

Assessment of scoring functions for computational models of protein-protein interfaces

Jacob Sumner^{1,2}, Grace Meng^{2,3}, Naomi Brandt^{2,4}, Alex T. Grigas^{1,2}, Andrés Córdoba^{5,6}, Lynne Regan⁷, Mark D. Shattuck⁸, and Corey S. O’Hern^{1,2,4,9,10}

¹*Graduate Program in Computational Biology and Bioinformatics, Yale University, New Haven, Connecticut, 06520, USA*

²*Integrated Graduate Program in Physical and Engineering Biology, Yale University, New Haven, Connecticut, 06520, USA*

³*Department of Chemistry, Yale University, New Haven, Connecticut 06520, USA*

⁴*Department of Physics, Yale University, New Haven, Connecticut 06520, USA*

⁵*Department of Chemistry, Duke University, Durham, North Carolina 27708, USA*

⁶*Department of Physics, Duke University, Durham, North Carolina 27708, USA*

⁷*Institute of Quantitative Biology, Biochemistry and Biotechnology, Centre for Synthetic and Systems Biology, School of Biological Sciences, University of Edinburgh*

⁸*Benjamin Levich Institute and Physics Department, City College of New York, New York, New York 10031, USA*

⁹*Department of Mechanical Engineering and Materials Science, Yale University, New Haven, Connecticut 06520, USA*

¹⁰*Department of Applied Physics, Yale University, New Haven, Connecticut 06520, USA*

Manuscript Pages: 21

Total Manuscript Figures: 7

Total Manuscript Tables: 0

Supporting Information Pages: 11

Total Supporting Information Figures/Tables: 7/1

Abstract: A goal of computational studies of protein-protein interfaces (PPIs) is to predict the binding site between two monomers that form a heterodimer. The simplest version of this problem is to rigidly re-dock the bound forms of the monomers, which involves generating computational models of the heterodimer and then scoring them to determine the most native-like models. Scoring functions have been assessed previously using rank- and classification-based metrics, however, these methods are sensitive to the number and quality of models in the scoring function training set. We assess the accuracy of seven PPI scoring functions by comparing their scores to a measure of structural similarity to the x-ray crystal structure (i.e. the DockQ score) for a non-redundant set of heterodimers from the Protein Data Bank. For each heterodimer, we generate re-docked models uniformly sampled over DockQ and calculate the Spearman correlation between the PPI scores and DockQ. For some targets, the scores and DockQ are highly correlated; however, for many targets, there are weak correlations. Several physical features can explain the difference between difficult- and easy-to-score targets. For example, strong correlations exist between the score and DockQ for targets with highly intertwined monomers and many interface contacts. We also develop a new score based on only three physical features that matches or exceeds the performance of current PPI scoring functions. These results emphasize that PPI prediction can be improved by focusing on correlations between the PPI score and DockQ and incorporating more discriminating physical features into PPI scoring functions.

Significance: We address the challenge of predicting binding interfaces for heterodimers. We evaluate PPI scoring functions by comparing their scores to a measure of similarity between the computational model and the corresponding x-ray crystal structure. For most targets, there are weak correlations between the score and ground truth, emphasizing the need for improved scoring functions. We identify physical features that can be combined into a new scoring function to predict the binding interface quality of computational models.

Availability: https://github.com/jakesummeryale/Protein_interface_assessment

Keywords: protein-protein interfaces | rigid-body and flexible docking | protein complex design

1 Introduction

Proteins interact and form interfaces in numerous important biological processes, such as enzymatic catalysis of reactions, cytoskeletal organization, and immune recognition of pathogens. Identifying the binding interface between two proteins is crucial for understanding the function of protein complexes. The structure of protein-protein interfaces (PPIs) can be resolved via x-ray crystallography, nuclear magnetic resonance spectroscopy, and cryo-electron microscopy. However, experimentally resolving PPI structures is expensive and time-consuming. For example, there are more than 10^5 PPIs in the human body, yet only $\sim 5\%$ of their structures have been experimentally resolved¹⁻⁴. An important goal in the field of computational protein design is to develop the ability to predict whether two proteins form an interface and identify the location of the interface⁵⁻¹⁴.

The problem of determining PPIs involves two steps: one must first know whether the two monomers interact and then determine the binding interface. In this work, we focus on the second step. In particular, we ask, given two monomers that are known to interact, can a near native structure be identified among a large pool of computational models of the bound structure? In general, this task involves flexible docking, i.e. predicting both the conformational changes in each monomer upon binding and the interface between the two proteins. In flexible docking, there is a large conformational space to search that includes changes in the conformation of the protein backbone and side chains to find a near-native interface structure. Therefore, flexible docking depends on two main factors: sufficient sampling of the conformation space of the docked proteins and accurately scoring these conformations. Several groups have created flexible docking programs^{5,6,9,15} and participate in the Critical Assessment of PRedicted Interactions competition (CAPRI), which asks groups to correctly predict the structures of novel experimentally determined PPIs from the amino acid sequence^{12,14}. While the performance of flexible docking software has improved over the past 20 years, it remains difficult to accurately determine the structure of heterodimers, which is a much simpler case than the large multimeric complexes that are frequently considered in CAPRI¹⁴. Even deep-learning PPI structure prediction methods, such as AlphaFold-Multimer^{13,16}, are unable to place $\sim 30\%$ of heterodimeric PPIs within 10 Å of their native structure¹⁷. In contrast, duplicate PPIs that have been crystallized by multiple groups possess root-mean-square (RMS) deviations in the C_α atoms that are less than 1 Å. Thus, the problem of flexible docking is not yet solved.

The flexible docking problem can be considered as a sequence of rigid-body docking calculations as the unbound conformations of the monomers are changed into the bound monomer conformations. In rigid-body docking, model generation is straightforward since one can simply generate all possible rigid-body translations and rotations of one monomer relative to the other. Thus, accurate scoring of PPI models is the key component of rigid-body docking. Rigid-body *re-docking*, or recapitulating the bound interface from the separated bound monomers of a heterodimer whose structure has been experimentally determined, removes the question of whether each monomer is in the correct conformation for binding. Rigid-body re-docking is typically considered a “solved” problem, with hit rates for *acceptable* computational models exceeding 50% in the top five scored models¹⁸ using a common docking benchmark dataset with ~ 200 PPIs^{19,20}. Given that there are $\gtrsim 10^5$ PPIs, this sample set is likely inadequate to capture the full heterogeneity of PPIs. In fact, a recent study²¹ has found that scoring functions for rigid-body docking software can identify *high-quality* models in the top 10 scored models only 3% of the time. This result suggests that rigid-body re-docking is not a solved problem and one must identify the shortcomings of current scoring functions and then correct them.

Simplifying the problem of PPI prediction to rigid-body re-docking allows us to assess the accuracy of scoring functions. Most computational studies of PPIs^{5,18,22-35} consider a scoring function as “successful” if any computational model in the top N ranked models (where N typically

ranges from 5 to 100) is above an arbitrary quality cutoff^{10,12,14,36–38}. However, such measurements of “hit rate” are sensitive to how the computational models are sampled. Another disadvantage of using hit rate is that it evaluates scoring performance over an entire dataset, instead of evaluating the scoring performance for each target within the dataset. In contrast, the receiver operator characteristic (ROC) curve can determine the classification performance for a single target by measuring the area under the ROC curve (AUC). For the ROC curve, all scored models are assigned a label of “true positive” or “true negative” according to an arbitrary quality cutoff. However, AUC is also sensitive to how the computational models are sampled. For random sampling of PPI models generated from a high-resolution x-ray crystal structure (or target), there are far fewer near-native models than low-quality models. (See Supplementary Material.) For example, we obtain an AUC of 0.959 when using a standard quality cutoff when randomly sampling computational models for a particular target³⁶. A more direct metric for assessing scoring accuracy is to correlate the predicted score to a ground truth measure of quality for each computational model (e.g. the DockQ score³⁹). For the same data, the Spearman correlation coefficient $|\rho| \approx 0.256$, which is low compared to the classification accuracy reported by the AUC. The contradictory prediction accuracies reported for ρ and AUC are caused by an imbalance of model quality in the data set, which is typical of docking software. Disparities in ρ and AUC can cause unreliable results when comparing the performance of existing PPI scoring functions. However, we have developed techniques to sample computational models uniformly over DockQ, which removes the disparities found in AUC and ρ and enables reliable comparisons of PPI scoring performance among different targets.

Here, we assess the current state-of-the-art PPI scoring functions and determine when they are able to discern near-native from low-quality models. We focus on a dataset of 84 high-resolution PPI heterodimer targets with no bound ligands obtained from the RCSB PDB⁴⁰. For each target, we generate 540,000 computational models using a template-free rigid-body re-docking protocol, ZDOCK3.0.2^{41,42}, and sample a uniform distribution of model quality based on their structural similarity to the target. We then use several scoring functions, such as ZRank2^{24,25}, ITScorePP²⁷, Rosetta^{9,28,29}, PyDOCK²⁶, VoromQA³⁰, Deeprank-GNN-ESM¹⁸, and GNN-DOVE³⁴, to score the models for each target. We assess the performance of these scoring functions by correlating each score with DockQ and measuring the AUC as a function of the DockQ cutoff. For each target, we find that the AUC and the Spearman correlation coefficient ρ obey $\text{AUC} \approx -0.5\rho + 0.5$ when the model quality is uniformly sampled. ($\rho \approx -1$ for accurate scoring functions because near-native models are typically associated with lower scores.) For half of the targets, the average magnitude of the Spearman correlation coefficient is less than 0.70, which shows that for a majority of the targets it is difficult to distinguish poor-quality, rigid-body docked models from native structures. Thus, the scoring of rigid-body re-docked PPI models is still an open problem. We also identify several important physical features for predicting model quality: the number of interfacial contacts, interface size, and degree to which the two monomers are intertwined at the interface. After combining these physical features into a scoring function using support vector regression, we show that its performance is comparable to the current best-performing scoring functions. In the future, we will develop a deep-learning PPI scoring function trained on our large dataset of models with balanced quality to improve the scoring accuracy of all PPI targets.

2 Results

Our goal is to assess the accuracy of current scoring functions for computational models of rigidly re-docked PPIs generated using the individual bound monomers of high-resolution x-ray crystal structure heterodimers. The most commonly used method to assess PPI scoring function accuracy

is the hit rate plot^{30,32–35}, which measures the fraction of targets with at least one “positive” (i.e. models with DockQ ≥ 0.23) in the top i ranked models (out of a total of N_t models) scored by a given scoring function averaged over all targets. (Models with DockQ < 0.23 are referred to as “negatives.”) However, the hit rate plot is extremely sensitive to N_t when there are weak correlations between the ground truth score (i.e. DockQ) and scores from PPI scoring functions. In Fig. 1 (a) and (c), we plot the probability distribution $\log_{10} P$ for computational models to have given DockQ and Rosetta scores for a large dataset of models generated by ZDOCK for two example targets, 3RCZ⁴³ and 3YGS⁴⁴, respectively. The magnitude of the Spearman correlation coefficient $|\rho|$ for 3YGS ($|\rho| = 0.35$) is much smaller than that for 3RCZ ($|\rho| = 0.62$).

To illustrate the sensitivity of the hit rate metric on the number of models that are scored, we calculate the *effective* hit rate fraction for each target. We first generate $N_t = 108,000$ computational models using ZDOCK for each target. We then randomly draw N_m models from this collection N_s times. The N_m models are ranked using the Rosetta scoring function from lowest to highest, $i = 1, \dots, N_m$. The effective hit rate fraction at i is the fraction of the N_s draws that have at least one positive model for ranks less than or equal to i . For target 3RCZ, we show in Fig. 1 (b) that the effective hit rate fraction for 3RCZ improves as N_m increases, which suggests that a larger sample size can improve the hit rate fraction. However, in Fig. 1 (d), the effective hit rate fraction decreases as N_m increases for 3YGS. In particular, as N_m approaches N_t , the effective hit rate fraction for 3RCZ approaches 1, while for 3YGS, it approaches 0. These results emphasize that the effective hit rate fraction is sensitive to the number of models N_m . Similar results for the hit rate fraction are found for all other scoring functions that we tested.

Why does the effective hit rate fraction tend to limiting values of 0 or 1 as $N_m \rightarrow \infty$? An important difference between the targets 3RCZ and 3YGS is that Rosetta assigns lower scores to negative models than those for positive models. Thus, it is important to quantify the fraction of computational models μ that are negative and have a PPI score (e.g. Rosetta score) that is lower than the minimum value for all positive models. To understand the dependence of the effective hit rate fraction on μ , we generated a large synthetic dataset of “docked models” for a PPI target with Pearson correlation $|\rho| \approx 0.1$ between “DockQ” and “PPI score”. We balance the synthetic dataset so that it has an equal number of positive and negative models, assuming again that the positive models have “DockQ” ≥ 0.23 . In Fig. S2, we show that the effective hit rate fraction tends toward zero for all $\mu > 0$, which emphasizes the unreliability of the hit rate fraction when considering the large number of samples limit.

We employ a similar method to assess the effective hit rate fraction for each target in our PPI dataset. Specifically, we take $N_s = 10^3$ samples for varying sample sizes N_m for each target with $N_t = 10^4$ ZDOCK-generated models. In Fig. 2 (a), we plot the effective hit rate fraction at $i = 100$ over N_s samples versus N_m/N_t ranking the models using ZRank2. The effective hit rate fractions for 81/84 targets approach one as N_m/N_t increases. For three targets, the effective hit rate fractions tend toward zero. (We find similar results for our PPI dataset for all other scoring functions that we tested.) These results highlight the problems associated with using hit rate fraction as a PPI scoring evaluation metric. First, the hit rate fraction depends on the value of N_m/N_t that is used. Second, even though the effective hit rate fraction tends toward 1 in the large- N_m limit for the majority of targets, we are unable to compare the relative performance of PPI scoring functions for different targets using the hit rate fraction.

Two other metrics that can be used to assess PPI scoring function performance are the area under the ROC curve (AUC) and Spearman correlation ρ . $0.5 < \text{AUC} < 1$ indicates how well a scoring function classifies models as positive versus negative. AUC values of 0.5 and 1 indicate random and perfect classification, respectively. ρ is a non-linear measure of the correlation between DockQ and the PPI score; $\rho = -1$ indicates a perfect negative correlation between DockQ and the

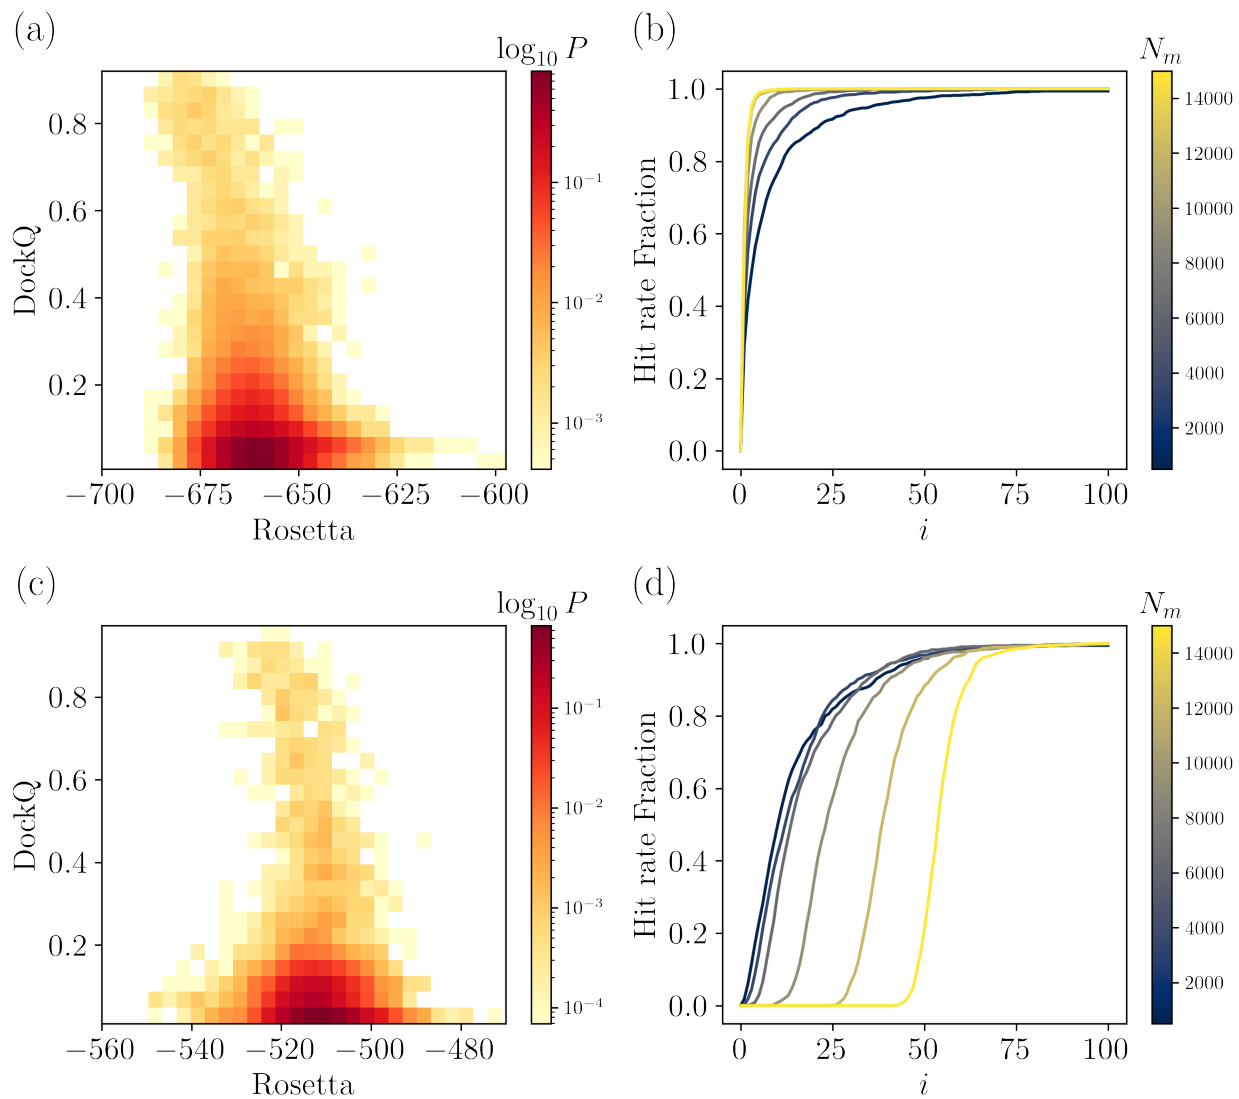


Figure 1: The probability distribution $P(\text{DockQ}, \text{Rosetta})$ for the computational models to have given DockQ and Rosetta scores when exhaustively sampling the computational models for heterodimer targets (a) 3RCZ and (c) 3YGS. The color scale from light yellow to dark red indicates increasing probability. The effective hit rate fraction is plotted versus the index i of the top-ranking models for increasing numbers of computational models N_m (from blue to yellow) for (b) 3RCZ and (d) 3YGS.

PPI score. In Fig. 2 (b), we show a scatterplot of DockQ versus ZRank2 scores for all models of PDB 3U82 with $\rho = -0.191$. The ROC curve calculated for the same target (with $\text{AUC} = 0.707$) is shown in Fig. 2 (c). We compare ρ and AUC for all 84 targets in Fig. 2 (d). These results show inconsistency between the classification and correlation metrics. For many targets, the $\text{AUC} \sim 1$, while $|\rho| \leq 0.5$. For example, GNN-DOVE inconsistently scores high DockQ models, either assigning them the highest or lowest scores. Thus, we omit GNN-DOVE from further analyses. For the rest of the scoring functions, why is the magnitude of the Spearman correlation coefficient so low, yet the AUC indicates accurate classification?

The disparity between ρ and AUC likely stems from the overabundance of poor-quality models and lack of high-quality models typically generated by PPI docking methods. To address both issues, we generated a dataset of ~ 50 models (sampled from a larger pool of 540,000 models per target) for each DockQ bin with width 0.05 (Fig. 3 (a)). After energy minimizing the models, the distribution of model quality is nearly uniform except for the highest few bins of DockQ. We then re-balanced the dataset so that half of the models are positive and half are negative by adding negative models to the energy minimized dataset. This uniform sampling procedure gives ~ 1400 models per target. We then calculate ρ and the AUC for the uniformly sampled dataset and find that the AUC obeys $\text{AUC} \approx -0.5\rho + 0.5$ as shown in Fig. 3 (b). Since ρ and AUC are linearly related for the uniformly sampled data, we focus on ρ for the remaining analyses.

To quantify the change in the evaluation metrics of the scoring functions after uniform sampling, we compared ρ and AUC obtained from exhaustive versus uniform sampling. We define $\Delta\rho = \rho_u - \rho_e$ and $\Delta\text{AUC} = \text{AUC}_u - \text{AUC}_e$, where “u” and “e” represent the uniform and exhaustively sampled datasets. In Fig. 3 (c), we show a scatterplot of $\Delta\rho$ versus ΔAUC . We find that most of the data points are found in the second quadrant, which indicates improvements in both correlation and classification metrics following uniform sampling. (Increased correlations correspond to $\Delta\rho < 0$ and improved classification corresponds to $\Delta\text{AUC} > 0$.) Some targets have a negative ΔAUC or a positive $\Delta\rho$, which likely stems from under-sampling of positive models for targets in the exhaustively sampled dataset. These results emphasize that after uniformly sampling models in DockQ, ρ is a consistent metric for evaluating the performance of PPI scoring functions.

In Fig. 4, we show the results for the Spearman correlation ρ between DockQ and six scoring functions (ZRank2, ITScorePP, Rosetta, PyDock, VoromQA, and DeepRank-GNN-ESM) for all 84 PPI targets. The targets are ordered according to increasing $\langle\rho\rangle$, averaged over the six scoring functions. (Note that since $\rho < 0$ increasing ρ implies decreasing magnitude of the correlations.) Some targets have $\rho \approx -1$, e.g. 3WHQ, 1UGQ, and 1N1J. Approximately 20 targets are “easy” with $|\langle\rho\rangle| > 0.8$. Twenty-three targets are “hard” with $|\langle\rho\rangle| < 0.6$. The hard targets have greater variability in ρ across the six scoring functions. The remaining 41 targets with $0.8 \leq |\langle\rho\rangle| \leq 0.6$ are medium difficulty. The average $\langle\rho\rangle_t$ over targets for each scoring function is shown in Fig. S3. ZRank2 has the largest $|\langle\rho\rangle_t| \approx 0.78$. In contrast, VoromQA has the smallest $|\langle\rho\rangle_t| \approx 0.56$ (excluding GNN-DOVE) and performed well only on the easy targets. Current scoring functions give accurate predictions ($|\rho| > 0.8$) of model quality for only $\sim 25\%$ of PPI targets. A highly accurate scoring function would generate correlations $|\rho| \geq 0.9$ between the score and DockQ for all re-docked targets. Thus, a large portion of targets are still difficult to score, with scoring functions unable to distinguish between low- and high-quality models.

What makes a PPI target easy or hard to score? Rigid-body re-docking involves generating models in a six-dimensional configuration space. If the receptor’s center of mass is fixed at the origin, the ligand’s center of mass can be placed at any position (x, y, z) (three degrees of freedom) and three angles are required to specify its orientation. We can calculate the PPI score and DockQ score at any point in this six-dimensional space, and a large $|\rho|$ indicates strong correlations between the PPI score and DockQ in this six-dimensional configuration space. However, it is difficult to

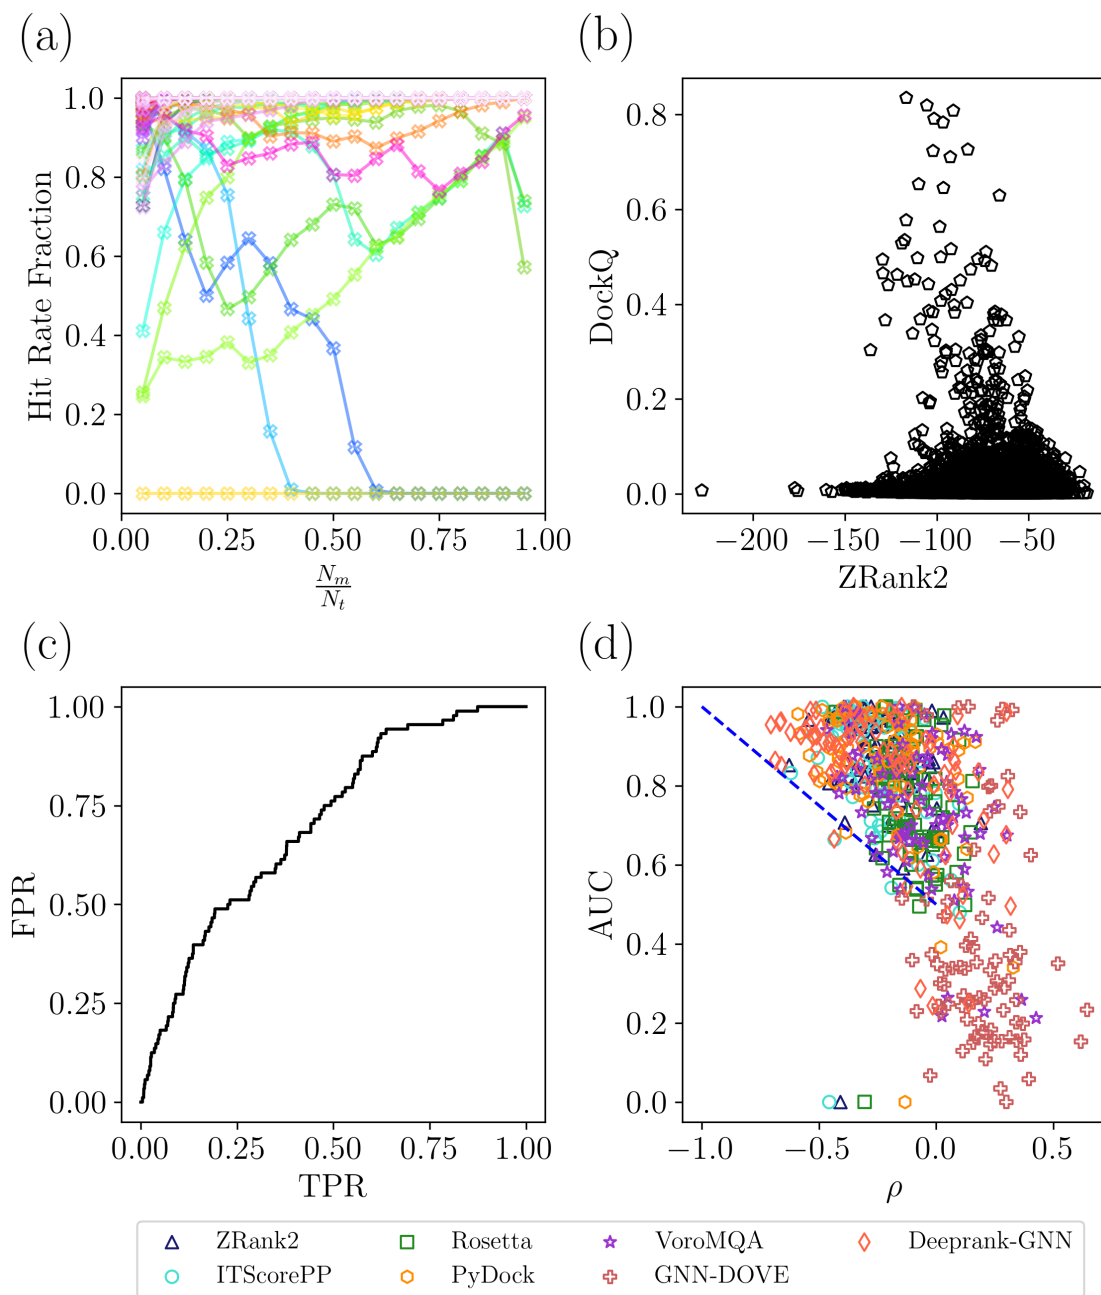


Figure 2: Metrics to evaluate the performance of PPI scoring of re-docked models generated without uniform sampling of DockQ. (a) The effective hit rate fraction at rank $i = 100$ generated for all targets obtained by randomly sampling models for each target plotted as a function of the ratio of the number of models in each sample to the total number of models, N_m/N_t . (b) A scatterplot of ZRank2 versus DockQ for PDB 3U82 with Spearman correlation $\rho = -0.191$. (c) ROC curve for the data in (b) with true positive cutoff $\text{DockQ} \geq 0.23$ and $\text{AUC} = 0.707$. (d) The Spearman correlation ρ between DockQ and each of the seven scoring functions plotted versus the AUC for all targets. Each color represents a different target and the blue dotted line represents $\text{AUC} = -0.5\rho + 0.5$.




Figure 3: Improvements in the evaluation of PPI scoring functions due to uniform sampling of DockQ. (a) Probability distribution $P(\text{DockQ})$ of the DockQ for computational models for all targets obtained by exhaustive sampling of models (gray circles), uniformly sampling DockQ before energy minimization (blue triangles), and uniformly sampling DockQ after energy minimization (black crosses). (b) AUC plotted versus ρ for models that uniformly sample DockQ (after energy minimization) for each scoring function. The blue dotted line gives $\text{AUC} = -0.5\rho + 0.5$. (c) $\Delta\rho = \rho_u - \rho_e$ plotted versus $\Delta\text{AUC} = \text{AUC}_u - \text{AUC}_e$, where ρ_u (ρ_e) is the Spearman correlation for the models that were uniformly (exhaustively) sampled in DockQ and AUC_u (AUC_e) is the area under the ROC curve for the models that were uniformly (exhaustively) sampled.

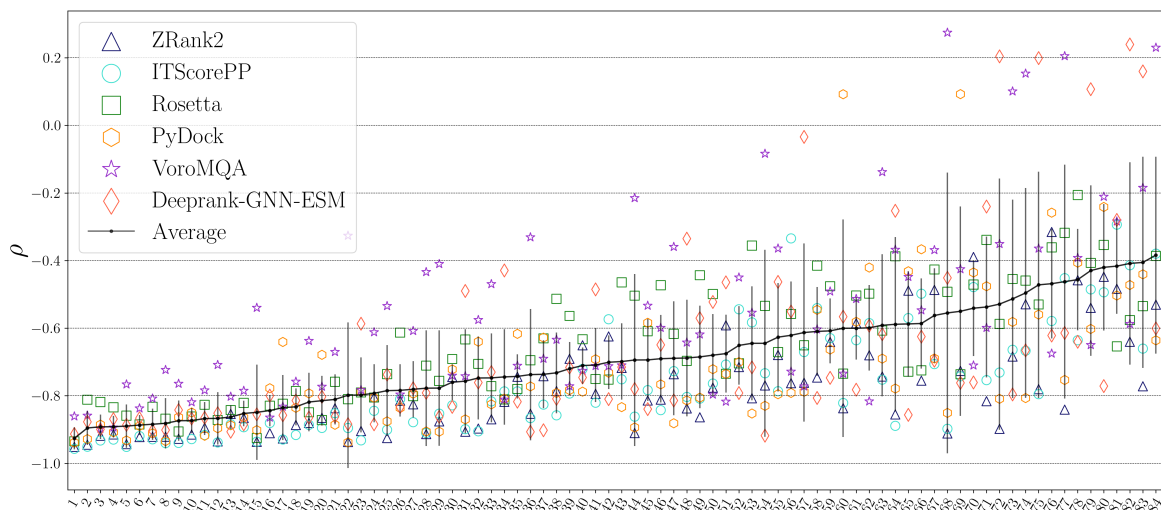


Figure 4: The Spearman correlation ρ between DockQ and PPI score for each target and scoring function. We order the targets by increasing $\langle\rho\rangle$ (black line), averaged over the scoring functions, from left to right. (Note that since $\rho < 0$ increasing ρ implies decreasing the magnitude of the correlations.) The standard deviations of the scores for each target are indicated by the black vertical lines. The numbers on the horizontal axis correspond to PPI targets listed in Table S1.

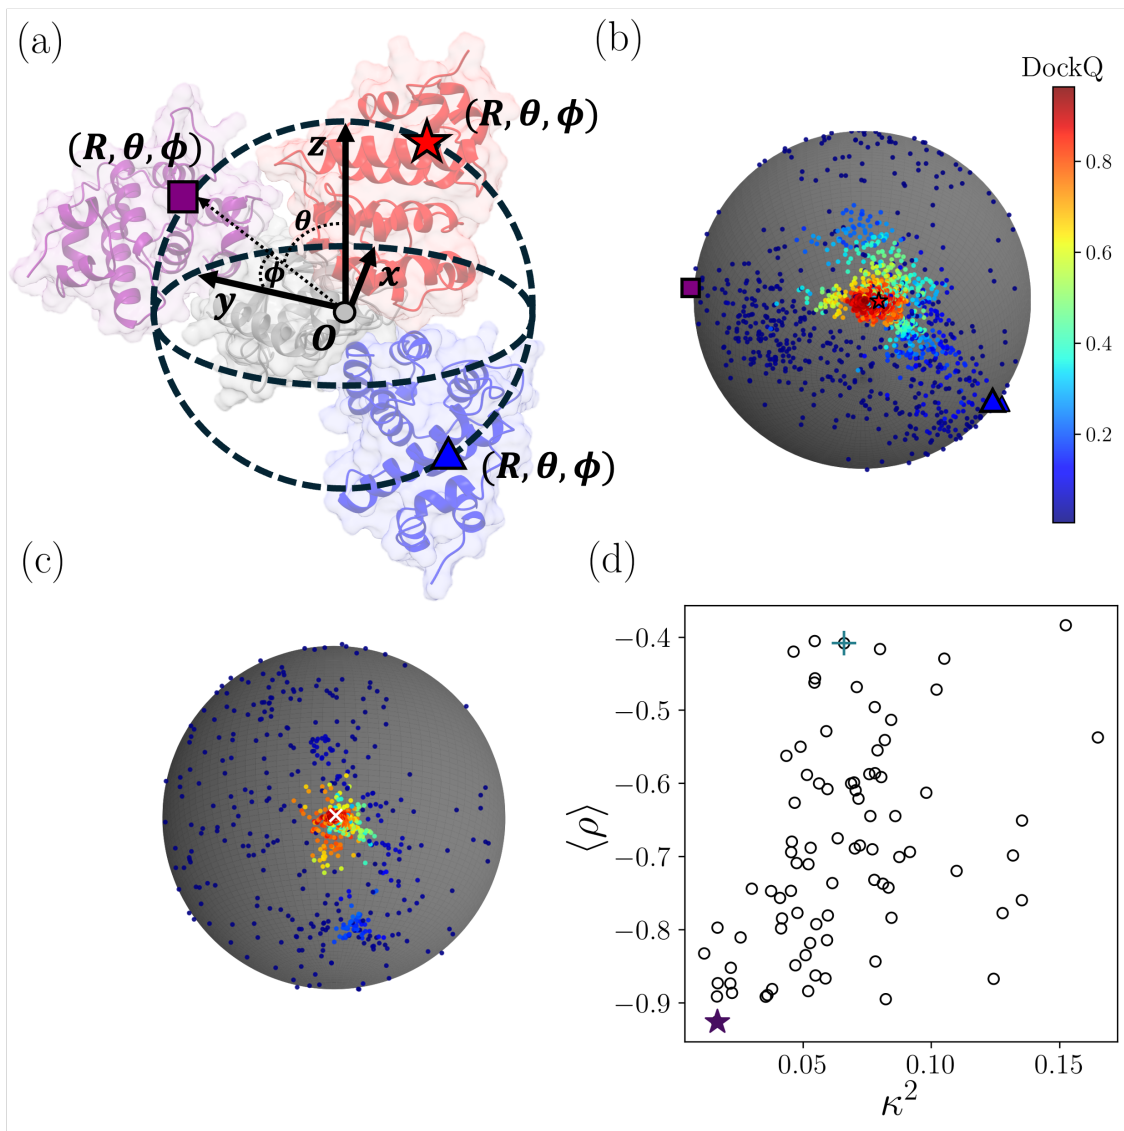


Figure 5: The shape of the DockQ landscape provides insight into the effectiveness of scoring computational models. (a) A schematic of re-docked models of heterodimer PDB 2GRN, which illustrates the variables that define the DockQ landscape. The receptor (gray shading) is located at the origin O with the same orientation as that in the x-ray crystal structure. The location of the ligand with the same orientation as that in the x-ray crystal structure is denoted using spherical coordinates, the distance from the origin r , the polar angle θ , and the azimuthal angle ϕ , for two models (blue triangle and purple square) and the x-ray crystal structure (red star). DockQ increasing from dark blue to dark red is plotted for all computational models for (b) 2GRN and (c) 3WHQ with Spearman correlations between DockQ and ZRank2 $\langle \rho \rangle = -0.41$ and -0.93 , respectively. The models were scored at arbitrary x -, y -, and z -coordinates, but plotted at $r = R$, where R is the separation between centers of mass of the receptor and ligand in the x-ray crystal structure. (d) The Spearman correlation $\langle \rho \rangle$ (between DockQ and PPI score) averaged over scoring functions plotted versus the relative anisotropy κ^2 of the DockQ landscape for all 84 targets. The values of ρ and κ^2 are highlighted for 2GRN (cross) and 3WHQ (star).

visualize and characterize the DockQ and PPI score landscapes in six dimensions, and thus we seek lower-dimensional representations. One approach is to fix the orientation of the ligand (e.g. to that in the x-ray crystal structure) and only consider the x -, y -, and z -coordinates of the ligand’s center of mass (given a fixed position and orientation of the receptor). We can remove another degree of freedom by plotting the scores at fixed radial distance $r = R$, where R is the distance between receptor and ligand centers of mass in the x-ray crystal structure. We can then visualize the DockQ score (and the PPI score) for each model as function of the spherical coordinate representation of the ligand’s position (R, θ, ϕ) , where θ and ϕ are the polar and azimuthal angles (Fig. 5 (a)).

We show the DockQ landscape for all computational models for the targets PDB 2GRN and 3WHQ in Fig. 5 (b) and (c). The models with the largest DockQ values in red are close to the native binding sites. The DockQ score decays rapidly (changing color from red to yellow to dark blue) as the models move further from the native binding site in θ and ϕ . As shown in Fig. 4, 2GRN is classified as a “hard” target with $\langle \rho \rangle = -0.41$ and 3WHQ is classified as an “easy” target with $\langle \rho \rangle = -0.93$. The DockQ landscape is more isotropic for the easy target 3WHQ compared to the landscape for the hard target 2GRN. To characterize the shape of the DockQ landscape, we calculate the moment of inertia tensor $I_{\alpha\beta}$ for the x -, y - and z -coordinates of the ligands for all models weighted by DockQ for each target. We define the relative anisotropy, $\kappa^2 = 1 - 3(\lambda_1\lambda_2 + \lambda_2\lambda_3 + \lambda_1\lambda_3)/(\lambda_1 + \lambda_2 + \lambda_3)^2$, of the landscape, where $\lambda_1 \geq \lambda_2 \geq \lambda_3$ are the eigenvalues of $I_{\alpha\beta}$. We show in Fig. 5 (d) that there is a positive correlation between κ^2 and $\langle \rho \rangle$ averaged over all scoring functions, i.e. $\langle \rho \rangle$ decreases (becomes more negative) as the large DockQ landscape becomes more isotropic. (See the methods section for more details.)

The magnitude of the Spearman correlation between DockQ and the PPI score is controlled by the differences in the DockQ and PPI score landscapes. In Fig. S4, we compare the DockQ and ZRank2 landscapes of PDB 3WHQ and 2GRN. For the easy target, 3WHQ, the ZRank2 landscape closely resembles the DockQ landscape, while the ZRank2 landscape does not overlap with the DockQ landscape for the hard target, 2GRN. These results emphasize that the isotropy of the DockQ landscape and the overlap between the DockQ and PPI score landscape influence $\langle \rho \rangle$ for each target. However, are there *physical* features of the binding interface that can provide information about the difficulty in accurately scoring PPI models?

We identify two important physical properties of protein-protein interfaces that influence scoring accuracy: the interface separability S and relative surface area f_a of the interface. (See Methods for more details.) The interface separability S is calculated using a support vector machine (SVM) to best separate the receptor and ligand by an order-three polynomial surface. The separability satisfies $0.5 \leq S \leq 1.0$ and the larger the separability, the flatter the interface. In Fig. 6 (a), we show that as the interfaces become more intertwined (S decreases), $|\langle \rho \rangle|$ averaged over scoring functions increases. In contrast, as $S \rightarrow 1$, $|\langle \rho \rangle|$ generally decreases, but the data for different targets are more scattered.

We define the relative surface area of the interface as

$$f_a = \frac{(SASA_r + SASA_l) - SASA_{r+l}}{SASA_r + SASA_l}, \quad (1)$$

where $SASA_r$ is the solvent accessible surface area of the receptor in its bound form in the absence of the ligand, $SASA_l$ is the solvent accessible surface area of the ligand in its bound form in the absence of the receptor, and $SASA_{r+l}$ is the solvent accessible surface area of the ligand bound to the receptor. In Fig. 6 (b), we show that targets with larger interfaces (i.e. $f_a \gtrsim 0.15$) possess large $|\langle \rho \rangle|$, whereas smaller interfaces with $f_a \lesssim 0.15$ typically possess smaller $|\langle \rho \rangle|$, but there is large scatter in the data. Since “easy” targets have a more isotropic DockQ landscape, and easy targets have low values of the separability and high values of the relative surface area of the interface f_a ,




Figure 6: The Spearman correlation between the DockQ and PPI scores $\langle \rho \rangle$, averaged over all scoring functions, plotted versus the (a) interface separability S and (b) fraction f_a of the surface area that is part of the interface for each target.

the anisotropy κ^2 is typically smaller for targets with low S and large f_a as shown in Fig. S5 (a) and (b). These results show that the physical properties of the interface, such as the separability and relative interface size, enable us to separate the “easy” from the “hard” to score interfaces.

The physical interface properties, such as S and f_a , can separate targets according to their scoring difficulty. Therefore, it is interesting to determine how these two physical interface properties, as well as the number of interface contacts N_c (i.e. heavy atoms on different protein chains with separations less than 4.5\AA), perform as scoring functions for computational models. We first calculate these physical properties on all models for each target that are uniformly sampled over DockQ. We then plot the Spearman correlation between DockQ and the PPI scores $\langle \rho \rangle_t$ averaged over all targets for each scoring function, including scores that are based on S , f_a , and N_c individually (Fig. 7 (a)). The interface separability score is the worst performing score out of S , f_a , and N_c with $\langle \rho \rangle_t = -0.577$, however, it still outperforms VoroMQA on average. The number of interface contacts had the best performance of the physical-metric-based scores with $\langle \rho \rangle_t = -0.722$. The success of the three physical interface properties individually in scoring computational models suggests that a combined scoring function that includes all three physical metrics can outperform the individual physical metrics. In Fig 7, we show results for a non-linear support vector regression (SVR) score that uses a gaussian kernel to combine the three physical metrics. (See Methods for more details.) The SVR score achieves $|\langle \rho \rangle_t| = 0.763$ across all targets, only slightly lower than ZRank2 ($|\langle \rho \rangle_t| = 0.778$), which is the highest performing PPI score. Thus, this simple three-feature SVR scoring function can match the performance of state-of-the-art PPI scoring functions.

We then measure the classification performance for the PPI scoring functions and the combined SVR score using the AUC of the ROC curve. Because the ROC requires an arbitrary cutoff between positives and negatives, we calculate AUC over a range of DockQ cutoffs $0.2 \leq \text{DockQ}_0 \leq 0.8$. If the DockQ score of a model is greater than or equal to the cutoff, it is classified as positive. (See Methods for more details.) In Fig. 7 (b), we plot $\langle \text{AUC} \rangle_t$ averaged over targets versus the cutoff DockQ_0 . The ordering of the Spearman correlation mirrors the classification performance at the

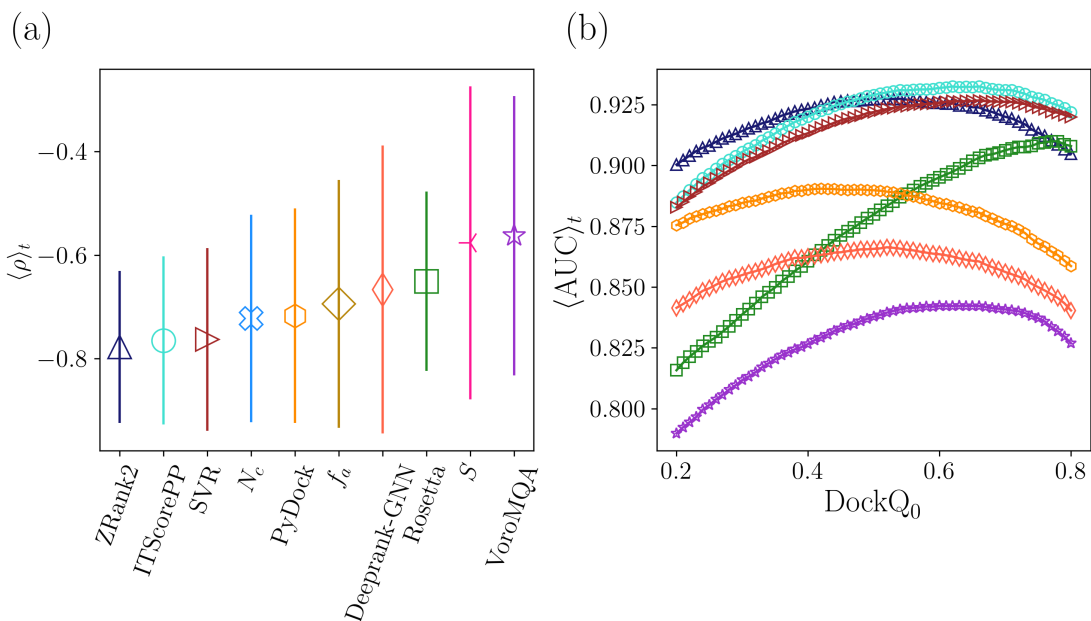


Figure 7: (a) The Spearman correlation between DockQ and the PPI score $\langle \rho \rangle_t$ (markers) averaged over targets and the standard deviation $\langle \Delta \rho \rangle_t$ (vertical lines) for each scoring function (including the individual physical features: f_a , N_c , and S). $\langle \rho \rangle_t$ for each scoring function is plotted in ascending order from left to right. (b) The AUC of the ROC plot $\langle \text{AUC} \rangle_t$ averaged over targets plotted versus the DockQ cutoff, DockQ_0 , for each scoring function: ZRank2 (blue triangles), ITScorePP (turquoise circles), Rosetta (green squares), PyDock (orange hexagons), VoronMQA (purple stars), Deepprank-GNN-ESM (red-orange diamonds), and the SVR scoring function with three physical features (burgundy rightward triangles).

lowest cutoff ($\text{DockQ}_0 = 0.2$) in Fig. 7 (b). As the DockQ cutoff increases, $\langle \text{AUC} \rangle_t$ the classification performance of most scoring functions does not change significantly. However, $\langle \text{AUC} \rangle_t$ for Rosetta increases by more than 0.1 as DockQ_0 increases, which suggests that Rosetta is better at identifying high-quality models (with $\text{DockQ}_0 \geq 0.5$). Notably, when $\text{DockQ}_0 = 0.8$, ITScorePP and the SVR score are the top classifiers of high-ranking models, both with $\langle \text{AUC} \rangle_t \approx 0.92$. We compare the performance of the SVR score and ZRank2 for each target in Fig. S6; we show that the SVR score has a larger $|\rho|$ than ZRank2 for 41 out of the 84 targets. In addition, the SVR scores are better than $\langle \rho \rangle$ averaged over the six scoring functions for 78 out of the 84 targets. Altogether, these results emphasize that the correlation and classification performance of the simple SVR score is comparable or better than all of the current PPI scoring functions.

3 Discussion

Predicting the interface between two protein monomers that form a heterodimer complex is an important, open problem. In this work, we focus on perhaps the simplest aspect of this problem: rigid-body re-docking of the bound forms of the monomers into the heterodimer, for which we have an experimentally-determined structure. When re-docking the bound forms, generation of the computational models is straightforward. Instead, the difficult task is to identify the model that

is closest to the native structure using a PPI scoring function. We evaluate the performance of seven PPI scoring functions (ZRank2, PyDock, VoromQA, Rosetta, ITScorePP, GNN-DOVE, and Deeprank-GNN-ESM) in carrying out this task. We test these scoring functions on a large dataset of models for 84 high-resolution x-ray crystal structure targets from the PDB. For each target, we generated an equal number of positive (DockQ ≥ 0.23) and negative (DockQ < 0.23) models. Our DockQ-balanced dataset gives consistent results for both classification and correlation measures between the PPI score and DockQ, i.e. we find that the Spearman correlations ρ between the PPI score and DockQ and the AUC of the ROC curve obey $AUC \approx -0.5\rho + 0.5$. We show that the targets in our dataset possess a range of difficulty in scoring, where some targets have significantly lower $|\rho|$ than others. We show that the scoring difficulty of the target increases as the anisotropy of the DockQ scoring landscape increases. Further, we identified several collective physical properties of the interface, such as interface separability S and the relative surface area of the interface f_a that can predict the scoring difficulty of the targets. Then, we constructed an SVR scoring function based on S , f_a , and the number of interfacial contacts N_c and show that it meets or exceeds the classification and correlation performance of the other state-of-the-art PPI scoring functions.

We find that an SVR with only three physical features matches or outperforms state-of-the-art PPI scoring functions with many more features, despite the long history of computational studies of PPIs^{24,26,30,45-49}. One possible reason for this lagging performance for current PPI scoring functions^{18,24-26,30,32-35} is that they have primarily been evaluated using classification-based metrics, such as hit rate plots, which employ quality cutoffs to define positive or negative models. However, using low cutoffs will make PPI scoring functions seem “successful,” even if the scoring functions possess very weak correlations with DockQ. Classification performance can be important, yet it should not be the sole metric upon which the PPI scoring function performance is assessed.

We have shown that the recent supervised GNN-based scoring functions, GNN-DOVE³⁴ and Deeprank-GNN-ESM¹⁸, have worse performance for the Spearman correlation between score and DockQ than all other scoring functions that we tested. One possible reason for the worse performance of these supervised GNN-based scoring functions is the quality imbalance in the datasets used for training. For example, GNN-DOVE is trained on the DockGround1.0 dataset⁵⁰, where only $\sim 9\%$ of models have DockQ ≥ 0.23 . Deeprank-GNN-ESM created their own training set, yet only $\sim 11\%$ of models have DockQ ≥ 0.23 . Supervised learning methods are most effective when trained on datasets with a uniform distribution of model quality. It is likely that GNN-based scoring functions will have improved performance if they are re-trained on datasets with more uniform distributions of model quality. Changing the architecture of the GNN to better represent physical features of the protein interface can also increase the scoring function performance. GNNs are composed of nodes (representing amino acids or atoms) with typically tens to thousands of features per node, yet they are currently unable to outperform the SVR in this study with only *three* physical features that describe the protein interface. In future studies, we will develop GNNs that first build-in successful physical features, such as the number of inter-chain contacts N_c . We will then add node features, such as hydrophobicity and charge of each amino acid, iteratively to improve the GNN scoring function performance and therefore identify important collective features of protein interfaces.

Template-free flexible docking is extremely difficult for PPI targets that undergo significant conformational changes upon binding. For example, recent studies were able to generate models where the interface C_α RMSD to the x-ray crystal structure was less than 5 Å for only 31% of the targets (i.e. 10 out of 32 targets that undergo significant conformational changes)⁹. To make more rapid progress in template-free flexible docking^{5,6,9,15}, it is important to develop PPI scoring functions with Spearman correlations with DockQ of $|\rho| \sim 1$ for rigid-body re-docking datasets.

Despite many years of development, all PPI scoring functions we tested have $|\langle \rho \rangle_t| < 0.8$ when averaged over all models and targets, but most of the scoring functions have $|\langle \rho \rangle_t| \leq 0.7$. Template-free flexible docking involves a series of rigid-body docking steps along a trajectory from the unbound to bound conformations of the monomers. The PPI score serves as the “potential energy landscape” that directs the system along this trajectory toward the native-like bound monomer conformations. Without scoring functions that achieve $|\rho| \approx 1$, models with both high and low DockQ values for a given target can end up with the same PPI score. In this case, flexible docking procedures will yield final protein conformations with a large C_α RMSD compared to the x-ray crystal structure target. Inconsistencies caused by inaccurate scoring functions are made worse when there are large differences between the unbound and bound forms of the monomers. Thus, we emphasize the importance of improving the Spearman correlations between PPI scores and DockQ for the field of computational modeling of PPIs.

4 Materials and Methods

Data Set Selection

The target data set is composed of 84 heterodimer x-ray crystal structures obtained from the Protein Data Bank (PDB). (See Table S1 in the Supplementary Material for a list of the structures.) The heterodimers were chosen according to the following criteria: all structures must have a resolution ≥ 3.5 Å; no non-protein polymers are allowed in the structure (e.g. nucleic acids or polysaccharides); no metal ions or bound small molecules are present in the structure; both protein chains must be at least 50 amino acids long; the proteins cannot be artificially engineered or include amino acid mutations; and there are no missing amino acids in both chains. In addition, all heterodimers were filtered so that the structures have less than 20% sequence similarity. Also, homodimers were excluded from the dataset to ensure that all targets occur as PPIs in a natural biological context and are not artifacts from crystallization⁵⁷.

In Fig. S7 (a), we show that there are very weak correlations between the resolution of the x-ray crystal structure and $\langle \rho \rangle$ up to 3.5Å. In addition, in Fig. S7 (b), we show that the average Spearman correlation for each scoring function does not depend strongly on resolution. Thus, the resolution of the x-ray crystal structure does not affect PPI scoring, at least for resolutions that are more accurate than 3.5 Å.

Model Generation and Sampling

The computational models were generated for each of the 84 heterodimer targets using ZDOCK 3.0.2 with 6-degree Euler angle increments, which yields 54,000 models per ZDOCK run. ZDOCK was run 10 times per target. For each run, the monomers were randomly rotated and translated in space to ensure that all computational models were distinct. We calculate $0 \leq \text{DockQ} \leq 1$ for all 540,000 models for each target and sort them into 20 evenly spaced DockQ bins. For each target, a maximum of 50 models were selected from each DockQ bin, yielding a maximum of 1000 models per target. In addition, to provide even sampling of negative models, 1000 models with $\text{DockQ} < 0.23$ were randomly selected from the models that were not already selected. All selected models had hydrogens added to them using `Reduce`⁵¹ from the PHENIX software package⁵² and were then relaxed using Rosetta v3.12^{28,29}. After relaxation, the set of models was balanced so that each target has an equal number of positive ($\text{DockQ} \geq 0.23$) and negative ($\text{DockQ} < 0.23$) models.

Scoring Protein Models

The ~ 1000 uniformly sampled models for each target were scored using ZRANK2^{24,25}, ITScorePP²⁷, Rosetta 3.12 (using the ref2015 weight function)^{9,29}, VoromQA version 1.26³⁰, PyDock3²⁶, Deeprank-GNN-ESM¹⁸, and GNN-DOVE³⁴.

Hit Rate Plots

The hit rate fraction h is defined as

$$h = \frac{1}{\mathcal{N}} \sum_{k=1}^{\mathcal{N}} \frac{1}{N_m^{k,+}} \sum_{i=1}^R \Theta(\text{DockQ}_k^i - \text{DockQ}_0), \quad (2)$$

where \mathcal{N} is the number of targets in the dataset, R is the maximum rank considered in the set of models ordered by a given PPI score, $N_m^{k,+}$ is the number of positives in the set of models up to rank R , DockQ_k^i is the DockQ value of the i th model for target k , $\text{DockQ}_0 = 0.23$ is the DockQ cutoff that defines a "positive" model, and $\Theta(\cdot)$ is the Heaviside step function which is nonzero for $\text{DockQ}_k^i > \text{DockQ}_0$.

The effective hit rate plots in Fig. 1 (b) and (d) were calculated for targets (b) 3RCZ and (d) 3YGS. $N_t = 108,000$ models were generated for 3RCZ, while $N_t = 20,000$ were generated for 3YGS. The effective hit rate fraction h_k for the k th target is defined as

$$h_k = \frac{1}{N_s} \sum_{l=1}^{N_s} \frac{1}{N_m^{k,l,+}} \sum_{i=1}^R \Theta(\text{DockQ}_{k,l}^i - \text{DockQ}_0). \quad (3)$$

For each target, multiple samples with sizes N_m are drawn from the total number of models N_t ranging from $N_m = 500$ to 15,000. $N_s = 1000$ samples are drawn for each value of N_m , and h_k is calculated at each rank i using Eq. 3, where $\text{DockQ}_{k,l}^i$ is the DockQ value for the l th sample at rank i for the k th target. We only consider up to rank $R = 100$, and the models are scored and ordered using the Rosetta scoring function.

In Fig. 2 (a), we calculated the effective hit rate h_k at $i = 100$ for each target k as a function of N_m/N_t using the exhaustively sampled model dataset, not the dataset that was uniformly sampled in DockQ. We collected $N_s = 1000$ random samples of size N_m from $N_t = 10,000$ models. All models in each set were ordered using the ZRank2 scoring function.

DockQ Landscape for Heterodimers

The DockQ landscape in Fig. 5 (b) and (c) contains a series of points (in spherical coordinates) that represent the centers of mass (COM) of the ligand C_α atoms of computational models for a given target. We assign the receptor and ligand using the same definition as that in the target PDB file. All models are aligned so that the receptor in each model has minimal C_α -RMSD to the receptor in the x-ray crystal structure of the target. Each model is translated so that the center of mass of the receptor is at the origin. In the DockQ landscape, the ligand COM position for each model is placed at distance R between the centers of mass of the receptor and ligand in the x-ray crystal structure of the target.

Relative Anisotropy of the DockQ Landscape

The anisotropy κ^2 of the DockQ landscape is calculated using the eigenvalues of the moment of inertia matrix formed from the centers of mass of the ligands of the computational models. The

moment of inertia matrix for a single target is defined as

$$I_{\alpha\beta} = \frac{1}{N_m} \sum_{i=1}^{N_m} \text{DockQ}_i^3 \left(r_i^2 \delta_{\alpha\beta} - r_\alpha^i r_\beta^i \right), \quad (4)$$

where N_m is the number of models, $\alpha, \beta = x, y, z$, $\delta_{\alpha\beta}$ is the Kronecker delta, r_α^i is α -component of the position of the i th point, and $r_i^2 = x_i^2 + y_i^2 + z_i^2$. The moment of inertia is weighted by DockQ_i^3 to emphasize native-like models. The components of the moment of inertia matrix \mathbf{I} are

$$\mathbf{I}_{\alpha\beta} = \begin{pmatrix} I_{xx} & I_{xy} & I_{xz} \\ I_{yx} & I_{yy} & I_{yz} \\ I_{zx} & I_{zy} & I_{zz} \end{pmatrix}. \quad (5)$$

\mathbf{I} can be expressed as

$$\mathbf{I} = \mathbf{U}^t \mathbf{\Lambda} \mathbf{U}, \quad (6)$$

where \mathbf{U} represents an orthogonal matrix whose columns are the normalized eigenvectors of \mathbf{I} and $\mathbf{\Lambda}$ is the diagonal matrix of eigenvalues,

$$\mathbf{\Lambda} = \begin{pmatrix} \lambda_1 & 0 & 0 \\ 0 & \lambda_2 & 0 \\ 0 & 0 & \lambda_3 \end{pmatrix}, \quad (7)$$

where $\lambda_1 \leq \lambda_2 \leq \lambda_3$. The eigenvalues of \mathbf{I} can be used to calculate the relative anisotropy:

$$\kappa^2 = 1 - 3 \frac{\lambda_1 \lambda_2 + \lambda_2 \lambda_3 + \lambda_1 \lambda_3}{(\lambda_1 + \lambda_2 + \lambda_3)^2}. \quad (8)$$

Physical Features of Proteins

We calculate the separability S of each interface using a degree-3 polynomial support vector machine (SVM) from the SciKit-Learn package in python⁵³. The separability satisfies $0.5 < S < 1.0$, where the SVM surface can perfectly separate the ligand and receptor for $S = 1.0$, while $S = 0.5$ corresponds to a highly intertwined interface.

The number of interface contacts N_c is defined as the number of heavy atom pairs in different protein chains with separations less than $d_c = 4.5 \text{ \AA}$. We also calculated the number of contacts for all possible distance cutoffs from $d_c = 1-30 \text{ \AA}$ with an increment of 0.1 \AA . We find that $d_c = 4.5 \text{ \AA}$ yielded the largest Spearman correlation between N_c and DockQ for the uniformly sampled dataset.

The fractional interface surface area f_a is defined by

$$f_a = \frac{(\text{SASA}_r + \text{SASA}_l) - \text{SASA}_{r+l}}{\text{SASA}_r + \text{SASA}_l}, \quad (9)$$

where SASA_{r+l} is the SASA of the bound model, and SASA_r and SASA_l are the SASA for the fully solvated subunits for the receptor and ligand, respectively. The SASA for all models was calculated using the FreeSASA Python API⁵⁴ with the Lee and Richards⁵⁵ algorithm and atom sizes specified in Gaines, *et al.*⁵⁶, which recapitulate the side chain dihedral angle distributions in high-resolution protein x-ray crystal structures.

Physical SVR Score

The SVR scores were calculated using the leave-one-out method, where each target and its models were tested on an SVR scoring function trained on the models for the 83 other targets. The input for the SVR were the three physical features: interface separability S , the number of interface contacts N_c , and the fractional interface surface area f_a . The input features were normalized before training to have zero mean and unit standard deviation. The kernel coefficient for the Gaussian kernel was set to $\gamma = 1/[N_f(\langle X^2 \rangle - \langle X \rangle^2)]$, where $N_f = 3$ is the number of features on which the SVR model was trained and X is the training vector with dimensions $\mathcal{N} \times N_m \times N_f$. The SVR scoring function was trained to predict DockQ from the input vector X .

References

- [1] Orchard S, Ammari M, Aranda B, Breuza L, Briganti L, Broackes-Carter F, Campbell NH, Chavali G, Chen C, del-Toro N, et al. (2014) The MIntAct project—IntAct as a common curation platform for 11 molecular interaction databases. *Nucleic Acids Res.* 42:D358-63
- [2] Luck K, Kim DK, Lambourne L, Spirohn K, Begg BE, Bian W, Brignall R, Cafarelli T, Campos-Laborie FJ, Charlotheaux B, et al. (2020) A reference map of the human binary protein interactome. *Nature* 580(7803):402-408
- [3] Drew K, Wallingford JB, Marcotte EM (2021) hu.MAP 2.0: Integration of over 15,000 proteomic experiments builds a global compendium of human multiprotein assemblies. *Mol. Syst. Biol.* 17(5):e10016
- [4] Burke DF, Bryant P, Barrio-Hernandez I, Memon D, Pozzati G, Shenoy A, Zhu W, Dunham AS, Albanese P, Keller A, et al. (2023) Towards a structurally resolved human protein interaction network. *Nat. Struct. Mol. Biol.* 30(2):216-225
- [5] Dominguez C, Boelens R, Bonvin AM (2003) HADDOCK: A protein-protein docking approach based on biochemical or biophysical information. *J. Am. Chem. Soc.* 125(7):1731-7
- [6] Torchala M, Moal IH, Chaleil RA, Fernandez-Recio J, Bates PA (2013) SwarmDock: A server for flexible protein-protein docking. *Bioinformatics* 29(6):807-9
- [7] Pierce BG, Wiehe K, Hwang H, Kim BH, Vreven T, Weng Z (2014) ZDOCK server: Interactive docking prediction of protein-protein complexes and symmetric multimers. *Bioinformatics* 30(12):1771-3
- [8] Yan Y, Zhang D, Zhou P, Li B, Huang SY (2017) HDock: A web server for protein-protein and protein-DNA/RNA docking based on a hybrid strategy. *Nucleic Acids Res.* 45(W1):W365-W373
- [9] Marze NA, Roy Burman SS, Sheffler W, Gray JJ (2018) Efficient flexible backbone protein-protein docking for challenging targets. *Bioinformatics* 34(20):3461-3469
- [10] Lensink MF, Velankar S, Baek M, Heo L, Seok C, Wodak SJ (2018) The challenge of modeling protein assemblies: The CASP12-CAPRI experiment. *Proteins* 86 Suppl 1:257-273
- [11] Gainza P, Sverrisson F, Monti F, Rodolà E, Boscaini D, Bronstein MM, Correia BE (2020) Deciphering interaction fingerprints from protein molecular surfaces using geometric deep learning. *Nat. Methods* 17(2):184-192
- [12] Lensink MF, Brysbaert G, Mauri T, Nadzirin N, Velankar S, Chaleil RAG, Clarence T, Bates PA, Kong R, Liu B, et al. (2021) Prediction of protein assemblies, the next frontier: The CASP14-CAPRI experiment. *Proteins* 89(12):1800-1823
- [13] Evans R, O'Neill M, Pritzel A, Antropova N, Senior A, Green T, Zidek A, Bates R, Blackwell S, Yim J, et al. (2022) Protein complex prediction with AlphaFold-Multimer. *bioRxiv* 2021.10.04.463034
- [14] Lensink MF, Brysbaert G, Raouraoua N, Bates PA, Giulini M, Honorato RV, van Noort C, Teixeira JMC, Bonvin AMJJ, Kong R, et al. (2023) Impact of AlphaFold on structure prediction of protein complexes: The CASP15-CAPRI experiment. *Proteins* 91(12):1658-1683

- [15] Gray JJ, Moughon S, Wang C, Schueler-Furman O, Kuhlman B, Rohl CA, Baker D (2003) Protein-protein docking with simultaneous optimization of rigid-body displacement and side-chain conformations. *J. Mol. Biol.* 331(1):281-99
- [16] Jumper J, Evans R, Pritzel A, Green T, Figurnov M, Ronneberger O, Tunyasuvunakool K, Bates R, Žídek A, Potapenko A, et al. (2021) Highly accurate protein structure prediction with AlphaFold. *Nature* 596(7873):583-589
- [17] Zhu W, Shenoy A, Kundrotas P, Elofsson A (2023) Evaluation of AlphaFold-Multimer prediction on multi-chain protein complexes. *Bioinformatics* 39(7):btad424
- [18] Xu X, Bonvin AMJJ (2024) DeepRank-GNN-esm: A graph neural network for scoring protein-protein models using protein language model. *Bioinform. Adv.* 4(1):vbad191
- [19] Vreven T, Moal IH, Vangone A, Pierce BG, Kastritis PL, Torchala M, Chaleil R, Jiménez-García B, Bates PA, Fernandez-Recio J, et al. (2015) Updates to the Integrated Protein-Protein Interaction Benchmarks: Docking Benchmark Version 5 and Affinity Benchmark Version 2. *J. Mol. Biol.* 427(19):3031-41
- [20] Guest JD, Vreven T, Zhou J, Moal I, Jeliazkov JR, Gray JJ, Weng Z, Pierce BG (2021) An expanded benchmark for antibody-antigen docking and affinity prediction reveals insights into antibody recognition determinants. *Structure* 29(6):606-621
- [21] Desta IT, Porter KA, Xia B, Kozakov D, Vajda S (2020) Performance and its limits in rigid body protein-protein docking. *Structure* 28(9):1071-1081
- [22] Tovchigrechko A, Vakser IA (2006) GRAMM-X public web server for protein-protein docking. *Nucleic Acids Res.* 34(Web Server issue):W310-4
- [23] Moal IH, Bates PA (2010) SwarmDock and the use of normal modes in protein-protein docking. *Int. J. Mol. Sci.* 11(10):3623-48
- [24] Pierce B, Weng Z (2007) ZRANK: Reranking protein docking predictions with an optimized energy function. *Proteins* 67(4):1078-86
- [25] Pierce B, Weng Z (2008) A combination of rescoring and refinement significantly improves protein docking performance. *Proteins* 72(1):270-9
- [26] Cheng TM, Blundell TL, Fernandez-Recio J (2007) pyDock: Electrostatics and desolvation for effective scoring of rigid-body protein-protein docking. *Proteins* 68(2):503-15
- [27] Huang SY, Zou X (2008) An iterative knowledge-based scoring function for protein-protein recognition. *Proteins* 72(2):557-79
- [28] Park H, Bradley P, Greisen P Jr, Liu Y, Mulligan VK, Kim DE, Baker D, DiMaio F (2016) Simultaneous optimization of biomolecular energy functions on features from small molecules and macromolecules. *J. Chem. Theory Comput.* 12(12):6201-6212
- [29] Alford RF, Leaver-Fay A, Jeliazkov JR, O’Meara MJ, DiMaio FP, Park H, Shapovalov MV, Renfrew PD, Mulligan VK, Kappel K, et al. (2017) The Rosetta all-atom energy function for macromolecular modeling and design. *J. Chem. Theory Comput.* 13(6):3031-3048
- [30] Olechnovič K, Venclovas Č (2017) VoromQA: Assessment of protein structure quality using interatomic contact areas. *Proteins* 85(6):1131-1145

- [31] Zhou H, Skolnick J (2011) GOAP: A generalized orientation-dependent, all-atom statistical potential for protein structure prediction. *Biophys. J.* 101(8):2043-52
- [32] Renaud N, Geng C, Georgievska S, Ambrosetti F, Ridder L, Marzella DF, Réau MF, Bonvin AMJJ, Xue LC (2021) DeepRank: A deep learning framework for data mining 3D protein-protein interfaces. *Nat. Commun.* 12(1):7068
- [33] Wang X, Terashi G, Christoffer CW, Zhu M, Kihara D (2020) Protein docking model evaluation by 3D deep convolutional neural networks. *Bioinformatics* 36(7):2113-2118
- [34] Wang X, Flannery ST, Kihara D (2021) Protein docking model evaluation by graph neural networks. *Front. Mol. Biosci.* 8:647915
- [35] Réau M, Renaud N, Xue LC, Bonvin AMJJ (2023) DeepRank-GNN: A graph neural network framework to learn patterns in protein-protein interfaces. *Bioinformatics* 39(1):btac759
- [36] Méndez R, Lepplae R, De Maria L, Wodak SJ (2003) Assessment of blind predictions of protein-protein interactions: Current status of docking methods. *Proteins* 52(1):51-67
- [37] Moal IH, Torchala M, Bates PA, Fernández-Recio J (2013) The scoring of poses in protein-protein docking: Current capabilities and future directions. *BMC Bioinformatics* 14:286
- [38] Huang SY (2015) Exploring the potential of global protein-protein docking: An overview and critical assessment of current programs for automatic ab initio docking. *Drug Discov. Today* 20(8):969-77
- [39] Basu S, Wallner B (2016) DockQ: A quality measure for protein-protein docking models. *PLoS One* 11(8):e0161879
- [40] Berman HM, Westbrook J, Feng Z, Gilliland G, Bhat TN, Weissig H, Shindyalov IN, Bourne PE (2000) The Protein Data Bank. *Nucleic Acids Res.* 28(1):235-42
- [41] Chen R, Li L, Weng Z (2003) ZDOCK: An initial-stage protein-docking algorithm. *Proteins* 52(1):80-7
- [42] Pierce BG, Hourai Y, Weng Z (2011) Accelerating protein docking in ZDOCK using an advanced 3D convolution library. *PLoS One* 6(9):e24657
- [43] Prudden J, Perry JJ, Nie M, Vashisht AA, Arvai AS, Hitomi C, Guenther G, Wohlschlegel JA, Tainer JA, Boddy MN (2011) DNA repair and global sumoylation are regulated by distinct Ubc9 noncovalent complexes. *Mol. Cell. Biol.* 31(11):2299-310
- [44] Qin H, Srinivasula SM, Wu G, Fernandes-Alnemri T, Alnemri ES, Shi Y (1999) Structural basis of procaspase-9 recruitment by the apoptotic protease-activating factor 1. *Nature* 399(6736):549-57
- [45] Wodak SJ, Janin J (1978) Computer analysis of protein-protein interaction. *J. Mol. Biol.* 124(2):323-42
- [46] Shoichet BK, Kuntz ID (1991) Protein docking and complementarity. *J. Mol. Biol.* 221(1):327-46
- [47] Cherfils J, Duquerroy S, Janin J (1991) Protein-protein recognition analyzed by docking simulation. *Proteins* 11(4):271-80

- [48] Camacho CJ, Gatchell DW, Kimura SR, Vajda S (2000) Scoring docked conformations generated by rigid-body protein-protein docking. *Proteins* 40(3):525-37
- [49] Olechnovič K, Venclovas Č (2023) VoroIF-GNN: Voronoi tessellation-derived protein-protein interface assessment using a graph neural network. *Proteins* 91(12):1879-1888
- [50] Kundrotas PJ, Anishchenko I, Dauzhenka T, Kotthoff I, Mnevets D, Copeland MM, Vakser IA (2018) Dockground: A comprehensive data resource for modeling of protein complexes. *Protein Sci.* 27(1):172-181
- [51] Davis IW, Leaver-Fay A, Chen VB, Block JN, Kapral GJ, Wang X, Murray LW, Arendall WB 3rd, Snoeyink J, Richardson JS, et al. (2007) MolProbity: All-atom contacts and structure validation for proteins and nucleic acids. *Nucleic Acids Res.* 35(Web Server issue):W375-83
- [52] Adams PD, Afonine PV, Bunkóczy G, Chen VB, Echols N, Headd JJ, Hung LW, Jain S, Kapral GJ, Grosse Kunstleve RW, et al. (2011) The Phenix software for automated determination of macromolecular structures. *Methods* 55(1):94-106
- [53] Pedregosa F, Varoquaux G, Gramfort A, Michel V, Thirion B, Grisel O, Blondel M, Prettenhofer P, Weiss R, Dubourg V, et al. (2011) Scikit-learn: Machine learning in Python. *Journal of machine learning research.* 2825–30
- [54] Mitternacht S (2016) FreeSASA: An open source C library for solvent accessible surface area calculations. *F1000 Res.* 5:189
- [55] Lee B, Richards FM (1971) The interpretation of protein structures: Estimation of static accessibility. *J. Mol. Biol.* 55(3):379-400
- [56] Gaines JC, Acebes S, Virrueta A, Butler M, Regan L, O'Hern CS (2018) Comparing side chain packing in soluble proteins, protein-protein interfaces, and transmembrane proteins. *Proteins* 86(5):581-591
- [57] Schweke H, Xu Q, Tauriello G, Pantolini L, Schwede T, Cazals F, Lhéritier A, Fernandez-Recio J, Rodríguez-Lumbreras LA, et al. (2023) Discriminating physiological from non-physiological interfaces in structures of protein complexes: A community-wide study. *Proteomics* 23(17):e2200323