
Lipophilicity Prediction with Multitask Learning and Molecular Substructures Representation

Nina Lukashina*

Alisa Alenicheva*

Elizaveta Vlasova†

Artem Kondiukov‡

Aigul Khakimova‡

Emil Magerramov‡

Nikita Churikov‡

Aleksei Shpilman§*

Abstract

Lipophilicity is one of the factors determining the permeability of the cell membrane to a drug molecule. Hence, accurate lipophilicity prediction is an essential step in the development of new drugs. In this paper, we introduce a novel approach to encoding additional graph information by extracting molecular substructures. By adding a set of generalized atomic features of these substructures to an established Direct Message Passing Neural Network (D-MPNN) we were able to achieve a new state-of-the-art result at the task of prediction of two main lipophilicity coefficients, namely logP and logD descriptors. We further improve our approach by employing a multitask approach to predict logP and logD values simultaneously. Additionally, we present a study of the model performance on symmetric and asymmetric molecules, that may yield insight for further research.

1 Introduction

Molecular property prediction is one of the fundamental issues in chemoinformatics. It is vital for many applications in chemistry and drug discovery [16, 18]. This paper focuses on the prediction of lipophilicity, a compound’s ability to dissolve in fats, oils, and lipids. Lipophilicity is a key physical property for developing small molecule drugs since it reflects the ability of a drug molecule to permeate cell membranes. It is represented by the descriptors logP (the partition coefficient) and logD (the distribution coefficient).

Most of the current state-of-the-art approaches employ the techniques of deep learning. These models solve two sub-tasks: representing a molecule with an appropriate numerical vector and predicting target property for this vector. A combination of trainable molecules’ representations with simple regression algorithms as molecular properties predictors demonstrates the most promising results according to the current state of the field.

The standard approach to molecular representations is based on molecular fingerprints, which have limited adaptability to a specific task [14].

Recently, many promising methods for generating molecule representations have been based on graph representation learning. The growth of the field started with Graph Convolutional Networks [6, 10, 13, 11, 17], which enabled convolutional neural networks to operate directly on graphs. Among

*JetBrains Research. Saint Petersburg, Russia. Corresponding to nina.lukashina@jetbrains.com.

†ITMO University. Saint Petersburg, Russia.

‡BIOCAD. Saint Petersburg, Russia.

§National Research University Higher School of Economics, Saint Petersburg, Russia.

other methods, one of the most popular is Message Passing Neural Networks [7, 18]. For example, the JtVAE model [9] uses the MPNN approach to split the molecular graph into valid chemical substructures. PAGTN model [4] uses path features in molecular graphs to create global attention layers. OT-GNN [3], the current state-of-the-art approach for the lipophilicity dataset, computes graph embeddings from optimal transport distances between the set of GNN node embeddings and "prototype" point clouds as free parameters.

In this paper, we improved the current state-of-the-art result for the prediction of logP/logD properties by introducing a novel approach for substructures representation in molecular graphs and by the application of multitask learning.

2 StructGNN

Figure 1 shows the overall network architecture of our method named StructGNN. The model consists of two parallel encoders. The first one is a D-MPNN encoder, which creates a molecular embedding h_d through passing edge-focused messages. For a detailed description of the D-MPNN encoder, we refer the reader to the original paper [18].




Figure 1: The overall architecture of our StructGNN approach.

The second encoder takes in generalized features of extracted substructures and produces embedding h_s via the feed-forward neural network (FFNN) and then averaging. We describe this process in detail in Section 2.1.

StructGNN concatenates both embeddings to produce the final molecular vector embedding. This vector could be used with different regression models in order to predict molecular properties. In this paper, we feed it into two layers of fully connected neural network that outputs the final prediction. For the loss function, we use a standard Root Mean Square Error (RMSE) loss.

2.1 Substructure encoder

Substructure encoder considers the molecule’s functional groups, namely rings, acids, amines, esters, and sulfonamides. The encoder represents these groups as hyper-atoms. Both hyper-atoms and atoms form a set of nodes for a particular molecule. We approach them as a set rather than a graph because our experiments have shown that introducing connections between nodes does not improve the model’s performance.

We encode each node in the set with a feature vector. We construct features for hyper-atoms by generalizing the features of atoms, originally used in D-MPNN model. Our substructure encoder doesn’t create new chemical features, but rather generalizes existing atom’s features in order to incorporate the information about the molecule’s functional groups. Table 1 shows the list of all used features.

Table 1: Features for atoms and hyper-atoms

Atom’s feature	Hyper-atom’s feature
Atom type	Number of each atom type in substructure
Number of bonds	Number of substructure’s <i>internal</i> edges
Formal charge	Substructure formal charge
Chirality	Number of hydrogen atoms
Number of Hs	Substructure’s aromaticity
Hybridization	Sum of atomic masses
Aromaticity	Substructure’s external valence
Atomic mass	One-hot encoding for the substructure type

Table 2: Data sources

LogP Dataset	Size (number of samples)
PhysProp [1]	13553
NCI Open Database Compounds [5]	2534
OChem [15]	773
DiverseDataset [12]	707
	17318 in total
logD Dataset	Size (number of samples)
lipophilicity [8]	4200

The feature vector for every node is passed through a one-layer feed-forward neural network. Finally, we average all resulting vectors to produce a single embedding h_s for the whole molecule.

2.2 Multitask learning

We also use multitask learning to predict logP and logD values. Multitask learning aims to improve predictions by solving multiple tasks simultaneously. While logP is used for non-ionizable compounds or the neutral form of ionizable compounds, and logD is used for ionizable compounds, there is an obvious similarity between the two. The Pearson Correlation Coefficient between logP and logD values is 0.66, so we consider descriptors as different targets despite their similarity.

We introduce a new loss function that balances between learning logP and logD properties. If there is only one known value for a molecule (either logP or logD), we calculate the standard RMSE loss function for that values while ignoring the other one. If both values are present, we use mean RMSE between logP and logD errors.

3 Experiments

We used a combined dataset of experimental logP and logD values from sources, listed in Table 2. Data preparation includes SMILES standardization and duplicates removal. The resulting logP and logD datasets contain 13688 and 4166 unique SMILES, respectively. For multitask learning, we merged two datasets and obtained 17603 unique SMILES with logP or logD values. Among them, 251 molecules have both logP and logD values.

In our experiments, we compare a set of different models to create molecules’ representations for logP/logD prediction. For a fair comparison, we use the same two layers of fully connected neural network as a regression model that takes in that generated representation vector. Our evaluation includes circular fingerprints [14] as a baseline method and OT-GNN as the current state-of-the-art approach. We also compare StructGNN with D-MPNN, as it was our base method for improvements, and JtVAE, as it also uses the idea of substructure encoding.

D-MPNN model originally uses an additional vector of 200 global molecular features from RDKit [2]. This approach helps the model to capture global features, while message-passing is fundamentally local in nature [18]. We followed this approach and used RDKit features in all experiments with

Table 3: Results of different models on the lipophilicity descriptors prediction

Model	logP (RMSE)	logP (R^2)	logD (RMSE)	logD (R^2)
Morgan Fingerprints	0.572+/-0.001	0.903+/-0.001	0.775+/-0.022	0.585+/-0.023
JtVAE	0.511+/-0.009	0.922+/-0.003	0.640+/-0.016	0.712+/-0.015
OTGNN	0.508+/-0.016	0.923+/-0.005	0.735+/-0.114	0.617+/-0.124
D-MPNN	0.464+/-0.002	0.936+/-0.001	0.601+/-0.008	0.749+/-0.007
StructGNN	0.453+/-0.004	0.939+/-0.001	0.591+/-0.014	0.758+/-0.012
D-MPNN+Multitask	0.442+/-0.004	0.944+/-0.001	0.554+/-0.004	0.805+/-0.003
StructGNN+Multitask	0.425+/-0.005	0.948+/-0.001	0.537+/-0.008	0.817+/-0.006

Table 4: Results of different models on the lipophilicity descriptors predictions for the symmetric and asymmetric molecules

Model	logP (RMSE)		logD (RMSE)	
	symmetric molecules 608 in total	asymmetric molecules 1380 in total	symmetric molecules 33 in total	asymmetric molecules 4133 in total
Morgan Fingerprints	0.823+/-0.021	0.561+/-0.001	1.322+/-0.200	0.767+/-0.018
JtVAE	0.677+/-0.048	0.504+/-0.008	0.953+/-0.186	0.636+/-0.014
OTGNN	0.797+/-0.056	0.495+/-0.016	1.025+/-0.125	0.731+/-0.116
D-MPNN	0.722+/-0.063	0.437+/-0.005	0.846+/-0.133	0.598+/-0.007
StructGNN	0.671+/-0.017	0.443+/-0.006	1.113+/-0.147	0.584+/-0.013
D-MPNN+Multitask	0.539+/-0.021	0.421+/-0.012	0.458+/-0.132	0.536+/-0.016
StructGNN+Multitask	0.541+/-0.005	0.419+/-0.006	0.367+/-0.113	0.538+/-0.008

D-MPNN and StructGNN models. We used all original features, except MolLogP, as it is calculated logP value from another algorithm. For the StructGNN model, we also excluded structural features, as they are mirrored by generalized hyper-atoms’ features.

We optimized the hyperparameters of every model 4 times on different 25% validation datasets and then measured the performance on a 20% test dataset. Table 3 shows the average RMSE results. Our method achieves new state-of-the-art results in both predicting logP and logD values.

We also analyzed models’ errors depending on the molecule’s symmetry. To distinguish between symmetric and asymmetric molecules, we used canonical atom ranking [2]. We’ve found out that all of the approaches had poor scores for the symmetric compounds, although our model improved the results significantly, as demonstrated in Table 4.

4 Conclusion

We have proposed StructGNN, an extension of the D-MPNN architecture, that encodes additional graph information by extracting molecules’ substructures. Our model shows new state-of-the-art results in lipophilicity prediction tasks for logP and logD targets.

We have further enhanced the performance of the model by applying the multitask approach by training the model to predict logP and logD targets simultaneously.

Finally, we analyzed predictions for symmetric and asymmetric molecules of all studied models. We show that all models struggle with symmetric molecules and no single model outperforms others both in logP and logD prediction for these molecules. Studies in underlying principles that lead to the disparity between the two can lead to further improvement.

The source code of our model is available at <https://github.com/jbr-ai-labs/lipophilicity-prediction>.

References

- [1] The Physical Properties Database (PHYSPROP) by Syracuse Research Corporation (SRC), North Syracuse, USA. URL <https://www.srcinc.com/what-we-do/environmental/scientific-databases.html>.
- [2] The RDKit Documentation — The RDKit 2020.03.1 documentation. URL <https://rdkit.org/docs/index.html>.
- [3] G. Bécigneul, O.-E. Ganea, B. Chen, R. Barzilay, and T. Jaakkola. Optimal Transport Graph Neural Networks. jun 2020. URL <http://arxiv.org/abs/2006.04804>.
- [4] B. Chen, R. Barzilay, and T. Jaakkola. Path-Augmented Graph Transformer Network. may 2019. URL <http://arxiv.org/abs/1905.12712>.
- [5] W. dietrich Ihlenfeldt, J. H. Voigt, B. Bienfait, F. Oellien, and M. C. Nicklaus. Enhanced CACTVS browser of the open NCI database. *Journal of Chemical Information and Computer Sciences*, 42(1):46–57, jan 2002. ISSN 00952338. doi: 10.1021/ci010056s. URL <https://pubmed.ncbi.nlm.nih.gov/11855965/>.
- [6] D. Duvenaud, D. Maclaurin, J. Aguilera-Iparraguirre, R. Gómez-Bombarelli, T. Hirzel, A. Aspuru-Guzik, and R. P. Adams. Convolutional networks on graphs for learning molecular fingerprints. *Advances in Neural Information Processing Systems*, 2015-January:2224–2232, 9 2015. URL <http://arxiv.org/abs/1509.09292>.
- [7] J. Gilmer, S. S. Schoenholz, P. F. Riley, O. Vinyals, and G. E. Dahl. Neural message passing for quantum chemistry. *34th International Conference on Machine Learning, ICML 2017*, 3: 2053–2070, 4 2017. URL <http://arxiv.org/abs/1704.01212>.
- [8] A. Hersey. ChEMBL Deposited Data Set - AZ_dataset. Technical report, EMBL-EBI, 2015. URL <https://www.ebi.ac.uk/chembl/doc/inspect/CHEMBL3301361>.
- [9] W. Jin, R. Barzilay, and T. Jaakkola. Junction Tree Variational Autoencoder for Molecular Graph Generation. *35th International Conference on Machine Learning, ICML 2018*, 5:3632–3648, feb 2018. URL <http://arxiv.org/abs/1802.04364>.
- [10] S. Kearnes, K. McCloskey, M. Berndl, V. Pande, and P. Riley. Molecular graph convolutions: moving beyond fingerprints. *Journal of Computer-Aided Molecular Design*, 30:595–608, 8 2016. ISSN 15734951. doi: 10.1007/s10822-016-9938-8. URL <https://pubmed.ncbi.nlm.nih.gov/27558503/>.
- [11] T. N. Kipf and M. Welling. Semi-supervised classification with graph convolutional networks. *5th International Conference on Learning Representations, ICLR 2017 - Conference Track Proceedings*, 9 2016. URL <http://arxiv.org/abs/1609.02907>.
- [12] S. Martel, F. Gillerat, E. Carosati, D. Maiarelli, I. V. Tetko, R. Mannhold, and P. A. Carrupt. Large, chemically diverse dataset of log P measurements for benchmarking studies. *European Journal of Pharmaceutical Sciences*, 48(1-2):21–29, jan 2013. ISSN 09280987. doi: 10.1016/j.ejps.2012.10.019.
- [13] M. Niepert, M. Ahmed, and K. Kutzkov. Learning convolutional neural networks for graphs. *33rd International Conference on Machine Learning, ICML 2016*, 4:2958–2967, 5 2016. URL <http://arxiv.org/abs/1605.05273>.
- [14] D. Rogers and M. Hahn. Extended-connectivity fingerprints. *Journal of Chemical Information and Modeling*, 50(5):742–754, may 2010. ISSN 15499596. doi: 10.1021/ci100050t. URL <https://pubmed.ncbi.nlm.nih.gov/20426451/>.
- [15] I. Sushko, S. Novotarskyi, R. Körner, A. K. Pandey, M. Rupp, W. Teetz, S. Brandmaier, A. Abdelaziz, V. V. Prokopenko, V. Y. Tanchuk, R. Todeschini, A. Varnek, G. Marcou, P. Ertl, V. Potemkin, M. Grishina, J. Gasteiger, C. Schwab, I. I. Baskin, V. A. Palyulin, E. V. Radchenko, W. J. Welsh, V. Kholodovych, D. Chekmarev, A. Cherkasov, J. Aires-De-Sousa, Q. Y. Zhang, A. Bender, F. Nigsch, L. Patiny, A. Williams, V. Tkachenko, and I. V. Tetko. Online chemical modeling environment (OCHEM): Web platform for data storage, model development and

- publishing of chemical information. *Journal of Computer-Aided Molecular Design*, 25(6): 533–554, jun 2011. ISSN 0920654X. doi: 10.1007/s10822-011-9440-2. URL <https://pubmed.ncbi.nlm.nih.gov/21660515/>.
- [16] J. Vamathevan, D. Clark, P. Czodrowski, I. Dunham, E. Ferran, G. Lee, B. Li, A. Madabhushi, P. Shah, M. Spitzer, and S. Zhao. Applications of machine learning in drug discovery and development. *Nature Reviews Drug Discovery*, 18:463–477, 6 2019. ISSN 14741784. doi: 10.1038/s41573-019-0024-5. URL <https://pubmed.ncbi.nlm.nih.gov/30976107/>.
- [17] X. Wang, Z. Li, M. Jiang, S. Wang, S. Zhang, and Z. Wei. Molecule Property Prediction Based on Spatial Graph Embedding. *Journal of Chemical Information and Modeling*, 59(9):3817–3828, sep 2019. ISSN 1549-9596. doi: 10.1021/acs.jcim.9b00410. URL <https://pubs.acs.org/doi/10.1021/acs.jcim.9b00410>.
- [18] K. Yang, K. Swanson, W. Jin, C. Coley, P. Eiden, H. Gao, A. Guzman-Perez, T. Hopper, B. Kelley, M. Mathea, A. Palmer, V. Settels, T. Jaakkola, K. Jensen, and R. Barzilay. Analyzing learned molecular representations for property prediction. *Journal of Chemical Information and Modeling*, 59:3370–3388, 8 2019. ISSN 15205142. doi: 10.1021/acs.jcim.9b00237. URL <https://pubs.acs.org/sharingguidelines>.