

# ProCLR: Supervised Contrastive Framework for Robust Protein Surface Classification

Wissam Ferroudj

*LIRIS Laboratory, UMR 5205 CNRS  
National Institute of Applied Sciences  
Lyon, France  
wissam.ferroudj@insa-lyon.fr*

Noura Faci

*LIRIS Laboratory, UMR 5205 CNRS  
Claude Bernard Lyon 1 University  
Lyon, France  
noura.faci@univ-lyon1.fr*

Hamamache Kheddouci

*LIRIS Laboratory, UMR 5205 CNRS  
Claude Bernard Lyon 1 University  
Lyon, France  
hamamache.kheddouci@univ-lyon1.fr*

**Abstract**—Protein classification is a key challenge in bioinformatics, essential for understanding biological functions, molecular interactions, and drug discovery. Yet, the functional landscape of proteins cannot always be inferred from sequence or fold similarity, as structurally divergent proteins may share common biological roles. Since most molecular interactions occur at the protein surface, capturing the geometric and biochemical features of these surfaces is essential for accurate functional inference. In this work, we introduce ProCLR, a graph-based framework that models protein surfaces as simplified, attribute-rich 3D geometric graphs, and applies supervised contrastive learning to learn discriminative embeddings. Our method integrates a geodesic-aware clustering strategy to reduce surface complexity while preserving functional features. We evaluate ProCLR using two encoder architectures—Graph Convolutional Networks (GCNs) and Graph Attention Networks (GATs)—on the SHREC’21 and SHREC’19 benchmarks. Experimental results show that ProCLR effectively captures functional patterns, achieving competitive performance with significantly reduced graph sizes, as measured by sensitivity, specificity, precision, accuracy, and F1-score.

**Index Terms**—Contrastive learning, GAT, GCN, graph, protein surface classification, surface simplification.

## I. INTRODUCTION

Protein classification is a challenging task in bioinformatics that helps understand how proteins function, interact, and supports the development of new drugs. Earlier approaches to protein classification often depend on sequence similarity and structural motifs, assuming that proteins with similar sequences or folds perform similar functions [1]–[3]. However, several studies [4], [5] have demonstrated that sequence and structural similarity do not always reliably match functional similarity. In fact proteins with divergent sequences or folds may exhibit convergent functions, while those with similar folds can perform distinct biological roles. This functional–structural decoupling highlights the inherent limitations of classification models based solely on sequence or backbone structure.

Focusing on protein surfaces provides a clearer understanding of a protein’s functional capabilities, as surfaces enable most biochemical interactions. Their complex topography and physicochemical properties play a key role in ligand binding, protein–protein interactions, and other biological processes. Despite this, surface-based classification remains relatively underexplored. Early efforts leveraging physicochemical and

geometric descriptors inspired by 3D object recognition have laid important groundwork [6]–[12]. However, existing methods based on surface similarity can struggle to distinguish proteins from different functional classes when they exhibit highly alike surface characteristics due to conformational dynamics and motions [13]. Addressing the challenge of distinguishing proteins with similar surface features requires methods that capture both intra-class and inter-class differences by incorporating biochemical properties that govern functional interactions. Since the molecular surface encodes unique chemical and geometric features serving as a fingerprint of a protein’s interaction behavior, our approach tackles this challenge by employing a geometric deep learning framework that models protein surfaces as graphs. This representation captures their irregular and non-Euclidean nature, curved manifolds with complex local connectivity and varying topology. Unlike traditional grid-based architectures such as CNNs, which are limited to regular Euclidean domains, graph neural networks (GNNs) offer the flexibility to operate on such data and can integrate both spatial structure and biochemical context. Geometric deep learning has demonstrated strong performance across various molecular tasks, including protein–protein interaction prediction, drug design [14], drug–target interaction [15] and protein structure prediction [16].

Graph-based modeling of protein surfaces enables the simultaneous encoding of local geometric features and contextual biochemical signals. However, dense molecular meshes typically contain tens of thousands of vertices, leading to prohibitively large graphs. To address this, we introduce a graph simplification procedure tailored to protein surfaces, which preserves structural and functional properties while reducing computational complexity.

To further improve the expressiveness of the learned representations, we employ supervised contrastive learning. This strategy encourages the model to cluster embeddings of functionally similar surfaces, while pushing apart dissimilar ones, even when geometric similarities exist due to structural convergence. By leveraging both structural and functional cues during training, contrastive learning promotes robust and discriminative embeddings that generalize across conformational variants. Recent work has demonstrated the effectiveness of contrastive learning on molecular graphs in other bioinformat-

ics tasks, such as drug–target affinity prediction, where hybrid-scale graph contrastive frameworks like HSGCL-DTA [17] have shown improved generalization by capturing complementary views of molecular relationships.

Building on this foundation, we introduce a novel framework for protein surface classification that relies on simplified graph representations combined with supervised contrastive learning. Our key contributions are: (i) a principled method for constructing graphs from protein surfaces that captures both geometric and biochemical features, (ii) a surface-specific graph simplification technique designed to enable scalable learning, and (iii) a contrastive learning approach that improves functional discrimination across a wide range of protein surfaces.

The rest of the paper is organized as follows. Section II examines related work in surface-based and geometric protein modeling. Section III describes our graph-based surface representation and learning framework. Section IV details the experiments and the obtained results. Section V summarizes the findings and outlines future directions.

## II. RELATED WORK

This section reviews advances in protein surface representation and analysis, focusing on methods for comparison and functional classification. Surface representations are crucial for understanding biological roles, drug discovery, and large-scale studies. Strategies can be grouped into four main categories: **Geometric deep learning-based approaches.** Approaches like MaSIF [18] and SurfaceID [19] exploit local surface patches using geodesic or geometric convolutional networks, showing strong potential for interaction prediction. However, they mainly target pairwise retrieval or interaction prediction and depend on large annotated datasets.

**3D Zernike descriptor-based methods.** such as 3D-SURFER [20], BioZernike [21] and FP-Zernike [22] represent protein shapes through rotationally invariant global descriptors. While efficient for retrieval, they often lose local functional information.

**Projection map and molecular cartography methods.** Tools like SURFMAP [11] or Structuprint [12] map surfaces into 2D to simplify feature visualization and comparison. These approaches simplify analysis but introduce distortions and orientation bias.

**Spectral methods.** Shape-DNA [6] and PLO3S [10] leverage Laplace–Beltrami eigenfunctions to capture intrinsic surface properties, providing descriptors naturally invariant to isometric transformations such as bending and twisting invariance to deformations. Despite their robustness, spectral descriptors can be affected by topological inconsistencies and distortions caused by spherical projections. Furthermore, their computational demands may restrict their use on high-resolution surfaces and large datasets.

**Graph-based representations** While many geometric deep learning methods operate on non-Euclidean domains, few explicitly construct graphs from entire protein surfaces. Here, we focus on such surface-level graph representations. Graphs

have proven valuable for modeling residue connectivity and aiding structural alignment at the sequence and tertiary structure levels [23], [24], yet their use directly on protein molecular surfaces remains limited. Notable exceptions include SURFCOMP [25], which creates graphs from key surface points and applies subgraph matching for pairwise comparisons, and EquiPockets [26], which encode pocket-shaped point clouds as graphs and leverage equivariant graph neural networks to compare binding regions. Both focus mainly on retrieval or alignment rather than classification. Meanwhile, MPGCNN [27] employs simple k-NN graphs constructed on surface vertices for classification, but more sophisticated graph-based classification methods are still largely unexplored.

Our work explores this promising direction by modeling protein surfaces as simplified geometric graphs enriched with physicochemical attributes. We apply supervised contrastive learning to enhance class separability, aiming to combine the expressiveness of graph structures with the generalization ability of modern deep learning to tackle the complex challenge of protein surface classification. Unlike traditional graph embeddings that can lose distinction due to class similarity or structural noise, contrastive learning actively guides the model to bring functionally related surfaces closer together while separating those that differ, even when geometry alone is insufficient. This approach produces embeddings that are both compact and highly discriminative.

## III. GRAPH-BASED PROTEIN SURFACE CLASSIFICATION FRAMEWORK

To overcome challenges from protein surface complexity and functional diversity, we propose a novel framework named *ProCLR* that combines biochemical-aware surface simplification with graph-based supervised contrastive learning.

### A. Overview

Figure 1 depicts the overall workflow of our approach, integrating biochemical-aware surface simplification and supervised contrastive learning on graph representations to classify protein surfaces. Our approach consists of two main stages.

Stage 1 focuses on simplifying protein surface data to reduce complexity while preserving key biochemical information. It begins by generating high-resolution mesh surfaces from protein structures. To manage the complexity and emphasize important regions, the meshes are divided into geodesic patches. Each patch is then clustered to identify representative centroids, which capture essential features. These centroids, enriched with aggregated biochemical and geometric information, form a compact graph representation of the protein surface.

Stage 2 learns discriminative and functionally meaningful representations from the simplified graphs using a GNN trained with supervised contrastive learning. The resulting embeddings are used to train a linear classifier for functional prediction. This stage is further detailed in Figure 2.




Fig. 1. ProCLR’s overview

## B. Surface representation simplification

Proteins exhibit highly complex surfaces, often represented by meshes containing upwards of 50,000 vertices. Directly measuring similarities between these surfaces is challenging due to their large size and geometric complexity. Additionally, physicochemical properties such as hydrophobicity and electrostatic potential play crucial roles in protein functionality. Therefore, effective protein surface comparison requires capturing both geometric and biochemical information. Rather than relying on traditional geometric decimation methods, we propose a novel surface simplification technique that integrates both biochemical and geometric features.

Our proposed method’s workflow is outlined in Algorithm 1 and visualized in Figure 1. Algorithm 1 describes the complete simplification workflow, which starts after generating the molecular surface mesh from the PDB file. From this mesh, we first construct a dense surface graph where vertices correspond to mesh vertices and edges reflect local connectivity (step 1). The process then begins with geodesic patch extraction using a truncated Dijkstra search (Algorithm 2), ensuring that each patch preserves local coherence within a fixed radius of 3 Å (step 2). This radius is chosen to remain within the scale of a single atom, as hydrogen- the smallest atom- has an approximate size of 3 Å. Within each patch, hierarchical clustering merges vertices according to biochemical and geometric similarity, and each cluster is then represented by its centroid (step 3). Finally, spatial proximity between centroids is captured through a  $k$ -nearest neighbor graph, yielding a compact yet information-rich representation of the original surface (step 4) <sup>1</sup>.

This integrated pipeline enables a compact, attributed graph representation of the protein surface that retains important biochemical and geometric information while greatly reducing complexity. As illustrated in Table I, the number of nodes and edges in the simplified output graphs is drastically reduced compared to the original graphs, while preserving key structural properties such as average degree and graph density.

<sup>1</sup>The implementation is available at <https://github.com/WissamFER/ProCLR-mesh-simplification>

TABLE I  
GRAPH STATISTICS BEFORE AND AFTER SIMPLIFICATION

| Graph    | Nodes | Edges | Avg. Degree | Graph Density |
|----------|-------|-------|-------------|---------------|
| Original | 29964 | 89892 | 6.00        | 0.0002        |
| Reduced  | 4122  | 10193 | 4.94        | 0.0012        |

## Algorithm 1 Patch extraction followed by centroid-graph construction

**Require:** mesh, features, curvature, max\_radius, threshold  
**Ensure:**  $G_c$ : centroid-level surface graph  
 /\* Step 1: Construct the initial surface-level geodesic graph \*/  
 1:  $G \leftarrow \text{construct\_surface\_graph}(\text{mesh})$   
 /\* Step 2: Partition surface into geodesic patches \*/  
 2: remaining  $\leftarrow \{1, \dots, N\}$ , patches  $\leftarrow []$   
 3: **while** remaining  $\neq \emptyset$  **do**  
 4:   seed  $\leftarrow \text{pop}(\text{remaining})$   
 5:    $S \leftarrow \text{TruncatedDijkstra}(G, \text{seed}, \text{max\_radius})$   
 6:   patch  $\leftarrow S \cap \text{remaining}$   
 7:   patches.append(patch)  
 8:   remaining  $\leftarrow \text{remaining} \setminus \text{patch}$   
 9: **end while**  
 /\* Step 3: Perform feature-based clustering within each patch \*/  
 10: centroids  $\leftarrow []$   
 11: **for all** patch in patches **do**  
 12:   clusters  $\leftarrow \text{AggloCluster}(\text{patch}, \text{features}, \text{threshold})$   
 13:   **for all** cluster  $c$  in clusters **do**  
 14:     centroid  $\leftarrow \text{aggregate\_cluster\_features}(c)$   
 15:     centroids.append(centroid)  
 16:   **end for**  
 17: **end for**  
 /\* Step 4: Build centroid-level graph via spatial  $k$ -NN \*/  
 18: **if** centroids  $\neq \emptyset$  **then**  
 19:    $k \leftarrow \text{find\_optimal\_k}(\text{centroids}[:,4:], 10)$   
 20:    $G_c \leftarrow \text{build\_knn\_graph}(\text{centroids}[:,4:], k + 1)$   
 21: **end if**  
 22: **return**  $G_c$

## C. Graph-based protein surface classification

Our training strategy consists of two main phases: a **supervised contrastive pre-training** followed by a **fine-tuning step** for classification. The framework’s main components are:

- **Data augmentation module** helps the model handle noisy or variable graph structures. For each graph  $x_i$  in a batch of size  $N$ , we generate two distinct augmented views  $\tilde{x}_i^1$  and  $\tilde{x}_i^2$  through transformations such as node feature masking and edge perturbation. We obtain a batch of  $2N$  graphs, each with its label  $y_i$ .
- **Encoder** refers to a GNN,  $\text{Enc}(\cdot)$ , that maps an input

**Algorithm 2** Truncated dijkstra for geodesic patching

---

```

1: procedure TruncatedDijkstra( $G, start, max\_radius$ )
  /* Initialize distances for all vertices */
2:   for all  $v \in G.V$  do
3:      $dist[v] \leftarrow \infty$ 
4:   end for
5:    $dist[start] \leftarrow 0$ 
  /* Initialize priority queue with the seed node */
6:   Initialize min-heap  $Q \leftarrow [(0, start)]$ 
  /* Standard Dijkstra loop with truncation */
7:   while  $Q \neq \emptyset$  do
8:      $(d, u) \leftarrow Q.pop\_min()$ 
9:     if  $d > max\_radius$  then
10:      continue
11:    end if
    /* Relax neighbors within geodesic constraint */
12:    for all neighbor  $v$  of  $u$  do
13:       $alt \leftarrow d + G.weight(u, v)$ 
14:      if  $alt < dist[v]$  and  $alt \leq max\_radius$  then
15:         $dist[v] \leftarrow alt$ 
16:         $Q.insert(alt, v)$ 
17:      end if
18:    end for
19:  end while
  /* Return reachable vertices under max_radius */
20:  return  $\{v : dist[v] < \infty\}$ 
21: end procedure

```

---

graph to a vector representation ( $r \in \mathbb{R}^{D_E}$ ). This representation captures both the topological and feature-based information of the input.

- **Projection head** is a Multi-Layer Perceptron (MLP),  $Proj(\cdot)$ , that maps the encoder output  $r$  to a latent space  $z = Proj(r) \in \mathbb{R}^{D_P}$  where contrastive learning is applied to prevents the collapse in the representation space.
- **Supervised contrastive loss** (SupCon) [28] leverages label information by treating all samples from the same class as positives. SupCon thus generalizes the SimCLR framework by incorporating multiple positive examples per anchor based on class membership. Let  $z_i$  denote the anchor,  $P(i)$  the set of indices of positives for  $i$  (i.e., samples in the batch with the same label as  $z_i$ ), and  $A(i) = 1, \dots, 2N \setminus i$ . Equation 1 defines the supervised contrastive loss for an anchor  $z_i$ :

$$\mathcal{L}_i = -\frac{1}{|P(i)|} \sum_{p \in P(i)} \log \frac{\exp(z_i \cdot z_p / \tau)}{\sum_{a \in A(i)} \exp(z_i \cdot z_a / \tau)}, \quad (1)$$

where  $\tau$  is a temperature hyperparameter [29]. This loss encourages representations of samples from the same class to be pulled together in the projection space, fostering discriminative and class-aware feature learning.

#### IV. EXPERIMENTS

We evaluate our framework on the SHREC'21 dataset [30] using a set of baseline methods for comparison. The experiments include data preprocessing, model training, and performance evaluation using standard classification metrics.

##### A. Preprocessing pipeline

The SHREC'21 dataset consists of 209 PDB entries, each representing a protein with multiple conformations obtained

through Nuclear Magnetic Resonance (NMR) spectroscopy. This results in approximately 5,000 molecular surfaces, capturing various stable conformational states that a protein can adopt. Each surface is enriched with physicochemical properties evaluated at the mesh vertices, including electrostatic potential, hydrophobicity, and the presence of hydrogen bond donors and acceptors.

The molecular surfaces of the SHREC'21 were originally generated using NanoShaper [31], following the Connolly Solvent Excluded Surface (SES) model. The dataset was then split into two subsets: 70% for training and 30% for testing.

However, during the computation of geodesic distances, we observed that some surfaces generated by NanoShaper contained disconnected mesh components, which could hinder downstream analysis. To address this issue, we reconstructed the meshes using EDTSurf [32], a tool that ensures the generation of connected surface meshes. This reconstruction resulted in a different number of vertices compared to the original OFF files.

To preserve the original physicochemical annotations, we applied interpolation techniques to map the property values from the original vertices onto the new mesh generated by EDTSurf. In addition, we computed a fourth geometric property, the *mean curvature*, estimated following the method described in [33].

##### B. Experimental setup

We evaluate our framework using two different GNN architectures as encoders: the standard Graph Convolutional Network (GCN) and the Graph Attention Network (GAT). Both are trained using the same supervised contrastive learning strategy and later fine-tuned with a linear classifier for protein surface classification.

We use a 3 layer Graph for each encoder, with hidden\_channels 128 and specifically 4 attention heads per layer for the GAT encoder. The output embeddings are pooled via both global mean and max pooling and then passed to a projection head composed of two fully connected layers with ReLU activation and L2 normalization. The classifier is a single linear layer.

The contrastive model is trained for up to 150 epochs with early stopping (patience = 10), using the Adam optimizer and an initial learning rate of 1e3. A linear warm-up is applied during the first 10 epochs. The batch size is set to 128, and we use two types of graph augmentations: node feature masking (with varying masking ratios) and edge perturbation. The best model is selected based on the validation loss.

The baseline methods, taken from the SHREC'21 benchmark, serve as a solid reference for evaluating our approach. These methods fall into two main categories: handcrafted descriptors and learned descriptors based on supervised learning, as summarized in Table II. We compare their performance across various masking levels using several evaluation metrics, including True Positive Rate or Sensitivity (TPR), True Negative Rate or Specificity (TNR), Accuracy (ACC), and F1-score (F1) [34].




Fig. 2. Supervised contrastive learning pipeline for protein surface graphs.

TABLE II  
SUMMARY OF BASELINE METHODS EVALUATED ON SHREC’21

| Type                    | Method                                  | Brief description                                                                                                                                                |
|-------------------------|-----------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Handcrafted Descriptors | P1 (Joint Histograms of Curvature) [30] | Multi-scale histograms of minimum/maximum curvatures combined with area projection; similarity measured via Jeffreys divergence, enhanced by LDA and clustering. |
|                         | P3 (HAPPS) [30]                         | Fusion of local point pair features and global histograms; biochemical textures encoded and compared with EMD/KLD.                                               |
|                         | P4 (GLoFe) [30]                         | Extraction of global and local features (surface unfolding, shape indices, volume); distances combined with supervised weights.                                  |
| Learned Descriptors     | P2 (3DZD) [30]                          | Rotation-invariant 3D Zernike descriptors on voxelized surfaces; similarity scored by MLP or end-to-end learning.                                                |
|                         | P5 (MPGCNN + PointNet) [30]             | k-NN graphs on enriched point clouds; embeddings learned with GNN or PointNet; compared via L2 distance or ensemble fusion.                                      |

All trainings experiments are implemented using PyTorch 2.2.0 and PyTorch Geometric, running on a Python 3.11 environment. The code is executed on a supercomputer with two NVIDIA A100 GPUs using PyTorch DistributedDataParallel.

### C. Results and discussion

Table III presents a performance comparison between our GCN and GAT models and the baseline methods from the SHREC’21 challenge. Our results show that GAT consistently outperforms our GCN across all metrics, confirming the advantage of attention-based aggregation for protein surface graphs. Although GAT does not surpass state-of-the-art methods from the challenge, it achieves competitive performance while operating on significantly simplified graphs. This demonstrates that our mesh reduction and graph construction strategy successfully reduces computational complexity without compromising performances.

Furthermore, Table IV analyzes the effect of different masking ratios during contrastive pretraining, highlighting that a 0.4. (40%) masking ratio yields the best performance.

To illustrate training behavior, we also plotted the supervised contrastive loss curves of both GCN and GAT models

under their best-performing masking configurations. As shown in Figure 3, the GAT model exhibits faster convergence.

TABLE III  
PERFORMANCE COMPARISON OF SHREC’21 BASELINE METHODS AND OUR APPROACH

| Method              | <i>TPR</i> | <i>TNR</i> | <i>ACC</i> | <i>F1</i> |
|---------------------|------------|------------|------------|-----------|
| P1                  | 0.9891     | 1.0000     | 1.0000     | 0.9980    |
| P2                  | 0.9987     | 0.9999     | 0.9999     | 0.9987    |
| P3                  | 0.9942     | 0.9996     | 0.9975     | 0.9942    |
| P4                  | 0.9903     | 0.9995     | 0.9992     | 0.9904    |
| P5                  | 0.9501     | 0.9975     | 0.9968     | 0.9491    |
| <b>ProCLR (GCN)</b> | 0.9502     | 0.9997     | 0.9607     | 0.9441    |
| <b>ProCLR (GAT)</b> | 0.9907     | 1.0000     | 0.9831     | 0.9774    |

To further validate our pipeline, we evaluated it on the SHREC’19 dataset, which contains mesh data only without electrostatic properties. The meshes were converted into graphs, and mean curvature values were computed at each surface vertex as features. The GCN and GAT models were then applied to both species-level and protein-level classification tasks. As shown in Table V, our approach outperforms several state-of-the-art methods evaluated on this dataset. In addition,

TABLE IV  
PERFORMANCES OF OUR METHOD ON SHREC’21 UNDER DIFFERENT  
FEATURE MASKING RATIOS

| Mask Ratio | ACC           | F1            | TPR           | TNR           |
|------------|---------------|---------------|---------------|---------------|
| 0.2        | <b>0.9831</b> | 0.9768        | 0.9802        | <b>0.9999</b> |
| 0.4        | <b>0.9831</b> | <b>0.9774</b> | <b>0.9807</b> | <b>0.9999</b> |
| 0.5        | 0.9757        | 0.9695        | 0.9730        | 0.9998        |
| 0.6        | 0.9719        | 0.9635        | 0.9718        | 0.9998        |
| 0.8        | 0.9813        | 0.9702        | 0.9746        | <b>0.9999</b> |




Fig. 3. Supervised contrastive losses for GCN and GAT-head3 on the SHREC’21 dataset during the pretraining stage.

Table VI provides a detailed performance breakdown across different masking ratios for GAT, revealing that the optimal ratio is 0.5 at the protein level and 0.8 at the species level.

TABLE V  
COMPARISON OF METHODS ON SHREC’19 – PROTEIN AND SPECIES  
LEVELS

| Method               | Protein Level |              | Species Level |              |
|----------------------|---------------|--------------|---------------|--------------|
|                      | MAP           | F1           | MAP           | F1           |
| Triangles Stars [35] | 0.880         | 0.830        | –             | –            |
| GEmb3D [36]          | 0.857         | 0.800        | –             | –            |
| Fraction-GauB [37]   | 0.928         | 0.910        | –             | –            |
| 3DZD [38]            | 0.720         | 0.490        | 0.610         | –            |
| 3DZM [38]            | 0.649         | 0.350        | 0.604         | –            |
| ConvLDSnet [38]      | 0.329         | 0.270        | 0.355         | –            |
| Ft-PSSC [38]         | 0.417         | 0.280        | 0.405         | –            |
| HAPT [38]            | 0.666         | 0.470        | 0.685         | –            |
| GF-GAT [39]          | 0.986         | 0.978        | 0.957         | 0.956        |
| <b>ProCLR (GCN)</b>  | 0.891         | 0.695        | 0.894         | 0.787        |
| <b>ProCLR (GAT)</b>  | <b>0.983</b>  | <b>0.981</b> | <b>0.962</b>  | <b>0.966</b> |

## V. CONCLUSION AND PERSPECTIVE

In this study, we propose a novel graph-based framework for protein surface classification that synergistically integrates surface simplification, geodesic-aware clustering, and supervised contrastive learning. By converting high-resolution molecular meshes into compact, attributed graphs, our method effectively encodes both geometric and physicochemical properties in a representation well suited for GNN architectures.

Comprehensive experimental evaluations on the SHREC’21 and SHREC’19 benchmarks demonstrate the effectiveness of our approach, particularly when employing attention-based models such as GAT. These findings reflect the framework’s ability to capture rich, discriminative features from protein surfaces. This capability naturally extends to more granular, localized analysis, as the graph representation preserves spatial coherence by modeling protein surfaces as patches. Such localized embeddings enable the detection of functional sites and binding pockets by facilitating direct comparison of region-specific features, which are often critical for understanding protein function but may be overlooked by global descriptors. We aim to extend this work by focusing on local similarity retrieval through patch-level embeddings to identify analogous surface regions across proteins, alongside contrastive pretraining guided by binding site supervision using known ligand or interaction partner contacts. These directions are expected to further enhance the interpretability and applicability of our graph-based approach to protein surface analysis.

## VI. ACKNOWLEDGMENTS

This work has received funding from the French Agence Nationale pour la Recherche (ANR) under the grant reference GRADIENT ANR-22-CE23-0009. This work was granted access to the HPC resources of IDRIS under the allocation 2024-AD011014941 made by GENCI.

## REFERENCES

- [1] W. R. Pearson and D. J. Lipman, “Improved tools for biological sequence comparison,” *Proceedings of the National Academy of Sciences*, vol. 85, no. 8, pp. 2444–2448, 1988.
- [2] Y. Wang, H. Wu, and Y. Cai, “A benchmark study of sequence alignment methods for protein clustering,” *BMC bioinformatics*, vol. 19, pp. 95–104, 2018.
- [3] M. Carpentier and J. Chomilier, “Protein multiple alignments: sequence-based versus structure-based programs,” *Bioinformatics*, vol. 35, no. 20, pp. 3970–3980, 2019.
- [4] A. Alcami, “Viral mimicry of cytokines, chemokines and their receptors,” *Nature Reviews Immunology*, vol. 3, no. 1, pp. 36–50, 2003.
- [5] M. L. Chiu, D. R. Goulet, A. Teplyakov, and G. L. Gilliland, “Antibody structure and function: the basis for engineering therapeutics,” *Antibodies*, vol. 8, no. 4, p. 55, 2019.
- [6] M. Reuter, F.-E. Wolter, and N. Peinecke, “Laplace-beltrami spectra as ‘shape-dna’ of surfaces and solids,” *Computer-Aided Design*, vol. 38, no. 4, pp. 342–366, 2006.
- [7] P. Papadakis, I. Pratikakis, T. Theoharis, and S. Perantonis, “Panorama: A 3d shape descriptor based on panoramic views for unsupervised 3d object retrieval,” *International Journal of Computer Vision*, vol. 89, pp. 177–192, 2010.
- [8] I. Budowski-Tal, R. Kolodny, and Y. Mandel-Gutfreund, “A novel geometry-based approach to infer protein interface similarity,” *Scientific reports*, vol. 8, no. 1, p. 8192, 2018.
- [9] D. La, J. Esquivel-Rodríguez, V. Venkatraman, B. Li, L. Sael, S. Ueng, S. Ahrendt, and D. Kihara, “3d-surfer: software for high-throughput protein surface comparison and analysis,” *Bioinformatics*, vol. 25, no. 21, pp. 2843–2844, 2009.
- [10] L. Sirugue, F. Langenfeld, N. Lagarde, and M. Montes, “Plo3s: Protein local surficial similarity screening,” *Computational and Structural Biotechnology Journal*, vol. 26, pp. 1–10, 2024.
- [11] H. Schweke, M.-H. Mucchielli, N. Chevrollier, S. Gosset, and A. Lopes, “Surfmap: a software for mapping in two dimensions protein surface features,” *Journal of Chemical Information and Modeling*, vol. 62, no. 7, pp. 1595–1601, 2022.
- [12] D. G. Kontopoulos, D. Vlachakis, G. Tsiliki, and S. Kossida, “Struc-tuprint: a scalable and extensible tool for two-dimensional representation of protein surfaces,” *BMC Structural Biology*, vol. 16, pp. 1–8, 2016.

TABLE VI  
PERFORMANCE ON SHREC'19 AT PROTEIN AND SPECIES LEVELS FOR DIFFERENT MASKING RATIOS

| Mask Ratio | Protein Level |               |               | Species Level |               |               |
|------------|---------------|---------------|---------------|---------------|---------------|---------------|
|            | Acc           | F1            | MAP           | Acc           | F1            | MAP           |
| 0.2        | 0.9723        | 0.9651        | <b>0.9838</b> | <b>0.9861</b> | 0.9519        | 0.9511        |
| 0.4        | 0.9421        | 0.9202        | 0.9794        | <b>0.9861</b> | 0.9508        | 0.9506        |
| 0.5        | <b>0.9849</b> | <b>0.9808</b> | 0.9828        | 0.9824        | 0.9456        | 0.9544        |
| 0.6        | 0.9748        | 0.9106        | 0.9701        | 0.9635        | 0.9463        | 0.9464        |
| 0.8        | 0.9698        | 0.9158        | 0.9820        | 0.9811        | <b>0.9661</b> | <b>0.9619</b> |

- [13] C. A. Orengo, A. E. Todd, and J. M. Thornton, "From protein structure to function," *Current opinion in structural biology*, vol. 9, no. 3, pp. 374–382, 1999.
- [14] C. Isert, K. Atz, and G. Schneider, "Structure-based drug design with geometric deep learning," *Current Opinion in Structural Biology*, vol. 79, p. 102548, 2023.
- [15] K. Shao, Z. Zhang, S. He, and X. Bo, "Dtigccn: Prediction of drug-target interactions based on gcnn and cnn," in *2020 IEEE 32nd International Conference on Tools with Artificial Intelligence (ICTAI)*. IEEE, 2020, pp. 337–342.
- [16] S. C. Pakhrin, B. Shrestha, B. Adhikari, and D. B. Kc, "Deep learning-based advances in protein structure prediction," *International journal of molecular sciences*, vol. 22, no. 11, p. 5553, 2021.
- [17] H. Ye, Y. Song, B. Wang, L. Wu, S. He, X. Bo, and Z. Zhang, "Hsgcl-dta: Hybrid-scale graph contrastive learning based drug-target binding affinity prediction," in *2023 IEEE 35th International Conference on Tools with Artificial Intelligence (ICTAI)*, 2023, pp. 947–954.
- [18] P. Gainza, F. Sverrisson, F. Monti, E. Rodola, D. Boscaini, M. M. Bronstein, and B. E. Correia, "Deciphering interaction fingerprints from protein molecular surfaces using geometric deep learning," *Nature Methods*, vol. 17, no. 2, pp. 184–192, 2020.
- [19] S. Riahi, J. H. Lee, T. Sorenson, S. Wei, S. Jager, R. Olfati-Saber, Y. Zhou, A. Park, M. Wendt, H. Minoux *et al.*, "Surface id: a geometry-aware system for protein molecular surface comparison," *Bioinformatics*, vol. 39, no. 4, p. btad196, 2023.
- [20] D. La, J. Esquivel-Rodríguez, V. Venkatraman, B. Li, L. Sael, S. Ueng, S. Ahrendt, and D. Kihara, "3d-surfer: software for high-throughput protein surface comparison and analysis," *Bioinformatics*, vol. 25, no. 21, pp. 2843–2844, 2009.
- [21] D. Guzenko, S. K. Burley, and J. M. Duarte, "Real time structural search of the protein data bank," *PLoS computational biology*, vol. 16, no. 7, p. e1007970, 2020.
- [22] J. Qi, C. Feng, Y. Shi, J. Yang, F. Zhang, G. Li, and R. Han, "Fp-zernike: an open-source structural database construction toolkit for fast structure retrieval," *Genomics, Proteomics and Bioinformatics*, vol. 22, no. 1, p. qzae007, 2024.
- [23] B. Mészáros, G. Erdős, and Z. Dosztányi, "Iupred2a: context-dependent prediction of protein disorder as a function of redox state and protein binding," *Nucleic acids research*, vol. 46, no. W1, pp. W329–W337, 2018.
- [24] P. Zhang, L. Tao, X. Zeng, C. Qin, S. Chen, F. Zhu, Z. Li, Y. Jiang, W. Chen, and Y.-Z. Chen, "A protein network descriptor server and its use in studying protein, disease, metabolic and drug targeted networks," *Briefings in bioinformatics*, vol. 18, no. 6, pp. 1057–1070, 2017.
- [25] C. Hofbauer, H. Lohninger, and A. Aszódi, "Surfcomp: a novel graph-based approach to molecular surface comparison," *Journal of chemical information and computer sciences*, vol. 44, no. 3, pp. 837–847, 2004.
- [26] Y. Zhang, Z. Wei, Y. Yuan, C. Li, and W. Huang, "Equipocket: an e (3)-equivariant geometric graph neural network for ligand binding site prediction," *arXiv preprint arXiv:2302.12177*, 2023.
- [27] 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)*, vol. 38, no. 5, pp. 1–12, 2019.
- [28] P. Khosla, P. Teterwak, C. Wang, A. Sarna, Y. Tian, P. Isola, A. Maschinot, C. Liu, and D. Krishnan, "Supervised contrastive learning," *Advances in neural information processing systems*, vol. 33, pp. 18 661–18 673, 2020.
- [29] A. Bardes, J. Ponce, and Y. LeCun, "Vicregl: Self-supervised learning of local visual features," *Advances in Neural Information Processing Systems*, vol. 35, pp. 8799–8810, 2022.
- [30] A. Raffo, U. Fugacci, S. Biasotti, W. Rocchia, Y. Liu, E. Otu, R. Zwigge-laar, D. Hunter, E. I. Zacharaki, E. Psatha *et al.*, "Shrec 2021: Retrieval and classification of protein surfaces equipped with physical and chemical properties," *Computers & Graphics*, vol. 99, pp. 1–21, 2021.
- [31] S. Decherchi, A. Spitaleri, J. Stone, and W. Rocchia, "Nanoshaper-vmd interface: computing and visualizing surfaces, pockets and channels in molecular systems," *Bioinformatics*, vol. 35, no. 7, pp. 1241–1243, 08 2018. [Online]. Available: <https://doi.org/10.1093/bioinformatics/bty761>
- [32] D. Xu and Y. Zhang, "Edtsurf: Quick and accurate construction of macromolecular surfaces," *PLoS One*, 2009.
- [33] D. Cohen-Steiner and J.-M. Morvan, "Restricted delaunay triangulations and normal cycle," in *Proceedings of the nineteenth annual symposium on Computational geometry*, 2003, pp. 312–321.
- [34] M. Grandini, E. Bagli, and G. Visani, "Metrics for multi-class classification: an overview," *arXiv preprint arXiv:2008.05756*, 2020.
- [35] K. Madi and E. Paquet, "3d deformable protein shapes classification based on triangles-stars and composite deep neural networks," in *2022 IEEE/ACS 19th International Conference on Computer Systems and Applications (AICCSA)*. IEEE, 2022, pp. 1–8.
- [36] K. Madi *et al.*, "New graph embedding approach for 3d protein shape classification," in *2020 Joint 9th International Conference on Informatics, Electronics 'I&' Vision (ICIEV) and 2020 4th International Conference on Imaging, Vision 'I&' Pattern Recognition (icIVPR)*, 2020, pp. 1–7.
- [37] E. Paquet, H. L. Viktor, K. Madi, and J. Wu, "Deformable protein shape classification based on deep learning, and the fractional fokker-planck and kähler-dirac equations," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 45, no. 1, pp. 391–407, 2022.
- [38] F. Langenfeld, A. Axenopoulos, H. Benhabiles, P. Daras, A. Giachetti, X. Han, K. Hammoudi, D. Kihara, T. M. Lai, H. Liu *et al.*, "Shrec19 protein shape retrieval contest," in *Eurographics Workshop on 3D Object Retrieval, Delft, The Netherlands*, 2019, pp. 25–31.
- [39] W. Ferroudj, N. Faci, and H. Kheddouci, "Geometric features and gat neural network for protein surface classification," in *International Conference on Database and Expert Systems Applications*. Springer, 2024, pp. 320–325.