

BENCHMARKING ACCURACY AND GENERALIZABILITY OF FOUR GRAPH NEURAL NETWORKS USING LARGE IN VITRO ADME DATASETS FROM DIFFERENT CHEMICAL SPACES

Fabio Broccatelli ¹, Richard Trager ¹, Michael Reutlinger ², George Karypis ^{3,4}, Mufei Li ⁵

¹ Genentech, ² Roche Innovation Center Basel, ³ AWS AI

⁴ University of Minnesota, ⁵ AWS Shanghai AI Lab

ABSTRACT

In this work, we benchmark a variety of single- and multi-task graph neural network (GNN) models against lower-bar and higher-bar traditional machine learning approaches employing human engineered molecular features. We consider four GNN variants – Graph Convolutional Network (GCN), Graph Attention Network (GAT), Message Passing Neural Network (MPNN), and Attentive Fingerprint (AttentiveFP). So far deep learning models have been primarily benchmarked using lower-bar traditional models solely based on fingerprints, while more realistic benchmarks employing fingerprints, whole-molecule descriptors and predictions from other related endpoints (e.g., LogD7.4) appear to be scarce for industrial ADME datasets. In addition to time-split test sets based on Genentech data, this study benefits from the availability of measurements from an external chemical space (Roche data). We identify GAT as a promising approach to implementing deep learning models. While all the deep learning models significantly outperform lower-bar benchmark traditional models solely based on fingerprints, only GATs seem to offer a small but consistent improvement over higher-bar benchmark traditional models. Finally, the accuracy of in vitro assays from different laboratories predicting the same experimental endpoints appears to be comparable with the accuracy of GAT single-task models, suggesting that most of the observed error from the models is a function of the experimental error propagation.

1 INTRODUCTION

Absorption, distribution, metabolism and elimination (ADME) of new chemical entities were recognized in the 90’ as the leading cause for clinical attribution of drug candidates [12]. This challenge was significantly reduced thanks to the advancements in the area of in vitro ADME modeling, which allows predicting the pharmacokinetic profile of novel compounds based on in vitro biological assays. These advancements were in turn largely driven by the availability of reliable in vitro reagents and advancement in LC-MS instrumentation. At Genentech every new project compound is tested in three “first tier” ADME assays (kinetic solubility, metabolic stability in microsomes, lipophilicity). A fraction of the most promising compounds is additionally tested in more biologically relevant “second tier” assays (e.g., metabolic stability in hepatocytes and permeability) and “third tier” assays that are less amenable for high-throughput screening (e.g., long term culture stability assays).

Quantitative structure activity relationship (QSAR) models are routinely utilized across industry to prioritize compounds to be synthesized or tested [17]. These machine learning (ML) models are employed as “zero tier” assays to prioritize ideas based on the chemical structures. QSAR models allow exploring and de-risking a large portion of neighboring chemical space, which is beneficial due to the high cost of chemical synthesis. Before the use of deep learning (DL) was popularized in chemistry, the choice of the human engineered molecular features or descriptors was typically found to be the most important consideration to achieve accurate ML ADME models [17]. This was particularly true for small datasets for which human engineered descriptors bespoke to the predicted property [4]. With the availability of larger datasets and GPU based technologies, reliance

on "chemically meaningful" descriptors has decreased and machine engineered features produced by graph neural networks (GNNs) have gained popularity due to the ability to develop data driven representations leading to apparent improvement in accuracy and scalability [6, 10, 26].

The strategy behind in vitro ADME data gathering across different companies is guided by trying to maximize assay biological relevance versus assay cost. Frequently the more biologically relevant assays require more resources and are less amenable for high-throughput experiments. Consequently, datasets for the more biologically relevant assays tend to be smaller than datasets for workhorse "first tier" assays, which often aim to predict related endpoints (e.g. metabolic turnover). This sequential approach to compound characterization offers opportunities for transfer learning. Model architectures capable of leveraging information contained in higher volume assays in models for lower volume assays are desirable due to the potential increase in accuracy and ability to generalize.

Several benchmark studies have explored the usefulness of employing neural networks and graph convolution in QSAR models [19, 26, 7, 10, 13]. These studies demonstrated that single- and multi-task DL models substantially outperform single-task traditional QSAR models employing molecular fingerprints over a variety of tasks. Internally at Genentech we have observed that models exclusively based on fingerprints tend to underperform, and we have historically adopted a tiered "dependency based" model architecture. In this architecture, models are trained sequentially, and predictions of some models are used as molecular features in the following tier of models (e.g. calculated lipophilicity is utilized as a feature when predicting solubility). This idea is called feature nets (FN) in the context of neural networks [22, 5]. However, it can be applied to other machine learning models like tree models in this work. While previously published works applying DL to ADME datasets are encouraging, additional studies which better reflect industry practices may be required to adequately benchmark the performance of such models.

In this study, we evaluate the utility of DL based GNN models to predict ADME properties solely based on molecular structures. We evaluate these models on four internal datasets, which are selected based on their relatively large size, importance in the decision-making across all Genentech projects, and assay interdependency (to better address the potential for transfer learning). Two of the four datasets are accompanied with an additional subset from the external chemical space (Roche data), enabling us to evaluate the ability of the models to generalize outside of the neighboring chemical space. In this study, four GNN variants are evaluated in single- and multi-task mode [19]. We compare these results against both lower-bar and higher-bar traditional ML approaches. Both of these benchmarks employ single task dependency free extremely randomized tree models. The lower-bar models are only based on fingerprints, whereas the higher-bar models consider additional RDKit descriptors. The selection of extremely randomized trees as non-DL benchmark, as opposed to more traditional random forest models, emerges from extensive in house evaluations (not discussed in this work).

2 EXPERIMENTAL SECTION

2.1 DATASETS

Four internal datasets are used to train and test models: lipophilicity (LogD7.4), clearance intrinsic in human liver microsomes (HLM CL_{int}), clearance intrinsic in hepatocytes (HH CL_{int}), kinetic solubility phosphate buffer (KinSol) [1, 9, 15, 16]. The internal datasets are split into training and test sets based on the date of testing (time-split). The test sets consist of the compounds tested across a time period of approximately 3 months with the chronological cut-off applied [21]. Genentech and Roche are part of the same organization, although they have historically operated independently. Recently large datasets of ADME data (LogD7.4 and HLM CL_{int}) have been shared. This creates the opportunity to test internal Genentech models on a large portion of unexplored chemical space, which in turn allows investigating the ability of the different models to generalize.

All the endpoints are transformed in logarithmic scale. Due to underlying relationships between the different assays, some models may benefit from incorporating the predictions of other models into the input features during training and inference. For example, LogD7.4 is an important component of water solubility and is a driver for metabolic clearance in both HLM and HH [3]. Similarly, HLM contains most of the drug metabolizing enzymes present in HH. These relationships, defined as "dependencies", are human engineered (established based on the biological understanding of the

Table 1: Datasets Used (# Means Set Size)

Assay	# Train	# Temporal Test	# Roche	Dependencies
LogD7.4	111,144	3,351	58,820	
Kinetic Solubility (uM) at pH 7.4 (KinSol)	99,999	2,915		LogD7.4
Metabolic Stability in human liver microsome (HLM CL _{int})	76,556	2,984	267,120	LogD7.4
Metabolic Stability in isolated human hepatocytes (HH CL _{int})	18,862	476		LogD7.4, HLM CL _{int}

Table 2: Atom and Bond Descriptors for Featurization

Descriptor	Value Range
Atom degree	1, 2, 3, 4, 6
Atom type	B, Br, C, Cl, F, H, I, N, O, P, S, Se, Si
Atom chiral tag	CCW, CW, unspecified, other
Atom formal charge	-1, 0, 1
Atom hybridization	S, SP, SP2, SP3, SP3D2
Atom implicit valence	0, 1, 2, 3
Aromatic atom	True, False
Atom mass	real numbers
Bond type	aromatic, single, double, triple
Conjugated bond	True, False
Bond is in ring	True, False
Bond stereo configuration	OE, NONE, OZ, CIS, TRANS, ANY

mechanistic processes measured by the assay) and have been validated by years of experience in model deployment at Genentech [1]. Table 1 describes the different datasets characterized in this study. The datasets employed in this study are of proprietary nature and therefore not disclosed.

In order to quantify the chemical diversity between the temporal test set and the Roche set as compared against the training set, we calculated the maximum Tanimoto similarity based on chemical fingerprints using a subset of the compounds (20% random selection stratified by the 3 different sets). This analysis was performed for both LogD7.4 and HLM CL_{int}. The average maximum similarity against the training set was 0.64 (standard deviation equals to 0.14) for the temporal test set and 0.42 (standard deviation equals to 0.06) for the Roche set.

Several methods have been proposed as a measure of model applicability domain. However, this is still an open area of research [17]. This work does not provide direct quantification of applicability domain. Instead, it is assumed that the fraction of molecules within the theoretical applicability domain is in aggregate higher in datasets originated from the same company site. This is quantitatively reflected by a significant delta in the average Tanimoto similarity against the training set between Roche and Genentech temporal test set. We acknowledge that Tanimoto similarity is an imperfect way to quantify applicability domain. However, in our experience all the metrics to quantify applicability domain are imperfect, hence the choice to adopt a qualitative assessment.

2.2 METHODS

Graph Construction and Featurization for GNNs. We construct molecular graphs for GNNs, with atoms being nodes and bonds being edges. RDKit [20] was used to generate the set of atom and bond descriptors in Table 2, which were used as input node and edge features for GNNs.

GNNs. We develop the models with Deep Graph Library (DGL) [24] and DGL-LifeSci [14]. We consider four GNN variants, including Graph Convolutional Network (GCN) [11], Graph Attention Network (GAT) [23], Message Passing Neural Network (MPNN) [8], and Attentive Fingerprint (AttentiveFP) [26]. While GCN and GAT were originally developed for node classification, DGL-LifeSci extends them for graph property prediction.

Multi-task DL Architectures. We considered two model architectures for multi-task learning, inspired by Ramsundar et al. [19]. The **parallel architecture** is the same as the architecture for single-task prediction, except that the multilayer perceptron (MLP) has multiple outputs, one per task. The **bypass** architecture employs $T + 1$ GNNs for prediction on T tasks, one shared by all tasks and one additional GNN for each task. For prediction on task t , the input is passed to both the shared GNN and the task-specific GNN. Output representations of the two GNNs are concatenated and then passed on to the prediction pipeline as in the parallel architecture, except that there is no parameter sharing across tasks. Different from the classical parallel architecture, the bypass architecture allows developing both task-shared chemical representations and task-specific ones.

Training of DL Models. Smooth L1 loss was used as the loss function. Adam with weight decay was used for optimization. Early stopping was performed based on the coefficient of determination averaged over tasks for the internal validation set. The hyperparameters were chosen based on the best internal validation performance observed. We used Bayesian optimization with hyperopt [2] for hyperparameter search, with the number of search trials being 20.

2.3 BASELINE MODELS AND ASSAYS

Baseline models were created using ExtraTreesRegressor (XRT)¹ from scikit-learn [18]. The number of trees was 100, the minimum number of samples required to form a leaf was 1, and the maximum tree depth was None. The hyperparameters were chosen based on extensive internal evaluations not described in this work.

We trained two types of baseline models. The **lower-bar** variant (XRT_FP) was only based on RDKit Morgan fingerprints (4096 bits, radius 2). The **higher-bar** variant (XRT_All) was based on both Morgan fingerprints and additional RDKit descriptors. The higher-bar models were sequentially trained with the dependencies described in Table 1. The additional descriptors were Chi0, Chi0n, Chi0v, Chi1, Chi1n, Chi1v, Chi2n, Chi2v, Chi3n, Chi3v, Chi4n, Chi4v, EState_VSA1, EState_VSA2, EState_VSA3, EState_VSA4, EState_VSA5, EState_VSA6, EState_VSA7, EState_VSA8, EState_VSA9, EState_VSA10, EState_VSA11, FractionCSP3, HallKierAlpha, HeavyAtomCount, Kappa1, Kappa2, Kappa3, LabuteASA, MolLogP, MolMR, MolWt, NHOHCount, NOCount, NumAliphaticCarbocycles, NumAliphaticHeterocycles, NumAliphaticRings, NumAromaticCarbocycles, NumAromaticHeterocycles, NumAromaticRings, NumHAcceptors, NumHDonors, NumHeteroatoms, NumRotatableBonds, NumSaturatedCarbocycles, NumSaturatedHeterocycles, NumSaturatedRings, PEOE_VSA1, PEOE_VSA2, PEOE_VSA3, PEOE_VSA4, PEOE_VSA5, PEOE_VSA6, PEOE_VSA7, PEOE_VSA8, PEOE_VSA9, PEOE_VSA10, PEOE_VSA11, PEOE_VSA12, PEOE_VSA13, PEOE_VSA14, RingCount, SMR_VSA1, SMR_VSA2, SMR_VSA3, SMR_VSA4, SMR_VSA5, SMR_VSA6, SMR_VSA7, SMR_VSA8, SMR_VSA9, SMR_VSA10, SlogP_VSA1, SlogP_VSA2, SlogP_VSA3, SlogP_VSA4, SlogP_VSA5, SlogP_VSA6, SlogP_VSA7, SlogP_VSA8, SlogP_VSA9, SlogP_VSA10, SlogP_VSA11, SlogP_VSA12, TPSA, VSA_EState1, VSA_EState2, VSA_EState3, VSA_EState4, VSA_EState5, VSA_EState6, VSA_EState7, VSA_EState8, VSA_EState9, VSA_EState10.

Additional available RDKit descriptors were excluded from the selection due to the observed potential to generate numerical and non-numerical values that were problematic to handle within the modeling framework.

For a limited number of compounds in the LogD7.4 (139) and HLM CL_{int} (45) datasets, experimental data was available in both the Genentech assay and the Roche assay. These sets were utilized to quantify the extent of variability deriving from using related but distinct assay protocols.

¹<https://scikit-learn.org/stable/modules/generated/sklearn.ensemble.ExtraTreesRegressor.html>

Table 3: Performance of Single-Task Models (AvgError for average error, R^2 for coefficient of determination, ILog for the fraction of predictions within 1 log unit)

Model	Dataset	AvgError	R^2	ILog
GCN	LogD7.4 Temporal	0.58	0.65	0.87
GAT	LogD7.4 Temporal	0.66	0.63	0.82
MPNN	LogD7.4 Temporal	0.61	0.62	0.84
AttentiveFP	LogD7.4 Temporal	0.58	0.56	0.87
GCN	LogD7.4 Roche	0.62	0.65	0.81
GAT	LogD7.4 Roche	0.60	0.63	0.83
MPNN	LogD7.4 Roche	0.67	0.60	0.79
AttentiveFP	LogD7.4 Roche	0.68	0.61	0.77
GCN	HLM CL_{int} Temporal	0.39	0.44	0.94
GAT	HLM CL_{int} Temporal	0.38	0.48	0.95
MPNN	HLM CL_{int} Temporal	0.40	0.41	0.93
AttentiveFP	HLM CL_{int} Temporal	0.47	0.53	0.89
GCN	HLM CL_{int} Roche	0.45	0.19	0.93
GAT	HLM CL_{int} Roche	0.42	0.24	0.95
MPNN	HLM CL_{int} Roche	0.49	0.16	0.91
AttentiveFP	HLM CL_{int} Roche	0.46	0.19	0.92
GCN	KinSol Temporal	0.56	0.22	0.86
GAT	KinSol Temporal	0.49	0.26	0.89
MPNN	KinSol Temporal	0.53	0.24	0.86
AttentiveFP	KinSol Temporal	0.56	0.25	0.85
GCN	HH CL_{int} Temporal	0.42	0.39	0.93
GAT	HH CL_{int} Temporal	0.39	0.45	0.98
MPNN	HH CL_{int} Temporal	0.37	0.41	0.96
AttentiveFP	HH CL_{int} Temporal	0.39	0.37	0.95
GCN	Averaged	0.50	0.42	0.89
GAT	Averaged	0.49	0.45	0.90
MPNN	Averaged	0.51	0.41	0.88
AttentiveFP	Averaged	0.52	0.42	0.88
GCN	Averaged, excluding LogD7.4	0.46	0.31	0.92
GAT	Averaged, excluding LogD7.4	0.42	0.36	0.94
MPNN	Averaged, excluding LogD7.4	0.45	0.30	0.92
AttentiveFP	Averaged, excluding LogD7.4	0.47	0.34	0.90

3 RESULTS AND DISCUSSION

The performance of the different models across the temporal test set and the Roche set was evaluated according to three statistical metrics – average error (AvgError), coefficient of determination (R^2), and the fraction of predictions within 1 log unit (ILog).

In order to reduce the combinatorial explosion resulting from four distinct GNN variants, two multi-task model architectures, and four datasets, single-task models were approached initially as in Table 3. GATs performed overall better than the other models and were therefore utilized in multi-task models to evaluate the ability of inter-task transfer learning.

Table 4 characterizes the performance of all GAT and XRT benchmark models analyzed in this study across 6 test sets (4 temporal, 2 external). GAT ST, GAT MTP and GAT MTB are respectively single-task GAT, multi-task GAT with parallel architecture, and multi-task GAT with bypass architecture. While some differences can be observed from dataset to dataset, the conclusions are fairly consistent overall. Fingerprint-only models perform significantly worse than all the other methodologies ($p < 0.05$ for all the metrics) and thus are not an adequate benchmark for DL models. The average error of fingerprint-only models across all datasets equals 0.73, which is substantially higher than that of the other traditional and DL models (average error between 0.48 and 0.52). The gap in performance between the best overall methodologies (GAT MTB and GAT MTP) and traditional models using

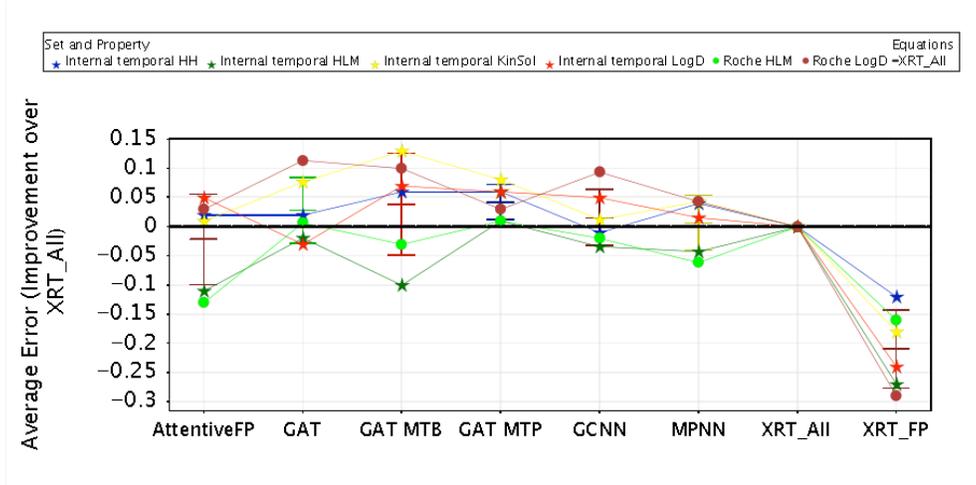


Figure 1: Average error of the different models analyzed across multiple datasets as compared with respect to XRT_All. Stars represent temporal test sets, while circles represent external Roche sets. Green: HLM CL_{int} , Red: LogD7.4, Blue: HH CL_{int} , Yellow: KinSol

additional descriptors and dependencies (XRT_All) is significantly narrower (average error 0.48 vs 0.52). By analyzing the performance difference between all methods with respect to XRT_All as in Figure 1, GAT MTP is the only model for which a statistically significant improvement is observed across all metrics. GAT MTB appears to be more variable and requires more time for training and hyperparameter tuning. As a result, this architecture is less reliable. While GAT ST models are on average superior to XRT_All models, the improvement appears incremental and is not sufficient to achieve statistical significance. Interestingly, these conclusions do not hold when the two external datasets (Roche data) are analyzed.

These datasets are a higher bar because a higher accuracy for them implies a stronger ability to generalize. In addition, their size is considerably larger than the temporal test sets. GAT ST models appear to display the best compromise between local accuracy and generalizability, and are therefore used as the preferred method for subsequent analysis in this work. More comprehensive analysis including additional endpoints and external datasets might be required to cement these conclusions.

The two datasets for LogD7.4 allow an interesting case study. Lipophilicity is an important component of, and is not influenced by, all the other properties modeled in this study. Consequently, the benefits of multi-task learning as a means to enabling transfer learning should theoretically be less apparent when modeling this endpoint. Interestingly, single-task GAT performed worse than the multi-task GAT and the high-bar XRT benchmark on the LogD7.4 temporal test set, while it was the best model on LogD7.4 Roche set, which is composed of over 59k compounds. The gap in average error observed on the temporal test set between single-task and multi-task GAT models was the broadest across all the datasets analyzed. While some degree of underfitting was observed in the local chemical space for single-task GAT, these shortcomings had little influence on the ability of the model to generalize to a different chemical space. Intrigued by this observation, we increased the number of trials in hyperparameter search from 20 to 40 for the special case of single-task GAT on LogD7.4. This improved the performance on both the temporal test set (average error from 0.66 to 0.61, same R^2 , and fraction within 1 log unit from 0.82 to 0.86) and the Roche set (average error from 0.60 to 0.58, R^2 from 0.63 to 0.66, and fraction within 1 log unit from 0.83 to 0.84). For these two different chemical spaces, single-task GAT exhibited both similar accuracy and competitive performance, indicating a superior ability to generalize and a lower susceptibility to overfit to the local space. The performance of single-task GAT trained on Genentech data was additionally compared with the Genentech’s assay ability of Genentech to predict data for compounds tested in the Roche assay as in Table 5. This comparison allows quantifying the composite experimental variability resulting from measuring compounds in a different laboratory with related but distinct protocol. Surprisingly the experimental assay and single-task GAT exhibited comparable accuracy for this dataset.

Table 4: Performance of All GAT and XRT Models (AvgError for average error, R^2 for coefficient of determination, ILog for the fraction of predictions within 1 log unit)

Model	Dataset	AvgError	R^2	ILog
GAT ST	LogD7.4 Temporal	0.66	0.63	0.82
GAT MTP	LogD7.4 Temporal	0.56	0.57	0.89
GAT MTB	LogD7.4 Temporal	0.56	0.56	0.88
XRT_ FP	LogD7.4 Temporal	0.87	0.25	0.68
XRT_ All	LogD7.4 Temporal	0.63	0.55	0.83
GAT ST	LogD7.4 Roche	0.60	0.63	0.83
GAT MTP	LogD7.4 Roche	0.68	0.63	0.77
GAT MTB	LogD7.4 Roche	0.61	0.63	0.81
XRT_ FP	LogD7.4 Roche	1.00	0.14	0.52
XRT_ All	LogD7.4 Roche	0.71	0.50	0.76
GAT ST	HLM CL _{int} Temporal	0.38	0.48	0.95
GAT MTP	HLM CL _{int} Temporal	0.35	0.48	0.96
GAT MTB	HLM CL _{int} Temporal	0.46	0.55	0.89
XRT_ FP	HLM CL _{int} Temporal	0.63	0.21	0.80
XRT_ All	HLM CL _{int} Temporal	0.36	0.45	0.95
GAT ST	HLM CL _{int} Roche	0.42	0.24	0.95
GAT MTP	HLM CL _{int} Roche	0.42	0.24	0.96
GAT MTB	HLM CL _{int} Roche	0.46	0.23	0.94
XRT_ FP	HLM CL _{int} Roche	0.59	0.03	0.82
XRT_ All	HLM CL _{int} Roche	0.43	0.17	0.96
GAT ST	KinSol Temporal	0.49	0.26	0.89
GAT MTP	KinSol Temporal	0.49	0.27	0.91
GAT MTB	KinSol Temporal	0.44	0.30	0.93
XRT_ FP	KinSol Temporal	0.75	0.02	0.73
XRT_ All	KinSol Temporal	0.43	0.17	0.96
GAT ST	HH CL _{int} Temporal	0.39	0.45	0.98
GAT MTP	HH CL _{int} Temporal	0.35	0.50	0.98
GAT MTB	HH CL _{int} Temporal	0.35	0.54	0.95
XRT_ FP	HH CL _{int} Temporal	0.53	0.25	0.85
XRT_ All	HH CL _{int} Temporal	0.41	0.47	0.93
GAT ST	Averaged	0.49	0.45	0.90
GAT MTP	Averaged	0.48	0.45	0.91
GAT MTB	Averaged	0.48	0.47	0.90
XRT_ FP	Averaged	0.73	0.15	0.73
XRT_ All	Averaged	0.52	0.40	0.86
GAT ST	Averaged, excluding LogD7.4	0.42	0.36	0.94
GAT MTP	Averaged, excluding LogD7.4	0.40	0.37	0.95
GAT MTB	Averaged, excluding LogD7.4	0.43	0.40	0.93
XRT_ FP	Averaged, excluding LogD7.4	0.63	0.13	0.80
XRT_ All	Averaged, excluding LogD7.4	0.44	0.33	0.93

The same analysis was performed for the HLM CL_{int} dataset. In this case, the model exhibited a better average error (0.42 vs 0.44) but a significantly lower R^2 (0.24 vs 0.51), reflecting the tendency of the model to populate the middle of the assay dynamic range and a lower ability to rank-order observations at the extremes. Nevertheless, the model was remarkably better than the assay in avoiding large mispredictions (fraction within 1 log unit 0.95 vs 0.91). Taken together, these observations point to that the in-silico model error is largely a product of the in vitro variability, in agreement with the work presented by Wenlock and Carlsson [25].

Table 5: Performance of Genentech In-Silico Models and In Vitro Assays in Predicting Roche In Vitro Assay Data (AvgError for average error, R^2 for coefficient of determination, lLog for the fraction of predictions within 1 log unit)

Property	Method	AvgError	R^2	lLog	N
LogD7.4	Single-Task GAT	0.58	0.66	0.84	61,427
LogD7.4	Experimental	0.55	0.63	0.86	139
HLM CL _{int}	Single-Task GAT	0.42	0.24	0.95	26,719
HLM CL _{int}	Experimental	0.44	0.51	0.91	45

4 CONCLUSIONS

This study identifies multi-task GAT as a state of the art approach across most tasks. All the other single-task DL approaches (GCN, MPNN, AttentiveFP) evaluated did not appear to outperform traditional XRT models utilizing fingerprints, general descriptors and dependencies (e.g., using predictions from other models as descriptors). Similarly, single-task GAT models were found to be only incrementally better than traditional XRT models utilizing fingerprints, general descriptors and dependencies (XRT_ All). In contrast, XRT models exclusively based on fingerprints performed significantly worse compared with all the other approaches across all tasks and datasets. DL models, GAT in particular, exhibited a higher accuracy compared with XRT models on the LogD7.4 datasets. For this task, the availability of the independent dataset composed of Roche data allowed differentiating single-task GAT in terms of the ability to generalize to an external chemical space. Strikingly, the Genentech assay for LogD7.4 was only slightly more accurate than the GAT model, suggesting that the model accuracy is approaching the limit imposed by the experimental assay error propagation. The other prediction tasks approached in this work (KinSol, HLM CL_{int}, HH CL_{int}) are influenced by LogD7.4, and were therefore utilized as benchmark for transfer learning. Surprisingly, multi-task models trained across different tasks in parallel outperform single-task GAT models by a minimal margin, which appears to be erased and reversed for the considerably larger and more diverse external Roche datasets. In aggregate, considering the limited improvement in performance for multi-task models in light of the high potential for transfer learning on these datasets, these findings suggest that transfer learning for GAT could be improved by alternative approaches. One potential approach could involve designing networks explicitly reproducing the hierarchy of inter-task relationships based on the knowledge of the assays. Furthermore, given the remarkable similarity in performance between assays and in silico models, it is possible that the lack of clear differentiation between most approaches explored in this study is a reflection of the proximity to the theoretical accuracy limit dictated by the assay error propagation.

This study highlights the importance of utilizing appropriate benchmarking in order to understand strengths and areas for opportunities for new technologies. In light of these findings, it is recommended that models solely based on fingerprints are not utilized as benchmarks for deep learning approaches in ADME. In addition, this study emphasizes the importance of using orthogonal approaches to model validation where possible. Despite our best efforts to validate models using a higher bar splitting paradigm (time-split), it is the availability of two large external datasets (Roche set for LogD7.4 and HLM CL_{int}) that ultimately allowed us to differentiate single-task GAT models for their ability to generalize (particularly for LogD7.4). From an algorithmic standpoint, these observations might be challenging to address since the opportunity to consistently test models over external chemical space may be missing for many experimental endpoints. From an operational standpoint, these findings reinforce the importance of in-depth outlier analysis, open communication with model users, and partnership with experimental scientists as tools to assess and improve model impact.

This study is based on the use of DGL libraries. They offer convenient APIs, fast training process, and optimal scalability in both the HPC and the cloud. The availability of such libraries for single-task and multi-task training across different GNN variants is an important step towards deep learning democratization and reproducibility in drug discovery.

Natural next steps for this work include the extension of the analysis to ADME assays that are less data rich and the experimentation of different approaches to transfer learning. The code and chosen hyperparameters are released on github here.

REFERENCES

- [1] I. Aliagas, A. Gobbi, T. Heffron, M.-L. Lee, D. F. Ortwine, M. Zak, and S. C. Khojasteh. A probabilistic method to report predictions from a human liver microsomes stability qsar model: a practical tool for drug discovery. *J. Comput.-Aided Mol. Des.*, 29(4):327–338, 2015.
- [2] J. Bergstra, D. Yamins, and D. D. Cox. Making a science of model search: Hyperparameter optimization in hundreds of dimensions for vision architectures. In *Proceedings of the 30th International Conference on Machine Learning*, pages I–115–I–123, 2012.
- [3] F. Broccatelli, I. Aliagas, and H. Zheng. Why decreasing lipophilicity alone is often not a reliable strategy for extending iv half-life. *ACS Med. Chem. Lett.*, 9(6):522–527, 2018.
- [4] P. Crivori, G. Cruciani, P.-A. Carrupt, and B. Testa. Predicting blood-brain barrier permeation from three-dimensional molecular structure. *J. Med. Chem.*, 43(11):2204–2216, 2000.
- [5] I. Davis and A. Stentz. Sensor fusion for autonomous outdoor navigation using neural networks. In *Proceedings of IEEE's Intelligent Robots and Systems Conference*, volume 3, pages 338–343, 1995.
- [6] E. N. Feinberg, R. Sheridan, E. Joshi, V. S. Pande, and A. C. Cheng. Step change improvement in admet prediction with potentialnet deep featurization. *arXiv preprint arXiv:1903.11789*, 2019.
- [7] E. N. Feinberg, E. Joshi, V. S. Pande, and A. C. Cheng. Improvement in admet prediction with multitask deep featurization. *J. Med. Chem.*, 63(16):8835–8848, 2020.
- [8] J. Gilmer, S. S. Schoenholz, P. F. Riley, O. Vinyals, and G. E. Dahl. Neural message passing for quantum chemistry. In D. Precup and Y. W. Teh, editors, *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pages 1263–1272. PMLR, 06–11 Aug 2017.
- [9] J. S. Halladay, S. Wong, S. M. Jaffer, A. K. Sinhababu, and S. C. Khojasteh-Bakht. Metabolic stability screen for drug discovery using cassette analysis and column switching. *Drug Metab. Lett.*, 1(1):67–72, 2007.
- [10] S. Kearnes, K. McCloskey, M. Berndl, V. Pande, and P. Riley. Molecular graph convolutions: moving beyond fingerprints. *J. Comput.-Aided Mol. Des.*, 30(8):595–608, 2016.
- [11] T. N. Kipf and M. Welling. Semi-supervised classification with graph convolutional networks. In *International Conference on Learning Representations*, 2017.
- [12] I. Kola. The state of innovation in drug development. *Clin. Pharmacol. Ther.*, 83(2):227–230, 2008.
- [13] E. B. Lenselink, N. ten Dijke, B. Bongers, G. Papadatos, H. W. T. van Vlijmen, W. Kowalczyk, A. P. IJzerman, and G. J. P. van Westen. Beyond the hype: deep neural networks outperform established methods using a chembl bioactivity benchmark set. *J. Cheminf.*, 9(1), 2017.
- [14] M. Li, J. Zhou, J. Hu, W. Fan, Y. Zhang, Y. Gu, and G. Karypis. Dgl-lifesci: An open-source toolkit for deep learning on graphs in life science. *ACS Omega*, 2021.
- [15] B. Lin and J. H. Pease. A novel method for high throughput lipophilicity determination by microscale shake flask and liquid chromatography tandem mass spectrometry. *Comb. Chem. High T. Scr.*, 16(10):817–825, 2013.
- [16] B. Lin and J. H. Pease. A high throughput solubility assay for drug discovery using microscale shake-flask and rapid uhplc-uv-clnd quantification. *J. Pharm. Biomed. Anal.*, 122:126–140, 2016.
- [17] F. Lombardo, P. V. Desai, R. Arimoto, K. E. Desino, H. Fischer, C. E. Keefer, C. Petersson, S. Winiwarer, and F. Broccatelli. In silico absorption, distribution, metabolism, excretion, and pharmacokinetics (adme-pk): Utility and best practices. an industry perspective from the international consortium for innovation through quality in pharmaceutical development. *J. Med. Chem.*, 60(22):9097–9113, 2017.

- [18] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.
- [19] B. Ramsundar, B. Liu, Z. Wu, A. Verras, M. Tudor, R. P. Sheridan, and V. Pande. Is multitask deep learning practical for pharma? *J. Chem. Inf. Model.*, 57(8):2068–2076, 2017.
- [20] RDKit. RDKit: Open-source cheminformatics. <http://www.rdkit.org>. [Online; accessed 7-Sep-2021].
- [21] R. P. Sheridan. Time-split cross-validation as a method for estimating the goodness of prospective prediction. *J. Chem. Inf. Model.*, 53(4):783–790, 2013.
- [22] A. Varnek, C. Gaudin, G. Marcou, I. Baskin, A. K. Pandey, and I. V. Tetko. Inductive transfer of knowledge: Application of multi-task learning and feature net approaches to model tissue-air partition coefficients. *J. Chem. Inf. Model.*, 49(1):133–144, 2009.
- [23] P. Veličković, G. Cucurull, A. Casanova, A. Romero, P. Liò, and Y. Bengio. Graph attention networks. In *International Conference on Learning Representations*, 2018.
- [24] M. Wang, D. Zheng, Z. Ye, Q. Gan, M. Li, X. Song, J. Zhou, C. Ma, L. Yu, Y. Gai, T. Xiao, T. He, G. Karypis, J. Li, and Z. Zhang. Deep graph library: A graph-centric, highly-performant package for graph neural networks. *arXiv preprint arXiv:1909.01315*, 2020.
- [25] M. C. Wenlock and L. A. Carlsson. How experimental errors influence drug metabolism and pharmacokinetic qsar/qspr models. *J. Chem. Inf. Model.*, 55(1):125–134, 2014.
- [26] Z. Xiong, D. Wang, X. Liu, F. Zhong, X. Wan, X. Li, Z. Li, X. Luo, K. Chen, H. Jiang, and M. Zheng. Pushing the boundaries of molecular representation for drug discovery with the graph attention mechanism. *J. Med. Chem.*, 63(16):8749–8760, 2020.