Supported Policy Optimization for Offline Reinforcement Learning

Anonymous Author(s)
Affiliation
Address
email

Abstract

Policy constraint methods to offline reinforcement learning (RL) typically utilize parameterization or regularization that constrains the policy to perform actions within the support set of the behavior policy. The elaborative designs of parameterization methods usually intrude into the policy networks, which may bring extra inference cost and cannot take full advantage of well-established online methods. Regularization methods reduce the divergence between the learned policy and the behavior policy, which may mismatch the inherent density-based definition of support set thereby failing to avoid the out-of-distribution actions effectively. This paper presents **S**upported **P**olicy **O**p**T**imization (SPOT), which is directly derived from the theoretical formalization of the density-based support constraint. SPOT adopts a VAE-based density estimator to explicitly model the support set of behavior policy and presents a simple but effective density-based regularization term, which can be plugged non-intrusively into off-the-shelf off-policy RL algorithms. SPOT achieves the state-of-the-art performance on standard benchmarks for offline RL. Benefiting from the pluggable design, offline pretrained models from SPOT can also be applied to perform online fine-tuning seamlessly.

1 Introduction

Offline RL [23, 24], where the agent learns from a fixed dataset, collected by arbitrary process, not only provides a bridge between RL and the data-driven paradigm but also eliminates the need to interact with the live environment, which is always expensive or risky in practical scenarios [14, 26, 16]. Unfortunately, the absence of environment interaction also raises a number of challenges. Previous work has shown that the extrapolation error of the Q-function queried by out-of-distribution actions significantly degrades the performance of off-policy algorithms [8].

Avoiding out-of-distribution actions, namely to constrain the learned policy to perform actions within the support set of the behavior policy, is essential to mitigate extrapolation error. To meet this support constraint, policy constraint methods [24] to offline RL utilize either *parameterization* [8, 43, 10] or *regularization* [13, 21, 40] techniques. However, there are still several drawbacks in previous methods of policy constraint, limiting their performance and applications. Firstly, parameterization methods involve elaborate designs of the policy parameterization, typically coupled to generative models of the behavior policy, to *directly* constrain actions taken by the learned policy. But these designs intrude into the architecture of policy networks, which may bring extra inference cost and supplementary difficulties to implement and tune offline RL algorithms. Furthermore, as criticized by Fujimoto and Gu [6], these *intrusive* designs complicate causal attribution of performance gains and transfer of techniques between offline RL algorithms or from well-established online RL algorithms. In contrast, for the second category, regularization methods are designed with a non-intrusive or *pluggable* manner, which is done by simply augmenting a penalty to the actor loss that measures the divergence

Submitted to 36th Conference on Neural Information Processing Systems (NeurIPS 2022). Do not distribute.

of the learned policy from the behavior policy [13, 21, 40, 19]. However, these divergence-based regularization methods may mismatch the inherent density-based definition of support set, which applies support constraint *indirectly*. As we show in experiments (Section 5.1), the divergence-based design prevents regularization methods from effectively avoiding out-of-distribution actions and thus limits their performance.

In this work, we aim to design a *pluggable* offline RL method that also *directly* meets the standard formalization of the support constraint based on the density of the behavior policy. We introduce **S**upported **P**olicy **O**p**T**imization (SPOT), a regularization method which can be plugged non-intrusively into off-the-shelf off-policy RL algorithms. SPOT involves a regularization term from the new perspective of explicit estimation of the behavior density. Concretely, SPOT adopts a VAE-based density estimator [17, 31] to explicitly model the support set of behavior policy and presents a simple yet effective regularization term directly applied to the estimated density.

Our method benefits from a closer connection between theory and algorithm, thereby achieving superior performance compared to previous methods. It also profits from the pluggable algorithmic design, which leads to efficient inference and minor algorithmic modifications on top of standard off-policy algorithms. Moreover, a minimal gap between offline learning objectives and standard online learning objectives also enables SPOT to take full advantage of existing online RL algorithms and attain strong online fine-tuning performance after offline RL, exceeding state-of-the-art methods.

The main contributions of this work are three-fold:

- We derive a regularization term for constrained policy optimization in offline RL, based on the theoretical formalization of the density-based support constraint, which directly regularizes the behavior density of actions taken by the learned policy.
- We propose Supported Policy OpTimization (SPOT), a practical algorithm with a neural VAE-based density estimator to implement the proposed regularization term.
- Compared to several strong baselines, SPOT achieves state-of-the-art results on standard offline RL benchmarks [5] and also outperforms previous methods when online fine-tuned after offline RL initialization.

2 Related Work

Our work belongs to the family of policy constraint methods in offline RL, where parameterization or regularization is typically utilized. Comparison between previous policy constraint methods and ours has been summarized in Table 1. In addition to the policy constraint methods, we also review behavior policy modeling in offline RL, which is commonly necessary to both parameterization and regularization methods. Lastly, we briefly discuss a broader range of competitive offline RL approaches.

Table 1: Comparison among policy constraint methods.

Method	Support Constraint	Implementation
BCQ [8] PLAS [43] EMaQ [10]	Explicit Density Constrained	Intrusive Parameterization
BEAR [21] BRAC-p [40] TD3+BC [6]	Implicit Density Constrained	Pluggable Regularization
SPOT (Ours)	Explicit Density Constrained	Pluggable Regularization

Policy Constraint via Parameterization. Careful parameterization of the learned policy can naturally satisfy the support constraint. For example, BCQ [8] learns a generative model of the behavior policy and trains the actor to perturb randomly generated actions. The policy parameterized by BCQ is to greedily select the one maximizing Q function among a large number of perturbed sampled actions. EMaQ [10] simplifies BCQ by discarding the perturbation models. PLAS [43] learns a policy in the latent space of the generative model and parameterizes the policy using the decoder of the generative model on top of the latent policy.

Policy Constraint via Regularization. An alternative approach is to use a divergence penalty in order to compel the learned policy to stay close to the behavior policy, such as KL-divergence [13, 40], maximum mean discrepancy (MMD) [21], Fisher divergence [19] and Wasserstein distance [40]. Using divergence penalties alleviates the hard constraint of explicit parameterization but may bear the risk of out-of-distribution actions. Although BEAR [21] attempts to constrain the policy to the support of the behavior policy, it heavily relies on the empirically-found approximate property of

low-sampled MMD. TD3+BC [6] simply adds a behavior cloning term to the policy update and presents competitive performance on simple locomotion tasks. Note that recent SBAC [41] proposes a new policy learning objective based on performance difference lemma [15], along with a density-based regularization term similar to ours, but our work simply plugs the term into the standard policy training objective to enjoy minimal algorithmic modifications.

Behavior Policy Modeling in Offline RL. Most policy constraint methods need to fit an accurate generative model of the behavior policy, to sample in-distribution actions or estimate behavior density. Conditional variational auto-encoders (CVAE) [17, 34] are typically used by past works [8, 43] to sample actions, while EMaQ [10] opt for using an autoregressive model [9] which enables more accurate sampling. On the other hand, policy class of Gaussian [21, 40] or Gaussian mixture models [19] are commonly used to fit and estimate the density of behavior policy. Instead of explicitly fitting the behavior policy, recent approaches have utilized implicit constraint without ever querying the values of any out-of-sample actions [42, 20]. Although MBS-QL [25] uses ELBO to estimate marginal state distribution, to the best of our knowledge, we are unique in estimating the density of behavior policy based on VAE, for its flexibility to capture almost arbitrary class of distributions [18].

Broader Range of Offline RL Approaches. Besides policy constraint methods based on parameterization or regularization, there exist more types of competitive offline RL approaches. IQL [20] designs a multi-step dynamic programming procedure based on expectile regression, which completely avoids any queries to values of out-of-sample actions. Pessimistic value methods, such as CQL [22], produce a lower bound on the value of the current policy to effectively alleviate overestimation, but their performance may suffer from excessive pessimism. Advantage-weighted regression [29, 28, 38] improves upon behavior policy, while simultaneously enforcing an implicit KL-divergence constraint. Recent sequence modeling methods based on Transformers [37] also show competitive performance in both model-free [4] or model-based [12] paradigm.

3 Background

The reinforcement learning problem [35] is formulated as decision making in the environment represented by a Markov Decision Process (MDP) $\mathcal{M} = (\mathcal{S}, \mathcal{A}, \rho_0, p, r, \gamma)$, where $\mathcal{S}$ is the state space, $\mathcal{A}$ is the action space, $\rho_0(s_0)$ is the initial state distribution, $p(s'|s,a)$ is the transition distribution, $r(s,a)$ is the reward function, and γ is the discount factor. The goal in RL is to find a policy $\pi(a|s)$ maximizing the expected return: $\mathbb{E}_\pi\left[\sum_{t=0}^{\infty} \gamma^t r(s_t, a_t)\right]$.

The optimal state-action value function or Q function $Q^*(s,a)$ measures the expected return starting in state s taking action a and then acting optimally thereafter. A corresponding optimal policy can be obtained through greedy action choices $\pi^*(s) = \arg\max_a Q^*(s,a)$. The Q-learning algorithm [39] learns Q^* via iterating the Bellman optimality operator $\mathcal{T}$, defined as: $\mathcal{T}Q(s,a) = \mathbb{E}_{s'}\left[r(s,a) + \gamma \max_{a'} Q(s',a')\right]$. For large or continuous state space, the value can be represented by function approximators $Q_\theta(s,a)$ with parameters θ. In practice, the parameters θ are updated by minimizing the mean squared Bellman error with an experience replay dataset $\mathcal{D}$ and a target function $Q_{\bar{\theta}}$ [27]: $J_Q(\theta) = \mathbb{E}_{(s,a,r,s')\sim\mathcal{D}}\left[Q_\theta(s,a) - r - \gamma \max_{a'} Q_{\bar{\theta}}(s',a')\right]^2$.

In a continuous action space, the analytic maximum is intractable. Actor-Critic methods [35, 7, 11] perform action selection with a separate policy function π_ϕ maximizing the value function:

$$J_Q(\theta) = \mathbb{E}_{(s,a,r,s')\sim\mathcal{D}}\left[Q_\theta(s,a) - r - \gamma Q_{\bar{\theta}}(s', \pi_\phi(s'))\right]^2. \tag{1}$$

The policy can be updated following the deterministic policy gradient (DPG) theorem [33]:

$$J_\pi(\phi) = \mathbb{E}_{s\sim\mathcal{D}}\left[-Q_\theta(s, \pi_\phi(s))\right]. \tag{2}$$

3.1 Offline Reinforcement Learning

In contrast to online RL methods, which interact with environment to collect experience data, offline RL [23, 24] methods learn from a finite and fixed dataset $\mathcal{D} = \{(s,a,r,s')\}$ which has been collected by some unknown behavior policy π_β. Direct application of off-policy methods on offline setting suffers from *extrapolation error* [8, 21], which means that an *out-of-distribution* action a in state s can produce erroneously overestimated values $Q_\theta(s,a)$.

3.2 Support Constraints in Offline RL

To avoid *out-of-distribution* actions from function approximation, support constraint [8, 21, 10] is commonly used, which means in state s to only allow action that has ϵ-support under behavior policy: $\{a : \pi_\beta(a|s) > \epsilon\}$. Kumar *et al.* [21] first introduce the *distribution-constrained operator*, which is instantiated to the *supported operator* in this work:

Definition 3.1. Given behavior policy π_β and threshold ϵ, the *supported backup operator* is defined as

$$\mathcal{T}_\epsilon Q\left(s,a\right) = \mathbb{E}_{s'}\left[r(s,a) + \gamma \max_{a':\pi_\beta(a'|s')>\epsilon} Q\left(s',a'\right)\right] \tag{3}$$

with its fixed point $Q^*_\epsilon(s,a)$ named as the *supported optimal Q function*.

Following the theoretical analysis of Kumar *et al.* [21], we can obtain a bound of how suboptimal Q^*_ϵ may be with respect to the optimal Q function Q^*:

Corollary 3.2. *Let* $\alpha(\epsilon) = \|\mathcal{T}Q^* - \mathcal{T}_\epsilon Q^*\|_\infty = \max_{s,a} |\mathcal{T}Q^*(s,a) - \mathcal{T}_\epsilon Q^*(s,a)|$. *The suboptimality of* $Q^*_\epsilon(s,a)$ *can be upper-bounded as* $\|Q^* - Q^*_\epsilon\|_\infty \leq \frac{1}{1-\gamma}\alpha(\epsilon)$.

Note that the complete theoretical results from Kumar *et al.* [21] also includes a term w.r.t. the bootstrapping error. We refer readers to Kumar *et al.* [21] for more details. Similar results are also presented by Ghasemipour *et al.* [10].

4 Supported Policy Optimization

As aforementioned, performing support constraint is the typical method to mitigate extrapolation error in offline RL. Noticing that the support constraint can be formalized based on the density of behavior policy, we propose the **S**upported **P**olicy **O**p**T**imization (SPOT) as a regularization method from the new perspective of explicit density estimation. Concretely, SPOT involves a new regularization term, which is directly derived from the theoretical formalization of the support constraint. Besides, a conditional VAE is adopted to explicitly estimate the behavior density in the regularization term. Plugged into off-policy RL algorithms, we will finally arrive at the practical algorithm of SPOT.

4.1 Support Constraint via Behavior Density

Similar to how the optimal policy can be extracted from the optimal Q function, the *supported optimal policy* can also be recovered via greedy selection: $\pi^*_\epsilon(s) = \arg\max_{a:\pi_\beta(a|s)>\epsilon} Q^*_\epsilon\left(s,a\right)$. For the case of function approximation, it corresponds to a constrained policy optimization problem.

While prior works use specific parameterization of π [8, 10] or divergence penalty [21, 40] to perform support constraint, we propose to directly use behavior density $\pi_\beta(\cdot|s)$ as constraint:

$$\begin{aligned} &\max_\phi \ \mathbb{E}_{s\sim\mathcal{D}}\left[Q_\theta(s,\pi_\phi(s))\right] \\ &\text{s.t.} \ \min_s \ \log \pi_\beta(\pi_\phi(s)|s) > \hat{\epsilon}, \end{aligned} \tag{4}$$

where $\hat{\epsilon} = \log\epsilon$ for notational simplicity. Constraint via behavior density is simple and straightforward in the context of support constraint. We adopt log-likelihood instead of raw likelihood because of its mathematical convenience.

This problem imposes a constraint that the density of behavior policy is lower-bounded at every point in the state space, which is impractical to solve due to the large even infinite number of constraints. Following previous works from both online RL [32] and offline RL [21, 29] w.r.t. constrained policy optimization, we instead use a heuristic approximation that considers the average behavior density:

$$\begin{aligned} &\max_\phi \ \mathbb{E}_{s\sim\mathcal{D}}\left[Q_\theta(s,\pi_\phi(s))\right] \\ &\text{s.t.} \ \mathbb{E}_{s\sim\mathcal{D}}\left[\log \pi_\beta(\pi_\phi(s)|s)\right] > \hat{\epsilon}. \end{aligned} \tag{5}$$

Converting the constrained optimization problem into an unconstrained one by treating the constraint term as a penalty, we finally get the policy learning objective as

$$J_\pi(\phi) = \mathbb{E}_{s\sim\mathcal{D}}\left[-Q_\theta\left(s,\pi_\phi(s)\right) - \lambda \log \pi_\beta\left(\pi_\phi(s)|s\right)\right], \tag{6}$$

where λ is a Lagrangian multiplier.

4.2 Explicit Estimation of Behavior Density

The straightforward regularization term in Eq. (6) requires access to π_β. While we only have offline data generated by π_β, we can explicitly estimate the probability density at an arbitrary point with the density estimation methods [1].

The variational autoencoder (VAE) [17] is among the best performing neural density-estimation models and we opt to use a conditional variantional autoencoder [34] as our density estimator. Typically, $\pi_\beta(a|s)$ can be approximated by a Deep Latent Variable Model $p_\psi(a|s) = \int p_\psi(a|z,s)p(z|s)\mathrm{d}z$ with a fixed prior $p(z|s) = \mathcal{N}(\mathbf{0}, I)$. While the marginal likelihood $p_\psi(a|s)$ is intractable, VAE additionally uses an approximate posterior $q_\varphi(z|a,s) \approx p_\psi(z|a,s)$ and parameters ψ and φ can be optimized jointly with evidence lower bound (ELBO):

$$\begin{aligned}\log p_\psi(a|s) &\geq \mathbb{E}_{q_\varphi(z|a,s)}\left[\log \frac{p_\psi(a,z|s)}{q_\varphi(z|a,s)}\right] \\ &= \mathbb{E}_{q_\varphi(z|a,s)}\left[\log p_\psi(a|z,s)\right] - \mathrm{KL}\left[q_\varphi(z|a,s)\|p(z|s)\right] \\ &\stackrel{\text{def}}{=} -\mathcal{L}_{\mathrm{ELBO}}(s,a;\varphi,\psi).\end{aligned} \tag{7}$$

After training a VAE, we can simply use $-\mathcal{L}_{\mathrm{ELBO}}$ to lower-bound $\log p_\psi(a|s)$ and thus approximately lower-bound $\log \pi_\beta$ in Eq. (6). However, there theoretically exists a bias between them, as we know $\log p_\psi(a|s) = -\mathcal{L}_{\mathrm{ELBO}} + \mathrm{KL}(q_\varphi(z|a,s)||p_\psi(z|a,s))$. To obtain an estimation with lower bias, we can use the importance sampling technique [31, 18]:

$$\begin{aligned}\log p_\psi(a|s) &= \log \mathbb{E}_{q_\varphi(z|a,s)}\left[\frac{p_\psi(a,z|s)}{q_\varphi(z|a,s)}\right] \\ &\approx \mathbb{E}_{z^{(l)}\sim q_\varphi(z|a,s)}\left[\log \frac{1}{L}\sum_{l=1}^{L}\frac{p_\psi(a,z^{(l)}|s)}{q_\varphi(z^{(l)}|a,s)}\right] \\ &\stackrel{\text{def}}{=} \widehat{\log \pi_\beta}(a|s;\varphi,\psi,L).\end{aligned} \tag{8}$$

Burda *et al.* [3] show that $\widehat{\log \pi_\beta}(a|s;\varphi,\psi,L)$ gives a lower bound of $\log p_\psi(a|s)$ and the bound becomes tighter as L increases. Note that here we adopt sampling of VAE to directly estimate the density of the behavior policy instead of to estimate the divergence [21].

In summary, the loss function in Eq. (6) can be implemented with the explicit density estimator as follows:

$$J_\pi(\phi) = \mathbb{E}_{s\sim\mathcal{D}}\left[-Q_\theta\left(s,\pi_\phi(s)\right) - \lambda\widehat{\log \pi_\beta}(\pi_\phi(s)|s;\varphi,\psi,L)\right]. \tag{9}$$

4.3 Practical Algorithm

The general framework derived above can be built on top of off-policy algorithms with minimal modifications. We choose TD3 [7] as our base algorithm, which recently shows strong resistance to overestimation in offline RL [6] (see Section 5.2 for a detailed discussion).

Q normalization. Following TD3+BC [6], we add a normalization term to policy loss as a default option for better balance between Q value objective and regularization: $J_\pi(\phi) = \mathbb{E}_{s\sim\mathcal{D}}\left[\frac{-Q_\theta(s,\pi_\phi(s))}{\alpha} - \lambda\widehat{\log \pi_\beta}(\pi_\phi(s)|s;\varphi,\psi,L)\right]$, where $\alpha = \frac{1}{N}\sum_{s_i}\left|Q\left(s_i,\pi_\phi(s_i)\right)\right|$ is the normalization term based on the minibatch $\{s_i\}_{i=1}^N$ with size N.

Simpler density estimator. While $\widehat{\log \pi_\beta}(a|s;\varphi,\psi,L)$ with large L is much tighter, we empirically find there is no further improvement with larger L compared to $L=1$ (see supplementary materials for results of ablation study). To make our algorithm simpler, we choose to only use $L=1$ for a practical estimator, which is just the ELBO estimator of the VAE: $\widehat{\log \pi_\beta}(a|s;\varphi,\psi,L=1) = -\mathcal{L}_{\mathrm{ELBO}}(s,a;\varphi,\psi)$. Note that with $L=1$, we can analytically separate out the KL divergence as Eq. (7) to enjoy a simpler and lower-variance update.

Overall algorithm. Putting everything together, the full algorithm is summarized in Algorithm 1. Our algorithm first trains VAE using $\mathcal{L}_{\mathrm{ELBO}}(s,a;\varphi,\psi)$ to obtain a density estimator with sufficient accuracy. Then it turns to policy training analogous to common Actor-Critic methods except that we plug the regularization term computed by the density estimator into the policy loss $J_\pi(\phi)$.

Algorithm 1 Supported Policy Optimization (SPOT)

Input: Dataset $\mathcal{D} = \{(s, a, r, s')\}$
// VAE Training
Initialize VAE with parameters ψ and φ
for $t = 1$ **to** T_1 **do**
  Sample minibatch of transitions $(s, a) \sim \mathcal{D}$
  Update ψ, φ minimizing $\mathcal{L}_{\text{ELBO}}(s, a; \varphi, \psi)$ in Eq. (7)
end for
// Policy Training
Initialize the policy network π_ϕ, critic network Q_θ and target network $Q_{\bar{\theta}}$ with $\bar{\theta} \leftarrow \theta$
for $t = 1$ **to** T_2 **do**
  Sample minibatch of transitions $(s, a, r, s') \sim \mathcal{D}$
  Update θ minimizing $J_Q(\theta)$ in Eq. (1)
  Update ϕ minimizing $J_\pi(\phi)$ in Eq. (9)
  Update target network: $\bar{\theta} \leftarrow \tau\theta + (1 - \tau)\bar{\theta}$
end for

5 Experimental Evaluation

Our experiments aim to evaluate our method comparatively, in contrast to prior offline RL methods, focusing on both offline training and online fine-tuning. We first demonstrate the effect of λ on applying support constraint and show that our method is able to learn a policy with the strongest performance at the same level of constraint strength, compared to previous policy constraint methods. We then evaluate SPOT on D4RL benchmark [5], studying how effective our method is in contrast to a broader range of state-of-the-art offline RL methods. Finally, we study how SPOT compares to prior methods when fine-tuning with online RL from an offline RL initialization, and investigate the computational efficiency of different methods. See supplemental materials for the code.

5.1 Analysis of Support Constraint in SPOT

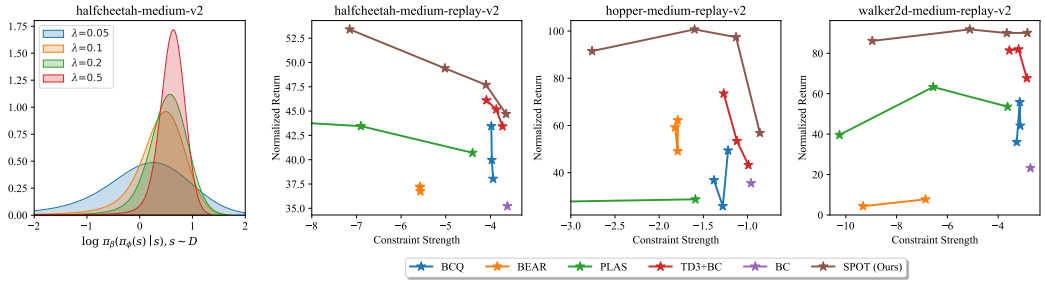

(a) Effect of λ.

(b) Tradeoff between constraint strength and optimality.

Figure 1: Analysis on constraint strength: (a) With varying values of the coefficient λ, SPOT applies support constraint with different strength, which is demonstrated by the behavior density of actions taken by the learned policy: $\log \pi_\beta(\pi_\phi(s)|s), s \sim \mathcal{D}$. (b) When evaluating the performance at varying levels of constraint strength, SPOT takes shape of the "upper envelope" of all methods, showing that SPOT can always achieve the strongest performance among extensive policy constraint methods. Constraint strength is captured approximately by the 5th-percentile of the distribution $\log \pi_\beta(\pi_\phi(s)|s), s \sim \mathcal{D}$. See supplementary materials for extended results.

Effect of λ on constraint strength. The coefficient λ in SPOT is essential and corresponds to a specific constraint strength in the constrained policy optimization problem formalized in Eq. (4). To illustrate how λ effects the learned policy, we evaluate behavior density of actions taken by the policy learned with varying values of $\lambda \in [0.05, 0.1, 0.2, 0.5]$ on standard D4RL [5] Gym-Mujoco domains. Concretely, we plot the distribution of behavior density $\log \pi_\beta(\pi_\phi(s)|s), s \sim \mathcal{D}$ in Figure 1a, where $\log \pi_\beta$ is estimated by our learned density estimator (Eq. (8)) with L set to a sufficient large number 500 for more accurate estimation. As we show, with smaller λ, the learned policy is much more possible to perform actions with low behavior density $\log \pi_\beta(\pi_\phi(s)|s)$. On the other hand, policies learned by higher λ are restricted to take only high-density actions.

Tradeoff between constraint strength and optimality. It has been shown by Kumar *et al.* [21] that the optimality of approximate supported optimal policy is lower-bounded by a tradeoff between keeping the learned policy supported by the behavior policy (controlling extrapolation error) and

Table 2: Performance of SPOT and prior methods on Gym-MuJoCo tasks. m = medium, m-r = medium-replay, m-e = medium-expert. For baselines, we report numbers directly from the IQL paper [20], which provides an unified comparison for "-v2" datasets. For SPOT, we report the mean and standard deviation for 10 seeds.

	BC	AWAC	DT	Onestep	TD3+BC	CQL	IQL	SPOT (Ours)
HalfCheetah-m-e-v2	55.2	42.8	86.8	**93.4**	90.7	91.6	86.7	86.9±4.3
Hopper-m-e-v2	52.5	55.8	**107.6**	103.3	98.0	105.4	91.5	99.3±7.1
Walker-m-e-v2	107.5	74.5	108.1	**113.0**	110.1	108.8	109.6	112.0±0.5
HalfCheetah-m-v2	42.6	43.5	42.6	48.4	48.3	44.0	47.4	**58.4**±1.0
Hopper-m-v2	52.9	57.0	67.6	59.6	59.3	58.5	66.2	**86.0**±8.7
Walker-m-v2	75.3	72.4	74.0	81.8	83.7	72.5	78.3	**86.4**±2.7
HalfCheetah-m-r-v2	36.6	40.5	36.6	38.1	44.6	45.5	44.2	**52.2**±1.2
Hopper-m-r-v2	18.1	37.2	82.7	97.5	60.9	95.0	94.7	**100.2**±1.9
Walker-m-r-v2	26.0	27.0	66.6	49.5	81.8	77.2	73.8	**91.6**±2.8
Gym-MuJoCo total	466.7	450.7	672.6	684.6	677.4	698.5	692.6	**773.0**±30.2

keeping the supported policy set large enough to capture well-performing policies. If the constraint in Eq. (4) is strong (by a large $\log \epsilon$), the extrapolation error is restrained to be small but the optimal policy under constraint may have poor performance. Otherwise, if the constraint is weak, well-performing policies can be learned though at the risk of the extrapolation error.

We aim to answer the question that *at the same level of constraint being satisfied, is SPOT able to learn a policy with the strongest performance compared to previous policy constraint methods?* We compare SPOT with BC [30], BCQ [8], BEAR [21], PLAS [43] and TD3+BC [6]. Hyperparameters to control constraint strength of various methods are adjusted to several values to form a spectrum of constraint strength (see supplementary materials for details). We approximate the satisfied constraint strength $\log \epsilon$ by the 5th-percentile of the distribution $\log \pi_\beta(\pi_\phi(s)|s), s \sim \mathcal{D}$. As shown in Figure 1b, our method SPOT takes shape of the "upper envelope" of all methods, demonstrating that taking advantage of exact standard formulation of support constraint, SPOT is flexible to learn the supported optimal policy and resist extrapolation error at the same time. Note that the divergence-based regularization method BEAR [21] yields a poor performance in our experiments and only satisfies a loose constraint in contrast to other baselines, showing that it cannot prevent out-of-distribution actions effectively and suffers from extrapolation error with indirect divergence regularization.

5.2 Comparisons on Offline RL Benchmarks

Next, we evaluate our approach on the D4RL benchmark [5] in comparison to state-of-the-art methods. We focus on Gym-MuJoCo locomotion domains and much more challenging AntMaze domains, which consists of sparse-reward tasks and requires "stitching" fragments of suboptimal trajectories traveling undirectedly in order to find a path from the start to the goal of the maze.

Baselines. We select the classic BC [30] and state-of-the-art offline RL methods as baselines. For methods based on dynamic programming, we compare to AWAC [28], Onestep RL [2], TD3+BC [6], CQL [22], and IQL [20]. We also include the sequence-modeling method Decision Transformer [4].

Hyperparameter tuning. The weight λ in Eq. (9) is essential for policy constraint. Following prior works [2, 21, 40], we allow access to the environment to tune a small (< 10) set of the hyperparameter λ, which is reasonable for practical applications. See supplementary materials for more discussions.

Gym-MuJoCo domains. Results for the Gym-MuJoCo domains are shown in Table 2. As we show, SPOT substantially outperforms state-of-the-art methods, especially in suboptimal "medium" and "medium-replay" datasets with a large margin, which further demonstrates the advantages of direct constraint on behavior density proposed by SPOT.

AntMaze domains. Results for the AntMaze domains are shown in Table 3. Note that D4RL recently releases a bug-fixed "-v2" version of AntMaze datasets, and thus we select competitive baselines and rerun their author-provided implementations for comparison. We also include results for BCQ and BEAR trained on "-v0" datasets, directly from [5]. It is not suitable to compare them with reproduced baselines but we include them in order to highlight that previous policy constraint methods struggle to succeed in training on challenging AntMaze domains.

As shown in Table 3, SPOT performs slightly worse than IQL but outperforms remaining modern offline RL baselines, including the pessimistic value method CQL and the sequence modeling method

Table 3: Performance of SPOT and prior methods on AntMaze tasks. For baselines, we obtain the results using author-provided implementations on "-v2" datasets. For BCQ and BEAR, we report numbers from D4RL whitepaper [5]. For SPOT, we report the mean and standard deviation for 10 seeds.

	BCQ	BEAR	BC	DT	TD3+BC	CQL	IQL	SPOT (Ours)
Umaze-v2	78.9	73.0	49.2±9.0	54.2±4.1	73.0±34.0	82.6±5.7	89.6±4.2	**93.5**±2.4
Umaze-diverse-v2	55.0	61.0	41.8±4.3	41.2±11.4	47.0±7.3	10.2±6.7	**65.6**±8.3	40.7±5.1
Medium-play-v2	0.0	0.0	0.4±0.5	0.0±0.0	0.0±0.0	59.0±1.6	**76.4**±2.7	74.7±4.6
Medium-diverse-v2	0.0	8.0	0.2±0.4	0.0±0.0	0.2±0.4	46.6±24.0	72.8±7.0	**79.1**±5.6
Large-play-v2	6.7	0.0	0.0±0.0	0.0±0.0	0.0±0.0	16.4±17.1	**42.0**±3.8	35.3±8.3
Large-diverse-v2	2.2	0.0	0.0±0.0	0.0±0.0	0.0±0.0	3.2±4.1	**46.0**±4.5	36.3±13.7
AntMaze total	142.8	142.0	91.6±14.1	95.4±15.5	120.2±41.7	218.0±59.4	**392.4**±30.5	359.6±39.7

Decision Transformer. Note that IQL ingeniously designs multi-step dynamic programming and policy extraction steps to apply an implicit constraint for offline RL, but when online fine-tuned after offline RL initialization, IQL is inferior to SPOT. SPOT's pluggable design can take full advantage of existing online RL algorithms (see Section 5.3 and Table 4). To the best of our knowledge, our algorithm is the first to train successfully in challenging AntMaze domains with pluggable modification on top of off-policy RL methods for offline RL.

Ablation. In Figure 2, we perform an ablation study over the components in our method. First, we replace the VAE-based density estimator with a behavioral-cloned Gaussian policy model [21, 40]. As expected, it degrades the performance due to the lack of flexibility to model complex distribution, especially on Gym-MuJoCo medium-expert datasets and AntMaze datasets. Then, we evaluate SPOT without Q normalization and find that the removal provides some benefits as well as damages on different datasets, with only an insignificant impact on total performance. Nevertheless, we include it as a default option following TD3+BC [6].

Lastly, we investigate how the base off-policy algorithm matters. Since SAC [11] has similar off-policy performance with TD3, we attempt to adopt SAC as our base method. As Xu *et al.* [41] argue that the entropy term will do harm to the offline setting, we also implement a variant of SAC without maximum entropy regularization. Surprisingly, both variants are vulnerable to pathological extrapolation error on AntMaze domains and provide poor performance (see Figure 2). We argue that several key native designs of TD3 are beneficial to offline RL. **(1).** For the case of learning a stochastic policy *without* entropy regularization, we find that the learned policy quickly degenerates into a deterministic one (with near-zero standard deviation). A concern with deterministic policies is that they are prone to overfit overestimated actions and propagate the estimation error through Bellman backups, which is even serious in offline RL. Hopefully, TD3 introduces Target Policy Smoothing [7, 36] into Eq. 1, which adds random noise to target actions and can alleviate the effect of error propagation. **(2).** For the case *with* entropy regularization, Bellman backup of a stochastic policy resembles Target Policy Smoothing, then the primary distinction may come from the actor learning objective. While TD3 produces a deterministic action minimizing Eq. (9), stochastic policies are more likely to produce out-of-distribution actions with erroneous estimated values, thus the policy gradient may become biased as well as with high variance. On much easier Gym-MuJoCo domains, SAC-style SPOT variants are comparable to the TD3 variant on most of the datasets, but we remark that TD3 indeed extends the limit of policy constraint methods on some tasks, such as hopper-medium. Our analysis suggests that TD3 may be preferable to offline RL as the base off-policy method with native

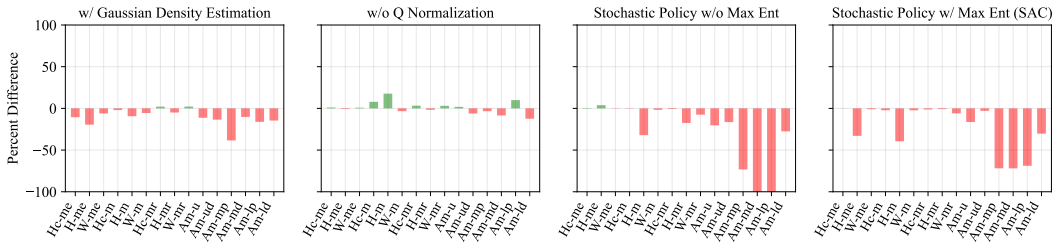

Figure 2: Percent difference of the performance of an ablation of SPOT, compared with the original algorithm. Hc = HalfCheetah, H = Hopper, W = Walker, me = medium-expert, m = medium, mr = medium-replay. Am = AntMaze, u = umaze, m = medium, l = large, d = diverse, p = play. Implementation details and quantitative results can be found in supplementary materials.

designs (such as "stochastic" critic training and deterministic actor training) addressing function approximation error not only in the online setting but also in the offline setting [6].

5.3 Online Fine-tuning after Offline RL

Pluggable SPOT can be online fine-tuned seamlessly, which means that we only need to gradually decay the constraint strength λ in the online phase in order to avoid excessive conservation. As AWAC [28] shows that behavior models are hard to update online, we fix the VAE during online fine-tuning. Note that when λ is zero, our algorithm is exactly the standard off-policy RL algorithm that SPOT builds upon. It is beneficial since we enjoy a minimal gap between offline RL and well-established online RL methods and can take full advantage of them for online fine-tuning.

Since IQL [20] is the strongest baseline in our offline experiments, which also has shown superior online performance than prior methods [28, 22] in its paper, and most of the other baselines fail to learn meaningful results, we follow the experiments of IQL and compare to IQL in online fine-tuning. We also compare to our base RL method TD3 [7] trained online from scratch. We use the challenging AntMaze domains [5]. Results are shown in Table 4. While online training from scratch fails in the challenging sparse reward tasks on AntMaze domains, SPOT initialized with offline RL succeeds to learn nearly optimal policies and performs significantly better than the strongest baseline IQL, especially in the most difficult large maze.

Table 4: Online fine-tuning results on AntMaze tasks, showing initial performance after offline RL and performance after 1M steps of online RL. All numbers are reported by the mean of 5 seeds.

	IQL	SPOT (Ours)
Umaze-v2	85.4 → 96.2	**93.2** → **99.2** (+3.0)
Umaze-diverse-v2	**70.8** → 62.2	41.6 → **96.0** (+33.8)
Medium-play-v2	68.6 → 89.8	**75.2** → **97.4** (+7.6)
Medium-diverse-v2	**73.4** → 90.2	73.0 → **96.2** (+6.0)
Large-play-v2	**40.0** → 78.6	40.8 → **89.4** (+10.8)
Large-diverse-v2	40.4 → 73.4	**44.0** → **90.8** (+17.4)
AntMaze total	378.6 → 490.4	367.8 → **569.0** (+78.6)

5.4 Computation Cost

Regularization methods, including our SPOT, benefiting from the pluggable design, only need one forward pass of the policy network to do inference, while parameterization methods always need inference through secondary components, such as the generative model or the critic network, which may bring extra time or memory cost. As demonstrated empirically in Figure 3, parameterization methods are usually slower than regularization methods. Note that BEAR also includes a greedy selection step that executes the Q-value maximizing action out of sampled actions like BCQ and EMaQ, leading to extra inference time. Pure regularization methods, including TD3+BC and our SPOT, run more than two times faster compared to the most time-consuming BCQ.

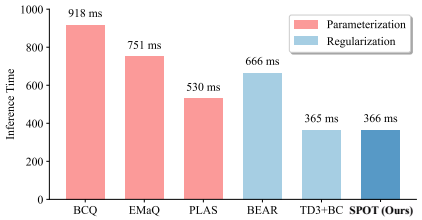

Figure 3: Runtime of various offline RL algorithms interacting with the HalfCheetah environment to produce a 1000-steps trajectory. See supplementary materials for the details.

SPOT indeed adds some training overhead due to the VAE-based density estimator, but is still much simpler than most methods. A comparison of training time is provided in supplementary materials. SPOT lies in the second tier, only slightly worse than PLAS and beaten by the simplest TD3+BC.

6 Conclusion

We present Supported Policy OpTimization (SPOT), a policy constraint method to offline RL built upon off-the-shelf off-policy RL algorithms. SPOT introduces a pluggable regularization term applied directly to the estimated behavior density, which brings a number of important benefits. First, our algorithm is computationally efficient at inference, which only needs one forward process of the policy network for action selection. Second, capturing the standard formulation of the support constraint based on behavior density, it obtains excellent performance across different tasks in the D4RL benchmarks, including standard Gym-MuJoCo tasks and much more challenging AntMaze tasks. Finally, the pluggable design of our algorithm makes it seamless to apply online fine-tuning after offline RL pre-training. Taking full advantage of well-established online methods, SPOT exceeds the state-of-the-art online fine-tuning performance on the challenging AntMaze domains.

References

[1] Christopher M. Bishop. *Pattern Recognition and Machine Learning*. Springer, 2006.

[2] David Brandfonbrener, William F Whitney, Rajesh Ranganath, and Joan Bruna. Offline rl without off-policy evaluation. In *NeurIPS*, 2021.

[3] Yuri Burda, Roger Grosse, and Ruslan Salakhutdinov. Importance weighted autoencoders. *arXiv preprint arXiv:1509.00519*, 2015.

[4] Lili Chen, Kevin Lu, Aravind Rajeswaran, Kimin Lee, Aditya Grover, Michael Laskin, Pieter Abbeel, Aravind Srinivas, and Igor Mordatch. Decision transformer: Reinforcement learning via sequence modeling. In *NeurIPS*, 2021.

[5] Justin Fu, Aviral Kumar, Ofir Nachum, George Tucker, and Sergey Levine. D4rl: Datasets for deep data-driven reinforcement learning. *arXiv preprint arXiv:2004.07219*, 2020.

[6] Scott Fujimoto and Shixiang Shane Gu. A minimalist approach to offline reinforcement learning. In *NeurIPS*, 2021.

[7] Scott Fujimoto, Herke Hoof, and David Meger. Addressing function approximation error in actor-critic methods. In *ICML*, 2018.

[8] Scott Fujimoto, David Meger, and Doina Precup. Off-policy deep reinforcement learning without exploration. In *ICML*, 2019.

[9] Mathieu Germain, Karol Gregor, Iain Murray, and Hugo Larochelle. MADE: Masked autoencoder for distribution estimation. In *ICML*, 2015.

[10] Seyed Kamyar Seyed Ghasemipour, Dale Schuurmans, and Shixiang Shane Gu. EMaQ: Expected-max q-learning operator for simple yet effective offline and online rl. In *ICML*, 2021.

[11] Tuomas Haarnoja, Aurick Zhou, Kristian Hartikainen, George Tucker, Sehoon Ha, Jie Tan, Vikash Kumar, Henry Zhu, Abhishek Gupta, Pieter Abbeel, et al. Soft actor-critic algorithms and applications. *arXiv preprint arXiv:1812.05905*, 2018.

[12] Michael Janner, Qiyang Li, and Sergey Levine. Offline reinforcement learning as one big sequence modeling problem. In *NeurIPS*, 2021.

[13] Natasha Jaques, Asma Ghandeharioun, Judy Hanwen Shen, Craig Ferguson, Agata Lapedriza, Noah Jones, Shixiang Gu, and Rosalind Picard. Way off-policy batch deep reinforcement learning of implicit human preferences in dialog. *arXiv preprint arXiv:1907.00456*, 2019.

[14] Alistair EW Johnson, Tom J Pollard, Lu Shen, H Lehman Li-Wei, Mengling Feng, Mohammad Ghassemi, Benjamin Moody, Peter Szolovits, Leo Anthony Celi, and Roger G Mark. MIMIC-III, a freely accessible critical care database. *Scientific data*, 3(1):1–9, 2016.

[15] Sham Kakade and John Langford. Approximately optimal approximate reinforcement learning. In *ICML*, 2002.

[16] Dmitry Kalashnikov, Alex Irpan, Peter Pastor, Julian Ibarz, Alexander Herzog, Eric Jang, Deirdre Quillen, Ethan Holly, Mrinal Kalakrishnan, Vincent Vanhoucke, et al. QT-Opt: Scalable deep reinforcement learning for vision-based robotic manipulation. In *CoRL*, 2018.

[17] Diederik P Kingma and Max Welling. Auto-encoding variational bayes. In *ICLR*, 2014.

[18] Diederik P Kingma and Max Welling. An introduction to variational autoencoders. *arXiv preprint arXiv:1906.02691*, 2019.

[19] Ilya Kostrikov, Rob Fergus, Jonathan Tompson, and Ofir Nachum. Offline reinforcement learning with fisher divergence critic regularization. In *ICML*, 2021.

[20] Ilya Kostrikov, Ashvin Nair, and Sergey Levine. Offline reinforcement learning with implicit q-learning. In *ICLR*, 2022.

[21] Aviral Kumar, Justin Fu, George Tucker, and Sergey Levine. Stabilizing off-policy q-learning via bootstrapping error reduction. In *NeurIPS*, 2019.

[22] Aviral Kumar, Aurick Zhou, George Tucker, and Sergey Levine. Conservative q-learning for offline reinforcement learning. In *NeurIPS*, 2020.

[23] Sascha Lange, Thomas Gabel, and Martin Riedmiller. Batch reinforcement learning. In *Reinforcement learning*, pages 45–73. Springer, 2012.

[24] Sergey Levine, Aviral Kumar, George Tucker, and Justin Fu. Offline reinforcement learning: Tutorial, review, and perspectives on open problems. *arXiv preprint arXiv:2005.01643*, 2020.

[25] Yao Liu, Adith Swaminathan, Alekh Agarwal, and Emma Brunskill. Provably good batch reinforcement learning without great exploration. In *NeurIPS*, 2020.

[26] Will Maddern, Geoffrey Pascoe, Chris Linegar, and Paul Newman. 1 year, 1000 km: The oxford robotcar dataset. *The International Journal of Robotics Research*, 36(1):3–15, 2017.

[27] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A Rusu, Joel Veness, Marc G Bellemare, Alex Graves, Martin Riedmiller, Andreas K Fidjeland, Georg Ostrovski, et al. Human-level control through deep reinforcement learning. *nature*, 518(7540):529–533, 2015.

[28] Ashvin Nair, Abhishek Gupta, Murtaza Dalal, and Sergey Levine. AWAC: Accelerating online reinforcement learning with offline datasets. *arXiv preprint arXiv:2006.09359*, 2020.

[29] Xue Bin Peng, Aviral Kumar, Grace Zhang, and Sergey Levine. Advantage-weighted regression: Simple and scalable off-policy reinforcement learning. *arXiv preprint arXiv:1910.00177*, 2019.

[30] Dean A Pomerleau. Alvinn: An autonomous land vehicle in a neural network. Technical report, Carnegie Melon Univ. Pittsburgh, PA. Artificial Intelligence and Psychology., 1989.

[31] Danilo Jimenez Rezende, Shakir Mohamed, and Daan Wierstra. Stochastic backpropagation and approximate inference in deep generative models. In *ICML*, 2014.

[32] John Schulman, Sergey Levine, Pieter Abbeel, Michael Jordan, and Philipp Moritz. Trust region policy optimization. In *ICML*, 2015.

[33] David Silver, Guy Lever, Nicolas Heess, Thomas Degris, Daan Wierstra, and Martin Riedmiller. Deterministic policy gradient algorithms. In *ICML*, 2014.

[34] Kihyuk Sohn, Honglak Lee, and Xinchen Yan. Learning structured output representation using deep conditional generative models. In *NeurIPS*, 2015.

[35] Richard S Sutton and Andrew G Barto. *Reinforcement learning: An introduction*. MIT press, 2018.

[36] Harm van Seijen, Hado van Hasselt, Shimon Whiteson, and Marco Wiering. A theoretical and empirical analysis of expected sarsa. In *Proceedings of the IEEE Symposium on Adaptive Dynamic Programming and Reinforcement Learning: ADPRL*, 2009.

[37] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *NeurIPS*, 2017.

[38] Ziyu Wang, Alexander Novikov, Konrad Zolna, Jost Tobias Springenberg, Scott Reed, Bobak Shahriari, Noah Siegel, Josh Merel, Caglar Gulcehre, Nicolas Heess, et al. Critic regularized regression. In *NeurIPS*, 2020.

[39] Christopher JCH Watkins and Peter Dayan. Q-learning. *Machine learning*, 8(3-4):279–292, 1992.

[40] Yifan Wu, George Tucker, and Ofir Nachum. Behavior regularized offline reinforcement learning. *arXiv preprint arXiv:1911.11361*, 2019.

[41] Haoran Xu, Xianyuan Zhan, Jianxiong Li, and Honglei Yin. Offline reinforcement learning with soft behavior regularization. *arXiv preprint arXiv:2110.07395*, 2021.

[42] Yiqin Yang, Xiaoteng Ma, Chenghao Li, Zewu Zheng, Qiyuan Zhang, Gao Huang, Jun Yang, and Qianchuan Zhao. Believe what you see: Implicit constraint approach for offline multi-agent reinforcement learning. In *NeurIPS*, 2021.

[43] Wenxuan Zhou, Sujay Bajracharya, and David Held. PLAS: Latent action space for offline reinforcement learning. In *CoRL*, 2020.

Checklist

1. For all authors...
   (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes]
   (b) Did you describe the limitations of your work? [Yes] See Section 4 in supplementary materials.
   (c) Did you discuss any potential negative societal impacts of your work? [Yes] See Section 4 in supplementary materials.
   (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
   (a) Did you state the full set of assumptions of all theoretical results? [Yes]
   (b) Did you include complete proofs of all theoretical results? [Yes] See Section 1 in supplementary materials.
3. If you ran experiments...
   (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [Yes] We have uploaded the code to reproduce the main results of our methods (Table 2, 3, 4 in the main text) in supplementary materials.
   (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes] See Section 3 in supplementary materials.
   (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [Yes] See Section 5 in the main text and Section 3 in supplementary materials.
   (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [Yes] See Section 5.4 in the main text and Section 3.5 in supplementary materials.
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
   (a) If your work uses existing assets, did you cite the creators? [Yes]
   (b) Did you mention the license of the assets? [Yes]
   (c) Did you include any new assets either in the supplemental material or as a URL? [Yes]
   (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [Yes] All the datasets we used are open-sourced and we comply with the license.
   (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [Yes] All the data and experiments are simulated without any personally identifiable information.
5. If you used crowdsourcing or conducted research with human subjects...
   (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]
   (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
   (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]