



HAL
open science

6DCNN with roto-translational convolution filters for volumetric data processing

Dmitrii Zhemchuzhnikov, Ilia Igashov, Sergei Grudinin

► **To cite this version:**

Dmitrii Zhemchuzhnikov, Ilia Igashov, Sergei Grudinin. 6DCNN with roto-translational convolution filters for volumetric data processing. 2021. hal-03298699v1

HAL Id: hal-03298699

<https://hal.science/hal-03298699v1>

Preprint submitted on 26 Jul 2021 (v1), last revised 28 Jul 2021 (v2)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

6DCNN WITH ROTO-TRANSLATIONAL CONVOLUTION FILTERS FOR VOLUMETRIC DATA PROCESSING

Dmitrii Zhemchuzhnikov, Ilia Igashov, and Sergei Grudin
Univ. Grenoble Alpes, CNRS, Grenoble INP, LJK
38000 Grenoble, France
sergei.grudin@univ-grenoble-alpes.fr

ABSTRACT

In this work, we introduce 6D Convolutional Neural Network (6DCNN) designed to tackle the problem of detecting relative positions and orientations of local patterns when processing three-dimensional volumetric data. 6DCNN also includes SE(3)-equivariant message-passing and nonlinear activation operations constructed in the Fourier space. Working in the Fourier space allows significantly reducing the computational complexity of our operations. We demonstrate the properties of the 6D convolution and its efficiency in the recognition of spatial patterns. We also assess the 6DCNN model on several datasets from the recent CASP protein structure prediction challenges. Here, 6DCNN improves over the baseline architecture and also outperforms the state of the art.

Keywords Equivariant convolutional networks · Roto-translational convolutions · Local equivariance · 6DCNN · Protein structure prediction

1 Introduction

Methods of deep learning have recently made a great leap forward in the spatial data processing. This domain contains various tasks from different areas of industry and natural sciences, including three-dimensional (3D) data analysis. For a long time, convolutional neural networks (CNNs) remained the main tool in this domain. CNNs helped to solve many real-world challenges, especially in computer vision. However, these architectures have rather strict application restrictions. Unfortunately, real-world raw data rarely have standard orientation and size, which limits the efficiency of translational convolutions. This circumstance has created an increased interest in the topic of SE(3)-equivariant operations in recent years.

Moreover, volumetric data may contain arbitrarily-positioned and -oriented sub-patterns. Their detection makes 3D pattern recognition particularly challenging. Indeed, a volumetric sub-pattern in 3D has six degrees of freedom (DOFs), three to define a rotation, and three for a translation. Thus, a classical convolution technique would require scanning through all these six DOFs and scale as $O(N^3M^6)$ if a brute-force computation is used, where N is the linear size of the volumetric data, and M is the linear size of the sub-pattern.

This work proposes a set of novel operations with the corresponding architecture based on six-dimensional (6D) *roto-translational* convolutional filters. For the first time, thanks to the polynomial expansions in the Fourier space, we demonstrate the feasibility of the 6D roto-translational-based convolutional network with the leading complexity of $O(N^2M^4)$ operations. We tested our method on simulated data and also on protein structure prediction datasets, where the overall accuracy of our predictions is on par with the state-of-the-art methods. Proteins play a crucial role in most biological processes. Despite their seeming complexity, structures of proteins attract more and more attention from the data science community (Senior et al., 2020; Jumper et al., 2021; Laine et al., 2021). In particular, the task of protein structure prediction and analysis raises the challenge of constructing rotational and translational equivariant architectures.

2 State of the art / Related work

Equivariant operations. The first attempt of learning rotation-equivariant representations was made in Harmonic Networks (Worrall et al., 2017) in application to 2D images. Further, this idea was transferred to the 3D space with the corresponding architecture known as 3D Steerable CNNs (Weiler et al., 2018). In Spherical CNNs (Cohen et al., 2018b), the authors introduced a correlation on the rotation group and proposed a concept of rotation-equivariant CNNs on a sphere. Spherical harmonics kernels that provide rotational invariance have also been applied to point-cloud data (Poulenard et al., 2019). A further effort on leveraging compact group representations resulted in the range of methods based on Clebsh-Gordan coefficients (Kondor, 2018; Kondor et al., 2018; Anderson et al., 2019). This approach was finally generalized in Tensor field networks (Thomas et al., 2018), where rotation-equivariant operations were applied to vector and tensor fields. Later, SE(3)-Transformers (Fuchs et al., 2020) were proposed to efficiently capture the distant spatial relationships. More recently, Hutchinson et al. (2021) continued to develop the theory of equivariant convolution operations for homogeneous spaces and proposed Lie group equivariant transformers, following works on the general theory of group equivariant operations in SO(2) (Romero et al., 2020; Romero and Cordonnier, 2021) and SO(3) (Cohen et al., 2018a). We should mention that the most common data representation in this domain is a 3D point cloud, however, several approaches operate on regular 3D grids (Weiler et al., 2018; Pagès et al., 2019). We should also add that some of the above-mentioned methods (Cohen et al., 2018b; Weiler et al., 2018; Kondor, 2018; Anderson et al., 2019) employ Fourier transform in order to learn rotation-equivariant representations.

Geometric learning on molecules. As graphs and point clouds are natural structures for representing molecules, it is reasonable that geometric learning methods have been actively evolving especially in application to biology, chemistry, and physics. More classical graph-learning methods for molecules include MPNNs (Gilmer et al., 2017), SchNet (Schütt et al., 2017), and MEGNet (Chen et al., 2019). Ideas for efficient capturing of spatial relations in molecules resulted in rotation-invariant message-passing methods DimeNet (Klicpera et al., 2020b) and DimeNet++ (Klicpera et al., 2020a). Extending message-passing mechanism with rotationally equivariant representations, polarizable atom interaction neural networks (Schütt et al., 2021) managed to efficiently predict tensorial properties of molecules. Satorras et al. (2021b) proposed E(n) equivariant GNNs for predicting molecular properties and later used them for developing generative models equivariant to Euclidean symmetries (Satorras et al., 2021a).

Proteins are much bigger and more complex systems than small molecules but are composed of repeating blocks. Therefore, more efficient and slightly different methods are required to operate on them. A very good example is the two recent and very powerful methods AlphaFold2 (Jumper et al., 2021) and RoseTTAFold (Baek et al., 2021). Most recent geometric learning methods designed for proteins include deep convolutional networks processing either volumetric data in local coordinate frames (Pagès et al., 2019; Hiranuma et al., 2021), graph neural networks (Ingraham et al., 2019; Sanyal et al., 2020; Baldassarre et al., 2021; Igashov et al., 2021a,b), deep learning methods on surfaces and point clouds (Gainza et al., 2020; Sverrisson et al., 2020), and geometric vector perceptrons (Jing et al., 2021b,a). In addition, several attempts were made to scale tensor-field-like SE(3)-equivariant methods to proteins (Derevyanko and Lamoureux, 2019; Eismann et al., 2020; Townshend et al., 2020; Baek et al., 2021).

3 Model/Method

3.1 Workflow

Here, we give a brief description of all steps of our method that are described in more detail below. Firstly, for each residue in the input protein molecule, we construct a function $\mathbf{f}(\vec{r})$ that describes its local environment. More technically, this function is a set of 3D Gaussian-shaped features centered on the location of atoms within a certain distance R_{max} from the C_α atom of the corresponding residue (see Fig. 1A).

Then, for each function $\mathbf{f}(\vec{r})$, we compute its spherical Fourier expansion coefficients $\mathbf{F}_l^k(\rho)$. The angular resolution of the expansion is determined by the maximum order of spherical harmonics L . The radial resolution of the expansion corresponds to the maximum reciprocal distance ρ_{max} and is inversely proportional to the resolution σ of the real-space Gaussian features as $\rho_{max} = \pi/\sigma$ (see Fig. 1B). Similarly, the radial spacing between the reciprocal points is inversely proportional to the linear size of the data, $\Delta\rho = \pi/(2R_{max})$. Without loss of generality, we can set the number of reciprocal radial points to be equal L , such that $\rho_{max}/\Delta\rho = L + 1 = 2R_{max}/\sigma$.

Spherical Fourier coefficients $\mathbf{F}_l^k(\rho)$ constitute the input for our network, along with the information about the transition from the coordinate system of one residue to another. We start the network with the embedding block that reduces the dimensionality of the feature space. Then, we apply a series of 6D convolution blocks that consist of 6D convolution, normalization, and activation layers, followed by a message-passing layer. After a series of operations on continuous data, we switch to the discrete representation and continue the network with graph convolutional layers (see Fig. 1C-D).

In the graph representation, each node corresponds to a protein residue, and a graph edge links nodes if the distance between the corresponding C_α atoms is smaller than a certain threshold R_n . We should also mention that the backbone structure of a protein residue can be used to unambiguously define its local coordinate system (Pagès et al., 2019; Jumper et al., 2021) using the Gram–Schmidt orthogonalization process starting from $C_\alpha - N$ and $C_\alpha - C$ vectors.




Figure 1: **A.** Six-dimensional (6D) convolution between a filter $w(\vec{r})$ and a function $f(\vec{r}_0 + \vec{r})$. The function $f(\vec{r}_0 + \vec{r})$ describes the local environment of a protein residue and is defined within a certain radius R_{max} from the corresponding C_α atom. The local coordinate system xyz is built on the backbone atoms C_α , C , N of each protein residue. R denotes the location of a residue’s side-chain. **B.** The spherical Fourier space with the reciprocal spacing $\Delta\rho$ and the maximum resolution of ρ_{max} . Grey dots schematically illustrate points where the Fourier image is stored. **C.** An illustration of a protein chain representation. Each protein residue has its own coordinate system xyz and the corresponding local volumetric description $F_l^k(\rho)$ within a certain sphere of R_{max} radius. Spheres of different residues may overlap. Two residues are considered as neighbors in the graph representation if their C_α atoms are located within a certain threshold R_n . **D.** The graph representation of the protein structure. The node features are learned by the network and are represented with colored rectangles. The edge features are assigned based on the types of the corresponding residues and the topological distance of the protein graph.

3.2 Representation of volumetric data

Let us consider a function $f(\vec{r}) : \mathcal{R}^3 \rightarrow \mathcal{R}^{d_f}$ that describes a distribution of d_f -dimensional features in the 3D space. Very often, initial data is given as a point cloud, as it is typically the case for protein structures. Let us consider a set of points located within a maximum radius R_{max} at positions $\vec{r}_1, \dots, \vec{r}_n, \dots, \vec{r}_N$ with the corresponding feature vectors $\mathbf{t}_1, \dots, \mathbf{t}_n, \dots, \mathbf{t}_N$. To convert this representation to a continuous one, we assume that each point feature has a Gaussian shape with a standard deviation σ . Then, the continuous function characterizing this area will have the form,

$$\mathbf{f}(\vec{r}) = \sum_{n=1}^N \mathbf{f}_n(\vec{r}) = \sum_{n=1}^N \mathbf{t}_n \exp\left(-\frac{(\vec{r} - \vec{r}_n)^2}{2\sigma^2}\right), \quad (1)$$

where σ is the *spatial resolution* of the features. It is very convenient to use the Fourier description of this function in spherical coordinates. The spherical harmonics expansion of the Fourier transform of function $f(\vec{r})$ will be

$$\mathbf{F}_l^k(\rho) = 4\pi(-i)^l \sum_{n=1}^N j_l(\rho r_n) \overline{Y_l^k(\Omega_{r_n})} (\sqrt{2\pi}\sigma)^3 \exp\left(-\frac{\sigma^2 \rho^2}{2}\right) \mathbf{t}_n, \quad (2)$$

where ρ is the *reciprocal distance*. Please see Appendices A-E for more details.

3.3 6D convolution operation

The initial idea that prompted us to consider non-traditional types of convolution is the intention to perceive spatial patterns in whatever orientations they have. In a classical 3D convolution, when a filter is applied to the image, the result is sensitive to the orientation of the filter. Keeping this in mind, we came with the idea to extend the convolution with an integration of all possible filter rotations. Let $f(\vec{r}) : \mathcal{R}^3 \rightarrow \mathcal{R}^{d_i}$ and $w(\vec{r}) : \mathcal{R}^3 \rightarrow \mathcal{R}^{d_o}$ be the initial signal and a spatial filter, correspondingly. We propose to extend the classical convolution as follows,

$$\int_{\vec{r}} d\vec{r} f(\vec{r}_0 + \vec{r}) w(\vec{r}) \rightarrow \int_{\Lambda} d\Lambda \int_{\vec{r}} d\vec{r} f(\vec{r}_0 + \Lambda^{-1}\vec{r}) w(\Lambda\vec{r}), \quad (3)$$

where $\Lambda \in \text{SO}(3)$ is a 3D rotation. Please see Fig. 1A for an illustration. Let the functions $f(\vec{r})$ and $w(\vec{r})$ be *finite-resolution* and have spherical Fourier expansion coefficients $F_l^k(\rho)$ and $W_l^k(\rho)$, correspondingly, which are

nonzero for $l \leq$ than some maximum expansion coefficient L . Then, the result of the 6D convolution has the following coefficients,

$$[\mathbf{F}_{\text{out}}]_l^k(\rho) = \sum_{l_1=0}^L \sum_{k_1=-l_1}^{l_1} \frac{8\pi^2}{2l_1+1} \mathbf{W}_{l_1}^{-k_1}(\rho) \sum_{l_2=|l-l_1|}^{l+l_1} c^{l_2}(l, k, l_1, -k_1) \mathbf{F}_{l_2}^{k+k_1}(\rho), \quad (4)$$

where c^l are the products of three spherical harmonics, see Eq. 27. The proof can be found in Appendix F. For a single reciprocal distance ρ , the complexity of this operation is $O(L^5)$, where L is the maximum order of the spherical harmonics expansion.

We should specifically draw the reader's attention to the fact of opposite rotations Λ and Λ^{-1} in the convolution equation 3. Indeed, if we leave only one rotation on $\mathbf{f}(\vec{r})$ or $\mathbf{w}(\vec{r})$, after the integration over rotations in $\text{SO}(3)$, we will obtain the resultant convolution of the input function that will depend only on the radial components of the filter. Alternatively, if we apply the same rotation on both $\mathbf{f}(\vec{r})$ and $\mathbf{w}(\vec{r})$, it will be integrated out, and we will obtain a common 3D translational convolution as a result.

3.4 Nonparametric message passing of continuous data

Let us assume that our spatial data can be represented with overlapping spatial fragments, each having its own local coordinate frame, as it is shown in Fig. 1C for the continuous representation of a protein molecule. Then, we can recompute the fragment representation in the neighboring coordinate frames using spatial transformation operators. For this purpose, we designed a message passing operation in a form that recomputes the spherical Fourier coefficients in a new reference frame. We should specifically note that this operation is $\text{SE}(3)$ -equivariant by construction, because the spatial relationship between molecular fragments remain the same when rotating and shifting the global coordinate system. We decompose such a spatial transformation in a sequence of a rotation followed by a z -axis translation followed by the second rotation. Indeed, the spherical Fourier basis provides low computational complexity for the z -translation and rotation operations. Let function $\mathbf{f}_z(\vec{r})$ be the results of translating function $\mathbf{f}(\vec{r})$ along the z -axis by an amount Δ . Then, the expansion coefficients of these two functions will have the following relation,

$$[\mathbf{F}_z]_l^k(\rho) = \sum_{l'=k}^L T_{l,l'}^k(\rho, \Delta) \mathbf{F}_{l'}^k(\rho) + O\left(\frac{1}{(L-l)!} \left(\frac{\rho\Delta}{2}\right)^{L-l}\right), \quad (5)$$

where $T_{l,l'}^k$ is a translation tensor specified in Appendix D. The update of all the expansion coefficients costs $O(L^4)$ operations.

Similarly, let function $\mathbf{f}_\Lambda(\vec{r})$ be the rotation of function $\mathbf{f}(\vec{r})$ by an amount $\Lambda \in \text{SO}(3)$, $\mathbf{f}_\Lambda(\vec{r}) = \mathbf{f}(\Lambda\vec{r})$. The corresponding expansion coefficients are then related as

$$[\mathbf{F}_\Lambda]_l^k(\rho) = \sum_{k'=-l}^l D_{k',k}^l(\Lambda) \mathbf{F}_l^{k'}(\rho), \quad (6)$$

where $D_{k',k}^l$ is a rotation Wigner matrix specified in Appendix C. The update of all the expansion coefficients will again cost $O(L^4)$ operations.

3.5 Normalization

Since we are working with continuous data in the Fourier space, we have to introduce our own activation and normalization functions. When developing the normalization, we proceeded from very basic premises. More precisely, we normalize the signal $\mathbf{f}(\vec{r})$ by setting its mean to zero and its variance to unity. This can be achieved if the following operations are performed on the spherical Fourier expansion coefficients of the initial function,

$$[\mathbf{F}_n]_l^k(\rho) = \begin{cases} 0, & \text{if } l = k = \rho = 0, \\ [\mathbf{F}]_l^k(\rho) / \mathbf{S}_2, & \text{otherwise,} \end{cases} \quad (7)$$

where $\mathbf{S}_1 = \int_{\mathcal{R}^3} \mathbf{f}(\vec{r}) d\vec{r}$, $\mathbf{S}_2 = \int_{\mathcal{R}^3} (\mathbf{f}(\vec{r}) - \mathbf{S}_1)^2 d\vec{r}$. We should also notice that we apply the element-wise division in Eq. 7. The proof can be found in Appendix G.

3.6 Activation

The concept of our activation operation coincides with the idea of the classical activation in neural networks, i.e. to nonlinearly transform the initial signal depending on how it differs from the bias. We aimed to extend Let the initial signal be $\mathbf{f}(\vec{r})$, and the bias signal with trainable Fourier coefficients be $\mathbf{b}(\vec{r})$. Then, we propose the following output of the normalization-activation block,

$$\mathbf{f}_a(\vec{r}) = \left(\frac{1}{4} \int_{\mathbb{R}^3} (N(\mathbf{f}(\vec{r}) + \mathbf{b}(\vec{r})) - N(\mathbf{b}(\vec{r})))^2 d^3\vec{r} \right) N(\mathbf{f}(\vec{r}) + \mathbf{b}(\vec{r})), \quad (8)$$

where $N()$ is the normalization operation defined in Eq. 7 such that the value $\frac{1}{4} \int_{\mathbb{R}^3} (N(\mathbf{f}(\vec{r}) + \mathbf{b}(\vec{r})) - N(\mathbf{b}(\vec{r})))^2 d^3\vec{r}$ lies in the interval $[0, 1]$. If the signal $\mathbf{f}(\vec{r})$ is 'coherent' to $\mathbf{b}(\vec{r})$, $\mathbf{f}(\vec{r}) = K\mathbf{b}(\vec{r})$, $K > 0$, then it does not pass the block. The amplification factor reaches its maximum value when the two signals are anti-coherent. Parseval's theorem allows us to translate these formulas into operations on their expansion coefficients,

$$[\mathbf{F}_a]_l^k(\rho) = \frac{1}{4} \left(\sum_{l'=0}^L \sum_{k'=-l'}^{l'} \int_0^\infty (N(\mathbf{F}_{l'}^{k'}(\rho) + \mathbf{B}_{l'}^{k'}(\rho)) - N(\mathbf{B}_{l'}^{k'}(\rho)))^2 \rho^2 d\rho \right) N(\mathbf{F}_l^k(\rho) + \mathbf{B}_l^k(\rho)), \quad (9)$$

where $[\mathbf{F}_a]_l^k(\rho_p)$ and $\mathbf{B}_l^k(\rho_p)$ are the expansion coefficients of functions $\mathbf{f}(\vec{r})$ and $\mathbf{b}(\vec{r})$, correspondingly.

3.7 Switching from the continuous to the vector representation

For most of the real-world tasks that may require the proposed architecture, a transition from a continuous functional representation $\mathbf{f}(\vec{r}) : \mathcal{R}^3 \rightarrow \mathcal{R}^{d_f}$ to a discrete vector representation $\mathbf{h} \in \mathcal{R}^{d_f}$ is necessary. There can be several ways to achieve it. In the simplest case, when the input function $\mathbf{f}(\vec{r})$ can be unambiguously associated with some reference coordinate system, as in the case of protein's peptide chain, we may use the following operation,

$$\mathbf{h} = \int_{\mathcal{R}^3} \mathbf{f}(\vec{r}) \mathbf{w}(\vec{r}) d\vec{r}, \quad (10)$$

where $\mathbf{w}(\vec{r}) : \mathcal{R}^3 \rightarrow \mathcal{R}^{d_f}$ is a filter function element-wise multiplied by the input function. If the functions $\mathbf{f}(\vec{r})$ and $\mathbf{w}(\vec{r})$ have corresponding expansion coefficients $\mathbf{F}_l^k(\rho_p)$ and $\mathbf{W}_l^k(\rho_p)$, then this operation will have the following form (please see more details in Appendix H),

$$\mathbf{h} = \sum_{l=0}^L \sum_{k=-l}^l \int_0^\infty \mathbf{F}_l^k(\rho) \overline{\mathbf{W}_l^k(\rho)} \rho^2 d\rho. \quad (11)$$

3.8 Graph convolutions layers

In our model, the continuous representation is followed by the classical graph convolutional layers (see Fig. 1C-D). Indeed, protein structures, on which we assess our model, allow us to use the graph representation. In such a graph, each node corresponds to an amino acid residue and characterizes the 3D structure of its neighborhood, and each edge between two nodes indicates their spatial proximity, i.e. the distance between the corresponding C-alpha atoms within a certain threshold R_n (see Fig. 1D).

Let us consider a graph \mathcal{G} that is described by the feature matrix $\mathbf{H} \in \mathcal{R}^{N \times d_v}$, where N is the number of graph nodes and d_v is the dimensionality of the node feature space, and the adjacency matrix $\mathbf{A} \in \mathcal{R}^{N \times N \times d_e}$, where d_e is the dimensionality of the edge feature space. We decided to use one-hot edge features that would encode the types of amino acids of the associated nodes. To reduce the dimensionality of the edge feature space to d'_e , we use the following trainable embedding $\mathbf{R}_e \in \mathcal{R}^{d_e \times d'_e}$, and a reduced adjacency matrix $\mathbf{A}_r = \mathbf{A}\mathbf{R}_e$. Finally, the graph convolution step is defined as

$$\mathbf{H}_{k+1} = \sigma_a(\mathbf{A}_r \mathbf{H}_k \mathbf{W}_k + \mathbf{H}_k \mathbf{W}_k^s + \mathbf{b}_k), \quad (12)$$

where $\mathbf{W}_k \in \mathcal{R}^{d_k \times d_{k+1} \times d'_e}$ and $\mathbf{W}_k^s \in \mathcal{R}^{d_k \times d_{k+1}}$ are trainable matrices.

4 Experiments

4.1 6D filters

Our first step was to study the properties of the 6D roto-translational convolution operation. To do so, we generated a small 3D pattern, $\mathbf{f}(\vec{r})$, composed of six 3D Gaussians with $\sigma = 0.4 \text{ \AA}$ shown in Fig. 2A-B. We then created a function

$\mathbf{h}(\vec{r})$, rotated and translated $\mathbf{f}(\vec{r})$ to a new location. Figure 2C-D shows the result of the 6D convolution between $\mathbf{h}(\vec{r})$ and $\mathbf{f}(\vec{r})$ given by Eq. 4. As we can expect, the maximum of this convolution corresponds to the position of the center of mass of function $\mathbf{h}(\vec{r})$, and the value of the convolution reduces as we go further from this point. We used the following parameters for this experiment, $\sigma = 0.4 \text{ \AA}$, $L = 4$, $\rho_{\max} = 0.6\pi \text{ \AA}^{-1}$, $\Delta\rho = 0.2\pi \text{ \AA}^{-1}$.




Figure 2: **A-B.** Six Gaussian volumetric features of $\sigma = 0.4 \text{ \AA}$ shown in the xy -plane (A) and xz -plane (B). The center of mass of the whole pattern is shown with the pink star. **C-D.** The density map of the resultant 6D convolution between an original pattern and its translated and rotated copy to the pink star position, shown in the xy -plane (C) and xz -plane (D). **E.** The maximum error in determining the center of the translated and rotated pattern as a function of the expansion order L divided by the Gaussian feature σ for 3D and 6D convolutions.

We then compared the result of the classical 3D convolution (see Appendix B) with the proposed 6D convolution. Using the same initial pattern, randomly rotated multiple times, we recorded the maximum positional error in determining the center of mass of the translated pattern with respect to the maximum expansion order L . As we can see in Fig. 2E, the 6D convolution detects the position of the shifted pattern more accurately compared to its classical 3D counterpart, and the accuracy increases with the maximum expansion order L .

4.2 Message passing and translation operator

Our next experiment was the assessment of the message-passing step. Our main goal was to study the conditions of validity for the translation operator 29, which is also implicitly used in the 6D convolution part. Specifically, we were interested in whether the result of the 6D convolution will be preserved after changing the coordinate systems. For this experiment, we used the same volumetric pattern $\mathbf{f}(\vec{r})$ described above and shown in Fig. 2A-B. We then shifted this pattern to a new location and recorded the value of the 6D convolution, as described above, shown in Fig. 3A-B. For the comparison, we computed the same 6D convolution from a different coordinated system, shifted by $3\sqrt{3} \text{ \AA}$ from the original one. We used parameters from the previous experiment. The result is shown in Fig. 3C-D. Both convolution functions have their maximums near the location of the center of mass of the shifted pattern, however, their volumetric shapes are slightly different.

For a more rigorous experiment, we examined the relative error of the translation operator 29 as a function of the displacement of the coordinate system and the expansion order. Here, we fixed parameters to $\sigma = 0.4 \text{ \AA}$, $\rho_{\max} = 0.6\pi \text{ \AA}^{-1}$, $\Delta\rho = 0.2\pi \text{ \AA}^{-1}$, and varied the value of L . From the results shown in Fig. 3E we can see that for the displacements within about $\sigma L/2$, the error is negligibly small, which is the consequence of the Nyquist–Shannon theorem.

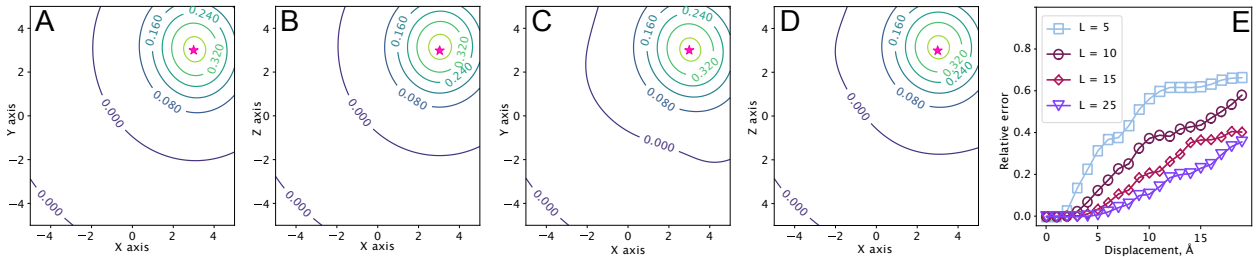


Figure 3: **A-B.** Results of the 6D convolution in the original coordinate system, shown in the xy -plane (A) and xz -plane (B). **C-D.** Results of the 6D convolution in the translated coordinate system, by $3\sqrt{3} \text{ \AA}$, shown in the xy -plane (C) and xz -plane (D). **E.** Relative error of the recovering volumetric patterns as a function of the translation amplitude and the expansion order.

4.3 Datasets

In order to train and assess our models, we used data from the Critical Assessment of Structure Prediction (CASP) benchmarks (Kryshtafovych et al., 2019; Moult et al., 1995). Each model from the benchmarks has a quality score that indicates how close it is to the experimentally determined solution. There are several quality metrics designed for protein models, e.g., GDT-TS (Zemla et al., 1999), CAD-score (Olechnovič et al., 2013), or IDDT (Mariani et al., 2013). We trained our networks to predict the local per-residue IDDT score, which shows the precision of the local 3D structure and topology around each amino acid. We also considered the average value of predicted local scores as the global score for the model.

We used models from CASP8-11 stage 2 submissions as the training and validation sets. We randomly split the targets into two sets. The training set contained 301 target structures with 104,985 models. The validation set contained 32 target structures with 10,872 models. Finally, we tested the networks on CASP12 stage2 dataset that has 39 native structures and 5,462 models, and also on CASP13 stage2 dataset with 18 native structures and 2,583 models. All the sets included only the models for which the native structure is publicly available.

4.4 Evaluation metrics

We assessed our networks on both global and local IDDT scores. In order to evaluate the efficiency of our networks, we engaged the following metrics: z-scores, determination coefficient R^2 , and Pearson and Spearman correlations. Z-scores are equal to the ground-truth score of the top-predicted model divided by the standard deviation of the models for a particular target. We computed R^2 and correlations in two modes, global and per-target. The per-target approach assumed calculating metrics for each target separately and averaging results over all the targets. We used the Fisher transformation to extract average values. For the global approach, we calculated metrics for all model structures of regardless which target they belong to. For the local structure quality predictions, as CASP submissions do not explicitly store local IDDT scores, we assessed our models using Spearman rank correlation of the submitted distance deviations with the ground-truth data.

4.5 Loss function

Our loss function contains two terms, the first is the MSE between the predictions and the ground-truth local scores, and the second is the MSE of the z-scores between these two arrays,

$$L(s_i, p_i) = \alpha \frac{1}{N} \sum_{i=1}^N (s_i - p_i)^2 + (1 - \alpha) \sigma_s^2 \frac{1}{N} \sum_{i=1}^N \left(\frac{s_i - \mu_s}{\sigma_s} - \frac{p_i - \mu_p}{\sigma_p} \right)^2, \quad (13)$$

where N is the number of residues in the protein model; $\mu_s = \frac{1}{N} \sum_i^N s_i$, $\sigma_s = \sqrt{\frac{1}{N} \sum_i^N (s_i - \mu_s)^2}$; μ_p and σ_p have the same formulas but for the array p_i . To make two terms in the loss function to have approximately the same amplitude, we introduced an additional variance of the ground-truth local scores factor σ_s^2 . The weight α is a hyperparameter. We need both terms in the loss function to simultaneously increase the correlation of the local predictions and to also predict the absolute values of the scores.

4.6 Technical details

Amino-acid residues can contain atoms of 167 distinguishable types. We have also included one additional type for the solvent molecules. Overall, the dimensionality of the volumetric function characterizing the protein model $N_a = 168$. We choose some maximum value of the radius R_{max} , which limits the set of atoms that fall in the amino-residue neighborhood. We also choose some parameter R_n , which is the maximum distance between amino residues that considered as neighbors in the graph.

We use 20 amino acids types. The edge between two amino acids can be determined by a vector of the size $20 \times 20 + d_t$, where the first 20×20 elements are a one-hot representation of the amino acids pair, with the distinguishable order of residues in a pair. The last d_t elements in this vector is a one-hot representation of the topological distance between the residues in the protein graph. The values of R_{max} , R_n , d_t , and L are the hyper-parameters of the network that were optimized with a grid search.

During the training process, we preprocessed all data using C++ code before feeding them to the network. For each residue of each protein model, we precomputed its Fourier representation. In the case of the architectures with the message-passing step, we have also generated tensors for the transformation of Fourier coefficients between coordinate system frames along the edges of the graph.

4.7 Baseline architecture

For the comparison, we introduced a baseline architecture that would help us to assess the novel layers. It begins with trainable embedding in the feature space. We then applied the transition from the continuous to the discrete representation using operation 11 that is followed by three graph convolutional layers described in Eq. 12. For the activation, we used the tanh function in the last layer and the LReLU function with the 'leak' parameter of 0.05 in all other layers. We also introduced two trainable parameters μ_t and σ_t for the mean and the standard deviation of the local rates in the training sample. The relationship between the output of the last layer l_N and the output of the network o is $o = \sigma_t l_N + \mu_t$. Overall, the baseline architecture had 21,026 trainable parameters.

Training. This network was trained on CASP 8-11 datasets in 1,280 iterations. At each iteration, the network was fed by 16 input protein models. One training iteration took ≈ 5 minutes on Intel ©Xeon(R) CPU E5-2630 v4 @ 2.20GHz.

Hyperparameters. The network has the following hyperparameters, $\sigma = 2 \text{ \AA}$, $R_n = 12 \text{ \AA}$, $d_t = 10$, $L = 4$, and $R_{max} = 8 \text{ \AA}$.

4.8 6DCNN networks

The main difference between the baseline architecture and the 6DCNN networks is the presence of 6D convolution layers in the latter. Our first architecture (6DCNN-1) contains only one 6D convolution layer. The second architecture (6DCNN-2) has two 6DCNN layers. The 6DCNN layer is composed of the following operations: 6D convolution, followed by normalization and activation. Consecutive 6D convolution layers are linked with the message passing step. Table 3 lists the architectures of the networks. Overall, the 6DCNN-1 and 6DCNN-2 architectures had 187,326 and 283,426 trainable parameters, correspondingly. Figure 4A shows real-space projections of two 6D convolution filters learned by 6DCNN-1.

Training. The two networks were trained on CASP 8-11 datasets in 1,280 iterations. At each iteration, the networks were fed by 16 protein models. One training iteration took ≈ 6 minutes for 6DCNN-1 and ≈ 15 minutes for 6DCNN-2 on Intel ©Xeon(R) CPU E5-2630 v4 @ 2.20GHz. Figure 4D demonstrates the learning curves on the validation dataset of three architectures, baseline, 6DCNN-1, and 6DCNN-2.

Hyperparameters. The networks have the following hyperparameters, $\sigma = 2 \text{ \AA}$, $R_n = 12 \text{ \AA}$, $d_t = 10$, $L = 4$, and $R_{max} = 8 \text{ \AA}$.

4.9 CASP results

In order to assess the 6DCNN architectures, we compared their performance on the CASP12 (Table 1) and CASP13 (Table 2) datasets with the baseline model and also with the state-of-the-art single-model quality assessment methods SBROD, SVMQA, VoroCNN, Ornate, ProQ3, and VoroMQA (Cheng et al., 2019). SBROD is a linear regression model that uses geometric features of the protein backbone (Karasikov et al., 2019). SVMQA is a support-vector-machine-based method that also uses structural features (Manavalan and Lee, 2017). VoroMQA engages statistical features of 3D Voronoi tessellation of the protein structure (Olechnovič and Venclovas, 2017). VoroCNN is a graph neural network built on the 3D Voronoi tessellation of protein structures (Igashov et al., 2021a). Ornate is a convolutional neural network that uses 3D volumetric representation of protein residues in their local reference frames (Pagès et al., 2019). ProQ3 is a fully connected neural network operating on the precomputed descriptors (Uziela et al., 2016). We computed the ground-truth IDDT values ourselves. Therefore, we were forced to limit the datasets to only those models that had publicly available target structures. As a result, the CASP12 dataset turned out to be significantly bigger than CASP13, with more demonstrative and representative results.

On the CASP12 test set, we achieved a noticeable improvement in comparison with the state-of-the-art methods. Even though the difference between the 6DCNN networks and the baseline model performance is not big, one of the 6DCNN architectures outperforms the baseline in every metric except for the z-score. We can also notice that the 6DCNN-2 method gives significantly higher global correlations and R^2 metrics on the CASP12 dataset than 6DCNN-1 and all other methods. However, 6DCNN-1 demonstrates better per-target correlations on CASP12 data than 6DCNN-2. Both of the networks have higher per-target correlations than most of the state-of-the-art methods. Unfortunately, we did not manage to achieve satisfying performance on the z-score metric. However, z-scores are rather noisy compared to correlations, and not directly linked to the optimized loss function. The fact that 6DCNN-2 has better global correlation scores confirms the importance of the additional 6D correlation block. Figure 4 (B-C) shows correlations between the ground-truth global scores from the CASP12 dataset and the corresponding predictions by the two 6DCNN models. The 6DCNN-2 map has a higher density near the diagonal, indicating a better absolute predictions of global scores and a better R^2 metric.

On the CASP13 dataset, we did not greatly outperform the state-of-the-art methods (see Table 2). However, we reached a performance that is on par with the state of the art. Moreover, we should notice that 6DCNN-2 did not outperform 6DCNN-1. This can be explained by the fact that we trained our models on CASP[8-11] datasets, which are rather different from CASP13, and also that the CASP13 dataset is less representative than CASP12.




Figure 4: **A.** Examples of two filters learned by 6DCNN-1, two rows represent the two filters. Each column demonstrates different filter projections. The red color represents the positive filter values, and the blue color corresponds to the negative values. **B-C.** Density maps comparing the ground-truth global IDDT scores with the predictions of the 6DCNN networks on the CASP12 dataset, 6DCNN-2 (left, B) and 6DCNN-1 (right, C). **D.** Loss function on the validation dataset during the training of the 6DCNN networks and the baseline model.

Table 4 from Appendix J lists Spearman rank correlations of local quality predictions with the corresponding ground-truth values of our networks and the state-of-the-art methods on model structures of 11 targets from CASP13. For the comparison, we chose only those models that had both publicly available target structures and local score predictions by all other methods. As we did not have these predictions for the CASP12 dataset, we limited local score evaluation by CASP13 data only. Here, we did not achieve the best results that could be explained by the small size of the dataset.

Method	z-score	Global			Per-target		
		R^2	Pearson, r	Spearman, ρ	R^2	Pearson, r	Spearman, ρ
SBROD	1,29	-33,66	0,55	0,54	-325,18	0,76	0,67
SVMQA	1,48	0,32	0,82	0,80	-2,44	0,76	0,73
VoroCNN	1,39	0,61	0,80	0,80	-2,79	0,73	0,69
Ornate	1,42	0,36	0,78	0,77	-5,14	0,73	0,69
ProQ3	1,18	0,26	0,74	0,77	-4,85	0,73	0,68
VoroMQA	1,18	-0,25	0,59	0,62	-6,23	0,74	0,70
Baseline	1,34	0,55	0,82	0,82	-2,35	0,78	0,71
6DCNN-1	1,26	0,57	0,81	0,81	-2,29	0,80	0,73
6DCNN-2	1,19	0,63	0,85	0,84	-1,95	0,79	0,71

Table 1: Comparison of the 6DCNN networks with the baseline architecture and the state-of-the-art methods on the unrefined CASP12 stage2 dataset. The best value for each metric is highlighted in bold.

Method	z-score	Global			Per-target		
		R^2	Pearson, r	Spearman, ρ	R^2	Pearson, r	Spearman, ρ
SBROD	1,23	0,07	0,72	0,69	-1,44	0,81	0,74
VoroCNN	1,15	0,67	0,84	0,82	0,03	0,79	0,77
VoroMQA	1,32	0,20	0,77	0,79	-1,10	0,79	0,75
ProQ3D	1,39	-0,03	0,75	0,75	-2,07	0,76	0,72
Ornate	0,93	0,26	0,62	0,67	-2,10	0,78	0,77
Baseline	1,35	0,44	0,78	0,77	-0,30	0,82	0,77
6DCNN-1	1,03	0,59	0,82	0,80	-0,02	0,83	0,79
6DCNN-2	1,30	0,56	0,79	0,78	-0,12	0,82	0,77

Table 2: Comparison of the 6DCNN networks with the baseline architecture and the state-of-the-art methods on the unrefined CASP13 stage2 dataset. The best value for each metric is highlighted in bold.

5 Conclusion

This work presents a theoretical foundation for 6D roto-translational spatial patterns detection and the construction of neural network architectures for learning on spatial continuous data in 3D. We built several networks that consisted of 6DCNN blocks followed by GCNN layers specifically designed for 3D models of protein structures. We then tested them on the CASP datasets from the community-wide protein structure prediction challenge. Our results demonstrate that 6DCNN blocks are able to accurately learn local spatial patterns and improve the quality prediction of protein models. The current network architecture can be extended in multiple directions, for example, including the attention mechanism.

References

- Brandon Anderson, Truong-Son Hy, and Risi Kondor. Cormorant: Covariant molecular neural networks. *arXiv preprint arXiv:1906.04015*, 2019.
- Natalie Baddour. Operational and convolution properties of three-dimensional fourier transforms in spherical polar coordinates. *J. Opt. Soc. Am. A*, 27(10):2144–2155, Oct 2010.
- Minkyung Baek, Frank DiMaio, Ivan Anishchenko, Justas Dauparas, Sergey Ovchinnikov, Gyu Rie Lee, Jue Wang, Qian Cong, Lisa N Kinch, R Dustin Schaeffer, et al. Accurate prediction of protein structures and interactions using a three-track neural network. *Science*, 2021.
- Federico Baldassarre, David Menéndez Hurtado, Arne Elofsson, and Hossein Azizpour. Graphqa: protein model quality assessment using graph convolutional networks. *Bioinformatics*, 37(3):360–366, 2021.
- Chi Chen, Weike Ye, Yunxing Zuo, Chen Zheng, and Shyue Ping Ong. Graph networks as a universal machine learning framework for molecules and crystals. *Chemistry of Materials*, 31(9):3564–3572, 2019.
- Jianlin Cheng, Myong-Ho Choe, Arne Elofsson, Kun-Sop Han, Jie Hou, Ali H A Maghrabi, Liam J McGuffin, David Menéndez-Hurtado, Kliment Olechnovič, Torsten Schwede, Gabriel Studer, Karolis Uziela, Česlovas Venclovas, and Björn Wallner. Estimation of model accuracy in casp13. *Proteins*, 87(12):1361–1377, 12 2019. doi: 10.1002/prot.25767.
- Taco Cohen, Mario Geiger, and Maurice Weiler. A general theory of equivariant CNNs on homogeneous spaces. *CoRR*, abs/1811.02017, 2018a. URL <http://arxiv.org/abs/1811.02017>.
- Taco S Cohen, Mario Geiger, Jonas Köhler, and Max Welling. Spherical CNNs. *arXiv preprint arXiv:1801.10130*, 2018b.
- Georgy Derevyanko and Guillaume Lamoureaux. Protein-protein docking using learned three-dimensional representations. *bioRxiv*, page 738690, 2019.
- Stephan Eismann, Patricia Suriana, Bowen Jing, Raphael JL Townshend, and Ron O Dror. Protein model quality assessment using rotation-equivariant, hierarchical neural networks. *arXiv preprint arXiv:2011.13557*, 2020.
- Fabian B Fuchs, Daniel E Worrall, Volker Fischer, and Max Welling. SE(3)-transformers: 3D roto-translation equivariant attention networks. *arXiv preprint arXiv:2006.10503*, 2020.
- Pablo Gainza, Freyr Sverrisson, Frederico Monti, Emanuele Rodola, D Boscaini, MM Bronstein, and BE Correia. Deciphering interaction fingerprints from protein molecular surfaces using geometric deep learning. *Nature Methods*, 17(2):184–192, 2020.
- Justin Gilmer, Samuel S Schoenholz, Patrick F Riley, Oriol Vinyals, and George E Dahl. Neural message passing for quantum chemistry. In *International Conference on Machine Learning*, pages 1263–1272. PMLR, 2017.
- Naozumi Hiranuma, Hahnbeom Park, Minkyung Baek, Ivan Anishchenko, Justas Dauparas, and David Baker. Improved protein structure refinement guided by deep learning based accuracy estimation. *Nature communications*, 12(1):1–11, 2021.
- Michael J Hutchinson, Charline Le Lan, Sheheryar Zaidi, Emilien Dupont, Yee Whye Teh, and Hyunjik Kim. Lieformer: Equivariant self-attention for lie groups. In Marina Meila and Tong Zhang, editors, *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pages 4533–4543. PMLR, 18–24 Jul 2021. URL <http://proceedings.mlr.press/v139/hutchinson21a.html>.
- Iliia Igashov, Liment Olechnovič, Maria Kadukova, Česlovas Venclovas, and Sergei Grudinin. VoroCNN: Deep convolutional neural network built on 3D voronoi tessellation of protein structures. *Bioinformatics*, Feb 2021a. doi: 10.1093/bioinformatics/btab118.

- Iliia Igashov, Nikita Pavlichenko, and Sergei Grudinin. Spherical convolutions on molecular graphs for protein model quality assessment. *Machine Learning: Science and Technology*, 2:045005, 2021b.
- John Ingraham, Vikas K. Garg, Regina Barzilay, and Tommi S. Jaakkola. Generative models for graph-based protein design. In *Deep Generative Models for Highly Structured Data, ICLR 2019 Workshop, New Orleans, Louisiana, United States, May 6, 2019*. OpenReview.net, 2019. URL <https://openreview.net/forum?id=SJgxrLLK0E>.
- Bowen Jing, Stephan Eismann, Pratham N. Soni, and Ron O. Dror. Equivariant graph neural networks for 3D macromolecular structure. *CoRR*, abs/2106.03843, 2021a. URL <https://arxiv.org/abs/2106.03843>.
- Bowen Jing, Stephan Eismann, Patricia Suriana, Raphael John Lamarre Townshend, and Ron O. Dror. Learning from protein structure with geometric vector perceptrons. In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net, 2021b. URL <https://openreview.net/forum?id=1YLJDvSx6J4>.
- John Jumper, Richard Evans, Alexander Pritzel, Tim Green, Michael Figurnov, Olaf Ronneberger, Kathryn Tunyasuvunakool, Russ Bates, Augustin Židek, Anna Potapenko, Alex Bridgland, Clemens Meyer, Simon A. A. Kohl, Andrew J. Ballard, Andrew Cowie, Bernardino Romera-Paredes, Stanislav Nikolov, Rishub Jain, Jonas Adler, Trevor Back, Stig Petersen, David Reiman, Ellen Clancy, Michal Zielinski, Martin Steinegger, Michalina Pacholska, Tamas Berghammer, Sebastian Bodenstein, David Silver, Oriol Vinyals, Andrew W. Senior, Koray Kavukcuoglu, Pushmeet Kohli, and Demis Hassabis. Highly accurate protein structure prediction with AlphaFold. *Nature*, jul 2021. doi: 10.1038/s41586-021-03819-2. URL <https://doi.org/10.1038/s41586-021-03819-2>.
- Mikhail Karasikov, Guillaume Pagès, and Sergei Grudinin. Smooth orientation-dependent scoring function for coarse-grained protein quality assessment. *Bioinformatics*, 35(16):2801–2808, 08 2019. doi: 10.1093/bioinformatics/bty1037.
- Johannes Klicpera, Shankari Giri, Johannes T Margraf, and Stephan Günnemann. Fast and uncertainty-aware directional message passing for non-equilibrium molecules. *arXiv preprint arXiv:2011.14115*, 2020a.
- Johannes Klicpera, Janek Groß, and Stephan Günnemann. Directional message passing for molecular graphs. *arXiv preprint arXiv:2003.03123*, 2020b.
- Risi Kondor. N-body networks: a covariant hierarchical neural network architecture for learning atomic potentials. *arXiv preprint arXiv:1803.01588*, 2018.
- Risi Kondor, Zhen Lin, and Shubhendu Trivedi. Clebsch-gordan nets: a fully fourier space spherical convolutional neural network. *arXiv preprint arXiv:1806.09231*, 2018.
- Andriy Kryshchak, Torsten Schwede, Maya Topf, Krzysztof Fidelis, and John Moult. Critical assessment of methods of protein structure prediction (casp)-round xiii. *Proteins*, 87(12):1011–1020, 12 2019. doi: 10.1002/prot.25823.
- Elodie Laine, Stephan Eismann, Arne Elofsson, and Sergei Grudinin. Protein sequence-to-structure learning: Is this the end(-to-end revolution)? *CoRR*, abs/2105.07407, 2021. URL <https://arxiv.org/abs/2105.07407>.
- Balachandran Manavalan and Jooyoung Lee. Svmqa: support-vector-machine-based protein single-model quality assessment. *Bioinformatics*, 33(16):2496–2503, Aug 2017. doi: 10.1093/bioinformatics/btx222.
- Valerio Mariani, Marco Biasini, Alessandro Barbato, and Torsten Schwede. Iddt: a local superposition-free score for comparing protein structures and models using distance difference tests. *Bioinformatics*, 29(21):2722–8, Nov 2013. doi: 10.1093/bioinformatics/btt473.
- John Moult, Jan T Pedersen, Richard Judson, and Krzysztof Fidelis. A large-scale experiment to assess protein structure prediction methods. *Proteins: Structure, Function, and Genetics*, 23(3):ii–v, Nov 1995. doi: 10.1002/prot.340230303.
- Kliment Olechnovič and Česlovas Venclovas. Voromqa: Assessment of protein structure quality using interatomic contact areas. *Proteins*, 85(6):1131–1145, 06 2017. doi: 10.1002/prot.25278.
- Kliment Olechnovič, Eleonora Kulberkytė, and Česlovas Venclovas. CAD-score: a new contact area difference-based function for evaluation of protein structural models. *Proteins*, 81(1):149–62, Jan 2013. doi: 10.1002/prot.24172.
- Guillaume Pagès, Benoit Charmettant, and Sergei Grudinin. Protein model quality assessment using 3D oriented convolutional neural networks. *Bioinformatics*, 35(18):3313–3319, 2019.
- Adrien Poulenc, Marie-Julie Rakotosaona, Yann Ponty, and Maks Ovsjanikov. Effective rotation-invariant point CNN with spherical harmonics kernels. In *2019 International Conference on 3D Vision (3DV)*, pages 47–56. IEEE, 2019.
- David Romero, Erik Bekkers, Jakub Tomczak, and Mark Hoogendoorn. Attentive group equivariant convolutional networks. In Hal Daumé III and Aarti Singh, editors, *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pages 8188–8199. PMLR, 13–18 Jul 2020. URL <http://proceedings.mlr.press/v119/romero20a.html>.

- David W. Romero and Jean-Baptiste Cordonnier. Group equivariant stand-alone self-attention for vision. In *International Conference on Learning Representations*, 2021. URL <https://openreview.net/forum?id=JkfYjn0Eo6M>.
- Soumya Sanyal, Ivan Anishchenko, Anirudh Dagar, David Baker, and Partha Talukdar. Proteingcn: Protein model quality assessment using graph convolutional networks. *BioRxiv*, 2020.
- Victor Garcia Satorras, Emiel Hoogeboom, Fabian B Fuchs, Ingmar Posner, and Max Welling. E (n) equivariant normalizing flows for molecule generation in 3D. *arXiv preprint arXiv:2105.09016*, 2021a.
- Victor Garcia Satorras, Emiel Hoogeboom, and Max Welling. E (n) equivariant graph neural networks. *arXiv preprint arXiv:2102.09844*, 2021b.
- Kristof T Schütt, Pieter-Jan Kindermans, Huziel E Sauceda, Stefan Chmiela, Alexandre Tkatchenko, and Klaus-Robert Müller. Schnet: A continuous-filter convolutional neural network for modeling quantum interactions. *arXiv preprint arXiv:1706.08566*, 2017.
- Kristof T Schütt, Oliver T Unke, and Michael Gastegger. Equivariant message passing for the prediction of tensorial properties and molecular spectra. *arXiv preprint arXiv:2102.03150*, 2021.
- Andrew W Senior, Richard Evans, John Jumper, James Kirkpatrick, Laurent Sifre, Tim Green, Chongli Qin, Augustin Židek, Alexander WR Nelson, Alex Bridgland, et al. Improved protein structure prediction using potentials from deep learning. *Nature*, 577(7792):706–710, 2020.
- Freyr Sverrisson, Jean Feydy, Bruno Correia, and Michael Bronstein. Fast end-to-end learning on protein surfaces. *bioRxiv*, 2020.
- Nathaniel Thomas, Tess Smidt, Steven Kearnes, Lusann Yang, Li Li, Kai Kohlhoff, and Patrick Riley. Tensor field networks: Rotation-and translation-equivariant neural networks for 3D point clouds. *arXiv preprint arXiv:1802.08219*, 2018.
- Raphael JL Townshend, Brent Townshend, Stephan Eismann, and Ron O Dror. Geometric prediction: Moving beyond scalars. *arXiv preprint arXiv:2006.14163*, 2020.
- Karolis Uziela, Nanjiang Shu, Björn Wallner, and Arne Elofsson. Proq3: Improved model quality assessments using rosetta energy terms. *Sci Rep*, 6:33509, 10 2016. doi: 10.1038/srep33509.
- Maurice Weiler, Mario Geiger, Max Welling, Wouter Boomsma, and Taco Cohen. 3D steerable CNNs: Learning rotationally equivariant features in volumetric data. *arXiv preprint arXiv:1807.02547*, 2018.
- Daniel E Worrall, Stephan J Garbin, Daniyar Turmukhambetov, and Gabriel J Brostow. Harmonic networks: Deep translation and rotation equivariance. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5028–5037, 2017.
- Adam Zemla, Česlovas Venclovas, John Moult, and Krzysztof Fidelis. Processing and analysis of CASP3 protein structure predictions. *Proteins: Structure, Function, and Bioinformatics*, 37(S3):22–29, 1999.

Appendix

A Mathematical prerequisites

A.1 Spherical harmonics and spherical Bessel transform

Spherical harmonics are complex functions defined on the surface of a sphere that constitute a complete set of orthogonal functions and thus an orthonormal basis. They are generally defined as

$$Y_l^k(\Omega) = \sqrt{\frac{(2l+1)(l-k)!}{4\pi(l+k)!}} P_l^k(\cos(\psi_\Omega)) \exp ik\theta_\Omega, \quad (14)$$

where $0 \leq \psi_\Omega \leq \pi$ is the colatitude of the point Ω , and $0 \leq \theta_\Omega \leq 2\pi$ is the longitude. As mentioned above, these functions are orthonormal,

$$\int_{4\pi} Y_l^k(\Omega) \overline{Y_{l'}^{k'}(\Omega)} d\Omega = \delta_{ll'} \delta_{kk'}, \quad (15)$$

where the second function $\overline{Y_{l'}^{k'}(\Omega)}$ is complex-conjugated. An expansion of a function $f(\vec{r})$ in spherical harmonics has the following form,

$$f_l^k(r) = \int_{4\pi} f(\vec{r}) \overline{Y_l^k(\Omega_r)} d\Omega_r. \quad (16)$$

Spherical Bessel transform (SBT, sometimes referred to as spherical Hankel transform) of order l computes Fourier coefficients of spherically symmetric functions in 3D,

$$F_l(\rho) = \text{SBT}_l(f(r)) = \int f(r) j_l(\rho r) r^2 dr, \quad (17)$$

where $j_l(r)$ are spherical Bessel functions of order l . The inverse transform has the following form,

$$f(r) = \text{SBT}_l^{-1}(F_l(\rho)) = \frac{2}{\pi} \int F_l(\rho) j_l(\rho r) \rho^2 d\rho. \quad (18)$$

Spherical Bessel functions of the same order are orthogonal with respect to the argument,

$$\int_0^\infty j_l(\rho_1 r) j_l(\rho_2 r) r^2 dr = \frac{\pi}{2\rho_1 \rho_2} \delta(\rho_1 - \rho_2). \quad (19)$$

A.2 Plane wave expansion

The plane wave expansion is the decomposition of a plane wave to a linear combination of spherical waves. According to the spherical harmonic addition theorem, a plane wave can be expressed through spherical harmonics and spherical Bessel functions,

$$\exp i\vec{\rho}\vec{r} = 4\pi \sum_{l=0}^{\infty} \sum_{k=-l}^l i^l j_l(\rho r) Y_l^k(\Omega_\rho) \overline{Y_l^k(\Omega_r)}. \quad (20)$$

B 3D Fourier transforms in spherical coordinates

The 3D Fourier transform of a function $f(\vec{r})$ is defined as

$$F(\vec{\rho}) = \int_{\mathbb{R}^3} f(\vec{r}) \exp i\vec{r}\vec{\rho} d\vec{r}. \quad (21)$$

The spherical harmonics expansion of this transform has the following form,

$$F_l^k(\rho) = \int_{4\pi} F(\vec{\rho}) Y_l^k(\Omega_\rho) d\Omega_\rho. \quad (22)$$

The Fourier spherical harmonics expansion coefficients relate to the real-space spherical harmonics coefficients through SBT,

$$F_l^k(\rho) = 4\pi(-i)^l \text{SBT}_l(f_l^k(r)). \quad (23)$$

This equation can be also rewritten as follows,

$$F_l^k(\rho) = 4\pi(-i)^l \int_{\mathbb{R}^3} f(\vec{r}) j_l(\rho r) \overline{Y_l^k(\Omega_r)} d\vec{r}. \quad (24)$$

C Rotation of spherical harmonics and Wigner D -matrices

The Wigner D -matrices \mathbf{D}_l are the irreducible representations of $\text{SO}(3)$ that can be applied to spherical harmonics functions to express the rotated functions with tensor operations on the original ones,

$$Y_l^k(\Lambda\Omega) = \sum_{k'=-l}^l D_{kk'}^l(\Lambda) Y_l^{k'}(\Omega), \quad (25)$$

where $\Lambda \in \text{SO}(3)$.

D Translation operator

Let us translate a 3D function $f(\vec{r})$ along the z -axis by an amount Δ . The expansion coefficients of a translated function will be

$$[F_z]_l^k(\rho) = \int F(\vec{\rho}) e^{-i\Delta\vec{\rho}\cdot\vec{e}_z} \overline{Y_l^k(\Omega)} d\Omega. \quad (26)$$

Using the plane-wave expansion and triple spherical harmonics integrals defined though Slater coefficients $c^{l_2}(l, k, l_1, k_1)$,

$$c^{l_2}(l, k, l_1, k_1) = \int_{4\pi} \overline{Y_l^k(\Omega)} Y_{l_1}^{k_1}(\Omega) Y_{l_2}^{k-k_1}(\Omega) d\Omega, \quad (27)$$

we obtain

$$\begin{aligned} [F_z]_l^k(\rho) &= \sum_{p=0}^{\infty} \sum_{l'=\max(|l-p|, |k|)}^{l+p} i^p j_p(\rho\Delta) F_{l'}^k(\rho) 4\pi \sqrt{\frac{2p+1}{4\pi}} \int_{4\pi} Y_l^k(\Omega) Y_p^0(\Omega) \overline{Y_{l'}^k(\Omega)} d\Omega \\ &= \sum_{p=0}^{\infty} \sum_{l'=\max(|l-p|, |k|)}^{l+p} i^p j_p(\rho\Delta) F_{l'}^k(\rho) 4\pi \sqrt{\frac{2p+1}{4\pi}} c^p(l', k, l, k). \end{aligned} \quad (28)$$

Changing the summation order and introducing the maximum expansion order L , we arrive at

$$[F_z]_l^k(\rho) = \sum_{l'=|k|}^L T_{l,l'}^k(\rho, \Delta) F_{l'}^k(\rho) + O\left(\frac{1}{(L-l)!} \left(\frac{\rho\Delta}{2}\right)^{L-l}\right), \quad (29)$$

where

$$T_{l,l'}^k(\rho, \Delta) = \sum_p i^p j_p(\rho\Delta) 4\pi \sqrt{\frac{2p+1}{4\pi}} c^p(l', k, l, k). \quad (30)$$

The error estimation in Eq. 29 is based on the equivalence relation $j_p(x) \sim \frac{1}{\sqrt{xp}} \left(\frac{ex}{2p}\right)^p$ as $p \rightarrow \infty$ that allows us to find the upper bound of the error,

$$\sum_{p=L+1}^{\infty} \sum_{l'=\max(|l-p|, |k|)}^{l+p} i^p j_p(\rho\Delta) F_{l'}^k(\rho) 4\pi \sqrt{\frac{2p+1}{4\pi}} c^p(l', k, l, k) \leq C \frac{(\frac{\rho\Delta}{2})^2}{(L-l-1)!}, \quad (31)$$

where C is a constant independent of L .

E Fourier coefficients for a 3D Gaussian function

Let us consider a function of the Gaussian type in 3D,

$$f(\vec{r}) = \exp\left(-\frac{(\vec{r} - \vec{r}_0)^2}{2\sigma^2}\right). \quad (32)$$

The Fourier transform of $f(\vec{r})$ will be

$$\begin{aligned} F(\vec{\rho}) &= \int_{\mathcal{R}^3} \exp\left(-\frac{(\vec{r} - \vec{r}_0)^2}{2\sigma^2}\right) \exp(-i\vec{r}\vec{\rho}) d^3\vec{r} \\ &= \exp\left(-\frac{i2\sigma^2\vec{r}_0\vec{\rho} + \sigma^4\rho^2}{2\sigma^2}\right) \int \exp\left(-\frac{r^2 + 2\vec{r}(i\sigma^2\vec{\rho} - \vec{r}_0) + r_0^2 - 2i\sigma^2\vec{r}_0\vec{\rho} - \sigma^4\rho^2}{2\sigma^2}\right) d^3\vec{r} \\ &= \exp(-i\vec{r}_0\vec{\rho}) \exp\left(-\frac{\sigma^2\rho^2}{2}\right) (\sqrt{2\pi}\sigma)^3. \end{aligned} \quad (33)$$

Using the plane wave expansion (20), we obtain

$$F(\vec{\rho}) = 4\pi \exp\left(-\frac{\sigma^2\rho^2}{2}\right) (\sqrt{2\pi}\sigma)^3 \sum_{l=0}^{\infty} \sum_{k=-l}^l (-i)^l j_l(\rho r_0) \overline{Y_l^k(\Omega_{r_0})} Y_l^k(\Omega_{\rho}) = \sum_{l=0}^{\infty} \sum_{k=-l}^l F_l^k(\rho) Y_l^k(\Omega_{\rho}). \quad (34)$$

Consequently,

$$F_l^k(\rho) = 4\pi (-i)^l \exp\left(-\frac{\sigma^2\rho^2}{2}\right) (\sqrt{2\pi}\sigma)^3 j_l(\rho r_0) \overline{Y_l^k(\Omega_{r_0})}. \quad (35)$$

F Proof of Eq. (4)

Proof. We start in the proof from the translation operator in 3D (Baddour, 2010),

$$f(\vec{r} - \vec{r}_0) = \sum_{l=0}^{\infty} \sum_{k=-l}^l 8(i)^l Y_l^k(\Omega_{r_0}) \sum_{l_1=0}^{\infty} \sum_{k_1=-l_1}^{l_1} (-i)^{l_1} \overline{Y_{l_1}^{k_1}(\Omega_r)} \sum_{l_2=\|l-l_1\|}^{l+l_1} (-i)^{l_2} c^{l_2}(l, k, l_1, k_1) \int_0^{\infty} f_{l_2}^{k-k_1}(u) S_{l_2}^{l_1, l_1}(u, r_0, r) u^2 du, \quad (36)$$

where $c^{l_2}(l, k, l_1, k_1)$ are Slater coefficients defined in Eq. 27. They are nonzero only for $\|l - l_1\| \leq l_2 \leq \|l + l_1\|$. $S_{l_2}^{l_1, l_1}(u, r, r_0)$ is a triple Bessel product which is defined as

$$S_{l_2}^{l_1, l_1}(u, r, r_0) = \int_0^{\infty} j_{l_2}(\rho u) j_{l_1}(\rho r_0) j_l(\rho r) \rho^2 d\rho \quad (37)$$

Let us consider the function $f(\vec{r}_0 - \Lambda^{-1}\vec{r})$, where $\Lambda \in \text{SO}(3)$,

$$f(\vec{r}_0 - \Lambda^{-1}\vec{r}) = \sum_{l_1=0}^{\infty} \sum_{k_2=-l_1}^{l_1} 8(i)^{l_1} Y_{l_1}^{k_2}(\Omega_{r_0}) \sum_{l_2=0}^{\infty} \sum_{k_3=-l_2}^{l_2} \sum_{k_4=-l_2}^{l_2} (-i)^{l_2} \overline{Y_{l_2}^{k_3}(\Omega_r)} D_{k_4 k_3}^{l_2}(\Lambda^{-1}) \sum_{l_3=|l_1-l_2|}^{l_1+l_2} (-i)^{l_3} c^{l_3}(l_1, k_2, l_2, k_4) \int_0^{\infty} f_{l_3}^{k_2-k_4}(u) S_{l_3}^{l_1, l_2}(u, r_0, r) u^2 du. \quad (38)$$

Using the symmetry properties of Wigner D -matrices,

$$D_{k_4 k_3}^{l_2}(\Lambda^{-1}) = (-1)^{(-k_4+k_3)} D_{-k_3 -k_4}^{l_2}(\Lambda), \quad (39)$$

we can retrieve the following equation for the convolution operation in 6D,

$$\begin{aligned} h(\vec{r}_0) &= \int_{\vec{r}} \int_{\Lambda} f(\vec{r}_0 - \Lambda^{-1}\vec{r}) g(\Lambda\vec{r}) d\vec{r} d\Lambda \\ &= \int_{\vec{r}} \int_{\Lambda} \sum_{l=0}^{\infty} \sum_{k=-l}^l \sum_{k_1=-l}^l g_{l_1}^{k_1}(r) \overline{D_{k_1 k}^l(\Lambda)} Y_l^k(\Omega_r) \sum_{l_1=0}^{\infty} \sum_{k_2=-l_1}^{l_1} 8(i)^{l_1} Y_{l_1}^{k_2}(\Omega_{r_0}) \\ &\quad \sum_{l_2=0}^{\infty} \sum_{k_3=-l_2}^{l_2} \sum_{k_4=-l_2}^{l_2} (-i)^{l_2} \overline{Y_{l_2}^{k_3}(\Omega_r)} (-1)^{(-k_3+k_4)} D_{-k_3 -k_4}^{l_2}(\Lambda) \\ &\quad \sum_{l_3=|l_1-l_2|}^{l_1+l_2} (-i)^{l_3} c^{l_3}(l_1, k_2, l_2, k_4) \int_0^{\infty} f_{l_3}^{k_2-k_4}(u) S_{l_3}^{l_1, l_2}(u, r_0, r) u^2 du d\vec{r} d\Lambda. \end{aligned} \quad (40)$$

Using the orthogonality of Wigner D -matrices,

$$\int_{\text{SO}(3)} \overline{D_{k_1 k}^l(\Lambda)} D_{k_3 k_2}^{l_1}(\Lambda) d\Lambda = \frac{8\pi^2}{2l+1} \delta_{l_1 l} \delta_{k_2 k_3} \delta_{k_1 k_3}, \quad (41)$$

and the orthogonality of spherical harmonics, we obtain

$$\begin{aligned} h(\vec{r}_0) &= \int_0^{\infty} \int_0^{\infty} \sum_{l=0}^{\infty} \sum_{k=-l}^l \sum_{k_1=-l}^l \\ &\quad \sum_{l_1=0}^{\infty} \sum_{k_2=-l_1}^{l_1} \sum_{l_2=0}^{\infty} \sum_{k_3=-l_2}^{l_2} \sum_{k_4=-l_2}^{l_2} \sum_{l_3=|l_1-l_2|}^{l_1+l_2} 8(i)^{l_1} (-i)^{l_2} (-i)^{l_3} \frac{8\pi^2}{2l+1} \delta_{l_2 l} \delta_{k_1 -k_3} \delta_{k -k_4} \delta_{k k_3} \delta_{l l_2} \\ &\quad Y_{l_1}^{k_2}(\Omega_{r_0}) g_{l_1}^{k_1}(r) c^{l_3}(l_1, k_2, l_2, k_4) (-1)^{(k_3-k_4)} f_{l_3}^{k_2-k_4}(u) S_{l_3}^{l_1, l_2}(u, r_0, r) u^2 du r^2 dr \\ &= \int_0^{\infty} \int_0^{\infty} \sum_{l=0}^{\infty} \sum_{k=-l}^l \sum_{l_1=0}^{\infty} \sum_{k_2=-l_1}^{l_1} \sum_{l_3=|l_1-l_2|}^{l_1+l_2} 8(-i)^l (i)^{l_1} (-i)^{l_3} \frac{8\pi^2}{2l+1} g_l^{-k}(r) \\ &\quad c^{l_3}(l_1, k_2, l, -k) Y_{l_1}^{k_2}(\Omega_{r_0}) f_{l_3}^{k_2+k}(u) S_{l_3}^{l_1, l}(u, r_0, r) u^2 du r^2 dr. \end{aligned} \quad (42)$$

Let us change indices: $l_1 \rightarrow l, l \rightarrow l_1, k_2 \rightarrow k, k \rightarrow k_1, l_3 \rightarrow l_2$,

$$\begin{aligned}
 h(\vec{r}_0) &= \sum_{l=0}^{\infty} \sum_{k=-l}^l Y_l^k(\Omega_{r_0}) 8(i)^l \int_0^{\infty} \sum_{l_1=0}^{\infty} \sum_{k_1=-l_1}^{l_1} \frac{8\pi^2}{2l_1+1} (-i)^{l_1} \left(\int_0^{\infty} g_{l_1}^{-k_1}(r) j_l(\rho r) r^2 dr \right) \\
 &\quad \sum_{l_2=|l-l_1|}^{l+l_1} (-i)^{l_2} c^{l_2}(l, k, l_1, -k) \left(\int_0^{\infty} f_{l_2}^{k+k_1}(u) j_{l_2}(\rho u) u^2 du \right) j_l(\rho r_0) \rho^2 d\rho = \\
 &= \sum_{l=0}^{\infty} \sum_{k=-l}^l Y_l^k(\Omega_{r_0}) \frac{2}{\pi^2} (i)^l \int_0^{\infty} \left(\sum_{l_1=0}^{\infty} \sum_{k_1=-l_1}^{l_1} \frac{8\pi^2}{2l_1+1} G_{l_1}^{-k_1}(\rho) \sum_{l_2=|l-l_1|}^{l+l_1} c^{l_2}(l, k, l_1, -k_1) F_{l_2}^{k+k_1}(\rho) \right) j_l(\rho r_0) \rho^2 d\rho.
 \end{aligned} \tag{43}$$

Using the spherical harmonics expansion,

$$h(\vec{r}_0) = \sum_{l=0}^{\infty} \sum_{k=-l}^l Y_l^k(\Omega_{r_0}) \frac{1}{2\pi^2} (i)^l \int_0^{\infty} H_l^k(\rho) j_l(\rho r_0) \rho^2 d\rho, \tag{44}$$

we obtain the final result,

$$H_l^k(\rho) = \sum_{l_1=0}^{\infty} \sum_{k_1=-l_1}^{l_1} \frac{8\pi^2}{2l_1+1} G_{l_1}^{-k_1}(\rho) \sum_{l_2=|l-l_1|}^{l+l_1} c^{l_2}(l, k, l_1, -k_1) F_{l_2}^{k+k_1}(\rho). \tag{45}$$

□

G Normalization

Let us consider a function $f(\vec{r}) : \mathcal{R}^3 \rightarrow \mathcal{R}$ and its integral over \mathcal{R}^3 , $\int_{\mathcal{R}^3} f(\vec{r}) d^3 \vec{r}$. Using the fact that $\forall \Omega, Y_0^0(\Omega) = \frac{1}{\sqrt{4\pi}}$ and $j_0(0) = 1$ and the properties of orthogonality for spherical harmonics and spherical Bessel functions we rewrite this expression as

$$\int_{\mathcal{R}^3} f(\vec{r}) d^3 \vec{r} = \int_0^{\infty} \int_{4\pi} \sum_{l=0}^{\infty} \sum_{k=-l}^l \frac{(i)^l}{4\pi} \frac{2}{\pi} \int_0^{\infty} F_l^k(\rho) j_l(r\rho) \rho^2 d\rho \sqrt{4\pi} Y_l^k(\Omega) Y_0^0(\Omega) d\Omega j_0(0r) r^2 dr = \frac{1}{\sqrt{4\pi}} F_0^0(0). \tag{46}$$

If we set $F_0^0(0)$ to zero, we obtain the Fourier expansion of a function $f_m(\vec{r})$ that can be formulated as follows,

$$f_m(\vec{r}) = f(\vec{r}) - \int_{\mathcal{R}^3} f(\vec{r}) d^3 \vec{r}. \tag{47}$$

Below we also explain how to normalize the integral of a square of a function to unity using Parseval's theorem.

H Overlap integrals in 3D

H.1 Parseval's theorem

Let us consider two finite-resolution complex-valued functions $g(\vec{r})$ and $f(\vec{r})$, meaning that their spherical Fourier expansion coefficients $F_l^k(\rho)$ and $G_l^k(\rho)$ are nonzero only for $l \leq$ than some maximum expansion order L . Parseval's theorem for these coefficients states the following (Baddour, 2010),

$$\int_{\mathcal{R}^3} g(\vec{r}) \overline{f(\vec{r})} d\vec{r} = \frac{1}{(2\pi)^3} \sum_{l=0}^{\infty} \sum_{k=-l}^l \int_0^{\infty} G_l^k(\rho) \overline{F_l^k(\rho)} \rho^2 d\rho = \frac{1}{(2\pi)^3} \sum_{l=0}^L \sum_{k=-l}^l \int_0^{\infty} G_l^k(\rho) \overline{F_l^k(\rho)} \rho^2 d\rho, \tag{48}$$

In the case $g(\vec{r}) = f(\vec{r})$, it reduces to

$$\int_{\mathcal{R}^3} \|f(\vec{r})\|^2 d\vec{r} = \frac{1}{(2\pi)^3} \sum_{l=0}^{\infty} \sum_{k=-l}^l \int_0^{\infty} \|F_l^k(\rho_p)\|^2 \rho_p^2 d\rho = \frac{1}{(2\pi)^3} \sum_{l=0}^L \sum_{k=-l}^l \int_0^{\infty} \|F_l^k(\rho_p)\|^2 \rho_p^2 d\rho. \tag{49}$$

H.2 Switching from continuous to vector representation

Let $\mathbf{f}(\vec{r}) : \mathcal{R}^3 \rightarrow \mathcal{R}^{d_f}$ and $\mathbf{w}(\vec{r}) : \mathcal{R}^3 \rightarrow \mathcal{R}^{d_f}$ be functions with a finite resolution, i.e., their Fourier expansion coefficients are zero for all $l \leq$ than some maximum expansion order L . In order to obtain a vector $\mathbf{h} \in \mathcal{R}^{d_f}$, we integrate element-wise product of two functions over the 3D space. We are using the following relation to compute our vector representation,

$$\mathbf{h} = \int_{\mathcal{R}^3} \mathbf{f}(\vec{r})\mathbf{w}(\vec{r})d\vec{r} = \frac{1}{(2\pi)^3} \sum_{l=0}^L \sum_{k=-l}^l \int_0^\infty \mathbf{F}_l^k(\rho)\overline{\mathbf{W}}_l^k(\rho)\rho^2 d\rho, \quad (50)$$

where we approximate the last integral with a numerical integration.

I Network architectures

Table 3 lists three network architectures used in our experiments.

Layers	Baseline	6DCNN-1	6DCNN-2
Edges features embedding	Filter (410, 10); $N_p = 4,100$	Filter (410, 10); $N_p = 4,100$	Filter (410, 10); $N_p = 4,100$
Nodes features embedding	Filter (168, 40); $N_p = 6,720$	Filter (168, 40); $N_p = 6,720$	Filter (168, 40); $N_p = 6,720$
6D conv layer	-	-	Filter $\forall l(l, 4, 40, 40)$; $N_p = 131,200$
Message-passing	-	-	No trainable parameters
6D conv layer	-	Filter $\forall l(l, 4, 40, 40)$; $N_p = 131,200$	Filter $\forall l(l, 4, 40, 40)$; $N_p = 131,200$
Continuous to discrete	Filter $\forall l(l, 4, 40)$; $N_p = 3,200$	Filter $\forall l(l, 4, 40)$; $N_p = 3,200$	Filter $\forall l(l, 4, 40)$; $N_p = 3,200$
GC layer	Filters (40, 14, 10) and (40, 14); $N_p = 6,174$	Filters (40, 14, 10) and (40, 14); $N_p = 6,174$	Filters (40, 14, 10) and (40, 14); $N_p = 6,174$
GC layer	Filters (14, 5, 10) and (14, 5); $N_p = 775$	Filters (14, 5, 10) and (14, 5); $N_p = 775$	Filters (14, 5, 10) and (14, 5); $N_p = 775$
GC layer	Filters (5, 1, 10) and (5, 1); $N_p = 55$	Filters (5, 1, 10) and (5, 1); $N_p = 55$	Filters (5, 1, 10) and (5, 1); $N_p = 55$
The total number of parameters	21,026	187,326	283,426

Table 3: Comparison of the 6DCNN and the baseline network architectures with the number of trainable parameters N_p on each layer.

J Local scores' prediction results.

Table 4 lists Spearman rank correlations of local quality predictions with the corresponding ground-truth values of our three networks and the state-of-the-art methods on model structures of 11 targets from CASP13. For the comparison, we chose only model structures that had both publicly available target structures and predictions of local scores from other methods.

Methods	$ \rho $
VoroCNN	0.795
ModFOLD7	0.776
VoroMQA-A	0.665
ProQ3	0.744
ProQ4	0.728
Ornate	0.540
Baseline	0.735
6DCNN-1	0.774
6DCNN-2	0.763

Table 4: Spearman rank correlations between local scores and its predictions of the 6DCNN networks, the baseline architecture, and the state-of-the-art methods on the unrefined CASP13 stage2 dataset. Only 11 targets with the publicly available structure are included here. The best value is highlighted in bold.