Exploring the Protein-Protein Interactions by Exploiting Deep Learning Techniques

A thesis submitted to the
Indian Institute of Technology Kharagpur
for the award of the degree

of

Doctor of Philosophy

by

Shubhrangshu Ghosh

Under the guidance of

Prof. Pralay Mitra

Department of Computer Science and Engineering
Indian Institute of Technology Kharagpur
Kharagpur - 721 302, India

April 2026

Abstract

Proteins are amino-acid polymers that power essential cellular functions—including catalysis, signaling, transport, immunity, and regulation—and are therefore central to understanding health and disease. As a result, computational biology has increasingly focused on predicting protein–protein interactions (PPIs); however, predicting and interpreting these interactions—and designing new ones—remains difficult because the space of possible protein pairs is vast, while experimentally validated PPIs and resolved complex structures cover only a small fraction of the interactome. Classical structure-based docking and physics-driven design provide insight, but are often limited by structural availability and computational cost. In contrast, sequence-based approaches promise proteome-scale coverage, but must overcome persistent obstacles: (i) robust generalization beyond training distribution, (ii) interpretability that links sequence-level predictions to putative binding mechanisms, and (iii) ability to move from prediction to actionable design. Motivated by recent advances in deep learning, protein language models (PLMs), explainable AI (XAI), and GPU-accelerated computing, this thesis explores protein-protein interactions through a unified, sequence-first lens, and develops a coherent pipeline that progresses from predicting interactions, to explaining them, and finally to engineering new interactions *in silico*.

The thesis explores protein–protein interactions from various perspectives using different Deepl Learning (DL) techniques. First, for a given pair of protein sequences, we ask whether PPIs can be predicted reliably at a quality suitable for interactome screening. Then, we extend the task beyond a binary decision and ask whether a sequence-based model can propose an interpretable residue–residue interaction hypothesis that approximates a protein-protein interaction map, bridging the gap between sequence-only inference and structural intuition. Next, starting from a known protein complex, we try to explore other possible altered/mutated protein sequence that has the same protein interaction. Across all these explorations, a central theme is that modern representation learning can extract latent interaction signals from sequences, while designed learning and inference mechanisms can translate these signals into models that are accurate, interpretable, and usable in design workflows.

First, we introduce *MaTPIP*, a feature-mixed deep learning architecture for sequence-driven PPI prediction with explainability analysis. It combines pre-trained PLM embeddings with manually curated sequence descriptors, representing information at both residue ("part") and whole-protein levels. This architecture integrates convolutional components that capture local amino-acid motifs with Transformer-style sequence-context components, enabling model to learn complementary local and global interaction cues. We incorporate XAI techniques to quantify contribution of feature families and to probe model behavior across evaluation settings. On a human PPI benchmark, it improves area under the ROC curve (AUC) from 74.1% to 78.6% and average precision (AP) from 23.2% to 32.8%, doubling precision at 3% recall (4.9%

→ 9.5%). In cross-species prediction, it raises area under the precision–recall curve (AUPR) from 60.9% to 81.1% for mouse and from 56–58% to ≈78–81% for worm, fly, yeast, and *E. coli*, indicating reliability suitable for large-scale interactome prioritization.

Building on this foundation, we move beyond binary PPI prediction to estimate coarse-grained interaction maps using *PEPpip*, a vision-inspired deep learning framework that couples binary PPI prediction with residue-level interaction mapping via XAI. It encodes a pair of protein sequences into an image-like, three-dimensional feature representation using multiple pre-trained PLMs, enabling the use of computer vision architectures on sequence-derived tensors. Two classifiers—a ResNet-based convolutional model and a Vision Transformer—are trained to learn PPI patterns from these feature maps, and their probabilistic outputs are fused through a *Post-hoc Inference Combiner* to leverage complementary strengths. Beyond classification, this framework derives coarse interaction maps as a secondary output by integrating attribution-based signals from integrated gradients with attention-derived signals from the transformer, followed by a modular noise-reduction pipeline; a residue–residue interaction-aware, class-discriminative attention mechanism that injects biologically meaningful priors into attention computation, further improves contact relevance. Empirically, it establishes new cross-species AUPR benchmarks of 91.1% (mouse), 89.1% (fly), 87.7% (worm), 68.1% (yeast), and 71.9% (*E. coli*)—about 2–3 points above prior best results—while remaining competitive on human PPI tasks. For interaction-map estimation, although maps are not intended to replace dedicated contact predictors or structure-based modeling, It improves contact precision over prior sequence-based classifiers, providing a practical hypothesis-generation tool in settings where structural, evolutionary, or co-evolutionary annotations are unavailable.

Finally, we extend our understanding on PPI to protein interaction design. We present *PinMoM*, an AI-guided pipeline that integrates stochastic search with learned sequence-based scoring to design protein interactions. It reframes interface design as a Monte Carlo optimization in sequence space, where candidate binders are mutated and accepted or rejected using a Metropolis criterion guided by a PPI stability score derived from MaTPIP. This substitutes physics-based energy evaluation with a data-driven scoring function while retaining an interpretable acceptance criterion characteristic of classical sampling. To scale exploration, this pipeline supports replica-exchange variants of Monte Carlo and employs clustering-based candidate selection. Designed sequences are then screened and validated using AlphaFold-based complex modeling and molecular dynamics (MD) simulation. To reduce expense of verification and prioritize the promising designs under limited labels, we incorporate positive–unlabeled (PU) learning via Kolmogorov–Arnold Networks (KAN), improving design selection (F_1 +5%, precision +11%, AUC +2% over a multi-layer perceptron (MLP) baseline). Across benchmark complexes, AlphaFold2 predictions for designed binders achieve a median interface root-mean-square deviation (RMSD) of 1.17 Å, a predicted template modelling score (pTM) of 0.72, and an interface pTM (ipTM) of 0.88, demonstrating near-native quality and compares favorably to traditional evolutionary and recent simulation-driven baselines while remaining competitive with modern data-driven design approaches.

Overall, in this thesis, we establish that combining PLM-based sequence embeddings with

hand-crafted descriptors yields better sequence-based PPI prediction. Next we show that accuracy can be associated with interpretability, producing residue-level interaction hypotheses that help translate sequence-only decisions into mechanistic insight. Then the proposed pipeline closes the loop by using learned interaction scoring to guide stochastic sequence design and by integrating filtering and verification to make interface design tractable at scale. Beyond individual models, this thesis contributes a broader methodology for sequence-based protein engineering: use foundation-model representations to capture latent biological signals, use interpretability to expose contact-relevant structure in those signals, and use fast learned scoring to enable iterative design. Looking forward, several directions can advance sequence-based docking and design. Promising avenues include unified sequence–structure models with outputs, stronger and benchmarked interpretability, and design loops that combine generative modeling with physical validation and experimental feedback to close the loop between prediction, explanation, and real-world protein engineering.

Keywords: Protein-protein interaction prediction; Protein interaction design; Protein Language Model; Transfer Learning; Deep Learning Architecture; eXplainable AI; Vision Transformer; Monte Carlo simulation; AlphaFold2; PU learning;