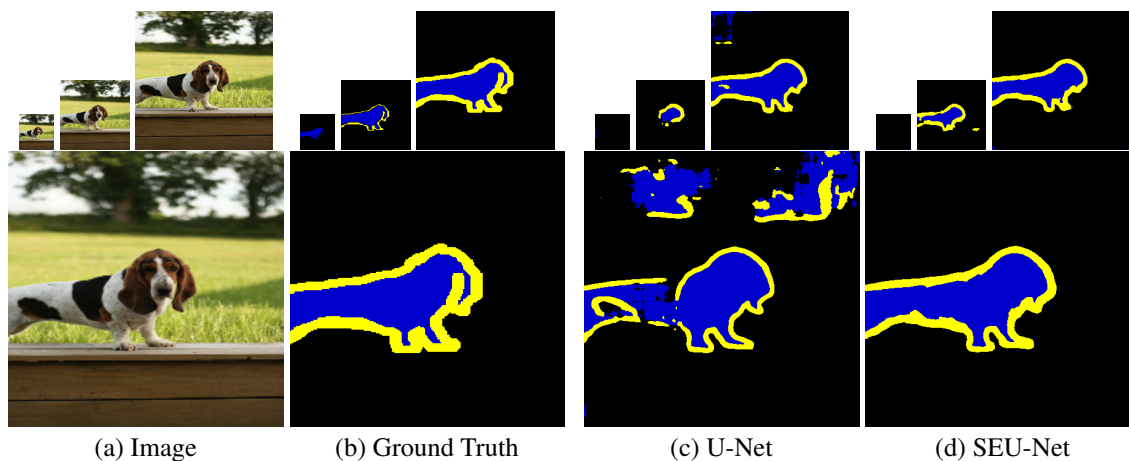


Figure 5: Sample test image at different scales and ground truth from the Oxford-IIIT Pet dataset, along with the U-Net and SEU-Net predictions. The scales present are 0.25, 0.5, 1 and 2 times the training scale.

Results. Figure 6 (a) shows the IoU of different models on the re-scaled test sets. Figure 7 shows some segmentation examples. Again, the SEU-Net outperforms the U-Net. The poor results of the SResNet for smaller scales is possibly due to the cell images containing more high-frequency information, compared to the pets images. In contrast to the previous experiment, dropout did not seem to significantly increase performance of the SEU-Net, neither in the train scale nor the test scales. Moreover the SResNet results were greatly decreased due to dropout. This is likely a result of the augmented dataset being more difficult to segment than the original and not being representative of the dataset at base scale. The gain in generalization is only better than the SEU-Net for the smallest scales. The jittering with range 1.5 does not have a very noticeable effect. On the other hand the SEU-Net with 1.5 jittering has a noticeable gain in generalization to larger scales.

6 Conclusions and Future Work

In this paper we revisited the framework of scale semigroup-equivariant neural networks and applied it to the definition of a Scale Equivariant U-Net for semantic segmentation. Experimental results show that the SEU-Net can greatly improve the generalization to new scales and even the performance in the training scale. Moreover, the results lead us to conjecture that the U-Net with scale jittering would need to have more parameters to have a good performance in all the range of scales, while the SEU-Net achieves good results without increasing its size. The results suggest that implementing that the improvement comes not only from the scale-equivariant cross-correlations, but also from the SEU-Net global architecture and applying the pooling operators inside the equivariant pipeline. The proposed scale




Figure 6: IoUs of the cell segmentation experiment with comparisons with U-Net, SResNet and data augmentation. U-Net (aug. 4) refers to the U-Net trained with scale jittering with range 4 and U-Net (aug. 1.5) refers to the U-Net trained with jittering with range 1.5. The same for SEU-Net (aug. 1.5).




Figure 7: Predictions from DIC-HeLa at different scales, namely scales 0.5, 1 and 2. Although the U-Net suffers from the scale change, especially the change to a larger scale, the SEU-Net can still capture the overall structure of the cells.

dropout was also shown to have the potential to increase scale-equivariant models' performance. In future works it would be interesting to study an equivariant regularization term such as in [12] in addition to the scale-dropout.

Acknowledgements

This work was granted access to the HPC resources of IDRIS under the allocations AD I011013637 and AD 011012212R1.

References

- [1] B. Chidester, T.-V. Ton, M.-T. Tran, J. Ma, and M. N. Do. Enhanced rotation-equivariant u-net for nuclear segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, June 2019.
- [2] T. Cohen and M. Welling. Group equivariant convolutional networks. In *International conference on machine learning*, pages 2990–2999, 2016.
- [3] R. Ghosh and A. K. Gupta. Scale steerable filters for locally scale-invariant convolutional neural networks. *arXiv preprint arXiv:1906.03861*, 2019.
- [4] H. J. Heijmans and R. van den Boomgaard. Algebraic framework for linear and morphological scale-spaces. *Journal of Visual Communication and Image Representation*, 13(1-2):269–301, 2002.

- [5] Y. Jansson and T. Lindeberg. Exploring the ability of cnns to generalise to previously unseen scales over wide scale ranges. *arXiv preprint arXiv:2004.01536*, 2020.
- [6] D. P. Kingma and J. Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [7] T. Lindeberg. Scale-covariant and scale-invariant gaussian derivative networks. In *International Conference on Scale Space and Variational Methods in Computer Vision*, pages 3–14. Springer, 2021.
- [8] I. Loshchilov and F. Hutter. Decoupled weight decay regularization. In *International Conference on Learning Representations*, 2018.
- [9] O. M. Parkhi, A. Vedaldi, A. Zisserman, and C. V. Jawahar. Cats and dogs. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2012.
- [10] O. Ronneberger, P. Fischer, and T. Brox. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, pages 234–241. Springer, 2015.
- [11] M. Sangalli, S. Blusseau, S. Velasco-Forero, and J. Angulo. Scale equivariant neural networks with morphological scale-spaces. In *International Conference on Discrete Geometry and Mathematical Morphology*, pages 483–495. Springer, 2021.
- [12] I. Sosnovik, A. Moskalev, and A. Smeulders. Disco: accurate discrete scale convolutions. In *Proceedings of the 32nd British Machine Vision Conference*, 2021.
- [13] I. Sosnovik, M. Szmaja, and A. Smeulders. Scale-equivariant steerable networks. In *International Conference on Learning Representations*, 2019.
- [14] N. Thomas, T. Smidt, S. Kearnes, L. Yang, L. Li, K. Kohlhoff, and P. Riley. Tensor field networks: Rotation-and translation-equivariant neural networks for 3d point clouds. *arXiv preprint arXiv:1802.08219*, 2018.
- [15] V. Ulman, M. Maška, K. E. Magnusson, O. Ronneberger, C. Haubold, N. Harder, P. Matula, P. Matula, D. Svoboda, M. Radojevic, et al. An objective comparison of cell-tracking algorithms. *Nature methods*, 14(12):1141–1152, 2017.
- [16] R. Van Den Boomgaard and A. Smeulders. The morphological structure of images: The differential equations of morphological scale-space. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 16(11):1101–1113, 1994.
- [17] M. Weiler, M. Geiger, M. Welling, W. Boomsma, and T. S. Cohen. 3d steerable cnns: Learning rotationally equivariant features in volumetric data. In *Advances in Neural Information Processing Systems*, pages 10381–10392, 2018.
- [18] 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.
- [19] D. Worrall and M. Welling. Deep scale-spaces: Equivariance over scale. In *Advances in Neural Information Processing Systems*, pages 7364–7376, 2019.
- [20] W. Zhu, Q. Qiu, R. Calderbank, G. Sapiro, and X. Cheng. Scale-equivariant neural networks with decomposed convolutional filters. *arXiv preprint arXiv:1909.11193*, 2019.

A Semigroup cross-correlation

Let $\mathcal{F} = \mathbb{R}^G$ denote the set of functions mapping G to \mathbb{R} . Bearing in mind the final purpose of defining equivariant CNN layers, we focus on linear operators on \mathcal{F} . Let the semigroup right action $(R_u)_{u \in G}$ on \mathcal{F} , defined by

$$\forall u, g \in G, \forall f \in \mathcal{F}, \quad R_u(f)(g) = f(u \cdot g). \quad (6)$$

Then for any $h \in \mathcal{F}$, the linear operator defined by

$$\forall u \in G, \quad H(f)(u) = (f \star_G h) = \sum_{g \in G} R_u(f)(g)h(g) \quad (7)$$

is equivariant to the semigroup action $(R_u)_{u \in G}$, as $H(R_u(f)) = R_u(H(f))$. This class of semigroup equivariant linear operators is the *semigroup cross-correlation* proposed in [19] as the key element to define scale-equivariant convolutional layers. Note that when G is the group of image translations (groups are special semigroups), (7) corresponds to the classic discrete convolution with the reversed filter $h^*(g) = h(g^{-1})$. We use the notation $f \star_G h$ remarking however that this operation is not symmetrical in f and h even when the law \cdot on G is commutative. Also, contrary to the group case, we do not have the property that every linear and equivariant operator can be written as a semigroup-cross-correlation.




Figure 8: Overall results in terms of IoU and Consistency for each scale of the Pet dataset. The SEU-Net has been computed with different pooling functions.

B Different Pooling Operators

Besides the strided scale-crosscorrelations we used, we can define another class of pooling operators, inspired by classical max-pooling. Let us place ourselves in a slightly different case of pooling a function in a continuous domain $f : \mathcal{S} \times \mathbb{R}^2 \rightarrow \mathbb{R}$, with $\mathcal{S} \times \mathbb{R}^2$ acting on it by $R_{\gamma^l, z} f(\gamma^k, x) = f(\gamma^{k+l}, \gamma^l x + z)$, $k, l \in \mathbb{N}, x \in \mathbb{R}^2, z \in \mathbb{Z}^2$. We define the pooling of f as an operator F followed by a downsampling $D_{\gamma^l}[f](\gamma^k, z) = f(\gamma^k, \gamma^l z)$

$$P[f] = D_{\gamma^l} F f. \quad (8)$$

If F commutes with $R_{\gamma^k, x}$, then so does P . We consider three pooling functions: $F_{\text{id}} = \text{id}$ (strides) and two dilation scale-spaces [4]:

- The max-pooling of scale-semigroup-valued images is given by a re-scaled max-pooling

$$F_{\text{max}}[f](\gamma^k, z) = \sup_{y \in N_k \times N_k} \{f(z - y)\} \quad (9)$$

where $N_k = \{\gamma^k x | x \in N\}$ and N is for example a γ^l -sided square in \mathbb{R}^2 .

- The quadratic dilation (quadpool) scale-space is a morphological counterpart to the Gaussian scale-space [16] defined by

$$F_{\text{quad}}[f](\gamma^k, z) = \sup_{y \in \mathbb{R}^2} \left\{ f(z - y) - \frac{\|y\|^2}{c\gamma^{2k}} \right\}, \quad (10)$$

where $c > 0$ is some constant.

In contrast to the strided scale-cross-correlations given by F_{id} , the functions F_{max} and F_{quad} are scale-equivariant only in this continuous setting, their discretized versions are not actually equivariant. Nonetheless, a network employing scale-cross-correlations and these poolings would be equivariant when applied to signals in the domain $\mathcal{S} \times \mathbb{R}^2$.

In Figure 8 we extend the experiments from Section 5.2. Using different pooling functions did not improve the performance of the SEU-Net compared to its performance using strided scale-cross-correlations.

C Non-equivariance of the upsampling

In this section we show that $U_{\gamma^l} R_{\gamma^k, x} f \neq R_{\gamma^k, \gamma^l x} U_{\gamma^l} f$ for at least one lifted image f , one couple of integers (k, l) and a point $x \in \mathbb{Z}^2$. Note that U_{γ^l} is an upsampling defined in the associated paper.

Given any $k \in \mathbb{N}$, take $l = k$, $x = (0, 0)$ and any two lifted images f_1 and f_2 that coincide on certain points,

$$f_1(s, \gamma^k y) = f_2(s, \gamma^k y) \quad \forall s \in \mathcal{S}, y \in \mathbb{Z}^2,$$




Figure 9: Example illustrating the problem with upsampling in a scale-equivariant architecture. We have images f_1 and f_2 such that when both are downsampled and then upsampled they yield the same result, but if both are upsampled and then downsampled they yield different results.

and are different elsewhere, as illustrated in Figure 9. Let us show that $U_{\gamma^l} R_{\gamma^k, x} f_i \neq R_{\gamma^k, \gamma^l x} U_{\gamma^l} f_i$ either for $i = 1$ or $i = 2$ or both. The set of points where f_1 and f_2 coincide implies in particular that $R_{\gamma^k, 0} f_1 = R_{\gamma^k, 0} f_2$. Then we have

$$R_{\gamma^k, 0} U_{\gamma^k} f_1 \neq R_{\gamma^k, 0} U_{\gamma^k} f_2,$$

as $R_{\gamma^k, 0} U_{\gamma^k} f_i : (s, y) \mapsto f_i(\gamma^k s, y)$, and $f_1(\gamma^k s, y) \neq f_2(\gamma^k s, y)$ for $y \notin k\mathbb{Z}^2$. Note that $R_{\gamma^k, 0} U_{\gamma^k}$ is nothing else than an upsampling followed by a downsampling, as in Figure 9.

Since, on the other hand, $R_{\gamma^k, 0} f_1 = R_{\gamma^k, 0} f_2$, we get

$$U_{\gamma^k} R_{\gamma^k, 0} f_1 = U_{\gamma^k} R_{\gamma^k, 0} f_2.$$

Hence, either $R_{k, 0} U_k f_1 \neq U_k R_{k, 0} f_1$ or $R_{k, 0} U_k f_2 \neq U_k R_{k, 0} f_2$ or both, proving our point.

D Proof of Proposition 1

Proof. First, consider $k < m$

$$\begin{aligned} R_{\gamma^k, 0} \circ U_{\gamma^k} f(\gamma^p, y) &= (U_{\gamma^k} f)(\gamma^k \gamma^p, \gamma^k y) \\ &= f(\gamma^k \gamma^p, y) \end{aligned}$$

so $R_{\gamma^k, 0} U_{\gamma^m} f(\gamma^p, y) = U_{\gamma^{m-k}} f(\gamma^{p+m}, y)$.

Now, let $f = U_{\gamma^m} f_i \in F$, $k \leq \min\{n_i | i = 1, \dots, N\} \leq m$, we have

$$\begin{aligned} U_{\gamma} R_{\gamma^k, x} f(\gamma^p, y) &= U_{\gamma} R_{\gamma^k, x} U_{\gamma^m} f_i(\gamma^p, y) \\ &= U_{\gamma} R_{1, x} R_{\gamma^k, 0} U_{\gamma^m} f_i(\gamma^p, y) \\ &= R_{1, \gamma x} U_{\gamma} U_{\gamma^{m-k}} f_i(\gamma^{p+m} r, y) \\ &= U_{\gamma^{m-k+1}} f_i(\gamma^{m+p}, y + \gamma x) \end{aligned}$$

and, on the other hand

$$\begin{aligned} R_{\gamma^k, \gamma x} U_{\gamma} f(\gamma^p, y) &= R_{\gamma^k, \gamma x} U_{\gamma} U_{\gamma^m} f_i(\gamma^p, y) \\ &= R_{1, \gamma x} R_{\gamma^k, 0} U_{\gamma^{m+1}} f_i(\gamma^p, y) \\ &= R_{1, \gamma x} U_{\gamma^{m-k+1}} f_i(\gamma^{p+k}, y) \\ &= U_{\gamma^{m-k+1}} f_i(\gamma^{p+m}, y + \gamma x) \\ &= U_{\gamma} R_{\gamma^k, x} f(\gamma^p, y), \end{aligned}$$

implying $U_{\gamma} R_{\gamma^k, x} f = R_{\gamma^k, \gamma x} U_{\gamma} f$. Repeated application gives us the desired result. \square

E Details of the Experiments

E.1 Oxford IIT Pet

All models, except for the augmented U-Net are trained for 300 epochs. The augmented U-Net is trained for four times as many epochs. To train all models we apply data augmentation consisting of, rotations by a uniformly sampled angles in $[-10^\circ, 10^\circ]$, linear contrast changes by values in the range $[0.9, 1.1]$, random horizontal flipping and random cropping to size 112×112 . Learning rate starts at 10^{-3} and is reduced by 10 when the validation loss does not improve for 30 epochs. We use a batch size of 8.

E.2 DIC-C2DH-HeLa

All models, except for the U-Net with jittering 4 are trained for 200 epochs. The augmented U-Net with jittering 4 is trained for four times as many epochs. To train all models we apply data augmentation consisting of, rotations by a uniformly sampled angles in $[-10^\circ, 10^\circ]$, linear contrast changes by values in the range $[0.9, 1.1]$, random horizontal and vertical flipping and elastic transformations. Learning rate starts at 10^{-3} , weight decay starts at 10^{-4} and both are reduced by exponential decay such that they are divided by 10 every 100 epochs (the decay stops at epoch 300 for the U-Net with size 4 jittering). We use a batch size of 1.

F More Examples of Predictions

In Figure 10 we can see some more examples of predictions from the Oxford Pet dataset, particularly when the U-Net struggles to generalize to new scales. Similarly Figure 11 showcases some extra examples from the experiment from the DIC-HeLa experiment.



Figure 10: Sample test images at different scales and ground truth from the Oxford-IIIT Pet dataset, along with the U-Net and SEU-Net predictions. The scales present are 0.25, 0.5, 1 and 2 times the training scale.




Figure 11: Predictions from DIC-HeLa at different scales, namely scales 0.5, 1 and 2.