

Offline Safe Reinforcement Learning Using Trajectory Classification

Ze Gong, Akshat Kumar*, Pradeep Varakantham*

School of Computing and Information Systems
Singapore Management University
{zegong, akshatkumar, pradeepv}@smu.edu.sg

Abstract

Offline safe reinforcement learning (RL) has emerged as a promising approach for learning safe behaviors without engaging in risky online interactions with the environment. Most existing methods in offline safe RL rely on cost constraints at each time step (derived from global cost constraints) and this can result in either overly conservative policies or violation of safety constraints. In this paper, we propose to learn a policy that generates *desirable* trajectories and avoids *undesirable* trajectories. To be specific, we first partition the pre-collected dataset of state-action trajectories into *desirable* and *undesirable* subsets. Intuitively, the desirable set contains high reward and safe trajectories, and undesirable set contains unsafe trajectories and low-reward safe trajectories. *Second*, we learn a policy that generates desirable trajectories and avoids undesirable trajectories, where (un)desirability scores are provided by a classifier learnt from the dataset of *desirable* and *undesirable* trajectories. This approach bypasses the computational complexity and stability issues of a min-max objective that is employed in existing methods. Theoretically, we also show our approach’s strong connections to existing learning paradigms involving human feedback. *Finally*, we extensively evaluate our method using the DSRL benchmark for offline safe RL. Empirically, our method outperforms competitive baselines, achieving higher rewards and better constraint satisfaction across a wide variety of benchmark tasks.

1 Introduction

Reinforcement learning (RL) has achieved remarkable success in autonomous decision-making tasks (Mnih et al. 2015; Ibarz et al. 2021; Kiran et al. 2021). However, ensuring safety during policy training and deployment remains a significant challenge, especially for safety-critical tasks where unsafe behavior can lead to unexpected consequences. To address the problem, safe RL has been extensively studied to develop policies that satisfy safety constraints, typically formulated within the framework of constrained Markov Decision Processes (Altman 2021; Gu et al. 2022). Safe RL agents generally interact with the environment to gather information on rewards and costs, and learn

policies through constrained optimization. Nonetheless, on-line safe RL may inevitably violate safety constraints during training and deployment due to the need for environmental interaction. To mitigate this issue, offline safe RL (Liu et al. 2023a) has emerged as a promising learning paradigm, enabling policy learning using pre-collected offline datasets, thus avoiding direct interactions with the environment.

Existing offline safe RL research has explored several directions (Liu et al. 2023a), such as enforcing constrained optimization into offline RL (Xu, Zhan, and Zhu 2022), using distribution correction-based methods (Lee et al. 2022), and employing sequence modeling frameworks (Liu et al. 2023b). Though promising, several limitations remain. First, most existing methods rely on local cost constraints at each time step, which can lead to loss of sequential information related to global cost constraints at the trajectory level (Xu, Zhan, and Zhu 2022; Lee et al. 2022). This often results in either overly conservative policies with low rewards or the failure to generate a safe policy. Additionally, most existing works employ min-max optimization within the conventional RL framework (Xu, Zhan, and Zhu 2022), resulting in an intertwined training process involving reward maximization and safety compliance resulting in stability issues. Lastly, some approaches require auxiliary models to be learned alongside the policy (Lee et al. 2022), or depend on complex model architectures (Liu et al. 2023b; Zheng et al. 2024a), which contribute to high computational complexity.

To address such challenges, we propose a novel approach that solves the offline safe RL problem by utilizing a binary trajectory classifier. This classifier allows us to directly optimize a policy by leveraging both *desirable* and *undesirable* samples at the trajectory level, without relying on conventional min-max optimization or traditional RL techniques. We apply a two-phase algorithm. In the first phase, we divide the offline dataset into two subsets, one containing desirable trajectories and the other containing undesirable ones. The desirable set includes *safe trajectories* that satisfy the constraints, with each trajectory weighted according to its cumulative reward to reflect its level of desirability. The undesirable set comprises (a) all unsafe trajectories, and (b) selection of safe trajectories with low rewards. We also develop a concrete empirical methodology to construct such subsets. In the second phase, we optimize the policy within a classifier designed to distinguish between desirable and un-

*Equal advising.

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

desirable trajectories. The classifier is specifically designed to provide a high score (derived using the principle of maximum entropy RL) to desirable trajectories, and low score to undesirable ones. By solving this trajectory classification problem, our algorithm learns the policy directly, bypassing the need for min-max optimization and traditional RL methods. Additionally, we identify close connections with the Reinforcement Learning from Human Feedback (RLHF) paradigm (Hejna et al. 2024), providing theoretical support for our method and a promising direction for future research.

For evaluation, we adopt the well-established DSRL benchmark (Liu et al. 2023a), designed specifically for offline safe RL approaches. We conduct extensive experiments by comparing our method with several state-of-the-art baselines in terms of normalized reward and normalized cost across 38 continuous control tasks in three popular domains. The results suggest that our method successfully learns policies that achieve high rewards while satisfying safety constraints in most tasks, outperforming state-of-the-art baselines.

The main contributions of the work are three-fold:

- We introduce a novel approach that solves offline safe RL by utilizing a trajectory classifier, that enables direct policy learning within a standard classification framework and eliminating the need for challenging min-max optimization.
- Furthermore, to ensure stability and handling trajectory level safety, we provide a contrastive trajectory classifier that employs a novel score function built from regret based preference models.
- Extensive experiments demonstrate that our method surpasses SOTA baselines in both reward maximization and constraint satisfaction within a well-established benchmark dataset.

2 Related Work

2.1 Safe RL

Ensuring safety remains a critical challenge in RL during both policy training and deployment phases (Garcia and Fernández 2015; Achiam et al. 2017; Gu et al. 2022; Ji et al. 2023). Safe RL aims to learn a policy that maximizes the cumulative reward while satisfying the safety constraints through interactions with the environment. Various techniques have been explored to address safe RL, such as Lagrangian-based optimization (Chow et al. 2018; Tessler, Mankowitz, and Mannor 2018; Stooke, Achiam, and Abbeel 2020; Chen, Dong, and Wang 2021; Ding et al. 2020), and correction-based approaches (Zhao, He, and Liu 2021; Luo and Ma 2021). However, ensuring zero constraint violations during training and deployment remains a significant challenge.

2.2 Offline Safe RL

Offline safe RL has emerged as a new paradigm for learning safe policies using pre-collected offline datasets (Le, Voloshin, and Yue 2019; Guan et al. 2024). Existing methods include behavior cloning (BC) (Liu et al. 2023b), distri-

bution correction-based approach (Lee et al. 2022), and incorporating constraints into existing offline RL techniques (Xu, Zhan, and Zhu 2022). Recent research has also explored decision transformer-based approaches (Liu et al. 2023b) and diffusion-based methods (Zheng et al. 2024b). Despite these advancements, several limitations persist, such as performance degradation, unstable training due to min-max optimization, and computational challenges associated with complex model architectures.

3 Preliminaries

Safe RL problems are often formulated using Constrained Markov Decision Processes (CMDPs) (Altman 2021), defined by a tuple $\mathcal{M} = (\mathcal{S}, \mathcal{A}, T, r, c, \mu_0, \gamma)$ where \mathcal{S} is the state space and \mathcal{A} is the actions space; $T : \mathcal{S} \times \mathcal{A} \rightarrow \Delta(\mathcal{S})$ describes the transition dynamics; $r : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$ is the reward function; $c : \mathcal{S} \times \mathcal{A} \rightarrow [0, C_{\max}]$ is cost function where C_{\max} is the maximum cost; $\mu_0 : \mathcal{S} \rightarrow [0, 1]$ represents the initial state distribution; γ is the discount factor. Let $\pi(a|s)$ denote the probability of taking action a in state s . The state-value function for a policy π is defined as $V_r^\pi(s) = \mathbb{E}_\pi[\sum_{t=0}^{\infty} \gamma^t r(s_t, a_t) | s_0 = s]$ with discounted cumulative rewards. Similarly, the cost state-value function is defined as $V_c^\pi(s) = \mathbb{E}_\pi[\sum_{t=0}^{\infty} \gamma^t c(s_t, a_t) | s_0 = s]$ based on discounted cumulative costs.

In offline safe RL, we assume access to a pre-collected offline dataset $D = \{\tau_1, \tau_2, \dots\}$ of trajectories. Each trajectory τ contains a sequence of tuples $(s_t, a_t, s_{t+1}, r_t, c_t)$ from start till episode end, where r_t and c_t represent the reward and cost at time step t . The goal in offline safe RL is to maximize the expected cumulative rewards $V_r^\pi(s)$ while satisfying safety constraints determined by the cumulative costs V_c^π and a specified cost threshold. The formal constrained objective function is written as follows:

$$\max_{\pi} \mathbb{E}_s [V_r^\pi(s)], \text{ s.t., } \mathbb{E}_s [V_c^\pi(s)] \leq l; \mathcal{D}(\pi || \pi_\beta) \leq \epsilon \quad (1)$$

where $l \geq 0$ is the cost threshold and π_β is the underlying behavioral policy used to collect the dataset. $\mathcal{D}(\pi || \pi_\beta)$ denotes the KL divergence between the policy π and the behavioral policy π_β which is used to address distributional drift in the offline setting.

4 Our Approach

The constraint term in Equation 1 serves as a criterion for distinguishing between safe and unsafe trajectories. As shown in Figure 1, given a cost threshold, the trajectories in the offline datasets can be divided into two categories: safe and unsafe. According to the objective function in Equation 1, all unsafe trajectories are considered undesirable, while all safe trajectories are viewed as *potential* candidates for desirable samples. Additionally, the maximization term in Equation 1 reflects the level of desirability of trajectories, with those yielding higher cumulative rewards deemed more desirable by the objective function. Building on this insight, we aim to learn a policy capable of generating desirable trajectories with high probability. We propose a two-phase approach – (a) we first construct two contrastive datasets (desirable, undesirable), and (b) then we transform the offline

safe RL problem into a direct policy learning within a trajectory classifier.

4.1 Contrastive Dataset Construction

Given the pre-collected offline dataset D , we create two new subdatasets at the trajectory level: one containing desirable trajectories and the other containing undesirable ones. *Desirable* refers to trajectories with high cumulative rewards and being safe, while *undesirable* refers to those with low cumulative rewards or unsafe.

Using the predefined cost threshold l , we first split the dataset into two categories based on the cumulative cost, i.e., safe and unsafe. Within the safe trajectories, we further rank them according to cumulative rewards. The top $x\%$ of these safe trajectories are selected as desirable. Moreover, we identify the bottom $y\%$ of the safe trajectories, along with all unsafe trajectories as undesirable (x, y are hyperparameters that we show how to set empirically). To reflect the level of desirability, we assign normalized weights w ($w \in [\delta, 1]$, and $\delta = 0.5$) to the trajectories in the desirable datasets, based on their cumulative rewards:

$$w_\tau = \frac{v(\tau) - v_{\min}}{v_{\max} - v_{\min}} \cdot (1 - \delta) + \delta \quad (2)$$

where $v(\tau)$ denotes the return of the trajectory $\tau = (s_1, a_1, s_2, a_2, \dots, s_k, a_k)$ with length k . The terms v_{\max} and v_{\min} represent the empirical maximum and minimum values of trajectory returns over all samples in the original dataset. Equation 2 is used to normalize the original returns to the range of $[\delta, 1]$. Furthermore, we assign weights to the unsafe trajectories to be 1, since we aim to avoid unsafe outcomes irrespective of their return values, treating all unsafe trajectories equally. Moreover, undesirable safe trajectories within the undesirable set are weighted reversely according to their returns, meaning that safe trajectories with lower returns receive higher weights during training:

$$w_\tau = \left(1 - \frac{v(\tau) - v_{\min}}{v_{\max} - v_{\min}}\right) \cdot (1 - \delta) + \delta \quad (3)$$

4.2 Contrastive Trajectory Classification (TraC)

With the desirable and undesirable datasets, our goal is to learn a policy that generates desirable behavior with high probability and undesirable behavior with very low probability. Intuitively, we can formulate the following objective function¹:

$$\max_{\pi_\theta} \lambda_d \mathbb{E}_{\tau \sim D^d} [w_\tau f(\tau; \pi_\theta)] - \lambda_u \mathbb{E}_{\tau \sim D^u} [w_\tau f(\tau; \pi_\theta)] \quad (4)$$

where D^d represents the dataset containing all desirable trajectories, and D^u represents the dataset containing all undesirable trajectories. The function $f(\tau; \pi_\theta)$ intuitively should return high value for desirable τ under π_θ , and low value

¹In Equation 4, our objective emphasizes training a policy to favor desirable trajectories and avoid undesirable ones, initially ignoring the KL term. Later, in Equation 7, the score function integrates the reference (behavioral) policy through the KL-regularized RL framework, addressing the distribution drift in offline RL.




Figure 1: Visualization of trajectories in the reward-cost return space with a cost threshold of 40. The x-axis is the total cost and the y-axis is the total reward. Each blue dot in the figure corresponds to a trajectory in the offline dataset. Based on the cost threshold, trajectories are categorized into safe (in the green area) and unsafe (in the pink area).

for undesirable trajectories. The λ_d and λ_u are the factors that represent the relative importance of desirable and undesirable data during training, respectively. The influence of these factors may vary depending on the dataset ratios and specific tasks involved in addressing dataset imbalance.

Specifically, the function $f(\tau; \pi)$ could be modeled using Behavior Cloning (BC) as follows:

$$f(\tau; \pi) = \sum_t \log \pi(a_t | s_t)$$

However, there are several limitations of BC approach. First, it loses sequential information by optimizing the policy for each (s_t, a_t) pair individually. Additionally, using the BC formulation, minimizing the logarithmic terms in the second negative term of Equation 4 can be unstable because the logarithm has no lower bound as $\pi(a_t | s_t) \in [0, 1]$.

Our method To address above issues, we make two key contributions:

- Instead of direct log likelihood maximization, we consider a Contrastive Trajectory Classification (TraC) problem that avoids the stability issues, and also design $f(\cdot)$ that depends on entire trajectory and not just on individual (s_t, a_t) pairs.
- More importantly, to enable imitation of desirable trajectories and avoidance of undesirable trajectories, we propose $f(\cdot)$ to be sigmoid over a score function that assigns higher value to a *desirable* trajectory and lower values to *undesirable* trajectories.

First, the TraC problem formulation is defined as:

$$\max_{\pi_\theta} \lambda_d \mathbb{E}_{\tau \sim D^d} [w_\tau f(y_\tau, \tau; \pi_\theta)] + \lambda_u \mathbb{E}_{\tau \sim D^u} [w_\tau f(y_\tau, \tau; \pi_\theta)]$$

where $y \in \{0, +1\}$ is the binary desirability label, where $+1$ is for desired trajectories, and 0 is for undesired trajectories. The parameters λ_d and λ_u are determined under the constraints $\lambda_d + \lambda_u = 1$ and $\frac{\lambda_d N_d}{\lambda_u N_u} = \eta$, where N_d and N_u are the numbers of desirable and undesirable trajectories in our

constructed datasets. The parameter η is predefined to control the relative values of λ_d and λ_u ². We define $f(y_\tau, \tau; \pi_\theta)$ as the probability of predicting the binary desirability label y_τ for the corresponding trajectory τ under policy π_θ :

$$f(y_\tau, \tau; \pi_\theta) = p(Y = y_\tau | \tau, \pi_\theta) := \sigma(\psi(\tau, \pi_\theta)) \quad (5)$$

where σ represents the sigmoid function that maps the input to the range of $[0, 1]$. The function $\psi(\tau, \pi_\theta)$ computes a score for the trajectory τ under policy π_θ which is used to calculate the probability.

Trajectory score function ψ Next, we define the score function $\psi(\tau, \pi_\theta)$, and establish some of its properties. Since, we have to ensure ψ captures a preference for desirable trajectories over undesirable trajectories, we build on key insights from regret-based preference model (Knox et al. 2024) that is expressed as:

$$P[\tau^+ \succ \tau^-] = \frac{\exp \sum_{\tau^+} \gamma^t A_r^*(s_t^+, a_t^+)}{\exp \sum_{\tau^+} \gamma^t A_r^*(s_t^+, a_t^+) + \exp \sum_{\tau^-} \gamma^t A_r^*(s_t^-, a_t^-)}$$

where preference between two trajectories is denoted as $\tau^+ \succ \tau^-$; τ^+ indicates the preferred trajectory, and τ^- represents the less preferred trajectory. The optimal advantage function $A_r^*(s, a) := Q_r^*(s, a) - V_r^*(s)$ measures how much worse taking action a in state s is than acting according to optimal policy π^* , where $Q_r^*(s, a)$ denotes the optimal state-action value function and $V_r^*(s)$ denotes the optimal state value function under optimal policy π^* .

A reason for choosing this regret based preference model is that as discussed in CPL (Hejna et al. 2024), this preference learning framework shares similarities with contrastive learning approaches. The discounted sum of advantage function (i.e., $\sum_{\tau} \gamma^t A_r^*(s_t, a_t)$) can be considered as a trajectory’s score from the Noise Contrastive Estimation (NCE) (Gutmann and Hyvärinen 2010) perspective. Furthermore, according to the principle of maximum entropy (Ziebart et al. 2008; Ziebart 2010; Hejna et al. 2024), the optimal advantage function $A_r^*(s, a)$ can be equivalently represented:

$$A_r^*(s, a) = \alpha \log \frac{\pi^*(a|s)}{\pi_{\text{ref}}(a|s)} \quad (6)$$

where $\alpha (> 0)$ is a temperature parameter. The reference policy π_{ref} is introduced to regularize π^* , a common practice in KL-regularized RL (Galashov et al. 2019). Typically, standard behavior cloning on the offline dataset is used as π_{ref} . In our context, it can also be seen as a behavioral policy that mitigates the distribution drift issue in offline settings. We define the score for a trajectory in our setting as:

$$\psi(\tau, \pi^*) = \sum_{t=0}^T \gamma^t A_r^*(s_t, a_t) = \sum_{t=0}^T \gamma^t \alpha \log \frac{\pi^*(a_t|s_t)}{\pi_{\text{ref}}(a_t|s_t)} \quad (7)$$

Score function justification We next justify why having a high score for a trajectory τ also implies it is more likely

²The formulation of the constraints for defining λ_d and λ_u is motivated by the work of Kahneman-Tversky Optimization (KTO) (Ethayarajh et al. 2024), and we adapt it to suit our setting.

under π^* . For simplicity, we ignore α . From Equation 6:

$$\log \pi^*(a|s) = A_r^*(s, a) + \log \pi_{\text{ref}}(a|s) \quad (8)$$

$$\log p(\tau; \pi^*) \propto \sum_{t=0}^T \log \pi^*(a_t|s_t) \quad (9)$$

$$= \sum_{t=0}^T [A_r^*(s_t, a_t) + \log \pi_{\text{ref}}(a_t|s_t)] \quad (10)$$

where $p(\tau; \pi^*)$ is the probability of trajectory under π^* . The ignored terms in Equation 9 are the log probabilities of the transition function, which are independent of the policy. Thus, the log probability of observing τ under π^* is directly proportional to its score function $\sum_{t=0}^T A_r^*(s_t, a_t)$ (in practice, to reduce variance, we use γ^t to discount $A_r^*(s_t, a_t)$). This implies that assigning higher score (in Equation 7) to a trajectory makes it more likely (which should be the case for desirable trajectories).

Overall loss function Using the score function in Equation 7, we formulate a new loss function that allows us to directly optimize the parameterized policy π_θ to align with the classification objective as follows:

$$L(\pi_\theta, D) = -\mathbb{E}_{\tau \sim D} \left[\lambda_d w_\tau \cdot y_\tau \log(\sigma(\psi(\tau, \pi_\theta))) + \lambda_u w_\tau \cdot (1 - y_\tau) \log(1 - \sigma(\psi(\tau, \pi_\theta))) \right] \quad (11)$$

In the first term of Equation 11, the objective is to learn a policy that assigns high values to desirable trajectories. This means the learning policy should be able to closely align with safe and high-reward trajectories, where desirability is determined by the discounted cumulative rewards of each safe trajectory. The second term focuses on ensuring the policy assigns low values to undesirable trajectories, enabling it to effectively avoid unsafe behaviors. Additionally, incorporating π_{ref} in each term helps mitigate the distribution drift problem in offline setting and can be pretrained using Behavior Cloning (BC) across the entire dataset. Compared to Equation 1, our loss function implicitly captures the objective, the constraints, and behavior regularization, which are essential components in the conventional offline safe RL objective. Our objective also does not require challenging min-max optimization, and only a single neural net to parameterize policy π is used.

5 Evaluation

In this section, we present experiments designed to evaluate the performance of TraC in maximizing the cumulative rewards while satisfying safety constraints. We specifically address the following questions:

- Can TraC effectively improve policy performance and strengthen safety constraints compared to state-of-the-art offline safe RL approaches?
- How do different compositions of desirable and undesirable datasets affect the performance of TraC?
- What ingredients of TraC are important for achieving high performance while adhering to safety constraints?

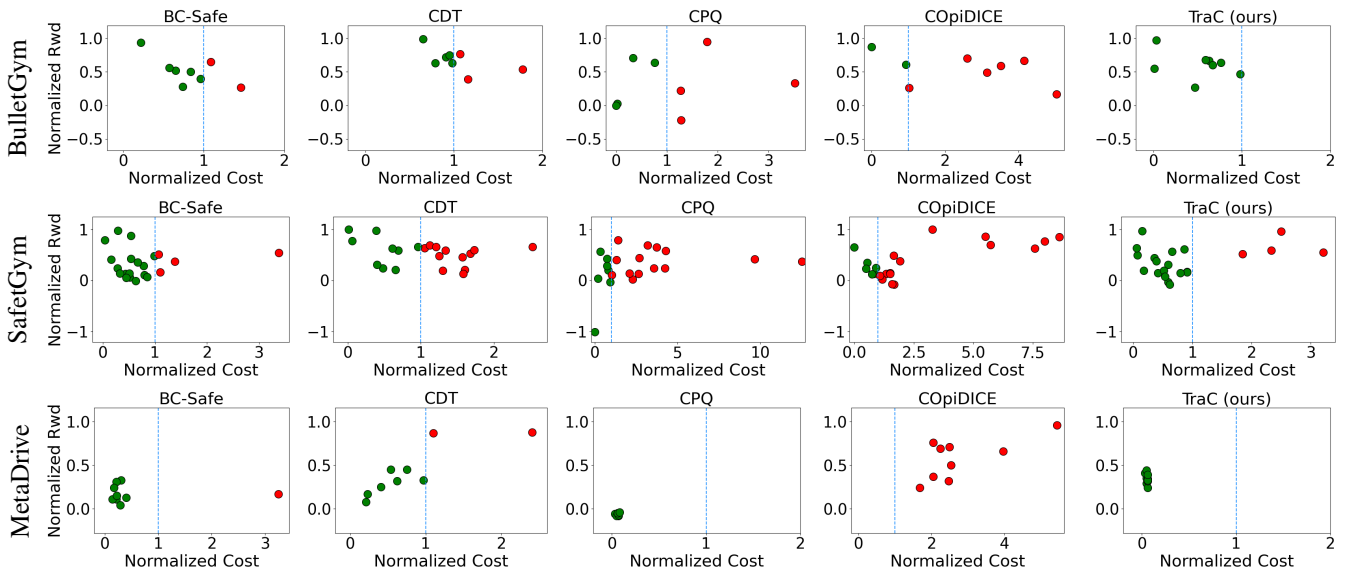


Figure 2: Visualization of normalized reward and cost for each task. The dotted blue vertical lines represent the cost threshold of 1. Each round dot represents a task, with green indicating safety and red indicating constraint violation.

5.1 Experimental Setup

We conduct experiments using the well-established DSRL benchmark (Liu et al. 2023a), specifically designed to evaluate offline safety RL approaches. The offline datasets were precollected for 38 tasks across three widely recognized environments, SafetyGymnasium (Ray, Achiam, and Amodei 2019; Ji et al. 2023), BulletSafetyGym (Gronauer 2022) and MetaDrive (Li et al. 2022). To assess performance, we adopt the constraint variation evaluation (Liu et al. 2023a) which is introduced in DSRL to evaluate algorithm versatility. Each algorithm is tested on each dataset using three distinct cost thresholds and three random seeds to ensure a fair comparison. We use normalized reward and normalized cost as evaluation metrics (Liu et al. 2023a; Fu et al. 2020), where a normalized cost below 1 indicates safety. For the practical implementation of TraC, we first pretrain the policy using behavior cloning (BC) with the offline dataset, which we then maintain as the reference policy π_{ref} . We subsequently learn a policy using our method on the newly constructed desirable and undesirable datasets.

5.2 Baselines

We compare TraC with several state-of-the-art offline safe RL baselines to demonstrate the effectiveness of our approach: 1) BC-All: Behavior cloning trained on the entire datasets. 2) BC-Safe: Behavior cloning trained exclusively on safe trajectories that satisfy the safety constraints. 3) CDT (Liu et al. 2023b): A sequence modeling approach that incorporates safety constraints into Decision Transformer architecture. 4) BCQ-Lag: A Lagrangian-based method that incorporates safety constraints into BCQ (Fujimoto, Meger, and Precup 2019). 5) BEAR-Lag: A Lagrangian-based method that incorporates safety constraints into BEAR (Kumar et al. 2019). 6) CPQ (Xu, Zhan, and Zhu 2022): a mod-




Figure 3: Ratio of tasks solved regarding safety.

ified Q-learning approach that penalizes unsafe actions by treating them as out-of-distribution actions. 7) COptiDICE (Lee et al. 2022): A DICE (distribution correction estimation) based safe offline RL method built on OptiDICE (Lee et al. 2021).

5.3 Results

How does TraC perform against SOTA offline safe RL baselines? We conducted extensive experiments using DSRL and compared TraC against baselines. The main results are presented in Table 1, which shows the average performance of all approaches in terms of normalized reward and normalized cost across 38 tasks in three environments. TraC demonstrates satisfactory safety performance and high rewards in all environments, achieving the highest rewards in SafetyGym and BulletGym. In contrast, other baselines either suffer from low rewards or fail to meet safety constraints. On average, BC-All, BCQ-Lag, COpiDICE fail to learn safe policies in any of the three environments. While CDT, BEAR-Lag, and CPQ satisfy safety

Benchmark	BC-All		BC-Safe		CDT		BCQ-Lag		BEAR-Lag		CPQ		COptiDICE		TraC (ours)	
	reward \uparrow	cost \downarrow	reward \uparrow	cost \downarrow	reward \uparrow	cost \downarrow	reward \uparrow	cost \downarrow	reward \uparrow	cost \downarrow	reward \uparrow	cost \downarrow	reward \uparrow	cost \downarrow	reward \uparrow	cost \downarrow
SafetyGym	0.46	3.03	0.34	0.75	0.54	1.06	0.5	3.29	0.39	2.7	0.27	2.79	0.37	2.65	0.4	0.92
BulletGym	0.64	3.36	0.52	0.82	0.68	1.04	0.74	3.11	0.48	3.8	0.33	1.12	0.55	2.55	0.61	0.52
MetaDrive	0.42	2.06	0.18	0.58	0.42	0.8	0.54	2.05	0.02	0.39	-0.06	0.06	0.58	2.77	0.35	0.05

Table 1: Results of normalized reward and cost averaged over tasks in each environment. \uparrow means the higher the better. \downarrow means the lower the better. Each value is averaged over 3 distinct cost thresholds, 20 evaluation episodes, and 3 random seeds. **Bold**: Safe agents whose normalized cost is smaller than 1. **Blue**: Safe agents with the highest reward.




Figure 4: Results of normalized reward and normalized cost with various $x\%$ in two task.

constraints in MetaDrive, they fail in the other two environments. BC-Safe also learns safe policies across all environments but results in conservative policies with low rewards. Additionally, although CDT can achieve high rewards across all three environments, it does so at the cost of violating safety constraints, further confirmed in Figure 3 also.

Regarding safety constraint satisfaction, we present the proportion of tasks solved safely by each approach in Figure 3. TraC successfully learned the most safe policies for tasks across all three environments, especially achieving safety in all tasks within BulletGym and MetaDrive. While some baselines, such as BC-Safe and CDT, demonstrate relatively low average normalized costs (as noted in Table 1), but they still fail to learn policies that meet the required cost thresholds for several tasks as shown in Figure 3.

Additionally, we illustrate the performance of each approach by visualizing their results in the normalized reward-cost space for each task in Figure 2. For some baselines, such as CDT, CPQ, and COptiDICE, high average rewards, as shown in Table 1, do not necessarily indicate good performance. This is because these rewards often stem from unsafe policies across various tasks. In particular, while CDT achieves the highest average reward in MetaDrive, this high average is due to two unsafe tasks that produce high rewards.

How do different compositions of desirable and undesirable datasets affect the performance of TraC? As a key component of TraC, the trajectory construction phase is expected to yield desirable and undesirable datasets that enable the learning of policies with high rewards while satisfying safety constraints. As outlined in Section 4.1, we select the top $x\%$ of safe trajectories as desirable, and the bottom $y\%$ as undesirable, based on their rewards. We conducted experi-




Figure 5: Results of normalized reward and normalized cost with various $y\%$ in two tasks.

ments with various selections of $x\%$ and $y\%$ to examine how different compositions influence the performance of TraC.

In Figure 4, we present the results of varying $x\%$ while keeping $y\%$ fixed in two tasks. On the x-axis, the notation $+x\% - y\%$ represents the experimental setting. In AntCircle task, the normalized reward remains consistent across difference $x\%$ selections, while the normalized cost decreases as $x\%$ increases. In MetaDrive-mediummean task, remains relatively unchanged with varying $x\%$. The highest normalized reward is observed when $x\%$ is set to 75%, although the differences are minor. Additionally, we conducted experiments with varying $y\%$ while keeping $x\%$ fixed, as shown in Figure 5. Similarly, altering $y\%$ results in only minor changes in both test tasks. These findings suggest that TraC is robust and exhibits little sensitivity to different selections of $x\%$ and $y\%$.

As we partition the original offline dataset into two subsets, desirable and undesirable, and train a policy using both, the goal is to encourage behavior that aligns more closely with desirable trajectories while avoiding undesirable behaviors. To verify the effectiveness of using both the subsets for training, we also conducted experiments using only desirable or only undesirable trajectories. Figure 6 presents the ablation results across three tasks. Training with only undesirable trajectories typically results in a policy with very low rewards and, in some cases, very high costs, as seen in the AntCircle task. This outcome occurs because, with only undesirable trajectories, TraC attempts to learn a policy that avoids undesirable behavior, but lacks information about what desirable behavior looks like. Conversely, when training with only desirable trajectories, our algorithm learns a policy with high rewards and relatively low costs, as the agent is expected to imitate the desirable behaviors in those

Task	$\delta = 0$		$\delta = 0.4$		$\delta = 0.7$		$\delta = 1.0$		$\eta = 0.25$		$\eta = 0.5$		$\eta = 1.0$		$\eta = 2.0$	
	reward \uparrow	cost \downarrow	reward \uparrow	cost \downarrow	reward \uparrow	cost \downarrow	reward \uparrow	cost \downarrow	reward \uparrow	cost \downarrow	reward \uparrow	cost \downarrow	reward \uparrow	cost \downarrow	reward \uparrow	cost \downarrow
AntCircle	0.46	0.87	0.47	0.95	0.47	0.98	0.46	0.9	0.47	0.98	0.47	0.98	0.46	1.18	0.46	0.92
CarPush1	0.18	0.18	0.2	0.14	0.2	0.22	0.14	0.2	0.21	0.33	0.2	0.22	0.16	0.29	0.14	0.17
mediummean	0.29	0.06	0.28	0.06	0.31	0.05	0.3	0.06	0.31	0.05	0.31	0.05	0.3	0.06	0.24	0.06

Table 2: Ablation study of varying values of δ and η in three tasks.

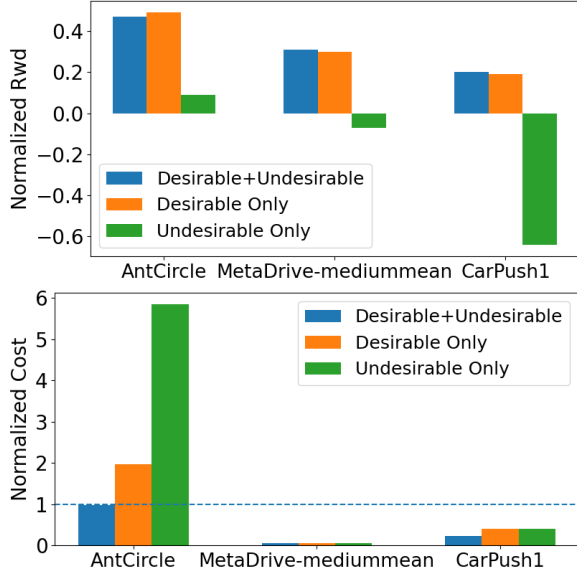


Figure 6: Ablation on training with only desirable or only undesirable trajectories.

trajectories. However, solely learning from desirable trajectories limits the agent from imitating behaviors in the desirable dataset, potentially overlooking how to avoid undesirable actions. As shown in Figure 6, training with both desirable and undesirable trajectories yields the best performance, as it allows the agent to learn both how to behave desirably and how to avoid undesirable behaviors.

5.4 Ablation Study

We evaluate the sensitivity of hyperparameter selections of δ and η in TraC to demonstrate its robustness and effectiveness. We tested four different values for each hyperparameter, and the results are shown in Table 2. The results indicate that TraC exhibits little sensitivity to the choice of δ , maintaining robust performance across different selections. For η , we observe that the normalized reward decreases as η increases while the normalized cost remains stable. That is because, a higher η , reduces the weights on desirable trajectories during training, potentially leading to a policy with lower rewards that does not strictly imitate desirable behaviors but instead explores how to avoid undesirable ones.

Furthermore, in the loss function of TraC, the reference policy π_{ref} which is pretrained on the entire dataset using behavior cloning, serves as an approximation of the underlying behavioral policy of the offline dataset. Incorporating π_{ref} helps mitigate the distributional drift issues in offline set-

Task	w/o π_{ref} pretraining		TraC	
	reward \uparrow	cost \downarrow	reward \uparrow	cost \downarrow
AntCircle	0.0	0.0	0.47	0.98
AntRun	0.11	0.1	0.67	0.63
BallCircle	0.02	6.68	0.68	0.59
BallRun	0.26	0.15	0.27	0.47
CarCircle	-0.01	3.81	0.64	0.76
CarRun	0.66	2.62	0.97	0.03
DroneCircle	-0.27	1.79	0.6	0.67
DroneRun	0.16	6.53	0.55	0.01
Average	0.12	2.71	0.61	0.52

Table 3: Ablation on π_{ref} pretraining for all task in BulletGym.

tings. To demonstrate the effectiveness of reference policy π_{ref} , we conduct an ablation study across all tasks in BulletGym as presented in Table 3. The results indicate a significant decrease in performance, both in normalized reward and normalized cost, when π_{ref} pretraining is omitted. This demonstrates that a pretrained reference policy is crucial for learning a more effective policy.

6 Conclusion

In this paper, we propose a novel approach that reformulates the offline safe RL problem into a trajectory classification setting. Our goal is to develop policies that generate "desirable" trajectories and avoid "undesirable" ones. Specifically, the approach involves partitioning a pre-collected dataset of trajectories into desirable and undesirable subsets. The desirable subset contains high-reward, safe trajectories, while the undesirable subset includes low-reward safe and unsafe trajectories. We then optimize a policy using a classifier that evaluates each trajectory based on the learning policy. This approach bypasses the computational and stability challenges of traditional min-max objectives. Empirical results from the DSRL benchmark show that our method surpasses competitive baselines in terms of reward and constraint satisfaction across a range of tasks. Additional ablation studies further confirm the robustness and effectiveness of our approach.

Acknowledgements

This research/project is supported by the National Research Foundation Singapore and DSO National Laboratories under the AI Singapore Programme (Award Number: AISG2-RP-2020-016).

References

- Achiam, J.; Held, D.; Tamar, A.; and Abbeel, P. 2017. Constrained policy optimization. In *ICML*, 22–31. PMLR.
- Altman, E. 2021. *Constrained Markov decision processes*. Routledge.
- Chen, Y.; Dong, J.; and Wang, Z. 2021. A primal-dual approach to constrained markov decision processes. *arXiv preprint arXiv:2101.10895*.
- Chow, Y.; Ghavamzadeh, M.; Janson, L.; and Pavone, M. 2018. Risk-constrained reinforcement learning with percentile risk criteria. *Journal of Machine Learning Research*, 18(167): 1–51.
- Ding, D.; Zhang, K.; Basar, T.; and Jovanovic, M. 2020. Natural policy gradient primal-dual method for constrained markov decision processes. *Advances in Neural Information Processing Systems*, 33: 8378–8390.
- Ethayarajh, K.; Xu, W.; Muennighoff, N.; Jurafsky, D.; and Kiela, D. 2024. Kto: Model alignment as prospect theoretic optimization. *arXiv preprint arXiv:2402.01306*.
- Fu, J.; Kumar, A.; Nachum, O.; Tucker, G.; and Levine, S. 2020. D4rl: Datasets for deep data-driven reinforcement learning. *arXiv preprint arXiv:2004.07219*.
- Fujimoto, S.; Meger, D.; and Precup, D. 2019. Off-policy deep reinforcement learning without exploration. In *ICML*, 2052–2062. PMLR.
- Galashov, A.; Jayakumar, S. M.; Hasenclever, L.; Tirumala, D.; Schwarz, J.; Desjardins, G.; Czarnecki, W. M.; Teh, Y. W.; Pascanu, R.; and Heess, N. 2019. Information asymmetry in KL-regularized RL. *arXiv preprint arXiv:1905.01240*.
- Garcia, J.; and Fernández, F. 2015. A comprehensive survey on safe reinforcement learning. *Journal of Machine Learning Research*, 16(1): 1437–1480.
- Gronauer, S. 2022. Bullet-safety-gym: A framework for constrained reinforcement learning.
- Gu, S.; Yang, L.; Du, Y.; Chen, G.; Walter, F.; Wang, J.; and Knoll, A. 2022. A review of safe reinforcement learning: Methods, theory and applications. *arXiv preprint arXiv:2205.10330*.
- Guan, J.; Chen, G.; Ji, J.; Yang, L.; Li, Z.; et al. 2024. VOCE: Variational optimization with conservative estimation for offline safe reinforcement learning. *Advances in Neural Information Processing Systems*, 36.
- Gutmann, M.; and Hyvärinen, A. 2010. Noise-contrastive estimation: A new estimation principle for unnormalized statistical models. In *Proceedings of the thirteenth international conference on artificial intelligence and statistics*, 297–304. JMLR Workshop and Conference Proceedings.
- Hejna, J.; Rafailov, R.; Sikchi, H.; Finn, C.; Niekum, S.; Knox, W. B.; and Sadigh, D. 2024. Contrastive Preference Learning: Learning from Human Feedback without Reinforcement Learning. In *ICLR*.
- Ibarz, J.; Tan, J.; Finn, C.; Kalakrishnan, M.; Pastor, P.; and Levine, S. 2021. How to train your robot with deep reinforcement learning: lessons we have learned. *The International Journal of Robotics Research*, 40(4-5): 698–721.
- Ji, J.; Zhou, J.; Zhang, B.; Dai, J.; Pan, X.; Sun, R.; Huang, W.; Geng, Y.; Liu, M.; and Yang, Y. 2023. Omnisafe: An infrastructure for accelerating safe reinforcement learning research. *arXiv preprint arXiv:2305.09304*.
- Kiran, B. R.; Sobh, I.; Talpaert, V.; Mannion, P.; Al Sallab, A. A.; Yogamani, S.; and Pérez, P. 2021. Deep reinforcement learning for autonomous driving: A survey. *IEEE Transactions on Intelligent Transportation Systems*, 23(6): 4909–4926.
- Knox, W.; Hatgis-Kessell, S.; Booth, S.; Niekum, S.; Stone, P.; and Allievi, A. 2024. Models of Human Preference for Learning Reward Functions. *Transactions on Machine Learning Research*.
- Kumar, A.; Fu, J.; Soh, M.; Tucker, G.; and Levine, S. 2019. Stabilizing off-policy q-learning via bootstrapping error reduction. *Advances in neural information processing systems*, 32.
- Le, H.; Voloshin, C.; and Yue, Y. 2019. Batch policy learning under constraints. In *ICML*, 3703–3712. PMLR.
- Lee, J.; Jeon, W.; Lee, B.; Pineau, J.; and Kim, K.-E. 2021. Optidice: Offline policy optimization via stationary distribution correction estimation. In *ICML*, 6120–6130. PMLR.
- Lee, J.; Paduraru, C.; Mankowitz, D. J.; Heess, N.; Precup, D.; Kim, K.-E.; and Guez, A. 2022. COptiDICE: Offline Constrained Reinforcement Learning via Stationary Distribution Correction Estimation. In *ICLR*.
- Li, Q.; Peng, Z.; Feng, L.; Zhang, Q.; Xue, Z.; and Zhou, B. 2022. Metadrive: Composing diverse driving scenarios for generalizable reinforcement learning. *IEEE transactions on pattern analysis and machine intelligence*, 45(3): 3461–3475.
- Liu, Z.; Guo, Z.; Lin, H.; Yao, Y.; Zhu, J.; Cen, Z.; Hu, H.; Yu, W.; Zhang, T.; Tan, J.; et al. 2023a. Datasets and benchmarks for offline safe reinforcement learning. *arXiv preprint arXiv:2306.09303*.
- Liu, Z.; Guo, Z.; Yao, Y.; Cen, Z.; Yu, W.; Zhang, T.; and Zhao, D. 2023b. Constrained decision transformer for offline safe reinforcement learning. In *ICML*, 21611–21630. PMLR.
- Luo, Y.; and Ma, T. 2021. Learning barrier certificates: Towards safe reinforcement learning with zero training-time violations. *Advances in Neural Information Processing Systems*, 34: 25621–25632.
- Mnih, V.; Kavukcuoglu, K.; Silver, D.; Rusu, A. A.; Veness, J.; Bellemare, M. G.; Graves, A.; Riedmiller, M.; Fidjeland, A. K.; Ostrovski, G.; et al. 2015. Human-level control through deep reinforcement learning. *nature*, 518(7540): 529–533.
- Ray, A.; Achiam, J.; and Amodei, D. 2019. Benchmarking safe exploration in deep reinforcement learning. *arXiv preprint arXiv:1910.01708*, 7(1): 2.
- Stooke, A.; Achiam, J.; and Abbeel, P. 2020. Responsive safety in reinforcement learning by pid lagrangian methods. In *ICML*, 9133–9143. PMLR.
- Tessler, C.; Mankowitz, D. J.; and Mannor, S. 2018. Reward constrained policy optimization. *arXiv preprint arXiv:1805.11074*.
- Xu, H.; Zhan, X.; and Zhu, X. 2022. Constraints penalized q-learning for safe offline reinforcement learning. In *AAAI*, volume 36, 8753–8760.
- Zhao, W.; He, T.; and Liu, C. 2021. Model-free safe control for zero-violation reinforcement learning. In *5th Annual Conference on Robot Learning*.
- Zheng, Y.; Li, J.; Yu, D.; Yang, Y.; Li, S. E.; Zhan, X.; and Liu, J. 2024a. Safe offline reinforcement learning with feasibility-guided diffusion model. *arXiv preprint arXiv:2401.10700*.
- Zheng, Y.; Li, J.; Yu, D.; Yang, Y.; Li, S. E.; Zhan, X.; and Liu, J. 2024b. Safe Offline Reinforcement Learning with Feasibility-Guided Diffusion Model. In *ICLR*.
- Ziebart, B. D. 2010. *Modeling purposeful adaptive behavior with the principle of maximum causal entropy*. Carnegie Mellon University.
- Ziebart, B. D.; Maas, A. L.; Bagnell, J. A.; Dey, A. K.; et al. 2008. Maximum entropy inverse reinforcement learning. In *Aaai*, volume 8, 1433–1438. Chicago, IL, USA.