
Generative Flow Networks as Entropy-Regularized RL

Daniil Tiapkin*

CMAP - CNRS - École polytechnique
- Institut Polytechnique de Paris,
LMO, CNRS, Université Paris-Saclay,
HSE University

Nikita Morozov*

HSE University

Alexey Naumov

HSE University,
Steklov Mathematical Institute
Russian Academy of Sciences

Dmitry Vetrov

Constructor University, Bremen

Abstract

The recently proposed generative flow networks (GFlowNets) are a method of training a policy to sample compositional discrete objects with probabilities proportional to a given reward via a sequence of actions. GFlowNets exploit the sequential nature of the problem, drawing parallels with reinforcement learning (RL). Our work extends the connection between RL and GFlowNets to a general case. We demonstrate how the task of learning a generative flow network can be efficiently redefined as an entropy-regularized RL problem with a specific reward and regularizer structure. Furthermore, we illustrate the practical efficiency of this reformulation by applying standard soft RL algorithms to GFlowNet training across several probabilistic modeling tasks. Contrary to previously reported results, we show that entropic RL approaches can be competitive against established GFlowNet training methods. This perspective opens a direct path for integrating RL principles into the realm of generative flow networks.

complex discrete space according to a probability mass function given up to an unknown normalizing constant. A fundamental assumption underlying GFlowNets is an additional sequential structure of the space: any object can be generated by following a sequence of actions starting from an initial state, which induces a directed acyclic graph over "incomplete" objects. In essence, GFlowNet aims to learn a stochastic policy to select actions and navigate in this graph in order to match a distribution over final objects with the desired one. A motivational example of the work [Bengio et al. \(2021\)](#) was molecule generation: GFlowNet starts from an empty molecule and learns a policy that adds various blocks sequentially in order to match a desired distribution over a certain molecule space.

The problem statement shares a lot of similarities with reinforcement learning (RL, [Sutton and Barto, 2018](#)) and indeed existing generative flow network algorithms often borrow ideas from RL. In particular, the original GFlowNet training method called flow matching ([Bengio et al., 2021](#)) was inspired by temporal-difference learning ([Sutton, 1988](#)). Next, established techniques in RL were a large source of inspiration for developing novel GFlowNet techniques: subtrajectory balance loss ([Madan et al., 2023](#)) was inspired by TD(λ); a prioritized experience replay idea ([Lin, 1992](#); [Schaul et al., 2016](#)) was actively used in ([Shen et al., 2023](#); [Vemgal et al., 2023](#)); distributional perspective on RL ([Belle-mare et al., 2023](#)) also appears in GFlowNet literature ([Zhang et al., 2023](#)). Additionally, various RL exploration techniques ([Osband et al., 2016](#); [Ostrovski et al., 2017](#); [Burda et al., 2019](#)) were recently adapted into the GFlowNet framework ([Pan et al., 2023](#); [Rector-Brooks et al., 2023](#)).

1 INTRODUCTION

Generative flow networks (GFlowNet, [Bengio et al., 2021](#)) are models designed to learn a sampler from a

*Equal contribution. Proceedings of the 27th International Conference on Artificial Intelligence and Statistics (AISTATS) 2024, Valencia, Spain. PMLR: Volume 238. Copyright 2024 by the author(s).

This large source of techniques results in a question about a more rigorous connection between RL and GFlowNets. Unfortunately, RL with the classical goal of return maximization alone cannot adequately address the sampling problem since the final optimal policies are deterministic or close to deterministic. However, there exists a paradigm of entropy-regularized RL (Fox et al., 2016; Haarnoja et al., 2017; Schulman et al., 2017a) (also known as MaxEnt RL and soft RL), where a final goal is to find a policy that will not only maximize return but also be close to a uniform one, resulting in an energy-based representation of the optimal policies (Haarnoja et al., 2017).

The authors of Bengio et al. (2021) have drawn a direct connection between MaxEnt RL and GFlowNets in the setup of autoregressive generation, or, equivalently, in the setting where the corresponding generation graph is a tree. However, the presented reduction does not hold in the setting of general graphs and results in generation from a highly biased distribution. After that, the authors of Bengio et al. (2023) concluded that the direct connection holds *only* in the tree setup.

In our work, we refute this claim and show that it is possible to reduce the GFlowNet learning problem to entropy-regularized RL beyond the setup of tree-structured graphs.

The main corollary of our result is the possibility to directly apply existing RL techniques without the need to additionally adapt them within GFlowNet framework.

We highlight our main contributions:

- We establish a connection between GFlowNets and entropy-regularized RL by direct reduction;
- We represent detailed balance (Bengio et al., 2023) and trajectory balance (Malkin et al., 2022) algorithms as existing soft RL algorithms;
- We show how to apply the proposed reduction practically by transmitting SoftDQN (Haarnoja et al., 2017) and MunchausenDQN (Vieillard et al., 2020b) into the realm of GFlowNets;
- Finally, we provide experimental validation of our theoretical results and show that MaxEnt RL algorithms can be competitive or even outperform established GFlowNet training methods, contrary to some previously reported results (Bengio et al., 2021).

Source code: github.com/d-tiapkin/gflownet-rl.

2 BACKGROUND

2.1 Generative Flow Networks

In this section we provide some essential definitions and background on GFlowNets. All proofs of the fol-

lowing statements can be found in Bengio et al. (2021) and Bengio et al. (2023).

Suppose we have a finite discrete space \mathcal{X} with a black-box non-negative reward function $\mathcal{R}(x)$ given for every $x \in \mathcal{X}$. Our goal is to construct and train a model that will sample objects from \mathcal{X} from the probability distribution $\mathcal{R}(x)/Z$ proportional to the reward, where Z is an unknown normalizing constant.

Flows over Directed Acyclic Graphs The generation process in GFlowNets is viewed as a sequence of constructive actions where we start with an "empty" object and at each step add some new component to it. To formally describe this process, a finite directed acyclic graph (DAG) $\mathcal{G} = (\mathcal{S}, \mathcal{E})$ is introduced, where \mathcal{S} is a state space and $\mathcal{E} \subseteq \mathcal{S} \times \mathcal{S}$ is a set of edges. For each state $s \in \mathcal{S}$ we define a set of actions \mathcal{A}_s as possible next states: $s' \in \mathcal{A}_s \Leftrightarrow s \rightarrow s' \in \mathcal{E}$. Each action corresponds to adding some new component to an object s , transforming it into a new object s' . If $s' \in \mathcal{A}_s$ we say that s' is a child of s and s is a parent of s' . There is a special initial state $s_0 \in \mathcal{S}$ associated with an "empty" object, and it is the only state that has no incoming edges. All other states can be reached from s_0 , and the set of terminal states (states with no outgoing edges) directly corresponds to the initial space of interest \mathcal{X} .

Denote \mathcal{T} as a set of all complete trajectories in the graph $\tau = (s_0 \rightarrow s_1 \rightarrow \dots \rightarrow s_{n_\tau})$, where n_τ is the trajectory length, $s_i \rightarrow s_{i+1} \in \mathcal{E}$, s_0 is always the initial state, and s_{n_τ} is always a terminal state, thus $s_{n_\tau} \in \mathcal{X}$. As a result, any complete trajectory can be viewed as a sequence of actions that constructs the object corresponding to s_{n_τ} starting from s_0 . A trajectory flow is any nonnegative function $\mathcal{F}: \mathcal{T} \rightarrow \mathbb{R}_{\geq 0}$. If \mathcal{F} is not identically zero, it can be used to introduce a probability distribution over complete trajectories:

$$\mathcal{P}(\tau) \triangleq \frac{1}{Z_{\mathcal{F}}} \mathcal{F}(\tau), \quad Z_{\mathcal{F}} \triangleq \sum_{\tau \in \mathcal{T}} \mathcal{F}(\tau).$$

After that, flows for states and edges are introduced as

$$\mathcal{F}(s) \triangleq \sum_{\tau \ni s} \mathcal{F}(\tau), \quad \mathcal{F}(s \rightarrow s') \triangleq \sum_{\tau \ni (s \rightarrow s')} \mathcal{F}(\tau).$$

Then, $\mathcal{F}(s)/Z_{\mathcal{F}}$ corresponds to the probability that a random trajectory contains the state s , and $\mathcal{F}(s \rightarrow s')/Z_{\mathcal{F}}$ corresponds to the probability that a random trajectory contains the edge $s \rightarrow s'$. These definitions also imply $Z_{\mathcal{F}} = \mathcal{F}(s_0) = \sum_{x \in \mathcal{X}} \mathcal{F}(x)$.

Denote $\mathcal{P}(x)$ as the probability that x is the terminal state of a trajectory sampled from $\mathcal{P}(\tau)$. If the *reward matching constraint* is satisfied for all terminal states $x \in \mathcal{X}$:

$$\mathcal{F}(x) = \mathcal{R}(x), \quad (1)$$

$\mathcal{P}(x)$ will be equal to $\mathcal{R}(x)/Z$. Thus by sampling trajectories we can sample objects from the distribution of interest.

Markovian Flows To allow efficient sampling of trajectories, the family of trajectory flows considered in practice is narrowed down to Markovian flows. A trajectory flow \mathcal{F} is said to be Markovian if $\forall s \in \mathcal{S} \setminus \mathcal{X}$ there exist distributions $\mathcal{P}_F(-|s)$ over the children of s such that \mathcal{P} can be factorized as

$$\mathcal{P}(\tau = (s_0 \rightarrow \dots \rightarrow s_{n_\tau})) = \prod_{t=1}^{n_\tau} \mathcal{P}_F(s_t | s_{t-1}).$$

This means that trajectories can be sampled by iteratively sampling next states from $\mathcal{P}_F(-|s_i)$ (sequential construction of an object). Equivalently, \mathcal{F} is Markovian if $\forall s \in \mathcal{S} \setminus \{s_0\}$ there exist distributions $\mathcal{P}_B(-|s)$ over the parents of s such that $\forall x \in \mathcal{X}$

$$\mathcal{P}(\tau = (s_0 \rightarrow \dots \rightarrow s_{n_\tau}) | s_{n_\tau} = x) = \prod_{t=1}^{n_\tau} \mathcal{P}_B(s_{t-1} | s_t),$$

which means that trajectories ending at x can be sampled by iteratively sampling previous states from $\mathcal{P}_B(-|s_i)$. If \mathcal{F} is Markovian, then \mathcal{P}_F and \mathcal{P}_B can be uniquely defined as

$$\mathcal{P}_F(s'|s) = \frac{\mathcal{F}(s \rightarrow s')}{\mathcal{F}(s)}, \quad \mathcal{P}_B(s|s') = \frac{\mathcal{F}(s \rightarrow s')}{\mathcal{F}(s')}. \quad (2)$$

We call \mathcal{P}_F and \mathcal{P}_B the forward policy and the backward policy respectively. *The detailed balance constraint* is a relation between them:

$$\mathcal{F}(s)\mathcal{P}_F(s'|s) = \mathcal{F}(s')\mathcal{P}_B(s|s'). \quad (3)$$

The trajectory balance constraint states that for any complete trajectory

$$\prod_{t=1}^{n_\tau} \mathcal{P}_F(s_t | s_{t-1}) = \frac{\mathcal{F}(s_{n_\tau})}{Z_{\mathcal{F}}} \prod_{t=1}^{n_\tau} \mathcal{P}_B(s_{t-1} | s_t). \quad (4)$$

Thus any Markovian flow can be uniquely determined from $Z_{\mathcal{F}}$ and $\mathcal{P}_F(-|s) \forall s \in \mathcal{S} \setminus \mathcal{X}$ or from $\mathcal{F}(x) \forall x \in \mathcal{X}$ and $\mathcal{P}_B(-|s) \forall s \in \mathcal{S} \setminus \{s_0\}$.

Training GFlowNets In essence, GFlowNet is a model that parameterizes Markovian flows over some fixed DAG and is trained to minimize some objective. The objective should be chosen in such a way that, if the model is capable of parameterizing any Markovian flow, global minimization of the objective will lead to producing a flow that satisfies (1). Thus the forward policy of a trained GFlowNet can be used as a sampler from the reward distribution. Two of the widely used training objectives are:

Trajectory Balance (TB). Introduced in Malkin et al. (2022), based on the trajectory balance constraint (4).

A model with parameters θ predicts the normalizing constant Z_θ , forward policy $\mathcal{P}_F(-|s, \theta)$ and backward policy $\mathcal{P}_B(-|s, \theta)$. The trajectory balance objective is defined for each complete trajectory $\tau = (s_0 \rightarrow s_1 \rightarrow \dots \rightarrow s_{n_\tau} = x)$:

$$\mathcal{L}_{\text{TB}}(\tau) = \left(\log \frac{Z_\theta \prod_{t=1}^{n_\tau} \mathcal{P}_F(s_t | s_{t-1}, \theta)}{\mathcal{R}(x) \prod_{t=1}^{n_\tau} \mathcal{P}_B(s_{t-1} | s_t, \theta)} \right)^2. \quad (5)$$

Detailed Balance (DB). Introduced in Bengio et al. (2023), based on the detailed balance constraint (3). A model with parameters θ predicts state flows $\mathcal{F}_\theta(s)$, forward policy $\mathcal{P}_F(-|s, \theta)$ and backward policy $\mathcal{P}_B(-|s, \theta)$. It is important to note that such parameterization is excessive and the components are not necessarily compatible with each other during the course of training. The detailed balance objective is defined for each edge $s \rightarrow s' \in \mathcal{E}$:

$$\mathcal{L}_{\text{DB}}(s, s') = \left(\log \frac{\mathcal{F}_\theta(s)\mathcal{P}_F(s'|s, \theta)}{\mathcal{F}_\theta(s')\mathcal{P}_B(s|s', \theta)} \right)^2. \quad (6)$$

For terminal states $\mathcal{F}_\theta(x)$ is substituted with $\mathcal{R}(x)$ in the objective.

During the training, the objective of choice is stochastically optimized across complete trajectories sampled from a training policy (which is usually taken to be \mathcal{P}_F , a tempered version of it or a mixture of \mathcal{P}_F and a uniform distribution). An important note is that for all objectives $\mathcal{P}_B(-|s)$ can either be trained or fixed, e.g. to be uniform over the parents of s . When the backward policy is fixed, the corresponding forward policy is uniquely determined by \mathcal{P}_B and \mathcal{R} .

Among other existing objectives, the originally proposed flow matching objective Bengio et al. (2021) was shown to have worse performance than TB and DB (Malkin et al., 2022; Madan et al., 2023). Recently, *subtrajectory balance* (SubTB) (Madan et al., 2023) was introduced as a generalization of TB and DB and shown to have better performance when appropriately tuned.

2.2 Soft Reinforcement Learning

In this section we describe the main aspects of usual reinforcement learning and its entropy-regularized version, also known as MaxEnt and Soft RL.

Markov Decision Process The main theoretical object of RL is Markov decision process (MDP) that is defined as a tuple $\mathcal{M} = (\mathcal{S}, \mathcal{A}, P, r, \gamma, s_0)$, where \mathcal{S} and \mathcal{A} are finite state and action spaces, P is a Markovian transition kernel, r is a bounded deterministic reward function, $\gamma \in [0, 1]$ is a discount factor, and s_0 is a fixed initial state. Additionally, for each state s we define a set of possible actions as $\mathcal{A}_s \subseteq \mathcal{A}$. An RL agent interacts with MDP by a policy π that maps each state s to a probability distribution over possible actions.

Entropy-Regularized RL The performance measure of the agent in classic RL is a value that is defined as an expected (discounted) sum of rewards over a trajectory generated by following a policy π starting from a state s . In contrast, in entropy-regularized reinforcement learning (Neu et al., 2017; Geist et al., 2019) the value is augmented by Shannon entropy:

$$V_\lambda^\pi(s) \triangleq \mathbb{E}_\pi \left[\sum_{t=0}^{\infty} \gamma^t (r(s_t, a_t) + \lambda \mathcal{H}(\pi(s_t))) \mid s_0 = s \right], \quad (7)$$

where λ is a regularization coefficient, \mathcal{H} is a Shannon entropy function, $a_t \sim \pi(s_t)$, $s_{t+1} \sim P(s_t, a_t)$ for all $t \geq 0$, and expectation is taken with respect to these random variables. In a similar manner we can define regularized Q-values $Q_\lambda^\pi(s, a)$ as an expected (discounted) sum of rewards augmented by Shannon entropy given a fixed initial state $s_0 = s$ and an action $a_0 = a$. A regularized optimal policy π_λ^* is a policy that maximizes $V_\lambda^\pi(s)$ for any initial state s .

Soft Bellman Equations Let V_λ^* and Q_λ^* be the value and the Q-value of the optimal policy π_λ^* correspondingly. Then Theorem 1 and 2 by Haarnoja et al. (2017) imply the following system relations for any $s \in \mathcal{S}$, $a \in \mathcal{A}_s$

$$\begin{aligned} Q_\lambda^*(s, a) &\triangleq r(s, a) + \gamma \mathbb{E}_{s' \sim P(s, a)} [V^*(s')] \\ V_\lambda^*(s) &\triangleq \text{LogSumExp}_\lambda(Q_\lambda^*(s, \cdot)), \end{aligned} \quad (8)$$

where $\text{LogSumExp}_\lambda(x) \triangleq \lambda \log(\sum_{i=1}^d e^{x_i/\lambda})$, and the optimal policy is defined as a λ -softmax policy with respect to optimal Q-values $\pi_\lambda^*(s) = \text{Softmax}_\lambda(Q_\lambda^*(s, \cdot))$, where $[\text{Softmax}_\lambda(x)]_i \triangleq e^{x_i/\lambda} / (\sum_{j=1}^d e^{x_j/\lambda})$, or, alternatively

$$\pi_\lambda^*(a|s) = \exp\left(\frac{1}{\lambda}(Q_\lambda^*(s, a) - V_\lambda^*(s))\right). \quad (9)$$

Notice that as $\lambda \rightarrow 0$ we recover the well-known Bellman equations (see e.g. Sutton and Barto (2018)) since LogSumExp_λ approximates maximum and Softmax_λ approximates a uniform distribution over $\arg \max$.

Soft Deep Q-Network One of the classic ways to solve RL problem for discrete action spaces is the well-known Deep Q-Network algorithm (DQN) (Mnih et al., 2015). Below we present the generalization of this algorithm to entropy-regularized RL problems called **SoftDQN** (also known as **SoftQLearning** (Fox et al., 2016; Haarnoja et al., 2017; Schulman et al., 2017a) in a continuous control setting).

The goal of **SoftDQN** algorithm is to approximate a solution to soft Bellman equations (8) using a Q-value parameterized by a neural network. Let Q_θ be an online Q-network that represents the current approximation of the optimal regularized Q-value. Also let $\pi_\theta = \text{Softmax}_\lambda(Q_\theta)$ be a policy that approximates the optimal policy. Then, akin to DQN, agent interacts with MDP using policy π_θ (with an additional

ε -greedy exploration or even without it) and collects transitions (s_t, a_t, r_t, s_{t+1}) for $r_t = r(s_t, a_t)$ in a replay buffer \mathcal{B} . Then **SoftDQN** performs stochastic gradient descent on the regression loss $\mathcal{L}_{\text{SoftDQN}}(\mathcal{B}) = \mathbb{E}_{\mathcal{B}}[(Q_\theta(s_t, a_t) - y_{\text{SoftDQN}}(r_t, s_{t+1}))^2]$ with the target Q-value

$$y_{\text{SoftDQN}}(r_t, s_{t+1}) = r_t + \gamma \text{LogSumExp}_\lambda(Q_{\bar{\theta}}(s_{t+1}, \cdot)), \quad (10)$$

where $Q_{\bar{\theta}}$ is a target Q-network with parameters $\bar{\theta}$ that are periodically copied from the parameters of online Q-network θ . Additionally, there exist other algorithms that were generalized from traditional RL to soft RL, such as celebrated **SAC** (Haarnoja et al., 2018) and its discrete version (Christodoulou, 2019).

It is worth mentioning another algorithm such as Munchausen DQN (M-DQN, Vieillard et al., 2020b) that exploits the regularization structure similarly to **SoftDQN** but augments the target using the current policy:

$$y_{\text{M-DQN}}(s_t, a_t, r_t, s_{t+1}) = y_{\text{SoftDQN}}(r_t, s_{t+1}) + \alpha \lambda \log \pi_{\bar{\theta}}(a_t | s_t).$$

It turns out that this scheme is equivalent (up to a reparametrization) to approximate mirror descent value iteration (Vieillard et al., 2020a) method for solving $(1 - \alpha)\lambda$ entropy-regularized MDP with an additional KL-regularization with respect to the target policy with coefficient $\alpha\lambda$ (Vieillard et al., 2020b, Theorem 1).

3 REPRESENTATION OF GFLOWNETS AS SOFT RL

In this section we describe a deep connection between GFlowNets and entropy-regularized reinforcement learning.

3.1 Prior Work

The sequential type of the problem that GFlowNet is solving creates the parallel automatically, and the main inspiration for flow-matching loss arises from temporal-difference learning (Bengio et al., 2021).

Beyond inspiration, one formal link between GFlowNets and RL was established in the seminal paper (Bengio et al., 2021). In particular, the authors showed that it is possible to learn the GFlowNet forward policy by applying MaxEnt RL in the setup of tree DAG. However, for a general setting the authors claimed that MaxEnt RL is sampling with probability proportional to $n(x)\mathcal{R}(x)$ where $n(x)$ is a number of paths to a terminal state $x \in \mathcal{X}$, thus these methods should be highly biased. Therefore, the authors of (Bengio et al., 2023) claimed that equivalence between MaxEnt RL and GFlowNet holds only if the underlying DAG has a tree structure.

In addition, it is possible to treat the GFlowNet learning problem as a convex MDP problem (Zahavy et al., 2021) with a goal to directly minimize $\text{KL}(d^\pi \|\mathcal{R}/\mathcal{Z})$, where d^π is a distribution induced by a policy π over terminal states. However, as it was claimed by Bengio et al. (2023), RL-based algorithms for solving this problem require high-quality estimation of a distribution d^π , and the scheme by Zahavy et al. (2021) needs to do it for changing policies during the learning process. A similar problem arises in the generalization of count-based exploration in deep RL (Bellemare et al., 2016; Ostrovski et al., 2017; Saade et al., 2023).

3.2 Direct Reduction by MDP Construction

In this section we show how to recast the problem of learning GFlowNet forward policy \mathcal{P}_F given a reward function \mathcal{R} and a backward policy \mathcal{P}_B as a soft RL problem.

First, for any DAG $\mathcal{G} = (\mathcal{S}, \mathcal{E})$ with a set of terminal states \mathcal{X} we can define the corresponding deterministic controlled Markov process $(\mathcal{S}', \mathcal{A}, \mathcal{P})$ with an additional absorbing state s_f . Formally, the construction is as follows. Firstly, choose the state space \mathcal{S}' as the state space of the corresponding DAG \mathcal{G} with an additional absorbing state s_f : $\mathcal{S}' = \mathcal{S} \cup \{s_f\}$. Next, instead of defining the full action space \mathcal{A} , for any s we select $\mathcal{A}_s \subseteq \mathcal{S}$ as a set of actions in state s in $\mathcal{S} \setminus \mathcal{X}$, and define $\mathcal{A}_x \triangleq \{s_f\}$ for any $x \in \mathcal{X} \cup \{s_f\}$. In other words, each terminal state only has an action that leads to s_f and s_f itself has a loop. Then we can define \mathcal{A} as a union of all \mathcal{A}_s for any $s \in \mathcal{S}'$. Finally, the construction of a deterministic Markov kernel is simple: $\mathcal{P}(s'|s, a) = \mathbf{1}\{s' = a\}$ since a corresponds to the next state in \mathcal{S} given the corresponding action. Subsequently, we will write s' instead of a to emphasize this connection.

Theorem 1. *Let $\mathcal{G} = (\mathcal{S}, \mathcal{E})$ be a DAG with a set of terminal states \mathcal{X} , let \mathcal{P}_B be a fixed backward policy and \mathcal{R} be a GFlowNet reward function.*

Let $\mathcal{M}_G = (\mathcal{S}', \mathcal{A}, \mathcal{P}, r, \gamma, s_0)$ be an MDP with $\mathcal{S}', \mathcal{A}, \mathcal{P}$ constructed from \mathcal{G} with an additional absorbing state s_f , $\gamma = 1$, and MDP rewards defined as follows

$$r(s, s') = \begin{cases} \log \mathcal{P}_B(s|s') & s \notin \mathcal{X} \cup \{s_f\}, \\ \log \mathcal{R}(s) & s \in \mathcal{X}, \\ 0 & s = s_f \end{cases} \quad (11)$$

Then the optimal policy $\pi_1^(s'|s)$ for the regularized MDP with coefficient $\lambda = 1$ is equal to $\mathcal{P}_F(s'|s)$ for all $s \in \mathcal{S} \setminus \mathcal{X}, s' \in \mathcal{A}_s$;*

Moreover, $\forall s \neq s_f, s' \neq s_f$ regularized optimal value $V_1^(s)$ and Q -value $Q_1^*(s, s')$ coincide with $\log \mathcal{F}(s)$ and $\log \mathcal{F}(s \rightarrow s')$ respectively, where \mathcal{F} is the Markovian flow defined by \mathcal{P}_B and the GFlowNet reward \mathcal{R} .*

The complete proof is given in Appendix B and basically uses soft Bellman equations (8) and backward induction over a topological ordering of the graph.

Intuition. Given a fixed GFlowNet backward policy, the optimal policy of MaxEnt RL with $\lambda = 1$ and adjusted rewards is identical to the GFlowNet forward policy. The key insight lies in introducing intermediate rewards through the fixed backward policy in order to normalize probabilities of different paths leading to the same terminal state.

Remark 1. In the setting of tree-structures DAG \mathcal{G} we notice that our result shows that it is enough to define the reward $\log \mathcal{R}$ for terminal states since any backward policy \mathcal{P}_B satisfies $\log \mathcal{P}_B(s|s') = 0$. Thus, we reproduce the known result of Proposition 1(a,b) by Bengio et al. (2021), see also (Bengio et al., 2023). Regarding the results of a part (c), the key difference between our construction of MDP and construction by Bengio et al. (2021) is non-trivial rewards for non-terminal states that "compensate" a multiplier $n(x)$ that appears in (Bengio et al., 2021).

Remark 2. Notably, our derivation explains a logarithmic scaling that appears in TB (see (5)), DB (see (6)) and SubTB (Madan et al., 2023) losses. Previously, the appearance of these logarithms was explained solely by numerical reasons (Bengio et al., 2021, 2023).

Remark 3. Taking $\lambda \neq 1$ under the choice of rewards established in Theorem 1 will lead to a biased policy. To see it, notice that solving soft RL with coefficient $\lambda > 0$ and rewards r is equivalent to solving soft RL with coefficient $\lambda' = 1$ and new rewards $r' = r/\lambda$ since $V_\lambda^*(s_0; r) = \mathbb{E}[\sum_t r_t | s'] = \lambda \mathbb{E}[\sum_t r_t / \lambda + \mathcal{H}(\pi(s_t))] = \lambda V_{\lambda'}^*(s_0; r/\lambda)$. For the terminal rewards it will lead to rescaling the GFlowNet reward $R(x) \mapsto R^{1/\lambda}(x)$, but all intermediate rewards will also be rescaled $\log \mathcal{P}_B(s|s') \mapsto \log \mathcal{P}_B(s|s')/\lambda$, leading to a bias. However, additionally multiplying intermediate rewards by λ in Theorem 1 for $\lambda \neq 1$ will lead to a policy that samples from the tempered GFlowNet reward distribution.

Remark 4. We highlight two major differences with the experimental setup of Bengio et al. (2021): they applied PPO (Schulman et al., 2017b) with regularization coefficient $\lambda < 1$ and SAC (Haarnoja et al., 2018) with adaptive regularization coefficient (whilst our construction requires $\lambda = 1$), and set intermediate rewards to zero (while we set them to $\log \mathcal{P}_B(s | s')$) This resulted in convergence to a biased final policy.

Interpretation of Value Notably, in the MDP \mathcal{M}_G we have a clear interpretation of the value function. But first we need to introduce several important definitions. For a policy π in a deterministic MDP \mathcal{M}_G with all trajectories of length $n_\tau \leq N$ we can define an induced trajectory distribution as follows

$$q^\pi(\tau) = \prod_{i=1}^N \pi(s_i | s_{i-1}) = \prod_{i=1}^{n_\tau} \pi(s_i | s_{i-1}),$$

where $\tau = (s_0 \rightarrow \dots \rightarrow s_N)$ with $s_N = s_f$ and $s_{n_\tau} \in \mathcal{X}$. The last identity holds since $\pi(s_f | s_f) = \pi(s_f | x) = 1$ for any $x \in \mathcal{X}$. Additionally, we define a marginal distribution of sampling by policy π as $d^\pi(x) = \mathbb{P}_\pi[s_{n_\tau} = x]$. Additionally, for a fixed GFlowNet reward \mathcal{R} and a backward policy \mathcal{P}_B we define a corresponding distribution over trajectories with a slight abuse of notation

$$\mathcal{P}_B(\tau) = \frac{\mathcal{R}(s_{n_\tau})}{Z} \prod_{i=1}^{n_\tau} \mathcal{P}_B(s_{i-1} | s_i).$$

Notice that by the trajectory balance constraint (4) and by Theorem 1 $\mathcal{P}_B(\tau)$ is equal to $q^{\pi^*}(\tau)$.

Proposition 1. *Let \mathcal{M}_G be an MDP constructed in Theorem 1 by a DAG \mathcal{G} given a fixed backward policy \mathcal{P}_B and a GFlowNet reward function \mathcal{R} . Then for any policy π its regularized value V_1^π in the initial state s_0 is equal to $V_1^\pi(s_0) = \log Z - \text{KL}(q^\pi \| \mathcal{P}_B)$.*

The proof idea is to apply the direct definition of value (7). The complete proof is given in Appendix B.

The main takeaway of this proposition is that maximization of the value function in the prescribed MDP \mathcal{M} is equivalent to minimization of the KL-divergence between trajectory distribution induced by the current policy and by the backward policy. Moreover, it shows that all near-optimal policies in soft RL-sense introduce a flow with close trajectory distributions, and, moreover, close marginal distribution:

$$V_1^*(s_0) - V_1^\pi(s_0) = \text{KL}(q^\pi \| \mathcal{P}_B) \geq \text{KL}(d^\pi \| \mathcal{R}/Z),$$

where $d^\pi(x)$ is a marginal distribution over terminal states induced by π and $\mathcal{R}(x)/Z$ is the reward distribution we want to sample from, see (Cover, 1999, Theorem 2.5.3). In other words, the policy error controls the distribution approximation error, additionally justifying the presented MDP construction.

3.3 Soft DQN in the realm of GFlowNets

In this section we explain how to apply SoftDQN algorithm to GFlowNet training.

As described in Section 2.2, in SoftDQN an agent uses a network Q_θ that approximates optimal Q -value or, equivalently, the logarithm of the edge flow. To interact with DAG the agent uses a softmax policy $\pi_\theta(s) = \text{Softmax}_1(Q_\theta(s, \cdot))$. During the interaction, the agent collects transitions (s_t, a_t, r_t, s_{t+1}) into a replay buffer \mathcal{B} and optimizes a regression loss $\mathcal{L}_{\text{SoftDQN}}$ using targets (10) with rewards $r_t = \log \mathcal{P}_B(s_t | s_{t+1})$ for any $s_t \notin \mathcal{X}$. For terminal transitions with $s_t \in \mathcal{X}$ and $s_{t+1} = s_f$ the target is equal to $\log \mathcal{R}(s_t)$. In particular, using the interpretation $Q_\theta(s, s') = \log \mathcal{F}_\theta(s \rightarrow s')$ we can

obtain the following loss

$$\mathcal{L}_{\text{SoftDQN}}(\mathcal{B}) = \mathbb{E}_{\mathcal{B}} \left[\left(\log \frac{\mathcal{F}_\theta(s_t \rightarrow s_{t+1})}{\mathcal{P}_B(s_t | s_{t+1}) \mathcal{F}_{\bar{\theta}}(s_{t+1})} \right)^2 \right], \quad (12)$$

where $\mathcal{F}_{\bar{\theta}}(x) = \mathcal{R}(x) \forall x \in \mathcal{X}$, and otherwise $\mathcal{F}_{\bar{\theta}}(s) \triangleq \sum_{s'} \mathcal{F}_{\bar{\theta}}(s \rightarrow s')$. $\bar{\theta}$ are parameters of a target network that is updated with weights θ from time to time. In particular, this loss may be "derived" from the definition of backward policy (2) akin to deriving DB and TB losses from the corresponding identities (3)-(4) with the differences of fixed backward policy and usage of a target network $\bar{\theta}$ and a replay buffer \mathcal{B} . We refer to Appendix C for discussion on technical algorithmic choices.

Munchausen DQN Using exactly the same representation, we can apply M-DQN (Vieillard et al., 2020b) to train GFlowNets. However, unlike SoftDQN, it requires a more careful choice of λ since M-DQN solves a regularized problem with entropy coefficient $(1 - \alpha)\lambda$. Thus, one needs to opt for a non-trivial choice of $\lambda = 1/(1 - \alpha)$ to compensate the effect of Munchausen penalty. We refer the reader to Appendix C for a more detailed description.

Learnable backward policy The aforementioned approach works with fixed backward policies, which is a viable and widely utilized option for GFlowNet training (Malkin et al., 2022; Zhang et al., 2022b; Deleu et al., 2022; Malkin et al., 2023; Madan et al., 2023) that has benefits such as more stable training. At the same time having a trainable backward policy can lead to faster convergence (Malkin et al., 2022). Our approach still lacks understanding of the effect of optimization of the backward policy, which is a limitation.

Nevertheless, it seems to be possible to interpret this optimization in a game-theoretical fashion: one player is selecting rewards (defined by a backward policy) to optimize its objective and the other player is trying to guess a good forward policy for these rewards, see (Zahavy et al., 2021; Tiapkin et al., 2023) for a similar idea applied to convex MDPs and maximum entropy exploration. We leave the formal derivation as a promising direction for future research.

3.4 Existing GFlowNets through Soft RL

In this section we show that DB and TB GFlowNet losses can be derived from existing RL literature up to technical details.

Detailed Balance as Dueling Soft DQN To show a more direct connection between SoftDQN and DB algorithms, we first describe a well-known technique in RL known as a *dueling architecture* (Wang et al., 2016). Instead of parameterizing Q -values solely, the dueling

architecture proposes to have a neural network with two streams for values $V_\theta(s)$ and action advantages $A_\theta(s, a)$. These two streams are connected to the original Q-value as $Q_\theta(s, a) = V_\theta(s) + A_\theta(s, a)$ by the definition of advantage function.

The main idea behind this technique is that it is believed that value could be learnt faster and then accelerate the learning of advantage by bootstrapping (Tang et al., 2023). However, as it was acknowledged in the seminal work (Wang et al., 2016), the reparametrization presented above is unidentifiable since the advantage is not uniquely defined by a Q-value. In the setup of classic RL it was suggested to subtract maximum of the advantage. In the setup of soft RL it makes sense to subtract log-sum-exp of the advantage and obtain the following identity

$$Q_\theta(s, a) = V_\theta(s) + A_\theta(s, a) - \text{LogSumExp}_1(A_\theta(s, \cdot)).$$

The key properties of this representation are the following relations (for $\lambda = 1$ in the GFlowNet setup):

$$\begin{aligned} \text{LogSumExp}_1(Q_\theta(s, \cdot)) &= V_\theta(s), \\ \log \pi_\theta(a|s) &= A_\theta(s, a) - \text{LogSumExp}_1(A_\theta(s, \cdot)). \end{aligned}$$

Using GFlowNet interpretation of $V_\theta(s) = \log \mathcal{F}(s)$ and $\pi_\theta(s'|s) = \mathcal{P}_F(s'|s, \theta)$, the loss (12) transforms into

$$\mathcal{L}_{\text{SoftDQN}}^{\text{Dueling}}(\mathcal{B}) = \mathbb{E}_{\mathcal{B}} \left[\left(\log \frac{\mathcal{P}_F(s_{t+1}|s_t, \theta) \mathcal{F}_\theta(s_t)}{\mathcal{P}_B(s_t|s_{t+1}) \mathcal{F}_\theta(s_{t+1})} \right)^2 \right],$$

that coincides with DB-loss (6) up to the lack of a learnable backward policy and usage of a replay buffer and a target network.

Trajectory Balance as Policy Gradient Let us consider the setting of policy gradient methods: we parameterize the policy space π_θ by parameter θ and our goal is to directly maximize $V^{\pi_\theta}(s_0)$. The essence of policy gradient methods is to apply stochastic gradient ascent over θ to optimize this objective.

However, a combination of Proposition 1 and Proposition 1 by (Malkin et al., 2023) implies that

$$-\nabla_\theta V_1^{\pi_\theta}(s_0) = \nabla_\theta \text{KL}(q^{\pi_\theta} \parallel \mathcal{P}_B) = \frac{1}{2} \mathbb{E}_{\tau \sim \pi_\theta} [\nabla_\theta \mathcal{L}_{\text{TB}}(\tau)],$$

where $\mathcal{L}_{\text{TB}}(\tau)$ is TB-loss (5). This means that optimization of *on-policy* TB-loss is equivalent to policy gradient with using $\log Z_\theta$ as a baseline function. Additionally, this connection is very close to a connection between soft Q-learning and policy gradient methods established by Schulman et al. (2017a).

4 EXPERIMENTS

In this section we evaluate the soft RL approach to the task of GFlowNet training under the proposed

paradigm and compare it to prior GFlowNet training approaches: trajectory balance (TB, Malkin et al., 2022), detailed balance (DB, Bengio et al., 2023) and subtrajectory balance (SubTB, Madan et al., 2023). Among soft RL methods, for the sake of brevity this section only presents the results for Munchausen DQN (M-DQN, Vieillard et al., 2020b) since we found it to show better performance in comparison to other considered techniques. For ablation studies and additional comparisons we refer the reader to Appendix E.

4.1 Hypergrid Environment

We start with synthetic hypergrid environments introduced in Bengio et al. (2021). These environments are small enough that the normalizing constant Z can be computed exactly and the trained policy can be efficiently evaluated against the reward distribution.

The set of non-terminal states is a set of points with integer coordinates inside a D -dimensional hypercube with side length H : $\{(s^1, \dots, s^D) \mid s^i \in \{0, 1, \dots, H-1\}\}$. The initial state is $(0, \dots, 0)$, and for each non-terminal state s there is its copy s^\top which is terminal. For a non-terminal state s the allowed actions are incrementing one coordinate by 1 without exiting the grid and the terminating action $s \rightarrow s^\top$. The reward has modes near the corners of the grid which are separated by wide troughs with a very small reward. We study 2-dimensional and 4-dimensional grids with reward parameters taken from Malkin et al. (2022). To measure the performance of GFlowNet during training, L^1 -distance is computed between the true reward distribution $\mathcal{R}(x)/Z$ and the empirical distribution $\pi(x)$ of last $2 \cdot 10^5$ samples seen in training (endpoints of trajectories sampled from \mathcal{P}_F): $\frac{1}{|\mathcal{X}|} \sum_{x \in \mathcal{X}} |\mathcal{R}(x)/Z - \pi(x)|$. Remember, that in case of M-DQN we have $\mathcal{P}_F(\cdot | s) = \text{Softmax}_\lambda(Q_\theta(s, \cdot))$ for $\lambda = 1/(1 - \alpha)$. All models are parameterized by MLP, which accepts a one-hot encoding of s as input. Additional details can be found in Appendix D.1.

We consider two cases for the baselines: uniform and learnt \mathcal{P}_B . Since our theoretical framework does not support learnt \mathcal{P}_B , we fix it to be uniform for M-DQN in all cases. Figure 1 presents the results. M-DQN has faster convergence than all baselines in case of uniform \mathcal{P}_B . Training \mathcal{P}_B improves the speed of convergence for the baselines, with SubTB showing better performance than M-DQN in some cases. However, even with learnt \mathcal{P}_B , TB and DB converge slower than M-DQN.

4.2 Small Molecule Generation

Next, we consider the molecule generation task from Bengio et al. (2021). The goal is to generate binders of the sEH (soluble epoxide hydrolase) protein. The molecules represented by graphs are generated by

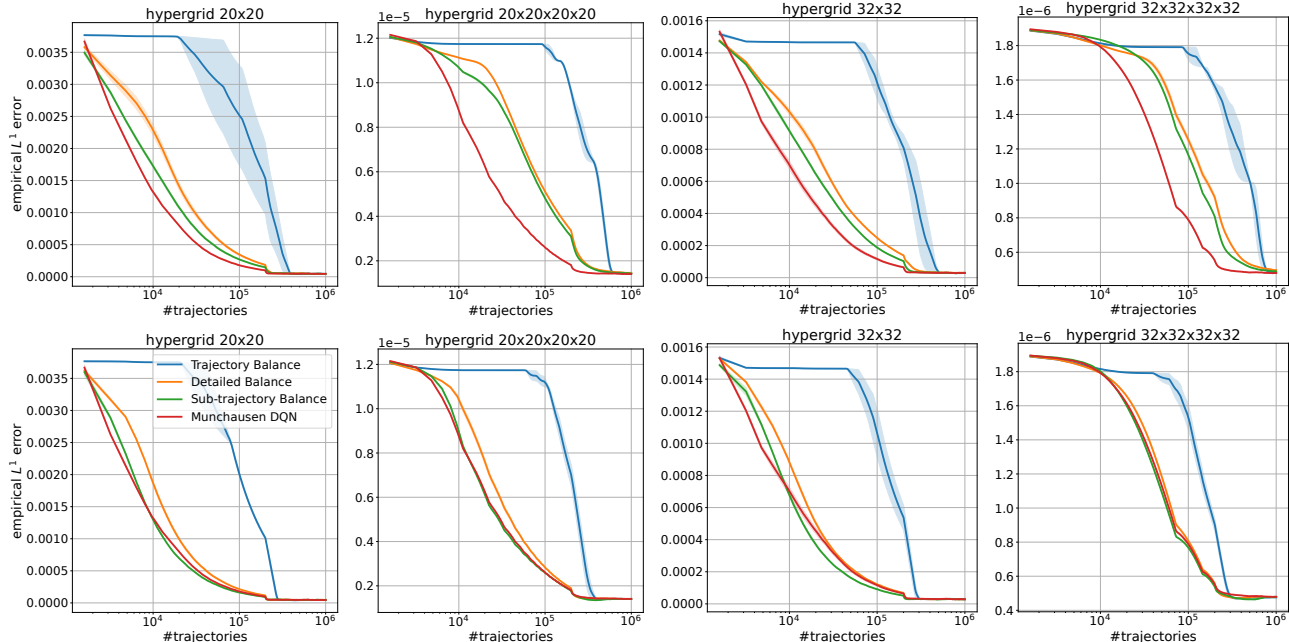


Figure 1: L^1 distance between target and empirical GFlowNet distributions over the course of training on the hypergrid environment. *Top row*: \mathcal{P}_B is fixed to be uniform for all methods. *Bottom row*: \mathcal{P}_B is learnt for the baselines and fixed to be uniform for M-DQN. Mean and std values are computed over 3 runs.

sequentially joining parts from a predefined vocabulary of building blocks (Jin et al., 2020). The state space size is approximately 10^{16} and each state has between 100 and 2000 possible actions. The reward is based on a docking prediction (Trott and Olson, 2010) and is given in a form of pre-trained proxy model $\tilde{\mathcal{R}}$ from Bengio et al. (2021). Additionally, the reward exponent hyperparameter β is introduced, so the GFlowNet is trained using the reward $\mathcal{R}(x) = \tilde{\mathcal{R}}(x)^\beta$.

We follow the setup from Madan et al. (2023). Each model is trained with different values of learning rate and β . \mathcal{P}_B is fixed to be uniform for all methods. To measure how well the trained models match the target distribution, Pearson correlation is computed on a fixed test set of molecules between $\log \mathcal{R}(x)$ and $\log \mathcal{P}_\theta(x)$, where the latter is log probability that x is the end of a trajectory sampled from \mathcal{P}_F . $\log \mathcal{P}_\theta(x)$ is estimated the same way as in Madan et al. (2023). Additionally, we track the number of Tanimoto-separated modes captured by each method over the course of training, as proposed in Bengio et al. (2021). Further experimental details are given in Appendix D.2.

The results are presented in Figure 2. We find that M-DQN outperforms TB and DB in terms of reward correlation and shows similar (or slightly worse when $\beta = 4$) performance when compared to SubTB. It also has a similar robustness to the choice of learning rate when compared to SubTB. However, M-DQN discovers more modes over the course of training than all baselines.

4.3 Non-Autoregressive Sequence Generation

Finally, we consider the bit sequence generation task from Malkin et al. (2022) and Madan et al. (2023). To create a harder GFlowNet learning task with a non-tree DAG structure, we modify the state and action spaces, allowing non-autoregressive generation akin to the approach used in Zhang et al. (2022b).

The goal is to generate binary strings of a fixed length $n = 120$. The reward is specified by the set of modes $M \subset \mathcal{X} = \{0,1\}^n$ and is defined as $\mathcal{R}(x) = \exp(-\min_{x' \in M} d(x, x'))$, where d is Hamming distance. We construct M the same way as in Malkin et al. (2022) and use $|M| = 60$. Parameter k is introduced as a tradeoff between trajectory length and action space sizes. For varying $k \mid n$, the vocabulary is defined as a set of all k -bit words. The generation starts at a sequence of n/k special "empty" words \emptyset , and at each step the allowed actions are to pick any position with an empty word and replace it with some word from the vocabulary. Thus the state space consists of sequences of n/k words, having one of $2^k + 1$ possible words at each position. The terminal states are the states with no empty words, thus corresponding to binary strings of length n .

To evaluate the model performance, we utilize two metrics used in Malkin et al. (2022) and Madan et al. (2023) for this task: 1) test set Spearman correlation between $\mathcal{R}(x)$ and $\mathcal{P}_\theta(x)$ 2) number of modes captured during

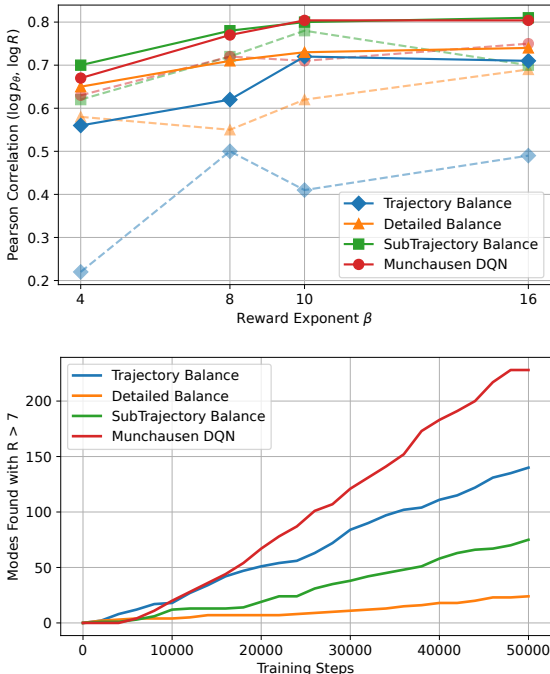


Figure 2: Small molecule generation results. *Above*: Pearson correlation between $\log \mathcal{R}$ and $\log \mathcal{P}_\theta$ on a test set for each method and varying $\beta \in \{4, 8, 10, 16\}$. Solid lines represent the best results over choices of learning rate, dashed lines — mean results. *Below*: Number of Tanimoto-separated modes with $\tilde{\mathcal{R}} > 7.0$ found over the course of training for $\beta = 10$.

training (number of bit strings from M for which a candidate within a distance $\delta = 30$ has been generated). In contrast to the setup in Malkin et al. (2022), the exact probability of sampling x from GFlowNet is intractable in our case due to a large number of paths leading to each x , so we utilize a Monte Carlo estimate proposed in Zhang et al. (2022b). We fix \mathcal{P}_B to be uniform in our experiments. Additional details can be found in Appendix D.3.

Figure 3 presents the results. M-DQN has similar or better reward correlation than all baselines except in the case of $k = 6$, where it outperforms TB and DB but falls behind SubTB. In terms of discovered modes, it shows better performance than DB and similar performance to TB and SubTB.

5 CONCLUSION

In this paper we established a direct connection between GFlowNets and entropy-regularized RL in the general case and experimentally verified it. If a GFlowNet backward policy is fixed, an MDP is appropriately constructed from a GFlowNet DAG and specific rewards are set for all transitions, we showed that learning the GFlowNet forward policy \mathcal{P}_F is equivalent to learn-

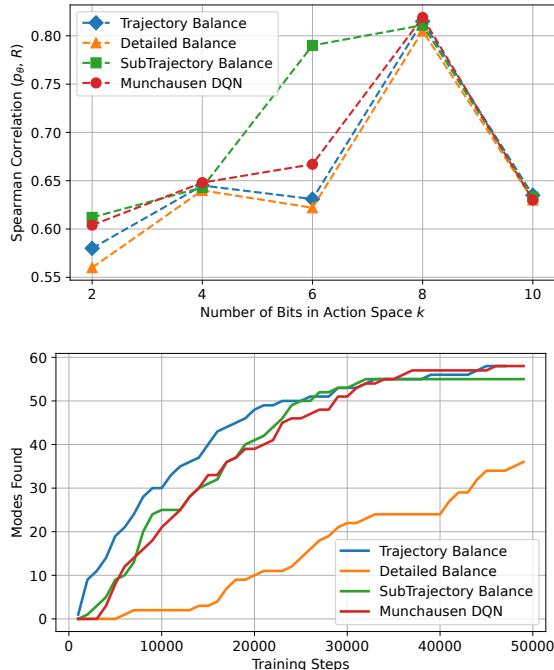


Figure 3: Bit sequence generation results. *Above*: Spearman correlation between \mathcal{R} and \mathcal{P}_θ on a test set for each method and varying $k \in \{2, 4, 6, 8, 10\}$. *Below*: The number of modes discovered over the course of training for $k = 8$.

ing the optimal RL policy π_λ^* when entropy regularization coefficient λ is set to 1. In our experiments, we demonstrated that M-DQN (Vieillard et al., 2020b) shows competitive results when compared to established GFlowNet training methods (Malkin et al., 2022; Bengio et al., 2023; Madan et al., 2023) and even outperforms them in some cases.

A promising direction for further research is to explore theoretical properties of GFlowNets by extending the analysis of Wen and Van Roy (2017). Additionally, it would be valuable to generalize the established connection beyond the discrete setting, following the work of Lahlou et al. (2023a). Furthermore, from a practical perspective there is an interesting direction of extending Monte-Carlo Tree Search-based (MCTS) algorithms to GFlowNets such as AlphaZero (Silver et al., 2018) since known deterministic environments allow to perform MCTS rollouts. Finally, in the context of existing connections between GFlowNets and 1) generative modeling (Zhang et al., 2022a), 2) Markov Chain Monte Carlo methods (Deleu and Bengio, 2023), 3) variational inference (Zimmermann et al., 2022; Malkin et al., 2023), with the addition of the presented link to 4) reinforcement learning, another potential direction can be to view GFlowNets as a unifying point of different ML areas and models, which can lead to further establishment of valuable connections.

Acknowledgements

The experimental results from Section 4 were obtained by Nikita Morozov and Alexey Naumov with the support of the grant for research centers in the field of AI provided by the Analytical Center for the Government of the Russian Federation (ACRF) in accordance with the agreement on the provision of subsidies (identifier of the agreement 000000D730321P5Q0002) and the agreement with HSE University No. 70-2021-00139. Theoretical results from Section 3 were obtained by Daniil Tiapkin with the support by the Paris Île-de-France Région in the framework of DIM AI4IDF. This research was supported in part through computational resources of HPC facilities at HSE University (Kostenetskiy et al., 2021).

References

- Bellemare, M., Srinivasan, S., Ostrovski, G., Schaul, T., Saxton, D., and Munos, R. (2016). Unifying count-based exploration and intrinsic motivation. *Advances in neural information processing systems*, 29.
- Bellemare, M. G., Dabney, W., and Rowland, M. (2023). *Distributional Reinforcement Learning*. MIT Press. <http://www.distributional-rl.org>.
- Bengio, E., Jain, M., Korablyov, M., Precup, D., and Bengio, Y. (2021). Flow network based generative models for non-iterative diverse candidate generation. *Advances in Neural Information Processing Systems*, 34:27381–27394.
- Bengio, Y., Lahlou, S., Deleu, T., Hu, E. J., Tiwari, M., and Bengio, E. (2023). Gflownet foundations. *Journal of Machine Learning Research*, 24(210):1–55.
- Bou, A., Bettini, M., Dittert, S., Kumar, V., Sodhani, S., Yang, X., De Fabritiis, G., and Moens, V. (2023). Torchrl: A data-driven decision-making library for pytorch. *arXiv preprint arXiv:2306.00577*.
- Burda, Y., Edwards, H., Storkey, A., and Klimov, O. (2019). Exploration by random network distillation. In *International Conference on Learning Representations*.
- Christodoulou, P. (2019). Soft actor-critic for discrete action settings. *arXiv preprint arXiv:1910.07207*.
- Cover, T. M. (1999). *Elements of information theory*. John Wiley & Sons.
- Deleu, T. and Bengio, Y. (2023). Generative flow networks: a markov chain perspective. *arXiv preprint arXiv:2307.01422*.
- Deleu, T., Góis, A., Emezue, C., Rankawat, M., Lacoste-Julien, S., Bauer, S., and Bengio, Y. (2022). Bayesian structure learning with generative flow networks. In *Uncertainty in Artificial Intelligence*, pages 518–528. PMLR.
- Fox, R., Pakman, A., and Tishby, N. (2016). Taming the noise in reinforcement learning via soft updates. In Ihler, A. and Janzing, D., editors, *Proceedings of the Thirty-Second Conference on Uncertainty in Artificial Intelligence, UAI 2016, June 25-29, 2016, New York City, NY, USA*. AUAI Press.
- Geist, M., Scherrer, B., and Pietquin, O. (2019). A theory of regularized markov decision processes. In *International Conference on Machine Learning*, pages 2160–2169. PMLR.
- Gilmer, J., Schoenholz, S. S., Riley, P. F., Vinyals, O., and Dahl, G. E. (2017). Neural message passing for quantum chemistry. In *International conference on machine learning*, pages 1263–1272. PMLR.
- Haarnoja, T., Tang, H., Abbeel, P., and Levine, S. (2017). Reinforcement learning with deep energy-based policies. In *International conference on machine learning*, pages 1352–1361. PMLR.
- Haarnoja, T., Zhou, A., Abbeel, P., and Levine, S. (2018). Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. In Dy, J. and Krause, A., editors, *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pages 1861–1870. PMLR.
- Huang, S., Dossa, R. F. J., Ye, C., Braga, J., Chakraborty, D., Mehta, K., and Araújo, J. G. (2022). Cleanrl: High-quality single-file implementations of deep reinforcement learning algorithms. *Journal of Machine Learning Research*, 23(274):1–18.
- Jin, W., Barzilay, R., and Jaakkola, T. (2020). Junction tree variational autoencoder for molecular graph generation. In *Artificial Intelligence in Drug Discovery*, pages 228–249. The Royal Society of Chemistry.
- Kostenetskiy, P., Chulkevich, R., and Kozyrev, V. (2021). Hpc resources of the higher school of economics. In *Journal of Physics: Conference Series*, volume 1740, page 012050. IOP Publishing.
- Lahlou, S., Deleu, T., Lemos, P., Zhang, D., Volokhova, A., Hernández-García, A., Ezzine, L. N., Bengio, Y., and Malkin, N. (2023a). A theory of continuous generative flow networks. In *International Conference on Machine Learning*, pages 18269–18300. PMLR.
- Lahlou, S., Viviano, J. D., and Schmidt, V. (2023b). torchgfn: A pytorch gflownet library. *arXiv preprint arXiv:2305.14594*.
- Lillicrap, T. P., Hunt, J. J., Pritzel, A., Heess, N., Erez, T., Tassa, Y., Silver, D., and Wierstra, D. (2015).

- Continuous control with deep reinforcement learning. *arXiv preprint arXiv:1509.02971*.
- Lin, L.-J. (1992). Self-improving reactive agents based on reinforcement learning, planning and teaching. *Machine learning*, 8:293–321.
- Madan, K., Rector-Brooks, J., Korablyov, M., Bengio, E., Jain, M., Nica, A. C., Bosc, T., Bengio, Y., and Malkin, N. (2023). Learning gflownets from partial episodes for improved convergence and stability. In *International Conference on Machine Learning*, pages 23467–23483. PMLR.
- Malkin, N., Jain, M., Bengio, E., Sun, C., and Bengio, Y. (2022). Trajectory balance: Improved credit assignment in gflownets. *Advances in Neural Information Processing Systems*, 35:5955–5967.
- Malkin, N., Lahlou, S., Deleu, T., Ji, X., Hu, E. J., Everett, K. E., Zhang, D., and Bengio, Y. (2023). GFlownets and variational inference. In *The Eleventh International Conference on Learning Representations*.
- Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A. A., Veness, J., Bellemare, M. G., Graves, A., Riedmiller, M., Fidjeland, A. K., Ostrovski, G., et al. (2015). Human-level control through deep reinforcement learning. *nature*, 518(7540):529–533.
- Neu, G., Jonsson, A., and Gómez, V. (2017). A unified view of entropy-regularized markov decision processes. *arXiv preprint arXiv:1705.07798*.
- Osband, I., Blundell, C., Pritzel, A., and Van Roy, B. (2016). Deep exploration via bootstrapped dqn. *Advances in neural information processing systems*, 29.
- Ostrovski, G., Bellemare, M. G., Oord, A., and Munos, R. (2017). Count-based exploration with neural density models. In *International conference on machine learning*, pages 2721–2730. PMLR.
- Pan, L., Zhang, D., Courville, A., Huang, L., and Bengio, Y. (2023). Generative augmented flow networks. In *The Eleventh International Conference on Learning Representations*.
- Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., et al. (2019). Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems*, 32.
- Rector-Brooks, J., Madan, K., Jain, M., Korablyov, M., Liu, C.-H., Chandar, S., Malkin, N., and Bengio, Y. (2023). Thompson sampling for improved exploration in gflownets. *arXiv preprint arXiv:2306.17693*.
- Saade, A., Kapturowski, S., Calandriello, D., Blundell, C., Sprechmann, P., Sarra, L., Groth, O., Valko, M., and Piot, B. (2023). Unlocking the power of representations in long-term novelty-based exploration. *arXiv preprint arXiv:2305.01521*.
- Schaul, T., Quan, J., Antonoglou, I., and Silver, D. (2016). Prioritized experience replay. In Bengio, Y. and LeCun, Y., editors, *4th International Conference on Learning Representations, ICLR 2016, San Juan, Puerto Rico, May 2-4, 2016, Conference Track Proceedings*.
- Schulman, J., Chen, X., and Abbeel, P. (2017a). Equivalence between policy gradients and soft q-learning. *arXiv preprint arXiv:1704.06440*.
- Schulman, J., Wolski, F., Dhariwal, P., Radford, A., and Klimov, O. (2017b). Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*.
- Shen, M. W., Bengio, E., Hajiramezani, E., Loukas, A., Cho, K., and Biancalani, T. (2023). Towards understanding and improving gflownet training. In *Proceedings of the 40th International Conference on Machine Learning, ICML’23*. JMLR.org.
- Silver, D., Hubert, T., Schrittwieser, J., Antonoglou, I., Lai, M., Guez, A., Lanctot, M., Sifre, L., Kumaran, D., Graepel, T., Lillicrap, T., Simonyan, K., and Hassabis, D. (2018). A general reinforcement learning algorithm that masters chess, shogi, and go through self-play. *Science*, 362(6419):1140–1144.
- Silver, D., Lever, G., Heess, N., Degris, T., Wierstra, D., and Riedmiller, M. (2014). Deterministic policy gradient algorithms. In *International conference on machine learning*, pages 387–395. Pmlr.
- Sutton, R. S. (1988). Learning to predict by the methods of temporal differences. *Machine learning*, 3:9–44.
- Sutton, R. S. and Barto, A. G. (2018). *Reinforcement learning: An introduction*. MIT press.
- Tang, Y., Munos, R., Rowland, M., and Valko, M. (2023). VA-learning as a more efficient alternative to q-learning. In Krause, A., Brunskill, E., Cho, K., Engelhardt, B., Sabato, S., and Scarlett, J., editors, *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pages 33739–33757. PMLR.
- Tiapkin, D., Belomestny, D., Calandriello, D., Moulines, E., Munos, R., Naumov, A., Perrault, P., Tang, Y., Valko, M., and Ménard, P. (2023). Fast rates for maximum entropy exploration. In *Proceedings of the 40th International Conference on Machine Learning, ICML’23*. JMLR.org.
- Trott, O. and Olson, A. J. (2010). 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.

- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., and Polosukhin, I. (2017). Attention is all you need. *Advances in neural information processing systems*, 30.
- Vemgal, N., Lau, E., and Precup, D. (2023). An empirical study of the effectiveness of using a replay buffer on mode discovery in gflownets. *arXiv preprint arXiv:2307.07674*.
- Vieillard, N., Kozuno, T., Scherrer, B., Pietquin, O., Munos, R., and Geist, M. (2020a). Leverage the average: an analysis of kl regularization in reinforcement learning. *Advances in Neural Information Processing Systems*, 33:12163–12174.
- Vieillard, N., Pietquin, O., and Geist, M. (2020b). Munchausen reinforcement learning. *Advances in Neural Information Processing Systems*, 33:4235–4246.
- Wang, Z., Schaul, T., Hessel, M., Hasselt, H., Lanctot, M., and Freitas, N. (2016). Dueling network architectures for deep reinforcement learning. In *International conference on machine learning*, pages 1995–2003. PMLR.
- Wen, Z. and Van Roy, B. (2017). Efficient reinforcement learning in deterministic systems with value function generalization. *Mathematics of Operations Research*, 42(3):762–782.
- Zahavy, T., O’Donoghue, B., Desjardins, G., and Singh, S. (2021). Reward is enough for convex mdps. *Advances in Neural Information Processing Systems*, 34:25746–25759.
- Zhang, D., Chen, R. T., Malkin, N., and Bengio, Y. (2022a). Unifying generative models with gflownets. *arXiv preprint arXiv:2209.02606*.
- Zhang, D., Malkin, N., Liu, Z., Volokhova, A., Courville, A., and Bengio, Y. (2022b). Generative flow networks for discrete probabilistic modeling. In *International Conference on Machine Learning*, pages 26412–26428. PMLR.
- Zhang, D., Pan, L., Chen, R. T., Courville, A., and Bengio, Y. (2023). Distributional gflownets with quantile flows. *arXiv preprint arXiv:2302.05793*.
- Zimmermann, H., Lindsten, F., van de Meent, J.-W., and Naesseth, C. A. (2022). A variational perspective on generative flow networks. *arXiv preprint arXiv:2210.07992*.

Checklist

1. For all models and algorithms presented, check if you include:
 - (a) A clear description of the mathematical setting, assumptions, algorithm, and/or model. [Yes, see Section 2 and 3, also Appendix C]
 - (b) An analysis of the properties and complexity (time, space, sample size) of any algorithm. [Yes, see Appendix C]
 - (c) (Optional) Anonymized source code, with specification of all dependencies, including external libraries. [Yes, see Section 1]
2. For any theoretical claim, check if you include:
 - (a) Statements of the full set of assumptions of all theoretical results. [Yes, see Section 3]
 - (b) Complete proofs of all theoretical results. [Yes, see Appendix B]
 - (c) Clear explanations of any assumptions. [Yes, see Section 3]
3. For all figures and tables that present empirical results, check if you include:
 - (a) The code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL). [Yes, see Section 1 and Appendix D]
 - (b) All the training details (e.g., data splits, hyperparameters, how they were chosen). [Yes, see Section 4 and Appendix D]
 - (c) A clear definition of the specific measure or statistics and error bars (e.g., with respect to the random seed after running experiments multiple times). [Yes, see Section 4 and Appendix D]
 - (d) A description of the computing infrastructure used. (e.g., type of GPUs, internal cluster, or cloud provider). [Yes, see Appendix D]
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets, check if you include:
 - (a) Citations of the creator If your work uses existing assets. [Yes, see Appendix D]
 - (b) The license information of the assets, if applicable. [Not Applicable]
 - (c) New assets either in the supplemental material or as a URL, if applicable. [Not Applicable]
 - (d) Information about consent from data providers/curators. [Not Applicable]
 - (e) Discussion of sensible content if applicable, e.g., personally identifiable information or offensive content. [Not Applicable]
5. If you used crowdsourcing or conducted research with human subjects, check if you include:
 - (a) The full text of instructions given to participants and screenshots. [Not Applicable]
 - (b) Descriptions of potential participant risks, with links to Institutional Review Board (IRB) approvals if applicable. [Not Applicable]
 - (c) The estimated hourly wage paid to participants and the total amount spent on participant compensation. [Not Applicable]

A NOTATION

Table 1: Table of notation used throughout the paper.

Notation	Meaning
\mathcal{S}	state space
\mathcal{A}	action space
\mathcal{X}	sampling space
\mathcal{R}	GFlowNet reward function
$\mathcal{F}(\tau), \mathcal{F}(s), \mathcal{F}(s \rightarrow s')$	flow function
$\mathcal{P}(\tau)$	distribution over trajectories induced by a flow
$\mathcal{P}(x)$	distribution over terminal states induced by a flow
$\mathcal{P}_F(s' s)$	GFlowNet forward policy
$\mathcal{P}_B(s s')$	GFlowNet backward policy
$Z, Z_{\mathcal{F}}$	normalizing constant of the reward distribution and a flow \mathcal{F}
$\mathbb{P}(s' s, a)$	Markovian transition kernel for MDP
$r(s, a)$	MDP reward function
λ	regularization coefficient
$V_{\lambda}^{\pi}(s), Q_{\lambda}^{\pi}(s, a)$	regularized value and Q-value of a given policy π
$V_{\lambda}^{\star}(s), Q_{\lambda}^{\star}(s, a)$	regularized value and Q-value of an optimal policy
$V_1^{\pi}(s), Q_1^{\pi}(s, a)$	regularized value and Q-value of a given policy π with coefficient 1
$V_1^{\star}(s), Q_1^{\star}(s, a)$	regularized value and Q-value of an optimal policy with coefficient 1
$\pi_{\lambda}^{\star}(a s)$	regularized optimal policy
$q^{\pi}(\tau)$	distribution over trajecories induced by a policy π
$d^{\pi}(x)$	distribution over terminal states induced by a policy π
Softmax_{λ}	softmax function with a temperature $\lambda > 0$, for $x \in \mathbb{R}^d$ $[\text{Softmax}_{\lambda}(x)]_i \triangleq e^{x_i/\lambda} / (\sum_{j=1}^d e^{x_j/\lambda}) \forall i \in [d]$
$\text{LogSumExp}_{\lambda}$	log-sum-exp function with a temperature $\lambda > 0$, for $x \in \mathbb{R}^d$ $\text{LogSumExp}(x) \triangleq \lambda \log(\sum_{i=1}^d e^{x_i/\lambda})$.
\mathcal{H}	Shannon entropy function, $\mathcal{H}(p) \triangleq \sum_{i=1}^n p_i \log(1/p_i)$.

Let $(\mathcal{X}, \mathcal{X})$ be a measurable space and $\mathcal{P}(\mathcal{X})$ be the set of all probability measures on this space. For $p \in \mathcal{P}(\mathcal{X})$ we denote by \mathbb{E}_p the expectation w.r.t. p . For random variable $\xi : \mathcal{X} \rightarrow \mathbb{R}$ notation $\xi \sim p$ means $\text{Law}(\xi) = p$.

For any $p, q \in \mathcal{P}(\mathcal{X})$ the Kullback-Leibler divergence $\text{KL}(p, q)$ is given by

$$\text{KL}(p, q) \triangleq \begin{cases} \mathbb{E}_p \left[\log \frac{dp}{dq} \right], & p \ll q, \\ +\infty, & \text{otherwise.} \end{cases}$$

B MISSING PROOFS

B.1 Proof of Theorem 1

Let $\mathcal{F} : \mathcal{T} \rightarrow \mathbb{R}_{\geq 0}$ be a Markovian flow uniquely defined by a backward policy \mathcal{P}_B and a GFlowNet reward \mathcal{R} that satisfies reward matching constraint (1). By soft Bellman equations (8), the value and the Q-value of the optimal policy satisfy for all $s \in \mathcal{S} \setminus \mathcal{X}, s' \in \mathcal{A}_s$

$$\begin{aligned} Q_1^{\star}(s, s') &= r(s, s') + V_1^{\star}(s'), \\ V_1^{\star}(s) &= \log \left(\sum_{s' \in \mathcal{A}_s} \exp(Q_1^{\star}(s, s')) \right) \end{aligned} \tag{13}$$

and, additionally, $V_1^{\star}(s_f) = 0$ and $V_1^{\star}(x) = \log \mathcal{R}(x)$ for any $x \in \mathcal{X}$.

Next, we claim that for any $s, s' \in \mathcal{S}$ it holds $V_1^{\star}(s) = \log \mathcal{F}(s)$, $Q_1^{\star}(s, s') = \log \mathcal{F}(s \rightarrow s')$ and prove it by backward induction over the topological ordering of vertices of \mathcal{G} . The base of induction already holds for $V_1^{\star}(x)$

for all $x \in \mathcal{X}$. Let $s \in \mathcal{S} \setminus \mathcal{X}$ be a fixed state and assume that the claim is proven for all its children $s' \in \mathcal{A}_s$. By soft Bellman equations (13), we have for any $s' \in \mathcal{A}_s$

$$\begin{aligned} Q_1^*(s, s') &= \log \mathcal{P}_B(s|s') + V_1^*(s') \\ &= \log \frac{\mathcal{F}(s \rightarrow s')}{\mathcal{F}(s')} + \log \mathcal{F}(s') = \log \mathcal{F}(s \rightarrow s'), \end{aligned}$$

where we used the choice of MDP rewards (11) and the definition of \mathcal{P}_B through the flow function (2) combined with the induction hypothesis. Next, soft Bellman equations imply that

$$\exp(V_1^*(s)) = \sum_{s' \in \mathcal{A}_s} \exp(Q_1^*(s, s')) = \sum_{s' \in \mathcal{A}_s} \mathcal{F}(s \rightarrow s') = \mathcal{F}(s),$$

where the induction hypothesis for Q-values and flow machining constraint (Bengio et al., 2023, Proposition 19) are applied. Finally, using (9) we have the following expression for the optimal policy

$$\pi_1^*(s'|s) = \exp(Q_1^*(s, s') - V_1^*(s)) = \frac{\mathcal{F}(s \rightarrow s')}{\mathcal{F}(s)}$$

that coincides with the forward policy \mathcal{P}_F (2) corresponding to the flow.

B.2 Proof of Proposition 1

Let us take N equal to the maximal length of trajectory $\max_\tau n_\tau$, i.e. at step N agent is guaranteed to reach the absorbing state with zero reward. Then, by the definition (7) of $V_1^\pi(s_0)$ in a regularized MDP with $\gamma = \lambda = 1$ we have

$$V_1^\pi(s_0) = \mathbb{E}_\pi \left[\sum_{t=0}^{N-1} r(s_t, s_{t+1}) - \log \pi(s_{t+1}|s_t) \right],$$

where we used the tower property of conditional expectation to replace entropy with negative logarithm. By substituting the expression for rewards we have

$$\begin{aligned} V_1^\pi(s_0) &= \mathbb{E}_{\tau \sim q^\pi} \left[\sum_{t=1}^{n_\tau} \log \frac{\mathcal{P}_B(s_{t-1}|s_t)}{\pi(s_t|s_{t-1})} + \log \mathcal{R}(s_{n_\tau}) \right] \\ &= \mathbb{E}_{\tau \sim q^\pi} \left[\log \frac{\prod_{t=1}^{n_\tau} \mathcal{P}_B(s_{t-1}|s_t) \cdot \mathcal{R}(s_{n_\tau}) \cdot Z}{\prod_{t=1}^{n_\tau} \pi(s_t|s_{t-1}) \cdot Z} \right] = -\text{KL}(q^\pi \| \mathcal{P}_B) + \log Z. \end{aligned}$$

C DETAILED ALGORITHM DESCRIPTION

In this section we provide a detailed algorithmic description of SoftDQN and MunchausenDQN and discuss the effects of several hyperparameters.

C.1 Soft DQN

The pipeline of this algorithm is as follows. In the beginning of the learning procedure we have two neural networks: Q_θ and $Q_{\bar{\theta}}$, and $\bar{\theta} = \theta$.

On each iteration of the algorithm, several trajectories (usually 4 or 16) are sampled by following a softmax policy $\pi_\theta(s) = \text{Softmax}_\lambda(Q_\theta(s))$ (here $\lambda = 1$) using an online network Q_θ with (optionally) an additional ε -greedy exploration.¹

¹Notably, softmax parameterization naturally induces Boltzman exploration. However, the effects of Boltzman and ε -exploration are different and in general cannot substitute each other, see discussion in (Veillard et al., 2020b).

Algorithm 1 SoftDQN (Haarnoja et al., 2017) / MunchausenDQN (Veillard et al., 2020b)

- 1: **Input:** trajectory budget N , number of trajectories on one update M , exploration parameter ε , backward kernel \mathcal{P}_B , PER \mathcal{B} , PER batch size B , hard update parameter T , soft update parameter τ , Munchausen parameters α, l_0 ;
- 2: Initialize parameters of online network θ and target network $\bar{\theta} = \theta$;
- 3: Define entropy coefficient $\lambda = 1/(1 - \alpha)$
- 4: **for** $t = 1, \dots, \lceil N/M \rceil$ **do**
- 5: Sample M trajectories interaction with ε -greedy version of policy $\pi_\theta = \text{Softmax}_\lambda(Q_\theta)$, and place them into replay buffer \mathcal{B} ;
- 6: Sample B transitions $\{(s_k, a_k, R_k, s'_k, d_k)\}_{k=1}^B$ from PER \mathcal{B} , where d_k is a DONE flag: $d_k = \mathbb{1}\{s_k \in \mathcal{X}\}$, and R_k is GflowNet reward is $d_k = 1$ and 0 otherwise;
- 7: Compute MDP rewards: $r_k = d_k \log R_k + (1 - d_k) \log \mathcal{P}_B(s_k|s'_k)$ with a convention $0 \log 0 = 0$ for all $k \in \{1, \dots, B\}$.
- 8: Compute TD targets

$$y_k = r_k + (1 - d_k) \text{LogSumExp}_\lambda(Q_{\bar{\theta}}(s'_k)) + \alpha \max\{\lambda \log \pi_{\bar{\theta}}(a_k|s_k), l_0\}$$

optionally using a convention $\text{LogSumExp}_\lambda(Q_{\bar{\theta}}(s'_k)) = \log \mathcal{R}(s'_k) \forall s'_k \in \mathcal{X}$ (see explanation below);

- 9: Use y_k to update priorities in PER (Schaul et al., 2016) as

$$p_k = \ell(Q_\theta(s_k, a_k), y_k),$$

where ℓ is a Huber loss (15).

- 10: Compute TD loss

$$\mathcal{L}(\theta) = \frac{1}{B} \sum_{k=1}^B \ell(Q_\theta(s_k, a_k), y_k),$$

where ℓ is a Huber loss (15).

- 11: Perform gradient update by $\nabla_\theta \mathcal{L}(\theta)$;
- 12: **if** $t \bmod T = 0$ **then**
- 13: Update target network $\bar{\theta} := (1 - \tau)\bar{\theta} + \tau\theta$;
- 14: **end if**
- 15: **end for**
- 16: **Output** policy π_θ .

All these sampled trajectories are separated onto transitions $(s_t, a_t, s_{t+1}, R_t, d_t)$, where $R_t = \log \mathcal{R}(s_t)$ if $s_t \in \mathcal{X}$ and zero otherwise, $d_t = \mathbb{1}\{s_t \in \mathcal{X}\}$ is a DONE float. Note that we always have $a_t = s_{t+1}$, and a_t is used solely for the purpose of consistency with common RL notation. These separated transitions are placed into a replay buffer \mathcal{B} (Lin, 1992) with a prioritization mechanism following Schaul et al. (2016).

Next, a batch of transitions $\{(s_k, a_k, s'_k, R_k, d_k)\}_{k=1}^B$ is sampled from the replay buffer \mathcal{B} using prioritization. MDP rewards are computed as $r_k = d_k \log \mathcal{R}(s_k) + (1 - d_k) \log \mathcal{P}_B(s_k|s'_k)$, and TD-targets are computed as follows:

$$\forall k \in [B] : y_k = r_k + (1 - d_k) \text{LogSumExp}_\lambda(Q_{\bar{\theta}}(s'_k)), \quad (14)$$

where $\lambda = 1$.

Alternatively, one can fully skip transitions that lead to the sink state s_f (transitions for which $d_k = 1$), and instead substitute $\text{LogSumExp}_\lambda(Q_{\bar{\theta}}(s'_k)) = \log \mathcal{R}(s'_k) \forall s'_k \in \mathcal{X}$. This can be done since we know the exact value of $V_\lambda^*(s)$ for all $s \in \mathcal{X}$. We follow this choice in our experiments.

Given targets, we update priorities for a sampled batch in the replay buffer and compute the regression TD-loss as follows:

$$\mathcal{L}(\theta) = \frac{1}{B} \sum_{k=1}^B \ell(Q_\theta(s_k, a_k), y_k),$$

where ℓ is a Huber loss

$$\ell(x, y) = \begin{cases} \frac{1}{2}(x - y)^2 & |x - y| \leq 1 \\ |x - y| - 1/2 & |x - y| \geq 1. \end{cases} \quad (15)$$

Finally, at the end of each iteration, we perform either a hard update by copying weights of an online network to weights of the target network following the original DQN approach (Mnih et al., 2015) or a soft update by applying Polyak averaging with a parameter $\tau < 1$ following DDPG approach (Silver et al., 2014; Lillicrap et al., 2015).

The detailed description of this algorithm is also provided in Algorithm 1. Notice that the complexity of this algorithm is linear in the number of sampled transitions since the size of the buffer is fixed during the training.

Next, we discuss some elements of this pipeline in more detail.

Prioritized experience replay The essence of prioritized experience replay by (Schaul et al., 2016) is to sample observations from a replay buffer not uniformly but according to priorities that are computed as follows:

$$P(i) = \frac{|p_i|^\alpha}{\sum_{i \in \mathcal{B}} |p_i|^\alpha},$$

where $P(i)$ is a probability to sample i -th transition in the replay buffer \mathcal{B} , p_i is a priority that is equal to the last computed TD-error for this transition, and α is a temperature parameter that interpolates between uniform sampling $\alpha = 0$ and pure prioritization $\alpha = +\infty$. A typically chosen value of α is between 0.4 and 0.6, however, for two experiments (molecule and bit sequence generation) we found it more beneficial to select an aggressive value $\alpha = 0.9$ to collect hard examples faster.

Additionally, authors of (Schaul et al., 2016) applied importance-sampling (IS) correction since the theory prescribes to sample uniformly over the buffer due to stochastic next state s'_k . Smoothed correction is controlled with a coefficient β that should be annealed to 1 over the course of training.

However, in the setup of GFlowNets we noticed that this correction is not needed since the next transition and reward are always deterministic; therefore it is possible to set β to be a small constant without any loss of convergence properties. We set $\beta = 0.0$ in the hypergrid experiment and $\beta = 0.1$ in the molecule and bit sequence generation experiments and found it more beneficial than doing higher IS-correction.

Choice of loss We examine two types of regression losses: classic MSE loss and Huber loss that was applied in DQN (Mnih et al., 2015) and MunchausenDQN (Vieillard et al., 2020b). The main idea behind the Huber loss is automatic gradient clipping for very large updates. We observed such benefits as more stable training while using Huber loss, so we utilized it in all our experiments except some ablation studies in Appendix E.

C.2 Munchausen DQN

The differences between SoftDQN and M-DQN in the context of GFlowNets are:

- Two additional parameters: α and truncation parameter l_0 ;
- Using entropy coefficient λ equal to $1/(1 - \alpha)$ instead of just $\lambda = 1$. In particular, it changes how policy $\pi_\theta(s) = \text{Softmax}_{1/(1-\alpha)}(Q_\theta(s))$ and value $V_\theta(s) = \text{LogSumExp}_{1/(1-\alpha)}(Q_\theta(s))$ are computed;
- Additional term $\alpha \max\{\lambda \log \pi_\theta(a_k | s_k), l_0\}$ in the computation of TD-target (14).

All these differences are highlighted in Algorithm 1 with a blue color. The use of a different entropy coefficient is justified by Theorem 1 of Vieillard et al. (2020b): M-DQN with entropy coefficient λ and Munchausen coefficient α is equivalent to solving entropy-regularized RL with coefficient $(1 - \alpha)\lambda$.

D EXPERIMENTAL DETAILS

In this section we provide additional details for the experiments from Section 4.

D.1 Hypergrids

The reward at $s^\top = (s^1, \dots, s^D)^\top$ is formally defined as

$$\begin{aligned} \mathcal{R}(s^\top) = & R_0 + R_1 \prod_{i=1}^D \mathbb{I} \left[0.25 < \left| \frac{s^i}{H-1} - 0.5 \right| \right] \\ & + R_2 \prod_{i=1}^D \mathbb{I} \left[0.3 < \left| \frac{s^i}{H-1} - 0.5 \right| < 0.4 \right], \end{aligned}$$

where $0 < R_0 \ll R_1 < R_2$. This reward has modes near the corners of the grid which are separated by wide troughs with a very small reward of R_0 . In Section 4 we use reward parameters ($R_0 = 10^{-3}$, $R_1 = 0.5$, $R_2 = 2.0$) taken from Malkin et al. (2022).

All models are parameterized by MLP of the same architecture as in Bengio et al. (2021) with 2 hidden layers and 256 hidden size. Following Malkin et al. (2022), we train all models with Adam optimizer and use a learning rate of 10^{-3} and a batch size 16 (number of trajectories sampled at each training step). In case of TB, we use a learning rate of 10^{-1} for Z_θ following Malkin et al. (2022). For SubTB parameter λ (not to be confused with soft RL regularization coefficient) we use the value of 0.9 following Madan et al. (2023). All models are trained until 10^6 trajectories are sampled and the empirical sample distribution $\pi(x)$ is computed over the last $2 \cdot 10^5$ samples seen in training, which explains the drop in plots at the $2 \cdot 10^5$ mark (see Figure 1).

For M-DQN we use Huber loss instead of MSE, following Mnih et al. (2015); Vieillard et al. (2020b) and prioritized replay buffer (Schaul et al., 2016) from TorchRL library (Bou et al., 2023) of size 100,000, a standard constant $\alpha = 0.5$ and without IS correction $\beta = 0$. At each training step, we sample 16 trajectories, put all transitions to the buffer, and then sample 256 transitions from the buffer to calculate the loss. Thus M-DQN uses the same number of trajectories and the same number of gradient steps for training as the baselines. Regarding specific Munchausen parameters (Vieillard et al., 2020b), we use $\alpha = 0.15$ and unusually large $l_0 = -100$, since the entropy coefficient for our setup is much larger than one that was used in the original M-DQN. To update the target network, we utilize soft updates akin to celebrated DDPG (Silver et al., 2014; Lillicrap et al., 2015) with a large coefficient $\tau = 0.25$ due to sampling 16 trajectories prior to obtaining a batch.

For hypergrid environment experiments we utilized torchgfn library (Lahlou et al., 2023b). All experiments were performed on CPUs.

D.2 Molecules

All models are parameterized with MPNN (Gilmer et al., 2017), for which we use the same parameters as in previous works (Bengio et al., 2021; Malkin et al., 2022; Madan et al., 2023). The reward proxy and the test set are also taken from Bengio et al. (2021). Following Madan et al. (2023), we consider reward exponents from $\{4, 8, 10, 16\}$, learning rates from $\{5 \times 10^{-5}, 10^{-4}, 5 \times 10^{-4}, 10^{-3}\}$ and train all models for 50,000 iterations with a batch size of 4 using Adam optimizer. For SubTB we use the same value $\lambda = 1.0$ as in Madan et al. (2023). Reward correlations for the baselines in Figure 2 are taken directly from Madan et al. (2023). We track Tanimoto-separated modes as described in Bengio et al. (2021), using a raw reward threshold of 7.0 and a Tanimoto similarity threshold of 0.7.

Regarding M-DQN, we use Huber loss instead of MSE, following Mnih et al. (2015); Vieillard et al. (2020b) and find great benefits from utilizing dueling architecture (Wang et al., 2016), which builds similarity with DB (see Section 3.4). We also utilize prioritized replay buffer (Schaul et al., 2016) from TorchRL library (Bou et al., 2023) of size 1,000,000, and unusual parameters $\alpha = 0.9$, $\beta = 0.1$ (see Appendix C for details on this choice), and use Munchausen parameters $\alpha = 0.15$ and $l_0 = -2500$. At each training step, we sample 4 trajectories, put all transitions to the buffer, and then sample 256 transitions from the buffer to calculate the loss. To update the target network, we utilize soft updates akin to DDPG (Silver et al., 2014; Lillicrap et al., 2015) with a coefficient $\tau = 0.1$. Notice that M-DQN samples exactly the same number of trajectories for training as the baselines, thus the comparison in terms of discovered modes is fair.

Our experiments are based upon the published code of Bengio et al. (2021) and Malkin et al. (2022). We utilized a cluster with NVIDIA V100 and NVIDIA A100 GPUs for molecule generation experiments.

D.3 Bit Sequences

Mode set M and the test set are constructed as described in Malkin et al. (2022). We also use the same Transformer (Vaswani et al., 2017) neural network architecture for all models with 3 hidden layers, 64 hidden dimension and 8 attention heads, with the only difference from Malkin et al. (2022) that it outputs logits for a larger action space since we generate sequences non-autoregressively.

To approximate $\mathcal{P}_\theta(x)$ we use the Monte Carlo estimate proposed in Zhang et al. (2022b):

$$\mathcal{P}(x) = \mathbb{E}_{\mathcal{P}_B(\tau|x)} \frac{\mathcal{P}_F(\tau)}{\mathcal{P}_B(\tau|x)} \approx \frac{1}{N} \sum_{i=1}^N \frac{\mathcal{P}_F(\tau^i)}{\mathcal{P}_B(\tau^i|x)},$$

where $\mathcal{P}_F(\tau) = \prod_{t=1}^{n_\tau} \mathcal{P}_F(s_t|s_{t-1}, \theta)$, $\mathcal{P}_B(\tau|x) = \prod_{t=1}^{n_\tau} \mathcal{P}_B(s_{t-1}|s_t)$ and $\tau^i \sim \mathcal{P}_B(\tau|x)$. Zhang et al. (2022b) showed that even with $N = 10$ this approach can provide an adequate estimate, so we use this value in our experiments.

We train all models for 50,000 iterations with a batch size 16 using Adam optimizer. For all methods the trajectories are sampled using a mixture of \mathcal{P}_F and a uniform distribution over next possible actions, the latter having a weight of 10^{-3} (ϵ -greedy exploration). We set the reward exponent value to 2. For M-DQN and all baselines we pick the best learning rate from $\{5 \times 10^{-4}, 10^{-3}, 2 \times 10^{-3}\}$. For SubTB we pick the best λ from $\{0.9, 1.1, 1.5, 1.9\}$.

For M-DQN we use Huber loss instead of MSE, following (Mnih et al., 2015; Vieillard et al., 2020b) and prioritized replay buffer (Schaul et al., 2016) from TorchRL library (Bou et al., 2023) of size 100,000, and constants $\alpha = 0.9, \beta = 0.1$. At each training step, we sample 16 trajectories and put all transitions to the buffer, and then sample 256 transitions from the buffer to calculate the loss. Regarding specific Munchausen parameters (Vieillard et al., 2020b), we use $\alpha = 0.15$ and $l_0 = -25$. We only use hard updates for the target network (Mnih et al., 2015) with a frequency of 5 iterations. Notice that M-DQN samples exactly the same number of trajectories for training as the baselines, thus the comparison in terms of discovered modes is fair. Additionally, for M-DQN we found it helpful to increase the weight of the loss for edges leading to terminal states, which was also done in the original GFlowNet paper (Bengio et al., 2021) for flow matching objective, and pick the best weight coefficient from $\{2, 5\}$ in our experiments.

For bit sequence experiments we used our own code implementation in PyTorch (Paszke et al., 2019). We also utilized a cluster with NVIDIA V100 and NVIDIA A100 GPUs.

E ADDITIONAL COMPARISONS

In this section we provide additional experimental comparisons and ablation studies.

E.1 Hypergrids with Hard Reward

We provide experimental results for hypergrids with a harder reward variant proposed in Madan et al. (2023), which has parameters ($R_0 = 10^{-4}, R_1 = 1.0, R_2 = 3.0$). We follow the same experimental setup as in Section 4. The results are presented in Figure 4. We make similar observations to the ones described in Section 4: when \mathcal{P}_B is fixed to be uniform for all methods, M-DQN converges faster than the baselines, and when \mathcal{P}_B is learnt for the baselines, M-DQN shows comparable performance to SubTB and outperforms TB and DB. Notably, TB fails to converge in most cases, which coincides with the results presented in Madan et al. (2023).

E.2 Training Backward Policy

Our theoretical framework does not support running soft RL with a learnable backward policy, thus in Section 4 we used M-DQN with uniform \mathcal{P}_B in all cases. Nevertheless, one can try to use the same M-DQN training loss to calculate gradients and update \mathcal{P}_B , thus creating a setup with a learnable backward policy. Figure 5 presents the results of comparison on the hypergrid environment (with a standard reward, the same as in Section 4). We find that having a trainable \mathcal{P}_B does not improve the speed of convergence for M-DQN, in contrast to the behavior of a learnt \mathcal{P}_B in the case of the previously proposed GFlowNet training methods presented in Figure 1.

Another possible approach can be to use some other GFlowNet training objective to calculate a loss for \mathcal{P}_B , for example TB, while using soft RL to train \mathcal{P}_F . We leave this case for further study.

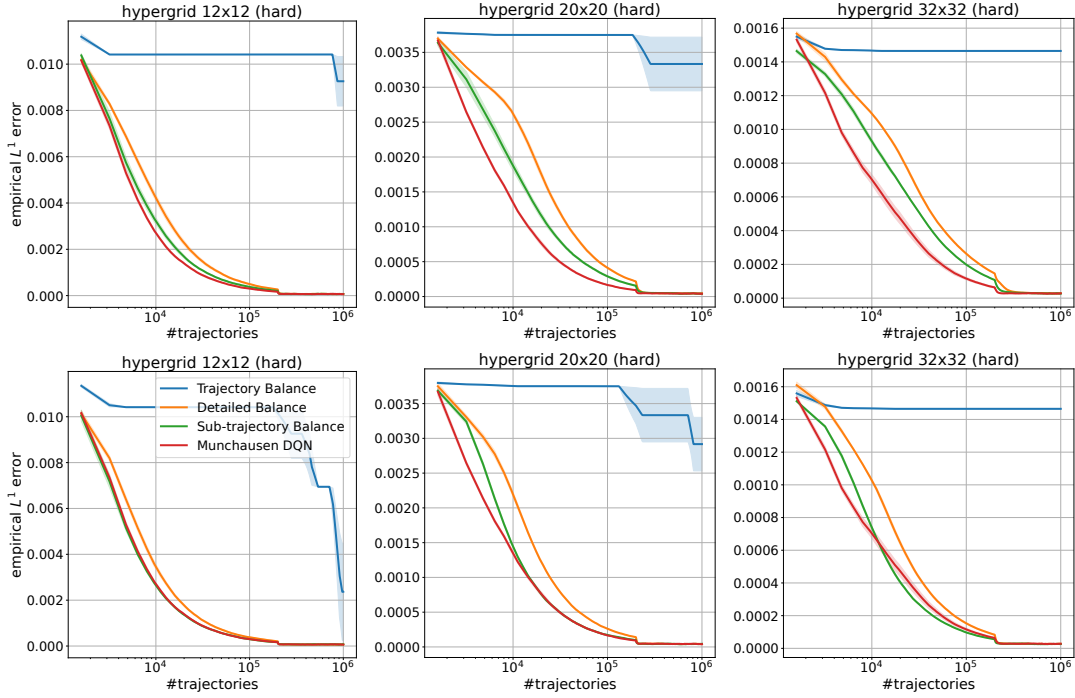


Figure 4: L^1 distance between target and empirical GFlowNet distributions over the course of training on the hypergrid environment for the hard reward variant. *Top row*: \mathcal{P}_B is fixed to be uniform for all methods. *Bottom row*: \mathcal{P}_B is learnt for the baselines and fixed to be uniform for M-DQN.

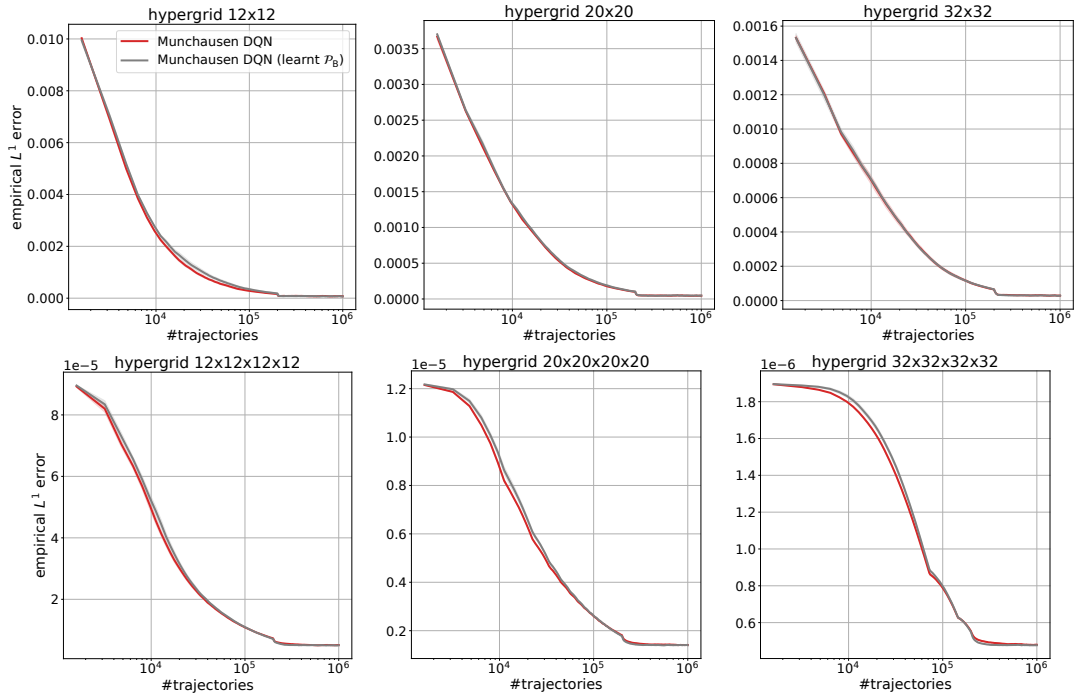


Figure 5: L^1 distance between target and empirical GFlowNet distributions over the course of training on the hypergrid environment. Here we present two versions of M-DQN: one with a fixed uniform \mathcal{P}_B and another one with a learnt \mathcal{P}_B during the training.

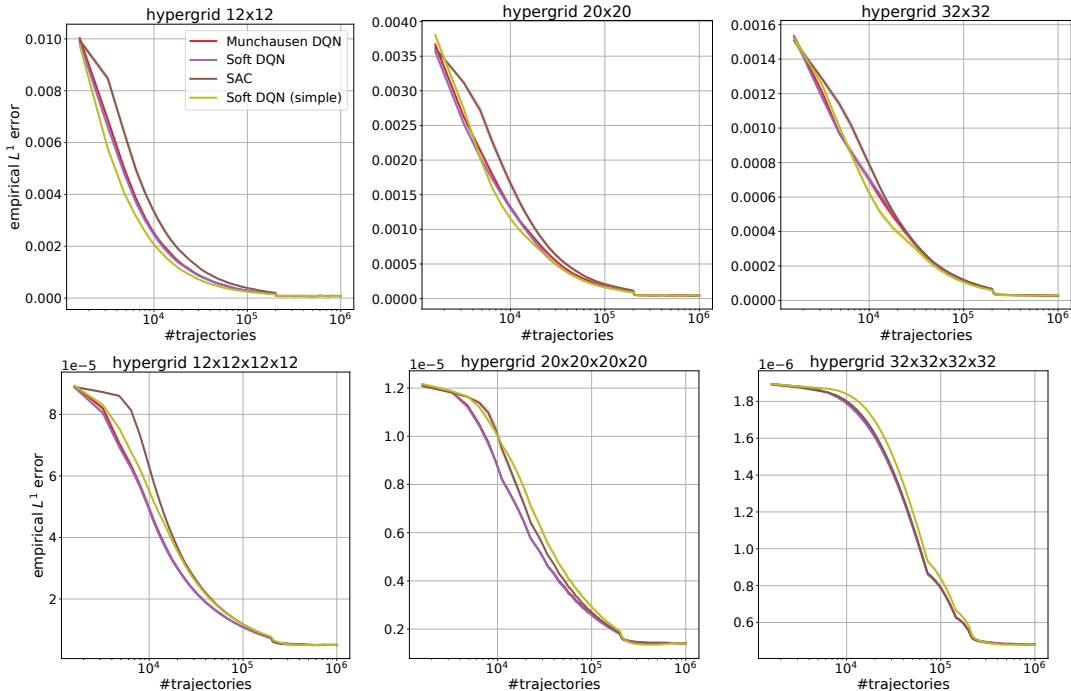


Figure 6: L^1 distance between target and empirical GFlowNet distributions over the course of training on the hypergrid environment. \mathcal{P}_B is fixed to be uniform for all methods. Here we compare different soft RL algorithms.

E.3 Different Soft RL algorithms

E.3.1 Hypergrid Environment

We also compare M-DQN to other known off-policy soft RL algorithms:

Soft DQN To ablate the effect of Munchausen penalty in M-DQN we compare it to a classical version of **SoftDQN**, which is equivalent to setting the Munchausen constant $\alpha = 0$. All other hyperparameters are chosen exactly the same as for M-DQN (see Appendix D).

Soft DQN (simple) To ablate the effect of using a replay buffer and Huber loss, we also train a simplified version of **SoftDQN**. At each training iteration, it samples a batch of trajectories and uses the sampled transitions to calculate the loss, similarly to the way GFlowNet is trained with DB objective. It also uses MSE loss instead of Huber.

Soft Actor Critic We also considered **SAC** (Haarnoja et al., 2018) algorithm, namely its discrete counterpart presented by Christodouλου (2019). We follow the implementation presented in CleanRL library (Huang et al., 2022) with two differences: usage of Huber loss (Mnih et al., 2015) and lack of entropy regularization parameter tuning (since our theory prescribes to set its value to 1). All other hyperparameters are chosen exactly the same as for M-DQN and **SoftDQN** (see Appendix D).

Figure 6 presents the results of comparison on the hypergrid environment (with standard reward, the same as in Section 4). We found that M-DQN and **SoftDQN** generally have very similar performance, while converging faster than SAC in most cases. Interestingly, the simplified version of **SoftDQN** holds its own against other algorithms, converging faster on smaller 2-dimensional grids while falling behind M-DQN and **SoftDQN** on 4-dimensional ones. This indicates that the success of soft RL is not only due to using replay buffer and Huber loss instead of MSE, but the training objective itself is well-suited for the GFlowNet learning problem.

E.3.2 Small Molecule Generation

We also compare M-DQN and SoftDQN on the molecule generation task. The experimental setup is the same as the one used in Section 4. Figure 7 presents the results. SoftDQN shows worse performance when $\beta = 16$ in comparison to M-DQN. But in all other cases, it shows comparable or better reward correlations and is more robust to the choice of learning rate. However, M-DQN discovers modes faster.




Figure 7: Small molecule generation results. *Left*: Pearson correlation between $\log \mathcal{R}$ and $\log \mathcal{P}_\theta$ on a test set for each method and varying $\beta \in \{4, 8, 10, 16\}$. Solid lines represent the best results over choices of learning rate, dashed lines — mean results. *Right*: Number of Tanimoto-separated modes with $\tilde{\mathcal{R}} > 7.0$ found over the course of training for $\beta = 10$.

E.3.3 Bit Sequence Generation

Finally, we compare M-DQN and SoftDQN on the bit sequence generation task. The experimental setup is the same as the one used in Section 4. Figure 8 presents the results. SoftDQN has comparable or worse reward correlations in comparison to M-DQN and discovers modes slightly slower.




Figure 8: Bit sequence generation results. *Left*: Spearman correlation between \mathcal{R} and \mathcal{P}_θ on a test set for each method and varying $k \in \{2, 4, 6, 8, 10\}$. *Right*: The number of modes discovered over the course of training for $k = 8$.