
Convolutional Networks for Spherical Signals

Taco S. Cohen¹ Mario Geiger¹ Jonas Köhler¹ Max Welling¹

Abstract

The success of convolutional networks in learning problems involving planar signals such as images is due to their ability to exploit the translation symmetry of the data distribution through weight sharing. Many areas of science and engineering deal with signals with other symmetries, such as rotation invariant data on the sphere. Examples include climate and weather science, astrophysics, and chemistry.

In this paper we present spherical convolutional networks. These networks use convolutions on the sphere and rotation group, which results in rotational weight sharing and rotation equivariance. Using a synthetic spherical MNIST dataset, we show that spherical convolutional networks are very effective at dealing with rotationally invariant classification problems.

1. Introduction

Today, neural network architectures are usually found by trial and error. In an ideal world, we would have a few simple design principles that would guide the design of neural networks for various tasks. Although we are currently still far from this goal, over the last few years, *equivariance to symmetry transformations* has emerged as a powerful and widely applicable design principle.

Many learning problems have symmetries. For instance, in image classification, the class label of an image remains the same when the image is shifted, so translations are symmetries of the label function. Convolutional networks exploit this symmetry through weight sharing, thus achieving excellent statistical efficiency. Key to this success is a property of convolutions called equivariance: a shift of the input leads to a shift of the output. Because convolution layers are equivariant, they preserve the symmetry and can thus be stacked to create arbitrarily deep networks.

¹University of Amsterdam. Correspondence to: Taco S. Cohen <taco.cohen@gmail.com>.

This idea, implicit in the classical ConvNet literature, can be generalized to other symmetries that the data may have. A general theory of such group equivariant convolutional networks was presented by (Cohen & Welling, 2016; 2017). By designing the network so that every layer is equivariant to transformations from the group of interest, the symmetry is preserved throughout the network and can be exploited by (generalized) convolutional weight sharing.

In this paper we show how a group equivariant network for the rotation group $SO(3)$ acting on the 2-sphere S^2 can be constructed. This is challenging for two reasons. First, S^2 and $SO(3)$ are continuous manifolds that do not admit symmetric sampling grids, which makes it somewhat tricky to rotate filters numerically (interpolation would be required). Second, the group $SO(3)$ is much bigger than the finite groups considered in earlier work on equivariant networks, making computational and memory efficiency paramount. We address both of these challenges by using generalized Fast Fourier Transform (FFT) algorithms to compute the convolution.

The rest of the paper is structured as follows. We first discuss related work on equivariant deep networks and generalized Fourier analysis. Then, in section 3 and 4 we define spherical and $SO(3)$ convolution and the generalized Fourier transform. Section 5 puts the pieces together and describes the spherical convolutional network. Finally, in section 6 we describe our experiments that verify the mathematical properties expected of the network, and demonstrate the effectiveness of the inductive bias of spherical ConvNets for rotation-invariant classification on a synthetic dataset of spherical MNIST digits.

2. Related Work

A lot of recent work has focussed on exploiting symmetries for data-efficient deep learning. Theoretical investigations of (Anselmi et al., 2013) point out the potential for significant improvements in sample efficiency from deep networks that incorporate geometrical prior knowledge. One of the first works aimed at generalizing convolutional networks to exploit non-translational symmetries was done by Gens & Domingos (2014). Implementations of equivariant neural networks include (Worrall et al., 2017; Jacobsen et al., 2017; Zhou et al., 2017; Ravanbakhsh et al.,

2017; Dieleman et al., 2016; Cohen & Welling, 2016; 2017). Group invariant scattering networks were explored by (Sifre & Mallat, 2013; Oyallon & Mallat, 2015).

3. Spherical and SO(3)-Convolution

By analogy to a planar convolution, which is computed by taking inner products between a signal and a shifted filter, we define the spherical convolution in terms of the inner product between a spherical signal and a rotated spherical filter. Both the signal f and filter ψ are considered to be vector-valued functions on the sphere, i.e. $f : S^2 \rightarrow \mathbb{R}^K$ for K channels. We define S^2 convolution as:

$$\psi * f(R) = \int_{S^2} \sum_{k=1}^K \psi_k(R^{-1}x) f_k(x) dx \quad (1)$$

It is important to note that in our definition, the result of convolution, $\psi * f$ is a function on the rotation group SO(3) and not a function on the sphere S^2 . As such, the convolution operation used in the following layers is slightly different; it is the SO(3) convolution:

$$\psi * f(R) = \int_{\text{SO}(3)} \sum_{k=1}^K \psi_k(R^{-1}R') f_k(R') dR' \quad (2)$$

Our definition for spherical convolution differs from the more common one given by (Driscoll & Healy, 1994), where the result of convolution is a function on the sphere. We found this definition too limiting, because it amounts to convolving with a filter that is rotationally symmetric around the north pole. Other constructions that have as output a function on S^2 cannot be made to be equivariant.

Rotating a fixed filter over the sphere to detect patterns at every position and orientation on the sphere makes sense, because (by hypothesis) patterns appear with equal likelihood in every position and orientation. For this logic to hold also for the internal representations in the network, it is critical to show that the network layers are equivariant.

To this end we define the rotation operator $L_R f(x) = f(R^{-1}x)$ for functions on the sphere or SO(3). It is easy to show, using definitions above and the invariance of the Haar measure (Nachbin, 1965) that $\psi * (L_R f) = L_R(\psi * f)$ for both the S^2 and SO(3) convolution. In other words, these operations are equivariant.

4. Spherical FFT and Convolution Theorem

It is well known that planar convolutions can be computed efficiently using the Fast Fourier Transform (FFT). The Fourier theorem states that the Fourier transform of the convolution equals the element-wise product of the Fourier

transforms, i.e. $\widehat{\psi * f} = \hat{\psi} \cdot \hat{f}$. Since the FFT can be computed in $O(n \log n)$ time and the element-wise product has linear complexity, implementing the convolution using FFTs is asymptotically faster than the naive $O(n^2)$ spatial implementation.

For functions on the sphere and rotation group, there is an analogous transform, which we will refer to as the generalized Fourier transform (GFT) and a corresponding fast algorithm (GFFT). This transform finds its roots in the representation theory of groups, but due to space constraints we will not go into details here and instead refer to interested reader to Sugiura (1990); Folland (1995).

Abstractly, the GFT is nothing more than the projection of a function on a set of orthogonal basis functions called ‘‘matrix element of irreducible unitary representations’’. For SO(3), these are the Wigner D-functions $D_{mn}^l(R)$ indexed by $l \geq 0$ and $-l \leq m, n \leq l$. For S^2 , these are the spherical harmonics¹ $Y_m^l(x)$ indexed by $l \geq 0$ and $-l \leq m \leq l$.

Denoting the manifold (e.g. S^2 or SO(3)) by X and the basis functions by U^l (which is either vector-valued or matrix-valued), we can write the GFT of a function $f : X \rightarrow \mathbb{R}$ as

$$\hat{f}^l = \int_X f(x) \overline{U^l(x)} dx \quad (3)$$

This integral can be computed efficiently using a G-FFT algorithm (Kostelec & Rockmore, 2007; 2008; Kunis & Potts, 2003; Driscoll & Healy, 1994; Maslen, 1998; Potts et al., 2009). For algorithmic details we refer the interested reader to (Kostelec & Rockmore, 2007).

The inverse SO(3) Fourier transform is defined as:

$$f(R) = \sum_{l=0}^b (2l+1) \sum_{m,n=-l}^l \hat{f}_{mn}^l U_{mn}^l(R), \quad (4)$$

and similarly for S^2 . The maximum frequency b is known as the bandwidth, and is related to the resolution of the spatial grid (Kostelec & Rockmore, 2007).

Using the well-known (in fact, defining) property of the Wigner D-functions that $D^l(R)D^l(R') = D^l(RR')$, it can be shown that the SO(3) Fourier transform satisfies a convolution theorem: $\widehat{\psi * f} = \hat{f} \cdot \hat{\psi}^\dagger$, where \cdot denotes matrix multiplication of two block-diagonal matrices \hat{f} and $\hat{\psi}$. Similarly, using $Y(Rx) = D(R)Y(x)$ and $Y_m^l = D_{m0}^l|_{S^2}$, one can derive an analogous S^2 convolution theorem: $\widehat{\psi * f}^l = \hat{f}^l \cdot \hat{\psi}^{l\dagger}$. That is, the SO(3)-FT of the S^2 convolution (as we have defined it) of two spherical

¹Technically, S^2 is not a group and therefore does not have irreducible representations, but it is a quotient of SO(3) and we have the relation $Y_m^l = D_{m0}^l|_{S^2}$

signals can be computed by taking the outer product of the S^2 -FTs of the signals. This is shown in figure 1.

We were unable to find a reference for the latter version of the S^2 Fourier theorem, but its derivation is analogous to that of the well-known $SO(3)$ Fourier theorem. The simplicity of the result shows that our definition of spherical convolution is natural.

5. Spherical Convolutional Networks

A spherical ConvNet is constructed as follows. The input signal $f : S^2 \rightarrow \mathbb{R}^{K_0}$ is convolved with a set of learnable filters $\psi_j^1, j = 1, \dots, K_1$ using the spherical convolution defined in eq. 1. This produces K_1 feature maps on $SO(3)$, which together form a vector valued function $S^2 \rightarrow \mathbb{R}^{K_1}$. The feature maps are composed with a non-linearity and then convolved again with a set of $SO(3)$ filters ψ_j^2 , etc. At the end, we use a spatial pooling operation followed by a learned linear map and softmax nonlinearity to produce a distribution over classes.

6. Experiments

6.1. Numerical tests of Equivariance

We have shown mathematically that the spherical convolution is equivariant, but this proof assumes that we are dealing with well-behaved (e.g. continuous, differentiable) functions. In reality, we only have a set of samples $f(x^i)$ on some sampling set $\{x^i\}_i$. It is therefore reasonable to ask to what degree the computation we actually perform is equivariant. If there are severe artefacts that increase with network depth, we may find that equivariance is eventually lost and rotational weight sharing loses its effectiveness in a deep network.

To test the equivariance of the $SO(3)$ convolution layer, we sample $n = 500$ random rotations R_i and n feature maps f_i with $K = 10$ channels. We then compute the average discrepancy $\Delta = \frac{1}{n} \sum_{i=1}^n \|L_{R_i} \Phi(f_i) - \Phi(L_{R_i} f_i)\|^2 / \|\Phi(f_i)\|^2$, where Φ denotes an $SO(3)$ convolution layer. As shown in figure 2, the discrepancy remains manageable for a range of signal resolutions of interest.

In our second experiment we replace the group convolution layer with a ReLU network with $L = 1, \dots, 10$ layers. The error increases rapidly to about 10^{-2} in the first two layers and then stays there.

6.2. Spherical MNIST

To find out if the inductive bias of a spherical CNN is indeed useful for classifying data with a rotationally invariant label function, we constructed a spherical MNIST dataset. The dataset is constructed by projecting digits onto

the sphere, followed by an optional random rotation.

We use a conventional planar convolutional network as a baseline. The network consists of two conv+relu layers, followed by a fully connected layer and a softmax. We use 5×5 filters, $K = 58, 114, 10$ channels and stride 3 in both convolution layers. Our spherical CNN has a similar architecture: S2Conv-relu-SO3Conv-relu-fc-softmax, with bandwidth $b = 30, 10, 5$ and $k = 100, 200, 10$. Both models have about 165000 parameters.

We evaluate both architectures in three regimes: NR/NR, in which neither the train nor the test data is rotated, R/R in which both are randomly rotated, and NR/R, where only the test data is rotated. Table 1 shows that the spherical CNN shows a slight decrease in performance for non-rotated data, but drastically outperforms the planar CNN on rotated data. Most interestingly, the accuracy of the spherical CNN hardly drops in the NR / R regime, whereas the planar CNN deteriorates to chance level.

	NR / NR	R / R	NR / R
planar	0.99	0.45	0.09
spherical	0.91	0.91	0.85

Table 1. Test accuracy for the networks evaluated on the spherical MNIST dataset. Here R = rotated, NR = non-rotated and X / Y denotes, that the network was trained on X and evaluated on Y.

7. Conclusion

In this paper we have presented spherical convolutional networks, which exploit prior knowledge about rotational symmetry of spherical signals in the same way a conventional convolutional network exploits translation symmetry of planar signals. We have demonstrated mathematically as well as empirically that our spherical and $SO(3)$ convolution layer is equivariant, so that it can be used effectively in deep networks. Furthermore, we have developed an efficient convolution algorithm based on the generalized FFT. Finally, we have demonstrated the effectiveness of the spherical CNN architecture for rotation-invariant classification of spherical signals.

In future work, we aim to apply the spherical CNN to problems in several important scientific problems such as molecular prediction problems (Eickenberg et al., 2017) and global climate and meteorological data. In addition, we think spherical CNNs will be useful for analyzing visual data from omnidirectional cameras as well as 3D sensors.




Figure 1. Spherical convolution in the spectrum. The signal f and the locally-supported filter ψ are Fourier transformed, block-wise tensored, summed over input channels, and finally inverse transformed. Note that because the filter is locally supported, it is faster to use a matrix multiplication (DFT) than an FFT algorithm for it. We parameterize the sphere using spherical coordinates α, β , and $\text{SO}(3)$ with ZYZ-Euler angles α, β, γ .




Figure 2. Δ as a function of the resolution (expressed in terms of bandwidth b) or the number of layers (separated by ReLU activation functions). In both cases, we use $k = 10$ channels. The error increases linearly with the resolution. The error increases sharply after the first ReLU layer, and then stays relatively constant.

References

- Anselmi, Fabio, Leibo, Joel Z, Rosasco, Lorenzo, Mutch, Jim, Tacchetti, Andrea, and Poggio, Tomaso. Magic Materials: A Theory of Deep Hierarchical Architectures for Learning Sensory Representations. *CBCL Paper*, 2013.
- Cohen, Taco S. and Welling, Max. Group equivariant convolutional networks. In *Proceedings of The 33rd International Conference on Machine Learning (ICML)*, volume 48, pp. 2990–2999, 2016.
- Cohen, Taco S and Welling, Max. Steerable CNNs. In *ICLR*, 2017.
- Dieleman, S., De Fauw, J., and Kavukcuoglu, K. Exploiting Cyclic Symmetry in Convolutional Neural Networks. In *International Conference on Machine Learning (ICML)*, 2016.
- Driscoll, J.R. and Healy, D.M. Computing Fourier Transforms and Convolutions on the 2-Sphere. *Advances in Applied Mathematics*, 15(2):202–250, 1994. ISSN 01968858. doi: 10.1006/aama.1994.1008.
- Eickenberg, Michael, Exarchakis, Georgios, Hirn, Matthew, and Mallat, Stephane. Solid Harmonic Wavelet Scattering for Molecular Energy Regression. 2017.
- Folland, G. B. *A Course in Abstract Harmonic Analysis*. CRC Press, 1995.
- Gens, R. and Domingos, P. Deep Symmetry Networks. In *Advances in Neural Information Processing Systems (NIPS)*, 2014.
- Jacobsen, Jörn-Henrik, de Brabandere, Bert, and Smeulders, Arnold W. M. Dynamic Steerable Blocks in Deep Residual Networks. In *BMVC*, 2017.
- Kostelec, Peter J and Rockmore, Daniel N. SOFT: $\text{SO}(3)$ Fourier Transforms. 2007.

- Kostelec, Peter J. and Rockmore, Daniel N. FFTs on the rotation group. *Journal of Fourier Analysis and Applications*, 14(2):145–179, 2008.
- Kunis, Stefan and Potts, Daniel. Fast spherical Fourier algorithms. *Journal of Computational and Applied Mathematics*, 161:75–98, 2003.
- Maslen, David K. Efficient Computation of Fourier Transforms on Compact Groups. *Journal of Fourier Analysis and Applications*, 4(1), 1998.
- Nachbin, L. *The Haar Integral*. 1965.
- Oyallon, E. and Mallat, S. Deep Roto-Translation Scattering for Object Classification. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 2865—2873, 2015.
- Potts, Daniel, Prestin, J, and Vollrath, A. A fast algorithm for nonequispaced Fourier transforms on the rotation group. *Numerical Algorithms*, pp. 1–28, 2009.
- Ravanbakhsh, Siamak, Schneider, Jeff, and Poczos, Barnabas. Deep Learning with Sets and Point Clouds. In *International Conference on Learning Representations (ICLR) – workshop track*, 2017.
- Sifre, Laurent and Mallat, Stephane. Rotation, Scaling and Deformation Invariant Scattering for Texture Discrimination. *IEEE conference on Computer Vision and Pattern Recognition (CVPR)*, 2013.
- Sugiura, Mitsuo. *Unitary Representations and Harmonic Analysis*. John Wiley & Sons, New York, London, Sydney, Toronto, 2nd edition, 1990.
- Worrall, Daniel E, Garbin, Stephan J, Turmukhambetov, Daniyar, and Brostow, Gabriel J. Harmonic Networks: Deep Translation and Rotation Equivariance. In *CVPR*, 2017.
- Zhou, Yanzhao, Ye, Qixiang, Qiu, Qiang, and Jiao, Jianbin. Oriented Response Networks. In *CVPR*, 2017.