FusionProt: Fusing Sequence and Structural Information for Unified Protein Representation Learning

Dan Kalifa
Technion Israel Institute of Technology
kalifadan@cs.technion.ac.il

Uriel Singer
Meta AI
urielsinger@meta.com

Kira Radinsky
Technion Israel Institute of Technology
kirar@cs.technion.ac.il

ABSTRACT

Accurate protein representation is vital for diverse biological and biomedical applications. While three-dimensional (3D) structural context is central to protein function, most computational approaches either ignore it or fuse it with sequence information in a single late step, yielding limited benefits. We present FusionProt, a unified representation learning framework that iteratively exchanges information between a protein language model and a graph-based structure encoder via a single learnable fusion token. This early, bidirectional conditioning preserves structural cues across layers while maintaining near-constant complexity. Across EC and GO benchmarks, FusionProt achieves state-of-the-art results, improving F_{max} by up to 3% over strong joint baselines; on mutation stability prediction it boosts AUROC by 5.1% versus the best structure model, with only 2–5% runtime overhead. We further demonstrate how analysing the specific gains in predictive capability can help focus attention and generate hypotheses about the underlying biological mechanisms.

1 INTRODUCTION

Proteins are essential for biological processes and for understanding complex mechanisms in living organisms. They are comprised of linear chains of amino acids that fold into a specific three-dimensional (3D) structure, which underscores their functional diversity and dynamic behaviors [58]. An effective understanding of proteins is essential for understanding disease mechanisms and synthetic biology and for advancing drug development [23].

Current methodologies for protein representation primarily emphasize the exploration of proteins' one-dimensional (1D) structures, specifically the relationships between amino acids. These approaches, such as ProteinBERT [1] and ESM [22, 33], often utilize text-based techniques and transformer architectures [49], which are trained on extensive protein sequence datasets. These models take an amino acid sequence as input and typically produce protein representations by averaging the representations of individual amino acids. However, this narrow focus on amino acid sequences neglects crucial protein structure details, thereby limiting the effectiveness of these methods.

The intricate 3D structure of proteins is crucial, as the conformation plays a pivotal role in determining their activities [57]. Proteins possess specific active sites for interactions with other molecules, which are defined by the 3D arrangement of amino acids. This 3D structure determines the specificity and affinity of binding interactions, aspects that a 1D representation cannot adequately capture [15, 55]. Furthermore, drug design relies on an understanding of 3D structures to identify binding sites and find molecules that modulate protein function [24, 26]. One of the state-of-the-art (SOTA) techniques today for 3D protein representation is GearNet [58]. This approach converts the 3D structure of a protein into a graph that captures its biological characteristics. Subsequently, graph neural network techniques [16, 37] are applied to this graph, facilitating the creation of comprehensive protein representations.

Recent research emphasizes the importance of comprehensive protein representation that includes both 1D and 3D structures to capture the protein's functional and interactional properties accurately. ESM-GearNet [57] was one of the first approaches to integrate these modalities. Although the study explored various fusion strategies, empirical results showed that the most effective method is using a large protein language model (PLM) such as ESM [22] to generate representations, which were then used as context for a graph encoder like GearNet [58]. Other approaches, such as SaProt [41], leverage an AlphaFold-based model [46] to reduce the 3D structure to tokens and train them along with amino acid tokens using a traditional PLM. However, these approaches are limited as they reduce one modality into context for another model, which potentially leads to the loss of critical structure information.

In this study, we introduce FusionProt (See Figure 1), a novel approach designed to learn a unified representation of the 1D and 3D structures of proteins simultaneously. Despite the vast number of proteins in nature, the number of known 3D structures remains limited [48]. To address this, we leverage an AlphaFold model [14] for accurate protein structure predictions. We introduce an innovative learnable fusion token that serves as an adaptive bridge, enabling an iterative exchange of information between a PLM and the protein's 3D structure graph. This token is integrated into the training process of both modalities, enabling seamless propagation of information and facilitating comprehensive representation through iterative learning cycles. In practice, this token is concatenated to the sequence, allowing attention mechanisms to query the unique fusion token alongside the amino acids. This process extracts and integrates valuable information, enhancing the learning of amino acid representations. Then, the fusion token is incorporated as an additional node in the graph representing the protein's 3D structure, connected to all nodes. A graph encoder (i.e., a structure model) processes this graph, generating a new representation for the fusion token, which is subsequently used in the PLM training over the amino acids. Through this iterative process, the model representations are combined to form a refined protein representation.

Unlike prior models such as ESM-GearNet [57], which perform static or one-shot fusion, such as simple concatenation after independent encoding, FusionProt introduces a dynamic mechanism where the fusion token continuously evolves through repeated interaction across layers. This design allows sequence and structure

modalities to co-adapt throughout training, resulting in richer, functionally informed representations that better capture the complexity of protein behavior.

We perform an empirical evaluation over several protein tasks, spanning a broad spectrum of biological domains, comparing numerous methods of protein representations and achieving SOTA performance across various benchmarks, with statistical significance improvements. We present ablation tests to better study the performance of the algorithm. Then, we show via a set of case studies how gains in predictive power can be harnessed to derive insights into diverse biochemical phenomena.

The contributions of this study are threefold: (1) We introduce FusionProt, a novel approach that learns a unified representation for both the 1D and 3D structures of proteins simultaneously. Our main focus is on the fusion of 1D and 3D models in an effective manner. Our method enhances the accuracy of capturing functional and interactional properties of proteins, addressing limitations of previous methods that treated these structures separately; (2) We propose a novel fusion architecture that utilizes a specialized learnable fusion token, enabling an iterative exchange of information between a PLM and the protein's 3D structure graph. This token is integrated into the training process of both modalities, enabling seamless propagation of information and facilitating comprehensive representation through iterative learning cycles. The iterative process facilitates the exchange of contextually relevant structure and sequential features, improving the model's ability to capture both 1D and 3D protein characteristics; (3) We conduct an empirical evaluation of our work over several protein tasks establishing SOTA results and presenting biological insights that further demonstrates the model's strengths. Also, we contribute our code to the community [1].

2 RELATED WORK

2.1 Sequence-based Representation Learning

Proteins are comprised of sequences of amino acids, that establish a natural analogy to tokens in natural language processing. Recently, the adoption of unsupervised deep learning techniques has become prevalent in modeling protein sequence data.

The advent of Transformers [49] has led to the development of numerous PLMs such as MSA Transformer [32], ProteinBERT [1], ProteinLM [54], ProtBERT-BFD [5], ProtTrans [6], ESM-1b [33], and ESM-2 [22] which is considered as the SOTA PLM. These models were trained on data from UniRef [42] which contains hundreds of billions of protein sequences, via masked language modeling [3].

A protein's biological function hinges on its 3D native structure [4]. However, many PLMs do not explicitly encode protein 3D structures, which are pivotal in understanding protein functions.

In this work, we aim to overcome this limitation by enhancing a PLM through novel fusion algorithms that integrate protein 3D structure models into the protein embedding. This approach seeks to capture both sequential and structure attributes of proteins, thereby advancing the capabilities of existing PLMs in biological research and applications.

[1]https://github.com/kalifadan/FusionProt

2.2 Structure-based Representation Learning

The rising success of AlphaFold [14, 38] in predicting the 3D structure of proteins has led to deeper insights into their functional roles. Moreover, the release of more than 200 million protein structures in AlphaFoldDB [47, 48] has significantly advanced the development of large-scale protein structure models [45].

Protein structures are commonly represented as graphs, where amino acids serve as the nodes. Therefore, utilizing protein structure models on these graphs is a common practice. Models such as GVP [13], CDConv [7] and GearNet [58] have shown promising results, with different learning techniques. GearNet, which incorporates a Multiview Contrastive pre-training algorithm, is considered SOTA [58] and has outperformed the IEConv model [11], which proposed to apply a learnable kernel function on edge features. Furthermore, CDConv [7] has outperformed HoloProt [39] and ProNet [50] in a recent study [25].

Foldseek [46] suggested a different approach, optimized for protein structural search, which utilizes a VQ-VAE [44] to encode protein structures into informative tokens. Then, SaProt [41] integrates residue and structure tokens during training, derived from encoding 3D protein structures using Foldseek.

In this study, we enhance the structure-based representation learning approach by integrating sequential information, thus capturing both sequential and structure attributes of proteins.

2.3 Joint Representation Learning

The integration of protein sequence-based models with protein structure models has gained popularity in recent years [17–19, 31, 53]. Early efforts, such as LM-GVP [51] or MIF-ST [56], aimed to combine PLMs with Graph Neural Networks (GNNs) [36]. More recently, ESM-GearNet [57] proposed to incorporated sequential information into distinct residue-level models such as GearNet, GVP, and CDConv [57]. SaProt-GearNet [41] also integrates GearNet and achieves similar results. However, these approaches are limited by reducing one modality into context for another model, potentially losing critical structure information.

Unlike earlier approaches, we introduce a novel fusion architecture that simultaneously learns a unified representation for both the 1D and 3D structures of proteins. To the best of our knowledge, no existing PLMs are based on an iterative fusion of structure and sequential models simultaneously.

3 METHODS

We introduce the **FusionProt** model (See Figure 1). Given a protein $p = (S, R)$, which combines its amino acid sequence S and its 3D structure R, FusionProt simultaneously learns a unified representation for both the 1D and 3D structures of the protein. We propose a specialized learnable fusion token designed to enable an iterative exchange of information between a PLM and the protein's 3D structure graph.

In this section, we formally define sequential (1D) and structure (3D) protein layers, which are used in the model as part of the protein representation learning, and present the FusionProt algorithm. Information from both layers is fused simultaneously during learning throughout the proposed specialized fusion token.

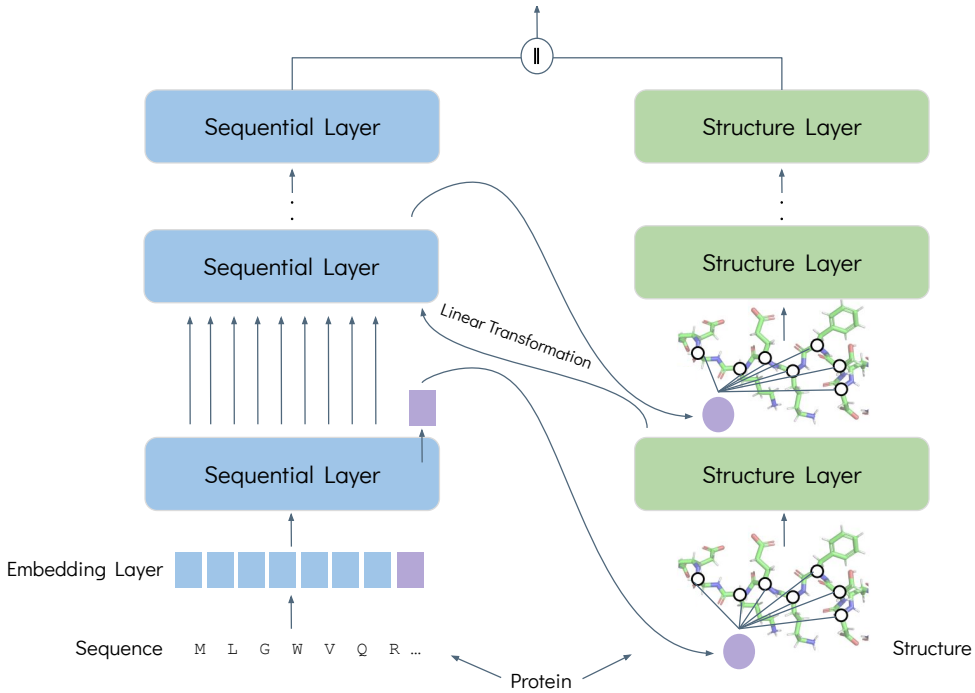

Figure 1: The FusionProt pre-training architecture. The model is trained on amino acid sequences and their corresponding 3D structures, utilizing the SOTA AlphaFold2 [14] for accurate protein structure predictions. We introduce an innovative learnable fusion token that serves as an adaptive bridge, enabling an iterative exchange of information between a PLM and the protein's 3D structure graph. This token is concatenated to the protein sequence, allowing attention mechanisms to query the unique fusion token alongside the amino acids. Then, the fusion token is incorporated as an additional node in the graph representing the protein's 3D structure, connected to all nodes. A graph encoder (i.e., a structure model) processes this graph, generating a new representation for the fusion token. This representation is subsequently used in the PLM's sequential layers. A learnable linear transformation is applied between each sequential and structure layer pair to align and adapt their distinct modality spaces. Through this iterative process, the model representations are combined to form a refined protein representation.

3.1 Protein Sequential Layer

3.1.1 Protein Sequence. The simplest level of protein structure, known as the primary structure, is a 1D structure, consisting of a linear sequence of amino acids, referred to as residues. The protein sequence exhibits similarities with natural language sequences, making the application of language models a common practice in this domain [1, 54].

Given a protein sequence $S = [s_1, s_2, ..., s_n]$, where n denotes the number of residues, the sequential model aims to capture the essential features of the protein sequence and outputs a protein representation denoted by $z = [z_1, z_2, ..., z_n] \in \mathbb{R}^{n \times D}$, where D is the embedding dimension.

3.1.2 Sequential Layer Definition. We utilize ESM-2 [22] as our sequential model, leveraging its superior performances as a PLM. Alternatively, other advanced sequential models (i.e., PLMs) can be used.

Within the context of a given sequential layer denoted by l, we denote $z^{(l)}$ as the output representation of this layer, initialized with $z_i^{(0)} = \text{Embedding}(s_i) \in \mathbb{R}^D$, where D denotes the embedding representation dimension. The layer performs the following update:

$$z^{(l)} = Attention\left(z_1^{(l-1)}, z_2^{(l-1)}, ..., z_n^{(l-1)}\right)$$

where n denotes the number of residues in the protein, and *Attention* is a multi-head self-attention layer [49].

3.2 Protein Structure Layer

3.2.1 Protein 3D Structure. A protein 3D structure is uniquely determined by its primary structure (amino acid sequences) [4]. There are only 20 standard residue (amino acid) types, each containing multiple components connected to a central carbon atom known as alpha carbon. Following GearNet [58], we use only alpha carbons to represent the main backbone structure of each protein. Therefore, we define a protein 3D structure as $R = [r_1, r_2 ... r_n] \in \mathbb{R}^{n \times 3}$, where r_i represents the Cartesian coordinates of the $i - th$ alpha carbon atom of each amino acid, and n denotes the number of residues, which is the sequence length.

3.2.2 Protein Structure Graph. We represent proteins using a multi-relational residue graph $G = (V, E, T)$, where V is the set of residues, E is the set of edges, and T is the edge types. The set of edges E consists of three directed edge types, namely sequential edges,

radius edges, and KNN edges:

$$E_{\text{seq}} = \{(i,j) \mid i,j \in V, |j-i| < d_{\text{seq}}\}$$
$$E_{\text{radius}} = \{(i,j) \mid i,j \in V, |r_j - r_i| < d_{\text{radius}}\}$$
$$E_{\text{KNN}} = \{(i,j) \mid i,j \in V, j \in \text{KNN}(i)\}$$
$$E = E_{\text{seq}} \cup E_{\text{radius}} \cup E_{\text{KNN}}$$

where $d_{\text{seq}} = 3$ defines the sequential distance threshold, $d_{\text{radius}} = 10\text{Å} = 0.1$ [nm] defines the spatial distance threshold, and $\text{KNN}(i)$ indicates the K-nearest neighbors [30] of node i with $k = 10$ (all parameters were set as reported in GearNet [58]).

3.2.3 Structure Layer Definition. We utilize GearNet [58] as the structure model, leveraging its contextual understanding of protein structures. Alternatively, other advanced structure models can be used in the algorithm (See Section 6.2).

Given a protein's 3D structure, R, we construct its corresponding protein structure graph $G = (V, E, T)$ (See Section 3.2.2). Then, the structure layer employs a relational message passing procedure, based on a relational graph convolutional neural network [37].

Within the context of a given structure layer denoted by l, we denote $u^{(l)}$ as the output representations of this layer, initialized with $u_i^{(0)} = \text{Embedding}(v_i) \in \mathbb{R}^D$, where $v_i \in V$, and D denotes the embedding representation dimension. The layer performs the following update:

$$u_i^{(l)} = u_i^{(l-1)} + \sigma\left(\sum_{t \in T} W_t \sum_{j \in \mathcal{N}_t(i)} u_j^{(l-1)}\right)$$

where $\mathcal{N}_t(i)$ is the set of neighbors of i with edge type t, $\sigma(\cdot)$ is a ReLU activation function, and the weight matrix W_t is learned per edge type t. This approach allows the model to incorporate various types of relational information, thereby enhancing its ability to learn comprehensive protein representations.

3.3 FusionProt

3.3.1 Fusion Design. We introduce a specialized learnable fusion token, designed to enable an iterative exchange of information between a PLM and the protein's 3D structure graph. This fusion token serves as a dynamic bridge, enabling iterative information exchange between the two modalities. During training (See Figure 1), the specialized fusion token, holds information from both models, the sequential and the structure. In practice, the amino acids query the associated unique fusion token to extract and integrate valuable information for learning amino acid representations via attention mechanisms. Additionally, the fusion token is incorporated as an additional node in the graph representing the protein's 3D structure. We connect it to all nodes, with a new edge type, thus enabling the structure model to learn important features from the sequential model. Hence, the token enables a novel fusion between both models, leading to an enhanced protein representation.

The choice to connect the fusion token to all nodes in the 3D graph is informed by principles from GNNs [36], where nodes exchange information to capture global and local dependencies. The fusion token is not merely an addition to the sequence or structure; rather, it mediates the dynamic interaction between the two modalities, ensuring iterative refinement of the protein representation. This approach contrasts with previous methods that simply concatenate or independently process the sequence and structure information. This novel fusion mechanism leverages multi-modal learning and graph-based information propagation. By iteratively combining the 1D protein sequence and 3D structure features, FusionProt captures both local and global dependencies, enabling the model to learn richer and more holistic representations of proteins.

Algorithm 1 FusionProt Algorithm

Input: Protein $p = (S, R)$ with sequence $S = [s_1, ..., s_n]$, and 3D structure $R = [r_1, ..., r_n] \in \mathbb{R}^{n \times 3}$, Sequential encoder with L_1 layers, Structure encoder L_2 layers.
Output: Unified protein representation h

Create a learnable fusion token $s_{n+1} = f$
Augment sequence: $S' = [s_1, ..., s_n, s_{n+1}]$
Initialize sequential embeddings: $z_i^{(0)} = \text{Embedding}(s_i)\ \forall i \in [1, ..., n+1]$
Construct graph $G = (V, E, T)$ from R
Add node v_{n+1} to G, connect $v_{n+1} \leftrightarrow v_i$ with edge type $t_f\ \forall i \in [1, ..., n]$
Set $u_i^{(0)} = \text{Embedding}(v_i)$ for all $v_i \in V$
for $l = 1$ to L_2 **do**
  for $j = 1$ to $L_1 \div L_2$ **do**
    $l_l = j + (l-1) \cdot (L_1 \div L_2)$
    $z^{(l_j)} = Attention(z_1^{(l_j-1)}, ..., z_n^{(l_j-1)}, z_{n+1}^{(l_j-1)})$
  end for
  $u_{n+1}^{(l-1)} = \text{Linear}_l(z_{n+1}^{(l_j)})$
  Update node v_{n+1} in graph G with $u_{n+1}^{(l-1)}$
  $u_i^{(l)} = u_i^{(l-1)} + \sigma\left(\sum_{t \in T} W_t \sum_{j \in \mathcal{N}_t(i)} u_j^{(l-1)}\right)\ \forall u_i \in V$
  $z_{n+1}^{(l_j)} = \text{Linear}'_l(u_{n+1}^{(l)})$
end for
Concatenate final embeddings: $h = [z_{1:n}^{(L_1)}; u_{1:n}^{(L_2)}]$
return h

3.3.2 Fusion Algorithm. We provide a formal description of the FusionProt algorithm in Algorithm 1.

Given a protein $p = (S, R)$, which combines its amino acid sequence S and its 3D structure R, FusionProt concatenates the learnable fusion token to the sequence, resulting in a new sequence $S' = [s_1, s_2, ..., s_n, f]$ with a length of $n + 1$. Then, it is used to initialize the sequential embedding layer:

$$\forall s_i \in S' : z_i^{(0)} = \text{Embedding}(s_i)$$

Within the context of a given sequential layer denoted by l, the algorithm yields a sequence representation:

$$z^{(l)} = Attention\left(z_1^{(l-1)}, z_2^{(l-1)}, ..., z_n^{(l-1)}, z_{n+1}^{(l-1)}\right)$$

where $z_{n+1}^{(l-1)}$ is the intermediate representation of the fusion token. Then, the fusion token representation is passed into the corresponding structure layer l, represented as node $n + 1$, while connecting to all nodes in the 3D structure graph, yielding a structure representation:

$$u_i^{(l)} = u_i^{(l-1)} + \sigma\left(\sum_{t\in T'} W_t \sum_{j\in\mathcal{N}_t(i)} u_j^{(l-1)}\right)$$

where $u_{n+1}^{(l-1)} = \text{Linear}_l\left(z_{n+1}^{(l-1)}\right)$, $T' = T \cup \{t_f\}$, $\mathcal{N}_{t_f}(n+1) = V$, and $\forall i \in [1, ..., n] : \mathcal{N}_{t_f}(i) = \{n+1\}$, as we connected the fusion token to all residues in the protein structure graph with a new type of edge t_f (in both directions). To bridge the gap between the sequential and structure representation space, we utilize an affine transformation Linear_l (a linear layer) per each structure layer l to project these representations into the same space. Subsequently, we set the fusion token representation $u_{n+1}^{(l)}$ as the input for the following sequential layer $l+1$, resulting in:

$$z_{n+1}^{(l)} = \text{Linear}_{l'}\left(u_{n+1}^{(l)}\right)$$

where $\text{Linear}_{l'}$ is the affine transformation in the opposite direction (i.e., structure to sequential space or vice versa).

The protein representation undergoes transformations across the structure and sequential layers, where L_1 and L_2 are the number of sequential and structure layers, respectively, and $L_1 \geq L_2$ (note that L_1 need not be equal to L_2). Our aim is to uniformly integrate 3D structure information across the L_1 sequential and L_2 structure layers. Therefore, the fusion token is passed after $L_1 \div L_2$ sequential layers into the structure model. Then, the fusion token is passed back into the sequential model after every structure layer.

Lastly, the final hidden states of both the sequential and structure layers are concatenated to form the final protein output representation, which is:

$$h = \left[z^{(L_1)}, u^{(L_2)}\right]$$

3.4 Pre-training Objective

Following GearNet [58] and ESM-GearNet [57], we use Multiview Contrastive learning as our pre-training objective, for a fair comparison. Other pre-training algorithms could also be applied. We compare different pre-training algorithms in Section 6.3.

The Multiview Contrastive objective is to preserve the similarity between correlated protein subcomponents when mapped to a lower-dimensional latent space. Therefore, for a protein structure graph, we randomly select consecutive subsequences and apply random edge masking to hide 15% of edges in the protein graph, generating diverse views. Then, we align their representations in the latent space with an InfoNCE loss [43].

4 EMPIRICAL EVALUATION

In this section, we outline our empirical setting. We performed analysis on several downstream tasks, demonstrating the value of the representation we provide in real-world bioinformatics tasks.

4.1 Protein Structure Dataset

To reduce reliance on the availability and quality of external structural data, we leverage AlphaFold2 [14], a well-known SOTA model, which enables us to generate 3D structures and eliminate the need for external sources. Furthermore, many recent high-performing methods, such as SaProt [41] and ESM-GearNet [57], also used it for structure generation, ensuring consistency in comparisons. Therefore, we use the AlphaFold protein structure database [47, 48] for pre-training, which contains 805K protein structures predicted by AlphaFold2 [14]. Alternatively, other models, such as ESMFold [21], can be applied. Since AlphaFold predictions may include local inaccuracies, we assess FusionProt's robustness to structure noise in Section 6.4.

4.2 Tasks

We evaluated our proposed method on several common downstream tasks, selected based on SOTA works of GearNet [57] and SaProt [41], to assess its effectiveness for protein representation learning. These tasks include protein annotations such as Enzyme Commission (EC) prediction, Gene Ontology (GO) prediction, and Mutation Stability Prediction. The GO prediction task includes three sub-tasks: predicting a protein's biological processes (BP), molecular functions (MF), and cellular components (CC).

To ensure the validity of our results, we validate all downstream task splits by ensuring that sequences in one set share no more than 30% Needleman–Wunsch sequence identity with any sequence in the other sets.

4.2.1 Enzyme Commission Number Prediction. Annotation of enzyme function has a wide range of real-world applications, including metagenomics, diagnosis of enzyme-deficiency-related diseases [20], and cellular metabolism [34]. This task focuses on determining enzyme function by predicting EC numbers, which characterize a protein's catalytic activity in biochemical reactions. It involves 538 binary classification problems derived from the third and fourth levels of the EC classification tree [52]. We used dataset splits from DeepFRI [9], and the evaluation metric is the $F_{\max}$ score.

4.2.2 Gene Ontology Term Prediction. The GO knowledgebase [10] provides a set of structured and controlled terms describing gene products and their molecular properties. Many real-world biological applications, such as predictions of protein-protein interactions, rely on GO term-protein annotations [12, 29]. This benchmark includes three tasks: predicting a protein's biological processes, molecular functions, and cellular components. Each task involves multiple binary classification problems based on GO term annotations. We used dataset splits from DeepFRI [9], and the evaluation metric is the $F_{\max}$ score.

4.2.3 Mutation Stability Prediction. Mutation stability prediction (MSP) is crucial in bioinformatics for understanding how genetic mutations affect protein stability, which plays a key role in disease mechanisms and drug development. This task aims to predict whether a mutation enhances a protein complex's stability. We utilize the datasets and hyperparameters from ESM-GearNet [57]. Evaluation is based on AUROC.

4.3 Baselines

We compare with numerous baseline models, including PLMs and structure models. We used ProtBERT-BFD [5], DeepFRI [9], ESM-1b [33], and ESM-2 [22] as SOTA PLM sequential models. GearNet [58], GVP [13], and CDConv [7] are used as SOTA structure models. Furthermore, we include MIF-ST [56], ESM-GearNet [57], SaProt [41] and SaProt-GearNet [41] as SOTA in joint learning of structure models with sequential models. For consistency, when using ESM-2 we used the ESM-2-650M variant [22]. In addition,

Table 1: Evaluation results on EC and GO prediction under various pre-trained baseline models. "PLM" and "Structure Info." indicate the usage of protein language models and structure information in the model, respectively. The $F_{\max}$ score is the evaluation metric. Statistically significant results with $p < 0.05$ using a t-test are marked with an asterisk (*). For FusionProt, we report the mean and standard deviation over 5 independent fine-tuning runs (with different seeds). The best result is highlighted in bold.

Method	PLM	Structure Info.	EC	GO-BP	GO-MF	GO-CC
			$F_{\max}$	$F_{\max}$	$F_{\max}$	$F_{\max}$
ProtBERT-BFD [5]	✓	✗	0.838	0.279	0.456	0.408
DeepFRI [9]	✓	✗	0.631	0.399	0.465	0.460
ESM-1b [33]	✓	✗	0.859	0.320	0.661	0.392
ESM-2 [22]	✓	✗	0.877	0.345	0.668	0.411
GVP [13]	✗	✓	0.886	0.495	0.672	0.420
CDConv [7]	✗	✓	0.820	0.453	0.654	0.479
GearNet [58]	✗	✓	0.871	0.481	0.650	0.476
MIF-ST [56]	✓	✓	0.803	0.239	0.627	0.322
ESM-GearNet [57]	✓	✓	0.886	0.512	0.670	0.495
SaProt [41]	✓	✓	0.884	0.486	0.678	0.479
SaProt-GearNet [41]	✓	✓	0.886	0.512	0.672	0.504
FusionProt	✓	✓	**0.904***±0.003	**0.524***±0.004	**0.689***±0.002	**0.518***±0.004

GearNet [58], ESM-GearNet [57], and FusionProt are pre-trained with the same objective, which is the Multiview Contrast [58] objective, as prior work showed its superior performance.

4.4 Implementation Details

4.4.1 Pre-Training Phase. We follow ESM-GearNet [57] in adopting the same training objective and model selection strategy. Our model uses a pre-trained ESM-2-650M [22] as the base PLM, with 33 layers ($L_1 = 33$), and GearNet [58] with 6 layers ($L_2 = 6$) and 512 hidden dimensions as the structure encoder. The embedding dimension D is set to 1280. Multiview Contrast [58] is used as the pre-training objective. Hyperparameters were tuned using the same search procedure and ranges reported in ESM-2 [22] and ESM-GearNet [57], with the best configuration selected based on validation performance. FusionProt is trained for 50 epochs with a learning rate of 2e-4 and a global batch size of 256 proteins. To accommodate long sequences, inputs are truncated to a maximum of 1,024 tokens. All implementations use the TorchDrug library.

4.4.2 Fine-Tuning Phase. During inference, we incorporate task-specific classification heads to generate predictions for each downstream task. Following the protocol of the recent SOTA model SaProt [41], we evaluated our model and baselines under a consistent hyperparameter tuning procedure to ensure fair comparison. For all methods, we performed tuning within the hyperparameter ranges reported in SaProt, with the best configuration for each model selected based on validation performance. Specifically, we used AdamW optimizer with $\beta_1 = 0.9$, $\beta_2 = 0.98$, and weight decay of $L_2 = 0.01$, trained with a batch size of 64 and an initial learning rate of 2e-5. All model parameters were fine-tuned until convergence, and the final checkpoint was chosen according to the highest validation score.

4.4.3 Computational Complexity. FusionProt was trained using 4× NVIDIA A100 80GB GPUs over 48 hours, totaling 192 GPU hours. It retains the underlying architecture of the base sequence and structure encoders (e.g., ESM-2 and GearNet), applying each layer only once, similar to prior baselines such as ESM-GearNet [57]. The proposed fusion mechanism introduces only lightweight additions and arithmetic operations via a token exchange module between encoders, contributing a bounded $O(1)$ overhead per layer. At inference time, FusionProt incurs a marginal runtime increase of approximately 2–5% relative to SOTA baselines, with latency rising from approximately 0.012 seconds to 0.014 seconds per 1,000 residues. Consequently, the overall computational complexity of FusionProt remains comparable to existing methods in both theoretical and empirical terms.

5 EMPIRICAL RESULTS

In this section, we present our empirical results. We report the mean and standard deviation of FusionProt's performance over three independent fine-tuning runs (with different random seeds) to ensure reproducibility. Additionally, we validate the statistical significance of performance differences, using a two-tailed paired t-test with a 95% confidence level ($p < 0.05$), comparing observations from the tested models. Significant results are marked with an asterisk (*) in the tables.

5.1 Task 1: EC Number Prediction

In Table 1, we compared the performance of FusionProt with eleven baseline methods on the EC number prediction task. In particular, FusionProt significantly outperformed all baseline methods, achieving the highest $F_{\max}$ score.

PLMs such as ProtBERT-BFD [5] or ESM-2 [22], which rely solely on sequence data, produced substantially lower $F_{\max}$ scores, compared to FusionProt. This highlights the benefit of incorporating 3D structure information into representation learning.

Similarly, FusionProt, which leverages both structure and sequential information, performed better than SOTA structure models such as GearNet [58] or CDConv [7], which trained only on 3D structure data. This indicates that while structure models are strong, they are not sufficient alone, as they require PLMs for optimal performance, probably due to their effective self-attention layers [49].

While models such as MIF-ST [56], ESM-GearNet [57], and SaProt-GearNet [41] also attempt to utilize both types of information, they reduce one modality into a context for another model, leading to a loss of critical structure information. In addition, the importance of the fusion technique is emphasized by the significant underperformance of MIF-ST compared to other fusion models. Similarly to ESM-GearNet, MIF-ST uses outputs from a pre-trained sequence-only PLM as input to a structure model.

By integrating 1D and 3D protein structure information synergistically, FusionProt attempts to capture subtle structure features that influence enzyme specificity and activity, which are crucial for EC prediction. This performance improvement is essential for real-world applications, such as the diagnosis of enzyme deficiency-related diseases [20].

5.2 Task 2: GO Term Prediction

Table 1 presents the results of the GO term prediction tasks compared to eleven baseline methods, including PLMs, structure models, and ensemble models. FusionProt showed strong performance across the board, achieving the highest $F_{\max}$ scores in all tasks, with a statistical significance.

The results demonstrate that sequential or structure information alone is insufficient, similarly to the EC prediction task (See Section 5.1). This indicates that the novel fusion of FusionProt is critical for achieving superior performance.

In the GO-BP task, which involves predicting a protein's role in biological processes, FusionProt achieves an $F_{\max}$ score of 0.524, significantly outperforming the next-best method. Biological processes often rely on long-range structural interactions and cooperative protein functions, which makes this task particularly sensitive to accurate 3D structure representations. The substantial performance gap between SaProt and SaProt-GearNet further emphasizes the critical role of structure in this setting, a factor that FusionProt leverages more effectively through early integration.

In the GO-CC task, which predicts the cellular component of proteins, FusionProt achieves an $F_{\max}$ of 0.518, outperforming all baselines including ESM-GearNet [57], probably due to the more direct connection between 3D structure characteristics and subcellular localization [8]. The prediction of cellular components is strongly based on the spatial organization of proteins within the cell, which is well represented by detailed 3D structures [40].

This reliance on structural information benefits models such as CDConv [7], although it exhibits weaker performance in EC prediction. In contrast, predicting molecular functions involves complex interactions and dynamic changes that are less directly captured by 3D structures [35], resulting in a comparatively smaller performance gap than in the other tasks.

5.3 Task 3: Mutation Stability Prediction

Table 2 reports the results for the MSP task, which evaluates a model's ability to assess the effect of amino acid substitutions on protein stability, a critical problem in protein engineering, variant effect prediction, and drug design.

Following the approach in ESM-GearNet [57], we demonstrate the robustness of our proposed method by comparing its performance with the SOTA method, specifically GVP [13] for the MSP task. Due to space limitations, we present a comparison to the SOTA and refer the reader to additional baselines in [57].

FusionProt achieves the highest AUROC among all evaluated methods, reaching 0.745 with statistical significance ($p < 0.05$), outperforming both structure-aware (GVP) and sequence-structure fused (ESM-GearNet) baselines. Notably, FusionProt improves upon ESM-GearNet by approximately 24.37%, and surpasses the current SOTA GVP by 5.1%. This highlights the effectiveness of our iterative fusion mechanism in preserving and integrating long-range sequence and structural dependencies, which are especially important for modeling stability changes that may arise from distal or context-dependent mutations.

Compared to GVP, which is limited to local spatial interactions, and ESM-GearNet, which performs shallow one-shot fusion, FusionProt allows multi-layer cross-modal refinement via a learnable token and bidirectional interaction path. This enables a more expressive and biologically grounded representation, particularly for capturing distributed compensatory effects in allosteric regions, commonly overlooked in residue-centric models.

Table 2: Evaluation results on MSP prediction across different models. "PLM" and "Structure" indicate the usage of protein language models and structure information in the model, respectively. The AUROC score is used as the evaluation metric. Statistically significant results ($p < 0.05$, t-test) are marked with an asterisk (*). For FusionProt, we report the mean and standard deviation over 5 independent fine-tuning runs (with different seeds). The best result is in bold.

Method	PLM	Structure	MSP (AUROC)
ESM-GearNet [57]	✓	✓	0.599
GVP [13] (SOTA)	✗	✓	0.709
FusionProt	✓	✓	**0.745***±0.006

6 ABLATION AND ANALYSIS

We conduct a series of ablation and diagnostic studies to evaluate the key design decisions behind FusionProt. These include varying the fusion-injection frequency, comparing different structure encoders, and testing alternative pre-training objectives. We also assess the model's robustness to noise in predicted 3D structures. Finally, we present biological insights derived from model predictions, highlighting FusionProt's practical utility in real-world tasks such as drug discovery and disease research.

6.1 Ablation on Fusion-Injection Frequency

We aimed to determine the optimal number of fusion injections in our FusionProt model, which fused information between the structure and sequential models using a specialized learnable fusion token. In our standard method, given a sequential model with L_1 layers and a structure model with L_2 layers, where $L_1 \geq L_2$, the fusion token is passed after $L_1 \div L_2$ sequential layers into the structure model (See Section 3.3.2). Then, the fusion token is passed back into the sequential model after every structure layer. In this ablation, we experimented with different injection frequencies: half the regular number (every $2 \cdot (L_1 \div L_2)$ sequential layers and two structure layers) and one-third the regular number (every $3 \cdot (L_1 \div L_2)$ sequential layers and three structure layers). The results in Table 3 indicate that the standard fusion injection frequency consistently provides the best performance across tasks. Reducing fusion injections negatively impacted performance, highlighting the importance of the fusion token. By decreasing the frequency of injections, the model's ability to integrate these two types of information is compromised, leading to suboptimal predictions.

Table 3: Ablation test for the number of fusion injections in FusionProt. Statistically significant results with $p < 0.05$ using a t-test are marked with an asterisk (*). The best result is highlighted in bold.

Number of Fusion	EC	GO-BP	GO-MF	GO-CC
	$F_{\max}$	$F_{\max}$	$F_{\max}$	$F_{\max}$
Few (One-Third)	0.876	0.511	0.659	0.498
Medium (Half)	0.882	0.514	0.669	0.504
Full (Standard)	**0.904***	**0.524***	**0.689***	**0.518***

6.2 3D Structure Models Comparison

FusionProt is designed to be flexible and modular, allowing the structure encoder component to be replaced with alternative architectures. The ablation study for various SOTA 3D structure models is detailed in Table 4. We evaluated several structure models, including GearNet [58], GVP [13], and CDConv [7]. For each structure model, we applied our fusion technique and compared their performance with FusionProt, which utilizes the GearNet [58] model (See Section 3.2.3). The GVP model replaces standard MLPs [28] in GNN [36] layers with generalized vector perceptrons, which handle scalar and geometric features as vectors that adapt to spatial rotations. In contrast, CDConv utilizes GearNet's multi-type message passing to capture sequential and spatial interactions among residues. For consistency, we adopted the hyperparameters of ESM-GearNet [57]. Specifically, we selected 8 layers of CDConv [7] with an initial hidden dimension of 256, and 5 layers of GVP [13] with scalar features of 256 and vector features of 16 dimensions. Our results demonstrate that the FusionProt model, utilizing the GearNet structure, consistently achieves superior performance compared to other models. This aligns with previous studies [57], which have also established GearNet as a SOTA approach for joint representation learning.

Table 4: Ablation test for various structure models used by FusionProt. For all models, the PLM is the ESM-2 [22]. Statistically significant results with $p < 0.05$ using a t-test are marked with an asterisk (*). The best result is highlighted in bold.

Method	EC	GO-BP	GO-MF	GO-CC
	$F_{\max}$	$F_{\max}$	$F_{\max}$	$F_{\max}$
FusionProt (GVP)	0.889	0.507	0.677	0.462
FusionProt (CDConv)	0.837	0.478	0.665	0.497
FusionProt (GearNet)	**0.904***	**0.524***	**0.689***	**0.518***

6.3 Pre-training Algorithm Ablation

Table 5 compares FusionProt, trained with Multiview Contrastive learning (Section 3.4), with other pre-training algorithms. First, we tested self-prediction methods [58], which aim to predict one part of the protein given the remaining context. Specifically, we utilized the Residue Type Prediction method, a self-supervised task that performs masked prediction on individual residues. Next, we evaluated diffusion-based methods [59], inspired by the success of diffusion models in capturing the relationship between sequences and structures. During training, noise levels are added to structures and sequences, with higher noise levels indicating more distortion. We tested the SiamDiff objective, which refines structures via torsional adjustments and noise reduction. The results show that the Multiview Contrastive approach outperformed both Residue Type Prediction (using only sequence data) and SiamDiff (which integrates sequence and structure separately). Unlike methods that treat sequence and structure separately, Multiview Contrast integrates both, aligning subsequence representations from the same protein to capture interrelated sequences and structure patterns.

Table 5: Ablation test for the FusionProt model with different pre-training algorithms. Statistically significant results with $p < 0.05$ using a t-test are marked with an asterisk (*). The best result is highlighted in bold.

Method	EC	GO-BP	GO-MF	GO-CC
	$F_{\max}$	$F_{\max}$	$F_{\max}$	$F_{\max}$
Residue Type	0.890	0.524	0.673	0.506
SiamDiff	0.874	0.509	0.645	0.512
Multiview Contrast	**0.904***	**0.524***	**0.689***	**0.518***

6.4 Robustness to Noise in 3D Structures

Robustness in protein modeling often refers to a model's ability to tolerate minor perturbations [2]—an important property given that predicted structures (e.g., from AlphaFold2) may contain local inaccuracies, especially in flexible regions. To evaluate FusionProt's robustness to such noise, we injected Gaussian perturbations into the alpha-carbon backbone coordinates of AlphaFold2-predicted structures. The noise was zero-mean, with standard deviations

(a) (b)

Figure 2: 3D structure of proteins: (a) DNA-directed RNA polymerase subunit omega and (b) D-alanine–D-alanine ligase.

ranging from 0.1 Å to 1.5 Å, simulating local structure deviations commonly observed in predicted or low-confidence regions. FusionProt maintained stable performance up to a noise level of 0.9 Å, with only a modest drop of 3–5 $F_{\max}$ points observed in the EC and GO-BP tasks beyond this threshold. These results demonstrate that FusionProt is resilient to realistic structure noise and remains effective even when relying on imperfect 3D structures, as commonly encountered in high-throughput or large-scale prediction settings.

6.5 Biological Insights

FusionProt consistently outperforms SOTA baselines across all downstream tasks (Table 1, Table 2). When such performance gains are driven by biologically meaningful inputs, such as detailed 3D structural features, they can indicate that the model is capturing relationships of functional relevance. However, we note that one must avoid conflating improved prediction with causation: a model's ability to predict biological outcomes with more precision does not in itself support causal inferences. Nevertheless, significant increases in specific classes of predictions can provide valuable hints that frame hypotheses, guide experimental questions, and motivate follow-up studies.

For FusionProt, the largest improvements are often observed in proteins whose function, stability, or interactions depend strongly on complex structural organization or long-range spatial effects. In contrast to methods such as ESM-GearNet [57] or SaProt [41], which integrate sequence and structure in a single pass, FusionProt employs an iterative fusion token mechanism that enables multiple rounds of bidirectional information exchange between sequence and structure encoders. This repeated refinement allows the model to amplify subtle, spatially localized features that might otherwise be diluted in one-shot fusion.

To better understand the nature of these improved learned protein representations, we evaluated the model's ability to predict EC numbers, as outlined in Section 4.2.1, using pre-trained representations without any fine-tuning. This approach enables an assessment of the inherent quality of the learned embeddings, independent of task-specific training. We then compared FusionProt's predictions with those of ESM-GearNet, focusing on cases with the largest discrepancies. These examples were organized into distinct biological groups based on shared mechanistic characteristics, two of which are described below.

Structural subunits for macromolecular assembly. The bacterial RNA polymerase ω subunit (Figure 2a) is small and poorly conserved, yet it contributes structurally to holoenzyme assembly by helping recruit and stabilize the largest subunit [27]. Its weak sequence signal and size make it difficult for sequence-centric models to classify. FusionProt's iterative sequence–structure fusion appears to preserve interface features at the β' contact surface that are tied to assembly and stability, whereas one-pass fusion may underweight such localized cues. Consistent with this interpretation, FusionProt assigned the correct complex-level EC number with high confidence (0.90), while ESM-GearNet did not (0.21). Because ω is noncatalytic and the protein's EC number reflects RNA polymerase activity at the holoenzyme level, we hypothesize that FusionProt is recognizing quaternary-structure context, specifically an interface-centric fingerprint of local shape, charge, and hydrophobic patterning at the ω–β' surface—rather than active-site chemistry. This implies that weakening these interface features would reduce assembly propensity and lower model confidence, providing a concrete avenue for validation.

Conserved ATP-dependent peptide ligases. D-alanine–D-alanine ligase (Figure 2b) is a well-characterized member of the ATP-grasp enzyme superfamily, catalyzing formation of the D-Ala–D-Ala dipeptide, an essential step in bacterial peptidoglycan biosynthesis. Its conserved 3D fold encloses an ATP-binding site and supports relatively simple, well-understood catalytic chemistry. Both FusionProt and ESM-GearNet correctly predicted its EC class, but FusionProt did so with higher confidence (0.88 vs. 0.72). We hypothesize that the iterative fusion mechanism further amplified local structural features of the mobile loop regions surrounding the ligand-binding pocket, which mediate precise substrate recognition and catalysis, beyond what the global fold alone conveys. This suggests that repeated sequence–structure refinement can sharpen the representation of flexible, function-critical motifs even in enzymes where the catalytic core is already highly conserved, implying that perturbations to these loops could directly influence both enzymatic activity and model confidence. Similar effects may extend to other loop-gated enzyme families, where subtle conformational changes in peripheral regions govern specificity and turnover rates.

Across categories, FusionProt's unified embeddings capture structural determinants often missed in less iterative fusion frameworks, such as allosteric loops, oligomerization interfaces, and convergent

active-site geometries. These enriched representations promise advances in function prediction, mutagenesis design, drug discovery, and the annotation of orphan proteins from metagenomic datasets.

7 DISCUSSION

In this paper, we introduced FusionProt, a novel architecture designed to learn a unified representation of the 1D and 3D structures of proteins simultaneously. FusionProt incorporates a specialized learnable fusion token that enables an iterative exchange of information between a PLM and the protein's 3D structure graph, facilitating a more comprehensive representation through iterative learning cycles between a PLM and a structure model.

We propose a novel fusion algorithm that enables effective propagation of information between a PLM and a structure model. This fusion technique outperforms previous methods, which often convert one modality into context for another model, potentially leading to the loss of crucial structural information.

We evaluated our proposed method on several protein-level tasks to assess its effectiveness in protein representation learning. These tasks include protein annotations, which are essential for real-world medical applications. FusionProt significantly outperformed all baseline methods, including SOTA models, across all tasks, with statistically significant improvements.

We perform ablation tests to examine the contribution of our novel approach of learning unified representations of both 1D and 3D structures of proteins simultaneously. The tests confirm that our approach significantly enhances FusionProt's ability to capture the intricate relationships between protein sequences and their 3D structures, leading to improved performance across various real-world tasks.

Although the proposed fusion token mechanism is tailored for bidirectional interaction between sequence and structure modalities, it is conceptually extensible to more than two information channels. For example, additional modalities such as ligand descriptors, protein dynamics, or expression context could, in principle, be integrated into the framework by assigning separate fusion tokens or adopting a shared token passed through each encoder block in sequence. We leave the systematic evaluation of such multi-modal extensions to future work.

REFERENCES

[1] Nadav Brandes, Dan Ofer, Yam Peleg, and et al. 2021. ProteinBERT: a universal deep-learning model of protein sequence and function. *Bioinformatics* 38 (2021), 2102 – 2110.

[2] Yehlin Cho, Sergey Ovchinnikov, and Christopher Frank. 2024. Enhancing Protein Design Robustness through Noise-Informed Sequence Design. In *ICML 2024 AI for Science Workshop*. https://openreview.net/forum?id=1tUTPheba8

[3] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *North American Chapter of the Association for Computational Linguistics*. https://api.semanticscholar.org/CorpusID:52967399

[4] Ken A. Dill and Justin L. MacCallum. 2012. The Protein-Folding Problem, 50 Years On. *Science* 338 (2012), 1042 – 1046. https://api.semanticscholar.org/CorpusID:5756068

[5] Ahmed Elnaggar, Michael Heinzinger, and et al. 2021. ProtTrans: Towards Cracking the Language of Lifes Code Through Self-Supervised Deep Learning and High Performance Computing. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2021), 1–1. https://doi.org/10.1109/TPAMI.2021.3095381

[6] Ahmed Elnaggar, Michael Heinzinger, and et al. 2022. ProtTrans: Toward Understanding the Language of Life Through Self-Supervised Learning. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 44 (2022), 7112–7127.

[7] Hehe Fan, Zhangyang Wang, Yi Yang, and Mohan Kankanhalli. 2023. Continuous-Discrete Convolution for Geometry-Sequence Modeling in Proteins. In *The Eleventh International Conference on Learning Representations*.

[8] Maryam Gillani and Gianluca Pollastri. 2024. Protein subcellular localization prediction tools. *Computational and Structural Biotechnology Journal* 23 (2024), 1796–1807. https://doi.org/10.1016/j.csbj.2024.04.032

[9] Vladimir Gligorijević, P. Douglas Renfrew, and et al. 2021. Structure-based protein function prediction using graph convolutional networks. *Nature Communications* 12 (2021). https://api.semanticscholar.org/CorpusID:235218243

[10] Midori A. Harris, Jennifer I. Clark, Amelia Ireland, Jane Lomax, Michael Ashburner, and et al. 2004. Gene Ontology Consortium: The Gene Ontology (GO) database and informatics resource. *Nucleic acids research* 32 Database issue (2004), D258–61. https://api.semanticscholar.org/CorpusID:22565487

[11] Pedro Hermosilla, Marco Schäfer, Matěj Lang, and et al. 2021. Intrinsic-Extrinsic Convolution and Pooling for Learning on 3D Protein Structures. *International Conference on Learning Representations* (2021).

[12] Jay J. Jiang and David W. Conrath. 1997. Semantic Similarity Based on Corpus Statistics and Lexical Taxonomy. In *ROCLING/IJCLCLP*. https://api.semanticscholar.org/CorpusID:1359050

[13] Bowen Jing, Stephan Eismann, Pratham N. Soni, and et al. 2021. Learning from Protein Structure with Geometric Vector Perceptrons. In *International Conference on Learning Representations*. https://openreview.net/forum?id=1YLJDvSx6J4

[14] John M. Jumper, Richard Evans, Alexander Pritzel, Tim Green, and et al. 2021. Highly accurate protein structure prediction with AlphaFold. *Nature* 596 (2021), 583 – 589. https://api.semanticscholar.org/CorpusID:235959867

[15] Panagiotis L. Kastritis and Alexandre M.J.J. Bonvin. 2013. On the binding affinity of macromolecular interactions: daring to ask why proteins interact. *Journal of the Royal Society Interface* 10 (2013). https://api.semanticscholar.org/CorpusID:14659814

[16] Thomas N. Kipf and Max Welling. 2017. Semi-Supervised Classification with Graph Convolutional Networks. In *International Conference on Learning Representations*. https://openreview.net/forum?id=SJU4ayYgl

[17] Eunji Ko, Seul Lee, Minseon Kim, Dongki Kim, and Sung Ju Hwang. 2024. Protein Representation Learning by Capturing Protein Sequence-Structure-Function Relationship. In *ICLR 2024 Workshop on Machine Learning for Genomics Explorations*. https://openreview.net/forum?id=1YBM58iaNt

[18] Mingchen Li, Yang Tan, Xinzhu Ma, Bozitao Zhong, Huiqun Yu, Ziyi Zhou, Wanli Ouyang, Bingxin Zhou, Pan Tan, and Liang Hong. 2024. ProSST: Protein Language Modeling with Quantized Structure and Disentangled Attention. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*. https://openreview.net/forum?id=4Z7RZixpJQ

[19] Pan Li, Xingyi Cheng, Le Song, and Eric P. Xing. 2025. Retrieval Augmented Protein Language Models for Protein Structure Prediction. In *ICML 2025 Generative AI and Biology (GenBio) Workshop*. https://openreview.net/forum?id=uuZtbiqWdn

[20] Yu Li, Sheng Wang, Ramzan Umarov, and et al. 2017. DEEPre: sequence-based enzyme EC number prediction by deep learning. *Bioinformatics* (2017). https://api.semanticscholar.org/CorpusID:3790226

[21] Zeming Lin, Halil Akin, Roshan Rao, and et al. 2022. Language models of protein sequences at the scale of evolution enable accurate structure prediction. *bioRxiv* (2022). https://doi.org/10.1101/2022.07.20.500902 arXiv:https://www.biorxiv.org/content/early/2022/07/21/2022.07.20.500902.full.pdf

[22] Zeming Lin, Halil Akin, Roshan Rao, and et al. 2023. Evolutionary-scale prediction of atomic-level protein structure with a language model. *Science* 379, 6637 (2023), 1123–1130. https://doi.org/10.1126/science.ade2574 arXiv:https://www.science.org/doi/pdf/10.1126/science.ade2574

[23] Jiajia Liu, Mengyuan Yang, Yankai Yu, Haixia Xu, Kang Li, and Xiaobo Zhou. 2024. Large language models in bioinformatics: applications and perspectives. *ArXiv* (2024). https://api.semanticscholar.org/CorpusID:266899789

[24] Meng Liu, Youzhi Luo, Kanji Uchino, Koji Maruhashi, and Shuiwang Ji. 2022. Generating 3D Molecules for Target Protein Binding. In *International Conference on Machine Learning*.

[25] Shengchao Liu, weitao Du, Yanjing Li, and et al. 2023. Symmetry-Informed Geometric Representation for Molecules, Proteins, and Crystalline Materials. In *Thirty-seventh Conference on Neural Information Processing Systems Datasets and Benchmarks Track*. https://openreview.net/forum?id=ygXSNrIU1p

[26] Shitong Luo. 2022. A 3D Generative Model for Structure-Based Drug Design. In *Neural Information Processing Systems*. https://api.semanticscholar.org/CorpusID:244906742

[27] Renjith Mathew and Dipankar Chatterji. 2006. The evolving story of the omega subunit of bacterial RNA polymerase. *Trends in microbiology* 14 10 (2006), 450–5. https://api.semanticscholar.org/CorpusID:27701631

[28] Fionn Murtagh. 1991. Multilayer perceptrons for classification and regression. *Neurocomputing* 2 (1991), 183–197. https://api.semanticscholar.org/CorpusID:7949799

[29] Catia Pesquita, Daniel Faria, Hugo Bastos, André Falcao, and Francisco Couto. 2007. Evaluating GO-based semantic similarity measures. In *Proc. 10th Annual Bio-Ontologies Meeting*, Vol. 37. 38.

[30] Leif E. Peterson. 2009. K-nearest neighbor. *Scholarpedia* 4 (2009), 1883. https://api.semanticscholar.org/CorpusID:29611121

[31] Ruijie Quan, Wenguan Wang, Fan Ma, Hehe Fan, and Yi Yang. 2024. Clustering for Protein Representation Learning. https://openreview.net/forum?id=IWfqn6Ythj

[32] Roshan M Rao, Jason Liu, Robert Verkuil, and et al. 2021. MSA Transformer. In *Proceedings of the 38th International Conference on Machine Learning*, Marina Meila and Tong Zhang (Eds.). https://proceedings.mlr.press/v139/rao21a.html

[33] Alexander Rives, Siddharth Goyal, Joshua Meier, and et al. 2019. Biological structure and function emerge from scaling unsupervised learning to 250 million protein sequences. *Proceedings of the National Academy of Sciences of the United States of America* 118 (2019).

[34] Jae Yong Ryu, Hyun Uk Kim, and Sang Yup Lee. 2019. Deep learning enables high-quality and high-throughput prediction of enzyme commission numbers. *Proceedings of the National Academy of Sciences* 116 (2019), 13996 – 14001. https://api.semanticscholar.org/CorpusID:195192023

[35] Ömer Sinan Saraç, Volkan Atalay, and Rengul Cetin-Atalay. 2010. GOPred: GO Molecular Function Prediction by Combined Classifiers. *PLoS ONE* 5 (2010). https://api.semanticscholar.org/CorpusID:1258725

[36] Franco Scarselli, Marco Gori, Ah Chung Tsoi, Markus Hagenbuchner, and Gabriele Monfardini. 2009. The Graph Neural Network Model. *IEEE Transactions on Neural Networks* 20 (2009), 61–80. https://api.semanticscholar.org/CorpusID:206756462

[37] M. Schlichtkrull, Thomas Kipf, Peter Bloem, Rianne van den Berg, Ivan Titov, and Max Welling. 2017. Modeling Relational Data with Graph Convolutional Networks. In *Extended Semantic Web Conference*. https://api.semanticscholar.org/CorpusID:5458500

[38] Andrew W. Senior, Richard Evans, John M. Jumper, and et al. 2020. Improved protein structure prediction using potentials from deep learning. *Nature* 577 (2020), 706–710.

[39] Vignesh Ram Somnath, Charlotte Bunne, and Andreas Krause. 2022. Multi-Scale Representation Learning on Proteins. In *Neural Information Processing Systems*. https://api.semanticscholar.org/CorpusID:236317531

[40] Bosheng Song, Xiaoyan Luo, Xiaoli Luo, Yuansheng Liu, Zhangming Niu, and Xiangxiang Zeng. 2022. Learning spatial structures of proteins improves protein–protein interaction prediction. *Briefings in Bioinformatics* 23, 2 (01 2022), bbab558. https://doi.org/10.1093/bib/bbab558 arXiv:https://academic.oup.com/bib/article-pdf/23/2/bbab558/42805821/bbab558.pdf

[41] Jin Su, Chenchen Han, Yuyang Zhou, Junjie Shan, Xibin Zhou, and Fajie Yuan. 2024. SaProt: Protein Language Modeling with Structure-aware Vocabulary. In *The Twelfth International Conference on Learning Representations*. https://openreview.net/forum?id=6MRm3G4NiU

[42] Baris E. Suzek, Hongzhan Huang, Peter B. McGarvey, Raja Mazumder, and Cathy H. Wu. 2007. UniRef: comprehensive and non-redundant UniProt reference clusters. *Bioinformatics* 23 10 (2007), 1282–8.

[43] Aäron van den Oord, Yazhe Li, and Oriol Vinyals. 2018. Representation Learning with Contrastive Predictive Coding. *ArXiv* abs/1807.03748 (2018). https://api.semanticscholar.org/CorpusID:49670925

[44] Aaron van den Oord, Oriol Vinyals, and Koray Kavukcuoglu. 2017. Neural discrete representation learning. In *Proceedings of the 31st International Conference on Neural Information Processing Systems* (Long Beach, California, USA) *(NIPS'17)*. Curran Associates Inc., Red Hook, NY, USA, 6309–6318.

[45] Karel van der Weg, Erinc Merdivan, Marie Piraud, and Holger Gohlke. 2025. TopEC: prediction of Enzyme Commission classes by 3D graph neural networks and localized 3D protein descriptor. *Nature Communications* 16 (2025). https://api.semanticscholar.org/CorpusID:277148295

[46] Michel van Kempen, Stephanie Kim, Charlotte Tumescheit, and et al. 2022. Fast and accurate protein structure search with Foldseek. *Nature Biotechnology* 42 (2022), 243 – 246. https://api.semanticscholar.org/CorpusID:257837735

[47] Mihály Váradi, Stephen Anyango, Mandar S. Deshpande, and et al. 2021. AlphaFold Protein Structure Database: massively expanding the structural coverage of protein-sequence space with high-accuracy models. *Nucleic Acids Research* 50 (2021), D439 – D444. https://api.semanticscholar.org/CorpusID:245770129

[48] Mihály Váradi, Damian Bertoni, Paulyna Magaña, and et al. 2023. AlphaFold Protein Structure Database in 2024: providing structure coverage for over 214 million protein sequences. *Nucleic Acids Research* 52 (2023), D368 – D375. https://api.semanticscholar.org/CorpusID:265042036

[49] Ashish Vaswani, Noam M. Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. Attention is All you Need. In *Neural Information Processing Systems*. https://api.semanticscholar.org/CorpusID:13756489

[50] Limei Wang, Haoran Liu, Yi Liu, Jerry Kurtin, and Shuiwang Ji. 2022. Learning Hierarchical Protein Representations via Complete 3D Graph Networks. In *International Conference on Learning Representations*. https://api.semanticscholar.org/CorpusID:257364758

[51] Zichen Wang, Steven A. Combs, Ryan Brand, and et al. 2022. LM-GVP: an extensible sequence and structure informed deep learning framework for protein property prediction. *Scientific Reports* 12 (2022). https://api.semanticscholar.org/CorpusID:248415326

[52] Edwin C. Webb. 1992. Enzyme nomenclature 1992. Recommendations of the Nomenclature Committee of the International Union of Biochemistry and Molecular Biology on the Nomenclature and Classification of Enzymes. https://api.semanticscholar.org/CorpusID:83964422

[53] Ruidong Wu, Fan Ding, Rui Wang, Rui Shen, Xiwen Zhang, Shitong Luo, Chenpeng Su, Zuofan Wu, Qi Xie, Bonnie Berger, Jianzhu Ma, and Jian Peng. 2022. High-resolution de novo structure prediction from primary sequence. *bioRxiv* (2022). https://doi.org/10.1101/2022.07.21.500999 arXiv:https://www.biorxiv.org/content/early/2022/07/22/2022.07.21.500999.full.pdf

[54] Yijia Xiao, Jiezhong Qiu, Ziang Li, Chang-Yu Hsieh, and Jie Tang. 2021. Modeling Protein Using Large-scale Pretrain Language Model. *ArXiv* abs/2108.07435 (2021).

[55] Zhiqiang Yan, Liyong Guo, Liang Hu, and Jin Wang. 2013. Specificity and affinity quantification of protein-protein interactions. *Bioinformatics* 29 9 (2013), 1127–33. https://api.semanticscholar.org/CorpusID:1339389

[56] Kevin K Yang, Niccolò Zanichelli, and Hugh Yeh. 2022. Masked inverse folding with sequence transfer for protein representation learning. *Protein Engineering, Design and Selection* 36 (10 2022), gzad015. https://doi.org/10.1093/protein/gzad015 arXiv:https://academic.oup.com/peds/article-pdf/doi/10.1093/protein/gzad015/53269438/gzad015.pdf

[57] Zuobai Zhang, Chuanrui Wang, Minghao Xu, and et al. 2023. A Systematic Study of Joint Representation Learning on Protein Sequences and Structures. https://api.semanticscholar.org/CorpusID:257496831

[58] Zuobai Zhang, Minghao Xu, Arian Rokkum Jamasb, and et al. 2023. Protein Representation Learning by Geometric Structure Pretraining. In *The Eleventh International Conference on Learning Representations*. https://openreview.net/forum?id=to3qCB3tOh9

[59] Zuobai Zhang, Minghao Xu, Aurélie Lozano, and et al. 2023. Pre-Training Protein Encoder via Siamese Sequence-Structure Diffusion Trajectory Prediction. In *Advances in Neural Information Processing Systems*.