

---

# Cell Morphology-Guided Small Molecule Generation with GFlowNets

---

Stephen Zhewen Lu<sup>1,2</sup> Ziqing Lu<sup>3</sup> Ehsan Hajiramezani<sup>3</sup> Tommaso Biancalani<sup>3</sup> Yoshua Bengio<sup>1,4</sup>  
Gabriele Scalia<sup>3</sup> Michał Koziarski<sup>1,4</sup>

## Abstract

High-content phenotypic screening, including high-content imaging (HCI), has gained popularity in the last few years for its ability to characterize novel therapeutics without prior knowledge of the protein target. This work focuses on the novel task of HCI-guided molecular design. We consider an approach in which we leverage an unsupervised multimodal joint embedding to define a latent similarity as a reward for GFlowNets. The proposed model learns to generate new molecules that could produce phenotypic effects similar to those of the given image target, without relying on pre-annotated phenotypic labels. We demonstrate that our method generates molecules with high morphological and structural similarity to the target, increasing the likelihood of similar biological activity.

## 1. Introduction

Phenotypic screening allows assessing drug efficacy based on observed biological effects, without detailed knowledge of the underlying mode of action. High-content phenotypic screening, especially through high-content imaging (HCI) (Bray et al., 2016), provides rich data such as morphological changes in cell shape and structure. This can potentially elucidate broad biological effects, targets, and mechanisms of action (Moffat et al., 2017). Despite its potential to guide the drug discovery process, challenges persist in effectively utilizing this data. One common approach to utilize this information is training a supervised classification model to recognize the phenotype of interest, and then use it to virtually screen existing libraries (Krentzel et al., 2023). However, this can require a substantial amount of labeled data and prior knowledge of the phenotype of interest, often

---

<sup>1</sup>Mila – Québec AI Institute <sup>2</sup>McGill University <sup>3</sup>Biology Research | AI Development (BRAID), Genentech <sup>4</sup>Université de Montréal. Correspondence to: Michał Koziarski <michal.koziarski@mila.quebec>.

Accepted by the *Structured Probabilistic Inference & Generative Modeling workshop* of ICML 2024, Vienna, Austria. Copyright 2024 by the author(s).

unavailable for poorly understood diseases. A preferred scenario would be to have a method capable of designing new compounds based on only a few, or even a single, target morphological readout. Another challenge lies in the screening process itself. Virtual screening limits the search space to existing screening libraries, which are significantly smaller than the whole drug-like space, often estimated to contain more than  $10^{60}$  (Lipinski et al., 1997) molecules. This limitation can become particularly important when structurally and functionally novel molecules are desirable, for example, to improve the potency and diversity of the leads or to overcome unwanted secondary effects (Jain et al., 2022; Gao et al., 2024; Song & Li, 2023; Ghari et al., 2023).

In this paper, we address both of these challenges by proposing a generative method guided by the target cell morphology. Specifically, we propose a reward function that utilizes the latent similarity between the generated molecule and the target morphology. For this, we utilize a multi-modal contrastive learning model that aligns small molecules to morphological readouts. Then, we use this reward to train a Generative Flow Network (GFlowNet), as recently proposed by Bengio et al. (2023), to generate molecules capable of inducing similar morphological outcomes. We demonstrate that our approach is capable of generating diverse molecules with high latent similarity to the provided morphology, which translates into a higher likelihood of obtaining similar predicted biological activity. We also show that the generation process can be structurally conditioned by using joint latent embeddings, which combine both target readout and molecule, further improving the performance. Practical use cases of the proposed method include generating molecules that induce a desired cell morphology obtained through gene perturbations (Rohban et al., 2022), scaffold hopping, where novel molecular structures with similar effects to a target one are desired (Hu et al., 2017), and, more generally, molecular design guided by phenotypic readouts (Krentzel et al., 2023).

## 2. Related work

**Generative models for drug discovery.** There is a plethora of methods for molecular generation in the literature (Meyers et al., 2021; Bilodeau et al., 2022). They can be broadly

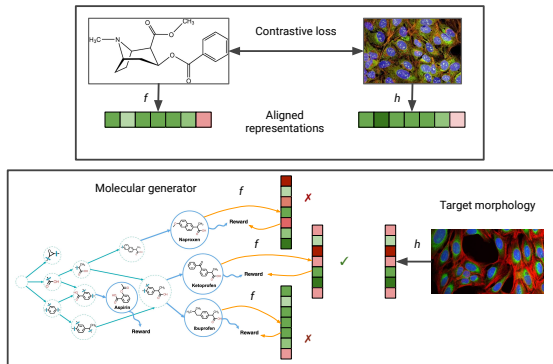


Figure 1. Graphical illustration of the proposed approach. In the first stage (top), cross-modal contrastive learning is used to train latent encoders  $f$  and  $h$  that produce aligned representations between molecules and readouts. Then, in the second stage (bottom), the target morphology readout is first converted into a latent vector, and the similarity between that and latents from molecules output by a generator is used as a reward. The model learns to sample molecules capable of inducing similar latents to the target.

categorized based on the molecular representation used: textual representation such as SMILES (Kang & Cho, 2018; Arús-Pous et al., 2020; Kotsias et al., 2020), molecular graphs (Jin et al., 2018; Maziarka et al., 2020; Pedawi et al., 2022; Diamant et al., 2023) or 3D atom coordinate representations (O Pinheiro et al., 2024); as well as the underlying methodology, e.g. variational autoencoders (Jin et al., 2018; Maziarka et al., 2020), reinforcement learning (Loeffler et al., 2024) or diffusion models (Runcie & Mey, 2023; Uehara et al., 2024). These methods have found applications in drug discovery, reporting successes in several areas such as immunology and infectious diseases (Godinez et al., 2022; Moret et al., 2023). Recently, Generative Flow Networks (GFlowNets) (Bengio et al., 2021; Nica et al., 2022; Shen et al., 2023; Volokhova et al., 2024; Koziarski et al., 2024a; Gaiński et al., 2024; Koziarski et al., 2024b) emerged as a promising paradigm for molecular generation due to the ability to sample diverse candidate molecules, which is crucial in the drug discovery process. Importantly, similar to RL, GFlowNets can be trained based on the specified reward function, which makes them suitable for phenotypic discovery. Compared to existing methods, which focus on conditional generation based on a single property or multiple properties of interest, we tackle generation guided by high-content readouts, which we achieve through a multi-modal latent joint representation.

## 3. Method

### 3.1. Generative Flow Networks

GFlowNets are amortized variational inference algorithms that are trained to sample from an unnormalized target distribution over compositional objects. GFlowNets aim to sam-

ple objects from a set of terminal states  $\mathcal{X}$  proportionally to a reward function  $\mathcal{R} : X \rightarrow \mathbb{R}^+$ . GFlowNets are defined on a pointed directed acyclic graph (DAG),  $G = (S, A)$ , where:

- $s \in S$  are the nodes, referred to as states in our setting, with the special starting state  $s_0$  being the only state with no incoming edges, and the terminal states  $\mathcal{X}$  have no outgoing edges,
- $a = s \rightarrow s' \in A$  are the edges, referred to as actions in our setting, and correspond to applying an action while in a state  $s$  and landing in state  $s'$ .

A state sequence  $\tau = (s_0 \rightarrow s_1 \rightarrow \dots \rightarrow s_n = x)$ , with  $s_n = x \in \mathcal{X}$  and  $a_i = (s_i \rightarrow s_{i+1}) \in A$  for all  $i$ , is called a complete trajectory. We denote the set of trajectories as  $\mathcal{T}$ .

**Trajectory balance.** Several training losses were explored to train a GFlowNet. Among those, trajectory balance (Malkin et al., 2022a) has been shown to improve credit assignment. In addition to learning a forward policy  $P_F$ , we also learn a backwards policy  $P_B$  and a scalar  $Z_\theta$ , such that, for every trajectory  $\tau = (s_0 \rightarrow s_1 \rightarrow \dots \rightarrow s_n = x)$ , they satisfy  $Z_\theta \prod_{t=1}^n P_F(s_t | s_{t-1}) = R(x) \prod_{t=1}^n P_B(s_{t-1} | s_t)$ .

### 3.2. Multi-modal contrastive learning

Contrastive learning is a self-supervised approach that learns embeddings by maximizing agreement between similar samples and minimizing it between dissimilar ones, using contrastive loss functions like InfoNCE (Oord et al., 2018). Multi-modal contrastive learning effectively learns joint representations from diverse data modalities. A notable instance of this approach is CLIP (Radford et al., 2021), which aligns textual descriptions with visual representations, enabling robust cross-modal understanding. We leverage a multi-modal contrastive model to learn a joint embedding of molecules and molecular effects. This choice avoids high-dimensional supervised losses and promotes learning “informative” features for the task (i.e., features that relate the two modalities to each other).

Let  $\{(x_i, y_i)\}_{i=1}^N$  be a batch of  $N$  pairs of molecular graphs ( $x_i$ ) and their corresponding morphology images ( $y_i$ ). Let  $f$  and  $h$  be the molecular and morphology encoders, respectively. Let the similarity between the molecular graph and morphology image embeddings be defined as  $S_{ij} = \exp\left(\frac{\text{sim}(f(x_i), h(y_j))}{\tau}\right)$ , where  $\text{sim}(f(x_i), h(y_j))$  denotes the cosine similarity between embeddings. The CLIP loss is defined as follows:

$$L_{\text{CLIP}} = \frac{1}{N} \sum_{i=1}^N \left[ -\log \frac{S_{ii}}{\sum_{j=1}^N S_{ij}} - \log \frac{S_{ii}}{\sum_{j=1}^N S_{ji}} \right], \quad (1)$$

where  $\tau$  is a temperature parameter.

Instead of relying on CLIP, in this work, we leverage the closely related Geometric Multimodal Contrastive (GMC) loss (Poklukar et al., 2022). GMC leverages the same loss function in (1) to align modality-specific encoders to a joint encoder, which takes as input all modalities. Therefore, in addition to providing modality-specific embeddings  $f$  and  $h$ , it also provides a *joint* embedding  $fh$  that we leverage when both the target molecule and its readout are available.

### 3.3. GFlowNets for morphology-guided molecular design

The proposed approach exploits recent developments in multi-modal contrastive learning and molecular generation with GFlowNets into a unified pipeline, as illustrated in Figure 1. The method relies on first training a contrastive learning model capable of producing aligned latent representations, and then using these representations as a guiding signal for GFlowNet. We define the reward function for the GFlowNet from the embeddings of the GMC model as described in Section 3.2:

$$R(x|y) = 1 + \frac{f(x)h(y)}{2\|f(x)\| \|h(y)\|}. \quad (2)$$

This normalized cosine similarity between the target morphology latent and the generated structure latent is crucial to enforce the non-negativity of the GFlowNet reward. During the training of the GMC model, we impose early-stopping using the correlation between the cross-modality distance metric. We observe that early-stopping reduces the variance of cosine similarity between the multimodal GMC embeddings, while it does not affect the performance of GFlowNet (see Appendix A.4 for more details). Inspired by using replay buffer in reinforcement learning (Vemgal et al., 2023), we leverage replay buffer with known decomposed structure when training on joint morphology and structure-guided generations. This increases the structural similarity of generated samples to the known target.

## 4. Experimental study

We evaluate the proposed approach, first by examining the quality and diversity of generated samples, and then their predicted biological activity in downstream tasks. Experiments verifying the underlying assumption of our method, namely that the similarity in latent space produced by the contrastive model correlates with the morphological similarity, can be found in Appendix A.1. Experiments on structurally conditioned generation, where we base our target latent on the combination of target morphology and the associated molecular structure, can be found in Appendix A.2.

We perform experiments on the Cell Painting dataset introduced by (Bray et al., 2016; 2017), as further pre-processed




Figure 2. Comparison of examined methods in terms of reward optimization: a) distribution of rewards from generated samples, 10,000 samples generated by a trained model (left) and b) number of modes discovered (right). Both plots aggregate across all examined targets.

in (Moshkov et al., 2023). The dataset includes 16,170 molecules and associated cell morphology images. Each image includes five color channels that describe the morphology of five cellular compartments. Images have been pre-processed using CellProfiler 3.0 (McQuin et al., 2018; Moshkov et al., 2023).

Additionally, to support the evaluation of the generated molecules, we leverage oracle models independently trained on data from multiple assays released by the Broad Institute (Moshkov et al., 2023) (see Appendix B.3 for more details). This allows us to evaluate the ability of the model to generate molecules with biochemical and cellular effects similar to the (unknown) target molecule.

### 4.1. Generating high reward and diverse samples

We are interested in evaluating the capabilities of a GFlowNet in optimizing the specified reward function. Given our primary focus on its application in the initial discovery stage, it is essential to generate not only high-reward outcomes but also a diverse set of samples.

We compare GFlowNet with random sampling (RND) and two standard RL-based baselines: soft Q-learning (SQL) (Haarnoja et al., 2017) and soft actor-critic (SAC) (Haarnoja et al., 2018). Note that since, to the best of our knowledge, this is the first published attempt at guiding the generative molecular model with expected image morphology outcome, we focus specifically on benchmarking against other potential molecular generation methods.

We show the distribution of rewards for generated samples and the number of discovered modes (defined as molecules with reward  $\geq 90$ th percentile and Tanimoto similarity to other modes  $\leq 0.3$ ) in Figure 2, and the distribution of similarities between the top-100 generated samples in Figure 3 (each figure aggregated across all considered targets). As shown, GFlowNet learns to sample high-reward candidate molecules (with a significantly higher average reward than random sampling and SQL, comparable to SAC) while also significantly improving the diversity compared to SAC (with a lower similarity between top candidates). Both of the above translate into a significantly higher number of




Figure 3. Comparison of examined methods in terms of diversity: a) distribution of Tanimoto similarities between top-100 generated samples (left; aggregated across all examined targets, lower is better) and b) structural diversity of top-1000 highest reward samples to Target #8636 (right).

Table 1. Max. Tanimoto similarity to the target in the last 10,000 samples, averaged across all targets.

	Random	Soft Actor-Critic	Soft Q-Learning	GFlowNet
Morphology Target	0.305 ( $\pm$ 0.057)	0.261 ( $\pm$ 0.068)	0.329 ( $\pm$ 0.079)	<b>0.337 (<math>\pm</math> 0.092)</b>
Joint Target	0.311 ( $\pm$ 0.064)	0.388 ( $\pm$ 0.163)	0.309 ( $\pm$ 0.064)	<b>0.451 (<math>\pm</math> 0.163)</b>

discovered modes than the baseline methods. The diversity of the generated samples is further illustrated in Figure 3, where a UMAP visualization of the molecular structures produced by different methods is presented for a specific target. As can be seen, GFlowNet displays sample coverage similar to random sampling, which is a desirable outcome.

## 4.2. Biological activity estimation

So far, we have established that the proposed approach is capable of generating diverse samples with high reward and that there is a moderate correlation between the reward and the similarity to the target. However, what is critical in the end is whether this will translate into generated molecules inducing similar biological effects to the original target perturbation (ground truth). Ideally, we wish to evaluate this based on experiments comparing generated molecules and the ground truth, but this can be costly and time-consuming. Instead, we estimate the similarity of the biological effect based on the available data. We consider two approaches.

First of all, we would expect a perfectly optimized generator to be able to sample known molecules that induced target morphological profiles. In practice, this might be unattainable: not only is the problem itself heavily under-constrained (we expect multiple molecules, likely a very large number, to be able to induce a given morphological outcome), but also our morphological similarity estimation is intrinsically noisy. Because of the above, what we try to achieve is the highest possible structural similarity of generated samples to the known molecule that induced the given morphology. The maximum Tanimoto similarity to the known molecular target, averaged across all considered targets, is presented in Table 1. As shown, the proposed approach generally recovers the underlying targets more effectively, suggesting its utility in identifying molecules with expected biological activity. Unsurprisingly, this effect




Figure 4. Number of samples with predicted assay probability  $\geq$  0.7 in the top 1000 modes with Tanimoto similarity  $\leq$  0.3. Results are aggregated across multiple targets.

is particularly pronounced when conditioning on specific molecular structures.

The second approach we consider involves utilizing oracle models for predicting biological activity. We train an MLP using molecular structures, specifically their extracted molecular fingerprints, as inputs to predict the outcomes of biological assays (details of the model training are provided in Appendix B.4). For each target molecule, at least one associated assay has a positive outcome. The objective is to generate molecules with a high predicted probability of producing a positive outcome in the specific assay for which the target molecule has known activity. The number of generated samples with high predicted assay probability ( $\geq$  0.7) is presented in Figure 4. Note that, due to the high uncertainty of the oracle model, we are primarily interested in quantifying the number of high-likelihood samples. As can be seen, using guided generation helps improve the proportion of generated molecules with high predicted assay probability, which serves as a proxy for generating molecules with similar biological activity. It is worth noting that while we observe a higher predicted probability for SAC when considering the top-1000 molecules by reward, this trend reverses when considering different modes, once again highlighting the higher diversity combined with high reward of GFlowNet samples. It is also worth noting that we observe a large variance across targets. While this can be partially attributed to the uncertainty of the oracle, further investigation of factors determining whether the proposed approach is helpful or not for a given target is an important future research direction.

## 5. Conclusions

In this paper, we consider the task of designing a generative model able to produce molecules that can induce a cell morphology profile similar to a given target. Such framework is broadly applicable, for example to design drugs mimicking the effect of a genetic perturbation, designing drug analogs, or, more generally, molecular design guided by phenotypic readouts. The proposed approach relies on GFlowNet for molecular generation, and uses a reward based on the latent similarity of representations from a multi-modal contrastive



learning model. To the best of our knowledge, this is the first published attempt at the challenging task of guiding the generative molecular model with the expected image morphology outcome.

## Acknowledgements

The research was supported by funding from CQDM Fonds d'Accélération des Collaborations en Santé (FACS) / Acuité Québec and Genentech.

## References

- Arús-Pous, J., Patronov, A., Bjerrum, E. J., Tyrchan, C., Reymond, J.-L., Chen, H., and Engkvist, O. Smiles-based deep generative scaffold decorator for de-novo drug design. *Journal of cheminformatics*, 12:1–18, 2020.
- Bengio, E., Jain, M., Korablyov, M., Precup, D., and Bengio, Y. Flow network based generative models for non-iterative diverse candidate generation. *Advances in Neural Information Processing Systems*, 34:27381–27394, 2021.
- Bengio, Y., Lahlou, S., Deleu, T., Hu, E. J., Tiwari, M., and Bengio, E. GFlowNet foundations. *Journal of Machine Learning Research*, 24(210):1–55, 2023.
- Bilodeau, C., Jin, W., Jaakkola, T., Barzilay, R., and Jensen, K. F. Generative models for molecular discovery: Recent advances and challenges. *Wiley Interdisciplinary Reviews: Computational Molecular Science*, 12(5):e1608, 2022.
- Bray, M.-A., Singh, S., Han, H., Davis, C. T., Borgeson, B., Hartland, C., Kost-Alimova, M., Gustafsdottir, S. M., Gibson, C. C., and Carpenter, A. E. Cell painting, a high-content image-based assay for morphological profiling using multiplexed fluorescent dyes. *Nature protocols*, 11(9):1757–1774, 2016.
- Bray, M.-A., Gustafsdottir, S. M., Rohban, M. H., Singh, S., Ljosa, V., Sokolnicki, K. L., Bittker, J. A., Bodycombe, N. E., Dančák, V., Hasaka, T. P., et al. A dataset of images and morphological profiles of 30 000 small-molecule treatments using the cell painting assay. *Gigascience*, 6(12):giw014, 2017.
- Chandrasekaran, S. N., Ceulemans, H., Boyd, J. D., and Carpenter, A. E. Image-based profiling for drug discovery: due for a machine-learning upgrade? *Nature Reviews Drug Discovery*, 20(2):145–159, 2021.
- Diamant, N. L., Tseng, A. M., Chuang, K. V., Biancalani, T., and Scalia, G. Improving graph generation by restricting graph bandwidth. In *International Conference on Machine Learning*, pp. 7939–7959. PMLR, 2023.
- Gaiński, P., Koziarski, M., Maziarz, K., Segler, M., Tabor, J., and Śmieja, M. Retrofn: Diverse and feasible retrosynthesis using gflownets. *arXiv preprint arXiv:2406.18739*, 2024.
- Gao, Z., Tan, C., Zhang, Y., Chen, X., Wu, L., and Li, S. Z. Proteininvbench: Benchmarking protein inverse folding on diverse tasks, models, and metrics. *Advances in Neural Information Processing Systems*, 36, 2024.
- Gavriilidis, G. I., Vasileiou, V., Orfanou, A., Ishaque, N., and Psomopoulos, F. A mini-review on perturbation modelling across single-cell omic modalities. *Computational and Structural Biotechnology Journal*, 2024.
- Ghari, P. M., Tseng, A., Eraslan, G., Lopez, R., Biancalani, T., Scalia, G., and Hajiramezani, E. Generative flow networks assisted biological sequence editing. In *NeurIPS 2023 Generative AI and Biology (GenBio) Workshop*, 2023.
- Godinez, W. J., Ma, E. J., Chao, A. T., Pei, L., Skewes-Cox, P., Canham, S. M., Jenkins, J. L., Young, J. M., Martin, E. J., and Guiguemde, W. A. Design of potent antimalarials with generative chemistry. *Nature Machine Intelligence*, 4(2):180–186, 2022.
- Haarnoja, T., Tang, H., Abbeel, P., and Levine, S. Reinforcement learning with deep energy-based policies. In *International conference on machine learning*, pp. 1352–1361. PMLR, 2017.
- Haarnoja, T., Zhou, A., Abbeel, P., and Levine, S. Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. In *International conference on machine learning*, pp. 1861–1870. PMLR, 2018.
- Hetzel, L., Boehm, S., Kilbertus, N., Günnemann, S., Theis, F., et al. Predicting cellular responses to novel drug perturbations at a single-cell resolution. *Advances in Neural Information Processing Systems*, 35:26711–26722, 2022.
- Hu, Y., Stumpfe, D., and Bajorath, J. Recent advances in scaffold hopping: miniperspective. *Journal of medicinal chemistry*, 60(4):1238–1246, 2017.
- Jain, M., Bengio, E., Hernandez-Garcia, A., Rector-Brooks, J., Dossou, B. F., Ekbote, C. A., Fu, J., Zhang, T., Kilgour, M., Zhang, D., et al. Biological sequence design with gflownets. In *International Conference on Machine Learning*, pp. 9786–9801. PMLR, 2022.
- Jin, W., Barzilay, R., and Jaakkola, T. Junction tree variational autoencoder for molecular graph generation. In *International conference on machine learning*, pp. 2323–2332. PMLR, 2018.

- Kang, S. and Cho, K. Conditional molecular design with deep generative models. *Journal of chemical information and modeling*, 59(1):43–52, 2018.
- Kotsias, P.-C., Arús-Pous, J., Chen, H., Engkvist, O., Tyrchan, C., and Bjerrum, E. J. Direct steering of de novo molecular generation with descriptor conditional recurrent neural networks. *Nature Machine Intelligence*, 2(5): 254–265, 2020.
- Koziarski, M., Abukalam, M., Shah, V., Vaillancourt, L., Schuetz, D. A., Jain, M., van der Sloot, A. M., Bourgey, M., Marinier, A., and Bengio, Y. Towards dna-encoded library generation with gflownets. In *ICLR 2024 Workshop on Generative and Experimental Perspectives for Biomolecular Design*, 2024a.
- Koziarski, M., Rekes, A., Shevchuk, D., van der Sloot, A., Gaiński, P., Bengio, Y., Liu, C.-H., Tyers, M., and Batey, R. A. Rgfn: Synthesizable molecular generation using gflownets. *arXiv preprint arXiv:2406.08506*, 2024b.
- Krentzel, D., Shorte, S. L., and Zimmer, C. Deep learning in image-based phenotypic drug discovery. *Trends in Cell Biology*, 33(7):538–554, 2023.
- Lipinski, C. A., Lombardo, F., Dominy, B. W., and Feeney, P. J. Experimental and computational approaches to estimate solubility and permeability in drug discovery and development settings. *Advanced drug delivery reviews*, 23(1-3):3–25, 1997.
- Loeffler, H. H., He, J., Tibo, A., Janet, J. P., Voronov, A., Mervin, L. H., and Engkvist, O. Reinvent 4: Modern ai-driven generative molecule design. *Journal of Cheminformatics*, 16(1):20, 2024.
- Lotfollahi, M., Wolf, F. A., and Theis, F. J. scgen predicts single-cell perturbation responses. *Nature methods*, 16(8):715–721, 2019.
- Malkin, N., Jain, M., Bengio, E., Sun, C., and Bengio, Y. Trajectory balance: Improved credit assignment in gflownets. *Advances in Neural Information Processing Systems*, 35:5955–5967, 2022a.
- Malkin, N., Lahlou, S., Deleu, T., Ji, X., Hu, E., Everett, K., Zhang, D., and Bengio, Y. Gflownets and variational inference. *arXiv preprint arXiv:2210.00580*, 2022b.
- Maziarka, Ł., Pocha, A., Kaczmarczyk, J., Rataj, K., Danel, T., and Warchoń, M. Mol-cycleGAN: a generative model for molecular optimization. *Journal of Cheminformatics*, 12(1):1–18, 2020.
- McQuin, C., Goodman, A., Chernyshev, V., Kamentsky, L., Cimini, B. A., Karhohs, K. W., Doan, M., Ding, L., Rafelski, S. M., Thirstrup, D., et al. Cellprofiler 3.0: Next-generation image processing for biology. *PLoS biology*, 16(7):e2005970, 2018.
- Meyers, J., Fabian, B., and Brown, N. De novo molecular design and generative models. *Drug Discovery Today*, 26(11):2707–2715, 2021.
- Moffat, J. G., Vincent, F., Lee, J. A., Eder, J., and Prunotto, M. Opportunities and challenges in phenotypic drug discovery: an industry perspective. *Nature reviews Drug discovery*, 16(8):531–543, 2017.
- Moret, M., Pachon Angona, I., Cotos, L., Yan, S., Atz, K., Brunner, C., Baumgartner, M., Grisoni, F., and Schneider, G. Leveraging molecular structure and bioactivity with chemical language models for de novo drug design. *Nature Communications*, 14(1):114, 2023.
- Moshkov, N., Becker, T., Yang, K., Horvath, P., Dancik, V., Wagner, B. K., Clemons, P. A., Singh, S., Carpenter, A. E., and Caicedo, J. C. Predicting compound activity from phenotypic profiles and chemical structures. *Nature Communications*, 14(1):1967, 2023.
- Nguyen, C. Q., Pertusi, D., and Branson, K. M. Molecule-morphology contrastive pretraining for transferable molecular representation. *ICML 2023 Workshop on Computational Biology*, pp. 2023–05, 2023.
- Nica, A. C., Jain, M., Bengio, E., Liu, C.-H., Korablyov, M., Bronstein, M. M., and Bengio, Y. Evaluating generalization in gflownets for molecule design. In *ICLR 2022 Machine Learning for Drug Discovery*, 2022.
- O Pinheiro, P. O., Rackers, J., Kleinhenz, J., Maser, M., Mahmood, O., Watkins, A., Ra, S., Sresht, V., and Saremi, S. 3d molecule generation by denoising voxel grids. *Advances in Neural Information Processing Systems*, 36, 2024.
- Oord, A. v. d., Li, Y., and Vinyals, O. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*, 2018.
- Palma, A., Theis, F. J., and Lotfollahi, M. Predicting cell morphological responses to perturbations using generative modeling. *bioRxiv*, pp. 2023–07, 2023.
- Pedawi, A., Gniewek, P., Chang, C., Anderson, B., and van den Bedem, H. An efficient graph generative model for navigating ultra-large combinatorial synthesis libraries. *Advances in Neural Information Processing Systems*, 35:8731–8745, 2022.
- Piran, Z., Cohen, N., Hoshen, Y., and Nitzan, M. Disentanglement of single-cell data with biolord. *Nature Biotechnology*, pp. 1–6, 2024.

- Poklukar, P., Vasco, M., Yin, H., Melo, F. S., Paiva, A., and Kragic, D. Geometric multimodal contrastive representation learning. In *International Conference on Machine Learning*, pp. 17782–17800. PMLR, 2022.
- Radford, A., Kim, J. W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pp. 8748–8763. PMLR, 2021.
- Rohban, M. H., Fuller, A. M., Tan, C., Goldstein, J. T., Syangtan, D., Gutnick, A., DeVine, A., Nijsure, M. P., Rigby, M., Sacher, J. R., et al. Virtual screening for small-molecule pathway regulators by image-profile matching. *Cell systems*, 13(9):724–736, 2022.
- Runcie, N. T. and Mey, A. S. Silvr: Guided diffusion for molecule generation. *Journal of Chemical Information and Modeling*, 63(19):5996–6005, 2023.
- Sanchez-Fernandez, A., Rumetshofer, E., Hochreiter, S., and Klambauer, G. Contrastive learning of image-and structure-based representations in drug discovery. *ICLR 2022 Machine Learning for Drug Discovery*, 2022.
- Shen, T., Pandey, M., and Ester, M. Tacogfn: Target conditioned gflownet for structure-based drug design. *arXiv preprint arXiv:2310.03223*, 2023.
- Song, Z. and Li, L. Importance weighted expectation-maximization for protein sequence design. In *International Conference on Machine Learning*, pp. 32349–32364. PMLR, 2023.
- Uehara, M., Zhao, Y., Black, K., Hajiramezanali, E., Scalia, G., Diamant, N. L., Tseng, A. M., Levine, S., and Biancalani, T. Feedback efficient online fine-tuning of diffusion models. *arXiv preprint arXiv:2402.16359*, 2024.
- Veličković, P., Cucurull, G., Casanova, A., Romero, A., Lio, P., and Bengio, Y. Graph attention networks. *arXiv preprint arXiv:1710.10903*, 2017.
- Vemgal, N. M., Lau, E., and Precup, D. An empirical study of the effectiveness of using a replay buffer on mode discovery in GFlowNets. In *ICML 2023 Workshop on Structured Probabilistic Inference & Generative Modeling*, 2023. URL <https://openreview.net/forum?id=pBk1cRPKBv>.
- Volokhova, A., Koziarski, M., Hernández-García, A., Liu, C.-H., Miret, S., Lemos, P., Thiede, L., Yan, Z., Aspuru-Guzik, A., and Bengio, Y. Towards equilibrium molecular conformation generation with gflowNets. *Digital Discovery*, 2024.
- Zheng, S., Rao, J., Zhang, J., Cohen, E., Li, C., and Yang, Y. Cross-modal graph contrastive learning with cellular images. *bioRxiv*, pp. 2022–06, 2022.

## A. Additional experiments and plots

### A.1. Relation between proposed reward and morphological distance

The underlying assumption behind our method is that the latent representations produced by the contrastive learning model are of a high enough quality to reliably compute similarity. Specifically, given a target pair of a molecule and its associated morphology  $(\hat{x}, \hat{y})$ , an arbitrary pair of a generated molecule and the morphology that would be induced by that molecule  $(x, y)$ , distance  $d$  in the input modality space, distance  $\hat{d}$  in the latent space, and models  $f$  and  $h$  that transforms our input into a latent representation, we require that

$$R(x|\hat{y}) = \hat{d}(h(\hat{y}), f(x)) \sim d(\hat{y}, y). \quad (3)$$

It is worth noting that during generation we always know  $x$ , which is simply the molecule generated by our model, and never know  $y$ , which is why we need  $f$ . Similarly, we always know the target morphology  $\hat{y}$ , and in some settings, such as finding drug analogs, we might also have access to the associated molecule  $\hat{x}$ . In particular, in the latter setting, we might consider conditioning on the joint latent, produced based on the pair of  $(\hat{x}, \hat{y})$ :  $R(x|\hat{x}, \hat{y}) = \hat{d}(fh(\hat{x}, \hat{y}), f(x))$ .

To evaluate this assumption, for every pair of observations from the dataset, we measure the correlation between the similarity of morphological features and the similarity between the latent representation of the structure from the first observation and the latent representation of the target morphology from the second observation. Additionally, we also do the same for the joint target latent, computed based on both the morphology and the associated molecule structure. The results are presented in Figure 5. As can be seen, in both cases we observe a medium level of correlation. While not perfect, we argue that this can be sufficient for screening purposes (where we are interested in improving the hit ratio, but due to the difficulty of the task do not expect very precise outcomes).

### A.2. Joint morphology and structure-guided generation

Our original assumption is that generating a molecule capable of inducing specified morphological outcome would be sufficient to produce desirable drug candidates. However, it is worth noting that in practice this problem can be significantly under-constrained, meaning that there might be a large number of not sufficiently druglike, too toxic, or otherwise not useful molecules capable of inducing the




Figure 5. Correlation between the similarity in morphological feature space and latent similarity between target structure and a) morphology (left) or b) joint morphology and structure (right).

same response. There are multiple ways of constraining generated molecules to a particular subspace of chemical space with desirable properties, for instance by including additional reward terms during GFlowNet training. However, one particularly well suited for our approach is, assuming that we have the molecular structure associated with target morphology available (e.g. in the drug analog search setting), conditioning on the joint latent representation. The aim of this is anchoring generated molecules to the known, desirable molecular structure.

We evaluate the capabilities of the proposed approach in constraining searchable space based on the given structure by replacing target latent with a joint representation, generated based on both morphology and associated structure. We evaluate the number of discovered modes (Figure 6), molecular similarity to the given target (Table 1) and number of samples with high predicted assay probability (Figure 7). As can be seen, using joint representations does not impact the ability of GFlowNet to generate high-reward and diverse samples, and actually increases the average number of discovered modes. Crucially, as expected, it also leads to producing more structurally similar to the target molecules, effectively constraining generation process. However, somehow surprisingly, conditioning based on the joint does not seem to increase the proportion of molecules with high assay probability.

Our leading hypothesis is that the joint embedding space learned by the GMC model acquires a stronger structural signal than a morphological signal, thus leading the GFlowNet to sample molecules that are more similar to the target,




Figure 6. Number of modes discovered for joint morphology and structure-guided generation.

but that don't necessarily trigger the desired morphological profile in the target cell. This claim is supported by Appendix A.6, where the GMC alignment between joint and structural latent space ( $r = 0.94$ ) is greater than the alignment between joint and morphology ( $r = 0.88$ ). Although the alignment between structure and morphology embeddings is even lower ( $r = 0.75$ ), we argue that this setting provides a stronger and more direct signal that guides the GFlowNet towards more diverse molecules that exhibit the desired morphological profile of the target. This is consistent with our results in Figure 7 which suggest that the joint setting yields a higher number of molecules that are structurally similar to the target, but finds less molecules with high assay probabilities than the morphology-only setting.






Figure 7. Comparison of morphology-only versus joint morphology and structure guided generation: a) number of samples with prediction assay probability  $\geq 0.7$  in top 1000 modes (left; higher is better) and b) number of samples in top 1000 modes with Tanimoto similarity to target  $\geq 0.2$  (right). Modes are defined as high reward samples with mutual Tanimoto similarity  $\leq 0.3$

### A.3. Target selection

To evaluate the generative models, we selected a total of 17 targets from the Cell Painting dataset among which 9 were used for the cluster identity experiments and 8 for the assay activity experiments. All targets were verified to be decomposable into the molecular fragment set used for generative modelling and were chosen to represent a diverse set of morphological clusters and active biological assay.

### A.4. GMC model selection

While training the GMC model, we were faced with the choice of which metric to use for model selection. Other than the standard GMC contrastive loss, we used the correlation between the cross modality distance metric presented in section 4.2 for early-stopping. While there wasn’t a significant difference in GFlowNet performance when using these two GMC variants, we found that early-stopping on the correlation metric reduces the variance of cosine similarity between the multimodal GMC embeddings associated

from the same sample when measured in the test split.

### A.5. High Reward Samples

Here we plot some modes sampled by GFlowNet for some targets in a) morphology only guided sampling (left) and b) joint morphology and structure guided sampling (right)

### A.6. GMC cross-modality alignment

Here we plot the correlation between latent distances in one modality versus latent distances in a second modality of the GMC representation space. A high correlation value indicates that the GMC model learns to properly align the modality specific inputs of the data such that associated inputs are aligned closely and distinct inputs aren’t. As expected, GMC achieves higher correlation when the joint modality is included in one of the axes since the joint latent space integrates signals from both the structure and morphology features.

## B. Experimental details

In this section, we present the experiment details for the results obtained in the main paper.

### B.1. GMC model training

We follow the specification in the original GMC paper (Poklugar et al., 2022) and select a single model checkpoint for all our experiments by early-stopping on the GMC validation loss. We employ a Graph Convolutional Network (GCN) for the structure encoder and a simple Multilayer Perceptron (MLP) for the cell morphology inputs. MLPs are also used for the projector architecture for all modalities. See table Table 2 for a full breakdown of the hyperparameters we used.

Parameter	Value
Batch size $\beta$	128
Number of epochs	200
Optimizer	Adam
Learning rate	$1 \times 2e^{-6}$
Non-Linearity	ReLU
Temperature $\tau$	0.4
Intermediate Dim. Size $d$	1024
Latent Dim. Size $s$	1024

Table 2. Hyperparameters of the Geometric Multimodal Contrastive proxy model

## Cell Morphology-Guided Small Molecule Generation with GFlowNets




Figure 8. Active Assay Targets (left) and Morphological Cluster Targets (right).




Figure 9. Cosine similarity between a) structure and morphology embeddings (left) and b) structure and joint embeddings (right) produced by GMC model on test split.

### B.2. Fragment-based molecule generation

Parameter	Value
Batch size	64
Number of steps	10,000
Optimizer	Adam
Number of Layers	4
Hidden Dim. Size	128
Number of Heads	2
Positional Embeddings	Rotary
Reward scaling $\beta$ in $R^\beta$	64
Learning rate	$1 \times 10^{-4}$
$Z$ Learning rate	$1 \times 10^{-3}$

Table 3. Hyperparameters of the Graph Attention Transformer used across all models in fragment-based molecule generation.

In this section, we provide details on the molecule generation experiments and the hyperparameters we used for the methods presented in the paper. In our experimental setup, we follow the same environment specifications and implementations provided in (Malkin et al., 2022b) with

the exception of a different proxy model (GMC) and reward function. The architecture of the GFlowNet, SAC and SQL models is based on a graph attention transformer (Veličković et al., 2017) whose specification is detailed in table Table 3. For SAC, we use a fixed  $\alpha$  value of 0.2 chosen from 0.1, 0.2, 0.5. For SQL, we use a fixed  $\alpha$  value of 0.1. All methods use discount factor  $\gamma$  value of 1.0.

### B.3. Assay selection

For the selection of assays to train the oracle, we primarily focus on assays that have been linked to morphological features and/or combinations of morphological and chemical properties. We select 37 assays identified in (Moshkov et al., 2023) as predictable from morphological features or combined chemical and morphological features with high accuracy (AUROC > 0.9).

### B.4. Oracle training

We trained two MLPs on molecular fingerprints to predict active biological assays (in a multi-label classification setting), and active morphological clusters (in a multi-class classification setting) for the targets in Appendix A.3. For

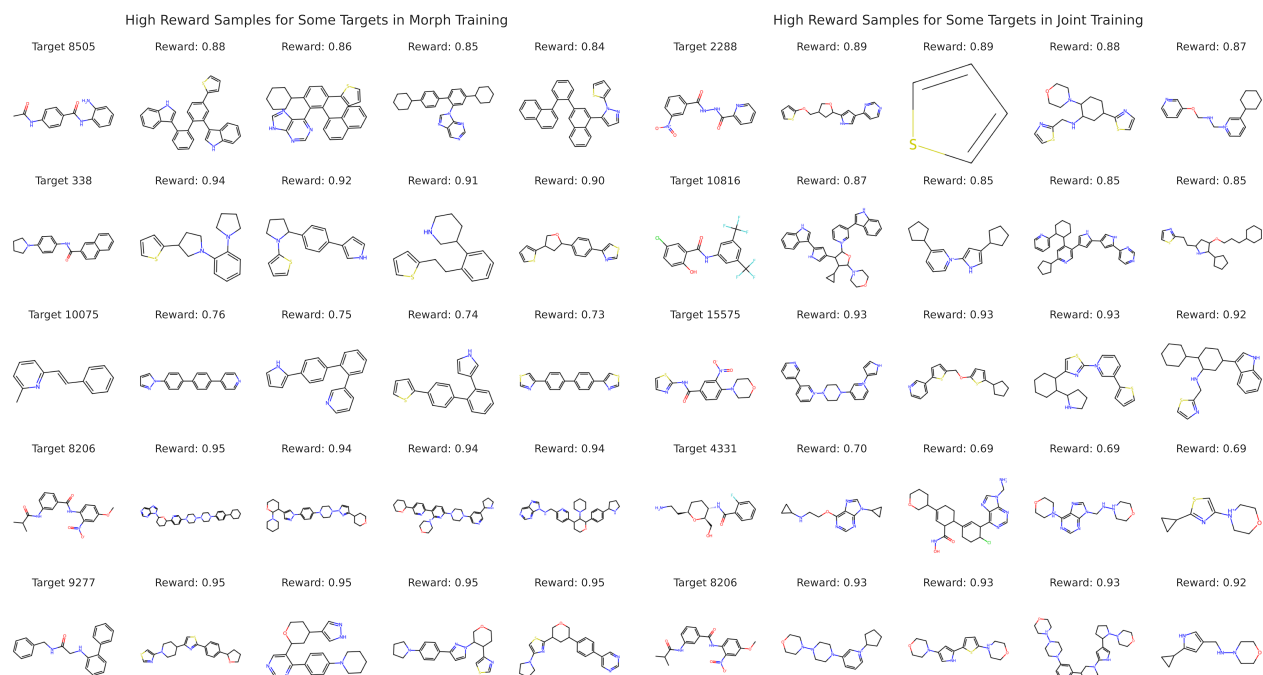


Figure 10. High reward samples generated by GFlowNet.




Figure 11. Cross-modal alignment of samples.

the input, we use Morgan fingerprints with radius 3 and dimension 2048. The MLP has two 64 dim. hidden layers and uses ReLU activation. Both methods are trained with a learning rate of  $1e^{-4}$  with Adam optimizer for 200 epochs. We perform model selection based average precision score on the validation set for both models.

### C. Additional related work

**Deep learning for high-content molecular perturbations.** High-content phenotypic screening, particularly high-

content imaging (HCI), has become crucial in drug discovery for characterizing molecular effects in cells and elucidating targets, gene programs, and biological functions (Moffat et al., 2017; Chandrasekaran et al., 2021). Recent advances propose deep learning techniques to accelerate and enhance these processes (Gavriilidis et al., 2024). Predictive models to infer the outcome of molecular effects have been developed both for transcriptomic readouts (Lotfollahi et al., 2019; Hetzel et al., 2022; Piran et al., 2024) and HCI readouts (Palma et al., 2023). In these models, the output is highly multi-dimensional, capturing the full readout of

the high-throughput experiment, thus potentially requiring a large amount of data for training and making it challenging to separate biological effects from background signals. In contrast to these works, we focus on the inverse problem, designing molecules leading to a specific (target) readout.

**Joint representation of molecules and high-content readouts.** Multi-modal contrastive models (Radford et al., 2021) have been used to align molecular representations to high-dimensional readouts in latent space, thus capturing shared features while avoiding high-dimensional supervised losses (Sanchez-Fernandez et al., 2022; Zheng et al., 2022; Nguyen et al., 2023). While these models can be used for screening tasks, reporting improved generalization ability compared to molecule-only models, they are unable to perform generative tasks. Compared to existing works in this area, we focus on the novel task of HCI-guided molecular design, while relying on a joint representation to guide the generation.