

000 001 002 003 004 005 006 007 008 009 010 011 012 013 014 015 016 017 018 019 020 021 022 023 024 025 026 027 028 029 030 031 032 033 034 035 036 037 038 039 040 041 042 043 044 045 046 047 048 049 050 051 052 053 EPIFORMER: A TRANSFORMER-BASED MULTI- RELATIONAL EQUIVARIANT GRAPH NEURAL NETWORK FOR ANTIBODY-AWARE EPITOPE PREDICTION

006 **Anonymous authors**

007 Paper under double-blind review

011 ABSTRACT

013 Antibodies are essential components of the immune system, neutralizing foreign
014 antigens such as viruses by binding to specific regions called epitopes. Com-
015 putational prediction of epitopes is critical for antibody design and therapeutic
016 development. Current approaches for epitope prediction still remain challenging
017 due to: (1) lack of sophisticated architectures to model the complex interaction
018 patterns; (2) ineffective protein representations; (3) antibody-agnostic modeling
019 despite antibody specificity; (4) severe class imbalance; and (5) scarcity of known
020 antigen-antibody complexes. In order to overcome these challenges, we propose
021 *EpiFormer*, an encoder-decoder-based architecture that utilizes an $E(3)$ -equivariant
022 multi-relational graph neural network (GNN) coupled with cross-attention to
023 model antigen-antibody interactions. Our contributions are an $E(3)$ -equivariant
024 multi-relational GNN, a Transformer-style cross-attention mechanism, and tailored
025 losses for severe class imbalance and data scarcity. Our method significantly out-
026 performs existing baselines on the Antibody-specific Epitope Prediction (AsEP)
027 dataset by achieving an overall $\approx 1.7x$ performance improvement on multiple
028 classification metrics. This work advances the state-of-the-art in antibody-aware
029 epitope prediction, providing a robust framework for therapeutic antibody design
030 and vaccine development.

031 1 INTRODUCTION

032 Antibodies are large, Y-shaped proteins produced by B-cells that play a critical role in the immune
033 system by identifying and neutralizing foreign substances such as toxins, bacteria, and viruses,
034 collectively known as antigens. They are currently known to be the largest class of biotherapeutics,
035 where five of the current top 10 blockbuster drugs are monoclonal antibodies (Norman et al., 2020;
036 Joubbi et al., 2024). Recently, computational approaches have been proposed to design antibodies to
037 aid the existing traditional approaches that are time-consuming, expensive, and laborious (Fischman
038 & Oftran, 2018; Krishnan et al., 2024; Hummer et al., 2022). An important step in computational
039 antibody design is antigen binding site or epitope prediction, which involves identifying the residues
040 on the surface of an antigen that are recognized and bound by an antibody (Zeng et al., 2023).
041 Accurate epitope prediction is also essential for understanding antibody-antigen interactions in
042 biomedical research (Krishnan et al., 2024).

043 Despite significant advances in deep learning-based protein binding site prediction methods, cur-
044 rent approaches for epitope prediction encounter limitations that severely constrain their effective-
045 ness (Wang et al., 2024a; Fang et al., 2023). (1) Existing architectures lack the sophistication to
046 model complex interaction patterns, with standard GNN struggling to differentiate and learn the
047 distinct geometries of antigens and antibodies while missing essential 3D-related inductive biases
048 like translational invariance and rotational/reflectional equivariance (Zhang et al., 2022). 2) Most
049 methods rely on ineffective protein representations, predominantly using sequence-based approaches
050 that fail to capture the complex three-dimensional spatial arrangements of antigen binding sites,
051 despite epitopes being inherently non-linear and conformationally diverse (Hummer et al., 2022). 3)
052 Most approaches remain agnostic to pre-conditioned antibodies and treat epitope prediction as an
053 antibody-independent problem despite the fact that epitopes are antibody-specific (Norman et al.,
2020). 4) The epitope prediction problem suffers from severe data imbalance, as epitopic residues

054 comprise merely 10% of all residues in an antigen (Liu et al., 2024). 5) The sparsity of known
 055 antigen-antibody complexes creates a fundamental data limitation, with only approximately 2,000
 056 available interaction pairs compared to millions of general protein structures (Joubbi et al., 2024).
 057

058 To address these fundamental challenges, we propose *EpiFormer*, an encoder-decoder architecture
 059 that utilizes E(3)-equivariant graph neural networks (EGNN) on multi-relational protein graphs
 060 coupled with cross-attention mechanisms to model antibody-antigen interactions. We represent
 061 antigens and antibodies as multi-relational graphs and introduce essential 3D-related inductive
 062 biases, including translational invariance and rotational equivariance via equivariant message passing
 063 to effectively handle protein geometric constraints that standard GNNs cannot capture (Liao &
 064 Smidt, 2022). We train *EpiFormer* with a custom joint objective that addresses the severe epitope
 065 class imbalance, predicts antigen-antibody interaction maps, and enforces geometric consistency
 066 through inter-chain distance constraints. The framework operates in an antibody-aware manner by
 067 explicitly incorporating antibody structure and binding context through bidirectional cross-attention
 068 mechanisms that enable dynamic modeling of both intra-chain geometric relationships and inter-chain
 069 interaction patterns (Lim et al., 2025). Our main contributions are as follows:
 070

- 071 1. We develop a novel transformer-based GNN architecture that achieves $\approx 1.7x$ performance
 072 improvement over existing baselines on the antibody-aware epitope prediction task on
 073 multiple classification metrics.
- 074 2. We introduce a multi-relational E(3)-equivariant message passing (EGNN-R) framework
 075 that handles multiple edge relations for robust epitope prediction.
- 076 3. We develop *EpiFormer* with a novel joint loss function designed for: (a) severely imbalanced
 077 epitope data, (b) interaction map prediction, and (c) inter-chain geometric consistency.

079 2 RELATED WORK

080 GNN have emerged as a powerful approach for epitope prediction by modeling the spatial and sequential
 081 relationships in protein structures. Several methods demonstrate this approach: *PECAN* (Pittala & Bailey-Kellogg, 2020), *PInet* (Dai & Bailey-Kellogg, 2021), and related work (Jha et al., 2022)
 082 use GNNs with attention mechanisms for protein-protein interaction prediction. *EPMP* (Vecchio et al., 2021) uses a neural message-passing framework with asymmetrical architectures for paratope-
 083 epitope prediction. Recent advances combine protein language model (PLM) embeddings with
 084 graph-based architectures, with *EpiGraph* (Choi & Kim, 2024) using GAT with ESM-2 embeddings,
 085 *AsEP* (Liu et al., 2024) employing the *WALLE* method with ESM-2 and AntiBERTy embeddings, and
 086 *GraphBepi* (Zeng et al., 2023) leveraging ESM-2 representations.
 087

088 These graph-based methods can be categorized based on whether they use antibody-specific information.
 089 Antibody-agnostic approaches, such as *epitope1D* (Silva et al., 2023), *GraphBepi* (Zeng et al.,
 090 2023), and *EpiGraph* (Choi & Kim, 2024), rely on sequential and structural features but lack specificity
 091 for antibody-specific applications (Vecchio et al., 2021). In contrast, antibody-aware methods
 092 like *EpiScan* (Wang et al., 2024a), *PECAN* (Pittala & Bailey-Kellogg, 2020), and *EPMP* (Vecchio
 093 et al., 2021) explicitly incorporate antibody structure or sequence information. Some approaches
 094 like (Lu et al., 2022) combine GNNs with attention-based bidirectional LSTM networks to capture
 095 both local spatial information and global sequence information from antigens.
 096

097 EGNN have gained attention for protein structure modeling because they preserve geometric properties
 098 under rotations and translations, which are essential for capturing 3D protein conformations
 099 (Satorras et al., 2021b; Schütt et al., 2018). Traditional GNNs often fail to maintain these
 100 geometric constraints when processing protein structures, leading to suboptimal representations of
 101 spatial relationships. E(3)-equivariant approaches like EGNN (Satorras et al., 2021b) and Gear-
 102 Net (Zhang et al., 2022) address this limitation by incorporating equivariance directly into the
 103 message-passing framework. Multi-relational graphs further improve protein modeling by representing
 104 different types of interactions through distinct edge types (Zhang et al., 2022). Recent work has
 105 applied these concepts to protein-protein interactions, with methods like (Liao & Smidt, 2022) using
 106 equivariant transformers for molecular modeling and (Lim et al., 2025) employing multi-relational
 107 representations for protein-ligand binding affinity prediction.
 108

108 **3 METHODS**
 109

110 In this section, we present the graph construction, problem formulation, and the architecture of
 111 *EpiFormer*, a model designed for antibody-aware epitope prediction. *EpiFormer* takes as input an
 112 antigen and an antibody, and predicts their binding sites by dynamically modeling their interaction
 113 using geometric message passing and cross-attention. We then present the customized loss functions
 114 tailored for antigen-antibody interaction prediction to train *EpiFormer*.
 115

116 **3.1 PRELIMINARIES**
 117

118 **Graph construction** The protein 3D structure is described as a point cloud of atoms
 119 $\{v_{i,k}\}_{1 \leq i \leq p, 1 \leq k \leq p_i}$, where p_i is the number of atoms in residue v_i and p represents the number of
 120 amino acid residues in the protein. The first four atoms in any residue correspond to its backbone
 121 atoms (N, C_α, C_β, O) and the rest are its side chain atoms. The 3D coordinate of an atom $v_{i,k}$ is
 122 denoted as $x(v_{i,k}) \in \mathbb{R}^3$. Since we work with the *unbound* structures or point clouds of antigen(ag)
 123 and antibody(ab), we build two completely independent residue graphs $\mathcal{G}_{\text{ag}} = (\mathcal{V}_{\text{ag}}, \mathcal{E}_{\text{ag}}, \mathcal{R})$, and
 124 $\mathcal{G}_{\text{ab}} = (\mathcal{V}_{\text{ab}}, \mathcal{E}_{\text{ab}}, \mathcal{R})$. Vertex $v_i \in \mathcal{V}$ represents residue i , centered on C_α at coordinate $\mathbf{x}_i \in \mathbb{R}^3$.
 125 $|\mathcal{V}_{\text{ag}}| = n$, $|\mathcal{V}_{\text{ab}}| = m$, and edges $e_{i,j} \in \mathcal{E}$ encode structural/functional relationships between residues.
 126

127 Each node $v_i \in \mathcal{V}$ is attributed a node feature vector $\mathbf{h}_i \in \mathbb{R}^{d_h}$ and a node coordinate matrix $\mathbf{X}_i \in$
 128 $\mathbb{R}^{3 \times 4}$ consisting of four backbone atoms $\xi = \{\text{N}, \text{C}_\alpha, \text{C}_\beta, \text{O}\}$ (\mathbf{x}_i is short for $\mathbf{x}_{i,\text{C}_\alpha}$). Specifically, the
 129 node feature vector \mathbf{h}_i constitutes handcrafted geometric features and PLM-derived embeddings to
 130 capture both structural and evolutionary information. In addition, each edge $e_{i,j}$ is attributed an edge
 131 feature vector $\mathbf{f}_{i,j} \in \mathbb{R}^{d_f}$ and a tuple of edge relations $\mathbf{r}_{i,j} \subseteq \mathcal{R}$. The edge vector $\mathbf{f}_{i,j}$ encodes features
 132 such as distances and angles to capture both local geometry and global structural context. The set
 133 of edge relations $\mathcal{R} = \{\rho_1, \rho_2, \rho_3, \rho_4\}$ captures distinct protein interactions: sequential relations for
 134 peptide bonds (ρ_1) and short-range coupling (ρ_2), plus spatial relations for local packing shells via
 135 K-nearest neighbors (ρ_3) and medium-range contacts within 8 Å (ρ_4). Please refer to Appendix A.3
 136 for further details. We extend the notation of these attributes to refer to the residue graph \mathcal{G} of the
 137 antigen (or antibody) as $(\mathbf{H}, \mathbf{X}, \mathbf{F}, \mathbf{R})$.
 138

Problem Formulation We formulate the problem as the following two tasks:

Epitope node prediction: A binary node classification task where a residue $v \in \mathcal{V}_{\text{ag}}$ is labeled as an
 139 epitope (1) if it is within 4.5 Å of any residue in \mathcal{V}_{ab} ; otherwise, it is labeled as a non-epitope (0). The
 140 classifier predicts the epitope node labels \hat{y}_{ag} using $f : \mathcal{V}_{\text{ag}} \rightarrow \{0, 1\}$ and is defined as:
 141

$$\hat{y}_{\text{ag}} = f(v_{\text{ag}}; \mathcal{G}_{\text{ag}}, \mathcal{G}_{\text{ab}}) = \begin{cases} 1 & \text{if } v_{\text{ag}} \text{ is an epitope,} \\ 0 & \text{otherwise.} \end{cases} \quad (1)$$

Bipartite graph link prediction: This task predicts the bipartite adjacency matrix $\hat{\mathcal{E}}_{\text{bg}}$ between an-
 142 tibody and antigen in the bipartite graph $\mathcal{G}_{\text{bg}} = (\mathcal{V}_{\text{ag}} \cup \mathcal{V}_{\text{ab}}, \mathcal{E}_{\text{bg}})$, where \mathcal{V}_{ag} and \mathcal{V}_{ab} are disjoint
 143 vertex sets, and $\mathcal{E}_{\text{bg}} \subseteq \mathcal{V}_{\text{ag}} \times \mathcal{V}_{\text{ab}} \in \{0, 1\}^{n \times m}$ denotes inter-molecular contacts between antigen and
 144 antibody. An edge $e_{\text{bg}} \in \mathcal{E}_{\text{bg}}$ is a contact (labeled as 1) if the corresponding residues $(v_{\text{ag}}, v_{\text{ab}})$ are
 145 within 4.5 Å of each other and 0 otherwise. The edge classifier $g : \mathcal{E}_{\text{bg}} \rightarrow \{0, 1\}$ is defined as:
 146

$$\hat{\mathcal{E}}_{\text{bg}} = g(e_{\text{bg}}; \mathcal{G}_{\text{bg}}) = \begin{cases} 1 & \text{if } e_{\text{bg}} \text{ is a contact,} \\ 0 & \text{otherwise.} \end{cases} \quad (2)$$

Equivariance and Invariance in E(3) Space Traditional graph representations of proteins capture
 153 connectivity but ignore crucial 3D geometric information. Recently, proteins have been naturally
 154 modeled as geometric graphs that encode both topological connectivity and 3D spatial coordinates of
 155 atoms. Since molecular properties remain unchanged under rigid body transformations (rotations,
 156 translations, reflections), geometric GNN incorporate E(3)-equivariance as an inductive bias to respect
 157 these fundamental symmetries (Jiao et al., 2023).
 158

159 For a protein with coordinates $\mathbf{X} \in \mathbb{R}^{3 \times m}$ and scalar features $\mathbf{h} \in \mathbb{R}^d$, an E(3)-equivariant function
 160 f satisfies:
 161

$$f(g \cdot \mathbf{X}, \mathbf{h}) = g \cdot f(\mathbf{X}, \mathbf{h}), \quad \forall g \in \text{E}(3) \quad (3)$$

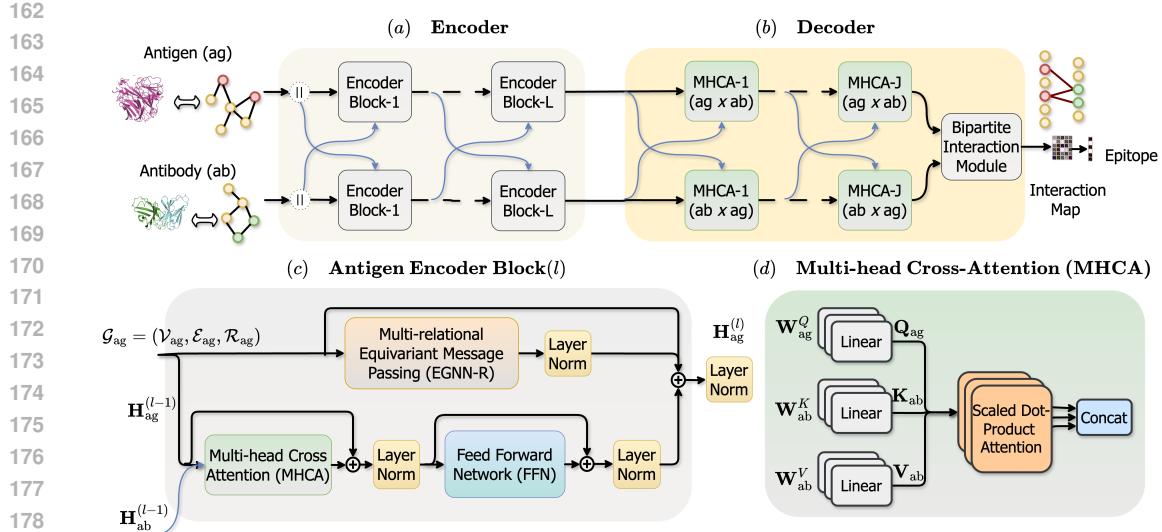


Figure 1: Overview of *EpiFormer*. The inputs are an antigen multi-relational graph $\mathcal{G}_{\text{ag}} = (\mathcal{V}_{\text{ag}}, \mathcal{E}_{\text{ag}}, \mathcal{R}_{\text{ag}})$ and an antibody multi-relational graph $\mathcal{G}_{\text{ab}} = (\mathcal{V}_{\text{ab}}, \mathcal{E}_{\text{ab}}, \mathcal{R}_{\text{ab}})$, while the outputs are the bipartite adjacency matrix and the binary epitope node labels. (a) Antigen and antibody graphs are encoded with parallel multi-relational equivariant message passing layers (EGNN-R) and cross-attention blocks. "||" is a small gating network determines the relative importance of geometric and language features for every residue. (b) A bi-directional cross-attention decoder produces the interaction map. (c) Antigen Encoder Block schematic (Antibody Encoder Block is analogous) where “ \oplus ” denotes addition. (d) An example of MHCA between antigen and antibody.

where group actions are defined as translations $g \cdot \mathbf{X} = \mathbf{X} + \mathbf{b}$ or rotations/reflections $g \cdot \mathbf{X} = \mathbf{O}\mathbf{X}$ with $\mathbf{O} \in \text{O}(3)$. This is contrast to $\text{E}(3)$ -invariant functions, which satisfy $f(g \cdot \mathbf{X}, \mathbf{h}) = f(\mathbf{X}, \mathbf{h})$, producing outputs unchanged by coordinate transformations.

3.2 EPIFORMER

In this section, we present the architecture of *EpiFormer*, an encoder-decoder framework for antibody-antigen binding-site prediction. The model receives two disjoint multi-relational residue graphs, \mathcal{G}_{ag} and \mathcal{G}_{ab} , processes them with independent encoders that produce residue-level embeddings, and then passes these embeddings to a decoder to reconstruct the bipartite adjacency matrix $\hat{\mathcal{E}}_{\text{bg}} \in \{0, 1\}^{n \times m}$. A desirable property of our proposed framework is its $\text{E}(3)$ -equivariance to address a broader range of symmetries in antigen-antibody interactions and preserve the geometry of these proteins. The overall workflow is presented in Figure 1 while the algorithm is provided in the Appendix 1.

Encoder The *EpiFormer* contains two parallel encoders with no shared parameters, one dedicated to the antigen chain and the other to the antibody chain, as shown in Figure 1 (a). Both encoders operate on heterogeneous residue graphs \mathcal{G}_{ag} and \mathcal{G}_{ab} whose nodes encode Cartesian coordinates $\mathbf{x}_i \in \mathbb{R}^3$, geometric descriptors $\mathbf{h}_i^{\text{geo}} \in \mathbb{R}^{d_{\text{geo}}}$ and PLM embeddings $\mathbf{h}_i^{\text{plm}} \in \mathbb{R}^{d_{\text{plm}}}$. Before message passing begins, a small gating network determines the relative importance of geometric and language features for every residue. The gate first concatenates the two feature vectors, applies a linear projection, and normalises the result with a softmax, $g_i = \text{Softmax}(\mathbf{W}_g[\mathbf{h}_i^{\text{geo}} \parallel \mathbf{h}_i^{\text{plm}}])$, where $\mathbf{W}_g \in \mathbb{R}^{2 \times d_h}$ is the weight matrix of the gate network with $d_h = d_{\text{geo}} + d_{\text{plm}}$. It then combines the inputs through feature-specific projections to the working width d_h :

$$\mathbf{h}_i^0 = \sum_{k \in \{\text{geo, plm}\}} g_{ik} \mathbf{W}_k \mathbf{h}_i^{(k)} \in \mathbb{R}^{d_h}. \quad (4)$$

The vector \mathbf{h}_i^0 serves as the initial node state for the first *EpiFormer* encoder block. The schematic of an *EpiFormer* block is shown in Figure 1 (c). Let $\mathbf{H}_{\text{ag}}^l \in \mathbb{R}^{n \times d_h}$ and $\mathbf{H}_{\text{ab}}^l \in \mathbb{R}^{m \times d_h}$ be the current embeddings, which are passed in parallel to their EGNN-R and MHCA layers.

216 *Relation-aware EGNN (EGNN-R) layer:* We develop a relation-aware variant of EGNN (Satorras
 217 et al., 2021a) to propagate structural and geometric information within each chain. Let $\mathbf{h}_i^\ell \in \mathbb{R}^{d_h}$
 218 and $\mathbf{x}_i^\ell \in \mathbb{R}^3$ denote the feature and coordinate of residue i after the ℓ -th EGNN-R layer. Every
 219 undirected edge $e_{i,j}$ carries a tuple $\mathbf{r}_{i,j} \subseteq \mathcal{R}$ that encodes sequential and spatial relations. With the
 220 squared distance $d_{ij} = \|\mathbf{x}_i^\ell - \mathbf{x}_j^\ell\|_2^2$ and the displacement vector $\delta_{ij} = \mathbf{x}_i^\ell - \mathbf{x}_j^\ell$, the layer performs
 221 the following computations:
 222

$$224 \quad m_{ij}^\rho = \phi_m^\rho(\mathbf{h}_i^\ell, \mathbf{h}_j^\ell, \gamma(d_{ij}), \mathbf{f}_{ij}), \quad \mathbf{h}_i^{(\ell+1)} = \mathbf{h}_i^\ell + \phi_h\left(\mathbf{h}_i^\ell, \sum_{j \in \mathcal{N}(i)} \sum_{\rho \in \mathbf{r}_{ij}} m_{ij}^\rho\right), \quad (5)$$

$$227 \quad s_{ij}^\rho = \phi_x^\rho(m_{ij}^\rho), \quad \mathbf{x}_i^{(\ell+1)} = \mathbf{x}_i^\ell + \sum_{j \in \mathcal{N}(i)} \sum_{\rho \in \mathbf{r}_{ij}} \frac{\delta_{ij}}{\sqrt{d_{ij} + \varepsilon}} s_{ij}^\rho. \quad (6)$$

230 Here, $\gamma(\cdot)$ denotes a 16-term radial basis function, \mathbf{f}_{ij} is the edge’s attribute vector, and each mapping
 231 $\phi_{\{m,x\}}^\rho$ is realized as a two-layer multilayer perceptron whose parameters are shared by all edges
 232 with the same relation label ρ , and $\varepsilon = 10^{-8}$. Specifically, we have four relation-specific message
 233 MLPs $\phi_m^\rho : \mathbb{R}^{2d_h+d_f+16} \rightarrow \mathbb{R}^{d_q}$ and coordinate MLPs $\phi_x^\rho : \mathbb{R}^{d_q} \rightarrow \mathbb{R}^3$, and a node update MLP
 234 $\phi_h : \mathbb{R}^{d_h+d_q} \rightarrow \mathbb{R}^{d_h}$ shared across all relations, where d_q represents hidden layer dimension.
 235 Applying residual connections and layer normalization produces output embeddings at layer ℓ as:

$$236 \quad \mathbf{H}_{\text{ag}}^{\text{intra}} = \{W_{\text{ag}}^\ell \mathbf{h}_i^\ell \mid v_i \in \mathcal{V}_{\text{ag}}\}, \quad \mathbf{H}_{\text{ab}}^{\text{intra}} = \{W_{\text{ab}}^\ell \mathbf{h}_j^\ell \mid v_j \in \mathcal{V}_{\text{ab}}\}, \quad (7)$$

238 where W^ℓ represents the trainable parameters for the EGNN-R layer ℓ for each *EpiFormer* encoder
 239 block and $\mathbf{H}_{\{\text{ag,ab}\}}^{\text{intra}}$ represents the output residue embeddings of antigen and antibody after passing
 240 through their respective EGNN-R layer ℓ . The layer remains $E(3)$ -equivariant by construction
 241 because the only vector quantity entering the coordinate update is the displacement δ_{ij} , while cross-
 242 attention works with rotation and translation-invariant features (Liao & Smidt, 2022) (please refer to
 243 Appendix A.1 for the formal proof).

244 *Multi-head cross-attention (MHCA) layer with feed-forward network:* In parallel to geometric mes-
 245 sage passing, each encoder block applies bidirectional multi-head cross-attention (MHCA) (Vaswani
 246 et al., 2017) to enable inter-chain communication. The MHCA mechanism shown in Figure 1 (d),
 247 produces cross-chain context representations $\tilde{\mathbf{H}}_{\text{ag}}$ and $\tilde{\mathbf{H}}_{\text{ab}}$. A learnable scalar gate α balances
 248 intra-chain geometry with cross-chain context:

$$249 \quad \mathbf{H}_{\text{ag}}^{(\ell+1)} = \mathbf{H}_{\text{ag}}^\ell + \mathbf{H}_{\text{ag}}^{\text{intra}} + \alpha_{\text{ag}} \text{FFN}(\tilde{\mathbf{H}}_{\text{ag}}), \quad \mathbf{H}_{\text{ab}}^{(\ell+1)} = \mathbf{H}_{\text{ab}}^\ell + \mathbf{H}_{\text{ab}}^{\text{intra}} + \alpha_{\text{ab}} \text{FFN}(\tilde{\mathbf{H}}_{\text{ab}}), \quad (8)$$

251 where $\alpha_{\text{ag}}, \alpha_{\text{ab}} \in \mathbb{R}^+$ are learnable parameters, $\tilde{\mathbf{H}} = \text{MHCA}(\mathbf{H})$, and FFN is a two-layer Feed
 252 Forward Network. The MHCA is detailed in Appendix A.2.

254 **Decoder** The decoder refines the residue embeddings \mathbf{H}_{ag}^L and \mathbf{H}_{ab}^L and performs bipartite interac-
 255 tion prediction. The decoder has J identical layers, each containing: (i) bidirectional MHCA with
 256 FFN, and (ii) layer normalization with residual connections, followed by a bipartite interaction head.
 257 The embeddings \mathbf{H}_{ag}^J and \mathbf{H}_{ab}^J serve as inputs to the bipartite interaction module.

258 **Bipartite interaction prediction module** The bipartite adjacency matrix is obtained by projecting
 259 the embeddings into queries and keys of width d_k in both directions, forming scaled dot-product
 260 similarities:

$$261 \quad \mathbf{S}_{\text{ag} \rightarrow \text{ab}} = \frac{(\mathbf{H}_{\text{ag}}^L \mathbf{W}_Q^{\text{out}})(\mathbf{H}_{\text{ab}}^L \mathbf{W}_K^{\text{out}})^\top}{\sqrt{d_k}}, \quad \mathbf{S}_{\text{ab} \rightarrow \text{ag}} = \frac{(\mathbf{H}_{\text{ab}}^L \mathbf{W}_Q^{\text{out}})(\mathbf{H}_{\text{ag}}^L \mathbf{W}_K^{\text{out}})^\top}{\sqrt{d_k}}. \quad (9)$$

264 The two score maps are fused via a learnable mixing vector $\mathbf{w} \in \mathbb{R}^2$ and bias $b \in \mathbb{R}$ to produce logits
 265 $\mathbf{Z} = \mathbf{w}^\top [\mathbf{S}_{\text{ag} \rightarrow \text{ab}} (\mathbf{S}_{\text{ab} \rightarrow \text{ag}})^\top] + b$, and the interaction probabilities are $\hat{\mathcal{E}}_{\text{bg}} = \sigma(\mathbf{Z}) \in \mathbb{R}^{n \times m}$.

267 3.3 JOINT OBJECTIVE

269 *EpiFormer* is trained with a joint objective that consists of a bipartite edge reconstruction loss, epitope
 270 node classification loss, and an auxiliary inter-chain distance classification objective. The overall

270 training objective is a weighted sum of these loss components:
 271

$$272 \quad \mathcal{L} = \lambda_{\text{edge}} \mathcal{L}_{\text{edge}} + \lambda_{\text{node}} \mathcal{L}_{\text{node}} + \lambda_{\text{geo}} \mathcal{L}_{\text{geo}}. \quad (10)$$

273
 274 **Edge Prediction Loss ($\mathcal{L}_{\text{edge}}$)** This loss applies positive-class-reweighted binary cross-entropy
 275 over all antigen-antibody residue pairs:

$$276 \quad \mathcal{L}_{\text{edge}} = -\frac{1}{nm} \sum_{i=1}^n \sum_{j=1}^m \left[\pi_{\text{edge}} (\mathcal{E}_{\text{bg}})_{ij} \log(\hat{\mathcal{E}}_{\text{bg}})_{ij} + (1 - (\mathcal{E}_{\text{bg}})_{ij}) \log(1 - (\hat{\mathcal{E}}_{\text{bg}})_{ij}) \right], \quad (11)$$

279 where $\mathcal{E}_{\text{bg}} \in \{0, 1\}^{n \times m}$ is the ground-truth interaction matrix per complex, and π_{edge} compensates
 280 for the extreme sparsity of positives. This loss directly supervises the bipartite interaction prediction,
 281 which serves as the foundation for deriving epitope probabilities.
 282

283 **Node Classification Loss ($\mathcal{L}_{\text{node}}$)** The node classification loss supervises epitope nodes only and
 284 combines three complementary objectives to handle class imbalance and enforce structural priors:
 285

$$286 \quad \mathcal{L}_{\text{node}} = \beta_{\text{BCE}} \mathcal{L}_{\text{BCE}}^{\text{epi}} + \beta_{\text{Dice}} \mathcal{L}_{\text{Dice}}^{\text{epi}} + \beta_{\text{sparsity}} \mathcal{L}_{\text{sparsity}}^{\text{epi}}, \quad (12)$$

287 where $\beta_{\{\cdot\}}$ weight the different terms. The probability that node v_{ag} is an epitope is derived from the
 288 bipartite interaction matrix using a top- k pooling strategy which captures the relationship between
 289 v_{ag} and nodes of the antibody :

$$290 \quad (\hat{y}_{\text{ag}})_i = \frac{1}{k} \sum_{j \in \text{top-}k(\hat{\mathcal{E}}_{\text{bg}})_i} (\hat{\mathcal{E}}_{\text{bg}})_{ij}, \quad (13)$$

294 where $(\hat{\mathcal{E}}_{\text{bg}})_{i:}$ denotes the i -th row, and k is determined using cross-validation.
 295

296 **Class-Reweighted Binary Cross-Entropy:** The primary classification loss applies positive class
 297 reweighting (with $\pi_{\text{epi}} > 1$) to address the severe class imbalance in epitope prediction:
 298

$$299 \quad \mathcal{L}_{\text{BCE}}^{\text{epi}} = -\frac{1}{n} \sum_{i=1}^n [\pi_{\text{epi}} (y_{\text{ag}})_i \log(\hat{y}_{\text{ag}})_i + (1 - (y_{\text{ag}})_i) \log(1 - (\hat{y}_{\text{ag}})_i)]. \quad (14)$$

301 **Dice Loss for Graph Segmentation:** The Dice loss treats epitope prediction as a segmentation problem
 302 which is effective for highly imbalanced image segmentation (Sudre et al., 2017):

$$303 \quad \mathcal{L}_{\text{Dice}}^{\text{epi}} = 1 - \frac{2 \sum_{i=1}^n (\hat{y}_{\text{ag}})_i (y_{\text{ag}})_i + \alpha}{\sum_{i=1}^n (\hat{y}_{\text{ag}})_i + \sum_{i=1}^n (y_{\text{ag}})_i + \alpha}, \quad (15)$$

306 where $\alpha > 0$ is a small smoothing constant for numerical stability. The Dice coefficient measures the
 307 overlap between predicted and true epitope regions, with the loss being $1 - \text{Dice}$.

308 **Per-Graph Sparsity Regularization:** The sparsity term enforces cardinality matching between pre-
 309 dicted and true epitope counts for each complex in the mini-batch:
 310

$$311 \quad \mathcal{L}_{\text{sparsity}}^{\text{epi}} = \|\hat{y}_{\text{ag}} - y_{\text{ag}}\|_1. \quad (16)$$

312 This regularizer is crucial for calibrating predictions across complexes of varying sizes.
 313

314 **Auxiliary Distance Classification Loss (\mathcal{L}_{geo})** The auxiliary geometric term provides additional su-
 315 pervision by classifying inter-chain distances into discrete bins, helping the model learn geometrically
 316 meaningful representations. The loss focuses on near-contact pairs and ignores distant residue pairs
 317 that are unlikely to interact. This auxiliary supervision encourages the model to learn distance-aware
 318 representations while still maintaining focus on the primary epitope prediction task.
 319

320 Let $\mathcal{M} = \{(i, j) : d_{ij} \leq D_{\text{max}}\}$ be the set of antigen-antibody residue pairs within the maximum
 321 distance cutoff, where d_{ij} is the Euclidean distance between residues i and j . The bins are defined by
 322 distances $\{d_0, d_1, d_2, d_3, d_4\} = \{0, 4, 8, 16, 32\}$ Å, creating $B = 4$ bins:
 323

$$323 \quad b(i, j) = \arg \max_{b \in \{1, \dots, 4\}} \mathbf{1}[d_{b-1} \leq d_{ij} < d_b]. \quad (17)$$

The network predicts per-pair distance logits $\Delta_{ij} \in \mathbb{R}^5$, but only the first $B = 4$ components $\hat{\Delta}_{ij} \in \mathbb{R}^4$ are used for pairs in \mathcal{M} , ignoring the “far” class beyond $D_{\max} = 32 \text{ \AA}$. The class probabilities are:

$$p_{ijb} = \frac{\exp(\hat{\Delta}_{ijb})}{\sum_{b'=1}^4 \exp(\hat{\Delta}_{ijb'})}. \quad (18)$$

The loss combines class balancing with distance-aware weighting:

$$\mathcal{L}_{\text{geo}} = -\frac{1}{|\mathcal{M}|} \sum_{(i,j) \in \mathcal{M}} w_{ij} \sum_{b=1}^4 \alpha_b \mathbf{1}[b(i,j) = b] \log p_{ijb}, \quad (19)$$

where $\alpha_b > 0$ are class-balance weights computed from empirical bin frequencies within \mathcal{M} and $w_{ij} > 0$ are distance weights inversely proportional to d_{ij} , normalized to unit mean over \mathcal{M} .

4 EXPERIMENTS AND RESULTS

Dataset We utilized the AsEP dataset (Liu et al., 2024), a novel benchmark dataset of antibody-antigen complexes designed specifically for epitope prediction tasks. After preprocessing, we retain 1,723 unique antibody-antigen complexes; details are in Appendix A.3.4. We excluded two complexes (5nj6_0P and 5ies_0P) from the AsEP dataset due to sequence alignment inconsistencies and unresolved residues, with the final dataset containing 1,721 complexes.

Our EDA revealed several key insights into the dataset and is shown in Figure 2. The distribution of epitope residues showed a mean of 19 ± 4.7 , while the antigen surface residues numbered in the hundreds. The contact distribution between residues in the bipartite graph had a mean of 43.7 contacts with a standard deviation of 12.8. Additionally, the dataset includes 641 unique antigens and 973 epitope groups, highlighting the diversity and complexity of the antibody-antigen interactions captured in the AsEP dataset.

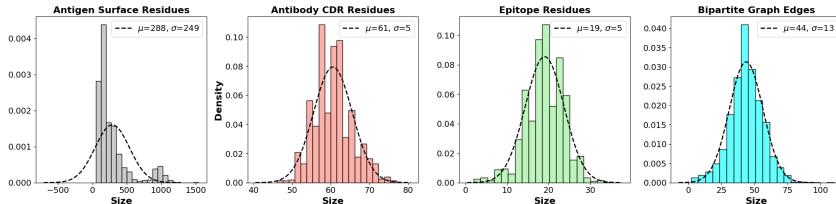


Figure 2: The size distribution of the antigen surface residues, antibody CDR residues, epitope residues, and antibody-antigen bipartite graph edges in the AsEP dataset.

Stratified Splits: We adopt two stratified splitting strategies from the AsEP benchmark dataset (Liu et al., 2024): epitope-to-antigen surface ratio split and epitope-group split. The first approach stratifies complexes by the ratio ($\# \text{epitope_nodes} / \# \text{antigen_nodes}$) to balance the class imbalance between interface and non-interface residues across train, validation, and test sets. Given that epitopes are typically limited in size (approximately 14.6 ± 4.9 residues) whereas antigen surfaces often contain several hundred residues, this stratification controls task difficulty by matching the distribution of epitope-to-surface ratios across splits.

The epitope-group split employs a different strategy by clustering complexes by antigen epitope and completely excluding test epitopes from training and validation data to evaluate model performance on novel binding sites. The dataset also includes multi-epitope antigens for which different antibodies bind distinct locations on the same antigen, and the split follows an 80/10/10 allocation by complexes. Both dataset splits result in 1,381 training complexes and 170 complexes each for validation and testing.

4.1 BASELINE COMPARISON

We evaluate the model performance using standard classification metrics such as Matthews Correlation Coefficient (MCC), Area Under the Receiver Operating Characteristic Curve (AUC-ROC), Area

Under the Precision-Recall Curve (AUPRC), accuracy, precision, recall, and F1 score. Table 1 presents the performance comparison of *EpiFormer* with the existing baseline methods for epitope prediction. We trained several baseline methods for epitope prediction on the AsEP dataset and report their results to establish a fair comparison with our model using their reported training configurations. Methods such as WALL-E (Liu et al., 2024), MIPE (Wang et al., 2024b), and EpiScan (Wang et al., 2024a) are antibody-aware while others such as EpiGraph (Choi & Kim, 2024) are not. *EpiFormer* achieves the strongest overall performance, with best AUC/AUPRC/F1/MCC across both evaluation settings (epitope-ratio and epitope-group splits), outperforming these antibody-agnostic and antibody-aware epitope prediction methods. It can also be seen that, unlike our method, the existing baselines do not provide consistent overall performance on the classification metrics. We also evaluated *EpiFormer* on the challenging epitope-group split and achieved F1/MCC scores of 0.228/0.168, compared to next best-performing baseline model, WALL-E, which achieved F1/MCC scores of 0.145/0.077. This performance gap highlights the importance of modeling both geometric constraints and dynamic antibody-antigen interactions for accurate epitope prediction. The table 2 summarizes whether each baseline conditions on antibody inputs, uses structural topology, leverages PLM representations, and adopts a graph representation, and additionally whether it incorporates explicit geometric surface/3D features, models multi-relational edges, employs equivariant GNNs, or cross-attention.

Table 1: Performance comparison of epitope prediction baseline methods with *EpiFormer* on the AsEP dataset using the epitope-to-surface ratio stratified split. The best values are represented in bold, while the second-best values are underlined.

Method	AUC	AUPRC	F1	MCC	Precision	Recall
EpiGraph	<u>0.819</u>	<u>0.279</u>	0.247	<u>0.240</u>	0.145	<u>0.852</u>
EpiScan	0.593	0.229	0.197	0.043	0.115	0.912
MIPE	0.774	0.213	0.169	0.176	<u>0.317</u>	0.248
WALLE	0.635	0.2195	<u>0.258</u>	0.210	0.235	0.422
<i>EpiFormer (ours)</i>	0.889	0.443	0.433	0.404	0.329	0.633

Table 2: Summary of features and modeling choices in baseline methods. Antibody: uses antibody information for epitope prediction; Structure: uses structure/topology as model input; PLM: uses pre-trained PLM embeddings; Graph: uses a graph representation; Geom.: explicit geometric surface/3D features; Multi-rel.: uses relation-aware/multi-edge types; E(3)-Eq.: uses an E(3)-equivariant GNN; Cross-Attn.: employs cross-attention between antibody and antigen representations.

Method	Antibody	Structure	PLM	Graph	Geom.	Multi-rel.	E(3)-Eq.	Cross-Attn.
EpiGraph	✗	✓	✓	✓	✓	✗	✗	✗
WALLE	✓	✓	✓	✓	✗	✗	✗	✗
EpiScan	✓	✗	✗	✗	✗	✗	✗	✗
MIPE	✓	✓	✗	✓	✓	✗	✗	✓
<i>EpiFormer</i>	✓	✓	✓	✓	✓	✓	✓	✓

Discussion: *EpiFormer* combines relation-aware EGNN-R message passing with early cross-attention to capture local structural detail and long-range inter-chain interactions. EGNN-R maintains E(3)-equivariance while encoding multi-relational protein structures, enabling invariant, geometry-aware representations under rigid-body transformations. Cross-attention in the encoder supports dynamic information exchange between antigen and antibody, providing binding context unavailable to antibody-agnostic approaches. The architecture jointly models intra-chain geometry and inter-chain binding dynamics using parallel processing streams, addressing a common limitation of prior methods. Together, these components allow simultaneous reasoning over structure and interaction without sacrificing equivariance. Performance on the epitope-group split suggests improved generalization to unseen binding sites, indicating that the model captures principles of antibody-antigen recognition rather than memorizing specific patterns. The method also produces interpretable antigen-antibody interaction maps: by modeling the full contact interface rather than only epitope residues, it predicts how binding is distributed across the paratope-epitope interface and highlights potential interaction hotspots.

432 **Limitations:** Despite the promising results by *EpiFormer*, there remain various ways to improve
 433 our model. Though we employ an $E(3)$ -equivariant GNN in the encoder, exploring other alternatives
 434 such as $SE(3)$ -equivariant GNNs (Fuchs et al., 2020) could improve its ability to handle global and
 435 local 3D symmetries. Our model can also be extended by performing self-supervised warm-up and
 436 transfer learning from general protein complexes that could boost its generalization capability (Zhang
 437 et al., 2022).

438
 439
 440 **4.2 ABLATIONS**
 441

442 We conducted extensive ablations to isolate the contribution of each model component (please refer
 443 to Appendix A.5 for further details). Our analysis demonstrated that multi-relational graph structures
 444 substantially exceed the performance of basic proximity graphs (Table 4). Among the tested PLMs,
 445 ESM2-650M achieved the best results, outperforming both smaller and larger parameter variants
 446 (Table 5). The cross-attention-based decoders achieve 7.5% higher AUC than simple dot-product
 447 alternatives and maintain a better precision-recall balance (Table 6). The top-2 pooling strategy
 448 achieved superior performance over hierarchical (0.836), max (0.830), mean (0.834), and larger
 449 top- k variants (Table 7). Our joint loss formulation which includes edge reconstruction, node
 450 classification, and auxiliary distance supervision achieves the best overall performance, while the
 451 failure of contrastive learning illustrates the challenges of multi-objective optimization in node
 452 classification tasks (Table 8). The most effective architectural configuration consisted of EGNN-R
 453 encoders paired with cross-attention decoders and the top-2 pooling strategy.

454
 455 **GNNs:** To assess the impact of geometric message passing on epitope prediction performance,
 456 we systematically replaced the EGNN-R layers in the encoder of *EpiFormer* with alternative
 457 GNN architectures. We evaluated standard GNN variants including graph convolutional network
 458 (GCN) (Kipf, 2016), graph isomorphism network (GIN) (Xu et al., 2018), graph attention transformer
 459 (GAT) (Veličković et al., 2017), as well as more sophisticated approaches such as relational graph
 460 convolutional network (RGCN) (Zhang et al., 2022), and relation-aware equivariant graph network
 461 (REGNN) (Wu et al., 2025). As shown in Table 3, EGNN-R achieves superior performance across
 462 all metrics, with particularly notable improvements in AUPRC (0.443 vs 0.334 for REGNN) and F1
 463 score (0.433 vs 0.343 for REGNN). While traditional GNNs like GCN, GIN, and GAT perform com-
 464 petitively but below EGNN-R, which highlights the critical importance of incorporating geometric
 465 equivariance for accurate modeling of three-dimensional protein binding interfaces.

466
 467 Table 3: Performance comparison of different GNNs used in the *EpiFormer* encoder blocks on
 468 epitope prediction tasks. The best values are represented in bold, while the second-best values are
 469 underlined.

Model	AUC	AUPRC	F1	MCC	Precision	Recall
EGNN-R	0.889 ± 0.045	0.443 ± 0.130	0.433 ± 0.014	0.404 ± 0.235	0.329 ± 0.067	0.633 ± 0.030
GAT	0.827 ± 0.006	0.308 ± 0.021	0.326 ± 0.010	0.276 ± 0.012	0.263 ± 0.016	0.435 ± 0.062
GCN	0.831 ± 0.006	0.325 ± 0.009	0.337 ± 0.010	0.290 ± 0.010	0.264 ± 0.014	0.467 ± 0.016
GIN	0.826 ± 0.007	0.310 ± 0.022	0.333 ± 0.016	0.284 ± 0.019	0.270 ± 0.004	0.437 ± 0.043
REGNN	0.833 ± 0.005	0.334 ± 0.015	0.343 ± 0.015	0.294 ± 0.015	0.276 ± 0.025	0.453 ± 0.016
RGCN	0.824 ± 0.004	0.314 ± 0.016	0.325 ± 0.008	0.276 ± 0.009	0.255 ± 0.018	0.452 ± 0.042

476
 477
 478
 479 **5 CONCLUSION**
 480

481 We presented *EpiFormer*, an encoder-decoder architecture for antibody-aware epitope prediction.
 482 Under comparable experimental conditions, *EpiFormer* outperforms prior methods on the AsEP
 483 benchmark and on the epitope-group split. Our experiments suggest that coupling multi-relational
 484 geometric message passing with cross-attention at different levels is a promising direction for
 485 antibody-specific epitope prediction. Extensive ablations demonstrate the robustness of our work.

486 REPRODUCIBILITY STATEMENT
487488 We will make the code publicly available on GitHub and provide installation scripts to address
489 libraries' complex dependency issue. We hope that this will support and accelerate future research
490 and development.
491492 REFERENCES
493

494 Mansoor Ahmed, Sarwan Ali, Avais Jan, Imdad Ullah Khan, and Murray Patterson. Improved
495 graph-based antibody-aware epitope prediction with protein language model-based embeddings.
496 *bioRxiv*, 2025. doi: 10.1101/2025.02.12.637989. 16

497 Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for
498 contrastive learning of visual representations. In *International conference on machine learning*, pp.
499 1597–1607. PMLR, 2020. 19

500 Sungjin Choi and Dongsup Kim. B cell epitope prediction by capturing spatial clustering property of
501 the epitopes using graph attention network. *Scientific Reports*, 14(1):27496, 2024. 2, 8, 17

502 Bowen Dai and Chris Bailey-Kellogg. Protein interaction interface region prediction by geometric
503 deep learning. *Bioinformatics*, 37(17):2580–2588, 03 2021. ISSN 1367-4803. 2

504 Stefan Elfwing, Eiji Uchibe, and Kenji Doya. Sigmoid-weighted linear units for neural network
505 function approximation in reinforcement learning. *Neural networks*, 107:3–11, 2018. 16

506 Yitian Fang, Yi Jiang, Leyi Wei, Qin Ma, Zhixiang Ren, Qianmu Yuan, and Dongqing Wei. Deep-
507 ProSite: structure-aware protein binding site prediction using ESMFold and pretrained language
508 model. *Bioinformatics*, 39(12):btad718, 2023. 1

509 Matthias Fey and Jan Eric Lenssen. Fast graph representation learning with pytorch geometric. *arXiv
510 preprint arXiv:1903.02428*, 2019. 16

511 Sharon Fischman and Yanay Oftran. Computational design of antibodies. *Current opinion in structural
512 biology*, 51:156–162, 2018. 1

513 Fabian Fuchs, Daniel Worrall, Volker Fischer, and Max Welling. Se (3)-transformers: 3d roto-
514 translation equivariant attention networks. *Advances in neural information processing systems*, 33:
515 1970–1981, 2020. 9

516 Thomas Hayes, Roshan Rao, Halil Akin, Nicholas J Sofroniew, Deniz Oktay, Zeming Lin, Robert
517 Verkuil, Vincent Q Tran, Jonathan Deaton, Marius Wigert, et al. Simulating 500 million years of
518 evolution with a language model. *Science*, 387(6736):850–858, 2025. 18

519 Alissa Hummer, Brennan Abanades, and Charlotte Deane. Advances in computational structure-based
520 antibody design. *Current Opinion in Structural Biology*, 74:102379, 2022. 1

521 Kanchan Jha, Sriparna Saha, and Hiteshi Singh. Prediction of protein–protein interaction using graph
522 neural networks. *Scientific Reports*, 12(1):8360, 2022. 2

523 Cheng Ji, Zixuan Huang, Qingyun Sun, Hao Peng, Xingcheng Fu, Qian Li, and Jianxin Li. Regcl:
524 rethinking message passing in graph contrastive learning. In *Proceedings of the AAAI Conference
525 on Artificial Intelligence*, volume 38, pp. 8544–8552, 2024. 20

526 Rui Jiao, Jiaqi Han, Wenbing Huang, Yu Rong, and Yang Liu. Energy-motivated equivariant pre-
527 training for 3d molecular graphs. In *Proceedings of the AAAI Conference on Artificial Intelligence*,
528 volume 37, pp. 8096–8104, 2023. 3

529 Bowen Jing, Stephan Eismann, Patricia Suriana, Raphael JL Townshend, and Ron Dror. Learning
530 from protein structure with geometric vector perceptrons. *arXiv preprint arXiv:2009.01411*, 2020.
531 14, 15

532 Sara Joubbi, Alessio Micheli, Paolo Milazzo, Giuseppe Maccari, Giorgio Ciano, Dario Cardamone,
533 and Duccio Medini. Antibody design using deep learning: from sequence and structure design to
534 affinity maturation. *Briefings in Bioinformatics*, 25(4):bbae307, 2024. 1, 2

540 John Jumper, Richard Evans, Alexander Pritzel, Tim Green, Michael Figurnov, Olaf Ronneberger,
 541 Kathryn Tunyasuvunakool, Russ Bates, Augustin Žídek, Anna Potapenko, et al. Highly accurate
 542 protein structure prediction with alphafold. *Nature*, 596(7873):583–589, 2021. 14
 543

544 TN Kipf. Semi-supervised classification with graph convolutional networks. *arXiv preprint*
 545 *arXiv:1609.02907*, 2016. 9

546 Sowmya Krishnan, Divya Sharma, Yasin Nazeer, Mayilvahanan Bose, Thangarajan Rajkumar, Guhan
 547 Jayaraman, Narayanan Madabousi, and Michael Gromiha. rAbDesFlow: a novel workflow for
 548 computational recombinant antibody design for healthcare engineering. *Antibody Therapeutics*, 7
 549 (3):256–265, 2024. 1

550 Yi-Lun Liao and Tess Smidt. Equiformer: Equivariant graph attention transformer for 3d atomistic
 551 graphs. *arXiv preprint arXiv:2206.11990*, 2022. 2, 5

552 Hyukjun Lim, Sun Kim, and Sangseon Lee. Cheapnet: Cross-attention on hierarchical representations
 553 for efficient protein-ligand binding affinity prediction. In *The Thirteenth International Conference*
 554 *on Learning Representations*, 2025. 2

555 Zeming Lin, Halil Akin, Roshan Rao, Brian Hie, Zhongkai Zhu, Wenting Lu, Nikita Smetanin,
 556 Robert Verkuil, Ori Kabeli, Yaniv Shmueli, et al. Evolutionary-scale prediction of atomic-level
 557 protein structure with a language model. *Science*, 379(6637):1123–1130, 2023. 15, 18

558 Chunan Liu, Lilian Denzler, Yihong Chen, Andrew Martin, and Brooks Paige. AsEP: Benchmarking
 559 deep learning methods for antibody-specific epitope prediction. *arXiv preprint arXiv:2407.18184*,
 560 2024. 2, 7, 8

561 Shuai Lu, Yuguang Li, Qiang Ma, Xiaofei Nan, and Shoutao Zhang. A structure-based B-cell epitope
 562 prediction model through combing local and global features. *Frontiers in immunology*, 13:890943,
 563 2022. 2

564 Richard Norman, Francesco Ambrosetti, Alexandre Bonvin, Lucy Colwell, Sebastian Kelm, Sandeep
 565 Kumar, and Konrad Krawczyk. Computational approaches to therapeutic antibody design: es-
 566 tablished methods and emerging trends. *Briefings in bioinformatics*, 21(5):1549–1567, 2020.
 567 1

568 Srivamshi Pittala and Chris Bailey-Kellogg. Learning context-aware structural representations to
 569 predict antigen and antibody binding interfaces. *Bioinformatics*, 36(13):3996–4003, 2020. 2

570 Jeffrey A Ruffolo, Lee-Shin Chu, Sai Pooja Mahajan, and Jeffrey J Gray. Fast, accurate antibody struc-
 571 ture prediction from deep learning on massive set of natural antibodies. *Nature communications*,
 572 14(1):2389, 2023. 16

573 Victor Garcia Satorras, Emiel Hoogeboom, and Max Welling. E (n) equivariant graph neural networks.
 574 In *International conference on machine learning*, pp. 9323–9332. PMLR, 2021a. 5

575 Victor Garcia Satorras, Emiel Hoogeboom, and Max Welling. E (n) equivariant graph neural networks.
 576 In *International conference on machine learning*, pp. 9323–9332. PMLR, 2021b. 2

577 Kristof T Schütt, Huziel E Sauceda, P-J Kindermans, Alexandre Tkatchenko, and K-R Müller.
 578 Schnet—a deep learning architecture for molecules and materials. *The Journal of chemical physics*,
 579 148(24), 2018. 2

580 Fabian Sievers, Andreas Wilm, David Dineen, Toby J Gibson, Kevin Karplus, Weizhong Li, Rodrigo
 581 Lopez, Hamish McWilliam, Michael Remmert, Johannes Söding, et al. Fast, scalable generation
 582 of high-quality protein multiple sequence alignments using clustal omega. *Molecular systems*
 583 *biology*, 7(1):539, 2011. 16

584 Bruna Silva, David Ascher, and Douglas Pires. epitope1d: accurate taxonomy-aware B-cell linear
 585 epitope prediction. *Briefings in Bioinformatics*, 24(3):bbad114, 2023. 2

586 Carole H Sudre, Wenqi Li, Tom Vercauteren, Sébastien Ourselin, and M Jorge Cardoso. Generalised
 587 dice overlap as a deep learning loss function for highly unbalanced segmentations. In *International*
 588 *Workshop on Deep Learning in Medical Image Analysis*, pp. 240–248. Springer, 2017. 6

594 Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz
 595 Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing*
 596 systems, 30, 2017. 5

597

598 Alice Vecchio, Andreea Deac, Pietro Liò, and Petar Veličković. Neural message passing for joint
 599 paratope-epitope prediction. *arXiv preprint arXiv:2106.00757*, 2021. 2

600 Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Lio, and Yoshua
 601 Bengio. Graph attention networks. *arXiv preprint arXiv:1710.10903*, 2017. 9

602

603 Chuan Wang, Jiangyuan Wang, Wenjun Song, Guanzheng Luo, and Taijiao Jiang. EpiScan: accurate
 604 high-throughput mapping of antibody-specific epitopes using sequence information. *NPJ Systems*
 605 *Biology and Applications*, 10(1):101, 2024a. 1, 2, 8

606

607 Zhiwei Wang, Yongkang Wang, and Wen Zhang. Improving paratope and epitope prediction by
 608 multi-modal contrastive learning and interaction informativeness estimation. *arXiv preprint*
 609 *arXiv:2405.20668*, 2024b. 8

610

611 Lirong Wu, Haitao Lin, Yufei Huang, Zhangyang Gao, Cheng Tan, Yunfan Liu, Tailin Wu, and
 612 Stan Z Li. Relation-aware equivariant graph networks for epitope-unknown antibody design
 613 and specificity optimization. In *Proceedings of the AAAI Conference on Artificial Intelligence*,
 614 volume 39, pp. 895–904, 2025. 9, 14, 15, 16, 17

615

616 Keyulu Xu, Weihua Hu, Jure Leskovec, and Stefanie Jegelka. How powerful are graph neural
 617 networks? *arXiv preprint arXiv:1810.00826*, 2018. 9

618

619 Yuansong Zeng, Zhuoyi Wei, Qianmu Yuan, Sheng Chen, Weijiang Yu, Yutong Lu, Jianzhao Gao, and
 620 Yuedong Yang. Identifying B-cell epitopes using AlphaFold2 predicted structures and pretrained
 621 language model. *Bioinformatics*, 39(4):btad187, 2023. 1, 2

622

623 Zuobai Zhang, Minghao Xu, Arian Jamasb, Vijil Chenthamarakshan, Aurelie Lozano, Payel Das,
 624 and Jian Tang. Protein representation learning by geometric structure pretraining. *arXiv preprint*
 625 *arXiv:2203.06125*, 2022. 1, 2, 9, 15, 17

626

627

628

629

630

631

632

633

634

635

636

637

638

639

640

641

642

643

644

645

646

647

648 **A APPENDIX**649 **A.1 E(3)-EQUIVARIANCE OF THE EGNN-R LAYER**650 **Theorem 1** (E(3)-equivariance of the EGNN-R layer). *Consider the EGNN-R layer in §3.2 with*
651 *updates*

652
$$m_{ij}^\rho = \phi_m^\rho(\mathbf{h}_i^\ell, \mathbf{h}_j^\ell, \gamma(d_{ij}), \mathbf{f}_{ij}), \quad (20)$$

653
$$s_{ij}^\rho = \phi_x^\rho(m_{ij}^\rho), \quad (21)$$

654
$$\mathbf{h}_i^{(\ell+1)} = \mathbf{h}_i^\ell + \phi_h\left(\mathbf{h}_i^\ell, \sum_{j \in \mathcal{N}(i)} \sum_{\rho \in \mathbf{r}_{ij}} m_{ij}^\rho\right), \quad (22)$$

655
$$\mathbf{x}_i^{(\ell+1)} = \mathbf{x}_i^\ell + \sum_{j \in \mathcal{N}(i)} \sum_{\rho \in \mathbf{r}_{ij}} \frac{\delta_{ij}}{\sqrt{d_{ij} + \varepsilon}} s_{ij}^\rho, \quad (23)$$

656 where $\delta_{ij} = \mathbf{x}_i^\ell - \mathbf{x}_j^\ell$, $d_{ij} = \|\delta_{ij}\|_2^2$, and $\varepsilon > 0$. Assume: (i) node features $\mathbf{h}_i^\ell \in \mathbb{R}^{d_h}$ are scalar
657 channels, (ii) \mathbf{h}_{ij} and \mathbf{r}_{ij} are categorical and independent of coordinates, (iii) γ is any scalar function
658 of d_{ij} , (iv) each $\phi_{\{m,x\}}^\rho$ is an MLP from scalars to scalars. Let the E(3) action be $g = (R, t)$ with
659 $R \in O(3)$ and $t \in \mathbb{R}^3$, acting as $\mathbf{x}_i^\ell \mapsto R\mathbf{x}_i^\ell + t$ and $\mathbf{h}_i^\ell \mapsto \mathbf{h}_i^\ell$. Then the layer is E(3)-equivariant:

660
$$\{\mathbf{x}_i^\ell, \mathbf{h}_i^\ell\}_{i=1}^n \mapsto \{R\mathbf{x}_i^\ell + t, \mathbf{h}_i^\ell\}_{i=1}^n \implies \{\mathbf{x}_i^{(\ell+1)}, \mathbf{h}_i^{(\ell+1)}\}_{i=1}^n \mapsto \{R\mathbf{x}_i^{(\ell+1)} + t, \mathbf{h}_i^{(\ell+1)}\}_{i=1}^n.$$

661 Consequently, any stack of such layers is E(3)-equivariant by composition.

662 *Proof.* Let $g = (R, t) \in E(3)$ act as stated. Edge data \mathbf{f}_{ij} and \mathbf{r}_{ij} are unchanged.663 *Invariants.* Relative displacement and distance transform as

664
$$d_{ij} \mapsto R\delta_{ij}, \quad d_{ij} = \|\delta_{ij}\|^2 \mapsto \|R\delta_{ij}\|^2 = d_{ij}. \quad (24)$$

665 Hence d_{ij} , $\gamma(d_{ij})$, and $(d_{ij} + \varepsilon)^{-1/2}$ are invariant scalars.666 *Scalar messages and coefficients.* Each message $m_{ij}^\rho = \phi_m^\rho(\mathbf{h}_i^\ell, \mathbf{h}_j^\ell, \gamma(d_{ij}), \mathbf{f}_{ij})$ depends only on
667 scalars that are invariant under g , so m_{ij}^ρ is invariant. Then $s_{ij}^\rho = \phi_x^\rho(m_{ij}^\rho)$ is also invariant.668 *Feature update.* The update

669
$$\mathbf{h}_i^{(\ell+1)} = \mathbf{h}_i^\ell + \phi_h\left(\mathbf{h}_i^\ell, \sum_{j \in \mathcal{N}(i)} \sum_{\rho \in \mathbf{r}_{ij}} m_{ij}^\rho\right) \quad (25)$$

670 uses only invariant scalars, so $\mathbf{h}_i^{(\ell+1)}$ is invariant. This matches the scalar action on features.671 *Coordinate update.* The increment

672
$$\Delta\mathbf{x}_i = \sum_{j \in \mathcal{N}(i)} \sum_{\rho \in \mathbf{r}_{ij}} \frac{\delta_{ij}}{\sqrt{d_{ij} + \varepsilon}} s_{ij}^\rho \quad (26)$$

673 is a sum of relative vectors scaled by invariant scalars. Under g each term becomes

674
$$\frac{\delta_{ij}}{\sqrt{d_{ij} + \varepsilon}} s_{ij}^\rho \mapsto \frac{R\delta_{ij}}{\sqrt{d_{ij} + \varepsilon}} s_{ij}^\rho = R\left(\frac{\delta_{ij}}{\sqrt{d_{ij} + \varepsilon}} s_{ij}^\rho\right), \quad (27)$$

675 so $\Delta\mathbf{x}_i \mapsto R\Delta\mathbf{x}_i$. Therefore

676
$$\mathbf{x}_i^{(\ell+1)} = \mathbf{x}_i^\ell + \Delta\mathbf{x}_i \mapsto R\mathbf{x}_i^\ell + t + R\Delta\mathbf{x}_i = R(\mathbf{x}_i^\ell + \Delta\mathbf{x}_i) + t = R\mathbf{x}_i^{(\ell+1)} + t. \quad (28)$$

677 *Composition.* The composition of equivariant maps is equivariant. Hence any stack of EGNN-R
678 layers is E(3)-equivariant. \square

702 A.2 MULTI-HEAD CROSS-ATTENTION WITH FEED-FORWARD NETWORK (MHCA)
703

704 The bidirectional multi-head cross-attention mechanism enables information exchange between
705 antigen and antibody chains. Let n_{head} be the number of heads with per-head width $d_a = d_h/n_{\text{head}}$.
706 For layer ℓ , independent linear projections produce queries, keys, and values:

$$707 \mathbf{Q}_{\text{ag}}^{\ell} = \mathbf{H}_{\text{ag}}^{(\ell-1)} \mathbf{W}_{\text{ag}}^{Q(\ell)}, \quad (29)$$

$$709 \mathbf{K}_{\text{ab}}^{\ell} = \mathbf{H}_{\text{ab}}^{(\ell-1)} \mathbf{W}_{\text{ab}}^{K(\ell)}, \quad (30)$$

$$711 \mathbf{V}_{\text{ab}}^{\ell} = \mathbf{H}_{\text{ab}}^{(\ell-1)} \mathbf{W}_{\text{ab}}^{V(\ell)}, \quad (31)$$

712 with analogous expressions for the reverse direction. After reshaping to n_{head} heads of width d_h ,
713 scaled dot-product attention computes the affinity matrices:

$$715 \mathbf{A}_{\text{ag} \leftarrow \text{ab}}^{\ell} = \text{softmax}\left(\frac{1}{\sqrt{d_h}} \mathbf{Q}_{\text{ag}}^{\ell} \mathbf{K}_{\text{ab}}^{\ell \top} + \mathbf{M}\right), \quad (32)$$

717 where \mathbf{M} is a batch mask (applied only in decoder) that assigns $-\infty$ to residue pairs from different
718 complexes. The resulting context vectors are:

$$720 \tilde{\mathbf{H}}_{\text{ag}}^{\ell} = [\mathbf{A}_{\text{ag} \leftarrow \text{ab}}^{\ell} \mathbf{V}_{\text{ab}}^{\ell}] \mathbf{W}_{O,\text{ag}}^{\ell}, \quad (33)$$

$$722 \tilde{\mathbf{H}}_{\text{ab}}^{\ell} = [\mathbf{A}_{\text{ab} \leftarrow \text{ag}}^{\ell} \mathbf{V}_{\text{ag}}^{\ell}] \mathbf{W}_{O,\text{ab}}^{\ell}. \quad (34)$$

723 Each direction then applies a feed-forward network $\text{FFN}(\mathbf{x}) = \mathbf{W}_2 \sigma(\mathbf{W}_1 \mathbf{x} + \mathbf{b}_1) + \mathbf{b}_2$ with dropout,
724 residual connections, and layer normalization.

726 A.3 GRAPH CONSTRUCTION
727728 A.3.1 NODE FEATURES
729

730 Each residue node in our protein graph incorporates two complementary information sources that
731 together provide a rich representation of both local structural properties and evolutionary context:

732 **Local geometry & physicochemistry:** Each residue $v_i \in \mathcal{V}$ is annotated with a 105-dimensional
733 geometric and biochemical feature vector $\mathbf{h}_i^{\text{geo}} \in \mathbb{R}^{d_{\text{geo}}}$ that encodes the type, position, distance,
734 direction, angle, and orientation of each residue. Such residue-level descriptors are widely employed
735 in diverse protein-related studies in structural bioinformatics (Wu et al., 2025; Jing et al., 2020;
736 Jumper et al., 2021). This vector is constructed as follows:

$$738 \mathbf{h}_i^{\text{geo}} = \left[E_{\text{type}}(v_i), E_{\text{pos}}(i), \sin(\eta_i), \cos(\eta_i), \text{RBF}(\|\mathbf{x}_{i,C_\alpha} - \mathbf{x}_{i,\xi}\|), Q_i^\top \frac{\mathbf{x}_{i,\xi} - \mathbf{x}_{i,C_\alpha}}{\|\mathbf{x}_{i,\xi} - \mathbf{x}_{i,C_\alpha}\|} \right], \quad (35)$$

741 where:

- 743 • E_{type} : Embedding for amino acid residue type (e.g., arginine, glycine).
- 744 • E_{pos} : Positional encoding of residue index in the sequence, enabling the model to distinguish
745 between identical amino acids based on their sequence context. This positional information
746 is crucial for understanding long-range dependencies and structural motifs, as amino acids
747 at different sequence positions (N-terminus vs. C-terminus, loop regions vs. secondary
748 structures) often play different functional roles even if they are the same amino acid type.
- 749 • η_i : Local backbone geometry encoded through six fundamental angles that determine how
750 the protein chain folds at each residue v_i and are encoded by their sine and cosine (12
751 scalars). Bond angles ($\alpha_i, \beta_i, \gamma_i$) describe the geometric constraints of covalent bonds,
752 while dihedral angles (ψ_i, ϕ_i, ω_i) capture the rotational freedom that gives rise to secondary
753 structures like helices and sheets.
- 754 • RBF(\cdot): Radial basis function encoding distances between C_α and other backbone atoms
755 ($\xi \in \{C_\beta, N, O\}$), with each distance represented by 16 Gaussian basis functions.

756 • $Q_i^\top \mathbf{u}_i$: Here, $Q_i \in \mathbb{R}^{3 \times 3}$ is the orthonormal rotation matrix defining the local coordinate sys-
 757 tem constructed from the C_α , C_β , and N atoms of residue i , and $\mathbf{u}_i = [\mathbf{u}_i^1, \mathbf{u}_i^2, \mathbf{u}_i^3] \in \mathbb{R}^{3 \times 3}$
 758 contains the normalized direction vectors between these atoms (e.g., $\mathbf{u}_i^1 = \frac{\mathbf{x}_{i,C_\beta} - \mathbf{x}_{i,C_\alpha}}{\|\mathbf{x}_{i,C_\beta} - \mathbf{x}_{i,C_\alpha}\|}$).
 759 The matrix product $Q_i^\top \mathbf{u}_i$ transforms these direction vectors into the local coordinate frame
 760 and is flattened to yield a 9-dimensional feature vector. Note that the oxygen atom is stored
 761 in the coordinate matrix for other calculations (like the RBF distance features), but isn't
 762 used for the local coordinate frame construction.
 763

764 • The coordinates are held in a 3×4 matrix which is used in the calculation of node and edge
 765 features.

766
$$\mathbf{X}_i = [\mathbf{x}_{i,N} \quad \mathbf{x}_{i,C_\alpha} \quad \mathbf{x}_{i,C_\beta} \quad \mathbf{x}_{i,O}] \in \mathbb{R}^{3 \times 4}, \quad \text{where } \mathbf{x}_{i,\xi} \in \mathbb{R}^3$$

 767

768 **Frozen protein-language-model (PLM) embeddings** We extract embeddings for the antigen and
 769 antibody sequences $\mathbf{z}_i^{\text{plm}} \in \mathbb{R}^{d_c}$ using pre-trained protein-language models (e.g., ESM-2 (Lin et al.,
 770 2023)) to provide the model an orthogonal information source (evolutionary + biochemical context).
 771 Since the original PLM embeddings are high-dimensional (for example, $d_c = 1280$ for ESM2-650M),
 772 we project them to a lower-dimensional representation suitable for our architecture:
 773

774
$$\mathbf{h}_i^{\text{plm}} = \mathbf{W}_{\text{plm}} \mathbf{z}_i^{\text{plm}}, \quad \text{where } \mathbf{W}_{\text{plm}} \in \mathbb{R}^{d_{\text{plm}} \times d_c}. \quad (36)$$

 775

776 Here, d_{plm} is the target dimensionality for the compressed PLM features, and \mathbf{W}_{plm} serves as a
 777 learnable bottleneck that adapts the frozen PLM representations to our specific task.
 778

A.3.2 EDGE FEATURES

779 We compute a 100-dimensional edge feature vector $\mathbf{f}_{i,j} \in \mathbb{R}^{d_h}$ that describes the spatial and sequen-
 780 tial relationship between two residues v_i and v_j . This vector integrates multiple complementary
 781 descriptors to provide a rich representation of inter-residue interactions (Jing et al., 2020) and is
 782 defined as follows:
 783

784
$$\mathbf{f}_{i,j} = \left\{ E_{\text{type}}(e_{i,j}), E_{\text{pos}}(i - j), \text{RBF}(\|\mathbf{x}_{i,C_\alpha} - \mathbf{x}_{j,\xi}\|), Q_i^\top \frac{\mathbf{x}_{j,\xi} - \mathbf{x}_{i,C_\alpha}}{\|\mathbf{x}_{j,\xi} - \mathbf{x}_{i,C_\alpha}\|}, q(Q_i^\top Q_j) \mid \xi \right\}, \quad (37)$$

 785

786 where $E_{\text{type}}(e_{i,j})$ is the one-hot encoding of relations $\mathbf{r}_{i,j}$ of length 4 between two residues, and
 787 the positional encoding $E_{\text{pos}}(i - j)$ encodes the relative sequential position sinusoidally to 16
 788 scalars. The third and fourth terms are distance and direction encodings of four backbone atoms
 789 ξ in residue v_j in the local coordinate frame Q_i . These four inter-residue distances $\{d(C_\alpha, C_\beta),$
 790 $d(C_\alpha, N), d(C_\alpha, O), d(C_\alpha, C_\alpha)\}$ are each represented by 16 Gaussian basis functions. The last term
 791 $q(Q_i^\top Q_j)$ is the quaternion representation $q(\cdot)$ of $Q_i^\top Q_j$. By integrating sequence position, local
 792 geometry, and orientation, the model understands the residue identity from global pose and enables
 793 robust generalization across structures. These node and edge features are visualized in Figure 3(a).
 794

A.3.3 EDGE RELATIONS

795 Since spatial proximity between residues alone cannot capture hydrogen bonding's directional
 796 specificity or electrostatic complementarity's charge-based selectivity, we use multi-relational edges
 797 to capture distinct interaction types (Zhang et al., 2022). By treating each relation separately, the
 798 model learns complex interaction patterns within the protein. Hence, to expand the contexts of
 799 these interactions, we divide the edges into four different types of relations $\mathcal{R} = \{\rho_1, \rho_2, \rho_3, \rho_4\}$,
 800 including (i) **sequential relations** ρ_1 and ρ_2 between two residues with relative sequential distance
 801 equal to 1 (peptide bond) and 2 (short-range torsion coupling); (ii) **spatial relations** between residues
 802 that are from the same component and spatially connected due to K -nearest neighbors (relation ρ_3
 803 that captures local packing shell) or with a Euclidean distance less than 8 Å (relation ρ_4) capturing
 804 medium-range contact between residues within the protein structure (Wu et al., 2025).
 805

806 To illustrate the importance of edge relations, consider a discontinuous epitope spanning two antigen
 807 loops: sequential edges (ρ_1, ρ_2) maintain the structural integrity of each loop, while spatial edges
 808 (ρ_3, ρ_4) capture the three-dimensional proximity between residues from different loops, enabling the
 809 model to understand how distant sequence regions come together to form a cohesive binding interface.
 We provide a schematic of edge relations in Fig. 3 (b), where each edge $e_{i,j} \in \mathcal{E}$ is associated with a

set of relations $\mathbf{r}_{i,j} \in \mathcal{R}$. Besides, two relations ρ_1 (with sequence distance equal to 1) can derive a relation ρ_2 (with sequence distance equal to 2), while an edge may connect two nodes (residues) due to both relations ρ_3 and ρ_4 .

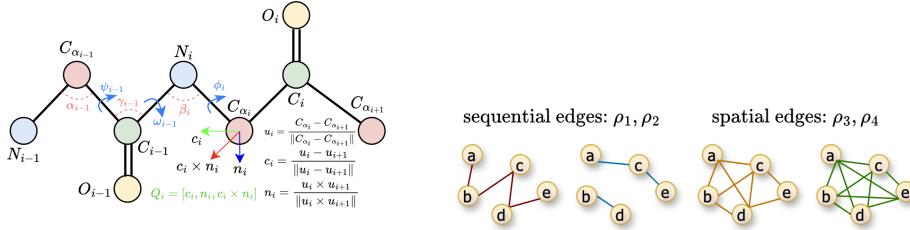


Figure 3: (a) Node and edge features encoding position, distance, direction, angle, and orientation (Figure credit: (Wu et al., 2025)). (b) Four edge relations (sequential ρ_1, ρ_2 ; spatial ρ_3, ρ_4). To avoid complexity, we visualize only some edges.

A.3.4 PREPROCESSING

For each complex, we first separated the paired antigen and antibody chains into individual structure files. We then performed sequence-structure alignment using Clustal Omega (Sievers et al., 2011) to establish correspondence between SEQRES (complete sequence) and ATOMSEQ (resolved atoms) records. This alignment generated binary masks that enable reliable mapping of sequences to structural residues (seqres2surf and seqres2cdr) while preserving the native crystallographic ordering.

For antibody chains, we applied the alignment masks to reindex heavy (H) and light (L) chains by removing insertion codes to enforce consecutive 1-based residue numbering required for graph construction. Antigen chains underwent similar processing to maintain parity between sequences and structures. This step ensures that each residue in the protein sequence corresponds exactly to its structural counterpart during the graph representation. Then, we applied solvent-accessibility filters to retain only antigen surface residues, using the original AsEP seqres2surf masks to define the node set for antigen residue graphs. The binary epitope labels were projected onto the surface ATOMSEQ via alignment masks, while paratope labels were preserved for antibody residue nodes. This surface filtering step prevents non-surface residues from confounding epitope supervision while maintaining all necessary information for cross-chain interaction modeling.

To incorporate evolutionary and semantic information, we integrated embeddings from state-of-the-art PLMs. For antigens, we extracted embeddings using the ESM model family, while, for antibodies, we incorporated AntiBERTy embeddings (Ruffolo et al., 2023; Ahmed et al., 2025), a transformer model specialized for antibody sequences, providing better functional and evolutionary context for paratope regions. These embeddings were mapped to graph nodes using the seqres2atmseq alignment masks. Finally, we used these preprocessed structures to generate HeteroData objects for the multi-relational graphs using PyTorch Geometric (Fey & Lenssen, 2019).

A.4 IMPLEMENTATION DETAILS

The model is trained with an Adam optimizer and a ReduceLROnPlateau learning-rate schedule with decoupled weight decay. The learning rate is selected from the sweep-defined range and fixed at approximately $9.1e-5$ in the best configuration. A ReduceLROnPlateau scheduler monitors validation performance and decays the learning rate on stagnation, while an early stopping with patience of 10 epochs prevents overfitting and reduces variance in final selection. We used SiLU activation functions (Elfwing et al., 2018) throughout the model because they provide stable gradients via their smooth, non-monotonic curve, which are crucial for training deep graph networks. The hyperparameter tuning was performed via a Bayesian optimization sweep in Weights & Biases to maximize validation F1 score, and the best hyperparameters were chosen within a predefined search space using bounded uniform and log-uniform distributions.

- The model weight decay was sampled log-uniformly over $[1e-5, 1e-6]$ to prevent overfitting by penalizing large weights.

- The model dropout was sampled log-uniformly over $[0.05, 0.5]$ to improve the generalizability of the model, and the best performing configuration used a dropout of 0.132.
- The number of layers in the encoder module is treated as a hyperparameter and was chosen from the set $[3, 4, 5]$ while for the decoder, the number of layers was chosen from the $[2, 3, 4]$. We experimented with different encoder hidden dimensions and the best configuration of 128 was picked from $[64, 128, 256, 512]$ across different runs.
- We also experimented with different number of attention heads for the encoder and decoder MHCA (2,4,8,16) and picked the best model with 8 attention heads.
- A batch size of 8 was chosen from $[4, 8, 16, 32]$ across different runs.
- α_{ag} and α_{ab} are initialised to 0.05

For the loss coefficients, the best run uses $\lambda_{\text{edge}} = 1.0$, $\lambda_{\text{node}} = 0.4816$, $\lambda_{\text{geo}} = 0.0514$, $\beta_{\text{BCE}} = 9.3249$, $\beta_{\text{Dice}} = 2.2966$, $\beta_{\text{sparsity}} = 0.3068$, $\pi_{\text{epi}} = 15.2856$, $\pi_{\text{edge}} = 58.7077$, label smoothing $\epsilon = 0.1$, and a distance cutoff of 32 Å for \mathcal{L}_{geo} .

- The bipartite edge positive-class weight π_{edge} for the BCE-with-logits interaction loss was sampled log-uniformly over [30, 150], accommodating variation in pairwise sparsity across complexes.
- The node objective weight λ_{node} was sampled uniformly over [0.05, 0.5], exploring the trade-off between residue supervision and the other objectives.
- The binary cross-entropy multiplier within the node objective β_{BCE} was drawn uniformly over [2, 10], spanning weak to strong emphasis on classification error.
- The Dice multiplier β_{Dice} was drawn uniformly over [0.1, 3.0], reflecting its role as a secondary calibrator under class imbalance.
- The epitope positive-class weight π_{epi} was sampled log-uniformly over [10, 60], covering roughly an order of magnitude in imbalance without biasing toward either extreme.
- The per-graph epitope count-regularizer weight β_{sparsity} was sampled uniformly over [0.05, 1.0], enabling calibration of predicted positive counts at the complex level.
- The auxiliary distance-classification weight λ_{geo} was sampled uniformly over [0.05, 0.3], with class balancing across distance bins and distance-aware pair weighting kept enabled and the maximum distance fixed at 32 Å for all trials.

The experiments were performed on an NVIDIA RTX 6000 GPU and it took around 35-60 minutes for a single hyperparameter sweeping experiment of around 50 epochs. To ensure full reproducibility of our experiments, we implement random seed management across all computational components including NumPy (`numpy.random`), Python (`random`), PyTorch (`torch.manual_seed`), and CUDA operations (`torch.cuda.manual_seed_all`), while additionally controlling worker initialization in data loaders and disabling non-deterministic algorithms (`torch.backends.cudnn.deterministic=True`).

A.5 ABLATION STUDIES

We performed ablation studies on the different protein graph representations, model components such as encoder and decoder architectures, pooling strategies, and loss functions. The results are reported as mean \pm standard deviation over 3 random seeds.

A.5.1 GRAPH CONSTRUCTION

This ablation isolates how residue-level graph design affects *EpiFormer*’s antibody-specific epitope prediction by holding node/edge features, PLM inputs, and training configuration fixed while swapping the underlying graph topology. Specifically, we compared three protein graph representations: a simple residue-only graph that collapses relations into proximity edges (Choi & Kim, 2024), a RAAD-style multi-relational graph with four edge types (sequential and spatial) (Wu et al., 2025), and a GearNet (Zhang et al., 2022) variant with seven relation types constructed to capture finer-grained structural neighborhoods. The node and edge features were fixed for all three graph types, and

918 the edge relations were only varied. This design quantifies the contribution of relation granularity
 919 and edge semantics of the proteins to the downstream performance of epitope prediction. Table 4
 920 compares the epitope prediction performance of *EpiFormer* using the three graph representations.
 921
 922

923 Table 4: Performance metrics for different protein graph representation architectures on epitope
 924 prediction tasks. All values are reported for the epitope-to-surface ratio split. The best values are
 925 represented in bold, while the second-best values are underlined.

Graph type	AUC	AUPRC	F1	MCC	Precision	Recall
Simple	<u>0.821</u>	<u>0.355</u>	0.333	<u>0.294</u>	0.240	<u>0.543</u>
GearNet	0.812	0.315	<u>0.337</u>	0.286	0.290	0.401
Multi-relational	0.888	0.443	0.433	0.404	<u>0.329</u>	<u>0.633</u>

926
 927
 928
 929
 930 We also performed experiments by using different sequence embeddings from the Evolutionary Scale
 931 Modeling (ESM) family to explore their contribution to the epitope prediction task. We used three
 932 variants of the ESM2 (Lin et al., 2023) model family (35M, 650M, and 3B parameters) as well as
 933 the newer ESM3-small (Hayes et al., 2025) model (1.4B parameters). Our experiments in Table 5
 934 show that ESM2-650M produces the best contextual features for the antigen-antibody binding site
 935 prediction task.
 936
 937

938 Table 5: Performance metrics for different PLM embeddings on the epitope prediction tasks. 4
 939 models from the Evolutionary Scale Modeling (ESM) family were used to generate embeddings for
 940 antigens, while AntiBERTy (IgFold) was used to generate embeddings for the antibodies. All values
 941 are reported for the epitope-to-surface ratio split. The best values are represented in bold, while the
 942 second-best values are underlined.

PLM	AUC	AUPRC	F1	MCC	Precision	Recall
ESM2-35M	0.815	0.330	0.334	0.283	0.287	0.399
ESM2-650M	0.888	0.443	0.433	0.404	<u>0.329</u>	<u>0.633</u>
ESM2-3B	0.826	0.331	0.349	0.300	0.285	0.449
ESM3-small	<u>0.840</u>	<u>0.374</u>	<u>0.377</u>	<u>0.330</u>	<u>0.331</u>	<u>0.437</u>

943 A.5.2 MODEL

944 We also replaced the cross-attention decoder with dot-product and dual alternatives. The **dot-product**
 945 decoder computes the interaction matrix as a plain inner product between antigen and antibody
 946 embeddings and produces a fast and parameter-free similarity score. The **dual** decoder architecture
 947 integrates two parallel processing paths: a dot-product similarity route and a sparse cross-attention
 948 mechanism, and merges their outputs via a learnable weight α . The ablation studies show lower AUC,
 949 AUPRC, and F1 metrics for dot product decoders compared to cross-attention and dual decoders.
 950 Dot-product decoding favors precision but substantially reduces recall, whereas cross-attention
 951 preserves a stronger precision-recall balance as shown in Table 6.

952 Table 6: Performance comparison of different decoder blocks for epitope prediction. The best values
 953 are represented in bold, while the second-best values are underlined.

Decoder	AUC	AUPRC	F1	MCC	Precision	Recall
Cross Attn.	0.889 ± 0.045	0.443 ± 0.130	0.433 ± 0.014	0.404 ± 0.235	0.329 ± 0.067	0.633 ± 0.030
Dot Product	0.827 ± 0.009	0.315 ± 0.034	0.326 ± 0.011	0.278 ± 0.015	0.252 ± 0.009	0.464 ± 0.053
Dual	0.834 ± 0.008	0.339 ± 0.030	0.334 ± 0.014	0.286 ± 0.017	0.266 ± 0.008	0.450 ± 0.033

954
 955 We performed ablation studies over different pooling strategies. We map the bipartite interaction
 956 matrix $\hat{\mathcal{E}}_{bg}$ to per-residue probabilities by aggregating across the partner dimension (row-wise for
 957 epitopes, column-wise for paratopes): **Max pooling** assigns the maximum interaction per residue;
 958 **Mean pooling** averages interactions over all partners; **Top- k mean pooling** averages the largest k
 959 interactions (small k , e.g., 2) to reflect a few key partners; **Noisy-OR** aggregates as $1 - \prod_j (1 - Y_{ij})$,
 960 modeling the probability that at least one partner induces a positive signal; **Softmax-attention**

972 converts interactions to attention weights via a softmax along the partner dimension and returns the
 973 weighted sum; **Hierarchical pooling** takes a convex combination of top-2 mean (local specificity)
 974 and global mean (context) with a mixing weight α . Empirically (Table 7), Top-2 pooling yields the
 975 highest AUC/AUPRC/F1, hierarchical pooling is competitive, while max/mean/softmax-attention
 976 and larger k underperform and tend to over-concentrate probability mass and impair calibration.
 977

978
 979 Table 7: Performance comparison of different pooling methods for epitope prediction. The best
 980 values are represented in bold, while the second-best values are underlined.

Pooling Method	AUC	AUPRC	F1	MCC	Precision	Recall
Hierarchical Pooling	0.836 ± 0.004	0.338 ± 0.012	0.341 ± 0.009	0.295 ± 0.006	0.268 ± 0.022	0.476 ± 0.038
Max	0.830 ± 0.005	0.321 ± 0.021	0.326 ± 0.006	0.279 ± 0.011	0.265 ± 0.029	0.441 ± 0.085
Mean	0.834 ± 0.006	0.324 ± 0.016	0.332 ± 0.004	0.283 ± 0.004	0.281 ± 0.024	0.414 ± 0.048
Pool Top-2	0.889 ± 0.045	0.443 ± 0.130	0.433 ± 0.014	0.404 ± 0.235	0.329 ± 0.067	0.633 ± 0.030
Pool Top-3	<u>0.851 ± 0.030</u>	<u>0.370 ± 0.062</u>	<u>0.370 ± 0.048</u>	<u>0.330 ± 0.059</u>	<u>0.286 ± 0.034</u>	<u>0.529 ± 0.103</u>
Pool Top-4	0.836 ± 0.008	0.342 ± 0.019	0.340 ± 0.018	0.295 ± 0.019	0.260 ± 0.020	0.493 ± 0.020
Softmax Attn.	0.832 ± 0.007	0.329 ± 0.018	0.332 ± 0.004	0.285 ± 0.005	0.256 ± 0.006	0.472 ± 0.020

A.5.3 LOSS

We performed ablations to evaluate the contribution of the loss function/s (primary, auxiliary, and regularizers) on the epitope prediction task, as shown in Table 8.

Contrastive Learning Loss ($\mathcal{L}_{\text{InfoNCE}}$) We also performed contrastive learning with the SimCLR InfoNCE (Information Noise Contrastive Estimation) loss (Chen et al., 2020) to learn discriminative representations by contrasting positive and negative residue pairs within and across protein chains. The contrastive loss combines intra-chain and inter-chain objectives:

$$\mathcal{L}_{\text{contrastive}} = \lambda_{\text{intra}} \mathcal{L}_{\text{intra}} + \lambda_{\text{inter}} \mathcal{L}_{\text{inter}}, \quad (38)$$

where λ_{intra} and λ_{inter} balance the relative importance of within-chain and cross-chain contrastive learning.

INTRA-CHAIN CONTRASTIVE LOSS ($\mathcal{L}_{\text{INTRA}}$) The intra-chain loss encourages similar representations for residues with the same label (epitope/non-epitope or paratope/non-paratope) within each protein chain:

$$\mathcal{L}_{\text{intra}} = \mathcal{L}_{\text{intra}}^{\text{ag}} + \mathcal{L}_{\text{intra}}^{\text{ab}}. \quad (39)$$

For each chain (antigen or antibody), the loss is computed as:

$$\mathcal{L}_{\text{intra}}^{\text{chain}} = -\frac{1}{|\mathcal{P}|} \sum_{i \in \mathcal{P}} \log \frac{\sum_{j \in \mathcal{P}_{i+}} \exp(\mathbf{h}_i^T \mathbf{h}_j / \tau)}{\sum_{k \in \mathcal{N}_i} \exp(\mathbf{h}_i^T \mathbf{h}_k / \tau)}, \quad (40)$$

where $\mathcal{P} = \{i : y_i = 1\}$ is the set of positive (binding) residues, $\mathcal{P}_{i+} = \{j \in \mathcal{P} : j \neq i\}$ are other positive residues sharing the same label as anchor i , \mathcal{N}_i includes all negative residues for anchor i , $\mathbf{h}_i, \mathbf{h}_j$ are L_2 -normalized residue embeddings, and τ is the temperature parameter controlling concentration.

INTER-CHAIN CONTRASTIVE LOSS ($\mathcal{L}_{\text{INTER}}$) The inter-chain loss promotes alignment between epitope and paratope representations across antigen-antibody pairs:

$$\mathcal{L}_{\text{inter}} = \mathcal{L}_{\text{ag} \rightarrow \text{ab}} + \mathcal{L}_{\text{ab} \rightarrow \text{ag}}. \quad (41)$$

The bidirectional formulation ensures symmetric learning:

$$\mathcal{L}_{\text{ag} \rightarrow \text{ab}} = -\frac{1}{|\mathcal{P}_{\text{ag}}|} \sum_{i \in \mathcal{P}_{\text{ag}}} \log \frac{\sum_{j \in \mathcal{P}_{\text{ab}}} \exp(\mathbf{h}_i^{\text{ag}T} \mathbf{h}_j^{\text{ab}} / \tau)}{\sum_{k \in \mathcal{N}_{\text{cross}}} \exp(\mathbf{h}_i^{\text{ag}T} \mathbf{h}_k / \tau)}, \quad (42)$$

where $\mathcal{P}_{\text{ag}}, \mathcal{P}_{\text{ab}}$ are epitope and paratope residue sets, $\mathcal{N}_{\text{cross}}$ includes negative residues from both chains, and the loss pulls epitope embeddings closer to paratope embeddings while pushing them away from non-binding residues. Our experiments show that contrastive learning didn't contribute

1026 to improving the classification performance. We attribute this to conflicting optimization objectives
 1027 between BCE loss and standard InfoNCE loss, a phenomenon demonstrated in a recent work (Ji et al.,
 1028 2024).

1029

1030

1031 Table 8: Performance comparison of different loss function configurations for epitope prediction. All
 1032 metrics are reported for epitope prediction tasks. The best values are represented in bold, while the
 1033 second-best values are underlined.

Loss Configuration	AUC	AUPRC	F1	MCC	Precision	Recall
\mathcal{L}_{bce}	0.822 ± 0.013	0.274 ± 0.038	0.199 ± 0.017	0.205 ± 0.022	0.111 ± 0.011	0.946 ± 0.016
$\mathcal{L}_{\text{edge}}$	0.581 ± 0.006	0.098 ± 0.002	0.142 ± 0.010	0.086 ± 0.007	0.154 ± 0.002	0.132 ± 0.016
$\mathcal{L}_{\text{bce}} + \mathcal{L}_{\text{geo}}$	0.822 ± 0.009	0.266 ± 0.025	0.205 ± 0.012	0.214 ± 0.013	0.115 ± 0.008	0.941 ± 0.019
$\mathcal{L}_{\text{bce}} + \mathcal{L}_{\text{edge}}$	0.826 ± 0.006	0.296 ± 0.018	0.220 ± 0.008	0.230 ± 0.009	0.125 ± 0.005	0.914 ± 0.011
$\mathcal{L}_{\text{bce}} + \mathcal{L}_{\text{edge}} + \mathcal{L}_{\text{dice}}$	0.818 ± 0.011	0.268 ± 0.029	0.205 ± 0.017	0.210 ± 0.021	0.115 ± 0.011	0.930 ± 0.026
$\mathcal{L}_{\text{bce}} + \mathcal{L}_{\text{edge}} + \mathcal{L}_{\text{geo}}$	0.826 ± 0.015	0.299 ± 0.050	0.214 ± 0.020	0.223 ± 0.022	0.121 ± 0.013	0.926 ± 0.033
$\mathcal{L}_{\text{edge}} + \mathcal{L}_{\text{node}} + \mathcal{L}_{\text{geo}}$	$\mathbf{0.889 \pm 0.045}$	$\mathbf{0.443 \pm 0.130}$	$\mathbf{0.433 \pm 0.014}$	$\mathbf{0.404 \pm 0.235}$	$\mathbf{0.329 \pm 0.067}$	$\mathbf{0.633 \pm 0.030}$
$\mathcal{L}_{\text{edge}} + \mathcal{L}_{\text{node}} + \mathcal{L}_{\text{geo}} + \mathcal{L}_{\text{InfoNCE}}$	0.850 ± 0.031	0.362 ± 0.064	0.361 ± 0.051	0.323 ± 0.064	0.270 ± 0.034	0.550 ± 0.111
$\mathcal{L}_{\text{bce}} + \mathcal{L}_{\text{edge}} + \mathcal{L}_{\text{InfoNCE}}$	0.837 ± 0.002	0.345 ± 0.007	0.338 ± 0.008	0.296 ± 0.005	0.254 ± 0.020	0.511 ± 0.049
$\mathcal{L}_{\text{bce}} + \mathcal{L}_{\text{edge}} + \mathcal{L}_{\text{sparsity}}$	0.835 ± 0.002	0.336 ± 0.013	0.334 ± 0.006	0.288 ± 0.007	0.270 ± 0.026	0.453 ± 0.078
$\mathcal{L}_{\text{edge}} + \mathcal{L}_{\text{node}}$	0.835 ± 0.003	0.325 ± 0.012	0.329 ± 0.001	0.283 ± 0.004	0.260 ± 0.026	0.462 ± 0.071
$\mathcal{L}_{\text{edge}} + \mathcal{L}_{\text{node}} + \mathcal{L}_{\text{InfoNCE}}$	0.829 ± 0.006	0.305 ± 0.015	0.326 ± 0.010	0.276 ± 0.012	0.261 ± 0.006	0.435 ± 0.023

1042

1043

1044 LLM USAGE CLAIM

1045

1046 LLMs were used in a limited capacity for the retrieval and discovery of related work. During paper
 1047 writing, LLMs were used for the purpose of improving grammar and wording. All technical content,
 1048 experimental design, implementation, analysis, and scientific contributions are entirely the authors'
 1049 original work.

1050

1051

1052

1053

1054

1055

1056

1057

1058

1059

1060

1061

1062

1063

1064

1065

1066

1067

1068

1069

1070

1071

1072

1073

1074

1075

1076

1077

1078

1079

Algorithm 1: EpiFormer: High-Level Architecture

Input: Antigen graph \mathcal{G}_{ag} and antibody graph \mathcal{G}_{ab} with coordinates \mathbf{X} , features $\mathbf{h}^{\text{geo}}, \mathbf{h}^{\text{plm}}$
Output: Bipartite interaction matrix $\hat{\mathcal{E}}_{\text{bg}} \in [0, 1]^{n \times m}$

// Feature Initialization

1 **foreach** $\text{chain} \in \{\text{ag}, \text{ab}\}$ **do**
2 | Apply gating network to combine geometric and PLM features;
3 | $\mathbf{h}_i^0 \leftarrow \text{Gate}(\mathbf{h}_i^{\text{geo}}, \mathbf{h}_i^{\text{plm}})$ for each residue i ;
4 **end**

// Encoder: Parallel Processing

5 **for** $\text{layer } \ell = 1$ **to** L **do**
6 | // Intra-chain geometric message passing
7 | $(\mathbf{H}_{\text{ag}}^{\text{intra}}, \mathbf{X}^{\text{ag}}) \leftarrow \text{EGNN-R}(\mathcal{G}_{\text{ag}}, \mathbf{H}_{\text{ag}}^{(\ell-1)}, \mathbf{X}^{\text{ag}});$
8 | $(\mathbf{H}_{\text{ab}}^{\text{intra}}, \mathbf{X}^{\text{ab}}) \leftarrow \text{EGNN-R}(\mathcal{G}_{\text{ab}}, \mathbf{H}_{\text{ab}}^{(\ell-1)}, \mathbf{X}^{\text{ab}});$
9 | // Inter-chain cross-attention
10 | $\tilde{\mathbf{H}}_{\text{ag}} \leftarrow \text{MHCA}(\mathbf{H}_{\text{ag}}^{\text{intra}}, \mathbf{H}_{\text{ab}}^{\text{intra}}, \mathbf{H}_{\text{ab}}^{\text{intra}});$
11 | $\tilde{\mathbf{H}}_{\text{ab}} \leftarrow \text{MHCA}(\mathbf{H}_{\text{ab}}^{\text{intra}}, \mathbf{H}_{\text{ag}}^{\text{intra}}, \mathbf{H}_{\text{ag}}^{\text{intra}});$
12 | // Combine intra-chain and cross-chain information
13 | $\mathbf{H}_{\text{ag}}^{\ell} \leftarrow \mathbf{H}_{\text{ag}}^{(\ell-1)} + \mathbf{H}_{\text{ag}}^{\text{intra}} + \alpha_{\text{ag}} \text{FFN}(\tilde{\mathbf{H}}_{\text{ag}});$
14 | $\mathbf{H}_{\text{ab}}^{\ell} \leftarrow \mathbf{H}_{\text{ab}}^{(\ell-1)} + \mathbf{H}_{\text{ab}}^{\text{intra}} + \alpha_{\text{ab}} \text{FFN}(\tilde{\mathbf{H}}_{\text{ab}});$
15 **end**

// Decoder: Cross-Attention Refinement

16 Initialize decoder embeddings: $\mathbf{H}_{\text{ag}}^{\text{dec}} \leftarrow \mathbf{H}_{\text{ag}}^L, \mathbf{H}_{\text{ab}}^{\text{dec}} \leftarrow \mathbf{H}_{\text{ab}}^L$;
17 **for** $\text{layer } \ell = 1$ **to** L **do**
18 | // Inter-chain cross-attention
19 | $\tilde{\mathbf{H}}_{\text{ag}}^{\text{dec}} \leftarrow \text{MHCA}(\mathbf{H}_{\text{ag}}^{\text{dec}}, \mathbf{H}_{\text{ab}}^{\text{dec}}, \mathbf{H}_{\text{ab}}^{\text{dec}});$
20 | $\tilde{\mathbf{H}}_{\text{ab}}^{\text{dec}} \leftarrow \text{MHCA}(\mathbf{H}_{\text{ab}}^{\text{dec}}, \mathbf{H}_{\text{ag}}^{\text{dec}}, \mathbf{H}_{\text{ag}}^{\text{dec}});$
21 | // Combine intra-chain and cross-chain information
22 | $\mathbf{H}_{\text{ag}}^{\text{dec}(\ell)} \leftarrow \mathbf{H}_{\text{ag}}^{\text{dec}(\ell-1)} + \text{FFN}(\tilde{\mathbf{H}}_{\text{ag}}^{\text{dec}});$
23 | $\mathbf{H}_{\text{ab}}^{\text{dec}(\ell)} \leftarrow \mathbf{H}_{\text{ab}}^{\text{dec}(\ell-1)} + \text{FFN}(\tilde{\mathbf{H}}_{\text{ab}}^{\text{dec}});$
24 **end**

// Bipartite Interaction Prediction

25 Compute bidirectional attention scores::;
26 $\mathbf{S}_{\text{ag} \rightarrow \text{ab}} \leftarrow \frac{(\mathbf{H}_{\text{ag}}^{\text{dec}} \mathbf{W}_Q^{\text{out}})(\mathbf{H}_{\text{ab}}^{\text{dec}} \mathbf{W}_K^{\text{out}})^{\top}}{\sqrt{d_k}};$
27 $\mathbf{S}_{\text{ab} \rightarrow \text{ag}} \leftarrow \frac{(\mathbf{H}_{\text{ab}}^{\text{dec}} \mathbf{W}'_Q^{\text{out}})(\mathbf{H}_{\text{ag}}^{\text{dec}} \mathbf{W}'_K^{\text{out}})^{\top}}{\sqrt{d_k}};$
28 Fuse scores and apply sigmoid::;
29 $\mathbf{Z} \leftarrow \mathbf{w}^{\top} [\mathbf{S}_{\text{ag} \rightarrow \text{ab}} (\mathbf{S}_{\text{ab} \rightarrow \text{ag}})^{\top}] + b;$
30 $\hat{\mathcal{E}}_{\text{bg}} \leftarrow \sigma(\mathbf{Z});$

// Epitope Extraction

31 Extract per-residue epitope probabilities via top- k pooling::;
32 $(\hat{y}_{\text{ag}})_i = \frac{1}{k} \sum_{j \in \text{top-}k} (\hat{\mathcal{E}}_{\text{bg}})_{ij};$

Function MHCA(Q, K, V):

33 $Q_h \leftarrow QW_Q^h, K_h \leftarrow KW_K^h, V_h \leftarrow VW_V^h;$ // Project per head h
34 $\alpha_{ij}^h \leftarrow \text{Softmax}_j \left(\frac{Q_{h,i} \cdot K_{h,j}^{\top}}{\sqrt{d_h}} \right);$ // Attention scores
35 $C_i^h \leftarrow \sum_j \alpha_{ij}^h V_{h,j};$ // Context vector
36 **Result:** $\text{Concat}(C^1, \dots, C^H)W_O$
37 ; // Combine heads

end

Function FFN(X):

38 $\hat{\mathcal{E}}_{\text{bg}} \leftarrow \text{SiLU}(XW_1 + b_1)W_2 + b_2;$ // $W_1 \in \mathbb{R}^{d \times d_{ff}}, W_2 \in \mathbb{R}^{d_{ff} \times d}$
39 **Result:** $\hat{\mathcal{E}}_{\text{bg}}$

end