RGFN: Synthesizable Molecular Generation using GFlowNets

Michał Koziarski*, Andrei Rekesh*, Dmytro Shevchuk*, Almer van der Sloot, Piotr Gaiński, Yoshua Bengio, Cheng-Hao Liu, Mike Tyers, Robert A. Batey

* Equal contribution

Motivation

Most generative methods (including GFlowNets) produce candidates difficult to synthesize in a wet lab. We ensure synthesizability out-of-the-box within the GFN framework by operating directly in the action space of chemical reactions.

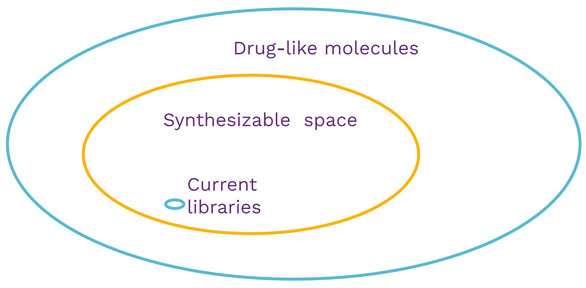

GFlowNets

GFlowNets are amortized variational inference algorithms trained to sample from an unnormalized distribution over compositional objects. They aim to sample objects from a set of terminal states X proportionally to a reward function $R : X \rightarrow R^+$.

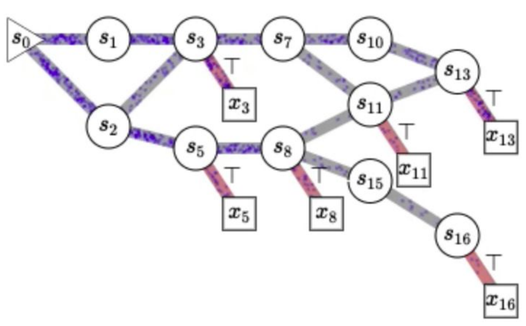

Reaction - GFlowNet

RGFN learns a forward policy over fragments and reactions, parametrized by a multi-headed transformer. We combine **cheap building blocks** along **high-yield compatible reactions** to synthesize molecules as follows:

1. Select initial building block
2. Select compatible reaction
3. Select second reactant
4. Run reaction (RDKit) and select product
5. Repeat 2-4 until stop action is chosen

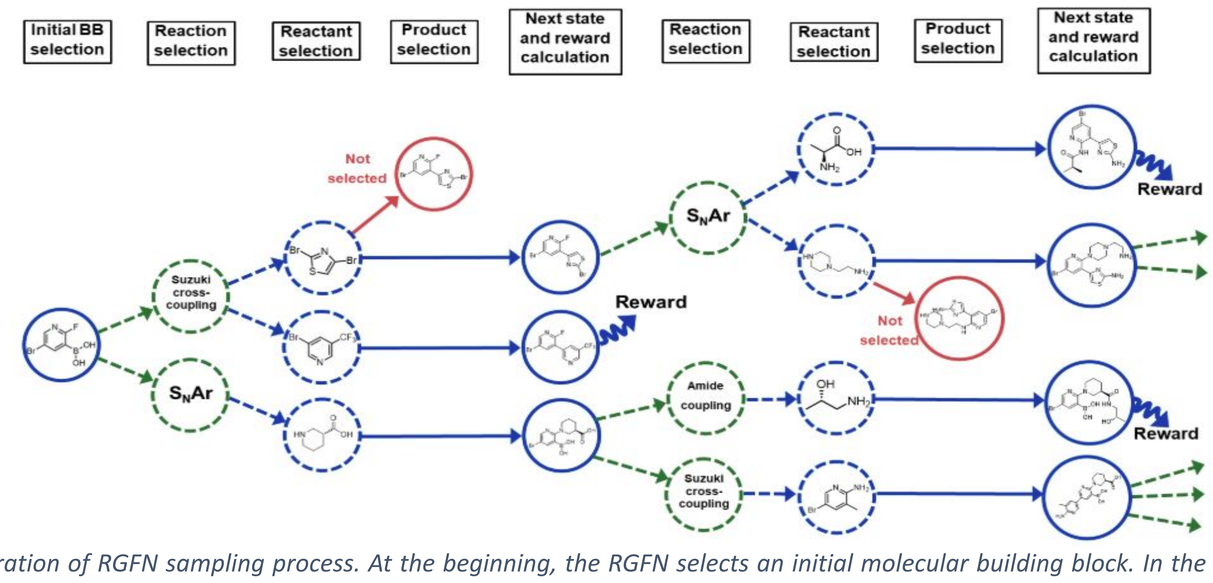

Illustration of RGFN sampling process. At the beginning, the RGFN selects an initial molecular building block. In the next two steps, a reaction and a proper reactant are chosen. Then the reaction is simulated with RDKit's RunReactants functionality and one of the resulting molecules is selected. The process is repeated until the stop action is chosen.

Chemical Language

Only affordable reagents (≤$200/g) are considered. We select a total of 350 reactants and 17 high-yield reactions, which, when combined, still produce a massive search space.

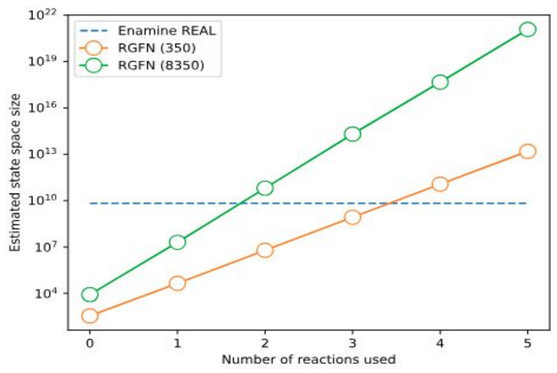

Action Embeddings

As the standard GFlowNet implementation does not encode similarities between actions (BBs), the model would typically have to learn them from scratch. This poses a challenge when scaling to larger action spaces. Instead, we leverage MACCS fingerprints to encode structural similarity of each BB selection action. We learn logits for selecting BB m_i given previously chosen BB m and reaction r:

$$p(m_i|m,r) = \sigma^{|M|}(\mathbf{s})_i,\ s_i = \phi(Wf(m,r))^T g(m_i)$$

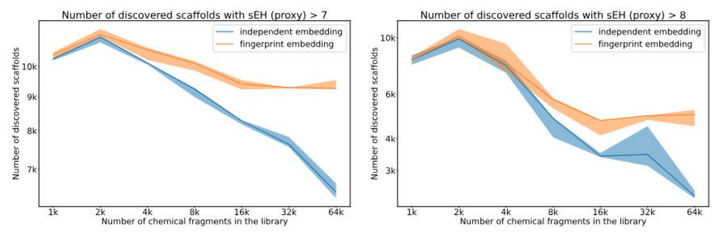

(a) sEH (proxy) > 7 (b) sEH (proxy) > 8

The number of discovered Murcko scaffolds with sEH proxy value above 7 (a) and 8 (b) as a function of fragment library size. We compare one-hot embeddings for reactants (blue) with MACSS fingerprint-based embeddings (orange).

Product Selection

Often, reactions yield multiple products. This introduces stochasticity, which can be problematic for GFNs. We handle this by selecting the product m_i with probability

$$p(m_i'|r) = \sigma^{|M'|}(\mathbf{s})_i,\ s_i = \mathrm{MLP}_{M'}(f(m_i', r))$$

Oracle Models

We evaluate RGFN on molecule generation guided by 3 oracles: a sEH binding affinity proxy, a senolytic activity proxy, and directly calculated docking scores against 4 protein targets. As a faster alternative to traditional docking algorithms, we use **GPU-accelerated docking**, allowing flexibility in protein target selection.

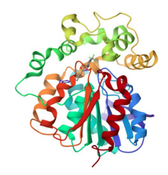

sEH (PDB ID: 4JNC)

ClpP (PDB ID: 7UVU)

Mpro (PDB ID: 6W63)

TBLR1 (PDB ID: 5NAF)

Results

We use SAScore and AiZynthFinder scores to evaluate synthesizability.

Table 1: Average values of synthesizability-related metrics for top-k modes.

Task	Method	Mol. weight ↓	QED ↑	SAScore ↓	AiZynth ↑
sEH	GraphGA	528.6 ± 42.3	0.21 ± 0.06	3.87 ± 0.24	0.04
	SyntheMol	**411.1 ± 66.7**	**0.57 ± 0.18**	**2.85 ± 0.55**	**0.80**
	FGFN	473.4 ± 58.9	0.39 ± 0.13	3.43 ± 0.48	0.14
	RGFN	495.2 ± 49.6	0.29 ± 0.10	3.09 ± 0.39	0.56
Seno.	GraphGA	485.7 ± 75.6	0.09 ± 0.05	2.92 ± 0.26	0.05
	SyntheMol	**441.4 ± 83.5**	**0.48 ± 0.19**	**2.77 ± 0.40**	0.53
	FGFN	467.9 ± 57.3	0.41 ± 0.14	3.74 ± 0.54	0.01
	RGFN	558.7 ± 62.8	0.21 ± 0.09	3.24 ± 0.32	**0.58**
ClpP	GraphGA	521.0 ± 31.8	0.32 ± 0.07	4.14 ± 0.51	0.00
	SyntheMol	**458.2 ± 60.7**	**0.45 ± 0.16**	2.86 ± 0.56	0.56
	FGFN	548.6 ± 42.9	0.22 ± 0.03	2.94 ± 0.54	0.25
	RGFN	526.2 ± 37.6	0.23 ± 0.04	**2.83 ± 0.22**	**0.65**

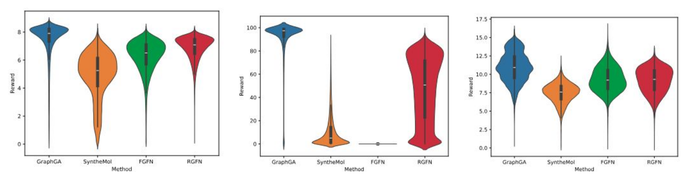

(a) sEH (proxy) (b) senolytics (proxy) (c) ClpP (docking)

RGFN greatly outperforms GraphGA and FGFN on routes found by AiZynthFinder while producing competitive oracle scores.

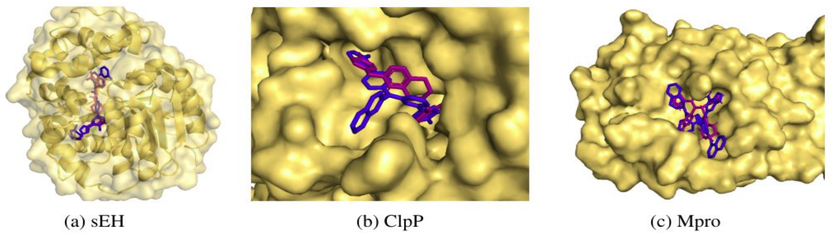

(a) sEH (b) ClpP (c) Mpro

Top docked RGFN ligands after filtering steps (blue) overlaid with the PDB-derived ligand (purple) for each of sEH, ClpP, and Mpro.

RGFN produces molecules with poses similar to known ligands and an average synthesis cost of under $3 per 0.1 mmol of compound.

Ligand	Cost per 0.1 mmol, $
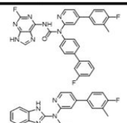	2.07
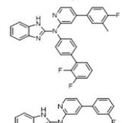	1.90
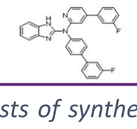	1.76

Costs of synthesis for top 3 ClpP hits generated by RGFN.

Position	Ligand	Cost per 0.1 mmol, $
1	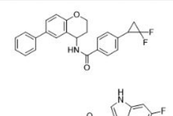	185.79
2	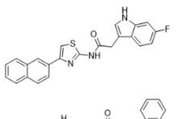	17.68
3	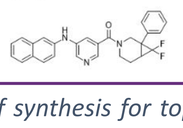	260.33

Costs of synthesis for top 3 ClpP hits generated by SyntheMol.