
Clebsch–Gordan Nets: a Fully Fourier Space Spherical Convolutional Neural Network

Risi Kondor^{1*} Zhen Lin^{1*} Shubhendu Trivedi^{2*}

¹The University of Chicago ²Toyota Technological Institute
{risi, zlin7}@uchicago.edu, shubhendu@ttic.edu

Abstract

Recent work by Cohen *et al.* [1] has achieved state-of-the-art results for learning spherical images in a rotation invariant way by using ideas from group representation theory and noncommutative harmonic analysis. In this paper we propose a generalization of this work that generally exhibits improved performance, but from an implementation point of view is actually simpler. An unusual feature of the proposed architecture is that it uses the Clebsch–Gordan transform as its only source of nonlinearity, thus avoiding repeated forward and backward Fourier transforms. The underlying ideas of the paper generalize to constructing neural networks that are invariant to the action of other compact groups.

1 Introduction

Despite the many recent breakthroughs in deep learning, we still do not have a satisfactory understanding of how deep neural networks are able to achieve such spectacular performance on a wide range of learning problems. One thing that is clear, however, is that certain architectures pick up on natural invariances in data, and this is a key component to their success. The classic example is of course Convolutional Neural Networks (CNNs) for image classification [2]. Recall that, fundamentally, each layer of a CNN realizes two simple operations: a linear one consisting of convolving the previous layer’s activations with a (typically small) learnable filter, and a nonlinear but pointwise one, such as a ReLU operator². This architecture is sufficient to guarantee *translation equivariance*, meaning that if the input image is translated by some vector t , then the activation pattern in each higher layer of the network will translate by the same amount. Equivariance is crucial to image recognition for two closely related reasons: (a) It guarantees that exactly the same filters are applied to each part of the input image regardless of position. (b) Assuming that finally, at the very top of the network, we add some layer that is *translation invariant*, the entire network will be invariant, ensuring that it can detect any given object equally well regardless of its location.

Recently, a number of papers have appeared that examine equivariance from the theoretical point of view, motivated by the understanding that the natural way to generalize convolutional networks to other types of data will likely lead through generalizing the notion of equivariance itself to other transformation groups [3, 4, 5, 6, 7]. Letting f^s denote the activations of the neurons in layer s of a hypothetical generalized convolution-like neural network, mathematically, equivariance to a group G means that if the inputs to the network are transformed by some transformation $g \in G$, then f^s transforms to $T_g^s(f^s)$ for some fixed set of linear transformations $\{T_g^s\}_{g \in G}$. (Note that in some contexts this is called “covariance”, the difference between the two words being only one of emphasis.)

* Authors are arranged alphabetically

²Real CNNs typically of course have multiple channels, and correspondingly multiple filters per layer, but this does not fundamentally change the network’s invariance properties.

A recent major success of this approach are Spherical CNNs [1][8], which are an $SO(3)$ -equivariant neural network architecture for learning images painted on the sphere³. Learning images on the sphere in a rotation invariant way has applications in a wide range of domains from 360 degree video through drone navigation to molecular chemistry [9, 10, 11, 12, 13, 14, 15]. The key idea in Spherical CNNs is to generalize convolutions using the machinery of noncommutative harmonic analysis: employing a type of generalized $SO(3)$ Fourier transform [16, 17], Spherical CNNs transform the image to a sequence of matrices, and compute the spherical analog of convolution in Fourier space. This beautiful construction guarantees equivariance, and the resulting network attains state of the art results on several benchmark datasets.

One potential drawback of Spherical CNNs of the form proposed in [1], however, is that the nonlinear transform in each layer still needs to be computed in “real space”. Consequently, each layer of the network involves a forward and a backward $SO(3)$ Fourier transform, which is relatively costly, and is a source of numerical errors, especially since the sphere and the rotation group do not admit any regular discretization similar to the square grid for Euclidean space.

Spherical CNNs are not the only context in which the idea of Fourier space neural networks has recently appeared [18, 19, 5, 7]. From a mathematical point of view, the relevance of Fourier theoretic ideas in all these cases is a direct consequence of equivariance, specifically, of the fact that the $\{T_g^s\}_{g \in G}$ operators form a *representation* of the underlying group, in the algebraic sense of the word [20]. In particular, it has been shown that whenever there is a compact group G acting on the inputs of a neural network, there is a natural notion of Fourier transformation with respect to G , yielding a sequence of Fourier matrices $\{F_\ell^s\}_\ell$ at each layer, and the linear operation at layer s will be equivariant to G if and only if it is equivalent to multiplying each of these matrices from the right by some (learnable) filter matrix H_ℓ^s [7]. Any other sort of operation will break equivariance. The spherical convolutions employed in [1] are a special case of this general setup for $SO(3)$, and the ordinary convolutions employed in classical CNNs are a special case for the integer translation group \mathbb{Z}^2 . In all of these cases, however, the issue remains that the nonlinearities need to be computed in “real space”, necessitating repeated forward and backward Fourier transforms.

In the present paper we propose a spherical CNN that differs from [1] in two fundamental ways:

1. While retaining the connection to noncommutative Fourier analysis, we relax the requirement that the activation of each layer of the network needs to be a (vector valued) function on $SO(3)$, requiring only that it be expressible as a collection of some number of $SO(3)$ -covariant vectors (which we call fragments) corresponding to different irreducible representations of the group. In this sense, our architecture is strictly more general than [1].
2. Rather than a pointwise nonlinearity in real space, our network takes the tensor (Kronecker) product of the activations in each layer followed by decomposing the result into irreducible fragments using the so-called *Clebsch–Gordan decomposition*. This way, we get a “fully Fourier space” neural network that avoids repeated forward and backward Fourier transforms.

The resulting architecture is not only more flexible and easier to implement than [1], but our experiments show that it can also perform better on some standard datasets.

The Clebsch–Gordan transform has recently appeared in two separate preprints discussing neural networks for learning physical systems [21, 22]. However, to the best of our knowledge, it has never been proposed as a general purpose nonlinearity for covariant neural networks. In fact, any compact group has a Clebsch–Gordan decomposition (although, due to its connection to angular momentum in physics, the $SO(3)$ case is by far the best known), so, in principle, the methods of the present paper could be applied much broadly, in any situation where one desires to build a neural network that is equivariant to some class of transformations captured by a compact group.

2 Convolutions on the sphere

The simplest example of a covariant neural network is a classical $S + 1$ layer CNN for image recognition. In each layer of a CNN the neurons are arranged in a rectangular grid, so (assuming for simplicity that the network has just one channel) the activation of layer s can be regarded as a function $f^s: \mathbb{Z}^2 \rightarrow \mathbb{R}$, with f^0 being the input image. The neurons compute f^s by taking the

³ $SO(3)$ denotes the group of three dimensional rotations, i.e., the group of 3×3 orthogonal matrices.

cross-correlation⁴ of the previous layer’s output with a small (learnable) filter h^s ,

$$(h^s \star f^{s-1})(x) = \sum_y h^s(y-x) f^{s-1}(y), \quad (1)$$

and then applying a nonlinearity σ , such as the Re-LU operator:

$$f^s(x) = \sigma((h^s \star f^{s-1})(x)). \quad (2)$$

Defining $T_x(h^s)(y) = h^s(y-x)$, which is nothing but h^s translated by x , allows us to equivalently write (1) as

$$(h^s \star f^{s-1})(x) = \langle f^{s-1}, T_x(h^s) \rangle, \quad (3)$$

where the inner product is $\langle f^{s-1}, T_x(h^s) \rangle = \sum_y f^{s-1}(y) T_x(h^s)(y)$. What this formula tells us is that fundamentally each layer of the CNN just does pattern matching: $f^s(x)$ is an indication of how well the part of f^{s-1} around x matches the filter h^s .

Equation 3 is the natural starting point for generalizing convolution to the unit sphere, S^2 . An immediate complication that we face, however, is that unlike the plane, S^2 cannot be discretized by any regular (by which we mean rotation invariant) arrangement of points. A number of authors have addressed this problem in different ways [14, 15]. Instead of following one of these approaches, similarly to recent work on manifold CNNs [23, 24], in the following we simply treat each f^s and the corresponding filter h^s as continuous functions on the sphere, $f^s(\theta, \phi)$ and $h^s(\theta, \phi)$, where θ and ϕ are the polar and azimuthal angles. We allow both these functions to be complex valued, the reason for which will become clear later.

The inner product of two complex valued functions on the surface of the sphere is given by the formula

$$\langle g, h \rangle_{S^2} = \frac{1}{4\pi} \int_0^{2\pi} \int_{-\pi}^{\pi} [g(\theta, \phi)]^* h(\theta, \phi) \cos \theta d\theta d\phi, \quad (4)$$

where $*$ denotes complex conjugation. Further, h (dropping the layer index for clarity) can be moved to any point (θ_0, ϕ_0) on S^2 by taking $h'(\theta, \phi) = h(\theta - \theta_0, \phi - \phi_0)$. This suggests that the generalization of 3 to the sphere should be

$$(h \star f)(\theta_0, \phi_0) = \frac{1}{4\pi} \int_0^{2\pi} \int_{-\pi}^{\pi} [h(\theta - \theta_0, \phi - \phi_0)]^* f(\theta, \phi) \cos \theta d\theta d\phi. \quad (5)$$

Unfortunately, this generalization would be *wrong*, because it does not take into account that h can also be rotated around a third axis. The correct way to generalize cross-correlations to the sphere is to define $h \star f$ as a function *on the rotation group itself*, i.e., to set

$$(h \star f)(R) = \frac{1}{4\pi} \int_0^{2\pi} \int_{-\pi}^{\pi} [h_R(\theta, \phi)]^* f(\theta, \phi) \cos \theta d\theta d\phi \quad R \in \text{SO}(3), \quad (6)$$

where h_R is h rotated by R , expressible as

$$h_R(x) = h(R^{-1}x), \quad (7)$$

with x being the point on the sphere at position (θ, ϕ) (c.f. [25][1]).

2.1 Fourier space filters and activations

Cohen et al.[1] observe that the double integral in (6) would be extremely inconvenient to compute in a neural network. As mentioned, in the case of the sphere, just finding the right discretizations to represent f and h is already problematic. As an alternative, it is natural to represent both these functions in terms of their spherical harmonic expansions

$$f(\theta, \phi) = \sum_{\ell=0}^{\infty} \sum_{m=-\ell}^{\ell} \hat{f}_{\ell}^m Y_{\ell}^m(\theta, \phi) \quad h(\theta, \phi) = \sum_{\ell=0}^{\infty} \sum_{m=-\ell}^{\ell} \hat{h}_{\ell}^m Y_{\ell}^m(\theta, \phi). \quad (8)$$

⁴Convolution and cross-correlation are closely related mathematical concepts that are somewhat confounded in the deep learning literature. In this paper we are going to be a little more precise and focus on cross-correlation, because, despite their name, that is what CNNs actually compute.

Here, $Y_\ell^m(\theta, \phi)$ are the well known spherical harmonic functions indexed by $\ell = 0, 1, 2, \dots$ and $m \in \{-\ell, -\ell + 1, \dots, \ell\}$. The spherical harmonics form an orthonormal basis for $L_2(S^2)$, so (8) can be seen as a kind of Fourier series on the sphere, in particular, the elements of the f_0, f_1, f_2, \dots coefficient vectors can be computed relatively easily by

$$\widehat{f}_\ell^m = \frac{1}{4\pi} \int_0^{2\pi} \int_{-\pi}^{\pi} f(\theta, \phi) Y_\ell^m(\theta, \phi) \cos \theta \, d\theta \, d\phi,$$

and similarly for h . Similarly to usual Fourier series, in practical scenarios spherical harmonic expansions are computed up to some limiting “frequency” L , which depends on the desired resolution.

Noncommutative harmonic analysis [26, 27] tells us that functions on the rotation group also admit a type of generalized Fourier transform. Given a function $g: \text{SO}(3) \rightarrow \mathbb{C}$, the Fourier transform of g is defined as the collection of *matrices*

$$G_\ell = \frac{1}{4\pi} \int_{\text{SO}(3)} g(R) \rho_\ell(R) \, d\mu(R) \quad \ell = 0, 1, 2, \dots, \quad (9)$$

where $\rho_\ell: \text{SO}(3) \rightarrow \mathbb{C}^{(2\ell+1) \times (2\ell+1)}$ are fixed matrix valued functions called the irreducible representations of $\text{SO}(3)$, sometimes also called Wigner D-matrices. Here μ is a fixed measure called the Haar measure that just hides factors similar to the $\cos \theta$ appearing in (4). For future reference we also note that one dimensional irreducible representation ρ_0 is the constant representation $\rho_0(R) = (1)$. The inverse Fourier transform is given by

$$g(R) = \sum_{\ell=0}^{\infty} \text{tr} [G_\ell \rho_\ell(R^{-1})] \quad R \in \text{SO}(3).$$

While the spherical harmonics can be chosen to be real, the $\rho_\ell(R)$ representation matrices are inherently complex valued. This is the reason that we allow all other quantities, including the f^s activations and h^s filters to be complex, too.

Remarkably, the above notions of harmonic analysis on the sphere and the rotation group are closely related. In particular, it is possible to show that each Fourier component of the spherical cross correlation (6) that we are interested in computing is given simply by the outer product

$$[\widehat{h \star f}]_\ell = \widehat{f}_\ell \cdot \widehat{h}_\ell^\dagger \quad \ell = 0, 1, 2, \dots, L, \quad (10)$$

where \dagger denotes the conjugate transpose (Hermitian conjugate) operation. Cohen et al.’s Spherical CNNs [1] are essentially based on this formula. In particular, they argue that instead of the continuous function f , it is more expedient to regard the components of the $\widehat{f}_0, \widehat{f}_1, \dots, \widehat{f}_L$ vectors as the “activations” of their neural network, while the learnable weights or filters are the $\widehat{h}_0, \widehat{h}_1, \dots, \widehat{h}_L$ vectors. Computing spherical convolutions in Fourier space then reduces to just computing a few outer products. Layers $s = 2, 3, \dots, S$ of the Spherical CNN operate similarly, except that f^{s-1} is a function on $\text{SO}(3)$, so (6) must be replaced by cross-correlation on $\text{SO}(3)$ itself, and h must also be a function on $\text{SO}(3)$ rather than just the sphere. Fortunately, the resulting cross-correlation formula is almost exactly the same:

$$[\widehat{h \star f}]_\ell = F_\ell \cdot H_\ell^\dagger \quad \ell = 0, 1, 2, \dots, L, \quad (11)$$

apart from the fact that now F_ℓ and H_ℓ are matrices (see [1] for details).

3 Generalized spherical CNNs

The starting point for our Generalized Spherical CNNs is the Fourier space correlation formula (10). In contrast to [1], however, rather than the geometry, we concentrate on its algebraic properties, in particular, its behavior under rotations. It is well known that if we rotate a spherical function by some $R \in \text{SO}(3)$ as in (7), then each vector of its spherical harmonic expansion just gets multiplied with the corresponding Wigner D-matrix:

$$\widehat{f}_\ell \mapsto \rho_\ell(R) \cdot \widehat{f}_\ell. \quad (12)$$

For functions on $\text{SO}(3)$, the situation is similar. If $g: \text{SO}(3) \rightarrow \mathbb{C}$, and g' is the rotated function $g'(R') = g(R^{-1}R')$, then the Fourier matrices of g' are $G'_\ell = \rho_\ell(R) G_\ell$. The following proposition shows that the matrices output by the (10) and (11) cross-correlation formulae behave analogously.

Proposition 1 Let $f: S^2 \rightarrow \mathbb{C}$ be an activation function that under the action of a rotation R transforms as (7), and let $h: S^2 \rightarrow \mathbb{C}$ be a filter. Then, each Fourier component of the cross correlation (6) transforms as

$$[\widehat{h \star f}]_\ell \mapsto \rho_\ell(R) \cdot [\widehat{h \star f}]_\ell. \quad (13)$$

Similarly, if $f', h': \text{SO}(3) \rightarrow \mathbb{C}$, then $\widehat{h' \star f'}$ (as defined in (11)) transforms the same way.

Equation (12) describes the behavior of spherical harmonic *vectors* under rotations, while (15) describes the behavior of Fourier *matrices*. However, the latter is equivalent to saying that each column of the matrices separately transforms according to (12). One of the key ideas of the present paper is to take this property as the basis for the definition of covariance to rotations in neural nets. Thus we have the following definition.

Definition 1 Let \mathcal{N} be an $S+1$ layer feed-forward neural network whose input is a spherical function $f^0: S^2 \rightarrow \mathbb{C}^d$. We say that \mathcal{N} is a **generalized SO(3)-covariant spherical CNN** if the output of each layer s can be expressed as a collection of vectors

$$\widehat{f}^s = \left(\underbrace{\widehat{f}_{0,1}^s, \widehat{f}_{0,2}^s, \dots, \widehat{f}_{0,\tau_0^s}}_{\ell=0}, \underbrace{\widehat{f}_{1,1}^s, \widehat{f}_{1,2}^s, \dots, \widehat{f}_{1,\tau_1^s}}_{\ell=1}, \dots, \dots, \underbrace{\widehat{f}_{L,\tau_L^s}}_{\ell=L} \right), \quad (14)$$

where each $\widehat{f}_{\ell,j}^s \in \mathbb{C}^{2\ell+1}$ is a ρ_ℓ -covariant vector in the sense that if the input image is rotated by some rotation R , then $\widehat{f}_{\ell,j}^s$ transforms as

$$\widehat{f}_{\ell,j}^s \mapsto \rho_\ell(R) \cdot \widehat{f}_{\ell,j}^s. \quad (15)$$

We call the individual $\widehat{f}_{\ell,j}^s$ vectors the **irreducible fragments** of \widehat{f}^s , and the integer vector $\tau^s = (\tau_0^s, \tau_1^s, \dots, \tau_L^s)$ counting the number of fragments for each ℓ the **type** of \widehat{f}^s .

There are a few things worth noting about Definition 1. First, since the (15) maps are linear, clearly any SO(3)-covariant spherical CNN is equivariant to rotations, as defined in the introduction. Second, since in [1] the inputs are functions on the sphere, whereas in higher layers the activations are functions on SO(3), their architecture is a special case of Definition 1 with $\tau^0 = (1, 1, \dots, 1)$ and $\tau^s = (1, 3, 5, \dots, 2L+1)$ for $s \geq 1$.

Finally, by the theorem of complete reducibility of representations of compact groups, *any* f^s that transforms under rotations linearly is reducible into a sequence of irreducible fragments as in (14). This means that (14) is really the most general possible form for an SO(3) equivariant neural network. As we remarked in the introduction, technically, the terms “equivariant” and “covariant” map to the same concept. The difference between them is one of emphasis. We use the term “equivariant” when we have the same group acting on two objects in a way that is qualitatively similar, as in the case of the rotation group acting on functions on the sphere and on cross-correlation functions on SO(3). We use the term “covariant” if the actions are qualitatively different, as in the case of rotations of functions on the sphere and the corresponding transformations (15) of the irreducible fragments in a neural network.

To fully define our neural network, we need to describe three things: 1. The form of the linear transformations in each layer involving learnable weights, 2. The form of the nonlinearity in each layer, 3. The way that the final output of the network can be reduced to a vector that is rotation *invariant*, since that is our ultimate goal. The following subsections describe each of these components in turn.

3.1 Covariant linear transformations

In a covariant neural network architecture, the linear operation of each layer must be covariant. As described in the Introduction, in classical CNNs, convolution automatically satisfies this criterion. In the more general setting of covariance to the action of compact groups, the problem was studied in [7]. The specialization of their result to our case is the following.

Proposition 2 Let \widehat{f}^s be an SO(3)-covariant activation function of the form (14), and $\widehat{g}^s = \mathcal{L}(\widehat{f}^s)$ be a linear function of \widehat{f}^s written in a similar form. Then \widehat{g}^s is SO(3)-covariant if and only if each $\widehat{g}_{\ell,j}^s$ fragment is a linear combination of fragments from \widehat{f}^s with the same ℓ .

Proposition 2 can be made more clear by stacking all fragments of \widehat{f} corresponding to ℓ into a $(2\ell + 1) \times \tau_\ell^s$ dimensional matrix F_ℓ^s , and doing the same for \widehat{g} . Then the proposition tells us simply that

$$G_\ell^s = F_\ell^s W_\ell^s \quad \ell = 0, 1, 2, \dots, L \quad (16)$$

for some sequence of complex valued matrices W_0^s, \dots, W_L^s . Note that W_ℓ^s does not necessarily need to be square, i.e., the number of fragments in f and \widehat{g} corresponding to ℓ might be different. In the context of a neural network, the entries of the W_ℓ^s matrices are learnable parameters.

Note that the Fourier space cross-correlation formulae (10) and (11) are special cases of (16) corresponding to taking $W_\ell = \widehat{h}_\ell^\dagger$ or $W_\ell = H_\ell^\dagger$. The case of general W_ℓ does not have such an intuitive interpretation in terms of cross-correlation. What the (16) lacks in interpretability it makes up for in terms of generality, since it provides an extremely simple and flexible way of inducing SO(3)-covariant linear transformations in neural networks.

3.2 Covariant nonlinearities: the Clebsch–Gordan transform

Differentiable nonlinearities are essential for the operation of multi-layer neural networks. Formulating covariant nonlinearities in Fourier space, however, is more challenging than formulating the linear operation. This is the reason that most existing group equivariant neural networks perform this operation in “real space”. However, as discussed above, moving back and forth between real space and the Fourier domain comes at a significant cost and leads to a range of complications involving quadrature on the transformation group and numerical errors.

One of the key contributions of the present paper is to propose a fully Fourier space nonlinearity based on the Clebsch–Gordan transform. In representation theory, the Clebsch–Gordan decomposition arises in the context of decomposing the tensor (i.e., Kronecker) product of irreducible representations into a direct sum of irreducibles. In the specific case of SO(3), it takes form

$$\rho_{\ell_1}(R) \otimes \rho_{\ell_2}(R) = C_{\ell_1, \ell_2} \left[\bigoplus_{\ell=|\ell_1-\ell_2|}^{\ell_1+\ell_2} \rho_\ell(R) \right] C_{\ell_1, \ell_2}^\top, \quad R \in \text{SO}(3),$$

where C_{ℓ_1, ℓ_2} are fixed matrices. Equivalently, letting $C_{\ell_1, \ell_2, \ell}$ denote the appropriate block of columns of C_{ℓ_1, ℓ_2} ,

$$\rho_\ell(R) = C_{\ell_1, \ell_2, \ell}^\top [\rho_{\ell_1}(R) \otimes \rho_{\ell_2}(R)] C_{\ell_1, \ell_2, \ell}.$$

The CG-transform is well known in physics, because it is intimately related to the algebra of angular momentum in quantum mechanics, and the entries of the $C_{\ell_1, \ell_2, \ell}$ matrices can be computed relatively easily. The following Lemma explains why this construction is relevant to creating Fourier space nonlinearities.

Lemma 3 *Let \widehat{f}_{ℓ_1} and \widehat{f}_{ℓ_2} be two ρ_{ℓ_1} resp. ρ_{ℓ_2} covariant vectors, and ℓ be any integer between $|\ell_1 - \ell_2|$ and $\ell_1 + \ell_2$. Then*

$$\widehat{g}_\ell = C_{\ell_1, \ell_2, \ell}^\top [\widehat{f}_{\ell_1} \otimes \widehat{f}_{\ell_2}] \quad (17)$$

is a ρ_ℓ -covariant vector.

Exploiting Lemma 3, the nonlinearity used in our generalized Spherical CNNs consists of computing (17) between all pairs of fragments. In matrix notation,

$$G_\ell^s = \bigsqcup_{|\ell_1-\ell_2| \leq \ell \leq \ell_1+\ell_2} C_{\ell_1, \ell_2, \ell}^\top [F_{\ell_1}^s \otimes F_{\ell_2}^s], \quad (18)$$

where \bigsqcup denotes merging matrices horizontally. Note that this operation increases the size of the activation substantially: the total number of fragments is squared, which can potentially be problematic, and is addressed in the following subsection.

The Clebsch–Gordan decomposition has recently appeared in two preprints discussing neural networks for learning physical systems [21, 22]. However, to the best of our knowledge, in the present context of a general purpose nonlinearity, it has never been proposed before. At first sight, the computational cost it would appear that the computational cost of (17) (assuming that $C_{\ell_1, \ell_2, \ell}$ has been precomputed) is $(2\ell_1 + 1)(2\ell_2 + 1)(2\ell + 1)$. However, $C_{\ell_1, \ell_2, \ell}$ is actually sparse, in particular $[C_{\ell_1, \ell_2, \ell}]_{(m_1, m_2), m} = 0$ unless $m_1 + m_2 = m$. Denoting the total number of scalar entries in the F_{ℓ}^s

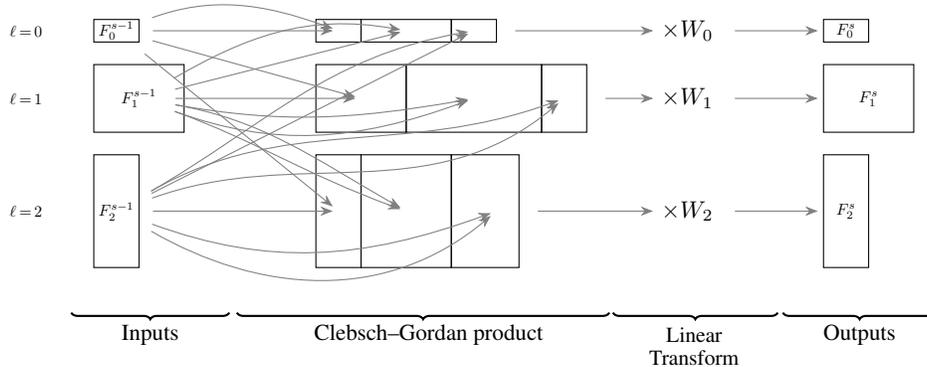


Figure 1: Schematic of a single layer of the Clebsch–Gordan network.

matrices by N , this reduces the complexity of computing (18) to $O(N^2L)$. While the CG transform is not currently available as a differentiable operator in any of the major deep learning software frameworks, we have developed and will publicly release a C++ PyTorch extension for it.

A more unusual feature of the CG nonlinearity is that its essentially quadratic nature. Quadratic nonlinearities are not commonly used in deep neural networks. Nonetheless, our experiments indicate that the CG nonlinearity is effective in the context of learning spherical images. It is also possible to use higher CG powers, although the computational cost will obviously increase.

3.3 Limiting the number of channels

In a covariant network, each individual \hat{f}_ℓ^s fragment is effectively a separate channel. In this sense, the quadratic increase in the number of channels after the CG-transform can be seen as a natural broadening of the network to capture more complicated features. Naturally, allowing the number of channels to increase quadratically in each layer would be untenable, though.

Following the results of Section 3.1, the natural way to counteract the exponential increase in the number of channels is follow the CG-transform with another learnable linear transformation that reduces the number of fragments for each ℓ to some fixed maximum number $\bar{\tau}_\ell$. In fact, this linear transformation can *replace* the transformation of Section 3.1. Whereas in conventional neural networks the linear transformation always precedes the nonlinear operation, in Clebsch–Gordan networks it is natural to design each layer so as to perform the CG-transform first, and then the convolution like step (16), which will limit the number of fragments.

3.4 Final invariant layer

After the $S - 1$ 'th layer, the activations of our network will be a series of matrices $F_0^{S-1}, \dots, F_L^{S-1}$, each transforming under rotations according to $F_\ell^{S-1} \mapsto \rho_\ell(R) F_\ell^{S-1}$. Ultimately, however, the objective of the network is to output a vector that is *invariant* with respect rotations, i.e., a collection of *scalars*. In our Fourier theoretic language, this simply corresponds to the $\hat{f}_{0,j}^S$ fragments, since the $\ell = 0$ representation is constant, and therefore the elements of F_0^S are invariant. Thus, the final layer can be similar to the earlier ones, except that it only needs to output this single (single row) matrix.

Note that in contrast to other architectures such as [1] that involve repeated forward and backward transforms, thanks to their fully Fourier nature, for Clebsch–Gordan nets, in both training and testing, the elements of F_0^S are guaranteed to be invariant to rotations of arbitrary magnitude not just approximately, but in the *exact* sense, up to limitations of finite precision arithmetic. This is a major advantage of Clebsch–Gordan networks compared to other covariant architectures.

3.5 Summary of algorithm

In summary, our Spherical Clebsch–Gordan network is an $S + 1$ layer feed-forward neural network in which apart from the initial spherical harmonic transform, every other operation is a simple matrix operation. In the forward pass, the network operates as follows.

1. The inputs to the network are n_{in} functions $f_1^0, \dots, f_{n_{\text{in}}}^0: S^2 \rightarrow \mathbb{C}$. For example, for spherical color images, f_1^0, f_2^0 and f_3^0 might encode the red, green and blue channels. For generality, we treat these functions as complex valued, but of course they may also be real. The activation of layer $s = 0$ is the union of the spherical transforms of these functions $f_1^0, \dots, f_{n_{\text{in}}}^0$ up to some band limit L , i.e.,

$$[\widehat{f}_{\ell,j}^0]_m = \frac{1}{4\pi} \int_0^{2\pi} \int_{-\pi}^{\pi} f_j^0(\theta, \phi)^* Y_\ell^m(\theta, \phi) \cos \theta d\theta d\phi. \quad (19)$$

Therefore, the type of \widehat{f}^0 is $\tau^0 = (n_{\text{in}}, n_{\text{in}}, \dots, n_{\text{in}})$, and \widehat{f}^0 is stored as a collection of $L + 1$ matrices $\{F_0^0, F_1^0, \dots, F_L^0\}$ of sizes $1 \times n_{\text{in}}, 3 \times n_{\text{in}}, 5 \times n_{\text{in}}, \dots$

2. For layers $s = 1, 2, \dots, S - 1$, the Fourier space activation $\widehat{f}^s = (F_0^s, F_1^s, \dots, F_L^s)$ is computed as follows:

- (a) We form all possible Kronecker products

$$G_{\ell_1, \ell_2}^s = F_{\ell_1}^{s-1} \otimes F_{\ell_2}^{s-1} \quad 0 \leq \ell_1 \leq \ell_2 \leq L. \quad (20)$$

Note that the size of G_{ℓ_1, ℓ_2}^s is $(2\ell_1 + 1)(2\ell_2 + 1) \times (\tau_{\ell_1}^{s-1} \tau_{\ell_2}^{s-1})$.

- (b) Each G_{ℓ_1, ℓ_2}^s is decomposed into ρ_ℓ -covariant blocks by

$$[G_{\ell_1, \ell_2}^s]_\ell = C_{\ell_1, \ell_2, \ell}^\dagger G_{\ell_1, \ell_2}^s, \quad (21)$$

where $C_{\ell_1, \ell_2, \ell}^\dagger$ is the inverse Clebsch–Gordan matrix as in (17).

- (c) All $[G_{\ell_1, \ell_2}^s]_\ell$ blocks with the same ℓ are concatenated into a large matrix $H_\ell^s \in \mathbb{C}^{(2\ell+1) \times \tau_\ell^s}$, and this is multiplied by the weight matrix $W_\ell^s \in \mathbb{C}^{\tau_\ell^s \times \tau_\ell^s}$ to give

$$F_\ell^s = H_\ell^s W_\ell^s \quad \ell = 0, 1, \dots, L. \quad (22)$$

3. The operation of layer S is similar, except that the output type is $\tau^S = (n_{\text{out}}, 0, 0, \dots, 0)$, so components with $\ell > 0$ do not need to be computed. By construction, the entries of $F_0^s \in \mathbb{C}^{1 \times n_{\text{out}}}$ are SO(3)-invariant scalars, i.e., they are invariant to the simultaneous rotation of the $f_1^0, \dots, f_{n_{\text{in}}}^0$ inputs. These scalars may be passed on to a fully connected network or plugged directly into a loss function.

The learnable parameters of the network are the $\{W_\ell^s\}_{s, \ell}$ weight matrices. The matrices are initialized with random complex entries. As usual, the network is trained by backpropagating the gradient of the loss from layer S to each of the weight matrices.

4 Experiments

In this section we describe experiments that give a direct comparison with those reported by Cohen *et al.* [1]. We choose these experiments as the Spherical CNN proposed in [1] is the only direct competition to our method. Besides, the comparison is also instructive for two different reasons: Firstly, while the procedure used in [1] is exactly equivariant in the discrete case, for the continuous case they use a discretization which causes their network to partially lose equivariance with changing bandwidth and depth, whereas our method is always equivariant in the exact sense. Secondly, owing to the nature of their architecture and discretization, [1] use a more traditional non-linearity i.e. the ReLU, which is also quite powerful. In our case, to maintain full covariance and to avoid the quadrature, we use an unconventional quadratic non-linearity in fourier space. Because of these two differences, the experiments will hopefully demonstrate the advantages of avoiding the quadrature and maintaining full equivariance despite using a purportedly weaker nonlinearity.

Rotated MNIST on the Sphere We use a version of MNIST in which the images are painted onto a sphere and use two instances as in [1]: One in which the digits are projected onto the northern hemisphere and another in which the digits are projected on the sphere and are also randomly rotated.

The baseline model is a classical CNN with 5×5 filters and 32, 64, 10 channels with a stride of 3 in each layer (roughly 68K parameters). This CNN is trained by mapping the digits from the sphere back onto the plane, resulting in nonlinear distortions. The second model we use to compare to is the

Spherical CNN proposed in [1]. For this method, we use the same architecture as reported by the authors i.e. having layers S^2 convolution – ReLU – $SO(3)$ convolution – ReLU – Fully connected layer with bandwidths 30, 10 and 6, and the number of channels being 20, 40 and 10 (resulting in a total of 58K parameters).

For our method we use the following architecture: We set the bandlimit $L_{max} = 10$, and keep $\tau_l = \lceil \frac{12}{\sqrt{2l+1}} \rceil$, using a total of 5 layers as described in section 3.5, followed by a fully connected layer of size 256 by 10. We use a variant of batch normalization that preserves covariance in the Fourier layers. This method takes an expanding average of the standard deviation for a particular fragment for all examples seen during training till then and divide the fragment by it (in testing, use the average from training); the parameter corresponding to the mean in usual batch normalization is kept to be zero as anything else will break covariance. Finally, we concatenate the output of each F_0^s in each internal layer (length 24 each, as each is $\tau_0 = 12$ complex numbers), as well as the original coefficient at $l = 0$ (length 2), into a $SO(3)$ invariant vector of length 122. (We observed that having these skip connections was crucial to facilitate smooth training.) After that, we use the usual batch normalization [28] on the concatenated results before feeding it into fully connected layers of length 256, a dropout layer with dropout probability 0.5, and finally a linear layer to 10 output nodes. The total number of parameters was 285772, the network was trained by using the ADAM optimization procedure [29] with a batch size of 100 and a learning rate of 5×10^{-4} . We also used L2 weight decay of 1×10^{-5} on the trainable parameters.

We report three sets of experiments: For the first set both the training and test sets were not rotated (denoted NR/NR), for the second, the training set was not rotated while the test was randomly rotated (NR/R) and finally when both the training and test sets were rotated (denoted R/R).

Method	NR/NR	NR/R	R/R
Baseline CNN	97.67	22.18	12
Cohen <i>et al.</i>	95.59	94.62	93.4
Ours (FFS2CNN)	96.4	96	96.6

We observe that the baseline model’s performance deteriorates in the three cases, more or less reducing to random chance in the R/R case. While our results are better than those reported in [1], they also have another characteristic: they remain roughly the same in the three regimes, while those of [1] slightly worsen. We think this might be a result of the loss of equivariance in their method.

Atomization Energy Prediction Next, we apply our framework to the QM7 dataset [30, 31], where the goal is to regress over atomization energies of molecules given atomic positions (p_i) and charges (z_i). Each molecule contains up to 23 atoms of 5 types (C, N, O, S, H). More details about the representations used, baseline models, as well as the architectural parameters are provided in the appendix. We use the Coulomb Matrix (CM) representation proposed by [31], which is rotation and translation invariant but not permutation invariant. The Coulomb matrix $C \in \mathbb{R}^{N \times N}$ is defined such that for a pair of atoms $i \neq j$, $C_{ij} = (z_i z_j) / (|p_i - p_j|)$, which represents the Coulomb repulsion, and for atoms $i = j$, $C_{ii} = 0.5 z_i^{2.4}$, which denotes the atomic energy due to charge. To test our algorithm we use the same set up as in [1]: We define a sphere S_i around p_i for each atom i . Ensuring uniform radius across atoms and molecules and ensuring no intersections amongst spheres during training, we define potential functions $U_z(x) = \sum_{j \neq i, z_j = z} \frac{z_i z_j}{|x - p_i|}$ for every z and for every x on S_i . This yields a T channel spherical signal for each atom in a molecule. This signal is then discretized using Driscoll-Healy [32] grid using a bandwidth of $b = 10$. This gives a sparse tensor representation of dimension $N \times T \times 2b \times 2b$ for every molecule.

Our spherical CNN architecture has the same parameters and hyperparameters as in the previous subsection except that $\tau_l = 15$ for all layers, increasing the number of parameters to 1.1 M. Following [1], we share weights amongst atoms and each molecule is represented as a $N \times F$ tensor where F represents F_0^s scalars concatenated together. Finally, we use the approach proposed in [33] to ensure permutation invariance. The feature vector for each atom is projected onto 150 dimensions using a MLP. These embeddings are summed over atoms, and then the regression target is trained using another MLP having 50 hidden units. Both of these MLPs are jointly trained. The final results are presented below, which show that our method outperforms the Spherical CNN of Cohen *et al.*. The only method that delivers better performance is a MLP trained on randomly permuted Coulomb matrices [34], and as [1] point out, this method is unlikely to scale to large molecules as it needs a large sample of random permutations, which grows rapidly with N .

Method	RMSE
MLP/Random CM [34]	5.96
LGKA (RF) [35]	10.82
RBF Kernels/Random CM [34]	11.42
RBF Kernels/Sorted CM [34]	12.59
MLP/Sorted CM [34]	16.06
Spherical CNN [1]	8.47
Ours (FFS2CNN)	7.97

3D Shape Recognition Finally, we report results for shape classification using the SHREC17 dataset [36], which is a subset of the larger ShapeNet dataset [37] having roughly 51300 3D models spread over 55 categories. It is divided into a 70/10/20 split for train/validation/test. Two versions of this dataset are available: A regular version in which the objects are consistently aligned and another where the 3D models are perturbed by random rotations. Following [1] we focus on the latter version, as well as represent each 3D mesh as a spherical signal by using a ray casting scheme. For each point on the sphere, a ray towards the origin is sent which collects the ray length, cosine and sine of the surface angle. In addition to this, ray casting for the convex hull of the mesh gives additional information, resulting in 6 channels. The spherical signal is discretized using the Discroll-Healy grid [32] with a bandwidth of 128. We use the code provided by [1] for generating this representation.

We use a ResNet style architecture, but with the difference that the full input is not fed back but rather different frequency parts of it. We consider $L_{max} = 14$, and first train a block only till $L = 8$ using $\tau_l = 10$ using 3 layers. The next block consists of concatenating the fragments obtained from the previous block and training for two layers till $L = 10$, repeating this process till L_{max} is reached. These later blocks use $\tau_l = 8$. As earlier, we concatenate the F_0^s scalars from each block to form the final output layer, which is connected to 55 nodes forming a fully connected layer. We use Batch Normalization in the final layer, and the normalization discussed in 4 in the Fourier layers. The model was trained with ADAM using a batch size of 100 and a learning rate of 5×10^{-4} , using L2 weight decay of 0.0005 for regularization. We compare our results to some of the top performing models on SHREC (which use architectures specialized to the task) as well as the model of Cohen *et al.*. Our method, like the model of Cohen *et al.* is task agnostic and uses the same representation. Despite this, it is able to consistently come second or third in the competition, showing that it affords an efficient method to learn from spherical signals.

Method	P@N	R@N	F1@N	mAP	NDCG
Tatsuma_ReVGG	0.705	0.769	0.719	0.696	0.783
Furuya_DLAN	0.814	0.683	0.706	0.656	0.754
SHREC16-Bai_GIFT	0.678	0.667	0.661	0.607	0.735
Deng_CM-VGG5-6DB	0.412	0.706	0.472	0.524	0.624
Spherical CNNs [1]	0.701	0.711	0.699	0.676	0.756
FFS2CNNs (ours)	0.707	0.722	0.701	0.683	0.756

5 Conclusion

We have presented an $SO(3)$ -equivariant neural network architecture for spherical data that operates completely in Fourier space, circumventing a major drawback of earlier models that need to switch back and forth between Fourier space and “real” space. We achieve this by – rather unconventionally – using the Clebsch-Gordan decomposition as the only source of nonlinearity. While the specific focus is on spheres and $SO(3)$ -equivariance, the approach is more widely applicable, suggesting a general formalism for designing fully Fourier neural networks that are equivariant to the action of any compact continuous group.

References

- [1] T. S. Cohen, M. Geiger, J. Köhler, and M. Welling. Spherical CNNs. *International Conference on Learning Representations*, 2018.
- [2] Y LeCun, B Boser, J. S. Denker, D. Henderson, R. E. Howard, W. Hubbard, and L. D. Jackel. Backpropagation applied to handwritten zip code recognition. *Neural Computation*, 1:541–551, 1989.
- [3] R. Gens. Deep Symmetry Networks. *NIPS 2014*, pages 1–9, 2014.
- [4] T. S. Cohen and M. Welling. Group equivariant convolutional networks. *Proceedings of The 33rd International Conference on Machine Learning*, 48:2990–2999, 2016.
- [5] T. S. Cohen and M. Welling. Steerable cnns. In *ICLR*, 2017.
- [6] S. Ravanbakhsh, J. Schneider, and B. Póczos. Equivariance through parameter-sharing. In *Proceedings of International Conference on Machine Learning*, 2017.
- [7] R. Kondor and S. Trivedi. On the generalization of equivariance and convolution in neural networks to the action of compact groups. 2018.
- [8] C. Esteves, C. Allen-Blanchette, A. Makadia, and K. Daniilidis. Learning SO(3) Equivariant Representations with Spherical CNNs. 2017.
- [9] L. Zelnik-Manor, G. Peters, and P. Perona. Squaring the Circles in Panoramas. *Ieee Iccv*, pages 1292–1299, 2005.
- [10] J. Cruz-Mota, I. Bogdanova, B. Paquier, M. Bierlaire, and J-P Thiran. Scale invariant feature transform on the sphere: Theory and applications. *International Journal of Computer Vision*, 98(2):217–241, 2012.
- [11] Y-C Su, D. Jayaraman, and K. Grauman. Pano2vid: Automatic cinematography for watching 360° videos. *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, 10114 LNCS(1):154–171, 2017.
- [12] W-S Lai, Y. Huang, N. Joshi, C. Buehler, M-H Yang, and S-B Kang. Semantic-driven Generation of Hyperlapse from 360 @BULLET Video. pages 1–12.
- [13] R. Khasanova and P. Frossard. Graph-Based Classification of Omnidirectional Images. 2017.
- [14] W. Boomsma and J. Frellsen. Spherical convolutions and their application in molecular modelling. (Nips):3436–3446, 2017.
- [15] Y-C Su and K. Grauman. Making 360° Video Watchable in 2D: Learning Videography for Click Free Viewing. 2017.
- [16] D. M. Healey, D. N. Rockmore, and S. B. Moore. An fft for the 2-sphere and applications. In *1996 IEEE International Conference on Acoustics, Speech, and Signal Processing Conference Proceedings*, volume 3, pages 1323–1326 vol. 3, May 1996.
- [17] P. J. Kostelec and D. N. Rockmore. Ffts on the rotation group. *Journal of Fourier Analysis and Applications*, 14(2):145–179, Apr 2008.
- [18] D. E. Worrall, S. J. Garbin, D. Turmukhambetov, and G. J. Brostow. Harmonic Networks: Deep Translation and Rotation Equivariance. 2016.
- [19] C. Esteves, C. Allen-Blanchette, X. Zhou, and K. Daniilidis. Polar Transformer Networks. 2017.
- [20] J-P. Serre. *Linear Representations of Finite Groups*, volume 42 of *Graduate Texts in Mathematics*. Springer-Verlag, 1977.
- [21] N. Thomas, T. Smidt, S. Kearnes, L. Yang, L. Li, K. Kohlhoff, and P. Riley. Tensor field networks: Rotation- and translation-equivariant neural networks for 3d point clouds. *arXiv:1802.08219 [cs]*, Feb 2018. arXiv: 1802.08219.
- [22] R. Kondor. N-body networks: a covariant hierarchical neural network architecture for learning atomic potentials. *CoRR*, abs/1803.01588, 2018.
- [23] J. Masci, D. Boscaini, M. M. Bronstein, and P. Vandergheynst. Geodesic convolutional neural networks on riemannian manifolds. 2015.
- [24] F. Monti, D. Boscaini, J. Masci, E. Rodola, J. Svoboda, and M. M. Bronstein. Geometric deep learning on graphs and manifolds using mixture model cnns. 2016.
- [25] B. Gutman, Y. Wang, T. Chan, P. M. Thompson, and A. W. Toga. Shape Registration with Spherical Cross Correlation. *2nd MICCAI Workshop on Mathematical Foundations of Computational Anatomy*, pages 56–67, 2008.
- [26] A. Terras. *Fourier analysis on finite groups and applications*, volume 43 of *London Mathematical Society Student Texts*. Cambridge Univ. Press, 1999.
- [27] P. Diaconis. *Group Representation in Probability and Statistics*, volume 11 of *IMS Lecture Series*. Institute of Mathematical Statistics, 1988.
- [28] S. Ioffe and C. Szefedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. *International Conference on Machine Learning*, 2015.

- [29] D. P. Kingma and J. Ba. Adam: A method for stochastic optimization. *International Conference on Learning Representations*, 2015.
- [30] L. C. Blum and J.-L. Reymond. 970 million druglike small molecules for virtual screening in the chemical universe database gdb-13. *Journal of the American Chemical Society*, 2009.
- [31] M. Rupp, A. Tkatchenko, K.-R. Müller, and O. A. von Lilienfeld. Fast and accurate modeling of molecular atomization energies with machine learning. *Physical Review Letters*, 2012.
- [32] J. R. Driscoll and D. M. Healy. Computing fourier transforms and convolutions on the 2-sphere. *Advances in Applied Mathematics*, 1994.
- [33] M. Zaheer, S. Kottur, S. Ravanbakhsh, B. Póczos, R. Salakhutdinov, , and A. Smola. Deep sets. 2017.
- [34] G. Montavon, K. Hansen, S. Fazli, M. Rupp, F. Biegler, A. Ziehe, A. Tkatchenko, O.A. von Lilienfeld, and K. Müller. Learning invariant representations of molecules for atomization energy prediction. In *NIPS*, 2012.
- [35] A. Raj, A. Kumar, Y. Mroueh, and P.T. Fletcher et al. Local group invariant representations via orbit embeddings. 2016.
- [36] M. Savva, F. Yu, H. Su, A. Kanezaki, T. Furuya, R. Ohbuchi, Z. Zhou, R. Yu, S. Bai, X. Bai, M. Aono, A. Tatsuma, S. Thermos, A. Axenopoulos, G. Th. Papadopoulos, P. Daras, X. Deng, Z. Lian, B. Li, H. Johan, Y. Lu, and S. Mk. Large-scale 3d shape retrieval from shapenet core55. *Eurographics Workshop on 3D Object Retrieval*, 2017.
- [37] A. X. Chang, T. Funkhouser, L. Guibas, P. Hanrahan, Q. Huang, Z. Li, S. Savarese, M. Savva, S. Song, H. Su, J. Xiao, L. Yi, and F. Yu. Shapenet: An information-rich 3d model repository. 2015.