

---

# Steerable 3D Spherical Neurons

---

Pavlo Melnyk<sup>1</sup> Michael Felsberg<sup>1</sup> Mårten Wadenbäck<sup>1</sup>

## Abstract

Emerging from low-level vision theory, steerable filters found their counterpart in prior work on steerable convolutional neural networks equivariant to rigid transformations. In our work, we propose a steerable feed-forward learning-based approach that consists of neurons with spherical decision surfaces and operates on point clouds. Such spherical neurons are obtained by conformal embedding of Euclidean space and have recently been revisited in the context of learning representations of point sets. Focusing on 3D geometry, we exploit the isometry property of spherical neurons and derive a 3D steerability constraint. After training spherical neurons to classify point clouds in a canonical orientation, we use a tetrahedron basis to quadruplicate the neurons and construct rotation-equivariant spherical filter banks. We then apply the derived constraint to interpolate the filter bank outputs and, thus, obtain a rotation-invariant network. Finally, we use a synthetic point set and real-world 3D skeleton data to verify our theoretical findings. The code is available at <https://github.com/pavlo-melnyk/steerable-3d-neurons>.

## 1. Introduction

We present a feed-forward model consisting of steerable 3D neurons for point cloud classification, an important and challenging problem with many applications such as autonomous driving, human-robot interaction, and mixed-reality installations. Constructing rotation-equivariant (steerable) models allows us to use the feature of a given point cloud to synthesize features of the same point cloud with different orientations. Further, rotation-equivariant networks

---

<sup>1</sup>Computer Vision Laboratory, Department of Electrical Engineering, Linköping University, SE-581 83 Linköping, Sweden. Correspondence to: Pavlo Melnyk <pavlo.melnyk@liu.se>, Michael Felsberg <michael.felsberg@liu.se>.

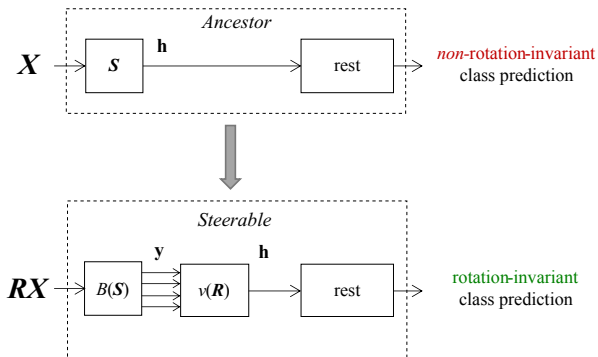


Figure 1. Our approach overview: First, we train a classifier (*Ancestor*) consisting of the spherical neurons (necessarily in the first layer) to classify 3D point clouds in a canonical orientation. We then fix all the learned parameters and construct rotation-equivariant spherical filter banks  $B(S)$  (8) from the first layer parameters. Next, we compute the interpolation coefficients  $v(R)$  (15) to fulfill the steerability constraint (12). The result is a *Steerable* classifier, where  $y$  is equivariant and  $h$  is invariant to rotations of the model input.

enable producing rotation-invariant predictions and, therefore, lowered data augmentation requirements for learning.

In this paper, we achieve the steerability using a conformal embedding to obtain higher-order decision surfaces. Following the motivation in the recent work of Melnyk et al. (2021), we focus on 3D geometry and spherical decision surfaces, arguing for their natural suitability for problems in Euclidean space. We show how a *spherical neuron*, i.e., the hypersphere neuron (Baner et al., 2003b) or its generalization for 3D input point sets — the geometric neuron (Melnyk et al., 2021) — can be turned into a steerable neuron. Making spherical neurons *steerable* adds to the practical value of the prior work by Melnyk et al. (2021).

We prove that the aforementioned spherical neurons in any dimension require only up to first-degree spherical harmonics to accommodate the effect of rotation. This allows us to derive a 3D steerability constraint for such neurons and to describe a recipe for creating a steerable model from a pre-trained classifier (see Figure 1). Using the synthetic Tetris dataset (Thomas et al., 2018) and the skeleton data from the UTKinect-Action3D dataset (Xia et al., 2012), we verify the derived constraint and check its stability with respect to

perturbations in the input.

Importantly, we focus on theoretical aspects of our steerable method to produce rotation-invariant predictions, and show that it is attainable by using both synthetic and real 3D data (Section 5.3). What remains to do is to devise a practical way to achieve it.

The core of our contributions are the novel theoretical results on steerability and equivariance and can be summarized as follows:

(a) We prove that the activation of spherical neurons (Baner et al., 2003b; Melnyk et al., 2021) on rotated input only varies by up to first-degree spherical harmonics in the rotation angle (Section 4.1).

(b) Based on a minimal set of four spherical neurons that are rotated to the corresponding vertices of a regular tetrahedron, we construct a rotation-equivariant spherical filter bank (Section 4.2) and derive the main result of our paper (Section 4.3) — a 3D steerability constraint for spherical neurons (12) and (16).

## 2. Related work

### 2.1. Steerability and equivariance

Steerability is a powerful concept from early vision and image processing (Freeman et al., 1991; Knutsson et al., 1992; Simoncelli et al., 1992; Perona, 1995; Simoncelli & Freeman, 1995; Teo & Hel-Or, 1998) that resonates in the era of deep learning. The utility of steerable filters and the main theorems for their construction are presented in the seminal work of Freeman et al. (1991).

Geometric equivariance in the context of computer vision has been an object of extensive research over the last decades (Van Gool et al., 1995). Equivariance is a necessary property for steerability since steerability requires changing the function output depending on the actions of a certain group. Thus, the output needs to contain a representation of the group that acts on the input. Therefore, the operator must commute with the group (note that the representation might change from input to output). For this reason, Lie theory can be used to study steerable filters for equivariance and invariance (Reisert, 2008).

Nowadays, equivariance-related research extends into deep learning, e.g., the  $SE(3)$ -equivariant models of Fuchs et al. (2020), Thomas et al. (2018), and Zhao et al. (2020), and the  $SO(3)$ -equivariant network of Anderson et al. (2019), as well as Marcos et al. (2017) and Kondor et al. (2018). Also steerable filter concepts are increasingly used in current works. For example, the work of Cohen & Welling (2016) considered image data and proposed a CNN architecture to produce equivariant representations with steerable features,

which involves fewer parameters than traditional CNNs. They considered the dihedral group (*discrete* rotation, translation, and reflection), and the steerable representations in their work are proposed as formation of elementary feature types. One limitation of their approach is that rotations are restricted to four orientations, i.e., by multiples  $\pi/2$ . More recently, Weiler et al. (2018b) utilized group convolutions and introduced steerable filter convolutional neural networks (SFCNNs) operating on images to jointly attain equivariance under translations and discrete rotations. In their work, the filter banks are learned rather than fixed. Further, the work of Weiler & Cesa (2019) proposed a unified framework for  $E(2)$ -equivariant steerable CNNs and presented their general theory.

The steerable CNNs for 3D data proposed by Weiler et al. (2018a) are closely related to our work. The authors employed a combination of scalar, vector and tensor fields as features transformed by  $SO(3)$  representations and presented a model that is equivariant to  $SE(3)$  transformations. They also considered different types of nonlinearities suitable for nonscalar components of the feature space. The novel  $SE(3)$ -equivariant approach by Fuchs et al. (2020) introduced a self-attention mechanism that is invariant to global rotations and translations of its input and solves the limitation of angularly constrained filters in other equivariant methods, e.g., Thomas et al. (2018). Noteworthy, the work of Jing et al. (2020) proposed the geometric vector perceptron (GVP) consisting of two linear transformations for the input scalar and vector features, followed by nonlinearities. The GVP scalar and vector outputs are invariant and equivariant, respectively, with respect to an arbitrary composition of rotations and reflections in the 3D Euclidean space.

### 2.2. Conformal modeling and the hypersphere neuron

The utility of conformal embedding for Euclidean geometry and the close connection to Minkowski spaces are thoroughly discussed by Li et al. (2001a). An important result is that one can construct hyperspherical decision surfaces using representations in the conformal space (Li et al., 2001b), as done in the work of Perwass et al. (2003). The hypersphere neuron proposed by Baner et al. (2003b) is such a spherical classifier. Remarkably, since a hypersphere can be seen as a generalization of a hyperplane, the standard neuron can be considered as a special case of the hypersphere neuron.

Stacking multiple hypersphere neurons in a feed-forward network results in a multilayer hypersphere perceptron (MLHP), which was shown by Baner et al. (2003a) to outperform the standard MLP for some classification tasks. However, its application to point sets was not discussed. This motivated the work of Melnyk et al. (2021) on the ge-

ometric neuron, where the learned parameters were shown to represent a combination of spherical decision surfaces. Moreover, the geometric (and hypersphere) neuron activations were proved to be isometric in the 3D Euclidean space. In our work, we use the latter observation as the necessary condition for deriving the steerability constraint.

### 2.3. Comparison to other methods

Steerability requires filters that either are (combinations of) spherical harmonics (see, e.g., Fuchs et al. (2020)) or are constructed based on *learned* neurons and that behave as such (our work).

The key point distinguishing our approach from other equivariant networks is that we do not constrain the space of learnable parameters (as opposed to, e.g., Thomas et al. (2018), Weiler et al. (2018a), and Fuchs et al. (2020)), but construct our steerable model from a freely trained base network, as we discuss in detail in Section 4. That is, the related work uses spherical harmonics as atoms, i.e., a hand-designed basis, and learns *only* the linear coefficients under constraints. In contrast, the only thing we inherit from the hand design is the constraint of first-degree harmonics (see Theorem 4.1) – all other degrees of freedom are learnable.

In addition, even though we only consider scalar fields in the present work, the theoretical results of Section 4 can be applied to broader classes of feature fields.

## 3. Background

### 3.1. Steerability

As per Freeman et al. (1991), a 2D function  $f(x, y)$  is said to steer if it can be written as a linear combination of rotated versions of itself, i.e., when it satisfies the constraint

$$f^\theta(x, y) = \sum_{j=1}^M v_j(\theta) f^{\theta_j}(x, y), \quad (1)$$

where  $v_j(\theta)$  are the interpolation functions,  $\theta_j$  are the basis function orientations, and  $M$  is the number of basis function terms required to steer the function. An alternative formulation can be found in the work of Knutsson et al. (1992). In 3D, the steering equation becomes

$$f^{\mathbf{R}}(x, y, z) = \sum_{j=1}^M v_j(\mathbf{R}) f^{\mathbf{R}_j}(x, y, z), \quad (2)$$

where  $f^{\mathbf{R}}(x, y, z)$  is  $f(x, y, z)$  rotated by  $\mathbf{R} \in \text{SO}(3)$ , and each  $\mathbf{R}_j \in \text{SO}(3)$  orients the corresponding  $j$ th basis function.

Theorems 1, 2 and 4 in Freeman et al. (1991) describe the conditions under which the steerability constraints (1) and

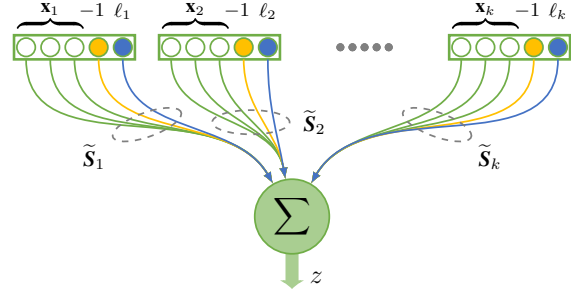


Figure 2. The geometric neuron ( $l_k = -\frac{1}{2}\|\mathbf{x}_k\|^2$ ).

(2) hold, and how to determine the minimum number of basis functions for the 2D and 3D case, respectively.

### 3.2. Conformal embedding

We refer the reader to Section 3 in the work of Melnyk et al. (2021) for more details, and only briefly introduce important notation in this section. The conformal space for the Euclidean  $\mathbb{R}^n$  counterpart can be formed as  $\mathbb{MIE}^n \equiv \mathbb{R}^{n+1,1} = \mathbb{R}^n \oplus \mathbb{R}^{1,1}$ , where  $\mathbb{R}^{1,1}$  is the Minkowski plane (Li et al., 2001a) with orthonormal basis defined as  $\{e_+, e_-\}$  and null basis  $\{e_0, e_\infty\}$  representing the origin  $e_0 = \frac{1}{2}(e_- - e_+)$  and point at infinity  $e_\infty = e_- + e_+$ . Thus, a Euclidean vector  $\mathbf{x} \in \mathbb{R}^n$  can be embedded in the conformal space  $\mathbb{MIE}^n$  as

$$X = \mathcal{C}(\mathbf{x}) = \mathbf{x} + \frac{1}{2}\|\mathbf{x}\|^2 e_\infty + e_0, \quad (3)$$

where  $X \in \mathbb{MIE}^n$  is called *normalized*. The conformal embedding (3) represents the stereographic projection of  $\mathbf{x}$  onto a projection sphere in  $\mathbb{MIE}^n$  and is *homogeneous*, i.e., all embedding vectors in the equivalence class  $[X] = \{Z \in \mathbb{R}^{n+1,1} : Z = \gamma X, \gamma \in \mathbb{R} \setminus \{0\}\}$  are taken to represent the same vector  $\mathbf{x}$ . Importantly, given the conformal embedding  $X$  and some  $Y = \mathbf{y} + \frac{1}{2}\|\mathbf{y}\|^2 e_\infty + e_0$ , their scalar product in the conformal space corresponds to their Euclidean distance,  $X \cdot Y = -\frac{1}{2}\|\mathbf{x} - \mathbf{y}\|^2$ . This interpretation of the scalar product in the conformal space is the main motivation in constructing spherical neurons.

### 3.3. Spherical neurons

By *spherical neurons*, we collectively refer to hypersphere (Banarer et al., 2003a) and geometric (Melnyk et al., 2021) neurons, which have spherical decision surfaces.

As discussed by Banarer et al. (2003a), by embedding both a data vector  $\mathbf{x} \in \mathbb{R}^n$  and a hypersphere  $S \in \mathbb{MIE}^n$  in  $\mathbb{R}^{n+2}$  as

$$\begin{aligned} \mathbf{X} &= (x_1, \dots, x_n, -1, -\frac{1}{2}\|\mathbf{x}\|^2) \in \mathbb{R}^{n+2}, \\ \mathbf{S} &= (c_1, \dots, c_n, \frac{1}{2}(\|\mathbf{c}\|^2 - r^2), 1) \in \mathbb{R}^{n+2}, \end{aligned} \quad (4)$$

where  $\mathbf{c} = (c_1, \dots, c_n) \in \mathbb{R}^n$  is the hypersphere center and  $r \in \mathbb{R}$  is the radius, their scalar product  $X \cdot S$  in the conformal space  $\mathbb{M}\mathbb{E}^n$  can be computed equivalently in  $\mathbb{R}^{n+2}$  as  $\mathbf{X}^\top \mathbf{S}$ :

$$X \cdot S = \mathbf{X}^\top \mathbf{S} = -\frac{1}{2} \|\mathbf{x} - \mathbf{c}\|^2 + \frac{1}{2} r^2. \quad (5)$$

This result enables the implementation of a hypersphere neuron in  $\mathbb{M}\mathbb{E}^n$  using the standard dot product in  $\mathbb{R}^{n+2}$ . The hypersphere vector components are treated as independent learnable parameters during training. Thus, a spherical classifier effectively learns *non-normalized* hyperspheres of the form  $\tilde{\mathbf{S}} = (s_1, \dots, s_{n+2}) \in \mathbb{R}^{n+2}$ . Due to the homogeneity of the representation, both normalized and non-normalized hyperspheres represent the same decision surface. More details can be found in Section 3.2 in the work of Melnyk et al. (2021).

The geometric neuron is a generalization of the hypersphere neuron for point sets as input, see Figure 2. A single geometric neuron output is thus the sum of the signed distances of  $K$  input points to  $K$  learned hyperspheres

$$z = \sum_{k=1}^K \gamma_k \mathbf{X}_k^\top \mathbf{S}_k, \quad (6)$$

where  $z \in \mathbb{R}$ ,  $\mathbf{X}_k \in \mathbb{R}^5$  is a properly embedded 3D input point,  $\gamma_k \in \mathbb{R}$  is the scale factor, i.e., the last element of the learned parameter vector  $\tilde{\mathbf{S}}_k$ , and  $\mathbf{S}_k = \tilde{\mathbf{S}}_k / \gamma_k \in \mathbb{R}^5$  are the corresponding normalized learned parameters (spheres).

Furthermore, Melnyk et al. (2021) demonstrated that the hypersphere (and geometric) neuron activations are isometric in 3D. That is, rotating the input is equivalent to rotating the decision spheres. This result is a necessary condition to consider rotation and translation equivariance of models constructed with spherical neurons and forms the foundation of our methodology.

In the following sections, we use the same notation for a 3D rotation  $\mathbf{R}$  represented in the Euclidean space  $\mathbb{R}^3$ , the homogeneous (projective) space  $P(\mathbb{R}^3)$ , and  $\mathbb{M}\mathbb{E}^3 \cong \mathbb{R}^5$ , depending on the context. This is possible since we can append the required number of ones to the diagonal of the original rotation matrix without changing the transformation representation.

## 4. Method

In this section, we identify the conditions under which a spherical neuron as a function of its 3D input can be steered. In other words, we derive an expression that gives us the response of a *hypothetical* spherical neuron for some input, using rotated versions of the learned spherical neuron parameters. We start by considering the steerability conditions for a single sphere classifying the corresponding input

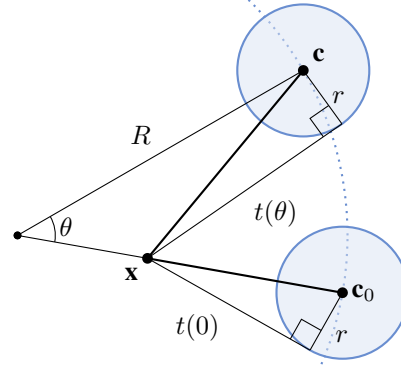


Figure 3. The effect of rotation on the spherical neuron activation in 2D;  $t(\theta)$  denotes the tangent length as a function of rotation angle, i.e.,  $t(0)^2 = \|\mathbf{x} - \mathbf{c}_0\|^2 - r^2 = -2\mathbf{X}^\top \mathbf{S}$ .

point  $\mathbf{X}$ , i.e.,  $f(\mathbf{X}) = \mathbf{X}^\top \mathbf{S}$ , where  $\mathbf{X}$  and  $\mathbf{S}$  are embedded in  $\mathbb{M}\mathbb{E}^3 \cong \mathbb{R}^5$  according to (4).

### 4.1. Basis construction

To formulate a steerability constraint for a spherical neuron that has one sphere as a decision surface, we first need to determine the minimum number of basis functions, i.e., the number of terms  $M$  in (2). This number only depends on the degree of the spherical harmonics that are required to compute the steered result (Freeman et al., 1991). Thus, we need to determine the required degrees.

**Theorem 4.1.** *Let  $\mathbf{S} \in \mathbb{R}^{n+2}$  be an  $nD$  spherical classifier with center  $\mathbf{c}_0 \in \mathbb{R}^n$  ( $R := \|\mathbf{c}_0\|$ ) and radius  $r$ , and  $\mathbf{x} \in \mathbb{R}^n$  be a point represented by  $\mathbf{X} \in \mathbb{R}^{n+2}$ , see (4). Let further  $\mathbf{S}'$  be the classifier that is obtained by rotating  $\mathbf{S}$  in  $nD$  space (i.e., using an element of  $\text{SO}(n)$ ), then  $\mathbf{X}^\top \mathbf{S}'$  and  $\mathbf{X}^\top \mathbf{S}$  are related by spherical harmonics up to first degree.*

*Proof.* Without loss of generality, the rotation is defined by the plane of rotation  $\pi$  and the angle  $\theta$ . Denote the projection of a vector  $\mathbf{v} \in \mathbb{R}^n$  on  $\pi$  by  $\mathbf{v}_\pi$  and define  $\mathbf{v}_{\perp\pi} = \mathbf{v} - \mathbf{v}_\pi$ . From (5) we obtain

$$\begin{aligned} 2\mathbf{X}^\top \mathbf{S} &= r^2 - \|\mathbf{x} - \mathbf{c}_0\|^2 \\ &= r^2 - \|(\mathbf{x} - \mathbf{c}_0)_{\perp\pi}\|^2 - \|(\mathbf{x} - \mathbf{c}_0)_\pi\|^2. \end{aligned}$$

A rotation in  $\pi$  only affects the rightmost term above and there exists a  $\phi \in [0, 2\pi)$  such that

$$\begin{aligned} \|(\mathbf{x} - \mathbf{c}_0)_\pi\|^2 &= \|\mathbf{x}_\pi - \mathbf{c}_{0\pi}\|^2 \\ &= \|\mathbf{x}_\pi\|^2 + \|\mathbf{c}_{0\pi}\|^2 - 2\|\mathbf{x}_\pi\| \|\mathbf{c}_{0\pi}\| \cos \phi. \end{aligned}$$

With a similar argument, we obtain

$$\begin{aligned} 2\mathbf{X}^\top \mathbf{S}' &= r^2 - \|(\mathbf{x} - \mathbf{c}_0)_{\perp\pi}\|^2 \\ &\quad - \|\mathbf{x}_\pi\|^2 - \|\mathbf{c}_{0\pi}\|^2 + 2\|\mathbf{x}_\pi\| \|\mathbf{c}_{0\pi}\| \cos(\phi + \theta). \end{aligned}$$

□

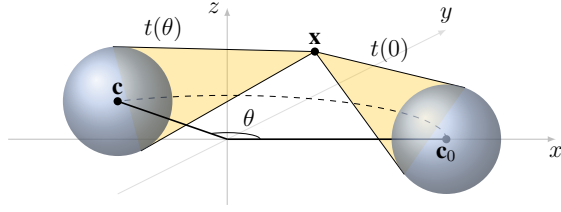


Figure 4. The effect of rotation on the spherical neuron activation in 3D;  $t(\theta)$  denotes the tangent length as a function of rotation angle.

This result is valid in any dimension, but we are primarily interested in  $n = 2$ , as illustrated in Figure 3 for the case  $\phi = 0$ , and  $n = 3$ , shown in Figure 4.

Following the result of Theorem 4 in Freeman et al. (1991) and using  $N = 1$ , we have that  $M = (N + 1)^2 = 4$  basis functions suffice in the 3D case (2).

#### 4.2. Spherical filter banks in 3D

In 3D, we thus select four rotated versions of the function  $f(\mathbf{X})$ , as the basis functions. The rotations  $\{\mathbf{R}_j\}_{j=1}^4$  must be chosen to satisfy the condition (b) in Theorem 4 (Freeman et al., 1991). Therefore, we transform  $f(\mathbf{X})$  such that the resulting four spheres are spaced in three dimensions equally, i.e., form a regular tetrahedron with the vertices  $(1, 1, 1)$ ,  $(1, -1, -1)$ ,  $(-1, 1, -1)$ , and  $(-1, -1, 1)$ , as shown in Figure 5. We stack the homogeneous coordinates of the tetrahedron vertices  $\mathbf{m}_j$  in a matrix column-wise (scaled by  $1/2$ ) to get the orthogonal matrix

$$\mathbf{M} = [\mathbf{m}_1 \quad \mathbf{m}_2 \quad \mathbf{m}_3 \quad \mathbf{m}_4] = \frac{1}{2} \begin{bmatrix} 1 & 1 & -1 & -1 \\ 1 & -1 & 1 & -1 \\ 1 & -1 & -1 & 1 \\ 1 & 1 & 1 & 1 \end{bmatrix}. \quad (7)$$

We will use this matrix operator  $\mathbf{M}$  to compute the linear coefficients in the vector space generated by the vertices of the regular tetrahedron (Granlund & Knutsson, 1995). This will be necessary to find the appropriate interpolation functions and formulate the steerability constraint in Section 4.3.

The four rotated versions of the function  $f(\mathbf{X})$  will constitute the basis functions that we call a spherical filter bank. To construct this filter bank, we choose the following convention. The originally learned spherical classifier  $f(\mathbf{X}) = \mathbf{X}^\top \mathbf{S}$  is first rotated to  $\|\mathbf{c}_0\| \cdot (1, 1, 1)$  (see Figure 5) with the corresponding (geodesic) transformation denoted as  $\mathbf{R}_O$ . Next, we rotate the transformed sphere into the other three vertices of the regular tetrahedron and transform back to the original coordinate system (see Figure 6 for the case of  $(1, -1, -1)$ ).

The resulting filter bank for one spherical classifier is thus

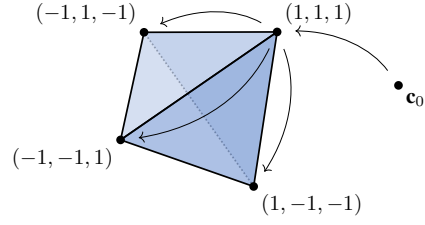


Figure 5. A regular tetrahedron as a basis. Without loss of generality, assume  $\|\mathbf{c}_0\| = \sqrt{3}$ .

composed as the following  $4 \times 5$  matrix:

$$B(\mathbf{S}) = \left[ \mathbf{R}_O^\top \mathbf{R}_{T_i} \mathbf{R}_O \mathbf{S} \right]_{i=0 \dots 3}, \quad (8)$$

where each of  $\{\mathbf{R}_{T_i}\}_{i=0}^3$  is the rotation isomorphism in  $\mathbb{R}^5$  corresponding to a 3D rotation from  $(1, 1, 1)$  to the vertex  $i + 1$  of the regular tetrahedron. Therefore,  $\mathbf{R}_{T_0} = \mathbf{I}_5$ .

As we show by Theorem 4.2, the spherical filter bank  $B(\mathbf{S})$  is equivariant under 3D rotations.

**Theorem 4.2.** *Let  $\mathbf{S} \in \mathbb{R}^5$  be a 3D spherical classifier (i.e., the hypersphere neuron) and  $\mathbf{x} \in \mathbb{R}^3$  be an input point represented by  $\mathbf{X} \in \mathbb{R}^5$ , see (4). Let further  $B(\mathbf{S}) \in \mathbb{R}^{4 \times 5}$  be a filter bank obtained according to (8). Then the filter bank output  $\mathbf{y} = B(\mathbf{S})\mathbf{X} \in \mathbb{R}^4$  is equivariant to 3D rotations of  $\mathbf{x}$ .*

*Proof.* Since rotations in  $\mathbb{R}^3$  are embedded into  $\mathbb{R}^5$  acting on the first three components (for details, see Melnyk et al. (2021)), we need to show that the left sub-matrix  $B_3(\mathbf{S}) \in \mathbb{R}^{4 \times 3}$  of  $B(\mathbf{S})$  is of rank 3. By construction, the four rows of  $B_3(\mathbf{S})$  form the vertices of a tetrahedron, thus span  $\mathbb{R}^3$ .  $\square$

Note that the right sub-matrix  $B_2(\mathbf{S}) \in \mathbb{R}^{4 \times 2}$  of  $B(\mathbf{S})$  results in a constant vector  $B_2(\mathbf{S})[x_4, x_5]^\top \in \mathbb{R}^4$  independent of the rotation, where  $x_4$  and  $x_5$  are the last two components of  $\mathbf{X}$  (see (4)). This additional constant vector requires the use of the fourth vertex of the tetrahedron.

Furthermore, we can explicitly show how the representation  $V_{\mathbf{R}} \in \mathbb{R}^{4 \times 4}$  of the 3D rotation  $\mathbf{R}$  can be obtained in the filter bank  $B(\mathbf{S})$  output space.

**Proposition 4.3.** *Let  $\mathbf{M}$  be the  $4 \times 4$  orthogonal matrix defined in (7),  $\mathbf{R}_O^k \in \mathbb{R}^{4 \times 4}$  be the  $P(\mathbb{R}^3)$  representation (constructed by appending ones to the main diagonal – see the last paragraph in Section 3.3) of the 3D rotation (geodesic) from the center  $\mathbf{c}_0^k$  of the learned  $k$ th spherical classifier  $\mathbf{S}_k$  to  $\|\mathbf{c}_0^k\| \cdot (1, 1, 1)$ , and, slightly abusing notation,  $\mathbf{R} \in \mathbb{R}^{4 \times 4}$  be the  $P(\mathbb{R}^3)$  representation of the 3D rotation  $\mathbf{R}$ .*

*Then  $V_{\mathbf{R}}^k \in \mathbb{R}^{4 \times 4}$  defined below is the representation of the 3D rotation  $\mathbf{R}$  in the filter bank  $B(\mathbf{S}_k)$  output space:*

$$V_{\mathbf{R}}^k = \mathbf{M}^\top \mathbf{R}_O^k \mathbf{R} \mathbf{R}_O^{k \top} \mathbf{M}. \quad (9)$$

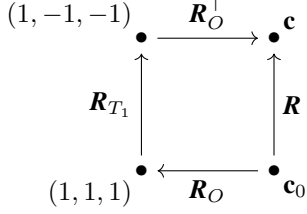


Figure 6. A rotation from  $\mathbf{c}_0$  to  $\mathbf{c}$  described by a tetrahedron rotation. Without loss of generality, assume  $\|\mathbf{c}_0\| = \sqrt{3}$ .

*Proof.* Since  $\det \mathbf{M} = 1$ , we have that  $\mathbf{M} \in \text{SO}(4)$ , and thus, all terms on the right hand-side of (9) are  $4 \times 4$  rotation matrices with only one variable —  $\mathbf{R}$ . Therefore,  $V_{\mathbf{R}}^k$  is also a rotation matrix and is a unique map. The inverse transformation can be computed straightforwardly as

$$\mathbf{R} = \mathbf{R}_O^k \mathbf{M} V_{\mathbf{R}}^k \mathbf{M}^\top \mathbf{R}_O^k \quad (10)$$

with subsequent extraction of the upper-left  $3 \times 3$  sub-matrix as the original 3D rotation matrix  $\mathbf{R}$ . Hence, there is a one-to-one mapping between  $V_{\mathbf{R}}^k$  and  $\mathbf{R}$ .  $\square$

Thus, both  $V_{\mathbf{R}}$  and  $\mathbf{R}$  are representations of the 3D rotation, and we can write the rotation equivariance property of the spherical filter bank  $B(\mathbf{S})$  as

$$V_{\mathbf{R}} B(\mathbf{S}) \mathbf{X} = B(\mathbf{S}) \mathbf{R} \mathbf{X}. \quad (11)$$

### 4.3. 3D steerability constraint

The steerability constraint can be formulated as follows. For an arbitrary rotation  $\mathbf{R}$  applied to the input of the function  $f(\mathbf{X})$ , we want the output of the spherical filter bank  $B(\mathbf{S})$  in (8) to be interpolated with  $v_j(\mathbf{R})$  such that the response is equal to the original function output, i.e.,

$$\begin{aligned} f(\mathbf{X}) &= f^{\mathbf{R}}(\mathbf{R}\mathbf{X}) = \sum_{j=1}^M v_j(\mathbf{R}) f^{\mathbf{R}_j}(\mathbf{R}\mathbf{X}) \\ &= v(\mathbf{R})^\top B(\mathbf{S}) \mathbf{R} \mathbf{X}, \end{aligned} \quad (12)$$

where  $\mathbf{X} \in \mathbb{R}^5$  is a single, appropriately embedded, 3D point, and  $v(\mathbf{R}) \in \mathbb{R}^4$  is a vector of the interpolation coefficients.

The coefficients  $v(\mathbf{R})$  should conform to the basis function construction (condition (b) in Theorem 4 in Freeman et al. (1991)), which is why they are computed with  $\mathbf{M}$  defined in (7). Given  $\mathbf{X} \in \mathbb{R}^5$  as input and an unknown rotation  $\mathbf{R}$  acting on it, the steering equation (12) implies

$$v(\mathbf{R})^\top B(\mathbf{S}) \mathbf{R} \mathbf{X} = v(\mathbf{I})^\top B(\mathbf{S}) \mathbf{X}. \quad (13)$$

Given a tetrahedron rotation, e.g.,  $\mathbf{R}_{T_1}$ , as shown in the diagram in Figure 6, we can define the unknown rotation

accordingly as  $\mathbf{R} = \mathbf{R}_O^\top \mathbf{R}_{T_1} \mathbf{R}_O$ . In this case, it is easy to see that to satisfy the constraint (13),  $v(\mathbf{R})$  must be  $(0, 1, 0, 0)$ , i.e., the second filter in the filter bank  $B(\mathbf{S})$  must be chosen. This can be achieved by transforming a constant vector  $\mathbf{m}_1 = (\frac{1}{2}, \frac{1}{2}, \frac{1}{2}, \frac{1}{2})$  by rotation  $\mathbf{R}_{T_1}$  and multiplying it by the basis matrix  $\mathbf{M}$  as follows:

$$\begin{aligned} v(\mathbf{R}) &= \mathbf{M}^\top (\mathbf{R}_{T_1} \mathbf{m}_1) \\ &= \mathbf{M}^\top \left( \frac{1}{2}, -\frac{1}{2}, -\frac{1}{2}, \frac{1}{2} \right) = (0, 1, 0, 0). \end{aligned} \quad (14)$$

Note that, in general, a geometric neuron (6) takes a set of embedded points as input. Therefore, with the setup above, the rotation of  $\mathbf{m}_1$  will be different for each input shape point  $k$  if the same  $v$  is used for all  $k$ , which contradicts that the shape is transformed by a rigid body motion, i.e., the same  $\mathbf{R}$  for all  $k$ . Thus, we need to consider a suitable vector  $v^k$  for each input shape point  $k$ , such that the resulting  $\mathbf{R}$  is the same for all  $k$ . This can be achieved by recalling how we construct the basis functions in the spherical filter bank (8): we need to consider the respective initial rotation  $\mathbf{R}_O^k$ . The desired interpolation coefficients  $v^k$  are thus computed as

$$v^k(\mathbf{R}) = \mathbf{M}^\top (\mathbf{R}_O^k \mathbf{R} \mathbf{R}_O^{k\top} \mathbf{m}_1). \quad (15)$$

The resulting  $v^k(\mathbf{R}) \in \mathbb{R}^4$  interpolate the responses of the tetrahedron-copies of the originally learned sphere  $S_k$  to replace the rotated sphere.

Remarkably, the interpolation vector  $v^k(\mathbf{R})$  is the first column of  $V_{\mathbf{R}}^k$  (9). Note that  $v^k(\mathbf{R})$  has only three degrees of freedom since the four vertices of the regular tetrahedron sum to zero (see Figure 5). Thus,  $v^k(\mathbf{R})$  is equivariant under 3D rotations.

By plugging (12) into (6), we can now establish the steerability constraint for a geometric neuron,  $f_{\text{GN}}$ , which takes a set of  $K$  embedded points  $\{\mathbf{X}_1, \dots, \mathbf{X}_k\}$  as input:

$$\begin{aligned} f_{\text{GN}}^{\mathbf{R}}(\mathbf{R}\mathbf{X}) &= \sum_{k=1}^K \gamma_k f^{k\mathbf{R}}(\mathbf{R}\mathbf{X}_k) \\ &= \sum_{k=1}^K \gamma_k v^k(\mathbf{R})^\top B(\mathbf{S}_k) \mathbf{R} \mathbf{X}_k. \end{aligned} \quad (16)$$

### 4.4. Steerable model overview

To build a 3D steerable model, we perform the following steps (see Figure 1): We first train an ancestor model, which consists of spherical neurons. After training the model parameters to classify data in a canonical orientation, we freeze them and transform the first layer weights according to the 3D steerability constraint (12) (for hypersphere neurons) or (16) (for geometric neurons). Finally, by combining the resulting parameters in spherical filter banks according to (8) and adding interpolation coefficients (15) as free parameters, we create a steerable model. If the interpolation

coefficients are computed correctly, the model will produce rotation-invariant predictions.

## 5. Experiments

In this section, we describe the experiments we conducted to confirm our findings presented in Section 4.

### 5.1. Datasets

**3D Tetris** Following the experiments reported by Thomas et al. (2018), Weiler et al. (2018a), and Melnyk et al. (2021), we use the following synthetic point set of eight 3D Tetris shapes (Thomas et al., 2018) consisting of four points each (see, e.g., Figure 3 in Melnyk et al. (2021)):

```
chiral_shape_1: [(0, 0, 0), (0, 0, 1), (1, 0, 0), (1, 1, 0)],
chiral_shape_2: [(0, 0, 0), (0, 0, 1), (1, 0, 0), (1, -1, 0)],
square: [(0, 0, 0), (1, 0, 0), (0, 1, 0), (1, 1, 0)],
line: [(0, 0, 0), (0, 0, 1), (0, 0, 2), (0, 0, 3)],
corner: [(0, 0, 0), (0, 0, 1), (0, 1, 0), (1, 0, 0)],
L: [(0, 0, 0), (0, 0, 1), (0, 0, 2), (0, 1, 0)],
T: [(0, 0, 0), (0, 0, 1), (0, 0, 2), (0, 1, 1)],
zigzag: [(0, 0, 0), (1, 0, 0), (1, 1, 0), (2, 1, 0)].
```

**3D skeleton data** We also perform experiments on real-world data to substantiate the validity of our theoretical results. We use the UTKinect-Action3D dataset introduced by Xia et al. (2012), in particular, the 3D skeletal joint locations extracted from Kinect depth maps. For each action sequence and from each frame, we extract the skeleton consisting of twenty points and assign it to the class of actions (ten categories) this frame is part of. Therefore, we formulate the task as shape recognition (i.e., a kind of action recognition from a static point cloud), where each shape is of size  $20 \times 3$ .

Since the orientations of the shapes vary significantly across the sequences, we perform the following standardization: We first center each shape at the origin, and then compute the orientation from its three hip joints and de-rotate the shape in the  $xy$ -plane (viewer coordinate system) accordingly. We illustrate the effect of de-rotation in Figure 7. From each action sequence, we randomly select 50% of the skeletons for the test set and 20% of the remainder as validation data. The resulting data split is as follows: 2295 training shapes, 670 shapes for validation, and 3062 test shapes, approximately corresponding to 38%, 11%, 51% of the total amount of the skeletons, respectively.

### 5.2. Steerable model construction

To construct and test steerable models, we perform the same steps for both datasets. The minor differences are the choice of training hyperparameters and the presence of validation and test subsets in the 3D skeleton dataset.

We first train a two-layer (*ancestor*) multilayer geometric perceptron (MLGP) model (Melnyk et al., 2021), where the

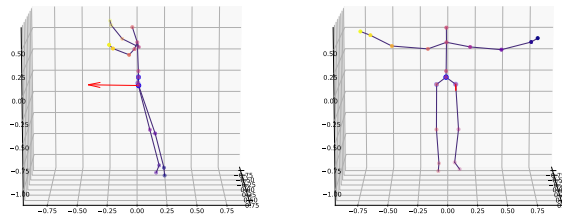


Figure 7. The effect of standardizing the orientation of the 3D skeleton representing *waveHands* action: the original (left) and the de-rotated (right) shape; the arrow is the normal vector of the plane formed by the three hip joints.

first layer consists of geometric neurons and the output layer of hypersphere neurons, to classify the shapes. Since the architecture choice is not the objective of the experiments, when building the MLGP, we use only one configuration with five hidden units for the Tetris data and twelve hidden units (determined by using the validation data) for the 3D skeleton dataset throughout the experiments.

Similar to Melnyk et al. (2021), we do not use any activation function in the first layer due to the nonlinearity of the conformal embedding. We implement both MLGP models in PyTorch (Paszke et al., 2019) and keep the default parameter initialization for the linear layers. We train both models by minimizing the cross-entropy loss function and use the Adam optimizer (Kingma & Ba, 2015) with the default hyperparameters (the learning rate is 0.001). The Tetris MLGP learns to classify the eight shapes in the canonical orientation perfectly after 2000 epochs, whereas the Skeleton model trained for 10000 epochs achieves a test set accuracy of 92.9%.

For both, we then freeze the trained parameters and construct a steerable model. Note that we form steerable units only in the first layer and keep the output, i.e., classification, layer hypersphere neurons as they are. The steerability is not required for the subsequent layers as the output of the first layer becomes rotation-independent. The steerable units are formed from the corresponding frozen parameters as the (fixed) filter banks according to (8).

The only free parameters of this constructed steerable model are interpolation coefficients  $v^k(\mathbf{R}) \in \mathbb{R}^4$  defined in (15), where  $k$  indexes the learned first layer parameters (spheres) in the ancestor MLGP model.

Note that the interpolation coefficients (15) are all parameterized by the input orientation  $\mathbf{R}$ . If the interpolation coefficients are computed correctly, the model output will be rotation-invariant. However, in this paper, we focus on the steerability as such, and a practical way of determining  $v^k(\mathbf{R})$  remains to be done in future work, including the experimental comparison with rotation-invariant classifiers from related work. Therefore, the following experiment

Table 1. The steerable model classification accuracy for the distorted (the noise units are specified in the square brackets) rotated shapes and the ancestor accuracy for the distorted shapes in their canonical orientation (mean and std over 1000 runs, %).

3D Tetris			3D skeleton data ( <i>test set</i> )		
Noise ( $a$ ), [1]	<i>Steerable</i>	<i>Ancestor</i>	Noise ( $a$ ), [m]	<i>Steerable</i>	<i>Ancestor</i>
0.00	100.0 ± 0.0	100.0 ± 0.0	0.000	92.9 ± 0.0	92.9 ± 0.0
0.05	100.0 ± 0.0	100.0 ± 0.0	0.005	92.4 ± 0.2	92.4 ± 0.2
0.10	100.0 ± 0.0	100.0 ± 0.0	0.010	91.1 ± 0.3	91.1 ± 0.3
0.20	100.0 ± 0.4	100.0 ± 0.0	0.020	87.1 ± 0.5	87.1 ± 0.5
0.30	99.7 ± 1.9	99.8 ± 1.6	0.030	82.3 ± 0.6	82.2 ± 0.6
0.50	94.9 ± 7.7	95.0 ± 7.9	0.050	72.0 ± 0.7	71.9 ± 0.7

Table 2. The L1 distance between the steerable model hidden activations and the ground truth activations given the distorted (the noise units are specified in the square brackets) rotated shapes and between the ancestor hidden activations given the distorted shapes in their canonical orientation (mean and std over 1000 runs).

3D Tetris			3D skeleton data ( <i>test set</i> )		
Noise ( $a$ ), [1]	<i>Steerable</i>	<i>Ancestor</i>	Noise ( $a$ ), [m]	<i>Steerable</i>	<i>Ancestor</i>
0.00	0.00 ± 0.00	0.00 ± 0.00	0.000	0.00 ± 0.00	0.00 ± 0.00
0.05	0.33 ± 0.05	0.33 ± 0.05	0.005	0.53 ± 0.00	0.53 ± 0.00
0.10	0.66 ± 0.10	0.66 ± 0.10	0.010	1.06 ± 0.01	1.06 ± 0.01
0.20	1.32 ± 0.19	1.32 ± 0.19	0.020	2.12 ± 0.01	2.12 ± 0.01
0.30	2.00 ± 0.31	2.00 ± 0.29	0.030	3.18 ± 0.01	3.18 ± 0.01
0.50	3.33 ± 0.48	3.32 ± 0.48	0.050	5.30 ± 0.02	5.30 ± 0.02

serves as an empirical validation of the derived constraint rather than rotation-invariant predictions.

### 5.3. Known rotation experiment

Using the trained ancestor MLGP, we verify the correctness of (16) (and, therefore, (12)). We first rotate the original data and then use this ground truth rotation to compute the interpolation coefficients of the constructed steerable model according to (15). Our intuition is that if the steerability constraint (16) is correct, then, given the transformed point set, the activations of the steerable units in the steerable model will be equal to the activations of the geometric neurons in the ancestor MLGP model fed with the point set in the canonical orientation. Hence, the classification accuracies of the ancestor and steerable models on the original and transformed datasets, respectively, should be equal.

We run this experiment 1000 times. Each time, we generate a random rotation and apply it to the original point set (in case of the 3D skeleton data, we use the *test* split). We use this ground-truth rotation information to compute the interpolation coefficients for the steerable model, which we then evaluate on the transformed point set.

To verify the stability of the steerable unit activations, we add uniform noise to the transformed points,  $\mathbf{n} \sim U(-a, a)$ , where the range of  $a$  is motivated by the magnitude of the points in the datasets: for the Tetris data, the highest  $a$

is chosen to be 0.5, making it impossible for the noise to perturb a shape of one category into a shape of another class; for the skeleton data, the highest  $a$  is set to 0.05 m – a reasonable amount of distortion for a human of average size (see, e.g., Figure 7), which should be insufficient to completely change the representation of the skeleton action.

For reference, we also present the accuracy of the ancestor model classifying the data in the canonical orientation. We summarize the results in Table 1. Additionally, by computing the L1 distance, we compare the hidden unit activations of the ancestor MLGP fed with the shapes in the canonical orientation (*ground truth* activations) and those of the constructed steerable models fed with the rotated data (see Table 2).

## 6. Discussion and Conclusion

Enabled by the complete understanding of the geometry of the spherical neurons, we show in Section 4 that we only need the spherical harmonics of degree up to  $N = 1$  to determine the effect of rotation on the activations in 3D. Using this result, we derive a novel 3D steerability constraint (12) (and (16)), adding to the practical value of prior work (Banarer et al., 2003b; Perwass et al., 2003; Melnyk et al., 2021).

The experiment conducted in Section 5.3 shows that the derived constraint is correct since the constructed steerable



model produces equally (up to a numerical precision) accurate predictions for the rotated shapes as the ancestor for the shapes in the canonical orientation, provided that the interpolation coefficients are computed with the *known* rotation. From Table 1 and 2, we can see that the steerable model classification error and the L1 distance to the ground truth activations only moderately increase with the level of noise in the input data, which is a clear indication of the robustness of the classifier.

Noteworthy, invariance to point permutations in the input is not attained with the ancestor MLGP model that we used in the experiments. However, one can address the problem of permutation invariance quite straightforwardly: For example, one could follow the construction of, e.g., PointNet (Qi et al., 2017), and replace the shared MLPs applied to each point in isolation with shared spherical neurons taking one point as input (i.e., hypersphere neurons), and then apply a global aggregation function to the extracted features. One would then apply the steerability constraint (12) to the first layer learned parameters to create a steerable model.

In general, one can build an ancestor model containing spherical neurons only in the first layer, with subsequent layers of various flavors (see Figure 1). As we mentioned in Section 5.2, steerability is not required for the subsequent layers as the output of the first layer ( $\mathbf{h}$  in Figure 1) becomes rotation-independent, given that the interpolation coefficients are computed correctly.

The interpolation coefficients  $v^k$  could be learned directly, e.g., by using a regression network. However, this would be a much more complex learning problem than determining  $v^k$  indirectly by regressing  $\mathbf{R}$ , since all  $K$  interpolation vectors  $v^k(\mathbf{R})$  are by definition uniquely parameterized by  $\mathbf{R}$ .

Rotation regression with NNs is per se an important and continuously studied problem that is often considered in the context of pose estimation (Zhou et al., 2019; Chen et al., 2022). Therefore, it is more efficient and intuitive to combine a rotation regressor with our steerable approach for creating a rotation-invariant classifier as compared to the approaches where the rotation information is naively encompassed in the learned parameters of a generic network by training on rotation-augmented data.

Finally, the main theoretical results of this paper provide a rigorous and geometrically lucid mechanism for constructing steerable feature extractors, the practical utility of which can be fully unveiled in future work.

## Acknowledgments

This work was supported by the Wallenberg AI, Autonomous Systems and Software Program (WASP), by the Swedish Research Council through a grant for the project AI-

gebraically Constrained Convolutional Networks for Sparse Image Data (2018-04673), and the strategic research environment ELLIIT.

## References

- Anderson, B., Hy, T. S., and Kondor, R. Cormorant: Covariant molecular neural networks. In *Advances in Neural Information Processing Systems*, pp. 14510–14519, 2019.
- Banarer, V., Perwass, C., and Sommer, G. Design of a multilayered feed-forward neural network using hypersphere neurons. In *International Conference on Computer Analysis of Images and Patterns*, pp. 571–578. Springer, 2003a.
- Banarer, V., Perwass, C., and Sommer, G. The hypersphere neuron. In *ESANN*, pp. 469–474, 2003b.
- Chen, J., Yin, Y., Birdal, T., Chen, B., Guibas, L. J., and Wang, H. Projective manifold gradient layer for deep rotation regression. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 6646–6655, June 2022.
- Cohen, T. S. and Welling, M. Steerable CNNs. *arXiv preprint arXiv:1612.08498*, 2016. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2017.
- Freeman, W. T., Adelson, E. H., et al. The design and use of steerable filters. *IEEE Transactions on Pattern analysis and machine intelligence*, 13(9):891–906, 1991.
- Fuchs, F., Worrall, D., Fischer, V., and Welling, M. Se(3)-transformers: 3d roto-translation equivariant attention networks. In Larochelle, H., Ranzato, M., Hadsell, R., Balcan, M. F., and Lin, H. (eds.), *Advances in Neural Information Processing Systems*, volume 33, pp. 1970–1981. Curran Associates, Inc., 2020.
- Granlund, G. and Knutsson, H. (eds.). *Signal Processing for Computer Vision*. Kluwer, Dordrecht, 1995.
- Jing, B., Eismann, S., Suriana, P., Townshend, R. J., and Dror, R. Learning from protein structure with geometric vector perceptrons. *arXiv preprint arXiv:2009.01411*, 2020. Published as a conference paper at ICLR 2021.
- Kingma, D. P. and Ba, J. Adam: A method for stochastic optimization. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2015.
- Knutsson, H., Haglund, L., Bårman, H., and Granlund, G. H. A framework for anisotropic adaptive filtering and analysis of image sequences and volumes. In *Proceedings ICASSP-92: 1992 IEEE International Conference on Acoustics, Speech, and Signal Processing*, volume 3, pp. 469–472 vol.3, 1992. doi: 10.1109/ICASSP.1992.226174.

- Kondor, R., Lin, Z., and Trivedi, S. Clebsch–Gordan Nets: a Fully Fourier Space Spherical Convolutional Neural Network. In Bengio, S., Wallach, H., Larochelle, H., Grauman, K., Cesa-Bianchi, N., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems*, volume 31. Curran Associates, Inc., 2018.
- Li, H., Hestenes, D., and Rockwood, A. Generalized homogeneous coordinates for computational geometry. In *Geometric Computing with Clifford Algebras*, pp. 27–59. Springer, 2001a.
- Li, H., Hestenes, D., and Rockwood, A. A universal model for conformal geometries of Euclidean, spherical and double-hyperbolic spaces. In *Geometric computing with Clifford algebras*, pp. 77–104. Springer, 2001b.
- Marcos, D., Volpi, M., Komodakis, N., and Tuia, D. Rotation Equivariant Vector Field Networks. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, Oct 2017.
- Melnyk, P., Felsberg, M., and Wadenbäck, M. **Embed Me if You Can: A Geometric Perceptron**. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 1276–1284, October 2021.
- Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., et al. Pytorch: An imperative style, high-performance deep learning library. In *Advances in Neural Information Processing Systems*, pp. 8024–8035, 2019.
- Perona, P. Deformable kernels for early vision. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 17(5):488–499, 1995. doi: 10.1109/34.391394.
- Perwass, C., Banarer, V., and Sommer, G. Spherical decision surfaces using conformal modelling. In *Joint Pattern Recognition Symposium*, pp. 9–16. Springer, 2003.
- Qi, C. R., Su, H., Mo, K., and Guibas, L. J. Pointnet: Deep learning on point sets for 3d classification and segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 652–660, 2017.
- Reisert, M. *Group integration techniques in pattern analysis - a kernel view*. PhD thesis, University of Freiburg, Freiburg im Breisgau, Germany, 2008.
- Simoncelli, E. P. and Freeman, W. T. The steerable pyramid: a flexible architecture for multi-scale derivative computation. In *Proceedings., International Conference on Image Processing*, volume 3, pp. 444–447 vol.3, 1995. doi: 10.1109/ICIP.1995.537667.
- Simoncelli, E. P., Freeman, W. T., Adelson, E. H., and Heeger, D. J. Shiftable multiscale transforms. *IEEE Transactions on Information Theory*, 38(2):587–607, 1992. doi: 10.1109/18.119725.
- Teo, P. C. and Hel-Or, Y. Lie generators for computing steerable functions. *Pattern Recognition Letters*, 19(1): 7–17, 1998. ISSN 0167-8655. doi: [https://doi.org/10.1016/S0167-8655\(97\)00156-6](https://doi.org/10.1016/S0167-8655(97)00156-6).
- Thomas, N., Smidt, T., Kearnes, S., Yang, L., Li, L., Kohlhoff, K., and Riley, P. Tensor field networks: Rotation-and translation-equivariant neural networks for 3D point clouds. *arXiv preprint arXiv:1802.08219*, 2018.
- Van Gool, L., Moons, T., Pauwels, E., and Oosterlinck, A. Vision and Lie’s approach to invariance. *Image and Vision Computing*, 13(4):259–277, 1995. ISSN 0262-8856. doi: [https://doi.org/10.1016/0262-8856\(95\)99715-D](https://doi.org/10.1016/0262-8856(95)99715-D).
- Weiler, M. and Cesa, G. General  $E(2)$ -Equivariant Steerable CNNs. *Advances in Neural Information Processing Systems*, 32, 2019.
- Weiler, M., Geiger, M., Welling, M., Boomsma, W., and Cohen, T. S. 3D steerable CNNs: Learning rotationally equivariant features in volumetric data. In *Advances in Neural Information Processing Systems*, pp. 10381–10392, 2018a.
- Weiler, M., Hamprecht, F. A., and Storath, M. Learning steerable filters for rotation equivariant CNNs. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 849–858, 2018b.
- Xia, L., Chen, C.-C., and Aggarwal, J. K. View invariant human action recognition using histograms of 3d joints. In *2012 IEEE Computer Society Conference on Computer Vision and Pattern Recognition Workshops*, pp. 20–27, 2012. doi: 10.1109/CVPRW.2012.6239233.
- Zhao, Y., Birdal, T., Lenssen, J. E., Menegatti, E., Guibas, L., and Tombari, F. Quaternion equivariant capsule networks for 3d point clouds. In *European Conference on Computer Vision*, pp. 1–19. Springer, 2020.
- Zhou, Y., Barnes, C., Lu, J., Yang, J., and Li, H. On the continuity of rotation representations in neural networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 5745–5753, 2019.