

Protein 3D Graph Structure Learning for Robust Structure-Based Protein Property Prediction

Yufei Huang^{1,2}, Siyuan Li^{1,2}, Lirong Wu^{1,2}, Jin Su^{1,2}, Haitao Lin^{1,2},
Odin Zhang¹, Zihan Liu^{1,2}, Zhangyang Gao^{1,2}, Jiangbin Zheng^{1,2}, Stan Z. Li^{2*}

¹ Zhejiang University, Hangzhou

² AI Lab, Research Center for Industries of the Future, Westlake University

huangyufei, lisiyuan, sujian, wulirong, linhaitao, zhengjiangbin, Stan.ZQ.Li@westlake.edu.cn, haotianzhang@zju.edu.cn

Abstract

Protein structure-based property prediction has emerged as a promising approach for various biological tasks, such as protein function prediction and sub-cellular location estimation. The existing methods highly rely on experimental protein structure data and fail in scenarios where these data are unavailable. Predicted protein structures from AI tools (e.g., AlphaFold2) were utilized as alternatives. However, we observed that current practices, which simply employ accurately predicted structures during inference, suffer from notable degradation in prediction accuracy. While similar phenomena have been extensively studied in general fields (e.g., Computer Vision) as model robustness, their impact on protein property prediction remains unexplored. In this paper, we first investigate the reason behind the performance decrease when utilizing predicted structures, attributing it to the structure embedding bias from the perspective of structure representation learning. To study this problem, we identify a Protein 3D Graph Structure Learning Problem for Robust Protein Property Prediction (PGSL-RP3), collect benchmark datasets, and present a protein Structure embedding Alignment Optimization framework (**SAO**) to mitigate the problem of structure embedding bias between the predicted and experimental protein structures. Extensive experiments have shown that our framework is model-agnostic and effective in improving the property prediction of both predicted structures and experimental structures.

Introduction

Proteins are workhorses of the cell, involved in various biological processes, such as immune response and DNA replication. Understanding the properties of proteins is important for deciphering the mystery of life (Degn et al. 2023; Hu et al. 2022) and treating various diseases (Rossi Sebastiano et al. 2022; Zheng et al. 2023). As most protein properties are governed by their folded structures, protein structure-based property prediction has emerged as a promising approach for various biological tasks, such as protein function prediction (Huang et al. 2023; Liu et al. 2023), sub-cellular location estimation (Zhang et al. 2022), and structure-based drug estimation (Lin et al. 2022a, 2023a).

*Corresponding Author

Copyright © 2024, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

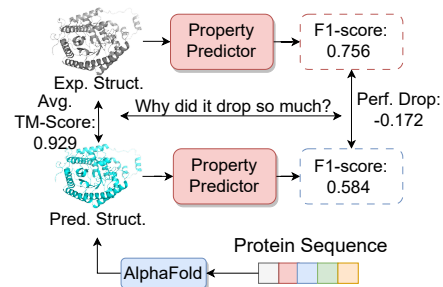


Figure 1: Performance drop on EC task. TM-score is a widely used metric for structure prediction accuracy.

The existing methods' reliance on experimental protein structures poses a challenge when such structures are unavailable. Predicted protein structures from tools like AlphaFold2 have been utilized as alternatives. However, for various structure-based protein property prediction methods, even if these predicted structures are **accurate**, using them during inference also results in a notable decrease in property prediction accuracy. (as illustrated in Fig.1). This suggests that there is a deeper factor behind the accuracy of predicted structures that misleads structure-based predictors. Similar phenomena also occur in other fields where new samples with small and hidden differences from the original ones can easily fool networks in making predictions for downstream tasks. In Computer Vision, when we apply small and invisible perturbations to a panda picture, the neural network misclassifies it to gibbon (Goodfellow, Shlens, and Szegedy 2015). Graph Neural Network is also known to be vulnerable to small perturbations like adding or deleting a few edges (Jin et al. 2020; Wu et al. 2021, 2022c). While developing robust algorithms (e.g., Graph Structure Learning, GSL (Liu et al. 2022; Jin et al. 2020)) to resist permutations has been well studied in general domains, model robustness in structure-based property predictions remains unexplored.

Therefore, in this paper, we investigate the decrease in prediction accuracy when using accurately predicted protein structures. It is attributed to the **structure embedding bias** of predictors from the perspective of structure representation learning, i.e., a distribution gap between embedding of accurately predicted structure and that of experimental struc-




Figure 2: The illustration of vanilla and our learned protein structure embedding of a subset of GOCC dataset by t-SNE. The Red is experimental structure embedding, and the blue is predicted structure embedding. There is a clear bias between the predicted and experimental structures in embeddings from vanilla encoder but embeddings from SAO-pretrained encoder tend to be smoother and bias-free.

ture as shown in Fig.2. We further formulate this problem as Protein 3D Graph Structure Learning for Robust Protein Property Prediction (*PGSL-RP3*). Improved predicted structures don’t necessarily lead to better prediction results due to the structure embedding bias; the strategy of structure refinement, i.e., further improving the structure prediction accuracy, can’t completely alleviate the performance decrease.

To address these issues, we propose to align the representation of the predicted structure to that of the experimental structure rather than directly improving the similarity (e.g., structure prediction accuracy) between predicted and experimental structures in data space. To achieve this, we present the protein Structure embedding Alignment Optimization framework (*SAO*). In *SAO*, we create pairs of predicted and experimental structures and train the encoder to align the predicted structure representation with the corresponding experimental structure representation using a bootstrap and denoising approach. One advantage of our alignment-based framework is its ability to leverage paired information from both predicted and experimental structures, resulting in improved representation learning performance compared to using a single data source alone. Additionally, our framework can utilize low-precision predicted structures that are often overlooked, further enhancing its effectiveness.

Related Work

Protein Structure Representation Model. Graph neural networks (Wu et al. 2022b, 2023a) were applied to protein structure representation (Baldassarre et al. 2021; Gligorić et al. 2021), with promising results in protein structure design (Ingraham et al. 2019; Dauparas et al. 2022). With the rise of geometric deep learning (Lin et al. 2023b; Wu et al. 2022a, 2023b), Equivariant Neural Networks(ENN) began to be applied to protein representation, allowing us to directly learn protein structures end-to-end and achieve more powerful results (Hsu et al. 2022; Jing et al. 2020; Wu and Cheng 2022). Inspired by AlphaFold (Jumper et al. 2021), many works have tried to combine sequence and structure representation, expecting that co-modeling can combine the

advantages of both modalities (Lin et al. 2022b; Mansoor et al. 2021; You and Shen 2022; Wu et al. 2023a).

Protein Data Mining With Predicted Structure. With the introduction of AlphaFold2, many unknown structures have been solved predictively, which has fuelled the enthusiasm to transform large amounts of predicted structure into valuable knowledge (Shi et al. 2019; Degn et al. 2023). Some progress has been made in certain areas (Al-Masri et al. 2022; Rossi Sebastiano et al. 2022; Degn et al. 2023). In particular, areas with less mainstream attention, such as rare disease research, are expected to benefit greatly from AlphaFold2’s predicted structures (Rossi Sebastiano et al. 2022). At the same time, there is also a growing interest in the quality of knowledge generated by predictive structures, and there are now a number of relevant evaluation studies (Shi et al. 2019; Degn et al. 2023; Pan et al. 2022). However, current studies still focus on the accuracy of the prediction and ignore the structure embedding bias that the predicted structure itself may carry. EquiPPIS (Roche et al. 2022) has considered the problem of generalization to predicted structures and designed an equivariant neural network-based model for a specific task to alleviate this issue. We go further to reveal the nature of the problem and formulate it as the PGSL problem, thus proposing a more efficient and universal framework for its solution. More related work can be referred to the Appendix A.

Preliminaries

Notions Protein data can be modeled at multiple levels: sequence, amino acid level, full atom level, etc. Here we model proteins uniformly as an Attributed Relational Graph: $\mathcal{G} = (\mathcal{V}, \mathcal{E}, \mathcal{N}, \mathcal{R})$, where \mathcal{V} represents the ordered set of graph nodes (can be amino acids or atoms) and $\mathcal{E} \in \mathcal{V} \times \mathcal{V}$ represents the corresponding set of edges connecting the nodes (some relationship between nodes, e.g., distance less than 4 Å). Every vertex $v \in \mathcal{V}$ in \mathcal{G} can have both scalar and vector attributes $\mathbf{n}_v = (S_v, V_v) \in \mathcal{N}$, where $S_v \in \mathbb{R}^S$ and $V_v \in \mathbb{R}^{3 \times V}$. Similarly, each edge $e \in \mathcal{E}$ have attributes $\mathbf{r}_v = (S_e, V_e) \in \mathcal{R}$, where $S_e \in \mathbb{R}^N$ and $V_e \in \mathbb{R}^{3 \times T}$. \mathcal{G} can contain empty sets. When the sets \mathcal{E} and \mathcal{R} are empty sets, \mathcal{G} degenerates to a single sequence representation. Furthermore, if \mathcal{N} contains only amino acid composition, \mathcal{G} degenerates to the amino acid sequence.

PGSL for Protein Property Prediction. The aim of structure-based protein property prediction is to predict (classify or regress) the property of a protein given its structure. The Protein 3D Graph Structure Learning problem (PGSL) extends protein structures from experimental to predicted structures. Its goal is to enable protein structure representation models to align predicted structures to experimental structures, thus enabling representation models to better handle large numbers of predicted protein structures and improve humanity’s understanding of unknown proteins. Specifically for Protein Property Prediction, PGSL for Robust Protein Property Prediction (PGSL-RP3) requires that structure-based protein property prediction models can make correct annotations on predicted protein structures and experimental structures.

We propose two different views for understanding PGSL-RP3 in Fig.3, which also correspond to two types of solution ideas. The second one is adopted in our SAO framework. Further preliminaries can be found in Appendix B.

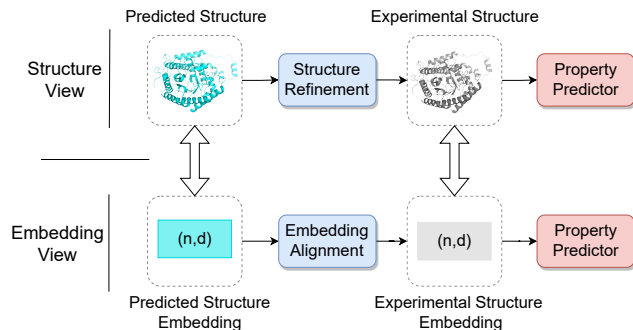


Figure 3: An illustration of PGSL-RP3 and two different views for understanding the problem. The alignment (or protein 3D graph structure learning) from the predicted structure to the experimental structure can occur in either the structure space or the embedding space.

Protein Structure Embedding Alignment Optimization Framework

Problem Analysis and Motivation

The comparisons between PGSL and GSL motivate us to analyze PGSL from the perspective of embedding. PGSL is similar to GSL in that they both expect a cleaner graph structure to improve performance on downstream tasks. GSL usually focuses on node classification tasks on 2D graphs, where the modifications of the graph structure are generally discrete operations such as the addition and deletion of edges. Different from GSL, PGSL focuses on graph classification tasks on 3D graphs, similar to 3D point clouds, where the modifications of the graph structure are continuous operations, such as the modification of node coordinates. As a result, it is difficult to migrate methods directly from the GSL domain to PSGSL. However, the idea of node embedding alignment in unsupervised graph structure learning inspires us to explore PGSL from the embedding view.

As a result, we propose the SAO framework for alignment in Embedding Space as illustrated in Fig.4(a). Instead of optimizing the structure directly, our framework used the idea of representation learning to move the embedding of predicted structures toward that of the experimental structures. In practice, we first pretrain the protein encoder based on the SAO framework to give it the ability to migrate representations of predicted structures towards that of experimental structures. Then, by fine-tuning the encoder on specific tasks, as in Fig.4(b), the model learns to associate representations and labels together. Since the property prediction problem is essentially carried out based on the structure embedding, such an idea can instead improve the model robustness more directly. Our proposed framework will be discussed in the following subsections.

Directional Embedding Alignment

The motivation of the first module is to help the model achieve *Direction Embedding Alignment*. Noting that there is actually only some natural mapping relationship between the experimental protein structure and its properties, we should learn the correlation between protein structure and properties based on the experimental structure. Therefore, we need to align the embedding of the predicted structure with the embedding of the experimental structure.

To achieve this, we propose a one-way embedding alignment method, drawing on bootstrap-based contrastive learning methods (Grill et al. 2020; Chen and He 2020). Our method is based on the student-teacher architecture, where the parameters ϕ of the teacher network (a.k.a. Target network) are the exponential moving average of the parameters θ of the student network (also known as Online network) and only the parameters of the student encoder are trainable. More precisely, given a target decay rate λ , the following updates are performed after each training iteration,

$$\phi \leftarrow \lambda\phi + (1 - \lambda)\theta. \quad (1)$$

The online network is comprised of three stages: an encoder f_θ , a projector g_θ , and a predictor q_θ . The target network and online network have the same architecture, except that there is no predictor. The embedding outputs of both the student encoder and teacher encoder are fed into an embedding projector as shown in Figure 4(a).

Given a set \mathcal{S} of paired protein structures, a protein structure pair sampled uniformly from \mathcal{S} , the teacher network takes the experimental structure as input and outputs a representation z_ϕ of the experimental structure as the prediction target. The student network takes the predicted structure as input. It outputs a representation z_θ of the predicted structure that goes through the predictor q_θ to be aligned with the experimental structure representation. We can ℓ_2 -normalize both $q_\theta(z_\theta)$ and z_ϕ to $\bar{q}_\theta(z_\theta) \triangleq q_\theta(z_\theta) / \|q_\theta(z_\theta)\|_2$ and $\bar{z}_\phi \triangleq (z_\phi) / \|z_\phi\|_2$. The loss function is as follow:

$$\mathcal{L}_{align}^{\theta, \phi} \triangleq \|\bar{q}_\theta(z_\theta) - \bar{z}_\phi\|_2^2 = 2 - 2 \cdot \frac{\langle q_\theta(z_\theta), z_\phi \rangle}{\|q_\theta(z_\theta)\|_2 \cdot \|z_\phi\|_2}, \quad (2)$$

where the first part (i.e., ℓ_2 -norm loss) is equal to cosine similarity loss.

The loss $\mathcal{L}_{align}^{\theta, \phi}$ is asymmetric to ensure directional embedding alignment, and networks always use the representation of the predicted structure to predict that of the experimental structure. It is worth noting that we only optimize the loss $\mathcal{L}_{align}^{\theta, \phi}$ according to θ , for the target network with ϕ we use a gradient-stop technique. After pretraining, we only keep the online encoder shown in Figure 4(b).

Embedding Bootstrap Mechanism To give the encoder persistent guidance and avoid over-fitting in the experimental representation, we use the embedding bootstrap mechanism to provide a self-enhanced experimental structure embedding as the learning target. Considering that the embedding is ultimately reflected in the parameters, our core idea is to update the target encoder with an exponential moving average of the learned online encoder as in equation 1 instead of keeping it unchanged.

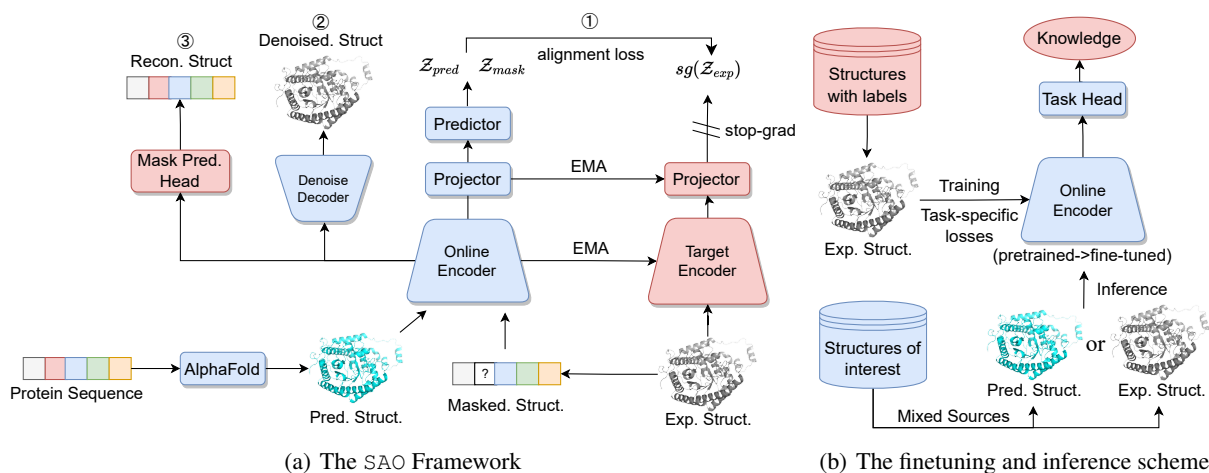


Figure 4: The illustration of SAO Framework. (a) We first pretrain the encoder in SAO framework for the ability of directional embedding alignment (b). Then we finetune the encoder on specific downstream tasks. Finally, we can turn structures of interest (predicted or experimental) into new knowledge(predict its various properties) to guide scientific discovery.

Mask View as Embedding Augmentation

The feasible solutions in the above framework include collapsed representation (e.g., constant representation across sources is always fully predictive of itself). However, $\mathcal{L}^{\theta, \phi}_{align}$ is not a loss such that SAO’s dynamics is a gradient descent on \mathcal{L} jointly over θ, ϕ . There is, therefore, no *a priori* reason why SAO’s parameters would converge to an undesirable minimum of $\mathcal{L}^{\theta, \phi}_{align}$ (Grill et al. 2020).

Furthermore, to avoid the undesirable equilibria in SAO’s dynamics, we introduce the Mask view as Embedding Augmentation (as shown in Fig.4(a)). We mask the protein sequence of the experimental structure and keep the structure visible to create a mask view. Then we align the embedding of the masked structure toward that of the experimental structure following the same procedure we operate on the predicted structure. It has the following two advantages.

Increase Variability for Avoiding Collapse The technical consideration behind this is to propagate new sources of variability captured by the online projection into the target projection. As shown in (Grill et al. 2020), increasing variability across sources could make these collapsed constant equilibria unstable. The variability between the predicted and experimental structure is limited because, as we mentioned above, the difference between the two is mainly in **embedding bias** when the prediction accuracy is high.

Embedding Augmentation for Better Alignment Another motivation behind introducing the mask view is to increase the ability of the framework to align embedding from other sources to the experimental structural embedding. Due to the small variability between the predicted and experimental structures, the model is more prone to overfitting and loss of generalizability. The predicted structure mainly carries global variability, so we consider proposing augmentation with **contextual variability** to enhance the ability of

directional embedding alignment of our framework.

Structure Denoising Guidance

To ensure meaningful representations that can be decoded into cleaner structures and guide the directional embedding alignment, we incorporate a structure denoising task.

Instead of reducing the amino acid molecule to a single C_{α} atom, we pursued a finer-grained modeling, so we considered all the amino acid backbone atoms (C, N, O, C_{α}). Considering physical plausibility (the bond lengths and bond angles between backbone atoms are relatively fixed, i.e., the relative positions between backbone atoms can be considered fixed), we modeled the backbone atoms as a frame, such that an amino acid backbone structure is determined by two vector properties: translation \mathbf{t} (coordinates of the C_{α} atom) and orientation \mathcal{O} (which determines the final coordinates):

$$\mathcal{P} = \{(\mathbf{t}_i, \mathcal{O}_i)\}_{i=1}^N \text{ and } x_i^a = \mathcal{O}_i x_{stand}^a + \mathbf{t}_i. \quad (3)$$

where \mathcal{P} is the protein backbone structure, N is the length of the protein sequence, x_i^a is the a -type backbone atom of residue i ($a \in C, N, O, C_{\alpha}$). x_{stand}^a is the standard coordinate of the a -type amino acid backbone atom when the C_{α} atom is at the origin and under the unit orthogonal group (Jumper et al. 2021).

Given a predicted protein structure \mathcal{P} , the online encoder maps it into embedding. A denoise decoder head predicts the direction of conformational change for further optimization. Formally, we can formulate this process as follows:

$$\begin{aligned} \{(\Delta \mathbf{t}_i, \Delta \mathcal{O}_i)\}_{i=1}^N &= \mathcal{E}(\mathcal{P}), \\ \mathbf{t}_i &= \mathbf{t}_i + \Delta \mathbf{t}_i \text{ and } \mathcal{O}_i = \mathcal{O}_i \circ \Delta \mathcal{O}_i. \end{aligned} \quad (4)$$

where \mathcal{E} is the equivariant neural network and \circ corresponds to the composition of elements in SO(3) group. we use the MSE loss under the local frame (Jumper et al. 2021) as the training objective:

$$\mathcal{L}_{mse} = \text{Mean}_{i,j} \sqrt{\|T_i^{-1} \circ x_j - T_i^{true-1} \circ x_j^{true}\|^2}. \quad (5)$$

Models	Inference Setting	EC	GO-MF	GO-CC	GO-BP
IPA-Encoder	Predicted (TM \geq 0.5)	0.568 (+0.156)	0.492 (+0.055)	0.383 (+0.061)	0.311 (+0.039)
	Predicted (pLDDT \geq 70)	0.579 (+0.158)	0.494 (+0.056)	0.362 (+0.084)	0.312 (+0.040)
	Experimental Structure	0.711 (+0.054)	0.522 (+0.025)	0.433 (+0.024)	0.338 (+0.023)
	The Performance Gap	-0.132 (+0.104)	-0.028 (+0.031)	-0.071 (+0.060)	-0.026 (+0.017)
Uni-Mol	Predicted (TM \geq 0.5)	0.520 (+0.315)	0.385 (+0.195)	0.333 (+0.078)	0.267 (+0.115)
	Predicted (pLDDT \geq 70)	0.528 (+0.305)	0.397 (+0.187)	0.343 (+0.068)	0.267 (+0.115)
	Experimental Structure	0.657 (+0.155)	0.415 (+0.156)	0.402 (+0.030)	0.314 (+0.070)
	The Performance Gap	-0.129 (+0.150)	-0.018 (+0.031)	-0.059 (+0.038)	-0.047 (+0.045)
REINet	Predicted (TM \geq 0.5)	0.584 (+0.147)	0.519 (+0.028)	0.337 (+0.105)	0.340 (+0.032)
	Predicted (pLDDT \geq 70)	0.599 (+0.153)	0.526 (+0.030)	0.335 (+0.104)	0.340 (+0.033)
	Experimental Structure	0.756 (+0.038)	0.535 (+0.013)	0.427 (+0.021)	0.360 (+0.019)
	The Performance Gap	-0.157 (+0.115)	-0.009 (+0.017)	-0.092 (+0.083)	-0.020 (+0.014)

Table 1: Downstream task performance of three inference settings across three encoders. Results in parentheses (bold) are performance gain training with SAO framework, the performance gap is calculated by values of Predicted(Plddt \geq 70) - values of Experimental Structure. Results clearly show SAO is effective across tasks and model-agnostic.

where $i \in \{1, \dots, N_{res}\}$, $j \in \{C, N, O, C_\alpha\}$, T_i and T_i^{true} corresponds to the predicted and ground truth frame (\hat{O}, \mathbf{t}) of residue i . And $T_i^{-1} \circ x_j = \mathcal{O}_i^{-1}(x_j - t_i)$ converts the coordinates of backbone atoms from the global to local frame. More details can be referred to Appendix. C.

Overall Framework

In this subsection, we first illustrate the training process of SAO, and then briefly introduce our model architecture. **Model Training.** In our training process, we first pretrain the protein structure encoder in SAO framework on the corresponding downstream task dataset (i.e., no new sample is added during pretrain). The overall pretraining objective is $L_{SAO} = \gamma_1 L_{align} + \gamma_2 L_{mlm} + \gamma_3 L_{mse}$ as figure 4(a) shows. L_{mlm} is the standard mask language modeling loss (Devlin et al. 2019), and γ is the loss weight (value setups are shown in Appendix. E). Then we fine-tune the encoder on the downstream task using task-specific losses(as shown in fig.4(b)). The algorithmic description and more details, including the mask ratio, are provided in Appendix E.

Frame-Aware 3D Graphformer. Our SAO framework is model-agnostic, allowing for compatibility with various models. Additionally, we propose a simple yet effective encoder named **REINet** that can co-model the 3D geometry of amino acid coordinates and the 1D protein sequences. As the model architecture is not our main concern, we defer its detailed introduction to the Appendix. D.

Experiments

In this section, we first introduce the experimental setup for four standard protein property prediction tasks, including Enzyme Commission number prediction, and Gene Ontology term prediction following (Gligorijević et al. 2021), model architectures, and robust training framework base-lines. We then conduct empirical experiments to demonstrate the effectiveness of the proposed framework SAO. We aim to answer five research questions as follows: **Q1:** Does the protein structure embedding bias problem generally exist in various protein property prediction tasks across different predictors? **Q2:** Are our proposed framework SAO model-

agnostic? **Q3:** How effective is SAO for PGSL-RPA? **Q4:** How do key framework components impact the performance of SAO? **Q5:** How robust is SAO to less accurate or non-AlphaFold predicted protein structures?

Experimental Setups

Downstream Tasks for Evaluation We adopt four tasks proposed in (Gligorijević et al. 2021) as downstream tasks for evaluation. **Enzyme Commission (EC) number prediction** aims to forecast the EC numbers of various proteins, which describe their catalysis of biological activities in a tree structure. It’s a multi-label classification task with 538 categories. **Gene Ontology(GO) term prediction** aims to annotate a protein with GO terms that describe the Molecular Function (MF) of gene products, the Biological Processes (BP) in which those actions occur, and the Cellular Component (CC) where they are present. Thus, GO term prediction is actually composed of three different sub-tasks: **GO-MF**, **GO-BP**, **GO-CC**. And two metrics are employed, including the maximum F-score (Fmax) and AUPR.

Inference Setting To evaluate SAO in realistic scenarios, we choose three inference settings: 1) inference with the experimental structure; 2) inference with predicted protein structure with TMScore \geq 0.5 3) inference with predicted protein structure with pLDDT \geq 70. TMScore and pLDDT are commonly used criteria for selecting accurately predicted protein structures for downstream applications.

Model Architectures We explore the PGSL-RAP problem in various protein structure encoder architectures, including graph neural network-based GearNet_Edge (Zhang et al. 2022), graphormer-based Uni-Mol (Zhou et al. 2022), SE(3)-invariant module IPA (Jumper et al. 2021) and 3D Graphormer REINet. Implementation details can be referred to Appendix. For simplicity, we use REINet as the encoder in subsequent experiments unless otherwise specified.

Baselines and Training We mainly compare SAO with two categories of methods, including three directly training-based methods: 1. vanilla method of training on experimental protein structures of downstream tasks(*Vanilla*); 2.

Inference Setting	Robust Training Framework	EC		GO-MF		GO-CC		GO-BP	
		Fmax	AUPR	Fmax	AUPR	Fmax	AUPR	Fmax	AUPR
TMscore \geq 0.5	Vanilla	0.584	0.506	0.519	0.521	0.337	0.171	0.340	0.227
	TonP	0.731	0.693	0.537	0.547	0.414	0.266	0.371	0.263
	Mixed	0.667	0.620	0.534	0.537	0.440	0.234	0.365	0.257
	RefthenPred	0.562	0.616	0.524	0.524	0.341	0.184	0.345	0.231
	RefandPred	0.667	0.623	0.501	0.505	0.396	0.194	0.337	0.224
	SAO(ours)	0.731	0.699	0.547	0.554	0.442	0.225	0.372	0.247
pLDDT \geq 70	Vanilla	0.599	0.546	0.526	0.518	0.335	0.177	0.340	0.229
	TonP	0.750	0.716	0.544	0.547	0.419	0.273	0.371	0.269
	Mixed	0.675	0.635	0.542	0.534	0.433	0.224	0.366	0.265
	RefthenPred	0.627	0.587	0.531	0.523	0.336	0.188	0.347	0.235
	RefandPred	0.680	0.644	0.511	0.508	0.400	0.212	0.337	0.229
	SAO(ours)	0.752	0.730	0.556	0.557	0.439	0.225	0.373	0.253
Experimental Structure	Vanilla	0.746	0.745	0.535	0.540	0.427	0.251	0.360	0.248
	TonP	0.543	0.430	0.388	0.357	0.334	0.165	0.279	0.173
	Mixed	0.676	0.575	0.520	0.510	0.432	0.215	0.364	0.247
	RefthenPred	0.746	0.745	0.535	0.54	0.427	0.251	0.36	0.248
	RefandPred	0.742	0.730	0.512	0.506	0.425	0.221	0.345	0.227
	SAO(ours)	0.784	0.771	0.548	0.551	0.448	0.241	0.379	0.255

Table 2: Performance comparison across different robust training framework baselines. Our SAO surpasses most baselines across different tasks and inference settings, which shows its effectiveness and advantage of using paired protein structure data. The best results are marked by bold.

training on predicted protein structure (*TonP*); 3. training on the mixture of predicted and experimental protein structure for direct adaptation (*Mixed*). And two protein structure refinement (e.g., improving the similarity of the predicted and experimental structures directly in Euclidean space) based methods are further included: 1. *RefthenPred* which refine the protein structure before predicting its property; 2. *RefandPredict* which refine the protein structure while predicting its property. The *RefthenPred* is like unsupervised GSL, which obtains clean graph structure without supervision signals; we use the SOTA protein structure refinement method ATOMRefine (Wu and Cheng 2022) for refinement. And the *RefandPred* is like supervised GSL, which refines the graph structure under supervision. We add one layer of simplified Structure Module (Jumper et al. 2021) after the structure encoder to refine the structure while predicting properties.

Following (Gligorijević et al. 2021), we use the multi-cutoff split methods for EC and GO tasks to guarantee the test set contains only PDB chains with sequence identity less than 95% to the training set. We pretrain encoders under SAO for 400 epochs on structures of corresponding downstream tasks and then fine-tune them on downstream tasks for specific epochs (EC: 100 epochs, GO-CC: 45 epochs, GO-MF and GO-BP: 100 epochs). Warmup and the exponential learning rate decay schedule are used with a start learning rate of 0.0, a max learning rate of $1e-4$, and a decay factor of 0.99. For other experimental details, interested readers can refer to the Appendix.

Empirical Studies (Q1)

Table 3 shows how well different methods predict certain properties of proteins in various tasks. The methods are trained on experimental protein structures, following

(Zhang et al. 2022). Then, they are tested in three different inference settings, including two settings with predicted protein structures. As can be observed, all three models perform worse when making predictions with predicted protein structures. This suggests a general issue of protein structure embedding bias, which affects protein property prediction tasks across different prediction methods.

Models	Inference Setting	EC	GO-CC
GearNet_Edge	Predicted (TMscore \geq 0.5)	0.769	0.399
	Predicted (pLDDT \geq 70)	0.763	0.409
	Experimental Structure	0.810	0.460
Uni-Mol	Predicted (TMscore \geq 0.5)	0.520	0.333
	Predicted (pLDDT \geq 70)	0.528	0.343
	Experimental Structure	0.657	0.402
IPA-Encoder	Predicted (TMscore \geq 0.5)	0.568	0.383
	Predicted (pLDDT \geq 70)	0.579	0.362
	Experimental Structure	0.711	0.433

Table 3: Fmax on EC and GO-CC across different architectures in different settings. Representative encoders suffer from performance degradation on predicted structures.

Performance for Different Encoders (Q2)

Table 1 shows the performance of SAO framework across different protein structure encoders. It indicates that SAO framework is model-agnostic and effective in improving the prediction accuracy of both experimental and predicted structures. With SAO, encoders can learn to align the embeddings of predicted structures to that of experimental structures and output meaningful embeddings. It’s worth noting that the property prediction performance with predicted

protein structures surpasses that with experimental protein structures in some settings, which shows the value of the idea of directional embedding alignment.

Performance Comparison (Q3)

Table 2 reports the classification performance of our framework and other baselines in experimental and predicted protein structure inference scenarios. We can observe that our proposed SAO outperforms all baselines on 8 out of 8 metrics (2 for each property prediction task). This superior performance benefits from the novel idea of guiding PGSL in embedding space with a self-enhanced learning target by directional embedding alignment.

We make other observations as follows. Firstly, the performances of training predictors on experimental or predicted protein structures are both notably biased. It is even worse when training on predicted protein structures. This observation provides another evidence of protein structure embedding bias. Secondly, SAO can surpass all three directly training-based baselines, indicating SAO can make better use of paired experimental-predicted data. Thirdly, compared to refinement-based baselines (whether supervised or not), our unsupervised embedding-based methods also achieve better results, which shows our effectiveness.

Inference Setting	Robust Training Framework	GO-CC	
		Fmax	AUPR
TMscore \geq 0.5	Vanilla	0.337	0.171
	SAO	0.442	0.226
	w/o structure denoising	0.429	0.206
	w/o mask view augmentation	0.388	0.201
	w/o embedding alignment	0.389	0.203
Experimental Structure	Vanilla	0.427	0.251
	SAO	0.449	0.252
	w/o structure denoising	0.432	0.218
	w/o mask view augmentation	0.412	0.240
	w/o embedding alignment	0.438	0.234
	w/o mask language modeling	0.419	0.238

Table 4: Ablation study for designed components in two inference scenarios. The best metrics are marked in bold.

Ablation Study (Q4)

To study the importance of every component in SAO, we perform an ablation study on every loss term. As shown in Table 4, without structure denoising guidance, the classification performance decrease by 2% on average, indicating this component helps improve the quality of learned embedding and the alignment from predicted embeddings to experimental embeddings. Without embedding alignment, we can find obvious drops in classification performance, especially when performing inference on predicted protein structures. It shows the effectiveness of our directional embedding alignment mechanisms. We can further notice a clear decline without the mask view augmentation component,

indicating its effectiveness in avoiding embedding collapse and improving directional alignment as embedding augmentation. Finally, it is unsurprising that mask language modeling tasks can improve performance on experimental protein structures. Due to the page limit, more ablation studies, including task weights, can be referred to the Appendix F.




Figure 5: Fmax on GO-CC task in different inference settings. AtomRefine refers to inferencing with ATOMRefine-predicted protein structures. Encoders trained under SAO framework can better handle less accurately predicted protein structures and non-AlphaFold predicted structures.

More Results on Predicted Structure (Q5)

Figure 5 shows the performance under less accurate predicted protein structure and non-AlphaFold predicted structure. The results indicate that SAO can make the encoder more robust to less accurate predicted protein structures and generalize well to non-AlphaFold predicted structures.

We attribute its comparative performance to our design of directional embedding alignment. Directly training-based baselines may struggle to generalize to other inference settings, although have comparative performances in accurate AF-predicted structures.

Conclusion

This paper investigates 3D protein graph structure learning and designs a novel framework, SAO, which is capable of leveraging pair data to perform directional embedding alignment. To learn the optimal alignment, we employ bootstrap-based contrastive learning to maximize the agreement between predicted structure embeddings and self-enhanced learning targets. Extensive experiments demonstrate the superiority of SAO.

Social Impact and Limitations This paper identifies an underlying problem of current common practices (i.e., using predicted protein structures from tools like AlphaFold2 as alternatives when experimental protein structures are missing). It might help the community to rethink the usage of massive AlphaFold2 predicted structures. Our study helps improve the property prediction on predicted protein structures. It contributes to the global efforts to transform large amounts of predicted structures into *high-quality* and valuable knowledge for human well-being.

Acknowledgments

This work was supported by National Key R&D Program of China (No. 2022ZD0115100), National Natural Science Foundation of China Project (No. U21A20427), and Project (No. WU2022A009) from the Center of Synthetic Biology and Integrated Bioengineering of Westlake University.

References

- Al-Masri, C.; et al. 2022. Investigating the conformational landscape of AlphaFold2-predicted protein kinase structures. *bioRxiv*.
- Baldassarre, F.; Menéndez Hurtado, D.; Elofsson, A.; and Azizpour, H. 2021. GraphQA: protein model quality assessment using graph convolutional networks. *Bioinformatics*, 37(3): 360–366.
- Chen, X.; and He, K. 2020. Exploring Simple Siamese Representation Learning. arXiv:2011.10566.
- Dauparas, J.; Anishchenko, I.; Bennett, N.; Bai, H.; Ragotte, R. J.; Milles, L. F.; Wicky, B. I.; Courbet, A.; de Haas, R. J.; Bethel, N.; et al. 2022. Robust deep learning-based protein sequence design using ProteinMPNN. *Science*, 378(6615): 49–56.
- Degn, K.; et al. 2023. PDBminer to Find and Annotate Protein Structures for Computational Analysis. *bioRxiv*.
- Devlin, J.; Chang, M.-W.; Lee, K.; and Toutanova, K. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. arXiv:1810.04805.
- Gligorijević, V.; Renfrew, P. D.; Kosciolk, T.; Leman, J. K.; Berenberg, D.; Vatanen, T.; Chandler, C.; Taylor, B. C.; Fisk, I. M.; Vlamakis, H.; et al. 2021. Structure-based protein function prediction using graph convolutional networks. *Nature communications*, 12(1): 1–14.
- Goodfellow, I. J.; Shlens, J.; and Szegedy, C. 2015. Explaining and Harnessing Adversarial Examples. arXiv:1412.6572.
- Grill; et al. 2020. Bootstrap your own latent—a new approach to self-supervised learning. *NeurIPS*, 33: 21271–21284.
- Hsu, C.; Verkuil, R.; Liu, J.; Lin, Z.; Hie, B.; Sercu, T.; Lerer, A.; and Rives, A. 2022. Learning inverse folding from millions of predicted structures. *ICML*.
- Hu, B.; Xia, J.; Zheng, J.; Tan, C.; Huang, Y.; Xu, Y.; and Li, S. Z. 2022. Protein Language Models and Structure Prediction: Connection and Progression. arXiv:2211.16742.
- Huang, Y.; Wu, L.; Lin, H.; Zheng, J.; Wang, G.; and Li, S. Z. 2023. Data-Efficient Protein 3D Geometric Pretraining via Refinement of Diffused Protein Structure Decoy. *arXiv preprint arXiv:2302.10888*.
- Ingraham, J.; Garg, V.; Barzilay, R.; and Jaakkola, T. 2019. Generative models for graph-based protein design. *Advances in neural information processing systems*, 32.
- Jin, W.; Ma, Y.; Liu, X.; Tang, X.; Wang, S.; and Tang, J. 2020. Graph Structure Learning for Robust Graph Neural Networks. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, KDD '20, 66–74. New York, NY, USA: Association for Computing Machinery. ISBN 9781450379984.
- Jing, B.; Eismann, S.; Suriana, P.; Townshend, R. J.; and Dror, R. 2020. Learning from protein structure with geometric vector perceptrons. *arXiv preprint arXiv:2009.01411*.
- Jumper, J.; Evans, R.; Pritzel, A.; Green, T.; Figurnov, M.; Ronneberger, O.; Tunyasuvunakool, K.; Bates, R.; Židek, A.; Potapenko, A.; et al. 2021. Highly accurate protein structure prediction with AlphaFold. *Nature*, 596(7873): 583–589.
- Lin, H.; Huang, Y.; Liu, M.; Li, X.; Ji, S.; and Li, S. Z. 2022a. DiffBP: Generative Diffusion of 3D Molecules for Target Protein Binding. *arXiv preprint arXiv:2211.11214*.
- Lin, H.; Huang, Y.; Zhang, H.; Wu, L.; Li, S.; Chen, Z.; and Li, S. Z. 2023a. Functional-Group-Based Diffusion for Pocket-Specific Molecule Generation and Elaboration. arXiv:2306.13769.
- Lin, H.; Wu, L.; Xu, Y.; Huang, Y.; Li, S.; Zhao, G.; and Li, S. Z. 2023b. Non-equispaced Fourier Neural Solvers for PDEs. arXiv:2212.04689.
- Lin, Z.; Akin, H.; Rao, R.; Hie, B.; Zhu, Z.; Lu, W.; Smetanin, N.; dos Santos Costa, A.; Fazel-Zarandi, M.; Sercu, T.; Candido, S.; et al. 2022b. Language models of protein sequences at the scale of evolution enable accurate structure prediction. *bioRxiv*.
- Liu, Y.; Zheng, Y.; Zhang, D.; Chen, H.; Peng, H.; and Pan, S. 2022. Towards Unsupervised Deep Graph Structure Learning. In *Proceedings of the ACM Web Conference 2022*, WWW '22, 1392–1403. New York, NY, USA: Association for Computing Machinery. ISBN 9781450390965.
- Liu, Z.; Wang, G.; Wang, J.; Zheng, J.; and Li, S. Z. 2023. Co-modeling the Sequential and Graphical Routes for Peptide Representation Learning. arXiv:2310.02964.
- Mansoor, S.; Baek, M.; Madan, U.; and Horvitz, E. 2021. Toward more general embeddings for protein design: Harnessing joint representations of sequence and structure. *bioRxiv*.
- Pan, Q.; Nguyen, T. B.; Ascher, D. B.; and Pires, D. E. V. 2022. Systematic evaluation of computational tools to predict the effects of mutations on protein stability in the absence of experimental structures. *Briefings in Bioinformatics*, 23(2): bbac025.
- Roche, R.; Moussad, B.; Shuvo, M. H.; and Bhattacharya, D. 2022. E(3) equivariant graph neural networks for robust and accurate protein–protein interaction site prediction. *bioRxiv*.
- Rossi Sebastiano, M.; Ermondi, G.; Hadano, S.; and Caron, G. 2022. AI-based protein structure databases have the potential to accelerate rare diseases research: AlphaFoldDB and the case of IAHSPL/Alsin. *Drug Discovery Today*, 27(6): 1652–1660.
- Shi, Q.; Chen, W.; Huang, S.; Wang, Y.; and Xue, Z. 2019. Deep learning for mining protein data. *Briefings in Bioinformatics*, 22(1): 194–218.
- Wu, L.; Huang, Y.; Lin, H.; Liu, Z.; Fan, T.; and Li, S. Z. 2022a. Automated Graph Self-supervised Learning via Multi-teacher Knowledge Distillation. arXiv:2210.02099.
- Wu, L.; Lin, H.; Huang, Y.; Fan, T.; and Li, S. Z. 2023a. Extracting Low-/High-Frequency Knowledge from Graph

- Neural Networks and Injecting it into MLPs: An Effective GNN-to-MLP Distillation Framework. *arXiv:2305.10758*.
- Wu, L.; Lin, H.; Huang, Y.; and Li, S. Z. 2022b. Knowledge Distillation Improves Graph Structure Augmentation for Graph Neural Networks. In Koyejo, S.; Mohamed, S.; Agarwal, A.; Belgrave, D.; Cho, K.; and Oh, A., eds., *Advances in Neural Information Processing Systems*, volume 35, 11815–11827. Curran Associates, Inc.
- Wu, L.; Lin, H.; Huang, Y.; and Li, S. Z. 2023b. Quantifying the Knowledge in GNNs for Reliable Distillation into MLPs. *arXiv:2306.05628*.
- Wu, L.; Lin, H.; Tan, C.; Gao, Z.; and Li, S. Z. 2021. Self-supervised learning on graphs: Contrastive, generative, or predictive. *IEEE Transactions on Knowledge and Data Engineering*.
- Wu, L.; Xia, J.; Gao, Z.; Lin, H.; Tan, C.; and Li, S. Z. 2022c. Graphmixup: Improving class-imbalanced node classification by reinforcement mixup and self-supervised context prediction. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, 519–535. Springer.
- Wu, T.; and Cheng, J. 2022. Atomic protein structure refinement using all-atom graph representations and SE (3)-equivariant graph neural networks. *bioRxiv*.
- You, Y.; and Shen, Y. 2022. Cross-modality and self-supervised protein embedding for compound–protein affinity and contact prediction. *Bioinformatics*, 38(Supplement_2): ii68–ii74.
- Zhang, Z.; et al. 2022. Protein representation learning by geometric structure pretraining. *arXiv*.
- Zheng, J.; Wang, G.; Huang, Y.; Hu, B.; Li, S.; Tan, C.; Fan, X.; and Li, S. Z. 2023. Lightweight Contrastive Protein Structure-Sequence Transformation. *arXiv:2303.11783*.
- Zhou, G.; Gao, Z.; Ding, Q.; Zheng, H.; Xu, H.; Wei, Z.; Zhang, L.; and Ke, G. 2022. Uni-Mol: A Universal 3D Molecular Representation Learning Framework. *ChemRxiv*.