

## EDGE ARTICLE

[View Article Online](#)  
[View Journal](#) | [View Issue](#)



Cite this: *Chem. Sci.*, 2025, **16**, 19876

All publication charges for this article have been paid for by the Royal Society of Chemistry

Received 20th November 2024  
Accepted 22nd September 2025

DOI: 10.1039/d4sc07875f  
[rsc.li/chemical-science](http://rsc.li/chemical-science)

## Pose ensemble graph neural networks to improve docking performances

Thanawat Thaingtamtanha,<sup>†,a</sup> Jordane Preto<sup>†,b</sup> and Francesco Gentile  <sup>†,ac</sup>

Predicting the geometry and strength governing small molecule–protein interactions remains a paramount challenge in drug discovery due to their complex and dynamic nature. Several machine learning (ML) methods have been proposed to complement and improve on physics-based tools such as molecular docking, usually by mapping three dimensional features of poses to their closeness to experimental structures and/or to binding affinities. Here, we introduce Dockbox2 (DBX2), a novel approach that encodes ensembles of computational poses within a graph neural network framework *via* energy-based features derived from molecular docking. The model was jointly trained to predict binding pose likelihood as a node-level task and binding affinity as a graph-level task using the PDBbind dataset and demonstrated significant performance in comprehensive, retrospective docking and virtual screening experiments, compared with state-of-the-art physics- and ML-based tools. Our results encourage further exploration of ML models learning from conformational ensembles to accurately model small molecule–protein interactions and thermodynamics. The DBX2 code is available at <https://github.com/jp43/DockBox2>.

## Introduction

Drugs exert their therapeutic effects by binding to specific biomolecular targets, typically proteins or nucleic acids, and modulating their function, thereby inhibiting or restoring processes related to various diseases. The initial step in the drug discovery pipeline involves identifying molecules binding to the target of interest with high affinity and specificity,<sup>1</sup> hence making the accurate prediction of both a crucial aspect for therapeutic development.<sup>2</sup> Binding affinity, which reflects the strength of the interaction between a drug and its protein target, is commonly expressed in terms of dissociation constant ( $K_d$ ), measurable *via* a plethora of experimental techniques.<sup>3</sup> However, these techniques are usually time-consuming and resource intensive,<sup>4,5</sup> especially at high throughput rates required to explore vast chemical spaces.<sup>6</sup> Consequently, *in silico* screening methods have gained significant momentum, especially in the recent years.<sup>7</sup>

Although the estimation of ligand–protein affinities and interactions is essential, significant challenges arise due to the dynamic nature of these complexes. Molecular dynamics (MD) simulations can provide valuable insights into the nature of these interactions, *e.g.*, by considering an ensemble of bound conformations to compute thermodynamically accurate energies.<sup>8</sup> This is usually done by simulating the complexes in their

thermodynamic equilibrium and considering the time spent in the various microstates. Therefore, MD has the potential to connect the chemical world to physical observables, aiding in the determination of state variables (free energy, enthalpy, entropy, ...), kinetics, and the exploration of biomolecular mechanisms driven by rare events.<sup>9</sup> For instance, the ligand Gaussian accelerated MD (LGMD) method, an enhanced sampling technique pioneered by Miao *et al.*,<sup>10</sup> was employed to forecast the binding affinity of nirmatrelvir with the coronavirus 3C-like protease, yielding predictions consistent with experimental observations.<sup>11,12</sup> Likewise, Wolf *et al.*<sup>13</sup> harnessed the power of Langevin simulations an extended MD approach that delves into the intricate low-frequency motions governing large conformational shifts,<sup>14</sup> to estimate the binding affinity of the benzamidine–trypsin complex. However, both standard and biased MD methods require significant computational power that makes these techniques unsuited for high-throughput screening purposes. Consequently, faster and less accurate methods such as molecular docking and machine learning (ML) approaches have been proposed as alternatives.

Molecular docking methods generate bound conformations of a ligand within a rigid binding pocket and then rank the poses using a scoring function, both to identify the most probable pose and to estimate the binding affinity.<sup>15</sup> Despite its simplicity, docking has shown great potential for the identification active molecules from vast backgrounds of inactive compounds,<sup>16,17</sup> with its impact extending across numerous therapeutic areas. Manglik *et al.*, for example, docked over 3 million molecules against the  $\mu$ -opioid receptor ( $\mu$ OR), leading to the discovery of PZM21, a G protein-biased  $\mu$ OR agonist.<sup>18</sup>

<sup>a</sup>Department of Chemistry and Biomolecular Sciences, University of Ottawa, Ottawa, ON K1N 6N5, Canada. E-mail: [fgentile@uottawa.ca](mailto:fgentile@uottawa.ca)

<sup>b</sup>Aix-Marseille University, Université de Toulon, CNRS, Centre de Physique Théorique UMR 7332, 13288 Marseille Cedex 09, France

<sup>c</sup>Ottawa Institute of Systems Biology, Ottawa, ON K1H 8M5, Canada

<sup>†</sup> Equal contribution.



Zernov *et al.* discovered a compound targeting the transient receptor potential cation channel 6 as a potential starting point to develop anti-Alzheimer's therapies, with *in vitro* studies confirming its efficacy, stability, and target specificity without adverse effects.<sup>19</sup> Stein *et al.* employed docking to screen over 150 million molecules targeting melatonin receptor 1 (MT1) in the search for therapeutics addressing sleep disorders and depression, reporting a novel chemotype with experimentally validated, selective MT1 agonist activity.<sup>20</sup> Fink *et al.* utilized large-scale docking to identify novel  $\alpha$ 2A-adrenergic receptor ( $\alpha$ 2AAR) agonists with fewer adverse effects compared to earlier treatments, as new starting points to develop nonopioid analgesics.<sup>21</sup> These and many other studies underscore the important role of docking in advancing drug discovery.

However, several limitations remain in docking, mainly due to the approximative nature of scoring functions and the neglection of flexibility.<sup>15,22</sup> Thus, ML methods have been introduced in the last decade to tackle molecular docking challenges.<sup>15</sup> For example, Graph Neural Networks (GNNs) have been widely explored to characterize ligand–protein interactions.<sup>23</sup> Several models have been proposed, such as CurvAGN,<sup>24</sup> PIGNet,<sup>25</sup> GenScore<sup>26</sup> and SS-GNN,<sup>27</sup> reporting strong correlations between predicted and experimental affinities.<sup>23,28,29</sup> Additionally, GNNs have been applied in generative settings to replace physics-based sampling in generate and scoring ligand–protein poses, such as in DiffDock<sup>30</sup> and MedusaGraph.<sup>31</sup> Although these architectures have shown promising results, an increasing number of studies suggest that GNNs tend to memorize ligand and protein patterns instead of learning the phcial chemistry of the interactions.<sup>23,29</sup> Moreover, these methods generally map single pose graphs to binding affinities, thus neglecting full thermodynamic profile and dynamics of ligand–protein interactions that depends on multiple conformations.<sup>23</sup> Notably, recent efforts have been made to consider multiple conformations in training GNNs for binding affinity predictions, such as Dynaformer, a method that encode each MD-derived binding conformation into a graph within a framework to provide better affinity estimates.<sup>32</sup> Notably, Dynaformer still relies on mapping each conformation to a single affinity value, and requires the use of costly simulations, hence limiting its scalability.

In this work, we introduce DockBox2 (DBX2), a GNNs framework enabling to encode multiple ligand–protein conformations derived from docking within individual graph neural networks in order to leverage ensemble representations for jointly predicting pose likelihood at the node level and binding affinities at the graph level. In a series of retrospective experiments, DBX2 demonstrated significant improved performances both for docking and virtual screening (VS) tasks compared with physics-based and ML methods, warranting further investigation of ensemble-based ML models in computer-aided drug discovery.

## Material and methods

### Datasets

The DBX2 model was trained and evaluated using the PDBbind database.<sup>33</sup> The refined set of PDBbind v2016 (4057 complexes)<sup>34</sup>

was used to train the model. PDBbind is a comprehensive and widely adopted benchmark for protein–ligand binding, and several widely used benchmark datasets, such as CASF-2016,<sup>35</sup> are derived from this refined set. The PDBbind v2019-based hold-out test set built by Volkov *et al.*<sup>29</sup> and the Runs N' Pose database from Škrinjar *et al.*,<sup>36</sup> consisting of 3393 and 2600 complexes respectively, were used as external test sets. Volkov's dataset is curated to mitigate latent biases, such as structural patterns in ligands or proteins, which can favor GNN memorization rather than protein–ligand interaction learning. As highlighted in previous studies,<sup>23,29</sup> this memorization often arises from significant redundancies between training and test sets, resulting in data leakage. The Runs N' Poses dataset is a recently developed dataset containing high-resolution protein–ligand systems released after the publication of PDBbind v2020 and the training date cutoff of several protein–ligand co-folding models (e.g., AlphaFold3,<sup>37</sup> Chai-1,<sup>38</sup> Proteinix,<sup>39</sup> and Boltz-1 (ref. 40)). A subset of the LIT-PCBA database<sup>41</sup> was used to perform retrospective VS experiments.

### Protein and ligand preparation

Complexes from PDBbind were prepared following the same procedure of our previous work.<sup>42</sup> For retrospective VS, dominant protonation and tautomerization states of small molecules were computed from the SMILES using Openeye's QUACPAC<sup>43</sup> and converted into low-energy 3D conformations (mol2 format) using Openeye's OMEGA tool.<sup>43</sup> The target proteins were prepared by removing redundant protein chains, along with non-essential ions, waters, and heteroatoms. The resulting protein structures were prepared using the Molecular Operating Environment (MOE) QuickPrep tool,<sup>44</sup> to automatically add missing loops and assign reasonable conformations to the residues with alternate orientation. Subsequently, protonation states were generated using the Protonate 3D tool from MOE (at pH 7.4). Finally, the structures were energy-minimized using the AMBER10:EHT forcefield implemented in MOE, and saved in pdb format.

### Molecular docking and rescoring

The first Dockbox package (DBX)<sup>42</sup> was utilized to generate binding poses with AutoDock,<sup>45</sup> Vina<sup>46</sup> and DOCK 6 (DOCK),<sup>47</sup> and rescoring with their scoring function in addition to Gnina<sup>48</sup> and DSX.<sup>49</sup> The DBX configuration file used for this purpose on PDBbind v2016 and the test sets is illustrated in Fig. S1; a maximum of 140 binding poses were generated for each system, 60 from AutoDock, 20 from Vina, and 60 from DOCK. For AutoDock, grid spacing was set to 0.3 Å, and the Lamarckian genetic algorithm<sup>50</sup> was employed to generate poses. For Vina, the energy range for final poses was set to 3 kcal mol<sup>-1</sup>. In DOCK, a grid-based scoring method was applied with a spacing of 0.3 Å. All other parameters were left as default. Docking with any of the above programs was followed by energy minimization, starting with 500 steps of the steepest descent method followed by 1000 steps combining steepest descent and conjugate gradient methods. Energy minimization was performed using AmberTools 17 (ref. 51) to prevent structural clashes and



ensure appropriate rescoring with different programs. Rescoring was then conducted with AutoDock, Vina, DOCK, Gnina's CNNScore, and DSX scoring functions.

### Dockbox2 architecture

DBX2 architecture is based on the GraphSAGE model<sup>52</sup> as shown in Fig. 1. The ensemble of poses generated by docking a given ligand–protein pair is used to construct a graph (Fig. 1A), with each node encoding an individual binding pose represented by categorical and energetic features, listed in Table 1.

All available scoring terms provided by Vina<sup>46</sup> were included as node features, with the exception of the entropy term which is determined solely by the ligand structure and therefore remains constant across different poses of the same ligand. In the constructed graph, pairwise root mean square deviation (RMSD) values are calculated between all poses. Two nodes are connected by an edge if the RMSD between the two poses is below a predefined threshold while the RMSD value is kept as edge feature. Graphs were generated using the `create_graphs` script available in the DBX2 package. In the shared layers, the DBX2 model uses the message passing (MP) framework,<sup>53</sup> *i.e.*, for each node  $i$ , information from its neighbors  $j \in \mathcal{N}(i)$  is gathered and aggregated using the symmetric mean (symmean) aggregation:

$$\mathbf{m}_{\mathcal{N}(i)}^{(k-1)} = \text{SYMMEAN} \left\{ \mathbf{s}_j^{(k-1)} \oplus \text{RMSD}_{ij}, \forall j \in \mathcal{N}(i) \right\} \quad (1)$$

where  $\mathbf{m}_{\mathcal{N}(i)}^{(k-1)}$  is the aggregated message for node  $i$  from its neighbors,  $\mathbf{s}_j^{(k-1)}$  is the feature vector of neighbor node  $j$ ,  $\text{RMSD}_{ij}$  is the RMSD between node  $i$  and  $j$ . The feature vector is concatenated with the RMSD between nodes  $i$  and  $j$ . The

aggregation function then combines these concatenated vectors to produce a single aggregation message vector. The node feature vector is then updated:

$$\mathbf{s}_i^{(k)} = \sigma \left( \mathbf{W}_{\text{self}}^{(k)} \mathbf{s}_i^{(k-1)} \oplus \mathbf{W}_{\text{neigh}}^{(k)} \mathbf{m}_{\mathcal{N}(i)}^{(k-1)} \right) \quad (2)$$

where  $\mathbf{s}_i^{(k-1)}$  is the feature vector of node  $i$  at layer  $k$ .  $\mathbf{s}_i^{(k-1)}$  is the feature vector of node  $i$  from the previous layer  $k-1$ .  $\mathbf{W}_{\text{self}}^{(k)}$  and  $\mathbf{W}_{\text{neigh}}^{(k)}$  are learnable weight matrices that apply to the feature vector of the current node and to the aggregated message vector from neighbor nodes, respectively.  $\mathbf{m}_{\mathcal{N}(i)}^{(k-1)}$  is the aggregated message from the neighbors  $\mathcal{N}(i)$  of node  $i$ . The MP layers are followed by multilayer perceptron (MLP) layers to predict pose correctness (node-level task) and the  $\text{pK}_d/\text{pK}_i$  (graph-level task) as illustrated in Fig. 1B. For node-level predictions, aggregated information from the MP layers is passed to an MLP with Rectified Linear Unit (ReLU) and sigmoid activation function for hidden layers and final layer of MLP, respectively. For graph-level predictions, aggregated information is passed to a readout layer corresponding to a MeanMax pooling and then passed to a two-layers MLP, with ReLu activation function for the hidden layer and linear activation function for the output layer.

### Model training and evaluation

The total loss function of DBX2 consists of three components  $\text{Loss}_n$ ,  $\text{Loss}_g$ , and  $\text{Loss}_{\text{reg}}$

$$\text{Total loss} = \text{Loss}_n + w_1 \text{Loss}_g + \text{Loss}_{\text{reg}} \quad (3)$$

$\text{Loss}_n$  is the loss function for node-level task, where the binary focal cross entropy<sup>54</sup> is used as loss function applied to each node in the batch and averaged:

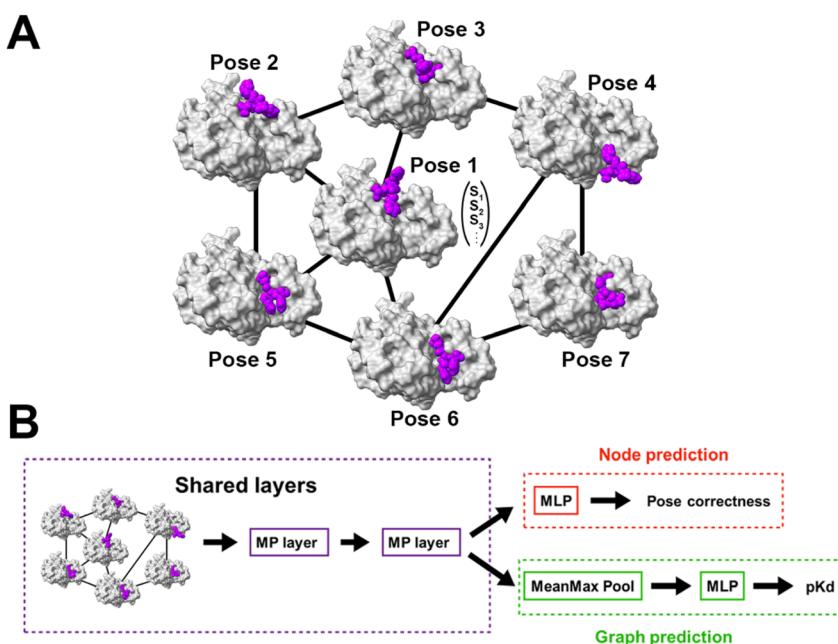


Fig. 1 Architecture of DBX2. (A) Binding poses are represented as nodes. Two pose nodes are connected by an edge based on the root mean square deviation (RMSD) between them. Docking-derived energies and categorical features of each binding pose, here referred as  $s_1, s_2, s_3, \dots$  are used as node features. (B) Schematic of the DBX2 architecture; pose correctness and  $\text{pK}_d$  are jointly learned as node- and graph-level tasks, respectively.



Table 1 DBX2 node features

Features	Description
Instance Score	Docking software utilized to generate the binding pose
Rescoring score (AutoDock, Vina, Dock, DSX, Gnina)	Docking score from original docking program
Gaussian terms (gauss1_inter, gauss2_inter, gauss1_intra, gauss2_intra)	Docking score obtained by rescoring the pose with another scoring function
Gaussian terms (gauss1_inter, gauss2_inter, gauss1_intra, gauss2_intra)	Gaussian terms of the binding pose, as provided by Vina <sup>46</sup>
Hydrophobic interactions (hydrophobic_inter, hydrophobic_intra)	Hydrophobic terms evaluated by Vina <sup>46</sup>
Hydrogen bonding (hydrogenbonding_inter, hydrogenbonding_intra)	Hydrogen bond terms evaluated by Vina <sup>46</sup>
Repulsion (repulsion_inter, repulsion_intra)	Repulsive Lennard-Jones energies from Vina <sup>46</sup>

$$\text{Loss}_n = \frac{1}{N} \sum_{i=1}^N -\alpha_t^{(i)} \cdot (1 - p_t^{(i)})^\gamma \cdot \log(p_t^{(i)}) \quad (4)$$

where  $N$  is the number of nodes in the batch,  $\gamma$  is the focusing parameter (set to 1.0 in this study), and  $\alpha_t^{(i)}$  is the weighting factor for each  $i$ -th sample:

$$\alpha_t^{(i)} = \begin{cases} \alpha & \text{if } y^{(i)} = 1 \\ 1 - \alpha & \text{if } y^{(i)} = 0 \end{cases} \quad (5)$$

where  $\alpha$  is computed as:

$$\alpha = \frac{1}{1 + \frac{1}{G_t} \sum_{i=1}^{G_t} \frac{C_i}{I_i}} \quad (6)$$

where  $G_t$  is the number of graphs in the training set, and  $C_i$  and  $I_i$  are the number of correct poses and incorrect poses in the  $i$ -th graph, respectively. A pose was considered as correct if it was 2 Å or less of RMSD from the experimental one.  $p_t^{(i)}$  is the predicted probability output by the model for the correct class label of each  $i$ -th node:

$$p_t^{(i)} = \begin{cases} p^{(i)} & \text{if } y^{(i)} = 1 \\ 1 - p^{(i)} & \text{if } y^{(i)} = 0 \end{cases} \quad (7)$$

where  $p^{(i)}$  is the model output for each pose.

$\text{Loss}_g$  and  $w_1$  are the loss function for the graph-level task and its weight, respectively. The optimal value of  $w_1$  was determined through hyperparameter optimization (Table S1).  $\text{Loss}_g$  corresponds to the root mean square error (RMSE):<sup>55</sup>

$$\text{Loss}_g = \sqrt{\frac{1}{G} \sum_{i=1}^G (y_i - \hat{y}_i)^2} \quad (8)$$

here  $G$  denotes the number of ligand–protein complexes in the batch,  $y_i$  is the actual value of binding affinity for each complex and  $\hat{y}_i$  is the predicted binding affinity for each ligand–protein complex. Minimizing  $\text{Loss}_g$  contributes to correctly predicting the ligand–protein affinity, in which all poses within a graph are processed through message passing and readout, then used to predict the binding affinity.  $\text{Loss}_{\text{reg}}$  is the regularization loss, while L2 regularization loss<sup>56</sup> was here used to prevent overfitting of model:

$$\text{Loss}_{\text{reg}} = \frac{1}{2} \sum_{i=1}^n \theta_i^2 \quad (9)$$

where  $\theta_i$  represent the model parameter,  $n$  is the number of model parameter. The model was trained using the `traindbx2` routine (example of a configuration file for `traindbx2` in the INI format is provided in Fig. S2). Training was performed with a maximum of 200 epochs and early stopping was used by monitoring the total loss on the validation sets for 3 consecutive epochs. The model was trained with mini-batch gradient descent (batch size of 100) and the adaptive moment estimation (ADAM) optimizer with a learning rate of  $5 \times 10^{-4}$  and a decay rate of 0.99.

Hyperparameter optimization was performed using a grid search, considering the following hyperparameters: RMSD cutoff value to define an edge (RMSD cutoff), number of adjacent nodes to randomly sample for aggregation (nrof-neigh), and graph loss weight ( $w_1$ ), for a total of 30 combinations (Table S1). Training and validation sets were prepared using the `split_train_val_dbx2` routine of the DBX2 package. The generated graphs were split for stratified 5-fold cross-validation, keeping a consistent distribution of protein families across all folds. Node and edge features for each graph were standardized using scikit-learn's StandardScaler.<sup>57</sup> For node-level predictions, success rate, accuracy, and area under the curve (AUC) were used as evaluation metrics. For graph-level predictions, RMSE was used.

### Model testing

Models were compared for docking and scoring tasks with other methods on the hold-out and Runs N' Poses test sets. To evaluate docking power, the success rate was computed as the ratio of top-ranked poses with an RMSD equal or lower than a predefined threshold with respect to the experimental pose. Five different thresholds were tested, 1, 1.5, 2, 2.5 and 3 Å. For DBX2, the success rate was evaluated by considering the top-ranked poses from node-level predictions.

Next, the scoring power was assessed to evaluate the model's ability to predict experimental binding affinities using linear and multiple linear regression. The correlation between experimental binding affinities and scores of the best poses from different scoring functions was analyzed through linear regression, and the  $R^2$  values were calculated. For DBX2, graph-level predictions were utilized to evaluate the correlation with experimental binding affinities. Additionally, multiple linear regression was conducted to correlate experimental binding affinities with predicted values derived from various



combinations of scoring functions, as described in our previous study.<sup>42</sup>

Scoring power was also evaluated using Pearson correlation coefficient and the predictive index (PI) as before.<sup>42</sup> Proposed by Pearlman *et al.*,<sup>58</sup> PI measures the reliability of a scoring function in identifying the most potent binder between two compounds. It is calculated as follows:

$$\text{PI} = \sum_{j>i} \sum_i w_{ij} C_{ij} \quad (10)$$

with

$$w_{ij} = |E_j - E_i| \quad (11)$$

$$C_{ij} = \begin{cases} 1 & \text{if } \frac{E_j - E_i}{S_j - S_i} < 0 \\ -1 & \text{if } \frac{E_j - E_i}{S_j - S_i} > 0 \\ 0 & \text{if } S_j - S_i = 0 \end{cases} \quad (12)$$

where  $E_i$  is the experimental binding affinity of compound  $i$ , and  $S_i$  is the score of compound  $i$ . Predictive index gives values in range from  $-1$  (wrong prediction) to  $1$  (perfect prediction), with  $0$  being random prediction.  $w_{ij}$  is the weighting term which underscores the accurate ranking of compounds exhibiting substantial disparities in experimental binding affinities.

### Retrospective virtual screening

VS experiments were conducted on the three target proteins from the LIT-PCBA database<sup>41</sup> that were not present in the DBX2 training set: Flap structure-specific Endonuclease 1 (FEN1, PDB id: 5FV7),<sup>59</sup> Glucocerebrosidase (GBA, PDB id: 2XWE),<sup>60</sup> and Mammalian Target of Rapamycin Complex 1 (MTORC1, PDB id: 5GPG).<sup>61</sup> Initially, Vina was used to screen active-inactive sets derived from LIT-PCBA against each corresponding structure. The top 20 000 compounds based on the Vina ranking were then docked also with AutoDock to their respective targets. 80 binding poses (60 from AutoDock and 20 from Vina) were generated for each ligand-protein complex (Fig. S3). Rescoring was performed with AutoDock, Vina, DOCK, and Gnina (considering the CNNAffinity of the pose with the highest CNNScore).<sup>48</sup> VS performances were evaluated by computing the logarithmic area under the curve (logAUC),<sup>62</sup> enrichment factors (EF) and Boltzmann-enhanced discrimination of receiver operating characteristic (BEDROC) with adjust parameter ( $\alpha$ ) values of 20 and 80.5 using the CROC Python package.<sup>63–65</sup>

The logAUC quantifies the performance of a VS method by assessing its ability to distinguish active compounds from decoys across the ranked list. By applying a logarithmic scale to the false positive rate axis, it places greater emphasis on the early retrieval of active compounds, which is critical in VS.

EF measures how effectively a VS method identifies active compounds within a specific fraction of the ranked list.<sup>66</sup> EF at a given cutoff ( $x$ ) is calculated from the ratio of true active compounds in the top  $x$  ranked compounds in relation to the ratio of true active compounds in the entire dataset:

$$\text{EF}(x) = \frac{\text{TP}/(\text{TP} + \text{FP})}{[(\text{TP} + \text{FN})/(\text{TP} + \text{TN} + \text{FP} + \text{FN})]} = \frac{N \times n_s}{n \times N_s} \quad (13)$$

where TP and TN are true positives and true negatives, FP and FN are false positives and false negatives.  $N$  is a total number of compounds in the entire dataset,  $N_s$  is a total number of predicted active compounds in the selection set ( $x$ ),  $n$  is a total number of true active compounds in the entire dataset,  $n_s$  is the number of true active compounds in the selection set ( $x$ ). EF was computed by considering the top 2% of the ranked compounds for each scoring functions and for both graph-level and node-level predictions in DBX2 (EF2).

Normalized enrichment factor (NEF) rescales EF values into a range from 0 (bad prediction) to 1 (perfect prediction),<sup>67</sup> with the goal of standardizing comparison across different datasets. NEF is calculated as follow:

$$\text{NEF}(x) = \frac{\text{EF}(x)}{\text{EF}(x)_{\max}} \quad (14)$$

with

$$\text{EF}(x)_{\max} = \frac{\min\{n_s, N \times x\}}{n \times x} \quad (15)$$

where  $\text{EF}(x)_{\max}$  denotes the maximum enrichment factor achievable within a selection set ( $x$ ).  $n_s$  is the number of true active compounds in the selection set ( $x$ ),  $N$  is the number of compounds in the entire dataset.

BEDROC metric emphasizes the concentration of active compounds at several range of ranked data sets<sup>64,67</sup> through a scaling function ( $\alpha$ ). This metric is defined as:

$$\text{BEDROC} = \frac{\text{RIE} - \text{RIE}_{\min}}{\text{RIE}_{\min} - \text{RIE}_{\max}} \quad (16)$$

with

$$\text{RIE}_{\min} = \frac{1 - e^{\alpha R_{\alpha}}}{R_{\alpha}(1 - e^{\alpha})} \quad (17)$$

$$\text{RIE}_{\max} = \frac{1 - e^{-\alpha R_{\alpha}}}{R_{\alpha}(1 - e^{-\alpha})} \quad (18)$$

$$\text{RIE} = \frac{\frac{1}{n} \sum_{i=1}^n e^{\alpha x_i}}{\frac{1}{n} \left( \frac{1 - e^{\alpha}}{e^{\alpha} / N_{-1}} \right)} \quad (19)$$

the robust initial enhancement proposed by Sheridan *et al.*,<sup>68</sup>  $x_i$  is a relative ranking of active compound  $i$ .  $R_{\alpha}$  is the fraction of active compound ( $R_{\alpha} = \frac{n}{N}$ ),  $\alpha$  is the scaling function.

We also investigated the potential of DBX2 to improve VS performance of individual docking programs (rather than on pose pools deriving from different software) and by using different docking setups to generate poses. The top 20 000 LIT-PCBA compounds docked and scored with Vina against FEN1 were redocked with AutoDock and Vina using several combinations of docking parameters for each program (Table S2). The



resulting poses were subsequently subjected to DBX2, using the same settings used in the retrospective VS experiments. The same metrics were calculated to assess the effectiveness of DBX2 in this specific scenario.

### Baseline models

We compared DBX2 model with other methods, including docking and rescoring tools either physics- or ML-based, using the following protocol:

- AutoDock, Vina, DOCK, Gnina, KarmaDock,<sup>69</sup> RTMscore<sup>70</sup> and DBX2 were compared both in terms of docking and scoring power, as well as for retrospective VS (CarsiDock was excluded from the VS experiments due to the computational cost).
- CarsiDock,<sup>71</sup> DSX and DBX2 were compared for docking and scoring power.

Default settings were used for all programs. To evaluate the docking, scoring and VS capabilities of RTMscore and Gnina on the hold-out and Runs N' Poses sets, the binding poses used in DBX2 were also utilized for rescoring with these tools. For Gnina, the success rate was evaluated using CNNScore, and the scoring power was evaluated using the CNNAffinity and Minimized Affinity scores of the pose with the best CNNscore for each system. KarmaDock and CarsiDock, both generative models, automatically generated their own protein-ligand poses and associated scores.

## Results and discussion

### Hyperparameter optimization

The results of hyperparameter optimization for the DBX2 model are summarized in Table S3. The best performing set of hyperparameters included a RMSD cutoff of 10 Å to define edges, a nrof-neigh of 30, and a graph-level loss weight ( $w_1$ ) of 0.02, yielding an average success rate of 60% on 5-fold cross validation. The model with the highest performance was then retained and used in subsequent testing.

### Docking and scoring power

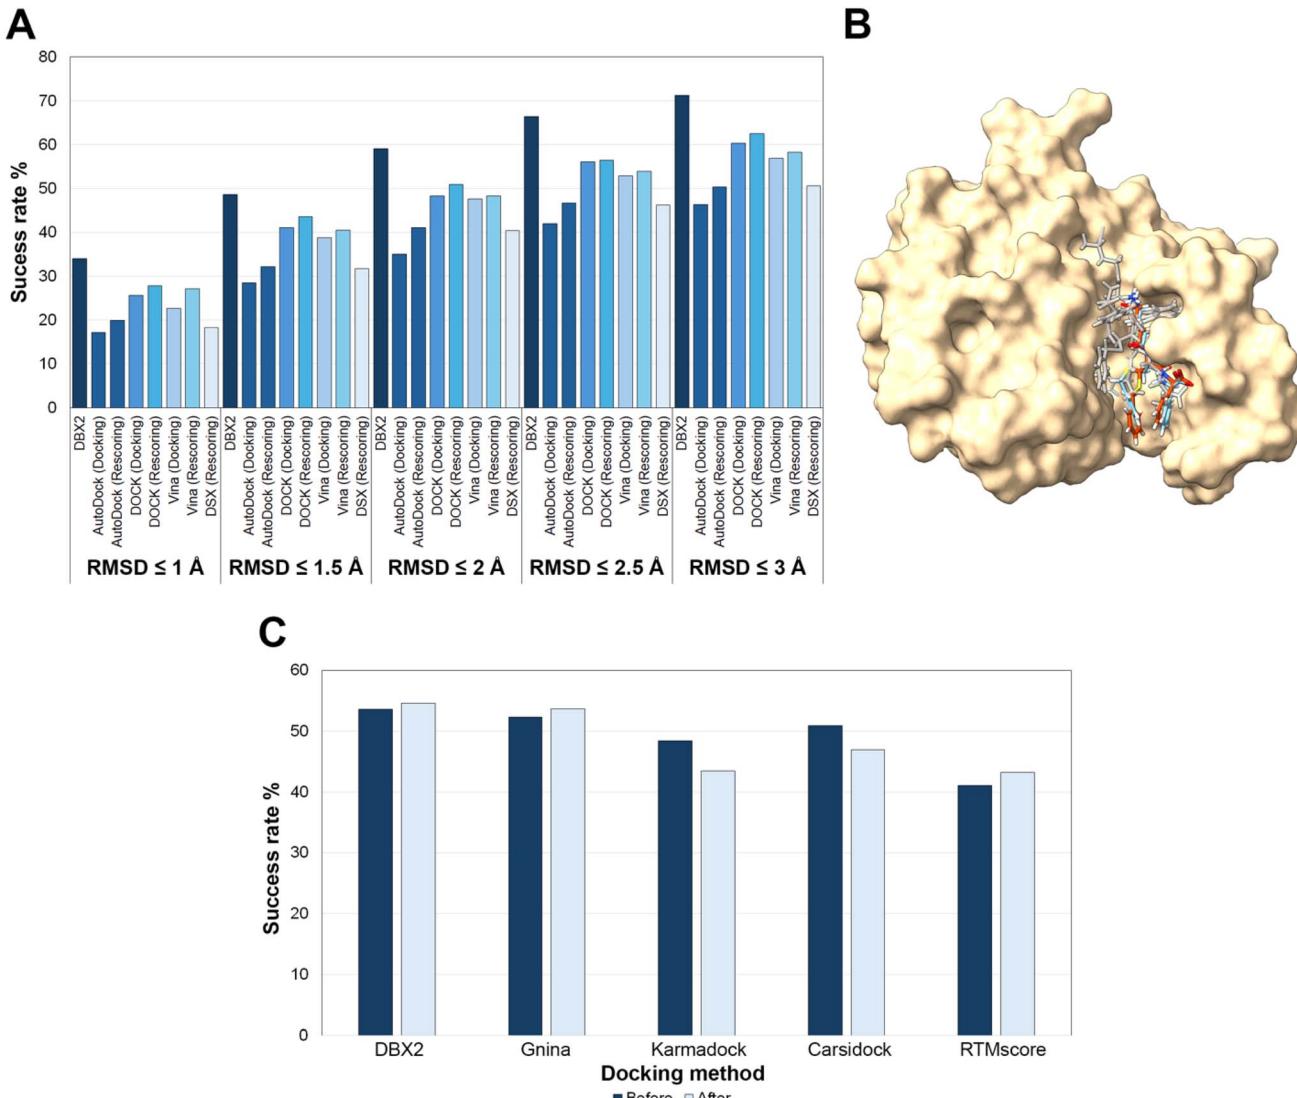
We compared the success rate of DBX2 and other physics-based methods for the docking and rescoring tasks on the hold-out test set, as described in the Material and methods section (Fig. 2A). As expected, rescoring ensembles of docking poses with different scoring functions led to significantly improved performance due to enhanced pose sampling, as observed in previous studies.<sup>42</sup> Noticeably, the node-level pose classification method implemented in DBX2 significantly outperformed all docking and rescoring schemes at all the tested RMSD thresholds. These findings suggest that by leveraging neighbor information *via* the GNN framework, DBX2 offers a significant advantage in accurately identifying native near-to-native ligand binding poses compared with docking methods that score each pose independently. Fig. 2B illustrates an example of successful application of DBX2 for identifying the native pose of the potent TER-117 inhibitor bound to its target, the human glutathione S-transferase P1-1 (PDB id: 10gs).<sup>72</sup> Additionally, we compared

DBX2 against four ML-based docking methods, Gnina, KarmaDock, CarsiDock, and RTMscore, using a 2 Å cutoff on the hold-out dataset (Fig. S4A) and Runs N's Poses dataset (Fig. 2C). Unsurprisingly, KarmaDock, CarsiDock, and RTMscore outperformed both DBX2 and Gnina on the PDBbind v2019-based hold-out test set, which was part of the PDBbind v2020 general set used to train these models.<sup>69–71</sup> Nevertheless, DBX2 displayed encouraging performance despite the limited size of the training set (4057 complexes) compared with the other methods. Next, we performed the same comparison on the Runs N's Poses dataset, which was completely unseen by all five investigated methods during training. Moreover, we investigated the performance of docking before and after removing the Runs N's Poses protein families that overlapped with v2016 and v2020. Notably, DBX2 demonstrated superior performance compared to all other models on the Runs N's Poses dataset, followed by Gnina, both before and after the removal of overlapping protein families (Fig. 2C). Interestingly, upon overlap removal, the success rates for RTMscore, DBX2, and Gnina experienced a slight increase. In contrast, the success rates for KarmaDock and CarsiDock slightly declined. Moreover, the impact of node count per graph on DBX2 prediction performance was further examined by generating additional graphs from the PDBbind v2016 and the hold-out set with reduced node counts: 70 nodes (30 poses from AutoDock and DOCK, 10 from Vina) and 35 nodes (15 poses from AutoDock and DOCK, 5 from Vina). For each setting, the model was retrained and reevaluated on the hold-out test set. The success rate was then compared to the default 140-node configuration. While DBX2 achieved its highest performance with the default setting, the prediction accuracy did not decline dramatically with fewer nodes (Fig. S4B). These results suggest that when generating or training with a large number of poses is challenging, DBX2 can still achieve reasonable performance using ensembles of limited size.

Next, we evaluated the ability of the scoring functions to reproduce experimentally determined binding affinities in the hold-out test set (Table 2). Notably, DBX2 directly computes the binding affinity from an ensemble of poses, so it does not require selecting a specific docking pose as input, unlike other scoring functions. Thus, since DOCK showed the best success rate among classical docking programs, we focused only on poses with the best DOCK scores (after rescoring) in order to compute binding affinities, similarly to our previous work.<sup>42</sup> Thus, linear regression was performed to compare binding affinities from the hold-out dataset with the scores of the best DOCK poses using different scoring functions and their linear combinations.<sup>42</sup> For DBX2, the affinity values for each protein-ligand complex in the hold-out dataset were predicted as graph-level tasks, hence as readouts of pose ensembles *via* docking rather than relying on a single pose.

Interestingly, DBX2 exhibited the highest correlation with experimental binding affinities on the hold-out dataset, outperforming other tested scoring functions. In contrast, DOCK, despite showing the best prediction of binding poses, had the





**Fig. 2** (A) Comparison of success rates of identification of the correct pose on hold-out test set between AutoDock, DOCK, Vina, DSX, and DBX2, comparing docking and rescoring strategies. Rescoring improved the performance of each docking program compared to standard docking, emphasizing the advantage of refining initial pose predictions by evaluating them with additional scoring functions. DBX2 node-level classification outperformed all the other tested methods. (B) Crystal structure of human glutathione S-transferase (PDB id: 10gs) with bound TER117 inhibitor (cyan). The binding pose predicted by DBX2 (orange) aligns closely with the crystallographic structure, in contrast to the poses predicted as native by other docking software (grey). (C) Success rate of identification of the pose correctness on Runs N's Poses dataset before (light) and after (dark) removing overlapping protein families with PDBbind v2020 for DBX2, Gnina, KarmaDock, CarsiDock, and RTMscore.

lowest correlation ( $R^2 = 0.16$ ). DBX2 scoring function also displayed a significantly higher predictive index (0.79) than other methods, indicating its potential suitability in ranking active molecules based on their binding affinities to a target of interest. Likewise, the Pearson coefficient of DBX2 (0.61) indicated a good predictive power based on pharmaceutical industry standards.<sup>73</sup> Nevertheless, the  $R^2$  value, while indicating positive correlation as well as an improvement compared with other methods, remained low (0.38), underscoring remaining challenges in accurate thermodynamics predictions *via* docking-based sampling. Indeed, while our results suggest that docking poses ensembles appear to be more suitable than single poses for binding affinity predictions, they likely fail to

provide a comprehensive thermodynamic picture of binding processes, due to the approximations necessary to ensure the high throughput required in docking. Additionally, DBX2 also outperforms other ML models (KarmaDock, CarsiDock, and RTMscore) in this task, despite being trained on fewer protein-ligand complexes, highlighting the challenges that these methods may face in VS due to the neglection of experimental affinities in their training.<sup>69,71</sup> Correlation plots between experimental and computational affinities are shown in Fig. S5.

Moreover, the DBX2 scoring power on the hold-out set was compared with established methods that were trained and tested on the same splits or supersets of them. Thus, DBX2 was compared with GNN-MP neural network (MPNN) models from



**Table 2**  $R^2$ , Pearson correlation coefficients and predictive index values between experimental binding affinities and the scores provided by tested scoring functions. Best values are indicated in bold

Number of functions	Scoring function/combination	$R^2$	Pearson coefficient	Predictive index
1	DBX2	<b>0.38</b>	<b>0.61</b>	<b>0.79</b>
1	AutoDock	0.20	0.45	0.45
1	DOCK	0.16	0.41	0.42
1	Vina	0.25	0.52	0.48
1	DSX	0.22	0.47	0.46
1	KarmaDock	0.03	0.18	-0.79
1	CarsiDock	0.03	0.17	-0.68
1	RTMscore	0.22	0.46	-0.36
1	Gnina CNNAffinity	0.36	<b>0.61</b>	0.55
1	Gnina MinimizedAffinity	0.25	0.44	0.18
2	AutoDock, Vina	0.25	0.50	0.49
3	AutoDock, Vina, DOCK	0.18	0.44	0.43
3	AutoDock, Vina, DSX	0.23	0.49	0.48
4	AutoDock, Vina, DSX, DOCK	0.22	0.47	0.47

Volkov *et al.*<sup>29</sup> and Pafnucy model from Stepniewska-Dziubinska *et al.*<sup>74</sup> The first class of models are GNNs mapping protein- (P), ligand- (L) and protein–ligand interactions (I) graph representations to ligand–protein affinities. The Pafnucy model is a convolutional neural network utilizing 3D convolution to produce a feature map for protein and ligand atoms to predict ligand–protein affinity. Notably, these models were already trained and tested on the same datasets used in DBX2 (PDBbind v2016 dataset and the hold-out test set, respectively) as previously reported.<sup>29</sup> The comparison of Pearson coefficient and RMSE is summarized in Table S4. Even in this case, DBX2 exhibited significantly improved performances in predicting binding affinity against hold-out set with respect to GNN-MPNN pure interaction (I) models from Volkov *et al.*<sup>29</sup> and Pafnucy model,<sup>74</sup> as evident from the Pearson coefficient and RMSE values, and comparable performances with GNN models that included protein and ligand structural information explicitly, while being based entirely on energetic representations without taking into account any structural information. This observation suggests that DBX2 could (at least partially) overcome the hidden biases causing memorization of 2D molecular patterns that these models display, as described in the study by Volkov *et al.*,<sup>29</sup> while significantly outperforming the success rate of pure interaction models.

### Retrospective virtual screening

To test the VS power of DBX2 in realistic scenarios, we focused on the three LIT-PCBA targets that were not present in our training set: FEN1, GBA, and MTORC1. LIT-PCBA is a small molecule bioactivity dataset to mitigate biases and avoid overestimating VS performances. Derived from bioassays, it mimics experimental active and potency distributions within screening libraries, spans diverse protein targets, and has been validated across multiple screening methods, making it suitable for both structure- and ligand-based VS retrospective experiments.<sup>41</sup> The numbers of active and inactive compounds for each LIT-PCBA

protein target at the beginning of the retrospective VS experiment and after the first round of Vina docking (with the top 20 000 molecules brought forward) are reported in Table S5.

After generating additional poses with AutoDock for molecules endowed by the Vina docking step, rescoring with different scoring functions (including DBX2) was performed and the result evaluated by computing top-100 hit rate, EF2, and NEF (Fig. 3A–C). DBX2 demonstrated superior performance across all metrics when compared to other scoring functions, on the three target proteins. Surprisingly, DBX2's node-level predictions, which assess the likelihood of each binding pose to be the correct one within a specific graph, consistently matched the screening power of graph-level predictions of binding affinities. Gnina, a ML-based tool that recently demonstrated state-of-the-art performance in prospective drug discovery challenges,<sup>75,76</sup> and the other ML-based tools (KarmaDock and RTMscore) also performed well, further validating the potential of data-driven models in VS tasks. Additionally, logAUC (Fig. 3D–F) and BEDROC (Table S6) were calculated to further assess each scoring functions' ability to distinguish between active and inactive compounds. DBX2 demonstrates superior performance across both these metrics as well, suggesting a robust efficacy in prioritizing active compounds throughout top and broad ranks of compounds. Node-level predictions showed the highest performance, followed by graph-level predictions, KarmaDock, CarsiDock, and Gnina's CNNAffinity scoring function.

Lastly, since the use of multiple programs may result computationally expensive in large-scale screens, we investigated the effect of DBX2 in enhancing the VS performance of single docking programs, focusing on FEN1 as the target, we used DBX2 to rescore the top 20 000 Vina-scored molecules from LIT-PCBA, computing top-100 hit rate, EF2 and NEF metrics as well as logAUC before and after the application of DBX2 (Fig. S6 and S7). The results clearly indicated that also in this case, DBX2 significantly improved upon both AutoDock and Vina outcomes across different sets of docking parameters.



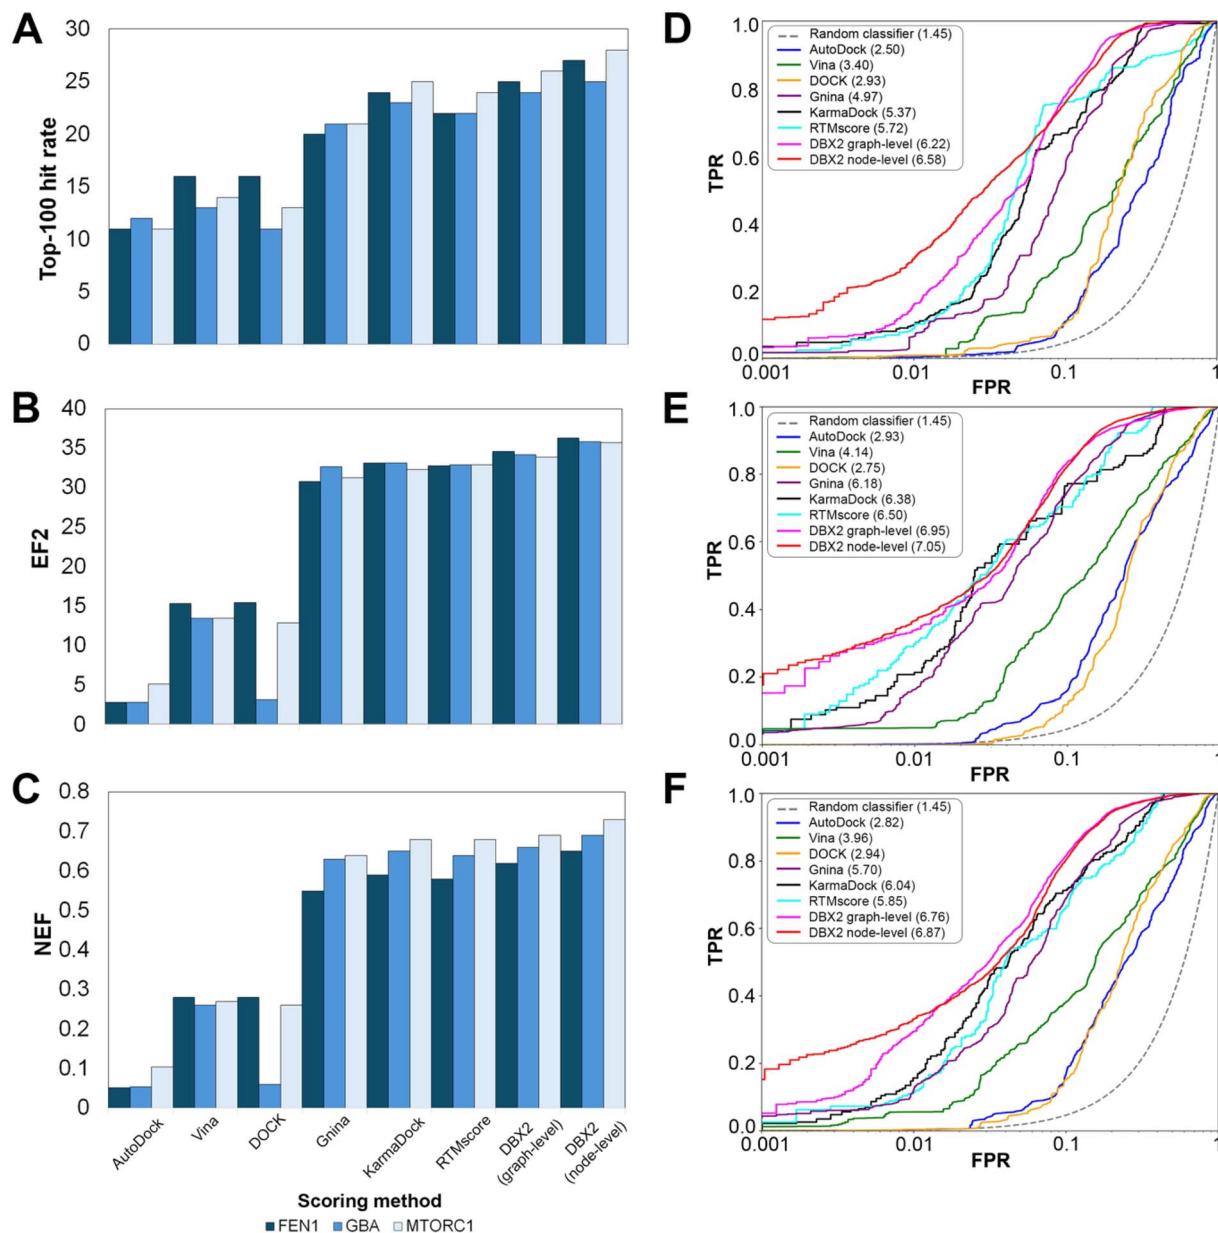


Fig. 3 Retrospective VS results of different scoring functions on three proteins from the LIT-PCBA database, (A) top-100 hit rate (B) EF2 (C) NEF, illustrating the significant performances of DBX2 node- and graph-level scores across different targets. LogAUC plots computed for (D) Flap structure-specific Endonuclease 1 (FEN1), (E) Glucocerebrosidase (GBA), and (F) Mechanistic Target of Rapamycin (MTORC1) confirmed the promising performance of the two DBX2 scores.

## Conclusions

We introduced DBX2, a novel GNN framework that enables to represent computational ensembles of small molecule-protein conformations as single graphs to jointly predict binding modes and affinities. The model relies solely on simple energetic features derived directly from docking, thus without requiring additional costly sampling steps. We comprehensively evaluated DBX2 across various metrics for docking and VS tasks, underscoring its effectiveness as a robust tool with superior performances compared to conventional scoring functions and ML models relying on single pose. At the same

time, some caveats associated with the newly proposed ensemble-based method emerged, especially reflected in the relatively poor correlation between graph-level prediction and experimental binding affinities. We reasoned that these constraints can be ascribed to the limitations of the data generating process, *i.e.*, docking, both in sampling the free energy landscape of binding and estimating the binding energy contributions that are used as features. Nevertheless, the performances observed for DBX2 not only advocate for its adoption in prospective VS campaigns relying on high throughput VS but encourages also further exploration of ML models learning from computationally generated ensembles

that can represent the thermodynamics of binding better than single poses. In this context, an exciting venue for further investigation could be the adaptation of the DBX2 architecture to MD-derived conformational ensembles of small molecule–protein complexes, to take into consideration protein flexibility, induced fit effect, solvation, and overall equilibrium ensembles.

## Author contributions

TT: methodology, data curation, investigation, formal analysis, writing – original draft, writing – review & editing. JP: conceptualization, data curation, software, writing – original draft. FG: conceptualization, funding acquisition, supervision, writing – original draft, writing – review & editing.

## Conflicts of interest

The authors declare no conflict of interest.

## Data availability

The DBX2 code is available at <https://github.com/jp43/DockBox2>. Trained models and training data are available at <https://doi.org/10.5281/zenodo.14181651>.

Supplementary information is available. See DOI: <https://doi.org/10.1039/d4sc07875f>.

## Acknowledgements

This work was supported by a Natural Sciences and Engineering Research Council of Canada Discovery Grant (RGPIN-2023-04129) and a uOttawa start-up grant awarded to FG. Computations were performed on the resources of the Digital Research Alliance of Canada (RRG ID 4879 awarded to FG) and the University of Ottawa's Wooki supercomputing cluster. We thank Cadence Molecular Sciences for providing an academic license for Openeye suite.

## References

- 1 S.-F. Zhou and W.-Z. Zhong, Drug Design and Discovery: Principles and Applications, *Molecules*, 2017, **22**(2), 279, DOI: [10.3390/molecules22020279](https://doi.org/10.3390/molecules22020279).
- 2 X. Zeng, S.-J. Li, S.-Q. Lv, M.-L. Wen and Y. Li, A comprehensive review of the recent advances on predicting drug–target affinity based on deep learning, *Front. Pharmacol.*, 2024, **15**, DOI: [10.3389/fphar.2024.1375522](https://doi.org/10.3389/fphar.2024.1375522).
- 3 X. Du, *et al.*, Insights into Protein–Ligand Interactions: Mechanisms, Models, and Methods, *Int. J. Mol. Sci.*, 2016, **17**(2), 144, DOI: [10.3390/ijms17020144](https://doi.org/10.3390/ijms17020144).
- 4 D. J. Newman and G. M. Cragg, Natural Products as Sources of New Drugs over the Nearly Four Decades from 01/1981 to 09/2019, *J. Nat. Prod.*, 2020, **83**(3), 770–803, DOI: [10.1021/acs.jnatprod.9b01285](https://doi.org/10.1021/acs.jnatprod.9b01285).
- 5 T. Takebe, R. Imai and S. Ono, The Current Status of Drug Discovery and Development as Originated in United States Academia: The Influence of Industrial and Academic Collaboration on Drug Discovery and Development, *Clin. Transl. Sci.*, 2018, **11**(6), 597–606, DOI: [10.1111/cts.12577](https://doi.org/10.1111/cts.12577).
- 6 J. Kuan, M. Radaeva, A. Avenido, A. Cherkasov and F. Gentile, Keeping pace with the explosive growth of chemical libraries with structure-based virtual screening, *Wiley Interdiscip. Rev.: Comput. Mol. Sci.*, 2023, **13**(6), e1678, DOI: [10.1002/wcms.1678](https://doi.org/10.1002/wcms.1678).
- 7 B. Shaker, S. Ahmad, J. Lee, C. Jung and D. Na, In silico methods and tools for drug discovery, *Comput. Biol. Med.*, 2021, **137**, 104851, DOI: [10.1016/j.combiomed.2021.104851](https://doi.org/10.1016/j.combiomed.2021.104851).
- 8 M. De Vivo, M. Masetti, G. Bottegoni and A. Cavalli, Role of Molecular Dynamics and Related Methods in Drug Discovery, *J. Med. Chem.*, 2016, **59**(9), 4035–4061, DOI: [10.1021/acs.jmedchem.5b01684](https://doi.org/10.1021/acs.jmedchem.5b01684).
- 9 S. Decherchi and A. Cavalli, Thermodynamics and Kinetics of Drug–Target Binding by Molecular Simulation, *Chem. Rev.*, 2020, **120**(23), 12788–12833, DOI: [10.1021/acs.chemrev.0c00534](https://doi.org/10.1021/acs.chemrev.0c00534).
- 10 Y. Miao, A. Bhattacharai and J. Wang, Ligand Gaussian Accelerated Molecular Dynamics (LiGaMD): Characterization of Ligand Binding Thermodynamics and Kinetics, *J. Chem. Theory Comput.*, 2020, **16**(9), 5526–5547, DOI: [10.1021/acs.jctc.0c00395](https://doi.org/10.1021/acs.jctc.0c00395).
- 11 Y.-T. Wang, *et al.*, Structural insights into Nirmatrelvir (PF-07321332)-3C-like SARS-CoV-2 protease complexation: a ligand Gaussian accelerated molecular dynamics study, *Phys. Chem. Chem. Phys.*, 2022, **24**(37), 22898–22904, DOI: [10.1039/D2CP02882D](https://doi.org/10.1039/D2CP02882D).
- 12 D. W. Kneller, *et al.*, Covalent nirlaprevir- and boceprevir-derived hybrid inhibitors of SARS-CoV-2 main protease, *Nat. Commun.*, 2022, **13**(1), 2268, DOI: [10.1038/s41467-022-29915-z](https://doi.org/10.1038/s41467-022-29915-z).
- 13 S. Wolf, B. Lickert, S. Bray and G. Stock, Multisecond ligand dissociation dynamics from atomistic simulations, *Nat. Commun.*, 2020, **11**(1), 2918, DOI: [10.1038/s41467-020-16655-1](https://doi.org/10.1038/s41467-020-16655-1).
- 14 E. Paquet and H. L. Viktor, Molecular Dynamics, Monte Carlo Simulations, and Langevin Dynamics: A Computational Review, *BioMed Res. Int.*, 2015, **2015**, 1–18, DOI: [10.1155/2015/183918](https://doi.org/10.1155/2015/183918).
- 15 K. Crampon, A. Giorkallos, M. Deldossi, S. Baud and L. A. Steffanel, Machine-learning methods for ligand–protein molecular docking, *Drug Discovery Today*, 2022, **27**(1), 151–164, DOI: [10.1016/j.drudis.2021.09.007](https://doi.org/10.1016/j.drudis.2021.09.007).
- 16 F. Liu, *et al.*, Large library docking identifies positive allosteric modulators of the calcium-sensing receptor, *Science*, 2024, **385**(6715), eado1868, DOI: [10.1126/science.ado1868](https://doi.org/10.1126/science.ado1868).
- 17 J. Lyu, *et al.*, Ultra-large library docking for discovering new chemotypes, *Nature*, 2019, **566**(7743), 224–229, DOI: [10.1038/s41586-019-0917-9](https://doi.org/10.1038/s41586-019-0917-9).
- 18 A. Manglik, *et al.*, Structure-based discovery of opioid analgesics with reduced side effects, *Nature*, 2016, **537**(7619), 185–190, DOI: [10.1038/nature19112](https://doi.org/10.1038/nature19112).



19 N. Zernov, V. Ghamaryan, D. Melenteva, A. Makichyan, L. Hunanyan and E. Popugaeva, Discovery of a novel piperazine derivative, cmp2: a selective TRPC6 activator suitable for treatment of synaptic deficiency in Alzheimer's disease hippocampal neurons, *Sci. Rep.*, 2024, **14**(1), 23512, DOI: [10.1038/s41598-024-73849-z](https://doi.org/10.1038/s41598-024-73849-z).

20 R. M. Stein, *et al.*, Virtual discovery of melatonin receptor ligands to modulate circadian rhythms, *Nature*, 2020, **579**(7800), 609–614, DOI: [10.1038/s41586-020-2027-0](https://doi.org/10.1038/s41586-020-2027-0).

21 E. A. Fink, *et al.*, Structure-based discovery of nonopioid analgesics acting through the  $\alpha_{2A}$ -adrenergic receptor, *Science*, 2022, **377**(6614), eabn7065, DOI: [10.1126/science.abn7065](https://doi.org/10.1126/science.abn7065).

22 K. M. Elokely and R. J. Doerkens, Docking Challenge: Protein Sampling and Molecular Docking Performance, *J. Chem. Inf. Model.*, 2013, **53**(8), 1934–1945, DOI: [10.1021/ci400040d](https://doi.org/10.1021/ci400040d).

23 A. Mastropietro, G. Pasculli and J. Bajorath, Learning characteristics of graph neural networks predicting protein-ligand affinities, *Nature Machine Intelligence*, 2023, **5**(12), 1427–1436, DOI: [10.1038/s42256-023-00756-9](https://doi.org/10.1038/s42256-023-00756-9).

24 J. Wu, H. Chen, M. Cheng and H. Xiong, CurvAGN: Curvature-based Adaptive Graph Neural Networks for Predicting Protein-Ligand Binding Affinity, *BMC Bioinf.*, 2023, **24**(1), 378, DOI: [10.1186/s12859-023-05503-w](https://doi.org/10.1186/s12859-023-05503-w).

25 S. Moon, W. Zhung, S. Yang, J. Lim and W. Y. Kim, PIGNet: a physics-informed deep learning model toward generalized drug–target interaction predictions, *Chem. Sci.*, 2022, **13**(13), 3661–3673, DOI: [10.1039/D1SC06946B](https://doi.org/10.1039/D1SC06946B).

26 C. Shen, *et al.*, A generalized protein–ligand scoring framework with balanced scoring, docking, ranking and screening powers, *Chem. Sci.*, 2023, **14**(30), 8129–8146, DOI: [10.1039/D3SC02044D](https://doi.org/10.1039/D3SC02044D).

27 S. Zhang, *et al.*, SS-GNN: A Simple-Structured Graph Neural Network for Affinity Prediction, *ACS Omega*, 2023, **8**(25), 22496–22507, DOI: [10.1021/acsomega.3c00085](https://doi.org/10.1021/acsomega.3c00085).

28 H. Shen, Y. Zhang, C. Zheng, B. Wang and P. Chen, A Cascade Graph Convolutional Network for Predicting Protein–Ligand Binding Affinity, *Int. J. Mol. Sci.*, 2021, **22**(8), 4023, DOI: [10.3390/ijms22084023](https://doi.org/10.3390/ijms22084023).

29 M. Volkov, *et al.*, On the Frustration to Predict Binding Affinities from Protein–Ligand Structures with Deep Neural Networks, *J. Med. Chem.*, 2022, **65**(11), 7946–7958, DOI: [10.1021/acs.jmedchem.2c00487](https://doi.org/10.1021/acs.jmedchem.2c00487).

30 G. Corso, H. Stärk, B. Jing, R. Barzilay and T. Jaakkola, DiffDock: Diffusion Steps, Twists, and Turns for Molecular Docking, *arXiv*, 2022, preprint, arXiv:2210.01776, DOI: [10.48550/ARXIV.2210.01776](https://doi.org/10.48550/ARXIV.2210.01776).

31 H. Jiang, *et al.*, Predicting Protein–Ligand Docking Structure with Graph Neural Network, *J. Chem. Inf. Model.*, 2022, **62**(12), 2923–2932, DOI: [10.1021/acs.jcim.2c00127](https://doi.org/10.1021/acs.jcim.2c00127).

32 Y. Min, *et al.*, From Static to Dynamic Structures: Improving Binding Affinity Prediction with Graph-Based Deep Learning, *arXiv*, 2022, preprint, arXiv:2208.10230, DOI: [10.48550/ARXIV.2208.10230](https://doi.org/10.48550/ARXIV.2208.10230).

33 R. Wang, X. Fang, Y. Lu and S. Wang, The PDBbind Database: Collection of Binding Affinities for Protein–Ligand Complexes with Known Three-Dimensional Structures, *J. Med. Chem.*, 2004, **47**(12), 2977–2980, DOI: [10.1021/jm030580l](https://doi.org/10.1021/jm030580l).

34 Z. Liu, *et al.*, Forging the Basis for Developing Protein–Ligand Interaction Scoring Functions, *Acc. Chem. Res.*, 2017, **50**(2), 302–309, DOI: [10.1021/acs.accounts.6b00491](https://doi.org/10.1021/acs.accounts.6b00491).

35 M. Su, *et al.*, Comparative Assessment of Scoring Functions: The CASF-2016 Update, *J. Chem. Inf. Model.*, 2019, **59**(2), 895–913, DOI: [10.1021/acs.jcim.8b00545](https://doi.org/10.1021/acs.jcim.8b00545).

36 P. Škrinjar, J. Eberhardt, J. Durairaj and T. Schwede, Have protein–ligand co-folding methods moved beyond memorisation?, *bioRxiv*, 2025, preprint, DOI: [10.1101/2025.02.03.636309](https://doi.org/10.1101/2025.02.03.636309).

37 J. Abramson, *et al.*, Accurate structure prediction of biomolecular interactions with AlphaFold 3, *Nature*, 2024, **630**(8016), 493–500, DOI: [10.1038/s41586-024-07487-w](https://doi.org/10.1038/s41586-024-07487-w).

38 J. Boitraud, *et al.*, Chai-1: decoding the molecular interactions of life, *bioRxiv*, 2024, preprint, DOI: [10.1101/2024.10.10.615955](https://doi.org/10.1101/2024.10.10.615955).

39 X. Chen, *et al.*, Protenix - Advancing Structure Prediction through a Comprehensive AlphaFold3 Reproduction, *bioRxiv*, 2025, preprint, DOI: [10.1101/2025.01.08.631967](https://doi.org/10.1101/2025.01.08.631967).

40 J. Wohlwend, *et al.*, Boltz-1: Democratizing Biomolecular Interaction Modeling, *bioRxiv*, 2024, preprint, DOI: [10.1101/2024.11.19.624167](https://doi.org/10.1101/2024.11.19.624167).

41 V.-K. Tran-Nguyen, C. Jacquemard and D. Rognan, LIT-PCBA: An Unbiased Data Set for Machine Learning and Virtual Screening, *J. Chem. Inf. Model.*, 2020, **60**(9), 4263–4273, DOI: [10.1021/acs.jcim.0c00155](https://doi.org/10.1021/acs.jcim.0c00155).

42 J. Preto and F. Gentile, Assessing and improving the performance of consensus docking strategies using the DockBox package, *J. Comput.-Aided Mol. Des.*, 2019, **33**(9), 817–829, DOI: [10.1007/s10822-019-00227-7](https://doi.org/10.1007/s10822-019-00227-7).

43 OpenEye, *OpenEye Toolkits*, Cadence Molecular Sciences, Santa Fe, NM, available: <http://www.eyesopen.com>.

44 *Molecular Operating Environment (MOE)*, Chemical Computing Group ULC, Montreal, QC.

45 G. M. Morris, *et al.*, AutoDock4 and AutoDockTools4: automated docking with selective receptor flexibility, *J. Comput. Chem.*, 2009, **30**(16), 2785–2791, DOI: [10.1002/jcc.21256](https://doi.org/10.1002/jcc.21256).

46 O. Trott and A. J. Olson, AutoDock Vina: improving the speed and accuracy of docking with a new scoring function, efficient optimization, and multithreading, *J. Comput. Chem.*, 2010, **31**(2), 455–461, DOI: [10.1002/jcc.21334](https://doi.org/10.1002/jcc.21334).

47 T. E. Baliaus, S. Mukherjee and R. C. Rizzo, Implementation and evaluation of a docking–rescoring method using molecular footprint comparisons, *J. Comput. Chem.*, 2011, **32**(10), 2273–2289, DOI: [10.1002/jcc.21814](https://doi.org/10.1002/jcc.21814).

48 A. T. McNutt, *et al.*, GNINA 1.0: molecular docking with deep learning, *J. Cheminf.*, 2021, **13**(1), 43, DOI: [10.1186/s13321-021-00522-2](https://doi.org/10.1186/s13321-021-00522-2).

49 G. Neudert and G. Klebe, DSX: A Knowledge-Based Scoring Function for the Assessment of Protein–Ligand Complexes, *J. Chem. Inf. Model.*, 2011, **51**(10), 2731–2745, DOI: [10.1021/ci200274q](https://doi.org/10.1021/ci200274q).

50 G. M. Morris, *et al.*, Automated docking using a Lamarckian genetic algorithm and an empirical binding free energy



function, *J. Comput. Chem.*, 1998, **19**(14), 1639–1662, DOI: [10.1002/\(SICI\)1096-987X\(19981115\)19:14%3C1639::AID-JCC10%3E3.0.CO;2-B](https://doi.org/10.1002/(SICI)1096-987X(19981115)19:14%3C1639::AID-JCC10%3E3.0.CO;2-B).

51 R. Salomon-Ferrer, D. A. Case and R. C. Walker, An overview of the Amber biomolecular simulation package, *Wiley Interdiscip. Rev.: Comput. Mol. Sci.*, 2013, **3**(2), 198–210, DOI: [10.1002/wcms.1121](https://doi.org/10.1002/wcms.1121).

52 W. L. Hamilton, R. Ying and J. Leskovec, Inductive Representation Learning on Large Graphs, *arXiv*, 2018, preprint, arXiv:1706.02216, DOI: [10.48550/arXiv.1706.02216](https://doi.org/10.48550/arXiv.1706.02216), accessed: October 01, 2024.

53 D. Duvenaud, *et al.*, Convolutional Networks on Graphs for Learning Molecular Fingerprints, *arXiv*, 2015, preprint, arXiv:1509.09292, DOI: [10.48550/arXiv.1509.09292](https://doi.org/10.48550/arXiv.1509.09292), accessed: October 01, 2024.

54 T.-Y. Lin, P. Goyal, R. Girshick, K. He and P. Dollár, Focal Loss for Dense Object Detection, *arXiv*, 2018, preprint, arXiv:1708.02002, DOI: [10.48550/arXiv.1708.02002](https://doi.org/10.48550/arXiv.1708.02002), accessed: October 19, 2024.

55 X. Zhang, Y. Li, J. Wang, G. Xu and Y. Gu, A Multi-perspective Model for Protein–Ligand-Binding Affinity Prediction, *Interdiscip. Sci.: Comput. Life Sci.*, 2023, **15**(4), 696–709, DOI: [10.1007/s12539-023-00582-y](https://doi.org/10.1007/s12539-023-00582-y).

56 C. Cortes, M. Mohri and A. Rostamizadeh, L2 Regularization for Learning Kernels, *arXiv*, 2012, preprint, arXiv:1205.2653, DOI: [10.48550/arXiv.1205.2653](https://doi.org/10.48550/arXiv.1205.2653), accessed: October 20, 2024.

57 F. Pedregosa, *et al.*, Scikit-learn: Machine Learning in Python, *Journal of Machine Learning Research*, 2011, **12**, 2825–2830.

58 D. A. Pearlman and P. S. Charifson, Are Free Energy Calculations Useful in Practice? A Comparison with Rapid Scoring Functions for the p38 MAP Kinase Protein System, *J. Med. Chem.*, 2001, **44**(21), 3417–3423, DOI: [10.1021/jm0100279](https://doi.org/10.1021/jm0100279).

59 J. C. Exell, *et al.*, Cellularly active N-hydroxyurea FEN1 inhibitors block substrate entry to the active site, *Nat. Chem. Biol.*, 2016, **12**(10), 815–821, DOI: [10.1038/nchembio.2148](https://doi.org/10.1038/nchembio.2148).

60 B. Brumshtein, *et al.*, Cyclodextrin-mediated crystallization of acid  $\beta$ -glucosidase in complex with amphiphilic bicyclic nojirimycin analogues, *Org. Biomol. Chem.*, 2011, **9**(11), 4160, DOI: [10.1039/c1ob05200d](https://doi.org/10.1039/c1ob05200d).

61 S.-Y. Lee, *et al.*, Proximity-Directed Labeling Reveals a New Rapamycin-Induced Heterodimer of FKBP25 and FRB in Live Cells, *ACS Cent. Sci.*, 2016, **2**(8), 506–516, DOI: [10.1021/acscentsci.6b00137](https://doi.org/10.1021/acscentsci.6b00137).

62 K. Palacio-Rodríguez, I. Lans, C. N. Cavasotto and P. Cossio, Exponential consensus ranking improves the outcome in docking and receptor ensemble docking, *Sci. Rep.*, 2019, **9**(1), 5142, DOI: [10.1038/s41598-019-41594-3](https://doi.org/10.1038/s41598-019-41594-3).

63 J.-F. Truchon and C. I. Bayly, Evaluating Virtual Screening Methods: Good and Bad Metrics for the ‘Early Recognition’ Problem, *J. Chem. Inf. Model.*, 2007, **47**(2), 488–508, DOI: [10.1021/ci600426e](https://doi.org/10.1021/ci600426e).

64 Y. Perez-Castillo, *et al.*, Fusing Docking Scoring Functions Improves the Virtual Screening Performance for Discovering Parkinson’s Disease Dual Target Ligands, *Curr. Neuropharmacol.*, 2017, **15**(8), 1107–1116, DOI: [10.2174/1570159X15666170109143757](https://doi.org/10.2174/1570159X15666170109143757).

65 G.-L. Xiong, W.-L. Ye, C. Shen, A.-P. Lu, T.-J. Hou and D.-S. Cao, Improving structure-based virtual screening performance via learning from scoring function components, *Briefings Bioinf.*, 2021, **22**(3), bbaa094, DOI: [10.1093/bib/bbaa094](https://doi.org/10.1093/bib/bbaa094).

66 J. C. D. Lopes, F. M. Dos Santos, A. Martins-José, K. Augustyns and H. De Winter, The power metric: a new statistically robust enrichment-type metric for virtual screening applications with early recovery capability, *J. Cheminf.*, 2017, **9**(1), 7, DOI: [10.1186/s13321-016-0189-4](https://doi.org/10.1186/s13321-016-0189-4).

67 S. Liu, *et al.*, Practical Model Selection for Prospective Virtual Screening, *J. Chem. Inf. Model.*, 2019, **59**(1), 282–293, DOI: [10.1021/acs.jcim.8b00363](https://doi.org/10.1021/acs.jcim.8b00363).

68 R. P. Sheridan, S. B. Singh, E. M. Fluder and S. K. Kearsley, Protocols for Bridging the Peptide to Nonpeptide Gap in Topological Similarity Searches, *J. Chem. Inf. Comput. Sci.*, 2001, **41**(5), 1395–1406, DOI: [10.1021/ci0100144](https://doi.org/10.1021/ci0100144).

69 X. Zhang, *et al.*, Efficient and accurate large library ligand docking with KarmaDock, *Nat. Comput. Sci.*, 2023, **3**(9), 789–804, DOI: [10.1038/s43588-023-00511-5](https://doi.org/10.1038/s43588-023-00511-5).

70 C. Shen, *et al.*, Boosting Protein–Ligand Binding Pose Prediction and Virtual Screening Based on Residue–Atom Distance Likelihood Potential and Graph Transformer, *J. Med. Chem.*, 2022, **65**(15), 10691–10706, DOI: [10.1021/acs.jmedchem.2c00991](https://doi.org/10.1021/acs.jmedchem.2c00991).

71 H. Cai, *et al.*, CarsiDock: a deep learning paradigm for accurate protein–ligand docking and screening based on large-scale pre-training, *Chem. Sci.*, 2024, **15**(4), 1449–1471, DOI: [10.1039/D3SC05552C](https://doi.org/10.1039/D3SC05552C).

72 A. J. Oakley, *et al.*, The structures of human glutathione transferase P1-1 in complex with glutathione and various inhibitors at high resolution, *J. Mol. Biol.*, 1997, **274**(1), 84–100, DOI: [10.1006/jmbi.1997.1364](https://doi.org/10.1006/jmbi.1997.1364).

73 E. J. Martin, V. R. Polyakov, X.-W. Zhu, L. Tian, P. Mukherjee and X. Liu, All-Assay-Max2 pQSAR: Activity Predictions as Accurate as Four-Concentration IC<sub>50</sub>s for 8558 Novartis Assays, *J. Chem. Inf. Model.*, 2019, **59**(10), 4450–4459, DOI: [10.1021/acs.jcim.9b00375](https://doi.org/10.1021/acs.jcim.9b00375).

74 M. M. Stepniewska-Dziubinska, P. Zielenkiewicz and P. Siedlecki, Development and evaluation of a deep learning model for protein–ligand binding affinity prediction, *Bioinformatics*, 2018, **34**(21), 3666–3674, DOI: [10.1093/bioinformatics/bty374](https://doi.org/10.1093/bioinformatics/bty374).

75 F. Li, *et al.*, CACHE Challenge #1: targeting the WDR domain of LRRK2, a Parkinson’s disease associated protein, *J. Chem. Inf. Model.*, 2024, **64**, 8521–8536, DOI: [10.1021/acs.jcim.4c01267](https://doi.org/10.1021/acs.jcim.4c01267).

76 I. Dunn, S. Pirhadi, Y. Wang, S. Ravindran, C. Concepcion and D. R. Koes, CACHE Challenge #1: Docking with GNINA Is All You Need, *J. Chem. Inf. Model.*, 2024, **64**(24), 9388–9396, DOI: [10.1021/acs.jcim.4c01429](https://doi.org/10.1021/acs.jcim.4c01429).

