Deep contrastive learning enables genome-wide virtual screening

Yinjun Jia[1,2,3,4,7,*], Bowen Gao[1,5,*], Jiaxin Tan[2,4,6,7,*], Xin Hong[1,*], Wenyu Zhu[1], Haichuan Tan[1,5], Yuan Xiao[2,4,6,7], Yanwen Huang[8], Yue Jin[2,3,4,7], Yafei Yuan[2,4,6,7], Jiekang Tian[9], Weiying Ma[1], Yaqin Zhang[1], Chuangye Yan[2,4,6,7,#], Wei Zhang[2,3,4,7,#], Yanyan Lan[1,6,10,#]

[1] Institute for AI Industry Research (AIR), Tsinghua University, Beijing, China.

[2] School of Life Sciences, Tsinghua University, Beijing, China.

[3] IDG/McGovern Institute for Brain Research, Tsinghua University, Beijing, China.

[4] Tsinghua-Peking Center for Life Sciences, Beijing, China.

[5] Department of Computer Science and Technology, Tsinghua University, Beijing, China.

[6] Beijing Frontier Research Center for Biological Structure, Tsinghua University, Beijing, China

[7] State Key Laboratory of Membrane Biology, Tsinghua University, Beijing, China

[8] State Key Laboratory of Natural and Biomimetic Drugs, School of Pharmaceutical Sciences, Peking University, Beijing 100191, China

[9] Center of Pharmaceutical Technology, School of Pharmaceutical Sciences, Tsinghua University, Beijing 100084, China

[10] Beijing Academy of Artificial Intelligence, Beijing, China.

* Contribute equally to this work.

Correspondence and requests for materials should be addressed to Y.L. (lanyanyan@air.tsinghua.edu.cn), C.Y. (yancy2019@tsinghua.edu.cn) and W.Z. (wei_zhang@mail.tsinghua.edu.cn).

Abstract

Numerous protein-coding genes are associated with human diseases, yet approximately 90% of them lack targeted therapeutic intervention. While conventional computational methods such as molecular docking have facilitated the discovery of potential hit compounds, the development of genome-wide virtual screening against the expansive chemical space remains a formidable challenge. Here we introduce DrugCLIP, a novel framework that combines contrastive learning and dense retrieval to achieve rapid and accurate virtual screening. Compared to traditional docking methods, DrugCLIP improves the speed of virtual screening by several orders of magnitude. In terms of performance, DrugCLIP not only surpasses docking and other deep learning-based methods across two standard benchmark datasets but also demonstrates high efficacy in wet-lab experiments. Specifically, DrugCLIP successfully identified agonists with < 100 nM affinities for $5HT_{2A}R$, a key target in psychiatric diseases. For another target NET, whose structure is newly solved and not included in the training set, our method achieved a hit rate of 15%, with 12 diverse molecules exhibiting affinities better than Bupropion. Additionally, two chemically novel inhibitors were validated by structure determination with Cryo-EM. Building on this foundation, we present the results of a pioneering trillion-scale genome-wide virtual screening, encompassing approximately 10,000 AlphaFold2 predicted proteins within the human genome and 500 million molecules from the ZINC and Enamine REAL database. This work provides an innovative perspective on drug discovery in the post-AlphaFold era, where comprehensive targeting of all disease-related proteins is within reach.

Introduction

The human genome comprises approximately 20,000 protein-coding genes (*1*), many of which are related to a variety of diseases. Despite this, only about 10% of these genes have been successfully targeted by FDA-approved drugs or have documented small-molecule binders in the literature (*2*). This leaves a substantial portion of the druggable genome largely unexplored, representing a promising opportunity for therapeutic innovation. The scientific community is eager to translate biological relevant targets into pharmaceutical breakthroughs. However, most researchers lack access to advanced high-throughput screening equipment or sufficient computational power to perform comprehensive virtual screenings. Additionally, proteins often function as parts of families or pathways, indicating that targeting single proteins may not always be the most effective strategy (*3, 4*). These limitations can significantly reduce the success rate of drug discovery, especially for new targets. Therefore, developing a comprehensive chemical database containing genome-wide virtual screening results would be an invaluable asset for the biomedical research community, with the potential to significantly accelerate the discovery of new drugs.

Given the impracticality of experimentally screening all human proteins, virtual screening has emerged as the only viable approach to tackle the vast number of potential targets. In classical computer-aided drug discovery (CADD), molecular docking serves as a foundational technique for target-based virtual screening. Despite advancements in simplified scoring functions, optimized algorithms, and hardware acceleration (*5-9*), molecular docking remains time-intensive, often requiring several seconds to minutes to evaluate each protein-ligand pair. For example, a recent large-scale docking campaign took two weeks to screen 1 billion molecules against a single target, even with the use of 10,000 CPU cores (*10*). As a result, the computational demands for genome-wide virtual screening are prohibitively high, rendering such efforts impractical with existing technologies.

Artificial intelligence holds great promise for drug discovery. Various deep learning methods have been developed for virtual screening, focusing on predicting ligand-receptor affinities (*11-13*). Yet, applying these methods to large-scale virtual screening still face significant challenges. A primary issue is the inconsistency of affinity values due to heterogeneous experimental conditions (*14, 15*), which may negatively impact the performance of the trained model. Moreover, a notable distribution shift between training datasets and real-world testing scenarios hinders the generalizability of AI models, as real-world virtual screenings often involve a larger proportion of inactive molecules than those represented in the curated training sets (*16*). Additionally, the computational demands of deep learning models, with millions of parameters, pose a crucial bottleneck in inference speed, especially as chemical libraries and target numbers grow. Consequently, there is an urgent need for the development of more efficient and robust AI methodologies to effectively address these challenges.

In this work, we introduce DrugCLIP, a novel contrastive learning approach for virtual screening. Contrastive learning has demonstrated significant success in various applications like image-text retrieval (*17*), enzyme function annotation (*18*), and protein homology detection (*19*). The core innovation of DrugCLIP lies in its ability to distinguish potent binders from non-binding molecules with a given protein pocket by aligning their representations. This approach effectively mitigates the impact of noisy affinity labels and chemical library imbalances that have traditionally challenged virtual screening efforts. Moreover, the inference of DrugCLIP is highly efficient, approximately 100,000 times faster than traditional docking methods and 100 times faster than current machine learning-based approaches.

Comprehensive *in silico* and wet-lab evaluations were conducted to assess the accuracy of the DrugCLIP model. Our model achieved state-of-the-art performance on two widely recognized virtual screening benchmarks, DUD-E (*20*) and LIT-PCBA (*21*), outperforming traditional docking-based screening methods and other deep neural

networks. To further validate its performance, DrugCLIP was applied to screen molecules for two real-world targets: $5HT_{2A}R$ (5-hydroxytryptamine receptor 2A) and NET (norepinephrine transporter). Remarkably, our model identified chemically diverse binders with adequate affinities, which were further validated through functional assays and structure determination. These results provide compelling evidence of the efficacy of our virtual screening method.

Finally, a genome-wide virtual screening was conducted using DrugCLIP on all human proteins predicted by AlphaFold2 (*22, 23*). In this process, we first define pockets for AlphFold predictions with structure alignment (*24*), pocket detection software (*25*), and generative AI models. Next, we screened over 500 million drug-like molecules from the ZINC (*26, 27*) and Enamine REAL (*28*) databases against identified pockets. Notably, this unprecedented large-scale virtual screening was completed in just 24 hours on a single computing node equipped with 8 A100 GPUs. Lastly, we applied a CADD cluster-docking pipeline to select chemically diverse and physically proper molecules for each pocket. These result a dataset containing over 2 million potential hits targeting more than 20,000 pockets from around 10,000 human proteins. To the best of our knowledge, this is the first virtual screening campaign to perform more than 10 trillion of scoring operations on protein-ligand pairs, covering nearly half of the human genome. All molecules, scores, and poses have been made freely accessible at https://drug-the-whole-genome.yanyanlan.com, facilitating further research in drug discovery on a genome-wide scale.

Results

The design and *in-silico* evaluation of the DrugCLIP model

Unlike previous machine learning models that relied on regression to directly predict protein-ligand affinity values, DrugCLIP (Fig. 1) redefines virtual screening as a dense retrieval task. The key innovation lies in its training objective, which aims to learn an aligned embedding space for protein pockets and molecules, encoded by separated neural networks. Vector similarity metrics can then be employed to reflect their binding probability. Using contrastive loss during training, the similarity between protein pockets and their binders (positive protein-ligand pair) is maximized, whereas the similarity between protein pockets and molecules binding to other targets (negative protein-ligand pairs) is minimized.

The training process of DrugCLIP includes two stages: pretraining and fine-tuning. The molecule and pocket encoders are pretrained with large-scale synthetic data and are further refined using experimentally determined protein-ligand complex structures during fine-tuning.

In the pretraining stage, the molecule encoder is initialized with Uni-Mol, a well-established molecule encoder. With the molecule encoder frozen, the pocket encoder is randomly initialized and trained to align with the molecule encoder using contrastive learning (Fig. 1B). we developed a Protein Fragment-Surrounding Alignment (ProFSA) framework (Fig. 1A) to generate large-scale synthetic data specifically tailored for contrastive pretraining. In this approach, short peptide fragments are extracted from protein-only structures to serve as pseudo-ligands, while their surrounding regions are designated as pseudo-pockets. Appling ProFSA framework to PDB (*29*) data yielded 5.5 million pseudo-pocket and ligand pairs to facilitate the pretraining. The trained pocket encoder has been evaluated across various downstream tasks such as pocket property prediction (Table S1), pocket matching (Table S2) and protein-ligand affinity

prediction (Table S3). Experimental results demonstrate that our pretrained pocket encoder exhibits strong performance, even in a zero-shot setting, outperforming many supervised learning-based models as well as physical and knowledge-based models. These results underscore the success of the pretraining stage to obtaining meaningful molecule and pocket representations.

After pretraining, the molecule and pocket encoders are further fine-tuned (Fig. 1D) using 40,000 experimentally determined protein-ligand complex structures collected by BioLip2 database (*30*). Given that the binding conformations of molecules are unknown and only their topologies are provided in virtual screening, we implemented a random conformation sampling strategy for data augmentation by using RDKit (*31*) for conformation generation. This augmentation allows DrugCLIP to train on data that more accurately reflect the variability of real-world screenings, thereby enhancing the model's performance and generalization ability.

In the screening process (Fig. 1E), we first use our trained encoders to represent molecules and pockets to vectors. Cosine similarities between the pocket and molecule embeddings are then computed, and candidate molecules are ranked according to these similarity scores. Since the molecule representations can be computed offline, DrugCLIP screening is highly efficient, requiring only the calculation of a simple cosine similarity and subsequent ranking. Compared with traditional computational method like docking, DrugCLIP is more than 100,000 times faster.

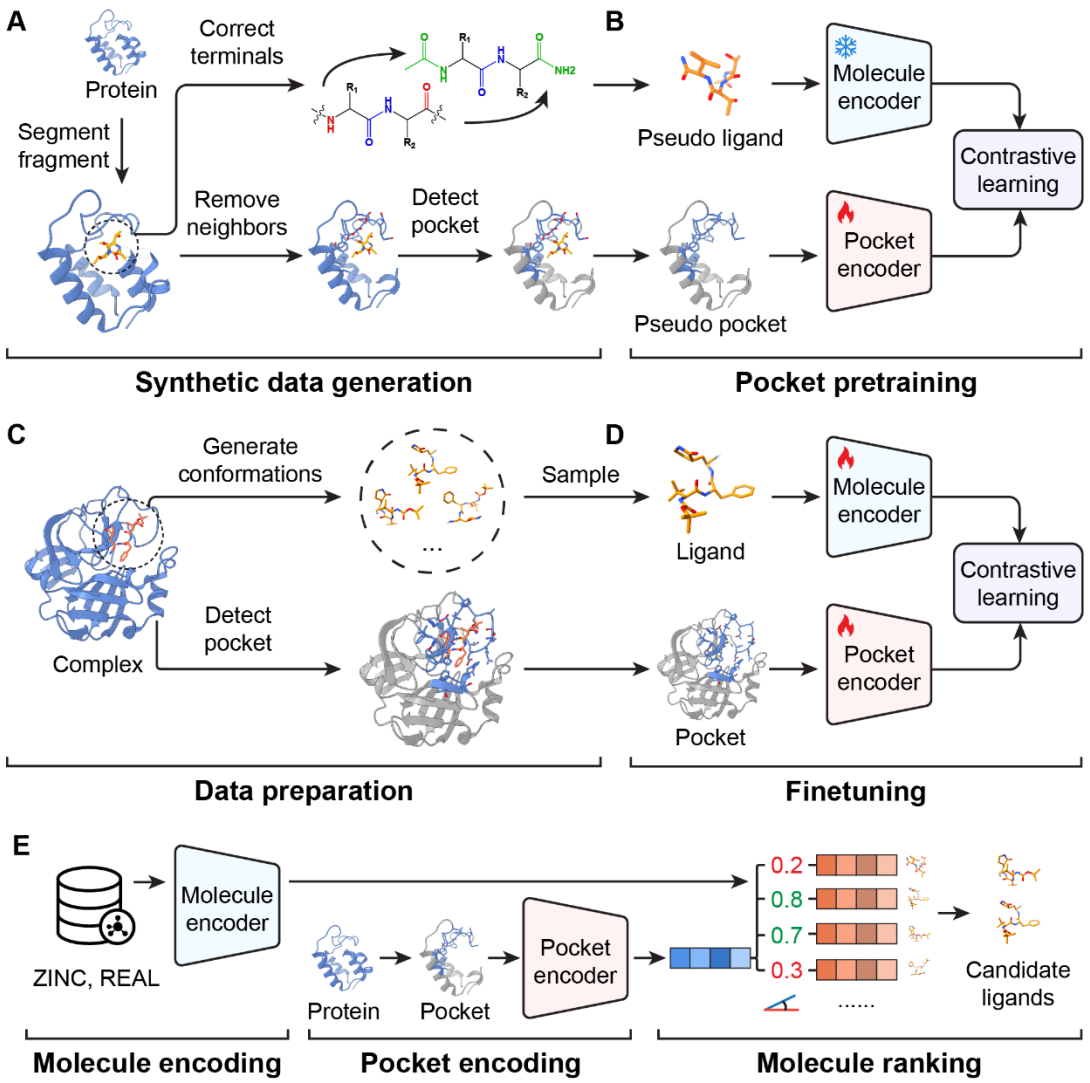

Fig. 1 The framework of DrugCLIP. **(A)** In the pretraining stage, a large-scale synthetic data was created using the ProFSA strategy. Specifically, pseudo pocket-ligand pairs were constructed through a series of operations, including fragment segmentation, terminal correction, neighbor removal and pocket detection, on protein data. **(B)** The pocket encoder is pretrained with pseudo pocket-ligand pairs in a contrastive distillation manner to transfer knowledge from a well-established molecular encoder to the pocket encoder. **(C)** During the fine-tuning process, experimentally determined protein-ligand pairs were used as training data, with multiple ligand conformations generated by RDKit. **(D)** In the fine-tuning stage, both the pocket and molecule encoders were updated using a contrastive loss, which maximizes the similarity between positive pairs and minimize it between negative pairs. **(E)** The pipeline for virtual screening with DrugCLIP. The candidate molecules from the library were pre-encoded with the trained molecular encoder. For a given pocket, the trained pocket encoder converts it to a vector, and the cosine similarity is then utilized to select top ligands with the highest scores.

We benchmarked DrugCLIP on two widely used virtual screening datasets, DUD-E (*20*) and LIT-PCBA (*21*). The DUD-E dataset contains 22,886 active compounds of 102 protein targets. For each active compounds, 50 decoys with similar physical properties but different structures are generated. In contrast, LIT-PCBA comprises approximately 8,000 active and 2.64 million inactive compounds across 15 targets, derived from experimental results of the PubChem BioAssay database. DrugCLIP was compared with established physical-informed docking software including Glide-SP (*5*), Autodock Vina (*6*), Surflex (*32*), and regression-oriented machine learning models including NNscore (*13*), RFscore (*33*), Pafnucy (*34*), OnionNet (*12*), PLANET (*11*), Gnina (*35*), BigBind (*36*). In both sets of results (Fig. 2A and 2B), DrugCLIP demonstrated a substantial improvement over all baseline methods in terms of EF1%, measuring the recall capacity of virtual screening models.

We also investigated the influence of homology information and protein structure accuracy on DrugCLIP's performance. Remarkably, even when test protein families were entirely excluded from the training set, DrugCLIP still outperform one of the most popular virtual screening method AutoDock Vina (Fig. 2C), highlighting its strong generalization capability to new targets. Moreover, DrugCLIP shows exceptional robustness by outperforming AutoDock Vina even with a 3 Å RMSD error in the side chain conformations of protein pockets (Fig. 2D), indicating its robustness to structural inaccuracies.

Furthermore, DrugCLIP is exceptionally efficient (Fig. 2E), making it highly suitable for large-scale screening tasks. For instance, DrugCLIP completed the screening for LIT-PCBA in merely 38 seconds with the aid of pre-encoding acceleration, significantly faster than Glide docking (3 days), Uni-Dock (22 hours) (*8*), and another machine learning method PLANET (3 hours) (*11*). Moreover, the time consumption of DrugCLIP screening scales linearly with the simultaneously increase of target and molecule numbers (Fig. 2F), which can facilitate multi-target virtual screening.

These *in silico* results confirm that DrugCLIP possesses superior virtual screening capabilities, combining high performance, generalizability, robustness, and efficiency.

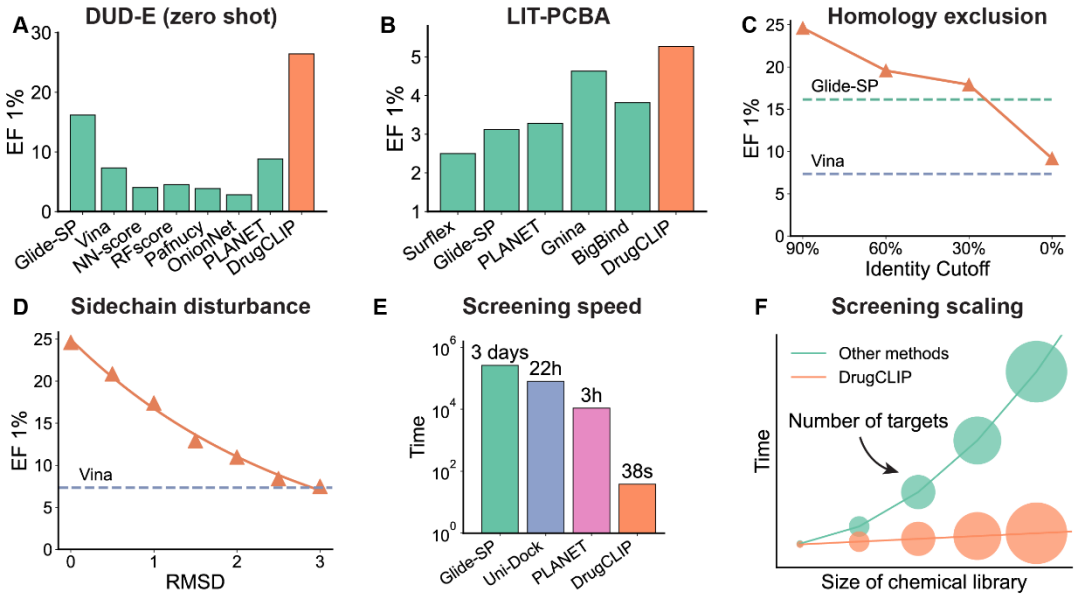

Fig. 2 *In silico* benchmarking results of DrugCLIP. **(A)** The evaluation of DrugCLIP on the DUD-E dataset using the EF1% to assess model performance. The results of baseline models are token from previous studies (*11, 37*). **(B)** The evaluation of DrugCLIP on the LIT-PCBA dataset, also using the EF1% for performance measurement. The results of baseline models are token from previous studies (*11, 21, 36, 38*). **(C)** The assessment of DrugCLIP's generalization ability was conducted by varying the identity cutoffs between testing targets and training data in DUD-E, with Glide-SP and Vina represented as dashed lines. **(D)** The evaluation of DrugCLIP's robustness regarding errors in pocket side-chain conformations, was conducted by using RMSD values ranging from 0 Å to 3 Å, with Vina shown as a dashed line for reference. **(E)** The screening speed on the LIT-PCBA dataset, compared with docking methods like Glide-SP and Uni-Dock, and the machine learning model PLANET. Speeds of baseline methods are token from previous studies (*8, 11*). **(F)** An illustration of time consumption as the screening scale increases, with the x-axis representing the size of the molecule library and the size of the circles representing the numbers of targets. DrugCLIP (the orange line) has a computational complexity of O(M+N), where M is the number of targets and N is the number of compounds, whereas most existing methods (the green line) have a complexity of O(MN).

234 **Validating the performance of DrugCLIP with wet-lab experiments**

235 In addition to *in silico* evaluation, we tested the DrugCLIP model on real-world
236 targets using wet-lab experiments. We focused on two well-established targets for
237 psychiatric diseases: the serotonin receptor 2A ($5HT_{2A}R$) and the norepinephrine
238 transporter (NET).

239 $5HT_{2A}R$ is an emerging target for antidepressant development. Psychedelics such as
240 lysergic acid diethylamide (LSD) and psilocybin act as agonists of $5HT_{2A}R$. Despite
241 their potent hallucinogenic effects, these psychedelics have demonstrated strong and
242 long-lasting antidepressant effects in both rodent models and humans (*39, 40*), making
243 $5HT_{2A}R$ a crucial target for the development of next-generation antidepressants.
244 Previous research suggests that the recruitment of β-arrestin2 following $5HT_{2A}R$
245 activation is a key biochemical mechanism underlying these antidepressant effects (*41,*
246 *42*).

247 Since the structure of $5HT_{2A}R$ was included in our training data, we trained a
248 customized DrugCLIP model more specifically for $5HT_{2A}R$ virtual screening. We
249 excluded similar experimental structures (90% homology) from the training dataset to
250 prevent potential data leakage. A library of 1,648,137 in stock molecules from
251 ChemDiv were screened, and 78 compounds from the top results were ordered,
252 selecting with criteria including drug-likeness, docking score, and chemical diversity.

253 Firstly, a calcium flux assay was used for a primary screening of these compounds in
254 an agonist model at a concentration of 10 μM. Eight of the 78 molecules were identified
255 as positive agonists, exhibiting a minimal activity of 10% compared to serotonin. The
256 affinities of these molecules to $5HT_{2A}R$ were further assessed using [^{3}H]-labeled
257 ketanserin competitive binding assays, with six showing a K_i of less than 10 μM (Table
258 S4, Fig. S5). We then evaluated the cellular function of these hit molecules using

259 NanoBit assays for β-arrestin2 recruitment, and all six molecules achieved an EC_{50} of
260 less than 1 μM (Table S4, Fig. S5).

261 Among the six validated agonists, two chemically novel molecules are particularly
262 interesting. The molecule with the highest affinity, V008-4481, binds to the orthosteric
263 (*43*), extended (*44*) and side-extended (*41*) pocket of $5HT_{2A}R$ with its indole ring,
264 phenylpiperazine part and linear tail respectively. In previous researches, molecules that
265 bind to all three regions are rarely seen. This molecule achieves an affinity of 21.0 nM
266 and exhibits an EC_{50} of 60.3 nM with an E_{max} of 35.8% in the NanoBit assay. Another
267 molecule, V006-3328, features a unique branched topology uncommon among $5HT_{2A}R$
268 binders. It has an affinity of 3.51 μM, with an EC_{50} of 599 nM and an E_{max} of 23.4% in
269 the NanoBit assay. These two molecules present strong potential as starting points for
270 the discovery of new serials of $5HT_{2A}R$ agonists.

271 Unlike $5HT_{2A}R$, the norepinephrine transporter (NET) is a well-established drug
272 target for depression and attention deficit hyperactivity disorder (ADHD), with multiple
273 FDA-approved inhibitors (*45*). However, the structures of NET with or without its
274 inhibitors in complexes were not solved until 2024 (*46-48*). The closest protein
275 structure in our dataset is the dopamine transporter from *Drosophila* (*49*), which shares
276 less than 60% similarity with NET. Therefore, screening against NET provides a more
277 challenging test of our model's ability to generalize to relatively new targets.

278 For this target, we directly utilized the DrugCLIP model trained with all collected
279 complex structures to screen the ChemDiv in-stock library. We ultimately selected 100
280 compounds considering chemical novelty and diversity. We tested their inhibition of
281 NET protein by measuring the transport of [^{3}H]-labeled norepinephrine in NET-
282 containing liposomes. Among these compounds, 15% of them exhibited more than 60%
283 inhibition of NET, with 12 compounds demonstrating greater potency than the widely
284 used antidepressant bupropion.

Unlike previous NET inhibitors that typically feature aliphatic nitrogen atoms capable of forming a salt bridge interaction with ASP75 of NET (*46-48*), our screening identified several hits with positively charged aromatic nitrogen atoms. Notably, two such molecules, 0086-0043 and Y510-9709, demonstrated better IC_{50} (with values of 1.14 μM and 0.31 μM, respectively) than bupropion (1.5 μM). Structural analysis of the complexes between these compounds and the NET protein revealed that the aromatic rings indeed form more favorable interactions with NET: the isoquinoline ring of 0086-0043 engages in a T-shaped π-π interaction with PHE72, and the thiazole ring of Y510-9709 likely interacts with surrounding aromatic side chains like PHE323. These findings highlight the potential of DrugCLIP model to providing new chemical insights for drug discovery.

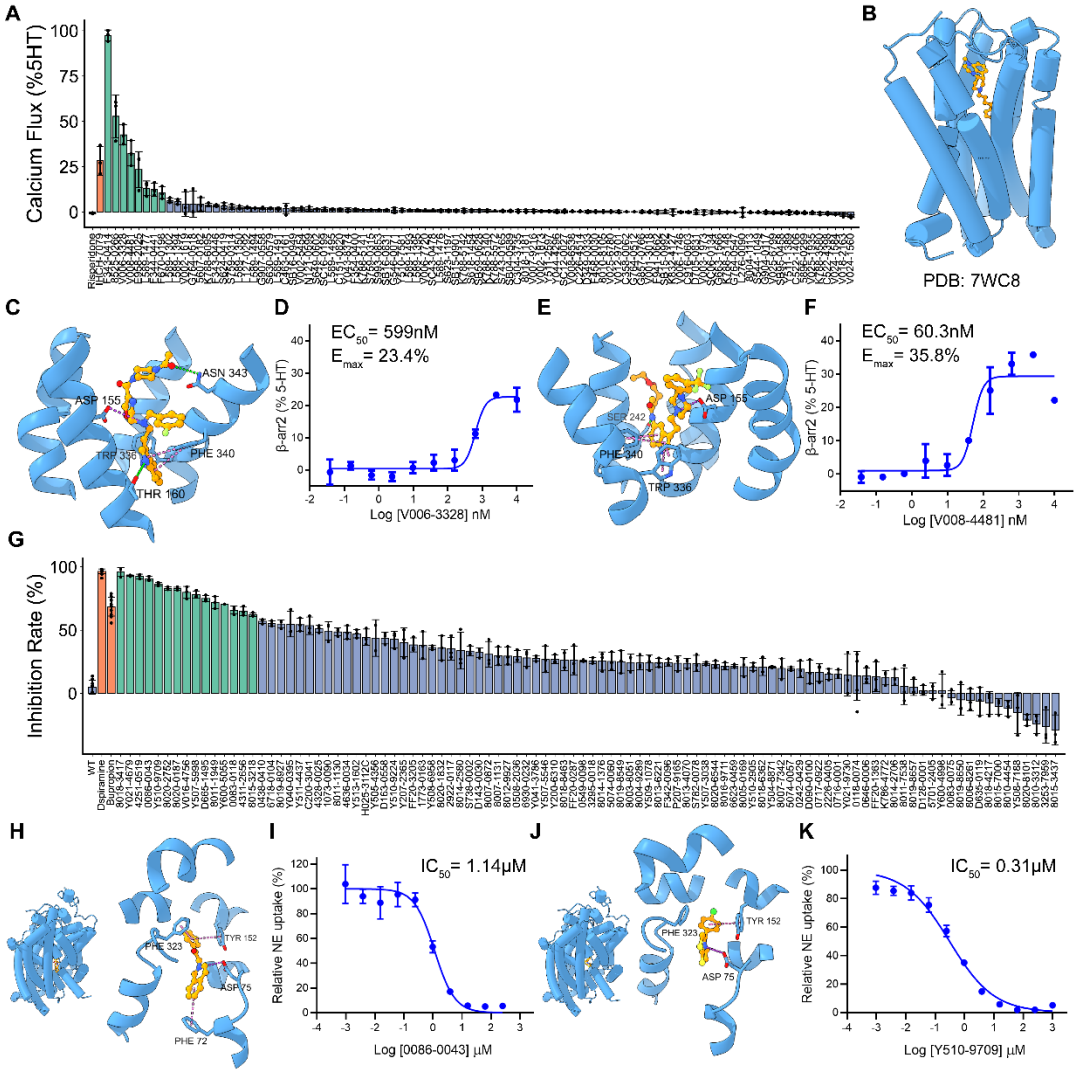

Fig. 3 Wet-lab validations of DrugCLIP. **(A)** The screening results of 78 DrugCLIP identified molecules using calcium flux assays for $5HT_{2A}R$ agonist at a concentration of 10 μM. Eight molecules showed signals larger than 10%. **(B)** The 7WC8 structure was used as the receptor for docking analysis. **(C)** The docking pose of V006-3328 to $5HT_{2A}R$. **(D)** The dose response curve of V006-3328 in the NanoBit assay measuring the β-arrestin2 recruitment. **(E)** The docking pose of V008-4481 to $5HT_{2A}R$. **(F)** The dose response curve of V008-4481 in the NanoBit assay measuring the β-arrestin2 recruitment. **(G)** The evaluation of 100 DrugCLIP identified molecules with radio-ligand transportation assays for NET inhibitor at a concentration of 10 μM, and 15 molecules showed inhibition larger than 60% **(H)** The complex structure of 0086-0043 and NET determined with Cryo-EM. **(I)** The dose response curve of 0086-0043 in the radio-ligand transportation assay. **(J)** The complex structure of Y510-9709 and NET determined with Cryo-EM. **(K)** The dose response curve of Y510-9709 in the radio-ligand transportation assay.

312 **Genome-wide virtual screening with DrugCLIP**

313 After validating the DrugCLIP model through both *in silico* and wet-lab experiments,
314 we apply it to a genome-wide virtual screening campaign. The effort is grounded on
315 significant advancements in protein structure prediction, specifically targeting all
316 human protein structures predicted by AlphaFold2 (*22, 23*).

317 Virtual screening using AlphaFold-predicted structures remains a topic of debate.
318 The primary concern is that these predicted structures may lack the accuracy needed to
319 replicate experimental conformations and effectively filter out inactive molecules (*50,*
320 *51*). Despite this, some studies have shown that virtual screening with AlphaFold-
321 predicted structures can still yield reasonable results for certain targets (*52, 53*). Despite
322 of the robustness of DrugCLIP to sidechain inaccuracies (Fig. 2D), we further assess
323 the influence of predicted structure by replacing the structures in the DUD-E dataset
324 with AlphaFold predictions and using Fpocket (*25*) to identify the nearest pockets. This
325 substitution resulted in a significant performance drop for DrugCLIP, with the EF1%
326 value decreasing from 27.43% to 13.25% (Fig. 4C), reflecting similar challenges
327 observed with docking methods in both virtual screening (*50*) and conformation
328 prediction (Fig. 4B).

329 To improve the utilities of AlphaFold-predicted structures, we developed a strategy
330 called GenPack (Generation-Packing, Fig. 4A). This strategy involves training
331 molecular generative models conditioned on the backbone structures of protein pockets.
332 While the generated molecules may not always be synthesizable, they help to localize
333 pockets more precisely and induce the pocket conformation into a more suitable state.
334 After this generation step, side chains are reintroduced, and the overall conformation is
335 refined using physical force fields. With the GenPack strategy, we significantly
336 enhanced the screening power of AlphaFold-predicted structures, increasing EF1%
337 value on the DUD-E benchmark from 13.25% to 19.54% (Fig. 4C) and improving the
338 docking success rate from 19.10% to 38.71% (Fig. 4B).

Finally, we introduced a genome-wide virtual screening pipeline to facilitate future drug discovery. We began with splitting AlphaFold predictions into high-confidence regions based on plDDT and PAE scores. For each region, we used homology alignment and Fpocket (*25*) along with GenPack to detect potential pockets. The DrugCLIP model was then employed to screen over 500 million drug-like molecules from the ZINC (*26, 27*) and Enamine REAL (*28*) databases. The screening process, which involved more than 10 trillion of scoring operations on protein-ligand pairs, was completed in about 24 hours on a single computing node equipped with 8 A100 GPUs. The top-ranked molecules were then clustered and further evaluated using molecular docking, filtering out poor poses with Glide score < -6 kcal/mol. The final database contains over 2 million potential hits molecules for more than 20,000 pockets from 10,000 human targets. All molecules, docking scores and poses have been made freely accessible at https://drug-the-whole-genome.yanyanlan.com, facilitating further research and drug discovery processes.

Our genome-wide screening results cover a more extensive range of targets than ChEMBL (*54*), one of the most comprehensive databases for bioactive molecules. While UniProt (*1*) contains 20,436 reviewed human proteins, the latest ChEMBL release (ChEMBL 34) covers 4,810 of them. Moreover, not all targets in the ChEMBL database have high-affinity small molecule binders; some targets only have peptide or antibody binders, or merely vague results from low-quality assays. In contrast, our database spans 9,908 targets, more than twice the number in ChEMBL and covers nearly half of the human genome. To visualize the difference between the two protein spaces, we encoded all protein sequences using the ESM1b model (*55*). The t-SNE plot shows that our space encompasses a broader range of proteins, including many that are not closely related to those in ChEMBL.

Our database includes a diverse range of targets, from well-studied proteins to less-explored members of well-known families, as well as proteins with limited

pharmacological understanding. For example, the c-Jun N-terminal kinase 3 (JNK3) is a classical kinase target with many ligand-bound crystal structures (*56, 57*). DrugCLIP identified molecules binds to the ATP-binding pockets, forming H-bonds with backbone atoms of MET149 in the hinge region. SLC45A2 belongs to the solute carrier (SLC) superfamily, many of whose members are important drug targets. Nevertheless, SLC45A2 has limited pharmacological studies. This gene plays a crucial role in pigmentation (*58*) and is widely expressed in cutaneous melanomas (*59*), with evidence suggesting its oncogenic potential (*60*). All molecules in the database could bind near L374, which is an important site for protein stability (*58*), thus having potential modulatory effects. Another interesting example OR6A2 belongs to the olfactory receptor family, whose members are mainly found to be expressed in olfactory receptor neurons, yet many of them are expressed in various other tissues with unexplored pharmaceutical potentials (*61*). OR6A2 is expressed in macrophages, sensing blood octanal and promoting the formation of atherosclerotic plaques (*62*). Our predicted molecules fit the orthosteric pocket of OR6A2 and can serve as potential inhibitors for treating atherosclerosis. The final example Sestrin-2 can sense leucine (*63*) and promote drug resistance of cancer cells (*64*), which belongs to a unique highly-conserved stress-inducible protein family (PF04636 or IPR006730) with only three members in the human genome. Our database contains predicted molecules bind to the same pocket of leucine (*65*) that may serve as good starting points for anti-cancer therapies. These examples highlight the potential of our database as a valuable resource for exploring the undrugged genome and facilitate future drug discovery.

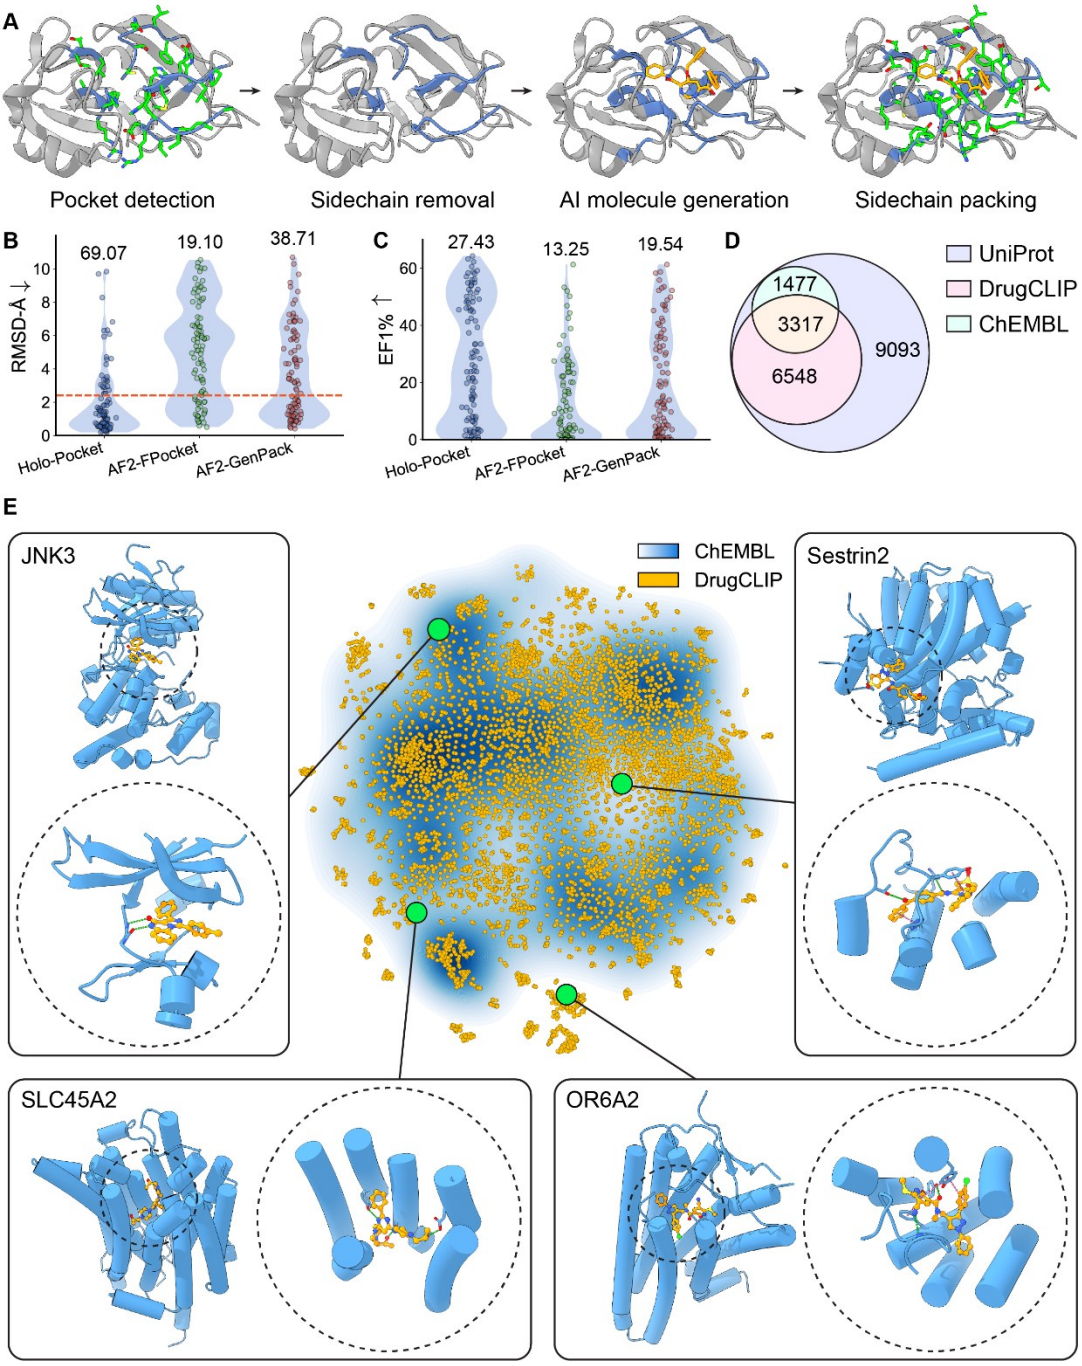

Fig. 4 DrugCLIP enables genome-wide virtual screening. **(A)** The GenPack (Generation-Packing) process for extracting pockets from AlphaFold2-predicted structures involves using Fpocket to detect initial pockets, removing sidechains, applying an AI-generative model to create molecules based on the backbone structure, and then performing sidechain packing with the generated molecules. **(B)** The redocking RMSD comparisons for different pocket definitions: holo-pocket, pockets detected by Fpocket on AF2-predicted structures, and pockets generated by GenPack. The red dash line indicates the RMSD threshold of 2 Å, and corresponding docking success rates are labeled above each column. **(C)** The EF1% comparisons for virtual screening on the DUD-E dataset using different pocket definitions:

holo-pocket, pockets detected by Fpocket on AlphaFold2-predicted structures, and pockets generated by GenPack on AlphaFold2 predictions, and averaged EF1% values are labeled above each column. **(D)** The Venn diagram of target numbers in different database, with UniProt, DrugCLIP, and ChEMBL shown as different circles. **(E)** The t-SNE visualization and examples for the genome-wide virtual screening results. Yellow dots indicate targets in our database while the blue-white gradient represents targets in the ChEMBL database, with density ranging from high (blue) to low (white).

Conclusions and Discussions

With the rapid advancement of protein structure prediction methods and the availability of a comprehensive atlas of predicted protein structures for human and disease related species (*23, 66*), we have entered a new era where effective drug discovery for all disease-related targets is within reach. In this paper, we introduce DrugCLIP, a groundbreaking contrastive learning based virtual screening approach that aims to achieve genome-wide drug discovery. The efficacy of DrugCLIP has been rigorously validated through both *in silico* benchmarks and wet-lab experiments. In well-established benchmarks, DrugCLIP consistently outperformed traditional docking software and contemporary machine learning models. Notably, for the $5HT_{2A}R$ and NET targets, DrugCLIP identified diverse high-affinity binders and novel chemical entities. These findings underscore the potential of DrugCLIP model as a reliable tool for virtual screening in real-world drug development. We demonstrate its application through a genome-wide virtual screening campaign, encompassing more than 20,000 pockets across approximately 10,000 human proteins, using a chemical library of 500 million molecules from ZINC and Enamine REAL. Remarkably, DrugCLIP completes this trillion-level virtual screening campaign in just 24 hours using just a single computational node with 8 GPU accelerators. Beyond the screening results, we have generated over 2 million high-confidence protein-ligand complex structures accompanied with their docking score. By making this extensive database freely accessible, we aim to make a substantial contribution to the research community, accelerating drug discovery and fostering innovation in therapeutic development.

DrugCLIP is more than just a new tool. It represents a transformative shift in the development of new therapeutics, heralding a new paradigm in drug discovery. Its genome-wide virtual screening capability opens the door to truly end-to-end drug discovery on a genomic scale, allowing researchers to screen all relevant targets simultaneously, rather than focusing on a few promising targets. This expansive

434 approach facilitates the creation of customized chemical libraries for advanced
435 phenotypic screening with high-fidelity models such as organoids (*67-69*) or
436 humanized mice (*70-72*), potentially reducing failure rates in drug development.

437 DrugCLIP paves the way for new advancements in AI driven drug discovery. Its
438 outstanding efficiency allows the screening scale to the largest ultra-large chemical
439 library available today, e.g., 48 billion-compound Enamine REAL Space library. This
440 effort pushes the boundaries of what virtual screening can achieve in drug discovery.
441 Moreover, the release of these genome-wide virtual screening results could serve as a
442 valuable resource for molecular generation, particularly through a retrieval-augmented
443 generation approach (*73, 74*), enhancing our capacity for drug discovery and design.

444 Acknowledgement

445 We would like to take this opportunity to thank Dr. Xiaobing Deng, Prof. Zhenming
446 Liu, Dr. Qiwei Ye and Bo Qiang for their valuable discussions and insightful comments.
447 The manuscript is polished and revised with ChatGPT-3.5 and ChatGPT-4. IHCH-7079
448 is a gift from Prof. Sheng Wang and Prof. Jianjun Cheng, and we sincerely thank their
449 generous help. This work is supported by the National Key R&D Program of China
450 (No.2021YFF1201600 to Y.L. and No.2020YFA0509301 to C.Y.), the National Natural
451 Science Foundation of China (32341016 and 32171204 to C.Y.), Beijing Academy of
452 Artificial Intelligence, and Beijing Frontier Research Center for Biological Structure
453 Fundings.

References

1. The-UniProt-Consortium, UniProt: the Universal Protein Knowledgebase in 2023. *Nucleic Acids Research* **51**, D523-D531 (2022).

2. T. I. Oprea *et al.*, Unexplored therapeutic opportunities in the human genome. *Nature Reviews Drug Discovery* **17**, 317-332 (2018).

3. P. R. Caron *et al.*, Chemogenomic approaches to drug discovery. *Current Opinion in Chemical Biology* **5**, 464-470 (2001).

4. A.-L. Barabási, N. Gulbahce, J. Loscalzo, Network medicine: a network-based approach to human disease. *Nature Reviews Genetics* **12**, 56-68 (2011).

5. T. A. Halgren *et al.*, Glide: a new approach for rapid, accurate docking and scoring. 2. Enrichment factors in database screening. *J Med Chem* **47**, 1750-1759 (2004).

6. O. Trott, A. J. Olson, AutoDock Vina: improving the speed and accuracy of docking with a new scoring function, efficient optimization, and multithreading. *J Comput Chem* **31**, 455-461 (2010).

7. O. Korb, T. Stützle, T. E. Exner, Empirical scoring functions for advanced protein-ligand docking with PLANTS. *J Chem Inf Model* **49**, 84-96 (2009).

8. Y. Yu *et al.*, Uni-Dock: GPU-Accelerated Docking Enables Ultralarge Virtual Screening. *Journal of Chemical Theory and Computation* **19**, 3336-3345 (2023).

9. M. Ling *et al.*, Vina-FPGA: A Hardware-Accelerated Molecular Docking Tool With Fixed-Point Quantization and Low-Level Parallelism. *IEEE Transactions on Very Large Scale Integration (VLSI) Systems* **31**, 484-497 (2023).

10. C. Gorgulla *et al.*, An open-source drug discovery platform enables ultra-large virtual screens. *Nature* **580**, 663-668 (2020).

11. X. Zhang *et al.*, PLANET: A multi-objective graph neural network model for protein–ligand binding affinity prediction. *Journal of Chemical Information and Modeling* **64**, 2205-2220 (2024).

12. L. Zheng, J. Fan, Y. Mu, OnionNet: A multiple-layer intermolecular-contact-based convolutional neural network for protein–ligand binding affinity prediction. *ACS Omega* **4**, 15956-15965 (2019).

13. J. D. Durrant, J. A. McCammon, NNScore 2.0: a neural-network receptor-ligand scoring function. *J Chem Inf Model* **51**, 2897-2903 (2011).

14. B. T. Burlingham, T. S. Widlanski, An intuitive look at the relationship of K_i

and IC_{50}: A more general use for the Dixon plot. *Journal of Chemical Education* **80**, 214 (2003).

15. G. A. Landrum, S. Riniker, Combining IC50 or Ki values from different sources is a source of significant noise. *Journal of Chemical Information and Modeling* **64**, 1560-1567 (2024).

16. R. Rodríguez-Pérez, M. Vogt, J. Bajorath, Influence of Varying Training Set Composition and Size on Support Vector Machine-Based Prediction of Active Compounds. *J Chem Inf Model* **57**, 710-716 (2017).

17. A. Radford *et al.*, Learning transferable visual models from natural language supervision. *International Conference on Machine Learning*, (2021).

18. T. Yu *et al.*, Enzyme function prediction using contrastive learning. *Science* **379**, 1358-1363 (2023).

19. L. Hong *et al.*, Fast, sensitive detection of protein homologs using deep dense retrieval. *Nature Biotechnology*, (2024).

20. M. M. Mysinger, M. Carchia, J. J. Irwin, B. K. Shoichet, Directory of useful decoys, enhanced (DUD-E): Better ligands and decoys for better benchmarking. *Journal of Medicinal Chemistry* **55**, 6582-6594 (2012).

21. V.-K. Tran-Nguyen, C. Jacquemard, D. Rognan, LIT-PCBA: An Unbiased Data Set for Machine Learning and Virtual Screening. *Journal of Chemical Information and Modeling* **60**, 4263-4273 (2020).

22. J. Jumper *et al.*, Highly accurate protein structure prediction with AlphaFold. *Nature* **596**, 583-589 (2021).

23. M. Varadi *et al.*, AlphaFold Protein Structure Database: massively expanding the structural coverage of protein-sequence space with high-accuracy models. *Nucleic Acids Research* **50**, D439-D444 (2021).

24. Y. Zhang, J. Skolnick, TM-align: a protein structure alignment algorithm based on the TM-score. *Nucleic Acids Res* **33**, 2302-2309 (2005).

25. V. Le Guilloux, P. Schmidtke, P. Tuffery, Fpocket: An open source platform for ligand pocket detection. *BMC Bioinformatics* **10**, 168 (2009).

26. J. J. Irwin *et al.*, ZINC20—A Free Ultralarge-Scale Chemical Database for Ligand Discovery. *Journal of Chemical Information and Modeling* **60**, 6065-6073 (2020).

27. J. J. Irwin, B. K. Shoichet, ZINC--a free database of commercially available compounds for virtual screening. *J Chem Inf Model* **45**, 177-182 (2005).

28. O. O. Grygorenko *et al.*, Generating Multibillion Chemical Space of Readily Accessible Screening Compounds. *iScience* **23**, 101681 (2020).

29. H. M. Berman *et al.*, The Protein Data Bank. *Nucleic Acids Research* **28**, 235-242 (2000).

30. C. Zhang, X. Zhang, Peter L. Freddolino, Y. Zhang, BioLiP2: an updated structure database for biologically relevant ligand–protein interactions. *Nucleic Acids Research* **52**, D404-D412 (2023).

31. G. Landrum *et al.*, RDKit: Open-source cheminformatics. *Zenodo*, 13469390 (2024).

32. A. N. Jain, Surflex: Fully Automatic Flexible Molecular Docking Using a Molecular Similarity-Based Search Engine. *Journal of Medicinal Chemistry* **46**, 499-511 (2003).

33. M. Wójcikowski, P. J. Ballester, P. Siedlecki, Performance of machine-learning scoring functions in structure-based virtual screening. *Scientific Reports* **7**, 46710 (2017).

34. M. M. Stepniewska-Dziubinska, P. Zielenkiewicz, P. Siedlecki, Development and evaluation of a deep learning model for protein–ligand binding affinity prediction. *Bioinformatics* **34**, 3666-3674 (2018).

35. A. T. McNutt *et al.*, GNINA 1.0: molecular docking with deep learning. *Journal of Cheminformatics* **13**, 43 (2021).

36. M. Brocidiacono *et al.*, BigBind: Learning from Nonstructural Data for Structure-Based Virtual Screening. *Journal of Chemical Information and Modeling* **64**, 2488-2495 (2024).

37. C. Shen *et al.*, Beware of the generic machine learning-based scoring functions in structure-based virtual screening. *Briefings in Bioinformatics* **22**, (2020).

38. J. Sunseri, D. R. Koes, Virtual Screening with Gnina 1.0. *Molecules* **26**, (2021).

39. C. L. Raison *et al.*, Single-dose psilocybin treatment for major depressive disorder: A randomized clinical trial. *JAMA* **330**, 843-853 (2023).

40. W. Duan, D. Cao, S. Wang, J. Cheng, Serotonin 2A Receptor (5-HT2AR) Agonists: Psychedelics and Non-Hallucinogenic Analogues as Emerging Antidepressants. *Chemical Reviews* **124**, 124-163 (2024).

41. D. Cao *et al.*, Structure-based discovery of nonhallucinogenic psychedelic analogs. *Science* **375**, 403-411 (2022).

42. J. Wallach *et al.*, Identification of 5-HT2A receptor signaling pathways associated with psychedelic potential. *Nature Communications* **14**, 8221 (2023).

43. K. Kim *et al.*, Structure of a Hallucinogen-Activated Gq-Coupled 5-HT(2A) Serotonin Receptor. *Cell* **182**, 1574-1588.e1519 (2020).

44. D. Wacker *et al.*, Crystal Structure of an LSD-Bound Human Serotonin Receptor. *Cell* **168**, 377-389.e312 (2017).

45. J. Zhou, Norepinephrine transporter inhibitors and their therapeutic potential. *Drugs Future* **29**, 1235-1244 (2004).

46. J. Tan *et al.*, Molecular basis of human noradrenaline transporter reuptake and inhibition. *Nature* **632**, 921-929 (2024).

47. H. Zhang *et al.*, Dimerization and antidepressant recognition at noradrenaline transporter. *Nature* **630**, 247-254 (2024).

48. T. Hu *et al.*, Transport and inhibition mechanisms of the human noradrenaline transporter. *Nature* **632**, 930-937 (2024).

49. S. Pidathala, A. K. Mallela, D. Joseph, A. Penmatsa, Structural basis of norepinephrine recognition and transport inhibition in neurotransmitter transporters. *Nature Communications* **12**, 2199 (2021).

50. Y. Zhang *et al.*, Benchmarking Refined and Unrefined AlphaFold2 Structures for Hit Discovery. *J Chem Inf Model* **63**, 1656-1667 (2023).

51. M. Karelina, J. J. Noh, R. O. Dror, How accurately can one predict drug binding modes using AlphaFold models? *eLife*, (2023).

52. J. Lyu *et al.*, AlphaFold2 structures guide prospective ligand discovery. *Science* **384**, eadn6354 (2024).

53. A. Díaz-Holguín *et al.*, AlphaFold accelerated discovery of psychotropic agonists targeting the trace amine–associated receptor 1. *Science Advances* **10**, eadn1524 (2024).

54. A. Gaulton *et al.*, ChEMBL: a large-scale bioactivity database for drug discovery. *Nucleic Acids Res* **40**, D1100-1107 (2012).

55. A. Rives *et al.*, Biological structure and function emerge from scaling unsupervised learning to 250 million protein sequences. *Proc Natl Acad Sci U S A* **118**, (2021).

56. T. Kamenecka *et al.*, Structure-Activity Relationships and X-ray Structures Describing the Selectivity of Aminopyrazole Inhibitors for c-Jun N-terminal

591 Kinase 3 (JNK3) over p38*. *Journal of Biological Chemistry* **284**, 12853-
592 12861 (2009).

593 57. K. Zheng *et al.*, Design and Synthesis of Highly Potent and Isoform Selective
594 JNK3 Inhibitors: SAR Studies on Aminopyrazole Derivatives. *Journal of*
595 *Medicinal Chemistry* **57**, 10013-10030 (2014).

596 58. L. Le *et al.*, SLC45A2 protein stability and regulation of melanosome pH
597 determine melanocyte pigmentation. *Mol Biol Cell* **31**, 2687-2702 (2020).

598 59. J. Park *et al.*, SLC45A2: A Melanoma Antigen with High Tumor Selectivity
599 and Reduced Potential for Autoimmune Toxicity. *Cancer Immunol Res* **5**, 618-
600 629 (2017).

601 60. Z. H. Zuo *et al.*, Oncogenic Activity of Solute Carrier Family 45 Member 2
602 and Alpha-Methylacyl-Coenzyme A Racemase Gene Fusion Is Mediated by
603 Mitogen-Activated Protein Kinase. *Hepatol Commun* **6**, 209-222 (2022).

604 61. R. G. Naressi, D. Schechtman, B. Malnic, Odorant receptors as potential drug
605 targets. *Trends in Pharmacological Sciences* **44**, 11-14 (2023).

606 62. M. Orecchioni *et al.*, Olfactory receptor 2 in vascular macrophages drives
607 atherosclerosis by NLRP3-dependent IL-1 production. *Science* **375**, 214-221
608 (2022).

609 63. R. L. Wolfson *et al.*, Sestrin2 is a leucine sensor for the mTORC1 pathway.
610 *Science* **351**, 43-48 (2016).

611 64. J. Qu *et al.*, A paradoxical role for sestrin 2 protein in tumor suppression and
612 tumorigenesis. *Cancer Cell International* **21**, 606 (2021).

613 65. R. A. Saxton *et al.*, Structural basis for leucine sensing by the Sestrin2-
614 mTORC1 pathway. *Science* **351**, 53-58 (2016).

615 66. Z. Lin *et al.*, Evolutionary-scale prediction of atomic-level protein structure
616 with a language model. *Science* **379**, 1123-1130 (2023).

617 67. L. Thorel *et al.*, Patient-derived tumor organoids: a new avenue for preclinical
618 research and precision medicine in oncology. *Experimental & Molecular*
619 *Medicine* **56**, 1531-1551 (2024).

620 68. F. Weeber, S. N. Ooft, K. K. Dijkstra, E. E. Voest, Tumor Organoids as a Pre-
621 clinical Cancer Model for Drug Discovery. *Cell Chemical Biology* **24**, 1092-
622 1100 (2017).

623 69. J. J. Vandana, C. Manrique, L. A. Lacko, S. Chen, Human pluripotent-stem-
624 cell-derived organoids for drug discovery and evaluation. *Cell Stem Cell* **30**,
625 571-591 (2023).

70. J. Chuprin *et al.*, Humanized mouse models for immuno-oncology research. *Nature Reviews Clinical Oncology* **20**, 192-206 (2023).

71. G. Peltz, Can humanized mice improve drug development in the 21st century? *Trends in Pharmacological Sciences* **34**, 255-260 (2013).

72. N. Legrand *et al.*, Humanized Mice for Modeling Human Infectious Disease: Challenges, Progress, and Outlook. *Cell Host & Microbe* **6**, 5-9 (2009).

73. P. Lewis *et al.*, Retrieval-augmented generation for knowledge-intensive NLP tasks. *Proceedings of the 34th International Conference on Neural Information Processing Systems*, Article 793 (2020).

74. Z. Wang *et al.*, Retrieval-based Controllable Molecule Generation. *ArXiv* **abs/2208.11126**, (2022).

Supplementary Results

Benchmarking the performance of pocket pretraining with ProFSA

To test the performance of the pretrained pocket encoder, we benchmark the encoder on three major benchmarks. The first task is about the pocket druggability prediction. We assess the effectiveness of ProFSA in predicting various physical and pharmaceutical properties of protein pockets, utilizing the druggability prediction dataset created by Uni-Mol. This dataset comprises four separate regression tasks: Fpocket score, Druggability score, Total Solvent Accessible Surface Area (SASA), and Hydrophobicity score. The evaluation metric employed for these tasks is the Root Mean Square Error (RMSE), which measures the accuracy of the predictions. The baseline model we compared is the pocket encoder from the Uni-Mol [1]. The result is shown in **Table S1**.

The second task is the zero-shot pocket matching, for which we use two datasets: the Kahraman dataset [2] and the TOUGH-M1 dataset [3]. The Kahraman dataset contains matched pockets from two non-homologous proteins that bind to the same ligand. It consists of 100 proteins binding to 9 different ligands. We use a reduced version of this dataset, excluding 20 PO_4 binding pockets due to their low number of interactions. The TOUGH-M1 dataset, on the other hand, involves relaxing identical ligands to identify similar pockets and comprises 505,116 positive and 556,810 negative protein pocket pairs derived from 7,524 protein structures. The baseline models we employed encompass various approaches, including PocketMatch [4], DeeplyTough [5] and IsoMIF [6]. Additionally, we consider established software tools like SiteEngine [7] and TM-align [8]. We also incorporate pretraining strategies, such as Uni-Mol [1] and CoSP [9]. The result is shown in **Table S2**.

The third task is binding affinity prediction. We use the widely recognized PDBBind dataset (v2019) for predicting ligand binding affinity (LBA), following the strict 30% or 60% protein sequence identity splits and preprocessing protocols specified by Atom3D. These strict data splits are crucial for providing reliable and meaningful comparisons, especially in evaluating the robustness and generalization capabilities of the models. For each protein-ligand pair, we concatenate the protein embedding from our pretrained pocket encoder with the molecular embedding from the Uni-Mol molecular encoder and pass this combined representation through a multilayer perceptron (MLP) to generate the final binding affinity prediction. For our baseline models, we utilize a diverse range of methods including DeepDTA [10], B&B [11], TAPE [12], ProtTrans [13], Holoprot [14], IEConv [15], MaSIF [16], and several ATOM3D variants—3DCNN, ENN, and GNN [17]. Additionally, we incorporate ProNet [18] and pretraining approaches such as GeoSSL [19], EGNN-PLM [20], DeepAffinity [21], and Uni-Mol [1]. The result is shown in **Table S3**.

Materials and Methods

The design of DrugCLIP

The DrugCLIP model has a molecule encoder and a pocket encoder. These two encoders are aligned by contrastive learning.

Both encoders are based on the Uni-Mol architecture [1], a transformer architecture that takes 3D atomic features as input. For the molecule encoder, we directly utilize the pretrained weights from Uni-Mol for initialization, leveraging its learned representations for small molecules. The pocket encoder is pretrained to be aligned with the molecule encoder in a contrastive distillation manner [22] with the ProFSA dataset.

The training of the DrugCLIP model is under a contrastive learning framework. Given a batch of encoded protein-ligand pairs $\{(p_1, m_1), (p_2, m_2), \ldots, (p_n, m_n)\}$, where p_i is the embedding of the protein pocket i obtained from the pocket encoder. m_i is the embedding of the corresponding ligand i encoded by the molecular encoder. The objective is to learn embeddings such that the representations of true (positive) protein-ligand pairs are closer together in the embedding space, while the representations of incorrect (negative) pairs are further apart.

To accomplish this, we use a contrastive learning framework with a batch softmax approach, which involves two main loss functions.

The first loss is designed to find the correct ligand m_i for a given protein pocket p_i. The loss function for this objective can be written as:

$$\mathcal{L}_{\text{p2m}} = -\frac{1}{N}\sum_{i=1}^{N} \log \frac{\exp(\text{sim}(p_i, m_i)/\tau)}{\sum_{j=1}^{N} \exp(\text{sim}(p_i, m_j)/\tau)}. \tag{1}$$

$\text{sim}(p_i, m_j)$ represents a similarity measure between the protein pocket embedding p_i and ligand embedding m_j. Here we use the cosine similarity. τ is the temperature parameter controlling the sharpness of the softmax distribution.

The second loss aims to find the correct protein pocket p_i from a batch of pocket candidates given a ligand m_i.

$$\mathcal{L}_{\text{m2p}} = -\frac{1}{N}\sum_{i=1}^{N} \log \frac{\exp(\text{sim}(m_i, p_i)/\tau)}{\sum_{j=1}^{N} \exp(\text{sim}(m_i, p_j)/\tau)}. \tag{2}$$

The final contrastive loss for training the model is the sum of the two losses:

$$\mathcal{L} = \mathcal{L}_{\text{p2m}} + \mathcal{L}_{\text{m2p}}. \tag{3}$$

The pretraining of the pocket encoder

The pocket pertaining uses the protein fragment-surroundings alignment (ProFSA) framework. The Protein Data Bank (PDB) [23] contains a vast amount of protein-only data. Interestingly, small molecule-protein interactions often mirror the non-covalent interactions found within proteins themselves[24]. Such similarity is shown in **Fig. S1**. Leveraging this similarity, we first extract fragments from protein structures that closely resemble known ligands and define the surrounding regions as the associated pockets of these pseudo-ligands.

In the initial phase, we iteratively isolate protein fragments ranging from 1 to 8 residues, ensuring these segments are continuous from the N-terminal to the C-terminal while excluding any discontinuous sites or non-standard amino acids. To minimize artifacts introduced by the cleavage of peptide bonds during fragment segmentation, we apply terminal modifications: acetylation at the N-terminus and amidation at the C-terminus. These modifications result in the formation of pseudo-ligands.

In the subsequent phase, to focus on long-range interactions, we exclude the five nearest residues on each side of the fragment. We then designate the pocket as the surrounding residues that have at least one heavy atom within a 6 Å distance from the fragment.

The derived pseudo-complexes undergo stratified sampling based on the distribution observed in the PDBbind2020

dataset [25, 26], considering critical parameters such as pocket sizes (measured by the number of residues) and ligand sizes (expressed as effective residue numbers, calculated by dividing the molecular weight by 110 Da). Another key metric is the relative solvent-accessible surface area (rBSA), which we calculate using the FreeSASA package [27]. The pseudo-complexes are sampled to approximate the distributions seen in the PDBbind dataset [25, 26], particularly in terms of rBSA and the joint distribution of pocket-ligand size. This ensures the dataset's representativeness and its suitability for training ligand-oriented contrastive learning models, as shown in **Fig. S2** and **Fig. S3**.

The final dataset comprises 5.5 million ligand-protein pairs, significantly larger than any existing protein-ligand complex structure dataset.

The ProFSA pretraining objective is also a batch softmax loss, where the Uni-Mol molecular encoder is used for the pseudo-ligands. During the training, the weights of the molecular encoder are frozen. This setup allows us to distill knowledge from the pretrained molecular encoder into the pocket encoder, enhancing its ability to learn interaction-aware representations of protein pockets. During the pretraining phase, the batch size is 4 × 48 on 4 Nvidia A100 GPUs. We use the Adam optimizer with a learning rate of 0.0001. The max training epochs is 100. We use polynomial decay for the learning rate with a warmup ratio of 0.06.

The fine-tuning process of DrugCLIP

We use ligand-receptor complex data from the BioLip2 [28] database, removing redundant entries (proteins with a sequence identity > 90% and binding to the same ligand) and cleaning the dataset to obtain around 43,980 high-quality protein-ligand complexes (a list of all PDB id in the training set is included the **Supplementary Materials 1**). The binding pocket for each protein is defined as the set of residues with at least one atom within 6 Å of any ligand atom. During training, we use ligand conformations sampled by RDKit rather than their co-crystal conformations to minimize the discrepancy between training and actual virtual screening conditions, as the true conformations of candidate molecules are unknown during screening.

We use an ensemble model for most applications unless mentioned otherwise, including *in silico* validations on the DUD-E and LIT-PCBA benchmarks, wet-lab validations with the NET target, and the final genome-wide virtual screening. We adopt a k-fold cross-validation strategy: the dataset is divided into k folds, and the model is trained on k-1 folds while validated on the remaining fold in each iteration. For validations on the DUD-E and LIT-PCBA benchmarks, as well as for wet-lab validations using the NET target and the final genome-wide virtual screening, we use a 6-fold cross-validation approach.

A few experiments are conducted under slightly different settings. Only one model is trained and used in the homology removal test on the DUD-E benchmark. Homology removals at a sequence identity cutoff of 30%, 60%, and 90% are achieved with MMSeqs2 [29], while the most strict homology removal is conducted with HMMER [30, 31] and the Pfam database [32]. For the results related to the $5HT_{2A}R$ target, we adopt an 8-fold cross-validation and apply data augmentation strategies, including HomoAug and ligand augmentation using the ChEMBL dataset [33], following the method outlined in the conference paper of DrugCLIP method [34].

We train the model with a batch size of 48 on 4 NVIDIA A100 GPUs. The optimizer is Adam with a learning rate of 1e-3. adam betas are 0.9 and 0.999, adam eps is 1e-8. The max epochs is set to be 200. We use polynomial decay for the learning rate and the warm-up ratio is 0.06.

Ensembling multiple pockets and models during screening

As described above, we obtain six model weights through 6-fold cross-validation. During virtual screening, these six model weights are used to generate six different predictions, which are then combined using mean pooling to achieve a robust virtual screening result.

During virtual screening, a target of interest may have multiple pocket conformations. For any candidate molecule, we use a max pooling approach to determine the maximum score between the molecule and the different pockets. However, because different pockets may have varying score ranges, this can introduce bias when applying max

pooling. To address this, we normalize the scores using an adjusted robust z-score before performing the max pooling. Specifically, for a list of scores X:

$$\text{Adjusted Robust Z-Score} = \frac{x_i - \text{Median}(X)}{\frac{\text{MAD}(X)}{0.675}}, \tag{4}$$

$$\text{MAD}(X) = \text{Median}(|x_i - \text{Median}(X)|). \tag{5}$$

In silico validation with DUD-E and LIT-PCBA dataset

The DUD-E (Directory of Useful Decoys: Enhanced) dataset [35] is a widely used resource in drug discovery research, particularly for evaluating the performance of virtual screening methods. It includes data on 102 protein targets with 22,886 active compounds known to bind to these proteins, along with a set of decoy molecules that are similar in physical properties but different in structure from the active compounds.

LIT-PCBA [36] is a benchmark dataset derived from the PubChem BioAssay database, designed for evaluating machine learning models in virtual screening and drug discovery. In the LIT-PCBA dataset, actives and decoys are defined based on experimental results from the PubChem BioAssay database. The dataset contains approximately 1.5 million compounds across 15 targets.

For each target in the DUD-E or LIT-PCBA dataset, we rank candidate molecules (including both actives and decoys) based on their cosine similarity score. This score is calculated between the encoded embeddings of the pocket and molecule using the DrugCLIP model. The Enrichment Factor (EF) is then calculated to evaluate the ability of the model to prioritize active compounds over decoys. The Enrichment Factor (EF) is typically calculated as:

$$\text{EF}_\alpha = \frac{\text{NTB}_\alpha}{\text{NTB}_t \times \alpha}. \tag{6}$$

Where NTB_α is the number of true active compounds (True Binders) identified within the top α fraction of the screened list. NTB_t is the total number of true active compounds in the entire dataset. α is the fraction of the dataset considered. In this manuscript, we use $\alpha = 1\%$.

Molecule selection for wet-lab experiments of $5\text{HT}_{2\text{A}}$R and NET

The virtual screening for $5\text{HT}_{2\text{A}}$R utilizes experimentally determined structures including 6A93, 6A94 [37], 6WGT, 6WH4, 6WHA [38], 7RAN [39], 7VOD, 7VOE [40], 7WC4, 7WC5, 7WC6, 7WC7, 7WC8, 7WC9 [41]. As for NET, structures used for virtual screening include 8HFE, 8HFF, 8HFG, 8HFI, 8HFL, 8I3V [42], where 8HFE is modified to ligand-bound complex structures using human serotonin transporter structures as templates [43, 44]. The in-stock molecules from ChemDiv company are used for screening. Top-ranked molecules identified by the DrugCLIP model are selected manually based on drug-likeness, chemical novelty, and diversity.

For $5\text{HT}_{2\text{A}}$R, the top 2% molecules are extracted, and for NET, the top 1% molecules are extracted. Then, simple drug-likeness filters are applied, with a molecular weight threshold of 550 and a QED [45] threshold of 0.5. The novelty filter excludes molecules that have large ECFP4 similarities to known actives. Known actives are obtained from the ChEMBL database [33], and defined as molecules with a pChEMBEL value > 5, or comments like "active". The ECFP4 similarity thresholds are set to 0.45 and 0.35 for $5\text{HT}_{2\text{A}}$R and NET, respectively.

From all molecules passed the above filters, human experts select around 100 molecules for wet-lab experiments. All molecules are purchased from the TopScience (Tao Shu) Company.

Functional assays of $5\text{HT}_{2\text{A}}$R

The primary screening was conducted via calcium flux assays. All molecules were dissolved in DMSO at 10mM, including the positive control IHCH-7079 [41] and the negative control Risperidone. Calcium flux assays in the agonist mode were conducted by Pharmaron, Beijing, China.

Briefly, Flp-In-CHO-5HT2A cells used in the experiment were cultured in complete medium composed of Ham's

F-12K (Hyclone, SH30526.01), 10% FBS (Gibco, 10999141), Penicillin-Streptomycin (Gibco, 15140122), and Hygromycin B (Invivogen, ant-hg-5) at a final concentration of 600 μg/mL. The cells were maintained under standard conditions at 37°C with 5% CO_2 to ensure optimal cell density. On the first day of the experiment, the cultured cells were centrifuged and resuspended in an antibiotic-free medium consisting of Ham's F-12K (Hyclone, SH30526.01) and 10% DFBS (ThermoFisher Scientific, 30067334). Approximately 7,000 cells per well were then seeded into 384-well plates (Corning, 3764) and incubated overnight. The following day, the medium in the 384-well plates was removed, and the cells were thoroughly washed with an assay buffer composed of Hank's Balanced Salt Solution (HBSS) (Gibco, 14025076) supplemented with 20 mM HEPES (Gibco, 15630080). After washing, 20 μL of assay buffer was left in each well. The 20x Component A from the FLIPR Calcium 6 Assay Kit (Molecular Devices, R8191) was diluted to 2x, and 5 mM probenecid was added. A 20 μL aliquot of this dilution was then added to each well, and the plate was incubated at 37°C for 2 hours. Subsequently, 5x concentrated test solutions of the compounds of interest and a serotonin reference solution were prepared. Using the FLIPR Tetra (Molecular Devices) system, 10 μL of each test compound solution was transferred to the respective wells of the 384-well plate, and the assay results were recorded. Calcium flux assays were repeated three times and recorded relative values were averaged.

Primary hits were defined as molecules that induced > 10% response of the 5-HT reference. These molecules were then verified with radio-ligand comparative binding assays, which were conducted by WuXi Biology. First, 5HT2A-HEK293 cells were cultured, and the cell membranes were harvested to serve as the source of $5HT_{2A}R$ protein, hereafter referred to as the membrane solution, at a concentration of 2.55 mg/mL. According to the experimental design, the test compounds and the reference compound, ketanserin (Sigma-S006), were diluted and 1 μL of each was added to the respective reaction wells. Following this, 100 μL of the membrane solution was added to each well. Next, 100 μL of ^{3}H-ketanserin was added to each well to achieve a final concentration of 1 nM. The plates were then sealed and incubated on a shaker at 300 rpm for 1 hour at room temperature. After incubation, 50 μL of 0.3% PEI (Sigma, P3143) solution was added to the Unifilter-96 GF/B filter plates (Perkin Elmer) and incubated for 30 minutes at room temperature. The reaction mixture from each well was then transferred to the filter plates, followed by filtration using a Perkin Elmer Filtermate Harvester. The wells were washed four times with 50 mM Tris-HCl buffer. Subsequently, the filter plates were dried at 50°C for 1 hour. Once dried, the filter plates were sealed at the bottom using Unifilter-96 backing tape (Perkin Elmer), and 50 μL of Microscint 20 cocktail (Perkin Elmer, 6013329) was added to each well. Finally, the top of the plates was sealed with TopSeal-A film (Perkin Elmer). The prepared plates were then placed in a MicroBeta2 Reader (Perkin Elmer) for counting. Radio-ligand comparative binding assays were replicated twice.

Molecules that showed adequate affinities to $5HT_{2A}R$ were further tested with NanoBit assays measuring the recruitment of the β-arrestin2 protein. NanoBit assays were also conducted by Wuxi Biology. On the first day of the experiment, cultured 5HT2A-HEK293 cells were collected. The HEK293 cells were first washed with DPBS solution and then treated with an appropriate amount of 0.25% trypsin-EDTA solution for 5 minutes at 37°C. After digestion, the reaction was quenched by adding an appropriate amount of complete medium, and the mixture was gently mixed. The cells were then centrifuged at 1000 rpm at room temperature to collect the cell pellet. The cells were resuspended to a concentration of 750,000 cells/mL. A 40 μL aliquot of the cell suspension was added to each well of a 384-well plate (Greiner, 781090) and incubated overnight. On the following day, 5 μL of appropriately diluted test samples and control samples were added to each well, followed by the addition of diluted NanoBit assay solution (Promega, N2012). The reaction mixture was incubated at 37°C for 30 minutes. After incubation, the experimental data were read using the Envision2104 (PerkinElmer, 2814243) system. NanoBit assays were replicated twice.

All IC_{50} and K_d values were fitted with GraphPad Prism.

For structural analysis of hit molecules, molecules are docked to 7WC8 [41] with Glide SP, and a template of OLC is used for V008-4481 with a RMSD tolerance of 5 Å

<u>Functional assays of NET</u>

Cells used for NET functional assays included *Escherichia coli* and HEK293F. The *Escherichia coli* strain DH5α was cultured in LB medium (Sigma) at 37 ℃ to generate and amplify plasmids for NET. Mammalian HEK293F cells were maintained in SMM 293-TII medium (Sino Biological) at 37°C with 5% CO2 for protein expression.
The full-length human wild-type NET cDNA (UniProt ID: P23975) was inserted into the pCAG vector using the KpnI and XhoI restriction sites, with an N-terminal FLAG tag. NET overexpression was achieved in HEK293F cells. For transfection, 2 mg of plasmid DNA and 4 mg of polyethylenimine (Polysciences) were pre-incubated in 50 ml of fresh SMM 293-TII medium for 15 minutes before being added to one liter of HEK293F cells at a density of 2.0×10^6 cells/ml. After 48 hours of shaking at 37°C, 5% CO2, and 220 rpm, the cells were collected via centrifugation, resuspended in lysis buffer (20 mM Tris-HCl pH 8.0, 150 mM NaCl), frozen in liquid nitrogen, and stored at -80°C for later use.
For protein purification, the thawed cell pellet was solubilized in lysis buffer containing protease inhibitors (5 µg/ml aprotinin, 1 µg/ml pepstatin, 5 µg/ml leupeptin; Amresco) and 2% (w/v) DDM (Anatrace) at 4°C for 2 hours, followed by centrifugation at 20,000 g at 4°C for 1 hour. The resulting supernatant was applied to anti-FLAG M2 resin (Sigma), which was washed with 15 column volumes (CV) of buffer (20 mM Tris-HCl pH 8.0, 150 mM NaCl, 0.02% (w/v) DDM). The protein was eluted with 6 CV of the wash buffer containing 0.4 mg/ml FLAG peptide at 4°C. The eluted protein was concentrated and further purified by size-exclusion chromatography using a Superose 6 Increase 10/300 GL column (GE Healthcare) in buffer (20 mM Tris-HCl pH 8.0, 150 mM NaCl, 0.02% (w/v) DDM). The peak fractions were collected and concentrated for subsequent experiments.
Then, purified NET protein was reconstructed into liposomes to form proteoliposomes. The *E. coli* polar lipid extract (Avanti), with 20% (wt %) cholesterol added, was resuspended to 20 mg/ml in buffer A (25 mM HEPES pH 7.4, 150 mM KCl). This mixture underwent ten freeze-thaw cycles using liquid nitrogen and was then extruded 21 times through 0.4 µm polycarbonate membranes (GE Healthcare). The resulting liposomes were pre-treated with 1% n-octyl-β-D-glucoside (β-OG; Anatrace) for 30 minutes at 4°C. They were then incubated with 200 µg/ml of purified NET protein (wild-type or mutants) for 1 hour at 4°C. To remove the detergents, the mixture was treated overnight with 250 mg/ml Bio-Beads SM2 (Bio-Rad) at 4°C, followed by an additional 1-hour incubation with 100 mg/ml Bio-Beads SM2. After five more freeze-thaw cycles and 21 additional extrusion passes, the proteoliposomes were collected by ultracentrifugation at 100,000 g for 1 hour at 4°C, washed twice, and resuspended to 100 mg/ml in buffer A for the subsequent uptake assay.
Each uptake assay was conducted by adding 2 µl of proteoliposomes to 96.5 µl of buffer B (25 mM HEPES pH 7.4, 150 mM NaCl) along with 0.5 µl (0.5 µCi, 12.3 Ci/mmol) of Levo-[7-3H]-Norepinephrine and 1 µl of 50 µM valinomycin. To assess the single-point inhibitory activity of the screened small molecules, proteoliposomes were incubated with these compounds, while Desipramine and Bupropion were used as positive controls for NET inhibition. All inhibitors were added at a concentration of 1 µM in a volume of 1 µl. The uptake of the radiolabeled substrates was halted after 60 seconds by rapidly filtering the solution through 0.22 µm GSTF filters (Millipore) and washing with 2.5 ml of ice-cold buffer B. Filters were then incubated with 0.5 ml of Optiphase HISAFE 3 (PerkinElmer) overnight, and radioactivity was measured using a MicroBeta2® Microplate Counter (PerkinElmer).
For IC_{50} determination of antidepressants, proteoliposomes were pre-incubated with varying concentrations of the drugs for 30 minutes before the addition of isotope-labeled substrates. IC_{50} values were calculated using GraphPad Prism 8, applying non-linear regression to fit the data to the equation:

$$Y = \frac{100}{1 + 10^{(\log \mathrm{IC}_{50} - X) \cdot \mathrm{HillSlope}}}. \tag{7}$$

with option: 'log(inhibitor) vs. normalized response—Variable slope'. Here, X represents the log of the inhibitor concentration, Y represents the normalized response (ranging from 100% to 0%), and HillSlope starts with an initial value of -1.
All experiments were conducted in triplicate using biologically independent samples. Data were normalized to the wild-type protein to express values relative to 100%. Non-specific binding was accounted for by using control

liposomes without protein insertion, ensuring that only specific interactions were measured.

Synthesis of 0086-0043 and Y510-9709

Both molecules were synthesized by Bellen Chemistry Company.
For Y510-9709 (**5-(4-chlorophenyl)-2,3-dihydrothiazolo[2,3-b]thiazol-4-ium bromide**), first synthesize compound 2 (**1-(4-chlorophenyl)-2-((4,5-dihydrothiazol-2-yl)thio)ethan-1-one**). To a solution of compound 1 (**2-bromo-1-(4-chlorophenyl)ethan-1-one**) (10.0 g, 42.8 mol, 1.0 eq) and **thiazolidine-2-thione** (5.1 g, 42.8 mmol, 1.0 eq) in EtOH (150 mL) and DMF (50 mL) was added TEA (4.3 g, 42.8 mol, 1.0 eq). The reaction mixture was stirred at room temperature for 2 h. HPLC showed no compound 1 remained. The reaction mixture was poured into crushed ice and filtered to give compound 2 (9.6 g, 82.5%) as a yellow solid. 1H NMR (300 MHz, CDCl3): δ ppm 8.00 – 7.90 (m, 2H), 7.50 – 7.40 (m, 2H), 4.62 (s, 2H), 4.17 (t, J = 8.1 Hz, 2H), 3.43 (t, J = 7.8 Hz, 2H). LCMS: 272.0 ([M+H]+).
Then, The solution of compound 2 (2.5 g, 9.2 mmol, 1.0 eq) in 30% HBr in AcOH (25 mL) was stirred at 120 °C for 3 h. TLC and HPLC showed no compound2 remained. The reaction was allowed to be cooled to room temperature and concentrated in vacuo to give the residue, which was triturated with MeOH (7.5 mL) and filtered to give Y510-9709 (1.1 g, 35.7%) as an off-white solid. 1H NMR (400 MHz, DMSO-d6): δ ppm 7.90 (s, 1H), 7.68 (s, 4H), 4.70 (t, J = 8.0 Hz, 2H), 4.10 (t, J = 8.4 Hz, 2H). LCMS: 254.0 ([M-Br]+).
For 0086-0043(**2-(2-oxo-2-phenylethyl)isoquinolin-2-ium chloride**), The solution of **2-chloro-1-phenylethan-1-one** (2.0 g, 12.9 mol, 1.0 eq) and **isoquinoline** (1.7 g, 12.9 mmol, 1.0 eq) in ACN (12 mL) was stirred at room temperature for 16 h. HPLC showed no **2-chloro-1-phenylethan-1-one** remained. The reaction mixture was filtered to give 0086-0043 (1.3 g, 35.4%) as an off-white solid. 1H NMR (400 MHz, DMSO-d6): δ ppm 10.06 (s, 1H), 8.76 (d, J = 6.8 Hz, 1H), 8.69 (d, J = 6.8 Hz, 1H), 8.56 (d, J = 8.4 Hz, 1H), 8.43 (d, J = 8.4 Hz, 1H), 8.39 – 8.28 (m, 1H), 8.20 – 8.04 (m, 3H), 7.81 (t, J = 7.6 Hz, 1H), 7.69 (t, J = 7.6 Hz, 2H), 6.66 (s, 2H). LCMS: 248.1 ([M-Cl]+).

The structure determination of NET and its inhibitors

For cryo-EM samples, 4µl purified NET protein was applied to glow-discharged Quantifoil holey carbon grids (Quantifoil Au R1.2/1.3, 300 mesh). Protein was concentrated to approximately 10 mg/ml and separately incubated with 2 mM Y510-9709 or 0086-0043 for 30 min before freezing. After applying the protein, the grids were blotted for 3 s with 100% humidity at 4 °C and plunge frozen in liquid ethane cooled by liquid nitrogen with Vitrobot (Mark IV, Thermo Fisher Scientific).
Cryo-EM data was collected on a 300 kV Titan Krios G3i equipped with a Gatan K3 detector and a GIF Quantum energy filter (slit width 20 eV). The defocus values ranged from -1.5 to -2.0 µm. Each stack of 32 frames was exposed for 2.56 s, and the exposure time of each frame was 0.08 s. The micrographs were automatically collected with AutoEMation program [46] in super-resolution counting mode with a binned pixel size of 1.083 Å. The total dose of each stack was about 50 $e^-/Å^2$. All 32 frames in each stack were aligned and summed using the whole-image motion correction program MotionCor2 [47].
All dose-weighted micrographs were manually inspected and imported into cryoSPARC [48]. Micrographs with an estimated CTF resolution worse than 4 Å were excluded during exposure curation. CTF parameters were estimated using patch-CTF. The was used for initial good templates generation via 2D classification. Initial good templates were generated via 2D classification, using the previously reported NET structure [49] (NET–DSP, PDB code: 8FHI) as a reference. The Template Picker tool was used for all particle picking tasks. For the NET_Y510-9709 and NET_0086-0043 datasets, 3,204,486 and 9,008,886 particles were extracted from 2,918 and 4,687 micrographs, respectively. Particles were initially extracted with a box size of 192 and then cropped to 128 to speed up calculations. The initial good reference for 3D classification was derived from the NET-DSP dataset, while bad references were generated using the graphical user interface (GUI) of UCSF ChimeraX [50]. Global pose estimation was performed using Non-uniform refinement, followed by local refinement for the first round of local pose assignment. A second round of local pose estimation was conducted using 3D classification (without image alignment), followed by another

round of local refinement (**Fig.S7**). This process yielded 507,444 and 506,286 particles representing the inward-open conformation, resulting in resolutions of 2.87 Å for NET_Y510-9709 and 2.98 Å for NET_0086-0043, respectively. The atomic coordinates of NET in the presence of Y510-9709 or 0086-0043 have been deposited in the Protein Data Bank (http://www.rcsb.org) under accession codes 9JEL and 9JF3. The corresponding electron microscopy maps are available in the Electron Microscopy Data Bank (https://www.ebi.ac.uk/pdbe/emdb/) under accession codes EMD-61420 and EMD-61426.

The training and inference of the GenPack generative model

We have developed a GenPack model that operates within a continuous parameter space, incorporating a noise-reduced sampling strategy inspired by MOLCRAFT[51]. Unlike full-atom approaches, our method focuses solely on the given backbone atoms to minimize the impact of potential structural variations between *apo* and *holo* states of the proteins. We meticulously curate a dataset comprising 14,616 protein-ligand pairs from the PDBbind database, which we divide into a training set of 13,137 pairs and a validation set of 1,479 pairs (**Supplementary Materials 2**). Additionally, we use 101 protein-ligand pairs from the DUD-E database as our test set. Pockets are defined by extracting backbone atoms within a 10 Å radius of the ligands. The training is conducted on a single Nvidia A100 GPU with a learning rate of 5e-4 for 60 epochs, resulting in our pocket location optimization model.
During inference, Fpocket [52] is initially employed to detect pockets approximately 10 Å in size, after which our SBDD model generates potential ligand molecules conditioned on backbone atoms only. Subsequently, side-chain atoms are introduced to the complex structure, and the complex structures are relaxed with Prime software in the Schrodinger Suite. The protein residues with at least one heavy atom within a 6 Å radius of the generated ligands are selected as the final pocket region. This approach ensures a focus on critical interactions within the binding site while reducing noise and irrelevant structures, thereby facilitating accurate pocket detection.

Pocket Detection for all AlphaFold2 predicted human proteins

The AlphaFold DB [53, 54] contains predicted structures for 20,504 human proteins identified by UniProt accessions. Among these, 208 proteins are larger than 2500 amino acids (AAs), and their Pairwise Alignment Error (PAE) cannot be accessed through the official website. Consequently, only 20,296 proteins are used for pocket detection. Not all AlphaFold2 predictions are accurate. Two types of inaccuracies can be avoided by examining the plDDT and PAE scores. First, we remove all residues with a plDDT score below 50. The remaining structures exhibit high local accuracy, but the interactions between protein domains may still be incorrect. To address this, the PAE is symmetrized and used as precomputed metrics for agglomerative clustering. The average linkage method is applied, and the PAE threshold for clustering is set at 15 Å. Each cluster is then regarded as a confidently predicted protein super-domain, and protein fragments shorter than 10 AAs are removed to ensure stability during refinement. From the 20,296 proteins, we have identified 24,692 super-domains, covering 17,188 proteins (69.6%).
For each super-domain, we utilize two methods to detect potential pockets. First, we implement a template-based structural alignment approach. Each super-domain is aligned with proteins from the PDBbind database [25, 26]. When a local structure of the super-domain exhibited high structural similarity to a known pocket from PDBbind, it is considered a likely pocket. Specifically, TM-align [55] is used for structural alignment, with a TM-score threshold of 0.6 to ensure significant overall similarity. The corresponding ligands from PDBbind are mapped to the identified pocket location in the super-domain using a rotation matrix, thereby confirming the pocket. We then calculate the local alignment IoU (intersection over union) for the pocket, defined as the ratio of the number of aligned amino acids in the pocket to the number of the union of amino acids in both the super-domain and the PDBbind protein pockets. Alignments with an IoU exceeding 0.6 are retained. Since all super-domains are single-chain proteins, only proteins from PDBbind with single-chain pockets are used for template matching. We also exclude ligand-receptor pairs from the PDBbind database where the ligand contains more than 800 atoms. In addition to the approach above, for each super-domain, Fpocket software [52] is used for pocket detection. However, the accuracy of pocket detection using Fpocket alone is limited, and the side-chain conformation of the *apo* pocket is not suitable for molecular

docking. To address this, we adopt the proposed GenPack method to refine the pocket.

The chemical library for the genome-wide virtual screening

ZINC database is pre-filtered by anodyne reactivity and lead-like properties (molecular weight is no less than 200 and up to 500, logP is up to 5). The resulting subset contains 2,782 tranches, and over 609 million protomers are downloaded from ZINC20 [56]. Enamine REAL database is downloaded from VirtualFlow [57] in the format of PDBQT. The whole database contains 46570 tranches, over 1337 million protomers. Both databases are filtered by cutoff rules for molecular properties calculated from SMILES and structural alert patterns using RDKit. Molecules of properties meeting the rules in **Table S6** are kept for subsequential research. For ZINC, SMILES strings are matched to 3D structures in PDBQT by ZINC id. For REAL, SMILES strings are first extracted from remarks in PDBQT files; if errors like syntax errors due to the letter 'q' in SMILES occurred, they are then converted from PDBQT structures via Open Babel. A regular expression filter is applied to REAL to exclude PDBQT files with overflowed atom coordinate digits.

The genome-wide virtual screening

All pockets and molecules are pre-encoded with DrugCLIP models. Then cosine similarities of their embeddings are calculated with Pytorch [58] with 8 A100 GPUs. Then, scores from 6 models and multiple pocket replicas are ensembled as discussed previously. The top 100,000 molecules for each pocket are obtained, and clustered into around 100 clusters with an ECFP4 cut-off of 0.15. Finally, the remaining molecules are docked to the pocket replica with the highest fitness with Glide SP software from the Schrodinger Suite. Only molecules with a DrugCLIP Zscore > 4 and Glide Score < -6 are included in the final database.

Supplementary Tables and Figures.

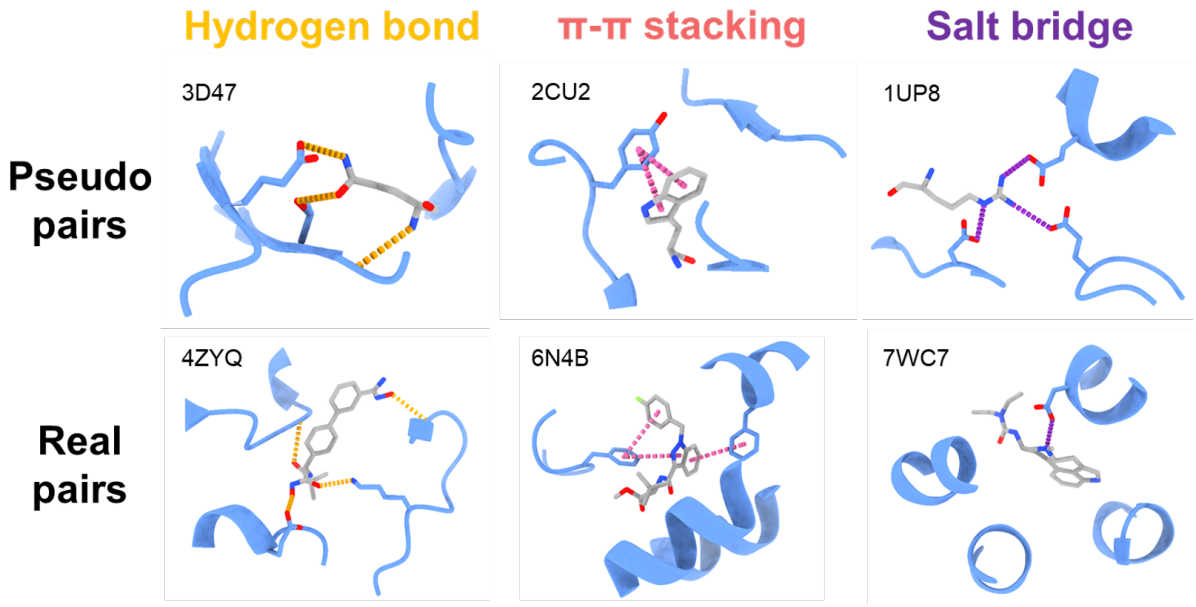

Fig. S1. Visualizations of non-covalent interactions shared by both real protein-ligand pairs and pseudo protein-ligand pairs.

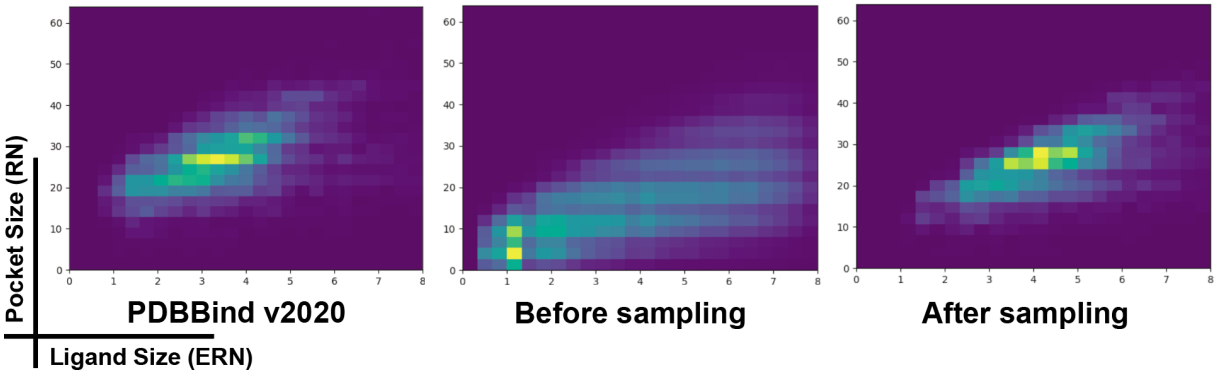

Fig. S2. The joint distributions of pocket size and ligand size are examined for the PDBBind dataset, our ProFSA dataset before applying stratified sampling, and the ProFSA dataset after stratified sampling.

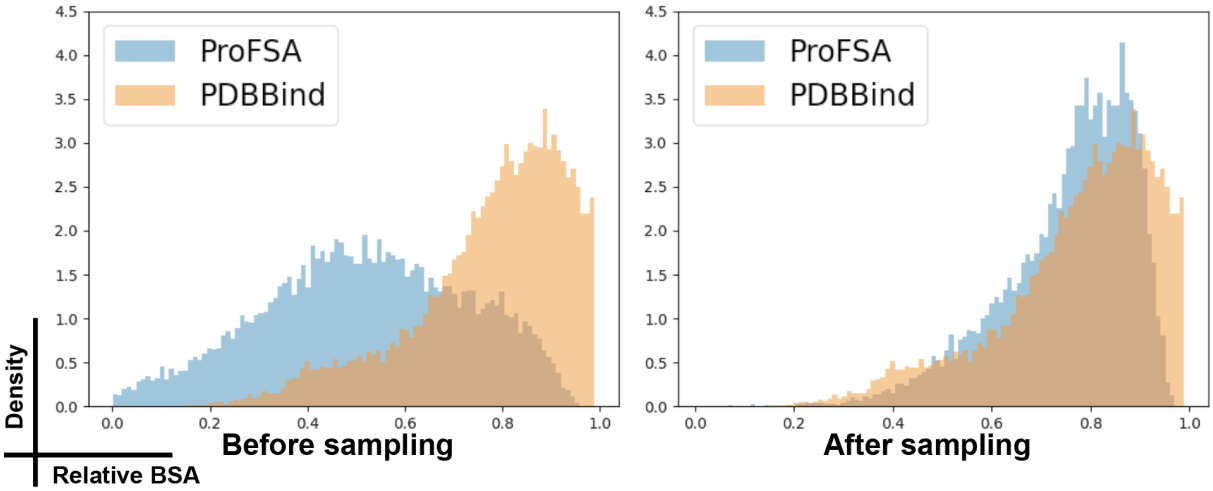

Fig. S3. Comparisons between the ProFSA dataset and the PDBBind dataset are made based on the distributions of relative Binding Surface Area (rBSA) for ligand-pocket pairs

Hit compound structure	Hit compound name	Most similar active	ChEMBL ID, Similarity & Activity records
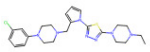	E958-2025	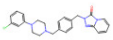	CHEMBL5186688 sim = 0.3106 Ki = 1067.0 nM
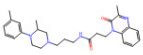	F670-0198	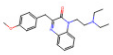	CHEMBL1729803 sim = 0.2806 AC50 = 9699.7 nM
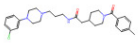	F344-0441	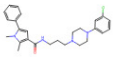	CHEMBL1080726 sim = 0.4286 IC50 = 92.0 nM IC50 = 92.0 nM
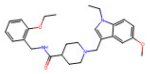	L589-1477	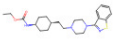	CHEMBL3923240 sim = 0.2414 Ki = 0.18 nM
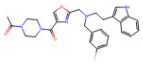	V006-3328	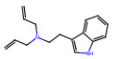	CHEMBL3752576 sim = 0.35 Ki = 707.95 nM
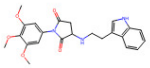	8525-0266	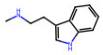	CHEMBL348588 sim = 0.3696 EC50 = 50.7 nM
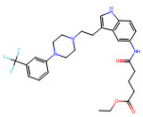	V008-4481	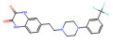	CHEMBL2298807 sim = 0.4044 Ki = 1548.82 nM
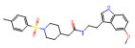	F343-0414	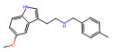	CHEMBL294216 sim = 0.4035 Ki = 110.0 nM Ki = 30.0 nM

Fig. S4. Primary hit molecules of $5HT_{2A}R$ and the known actives with the largest similarity scores. All similarity scores were calculated with Canavs software from the Schrodinger Suite with the ECFP4 fingerprint.

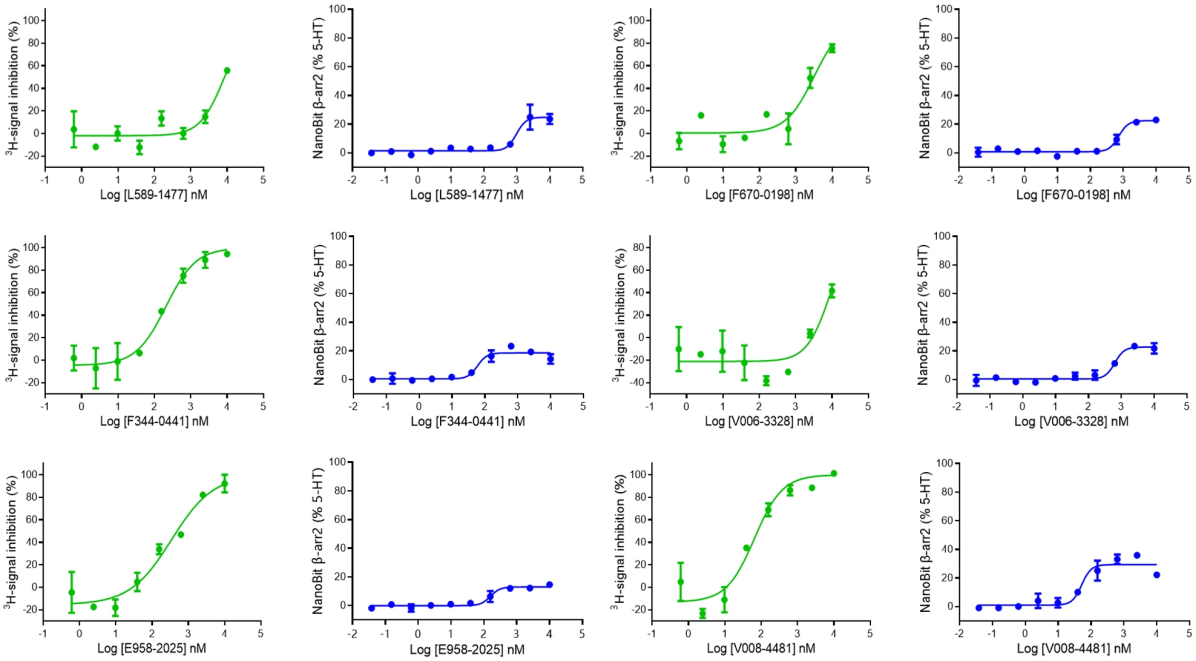

Fig. S5. Dosage response curves of primary hits of $5HT_{2A}R$ in radio-ligand competitive binding assays and NanoBit assays.

Hit compound structure	Hit compound name	Most similar active	ChEMBL ID, Similarity & Activity records
	8018-3417 Inhibition = 95.89%		CHEMBL405 sim = 0.2558 AC50 = 8840.0 nM
	Y021-4679 Inhibition = 93.05%		CHEMBL1233879 sim = 0.2584 Ki = 2757.0 nM
	4251-0519 Inhibition = 92.48%		CHEMBL4227573 sim = 0.169 IC50 = 6.0 nM
	0086-0043 Inhibition = 90.55%		CHEMBL573667 sim = 0.2209 IC50 = 34.0 nM
	Y510-9709 Inhibition = 86.54%		CHEMBL2326687 sim = 0.194 Ki = 250.0 nM
	8020-2752 Inhibition = 82.86%		CHEMBL371726 sim = 0.2472 IC50 = 3956.0 nM
	8020-0187 Inhibition = 82.77%		CHEMBL30713 sim = 0.3 Ki = 642.0 nM IC50 = 3715.35 nM
	V016-4756 Inhibition = 79.83%		CHEMBL4167315 sim = 0.2268 Ki = 100.0 nM
	Y507-5998 Inhibition = 78.05%		CHEMBL402851 sim = 0.3387 IC50 = 15.0 nM
	D665-1495 Inhibition = 75.02%		CHEMBL478032 sim = 0.2125 Activity = 19.0 nM
	8011-1949 Inhibition = 71.87%		CHEMBL242920 sim = 0.2235 Ki = 2300.0 nM
	Y600-5055 Inhibition = 70.5%		CHEMBL1213033 sim = 0.2759 AC50 = 2300.0 nM
	0083-0118 Inhibition = 65.79%		CHEMBL3323185 sim = 0.1525 IC50 = 110.0 nM Ki = 80.6 nM
	4311-2656 Inhibition = 64.92%		CHEMBL30713 sim = 0.2692 Ki = 642.0 nM IC50 = 3715.35 nM
	8015-3218 Inhibition = 62.2%		CHEMBL811 sim = 0.1918 AC50 = 9700.0 nM AC50 = 5640.4 nM

Fig. S6. Primary hit molecules of NET and the known actives with the largest similarity scores. All similarity scores were calculated with Canavs software from the Schrodinger Suite with the ECFP4 fingerprint.

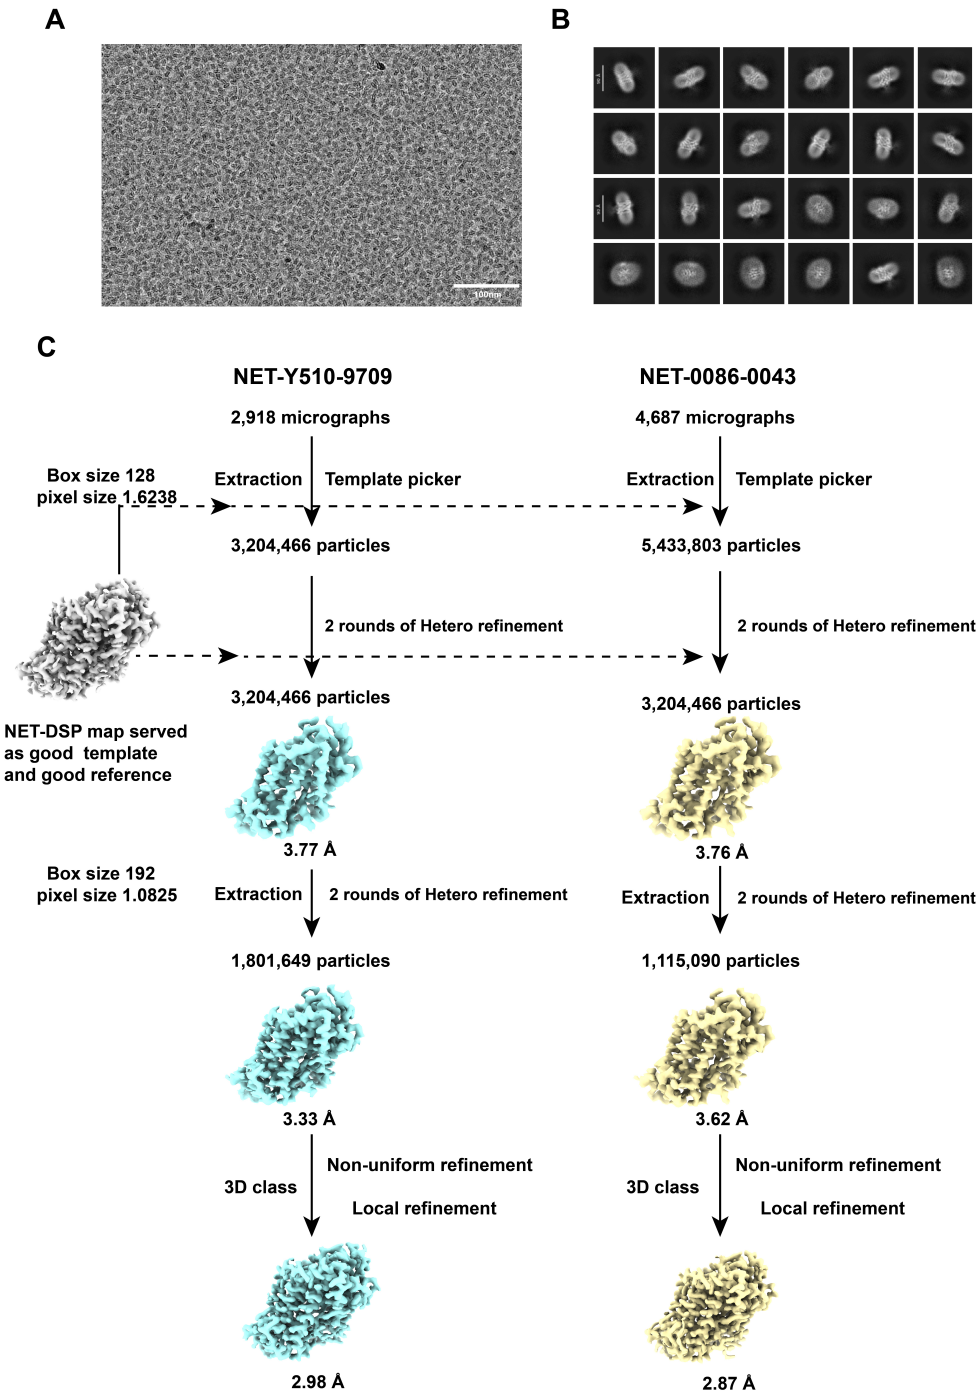

Fig. S7. Data processing of NET datasets. A-B, Representative micrograph and 2D class averages of NET. C, The flowchart for the data processing of NET bound to Y510-9709 or 0086-0043.

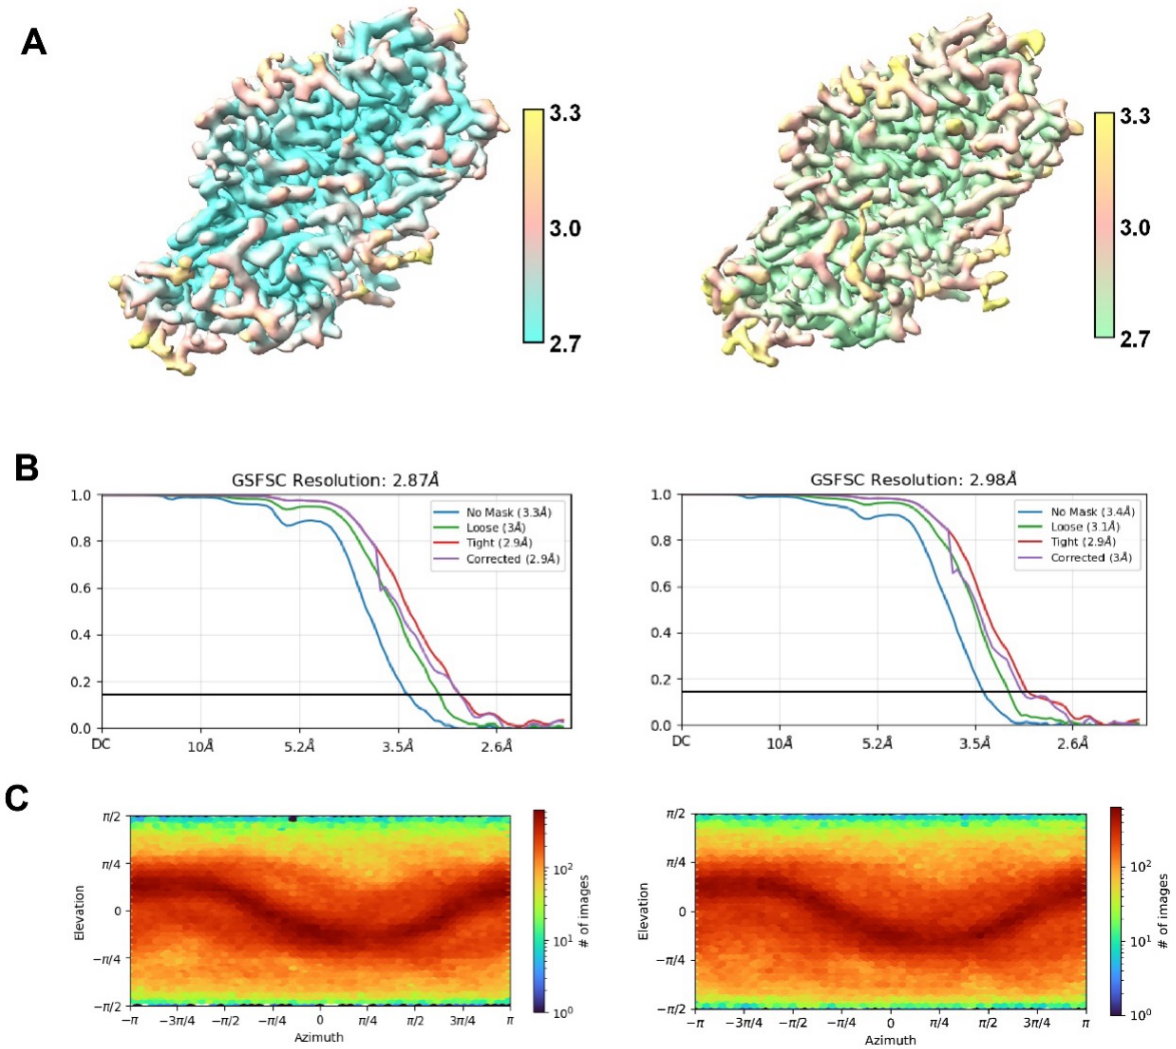

Fig. S8. Cryo-EM analysis of NET datasets. Left panel: NET bound to Y510-9709; Right panel: NET bound to 0086-0043. Various assessments of the cryo-EM reconstruction are presented. These include A, local resolution maps; B, gold-standard Fourier shell correlation (FSC) curves; C, angular distribution of the particles used for the final reconstruction.

Table S1. Druggability prediction results for pocket pretrining, using the RMSE metric.

		Fpocket ↓	Druggability ↓	Total Sasa ↓	Hydrophobicity ↓
Finetuning	Uni-Mol	0.1140	0.1001	20.73	1.285
	ProFSA	**0.1077**	**0.0934**	**20.01**	**1.275**
zero-shot	Uni-Mol	0.1419	0.1246	49.00	17.03
	ProFSA	**0.1228**	**0.1106**	**30.50**	**13.07**

Table S2. Pocket matching results for pocket pretraining, using the AUC metric.

	Methods	Kahraman(w/o PO_4) ↑	TOUGH-M1 ↑
Traditional	SiteEngine	0.64	0.73
	IsoMIF	0.75	-
Zero-shot	Uni-Mol	0.66	0.76
	ProFSA	**0.80**	**0.82**
Finetuning	DeeplyTough	0.67	0.91
	ProFSA	**0.85**	**0.94**

Table S3. Results on LBA prediction task for pocket pertaining, using pearson and spearman correlation

	Method	**Sequence Identity 30%**			**Sequence Identity 60%**		
		RMSE ↓	Pearson ↑	Spearman ↑	RMSE ↓	Pearson ↑	Spearman ↑
SequenceBased	DeepDTA	1.866	0.472	0.471	1.762	0.666	0.663
	B&B	1.985	0.165	0.152	1.891	0.249	0.275
	TAPE	1.890	0.338	0.286	1.633	0.568	0.571
	ProtTrans	1.544	0.438	0.434	1.641	0.595	0.588
StructureBased	Holoprot	1.464	0.509	0.500	1.365	0.749	0.742
	ATOM3D-3DCNN	1.416	0.550	0.553	1.621	0.608	0.615
	ATOM3D-GNN	1.601	0.545	0.533	1.408	0.743	0.743
	ProNet	1.463	0.551	0.551	1.343	**0.765**	0.761
PretrainingBased	GeoSSL	1.451	0.577	0.572	-	-	-
	EGNN-PLM	1.403	0.565	0.544	1.559	0.644	0.646
	Uni-Mol	1.520	0.558	0.540	1.619	0.645	0.653
	ProFSA	**1.377**	**0.628**	**0.620**	**1.334**	0.764	**0.762**

Table S4. The biochemical and cellular parameters of initially screened positive compounds.

Compound number	K_i (nM)	β-arr2 NanoBiT	
		EC_{50} (nM)	E_{max} (%)
L589-1477	3201.5	961.5	24.9
F344-0441	68.4	65.0	23.4
8525-0266	-	-	-
E958-2025	138.5	163.8	14.6
F343-0414	-	-	-
F670-0198	1224.2	771.2	23.0
V006-3328	3510.6	599.3	23.4
V008-4481	21.0	60.3	35.8

Table S5. Cryo-EM data collection, refinement and validation statistics.

Category	**Y510-9709**	**0086-0043**
Data collection and processing		
Magnification	64,000	64,000
Voltage (kV)	300	300
Electron exposure (e^{-}/Å^2)	50	50
Defocus range (μm)	-1.5 to -2.0	-1.5 to -2.0
Pixel size (Å)	1.0825	1.0825
Symmetry imposed	C2	C2
Raw movies	2,918	2,687
Particle number	507 k	506 k
Map resolution (Å)	2.98	2.87
FSC threshold	0.143	0.143
Map resolution range (Å)	40–2.8	40–2.7
Refinement		
Protein residues	548	548
Ligand	Y510-9709:1 Cl-	0086-0043:1 Cl-
B factors (Å^2)		
Protein	25.76	50.53
Ligand	32.09	38.45
Water	30.28	48.95
R.m.s. deviations		
Bond lengths (Å)	0.004	0.003
Bond angles (°)	0.666	0.631
Validation		
MolProbity score	1.64	1.41
Clashscore	6.27	5.37
Ramachandran plot		
Favored (%)	96.32	97.24
Allowed (%)	3.68	2.76
Disallowed (%)	0.00	0.00
PDB code	9JEL	9JF3
EMDB code	EMD-61420	EMD-61426

Table S6. Molecular database filter rules. These rules were concluded based on druglike-ness rules, public structural alerts, and the world (drug) subset of ZINC quantile numbers. Additional constraints on flexibility-related properties were imposed to prevent a sharp increase in the computational cost of molecular docking.

Property	**Limitation**
Molecular weight	(0, 500]
Number of rings	[1, 7]
Number of H-bond donors	[0, 5]
Number of H-bond acceptors	[0, 10]
ClogP	[-3, 5]
Topological polar surface area (TPSA)	[0, 140]
Number of rotatable bonds	[0, 10]
Number of aromatic rings	[1, 7]
Max size of ring	[3, 8]
Number of isomers	[1, 4]
Fraction of N or O	[0.001, 0.4]
Fraction of heteroatoms	[0.001, 0.5]
Number of contiguous rotatable bonds	[0, 4]
Number of contiguous non-ring bonds	[0, 6]
Allowed atom types	{H, C, N, O, F, Cl, Br, I, S, P}
No matching structural alert catalogs	PAINS, ZINC, CHEMBL_Glaxo, CHEMBL_BMS, CHEMBL_SureChEMBL, CHEMBL_Inpharmatica, NIH
No matching patterns	Multi-ether-ester (#[6]-#[8,#16;!a]-#[6].#[6]-#[8,#16;!a]-#[6])
	Di-guanidine (#[7]~#[6](~#[7])~#[7]~#[6](~#[7])~#[7])

References

[1] G. Zhou, Z. Gao, Q. Ding, *et al.*, "Uni-mol: A universal 3d molecular representation learning framework," in *International Conference on Learning Representations,* (2023).

[2] A. Kahraman, R. J. Morris, R. A. Laskowski, *et al.*, "On the diversity of physicochemical environments experienced by identical ligands in binding pockets of unrelated proteins," Proteins: Struct. **78** (2010).

[3] R. G. Govindaraj and M. Brylinski, "Comparative assessment of strategies to identify similar ligand-binding pockets in proteins," BMC Bioinform. **19** (2018).

[4] K. Yeturu and N. Chandra, "Pocketmatch: a new algorithm to compare binding sites in protein structures," BMC bioinformatics **9**, 1–17 (2008).

[5] M. Simonovsky and J. Meyers, "Deeplytough: learning structural comparison of protein binding sites," J. chemical information modeling **60**, 2356–2366 (2020).

[6] M. Chartier and R. Najmanovich, "Detection of binding site molecular interaction field similarities," J. chemical information modeling **55**, 1600–1615 (2015).

[7] A. Shulman-Peleg, R. Nussinov, and H. J. Wolfson, "Siteengines: recognition and comparison of binding sites and protein–protein interfaces," Nucleic acids research **33**, W337–W341 (2005).

[8] Y. Zhang and J. Skolnick, "Tm-align: a protein structure alignment algorithm based on the tm-score," Nucleic acids research **33**, 2302–2309 (2005).

[9] Z. Gao, C. Tan, L. Wu, and S. Z. Li, "Cosp: Co-supervised pretraining of pocket and ligand," arXiv preprint arXiv:2206.12241 (2022).

[10] H. Öztürk, A. Özgür, and E. Ozkirimli, "Deepdta: deep drug–target binding affinity prediction," Bioinformatics **34**, i821–i829 (2018).

[11] T. Bepler and B. Berger, "Learning protein sequence embeddings using information from structure," in *International Conference on Learning Representations,* (2019).

[12] R. Rao, N. Bhattacharya, N. Thomas, *et al.*, "Evaluating protein transfer learning with tape," in *Advances in Neural Information Processing Systems,* (2019).

[13] A. Elnaggar, M. Heinzinger, C. Dallago, *et al.*, "Prottrans: Towards cracking the language of lifes code through self-supervised deep learning and high performance computing," IEEE Trans. on Pattern Anal. Mach. Intell. pp. 1–1 (2021).

[14] V. R. Somnath, C. Bunne, and A. Krause, "Multi-scale representation learning on proteins," Adv. Neural Inf. Process. Syst. **34**, 25244–25255 (2021).

[15] P. Hermosilla, M. Schäfer, M. Lang, *et al.*, "Intrinsic-extrinsic convolution and pooling for learning on 3d protein structures," Int. Conf. on Learn. Represent. (2021).

[16] P. Gainza, F. Sverrisson, F. Monti, *et al.*, "Deciphering interaction fingerprints from protein molecular surfaces using geometric deep learning," Nat. Methods **17**, 184–192 (2020).

[17] R. J. L. Townshend, M. Vögele, P. Suriana, *et al.*, "Atom3d: Tasks on molecules in three dimensions," (2022).

[18] L. Wang, H. Liu, Y. Liu, *et al.*, "Learning hierarchical protein representations via complete 3d graph networks," in *The Eleventh International Conference on Learning Representations,* (2022).

[19] S. Liu, H. Guo, and J. Tang, "Molecular geometry pretraining with SE(3)-invariant denoising distance matching," in *The Eleventh International Conference on Learning Representations,* (2023).

[20] F. Wu, S. Li, L. Wu, *et al.*, "Discovering the representation bottleneck of graph neural networks from multi-order interactions," arXiv preprint arXiv:2205.07266 (2022).

[21] M. Karimi, D. Wu, Z. Wang, and Y. Shen, "Deepaffinity: interpretable deep learning of compound–protein affinity through unified recurrent and convolutional neural networks," Bioinformatics **35**, 3329–3338 (2019).

[22] Y. Tian, D. Krishnan, and P. Isola, "Contrastive representation distillation," ArXiv **abs/1910.10699** (2019).

[23] H. M. Berman, J. D. Westbrook, Z. Feng, *et al.*, "The protein data bank," Nucleic acids research **28 1**, 235–42 (2000).

[24] N. F. Polizzi and W. F. DeGrado, "A defined structural unit enables de novo design of small-molecule–binding proteins," Science **369**, 1227 – 1233 (2020).

[25] Z. Liu, Y. Li, L. Han, *et al.*, "Pdb-wide collection of binding data: current status of the pdbbind database," Bioinformatics **31 3**, 405–12 (2015).

[26] R. Wang, X. Fang, Y. Lu, and S. Wang, "The pdbbind database: collection of binding affinities for protein-ligand complexes with known three-dimensional structures." J. medicinal chemistry **47 12**, 2977–80 (2004).

[27] S. Mitternacht, "Freesasa: An open source c library for solvent accessible surface area calculations," F1000Research **5** (2016).

[28] C. Zhang, X. Zhang, P. L. Freddolino, and Y. Zhang, "Biolip2: an updated structure database for biologically relevant ligand–protein interactions," Nucleic Acids Res. **52**, D404 – D412 (2023).

[29] M. Steinegger and J. Söding, "Mmseqs2: sensitive protein sequence searching for the analysis of massive data sets," bioRxiv (2017).

[30] S. R. Eddy, "A probabilistic model of local sequence alignment that simplifies statistical significance estimation," PLoS Comput. Biol. **4** (2008).

[31] S. R. Eddy, "Accelerated profile hmm searches," PLoS Comput. Biol. **7** (2011).

[32] R. D. Finn, J. Mistry, J. G. Tate, *et al.*, "The pfam protein families database," Nucleic Acids Res. **38**, D211 – D222 (2007).

[33] A. Gaulton, L. J. Bellis, A. P. Bento, *et al.*, "Chembl: a large-scale bioactivity database for drug discovery," Nucleic acids research **40**, D1100–D1107 (2012).

[34] B. Gao, B. Qiang, H. Tan, *et al.*, "Drugclip: Contrasive protein-molecule representation learning for virtual screening," in *NeurIPS 2023,* (2023).

[35] M. M. Mysinger, M. Carchia, J. J. Irwin, and B. K. Shoichet, "Directory of useful decoys, enhanced (dud-e): Better ligands and decoys for better benchmarking," J. Med. Chem. **55**, 6582 – 6594 (2012).

[36] V.-K. Tran-Nguyen, C. Jacquemard, and D. Rognan, "Lit-pcba: An unbiased data set for machine learning and virtual screening," J. chemical information modeling (2020).

[37] K. T. Kimura, H. Asada, A. Inoue, *et al.*, "Structures of the 5-ht2a receptor in complex with the antipsychotics risperidone and zotepine," Nat. Struct. & Mol. Biol. **26**, 121 – 128 (2019).

[38] K. Kim, T. Che, O. Panova, *et al.*, "Structure of a hallucinogen-activated gq-coupled 5-ht2a serotonin receptor," Cell **182**, 1574–1588.e19 (2020).

[39] A. L. Kaplan, D. N. Confair, K. Kim, *et al.*, "Bespoke library docking for 5-ht2a receptor agonists with antidepressant activity," Nature **610**, 582–591 (2022).

[40] Z. Chen, L. Fan, H. Wang, *et al.*, "Structure-based design of a novel third-generation antipsychotic drug lead with potential antidepressant properties," Nat. Neurosci. **25**, 39 – 49 (2021).

[41] D. Cao, J. Yu, H. Wang, *et al.*, "Structure-based discovery of nonhallucinogenic psychedelic analogs," Science **375**, 403 – 411 (2022).

[42] J. Tan, Y. Xiao, F. Kong, *et al.*, "Molecular basis of human noradrenaline transporter reuptake and inhibition." Nature (2024).

[43] J. A. Coleman, E. Green, and E. Gouaux, "X-ray structures and mechanism of the human serotonin transporter," Nature **532**, 334 – 339 (2016).

[44] J. A. Coleman, D. Yang, Z. Zhao, *et al.*, "Serotonin transporter–ibogaine complexes illuminate mechanisms of inhibition and transport," Nature **569**, 141 – 145 (2019).

[45] G. R. J. Bickerton, G. V. Paolini, J. Besnard, *et al.*, "Quantifying the chemical beauty of drugs." Nat. chemistry **4 2**, 90–8 (2012).

[46] J. Lei and J. Frank, "Automated acquisition of cryo-electron micrographs for single particle reconstruction on an fei tecnai electron microscope," J. structural biology **150**, 69–80 (2005).

[47] S. Q. Zheng, E. Palovcak, J.-P. Armache, *et al.*, "Motioncor2: anisotropic correction of beam-induced motion for improved cryo-electron microscopy," Nat. methods **14**, 331–332 (2017).

[48] A. Punjani, J. L. Rubinstein, D. J. Fleet, and M. A. Brubaker, "cryosparc: algorithms for rapid unsupervised cryo-em structure determination," Nat. methods **14**, 290–296 (2017).

[49] J. Tan, Y. Xiao, F. Kong, *et al.*, "Molecular basis of human noradrenaline transporter reuptake and inhibition," Nature pp. 1–9 (2024).

[50] E. F. Pettersen, T. D. Goddard, C. C. Huang, *et al.*, "Ucsf chimerax: Structure visualization for researchers, educators, and developers," Protein science **30**, 70–82 (2021).

[51] Y. Qu, K. Qiu, Y. Song, *et al.*, "Molcraft: Structure-based drug design in continuous parameter space," arXiv preprint arXiv:2404.12141 (2024).

[52] V. L. Guilloux, P. Schmidtke, and P. Tufféry, "Fpocket: An open source platform for ligand pocket detection," BMC Bioinform. **10**, 168 – 168 (2009).

[53] M. Váradi, D. Bertoni, P. Magaña, *et al.*, "Alphafold protein structure database in 2024: providing structure coverage for over 214 million protein sequences," Nucleic Acids Res. **52**, D368 – D375 (2023).

[54] J. M. Jumper, R. Evans, A. Pritzel, *et al.*, "Highly accurate protein structure prediction with alphafold," Nature **596**, 583 – 589 (2021).

[55] Y. Zhang and J. Skolnick, "Tm-align: a protein structure alignment algorithm based on the tm-score," Nucleic Acids Res. **33**, 2302 – 2309 (2005).

[56] J. J. Irwin, K. G. Tang, J. Young, *et al.*, "Zinc20 - a free ultralarge-scale chemical database for ligand discovery," J. chemical information modeling (2020).

[57] C. Gorgulla, A. Boeszoermenyi, Z.-F. Wang, *et al.*, "An open-source drug discovery platform enables ultra-large virtual screens," Nature **580**, 663 – 668 (2020).

[58] A. Paszke, S. Gross, F. Massa, *et al.*, "Pytorch: An imperative style, high-performance deep learning library," ArXiv **abs/1912.01703** (2019).