

Supplementary Material

Harmonic Networks: Deep Translation and Rotation Equivariance

Daniel E. Worrall, Stephan J. Garbin, Daniyar Turmukhambetov and Gabriel J. Brostow
 {d.worrall, s.garbin, d.turmukhambetov, g.brostow}@cs.ucl.ac.uk
 Computer Science Department, University College London.

Abstract

We include some proofs and derivations of the rotational equivariance properties of the circular harmonics, along with a demonstration of how we calculate the number of parameters for various network architectures.

1. Equivariance properties

In Section 3.2 we mentioned that cross-correlation with the circular harmonics is a 360° -rotation equivariant feature transform. Here we provide the proof, and some of the properties mentioned in **Arithmetic and Equivariance Condition**.

1.1. Equivariance of the Circular Harmonics

Consider an image $\mathbf{F}(r, \phi)$, where r and ϕ are polar coordinates. We are interested in proving that there exists a filter \mathbf{W}_m , such that cross-correlation of \mathbf{F} with \mathbf{W}_m yields a rotationally equivariant feature map. We show that \mathbf{W}_m takes the form of a circular harmonic, $\mathbf{W}_m = R(r)e^{i(m\phi+\beta)}$, for some $m \in \mathbb{Z}$. We begin from the fact that a rotation by θ about the origin leads to a new image $\mathbf{F}(r, \phi - \theta)$. In the notation of the paper, we also write this as $\mathbf{F}(r, \phi^\theta[\phi]) := \mathbf{F}(r, \phi - \theta)$. This means that the cross-correlation is

$$[\mathbf{W} \star \mathbf{F}(r, \pi^\theta[\phi])] = \int \mathbf{W}(r, \phi) \mathbf{F}(r, \pi^\theta[\phi]) \, dr d\phi \quad (1)$$

$$= \int \mathbf{W}(r, \phi) \mathbf{F}(r, \phi - \theta) \, dr d\phi \quad (2)$$

$$= \int \mathbf{W}(r, \phi' + \theta) \mathbf{F}(r, \phi') \, dr d\phi', \quad (3)$$

$$(4)$$

where we have used the substitution $\phi' := \phi - \theta$. If we assume that \mathbf{W}_m is of the form $\mathbf{W}_m = R(r)e^{i(m\phi+\beta)}$, then the integral becomes

$$[\mathbf{W} \star \mathbf{F}(r, \pi^\theta[\phi])] = \int R(r)e^{i(m(\phi'+\theta)+\beta)} \mathbf{F}(r, \phi') \, dr d\phi' \quad (5)$$

$$= e^{im\theta} \int R(r)e^{i(m\phi'+\beta)} \mathbf{F}(r, \phi') \, dr d\phi' \quad (6)$$

$$= e^{im\theta} [\mathbf{W} \star \mathbf{F}(r, \pi^0[\phi])] = e^{im\theta} [\mathbf{W} \star \mathbf{F}(r, \phi)] \quad (7)$$

Thus we see that cross-correlation of the rotated signal $\mathbf{F}(r, \pi^\theta[\phi])$ with the circular harmonic filter $\mathbf{W}_m = R(r)e^{i(m\phi+\beta)}$ is equal to the response at zero rotation $[\mathbf{W} \star \mathbf{F}(r, \phi)]$, multiplied by a complex phase shift $e^{im\theta}$. In the notation of the paper, we denote this multiplication by $e^{im\theta}$ as $\psi_m^\theta[\bullet] = e^{im\theta} \cdot \bullet$. Thus cross-correlation with \mathbf{W}_m yields a rotationally equivariant feature mapping.

1.2. Properties

1.2.1 Chained cross-correlation

We claimed in **Arithmetic and Equivariance Condition**, that the rotation order of a feature map resulting from chained cross-correlations is equal to the sum of the the rotation orders of the filters in the chain. We prove this for a chain of two filters, and the rest follows by induction. Consider cross-correlation of a θ -rotated image $\mathbf{F}(r, \phi - \theta)$, with a filter \mathbf{W}_m , followed by cross-correlation with \mathbf{W}_n . We write the response of the first convolution as $e^{im\theta}[\mathbf{W} \star \mathbf{F}]$ when we are considering a single point in the plane. This is a function of θ only. If instead we considered the response at every point in the first layer feature map, then we would write the response as $e^{im\theta}[\mathbf{W} \star \mathbf{F}](\hat{r}, \hat{\phi} - \theta)$, where $\hat{r}, \hat{\phi}$ are the polar coordinates indexing the spatial dimensions of the first layer response. Then we can write the chained convolution as

$$\mathbf{W}_n(\hat{r}, \hat{\phi}) \star [\mathbf{W}_m(r, \phi) \star \mathbf{F}(r, \phi - \theta)] = \mathbf{W}_n(\hat{r}, \hat{\phi}) \star e^{im\theta}[\mathbf{W}_m \star \mathbf{F}](\hat{r}, \hat{\phi} - \theta) \quad (8)$$

$$= e^{im\theta} \left(\mathbf{W}_n(\hat{r}, \hat{\phi}) \star [\mathbf{W}_m \star \mathbf{F}](\hat{r}, \hat{\phi} - \theta) \right) \quad (9)$$

$$= e^{im\theta} e^{in\theta} [\mathbf{W}_n \star [\mathbf{W}_m \star \mathbf{F}]] \quad (10)$$

$$= e^{i(m+n)\theta} [\mathbf{W}_n \star [\mathbf{W}_m \star \mathbf{F}]]. \quad (11)$$

From the first to the second line we have used the property that the cross-correlation is linear and that we may pull the scalar factor $e^{im\theta}$ outside. Thus we see that the chained cross-correlation results in a summation of the rotation orders of the individual filters \mathbf{W}_m and \mathbf{W}_n .

1.2.2 Magnitude nonlinearities

Point-wise nonlinearities acting on the magnitude of a feature map maintain rotational equivariance. Consider that we have a point on a feature map of rotation order m , which we can write as $F e^{im\theta}$, where $F \geq 0$ is the magnitude of the feature map and $e^{im\theta}$ is the phase component. The output of the nonlinearity $g : \mathbb{R}_+ \rightarrow \mathbb{R}$ is

$$g(F e^{im\theta}) = g(F) e^{im\theta}, \quad (12)$$

since g only acts on magnitudes. Since for fixed F the output is a function of m and θ only, the point-wise magnitude-acting nonlinearity preserves rotational equivariance.

1.2.3 Summation of feature maps

The summation of feature maps of the same rotation order is a new feature map of the same rotation order. Consider two feature maps \mathbf{F}_1 and \mathbf{F}_2 of rotation order m . Summation is a pointwise operation, so we only consider two corresponding points in the feature maps, which we denote $F_1 e^{i(m\theta + \beta_1)}$ and $F_2 e^{i(m\theta + \beta_2)}$, where β_1 and β_2 are phase offsets. The sum is

$$F_1 e^{i(m\theta + \beta_1)} + F_2 e^{i(m\theta + \beta_2)} = e^{im\theta} (F_1 e^{i\beta_1} + F_2 e^{i\beta_2}), \quad (13)$$

which for fixed $F_1, F_2, \beta_1, \beta_2$ is a function of m and θ only and so also rotationally equivariant with order m .

2. Number of parameters

Here we list a break down of how we computed the number of parameters for the various network architectures in the experiments section. The networks architectures used are in Figure 1. Red boxes are cross-correlations, blue boxes are pooling (average for H-Nets, max for regular CNNs), green boxes are 1×1 -cross-correlations.




Figure 1. Networks used

2.1. Standard CNN

For a standard CNN layer with i input channels and o output channels, and $k \times k$ sized weights, the number of learnable parameters is $io k^2$. Since there is one bias per output layer, this increases to $io k^2 + o$. If using batch normalization, then there is an extra per-channel scaling factor, which increases the number of learnable parameters to $io k^2 + 2o$. The standard CNN for the rotated MNIST experiments has 6 layers of 3×3 cross-correlations, and 1 layer of 4×4 -cross-correlations, with 20 feature maps per layer and 3 batch normalization layers so the number of learnable parameters is 21570. The calculations are shown in Table 1.

Layer	Weights	Batch Norm/Bias	#Params
1	$3 \cdot 3 \cdot 1 \cdot 20$	20	200
2	$3 \cdot 3 \cdot 20 \cdot 20$	$2 \cdot 20$	3640
3	$3 \cdot 3 \cdot 20 \cdot 20$	20	3620
4	$3 \cdot 3 \cdot 20 \cdot 20$	$2 \cdot 20$	3640
5	$3 \cdot 3 \cdot 20 \cdot 20$	20	3620
6	$3 \cdot 3 \cdot 20 \cdot 20$	$2 \cdot 20$	3640
7	$4 \cdot 4 \cdot 20 \cdot 10$	10	3210
Total			21570

Table 1. Number of parameters for a regular CNN.

2.2. Harmonic networks

The learnable parameters of a Harmonic Network are the radial profile and the the per-filter phase offset. For a $k \times k$ filter, the number of radial profile elements is equal to the number of rings of equal distance from the center of the filter. For example, consider the Figure 2, which is an excerpt from the main paper. This is a 5×5 filter, with 6 rings of equal distance from the center of the filter (the smallest ring is just a single point). So this filter has 6 radial profile terms and 1 phase offset to learn. This contrasts with a regular filter, which would have 25 learnable parameters. Note, that for filters with rotation order $m \neq 0$, the center pixel of the filter is in fact always zero, and so for a 5×5 rotation order $m \neq 0$ filter, the number of radial profile terms is $6 - 1 = 5$. So for the H-Net in the main paper with 5×5 filters and batch normalization in layers 2, 4, & 6, the number of learnable parameters is 33347. The calculations are in Table 2. Note that the final layer contains just one set of biases and no phase offsets. A similar set of calculations can be performed for the deeply supervised networks.




Figure 2. Mapping the radial profile to the elements of a square filter.

Layer	$m = 0$	$m = 1$	Batch Norm/Bias	#Params
1	$6 \cdot 1 \cdot 8 + 1 \cdot 8$	$5 \cdot 1 \cdot 8 + 1 \cdot 8$	$2 \cdot 8$	120
2	$6 \cdot 8 \cdot 8 + 8 \cdot 8$	$5 \cdot 8 \cdot 8 + 8 \cdot 8$	$2 \cdot 16$	864
3	$6 \cdot 8 \cdot 16 + 8 \cdot 16$	$5 \cdot 8 \cdot 16 + 8 \cdot 16$	$2 \cdot 16$	1696
4	$6 \cdot 16 \cdot 16 + 16 \cdot 16$	$5 \cdot 16 \cdot 16 + 16 \cdot 16$	$2 \cdot 32$	3392
5	$6 \cdot 16 \cdot 35 + 16 \cdot 35$	$5 \cdot 16 \cdot 35 + 16 \cdot 35$	$2 \cdot 35$	7350
6	$6 \cdot 35 \cdot 35 + 35 \cdot 35$	$5 \cdot 35 \cdot 35 + 35 \cdot 35$	$2 \cdot 70$	16065
7	$6 \cdot 35 \cdot 10$	$5 \cdot 35 \cdot 10$	10	3860
Total				33347

Table 2. Number of parameters for H-Net.