

Large-scale prospective evaluation of co-folding across hundreds of Mac1–ligand complexes and three virtual screens

Jongbin Kim^{1†}, Galen J Correy^{2†}, Brendan W Hall^{1†}, Moira M Rachman¹, Olivier Mailhot¹, Takaya Togo¹, Ryan L Gonciarz¹, Priyadarshini Jaishankar¹, R Jeffrey Neitz^{1,3}, Eric R Hantz¹, Yagmur U Doruk⁴, Maisie GV Stevens⁴, Morgan E Diolaiti⁴, Rashad Reid⁵, Saumya Gopalkrishnan⁶, Nevan J Krogan^{2,6,7}, Adam R Renslo¹, Alan Ashworth⁴, Brian K Shoichet^{1*}, James S Fraser^{2*}

¹Department of Pharmaceutical Chemistry, University of California, San Francisco, San Francisco, United States; ²Department of Bioengineering and Therapeutic Sciences, University of California, San Francisco, San Francisco, United States; ³Small Molecule Discovery Center, Department of Pharmaceutical Chemistry, University of California, San Francisco, San Francisco, United States; ⁴Helen Diller Family Comprehensive Cancer Center, University of California, San Francisco, San Francisco, United States; ⁵Department of Biological and Medical Informatics, University of California, San Francisco, San Francisco, United States; ⁶Quantitative Biosciences Institute (QBI), University of California, San Francisco, San Francisco, United States; ⁷Data Science and Biotechnology Institute, Gladstone Institutes, San Francisco, United States

*For correspondence:
bshoichet@gmail.com (BKS);
jfraser@fraserlab.com (JSF)

[†]These authors contributed
equally to this work

Competing interest: See page
18

Funding: See page 18

Preprint posted
29 December 2025

Sent for Review
14 January 2026

Reviewed preprint v1
18 March 2026

Reviewed preprint v2
24 September 2026

Version of Record
12 October 2026

Reviewing Editor: Fabrizio
Marinelli, Medical College of
Wisconsin, United States

© Copyright Kim, Correy, Hall
et al. This article is distributed
under the terms of the [Creative
Commons Attribution License](#),
which permits unrestricted use
and redistribution provided that
the original author and source
are credited.

eLife Assessment

This study addresses an **important** gap in drug discovery by delivering a rigorous, large-scale evaluation of widely used co-folding methods for predicting ligand-bound protein complexes and virtual screening. A key strength is the comprehensive benchmarking framework, which leverages structures and chemical compounds that were absent from the AI models training set, thereby providing particularly **compelling** and unbiased evidence of co-folding performance. The findings clearly delineate the complementary roles of deep learning-based co-folding and physics-based docking, offering practical guidance for their rational integration into drug discovery workflows. Overall, the conclusions are well supported by thorough analyses across a representative set of cases.

Abstract Accurate prediction of ligand-bound protein complexes and ranking them by affinity are central problems in drug discovery. While deep-learning co-folding methods can help address these challenges, their evaluation has been hampered by the difficulties in assessing independence from training data and insufficiently large test sets. Here, we test the ability of co-folding methods to predict the structures of 551 ligands, 489 of which are of sufficient quality for evaluation, bound to the SARS-CoV-2 NSP3 macrodomain (Mac1) that were determined after the training cutoff dates. AlphaFold3 (AF3), Boltz-2, and Chai-1 each reproduced >50% of the Mac1 ligand poses to better than 2 Å root-mean-square deviation of experiment. Despite the potential for co-folding to describe protein conformational changes that stabilize ligand binding, we did not find that common conformational rearrangements, including peptide flip and a large loop opening, were recapitulated by the co-folding prediction. For AF3 and Chai-1, ligand pose prediction confidence weakly, but significantly, tracked experimental potency, while DOCK3.7 energies were only weakly correlated. Boltz-2 affinity predictions showed the strongest correlation with measured potency and, after calibration,

achieved lower mean absolute error than a baseline predictor. We next assessed whether co-folding scores could rescore docking hit lists to distinguish true ligands from non-binders among hundreds of molecules prospectively experimentally tested against AmpC β -lactamase, the dopamine D4, and the σ_2 receptors. AF3 ligand pose confidence values did not separate true ligands from high-scoring false positives as effectively as docking scores or Boltz-2 affinity predictions did. Taken together, the modest, but independent correlations of docking score and co-folding confidence or affinity suggest that integrating physics-based and deep-learning approaches may help with hit prioritization and subsequent optimization in structure-based ligand discovery.

Introduction

Deep-learning methods such as AlphaFold2 ([Jumper et al., 2021](#)) and RosettaFold ([Baek et al., 2021](#)) revolutionized protein structure prediction. Key to their wide adoption has been their extensive testing, first in well-controlled prospective tests like CASP ([Kryshtafovych et al., 2021](#); [Ashizawa et al., 2026](#)) and subsequently in many studies where the methods found new applications and new experimental validation ([Scardino et al., 2023](#); [Goverde et al., 2023](#)). Notably, these methods can provide protein structural models for small molecule ligand discovery by docking/virtual screening, leveraging physics-inspired functions to score the resulting complexes ([Seok and Tiwary, 2025](#); [Brocidiaccono et al., 2025](#); [Huang et al., 2006](#); [Michino et al., 2025](#); [Xie et al., 2025](#)). Retrospective analyses of AlphaFold2 models suggested that the accuracy of the predicted ligand binding poses was lower than experimental structures ([Karelina et al., 2023](#); [Holcomb et al., 2023](#)). However, prospective use of AlphaFold2 models and experimental structures (determined after the training cutoff) to identify new ligands revealed that both approaches yielded similar quality and quantity of ligands ([Lyu et al., 2024](#); [Díaz-Holguín et al., 2024](#)). These results indicate that AF2 models may sample alternative low-energy conformations that are compatible with distinct ligands from the experimentally resolved conformations, reconciling the poor retrospective recall with the prospective successes.

More recently, the advent of co-folding methods like AlphaFold3 (AF3) ([Abramson et al., 2024](#)), Chai ([Boitreaud et al., 2024](#)), and Boltz ([Wohlwend et al., 2025](#)) that leverage diffusion-based approaches for multi-component (e.g. protein and small molecule) complex predictions have promised to extend the impact of deep-learning methods from protein structure prediction to small molecule ligand-bound protein complexes ([He et al., 2025](#)). The advantages co-folding may have over traditional docking include its abilities to explore conformational variability of the protein on-the-fly and to predict ligand affinity, reportedly as well as or better than a gold-standard physics-based method, free energy perturbation ([Passaro et al., 2025](#)). However, multiple studies have highlighted how co-folding success can track with structural similarity to training structures, consistent with substantial structural memorization ([Xu et al., 2025](#); [Škrinjar et al., 2025](#)). An innovative approach introduced adversarial mutations to fill binding sites to show that co-folding models will still produce high-confidence poses even when interactions are physically impossible ([Masters et al., 2025](#)). These results highlight pervasive memorization and hallucination in current co-folding methods and suggest that alternative computational approaches are needed to quantify the physical basis of the protein–ligand interactions.

Still, co-folding has potential for synergizing with many aspects of structure-based design and more traditional computational chemistry approaches. Two attractive applications are the prediction of ligand-bound complexes, around which structure-based ligand design is articulated ([Erlanson et al., 2025](#); [Backus et al., 2016](#)), and the re-scoring of large library docking hit lists. Co-folding is orders of magnitude slower than most docking programs, which can address libraries of hundreds of millions to billions of compounds on CPU clusters. The method nevertheless is fast enough to re-rank virtual screening ‘hit lists’ and may be better at categorizing certain types of ligands (e.g. those that stabilize alternative conformations) than the simpler physics-based methods. Moreover, the errors and biases from co-folding affinity estimation methods ([Passaro et al., 2025](#)) may be independent from those in docking, suggesting that each method may have different strengths in distinguishing likely from unlikely ligands. However, partly because of the lack of widely adopted community prospective tests, akin to CASP ([Gathiaka et al., 2016](#)), and partly because of the challenges of developing, testing, and potentially optimizing new chemical matter for ligand activity studies ([Ackloo et al., 2022](#)), co-folding

methods have not yet been subject to the rigorous, large-scale testing that helped establish the utility, domain of applicability, and ultimately optimization of the early protein structure-prediction methods.

It therefore seemed interesting to test co-folding predictions of the co-complexed structures and affinities of ligands in a drug-discovery series from outside of their training sets, and to test their ability to distinguish likely from unlikely ligands from large library virtual screens (*Cabeza de Vaca et al., 2025*). Here, we evaluated the performance of three widely used co-folding methods, AF3, Chai-1, and Boltz-2 against two benchmarks that speak to domain utility and are largely outside the benchmarks against which they trained. First, we tested the ability of the co-folding methods to recapitulate the crystallographic structures of 551 different ligands bound to the SARS-CoV-2 Macrodome (Mac1), an emerging antiviral target (*Suryawanshi et al., 2025*). This dataset represents the scale and depth typical of the early phase of industrial structure-based drug design campaigns. Second, we investigated the ability of the co-folding methods to distinguish true- from false positives in hit lists from recent large library, prospective docking against three well-behaved targets: AmpC β -lactamase, the σ_2 receptor, and the dopamine D4 receptor (*Lyu et al., 2019; Alon et al., 2021; Liu et al., 2025*). In the campaigns against these targets, between 500 and 1500 molecules were prospectively tested across a range of docking ranks identifying ligands with previously unexplored chemotypes as well as finding many docking false positives. Contemporaneous work has highlighted the importance of using AF3 in the same experimental datasets from ultra-large library docking campaigns (*Menon et al., 2025*). Our analyses reveal the strengths and limitations of different co-folding methods vs. traditional docking and point to potential synergies that can improve probe and lead discovery pipelines.

Results

A set of 551 previously unreported Mac1–ligand complexes

As part of a drug discovery effort (*Suryawanshi et al., 2025; Jaishankar et al., 2025; Correy et al., 2025; Flowers et al., 2025; Gahbauer et al., 2023; Schuller et al., 2021*), we tested thousands of novel small molecules as inhibitors of Mac1, a new target for SARS-CoV-2 and other viruses. This led to potent inhibitors that are active in vivo as antivirals (*Suryawanshi et al., 2025; Jaishankar et al., 2025*). As part of this work, we determined X-ray crystal structures for ~1000 small molecules bound to Mac1, typically at close to 1 Å resolution, but only a limited number of these were disclosed in the papers describing the more advanced candidates. Moreover, only up to 326 of our molecules, largely consisting of initial fragments and early hit-to-lead material, were available in public databases prior to the training cutoffs for the co-folding methods (e.g. 2021-09-30 for AF3, 2023-06-30 for Boltz-2). Apart from their role in driving inhibitor optimization for Mac1, we thought that these structures might be a resource for evaluating computational approaches against a set that is representative of the scale and depth often seen in industrial structure-based drug design campaigns.

From our internal dataset, we identified 551 ligands where we determined X-ray structures that have not been disclosed previously (*Figure 1a, Supplementary file 1*). These molecules all bind in the active site where they compete with the substrate, ADP-ribosylated (ADP-r) peptides. There are several stereotyped interactions across this set, which have been described in previous fragment and lead optimization efforts (*Gahbauer et al., 2023; Schuller et al., 2021; Correy et al., 2025*), including: interactions with backbone NH groups of Phe156 and Asp157 (the oxyanion site), aromatic packing against the side chain of Phe156 (mimicking the adenosine ring of ADP-r), and hydrogen bonds to Asp22 and Ile23 (the upper active site) (*Figure 1a*). Because these ligands were synthesized during an active campaign and many did not meet evolving thresholds for single point % inhibition, only 202 have fully measured dose-response IC_{50} values by a HTRF-based peptide displacement assay. For these compounds, the measured potencies range from sub-micromolar to >500 μ M (*Figure 1b*). This distribution is skewed towards weaker affinities as more potent molecules were progressed and disclosed in papers describing in vivo-active lead molecules (*Suryawanshi et al., 2025; Jaishankar et al., 2025*).

The diversity of the 551 ligands is important for understanding the performance of the co-folding methods. By topological similarity, using ECFP4-based pairwise Tanimoto coefficients (T_c), the ligands were diverse, with pairwise T_c clustering between 0.2 and 0.4 (*Figure 1—figure supplement 1a, Supplementary file 2*). For context, a scaffold-hop for the ECFP4-based fingerprint is thought to occur at T_c values below 0.35 (*Cramer et al., 2004; Muchmore et al., 2008*), while random similarity




Figure 1. Structure diversity of 551 previously unreported ligands co-folded with the SARS-CoV-2 Mac1 (NSP3 macrodomain 1). **(a)** Global superposition of all 551 ligands within the Mac1-binding pocket highlighting residues involved in key interactions. **(b)** The affinity distribution of 202 molecules determined by HTRF (dose-response curves at eight concentration points). **(c)** Principal component analysis projection of compounds represented by

Figure 1 continued on next page

Figure 1 continued

their full pairwise similarity vectors (by ECFP-4 Tc and MCS%). (d) The highest MCS% vs. ECFP-4 Tc values in the training set for each Mac1 compound. At right are examples of the low similarity, scaffold hop, and high similarity distributions, respectively.

The online version of this article includes the following figure supplement(s) for figure 1:

Figure supplement 1. Distribution of pairwise similarity for Mac1 ligands and an overlay of clusters of compounds tested on entire ligand set deposited in PDB.

Figure supplement 2. Benchmarked datasets are clustered based on TC and MCS% to predict the number of scaffold series.

occurs around Tc values of 0.25 (Hert et al., 2008). As the Tc metric has well-known liabilities that can mask what a chemist would see as similarity, we also calculated pairwise maximum common substructures (MCS) and found a distribution that centers around 44.8% (Figure 1—figure supplement 1b, Supplementary file 2). Clustering based on pairwise MCS >35% revealed only four distinct cluster heads and a limited number of outliers (Figure 1c, Supplementary file 2, Figure 1—figure supplement 2). To quantify the potential for model performance to be driven by memorization, for each ligand in our set, we retrieved the most similar ligand that was deposited prior to the training cutoff date by both MCS and Tc (Figure 1d). Here, the Tcs have a trimodal distribution with a small number having Tc >0.4 (indicating relatively high similarity for this fingerprint), and two larger distributions centered around 0.35 (indicating scaffold hops) and 0.2 (indicating low similarity to any ligand). The MCS similarities spanned a broad range centered around ~60%, consistent with the majority of the dataset representing a small number of conserved scaffolds bearing diverse peripheral substitutions. We could also confirm the chemical diversity of the Mac1 compounds relative to the broader chemical space represented by all ligands deposited in the PDB by both Tc and MCS% (Figure 1—figure supplement 1c, d).

Co-folding can accurately reproduce poses of ligands dissimilar to those trained

Next we assessed how well co-folding methods could predict the structures of the 551 Mac1-complexes. Three co-folding methods (AlphaFold3, Chai-1, and Boltz-2) were given the SMILES strings of each ligand and the sequence of the enzyme as inputs. In parallel, we performed classic docking with DOCK3.7 using the SMILES of the ligand and the three-dimensional structure of the enzyme bound to an early inhibitor (PDB: 5SQW, Gahbauer et al., 2023) that is roughly contemporaneous with the ligands in this dataset. However, 68 structures had weak density for part of the molecule and were not used as reference structures for comparing co-folded poses. To assess whether ligands were correctly predicted in the right area of the active site (Nittinger et al., 2025), we used a threshold of ligand center of mass <2.5 Å after alignment of 489 protein structures: 449 AlphaFold3 poses (92%), 440 Chai-1 poses (90%), and 444 Boltz-2 poses (91%). With DOCK3.7, only 80% (392 out of 489 poses) passed this cutoff. In addition, roughly 1.5% of co-folded or docked molecules did not pass the RDKit-based sanitization pipeline necessary for root-mean-square deviation (RMSD) comparisons. Therefore, the total number of poses compared for prospective prediction differed slightly among all co-folding and docking methods (Supplementary file 3).

Across the co-folding methods, AF3 and Chai-1 performed qualitatively similar with a roughly unimodal distribution centered around ~1 Å (the mode) RMSD values (Figure 2a). Using the field standard, 2 Å cutoff for correct predictions, AF3 and Chai-1 predicted roughly 70% of the ligands correctly. Boltz-2 performed less well, with more structures shifting to the 2 Å distribution, and an overall predictive accuracy of 53%. This result was surprising because the later training cutoff date for Boltz-2 meant that it could potentially benefit from additional Mac1 ligand-bound structures. The distribution for docking was more bimodal centered around an RMSD of 2 and 5 Å, yielding a predictive accuracy of 40%. At least for this set, the co-folding models outperformed DOCK3.7 in accurately predicting the poses of bound ligands.

Next we investigated the potential for co-folding methods to identify alternative conformations of the protein that are stabilized by ligand binding. This application is a potentially large differentiator relative to virtual screening, which generally uses a fixed structure. Although protein alignment gave generally low global C-alpha RMSDs (range 0.1–0.4 Å; Supplementary file 3), rare local and high magnitude conformational changes observed experimentally were not captured by co-folding




Figure 2. Co-folding outperforms DOCK3.7 in predicting the ligand pose for 489 ligands against Mac1 binding site. **(a)** Histogram distributions of ligand heavy-atom RMSD (L-RMSD) between predicted and crystallographic poses obtained with AF3, Chai-1, Boltz-2, and DOCK3.7, and the success rate defined as RMSD < 2 Å. **(b)** Demonstration of alternate conformations of Mac1 binding pocket (twisted backbone and open state) and ability of AF3 to capture these conformational changes (residues in green are from PDB ID: 5SQW, and residues in pink have altered conformations). **(c)** AF3 L-RMSD

Figure 2 continued on next page

Figure 2 continued

pose recovery is compared with TC ECFP4 and MCS% and accurate poses for test molecules with poor similarity to the closest molecule in the training set.

The online version of this article includes the following figure supplement(s) for figure 2:

Figure supplement 1. Detailed predictions of alternate residue conformations and interactions in co-folded poses.

methods. Previously published structures indicated two characteristic states with significant conformational changes: a twisted state, characterized by a peptide backbone flip around Phe156, and an open state, characterized by a ~8 Å loop rearrangement including Gly130 and Ala129 (**Figure 2b**, **Figure 2—figure supplement 1**, **Supplementary file 4**). The open state has also been observed in the human MacroD1 enzyme (PDB 2X47, **Chen et al., 2011** and PDB 6LH4, **Yang et al., 2020**), suggesting it might be a common feature across macrodomains. We observed that only 4/19 twisted states and 0/20 open states were correctly predicted by AF3. Interestingly, AF3 predicted eight false positive twisted states, where the ligand was experimentally observed in the normal ground state. This result suggests that there is only a limited ability to predict conformational adaptations, which admittedly may be, at least in part, stabilized by crystal contacts. We next asked how often the predictions recapitulated hydrogen bonding interactions in the Mac1 active site (**Figure 2—figure supplement 1a**). Among the dominant interactions, we found a low false positive rate for hydrogen bonds with Asp22 and Ile23 at the top of the active site, and mispredictions in both directions for hydrogen bonds with Phe156 and Asp157 at the oxyanion site (**Figure 2—figure supplement 1b**). Thus, despite failing to predict large conformational changes and aspects of the detailed hydrogen bonding interactions, AF3 still correctly placed these ligands with low RMSDs relative to the crystal structure pose (**Errington et al., 2025**). This result suggests that co-folding reliably recapitulates dominant ligand-binding interactions even in the absence of accurate protein conformational modeling, providing further support to the idea that they are learning specific interaction patterns rather than a deeper physics-based representation (**Masters et al., 2025**). Finally, we asked how much of co-folding success might be explained by memorization of Mac1 ligands deposited prior to the training date (**Figure 1d**). We measured the correlation between AF3 ligand RMSD and either the Tc (ECFP4) or %MCS of the most similar ligand within the training set, finding no meaningful signal (**Figure 2c**). The examples illustrated have low similarity in both metrics, and AVI-5083 (PDB: 15JC) and AVI-3363 (PDB: 15CH) are not from the commonly shared purine scaffold. Nonetheless, even for these cases, AF3 co-folded poses are recovered to within 1 Å accuracy. The weak correlations between RMSD and both similarity metrics illustrate the ability to predict binding modes for ligands dissimilar from the training set.

Model confidence and energy scores reflect pose reproduction and affinities

A central question for any co-folding or docking method is whether its internal confidence or energy scores track with experimental pose accuracy and affinity. We used the structure confidence scores reported by AF3, the per atom Predicted Local Distance Difference Test (pLDDT) scores for ligand atoms, and for Chai-1 the interface predictive template modeling (ipTM) score. While these metrics are intended to estimate the model confidence in accuracy, they are not explicitly designed to predict affinity. In contrast, the Boltz-2 affinity module predicts IC_{50} (pIC_{50}), which should correlate with affinity, while the DOCK3.7 score, though in kcal/mol, is not expected to correlate strongly with experimental affinity (**Shoichet, 2004; Irwin and Shoichet, 2016; Tirado-Rives and Jorgensen, 2006**). Boltz-2 affinity score and DOCK energies do not explicitly estimate the accuracy (RMSD) of the prediction.

We first generated receiver operating characteristic (ROC) curves to measure the retrieval of true positives (good score, RMSD <2 Å) vs. false positives (good score, RMSD >2 Å) (**Figure 3a**). We observed the best performance for AF3 L-pLDDT (AUC = 75.6) and Chai-1 ipTM (AUC = 74.6). This indicates that the internal model confidence has predictive power and, indeed, we observed that these scores are correlated with RMSD (**Figure 3—figure supplement 1**). Perhaps not surprisingly, the performance of the DOCK3.7 energy (AUC = 67.1) and Boltz-2 pIC_{50} (AUC = 53.8) scores, which are designed to predict affinity, was poorer than those designed to predict accuracy. The DOCK3.7 scores correlated with Mac1 ligand RMSD with a Pearson r of 0.33 ($p < 0.001$), which is in line with expectations from historical analyses of docking studies (**Wierbowski et al., 2020**). However, Boltz-2




Figure 3. Complementarity between co-folding confidence metrics and ligand-pose accuracy. (a) Receiver operating characteristic (ROC) curves for affinity/energy and confidence scores in AlphaFold3, Chai-1, and Boltz-2, with the success rate of binding pose prediction being defined as L-RMSD < 2 Å and both AUC and logAUC are shown in percentages. (b) Chemical structure of a compound that was challenging to model with all docking and co-folding methods (AVI-4225, PDB: 15DX), and one that was well-modeled with all methods (AVI-4129, PDB: 15DE), (c) Overlay of co-folded (AlphaFold3, Figure 3 continued on next page

Figure 3 continued

Chai-1, and Boltz-2) and docked ligand poses on ground-truth, for hits with similar IC_{50} , but different structures. Hydrogen bonds with co-folded, docked ligands are shown in red, and those with crystal ligands are shown in yellow.

The online version of this article includes the following figure supplement(s) for figure 3:

Figure supplement 1. Correlations between Mac1 pose recovery with co-folding/docking scoring metrics.

Figure supplement 2. Pose accuracy and co-folding, docking scores are compared for different similarity bins.

Figure supplement 3. Correlation between AF3 pose-prediction error and errors from other co-folding and docking methods for 551 Mac1 compounds.

Figure supplement 4. Pose recovery by methods of co-folding and docking against the ground-truth pose.

affinity scores did not correlate with the predicted pose (Pearson $r = -0.001$, $p > 0.5$), which likely results from the fact that its internal structure prediction confidence module is separated from the affinity module (Passaro et al., 2025). Indeed, using the Boltz-2 structure prediction confidence metric (Ligand ipTM) yields an AUC of 75.2, comparable to AF3 and Chai-1. Considering its early enrichment performance, Boltz-2 Ligand ipTM was the strongest predictor of pose accuracy based on normalized logAUC (27.1% above random, Figure 3a). In contrast, although Boltz-2 pIC_{50} showed poor overall discrimination, it overestimated its ability to enrich true positive poses at low false positive rates, despite having a weak early enrichment behavior. Since confidence scores correlated with RMSD, we also investigated how each method was able to predict poses and scores depending on the similarity of molecules to the training set (Figure 3—figure supplement 2). As before (Figure 2c), the performance of co-folding and docking scores did not show a strong dependence on novelty relative to the training set. This result suggests that co-folding can accurately reproduce poses of ligands dissimilar to those on which it trained, and that the model internal confidence scores reflect how well poses are reproduced.

Representative examples of correct and incorrect pose predictions across co-folding and docking methods are illustrated with AVI-4129 (PDB: 15DE) and AVI-4225 (PDB: 15DX), respectively (Figure 3b, c). Both compounds exhibited comparable experimental IC_{50} values 20 μ M for AVI-4225 (PDB: 15DX) and 42 μ M for AVI-4129 (PDB: 15DE) and similarity to molecules in the training set (Figure 3b). The molecule consistently predicted correctly AVI-4129 (PDB: 15DE) was a small, rigid fragment comprising a purine scaffold rigidly linked to a fluorobenzene moiety, favoring a well-defined binding mode. In contrast, the compound mispredicted by all methods AVI-4225 (PDB: 15DX) was larger and more conformationally flexible, with an alkyne linker prone to flipping, potentially altering the orientation of the fluorobenzyl ether substituent. Although this group packs against a crystal contact, the high predictive confidence/affinity scores across co-folding methods (AF3 L-pLDDT = 86.4, Chai-1 ipTM = 0.91) are likely explained by the adenosine mimic purine group where predictions recapitulate key interactions observed in the experimental structure.

To understand if prediction errors between methods are correlated, we directly compared the agreement between AF3 to Chai-1, Boltz-2, or DOCK3.7 predictions with the agreement between AF3 and experimental ground-truth (Figure 3—figure supplement 3). We observed that while there are a small number of structures where AF3 predicts much better than other methods (high AF3-other prediction RMSD, low AF3-experiment RMSD), most large errors are either missed in a similar way (low AF3-other prediction RMSD, but high AF3-experiment RMSD) or are both extreme mispredictions (high AF3-other prediction RMSD and AF3-experiment RMSD). Consistent with this result, in comparing RMSD relative to experimental structures for all methods (Figure 3—figure supplement 4), we observed that Chai-1 and AF3 exhibited the strongest agreement ($r = 0.72$), whereas poses from DOCK and Boltz-2 diverged more ($r = 0.45$ for AF3 vs. DOCK, and $r = 0.52$ for AF3 vs. Boltz-2). Collectively, these results suggest that the co-folding methods (AF3 and Chai-1) may be complementary to more physics-based approaches like DOCK3.7 as the errors are less correlated than observed for the different co-folding methods.

Second, we asked how well these scores correlate with experimentally measured potency (Figure 4a). The correlation was best for the Boltz-2 pIC_{50} ($r = 0.6$), significantly outperforming DOCK when tested with Friedman and Conover tests with Bonferroni corrections ($p = 0.0057$). Prior to any linear correction, the mean absolute error (MAE) in affinity prediction was ~ 1.6 kcal/mol (1.19 pIC units), with most affinities predicted as too potent. Empirically correcting with a linear fit reduced this

a) Co-folding scores as Mac1 affinity predictors


Friedman χ^2 (global comparison) = 16.68, $p = 2.39\text{e-}04$

Bonferroni-corrected Conover post-hoc tests (pairwise p-values for MAE):

Boltz-2 pIC_{50} vs DOCK score* ($p = 0.0057$)

AF3 L-pLDDT vs Boltz-2 pIC_{50} ($p = 0.089$)

AF3 L-pLDDT vs DOCK score ($p = 1.000$)

**b) Co-folded pose accuracy correlates with binding affinity**

Friedman χ^2 (global comparison) = 8.50, $p = 0.014$

Bonferroni-corrected Conover post-hoc tests (pairwise p-values for MAE):

Boltz-2 L-RMSD vs DOCK L-RMSD ($p = 0.023$)

AF3 L-RMSD vs Boltz-2 L-RMSD ($p = 1.000$)

AF3 L-RMSD vs DOCK L-RMSD ($p = 0.050$)




Figure 4. Co-folding scores and pose accuracy correlate with Mac1-binding affinity for 202 newly synthesized Mac1 ligands. (a) AF3 ligand pLDDT (L-pLDDT) and Boltz-2 predicted pIC_{50} correlate with experimental pIC_{50} more strongly than DOCK energy scores. (b) Lower co-folded ligand RMSDs associate with higher affinity, while docked pose RMSDs do not. Pearson's correlation coefficients and mean absolute errors (MAE) are shown. The baseline MAE was 0.68 pIC (or 0.93 kcal/mol), measured against a fixed average pIC_{50} ($\text{pIC}_{50} = 4.95$). Friedman tests with Conover post hoc comparisons and Bonferroni corrections are used to compare MAE between methods.

The online version of this article includes the following figure supplement(s) for figure 4:

Figure supplement 1. Actives from the benchmarked datasets are compared with resolved structures used to train models like AF3 and Chai-1.

error to ~ 0.7 kcal/mol (0.54 pIC units). To put these results in context, simply predicting the mean pIC_{50} from the dataset for every compound would yield a baseline MAE of 0.93 kcal/mol because of the relatively limited dynamic range of the dataset. The DOCK3.7 energy correlated more weakly with potency ($r = -0.225$), consistent with previous studies that have observed only weak correlations with affinity from virtual screening energies (Kinnings et al., 2011; Jones et al., 2021). The AF3 co-folding confidence metrics correlated slightly better than docking with potency ($r = 0.314$,

Figure 4a), but the post hoc Conover tests with Bonferroni corrections showed the difference was not statistically meaningful. Despite this, a correlation between pose accuracy and potency was more relevant with AF3 ($r = -0.233$, **Figure 4b**), as more accurate poses tend to be more confidently predicted (**Figure 3—figure supplement 1**). In contrast, there was no correlation between docking accuracy and potency (**Figure 4b**). Boltz pIC_{50} had the highest correlation between pose accuracy and affinity ($r = -0.279$). This result is interesting because in Boltz-2 affinity prediction and confidence are decoupled architecturally and have little correlation with each other across the Mac1 dataset (**Figure 3—figure supplement 1**).

Collectively, these results suggest that the categorization abilities of the co-folding confidence metrics, and their correlation with affinity, could mark an important advance. Still, the Mac1 set, as large as it is, represents only a single system with ligands that share a limited set of pharmacophores. We wanted to broaden testing to more targets and more diverse ligands, such as those encountered in virtual screening hit lists with their focus on new chemotype discovery.

Co-folding struggles to reliably re-rank docking hit lists

Encouraged by the ability of co-folding to reproduce ligand structures and affinities, we were interested to see if the methods could also improve hit identification from large library docking screens. In these screens, hundreds of millions to billions of molecules are docked against a target and scored for complementarity to a binding site. High-ranking molecules are then synthesized and tested for binding. While large library docking has successfully discovered potent novel chemotypes in recent studies (Lyu et al., 2019; Gorgulla et al., 2020; Stein et al., 2020; Sadybekov et al., 2022; Fink et al., 2022; Singh et al., 2023; Liu et al., 2025), the method is plagued by false positives (molecules that rank well but are inactive). We reasoned that if co-folding can predict the structures and activities of true ligands confidently, perhaps it could distinguish docking true binders from false positives.

To test this idea, we analyzed experimentally tested compounds from three large and heavily tested docking campaigns (**Supplementary file 5**): 506 experimentally tested compounds from a 490-million molecule screen against the σ_2 receptor (Alon et al., 2021), 1293 tested compounds from a 1.7-billion molecule screen against AmpC β -lactamase (Liu et al., 2025), and 541 tested compounds from a 138-million molecule screen against dopamine D4 receptors (Lyu et al., 2019) with AF3 (the best performing model from the Mac1 set), DOCK scoring function, and the Boltz-2 pIC_{50} score. These test sets differ from the Mac1 benchmark in three ways. First, they are a mix of molecules that would turn out to be true binders (247 in AmpC, 201 in σ_2 and 205 in dopamine D4) and high-ranking docking false positives as judged by measured potencies (worse than 400 μ M in AmpC, worse than 2.0 μ M against the σ_2 receptor, and worse than 7.8 μ M against the dopamine D4 receptor). For these molecules, the challenge is not only predicting the bound structures of true ligands, but distinguishing true ligands from decoy molecules that by many biophysical criteria fit the binding site well. Second, the docking molecules are more diverse than the Mac1 set (**Figure 1—figure supplement 2**), which largely falls into four substructure clusters representing conserved pharmacophores. The docking sets were selected, in the original studies, for their topological diversity, and their dissimilarity to previously known ligands. For each target, the ligands group into more than 10-times the number of clusters than the Mac1 dataset, consistent with a greater diversity. Indeed, the average pairwise ECFP4 Tc values within the three sets vary from 0.12 to 0.15, all well-below scaffold-hop cutoffs for this fingerprint (**Figure 1—figure supplement 2, Supplementary file 2**). Finally, the sets were all selected for novelty and there was little similarity between the ligands in the hit lists and known ligands, which ranged between 0.11 and 0.19 by Tc, and between 33 and 42 by MCS%. These ranges were lower than our prior Mac1 dataset (where Tc is distributed around 0.36, and MCS% is distributed around 60.1%), and the lack of similarity with the training set for target systems with hit lists could be explained by having relatively fewer ligand structures to compare to in the PDB (**Figure 4—figure supplement 1, Supplementary file 6**). These conditions, while representing new and difficult tests for co-folding, are the common case for structure-based ligand discovery campaigns. For validation, there are more measured IC_{50} values, but they are admittedly not as deeply supported by structural information as the Mac1 set.

We observed the strongest discrimination between active and inactive molecules in the σ_2 dataset, which contains 201 confirmed actives (apparent K_i values between 2.5 nM and 1.5 μ M) and 305 inactive false positives (**Figure 5a**). DOCK3.7 scores separated binders from non-binders (median




Figure 5. Classification of known active hits from docked false positives from initially screened docked molecules. The leftmost panel shows affinity scores from Boltz-2 (Boltz-2 pIC_{50}), the middle panel for AF3 L-pLDDT, and the right panel for the absolute value of DOCK energies (kcal/mol) from DOCK3.7. Docked hit lists from (a) σ_2 , (b) dopamine D4, (c) AmpC are grouped into known hits and non-hits, and score distributions are shown in violin plots. Kolmogorov–Smirnov test and its KS-statistic (KS stat) are used to indicate meaningful separation of two distribution curves. Asterisks indicate the p-value of KS statistic <0.001. Receiver operating characteristic (ROC) curves show an ability to distinguish known actives from the docked list with different scoring methods.

The online version of this article includes the following figure supplement(s) for figure 5:

Figure supplement 1. Discriminative power of co-folding and docking scores by correlation against experimentally measured pK_i across different targets: (a) σ_2 ($n = 201$ actives), (b) D4 ($n = 205$ actives), and (c) AmpC ($n = 247$ actives).

Figure supplement 2. Different methods of treating randomly assigned pK_i values for docked false positives in σ_2 , D4, and AmpC docked lists.

Figure supplement 3. Semi-logarithmic receiver operating characteristic (ROC) curves to differentiate actives over decoys using co-folding and docking scores for the three experimental benchmark datasets.

Figure supplement 4. Hit rate curves for co-folding and docking scores for the three experimental benchmark datasets.

Figure supplement 5. Correlation between DOCK scores and some co-folding scores (AF3 L-pLDDT and Boltz-2 pIC_{50}).

–55.0 vs. –42.5 kcal/mol; AUC = 78.8), though the correlation between docking score and pK_i for true ligands was relatively weak ($r = -0.211$). AF3 L-pLDDT showed a poorer ability to discriminate between true ligands and false positives (median 75.9 vs. 74.1, AUC = 56.1). Intriguingly, despite the worse discriminatory ability, the AF3 correlation with binders pK_i was not much worse than docking's ($r = 0.183$). Against the σ_2 set, Boltz-2 pIC_{50} performed better still, with high discriminatory power (median 6.70 vs. 5.38, AUC = 83.8) and the strongest correlation of affinity predictions against true positives ($r = 0.486$). Among active molecules, Boltz-2 pIC_{50} had an MAE of only 0.7 kcal/mol and a linear calibration reduced that error to 0.54 kcal/mol (**Figure 5—figure supplement 1**). This absolute estimation of IC_{50} was ~0.5 kcal/mol better than on the Mac1 dataset; however, it is restricted only to true positives, and the performance could be much worse if the false positives were accounted for in the comparison. Since inactive predictions were in the same pIC_{50} range as active molecules, as an optimistic limit of the performance that includes inactive molecules, we assigned an IC_{50} of 4 μ M, which is 2 $\times$ our threshold for declaring a non-binder, to all false positives and calculated the MAE. With this assumption, the MAE is 0.92 kcal/mol (0.56 kcal/mol after linear calibration). A less generous assumption would be to randomly assign an IC_{50} between the pK_i threshold (2 μ M) and $pK_i = 1$ (100 mM). Under this assumption, the MAE is 1.66 kcal/mol (1.55 kcal/mol after linear calibration) (**Figure 5—figure supplement 2**), which is in line with previously published estimates of the accuracy of this method (*Passaro et al., 2025*) and only a slight improvement over the MAE (1.84 kcal/mol) produced by predicting the dataset average affinity for each datapoint.

For AmpC (**Figure 5b**), a total of 1293 high-ranking molecules were experimentally tested for inhibition (*Liu et al., 2025*). Among these were 247 molecules that turned out to be true inhibitors. We observed that docking (AUC = 76.5, median –85.6 to –82.8 kcal/mol) was much stronger than Boltz-2 (AUC = 68.1, median 5.16 vs. 4.87) or AF3 (AUC = 60.5, median 76.6 vs. 72.0) in discriminating actives from inactives. The correlation of metrics to pK_i among actives was relatively weak for DOCK3.7, AF3, and Boltz-2 in this case. The Boltz-2 pIC_{50} values for AmpC were particularly high in absolute error 1.81 kcal/mol; however, linear correction reduced this error to only 0.42 kcal/mol. Including inactive molecules degraded performance to 2.33 kcal/mol (0.32 kcal/mol with linear correction) under the optimistic and 3.29 kcal/mol (1.04 kcal/mol with linear correction) under the random IC_{50} assumptions (**Figure 5—figure supplement 2**). Similarly, this represents a small reduction relative to the MAE (1.07 kcal/mol) produced by predicting the dataset average affinity for each datapoint.

The dopamine receptor also proved to be more challenging than σ_2 , particularly for AF3. With D4 with 205 confirmed actives (apparent K_i values between 3.1 nM and 7.8 μ M) and 336 inactive false positives, we observed that actives were distinguished from docked false positives by docking (–63.9 to –52.0 kcal/mol) and Boltz-2 (6.15 vs. 5.72) scores, both with AUC ~ 71. In contrast, AF3 L-pLDDT values showed little separation (65.4 vs. 66.2) and an AUC = 46.4, which is slightly worse than random (**Figure 5c**). There was also no clear evidence of early ligand enrichment with co-folding methods over docking methods from semi-logarithmic ROC plots across the three virtual screen targets (**Figure 5—figure supplement 3**). Among active molecules, the docking correlation was slightly stronger for D4 than it was for σ_2 ($r = -0.247$), but performance degraded for AF3 ($r = -0.109$, notably a positive correlation is expected) and Boltz-2 pIC_{50} ($r = 0.284$). Despite this reduced linear correlation for the Boltz-2 pIC_{50} , the MAE was 0.89 kcal/mol (0.55 kcal/mol after correction). The performance degraded to 1.15 kcal/mol (0.57 kcal/mol after linear correction) under the optimistic assumption of 2 $\times$ IC_{50} for all false positives and to 1.97 kcal/mol (1.4 kcal/mol after linear correction) under the randomly assigned assumption (**Figure 5—figure supplement 2**). This result is also a small improvement over the MAE (1.46 kcal/mol) produced by predicting the dataset average affinity for each datapoint.

An alternative way of looking at how well the scores can identify active ligands is to perform a hit enrichment analysis using normalized ranking scores (pProp, the negative logarithm of the percentage rank *Liu et al., 2025*: a molecule that ranks among the top 0.1% of the docked library would have a pProp of 3, one that ranked among the top 0.01% of the docked library would have a pProp of 4, and so forth). pProp assigned here is based on the ranking within the docked list, for direct comparison between scoring methods and largely confirmed the trends above (**Figure 5—figure supplement 4**). Both DOCK3.7 and Boltz-2 clearly enriched actives across most targets, with hit rates plateauing at 40–50% for high-ranking compounds. AF3 L-pLDDT showed meaningful and comparable enrichment only for AmpC, while providing no enrichment signal for σ_2 and D4. Overall, the success of the co-folding methods at predicting structures and correlating with affinities among true Mac1 inhibitors

was far less apparent in their ability to distinguish true- from false positives among docking hit lists, and to correlate with affinities among these highly diverse molecules. Indeed, the methods showed little advantage over docking, a method that is about four orders of magnitude faster. An important caveat is that the hit lists were composed of molecules prioritized by docking in the first place, giving it an advantage on these particular sets.

Discussion

A striking observation from this study was the ability of the co-folding methods to predict ligand structures across 551 previously unpublished Mac1 complexes, and for the co-folding confidence metrics to correlate with ligand affinities. By both criteria all co-folding methods were meaningfully better than docking, at least as represented by DOCK3.7. Affinity prediction by Boltz-2 also performed well with MAE that was ~ 0.2 kcal/mol better than a baseline expectation of simply using the mean IC_{50} value. This small improvement was a trend that also held up in the large docking sets as well. Conversely, when confronted with high-ranking molecules from large-library docking, a diverse mix of true and false positives with thousands of novel molecules prospectively tested against three unrelated targets, co-folding was not meaningfully better than docking at ligand categorization or at affinity prediction, indeed often worse. This despite the four log-orders greater speed of the docking and its well-known approximations and errors, long a topic of regret in the community (*Jorgensen and Tirado-Rives, 2005; Irwin and Shoichet, 2016*). These apparently contradictory results may be reconciled by the different tasks of predicting geometries for molecules that were all true ligands vs. that of categorizing binders and non-binders within a highly diverse docking hit list intentionally dissimilar to known ligands.

It may be that success and relative failure in these different tasks points also to a complementarity between physics-based docking and co-folding in the related, but different tasks of ligand discovery and optimization. The weak correlation between the methods (**Figure 5—figure supplement 5**) further suggests complementarity between deep-learning and physics-based approaches. Docking and related methods are designed for categorization among diverse, novel sets of candidate molecules, drawn from the vast new make-on-demand space (*Gentile et al., 2020; Radaeva et al., 2023; Sadybekov et al., 2022; Sadybekov and Katritch, 2023; Gorgulla et al., 2021; Lyu et al., 2019; Lyu et al., 2023; Liu et al., 2025; Liu et al., 2024*), whereas the co-folding methods, relying on learned patterns, may be relatively less effective. Where they have complementary strengths is in predicting the accurate geometries of ligands, once known, and their ranking within a series, both of which are crucial for affinity maturation once a series is discovered and with both of which docking has historically struggled. It is in these two linked but different aspects of ligand discovery and ligand optimization where the true complementarity between the methods may lie, at least for now.

Several caveats merit airing. The Mac1 benchmark, while diverse by ECFP4 Tc, is nevertheless dominated by four groups when clustered by MCS and benefits from a large training set of fragment molecules bound to the protein in the training set (*Schuller et al., 2021*). This may inflate apparent pose recovery rates, although Boltz-2 predicted poses were poorly recapitulated even with its much later training cutoff dates. Unexpectedly, even when the co-folding methods got important conformations of the enzyme binding site wrong, the ligands were still often posed correctly, speaking to a non-physical aspect of the ML methods that others have also described (*Masters et al., 2025; Škrinjar et al., 2025*). Even though the prospective docking hit lists included thousands of diverse molecules, they still represent only three targets, and so offer a limited view of molecular recognition. Finally, comparing co-folding to docking based on hit lists themselves selected by docking is arguably unfair to co-folding. Counter-balancing this is the inclusion, in each of the three hit lists, of molecules that had mediocre and poor docking scores intentionally selected to test the correlation between docking score and hit rate. Here too, the correlation between co-folding score and likelihood to bind, what we sometimes call a 'dock-response curve' was no better than docking's, often worse (**Figure 5—figure supplement 4**).

These caveats should not obscure the central lessons from this study. Against a benchmark of 551 topologically diverse Mac1 ligand-bound structures, determined to atomic resolution and largely dissimilar to structures already in the PDB, each of the co-folding methods predicted ligand poses and interactions with high-fidelity to the experimental structure. For Boltz-2, there was a strong and highly significant correlation between predicted affinity and experimental IC_{50} (compared to a modest

correlation for AF3 L-pLDDT and experimental IC₅₀). In both cases, co-folding performance was much better than that of DOCK3.7, although the correlation was only significantly different for Boltz-2 pIC₅₀ vs. experimental pIC₅₀. Conversely, when the more homogeneous benchmark of true ligands was replaced by a diverse set of docked molecules, composed of what would turn out to be hundreds of true ligands and thousands of false positives, the co-folding methods did no better than docking and often did worse. These observations are supported by a fascinating study on some of the same ligand sets as investigated here, using AlphaFold3, reaching similar conclusions (Menon *et al.*, 2025). The relatively strong performance by docking is surprising given its well-known weaknesses and its much reduced computational cost (and therefore greater speed). What emerges are approaches that may be complementary in their roles in ligand discovery (docking) and structure-based ligand optimization (co-folding), areas that will merit future research and application.

Methods

Mac1 X-ray crystallography

Ligands were soaked into SARS-CoV-2 Mac1 crystals at nominal concentrations ranging from 2.5–20 mM using the P4₃ crystal form (Schuller *et al.*, 2021). Data were collected and reduced as described previously (Gahbauer *et al.*, 2023; Correy *et al.*, 2025). We modeled ligands into PanDDA event maps (Pearce *et al.*, 2017) using ligand coordinates and restraints generated with phenix.elbow (Moriarty *et al.*, 2009). Ligands were modeled in COOT (Emsley *et al.*, 2010) and multi-state models were refined with phenix.refine (Liebschner *et al.*, 2019). Data collection and refinement statistics, and ligand formation for each Mac1-ligand complex, are presented in **Supplementary file 1**.

Ligand similarity, chemical diversity, and memorization analyses

To test whether pose-prediction success reflected latent memorization, we quantified chemical similarity between test and training ligands using two complementary metrics. First, extended-connectivity fingerprints (ECFP4) were computed for all molecules, and TC were used to identify, for each test ligand, the most similar compound present in the training corpus (Chai-1 entries deposited before 1 December 2021 and AlphaFold3 entries before 30 September 2021). The training date cutoff for Boltz-2 is 1 June 2023 (**Supplementary file 6**).

As another measure of structural similarity, the maximum common substructure percentage (MCS%) was calculated using the rdFMCS algorithm in RDKit, reporting the proportion of heavy atoms shared between two ligands. To evaluate chemical diversity within benchmark datasets, pairwise TC or MCS% values were computed across all ligands. Best-First Clustering based on these similarity metrics was used to estimate the number of unique chemotypes represented in each dataset. Principal component analysis (PCA) was used to visualize the resulting chemical space, and to identify an appropriate threshold for MCS-based scaffold grouping. Within the Mac1 benchmark set, applying an MCS cutoff of 35% yielded four distinct scaffolds, in agreement with PCA-derived cluster separation. The same clustering threshold was subsequently applied to define scaffold families for the AmpC, σ_2 , and D4 benchmark datasets.

Molecular docking of Mac1 compounds

The modeled orthosteric site of the Mac1 model was available from the PanDDA analysis group deposition in the Protein Data Bank (PDB), with PDB ID: 5SQW (Gahbauer *et al.*, 2023). We used this structure because the inhibitor (Z5014193706) was the most potent molecule with a structure determined around the same time as the ligands in this dataset were tested. This structure was used to generate matching spheres, which are later used by DOCK3.7 to pre-sample ligand conformations into binding sites by superimposing ligand atoms on spheres in the site. 45 matching spheres were used, and the docked ligands were scored by summing receptor–ligand electrostatic and van der Waals energies and corrected for context-dependent ligand desolvation. Receptor structures were protonated using Reduce. Partial charges from the united-atom AMBER force field were used for all receptor atoms. Potential energy grids for the different energy terms of the scoring function were precalculated using AMBER for the van der Waals term and the Poisson–Boltzmann method QNIFFT for electrostatics. Context-dependent ligand desolvation was calculated from the Generalized Born method. Ligands were protonated at pH 7.4 (ChemAxon), and each protomer was rendered into 3D

using Corina, followed by conformational sampling using Omega. For the docking campaign of 551 novel compounds, each molecule was sampled in about 2,719 orientations and, on average, 105 conformations. Overall, over 290 million complexes were sampled and scored.

Co-folding

All three co-folding methods (AlphaFold3, Chai-1, and Boltz-2) were available from the following sources:

- AlphaFold3: <https://github.com/google-deepmind/alphafold3>, Židek et al., 2026
- Chai-1: <https://github.com/chaidiscovery/chai-lab>, Rogozhnikov et al., 2026
- Boltz-2: <https://github.com/jwohlwend/boltz>, Wohlwend et al., 2026

Ligand SMILES and amino acid sequences for the target protein were used to perform co-folding on an array of NVIDIA A40 GPUs. Multiple Sequence Alignment was performed with the Jackhmmer module. Co-folding tools will generate predicted protein–ligand complex poses (mmCIF files), and different confidence metrics for these predictions. Chai-1 and AlphaFold3 generated global confidence scores (pTM and ipTM), while AlphaFold3 had local atom-specific metrics, such as pLDDT, PAE, and contact probabilities in large vectors. To put an emphasis on the prediction of ligands, we computed a ligand-centric pLDDT (L-pLDDT) by taking an arithmetic mean of pLDDT on ligand atoms, and a ligand-centric PAE (L-PAE) by taking an average of all PAE elements that contain ligand atoms. Minimum PAE (mPAE) was obtained by scanning all PAE elements that involve ligand atoms and finding the smallest error. Boltz-2 generated all the confidence metrics in Chai-1 and AlphaFold3, but its affinity module generated an output of the binding likelihood (probability metric) and the predicted binding affinity (Boltz-2 pIC₅₀) of compounds. Scripts to prepare input.json, .fasta files, post-process output files, are available in <https://github.com/jongbin99/cofold-em> (copy archived at Kim, 2026).

Benchmarking co-folding on experimentally validated compounds

With novel compounds discovered by fragment-based screening, 551 Mac1 complexes were resolved by the crystallography work. The same sequences were used to perform co-folding with AlphaFold3, Chai-1, and Boltz-2. For further sets of analyses with docked hit lists on targets from the Large-Scale Docking (source: lsd.docking.org) (Hall et al., 2025), only AlphaFold3 and Boltz-2 were used. For each system, we redefined the threshold for true hits and non-binders as described herein:

- AmpC β-lactamase: Among 1.7 billion docked molecules, the top scoring 1% were selected and compounds were sampled in different scoring bins for experimental validation ($n = 1293$) (Liu et al., 2025). We selected ones with apparent $K_i < 400 \mu\text{M}$ ($n = 247$) as true hits, and the rest were regarded as non-binders.
- σ_2 : True hits among these docked molecules were previously defined as displacing greater than 50% of ³H ditolylguanidine ([³H]DTG) to σ_2 (Alon et al., 2021), but in our work, the cutoff is adjusted to 25% displacement, giving a less stringent threshold to include more molecules into the space of true hits. Single concentration point measurement (at 1 μM) is converted to an apparent K_i value using the Cheng–Prusoff equation ($K_i = \frac{IC_{50} [S]}{1 + \frac{[S]}{K_m}}$), with pre-calculated values of $\frac{[S]}{K_m}$ based on radioligand displacement (Alon et al., 2021). This yields a non-binding cutoff of an apparent $K_i < 2.0 \mu\text{M}$.
- Dopamine D4: For the dopamine receptor D4, 138 million molecules were docked and then prioritized synthesizable high-ranking docked molecules ($n = 541$). True hits from these molecules were defined as displacing more than 50% of ³H-N-methylspiperone (Lyu et al., 2019), but similarly to σ_2 , we adjusted the cutoff to 25% displacement. Apparent K_i values are calculated similarly to the σ_2 case, but at 10 μM . This yields a non-binding cutoff of an apparent $K_i < 7.8 \mu\text{M}$.

Ligand pose evaluation via RMSD against crystal structures

Predicted ligand binding modes were benchmarked against experimentally determined reference poses from ultra-high-resolution (~1.0 Å) crystal structures. The Mac1 protein receptor was first superimposed by least-squares fitting of all Cα atoms using PyMol. The heavy-atom RMSD between the aligned co-folded ligand and its crystallographic counterpart was then calculated. Before RMSD

evaluation, both predicted and reference ligands were confirmed using a uniform RDKit-based sanitization pipeline (*Buttenschoen et al., 2024*). SMILES strings were converted to RDKit Mol objects, chemically validated (valence, aromaticity), and assigned bond orders from a canonical SMILES-derived template. We then check that ground-truth (experimental) and co-folded ligand coordinates for all molecules have valid chemistry as described with the SMILES-derived template. A substructure match then enabled a one-to-one correspondence of heavy atoms between the ground-truth and co-folded poses for comparison. The crystal ligand and co-folded ligand had the center-of-mass (COM) calculated separately, and measured the distance between each COM. If the distance exceeds 2.5 Å, we would expect the RMSD to be large, indicating that ligand poses are not modeled accurately in an orthosteric pocket (*Nittinger et al., 2025*). For all the predicted ligand poses, `rdMolAlign.CalcRMS` from `Rdtk.Chem` is used for calculating the ligand RMSD between co-folding and the ground-truth crystal pose.

Hit rate curves for different scoring methods

To get hit rate curves for co-folding and docking scores across three targets of docked hit lists (AmpC, σ_2 , and D4), molecules in the list are ordered by increasing score (for DOCK) and decreasing score (for L-pLDDT and Boltz-2 pIC_{50}). A rolling window of 100 was used for AmpC and σ_2 , and 50 for D4 (*Liu et al., 2025*), calculating the hit rate as a fraction of molecules with experimental binding affinity that match our definition for true hits described above (refer to Benchmarking co-folding on experimentally validated compounds in the Method section). The docked hit list is a mixture of known actives and non-binders that include false positives based on docking scores. For direct comparison of hit classification performance between DOCK energies, AF3 L-pLDDT, and Boltz-2 pIC_{50} , the `pProp` (the negative base 10 logarithm of the fractional rank) was assigned within the list, instead of the billion-scale full docking list. For each window, we used a 95% Wilson confidence interval with normal approximation for the hit rate.

Statistical analysis

Pearson's correlation coefficients were used throughout the paper to quantify linear associations between two variables in the study. We have also reported an MAE of datapoints with respect to the regression line for Boltz-2 affinity pIC_{50} against the measured pIC_{50} . The MAEs are also calculated for the perfect fit (Pearson $r = 1$) and for the fixed measured pIC_{50} (Pearson $r = 0$), to give an indication of how accurate the Boltz-2 scores are with respect to actual binding energies (in kcal/mol), in terms of correlations and absolute errors. Friedman and post hoc Conover tests were used to compare how each of AF3, Boltz-2, and DOCK3.7 score predicts for actual binding affinity, and Holm–Bonferroni correction for pairwise comparison p-values (*Ash et al., 2025*). The distribution of scores for known actives and docked false positives were compared using the Kolmogorov–Smirnov (KS) non-parametric statistical test, as shown in violin plots, and the KS statistic value indicates how distant the two distributions are, and p-values show the extent of statistically detectable differences. All analyses were performed using Python packages.

Acknowledgements

We acknowledge funding from NIH GM122481 (to BKS), GM145238 (to JSF), and U19AI171110 (to NJK). We thank Pat Walters for helpful comments. The synchrotron X-ray diffraction data used to determine Mac1 structures were collected at beamline 8.3.1 of the Advanced Light Source (ALS) and beamlines 9-2, 12-1, and 12-2 of the Stanford Synchrotron Radiation Lightsource (SSRL). The ALS, a U.S. DOE Office of Science User Facility under Contract No. DE-AC02-05CH11231, is supported in part by the ALS-ENABLE program funded by the NIH, National Institute of General Medical Sciences, grant P30 GM124169. Use of the SSRL, SLAC National Accelerator Laboratory, is supported by the U.S. Department of Energy, Office of Science, Office of Basic Energy Sciences under Contract No. DE-AC02-76SF00515. The SSRL Structural Molecular Biology Program is supported by the DOE Office of Biological and Environmental Research, and by the National Institutes of Health, National Institute of General Medical Sciences (P30GM133894).

Additional information

Competing interests

Nevan J Krogan: The Krogan laboratory has received research support from Vir Biotechnology, F Hoffmann-La Roche and Rezo Therapeutics; NJK has financially compensated consulting agreements with Maze Therapeutics and Interline Therapeutics; he is on the Board of Directors and is President of Rezo Therapeutics and is a shareholder in Tenaya Therapeutics, Maze Therapeutics, Rezo Therapeutics, GEn1E Lifesciences and Interline Therapeutics; and he has monetary and/or stock compensation with the following companies: GEn1E Lifesciences, Maze Therapeutics, Mreza Therapeutics, Rezo Therapeutics, and Tenaya Therapeutics. Adam R Renslo: Co-founder and holds equity in BridgeBio Oncology Therapeutics, Tatara Therapeutics and Elgia Therapeutics. Alan Ashworth: Co-founder of Tango Therapeutics, Azkarra Therapeutics and Kytarro; serves on boards and advisory boards as listed in the manuscript; receives research support from SPARC; and holds patents on the use of PARP inhibitors held jointly with AstraZeneca. Brian K Shoichet: Co-founder of BlueDolphin LLC, Epiodyne Inc and Deep Apple Therapeutics, Inc, and serves on the SRB of Genentech and the SAB of Schrodinger LLC and Vilya Therapeutics. James S Fraser: Holds equity in Relay Therapeutics, Impossible Foods, Arda Therapeutics, Profluent Bio, Interdict Bio (co-founder), Vilya Therapeutics and Edison Scientific Inc, and is a paid consultant for Relay Therapeutics, Profluent Bio, Vilya Therapeutics and Monimoi Therapeutics. The other authors declare that no competing interests exist.

Funding

Funder	Grant reference number	Author
National Institutes of Health	GM145238	James S Fraser
National Institutes of Health	GM122481	Brian K Shoichet
National Institutes of Health	U19AI171110	Nevan J Krogan

The funders had no role in study design, data collection, and interpretation, or the decision to submit the work for publication.

Author contributions

Jongbin Kim, Data curation, Formal analysis, Investigation, Writing – original draft, Writing – review and editing; Galen J Correy, Data curation, Validation, Investigation, Visualization, Methodology, Writing – original draft, Writing – review and editing; Brendan W Hall, Formal analysis, Investigation; Moira M Rachman, Formal analysis, Methodology, Writing – original draft, Writing – review and editing; Olivier Mailhot, Formal analysis, Investigation, Methodology; Takaya Togo, Ryan L Gonciarz, Priyadarshini Jaishankar, R Jeffrey Neitz, Eric R Hantz, Yagmur U Doruk, Maisie GV Stevens, Morgan E Diolaiti, Rashad Reid, Investigation, Methodology; Saumya Gopalkrishnan, Supervision, Project administration; Nevan J Krogan, Alan Ashworth, Supervision, Funding acquisition, Writing – review and editing; Adam R Renslo, Funding acquisition, Methodology, Writing – original draft, Writing – review and editing; Brian K Shoichet, James S Fraser, Conceptualization, Supervision, Funding acquisition, Writing – original draft, Writing – review and editing

Author ORCIDs

Jongbin Kim  <https://orcid.org/0009-0007-3405-7880>
Galen J Correy  <https://orcid.org/0000-0001-5155-7325>
Olivier Mailhot  <https://orcid.org/0000-0002-7533-4561>
Takaya Togo  <https://orcid.org/0000-0003-0243-0760>
Ryan L Gonciarz  <https://orcid.org/0000-0002-6600-7032>
Priyadarshini Jaishankar  <https://orcid.org/0009-0005-2013-8941>
R Jeffrey Neitz  <https://orcid.org/0000-0002-2247-9345>
Eric R Hantz  <https://orcid.org/0000-0002-2588-7619>
Maisie GV Stevens  <https://orcid.org/0009-0004-9732-0349>
Morgan E Diolaiti  <https://orcid.org/0000-0001-5900-3060>

Rashad Reid  <https://orcid.org/0009-0001-1513-8614>
Saumya Gopalkrishnan  <https://orcid.org/0009-0003-6713-4492>
Brian K Shoichet  <https://orcid.org/0000-0002-6098-7367>
James S Fraser  <https://orcid.org/0000-0002-5080-2859>

Peer review material

Version of Record:

Reviewer #1 (Public review): <https://doi.org/10.7554/eLife.110475.3.sa1>
Reviewer #3 (Public review): <https://doi.org/10.7554/eLife.110475.3.sa2>
Author response <https://doi.org/10.7554/eLife.110475.3.sa3>

Reviewed preprint v2

Reviewer #1 (Public review): <https://doi.org/10.7554/eLife.110475.2.sa2>
Reviewer #3 (Public review): <https://doi.org/10.7554/eLife.110475.2.sa1>
Author response: <https://doi.org/10.7554/eLife.110475.2.sa0>
eLife Assessment <https://doi.org/10.7554/eLife.110475.2.sa3>

Reviewed preprint v1

Reviewer #1 (Public review): <https://doi.org/10.7554/eLife.110475.1.sa2>
Reviewer #2 (Public review): <https://doi.org/10.7554/eLife.110475.1.sa1>
Reviewer #3 (Public review): <https://doi.org/10.7554/eLife.110475.1.sa0>
eLife Assessment <https://doi.org/10.7554/eLife.110475.1.sa3>

Additional files

Supplementary files

Supplementary file 1. X-ray data statistics.

Supplementary file 2. Pairwise similarity between each dataset (Mac1, AmpC, σ_2 , and D4) based on ECFP-4 TC and MCS%.

Supplementary file 3. Mac1 ligands with all co-folded, docked poses and scores.

Supplementary file 4. Ligand interactions (hydrogen bonds only) on ground-truth and AF3 predicted models influenced by residue flexibility.

Supplementary file 5. Data collection of all docked hits on AmpC, σ_2 , D4.

Supplementary file 6. PDB accessions by cutoff date.

MDAR checklist

Data availability

Coordinates and structure factors for the Mac1–ligand complexes determined in this work have been deposited in the Protein Data Bank. The 468 accession codes are listed in **Supplementary file 1**, together with the data collection and refinement statistics for each structure. A further 22 Mac1 complexes analyzed here were deposited as part of earlier studies. Their accession codes and source publications are listed in the same file. The Mac1 structure used to build the docking receptor is PDB 5SQW. Co-folded and docked poses, confidence metrics, and scores for every ligand are provided in **Supplementary file 3**. The experimentally tested compounds and measured potencies for the AmpC β -lactamase, σ_2 , and dopamine D4 docking campaigns are provided in **Supplementary file 5**. These data were taken from previously published screens (Alon et al., 2021; Liu et al., 2025; Lyu et al., 2019), and the corresponding hit lists are available from <https://lsd.docking.org/>. Scripts used to prepare the co-folding inputs and to analyze the resulting poses and scores are available at <https://github.com/jongbin99/cofold-em> (copy archived at Kim, 2026).

The following dataset was generated:

Author(s)	Year	Dataset title	Dataset URL	Database and Identifier
Correy GJ, Fraser JS	2022	PanDDA analysis group deposition -- Crystal structure of SARS-CoV-2 NSP3 macrodomain in complex with Z5014193706 - (R,R) and (S,S) isomers	https://doi.org/10.2210/pdb5SQW/pdb	Worldwide Protein Data Bank, 10.2210/pdb5SQW/pdb

References

- Abramson J**, Adler J, Dunger J, Evans R, Green T, Pritzel A, Ronneberger O, Willmore L, Ballard AJ, Bambrick J, Bodenstein SW, Evans DA, Hung C-C, O'Neill M, Reiman D, Tunyasuvunakool K, Wu Z, Žemgulytė A, Arvaniti E, Beattie C, et al. 2024. Accurate structure prediction of biomolecular interactions with AlphaFold 3. *Nature* **630**:493–500. DOI: <https://doi.org/10.1038/s41586-024-07487-w>, PMID: 38718835
- Ackloo S**, Al-Awar R, Amaro RE, Arrowsmith CH, Azevedo H, Batey RA, Bengio Y, Betz UAK, Bologa CG, Chodera JD, Cornell WD, Dunham I, Ecker GF, Edfeldt K, Edwards AM, Gilson MK, Gordijo CR, Hessler G, Hillisch A, Hogner A, et al. 2022. CACHE (Critical Assessment of Computational Hit-finding Experiments): a public-private partnership benchmarking initiative to enable the development of computational methods for hit-finding. *Nature Reviews Chemistry* **6**:287–295. DOI: <https://doi.org/10.1038/s41570-022-00363-z>, PMID: 35783295
- Alon A**, Lyu J, Braz JM, Tummino TA, Craik V, O'Meara MJ, Webb CM, Radchenko DS, Moroz YS, Huang X-P, Liu Y, Roth BL, Irwin JJ, Basbaum AI, Shoichet BK, Kruse AC. 2021. Structures of the σ_2 receptor enable docking for bioactive ligand discovery. *Nature* **600**:759–764. DOI: <https://doi.org/10.1038/s41586-021-04175-x>, PMID: 34880501
- Ash JR**, Wognum C, Rodríguez-Pérez R, Aldeghi M, Cheng AC, Clevert DA, Engkvist O, Fang C, Price DJ, Hughes-Oliver JM, Walters WP. 2025. Practically significant method comparison protocols for machine learning in small molecule drug discovery. *Journal of Chemical Information and Modeling* **65**:9398–9411. DOI: <https://doi.org/10.1021/acs.jcim.5c01609>, PMID: 40932128
- Ashizawa R**, Kotelnikov S, Khan O, Li SX, Glukhov E, Cao X, Lazou M, Bekar-Cesaretli A, Hailegeorgis D, Averkava V, Zhu Y, Jones G, Yu H, Kalitin D, Stepanenko D, Koirala K, Patsahan T, Beglov D, Lukin M, Joseph-McCarthy D, et al. 2026. Modeling protein-protein and protein-ligand interactions by the ClusPro team in CASP16. *Proteins* **94**:183–191. DOI: <https://doi.org/10.1002/prot.70066>, PMID: 41115690
- Backus KM**, Correia BE, Lum KM, Forli S, Horning BD, González-Páez GE, Chatterjee S, Lanning BR, Teijaro JR, Olson AJ, Wolan DW, Cravatt BF. 2016. Proteome-wide covalent ligand discovery in native biological systems. *Nature* **534**:570–574. DOI: <https://doi.org/10.1038/nature18002>, PMID: 27309814
- Baek M**, DiMaio F, Anishchenko I, Dauparas J, Ovchinnikov S, Lee GR, Wang J, Cong Q, Kinch LN, Schaeffer RD, Millán C, Park H, Adams C, Glassman CR, DeGiovanni A, Pereira JH, Rodrigues AV, van Dijk AA, Ebrecht AC, Opperman DJ, et al. 2021. Accurate prediction of protein structures and interactions using a three-track neural network. *Science* **373**:871–876. DOI: <https://doi.org/10.1126/science.abj8754>, PMID: 34282049
- Boitreaud J**, Dent J, McPartlon M. 2024. Chai-1: Decoding the Molecular Interactions of Life. *bioRxiv*. DOI: <https://doi.org/10.1101/2024.10.10.615955>
- Brocchiacono M**, Wellnitz J, Popov KI, Tropsha A. 2025. Look mom, no experimental data! Learning to score protein-ligand interactions from simulations. *arXiv*. DOI: <https://doi.org/10.48550/arXiv.2506.00593>
- Buttenschoen M**, Morris GM, Deane CM. 2024. PoseBusters: AI-based docking methods fail to generate physically valid poses or generalise to novel sequences. *Chemical Science* **15**:3130–3139. DOI: <https://doi.org/10.1039/D3SC04185A>
- Cabeza de Vaca I**, Trapkov B, Shen L, Vo DD, Zhang X, Yang Y, Pezeshki M, Zhang X, Bällgren F, Saleh A, Tarnovskiy AV, Radchenko DS, Moroz YS, Bräuner-Osborne H, Svenningsson P, Kihlberg J, Liu ZJ, Hauser AS, Carlsson J. 2025. Ultra-large virtual screening unveils potent agonists of the neuromodulatory orphan receptor GPR139. *Nature Communications* **17**:129. DOI: <https://doi.org/10.1038/s41467-025-66845-y>, PMID: 41365886
- Chen D**, Vollmar M, Rossi MN, Phillips C, Kraehenbuehl R, Slade D, Mehrotra PV, von Delft F, Crosthwaite SK, Gileadi O, Denu JM, Ahel I. 2011. Identification of macrodomain proteins as novel O-acetyl-ADP-ribose deacetylases. *The Journal of Biological Chemistry* **286**:13261–13271. DOI: <https://doi.org/10.1074/jbc.M110.206771>, PMID: 21257746
- Correy GJ**, Rachman MM, Togo T, Gahbauer S, Doruk YU, Stevens MGV, Jaishankar P, Kelley B, Goldman B, Schmidt M, Kramer T, Radchenko DS, Moroz YS, Ashworth A, Riley P, Shoichet BK, Renslo AR, Walters WP, Fraser JS. 2025. Exploration of structure-activity relationships for the SARS-CoV-2 macrodomain from shape-based fragment linking and active learning. *Science Advances* **11**:eads7187. DOI: <https://doi.org/10.1126/sciadv.ads7187>, PMID: 40435250
- Cramer RD**, Jilek RJ, Guessregen S, Clark SJ, Wendt B, Clark RD. 2004. "Lead hopping". Validation of topomer similarity as a superior predictor of similar biological activities. *Journal of Medicinal Chemistry* **47**:6777–6791. DOI: <https://doi.org/10.1021/jm049501b>, PMID: 15615527

- Díaz-Holguín A**, Saarinen M, Vo DD, Sturchio A, Branzell N, Cabeza de Vaca I, Hu H, Mitjavila-Domènech N, Lindqvist A, Baranczewski P, Millan MJ, Yang Y, Carlsson J, Svenningsson P. 2024. AlphaFold accelerated discovery of psychotropic agonists targeting the trace amine-associated receptor 1. *Science Advances* **10**:eadn1524. DOI: <https://doi.org/10.1126/sciadv.adn1524>, PMID: 39110804
- Emsley P**, Lohkamp B, Scott WG, Cowtan K. 2010. Features and development of Coot. *Acta Crystallographica. Section D, Biological Crystallography* **66**:486–501. DOI: <https://doi.org/10.1107/S0907444910007493>, PMID: 20383002
- Erlanson DA**, Burley SK, Fearon D, Fraser JS, Kreitler D, Nonato MC, Sakai N, Wollenhaupt J, Weiss MS. 2025. Where and how to house big data on small fragments. *Nature Communications* **16**:4179. DOI: <https://doi.org/10.1038/s41467-025-59233-z>, PMID: 40325009
- Errington D**, Schneider C, Bouysset C, Dreyer FA. 2025. Assessing interaction recovery of predicted protein-ligand poses. *Journal of Cheminformatics* **17**:76. DOI: <https://doi.org/10.1186/s13321-025-01011-6>, PMID: 40389970
- Fink EA**, Xu J, Hübner H, Braz JM, Seemann P, Avet C, Craik V, Weikert D, Schmidt MF, Webb CM, Tolmachova NA, Moroz YS, Huang XP, Kalyanaraman C, Gahbauer S, Chen G, Liu Z, Jacobson MP, Irwin JJ, Bouvier M, et al. 2022. Structure-based discovery of nonopioid analgesics acting through the α_{2A} -adrenergic receptor. *Science* **377**:eabn7065. DOI: <https://doi.org/10.1126/science.abn7065>, PMID: 36173843
- Flowers J**, Echols N, Correy G, Jaishankar P, Togo T, Renslo AR, van den Bedem H, Fraser JS, Wankowicz SA. 2025. Expanding automated multiconformer ligand modeling to macrocycles and fragments. *bioRxiv*. DOI: <https://doi.org/10.1101/2024.09.20.613996>, PMID: 39386683
- Gahbauer S**, Correy GJ, Schuller M, Ferla MP, Doruk YU, Rachman M, Wu T, Diolaiti M, Wang S, Neitz RJ, Fearon D, Radchenko DS, Moroz YS, Irwin JJ, Renslo AR, Taylor JC, Gestwicki JE, von Delft F, Ashworth A, Ahel I, et al. 2023. Iterative computational design and crystallographic screening identifies potent inhibitors targeting the Nsp3 macrodomain of SARS-CoV-2. *PNAS* **120**:e2212931120. DOI: <https://doi.org/10.1073/pnas.2212931120>, PMID: 36598939
- Gathiaka S**, Liu S, Chiu M, Yang H, Stuckey JA, Kang YN, Delproposto J, Kubish G, Dunbar JB, Carlson HA, Burley SK, Walters WP, Amaro RE, Feher VA, Gilson MK. 2016. D3R grand challenge 2015: Evaluation of protein-ligand pose and affinity predictions. *Journal of Computer-Aided Molecular Design* **30**:651–668. DOI: <https://doi.org/10.1007/s10822-016-9946-8>, PMID: 27696240
- Gentile F**, Agrawal V, Hsing M, Ton AT, Ban F, Norinder U, Gleave ME, Cherkasov A. 2020. Deep docking: a deep learning platform for augmentation of structure based drug discovery. *ACS Central Science* **6**:939–949. DOI: <https://doi.org/10.1021/acscentsci.0c00229>, PMID: 32607441
- Gorgulla C**, Boeszoermenyi A, Wang ZF, Fischer PD, Coote PW, Padmanabha Das KM, Malets YS, Radchenko DS, Moroz YS, Scott DA, Fackeldey K, Hoffmann M, Iavniuk I, Wagner G, Arthanari H. 2020. An open-source drug discovery platform enables ultra-large virtual screens. *Nature* **580**:663–668. DOI: <https://doi.org/10.1038/s41586-020-2117-z>, PMID: 32152607
- Gorgulla C**, Padmanabha Das KM, Leigh KE, Cesugli M, Fischer PD, Wang ZF, Tesseyre G, Pandita S, Shnapir A, Calderaio A, Gechev M, Rose A, Lewis N, Hutcheson C, Yaffe E, Luxemburg R, Herce HD, Durmaz V, Halazonetis TD, Fackeldey K, et al. 2021. A multi-pronged approach targeting SARS-CoV-2 proteins using ultra-large virtual screening. *iScience* **24**:102021. DOI: <https://doi.org/10.1016/j.isci.2020.102021>, PMID: 33426509
- Goverde CA**, Wolf B, Khakzad H, Rosset S, Correia BE. 2023. De novo protein design by inversion of the AlphaFold structure prediction network. *Protein Science* **32**:e4653. DOI: <https://doi.org/10.1002/pro.4653>, PMID: 37165539
- Hall BW**, Tummino TA, Tang K, Mailhot O, Castanon M, Irwin JJ, Shoichet BK. 2025. A database for large-scale docking and experimental results. *Journal of Chemical Information and Modeling* **65**:4458–4467. DOI: <https://doi.org/10.1021/acs.jcim.5c00394>, PMID: 40273444
- He XH**, Li JR, Shen SY, Xu HE. 2025. AlphaFold3 versus experimental structures: assessment of the accuracy in ligand-bound G protein-coupled receptors. *Acta Pharmacologica Sinica* **46**:1111–1122. DOI: <https://doi.org/10.1038/s41401-024-01429-y>, PMID: 39643640
- Hert J**, Keiser MJ, Irwin JJ, Oprea TI, Shoichet BK. 2008. Quantifying the relationships among drug classes. *Journal of Chemical Information and Modeling* **48**:755–765. DOI: <https://doi.org/10.1021/ci8000259>, PMID: 18335977
- Holcomb M**, Chang YT, Goodsell DS, Forli S. 2023. Evaluation of AlphaFold2 structures as docking targets. *Protein Science* **32**:e4530. DOI: <https://doi.org/10.1002/pro.4530>, PMID: 36479776
- Huang N**, Kalyanaraman C, Irwin JJ, Jacobson MP. 2006. Physics-based scoring of protein-ligand complexes: enrichment of known inhibitors in large-scale virtual screening. *Journal of Chemical Information and Modeling* **46**:243–253. DOI: <https://doi.org/10.1021/ci0502855>, PMID: 16426060
- Irwin JJ**, Shoichet BK. 2016. Docking screens for novel ligands conferring new biology. *Journal of Medicinal Chemistry* **59**:4103–4120. DOI: <https://doi.org/10.1021/acs.jmedchem.5b02008>, PMID: 26913380
- Jaishankar P**, Correy GJ, Matsui Y, Togo T, Rachman MM, Stevens MGV, Hantz ER, Zheng J, Diolaiti ME, Montano M, Taha TY, Rosecrans J, Pampel J, Krogan NJ, Shoichet BK, Ashworth A, Ott M, Fraser JS, Renslo AR. 2025. Discovery of AVI-6451, a potent and selective inhibitor of the SARS-CoV-2 ADP-ribosylhydrolase mac1 with oral efficacy in vivo. *bioRxiv*. DOI: <https://doi.org/10.1101/2025.10.11.681833>
- Jones D**, Kim H, Zhang X, Zemla A, Stevenson G, Bennett WFD, Kirshner D, Wong SE, Lightstone FC, Allen JE. 2021. Improved protein-ligand binding affinity prediction with structure-based deep fusion inference. *Journal*

- of *Chemical Information and Modeling* **61**:1583–1592. DOI: <https://doi.org/10.1021/acs.jcim.0c01306>, PMID: 33754707
- Jorgensen WL, Tirado-Rives J. 2005. Molecular modeling of organic and biomolecular systems using BOSS and MCPRO. *Journal of Computational Chemistry* **26**:1689–1700. DOI: <https://doi.org/10.1002/jcc.20297>, PMID: 16200637
- Jumper J, Evans R, Pritzel A, Green T, Figurnov M, Ronneberger O, Tunyasuvunakool K, Bates R, Židek A, Potapenko A, Bridgland A, Meyer C, Kohl SAA, Ballard AJ, Cowie A, Romero-Paredes B, Nikolov S, Jain R, Adler J, Back T, et al. 2021. Highly accurate protein structure prediction with AlphaFold. *Nature* **596**:583–589. DOI: <https://doi.org/10.1038/s41586-021-03819-2>, PMID: 34265844
- Karelina M, Noh JJ, Dror RO. 2023. How accurately can one predict drug binding modes using AlphaFold models? *eLife* **12**:RP89386. DOI: <https://doi.org/10.7554/eLife.89386>, PMID: 38131311
- Kim J. 2026. Cofolding. swl:1:rev:921b74cfb6fd83e5e1953a6c971e265111f5e252. Software Heritage. <https://archive.softwareheritage.org/swl:1:dir:1fa2d87cf999afac75ffd046d629c1299498516e;origin=https://github.com/jongbin99/cofold-em;visit=swl:1:snp:db81daea6dcc0700570f6e64ccc8f8a839f8f523;anchor=swl:1:rev:921b74cfb6fd83e5e1953a6c971e265111f5e252>
- Kinnings SL, Liu N, Tonge PJ, Jackson RM, Xie L, Bourne PE. 2011. Correction to “Machine learning-based method to improve docking scoring functions and its application to drug repurposing”. *Journal of Chemical Information and Modeling* **51**:1195–1197. DOI: <https://doi.org/10.1021/ci2001346>, PMID: 21526828
- Kryshtafovych A, Schwede T, Topf M, Fidelis K, Mout R. 2021. Critical assessment of methods of protein structure prediction (CASP)-Round XIV. *Proteins* **89**:1607–1617. DOI: <https://doi.org/10.1002/prot.26237>, PMID: 34533838
- Liebschner D, Afonine PV, Baker ML, Bunkóczi G, Chen VB, Croll TI, Hintze B, Hung LW, Jain S, McCoy AJ, Moriarty NW, Oeffner RD, Poon BK, Prisant MG, Read RJ, Richardson JS, Richardson DC, Sammito MD, Sobolev OV, Stockwell DH, et al. 2019. Macromolecular structure determination using X-rays, neutrons and electrons: recent developments in Phenix. *Acta Crystallographica. Section D, Structural Biology* **75**:861–877. DOI: <https://doi.org/10.1107/S2059798319011471>, PMID: 31588918
- Liu F, Wu CG, Tu CL, Glenn I, Meyerowitz J, Kaplan AL, Lyu J, Cheng Z, Tarkhanova OO, Moroz YS, Irwin JJ, Chang W, Shoichet BK, Skiniotis G. 2024. Large library docking identifies positive allosteric modulators of the calcium-sensing receptor. *Science* **385**:1868. DOI: <https://doi.org/10.1126/science.adb1868>, PMID: 39298584
- Liu F, Mailhot O, Glenn IS, Vigneron SF, Bassim V, Xu X, Fonseca-Valencia K, Smith MS, Radchenko DS, Fraser JS, Moroz YS, Irwin JJ, Shoichet BK. 2025. The impact of library size and scale of testing on virtual screening. *Nature Chemical Biology* **21**:1039–1045. DOI: <https://doi.org/10.1038/s41589-024-01797-w>, PMID: 39753705
- Lyu J, Wang S, Balus TE, Singh I, Levit A, Moroz YS, O'Meara MJ, Che T, Algaa E, Tolmachova K, Tolmachev AA, Shoichet BK, Roth BL, Irwin JJ. 2019. Ultra-large library docking for discovering new chemotypes. *Nature* **566**:224–229. DOI: <https://doi.org/10.1038/s41586-019-0917-9>, PMID: 30728502
- Lyu J, Irwin JJ, Shoichet BK. 2023. Modeling the expansion of virtual screening libraries. *Nature Chemical Biology* **19**:712–718. DOI: <https://doi.org/10.1038/s41589-022-01234-w>, PMID: 36646956
- Lyu J, Kopolka N, Gumpert R, Alon A, Wang L, Jain MK, Barros-Álvarez X, Sakamoto K, Kim Y, DiBerto J, Kim K, Glenn IS, Tummino TA, Huang S, Irwin JJ, Tarkhanova OO, Moroz Y, Skiniotis G, Kruse AC, Shoichet BK, et al. 2024. AlphaFold2 structures guide prospective ligand discovery. *Science* **384**:eadn6354. DOI: <https://doi.org/10.1126/science.adn6354>, PMID: 38753765
- Masters MR, Mahmoud AH, Lill MA. 2025. Investigating whether deep learning models for co-folding learn the physics of protein-ligand interactions. *Nature Communications* **16**:8854. DOI: <https://doi.org/10.1038/s41467-025-63947-5>, PMID: 41053181
- Menon KM, Datasam A, Chen G, Bryant C, Deng Z, Lam B, Yang C, Barcelos D, Liu F, Alon A, Lyu J. 2025. AlphaFold3 for structure-guided ligand discovery. *bioRxiv*. DOI: <https://doi.org/10.64898/2025.12.04.692352>
- Michino M, Vendome J, Kufareva I. 2025. AI meets physics in computational structure-based drug discovery for GPCRs. *NPJ Drug Discovery* **2**:16. DOI: <https://doi.org/10.1038/s44386-025-00019-0>, PMID: 40621551
- Moriarty NW, Grosse-Kunstleve RW, Adams PD. 2009. electronic ligand builder and optimization workbench (eLBOW): a tool for ligand coordinate and restraint generation. *Acta Crystallographica. Section D, Biological Crystallography* **65**:1074–1080. DOI: <https://doi.org/10.1107/S0907444909029436>, PMID: 19770504
- Muchmore SW, Debe DA, Metz JT, Brown SP, Martin YC, Hajduk PJ. 2008. Application of belief theory to similarity data fusion for use in analog searching and lead hopping. *Journal of Chemical Information and Modeling* **48**:941–948. DOI: <https://doi.org/10.1021/ci7004498>, PMID: 18416545
- Nittinger E, Yolk Ö, Tibo A, Olanders G, Tyrchan C. 2025. Co-folding, the future of docking – prediction of allosteric and orthosteric ligands. *Artificial Intelligence in the Life Sciences* **8**:100136. DOI: <https://doi.org/10.1016/j.aillsci.2025.100136>
- Passaro S, Corso G, Wohlwend J, Reveiz M, Thaler S, Somnath VR, Getz N, Portnoi T, Roy J, Stark H, Kwabi-Addo D, Beaini D, Jaakkola T, Barzilay R. 2025. Boltz-2: towards accurate and efficient binding affinity prediction. *bioRxiv*. DOI: <https://doi.org/10.1101/2025.06.14.659707>, PMID: 40667369
- Pearce NM, Krojer T, Bradley AR, Collins P, Nowak RP, Talon R, Marsden BD, Kelm S, Shi J, Deane CM, von Delft F. 2017. A multi-crystal method for extracting obscured crystallographic states from conventionally uninterpretable electron density. *Nature Communications* **8**:15123. DOI: <https://doi.org/10.1038/ncomms15123>, PMID: 28436492
- Radaeva M, Morin H, Pandey M, Ban F, Guo M, LeBlanc E, Lallous N, Cherkasov A. 2023. Novel Inhibitors of androgen receptor's DNA binding domain identified using an ultra-large virtual screening. *Molecular Informatics* **42**:e2300026. DOI: <https://doi.org/10.1002/minf.202300026>, PMID: 37193651

- Rogozhnikov A**, Wu K, Dent J. 2026. Chai-lab. 66c38d1. Github. <https://github.com/chaidiscovery/chai-lab>
- Sadybekov AA**, Sadybekov AV, Liu Y, Iliopoulos-Tsoutsouvas C, Huang XP, Pickett J, Houser B, Patel N, Tran NK, Tong F, Zvonok N, Jain MK, Savych O, Radchenko DS, Nikas SP, Petasis NA, Moroz YS, Roth BL, Makriyannis A, Katritch V. 2022. Synthon-based ligand discovery in virtual libraries of over 11 billion compounds. *Nature* **601**:452–459. DOI: <https://doi.org/10.1038/s41586-021-04220-9>, PMID: 34912117
- Sadybekov AV**, Katritch V. 2023. Computational approaches streamlining drug discovery. *Nature* **616**:673–685. DOI: <https://doi.org/10.1038/s41586-023-05905-z>, PMID: 37100941
- Scardino V**, Di Filippo JL, Cavasotto CN. 2023. How good are AlphaFold models for docking-based virtual screening? *iScience* **26**:105920. DOI: <https://doi.org/10.1016/j.isci.2022.105920>, PMID: 36686396
- Schuller M**, Correy GJ, Gahbauer S, Fearon D, Wu T, Díaz RE, Young ID, Carvalho Martins L, Smith DH, Schulze-Gahmen U, Owens TW, Deshpande I, Merz GE, Thwin AC, Biel JT, Peters JK, Moritz M, Herrera N, Kratochvil HT, Aimon A, et al. 2021. Fragment binding to the Nsp3 macrodomain of SARS-CoV-2 identified through crystallographic screening and computational docking. *Science Advances* **7**:eabf8711. DOI: <https://doi.org/10.1126/sciadv.abf8711>, PMID: 33853786
- Seok C**, Tiwary P. 2025. Editorial overview: artificial intelligence methodologies in structural biology. *Current Opinion in Structural Biology* **95**:103156. DOI: <https://doi.org/10.1016/j.sbi.2025.103156>, PMID: 40946380
- Shoichet BK**. 2004. Virtual screening of chemical libraries. *Nature* **432**:862–865. DOI: <https://doi.org/10.1038/nature03197>, PMID: 15602552
- Singh I**, Seth A, Billesbølle CB, Braz J, Rodriguiz RM, Roy K, Bekele B, Craik V, Huang X-P, Boytsov D, Pogorelov VM, Lak P, O'Donnell H, Sandtner W, Irwin JJ, Roth BL, Basbaum AI, Wetsel WC, Manglik A, Shoichet BK, et al. 2023. Structure-based discovery of conformationally selective inhibitors of the serotonin transporter. *Cell* **186**:2160–2175. DOI: <https://doi.org/10.1016/j.cell.2023.04.010>, PMID: 37137306
- Škrinjar P**, Eberhardt J, Tauriello G, Schwede T, Durairaj J. 2025. Have protein-ligand cofolding methods moved beyond memorisation?. *bioRxiv*. DOI: <https://doi.org/10.1101/2025.02.03.636309>
- Stein RM**, Kang HJ, McCorvy JD, Glatfelter GC, Jones AJ, Che T, Slocum S, Huang XP, Savych O, Moroz YS, Stauch B, Johansson LC, Cherezov V, Kenakin T, Irwin JJ, Shoichet BK, Roth BL, Dubocovich ML. 2020. Virtual discovery of melatonin receptor ligands to modulate circadian rhythms. *Nature* **579**:609–614. DOI: <https://doi.org/10.1038/s41586-020-2027-0>, PMID: 32040955
- Suryawanshi RK**, Jaishankar P, Correy GJ, Rachman MM, O'Leary PC, Taha TY, Zapatero-Belinchón FJ, McCavitt-Malvido M, Doruk YU, Stevens MG, Diolaiti ME, Jogalekar MP, Richards AL, Montano M, Rosecrans J, Matthay M, Togo T, Gonciarz RL, Gopalkrishnan S, Neitz RJ, et al. 2025. The Mac1 ADP-ribosylhydrolase is a therapeutic target for SARS-CoV-2. *eLife* **14**:3484. DOI: <https://doi.org/10.7554/eLife.103484.1>
- Tirado-Rives J**, Jorgensen WL. 2006. Contribution of conformer focusing to the uncertainty in predicting free energies for protein-ligand binding. *Journal of Medicinal Chemistry* **49**:5880–5884. DOI: <https://doi.org/10.1021/jm060763i>, PMID: 17004703
- Wierbowski SD**, Wingert BM, Zheng J, Camacho CJ. 2020. Cross-docking benchmark for automated pose and ranking prediction of ligand binding. *Protein Science* **29**:298–305. DOI: <https://doi.org/10.1002/pro.3784>, PMID: 31721338
- Wohlwend J**, Corso G, Passaro S, Getz N, Reveiz M, Leidal K, Swiderski W, Atkinson L, Portnoi T, Chinn I, Silterra J, Jaakkola T, Barzilay R. 2025. Boltz-1 democratizing biomolecular interaction modeling. *bioRxiv*. DOI: <https://doi.org/10.1101/2024.11.19.624167>, PMID: 39605745
- Wohlwend J**, Corso G, Anastopoulos I, Getz N. 2026. Boltz. b1ebfc4. Github. <https://github.com/jwohlwend/boltz>
- Xie S**, Zhu H, Huang N. 2025. AI-designed molecules in drug discovery, structural novelty evaluation, and implications. *Journal of Chemical Information and Modeling* **65**:8924–8933. DOI: <https://doi.org/10.1021/acs.jcim.5c00921>, PMID: 40820734
- Xu S**, Feng Q, Qiao L, Wu H, Shen T, Cheng Y, Zheng S, Sun S. 2025. Foldbench: an all-atom benchmark for biomolecular structure prediction. *bioRxiv*. DOI: <https://doi.org/10.1101/2025.05.22.655600>
- Yang X**, Ma Y, Li Y, Dong Y, Yu LL, Wang H, Guo L, Wu C, Yu X, Liu X. 2020. Molecular basis for the MacroD1-mediated hydrolysis of ADP-ribosylation. *DNA Repair* **94**:102899. DOI: <https://doi.org/10.1016/j.dnarep.2020.102899>, PMID: 32683309
- Židek A**, Joo H, Kelly J, Elsheikh AE, Hawkins P, andycowie. 2026. Alphafold3. b096023. Github. <https://github.com/google-deepmind/alphafold3>