
MULTI-OBJECTIVE GENERATIVE AI FOR DESIGNING NOVEL BRAIN-TARGETING SMALL MOLECULES

Ayush Noori[†] & Iñaki Arango[†]

Harvard University
Cambridge, MA 02138, USA
{anoori, inakiarango}@college.harvard.edu

William E. Byrd

Hugh Kaul Precision Medicine Institute
Heersink School of Medicine
University of Alabama at Birmingham
Birmingham, AL 35233, USA
webyrd@uab.edu

Nada Amin*

Harvard John A. Paulson School of Engineering and Applied Sciences
Cambridge, MA 02139, USA
namin@seas.harvard.edu

ABSTRACT

The strict selectivity of the blood-brain barrier (BBB) represents one of the most formidable challenges to successful central nervous system (CNS) drug delivery, preventing the diagnosis and treatment of CNS disorders. Computational methods to generate BBB permeable lead compounds *in silico* may be valuable tools in the CNS drug design pipeline. However, in real-world applications, BBB penetration alone is insufficient; rather, after transiting the BBB, molecules must perform some desired function – such as binding to a specific target or receptor in the brain – and must also be safe and non-toxic for use in human patients. To discover small molecules that concurrently satisfy these constraints, we use multi-objective generative AI to synthesize drug-like blood-brain-barrier permeable small molecules that also have high predicted binding affinity to a disease-relevant CNS target. Specifically, we computationally synthesize molecules with predicted bioactivity against dopamine receptor D₂, the primary target for almost all clinically effective antipsychotic drugs. After training several graph neural network-based property predictors, we adapt SyntheMol (Swanson et al., 2024), a recently developed Monte Carlo Tree Search-based algorithm for antibiotic design, to perform a multi-objective guided traversal over an easily synthesizable molecular space. We design a library of 26,581 novel and diverse small molecules containing hits with high predicted BBB permeability and favorable predicted safety and toxicity profiles, and that could readily be synthesized for experimental validation in the wet lab. We also validate top scoring molecules with molecular docking simulation against the D₂ receptor and demonstrate predicted binding affinity on par with risperidone, a clinically prescribed D₂-targeting antipsychotic. In the future, the SyntheMol-based computational approach described here may enable the discovery of novel neurotherapeutics for currently intractable disorders of the CNS.

*Correspondence: namin@seas.harvard.edu; † These authors contributed equally to this work.

1 INTRODUCTION

In this paper, we design small molecule drug candidates in a multi-objective setting to transit the blood-brain-barrier (BBB) and target the human brain, while also satisfying other therapeutically relevant constraints. As the interface between the central nervous system (CNS) and peripheral blood circulation, the BBB is responsible for regulating CNS homeostasis, protecting the brain microenvironment, and preventing the infiltration of toxins or pathogens across all levels of the neurovascular tree (Daneman & Prat, 2015; Obermeier et al., 2013) (Figure S1). Although the strict selectivity of the BBB is critical to protecting CNS health and integrity, it also represents the most formidable challenge to successful CNS drug delivery in disease. The BBB blocks the access of most drugs to the brain, including approximately 100% of all large-molecule and 98% of all small-molecule neurotherapeutics, *i.e.*, those with molecular masses greater than 400-500 kDa or with low lipid solubilities (Pardridge, 2005; 2012). Further, the limited small molecule drugs which penetrate the BBB only target select brain disorders; among these are depression, schizophrenia, chronic pain, and epilepsy. Meanwhile, most CNS conditions – including neurodegenerative disorders such as Alzheimer’s, Parkinson’s, and Huntington’s diseases; multiple sclerosis; HIV-associated neurocognitive disorders; brain cancer; neurotrauma; and cerebrovascular disease – remain refractory to neurotherapeutic administration and often lack effective treatment options (Pardridge, 2005). BBB-imposed transport selectivity also prevents the uptake of radiotracers in positron emission tomography (PET) and single photon emission computed tomography (SPECT) studies of the brain, impeding the diagnosis of CNS disorders and hindering neuroimaging-based disease staging.

Few strategies to evade or circumvent BBB regulation exist. Current approaches include intracerebroventricular injection, intranasal administration, exosome or nanoparticle-based delivery systems, and focused ultrasound (fUS) with intravenous microbubble agents to cause localized BBB disruption (He et al., 2018). However, these methods are frequently invasive with poor safety and efficacy profiles and remain an area of active research. For example, nasal epithelial barriers and the arachnoid membrane – which separates the nasal submucosa and olfactory cerebrospinal fluid (CSF) – limit brain uptake of drugs administered intranasally (Merkus et al., 2003), while fUS-mediated BBB disruption may cause microhemorrhages, direct vascular rupture, ischemia, edema, thermocoagulative necrosis, or injury by microbubble-induced mechanical forces (Meng et al., 2019). Thus, the most reliable solution to drug targeting in the brain is to engineer drug candidates which can cross the BBB without secondary BBB disruption or evasion. To that end, tools to discover or design CNS therapeutics that are BBB permeable with desirable pharmacokinetic profiles would be valuable assets in the CNS drug discovery process.

Wet lab methods to produce BBB permeable drugs such as high-throughput drug screening in BBB models have encountered practical challenges and achieved limited success. As a basic model of the BBB, static monolayers of primary endothelial cells or immortalized human endothelial cell lines grown in cell culture inserts can be used to study signaling pathways or transporter kinetics and may also be co-cultured with other BBB-associated cell types such as astrocytes and pericytes (Avdeef et al., 2014; Bagchi et al., 2019). However, such models fail to accurately replicate cell-to-cell signaling at the BBB interface and do not account for shear stress generated by physiological blood flow which regulates TJ formation and barrier function. In response, iPSC-based, microfluidic, organ-on-chip, and multicellular organoid models of the BBB continue to be developed (Bagchi et al., 2019). Drug screening via these dynamic systems remains expensive and has yet to achieve widespread acceptance as an effective assay of BBB integrity. Further, all existing *in vitro* and *in vivo* models (*e.g.*, the rodent brain) lack sufficient throughput required of modern, automated drug discovery programs. That is, wet lab approaches cannot be used to characterize BBB penetration at scale across the vast, unexplored regions of the neurotherapeutic molecular space.

By contrast, *in silico* generation of BBB permeable drug candidates would allow for the rapid, economical, and high-throughput screening of CNS drug candidates, and may enable the discovery of novel neurotherapeutics for currently intractable CNS disorders. However, *in silico* methods to generate BBB-penetrant molecules alone are insufficient. In real-world applications, after transiting the BBB, molecules must perform some desired function, such as binding to a specific target or receptor in the brain. Molecules must also be drug-like: they must be safe and non-toxic for use in human patients. To discover small molecules that concurrently satisfy these constraints, we use multi-objective generative AI to synthesize blood-brain-barrier permeable small molecules that also

feature other therapeutically relevant properties, including safety, toxicity, and binding affinity to specific targets in the CNS, such as the dopamine receptor D₂.

2 METHODOLOGY

Due to the intricate dynamics of transport regulation across the BBB, the complexity of drug safety and toxicity in the human body, and the large space of potential CNS neurotherapeutics, the generation of BBB permeable drug-like molecules is a task well-suited to machine learning (ML) (Butler et al., 2018). By training on small molecules with known properties, supervised ML algorithms attempt to learn inductive biases about the underlying biochemical rules which govern BBB penetration capacity, which can then be leveraged to guide a molecular generation process and propose candidate molecules that satisfy bioactivity and pharmacokinetic constraints. Here, we sought to design molecules that meet three such objectives:

1. Molecules must be able to transit the BBB.
2. Molecules must exhibit binding affinity to a specific target in the brain. In this work, we select for binding affinity against the dopamine receptor D₂ (D2R), however, we emphasize that our approach is modular. Should a different target be of interest, our method can be adapted to generate small molecules with bioactivity against this new target.
3. Molecules must be safe and non-toxic.

For each objective, we obtained labeled data to act as a supervisory signal for the training of an ML model, which we refer to as a "property predictor." Briefly, for BBB permeability, we used the Blood-Brain-Barrier Dataset (B3DB), a curated resource of 4,956 BBB permeable (BBB+) and 2,851 BBB non-permeable small molecules (Meng et al., 2021). For D2R binding affinity, we retrieved 8,034 D2R-active molecules and sampled 12,000 D2R-inactive molecules from the ExCAPE-DB database (Sun et al., 2017). Finally, for safety and toxicity, we retrieved a range of absorption, distribution, metabolism, excretion, and toxicity (ADME-Tox) datasets from the Therapeutics Data Commons (TDC) (Huang et al., 2021a; 2022).

Next, we trained graph neural networks on these datasets (see Appendix B) to predict the properties of interest. Specifically, for each model, we used a graph convolutional network encoder and concatenated with structured molecular features to produce likelihood scores of BBB permeability or D2R bioactivity (Figures 1A, S2). To evaluate out-of-sample model performance in a data-efficient manner, property predictor models were trained with 10-fold cross-validation, and the average prediction across all 10 folds was taken as the final likelihood score. For the ADME-Tox property predictor, we use ADMET-AI, an ML platform for ADME-Tox prediction trained on the TDC datasets (Swanson et al., 2023b).

Finally, we combined the outputs of all three property predictors to guide a search across the small molecule space. As described in Appendix B, we adapted SyntheMol, a recently developed Monte Carlo Tree Search (MCTS)-based algorithm for antibiotic design, to perform a multi-objective guided traversal over 29.6 billion molecules in the Enamine REadily AccessibLe (REAL) space (Swanson et al., 2024; Grygorenko et al., 2020). Like in SyntheMol, in each iteration of the algorithm, we perform an MCTS rollout to construct a molecule using the 132,479 commercially available molecular building blocks and 13 chemical synthesis reactions; however, we then evaluate the molecule using both the BBB permeability and D2R bioactivity property predictors. Using this search strategy, we generated a molecular library of 26,581 small molecule drug candidates with putative BBB permeability, D2R bioactivity, safety, and non-toxicity, which we characterized to evaluate the success of our method.

3 RESULTS AND DISCUSSION

To predict BBB permeability, D2R bioactivity, and ADME-Tox characteristics, three separate property predictors were used, trained on the B3DB, pre-processed ExCAPE-DB, and TDC databases, respectively (Meng et al., 2021; Sun et al., 2017; Huang et al., 2021a; Swanson et al., 2023b). Encouragingly, our BBB and D2R property predictor models demonstrated strong out-of-sample performance, as visualized by AUROC curves in Figure S3. The BBB predictor achieved greater than

0.95 AUC across all folds, while the D2R bioactivity predictor achieved near-perfect performance of 1.00 AUC, again across all folds. We highlight the performance of the first cross-validation fold for the BBB and D2R predictors (Figure S3A and S3B) and also show the performance across a set of randomly sampled folds for both models (Figures S3C, S3D). These results suggest that our trained models are likely useful oracles for predicting both BBB permeability and binding affinity to D2R.

Guided by these high-performing property predictors, we performed a SyntheMol-based MCTS search over 29.6 billion molecules in the pre-processed subset of the Enamine REAL combinatorial molecular space to generate a molecular library of 26,581 small molecules over 20,000 MCTS rollouts (Swanson et al., 2024; Grygorenko et al., 2020). First, we investigated diversity in building blocks and reactions sampled by the SyntheMol-based search. Across all 26,581 molecules, no individual molecular fragment was employed more than 27 times; that is, at most, any individual building block will be contained in only 0.102% of the molecular library (Figure 1B). Further, all 13 chemical reactions are employed in the search process, and 11 of the 13 chemical reactions are used to generate more than 300 molecules (Figure 1C).




Figure 1: Results of SyntheMol-based guided Monte Carlo Tree Search for multi-objective small molecule design. Guided by predictors of BBB permeability, D2R binding, and toxicity, MCTS was performed over 29.6 billion compounds to generate a library of 26,581 molecules.

Next, we explored the similarity of the generated molecules to the labeled molecules in the B3DB (Figure 1D) and ExCAPE-DB training (Figure 1E) data, for the BBB permeability and D2R bioactivity prediction tasks, respectively. For each molecule in the library, we computed the maximum Tversky similarity (an asymmetric similarity measure) between the Morgan fingerprints of that molecule and the fingerprints of all other molecules in each of the training sets. We consider a Tversky similarity score greater than 0.9 to be undesirable, and find that the 73.61% and 99.47% of the molecules have similarity scores less than 0.9 against both the B3DB (Figure 1D) and ExCAPE-DB (Figure 1E) datasets, respectively. The diversity of generated molecules was more preferable for ExCAPE-DB as compared to B3DB, which was expected due to the increased size and diversity of the ExCAPE-DB database as compared to B3DB. These results indicate that the SyntheMol-based search did not merely memorize and reproduce molecules from the training data.

We also characterized the structural diversity within the 26,581 compounds that comprise our molecular library (Figure 1F). For each molecule in the library, we computed the maximum Tanimoto distance (a symmetric distance metric) between the Morgan fingerprints of that molecule and the fingerprints of all other molecules in the library. We find that 73.61% of the molecules have an internal similarity less than 0.8, and 99.29% of the molecules have an internal similarity less than 0.9. These results suggest that the molecular library generated here is diverse and amenable to experimental screening; that is, the SyntheMol-based search did not produce a structurally homogeneous library. Our results also align with those of Swanson et al. (2024), even though the prior work focused on antibiotic discovery for infectious diseases and bacterial pathogens, while our synthesis efforts concern a wholly different biological domain of neurotherapeutics for CNS disorders.

Finally, we use the trained BBB, D2R, and ADME-Tox property predictors to evaluate each of the molecules in the library and produce an aggregate likelihood score α of each biomolecular property (Figure 1H). We visualize the average joint likelihood and compare against an unguided MCTS baseline where building blocks were randomly sampled (Figure 1G). Not all molecules are predicted as BBB permeable, active against D2R, and non-toxic – indeed, there is a peak at $\alpha = 0.55$, suggesting that many molecules are predicted to have some, but not all, properties of interest. However, in the guided search, there is a second peak at $\alpha = 0.87$: 2904 molecules have $\alpha > 0.65$ and 1662 molecules have $\alpha > 0.80$. By contrast, in the unguided random baseline, only 573 molecules have $\alpha > 0.65$ and only 224 molecules have $\alpha > 0.80$. Thus, our multi-objective SyntheMol-based guided Monte Carlo Tree Search synthesized high-scoring molecules at a 5.07-fold and 7.42-fold higher hit rate for $\alpha > 0.65$ and $\alpha > 0.80$, respectively, as compared to an unguided baseline.

We further validated top scoring generated molecules with protein-ligand molecular docking simulation using AutoDock Vina (Trott & Olson, 2010; Eberhardt et al., 2021; Murail et al., 2021). We retrieved the crystal structure of the dopamine receptor D₂ (Protein Data Bank (PDB) ID: 6cm4) (Wang et al., 2018), and simulated the binding of D2R against the top 5 generated molecules ranked by D2R property predictor score, as well as the bottom 5 scoring molecules and risperidone (a second generation antipsychotic prescribed for the treatment of bipolar disorder and schizophrenia which acts by binding to D2R) as negative and positive controls, respectively. After docking simulation, the predicted binding affinity of risperidone was -11.40 kcal/mol, comparable to the mean predicted affinities of the top scoring molecules at -11.50 kcal/mol (Figure 1I, error bars show mean \pm SE). By contrast, under simulation, the bottom scoring molecules bound more weakly, with mean predicted affinities of -6.36 kcal/mol. These results indicate that our SyntheMol-based molecular design strategy discovered small molecules with strong predicted binding affinity to the dopamine receptor D₂ as well as high predicted BBB permeability and non-toxicity. These compounds may represent early drug candidates for challenging CNS conditions such as schizophrenia, Parkinson’s disease, and substance abuse disorders, and could merit further investigation and experimental characterization in the wet lab.

3.1 FUTURE DIRECTIONS

Various future directions exist to extend this work, address the limitations of this proof-of-concept study, and validate our findings. First, when training property predictors, rather than a uniform, random split, a scaffold split could be used to partition molecules based on their Bemis-Murcko scaffolds (Wu et al., 2018; Landrum et al., 2022). Such two-dimensional scaffolds simplify the chemical structure of a compound by ignoring side chains, atom type, hybridization, and bond order; only ring systems are preserved (Bemis & Murcko, 1996). By ensuring greater structural diversity among the training and test sets, scaffold splitting would improve the out-of-distribution generalization capacity of our trained property predictors and likely increase the quality of our generated molecular library. Other areas of improvement include evaluation of different neural architectures in the encoder of our property predictor models (e.g., graph transformers), addition of DeepPurpose-based ADME-Tox predictors over ADMET-AI (Huang et al., 2021b), provisioning of additional computational resources to enable more MCTS rollouts and a larger generated library, logic-based integration of the individual property predictor outputs in $P(N)$ rather than taking the average score, and parallelization of the MCTS search to improve performance and increase coverage of the MCTS rollouts over the Enamine REAL space (see Appendix B). In particular, the use of Transformer-based encoders would allow us to identify functional groups of synthesized molecules with the highest attention weights that may be most relevant to classification of BBB permeability and could inform

future rational drug design efforts in medicinal chemistry or suggest novel mechanisms for BBB transport and uptake into the brain.

Nonetheless, these results establish groundwork in multi-objective *in silico* discovery of novel CNS therapeutic candidates. By training multiple property predictors to drive a search over the Enamine REAL space, we jointly optimized for several important drug prerequisites – namely, BBB permeability, toxicity, synthesizability, and bioactivity against the important dopamine receptor D₂. Further, we note that the method described here could be used to discover small molecules with binding affinity to targets other than D₂R. Although D₂R is indeed a disease-relevant target of high priority, we emphasize that our approach is modular and not restricted to D₂R or the conditions in which D₂R is implicated. Rather, by leveraging the 70,850,163 bioactivity annotations in the chemogenomics database ExCAPE-DB, we could feasibly train predictors of bioactivity against 1666 additional protein targets beyond D₂R (Sun et al., 2017). Coupled with SyntheMol-based guided search, this strategy would allow us to develop molecular libraries relevant to a range of other conditions that could again be validated by molecular docking simulation. The approach advanced here could be used to profile unexplored regions of the CNS neurotherapeutic space and suggest new drug and radiotracer candidates to better diagnose and treat CNS disease. Finally, the most exciting and urgent future direction of this work is experimental validation in the wet lab. We are eager to partner with experimentalists to evaluate our molecular library in both *in vitro* and *in vivo* models of schizophrenia and Parkinson’s disease, in an effort to work towards pharmacological interventions for these currently intractable neurological disorders.

CODE AVAILABILITY

Code is available via GitHub at <https://github.com/ayushnoori/molecule-synthesis>.

ACKNOWLEDGMENTS

We gratefully acknowledge Kyle Swanson and James Zou for their support and collaboration, allowing us to extend their prior work. K.S. and J.Z. shared an unpublished copy of their manuscript with us, “Generative AI for designing and validating easily synthesizable and structurally novel antibiotics”; this work is now published at the Conference on Neural Information Processing Systems (NeurIPS) 2023 Generative AI and Biology Workshop and in *Nature Machine Intelligence*. K.S. and J.Z. also provided us with early access to their code base, now available at <https://github.com/swansonk14/SyntheMol> on GitHub. We thank K.S. and J.Z. for their critical review of our manuscript and for their constructive feedback. We also thank Venkatesh Murthy and Tyler Holloway for valuable discussions.

Figure S1 was created with BioRender.com.

REFERENCES

- N. Joan Abbott, Lars Rönnbäck, and Elisabeth Hansson. Astrocyte-endothelial interactions at the blood-brain barrier. *Nature Reviews Neuroscience*, 7(1):41–53, January 2006. ISSN 1471-003X. doi: 10.1038/nrn1824.
- N. Joan Abbott, Adjanie A. K. Patabendige, Diana E. M. Dolman, Siti R. Yusof, and David J. Begley. Structure and function of the blood–brain barrier. *Neurobiology of Disease*, 37(1):13–25, January 2010. ISSN 0969-9961. doi: 10.1016/j.nbd.2009.07.030. URL <https://www.sciencedirect.com/science/article/pii/S0969996109002083>.
- Vinicius M. Alves, Eugene Muratov, Denis Fourches, Judy Strickland, Nicole Kleinstreuer, Carolina H. Andrade, and Alexander Tropsha. Predicting chemically-induced skin reactions. Part I: QSAR models of skin sensitization and their application to identify potentially hazardous compounds. *Toxicology and Applied Pharmacology*, 284(2):262–272, April 2015. ISSN 0041-008X. doi: 10.1016/j.taap.2014.12.014. URL <https://www.sciencedirect.com/science/article/pii/S0041008X14004529>.
- Alex Avdeef, Maria A. Deli, and Winfried Neuhaus. In vitro assays for assessing BBB permeability: artificial membrane and cell culture models. In *Blood-Brain Barrier in Drug Discovery: Optimizing Brain Exposure of CNS Drugs and Minimizing Brain Side Effects for Peripheral Drugs*.

-
- John Wiley & Sons, December 2014. ISBN 978-1-118-78854-7. Google-Books-ID: WRUHB-gAAQBAJ.
- Sounak Bagchi, Tanya Chhibber, Behnaz Lahooti, Angela Verma, Vivek Borse, and Rahul Dev Jayant. In-vitro blood-brain barrier models for drug screening and permeation studies: an overview. *Drug Design, Development and Therapy*, 13:3591–3605, October 2019. ISSN 1177-8881. doi: 10.2147/DDDT.S218708. URL <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC6805046/>.
- Jean-Martin Beaulieu and Raul R. Gainetdinov. The Physiology, Signaling, and Pharmacology of Dopamine Receptors. *Pharmacological Reviews*, 63(1):182–217, March 2011. ISSN 0031-6997, 1521-0081. doi: 10.1124/pr.110.002642. URL <https://pharmrev-aspetjournals-org.ezp-prod1.hul.harvard.edu/content/63/1/182>. Publisher: American Society for Pharmacology and Experimental Therapeutics Section: Review Article.
- G. W. Bemis and M. A. Murcko. The properties of known drugs. 1. Molecular frameworks. *Journal of Medicinal Chemistry*, 39(15):2887–2893, July 1996. ISSN 0022-2623. doi: 10.1021/jm9602928.
- A. Patrícia Bento, Anna Gaulton, Anne Hersey, Louisa J. Bellis, Jon Chambers, Mark Davies, Felix A. Krüger, Yvonne Light, Lora Mak, Shaun McGlinchey, Michal Nowotka, George Papadatos, Rita Santos, and John P. Overington. The ChEMBL bioactivity database: an update. *Nucleic Acids Research*, 42(D1):D1083–D1090, January 2014. ISSN 0305-1048. doi: 10.1093/nar/gkt1031. URL <https://doi.org/10.1093/nar/gkt1031>.
- Fabio Broccatelli, Emanuele Carosati, Annalisa Neri, Maria Frosini, Laura Goracci, Tudor I. Oprea, and Gabriele Cruciani. A Novel Approach for Predicting P-Glycoprotein (ABCB1) Inhibition Using Molecular Interaction Fields. *Journal of Medicinal Chemistry*, 54(6):1740–1751, March 2011. ISSN 0022-2623. doi: 10.1021/jm101421d. URL <https://doi.org/10.1021/jm101421d>. Publisher: American Chemical Society.
- Keith T. Butler, Daniel W. Davies, Hugh Cartwright, Olexandr Isayev, and Aron Walsh. Machine learning for molecular and materials science. *Nature*, 559(7715):547–555, July 2018. ISSN 1476-4687. doi: 10.1038/s41586-018-0337-2. URL <http://www.nature.com/articles/s41586-018-0337-2>. Number: 7715 Publisher: Nature Publishing Group.
- Patricia Campos-Bedolla, Fruzsina R. Walter, Szilvia Veszelka, and Mária A. Deli. Role of the Blood–Brain Barrier in the Nutrition of the Central Nervous System. *Archives of Medical Research*, 45(8):610–638, November 2014. ISSN 0188-4409. doi: 10.1016/j.arcmed.2014.11.018. URL <https://www.sciencedirect.com/science/article/pii/S0188440914002677>.
- Miriam Carbon-Mangels and Michael C. Hutter. Selecting Relevant Descriptors for Classification by Bayesian Estimates: A Comparison with Decision Trees and Support Vector Machines Approaches for Disparate Data Sets. *Molecular Informatics*, 30(10):885–895, 2011. ISSN 1868-1751. doi: 10.1002/minf.201100069. URL <https://onlinelibrary.wiley.com/doi/abs/10.1002/minf.201100069>. eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1002/minf.201100069>.
- Adrià Cereto-Massagué, María José Ojeda, Cristina Valls, Miquel Mulero, Santiago Garcia-Vallvé, and Gerard Pujadas. Molecular fingerprint similarity search in virtual screening. *Methods*, 71:58–63, January 2015. ISSN 1046-2023. doi: 10.1016/j.ymeth.2014.08.005. URL <https://www.sciencedirect.com/science/article/pii/S1046202314002631>.
- Feixiong Cheng, Weihua Li, Yadi Zhou, Jie Shen, Zengrui Wu, Guixia Liu, Philip W. Lee, and Yun Tang. admetSAR: A Comprehensive Source and Free Tool for Assessment of Chemical ADMET Properties. *Journal of Chemical Information and Modeling*, 52(11):3099–3105, November 2012. ISSN 1549-9596. doi: 10.1021/ci300367a. URL <https://doi.org/10.1021/ci300367a>. Publisher: American Chemical Society.

-
- Richard Daneman and Alexandre Prat. The Blood–Brain Barrier. *Cold Spring Harbor Perspectives in Biology*, 7(1):a020412, January 2015. ISSN 1943-0264. doi: 10.1101/cshperspect.a020412. URL <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC4292164/>.
- Li Di, Christopher Keefer, Dennis O. Scott, Timothy J. Strelevitz, George Chang, Yi-An Bi, Yurong Lai, Jonathon Duckworth, Katherine Fenner, Matthew D. Troutman, and R. Scott Obach. Mechanistic insights from comparing intrinsic clearance values between human liver microsomes and hepatocytes to guide drug design. *European Journal of Medicinal Chemistry*, 57:441–448, November 2012. ISSN 0223-5234. doi: 10.1016/j.ejmech.2012.06.043. URL <https://www.sciencedirect.com/science/article/pii/S0223523412003959>.
- Gaetano Di Chiara, Valentina Bassareo, Sandro Fenu, Maria Antonietta De Luca, Liliana Spina, Cristina Cadoni, Elio Acquas, Ezio Carboni, Valentina Valentini, and Daniele Lecca. Dopamine and drug addiction: the nucleus accumbens shell connection. *Neuropharmacology*, 47:227–241, January 2004. ISSN 0028-3908. doi: 10.1016/j.neuropharm.2004.06.032. URL <https://www.sciencedirect.com/science/article/pii/S0028390804002199>.
- Fang Du, Haibo Yu, Beiyan Zou, Joseph Babcock, Shunyou Long, and Min Li. hERGCentral: a large database to store, retrieve, and analyze compound-human Ether-à-go-go related gene channel interactions to facilitate cardiotoxicity assessment in drug development. *Assay and Drug Development Technologies*, 9(6):580–588, December 2011. ISSN 1557-8127. doi: 10.1089/adt.2011.0425.
- Jerome Eberhardt, Diogo Santos-Martins, Andreas F. Tillack, and Stefano Forli. AutoDock Vina 1.2.0: New Docking Methods, Expanded Force Field, and Python Bindings. *Journal of Chemical Information and Modeling*, 61(8):3891–3898, August 2021. ISSN 1549-960X. doi: 10.1021/acs.jcim.1c00203.
- Wenhao Gao and Connor W. Coley. The Synthesizability of Molecules Proposed by Generative Models. *Journal of Chemical Information and Modeling*, 60(12):5714–5723, December 2020. ISSN 1549-9596. doi: 10.1021/acs.jcim.0c00174. URL <https://doi.org/10.1021/acs.jcim.0c00174>. Publisher: American Chemical Society.
- Anna Gaulton, Louisa J. Bellis, A. Patricia Bento, Jon Chambers, Mark Davies, Anne Hersey, Yvonne Light, Shaun McGlinchey, David Michalovich, Bissan Al-Lazikani, and John P. Overington. ChEMBL: a large-scale bioactivity database for drug discovery. *Nucleic Acids Research*, 40(D1):D1100–D1107, January 2012. ISSN 0305-1048. doi: 10.1093/nar/gkr777. URL <https://doi.org/10.1093/nar/gkr777>.
- Kaitlyn M. Gayvert, Neel S. Madhukar, and Olivier Elemento. A Data-Driven Approach to Predicting Successes and Failures of Clinical Trials. *Cell Chemical Biology*, 23(10):1294–1301, October 2016. ISSN 2451-9448. doi: 10.1016/j.chembiol.2016.07.023.
- Oleksandr O. Grygorenko, Dmytro S. Radchenko, Igor Dziuba, Alexander Chuprina, Kateryna E. Gubina, and Yurii S. Moroz. Generating Multibillion Chemical Space of Readily Accessible Screening Compounds. *iScience*, 23(11), November 2020. ISSN 2589-0042. doi: 10.1016/j.isci.2020.101681. URL [https://www.cell.com/iscience/abstract/S2589-0042\(20\)30873-7](https://www.cell.com/iscience/abstract/S2589-0042(20)30873-7). Publisher: Elsevier.
- Quanguo He, Jun Liu, Jing Liang, Xiaopeng Liu, Wen Li, Zhi Liu, Ziyu Ding, and Du Tuo. Towards Improvements for Penetrating the Blood–Brain Barrier—Recent Progress from a Material and Pharmaceutical Perspective. *Cells*, 7(4):24, April 2018. ISSN 2073-4409. doi: 10.3390/cells7040024. URL <https://www.mdpi.com/2073-4409/7/4/24>. Number: 4 Publisher: Multidisciplinary Digital Publishing Institute.
- Tingjun Hou, Junmei Wang, Wei Zhang, and Xiaojie Xu. ADME evaluation in drug discovery. 7. Prediction of oral absorption by correlation and classification. *Journal of Chemical Information and Modeling*, 47(1):208–218, 2007. ISSN 1549-9596. doi: 10.1021/ci600343x.
- Kexin Huang, Tianfan Fu, Wenhao Gao, Yue Zhao, Yusuf H. Roohani, Jure Leskovec, Connor W. Coley, Cao Xiao, Jimeng Sun, and Marinka Zitnik. Therapeutics Data Commons: Machine Learning Datasets and Tasks for Drug Discovery and Development. June 2021a. URL <https://openreview.net/forum?id=8nvgnORnoWr>.

-
- Kexin Huang, Tianfan Fu, Lucas M. Glass, Marinka Zitnik, Cao Xiao, and Jimeng Sun. Deep-Purpose: a deep learning library for drug-target interaction prediction. *Bioinformatics (Oxford, England)*, 36(22-23):5545–5547, April 2021b. ISSN 1367-4811. doi: 10.1093/bioinformatics/btaa1005.
- Kexin Huang, Tianfan Fu, Wenhao Gao, Yue Zhao, Yusuf Roohani, Jure Leskovec, Connor W. Coley, Cao Xiao, Jimeng Sun, and Marinka Zitnik. Artificial intelligence foundation for therapeutic science. *Nature Chemical Biology*, 18(10):1033–1036, October 2022. ISSN 1552-4469. doi: 10.1038/s41589-022-01131-2. URL <https://www.nature.com/articles/s41589-022-01131-2>. Number: 10 Publisher: Nature Publishing Group.
- Abdul Karim, Matthew Lee, Thomas Balle, and Abdul Sattar. CardioTox net: a robust predictor for hERG channel blockade based on deep learning meta-feature ensembles. *Journal of Cheminformatics*, 13(1):60, August 2021. ISSN 1758-2946. doi: 10.1186/s13321-021-00541-z. URL <https://doi.org/10.1186/s13321-021-00541-z>.
- Sunghwan Kim, Jie Chen, Tiejun Cheng, Asta Gindulyte, Jia He, Siqian He, Qingliang Li, Benjamin A Shoemaker, Paul A Thiessen, Bo Yu, Leonid Zaslavsky, Jian Zhang, and Evan E Bolton. PubChem 2023 update. *Nucleic Acids Research*, 51(D1):D1373–D1380, January 2023. ISSN 0305-1048. doi: 10.1093/nar/gkac956. URL <https://doi.org/10.1093/nar/gkac956>.
- Richard M. Kostrzewa, Karolina Wydra, Malgorzata Filip, Cynthia A. Crawford, Sanders A. McDougall, Russell W. Brown, Dasiel O. Borroto-Escuela, Kjell Fuxe, and Raul R. Gainetdinov. Dopamine D2 Receptor Supersensitivity as a Spectrum of Neurotoxicity and Status in Psychiatric Disorders. *Journal of Pharmacology and Experimental Therapeutics*, 366(3):519–526, September 2018. ISSN 0022-3565, 1521-0103. doi: 10.1124/jpet.118.247981. URL <https://jpet.aspetjournals.org/content/366/3/519>. Publisher: American Society for Pharmacology and Experimental Therapeutics Section: Minireviews.
- Nikolay Kutuzov, Henrik Flyvbjerg, and Martin Lauritzen. Contributions of the glycocalyx, endothelium, and extravascular compartment to the blood–brain barrier. *Proceedings of the National Academy of Sciences of the United States of America*, 115(40):E9429–E9438, October 2018. ISSN 0027-8424. doi: 10.1073/pnas.1802155115. URL <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC6176561/>.
- Alexey Lagunin, Dmitrii Filimonov, Alexey Zakharov, Wei Xie, Ying Huang, Fucheng Zhu, Tianxiang Shen, Jianhua Yao, and Vladimir Poroikov. Computer-Aided Prediction of Rodent Carcinogenicity by PASS and CISOC-PSCT. *QSAR & Combinatorial Science*, 28(8):806–810, 2009. ISSN 1611-0218. doi: 10.1002/qsar.200860192. URL <https://onlinelibrary.wiley.com/doi/abs/10.1002/qsar.200860192>. eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1002/qsar.200860192>.
- Greg Landrum, Paolo Tosco, Brian Kelley, Ric, sriniker, gedeck, Riccardo Vianello, NadineSchneider, Eisuke Kawashima, Andrew Dalke, Dan N, David Cosgrove, Brian Cole, Matt Swain, Samo Turk, AlexanderSavelyev, Gareth Jones, Alain Vaucher, Maciej Wójcikowski, Ichiru Take, Daniel Probst, Kazuya Ujihara, Vincent F. Scalfani, guillaume godin, Axel Pahl, Francois Berenger, JL-Varjo, strets123, JP, and DoliathGavid. RDKit: Open-source cheminformatics., April 2022. URL <https://www.rdkit.org>.
- Urs H. Langen, Swathi Ayloo, and Chenghua Gu. Development and Cell Biology of the Blood-Brain Barrier. *Annual Review of Cell and Developmental Biology*, 35(1):591–613, 2019. doi: 10.1146/annurev-cellbio-100617-062608. URL <https://doi.org/10.1146/annurev-cellbio-100617-062608>. eprint: <https://doi.org/10.1146/annurev-cellbio-100617-062608>.
- Franco Lombardo and Yankang Jing. In Silico Prediction of Volume of Distribution in Humans. Extensive Data Set and the Exploration of Linear and Nonlinear Methods Coupled with Molecular Interaction Fields Descriptors. *Journal of Chemical Information and Modeling*, 56(10):2042–2052, October 2016. ISSN 1549-9596. doi: 10.1021/acs.jcim.6b00044. URL <https://doi.org/10.1021/acs.jcim.6b00044>. Publisher: American Chemical Society.

-
- Chang-Ying Ma, Sheng-Yong Yang, Hui Zhang, Ming-Li Xiang, Qi Huang, and Yu-Quan Wei. Prediction models of human plasma protein binding rate and oral bioavailability derived by using GA–CG–SVM method. *Journal of Pharmaceutical and Biomedical Analysis*, 47(4): 677–682, August 2008. ISSN 0731-7085. doi: 10.1016/j.jpba.2008.03.023. URL <https://www.sciencedirect.com/science/article/pii/S0731708508001738>.
- Fanwang Meng, Yang Xi, Jinfeng Huang, and Paul W. Ayers. A curated diverse molecular database of blood-brain barrier permeability with chemical descriptors. *Scientific Data*, 8(1):289, October 2021. ISSN 2052-4463. doi: 10.1038/s41597-021-01069-5. URL <https://www.nature.com/articles/s41597-021-01069-5>. Number: 1 Publisher: Nature Publishing Group.
- Ying Meng, Christopher B. Pople, Harriet Lea-Banks, Agessandro Abrahao, Benjamin Davidson, Suganth Suppiah, Laura M. Vecchio, Nardin Samuel, Faiza Mahmud, Kullervo Hynynen, Clement Hamani, and Nir Lipsman. Safety and efficacy of focused ultrasound induced blood-brain barrier opening, an integrative review of animal and human studies. *Journal of Controlled Release*, 309: 25–36, September 2019. ISSN 0168-3659. doi: 10.1016/j.jconrel.2019.07.023. URL <https://www.sciencedirect.com/science/article/pii/S0168365919304109>.
- Paul Merkus, Henk-Jan Guchelaar, D. Andries Bosch, and Frans W. H. M. Merkus. Direct access of drugs to the human brain after intranasal drug administration? *Neurology*, 60(10):1669–1671, May 2003. ISSN 0028-3878, 1526-632X. doi: 10.1212/01.WNL.0000067993.60735.77. URL <https://n-neurology-org.ezp-prod1.hul.harvard.edu/content/60/10/1669>. Publisher: Wolters Kluwer Health, Inc. on behalf of the American Academy of Neurology Section: Brief Communications.
- Alain Minn, Jean-François Ghersi-Egea, Rachel Perrin, Brigitte Leininger, and Gérard Siest. Drug metabolizing enzymes in the brain and cerebral microvessels. *Brain Research Reviews*, 16(1): 65–82, January 1991. ISSN 0165-0173. doi: 10.1016/0165-0173(91)90020-9. URL <https://www.sciencedirect.com/science/article/pii/0165017391900209>.
- David L. Mobley and J. Peter Guthrie. FreeSolv: a database of experimental and calculated hydration free energies, with input files. *Journal of Computer-Aided Molecular Design*, 28(7):711–720, July 2014. ISSN 1573-4951. doi: 10.1007/s10822-014-9747-x.
- Harry L Morgan. The Generation of a Unique Machine Description for Chemical Structures-a Technique Developed at Chemical Abstracts Service. *Journal of chemical documentation*, 5(2): 107–113, 1965. Publisher: ACS Publications.
- Samuel Murail, Sjoerd J. de Vries, Julien Rey, Gautier Moroy, and Pierre Tufféry. SeamDock: An Interactive and Collaborative Online Docking Resource to Assist Small Compound Molecular Docking. *Frontiers in Molecular Biosciences*, 8, 2021. ISSN 2296-889X. URL <https://www.frontiersin.org/articles/10.3389/fmolb.2021.716466>.
- R. Scott Obach, Franco Lombardo, and Nigel J. Waters. Trend Analysis of a Database of Intravenous Pharmacokinetic Parameters in Humans for 670 Drug Compounds. *Drug Metabolism and Disposition*, 36(7):1385–1405, July 2008. ISSN 0090-9556, 1521-009X. doi: 10.1124/dmd.108.020479. URL <https://dmd.aspetjournals.org/content/36/7/1385>. Publisher: American Society for Pharmacology and Experimental Therapeutics Section: Article.
- Birgit Obermeier, Richard Daneman, and Richard M. Ransohoff. Development, maintenance and disruption of the blood-brain barrier. *Nature medicine*, 19(12):1584–1596, December 2013. ISSN 1078-8956. doi: 10.1038/nm.3407. URL <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC4080800/>.
- Marcus Olivecrona, Thomas Blaschke, Ola Engkvist, and Hongming Chen. Molecular de-novo design through deep reinforcement learning. *Journal of Cheminformatics*, 9(1):48, September 2017. ISSN 1758-2946. doi: 10.1186/s13321-017-0235-x.
- William M. Pardridge. The Blood-Brain Barrier: Bottleneck in Brain Drug Development. *NeuroRx*, 2(1):3–14, January 2005. ISSN 1545-5343. URL <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC539316/>.

-
- William M Pardridge. Drug transport across the blood–brain barrier. *Journal of Cerebral Blood Flow & Metabolism*, 32(11):1959–1972, November 2012. ISSN 0271-678X. doi: 10.1038/jcbfm.2012.126. URL <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC3494002/>.
- Ann M. Richard, Richard S. Judson, Keith A. Houck, Christopher M. Grulke, Patra Volarath, Intirany Thillainadarajah, Chihae Yang, James Rathman, Matthew T. Martin, John F. Wambaugh, Thomas B. Knudsen, Jayaram Kancherla, Kamel Mansouri, Grace Patlewicz, Antony J. Williams, Stephen B. Little, Kevin M. Crofton, and Russell S. Thomas. ToxCast Chemical Landscape: Paving the Road to 21st Century Toxicology. *Chemical Research in Toxicology*, 29(8):1225–1251, August 2016. ISSN 1520-5010. doi: 10.1021/acs.chemrestox.6b00135.
- Philip Seeman. Targeting the dopamine D2 receptor in schizophrenia. *Expert Opinion on Therapeutic Targets*, 10(4):515–531, August 2006. ISSN 1472-8222. doi: 10.1517/14728222.10.4.515. URL <https://doi.org/10.1517/14728222.10.4.515>. Publisher: Taylor & Francis eprint: <https://doi.org/10.1517/14728222.10.4.515>.
- Vishal Siramshetty, Jordan Williams, Dac-Trung Nguyen, Jorge Neyra, Noel Southall, Ewy Mathé, Xin Xu, and Pranav Shah. Validating ADME QSAR Models Using Marketed Drugs. *SLAS DISCOVERY: Advancing the Science of Drug Discovery*, 26(10):1326–1336, December 2021. ISSN 2472-5552. doi: 10.1177/24725552211017520. URL <https://doi.org/10.1177/24725552211017520>. Publisher: SAGE Publications Inc STM.
- Murat Cihan Sorkun, Abhishek Khetan, and Süleyman Er. AqSolDB, a curated reference set of aqueous solubility and 2D descriptors for a diverse set of compounds. *Scientific Data*, 6(1):143, August 2019. ISSN 2052-4463. doi: 10.1038/s41597-019-0151-1. URL <https://www.nature.com/articles/s41597-019-0151-1>. Number: 1 Publisher: Nature Publishing Group.
- Svetlana M Stamatovic, Allison M Johnson, Richard F Keep, and Anuska V Andjelkovic. Junctional proteins of the blood-brain barrier: New insights into function and dysfunction. *Tissue Barriers*, 4(1):e1154641, February 2016. ISSN 2168-8362. doi: 10.1080/21688370.2016.1154641. URL <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC4836471/>.
- Jiangming Sun, Nina Jeliakova, Vladimir Chupakhin, Jose-Felipe Golib-Dzib, Ola Engkvist, Lars Carlsson, Jörg Wegner, Hugo Ceulemans, Ivan Georgiev, Vedrin Jeliakov, Nikolay Kochev, Thomas J. Ashby, and Hongming Chen. ExCAPE-DB: an integrated large scale dataset facilitating Big Data analysis in chemogenomics. *Journal of Cheminformatics*, 9(1):17, March 2017. ISSN 1758-2946. doi: 10.1186/s13321-017-0203-5. URL <https://doi.org/10.1186/s13321-017-0203-5>.
- Mengying Sun, Jing Xing, Han Meng, Huijun Wang, Bin Chen, and Jiayu Zhou. MolSearch: Search-based Multi-objective Molecular Generation and Property Optimization. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining, KDD '22*, pp. 4724–4732, New York, NY, USA, August 2022. Association for Computing Machinery. ISBN 978-1-4503-9385-0. doi: 10.1145/3534678.3542676. URL <https://dl.acm.org/doi/10.1145/3534678.3542676>.
- Kyle Swanson, Gary Liu, Denise B. Catacutan, James Zou, and Jonathan M. Stokes. Generative AI for designing and validating easily synthesizable and structurally novel antibiotics. In *37th Conference on Neural Information Processing Systems, Generative AI and Biology Workshop*, December 2023a. URL <https://openreview.net/forum?id=8nvgnORnoWr>.
- Kyle Swanson, Parker Walther, Jeremy Leitz, Souhrid Mukherjee, Joseph C. Wu, Rabindra V. Shivaraine, and James Zou. ADMET-AI: A machine learning ADMET platform for evaluation of large-scale chemical libraries, December 2023b. URL <https://www.biorxiv.org/content/10.1101/2023.12.28.573531v1>. Pages: 2023.12.28.573531 Section: New Results.
- Kyle Swanson, Gary Liu, Denise B. Catacutan, Autumn Arnold, James Zou, and Jonathan M. Stokes. Generative AI for designing and validating easily synthesizable and structurally novel antibiotics. *Nature Machine Intelligence*, 6(3):338–353, March 2024. ISSN 2522-5839. doi: 10.1038/s42256-024-00809-7. URL <https://www.nature.com/articles/s42256-024-00809-7>. Publisher: Nature Publishing Group.

-
- Oleg Trott and Arthur J. Olson. AutoDock Vina: Improving the speed and accuracy of docking with a new scoring function, efficient optimization, and multithreading. *Journal of Computational Chemistry*, 31(2):455–461, 2010. ISSN 1096-987X. doi: 10.1002/jcc.21334. URL <https://onlinelibrary.wiley.com/doi/abs/10.1002/jcc.21334>. eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1002/jcc.21334>.
- Henrike Veith, Noel Southall, Ruili Huang, Tim James, Darren Fayne, Natalia Artemenko, Min Shen, James Inglese, Christopher P. Austin, David G. Lloyd, and Douglas S. Auld. Comprehensive Characterization of Cytochrome P450 Isozyme Selectivity across Chemical Libraries. *Nature biotechnology*, 27(11):1050–1055, November 2009. ISSN 1087-0156. doi: 10.1038/nbt.1581. URL <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC2783980/>.
- Nora D. Volkow and Marisela Morales. The Brain on Drugs: From Reward to Addiction. *Cell*, 162(4):712–725, August 2015. ISSN 0092-8674, 1097-4172. doi: 10.1016/j.cell.2015.07.046. URL [https://www.cell.com/cell/abstract/S0092-8674\(15\)00962-9](https://www.cell.com/cell/abstract/S0092-8674(15)00962-9). Publisher: Elsevier.
- Ning-Ning Wang, Jie Dong, Yin-Hua Deng, Min-Feng Zhu, Ming Wen, Zhi-Jiang Yao, Ai-Ping Lu, Jian-Bing Wang, and Dong-Sheng Cao. ADME Properties Evaluation in Drug Discovery: Prediction of Caco-2 Cell Permeability Using a Combination of NSGA-II and Boosting. *Journal of Chemical Information and Modeling*, 56(4):763–773, April 2016a. ISSN 1549-960X. doi: 10.1021/acs.jcim.5b00642.
- Sheng Wang, Tao Che, Anat Levit, Brian K. Shoichet, Daniel Wacker, and Bryan L. Roth. Structure of the D2 dopamine receptor bound to the atypical antipsychotic drug risperidone. *Nature*, 555(7695):269–273, March 2018. ISSN 1476-4687. doi: 10.1038/nature25758. URL <https://www.nature.com/articles/nature25758>. Number: 7695 Publisher: Nature Publishing Group.
- Shuangquan Wang, Huiyong Sun, Hui Liu, Dan Li, Youyong Li, and Tingjun Hou. ADMET Evaluation in Drug Discovery. 16. Predicting hERG Blockers by Combining Multiple Pharmacophores and Machine Learning Approaches. *Molecular Pharmaceutics*, 13(8):2855–2866, August 2016b. ISSN 1543-8384. doi: 10.1021/acs.molpharmaceut.6b00471. URL <https://doi.org/10.1021/acs.molpharmaceut.6b00471>. Publisher: American Chemical Society.
- David Weininger. SMILES, a chemical language and information system. *Journal of Chemical Information and Computer Sciences*, 28(1):31–36, February 1988. ISSN 0095-2338. doi: 10.1021/ci00057a005. URL <https://doi.org/10.1021/ci00057a005>. Publisher: American Chemical Society.
- Zhenqin Wu, Bharath Ramsundar, Evan N. Feinberg, Joseph Gomes, Caleb Geniesse, Aneesh S. Pappu, Karl Leswing, and Vijay Pande. MoleculeNet: a benchmark for molecular machine learning. *Chemical Science*, 9(2):513–530, January 2018. ISSN 2041-6539. doi: 10.1039/C7SC02664A. URL <https://pubs.rsc.org/en/content/articlelanding/2018/sc/c7sc02664a>. Publisher: The Royal Society of Chemistry.
- Congying Xu, Feixiong Cheng, Lei Chen, Zheng Du, Weihua Li, Guixia Liu, Philip W. Lee, and Yun Tang. In silico Prediction of Chemical Ames Mutagenicity. *Journal of Chemical Information and Modeling*, 52(11):2840–2847, November 2012. ISSN 1549-9596. doi: 10.1021/ci300400a. URL <https://doi.org/10.1021/ci300400a>. Publisher: American Chemical Society.
- Youjun Xu, Ziwei Dai, Fangjin Chen, Shuaishi Gao, Jianfeng Pei, and Luhua Lai. Deep Learning for Drug-Induced Liver Injury. *Journal of Chemical Information and Modeling*, 55(10):2085–2093, October 2015. ISSN 1549-9596. doi: 10.1021/acs.jcim.5b00238. URL <https://doi.org/10.1021/acs.jcim.5b00238>. Publisher: American Chemical Society.
- Kevin Yang, Wengong Jin, Kyle Swanson, Regina Barzilay, and Tommi Jaakkola. Improving molecular design by stochastic iterative target augmentation. In *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *ICML'20*, pp. 10716–10726. JMLR.org, July 2020.

Hao Zhu, Todd M. Martin, Lin Ye, Alexander Sedykh, Douglas M. Young, and Alexander Tropsha. Quantitative Structure-Activity Relationship Modeling of Rat Acute Toxicity by Oral Exposure. *Chemical Research in Toxicology*, 22(12):1913–1921, December 2009. ISSN 0893-228X. doi: 10.1021/tx900189p. URL <https://doi.org/10.1021/tx900189p>. Publisher: American Chemical Society.

Youwen Zhuang, Peiyu Xu, Chunyou Mao, Lei Wang, Brian Krumm, X. Edward Zhou, Sijie Huang, Heng Liu, Xi Cheng, Xi-Ping Huang, Dan-Dan Shen, Tinghai Xu, Yong-Feng Liu, Yue Wang, Jia Guo, Yi Jiang, Hualiang Jiang, Karsten Melcher, Bryan L. Roth, Yan Zhang, Cheng Zhang, and H. Eric Xu. Structural insights into the human D1 and D2 dopamine receptor signaling complexes. *Cell*, 184(4):931–942.e18, February 2021. ISSN 0092-8674, 1097-4172. doi: 10.1016/j.cell.2021.01.027. URL [https://www.cell.com/cell/abstract/S0092-8674\(21\)00070-2](https://www.cell.com/cell/abstract/S0092-8674(21)00070-2). Publisher: Elsevier.

A BLOOD-BRAIN-BARRIER

The blood-brain-barrier (BBB) is comprised of specialized endothelial cells and mural cells which form the continuous nonfenestrated vasculature of the brain and spinal cord. Unlike non-CNS endothelium, various barrier functions of the cerebral capillary endothelial cells allow selective transcellular but not paracellular transport across the BBB, which allows close regulation of molecular traffic entering the brain (Abbott et al., 2006). These functions include:

1. Physical barriers, such as continuous intercellular tight junctions (TJs) formed of occludin, claudins, tricellulins, and junctional adhesion molecules; other junctional complexes, such as adherens and gap junctions; the lack of membrane fenestra, or pores; and low rates of transcellular vesicle trafficking and other mechanisms of transcytosis (Figure S1) (Abbott et al., 2010; Langen et al., 2019; Stamatovic et al., 2016).
2. Molecular barriers, such as efflux transporters to counteract the diffusion of lipophilic substances (*e.g.*, drugs or xenobiotics; see Figure S1) and low expression of leukocyte adhesion molecules (LAMs) to hinder leukocyte extravasation into the CNS parenchyma (Daneman & Prat, 2015).
3. Metabolic barriers, including intracellular (*e.g.*, monoamine oxidase and cytochrome P450) and extracellular (*e.g.*, peptidases and nucleotidases) drug-metabolizing enzymes (Abbott et al., 2006; Minn et al., 1991).
4. Extravascular barriers formed by other cell types of the neurovascular unit (NVU), such as astrocytes and pericytes – which are embedded in the basement membrane of the abluminal surface of the BBB endothelium (Obermeier et al., 2013). For example, extravascular structures such as the endothelial glycocalyx and astrocyte endfeet (Figure S1) may inhibit the delivery of large molecules across the BBB (Kutuzov et al., 2018). Further, ECM receptors in the basement membrane, including dystroglycans and integrins, mediate cell-matrix adhesion and influence barrier function (*e.g.*, by regulating TJ organization) (Langen et al., 2019).

These barrier systems prevent the passive (*i.e.*, transendothelial) diffusion of all but small gaseous molecules (*e.g.*, O₂ and CO₂) and certain lipophilic agents (*e.g.*, barbiturates, ethanol) across the BBB (Figure S1) (Abbott et al., 2006); all other transport is restricted to facilitated or active pathways. These pathways – which include carrier-mediated transport (*i.e.*, ATP-independent solute carriers) and receptor-mediated transcytosis – provision select nutrients and metabolites to satisfy the energetic and chemical requirements of the brain (Campos-Bedolla et al., 2014). Nutrient delivery notwithstanding, the capillaries of the BBB remain the least permeable in the human body.

B SUPPLEMENTAL METHODS

B.1 TRAINING DATA

The labeled datasets used to train property predictors for each design objective are described below.

Blood-brain-barrier permeability. To train a property predictor for BBB permeability, we use the Blood-Brain-Barrier Dataset (B3DB), a curated resource of 7,807 small molecules classified as either BBB permeable (BBB+) or BBB non-permeable (BBB-), with 4,956 BBB+ and 2,851 BBB- molecules included (Meng et al., 2021). BBB permeability is measured the logarithm of the brain-plasma concentration ratio:

$$\log \text{BB} = \log \frac{C_{\text{brain}}}{C_{\text{blood}}}$$

Although numerical data is included for 1,058 of the 7,807 molecules in the dataset, we seek to produce a property predictor that can be trained across a maximally diverse set of molecules to guide the search across a similarly diverse and much larger space of possible small molecules. Therefore, we train a supervised classifier to predict the binary BBB+ or BBB- label, where $\log \text{BB} > 0 = \text{BBB+}$ and $\log \text{BB} < 0 = \text{BBB-}$. To represent each molecule, we use the simplified molecular-input line-entry system, or SMILES, a specification for describing the structure of chemical species using short ASCII strings (Figure S2A) (Weininger, 1988). The final dataset included the SMILES representation of each molecule and binary classification label (*i.e.*, 0 for BBB- and 1 for BBB+).

Bioactivity against target of interest. To discover molecules that exhibit binding affinity to a specific target in the brain, we use ExCAPE-DB, a large-scale database of chemogenomics data that includes 70,850,163 bioactivity annotations of 998,131 unique chemical compounds against 1667 protein targets (Sun et al., 2017). ExCAPE-DB was assembled from the PubChem (Kim et al., 2023) and ChEMBL (Gaulton et al., 2012; Bento et al., 2014) databases, and is a valuable resource for quantitative structure-activity relationship (QSAR) modeling. Of note, ExCAPE-DB is already filtered for drug-like compounds; therefore, all molecules in the database are organic (*i.e.*, lack metal atoms), have molecular weights less than 1,000 Da, and have more than 12 heavy atoms. We download Version 2.0 of ExCAPE-DB from the Zenodo repository at DOI: 10.5281/zenodo.2543724.

In this work, we select the dopamine receptor D₂ (D2R), encoded by the *DRD2* gene in humans, as our target of interest. D2R is one of a family of five G-protein-coupled receptors that facilitate the function of the catecholamine neurotransmitter dopamine; D2R does so by coupling to the inhibitory G protein G_{i/o}, likely to regulate presynaptic firing rate (Zhuang et al., 2021; Beaulieu & Gainetdinov, 2011). D2R is the primary target for almost all clinically effective antipsychotic drugs and its function can be affected by binding of both agonist or antagonist ligands, which activate or inhibit D2R, respectively. Importantly, D2R signaling has been implicated in various neurological and neuropsychiatric diseases, including schizophrenia, tardive dyskinesia, attention-deficit hyperactivity disorder, Tourette’s syndrome, Parkinson’s disease, and substance abuse disorders (Beaulieu & Gainetdinov, 2011; Kostrzewa et al., 2018; Di Chiara et al., 2004; Volkow & Morales, 2015; Seeman, 2006). Further, D2R is expressed across the brain, with highest expression in the striatum, nucleus accumbens, and olfactory tubercle, but also in the substantia nigra, ventral tegmental area, hypothalamus, cortical regions, septum, amygdala, and hippocampus (Beaulieu & Gainetdinov, 2011). Therefore, the discovery of novel small molecules that can bind to and modulate the function of D2R may be important tools in the clinical arsenal to restore homeostatic dopaminergic signaling in challenging disorders of the CNS.

In line with previous work (Olivecrona et al., 2017; Yang et al., 2020), the ExCAPE-DB database was subset for molecules with bioactivity data against D2R. After filtering for duplicate molecules by canonical SMILES structure, 8,034 active molecules and 342,204 inactive molecules remained. Bioactivity is often measured by pIC₅₀, where IC₅₀ is a measure of the concentration of a chemical species required to inhibit a biological target *in vitro*, and $\text{pIC}_{50} = -\log_{10}(\text{IC}_{50})$. In ExCAPE-DB, compounds were classified as active against a specific target if $\text{pIC}_{50} > 5$ and inactive if $\text{pIC}_{50} < 5$. However, numerical data was only available for 293 of the 342,204 inactive compounds; therefore, we train a supervised classifier to predict the binary active or inactive label. We randomly sample 12,000 compounds from the 342,204 inactive compounds, to produce a final dataset of approximately 20,000 molecules with labeled bioactivity against D2R. Like before, the final dataset

included the SMILES representation of each molecule and binary classification label (*i.e.*, 0 for inactive and 1 for active).

Safety and toxicity. To predict safety and toxicity, we retrieve data from the Therapeutics Data Commons, a centralized resource of curated datasets and benchmarks for the development of AI models for therapeutics tasks (Huang et al., 2021a; 2022). We retrieve a range of absorption, distribution, metabolism, and excretion (ADME) datasets as well as several toxicity datasets. ADME-related benchmarks include including data on human intestinal absorption and permeability (Wang et al., 2016a; Siramshetty et al., 2021; Hou et al., 2007), absorption-related P-glycoprotein inhibition (Broccatelli et al., 2011), bioavailability (Ma et al., 2008), lipophilicity (Wu et al., 2018), solubility (Sorkun et al., 2019), hydration free energy (Wu et al., 2018; Mobley & Guthrie, 2014), plasma protein binding rate, volume of distribution at steady state (Lombardo & Jing, 2016), inhibition of various metabolism-related cytochrome P450 genes (Veith et al., 2009; Carbon-Mangels & Hutter, 2011; Cheng et al., 2012), half-life (Obach et al., 2008), and clearance (Di et al., 2012). Toxicity-related benchmarks include data on acute toxicity (Richard et al., 2016; Zhu et al., 2009), blockade of the human ether-à-go-go related gene (hERG) channel responsible for coordinating cardiovascular activity (Wang et al., 2016b; Du et al., 2011; Karim et al., 2021), mutagenicity (Xu et al., 2012), drug-induced liver injury (Xu et al., 2015), chemically-induced skin reaction (Alves et al., 2015), carcinogenicity (Cheng et al., 2012; Lagunin et al., 2009), and toxicity-related clinical trial failures (Gayvert et al., 2016). All datasets were retrieved in a harmonized format from the Therapeutics Data Commons.

B.2 PROPERTY PREDICTORS

Feature engineering. To guide search over the molecular space, computational prediction of molecular properties was facilitated by encoding the chemical space into mathematical descriptors. First, we used RDKit, an open-source cheminformatics toolkit, to compute a global pharmacophore descriptor represented as 200-element vector from the SMILES representations of each molecule (Figure S2A) (Landrum et al., 2022). Next, the SMILES representations of each molecule in the training data were converted to molecular graphs, where nodes correspond to atoms and edges correspond to chemical bonds. Based on the molecular graphs, we then created molecular fingerprints, or vector representations of the structural properties of each molecule which encode the local chemical environment of a molecule by iteratively applying a hashing function to molecular substructures (Cereto-Massagué et al., 2015). In particular, we use the Morgan fingerprint, also known as the extended-connectivity fingerprint ECFP4, a 1024-bit vector where bits are assigned based on the presence of circular substructures around each atom in a molecule (Morgan, 1965). We use a radius of 2; therefore, the hashing function is applied to all substructures within 2 bonds of each atom in the molecule.

Graph neural networks. To predict blood-brain-barrier permeability and D2R bioactivity, two separate property predictors were trained (Figure S2A). As in Swanson et al. (2024), using the Chemprop framework, we developed graph neural networks (GNNs) with three layers of message-passing, followed by a two-layer feed-forward neural network that combines the output of the GCN with the 200-element RDKit chemical descriptor to predict the likelihood of either blood-brain-barrier permeability or D2R bioactivity. We included the RDKit chemical descriptors based on our observation that structured features like molecular weight are correlated with blood-brain-barrier permeability, as measured by correlation with numerical ratio for 1,058 of the 7,807 molecules in the B3DB dataset (Figure S2B and S2C). To evaluate out-of-sample model performance while maximizing usage of the available training data, models were trained with 10-fold cross-validation with 30 epochs of training per fold and 80%-10%-10% split across training, validation, and test sets, respectively. The average prediction across all 10 folds was taken as the final likelihood score, and performance was visualized in area under the receiver operating characteristic (AUROC) curves (Figure S3).

Next, to predict absorption, distribution, metabolism, excretion, and toxicity, we evaluated two approaches. First, we used the DeepPurpose library to train our own individual property predictors for each of the ADME and toxicity (ADME-Tox) datasets described in the "Training Data" section. DeepPurpose is a library for drug-target interaction prediction that supports 15 compound and protein encoders and over 50 neural architectures (Huang et al., 2021b). We trained 2-layer message passing GNNs on the molecular graph of each compound to predict ADME-Tox features. Models were trained for 30 epochs, with a learning rate of and a batch size of 128. After training, an ag-

gregate model was assembled to take as input a single molecular query and output the results of the ADME-Tox models in a single concatenated vector. We also considered ADMET-AI, a recently developed ML platform for ADME-Tox property prediction trained on the TDC datasets described above (Swanson et al., 2023b). For each molecule, computed the weighted average of the ADMET-AI prediction scores (with a weight of 1 for non-toxic properties and -1 for toxic properties), and linearly scaled the scores from 0 to 1. In this work, we use ADMET-AI as the ADME-Tox property prediction oracle; the inclusion of our own DeepPurpose-based ADME-Tox predictors represents an interesting avenue for future work.

B.3 SYNTHEMOL-BASED MONTE CARLO TREE SEARCH

After training the property predictors, we adapt SyntheMol, a recently developed Monte Carlo Tree Search (MCTS)-based algorithm for antibiotic design, to perform a multi-objective guided traversal over an easily synthesizable molecular space (Swanson et al., 2023a; 2024). SyntheMol was designed to overcome a key challenge in generative AI for molecular design: most generative models for *de novo* molecular design often produce compounds that are intractable to synthesize in the wet lab, and thus cannot be validated experimentally (Gao & Coley, 2020). To ensure synthesizability, SyntheMol assembles compounds using a pre-processed subset of the Enamine REadily Accessible (REAL) Space, composed of 132,479 commercially available molecular building blocks with known reactivities and 13 chemical synthesis reactions (Grygorenko et al., 2020). The resulting combinatorial chemical search space spans 29.6 billion molecules that are tractable to synthesize experimentally.

In this work, we adapt the SyntheMol algorithm in Swanson et al. (2024) to perform multi-objective molecular search guided by the property prediction models for BBB permeability, D2R bioactivity, and ADME-Tox described above. Like in SyntheMol, in each iteration of the algorithm, we perform an MCTS rollout to construct a molecule using the 132,479 building blocks and 13 chemical reactions; however, we then evaluate the molecule using *both* the BBB permeability and D2R bioactivity property predictors.

In the MCTS search, molecular building blocks and their possible combinations via chemical synthesis reactions are represented by nodes in a synthesis tree. During the MCTS rollouts, each potential child node N is evaluated by a scoring function that balances both exploitation of nodes that lead to high scoring molecules as well as exploration of nodes that have rarely been visited during the search. We use the SyntheMol scoring function as described in Swanson et al. (2024); however, we adapt the molecular property score $P(N)$ to perform guided molecular search in a multi-objective setting by concurrently using the property predictors for BBB permeability, D2R bioactivity, and ADME-Tox. Let M_{BBB} be the predictor of BBB permeability, let M_{D2R} be the predictor of D2R bioactivity, and let M_{Tox} be the ADME-Tox predictor. Then, $P(N)$ is given by:

$$P(N) = \frac{1}{|N_{\text{mols}}|} \sum_{i=1}^{|N_{\text{mols}}|} \frac{M_{\text{BBB}}(N_{\text{mols}}^i) + M_{\text{D2R}}(N_{\text{mols}}^i) + M_{\text{Tox}}(N_{\text{mols}}^i)}{3}$$

Here, N_{mols}^i is the i -th molecule in node N (in case reactions produce more than 1 molecule). This approach is also inspired by MolSearch, a multi-objective MCTS-based method for molecular generation with a two-stage design which alternates between optimizing for biological properties and optimizing for non-biological properties, instead of taking the average across the property predictors as we do here (Sun et al., 2022).

We use this SyntheMol-based guided search strategy to produce a molecular library of 26,581 novel and diverse molecules. To characterize the library, we visualize the distribution of molecular building blocks and chemical synthesis pathways employed. We also investigate the molecular diversity of the library using maximum internal Tanimoto similarity as well as maximum Tversky similarity with the training data. Finally, we evaluate each of the molecules using both the BBB and D2R predictors, compare the average score to an unguided random sampling baseline, and validate top scoring hits with protein-ligand molecular docking simulation against D2R.

C SUPPLEMENTAL FIGURES

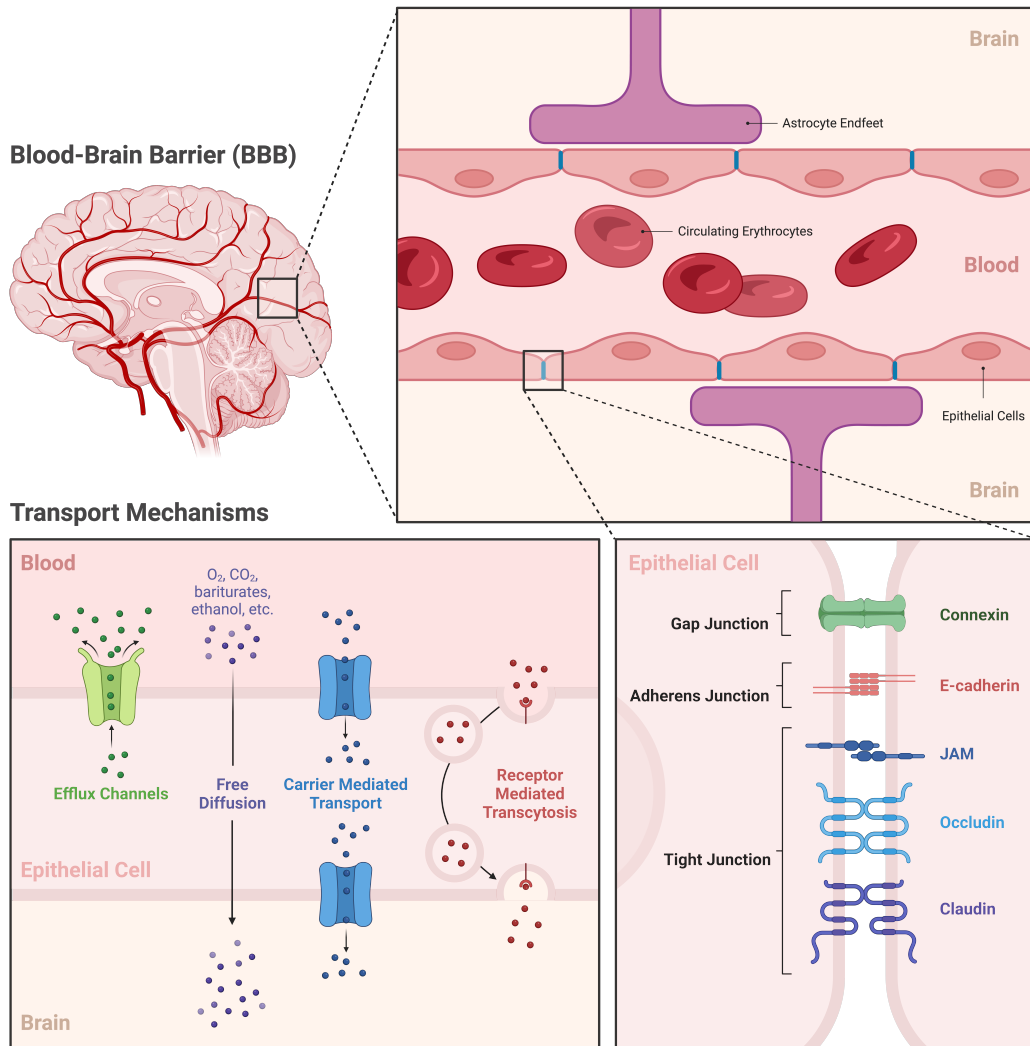


Figure S1: **Transport and barrier functions of the CNS epithelium.** Physical, molecular, metabolic, and extravascular barriers of the BBB regulate access to the CNS milieu. Created using Biorender.com.

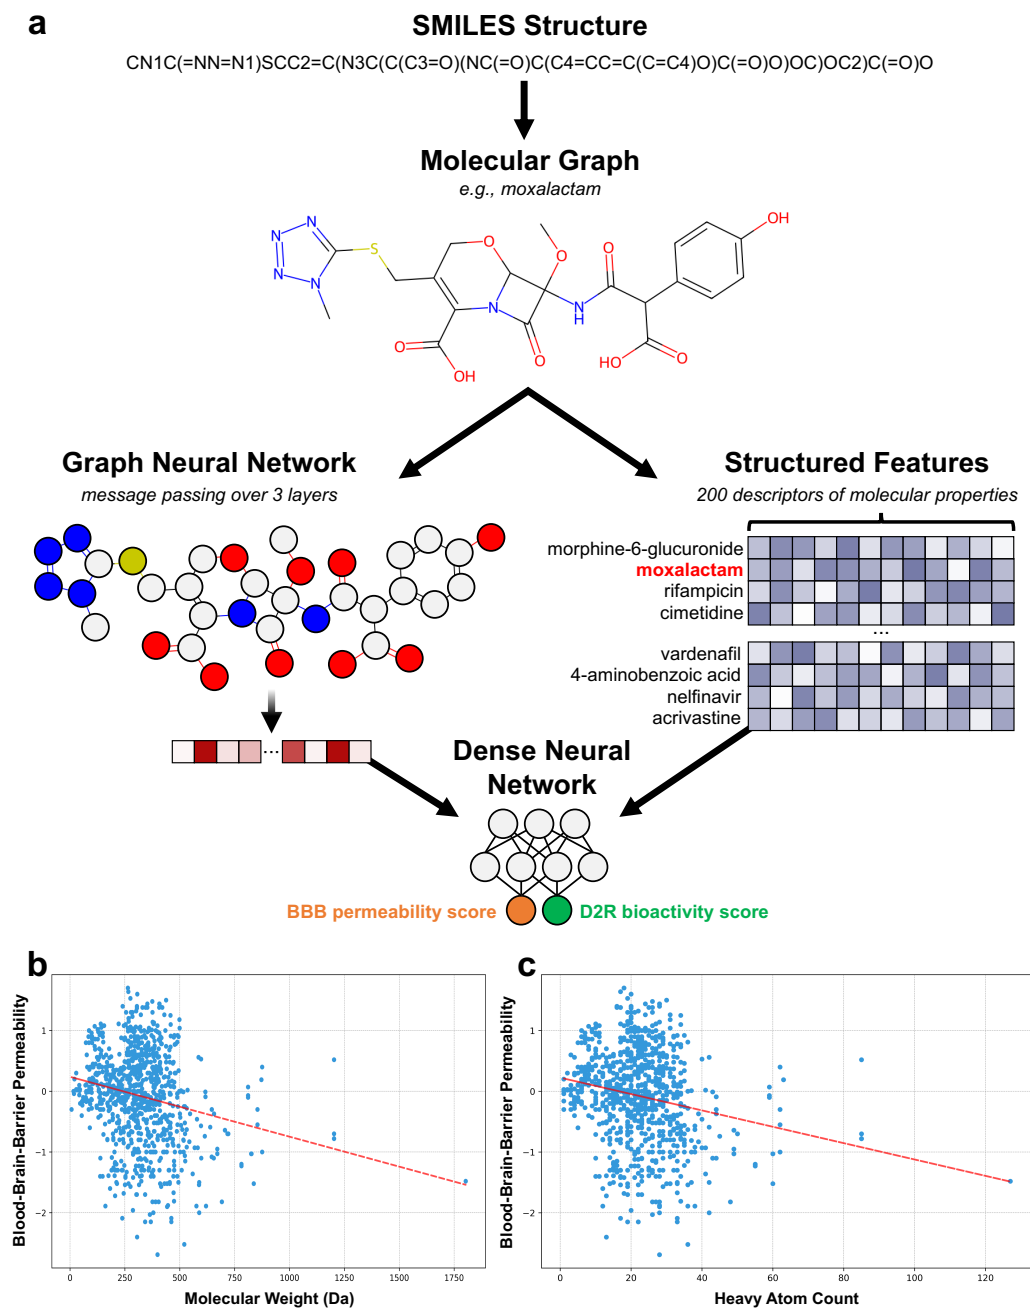


Figure S2: **Model architecture of property predictors.** To predict blood-brain-barrier permeability and D2R bioactivity, two separate property predictors were trained. **(A)** First, the SMILES representations of each molecule in the training data were converted to molecular graphs; for example, the SMILES structure and molecular graph of moxalactam, a third-generation cephalosporin antibiotic in the B3DB database, are shown here. Next, we developed graph neural networks (GNNs) with three layers of message-passing, followed by a two-layer feed-forward neural network that combined the output of the GCN with a 200-element RDKit chemical descriptor to predict the likelihood of either blood-brain-barrier permeability or D2R bioactivity. The structured chemical descriptors are included based on the observation that properties like **(B)** molecular weight and **(C)** number of heavy atoms are correlated with blood-brain-barrier permeability, as measured by correlation with the numerical blood-brain-barrier concentration ratios for 1,058 of the 7,807 molecules in the B3DB dataset for which the data was available.

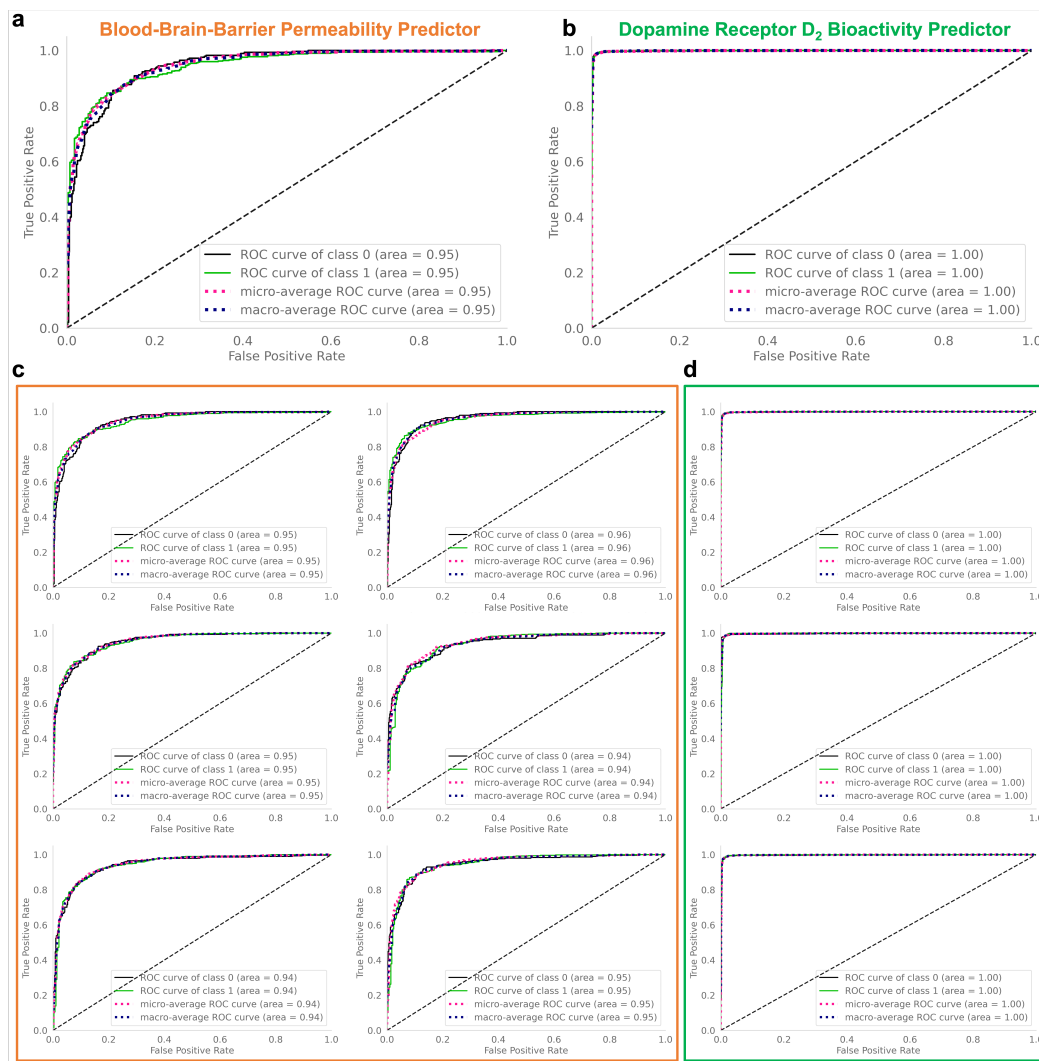


Figure S3: Property predictor performance. To evaluate out-of-sample model performance while maximizing usage of the available training data, property predictor models were trained with 10-fold cross-validation with 30 epochs of training per fold and 80%-10%-10% split across training, validation, and test sets, respectively. The average prediction across all 10 folds was taken as the final likelihood score, and model performance on the independent test set was visualized in area under the receiver operating characteristic (AUROC) curves. Here, we highlight the performance of the first cross-validation fold for the (A) blood-brain-barrier permeability predictor and (B) the dopamine receptor D₂ bioactivity predictor. We also show the performance across a set of randomly sampled folds for both the (C) BBB and (D) D2R predictors.