

---

# Hierarchical Attribute CNNs

---

Jörn-Henrik Jacobsen<sup>1</sup> Edouard Oyallon<sup>2</sup>  
Stéphane Mallat<sup>2</sup> Arnold W.M. Smeulders<sup>1</sup>

## Abstract

Deep neural network algorithms are difficult to analyze because they lack structure allowing to understand the properties of underlying transforms and invariants. Multiscale Hierarchical Convolutional Networks are a theoretical class of structured deep convolutional networks that constitute a framework to understand neural network classification properties. However, a naive implementation of such networks is infeasible due to the exponential growth of parameters with depth, so their empirical properties remain to be studied. We introduce a subclass of them that overcomes this issue, called Hierarchical Attribute Convolutional Networks, where individual layers are indexed by progressively higher-dimensional and increasingly invariant attributes. Each new layer is computed with multidimensional convolutions along spatial and attribute variables. The dimensionality is kept constant by averaging intermediate layers along attributes, allowing to control the size of the layers while reducing their parameters. This permits to train Hierarchical Attribute CNNs on CIFAR image databases where they obtain comparable accuracy to state of the art networks while having much fewer trainable parameters. We study some properties of the attributes learned from Cifar-10.

## 1. Introduction

Deep convolutional neural networks have demonstrated impressive performance for classification and regression tasks over a wide range of generic problems including images, audio signals, but also for game strategy, biological, medical, chemical and physics data (LeCun et al., 2015). However, their mathematical properties remain mysterious and

---

<sup>1</sup>Informatics Institute, University of Amsterdam, Amsterdam  
<sup>2</sup>Département Informatique, Ecole Normale Supérieure, Paris.  
Correspondence to: Jörn-Henrik Jacobsen <j.jacobsen@uva.nl>.

*Presented at the ICML 2017 Workshop on Principled Approaches to Deep Learning, Sydney, Australia, 2017. Copyright 2017 by the author(s).*

we are currently not able to relate their performance to the properties of the classification problem.

Classifying signals in high dimension requires to eliminate non-informative variables, and hence contract and reduce space dimensionality in appropriate directions. Convolutional Neural Networks (CNNs) discover these directions via backpropagation algorithms (LeCun et al., 1989). Several studies show numerically that linearization increases with depth (Zeiler & Fergus, 2014), but we do not know what type of information is preserved or eliminated. The variabilities which can be eliminated are mathematically defined as the group of symmetries of the classification function (Mallat, 2016). It is the group of transformations which preserves the labels of the classification problem. Translations usually belong to the symmetry group, and invariants to translations are computed with spatial convolutions, followed by a final averaging.

Much recent work has focused on incorporating geometrical knowledge and the notion of symmetries into deep network representations. Either in a fully handcrafted way (Bruna & Mallat, 2013; Oyallon & Mallat, 2015; Cohen & Welling, 2016; Worrall et al., 2016), or by learning invariance with respect to known geometrical transformations (Bruna et al., 2013; Gens & Domingos, 2014; Dai et al., 2017; Jaderberg et al., 2015; Jacobsen et al., 2017). However, understanding a deep neural network classifier requires specifying its symmetry group and invariants besides low-dimensional transformations, especially of non-geometrical nature.

To achieve this goal, we introduce Hierarchical Attribute CNNs, a computationally feasible subclass of Multiscale Hierarchical CNNs. The latter are a generic theoretical class of deep network algorithms introduced in (Mallat, 2016). Multiscale Hierarchical CNNs give explicit information on invariants by disentangling, progressively more signal attributes as depth increases. The CNN operators are multidimensional convolutions along attribute indices. As such, the number of parameters of the operators increases exponentially with depth of the network, making it challenging to learn general Multiscale Hierarchical CNNs in practice. We overcome this issue with our proposed Hierarchical Attribute CNN by indexing layers with progressively higher-dimensional attributes that are *increasingly*

*invariant* with respect to attributes not discriminative for the classification problem. This additional constraint is incorporated via cascades of intermediary averaging along attributes, permitting to keep the dimensionality constant as depth increases and allowing to control the number of trainable parameters.

Section 2 reviews important properties of generic CNN architectures (LeCun et al., 2015). Section 3 describes generic Multiscale Hierarchical CNNs, which are particular CNNs where linear operators are multidimensional convolutions along progressively more attributes. Section 4 describes Hierarchical Attribute CNNs, which reduce inner layers dimensions by computing invariants with an averaging along attributes, making it feasible to train Hierarchical CNNs in practice. Numerical experiments on the CIFAR database show that our proposed Hierarchical Attribute CNN obtains comparable performances to state of the art CNN architectures, while having organized channel dimensions (illustrated in figure 1) and a reduced number of parameters.

Our Hierarchical Attribute CNN (HCNN) presents the first empirical result of a Hierarchical CNN that learns to represent symmetry groups as multidimensional translations along non-linear attributes via generic organization of the channel dimensions in every layer. In section 5 we study the organization obtained by this deep network. The presented architecture provides a mathematical and experimental framework to understand deep neural network classification properties. The numerical results are reproducible and code is available online <sup>1</sup>.

## 2. Deep Convolutional Networks and Group Invariants

A classification problem associates a class  $y = f(x)$  to any vector  $x \in \mathbb{R}^N$  of  $N$  parameters. Deep convolutional networks transforms  $x$  into multiple layers  $x_j$  of coefficients at depths  $j$ , whose dimensions are progressively reduced after a certain depth (LeCun et al., 2010). We briefly review their properties.

We shall numerically concentrate on color images  $x(u, v)$  where  $u = (u_1, u_2)$  are the spatial coordinates and  $1 \leq v \leq 3$  is the index of a color channel. The input  $x(u, v)$  may, however, correspond to any other type of signals. For sounds,  $u = u_1$  is time and  $v$  may be the index of audio channels recorded at different spatial locations.

Each layer is an array of signals  $x_j(u, v)$  where  $u$  is the native index of  $x$ , and  $v$  is a 1-dimensional channel parameter. A deep convolutional network iteratively computes  $x_{j+1} = \rho W_{j+1} x_j$  with  $x_0 = x$ . Each  $W_{j+1}$  computes

sums over  $v$  of convolutions along  $u$ , with filters of small support. It usually also incorporates a batch normalization (Ioffe & Szegedy, 2015). The resolution of  $x_j(u, v)$  along  $u$  is progressively reduced by a subsampling as  $j$  increases until an averaging in the final output layer. The operator  $\rho(z)$  is a pointwise non-linearity. In this work, we shall use exponential linear units ELU (Clevert et al., 2015). It transforms each coefficient  $z(t)$  plus a bias  $c = z(t) + b$  into  $c$  if  $c < 0$  and  $e^c - 1$  if  $c \geq 0$ .

As the depth increases, the discriminative variations of  $x$  along  $u$  are progressively transferred to the channel index  $v$ . At the last layer  $x_J$ ,  $v$  stands for the class index and  $u$  has disappeared. An estimation  $\tilde{y}$  of the signal class  $y = f(x)$  is computed by applying a soft-max to  $x_J(v)$ . It is difficult to understand the meaning of this channel index  $v$  whose size and properties changes with depth. It mixes multiple unknown signal attributes with an arbitrary ordering. Multiscale Hierarchical CNNs will address this issue by imposing a high-dimensional hierarchical structure on  $v$ , with an ordering specified by the translation group.

In standard CNN, each  $x_j = \Phi_j x$  is computed with a cascade of convolutions and non-linearities

$$\Phi_j = \rho W_j \dots \rho W_1,$$

whose supports along  $u$  increase with the depth  $j$ . These multiscale operators replace  $x$  by the variables  $x_j$  to estimate the class  $y = f(x)$ . To avoid errors, this change of variable must be discriminative, despite the dimensionality reduction, in the sense that

$$\forall (x, x') \in \mathbb{R}^{2N} \quad \Phi_j(x) = \Phi_j(x') \Rightarrow f(x) = f(x'). \quad (1)$$

This is necessary and sufficient to guarantee that there exists a classification function  $f_j$  such that  $f = f_j \Phi_j$  and hence

$$\forall x \in \mathbb{R}^N, \quad f_j(x_j) = f(x).$$

The function  $f(x)$  can be characterized by its groups of symmetries. A group of symmetries of  $f$  is a group of operators  $g$  which transforms any  $x$  into  $x' = g.x$  which belong to the same class:  $f(x) = f(g.x)$ . The discriminative property (1) implies that if  $\Phi_j(x) = \Phi_j(g.x)$  then  $f(x) = f(g.x)$ . The discrimination property (1) is thus equivalent to impose that groups of symmetries of  $\Phi_j$  are groups of symmetries of  $f$ . Learning appropriate change of variables can thus be interpreted as learning progressively more symmetries of  $f$  (Mallat, 2016). The network must be sufficiently flexible to compute change of variables  $\Phi_j$  whose symmetries approximate the symmetries of  $f$ .

Deep convolutional networks are cascading convolutions along the spatial variable  $u$  so that  $\Phi_j$  is covariant to spatial translations. If  $x$  is translated along  $u$  then  $x_j = \Phi_j x$  is also translated along  $u$ . This covariance implies that for

<sup>1</sup><https://github.com/jhjacobson/HierarchicalCNN>




Figure 1. An illustration of the difference between the input to a vanilla CNN (left) and an HCNN layer (right). A CNN computes convolution along  $u$  and mixes feature maps via a linear combination along the channel index  $v$ . An HCNN performs joint convolution along space and channels  $(u, v_j)$ . This gives rise to an ordering of the feature maps along the attribute index  $v_j$ , as opposed to arbitrary ordering in a vanilla CNN layer. Note, that we only illustrate the 3-dimensional case here whereas HCNN layers are up to 5-dimensional.

all  $v$ ,  $\sum_u x_j(u, v)$  is invariant to translations of  $x$ . Next section explains how to extend this property to higher dimensional attributes with multidimensional convolutions.

### 3. Multiscale Hierarchical CNNs

Multiscale Hierarchical networks are highly structured convolutional networks defined in (Mallat, 2016). The one-dimensional index  $v$  is replaced by a multidimensional vector of attributes  $v = (v_1, \dots, v_j)$  and all linear operators  $W_j$  are convolutions over  $(u, v)$ . We explain their construction and a specific architecture adapted to an efficient learning procedure.

Each layer  $x_j(u, v)$  is indexed by a vector of multidimensional parameters  $v = (v_1, \dots, v_j)$  of dimension  $j$ . Each  $v_k$  is an ‘‘attribute’’ of  $x$  which is learned to discriminate classes  $y = f(x)$ . The operators  $W_j$  are defined as convolutions along a group which is a parallel transport in the index space  $(u, v)$ . With no loss of generality, in this implementation, the transport is a multidimensional translation along  $(u, v)$ . The operators  $W_j$  are therefore multidimensional convolutions, which are covariant to translations along  $(u, v)$ . As previously explained, this covariance to translations implies that the sum  $\sum_{v_k} x_j(u, v_0, \dots, v_j)$  is invariant to translations of previous layers along  $v_k$ . A convolution of  $z(u, v)$  by a filter  $w(u, v)$  of support  $S$  is written

$$z \star w(u, v) = \sum_{(u', v') \in S} z(u - u', v - v') w(u', v'). \quad (2)$$

Since  $z(u, v)$  is defined in a finite domain of  $(u, v)$ , boundary issues can be solved by extending  $z$  with zeros or as a periodic signal. We use zero-padding extensions for the next sections, except for the last section, where we use periodic convolutions. Both cases give similar accuracy.

The network takes as input a color image  $x(u, v_0)$ , or any type of multichannel signal indexed by  $v_0$ . The first layer computes a sum of convolutions of  $x(u, v_0)$  along  $u$ , with

filters  $w_{1, v_0, v_1}(u)$

$$x_1(u, v_1) = \rho \left( \sum_{v_0} x(\cdot, v_0) \star w_{1, v_0, v_1}(u) \right). \quad (3)$$

For any  $j \geq 2$ ,  $W_j$  computes convolutions of  $x_{j-1}(u, v)$  for  $v = (v_1, \dots, v_{j-1})$  with a family of filters  $\{w_{v_j}\}_{v_j}$  indexed by the new attribute  $v_j$ :

$$x_j(u, v, v_j) = \rho \left( x_{j-1} \star w_{v_j}(u, v) \right). \quad (4)$$

As explained in (Mallat, 2016),  $W_j$  has two roles. First, these convolutions indexed by  $v_j$  prepares the discriminability (1) of the next layer  $x_{j+1}$ , despite local or global summations along  $(u, v_1, \dots, v_{j-1})$  implemented at this next layer. It thus propagates discriminative variations of  $x_{j-1}$  from  $(u, v_1, \dots, v_{j-1})$  into  $v_j$ . Second, each convolution with  $w_{v_j}$  computes local or global invariants by summations along  $(u, v_1, \dots, v_{j-2})$ , in order to reduce dimensionality. This dimensionality reduction is implemented by a subsampling of  $(u, v)$  at the output (4), which we omitted here for simplicity. Since  $v_k$  is the index of multidimensional filters, a translation along  $v_k$  is a shift along an ordered set of multidimensional filters. For any  $k < j - 1$ ,  $\sum_{v_k} x_{j-1}(u, v_1, \dots, v_k, \dots, v_{j-1})$  is invariant to any such shift.

The final operator  $W_J$  computes invariants over  $u$  and all attributes  $v_k$  but the last one:

$$x_J(v_{J-1}) = \sum_{u, v_1, \dots, v_{J-1}} x_{J-1}(u, v_1, \dots, v_{J-1}). \quad (5)$$

The last attribute  $v_{J-1}$  corresponds to the class index, and its size is the number of classes. The class  $y = f(x)$  is estimated by applying a soft-max operator on  $x_J(v_{J-1})$ .

**Proposition 3.1** *The last layer  $x_J$  is invariant to translations of  $x_j(u, v_1, \dots, v_j)$  along  $(u, v_1, \dots, v_j)$ , for any  $j < J - 1$ .*




Figure 2. Implementation of a Multiscale Hierarchical CNN as a cascade of 5D convolutions  $W_j$ . The figure gives the size of the intermediate layers stored in 5D arrays. Dash dots lines indicate the parametrization of a layer  $x_j$  and its dimension. We only represent dimensions when the output has a different size from the input.

*Proof:* Observe that  $x_J = W_J \rho W_{J-1} \dots \rho W_j x_j$ . Each  $W_k$  for  $j < k < J$  is a convolution along  $(u, v_0, \dots, v_j, \dots, v_k)$  and hence covariant to translations along  $(u, v_0, \dots, v_j)$ . Since  $\rho$  is a pointwise operator, it is also covariant to translations. Translating  $x_j$  along  $(u, v_1, \dots, v_j)$  thus translates  $x_{J-1}$ . Since (7) computes a sum over these indices, it is invariant to these translations.  $\square$

This proposition proves that the soft-max of  $x_J$  approximates the classification function  $f_j(x_j) = f(x)$  by an operator which is invariant to translations along the high-dimensional index  $(u, v) = (u, v_1, \dots, v_j)$ . The change of variable  $x_j$  thus aims at mapping the symmetry group of  $f$  into a high-dimensional translation group, which is a flat symmetry group with no curvature. It means that classes of  $x_j$  where  $f_j(x_j)$  is constant define surfaces which are progressively more flat as  $j$  increases. However, this requires an important word of caution. A translation of  $x_j(u, v_1, \dots, v_j)$  along  $u$  corresponds to a translation of  $x(u, v_0)$  along  $u$ . On the contrary, a translation along the attributes  $(v_1, \dots, v_j)$  usually does not correspond to transformations on  $x$ . Translations of  $x_j$  along  $(v_1, \dots, v_j)$  is a group of symmetries of  $f_j$  but do not define transformations of  $x$  and hence do not correspond to a symmetry group of  $f$ . Next sections analyze the properties of translations along attributes computed numerically.

Let us give examples over images or audio signals  $x(u)$  having a single channel. The first layer (3) computes convolutions along  $u$ :  $x_1(u, v_1) = \rho(x \star w_{v_1}(u))$ . For audio signals,  $u$  is time. This first layer usually computes a wavelet spectrogram, with wavelet filters  $w_{v_1}$  indexed by a log-frequency index  $v_1$ . A frequency transposition corresponds to a log-frequency translation  $x_1(u, v_1 - \tau)$  along  $v_1$ . If  $x$  is a sinusoidal wave then this translation corresponds to a shift of its frequency and hence to a transformation of  $x$ . However, for more general signals  $x$ , there exists no  $x'$  such that  $\rho(x' \star w_{v_1}(u)) = x_1(u, v_1 - \tau)$ . It is indeed well known that a frequency transposition does not define an exact signal transformation. Other audio attributes such as timber are not either well defined transformations on  $x$  although important attributes to classify sounds.

For images,  $u = (u_1, u_2)$  is a spatial index. If  $w_{v_1} =$

$w(r_{v_1}^{-1}u)$  is a rotation of a filter  $w(u)$  by an angle  $v_1$  then

$$x_1(u, v_1 - \tau) = \rho(x_\tau \star w_{v_1}(r_\tau u)) \text{ with } x_\tau(u) = x(r_\tau^{-1}u).$$

However, there exists no  $x'$  such that  $\rho(x' \star w_{v_1}(u)) = x_1(u, v_1 - \tau)$  because of the missing spatial rotation  $r_\tau u$ . These examples show that translation  $x_j(u, v_1, \dots, v_j)$  along the attributes  $(v_1, \dots, v_j)$  usually do not correspond to a transformation of  $x$ .

## 4. Hierarchical Attribute CNNs

### 4.1. Progressive Dimensionality Reduction

Multiscale hierarchical network layers are indexed by two-dimensional spatial indices  $u = (u_1, u_2)$  and progressively higher dimensional attributes  $v = (v_1, \dots, v_j)$ . To avoid computing high-dimensional vectors and convolutions, we introduce an image classification architecture which eliminates the dependency relatively to all attributes but the last three  $(v_{j-2}, v_{j-1}, v_j)$ , for  $j > 2$ . Since  $u = (u_1, u_2)$ , all layers are stored in five dimensional arrays.

The network takes as an input a color image  $x(u, v_0)$ , with three color channels  $1 \leq v_0 \leq 3$  and  $u = (u_1, u_2)$ . Applying (3) and (4) up to  $j = 3$  computes a five-dimensional layer  $x_3(u, v_1, v_2, v_3)$ . For  $j > 3$ ,  $x_j$  is computed as a linear combination of marginal sums of  $x_{j-1}$  along  $v_{j-3}$ . Thus, it does not depend anymore on  $v_{j-3}$  and can be stored in a five-dimensional array indexed by  $(u, v_{j-2}, v_{j-1}, v_j)$ . This is done by convolving  $x_{j-1}$  with a filter  $w_{v_j}$  which does not depend upon  $v_{j-3}$ :

$$w_{v_j}(u, v_{j-3}, v_{j-2}, v_{j-1}) = w_{v_j}(u, v_{j-2}, v_{j-1}). \quad (6)$$

We indeed verify that this convolution is a linear combination of sums over  $v_{j-3}$ , so  $x_j$  depends only upon  $(u, v_{j-2}, v_{j-1}, v_j)$ . The convolution is subsampled by  $2^{s_j}$  with  $s_j \in \{0, 1\}$  along  $u$ , and a factor 2 along  $v_{j-1}$  and  $v_j$

$$x_j(u, v_{j-2}, v_{j-1}, v_j) = x_{j-1} \star w_{v_j}(2^{s_j}u, 2v_{j-2}, 2v_{j-1}),$$

At depth  $j$ , the array of attributes  $v = (v_{j-2}, v_{j-1}, v_j)$  is of size  $K/4 \times K/2 \times K$ . The parameters  $K$  and spatial sub-sampling factors  $s_j$  are adjusted with a trade-off between

computations, memory and classification accuracy. The final layer is computed with a sum (7) over all parameters but the last one, which corresponds to the class index:

$$x_J(v_{J-1}) = \sum_{u, v_{J-3}, v_{J-2}} x_{J-1}(u, v_{J-3}, v_{J-2}, v_{J-1}). \quad (7)$$

This architecture is illustrated in Figure 2.

#### 4.2. Filter Factorization for Training

Our newly introduced Hierarchical Attribute CNNs (HCNN) have been tested on CIFAR10 and CIFAR100 image databases. CIFAR10 has 10 classes, while CIFAR100 has 100 classes, which makes it more challenging. The train and test sets have 50k and 10k colored images of  $32 \times 32$  pixels. Images are preprocessed via a standardization along the RGB channels. No whitening is applied as we did not observe any improvement.

Our HCNN is trained in the same way as a classical CNN. We train it by minimizing a neg-log entropy loss, via SGD with momentum 0.9 for 240 epochs. An initial learning rate of 0.25 is chosen while being reduced by a factor 10 every 40 epochs. Each minibatch is of size 50. The learning is regularized by a weight decay of  $2 \cdot 10^{-4}$  (Krizhevsky et al., 2012). We incorporate a data augmentation with random translations of 6 pixels and flips (Krizhevsky & Hinton, 2010).

Just as in any other CNNs, the gradient descent is badly conditioned because of the large number of parameters (Goodfellow et al., 2014). We precondition and regularize the 4 dimensional filters  $w_{v_j}$ , by normalizing a factorization of these filters. We factorize  $w_{v_j}(u, v_{j-3}, v_{j-2}, v_{j-1})$  into a sum of  $Q$  separable filters:

$$w_{v_j}(u, v_{j-3}, v_{j-2}, v_{j-1}) = \sum_{q=1}^Q h_{j,q}(u) g_{v_j,q}(v_{j-2}, v_{j-1}), \quad (8)$$

and introduce an intermediate normalization before the sum. Let us write  $h_{j,q}(u, v) = \delta(u) h_{j,q}(u)$  and  $g_{v_j,q}(u, v) = \delta(u) g_{v_j,q}(v)$ . The batch normalization is applied to  $x_{j-1} \star h_{j,q}$  and subtracts a mean array  $m_{j,q}$  while normalizing the standard deviations of all coefficients  $\sigma_{j,q}$ :

$$\tilde{x}_{j,q}(u, v) = \frac{x_{j-1} \star h_{j,q} - m_{j,q}}{\sigma_{j,q}}.$$

This normalized output is retransformed according to (8) by a sum over  $q$  and a subsampling:

$$x_j(u, v) = \rho \left( \sum_{q=1}^Q \tilde{x}_{j,q} \star g_{v_j,q}(2^{s_j} u, 2v) \right).$$

The convolution operator  $W_j$  is thus subdivided into a first operator  $W_j^h$  which computes standardized convolutions

along  $u$  cascaded with  $W_j^g$  which sums  $Q$  convolutions along  $v$ . Since the tensor rank of  $W_j$  cannot be larger than 9, using  $Q \geq 9$  does not restrict the rank of the operators  $W_j$ . However, as reported in (Jacobsen et al., 2016), increasing the value of  $Q$  introduces an overparametrization which regularizes the optimization. Increasing  $Q$  from 9 to 16 and then from 16 to 32 brings a relative increase of the classification accuracy of 4.2% and then of 1.1%.

We also report a modification of our network (denoted by (+)) which incorporates an intermediate non-linearity:

$$x_j(u, v) = \rho(W_j^g \rho(W_j^h x_{j-1})).$$

Observe that in this case,  $x_j$  is still covariant with the actions of the translations along  $(u, v)$ , yet the factorization of  $w_{v_j}$  into  $(h_{j,q}, g_{v_j,q})$  does not hold anymore.

For classification of CIFAR images, the total depth is  $J = 12$  and a downsampling by 2 along  $u$  is applied at depth  $j = 5, 9$ . Figure 2 describes our model architecture as a cascade of  $W_j$  and  $\rho$ , and gives the size of each layer. Each attribute can take at most  $K = 16$  values.

The number of free parameters of the original architecture is the number of parameters of the convolution kernels  $w_{v_j}$  for  $1 \leq v_j \leq K$  and  $2 < j < J$ , although they are factorized into separable filters  $h_{j,q}(u)$  and  $g_{v_j,q}(v_{j-2}, v_{j-1})$  which involve more parameters. The filters  $w_{v_j}$  have less parameters for  $j = 2, 3$  because they are lower-dimensional convolution kernels. In CIFAR-10, for  $3 < j < J$ , each  $w_{v_j}$  has a spatial support of size  $3^2$  and a support of  $7 \times 11$  along  $(v_{j-2}, v_{j-1})$ . If we add the 10 filters which output the last layer, the resulting total number of network parameters is approximately 0.098M. In CIFAR-100, the filters rather have a support of  $11 \times 11$  along  $(v_{j-2}, v_{j-1})$  but the last layer has a size 100 which considerably increases the number of parameters which is approximately 0.25M.

The second implementation (+) introduces a non-linearity  $\rho$  between each separable filter, so the overall computations can not be reduced to equivalent filters  $w_{v_j}$ . There are  $Q = 32$  spatial filters  $h_{j,q}(u)$  of support  $3 \times 3$  and  $QK$  filters  $g_{v_j,q}(v_{j-2}, v_{j-1})$  of support  $7 \times 11$ . The total number of coefficients required to parametrize  $h_{j,q}, g_{v_j,q}$  is approximately 0.34M. In CIFAR-100, the number of parameters becomes 0.89M. The total number of parameters of the implementation (+) is thus much bigger than the original implementation which does not add intermediate non-linearities. Next section compares these number of parameters with architectures that have similar numerical performances.

## 5. An explicit structuration

This section shows that Hierarchical Attribute CNNs have comparable classification accuracies on the CIFAR image

dataset than state-of-the-art architectures, with much fewer parameters. We also investigate the properties of translations along the attributes  $v_j$  learned on CIFAR.

### 5.1. Classification Performance

We evaluate our Hierarchical CNN on CIFAR-10 (table 1) and CIFAR-100 (table 2) in the setting explained above. Our network achieves an error of 8.6% on CIFAR-10, which is comparable to recent state-of-the-art architectures. On CIFAR-100 we achieve an error rate of 38%, which is about 4% worse than the closely related all-convolutional network baseline, but our architecture has an order of magnitude fewer parameters.

Classification algorithms using a priori defined representations or representations computed with unsupervised algorithms have an accuracy which barely goes above 80% on CIFAR-10 (Oyallon & Mallat, 2015). On the contrary, supervised CNN have an accuracy above 90% as shown by Table 1. This is also the case for our structured hierarchical network which has an accuracy above 91%. Improving these results may be done with larger  $K$  and  $Q$  which could be done with faster GPU implementation of multi-dimensional convolutions, although it is a technical challenge (Budden et al., 2016). Our proposed architecture is based on “plain vanilla” CNN architectures to which we compare our results in Table 1. Applying residual connections (He et al., 2016), densely connected layers (Huang et al., 2016), or similar improvements, might overcome the 4% accuracy gap with the best existing architectures. In the following, we study the properties resulting from the hierarchical structuration of our network, compared with classical CNN.

### 5.2. Reducing the number of parameters

The structuration of a Deep neural network aims at reducing the number of parameters and making them easier to interpret in relation to signal models. Reducing the number of parameters means characterizing better the structures which govern the classification.

This section compares Hierarchical Attribute CNNs to other structured architectures and algorithms which reduce the number of parameters of a CNN during, and after training. We show that Hierarchical Attribute CNNs involve less parameters during and after training than other architectures in the literature.

We review various strategies to reduce the number of parameters of a CNN and compare them with our Hierarchical Attribute CNN. Several studies show that one can factorize CNN filters (Denton et al., 2014; Jaderberg et al., 2014) *a posteriori*. A reduction of parameters is obtained by computing low-rank factorized approximations of the

Table 1. Classification accuracy on CIFAR10 dataset.

MODEL	# PARAMETERS	% ACCURACY
HIEARCHICAL CNN	0.098M	91.23
HIEARCHICAL CNN (+)	0.34M	92.30
ALL-CNN	1.3M	92.75
RESNET 20	0.27M	91.25
NETWORK IN NETWORK	0.98M	91.20
WRN-STUDENT	0.17M	91.23
FITNET	2.5M	91.61

Table 2. Classification accuracy on CIFAR100 dataset.

MODEL	# PARAMETERS	% ACCURACY
HIEARCHICAL CNN	0.25M	62.01
HIEARCHICAL CNN (+)	0.89M	63.19
ALL-CNN	1.3M	66.29
NETWORK IN NETWORK	0.98M	64.32
FITNET	2.5M	64.96

filters calculated by a trained CNN. It leads to more efficient computations with operators defined by fewer parameters. Another strategy to reduce the number of network weights is to use teacher and student networks (Zagoruyko & Komodakis, 2016; Romero et al., 2014), which optimize a CNN defined by fewer parameters. The student network adapts a reduced number of parameters for data classification via the teacher network.

A parameter redundancy has also been observed in the final fully connected layers used by number of neural network architectures, which contain most of the CNN parameters (Cheng et al., 2015; Lu et al., 2016). This last layer is replaced by a circulant matrix during the CNN training, with no loss in accuracy, which indicates that last layer can indeed be structured. Other approaches (Jacobsen et al., 2016) represent the filters with few parameters in different bases, instead of imposing that they have a small spatial support. These filters are represented as linear combinations of a given family of filters, for example, computed with derivatives Gaussians. This approach is structuring jointly the channel and spatial dimensions. Finally, HyperNetworks (Ha et al., 2016) permits to drastically reducing the number of parameters used during the training step, to 0.097M and obtaining 91.98% accuracy. However, we do not report them as 0.97M corresponds to a non-linear number of parameters for the network.

Table 1 and 2 give the performance of different CNN architectures with their number of parameters, for the CIFAR10

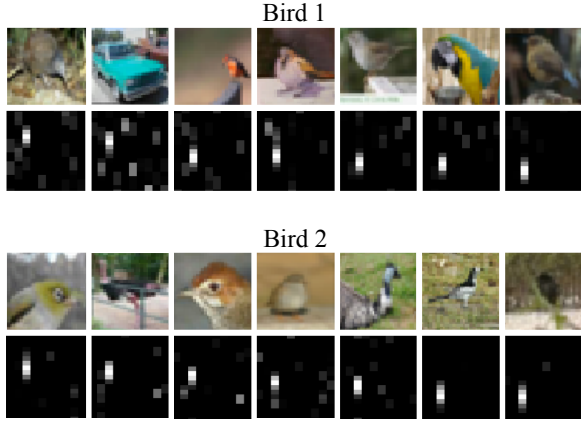


Figure 3. The first images of the first and third rows are the two input image  $x$ . Their invariant attribute array  $\bar{x}_j(v_{j-1}, v_j)$  is shown below for  $j = J - 1$ , with high amplitude coefficients appearing as white points. Vertical and horizontal axes correspond respectively to  $v_{j-1}$  and  $v_j$ , so translations of  $v_{j-1}$  by  $\tau$  are vertical translations. An image  $x^\tau$  in a column  $\tau + 1$  has an invariant attribute  $\bar{x}_j^\tau$  which is shown below. It is the closest to  $\bar{x}_j(v_{j-1} - \tau, v_j)$  in the databasis.

and CIFAR100 datasets. For Hierarchical Attribute CNNs, the convolution filters are invariant to translations along  $u$  and  $v$  which reduces the number of parameters by an important factor compared to other architectures. All-CNN (Springenberg et al., 2014) is an architecture based only on sums of spatial convolutions and ReLU non-linearities, which has a total of  $1.3M$  parameters, and a similar accuracy to ours. Its architecture is similar to our hierarchical architecture, but it has much more parameters because filters are not translation invariant along  $v$ . Interestingly, a ResNet (He et al., 2016) has more parameters and performs similarly whereas it is a more complex architecture, due to the shortcut connexions. WRN-student is a student resnet (Zagoruyko & Komodakis, 2016) with  $0.2M$  parameters trained via a teacher using  $0.6M$  parameters and which gets an accuracy of  $93.42\%$  on CIFAR10. FitNet networks (Romero et al., 2014) also use compression methods but need at least  $2.5M$  parameters, which is much larger than our network. Our architecture brings an important parameter reduction on CIFAR10 for accuracies around  $90\%$ . There is also a drastic reduction of parameters on CIFAR100.

### 5.3. Interpreting the translation

The structure of Hierarchical Attribute CNNs opens up the possibility of interpreting inner network coefficients, which is usually not possible for CNNs. A major mathematical challenge is to understand the type of invariants computed by deeper layers of a CNN. Hierarchical networks compute invariants to translations relatively to learned attributes  $v_j$ ,

which are indices of the filters  $w_{v_j}$ . One can try to relate these attributes translations to modifications of image properties. As explained in Section 3, a translation of  $x_j$  along  $v_j$  usually does not correspond to a well-defined transformation of the input signal  $x$  but it produces a translation of the next layers. Translating  $x_j$  along  $v_j$  by  $\tau$  translates  $x_{j+1}(u, v_{j-1}, v_j, v_{j+1})$  along  $v_j$  by  $\tau$ .

To analyze the effect of this translation, we eliminate variability along  $v_{j-2}$  and define an invariant attribute array by choosing the central spatial position  $u_0$ :

$$\bar{x}_j(v_{j-1}, v_j) = \sum_{v_{j-2}} x_j(u_0, v_{j-2}, v_{j-1}, v_j). \quad (9)$$

We relate this translation to an image in the training dataset by finding the image  $x^\tau$  in the dataset which minimizes  $\|\bar{x}_j(v_{j-1} - \tau, v_j) - \bar{x}_j^\tau(v_{j-1}, v_j)\|_2$ , if this minimum Euclidean distance is sufficiently small. To compute accurately a translation by  $\tau$  we eliminate the high frequency variations of  $x_j$  and  $x_j^\tau$  along  $v_{j-1}$  with a filter which averages consecutive samples, before computing their translation. The network used in this experiment is implemented with circular convolutions to avoid border effects, which have nearly the same classification performance.

Figure 3 shows the sequence of  $x^\tau$  obtained with a translation by  $\tau$  of  $\bar{x}_j$  at depth  $j = J - 1$ , for two images  $x$  in the “bird” class. Since we are close to the output, we expect that translated images belong to the same class. This is not the case for the second image of the first “Bird 1”. It is a “car” instead of a “bird”. This corresponds to a classification error but observe that  $\bar{x}_{j-1}^\tau$  is quite different from  $\bar{x}_{j-1}$  translated. We otherwise observe that in these final layers, translations of  $\bar{x}_{j-1}$  defines images in the same class.

Figure 4 gives sequences of translated attribute images  $x^\tau$ , computed by translating  $\bar{x}_j$  by  $\tau$  at different depth  $j$  and for different input  $x$ . As expected, at small depth  $j$ , translating an attribute  $v_{j-1}$  does not define images in the same class. These attribute rather correspond to low-level image properties which depend upon fine scale image properties. However, these low-level properties can not be identified just by looking at these images. Indeed, the closer images  $x^\tau$  identified in the databasis are obtained with a distance over coefficients which are invariant relatively to all other attributes. These images are thus very different and involve variabilities relatively to all other attributes. To identify the nature of an attribute  $v_j$ , a possible approach is to correlate the images  $x^\tau$  over a large set of images, while modifying known properties of  $x$ .

At deep layers  $j$ , translations of  $\bar{x}_j$  define images  $x^\tau$  which have a progressively higher probability to belong to the same class as  $x$ . These attribute transformations correspond to large scale image pattern related to modifications



Figure 4. The first columns give the input image  $x$ , from which we compute the invariant array  $\bar{x}_j$  at a depth  $3 \leq j \leq 11$  which increases with the row. The next images in the same row are the images  $x^\tau$  whose invariant arrays  $\bar{x}_j^\tau$  are the closest to  $\bar{x}_j$  translated by  $1 \leq \tau \leq 7$ , among all other images in the databasis. The value of  $\tau$  is the column index minus 1.

of the filters  $w_{v_{j-1}}$ . In this case, the attribute indices could be interpreted as addresses in organized arrays. The translation group would then correspond to translations of addresses. Understanding better the properties of attributes at different depth is an issue that will be explored in the future.

## 6. Conclusion

Hierarchical Attribute CNNs give a mathematical and experimental framework to study invariants computed by deep neural networks. Layers are parameterized in progressively higher dimensional spaces of hierarchical attributes, which are learned from training data. All network operators are multidimensional convolutions along attribute indices, so that invariants can be computed by summations along these attributes.

This paper introduces a computationally tractable implementation of such networks, progressively reducing dimensionality by calculating invariants in intermediate layers. This permits to reduce dimensionality and parameters of intermediate layers, making it possible to train so far theoretical Multiscale Hierarchical CNNs. We achieve good clas-

sification accuracies, while obtaining a CNN with generically organized channel dimensions and reduced number of parameters compared to common SOTA models.

Translations along attributes at shallow depth correspond to low-level image properties at fine scales like color, whereas attributes at deep layers correspond to modifications of large scale pattern structures like faces and eventually to classes in the last layer. Understanding better the multi-scale properties of these attributes and their relations to the symmetry group of  $f$  is an important issue, which can lead to a better mathematical understanding of CNN learning algorithms.

## Acknowledgements

This work is funded by STW project ImaGene, ERC grant InvariantClass 320959 and via a grant for PhD Students of the Conseil régional d'Île-de-France (RDMIdF).



## References

- Bruna, Joan and Mallat, Stéphane. Invariant scattering convolution networks. *IEEE transactions on pattern analysis and machine intelligence*, 35(8):1872–1886, 2013.
- Bruna, Joan, Szlam, Arthur, and LeCun, Yann. Learning stable group invariant representations with convolutional networks. *arXiv preprint arXiv:1301.3537*, 2013.
- Budden, David, Matveev, Alexander, Santurkar, Shibani, Chaudhuri, Shraman Ray, and Shavit, Nir. Deep tensor convolution on multicores. *arXiv preprint arXiv:1611.06565*, 2016.
- Cheng, Yu, Yu, Felix X, Feris, Rogerio S, Kumar, Sanjiv, Choudhary, Alok, and Chang, Shi-Fu. An exploration of parameter redundancy in deep networks with circulant projections. In *Proceedings of the IEEE International Conference on Computer Vision*, pp. 2857–2865, 2015.
- Clevert, Djork-Arné, Unterthiner, Thomas, and Hochreiter, Sepp. Fast and accurate deep network learning by exponential linear units (elus). *arXiv preprint arXiv:1511.07289*, 2015.
- Cohen, Taco S and Welling, Max. Steerable cnns. *arXiv preprint arXiv:1612.08498*, 2016.
- Dai, Jifeng, Qi, Haozhi, Xiong, Yuwen, Li, Yi, Zhang, Guodong, Hu, Han, and Wei, Yichen. Deformable convolutional networks. *arXiv preprint arXiv:1703.06211*, 2017.
- Denton, Emily L, Zaremba, Wojciech, Bruna, Joan, LeCun, Yann, and Fergus, Rob. Exploiting linear structure within convolutional networks for efficient evaluation. In *Advances in Neural Information Processing Systems*, pp. 1269–1277, 2014.
- Gens, Robert and Domingos, Pedro M. Deep symmetry networks. In *Advances in neural information processing systems*, pp. 2537–2545, 2014.
- Goodfellow, Ian J, Vinyals, Oriol, and Saxe, Andrew M. Qualitatively characterizing neural network optimization problems. *arXiv preprint arXiv:1412.6544*, 2014.
- Ha, David, Dai, Andrew, and Le, Quoc V. Hypernetworks. *arXiv preprint arXiv:1609.09106*, 2016.
- He, Kaiming, Zhang, Xiangyu, Ren, Shaoqing, and Sun, Jian. Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 770–778, 2016.
- Huang, Gao, Liu, Zhuang, Weinberger, Kilian Q, and van der Maaten, Laurens. Densely connected convolutional networks. *arXiv preprint arXiv:1608.06993*, 2016.
- Ioffe, Sergey and Szegedy, Christian. Batch normalization: Accelerating deep network training by reducing internal covariate shift. *arXiv preprint arXiv:1502.03167*, 2015.
- Jacobsen, Jorn-Henrik, van Gemert, Jan, Lou, Zhongyu, and Smeulders, Arnold WM. Structured receptive fields in cnns. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 2610–2619, 2016.
- Jacobsen, Jörn-Henrik, de Brabandere, Bert, and Smeulders, Arnold WM. Dynamic steerable blocks in deep residual networks. *arXiv preprint arXiv:1706.00598*, 2017.
- Jaderberg, Max, Vedaldi, Andrea, and Zisserman, Andrew. Speeding up convolutional neural networks with low rank expansions. *arXiv preprint arXiv:1405.3866*, 2014.
- Jaderberg, Max, Simonyan, Karen, Zisserman, Andrew, et al. Spatial transformer networks. In *Advances in Neural Information Processing Systems*, pp. 2017–2025, 2015.
- Krizhevsky, Alex and Hinton, G. Convolutional deep belief networks on cifar-10. *Unpublished manuscript*, 40, 2010.
- Krizhevsky, Alex, Sutskever, Ilya, and Hinton, Geoffrey E. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pp. 1097–1105, 2012.
- LeCun, Yann, Boser, Bernhard, Denker, John S, Henderson, Donnie, Howard, Richard E, Hubbard, Wayne, and Jackel, Lawrence D. Backpropagation applied to handwritten zip code recognition. *Neural computation*, 1(4): 541–551, 1989.
- LeCun, Yann, Kavukcuoglu, Koray, and Farabet, Clément. Convolutional networks and applications in vision. In *Circuits and Systems (ISCAS), Proceedings of 2010 IEEE International Symposium on*, pp. 253–256. IEEE, 2010.
- LeCun, Yann, Bengio, Yoshua, and Hinton, Geoffrey. Deep learning. *Nature*, 521(7553):436–444, 2015.
- Lu, Zhiyun, Sindhvani, Vikas, and Sainath, Tara N. Learning compact recurrent neural networks. In *Acoustics, Speech and Signal Processing (ICASSP), 2016 IEEE International Conference on*, pp. 5960–5964. IEEE, 2016.
- Mallat, Stéphane. Understanding deep convolutional networks. *Phil. Trans. R. Soc. A*, 374(2065):20150203, 2016.

- Oyallon, Edouard and Mallat, Stéphane. Deep roto-translation scattering for object classification. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 2865–2873, 2015.
- Romero, Adriana, Ballas, Nicolas, Kahou, Samira Ebrahimi, Chassang, Antoine, Gatta, Carlo, and Bengio, Yoshua. Fitnets: Hints for thin deep nets. *arXiv preprint arXiv:1412.6550*, 2014.
- Springenberg, Jost Tobias, Dosovitskiy, Alexey, Brox, Thomas, and Riedmiller, Martin. Striving for simplicity: The all convolutional net. *arXiv preprint arXiv:1412.6806*, 2014.
- Worrall, Daniel E, Garbin, Stephan J, Turmukhambetov, Daniyar, and Brostow, Gabriel J. Harmonic networks: Deep translation and rotation equivariance. *arXiv preprint arXiv:1612.04642*, 2016.
- Zagoruyko, Sergey and Komodakis, Nikos. Paying more attention to attention: Improving the performance of convolutional neural networks via attention transfer. *arXiv preprint arXiv:1612.03928*, 2016.
- Zeiler, Matthew D and Fergus, Rob. Visualizing and understanding convolutional networks. In *European conference on computer vision*, pp. 818–833. Springer, 2014.