
MULTIMODAL PRE-TRAINING MODEL FOR SEQUENCE-BASED PREDICTION OF PROTEIN-PROTEIN INTERACTION

Yang Xue
Baidu Inc.
Shenzhen, China
xueyang02@baidu.com

Zijing Liu
Baidu Inc.
Shenzhen, China
liuzijing01@baidu.com

Xiaomin Fang
Baidu Inc.
Shenzhen, China
fangxiaomin01@baidu.com

Fan Wang
Baidu Inc.
Shenzhen, China
wangfan04@baidu.com

ABSTRACT

Protein-protein interactions (PPIs) are essentials for many biological processes where two or more proteins physically bind together to achieve their functions. Modeling PPIs is useful for many biomedical applications, such as vaccine design, antibody therapeutics, and peptide drug discovery. Pre-training a protein model to learn effective representation is critical for PPIs. Most pre-training models for PPIs are sequence-based, which naively adopt the language models used in natural language processing to amino acid sequences. More advanced works utilize the structure-aware pre-training technique, taking advantage of the contact maps of known protein structures. However, neither sequences nor contact maps can fully characterize structures and functions of the proteins, which are closely related to the PPI problem. Inspired by this insight, we propose a multimodal protein pre-training model with three modalities: sequence, structure, and function (S2F). Notably, instead of using contact maps to learn the amino acid-level rigid structures, we encode the structure feature with the topology complex of point clouds of heavy atoms. It allows our model to learn structural information about not only the backbones but also the side chains. Moreover, our model incorporates the knowledge from the functional description of proteins extracted from literature or manual annotations. Our experiments show that the S2F learns protein embeddings that achieve good performances on a variety of PPIs tasks, including cross-species PPI, antibody-antigen affinity prediction, antibody neutralization prediction for SARS-CoV-2, and mutation-driven binding affinity change prediction.

1 Introduction

Physical protein-protein interactions (PPIs) refer to the physical contacts between proteins due to interactions such as electrostatic forces, hydrogen bonding, and hydrophobic effects. Many biological processes are driven by PPIs. For example, the severe acute respiratory syndrome coronavirus 2 (SARS-CoV-2) infects human cells starting from the PPI between the viral spike (S) protein and the receptor ACE2 [13]. Applications of PPIs range from vaccine design [30, 14], antibody therapeutics [37, 22], to peptide drug discovery [6, 24]. Furthermore, in many applications, assessing the strength of the *binding affinity* between proteins is important as well because the underlying biological mechanism involves the competitive binding between proteins. Also, in the example of the SARS-Cov-2, the vaccinated person obtained the antibody binding to the N-terminal domain of the viral S protein to impede the original PPIs related to virus entry [30]. However, experimentally determining PPIs can be labor-intensive and time-consuming. Thus robust and accurate computational models are needed as an alternative approach.

Most computational methods to predict PPIs can be briefly divided into two phases: the protein feature encoding and the predictive machine learning models for particular PPI tasks [19]. Appropriate protein encoding is crucial to predict PPIs accurately. The straightforward protein encoding method includes the one-hot embedding, or the k-mer of amino

acid residue [20]. Beyond the straightforward methods, other methods incorporate physicochemical features like the AAindex [17], and the coevolutionary information of proteins in the form of a positional-specific scoring matrix (PSSM) [26, 46]. More recently, using the protein language model to learning a protein representation is shown able to capture rich biological information [34], and has been applied in several PPIs models [49, 12, 40]. Protein language models model the probability of sequences which is found to imply the evolutionary pressure and originally used for homology search in proteomics [1, 2]. Evolutionary information is helpful to predict local structures (secondary structure, contact map) and properties like stability, fluorescence, subcellular localization etc [32, 11]. However, a model trained with only the sequences is not enough to fully characterize the structural information, such as variants of the protein folding patterns as shown in Figure 1 (a).




Figure 1: The missing but important information for PPIs prediction of existing protein encoding models. **(a)** An illustration of three different folding patterns of an intrinsically disordered protein (IDP) (HIV-1 capsid, PDB: 2M8N) which can interact with different proteins. For clearness, we align all the C-terminal domains in the boxes vertically. **(b)** A demonstration of PPI where the interaction involves the side chains. The side chains within the interaction interface are shown with the stick representation in different colors (PDB: 6M0J).

Recently, more attention has been paid on structure-aware protein pre-training models. For example, the protein family is used to incorporate some weak structural information [25]. Contact maps are also used as one pre-training task to learn the structure feature [3, 4]. However, contact maps only consider amino acid-level rigid structure (or C_{α} only), thus don't have explicit information on the side chains and only have limited information about protein-protein interaction. The detailed local structures, the side chains, play a more important role than the backbone structure in the PPI problem. As shown in Figure 1 (b), due to the physical scales, proteins interaction occurs mostly on the side chains, while contact maps cannot explicitly indicate the conformation of side chains. In addition, for the proteins with no fixed 3D structures, like the intrinsically disordered proteins (IDPs) [44], the protein function cannot be characterized only by one fixed 3D structure, and a high-level function description will be needed.

To alleviate the above limitations, we propose a multimodal protein model, which tries to extract a better protein representation by jointly learning the sequence-structure-function (**S2F**) features of proteins, inspired by the recent multimodal pre-training models in natural language processing and computer vision [41]. Besides the protein sequence, we also include structure and function information to learn the protein representation. For the structure encoding, instead of using the contact map, we utilize the point clouds of heavy atoms (C, N, O) that consist of both backbone atoms and side-chain atoms. Furthermore, the topology complex is adopted to model the flexibility of protein structure and extract the structure features, inspired by the works of Guo [45, 43]. Comparing to the contact map, the topology feature can model structures more precisely, especially for the side chains (See the analysis in the Appendix B). For the function modality, we use the textual description from related literature or manual annotations. The multimodal protein data is extracted from the protein domains of the CATH database [38], and the proteins from the UniProtKB/Swiss-Prot database [5]. For downstream PPI tasks, due to the limitation of structure and function data, we can mask the structure and function tokens to get the protein embedding with only protein sequences.

Our experimental results show that comparing to conventional protein language models and contact map-based structure-aware pre-training models, our multimodal S2F representation beats them in various PPI tasks, including cross-species PPI, antibody-antigen affinity prediction, antibody neutralization prediction for SARS-CoV-2, and mutation-driven affinity change prediction.

2 Related Work

Currently, the mainstream of the protein encoding methods for PPIs is to use a pre-training protein model to learn a protein representation. These works include MuPIPR [49], PaccMann-Proteomics [12], D-SCRIPT [40] etc. Specifically,

the embedding model of MuPIPR inspired by ELMo [28] pre-trains a protein language model using bidirectional LSTM on the STRING database [42]. PaccMann-Proteomics extends TAPE [32] to pre-train on pairs of interacting proteins in the STRING database. D-SCRIPT utilizes Bepler & Berger’s model [3] which incorporates structure information using auxiliary tasks including structural similarity prediction and residue-residue contact prediction. Recently, Bepler & Berger further develop their work and learn the structure knowledge and the protein language model in a multitask fashion [4]. For convenience, we use their project repository name ProSE to refer to their pre-training model. With the observations from natural language processing, it is anticipated that a larger pre-training model should also work better. Thus, in this work, we also investigate a super large protein model, ProtTrans [11], which implements Google’s T5 [31] on protein data with 3 billion parameters comparing to TAPE [32] with 34 million parameters.

The above-mentioned methods only need protein sequences to predict the PPIs. There are also methods that directly use protein structure as inputs for the PPI task, such as TopNetTree [43], GeoPPI [21], MutaBind2 [47]. TopNetTree utilizes the topological representation of proteins and auxiliary features extracted from PDB files as the inputs of gradient-boosting trees for the mutation-driven affinity change prediction task [43]. GeoPPI implements the idea of self-supervised pre-training on protein structures via introducing a side-chain perturbation reconstruction task. MutaBind2 shows that sophisticated feature engineering using the information in the PDB file can improve the performance on the PPI task. The major disadvantage of the structure-based PPI methods is that the protein structure data is relatively scarce, and many PPI tasks only have protein sequences data available. By employing multimodal pre-training, our model is able to obtain a representation from the sequence but containing structure and function information.

3 Method

3.1 Sequence-Structure-Function Transformer Model

On the whole, the proposed sequence-structure-function (**S2F**) transformer model is a single-stream multimodal model, in which three types of input tokens are fed into a single transformer and distinguished by input type embeddings.

Sequence. The sequence in the S2F model refers to the amino acid sequence of a protein. The raw amino acid sequence is first tokenized by a protein tokenizer then converted to the amino acids embeddings.

Function. The function modality in this work is the textual function description. Like the sequence modality, the function description is first tokenized by a function tokenizer then converted to the function tokens embeddings. Considering the large vocabulary of protein functions, we use BPE [36] to reduce the embedding matrix size.

Structure. Our S2F model utilizes a topology encoder adopting TopNetTree [43] to get the structure embedding. Specifically, the vectorized Rips complex barcodes and Alpha complex barcodes are feed into a convolutional network and fully connected network respectively, then concatenate them and apply a joint fully connected network to learn the ensemble structure embedding (Figure 2(b)). If we denote the point cloud generated by A-atom and B-atom as $\mathbf{P}_{A,B}$, the Rips barcodes calculation and vectorization by $\text{Rips}(\mathbf{P}_{A,B})$, and the Alpha barcodes calculation and vectorization by $\text{Alpha}(\mathbf{P}_{A,B})$, the progress of topology structure encoding can be represented by the following equations:

$$\begin{aligned} \mathbf{h}_0^{\text{str}} &= \text{Concat}_{A,B \in \{C,N,O\}} \text{Rips}(\mathbf{P}_{A,B}), \\ \mathbf{h}_{12}^{\text{str}} &= \text{Concat}_{A,B \in \{C,N,O\}} \text{Alpha}(\mathbf{P}_{A,B}), \\ \mathbf{e}^{\text{str}} &= \text{Concat}(\text{FC}_2(\text{Conv}(\mathbf{h}_0^{\text{str}}), \text{FC}_1(\mathbf{h}_{12}^{\text{str}}))), \end{aligned} \quad (1)$$

where $\text{FC}(\cdot)$ is a fully connected network, $\text{Conv}(\cdot)$ is a convolutional network, and \mathbf{e}^{str} is the topological structure token.

Like BERT [18], for each input token, the positional embedding is added such that the transformer model can have the sequential information. In our model, we apply two different positional embeddings for the sequence and function tokens and attach a segment embedding to indicate different types of tokens. This final inputs for the transformer can be obtained from Equation (2) to (4),

$$\begin{aligned} \mathbf{x}_{\text{CLS}} &= \mathbf{e}_{\text{CLS}} + \mathbf{p}_0^{\text{seq}} + \mathbf{s}_* \\ \mathbf{x}_i^{\text{seq}} &= \mathbf{e}_i^{\text{seq}} + \mathbf{p}_i^{\text{seq}} + \mathbf{s}_{\text{seq}}, \quad i = 1, \dots, N \end{aligned} \quad (2)$$

$$\begin{aligned} \mathbf{x}_{\text{SEP},1} &= \mathbf{e}_{\text{SEP}} + \mathbf{p}_{N+1}^{\text{seq}} + \mathbf{s}_* \\ \mathbf{x}^{\text{str}} &= \mathbf{e}^{\text{str}} + \mathbf{s}_{\text{str}} \end{aligned} \quad (3)$$

$$\begin{aligned} \mathbf{x}_{\text{SEP},2} &= \mathbf{e}_{\text{SEP}} + \mathbf{p}_0^{\text{fun}} + \mathbf{s}_* \\ \mathbf{x}_j^{\text{fun}} &= \mathbf{e}_j^{\text{fun}} + \mathbf{p}_j^{\text{fun}} + \mathbf{s}_{\text{fun}}, \quad j = 1, \dots, M \\ \mathbf{x}_{\text{SEP},3} &= \mathbf{e}_{\text{SEP}} + \mathbf{p}_{M+1}^{\text{fun}} + \mathbf{s}_*, \end{aligned} \quad (4)$$

where $\mathbf{e}_i^{\text{seq}}$ is the i -th amino acid token, $\mathbf{e}_j^{\text{fun}}$ is the j -th subword token of function, \mathbf{e}_{CLS} is the classifier token, \mathbf{e}_{SEP} is the separator token, $\mathbf{p}_i^{\text{seq}}$ and $\mathbf{p}_j^{\text{fun}}$ are positional embeddings of sequence and function tokens respectively, N is the length of sequence, M is the length of the function description, \mathbf{s}_* , \mathbf{s}_{seq} , \mathbf{s}_{str} and \mathbf{s}_{fun} are the segment embeddings to indicate four types of tokens. These input tokens are converted to feature vectors via a transformer network. In the pre-training phase, the outputs of the transformer are the feature vectors derived from the classifier token, the sequence tokens, and the function tokens, which can be represented by

$$\begin{aligned} &\mathbf{o}_{\text{CLS}}, \mathbf{o}_1^{\text{seq}}, \dots, \mathbf{o}_N^{\text{seq}}, \mathbf{o}_1^{\text{fun}}, \dots, \mathbf{o}_M^{\text{fun}} \\ &= \text{Transformer}(\mathbf{x}_{\text{CLS}}, \mathbf{x}_1^{\text{seq}}, \dots, \mathbf{x}_N^{\text{seq}}, \mathbf{x}_{\text{SEP},1}, \mathbf{x}^{\text{str}}, \mathbf{x}_{\text{SEP},2}, \mathbf{x}_1^{\text{fun}}, \dots, \mathbf{x}_M^{\text{fun}}, \mathbf{x}_{\text{SEP},3}). \end{aligned} \quad (5)$$

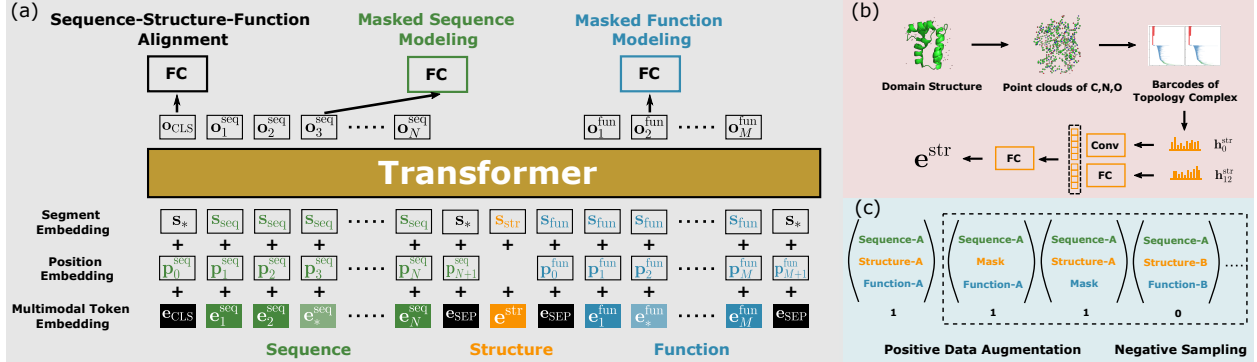


Figure 2: Overview of the Sequence-Structure-Function (S2F) protein pre-training model. Green, orange, and blue elements represent the sequence, structure, and function modality respectively. (a) The architecture of the single-stream multimodal transformer model and three pre-training tasks: sequence-structure-function alignment, masked sequence modeling and masked function modeling. Note that $\mathbf{e}_*^{\text{seq}}$ and $\mathbf{e}_*^{\text{fun}}$ mean the masked tokens. (b) The pipeline of the topology encoder that extracts the topological structure token \mathbf{e}^{str} from the domain structure. (c) An illustration of the positive data augmentation and negative sampling processes for pre-training tasks.

3.2 Pre-training tasks

We use three tasks when pre-training the S2F model (Figure 2). The first is the sequence-structure-function alignment, which is inspired by the previous work on the visual-text pre-training model [41]. In this task, the model learns to figure out whether the input tokens of different modalities are from the same example (Figure 2(c)). In the context of our S2F model, the model makes decisions using its **CLS** header to indicate whether the topological structure encoding corresponds to a specific amino acid sequence or whether the protein function matches the amino acid sequence. Aligning a sequence with its structure topology helps the model to encode meaningful structural information as there are already some works on deciphering protein sequences to low-level structures [32, 11]. Although the gap between a protein sequence and its function might be huge at first glance, we argue that even simply learning the co-occurrence probability of the protein sequence and its scientific function tokens is meaningful. The function descriptions contain important information about PPIs, such as where the protein acts and what other protein it cooperates with. Moreover, thanks to the growing biological research, we can collect more such function data and learn the newest biological knowledge reforming the current model radically.

The other two tasks are variants of the conventional masked language modeling [18], a self-supervised method on the sequential data. We randomly mask out individual residue tokens within the protein sequences or word tokens within the function texts and use a BERT-style objective to predict the masked tokens. The BERT objective is trained jointly with the above alignment task. We apply fully connected neural networks on the outputs of the transformer network to get the predictions of the pre-training tasks. The total training loss of the above three tasks can be written as Equation (6)

$$\mathcal{L} = \mathcal{L}_{ce}(y_{cls}, FC_3(\mathbf{o}_{CLS})) + \sum_{i \in \text{Masked}} \mathcal{L}_{ce}(y_i^{\text{seq}}, FC_4(\mathbf{o}_i^{\text{seq}})) + \sum_{j \in \text{Masked}} \mathcal{L}_{ce}(y_j^{\text{fun}}, FC_5(\mathbf{o}_j^{\text{fun}})), \quad (6)$$

where $\mathcal{L}_{ce}(\cdot)$ is the cross-entropy loss, y_{cls} is the 0-1 label indicating whether cross modalities are matched, y_i^{seq} is the class id (i.e. types of amino acids) of the masked i -th sequence token, and y_j^{fun} is the class id for the masked j -th function token.

3.3 Downstream PPI Tasks

Since we aim at developing a sequence-based PPI model and there are no structure or function data in many PPI problems, we use the **MASK** tokens as the structure and the function tokens when applying the pre-training model, which means we replace e^{str} and e^{fun} with the **MASK** tokens. After pre-training the S2F model, we take the feature vectors from the sequence $(\mathbf{o}_1^{\text{seq}}, \dots, \mathbf{o}_N^{\text{seq}}$ in Equation (5)) as the embedding of the protein, which is the **pre-trained embeddings** for the downstream PPI prediction models as shown in Figure 3.

On top of the protein representation learned by S2F, we utilize different neural networks based on the Residual RCNN from PIPR [8] to make the PPI prediction, according to different kinds of PPI tasks: cross-species PPI, antibody-antigen interaction, and mutation-driven affinity change prediction (Figure 3).

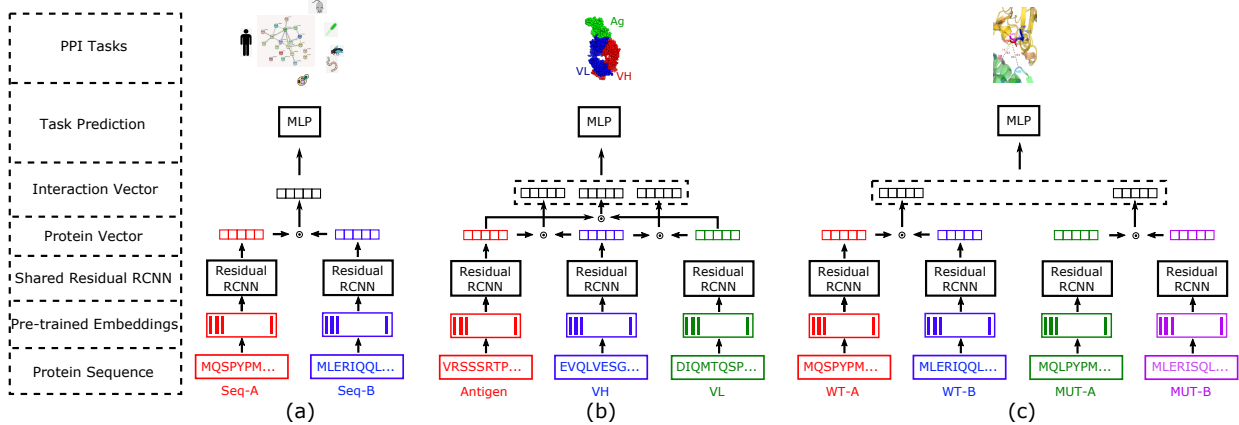


Figure 3: Variants of PIPR-like networks [8] for downstream PPI tasks. **(a)** The network for the cross-species PPI task. **(b)** The network for antibody-antigen interaction tasks. **(c)** The network for mutation-driven affinity change prediction task. Note that the weights of the Residual RCNN module are shared in each specific task. \odot means element-wise multiplication.

Cross-species PPI. In this task, we train a binary classifier on the human PPI data, then directly evaluate its generalization capability on new species, including worm, fly, yeast, E.coli, and mouse. The network architecture for this task is shown as Figure 3(a). As a canonical PPI task, it accepts two protein sequences as input, uses element-wise multiplication to mimic the interaction process, and applies the binary cross-entropy loss to train the model.

Antibody-Antigen interaction. In this work, we only study the cases that the antigen is a protein or a peptide. In such cases, the protein interactions are usually associated with three or more protein sequences. We consider two similar tasks here: the antibody-antigen affinity prediction and the antibody neutralization prediction. For the antibody-antigen affinity prediction task, we need to take the Fab domain (VH and VL chains) and the antigen into considerations. While in the antibody neutralization prediction task, it has an additional receptor protein that is the target of the antigen. In this work, we study the SARS-CoV-2 antibody neutralization task, which has a common receptor protein, such that it can share the same network architecture as the antibody-antigen affinity prediction task. The overall network architecture is illustrated in Figure 3(b). We use the mean square error loss for the affinity prediction task and the binary cross-entropy loss for the antibody neutralization prediction task.

Mutation-driven affinity change prediction. Mutations can dramatically change the properties of a protein by changing the binding affinity of PPIs. In this task, the network accepts the sequence pairs before and after the mutations and tries to predict the affinity change represented by free energy change, $\Delta\Delta G$. We use a network architecture similar to MuPIPR [49], as shown in Figure 3(c). It is a regression task, so the mean square error loss is used to train the network.

4 Results

4.1 Multi-modality Protein Data for Pre-Training

In order to perform the sequence-structure-function pre-training, we compile a multi-modality protein dataset from two databases, the CATH Protein Structure Classification database [38]¹ and the reviewed section of the UniProt Knowledgebase, UniProtKB/Swiss-Prot [5]². Specifically, for the structure data, we use the structures of protein domains because they are relatively stable and independent components comparing to the whole protein. We use PyMOL [35] to purify and extract domain PDB files from RCSB Protein Data Bank [7] based on the CATH annotations. The CATH annotations are also used to train a topology feature classifier, which is used to initialize the weights for the neural networks in Equation (1). For the function data, we collect the function entry from the Swiss-Prot database and the function of the CATH codes of the protein domain. Statistically, we collected 258,928 domains data in sequence-structure-function format, and 564,638 proteins data in sequence-function format. A more detailed description of the data cleaning and processing can be found in Appendix A.

4.2 Cross-species PPI

We first use a canonical PPI task to validate that the S2F learned embeddings are helpful to the PPI prediction. In the cross-species PPI task, we train a PPI classifier using a large PPI dataset on human and evaluate its performance on small PPI datasets from 5 different species: mouse, fly, worm, yeast, and E.coli. Specifically, the human PPI dataset has around 420K protein pairs. While for the five testing species, the E.coli dataset has only 22K pairs, and all other datasets have 55K protein pairs. All of these PPI datasets are obtained from Sledzieski et al.’s recent work and have a positive-to-negative ratio 1:10 [40].

Table 1 compares the performances of PIPR with S2F embeddings and the original PIPR, which uses one-hot encoding and the physicochemical features to embed the proteins [8]. The PIPR network with the S2F multimodal model embeddings significantly outperforms the original PIPR on five testing species. The experimental results of this cross-species PPI task show that the S2F learned embeddings can greatly enhance the representation and have a better generalization ability. Here, it is worth mentioning that for all of our downstream PPI tasks, no structure or function information is needed as inputs because we can use **MASK** tokens as the placeholders (See Section 3.3 for details).

Table 1: Results on Cross-Species PPI

Species	Methods	AUPR	Precision	Recall	AUROC
Mouse	PIPR	0.526	0.734	0.331	0.839
	PIPR+S2F	0.644	0.781	0.447	0.889
Fly	PIPR	0.278	0.521	0.121	0.728
	PIPR+S2F	0.450	0.652	0.286	0.795
Worm	PIPR	0.346	0.673	0.142	0.757
	PIPR+S2F	0.484	0.724	0.270	0.825
Yeast	PIPR	0.230	0.398	0.085	0.718
	PIPR+S2F	0.356	0.604	0.180	0.767
E. coli	PIPR	0.308	0.629	0.131	0.675
	PIPR+S2F	0.371	0.467	0.309	0.748
Human (5-fold CV)	PIPR	0.835	0.838	0.701	0.960
	PIPR+S2F	0.822	0.836	0.672	0.956

¹In this work, we downloaded the RCSB PDB database at 2021.3.15, <http://www.rcsb.org/>

²We collected UniProtKB/Swiss-Prot 2021-03 Release, <https://www.uniprot.org/downloads>

4.3 Antibody-antigen Interaction

We then compare the proposed S2F embeddings with various existing protein pre-training models, and the overall information is listed in Table 2. Regarding the training data size, the smallest one is the ELMo model [49], which only trains on protein sequences of *Homo sapiens*, *Bos taurus*, *Mus musculus*, and *Escherichia coli* in the STRING database [42]; the ProSE model [4] uses the largest training data size. Notably, ProSE incorporates structure information in form of contact maps and the structure classification from the SCOP database [27]. With respect to the network architectures, TAPE [32], ProtTrans [11], and our S2F are based on transformers; ELMo [49] and ProSE [4] are based on bidirectional LSTM.

Table 2: Details of Different Pre-trained Protein Models

Protein Model	Pre-training Dataset	Data Size	Parameters	Embedding Size
ELMo [49]	Subset of STRING	70K	200K	128
TAPE [32]	Pfam	31M	34M	512
ProtTrans [11]	Uniref50	45M	3B	1024
ProSE [4]	SCOP+UniRef90	28K+76M	97M	6165
S2F (Ours)	CATH+UniProtKB/SwissProt	260K+560K	34M	512

In this subsection, we consider the problem of antibody-antigen interaction, which is an important kind of PPIs in life processes. To evaluate our proposed pre-trained embeddings, we use two antibody-antigen tasks:

- **SAbDab**, an antibody-antigen affinity prediction task. The dataset is from the SAbDab database [10] with 494 examples. Each example has an antigen sequence, a variable heavy chain, a variable light chain, and the affinity label.
- **SARS-CoV-2 Antibody Neutralization**, an antibody neutralization prediction task, to predict whether the given antibody can neutralize the effect of S protein and hinder the infection of SARS-CoV-2. The dataset is from the CoV-AbDab database [33] with 747 positive examples and 330 negative examples. Note that the negative examples are selected from the antibodies that can bind to the receptor binding domain (RBD) of the S protein in SARS-CoV-2 but cannot neutralize the virus, which makes the task even harder.

For the antibody-antigen affinity prediction task, SAbDab, we use the root-mean-square error (RMSE) and Pearson’s correlation (Rp) to evaluate the performance. For the classification task, SARS-CoV-2 Antibody Neutralization, we use AUROC, AUPR, F1 score, precision, and recall to evaluate the model. 10-fold cross-validation is performed, and the average across all folds is reported in Table 3. It is easy to see that the proposed S2F model beats the other pre-trained embeddings in both regression and classification tasks.

Table 3: Results on Antibody-Antigen Tasks

Methods	SAbDab		SARS-CoV-2 Antibody Neutralization				
	RMSE	Rp	Precision	Recall	F1	AUPR	AUROC
PIPR+ELMo	0.304	0.317	0.745	0.657	0.698	0.792	0.603
PIPR+TAPE	0.277	0.517	0.735	0.788	0.761	0.801	0.637
PIPR+ProtTrans	0.274	0.521	0.728	0.807	0.766	0.794	0.623
PIPR+ProSE	0.273	0.518	0.721	0.704	0.712	0.794	0.612
PIPR+S2F (Ours)	0.272	0.526	0.742	0.861	0.797	0.836	0.693

4.4 Mutation-driven Affinity Change Prediction

Besides antibody-antigen recognition, other important roles of PPIs include virus entry and tumor cell immune escape. Considering these two functions become unpredictable and easy-perturbed by mutations, in this work, we also test the proposed S2F model on the mutation-driven affinity change prediction task. We perform the experiments on two datasets for this task:

- **S645** [29], an affinity change dataset for single-site mutations. The dataset is compiled from the AB-Bind database [39], with a total of 645 single-point mutations on 29 different Ab-antigen complexes.
- **M1535**, an affinity change dataset for multiple sites mutations. The dataset is filtered from SKEMPI 2.0 database [15]. It is a subset of the M1707 dataset with only two sequences in the PDB files [47].

The samples in the above two datasets are mutations, and the sequences may be very similar if we randomly split them. To minimize the effect of over-fitting and test the ability of generalization, we divide the collected datasets into groups that are different in both structure and homology, with the homology code from the ECOD database [9] and the structure id from the RCSB PDB database [7]. We then merge these initial clusters to organize a 5-fold split and make them as even as possible. We report the performance of different pre-training models with cross-validation on these 5-fold splits. We call such an evaluation method as **Homology&Structure-split Cross Validation (HSCV)**. As shown in Table 4, the HSCV evaluation makes the task more challenging, and the proposed S2F embedding achieves the best performance in terms of both RMSE and Pearson’s correlation.

Table 4: Results on Mutation-driven Affinity Change Prediction Tasks using Homology&Structure-Split Cross Validation

Methods	S645-HSCV		M1535-HSCV	
	RMSE	Rp	RMSE	Rp
PIPR+ELMo	1.913	0.289	3.229	0.220
PIPR+TAPE	1.895	0.273	3.260	0.155
PIPR+ProtTrans	1.888	0.284	3.240	0.196
PIPR+ProSE	1.912	0.257	3.279	0.142
PIPR+S2F (Ours)	1.866	0.318	3.199	0.264

5 Discussion

In this work, we introduce a multimodal protein pre-training model that incorporates sequence, structure, and function data of the protein to jointly train a protein representation for the PPI problem. Our idea is inspired by a more holistic view of proteins that interweaves sequence, structure, and function. By masking the structure and function related input tokens, we can adopt the proposed pre-training model for sequence-based PPI tasks. We assess our pre-trained protein embeddings on various PPI tasks. The first experiment on cross-species PPI shows that our pre-training protein embeddings outperform traditional one-hot and physicochemical feature-based representations. We also compare with several other protein pre-training embeddings. With the same downstream network, the experimental results show that our multimodal pre-training embeddings can boost the performance on different PPI tasks.

From the aspect of applications, we can easily adapt the protein embeddings learned by S2F to other protein problems, such as protein annotation and protein stability prediction. For future works, it is possible to use the high-quality computed protein structures from AlphaFold2 to increase the multimodality data size for pre-training our model [16]. Besides, one may use a pre-trained language model to encode the textual function data better for the function modality. We hope to learn richer knowledge and build better representation for proteins via larger and better datasets, then applying the representation to PPI-related drug discovery.

References

- [1] Stephen F Altschul and Eugene V Koonin. Iterated profile searches with psi-blast—a tool for discovery in protein databases. *Trends in biochemical sciences*, 23(11):444–447, 1998.
- [2] Alex Bateman, Lachlan Coin, Richard Durbin, Robert D Finn, Volker Hollich, Sam Griffiths-Jones, Ajay Khanna, Mhairi Marshall, Simon Moxon, Erik LL Sonnhammer, et al. The pfam protein families database. *Nucleic acids research*, 32(suppl_1):D138–D141, 2004.
- [3] Tristan Bepler and Bonnie Berger. Learning protein sequence embeddings using information from structure. In *International Conference on Learning Representations*, 2019.
- [4] Tristan Bepler and Bonnie Berger. Learning the protein language: Evolution, structure, and function. *Cell Systems*, 12(6):654–669, 2021.
- [5] Emmanuel Boutet, Damien Lieberherr, Michael Tognolli, Michel Schneider, and Amos Bairoch. Uniprotkb/swiss-prot. In *Plant bioinformatics*, pages 89–112. Springer, 2007.
- [6] Heriberto Bruzzoni-Giovanelli, Valerie Alezra, Nicolas Wolff, Chang-Zhi Dong, Pierre Tuffery, and Angelita Rebollo. Interfering peptides targeting protein–protein interactions: the next generation of drugs? *Drug Discovery Today*, 23(2):272–285, 2018.
- [7] Stephen K Burley, Charmi Bhikadiya, Chunxiao Bi, Sebastian Bittrich, Li Chen, Gregg V Crichlow, Cole H Christie, Kenneth Dalenberg, Luigi Di Costanzo, Jose M Duarte, et al. Rcsb protein data bank: powerful new tools for exploring 3d structures of biological macromolecules for basic and applied research and education in

- fundamental biology, biomedicine, biotechnology, bioengineering and energy sciences. *Nucleic acids research*, 49(D1):D437–D451, 2021.
- [8] Muhao Chen, Chelsea J-T Ju, Guangyu Zhou, Xuelu Chen, Tianran Zhang, Kai-Wei Chang, Carlo Zaniolo, and Wei Wang. Multifaceted protein–protein interaction prediction based on siamese residual rcnn. *Bioinformatics*, 35(14):i305–i314, 2019.
- [9] Hua Cheng, R Dustin Schaeffer, Yuxing Liao, Lisa N Kinch, Jimin Pei, Shuoyong Shi, Bong-Hyun Kim, and Nick V Grishin. Ecod: an evolutionary classification of protein domains. *PLoS computational biology*, 10(12):e1003926, 2014.
- [10] James Dunbar, Konrad Krawczyk, Jinwoo Leem, Terry Baker, Angelika Fuchs, Guy Georges, Jiye Shi, and Charlotte M Deane. Sabdab: the structural antibody database. *Nucleic acids research*, 42(D1):D1140–D1146, 2014.
- [11] Ahmed Elnaggar, Michael Heinzinger, Christian Dallago, Ghalia Rehawi, Wang Yu, Llion Jones, Tom Gibbs, Tamas Feher, Christoph Angerer, Martin Steinegger, Debsindhu Bhowmik, and Burkhard Rost. 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*, pages 1–1, 2021.
- [12] Modestas Filipavicius, Matteo Manica, Joris Cadow, and Maria Rodriguez Martinez. Pre-training protein language models with label-agnostic binding pairs enhances performance in downstream tasks. *arXiv preprint arXiv:2012.03084*, 2020.
- [13] Markus Hoffmann, Hannah Kleine-Weber, Simon Schroeder, Nadine Krüger, Tanja Herrler, Sandra Erichsen, Tobias S Schiergens, Georg Herrler, Nai-Huei Wu, Andreas Nitsche, et al. Sars-cov-2 cell entry depends on ace2 and tmprss2 and is blocked by a clinically proven protease inhibitor. *cell*, 181(2):271–280, 2020.
- [14] Lisa A Jackson, Evan J Anderson, Nadine G Rouphael, Paul C Roberts, Mamodikoe Makhene, Rhea N Coler, Michele P McCullough, James D Chappell, Mark R Denison, Laura J Stevens, et al. An mrna vaccine against sars-cov-2—preliminary report. *New England Journal of Medicine*, 2020.
- [15] Justina Jankauskaitė, Brian Jiménez-García, Justas Dapkūnas, Juan Fernández-Recio, and Iain H Moal. Skempi 2.0: an updated benchmark of changes in protein–protein binding energy, kinetics and thermodynamics upon mutation. *Bioinformatics*, 35(3):462–469, 2019.
- [16] John Jumper, Richard Evans, Alexander Pritzel, Tim Green, Michael Figurnov, Olaf Ronneberger, Kathryn Tunyasuvunakool, Russ Bates, Augustin Žídek, Anna Potapenko, et al. Highly accurate protein structure prediction with alphafold. *Nature*, 596(7873):583–589, 2021.
- [17] Shuichi Kawashima and Minoru Kanehisa. Aaindex: amino acid index database. *Nucleic acids research*, 28(1):374–374, 2000.
- [18] Jacob Devlin Ming-Wei Chang Kenton and Lee Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of NAACL-HLT*, pages 4171–4186, 2019.
- [19] Mst Khatun, Watshara Shoombuatong, Md Hasan, Hiroyuki Kurata, et al. Evolution of sequence-based bioinformatics tools for protein-protein interaction prediction. *Current Genomics*, 21(6):454–463, 2020.
- [20] Yiwei Li and Lucian Ilie. Sprint: ultrafast protein-protein interaction prediction of the entire human interactome. *BMC bioinformatics*, 18(1):1–11, 2017.
- [21] Xianggen Liu, Yunan Luo, Pengyong Li, Sen Song, and Jian Peng. Deep geometric representations for modeling effects of mutations on protein-protein binding affinity. *PLoS computational biology*, 17(8):e1009284, 2021.
- [22] Zhe Lv, Yong-Qiang Deng, Qing Ye, Lei Cao, Chun-Yun Sun, Changfa Fan, Weijin Huang, Shihui Sun, Yao Sun, Ling Zhu, et al. Structural basis for neutralization of sars-cov-2 and sars-cov by a potent therapeutic antibody. *Science*, 369(6510):1505–1509, 2020.
- [23] Valerio Mariani, Marco Biasini, Alessandro Barbato, and Torsten Schwede. Iddt: a local superposition-free score for comparing protein structures and models using distance difference tests. *Bioinformatics*, 29(21):2722–2728, 2013.
- [24] Susan Marqus, Elena Pirogova, and Terrence J Piva. Evaluation of the use of therapeutic peptides for cancer treatment. *Journal of biomedical science*, 24(1):1–15, 2017.
- [25] Seonwoo Min, Seunghyun Park, Siwon Kim, Hyun-Soo Choi, Byunghan Lee, and Sungroh Yoon. Pre-training of deep bidirectional protein sequence representations with structural information. *IEEE Access*, 9:123912–123926, 2021.
- [26] Yoichi Murakami and Kenji Mizuguchi. Homology-based prediction of interactions between proteins using averaged one-dependence estimators. *BMC bioinformatics*, 15(1):1–11, 2014.

- [27] Alexey G Murzin, Steven E Brenner, Tim Hubbard, and Cyrus Chothia. Scop: a structural classification of proteins database for the investigation of sequences and structures. *Journal of molecular biology*, 247(4):536–540, 1995.
- [28] Matthew E Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. Deep contextualized word representations. In *Proceedings of NAACL-HLT*, pages 2227–2237, 2018.
- [29] Douglas EV Pires and David B Ascher. mCSM-AB: a web server for predicting antibody–antigen affinity changes upon mutation with graph-based signatures. *Nucleic acids research*, 44(W1):W469–W473, 2016.
- [30] Fernando P Polack, Stephen J Thomas, Nicholas Kitchin, Judith Absalon, Alejandra Gurtman, Stephen Lockhart, John L Perez, Gonzalo Pérez Marc, Edson D Moreira, Cristiano Zerbini, et al. Safety and efficacy of the bnt162b2 mrna covid-19 vaccine. *New England Journal of Medicine*, 2020.
- [31] Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. Exploring the limits of transfer learning with a unified text-to-text transformer. *Journal of Machine Learning Research*, 21:1–67, 2020.
- [32] Roshan Rao, Nicholas Bhattacharya, Neil Thomas, Yan Duan, Xi Chen, John Canny, Pieter Abbeel, and Yun S Song. Evaluating protein transfer learning with tape. *Advances in neural information processing systems*, 32:9689, 2019.
- [33] Matthew IJ Raybould, Aleksandr Kovaltsuk, Claire Marks, and Charlotte M Deane. Cov-abdab: the coronavirus antibody database. *Bioinformatics*, 37(5):734–735, 2021.
- [34] Alexander Rives, Joshua Meier, Tom Sercu, Siddharth Goyal, Zeming Lin, Jason Liu, Demi Guo, Myle Ott, C Lawrence Zitnick, Jerry Ma, et al. Biological structure and function emerge from scaling unsupervised learning to 250 million protein sequences. *Proceedings of the National Academy of Sciences*, 118(15), 2021.
- [35] Schrödinger, LLC. The PyMOL molecular graphics system, version 1.8. November 2015.
- [36] Rico Sennrich, Barry Haddow, and Alexandra Birch. Neural machine translation of rare words with subword units. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1715–1725, 2016.
- [37] Balamurugan Shanmugaraj, Konlavat Siriwananon, Kittikhun Wangkanont, and Waranyoo Phoolcharoen. Perspectives on monoclonal antibody therapy as potential therapeutic intervention for coronavirus disease-19 (covid-19). *Asian Pacific journal of allergy and immunology*, 38(1):10–18, 2020.
- [38] Ian Sillitoe, Nicola Bordin, Natalie Dawson, Vaishali P Waman, Paul Ashford, Harry M Scholes, Camilla SM Pang, Laurel Woodridge, Clemens Rauer, Neeladri Sen, et al. Cath: increased structural coverage of functional space. *Nucleic acids research*, 49(D1):D266–D273, 2021.
- [39] Sarah Sirin, James R Apgar, Eric M Bennett, and Amy E Keating. Ab-bind: antibody binding mutational database for computational affinity predictions. *Protein Science*, 25(2):393–409, 2016.
- [40] Samuel Sledzieski, Rohit Singh, Lenore Cowen, and Bonnie Berger. Sequence-based prediction of protein-protein interactions: a structure-aware interpretable deep learning model. *bioRxiv*, 2021.
- [41] Weijie Su, Xizhou Zhu, Yue Cao, Bin Li, Lewei Lu, Furu Wei, and Jifeng Dai. Vi-bert: Pre-training of generic visual-linguistic representations. In *International Conference on Learning Representations*, 2019.
- [42] Damian Szklarczyk, John H Morris, Helen Cook, Michael Kuhn, Stefan Wyder, Milan Simonovic, Alberto Santos, Nadezhda T Doncheva, Alexander Roth, Peer Bork, Lars J. Jensen, and Christian von Mering. The STRING database in 2017: quality-controlled protein–protein association networks, made broadly accessible. *Nucleic Acids Research*, 45(D1):D362–D368, 10 2016.
- [43] Menglun Wang, Zixuan Cang, and Guo-Wei Wei. A topology-based network tree for the prediction of protein–protein binding affinity changes following mutation. *Nature Machine Intelligence*, 2(2):116–123, 2020.
- [44] Peter E Wright and H Jane Dyson. Intrinsically disordered proteins in cellular signalling and regulation. *Nature reviews Molecular cell biology*, 16(1):18–29, 2015.
- [45] Kelin Xia and Guo-Wei Wei. Persistent homology analysis of protein structure, flexibility, and folding. *International journal for numerical methods in biomedical engineering*, 30(8):814–844, 2014.
- [46] Jing-Xuan Zhai, Tian-Jie Cao, Ji-Yong An, and Yong-Tao Bian. Highly accurate prediction of protein self-interactions by incorporating the average block and pssm information into the general pseaac. *Journal of theoretical biology*, 432:80–86, 2017.
- [47] Ning Zhang, Yuting Chen, Haoyu Lu, Feiyang Zhao, Roberto Vera Alvarez, Alexander Goncareenco, Anna R Panchenko, and Minghui Li. Mutabind2: predicting the impacts of single and multiple mutations on protein-protein interactions. *Iscience*, 23(3):100939, 2020.

- [48] Yang Zhang and Jeffrey Skolnick. Scoring function for automated assessment of protein structure template quality. *Proteins: Structure, Function, and Bioinformatics*, 57(4):702–710, 2004.
- [49] Guangyu Zhou, Muhao Chen, Chelsea JT Ju, Zheng Wang, Jyun-Yu Jiang, and Wei Wang. Mutation effect estimation on protein–protein interactions using deep contextualized representation learning. *NAR genomics and bioinformatics*, 2(2):lqaa015, 2020.

Appendices

A Pre-training Data Preparation

Domain data processing. Extracting structures and functions of domains in the PDB file involves several steps. First, we remove solvent molecules and ligands in the PDB file. Because these molecules also have heavy atoms C, N, O, and unwanted heavy atoms will be included in the point clouds when we compute the topology complex features. We then split protein chains and remove the repeated structures. Finally, we apply sequence alignment over each protein sequence and domains sequences in the CATH database [38] to get the corresponding CATH code and function description.

Topology barcodes extraction and vectorization. Once the domains are extracted, we compute the topology complex features from the point clouds of the heavy atoms (C, N, O). Inspired by TopNetTree [43], we use both Alpha complex and Rips complex. Specifically, for each heavy atom pair A-B (one of C-C, N-N, O-O, C-N, C-O, N-O), we calculate the barcodes of the Rips and Alpha complex using the distance matrix of the point cloud represented by all A-atoms and B-atoms. Considering the physical size of protein domains, we select the cutoff 0.01 angstrom and 20 angstroms, meaning that the barcodes with length (death minus birth) smaller than 0.01 angstrom or with death value larger than 20 angstroms are ignored. Then we make bins to count the number of barcodes with interval 0.25 angstrom. These bins normalized by the mean and the standard deviation over 258,928 examples create a vector as a low-dimensional structure topology representation.

Topology feature classifier. Since the CATH codes describe a protein in four levels (class, architecture, topology, homologous superfamily), we treat them as classification labels of proteins. Using these labels and the topology encoder network (Figure 2 (b)), we can train a classifier that can act as good initialization weights for the proposed pre-training model. Statistically, there are 3,763 distinct CATH codes within our domains data.

B Contact Map and Topology Feature for IDPs

Intrinsically disordered proteins (IDPs) allow the same protein to undertake different interactions under different folding patterns. Therefore, IDPs act as important components of cellular signaling machinery and usually have a large degree in the PPI networks [44]. To better predict the PPIs related to IDPs, there is an unmet need to model their variants of structures precisely. Most existing works model the protein structure at the level of contact maps. However, for IDPs, different folding patterns cannot be precisely presented by contact map because of its characteristic of locality. For example, the contact maps in Figure 4 are corresponding to the three different folding patterns of HIV-1 capsid protein in Figure 1 (a). It is hard to characterize the significantly different folding patterns from the nuances between the contact maps.




Figure 4: Contact maps of HIV-1 capsid protein with different folding patterns. Note that the threshold of the contact map is 8 angstroms.

Furthermore, to assess the different capabilities of contact map and topology feature at modeling the side chains, we examine the correlations between the structural differences and the distances given by different structure encoding vectors. Specifically, we collect 100 different structures of the HIV-1 capsid protein from the PDB file 2m8n. Then we calculate the pairwise structure differences using TM-Score [48] and LDDT [23]. For both TM-Score and LDDT, the higher score (more closer to 1) means higher structure similarity. Since TM-Score only considers the C_α and LDDT scores are computed on all atoms, we can use $1 - \text{TMScore}$ as the structure difference without the side chains, and use $1 - \text{LDDT}$ as the structure difference with the side chains. Next, contact maps are flattened as encoding vectors and we calculate the pairwise Euclidean distances as the **contact map distance**. Similarly, we can get the **topology feature distance** from the normalized topology feature vectors. Finally, we compute two pairs of correlations shown in

Figure 5. As we can see, without the side chains information, both the contact map distance and the topology feature distance have the same Pearson correlation with the structural difference. While when considering the side chains, the topology feature distance has a higher correlation with the structure differences than the contact map distance. This result indicates that the topology feature can model the variants of structures with side chains more precisely.




Figure 5: Correlations of $1 - \text{TMScore}$ and $1 - \text{LDDT}$ with the contact map distance and the topology feature distance.