
Towards equilibrium molecular conformation generation with GFlowNets

Alexandra Volokhova*
Mila, Université de Montréal

Michał Koziarski*
Mila, Université de Montréal

Alex Hernández-García
Mila, Université de Montréal

Cheng-Hao Liu
Mila, McGill University

Santiago Miret
Intel Labs

Pablo Lemos
Mila, Université de Montréal,
Ciela Institute, Flatiron Institute

Luca Thiede
Vector Institute, University of Toronto

Zichao Yan
Mila, Université de Montréal

Alán Aspuru-Guzik
Vector Institute, University of Toronto

Yoshua Bengio
Mila, Université de Montréal

Abstract

Sampling diverse, thermodynamically feasible molecular conformations plays a crucial role in predicting properties of a molecule. In this paper we propose to use GFlowNet for sampling conformations of small molecules from the Boltzmann distribution, as determined by the molecule’s energy. The proposed approach can be used in combination with energy estimation methods of different fidelity and discovers a diverse set of low-energy conformations for highly flexible drug-like molecules. We demonstrate that GFlowNet can reproduce molecular potential energy surfaces by sampling proportionally to the Boltzmann distribution.

1 Introduction

Molecules exist in the three-dimensional space as a distribution of atomic positions, referred to as conformations. Given the temperature of the system, the probability of each conformation to occur is defined by its formation energy, and follows a Boltzmann distribution (McQuarrie & Simon, 1997). In many computational drug-discovery processes, it is crucial to know the set of most probable, i.e. low-energy, conformations to predict properties of interest (Boehr et al., 2009; Trott & Olson, 2010). In addition, exploring the potential energy surface by sampling proportionally to the Boltzmann distribution can give key chemical insights such as transition pathways and electron transfer (Schlegel, 2003; Benniston & Harriman, 2006).

Among computational chemistry methods, molecular dynamics simulation is the standard approach, where methods such as CREST have shown feasibility to accurately access numerous low-energy conformations (Pracht et al., 2020b). However, it remains computationally expensive for high throughput applications and large compounds. Faster alternatives with knowledge-based algorithms, such as distance-geometry methods like ETKDG (Riniker & Landrum, 2015), cannot sample in accordance to the Boltzmann distribution and quickly deteriorate with increasing molecular size.

Machine learning (ML) generative models are promising for conformation generation of molecules (Ganea et al., 2021; Xu et al., 2022; Jing et al., 2022). However, they are traditionally focused on

*Equal contribution.

maximum likelihood training on a dataset, which does not guarantee sampling proportionally to the Boltzmann distribution. Several recent works such as Boltzmann generators are approaching this problem (Noé et al., 2019; Köhler et al., 2021; Jing et al., 2022; Zheng et al., 2023), but none of them has yet demonstrated sufficient generality (see Section 2 for details).

In this paper, we use generative flow networks (GFlowNets) for sampling molecular equilibrium conformations from the Boltzmann distribution. We focus on sampling torsion angles of a molecule, as they contain most of the variance of the conformation space while bond lengths and angles can be efficiently generated by fast rule-based methods. Lahlou et al. (2023) presented a proof of concept to demonstrate the capability of GFlowNet to sample from a distribution defined on a two-dimensional torus. Here, we extend this work to a more realistic setting of an arbitrary number of torsion angles. Furthermore, we train GFlowNets with several energy estimation methods of varying fidelity. We experimentally demonstrate that the proposed approach can sample molecular conformations from the Boltzmann distribution, producing diverse, low-energy conformations for a wide range of drug-like molecules with varying number (2-12) of torsion angles.

2 Related works

2.1 Conventional approaches

The most accurate way to get a set of low-energy conformers is based on molecular-dynamics (MD) simulation (Pracht et al., 2020a), which is a computational method for studying the time evolution of physical systems. This method integrates a Newtonian equation of motion with forces obtained by differentiating the potential energy function of the system and employs metadynamics to discover multiple minima of the energy landscape. However, this accuracy comes with high computational costs, despite simplifications of the energy function, making this method not suitable for high-throughput applications and large molecules (Axelrod & Gomez-Bombarelli, 2022).

Alternatively, cheminformatics methods are a fast and popular approach for conformer generation. They utilize structures from experimental reference datasets, chemical rules and heuristics to generate plausible 3D structures given a molecular graph. While significantly faster than MD simulations, these methods tend to lack accuracy and generalization. ETKDG is the most widely used cheminformatics method for conformer generation Riniker & Landrum (2015) implemented in the open-source library RDKit and OMEGA Hawkins et al. (2010) is a popular commercial software implementing cheminformatics methods.

2.2 Machine learning and reinforcement learning approaches

Several machine learning approaches have been developed for conformer generation. Recent advancements include GeoMol (Ganea et al. (2021)), GeoDiff (Xu et al. (2022)), and Torsional Diffusion (Jing et al. (2022)). These methods show good results on the popular benchmark of drug-like molecules GEOM (Axelrod & Gomez-Bombarelli (2022)), demonstrating a decent performance and allowing for faster generation than MD simulations. However, a recently proposed simple clustering algorithm on top of the conformations generated by RDKit outperforms many of these approaches (Zhou et al., 2023). It is important to note that these machine learning methods are mainly trained for maximizing the likelihood of conformations present in the training dataset and therefore not suitable for sampling proportionally to the Boltzmann distribution. In contrast, the GFlowNet framework utilizes information about the Boltzmann weights of the conformations by querying the reward function and allows for generating conformations proportionally to the Boltzmann distribution.

Another promising avenue in the realm of conformer generation is the application of machine learning to model the force fields (Wang et al. (2019); Arts et al. (2023); Thölke & De Fabritiis (2021); Shi et al. (2021)). This approach involves training a machine learning model for predicting forces in the system. Then, molecular dynamics equations are unrolled to sample conformations from the local minima of the energy function based on the forces (gradients of the energy) provided by the model. Machine learning models allow for faster computation of the forces giving a speed advantage compared to MD with an energy function based on the first principles. However, molecular dynamics integration with estimated force fields presents significant stability challenges due to error accumulation during integration Fu et al. (2022), which limits the applicability of this approach.

Reinforcement learning (RL) is another class of machine learning algorithms which cast the generation of molecular structures into a Markov decision process. Well-known examples include TorsionNet (Gogineni et al., 2020) which proceeds by sequentially altering the molecular conformation via updating all torsion angles at every step. The algorithm employs proximal policy optimization (Schulman et al., 2017) and learns by querying a force-fields energy estimation. While TorsionNet and other RL methods do not rely on existing molecular conformation dataset — in a setup similar to our method, they lack the theoretical guarantee of exploring the broader ensemble space of molecular conformation, and indeed, Patel & Tewari (2022) shows that they fail at recovering a diverse set of conformations.

2.3 Boltzmann generators

A closely relevant direction of work is Boltzmann generators (Noé et al. (2019); Köhler et al. (2021)). These are machine learning models aimed at generating samples from the Boltzmann distribution. However, they are based on normalising flows Noé et al. (2019); Köhler et al. (2021), which are known to have limited expressivity and require training a separate model for each molecular graph. The torsional diffusion model Jing et al. (2022) incorporated annealed importance sampling into the training of the diffusion model which allowed for training the model to sample from the target distribution. However, importance sampling brings additional variance into the gradients and becomes challenging with higher dimensions which limits applicability of this approach. Recent work by Zheng et al. (2023) also attempted to make diffusion models sample from the Boltzmann distribution. They pretrained the score model, forcing it to both follow the Fokker-Plank equation and match the energy gradient and then trained it on the dataset with the standard maximum-likelihood objective. This approach allows to incorporate information about energies into the model but does not guarantee sampling from the Boltzmann distribution. Another recent approach (Swanson et al., 2023) approximates Boltzmann distribution as a product of independent distributions for each intrinsic coordinate (a mixture of von-Mises distribution for each torsion angle, unimodal Gaussian distribution for bond lengths and angles). This work shows promising results on the molecules with up to 64 atoms, but the independence assumption can be limiting for modeling larger molecules with long-range interactions.

3 GFlowNet for conformation generation

GFlowNets were originally introduced as a learning algorithm for amortized probabilistic inference in high-dimensional discrete spaces (Bengio et al., 2021) and a generalisation to continuous or hybrid spaces was recently introduced by Lahlou et al. (2023). The method is designed for sampling from an unnormalised probability density which is often represented as the reward function $R(x)$ over the sample space $x \in \mathcal{X}$.

The sampling process starts from a source state s_0 and continues with a trajectory of sequential updates $\tau = (s_0 \rightarrow s_1 \rightarrow \dots \rightarrow s_n = x)$ according to a trainable forward policy $p_F(s_t|s_{t-1}; \theta)$, which defines probability of the forward transition $s_{t-1} \rightarrow s_t$. Once the termination transition $s_{n-1} \rightarrow x$ is sampled, the reward function $R(x)$ provides a signal for computing the training objective. In addition to the trainable forward policy, GFlowNet can learn a backward policy $p_B(s_{t-1}|s_t; \theta)$ for modelling the probability of the backward transitions $s_t \rightarrow s_{t-1}$. It gives additional flexibility to the forward policy p_F allowing to model a rich family of distributions over \mathcal{X} .

In this paper, we propose the use of GFlowNets as a generative model to sample molecular conformations in the space of torsion angles. Formally, we can describe the space of d torsion angles of a molecule as the hyper-torus defined in $\mathcal{X} = [0, 2\pi)^d$. To sample trajectories τ with GFlowNets that satisfy the theoretical assumptions defined by (Lahlou et al., 2023), we establish a fixed number of steps T and include the step number into the state. This yields a state space $\mathcal{S} = \{s_0\} \cup [0, 2\pi)^d \times \{1, 2 \dots T\}$. In order to learn expressive distributions on the hyper-tori, we parameterized independently the forward and backward policies p_F and p_B and trained the GFlowNets with the Trajectory Balance objective (Malkin et al., 2022) defined as:

$$L_{TB}(\tau; \theta) = \left(\log \frac{Z_\theta \prod_{t=1}^n p_F(s_t|s_{t-1}; \theta)}{R(s_n) \prod_{t=1}^n p_B(s_{t-1}|s_t; \theta)} \right)^2. \quad (1)$$

In Equation (1), Z_θ is the (trainable) partition function and $R(s_n)$ is the reward function, evaluated on the terminating states of the trajectories, $s_n = x \in \mathcal{X}$.

Both forward and backward policies are parameterized with a multilayer perceptron (MLP) which outputs parameters for a mixture of von Mises distribution. Note that in this setting we need to train an individual GFlowNet for every molecule as different molecules may have different numbers of torsion angles.

In order to sample from the Boltzmann distribution, we define the reward function using the energy of the molecular conformation in the following way:

$$R(x) = e^{-\mathcal{E}(C(x))\beta}, \quad (2)$$

where $\mathcal{E}(C(x))$ is the potential energy of the molecular conformation $C(x)$ in vacuum and $C(x)$ is defined by the sampled torsion angles x . Other parameters of the conformation C such as bond lengths and angles are sampled with ETKDG (Riniker & Landrum, 2015) and fixed during the GFlowNet training. The positive scalar β corresponds to the inverse temperature of the molecule, but we treat it as a hyper-parameter of the method for the scope of this work.

The energy function is computed as an approximation of quantum mechanical density-functional theory (DFT) and we consider several estimators of different fidelity in our experiments. We use the molecule’s potential energy in vacuum for the scope of our work, but the method can be applied to any energy function of interest.

4 Empirical evaluation

We conducted an experimental study² aimed at addressing several research questions: whether the approach can sample conformers proportionally to the Boltzmann distribution, its capacity to generate diverse low-energy conformations, and how the choice of the energy estimator impacts performance. Additionally, we examined how the proposed approach scales with an increasing number of torsion angles. Firstly, we conducted experiments on molecules with only two considered torsion angles to perform a more in-depth analysis of the results which was not possible in higher dimensions. Then, we scaled up our study, considering a broader range of molecules with varying numbers of torsion angles. Training details can be found in Appendix A.

4.1 Set-up

Conformation generation is conditioned on an input molecular graph encoded as a SMILES string. We first process SMILES with RDKit library (Landrum, 2016) which uses ETKDG to generate an initial conformation. This defines bond lengths and angles and non-rotatable torsion angles. Then, GFlowNet generates rotatable torsion angles and their values are updated accordingly in the conformation.

Energy estimation We experimented with a representative set the existing approaches to energy estimation with different trade-offs between accuracy and computational costs. The most accurate of the considered methods is a semiempirical quantum chemical method, GFN2-xTB (Bannwarth et al., 2019), which is designed for estimating energies of molecular systems for accurate conformation generation. We also employ a faster and less precise force-field approach, GFN-FF (Bannwarth et al., 2021). Finally, we consider TorchANI (Gao et al., 2020), which implements a neural network potential called ANI (Devereux et al., 2020) for energy estimation of organic molecules. Its computational cost is comparable to GFN-FF, but applicability is limited to the specific domain of molecules it was trained on.

Data In the two-dimensional setting, we used alanine dipeptide, ibuprofen and ketorolac molecules. All of them have two main torsion angles largely affecting their energy. For experiments with multiple torsion angles, we used molecules from the GEOM-DRUGS dataset (Axelrod & Gomez-Bombarelli, 2022), a popular benchmark containing low-energy conformations of 304k drug-like molecules,

²One can reproduce our results using the code: <https://github.com/alexhernandezgarcia/gflownet> and the dataset: <https://github.com/learningmatter-mit/geom>.




Figure 1: KDE on samples from the reward function (left), GFlowNet (centre) and MCMC (right) for two molecules and two different proxies: ketorolac from GFN2-xTB (top) and ibuprofen from TorchANI (bottom).

which were generated with meta-dynamics simulation of potential energy estimated by GFN2-xTB (Pracht et al., 2020b). The average number of rotatable torsion angles for GEOM-DRUGS molecules is 7.9 and 92.8% of them have less than 13 rotatable torsion angles (Appendix B).

Non-rotatable torsion angles Torsion angles between the triplets of adjacent bonds defines the shape of the molecule together with bond lengths and angles. Due to the chemical constraints, some torsion angles do not vary between different conformations and can be accurately identified from a SMILES string by rule-based methods together with the bond lengths and angles. Specifically, torsion angles corresponding to non-single bonds and torsion angles within the rings are considered as non-rotatable. In our study we kept these torsion angles fixed and did not sample them with the GFlowNet.

4.2 Two-dimensional setting

We begin by investigating the performance of GFlowNet in simple, well-studied molecular systems in two dimensions: alanine dipeptide, ibuprofen, and ketorolac (Vargas et al., 2002; Zeng et al., 2023). In this experiment, we aim to assess how well the proposed approach can learn to sample from the target distribution and analyze the impact of the energy estimator.

The low dimensionality allows us to visualize the kernel density estimation and evaluate the performance numerically using the Jensen–Shannon divergence (JSD). We used nested sampling (Buchner, 2023) to produce reference ground-truth energy surfaces, and compared the proposed approach with MCMC as an example of an unamortized method for sampling from an unnormalized probability distribution.

In Figure 1, we present the obtained potential energy surfaces of ketorolac and ibuprofen. Comparable analysis for all three molecules and proxies, as well as the computed JSD values, can be found in Appendix E. As can be seen, while the choice of the energy estimator can influence the overall shape of the energy surface, both considered methods accurately reproduce the ground truth energy surface for all estimators. Interestingly, GFlowNet outperformed MCMC in some cases, producing energy surfaces more closely resembling the ground truth.

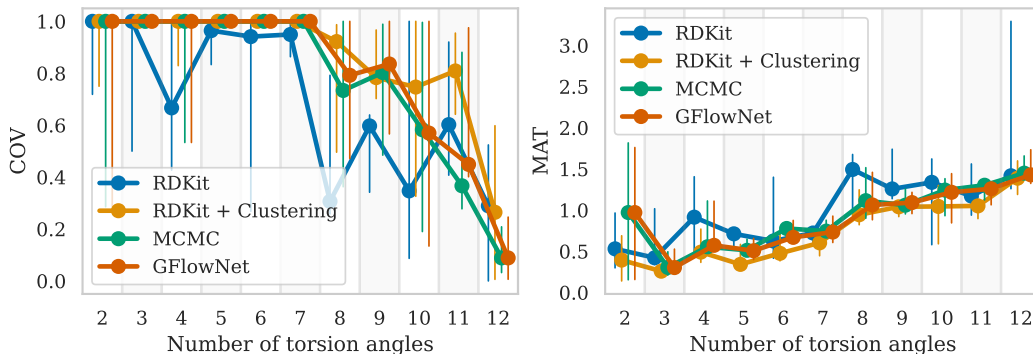


Figure 2: COV (higher is better) and MAT (lower is better) metrics for samples from GFlowNet, RDKit, RDKit+Clustering, MCMC. The markers indicate the median out of five molecules and the vertical lines the range between minimum and maximum values.

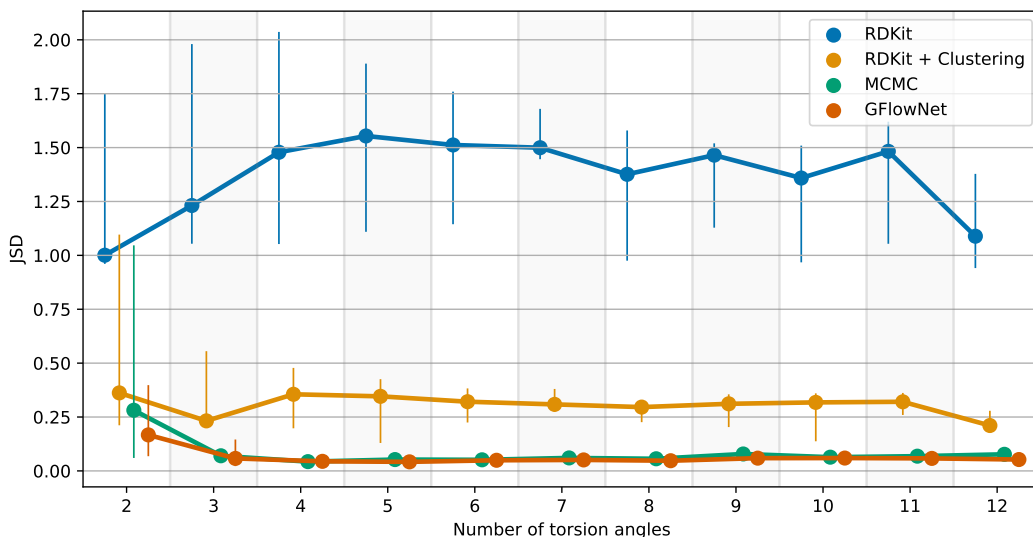


Figure 3: Estimated Jensen-Shannon divergence (JSD) (lower is better) between Boltzmann distribution (with GFN2-xTB energy function) and distributions of samples from GFlowNet and other methods. The markers indicate the median out of five molecules and the vertical lines the range between minimum and maximum values.

4.3 Scaling to multiple torsion angles

We then examined a setting with a higher number of torsion angles, up to 12. For each number of torsion angles we selected 5 different molecules from the GEOM-Drugs dataset, and trained GFlowNet using TorchANI energy estimator separately for each one of them. We compared our method with several baselines: MCMC, ETKDG (Riniker & Landrum, 2015) implemented by the RDKit library, as well as a recent approach combining ETKDG with clustering (Zhou et al., 2023), which was shown to outperform most existing machine learning methods in the low energy conformation generation task. Note that both ETKDG-based approaches perform a qualitatively different task of generating optimized, low-energy conformers, whereas GFlowNet and MCMC aim at sampling from the underlying Boltzmann distribution. In this experiment, we investigate whether diverse low-energy conformations are present among the GFlowNet samples.

In Figure 2, we report the COV and MAT metrics (Appendix C.1), which evaluate the proportion of the reference conformations among samples (COV) and the proximity of the sampled conformations to the reference ones (MAT). We use a set of 1000 generated conformations for all methods to compute the metrics. As can be seen, GFlowNet outperforms vanilla RDKit, and achieves performance comparable




Figure 4: Correlation between estimated sample probability and its reward, as a function of the number of torsion angles (higher is better).

to MCMC and RDKit-clustering, indicating that it can sample diverse low-energy conformations that closely match the ground-truth dataset even when using low-fidelity energy estimator (TorchANI).

4.4 Sampling from the Boltzmann distribution

To measure the discrepancy between model samples and the target distribution, we estimated Jensen-Shannon divergence (JSD) between samples and Boltzmann distribution given by GFN2-xTB energy function (see Appendix C.2 for more details). As one can see from the Figure 3 and Figure 5, both GFlowNet and MCMC samples are much closer to the Boltzmann distribution compared to RDKit and RDKit+Clustering methods. In this experiment, both GFlowNet and MCMC used TorchANI energy estimator and the obtained results suggest that utilising the low fidelity estimation is sufficient to sample according to the higher fidelity energy estimator.

Finally, to provide additional evaluation of GFlowNet’s scalability to higher dimensions, we examined the correlation between the estimated probability to sample a conformation with a GFlowNet (Appendix D) and its energy. We performed this analysis for GFlowNet exclusively because other methods do not provide a straightforward way for estimating sample probability. For the sake of this comparison, we also considered GFN-FF and GFN2-xTB as higher-fidelity energy estimators. The results are presented in Figure 4. As can be seen, while the correlation declines with the increasing number of torsion angles, it remains relatively high throughout the considered range, serving as a further indication that the proposed approach sample according to the Boltzmann distribution in the high dimensions. Interestingly, the higher correlation in higher dimensions with GFN2-xTB and GFN-FF may suggest that their energy landscape is easier to learn than TorchANI, possibly because of the more accurate modelling of the energy surface.




Figure 5: Marginal distributions over the torsion angles for O=C1COc2ccc(C(=O)COC(=O)CCSc3ccccc3)cc2N1 molecule. Notably, GFlowNet and MCMC distributions are much closer to the reward distribution compared to other methods

5 Discussion and future work

In this paper, we have proposed a method for sampling molecular conformations using GFlowNet, which samples proportionally to the Boltzmann distribution. We experimentally evaluated the proposed approach and demonstrated that it can sample diverse, low-energy conformations, can be used in combination with energy estimators of different fidelity, and scales well to a higher number of torsion angles.

While the proposed method still requires training individually for every molecule, the described work can be treated as a stepping stone towards developing a generalised GFlowNet model, which could be trained on a whole data set of molecules and sample torsion angles for an arbitrary molecule conditioned on its molecular graph. This can potentially allow us to amortize the computational costs of sampling proportionally to the Boltzmann distribution compared to methods such as MCMC.

Acknowledgments and Disclosure of Funding

We thank Léna Néhale Ezzine and Prudencio Tossou for their contributions to our discussions. We thank Will Hua for his contribution in the code development at the initial stages of the project. We acknowledge computational support from Intel Labs and Mila compute resources for necessary experiments. We acknowledge Mila’s IDT team support.

References

- Marloes Arts, Victor Garcia Satorras, Chin-Wei Huang, Daniel Zuegner, Marco Federici, Cecilia Clementi, Frank Noé, Robert Pinsler, and Rianne van den Berg. Two for one: Diffusion models and force fields for coarse-grained molecular dynamics. *arXiv preprint arXiv:2302.00600*, 2023.
- Simon Axelrod and Rafael Gomez-Bombarelli. GEOM, energy-annotated molecular conformations for property prediction and molecular generation. *Scientific Data*, 9(1):185, 2022.
- Christoph Bannwarth, Sebastian Ehlert, and Stefan Grimme. Gfn2-xtb—an accurate and broadly parametrized self-consistent tight-binding quantum chemical method with multipole electrostatics and density-dependent dispersion contributions. *Journal of chemical theory and computation*, 15(3):1652–1671, 2019.
- Christoph Bannwarth, Eike Caldeweyher, Sebastian Ehlert, Andreas Hansen, Philipp Pracht, Jakob Seibert, Sebastian Spicher, and Stefan Grimme. Extended tight-binding quantum chemistry methods. *Wiley Interdisciplinary Reviews: Computational Molecular Science*, 11(2):e1493, 2021.
- Emmanuel Bengio, Moksh Jain, Maksym Korablyov, Doina Precup, and Yoshua Bengio. Flow network based generative models for non-iterative diverse candidate generation. *Advances in Neural Information Processing Systems*, 34:27381–27394, 2021.
- Andrew C. Benniston and Anthony Harriman. Charge on the move: how electron-transfer dynamics depend on molecular conformation. *Chem. Soc. Rev.*, 35:169–179, 2006. doi: 10.1039/B503169A. URL <http://dx.doi.org/10.1039/B503169A>.
- David D Boehr, Ruth Nussinov, and Peter E Wright. The role of dynamic conformational ensembles in biomolecular recognition. *Nat. Chem. Bio.*, 5(11):789–796, 2009.
- Johannes Buchner. Nested sampling methods. *Statistic Surveys*, 17:169–215, 2023.
- Christian Devereux, Justin S Smith, Kate K Huddleston, Kipton Barros, Roman Zubatyuk, Olexandr Isayev, and Adrian E Roitberg. Extending the applicability of the ani deep learning molecular potential to sulfur and halogens. *Journal of Chemical Theory and Computation*, 16(7):4192–4202, 2020.
- Xiang Fu, Zhenghao Wu, Wujie Wang, Tian Xie, Sinan Keten, Rafael Gomez-Bombarelli, and Tommi Jaakkola. Forces are not enough: Benchmark and critical evaluation for machine learning force fields with molecular simulations. *arXiv preprint arXiv:2210.07237*, 2022.

- Octavian Ganea, Lagnajit Pattanaik, Connor Coley, Regina Barzilay, Klavs Jensen, William Green, and Tommi Jaakkola. Geomol: Torsional geometric generation of molecular 3d conformer ensembles. *Advances in Neural Information Processing Systems*, 34:13757–13769, 2021.
- Xiang Gao, Farhad Ramezanghorbani, Olexandr Isayev, Justin S Smith, and Adrian E Roitberg. Torchani: A free and open source pytorch-based deep learning implementation of the ani neural network potentials. *Journal of chemical information and modeling*, 60(7):3408–3415, 2020.
- Andrew Gelman and Donald B Rubin. Inference from iterative simulation using multiple sequences. *Statistical science*, 7(4):457–472, 1992.
- Tarun Gogineni, Ziping Xu, Exequiel Punzalan, Runxuan Jiang, Joshua Kammeraad, Ambuj Tewari, and Paul Zimmerman. Torsionnet: A reinforcement learning approach to sequential conformer search. *Advances in Neural Information Processing Systems*, 33:20142–20153, 2020.
- Paul CD Hawkins, A Geoffrey Skillman, Gregory L Warren, Benjamin A Ellingson, and Matthew T Stahl. Conformer generation with omega: algorithm and validation using high quality structures from the protein databank and cambridge structural database. *Journal of chemical information and modeling*, 50(4):572–584, 2010.
- Bowen Jing, Gabriele Corso, Jeffrey Chang, Regina Barzilay, and Tommi Jaakkola. Torsional diffusion for molecular conformer generation. *Advances in Neural Information Processing Systems*, 35:24240–24253, 2022.
- Jonas Köhler, Andreas Krämer, and Frank Noé. Smooth normalizing flows. *Advances in Neural Information Processing Systems*, 34:2796–2809, 2021.
- Salem Lahlou, Tristan Deleu, Pablo Lemos, Dinghuai Zhang, Alexandra Volokhova, Alex Hernández-García, Léna Néhale Ezzine, Yoshua Bengio, and Nikolay Malkin. A theory of continuous generative flow networks. In *International Conference on Machine Learning*, pp. 18269–18300. PMLR, 2023.
- G Landrum. Rdkit: open-source cheminformatics <http://www.rdkit.org>. *Google Scholar There is no corresponding record for this reference*, 3(8), 2016.
- Nikolay Malkin, Moksh Jain, Emmanuel Bengio, Chen Sun, and Yoshua Bengio. Trajectory balance: Improved credit assignment in gflownets. *Advances in Neural Information Processing Systems*, 35: 5955–5967, 2022.
- Donald A. McQuarrie and John D. Simon. *Physical Chemistry: A Molecular Approach*. University Science Books, Sausalito, CA, 1997. Chapter 17.
- Frank Noé, Simon Olsson, Jonas Köhler, and Hao Wu. Boltzmann generators: Sampling equilibrium states of many-body systems with deep learning. *Science*, 365(6457):eaaw1147, 2019.
- Yash Patel and Ambuj Tewari. Rl boltzmann generators for conformer generation in data-sparse environments. *arXiv preprint arXiv:2211.10771*, 2022.
- Philipp Pracht, Fabian Bohle, and Stefan Grimme. Automated exploration of the low-energy chemical space with fast quantum chemical methods. *Physical Chemistry Chemical Physics*, 22(14):7169–7192, 2020a.
- Philipp Pracht, Fabian Bohle, and Stefan Grimme. Automated exploration of the low-energy chemical space with fast quantum chemical methods. *Phys. Chem. Chem. Phys.*, 22:7169–7192, 2020b. doi: 10.1039/C9CP06869D. URL <http://dx.doi.org/10.1039/C9CP06869D>.
- Sereina Riniker and Gregory A Landrum. Better informed distance geometry: using what we know to improve conformation generation. *Journal of chemical information and modeling*, 55(12): 2562–2574, 2015.
- H. Bernhard Schlegel. Exploring potential energy surfaces for chemical reactions: An overview of some practical methods. *Journal of Computational Chemistry*, 24(12):1514–1527, 2003. doi: <https://doi.org/10.1002/jcc.10231>. URL <https://onlinelibrary.wiley.com/doi/abs/10.1002/jcc.10231>.

- John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *CoRR*, abs/1707.06347, 2017.
- Chence Shi, Shitong Luo, Minkai Xu, and Jian Tang. Learning gradient fields for molecular conformation generation. In *International conference on machine learning*, pp. 9558–9568. PMLR, 2021.
- Kirk Swanson, Jake Lawrence Williams, and Eric M Jonas. Von mises mixture distributions for molecular conformation generation. In *International Conference on Machine Learning*, pp. 33319–33342. PMLR, 2023.
- Philipp Thölke and Gianni De Fabritiis. Equivariant transformers for neural network based molecular potentials. In *International Conference on Learning Representations*, 2021.
- Jesús Torrado and Antony Lewis. Cobaya: Bayesian analysis in cosmology. *Astrophysics Source Code Library*, pp. ascl–1910, 2019.
- Jesus Torrado and Antony Lewis. Cobaya: Code for bayesian analysis of hierarchical physical models. *Journal of Cosmology and Astroparticle Physics*, 2021(05):057, 2021.
- Oleg Trott and Arthur J. Olson. Autodock vina: Improving the speed and accuracy of docking with a new scoring function, efficient optimization, and multithreading. *Journal of Computational Chemistry*, 31(2):455–461, 2010. doi: <https://doi.org/10.1002/jcc.21334>. URL <https://onlinelibrary.wiley.com/doi/abs/10.1002/jcc.21334>.
- Rubicelia Vargas, Jorge Garza, Benjamin P. Hay, and David A. Dixon. Conformational study of the alanine dipeptide at the mp2 and dft levels. *The Journal of Physical Chemistry A*, 106(13): 3213–3218, 2002. doi: 10.1021/jp013952f. URL <https://doi.org/10.1021/jp013952f>.
- Jiang Wang, Simon Olsson, Christoph Wehmeyer, Adrià Pérez, Nicholas E Charron, Gianni De Fabritiis, Frank Noé, and Cecilia Clementi. Machine learning of coarse-grained molecular dynamics force fields. *ACS central science*, 5(5):755–767, 2019.
- Minkai Xu, Shitong Luo, Yoshua Bengio, Jian Peng, and Jian Tang. Learning neural generative dynamics for molecular conformation generation. *arXiv preprint arXiv:2102.10240*, 2021.
- Minkai Xu, Lantao Yu, Yang Song, Chence Shi, Stefano Ermon, and Jian Tang. Geodiff: A geometric diffusion model for molecular conformation generation. *arXiv preprint arXiv:2203.02923*, 2022.
- Jinze Zeng, Yujun Tao, Timothy J. Giese, and Darrin M. York. Qd π : A quantum deep potential interaction model for drug discovery. *Journal of Chemical Theory and Computation*, 19(4):1261–1275, 2023. doi: 10.1021/acs.jctc.2c01172. URL <https://doi.org/10.1021/acs.jctc.2c01172>. PMID: 36696673.
- Shuxin Zheng, Jiyan He, Chang Liu, Yu Shi, Ziheng Lu, Weitao Feng, Fusong Ju, Jiayi Wang, Jianwei Zhu, Yaosen Min, et al. Towards predicting equilibrium distributions for molecular systems with deep learning. *arXiv preprint arXiv:2306.05445*, 2023.
- Gengmo Zhou, Zhifeng Gao, Zhewei Wei, Hang Zheng, and Guolin Ke. Do deep learning methods really perform better in molecular conformation generation? *arXiv preprint arXiv:2302.07061*, 2023.

A Experimental details

A.1 Energy computation

Due to perceived energy artifacts observed for GFN-FF, we performed energy normalization and clamping for all of the considered energy estimation methods. Specifically, for every molecule we computed 10,000 random conformations, based on them estimated the minimum and maximum energy values, and normalized all of the estimated energy values to a $[0; 1]$ range. Furthermore, to remove outliers we clamped the values between 1st and 99th quantile of the energy values observed in the initial random sample.




Figure 6: Distribution of torsion angles for molecules in GEOM-DRUGS.

A.2 GFlowNet training

GFlowNet used separate MLP for forward and backward policy, both consisting of 5 hidden layers with 512 neurons each, and using leaky ReLU activation function. Policy input was using positional encoding, with every torsion angle encoded by 10 different values using trigonometric functions. Policy output consisted of predicted von Mises distribution parameters, with 5 separate components per torsion angle. Trajectories of length 5 were used, meaning that every sample was constructed by making 5 separate steps, sampled from the predicted distribution parameters. Minimum concentration of 4 was used for the von Mises distribution. GFlowNet was trained using Boltzmann reward function with $\beta = 32$. Training of a single molecule lasted 40,000 iterations, and used Adam optimizer with learning rate of 0.0001 for the policy and 0.01 for Z_θ . Batches consisted of 80 on-policy trajectories with probability of random sampling equal to 0.1, and 20 trajectories from replay buffer of best trajectories seen so far, with a total buffer capacity of 1000.

A.3 MCMC sampling

For the Markov-Chain Monte Carlo (MCMC) experiments, we run a Metropolis-Hastings algorithm, with a self-tuning covariance matrix for the proposal step. We use the publicly available Cobaya implementation Torrado & Lewis (2019, 2021). We use 4 randomly initialized walkers, and run until reaching a Gelman-Rubin (Gelman & Rubin, 1992) statistic $R - 1 < 0.01$. We then burn-in the first 20% of the chains, and randomly select 1000 samples from the remaining set.

B GEOM-DRUGS torsion angle distribution

One important consideration for the experimental study was the range of evaluated torsion angles. To chose an appropriate range we analysed the distribution of torsion angles for molecules contained in GEOM-DRUGS, as presented in Figure 6. As can be seen, the considered range of 2 to 12 torsion angles faithfully represents the dataset, with majority of molecules contained in it.

C Evaluation

C.1 COVMAT metrics

Root Mean Square Deviation (RMSD) is a metric for evaluating the difference between 3D positions of two conformations, typically computed after first aligning the molecules using a function Φ . Formally, for two molecules R, \hat{R} it can be defined as:

$$\text{RMSD}(\mathbf{R}, \hat{\mathbf{R}}) = \min_{\Phi} \left(\frac{1}{n} \sum_{i=1}^n \left\| \Phi(\mathbf{R}_i) - \hat{\mathbf{R}}_i \right\|^2 \right)^{\frac{1}{2}} \quad (3)$$

Coverage (COV) and matching (MAT) are two metrics incorporating RMSD that, given a set of reference low-energy conformations S_r , measure to what extent generated samples S_g cover the modes in the reference dataset (COV), and how closely do they resemble specific reference conformations (MAT):

$$\text{COV}(S_g, S_r) = \frac{\left| \left\{ \mathbf{R} \in S_r \mid \text{RMSD}(\mathbf{R}, \hat{\mathbf{R}}) < \delta, \hat{\mathbf{R}} \in S_g \right\} \right|}{|S_r|}, \quad (4)$$

$$\text{MAT}(S_g, S_r) = \frac{1}{|S_r|} \sum_{\mathbf{R} \in S_r} \min_{\hat{\mathbf{R}} \in S_g} \text{RMSD}(\mathbf{R}, \hat{\mathbf{R}}). \quad (5)$$

Note that COVMAT metrics were traditionally used in the task of low-energy molecular conformation generation, with the assumption that generated conformations S_g are close to local energy optima. This does not hold when sampling from Boltzmann distribution, and can produce unfavourable results. Because of that we adjusted the procedure by sampling more conformations than typically seen in the literature, as described in the main text.

We follow the literature (Xu et al., 2021; Shi et al., 2021; Zhou et al., 2023) and set δ to 1.25Å in our experiments.

C.2 JSD estimation

Estimating Jensen-Shannon divergence (JSD) between target $q(x)$ and sampled $p(x)$ distribution in high dimensions $x \in \mathcal{X} = [0, 2\pi]^d$ is computationally expensive. Therefore, we estimated JSD separately for each marginal distribution of a torsion angle and then took the average.

For estimating marginal JSD, we first discretized each dimension into K bins. Then, we sampled M d -dimensional vectors of torsion angles from the uniform distribution over \mathcal{X} . Then, we computed their Boltzmann weights $w(x) = \exp(-\mathcal{E}(C(x))\beta)$ and used them to estimate Boltzmann probability for each bin:

$$q_k = \frac{1}{M\bar{C}} \sum_{i=1}^M w(x_i) [x_i \in \text{bin}_k] \quad (6)$$

where $\bar{C} = \frac{1}{M} \sum_{i=1}^M w(x_i)$ and the indicator function $[condition] = 1$ when $condition$ is true and $[condition] = 0$ otherwise. After that, we sampled M samples from the model and estimated sampling probability in each bin as a standard histogram:

$$p_k = \frac{1}{M} \sum_{i=1}^M [x_i \in \text{bin}_k] \quad (7)$$

Finally we computed discrete marginal JSD using estimated q_k and p_k :

$$\text{JSD}_{\text{marginal}} \approx \frac{1}{2} \sum_{k=1}^K q_k \log \left(\frac{2q_k}{p_k + q_k} \right) + \frac{1}{2} \sum_{k=1}^K p_k \log \left(\frac{2p_k}{p_k + q_k} \right) \quad (8)$$

And took the average over dimensions

$$\text{JSD}(q||p) \approx \frac{1}{d} \sum_{j=1}^d \text{JSD}_{\text{marginal}_j} \quad (9)$$

In our analysis, we used $M = 1000$ and $K = 25$.

D Estimation of the log-likelihood of sampling a terminating state x

The log-probability of sampling a terminating state x according to the GFlowNet policy $\pi(x)$ can be expressed as follows:

$$\log \pi(x) = \log \int_{\tau: s_{|\tau|-1} \rightarrow x \in \tau} \prod_{t=0}^{|\tau|-1} P_F(s_{t+1}|s_t; \theta) d\tau = \log \int_{\tau: x \in \tau} P_F(\tau) d\tau.$$

Note that we are abusing notation in favour of readability in the domain of the integral of the right-most equation, which should be identical to the previous integral's. While computing the integral (or even the corresponding sum, if the state space is discrete but very large) in the above expression is intractable in general, we can efficiently estimate it with importance sampling, by using the backward transition probability distribution $P_B(\tau|x)$ as the importance proposal distribution. Let us recall the core aspects of importance sampling, which is a kind of Monte Carlo simulation method.

Monte Carlo simulation methods can be used to estimate integrals $I(f) = \int_{\tau: x \in \tau} f(\tau) p(\tau) d\tau$, as well as very large sums, by drawing N independent and identically distributed (i.i.d.) samples $\{\tau^{(i)}\}_{i=1}^N$ and computing an estimate $I_N(f) = \frac{1}{N} \sum_{i=1}^N f(\tau^{(i)})$ which converges to $I(f)$ as $N \rightarrow \infty$. Importance sampling introduces an importance proposal distribution $q(\tau)$ whose support includes the support of the target distribution $p(\tau)$, such that we can express the integral $I(f)$ as an expectation over the proposal distribution:

$$I(f) = \int_{\tau: x \in \tau} f(\tau) \frac{p(\tau)}{q(\tau)} q(\tau) d\tau = \mathbb{E}_{q(\tau)} \left[f(\tau) \frac{p(\tau)}{q(\tau)} \right]$$

By drawing N i.i.d. samples from the proposal distribution $q(\tau)$ we can compute the estimate

$$\hat{I}_N(f) = \frac{1}{N} \sum_{i=1}^N f(\tau^{(i)}) \frac{p(\tau^{(i)})}{q(\tau^{(i)})}.$$

Returning to the original problem of estimating the log-probability of sampling x with a GFlowNet, we have that the target density is the forward transition probability distribution ($p(\tau) = P_F(\tau)$), the importance proposal distribution is the backward transition probability distribution given x ($q(\tau) = P_B(\tau|x)$), and the sample performance is simply one ($f(\tau) = 1$). Therefore:

$$\log \hat{\pi}(x) = \log \frac{1}{N} \sum_{i=1}^N \frac{P_F(\tau^{(i)})}{P_B(\tau^{(i)}|x)} = \log \sum_{i=1}^N \frac{P_F(\tau^{(i)})}{P_B(\tau^{(i)}|x)} - \log N. \quad (10)$$

E Two-dimensional results

Table 1: Comparison of JSD for model molecules in the two-dimensional setting.

Method	Proxy	Alanine Dipeptide	Ibuprofen	Ketorolac
MCMC	TorchANI	0.0056	0.0055	0.0172
	GFN-FF	0.0076	0.0175	0.0123
	GFN2-xTB	0.0090	0.0074	0.0143
GFlowNet	TorchANI	0.0077	0.0055	0.0070
	GFN-FF	0.0166	0.0182	0.0078
	GFN2-xTB	0.0073	0.0075	0.0059

E.1 Ketorolac




Figure 7: KDE on samples from the reward function (left) GFlowNet (centre) and MCMC (right) for ketorolac for the three proxies: GFN2-xTB (top), GFN-FF (middle) and TorchANI (bottom).

E.2 Ibuprofen

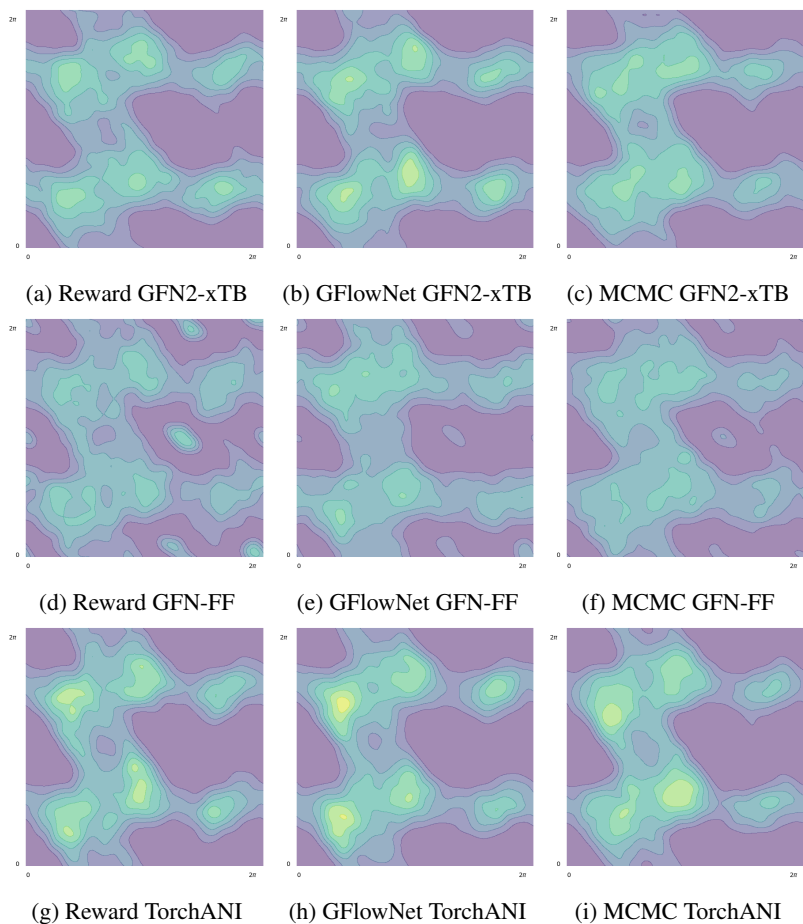


Figure 8: KDE on samples from the reward function (left) GFlowNet (centre) and MCMC (right) for ibuprofen for the three proxies: GFN2-xTB (top), GFN-FF (middle) and TorchANI (bottom).

E.3 Alanine dipeptide




Figure 9: KDE on samples from the reward function (left) GFlowNet (centre) and MCMC (right) for alanine dipeptide for the three proxies: GFN2-xTB (top), GFN-FF (middle) and TorchANI (bottom).