Exploiting Hierarchical Interactions for Protein Surface Learning

Yiqun Lin, Liang Pan, Yi Li, Ziwei Liu, and Xiaomeng Li*

arXiv:2401.10144v1 [q-bio.BM] 17 Jan 2024

Abstract—**Predicting interactions between proteins is one of the most important yet challenging problems in structural bioinformatics. Intrinsically, potential function sites in protein surfaces are determined by both geometric and chemical features. However, existing works only consider handcrafted or individually learned chemical features from the atom type and extract geometric features independently. Here, we identify two key properties of effective protein surface learning: 1) relationship among atoms: atoms are linked with each other by covalent bonds to form biomolecules instead of appearing alone, leading to the significance of modeling the relationship among atoms in chemical feature learning. 2) hierarchical feature interaction: the neighboring residue effect validates the significance of hierarchical feature interaction among atoms and between surface points and atoms (or residues). In this paper, we present a principled framework based on deep learning techniques, namely Hierarchical Chemical and Geometric Feature Interaction Network (HCGNet), for protein surface analysis by bridging chemical and geometric features with hierarchical interactions. Extensive experiments demonstrate that our method outperforms the prior state-of-the-art method by 2.3% in site prediction task and 3.2% in interaction matching task, respectively. Our code is available at https://github.com/xmed-lab/HCGNet.**

Index Terms—**Biology, Computer Vision, Point Cloud, Protein-Protein Interaction, Surface Learning.**

I. INTRODUCTION

PROTEINS composed of amino acids are large, complex molecules that play critical roles in all living organisms. Proteins can provide various functions with different structures for organisms, including causing biomedical reactions, acting as messengers, and balancing fluid. From a biological perspective, proteins can be described in four structural levels, including primary, secondary, tertiary, and quaternary structures. The amino acid sequence (primary) determines the three-dimensional (3D) structure (tertiary) of the protein, and its functions mainly depend on its 3D structure. Therefore, exploring the functions of different proteins' 3D structures is crucial for understanding their working mechanism, which benefits many applications, such as new drug development [1].

Fig. 1. (a-d) show a protein's amino acid sequence, surface, atoms, and cartoon structure (a simplified representation based on the secondary structure), respectively. (a) Due to the neighboring residue effect [2], multiscale relationships among atoms and between surface points and atoms should be considered in protein function analysis. (b-d) Our key idea is to model the hierarchical feature interactions between chemical (atom/residue) and geometric (surface) features for efficient protein surface learning.

Proteins carry out their functions by interacting with other proteins or molecules. Therefore, predicting protein-protein or protein-biomolecule interaction (e.g., interaction sites and interaction matching) becomes one of the most important and challenging problems in structural bioinformatics. Proteins often associate through hydrophobic patches on their surfaces. Amino acid residues in the interaction interfaces can give a general indication of hydrophobicity [3]. Various studies about the electrostatic nature of the protein-protein interface demonstrate that the associating surfaces interacting with one another have charge complementarity [4], [5] or electrostatic complementarity [6]–[8]. In addition, the size and shape of protein interfaces, measured with solvent accessible surface area (ΔASA) [9], varies in different types of protein complexes. Earlier work [3] shows that the range of ΔASA in the heterocomplexes is smaller than in the homodimers, and the larger molecules usually have larger interfaces. The interaction between proteins also relies on shape complementarity [10], such as convex bulges and concave pockets. Therefore, whether proteins interact with other molecules and where the interaction sites are located depends on both chemical features and 3D geometric shapes.

Existing research works on protein interaction prediction [8], [11]–[13] utilize handcrafted chemical features (e.g., hydrophobicity and charges) and adopt deep neural networks to learn geometric features (e.g., curvatures and normals) from 3D structures. Recently, dMaSIF [14] is proposed to

encode the atom type into a one-hot vector followed by multi-layer perceptrons (MLPs) to learn chemical features. Their experiments show that learned chemical features can perform as well as or better than handcrafted ones.

Nonetheless, dMaSIF [14] has two main limitations. Firstly, it does not consider the relationship among atoms. The chemical feature of each atom is learned individually only from the atom type. However, in a molecule, atoms are linked together by covalent bonds to form functional groups that cause chemical reactions of molecules [15]. This indicates that the chemical feature is not determined by a single atom, but by a group of atoms linked to each other. Secondly, it ignores hierarchical feature interactions between surface points and atoms. For each surface point, only the nearest 16 atoms around the surface point are gathered as the chemical feature. However, the neighboring residue effect [2] points out that amino acid residues can bring chemical shifts to their neighboring residues; see Figure 1.a. This reveals that the chemical properties of the protein surface are not only determined by atoms in the residues near the surface but also rely on atoms in some residues far away from the surface. The same goes for atoms in chemical feature learning.

To this end, we propose a novel **H**ierarchical **C**hemical and **G**eometric Feature Interaction **Net**work (**HCGNet**) to learn both chemical and geometric features in a hierarchical and interactive way. Our key idea is to capture the multiscale relationship among atoms to learn chemical features and hierarchically model their interactions with geometric features. Specifically, we introduce a dual hierarchical framework to extract chemical features and geometric features, respectively. By designing a chemical feature propagation module, the hierarchical chemical features from atoms (or residues) can be propagated to surface points for protein surface learning; see Figure 1.b-d. Our experiments show that HCGNet outperforms prior state-of-the-art (SoTA) methods in protein surface learning. To summarize, the main contributions are as follows:

- We highlight the importance of hierarchical relationships among atoms and between chemical and geometric features, which are new insights for protein surface learning and have been overlooked in previous works.
- We propose HCGNet to model hierarchical interactions in chemical feature learning and between chemical and geometric features for effective protein surface analysis.
- Our method outperforms prior SoTA method by 2.3% in site prediction task and 3.2% in interaction matching task. HCGNet is flexible and has the potential to be used in other protein-biomolecule interaction tasks, such as protein-ligand and protein-DNA/RNA.

II. RELATED WORK

Discovering the properties and functions of proteins at different structural levels will bring significant benefits to protein design and engineering. By leveraging large public protein databases, biological scientists are able to analyze proteins for a variety of downstream applications, for instance, protein structure prediction [16]–[19], protein-protein interaction [11], [13], [14], [20], [21], ligand binding affinity [22], [23], and protein docking [11], [24]–[26].

Feature learning on protein representation. The amino acid sequence is the simplest way to represent a protein. A straightforward way to process protein sequences is to directly apply techniques in natural language processing, including Word2Vec [27] and Doc2Vec [28]. Recently, UniRep [29] adopted multiplicative LSTMs [30] to learn features for each amino acid residue and averaged all residue features into a representative vector. However, sequence-based methods can not handle the detailed analysis of protein structures due to missing spatial information.

The three-dimensional (3D) structures of proteins are more important than sequences in protein function prediction for rich structural details. Atoms can be treated as points with chemical properties in 3D space. Compared with regular 1D sequences and 2D images, 3D representation learning has difficulties like spatial sparsity and disorder of indexing. Recent works [31], [32] are proposed to voxelize atom points into regular grids and adopt 3D convolution networks for analysis. However, voxel-based methods still suffer from high computational costs and limited resolution, resulting in insufficient analysis of structural details. Another effective way is regarding the 3D protein structure as a graph composed of vertices (atoms) and edges (chemical bonds). Fout A., *et al.* [33] handcraft features — distance and angle for edges, protrusion index, and residue depth for vertices, and utilize graph convolution networks, leading to better performance than SVM-based methods in interaction prediction. HOLO-PROT [22] jointly learns protein representation from both amino acid sequences and atom graphs to capture finer details.

In addition to amino acid sequences and atom graphs, researchers [11], [14] propose that surface can be a more natural way to analyze protein's functions because surfaces contain both chemical and geometric features, revealing information about protein interactions. Although they introduce chemical features to surface points in a handcrafting or learning way, neither of these works considers the neighboring residue effect and the multiscale relationship among atoms.

Deep learning on point clouds. Point clouds have become a more and more popular form of 3D data with the development of sensing devices and diverse learning techniques. To solve the problem of irregularity and disorder, Charles *et al.* [34] propose PointNet to apply shared MLPs on each point and aggregate features with a global max-pooling. PointNet can also be used as a local operator for hierarchical processing [35], and a lot of works [36]–[38] followed, developing various point convolution operators. Moreover, kNN graphs can be constructed from point clouds, and graph convolution networks can be applied for feature learning [39], [40]. Recent works expand a local patch of point clouds to an estimated [41] or a learned [42] 2D plane, and then regular 2D convolutions can be further adopted.

Although protein atoms can also be represented by a set of points, there are three main differences compared to general point clouds. *a.)* Points are regarded as homogeneous particles continuously distributed in the real world. Due to the limitation of sensing devices, points are collected discretely, with noise, and sometimes assigned with colors or reflection intensity.

(a) Protein   (b) Surface Points

(c) Atom Points   (d) Residue Points

Fig. 2. A protein can be represented by three pointsets, including surface points for geometric shape analysis, and atom and residue points for chemical property analysis.
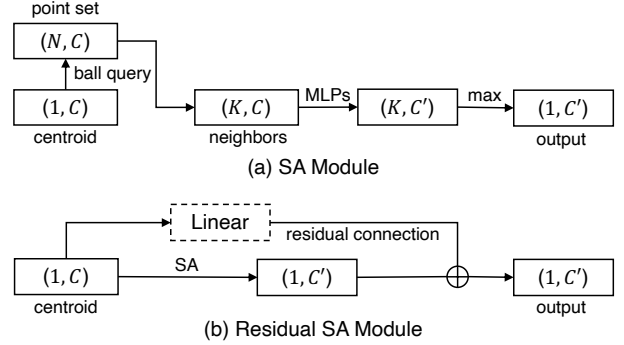

(a) SA Module

(b) Residual SA Module

Fig. 3. (a) In SA module, a set of neighbor points are grouped and followed by MLPs and max pooling for feature aggregation. (b) In residual SA module, SA is firstly used to extract the local feature for the centroid point, and then the residual connection is applied by adding input and output features. The linear layer is not necessary when C is equal to C'.

Atoms are indeed discretely distributed, free of noise, and have individual chemical properties and radii; *b.)* Points are captured from object surfaces, and there are almost no points inside, while atoms are distributed not only on protein surfaces but also inside the protein; *c.)* Geometric features play a decisive role in general point cloud analysis, while chemical features are as important as geometric features in protein structure analysis. Inspired by the above observations, we propose a novel learning architecture to bridge chemical and geometric features with hierarchical interactions, leading to SoTA performance in various protein analysis tasks.

III. OUR APPROACH

In this section, we first introduce protein data, then review some preliminaries about point cloud network design [35], and finally describe our proposed hierarchical chemical and geometric feature interaction network (HCGNet).

A. Protein Representation

Proteins are polymers composed of amino acid residues, which are formed by atoms. In the Protein Data Bank (PDB [43]), each residue is recorded with the amino acid type and a list of atoms, and each atom contains the spatial coordinates and atom type. In the preprocessing, each protein is generated into three pointsets, including atom, residue, and surface pointsets. A visual example is shown in Figure 2.

Atom pointset. We denote $\mathcal{A}$ as the atom pointset with M_1 points, where i^{th} atom point $(\mathbf{p}_{a_i}, \mathbf{f}_{a_i}) \in \mathbb{R}^3 \times \mathbb{R}^6$ consists of a 3-dim XYZ coordinate vector $\mathbf{p}_{a_i}$ to indicate the 3D spatial location, and a 6-dim one-hot vector $\mathbf{f}_{a_i}$ to represent the atom type (i.e., C, H, O, N, S, Se).

Residue pointset. Suppose that the protein is composed of M_2 amino acid residues, where i^{th} residue is formed by N_i atoms. Regarding the indices of atoms are denoted as $\{d_1, \ldots, d_{N_i}\}$, the center position of the residue is calculated by averaging

3-dim coordinates of all internal atoms:

$$\mathbf{p}_{r_i} = \frac{1}{N_i} \sum_{k=1}^{N_i} \mathbf{p}_{a_{d_k}}. \tag{1}$$

Hence, the residue pointset is represented by $\mathcal{R} = \left\{\mathbf{p}_{r_i}\right\}_{i=1}^{M_2} \subset \mathbb{R}^3$, which is used as the downsampled pointset to aggregate features from atom points for hierarchical chemical feature learning, which will be described in Section III-C.

Surface pointset. To perform effective learning on the protein surface, we follow [14] to generate $N (N \approx 2.5M_1)$ surface points from atoms with a fast sampling algorithm. Let $\mathcal{S}$ be the surface pointset, where i^{th} surface point $(\mathbf{p}_{s_i}, \mathbf{f}_{s_i}) \in \mathbb{R}^3 \times \mathbb{R}^{10}$ contains a 3-dim XYZ coordinate vector $\mathbf{p}_{s_i}$ and a 10-dim feature vector $\mathbf{f}_{s_i}$. $\mathbf{f}_{s_i}$ are Gaussian and mean curvatures estimated at 5 scales ranging from 1Å to 10Å.

B. Preliminaries of Point Cloud Processing

As mentioned, a protein can be represented by three pointsets, which can be processed in the way of point cloud learning. PointNet++ [35] is one of the most popular methods in point cloud processing and has been widely used in many computer vision tasks such as shape classification and scene understanding. It consists of two main modules, set abstraction (SA) module to encode features and downscale points and feature propagation (FP) module to upscale points and decode features. These two modules and the modified version are used as basic point (de)convolution operators in our framework.

Set abstraction (SA) module groups neighbor points around a centroid point and then uses a mini-PointNet [34] to encode local patterns into feature vectors. Let $\mathcal{P} = \left\{(\mathbf{p}_i, \mathbf{f}_i)\right\}_{i=1}^{N_p} \subset \mathbb{R}^3 \times \mathbb{R}^C$ be a pointset with N_p points and C refers to the dimension of the feature vector. Given a centroid point $(\mathbf{p}', \mathbf{f}')$, use ball query to group neighbor points' features as $\mathcal{N}(\mathbf{p}', \mathcal{P}) = \left\{[\mathbf{p} - \mathbf{p}'; \mathbf{f}] \mid \|\mathbf{p} - \mathbf{p}'\|_2 \le r, (\mathbf{p}, \mathbf{f}) \in \mathcal{P}\right\}$ with $r \in \mathbb{R}$ being the chosen radius and $[\cdot; \cdot]$ is the concatenation operator. Then we apply shared multi-layer perceptrons (MLPs) followed by a max pooling layer on $\mathcal{N}(\mathbf{p}', \mathcal{P})$ and
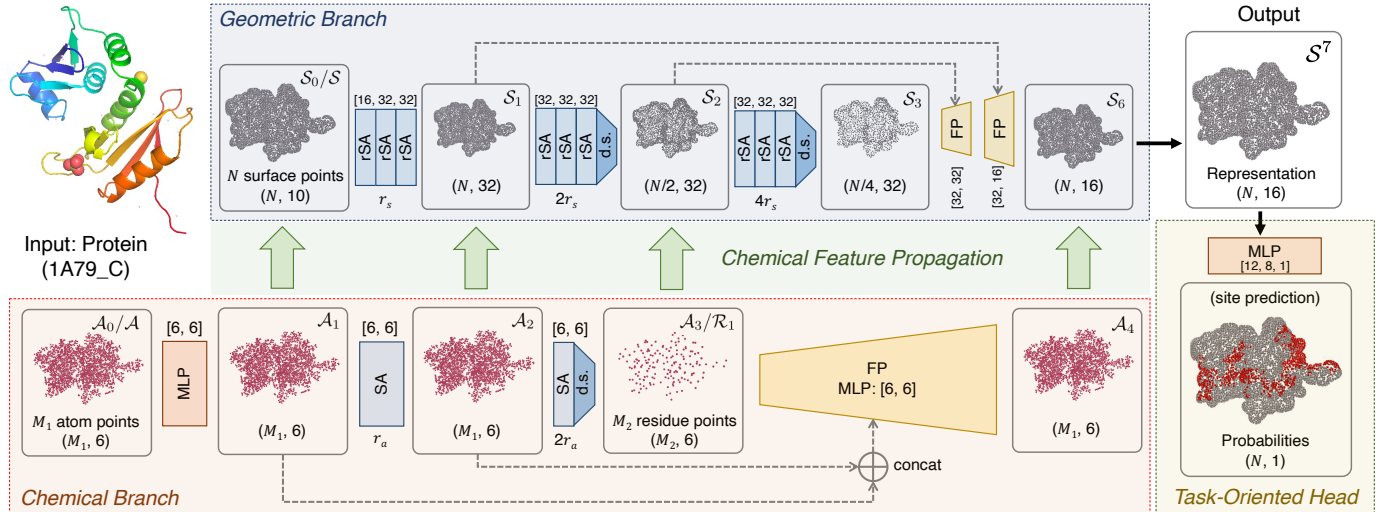
Fig. 4. The backbone of HCGNet is mainly composed of two branches for geometric (top) and chemical (bottom) feature learning. Two branches encode features from surface points and atom points in a multiscale way, respectively. Chemical feature propagation modules (middle) are proposed to propagate features from the chemical branch to the geometric branch also in a hierarchical way. Moreover, different task-oriented heads can be followed to handle different downstream tasks, such as site prediction and interaction matching. "d.s." indicates point cloud downsampling. r_s and r_a are the initial query radii for two branches, respectively.

output the local feature $\hat{\mathbf{f}}'$ at $\mathbf{p}'$:

$$
\begin{aligned}
\hat{\mathbf{f}}' &= \text{SA}(\mathbf{p}', \mathcal{P}) \\
&= \text{MAX}\left\{ \text{MLP}(\hat{\mathbf{f}}) \;\middle|\; \hat{\mathbf{f}} = [\mathbf{p} - \mathbf{p}'; \mathbf{f}] \in \mathcal{N}(\mathbf{p}', \mathcal{P}) \right\}.
\end{aligned} \tag{2}
$$

As shown in Figure 3, to build a deep learning network, we develop a residual version of SA (rSA) inspired by [44] by adding the original feature $\mathbf{f}'$ at $\mathbf{p}'$:

$$
\hat{\mathbf{f}}' = \text{rSA}(\mathbf{p}', \mathbf{f}', \mathcal{P}) = \text{SA}(\mathbf{p}', \mathcal{P}) + \mathbf{f}'. \tag{3}
$$

Feature propagation (FP) module. Points are downsampled in an encoder-decoder network, and FP module is used to recover the data size and propagate features from downsampled pointset $\mathcal{P}'$ to the original pointset $\mathcal{P}$. To be more specific, for each point $(\mathbf{p}, \mathbf{f}) \in \mathcal{P}$, we first select K nearest neighbor points from $\mathcal{P}'$ as $\mathcal{N}_K(\mathbf{p}, \mathcal{P}') = \left\{ (\mathbf{p}'_k, \mathbf{f}'_k) \right\}_{k=1}^{K}$. Then we interpolate feature values using their relative distances:

$$
\mathbf{f}_{\text{interp}} = \text{Interp}(\mathbf{p}, \mathcal{N}_K(\mathbf{p}, \mathcal{P}')) = \frac{\sum_k w_k \mathbf{f}'_k}{\sum_k w_k}, \tag{4}
$$

where $w_k = 1/\|\mathbf{p}'_k - \mathbf{p}\|_2 \in \mathbb{R}$ is the interpolation weight. Furthermore, we adopt skip connection to combine propagated features $\mathbf{f}_{\text{interp}}$ and original features $\mathbf{f}$ together and use a shared MLP to update each point's feature:

$$
\hat{\mathbf{f}} = \text{FP}(\mathbf{p}, \mathbf{f}, \mathcal{P}') = \text{MLP}\left([\mathbf{f}_{\text{interp}}; \mathbf{f}]\right). \tag{5}
$$

C. Hierarchical Feature Interaction Network

As addressed in Section I, the multiscale relationship among atoms and hierarchical feature interactions between surface points and atoms are significant and overlooked by previous works. In this section, we formally propose HCGNet to model the above hierarchical relationships. An overview of the network architecture is shown in Figure 4. Given a protein represented by three pointsets (atom, residue, and surface), we

feed surface points into the geometric branch to encode hierarchical geometry-related features and atom/residue points into the chemical branch to encode hierarchical chemistry-related features. More importantly, the chemical feature propagation module is introduced in a hierarchical way to enhance the feature interaction between two branches to improve feature representation learning. Finally, we obtain the representative features of each surface point and design task-oriented heads for different protein surface learning tasks.

Geometric branch is composed of three encoding layers and two decoding (upsampling) layers to learn geometric features from surface points in a hierarchical way, as shown at the top of Figure 4. To formulate, in i^{th} encoding layer, assume the input is $\mathcal{S}^i \in \mathbb{R}^{N_i \times C_i}$ with N_i points and a channel size of C_i. A stack of rSA modules (see Section III-B) are applied for local feature encoding. The channels of MLP in k^{th} rSA module are $[C_i^{k-1}, C_i^k, C_i^k, C_i^k]$, where C_i^{k-1} and C_i^k are the output channel sizes of $k-1^{\text{th}}$ and k^{th} rSA modules respectively, and $C_i^0 = C_i$. All rSA modules in i^{th} encoding layer have the same neighbor query radius r_i, which is twice the radius r_{i-1} of the previous encoding layer (i.e., $r_i = 2r_{i-1}$). Let $r_1 = r_s$ be the initial radius. In some encoding layers, farthest point sampling (FPS [45]) is adopted following the final rSA module to downsample half of surface points for efficient learning. In decoding layers, FP modules are used to propagate features from downsampled pointset to the original pointset layer by layer. Skip connections are used to combine low-level and high-level features to obtain a better representation. The initial radius r_s is set to 4Å in site prediction and 2Å in interaction matching. The output channel size of each module is given in Figure 4.

Chemical branch is composed of three encoding layers and one decoding layer for multiscale chemical feature learning, as shown at the bottom of Figure 4. Different from the geometric branch, the first encoding layer ($\mathcal{A}_1$) is a shared MLP to

Fig. 5. Implementation details of feature aggregation function $g(\cdot)$ in CFP. $g(\cdot)$ operates on each surface point to query chemical features from neighbor atoms (near) and residues (far). We use MLPs to transform features and summation to aggregate features from neighbors.

encode atom features individually from one-hot atom type vectors. In the rest of encoding layers ($\mathcal{A}_2$ and $\mathcal{R}_1$), the neighbor query radii are r_a and $2r_a$, where r_a is the initial radius and set to 4Å in site prediction and 3Å in interaction matching. We use SA modules without the residual connection for local feature encoding since the network is not as deep as the geometric branch. Instead of using FPS to downsample the pointset, residue points $\mathcal{R} = \{\mathbf{p}_{r_i}\}_i$ are used as centroids to group neighbor atom points and obtain $\mathcal{R}_1 = \{(\mathbf{p}_{r_i}, \mathbf{f}^1_{r_i})\}_i$:

$$\mathbf{f}^1_{r_i} = \mathrm{SA}(\mathbf{p}_{r_i}, \mathcal{A}_2), \tag{6}$$

where $\mathcal{A}_2$ are output atom points of the first SA module and $\mathbf{f}^1_{r_i}$ is the aggregated feature vector at residue $\mathbf{p}_{r_i}$. Then the FP module and skip connection are used to combine all levels of chemical features in the upsampling layer ($\mathcal{A}_4$). Hence, we have chemical features in four levels: individual, local, global, and mixed, from $\mathcal{A}_1$ to $\mathcal{A}_4$ (let $\mathcal{A}_3 = \mathcal{R}_1$).

Chemical feature propagation (CFP) module is introduced to propagate features from chemical branch to geometric branch. More specifically, to propagate features from $\mathcal{A}$ to each surface point $(\mathbf{p}_s, \mathbf{f}_s) \in \mathcal{S}$ (subscripts of $\mathcal{A}$ and $\mathcal{S}$ are ignored for simplification), we denote

$$\begin{aligned}\hat{\mathbf{f}}_s &= \mathrm{CFP}(\mathbf{p}_s, \mathbf{f}_s, \mathcal{A}) = h(\mathbf{c}_s, \mathbf{f}_s), \text{ and} \\ \mathbf{c}_s &= g\left(\mathcal{N}_{k_1}(\mathbf{p}_s, \mathcal{A}), \mathcal{N}_{k_2}(\mathbf{p}_s, \mathcal{A}')\right),\end{aligned} \tag{7}$$

where $\mathcal{N}_{k_1/k_2}(\cdot)$ are kNN selection functions with k equal to k_1 and k_2, respectively. $\mathcal{A}'$ is the sparse pointset sub-sampled from $\mathcal{A}$ and used for querying far away chemical points. $g(\cdot)$ is a feature aggregation function operated on chemical pointsets, $\mathbf{c}_s$ is the propagated chemical feature, $h(\cdot)$ is a multi-modality fusion function to combine chemical and geometric features, and $\hat{\mathbf{f}}_s$ is the fused feature. In practice, we regard residue points as the sub-sampled centroids and assign each residue feature by the feature of the nearest point in $\mathcal{A}$ to form $\mathcal{A}'$. If the chemical source is the residue set (e.g., $\mathcal{A}_3 \rightarrow \mathcal{S}_2$), then $\mathcal{A}' = \mathcal{A}$. The detailed design of $g(\cdot)$ is shown in Figure 5, and we directly concatenate the propagated chemical features and the original geometric features together:

$$\hat{\mathbf{f}}_s = h(\mathbf{c}_s, \mathbf{f}_s) = [\mathbf{c}_s; \mathbf{f}_s]. \tag{8}$$

As shown in the middle of Figure 4, we adopt CFP modules to bridge the chemical points (atoms or residues) and surface points that have similar receptive fields to establish hierarchical interactions between chemical and geometric branches.

D. Task-Oriented Heads for Downstream Tasks

The outputs of the backbone are representative point-wise features on the protein surface, which can be formulated as

$$\mathcal{S}^7 = \left\{(\mathbf{p}_{s_i}, \mathbf{f}^7_{s_i})\right\}_{i=1}^N = \mathrm{CFP}(\mathcal{S}_6, \mathcal{A}_4), \tag{9}$$

where $\mathcal{S}_6$ and $\mathcal{A}_4$ are the output pointsets of geometric and chemical branches, respectively. Hence we can design task-oriented heads to handle different tasks. In this work, two tasks are included, interface site prediction and interaction matching [11], [14].

Site prediction aims to identify interaction sites and non-interaction sites from surface points. As shown in Figure 4, a three-layer MLP is followed to predict the probability of being an interaction site for each surface point $(\mathbf{p}_s, \mathbf{f}^7_s) \in \mathcal{S}^7$. The full model is end-to-end optimized with binary cross entropy (BCE) loss and Dice loss [46]:

$$\begin{aligned}\mathcal{L}(\hat{y}_s, y_s) &= \lambda \mathcal{L}_{\mathrm{BCE}}(\hat{y}_s, y_s) + \\ &\quad (1-\lambda)\mathcal{L}_{\mathrm{Dice}}(\hat{y}_s, y_s), \text{ and} \\ \hat{y}_s &= \mathrm{MLP}(\mathbf{f}^7_s),\end{aligned} \tag{10}$$

where $\lambda \in [0, 1]$ is trade-off coefficient, $\hat{y}_s \in \mathbb{R}$ is final predicted probability, and $y_s \in \{0, 1\}$ is ground truth label.

Interaction matching aims to predict the interaction probability of each surface point pair, one from each protein involved in a complex, which is key to protein docking. Because we are about to measure feature complementarity instead of similarity, as shown in Figure 6.a, two branches are used to predict the representative feature (binder) and the corresponding complemental feature (target), respectively. To formulate, for each surface point $(\mathbf{p}_s, \mathbf{f}^7_s) \in \mathcal{S}^7$, the output binder feature $\mathbf{b}$ and target feature $\mathbf{t}$ are

$$\mathbf{b} = \mathrm{MLP}_b(\mathbf{f}^7_s) \text{ and } \mathbf{t} = \mathrm{MLP}_t(\mathbf{f}^7_s), \tag{11}$$

where MLP_b and MLP_t are MLP layers of binder and target branches. During training, the input is a pair of proteins (denoted as A and B, both are randomly rotated) from the same complex. The backbone network first produces binder and target features for the surface points of each protein. Given point a from protein A with output features $(\mathbf{t}_a, \mathbf{b}_a)$ and point b from protein B with $(\mathbf{t}_b, \mathbf{b}_b)$, the interaction probability of these two points is formulated as

$$p(a, b) = \mathrm{Sigmoid}\left(\frac{\mathbf{t}_a^T \mathbf{b}_b + \mathbf{t}_b^T \mathbf{b}_a}{2}\right). \tag{12}$$

If the relative distance between point a and point b is less than 1Å before the random rotation of proteins, the point pair (a, b) is labeled as a positive pair; otherwise, it is labeled as a negative pair. Therefore, the objective function is defined as

$$\mathcal{L}(a, b, y_{ab}) = \mathcal{L}_{\mathrm{BCE}}(p(a, b), y_{ab}), \tag{13}$$

where $y_{ab} \in \{0, 1\}$ is 1 if and only if (a, b) is a positive pair. All positive pairs are selected, and the same number of negative pairs are randomly sampled for training. Due to the memory limitation, gradients are only recorded during the forward pass of the first protein (i.e., *1A79_C* in Figure 6.a).

(a) Pipeline of Interaction Matching

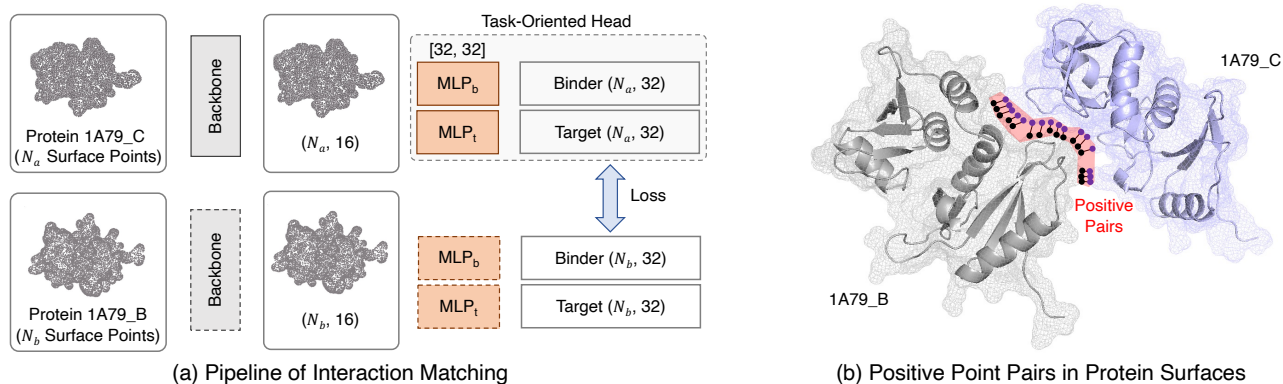(b) Positive Point Pairs in Protein Surfaces

Fig. 6. (a) Illustration of interaction matching network. The backbone network is shown in Figure 4. The task-oriented head comprises two MLPs to predict the representative feature (binder) and corresponding complementary feature (target). (b) An example: point pairs (colored in red) located very close in the formed complex are labeled as positive pairs.

IV. EXPERIMENTS

To validate the effectiveness of the proposed HCGNet, we conduct experiments on two tasks in the field of structural bioinformatics, which are introduced in [11] and aim to deal with interactions between proteins. We implement our method with PyTorch [47]. Momentum gradient descent optimizer is used to optimize loss function in Eqn. 10 and Eqn. 13 with an initial learning rate of 0.01 and a batch size of 8. Leaky ReLU and instance normalization are applied in MLP layers. Models are trained on a single NVIDIA GeForce RTX 3090 GPU for 300 epochs ($\sim$24h) in site prediction and 150 epochs ($\sim$36h) in interaction matching.

A. Site Prediction

Setting. In this task, the surface points of a protein are classified into interaction sites and non-interaction sites. Interaction sites are surface points that are more likely to interact with other proteins. This task can benefit protein engineering, such as drug discovery. We use the dataset proposed in [11], which contains 3,314 proteins collected from the Protein Data Bank [43] with 2,958 proteins for training and 356 for testing. We follow [11] to split 10% data from the training set for validation during training based on the pairwise matrix of TM-scores (refer to Supplementary Note 4 in [11] for more details). λ in Eqn. 10 is set to 0.5. To quantitatively compare our proposed HCGNet with previous works, we follow [14] to report ROC-AUC (area under the ROC curve) as the evaluation metric, since it measures the ability to distinguish between classes and is insensitive to class imbalance.

Results. As shown in Table I, HCGNet outperforms the SoTA method (dMaSIF [14]; see also Figure 8.a) by 2.3% ROC-AUC and is significantly ahead of other previous methods. In addition, we visualize the results in Figure 7.a. Compared with previous methods, our predicted interaction sites are more precise with the help of hierarchical chemical features.

B. Interaction Matching

Setting. In this task, two proteins involved in a complex are given to predict the interaction probability of point pairs (one

TABLE I
AREA UNDER THE ROC CURVE (ROC-AUC) IS EVALUATED TO COMPARE OUR PROPOSED ARCHITECTURE WITH PREVIOUS METHODS ON SITE PREDICTION AND INTERACTION MATCHING.

Method	Site Prediction	Interaction Matching
DGCNN [40]	0.710	-
PointNet [34]	0.816	0.747
PointNet++ [35]	0.848	0.785
PointConv [38]	0.855	0.790
MaSIF [11]	0.850	0.787
dMaSIF [14]	0.870	0.794
HCGNet (*ours*)	**0.893**	**0.826**

from each). This task is the key to protein docking, which aims to align the orientation of two proteins in a complex. We use the dataset proposed in [11], including 4,614 protein pairs for training and 912 for testing. 10% training data are used for validation based structural alignments (refer to Supplementary Note 6 in [11] for more details). In the evaluation stage, all positive pairs are selected, and the same number of negative pairs are randomly sampled for ROC-AUC calculation.

Results. As quantitative comparison shown in Table I, our HCGNet performs better than the SoTA (dMaSIF [14]; see also Figure 8.b) and all other previous methods with a remarkable margin. Note that dMaSIF [14] reported a ROC-AUC of 0.82 on interaction matching by selecting all positive pairs and 400 negative pairs for evaluation. Based on this evaluation strategy, our performance (0.837) is still better. In Figure 7.b, we highlight the matched points in red. Visual results show that our predicted pairs are more concentrated in the associating areas. Compared to site prediction, our proposed architecture has achieved more improvements in interaction matching (see also the second ablation study). The reason may be that interaction matching is aware of binding partners, and hierarchical chemical features can help the network to identify specific chemical complementarities such as alkaline and acidic, positively charged, and negatively charged. However, site prediction aims to predict general interaction sites where the functional group is active and the

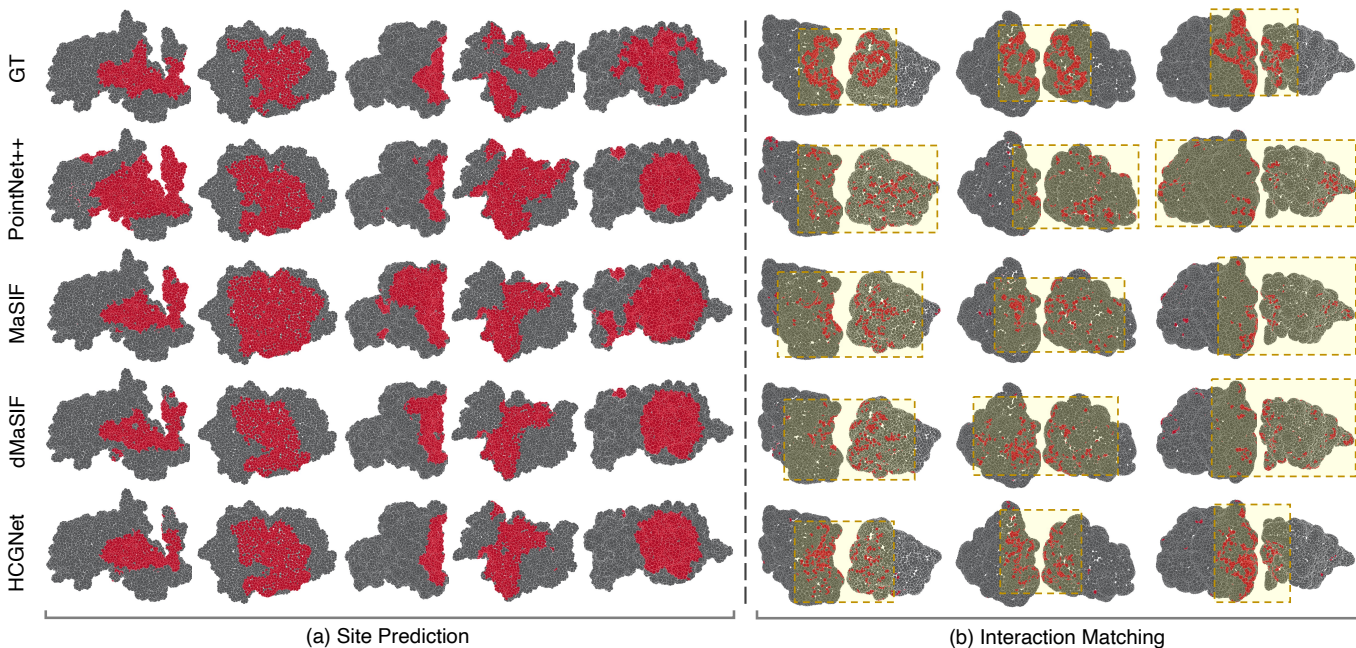|	(a) Site Prediction	(b) Interaction Matching

Fig. 7. Qualitative results on site prediction and interaction matching tasks. (a) Our HCGNet can perform more precise segmentation of interaction sites (red) than previous methods. (b) Matched points are highlighted in red. Compared with previous methods, our predicted point pairs can be more helpful for the orientation alignment since most matched points are located in associating areas; see bounded areas in yellow.

TABLE II

THE IMPORTANCE OF CHEMICAL AND GEOMETRIC FEATURES. ROC-AUC IS EVALUATED AND THE PERFORMANCE DROPS SHARPLY WHEN TRAINING WITH ONLY CHEMICAL (CHEM.) FEATURES OR ONLY GEOMETRIC (GEOM.) FEATURES IN BOTH TWO TASKS.

Method	Site Prediction	Interaction Matching
chem. only	0.673	0.692
geom. only	0.759	0.722
chem. & geom.	**0.893**	**0.826**

TABLE III

THE IMPORTANCE OF CHEMICAL FEATURE LEARNING AND HIERARCHICAL FEATURE INTERACTION. ROC-AUC IS EVALUATED FOR COMPARING THE FULL MODEL WITH REMOVING DIFFERENT PROPOSED PARTS, INCLUDING CFP, HIERARCHICAL CHEMICAL (CHEM.) FEATURE LEARNING, AND HIERARCHICAL FEATURE INTERACTION (INTER.) BETWEEN TWO BRANCHES.

Method	Site Prediction	Interaction Matching
w/o CFP	0.874 -1.9%	0.803 -2.3%
w/o hierarchical chem.	0.884 -0.9%	0.813 -1.3%
w/o hierarchical inter.	0.886 -0.7%	0.820 -0.6%
Full model (HCGNet)	**0.893**	**0.826**

geometric shape is obviously concave or convex.

C. Ablation Study

To further analyze our proposed architecture, ablative experiments are conducted to explore the importance of chemical and geometric features, the importance of different proposed parts, and different designs of hierarchical learning networks and hierarchical feature interaction.

Importance of chemical and geometric features. To illustrate the significance of chemical and geometric features, we design two backbones: *chem. only* to use only the chemical branch and propagate features in the final layer to surface points for point-wise classification, and *geom. only* to use only the geometric branch without any chemical feature propagation. Experiments are conducted on the tasks of site prediction and interaction matching. As shown in Table II, in the absence of either chemical features or geometric features, the performance drops dramatically in both two tasks, which validates the importance of the interaction between chemical features and geometric features.

Importance of different proposed parts. To demonstrate the effectiveness of different proposed parts, we conduct ablative experiments by removing CFP, hierarchical chemical feature learning, and hierarchical interaction between two branches, respectively. In Table III, the performance drops without any of the above parts, which implies the significance of modeling multiscale relationships among atoms and hierarchical feature interaction between surface points and atoms (or residues). In addition, we notice that by removing hierarchical chemical learning and CFP, the performance drops for interaction matching (-1.3%/-2.3%) are more than site prediction (-0.9%/-1.9%) since awareness of binding partner requires more precise chemical features to predict chemical complementarity.

Design on hierarchical branches. As mentioned, hierarchical analysis is significant for both chemical and geometric learning. We compare different radii of ball query in SA (or rSA) modules from 2Å to 7Å in two tasks, as shown in Table IV. We find that the best setting has similar initial radii in two

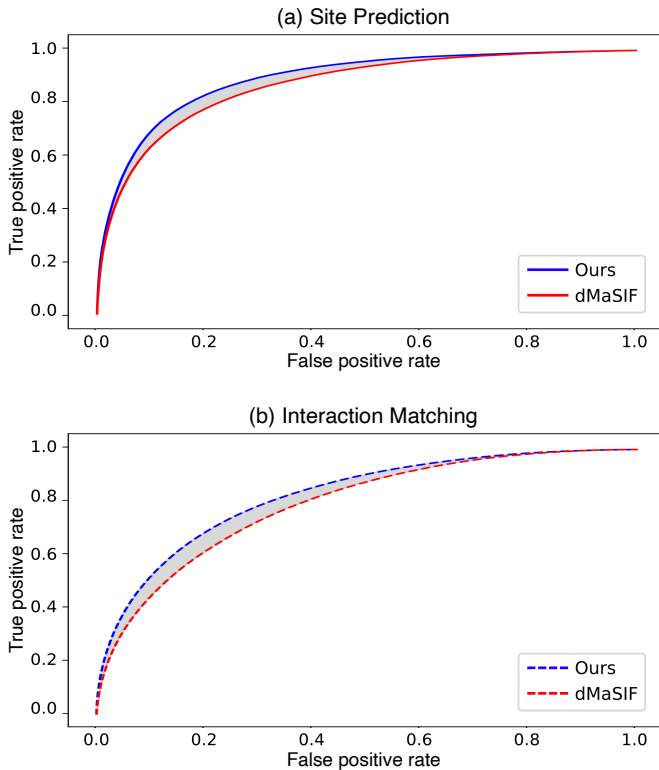(a) Site Prediction

(b) Interaction Matching

Fig. 8. ROC curves comparing the performance of our method (blue) and dMaSIF (red) [14] on the task of site prediction (a) and interaction matching (b).

TABLE IV

COMPARISON OF DIFFERENT INITIAL RADII r_s IN THE GEOMETRIC BRANCH AND r_a THE CHEMICAL BRANCH. ROC-AUC IS EVALUATED, AND THE BEST SETTING HAS SIMILAR r_s AND r_a IN BOTH TWO TASKS.

r_s (Å)	r_a (Å)	Site Prediction	Interaction Matching
2	3	0.889	**0.826**
4	2	0.888	0.814
4	4	**0.893**	0.825
4	7	0.887	0.821
7	4	0.883	0.813

learning branches, $(r_s, r_a) = (4, 4)$Å in site prediction and $(r_s, r_a) = (2, 3)$Å in interaction matching.

Design on hierarchical feature interaction. In HCGNet, the feature propagation from the chemical branch to the geometric branch is implemented as $(\mathcal{A}_1, \mathcal{A}_2, \mathcal{A}_3, \mathcal{A}_4) \rightarrow (\mathcal{S}_0, \mathcal{S}_1, \mathcal{S}_2, \mathcal{S}_6)$, which indicates $\mathcal{A}_1 \rightarrow \mathcal{S}_0$, $\mathcal{A}_2 \rightarrow \mathcal{S}_1$, $\mathcal{A}_3 \rightarrow \mathcal{S}_2$, and $\mathcal{A}_4 \rightarrow \mathcal{S}_6$. In Table V, we compare different chemical sources for the propagation process. Although the feature propagation between two branches can be arbitrary, we still recommend establishing the interaction in the way of Figure 4, because the linked surface points and chemical points (atoms or residues) have similar receptive fields. If not, the network may be confused by inconsistent receptive fields during training and bring worse performance.

TABLE V

COMPARISON OF DIFFERENT PROPAGATION STRATEGIES. ROC-AUC IS EVALUATED, AND THE CHEMICAL SOURCES REPRESENT THE ATOM POINTSET PROPAGATED TO $\mathcal{S}_0$, $\mathcal{S}_1$, $\mathcal{S}_2$ AND $\mathcal{S}_6$, RESPECTIVELY. - INDICATES NO CHEMICAL FEATURE PROPAGATION TO THIS SURFACE POINTSET.

Chemical Sources	Site Prediction	Interaction Matching
$(\mathcal{A}_1, -, -, -)$	0.884	0.813
$(\mathcal{A}_4, \mathcal{A}_3, \mathcal{A}_2, \mathcal{A}_1)$	0.888	0.818
$(\mathcal{A}_1, \mathcal{A}_3, \mathcal{A}_3, \mathcal{A}_4)$	0.887	0.820
$(\mathcal{A}_1, \mathcal{A}_2, \mathcal{A}_3, \mathcal{A}_4)$	**0.893**	**0.826**

TABLE VI

COMPARISON OF MODEL PARAMETERS. WE COMPARE THE MODEL PARAMETERS OF TRAINING ONLY WITH CHEMICAL FEATURES (CHEM. ONLY), TRAINING ONLY WITH GEOMETRIC FEATURES (GEOM. ONLY), REMOVING CFP (W/O CFP), REMOVING HIERARCHICAL CHEMICAL LEARNING (W/O HIERARCHICAL CHEM.), AND REMOVING HIERARCHICAL FEATURE INTERACTION (W/O HIERARCHICAL INTER.).

Method	Site Prediction		Interaction Matching	
chem. only	0.673	0.67K	0.692	5.08K
geom. only	0.759	24.4K	0.722	29.3K
w/o CFP	0.874	27.6K	0.803	32.9K
w/o hierarchical chem.	0.884	27.3K	0.813	32.5K
w/o hierarchical inter.	0.886	24.9K	0.820	29.9K
Full model (HCGNet)	**0.893**	28.3K	**0.826**	33.6K

V. DISCUSSION

A. Biomolecular Surface Learning

Biomolecules, including proteins, are composed of atoms and most chemical interactions occur on biomolecular surfaces. This means the surface features (e.g., chemistry, geometry) play an important role in interaction-related tasks. In this work, we point out that the relationship among atoms and the feature propagation from atoms to surfaces are significant for protein surface learning, which should be the same for other biomolecules' surface learning. More importantly, we propose a general framework HCGNet mainly focusing on the design of chemical learning and feature propagation. HCGNet can handle different downstream tasks, such as site prediction and interaction matching, by equipping different task-oriented heads. We believe that HCGNet can also handle tasks in other biomolecules, such as DNA/RNA and ligands, with only minor modifications to the proposed framework. The remaining difficulties are the task definition and dataset construction. We look forward to and will work towards the unification of frameworks for biomolecular surface learning.

B. Processing Efficiency

Compared with dMaSIF [14], HCGNet has more parameters (28.3K vs. 2.8K) and requires more GPU memories (407MB/protein vs. 132MB/protein) due to the design of deep network. However, due to the hierarchical design, the inference time of HCGNet (17ms) is comparable with dMaSIF [14]

(16ms), and much less than other compared methods[1], including DGCNN [40] (50ms), PointNet [34] (130ms), PointNet++ [35] (410ms), PointConv [38] (45ms), and MaSIF [11] (180ms). In addition, we should point out that 28.3K is not a big number in deep learning networks (e.g., ∼12M in ResNet18), and our proposed CFP, hierarchical chemical learning and interaction indeed do not introduce many training parameters but bring considerable improvements; see Section V-C. All the above results are tested under the same environment setting of Ubuntu 20.04, RTX 3090, CUDA 11.1, and PyTorch 1.8.

C. Model Parameters

In Table VI, we compare the performance and training parameters of networks in the ablative experiments, which shows that introducing CFP, hierarchical chemical learning and interaction can bring significant improvements (1.9%/0.9%/0.7%; 2.3%/1.3%/0.6%) with few increments of training parameters (0.7K/1.0K/3.4K; 0.7K/1.1K/3.7K) in both two tasks.

D. Scalability

The proposed framework is scalable to some extent. When processing larger molecules or more complicated downstream tasks, it would be better to increase the model capability by increasing the depth/width of chemical and geometric branches. However, increased computational requirements and overfitting could be possible issues when doing so.

VI. CONCLUSION

In this work, we highlight the importance of the multiscale relationship between atoms and the hierarchical interaction between chemical and geometric features. To this end, we propose HCGNet, a novel learning architecture for protein surface analysis. HCGNet takes atoms and surface points of a given protein as the input. Then two hierarchical branches are used to learn chemical features from atoms and geometric features from surface points in parallel. In addition, features are hierarchically propagated from the chemical branch to the geometric branch for multi-modality feature fusion. Our experiments demonstrate that HCGNet significantly improves the performance over the SoTA method in two challenging protein analysis tasks: site prediction and interaction matching.

Limitations and future work. Although HCGNet is a general framework for joint learning from chemical and geometric features, the hyperparameters and the model design should be specifically tuned for different downstream applications. Hence, it is more important to develop a unified framework and collect numerous data for general biomolecular surface learning, which will left as our future work.

---

[1]For the inference time of DGCNN, PointNet++, MaSIF, and dMaSIF, we report the results presented in [14].

REFERENCES

[1] C. Shen, J. Ding, Z. Wang, D. Cao, X. Ding, and T. Hou, "From machine learning to deep learning: Advances in scoring functions for protein–ligand docking," *Wiley Interdisciplinary Reviews: Computational Molecular Science*, vol. 10, no. 1, p. e1429, 2020.

[2] Y. Wang and O. Jardetzky, "Investigation of the neighboring residue effects on protein chemical shifts," *Journal of the American Chemical Society*, vol. 124, no. 47, pp. 14 075–14 084, 2002.

[3] S. Jones and J. M. Thornton, "Principles of protein-protein interactions," *Proceedings of the National Academy of Sciences*, vol. 93, no. 1, pp. 13–20, 1996.

[4] J. Novotny and K. Sharp, "Electrostatic fields in antibodies and antibody/antigen complexes," *Progress in biophysics and molecular biology*, vol. 58, no. 3, pp. 203–224, 1992.

[5] V. A. Roberts, H. Freeman, A. Olson, J. Tainer, and E. Getzoff, "Electrostatic orientation of the electron-transfer complex between plastocyanin and cytochrome c," *Journal of Biological Chemistry*, vol. 266, no. 20, pp. 13 431–13 441, 1991.

[6] B. C. Braden and R. J. Poljak, "Structural features of the reactions between antibodies and protein antigens," *The FASEB Journal*, vol. 9, no. 1, pp. 9–16, 1995.

[7] E. Demchuk, T. Mueller, H. Oschkinat, W. Sebald, and R. C. Wade, "Receptor binding properties of four-helix-bundle growth factors deduced from electrostatic analysis," *Protein Science*, vol. 3, no. 6, pp. 920–935, 1994.

[8] I. A. Vakser and C. Aflalo, "Hydrophobic docking: a proposed enhancement to molecular recognition techniques," *Proteins: Structure, Function, and Bioinformatics*, vol. 20, no. 4, pp. 320–329, 1994.

[9] B. Lee and F. M. Richards, "The interpretation of protein structures: estimation of static accessibility," *Journal of molecular biology*, vol. 55, no. 3, pp. 379–IN4, 1971.

[10] M. C. Lawrence and P. M. Colman, "Shape complementarity at protein/protein interfaces," 1993.

[11] P. Gainza, F. Sverrisson, F. Monti, E. Rodolà, D. Boscaini, M. Bronstein, and B. Correia, "Deciphering interaction fingerprints from protein molecular surfaces using geometric deep learning," *Nature Methods*, vol. 17, no. 2, pp. 184–192, 2020.

[12] G. Mai, K. Janowicz, B. Yan, R. Zhu, L. Cai, and N. Lao, "Multi-scale representation learning for spatial feature distributions using grid cells," 2020.

[13] N. B. Rego, E. Xi, and A. J. Patel, "Identifying hydrophobic protein patches to inform protein interaction interfaces," *Proceedings of the National Academy of Sciences*, vol. 118, no. 6, 2021.

[14] F. Sverrisson, J. Feydy, B. E. Correia, and M. M. Bronstein, "Fast end-to-end learning on protein surfaces," in *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2021.

[15] A. D. McNaught, A. Wilkinson *et al.*, *Compendium of chemical terminology*. Blackwell Science Oxford, 1997, vol. 1669.

[16] J. Jumper, R. Evans, A. Pritzel, T. Green, M. Figurnov, O. Ronneberger, K. Tunyasuvunakool, R. Bates, A. Žídek, A. Potapenko, A. Bridgland, C. Meyer, S. A. A. Kohl, A. J. Ballard, A. Cowie, B. Romera-Paredes, S. Nikolov, R. Jain, J. Adler, T. Back, S. Petersen, D. Reiman, E. Clancy, M. Zielinski, M. Steinegger, M. Pacholska, T. Berghammer, S. Bodenstein, D. Silver, O. Vinyals, A. W. Senior, K. Kavukcuoglu, P. Kohli, and D. Hassabis, "Highly accurate protein structure prediction with AlphaFold," *Nature*, vol. 596, no. 7873, pp. 583–589, 2021.

[17] A. W. Senior, R. Evans, J. Jumper, J. Kirkpatrick, L. Sifre, T. Green, C. Qin, A. Žídek, A. W. Nelson, A. Bridgland *et al.*, "Protein structure prediction using multiple deep neural networks in the 13th critical assessment of protein structure prediction (casp13)," *Proteins: Structure, Function, and Bioinformatics*, vol. 87, no. 12, pp. 1141–1148, 2019.

[18] A. W. Senior, R. Evans, Jumper *et al.*, "Improved protein structure prediction using potentials from deep learning," *Nature*, vol. 577, no. 7792, pp. 706–710, 2020.

[19] J. Xu, "Distance-based protein folding powered by deep learning," *Proceedings of the National Academy of Sciences*, vol. 116, no. 34, pp. 16 856–16 865, 2019.

[20] B. Dai and C. Bailey-Kellogg, "Protein interaction interface region prediction by geometric deep learning," *Bioinformatics*, 2021.

[21] A. R. Jamasb, B. Day, C. Cangea, P. Liò, and T. L. Blundell, "Deep learning for protein–protein interaction site prediction," in *Proteomics Data Analysis*. Springer, 2021, pp. 263–288.

[22] V. R. Somnath, C. Bunne, and A. Krause, "Multi-scale representation learning on proteins," *Advances in Neural Information Processing Systems*, vol. 34, 2021.

[23] G. Xiong, C. Shen, Z. Yang, D. Jiang, S. Liu, A. Lu, X. Chen, T. Hou, and D. Cao, "Featurization strategies for protein–ligand interactions and their applications in scoring function development," *Wiley Interdisciplinary Reviews: Computational Molecular Science*, p. e1567, 2021.

[24] B. G. Pierce, Y. Hourai, and Z. Weng, "Accelerating protein docking in zdock using an advanced 3d convolution library," *PloS one*, vol. 6, no. 9, p. e24657, 2011.

[25] N. Renaud, C. Geng, S. Georgievska, F. Ambrosetti, L. Ridder, D. F. Marzella, A. M. Bonvin, and L. C. Xue, "Deeprank: A deep learning framework for data mining 3d protein-protein interfaces," *Biorxiv*, 2021.

[26] D. Schneidman-Duhovny, Y. Inbar, R. Nussinov, and H. J. Wolfson, "Patchdock and symmdock: servers for rigid and symmetric docking," *Nucleic acids research*, vol. 33, no. suppl_2, pp. W363–W367, 2005.

[27] T. Mikolov, K. Chen, G. Corrado, and J. Dean, "Efficient estimation of word representations in vector space," *arXiv preprint arXiv:1301.3781*, 2013.

[28] Q. Le and T. Mikolov, "Distributed representations of sentences and documents," in *International conference on machine learning*. PMLR, 2014, pp. 1188–1196.

[29] E. C. Alley, G. Khimulya, S. Biswas, M. AlQuraishi, and G. M. Church, "Unified rational protein engineering with sequence-based deep representation learning," *Nature methods*, vol. 16, no. 12, pp. 1315–1322, 2019.

[30] B. Krause, L. Lu, I. Murray, and S. Renals, "Multiplicative lstm for sequence modelling," 2017.

[31] M. Simonovsky and J. Meyers, "Deeplytough: learning structural comparison of protein binding sites," *Journal of chemical information and modeling*, vol. 60, no. 4, pp. 2356–2366, 2020.

[32] R. J. L. Townshend, R. Bedi, P. A. Suriana, and R. O. Dror, "End-to-end learning on 3d protein structure for interface prediction," 2019.

[33] F. A., B. J., S. B., and B.-H. A., "Protein interface prediction using graph convolutional networks," in *Advances in Neural Information Processing System*, 2017.

[34] C. R. Qi, H. Su, K. Mo, and L. J. Guibas, "Pointnet: Deep learning on point sets for 3d classification and segmentation," *arXiv preprint arXiv:1612.00593*, 2016.

[35] C. R. Qi, L. Yi, H. Su, and L. J. Guibas, "Pointnet++: Deep hierarchical feature learning on point sets in a metric space," *arXiv preprint arXiv:1706.02413*, 2017.

[36] Y. Li, R. Bu, M. Sun, W. Wu, X. Di, and B. Chen, "Pointcnn: Convolution on x-transformed points," *Advances in neural information processing systems*, vol. 31, pp. 820–830, 2018.

[37] H. Thomas, C. R. Qi, J.-E. Deschaud, B. Marcotegui, F. Goulette, and L. J. Guibas, "Kpconv: Flexible and deformable convolution for point clouds," *Proceedings of the IEEE International Conference on Computer Vision*, 2019.

[38] W. Wu, Z. Qi, and L. Fuxin, "Pointconv: Deep convolutional networks on 3d point clouds," *arXiv preprint arXiv:1811.07246*, 2018.

[39] G. Li, M. Muller, A. Thabet, and B. Ghanem, "Deepgcns: Can gcns go as deep as cnns?" in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2019, pp. 9267–9276.

[40] Y. Wang, Y. Sun, Z. Liu, S. E. Sarma, M. M. Bronstein, and J. M. Solomon, "Dynamic graph cnn for learning on point clouds," *ACM Transactions on Graphics (TOG)*, 2019.

[41] M. Tatarchenko, J. Park, V. Koltun, and Q.-Y. Zhou, "Tangent convolutions for dense prediction in 3d," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018, pp. 3887–3896.

[42] Y. Lin, Z. Yan, H. Huang, D. Du, L. Liu, S. Cui, and X. Han, "Fpconv: Learning local flattening for point convolution," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020, pp. 4293–4302.

[43] H. Berman, K. Henrick, and H. Nakamura, "Announcing the worldwide protein data bank," *Nature Structural & Molecular Biology*, vol. 10, no. 12, pp. 980–980, 2003.

[44] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 770–778.

[45] Y. Eldar, M. Lindenbaum, M. Porat, and Y. Y. Zeevi, "The farthest point strategy for progressive image sampling," *IEEE Transactions on Image Processing*, vol. 6, no. 9, pp. 1305–1315, 1997.

[46] F. Milletari, N. Navab, and S.-A. Ahmadi, "V-net: Fully convolutional neural networks for volumetric medical image segmentation," in *2016 fourth international conference on 3D vision (3DV)*. IEEE, 2016, pp. 565–571.

[47] A. Paszke, S. Gross, S. Chintala, G. Chanan, E. Yang, Z. DeVito, Z. Lin, A. Desmaison, L. Antiga, and A. Lerer, "Automatic differentiation in pytorch," in *NIPS-W*, 2017.