CROSS-ATTENTION CONDITIONING FOR MULTI-TARGET MOLECULAR DESIGN IN DEEP GENERATIVE MODELS

ARTHUR CERVEIRA[1]; FREDERICO KREMER[2]; ULISSES CORRÊA[3]

[1]Universidade Federal de Pelotas – aacerveira@inf.ufpel.edu.br
[2]Universidade Federal de Pelotas – fred.s.kremer@gmail.com
[3]Universidade Federal de Pelotas – ulisses@inf.ufpel.edu.br

1. INTRODUCTION

Academic researchers and pharmaceutical companies increasingly rely on artificial intelligence (AI) techniques to reduce costs and accelerate the early stages of the drug discovery pipeline (DDP) (BLASS, 2021). Recent estimates suggest that developing a new therapeutic agent can take up to 10 years, with less than 10% of drug candidates advancing beyond the initial stages of the DDP (DOWNDEN, 2019).

The more traditional application of AI in drug discovery relies on high-throughput screening (HTS), which involves running a set of computational models and AI-based techniques to filter large virtual molecule databases (BLASS, 2021). For this reason, this approach is limited to identifying already known molecules. Although these databases may contain trillions of molecules, they still represent only a tiny fraction of the vast possible chemical space (POLISHCHUK, 2013).

In this context, de novo molecular design (dNMD) emerges as a promising alternative for applying AI in drug discovery. This approach aims to design novel therapeutic compounds from scratch, optimizing the generation process to design molecules with specific desired properties (BLASS, 2021). By doing so, dNMD methods can effectively explore and exploit the vast chemical space, leading to the discovery of new chemical structures (KORSHUNOVA, 2022).

An emerging paradigm in this research field is multi-target drug discovery, which seeks to identify therapeutic compounds capable of modulating multiple biological targets simultaneously. Such multi-target drugs are particularly effective in treating diseases with more complex pathomechanisms, including central nervous system disorders, cardiovascular conditions, and immune diseases MEDINA-FRANCO, 2013). However, the discovery of multi-target-directed ligands still poses numerous challenges and obstacles to researchers.

Few studies in the literature apply dNMD methods to generate multi-target compounds (BLASCHKE, 2022; CHEBROLU, 2023; MUNSON, 2024). The scarcity of open datasets containing molecules with multi-target activity poses a significant challenge in training deep learning models for this task. This restriction inspires the development of more specialized strategies that leverage alternative mechanisms and datasets to distinguish and design multi-target molecules. Considering this, there is a clear need for studies proposing solutions to the challenge of multi-target molecular design. Such techniques could advance the discovery of new multi-target therapeutic agents, leading to more effective treatments for complex diseases.

Considering this, we aim to introduce a new method to design multi-target molecules with transformer-based decoder models. This model architecture consists of a decoder module conditioned by the aggregated latent space representation of the desired properties. These latent space representations are

obtained through an embedding layer, which conditions the molecule generation process via a cross-attention computation with the decoder layers. Our proposed strategy involves using a pooling layer to aggregate different representations, enforcing the generation of molecules optimized for multiple properties.

Our study contributes to the field by (1) applying the conditional training approach to multi-target molecular design, (2) employing attention-based neural networks to generate the compounds, and (3) proposing a new method for multi-objective optimization in dNMD tasks. These advancements distinguish our research from existing works, which rely on fine-tuning (BLASCHKE, 2022) and reinforcement learning (CHEBROLU, 2023; MUNSON, 2024) methods with recurrent neural networks to design multi-target molecules.

2. METHODOLOGY

Our proposed model consists of a transformer-based decoder model conditioned through cross-attention with an aggregated latent space representation of the desired conditions. We employ a pooling layer to aggregate each learned condition representation into a multi-objective representation. Using this aggregated representation, we can simultaneously optimize the molecule generation process towards multiple properties. Fig. 1 depicts the proposed architecture diagram, including the conditional generation, attention-based decoder, and decoder head components.

Fig. 1: Proposed architecture diagram. Source: author.

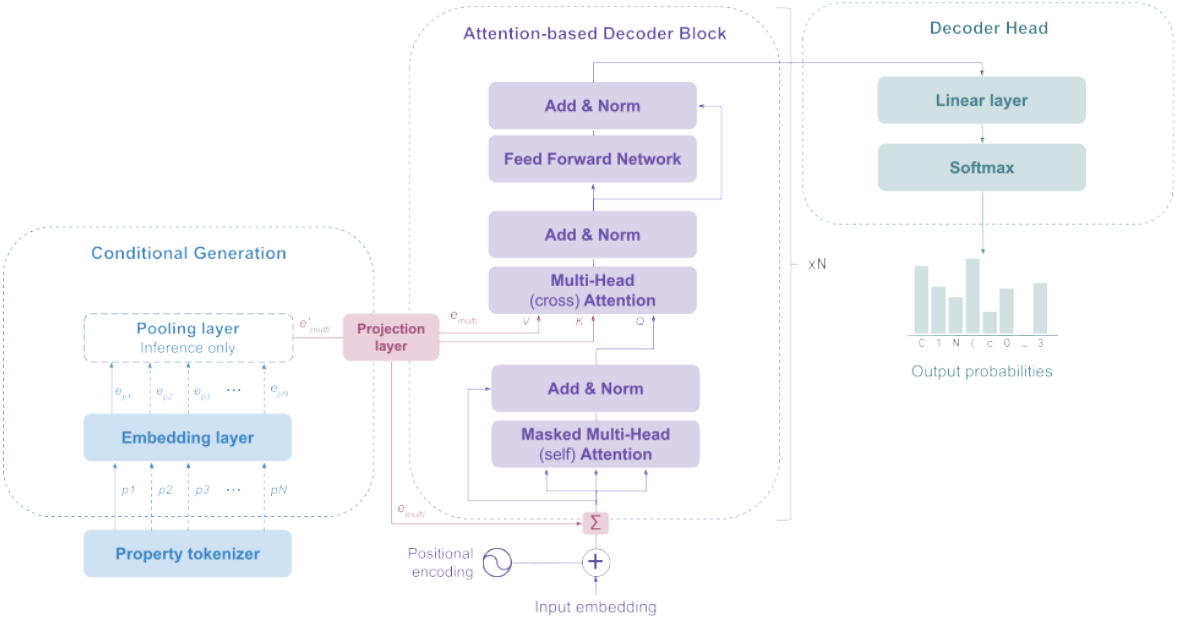

We gathered the data from ExCAPE-DB — a large-scale molecule database with target activity information. The molecules are represented as SMILES strings and go through standard preprocessing and filtering steps. Our training dataset comprises molecule-target pairs, filtering molecules with high activity for the defined target. In our proposed architecture, the molecule is used as input for the decoder block component, while the target information is input for the conditional generation component. During inference, we use multiple targets as input. The pooling layer aggregates the target representations, conditioning the generation of multi-target molecules.

3. RESULTS AND DISCUSSION

We demonstrate the application of our proposed model for multi-target molecular design using the schizophrenia disorder as a case study. For this disorder, we considered the dopamine D2 receptor (D2R) and the serotonin 5-hydroxytryptamine 2A receptor (5-HT2AR) targets. We selected this particular combination of targets based on a previous study on multi-target molecular design which derived target information for diseases from available commercial drugs data (CERVEIRA, 2024).

We analyzed the mean molecular activity distributions for molecules generated by the proposed model compared to two baseline approaches: unconditional generation and single-target conditional generation. We generated a set of 10,000 molecules for each scenario and predicted their molecular activities using machine learning models. We then plotted raincloud charts to visualize the predicted activity distribution for each method. These raincloud plots allow us to compare the proposed method's performance with the baselines, highlighting the efficacy of the proposed multi-target conditional generation approach. The predicted activity is measured in pXC50 units, which represents the degree of target response according to the dose of the therapeutic agent containing a certain substance. Higher pXC50 values indicate lower required therapeutic doses.

Fig. 2: Mean molecular activity distribution for schizophrenia biological targets.

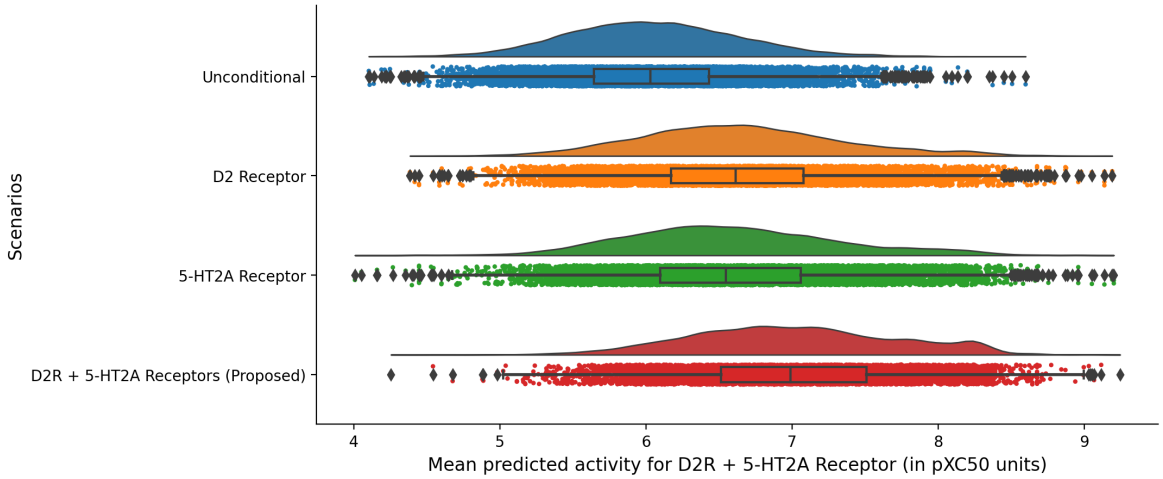

The proposed approach achieved an average pXC50 of 7.03 for mean molecular activity, while the target-specific conditional generation achieved 6.66 and 6.61 for the D2R and 5-HT2AR targets, respectively. The unconditional generation scenario resulted in an average of 6.04 mean molecular activity. Therefore, our model presented consistently higher mean molecular activity than all baseline approaches. To demonstrate the statistical significance of these results, we performed the Kruskal-Wallis H-test to verify the null hypothesis that the population median of all of the groups is equal, resulting in a $p\text{-value} < 0.01$ and an effect size of 20,079.57. We then performed Dunn's test as a posthoc analysis to verify if the pairwise comparison of the samples is also significant, which returned a $p\text{-value} < 0.01$ for every comparison except D2R and 5-HT2AR single-target conditional generation.

4. CONCLUSIONS

This work presented a novel approach for multi-target molecular design with deep generative models. Our proposed architecture exploits the cross-attention mechanism in transformer-based models to condition the generation of molecules with desired properties. To support multi-objective optimization, we aggregated the learned representations for the defined conditions, enforcing the generation of molecules optimized for multiple properties. Our proposed approach generated molecules with a 6.28% and 16.29% higher mean activity against target-specific conditional and unconditional generation, respectively. For future works, we aim to compare our proposed approach with other baselines and assess other target combinations and diseases. Our proposed architecture can also be extended and applied to different tasks in the field of drug discovery.

5. REFERENCES

BLASS, B.E. Basic Principles of Drug Discovery and Development. Amsterdam: Elsevier Science, 2021.

BLASCHKE, T.; BAJORATH, J. Fine-tuning of a generative neural network for designing multi-target compounds. Journal of Computer-Aided Molecular Design, v.36, p.363-371, 2022.

CHEBROLU, A. DeepLig: A De-Novo Computational Drug Design Approach to Generate Multi-Targeted Drugs. 2023 Fifth International Conference on Transdisciplinary AI (TransAI), p.34-41, 2023.

CERVEIRA, A.; KREMER, F.; LOURENÇO, D.; CORRÊA, U.B. Evaluation Framework for AI-driven Molecular Design of Multi-target Drugs: Brain Diseases as a Case Study. 2024 IEEE Congress on Evolutionary Computation (CEC), p.1-8, 2024.

DOWDEN, H.; MUNRO, J. Trends in clinical success rates and therapeutic focus. Nature Reviews Drug Discovery, Nature, v.18, n.7, p.495-496, 2019.

KORSHUNOVA, M.; HUANG, N.; CAPUZZI, S.J.; et al. Generative and reinforcement learning approaches for the automated de novo design of bioactive compounds. Communications Chemistry, v.5, p.129, 2022.

MEDINA-FRANCO, J.L.; GIULIANOTTI, M.A.; WELMAKER, G.S.; HOUGHTEN, R.A. Shifting from the single to the multitarget paradigm in drug discovery. Drug Discovery Today, v.18, n.9, p.495-501, 2013.

MUNSON, B.P.; CHEN, M.; BOGOSIAN, A.; et al. De novo generation of multi-target compounds using deep generative chemistry. Nature Communications, v.15, p.3636, 2024.

POLISHCHUK, P.G.; MADZHIDOV, T.I.; VARNEK, A. Estimation of the size of drug-like chemical space based on GDB-17 data. Journal of Computer-Aided Molecular Design, v.27, n.8, p.675-679, 2013.