Lecture Summary: Computational Modeling of Protein-Protein Interactions

[Module: 26s3 - Task 2: A combined comprehensive summary]
[Deadline: 4th May 11:59pm (adjusted by TA)]

Color code:
Red = Summary 1
Yellow = Summary 2
Green = Summary 3
Blue = My own additions and transitions

Background and Motivation

Proteins usually do not work alone. Many cellular functions depend on protein-protein interactions, where proteins form complexes or interaction networks involved in processes such as signaling, catalysis, and transport. The lecture emphasized that the possible interaction space is very large: the human proteome contains about 25,000 distinct proteins, which leads to an estimated tens of millions of possible dimeric interactions. However, only a much smaller fraction of these possible interactions have been structurally resolved. This gap motivates computational modeling of protein-protein interactions. In general, these models aim to ask whether two proteins bind, where they bind, what 3D complex they form, and how strong the interaction is. The main question is not only how to generate possible protein complex structures, but also how to evaluate which predicted structures are closest to the true bound form.

Objectives of Computational PPI Modeling

Computational PPI modeling is not a single problem, but a set of related questions about protein association. At its core, the field asks whether two proteins bind, what 3D complex they form, which residues control the interface, and how strong or specific the interaction is. These goals are connected, but they are usually treated as different computational tasks, including partner prediction, interface localization, complex structure prediction, and affinity estimation. This distinction matters because doing well on one task does not guarantee success on another. For example, a model may correctly identify a surface patch involved in binding, but still fail to orient the partner protein in the correct 3D arrangement. Similarly, a docking algorithm may generate a geometrically plausible complex, but the docking score is often not a calibrated measure of true binding affinity. As a result, a method may work well for stable protein assemblies but perform worse for transient signaling interactions. This task separation is important for understanding why the lecture later focuses specifically on docking poses and scoring functions.

Algorithmic Pillars: Sampling and Scoring

Docking can be understood through two main steps: sampling and scoring. Sampling means searching over possible rotations, translations, and sometimes conformational states to generate many candidate poses, often called decoys. Some docking programs can sample a very large number of poses; for example, ClusPro-like rigid-body docking can rotate the ligand through many orientations to search the possible binding space. After these candidate poses are generated, a scoring function is used to rank them. These scores often combine physical or geometric terms such as shape complementarity, electrostatics, and desolvation. This separation is important because a docking method can fail in two different ways: it may fail to generate a near-native pose, or it may generate one but fail to rank it highly. Therefore, the difficult part is not only making plausible protein contacts, but also distinguishing the biologically correct pose from many near-miss decoys. This helps explain why the lecture treats sampling and scoring as related but separate bottlenecks.

Metrics of Success: DockQ and Ground Truth

To evaluate a docking prediction, we need a reference structure, usually an experimentally resolved complex structure that can be treated as the ground truth. DockQ is used to measure how close a predicted docking model is to this true structure. It combines several CAPRI-style features, including native contacts, interface RMSD, and ligand RMSD, into one score between 0 and 1. In this scale, DockQ = 1 means the model is essentially the correct structure, while DockQ close to 0 means the model is very far from the native complex. The usual cutoffs are 0.23 for acceptable models, 0.49 for medium-quality models, and 0.80 for high-quality models. By comparing a scoring function's output with DockQ, we can test whether the scoring function actually gives better rankings to models that are closer to the experimental structure.

Critical Limitations of Current Scoring Functions

A central point of the lecture is that the bottleneck in docking is not always sampling, but often scoring. Even in the simplified setting of bound-form rigid-body redocking, where the monomers are already close to their bound conformations and the search space can be sampled more completely, scoring functions still do not always rank the correct models well. In this setting, a useful scoring function should show a strong relationship with DockQ. For energy-like scores, this relationship is expected to be strongly negative, because lower scores should correspond to higher DockQ values. However, many current scoring functions show weak correlations or inconsistent behavior across targets. In some examples, an incorrect decoy can receive a very favorable low-energy score, which means the scoring function is being misled by a wrong pose. This suggests that current scoring functions still do not fully capture the physical features needed to reliably distinguish near-native models from attractive but incorrect decoys.

Physical Determinants of Interface Scoring and Future Directions

The lecture also suggests that scoring difficulty can be partly explained by physical features of the binding interface. Interfaces with more interfacial contacts or stronger intertwinement, which means a more interlocking and non-flat interface geometry, provide clearer geometric

constraints, so the native-like pose is easier to distinguish from incorrect decoys. In contrast, flatter interfaces are harder to score because many rigid-body displacements can still look geometrically plausible. These results suggest that future scoring functions may benefit from including simple physical descriptors such as contact number, contact fraction, and interface flatness. Another future challenge is the transition from rigid-body docking to flexible docking. Once monomers deviate from their bound conformations, scoring performance can drop, as shown by the RMSD analysis in the lecture. Therefore, future PPI modeling methods may need to combine better interface-aware scoring with better ways to model unbound-to-bound conformational changes.

Sources Used

Red — Summary 1: Lecture summary focusing on protein simulation, protein packing, and basic PPI scoring concepts.
Yellow — Summary 2: Literature-style overview of computational PPI modeling tasks, docking methods, scoring functions, and validation.
Green — Summary 3: Lecture-specific summary focusing on bound-form redocking, DockQ correlation, interface physical features, and flexibility.
Blue — Own additions: Connecting sentences, wording changes, and synthesis across the three summaries.