Machine Learning Guided Drug Discovery: Comprehensive Application in Janus Kinase Inhibitor Design

by

Yingzi Bu

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Pharmaceutical Sciences and Scientific Computing)
in The University of Michigan
2024

Doctoral Committee:

        Professor Duxin Sun, Chair
        Professor Grace Chen
        Professor James Moon
        Professor David Sept
        Professor Peter Tessier

Yingzi Bu

yingzibu@umich.edu

ORCID iD: 0000-0002-2600-8946

To my family, Wen Li, Fanshi Bu and Lanlang Bu for all of their support and love.

ACKNOWLEDGEMENTS

Throughout this long journey, filled with both challenges and moments of joy, I am deeply grateful to many individuals who have contributed to my Ph.D. studies. First and foremost, I would like to express my sincere gratitude to Dr. Duxin Sun for serving as my advisor. Dr. Sun has shown remarkable patience in guiding me through various aspects of academic life, including scientific writing and presentation skills. I also want to extend my appreciation for the opportunity to explore diverse areas of study and for the freedom to engage in interdisciplinary research.

I would also like to extend my heartfelt gratitude to my esteemed committee members for their invaluable support and guidance in improving my presentation and scientific work. I am deeply thankful to Dr. David Sept as a committee member for my Ph.D. in scientific computing. I also want to express my appreciation to Dr. Peter Tessier and Dr. James Moon, who serve as committee members for my Ph.D. in Pharmaceutical Sciences. Additionally, I extend my thanks to Dr. Grace Chen for her role as the ex officio committee member. I am profoundly grateful for the time and effort each of you has devoted to helping a new researcher solidify her work. Your contributions and feedback have been instrumental.

Chapter 2. I extend my heartfelt thanks to Bohan Zhang for his invaluable contributions to the implementation of language models. Additionally, I want to express my special gratitude to Ruoxi Gao for not only being a wonderful friend but also for engaging in enlightening discussions on mathematical equations and machine learning papers. Ruoxi, your ability to simplify complex concepts is impressive. I also

appreciate the delicious home-made Chinese cuisines (and the occasional indulgence in Latiao junk food).

Chapter 3. My sincere gratitude goes out to all the members of the Sun Lab for their exceptional contributions to my research. In particular, I would like to thank Dr. Mady Traore for his assistance in drug synthesis and the patient explanations of the intricacies of overall drug design. I also want to acknowledge Dr. Lu Wang for invaluable help with metabolite identification. Dr. Zhongwei Liu, your meticulous approach to experimental procedures and guidance on western blot techniques have been invaluable. Dr. Hongxiang Hu, thank you for generously sharing your time to teach me about animal studies and for including me in festival gatherings. Meilin Wang, your assistance with LC-MS analysis, even on weekends, has been greatly appreciated. Chengyi Li, I want to express my thanks for your consistent support with H&E staining and for inviting me to social gatherings.

Chapter 4 and overall work. I want to express special thanks for Luchen Zhang, whose support has extended to all of my papers. Luchen, your assistance with writing and your talent in creating compelling data visualizations for the Journal of Biological Chemistry paper and Chapter 4 have been invaluable. I wholeheartedly endorse your remarkable skills in using Adobe Illustrator to craft beautiful figures, which I've had the pleasure of witnessing on LinkedIn. Additionally, I extend my heartfelt gratitude to Hanning Wen for improving my presentation. I also want to thank Fanshi Bu for catching the typo in the tables, and express my appreciation to Kuangzheng Li and Ruoxi Gao for correcting terminology in the field of machine learning.

I also want to acknowledge the invaluable assistance provided by staff members Antoinette Hopper, Patrina Hardy, and Erika Zucal in handling files and paperwork.

I am deeply honored to have been admitted to the University of Michigan for my Ph.D. studies in Pharmaceutical Sciences and Scientific Computing, and obtain a certificate of Advanced Studies in Engineering, Innovation and Entrepreneurship. My time here

TABLE OF CONTENTS

LIST OF FIGURES

Figure

LIST OF TABLES

LIST OF APPENDICES

Appendix

LIST OF ABBREVIATIONS

ADME-T absorption, distribution, metabolism, excretion and toxicity

AUC Area Under Curve

DAI disease activity index

DDPM Denoising Diffusion Probabilistic Model

DMPK Drug Metabolism and Pharmacokinetics

DNN Deep Neural Network

DRL Deep Reinforcement Learning

DSS dextran sulfate sodium

DTI Drug-Target Interaction

FDA Food and Drug Administration

GI gastrointestinal

GIN Graph Isomorphism Network

GNN Graph neural networks

GRU Gated Recurrent Unit

JAK Janus kinase

KL Kullback-Leibler

LSTM Long Short-Term Memory

MAE Mean Absolute Error

ML Machine Learning

MLP Multilayer Perceptron

MT Multi Task

NLP Natural Lanugage Processing

PCA principal component analysis

PK pharmacokinetics

QED Quantitative estimate of drug-likeness

QSPR quantitative structure-property relationship

RL Reinforcement Learning

RGCN Relational Graph Convolution Network

RNN Recurrent Neural Networks

SA Synthetic Accessibility

SELFIES self-referencing embedded strings

SMILES simplified molecular input line entry systems

SOTA state-of-the-art

ST Single Task

STAR structure-tissue selectivity-activity-relationship

STAT signal transducers and activators of transcription

t-SNE t-distributed stochastic neighbor embedding

UC ulcerative colitis

VAE variational autoencoder

ABSTRACT

Conventional drug discovery is often plagued by resource-intensive processes, leading to significant attrition rates due to challenges such as efficacy shortcomings, safety concerns and suboptimal pharmacokinetics (PK). Meantime, the development of Machine Learning (ML) is heralding a transformation in this field. ML empowers precise predictions and *de novo* drug design, mitigating the necessity for costly and time-consuming experimental endeavors. In response to the challenges in drug discovery field, we introduced CoGT, an ensemble ML method adept at distinguishing JAK inhibitors from non-inhibitors. CoGT integrates conventional ML models, a graph-based model GraphVAE and a transformer-based model chemBERTa, resulting in state-of-the-art performance in predicting JAK inhibition. Subsequently, we designed a novel gastrointestinal (GI) Janus kinase (JAK) inhibitor for ulcerative colitis (UC) treatment with the application of ML and structure-tissue selectivity-activity-relationship (STAR). Employing the STAR system, we successfully designed a Class III candidate, MMT3-72, characterized by high selectivity for GI tissues and moderate potency to JAK isoforms. This innovative approach mitigated the stringent requirement for JAK isoform specificity, promising an effective and safe treatment for UC. Employing CoGT, we identified MMT3-72-M2, major metabolite of MMT3-72, as a potent JAK inhibitor. Subsequent experimental validation corroborated these predictions, emphasizing MMT3-72's weak JAK inhibition and MMT3-72-M2's effectiveness against JAK1/2 and TYK2. *In vivo* investigations further demonstrated MMT3-72's superior efficacy, showcasing its targeted action and minimal systemic toxicity. Furthermore, we established a comprehensive ML framework to evaluate

absorption, distribution, metabolism, excretion and toxicity (ADME-T) profiles cost-effectively, addressing PK and toxicity concerns inherent in drug discovery. This ML approach facilitates concurrent prediction of multiple ADME-T properties, leveraging Graph neural networks (GNN)-based models to expedite drug candidate identification. Moreover, we harnessed ML to design JAK-specific inhibitors capable of targeting specific isoforms, aiming to mitigate toxicity concerns of JAK inhibitors. This ongoing work involves *in silico* drug design by using variational autoencoder and REINFORCE algorithm. Selection of potent drug candidates will be guided by the criteria set by the STAR, predicted ADME-T profiles, subsequently validation through experimental assays. In summary, our research underscores the profound potential of ML in accelerating drug discovery, aiming to make drug discovery practices more efficient and effective in delivering promising drug candidates at the forefront of pharmaceutical research.

CHAPTER I

Introduction

Formidable challenge in drug discovery: rising expenditures, time-intensive process, yet substantial attrition rates. The prevailing approach to drug discovery involves the selection of a pertinent drug target, followed by the pursuit of candidate drugs with optimal potency against a target, favorable safety/toxicity profiles, and finally balanced PK parameters [43, 7, 135]. However, this conventional method constitutes a time-consuming and resource-intensive process. Even with such endeavors, drug attrition remains an ever-present challenge. Predominantly, the leading causes of drug attrition consistently revolve around issues such as lack of efficacy, potential toxicology/safety risks, and sub-optimal drug PK profiles [43, 28]. Further examination reveals several inherent limitations and daunting tasks in the current drug discovery landscape:

1. Drug-Target Interaction (DTI) [55, 17]: In the realm of drug discovery, chemists traditionally dedicate substantial efforts to synthesizing an array of compounds. However, this approach can be akin to a blindfolded endeavor, often yielding compounds that fail to meet the desired outcomes. The challenge is exacerbated by the absence of guiding structures, leading to a lack of clear design direction.

2. Safety and Toxicity: Even when compounds demonstrate promising efficacy, the prediction of potential toxicity remains a formidable challenge[28]. Con-

ducting experiments to assess toxicity is a valuable but imperfect approach, as unpredictability often prevails. Unforeseen toxicological intricacies can surface, giving rise to unexpected issues during clinical trials and significant financial resource depletion.

3. Drug Metabolism and Pharmacokinetics (DMPK): The process of assessing DMPK is both time-consuming and costly, involving a series of labor-intensive *in vitro* and *in vivo* experiments [81]. These experiments are essential but contribute to the overall protracted timeline and substantial financial investments required in the drug development process[54].

ML, a transformative technology enabling data-driven insights and efficiency across various domains. ML stands as a transformative force applicable across diverse domains, revolutionizing industries and research fields alike. Its adaptability and capability to discern complex patterns from vast datasets have redefined how we approach problems and make decisions. In healthcare, ML has empowered early disease diagnosis, personalized treatment plans, and drug discovery, potentially saving countless lives [140, 1]. Additionally, ML's applications extend to Natural Lanugage Processing (NLP), computer vision, climate modeling, and countless other areas, promising groundbreaking innovations and efficiencies that continue to shape the future[48, 13, 72].

In the area of drug discovery, the rise of ML has ushered in a transformative era. For instance, ML excels in prediction tasks of DTI, thus facilitating the identification of potential drug candidates with a higher possibility of potency against targets [14, 29, 136]. Additionally, ML could also predict toxicity and PK parameters, thus reducing the need for time-consuming and costly laboratory experiments[138, 87, 34, 96]. Furthermore, ML could do a one-shot generation of novel drug candidates using architectures such as variational autoencoder (VAE) and Denoising Diffusion Probabilistic Model (DDPM)[65, 117, 109]. As ML continues to advance, it is poised

to play a pivotal role in ushering in a new era of precision medicine and accelerating the discovery of life-saving pharmaceuticals.

Leveraging ML to streamline drug discovery process through precise prediction and *de novo* drug design. This research endeavors to mitigate the resource-intensive nature of conventional drug discovery methods with high attrition rates, by leveraging ML for accurate predictions and *de novo* drug design, with the overarching goal of enhancing the efficiency and effectiveness of drug discovery processes [69, 96, 65, 42, 153].

In the first project (Chapter II), we address the feasibility of JAK inhibitor prediction using ML. JAK is a significant component in numerous immune disorders and serves as a validated drug target, yet conventional drug discovery is time-consuming and expensive. To tackle this problem, we established CoGT (Conventional models + Graph-based model + Transformer-based model), an ensemble ML method for molecular property prediction, specifically for JAK inhibition prediction task. Alongside conventional ML models, we incorporated 1) a graph-based model to extract non-Euclidean molecular structure data and 2) a pretrained RoBERTa model to explore the potential of using large pretrained models on chemical properties predictions [66, 18]. To fully utilize the advantages of those models, the ensemble model CoGT was built through Multilayer Perceptron (MLP). This novel co-model has achieved state-of-the-art (SOTA) performance on JAK inhibition prediction [10].

As a prime application of CoGT in drug development, we integrated CoGT with the structure-tissue selectivity-activity-relationship (STAR) system for the development of a gastrointestinal (GI) locally activating JAK inhibitor MMT3-72 to treat ulcerative colitis (Chapter III) [124, 10, 134, 22, 119]. The design ensured high local exposure/selectivity in the GI tissues while minimizing systemic toxicity. Powered by CoGT for predicting target specificity of MMT3-72 and its metabolites, confirmational *in vitro* and *in vivo* studies were carried out successfully with minimal synthesis

efforts. We found that MMT3-72 effectively ameliorated inflammation in the DSS-induced colitis model with favorable therapeutic efficacy. Our results demonstrate that combining STAR and ML can guide the process of drug optimization, providing a potential tool for developing colon-targeted therapies.

The encouraging efficacy of MMT3-72 due to its GI selectivity highlights the necessity of proper compound exposure at the action site. However, comprehensive, yet financially and temporally taxing, PK characterization is required to obtain such information. To tackle this issue, we have established a comprehensive ML on evaluation of ADME-T profiles (Chapter IV). Unlike conventional ML endeavors of concentrating on isolated tasks, our method enables simultaneous prediction of multiple ADME-T properties by using Multi Task (MT) model architecture with dynamic weight strategy to adjust weight between tasks [96, 61]. We experimented extensively on different ML methods and found GNN-based models could achieve SOTA on some tasks [149]. This work could facilitate an efficient and cost-effective pursuit of candidate drugs with acceptable ADME-T profiles.

To date, nine JAK inhibitors have been approved by the Food and Drug Administration (FDA) yet they have toxicities or side effects. These can be attributed to the unspecific binding to unwanted targets, or ineffective concentration levels in targeted tissue areas. To overcome this, besides designing local-activating compounds demonstrated in Chapter III, we aim to design JAK-specific inhibitors targeting a single isoform without affecting other JAK isoenzymes [131]. *In silico* design of JAK-specific inhibitors was done by using VAE to encode drugs in compressed latent space [65, 42]. We used REINFORCE algorithm to guide us to find the optimal molecules based on our definition [153]. We define the reward function using CoGT (method described in Chapter II), aiming to "reward" specific JAK inhibition and "penalize" other JAK isoforms inhibition. Theoretical proof, methodology and preliminary results are shown in Chapter V (this work is ongoing). After the *de novo* drug design

by ML, we will use ML models to predict ADME-T properties (methods developed in Chapter IV), and will utilize Lipinski's rule of 5 to help guide the selection of potent drugs based on the definition of STAR system. Experimental assays will also be carried out to confirm JAK inhibition and *in vitro/in vivo* efficacy.

In summary, this research emphasizes the transformative potential of incorporating ML in drug discovery, supporting a more time-efficient, cost-effective approach to pharmaceutical research. This study also underscores the critical part played by STAR in novel drug design methodologies.

CHAPTER II

CoGT: Ensemble Machine Learning Method and Its Application on JAK Inhibitor Discovery

Yingzi Bu[1,2,#], Ruoxi Gao[3,#], Bohan Zhang[4,#], Luchen Zhang[1] and Duxin Sun[1,*]

Abstract

The discovery of new drug candidate to inhibit its intended target is a complex and resource-consuming process. Machine learning (ML) method for predicting drug-target interactions (DTI) is a potential solution to improve the efficiency. However, traditional ML approaches have limitations in accuracy. In this study, we developed a novel ensemble model CoGT for DTI prediction using multilayer perceptron (MLP), which integrated graph-based models to extract non-Euclidean molecular structure and large pretrained models, specifically chemBERTa, to process simplified molecular input line entry system (SMILES) input. The performance of CoGT was evaluated

[1]Department of Pharmaceutical Sciences, College of Pharmacy, University of Michigan
[2]Michigan Institute for Computational Discovery & Engineering, University of Michigan
[3]Department of Electrical Engineering and Computer Science, University of Michigan
[4]School of Information, University of Michigan
# These authors contributed equally to the work.
* To whom correspondence should be addressed: duxins@umich.edu
This work has been published [10].

using compounds inhibiting four Janus kinases (JAKs). Results showed that the large pretrained model, chemBERTa, was better than other conventional ML models in predicting DTI across multiple evaluation metrics, while the Graph Neural Network (GNN) was effective for prediction on imbalanced data sets. To take full advantage of the strengths of these different models, we developed an ensemble model, CoGT, which outperformed other individual ML models in predicting compounds' inhibition on different isoforms of JAKs. Our data suggest that the ensemble model CoGT has the potential to accelerate the process of drug discovery.

2.1 Introduction

Janus kinase (JAK, including JAK1, JAK2, JAK3 and TYK2) are tyrosine kinases related to the regulation of immune responses and cellular metabolism. Triggered JAK phosphorylation by molecules such as cytokines could further activate STATs (signal transducer and activator of transcription) through JAK-STAT pathway, modulating immunity and other cellular mechanisms [3]. The disturbed regulation of JAK-STAT pathway could lead to a range of malignant and nonmalignant diseases, including myelofibrosis (MF) [46], inflammatory bowel disease (IBD) and psoriatic arthritis (PA) [56]. Thus, a series of JAK inhibitors were synthesized to treat JAK-related disorders. JAK plays a vital role in immune disease and thus more research needs to be done to synthesize JAK inhibitors with better efficacy.

Janus kinases (JAKs) are the tyrosine kinases associated with the intracellular domains of type I and type II cytokine receptors.[93] The ligands binding to these receptors comprise more than 50 cytokines, growth factors and hormones, thereby initiating fundamental biological processes of cell apoptosis, proliferation, and development.[105, 131]

However, all approved JAK inhibitors have commonly observed side effects, which may be due to their pan-inhibition of different JAK isoforms.[131] The JAK families,

JAK1, JAK2, JAK3, and TYK2, have seven homology domains (JH), where JH1 serves as the kinase domain that phosphorylates downstream signaling proteins.[131, 51] Most JAK inhibitors are designed to compete with adenosine triphosphate (ATP) for the binding site in the JH1 kinase domain. However, the JH1 is a highly evolutionarily conserved domain,[51]which makes it difficult to develop isoform-selective inhibitors. Thus, an *a priori* tool to predict JAK selectivity of designed molecules will be of valuable help to develop more isoform specific inhibitors to reduce their side effects. Machine learning can significantly improve the efficiency and accuracy of these processes.

To develop an isoform specific inhibitor, high throughput screening and lead compound optimization are usually performed, which are time-consuming and not economically efficient. On the other hand, machine learning methods, such as Random Forest (RF)[8], Support Vector Machine (SVM)[83], K Nearest Neighbors (KNN)[98], and extreme Gradient Boosting (XGBoost)[16], could be applied to accelerate these processes. For instance, XGBoost has shown promising prediction on JAK2 inhibitors, using fingerprint as drug molecule representation. [150] In our study, We would like to explore the ability of graph neural network (GNN) models, which directly used molecule graph as input. In addition, we attempted to experiment transformer-based model on JAK inhibition prediction using SMILES as input. By integrating different aspects of the ML methods, we developed an ensemble model CoGT (Conventional ML models + Graph-based models + Transformer-based models), aiming to leverage the predicting ability for drug-target interactions (DTI). This novel method could be further applied and validated on drug development of other molecular targets.

2.2   Materials and methods

Data Preparation. This data set was extracted from ChEMBL [38, 39], BindingDB [41], PubChem [133, 64] and Liu et al [76]. We removed duplicated drugs or drugs

Table 2.1: Number of molecules collected in each JAK category

molecule number	JAK1	JAK2	JAK3	TYK2
total	7373	10161	7722	2424
active (label 1)	5606	6846	5250	1627
inactive (label 0)	1767	3315	2472	797

with controversial labels (e.g., one drug with both active and inactive labels) in the data set based on Compound ID (CID) or compounds' SMILES (Simplified Molecular Input Line Entry System) strings. More than 2,130,000 compounds were extracted from ChEMBL without label and used to pretrain neural network structured models mentioned later. For four types of JAKs, the number of compounds collected was summarized in Table 2.1. The threshold of active drugs is those with IC_{50}, inhibition, EC_{50}, Ki to a certain JAK below 10μM. For model training, the data sets were randomly split into training, validation and test sets by a 8:1:1 ratio. To make sure the training, validation and test set are the same when training all categories of models, we set the random state seed as 42. Therefore, during all of our training processes including later comodel training, compound information from test set did not leak.

Molecular Fingerprints Calculation. A fingerprint of a molecule is a list of binary bits, which contains information of drug substructure. For instance, each bit of the fingerprint list could be a Boolean determination of certain element presence, ring structure, or atom pairing [63]. In our work, all molecules were represented by MACCS fingerprints (166 bits). Those fingerprints were calculated based on a compound's SMILES using the RDKit package.

Model building.

SVM. It defines a margin or decision plane to separate data from different classes. Here we used MACCS fingerprints as features. We tried different SVM methods: linear, poly, rbf and sigmoid. Model evaluation for SVM are summarized in table S1 and area under the curve (AUC) results are shown in figure S1. We found that SVM

poly performs best overall on 4 JAKs while SVM sigmoid shows worst performance (nearly random guessing). Thus, we chose SVM poly for model comparison and later co-model building.

Random Forest. RF consists of individual decision trees and each tree is trained on a subversion of the data set. We used fingerprints as features, and the number of trees was optimized for each JAK category. The n_estimator for JAK1, JAK2, JAK3, and TYK2 are 53, 91, 48, and 13, respectively.

Extreme Gradient Boosting. XGBoost [15] is a scalable machine learning system for tree boosting. Compared with RF, XGBoost has a range of adjustable parameters to optimize for each JAK category. We did a grid search on parameters listed in Table S2 for each JAK and built the final XGBoost model using the determined optimal parameters.

Graph Model. To leverage the natural structure of chemicals, we attempted to solve the problem by using Graph Neural Networks (GNN)[66, 130, 149] because of its performance and interpretability [152]. For each chemical molecule, we built one graph by taking atoms as nodes and chemical bonds as edges. For each node, we used 6 attributes of the atom: 1) atomic number 2) atom degree 3) formal charge 4) hybridization 5) aromatic 6) chiral tag. We removed all hydrogen atoms so the related nodes, edges and atom degrees would not be included in the graph construction. The architecture figure of our graph model is shown in Figure 2.1 A.

Apart from basic connection information between atoms, bond types (single, double, triple, and aromatic) are employed as edge relations between two nodes to supply more edge information for the constructed graphs, and relational Graph Convolutional Network (RGCN) is applied to adapt this data structure. The node embedding is initialized by embedding layer and the RGCN convoluton layers update the node embedding using the neighbours and relation information:
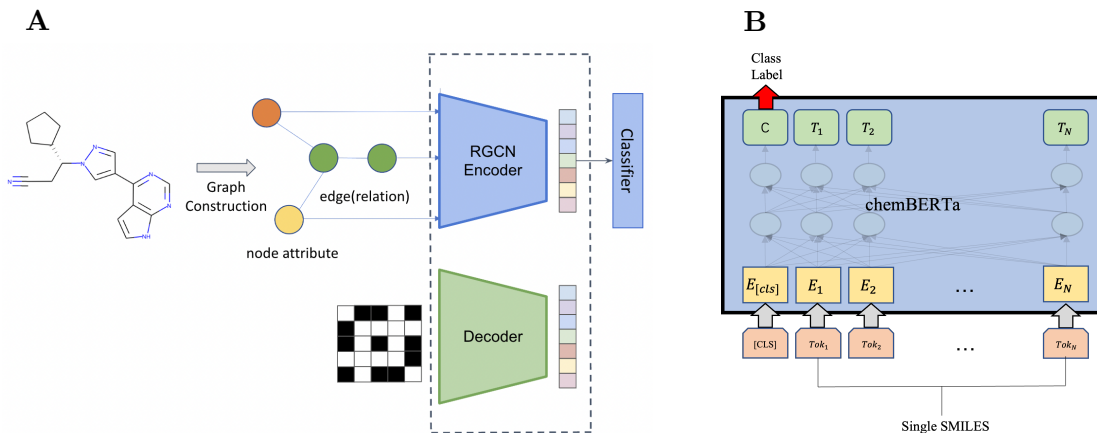
Figure 2.1: Model architecture. A: Graph Model, where encoder is a RGCN that maps drug to latent space Z, decoder is implemented by $\sigma(ZZ^T)$. B: ChemBERTa model.

$$\mathbf{h_i}^{(k+1)} = \sigma(\sum_{r\in R}\sum_{j\in\mathcal{N}_r(i)}\frac{1}{|\mathcal{N}_r(i)|}\mathbf{W_r}^{(k)}\mathbf{h_j}^{(k)} + \mathbf{W_0}\mathbf{h_i}^{(k)}), \tag{2.1}$$

where $\mathbf{h_i}^{(k)}$ is the node embedding of the ith node after the kth layer, R is the relation set, $\mathcal{N}_r(i)$ denotes the neighbours set that has r relation with of the ith node.

A two-layer RGCN[108] is experimented with embedding dimension 4 (embedding layer), hidden dimension 64 (the 1st RGCN convolutional layer), output dimension 128 (the 2nd RGCN convolutional layer).

Variational autoencoder (VAE)[65] is considered as a pretraining tool to eliminate effect of unbalanced data and train a more robust model. We modified variational graph auto-encoders (VGAE)[67] as a relation-employed graph auto-encoders (Graph-VAE). The RGCN is the simple inference model, i.e., encoder. The generative model is given by an inner product between latent variables to learn adjacency matrix. We optimize the variational bound on negative log likelihood:

$$\mathcal{L} = \mathbb{E}_x[-\mathbb{E}_{\mathbf{z}\sim q_\theta(\mathbf{z}|\mathbf{x})}\log(A\mid\mathbf{z}) + D_{KL}(q_\theta(\mathbf{z}\mid\mathbf{x}),\mathcal{N}(\mathbf{0},\mathbf{I}))], \tag{2.2}$$

where q_θ is the encoder distribution, $\mathbf{z}$ represents the drug representation on latent

11

space, and A denotes the adjacency matrix of a drug graph.

After GraphVAE training process, the pretrained encoder followed by a global attention pool layer and a linear layer is fine-tuned as a JAK classifier.

chemBERTa. Large pretrained neural networks especially transformer-based have made breakthrough in many domains like language[26], vision[101], as well as protein prediction[58]. However, the progress of chemical property prediction using large transformers is not significant compared to these domains. Previous work didn't fully take advantage of the capacity of large transformer models as they are either pretrained on smaller language models like recurrent neural networks or tuned on smaller datasets and narrow applications like reaction predictions[123] which may cause the over-fitting of models and unable to generalize to other tasks. In recent years, a new chemical transformer called chemBERTa[18] makes one of the first attempts to systematically evaluate large transformers on molecular property prediction tasks. As shown in Figure 2.1 B, the chemBERTa is originally pretrained on 77 million unique SMILES of chemicals from PubChem on RoBERTa[77] with a SMILE-based tokenizer[112] to predict corresponding Morgan fingerprints and then applied to several downstream properties prediction tasks. The SMILE-based tokenizer, first designed for another pretrained transformer model[112], tokenizes SMILES strings more reasonably than regular hard tokenization which turns SMILES into single letters but may lose information when two or more consecutive letters should stay in the integrity. The backbone model, RoBERTa, shares a similar architecture with BERT[26] but shows more robust performances specifically on classification tasks under different training strategies from BERT. The natural of RoBERTa can be a good fit for the chemical property predictions which are usually classifications.

Here, we fine-tune a chemBERTa with additional two million SMILES as mentioned in Data Preparation above where the inputs to the chemBERTa are SMILES of chemicals and the targets are MACCS fingerprints. Even though the original chem-

BERTa is pretrained to predict Morgan fingerprints, we believe the transformation from MACCS fingerprint to Morgan fingerprint can be handled easily by deep neural net models. Also, to be consistent and able to easily ensemble with other methods we explored in this paper, we need to use MACCS fingerprints in the downstream JAK classification task, so we decided to also use MACCS fingerprints in the pre-training stages. In this stage, the model is pretrained in 30 epochs as the loss starts to converge. The optimizer is AdamW[78] with $1e^{-4}$ learning rate and $1e^{-2}$ weight decay. We added a linear layer on top of the pretrained model to fine-tune and cross-validate the JAK classification dataset. In the stage of tuning for JAK prediction task, the model is trained in 20 epochs and the learning rate is $1e^{-5}$ while other settings keep the same as pretraining. For both stages, the batch size is 16 and a 0.5 dropout was applied before the final linear layer.

CNN. To compare with the large pretrained model, we implemented a Convolutional Neural Network (CNN) as a neural baseline model. The CNN model was also pretrained on the same set of data as the chemBERTa and then fine-tuned on the JAK dataset. In the CNN architecture, we had 3 convolution kernels with kernel sizes 1, 2, and 3 as unigram, bigram and trigram filters, which is analogous to common settings in language tasks. The output of 3 kernels after max-pooling were concatenated together to be fed into a final linear prediction layer. The embedding dimension of each character is 256. The output size of all 3 convolution layers is 128. The activation function is LeakyReLu and the dropout rate is 0.25. In pretraining, the model is trained in 20 epochs and the optimizer is stochastic gradient descent with 0.9 learning rate and $1e^{-2}$ weight decay. In the stage of tuning for JAK prediction task, the learning rate is 0.1 while other settings keep the same as pretraining. For both stages, the batch size is 1024.

Baseline Model. In this study, we chose K-nearest neighbor (KNN) to evaluate data set as a simple base model. By utilizing a simple model, one could examine

the data set and validation with rapid feedback. We would use the base model's performance to contrast with other models [57]. We used the MACCS fingerprints as inputs for fingerprint-based non-neural classification models.

CoGT. To fully utilize the advantages of conventional ML models, graph-based models and transformer-based models, we built a co-model CoGT using simplified multilayer perceptron (MLP). In detail, predicted probabilities of compounds calculated by SVM, RF, XGBoost, GraphVAE and chemBERTa were taken as input, and probability calculated by sigmoid function was the output by using SGD optimizer to minimize the weighted BCE loss.

Model Evaluation. All models listed above were evaluated on test sets. Model performance was evaluated based on accuracy, active recall (or sensitivity, SE), negative recall (or specificity, SP), weighted accuracy (average of SE and SP), Matthew's correlation coefficient (MCC), F1 score, AUC and average precision (AP). The equations to calculate each metric are listed below, in which TP is true positive, TN, true negative, FP, false positive, and FN is false negative. In AP formula, R_n and P_n denote the precision and recall at the nth threshold respectively.

$$\text{Accuracy} = \frac{\text{TP} + \text{TN}}{\text{TN} + \text{TP} + \text{FP} + \text{FN}} \tag{2.3}$$

$$\text{Precision} = \frac{\text{TP}}{\text{TP} + \text{FP}} \tag{2.4}$$

$$\text{Recall(SE)} = \frac{\text{TP}}{\text{TP} + \text{FN}} \tag{2.5}$$

$$\text{SP} = \frac{\text{TN}}{\text{TN} + \text{FP}} \tag{2.6}$$

$$\text{Weighted Accuracy} = \frac{\text{SE} + \text{SP}}{2} \tag{2.7}$$

$$\text{MCC} = \frac{\text{TP} \times \text{TN} - \text{FN} \times \text{FP}}{\sqrt{(\text{TP} + \text{FP})(\text{TP} + \text{FN})(\text{TN} + \text{FN})(\text{TN} + \text{FP})}} \tag{2.8}$$

$$F1 = \frac{2(\text{Precision} \times \text{Recall})}{\text{Precision} + \text{Recall}} \tag{2.9}$$

$$AP = \sum_n (R_n - R_{n-1})P_n \tag{2.10}$$

2.3 Results and discussion

2.3.1 Chemical Diversity Analysis

To visualize the chemical diversity of our data sets, principal component analysis (PCA) was performed on all molecules collected with MACCS fingerprint as input. If the features' relationship is nonlinear, PCA may not perform well to cluster data. Therefore, we also performed t-distributed stochastic neighbor embedding (t-SNE) to avoid PCA under fitting. Figure 2.2 shows that the distribution of active and inactive compounds for all 4 JAKs are still overlapped for both PCA and t-SNE, which suggests that in the chemical space, active and inactive compounds could not be separated easily since some of them share similar structure and properties.

In addition, we did similarity quantification for all JAKs using Tanimoto similarity [2]. MACCS fingerprints of each drug in the data set were used to calculate Tanimoto similarity index. As shown in Figure 2.3 A-D, the similarity for all molecules is relatively low in 4 JAKs, suggesting that molecules in our data set have a wide distribution with rather diverse structures. Furthermore, we examined active molecules for 4 JAKs and the similarity is also not high overall shown in Figure 2.3 E-H. These suggest that our data set is representative and our models have strong flexibility to identify active molecules from inactive ones.

2.3.2 Performance Evaluation and Comparison of Models

We summarized our models' evaluation in Table 2.2. Due to the imbalance of our data, we would focus more on weighted accuracy than accuracy when evaluating

15
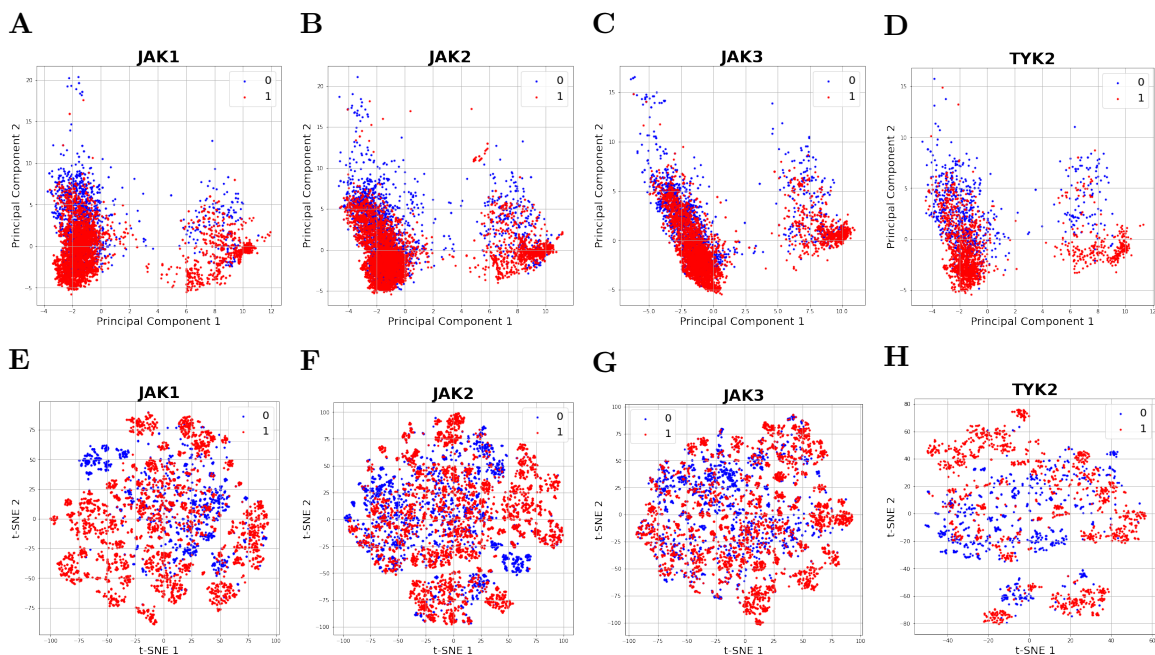
Figure 2.2: Data visualization based on MACCS fingerprint with PCA and t-SNE. (A, B, C, D) PCA for JAK1, JAK2, JAK3, and TYK2, respectively; (E, F, G, H) t-SNE for JAK1, JAK2, JAK3 and TYK2, respectively. Blue and red dots represent noninhibitors and inhibitors, respectively.

the performance of different models. For all 4 types of JAKs, neural models significantly outperforms other conventional models in more than half of all metrics and performs comparable to the best performances on other metrics. This implies the potential ability of transformer-based and graph-based models to better grasp chemical structure information and be applied to downstream chemical tasks. The pretrained CNN model performs poorly compared with chemBERTa and RGCN which are also pretrained. This suggests that CNN is not only too small to take advantage of a large volume of pretrained data and the lack of ability to extract helpful structural information from SMILE inputs.

Other than graph-based model and transformer-based model, traditional ML method SVM and tree-based models (RF and XGBoost) also perform well on JAK inhibition prediction. For instance, SVM performs well on JAK3 data sets and XG-Boost shows impressive performance on TYK2. Compared with the base model,

Table 2.2: Results of test set in JAK1, JAK2, JAK3 and TYK2. The best performances of each metric are shown in bold.

target	model	acc	w_acc	precision	recall	SP	F1	AUC	MCC	AP
JAK1	KNN	0.942	0.920	0.965	0.960	0.881	0.962	0.920	0.835	0.957
	SVM	0.958	0.939	0.972	0.974	0.905	0.973	0.974	0.880	0.990
	RF	0.954	0.932	0.969	0.972	0.893	0.970	0.986	0.868	0.996
	XGBoost	0.955	0.937	0.972	0.970	0.905	0.971	0.989	0.873	0.997
	CNN	0.744	0.720	0.887	0.765	0.674	0.821	0.765	0.392	0.888
	GraphVAE	0.902	0.924	0.988	0.884	0.964	0.933	0.948	0.770	0.986
	chemBERTa	0.957	0.938	0.972	0.972	0.905	0.972	0.989	0.877	0.997
	CoGT	**0.989**	**0.985**	**0.993**	**0.993**	**0.978**	**0.993**	**0.999**	**0.970**	**1.000**
JAK2	KNN	0.908	0.877	0.947	0.933	0.821	0.940	0.877	0.743	0.935
	SVM	0.923	0.893	0.952	0.947	0.839	0.950	0.943	0.782	0.981
	RF	0.924	0.896	0.954	0.947	0.845	0.951	0.948	0.786	0.979
	XGBoost	0.905	0.878	0.907	0.956	0.800	0.931	0.953	0.781	0.973
	CNN	0.668	0.623	0.747	0.760	0.486	0.753	0.646	0.248	0.748
	GraphVAE	0.901	0.900	0.948	0.902	0.899	0.924	0.965	0.783	0.981
	chemBERTa	0.896	0.887	0.930	0.913	0.860	0.922	0.950	0.766	0.973
	CoGT	**0.975**	**0.974**	**0.986**	**0.977**	**0.971**	**0.981**	**0.996**	**0.943**	**0.998**
JAK3	KNN	0.882	0.836	0.926	0.921	0.750	0.923	0.836	0.667	0.914
	SVM	0.879	0.836	0.927	0.916	0.756	0.921	0.912	0.662	0.972
	RF	0.878	0.824	0.920	0.923	0.726	0.921	0.926	0.652	0.978
	XGBoost	0.867	0.830	0.895	0.919	0.740	0.907	0.926	0.673	0.965
	CNN	0.696	0.500	0.696	1.000	0.000	0.821	0.503	N/A	0.699
	GraphVAE	0.894	0.889	0.946	0.901	0.877	0.923	0.956	0.755	0.972
	chemBERTa	0.875	0.849	0.912	0.910	0.789	0.911	0.943	0.698	0.976
	CoGT	**0.970**	**0.969**	**0.986**	**0.970**	**0.969**	**0.978**	**0.993**	**0.930**	**0.997**
TYK2	KNN	0.855	0.772	0.892	0.925	0.619	0.908	0.880	0.571	0.941
	SVM	0.866	0.833	0.931	0.893	0.774	0.911	0.893	0.638	0.957
	RF	0.882	0.808	0.907	0.944	0.673	0.925	0.923	0.651	0.967
	XGBoost	0.942	0.931	0.959	0.959	0.903	0.959	0.975	0.862	0.987
	CNN	0.718	0.593	0.716	0.960	0.225	0.820	0.733	0.289	0.812
	GraphVAE	0.951	0.945	0.970	0.959	0.931	0.965	0.977	0.883	0.991
	chemBERTa	0.926	0.891	0.923	0.977	0.806	0.949	0.981	0.819	0.993
	CoGT	**0.988**	**0.985**	**0.987**	**0.994**	**0.977**	**0.990**	**0.999**	**0.973**	**0.999**
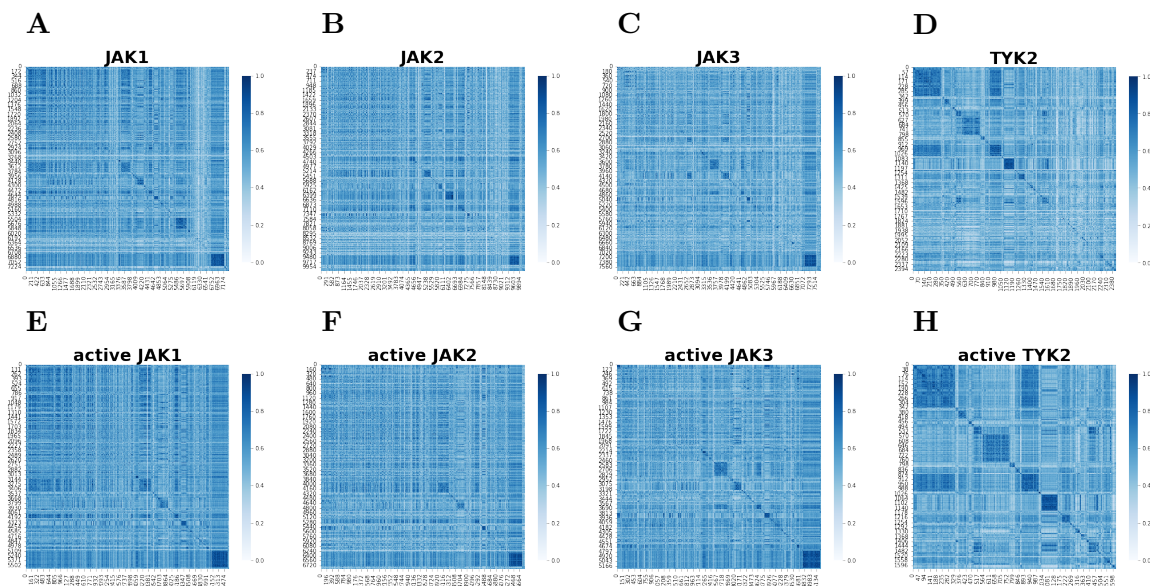
Figure 2.3: Data visualization based on Tanimoto similarity. (A, B, C, D) Tanimoto similarity for all compounds in JAK1, JAK2, JAK3, and TYK2 data sets, respectively; (E, F, G, H) Tanimoto similarity for JAK1, JAK2, JAK3 and TYK2 inhibitors, respectively.

Table 2.3: Results of test set in JAK1, JAK2, JAK3 and TYK2 on CoCM (Comodel using conventional models, i.e., SVM, RF, XGBoost), and CoGT (Comodel using conventional, graph and transformer-based models, i.e., SVM, RF, XGBoost, GraphVAE, chemBERTa).

target	model	acc	w_acc	precision	recall	SP	F1	AUC	MCC	AP
JAK1	CoCM	0.982	0.975	0.988	0.989	0.961	0.988	0.998	0.952	0.999
	CoGT	**0.989**	**0.985**	**0.993**	**0.993**	**0.978**	**0.993**	**0.999**	**0.970**	**1.000**
JAK2	CoCM	0.966	0.958	0.971	0.979	0.937	0.975	0.993	0.921	0.996
	CoGT	**0.975**	**0.974**	**0.986**	**0.977**	**0.971**	**0.981**	**0.996**	**0.943**	**0.998**
JAK3	CoCM	0.949	0.949	0.977	0.948	0.951	0.962	0.987	0.884	0.994
	CoGT	**0.970**	**0.969**	**0.986**	**0.970**	**0.969**	**0.978**	**0.993**	**0.930**	**0.997**
TYK2	CoCM	0.977	0.972	0.975	0.990	0.955	0.982	0.998	0.951	0.999
	CoGT	**0.988**	**0.985**	**0.987**	**0.994**	**0.977**	**0.990**	**0.999**	**0.973**	**0.999**
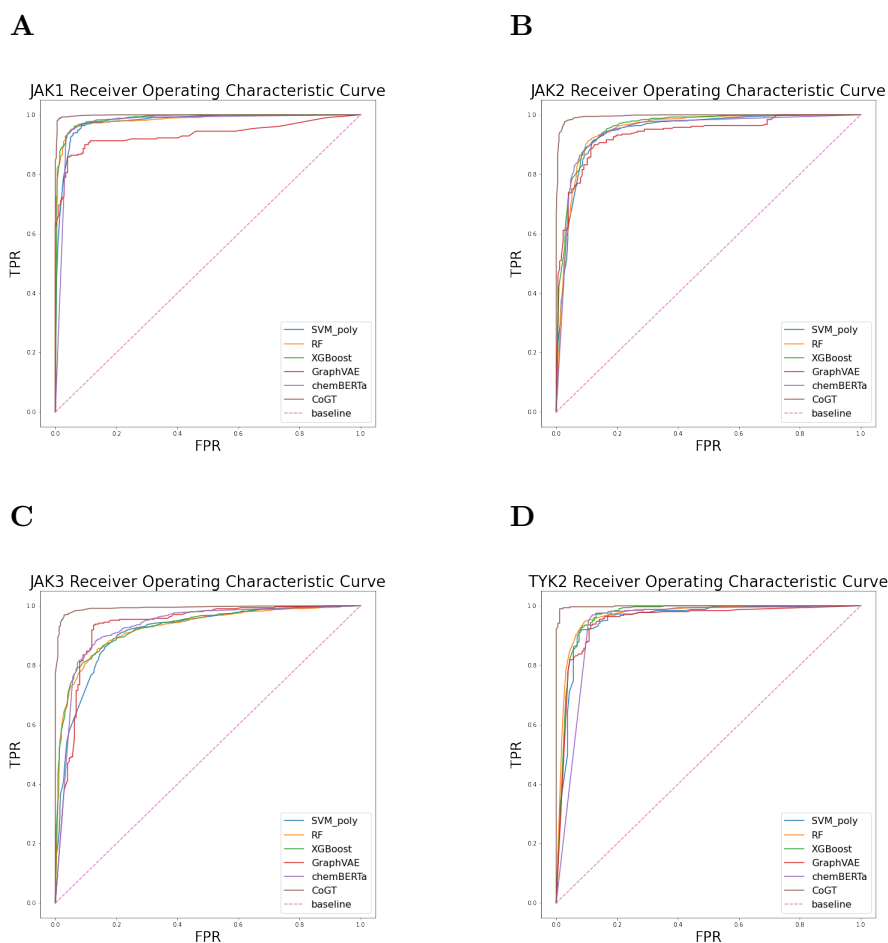
Figure 2.4: The AUC-ROC Curves of 5 selected models on (A) JAK1, (B) JAK2, (C) JAK3, and (D) TYK2 test sets.

SVM, RF and XGBoost achieved high weighted accuracy on all 4 JAK prediction tasks and those 3 algorithms were chosen for later comodel training.

To fully utilize the advantages of all different models, the ensemble models CoGT were built on all 4 JAKs with MLP as second-level model via stacking technique. The stacking technique builds a two-level model, first level contains SVM, RF, XGBoost, GraphVAE, and chemBERTa to estimate a probability of a drug being a JAK inhibitor as an intermediate prediction, the second level is a MLP which takes the prediction of three models in the first level as input to achieve final prediction [143]. We visualized the normalized weight of the co-model in Figure 2.5. Each model weight and bias
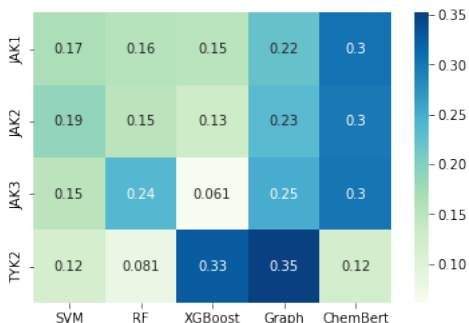
Figure 2.5: Normalized Weights Visualization for Co-Model CoGT

for each JAK could be found in Table S3. Our results show that CoGT performs impressively among all 4 JAK inhibition prediction tasks, scoring the highest for all metrics listed in Table 2.2. For AUC-ROC (Receiver Operating Characteristic) curves on test sets shown in Figure 2.4, Co-model CoGT outperforms among all other models, with AUC score nearly equals to 1 for all 4 JAKs.

As simple machine learning methods may already achieve similar performance compared with state-of-art ML methods [57], we examined the performance of CoCM (Comodel using conventional models only, i.e., SVM, RF and XGBoost). Results shown in Table 2.3 indicates that our comodel CoGT exceeds the accuracy of conventional ML comodels. This demonstrates that incorporating models which utilize different ways of molecule representation could extract more information than simply using fingerprint based conventional ML models.

We also examined the structure similarity between compounds that our model gives wrong prediction. Tanimoto similarity did not reveal a common substructure between wrongly predicted molecules, as shown in Figure S2.

2.3.3 Comparison with previous models for JAK-related ML methods

There are several works on JAK inhibitor prediction including deep learning models and traditional learning models as summarized in Table 2.4. Previous related work used XGBoost to predict JAK2 inhibition activity and RF models were also utilized

Table 2.4: Previous work and comparison with ours. FP: Fingerprint

Method	Molecular representation	Task category	Target category
MTATFP [134]	Graph	regression	JAK1, JAK2, JAK3, TYK2
MolGNN [76]	Graph	classification	JAK1, JAK2, JAK3
RF [21]	FP	classification	multi kinases
XGBoost[150]	FP	classification, regression	JAK2
CoGT (Ours)	FP, graph, SMILES	classification	JAK1, JAK2, JAK3, TYK2

to built RF models to predict JAK inhibition[150, 21]. We also used XGBoost and RF in our model building and our data showed that both models performed well compared with base model, yet our comodel CoGT shows better performance on all metrics.

Graph-based model methods has recently emerged in the field of JAK inhibitor discovery[134, 76]. These work demonstrates the promising predicting power of graph-based models on JAK inhibitor discovery and design. As a consequence, we also included graph-based model graphVAE in our comodel building.

Several recent studies[151, 102, 32] have investigated using machine learning techniques to address JAK related problems. However, these endeavors either involve a combination of the aforementioned methods or do not directly predict JAK types but rather do an effect exploration of JAK inhibitors. Consequently, we do not incorporate them in our experiment comparisons.

Overall, our comodel CoGT is the only model which tries to incorporate different representation information from a compound. Previous work mainly focused on using fingerprint or graph based representation, neglecting the possibility that more information could be extracted through different representation. Here, we not only include fingerprint and graph representation, but also utilize large pretrained transformer-based models to extract information directly from SMILES strings.
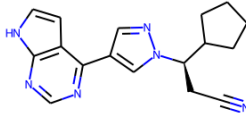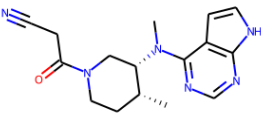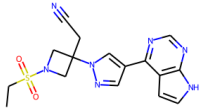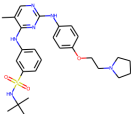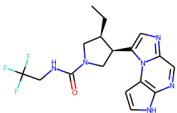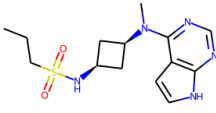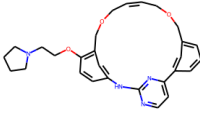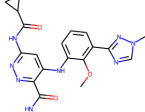
2.3.4   Co-Model CoGT Prediction on Approved Drugs

To further validate our ensemble model CoGT, we used approved drugs and drugs in clinical trials as input and predicted their JAK inhibition. Results are summarized in Table 2.5.

Our model prediction aligns well with real-world JAK inhibition for most drugs. For the four out of eight FDA-approved JAK inhibitors (i.e., Ruxolitinib, Tofacitinib, Baricitinib, Upadacitinib), our co-model gives accurate prediction on their inhibition profiles, which cannot be achieved by one single model. Besides, to further explore the phenomenon of active cliff, which is the phenomenon that compounds with similar structures may have significant efficacy difference [121], we examined our data set to see whether there are similar structures with FDA-approved drugs, and compounds highly similar to the approved drugs are shown in Figure 2.6. Results show that especially for JAK1-JAK3, there exists active cliff in our data set, i.e., compounds highly similar to approved drugs yet may have different levels of potency, and our model provides quite accurate prediction.

Among the remaining four approved inhibitors, the activity of two inhibitors with similar core structures (Fedratinib and Abrocitinib) is incorrectly predicted on JAK3 and JAK2, respectively. Our model also predicts that Pacritinib is a JAK1 noninhibitor based on the threshold of 10 μM. We can observe that the wrong predictions all happen at the values of $\sim$1 μM. Such discrepancy may be partially explained by the fact that most kinase-targeting small molecules are ATP-competitive inhibitors, and thus their measured IC$_{50}$ can be greatly affected by the concurrent ATP concentration. For example, in one previous work elucidating the JAK2 binding sites of Fedratinib, the authors showed that the IC$_{50}$ values of Fedratinib were measured to be 4.9 nM and 90 nM at the corresponding ATP concentrations of 10 μM and 100 μM, respectively. [62]

Table 2.5: Probability as inhibitors based on CoGT prediction and their IC_{50} on FDA-approved drugs for 4 JAKs (Drugs existing in the training sets are marked with *).

| Drug Name | Structure | $IC_{50}(\mu M)$ | | | |
| | | Probability of Inhibitor | | | |
		JAK1	JAK2	JAK3	TYK2
RUXOLITINIB [19]		6.4×10^{-3}	8.8×10^{-3}	0.487	0.0301
		0.931	0.896	0.885	0.834
TOFACITINIB [19]		0.0151*	0.0774	0.055	0.489
		0.905	0.689	0.899	0.848
BARICITINIB [19]		4×10^{-3}*	6.6×10^{-3}	0.787*	0.061
		0.941	0.643	0.895	0.83
FEDRATINIB [139]		0.105*	3×10^{-3}	1*	0.405
		0.951	0.911	0.103	0.841
UPADACITINIB [95]		0.047*	0.12*	2.304	4.69
		0.952	0.930	0.826	0.811
ABROCITINIB [148]		0.029*	0.803	>15	1.25
		0.958	0.094	0.108	0.840
PACRITINIB [141]		1.28	0.023*	0.52	0.05
		0.233	0.942	0.704	0.850
DEUCRAVACITINIB [144]		> 10	> 10	> 10	2×10^{-4}
		0.924	0.464	0.762	0.834

Therefore, inhibitors with weaker activity (i.e., IC_{50} values closer to the set threshold) may exhibit opposing categorizations depending on the testing conditions. Given

that the conditions applied in inhibition assays (i.e., ATP concentrations) can be slightly different across different research groups, while the collected IC_{50} values in the data sets do not necessarily include such information, a more consistent reporting format of IC_{50} values will be of valuable help in eliminating such uncertainty.

The remaining inhibitor Deucravacitinib shows least satisfactory prediction, where our model indicates it to be inhibitor on JAK1, JAK3 and TYK2, while it only inhibits TYK2. This discordance is most possible due to the unique incorporation of deuterium into the compound. Such tiny replacement of three hydrogen atoms into isotope deuterium may not be universally incorporated in the available data sets and thus the prediction accuracy suffers.[88]

To further analyze the 5 wrong predictions among all 32 predictions for approved drugs, we searched similar structure among wrongly predicted compounds in separate data set and most similar compounds are shown in Figure 2.7. Results show that there exists compounds wrongly predicted whose structures are similar to FDA-approved drugs. Especially for JAK1 data set, there is a compound highly similar to Pacritinib with Tanimoto similarity as high as 0.963. Thus, our model has difficulty giving accurate predictions for those moieties with high structural similarities, and more labeled compounds should be collected during training process.

2.4 Conclusions

In this research, we developed an ensemble model, called CoGT, which combined multiple machine learning models to achieve better accuracy than any individual model in predicting DTI for four JAK isoforms. We first compiled a comprehensive data set for JAK inhibitors. Using this data set, we compared different ML methods in predicting JAK inhibition, which includes a graph-based model (RGCN applied GraphVAE), a pre-trained RoBERTa model (chemBERTa), and traditional machine learning models. Our experiments revealed that the graph model was superior to
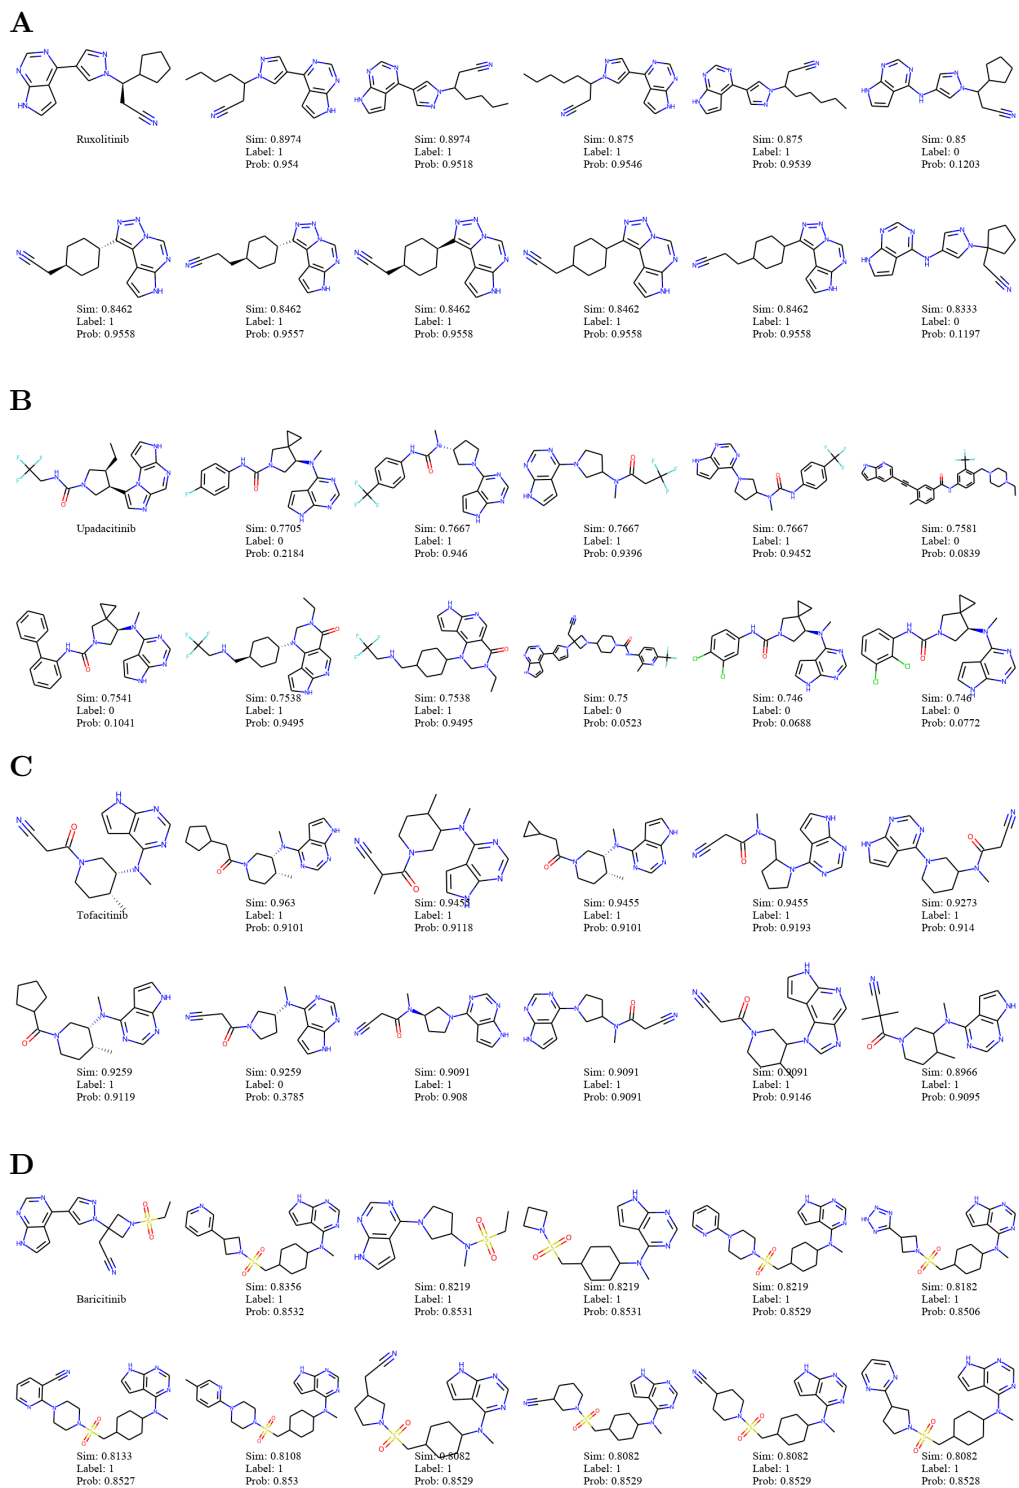
Figure 2.6: Compounds with high structure similarities compared to FDA-approved drugs. A, B, C, and D each represents the grouped compounds with high structure similarity to Ruxolitinib, Upadacitinib, Tofacitinib, Baricitinib, respectively. Sim: Tanimoto similarity; Label: True label; Prob: model predicting probability.
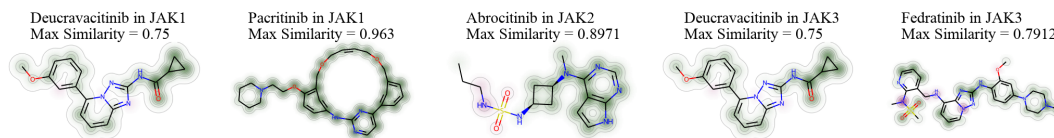
Figure 2.7: Similarity visualization between wrongly predicted molecules in Table 2.5 and the most similar drug compound in dataset.

conventional ML methods to effectively extract structural information for all JAK inhibitors. In addition, large pre-trained transformer-based model chemBERTa could also be effective for chemical predictions of these JAK inhibitory structures. Traditional models such as SVM, RF, and XGBoost performed well, despite their relatively low computational costs. By fully leveraging the strengths of various models, our ensemble mode CoGT performed best prediction accuracy on DTI of JAK inhibitors. Further improvement can be achieved by optimizing parameters in GraphVAE model for better description of the bond and atom types, as well as by utilizing the distribution of chemical fragments in the data set for the model training.

2.5 Data availability

Data sets, code, python package version are available at https://github.com/yingzibu/JAK_ML. Other data are available from the corresponding authors with reasonable request.

2.6 Supporting information

The details of SVM results and XGBoost parameters, the stacking parameters, the evaluation on the validation set, the Tanimoto similarity visualization of wrongly predicted molecules by CoGT.

2.7   Acknowledgement

Author Contributions

D.S. directed this study. Y.B., R.G. and B.Z. contributed equally to design the experiment and run the models. L.Z. performed data analysis. All authors have contributed on writing and reviewing the manuscript, and have given approval to the final version.

Supporting Information

Table S2.1: SVM preliminary results of test set in JAK1, JAK2, JAK3 and TYK2

target	model	acc	w_acc	precision	recall	SP	F1	AUC	MCC	AP
JAK1	SVM linear	0.899	0.839	0.921	0.951	0.727	0.935	0.934	0.707	0.975
	poly	0.950	0.934	0.971	0.964	0.903	0.967	0.977	0.860	0.992
	rbf	0.934	0.902	0.953	0.962	0.842	0.957	0.973	0.813	0.990
	sigmoid	0.771	0.504	0.770	1.000	0.009	0.870	0.547	0.082	0.772
JAK2	SVM linear	0.809	0.744	0.810	0.935	0.552	0.868	0.744	0.548	0.801
	poly	0.901	0.878	0.908	0.947	0.809	0.927	0.944	0.773	0.965
	rbf	0.898	0.874	0.905	0.946	0.801	0.925	0.937	0.767	0.959
	sigmoid	0.671	0.509	0.670	0.997	0.021	0.801	0.524	0.089	0.661
JAK3	SVM linear	0.801	0.725	0.818	0.918	0.532	0.865	0.847	0.501	0.914
	poly	0.871	0.847	0.907	0.907	0.787	0.907	0.912	0.694	0.955
	rbf	0.852	0.810	0.876	0.916	0.704	0.896	0.900	0.641	0.949
	sigmoid	0.696	0.500	0.696	1.000	0.000	0.821	0.503	N/A	0.699
TYK2	SVM linear	0.860	0.829	0.877	0.920	0.738	0.898	0.913	0.677	0.947
	poly	0.915	0.899	0.928	0.948	0.850	0.938	0.952	0.807	0.963
	rbf	0.909	0.891	0.922	0.945	0.838	0.933	0.956	0.793	0.971
	sigmoid	0.730	0.638	0.745	0.908	0.369	0.818	0.678	0.336	0.736

Table S2.2: XGBoost grid search parameters

XGBoost parameters	scalar used
max_depth	[3 , 6 , 1 0]
learning_rate	[0.01, 0.05, 0.1]
n_estimators	[10, 50, 100, 500, 1000]
colsample_bylevel	[0.3, 0.7, 1]
colsample_bytree	[0.3, 0.7, 1]
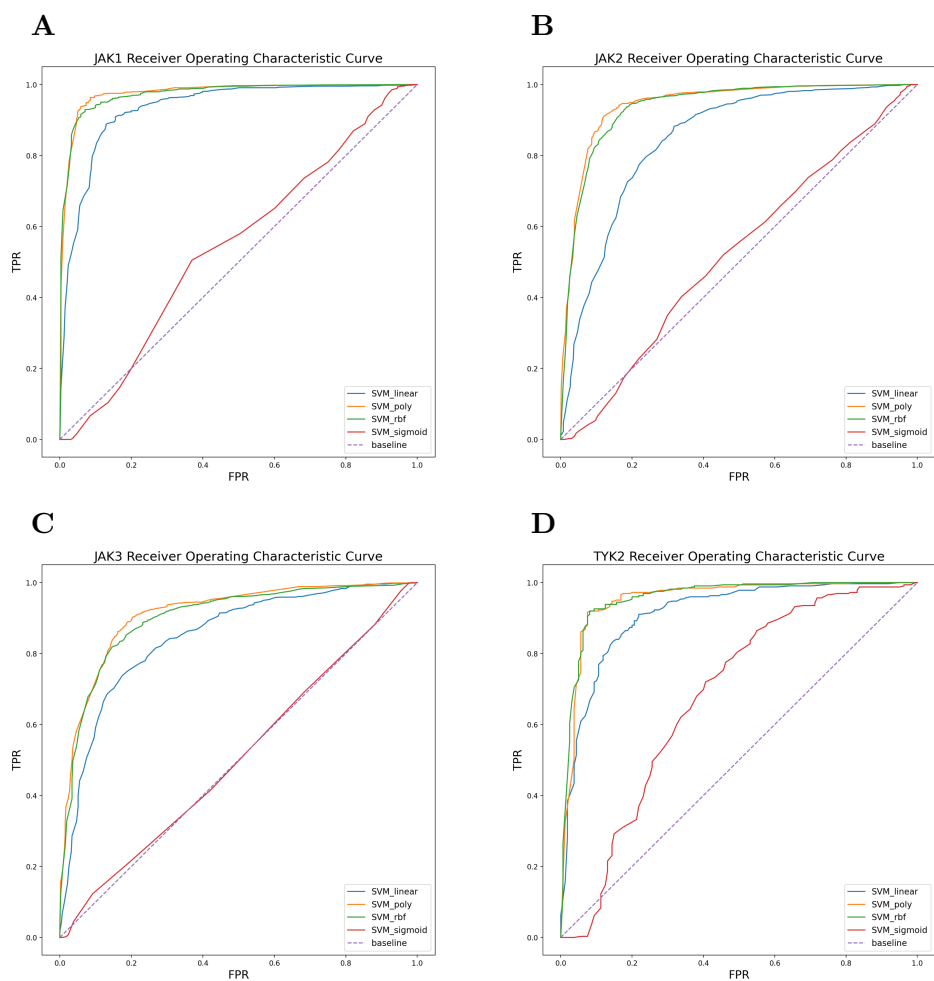
Figure S2.1: The AUC-ROC Curves of 5 selected models using SVM on (A) JAK1, (B) JAK2, (C) JAK3, and (D) TYK2 test sets.

Table S2.3: Co-Model CoGT construction: model weight and bias calculated by MLP

model weight	JAK1	JAK2	JAK3	TYK2
SVM poly	0.8725	1.1116	0.7654	0.4161
RF	0.8156	0.8958	1.1997	0.2859
XGBoost	0.8013	0.7605	0.3108	1.1439
GraphVAE	1.1736	1.3776	1.2645	1.2353
chemBERTa	1.5972	1.7598	1.5275	0.4284
bias	-2.1325	-2.955	-2.5899	-1.7482

Table S2.4: Results of **VALIDATION** set in JAK1, JAK2, JAK3 and TYK2. The best performances of each metric are shown in bold and the second best performances are italicized.

target	model	acc	w_acc	precision	recall	SP	F1	AUC	MCC	AP
JAK1	KNN	0.942	0.919	0.963	0.962	0.877	0.963	0.919	0.838	0.956
	SVM poly	0.950	0.934	0.971	0.964	0.903	0.967	0.977	0.860	0.992
	RF	0.953	0.929	0.965	*0.974*	0.883	0.970	0.987	0.867	0.996
	XGBoost	0.954	*0.939*	0.973	0.966	0.912	0.970	*0.989*	0.871	*0.997*
	CNN	0.744	0.720	0.887	0.765	0.674	0.821	0.765	0.392	0.888
	GraphVAE	0.899	0.918	*0.984*	0.881	*0.955*	0.930	0.942	0.767	0.983
	chemBERTa	*0.955*	0.936	0.970	0.972	0.900	*0.971*	0.987	*0.874*	0.995
	CoGT	**0.989**	**0.985**	**0.993**	**0.993**	**0.978**	**0.993**	**0.999**	**0.970**	**1.000**
JAK2	KNN	0.814	0.756	0.816	0.931	0.582	0.870	0.846	0.565	0.892
	SVM poly	0.901	0.878	0.908	0.947	0.809	0.927	0.944	0.773	0.965
	RF	0.908	0.882	0.907	*0.961*	0.803	0.933	0.952	0.790	0.966
	XGBoost	*0.912*	0.889	0.913	0.960	0.817	*0.936*	*0.959*	*0.801*	*0.976*
	CNN	0.668	0.623	0.747	0.760	0.486	0.753	0.646	0.248	0.748
	GraphVAE	0.864	0.854	0.914	0.882	0.826	0.898	0.928	0.695	0.955
	chemBERTa	0.897	*0.893*	*0.939*	0.902	*0.884*	0.921	0.957	0.773	*0.976*
	CoGT	**0.975**	**0.974**	**0.986**	**0.977**	**0.971**	**0.981**	**0.996**	**0.943**	**0.998**
JAK3	KNN	0.865	0.835	0.896	0.912	0.757	0.904	0.835	0.677	0.878
	SVM poly	0.871	0.847	0.907	0.907	0.787	0.907	0.912	0.694	0.955
	RF	0.860	0.828	0.891	0.911	0.745	0.901	0.924	0.665	0.963
	XGBoost	0.865	0.829	0.889	*0.920*	0.738	0.904	0.928	0.674	0.966
	CNN	0.696	0.500	0.696	1.000	0.000	0.821	0.503	N/A	0.699
	GraphVAE	0.850	0.839	0.907	0.871	0.807	0.889	0.912	0.663	0.946
	chemBERTa	*0.877*	*0.859*	*0.917*	0.905	*0.813*	*0.911*	*0.945*	*0.712*	*0.975*
	CoGT	**0.970**	**0.969**	**0.986**	**0.970**	**0.969**	**0.978**	**0.993**	**0.930**	**0.997**
TYK2	KNN	0.903	0.866	0.890	*0.975*	0.756	0.931	0.866	0.779	0.885
	SVM poly	0.915	0.899	0.928	0.948	0.850	0.938	0.952	0.807	0.963
	RF	0.928	0.908	0.929	0.966	0.850	0.947	0.966	0.835	0.974
	XGBoost	*0.936*	*0.917*	*0.935*	0.972	*0.863*	*0.953*	0.973	*0.854*	0.985
	CNN	0.718	0.593	0.716	0.960	0.225	0.820	0.733	0.289	0.812
	GraphVAE	0.915	0.903	0.921	0.948	0.858	0.935	0.963	0.816	0.980
	chemBERTa	0.915	0.880	0.899	*0.985*	0.775	0.940	*0.980*	0.808	*0.990*
	CoGT	**0.988**	**0.985**	**0.987**	**0.994**	**0.977**	**0.990**	**0.999**	**0.973**	**0.999**
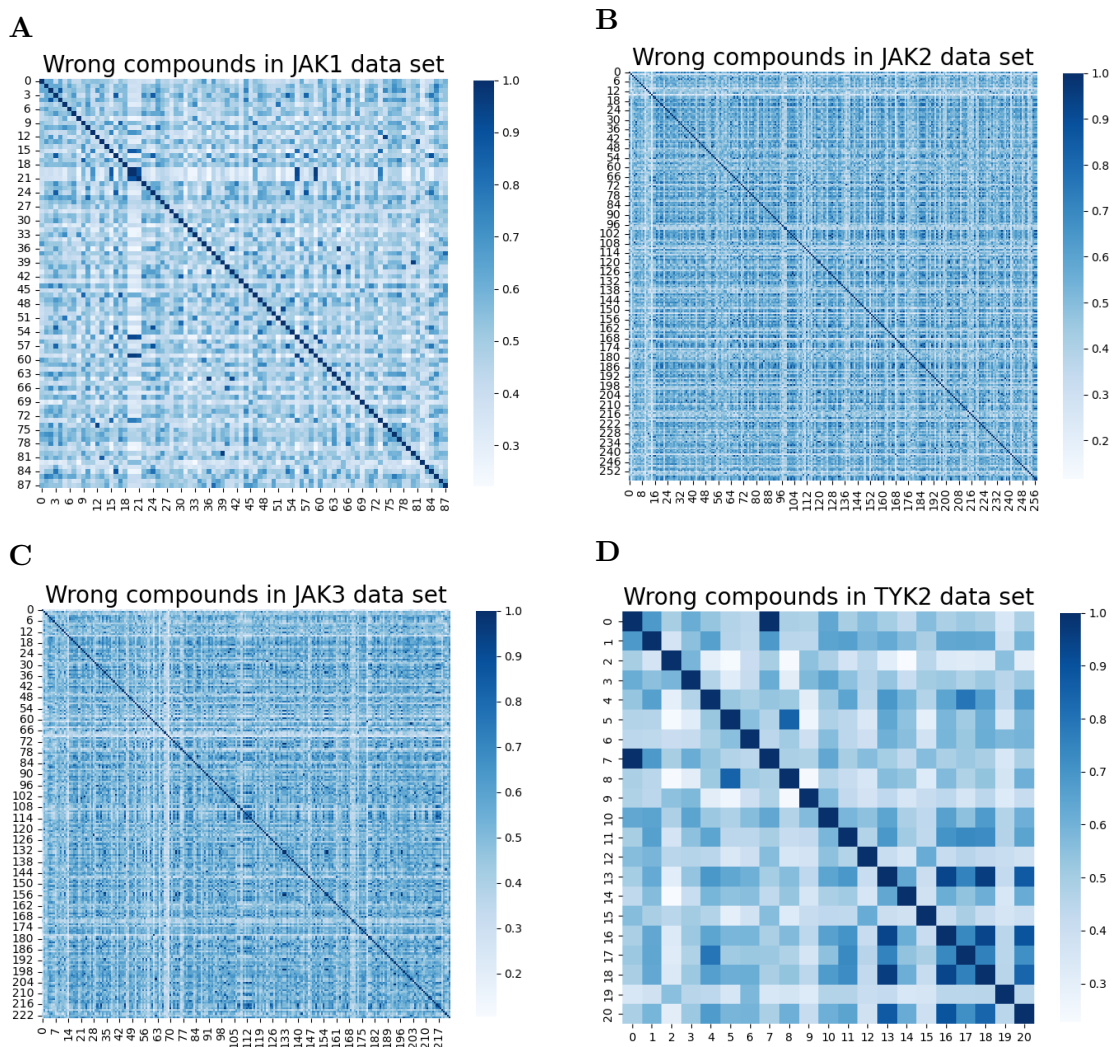
Figure S2.2: Data visualization based on Tanimoto similarity for compounds predicted wrongly by CoGT on dataset of (A) JAK1, (B) JAK2, (C) JAK3 and (D) TYK2.

CHAPTER III

A Gastrointestinal (GI) Locally-Activating Janus kinase (JAK) Inhibitor to Treat Ulcerative Colitis

Yingzi Bu[1,2,#], Mohamed Dit Mady Traore[1,#], Luchen Zhang[1,#], Lu Wang[1], Zhongwei Liu[1], Hongxiang Hu[1], Meilin Wang[1], Chengyi Li[1] and Duxin Sun[1,*]

Abstract

In this study, we integrated machine learning (ML), structure-tissue selectivity-activity-relationship (STAR), and wet lab synthesis/testing to design a gastrointestinal (GI) locally JAK inhibitor for ulcerative colitis treatment. The JAK inhibitor achieves site-specific efficacy through high local GI tissue selectivity, while minimizing the requirement for JAK isoform specificity to reduce systemic toxicity. We used the ML model (CoGT) to classify whether the designed compounds were inhibitors or noninhibitors. Then we used the regression ML model (MTATFP) to predict their IC_{50} against related JAK isoforms of predicted JAK inhibitors. The ML model predicted MMT3-72, which was retained in the GI tract, to be a weak JAK1 inhibitor,

---

[1]Department of Pharmaceutical Sciences, College of Pharmacy, University of Michigan
[2]Michigan Institute for Computational Discovery & Engineering, University of Michigan
# These authors contributed equally to the work.
* To whom correspondence should be addressed: duxins@umich.edu
This work has been published [11].

while MMT3-72-M2, which accumulated in only GI tissues, was predicted to be an inhibitor of JAK1/2 and TYK2. ML docking methods were applied to simulate their docking poses in JAK isoforms. Application of these ML models enabled us to limit our synthetic efforts to MMT3-72 and MMT3-72-M2 for subsequent wet lab testing. The kinase assay confirmed MMT3-72 weakly inhibited JAK1 and MMT3-72-M2 inhibited JAK1/2 and TYK2. We found that MMT3-72 accumulated in the GI lumen, but not in GI tissue or plasma; but released MMT3-72-M2 accumulated in colon tissue with minimal exposure in the plasma. MMT3-72 achieved superior efficacy and reduced p-STAT3 in DSS-induced colitis. Overall, the integration of ML, the STAR system, and wet lab synthesis/testing could minimize the effort in optimization of a JAK inhibitor to treat colitis. This site-specific inhibitor reduces systemic toxicity by minimizing the need for JAK isoform specificity.

3.1 Introduction

Ulcerative colitis (UC) affects more than 1-2 million patients in the US.[23, 113, 59] UC usually starts from the rectum and extends to the colon,[91] where the inflammation is restricted to the innermost layer of the intestine (mucosa), resulting in ulceration and bloody diarrhea.[113, 20, 30] UC needs life-long treatment,[113] and thus the treatment options require not only adequate efficacy but also minimal systemic toxicity.

Recently, inhibition of Janus kinases (JAK1, JAK2, JAK3, and tyrosine kinase 2 [TYK2]) has emerged as a new therapeutic approach for the treatment of UC.[105] The activation of JAKs enables the recruitment and phosphorylation of signal transducers and activators of transcription (STAT) family. Phosphorylated STATs (pSTATs) translocate to the nucleus to induce the expression of chemokines and cytokines in the UC.[89] Several orally bioavailable JAK inhibitors (tofacitinib and upadacitinib) have been developed and approved for the treatment of UC. For instance, tofacitinib,

a pan-JAK inhibitor, showed excellent efficacy in the treatment of moderate to severe UC.[73, 114] However, tofacitinib and all other JAK inhibitors have black box warnings for serious side effects,[33] which include 1) a high rate of major adverse cardiovascular events (MACE) (cardiovascular death, myocardial infarction, stroke), 2) arterial and venous thrombosis and pulmonary embolism, 3) malignancies of lymphomas and lung cancers, and 4) increased risk of serious infections leading to death.

Different approaches have been explored to reduce the systemic toxicity of JAK inhibitors in the treatment of UC. The first approach is to use a modified release formulation of JAK inhibitors to increase GI local concentration and reduce oral absorption. Although the modified release formulation of tofacitinib delays drug release, increases drug concentration in the GI tract and reduces drug exposure in the systemic circulation, it is still unable to alleviate the above serious side effects.[74] The second approach is to develop JAK isoform-specific inhibitors, such as upadacitinib and filgotinib targeting only one or two isoforms, to prevent their toxicity.[104, 125] Unfortunately, the safety review by the FDA has found that all JAK inhibitors, regardless of pan-JAK inhibitors or isoform-selective inhibitors, have similar severe adverse effects, and thus these black box warnings were added to the labels of multiple JAK inhibitors in 2021.[33] A third approach is to develop a TYK2-specific inhibitor for the treatment of UC. Deucravacitinib was developed to bind to the pseudokinase (JH2) domain of TYK2 for reduced toxicity. Although clinical testing of the TYK2-specific inhibitor showed positive efficacy in treating psoriasis[71, 94], this drug failed in treating UC.[5] Therefore, there is a critical need to develop JAK inhibitors with superior efficacy but minimal toxicity for the long-term treatment of UC. In this study, to overcome the above limitations, we integrated machine learning (ML), structure-tissue selectivity-activity relationship (STAR), and wet lab synthesis/testing to design a GI locally-activating JAK inhibitor to treat UC. The GI locally-activating JAK inhibitor achieves site-specific efficacy through high local gastrointestinal (GI) tissue

selectivity while minimizing the requirement for JAK isoform specificity to reduce systemic toxicity.

We utilized a STAR system to design JAK inhibitors for the treatment of UC, by concurrently taking into consideration two factors: the drug's tissue selectivity and target potency/specificity based on the drug's structure-tissue selectivity-activity relationship (STAR).[37, 50, 75, 124] The STAR system categorizes drug candidates into four classes, where class I/III candidates may be able to balance clinical dose/efficacy/toxicity for better clinical success. In contrast, class II/IV candidates may cause imbalanced clinical dose/efficacy/toxicity, which leads to dose-limiting toxicity or lack of efficacy. Specifically, we intended to design STAR class III JAK inhibitors to achieve high local exposure/selectivity in the GI tissues while minimizing the requirement for JAK isoform specificity since none of the isoform-specific JAK inhibitors could eliminate their systemic toxicity.

We first designed MMT3-72 with poor absorption potential, as modeled to have a low drug-likeness score (low QED score), to ensure the drug will be retained in the GI tract.[4] MMT3-72 was then studied to release the potential active compounds of JAK inhibitors in local GI tissues for potential efficacy in treating UC, but reduce their exposure in systemic circulation to minimize potential toxicity. Two machine learning models were used to classify inhibitors/non-inhibitors and estimate IC_{50} values of the predicted inhibitors. Based on the prediction from two machine learning models, we only needed to synthesize MMT3-72 and MMT3-72-M2 for *in vitro* and *in vivo* testing on drug's potency/specificity, tissue selectivity, and efficacy in colitis models. The integration of ML, the STAR system, and wet lab synthesis/testing could minimize the effort in the optimization of a JAK inhibitor to treat colitis. This site-specific inhibitor reduces systemic toxicity by minimizing the need for JAK isoform specificity. The overall workflow is summarized in Figure 3.1.
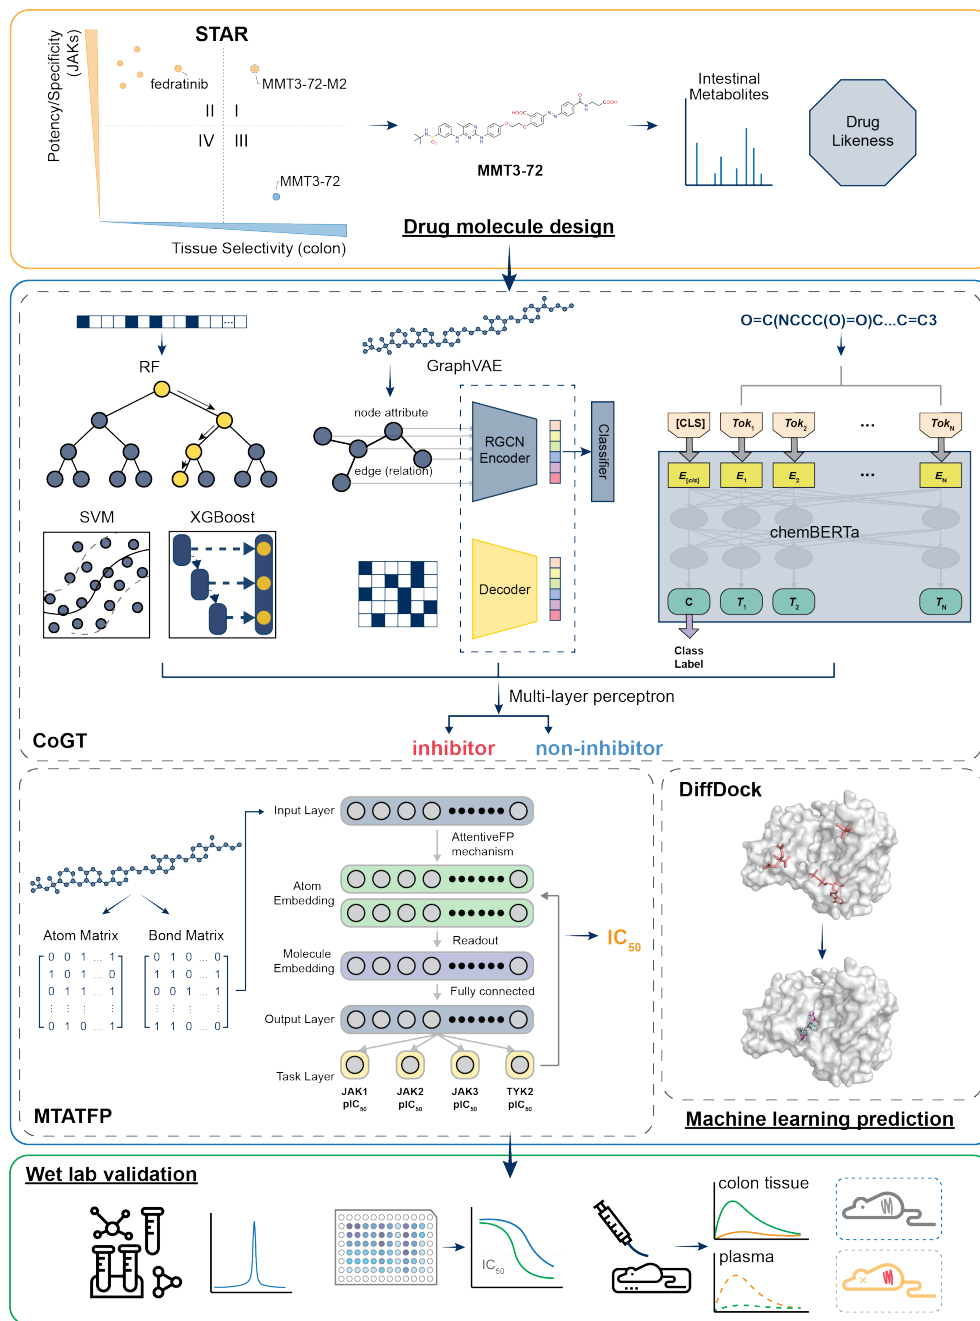
Figure 3.1: Overview for the machine learning-guided development of GI-locally activating JAK inhibitor for treatment of ulcerative colitis.

3.2 Materials and methods

General

All commercially available products and solvents were purchased from Sigma-Aldrich, AK Scientific, and Fisher Scientific. Solvents were used as received or dried over molecular sieves (4 Å). All water or air-sensitive reactions were performed under an argon atmosphere with dry solvents and anhydrous conditions. All reactions were monitored by thin-layer chromatography (TLC) that was performed on aluminum-backed silica plates (0.2 mm, 60 F254). Purification by flash chromatography was performed on Merck silica gel 60 (230 - 400 mesh). Yields refer to chromatographically and spectroscopically (^{1}H NMR) homogeneous materials unless otherwise stated.

NMR spectra were recorded on a Bruker instrument (500 or 300 MHz) and calibrated using a solvent peak as an internal reference. Spectra were processed using MestReNova software. Chemical shifts δ are given in ppm and coupling constants (J) in Hz. Peak multiplicities are described as follows: s, singlet, t, triplet, and m, multiplet. High-resolution mass spectra were obtained on an AB Sciex X500R QTOF spectrometer or an AB Sciex 6600+ Triple TOF mass spectrometer. The purity of all compounds subjected to biological tests was determined by analytical HPLC and was found to be $\geq$ 95 %.

Synthesis

The synthesis and characterization of compounds 1 and 4 have been previously described.[122, 79]

1-(2-chloroethoxy)-4-nitrobenzene (1): To a mixture of 4-nitrophenol (4 g, 28.754 mmol) and 1,2-dichloromethane (20 mL, 5 vol) in DMF (25 mL) was added K_2CO_3 (6 g, 43.131 mmol, 1.5 equiv), and the resulting mixture was stirred at 100 °C for 6h and monitored by TLC. Upon completion, the reaction mixture was quenched

with water and the product was extracted three times with CH_2Cl_2. The combined organic phase was washed with water, brine, dried over Na_2SO_4, and concentrated under vacuum to give compound 1 (5.35 g; yield, 92%). This intermediate was taken forward to the next step without further purification.

4-(2-chloroethoxy)aniline (2): To a mixture of compound 1 (1 g, 4.96 mmol) in EtOH (30 mL) was added $SnCl_2 \cdot 2H_2O$ (4.5 g, 19.84 mmol, 4 equiv) and the reaction mixture was stirred at 90 °C for overnight. Upon completion of the reaction, the solvent was evaporated under reduced pressure and the residue was taken into 5% aqueous NaOH and extracted three times with CH_2Cl_2. The combined organic phase was washed with 5% aqueous NaOH, water, brine and dried over Na_2SO_4. The solvent was concentrated under vacuum and the residue was purified by silica gel column chromatography to provide the compound 2 (532.4 mg, 63% yield). ^{1}H NMR (500 MHz, $CDCl_3$) δ 6.82 – 6.69 (m, 2H), 6.69 – 6.58 (m, 2H), 4.16 (t, $J = 5.9$ Hz, 2H), 3.77 (t, $J = 5.9$ Hz, 2H). HRMS (ESI): mass calcd. for $C_8H_{10}ClNO$, 171.05; m/z found, 172.0425 $[M + H]^+$.

N-(tert-butyl)-3-((2-((4-(2-chloroethoxy)phenyl)amino)-5-methylpyrimidin-4-yl)amino)benzenesulfonamide (3): To a mixture of compound 2 (400 mg, 2.330 mmol, 2 equiv) and N-tert-butyl-3-[(2-chloro-5-methylpyrimidin-4-yl)amino] benzenesulfonamide (413.53 mg, 1.165 mmol) in isopropanol (8 mL) was added 3 drops of concentrated HCl 37% and the reaction mixture was stirred at 80 °C for overnight. Upon completion of the reaction, the solvent was evaporated under reduced pressure and the residue was taken into aqueous $NaHCO_3$ and extracted three times with CH_2Cl_2. The combined organic phase was washed with water, brine and dried over Na_2SO_4. The solvent was concentrated under vacuum and the obtained solid was washed three times with EtOAc to provide the compound 3 (482 mg, 84% yield). ^{1}H NMR (500 MHz, DMSO-d_6) δ 8.82 (s, 1H), 8.55 (s, 1H), 8.12 (d, $J = 5.6$ Hz, 2H), 7.91 (d, $J = 1.0$ Hz, 1H), 7.56 (d, $J = 4.3$ Hz, 2H), 7.51 – 7.45 (m, 2H),

6.87 – 6.80 (m, 2H), 4.18 (d, $J = 6.0$ Hz, 2H), 3.92 (d, $J = 5.9$ Hz, 2H), 2.16 – 2.06 (m, 3H), 1.12 (s, 9H). HRMS (ESI): mass calcd. for $C_{23}H_{28}ClN_5O_3S$, 489.16; m/z found, 490.1506 $[M + H]^+$.

2-(4-nitrophenoxy)ethan-1-ol (4): To a mixture of 4-nitrophenol (3 g, 21.56 mmol) and 2-chloride ethanol (2.89 mL, 43.16 mmol, 2 equiv) in H_2O (10 mL) was added NaOH (1.73 g, 43.16 mmol, 2 equiv) and the reaction mixture was stirred at 80 °C for overnight. Upon completion of the reaction, the reaction mixture was cool down to room temperature, diluted with H_2O, and extracted three times with EtOAc. The combined organic phase was washed with water, brine and dried over Na_2SO_4. The solvent was concentrated under vacuum to give compound 4 (3.1 g, 79% yield), which was taken forward to the next step without further purification.

2-(4-aminophenoxy)ethan-1-ol (5): To a mixture of compound 4 (1 g, 5.460 mmol) in MeOH (20 mL) was added Pd/C (0.1 g, 10% equiv) and the reaction mixture was stirred at 50 °C for overnight under H_2 atmosphere. Upon completion of the reaction, the Pd/C was filtered off on celite and the solvent was evaporated under reduced pressure. The residue was purified by silica gel column chromatography to provide compound 5 (635.2 mg, 76% yield). ^{1}H NMR (300 MHz, CDCl$_3$) δ 6.82 – 6.72 (m, 2H), 6.70 – 6.59 (m, 2H), 4.01 (dd, $J = 5.1, 3.5$ Hz, 2H), 3.92 (dd, $J = 5.1, 3.5$ Hz, 2H). HRMS (ESI): mass calcd. for $C_8H_{11}NO_2$, 153.18; m/z found, 154.0770 $[M + H]^+$.

(E)-2-(2-(4-((4-((3-(N-(tert-butyl)sulfamoyl)phenyl)amino)-5-methylpyrimidin-2-yl)amino)phenoxy)ethoxy)-5-((4-((2-carboxyethyl)carbamoyl)phenyl)diazenyl)benzoic acid (MMT3-72): To a mixture of compound 3 (45 mg, 0.09 mmol, 1.5 equiv) and balsalazide disodium salt dehydrate (27.29 mg, 0.068 mmol, 1 equiv) in DMF (2 mL) was added K_2CO_3 (37.6 mg, 0.272 mmol, 4 equiv), and the resulting mixture was stirred at 100 °C for overnight and monitored by TLC. Upon completion, the solvent was evaporated under reduced pressure. The

residue was taken into H_2O and the solution was acidified with H_3PO_4 until pH 2-3. The precipitate was filtered and re-crystallization in CH_2Cl_2 provided the desired compound MMT3-72 (50.4 mg, 68%). ^1H NMR (500 MHz, DMSO-d_6) δ 8.68 (s, 1H), 8.29 (s, 1H), 8.08 (s, 2H), 7.96 (d, J = 8.3 Hz, 2H), 7.88 (d, J = 10.5 Hz, 2H), 7.82 (d, J = 8.3 Hz, 2H), 7.57 (s, 1H), 7.49 (t, J = 14.4 Hz, 4H), 6.86 (d, J = 9.7 Hz, 1H), 6.80 (d, J = 8.5 Hz, 2H), 4.36 (t, J = 4.5 Hz, 2H), 4.13 (t, J = 4.8 Hz, 2H), 3.61 – 3.48 (m, 2H), 2.66 (t, J = 6.9 Hz, 2H), 2.11 (s, 3H), 1.11 (s, 9H). HRMS (ESI): mass calcd. for $C_{40}H_{42}N_8O_9S$, 810.2; m/z found, 811.2758 $[M + H]^+$.

N-(tert-butyl)-3-((2-((4-(2-hydroxyethoxy)phenyl)amino)-5-methylpyrimidin-4-yl)amino)benzenesulfonamide (MMT3-72-M2): To a mixture of compound 5 (130 mg, 0.846 mmol, 3 equiv) and N-tert-butyl-3-[(2-chloro-5-methylpyrimidin-4-yl)amino]benzenesulfonamide (100 mg, 0.282 mmol) in isopropanol (2 mL) was added 3 drops of concentrated HCl 37% and the reaction mixture was stirred at 80 °C for overnight. Upon completion of the reaction, the solvent was evaporated under reduced pressure and the residue was taken into aqueous $NaHCO_3$ and extracted three times with CH_2Cl_2. The combined organic phase was washed with water, brine and dried over Na_2SO_4. The solvent was concentrated under vacuum and the obtained solid was washed three times with EtOAc to provide the compound MMT3-72-M2 (105 mg, 79% yield). ^1H NMR (300 MHz, DMSO-d_6) δ 8.79 (s, 1H), 8.55 (s, 1H), 8.12 (d, J = 3.4 Hz, 2H), 7.90 (s, 1H), 7.62 – 7.39 (m, 4H), 6.80 (d, J = 8.9 Hz, 2H), 3.92 (t, J = 5.1 Hz, 2H), 3.69 (t, J = 5.0 Hz, 2H), 2.12 (s, 3H), 1.12 (s, 9H). HRMS (ESI): mass calcd. for $C_{23}H_{29}N_5O_4S$, 471.58; m/z found, 472.1847 $[M + H]^+$.

Activation of MMT3-72 and metabolites identification

In-vivo metabolites identification was conducted using mouse plasma, colon, and feces samples that were collected at 6 h after oral administration of MMT3-72 (10

mg/kg). Liquid chromatography-tandem mass spectrometry was employed to separate and identify the possible metabolites. The LC–MS/MS method consisted of a Shimadzu LC-20AD HPLC system (Kyoto, Japan). Chromatographic separation of MMT3-72 and its metabolites was achieved using a Waters XBridage reverse phase C18 column (15 cm × 2.1 mm I.D., packed with 3.5 μm). A high-resolution AB Sciex X500R QTOF mass spectrometer (AB Sciex, Framingham, USA) in the positive-ion Information Dependent Acquisition (IDA) mode was used confirmation of accurate molecular weight. The mass range was recorded from m/z 100 to 1000 Da. The collision energy was set to 50 V for TOF MSMS. Data were collected with the software SCIEX OS and then processed with the software MetabolitePilot 2.0 (AB Sciex, Framingham, USA).

Machine learning models for JAK inhibition prediction

CoGT. We applied our previously trained ensemble model CoGT (Conventional ML models + Graph-based models + Transformer-based models) for JAK inhibition prediction of the compounds in this study.[10] CoGT is specifically trained for the JAK inhibitor prediction task. Briefly, a dataset $\mathcal{D} = \{(\mathbf{x}_i, y_i)\}_{i=1}^{N}$ was collected from ChEMBL [38, 39], BindingDB [41], PubChem [133, 63] and Liu et al [76], consisting of N molecule-label pairs and was assumed to be i.i.d. (Independently and identically distributed random variables). Tanimoto similarity analysis for the data set used is visualized in Supplementary Figure S3.2A. $\mathbf{x}_i$ is the feature of one molecule, and y_i is its true label (1 if the molecule is JAK inhibitor, 0 otherwise, 10 μM is the threshold for IC$_{50}$). For the feature $\mathbf{x}_i$, it could either be MACCS fingerprint, graph, or simplified molecular input line entry systems (SMILES) representation based on different ML models. The aim is to do binary classification: given an unknown drug compound $\mathbf{x}$, the model outputs the probability p of the compound being a JAK inhibitor (if p > 0.5, the molecule is predicted by CoGT as a JAK inhibitor with

$IC_{50} < 10\mu M$, else it is predicted to be a non-inhibitor).

For model architecture, CoGT is an ensemble model including graph-bsed model (GraphVAE), transformer-based model (chemBERTa) and conventional ML models Support Vector Machine (SVM), Random Forest (RF) and XGBoost. By integrating different ML methods with different input representation of compounds using multilayer perceptron (MLP), CoGT outperformed individual ML models and leveraged the predicting ability for JAK inhibition predicition.

To train CoGT, the dataset was first randomly split into training set, validation set, and test set with ratio 8:1:1. Training data were used to train the model, validation data was used for fine tuning and test set was used for model evaluation of the generalization ability. BCELoss was used as loss function for this binary classification task to calculate the difference between the model prediction and the compound labels (0 if noninhibitor; 1 if inhibitor for specific JAK). Training was terminated if loss did not decrease and model with lowest loss was saved as the optimal model.

Code is available at https://github.com/yingzibu/JAK_ML.

MTATFP. Wang et al. built a multitask regression model based on the attentive fingerprint framework (MTATFP) to predict compounds' IC_{50} for JAKs.[134] Briefly, data were collected from PubChem, ChEMBL, and Binding DB. Tanimoto similarity analysis for the data set used is visualized in Supplementary Figure S3.2B and the data distribution are summarized in Supplementary Figure S3.3. The model was trained as a regression task, aiming to predict compounds' pIC_{50} values based on the molecular graph constructed by atom matrix and bond matrix. The Attentive FP framework was used to learn atomic and molecular properties through the graph-attention network (GAT) mechanism. [146] MTATFP used a multitask learning strategy for four JAKs to enhance the performance by data augmentation through eavesdropping between tasks, allowing simultaneous prediction of compounds' IC50 values for four JAK isoforms. For model architecture, atomic features are extracted

by Attentive FP convolutional layers and the molecular embedding was obtained using a readout layer. Transformation of molecular embedding is performed using linear layers with LeakyReLU. To elaborate the training process, collected data set was randomly split into training set, validation set and test set with a ratio of 8:1:1. MSELoss was used as loss function to calculate the difference between the model prediction and the true pIC_{50} of compounds. If JAK pIC_{50} values were missing, masking strategy was used to ignore the loss computation for the missing data entry. Early stopping approach was applied to avoid overfitting. Code is available at https://github.com/Yimeng-Wang/JAK-MTATFP.

Molecular docking

DiffDock. We employed DiffDock for its speed and accuracy to accelerate this process.[22] Briefly, instead of considering ligand pose $\mathbf{x} \in \mathbb{R}^{3n}$ in which n is the atom number of the ligand like EquiBind, DiffDock constrains the ligand pose in an (m+6)-dimensional submanifold $\mathcal{M}_c \subset \mathbb{R}^{3n}$, in which m represents the number of rotatable bonds, and 6 comes from the rotation and translation of the ligand relative to the given protein position. DiffDock utilizes a generative diffusion model to gradually update the random-initialized molecules' torsion, rotation, and translation to fit into the fixed protein.

To elaborate on how to define the (m+6)-dimensional submanifold $\mathcal{M}_c \subset \mathbb{R}^{3n}$, for a seed conformation $\mathbf{c} \in \mathbb{R}^{3n}$, let $\mathbf{x}_i \in \mathbb{R}^3$ be the position of the ith atom, $\bar{\mathbf{x}} = \frac{1}{n}\sum \mathbf{x}_i$ be the mass center of the ligand, we have

$$\text{translation } A_{tr}(\mathbf{r}, \mathbf{x})_i = \mathbf{x}_i + \mathbf{r}, \qquad A_{tr} : \mathbb{T}(3) \times \mathbb{R}^{3n} \to \mathbb{R}^{3n} \qquad (3.1)$$

$$\text{rotation } A_{rot}(R, \mathbf{x})_i = R(\mathbf{x}_i - \bar{\mathbf{x}}) + \bar{\mathbf{x}}, \qquad A_{rot} : SO(3) \times \mathbb{R}^{3n} \to \mathbb{R}^{3n} \qquad (3.2)$$

Define $\text{RMSDAlign}(\mathbf{x}, \mathbf{x}') = \arg\min_{\mathbf{x}^\dagger \in \{g\mathbf{x}'|g \in SE(3)\}} \text{RMSD}(\mathbf{x}, \mathbf{x}^\dagger)$ and let $B_{k,\theta_k}(\mathbf{x})$

be the kth rotatable bond torsion update by θ_k and all updates for m rotatable bonds are $\boldsymbol{\theta} = (\theta_1, \cdots, \theta_m)$. To make sure the infinitesimal effect of a torsion is orthogonal to rototranslation, the torsion is defined as

$$A_{tor}(\boldsymbol{\theta}, \mathbf{x}) = \mathrm{RMSDAlign}(\mathbf{x}, (B_{1,\theta_1} \circ \cdots B_{m,\theta_m})(\mathbf{x})), A_{tor} : \mathrm{SO}(2)^m \times \mathbb{R}^{3n} \to \mathbb{R}^{3n}$$

(3.3)

Let the product space $\mathbb{P} = \mathbb{T}^3 \times \mathrm{SO}(3) \times \mathrm{SO}(2)^m$, and define the overall pose update as

$$A((\mathbf{r}, R, \boldsymbol{\theta}), \mathbf{x}) = A_{tr}(\mathbf{r}, A_{rot}(R, A_{tor}(\boldsymbol{\theta}, \mathbf{x}))), \quad A : \mathbb{P} \times \mathbb{R}^{3n} \to \mathbb{R}^{3n} \qquad (3.4)$$

Thus the space of ligand poses $\mathcal{M}_{\mathbf{c}} = \{A(g, \mathbf{c}) | g \in \mathbb{P}\}$ is formally defined.

To apply a diffusion model on the product manifold $\mathbb{P}$, it is sufficient enough to sample from the diffusion kernel and compute the score independently in each manifold and the tangent space $T_g\mathbb{P} = T_{\mathbf{r}}\mathbb{T}_3 \oplus T_R\mathrm{SO}(3) \oplus T_{\boldsymbol{\theta}}\mathrm{SO}(2)^m \cong \mathbb{R}^3 \oplus \mathbb{R}^3 \oplus \mathbb{R}^m$ is a direct sum, where $g = (\mathbf{r}, R, \boldsymbol{\theta})$. For ligand pose $\mathbf{x}$ and protein structure $\mathbf{y}$, the authors construct the score model $\mathbf{s}(\mathbf{x}, \mathbf{y}, t) \to T_{\mathbf{r}}\mathbb{T}_3 \oplus T_R\mathrm{SO}(3) \oplus T_{\boldsymbol{\theta}}\mathrm{SO}(2)^m \cong \mathbb{R}^3 \oplus \mathbb{R}^3 \oplus \mathbb{R}^m$ to output SE(3)-equivariant translation vectors, rotation vectors, and SE(3)-invariant torsion angles. An SE(3)-invariant confidence model $\mathbf{d}(\mathbf{x}, \mathbf{y})$ assigns scores for different ligand poses, and the top-ranked pose is then taken as DiffDock's top-1 prediction. We set the diffusion step to 20 and evaluated the docking poses of MMT3-72, M2, and fedratinib. The crystal structures of JAKs were retrieved from the Protein Data Bank (https://www.rcsb.org/): JAK1: 6BBU, JAK2: 6VNE, JAK3: 4HVD, and TYK2: 6DBK. The code is available at https://github.com/gcorso/DiffDock.

EquiBind. EquiBind was built to perform direct-shot prediction.[119] In detail, both ligands and proteins were represented as spatial k-nearest neighbor (k-NN) graphs $\mathcal{G} = (\mathcal{V}, \mathcal{E})$. $\mathcal{V}$ is the atom/residue information for ligand and protein, respectively, consisting of node 3D coordinates $\mathbf{X}$ and initial features $\mathbf{F}$ for each node. For edges $\mathcal{E}$, ligand edges include atom pairs < 4 Å and the protein edges consist of

connection of the closest 10 other nodes < 30 Å.

It is desired to predict the same binding complex regardless of the initial position and orientation of ligand/protein in 3D space, thus EquiBind incorporated Independent E(3)-Equivariant Graph Matching Network (IEGMN) to make sure the coordinate transformations are E(3)-equivariant while the feature embeddings are E(3)-invariant.[36] In other words, the output of the IEGMN is

$$\text{IEGMN}(\mathbf{X}, \mathbf{F}, \mathbf{X}', \mathbf{F}') = \mathbf{Z}, \mathbf{H}, \mathbf{Z}', \mathbf{H}' \tag{3.5}$$

in which $\mathbf{Z}, \mathbf{Z}'$ represent transformed coordinates and $\mathbf{H}, \mathbf{H}'$ are feature embeddings of molecules and targets, respectively. If applied a transformation in 3D space for both the ligand and/or target, the initial location of ligand now is $\mathbf{UX} + \mathbf{b}$ and the protein location is $\mathbf{U}'\mathbf{X}' + \mathbf{b}'$, in which rotation matrices $\mathbf{U}, \mathbf{U}' \in \mathbf{SO}(3)$ and translation vectors $\mathbf{b}, \mathbf{b}' \in \mathbb{R}^3$, IEGMN guarantees that

$$\text{IEGMN}(\mathbf{UX} + \mathbf{b}, \mathbf{F}, \mathbf{U}'\mathbf{X}' + \mathbf{b}', \mathbf{F}') = \mathbf{UZ} + \mathbf{b}, \mathbf{H}, \mathbf{U}'\mathbf{Z}' + \mathbf{b}', \mathbf{H}' \tag{3.6}$$

For updating transformed coordinates during training IEGMN, molecules were assumed to be flexible only with rotatable bonds, while bond lengths, adjacent bond angles and small rings were treated as rigid, which could be defined as minimizing the following function defined by the authors:

$$
\begin{aligned}
\mathcal{S}(\mathbf{Z}, \mathbf{X}) = &\sum_{\{(i,j) \in \mathcal{E}\}} (d_{\mathbf{X}}^2(i,j) - d_{\mathbf{Z}}^2(i,j))^2 \\
&+ \sum_{\{(i,j):\text{2-hops away in } \mathcal{G}\}} (d_{\mathbf{X}}^2(i,j) - d_{\mathbf{Z}}^2(i,j))^2 \\
&+ \sum_{\{(i,j):i \text{ in aromatic ring with } j\}} (d_{\mathbf{X}}^2(i,j) - d_{\mathbf{Z}}^2(i,j))^2
\end{aligned}
\tag{3.7}
$$

This constraints could be satisfied by incorporating T gradient descent layers and

update $\Psi_t(\mathbf{Z}) \leftarrow \mathbf{Z} - \eta \nabla_{\mathbf{Z}} \mathcal{S}(\mathbf{Z}, \mathbf{X}), \forall t \in \{1, \cdots, T\}$.

However, implausible conformers may be produced during this process, thus the authors came up with the idea of only changing the torsion angles of the conformer initialized by RDKit. In detail, the output $\mathbf{C} \in \mathbb{R}^{3 \times n}$ is initialized as $\mathbf{X}$ and thus $\mathcal{S}(\mathbf{C}, \mathbf{X}) = 0$. Then only the rotatable bonds' torsion angles are changed to match those of $\mathbf{Z}$ to avoid implausible bond angles/lengths. Instead of computing the gradient of a point cloud w.r.t. torsion angles which is computationally expensive, the authors compute the rotatable bonds' dihedral angles of $\mathbf{C}$ as maximum likelihood estimates of von Mises distributions on those of $\mathbf{Z}$'s without requiring optimization. I.e. to minimize the torsion angle difference between $\mathbf{Z}$ and $\mathbf{C}$, a closed form solution exists which could provide all dihedral angle values.

By using an SE(3)-equivariant multi-head attention mechanism, ligand and receptor keypoints $\mathbf{Y}, \mathbf{Y}' \in \mathbb{R}^{3 \times K}$ were trained to match the ground truth binding pocket points using the coordinate outputs $\mathbf{Z}, \mathbf{Z}'$ computed by IEGMN above. Once trained, performing SE(3) transformation to superimpose $\mathbf{Y}$ and $\mathbf{Y}'$ would be equivalent to performing ligand docking. Code is available at

https://github.com/HannesStark/EquiBind.

In vitro activity to inhibit JAK enzymes

JAK1, JAK2, JAK3, and TYK2 assay kits were obtained from BPS Bioscience (San Diego, CA, USA). The assays were conducted according to the manufacturer's protocols in 96-well microplates. Briefly, master mixtures (25 μL per well) were prepared for JAK1 and TYK2 assays (6 μL 5$\times$ kinase assay buffer + 1 μL ATP (500 μM) + 5 μL 10$\times$ IRS1-tide + 13 μL distilled water) or for JAK2 and JAK3 assays (6 μL 5$\times$ kinase assay buffer + 1 μL ATP (500 μM) + 1 μL PTK substrate Poly(Glu:Tyr 4:1) (10 mg/mL) + 17 μL distilled water), respectively. Then 5 μL of fedratinib, MMT3-72, MMT3-72-M2 solutions at different concentrations were added

to the above prepared master mixtures, which were followed by 20 μL of enzymes (JAK1 at 5 ng/μL, JAK2 at 2.5 ng/μL, JAK3 at 0.4 ng/μL, or TYK2 at 0.5 ng/μL), respectively. The reaction mixtures were incubated at 30°C for 40 minutes. Finally, 50 μL of Kinase-Glo Max reagent (Promega, Madison, WI, USA) were added to each well and the reactions were performed in darkness for 15 minutes at room temperature. The luminescence of the reaction mixture was read on a Synergy 2 microplate reader (Biotek).

LC-MS analysis of MMT3-72 and MMT3-72-M2 in biological samples

MMT3-72 and MMT3-72-M2 concentrations in plasma (ng/mL) and tissues (ng/g) were determined by the LC–MS/MS method that was developed and validated for this study. The HPLC method was conducted on a Shimadzu LC-20AD HPLC system (Kyoto, Japan), and chromatographic separation was achieved using a Waters XBridage reverse phase C18 column (5 cm $\times$ 2.1 mm I.D., packed with 3.5 μm). The flow rate of gradient elution was 0.4 ml/min with mobile phase A (0.1% formic acid in purified deionized water) and mobile phase B (0.1% formic acid in acetonitrile). An AB Sciex QTrap 4500 mass spectrometer (AB Sciex, Framingham, USA) in the positive-ion multiple reaction monitoring (MRM) mode was used for detection. Protonated molecular ions and the respective ion products were monitored at the transitions of m/z 811.3 > 737.4 for MMT3-72 and 472.3 > 416.0 for MMT3-72-M2. Data were processed with the software Analyst (version 1.6).

***In vitro* conversion of MMT3-72 to MMT3-72-M2 by active colon contents**

C57BL/6 female mice (aged 6-8 weeks) were purchased from Charles River Laboratories. Fresh colon contents were isolated from the mice and suspended in PBS. The resulting 10% (w/v) colon content suspension was used as the source of azoreductases. A stock solution of MMT3-72 was added to the suspension to achieve a

final concentration of 0.1 mM. The mixture was thoroughly combined and incubated on a shaker at 1000 rpm and 37 ℃. At 0, 1, 2, 4, 6, 24, and 48 h, a sample aliquot was collected, followed by 10-fold dilution and extraction using acetonitrile. Additionally, a parallel experiment was conducted using deactivated colon contents (boiled). The extracted supernatants were stored at -20 ℃ until quantitative analysis. The amounts of MMT3-72 and MMT3-72-M2 were determined via liquid chromatography-tandem mass spectrometry (LC-MS/MS) analysis using the method described above.

Animals

Female C57BL/6 mice were obtained from Charles River Laboratories at 4 – 6 weeks of age. Mice were housed in groups of five animals per cage on a 12:12 h light/dark cycle. Water and food were provided *ad libitum.* Mice aged 6 to 8 weeks were used for experiments. All procedures were approved by the Institutional Animal Care and Use Committee of the University of Michigan.

Pharmacokinetic study in mice

Briefly, C57BL/6 female mice (6-8 weeks old) were orally dosed with 10 mg/kg MMT3-72. At 0.5, 2, 4, 12, and 24 h, mice were sacrificed, and blood samples were collected by drawing directly from heart. Intestinal tissue samples were collected and homogenized in PBS to 10% homogenate. The contents of small and large intestines were collected and homogenized in PBS. Afterward, concentrations of MMT3-72, MMT3-72-M2 in the plasma, colon tissues, and colon content were determined by the quantitation method described above.

In vivo efficacy of MMT3-72 in the treatment of DSS-induced colitis in mice

C57BL/6 female mice (6-8 weeks old) were randomly divided into different treatment groups. Acute colitis was induced by administering 3% DSS (MP Biomedicals, CA, USA) in distilled water continuously for 5 days, and the control group received pure water.[116] MMT3-72 or tofacitinib was dissolved in a pH-adjusted solution comprising an 80% v/v mixture of beta-cyclodextrin and water (at a ratio of 1:4), 15% hyaluronic acid (HA), and 5% phosphate-buffered saline (PBS). Drugs were administered every other day orally by gavage in a volume of 0.1mL/10g body weight. During the model establishment, body weight, stool consistency, and gross blood in feces were monitored and recorded daily.[116] After 5 days, mice were sacrificed and blood was collected. Serum was obtained by centrifugation and stored at -80 ℃ until further assay. The colon was excised and the length was measured.

H&E staining of colon tissues

After dissecting and transecting the colon, a feeding needle and 5 ml syringe were used to incubate and flush the colon with ice-cold PBS until the stool was flushed out. Scissors were used to incise longitudinally from the distal to the proximal end of the colon and the colon tissue could then be expanded as a flat sheet. The edge of the distal colon was grasped using a pair of forceps and the colon tissue was rotated into a Swiss roll. The roll was firmly grasped and transected using a 27G 1/2 needle. Then the sample was placed in 4% Paraformaldehyde Fix Solution (Thermo Scientific, USA) at room temperature for 24 hours. The Swiss roll was then paraffin-embedded, sectioned, mounted, and stained with H&E to determine the extent of damage to the colon from the distal (inside end) to the proximal (outside end).

Western blot assay of colon tissues

The dissected colon tissue was rinsed in cold PBS and homogenized in RIPA buffer. The homogenates were subjected to ultrasonication and centrifugation for protein extraction. The resulting supernatants were collected, and the total protein was quantified using the BCA kit. For each treatment group, 200 μg of total protein was used for SDS-electrophoresis. The separated proteins were then transferred onto PVDF membranes. These membranes were incubated overnight at 4℃ with the indicated primary antibodies, followed by a one-hour incubation with horseradish peroxidase-labeled secondary antibodies at room temperature. The blots were visualized using the ChemiDoc Imaging System (Bio-Rad). Antibodies against JAK2, STAT3, Phospho-STAT3 (Tyr705), and GAPDH (Cell Signaling Technology) were used for immunoblotting.

Data analysis

Data are presented as mean $\pm$ SD. For statistical difference determination between treatment groups, one-way ANOVA or two-way ANOVA was performed with Tukey's post hoc correction (GraphPad Prism 9). Significance was accepted at $p < 0.05$. *, $p < 0.05$, **, $p < 0.01$, ***, $p < 0.005$, ****, $p \leq 0.001$; ns, not significant. IC_{50} values were derived from fits to Hill equations.

3.3 Results and discussion

Rational design of MMT3-72

We chose to design MMT3-72 based on the following three considerations: (A) The scaffold has intrinsic properties to be secreted into the GI tract. We used the scaffold of fedratinib because it could be preferentially secreted to the GI tract intact. For instance, a significant proportion of fedratinib-derived radioactivity ($\sim$77%) was ex-

creted in feces (23% unchanged) after a single oral dose of radiolabeled fedratinib.[90] This observation inspired us to further modify fedratinib to create a colon-selective compound for the treatment of colitis. (B) The designed compound needs to be retained in the GI tract without effective absorption to reduce systemic exposure for decreasing its toxicity. This is achieved by estimation of a low drug-likeness score (low QED score), which was modeled based on eight parameters, including molecular weight (MW), octanol-water partition coefficient (ALOGP), number of hydrogen bond donors (HBD), number of hydrogen bond acceptors (HBA), molecular polar surface area (PSA), number of rotatable bonds (ROTB), number of aromatic rings (AROM), and number of structural alerts (ALERTS).[4] (C) The designed compound could be activated by bacteria in the colon to release its active metabolite for achieving its site-specific efficacy.

Based on the X-ray structure of JAK2-bound fedratinib,[25] we replaced the solvent-exposed pyrrolidine moiety of fedratinib with 5-aminosalicylic acid (5-ASA) linked by an azo bond to N-4-(aminobenzoyl)-beta-alanine (Figure 3.2A, B). The modification is intended to increase the molecular weight and polarity of MMT3-72, which reduces its absorption potential from the GI tract and results in more accumulation in the colon tissues.

To evaluate its drug-like property (especially absorption potential), we estimated its drug-likeness (QED, 0-1) according to the published method.[4] The QED of MMT3-72 was predicted to be 0.0382, suggesting it has low or no absorption potential (the maximal value of QED is 1 suggesting best absorption potential) (Figure 3.2C). As a comparison, the QED of balsalazide is predicted to be 0.5598 (Table S3.1). When MMT3-72 reaches the colon region, the azo bond would be cleaved by bacterial azoreductases to release the active metabolite(s), which is then absorbed and accumulated in the colon tissues with minimal exposure to the systemic circulation.

Compound	MW	ALOGP	HBD	HBA	PSA	ROTB	AROM	ALERTS	QED
MMT3-72	810.89	7.12492	6	13	242.89	18	5	2	0.0382
M1	415.48	2.29072	4	8	139.46	8	3	0	0.4393
M2	471.58	3.72012	4	8	125.47	9	3	0	0.3721
M3	485.57	3.81242	4	9	142.54	9	3	0	0.3556
M4	754.78	5.69552	6	13	256.88	17	5	2	0.0493
M5	427.53	4.05462	4	7	116.24	6	3	1	0.4385
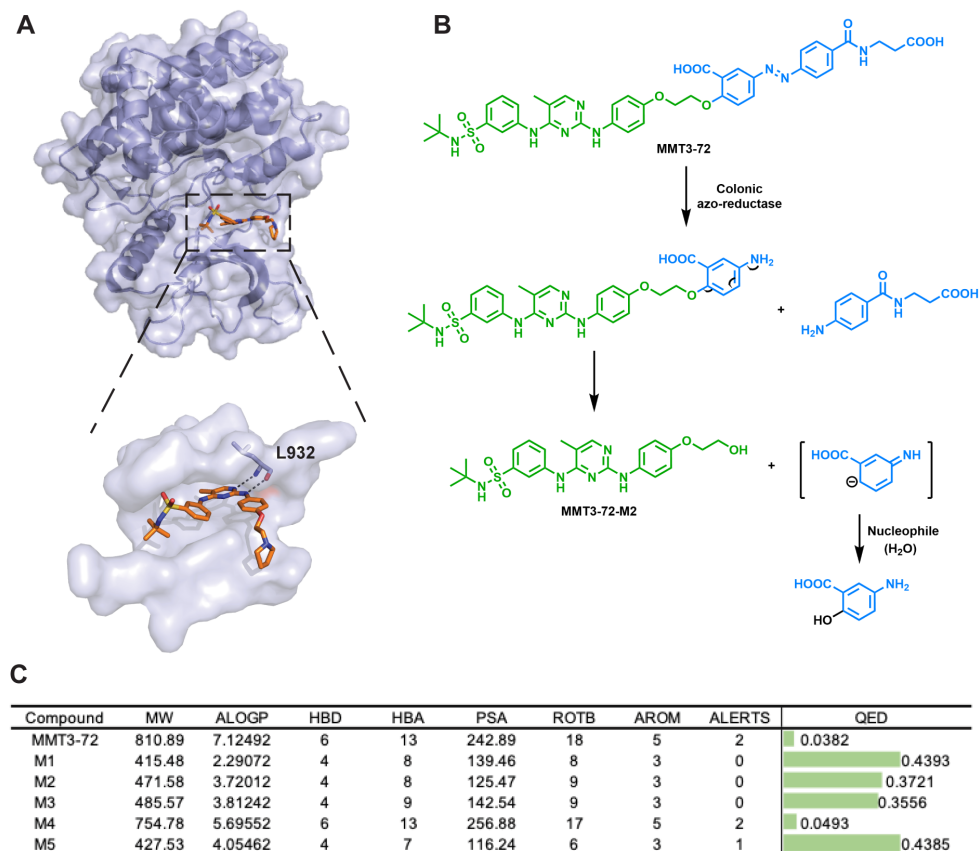
Figure 3.2: **Rational design of MMT3-72.** (A) Cocrystal structure of JAK2 with fedratinib (PDB code: 6VNE). (B) Rationale for MMT3-72. (C) Quantitative estimates of drug-likeness of MMT3-72 and its metabolites.

Chemistry: Synthesis of MMT3-72

The compound MMT3-72 was synthesized according to the synthetic route as shown in Figure 3.3A. First, 4-nitrophenol was condensed with DCE to produce intermediates 1. A subsequent nitro reduction of intermediates 1 with Tin chloride afforded the intermediates 2. Then, intermediate 2 was coupled with N-tert-butyl-3-[(2-chloro-5-methylpyrimidin-4-yl)amino]benzenesulfonamide in the presence of a few drops of HCl in isopropanol to generate compound 3. Finally, compound 3 was subject to a nucleophilic substitution reaction with balsalazide disodium salt dihydrate to yield the desired compound MMT3-72.

GI local activation of MMT3-72 and metabolites identification in GI contents, GI tissues, and plasma

To confirm the activation of our newly synthesized compound MMT3-72 in the colon, mice were dosed orally with 10 mg/kg MMT3-72 and sacrificed at 6 h to collect plasma, colon tissues, and colon contents (feces). Five metabolites (M1 to M5) were identified using LC-MS in the collected samples and their structures are shown in Fig. 3.3B. Interestingly, MMT3-72 was only detected in the feces with no detection in the plasma and the colon tissues. The major metabolite MMT3-72-M2 was only detected in the colon tissues and the feces with limited detection in the plasma.The minor metabolite M1 was detected in all three sampled tissues, with a majority in the feces. The other three minor metabolites M3, M4, and M5 were only identified in the feces with no detection in either the colon tissues or the plasma. Details are summarized in Fig. 3.3C.
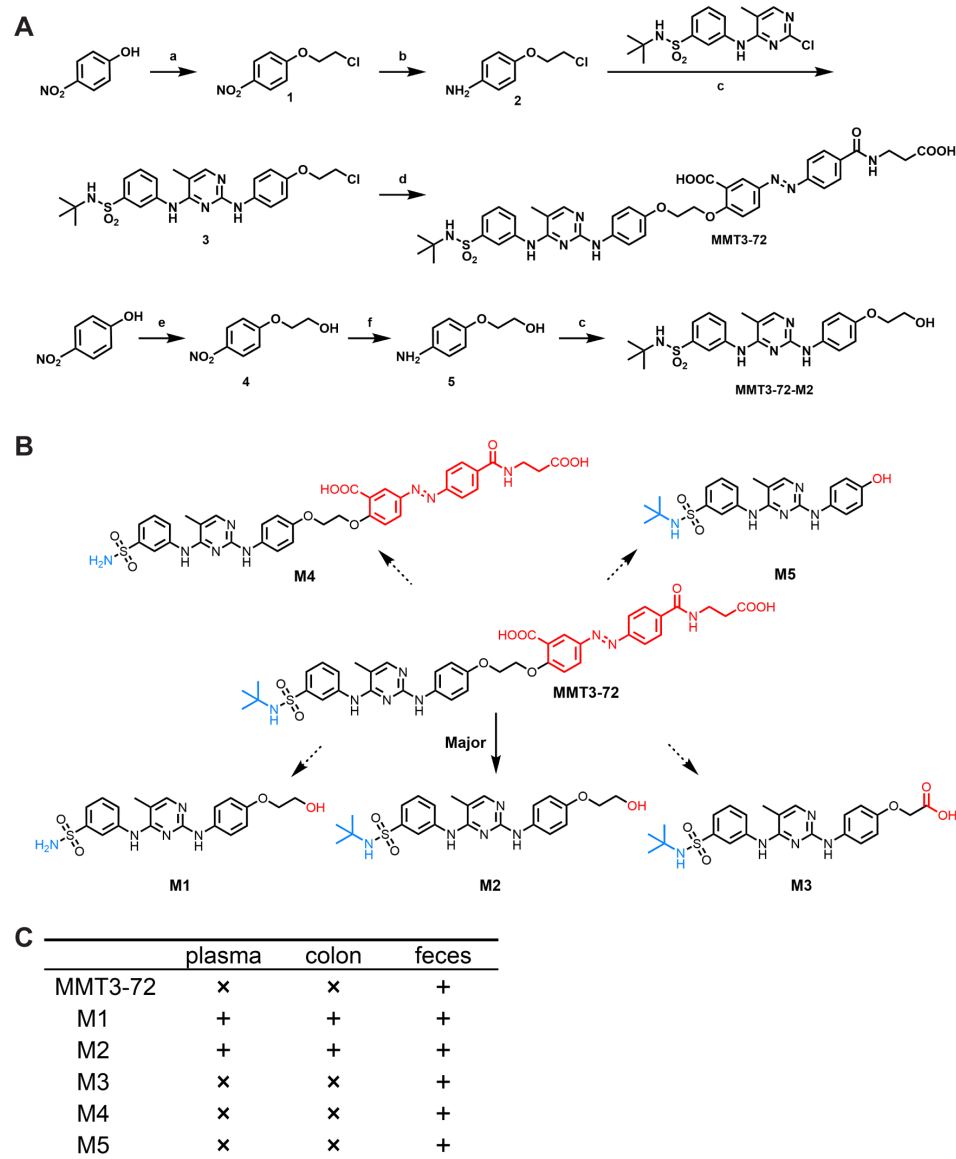
Figure 3.3: **Synthesis and metabolism study of MMT3-72**. (A) Synthesis of MMT3-72 and MMT3-72-M2. Reaction conditions: (a) DCE, K_2CO_3, DMF, 100°C, 6h, and 92%; (b) $SnCl_2 \cdot 2H_2O$, EtOH, rt, overnight, and 63%; (c) HCl conc, iPrOH, 80°C, and 79-84%; (d) balsalazide disodium, K_2CO_3, DMF, 100°C, 6h, and 68%; (e) 2-chloride ethanol, NaOH, H_2O, 8h, 80°C, and 79%. (f) H_2, Pd/C, MeOH, 50°C, overnight, and 76%. (B) GI locally activating of MMT3-72 and Metabolite identification *in vivo*. Structures of MMT3-72 and its 5 metabolites are shown as MMT3-72-M1, MMT3-72-M2, MMT3-72-M3, MMT3-72-M4, and MMT3-72-M5. (C) Concentrations of MMT3-72 and its metabolites *in vivo*. Mice were orally dosed with 10 mg/kg MMT3-72 and sacrificed at 6h. ×, not detected; +, detected.

54

	Predicted Kinase IC$_{50}$ (nM)			
	JAK1	JAK2	JAK3	TYK2
MMT3-72	67.8	-	-	-
M1	-	205.5	-	-
M2	19.1	34.2	-	65.5
M3	-	128.2	-	-
M4	-	-	-	-
M5	30.6	154.0	-	182.5
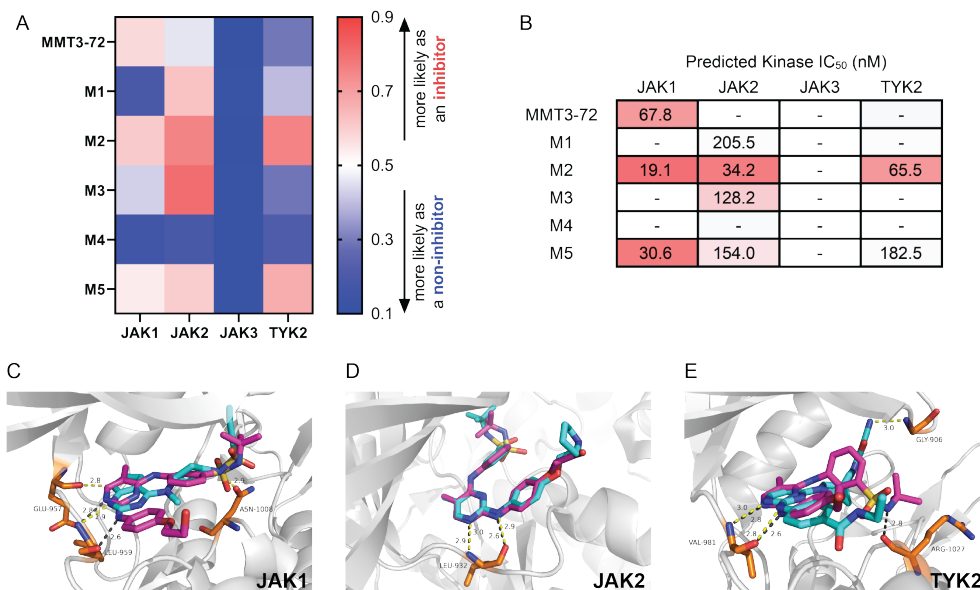
Figure 3.4: **Machine learning assisted analyses on the activities of MMT3-72 and its metabolites.** Prediction results on activities against JAKs of MMT3-72 and its metabolites (A) by CoGT and (B) by MTATFP. Docking simulation was performed using DiffDock. MMT3-72-M2 was docked in the binding sites in JAK1 (C), JAK2 (D), and TYK2 (E) as it is predicted to be inhibitors for these JAK isoforms. The key residues that formed hydrogen bonding contacts with the ligands are colored orange. The dashed lines represent the formed hydrogen bonding contacts between the amino acid residues and MMT3-72-M2 (dark grey) or ground-truth ligands (yellow). Magenta: MMT3-72-M2; cyan: ground-truth ligand (co-crystalized with the proteins). PDB code: JAK1: 6BBU; JAK2: 6VNE; TYK2: 6DBK. The ligand docking results were all visualized in PyMol.[110]

Machine learning (ML) models to predict IC$_{50}$ and binding of MMT3-72 and its five metabolites against JAK1, 2, 3 and TYK2

In order to know how these five metabolites may inhibit JAK1, 2, 3, and TYK2, we used two ML models to predict their activities against these isoforms. This will reduce the efforts to only synthesize the desired compounds for further testing. We first utilized CoGT model to predict whether the compound is an inhibitor or not, then used MTATFP model to predict their IC$_{50}$ values if the compound is predicted as an inhibitor by CoGT. Results are shown in Fig. 3.4A-B. Further details for CoGT and MTATFP are in Section 3.2. Both ML methods predicted that MMT3-72 is a

JAK1 inhibitor (with IC_{50} of 68 nM) but did not inhibit the other three isoforms. The main metabolite MMT3-72-M2 is predicted as a JAK1 (IC_{50} 19 nM), JAK2 (IC_{50} 34 nM), and TYK2 (IC_{50} 65 nM) inhibitor. This result affirmed our confidence to proceed with the synthesis and evaluation of MMT3-72-M2 for its activity. In addition, MMT3-72-M2 was also confirmed to have a high concentration in the colon tissue. The reason for not proceeding with the evaluation on M1 or M3-M5 is that they have low concentrations in the colon and did not show as good predicted inhibitory activities against different JAK isoforms.

To further understand how MMT3-72 and MMT3-72-M2 bind to different JAK isoforms, we also used ML-based docking models to predict molecular pose in JAK proteins. First, we used DiffDock, a diffusion generative model, to predict the binding pose of these two molecules during protein-ligand interaction.[22] DiffDock refines the ligand pose space as an (m + 6)-dimensional submanifold $\mathcal{M} \in \mathbb{R}^{3n}$, in which m and n are the rotatable bond number and the atom number of a ligand, respectively. Results of DiffDock prediction are shown in Fig. 3.4C-E. The docking simulations consistently revealed highly similar binding poses of MMT3-72-M2 (highlighted in magenta) compared to the co-crystalized ligand data (highlighted in cyan). Detailed analysis of the molecular interactions within the binding sites further confirmed the presence of similar hydrogen bonding contacts. Specifically, MMT3-72-M2 was able to form hydrogen bonding with the residue Leu959 in the binding site of JAK1 similar to its co-crystal structure with its ligand (compound 25[128])(Fig. 3.4C). In addition, MMT3-72-M2 exhibited a similar binding pose to that of fedratinib, and both compounds were capable of forming hydrogen bonding with the Leu932 residue (Fig. 3.4D). Further, MMT3-72-M2 was predicted to form hydrogen bonding with the residue Val981 as seen in the co-crystal structure with its ligand (Compound 8[35]) in the binding site of TYK2. In addition, MMT3-72-M2 also form additional interaction with the residue Arg1027 (Fig. 3.4E). These findings indicate the potential of

56

MMT3-72-M2 as an effective JAK inhibitor.

As a comparison, we also predict the binding of MMT3-72-M2 to JAKs using EquiBind.[119] Since EquiBind is a regression-based method to predict the ligand position during the process of minimizing the expected square error, the performance is inferior compared to that of DiffDock. Docking prediction results are shown in Supplementary Figure S3.1.

MMT3-72 was less active but MMT3-72-M2 was more potent against JAK1, 2, and TYK2 by *in vitro* kinase assays.

To confirm the correctness of the ML prediction of JAK inhibition, biological activities of MMT3-72 and its active metabolite MMT3-72-M2 were evaluated against JAK1, JAK2, JAK3, and TYK2 using kinase assays (Figure 3.5). MMT3-72 showed modest activities against JAK1 and JAK2 (367.7 nM and 630 nM, respectively) and poor activities against JAK3 and TYK2 (5237 nM and 4697 nM, respectively). On the other hand, the active metabolite MMT3-72-M2 showed potent inhibition of JAK1 (10.8 nM), JAK2 (26.3 nM), and TYK2 (91.6 nM), but weak inhibition of JAK3 (328.7 nM). These experimental results are in agreement with ML prediction.

MMT3-72 was readily converted to MMT3-72-M2 in the colon contents.

To quantify the activation process of the colon-activating MMT3-72, we investigated its release of MMT3-72-M2 in colon content suspensions. Figure 3.6A shows that fresh colon contents metabolized more than 90% of MMT3-72 within 48 hours of incubation. In contrast, most MMT3-72 remained intact in the aqueous environment when the colon contents had been heat-deactivated. Importantly, MMT3-72-M2 was readily converted from MMT3-72 in fresh colon contents (Figure 3.6B). Given that the IC_{50} of MMT3-72-M2 was measured to be 10.8 nM ($\sim$ 5.0 ng/mL) for JAK1, 26.3 nM ($\sim$ 12.1 ng/mL) for JAK2, 91.6 nm ($\sim$ 42.1 ng/mL) for TYK2 (Figure 3.5E), it
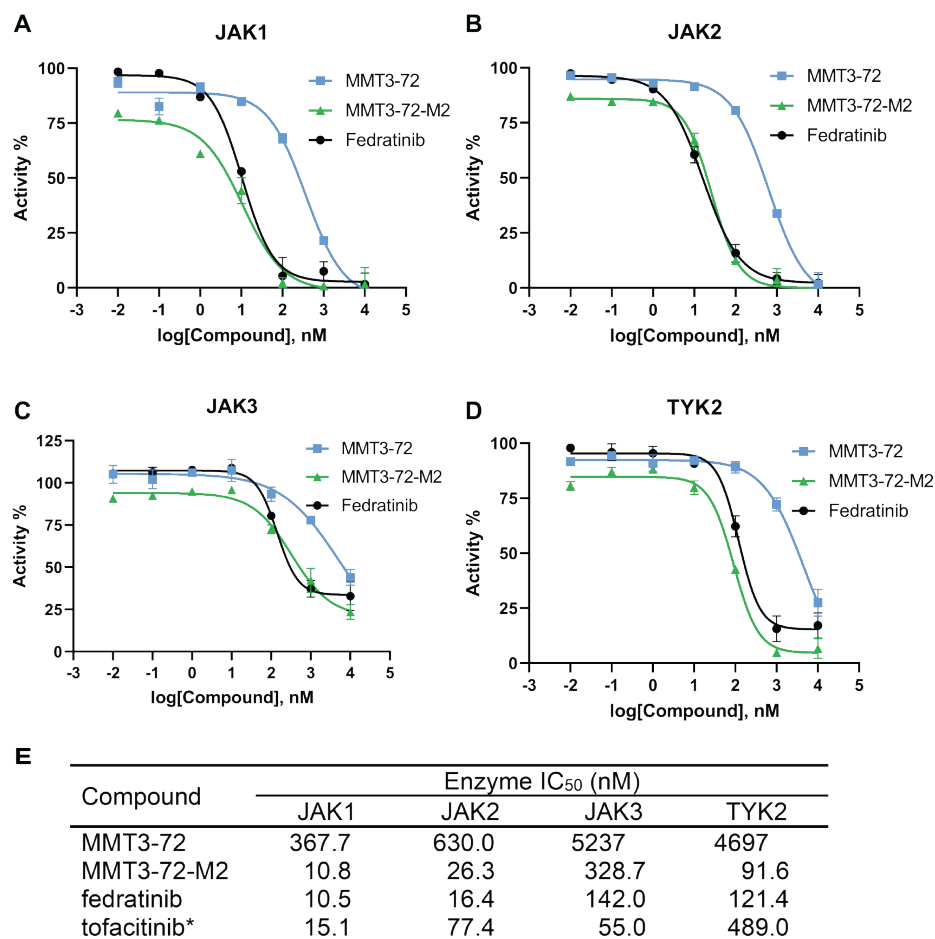
Figure 3.5: **Inhibition of different isoforms of JAK by MMT3-72 and active metabolite MMT3-72-M2.** Inhibition of JAK activity by MMT3-72 and MMT3-72-M2 (0.01 – 10,000 nM) was measured using Kinase-Glo Max assay against purified enzymes JAK1 (A), JAK2 (B), JAK3 (C), and TYK2 (D). Assays were run in the presence of 0.1 mM ATP. (E) *In vitro* inhibitory activities IC_{50} of MMT3-72 and active metabolite MMT3-72-M2 against different isoforms of JAKs. The IC_{50} values of compounds to inhibit different JAK isoforms were calculated using Prism 9. *The IC_{50} values of tofacitinib were cited from Ref. [19].

is reasonable to expect the efficient local conversion to MMT3-72-M2, allowing for a durable treatment after MMT3-72 reaches the colon region to inhibit all three JAK isoforms locally.

MMT3-72 was locally activated in the GI tract to release active metabolite MMT3-72-M2 which achieved high exposure in the GI tissues and minimized exposure in the plasma.

To investigate the GI local activation and pharmacokinetics of MMT3-72 and its active metabolite MMT3-72-M2 *in vivo*, mice were dosed orally with 10 mg/kg MMT3-72 and sacrificed to collect tissues at different time points from 0-24 h. As shown in Figure 3.6C, high concentration ($C_{max} > 50,000$ ng/g) of the MMT3-72 was observed in the colon content. However, MMT3-72 was not detected in either colon tissue or systemic circulation. In contrast, high levels of the active metabolite MMT3-72-M2 were detected in the colon tissue ($C_{max} > 1500$ ng/g) (In Figure 3.6D). The effective concentrations were last for 12 hrs. At 12 h, the concentration of MMT3-72-M2 was 131.3 ng/g ($\sim$ 285 nM) in the colon tissues, which exceeded the IC_{50} values for inhibiting JAK1, JAK2, and TYK2, the relevant targets for treating UC, but was below the IC_{50} for inhibiting JAK3 (Figure 3.5E). Further, the concentration of MMT3-72-M2 in plasma was minimal ($C_{max} = 8$ ng/mL) and was undetectable after 4 hours. These findings show that

1. MMT3-72 was not absorbed into the systemic circulation, was retained and activated in the colon region to release active metabolite MMT3-72-M2,

2. the active metabolite MMT3-72-M2 accumulated highly in colon tissues, which may inhibit JAK1, JAK2, and TYK2 for its therapeutic effects,

3. none of MMT3-72 and only low level of the active metabolite MMT3-72-M2 were detected in the systemic circulation, which has the potential to avoid systemic
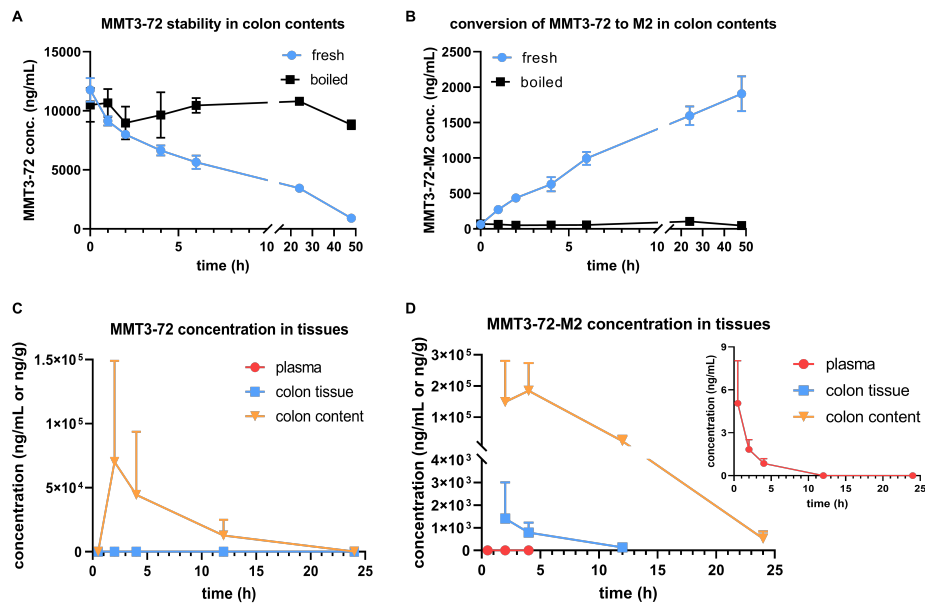
Figure 3.6: *In vitro* and *in vivo* activation of MMT3-72. (A) Concentrations of MMT3-72 at the indicated time points (0, 1, 2, 4, 6, 24, and 48 h) were measured after incubation with fresh and boiled colon contents. (B) Concentrations of MMT3-72-M2 were measured from the same collected samples in (A). (C) MMT3-72 concentrations in plasma, colon tissue, and colon content at 0.5, 2, 4, 12, and 24h. (D) MMT3-72-M2 concentrations in plasma, colon tissue, and colon content at 0.5, 2, 4, 12, and 24h. The inset figure shows the MMT3-72-M2 concentrations in the plasma.

toxicity of JAK inhibitions.

It is worth noting that our design of MMT3-72 is distinctly different from the design of izencitinib (TD-1473), which reduced the absorption potential to limit systemic exposure but without a local-activation mechanism.[45] The design of drugs (such as TD-1473) with only reduced absorption potential but without activation mechanism would reduce drug penetration in the colon tissue limiting its efficacy in human trials.[106] In contrast, MMT3-72 was designed to not only reduce the GI absorption potential but also have local activation properties to release the active form of MMT3-72-M2 that can easily penetrate colon tissue to reach a therapeutic concentration in the colon tissues while minimizing drug exposure in the systemic circulation.

MMT3-72 exhibited superior efficacy in treating UC in mice.

To evaluate the efficacy of MMT3-72 in the treatment of UC *in vivo*, we established a colitis model in mice using dextran sodium sulfate (DSS). DSS in drinking water could trigger colitis in mice. The DSS-induced colitis model is widely used because of its relatively easy administration and high similarity with human UC.[116] In this study, mice treated with 3% DSS water developed symptoms of colitis such as bloody stools and diarrhea on day 5. Disease activity index (DAI) was monitored for the severeness of disease in mice: Normal stool consistency with negative hemoccult: score 0; Soft stools with positive hemoccult: score 1; Very soft stools with traces of blood: score 2; Watery stools with visible rectal bleeding: score 3. To evaluate the efficacy of MMT3-72 in comparison with FDA-approved JAK inhibitor (tofacitinib) for UC treatment, mice were treated orally with 1mg/kg and 5mg/kg of both drugs (Figure 3.7A, 3.7B, 3.7C). MMT3-72 (5 mg/kg) improved DAI score by 5-fold in comparison with the DSS-induced colitis, while tofacitinib (5 mg/kg) did not show any improvement in DAI score (Figure 3.7A). In MMT3-72 (5 mg/kg) treatment group, no mouse developed severe colitis and only 10% mice (n = 10) developed moderate colitis (Figure 3.7B). In contrast, in the tofacitinib treatment group (5 mg/kg), 40% mice (n = 10) developed severe colitis and 40% developed moderate colitis (Figure 3.7B). Low dose (1 mg/kg) of both MMT3-72 (1 mg/kg) and tofacitinib (1 mg/kg) did not improve DAI score or disease severity in the DSS-induced colitis (Figure 3.7A, 3.7B).

Encouraged by these data, we tested high dose (10 mg/kg) of both MMT3-72 and tofacitinib for treatment of the DSS-induced UC (see Figure 3.7D, 3.7E, 3.7F). MMT3-72 (10 mg/kg) improved DAI score by 10-fold in the DSS-induced colitis model, and no mice (n = 10) developed moderate or severe colitis. In comparison, tofacitinib (10 mg/kg) also showed improvement in DAI score, and only 10% of mice developed severe disease with gross bleeding, and only 20% of mice developed mod-

erate colitis. High doses (10 mg/kg) of both MMT3-72 and tofacitinib recovered the colon length from the DSS-induced colitis (see Figure 3.7F). These data suggest that MMT3-72 has advantages in the treatment of UC.

To further evaluate the efficacy of MMT3-72 in reducing colon inflammation and tissue injury, H&E staining of colon tissues from the above *in vivo* studies were performed as shown in Figure 3.7G. The DSS-induced colitis showed severe and diffuse destruction of the epithelial layer with extensive immune cell infiltration in the epithelium. MMT3-72 (5, 10 mg/kg) reduced epithelial loss and decreased infiltration of immune cells in the DSS-induced colitis model. In contrast, tofacitinib (5 mg/kg) did not show improvement in epithelial cell loss and infiltration of immune cells in the DSS-induced colitis model while tofacitinib (10 mg/kg) showed moderate improvement. To investigate if the drug could inhibit JAKs and their downstream signaling, we also used western blot to detect the total STAT3 and phosphorylated STAT3. Consistent with the histology observation, treatment of MMT3-72 could effectively inhibit the phosphorylation of STAT3 (Figure 3.7H) in the inflammation tissue, thereby alleviating the production of inflammatory mediators in the colon region.

3.4   Conclusion

The JAK inhibitors have shown superior efficacy in the treatment of ulcerative colitis. Since UC needs life-long treatment, the therapeutic options require not only adequate efficacy but also minimal toxicity. However, all current JAK inhibitors, regardless of isoform specificity, have black box warnings for serious side effects including a high rate of major adverse cardiovascular events (MACE), thrombosis and pulmonary embolism, malignancies, and increased risk of serious infections.

In this study, we used machine learning-guided development of a GI locally activating JAK inhibitor for the treatment of ulcerative colitis. We designed STAR class
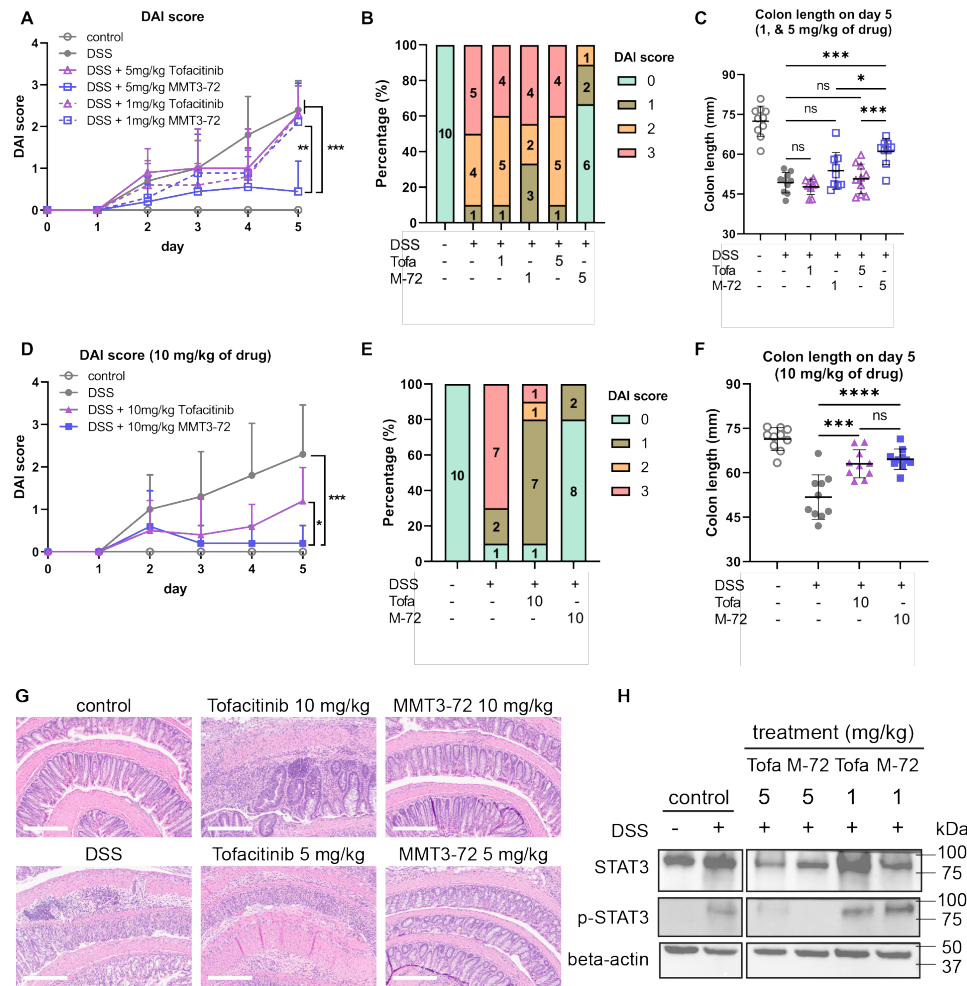
Figure 3.7: **In vivo efficacy of MMT3-72 in comparison with tofacitinib for UC treatment.** (A) Improvement of UC DAI score after treatment of MMT3-72 and tofacitinib (1, 5 mg/kg). (B) Percentage of mice within each DAI group on day 5 after treatment of MMT3-72 and tofacitinib (1, 5 mg/kg). (C) Recovery of colon length after treatment of MMT3-72 and tofacitinib (1, 5 mg/kg). (D) Improvement of UC DAI score after treatment of MMT3-72 and tofacitinib (10mg/kg). (E) Percentage of mice within each DAI group on day 5 after treatment of MMT3-72 and tofacitinib (10 mg/kg). (F) Recovery of colon length after treatment of MMT3-72 and tofacitinib (10 mg/kg). (G) H&E staining of colon tissues. Control was H&E staining of healthy mice colon tissue. The DSS-induced colitis showed disrupted epithelium and infiltration of immune cells in colon tissues. Treatment of MMT3-72 (5, 10 mg) reduced epithelium disruption and infiltration of immune cells in colon tissues in comparison with tofacitinib (5, 10 mg/kg) in the DSS-induced colitis model. The scale bar, indicated in the lower right corner, represents 300 μm (micrometers). (H) The expression level of JAK2 and STAT3 in colon tissues.

III JAK inhibitors (MMT3-72, MMT3-72-M2), which could achieve high local exposure/selectivity in the GI tissues while minimizing the requirement for JAK isoform specificity since none of the isoform-specific JAK inhibitors could eliminate their systemic toxicity. The molecular modeling showed that MMT3-72 has poor absorption potential with a low QED score and is expected to be retained in the GI tract. Two machine learning models predicted MMT3-72 is a JAK1 inhibitor (IC_{50} 67 nM) and MMT3-72-M2 is inhibitor of JAK1 (IC_{50} 19 nM), JAK2 (IC_{50} 34 nM), TYK2 (IC_{50} 65 nM).

Based on the machine learning prediction, we only needed to synthesize two compounds (MMT3-72 and MMT3-72-M2) for *in vitro* and *in vivo* testing on drug's potency/specificity, tissue selectivity, and efficacy in colitis models. *In vitro* kinase assays confirmed that MMT3-72 inhibited JAK1 (IC_{50} 368 nM), but was less active against JAK2, 3 and TYK2. MMT3-72-M2 inhibited JAK1 (IC_{50} 10.8 nM), JAK2 (IC_{50} 26.3 nM), TYK2 (IC_{50} 91.6 nM), and acted weakly on JAK3 (IC_{50} 328.7 nM). The pharmacokinetic study showed that a high level of MMT3-72 was accumulated in the lumen of the GI tract ($> 50,000$ ng/g), but was not detected in either GI tissues or plasma. However, MMT3-72 was locally activated primarily in the colon region to release the active form MMT3-72-M2, which showed a high accumulation in colon contents ($C_{max} > 50,000$ ng/g), colon tissues ($C_{max} > 1500$ ng/mL), but a minimal concentration ($C_{max} < 8$ ng/mL) in the plasma. Oral administration of MMT3-72 (5, 10 mg/kg, respectively) achieved superior efficacy, in comparison with tofacitinib, in DSS-induced colitis as measured by DAI score, bleeding, colon length, and pathological H&E staining of colon tissues. The oral dose of MMT3-72 also inhibited the phosphorylation of STAT3 in the colitis tissue. In summary, the integration of ML, the STAR system, and wet lab synthesis/testing could minimize the effort in optimization of a JAK inhibitor to treat colitis. This site-specific inhibitor reduces systemic toxicity by minimizing the need for JAK isoform specificity.

Supplementary information

EquiBind docking, Tanimoto similarity analysis for the datasets used in this study, visualized distribution of data used in MTATFP, western blot results for JAK2, QED of balsalazide, NMR spectra and HPLC chromatograms.

3.5   Declarations

Competing interests

The University of Michigan filed patent for these compounds, where some authors are inventors.

Authors' contributions

YB conducted experiments for machine learning, compounds stability, the *in vivo* pharmacokinetic and efficacy studies, and wrote the manuscript. MMT designed and synthesized compounds and wrote the manuscript. LZ performed the *in vitro* stability study, data visualization and wrote the manuscript. LW conducted metabolite identification studies and helped with LC-MS analysis. ZL tested the *in vitro* activity of compounds in inhibiting JAK enzymes. HH helped with the solvent selection, and supported the animal experiments and graphic abstract design. MW helped perform LC-MS analysis studies. CL helped conduct H&E staining. DS designed the experiments and wrote the manuscript.

Funding

Availability of data and materials

Datasets used to train models are all available and open to public. Information is provided in their github pages, respectively. Other data are available from the corresponding authors with reasonable request.

Supporting Information

QED of balsalazide

Table S3.1: Quantitative estimates of drug-likeness for balsalazide

Name	balsalazide
M.W.	357.32
ALOGP	2.71
HBD	4
HBA	6
PSA	148.65
ROTB	7
AROM	2
ALERTS	1
QED	0.5598

EquiBind docking
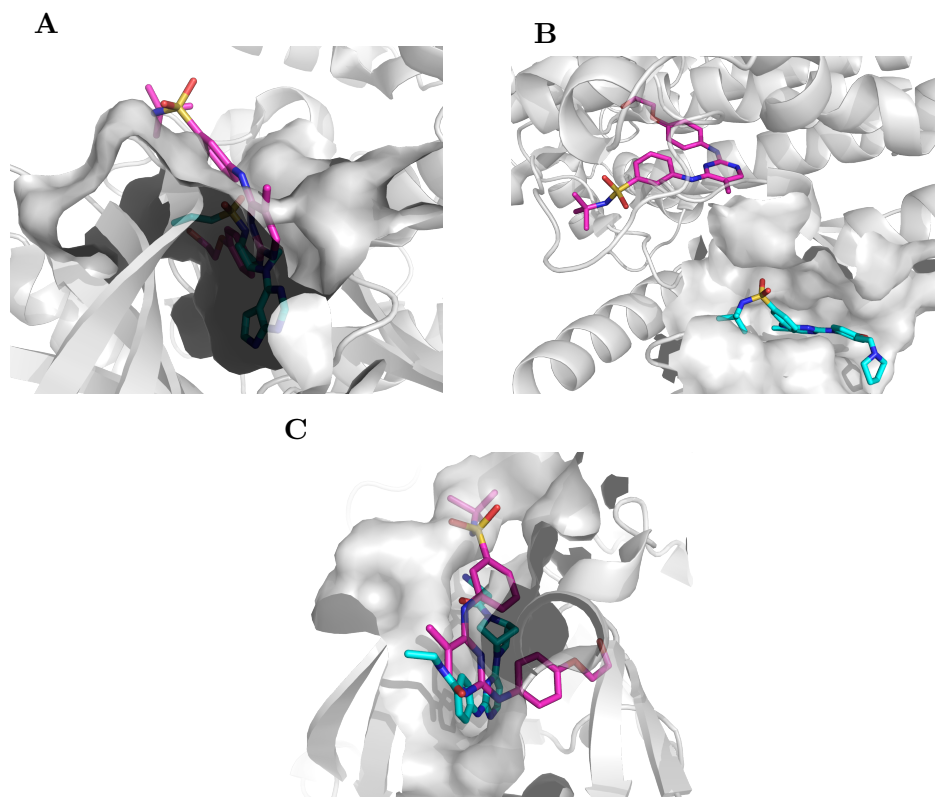
Figure S3.1: Docking results of MMT3-72-M2 to (A) JAK1, (B) JAK2, and (C) TYK2 by EquiBind. Magenta: MMT3-72-M2, cyan: co-crystalized ligand (PDB code: JAK1: 6BBU, JAK2: 6VNE, TYK2: 6DBK)

Tanimoto similarity analysis for the datasets used in this study

A

B
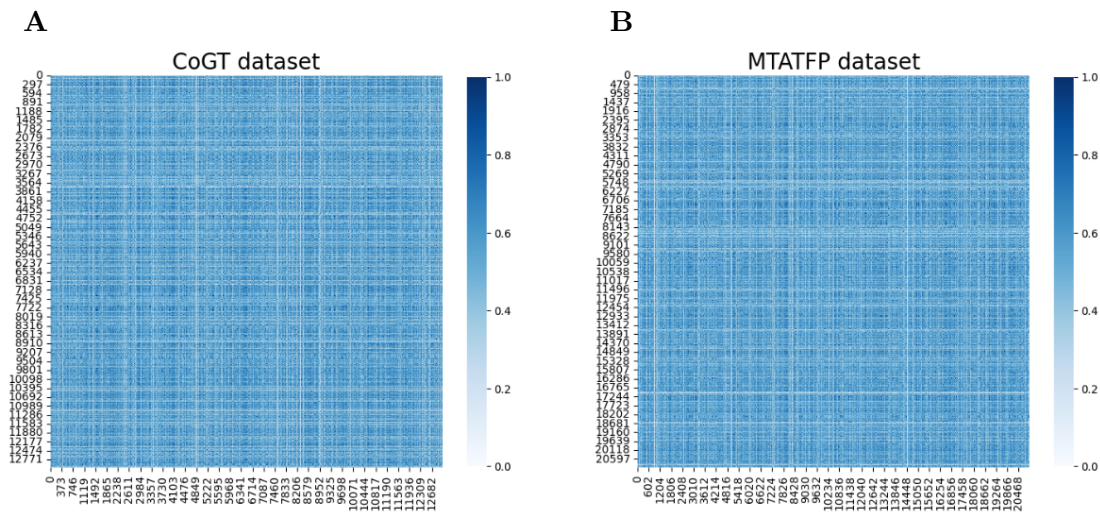
Figure S3.2: Tanimoto similarity for datasets used in (A) CoGT and (B) MTATFP.

Visualized distribution of data used in MTATFP
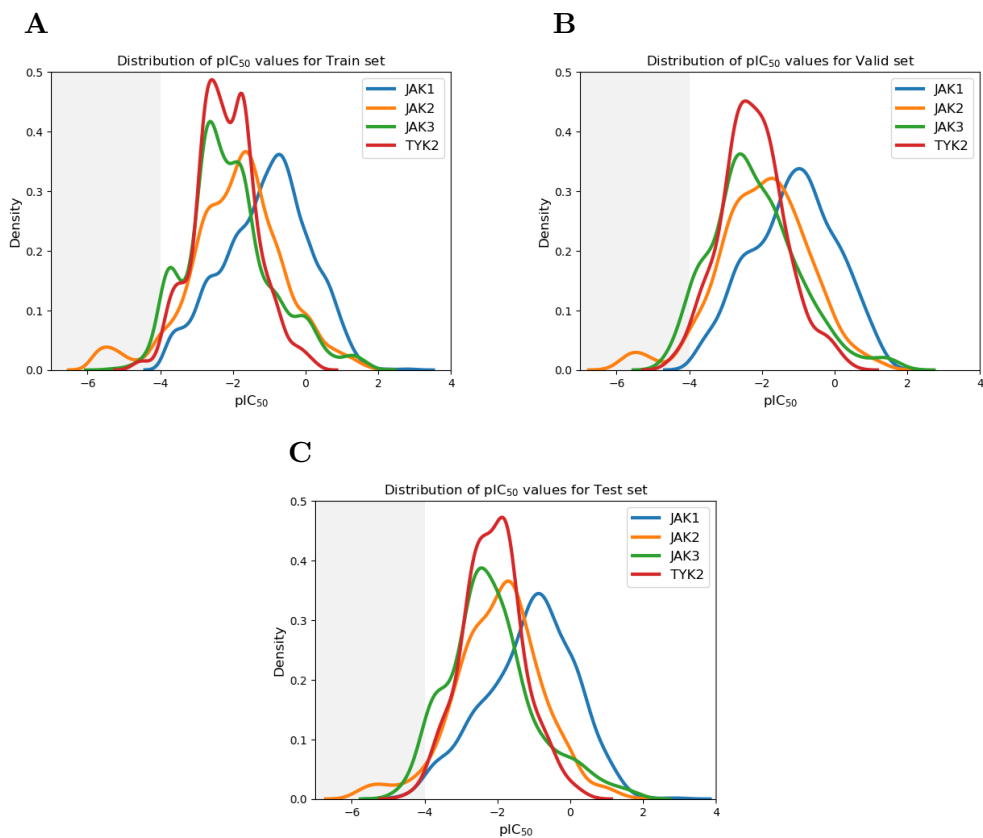
Figure S3.3: MTATFP data distribution for (A) training set, (B) validation set, and (C) testing set.

Western blot result (uncropped)
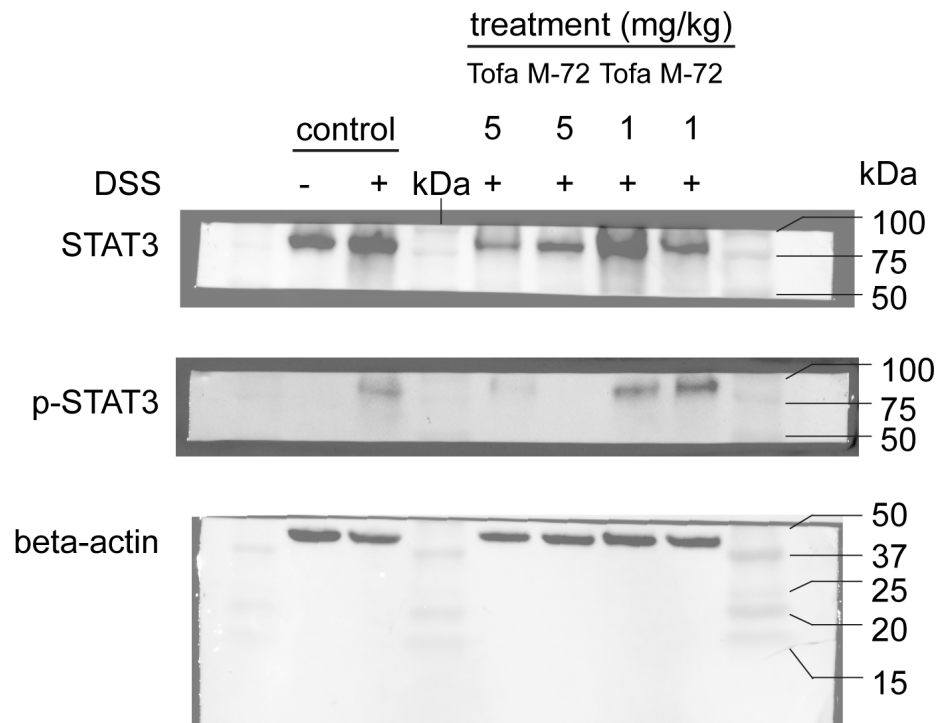
Figure S3.4: Uncropped western blots corresponding to the cropped western blots shown in the Figure 3.7.

NMR spectra and HPLC chromatograms

^{1}H-NMR spectrum of compound 2
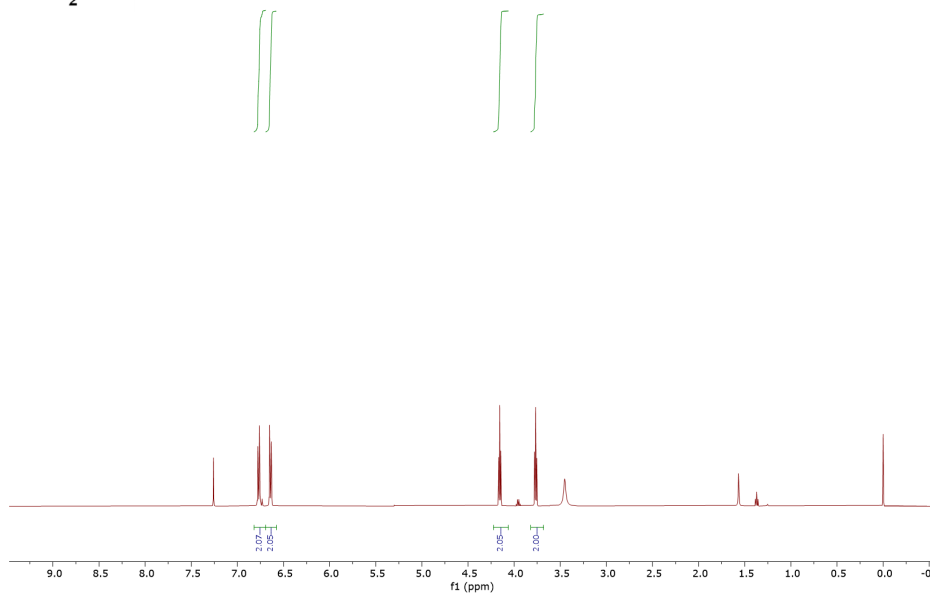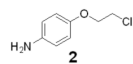
^{1}H-NMR spectrum of compound 3
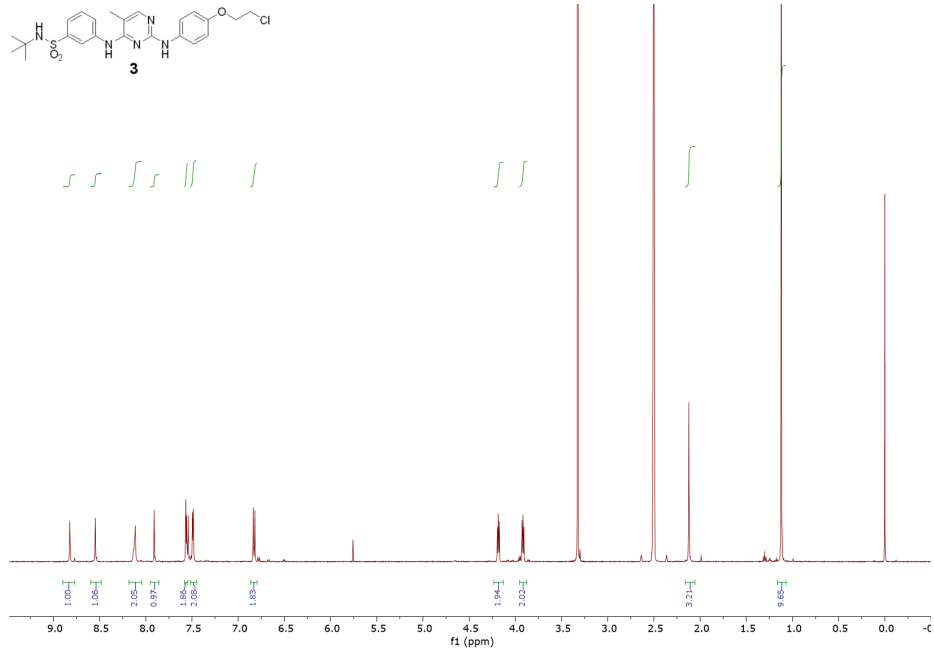
^{1}H-NMR spectrum of compound 5
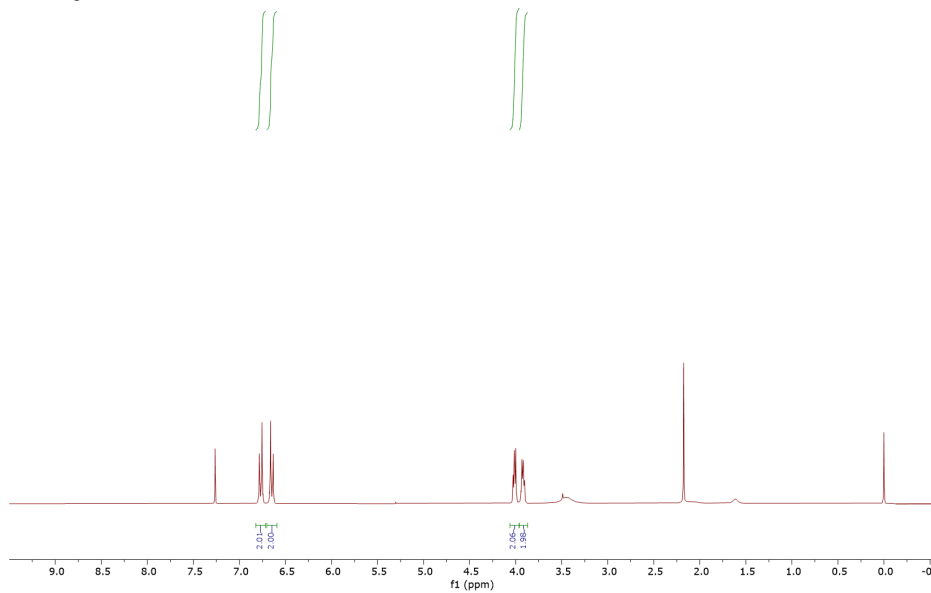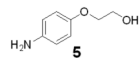
^{1}H-NMR spectrum of MMT3-72
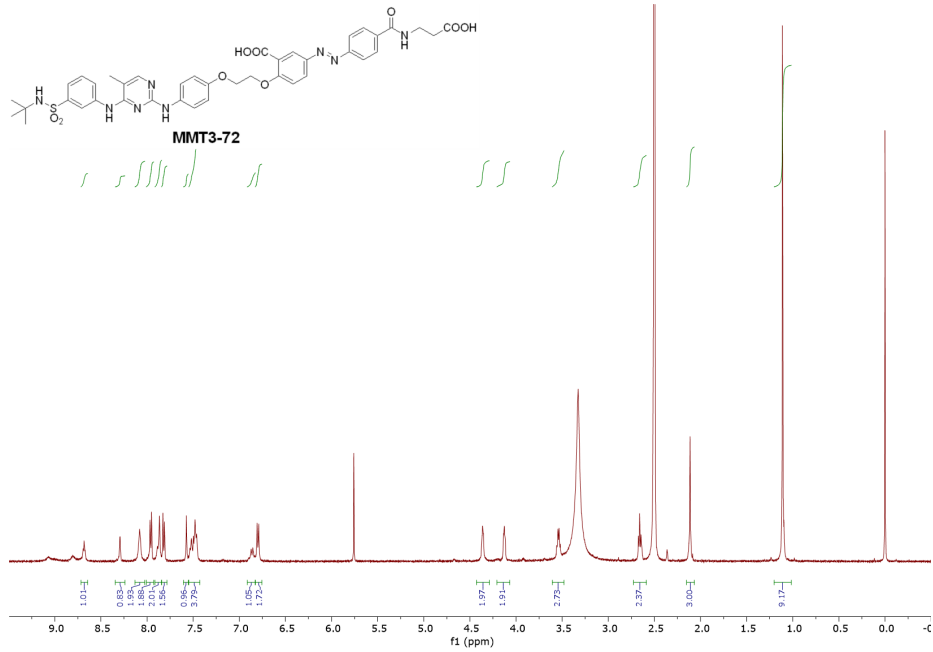
HPLC result of MMT3-72

HPLC: **MMT3-72**
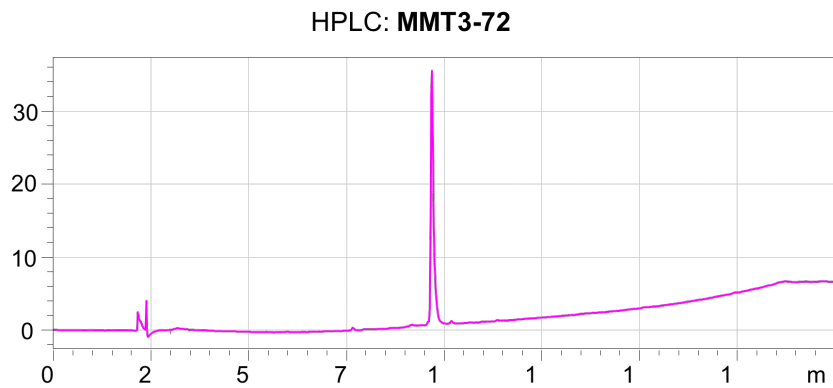
- Retention time: 9.7 min

- Purity $\sim 98\%$

- Concentration tested: 50 μg/mL in acetonitrile

- Detection wavelength: 254 nm

^{1}H-NMR spectrum of MMT3-72-M2
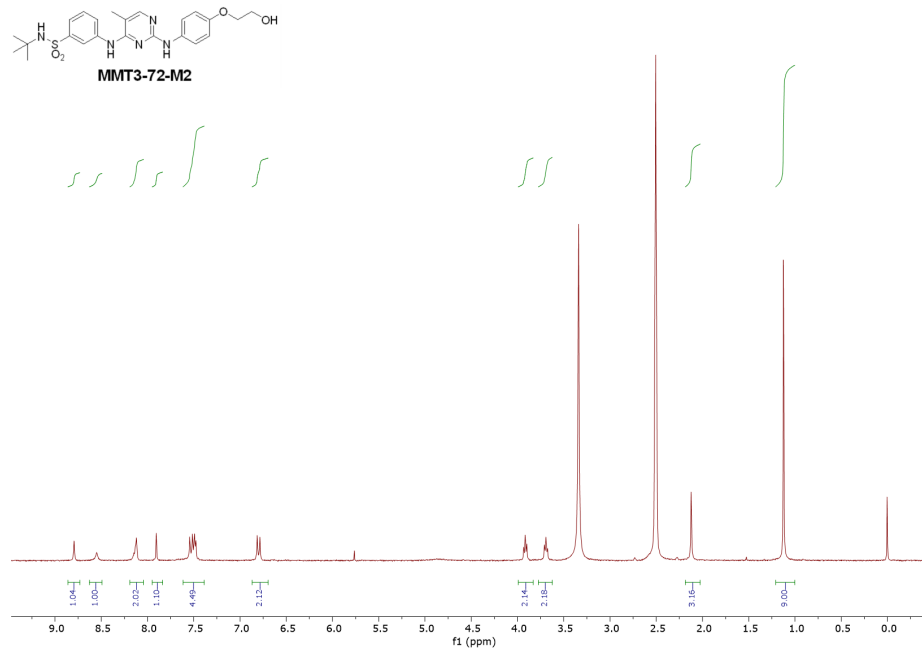
MMT3-72-M2

HPLC result of MMT3-72-M2

HPLC: **MMT3-72-M2**
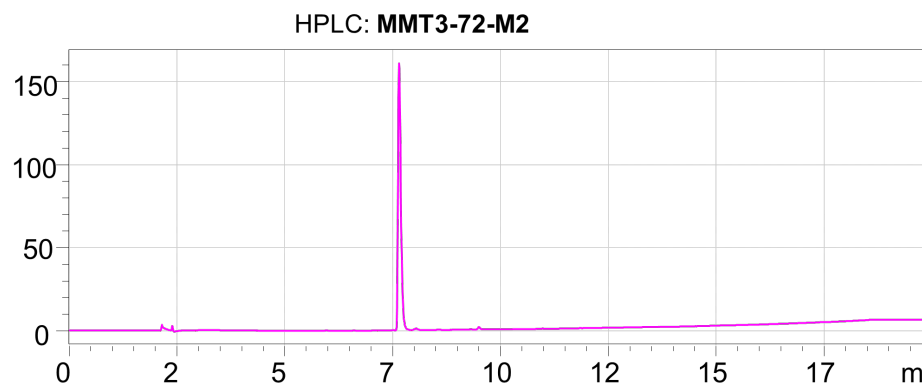
- Retention time: 7.7 min

- Purity ∼ 98%

- Concentration tested: 50 μg/mL in acetonitrile

- Detection wavelength: 254 nm

Mass spectra for MMT3-72 and its metabolites

MMT3-72
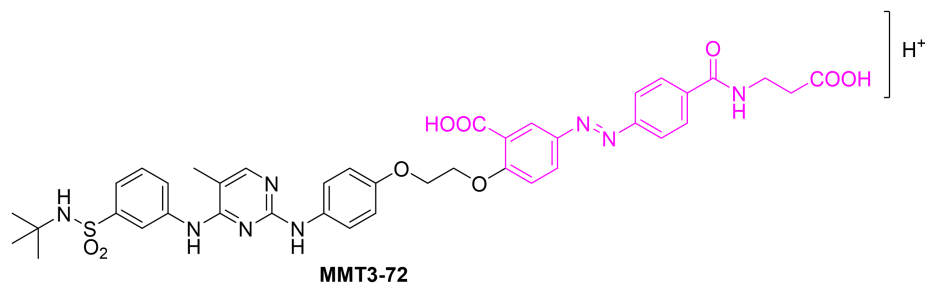
MMT3-72

Chemical Formula: $C_{40}H_{43}N_8O_9S^+$
Exact Mass: 811.2868
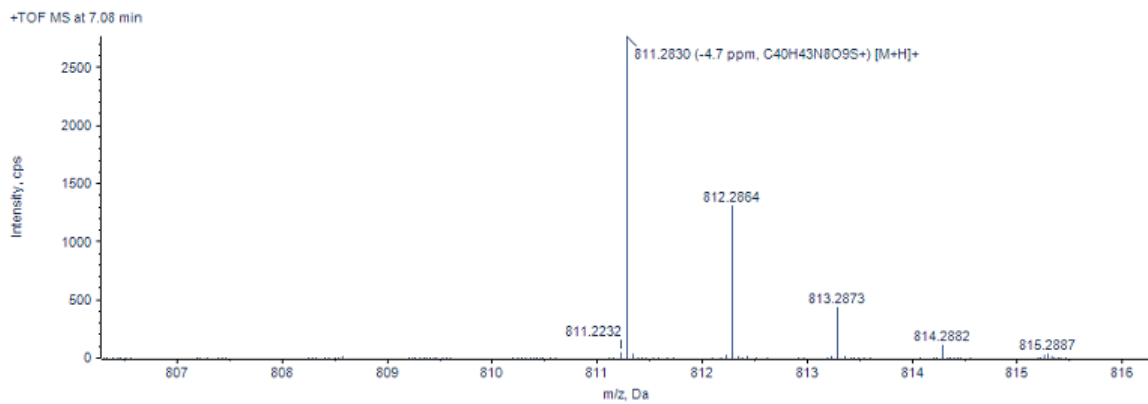
M1
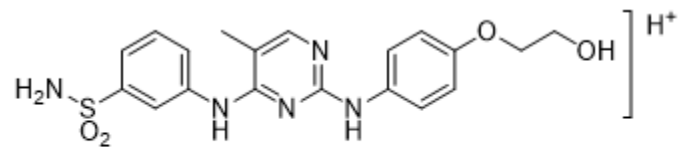
Chemical Formula: $C_{19}H_{22}N_5O_4S^+$
Exact Mass: 416.1387
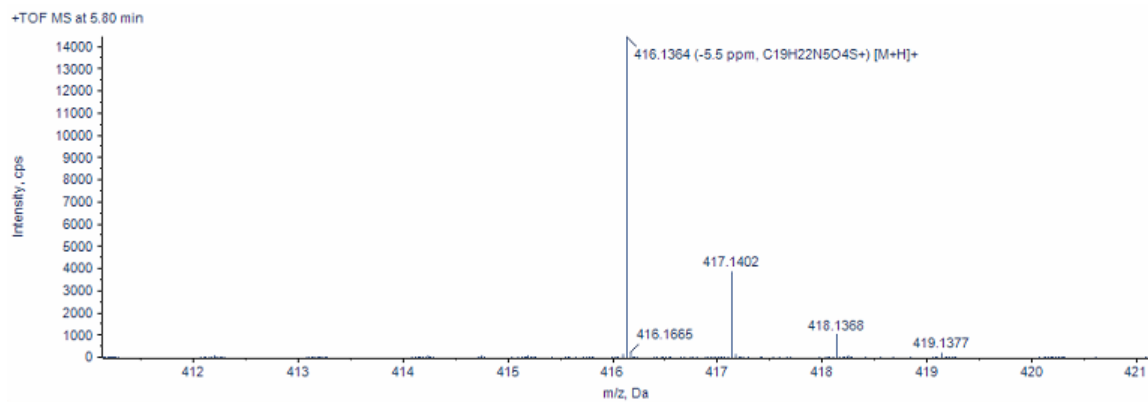
M2
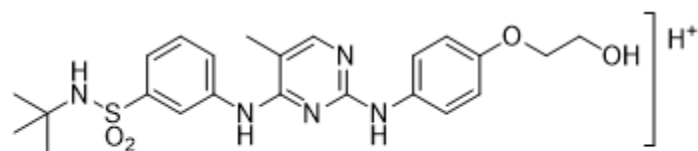
Chemical Formula: $C_{23}H_{30}N_5O_4S^+$
Exact Mass: 472.2013

+TOF MS at 5.80 min

472.1994 (-4.0 ppm, C23H30N5O4S+) [M+H]+

468.2477
467.2376
468.2142
469.2532
470.2617
471.2654
472.1607
473.2029
474.2002
475.1997

Intensity, cps

m/z, Da
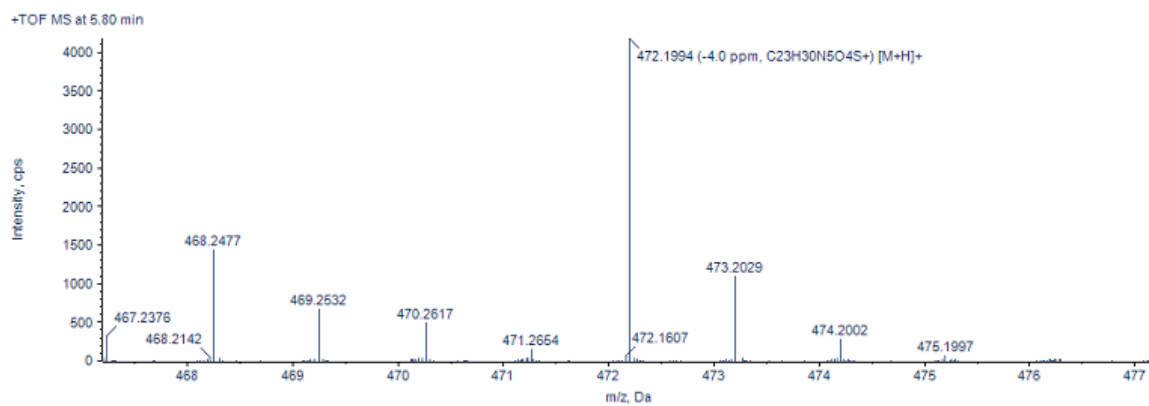
M3
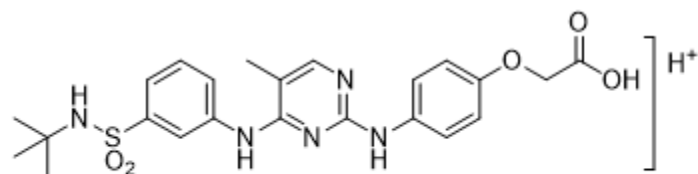
Chemical Formula: $C_{23}H_{28}N_5O_5S^+$
Exact Mass: 486.1806

M4

M4

Chemical Formula: $C_{38}H_{35}N_8O_9S^+$
Exact Mass: 755.2242
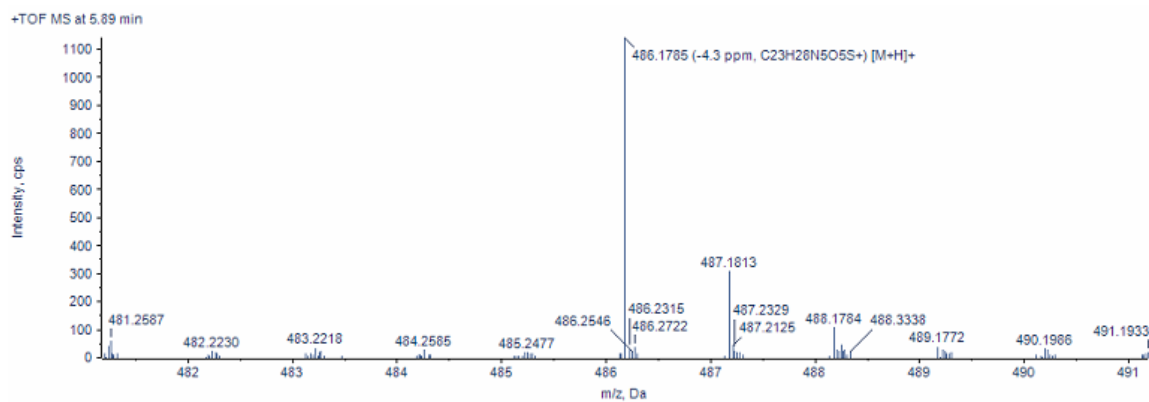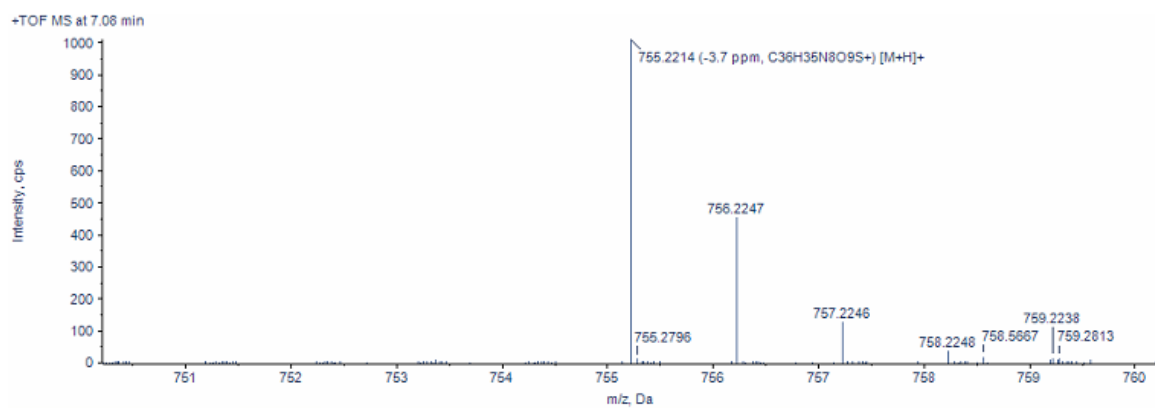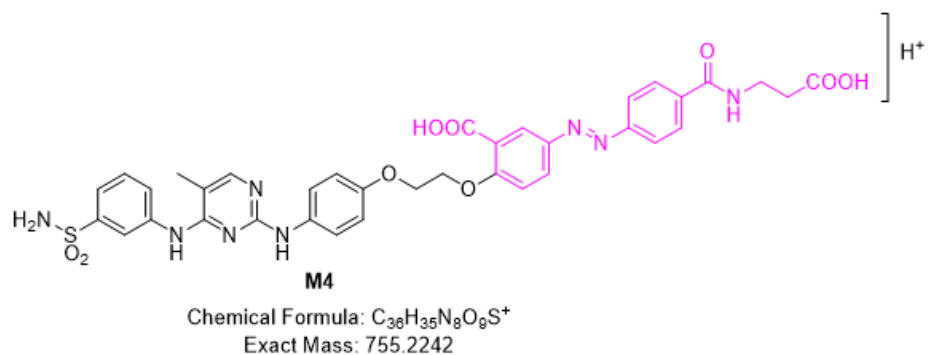
M5
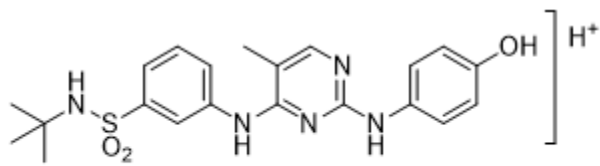
Chemical Formula: $C_{21}H_{26}N_5O_3S^+$
Exact Mass: 428.1751
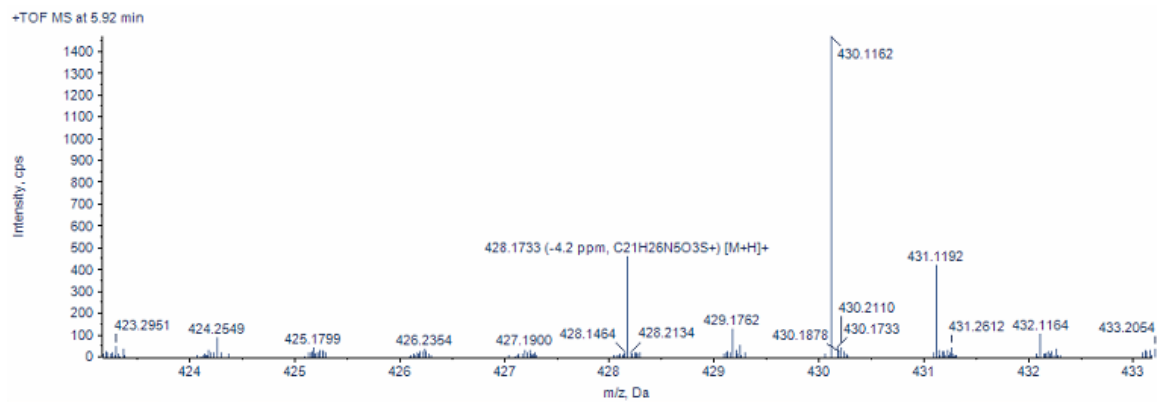
CHAPTER IV

Machine Learning-Enhanced Prediction of ADME-T Properties: Accelerating Drug Discovery and Optimization

Yingzi Bu[1,2], Luchen Zhang[1] and Duxin Sun[1,*]

Abstract

Traditionally, gaining insight into the ADME-T of a compound entails resource-intensive *in vitro* assays and animal studies for pre-clinical trials. We advocate for early PK and toxicity prediction in the drug design phase, leveraging ML models that utilize chemical structure data for a multitude of compounds. Our study offers a comprehensive comparison of ML techniques, including fingerprint-based, graph-based, and language-based algorithms. Notably, we highlight the superior performance of graph-based models in multi-objective learning, encompassing both classification and regression tasks. Additionally, we utilized a dynamic weight strategy, enhancing model efficiency and prediction performance by automatically calculating weights for

[1]Department of Pharmaceutical Sciences, College of Pharmacy, University of Michigan
[2]Michigan Institute for Computational Discovery & Engineering, University of Michigan
* To whom correspondence should be addressed: duxins@umich.edu

certain tasks. This innovative approach has the potential to streamline the process and cost-effectively identify promising drug candidates with optimal ADME-T profiles.

4.1 Introduction

Assessing ADME-T properties is critical in drug development and industrial chemical compound risk evaluation. High-quality ADME-T characteristics are essential for drug success, yet many candidates show unfavorable profiles, leading to lengthy timelines, high costs, attrition and even post-marketing withdrawals [43]. As *in vitro* and *in vivo* experiments are costly and time-consuming, *in silico* quantitative structure-property relationship (QSPR) models provide a valuable solution. They enable efficient evaluation and preliminary screening of prospective drug candidates, saving time and resources by identifying compounds likely to have unfavorable ADME-T properties before synthesis or experimentation.

In recent years, ML-driven models have advanced drug screening and discovery, employing methods such as linear regression, random forest (RF), support vector machine (SVM) and Deep Neural Network (DNN) to map compound structures to ADME-T properties [85, 92]. These ML-based models have become prominent, especially with sufficient data, as preferred tools for *in silico* drug evaluation. An exciting development in ML field is the rise of Graph Neural Network (GNN) that directly utilize graph-structured molecular representations rather than rely on various descriptors, demonstrating consistent superiority over classical ML or Deep Neural Network models [107]. GNN models exhibit robustness and over recent years and multiple GNN variants have consistently achieved state-of-the-art (SOTA) resulting predicting molecular properties. Particular noteworthy is the application of Graph Neural Network in predicting ADME-T properties, showcasing their potential to transform drug discovery [34, 87].

Multi-task learning has shown promise in improving the efficiency and accuracy of ADME-T tasks by training a model to handle multiple objectives simultaneously using a shared representation [12, 87, 34, 100, 138]. In conventional approaches, multi-task losses are combined using a simple weighted sum, often with uniform or manually adjusted loss weights [87, 96]. However, the effectiveness of these models largely depends on the choice of weights for each task's loss, a process that can be computationally expensive and challenging to fine-tune manually. In response to this challenge, Kendall et al. proposed the uncertainty weight strategy [61], which assigns weights to multiple tasks based on the inherent homoscedastic uncertainty of each task. This innovative approach enables the concurrent learning of various weights in both classification and regression scenarios.

Despite the growing prominence of ML in ADME-T prediction, the full potential of ML in this field has yet to be fully realized. MT learning, in particular, holds promise for addressing ADME-T challenges, but it has seldom explored the dynamic allocation of weights for different tasks. In our work, we seek to address this gap by comparing different ML architectures and assessing the impact of MT learning with and without the uncertainty weight strategy (as illustrated in Figure 4.1). Our primary contributions can be summarized as follows:

1. We propose a unified architecture suitable for different model types, including MLP using fingerprints, graph-based models, and sequence-based models, enabling simultaneous training on both regression and classification tasks.

2. We introduce the uncertainty weight strategy into this framework.

3. We conduct extensive performance evaluations of our approach on classification tasks, regression tasks, and a combination of both task types. The results demonstrate that incorporating different task information by sharing model parameters can lead to performance improvements, and the use of uncertainty

weights significantly enhances the prediction of ADME-T properties.

4.2 Methods

Data set and molecular representation

Data sets and data preparation. All data were extracted from Therapeutics Data Commons (TDC)[53], and each task data set is summarized in Table 4.1. The full information on each task is in supplementary material. For cleaning molecules, we used normalizer, remove salt, reionizer and uncharger to clean SMILES representations of molecules.

Table 4.1: Data set summary

Classification	
dataset	Total
CYP2C19_Veith[129]	12666
CYP2D6_Veith[129]	13130
CYP3A4_Veith[129]	9191
CYP1A2_Veith[129]	12580
CYP2C9_Veith[129]	12093
BBB_Martins[82]	2030
Bioavailability_Ma[80]	643
Pgp_Broccatelli[9]	1219
HIA_Hou[49]	580
PAMPA_NCATS[115]	2034
hERG_Karim[60]	13449
AMES[147]	7282
Regression	
dataset	Total
Caco2_Wang[132]	910
Lipophilicity_AstraZeneca[137]	4200
HydrationFreeEnergy_FreeSolv[86]	642
Solubility_AqSolDB[118]	9982
LD50_Zhu[154]	7385

Molecular representation. In this study, our model architecture takes drug information as input using different encoding schemes.

Graph. To represent the structure of the drug in a graph-wise approach, we extracted atom and bond information from the structure. Thus the graph for one compound is defined as $\mathcal{G} = (\mathcal{V}, \mathcal{E}) \in \mathbb{G}$, in which $\mathcal{V} = \{v_1, v_2, \cdots, v_N\}$ includes atoms' information N atoms and the bonds between the atoms are used as edges $\mathcal{E} \subseteq \mathcal{V} \times \mathcal{V}$. Node feature matrix $\mathbf{X} = \{x_1, x_2, \cdots, x_N\} \in \mathbb{R}^{N \times F}$, where $x_i \in \mathbb{R}^F$ represents node feature v_i with total feature number F. Adjacency matrix $\mathbf{A} \in \{0, 1\}^{N \times N}$ contains the adjacency information between atoms, in which $A_{ij} = 1$ if $(v_i, v_j) \in \mathcal{E}$. This input is suitable as input for graph-based algorithms.

Sequence. For sequence-based models such as RNN, we used drug simplified molecular input line entry systems (SMILES) or self-referencing embedded strings (SELFIES) [70], which is a textual representation of the chemical structure of drug. To convert SMILES or SELFIES into a suitable input format, we utilized a one-hot encoding approach. First we searched for all of the SMILES/SELFIES in our dataset and collected the alphabet $\mathcal{A} = \{a_1, a_2, \cdots, a_m\}$, in which a_i represents the ith unique token in SMILES/SELFIES. We also added a padding token, a "begin of string" token and a "end of string" token in the alphabet. Then the strings of drug representations could be converted into a vector of scalars with regards to the position of a certain token in the alphabet.

Fingerprint (FP). For the baseline model MLP, we used molecular fingerprints. The molecular fingerprints are projected as bit strings in most cases, although any vector of numerical values can be applied as a fingerprint in cheminformatics. MACCS (Molecular Access System by Molecular Design Limited), often referred to as the prototype of substructure key-based fingerprints, has two major types: the 166-key version and the 960-key version. The key-based fingerprints contain a group of molecular characteristics (e.g., atom environments) suitable for encoding the molecules. In this paper due to the size of the data set, we chose the 166-key version.

Model architecture

Baseline model MLP. MLP takes fingerprint as input and MLP with K layers is demonstrated below:

$$\mathbf{h}^{(1)} = \text{ReLU}(\mathbf{W}^{(1)} \cdot \mathbf{x}) \tag{4.1}$$

$$\mathbf{h}^{(2)} = \text{ReLU}(\mathbf{W}^{(2)} \cdot \mathbf{h}^{(1)}) \tag{4.2}$$

$$\vdots \tag{4.3}$$

$$\mathbf{h}^{(K)} = (\mathbf{W}^{(K)} \cdot \mathbf{h}^{(K-1)}) \tag{4.4}$$

$$\hat{\mathbf{y}}_i = \begin{cases} \text{Sigmoid}(\mathbf{h}_i^{(K)}) & \text{if task i is classification} \\ \mathbf{h}_i^{(K)} & \text{if task i is regression} \end{cases} \quad i \in \{1, \dots, n\} \tag{4.5}$$

In which $\mathbf{x}$ represents the drug (MACCS fingerprints). The output $\hat{\mathbf{y}}_i$ from MLP should either be a probability if the specific task is binary classification or the predicted value if the task is regression.

Attentive FP[146]. Attentive FP is a graph neural network architecture aimed at representing molecules, featuring an attention mechanism that effectively captures nonlocal interactions within the molecular context [127]. The fundamental concept behind employing the attention mechanism within a graph context is to acquire a context vector for the target node by concentrating on its neighboring nodes and the local surroundings [130]. The architecture of Attentive FP network involves atom embedding and molecular embedding to obtain the molecular representation for a specific compound.

For the k-th layer of atom embedding, the process involves:

$$e_{\mathbf{vu}}^k = \text{leaky\_relu}(\mathbf{W} \cdot [\mathbf{h}_{\mathbf{v}}^{k-1}, \mathbf{h}_{\mathbf{u}}^{k-1}]) \tag{4.6}$$

$$a_{\mathbf{vu}}^k = \text{softmax}(e_{\mathbf{vu}}^k) = \frac{\exp\left(e_{\mathbf{vu}}^k\right)}{\sum_{\mathbf{u}' \in \mathcal{N}(\mathbf{v})} \exp\left(e_{\mathbf{vu}'}^k\right)} \tag{4.7}$$

$$\mathbf{C}_{\mathbf{v}}^k = \text{elu}\left(\sum_{\mathbf{u} \in \mathcal{N}(\mathbf{v})} a_{\mathbf{vu}}^k \cdot \mathbf{W} \cdot \mathbf{h}_{\mathbf{u}}^{k-1}\right) \tag{4.8}$$

$$\mathbf{h}_{\mathbf{v}}^k = \text{GRU}^{k-1}(\mathbf{C}_{\mathbf{v}}^k, \mathbf{h}_{\mathbf{v}}^{k-1}) \tag{4.9}$$

where $\mathbf{v}$ represents the target node (a specific atom) and $\mathbf{u}$ is the node (a specific atom) in the neighborhood of $\mathbf{v}$, denoted as $\mathbf{u} \in \mathcal{N}(\mathbf{v})$. $\mathbf{h}_{\mathbf{v}}$ and $\mathbf{h}_{\mathbf{u}}$ is the node state vector for $\mathbf{v}$ and $\mathbf{u}$, respectively. To obtain the information for each target-neighbor pair $e_{\mathbf{vu}}$, the state vectors for both the target node and the neighbor node are concatenated, a linear transformation is performed with a trainable matrix $\mathbf{W}$, and the nonlinear activation function leaky_relu is performed. To obtain the importance (weight) $a_{\mathbf{vu}}$ for the target-neighbor pair $\mathbf{v}$-$\mathbf{u}$, which indicates the importance of neighbor node $\mathbf{u}$ to target node $\mathbf{v}$, $e_{\mathbf{vu}}$ is normalized by softmax over all nodes in the neighborhood of $\mathbf{v}$. For context vector $\mathbf{C}_{\mathbf{v}}$ of the target node $\mathbf{v}$, $\mathbf{h}_{\mathbf{u}}$ undergoes a linear transformation, followed by a subsequent step involving a weighted summation and the application of the elu activation function. During the readout phase, the GRU (gated recurrent unit) operates by taking as input the prior state vector of the target node, denoted as $\mathbf{h}_{\mathbf{v}}^{k-1}$, along with the message $\mathbf{C}_{\mathbf{v}}^{k-1}$. This input is then used to update the previous state to the current state, resulting in $\mathbf{h}_{\mathbf{v}}^k$.

Regarding molecule embedding, the procedure involves aggregating all atom embeddings by introducing a conceptual node that establishes connections with all individual atoms within the molecule. Specifically, the entire molecule is considered as a supervirtual node and is subject to embedding using the identical atom embedding attention mechanism. This operation is carried out across multiple attentive layers,

resulting in a state vector representing the entire molecule. The ultimate state vector serves as a learned representation encoding structural details of the molecular graph. Subsequently, a task-specific layer is applied for making predictions based on this representation the same as Eq. 4.5.

Graph Isomorphism Network (GIN) [149]. For popular GNN variants such as Graph Convolutional Networks [145] and GraphSAGE [44], the discriminative power of capturing different graph structures is not as powerful as GIN. Due to the "expensiveness" of GIN (Table 4.2), we used pretraining strategy [68] and fine-tuned pretrained GIN models on different downstream tasks instead of training GIN from scratch. Briefly, for a pretrained GNN model f, a learnable downstream network with parameters θ and downstream dataset $\mathcal{D} = \{(\mathcal{G}_i, y_i)\}_{i=1}^n$ with data size n, the training process will adjust the parameters in θ and/or f to maximize the likelihood, defined as $\max_{f,\theta} \sum_i^n P_{f,\theta}(y_i|\mathcal{G}_i)$.

Recurrent Neural Networks (RNN) [111]. RNNs are a class of artificial neural networks well-suited for sequential data analysis and Natural Language Processing (NLP) tasks, designed to process sequences of data by maintaining an internal hidden state that captures information from previous time steps. The hidden state allows RNNs to retain context and capture temporal dependencies in the data. Various advanced RNN architectures have been developed including Long Short-Term Memory (LSTM) networks and Gated Recurrent Unit (GRU) networks, improving gradient flow and are better at capturing long-range dependencies [40, 27]. We applied GRU layers within our RNN model.

Multi-objective learning strategy

In this approach [12], the model is concurrently trained for n tasks, sharing parameters to acquire knowledge that may be lacking in the original single-task setting. Even in cases where the final performance does not exhibit improvement compared to

that of a single-task model, the shared parameters contribute to enhanced efficiency, as exemplified in Table 4.2 and Table 4.3. To illustrate, when considering n tasks and the parameter requirements for a single-task model amount to M parameters with a corresponding runtime of t, training n separate models would necessitate roughly n × M parameters and n × t runtime. Conversely, employing a multi-task model requires only M + m′ parameters, where m′ accounts for the parameter disparity in the final layer of the single-task model architecture. Similarly, the runtime is significantly reduced. This multi-objective learning approach demonstrates its capability to simultaneously address multiple regression and classification tasks, and it can even accommodate a combination of regression and classification tasks within the same multi-objective framework.
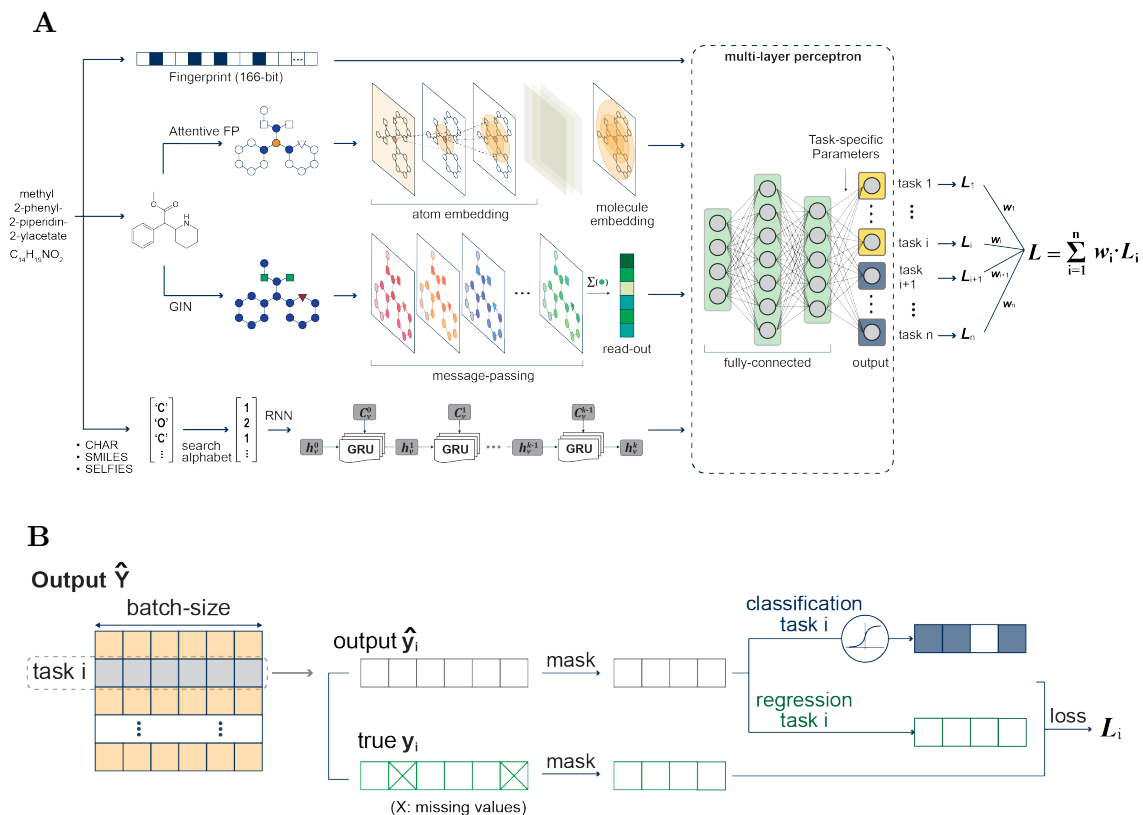
Figure 4.1: Model architecture. A, overall flow for MT learning with uncertainty weight strategy. B, thorough management of outputs for both regression and classification tasks.

Table 4.2: Runtime (ms) comparison multi-task vs single-task

Model #Task	MLP	Attentive FP	GIN	RNN
1	0.166 ± 0.005	0.340 ± 0.007	2.453 ± 0.041	1.131 ± 0.018
2	0.176 ± 0.006	0.347 ± 0.009	2.451 ± 0.034	1.143 ± 0.012
3	0.181 ± 0.005	0.357 ± 0.012	2.460 ± 0.032	1.160 ± 0.014
4	0.186 ± 0.005	0.364 ± 0.025	2.461 ± 0.030	1.163 ± 0.012
5	0.193 ± 0.004	0.365 ± 0.011	2.473 ± 0.031	1.166 ± 0.012

Table 4.3: model parameters comparison multi-task vs single-task

Model #Task	MLP	Attentive FP	GIN	RNN
1	30817	3823507	1977165	2953125
2	30834	3823808	1977182	2953382
3	30851	3824109	1977199	2953639
4	30868	3824410	1977216	2953896
5	30885	3824711	1977233	2954153

Training procedure

Training process. Training was done on one GPU of NVIDIA Tesla T4. Train: valid: test split = 7: 1: 2 is done by package TDC. We strictly followed split to avoid information leak (We only use training set to train the model, validation set to valid and fine-tune the parameters, and test set is only used for final performance evaluation).

Early Stopping. Early stopping was implemented to mitigate overfitting and expedite the training process. To achieve the best possible performance within the current hyperparameter configuration, we initiated a training process with early stopping strategy. If there was no improvement in the performance on the validation set for 30 consecutive epochs, the training process would be terminated early. This approach ensured that, based on the existing hyperparameters, early stopping would yield the most favorable performance on the validation set. A demonstration is shown in Figure 4.2.

Training protocol. For fairness of performance comparison, ll models were trained using the PyTorch framework, with optimization performed through the AdamW

92
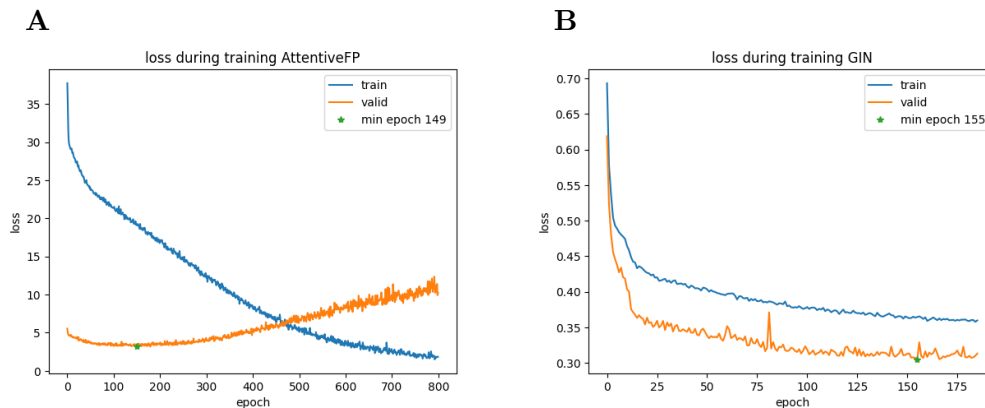
Figure 4.2: Early stopping. A, no early stopping on Attentive FP; B, early stopping applied, GIN

optimizer for gradient descent. In detail, given a training set $\{(x_i, y_i)|N = |X|, i \in \{1, \cdots, N\}\}$, where x_i and y_i are input data and the true labels for the i-th compound in training data, the model initially computes the output of the training set batch in the forward pass, and then adjusts it through supervised back-propagation to minimize the loss function, denoted as $\mathcal{L}(\cdot)$, for each epoch iteration. For regression tasks (e.g. Lipophilicity), we utilized the MSELoss as the loss function, measuring mean-squared error formulated as Equation (4.10). While for classification tasks (e.g. CYP2D6), we employed BCEWithLogitsLoss, which is a combination of the Sigmoid function and BCELoss to compute cross-entropy, shown in Equation (4.11).

$$\mathcal{L}(\theta) = \sum_{i=1}^{n}(y_i - (\hat{y}_i|x_i; \theta))^2, \tag{4.10}$$

$$\mathcal{L}(\theta) = \sum_{i=1}^{n} y_i \cdot \log(\sigma(\hat{y}_i|x_i, \theta)) \tag{4.11}$$

where n denotes the batch size, θ represents the model parameters, y_i and $\hat{y}_i$ are the ground truth and model raw prediction, respectively.

For multi-task learning, we used the weighted sum of the loss defined as:

$$\mathcal{L}(\theta) = \sum_{m=1}^{M} w_m \mathcal{L}_m(\theta) \tag{4.12}$$

where M denotes the number of tasks, and w_m denotes the fixed/trainable weight for the specific task m. Task weights were initialized as the same for different tasks. For trainable task weight, we adopted a dynamic weight strategy proposed by Kendall et al [61]. Briefly, For regression tasks, the likelihood, log likelihood are:

$$p(\mathbf{y}|\hat{\mathbf{y}}) = \mathcal{N}(\hat{\mathbf{y}}, \sigma^2) \tag{4.13}$$

$$\log p(\mathbf{y}|\hat{\mathbf{y}}) \propto -\frac{1}{2\sigma^2}\|\mathbf{y} - \hat{\mathbf{y}}\|^2 - \log \sigma \tag{4.14}$$

in which $\hat{\mathbf{y}}$ represents the output of the model, and σ is a noise. Then loss function of multi (task num k) regression tasks can be defined as minimizing the negative log likelihood:

$$
\begin{aligned}
-\log p(\mathbf{y}_1, \cdots, \mathbf{y}_k|\hat{\mathbf{y}}) &= -\sum_{i=1}^{k} \log p(\mathbf{y}_i|\hat{\mathbf{y}}) \\
&\propto \sum_{i=1}^{k} \left(\frac{1}{2\sigma_i^2}\|\mathbf{y}_i - \hat{\mathbf{y}}\|^2 + \log \sigma_i \right) \\
&= \sum_{i=1}^{k} \left(\frac{1}{2\sigma_i^2}\mathcal{L}_i + \log \sigma_i \right)
\end{aligned}
\tag{4.15}
$$

where Euclidean loss for regression task i is defined as $\mathcal{L}_i = \|\mathbf{y}_i - \hat{\mathbf{y}}\|^2$.

While for classification, the loss is defined as negative log Softmax:

$$\mathcal{L} = -\log \text{Softmax}(\mathbf{y}, \hat{\mathbf{y}}) = -\log \frac{\exp(\hat{\mathbf{y}}_c)}{\sum_{c_i} \exp(\hat{\mathbf{y}}_{c_i})} = -\hat{\mathbf{y}}_c + \log \left(\sum_{c_i} \exp(\hat{\mathbf{y}}_{c_i}) \right) \tag{4.16}$$

where $\hat{\mathbf{y}}_c$ is the c-th element in model output $\hat{\mathbf{y}}$. To introduce uncertainty for classifi-

cation tasks, the authors adapted a scaled version to compute classification likelihoods by a softmax function:

$$p(\mathbf{y}|\hat{\mathbf{y}}, \sigma) = \text{Softmax}(\frac{1}{\sigma^2} \cdot \hat{\mathbf{y}}) \tag{4.17}$$

Thus the negative log likelihood can be expressed using loss function:

$$
\begin{aligned}
-\log p(\mathbf{y} = c|\hat{\mathbf{y}}, \sigma) &= -\frac{1}{\sigma^2} \cdot \hat{\mathbf{y}}_c + \log \sum_{c'} \exp(\frac{1}{\sigma^2} \cdot \hat{\mathbf{y}}_{c'}) \\
&= \frac{1}{\sigma^2}\mathcal{L} - \frac{1}{\sigma^2}\log\left(\sum_{c_i} \exp(\hat{\mathbf{y}}_{c_i})\right) + \log \sum_{c'} \exp(\frac{1}{\sigma^2} \cdot \hat{\mathbf{y}}_{c'}) \\
&= \frac{1}{\sigma^2}\mathcal{L} + \log\sigma - \frac{1}{\sigma^2}\log\left(\sum_{c_i} \exp(\hat{\mathbf{y}}_{c_i})\right) + \log\frac{1}{\sigma} \sum_{c'} \exp(\frac{1}{\sigma^2} \cdot \hat{\mathbf{y}}_{c'}) \\
&\approx \frac{1}{\sigma^2}\mathcal{L} + \log\sigma \quad \text{when } \sigma \to 1
\end{aligned}
\tag{4.18}
$$

Finally the training is defined as minimizing the negative log likelihood:

$$
\begin{aligned}
\mathcal{L} &= -\log p(\cdots, \mathbf{y}_i, \cdots, \mathbf{y}_j = c, \cdots|\hat{\mathbf{y}}) \\
&= -\sum_t \log p(\mathbf{y}_t|\hat{\mathbf{y}}) \\
&= \sum_i (\frac{1}{2\sigma_i^2}\mathcal{L}_i + \log\sigma_i) + \sum_j (\frac{1}{\sigma_j^2}\mathcal{L}_j + \log\sigma_j)
\end{aligned}
\tag{4.19}
$$

in which task i is a regression task while task j is a classification task. In our work, the different weights were updated every epoch instead of every batch for model training.

Training continued until reaching an early termination criterion, indicating convergence in performance improvement.

Performance metrics and benchmarks

We included both regression task and classification task. For the classification task: the above-mentioned models were rigorously assessed using multiple evaluation metrics on designated test sets. Assessment metrics include Accuracy, Active Recall

(Sensitivity, SE), Negative Recall (Specificity, SP), Weighted Accuracy (mean of SE and SP), Matthew's Correlation Coefficient (MCC), F1 Score, Area Under Curve (AUC), and Average Precision (AP). These metrics involved calculations based on True Positive (TP), True Negative (TN), False Positive (FP), and False Negative (FN) values. Furthermore, in the AP calculation, R_n represents the precision at the n-th threshold, while P_n represents the recall at the same n-th threshold. Corresponding formulas for each metric are detailed below:

$$\text{Accuracy} = \frac{\text{TP} + \text{TN}}{\text{TN} + \text{TP} + \text{FP} + \text{FN}} \tag{4.20}$$

$$\text{Precision} = \frac{\text{TP}}{\text{TP} + \text{FP}} \tag{4.21}$$

$$\text{Recall(SE)} = \frac{\text{TP}}{\text{TP} + \text{FN}} \tag{4.22}$$

$$\text{SP} = \frac{\text{TN}}{\text{TN} + \text{FP}} \tag{4.23}$$

$$\text{Weighted Accuracy} = \frac{\text{SE} + \text{SP}}{2} \tag{4.24}$$

$$\text{MCC} = \frac{\text{TP} \times \text{TN} - \text{FN} \times \text{FP}}{\sqrt{(\text{TP} + \text{FP})(\text{TP} + \text{FN})(\text{TN} + \text{FN})(\text{TN} + \text{FP})}} \tag{4.25}$$

$$\text{F1} = \frac{2(\text{Precision} \times \text{Recall})}{\text{Precision} + \text{Recall}} \tag{4.26}$$

$$\text{AP} = \sum_n (\text{R}_n - \text{R}_{n-1})\text{P}_n \tag{4.27}$$

For the regression task: performance evaluation of the regression models was carried out using three key metrics: Mean Absolute Error (MAE), Root Mean Square Error (RMSE), and Coefficient of Determination (R^2). These were chosen for their efficacy in conveying predictive accuracy, averaging error magnitude, and correlation

96

strength respectively. The respective formulas for each metric are detailed below:

$$\text{MAE}(\mathbf{y}, \hat{\mathbf{y}}) = \frac{\sum_{i=1}^{n} |y_i - \hat{y}_i|}{n} \tag{4.28}$$

$$\text{RMSE}(\mathbf{y}, \hat{\mathbf{y}}) = \sqrt{\frac{\sum_{i=1}^{n} (y_i - \hat{y}_i)^2}{n}} \tag{4.29}$$

$$\text{R}^2(\mathbf{y}, \hat{\mathbf{y}}) = 1 - \frac{\sum_{i=1}^{n} (y_i - \hat{y}_i)^2}{\sum_{i=1}^{n} (y_i - \bar{y})^2} \tag{4.30}$$

where y_i is the true value for the i-th sample, $\hat{y}_i$ is the i-th predicted value, $\bar{y}$ is the mean of the sum of all true values, and n is the number of samples.

4.3 Results and Discussion

Chemical Diversity Analysis. In our pursuit of understanding the chemical diversity inherent in our datasets across various tasks, we conducted principal component analysis (PCA) and t-distributed stochastic neighbor embedding (t-SNE) using the MACCS fingerprint as the input representation. The results of these analyses, illustrated in Figure 4.3 for PCA and Figure 4.4 for t-SNE, reveal an intriguing overlap in the distribution of compounds with different label values. This overlap signifies that within the chemical space, compounds with different labels may share similar structures and properties, rendering their distinct separation challenging. Furthermore, to quantify the extent of similarity, we employed the Tanimoto similarity metric as depicted in Figure 4.5. The outcomes of this analysis highlight the wide distribution of molecules within the datasets, particularly evident in the Solubility, LD_{50}, and AMES datasets. This diversity underscores the representativeness of our data and underscores the robust flexibility and generalizability of our models in predicting task labels for a broad spectrum of compounds.

Performance Evaluation and Comparison of Models (ST). In our initial assessment, we focused on evaluating the performance of models concerning single

Table 4.4: Results of test set in Metabolism binary classification tasks. ST, single-objective; MT, multi-objective, models were trained on 5 tasks simultaneously.

task	model	acc	w_acc	SE	recall	SP	F1	AUC	MCC	AP
CYP2C19	ST-MLP	0.781	0.779	0.763	0.758	0.801	0.760	0.852	0.559	0.808
	MT-MLP	0.809	0.810	0.773	0.825	0.794	0.798	0.877	0.618	0.842
	ST-AttentiveFP	0.818	0.819	0.788	0.827	0.811	0.807	0.890	0.636	0.857
	MT-AttentiveFP	0.836	0.838	0.795	0.866	0.810	0.829	0.917	0.674	0.901
	ST-GIN	0.857	0.857	0.837	0.856	0.858	0.846	0.927	0.713	0.912
	MT-GIN	0.856	0.856	0.836	0.854	0.858	0.845	0.929	0.710	0.915
	ST-RNN	0.789	0.790	0.754	0.803	0.777	0.778	0.850	0.579	0.796
	MT-RNN	0.819	0.819	0.794	0.818	0.820	0.806	0.886	0.637	0.856
CYP2D6	ST-MLP	0.863	0.695	0.721	0.428	0.962	0.537	0.832	0.484	0.628
	MT-MLP	0.868	0.729	0.699	0.508	0.950	0.588	0.848	0.521	0.655
	ST-AttentiveFP	0.883	0.753	0.758	0.545	0.960	0.634	0.877	0.578	0.710
	MT-AttentiveFP	0.885	0.766	0.744	0.578	0.955	0.651	0.889	0.589	0.727
	ST-GIN	0.888	0.763	0.770	0.564	0.962	0.651	0.899	0.596	0.743
	MT-GIN	0.888	0.783	0.741	0.615	0.951	0.672	0.904	0.609	0.760
	ST-RNN	0.848	0.672	0.650	0.391	0.952	0.488	0.803	0.423	0.555
	MT-RNN	0.858	0.666	0.739	0.361	0.971	0.485	0.830	0.449	0.628
CYP3A4	ST-MLP	0.768	0.758	0.729	0.697	0.818	0.713	0.848	0.519	0.788
	MT-MLP	0.768	0.756	0.732	0.689	0.823	0.710	0.850	0.518	0.793
	ST-AttentiveFP	0.811	0.806	0.764	0.783	0.830	0.773	0.898	0.611	0.857
	MT-AttentiveFP	0.820	0.812	0.789	0.770	0.855	0.779	0.905	0.627	0.865
	ST-GIN	0.838	0.835	0.794	0.820	0.850	0.807	0.920	0.667	0.884
	MT-GIN	0.837	0.833	0.801	0.806	0.859	0.804	0.920	0.665	0.886
	ST-RNN	0.741	0.717	0.734	0.582	0.852	0.649	0.830	0.456	0.767
	MT-RNN	0.788	0.785	0.732	0.766	0.803	0.749	0.866	0.566	0.821
CYP1A2	ST-MLP	0.828	0.827	0.821	0.811	0.842	0.816	0.902	0.654	0.890
	MT-MLP	0.834	0.834	0.823	0.827	0.840	0.825	0.911	0.668	0.901
	ST-AttentiveFP	0.858	0.858	0.843	0.860	0.856	0.851	0.935	0.716	0.928
	MT-AttentiveFP	0.865	0.865	0.853	0.863	0.867	0.858	0.942	0.729	0.936
	ST-GIN	0.876	0.875	0.873	0.863	0.888	0.868	0.947	0.751	0.943
	MT-GIN	0.885	0.885	0.866	0.894	0.877	0.880	0.950	0.770	0.945
	ST-RNN	0.823	0.822	0.819	0.803	0.841	0.811	0.906	0.645	0.895
	MT-RNN	0.841	0.842	0.817	0.855	0.828	0.836	0.919	0.683	0.910
CYP2C9	ST-MLP	0.794	0.752	0.729	0.623	0.881	0.671	0.866	0.526	0.725
	MT-MLP	0.811	0.785	0.730	0.703	0.867	0.716	0.883	0.575	0.773
	ST-AttentiveFP	0.834	0.811	0.762	0.740	0.882	0.751	0.900	0.626	0.788
	MT-AttentiveFP	0.851	0.837	0.771	0.795	0.879	0.783	0.922	0.669	0.841
	ST-GIN	0.864	0.846	0.802	0.792	0.900	0.797	0.934	0.694	0.864
	MT-GIN	0.871	0.849	0.828	0.781	0.917	0.804	0.937	0.709	0.871
	ST-RNN	0.798	0.768	0.712	0.678	0.859	0.694	0.866	0.544	0.745
	MT-RNN	0.828	0.804	0.755	0.729	0.879	0.742	0.895	0.613	0.792

objective prediction. For classification tasks, our analysis revealed that GIN outperformed other models, exhibiting the highest AUC and AP values, as illustrated in Figure 4.6 and Figure 4.7, respectively. On the other hand, for regression tasks, Attentive FP demonstrated superior performance, characterized by high values of R^2, and low values of RMSE or MAE, as depicted in Figure 4.8. This observation underscores the effectiveness of graph-based representations in extracting vital molecular information. Interestingly, our analysis also unveiled that RNN performed sub-optimally, despite possessing comparable parameters to the graph-based models (Table 4.3), and exhibited performance similar to our baseline model MLP which utilizes fingerprints. This discrepancy suggests that RNN may lack the capability to extract valuable structural insights from inputs provided in SMILES/SELFIES format, rendering them less proficient in comparison to graph-based models. Training loss and valid loss change during training for both classification task and regression task are shown in Figure 4.9. As we observed previously, Attentive FP consistently outperforms others in regression tasks, as indicated by its lower validation set loss. Conversely, for classification tasks, GIN excels with lower validation losses. In contrast, MLP and RNN models exhibit a different behavior: although their training losses continue to decrease, the validation losses either remain stagnant or, in some cases, increase, suggesting potential overfitting. These results reaffirm the effectiveness of graph-based models in handling these tasks.

Performance Evaluation on MT tasks. Next we evaluated the MT learning for all 4 models. The detailed results of the test set in five binary classification tasks related to metabolism are presented in Table 2.3. Notably, MT training has consistently led to performance improvements across all four models, underscoring the effectiveness of MT learning. This approach facilitates information sharing between similar metabolism tasks by enabling the sharing of model parameters. In particular, MT-GIN stands out as the top-performing model. These findings serve to emphasize

the efficacy of MT-GIN in addressing these classification tasks.

Table 4.5: Uncertainty weight (UW) on multi-task training on 4 tasks simultaneously. Best model performance of each task is shown in bold. N: no UW, Y: UW applied.

Tasks, metrics and models		MT-MLP	MT-AttentiveFP	MT-GIN	MT-RNN
BBB AUC ↑	UW (N)	0.870±0.002	0.881±0.001	0.906±0.001	0.552±0.058
	UW (Y)	0.863±0.002	0.874±0.000	**0.910±0.004**	0.838±0.031
PAMPA AUC ↑	UW (N)	0.757±0.016	0.779±0.007	0.775±0.003	0.497±0.011
	UW (Y)	0.771±0.010	**0.808±0.012**	0.796±0.001	0.748±0.015
Lipophilicity RMSE ↓	UW (N)	1.052±0.031	0.943±0.034	0.917±0.019	1.221±0.005
	UW (Y)	0.874±0.006	**0.632±0.021**	0.699±0.008	0.840±0.037
Solubility RMSE ↓	UW (N)	1.838±0.069	1.306±0.020	1.602±0.015	2.348±0.024
	UW (Y)	1.316±0.025	**1.042±0.012**	1.297±0.013	1.148±0.019

Uncertainty weight strategy[61]. We illustrate the advantages of implementing this dynamic weight strategy in MT learning, employing four distinct tasks, including two classification tasks (BBB and PAMPA), and two regression tasks (lipophilicity and solubility), results shown in Table 4.5. Our findings demonstrate that models trained with this method can effectively learn task-specific weights in a dynamic manner, resulting in superior performance compared to models that rely on uniform weights in the context of MT learning.

Table 4.6: Comparison between current state-of-the-art (SOTA) and our model, MT-Attentive FP with uncertainty weight trained on 3 toxicity tasks simultaneously. Better performance is shown in bold.

data set	SOTA	MT-AttentiveFP (UW)
LD$_{50}$ MAE ↓	0.552±0.009 [52]	**0.445±0.003**
hERG_Karim AUC ↑	-	0.894±0.001
AMES AUC ↑	0.871±0.002 [126]	**0.884±0.004**

Comparison with current SOTA. We trained 3 toxicity tasks simultaneously with uncertainty weight strategy, and findings reveal the exceptional performance of our models, consistently outperforming the existing SOTA, as shown in Table 4.6. In the LD$_{50}$ regression task, our model achieved a substantially lower Mean Absolute

Error (MAE) of 0.445±0.003, outperforming the current SOTA model with an MAE of 0.552±0.009. Additionally, in the AMES classification task, our model exhibited a superior Area Under Curve (AUC) of 0.884±0.004, surpassing the SOTA model with an AUC of 0.871±0.002. These results underscore the remarkable performance enhancements offered by our approach, highlighting its potential to outshine existing methods in critical tasks.

4.4   Conclusion

Traditional methods involving *in vitro* and *in vivo* studies are not only resource-intensive but also time-consuming. To overcome these challenges, we have developed a robust ML framework for the evaluation of ADME-T profiles, addressing the critical need for a comprehensive PK characterization of novel drugs prior to clinical development. Unlike conventional ML approaches that focus on isolated tasks, our method leverages a MT model architecture equipped with a dynamic weight strategy to adaptively allocate weights between tasks. Through extensive experimentation with various ML methods, we have demonstrated the superiority of GNN-based models in achieving SOTA results in specific tasks. This research holds the potential to significantly streamline and economize the process of identifying promising drug candidates with favorable ADME-T profiles, addressing a critical aspect of drug discovery.

4.5   Data availability

Data are available from the corresponding authors with a reasonable request.

4.6 Supporting information

Details of dataset (train, valid, and test split) and description for each dataset. Results of testing sets in ADME-T ST binary classification tasks and regression tasks.

4.7 Acknowledgement

Author Contributions

D.S. directed this study. Y.B. designed the experiment and ran the models. L.Z. performed data analysis and data visualization. All authors have contributed to writing and reviewing the manuscript, and have given approval to the final version.
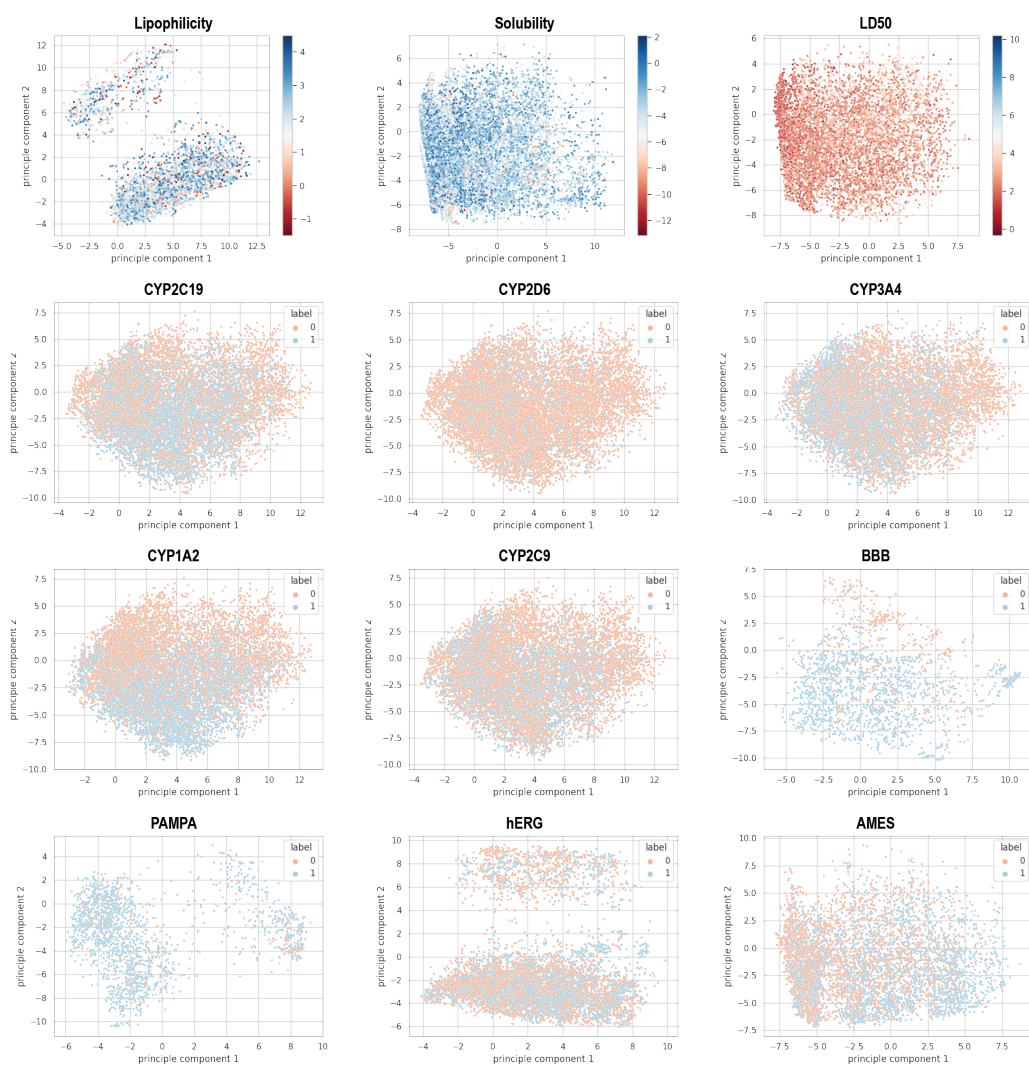
Figure 4.3: Data visualization based on MACCS fingerprint with PCA.

Figure 4.4: Data visualization based on MACCS fingerprint with t-SNE.
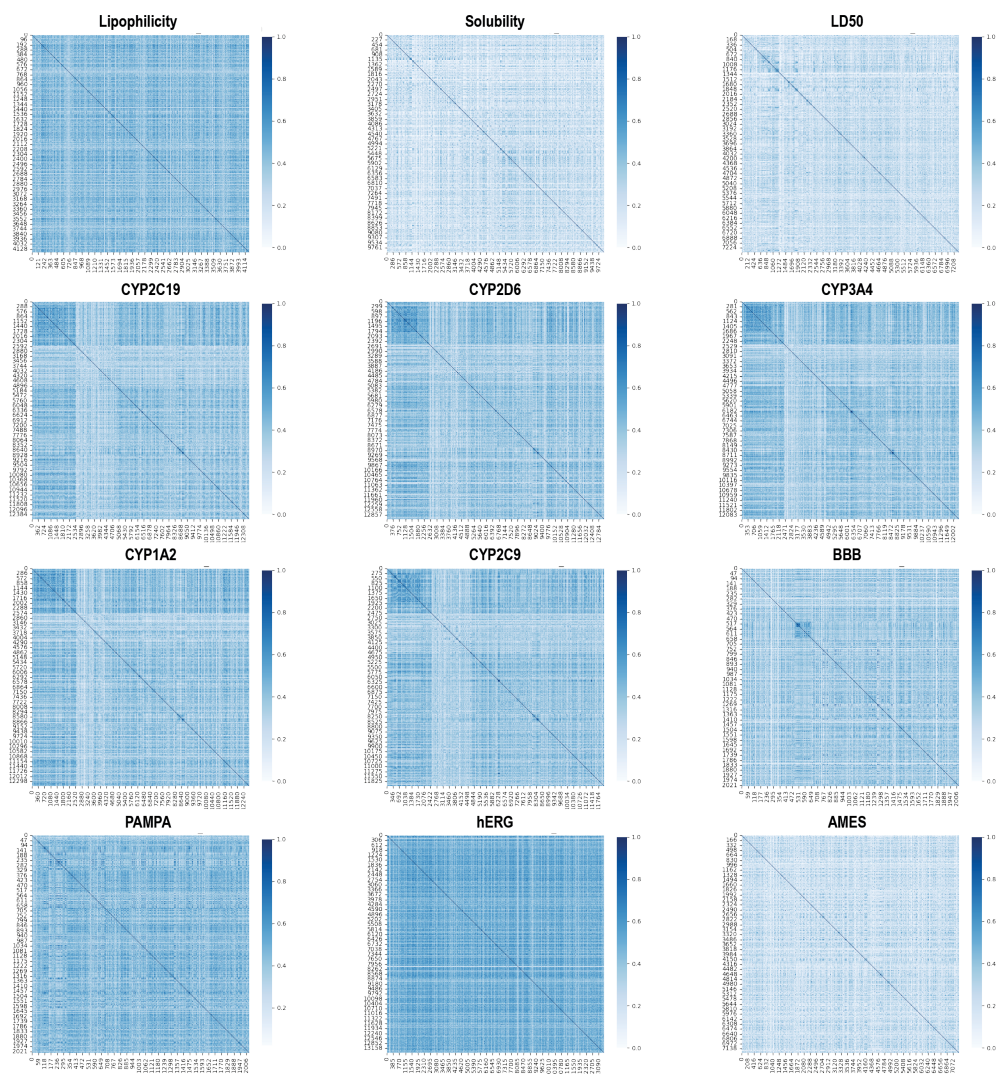
Figure 4.5: Data visualization based on MACCS fingerprint with Tanimoto similarity.
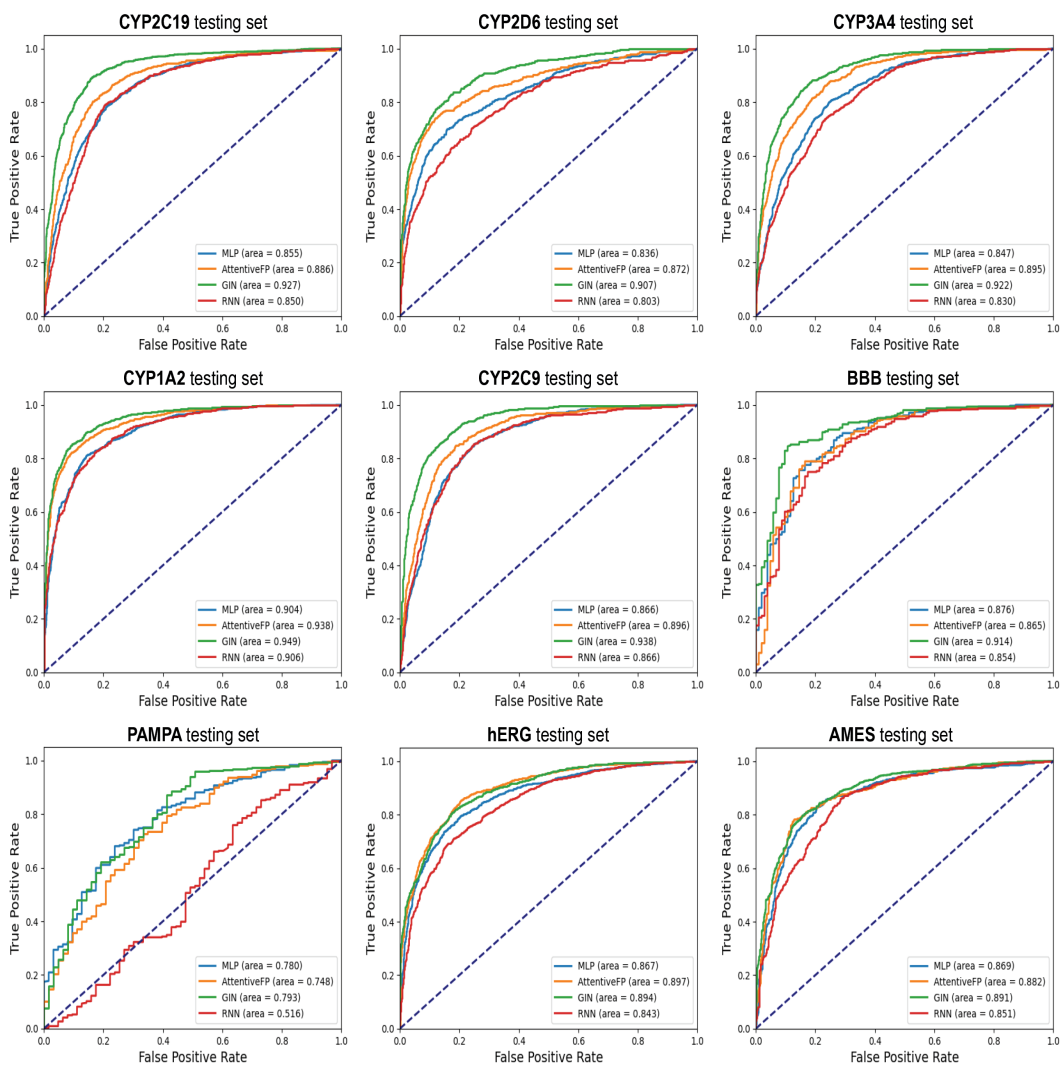
Figure 4.6: AUC curves of 4 models (single task) on classification task test sets.
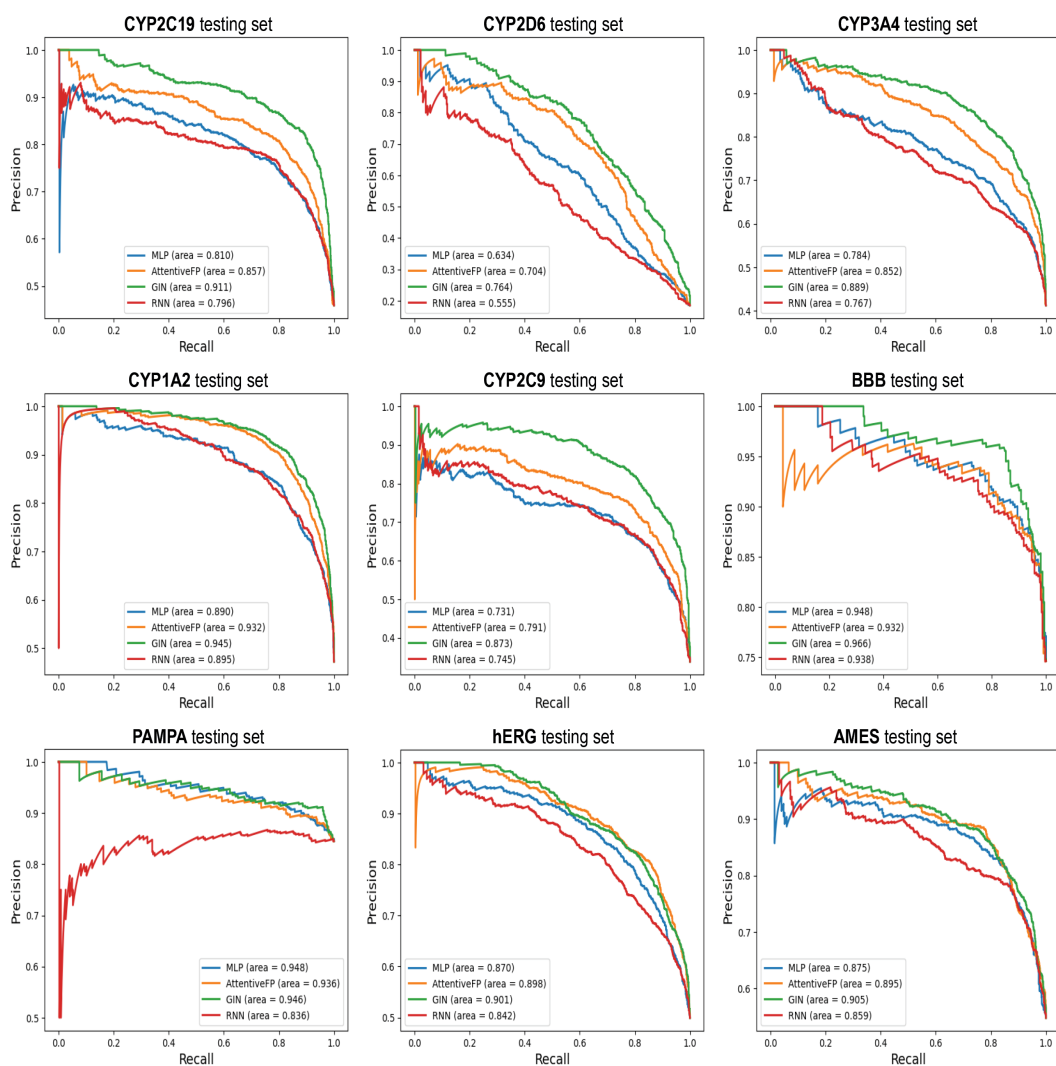
Figure 4.7: AP curves of 4 models(single task) on classification task test sets.

Performance of Lipophilicity testing set

MLP	Attentive FP	GIN	RNN
R2: 0.506 MAE: 0.654 RMSE: 0.855	R2: 0.776 MAE: 0.410 RMSE: 0.575	R2: 0.706 MAE: 0.498 RMSE: 0.660	R2: 0.549 MAE: 0.614 RMSE: 0.817
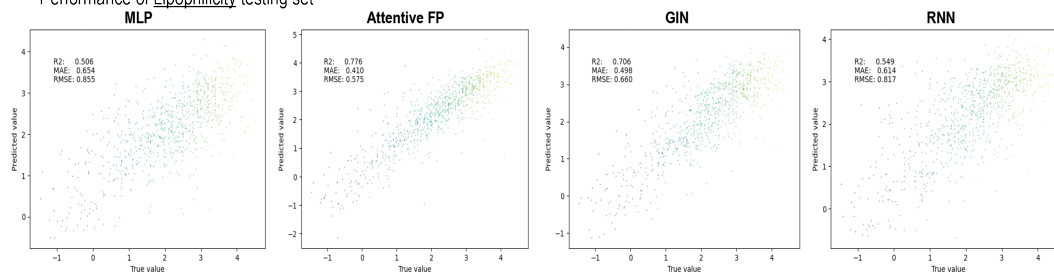
Performance of Solubility testing set

MLP	Attentive FP	GIN	RNN
R2: 0.711 MAE: 0.931 RMSE: 1.252	R2: 0.817 MAE: 0.682 RMSE: 0.997	R2: 0.772 MAE: 0.777 RMSE: 1.112	R2: 0.781 MAE: 0.772 RMSE: 1.089
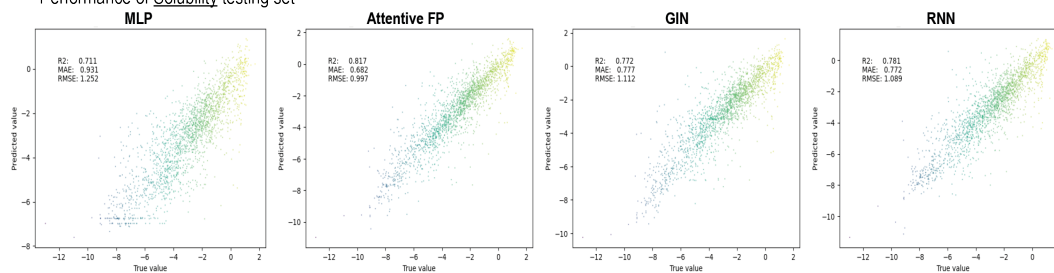
Performance of LD50 testing set

MLP	Attentive FP	GIN	RNN
R2: 0.578 MAE: 0.456 RMSE: 0.614	R2: 0.594 MAE: 0.450 RMSE: 0.602	R2: 0.615 MAE: 0.428 RMSE: 0.586	R2: 0.507 MAE: 0.498 RMSE: 0.664
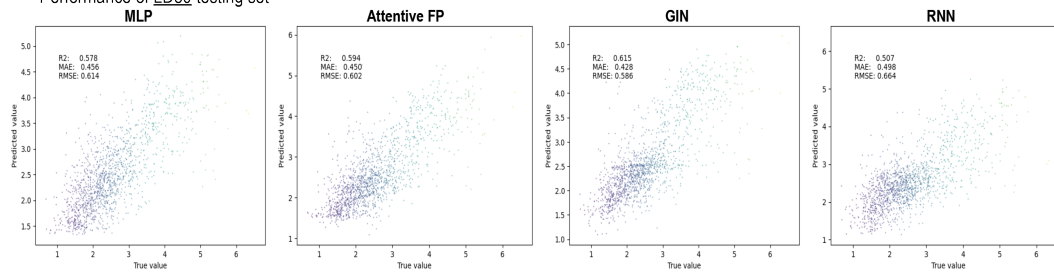
Figure 4.8: ST prediction by different models on regression tasks.
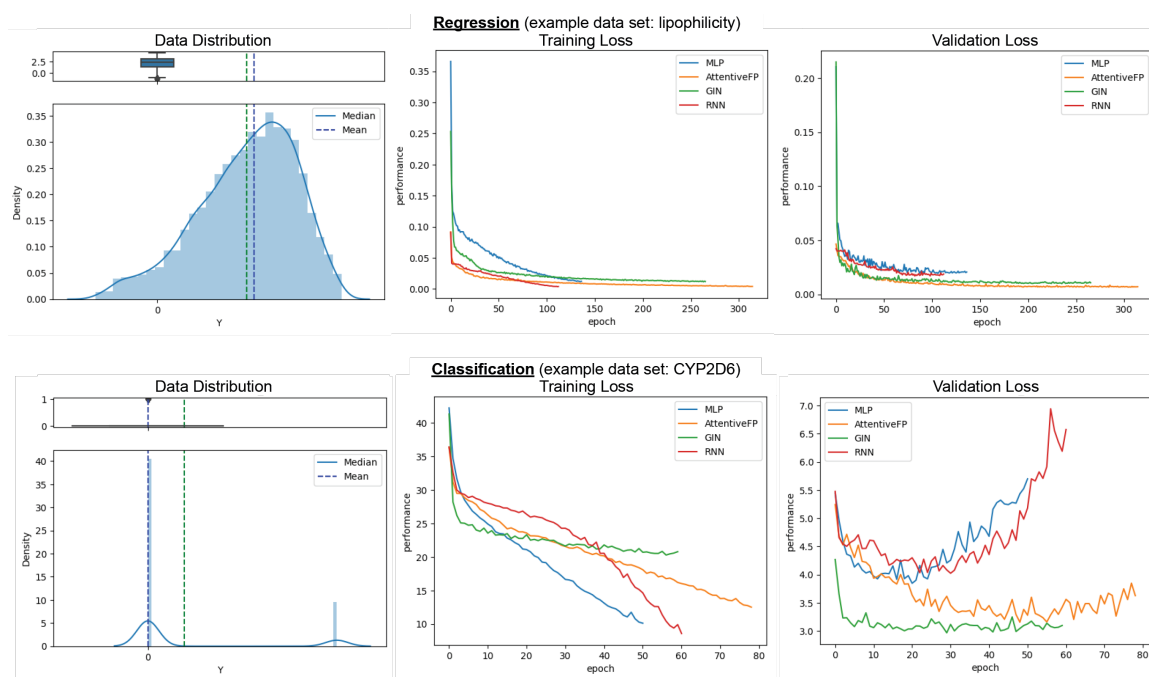
Figure 4.9: Model performance comparison on single task. Regression example: lipophilicity; Classification example: CYP2D6.

Supporting Information

Table S4.1: Data Sets Summary (classification)

data set	type	total	label 0	label 1
CYP2C19_Veith[129]	data	12666	6847	5819
	train	8867	4804	4063
	valid	1266	673	593
	test	2533	1370	1163
CYP2D6_Veith[129]	data	13130	10616	2514
	train	9191	7425	1766
	valid	1313	1053	260
	test	2626	2138	488
CYP3A4_Veith[129]	data	12329	7219	5110
	train	8630	5053	3577
	valid	1233	717	516
	test	2466	1449	1017
CYP1A2_Veith[129]	data	12580	6751	5829
	train	8806	4746	4060
	valid	1258	677	581
	test	2516	1328	1188
CYP2C9_Veith[129]	data	12093	8048	4045

Table S4.1 continued from previous page

data set	type	total	label 0	label 1
	train	8466	5653	2813
	valid	1209	796	413
	test	2418	1599	819
BBB_Martins[82]	data	2030	479	1551
	train	1421	325	1096
	valid	203	51	152
	test	406	103	303
Bioavailability_Ma[80]	data	643	148	495
	train	451	97	354
	valid	64	13	51
	test	128	38	90
Pgp_Broccatelli[9]	data	1219	569	650
	train	853	392	461
	valid	122	56	66
	test	244	121	123
HIA_Hou[49]	data	580	78	502
	train	406	53	353
	valid	58	6	52
	test	116	19	97
PAMPA_NCATS[115]	data	2034	295	1739
	train	1424	199	1225
	valid	203	33	170
	test	407	63	344
hERG_Karim[60]	data	13449	6731	6718
	train	9416	4702	4714

Table S4.1 continued from previous page

data set	type	total	label 0	label 1
	valid	1344	683	661
	test	2689	1346	1343
AMES[147]	data	7282	3307	3975
	train	5098	2338	2760
	valid	728	311	417
	test	1456	658	798

Table S4.2: Data Sets Summary (regression)

data set	total	train	valid	test
Caco2_Wang[132]	910	637	91	182
Lipophilicity_AstraZeneca[137]	4200	2940	420	840
HydrationFreeEnergy_FreeSolv[86]	642	450	64	128
Solubility_AqSolDB[118]	9982	6988	998	1996
LD50_Zhu[154]	7385	5170	738	1477

Table S4.3: Brief description of the data sets

data set	description
Caco-2	The Caco-2 (Cell Effective Permeability) test is an in vitro assay to assess the permeability of drugs and compounds through the intestinal epithelial cell layer. It measures how efficiently substances can cross this cell layer, which is used to approximate their potential for absorption in the human digestive system and bioavailability.
Lipophilicity	Lipophilicity is a property of a chemical compound indicating its ability to dissolve in and interact with lipids or non-polar solvents like fats and oils. A compound's lipophilicity affects its absorption, distribution, and overall pharmacokinetics within the body.
Hydration Free Energy	The hydration free energy of small molecules in water is a thermodynamic measure of the energy change associated with the dissolution of a small molecule in water. It quantifies the stability of the molecule-water interaction and can provide insights into a compound's solubility and its behavior in aqueous solutions.
Solubility	Aqueous solubility refers to the ability of a substance to dissolve in water. It measures the maximum amount of a compound that can dissolve in a given volume of water at a particular temperature and pressure. A compound's solubility can significantly influence its formulation, bioavailability, and efficacy.
LD50	The Acute Toxicity LD50 (Lethal Dose 50) is a standard measure used to determine the dose of a substance that is lethal to 50% of a population of test subjects.

Table S4.3 continued from previous page

data set	description
BBB	The blood-brain barrier (BBB) is a highly selective, semi-permeable barrier that separates the circulating blood from the brain and central nervous system (CNS) tissues. It serves to protect the brain from the passage of potentially harmful substances as well as most foreign drugs.
Bioavailability	Oral bioavailability is defined as "the rate and extent to which the active ingredient or active moiety is absorbed from a drug product and becomes available at the site of action".
CYP1A2 CYP2C9 CYP2C19 CYP2D6 CYP3A4	Cytochrome P450 (CYP) enzymes in the liver and other tissues are responsible for metabolizing a wide array of drugs, xenobiotics, and endogenous compounds in the body to make them more readily excreted. The CYP enzymes are found to metabolize approximately two-thirds of known drugs, with 80% of this can be attributed to the five subtypes, i.e., 1A2, 2C9, 2C19, 2D6, and 3A4.
Pgp	P-glycoprotein (Pgp) is a membrane protein that functions as an efflux pump in cells. Inhibition of Pgp can affect the absorption and distribution of drugs within the body.
HIA	Human Intestinal Absorption (HIA) is a measure of a drug's ability to be absorbed by the human gastrointestinal tract (i.e., pass through the intestinal wall and enter the bloodstream) after oral administration.
PAMPA	The Parallel Artificial Membrane Permeability Assay (PAMPA) is a high-throughput, in vitro method to predict the passive permeability of compounds through biological membranes.

Table S4.3 continued from previous page

data set	description
hERG	The hERG gene codes for a potassium ion channel in the human heart. This dataset includes information on compounds and their interactions with the hERG channel, especially concerning their potential to cause cardiac arrhythmias or other adverse effects.
AMES	The Ames mutagenicity test is a bacterial assay that assesses the mutagenic potential of chemical compounds by measuring their ability to induce mutations in specific strains of bacteria.

Table S4.4: Results of test set in ADME-T ST binary classification tasks. SO, single-objective.

task	model	acc	w_acc	SE	recall	SP	F1	AUC	MCC	AP
BBB	MLP MACCS (SO)	0.850	0.749	0.860	0.954	0.544	0.905	0.876	0.573	0.948
	AttentiveFP (SO)	0.867	0.794	0.888	0.941	0.647	0.914	0.901	0.629	0.948
	GIN (SO)	0.860	0.781	0.880	0.941	0.621	0.909	0.914	0.609	0.966
	RNN (SO)	0.835	0.720	0.845	0.954	0.485	0.896	0.854	0.524	0.938
F	MLP MACCS (SO)	0.719	0.564	0.733	0.944	0.184	0.825	0.665	0.202	0.784
	AttentiveFP (SO)	0.742	0.589	0.744	0.967	0.211	0.841	0.727	0.289	0.875
	GIN (SO)	0.711	0.574	0.739	0.911	0.237	0.816	0.685	0.199	0.811
	RNN (SO)	0.719	0.557	0.729	0.956	0.158	0.827	0.634	0.193	0.796
Pgp	MLP MACCS (SO)	0.791	0.791	0.777	0.821	0.760	0.798	0.856	0.583	0.886
	AttentiveFP (SO)	0.836	0.834	0.838	0.864	0.804	0.851	0.911	0.669	0.929
	GIN (SO)	0.820	0.820	0.832	0.805	0.835	0.818	0.882	0.640	0.904
	RNN (SO)	0.852	0.853	0.872	0.829	0.876	0.850	0.900	0.706	0.894
HIA	MLP MACCS (SO)	0.948	0.884	0.960	0.979	0.789	0.969	0.974	0.805	0.995
	AttentiveFP (SO)	0.931	0.814	0.962	0.962	0.667	0.962	0.981	0.628	0.998
	GIN (SO)	0.931	0.874	0.959	0.959	0.789	0.959	0.982	0.748	0.997
	RNN (SO)	0.879	0.780	0.928	0.928	0.632	0.928	0.954	0.559	0.991
PAMPA	MLP MACCS (SO)	0.845	0.513	0.849	0.994	0.032	0.916	0.780	0.095	0.948
	AttentiveFP (SO)	0.860	0.599	0.873	0.977	0.222	0.922	0.748	0.318	0.936
	GIN (SO)	0.875	0.667	0.893	0.968	0.365	0.929	0.793	0.435	0.946
	RNN (SO)	0.845	0.500	0.845	1.000	0.000	0.916	0.516	N/A	0.836
hERG	MLP MACCS (SO)	0.791	0.791	0.806	0.765	0.816	0.785	0.867	0.582	0.870
	AttentiveFP (SO)	0.809	0.809	0.827	0.780	0.837	0.803	0.897	0.619	0.898
	GIN (SO)	0.816	0.816	0.822	0.805	0.826	0.813	0.894	0.631	0.901
	RNN (SO)	0.765	0.765	0.796	0.711	0.819	0.751	0.843	0.533	0.842
AMES	MLP MACCS (SO)	0.809	0.804	0.804	0.861	0.746	0.832	0.869	0.614	0.875
	AttentiveFP (SO)	0.830	0.822	0.835	0.875	0.768	0.855	0.884	0.650	0.900
	GIN (SO)	0.812	0.809	0.817	0.848	0.769	0.832	0.891	0.621	0.905
	RNN (SO)	0.790	0.783	0.785	0.850	0.717	0.816	0.851	0.574	0.859

Table S4.5: Results of test set in ADME-T ST regression tasks.

task	model	MAE ↓	RMSE ↓	R^2 ↑
Caco2	MLP MACCS (SO)	0.385	0.493	0.618
	AttentiveFP (SO)	0.331	0.404	0.743
	GIN (SO)	0.384	0.492	0.619
	RNN (SO)	0.485	0.621	0.393
Lipophilicity	MLP MACCS (SO)	0.652	0.841	0.521
	AttentiveFP (SO)	0.353	0.489	0.834
	GIN (SO)	0.489	0.647	0.717
	RNN (SO)	0.614	0.817	0.549
HydrationFreeEnergy	MLP MACCS (SO)	1.376	1.774	0.808
	AttentiveFP (SO)	0.942	1.359	0.843
	GIN (SO)	1.274	1.578	0.848
	RNN (SO)	1.348	1.678	0.828
Solubility	MLP MACCS (SO)	0.931	1.252	0.711
	AttentiveFP (SO)	0.682	0.997	0.817
	GIN (SO)	0.777	1.112	0.772
	RNN (SO)	0.772	1.089	0.781
LD_{50}	MLP MACCS (SO)	0.456	0.614	0.578
	AttentiveFP (SO)	0.450	0.602	0.594
	GIN (SO)	0.428	0.586	0.615
	RNN (SO)	0.498	0.664	0.507

CHAPTER V

Deep Reinforcement Learning-Guided de novo Design of Janus Kinase Inhibitors

Abstract

We present a novel approach for *de novo* drug design by ML. This method leverages a deep neural network, employing VAE architecture with a MLP classifier, and REINFORCE algorithm. The VAE's encoder converts discrete molecular representations into continuous vectors, the decoder reverses this process to reconstruct discrete molecular forms, the classifier arranges compounds with similar properties in the latent space, and the REINFORCE algorithm further expands the latent manifold to uncover new inhibitors guided by reward functions. Initial results indicate the successful generation of pan-JAK inhibitors, with future work focusing on generating JAK-specific inhibitors, subjecting them to selection based on STAR and ADME-T profiles, and conducting *in vitro* and *in vivo* experiments to validate their efficacy.

5.1 Introduction

The process of drug discovery is well-known for its extended timelines, exorbitant costs and labor-intensive nature. However, the introduction of ML into this field has proven highly beneficial. ML offers the ability to accelerate drug discovery by

predicting essential properties of compounds, such as their efficacy, safety, and PK, thereby reducing the need for extensive and costly experimental work. Furthermore, deep generative models have the capability to create compounds with specific properties, such as activity against a particular target, making them particularly suitable for identifying potential drug candidates.

JAKs, a family of enzymes integral to cytokine receptor intracellular signaling, are implicated in malignancies and autoimmune disorders [93, 105, 131, 51]. Although several JAK inhibitors have gained approval for treating these conditions, they often exhibit common side effects, likely due to their broad-spectrum inhibition of various JAK isoforms [131]. Developing isoform-specific inhibitors typically involves time-consuming high-throughput screening and lead compound optimization, which is not cost-effective. Furthermore, most available JAK-specific inhibitors are non-isoform-specific, posing a challenge for lead compound optimization.

To tackle these challenges, ML can play a pivotal role in predicting whether a compound possesses JAK inhibition properties even before undergoing synthesis. CoGT method from Chapter II could be utilized to predict JAK inhibition based on chemical structures[10]. However, this approach may not necessarily result in the discovery of the most potent JAK-specific inhibitors. It requires the generation of numerous compounds for prediction, with no guarantee of finding one with the desired JAK inhibition; and even if it could find one, it may not be the most potent candidate.

Therefore, we introduce a generative approach employing pretrained VAE and REINFORCE, with the goal of designing novel JAK inhibitors *de novo*, directly providing us with structures that meet our requirements[65, 42, 153, 142]. Our findings demonstrate that a pretrained VAE effectively clusters compounds with similar characteristics in the latent space. Preliminary data indicates the capability of this method to generate diverse structures for pan-JAK inhibitors. Future work will leverage this approach to generate JAK-specific inhibitors, followed by ADME-T prediction, as

well as *in vitro* and *in vivo* experiments for validation.

5.2 Methods

Data

Data sets and data preparation. Chemical structures were extracted from ZINC data set[53, 120]. Similar to Chapter IV, we used normalizer, salt remover, reionizer and uncharger to clean SMILES representation before training.

Molecular representation. In this investigation, we employed RNN as layers within the VAE framework, utilizing SMILES or SELFIES strings to represent compounds. To facilitate the conversion of string into a suitable input format, we adopted a one-hot encoding approach. Initially, after cleaning all the molecules in dataset, we conducted a search to identify all SMILES/SELFIES present in our dataset, thereby creating an alphabet denoted as $\mathcal{A} = \{a_1, \cdots, a_i, \cdots, a_n\}$ with each a_i representing a unique token within drug representation. Additionally we incorporated a padding token to pad the sequence of drug strings to the same length, a "begin of string" token to indicate the beginning of the drug string, and an "end of string" token to indicate its end. Subsequently, the drug representation strings were transformed into a vector of scalar values, corresponding to the position of each token within the alphabet.

Model architecture

VAE [65]. VAE is a powerful stochastic variational inference and learning algorithm that excels at efficient inference and learning within directed probabilistic models. It demonstrates remarkable scalability to handle large datasets and can effectively handle scenarios involving continuous latent variables with intractable posterior distributions. In this study, consider the compound dataset without labels as $\mathbf{X} = \{\mathbf{x}^{(i)}\}_{i=1}^{N}$ of N i.i.d. samples of variable $\mathbf{x}$ which encodes drug structure infor-

mation, VAE helps to find a latent vector $\mathbf{z}$ from some prior distribution $p_{\boldsymbol{\theta}}(\mathbf{z})$ so that we can generate countless drug representation vector $\mathbf{x}$ from conditional distribution $p_{\boldsymbol{\theta}}(\mathbf{x}|\mathbf{z})$. This method offers the advantage of not relying on compound labels (e.g., whether a compound exhibits JAK inhibition) during the distribution learning process. Instead, it focuses on learning a more effective drug "representation" in the latent space. VAE lets the prior over the latent variables be $p_{\boldsymbol{\theta}}(\mathbf{z}) = \mathcal{N}(\mathbf{z}; \mathbf{0}, \mathbf{I})$. Since the true posterior $p_{\boldsymbol{\theta}}(\mathbf{z}|\mathbf{x})$ is intractable, Kingma et al. used multivariate Gaussian with a diagonal covariance $q_{\boldsymbol{\phi}}(\mathbf{z}|\mathbf{x}) = \mathcal{N}(\mathbf{z}; \boldsymbol{\mu}, \boldsymbol{\Sigma})$ to approximate. The loss function is defined as minimizing Kullback-Leibler (KL) divergence between $q_{\boldsymbol{\phi}}(\mathbf{z}|\mathbf{x})$ and $p_{\boldsymbol{\theta}}(\mathbf{z})$, and the reconstruction loss. The training process for VAE is thus minimizing the loss:

$$\min \text{Loss}(\mathbf{x}; \boldsymbol{\theta}, \boldsymbol{\phi}) = \min \left[\text{KL}(q_{\boldsymbol{\phi}}(\mathbf{z}|\mathbf{x})||p_{\boldsymbol{\theta}}(\mathbf{z})) - \frac{1}{L} \sum_{l=1}^{L} \log p_{\boldsymbol{\theta}}(\mathbf{x}|\mathbf{z}^{(l)}) \right] \tag{5.1}$$

Detailed proof derived from Kingma et al. for VAE is provided in Appendix A [65].

Pretraining strategy [103, 69]. The primary objective of small molecule drug discovery is to identify compounds capable of modulating a specific biological target and inducing a beneficial therapeutic response. Unfortunately, this process is exceptionally challenging and often unsuccessful due to the vast search space of potential drug-like molecules, estimated to range from 10^{23} to 10^{60} [99, 31, 6]. Compounding this challenge is the limited capacity of experimental medicinal chemistry to provide labels for only a tiny subset of compounds within this vast space. Consequently, there is considerable interest in leveraging supervised machine learning algorithms, trained on available data $\mathcal{D} = \{(\mathbf{x}_i, \mathbf{y}_i)|\mathbf{x}_i \in \mathcal{X}', \mathbf{y}_i \in \mathcal{Y}\}_{i=1}^{N}$, to predict the properties of compounds. Given the constraints of dataset size, we adopted a pretraining strategy to incorporate structural information from compounds without labels. Specifically, in the case of VAE, we pretrained it on the ZINC dataset to learn structural information and simultaneously predict compounds' Quantitative estimate of drug-likeness (QED)

and Synthetic Accessibility (SA) using an MLP to organize the latent space (Figure 5.1). The training process is defined as minimizing VAE loss (KL divergence and reconstruction loss), and the classification loss for $\mathbf{x}$ having labels $\mathbf{y}$.

$$\text{Loss}(\mathbf{x}, \mathbf{y}; \boldsymbol{\theta}, \phi, \mathbf{c}) = \text{KL}(q_{\phi}(\mathbf{z}|\mathbf{x})||p_{\boldsymbol{\theta}}(\mathbf{z})) - \frac{1}{L}\sum_{l=1}^{L}\log p_{\boldsymbol{\theta}}(\mathbf{x}|\mathbf{z}^{(l)}) - \frac{\alpha}{N}\sum_{i=1}^{N}\log p_{\mathbf{c}}(\mathbf{y}|\mathbf{z}^{(i)})$$

$$(5.2)$$

where $\mathbf{c}$ represents trainable parameters in classifier, and α is the weight for classification loss. Subsequently, we applied the pretrained VAE to downstream tasks JAK inhibitor generation, and focus on dataset with specific labels (dataset containing information about whether the compound is a JAK inhibitor or not).
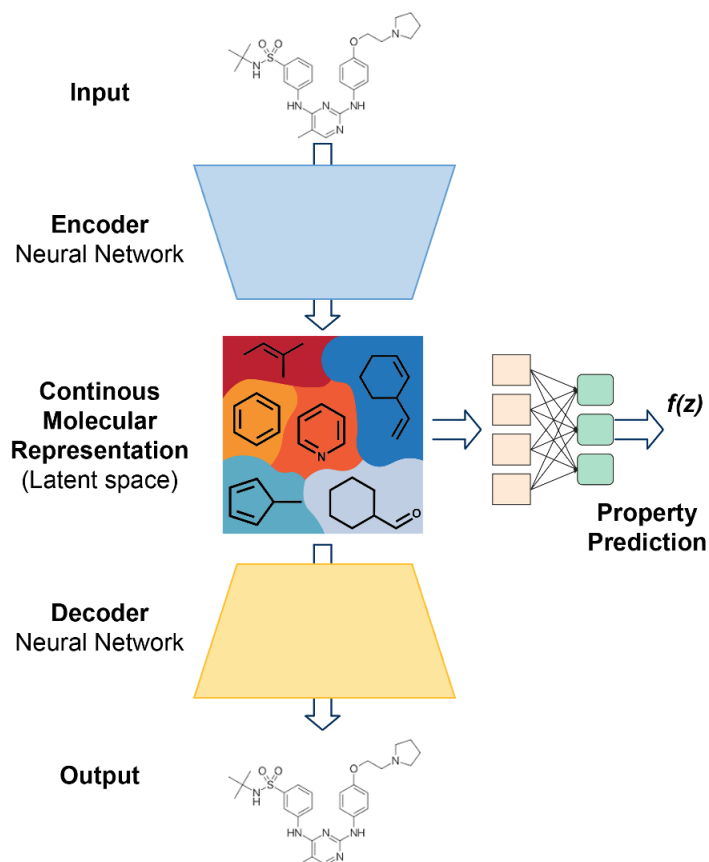
Figure 5.1: Architecture of model with pretraining strategy. VAE with a MLP classifier

REINFORCE algorithm [142]. It is a fundamental Reinforcement Learning (RL) technique used to optimize decision-making processes in an environment where actions lead to outcomes with associated rewards or penalties. By iteratively adjusting actions based on the observed rewards, REINFORCE allows a model to learn to make better decisions over time. This algorithm is widely applied in various domains, including robotics, game playing and NLP, and is particularly valuable for tasks where exploration and adaptation are crucial for achieving optimal results[97, 24]. REINFORCE algorithm is a powerful tool in the realm of drug discovery, particularly in the context of searching the latent space for *de novo* drug design [153, 84]. This algorithm is adept at optimizing and generating novel compounds by iteratively exploring and updating the latent manifold. After organizing the latent space by minimizing the loss defined in Equation 5.2, REINFORCE algorithm is applied to find the optimal JAK inhibitor with optimizing the "reward" function using CoGT:

$$R(\mathbf{z}) = \mathbb{E}_{\mathbf{x} \sim p_{\boldsymbol{\theta}}(\mathbf{x}|\mathbf{z})} \sum_i w_i R_{JAK_i}(\mathbf{x}) \tag{5.3}$$

in which R_{JAK_i} is the estimated probability of a compound $\mathbf{x}$ being a JAK_i inhibitor predicted by CoGT, $i \in \{1, 2, 3, 4\}$. We have the flexibility to influence the generation of either pan-JAK inhibitors or JAK-specific inhibitors by adjusting the weights, denoted as w_i, associated with each JAK_i. When all w_i values are greater than 0, it encourages the generation of pan-JAK inhibitors as the reward function "reward" all JAK inhibition. Conversely, if $w_i > 0$ while $w_j < 0$ for all other JAKs, it promotes the generation of JAK_i specific inhibitors as the reward function "penalizes" other JAK inhibition yet favors JAK_i inhibitor generation. Detailed proof for REINFORCE algorithm is provided in Appendix B.

5.3 Results and Discussion

Drug generation by VAE
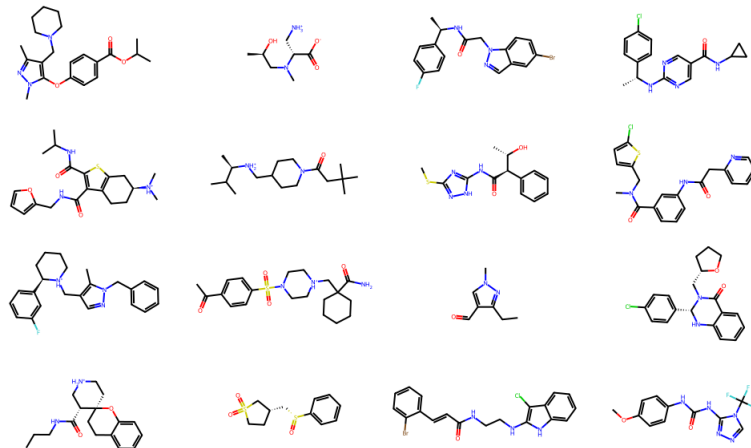
Figure 5.2: Drug generation by sampling $\mathbf{x} \sim \mathrm{p}_\theta(\mathbf{x}|\mathbf{z}), \mathbf{z} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$

We trained VAE on ZINC dataset and results show that it enables drug generation by mapping random latent vectors through the trained decoder, producing molecular structures. The sampling from distribution is defined as $\mathbf{x} \sim \mathrm{p}_\theta(\mathbf{x}|\mathbf{z})$, where $\mathbf{z}$ is sampled from a Gaussian distribution $\mathcal{N}(\mathbf{0}, \mathbf{I})$. Some generated structures are shown in Figure 5.2. The validity of generated compounds from VAE, as determined by RDKit filters, is indeed impressive, with a high validity rate of 91.967%. Additionally, a significant percentage of the generated compounds are unique, also at a rate of 91.967%. This underscores the effectiveness of the VAE-based approach in producing chemically valid and diverse molecules for drug discovery and design purposes.

Latent space visualization

The VAE was extended with a MLP classifier capable of simultaneously predicting property values based on the latent representations of molecules. This approach allows for the organization of molecules within the latent space according to their property values, as demonstrated in Figure 5.3. The two-dimensional representation

of the latent space, achieved through PCA, reveals a clear gradient in the distribution of molecules corresponding to their property values after training 29 epochs. This innovative approach highlights the ability of VAEs, when jointly trained with property prediction, to organize molecules in the latent space based on their desired properties [42].

Pan-JAK inhibitor generation

To assess the effectiveness of our approach, we employed pretrained VAE in combination with the REINFORCE algorithm and CoGT to generate a set of pan-JAK inhibitors. Our method successfully produced a diverse range of compounds, as indicated by the low Tanimoto similarity scores, reflecting the model's ability to generate chemically distinct molecules (Figure 5.4A). Furthermore, CoGT's predictions for JAK inhibition confirmed that the generated compounds indeed exhibit JAK inhibitory activity, with a broad-spectrum effect on four JAK isoforms, results shown in Figure 5.4B. The chemical structures of 36 such compounds are depicted in Figure 5.4C. Our future work will involve refining the reward function to enable the generation of JAK-specific inhibitors, thereby fine-tuning the method for more targeted drug design endeavors.

5.4 Conclusion

In this research, we delve into the theoretical foundations that support the effective utilization of Deep Reinforcement Learning (DRL) for the purpose of *de novo* drug design, particularly focusing on JAK inhibitors. This study provides insights into the rationale behind our choice to employ REINFORCE algorithm in conjunction with the pretrained VAE. Our methodology commences with the organization of latent space using metrics including QED and SA, facilitated by training a MLP classifier using the latent vector representations. Subsequently, we fine-tune this

pretrained model for the downstream task of generating JAK inhibitors. Leveraging the REINFORCE algorithm we strive to identify optimal JAK inhibitors based on a predefined reward function derived from ML predictions. Our initial experiments have demonstrated promise by producing a diverse array of pan-JAK inhibitors, as predicted by CoGT. Our future endeavors will be dedicated to the generation of JAK-specific inhibitors. Following the generation of molecular structures, our approach will extend to employing models for ADME-T predictions from Chapter IV, allowing us to assess the PK properties and toxicity profiles of the selected compounds. Subsequent phases will involve extensive *in vitro* and *in vivo* studies to validate the suitability of these compounds for further drug development endeavors.
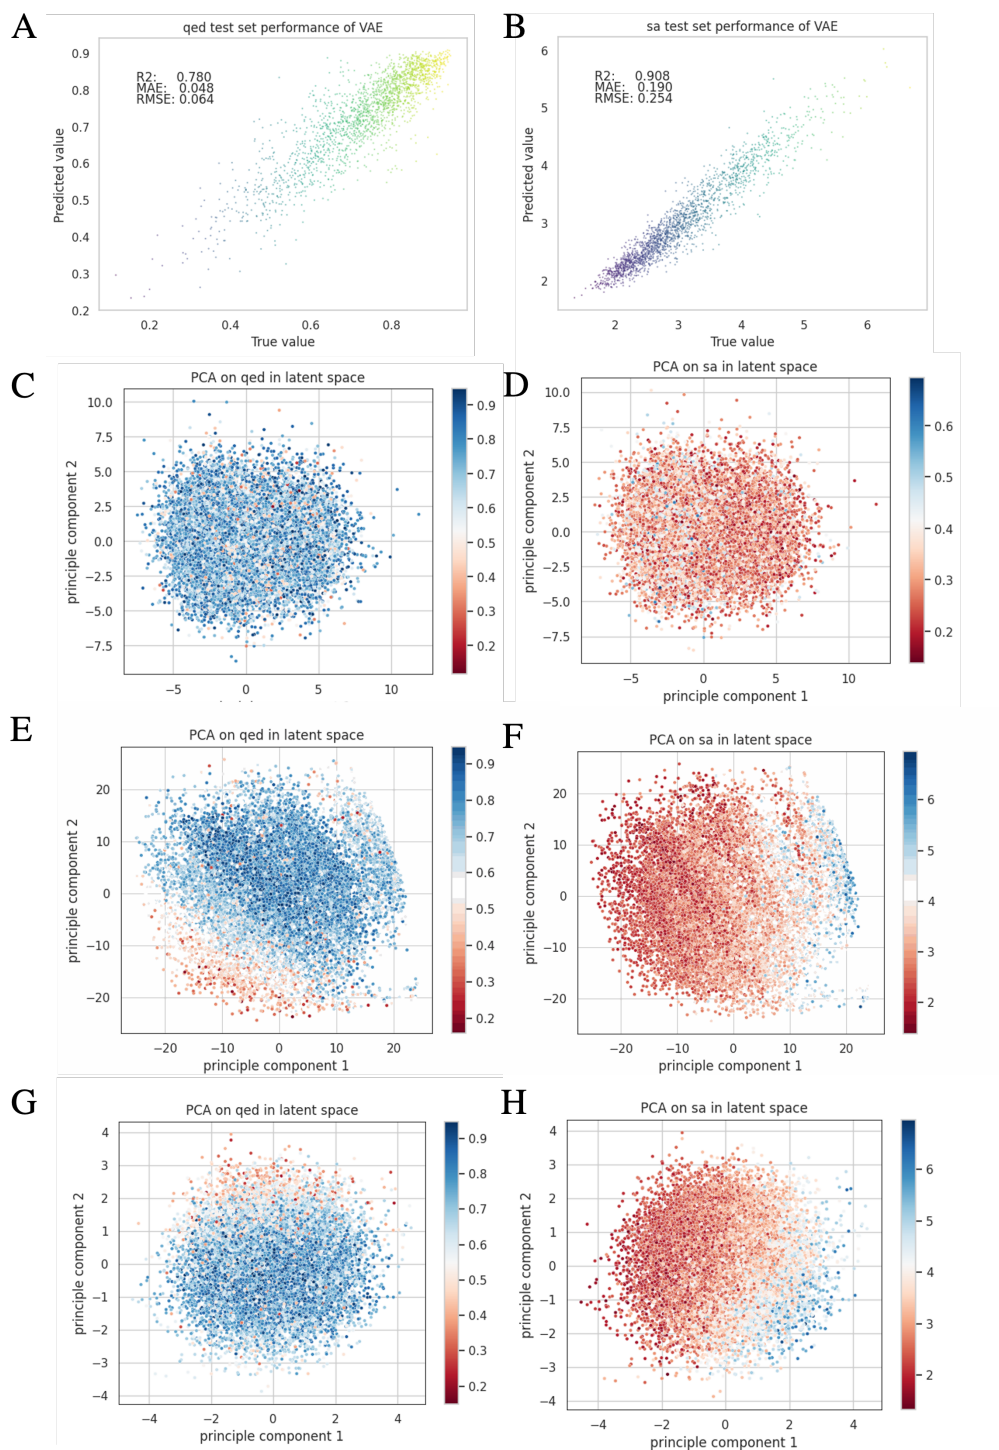
Figure 5.3: Latent space visualization during training. (A, B), predicted QED, SA vs real QED, SA values, respectively. (C, D), before training VAE, PCA analysis of latent space with labels QED and SA, respectively. (E, F), PCA analysis of latent space after 1 epoch for QED and SA, respectively. (G, H), PCA analysis of latent space after 29 epochs for QED and SA, respectively.
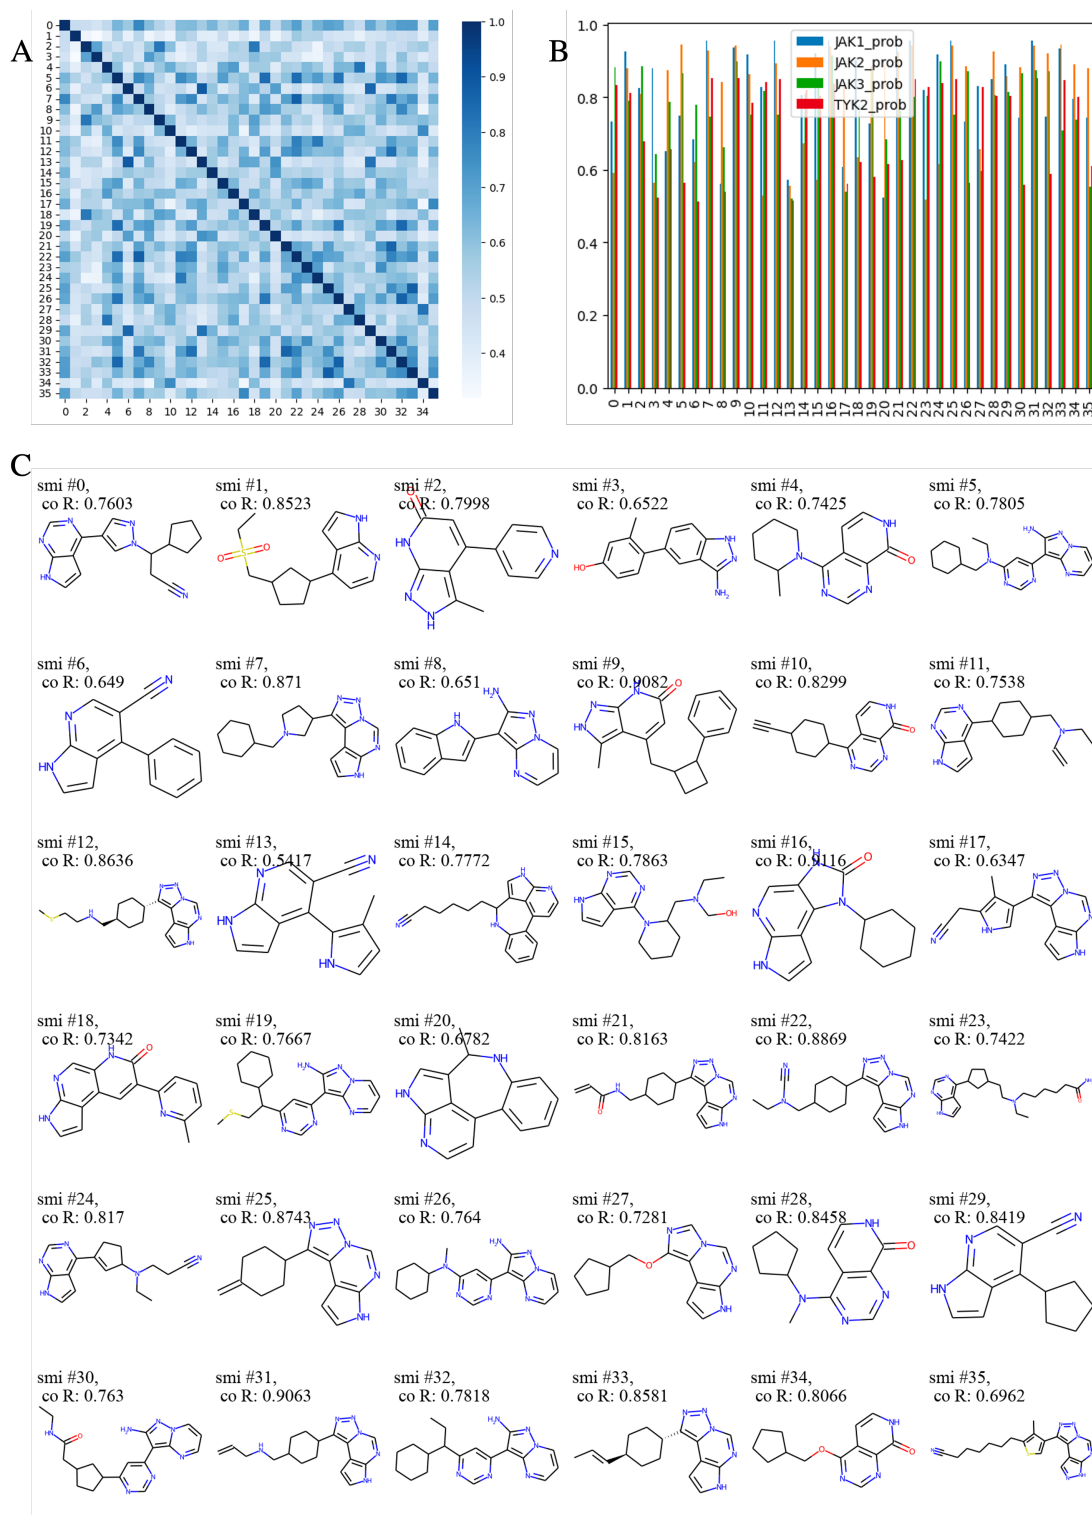
Figure 5.4: Pan-JAK inhibitor generation. 36 example structures were given. A, Tanimoto similarity of 36 compounds. B, JAK inhibition predicted by CoGT; C, structures for 36 generated compounds, co R: reward values for a certain compound, the higher the value, the better the compound satisfies the criteria outlined in our reward function.

CHAPTER VI

Conclusion

Conventional drug discovery is resource-intensive and prone to high attrition rates, mainly due to issues like lack of efficacy, safety concerns, and suboptimal PK. Meanwhile, the advent of ML is revolutionizing this field by enabling precise prediction of DTI, toxicity, and PK parameters, reducing the need for costly and time-consuming experiments. Besides, ML's one-shot generation capability offers a novel approach to drug design. In response to the challenges in drug discovery, we harnessed the power of ML for compound optimization and design. First, we utilized ML and STAR to design a GI JAK inhibitor tailored for UC. Utilizing the STAR system, we designed a class III candidate, MMT3-72, with high GI tissue selectivity and moderate potency, reducing the demand for strict JAK isoform specificity. Concurrently, we introduced CoGT, a novel ML-based method that excels in distinguishing JAK inhibitors from non-inhibitors. CoGT combines conventional ML models with a graph-based model and a pretrained RoBERTa model, achieving SOTA performance in JAK inhibition prediction. Using CoGT, we identified the metabolite of MMT3-72, MMT-72-M2, as a potent JAK inhibitor. Subsequent experiments confirmed these predictions, highlighting MMT3-72's limited JAK inhibition and MMT3-72-M2's effectiveness against JAK1/2 and TYK2. PK studies revealed that MMT3-72 accumulated in the GI lumen, while the metabolite preferentially accumulated in colon tissue with minimal

plasma exposure. *In vivo* studies showed that MMT3-72 exhibited superior efficacy and could reduce p-signal transducers and activators of transcription (STAT)3 levels in dextran sulfate sodium (DSS)-induced colitis, showcasing its site-specific action and minimized systemic toxicity. Next, we established a comprehensive ML framework for evaluating ADME-T profiles in a cost-effective manner, aiming to address the PK and toxicity issues in drug discovery. This ML approach enables simultaneous prediction of multiple ADME-T properties aided by GNN-based models, streamlining drug candidate identification. Furthermore, we utilized ML to design JAK-specific or pan-JAK inhibitors with the potential to target specific isoforms, aiming to mitigate toxicity concerns. This ongoing work involves *in silico* drug design using VAE and RL techniques. Selection of potent drugs will be guided based on the STAR system's criteria and the predicted ADME-T profiles, followed by experimental assays to confirm efficacy.

In essence, our research underscores ML's profound potential in expediting drug discovery and advancing pharmaceutical research. Chapter II describes CoGT, an ensemble ML method for classification tasks on JAK inhibition. Chapter III demonstrates the power of integrating ML with STAR to design GI locally activating JAK inhibitor for treatment of UC. Chapter IV addresses the issue of PK and toxicity, providing a tool for ADME-T property prediction. Chapter V is an ongoing work, describing *de novo* JAK inhibitor design by ML.

As described in Chapter II (Figure 6.1), *CoGT: Ensemble machine learning method and its application on JAK inhibitor discovery*, we introduced an ensemble model named CoGT, designed to harness the collective power of multiple ML models to enhance predictive accuracy in the context of DTI for four JAK isoforms. We meticulously compiled a comprehensive dataset of JAK inhibitors and non-inhibitors. Then we conducted a comparative analysis of various ML methods for predicting JAK inhibition. These methods encompassed a graph-based model (Relational Graph
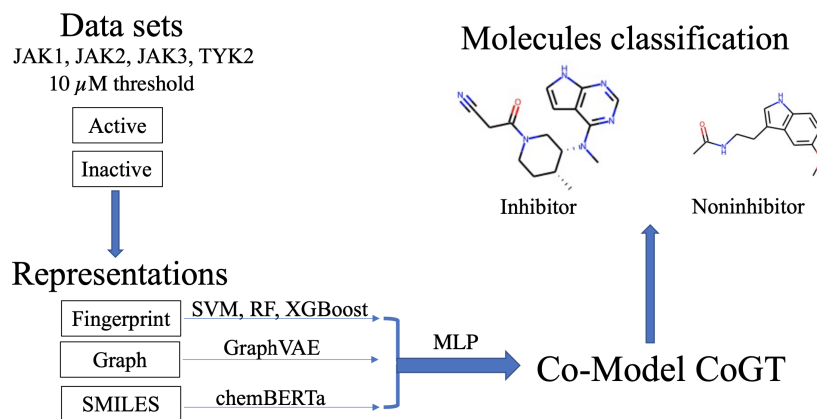
Figure 6.1: Summary of Chapter II

Convolution Network (RGCN) applied GraphVAE), a pre-trained RoBERTa model (chemBERTa), and conventional ML models. Our findings indicated that graph-based model excelled in effectively extracting structural information from JAK inhibitors, surpassing conventional ML approaches. Additionally, the large pre-trained transformer-based model, chemBERTa, demonstrated effectiveness in predicting the chemical properties of these JAK inhibitory structures. Even traditional models like SVM, RF and XGBoost exhibited strong performance, despite their relatively lower computational requirements. By leveraging the strengths of these diverse models, our ensemble model CoGT achieved the highest accuracy in predicting DTIs for JAK inhibitors.

As described in Chapter III (Figure 6.2), *a GI locally-activating JAK inhibitor to treat UC*, we combined ML, STAR and wet lab experiments to develop JAK inhibitor tailored for UC treatment. Our innovative approach led to the creation of STAR class III JAK inhibitor MMT3-72, engineered to achieve potent local exposure within GI tissues. Notably, molecular studies indicated that MMT3-72 exhibited limited absorption potential characterized by a low QED score, and was anticipated to be retained within the GI tract. Two ML models predicted MMT3-72 was not potent yet the major metabolite MMT3-72-M2 acted as JAK1, JAK2, and TYK2 inhibitor. Leveraging
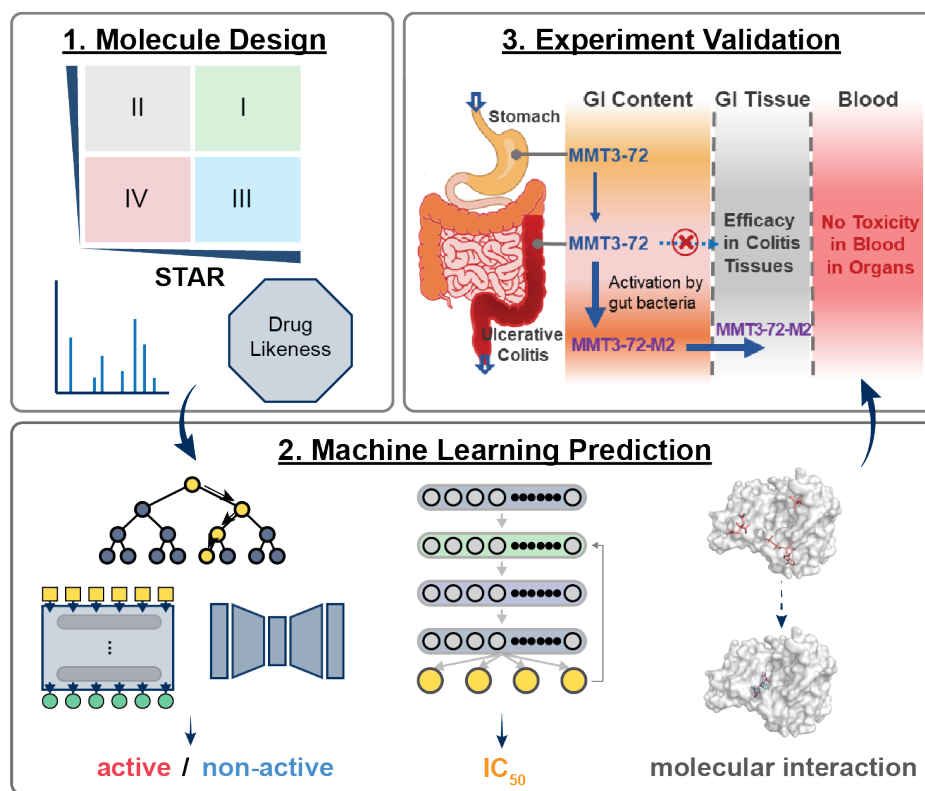
Figure 6.2: Summary of Chapter III

these ML predictions, we synthesized only these two compounds and subsequent *in vitro* kinase assays confirmed the ML prediction. PK investigations revealed that MMT3-72 predominantly accumulated within the GI tract lumen ($> 50,000$ ng/g) while remaining undetected in GI tissues or plasma. Notably, MMT3-72 underwent local activation in GI and released the active form MMT3-72-M2. The latter exhibited substantial accumulation in colon contents ($C_{max} > 50,000$ ng/g) and colon tissues ($C_{max} > 1500$ ng/mL) while maintaining minimal concentrations in the plasma ($C_{max} < 8$ ng/mL). Oral administration of MMT3-72 (at 5 or 10 mg/kg, respectively) yielded superior efficacy compared to tofacitinib in a DSS-induced colitis model. This superiority was evident in various parameters, including disease activity index (DAI) score, colon length, and histological assessments of colon tissues. Moreover, MMT3-72 administration effectively suppressed STAT3 phosphorylation within colitis tissues.

As described in Chapter IV (Figure 6.3), *Machine Learning-Enhanced Prediction*
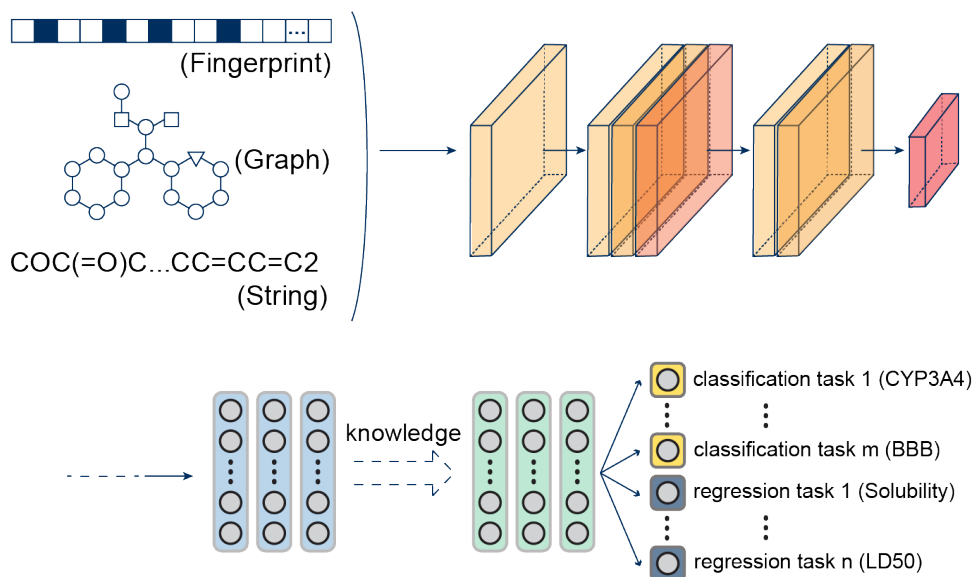
Figure 6.3: Summary of Chapter IV

of ADME-T Properties: Accelerating Drug Discovery and Optimization, we focus on the application of MT ML to predict ADME-T properties of drugs by training on large datasets. Our ML models leverage multi-objective optimization with dynamic weight strategy on various ML models, enabling concurrent training of classification and regression tasks. In our exploration of ST learning across various models, we observed that graph-based models demonstrated superior performance, with Attentive FP excelling in regression tasks and Graph Isomorphism Network (GIN) shining in classification tasks. Subsequently, we delved into MT training across these models, revealing that MT could effectively enhance performance in specific tasks by enabling parameter sharing between tasks, thereby facilitating information exchange among them. Our findings underscored the consistent excellence of graph-based models for MT learning. Additionally, we implemented a dynamic weight strategy to allocate task weight, yielding performance enhancements that were beneficial across models.

As described in Chapter V (ongoing work, Figure 6.4), *Deep Reinforcement Learning-Guided de novo Design of Janus Kinase Inhibitors*, we delve into the theoretical foundation underpinning the efficacy of employing DRL for *de novo* drug design in the
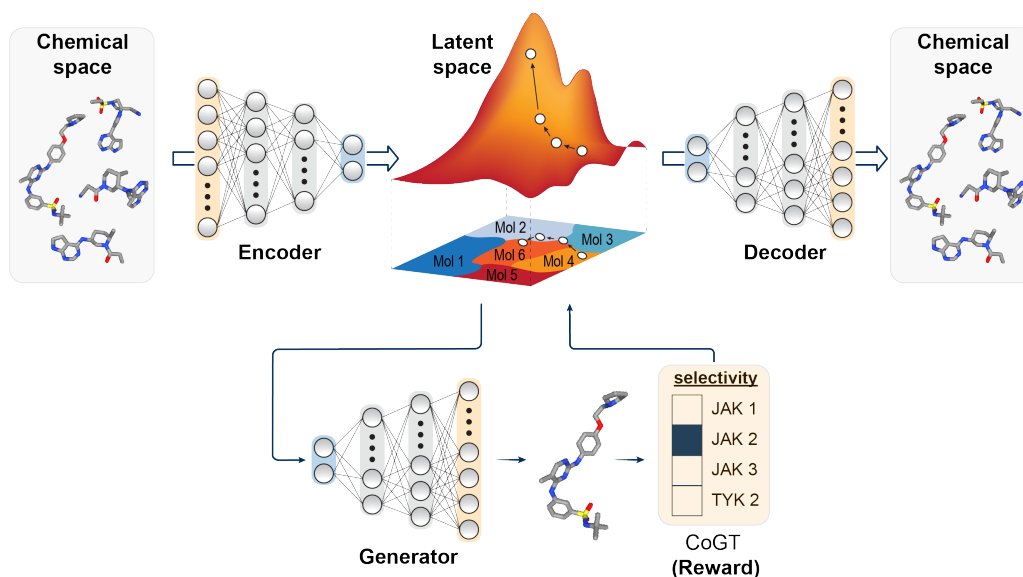
Figure 6.4: Summary of Chapter V

context of JAK inhibitors. This chapter elucidates the rationale behind the utilization of REINFORCE algorithm and VAE. Our approach begins with the organization of the latent space using QED and SA metrics, training on ZINC dataset. Subsequently, a pretrained model is fine-tuned for downstream task of JAK inhibitor generation by training with a MLP classifier using latent vector representation. Leveraging the REINFORCE algorithm, we aim to identify optimal JAK inhibitors based on a pre-defined reward function using ML prediction. Initial experiments have shown promise, generating diverse pan-JAK inhibitors predicted by the CoGT model. Future efforts will concentrate on the generation of JAK-specific inhibitors. Following structure generation, our approach will extend to the utilization of ML models for ADME-T predictions, enabling the assessment of PK properties and toxicity profiles for the se-lected compounds. Subsequent phases involve extensive *in vitro* and *in vivo* studies to validate the suitability of these compounds for further drug development endeavors.

APPENDICES

APPENDIX A

VAE Loss: Evidence Lower Bound

KL divergence is defined as $KL(q||p) = \int q \log \frac{q}{p} = \mathbb{E}_q[\log \frac{q}{p}] \geq 0$ for two distributions q and p [47]. This equation equals 0 if and only if the two distributions are equal, denoted as $q = p$.

Let p^* denote the true distribution of x and p_θ is the estimated distribution of x. We hope that $p_\theta(x) \approx p^*(x)$, which is equivalent to $\min KL(p^*(x)||p_\theta(x))$. Since

$$
\begin{aligned}
\log p_\theta(x) &= \mathbb{E}_{x \sim p^*(\cdot)}[\log p_\theta(x)] = \mathbb{E}_{x \sim p^*(\cdot)}[\log \frac{p^*(x)}{p^*(x)} p_\theta(x)] \\
&= \mathbb{E}_{x \sim p^*(\cdot)} \log p^*(x) + \mathbb{E}_{x \sim p^*(\cdot)}[\log \frac{p_\theta(x)}{p^*(x)}] \\
&= -H(p^*) - KL(p^*(x)||p_\theta(x))
\end{aligned}
$$

We can $\max \log p_\theta(x)$, which is equivalent to $\min KL(p^*(x)||p_\theta(x))$ as $H(p^*) = $ const.

Then to calculate $\log p_\theta(x)$, we have

$$p_\theta(x) = \int p_\theta(x,z)dz = \int \frac{p_\theta(x,z)}{q_\phi(z|x)}q_\phi(z|x)dz = \mathbb{E}_{z \sim q_\phi(z|x)}\frac{p_\theta(x,z)}{q_\phi(z|x)}$$

$$\log p_\theta(x) = \log \mathbb{E}_{z \sim q_\phi(z|x)}\frac{p_\theta(x,z)}{q_\phi(z|x)} \geq \mathbb{E}_{z \sim q_\phi(z|x)}\left[\log \frac{p_\theta(x,z)}{q_\phi(z|x)}\right]$$

$$\text{ELBO} := \mathbb{E}_{z \sim q_\phi(z|x)}\left[\log \frac{p_\theta(x,z)}{q_\phi(z|x)}\right]$$

The objective is then converted from $\max \mathbb{E}_{x \sim p^*(\cdot)}[\log p_\theta(\mathbf{x})]$ to $\max \text{ELBO}$. ELBO can be written as:

$$\text{ELBO} = \mathbb{E}_{z \sim q_\phi(z|x)}\left[\log \frac{p_\theta(x,z)}{q_\phi(z|x)}\right] = \mathbb{E}_{z \sim q_\phi(z|x)}\left[\log \frac{p_\theta(x|z)p_\theta(z)}{q_\phi(z|x)}\right]$$

$$= \mathbb{E}_{z \sim q_\phi(z|x)}\left[\log \frac{p_\theta(z)}{q_\phi(z|x)}\right] + \mathbb{E}_{z \sim q_\phi(z|x)}\left[\log p_\theta(x|z)\right]$$

$$= -\text{KL}(q_\phi(z|x)||p_\theta(z)) + \mathbb{E}_{z \sim q_\phi(z|x)}\left[\log p_\theta(x|z)\right]$$

$$\approx -\text{KL}(q_\phi(z|x)||p_\theta(z)) + \frac{1}{L}\sum_{l=1}^{L}\log p_\theta(x|z^{(l)})$$

To calculate $\text{KL}(q_\phi(z|x)||p_\theta(z))$, we have:

$$\int q_\phi(z|x)\log p_\theta(z)dz = \int \mathcal{N}(z;\mu,\sigma^2)\log \mathcal{N}(z;0,I)dz$$

$$= -\frac{J}{2}\log(2\pi) - \frac{1}{2}\sum_{j=1}^{J}(\mu_j^2 + \sigma_j^2)$$

$$\int q_\phi(z|x)\log q_\phi(z|x)dz = \int \mathcal{N}(z;\mu,\sigma^2)\log \mathcal{N}(z;\mu,\sigma^2)dz$$

$$= -\frac{J}{2}\log(2\pi) - \frac{1}{2}\sum_{j=1}^{J}(1 + \log \sigma_j^2)$$

Thus

$$\mathrm{KL}(\mathrm{q}_\phi(\mathrm{z}|\mathrm{x})||\mathrm{p}_\theta(\mathrm{z})) = \mathbb{E}_{\mathrm{z} \sim \mathrm{q}_\phi(\cdot|\mathrm{x})} \left[\log \frac{\mathrm{q}_\phi(\mathrm{z}|\mathrm{x})}{\mathrm{p}_\theta(\mathrm{z})} \right]$$

$$= \mathbb{E}_{\mathrm{z} \sim \mathrm{q}_\phi(\cdot|\mathrm{x})} \left[\log \mathrm{q}_\phi(\mathrm{z}|\mathrm{x}) \right] - \mathbb{E}_{\mathrm{z} \sim \mathrm{q}_\phi(\cdot|\mathrm{x})} \left[\log \mathrm{p}_\theta(\mathrm{z}) \right]$$

$$= \int \mathrm{q}_\phi(\mathrm{z}|\mathrm{x}) \log \mathrm{q}_\phi(\mathrm{z}|\mathrm{x}) \mathrm{dz} - \int \mathrm{q}_\phi(\mathrm{z}|\mathrm{x}) \log \mathrm{p}_\theta(\mathrm{z}) \mathrm{dz}$$

$$= -\frac{1}{2} \sum_{\mathrm{j}=1}^{\mathrm{J}} (1 + \log \sigma_\mathrm{j}^2 - \mu_\mathrm{j}^2 - \sigma_\mathrm{j}^2)$$

Thus, we obtain the loss for VAE:

$$\min \mathrm{Loss}(\mathrm{x}; \theta, \phi) = \min -\tilde{\mathcal{L}}_{\mathrm{VAE}}(\mathrm{x}; \theta, \phi)$$

$$= \min \mathrm{KL}(\mathrm{q}_\phi(\mathrm{z}|\mathrm{x})||\mathrm{p}_\theta(\mathrm{z})) - \frac{1}{\mathrm{L}} \sum_{\mathrm{l}=1}^{\mathrm{L}} \log \mathrm{p}_\theta(\mathrm{x}|\mathrm{z}^{(\mathrm{l})})$$

$$= \min \left[\mathrm{KL}(\mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma})||\mathcal{N}(\mathbf{0}, \mathbf{I})) - \frac{1}{\mathrm{L}} \sum_{\mathrm{l}=1}^{\mathrm{L}} \log \mathrm{p}_\theta(\mathrm{x}|\mathrm{z}^{(\mathrm{l})}) \right]$$

$$= \min \left[(-\frac{1}{2} \sum_{\mathrm{j}=1}^{\mathrm{J}} (1 + \log \sigma_\mathrm{j}^2 - \mu_\mathrm{j}^2 - \sigma_\mathrm{j}^2)) - \frac{1}{\mathrm{L}} \sum_{\mathrm{l}=1}^{\mathrm{L}} \log \mathrm{p}_\theta(\mathrm{x}|\mathrm{z}^{(\mathrm{l})}) \right]$$

which is equivalent to minimizing the KL divergence along with a reconstruction loss.

APPENDIX B

REINFORCE Algorithm

REINFORCE[142]

1: sample $\{\tau^i\}$ from $\pi_\theta(\mathbf{a}_t|\mathbf{s}_t)$ (run the policy)
2: $\nabla_\theta J(\theta) \approx \sum_i (\sum_t \nabla_\theta \log \pi_\theta(\mathbf{a}_t^i|\mathbf{s}_t^i))(\sum_t r(\mathbf{s}_t^i, \mathbf{a}_t^i) - \frac{1}{N}\sum_{i=1}^{N} r(\tau))$
3: $\theta \leftarrow \theta + \alpha \nabla_\theta J(\theta)$

The goal is to find parameters θ that could maximize the return expectation:

$$\theta^* = \arg\max_\theta \underbrace{\mathbb{E}_{\tau \sim p_\theta(\tau)}\left[\sum_t r(\mathbf{s}_t, \mathbf{a}_t)\right]}_{J(\theta)}$$

$$J(\theta) = \mathbb{E}_{\tau \sim p_\theta(\tau)}[\underbrace{r(\tau)}_{\sum_{t=1}^{T} r(\mathbf{s}_t, \mathbf{a}_t)}] = \int p_\theta(\tau) r(\tau) d\tau$$

$$\nabla_\theta J(\theta) = \int \nabla_\theta p_\theta(\tau) r(\tau) d\tau = \int p_\theta(\tau) \nabla_\theta \log p_\theta(\tau) r(\tau) d\tau$$

$$= \mathbb{E}_{\tau \sim p_\theta(\tau)}[\nabla_\theta \log p_\theta(\tau) r(\tau)]$$

in which $\mathbf{s}_t$ and $\mathbf{a}_t$ represent the state and action at time t, respectively. To calculate

$p_\theta(\tau)$ for the trajectory τ given current policy, we have

$$\underbrace{p_\theta(\mathbf{s}_1, \mathbf{a}_1, \cdots, \mathbf{s}_T, \mathbf{a}_T)}_{p_\theta(\tau)} = p(\mathbf{s}_1) \prod_{t=1}^{T} \pi_\theta(\mathbf{a}_t | \mathbf{s}_t) p(\mathbf{s}_{t+1} | \mathbf{s}_t, \mathbf{a}_t)$$

$$\log p_\theta(\tau) = \log p(\mathbf{s}_1) + \sum_{t=1}^{T} \log \pi_\theta(\mathbf{a}_t | \mathbf{s}_t) + \log p(\mathbf{s}_{t+1} | \mathbf{s}_t, \mathbf{a}_t)$$

$\nabla_\theta(\log p(\mathbf{s}_1)) = 0$ since $p(\mathbf{s}_1)$ does not rely on θ and $\log p(\mathbf{s}_{t+1} | \mathbf{s}_t, \mathbf{a}_t)$ is generated from environment and not rely on θ. Those terms can be get rid of.

$$\nabla_\theta J(\theta) = \mathbb{E}_{\tau \sim p_\theta(\tau)} [\nabla_\theta \log p_\theta(\tau) r(\tau)]$$

$$= \mathbb{E}_{\tau \sim p_\theta(\tau)} \left[\nabla_\theta \left(\log p(\mathbf{s}_1) + \sum_{t=1}^{T} \log \pi_\theta(\mathbf{a}_t | \mathbf{s}_t) + \log p(\mathbf{s}_{t+1} | \mathbf{s}_t, \mathbf{a}_t) \right) r(\tau) \right]$$

$$= \mathbb{E}_{\tau \sim p_\theta(\tau)} \left[\left(\sum_{t=1}^{T} \nabla_\theta \log \pi_\theta(\mathbf{a}_t | \mathbf{s}_t) \right) \left(\sum_{t=1}^{T} r(\mathbf{s}_t, \mathbf{a}_t) \right) \right]$$

$$\approx \frac{1}{N} \sum_{i=1}^{N} \left(\sum_{t=1}^{T} \nabla_\theta \log \pi_\theta(\mathbf{a}_{i,t} | \mathbf{s}_{i,t}) \right) \left(\sum_{t=1}^{T} r(\mathbf{s}_{i,t}, \mathbf{a}_{i,t}) \right)$$

To reduce variance, a baseline $b = \frac{1}{N} \sum_{i=1}^{N} r(\tau)$ can be subtracted without introducing bias in expectation since:

$$\mathbb{E}[\nabla_\theta \log p_\theta(\tau) b] = \int p_\theta(\tau) \nabla_\theta \log p_\theta(\tau) b \, d\tau = \int \nabla_\theta p_\theta(\tau) b \, d\tau$$

$$= b \nabla_\theta \int p_\theta(\tau) d\tau = b \nabla_\theta 1 = 0$$

Thus the final gradient can be written as:

$$\nabla_\theta J(\theta) \approx \frac{1}{N} \sum_{i=1}^{N} \nabla_\theta \log p_\theta(\tau)[r(\tau) - b], \quad b = \frac{1}{N} \sum_{i=1}^{N} r(\tau)$$

BIBLIOGRAPHY

BIBLIOGRAPHY

[1] Muhammad Aurangzeb Ahmad, Carly Eckert, and Ankur Teredesai. Interpretable machine learning in healthcare. In *Proceedings of the 2018 ACM International Conference on Bioinformatics, Computational Biology, and Health Informatics*, pages 559–560, 2018.

[2] Dávid Bajusz, Anita Rácz, and Károly Héberger. Why is tanimoto index an appropriate choice for fingerprint-based similarity calculations? *Journal of Cheminformatics*, 7(1):1–13, 2015.

[3] Shubhasree Banerjee, Ann Biehl, Massimo Gadina, Sarfaraz Hasni, and Daniella M Schwartz. Jak–stat signaling as a target for inflammatory and autoimmune diseases: current and future prospects. *Drugs*, 77(5):521–546, 2017.

[4] G Richard Bickerton, Gaia V Paolini, Jérémy Besnard, Sorel Muresan, and Andrew L Hopkins. Quantifying the chemical beauty of drugs. *Nature Chemistry*, 4(2):90–98, 2012.

[5] BMS. Bristol myers squibb provides update on phase 2 study of deucravacitinib in patients with moderate to severe ulcerative colitis, 2021.

[6] Regine S Bohacek, Colin McMartin, and Wayne C Guida. The art and practice of structure-based drug design: a molecular modeling perspective. *Medicinal Research Reviews*, 16(1):3–50, 1996.

[7] Joanne Bowes, Andrew J Brown, Jacques Hamon, Wolfgang Jarolimek, Arun Sridhar, Gareth Waldron, and Steven Whitebread. Reducing safety-related drug attrition: the use of in vitro pharmacological profiling. *Nature Reviews Drug Discovery*, 11(12):909–922, 2012.

[8] Leo Breiman. Random forests. *Machine Learning*, 45:5–32, 2001.

[9] 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, 2011.

[10] Yingzi Bu, Ruoxi Gao, Bohan Zhang, Luchen Zhang, and Duxin Sun. Cogt: Ensemble machine learning method and its application on jak inhibitor discovery. *ACS Omega*, 2023.

[11] Yingzi Bu, Mohamed Dit Mady Traore, Luchen Zhang, Lu Wang, Zhongwei Liu, Hongxiang Hu, Meilin Wang, Chengyi Li, and Duxin Sun. A gastrointestinal (gi) locally-activating janus kinase (jak) inhibitor to treat ulcerative colitis. *Journal of Biological Chemistry*, page 105467, 2023.

[12] Rich Caruana. Multitask learning. *Machine Learning*, 28:41–75, 1997.

[13] Junyi Chai, Hao Zeng, Anming Li, and Eric WT Ngai. Deep learning in computer vision: A critical review of emerging techniques and application scenarios. *Machine Learning with Applications*, 6:100134, 2021.

[14] Ruolan Chen, Xiangrong Liu, Shuting Jin, Jiawei Lin, and Juan Liu. Machine learning for drug-target interaction prediction. *Molecules*, 23(9):2208, 2018.

[15] Tianqi Chen and Carlos Guestrin. Xgboost: A scalable tree boosting system. In *Proceedings of the 22nd acm sigkdd international conference on knowledge discovery and data mining*, pages 785–794, 2016.

[16] Tianqi Chen, Tong He, Michael Benesty, Vadim Khotilovich, Yuan Tang, Hyunsu Cho, Kailong Chen, Rory Mitchell, Ignacio Cano, Tianyi Zhou, et al. Xgboost: extreme gradient boosting. *R package version 0.4-2*, 1(4):1–4, 2015.

[17] Xing Chen, Chenggang Clarence Yan, Xiaotian Zhang, Xu Zhang, Feng Dai, Jian Yin, and Yongdong Zhang. Drug–target interaction prediction: databases, web servers and computational models. *Briefings in Bioinformatics*, 17(4):696–712, 2016.

[18] Seyone Chithrananda, Gabriel Grand, and Bharath Ramsundar. Chemberta: large-scale self-supervised pretraining for molecular property prediction. *arXiv preprint arXiv:2010.09885*, 2020.

[19] James D Clark, Mark E Flanagan, and Jean-Baptiste Telliez. Discovery and development of janus kinase (jak) inhibitors for inflammatory diseases: Miniperspective. *Journal of Medicinal Chemistry*, 57(12):5023–5038, 2014.

[20] Robert H Collins Jr, Mark Feldman, and John S Fordtran. Colon cancer, dysplasia, and surveillance in patients with ulcerative colitis. *New England Journal of Medicine*, 316(26):1654–1658, 1987.

[21] Kelvin Cooper, Christopher Baddeley, Bernie French, Katherine Gibson, James Golden, Thiam Lee, Sadrach Pierre, Brent Weiss, and Jason Yang. Novel development of predictive feature fingerprints to identify chemistry-based features for the effective drug design of sars-cov-2 target antagonists and inhibitors using machine learning. *ACS Omega*, 6(7):4857–4877, 2021.

[22] Gabriele Corso, Hannes Stärk, Bowen Jing, Regina Barzilay, and Tommi Jaakkola. Diffdock: Diffusion steps, twists, and turns for molecular docking. *arXiv preprint arXiv:2210.01776*, 2022.

[23] Jacques Cosnes, Corinne Gower-Rousseau, Philippe Seksik, and Antoine Cortot. Epidemiology and natural history of inflammatory bowel diseases. *Gastroenterology*, 140(6):1785–1794, 2011.

[24] Constantinos Daskalakis, Dylan J Foster, and Noah Golowich. Independent policy gradient methods for competitive reinforcement learning. *Advances in neural information processing systems*, 33:5527–5540, 2020.

[25] Ryan R Davis, Baoli Li, Sang Y Yun, Alice Chan, Pradeep Nareddy, Steven Gunawan, Muhammad Ayaz, Harshani R Lawrence, Gary W Reuther, Nicholas J Lawrence, et al. Structural insights into jak2 inhibition by ruxolitinib, fedratinib, and derivatives thereof. *Journal of Medicinal Chemistry*, 64(4):2228–2241, 2021.

[26] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018.

[27] Rahul Dey and Fathi M Salem. Gate-variants of gated recurrent unit (gru) neural networks. In *2017 IEEE 60th international midwest symposium on circuits and systems (MWSCAS)*, pages 1597–1600. IEEE, 2017.

[28] Helen Dowden and Jamie Munro. Trends in clinical success rates and therapeutic focus. *Nature Reviews Drug Discovery*, 18(7):495–496, 2019.

[29] Sofia D'Souza, KV Prema, and Seetharaman Balaji. Machine learning models for drug–target interactions: current knowledge and future directions. *Drug Discovery Today*, 25(4):748–756, 2020.

[30] Michael Eisenstein et al. Ulcerative colitis: towards remission. *Nature*, 563(7730):S33, 2018.

[31] Peter Ertl. Cheminformatics analysis of organic substituents: identification of the most common substituents, calculation of substituent properties, and automatic identification of drug-like bioisosteric groups. *Journal of Chemical Information and Computer Sciences*, 43(2):374–380, 2003.

[32] Maria L Faquetti, Francesca Grisoni, Petra Schneider, Gisbert Schneider, and Andrea M Burden. Identification of novel off targets of baricitinib and tofacitinib by machine learning with a focus on thrombosis and viral infection. *Scientific Reports*, 12(1):7843, 2022.

[33] FDA. Fda requires warnings about increased risk of serious heart-related events, cancer, blood clots, and death for jak inhibitors that treat certain chronic inflammatory conditions, 2021.

[34] Evan N Feinberg, Elizabeth Joshi, Vijay S Pande, and Alan C Cheng. Improvement in admet prediction with multitask deep featurization. *Journal of Medicinal Chemistry*, 63(16):8835–8848, 2020.

[35] Andrew Fensome, Catherine M Ambler, Eric Arnold, Mary Ellen Banker, Matthew F Brown, Jill Chrencik, James D Clark, Martin E Dowty, Ivan V Efremov, Andrew Flick, et al. Dual inhibition of tyk2 and jak1 for the treatment of autoimmune diseases: discovery of ((s)-2, 2-difluorocyclopropyl)((1 r, 5 s)-3-(2-((1-methyl-1 h-pyrazol-4-yl) amino) pyrimidin-4-yl)-3, 8-diazabicyclo [3.2. 1] octan-8-yl) methanone (pf-06700841). *Journal of Medicinal Chemistry*, 61(19):8597–8612, 2018.

[36] Octavian-Eugen Ganea, Xinyuan Huang, Charlotte Bunne, Yatao Bian, Regina Barzilay, Tommi Jaakkola, and Andreas Krause. Independent se (3)-equivariant models for end-to-end rigid protein docking. *arXiv preprint arXiv:2111.07786*, 2021.

[37] Wei Gao, Hongxiang Hu, Lipeng Dai, Miao He, Hebao Yuan, Huixia Zhang, Jinhui Liao, Bo Wen, Yan Li, Maria Palmisano, et al. Structure–tissue exposure/selectivity relationship (str) correlates with clinical efficacy/safety. *Acta Pharmaceutica Sinica B*, 12(5):2462–2478, 2022.

[38] Anna Gaulton, Louisa J Bellis, A Patricia Bento, Jon Chambers, Mark Davies, Anne Hersey, Yvonne Light, Shaun McGlinchey, David Michalovich, Bissan Al-Lazikani, et al. Chembl: a large-scale bioactivity database for drug discovery. *Nucleic Acids Research*, 40(D1):D1100–D1107, 2012.

[39] Anna Gaulton, Anne Hersey, Michał Nowotka, A Patricia Bento, Jon Chambers, David Mendez, Prudence Mutowo, Francis Atkinson, Louisa J Bellis, Elena Cibrián-Uhalte, et al. The chembl database in 2017. *Nucleic Acids Research*, 45(D1):D945–D954, 2017.

[40] Felix A Gers, Jürgen Schmidhuber, and Fred Cummins. Learning to forget: Continual prediction with lstm. *Neural Computation*, 12(10):2451–2471, 2000.

[41] Michael K Gilson, Tiqing Liu, Michael Baitaluk, George Nicola, Linda Hwang, and Jenny Chong. Bindingdb in 2015: a public database for medicinal chemistry, computational chemistry and systems pharmacology. *Nucleic Acids Research*, 44(D1):D1045–D1053, 2016.

[42] Rafael Gómez-Bombarelli, Jennifer N Wei, David Duvenaud, José Miguel Hernández-Lobato, Benjamín Sánchez-Lengeling, Dennis Sheberla, Jorge Aguilera-Iparraguirre, Timothy D Hirzel, Ryan P Adams, and Alán Aspuru-Guzik. Automatic chemical design using a data-driven continuous representation of molecules. *ACS Central Science*, 4(2):268–276, 2018.

[43] Kenneth H Grime, Patrick Barton, and Dermot F McGinnity. Application of in silico, in vitro and preclinical pharmacokinetic data for the effective and efficient prediction of human pharmacokinetics. *Molecular Pharmaceutics*, 10(4):1191–1206, 2013.

[44] Will Hamilton, Zhitao Ying, and Jure Leskovec. Inductive representation learning on large graphs. *Advances in Neural Information Processing Systems*, 30, 2017.

[45] Rhiannon N Hardwick, Patrick Brassil, Ilaria Badagnani, Kimberly Perkins, Glenmar P Obedencio, Andrea S Kim, Michael W Conner, David L Bourdet, and Eric B Harstad. Gut-selective design of orally administered izencitinib (td-1473) limits systemic exposure and effects of janus kinase inhibition in nonclinical species. *Toxicological Sciences*, 186(2):323–337, 2022.

[46] Claire Harrison and Alessandro M Vannucchi. Ruxolitinib: a potent and selective janus kinase 1 and 2 inhibitor in patients with myelofibrosis. an update for clinicians. *Therapeutic Advances in Hematology*, 3(6):341–354, 2012.

[47] John R Hershey and Peder A Olsen. Approximating the kullback leibler divergence between gaussian mixture models. In *2007 IEEE International Conference on Acoustics, Speech and Signal Processing-ICASSP'07*, volume 4, pages IV–317. IEEE, 2007.

[48] Julia Hirschberg and Christopher D Manning. Advances in natural language processing. *Science*, 349(6245):261–266, 2015.

[49] 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.

[50] Hongxiang Hu, Mohamed Dit Mady Traore, Ruiting Li, Hebao Yuan, Miao He, Bo Wen, Wei Gao, Colleen B Jonsson, Elizabeth A Fitzpatrick, and Duxin Sun. Optimization of the prodrug moiety of remdesivir to improve lung exposure/selectivity and enhance anti-sars-cov-2 activity. *Journal of Medicinal Chemistry*, 65(18):12044–12054, 2022.

[51] Xiaoyi Hu, Jing Li, Maorong Fu, Xia Zhao, and Wei Wang. The jak/stat signaling pathway: From bench to clinic. *Signal Transduction and Targeted Therapy*, 6(1):1–33, 2021.

[52] James Y. Huang, Bangzheng Li, Jiashu Xu, and Muhao Chen. Unified semantic typing with meaningful label inference. In Marine Carpuat, Marie-Catherine de Marneffe, and Ivan Vladimir Meza Ruiz, editors, *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 2642–2654, Seattle, United States, July 2022. Association for Computational Linguistics.

[53] Kexin Huang, Tianfan Fu, Wenhao Gao, Yue Zhao, Yusuf 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. *Proceedings of Neural Information Processing Systems, NeurIPS Datasets and Benchmarks*, 2021.

[54] James P Hughes, Stephen Rees, S Barrett Kalindjian, and Karen L Philpott. Principles of early drug discovery. *British Journal of Pharmacology*, 162(6):1239–1249, 2011.

[55] Murat Iskar, Georg Zeller, Xing-Ming Zhao, Vera van Noort, and Peer Bork. Drug discovery in the age of systems biology: the rise of computational approaches for data integration. *Current Opinion in Biotechnology*, 23(4):609–616, 2012.

[56] Malemud C. J. The role of the jak/stat signal pathway in rheumatoid arthritis. *Therapeutic Advances in Musculoskeletal Disease, 10(5-6), 117–127*, 2018.

[57] Tiago Janela and Jürgen Bajorath. Simple nearest-neighbour analysis meets the accuracy of compound potency predictions using complex machine learning models. *Nature Machine Intelligence*, 4(12):1246–1255, December 2022.

[58] John Jumper, Richard Evans, Alexander Pritzel, Tim Green, Michael Figurnov, Olaf Ronneberger, Kathryn Tunyasuvunakool, Russ Bates, Augustin Žídek, Anna Potapenko, et al. Highly accurate protein structure prediction with alphafold. *Nature*, 596(7873):583–589, 2021.

[59] Gilaad G Kaplan. The global burden of ibd: from 2015 to 2025. *Nature Reviews Gastroenterology & Hepatology*, 12(12):720–727, 2015.

[60] 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):1–13, 2021.

[61] Alex Kendall, Yarin Gal, and Roberto Cipolla. Multi-task learning using uncertainty to weigh losses for scene geometry and semantics. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7482–7491, 2018.

[62] Meenu Kesarwani, Erika Huber, Zachary Kincaid, Chris R Evelyn, Jacek Biesiada, Mark Rance, Mahendra B Thapa, Neil P Shah, Jarek Meller, Yi Zheng, et al. Targeting substrate-site in jak2 kinase prevents emergence of genetic resistance. *Scientific Reports*, 5(1):1–19, 2015.

[63] S Kim, PA Thiessen, EE Bolton, J Chen, G Fu, A Gindulyte, L Han, J He, S He, BA Shoemaker, J Wang, B Yu, J Zhang, and SH. Bryant. Pubchem substance and compound databases. *Nucleic Acids Research*, 44(D1):D1202–D1213, 2016.

[64] Sunghwan Kim, Jie Chen, Tiejun Cheng, Asta Gindulyte, Jia He, Siqian He, Qingliang Li, Benjamin A Shoemaker, Paul A Thiessen, Bo Yu, et al. Pubchem 2019 update: improved access to chemical data. *Nucleic Acids Research*, 47(D1):D1102–D1109, 2019.

[65] Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.

[66] Thomas N Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. *arXiv preprint arXiv:1609.02907*, 2016.

[67] Thomas N Kipf and Max Welling. Variational graph auto-encoders. *arXiv preprint arXiv:1611.07308*, 2016.

[68] Leo Klarner, Tim G. J. Rudner, Michael Reutlinger, Torsten Schindler, Garrett M Morris, Charlotte Deane, and Yee Whye Teh. Drug discovery under covariate shift with domain-informed prior distributions over functions. In Andreas Krause, Emma Brunskill, Kyunghyun Cho, Barbara Engelhardt, Sivan Sabato, and Jonathan Scarlett, editors, *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pages 17176–17197. PMLR, 23–29 Jul 2023.

[69] Leo Klarner, Tim GJ Rudner, Michael Reutlinger, Torsten Schindler, Garrett M Morris, Charlotte Deane, and Yee Whye Teh. Drug discovery under covariate shift with domain-informed prior distributions over functions. 2023.

[70] Mario Krenn, Florian Häse, AkshatKumar Nigam, Pascal Friederich, and Alan Aspuru-Guzik. Self-referencing embedded strings (selfies): A 100% robust molecular string representation. *Machine Learning: Science and Technology*, 1(4):045024, 2020.

[71] James G Krueger, Iain B McInnes, and Andrew Blauvelt. Tyrosine kinase 2 and janus kinase–signal transducer and activator of transcription signaling and inhibition in plaque psoriasis. *Journal of the American Academy of Dermatology*, 86(1):148–157, 2022.

[72] Remi Lam, Alvaro Sanchez-Gonzalez, Matthew Willson, Peter Wirnsberger, Meire Fortunato, Alexander Pritzel, Suman Ravuri, Timo Ewalds, Ferran Alet, Zach Eaton-Rosen, et al. Graphcast: Learning skillful medium-range global weather forecasting. *arXiv preprint arXiv:2212.12794*, 2022.

[73] Christopher Andrew Lamb, Nicholas A Kennedy, Tim Raine, Philip Anthony Hendy, Philip J Smith, Jimmy K Limdi, Bu'Hussain Hayee, Miranda CE Lomer, Gareth C Parkes, Christian Selinger, et al. British society of gastroenterology consensus guidelines on the management of inflammatory bowel disease in adults. *Gut*, 68(Suppl 3):s1–s106, 2019.

[74] Manisha Lamba, Rong Wang, Tracey Fletcher, Christine Alvey, Joseph Kushner IV, and Thomas C Stock. Extended-release once-daily formulation of tofacitinib: evaluation of pharmacokinetics compared with immediate-release tofacitinib and impact of food. *The Journal of Clinical Pharmacology*, 56(11):1362–1371, 2016.

[75] Ann Lin, Christopher J Giuliano, Ann Palladino, Kristen M John, Connor Abramowicz, Monet Lou Yuan, Erin L Sausville, Devon A Lukow, Luwei Liu,

Alexander R Chait, et al. Off-target toxicity is a common mechanism of action of cancer drugs undergoing clinical trials. *Science Translational Medicine*, 11(509):eaaw8412, 2019.

[76] Yang Liu, You Wu, Xiaoke Shen, and Lei Xie. Covid-19 multi-targeted drug repurposing using few-shot learning. *Frontiers in Bioinformatics*, 1, 2021.

[77] Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*, 2019.

[78] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*, 2017.

[79] Guoshun Luo, Mingqi Chen, Weiting Lyu, Ruheng Zhao, Qian Xu, Qidong You, and Hua Xiang. Design, synthesis, biological evaluation and molecular docking studies of novel 3-aryl-4-anilino-2h-chromen-2-one derivatives targeting erα as anti-breast cancer agents. *Bioorganic & Medicinal Chemistry Letters*, 27(12):2668–2673, 2017.

[80] 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(4-5):677–682, 2008.

[81] Kit-Kay Mak, Ola Epemolu, and Mallikarjuna Rao Pichika. The role of dmpk science in improving pharmaceutical research and development efficiency. *Drug Discovery Today*, 27(3):705–729, 2022.

[82] Ines Filipa Martins, Ana L Teixeira, Luis Pinheiro, and Andre O Falcao. A bayesian approach to in silico blood-brain barrier penetration modeling. *Journal of Chemical Information and Modeling*, 52(6):1686–1697, 2012.

[83] M.E. Mavroforakis and S. Theodoridis. A geometric approach to support vector machine (svm) classification. *IEEE Transactions on Neural Networks*, 17(3):671–682, 2006.

[84] Eyal Mazuz, Guy Shtar, Bracha Shapira, and Lior Rokach. Molecule generation using transformers and policy gradient reinforcement learning. *Scientific Reports*, 13(1):8799, 2023.

[85] Filip Miljkovic, Anton Martinsson, Olga Obrezanova, Beth Williamson, Martin Johnson, Andy Sykes, Andreas Bender, and Nigel Greene. Machine learning models for human in vivo pharmacokinetic parameters with in-house validation. *Molecular Pharmaceutics*, 18(12):4520–4530, 2021.

[86] 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*, 28:711–720, 2014.

[87] Floriane Montanari, Lara Kuhnke, Antonius Ter Laak, and Djork-Arné Clevert. Modeling physico-chemical admet endpoints with multitask graph convolutional networks. *Molecules*, 25(1):44, 2019.

[88] Asher Mullard. First de novo deuterated drug poised for approval. *Nature Reviews Drug Discovery*, 2022.

[89] Markus F Neurath. Cytokines in inflammatory bowel disease. *Nature Reviews Immunology*, 14(5):329–342, 2014.

[90] Ken Ogasawara, Christine Xu, Vanaja Kanamaluru, Nicholas Siebers, Sekhar Surapaneni, Laurence Ridoux, Maria Palmisano, and Gopal Krishna. Excretion balance and pharmacokinetics following a single oral dose of [14 c]-fedratinib in healthy subjects. *Cancer Chemotherapy and Pharmacology*, 86:307–314, 2020.

[91] Ingrid Ordás, Lars Eckmann, Mark Talamini, Daniel C Baumgart, and William J Sandborn. Ulcerative colitis. *Lancet*, 380(9853):1606–1619, November 2012.

[92] Álmos Orosz, Károly Héberger, and Anita Rácz. Comparison of descriptor-and fingerprint sets in machine learning models for adme-tox targets. *Frontiers in Chemistry*, 10:852893, 2022.

[93] John J O'Shea, Daniella M Schwartz, Alejandro V Villarino, Massimo Gadina, Iain B McInnes, and Arian Laurence. The jak-stat pathway: impact on human disease and therapeutic intervention. *Annual Review of Medicine*, 66:311–328, 2015.

[94] Kim Papp, Kenneth Gordon, Diamant Thaçi, Akimichi Morita, Melinda Gooderham, Peter Foley, Ihab G Girgis, Sudeep Kundu, and Subhashis Banerjee. Phase 2 trial of selective tyrosine kinase 2 inhibition in psoriasis. *New England Journal of Medicine*, 379(14):1313–1321, 2018.

[95] Julie M Parmentier, Jeff Voss, Candace Graff, Annette Schwartz, Maria Argiriadi, Michael Friedman, Heidi S Camp, Robert J Padley, Jonathan S George, Deborah Hyland, et al. In vitro and in vivo characterization of the jak1 selectivity of upadacitinib (abt-494). *BMC Rheumatology*, 2(1):1–11, 2018.

[96] Yuzhong Peng, Yanmei Lin, Xiao-Yuan Jing, Hao Zhang, Yiran Huang, and Guang Sheng Luo. Enhanced graph isomorphism network for molecular admet properties prediction. *IEEE Access*, 8:168344–168360, 2020.

[97] Jan Peters and Stefan Schaal. Policy gradient methods for robotics. In *2006 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 2219–2225. IEEE, 2006.

[98] Leif E Peterson. K-nearest neighbor. *Scholarpedia*, 4(2):1883, 2009.

[99] Pavel G Polishchuk, Timur I Madzhidov, and Alexandre Varnek. Estimation of the size of drug-like chemical space based on gdb-17 data. *Journal of Computer-aided Molecular Design*, 27:675–679, 2013.

[100] Bharath Ramsundar, Steven Kearnes, Patrick Riley, Dale Webster, David Konerding, and Vijay Pande. Massively multitask networks for drug discovery. *arXiv preprint arXiv:1502.02072*, 2015.

[101] Joseph Redmon and Ali Farhadi. Yolov3: An incremental improvement. *arXiv preprint arXiv:1804.02767*, 2018.

[102] Steve Rodriguez, Clemens Hug, Petar Todorov, Nienke Moret, Sarah A Boswell, Kyle Evans, George Zhou, Nathan T Johnson, Bradley T Hyman, Peter K Sorger, et al. Machine learning identifies candidates for drug repurposing in alzheimer's disease. *Nature Communications*, 12(1):1033, 2021.

[103] Yu Rong, Yatao Bian, Tingyang Xu, Weiyang Xie, Ying Wei, Wenbing Huang, and Junzhou Huang. Self-supervised graph transformer on large-scale molecular data. *Advances in Neural Information Processing Systems*, 33:12559–12571, 2020.

[104] Andrea Rubbert-Roth, Jeffrey Enejosa, Aileen L Pangan, Boulos Haraoui, Maureen Rischmueller, Nasser Khan, Ying Zhang, Naomi Martin, and Ricardo M Xavier. Trial of upadacitinib or abatacept in rheumatoid arthritis. *New England Journal of Medicine*, 383(16):1511–1521, 2020.

[105] Azucena Salas, Cristian Hernandez-Rocha, Marjolijn Duijvestein, William Faubion, Dermot McGovern, Severine Vermeire, Stefania Vetrano, and Niels Vande Casteele. Jak–stat pathway targeting for the treatment of inflammatory bowel disease. *Nature Reviews Gastroenterology & Hepatology*, 17(6):323–337, 2020.

[106] William J Sandborn, Deanna D Nguyen, David T Beattie, Patrick Brassil, Whitney Krey, Jacky Woo, Eva Situ, Reuben Sana, Erik Sandvik, M Teresa Pulido-Rios, et al. Development of gut-selective pan-janus kinase inhibitor td-1473 for ulcerative colitis: a translational medicine programme. *Journal of Crohn's and Colitis*, 14(9):1202–1213, 2020.

[107] Franco Scarselli, Marco Gori, Ah Chung Tsoi, Markus Hagenbuchner, and Gabriele Monfardini. The graph neural network model. *IEEE Transactions on Neural Networks*, 20(1):61–80, 2008.

[108] Michael Schlichtkrull, Thomas N Kipf, Peter Bloem, Rianne van den Berg, Ivan Titov, and Max Welling. Modeling relational data with graph convolutional networks. In *European semantic web conference*, pages 593–607. Springer, 2018.

[109] Arne Schneuing, Yuanqi Du, Charles Harris, Arian Jamasb, Ilia Igashov, Weitao Du, Tom Blundell, Pietro Lió, Carla Gomes, Max Welling, et al. Structure-based drug design with equivariant diffusion models. *arXiv preprint arXiv:2210.13695*, 2022.

[110] LLC Schrödinger. The PyMOL molecular graphics system.

[111] Mike Schuster and Kuldip K Paliwal. Bidirectional recurrent neural networks. *IEEE Transactions on Signal Processing*, 45(11):2673–2681, 1997.

[112] Philippe Schwaller, Teodoro Laino, Théophile Gaudin, Peter Bolgar, Christopher A Hunter, Costas Bekas, and Alpha A Lee. Molecular transformer: a model for uncertainty-calibrated chemical reaction prediction. *ACS Central Science*, 5(9):1572–1583, 2019.

[113] Danese Silvio and Fiocchi Claudio. Ulcerative colitis. *New England Journal of Medicine*, 365:1713–25, 2011.

[114] Siddharth Singh, Antonio Facciorusso, Parambir S Dulai, Vipul Jairath, and William J Sandborn. Comparative risk of serious infections with biologic and/or immunosuppressive therapy in patients with inflammatory bowel diseases: a systematic review and meta-analysis. *Clinical Gastroenterology and Hepatology*, 18(1):69–81, 2020.

[115] 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, 2021.

[116] Ashley J Snider, Agnieszka B Bialkowska, Amr M Ghaleb, Vincent W Yang, Lina M Obeid, and Yusuf A Hannun. Murine model for colitis-associated cancer of the colon. *Mouse Models for Drug Discovery: Methods and Protocols*, pages 245–254, 2016.

[117] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. *arXiv preprint arXiv:2010.02502*, 2020.

[118] 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, 2019.

[119] Hannes Stärk, Octavian Ganea, Lagnajit Pattanaik, Regina Barzilay, and Tommi Jaakkola. Equibind: Geometric deep learning for drug binding structure prediction. In *International Conference on Machine Learning*, pages 20503–20521. PMLR, 2022.

[120] Teague Sterling and John J Irwin. Zinc 15–ligand discovery for everyone. *Journal of Chemical Information and Modeling*, 55(11):2324–2337, 2015.

[121] Dagmar Stumpfe, Ye Hu, Dilyana Dimova, and Jürgen Bajorath. Recent progress in understanding activity cliffs and their utility in medicinal chemistry: miniperspective. *Journal of Medicinal Chemistry*, 57(1):18–28, 2014.

[122] Vladimir Šukalović, Deana Andrić, Goran Roglić, Sladjana Kostić-Rajačić, A Schrattenholz, and Vukić Šoškić. Synthesis, dopamine d2 receptor binding studies and docking analysis of 5-[3-(4-arylpiperazin-1-yl) propyl]-1h-benzimidazole, 5-[2-(4-arylpiperazin-1-yl) ethoxy]-1h-benzimidazole and their analogs. *European Journal of Medicinal Chemistry*, 40(5):481–493, 2005.

[123] Chenhu Sun, Paul R Ohodnicki, and Emma M Stewart. Chemical sensing strategies for real-time monitoring of transformer oil: A review. *IEEE Sensors Journal*, 17(18):5786–5806, 2017.

[124] Duxin Sun, Wei Gao, Hongxiang Hu, and Simon Zhou. Why 90% of clinical drug development fails and how to improve it? *Acta Pharmaceutica Sinica B*, 2022.

[125] Paqui G Traves, Bernard Murray, Federico Campigotto, René Galien, Amy Meng, and Julie A Di Paolo. Jak selectivity and the implications for clinical inhibition of pharmacodynamic cytokine signalling by filgotinib, upadacitinib, tofacitinib and baricitinib. *Annals of the Rheumatic Diseases*, 80(7):865–875, 2021.

[126] Gemma Turon, Jason Hlozek, John G Woodland, Ankur Kumar, Kelly Chibale, and Miquel Duran-Frigola. First fully-automated ai/ml virtual screening cascade implemented at a drug discovery centre in africa. *Nature Communications*, 14(1):5736, 2023.

[127] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in Neural Information Processing Systems*, 30, 2017.

[128] Michael L Vazquez, Neelu Kaila, Joseph W Strohbach, John D Trzupek, Matthew F Brown, Mark E Flanagan, Mark J Mitton-Fry, Timothy A Johnson, Ruth E TenBrink, Eric P Arnold, et al. Identification of n-{cis-3-[Methyl (7 H-pyrrolo [2, 3-d] pyrimidin-4-yl) amino] cyclobutyl} propane-1-sulfonamide (pf-04965842): a selective jak1 clinical candidate for the treatment of autoimmune diseases. *Journal of Medicinal Chemistry*, 61(3):1130–1152, 2018.

[129] Henrike Veith, Noel Southall, Ruili Huang, Tim James, Darren Fayne, Natalia Artemenko, Min Shen, James Inglese, Christopher P Austin, David G Lloyd, et al. Comprehensive characterization of cytochrome p450 isozyme selectivity across chemical libraries. *Nature Biotechnology*, 27(11):1050–1055, 2009.

[130] Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Lio, and Yoshua Bengio. Graph attention networks. *arXiv preprint arXiv:1710.10903*, 2017.

[131] Alejandro V Villarino, Yuka Kanno, and John J O'Shea. Mechanisms and consequences of jak–stat signaling in the immune system. *Nature Immunology*, 18(4):374–384, 2017.

[132] 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, 2016.

[133] Yanli Wang, Jewen Xiao, Tugba O Suzek, Jian Zhang, Jiyao Wang, and Stephen H Bryant. Pubchem: a public information system for analyzing bioactivities of small molecules. *Nucleic Acids Research*, 37(suppl_2):W623–W633, 2009.

[134] Yimeng Wang, Yaxin Gu, Chaofeng Lou, Yuning Gong, Zengrui Wu, Weihua Li, Yun Tang, and Guixia Liu. A multitask gnn-based interpretable model for discovery of selective jak inhibitors. *Journal of Cheminformatics*, 14(1):16, 2022.

[135] Michael J Waring, John Arrowsmith, Andrew R Leach, Paul D Leeson, Sam Mandrell, Robert M Owen, Garry Pairaudeau, William D Pennie, Stephen D Pickett, Jibo Wang, et al. An analysis of the attrition of drug candidates from four major pharmaceutical companies. *Nature Reviews Drug Discovery*, 14(7):475–486, 2015.

[136] Ming Wen, Zhimin Zhang, Shaoyu Niu, Haozhi Sha, Ruihan Yang, Yonghuan Yun, and Hongmei Lu. Deep-learning-based drug–target interaction prediction. *Journal of Proteome Research*, 16(4):1401–1409, 2017.

[137] M Wenlock and N Tomkinson. Experimental in vitro dmpk and physicochemical data on a set of publicly disclosed compounds, 2015.

[138] Jan Wenzel, Hans Matter, and Friedemann Schmidt. Predictive multitask deep neural network models for adme-tox properties: learning from large data sets. *Journal of Chemical Information and Modeling*, 59(3):1253–1268, 2019.

[139] Gerlinde Wernig, Michael G Kharas, Rachel Okabe, Sandra A Moore, Dena S Leeman, Dana E Cullen, Maricel Gozo, Elizabeth P McDowell, Ross L Levine, John Doukas, et al. Efficacy of tg101348, a selective jak2 inhibitor, in treatment of a murine model of jak2v617f-induced polycythemia vera. *Cancer Cell*, 13(4):311–320, 2008.

[140] Jenna Wiens and Erica S Shenoy. Machine learning for healthcare: on the verge of a major shift in healthcare epidemiology. *Clinical Infectious Diseases*, 66(1):149–153, 2018.

[141] Anthony D William, Angeline C-H Lee, Stéphanie Blanchard, Anders Poulsen, Ee Ling Teo, Harish Nagaraj, Evelyn Tan, Dizhong Chen, Meredith Williams, Eric T Sun, et al. Discovery of the macrocycle 11-(2-pyrrolidin-1-yl-ethoxy)-14, 19-dioxa-5, 7, 26-triaza-tetracyclo [19.3. 1.1 (2, 6). 1 (8, 12)] heptacosa-1 (25), 2 (26), 3, 5, 8, 10, 12 (27), 16, 21, 23-decaene (sb1518), a potent janus kinase 2/fms-like tyrosine kinase-3 (jak2/flt3) inhibitor for the treatment of myelofibrosis and lymphoma. *Journal of Medicinal Chemistry*, 54(13):4638–4658, 2011.

[142] Ronald J Williams. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine Learning*, 8:229–256, 1992.

[143] David H Wolpert. Stacked generalization. *Neural Networks*, 5(2):241–259, 1992.

[144] Stephen T Wrobleski, Ryan Moslin, Shuqun Lin, Yanlei Zhang, Steven Spergel, James Kempson, John S Tokarski, Joann Strnad, Adriana Zupa-Fernandez, Lihong Cheng, et al. Highly selective inhibition of tyrosine kinase 2 (tyk2) for the treatment of autoimmune diseases: discovery of the allosteric inhibitor bms-986165. *Journal of Medicinal Chemistry*, 62(20):8973–8995, 2019.

[145] Felix Wu, Amauri Souza, Tianyi Zhang, Christopher Fifty, Tao Yu, and Kilian Weinberger. Simplifying graph convolutional networks. In *International conference on machine learning*, pages 6861–6871. PMLR, 2019.

[146] Zhaoping Xiong, Dingyan Wang, Xiaohong Liu, Feisheng Zhong, Xiaozhe Wan, Xutong Li, Zhaojun Li, Xiaomin Luo, Kaixian Chen, Hualiang Jiang, et al. Pushing the boundaries of molecular representation for drug discovery with the graph attention mechanism. *Journal of Medicinal Chemistry*, 63(16):8749–8760, 2019.

[147] 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, 2012.

[148] Hua Xu, Michael I Jesson, Uthpala I Seneviratne, Tsung H Lin, M Nusrat Sharif, Liang Xue, Chuong Nguyen, Robert A Everley, John I Trujillo, Douglas S Johnson, et al. Pf-06651600, a dual jak3/tec family kinase inhibitor. *ACS Chemical Biology*, 14(6):1235–1242, 2019.

[149] Keyulu Xu, Weihua Hu, Jure Leskovec, and Stefanie Jegelka. How powerful are graph neural networks? *arXiv preprint arXiv:1810.00826*, 2018.

[150] Minjian Yang, Bingzhong Tao, Chengjuan Chen, Wenqiang Jia, Shaolei Sun, Tiantai Zhang, and Xiaojian Wang. Machine learning models based on molecular fingerprints and an extreme gradient boosting method lead to the discovery of jak2 inhibitors. *Journal of Chemical Information and Modeling*, 59(12):5002–5012, 2019.

[151] Zhenwu Yang, Yujia Tian, Yue Kong, Yushan Zhu, and Aixia Yan. Classification of jak1 inhibitors and sar research by machine learning methods. *Artificial Intelligence in the Life Sciences*, 2:100039, 2022.

[152] Zehong Zhang, Lifan Chen, Feisheng Zhong, Dingyan Wang, Jiaxin Jiang, Sulin Zhang, Hualiang Jiang, Mingyue Zheng, and Xutong Li. Graph neural network approaches for drug-target interactions. *Current Opinion in Structural Biology*, 73:102327, 2022.

[153] Alex Zhavoronkov, Yan A Ivanenkov, Alex Aliper, Mark S Veselov, Vladimir A Aladinskiy, Anastasiya V Aladinskaya, Victor A Terentiev, Daniil A Polykovskiy, Maksim D Kuznetsov, Arip Asadulaev, et al. Deep learning enables rapid identification of potent ddr1 kinase inhibitors. *Nature Biotechnology*, 37(9):1038–1040, 2019.

[154] 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, 2009.