AI in Drug Discovery and Development

Prof. Rajnish Kumar
Dept. of Pharmaceutical Engineering and Technology
IIT-(BHU), Varanasi

Week-05

Lecture-23

Welcome to the course "AI in Drug Discovery and Development." In this session, we will talk about AI-assisted molecular docking. By the end of this lecture, you will be able to understand molecular recognition and molecular docking. Understand the role of AI in molecular docking and its advantages over traditional methods. Recognize AI's impact on post-prediction and virtual screening, and also analyze the challenges and future trends in AI-assisted docking. So, molecular recognition is one of the important phenomena in biochemistry, where we try to understand how a receptor or a target recognizes the molecule.

So, whether it is a substrate for that receptor or it is an inhibitor or an activator, how does this recognition happen? So, it can be between the enzymes and their substrates, or it can be between the receptors and the ligands, the agonist or antagonist. Or it can be the interaction between the recognition of the antigens and their antibodies. So, there are several approaches to investigate molecular recognition. So, one of them is molecular docking, but there are other methods, such as free energy calculations, like using MMGBSA, free energy perturbation, and thermodynamic integration.

or by using quantum mechanics and molecular mechanics methods, which are also known as hybrid methods. where a part of the system is defined using QM and the rest of the part is defined using the MM, and there are a lot more methods. But the idea is that we try to understand how this enzyme or receptor recognizes the molecule that is coming and binding to its binding pocket, and then further catalysis happens. So, we talk about binding energetics. So, basically, it is defined by the Gibbs energy of binding, where the delta G is equal to delta H minus T delta S.

Where delta H represents enthalpy and delta S represents entropy, it can also be written as delta G equals minus RT natural log Ki, which is equal to delta H minus T delta S. So, we can understand from this that molecular recognition depends on both enthalpy and entropy as well. So, where does the enthalpy come from, or what are the factors that contribute to the enthalpy? So, there are direct interactions between ligands, solvents, proteins, or ions. Because when the protein is in, you know, a buffer or a solution. So, the solvent can occupy all those binding pocket residues; they can interact with those binding pocket residues.

And this enthalpy contribution comes from ligand-protein interactions, ligand-solvent interactions, and solvent-protein interactions as well. and also, the conformational changes that happen during binding, so there might be conformational changes happening in the ligand itself, or there might be conformational changes happening in the side chains of the amino acid residues of those receptors or the targets. And the entropy contributions come from the rotational and translational entropy, conformational entropy, solvent reorganization, which is referred to as hydrophobicity, and the vibrational entropy. So, if you look at the biomolecular interactions, how does this happen? So, there are different kinds of factors, or we can say the forces, which are involved in defining this biomolecular interaction. So, there are intermolecular forces that exist between two molecules.

So, these can be electrostatic attractions or repulsions between the electrostatic charges, and then it can be a dipolar interaction as well. It can be H-bonding where a hydrogen bond donor and a hydrogen bond acceptor are involved in hydrogen bond formation. And then hydrophobicity can also be involved in intermolecular forces in this bimolecular interaction. And then we have vander Waals forces, which are also involved. And then, if you talk about intramolecular forces.

So, within the molecule. So, the two atoms are bound to each other with a chemical bond. And then the bond length, bond angle, and the dihedral angle define the intramolecular forces that are responsible for biomolecular interactions. We try to estimate the binding affinity. So, there they are; usually, the entropy is size dependent.

So, now, when we know that the binding affinity is composed of enthalpy and entropy, So, if we can determine both the enthalpy and entropy of the binding. So, we will be able to determine the binding affinity with confidence. So, entropy is usually size-dependent. So, rotational, translational, conformational, and vibrational. So, where is the water that was bound in the cavity released during the binding, and how tightly were they bound? So, if we can determine that that can help us to determine the entropy of the binding.

And then it can be measured using calorimetric methods, which have instruments like the biophysical methods, such as isothermal titration calorimetry or differential scanning calorimetry. So, in isothermal titration calorimetry, you determine the energy that is released or absorbed during a reaction when you add a ligand solution to the target solution. And, you measure the change in the energy actually, but the issue is that entropy is very difficult to estimate computationally. And then there is the hydrophobicity, which is connected to size. So, if the size is larger, the hydrophobicity will be larger.

And then we see that the enthalpy usually deals with the direct binding effects as well as solvent effects. So, we also need to take into account the conformational effects that are affecting the enthalpy components. So, if we talk about molecular docking, So, the docking studies are computational techniques for the exploration of possible binding modes of a substrate to a given receptor, enzyme, or other binding site. So, that is a, you know, IUPAC recommendation definition of docking. So, what we do here is try to identify the low-energy binding mode of a small molecule within the active site of a macromolecule or receptor whose structure is known.

And we have seen in the previous session that the structure can be determined either from the Protein Data Bank or by using the predicted structure from AlphaFold, RosettaFold, or homology modeling. And then, the computational determination of binding affinity between the protein structure and ligand is what we are doing in molecular docking. So, we have, for example, a ligand and a receptor, and then we try to see how this ligand can interact with this receptor or target, and if we can calculate the binding affinity; that is the objective of molecular docking. So, this is a co-crystalline structure of one of the inhibitors of the SARS-CoV-2 M-Pro main protease. So, you can see here it is bound to the active site.

So, this pocket is the active site of the SARS-CoV-2 main protease, and this is an inhibitor that is bound inside the active pocket. So, the idea of molecular docking is to get a kind of pose that closely resembles this co-crystal pose. So, if the obtained pose is very different from the co-crystal pose, then we can say that our docking algorithm is not very good. And if it is nicely aligned, the root means square deviation, which is a measure to determine how good the docking method is, should be considered. So, if the RMSD is very low, it means that it can replicate the co-crystal pose; then we can say that our docking method is good.

So, what do we do with docking? So, how do we use docking? So, we can reproduce the binding mode of a co-crystalline ligand, and then we can also predict the binding mode of known active ligands as well. So, we do not have the co-crystalline structure of those ligands, but we can use molecular docking to see how they fit into the binding pocket and determine their binding affinity. We can predict the binding affinities of related compounds from a known active series, as we have designed maybe 100 compounds based on one known inhibitor of an enzyme. And we wanted to see how good those 100 compounds could be. So, we can do that with the help of molecular docking.

And then, obviously, we can identify new ligands using virtual screening based on molecular docking methods. So, you can see here, for example, in this case, this ligand is fitting nicely into the binding pocket of this receptor. So, you can see a score of minus 9.1 it indicates a nice fitting while in this poach this ligand is not fitting very nicely and it has a score of minus 7.6 which is lower than the score shown above for nicely fitted ligand into the binding pocket.

And then predicting the binding affinity, you can compare this delta G prediction with the experimental measurement we are getting from the ITC. So that we can see how closely your docking method is able to reproduce the experimentally observed delta G, which will indicate the performance of the docking method. And then we can do the relative binding affinity prediction as well. So, we have two ligands; for example, ligand 1, which is shown here as having a docking score of minus 9.3, and another ligand that has a docking score of minus 8.

5. So, we can see that the ligand showing a docking score of minus 9.3 is better compared to the other ligand showing a docking score of minus 8.5. So, then there are multiple approaches for molecular docking. So, we have the rigid ligand-rigid protein, which is historically the first approach where we kept both the protein and the ligand rigid.

And so, what we were trying to do was search for the relative orientation of the two molecules with the lowest energy and then conduct the conformational analysis for both the ligand and the protein. So, FLOG flexible ligands are oriented on a grid where each ligand is represented by up to 25 low-energy conformations. So, then we can do the ligand; we can keep the ligand flexible and the protein rigid, which is the most commonly used method because of its high throughput and scalability. So, where can we, you know, keep the protein structure rigid or the target structure rigid; however, we keep the ligand flexible? Some of the tools that use this approach are GOLD, AUTODOCK, GLIDE, and there are methods where we can keep both the ligand and the protein flexible. But, in this case, if we are keeping the protein flexible, it means the binding pocket contains several residues.

And then the side chains of those residues will have a lot of degrees of freedom, which will take a lot of computational power. So, then we have computationally intensive problem to solve. So, it will take a lot of time, and that is why this method is usually avoided; usually, people go for the flexible ligand and rigid protein molecular docking method. And then there are induced fit methods; for example, in GLIDE or MOE, you can have this induced fit where, once your ligand is docked into the binding pocket of a protein or receptor, the receptor binding pocket residues undergo a conformational change to adapt to the ligand conformation. And then, in Surf-flex, Gold, or Autodock, you can have the possibility to keep the side chains of these amino acid residues in the binding pocket flexible, leading to a flexible protein and flexible ligand.

So, if you look at the steps, what are the steps that are required in molecular docking? So, first things first, we need to have the target structure and the ligand structure. So, the target structure we obtained again from either the PDB or Protein Data Bank. Or if it is not available as an experimentally resolved 3D structure, we can use the computationally predicted model from AlphaFold or RosettaFold. Or we can do it using homology modeling as well. And then we obtain the ligands from the zinc, which is a large database we have seen in earlier sessions, or Asinex or Vitas-M, and then we have to optimize the structure of those small molecules.

We have to prepare the target structure as well, which we will discuss soon. So, how do we optimize the structure? We try to minimize the energy, take care of the protonation states, etcetera, and then we prepare those ligands. So, once the target and ligand structures are prepared, the next step is to identify the binding site. So, now we can either go for global docking or go for local docking. And in global docking, we wanted to explore all the possible binding pockets of the ligand against the target.

But again, that is computationally very intensive because it has to search all the space around the target. But if we know, for example, that the target has an active site or catalytic site where we want to dock our ligand. So, then we can define that binding pocket using either a co-crystal structure that is available in the 3D structure, or we can define the binding pocket based on the amino acid residues that are involved in binding. So, once the binding site is identified, Now, we let the algorithm perform the docking to explore the binding of those ligands into that specified binding program. So, then do we perform the rigid or flexible docking based on what method we are going to use, and then we generate the multiple ligand poses.

So, once we generate those poses, which is the first step, the second step is the scoring and ranking of those poses. So, in the scoring, we evaluate those docking poses using scoring functions, and then you have a lot of different kinds of scoring functions. Some examples are glide score and auto dock scoring. And, once we get the score, we will sort those molecules on the basis of the scoring function, and then we will rank those molecules. And, after doing the docking, we can do the post-docking analysis where we can visually analyze the interactions.

It can be done manually or automatically, as well as through hydrogen bond interactions or hydrophobic contacts, which describe how the ligands are interacting with the residues in the binding. And then we can refine those results using molecular dynamics simulations or free energy calculations. So, let us take a look at how we prepare the protein structure. So, well, it depends on the docking program that is being used, but usually when we select a structure. So, we have to take care of the resolution.

So, the resolution of those structures is experimentally determined using X-ray crystallography. So, it should be 2 points less than 2.5 angstroms because the resolution of the crystal structure is directly related to the level at which we can see those atoms or, you know, functional groups or rings in the structure. So, if the resolution is more than 2.

5, meaning if it is 3.5 or 4.5, those structures are usually not very well resolved, and we cannot see all those side chains, so we cannot use them. So, we better keep the structures; we use the structures that are less than 2.5 angstroms resolution. And then we have to take care of the species as well, like the structures we are using if the target is human. If we are developing a drug for humans, then we have to identify and select the protein structure from humans only.

And then we select the site. So, that site may contain cofactors or metal ions. Based on our approach, we need to take care of them. So, either we keep them or we remove them; well, it depends upon the strategy that we are actually using. So, for example, if your enzyme is a metalloenzyme, then the metal ions are important for activity, and they are also involved in, you know, inhibition through another ligand. So, we have to keep those metal ions inside the binding pocket, and then we have to perform the molecular docking.

And then you have to add the charge depending on the pH. So, we need to determine the pH of the relevant assay that we are using, and then, based on that pH, we determine the charge on all those protein residues and then use them in that way. And then often we have to add hydrogens to the protonation of all those residues because usually in X-ray crystallography the structures do not contain hydrogens. Or even if they contain hydrogen, then they might be improperly placed. So, we need to carefully add hydrogens to all of those residues.

And then you have to remove those water molecules that are not involved in binding and not playing an important role in the binding of the ligand to the pocket. And then we remove the cofactors or metals again based on our research problem. And then we do a pre-docking refinement, and we also need to take care of the missing residues or atoms. So, this is very important. There are some missing residues in the protein structure because sometimes the residues that are very flexible are not visible in the X-ray crystallography.

So, there you will actually see a gap. So, mainly, it is related to the loops. So, loops are quite flexible. So, you cannot see them in the X-ray. So, in that case, we have to check if those missing residues are lying somewhere in the binding pocket or if your ligand might interact with them. So, we need to take care of them, and we need to build those and add those missing residues or atoms.

And then, you know, again, the water molecules that are not in the binding pocket need to be removed. So, normally we remove all the water molecules except when they play an important role in coordinating with the ligand; if they are forming some water bridges, etc., then we can keep them. If there is any missing hydrogen that we have already talked about, it is missing. So, ideally, we shall remove all the hydrogens and add hydrogens to the whole structure.

So, many docking programs require the protein to have explicit hydrogen; in general, these can be added unambiguously. An incorrect assignment of protonation state in the active site will give poor results because it will affect your binding affinity calculation. So, once we have prepared the protein, the next step is to prepare the small molecule or the ligand. So, we need to clean up the 2D structures by removing any counter ions, salts, or water molecules that might be present if the structure was determined by crystallography. So, then we have to also remove all reactive or otherwise undesirable compounds.

Sometimes, for example, you have those pan-assay interference compounds. So, some of those molecules contain functional groups that are interfering with almost every in vitro screening assay. So, we need to filter out those functional groups; we need to filter out those molecules that contain those functional groups. And then we need to possibly generate all the optical isomers, such as enantiomers, if there is a chiral carbon atom in your ligand. So, you have to generate all these isomers: stereoisomers or enantiomers, cis-trans isomers, tautomers, and the protonation state of the structures.

So, maybe one molecule that contains several chiral carbon atoms might have 20 to 50 different enantiomers, as well as protonation states or tautomers, based on the structure. So, the incorrect tautomer assignment may lead to docking errors, and that will actually be a problem. So, we need to assign the tautomers correctly as well. And then finally, we need to minimize the structures, and generally, we use a molecular mechanics' force field for that. So, the idea of generating the minimal energy structure is that it is considered when it binds to the receptor.

So, it is the least energetic conformation that is close to the bioactive conformation. So, that is why we minimize the structure and then use that structure for performing molecular docking. So, if we compare several programs, for example, we have those based on the search algorithm or the search algorithm being used in molecular docking programs. So, you can see that AutoDoc uses the Lamarckian genetic algorithm, Dock is using shape matching, and EUDOCK is also using shape matching. And then you have, for example, the ICM, which is using Monte Carlo minimization; you have ZDoc, which is using shape matching.

So, these are different search algorithms which generates the different poses of the molecules which we are docking. And then here on the right-hand side, we can see that these are the scoring functions, for example, DOCK, AutoDock, and Gold; they use force field-based scoring functions. And then we have Flex, Glide, Ludi, and ICM; they use the empirical scoring functions, and then we have other tools like IT Score, PMF, Drug Score, and Smog. So, they use the knowledge-based scoring functions.

So, now we have seen what molecular docking is and how it works. So, let us have a look at the limitations of this classical or traditional docking. So, one thing is that these scoring functions and optimization or search algorithms play a crucial role in molecular docking. So, all those scoring functions and search algorithm they have some limitations on their own and that is together that is one of the limitation of these traditional docking methods. And then the effective docking requires both accurate scoring function computation and efficient search strategies, and there are challenges in balancing accuracy and efficiency. So, if we are increasing the accuracy, then those are not very efficient, and they are not very, you know, scalable.

And then we cannot use them for screening ultra-large libraries, like billions of compounds. So, achieving high accuracy while maintaining computational costs is difficult. And, this challenge limits the ability to screen a large compound library effectively. And another major limitation of these, you know, traditional docking methods is that they neglect ligand and protein flexibility, which is a crucial aspect of molecular interaction. Especially the protein flexibility; because of the computational cost, we are not usually taking the protein flexibility into account, and that leads to, you know, reduced efficiency, I would say.

So, let us see the AI techniques that are used in docking. So, we use deep learning where we use neural nets, for example, convolutional neural nets, which extract the spatial features from protein-ligand interactions. We can also use the transformers that analyze complex relationships and improve docking pose accuracy. Tools such as DeepDock, AlphaFold, and Multimer are being used for structure-based docking. And then we also use the GNN, graph neural nets, which treat molecules as graphs with atoms as nodes and bonds as edges. So, it captures non-covalent interactions, like hydrogen bonds and pi stacking, for better predictions.

Some of the tools, like GraphDock or GNINA, which are used for enhanced docking in binding affinity scoring, are using these GNNs. And then we have reinforcement learning, which learns docking strategies by trial and error, improving over time. So, it is based on reward-based optimization to find low energy binding confirmation and we have tools such as deep bind RL or auto dock RL for improving binding accuracy and efficiency. So, there are multiple tools being developed nowadays that use all those AI-based methodologies to improve the docking and scoring. So, why do we need AI in molecular docking now? We have seen the limitations of traditional molecular docking, which are slow computation and limited scoring accuracy, as well as the failure to model solvent effects, entropy, and complex interactions.

and inefficient virtual screening, which takes weeks to screen a large compound library, and of course, biased scoring functions, which are usually empirical. So, the AI driven docking it can enhance speed, accuracy and flexibility compared to the traditional methods. So, if we look at the role of AI in molecular docking. So, we can again use AI for, you know, docking acceleration, where we are increasing the speed and efficiency. Or we can use the AI-based scoring functions, which will improve the binding affinity prediction.

Or we can use the generative AI-based docking, where we can have flexible and novel docking methods that use the generative models. And then we can also have AI-powered docking frameworks that are based on, you know, learning-based docking approaches. So, I would say that these are four important aspects where AI can contribute to molecular docking. And we will see examples of all these aspects now.

So, if we talk about the docking acceleration with AI. So, this is one of the first approaches in deep docking, which is a DL-based platform. So, it utilizes a small fraction of a chemical database along with its docking score to train a neural network in an AI loop. That is then used to predict the scores of the remaining molecules, resulting in a 100-fold increase in speed compared to the conventional docking methods. So, if I wanted to screen a library of, for example, 1 billion compounds. So, for that, I probably need, you know, thousands of nodes of a supercomputer, and that too for a couple of months actually to execute the docking of 1 billion compounds.

But with the help of tools such as deep docking, we can accelerate that process by using, you know, QSAR models, which can predict the docking score, and you can speed it up by at least 100-fold. So, this protocol uses binary classifiers in the form of a feedforward deep neural net trained on 1024-bit circular Morgan fingerprints. And then the fingerprint bits are used as features in the DNN model, which aims to learn which substructures are responsible for high predicted binding affinity in terms of the docking score. So, we take a small subsample, and then we use that sample to calculate the fingerprints and determine the docking score by docking those molecules into the binding pocket. And then we create a model between the docking score and the fingerprints, and then we use that validated model to predict the docking score of the remaining molecules.

So, by using this method, we can screen billions of molecules very easily. So, then we have AI-based docking scoring functions. So, these ML-based or AI-based scoring functions are a group of methods that use ML technologies to learn the functional form of binding affinity by associating patterns in the training data. So, earlier those SVM RF or GBT were used to improve the scoring performance. And nowadays, DL models are being applied to predict the protein-ligand scoring function development, protein-ligand interactions, or docking scores.

So, the NN score of 1.0 is employed as a simple neural net composed of only one hidden layer with 5 neurons to classify the active and inactive compounds based on 194 features, including both interaction and ligand-dependent terms. Comparatively, NN score 2.0, which included many more interaction terms, estimated the pKD rather than inactive or active classification. So, GNINA is one of the examples where it is using, you know, convolutional neural network types of deep learning to predict the docking score. So, it uses a combination of physics-based conformational search, which is, you know, taken from AutoDock Vina, with a 3D convolutional neural net to improve docking post-prediction and binding affinity scoring.

So, it supports rigid and flexible docking both significantly faster on GPU hardware making it efficient for large scale docking studies. So, what it does is that you have the receptor structure, you have the ligand structure, and using the exhaustiveness, it generates the poses. And then this rescoring is done using Monte Carlo simulation, followed by fast refinement and the Metropolis acceptance criterion. So, you have multiple options: you can do the full refinement, you can rescore and sort, or you can just do the RMSD filtering, and then you get the molecules, the docked pose, along with the CNN scores as well.

And then there is DeepBindRG, which is an AI-based binding affinity prediction tool. So, DeepBindRG is a deep learning model that predicts protein-ligand binding affinities. So, it uses residue ligand contact information to capture complex interaction patterns, and it outperforms traditional scoring functions like AutoDock Vina in affinity prediction, as it demonstrates lower RMSD and a higher correlation coefficient in benchmarks as well. So overall, it enhances drug discovery by improving accuracy in protein-ligand binding analysis. So, if we talk about the methodology of DeepBindRG. So, it works by using a dataset of protein-ligand complexes where it decodes the interface spatial information in 2D format, and then it uses the ResNet and CNN models by performing training, validation, and testing.

So, this model has been externally evaluated on that directory of decoys from the DuD.E web server. Then, it has been tested on the CASF benchmarking tool as well. And then it utilizes GNN to predict binding affinities by learning from protein-ligand interaction patterns. It has been trained on experimental binding data, improving affinity prediction and accuracy compared to traditional scoring functions.

So, then we have those generative AI models for docking. So, DiffDock is one of the generative models that has been used extensively right now, I would say. So, it is a blind docking method that uses utilizes diffusion generative models. So, fundamentally, these diffusion models work by destroying the training data through the successive addition of Gaussian noise and then learning to recover the data by reversing this noising process. So, after training, the diffusion models can be used to generate data by simply passing randomly sampled noise through the learned denoising process. So, it DiffDock achieves a remarkably high success rate to predict the experimental protein ligand binding poses from decoys significantly outperforming other docking methodologies like gnina and glide on blind docking task.

And then we can use that consensus scoring as well. So, it is a strategy that combines results from multiple docking methods or scoring functions to improve the accuracy and reliability of predicted binding affinities. So, there are types of consensus scoring approaches, like rank-based consensus, where we average or combine rankings from different docking methods. Score-based consensus where we use the weights or unweighted averages or binding scores from different scoring functions. Or we can use a classifier-based consensus where we employ machine learning models to integrate multiple docking scores for improved predictions. Every docking tool must be benchmarked, and benchmarking is just evaluating whether the method is good or not, and that is usually done by some benchmarking tools.

So, I have summarized here some of the benchmarking tools that are extensively used in evaluating these docking methods. So, the first one is CASF. It is a scoring function performance. So, its focus is on evaluating the performance of the scoring function.

So, the dataset contains protein-ligand complexes along with the binding affinities. So, one can use this dataset to evaluate and determine the binding affinity using docking and compare how those binding affinities are close to the experimentally determined binding affinities that are available in this dataset. So, the evaluation metrics that we can use are docking accuracy, ranking power, and scoring power. And then it has been updated regularly with diverse targets. So, every developer of this server is updating the database regularly.

And then we have the DUDE, which also evaluates the virtual screening performance. So, what it has is that it has the structure of actives along with the property of mass decoys, which means those molecules will be structurally similar to the active molecules, but they will be inactive in the assay. And then it is evaluated by enrichment, false positive, and early recognition, and the special feature of DUDE is that it has enhanced decoy selection for realistic benchmarking. And then we have PDBbind; again, the key focus of PDBbind is on docking and scoring accuracy. And then, the data includes protein-ligand complexes with experimental Kd, Ki, and IC50 values; it is similar to the CASF. And then it also evaluates the docking post-accuracy scoring reliability, and it covers diverse protein families and ligand types.

And then we have post DB, which is where the key focus is docking post-prediction. So, it has diverse protein-ligand pairs. And then it evaluates the docking algorithms on the basis of RMSD (root mean square deviation) for pose accuracy, and its focus is on pose generation rather than scoring. So, it is basically like typically evaluating the first step of docking, which is pose generation. And then we have LEADPEP, which is for peptide docking; it contains peptide-protein complexes where one can evaluate the docking accuracy and binding mode prediction of the docking algorithm.

Therefore, it is a specialized data set for peptide-protein interaction. And then we have this DUDeCOVE, which is specifically for covalent inhibitors docking, where the dataset includes covalent ligand-protein complexes. And then you can use it to evaluate the pose accuracy scoring performance of your molecular docking tool. And then it is tailored for benchmarking specifically covalent docking methods. Then we have DockGen, which is a kind of generative docking algorithm. Where we have AI-generated ligand libraries, you can evaluate the post-prediction chemical novelty and binding affinity, and then it is designed for generative AI-based docking models.

And then we have the AlphaFold docking challenge, where the idea is to dock with flexible target structures, and the data includes AI-predicted protein structures. So, here you can evaluate the post-prediction structure adaptation accuracy, etc. And then the focus is on docking against AI-predicted structures, where you can evaluate how well your model works on those AI-predicted structures, like the target structures coming from AlphaFold. And then, if you see the advantages of AI-assisted docking. So, now, I think we have understood that we can achieve a higher accuracy, which is an improvement in binding pose and affinity prediction.

faster computation, reducing a reduction in docking time from day to minutes, enabling high throughput screening, screening billions of molecules in a short period of time. With better flexibility in handling automated site detection, we can also take into account the induced fit effects, overcoming the limitations of rigid docking. And then we can have the improved scoring, where AI learns from experimental data, enhancing binding affinity predictions. So, how can we address the AI shortcomings in molecular docking? So, we can have the AI and integrate it with the experiment.

So, we can use a feedback loop to refine docking predictions and enhance the success rate. And then we can have active learning; we can continuously update AI models with experimental feedback, and we can use high-quality data because, in computational modeling, garbage in, garbage out. So, if you have good quality data in your hand, then only can you expect that your outcome will be useful. So, we can train the AI with diverse experimentally validated datasets, and then we can take into account the solvent effects. where we incorporate water mediated interactions for accurate binding predictions and we can keep into account the flexibility. We can use ensemble docking and MD to account for both protein-like dynamics, and we can improve the scoring by integrating machine learning with physics-based scoring functions.

So, if you look at the conclusion of this session, So, the AI driven models they improve docking precision and binding affinity prediction and AI accelerates large scale virtual screening for drug discovery. In combining AI with MD and QM, it can enhance reliability, and without any doubt, AI is transforming molecular docking; however, continuous improvement and experimental validations are essential for real-world impact. So, these are some of the papers that you can go through to learn more about how AI is being used in molecular docking. And then in the end, I have an open question for you: with AI rapidly improving molecular docking accuracy, how can we balance its predictive power with the need for experimental validation to ensure real-world reliability? And with that, thank you.