DISCRETE COMPOSITIONAL GENERATION VIA GENERAL SOFT OPERATORS AND ROBUST REINFORCE-MENT LEARNING

Anonymous authors

000

001

002

004

006

008 009 010

011 012

013

014

016

018

019

021

024

025

026

027

028

031

032

033

034

037

038

040

041

042

043

044

046

047

051

052

Paper under double-blind review

ABSTRACT

A major bottleneck in scientific discovery consists of narrowing an exponentially large set of objects, such as proteins or molecules, to a small set of promising candidates with desirable properties. While this process can rely on expert knowledge, recent methods leverage reinforcement learning (RL) guided by a proxy reward function to enable this filtering. By employing various forms of entropy regularization, these methods aim to learn samplers that generate diverse candidates that are highly rated by the proxy function. In this work, we make two main contributions. First, we show that these methods are liable to generate overly diverse, suboptimal candidates in large search spaces. To address this issue, we introduce a novel unified operator that combines several regularized RL operators into a general framework that better targets peakier sampling distributions. Secondly, we offer a novel, robust RL perspective of this filtering process. The regularization can be interpreted as robustness to a compositional form of uncertainty in the proxy function (i.e., the true evaluation of a candidate differs from the proxy's evaluation). Our analysis leads us to a novel, easy-to-use algorithm we name trajectory general mellowmax (TGM): we show it identifies higher quality, diverse candidates than baselines in both synthetic and real-world tasks.

1 Introduction

The task of scientific discovery centers on discovering novel objects with desirable properties such as antimicrobial resistance or binding affinity. These objects are often discrete and structured such that they can be constructed through a sequence of compositional steps, e.g., proteins as sequences of amino acids, molecules composed of various smaller fragments (Angermueller et al., 2019; Wang et al., 2023). Due to this compositionality, the number of possible objects grows exponentially with object size and becomes too large to search exhaustively. Additionally, experimentally evaluating potential objects can be too technical or too expensive. Nonetheless, in many cases, researchers have access to a proxy reward model that quantitatively approximates the satisfaction level of the property of interest (Davis et al., 2018; Du et al., 2022).

A large body of work leverages this proxy reward function to narrow the set of objects to a small subset of promising candidates to test in a lab (Currin et al., 2015; Gubernatis & Lookman, 2018; Wu et al., 2021; AI4Science et al., 2023). An alternative approach that has demonstrated growing popularity and effectiveness is to pose this problem as an RL task (De Cao & Kipf, 2018; Darvariu et al., 2021; Bou et al., 2024). In what we will denominate discrete compositional processes (DCP), object construction is viewed as a Markov decision process (MDP) where terminal states correspond to complete objects and yield a reward given by the evaluation of that object by the proxy reward function. Then, the learned policy can be used as a generative model that samples promising candidates.

Among these approaches, generative flow networks (GFNs) have shown considerable success (Bengio et al., 2021; Jain et al., 2023; Zhou et al., 2024; Cipcigan et al., 2024). GFNs learn an amortized policy and, at optimality, sample objects with probability proportional to the exponential of their reward (Eq. 2). Connections with maximum-entropy RL have been established, showing that this objective underpins the diversity of generated samples (Tiapkin et al., 2024; Mohammadpour et al., 2024a; Deleu et al., 2024). While this exact sampling goal is useful in certain applications, it is problematic in exponentially large sets: the probability of sampling from a small set of objects with high rewards is

Figure 1: (Left) Illustrated issue with sampling proportional to reward. The much larger reward of the optimal sequence (in green) is drowned out by all the rewards in the combinatorial explosion of longer subsequences. (Right) A DCP for a protein sequence generation task. (A) Starting from an empty string, amino acids are added sequentially until termination. (B) Then, the full sequence is evaluated by a proxy reward function $\Phi(x)$ whose value is given as reward for the termination action. (C) Over the course of protein generation, uncertainty accumulates through the δ_i . (D) The true reward depends on $\Phi(x)$ and the accumulated uncertainty.

dominated by the probability of sampling from an exponentially large set of objects with low rewards. In this work, we leverage the connection between GFNs and regularized RL to use soft RL operators for scientific discovery. By combining these approaches, we derive an improved operator for the task and offer a new perspective on regularization. Our contributions can be summarized as follows.

Our contributions: (1) We introduce general mellowmax, an interpolating operator that allows for a more flexible tradeoff between quality and diversity than previous operators; (2) We derive a trajectory constraint for this operator, which we refer to as **trajectory general mellowmax** (TGM). TGM enables us to propose a GFN-like algorithm for scientific discovery tasks that empirically finds diverse modes with higher reward than soft mellowmax, GFNs, and other RL baselines; (3) We provide a robust RL interpretation of these baselines in DCPs. Specifically, we show that the reward uncertainty sets induced by TGM are more interpretable than those associated with entropy regularization or GFNs. The analytical techniques we employ pave the way to more general trajectory constraints through an interpretable equivalence between regularization and reward robust sets.

2 RELATED WORK

The goal of sampling proportional to reward targeted by GFNs was historically tackled using Markov chain Monte Carlo (MCMC) (Gilks et al., 1995). These methods construct chains through a random walk: from a set of initial points, new points are proposed and either accepted or rejected based on their function value (here, their proxy reward). While theoretically appealing, these methods struggle in high dimensions, in particular when it comes to discovering separate modes (Bengio et al., 2021).

On the other hand, GFNs have demonstrated success in protein design (Jain et al., 2022), drug discovery (Shen et al., 2023), and material design (AI4Science et al., 2023). The standard way to tackle the excessive smoothness of GFN policies is through temperature conditioning (Zhang et al., 2023; Kim et al., 2023; Zhou et al., 2024): the GFN target can be posed as learning $p(x) \propto e^{\beta r(x)}$ where β is an inverse temperature parameter. Temperature conditioning involves learning a policy conditionally on β , with β varying throughout training. It complements our approach, as temperature conditioning can be used on top of our framework. Mohammadpour et al. (2024b) have also proposed a dynamic regularization coefficient based on the number of actions, for maximum entropy RL. This change can also be used to address excessive smoothness in GFNs or incorporated through a varying ω_s in our Eq. 6.

From an RL perspective, many soft RL methods fit in the regularized framework of Geist et al. (2019). Other forms of regularization and their corresponding operators include the mellowmax operator (Asadi & Littman, 2017), which takes the logmeanexp of Q-values instead of the logsumexp, and its extension, the soft mellowmax operator (Gan et al., 2021). Jiralerspong et al. (2024) also combine a GFN perspective with RL by developing an adversarial version of the mellowmax operator called AFlowNets. None of these works use their operators for DCPs, whereas we generalize them, integrate them with trajectory constraints, and successfully apply these constraints to real DCP tasks.

3 BACKGROUND AND MOTIVATION

MDPs. An MDP is a tuple $(\mathcal{S}, \mathcal{A}, \gamma, P, r)$ where \mathcal{S} and \mathcal{A} are finite state and action spaces respectively, $\gamma \in [0,1)$ is a discount factor, $P: \mathcal{S} \times \mathcal{A} \to \Delta_{\mathcal{S}}$ a transition kernel and $r: \mathcal{S} \times \mathcal{A} \to \mathbb{R}$ a reward function. Define Π as the set of policy mappings $\pi: \mathcal{S} \to \Delta_{\mathcal{A}}$. The goal is to find $\pi \in \Pi$ that maximizes the value function $v_r^\pi(s) := \mathbb{E}^\pi[\sum_{t=0}^\infty \gamma^t r(s_t, a_t) \mid s_0 = s], \ \forall s \in \mathcal{S}$. The dependence

of the value function on the reward function is made explicit through the subscript r, as clarified in the next paragraph. For any policy $\pi \in \Pi$, the expected reward is $r^\pi(s) := \sum_{a \in \mathcal{A}} \pi_s(a) r(s,a)$ and the expected transition is $P^\pi(s,s') := \sum_{a \in \mathcal{A}} \pi_s(a) P(s,a,s')$. The evaluation Bellman operator given by $T^\pi_r v := r^\pi + \gamma P^\pi v$ for all $v \in \mathbb{R}^{\mathcal{S}}$ is known to be contracting, admitting v^π_r as a fixed point. Similarly, the optimal value function $v^\pi_r(s) := \max_{\pi \in \Pi} v^\pi_r(s)$ is the fixed point of the optimal Bellman operator $T^*_r v(s) := \max_{\pi \in \Pi} T^\pi_r v(s), \forall v \in \mathbb{R}^{\mathcal{S}}, s \in \mathcal{S}$. An optimal policy can be derived from iterative Bellman updates, which form the building block of RL (Puterman, 2014).

Robust MDPs. The true reward or transition function of an MDP is rarely known in practice. It may be estimated from trajectory data, but a small error on the MDP model can alter policy performance (Mannor et al., 2004). The robust MDP setting addresses this issue by assuming that (P,r) is unknown, lying in a given uncertainty set. In our applicative setting of DCP where transitions are fully determined by the performed actions, we reasonably assume that the reward model $r \in \mathcal{R}$ is the only uncertain element. Then, we aim to maximize performance for the worst-case model, namely:

$$v_{\mathcal{R}}^{\pi}(s) := \min_{r \in \mathcal{R}} v_r^{\pi}(s), \quad \forall s \in \mathcal{S}.$$

To solve this max-min problem, one can resort to robust Bellman operators $T_{\mathcal{R}}^{\pi}v := \min_{r \in \mathcal{R}} T_r^{\pi}v$ and $T_{\mathcal{R}}^{*}v := \max_{\pi \in \Pi} T_{\mathcal{R}}^{\pi}v$, $\forall v \in \mathbb{R}^{\mathcal{S}}$. Indeed, both are contracting and admit the robust value $v_{\mathcal{R}}^{\pi}$ and the optimal robust value $v_{\mathcal{R}}^{*}$ as fixed points, respectively (Iyengar, 2005; Wiesemann et al., 2013).

Regularized MDPs provide a general framework for regularization in RL, recovering celebrated algorithms such as soft Q-learning (Haarnoja et al., 2017). A regularized MDP is an MDP (S, A, γ, P, r) combined with a family $\Omega := (\Omega_s)_{s \in S}$ of convex functions $\Omega_s : \Delta_{\mathcal{A}} \to \mathbb{R}$. At each state $s \in S$, Ω_s defines a policy regularizer $\Omega_s(\pi_s)$, for $\pi_s \in \Delta_{\mathcal{A}}$. The regularized Bellman operator is given by:

$$[T^{\pi,\Omega}v](s) := T_r^{\pi}v(s) - \Omega_s(\pi_s), \quad \forall v \in \mathbb{R}^{\mathcal{S}}, s \in \mathcal{S}.$$

and its greedy equivalent by $[T^{*,\Omega}v](s) := \max_{\pi_s \in \Delta_A} [T^{\pi,\Omega}v](s)$ (Geist et al., 2019). Derman et al. (2021) have established an equivalence between policy-regularized and robust Bellman operators, thus highlighting a formal motivation for regularized RL. The statement below is a direct reformulation of (Derman et al., 2021, Thm. 3.1).

Theorem 3.1 (Derman et al. (2021)). Assume that the reward function r is uncertain and satisfies $r_s \in \mathcal{R}_s := r_0(s,\cdot) + \tilde{\mathcal{R}}_s$, where $\tilde{\mathcal{R}}_s \subseteq [-R,R]^{\mathcal{A}}$ is closed and convex for all $s \in \mathcal{S}$. Then, for any $\pi \in \Pi$, the robust value function $v_{\mathcal{R}}^{\pi}$ is the fixed point of the regularized Bellman operator $T^{\pi,\Omega}$ with $\Omega_s(\pi_s) := \max_{r_s \in \tilde{\mathcal{R}}_s} \langle -\pi_s, r_s \rangle$. In other words, it holds that $v_{\mathcal{R}}^{\pi}(s) = T_{r_0}^{\pi} v_{\mathcal{R}}^{\pi}(s) - \Omega_s(\pi_s), \ \forall s \in \mathcal{S}$.

Discrete compositional processes (DCP) are special instances of standard MDPs with deterministic transitions. They can fully be described by a tuple (\mathcal{G}, Φ) , where \mathcal{G} is a directed acyclic graph with a single source state s_0 and a set of terminal states \mathcal{X} (see Fig. 1). Each node in the graph corresponds to a set of parts, s_0 being the empty set. An edge $(s \to s')$ represents either adding a part to the set s to get s' or a termination action if $s' \in \mathcal{X}$. Each completed object $s' \in \mathcal{X}$ has an associated reward $s' \in \mathcal{X}$ given by the proxy reward model $s' \in \mathcal{X}$ has an associated reward $s' \in \mathcal{X}$

In standard RL, the goal is to maximize the discounted cumulative reward. Instead, for scientific discovery, we are interested in finding the most promising set of k candidates that are sufficiently distinct from one another (Jain et al., 2022). In practice, an imperfect but computationally efficient proxy score Φ is used to estimate the usefulness of each candidate, so naively maximizing usefulness can be problematic. Instead, given some metric d between completed objects, practitioners aim to approximately find a diverse and novel (w.r.t. the training set) set of candidates $\{x_1, \ldots x_k\}$ with a high score (Jain et al., 2022). This goal can be formalized as maximizing the average mode reward (for simplicity, we only formalize the diversity constraint):

$$\max_{\{x_1,\dots x_k\}\subset\mathcal{X}} \frac{1}{k} \sum_{i=1}^k e^{\Phi(x_i)} \quad \text{subject to} \quad d(x_i, x_j) > \delta, \quad \forall i \neq j.$$
 (1)

3.1 Limitations of GFLowNets in scientific discovery

Under a GFN perspective, the exponential of the proxy reward function can be viewed as an unnormalized probability mass. Then, the goal is to learn a per-state policy π_s whose sequential

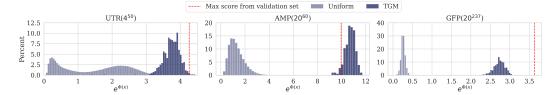


Figure 2: Distribution of exponential rewards from 1 million uniformly randomly drawn samples for various tasks compared to the distribution sampled by TGM at the end of training.

application samples objects proportionally to this function. Formally, if we denote by $\mathcal{T}(x)$ all trajectories in \mathcal{G} ending at terminal state $x \in \mathcal{X}$ and $\pi_{\tau(x)}$ the probability of picking all edges from a trajectory $\tau(x)$ under policy π , the objective is to match the following distribution:

$$p(x) := \sum_{\tau(x) \in \mathcal{T}(x)} \pi_{\tau(x)} \propto e^{\beta \Phi(x)}, \qquad \forall x \in \mathcal{X}.$$
 (2)

where β is the inverse temperature. We argue this objective is suboptimal for solving Eq. 1.

Motivating example. The main issue with Eq. 2 is that it yields a distribution assigning higher probability to a large number of suboptimal reward objects rather than to a small number of high-reward objects. To demonstrate this, consider a DCP (illustrated in Fig. 1) where the goal is to generate sequences of maximum length d from a vocabulary of size B. Suppose a sampler is at an optimal sequence s^* of length n < d. The sampler's policy then needs to decide whether to terminate and return x^* , or to add more tokens and harm the usefulness of the sequence. In doing so, it weighs $e^{\Phi(x^*)}$ against the reward of all sequences starting with s. There are $\Omega(B^{d-n})$ such sequences. Suppose the sampler perfectly matches (2) and these sequences have rewards lower bounded by r > 0. Then, $\pi(x^* \mid s) \leq e^{\Phi(x^*)}/r \cdot B^{d-n}$, and the probability of returning the optimal sequence at s decreases exponentially in d-n. The problem occurs in many applications, as almost all objects have non-negligible reward (see Fig. 2). When these empirical distributions are representative, GFNs may conservatively estimate the probability of sampling a promising candidate.

4 ALTERNATIVE OPERATORS

To address this limitation, we propose taking a regularized RL perspective and modifying the operator used by GFNs. To simplify the discussion, we examine the case where $\mathcal G$ in the DCP is a tree. In such settings, there is a direct equivalence between the GFN flow operator and the soft Bellman operator with temperature $\omega=1$. Denoting $Q_s[a]:=v_{s_a}+r(s,a)$ (with s_a being the state reached by taking action a from s) and the vector of Q-values at a state by $\mathbf{Q_s}$, the soft Bellman operator is given by:

$$T^{\mathrm{SB}}v_s := \frac{1}{\omega} \mathrm{logsumexp}(\omega \mathbf{Q_s}).$$
 (3)

Recursively applying this operator in a DCP, as you go from the leaves of \mathcal{G} to the root, causes value to accumulate. At optimality, the value of a state is the logsumexp of the rewards of all objects that can be reached from that state. The accumulation of the value of a large number of suboptimal objects can overwhelm one high-reward object.

Instead of taking the sum, we could take the *logmeanexp* of the Q-values. Doing so yields the mellowmax operator (Asadi & Littman, 2017):

$$T^{\text{MM}}v_s := \frac{1}{\omega} \log \sum_{a \in \text{Ch}(s)} \frac{1}{|\text{Ch}(s)|} e^{\omega Q_s(a)}.$$
 (4)

While the mellowmax operator solves the accumulation issue, it is liable to the opposite issue, dilution. Due to this averaging, the high reward of an object deep in the tree can get diluted by lower reward objects as we approach the root (since we divide by the number of children at each state).

Luckily, we can alleviate the dilution issue by weighting the elements of the sum proportionally to their value. In particular, we can use the softmax of the Q-values as weights (instead of an equal weighting) and thus obtain the soft mellowmax operator (Gan et al., 2021):

$$T^{\text{SMM}}v_s = \frac{1}{\omega}\log\langle\operatorname{softmax}(\alpha \mathbf{Q_s}), e^{\omega \mathbf{Q_s}}\rangle,$$
 (5)

where α controls the temperature of the softmax. Ultimately, we want to be able to trade off accumulation and dilution while maintaining a soft operator.

4.1 GENERAL MELLOWMAX

The operators mentioned above can be viewed through the lens of regularized MDPs. On one hand, the soft Bellman operator corresponds to regularizing by the Shannon entropy. On the other hand given a distribution $d_s \in \Delta_{\mathcal{A}}$, the mellowmax/soft mellowmax operators result from regularization by $\mathrm{KL}(\pi_s, d_s)$ where d_s is the uniform distribution or $\mathrm{softmax}(\alpha \mathbf{Q_s})$ respectively.

We define the **general mellowmax** (GM) regularizer by interpolating between these regularizers. Then, we can recover these operators as special cases and also trade off their respective effects. We denote the general regularizer by $\Omega_{d_a}^{\mathbf{q}}$ and the case $d_s = \operatorname{softmax}(\alpha \mathbf{Q_s})$ by $\Omega_{\alpha}^{\mathbf{q}}$.

$$\Omega_{d_s}^{\mathbf{q}}(\pi_s) := \frac{1}{\omega} \left[\mathsf{qKL}(\pi_s, d_s) + (1 - \mathsf{q})(-\mathsf{H}(\pi_s)) \right], \quad \forall s \in \mathcal{S}, \pi_s \in \Delta_{\mathcal{A}}, \omega \in \mathbb{R}_{>0}, \mathsf{q} \in [0, 1]. \tag{6}$$

Notably, unlike other regularizers, Ω_{α}^{q} depends on the Q-values of the current state. Consequently, the associated operator is not a convex conjugate and deviates from the framework of (Geist et al., 2019). Nonetheless, we can still derive the associated operator using similar techniques (see App. D):

$$\Omega_{\alpha}^{\mathbf{q},*}(\mathbf{Q_s}) = \frac{1}{\omega} \log \langle \operatorname{softmax}(\alpha \mathbf{Q_s}))^{\mathbf{q}}, e^{\omega Q_s} \rangle.$$
 (7)

Using the regularized interpretation, q trades off between maximizing the policy's entropy and having a policy that is close to the softmax of the Q-values. α allows us to put more or less weight on actions with high Q-values while ω controls the weight of regularization. In particular, setting $\mathbf{q}=0$ allows us to recover the standard entropy-regularized / GFN operators and policy. On the other hand, by setting $\mathbf{q}=1$, we recover the soft mellowmax operator (with $\alpha=0$ we recover the mellowmax operator). The interpolated operator captures three important soft RL operators and allows us to smoothly control the tension between accumulation/dilution using \mathbf{q} and α .

We also observe that $\Omega_{d_s}^{\mathbf{q}}$ is equivalent to a KL between a policy and a tilted softmax, as stated below.

Proposition 4.1. For all
$$q \in [0,1]$$
, $\Omega_{d_s}^q(\pi_s) = \frac{1}{\omega} \text{KL}(\pi_s, d_s^{(q)}) - \frac{1}{\omega} \log(Z_q(d_s))$, where $Z_q(d_s) := \sum_{a \in A} d_s(a)^q$ and $d_s^{(q)} = d_s^q/Z_q(d_s)$ is the q-tilted softmax distribution.

4.2 Trajectory general mellowmax

Instead of training on transitions, it is possible to leverage trajectory-level constraints that relate policy/value functions over multiple steps to recover training algorithms that use subtrajectories as data units. In soft RL, this connection was first formalized in Nachum et al. (2017) through path consistency learning (PCL) for maximum entropy RL. Path consistency objectives have since been derived for Tsallis entropy (Chow et al., 2018) as well as the general class of α -divergences (Brekelmans et al., 2022).

Trajectory-level constraints have been shown to be crucial to the performance of methods in DCP tasks. Trajectory balance (TB) (Malkin et al., 2022) and subtrajectory balance (Madan et al., 2023) are consistently used in GFNs, as DCP tasks only have terminal rewards. By connecting the policy at early states to this terminal reward, these constraints are noticeably better at propagating signal and thus improving credit assignment.

Theorem 4.1 (Trajectory GM). For any DCP (\mathcal{G}, Φ) , let (π^*, Q^*) be the unique optimal value/policy functions for $\Omega^{q,*}_{\alpha}$. Then, for a given Q-function Q^{θ} , the two following statements are equivalent

$$\operatorname{softmax}([\alpha \mathsf{q} + \omega]Q_s^{\theta}) = \arg\max_{\pi_s} \langle \pi_s, Q_s^* \rangle - \frac{1}{\omega} \Omega_{\alpha}^{\mathsf{q}}(\pi_s), \text{ for all states } s \in \mathcal{G}. \tag{8}$$

$$v_0^* + \sum_{i=0}^n \frac{1}{\omega} \left(\log \operatorname{softmax}([\alpha \mathsf{q} + \omega] Q_{s_i}^{\theta})[a_i] - \mathsf{q} \log \operatorname{softmax}(\alpha Q_{s_i}^{\theta})[a_i] \right) = r(x),$$
for all trajectories $s_0 \xrightarrow{a_0} \cdots \xrightarrow{a_{n-1}} s_n \xrightarrow{a_n} x$ in the DCP.

We derive a novel, equivalent trajectory constraint for the general mellowmax operator. In particular, Thm 4.1 explicits that satisfying the trajectory constraint Eq. 9 on all trajectories yields Q_{θ} such that softmax($[\alpha \mathbf{q} + \omega]Q_{\theta}$) is the optimal policy for the regularized problem. See App. D.2 for the proof.

Trajectory general mellowmax (TGM). We now instantiate the interpolated regularizer as a practical algorithm for use in DCPs. Specifically, we make three design decisions.

- (1) Using the general mellowmax operator with $d_s = \operatorname{softmax}(\alpha Q_s)$. Doing so allows us to use TGM both on-policy and off-policy and does not require a separate d_s . Instead, the policy seeks to maximize reward while being close to a softmax of the learned Q-values.
- (2) Aiming to satisfy a novel trajectory constraint derived in App. D.2. Doing so recovers the benefits of trajectory constraints for DCPs.
- (3) Training a single network Q^{θ} of Q-values through the VarGrad objective of Richter et al. (2020); Zhang et al. (2023). It can be shown that minimizing the trajectory constraint Eq. 9 is equivalent to minimizing the following loss (where σ_t denotes the softmax with inverse temperature t):

$$\mathcal{L}_{\text{TGM},Q_{\theta}}(\tau) = \operatorname{Var}\left[\frac{1}{\omega}\left(\sum_{i=1}^{n}\log\sigma_{\mathbf{q}\alpha+\omega}[Q_{s_{i}}^{\theta}(a_{i})] - \mathsf{q}\log\sigma_{\alpha}[Q_{s_{i}}^{\theta}(a_{i})]\right) - \beta r(x)\right].$$

From Q^{θ} , the optimal policy is easily computed. The resulting algorithm is easy to implement, efficient (particularly for transformers), and is equivalent to GFN training when q=0, $\omega=1$.

5 DISCRETE COMPOSITIONAL GENERATION VIA ROBUST RL

In this section, we offer a robust RL interpretation of regularized operators for DCPs. As described in Thm. 3.1, regularized MDPs are known to be equivalent to robust MDPs with uncertain reward. In the context of scientific discovery, since the reward is the evaluation of objects by a proxy function, we can conceptually consider the existence of a *hidden true reward* r^* that we would like to maximize¹. The filtering process can be interpreted as an attempt to robustly maximize the proxy reward Φ , accounting for the difference δ between Φ and r^* . Thus, the goal is to solve:

$$\max_{p \in \Delta_{\mathcal{X}}} \min_{\delta \in \mathcal{R}} \mathbb{E}_{x \sim p}[\Phi(x) + \delta(x)]$$
 (10)

where $\mathcal{R} \in \mathbb{R}^{|\mathcal{X}|}$ is the uncertainty set. The question then remains of how this uncertainty should be modeled in \mathcal{G} . Traditionally, the reward in a DCP is identically 0 everywhere except for terminating actions. However, assuming that uncertainty only exists at the last action misses the compositional nature of the task. Instead, we decompose $\delta(x)$ into a sum of perturbations occurring at each step of the generation process. More precisely, given a trajectory $s_0 \xrightarrow{a_0} \cdots \xrightarrow{a_{n-1}} s_n = x$, we have

$$\delta(x) = \sum_{i=0}^{n} \delta_i[a_i], \text{ where } \delta_i \in \mathcal{R}_{s_i}.$$
(11)

In this formulation, the uncertainty on Φ is *split between all actions taken to construct the object*, instead of only existing at the final action, as illustrated in Fig. 1. From a robust RL perspective, this uncertainty set structure corresponds to the common assumption of state-rectangularity (Wiesemann et al., 2013; Gadot et al., 2024). The following section makes this model of uncertainty more explicit and analyzes the uncertainty sets of different operators. In particular, we show that the reward uncertainty set entailed by GFNs is inadequate, whereas the set induced by the soft mellowmax operator provides a more meaningful notion of uncertainty.

5.1 FENCHEL-ROBUST FORMULATION OF REGULARIZED MDPs

In this subsection, we use the robust MDP notations introduced in § 3. The following result provides an explicit mapping between reward-robust MDPs, as defined in (11), and regularized MDPs, through equivalent value functions.

Theorem 5.1 (Fenchel-Robust MDP). For any state $s \in S$, let $\Omega_s : \Delta_A \to \overline{\mathbb{R}}$ be a proper convex regularization function. Denote by Ω_s^* its convex conjugate which is known to be proper and convex. For any $s \in S$, define the reward set

$$\mathcal{R}_s := r_0(s, \cdot) + \left\{ r_s \in \mathbb{R}^{\mathcal{A}} : \Omega_s^*(-r_s) \le 0 \right\}.$$

Then, denoting the robust operator for \mathcal{R}_s by $T_{\mathcal{R}_s}^{\pi}$ and the regularized operator for Ω_s by T^{π,Ω_s} :

$$T_{\mathcal{R}_s}^{\pi} v = T^{\pi,\Omega_s} v, \qquad \forall \pi \in \Pi, v \in \mathbb{R}^{\mathcal{S}}.$$

¹This reward would correspond to the effective properties of interest measured in a lab.

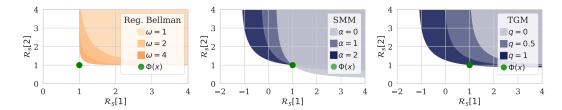


Figure 3: (Left) Regardless of ω , the uncertainty set *never contains* $\Phi(x)$. As a result, the soft Bellman/GFN operator is only robust to increases in reward. (Middle) For the soft mellowmax operator, for different values of α with $d_s[1] > d_s[2]$, the uncertainty set contains $\Phi(x)$. Thus, the operator is robust to decreases in reward of one object (but not both at the same time). When $\alpha=0$ (mellowmax), there is a symmetry in this tradeoff, while increasing α skews it such that the object with higher d_s only admits a small decrease in reward. (Right) The uncertainty sets for GM interpolate between the two effects. While the uncertainty set for 0 < q < 1 does not contain $\Phi(x)$, it does contain points corresponding to decreases in reward.

Based on this theorem, solving Eq. 10 is equivalent to solving the associated regularized MDP. Thm. 5.1 uses convex conjugacy in the context of dynamic programming. Similar results can be found in previous works (Eysenbach & Levine, 2021; Husain et al., 2021; Brekelmans et al., 2022; Derman et al., 2021), but as we carefully detail in Appx. A.5, they do not directly apply to our setting.

5.2 Robust sets induced by common regularizers

Given this perspective, we now analyze the uncertainty sets corresponding to the regularizers we considered above. Since each is a special case of GM, we first give the uncertainty set of GM:

$$\mathcal{R}_s := r_0(s, \cdot) + \left\{ r_s \in \mathbb{R}^{\mathcal{A}} : \frac{1}{\omega} \sum_{a \in \mathcal{A}} d_s(a)^{\mathsf{q}} e^{-\omega r_s(a)} \le 1 \right\}. \tag{12}$$

A proof is in App. A.4. As opposed to (Derman et al., 2021), the fact that our uncertainty sets are state-rectangular makes them independent of the executed policy, which results in a formulation that fits with the standard robust MDP setting (Wiesemann et al., 2013).

Fig. 3 shows the uncertainty sets induced by three regularizers: negative Shannon entropy, soft mellowmax, and GM. For this illustration, we consider a single-step generation process with 2 possible objects and associated proxy rewards [1, 1]. Similarly to (Brekelmans et al., 2022), we find that the entropy regularized Bellman operator (left) fails to capture a meaningful notion of robustness. Indeed, the uncertainty set never contains the proxy $\Phi(x)$, although higher values of ω do bring the set closer to the proxy reward. This problem is significantly exacerbated in a DCP with multiple steps: at each layer, the uncertainty set becomes further away from the proxy reward (see App. C).

6 EXPERIMENTS

Our experiments aim to answer three main questions:

- [Q1] How does the parameter q affect the peakiness of sampling in small and synthetic environments?
- [Q2] Does TGM find better candidates than standard methods in large biological design DCP tasks?
- **[Q3]** How robust is TGM to changes in hyperparameter settings?

6.1 IMPACT OF q IN SMALL AND SYNTHETIC ENVIRONMENTS [Q1]

TF-Bind-8. Originally proposed in (Trabucco et al., 2022), TF-Bind-8 generates DNA sequences of length 8 to find those with high binding activity to human transcription factors. The reward is the experimental binding activity of the sequence from (Barrera et al., 2016). As there are only 4^8 such sequences, the search space is small enough to compute the optimal value for each operator. We compare the optimal sampling densities of each algorithm in Fig. 4 for $\beta = 4$, $\alpha = 2$ and $\omega = 2$.

Bit sequence. Proposed in (Malkin et al., 2022), the bit sequence task is significantly larger with 2^{120} possible sequences. The goal is to generate bit sequences of length n by adding k bits at a time. M modes are selected semi-randomly and the reward is given by $r(x) = 1 - \min_{y \in M} d(x, y)/n$ where d is the Levenshtein edit distance (Levenshtein et al., 1966). Loosely, the reward of a sequence is the negative normalized edit distance to the nearest mode.

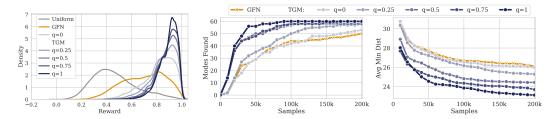


Figure 4: (Left) Comparison of the optimal sampling distribution of GFN and variants of TGM for TF-Bind-8 rewards. For the same $\beta=4$, TGM concentrates significantly more mass on the upper quantiles of the reward distribution. (Middle) Number of modes found by TGM and GFN. The increased peakiness of the TGM sampling does not harm its ability to find different modes. (Right) Average (over modes) of the distance of the closest sample found for each mode. On average, increasing q allows TGM to find closer samples to the true modes.

				TGM (ours)				Valid.	
	SAC	PPO	GFN	$\mathbf{q} = 0$	$\mathbf{q} = 0.25$	$\mathbf{q}=0.5$	$\mathbf{q}=0.75$	$\mathbf{q} = 1.0$	Max
UTR	$3.82 \pm .03$	$3.46 \pm .02$	$4.12 \pm .00$	$4.19 \pm .01$	$4.15 \pm .01$	$4.25 \pm .00$	$4.27{\pm}.01$	$4.17 \pm .01$	4.26
AMP	$5.69 \pm .35$	$9.79 {\pm} .02$	$10.04 \pm .04$	$10.13 \pm .04$	$9.93 \pm .03$	$9.82 {\pm} .02$	$10.43 {\pm}.07$	$10.29 \pm .05$	9.96
GFP	_	$0.66{\pm}.02$	$1.90 \pm .03$	$2.44 \pm .09$	$2.95 {\pm}.08$	$2.66 \pm .12$	$2.34 \pm .11$	$2.05 \pm .11$	3.63

Table 1: Average mode reward of TGM compared to baselines on the biological sequence design tasks. We report the average and standard error over 15 seeds. In all domains, variants of TGM either match or outperform GFN/PPO/SAC with a particularly pronounced difference in GFP, the largest domain. Valid. max refers to the object with the highest proxy reward in the validation set of the dataset used to train the proxy function.

As we have access to the ground-truth modes in this case, we evaluate methods according to the following metrics: (1) the number of modes found, where a mode is deemed found if a sample is generated within distance $\delta=28$ of that mode; (2) for each mode, we track the distance to the closest sample generated during training. We report the average of this number over the M modes for the best performing hyperparameters (see App. E.3 for further details). While the best GFN run matches q=0, the increased peakiness of non-zero q-values yields noticeable benefits. The best-performing runs of q>0 find samples that are much closer to the modes while also discovering more of them.

6.2 BIOLOGICAL SEQUENCE DESIGN [Q2]

We now focus on tasks based on real-world domains, where proxy rewards are trained on actual datasets. The tasks are based on the setup of Malkin et al. (2022) and use the datasets from Trabucco et al. (2022). For each, a transformer Φ is first trained on the training set (either for classification or regression). The transformer output is normalized before being used as a proxy reward (see App. E.4).

5' Untranslated Region Sequence (UTR). 5' UTR is an mRNA region that regulates transcription of the main coding sequence. The goal of the task is to find a UTR sequence of length 50 that maximizes predicted gene expression level. We take the dataset of 280 000 filtered sequences from Trabucco et al. (2022) and associated ribosome loads. We train a transformer as a regressor to predict the ribosome load from a sequence. The vocabulary consists of 4 nucleotides.

Antimicrobial Peptide (AMP). Antimicrobial peptides are short sequences of amino acids that have effects on microbes (bacteria, viruses, etc.). The goal is to discover novel peptides that are likely to have antimicrobial properties. The DBAASP database has collected known peptides with and without antimicrobial activity (Pirtskhalava et al., 2021). We use the dataset (sourced from Trabucco et al. (2022)) of 9222 non-AMP sequences and 6438 AMP sequences and train a binary classifier (predicting whether a sequence has antimicrobial properties or not). We use the normalized logit of the classifier as proxy reward. The vocabulary for this task consists of the 20 amino acids, plus a token corresponding to sequence termination. We mask the tokens such that the minimum sequence length is 14 and the maximum sequence length is 60. This is the only task where we allow variable-length sequences (as it is the only one whose dataset contains sequences of variable length).

Green Fluorescent Protein (GFP). The green fluorescent protein is a length-237 protein whose fluorescence has numerous applications in biology. The goal of this task is to discover other proteins with high predicted fluorescence. We take the processed dataset from Trabucco et al. (2022), which

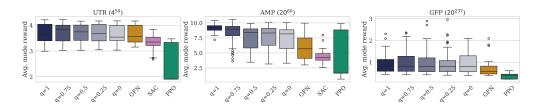


Figure 5: Spread of final average mode rewards for various algorithms from a grid sweep over learning rates, β and ω . TGM on average performs better in AMP and GFP and similarly in UTR.

contains 56086 variations of the original sequence and their measured fluorescence. We train a regressor to predict the measured fluorescence, and the vocabulary consists of the 20 amino acids.

Evaluation. Each method is trained for $100\,000$ generated samples during which we regularly evaluate the learned network. By varying the temperature coefficient of the policy, we can move along a quality/diversity curve for each sampler. To approximately evaluate Eq. 1, we generate a set of samples for a range of temperature coefficients. We then aggregate the samples and aim to approximately determine the best k distinct samples (i.e., such that the minimum distance is δ) found by the policy. While this problem is NP-complete Karp (2009) (it can be seen as an instance of finding a maximal independent set), a greedy approach appears to work well in practice. For each method, a sweep is performed over hyperparameters. The best performing setting (based on the final average mode reward) is then run with 15 different seeds. See App. E.4 for further details.

Results. As illustrated in Tab. 1, for all three tasks, all variants of TGM either roughly match or exceed the performance of GFN, soft actor-critic (SAC) (Haarnoja et al., 2018), and proximal policy optimization (PPO) (Schulman et al., 2017). Similarly to what was found in (Malkin et al., 2022; Madan et al., 2023), SAC and PPO perform relatively poorly. We hypothesize that the credit assignment problem from terminal rewards is a significant issue for these methods. In AMP, where generating shorter sequences is possible, PPO is able to achieve similar results to TGM/GFN. The difference between TGM and GFN is most pronounced in the largest environment GFP, where TGM finds modes with significantly higher reward than the mean in Fig. 2. We hope this result shows the potential scalability of TGM. Interestingly, the effect of q is more varied in these environments, showing the benefits of interpolation in the GM operator. The best performing variants balance dilution and accumulation by using $q \in (0, 1)$.

6.3 Hyperparameter robustness [Q3]

The best runs for Tab. 1 were selected from grid sweep over learning rates $\{0.00001, 0.0001, 0.001\}$, $\beta \in \{4, 16, 64, 256\}$ and inverse temperature parameters $\omega \in \{1, 4, 16\}$. For SAC and PPO, we set the entropy coefficient to $1/\omega$. For GFN, we set the sampling temperature of the policy to $1/\omega$. To ensure an equal number of runs per method, we fix $\alpha = 1$ and only vary ω for TGM. Varying α as well would likely be beneficial. We plot the performance distribution of this sweep in Fig. 5.

TGM variants seem relatively robust to these hyperparameter variations. The average SAC run is relatively stable but performs poorly. PPO is very sensitive to hyperparameter settings, with the best runs achieving strong performance in AMP. The mean average mode reward of TGM is higher than GFN across q and environments. Remarkably, TGM with q=1 in AMP has little variability, with almost all runs performing well despite the significantly different hyperparameters. For GFP, there is significantly more variability, with the best performances coming from outliers for each method.

7 CONCLUSION

In this paper, motivated by the inadequacy of sampling proportional to reward, we generalize various soft RL operators and propose GM as an interpolated regularizer. From this regularizer, we develop a novel algorithm (TGM) and show it is consistently able to outperform GFNs on a variety of DCP tasks. Finally, we adapt the reward-robust RL framework to DCP tasks where rewards are given for entire trajectories to offer a new perspective on regularization in these problems. Ultimately, we believe TGM has the potential to be a more general and effective framework for finding promising candidates in scientific discovery applications.

8 ETHICS STATEMENT

The purpose of this work is algorithmic in nature, and we do not aim to produce actual applicable models. Nonetheless, the ideas in this work could potentially be used to improve the training of models used in protein or other biological sequence design applications. While there are numerous beneficial applications for these types of models, they could also be used to design sequences that could eventually be harmful to humans. This harm could come from bad actors or even accidentally by over-optimizing to a single metric (e.g., antimicrobial resistance) and not considering potential harmful effects on the body. The extension of this work to multi-objective settings is a fruitful direction for future work.

In addition, the contribution of this work also has applications in language (another sequence design task). As such, it has the potential to improve the performance of large language models (LLMs) on both useful and nefarious tasks. Given the recent improvements in the abilities of these models, this potential effect is especially relevant.

9 REPRODUCIBILITY STATEMENT

To promote transparency and reproducibility, we share our code implementation as well as checkpoints of the proxy reward functions we use. These allow for the reproduction of our results and fair future comparisons with new methods. We describe in ample detail the hyperparameters we use and design decisions we made in App. E as well as Sec. 6. See the supplementary material for the code/checkpoints.

REFERENCES

- Mila AI4Science, Alex Hernandez-Garcia, Alexandre Duval, Alexandra Volokhova, Yoshua Bengio, Divya Sharma, Pierre Luc Carrier, Yasmine Benabed, Michał Koziarski, and Victor Schmidt. Crystal-GFN: sampling crystals with desirable properties and constraints. *arXiv preprint* 2310.04925, 2023.
- Christof Angermueller, David Dohan, David Belanger, Ramya Deshpande, Kevin Murphy, and Lucy Colwell. Model-based reinforcement learning for biological sequence design. In *International Conference on Learning Representations (ICLR)*, 2019.
- Kavosh Asadi and Michael L Littman. An alternative softmax operator for reinforcement learning. In *International Conference on Machine Learning (ICML)*, 2017.
- Luis A Barrera, Anastasia Vedenko, Jesse V Kurland, Julia M Rogers, Stephen S Gisselbrecht, Elizabeth J Rossin, Jaie Woodard, Luca Mariani, Kian Hong Kock, Sachi Inukai, et al. Survey of variation in human transcription factors reveals prevalent DNA binding changes. *Science*, 351 (6280):1450–1454, 2016.
- Emmanuel Bengio, Moksh Jain, Maksym Korablyov, Doina Precup, and Yoshua Bengio. Flow network based generative models for non-iterative diverse candidate generation. *Neural Information Processing Systems (NeurIPS)*, 2021.
- Dimitri Bertsekas. Convex optimization theory, volume 1. Athena Scientific, 2009.
- Albert Bou, Morgan Thomas, Sebastian Dittert, Carles Navarro, Maciej Majewski, Ye Wang, Shivam Patel, Gary Tresadern, Mazen Ahmad, Vincent Moens, et al. ACEGEN: Reinforcement learning of generative chemical agents for drug discovery. *Journal of Chemical Information and Modeling*, 64 (15):5900–5911, 2024.
- James Bradbury, Roy Frostig, Peter Hawkins, Matthew James Johnson, Chris Leary, Dougal Maclaurin, George Necula, Adam Paszke, Jake VanderPlas, Skye Wanderman-Milne, and Qiao Zhang. JAX: composable transformations of Python+NumPy programs, 2018. URL http://github.com/jax-ml/jax.

- Rob Brekelmans, Tim Genewein, Jordi Grau-Moya, Grégoire Delétang, Markus Kunesch, Shane Legg, and Pedro Ortega. Your policy regularizer is secretly an adversary. *arXiv preprint 2203.12592*, 2022.
- Yinlam Chow, Ofir Nachum, and Mohammad Ghavamzadeh. Path consistency learning in Tsallis entropy regularized MDPs. In *International Conference on Machine Learning (ICML)*, 2018.
- Petros Christodoulou. Soft actor-critic for discrete action settings. arXiv 2019. arXiv preprint 1910.07207, 2019.
- Flaviu Cipcigan, Jonathan Booth, Rodrigo Neumann Barros Ferreira, Carine Ribeiro Dos Santos, and Mathias Steiner. Discovery of novel reticular materials for carbon dioxide capture using GFlowNets. *Digital Discovery*, 3(3):449–455, 2024.
- Andrew Currin, Neil Swainston, Philip J Day, and Douglas B Kell. Synthetic biology for the directed evolution of protein biocatalysts: navigating sequence space intelligently. *Chemical Society Reviews*, 44(5):1172–1239, 2015.
- Victor-Alexandru Darvariu, Stephen Hailes, and Mirco Musolesi. Goal-directed graph construction using reinforcement learning. *Proceedings of the Royal Society A*, 477(2254):20210168, 2021.
- Sarah E Davis, Selen Cremaschi, and Mario R Eden. Efficient surrogate model development: impact of sample size and underlying model dimensions. In *Computer aided chemical engineering*, volume 44, pp. 979–984. Elsevier, 2018.
- Nicola De Cao and Thomas Kipf. MolGAN: An implicit generative model for small molecular graphs. *arXiv preprint 1805.11973*, 2018.
- Tristan Deleu, Padideh Nouri, Nikolay Malkin, Doina Precup, and Yoshua Bengio. Discrete probabilistic inference as control in multi-path environments. *Conference on Uncertainty in Artificial Intelligence (UAI)*, 2024.
- Esther Derman, Matthieu Geist, and Shie Mannor. Twice regularized MDPs and the equivalence between robustness and regularization. *Neural Information Processing Systems (NeurIPS)*, 2021.
- Yuanqi Du, Tianfan Fu, Jimeng Sun, and Shengchao Liu. Molgensurvey: A systematic survey in machine learning models for molecule design. *arXiv* preprint 2203.14500, 2022.
- Benjamin Eysenbach and Sergey Levine. Maximum entropy RL (provably) solves some robust RL problems. *arXiv* preprint 2103.06257, 2021.
- Uri Gadot, Esther Derman, Navdeep Kumar, Maxence Mohamed Elfatihi, Kfir Levy, and Shie Mannor. Solving non-rectangular reward-robust MDPs via frequency regularization. In *AAAI Conference on Artificial Intelligence*, 2024.
- Yaozhong Gan, Zhe Zhang, and Xiaoyang Tan. Stabilizing q learning via soft mellowmax operator. In *AAAI Conference on Artificial Intelligence*, 2021.
- Matthieu Geist, Bruno Scherrer, and Olivier Pietquin. A theory of regularized markov decision processes. In *International Conference on Machine Learning (ICML)*, 2019.
- Walter R Gilks, Sylvia Richardson, and David Spiegelhalter. *Markov chain Monte Carlo in practice*. CRC press, 1995.
- JE Gubernatis and TJPRM Lookman. Machine learning in materials design and discovery: Examples from the present and suggestions for the future. *Physical Review Materials*, 2(12):120301, 2018.
- Tuomas Haarnoja, Haoran Tang, Pieter Abbeel, and Sergey Levine. Reinforcement learning with deep energy-based policies. In *International Conference on Machine Learning (ICML)*, 2017.
- Tuomas Haarnoja, Aurick Zhou, Pieter Abbeel, and Sergey Levine. Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. In *International Conference on Machine Learning (ICML)*, 2018.

- Shengyi Huang, Rousslan Fernand Julien Dossa, Chang Ye, Jeff Braga, Dipam Chakraborty, Kinal Mehta, and João G.M. Araújo. CleanRL: High-quality single-file implementations of deep reinforcement learning algorithms. *Journal of Machine Learning Research*, 23(274):1–18, 2022.
 URL http://jmlr.org/papers/v23/21-1342.html.
 - Hisham Husain, Kamil Ciosek, and Ryota Tomioka. Regularized policies are reward robust. In *International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2021.
 - Garud N Iyengar. Robust dynamic programming. *Mathematics of Operations Research*, 30(2): 257–280, 2005.
 - Moksh Jain, Emmanuel Bengio, Alex Hernandez-Garcia, Jarrid Rector-Brooks, Bonaventure FP Dossou, Chanakya Ajit Ekbote, Jie Fu, Tianyu Zhang, Michael Kilgour, Dinghuai Zhang, et al. Biological sequence design with GFlowNets. In *International Conference on Machine Learning (ICML)*, 2022.
 - Moksh Jain, Tristan Deleu, Jason Hartford, Cheng-Hao Liu, Alex Hernandez-Garcia, and Yoshua Bengio. GFlowNets for AI-driven scientific discovery. *Digital Discovery*, 2(3):557–577, 2023.
 - Marco Jiralerspong, Bilun Sun, Danilo Vucetic, Tianyu Zhang, Yoshua Bengio, Gauthier Gidel, and Nikolay Malkin. Expected flow networks in stochastic environments and two-player zero-sum games. In *International Conference on Learning Representations (ICLR)*, 2024.
 - Richard M Karp. Reducibility among combinatorial problems. In 50 Years of Integer Programming 1958-2008: from the Early Years to the State-of-the-Art, pp. 219–241. Springer, 2009.
 - Minsu Kim, Joohwan Ko, Taeyoung Yun, Dinghuai Zhang, Ling Pan, Woochang Kim, Jinkyoo Park, Emmanuel Bengio, and Yoshua Bengio. Learning to scale logits for temperature-conditional GFlowNets. *arXiv* preprint 2310.02823, 2023.
 - Diederik P Kingma. Adam: A method for stochastic optimization. arXiv preprint 1412.6980, 2014.
 - Vladimir I Levenshtein et al. Binary codes capable of correcting deletions, insertions, and reversals. In *Soviet physics doklady*, volume 10, pp. 707–710. Soviet Union, 1966.
 - Kanika Madan, Jarrid Rector-Brooks, Maksym Korablyov, Emmanuel Bengio, Moksh Jain, Andrei Cristian Nica, Tom Bosc, Yoshua Bengio, and Nikolay Malkin. Learning GFlowNets from partial episodes for improved convergence and stability. In *International Conference on Machine Learning (ICML)*, 2023.
 - Nikolay Malkin, Moksh Jain, Emmanuel Bengio, Chen Sun, and Yoshua Bengio. Trajectory balance: Improved credit assignment in GFlowNets. *Neural Information Processing Systems (NeurIPS)*, 2022.
 - Shie Mannor, Duncan Simester, Peng Sun, and John N Tsitsiklis. Bias and variance in value function estimation. In *International Conference on Machine Learning (ICML)*, 2004.
 - Sobhan Mohammadpour, Emmanuel Bengio, Emma Frejinger, and Pierre-Luc Bacon. Maximum entropy GFlowNets with soft q-learning. In *International Conference on Artificial Intelligence and Statistics*, 2024a.
 - Sobhan Mohammadpour, Emma Frejinger, and Pierre-Luc Bacon. Decoupling regularization from the action space. In *International Conference on Learning Representations (ICLR)*, 2024b.
 - Ofir Nachum, Mohammad Norouzi, Kelvin Xu, and Dale Schuurmans. Bridging the gap between value and policy based reinforcement learning. *Neural Information Processing Systems (NeurIPS)*, 2017.
 - Malak Pirtskhalava, Anthony A Amstrong, Maia Grigolava, Mindia Chubinidze, Evgenia Alimbarashvili, Boris Vishnepolsky, Andrei Gabrielian, Alex Rosenthal, Darrell E Hurt, and Michael Tartakovsky. DBAASP v3: database of antimicrobial/cytotoxic activity and structure of peptides as a resource for development of new therapeutics. *Nucleic acids research*, 49(D1):D288–D297, 2021.

- Martin L Puterman. *Markov decision processes: Discrete stochastic dynamic programming*. John Wiley & Sons, 2014.
 - Lorenz Richter, Ayman Boustati, Nikolas Nüsken, Francisco Ruiz, and Omer Deniz Akyildiz. Vargrad: A low-variance gradient estimator for variational inference. *Neural Information Processing Systems (NeurIPS)*, 2020.
 - John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *arXiv preprint 1707.06347*, 2017.
 - Tony Shen, Mohit Pandey, and Martin Ester. Tacogfn: Target conditioned GFlowNet for drug design. In *NeurIPS 2023 Generative AI and Biology (GenBio) Workshop*, 2023.
 - Daniil Tiapkin, Nikita Morozov, Alexey Naumov, and Dmitry P Vetrov. Generative flow networks as entropy-regularized rl. In *International Conference on Artificial Intelligence and Statistics* (AISTATS), 2024.
 - Brandon Trabucco, Xinyang Geng, Aviral Kumar, and Sergey Levine. Design-bench: Benchmarks for data-driven offline model-based optimization. In *International Conference on Machine Learning (ICML)*, 2022.
 - Hanchen Wang, Tianfan Fu, Yuanqi Du, Wenhao Gao, Kexin Huang, Ziming Liu, Payal Chandak, Shengchao Liu, Peter Van Katwyk, Andreea Deac, et al. Scientific discovery in the age of artificial intelligence. *Nature*, 620(7972):47–60, 2023.
 - Michael L. Waskom. seaborn: statistical data visualization. *Journal of Open Source Software*, 6(60): 3021, 2021. doi: 10.21105/joss.03021. URL https://doi.org/10.21105/joss.03021.
 - Wolfram Wiesemann, Daniel Kuhn, and Berç Rustem. Robust Markov decision processes. *Mathematics of Operations Research*, 38(1):153–183, 2013.
 - Zachary Wu, Kadina E Johnston, Frances H Arnold, and Kevin K Yang. Protein sequence design with deep generative models. *Current opinion in chemical biology*, 65:18–27, 2021.
 - David W Zhang, Corrado Rainone, Markus Peschl, and Roberto Bondesan. Robust scheduling with GFlowNets. *International Conference on Learning Representations (ICLR)*, 2023.
 - Mingyang Zhou, Zichao Yan, Elliot Layne, Nikolay Malkin, Dinghuai Zhang, Moksh Jain, Mathieu Blanchette, and Yoshua Bengio. PhyloGFN: Phylogenetic inference with generative flow networks. *International Conference on Learning Representations (ICLR)*, 2024.

APPENDIX

A ROBUSTNESS-REGULARIZATION DUALITY

A.1 CONVEX ANALYSIS

Before proving our results, we briefly recall notions of convex analysis below Bertsekas (2009).

Definition A.1 (Convex conjugate). Let a function $f : \mathbb{R}^n \to \overline{\mathbb{R}}$ with domain $dom(f) \subseteq \mathbb{R}^n$. The convex conjugate of f is defined as

$$f^*(y) := \sup_{x \in \text{dom}(f)} \langle x, y \rangle - f(x).$$

Definition A.2 (Infimal convolution). Given two functions $f, g : \mathbb{R}^n \to \overline{\mathbb{R}}$, the infimal convolution of f and g is defined as:

$$[f\Box g](x) := \inf_{z \in \mathbb{R}^n} \{f(x-z) + g(z)\}.$$

Property A.1 (Operations on conjugate transforms). Let two functions $f_1, f_2 : \mathbb{R}^n \to \overline{\mathbb{R}}$ and a positive real number $\omega > 0$.

- (i) Defining $g(x) := \omega f_1(x)$, the convex conjugate of g satisfies $g^*(y) = \omega f_1^*(\frac{y}{\omega})$.
- (ii) For $\kappa \neq 0$ and $h(x) := f_1(\kappa x)$, the convex conjugate of h satisfies $h^*(y) = f_1^*(\frac{y}{\kappa})$.
- (iii) The convex conjugate of the sum $f_1 + f_2$ is the infimal convolution of their conjugates $f_1^* \Box f_2^*$, namely:

$$[f_1 + f_2]^*(y) = [f_1^* \square f_2^*](y), \quad \forall y \in \mathbb{R}^n.$$

A.2 APPLICATION TO POLICY DIVERGENCE

For Shannon entropy and KL divergence, the convex conjugate can be derived in closed form and is known to be a logsumexp function, so we omit the proof Geist et al. (2019); Derman et al. (2021).

Proposition A.1 (Shannon conjugate). Define the negative Shannon entropy $-H : \mathbb{R}^n \to \mathbb{R}$ as $[-H](x) := \sum_{i=1}^n x_i \log(x_i)$ over the simplex domain Δ_n . It is a convex function, and its convex conjugate is the logsumexp:

$$\mathrm{LSE}_n(y) := \log \left(\sum_{i=1}^n e^{y_i} \right), \qquad \forall y \in \mathbb{R}^n.$$

Proposition A.2 (KL conjugate). Let $d \in \mathbb{R}^n$ be such that d > 0 and define the KL-divergence function $\mathrm{KL}_d(x) := \sum_{i=1}^n x_i \log \left(\frac{x_i}{d_i}\right)$, $\forall x \in \Delta_n$. It is a convex function, and its convex conjugate is the weighted logsumexp:

$$\mathrm{WLSE}_{n,d}(y) := \log \left(\sum_{i=1}^n d_i e^{y_i} \right), \qquad \forall y \in \mathbb{R}^n.$$

We are now interested in deriving the Fenchel conjugate of a convex combination of Shannon entropy and KL divergence. Its explicit form is described below.

Proposition A.3 (Shannon-KL conjugate). For any $q \in [0, 1]$, the convex combination $(1-q)(-H) + qKL_d$ satisfies:

$$[(1-\mathsf{q})(-\mathsf{H})+\mathsf{qKL}_d](x)=\sum_{i=1}^n x_i\log\left(\frac{x_i}{(d_i)^\mathsf{q}}\right), \qquad \forall x\in\mathbb{R}^n.$$

and admits as convex conjugate the function:

$$[(1-\mathsf{q})(-\mathsf{H})+\mathsf{qKL}_d]^*(y)=\log\left(\sum_{i=1}^n(d_i)^\mathsf{q}e^{y_i}\right),\qquad\forall y\in\mathbb{R}^n.$$

Proof. The first statement comes from elementary algebra:

$$\begin{split} [(1-\mathbf{q})(-\mathbf{H}) + \mathbf{q} \mathbf{K} \mathbf{L}_d](x) &= (1-\mathbf{q}) \sum_{i=1}^n x_i \log(x_i) + \mathbf{q} \sum_{i=1}^n x_i \log\left(\frac{x_i}{d_i}\right) \\ &= (1-\mathbf{q}+\mathbf{q}) \sum_{i=1}^n x_i \log(x_i) - \mathbf{q} \sum_{i=1}^n x_i \log(d_i) \\ &= \sum_{i=1}^n x_i \log(x_i) - \sum_{i=1}^n x_i \log((d_i)^{\mathbf{q}}) \\ &= \sum_{i=1}^n x_i \log\left(\frac{x_i}{(d_i)^{\mathbf{q}}}\right) = \mathbf{K} \mathbf{L}_{d^{\mathbf{q}}}(x). \end{split}$$

For the second statement, we simply apply Prop. A.2 to obtain that

$$[(1-\mathsf{q})(-\mathsf{H}) + \mathsf{qKL}_d]^*(y) = \mathsf{WLSE}_{n,d^q}(y) = \log\left(\sum_{i=1}^n (d_i)^{\mathsf{q}} e^{y_i}\right).$$

A more involved proof would combine notions of infimal convolution with the respective conjugates of the two divergences. We omit it for brevity. \Box

A.3 PROOF OF THM. 5.1

We provide a slightly general proof below, with an arbitrary ϵ_s -level set instead of just 0.

Theorem A.1 (Fenchel-Robust Regularized MDP). For any $s \in \mathcal{S}$, define the reward set

$$\mathcal{R}_s := r_0(s, \cdot) + \{ r_s \in \mathbb{R}^{\mathcal{A}} : f_s^*(-r_s) \le \epsilon_s \},$$

and the regularization function $\Omega_s(\pi_s) := f_s(\pi_s) + \epsilon_s$, $\forall s \in \mathcal{S}$. Then, for any policy $\pi \in \Pi$, $T_{\mathcal{R}}^{\pi}v = T^{\pi,\Omega}v$, $\forall v \in \mathbb{R}^{\mathcal{S}}$, and $v_{\mathcal{R}}^{\pi} = v^{\pi,\Omega}$.

Proof. First, we establish the support function of a set $\tilde{\mathcal{R}}_s := \{r'_s \in \mathbb{R}^{\mathcal{A}} : f_s^*(r'_s) \leq \epsilon_s\}$ at any state $s \in \mathcal{S}$. By definition, for any policy $\pi_s \in \Delta_{\mathcal{A}}$, we have

$$\begin{aligned} \max_{r_s' \in \tilde{\mathcal{R}}_s} \langle \pi_s, r_s' \rangle &= \max_{\{r_s' \in \mathbb{R}^{\mathcal{A}} : f_s^*(r_s') \leq \epsilon_s\}} \langle \pi_s, r_s' \rangle \\ &= f_s(\pi_s) + \epsilon_s. \end{aligned}$$

Next, we compute the robust Bellman operator associated with the uncertainty set $\mathcal{R} = \times_s \mathcal{R}_s$:

$$\begin{split} T_{\mathcal{R}}^{\pi}v(s) &= \min_{r_s \in \mathcal{R}_s} r^{\pi}(s) + \gamma P^{\pi}v(s) \\ &= \min_{r_s \in r_0(s,\cdot) + \tilde{\mathcal{R}}_s} r^{\pi}(s) + \gamma P^{\pi}v(s) \\ &= r_0^{\pi}(s) + \min_{r_s' \in \tilde{\mathcal{R}}_s} \langle \pi_s, r_s' \rangle + \gamma P^{\pi}v(s) \\ &= r_0^{\pi}(s) - \max_{r_s' \in \tilde{\mathcal{R}}_s} \langle \pi_s, -r_s' \rangle + \gamma P^{\pi}v(s) \\ &= T_{r_0}^{\pi}v(s) - \max_{r_s' \in \tilde{\mathcal{R}}_s} \langle \pi_s, -r_s' \rangle. \end{split}$$

It remains to compute

$$\begin{split} \max_{r_s' \in \bar{\mathcal{R}}_s} \langle \pi_s, -r_s' \rangle &= \max_{r_s'} \langle \pi_s, -r_s' \rangle \text{ s. t. } f_s^*(-r_s') \leq \epsilon_s \\ &= \max_{\bar{r}} \langle \pi_s, \bar{r}_s \rangle \text{ s. t. } f_s^*(\bar{r}_s) \leq \epsilon_s \end{split}$$

Employing the above expression of the support function enables us to write:

$$\max_{\bar{r}_s} \langle \pi_s, \bar{r}_s \rangle \text{ s. t. } f_s^*(\bar{r}_s) \leq \epsilon_s = f_s(\pi_s) + \epsilon_s$$

so that $T_{\mathcal{R}}^{\pi}v=T^{\pi,\Omega}v, \quad \forall v\in\mathbb{R}^{\mathcal{S}}$. The unique fixed point of each operator is $v_{\mathcal{R}}^{\pi}$ and $v^{\pi,\Omega}$, respectively, which leads to the conclusion.

Equation 12 comes from combining Thm. 5.1 with Prop. A.3. We can similarly obtain the uncertainty set resulting from KL and negative Shannon regularizations, based on their corresponding dual described in Sec. A.1. This leads us to the following table, summarizing the different regularizers used in this paper along with their corresponding uncertainty sets.

A.4 PROOF OF PROP. 4.1

Proposition A.4. For all $\mathbf{q} \in [0,1]$, $\Omega_s^{\mathbf{q}}(\pi_s) = \frac{1}{\omega_s} \mathrm{KL}(\pi_s, d_s^{(\mathbf{q})}) - \frac{1}{\omega_s} \log(Z_{\mathbf{q}}(d_s))$, where $Z_{\mathbf{q}}(d_s) := \sum_{a \in \mathcal{A}} d_s(a)^{\mathbf{q}}$ and $d_s^{(\mathbf{q})} = d_s^{\mathbf{q}}/Z_{\mathbf{q}}(d_s)$ is the \mathbf{q} -tilted softmax distribution.

Proof. By definition,

$$\begin{split} \omega_s \Omega_s^{\mathbf{q}}(\pi_s) &= \left[\mathsf{qKL}(\pi_s, d_s) + (1 - \mathsf{q})(-\mathsf{H}(\pi_s)) \right] \\ &= \mathsf{q} \sum_{a \in \mathcal{A}} \pi_s(a) \log \left(\frac{\pi_s(a)}{d_s(a)} \right) + (1 - \mathsf{q}) \sum_{a \in \mathcal{A}} \pi_s(a) \log(\pi_s(a)) \\ &= -\mathsf{q} \sum_{a \in \mathcal{A}} \pi_s(a) \log(d_s(a)) + \sum_{a \in \mathcal{A}} \pi_s(a) \log(\pi_s(a)) \\ &= \sum_{a \in \mathcal{A}} \pi_s(a) \log \left(\frac{\pi_s(a)}{d_s(a)^{\mathbf{q}}} \right) \\ &= \sum_{a \in \mathcal{A}} \pi_s(a) \log \left(\frac{\pi_s(a)}{d_s(a)^{\mathbf{q}}/Z_{\mathbf{q}}(d_s)} \right) - \sum_{a \in \mathcal{A}} \pi_s(a) \log(Z_{\mathbf{q}}(d_s)) \\ &= \sum_{a \in \mathcal{A}} \pi_s(a) \log \left(\frac{\pi_s(a)}{d_s(a)^{\mathbf{q}}/Z_{\mathbf{q}}(d_s)} \right) - \log(Z_{\mathbf{q}}(d_s)), \end{split}$$

which yields the desired result.

Table 2: Summary table of policy regularizers.

	Neg. Shannon	KL divergence	GSM	General convex
Regularizer Ω_s	$- ext{H}(\pi_s)$	$ ext{KL}(\pi_s,d_s)$	$rac{1}{\omega_s}(extsf{qKL}(\pi_s,d_s) \ + (1- extsf{q})(- extsf{H}(\pi_s)))$	$f_s(\pi_s) + \epsilon_s$
Conjugate Ω_s^*	${\scriptscriptstyle LSE_\mathcal{A}(q_s)}$	$ ext{WLSE}_{\mathcal{A},d_s}(q_s)$	$\omega_s^{-1} ext{WLSE}_{\mathcal{A}, d_s^q} (\omega_s q_s)$	$f_s^*(q_s) - \epsilon_s$
Reward Uncer- tainty	$\{r_s: \Omega_s^*(-r_s) \le 0\}$	$\{r_s: \Omega_s^*(-r_s) \le 0\}$	$\{r_s: \Omega_s^*(-r_s) \le 0\}$	$\left\{r_s: f_s^*(-r_s) \le \epsilon_s\right\}$

A.5 COMPARISON WITH OTHER REGULARIZED RL WORKS

Eysenbach & Levine (2021) focuses on Shannon entropy, while we encompass a broad class of regularizers. Husain et al. (2021); Brekelmans et al. (2022) analyze robustness from the LP-dual perspective of RL. Although equivalent at optimum, this approach may not hold for any given policy. The notion of occupancy measure is also obscure in the context of DCP where transitions are fully determined by the agent's decisions and the time horizon is finite. Differently, Thm. 5.1 establishes

a robustness-regularization equivalence for any policy via Bellman evaluation operators. It thus provides a principled identity between robust dynamic programming in the sense of (Iyengar, 2005) and regularized MDPs in the sense of (Geist et al., 2019). Finally, the research motivation of Derman et al. (2021) being different, they proceed the opposite way from ours: they deduce a regularizer from generic uncertainty sets, whereas we deduce uncertainty sets from generic regularizers. This enables us to clarify the robustness properties caused by regularization.

B Additional results

B.1 IMPACT OF β ON GFN PERFORMANCE

Another potential solution to the issue of sampling proportional to $e^{\Phi(x)}$ consists of modifying the reward function. In particular, the reward exponent hyperparameter β could be used to arbitrarily increase the value of high-reward objects, hopefully overwhelming even an exponential amount of low-reward objects. To test this hypothesis, we perform an additional grid sweep for GFNs using higher values of β . We sweep over the same learning rates $\{0.00001, 0.0001, 0.001\}$ and sampling temperatures $\{1, 4, 16\}$ but also over $\beta \in \{512, 1024, 2048, 4096, 8192\}$. We plot the final average mode reward of the best performing setting for each β in Fig. 6.

Going past 256 worsens performance for UTR and AMP but does noticeably improve performance for GFP. However, for all values of β , GFNs do not manage to reach the best-performing TGM setting. This discrepancy indicates TGM is exploring meaningfully different peaky distributions with better quality and diversity. Surprisingly, GFNs still perform decently even with very high β values (up to 8192). It seems that gradient clipping and the VarGrad objective are enough to ensure some level of training stability at these high β values. These values yield equally high losses, reaching values in the millions for $\beta=8192$ on GFP.

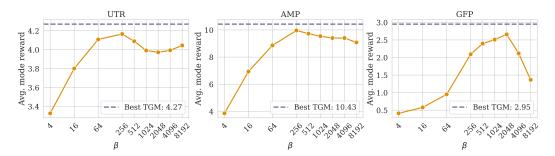


Figure 6: Average mode reward of best GFN setting for different values of β compared to the best performing setting of TGM.

B.2 COMPUTATIONAL COMPLEXITY

TGM is a generalization of TB but nearly identical in terms of computational complexity (essentially just requires an additional log-softmax in the loss computation). In particular, for sequence tasks with transformers, by acting on trajectories, TGM/TB are significantly faster than methods that act on transitions. The loss computation only requires a single forward pass per batch of trajectories instead of a forward pass per batch of transitions. We've included a table of training speed on an L40s GPU below. TGM/TB are roughly 10x faster than PPO/SAC when the latter splits generated trajectories into minibatches. TGM is only slightly more complicated to implement than GFNs, and simpler to implement than PPO (see the code submission for details).

Table 3: Wallclock training speed on an L40s for various methods on biological sequence tasks.

Algorithm	UTR (samples/s) ↑	AMP (samples/s) ↑	GFP (samples/s) ↑
SAC	115	92	5
PPO	191	144	8
TGM/TB	1680	1216	120

C MULTI-STEP UNCERTAINTY SETS

 While Fig. 3 examines the case of uncertainty sets for a single step, we now seek to explore the values that can be taken by δ_i throughout a trajectory and how this affects the final uncertainty set. For the entropic regularizer, if ω_s is the same for all s, then \mathcal{R}_{s_i} is the same at every state. However, the final uncertainty set \mathcal{R} for an object $x = s_1 s_2 \dots s_n$ is given by $\sum_i \mathcal{R}_{s_i}$ (i.e. the Minkowski sum of the individual uncertainty sets).

We extend the binary generation example to a DCP consisting of a sequence generation task with two tokens (i.e. a binary tree) and length d. Then, using Prop C.1, in Fig. 7, we illustrate how the sum of uncertainty sets change as the depth of the tree increases for different operators.

Proposition C.1. For a convex regularizer f and associated robust set $\mathcal{R}_f = \{\delta \in \mathbb{R}^{|\mathcal{A}|} : f^*(\delta) \leq 0\}$, we have that:

$$\sum_{i=1}^{k} \mathcal{R}_f = \{ r \in \mathbb{R}^{|\mathcal{A}|} : k f^* \left(\frac{r}{k} \right) \le 0 \}.$$
 (13)

Proof. We seek to show that, for a given $r \in \mathbb{R}^{|\mathcal{A}|}$, there exists $\delta_1, \delta_2 \dots \delta_k$ such that $\sum_{i=1}^k \delta_i = r$ and $\delta_i \in \mathcal{R}_f$ if and only if $kf^*\left(\frac{r}{k}\right) \leq 0$.

 $[\Rightarrow]$ Since f is convex, so is f^* . Thus, using Jensen's inequality, we get that:

$$kf^*\left(\frac{r}{k}\right) = kf^*\left(\frac{\sum_i \delta_i}{k}\right) \le \sum_i f^*(\delta_i) \le 0,$$

where the last inequality stems from $\delta_i \in \mathcal{R}_f \implies f^*(\delta_i) \leq 0$.

[
$$\Leftarrow$$
] Taking $\delta_i = \frac{r}{k}$ for all δ_i , we have that $\sum_i \delta_i = r$. Then, $kf^*\left(\frac{r}{k}\right) \leq 0 \implies \forall i: f^*(\delta_i) = f^*\left(\frac{r}{k}\right) \leq 0$ and thus $\delta_i \in \mathcal{R}_f$.

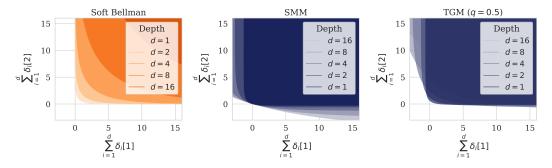


Figure 7: Illustration of sum of uncertainty sets along a trajectory. (Left) The final uncertainty set shifts farther away from (0,0) as we increase d. As such, for objects that are deeper in the tree, the operator is only robust to them having much higher reward than their proxy evaluation. This behavior provides an alternative viewpoint for the preference of the soft Bellman operator for longer objects. (Middle) On the other hand, the SMM operator stays centered around (0,0) and becomes increasingly robust to decreases in rewards. (Right) The final uncertainty set of the GM operator stays close to (0,0) and is robust to some decreases in reward.

D TRAJECTORY GENERAL MELLOWMAX

Notation: We denote by $\mathrm{Ch}(s)$ the set of possible actions from state s. $\sigma_{\tau}(Q_s) := \frac{e^{\tau Q_s}}{\sum_{a' \in \mathrm{Ch}(s)} e^{\tau Q_s[a']}}$ is the softmax with inverse temperature τ . For a DCP and a given state s, we denote by s'_a the state reached by taking action a (unique since transitions are deterministic). We denote by E the set of edges in \mathcal{G} , each of which represents a transition (s, a, s'_a) . Finally, for convenience, we denote

$$\begin{split} g(Q_s,a) &:= \frac{1}{\omega} \log \sigma_{\mathsf{q}\alpha+\omega}(Q_s)[a] - \mathsf{q} \log \sigma_\alpha(Q_s)[a], \\ g^*(Q_s) &:= \frac{1}{\omega} \log \langle \sigma_\alpha(Q_s)^\mathsf{q}, e^{\omega Q_s} \rangle. \end{split}$$

D.1 PROOF OF OPTIMALITY

Proposition D.1. In a DCP (G, r), for any state s,

$$v_s^* := \max_{\pi_s} \langle \pi_s, Q_s \rangle - \frac{1}{\omega} \Omega_\alpha^{\mathbf{q}}(\pi_s) = \frac{1}{\omega} \log \langle \sigma_\alpha(Q_s)^{\mathbf{q}}, e^{\omega Q_s} \rangle = g^*(Q_s), \tag{14}$$

and

$$\pi_s^* = \sigma_{\mathsf{q}\alpha + \omega}(Q_s). \tag{15}$$

is the unique maximizer of the state-wise regularized problem.

Proof. Unlike the case of a generic π , (14) is not a convex conjugate since Q_s appears in both the dot product and Ω . Nonetheless, Ω is still convex in π . Thus, similarly to (Nachum et al., 2017), we consider the Lagrangian:

$$L(\pi_s, \lambda) = \langle \pi_s, Q_s \rangle - \frac{1}{\omega} \Omega_{\alpha}^{\mathbf{q}}(\pi_s) + \lambda (1 - \langle \mathbf{1}, \pi_s \rangle).$$

Since Ω is convex, the overall function to optimize is concave. Given that Slater's condition holds, the KKT conditions are both sufficient and necessary for optimality. Due to the lack of cycle in DCPs, we have that $\nabla_{\pi_s} Q_s = 0$. Hence, the KKT conditions yield the following system of equations:

$$0 = Q_s - \frac{1}{\omega} \left(q[\log \pi_s - \log \pi_{Q_s} + 1] + (1 - q)[\log \pi_s + 1] \right) - \lambda \tag{16}$$

$$1 = \langle \pi_s, \mathbf{1} \rangle \tag{17}$$

where we denote $\pi_{Q_s} := \sigma_{\alpha}(Q_s)$. Solving for π_s in the first, we get that the optimal policy π_s^* is given by:

$$\pi_s^* = \exp([Q_s - \lambda]\omega + q \log \pi_{Q_s} - 1)$$

which ensures that $\pi_s \geq 0$. Replacing the above in (17), we have that:

$$\lambda^* = \frac{1}{\omega} \left(\log \sum_{a \in \operatorname{Ch}(s)} e^{(\alpha \mathsf{q} + \omega)Q_s[a]} - \mathsf{q} \log \sum_{a \in \operatorname{Ch}(s)} e^{\alpha Q_s[a]} \right) = \frac{1}{\omega} \left(\log \langle \sigma_\alpha(Q_s)^{\mathsf{q}}, e^{\omega_s Q_s} \rangle - 1 \right).$$

We note the similarity of λ^* and the maximal value (14). Plugging λ^* back in to π^* , we get

$$\pi_s^* = \sigma_{\mathsf{q}\alpha + \omega}(Q_s). \tag{15}$$

Finally, plugging π_s^* in $\langle \pi_s, Q_s \rangle - \frac{1}{\omega} \Omega_{\alpha}^{\mathbf{q}}(\pi_s)$ with some straightforward algebraic manipulations, we obtain (14). Since both the KL divergence and the negative entropy are strictly convex, (15) is the unique maximizer.

Corollary D.1. For any DCP (\mathcal{G}, Φ) , the GM operator has unique optimal value/policy functions (v^*, π^*) .

Proof. We show this by induction on an inverse topological sorting s_0, s_1, \ldots, s_n of the states of \mathcal{G} . For a state with no children, the optimal value is defined to be 0 and there is no policy. Then, assume each state $s_i, i < N$ has a unique optimal value/policy. By the properties of the topological sorting, all the children of s_N have a unique optimal value/policy. Then, by Prop. D.1, $(v_{s_N}^*, \pi_{s_N}^*)$ is also uniquely defined.

Lemma D.1. For any Q_s and a, we have that

$$g(Q_S, a) = Q_s[a] - g^*(Q_s).$$
 (18)

Proof. By simple algebra,

$$\begin{split} &g(Q_s,a) \\ &= \frac{1}{\omega} (\log \sigma_{\mathbf{q}\alpha + \omega}(Q_s)[a] - \mathbf{q} \log \sigma_{\alpha}(Q_s)[a]) \\ &= \frac{1}{\omega} \left((\mathbf{q}\alpha + \omega)Q_s[a] - \mathbf{q}\alpha Q_s[a] \right) \\ &- \frac{1}{\omega} \left(\log \sum_{a' \in \operatorname{Ch}(s)} e^{(\mathbf{q}\alpha + \omega)Q_s[a']} - \log \left[\sum_{a' \in \operatorname{Ch}(s)} e^{\alpha Q_s[a']} \right]^{\mathbf{q}} \right) \\ &= \frac{1}{\omega} (\omega Q_s[a]) - \frac{1}{\omega} \log \left[\sum_{a' \in \operatorname{Ch}(s)} \frac{e^{\mathbf{q}\alpha Q_s[a']} e^{\omega Q_s[a']}}{\left(\sum_{a' \in \operatorname{Ch}(s)} e^{\alpha Q_s[a']} \right)^{\mathbf{q}}} \right] \\ &= Q_s[a] - \frac{1}{\omega} \left(\log \langle \sigma_{\alpha}(Q_s)^{\mathbf{q}}, e^{\omega_s Q_s} \rangle \right) \\ &= Q_s[a] - g^*(Q_s) \end{split}$$

Lemma D.2. For any state s in a DCP, the following are equivalent:

$$\sigma_{\alpha \mathbf{q} + \omega}(Q_s^{\theta}) = \arg\max_{\pi_s} \langle \pi_s, Q_s^* \rangle - \frac{1}{\omega} \Omega_{\sigma_{\alpha}(Q_s^*), s}^{\mathbf{q}}(\pi_s), \tag{19}$$

$$\forall a \in \text{Ch}(s) : v_{s'_a}^* = v_s^* + g(Q_s^{\theta}, a) - r(s, a).$$
(20)

Proof.

[(19) \Rightarrow (20)] Since π_s^* is unique, by (19) and (15), by the properties of the softmax, we have that $\sigma_{\alpha q + \omega}(Q_s^{\theta}) = \sigma_{q\alpha + \omega}(Q_s^*)$ and $\sigma_{\alpha}(Q_s^{\theta}) = \sigma_{\alpha}(Q_s^*)$. Then, for any $a \in \mathrm{Ch}(s)$:

$$\begin{aligned} v_s^* + g(Q_s^{\theta}, a) \\ &= v_s^* + \frac{1}{\omega} \left(\log \sigma_{\mathbf{q}\alpha + \omega}(Q_s^{\theta})[a] - \mathbf{q} \log \sigma_{\alpha}(Q_s^{\theta})[a] \right) \\ &= \frac{1}{\omega} \left(\log \langle \sigma_{\alpha}(Q_s^*)^{\mathbf{q}}, e^{\omega_s Q_s^*} \rangle + \log \sigma_{\mathbf{q}\alpha + \omega}(Q_s^*)[a] - \mathbf{q} \log \sigma_{\alpha}(Q_s^*)[a] \right) \end{aligned} \qquad \text{Using (14)}.$$

$$&= \frac{1}{\omega} \left(\log \frac{\sum_{a' \in \text{Ch}(s)} e^{(\mathbf{q}\alpha + \omega)Q_s^*[a']}}{\sum_{a' \in \text{Ch}(s)} e^{\alpha Q_s^*[a']}} + \log \sigma_{\mathbf{q}\alpha + \omega}(Q_s^*)[a] - \mathbf{q} \log \sigma_{\alpha}(Q_s^*)[a] \right)$$

$$&= \frac{1}{\omega} \left(\log e^{(\mathbf{q}\alpha + \omega)Q_s^*[a]} - \mathbf{q} \log e^{\alpha Q_s^*[a]} \right)$$

$$&= Q_s^*[a]$$

$$&= v_{s'}^* + r(s, a) \qquad \qquad \text{Using the definition of } Q_s[a].$$

[(19) \Leftarrow (20)] For any $a \in Ch(s)$, we have that:

$$v_{s'_{a}}^{*} = v_{s}^{*} + g(Q_{s}^{\theta}, a) - r(s, a)$$

$$\iff v_{s'_{a}}^{*} = v_{s}^{*} + Q_{s}^{\theta}[a] - g^{*}(Q_{s}^{\theta}) - r(s, a)$$

$$\iff v_{s}^{*} - g^{*}(Q_{s}^{\theta}) = Q_{s}^{*}[a] - Q_{s}^{\theta}[a]$$
Using (18)

Since the above equality holds for all actions and v_s^* and $g^*(Q_s^\theta)$ do not depend on a, we have that $Q_s^*[a] - Q_s^\theta[a]$ must be equal for all actions. Thus, $\forall a \in \operatorname{Ch}(s) : Q_s^*[a] - Q_s^\theta[a] = c$ which implies

$$Q_s^{\theta}[a] = Q_s^*[a] - c, \tag{21}$$

for some action independent value c. Finally, we have

$$\sigma_{\alpha q + \omega}(Q_s^{\theta}) = \sigma_{\alpha q + \omega}(Q_s^* - c) = \sigma_{\alpha q + \omega}(Q_s^*),$$

by the properties of the softmax.

D.2 PROOF OF THEOREM 4.1

We note that it is straightforward to derive the equivalent subtrajectory constraint.

Theorem D.1 (Trajectory GM). For any DCP (\mathcal{G}, Φ) , let (v^*, π^*) be the unique optimal value/policy functions for the GM operator. Then, for a given Q-function Q^{θ} ,

$$\sigma_{\alpha \mathbf{q} + \omega}(Q_s^{\theta}) = \arg\max_{\pi_s} \langle \pi_s, Q_s \rangle - \frac{1}{\omega} \Omega_{\alpha}^{\mathbf{q}}(\pi_s), \tag{8}$$

holds for all states s if and only if

$$v_0^* + \sum_{i=0}^n \underbrace{\frac{1}{\omega} \left(\log \sigma_{\mathbf{q}\alpha+\omega}(Q_{s_i}^{\theta})[a_i] - \mathbf{q} \log \sigma_{\alpha}(Q_{s_i}^{\theta})[a_i] \right)}_{g(Q_{s_i}^{\theta}, a_i)} - r(s_i, a_i) = 0, \tag{9}$$

holds for all full trajectories $s_0 \xrightarrow{a_0} \cdots \xrightarrow{a_{n-1}} s_n \xrightarrow{a_n} x$ in the DCP.

Proof. Similarly to (Nachum et al., 2017), we aim to show that optimality implies consistency and vice versa. A refresher of relevant notation can be found at the start of App. D.

Optimality implies consistency [$(8) \Rightarrow (9)$] From Lemma D.2, for any given trajectory, we have that (20) holds for all transitions in the trajectory. Then, (9) follows straightforwardly by expanding (20) over that trajectory in a recursive manner.

Consistency implies optimality [(8) \Leftarrow (9)] Consider an inverse topological sorting $s^1, s^2 \dots s^M$ of \mathcal{G} such that for any state s^i , any child s^j of s^i has index j < i.

Now, we seek to show by induction over this sorting that the following holds for any state s and subtrajectory $s_0 \to s_1 \ldots \to s_n \to s$ ending in s:

$$v_0^* + \sum_{i=0}^n \frac{1}{\omega} \left(\log \sigma_{q\alpha+\omega}(Q_{s_i}^{\theta})[a_i] - q \log \sigma_{\alpha}(Q_{s_i}^{\theta})[a_i] \right) - r(s_i, a_i) = v_s^*.$$
 (22)

For the base case s^1 , s^1 must be a leaf and belong to \mathcal{X} . Then, by assumption $v_{s^1} = 0$ and (22) holds by (9). Now, suppose that (22) holds for any s^i , i < N. Then, by the definition of the sorting, for all $a \in \operatorname{Ch}(s^N)$, (22) holds for s_a where $s^N \stackrel{a}{\to} s_a$.

Let $s_0 \to s_1 \dots \to s_n \to s^N$ be a trajectory ending in s^N . From the induction hypothesis, $\forall a \in Ch(s)$:

$$v_0^* + \sum_{i=0}^n g(Q_{s_i}^{\theta}, a_i) - r(s_i, a_i) + g(Q_{s^N}^{\theta}, a) - r(s^N, a) = v_{s_a}^*$$

$$v_0^* + \sum_{i=0}^n g(Q_{s_i}^{\theta}, a_i) - r(s_i, a_i) - g^*(Q_{s^N}^{\theta}) = Q_{s_N}^*[a] - Q_{s_N}^{\theta}[a] \quad \text{Using (18)}.$$

Similarly to the proof of Lemma D.1, we have that the LHS is independent of a but must be equal to $Q_{s_N}^*[a] - Q_{s_N}^{\theta}[a]$ for all a. As such, we are guaranteed

$$Q_{s_N}^{\theta}[a] = Q_{s_N}^*[a] - c$$

for some action independent value c and all $a \in \operatorname{Ch}(s^N)$. Then, it is straightforward to compute $g^*(Q^\theta_{s^N}) = g^*(Q^*_{s_N} - c) = g^*(Q^*_{s_N}) - c$. Plugging both back in and canceling c, we get

$$v_0^* + \sum_{i=0}^n g(Q_{s_i}^\theta, a_i) - r(s_i, a_i) = g^*(Q_{s_N}^*) = v_{s_N}^*$$

and the induction is complete. Finally, it remains to show that $\sigma_{\alpha q+\omega}(Q_s^{\theta})$ is optimal for all s. If s has no children, optimality holds trivially. If s has at least one child a, take a trajectory

 $s_0 \to s_1 \dots \to s_n \to s$ ending in s. From (22), we have that both:

$$v_0^* + \sum_{i=0}^n g(Q_{s_i}^{\theta}, a_i) - r(s_i, a_i) = v_s^*$$

$$\forall a \in \operatorname{Ch}(s) : v_0^* + \sum_{i=0}^n g(Q_{s_i}^{\theta}, a_i) - r(s_i, a_i) + g(Q_s^{\theta}, a) - r(s, a) = v_{s_a}^*.$$

By subtracting the first from the second, we get that, $\forall a \in Ch(s)$:

$$g(Q_s^{\theta}, a) - r(s, a) = v_{sa}^* - v_s^*$$
$$v_s^* + g(Q_s^{\theta}, a) - r(s, a) = v_{sa}^*.$$

By Lemma D.2, $\sigma_{\alpha {\bf q} + \omega}(Q^{\theta}_s)$ is thus optimal at s and we are done.

E EXPERIMENTAL DETAILS

E.1 ALGORITHMS

We follow most of the hyperparameter decisions of (Malkin et al., 2022). For each method, we train a transformer with 3 layers, 8 attention heads, and an embedding dimension of 64. We set dropout to 0.1 and use causal masking. Each output head uses a fully-connected network with 2 hidden layers of dimensions 256.

We use weight decay of 10^{-4} and use gradient clipping with a threshold of 10 to help stabilize training at higher β values. When generating samples for training, we use a policy corresponding to a mixture of the method's policy and a uniform policy (with weight 0.01). Each method is trained with Adam (Kingma, 2014) with $\epsilon = 10^{-5}$.

TGM/GFN: Since GFNs correspond to a specific hyperparameter setting for TGM, training details are essentially identical between the two methods. Both methods are trained online with a batch size of 16 trajectories and have a single head that outputs logits. We sample trajectories by taking the tempered softmax of the logits, using the temperature $\alpha q + \omega$ from the optimal policy (15). For GFNs, ω affects the sampling policy but *not the actual training objective* as this ensures the network is learning the actual GFN objective.

SAC: For SAC, we use three networks: two Q networks (to reduce overestimation bias) and a policy network, following the implementation of (Christodoulou, 2019; Huang et al., 2022). We use a replay buffer of size 100 000 and a batch size of 1 024 transitions. Generation/training is adjusted to have a replay ratio of 1.0. The entropy coefficient is fixed at $1/\omega$, and we use target networks (for the Q networks), which get updated every 10 iterations.

PPO: For PPO, we use two networks: a policy network and a value network, following the implementation of (Huang et al., 2022). We follow the hyperparameter settings of (Huang et al., 2022) and use a generalized advantage estimator λ of 0.95 and $\gamma=1$ to avoid discounting. The clipping epsilon is 0.1 and the value coefficient is 0.5. During online training, 16 trajectories are generated, and the resulting transitions are split into 4 minibatches for training. The entropy coefficient is set to $1/\omega$.

E.2 SEQUENCES

For each sequence generation task, we add the following 3 tokens to the vocabulary: BOS, PAD, EOS. Sequences are padded to be of length equal to the maximum length for the task, with an added BOS, EOS. Sequences look like the following: $[BOS, x_1, \ldots, x_n, PAD, \ldots, PAD, EOS]$. When generating sequences, they start with BOS, and tokens are added until a terminating action is chosen, at which point PAD/EOS are added automatically. The terminating action is masked until the task's minimum length is reached.

E.3 SYNTHETIC TASKS

TF-Bind-8. Hyperparameters are as described in the paper. We sample 10 000 sequences for each method using their respective optimal policy. Fig. 4 is the kernel density estimate computed using these samples with the default parameters used by seaborn (Waskom, 2021).

Bit sequence. We use n=120 and k=8 with M=60 modes and force sequences to be of length 120 through action masking. The modes are sampled randomly and fixed across runs to ensure there is no variation from certain runs having harder-to-find modes. Each method is trained for 200 000 samples.

E.4 BIOLOGICAL SEQUENCE DESIGN

For each task, once the proxy reward model is trained, we compute the mean and standard deviation of the logits on the validation set. Then, we use the following normalized value as reward

1297
1298 $r(x) = \beta \Phi(x) = \beta \frac{\phi(x) - \mu_{\text{val}}}{\sigma_{\text{val}}}, \tag{23}$

where $\phi(x)$ is the logit output by the trained model. To train the proxy reward models, we follow the hyperparameters used by (Malkin et al., 2022).

UTR: Using the described dataset, we train a transformer with 4 layers, 8 attention heads, and an embedding dimension of 64. 80% of the data is used for training and the rest for validation. We train for 250 epochs with early stopping (using a patience of 15). The learning rate is set to 10^{-4} , and we use a batch size of 128 and weight decay of 1e-6. We set both the minimum and maximum length to 50 since all sequences in the dataset are of length 50.

AMP: We use the same hyperparameters for training the model as UTR. Since the dataset has sequences of variable length, we set the minimum length to 14 and the maximum length to 60. Unlike the other tasks, the model is trained for binary classification, and we use the logit passed to the sigmoid as $\phi(x)$.

GFP: GFP is the largest task. We train a transformer with 3 layers, 8 attention heads, and an embedding dimension of 128. Since all sequences are of length 237, we set both the minimum length and maximum length to 237.

E.4.1 EVALUATION

 For evaluation of the average mode reward, we sweep over inverse temperature modifiers $t \in \{0.005, 0.01, 0.02, 0.05, 0.1, 0.2, 0.5, 1, 2, 5\}$ (i.e. if the regular sampling policy was $\pi \propto e^{\tau x}$, we sample using $\pi' \propto e^{\tau tx}$) and generate 512 samples for each. We then aggregate the samples and greedily find the top 100 samples such that they are at least at a distance δ of each other. We use the Levenshtein edit distance as metric (Levenshtein et al., 1966) and set $\delta = 0.25 \cdot \frac{\text{minLen} + \text{maxLen}}{2}$ (i.e. $\delta = 13$ for UTR, $\delta = 10$ for AMP and $\delta = 60$ for GFP). The constraint can be roughly interpreted as requiring at least 25% of a sequence to be different for it to be considered distinct. In Tab. 1, we report the max of this value achieved over the course of training.

E.5 COMPUTE RESOURCES

Experiments were conducted on a cluster with a mix of GPUs. All experiments were run using jobs that requested 1xL40s, 24G of RAM, and 4 CPUs. The cluster allows for dozens of jobs to be run in parallel. Overall, the experiments (including all the hyperparameter sweeps, baselines, and various seeds) used roughly:

• Bit Sequence: 35 GPU hours.

UTR: 48 GPU hours.AMP: 48 GPU hours.GFP: 300 GPU hours.

for a total of roughly ≈ 20 GPU days. In particular, the entire codebase uses Jax (Bradbury et al., 2018) and a parallelized environment allowing for fast generation. Since we use transformers, training on a generated trajectory for TGM/GFNs only requires a single forward/backward pass. This significantly speeds up training (roughly 4x faster) compared to the PPO/SAC baselines, which operate on individual transitions (though there are potential engineering optimizations to improve the speed of these baselines).

E.6 LLM USAGE DECLARATION

LLMs were very lightly used by the authors for polishing/formatting as well as routine coding applications (e.g. plotting, creating scripts, etc.).