International Conference on Machine Learning, ICML 2022

Biological Sequence Design with GFlowNets

Moksh Jain, Emmanuel Bengio, Alex-Hernandez Garcia, Jarrid Rector-Brooks, Bonaventure F. P. Dossou, Chanakya Ekbote, Jie Fu, Tianyu Zhang, Michael Kilgour, Dinghuai Zhang, Lena Simine, Payel Das, Yoshua Bengio

Drug Discovery

The Drug Discovery Problem

Why diversity is key

- Novel and diverse antibiotics needed for tackling antimicrobial resistance [1]
- Diverse antibiotics to cover diversity in biological mechanisms [2]
- Diversity in candidate designs to account for underspecification of reward through various phases of discovery

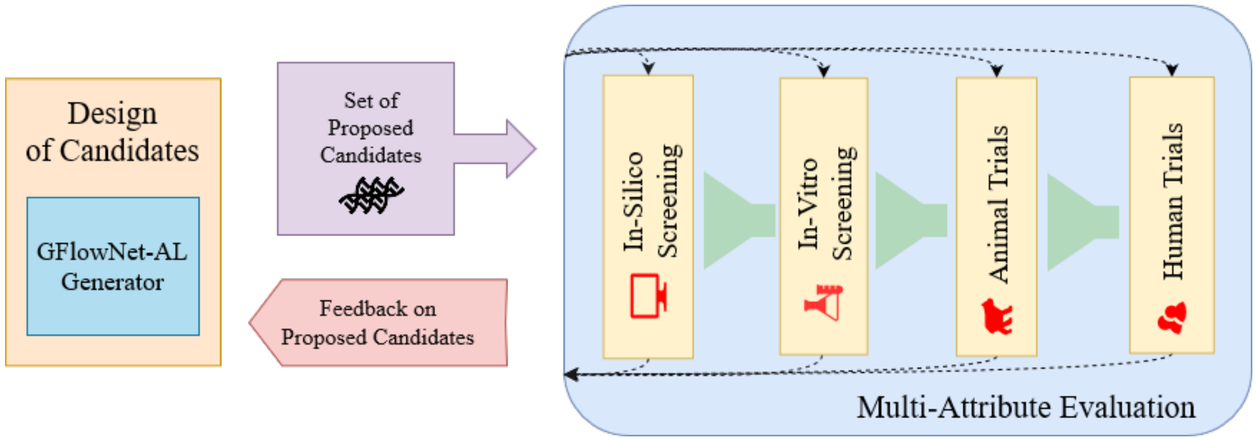

[1] "Global burden of bacterial antimicrobial resistance in 2019: a systematic analysis", Murray et al., 2022
[2] "Diversity, ecology, and prevalence of antimicrobials in nature", Mullis et al., 2019

Desiderata for Biological Sequence Design

Combination of different metrics

- **Performance:** Measure of the desired effect of the generated candidates

$$\text{Performance}(D) = \frac{\sum_{x \in D} \text{score}(x)}{|D|}$$

- **Diversity:** Measure of the variety in generated candidates

$$\text{Diversity}(D) = \frac{\sum_{x,x' \in D \times D; x \neq x'} \text{distance}(x, x')}{|D|(|D| - 1)}$$

- **Novelty:** Measure of novelty with respect to known candidates

$$\text{Novelty}(D) = \frac{\sum_{x \in D} \min_{x' \in D_{known}} \text{distance}(x, x')}{|D|}$$

GFlowNets

GFlowNets

A quick introduction

- Learns stochastic generative policy π for sampling discrete compositional objects x with probability proportional to a reward R, $\pi(x) \propto R(x)$

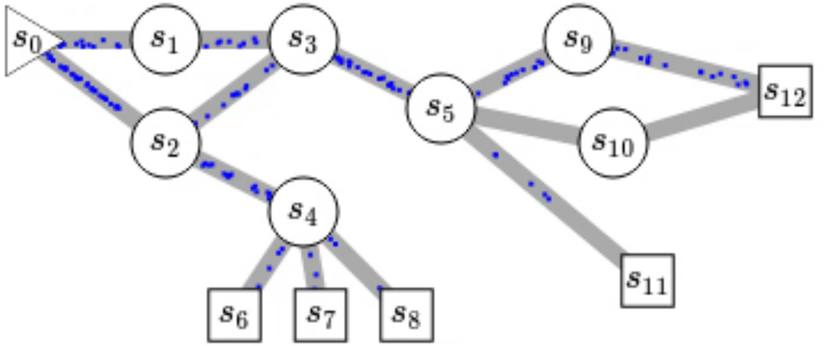

- Reward proportional sampling implicitly encourages diversity in the generated objects

- Framed as learning flows over a directed acyclic graph representing the construction of the object [1,2]

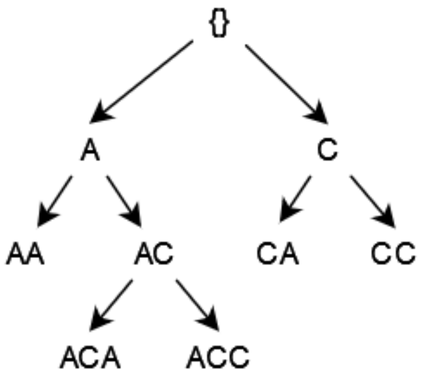

[1] "Flow network based generative models for non-iterative diverse candidate generation", Bengio et al., 2021
[2] "GFlowNet Foundations", Bengio et al., 2021

GFlowNets for Active Learning

GFlowNet-AL

Active learning with GFlowNets

- **Oracle:** Ground truth function governing the property of interest (black-box)
- **Generator:** GFlowNet policy for generating candidates to be queried
- **Proxy:** Approximate model of the oracle
- **Dataset:** Collection of evaluated candidates

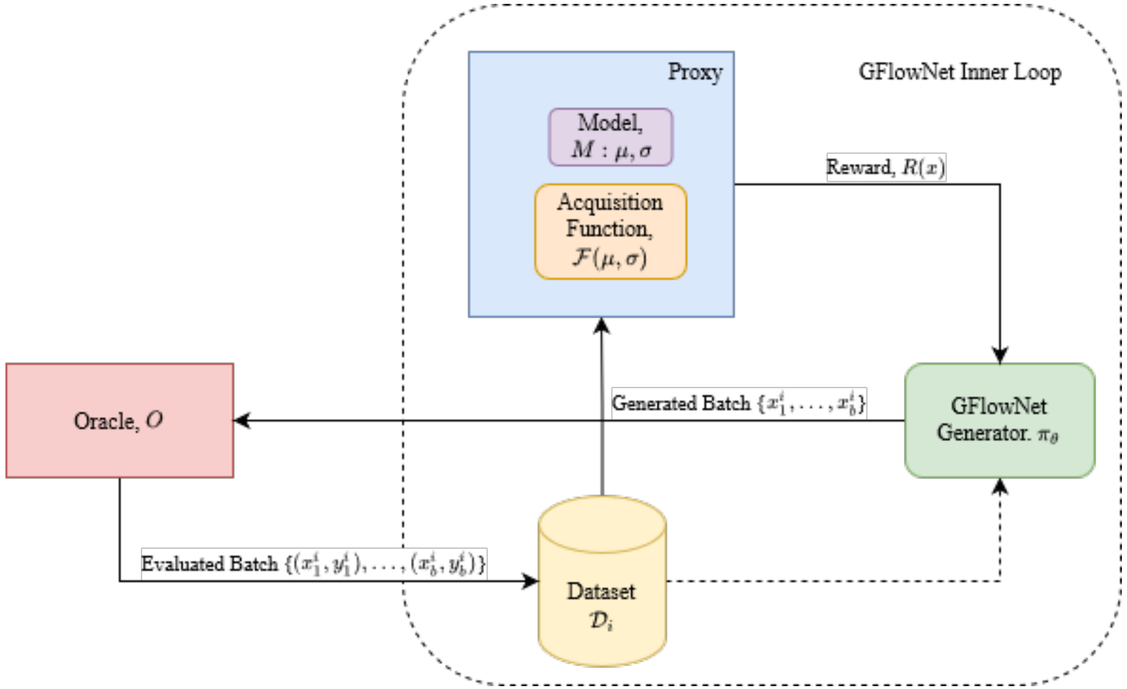

Using Offline Data

$$\mathscr{L}_{TB}(\tau;\theta) = \left(\log \frac{Z_\theta \prod_{s\to s'\in\tau} P_F(s'\,|\,s)}{R(x)}\right)^2$$

- **Key Observation:** GFlowNet objective is offline and off-policy
- Augment training batch with trajectories sampled from available data
  - ☑ Improves learning speed
  - ☑ Guaranteed exploration in known parts of the space

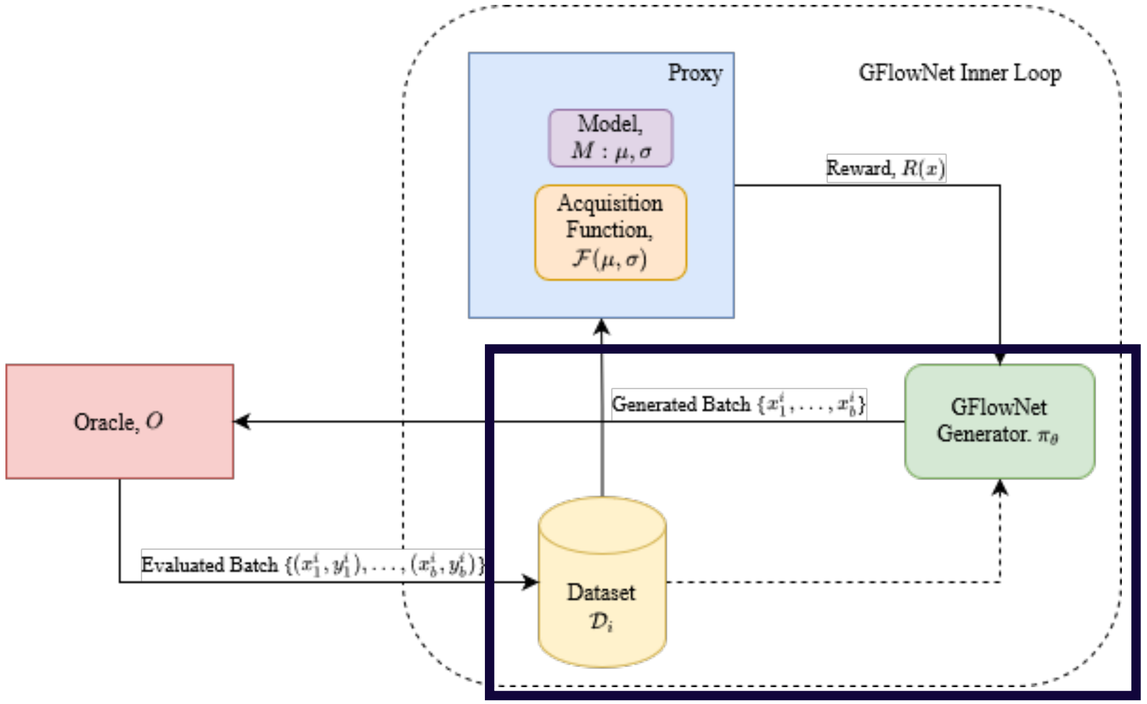

Epistemic Uncertainty from Proxy

- **Key Observation:** As the proxy is trained with limited data, it will have higher *epistemic uncertainty* away from it's training data
- The epistemic uncertainty of the proxy can be used as a signal to guide the generator to new parts of the state space
- An acquisition function (UCB, EI) can be used to trade-off exploration and exploitation

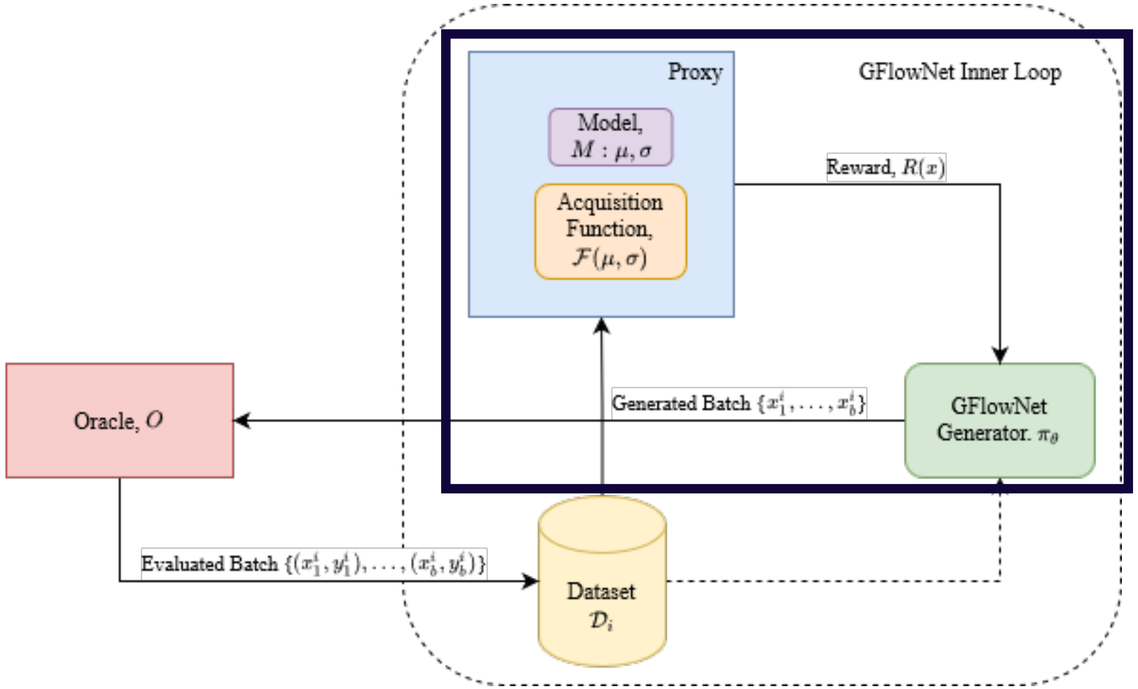

Empirical Results

Results

State-of-the-art Biological Sequence Design

	Performance	Diversity	Novelty
GFlowNet-AL	$\mathbf{0.932 \pm 0.002}$	$\mathbf{22.34 \pm 1.24}$	$\mathbf{28.44 \pm 1.32}$
DynaPPO	$\mathbf{0.938 \pm 0.009}$	12.12 ± 1.71	9.31 ± 0.69
COMs	0.761 ± 0.009	19.38 ± 0.14	26.47 ± 1.3

Antimicrobial Peptide Design

Goal: Designing peptide sequences with antimicrobial properties

Protein sequences with < 60 amino acids

10 rounds of active learning

	Performance	Diversity	Novelty
GFlowNet-AL	$\mathbf{0.84 \pm 0.05}$	4.53 ± 0.46	$\mathbf{2.12 \pm 0.04}$
DynaPPO	0.58 ± 0.02	5.18 ± 0.04	0.83 ± 0.03
COMs	0.74 ± 0.04	4.36 ± 0.24	1.16 ± 0.11
BO-qEI	0.44 ± 0.05	4.78 ± 0.17	0.62 ± 0.23
CbAS	0.45 ± 0.14	5.35 ± 0.16	0.46 ± 0.04
MINs	0.40 ± 0.14	$\mathbf{5.57 \pm 0.15}$	0.36 ± 0.00
CMA-ES	0.47 ± 0.12	4.89 ± 0.01	0.64 ± 0.21
AmortizedBO	0.62 ± 0.01	4.97 ± 0.06	1.00 ± 0.57

TFBind-8

Goal: Designing DNA sequences that bind to human transcription factors

Sequences of length 8

Single round of active learning

	Performance	Diversity	Novelty
GFlowNet-AL	$\mathbf{0.853 \pm 0.004}$	$\mathbf{211.51 \pm 0.73}$	$\mathbf{210.56 \pm 0.82}$
DynaPPO	0.794 ± 0.002	206.19 ± 0.19	203.20 ± 0.47
COMs	0.831 ± 0.003	204.14 ± 0.14	201.64 ± 0.42
BO-qEI	0.045 ± 0.003	139.89 ± 0.18	203.60 ± 0.06
CbAS	0.817 ± 0.012	5.42 ± 0.18	1.81 ± 0.16
MINs	0.761 ± 0.007	5.39 ± 0.00	2.42 ± 0.00
CMA-ES	0.063 ± 0.003	201.43 ± 0.12	203.82 ± 0.09
AmortizedBO	0.051 ± 0.001	205.32 ± 0.12	202.34 ± 0.25

GFP

Goal: Designing protein sequences with flourescence properties

Sequences with 237 amino acids

Single round of active learning

Ablations

- Augmenting GFlowNet training with **offline data** leads to learning speed-ups and improved performance
- Too much offline data, however, can lead to poor exploration and consequently poor performance

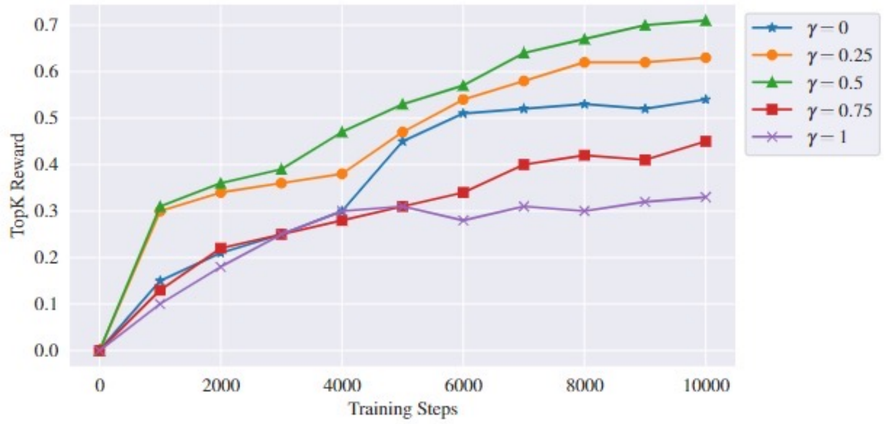

Figure 4. TopK ($K = 100$) scores over the training iterations for GFlowNet-AL in Round 1 of AMP Generation Task, with different values of γ, the proportion of trajectories sampled from the data.

- Using **epistemic uncertainty** from the proxy leads to improvements across all metrics
- Better quality uncertainty estimates result in better performance

Table 4. Results on the AMP Task with $K = 100$ for GFlowNet-AL with different methods for uncertainty estimation, with UCB as the acquisition function.

	Performance	Diversity	Novelty
GFlowNet-AL Ensemble	$\mathbf{0.932 \pm 0.002}$	$\mathbf{22.34 \pm 1.24}$	$\mathbf{28.44 \pm 1.32}$
GFlowNet-AL MC Dropout	0.921 ± 0.004	18.58 ± 1.78	19.58 ± 1.12
GFlowNet-AL None	0.909 ± 0.008	16.42 ± 0.74	17.24 ± 1.44

Conclusion and Future Works

- Using offline data and epistemic uncertainty from a proxy, GFlowNet-AL achieves state-of-the-art performance in generation of diverse and novel biological sequences such as proteins and DNA
- Future works should focus on:
  - Improving the computation efficiency of retraining the proxy
  - Better estimators of information gain
  - Multi-fidelity and multi-objective extensions