Multi-target optimization for drug discovery using generative models

Anirudh Jain [1] **Markus Heinonen** [1] **Samuel Kaski** [1 2]

Abstract

Graph generative models have been utilized for discovering novel molecules with desired chemical properties. Several of these methods use neural networks to map the latent space to chemical properties and explore the latent space efficiently. We propose using multi-target networks to jointly predict several molecular properties and learn better representations by exploiting auxiliary information. Our joint model outperforms existing methods in property prediction and molecular optimization tasks. We also propose a new benchmark to compare generative models for drug discovery.

1. Introduction

Drug discovery is a challenging and costly process with many challenges. Potential drug candidates should have specific biological activity while satisfying multiple reactivity and toxicity requirements. The first step in this process is hit finding. Millions of drug-like compounds are screened to find a few promising candidates for further optimization. Computational methods have been increasingly used to virtually screen large chemical spaces and find useful hits (Shoichet, 2004).

Recently, deep generative models have been utilized to accelerate the hit finding process and explore chemical latent spaces (Lavecchia, 2019). These models aim to discover novel molecules with desired chemical properties or optimize existing molecules to satisfy certain constraints. Existing state-of-the-art methods only focus on optimizing a single molecular property and are unlikely to perform as well in the real world setting (Zang & Wang, 2020; Shi et al., 2020; You et al., 2018).

We propose joint modelling of molecular properties from the chemical latent space. We use MoFlow (Zang & Wang, 2020) to learn the latent space over molecular graphs as the current SOTA method. We utilize multi-target neural networks to learn a joint representation over 200 molecular properties and show that the auxiliary information leads to better predictive model. We also show that the joint model improves results on existing molecular optimization tasks. Finally, we evaluate the performance of our joint model in discovering molecules with constraints over multiple properties as a more realistic benchmark for drug discovery.

2. Multi objective property prediction

Multi objective learning is a natural fit for predicting chemical properties. The relationship between properties can be utilized to train better predictive models(Ramsundar et al., 2015; Unterthiner et al., 2014). Several generative models utilize property prediction networks to map their chemical latent space to properties and discover new molecules (Honda et al., 2019; Madhawa et al., 2019; Zang & Wang, 2020). We propose multi objective setting to improve the performance of deep learning models in predicting molecular properties.

We use the ZINC 250k dataset (Irwin et al., 2012) and generate 212 regression targets using commercially available ADMET Predictor (Ghosh et al., 2016). We divide the dataset into trainset (200k), validation set (20k) and test set (30k). Validation set is used to pick optimal hyperparameters for all models. We use a residual deep model to predict properties from the latent vectors generated by a pretrained MoFlow model. The model architectures are described in Appendix A. We compare three models: the original MoFlow network used in (Zang & Wang, 2020), Single residual network to predict properties individually and Joint residual network to predict properties together.

We report the test set performance of all three models in Table 1. We compare the mean absolute error (MAE) and predictive correlation over chemically relevant properties. Joint predictive network outperforms all other models with a large improvement in complex properties such as QED, PLogP and Admet Risk. The multi objective setting allows us to significantly improve existing methods with a little cost by exploiting the relationship between molecular properties.

We evaluate the relationship between model performance and the number of tasks in Fig 1. We start with a joint

[1]Aalto University, Finland [2]University of Manchester, UK. Correspondence to: Anirudh Jain <anirudh.jain@aalto.fi>.

The 2021 ICML Workshop on Computational Biology.

Table 1. Mean absolute error(MAE) and correlation between predictive values and true values on the test set. We evaluate chemically relevant properties such as quantitative estimation of drug-likeness(QED), Penalized LogP(PLogP), ADMET Risk and Synthetic Accessibility Score(SA Score) along with reporting the mean performance over all properties.

Model	QED		PLogP		ADMET Risk		SA Score		All	
	MAE	Correlation	MAE	Correlation	MAE	Correlation	MAE	Correlation	MAE	Correlation
MoFlow	0.09	0.62	1.24	0.69	1.11	0.54	0.48	0.69	0.60 ± 0.23	0.68 ± 0.24
Single	0.07	0.77	0.96	0.81	1.07	0.60	0.42	0.76	0.27 ± 0.18	0.79 ± 0.20
Joint	0.05	0.83	0.65	0.88	0.79	0.77	0.33	0.83	0.24 ± 0.14	0.84 ± 0.16

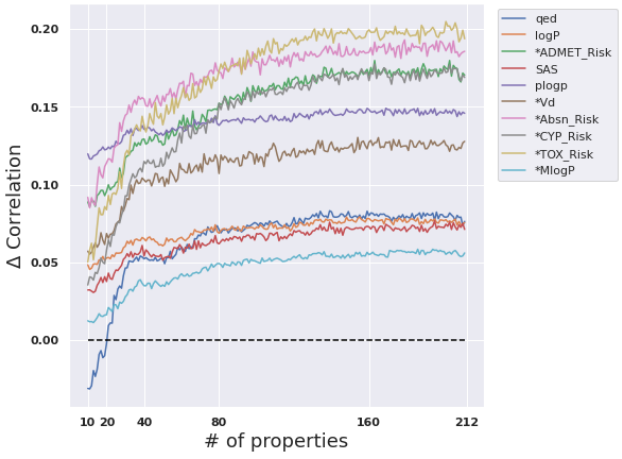

Figure 1. Test set growth curves. The y-axis shows the change in predictive correlation between joint predictive model and single predictive model as more properties are added. The properties are added in descending order of their correlation.

model trained over 10 properties of chemical relevance and add the highest correlated properties first. There is a sharp improvement initially in the joint model performance which gradually plateaus as more properties are added. The correlations between chemical properties provide a useful signal which is exploited by the deep network to significantly improve model performance.

3. Molecular optimization

The goal of molecular optimization tasks is to discover new molecules with desired properties. We explore the latent space using the trained predictive models by performing gradient ascent. Given a seed molecule M_0, we encode it via the MoFlow generative model $\mathbf{f}$ to get latent vector $\mathbf{z}_0 = \mathbf{f}(M_0)$. We then discover new latent vectors $\mathbf{z}_i$ via gradient ascent

$$\mathbf{z}_{i+1} = \mathbf{z}_i + \alpha \nabla_{\mathbf{z}} h(\mathbf{z}_i), \quad i = 0, \ldots, K-1 \qquad (1)$$

where h is the property prediction network and α is the step size. We obtain a new molecule by using the inverse transformation of the MoFlow generative model $M_i = \mathbf{f}^{-1}(\mathbf{z}_i)$ to get K candidate molecules. We compare the original Moflow property network, single predictive network and joint predictive network to discover new molecules with desired property. In all the experiments, we follow the same setting as (Zang & Wang, 2020).

3.1. Maximize QED

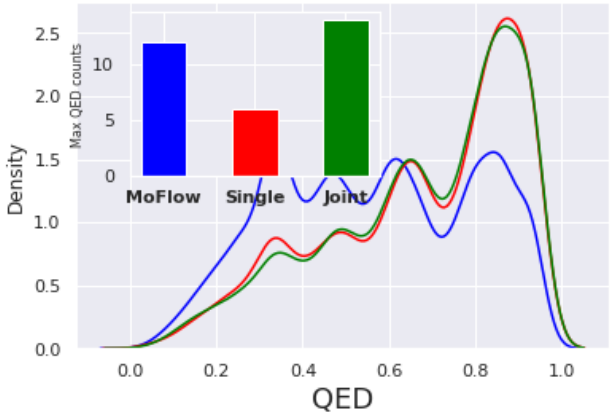

Figure 2. Novel molecules discovered by gradient ascent. Single and joint models discover a better distribution of novel molecules with a higher density at high QED scores. Joint model finds the most novel molecules with the best QED score.

The three models are compared in finding novel molecules with best possible quantitative estimate of drug-likeness (QED) scores (Kirkpatrick, 2012). The baseline MoFlow model is the current SOTA method in discovering molecules with the best QED score.

Fig 2 shows that the single and joint model discover a much better distribution of novel molecules than the original MoFlow model. The joint predictive model also discovers more novel molecules with the best QED score.

3.2. Constrained PLogP Optimization

We use the set of seed molecules used in (*Youet al.*, 2018) to find new molecules with improved Penalized LogP (PLogP) scores. The goal is to find molecules with better PLogP scores with molecular similarity greater than δ.

Table 2. Comparison of MoFlow, Single and Joint model for PLogP optimization. Joint model improves upon other approaches.

	GCPN			MoFlow		
δ	Improvement	Similarity	Success	Improvement	Similarity	Success
0.0	4.20 ± 1.28	0.32 ± 0.12	100%	3.32 ± 1.59	0.43 ± 0.26	96.50%
0.2	4.12 ± 1.19	0.34 ± 0.11	100%	2.84 ± 1.27	0.56 ± 0.18	94.00%
0.4	2.49 ± 1.30	0.48 ± 0.08	100%	2.68 ± 1.07	0.64 ± 0.12	89.12%
0.6	0.79 ± 0.63	0.68 ± 0.08	100%	2.35 ± 0.67	0.71 ± 0.06	70.62%
	Single			Joint		
δ	Improvement	Similarity	Success	Improvement	Similarity	Success
0.0	3.12 ± 1.48	0.46 ± 0.27	95.88%	4.01 ± 1.81	0.32 ± 0.24	97.88%
0.2	2.85 ± 1.22	0.59 ± 0.17	92.00%	3.36 ± 1.57	0.49 ± 0.19	94.50%
0.4	2.69 ± 1.02	0.65 ± 0.11	89.00%	2.87 ± 1.18	0.61 ± 0.12	90.88%
0.6	2.34 ± 0.68	0.71 ± 0.07	72.00%	2.42 ± 0.77	0.71 ± 0.07	73.00%

We use the Tanimoto index as the similarity measure (Bajusz et al., 2015) as used in previous works. Table 2 shows that the Joint model performs best for higher values of δ where the difficulty of finding novel molecules is increased. Surprisingly, the single model performs worse than the MoFlow baseline.

Fig 2 and Table 2 evaluate the performance of joint predictive model on existing molecular optimization tasks. The joint model learns a better representation over molecular properties and is more efficient in finding novel molecules with desired properties. Single predictive model is surprisingly worse than MoFlow baseline at exploration. We suspect that overfitting observed in the single model leads to poor generalization during exploration.

3.3. Joint Optimization

Table 3. Single target optimization maximizes the QED score with similarity greater than δ. Joint optimization maximizes the QED and PLogP while minimizing the ADMET risk

	Single target optimization			Joint target optimization		
δ	Δ QED	Δ PLogP	Δ ADMET	Δ QED	Δ PLogP	Δ ADMET
0.0	0.22 ± 0.15	-2.66 ± 4.55	3.61 ± 23.29	0.19 ± 0.15	1.50 ± 1.98	-1.40 ± 1.11
0.2	0.20 ± 0.14	-2.00 ± 3.37	-0.01 ± 2.88	0.18 ± 0.14	1.68 ± 1.95	-1.36 ± 1.04
0.4	0.15 ± 0.12	-1.30 ± 2.32	-0.27 ± 2.62	0.12 ± 0.12	1.73 ± 1.77	-1.03 ± 0.95
0.6	0.07 ± 0.03	-0.37 ± 0.35	-1.53 ± 0.83	0.03 ± 0.03	1.16 ± 0.64	-1.05 ± 0.87

We propose a novel benchmark to evaluate generative models in a more realistic setting. We aim to jointly optimize multiple property targets and find new molecules that satisfy multiple constraints. The goal of the joint optimization task is to maximize the QED and PLogP while simultaneously minimizing ADMET Risk. Additionally, the new molecules should have similarity greater than δ with the seed molecule.

Table 3 compares only optimizing QED and jointly optimizing all three properties jointly. As QED is improved PLogP goes down while ADMET Risk reduces for higher values of δ. With joint optimization, all three properties are in the desired range with QED and PLogP increasing and ADMET Risj decreasing for all values of δ.

4. Conclusion

In this paper, we propose using multi objective setting for drug discovery. We show that existing SOTA methods can significantly be improved by jointly predicting molecular properties. The relationship between properties can be exploited by deep networks to train better models and more efficiently perform molecular optimization. We also propose a new benchmark that is more realistic test of generative models for drug discovery. Jointly optimizing molecular properties to maximize desirable properties and minimize harmful properties allow us to naturally extend the multi objective setting.

Acknowledgements

We acknowledge the computational resources provided by the Aalto Science-IT project and CSC – IT Center for Sci-

ence, Finland. We would also like to thank Julius Sipilä and Heikki Käsnänen for useful discussions.

References

Bajusz, D., Rácz, A., and Héberger, K. Why is tanimoto index an appropriate choice for fingerprint-based similarity calculations? *Journal of cheminformatics*, 7(1):1–13, 2015.

Ghosh, J., Lawless, M. S., Waldman, M., Gombar, V., and Fraczkiewicz, R. Modeling admet. In *In Silico Methods for Predicting Drug Toxicity*, pp. 63–83. Springer, 2016.

He, K., Zhang, X., Ren, S., and Sun, J. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 770–778, 2016.

Honda, S., Akita, H., Ishiguro, K., Nakanishi, T., and Oono, K. Graph residual flow for molecular graph generation. *arXiv preprint arXiv:1909.13521*, 2019.

Irwin, J. J., Sterling, T., Mysinger, M. M., Bolstad, E. S., and Coleman, R. G. Zinc: a free tool to discover chemistry for biology. *Journal of chemical information and modeling*, 52(7):1757–1768, 2012.

Kirkpatrick, P. Shades of chemical beauty. *Nature Reviews Drug Discovery*, 11(2):107–107, 2012.

Lavecchia, A. Deep learning in drug discovery: opportunities, challenges and future prospects. *Drug discovery today*, 24(10):2017–2032, 2019.

Madhawa, K., Ishiguro, K., Nakago, K., and Abe, M. Graphnvp: An invertible flow model for generating molecular graphs. *arXiv preprint arXiv:1905.11600*, 2019.

Ramsundar, B., Kearnes, S., Riley, P., Webster, D., Konerding, D., and Pande, V. Massively multitask networks for drug discovery. *arXiv preprint arXiv:1502.02072*, 2015.

Shi, C., Xu, M., Zhu, Z., Zhang, W., Zhang, M., and Tang, J. Graphaf: a flow-based autoregressive model for molecular graph generation. *arXiv preprint arXiv:2001.09382*, 2020.

Shoichet, B. K. Virtual screening of chemical libraries. *Nature*, 432(7019):862–865, 2004.

Unterthiner, T., Mayr, A., Klambauer, G., Steijaert, M., Wegner, J. K., Ceulemans, H., and Hochreiter, S. Deep learning as an opportunity in virtual screening. In *Proceedings of the deep learning workshop at NIPS*, volume 27, pp. 1–9, 2014.

You, J., Liu, B., Ying, R., Pande, V., and Leskovec, J. Graph convolutional policy network for goal-directed molecular graph generation. *arXiv preprint arXiv:1806.02473*, 2018.

Zang, C. and Wang, F. Moflow: an invertible flow model for generating molecular graphs. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pp. 617–626, 2020.

A. Model Architectures

Table 4. Original MoFlow Predictive Network

Layer type	Activation	Output size
FC	ReLU	M x 6156
FC	ReLU	M x 16
FC	ReLU	M x 1

Table 5. Residual Block

Layer type	Activation	Output size
FC	-	M x N
BN	ReLU	M x N
FC	-	M x N
BN	-	M x N
Summation	ReLU	M x N

Table 6. Residual Predictive Network where T is the number of targets.

Layer type	Activation	Output size
Input	-	M x 6156
Residual	-	M x 6156
FC	ReLU	M x 128
Residual	-	M x 128
FC	-	M x T

Table 4 defines the model architecture used by (Zang & Wang, 2020) to predict molecular properties. We use a more complex network inspired by the success of residual neural network in other domains (He et al., 2016). We define the residual block we use in Table 5 and the model architecture used for single and joint models in Table 6. Here M is the mini-batch size, N is the size of model inputs and T is the number of targets predicted by the deep model. Single model use $T = 1$ and joint model use $T = 212$.