Symmetry-Aware GFlowNets

Hohyun Kim, Seunggeun Lee, Min-hwan Oh

Presented by Hohyun Kim

Graduate Schoold of Data Science
Seoul National University

GFlowNets

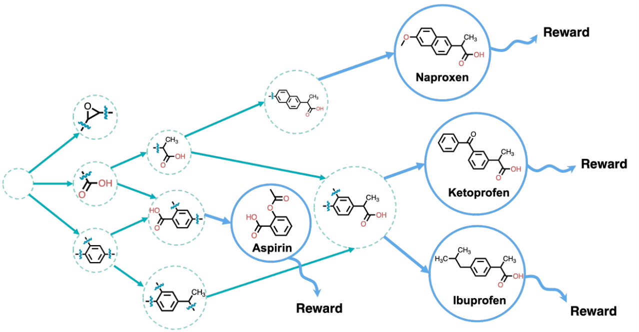

- GFlowNets: train a “flow” of probability mass through a directed graph of states
- Goal: sample objects in proportional to rewards, $p(x) \propto R(x)$

Image source: Jain el al., GFlowNets for AI-Driven Scientific Discovery, 2023.

Equivalent Actions

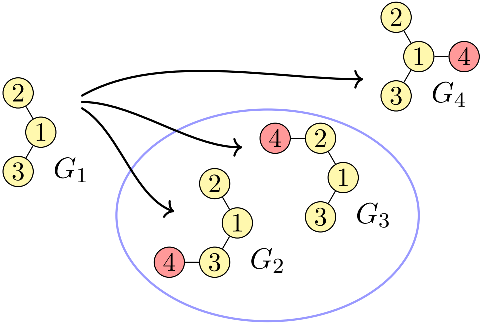

- Equivalent actions: actions that lead to isomorphic graphs
- E.g. G_2 and G_3 are isomorphic

Equivalent Actions

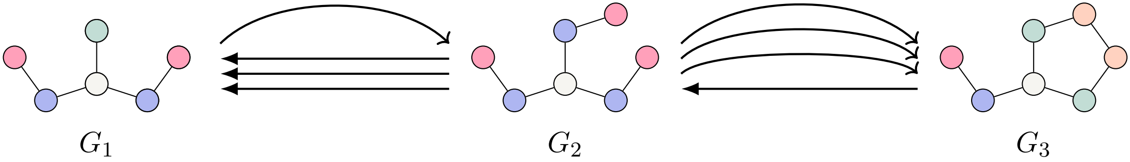

Equivalent Actions

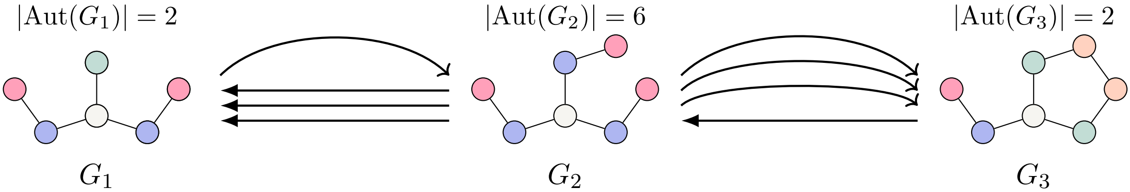

Equivalent Actions

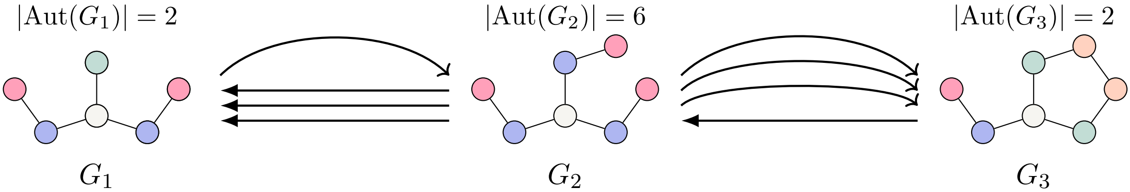

$$\frac{p(s_2|s_1)}{q(s_1|s_2)} = \underbrace{\frac{|\mathrm{Aut}(G_1)|}{|\mathrm{Aut}(G_2)|}}_{\text{correction term}} \cdot \frac{p(G_2|G_1)}{q(G_1|G_2)}$$

Method

Corollary (TB correction)

Assume that G_0 *is the empty graph or a single node, so that* $|\mathrm{Aut}(G_0)| = 1$. *Given the complete graph trajectory* $\tau = (G_0, G_1, \ldots, G_n)$, *the trajectory balance loss can be written as follows:*

$$\mathcal{L}_{\mathrm{TB}}(\tau) = \left(\log \frac{Z \prod_{t=0}^{n-1} p(G_{t+1}|G_t)}{|\mathrm{Aut}(G_n)| R(G_n) \prod_{t=0}^{n-1} q(G_t|G_{t+1})}\right)^2.$$

Method

Corollary (TB correction)

Assume that G_0 is the empty graph or a single node, so that $|\mathrm{Aut}(G_0)| = 1$. Given the complete graph trajectory $\tau = (G_0, G_1, \ldots, G_n)$, the trajectory balance loss can be written as follows:

$$\mathcal{L}_{\mathrm{TB}}(\tau) = \left(\log \frac{Z \prod_{t=0}^{n-1} p(G_{t+1}|G_t)}{|\mathrm{Aut}(G_n)| R(G_n) \prod_{t=0}^{n-1} q(G_t|G_{t+1})} \right)^2 .$$

- Implication: vanilla GFlowNets are biased toward less symmetric graphs

Method

Corollary (TB correction)

Assume that G_0 is the empty graph or a single node, so that $|\mathrm{Aut}(G_0)| = 1$. Given the complete graph trajectory $\tau = (G_0, G_1, \ldots, G_n)$, the trajectory balance loss can be written as follows:

$$\mathcal{L}_{\mathrm{TB}}(\tau) = \left(\log \frac{Z \prod_{t=0}^{n-1} p(G_{t+1}|G_t)}{|\mathrm{Aut}(G_n)| R(G_n) \prod_{t=0}^{n-1} q(G_t|G_{t+1})} \right)^2 .$$

- Implication: vanilla GFlowNets are biased toward less symmetric graphs
- Detailed Balance and Flow-matching objectives can also be adjusted through reward-scaling

Experiments: Illustrative Example

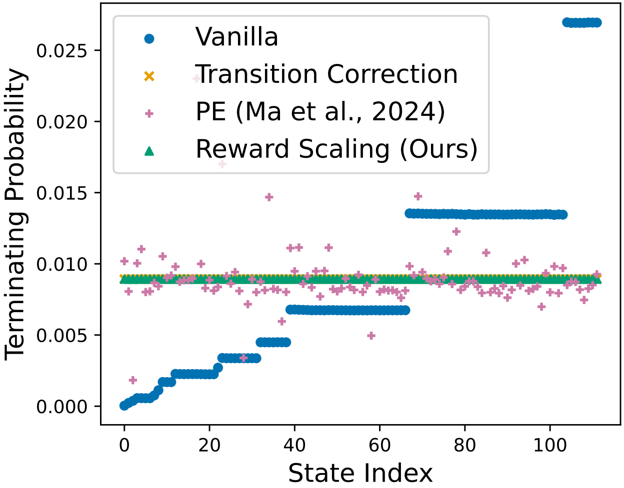

Experiments: Synthetic Graphs

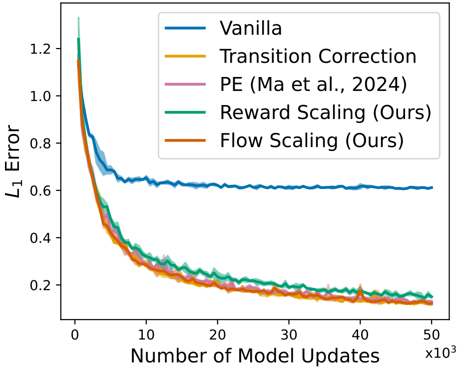

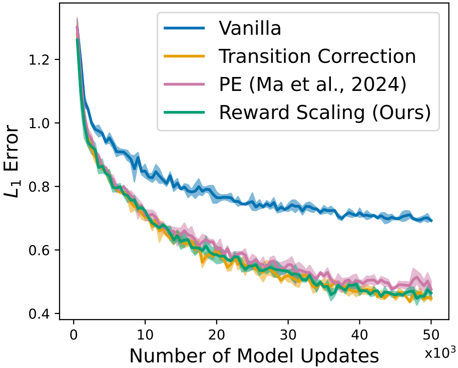

Experiments: Molecule Generaion

Task	Method	Diversity	Top K div.	Top K reward	Uniq. Frac.
Atom	Vanilla	$0.929_{\pm 0.024}$	$0.077_{\pm 0.022}$	$1.09_{\pm 0.02}$	$0.93_{\pm 0.077}$
	Ours (Exact)	$0.959_{\pm 0.01}$	$0.046_{\pm 0.006}$	$1.091_{\pm 0.013}$	$1.0_{\pm 0.0}$
Fragment	Vanilla	$0.877_{\pm 0.001}$	$0.153_{\pm 0.003}$	$0.941_{\pm 0.002}$	$1.0_{\pm 0.0}$
	Ours (Approx.)	$0.88_{\pm 0.001}$	$0.164_{\pm 0.008}$	$0.949_{\pm 0.006}$	$1.0_{\pm 0.0}$
	Ours (Exact)	$0.879_{\pm 0.0}$	$0.151_{\pm 0.002}$	$0.952_{\pm 0.003}$	$1.0_{\pm 0.0}$

Summary

- Without correction, highly symmetric graphs are less likely to be sampled, while symmetric fragments are more likely to be sampled
- Reward-scaling or flow-scaling can effectively eliminate the bias
- Experimental results show that unbiased methods allow the accurate modeling of the target distribution, which is essential for sampling high-reward molecules.