Mechanism-Aware Protein-Protein Interaction Prediction via Contact-Guided Dual Attention on Protein Language Models

Shuchen Deng, Xuanjun Wan, Zichun Mu, Sheng-you Huang*, Chengfei Yan*

School of Physics, Huazhong University of Science and Technology, China

Correspondence: chengfeiyan@hust.edu.cn, huangsy@hust.edu.cn

Abstract

Accurate prediction of protein-protein interactions (PPIs) is essential for understanding cellular mechanisms and disease. However, existing deep learning models often fail to generalize to novel proteins with low sequence similarity to training data. Here, we present a mechanism-aware deep learning framework to enhance generalization by incorporating residue-level interaction modeling. Our model integrates a dual-attention-based PPI prediction module with an inter-protein contact predictor based on protein language model-embedded geometric graphs. Joint training on structure-informed PPIs from the Protein Data Bank (PDB) enables the model to predict both interactions and the key residue pairs that mediate them. The model generalizes well to unseen proteins across species in the high-quality interactome dataset (HINT). Fine-tuning on *Homo sapiens* PPI data from HINT further improves the performance. Compared to D-SCRIPT, Topsy-Turvy, and TT3D, our model, referred to as PLMDA-PPI, achieves superior accuracy and robustness under strict sequence dissimilarity settings. This work provides a generalizable framework for biomolecular interaction prediction through the integration of structural and non-structural interaction data.

Introduction

Protein–protein interactions (PPIs) play essential roles in nearly all biological processes. Identifying PPIs involved in these processes is critical for understanding life at the molecular level, investigating disease mechanisms, and facilitating drug discovery[1-4]. However, experimental identification of PPIs is both time-consuming and costly, and currently, interaction data remain severely incomplete for most species. Even for humans, the most extensively studied organism, experimentally validated PPIs represent only a small fraction of the complete interactome[3,5-7]. Moreover, a significant portion of available PPI data originates from high-throughput screening techniques, which are often associated with high false-positive rates[8,9]. Therefore, developing effective computational methods for PPI prediction holds substantial practical importance.

In recent years, the application of deep learning has led to remarkable advances in this field[10-13]. Nevertheless, existing approaches still face critical limitations that hinder their broader applicability and accuracy. For instance, to train deep learning models for protein-protein interaction (PPI) prediction, a large amount of high-quality interaction data is required. Currently, most training data for these models are derived from literature-mined PPI databases such as STRING[14]. However, these databases contain a high proportion of false positives and are biased toward well-studied proteins, leading to biases in both model training and evaluation[15-17]. For

instance, evaluation results by Bernett et al.[15] showed that when sequence similarity between training and test sets is minimized, almost all previously reported PPI prediction models which claimed high predictive performance lose their predictive capability. Further analysis revealed that the previously reported high performance was largely attributed to information leakage caused by random splitting of PPI datasets into training and test sets, where proteins in both sets often exhibited significant overlap or high sequence similarity. Consequently, these models primarily relied on shortcuts , such as sequence similarity between proteins in the training and test sets or node degrees within the training PPI network, rather than truly learning generalizable features. As a result, they fail to generalize to novel proteins which have low sequence similarities to those seen during training.

A PPI prediction model with strong generalization ability is expected to capture the underlying mechanisms of protein-protein interactions (PPIs). Given that PPIs are predominantly driven by physical interactions at the residue level, we hypothesize that a model's ability to accurately uncover these residue-level interaction mechanisms reflects its understanding of the principles governing PPIs. Based on this assumption, we developed an inter-protein residue-residue contacts guided deep learning framework for PPI prediction. Specifically, given a pair of proteins, our model first infers the inter-protein residue-residue contact probability map using protein language model (PLM) embedded geometric

graphs[18], which is then fused with sequence-based PLM representations through a dual-attention network[19,20] to predict the interaction probability of the two proteins. We first trained our model using structure-informed PPIs collected from Protein Data Bank (PDB)[21], in which the inter-protein contact prediction module and PPI prediction module were jointly trained using a linear combination of the inter-protein contact loss and the PPI loss. Validated on structure-informed PPIs formed by proteins with low sequence similarities to the training set, we show the model not only can identify PPIs with high accuracy, but also the inter-protein contact prediction module can also identify key residue pairs for maintaining the PPIs. We further tested the model on high-quality interactome (HINT)[22] data for *Homo sapiens, Mus musculus, Arabidopsis thaliana, Drosophila melanogaster* and *Schizosaccharomyces pombe*, and the result shows that the model trained on the PDB dataset can well generalized to the HINT dataset. To address the data collection standards discrepancy between PDB and HINT, we further fine-tuned the model using the *Homo sapiens* PPIs from HINT. The evaluation results on multiple species show that the model performance can be significantly improved through the fine-tuning, and even when we minimize the sequence similarity between the training and test sets, the model performance can still be largely maintained. To provide a comparison, we also evaluated three state-of-the-art PPI prediction models including D-SCRIPT[23], Topsy-Turvy[24], and TT3D[25], the result

shows that our model referred to as PLMDA-PPI consistently outperforms these models by large margins.

Results

1. Overview of PLMDA-PPI

1.1 The model architecture of PLMDA-PPI

The model architecture of PLMDA-PPI is shown in Figure 1a~b. Given a pair of proteins, PLMDA-PPI requires sequences, structures (experimental or predicted structures) and multiple sequence alignments (MSAs) as the inputs. First, following the protocol of our previous work, the protein language model-embedded geometric graphs are constructed based the structures and MSAs of the proteins, from which the protein embeddings were obtained. Then, the pair representation of the two proteins were built through outer concatenation, followed by dimensional hybrid residual network[26] formed by 12 residual blocks to transform the pair representation to the predicted inter-protein contact map, which represents the residue-residue contact probability between the two proteins. Finally, the predicted inter-protein contact map is used to integrate the protein sequence embeddings (generated from ESM-1b[27]) to generate the PPI score (i.e. the interaction probability of the two proteins) using a dual-attention network modified from Visual Question Answering (VQA) tasks[20,28] (Figure 1b). A more detailed description of PLMDA-PPI can be found in the *Methods* section.

1.2 The training and evaluation protocols of PLMDA-PPI

In Figure 1c, we show the protocols we used to train and evaluate PLMDA-PPI. Specifically, we first trained the model using a non-redundant, structure-informed PPI dataset collected from the PDB. The inter-protein contact prediction module and the PPI prediction module were trained simultaneously using a linear combination of the inter-protein contact loss and the PPI loss. The structure-informed PPI dataset was divided into training and validation sets, ensuring that no protein in the training set shares more than 40% sequence identity with any protein in the validation set. This constraint was applied to minimize the risk of overfitting caused by data leakage (a more detailed discussion of the data leakage issue can be found in Section 3.1 of the Methods). We did not generate an additional test set for the structure-informed PPIs due to the limited number of non-redundant PPIs available from the PDB. Instead, we used literature-curated, high-quality binary PPIs from HINT for *Homo sapiens*, *Mus musculus*, *Arabidopsis thaliana*, *Drosophila melanogaster*, and *Schizosaccharomyces pombe* to further evaluate the model. Considering the discrepancy in PPI collection criteria between the PDB and HINT, we further fine-tuned the prediction model using *Homo sapiens* PPIs from HINT through transfer learning. The *Homo sapiens* binary PPIs from HINT were divided into training, validation, and test sets, ensuring that proteins in the three subsets do not share more than 40% sequence

identity. Finally, the *Homo sapiens* test set, along with the PPI datasets for *Mus musculus*, *Arabidopsis thaliana*, *Drosophila melanogaster*, and *Schizosaccharomyces pombe* from HINT, were used to evaluate the performance of the fine-tuned models. It is worth noting that for PPI data derived from the PDB database, we extracted input features using experimentally determined monomer structures from the corresponding complex structures during both training and testing. In contrast, for PPI data from the HINT dataset, where experimental monomer structures are not always available, we employed predicted monomer models from AlphaFoldDB[29] as substitutes.

For all datasets used in training, validation, and testing, negative samples (i.e., non-interacting protein pairs) were generated by randomly pairing protein monomers within each dataset, after filtering out protein pairs that share interacting partners. A 10:1 negative-to-positive sample ratio was adopted for each dataset. A more detailed description of our training and evaluation protocols is provided in the *Method*s section.

2. Performance of PLMDA-PPI with initial training on the structure-informed PPI dataset from PDB

2.1 Validation of the model using structure-informed PPIs

Following initial model training, we evaluated the performance of PLMDA-PPI on the validation set derived from structure-informed PPIs in the PDB. The evaluation involved two tasks: (1) PPI prediction as the

primary objective, and (2) inter-protein contact prediction for protein pairs forming physical interactions (i.e. positive samples) as an auxiliary task. For comparison, we trained three reference models: D-SCRIPT, Topsy-Turvy, and TT3D, using the same dataset. These models share a common architecture that includes an intermediate module for predicting inter-protein contact maps, which are subsequently summarized into a PPI score. The key differences lie in their input representations and training procedures: D-SCRIPT is purely sequence-based; Topsy-Turvy is also sequence-based but incorporates PPI network-derived index during training; and TT3D extends Topsy-Turvy by integrating a 21-letter discretized structural alphabet (3Di), derived using Foldseek[30], to provide richer structural representations.

These models differ from PLMDA-PPI in two important aspects in addition to input features. First, they compute the final PPI score using only a customized max-pooling operation over the predicted contact maps. In contrast, PLMDA-PPI employs dual-attention network to fuse the predicted residue-level contact maps with protein language model embeddings, resulting in more informative PPI representations. Second, during training, the three models only use a sparsity-inducing regularization term derived from the predicted contact maps in the loss , but they do not incorporate ground-truth contact maps as supervision. PLMDA-PPI, on the other hand, introduces a multi-task loss that explicitly

trains the contact prediction module using inter-protein contact maps calculated based experimental structures as labels.

In Table 1, we present the performance of PLMDA-PPI and three reference models on the validation set of structure-informed PPIs. The evaluation covers two aspects: PPI prediction and inter-protein contact prediction. For PPI prediction, we report the Area Under the Precision-Recall Curve (AUPRC) and the Area Under the Receiver Operating Characteristic Curve (AUROC). For contact prediction, we evaluate the mean precision of the top 50 predicted residue-residue contacts. Given the 10:1 negative-to-positive sample ratio in the validation set, a null model would achieve an AUPRC of 0.1 and an AUROC of 0.5. Since AUPRC is generally more informative for imbalanced datasets[31], we additionally visualize the Precision-Recall curves for all models in Figure 2a.

As shown in both Table 1 and Figure 2a, PLMDA-PPI shows encouraging performances for PPI prediction with an AUPRC 0.637 and AUROC 0.911, dramatically outperforming the reference models. For inter-protein contact prediction, PLMDA-PPI shows a mean precision of 0.287, although there is still substantial room for improvement, it clearly outperforms the reference models, which exhibit minimal capability in predicting inter-protein contacts. It is worth noting that in PLMDA-PPI, the inter-protein contact prediction serves only as an auxiliary task within an intermediate module of the model, primarily designed to enhance the

generalization ability of the main task — PPI prediction. Therefore, it is acceptable that its performance on inter-protein contact prediction does not match that of specialized models dedicated to this task. In Figure 2b, we compare the distribution of the mean contact probability among the top 50 predicted residue pairs for interacting and non-interacting protein pairs. Interacting pairs exhibit significantly stronger contact signals, further validating the residue-residue contact signals are employed to guide the PPI prediction (Figure 1b). Finally, Figures 2c and 2d illustrate an example of the experimental inter-protein contact map and the predicted contact map by PLMDA-PPI, respectively. The predicted map clearly identifies key residue-residue contacts that mediate the interaction, demonstrating the model's effectiveness in capturing physically meaningful contact patterns.

2.2 Evaluation of the model using PPIs of multiple species from HINT

We further evaluated performances of PLMDA-PPI and the three reference models on the high-quality literature-curated binary interactions for *Homo sapiens, Mus musculus, Arabidopsis thaliana, Drosophila melanogaster and Schizosaccharomyces pombe* from HINT database. It should be noted that most PPIs in HINT database do not have structural data available. Therefore, we only evaluated their performances on PPI prediction. In Table 2, we show the model performances in terms of

AUPRC and AUROC on the HINT PPI datasets, and in Figure 3a-d, we show the Precision-Recall curves. As it can be seen that PLMDA-PPI consistently outperforms the references models by large margins (also see Figure S1).

The HINT and PDB databases follow different criteria for collecting protein-protein interaction (PPI) data. For instance, the PDB focuses on interactions with experimentally resolved complex structures, whereas many interactions in HINT—such as transient interactions—are difficult to characterize structurally and thus lack associated structural information. As a result, most PPIs in the HINT database do not have corresponding entries in the PDB. To further evaluate our model's ability—trained on PDB-based data—to predict interactions lacking structural references, we excluded from the HINT dataset all interactions with any PDB reference. For each species, we then generated a new set of negative samples via random pairing at a 10:1 ratio (negatives to positives), and re-assessed the performances of these models using this revised dataset, which were also shown in Table 2 and Figure 3f. As it can be seen that although the predictive performance of PLMDA-PPI declined across datasets other than *Arabidopsis thaliana*, it still maintained a reasonable level of accuracy. The other three reference models exhibited similar trends. Nevertheless, PLMDA-PPI consistently outperformed these models on the revised datasets. These findings indicate that models trained solely on PDB-

derived data may not generalize well to PPIs whose structures are difficult to resolve experimentally. Therefore, incorporating PPI data without structural information into model training is necessary to further improve predictive performance.

Additionally, we applied another filtering step by excluding all PPIs in the HINT datasets that involve any protein with more than 40% sequence identity to any protein in the training set. For each species, we re-generated negative samples via random pairing, maintaining a 10:1 negative-to-positive ratio. The resulting filtered datasets were then used to evaluate model performance on novel proteins that were absent from the training set and shared low sequence similarity with training proteins. The results are also shown in Table 2 and Figure 3f. As observed, PLMDA-PPI does not show any systematic performance degradation on the filtered datasets; instead, modest improvements are observed on *Homo sapiens, Drosophila melanogaster, and Schizosaccharomyces pombe*. Similar patterns are seen across the other three reference models. These results suggest that the models do not rely on sequence similarity for interaction prediction. This is consistent with the fact that the PDB and HINT datasets differ substantially in their protein coverage. The performance gains may result from filtering out particularly challenging prediction cases.

3. Performance of PLMDA-PPI after transfer learning on the HINT *Homo sapiens* PPI dataset

3.1 Model evaluation on HINT PPI datasets

Given the discrepancy in data collection criteria between PDB and HINT, we adopted a transfer learning strategy to enhance the performance of our models in predicting protein-protein interactions (PPIs) under the HINT standard. Specifically, we fine-tuned the models initially trained on PDB data—including PLMDA-PPI and three reference models—using the *Homo sapiens* subset from the HINT dataset. To evaluate their performance after fine-tuning, we tested the models on the *Homo sapiens* test set, as well as on datasets from *Mus musculus*, *Arabidopsis thaliana*, *Drosophila melanogaster*, and *Schizosaccharomyces pombe* in HINT. As shown in Table 3 and Figure 4a–e, PLMDA-PPI consistently outperformed the three reference models across all five datasets, and transfer learning led to a substantial improvement in its predictive performance compared to the pre-transfer version (also see Figure S2). In addition, we analyzed the residue-level contact probability maps of the positive (interacting protein pairs) and negative (non-interacting protein pairs) samples across datasets from different species. Overall, the positive samples exhibit significantly stronger residue contact signals compared to the negative ones. This finding further supports that residue-level contact signals are leveraged to guide the PPI predictions (Figure S3).

3.2 Re-assessing model performance after excluding proteins with high sequence similarity to the training data

In the above evaluations, only the test set of *Homo sapiens* PPIs was strictly filtered to exclude any PPIs involving proteins with more than 40% sequence identity to those in the training set of transfer learning in the beginning of the dataset preparation. For the PPI data from other species, no such sequence similarity-based filtering was applied. Therefore, the model may still rely on sequence similarity to the training data when making predictions on these datasets. To provide a more objective evaluation of the models' performances in predicting interactions between novel proteins that were not present in the training set and share low sequence similarity with training proteins, we performed a stringent filtering on the HINT datasets. Specifically, we removed all PPIs involving any protein with a sequence identity greater than 40% to any protein in the training set of *Homo sapiens* PPIs. Based on the remaining interaction pairs of each species, we re-generated a negative set by randomly pairing proteins at a 10:1 ratio relative to the number of positive samples. The filtered datasets were then used to re-assess the models' performances, which were also shown in Table 3 and Figure 4g-h.

As shown in the table and figure, unlike the previous filtering strategy based on the PDB dataset, all models—including ours and the three reference models—exhibited a certain degree of performance drop on the filtered datasets. This may be attributed to the evolutionary relationships between some PPIs from non-*Homo sapiens* species in HINT and those

from *Homo sapiens*, which could allow models to rely on sequence similarity for prediction. Nevertheless, despite the decline in performance, our model still maintains considerable predictive power on the filtered datasets. In contrast, the performance of the other three reference models dropped sharply, approaching the level of a null model (AUPRC =0.1)，which is consistent with the findings of Bernett et al. in their previous study[15]. These results further demonstrate that our model possesses significantly better generalization ability compared to the reference models.

In addition, we also evaluated the model before transfer learning on the same dataset. The results show that the model after transfer learning still significantly outperforms the one without transfer learning in most cases. The only exception is on the *Schizosaccharomyces pombe* dataset, where the AUPRC slightly decreased (from 0.342 to 0.338); however, the AUROC improved from 0.794 to 0.826 (Supplementary Table S1). These findings suggest that the performance gains from transfer learning are not merely due to the presence of proteins in the test set that share high sequence similarity with those in the training set.

3.3 Re-assessing the model on the structure-informed PPI dataset

To assess the generalization ability of the models, we evaluated whether the models fine-tuned on the HINT dataset could still maintain their predictive performances on the original PDB-based structure-informed PPI dataset. Accordingly, we further twested the transfer-learned

models on the previously used validation set of structure-informed PPIs. In Table 5, we present the performance of the model after transfer learning on this dataset, covering both the primary task of PPI prediction and the auxiliary task of inter-protein contact prediction. As shown in the table, our model significantly outperforms the others on both tasks. Notably, our model even exhibits a slight improvement in PPI prediction performance, whereas the other models suffer a substantial drop in accuracy (Figure 5a-d). For the auxiliary task of contact prediction, our model experiences a moderate performance decline after transfer learning (with the mean precision of the top 50 predicted contacts dropping from 0.287 to 0.230), yet it still maintains considerable predictive power. In contrast, the reference models exhibit virtually no ability to predict residue-level contacts. These results further demonstrate the superior generalization capability of our model compared to the others.

4 Ablation studies

We further conducted ablation studies to systematically assess the contributions of the key components and training strategies of PLMDA-PPI, covering both the architectural design and training approach. Specifically, we retrained the model after removing major modules using the structure-informed PPI training set derived from PDB and evaluated performance on the corresponding validation set. When we removed the MSA and structural representations embedded in the geometric graph by

the protein language models—retaining only the geometric graph—and replaced the sequence-based protein language model embeddings in the interaction prediction module with geometric node features, the model's performance dropped substantially (AUPRC decreased from 0.637 to 0.229). Restoring the protein representations in the interaction module to language model embeddings led to a significant performance recovery (AUPRC = 0.593), although still inferior to the full model. These results highlight the crucial role of language model representations in our framework. Additionally, replacing the dual-attention network in our model with the customized max pooling module used in the baseline model also led to a notable performance decline (AUPRC = 0.411), underscoring the importance of the dual-attention network module (Fig. 6a).

Moreover, when training the model without the multitask loss for residue-residue contact prediction—i.e., optimizing only for the primary interaction prediction task—the performance on the PDB validation set showed a negligible improvement (AUPRC = 0.643). However, the same model exhibited consistently worse performance across the HINT multi-species datasets. This result indicates that enabling the model to predict inter-protein contacts contributes significantly to its generalization ability (Fig. 6b).

Finally, we trained the model directly on the training set of *Homo sapiens* PPIs from HINT without any pretraining on PDB data. Across all

species, this directly trained model performed significantly worse than the fine-tuned from the PDB-pretrained model, further demonstrating the importance of pretraining on structure-informed datasets for achieving strong cross-species generalization (Fig. 6c and Supplementary Table S2).

Conclusion

In this study, we sought to enhance the generalization ability of deep learning models for PPI prediction, addressing a key limitation of existing models: their inability to accurately predict novel proteins that are absent from the training set and exhibit low sequence similarity to training proteins. The central hypothesis guiding our work is that a robust model with strong generalization capabilities must not only predict interactions but also understand the underlying residue-level interaction mechanisms that sustain these interactions. Based on this hypothesis, we developed a model that integrates a dual-attention network-based PPI prediction module into our previously developed inter-protein contact prediction framework based on protein language mode-embedded geometric graphs. By simultaneously training both modules with structural information of PPIs, our model is able to predict PPIs and, crucially, identify the residue pairs that maintain these interactions.

Our experimental results confirm the validity of this hypothesis. The model, trained on PDB data, successfully predicts PPIs of multiple species from the HINT database. Fine-tuning on PPI data of *Homo sapiens* from

the HINT database further enhances the model's performance. When compared with other state-of-the-art models such as D-SCRIPT, TOPSY-TURVY, and TT3D, our approach not only achieves superior accuracy but also demonstrates remarkable generalization to novel proteins that were absent from the training data and exhibit low sequence similarity to known proteins.

The impact of this work extends beyond protein–protein interaction (PPI) prediction. By formalizing the integration of structural and non-structural data in the context of biomolecular interactions, our model offers a novel and adaptable framework for predicting other interaction types, such as protein–small molecule and protein–nucleic acid interactions. These interactions, like PPIs, often involve heterogeneous data—some with structural information, others without. As hypothesized, combining both data types enables deeper insights into the physical principles underlying these interactions, ultimately improving the generalization capability of predictive models for unseen cases.

While our model substantially outperforms existing approaches in terms of both accuracy and generalizability, there remain opportunities for further enhancement. A key challenge lies in narrowing the performance gap between our inter-protein contact prediction module and advanced end-to-end structure prediction tools including AlphaFold-Multimer[32] and AlphaFold3[33]. Future work could explore integrating AlphaFold's

structural modeling capabilities into our dual-attention-based PPI framework to further boost predictive power. However, due to AlphaFold's high computational cost, improving its efficiency will be critical for enabling large-scale interaction screening.

In summary, our study demonstrates that incorporating residue-level interaction mechanisms, guided by our central hypothesis, markedly improves both the accuracy and generalization of PPI prediction. Beyond advancing PPI prediction, this work lays the foundation for extending such frameworks to a broader range of complex biomolecular interactions.

Methods

1. Dataset preparation

1.1 PPI database selection

The PDB database contains a large number of experimentally resolved protein multimer structures, providing direct access to physically interacting protein pairs and their complex structural information. Therefore, we used the PDB database as a source of PPI data with known structures. However, due to the high cost and time-consuming nature of experimental structure determination, the coverage of non-redundant PPIs for a single species in the PDB database remains limited. This necessitates the integration of other high-quality PPI databases to expand the dataset.

Currently, multiple databases store protein interaction information categorized by species. For example, HuRi[3] and BioPlex[4] have

systematically screened human PPIs using yeast two-hybrid (Y2H) and affinity purification-mass spectrometry (AP-MS), respectively. However, these high-throughput screening methods often have a high false positive rate, and AP-MS can only detect co-complexed proteins without determining whether direct physical interactions occur. Additionally, computationally curated databases such as STRING[14] and HINT[22] extract PPI data from literature using various mining algorithms and filtering criteria. Due to differences in these approaches, the number and types of PPIs included in each database vary significantly for the same species. In particular, the *"Physical Interaction"* category of PPIs in the STRING database were widely used for training and testing PPI prediction models, but it does not distinguish between direct interactions and proteins that merely co-occur within a complex[14]. In contrast, the HINT database explicitly differentiates between direct and indirect interactions and provides high-quality PPI datasets that have been confirmed by at least two independent small-scale experimental studies (the high-quality literature-curated datasets). Analysis by Donald Petrey et al.[9] suggests that the high-quality literature-curated binary interaction datasets from the HINT database is the best available source for direct physical protein interactions. Therefore, we also leveraged these datasets for model development and testing in this study.

1.2 Structure-informed PPI collection from PDB

To construct a dataset of structure-informed PPIs, we filtered PDB entries deposited before October 24, 2022, using the following criteria: 1) Experimental Method: X-RAY DIFFRACTION, 2) Data Collection Resolution: $\leq$4Å, 3) Entry Polymer Composition: heteromeric protein, 4) Chain count: $\geq$2. This resulted in 43,983 PDB entries. We then extracted protein-protein pairs from their biological assemblies, yielding a total of 438,554 pairs. To remove redundancy, we clustered all the 877,108 protein sequences (from 438,554 pairs) using CD-HIT[34] with parameters "*-c 0.4 - n 0.2 -g 1*", grouping sequences at a 40% identity threshold, with which each protein was assigned a unique cluster ID (e.g. cluster 0, cluster 1, …), and each protein-protein pair was labeled with a cluster pair (e.g. cluster 0-cluster 1). The protein-protein pairs belonging to the same cluster pair (note: cluster n-cluster m and cluster m-cluster n were considered as the same pair) were considered as redundant. We computed inter-protein contact maps (defining residue contacts as heavy-atom distances $\leq$8Å) for all the protein-protein pairs, and only kept protein-protein pairs with more than 50 inter-protein contacts as protein-protein interactions (PPIs). Among redundant PPIs (i.e., those in the same cluster pair), we retained only the one with the highest number of contacts. We further filtered these PPIs with the following criteria:1) the chain length of each protein between 30 and 521, 2) $\leq$ 20% missing residues in the structure of each chain, 3) removal of self-interacting proteins (i.e., proteins from the same cluster) .

After these steps, we obtained 6,189 structure-informed, non-redundant PPIs, comprising 11,119 unique protein monomers from 5952 sequence clusters.

1.3 PPI collection for multiple species from HINT

We obtained high-quality, literature-curated binary PPI datasets for *Homo sapiens, Mus musculus, Arabidopsis thaliana*, *Drosophila melanogaster*, and *Schizosaccharomyces pombe* from the HINT database. We applied the following filtering steps: 1) removal of self-interacting proteins, 2) excluding PPIs involving proteins with chain lengths shorter than 30 or longer than 1024, 3) discarding PPIs formed by protein isoforms, to further refine the datasets. After filtering, the datasets contained 23,311 PPIs for *Homo sapiens,* 2,369 PPIs for *Mus musculus,* 2,789 PPIs *for Arabidopsis thaliana,* 1,336 PPIs *for Drosophila melanogaster and* 791 PPIs *for Schizosaccharomyces pombe*.

2. Model construction of PLMDA-PPI

2.1 Input features

2.1.1 Protein geometric graph

Given the 3D structure of a protein, its geometric graph is constructed following the same protocol as PLMGraph-Inter[18], our previous model for inter-protein contact prediction. Each residue is represented as a node, and edges are formed between residues with a $C\alpha$ distance ≤ 18Å. The graph incorporates geometric features for each node and edge in the graph,

including inter-atom distances (scalars) and orientations (vectors) derived from the 3D structure. To ensure SE(3) invariance, the geometric vectors are extracted using local coordinate systems. Further details on the geometric graph representation module can be found in our previous work.

2.1.2 Embeddings of sequence, MSA, and structure

ESM-1b[27] is employed to generate embeddings for protein sequences, and ESM-IF[35] is employed to generate embeddings for protein structures. Jackhmmer[36] is used to search the Uniref90[37] database to construct the multiple sequence alignment (MSA) (using default parameters) for a given protein sequence. The MSA embeddings are formed by two components: the embeddings generated by ESM-MSA-1b[38] and the position-specific scoring matrices (PSSM) derived from the statistics of multiple sequence alignment.

2.2 Protein encoder

The protocol used PLMGraph-Inter is also used to encode proteins in this study. Specifically, the scalar features of each node in the protein graph are formed by concatenating MSA embeddings, structure embeddings, and scalar geometric features of individual residues. The protein graph's scalar and vector features are then processed through three layers of message passing-based graph neural network formed by geometric vector perceptrons (GVP)[39], which consist of a scalar track and vector track, allowing rotationally invariant transformations of the scalar

features and rotationally equivariant transformations of the vector features. Finally, the scalar and vector features of each node are stitched to form a 1D representation of the protein. A more detailed description of the graph transformation module can also be found in our previous work[18]. A key difference between PLMDA-PPI and PLMGraph-Inter is that PLMDA-PPI does not incorporate protein sequence embedding (from ESM-1b) into the protein graph. Instead, the sequence embeddings are moved to the interaction prediction module, as we found that including them in both the protein graphs and the interaction module significantly increased the likelihood of false positive predictions.

2.3 Inter-protein contact prediction module

The inter-protein contact prediction module takes the protein 1D representation output by the protein encoder for each protein and models the interaction between residues of two proteins. For the two protein 1D representations $A \in \mathbb{R}^{L_1 \times d}$ and $B \in \mathbb{R}^{L_2 \times d}$ (L_1 and L_2 denote the protein lengths, d equaling to 448 denotes the representation dimension), We perform the following broadcast and concatenation operations (Equation (1)~(3)) to generate 2D representation $C \in \mathbb{R}^{2d \times L_1 \times L_2}$ for the protein pair:

$$A = Broadcast(Unsqueeze(\mathbf{A})) \in \mathbb{R}^{L_1 \times L_2 \times 448} \quad (1)$$

$$B = Broadcast(Unsqueeze(\mathbf{B})) \in \mathbb{R}^{L_1 \times L_2 \times 448} \quad (2)$$

$$C = Transpose(Concat(\mathbf{A}, \mathbf{B})) \in \mathbb{R}^{896 \times L_1 \times L_2} \quad (3)$$

The 2D representation of protein pair are then transformed by

dimensional hybrid residual block to generate the predicted inter-protein contact map. Specifically, we first use a convolution layer with kernel size of 1*1 to reduce the number of channels of C from 896 to 96, which are then transformed successively by 12 dimensional hybrid residual blocks[26,40], and another convolution layer with kernel size of 1*1 for channel reduction (from 96 to 1). Finally, the predicted inter-protein contact map is obtained through a sigmoid transformation of the feature map.

2.4 Interaction prediction module

Protein-protein interactions are maintained by physical interactions between residues. Compared to non-interacting protein pairs, the predicted inter-protein contact maps of interacting proteins should exhibit more significant contact signals. Therefore, the inter-protein contact map generated by the previous module is used to integrate the single sequence-base protein language model embeddings (using ESM-1b) to the protein interaction prediction module. Here, we use only the single sequence-based protein language model embeddings and exclude the representations used to generate the inter-protein contact maps, as our in-house tests indicated that directly incorporating these representations significantly increases the false positive rate. One possible reason is that protein-protein interactions exhibit a high degree of sequence specificity. Introducing explicit structural information and multiple sequence alignment-based representations may

reduce the model's sensitivity to sequence features. For example, the model may misidentify protein pairs with similar monomeric folding patterns to known interacting proteins as interacting, even when they do not actually interact.

We modify the Dual Attention Network for multimodal reasoning (r-DAN) proposed by Nam et al.[28] to model PPI prediction as a Visual Question Answering (VQA) task, in which the protein sequence embeddings are used to replace the visual and textual features in VQA, and the predicted inter-protein contact map outputted from previous module is leveraged as a weight matrix to calculate the information transferred between residues of the two proteins. It should be noted that the Dual Attention Networks were also previously used by Li et al. to build the compound-protein binding affinity prediction module in MONN[20]. The architecture details of our interaction prediction module are described as follows:

First, a linear transformation layer is applied to reduce the dimension of the sequence embeddings generated by ESM-1b from 1280 to 448 ($\boldsymbol{P_1} \in \mathbb{R}^{L_1 \times 448}$, $\boldsymbol{P_2} \in \mathbb{R}^{L_2 \times 448}$). Then two feed-forward network (FNN) blocks were successively applied to generate the protein context vectors ($\boldsymbol{p_1}^{(k)} \in \mathbb{R}^{1 \times 448}$, $\boldsymbol{p_1}^{(k)} \in \mathbb{R}^{1 \times 448}$) and update the memory vector ($\boldsymbol{m_p}^{(k)} \in \mathbb{R}^{1 \times 448}$) based on the protein embeddings ($\boldsymbol{P_1}, \boldsymbol{P_2}$) the memory vector ($\boldsymbol{m_p}^{(k-1)}$) in the previous step, and the predicted inter-protein contact map

(**Cmap**) through Equation (4)~(5):

$$p_1^{(k)}, p_2^{(k)} = P\_attention(P_1, P_2, m_p^{(k-1)}, Cmap) \quad (4)$$

$$m_p^{(k)} = Layernorm\left(GRUCell\left(m_p^{(k-1)}, p_1^{(k)} \odot p_2^{(k)}\right)\right) \quad (5)$$

where *GRUCell* is the computational unit that implements the GRU mechanism for processing sequential data in recurrent neural networks, *Layernorm* represents Layer Normalization, and *P_attention* represents Protein Attention, designed to generate protein context vectors by focusing on specific regions in the protein embeddings. Details of this block was described in the following section.

The Protein Attention block: First, two intermediate hidden states ($h_{p_1}^{(k)} \in \mathbb{R}^{L_2 \times 448}$, $h_{p_2}^{(k)} \in \mathbb{R}^{L_1 \times 448}$) correspond to information transferred from each protein to its partner are calculated using the predicted inter-protein contact map ($Cmap \in \mathbb{R}^{L_1 \times L_2}$) as the weight matrix through Equation (6)~(7):

$$h_{p_1}^{(k)} = Layernorm\left(Cmap^T \cdot \left(P_1 W_{p_1}^{(k)}\right)\right) \quad (6)$$

$$h_{p_2}^{(k)} = Layernorm\left(Cmap \cdot \left(P_2 W_{p_2}^{(k)}\right)\right) \quad (7)$$

where $W_{p_1}^{(k)}$, $W_{p_2}^{(k)}$ are the network parameters.

Then two protein attention weights ($\alpha_{p_1}^{(k)} \in \mathbb{R}^{L_1 \times 448}$, $\alpha_{p_2}^{(k)} \in \mathbb{R}^{L_2 \times 448}$) are calculated based on $m_p^{(k-1)}$, P_1, P_2, $h_{p_1}^{(k)}, h_{p_2}^{(k)}$ through Equation (8)~(9):

$$\alpha_{p_1}^{(k)} =$$

$$Sigmoid\left(\left(Layernorm\left(\left(P_1 W^{(k)}_{p_1,tmp}\right) \odot \left(m_p^{(k-1)} W^{(k)}_{p,m}\right)\right) W^{(k)}_{tmp_{p_1}} \odot h^k_{p_2}\right)\right)$$

(8)

$$\alpha^{(k)}_{p_2} =$$

$$Sigmoid\left(\left(Layernorm\left(\left(P_2 W^{(k)}_{p_2,tmp}\right) \odot \left(m_p^{(k-1)} W^{(k)}_{p,m}\right)\right) W^{(k)}_{tmp_{p_2}} \odot h^k_{p_1}\right)\right)$$

(9)

where $W^{(k)}_{tmp_{p_1}}$, $W^{(k)}_{tmp_{p_2}}$, $W^{(k)}_{p_1,tmp}$, $W^{(k)}_{p_2,tmp}$, $W^{(k)}_{p,m}$ are the network parameters, and $Sigmoid$ represents the Sigmoid Activation.

Finally, the protein context vectors are obtained through first elementwise multiplication of the protein embeddings and the attention weights, followed by summation and average over protein sequences and layer normalization through Equation (10)~(11):

$$p_1^{(k)} = Layernorm\left(\overline{P_1 \odot \alpha^{(k)}_{p_1}}\right) \quad (10)$$

$$p_2^{(k)} = Layernorm\left(\overline{P_2 \odot \alpha^{(k)}_{p_1}}\right) \quad (11)$$

Memory vector initialization and protein-protein interaction probability calculation:

The initial memory vector $m_p^{(0)}$ is calculated through Equation (12):

$$m_p^{(0)} = Layernorm(\overline{P_1} \odot \overline{P_2}) \quad (12)$$

The final predicted interaction probability between two proteins is calculated through Equation (13):

$$\textbf{PPI Score} = Sigmoid\left(Layernorm\left(p_1^{(2)} \cdot p_2^{(2)}\right)\right) \quad (13)$$

3. Training of PLMDA-PPI

3.1 The data leakage issue in PPI prediction

Previous studies often randomly split PPI datasets into training set and test set. Before splitting the dataset, protein-protein interactions (PPIs) are clustered to reduce redundancy. Specifically, if both proteins in a PPI exhibit sequence similarities exceeding a predefined threshold (typically 40%) to their respective counterparts in another PPI, the two interactions are considered redundant. In such cases, only one representative PPI is retained to ensure data diversity and minimize redundancy. This approach can mitigate data leakage to some extent but does not completely eliminate it. For example, although protein pairs in the test are not sequence similar to protein pairs in the training, proteins in the test protein pair can still occur in the training. Park et al. proposed three classes of test sets：C1 (both proteins in test protein pairs occur in training)，C2 (only one protein in test protein pairs occur in training) and C3 (no protein overlap between training and test), and found the PPI prediction performances dropped significantly between C1 and C2 as well as C2 and C3[16]. Bernett et al.'s analysis further reveals that the overlap between training and test sets causes models to rely on sequence similarity and node degrees in the network as shortcuts for prediction, leading to severe overfitting[15]. When they minimize sequence similarity between the training and test sets to prevent data leakage, the performance of many PPI prediction models—previously reported as highly effective—drops to nearly random levels. Therefore, careful

handling of data leakage is essential when preparing datasets for model training and evaluation.

3.2 Initial training using structure-informed PPIs from PDB

3.2.1 Datasets for training and validation

We randomly divided the 5,952 sequence clusters, formed by protein monomers in the structure-informed PPI dataset, into two groups: Group 1 with 5,500 clusters and Group 2 with 452 clusters. Positive samples in the training set consisted of 5,352 PPIs formed between proteins within Group I. The validation set included 211 PPIs formed within Group II . Additionally, 626 PPIs involving proteins from both groups were discarded to avoid data leakage. To train the PPI prediction models, we also generated negative samples (non-interacting protein pairs). Specifically, protein monomers within the training and validation sets were randomly paired to create candidate negative samples. We then filtered out any pairs that shared common interacting partners, as these may represent true interactions. For both training and validation sets, we adopted a 10:1 negative-to-positive ratio, following previous studies to reflect the sparsity of true PPIs among protein pairs. Notably, the training and validation sets do not share proteins with sequence similarity higher than 40%, thereby minimizing the risk of overfitting due to data leakage. For the structure-informed PPI datasets, we extracted input features using experimentally determined monomer structures from the corresponding complex

structures.

3.2.2 Training protocol

We employed a loss function formed by a linear combination of the inter-protein contact loss and the PPI loss calculated through Equation (14) to train the model:

$$Loss_{total} = \mu \cdot Loss_{contact} + v \cdot Loss_{interaction} \quad (14)$$

where μ and v are the combination weights, and μ is chosen as 100 and v is chosen as 1.

$Loss_{contact}$ is the inter-protein contact loss calculated through Equation (15):

$$Loss_{contact} = -\frac{1}{L_1 * L_2}\sum_{i=1}^{L_1}\sum_{j=1}^{L_2}\left[\hat{x}_{ij} \log x_{ij} + \left(1 - \hat{x}_{ij}\right) \log\left(1 - x_{ij}\right)\right]$$
(15)

where L_1 and L_2 are the protein lengths, x_{ij} represents the contact probability between residue i and j in the predicted contact map, and $\hat{x}_{ij}$ is the contact label (set as 1 if contact, and 0 if not) in the true contact map. For positive samples (i.e. interacting proteins pairs), the true contact maps are computed based on the heavy-atom-distance between residues in the experimental structures using 8Å as the threshold , and for negative samples (i.e. non-interacting proteins pairs), all residue pairs are considered as non-contact (i.e. $\hat{x}_{ij}$ is always set as 0).

$Loss_{interaction}$ is the protein-protein interaction loss calculated through Equation (16):

$$Loss_{interaction} = -[y \log \hat{y} + (1 - y) \log(1 - \hat{y})] \quad (16)$$

where y represents the interaction label (set as 1 for positive samples, and 0 for negative samples), and $\hat{y}$ represents the predicted interaction probability.

The model training was conducted using the PyTorch (v2.1.0) on a single NVIDIA A100 GPU. Specifically, the initial learning rate was set as 1×10^{-5}, the AdamW optimizer was used to update the model weights, and the batch size was set to 1 . After each epoch, we calculated the total loss on the validation set, and if the total loss on the validation set did not decrease over 3 epochs, the learning rate was reduced by a factor of 0.1 in the next epoch. After two learning rate reductions, we selected the model weights with the lowest loss on the validation set as the final model weights.

3.3 Transfer learning using *Homo sapiens* PPIs from HINT database

3.3.1 Datasets for training and validation

To further enhance the model, the *Homo sapiens* PPI dataset from the HINT database was used for fine-tuning, initialized with the weights obtained from training on the structure-informed PPI dataset. To minimize the possibility of data leakage in dataset splitting, following the same procedure as previously used in the structure-informed PPI dataset, we first clustered protein monomers in the 23,311 PPIs with CD-HIT using CD-

HIT with parameters "-c 0.4 -n 0.2 -g 1", grouping sequences at a 40% identity threshold, which yielded 5873 clusters. The 5873 clusters were then randomly categorized into three groups: group I containing 3800 clusters, group II containing 700 clusters, and group III containing 1373 clusters. We retained only the PPIs formed by proteins within each group and discarded the rest. In addition, to reduce redundancy, for PPIs belonging to the same cluster pair, we randomly selected one representative interaction. As a result, we obtained 4,585 PPIs formed by proteins within group I, 345 PPIs within group II, and 1,502 PPIs within group III. These PPIs were used as the positive samples for the training, validation, and test sets during model fine-tuning. To generate negative samples for each set, we also followed the previous procedure. Specifically, the negative samples were generated through random pairing protein monomers within each group after filtering out protein pairs sharing interacting partners, and a 10:1 negative-to-positive ratio was adopted for each set. For the HINT PPI datasets, where experimental monomer structures are not always available, we employed predicted monomer models from AlphaFoldDB as substitutes to generate the input features.

3.3.2 Training protocol

The training set prepared above was used to further update the model weights initialized with the weights obtained from training on the structure-informed PPI dataset. Since PPIs from HINT database generally do not

have their complex structure information available, we only employed the interaction labels to update the model weights following the $Loss_{interaction}$ defined in Equation (16) as the loss function. To preserve the model's ability to infer inter-protein contact information, the inter-protein contact maps of the structure-informed PPIs were incorporated as constraints to the inter-protein contact prediction module in the training using $\mu \cdot Loss_{contact}$ as the loss function, where $Loss_{contact}$ was the defined in Equation (15), and μ was still set as 100. It is worth noting that during the fine-tuning stage, the structure-informed PPIs were not used to train the interaction prediction module. Still the training was conducted using the PyTorch (v2.1.0) on a single NVIDIA A100 GPU. The initial learning rate was set as 1×10^{-5}, the AdamW optimizer was used to update the model weights, and the batch size was set to 1 . After each epoch, we calculated the total loss using $Loss_{interaction}$ as the loss function on the validation set, and if the total loss on the validation set did not decrease over 3 epochs, the learning rate was reduced by a factor of 0.1 in the next epoch. After two learning rate reductions, we selected the model weights with the lowest loss on the validation set as the final model weights. It should be noted that the inter-protein contact maps were incorporated solely during training to guide the inter-protein contact prediction module, without being involved in model validation or selection.

4. Training of D-SCRIPT, Topsy-Turvy and TT3D

We downloaded the training scripts for D-SCRIPT, Topsy-Turvy, and TT3D from https://dscript.csail.mit.edu/, and trained the models using the default settings provided in the official tutorial (We downloaded the training scripts for D-SCRIPT, Topsy-Turvy, and TT3D from https://dscript.csail.mit.edu/, and trained the models using the default settings provided in the official tutorial (https://d-script.readthedocs.io/en/stable/usage.html#training).

Data Availability

The training and test data for this work can be found in https://github.com/ChengfeiYan/PLMDA-PPI/tree/main/data.

Code Availability

The code for training and implementing PLMDA-PPI can be found in https://github.com/ChengfeiYan/PLMDA-PPI.

Acknowledgements

The work was supported by the National Natural Science Foundation of China (32101001), Major project of Guangzhou National Laboratory (GZNL2023A03007). The computation is completed in the HPC Platform of Huazhong University of Science and Technology.

Author Contributions

C.Y. and S.H. conceived and supervised the project. S.D., X.W., and C.Y. designed and performed the experiments. S.D., M.Z. and C.Y. analyzed the data. D.S., X.W., S.H. and C.Y. wrote the paper. All authors read and

approved the final version of the paper.

Ethics declarations

The authors declare no competing interests.

References

1. Stelzl, U. *et al.* A human protein-protein interaction network: A resource for annotating the proteome. *Cell* **122,** 975-68 (2005)

2. Mintseris, J. & Weng, Z. Structure, function, and evolution of transient and obligate protein-protein interactions. *Proc. Natl. Acad. Sci. U. S. A.* **102**, 10930-5 (2005).

3. Luck, K. *et al.* A reference map of the human binary protein interactome. *Nature* **580**, 402-8 (2020).

4. Huttlin, E. L. *et al.* The BioPlex Network: A Systematic Exploration of the Human Interactome. *Cell* **162**, 425-40 (2015).

5. Drew, K., Wallingford, J. B. & Marcotte, E. M. hu.MAP 2.0: integration of over 15,000 proteomic experiments builds a global compendium of human multiprotein assemblies. *Mol. Syst. Biol.* **17**, e10016 (2021).

6. Burke, D. F. *et al.* Towards a structurally resolved human protein interaction network. *Nat. Struct. Mol. Biol.* **30**, 216–5 (2023).

7. Vidal, M. How much of the human protein interactome remains to be mapped? *Science Signaling* **9**, eg7 (2016).

9. Petrey, D., Zhao, H., Trudeau, S. J., Murray, D. & Honig, B.

PrePPI: A Structure Informed Proteome-wide Database of Protein–Protein Interactions. *J. Mol. Biol.* **435**, 168052 (2023).

10. Zhang, J., Durham, J. & Qian Cong. Revolutionizing protein–protein interaction prediction with deep learning. *Curr. Opin. Struct. Biol.* **85**, 102775 (2024).

11. Humphreys, I. *et al.* Computed structures of core eukaryotic protein complexes. *Science* **374**, eabm4805 (2021).

12. Humphreys, I. R. *et al.* Protein interactions in human pathogens revealed through deep learning. *Nat. Microbiol.* **9**, 2642–2652 (2024).

13. Hu, L., Wang, X., Huang, Y. A., Hu, P. & You, Z. H. A survey on computational models for predicting protein-protein interactions. *Brief. Bioinform.* **22,** bbab036 (2021).

14. Szklarczyk, D. *et al.* The STRING database in 2021: Customizable protein-protein networks, and functional characterization of user-uploaded gene/measurement sets. *Nucleic Acids Res.* **49**, D605-12 (2021).

15. Bernett, J., Blumenthal, D. B. & List, M. Cracking the black box of deep sequence-based protein-protein interaction prediction. *Brief. Bioinform.* **25**, bbae076 (2024).

16. Park, Y. & Marcotte, E. M. Flaws in evaluation schemes for pair-input computational predictions. *Nat. Methods* **9,** 1134-6 (2012).

17. Lannelongue, L. & Inouye, M. Pitfalls of machine learning models for protein–protein interaction networks. *Bioinformatics* **40**, btae012 (2024).

18. Si, Y. & Yan, C. Protein language model-embedded geometric graphs power inter-protein contact prediction. *eLife* **12,** RP92184 (2023).

19. Jumper, J. *et al.* Highly accurate protein structure prediction with AlphaFold. *Nature* **596**, 583-9 (2021).

20. Li, S. *et al.* MONN: A Multi-objective Neural Network for Predicting Compound-Protein Interactions and Affinities. *Cell Syst.* **10**, 308-22(2020).

21. Berman, H. M. *et al.* The Protein Data Bank. *Nucleic Acids Res.* **28**, 235–242 (2000).

22. Das, J. & Yu, H. HINT: High-quality protein interactomes and their applications in understanding human disease. *BMC Syst. Biol.* **6**, 1-12 (2012).

23. Sledzieski, S., Singh, R., Cowen, L. & Berger, B. D-SCRIPT translates genome to phenome with sequence-based, structure-aware, genome-scale predictions of protein-protein interactions. *Cell Syst.* **12**, 969-82 (2021).

24. Singh, R., Devkota, K., Sledzieski, S., Berger, B. & Cowen, L. Topsy-Turvy: Integrating a global view into sequence-based PPI

prediction. *Bioinformatics* **38**, i264-72 (2022).

25. Sledzieski, S., Devkota, K., Singh, R., Cowen, L. & Berger, B. TT3D: Leveraging precomputed protein 3D sequence models to predict protein-protein interactions. *Bioinformatics* **39**, btad663 (2023).

26. Si, Y. & Yan, C. Improved protein contact prediction using dimensional hybrid residual networks and singularity enhanced loss function. *Brief. Bioinform.* **22**, bbad039 (2021).

27. Rives, A. *et al.* Biological structure and function emerge from scaling unsupervised learning to 250 million protein sequences. *Proc. Natl. Acad. Sci. U. S. A.* **118**, e2016239118 (2021).

28. Nam, H., Ha, J. W. & Kim, J. Dual attention networks for multimodal reasoning and matching. in *Proceedings - 30th IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017* (2017).

29. Varadi, M. *et al.* AlphaFold Protein Structure Database: Massively expanding the structural coverage of protein-sequence space with high-accuracy models. *Nucleic Acids Res.* **50**, D439-44 (2022).

30. Van Kempen, M. *et al.* Fast and accurate protein structure search with Foldseek. *Nat. Biotechnol.* **42**, 243-6 (2024).

31. Saito, T. & Rehmsmeier, M. The precision-recall plot is more informative than the ROC plot when evaluating binary classifiers on

imbalanced datasets. *PLoS One* **10**, e0118432 (2015).

32. Evans, R. *et al.* Protein complex prediction with AlphaFold-Multimer. *bioRxiv* (2021).

33. Abramson, J. *et al.* Accurate structure prediction of biomolecular interactions with AlphaFold 3. *Nature* **630**, 493–500 (2024).

34. Li, W. & Godzik, A. Cd-hit: A fast program for clustering and comparing large sets of protein or nucleotide sequences. *Bioinformatics* **22**, 1658-9 (2006).

35. Hsu, C. *et al.* Learning inverse folding from millions of predicted structures. *International conference on machine learning. PMLR 2022* (2022).

36. Johnson, L. S., Eddy, S. R. & Portugaly, E. Hidden Markov model speed heuristic and iterative HMM search procedure. *BMC Bioinformatics* **11**, 1-8 (2010).

37. Suzek, B. E., Huang, H., McGarvey, P., Mazumder, R. & Wu, C. H. UniRef: comprehensive and non-redundant UniProt reference clusters. *Bioinformatics* **23**, 1282–1288 (2007).

38. Rao, R. *et al.* MSA Transformer. *bioRxiv* (2021).

39. Jing, B., Eismann, S., Suriana, P., Townshend, R. J. L. & Dror, R. O. LEARNING FROM PROTEIN STRUCTURE WITH GEOMETRIC VECTOR PERCEPTRONS. in *ICLR 2021 - 9th International Conference on Learning Representations* (2021).

40. Si, Y. & Yan, C. Improved inter-protein contact prediction using dimensional hybrid residual networks and protein language models. *Brief. Bioinform.* **24**, bbad039 (2023).

Figures and Tables

Figure 1. The model architecture and training protocol of PLMDA-PPI. (a) The overall model architecture of PLMDA-PPI. (b) The interaction prediction module of PLMDA-PPI. (c) The workflow for training and evaluation of PLMDA-PPI.

Figure 2. Model validation after initial training on the structure-informed PPI dataset from PDB. (a) The Precision-Recall curves of PLMDA-PPI, D-SCRIPT, Topsy-Turvy and TT3D on the validation set of structure-informed PPIs. (b) Comparison of the distribution of the mean contact probability among the top 50 predicted residue pairs for interacting and non-interacting protein pairs. (c)~(d) An example (PDB 1QTL:A-B) of the experimental inter-protein contact map (c) and the predicted contact map by PLMDA-PPI (d).

Figure 3. Performances of PLMDA-PPI, D-SCRIPT, Topsy-Turvy and TT3D on the HINT datasets of multiple species after initial training on the structure-informed PPI dataset from PDB. (a)~(e) The Precision-Recall curves on *Homo sapiens* (a), *Mus musculus* (b), *Drosophila melanogaster* (c), *Arabidopsis thaliana* (d), *Schizosaccharomyces pombe* (e). (f) Performance variations of PLMDA-PPI (in terms of AUPRC) on the HINT datasets after using different criteria to filter the datasets. Blue bars denote the complete HINT datasets, red bars denote the datasets after removing PPIs with PDB references and green bars denote the datasets after excluding all PPIs involving proteins with >40% sequence identity to any protein in the training set.".

Figure 4. Performances of PLMDA-PPI, D-SCRIPT, Topsy-Turvy and TT3D on the HINT PPI datasets of multiple species after transfer learning on the HINT *Homo sapiens* PPI dataset. (a)~(e) The Precision-Recall curves on *Homo sapiens* (test set) (a), *Mus musculus* (b), *Drosophila melanogaster* (c), *Arabidopsis thaliana* (d), *Schizosaccharomyces pombe* (e). (f) Performance comparison of PLMDA-PPI (in terms of AUPRC) on the HINT datasets before (blue bars) and after transfer learning (red bars). (g)~(j) Performance variations (in terms of AUPRC) of PLMDA-PPI, D-SCRIPT, Topsy-Turvy and TT3D on the HINT PPI datasets of *Mus musculus* (g), *Drosophila melanogaster* (h), *Arabidopsis thaliana* (i), *Schizosaccharomyces pombe* (j) after excluding all PPIs involving proteins with >40% sequence identity to any protein in the *Homo sapiens* training set.

Figure 5. Re-assessment of the models after transfer learning on the HINT *Homo sapiens* PPI dataset on the validation set of the structure-informed PPIs from PDB. Comparisons of the performances (using Precision-Recall curves) before and after transfer learning for PLMDA-PPI (a), D-SCRIPT (b), Topsy-Turvy (c) and TT3D (d).

Figure 6. Ablation studies for PLMDA-PPI. (a) Performance variations (in terms of AUPRC) on the validation set of structure-informed PPIs when key components are removed from the model architecture. (b) Performance comparison (in terms of AUPRC) between models trained with and without inter-protein contact information. (c) Performance comparison (in terms of AUPRC) between models with and without the incorporation of structure-informed PPIs from the PDB for the initial training.

Table 1. Validation of the models trained from structure-informed PPIs from PDB.

Dataset	Model	AUPRC	AUROC	Mean Precision (Top50 Contact Prediction)
PDB dataset	D-SCRIPT	0.324	0.684	0.026
	Topsy-Turvy	0.352	0.756	0.029
	TT3D	0.448	0.826	0.036
	PLMDA-PPI	**0.637**	**0.911**	**0.287**

We show performances of PLMDA-PPI, D-SCRIPT, Topsy-Turvy, TT3D on the validation set of structure-informed PPIs from PDB after initial learning. The best-performing method and metrics are bolded. AUPRC, area under precision-recall curve; AUROC, area under receiver operating characteristic curve. Mean Precision(Top50 Contact Prediction), the mean precisions of the top 50 predicted contacts.

Table 2. Evaluation of the models trained from structure-informed PPIs using PPIs of multiple species from HINT dataset.

Species	Model	AUPRC	AUROC	AUPRC (No PDB ref)	AUROC (No PDB ref)	AUPRC (Seq id<40%)	AUROC (Seq id<40%)
Homo sapiens	D-SCRIPT	0.158	0.605	0.137	0.596	0.151	0.572
	Topsy-Turvy	0.172	0.67	0.149	0.653	0.189	0.661
	TT3D	0.237	0.646	0.223	0.632	0.213	0.635
	PLMDA-PPI	**0.378**	**0.801**	**0.338**	**0.792**	**0.456**	**0.816**
Mus musculus	D-SCRIPT	0.219	0.664	0.204	0.647	0.183	0.663
	Topsy-Turvy	0.206	0.62	0.123	0.588	0.189	0.616
	TT3D	0.237	0.69	0.158	0.684	0.222	0.679
	PLMDA-PPI	**0.485**	**0.804**	**0.346**	**0.777**	**0.488**	**0.844**
Drosophila melanogaster	D-SCRIPT	0.188	0.682	0.209	0.665	0.196	0.721
	Topsy-Turvy	0.196	0.629	0.183	0.642	0.241	0.674
	TT3D	0.249	0.709	0.185	0.676	0.327	0.745
	PLMDA-PPI	**0.476**	**0.851**	**0.389**	**0.812**	**0.533**	**0.889**
Arabidopsis thaliana	D-SCRIPT	0.239	0.697	0.230	0.681	0.205	0.661
	Topsy-Turvy	0.237	0.57	0.209	0.534	0.199	0.567
	TT3D	0.207	0.678	0.176	0.667	0.202	0.707
	PLMDA-PPI	**0.500**	**0.871**	**0.541**	**0.912**	**0.493**	**0.877**
Schizosaccharomyces pombe	D-SCRIPT	0.173	0.664	0.229	0.680	0.183	0.667
	Topsy-Turvy	0.19	0.592	0.171	0.604	0.161	0.603
	TT3D	0.223	0.712	0.222	0.708	0.255	0.744
	PLMDA-PPI	**0.415**	**0.816**	**0.304**	**0.765**	**0.524**	**0.865**

We show performances of PLMDA-PPI, D-SCRIPT, Topsy-Turvy, TT3D on PPIs of multiple species from HINT with initial training on structure-informed PPIs from PDB. The best-performing method for each species and the metrics are bolded. No PDB ref, PPIs in the HINT dataset without PDB reference tags; Seq id<40%, PPIs in the HINT dataset formed by proteins with less than 40% sequence identity to any protein the training dataset.

Table 3. Evaluation of the models after transfer learning on the HINT *Homo sapiens* PPI dataset

Species	Model	AUPRC	AUROC	AUPRC (Seq id < 40%)	AUROC (Seq id < 40%)
Homo sapiens (test set)	D-SCRIPT	-	-	0.245	0.713
	Topsy-Turvy	-	-	0.273	0.696
	TT3d	-	-	0.318	0.784
	PLMDA-PPI	-	-	**0.480**	**0.840**
Mus musculus	D-SCRIPT	0.313	0.719	0.138	0.661
	Topsy-Turvy	0.331	0.732	0.122	0.583
	TT3d	0.323	0.755	0.192	0.656
	PLMDA-PPI	**0.540**	**0.840**	**0.405**	**0.801**
Drosophila melanogaster	D-SCRIPT	0.309	0.67	0.129	0.643
	Topsy-Turvy	0.282	0.671	0.115	0.559
	TT3d	0.211	0.743	0.216	0.665
	PLMDA-PPI	**0.589**	**0.871**	**0.508**	**0.855**
Arabidopsis thaliana	D-SCRIPT	0.249	0.754	0.139	0.637
	Topsy-Turvy	0.263	0.721	0.133	0.578
	TT3d	0.259	0.724	0.149	0.593
	PLMDA-PPI	**0.601**	**0.898**	**0.585**	**0.899**
Schizosaccharomyces pombe	D-SCRIPT	0.234	0.682	0.114	0.570
	Topsy-Turvy	0.242	0.664	0.110	0.537
	TT3d	0.278	0.744	0.180	0.625
	PLMDA-PPI	**0.454**	**0.830**	**0.338**	**0.826**

We show performances of PLMDA-PPI, D-SCRIPT, Topsy-Turvy, TT3D on PPIs of multiple species from HINT after transfer learning on the HINT *Homo sapiens* PPI dataset. The best-performing method for each species and the metrics are bolded. Seq id<40%, PPIs in the HINT dataset formed by proteins with less than 40% sequence identity to any protein the training dataset.

Table 4. Reassessment of the models after transfer learning on the validation set of structure-informed PPIs.

Dataset	Model	AUPRC	AUROC	Mean Precision (Top50 Contact Prediction)
PDB dataset	D-SCRIPT	0.258	0.719	0.026
	Topsy-Turvy	0.271	0.697	0.029
	TT3d	0.296	0.757	0.036
	PLMDA-PPI	**0.642**	**0.907**	**0.230**

We re-evaluated the performance of PLMDA-PPI, D-SCRIPT, Topsy-Turvy, and TT3D after transfer learning on the HINT *Homo sapiens* PPI dataset, using the validation set of structure-informed PPIs derived from PDB. The best-performing method and metrics are bolded.

Supplementary

Figure S1. Distributions of PPI scores predicted by PLMDA-PPI, D-SCRIPT, Topsy-Turvy, and TT3D after initial training on structure-informed PPIs, for positive and negative samples in the validation set derived from structure-informed PPIs in PDB and the HINT datasets across multiple species.

Figure S2. Distributions of PPI scores predicted by PLMDA-PPI, D-SCRIPT, Topsy-Turvy, and TT3D after transfer learning, evaluated on positive and negative samples from the validation set of structure-informed PPIs derived from PDB and HINT datasets across multiple species.

Figure S3. Comparison of the mean contact probability distributions of the top 50 predicted residue pairs by PLMDA-PPI after transfer learning, for positive and negative samples in the validation set of structure-informed PPIs from PDB and HINT datasets across multiple species

Table S1. Performance comparison of PLMDA-PPI before and after transfer learning on HINT datasets across multiple species, following the exclusion of PPIs involving any protein sharing more than 40% sequence identity with proteins in the *Homo sapiens* training set.

Species	AUPRC (PDB)	AUROC (PDB)	AUPRC (Transfer)	AUROC (Transfer)
Homo sapiens	0.387	0.795	**0.480**	**0.840**
Mus musculus	0.389	0.789	**0.405**	**0.801**
Drosophila melanogaster	0.429	0.838	**0.508**	**0.855**
Arabidopsis thaliana	0.490	0.883	**0.585**	**0.899**
Schizosaccharomyces pombe	**0.342**	0.794	0.338	**0.826**

Table S2. Performances of the models when trained directly on the HINT *Homo sapiens* PPI training set

Species	Model	AUPRC	AUROC
Homo sapiens	D-SCRIPT	0.197	0.74
	Topsy-Turvy	0.212	0.709
	TT3d	0.227	0.706
	PLMDA-PPI	0.400	0.812
Mus musculus	D-SCRIPT	0.169	0.713
	Topsy-Turvy	0.16	0.63
	TT3d	0.196	0.688
	PLMDA-PPI	0.393	0.838
Arabidopsis thaliana	D-SCRIPT	0.134	0.635
	Topsy-Turvy	0.127	0.576
	TT3d	0.153	0.593
	PLMDA-PPI	0.313	0.829
Drosophila melanogaster	D-SCRIPT	0.123	0.615
	Topsy-Turvy	0.141	0.563
	TT3d	0.239	0.685
	PLMDA-PPI	0.378	0.828
Schizosaccharomyces pombe	D-SCRIPT	0.119	0.597
	Topsy-Turvy	0.141	0.573
	TT3d	0.204	0.66
	PLMDA-PPI	0.377	0.77