
Attentive Group Equivariant Convolutional Networks

David W. Romero¹ Erik J. Bekkers² Jakub M. Tomczak¹ Mark Hoogendoorn¹

Abstract

Although group convolutional networks are able to learn powerful representations based on symmetry patterns, they lack explicit means to learn meaningful relationships among them (e.g., relative positions and poses). In this paper, we present *attentive group equivariant convolutions*, a generalization of the group convolution, in which attention is applied during the course of convolution to accentuate meaningful symmetry combinations and suppress non-plausible, misleading ones. We indicate that prior work on visual attention can be described as special cases of our proposed framework and show empirically that our *attentive group equivariant convolutional networks* consistently outperform conventional group convolutional networks on benchmark image datasets. Simultaneously, we provide interpretability to the learned concepts through the visualization of equivariant attention maps.

1. Introduction

Convolutional Neural Networks (CNNs) (LeCun et al., 1989) have shown impressive performance in a wide variety of domains. The developments of CNNs as well as of many other machine learning approaches have been fueled by intuitions and insights into the composition and *modus operandi* of multiple biological systems (Wertheimer, 1938; Biederman, 1987; Delahunt & Kutz, 2019; Blake & Lee, 2005; Zhaoping, 2014; Delahunt & Kutz, 2019). Though CNNs have achieved remarkable performance increases on several benchmark problems, their training efficiency as well as generalization capabilities are still open for improvement. One concept being exploited for this purpose is that of *equivariance*, again drawing inspiration from human beings.

Humans are able to identify familiar objects despite modifi-

¹Vrije Universiteit Amsterdam, ²University of Amsterdam, The Netherlands. Correspondence to: David W. Romero <d.w.romeroguzman@vu.nl>.




Figure 1. Meaningful relationships among object symmetries. Though every figure is composed by the same elements, only the outermost examples resemble faces. The relative positions, orientations and scales of elements in the innermost examples do not match any meaningful face composition and hence, should not be labelled as such. Built upon Fig. 1 from Schwarzer (2000).

cations in location, size, viewpoint, lighting conditions and background (Bruce & Humphreys, 1994). In addition, we do not just recognize them but are able to describe in detail the type and amount of modification applied to them as well (von Helmholtz, 1868; Cassirer, 1944; Schmidt et al., 2016). Equivariance is strongly related to the idea of *symmetry*. As these modifications do not modify the essence of the underlying object, they should be treated (and learned) as a single concept. Recently, several approaches have embraced these ideas to preserve symmetries including translations (LeCun et al., 1989), planar rotations (Dieleman et al., 2016; Marcos et al., 2017; Worrall et al., 2017; Weiler et al., 2018b; Li et al., 2018; Cheng et al., 2018; Hooeboom et al., 2018; Bekkers et al., 2018; Veeling et al., 2018; Lenssen et al., 2018; Smets et al., 2020), spherical rotations (Cohen et al., 2018; Worrall & Brostow, 2018; Weiler et al., 2018a; Thomas et al., 2018; Cohen et al., 2019b), scaling (Marcos et al., 2018; Worrall & Welling, 2019; Sosnovik et al., 2020) and general symmetry groups (Cohen & Welling, 2016a; Kondor & Trivedi, 2018; Weiler & Cesa, 2019; Cohen et al., 2019a; Bekkers, 2020; Romero & Hoogendoorn, 2020; Venkataraman et al., 2020).

While group convolutional networks are able to learn powerful representations based on symmetry patterns, they lack any explicit means to learn meaningful relationships among them, e.g., relative positions, orientations and scales (Fig. 1). In this paper, we draw inspiration from another promising development in the machine learning domain driven by neuroscience and psychology (e.g., Pashler (2016)), *attention*, to learn such relationships. The notion of attention is related to the idea that not all components of an input signal are *per se* equally relevant for a particular task. As a consequence,

given a task and a particular input signal, task-relevant components of the input should be focused during its analysis while irrelevant, possibly misleading ones should be suppressed. Attention has been broadly applied to fields ranging from natural language processing (Bahdanau et al., 2014; Cheng et al., 2016; Vaswani et al., 2017) to visual understanding (Xu et al., 2015; Ilse et al., 2018; Park et al., 2018; Woo et al., 2018; Ramachandran et al., 2019; Diaconu & Worrall, 2019; Romero & Hoogendoorn, 2020) and graph analysis (Veličković et al., 2017; Zhang et al., 2020).

Specifically, we present *attentive group convolutions*, a generalization of the group convolution, in which attention is applied during convolution to accentuate meaningful symmetry combinations and suppress non-plausible, possibly misleading ones. We indicate that prior work on visual attention can be described as special cases of our proposed framework and show empirically that our *attentive group equivariant group convolutional networks* consistently outperform conventional group equivariant ones on rot-MNIST and CIFAR-10 for the $SE(2)$ and $E(2)$ groups. In addition, we provide means to interpret the learned concepts through the visualization of the predicted equivariant attention maps.

Contributions:

- We propose a general group theoretical framework for equivariant visual attention, *the attentive group convolution*, and show that prior works on visual attention are special cases of our framework.
- We introduce a specific type of network referred to as *attentive group convolutional networks* as an instance of this theoretical framework.
- We show that our *attentive group convolutional networks* consistently outperform plain group equivariant ones.
- We provide means to interpret the learned concepts via visualization of the predicted equivariant attention maps.

2. Preliminaries

Before describing our approach, we first define crucial prior concepts: (group) convolutions and attention mechanisms.

2.1. Spatial Convolution and Translation Equivariance

Let $f, \psi : \mathbb{R}^d \rightarrow \mathbb{R}^{N_{\bar{c}}}$ be a vector valued signal and filter on \mathbb{R}^d , such that $f = \{f_{\bar{c}}\}_{\bar{c}=1}^{N_{\bar{c}}}$ and $\psi = \{\psi_{\bar{c}}\}_{\bar{c}=1}^{N_{\bar{c}}}$. The spatial convolution ($\star_{\mathbb{R}^d}$) is defined as:

$$[f \star_{\mathbb{R}^d} \psi](y) = \sum_{\bar{c}=1}^{N_{\bar{c}}} \int_{\mathbb{R}^d} f_{\bar{c}}(x) \psi_{\bar{c}}(x - y) dx \quad (1)$$

Intuitively, Eq. 1 resembles a collection of \mathbb{R}^d inner products between the input signal f and y -translated versions of ψ . Since the continuous integration in Eq. 1 is usually performed on signals and filters captured in a discrete grid \mathbb{Z}^d ,

the integral on \mathbb{R}^d is reduced to a sum on \mathbb{Z}^d . In our derivations, however, we stick to the continuous case as to guarantee the validity of our theory for techniques defined on continuous spaces, e.g., steerable and Lie group convolutions (Cohen & Welling, 2016b; Worrall et al., 2017; Bekkers et al., 2018; Weiler et al., 2018b;a; Thomas et al., 2018; Weiler & Cesa, 2019; Bekkers, 2020; Sosnovik et al., 2020).

To study (and generalize) the properties of the convolution, we rewrite Eq. 1 using the translation operator \mathcal{L}_y :

$$[f \star_{\mathbb{R}^d} \psi](y) = \sum_{\bar{c}=1}^{N_{\bar{c}}} \int_{\mathbb{R}^d} f_{\bar{c}}(x) \mathcal{L}_y[\psi_{\bar{c}}](x) dx \quad (2)$$

where $\mathcal{L}_y[\psi_{\bar{c}}](x) = \psi_{\bar{c}}(x - y)$. Note that the translation operator \mathcal{L}_y is indexed by an amount of translation y . Resultantly, we actually consider a set of operators $\{\mathcal{L}_y\}_{y \in \mathbb{R}^d}$ that indexes the set of all possible translations $y \in \mathbb{R}^d$. A fundamental property of the convolution is that it commutes with translations:

$$\mathcal{L}_y[f \star_{\mathbb{R}^d} \psi](x) = [\mathcal{L}_y[f] \star_{\mathbb{R}^d} \psi](x), \quad x, y \in \mathbb{R}^d. \quad (3)$$

In other words, convolving a y -translated signal $\mathcal{L}_y[f]$ with a filter is equivalent to first convolving the original signal f with the filter ψ , and y -translating the obtained response next. This property is referred to as *translation equivariance* and, in fact, convolution (and reparametrizations thereof) is the *only linear translation equivariant mapping* (Kondor & Trivedi, 2018; Cohen et al., 2019a; Bekkers, 2020).

2.2. Group Convolution and Group Equivariance

The convolution operation can be extended to general transformations by utilizing a larger set of transformations $\{\mathcal{L}_g\}_{g \in G}$, s.t. $\{\mathcal{L}_y\}_{y \in \mathbb{R}^d} \subseteq \{\mathcal{L}_g\}_{g \in G}$. However, in order to preserve equivariance, we must restrict the class of transformations allowed in $\{\mathcal{L}_g\}_{g \in G}$. To formalize this intuition, we first present some important concepts from *group theory*.

2.2.1. PRELIMINARIES FROM GROUP THEORY

Groups. A *group* is a tuple (G, \cdot) consisting of a set G , $g \in G$, and a binary operation $\cdot : G \times G \rightarrow G$, referred to as the *group product*, that satisfies the following axioms:

- *Closure:* For all $h, g \in G$, $h \cdot g \in G$.
- *Identity:* There exists an $e \in G$, such that $e \cdot g = g \cdot e = g$.
- *Inverse:* For all $g \in G$, there exists an element $g^{-1} \in G$, such that $g \cdot g^{-1} = g^{-1} \cdot g = e$.
- *Associativity:* For all $g, h, k \in G$, $(g \cdot h) \cdot k = g \cdot (h \cdot k)$.

Group actions. Let G and X be a group and a set, respectively. The (left) *group action* of G on X is a function $\odot : G \times X \rightarrow X$ that satisfies the following axioms:

- *Identity:* If e is the identity of G , then, for any $x \in X$, $e \odot x = x$.

- *Compatibility*: For all $g, h \in G, x \in X, g \odot (h \odot x) = (g \cdot h) \odot x$.

In other words, the action of G on X describes how the elements $x \in X$ are transformed by $g \in G$. For brevity, we omit the operations \cdot and \odot and refer to the set G as a group, to elements $g \cdot h$ as gh and to actions $(g \odot x)$ as gx .

Semi-direct product and affine groups. In practice, one is mainly interested in the analysis of data (and hence convolutions) defined on \mathbb{R}^d . Consequently, groups of the form $G = \mathbb{R}^d \rtimes H$, resulting from the *semi-direct product* (\rtimes) between the translation group \mathbb{R}^d and an arbitrary (Lie) group H that acts on \mathbb{R}^d (e.g., rotation, scaling, mirroring), are of main interest. This family of groups is referred to as *affine groups* and their group product is defined as:

$$g_1 g_2 = (x_1, h_1)(x_2, h_2) = (x_1 + h_1 x_2, h_1 h_2) \quad (4)$$

where $g_1 = (x_1, h_1), g_2 = (x_2, h_2) \in G, x_1, x_2 \in \mathbb{R}^d$ and $h_1, h_2 \in H$. Some important affine groups are the roto-translation ($SE(d) = \mathbb{R}^d \rtimes SO(d)$), the scale-translation ($\mathbb{R}^d \rtimes \mathbb{R}^+$) and the euclidean ($E(d) = \mathbb{R}^d \rtimes O(d)$) groups.

Group representations. Let G be a group and $\mathbb{L}_2(X)$ be a space of functions defined on some vector space X . The (left) regular *group representation* of G on functions $f \in \mathbb{L}_2(X)$ is a transformation $\mathcal{L} : G \times \mathbb{L}_2(X) \rightarrow \mathbb{L}_2(X), (g, f) \mapsto \mathcal{L}_g[f]$, such that it shares the group structure via:

$$\mathcal{L}_g \mathcal{L}_h[f](x) = \mathcal{L}_{gh}[f](x) \quad (5)$$

$$\mathcal{L}_g[f](x) := f(g^{-1}x) \quad (6)$$

for any $g, h \in G, f \in \mathbb{L}_2(X), x \in X$. That is, concatenating two such transformations, parametrized by g and h , is equivalent to one transformation parametrized by $gh \in G$. Intuitively, the representation of G on a function $f \in \mathbb{L}_2(X)$ describes how the function as a whole, i.e., $f(x), \forall x \in X$, is transformed by the effect of group elements $g \in G$.

If the group G is affine, i.e., $G = \mathbb{R}^d \rtimes H$, the (left) group representation \mathcal{L}_g can be split as:

$$\mathcal{L}_g[f](x) = \mathcal{L}_y \mathcal{L}_h[f](x) \quad (7)$$

with $g = (y, h) \in G, y \in \mathbb{R}^d$ and $h \in H$. This property is key for the efficient implementation of functions on groups.

2.2.2. THE GROUP CONVOLUTION

Let $f, \psi : G \rightarrow \mathbb{R}^{N_\varepsilon}$ be a vector valued signal and kernel on G . The group convolution (\star_G) is defined as:

$$[f \star_G \psi](g) = \sum_{\tilde{c}=1}^{N_\varepsilon} \int_G f_{\tilde{c}}(\tilde{g}) \psi_{\tilde{c}}(g^{-1}\tilde{g}) d\tilde{g} \quad (8)$$

$$= \sum_{\tilde{c}=1}^{N_\varepsilon} \int_G f_{\tilde{c}}(\tilde{g}) \mathcal{L}_g[\psi_{\tilde{c}}](\tilde{g}) d\tilde{g} \quad (9)$$




Figure 2. Group convolution on the roto-translation group $SE(2)$ for discrete rotations by 90 degrees (also called the $p4$ group). The $p4$ group is defined as $H = \{e, h, h^2, h^3\}$, with h depicting a 90° rotation. The group convolution corresponds to $|H| = 4$ convolutions between the input f and h -transformations of the filter ψ , $\mathcal{L}_h[\psi], \forall h \in H$. Each of these convolutions is equal to the sum over group elements $\tilde{h} \in H$ and channels $\tilde{c} \in [N_\varepsilon]$ of the spatial channel-wise convolutions $[f_{\tilde{c}} \star_{\mathbb{R}^2} \mathcal{L}_h[\psi_{\tilde{c}}]]$ among f and $\mathcal{L}_h[\psi]$.

Differently to Eq. 2, the domain of the signal f , the filter ψ and the group convolution itself $[f \star_G \psi]$ are now defined on the group G .¹ Intuitively, the group convolution resembles a collection of inner products between the input signal f and g -transformed versions of ψ . A key property of the group convolution is that it generalizes equivariance (Eq. 3) to arbitrary groups, i.e., it commutes with g -transformations:

$$\mathcal{L}_{\bar{g}}[f \star_G \psi](g) = [\mathcal{L}_{\bar{g}}[f] \star_G \psi](g), \quad g, \bar{g} \in G. \quad (10)$$

In other words, group convolving a \bar{g} -transformed signal $\mathcal{L}_{\bar{g}}[f]$ with a filter ψ is equivalent to first convolving the original signal f with the filter ψ , and \bar{g} -transforming the obtained response next. This property is referred to as *group equivariance* and, just as for spatial convolutions, the group convolution (or reparametrizations thereof) is the *only* linear G -equivariant map (Kondor & Trivedi, 2018; Cohen et al., 2019a; Bekkers, 2020).

¹Note that Eq. 2 matches Eq. 9 with the substitution $G = \mathbb{R}^d$. It follows that $\mathcal{L}_g[f](x) = f(g^{-1}x) = f(x - y)$, where $g^{-1} = -y$ is the inverse of g in the translation group $(\mathbb{R}^d, +)$ for $g = y$.

Group convolution on affine groups. For affine groups, the group convolution (Eq. 9) can be decomposed, without modifying its properties, by taking advantage of the group structure and the representation decomposition (Eq. 7) as:

$$[f \star_G \psi](g) = \sum_{\tilde{c}=1}^{N_{\tilde{c}}} \int_H \int_{\mathbb{R}^2} f_{\tilde{c}}(\tilde{x}, \tilde{h}) \mathcal{L}_g[\psi_{\tilde{c}}](\tilde{x}, \tilde{h}) d\tilde{x} d\tilde{h} \quad (11)$$

$$= \sum_{\tilde{c}=1}^{N_{\tilde{c}}} \int_H \int_{\mathbb{R}^2} f_{\tilde{c}}(\tilde{x}, \tilde{h}) \mathcal{L}_x \mathcal{L}_h[\psi_{\tilde{c}}](\tilde{x}, \tilde{h}) d\tilde{x} d\tilde{h} \quad (12)$$

where $g = (x, h)$, $\tilde{g} = (\tilde{x}, \tilde{h}) \in G$, $x, \tilde{x} \in \mathbb{R}^d$ and $h, \tilde{h} \in H$. By doing so, the group convolution can be separated into $|H|$ spatial convolutions of the input signal f for each h -transformed filter $\mathcal{L}_h[\psi]$ (Fig. 2):

$$[f \star_G \psi](x, h) = \sum_{\tilde{c}=1}^{N_{\tilde{c}}} \int_H [f_{\tilde{c}} \star_{\mathbb{R}^2} \mathcal{L}_h[\psi_{\tilde{c}}]](x, \tilde{h}) d\tilde{h} \quad (13)$$

Resultantly, the computational cost of a group convolution is roughly equivalent to that of a spatial convolution with a filter bank of size $N_{\tilde{c}} \times |H|$ (Cohen & Welling, 2016a; Worrall & Welling, 2019; Cohen et al., 2019b).

2.3. Attention, Self-Attention and Visual Attention

Attention mechanisms find their roots in recurrent neural network (RNN) based machine translation. Let $\varphi(\cdot)$ be an arbitrary non-linear mapping (e.g., a neural network), $\underline{y} = \{y_j\}_{j=1}^m$ be a sequence of target vectors y_i , and $\underline{x} = \{x_i\}_{i=1}^n$ be a source sequence, whose elements influence the prediction of each value $y_j \in \underline{y}$. In early models (e.g., Kalchbrenner & Blunsom (2013); Cho et al. (2014)), features in the input sequence are aggregated into a context vector $c = \sum_i \varphi(x_i)$ which is used to augment the hidden state in RNN layers. These models assume that source elements x_i contribute equally to every target element y_j and hence, that the same context vector c can be utilized for all target positions y_j , which does not generally hold (Fig. 3).

Bahdanau et al. (2014) proposed the inclusion of *attention coefficients* $\alpha_i = \{\alpha_{i,j}\}$, $[n] = \{1, \dots, n\}$, $i \in [n]$, $j \in [m]$, $\sum_i \alpha_{i,j} = 1$, to modulate the contributions of the source elements x_i as a function of the current target element y_j by means of an adaptive context vector $c_j = \sum_i \alpha_{i,j} \varphi(x_i)$. Thereby, they obtained large improvements both in performance and interpretability. Recently, attention has been extended to several other machine learning tasks (e.g., Vaswani et al. (2017); Veličković et al. (2017); Park et al. (2018)). The main development behind these extensions was *self-attention* (Cheng et al., 2016), where, in contrast to conventional attention, the target and source sequences are equal, i.e., $\underline{x} = \underline{y}$. Consequently, the attention coefficients $\alpha_{i,j}$ encode correlations among input element pairs (x_i, x_j) . For vision tasks, self-attention has been proposed to encode visual co-occurrences in data (Hu et al., 2018; Wang et al.,




Figure 3. English to French translation. Brighter depicts stronger influence. Note how relevant parts of the input sentence are highlighted as a function of the current output word during translation. Taken from Bahdanau et al. (2014).

2018; Park et al., 2018; Woo et al., 2018; Cao et al., 2019; Bello et al., 2019; Ramachandran et al., 2019; Romero & Hoogendoorn, 2020). Unfortunately, its application on visual and, in general, on high-dimensional data is non-trivial.

2.3.1. VISUAL ATTENTION

In the context of visual attention, consider a feature map $f : X \rightarrow \mathbb{R}^{N_c}$ to be the source “sequence”². Self-attention then imposes the learning of a total $n^2 = |X|^2$ attention vectors $\alpha_{i,j} \in \mathbb{R}^{N_c}$, which rapidly becomes unfeasible with increasing feature map size. Interestingly, Cao et al. (2019) and Zhu et al. (2019) empirically demonstrated that, for visual data, the attention coefficients $\{\alpha_{i,j}\}$ are approximately invariant to changes in the target position x_j . Consequently, they proposed to approximate the attention coefficients $\{\alpha_{i,j}\} \in \mathbb{R}^{|X|^2 \times N_c}$ by a single vector $\{\alpha_i\} \in \mathbb{R}^{|X| \times N_c}$ which is independent of target position x_j . Despite this significant reduction in complexity, the dimensionality of $\{\alpha_i\}$ is still very large and further simplifications are mandatory. To this end, existing works (Hu et al., 2018; Woo et al., 2018) replace the input f with a much smaller vector of input statistics s that summarizes relevant information from f .

For instance, the SE-Net (Hu et al., 2018) utilizes global average pooling to produce a vector of channel statistics of f , $s^c \in \mathbb{R}^{N_c}$, $s^c = \frac{1}{|\mathbb{R}^d|} \int_{\mathbb{R}^d} f_{\tilde{c}}(x) dx$, which is subsequently passed to a small fully-connected network $\varphi^c(\cdot)$ to compute channel attention coefficients $\alpha^c = \{\alpha_{\tilde{c}}^c\}_{\tilde{c}=1}^{N_c} = \varphi^c(s^c)$. These attention coefficients are then utilized to modulate the corresponding input channels $f_{\tilde{c}}$.

Complementary to channel attention akin to that of the SE-Net, Park et al. (2018) utilize a similar strategy for spatial attention. Specifically, they utilize channel average pooling to generate a vector of spatial statistics of f , $s^x \in \mathbb{R}^d$, $s^x = \frac{1}{N_{\tilde{c}}} \sum_{\tilde{c}=1}^{N_{\tilde{c}}} f_{\tilde{c}}(x)$, which is subsequently passed to a small convolutional network $\varphi^x(\cdot)$ to compute spatial attention coefficients $\alpha^x = \{\alpha^x(x)\}_{x \in \mathbb{R}^2} = \varphi^x(s^x)$. These attention coefficients are then utilized to modulate the corresponding spatial input positions $f(x)$. Recent works include extra statistical information, e.g., max responses (Woo et al., 2018), or replace pooling by convolutions (Cao et al., 2019).

²In the machine translation context we can think of f as a sequence $\underline{x} = \{f(x_i)\}_{i=1}^n$, with $n = |X|$ number of elements.



Figure 4. Same colors depict equal weights. The first column of \mathcal{A}^C corresponds to ψ and the following ones to $\mathcal{L}_h[\psi]$, obtained via cyclic permutations. See how $\{\mathcal{L}_h[\psi]\}_{h \in H}$ resembles a circulant matrix. Taken from Romero & Hoogendoorn (2020).

3. Attentive Group Equivariant Convolution

In this section, we propose our generalization of visual self-attention, discuss its properties and relations to prior work.

Let $f, \psi : G \rightarrow \mathbb{R}^{N_{\tilde{e}}}$ be a vector valued signal and kernel on G , and let $\alpha : G \times G \rightarrow [0, 1]^{N_{\tilde{e}}}$ be an *attention map* that takes target and source elements $g, \tilde{g} \in G$, respectively, as input. We define the *attentive group convolution* (\star_G^α) as:

$$[f \star_G^\alpha \psi](g) = \sum_{\tilde{c}=1}^{N_{\tilde{e}}} \int_G \alpha_{\tilde{c}}(g, \tilde{g}) f_{\tilde{c}}(\tilde{g}) \mathcal{L}_g[\psi_{\tilde{c}}](\tilde{g}) d\tilde{g} \quad (14)$$

with $\alpha = \mathcal{A}[f]$ computed by some *attention operator* \mathcal{A} . As such, the attentive group convolution modulates the contributions of group elements $\tilde{g} \in G$ at different channels $\tilde{c} \in [N_{\tilde{e}}]$ during pooling.³ The properties and conditions on \mathcal{A} are summarized in Thm. 1. An extensive motivation as well as its proof are provided in the supplementary material.

Theorem 1. *The attentive group convolution is an equivariant operator if and only if the attention operator \mathcal{A} satisfies:*

$$\forall_{\tilde{g}, g, \tilde{g} \in G} : \mathcal{A}[\mathcal{L}_{\tilde{g}} f](g, \tilde{g}) = \mathcal{A}[f](\tilde{g}^{-1}g, \tilde{g}^{-1}\tilde{g}) \quad (15)$$

If, moreover, the maps generated by \mathcal{A} are invariant to one of its arguments, and, hence, exclusively attend to either the input or the output domain (Sec. 3.4), then \mathcal{A} satisfies Eq. 15 iff it is equivariant and thus, based on group convolutions.

3.1. Tying Together Equivariance and Visual Attention

Interestingly, and, perhaps in some cases unaware of it, *all* of the visual attention approaches outlined in Section 2.3.1, as well as all of those we are aware of (Xu et al., 2015; Hu et al., 2018; Park et al., 2018; Woo et al., 2018; Wang et al., 2018; Ilse et al., 2018; Hu et al., 2019; Ramachandran et al., 2019; Cao et al., 2019; Chen et al., 2019; Bello et al., 2019; Lin et al., 2019; Diaconu & Worrall, 2019; Romero & Hoogendoorn, 2020) *exclusively utilize translation (or group) equivariance preserving maps for the generation of the attention coefficients and, hence, constitute altogether group equivariant networks by which they satisfy Thm. 1.*

As will be explained in the following sections, all these works resemble special cases of Eq. 14 by substituting G with the corresponding group and modifying the specifications about how α is calculated (Sec. 3.2 - 3.4).

³Note that Eq. 14 is equal to Eq. 9 up to a multiplicative factor $\alpha_{\tilde{c}}(g, \tilde{g})^{-1}$, if $\alpha_{\tilde{c}}(g, \tilde{g})$ is constant for every $g, \tilde{g} \in G, \tilde{c} \in [N_{\tilde{e}}]$.

3.1.1. TRANSLATION EQUIVARIANT VISUAL ATTENTION

Since convolutions as well as popular pooling operations are translation equivariant, the visual attention approaches outlined in Sec. 2.3.1 are translation equivariant as well.⁴ One particular case worth emphasising is that of SE-Nets. Here, a fully-connected network φ^C , a non-translation equivariant map, is used to generate the channel attention coefficients α^C . However, φ^C is indeed translation equivariant. Recall that φ^C receives s^C as input, a signal obtained via global average pooling (a convolution-like operation). Resultantly, s^C can be interpreted as a $\mathbb{R}^{N_{\tilde{e}} \times 1 \times 1}$ tensor and hence, applying a fully connected layer to s^C equals a pointwise convolution between s^C and a filter $\psi_{\text{fully}} \in \mathbb{R}^{N_o \times N_{\tilde{e}} \times 1 \times 1}$ with N_o output channels.⁵

3.1.2. GROUP EQUIVARIANT VISUAL ATTENTION

To the best of our knowledge, the only work that provides a group theoretical approach towards visual attention is that of Romero & Hoogendoorn (2020). Here, the authors consider affine groups G with elements $g = (x, h)$, $x \in \mathbb{R}^d$, $h \in H$ and cyclic permutation groups H . Consequently, they utilize a cyclic permutation equivariant map, $\varphi^H(\cdot)$, to generate attention coefficients $\alpha^H(h)$, $h \in H$, with which the corresponding elements h are modulated. As a result, their proposed attention strategy is H -equivariant. To preserve translation equivariance, and hence, G -equivariance, φ^H is re-utilized at every spatial position $x \in \mathbb{R}^d$. This is equivalent to combining φ^H with a pointwise filter on \mathbb{R}^d . Romero & Hoogendoorn (2020) found that equivariance to cyclic groups H , can *only* be achieved by constraining φ^H to have a *circulant structure*. This is equivalent to a convolution with a filter ψ , whose group representations \mathcal{L}_h induce cyclical permutations of itself (Fig. 4) and hence, resembles a group convolution, by which Thm. 1 is satisfied.

The work of Romero & Hoogendoorn (2020) exclusively performs attention on the h component of the group elements $g = (x, h) \in G$ and is only defined for (block) cyclic groups. Consequently, it does not consider spatial relationships during attention (Fig. 1) and is not applicable to general groups. Conversely, our proposed framework allows for simultaneous attention on both components of the group elements $g = (x, h)$ in a G equivariance preserving manner.

3.2. Efficient Group Equivariant Attention Maps

Attentive group convolutions impose the generation of an additional attention map $\alpha : G \times G \rightarrow [0, 1]^{N_{\tilde{e}}}$, which is computationally demanding. To reduce this computational

⁴In fact, conventional pooling operations (e.g., max, average) can be written as combinations of convolutions and pointwise non-linearities, which are translation equivariant, as well.

⁵This resembles a depth-wise separable convolution (Chollet, 2017) with the first convolution given by global average pooling.

burden, we exploit the fact that visual data is defined on \mathbb{R}^d and, hence, relevant groups are affine, to provide an efficient factorization of the attention map α .

In Sec. 2.3.1 we indicated that attention coefficients α can be equivariantly factorized into spatial and channel components. We build upon this idea and factorize attention via:

$$\alpha_{\tilde{e}}(g, \tilde{g}) := \alpha^{\mathcal{X}}((x, h), (\tilde{x}, \tilde{h}))\alpha_{\tilde{e}}^{\mathcal{C}}(h, \tilde{h})$$

where $\alpha^{\mathcal{X}}$ attends for spatial relations without considering channel characteristics and $\alpha^{\mathcal{C}}$ attends for patterns in the channel- and H -axis, but ignores spatial patterns. We thus factorize α into a *spatial attention map* $\alpha^{\mathcal{X}} : G \times G \rightarrow [0, 1]$ and a *channel attention map* $\alpha^{\mathcal{C}} : H \times H \rightarrow [0, 1]^{N_{\tilde{e}}}$. Findings in literature have shown that, for visual data, attention maps are almost equivalent for different query positions and thus, only query-independent dependencies are learnt (Cao et al., 2019; Zhu et al., 2019). Based on this observation, we further simplify $\alpha^{\mathcal{X}}$ to be invariant over spatial positions either at the input or output space. Since separate convolutional filters ψ could possibly benefit from different attention maps, we omit spatial positions in the input space (see Sec. 3.2.1 for details). In other words, we replace $\alpha^{\mathcal{X}}(g, \tilde{g})$ with $\alpha^{\mathcal{X}}(g, \tilde{h})$, an spatial position invariant attention map over the input space: $\alpha^{\mathcal{X}} : G \times H \rightarrow [0, 1]$.

Conveniently, attention coefficients of type $\alpha : \mathbb{R}^d \times H \rightarrow [0, 1]^{N_{\tilde{e}}}$ can be interpreted as functions on \mathbb{R}^d with pointwise visualizations $\tilde{x} \mapsto \alpha(\tilde{x}, \tilde{h})$ for each $\tilde{x} \in \mathbb{R}^d$. Resultantly, we are able to aid the interpretability of the learned concepts and of the attended symmetries (e.g., Figs. 7, 8, 11).

3.2.1. THE ATTENTION OPERATOR \mathcal{A}

Recall that the attention map α is computed via an attention operator \mathcal{A} . In the most general case, α and, hence \mathcal{A} , is a function of both the input signal f and the filter ψ . In order to define \mathcal{A} as such, we generalize the approach of Woo et al. (2018) such that: (1) equivariance to general symmetry groups is preserved and (2) the attention maps depend on the filter ψ as well.

Let $\phi^{\mathcal{C}} : \tilde{f} \mapsto s^{\mathcal{C}} = \{s_{\text{avg}}^{\mathcal{C}}, s_{\text{max}}^{\mathcal{C}}\}$, $s_i^{\mathcal{C}} : H \times H \rightarrow \mathbb{R}^{N_{\tilde{e}}}$ and $\phi^{\mathcal{X}} : \tilde{f} \mapsto s^{\mathcal{X}} = \{s_{\text{avg}}^{\mathcal{X}}, s_{\text{max}}^{\mathcal{X}}\}$, $s_i^{\mathcal{X}} : G \times G \rightarrow \mathbb{R}$ be functions that generate channel ($s^{\mathcal{C}}$) and spatial statistics ($s^{\mathcal{X}}$), respectively, from an intermediary vector valued signal $\tilde{f} : G \times G \rightarrow \mathbb{R}^{N_{\tilde{e}}}$ containing information both from the input and output spaces. Analogously to Woo et al. (2018), we compute spatial and channel statistics to reduce the dimensionality of the input. However, in contrast to them, we compute these statistics from intermediary convolutional maps \tilde{f} rather than from the input signal f directly.⁶ As a result,

⁶This is why the statistics $s_i^{\mathcal{C}}$, $s_i^{\mathcal{X}}$ receive tuples (h, \tilde{h}) , (g, \tilde{g}) , respectively, as input, as opposed to single argument inputs which often emerge in several prior works on visual attention.

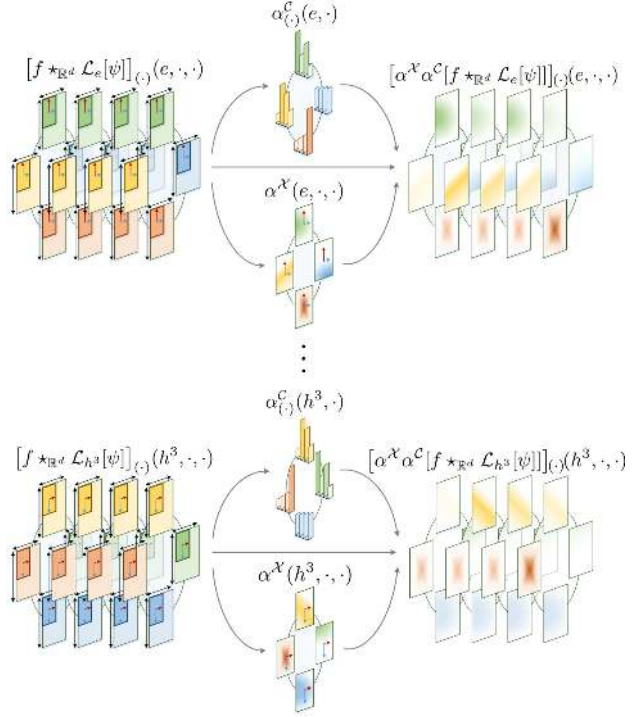


Figure 5. Attentive group convolution on the roto-translation group $SE(2)$. In contrast to group convolutions (Fig. 2, Eq. 13), attentive group convolutions utilize channel $\alpha^{\mathcal{C}}$ and spatial $\alpha^{\mathcal{X}}$ attention to modulate the intermediary convolutional responses $[f \star_{\mathbb{R}^2} \mathcal{L}_h[\psi]]$ before pooling over the \tilde{e} and \tilde{h} axes.

we take the influence of the filter ψ into account during the computation of the attention maps. Following the simplifications proposed in Sec. 3.2 for $\alpha^{\mathcal{X}}$, we can further reduce $s_i^{\mathcal{X}}$ and \tilde{f} to functions of the form $s_i^{\mathcal{X}} : G \times H \rightarrow \mathbb{R}$ and $\tilde{f} : G \times H \rightarrow \mathbb{R}^{N_{\tilde{e}}}$, respectively. Consequently, we define:

$$\tilde{f} = \{\tilde{f}_{\tilde{e}}\}_{\tilde{e}=1}^{N_{\tilde{e}}}, \tilde{f}_{\tilde{e}}(x, h, \tilde{h}) := [f_{\tilde{e}} \star_{\mathbb{R}^d} \mathcal{L}_h[\psi_{\tilde{e}}]](x, \tilde{h}), \quad (16)$$

which is the intermediary result of the convolution between the input f and the h -transformation of the filter ψ , $\mathcal{L}_h[\psi]$ before pooling over \tilde{e} and \tilde{h} (Fig. 5, Eq. 13).

Channel Attention. Let $\varphi^{\mathcal{C}} : s^{\mathcal{C}} \mapsto \alpha^{\mathcal{C}}$ be a function that generates a channel attention map $\alpha^{\mathcal{C}} : H \times H \rightarrow [0, 1]^{N_{\tilde{e}}}$ from a vector of channel statistics $s^{\mathcal{C}} : H \times H \rightarrow \mathbb{R}^{N_{\tilde{e}}}$ of the intermediate representation \tilde{f} . Our channel attention computation is analogous to that of Woo et al. (2018) based on two fully connected layers. However, in our case, each linear layer is parametrized by a *matrix-valued kernel* $\mathbf{W}_i : H \rightarrow \mathbb{R}^{N_{\text{out}} \times N_{\text{in}}}$, which we shift via left-regular representations $\mathcal{L}_h[\mathbf{W}_i](\tilde{h}) = \mathbf{W}_i(h^{-1}\tilde{h})$ in order to guarantee equivariance (Thm. 1):

$$\begin{aligned} \alpha^{\mathcal{C}}(h, \tilde{h}) &= \varphi^{\mathcal{C}}[s^{\mathcal{C}}](h, \tilde{h}) \\ &= \sigma\left([\mathbf{W}_2(h^{-1}\tilde{h}) \cdot [\mathbf{W}_1(h^{-1}\tilde{h}) \cdot s_{\text{avg}}^{\mathcal{C}}(h, \tilde{h})]^+] \right. \\ &\quad \left. + [\mathbf{W}_2(h^{-1}\tilde{h}) \cdot [\mathbf{W}_1(h^{-1}\tilde{h}) \cdot s_{\text{max}}^{\mathcal{C}}(h, \tilde{h})]^+] \right) \end{aligned} \quad (17)$$

with $[\cdot]^+$ the ReLU function, σ the sigmoid function, r a reduction ratio and $\mathbf{W}_1 : H \rightarrow \mathbb{R}^{\frac{N_c}{r} \times N_c}$, $\mathbf{W}_2 : H \rightarrow \mathbb{R}^{N_c \times \frac{N_c}{r}}$ filters defined on H .

Spatial Attention. Let $\varphi^{\mathcal{X}} : s^{\mathcal{X}} \mapsto \alpha^{\mathcal{X}}$ be a function that generates a spatial attention map $\alpha^{\mathcal{X}} : G \times H \rightarrow [0, 1]$ from channel statistics $s^{\mathcal{X}} : G \times H \rightarrow \mathbb{R}^2$, in which per input $\tilde{h} \in H$ and output $g \in G$, the mean and max value is taken over the channel axis. Similarly to Woo et al. (2018), spatial attention $\alpha^{\mathcal{X}}$ is then defined as:

$$\begin{aligned} \alpha^{\mathcal{X}}(x, h, \tilde{h}) &= \varphi^{\mathcal{X}}(s^{\mathcal{X}})(x, h, \tilde{h}) \\ &= \sigma \left([s^{\mathcal{X}} \star_{\mathbb{R}^d} \mathcal{L}_h[\psi^{\mathcal{X}}]] \right) (x, \tilde{h}) \end{aligned} \quad (18)$$

with $\psi^{\mathcal{X}} : G \rightarrow \mathbb{R}^2$ a group convolutional filter.

Full Attention. Woo et al. (2018) carried out extensive experiments to find the best performing configuration to combine channel and spatial attention maps for the \mathbb{R}^d case, e.g., in parallel, serially starting with channel attention, serially starting with spatial attention. Based on their results we adopt their best performing configuration, i.e., *serially starting with channel attention*, for the G case (Fig. 6).

Recall that \tilde{f} is the intermediary result from the convolution between the input f and the h -transformation of the filter ψ before pooling over \tilde{c} and \tilde{h} . We perform attention on top of \tilde{f} (Fig. 6), where $\alpha^{\mathcal{C}}$ and $\alpha^{\mathcal{X}}$ are computed by Eqs. 17, 18, respectively. Resultantly, the attentive group convolution is computed as:

$$[f \star_G^\alpha \psi](x, h) = \sum_{\tilde{c}=1}^{N_c} \int_H \alpha^{\mathcal{X}}(x, h, \tilde{h}) \alpha^{\mathcal{C}}(h, \tilde{h}) \tilde{f}(x, h, \tilde{h}) d\tilde{h} \quad (19)$$

3.3. The Residual Attention Branch

Based on the findings of He et al. (2016), several visual attention approaches propose to utilize residual blocks with direct connections during the course of attention to facilitate gradient flow (Hu et al., 2018; Park et al., 2018; Woo et al., 2018; Wang et al., 2018; Cao et al., 2019). However, these approaches calculate the final attention map α^+ as the sum of the direct connection $\mathbf{1}$ and the attention map obtained from the attention branch α , i.e., $\alpha^+ = \mathbf{1} + \alpha$. Consequently, the obtained attention map $\alpha^+ : \mathbb{R}^2 \rightarrow [1, 2]^{N_c}$ is *restricted* to the interval $[1, 2]$ and the network loses its ability to suppress input components. Inspired by the aforementioned works, we propose to calculate attention in what we call a *residual attention branch* (Fig. 6). Specifically, we utilize the attention branch to calculate a *residual attention map* defined as $\alpha^- = (\mathbf{1} - \alpha^+)$; $\alpha^- : G \times G \rightarrow [0, 1]$. Next, we subtract the residual attention map α^- from the direct connection $\mathbf{1}$ to obtain the resultant attention map α^+ , i.e., $\alpha^+ = \mathbf{1} - \alpha^-$. As a result, we are able to produce attention maps α^+ that span the $[0, 1]$ interval while preserving the benefits of the direct connections of He et al. (2016).




Figure 6. Sequential channel and spatial attention performed on a residual attention branch (Sec. 3.3).

3.4. The Attentive Group Convolution as a Sequence of Group Convolutions and Pointwise Non-linearities

CNNs are usually organized in layers and hence, the input f is usually convolved in parallel with a set of N_o filters $\{\psi_o\}_{o=1}^{N_o}$. As outlined in the previous section, this implies that the attention maps can change as a function of the current filter ψ_o . One assumption broadly utilized in visual attention is that these maps do not depend on the filters $\{\psi_o\}_{o=1}^{N_o}$, and, hence, that α is a sole function of the input signal f (Hu et al., 2018; Park et al., 2018; Woo et al., 2018; Diaconu & Worrall, 2019; Romero & Hoogendoorn, 2020). Consequently, the attention coefficients α are reduced from a function $\alpha : G \times G \rightarrow [0, 1]^{N_c}$ (c.f., Eq. 14) to a function $\alpha : G \rightarrow [0, 1]^{N_c}$. In other words, attention becomes only dependent on g (see Eqs. 17-19) and thus, the generation of the attention maps $\alpha^{\mathcal{C}}, \alpha^{\mathcal{X}}$ can be shifted to the input feature map f . Resultantly, the attentive group convolution is reduced to a sequence of conventional group convolutions and point-wise non-linearities (Thm. 1), which further reduces the computational cost of attention:

$$[f \star_G^\alpha \psi] = [f^\alpha \star_G \psi] = [(\alpha^{\mathcal{X}} \alpha^{\mathcal{C}} f) \star_G \psi] \quad (20)$$

4. Experiments

We validate our approach by exploring the effects of using attentive group convolutions in contrast to conventional ones. We compare the conventional group equivariant networks $p4$ - and $p4m$ -CNNs of Cohen & Welling (2016a) on the rotated MNIST and CIFAR-10 datasets with their corresponding attentive counterparts: α - $p4$ -CNNs and α - $p4m$ -CNNs, respectively; and the $p4$ - and $p4m$ -DenseNets of Veeling et al. (2018) on the PCam dataset with their corresponding attentive counterparts: α - $p4$ -DenseNet and α - $p4m$ -CNNs and DenseNets, respectively. Additionally, we explore the effects of only applying channel attention (e.g., α_{CH} - $p4$ -CNNs), spatial attention (e.g., α_{SP} - $p4$ -CNNs) and applying attention directly on the input (e.g., α_{F} - $p4$ -CNNs).⁷

We notice that the network architectures in Cohen & Welling (2016a) and Romero & Hoogendoorn (2020) used for the CIFAR-10 experiments are equivariant only approximately. This results from using odd-sized convolutional kernels with stride ≥ 1 on even-sized feature maps (see Appx. C for a

⁷Our code is publicly available at: https://github.com/dwromero/att_gconvs

complete discussion). Since this effect distorts the equivariance property of our equivariant attention maps, i.e., they also become equivariant only approximately (Figs. 10, 11), this issue must be fixed. We achieve this by replacing strided convolutions in such regimes by conventional convolutions followed by a max-pooling layer.

For all our experiments we replicate as close as possible the training and evaluation strategies of the corresponding baselines, replace approximately equivariant networks by exact equivariant ones, and initialize any additional parameter in the same way as the corresponding baseline. Extended implementation details are provided in Appx. B.

4.1. rot-MNIST

The rotated MNIST dataset (Larochelle et al., 2007) contains $62k$ gray-scale 28×28 handwritten digits uniformly rotated for $[0, 2\pi)$. The dataset is split into training, validation and test sets of $10k$, $2k$ and $50k$ images respectively. We compare $p4$ -CNNs with all the corresponding attention variants previously mentioned. For our attention models, we utilize a filter size of 7 and a reduction ratio r of 2 on the attention branch. Since attentive group convolutions impose the learning of additional parameters, we also instantiate bigger $p4$ -CNNs by increasing the number of channels uniformly at every layer to roughly match the number of parameters of the attentive versions. Furthermore, we compare our results with comparative attentive versions as defined in Romero & Hoogendoorn (2020) (α_{RH}), which perform attention exclusively over the axis of rotations. Our results show that (1) attentive versions consistently outperform non-attentive ones, and that (2) performing attention over the entire group is beneficial in terms of classification accuracy (Tab. 1).

4.2. CIFAR-10

The CIFAR-10 dataset (Krizhevsky et al., 2009) consists of $60k$ real-world 32×32 RGB images uniformly drawn from 10 classes. The dataset is split into training, validation and test sets of $40k$, $10k$ and $10k$ images, respectively. We compare the $p4$ and $p4m$ versions of the All-CNN (Springenberg et al., 2014) and the Resnet44 (He et al., 2016) in Cohen & Welling (2016a) with attentive variations. For all our attention models, we utilize a filter size of 7 and a reduction ratio r of 16 on the attention branch. Unfortunately, attentive group convolutions impose an unfeasible increment on the memory requirements for this dataset.⁸ Resultantly, we are only able to compare the α_F variations of the corresponding networks. Our results show that attentive α_F networks consistently outperform non-attentive ones (Tab. 2). Moreover,

⁸the $\alpha-p4$ All-CNN requires approx. 72GB of CUDA memory, as opposed to 5GBs for the $p4$ -All-CNN. This is due to the storage of the intermediary convolution responses required for the calculation of the attention weights (Eqs. 17-19)




Figure 7. Equivariant attention maps on the roto-translation group $SE(2)$. The predicted attention maps behave equivariantly for group symmetries. The arrows depict the strength of the filter responses at the corresponding orientations throughout the network.

Table 1. Test error rates on rot-MNIST (with standard deviation under 5 random seed variations).

| NETWORK | TEST ERROR (%) | PARAM. |
|-------------------------------|-------------------------------------|--------|
| $p4$ -CNN | 2.048 ± 0.045 | 24.61K |
| α_{RH} - $p4$ -CNN | 1.980 ± 0.032 | 24.85K |
| BIG ₁₉ - $p4$ -CNN | 1.796 ± 0.035 | 77.54K |
| α - $p4$ -CNN | 1.696 ± 0.021 | 73.13K |
| BIG ₁₅ - $p4$ -CNN | 1.848 ± 0.019 | 50.42K |
| α_{CH} - $p4$ -CNN | 1.825 ± 0.048 | 48.63K |
| α_{SP} - $p4$ -CNN | 1.761 ± 0.027 | 49.11K |
| BIG ₁₁ - $p4$ -CNN | 1.996 ± 0.083 | 29.05K |
| α_F - $p4$ -CNN | 1.795 ± 0.028 | 29.46K |

Table 2. Test error rates on CIFAR10 and augmented CIFAR10+.

| NETWORK | TYPE | CIFAR10 | CIFAR10+ | PARAM. |
|----------|--------------------|--------------|--------------|--------|
| ALL-CNN | $p4$ | 9.32 | 8.91 | 1.37M |
| | α_F - $p4$ | 8.8 | 7.05 | 1.40M |
| | $p4m$ | 7.61 | 7.48 | 1.22M |
| | α_F - $p4m$ | 6.93 | 6.53 | 1.25M |
| RESNET44 | $p4m$ | 15.72 | 15.4 | 2.62M |
| | α_F - $p4m$ | 10.82 | 10.12 | 2.70M |

we demonstrate that our proposed networks focus on relevant parts of the input and that the predicted attention maps behave equivariantly for group symmetries (Figs. 7, 11).

4.3. PCam

The PatchCamelyon dataset (Veeling et al., 2018) consists of $327k$ 96×96 RGB image patches of tumorous/non-tumorous breast tissues extracted from the Camelyon16 dataset (Bejnordi et al., 2017), where each patch was labelled as tumorous if the central region (32×32) contained at least one tumour pixel as given by the original annotation in Bejnordi et al. (2017). We compare the $p4$ and $p4m$ versions of the DenseNet (Huang et al., 2017) in Veeling et al. (2018) with attentive variants. For all our attention models, we utilize a filter size of 7 and a reduction ratio r of 16 on the attention branch. Similarly to the CIFAR-10 case, we restrict our




Figure 8. Equivariant attention maps on the PCam dataset. The predicted attention maps behave equivariantly for group symmetries. Additionally, the network seems to learn to focus on the nuclei of the cells and remove background elements during training.

Table 3. Test error rates on PCam.

| NETWORK | TYPE | TEST ERROR (%) | PARAM. |
|----------|----------------|----------------|---------|
| DENSENET | \mathbb{Z}^2 | 15.93 | 130.60K |
| | $p4$ | 12.45 | 129.65K |
| | α_F-p4 | 11.34 | 140.45K |
| | $p4m$ | 11.64 | 124.21K |
| | α_F-p4m | 10.88 | 141.22K |

experiments to α_F attentive networks due to computational constraints. Our results show that attentive α_F consistently outperform non-attentive ones (Tab. 3). Interestingly, the α_F-p4 -DenseNet is already able to outperform the $p4m$ -DenseNet without attention. Surprisingly, our equivariant attention maps reveal that the network learns to focus on the nuclei of the cells and to removes background elements during inference, all of this in a group equivariant way (Fig. 8).

5. Discussion and Future Work

Our results show that attentive group convolutions can be utilized as a drop-in replacement for standard and group equivariant convolutions that simultaneously facilitates the interpretability of the network decisions. Similarly to convolutional and group convolutional networks, attentive group convolutional networks also benefit of data augmentation. Interestingly, however, we also see that including additional symmetries reduces the effect of augmentations given by group elements. This finding supports the intuition that symmetry variants of the same concept are learned independently for non-equivariant networks (see Fig. 2 in (Krizhevsky et al., 2012)). The main shortcoming of our approach is its computational burden. As a result, the application of α -networks is computationally unfeasible for networks with several layers or channels. We believe, however, by extrapolation of our results on rot-MNIST, that further performance improvements are to be expected for α variations, should hardware requirements suffice.

Group convolutional networks have recently been proven very successful in medical imaging applications (Bekkers et al., 2018; Winkels & Cohen, 2018; Lafarge et al., 2020). Since explainability plays a crucial role here, we believe that our attentive maps could be of high relevance to aid the explainability of the network decisions. Moreover, since our attention maps are guaranteed to be equivariant to transformations in the considered group, it is ensured that the predicted attention maps will be consistent across group symmetries. We believe this to be of crucial importance for rotation invariant tasks. Illustratively, in contrast to vanilla attentive CNNs, a malignant tissue will be ensured to generate consistent attention maps regardless of the orientation at which it has been provided to the network.

In future work, we want to explore ways to reduce the computational cost of full attention networks. If successful, we consider feasible to obtain a direct performance boost over our CIFAR-10 and PCam experimental results, without extensive additional memory requirements. Furthermore, we want to extend our work to symmetry groups defined on 3D. By doing so, we expect the range of possible applications of our work to reach several other important applications such as 3D medical imaging applications like CT-scans and other voxel-based representations.

6. Conclusion

We introduced attentive group convolutions, a generalization of the group convolution in which attention is utilized to explicitly highlight meaningful relationships among symmetries. We provided a general mathematical framework for group equivariant visual attention and indicated that prior work on visual attention can be perfectly described as special cases of the attentive group convolution. Our experimental results indicate that attentive group convolutional networks consistently outperform conventional group convolutional ones and additionally provide equivariant attention maps that behave predictively for symmetries of the group, with which learned concepts can be visualized.

Acknowledgements

We gratefully acknowledge our anonymous reviewers for their helpful and valuable commentaries, and Hyunjik Kim for valuable remarks to improve the readability of our paper. This work is part of the Efficient Deep Learning (EDL) programme (grant number P16-25), partly funded by the Dutch Research Council (NWO) and Semiotic Labs, and the research programme VENI (grant number 17290), financed by the Dutch Research Council (NWO). This work was carried out on the Dutch national e-infrastructure with the support of SURF Cooperative.

References

- Bahdanau, D., Cho, K., and Bengio, Y. Neural machine translation by jointly learning to align and translate. *arXiv preprint arXiv:1409.0473*, 2014.
- Bejnordi, B. E., Veta, M., Van Diest, P. J., Van Ginneken, B., Karssemeijer, N., Litjens, G., Van Der Laak, J. A., Hermsen, M., Manson, Q. F., Balkenhol, M., et al. Diagnostic assessment of deep learning algorithms for detection of lymph node metastases in women with breast cancer. *Jama*, 318(22):2199–2210, 2017.
- Bekkers, E. J. B-spline {cnn}s on lie groups. In *International Conference on Learning Representations*, 2020. URL <https://openreview.net/forum?id=H1gBhkBFDH>.
- Bekkers, E. J., Lafarge, M. W., Veta, M., Eppenhof, K. A., Pluim, J. P., and Duits, R. Roto-translation covariant convolutional networks for medical image analysis. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pp. 440–448. Springer, 2018.
- Bello, I., Zoph, B., Vaswani, A., Shlens, J., and Le, Q. V. Attention augmented convolutional networks. *arXiv preprint arXiv:1904.09925*, 2019.
- Biederman, I. Recognition-by-components: a theory of human image understanding. *Psychological review*, 94(2):115, 1987.
- Blake, R. and Lee, S.-H. The role of temporal structure in human vision. *Behavioral and cognitive neuroscience reviews*, 4(1):21–42, 2005.
- Bruce, V. and Humphreys, G. W. Recognizing objects and faces. *Visual cognition*, 1(2-3):141–180, 1994.
- Cao, Y., Xu, J., Lin, S., Wei, F., and Hu, H. Gcnet: Non-local networks meet squeeze-excitation networks and beyond. *arXiv preprint arXiv:1904.11492*, 2019.
- Cassirer, E. The concept of group and the theory of perception. *Philosophy and phenomenological research*, 5(1): 1–36, 1944.
- Chen, Y., Rohrbach, M., Yan, Z., Shuicheng, Y., Feng, J., and Kalantidis, Y. Graph-based global reasoning networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 433–442, 2019.
- Cheng, J., Dong, L., and Lapata, M. Long short-term memory-networks for machine reading. *arXiv preprint arXiv:1601.06733*, 2016.
- Cheng, X., Qiu, Q., Calderbank, R., and Sapiro, G. Rotdcf: Decomposition of convolutional filters for rotation-equivariant deep networks. *arXiv preprint arXiv:1805.06846*, 2018.
- Cho, K., Van Merriënboer, B., Gulcehre, C., Bahdanau, D., Bougares, F., Schwenk, H., and Bengio, Y. Learning phrase representations using rnn encoder-decoder for statistical machine translation. *arXiv preprint arXiv:1406.1078*, 2014.
- Chollet, F. Xception: Deep learning with depthwise separable convolutions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 1251–1258, 2017.
- Cohen, T. and Welling, M. Group equivariant convolutional networks. In *International conference on machine learning*, pp. 2990–2999, 2016a.
- Cohen, T. S. and Welling, M. Steerable cnns. *arXiv preprint arXiv:1612.08498*, 2016b.
- Cohen, T. S., Geiger, M., Köhler, J., and Welling, M. Spherical cnns. *CoRR*, abs/1801.10130, 2018. URL <http://arxiv.org/abs/1801.10130>.
- Cohen, T. S., Geiger, M., and Weiler, M. A general theory of equivariant cnns on homogeneous spaces. In *Advances in Neural Information Processing Systems*, pp. 9142–9153, 2019a.
- Cohen, T. S., Weiler, M., Kicanaoglu, B., and Welling, M. Gauge equivariant convolutional networks and the icosahedral cnn. *arXiv preprint arXiv:1902.04615*, 2019b.
- Delahunt, C. B. and Kutz, J. N. Insect cyborgs: Bio-mimetic feature generators improve ml accuracy on limited data. 2019.
- Diaconu, N. and Worrall, D. E. Affine self convolution. *arXiv preprint arXiv:1911.07704*, 2019.
- Dieleman, S., De Fauw, J., and Kavukcuoglu, K. Exploiting cyclic symmetry in convolutional neural networks. *arXiv preprint arXiv:1602.02660*, 2016.
- He, K., Zhang, X., Ren, S., and Sun, J. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 770–778, 2016.
- Hoogeboom, E., Peters, J. W., Cohen, T. S., and Welling, M. Hexaconv. In *International Conference on Learning Representations*, 2018. URL <https://openreview.net/forum?id=r1vuQG-CW>.

- Hu, H., Zhang, Z., Xie, Z., and Lin, S. Local relation networks for image recognition. In *Proceedings of the IEEE International Conference on Computer Vision*, pp. 3464–3473, 2019.
- Hu, J., Shen, L., and Sun, G. Squeeze-and-excitation networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 7132–7141, 2018.
- Huang, G., Liu, Z., Van Der Maaten, L., and Weinberger, K. Q. Densely connected convolutional networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 4700–4708, 2017.
- Ilse, M., Tomczak, J. M., and Welling, M. Attention-based deep multiple instance learning. *ICML*, 2018.
- Kalchbrenner, N. and Blunsom, P. Recurrent continuous translation models. In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, pp. 1700–1709, 2013.
- Kingma, D. P. and Ba, J. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- Kondor, R. and Trivedi, S. On the generalization of equivariance and convolution in neural networks to the action of compact groups. *arXiv preprint arXiv:1802.03690*, 2018.
- Krizhevsky, A., Hinton, G., et al. Learning multiple layers of features from tiny images. Technical report, Citeseer, 2009.
- Krizhevsky, A., Sutskever, I., and Hinton, G. E. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pp. 1097–1105, 2012.
- Lafarge, M. W., Bekkers, E. J., Pluim, J. P. W., Duits, R., and Veta, M. Roto-translation equivariant convolutional networks: Application to histopathology image analysis. *arXiv preprint arXiv:2002.08725*, 2020.
- Larochelle, H., Erhan, D., Courville, A., Bergstra, J., and Bengio, Y. An empirical evaluation of deep architectures on problems with many factors of variation. In *Proceedings of the 24th international conference on Machine learning*, pp. 473–480. ACM, 2007.
- LeCun, Y., Boser, B., Denker, J. S., Henderson, D., Howard, R. E., Hubbard, W., and Jackel, L. D. Backpropagation applied to handwritten zip code recognition. *Neural computation*, 1(4):541–551, 1989.
- Lenssen, J. E., Fey, M., and Libuschewski, P. Group equivariant capsule networks. In *Advances in Neural Information Processing Systems*, pp. 8844–8853, 2018.
- Li, J., Yang, Z., Liu, H., and Cai, D. Deep rotation equivariant network. *Neurocomputing*, 290:26–33, 2018.
- Lin, X., Ma, L., Liu, W., and Chang, S.-F. Context-gated convolution, 2019.
- Marcos, D., Volpi, M., Komodakis, N., and Tuia, D. Rotation equivariant vector field networks. In *Proceedings of the IEEE International Conference on Computer Vision*, pp. 5048–5057, 2017.
- Marcos, D., Kellenberger, B., Lobry, S., and Tuia, D. Scale equivariance in cnns with vector fields. *arXiv preprint arXiv:1807.11783*, 2018.
- Park, J., Woo, S., Lee, J.-Y., and Kweon, I. S. Bam: Bottleneck attention module. *arXiv preprint arXiv:1807.06514*, 2018.
- Pashler, H. *Attention*. Psychology Press, 2016.
- Ramachandran, P., Parmar, N., Vaswani, A., Bello, I., Levskaya, A., and Shlens, J. Stand-alone self-attention in vision models. *arXiv preprint arXiv:1906.05909*, 2019.
- Romero, D. W. and Hoogendoorn, M. Co-attentive equivariant neural networks: Focusing equivariance on transformations co-occurring in data. In *International Conference on Learning Representations*, 2020. URL <https://openreview.net/forum?id=r1g6ogrtDr>.
- Schmidt, F., Spröte, P., and Fleming, R. W. Perception of shape and space across rigid transformations. *Vision research*, 126:318–329, 2016.
- Schwarzer, G. Development of face processing: The effect of face inversion. *Child development*, 71(2):391–401, 2000.
- Smets, B., Portegies, J., Bekkers, E., and Duits, R. Pde-based group equivariant convolutional neural networks. *arXiv preprint arXiv:2001.09046*, 2020.
- Sosnovik, I., Szmaja, M., and Smeulders, A. Scale-equivariant steerable networks. In *International Conference on Learning Representations*, 2020. URL <https://openreview.net/forum?id=HJgpugrKPS>.
- Springenberg, J. T., Dosovitskiy, A., Brox, T., and Riedmiller, M. Striving for simplicity: The all convolutional net. *arXiv preprint arXiv:1412.6806*, 2014.
- Thomas, N., Smidt, T., Kearnes, S., Yang, L., Li, L., Kohlhoff, K., and Riley, P. Tensor Field Networks: Rotation-and Translation-Equivariant Neural Networks for 3D Point Clouds. *arXiv preprint arXiv:1802.08219*, 2018.

- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., and Polosukhin, I. Attention is all you need. In *Advances in neural information processing systems*, pp. 5998–6008, 2017.
- Veeling, B. S., Linmans, J., Winkens, J., Cohen, T., and Welling, M. Rotation equivariant cnns for digital pathology. In *International Conference on Medical image computing and computer-assisted intervention*, pp. 210–218. Springer, 2018.
- Veličković, P., Cucurull, G., Casanova, A., Romero, A., Lio, P., and Bengio, Y. Graph attention networks. *arXiv preprint arXiv:1710.10903*, 2017.
- Venkataraman, S. R., Balasubramanian, S., and Sarma, R. R. Building deep equivariant capsule networks. In *International Conference on Learning Representations*, 2020. URL <https://openreview.net/forum?id=BJgNJgSFPS>.
- von Helmholtz, H. Über die Tatsachen, die der Geometrie zugrunde liegen. 1868.
- Wang, X., Girshick, R., Gupta, A., and He, K. Non-local neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 7794–7803, 2018.
- Weiler, M. and Cesa, G. General e (2)-equivariant steerable cnns. In *Advances in Neural Information Processing Systems*, pp. 14334–14345, 2019.
- Weiler, M., Geiger, M., Welling, M., Boomsma, W., and Cohen, T. S. 3d steerable cnns: Learning rotationally equivariant features in volumetric data. In *Advances in Neural Information Processing Systems*, pp. 10381–10392, 2018a.
- Weiler, M., Hamprecht, F. A., and Storath, M. Learning steerable filters for rotation equivariant cnns. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 849–858, 2018b.
- Wertheimer, M. Gestalt theory. 1938.
- Winkels, M. and Cohen, T. S. 3d g-cnns for pulmonary nodule detection. *arXiv preprint arXiv:1804.04656*, 2018.
- Woo, S., Park, J., Lee, J.-Y., and So Kweon, I. Cbam: Convolutional block attention module. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pp. 3–19, 2018.
- Worrall, D. and Brostow, G. Cubenet: Equivariance to 3d rotation and translation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pp. 567–584, 2018.
- Worrall, D. E. and Welling, M. Deep scale-spaces: Equivariance over scale. *arXiv preprint arXiv:1905.11697*, 2019.
- Worrall, D. E., Garbin, S. J., Turmukhambetov, D., and Brostow, G. J. Harmonic networks: Deep translation and rotation equivariance. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 5028–5037, 2017.
- Xu, K., Ba, J., Kiros, R., Cho, K., Courville, A., Salakhudinov, R., Zemel, R., and Bengio, Y. Show, attend and tell: Neural image caption generation with visual attention. In *International conference on machine learning*, pp. 2048–2057, 2015.
- Zhang, R., Zou, Y., and Ma, J. Hyper-{saggn}: a self-attention based graph neural network for hypergraphs. In *International Conference on Learning Representations*, 2020. URL <https://openreview.net/forum?id=ryeHuJBtPH>.
- Zhaoping, L. *The VI hypothesis—creating a bottom-up saliency map for preattentive selection and segmentation*, pp. 189–314. 05 2014. ISBN 9780199564668. doi: 10.1093/acprof:oso/9780199564668.003.0005.
- Zhu, X., Cheng, D., Zhang, Z., Lin, S., and Dai, J. An empirical study of spatial attention mechanisms in deep networks. *arXiv preprint arXiv:1904.05873*, 2019.