Conditional Pre-Trained Transformer for Multi-Target Molecular Design: A Case Study in Drug Discovery for Brain Diseases

Arthur A. Cerveira[1], **Frederico Kremer**[1],
Ricardo M. de Araújo[1], **Ulisses B. Corrêa**[1]

[1]Hub de Inovacão em Inteligência Artificial (H2IA)
Centro de Desenvolvimento Tecnológico (CDTec),
Universidade Federal de Pelotas (UFPel), Brazil

{aacerveira,ulisses}@inf.ufpel.edu.br

Abstract. *Neuropsychiatric disorders affect over one billion people worldwide and remain a major healthcare challenge, partly due to the difficulty of designing therapeutics capable of modulating multiple dysregulated biological targets. From a computational perspective, multi-target drug discovery can be formulated as a multi-objective optimization problem in high-dimensional chemical space. This paper presents CoMPO-GPT, a transformer-based conditional generative framework that leverages cross-attention to condition molecular generation on aggregated latent representations of target-specific properties. By formalizing multi-target molecular design as property-based optimization within a generative modeling paradigm, CoMPO-GPT bridges advances in foundational transformer architectures and applied computing in healthcare. We evaluate the framework on three brain disease benchmarks: Alzheimer's disease, schizophrenia, and Parkinson's disease, across mean molecular activity, distribution-learning metrics, multi-property optimization benchmarks, and docking simulations. Results show that CoMPO-GPT achieves state-of-the-art goal-directed performance in two of three scenarios while maintaining high validity, uniqueness, and novelty metrics. Our analysis further identifies distinct learning paradigms that explain performance differences for multi-objective generative modeling in computational drug discovery.*
Code and data available @ https://github.com/arthurcerveira/CoMPO-GPT.

Resumo. *Transtornos neuropsiquiátricos afetam mais de um bilhão de pessoas e representam um grande desafio terapêutico, em parte devido à necessidade de modular múltiplos alvos biológicos simultaneamente. Este trabalho apresenta o CoMPO-GPT, um modelo generativo condicional baseado em transformers que formula a descoberta de fármacos multi-alvo como um problema de otimização multiobjetivo no espaço químico. O método condiciona a geração molecular por meio de representações em espaço latente agregadas de propriedades farmacológicas. Avaliações em três cenários (doença de Alzheimer, esquizofrenia e Parkinson) mostram desempenho competitivo em otimização direcionada, mantendo validade, diversidade e novidade molecular. Os resultados também revelam padrões de aprendizado que ajudam a explicar o comportamento de modelos generativos para o design molecular multi-alvo.*

1. Introduction

Drug discovery remains a major computational challenge in healthcare, requiring exploration of vast chemical spaces under multiple biological and pharmacological constraints. Developing a new drug can take up to 10 years, with fewer than 10% of candidates advancing beyond early pipeline stages [Dowden and Munro 2019]. The challenge is particularly severe in central nervous system (CNS) drug development, where failure rates exceed 90%, partly because single-target therapies often fail to address the multifactorial nature of brain disorders [Danon et al. 2019]. *De novo* molecular design (dNMD) addresses this challenge by using machine learning to generate novel compounds while optimizing pharmacological properties.

Recent work has applied AI to multi-target drug discovery (MTDD), designing compounds that modulate multiple biological targets simultaneously [Bi et al. 2025]. This setting introduces a high-dimensional optimization problem requiring the simultaneous satisfaction of activity, selectivity, and drug-likeness constraints under limited supervision [Angelo et al. 2023]. However, datasets containing confirmed multi-target molecules are scarce, and most existing approaches rely on fine-tuning or reinforcement learning (RL), which present stability and efficiency limitations [Chebrolu 2023, Munson et al. 2024, Ai et al. 2024]. In contrast, conditional generative modeling offers a promising yet underexplored alternative for structured multi-property optimization (MPO) in molecular design.

To address this gap, this dissertation introduces CoMPO-GPT (**Co**nditional **M**ulti-**P**roperty **O**ptimization for **G**enerative **P**re-trained **T**ransformer), a conditional generative framework that extends transformer architectures to controllable multi-objective molecular design. The model conditions molecule generation on aggregated target-specific latent representations through cross-attention, enabling property-driven optimization without RL or task-specific fine-tuning. We evaluate CoMPO-GPT on three brain disease benchmarks: Alzheimer's disease (AD), schizophrenia (SCZ), and Parkinson's disease (PD), considering mean molecular activity (MMA), distribution-learning metrics, MPO benchmarks, and molecular docking. Beyond its pharmacological application, this work contributes to applied computing in healthcare by formalizing multi-target drug design as a conditional generative optimization problem and providing an open, reproducible framework for benchmarking multi-objective molecular generation.

2. CoMPO-GPT

CoMPO-GPT formulates multi-target drug discovery as a conditional generative modeling problem. Instead of relying on RL or post-hoc filtering, desired therapeutic properties are incorporated directly into molecule generation.

The model builds on a transformer decoder [Vaswani et al. 2017] and conditions generation through cross-attention over property embeddings (Fig. 1).

Decoder. Self-attention models dependencies in SMILES sequences, while cross-attention injects property constraints during autoregressive decoding. The model is trained with next-token prediction:

$$\mathcal{L}_{\text{CE}} = -\sum_{t=1}^{T} \log P(y_t \mid y_{<t}, \mathcal{C})$$

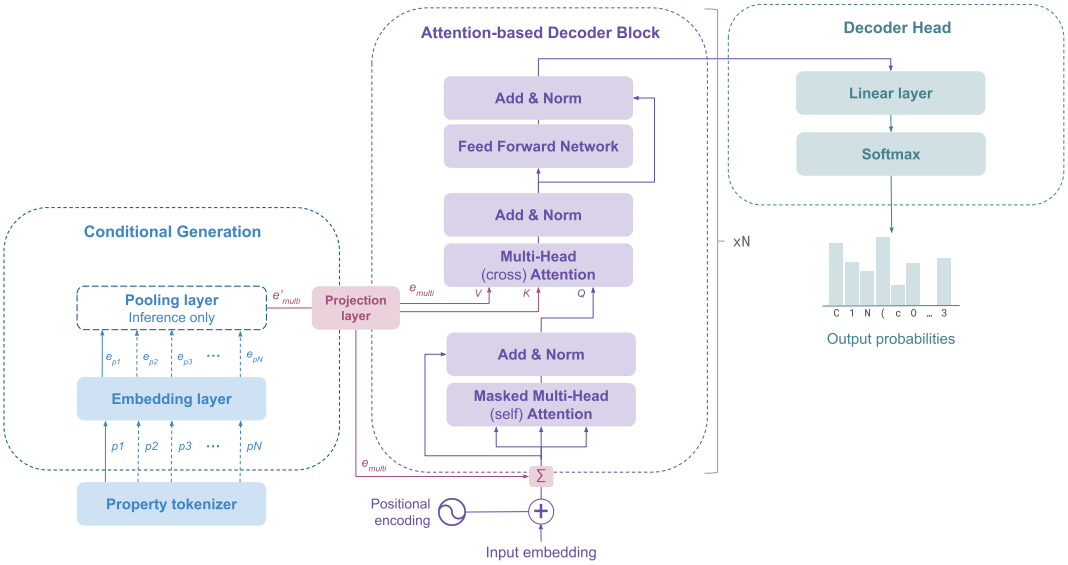

Figure 1. CoMPO-GPT architecture: conditional generation, attention-based decoder block, and decoder head.

Conditional Generation. Each property p is associated with a learned embedding e_p. Unlike prior approaches requiring fixed conditioning or joint multi-property training, CoMPO-GPT enables compositional conditioning by aggregating independently learned embeddings:

$$e_{\text{multi}} = \text{proj}(\text{pool}(e_{p_1}, ..., e_{p_n}))$$

which is injected both into token embeddings and cross-attention keys/values.

This design represents the core novelty of the method, allowing properties to be learned independently and composed at inference time, enabling MPO without explicit multi-property training data and generalization to unseen objective combinations.

Training. Pre-training uses 1.58M molecules from MOSES. Conditional training then introduces property embeddings using ExCAPE-DB subsets (357K molecules across 427 targets). The final model has 3 decoder layers, 1024 hidden dimensions, and 8 attention heads ($\sim$8.4M parameters).

Inference. Embeddings of desired properties are aggregated via sum pooling for autoregressive decoding: $P(y \mid \mathcal{C}) = \prod_{t=1}^{T} P(y_t \mid y_{<t}, e_{\text{multi}})$, achieving MPO through conditional likelihood modeling rather than external reward maximization.

From a computational perspective, CoMPO-GPT demonstrates how transformer-based conditional generation can be repurposed for structured compositional MPO under data scarcity, a setting common in healthcare applications.

3. Experimental Setup

The experimental setup is designed to evaluate CoMPO-GPT as a computational framework for MPO in healthcare-related molecular design, emphasizing reproducibility, benchmarking consistency, and comparative evaluation against state-of-the-art generative approaches. All experiments were conducted under identical data splits and evaluation criteria to ensure fair comparison and reproducibility.

Targets. Three disease settings selected via co-occurrence analysis to represent established polypharmacological mechanisms. For AD, AChE and MAO-B; for SCZ, D2R and 5-HT2AR; for PD, D2R and D3R.

Datasets. Pre-training uses MOSES (1.58M molecules) [Polykovskiy et al. 2020]. Conditional training uses ExCAPE-DB [Sun et al. 2017], selecting the top 1,000 most active molecules for each of 427 targets (357K molecules). Activity predictors are trained on 294 targets with $\geq$1,000 assays (614K molecules).

Baselines & Evaluation. Compared to RL-based baselines DeepLig [Chebrolu 2023], POLYGON [Munson et al. 2024], and MTMol-GPT [Ai et al. 2024] on (1) MMA using multi-task MPNN predictors [Cerveira et al. 2026]; (2) distribution-learning metrics (validity, uniqueness, novelty) [Polykovskiy et al. 2020] (3) MPO benchmarks including MMA, BBB permeability, CNS drug-likeness, and SAScore [Cerveira et al. 2024]; (4) molecular docking with AutoDock Vina [Eberhardt et al. 2021].

4. Results

4.1. Mean Molecular Activity & Distribution-Learning Metrics

Table 1 reports MMA, defined as the average predicted activity across targets and used as a proxy for therapeutic suitability. CoMPO-GPT achieves the highest MMA in SCZ (7.47) and PD (7.54), while POLYGON performs best in AD (6.45). Kruskal–Wallis tests indicate significant differences across models. Among latent representations aggregation strategies, sum pooling outperformed mean and max.

Table 1. Average MMA (pXC50) comparison between CoMPO-GPT and baselines.

Disease	Model	MMA (pXC50)
AD	CoMPO-GPT	5.58 ± 0.65
	DeepLig	5.95 ± 0.07
	MTMol-GPT	5.83 ± 0.58
	POLYGON	$\mathbf{6.45 \pm 0.18}$
SCZ	CoMPO-GPT	$\mathbf{7.47 \pm 0.83}$
	DeepLig	5.70 ± 0.13
	MTMol-GPT	6.95 ± 0.67
	POLYGON	6.50 ± 0.20
PD	CoMPO-GPT	$\mathbf{7.54 \pm 0.73}$
	DeepLig	5.19 ± 1.87
	MTMol-GPT	7.19 ± 0.69
	POLYGON	6.34 ± 0.36

Table 2 summarizes distribution-learning metrics. CoMPO-GPT achieves the highest uniqueness (0.70 – 0.88), indicating broader chemical space exploration. POLYGON and DeepLig show the highest validity and novelty. MTMol-GPT exhibits lower novelty (0.32 – 0.45) and higher KL/FCD, consistent with memorization of training actives, which is an undesirable behavior.

4.2. Multi-Property Optimization Benchmarks

Table 3 reports MPO performance using MMA, BBB permeability, CNS drug-likeness, and SAScore. CoMPO-GPT achieves the highest MPO scores for AD (0.742) and SCZ

Table 2. Summarized distribution-learning metrics for generated molecules from each baseline across the three disease scenarios.

Metric	CoMPO-GPT	MTMol-GPT	POLYGON	DeepLig
Validity	$\sim$90%	**>96%**	>94%	>92%
Uniqueness	**70% – 88%**	32% – 51%	37% – 44%	1% – 65%
Novelty	>95%	32% – 45%	**100%**	**100%**
KL Div	0.89 – 0.93	**$\sim$0.99**	0.00 – 0.06	0.22 – 0.23
FCD	0.26 – 0.31	**0.74 – 0.91**	0.00	0.00

(0.800), largely driven by strong CNS drug-likeness. For PD, MTMol-GPT marginally outperforms CoMPO-GPT (< 0.001). Integrating pharmacokinetic and drug-likeness constraints into the evaluation demonstrates that multi-objective optimization in healthcare cannot be reduced to activity maximization alone

Table 3. Best-performing model per metric on the MPO benchmarks for each considered disease.

Disease	Score	MMA	BBB	CNS MPO	SAScore
AD	CoMPO-GPT	POLYGON	DeepLig	CoMPO-GPT	DeepLig
SCZ	CoMPO-GPT	CoMPO-GPT	POLYGON	CoMPO-GPT	DeepLig
PD	MTMol-GPT	POLYGON	DeepLig	CoMPO-GPT	DeepLig

4.3. Docking and Two-Paradigm Analysis

Docking experiments confirmed that CoMPO-GPT generates ligands with binding affinities comparable to known actives, matching the reference set for the considered diseases. This alignment further supports the biological plausibility of the generated molecules.

Our analysis identifies two learning paradigms: (1) **distribution-matching** (CoMPO-GPT, MTMol-GPT), learning from active molecules through conditional or imitation learning; and (2) **objective-optimizing** (DeepLig, POLYGON), guided by external scoring functions via RL. Dataset analysis validates this: AD targets show minimal co-activity overlap (3.6%), favoring scorer-based models; PD exhibits the highest overlap (53 – 70%) and correlation (0.36 – 0.54), explaining the strong Group 1 performance.

From a computational standpoint, this distinction clarifies when likelihood-based conditional modeling is preferable to RL in healthcare settings: distribution-matching approaches benefit from correlated multi-property data, whereas scorer-based methods are advantageous when co-activity overlap is sparse. This provides a principled guideline for selecting generative strategies based on dataset structure rather than heuristic choice.

5. Conclusion

CoMPO-GPT introduces a conditional transformer framework that formulates multi-target drug design as controllable generative modeling in chemical space. By conditioning generation through cross-attention over aggregated property embeddings, the model enables MPO without RL or task-specific fine-tuning.

Experiments across three neuropsychiatric disease benchmarks show that CoMPO-GPT achieves competitive or superior goal-directed optimization while maintaining chemical validity, novelty, and CNS-relevant pharmacological constraints. Our

analysis further identifies two learning paradigms among the baselines, highlighting how dataset structure and target co-activity influence generative modeling strategies.

From an applied computing perspective, this work demonstrates how transformer architectures can support structured optimization problems in healthcare-related drug discovery. While results rely on computational evaluation, the framework provides a scalable and reproducible foundation for AI-assisted therapeutic design. Future work will explore hybrid training strategies, adaptive objective weighting, and higher-fidelity molecular representations.

References

Ai, C., Yang, H., Liu, X., Dong, R., Ding, Y., and Guo, F. (2024). MTMol-GPT: De novo multi-target molecular generation with transformer-based generative adversarial imitation learning. *PLOS Computational Biology*, 20(6):1–23.

Angelo, J. S., Guedes, I. A., Barbosa, H. J. C., and Dardenne, L. E. (2023). Multi-and many-objective optimization: present and future in de novo drug design. *Frontiers in Chemistry*, 11.

Bi, X., Wang, Y., Wang, J., and Liu, C. (2025). Machine learning for multi-target drug discovery. *Pharmaceutics*, 17(9):1186.

Cerveira, A., Kremer, F., Gomes, G. A., and Corrêa, U. B. (2026). Enhancing graph neural networks for multi-target activity prediction through multi-task learning and knowledge distillation. In *Intelligent Systems*, pages 464–477. Springer Nature Switzerland.

Cerveira, A., Kremer, F., Lourenço, D., and Corrêa, U. B. (2024). Evaluation framework for AI-driven molecular design of multi-target drugs: Brain diseases as a case study. In *2024 IEEE Congress on Evolutionary Computation (CEC)*, pages 1–8.

Chebrolu, A. (2023). DeepLig: A de-novo computational drug design approach to generate multi-targeted drugs. In *2023 Fifth International Conference on Transdisciplinary AI (TransAI)*, pages 34–41.

Danon, J. J., Reekie, T. A., and Kassiou, M. (2019). Challenges and opportunities in central nervous system drug discovery. *Trends in Chemistry*, 1(6):612–624.

Dowden, H. and Munro, J. (2019). Trends in clinical success rates and therapeutic focus. *Nature Reviews Drug Discovery*, 18(7):495–496.

Eberhardt, J., Santos-Martins, D., Tillack, A. F., and Forli, S. (2021). AutoDock Vina 1.2.0: New docking methods, expanded force field, and python bindings. *Journal of Chemical Information and Modeling*, 61(8):3891–3898.

Munson, B., Chen, M., Bogosian, A., et al. (2024). De novo generation of multi-target compounds using deep generative chemistry. *Nature Communications*, 15:3636.

Polykovskiy, D., Zhebrak, A., et al. (2020). Molecular Sets (MOSES): A Benchmarking Platform for Molecular Generation Models. *Frontiers in Pharmacology*.

Sun, J., Jeliazkova, N., et al. (2017). ExCAPE-DB: an integrated large scale dataset facilitating big data analysis in chemogenomics. *Journal of Cheminformatics*, 9(1):17.

Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., et al. (2017). Attention is all you need. In *Advances in Neural Information Processing Systems*, volume 30.