Generative Flow Networks for Lead Optimization in Drug Design (Student Abstract)

Akshat Santhana Gopalan[1], **Sowmya Ramaswamy Krishnan**[2]

[1]John P. Stevens High School, Edison, New Jersey - 08820, USA
[2]Department of Biotechnology, Indian Institute of Technology, Madras, Chennai - 600036, Tamil Nadu, India
akshatsg@outlook.com, krish.krishnan67@gmail.com

Abstract

This paper investigates the application of Generative Flow Networks (GFlowNets) to lead optimization in drug discovery. GFlowNets provide a novel framework for generating diverse molecular structures while optimizing for desired properties, addressing the limitations of traditional methods in exploring vast chemical spaces. We adapt GFlowNets to incrementally modify lead compounds, integrating domain-specific heuristics to guide the generation process. Our method employs the trajectory balance objective on a graph neural network (GNN), to learn a policy that samples fragments based on a multi-objective reward. The reward function ensures increase in cell permeability and similarity to the starting molecule. The results on benchmark datasets of activity cliffs demonstrate that GFlowNets can generate diverse modifications, producing optimized candidate molecules with improvement in cell permeability. This work can be extended with other pharmacokinetic properties for lead optimization in early-stage drug development, potentially accelerating the discovery of novel therapeutics.

Code — https://github.com/AkshatSG/GFN/
Datasets — https://tdcommons.ai/single_pred_tasks/adme
Extended version — NA

Introduction

Lead optimization represents a critical and complex phase in the drug discovery pipeline, involving the iterative refinement of promising compounds to enhance their pharmacological properties while maintaining drug-likeness (He et al., 2021). This process is fundamental to developing candidates with improved efficacy, reduced toxicity, and better pharmacokinetic profiles. Traditional approaches to lead optimization, including high-throughput screening, structure-based design, and medicinal chemistry intuition, often struggle to efficiently explore the vast and intricate chemical space. This limitation can lead to substantial time and resource investments, potentially overlooking promising candidates and slowing the overall drug discovery process.

Recent advances in artificial intelligence and machine learning, particularly in the domain of generative models, have shown significant promise in addressing these challenges. These computational approaches offer the potential for rapid in silico design and evaluation of candidate molecules, potentially accelerating the lead optimization process and reducing the reliance on costly experimental techniques. Among these emerging techniques, Generative-Flow Networks (GFlowNets), introduced by Bengio et al., offer a particularly promising framework for molecular design tasks (Bengio et al., 2021).

GFlowNets are designed to generate diverse samples proportional to a given reward function, making them well-suited for scenarios requiring the exploration of multi-modal distributions in complex spaces. Unlike traditional reinforcement learning methods that tend to converge to a single return-maximizing solution, GFlowNets can model entire distributions, a crucial capability in the context of lead optimization where maintaining a diverse set of candidates is often desirable. This research investigates the potential of GFlowNets for lead optimization tasks in drug design. By formulating the problem as learning a stochastic policy for sequential molecular modifications, we aim to improve the cell permeability while retaining the initial scaffold, similar to an analogue series. Using a fragment library and a lead molecule as input, a graph neural network (GFN agent) is trained to learn the policy for lead optimization using the trajectory balance loss. Results obtained indicate that the GFN agent generates diverse molecular modifications tackling multiple attachment points sequentially. The agent also learns to retain the initial scaffold and generates molecules with high validity and uniqueness.

Methodology

Dataset: The training dataset for the lead optimization task was obtained from the Therapeutics Data Commons (TDC) library (Huang et al., 2022). To train the proxy model for reward prediction, the Caco2 permeability prediction dataset was considered from the DeepDelta study (Fralish et al., 2023). It consists of maximum matched molecular pairs (MMPs) differing by a single substructure between them, which confers a change in the activity of the molecule. The dataset included 455 MMPs with their corresponding log-Papp values, indicative of their cellular permeability potential. A representative example of an MMP from the dataset is provided below (Fig. 1).

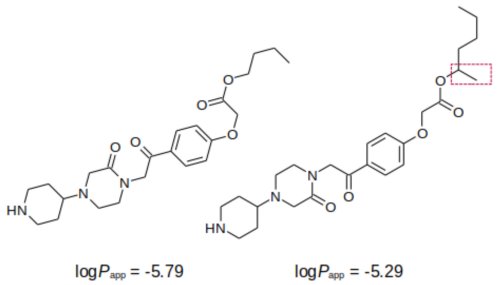

Figure 1: Example of an MMP from the Caco2 permeability dataset. The modification between the molecules is highlighted within the box.

GFN training: The GFN agent is a graph neural network (GNN) trained to predict the fragment modification (policy) proportional to the reward. The permeability predictions for reward calculation will be obtained from a pre-trained proxy model (random forest regressor). The graph neural network will consider an atomic graph of one of the molecules from an MMP as inputs (Wang et al., 2016). The GFN agent is provided with a fragment library (action space) to optimize the lead molecule (Yang et al., 2021). A forward policy and backward policy are trained using the trajectory balance loss.

Evaluation: The top 10 predicted modifications from the model were evaluated based on their ability to retain the starting scaffold provided and improvement in permeability observed.

Novelty and Limitations: While existing methods consider sampling of the entire molecule specific for a target protein of interest either at atom-scale or fragment-scale, this is the first study toward lead optimization with a complete molecule as input. However, the diversity of modifications upon inference is limited to the set of fragments from the fragment library. By expansion of the fragment space using libraries such as Enamine, this limitation can be addressed, although this might lead to an increase in training time of the agent.

Results

Proxy model pre-training: A random forest regressor was pre-trained for Caco2 permeability prediction with the TDC dataset using 1024-bit extended connectivity fingerprints (ECFP4) representations. The train-test split provided in TDC was directly considered for the model. Four metrics namely, regression coefficient (R2), correlation coefficient (r), mean absolute error (MAE) and mean squared error (MSE) were considered for model evaluation (Table 1).

GFN evaluation results: The agent could achieve a maximum molecular validity of 95% and uniqueness of 100% after training for 500 epochs on an 8-core CPU machine. An improvement in average logPapp of -1.5 was observed from the starting permeability value, as shown below (Fig. 2). The molecules had highly diverse modifications and the agent could explore multiple attachment points over iterations, till the maximum atom count (150) was reached. However, the molecules were also found to contain charges and complex functional groups, which will require modified reward functions in future.

R^2	r	MAE	MSE
0.72	0.83	0.5	0.25

Table 1: Metrics obtained by the pre-trained proxy model

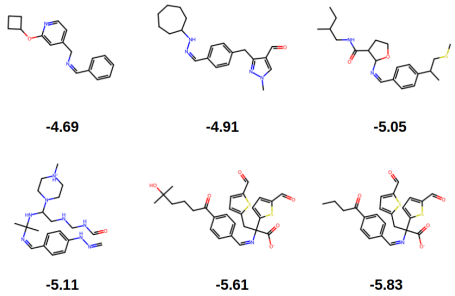

Figure 2: Sample molecules generated by the GFN agent retaining the common starting scaffold. The log*P*app values from the proxy model are also provided.

Conclusions

We have formulated the lead optimization problem using generative flow networks. A multi-objective reward function centred on Caco2 permeability is used as a case study, which can be further extended to any ADMET optimization tasks. Results indicate that with additional rewards to tackle toxicity issues and further finetuning, GFNs provide a promising avenue for lead optimization in drug discovery.

References

Bengio, E.; Jain, M.; Korablyov, M.; Precup, D.; and Bengio, Y. 2021. Flow Network based Generative Models for Non-Iterative Diverse Candidate Generation. arXiv:2106.04399.

Fralish, Z.; Chen, A.; Skaluba, P.; and Reker, D. 2023. DeepDelta: predicting ADMET improvements of molecular derivatives with deep learning. *Journal of Cheminformatics*, 15: 101.

He, J.; You, H.; Sandström, E.; Nittinger, E.; Bjerrum, E. J.; Tyrchan, C.; Czechtizky, W.; and Engkvist, O. 2021. Molecular optimization by capturing chemist's intuition using deep neural networks. *Journal of cheminformatics*, 13: 1–17.

Huang, K.; Fu, T.; Gao, W.; Zhao, Y.; Roohani, Y.; Leskovec, J.; Coley, C. W.; Xiao, C.; Sun, J.; and Zitnik, M. 2022. Artificial intelligence foundation for therapeutic science. *Nature chemical biology*, 18(10): 1033–1036.

Wang, N.-N.; Dong, J.; Deng, Y.-H.; Zhu, M.-F.; Wen, M.; Yao, Z.-J.; Lu, A.-P.; Wang, J.-B.; and Cao, D.-S. 2016. 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.

Yang, S.; Hwang, D.; Lee, S.; Ryu, S.; and Hwang, S. J. 2021. Hit and lead discovery with explorative rl and fragment-based molecule generation. *Advances in Neural Information Processing Systems*, 34: 7924–7936.