

SafeMIL: Learning Offline Safe Imitation Policy from Non-Preferred Trajectories

Returaj Burnwal¹, Nirav Pravinbhai Bhatt², Balaraman Ravindran²

¹Department of Computer Science and Engineering,

²Department of Data Science and AI, Wadhvani School of Data Science and AI

Indian Institute of Technology Madras, India

cs21d406@smail.iitm.ac.in, {niravbhatt, ravi}@dsai.iitm.ac.in

Abstract

In this work, we study the problem of offline safe imitation learning (IL). In many real-world settings, online interactions can be risky, and accurately specifying the reward and the safety cost information at each timestep can be difficult. However, it is often feasible to collect trajectories reflecting undesirable or risky behavior, implicitly conveying the behavior the agent should avoid. We refer to these trajectories as non-preferred trajectories. Unlike standard IL, which aims to mimic demonstrations, our agent must also learn to avoid risky behavior using non-preferred trajectories. In this paper, we propose a novel approach, SafeMIL, to learn a parameterized cost that predicts if the state-action pair is risky via *Multiple Instance Learning*. The learned cost is then used to avoid non-preferred behaviors, resulting in a policy that prioritizes safety. We empirically demonstrate that our approach can learn a safer policy that satisfies cost constraints without degrading the reward performance, thereby outperforming several baselines.

Introduction

Reinforcement Learning (RL) has demonstrated significant success across various challenging tasks (Kober, Bagnell, and Peters 2013; Mnih et al. 2015; Silver et al. 2016; Kiran et al. 2021). Nevertheless, its deployment in real-world applications is limited because most RL algorithms require many online interactions with the environment to learn a good policy. This requirement hinders its use in many real-world domains, especially where such interactions are risky, such as robotics or autonomous driving. Another challenge lies in selecting a suitable reward function for complex real-world tasks. RL agent learns a policy by maximizing the reward; however, crafting a reward function that encodes the desired behavior can be challenging in various practical applications. An ill-defined reward function can lead to unintended and potentially harmful behaviors (Dewey 2014; Baker et al. 2020). In contrast to RL, which relies on environmental rewards, Imitation Learning (IL) (Schaal 1996) has the distinct advantage of learning solely from expert demonstrations, eliminating the need to design a reward function.

In most real-world domains, agents should also adhere to safety constraints alongside reward maximization. Safety in RL is usually modeled as a Constrained Markov Decision Process (CMDP) (Altman 1999; Ray, Achiam, and Amodei 2019; Ji et al. 2023), where a policy is considered safe if it satisfies all the safety constraints. Offline Safe RL (Polosky et al. 2022; Liu et al. 2023) methods can learn safe policies from pre-collected datasets. However, these methods require access to constraint cost information at each timestep, which can be challenging to obtain. For example, assessing the toxicity in conversational agents or defining safety constraints in complex domains like surgical robotics can be difficult. In contrast, labeling a trajectory non-preferred is easier than defining constraint costs at each timestep. Furthermore, non-preferred trajectories that violate safety constraints are often naturally collected. For instance, vehicle accidents and toxic content reported by chatbot users can be used as non-preferred trajectories.

In this paper, we address the problem of offline safe imitation learning, where the agent learns a safe policy by avoiding non-preferred behaviors. Our approach relies on a small number of non-preferred trajectories, along with a large collection of unlabeled trajectories that contain a mix of preferred (i.e., high return, safety-constraint satisfying) and non-preferred (i.e., high return but safety-constraint violating) trajectories, where both per-timestep reward and cost information are unknown. For example, consider a real-world scenario of autonomous driving where we have access to limited non-preferred trajectories (e.g., red-light running, proximity to other vehicles) and a large unlabeled dataset of human-driven trajectories, which may include non-preferred trajectories. Given no further online data collection, the challenge is to learn a safe policy from these offline datasets.

We propose a novel *offline Safe imitation learning via Multiple Instance Learning*, SafeMIL, algorithm that can learn a safe policy using limited non-preferred trajectories and a large number of unlabeled trajectories where both reward and cost information are unavailable. Our key contributions are as follows: 1) We formulate the learning of a parameterized cost function that predicts if the state-action pair is risky as a Multiple Instance Learning (MIL) problem. Our work is the first to introduce the MIL formulation for offline safe IL setting. 2) Our method can learn the cost function even with a simple and intuitive score function, equation 6.

The learned cost is then used to identify preferred behavior in the unlabeled dataset. This allows the agent to learn a policy via behavior cloning that imitates behaviors likely to satisfy CMDP constraints. 3) We empirically demonstrate that our proposed algorithm learns significantly safer policies. It consistently outperforms or matches the state-of-the-art offline safe imitation learning methods in terms of safety, achieving final median performance that is $3.7\times$ better than the best baseline algorithm across all environments (See Table 3, 4 in Burnwal, Bhatt, and Ravindran (2025)).

Related Works

Offline Imitation Learning (Pomerleau 1991; Kim et al. 2022) primarily focuses on replicating actions of the demonstrations without explicitly considering safety. It implicitly assumes that demonstrations are safe, but if they contain non-preferred trajectories, simply imitating them can lead to learning policies that are risky. We address the problem of learning a safe imitation policy by seeking limited access to non-preferred trajectories and an abundance of unlabeled trajectories that contain both preferred and non-preferred trajectories. This specific scenario has received limited attention in the literature.

Learning Imitation Policy from Non-Preferred Trajectories : SafeDICE (Jang et al. 2023) addresses this challenge by directly estimating the stationary distribution of the preferred policy. The estimated distribution is then used to learn a safe policy that mimics the preferred behavior while effectively avoiding the non-preferred behaviors.

Learning Imitation Policy from Suboptimal Trajectories : T-REX (Brown et al. 2019), B-Pref (Lee et al. 2021), PEBBLE (Lee, Smith, and Abbeel 2021), OPRL (Shin, Dragan, and Brown 2023), learn a reward function from ranked trajectories. These methods focus on learning reward functions that assign a higher total reward to higher-ranked trajectories in the dataset. T-REX, B-Pref, PEBBLE then uses online-RL (Schulman et al. 2017; Haarnoja et al. 2018) algorithm, and OPRL uses offline-RL (Kumar et al. 2020; Qin et al. 2022) algorithm to learn the policy. Since ranking trajectories can be challenging, D-REX (Brown, Goo, and Niekum 2020) and SSR (Chen, Paleja, and Gombolay 2020) generate trajectories of varying quality by injecting noise into policies learned from suboptimal demonstrations. These approaches then use the resulting trajectories to learn a reward function. However, trajectory generation requires online interaction with the environment, which is not viable in an offline setting. Another approach, DWBC (Xu et al. 2022), uses positive-unlabeled (PU) learning (Elkan and Noto 2008; Żolna et al. 2021) to train a discriminator network, which is then incorporated as a weight in the BC loss function.

Preliminaries

Constrained Markov Decision Process

Consider a finite Constrained Markov Decision Process (CMDP) (Altman 1999) represented by the tuple $\mathcal{M} =$

$(\mathcal{S}, \mathcal{A}, \mathcal{P}, r, \rho_0, \gamma, \mathcal{C})$, where \mathcal{S} and \mathcal{A} represent state and action spaces, respectively. $\mathcal{P}(s_{t+1}|s_t, a_t)$ defines the probability of transitioning to state s_{t+1} after executing action a_t in state s_t at timestep t . $r : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$ denotes the immediate reward, ρ_0 is the initial state distribution, and $\gamma \in (0, 1)$ is the discount factor. $\mathcal{C} = \{(c_i, b_i)\}_{i=1}^m$ is a constraint set, where $c_i : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}_{\geq 0}$ is the i -th cost function and $b_i \in \mathbb{R}_{\geq 0}$ is the corresponding threshold. A policy $\pi : \mathcal{S} \rightarrow P(\mathcal{A})$ corresponds to a map from state to a probability distribution over actions.

The set of feasible stationary policies for a CMDP is:

$$\Pi_{\mathcal{C}} := \left\{ \pi \mid \forall i, \mathbb{E}_{\tau \sim \pi} \left[\sum_{t=0}^{T-1} \gamma^t c_i(s_t, a_t) \right] \leq b_i \right\} \quad (1)$$

where $\mathbb{E}_{\tau \sim \pi} \left[\sum_{t=0}^{T-1} \gamma^t c_i(s_t, a_t) \right] \leq b_i$ is the i -th constraint and $\tau = (s_0, a_0, s_1, a_1, \dots, s_{T-1}, a_{T-1}, s_T)$ is a T length trajectory sampled under policy π . The reinforcement learning problem in CMDP is then to find a safe optimal policy π^* that maximizes the expected discounted cumulative reward while satisfying all constraints:

$$\pi^* = \arg \max_{\pi \in \Pi_{\mathcal{C}}} \mathbb{E}_{\tau \sim \pi} \left[\sum_{t=0}^{T-1} \gamma^t r(s_t, a_t) \right] \quad (2)$$

Imitation learning does not rely on environmental rewards but rather on demonstrations. This work focuses on an offline setting, where the policy is learned solely from the pre-collected trajectories that do not contain reward and cost information explicitly, i.e., $\tau = (s_0, a_0, s_1, a_1, \dots, s_{T-1}, a_{T-1}, s_T)$. In this work, preferred and non-preferred trajectories are defined as follows:

Definition 1 (Preferred Trajectory). *A preferred trajectory is defined as a trajectory that achieves high returns and satisfies all the constraints.*

Definition 2 (Non-preferred Trajectory). *A non-preferred trajectory is defined as a trajectory that achieves high returns but violates constraints.*

We assume that we have access to a limited number of non-preferred trajectories, \mathcal{D}^N , and a large number of unlabeled trajectories, \mathcal{D}^U , containing both preferred and non-preferred trajectories. Additionally, we represent the empirical distribution of non-preferred trajectories as ρ^N . This means that a trajectory $\tau \in \mathcal{D}^N$ is a sample drawn from $\tau \sim \rho^N$. Likewise, the stationary distribution, ρ^U , characterizes the distribution of trajectories within \mathcal{D}^U .

Multiple Instance Learning

The goal of the standard binary supervised learning algorithm is to predict the target label, $y \in \{0, 1\}$, for a given instance, x . In Multiple Instance Learning (MIL) (Carbonneau et al. 2018) problem, instead of a single instance, there is a bag of instances, $\mathcal{B} = \{x_1, x_2, \dots, x_K\}$, that are neither dependent nor ordered among each other. In general, the size of the bag, K , can vary; however, in our work, we chose to fix K . We also assume that there exists a binary label for all the instances of the bag, i.e., y_1, \dots, y_K and

$y_k \in \{0, 1\} \forall k = 1, \dots, K$, however, we do not have access to these instance labels. Instead, we can access a single binary label, Y , associated with each bag. The bag label is determined by the presence or absence of at least one positive instance. A bag is negative, $Y = 0$; when all the instances of the bag are negative, i.e., $y_k = 0, \forall k$ and the bag is positive, $Y = 1$, when there is at least one positive instance in the bag, i.e., $y_k = 1$ for some k . The label Y is defined as:

$$Y = \begin{cases} 0, & y_k = 0, \forall k \text{ in bag } \mathcal{B} \\ 1, & \text{otherwise} \end{cases} \quad (3)$$

MIL has two primary classification tasks: bag-level and instance-level classification. Bag-level classification (Zhang and Goldman 2001; Zhou and Zhang 2007) focuses on determining the label of the entire bag of instances, and instance-level classification (Andrews, Tsochantaridis, and Hofmann 2002; Ilse, Tomczak, and Welling 2018) requires classifying each instance within the bag.

Methodology

This section presents *offline Safe imitation learning via Multiple Instance Learning* (SafeMIL) algorithm. We learn the cost function by formulating it as a MIL problem, and then use it to identify preferred behavior from unlabeled dataset. We learn a safe policy via behavior cloning on these preferred behaviors likely to satisfy CMDP constraints.

Formulating Cost Function Learning as MIL Problem

To frame our cost function learning problem as a MIL problem, we need to construct bags and assign positive/negative labels. Suppose we define our bag whose instances are trajectories, $\mathcal{B} = \{\tau_1, \dots, \tau_K\}$. Then, the negative bags can be created by sampling trajectories with replacement from the non-preferred trajectory dataset, as all trajectories within the bag are non-preferred. Similarly, we construct unlabeled bags by sampling trajectories with replacement from unlabeled dataset. For an unlabeled bag to be considered as positive bag, it must contain at least one preferred trajectory. For some bag size K , the probability of the unlabeled bag containing at least one preferred trajectory is defined as:

Lemma 1. *Let \mathcal{T}_p denote the set of all preferred trajectories. Let $\alpha \in (0, 1)$ represent the proportion of preferred trajectories within the unlabeled dataset \mathcal{D}^U . Consider a bag \mathcal{B} containing K trajectories sampled with replacement from \mathcal{D}^U . Then, the probability that bag \mathcal{B} contains at least one preferred trajectory is given by:*

$$P(\mathcal{B} \cap \mathcal{T}_p \neq \emptyset) = 1 - (1 - \alpha)^K$$

Proof. See Appendix A. in Burnwal, Bhatt, and Ravindran (2025). \square

Therefore, for sufficiently large bag size K , the probability of an unlabeled bag containing at least one preferred trajectory approaches 1.

We will overload the notation ρ^N to represent both trajectory ($\tau \sim \rho^N$) and negative bag ($\mathcal{B} \sim \rho^N$) sampling, and

similarly, overload ρ^U for trajectory ($\tau \sim \rho^U$) and unlabeled bag ($\mathcal{B} \sim \rho^U$) sampling.

As we only have access to bag-level labels, we can define a score for a bag \mathcal{B} such that it can capture the bag's underlying property. Since the trajectories within a bag are unordered and independent, the score should be invariant to any permutation of these trajectories. To ensure this permutation invariance, we adopt a score function based on the Fundamental Theorem of Symmetric Functions with monomials (Zaheer et al. 2017), defined as:

$$Score(\mathcal{B}) = g\left(\sum_{\tau \in \mathcal{B}} f(\tau)\right) \quad (4)$$

Here, f and g are suitable transformation functions. This score function is invariant to the permutation of the trajectories within the bag. The effectiveness of this score function relies on the appropriate choice of these transformation functions. We opt for intuitive functions for f and g . We define the function f parameterized by θ as:

$$f(\tau) = \frac{1}{K} \sum_{t=0}^{T-1} \gamma^t \hat{c}_\theta(s_t, a_t) \quad (5)$$

where (s_t, a_t) is the state-action pair at timestep t in the trajectory τ and $\hat{c}_\theta : \mathcal{S} \times \mathcal{A} \rightarrow (0, 1)$ is the parameterized cost function. Therefore, the function f estimates the cumulative cost of the trajectory τ . For the function g , we choose an identity function. This choice of g allows the bag score to be directly proportional to the sum of the cumulative costs of its trajectories. Given these choices for f and g , the score function for the bag \mathcal{B} is:

$$Score(\mathcal{B}) = \frac{1}{K} \sum_{\tau \in \mathcal{B}} \sum_{t=0}^{T-1} \gamma^t \hat{c}_\theta(s_t, a_t) \quad (6)$$

This score function has an intuitive interpretation. As the number of trajectories in the bag, K approaches infinity, the score function converges to the expected cumulative cost of a trajectory in bag \mathcal{B} :

$$\lim_{K \rightarrow \infty} Score(\mathcal{B}) = \mathbb{E}_{\tau \sim \mathcal{B}} \left[\sum_{t=0}^{T-1} \gamma^t \hat{c}_\theta(s_t, a_t) \right] \quad (7)$$

Then, the score function, equation 6, represents an empirical estimate of the expected cumulative cost of a trajectory associated with the bag \mathcal{B} . Assuming non-preferred trajectories have similar cost, then we can establish the following relationship:

$$\lim_{K \rightarrow \infty} Score(\mathcal{B} \sim \rho^N) > \lim_{K \rightarrow \infty} Score(\mathcal{B} \sim \rho^U) \quad (8)$$

This means that as the number of trajectories in negative and unlabeled bags approaches infinity, the expected cumulative cost of the negative bag's trajectory will be greater than that of the unlabeled bag. This is based on the intuition that the negative bag sampled from the non-preferred dataset will have more constraint-violating behavior than the unlabeled bag, which, according to Lemma 1, the unlabeled infinite-size bag is likely to contain at least one preferred trajectory.

Similarly, we can explain the relationship between the empirical scores of negative and unlabeled bags:

Theorem 1. Assuming non-preferred trajectories have similar cost. Then, for some bag size K , the score of negative (\mathcal{B}_n) and unlabeled (\mathcal{B}_u) bag satisfies $\text{Score}(\mathcal{B}_n) > \text{Score}(\mathcal{B}_u)$ with probability

$$P(\text{Score}(\mathcal{B}_n) > \text{Score}(\mathcal{B}_u)) = 1 - (1 - \alpha)^K$$

where α is the percentage of preferred trajectories in the unlabeled dataset.

Proof. The result follows directly from Lemma 1. The unlabeled bag \mathcal{B}_u contains at least one preferred trajectory with $1 - (1 - \alpha)^K$ probability. Therefore, the scores also satisfy the above relationship with the same probability. \square

Given the relationship between the empirical scores of negative and unlabeled bags, we can train the cost function \hat{c}_θ using Bradley-Terry model (Bradley and Terry 1952). Therefore, the loss function can be defined as:

$$\mathcal{L}_\theta = -\mathbb{E}_{\substack{\mathcal{B}_n \sim \rho^N \\ \mathcal{B}_u \sim \rho^U}} \left[\log \frac{\exp(\text{Score}(\mathcal{B}_n))}{\exp(\text{Score}(\mathcal{B}_n)) + \exp(\text{Score}(\mathcal{B}_u))} \right] \quad (9)$$

By minimizing the above loss function, we can to train the cost function \hat{c}_θ , which assigns a higher score to a negative bag than an unlabeled bag. In the next section, we discuss a method that uses this learned cost function to learn a preferred policy.

Policy Learning

The learned cost function \hat{c}_θ estimates whether a given state-action pair is risky, i.e., likely to violate constraints. Since \mathcal{D}^U contains both preferred and non-preferred trajectories, we can use this learned cost function to identify preferred trajectories that are likely to satisfy the CMDP constraints. Specifically, we identify the set of preferred trajectories as:

$$\mathcal{T}_{\hat{c}_\theta} := \left\{ \tau \in \mathcal{D}^U \mid \sum_{t=0}^{T-1} \gamma^t \hat{c}_\theta(s_t, a_t) \leq \hat{b} \right\} \quad (10)$$

where \hat{b} is a constraint threshold. We can then recover the preferred policy by optimizing the following loss function:

$$\min_{\pi} \sum_{\tau \in \mathcal{T}_{\hat{c}_\theta}} \left[\sum_{t=0}^{T-1} \mathcal{L}_\pi(s_t, a_t) \right] \quad (11)$$

where (s_t, a_t) is the state-action pair at timestep t in trajectory τ and $\mathcal{L}_\pi(\cdot, \cdot)$ is the behavior cloning loss function.

Alternatively, instead of applying a hard threshold, \hat{b} , to select trajectories, we can take a soft weighting approach by assigning weight to each trajectory based on its estimated cost. This results in the following weighted objective:

$$\min_{\pi} \sum_{\tau \in \mathcal{D}^U} \left[w(\tau) \sum_{t=0}^{T-1} \mathcal{L}_\pi(s_t, a_t) \right] \quad (12)$$

where the trajectory weight $w(\tau)$ is defined as:

$$w(\tau) = \exp \left(- \sum_{t=0}^{T-1} \gamma^t \hat{c}_\theta(s_t, a_t) / \beta \right) \quad (13)$$

Algorithm 1: SafeMIL

Input: non-preferred dataset \mathcal{D}^N , unlabeled dataset \mathcal{D}^U , bag size K , partial trajectory length H , β

Parameter: a parameterized cost function \hat{c}_θ and a parameterized policy network π

Output: preferred policy $\pi \approx \pi^*$

- 1: **for** $n = 1, 2, \dots$ **do**
 - 2: Sample bag $\mathcal{B}_n \sim \mathcal{D}^N$ of bag size K and trajectory length H
 - 3: Sample bag $\mathcal{B}_u \sim \mathcal{D}^U$ of bag size K and trajectory length H
 - 4: Update cost function \hat{c}_θ by optimizing the Eq. 9
 - 5: Update policy π on \mathcal{D}^U using the Eq. 12
 - 6: **end for**
-

and $\beta > 0$ is a hyperparameter. For small value of β , this weighting scheme assigns higher weight to constraint-satisfying preferred behaviors over non-preferred ones. This approach also allows us to leveraging all available data while prioritizing policy to mimic preferred behaviors while discouraging non-preferred behaviors.

Extension to Learning Cost Function with Partial Trajectories

The method so far assumes the use of full trajectory length; however, using the entire trajectory can be computationally expensive or sometimes impractical. Therefore, we modify our algorithm to construct bags with partial trajectories instead of full trajectory lengths. This modification can lead to incorrect sampling of negative bags from non-preferred dataset. The partial trajectories sampled from non-preferred dataset can exhibit preferred behavior, resulting in incorrect labeling of these bags. However, even if some negative bags are mislabeled, the relationship between the scores of the negative bag and unlabeled bag, equation 8, can still hold provided the number of trajectories, K , in each bag, is sufficiently large. This is because the score typically reflects the average cost of the trajectories within a bag. Even with some preferred segments in the negative bag, the overall average cost in the negative bag is expected to be higher than that in the unlabeled bag, given a sufficiently large bag size.

Training using partial trajectories allows the SafeMIL algorithm to learn an expressive cost function from only a limited number of non-preferred trajectories. During training, we randomly sample pairs of negative and unlabeled bags, each containing K partial trajectories of length H . The score for each bag is then used to compute “logit” in a binary cross-entropy loss function. The learned cost function is used to identify preferred behavior from unlabeled dataset. We learn a safe policy via behavior cloning on these preferred behaviors likely to satisfy CMDP constraints. The training procedure is shown in Algorithm 1.

Experiments

We evaluate our algorithm against the state-of-the-art offline safe IL algorithm using the Datasets for offline Safe RL (DSRL) on a set of benchmark tasks (Liu et al.

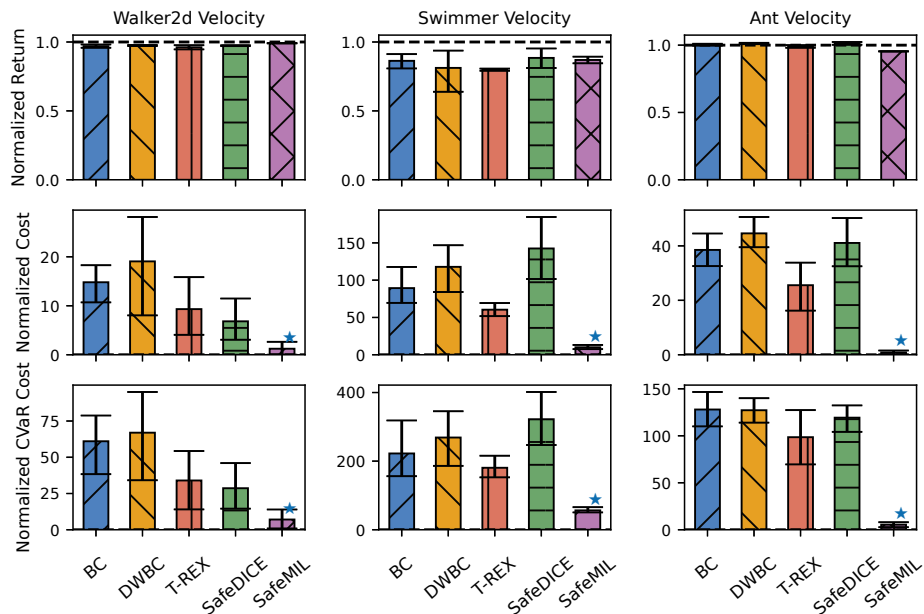


Figure 1: Performance Comparison. We report the final bootstrapped mean performance of the algorithm on Walker2d-Velocity, Swimmer-Velocity, Ant-Velocity task after 1 million training steps. Mean and 95% CIs over 5 seeds. We observe that our method outperforms all the baselines and can recover low cost safe policies without compromising reward performance. Fig. 8 (Burnwal, Bhatt, and Ravindran 2025) reports the learning curves for all the algorithms.

2024). We conduct experiments on the following tasks: i) MuJoCo-based velocity-constrained tasks (Walker-Velocity, Swimmer-Velocity, Ant-Velocity), where agents are required to move as fast as possible while adhering to the velocity limits; and ii) navigation tasks (Point-Circle2, Point-Goal1, Point-Button1), that require agents to maximize performance while avoiding collisions or contact with hazardous regions (illustrated in Fig. 5, 6 in Burnwal, Bhatt, and Ravindran (2025)). These tasks collectively provide comprehensive and realistic scenarios for evaluating the safety and the performance of offline safe IL algorithms.

We select a limited number of non-preferred trajectories (\mathcal{D}^N) and a large number of unlabeled trajectories (\mathcal{D}^U) from the DSRL dataset. \mathcal{D}^U contain a mix of preferred and non-preferred trajectories. We removed all the reward and cost information from \mathcal{D}^N and \mathcal{D}^U trajectory dataset. Our objective is to recover low-cost behavior while preserving high-reward performance from the unlabeled data. We have 50 trajectories in \mathcal{D}^N and 200 trajectories in \mathcal{D}^U . Appendix D. (Burnwal, Bhatt, and Ravindran 2025) contains all the other details regarding \mathcal{D}^N and \mathcal{D}^U dataset.

Baselines: We compare our algorithm against four baselines. **(1) BC**, a BC policy on the unlabeled dataset \mathcal{D}^U . This serves as a baseline to evaluate the performance of BC when the dataset contains a mixture of preferred and non-preferred trajectories. **(2) SafeDICE**, Jang et al. (2023), directly estimates the stationary distribution corrections for the preferred behavior and then trains a weighted BC policy. **(3) DWBC**, a variant of the discriminator-weighted behavior cloning algorithm (Xu et al. 2022) that uses *negative-unlabeled* learning

to train the discriminator model. The trained discriminator is then used as weight in the weighted BC loss function. **(4) T-REX**, Brown et al. (2019), a preference-based method that learns a reward function to prefer unlabeled trajectories over non-preferred ones. The learned reward function is then used as a weight in the weighted BC loss function. See Appendix C. (Burnwal, Bhatt, and Ravindran 2025) for implementation details of baseline algorithms.

We also train an offline constrained RL policy, COpTiDICE (Lee et al. 2022), on an unlabeled dataset augmented with ground-truth reward and cost annotations for 1 million training steps. Constrained-RL policy represents the optimal performance achievable by a policy that simultaneously achieves high reward and low cost. We treat this Constrained-RL policy as the preferred low-cost policy. No other methods have access to reward and cost information. We do not compare our algorithm with other constrained RL algorithms, as our setting, offline safe IL, fundamentally differs from offline safe RL.

We compare the performance of SafeMIL and other baseline algorithms in terms of task performance (i.e., expected episodic return) and safety (i.e., expected episodic cost). We report the results with the following metrics. **(1) Normalized Return** scales the episodic return of a given policy, such that 0 represents the episodic return from a random policy, while 1 represents the episodic return achieved by the Constrained-RL policy. **(2) Normalized Cost** scales a given policy’s episodic cost by subtracting the Constrained-RL policy’s episodic cost. It ensures that 0 represents the cost achieved by the Constrained-RL policy. **(3) Normal-**

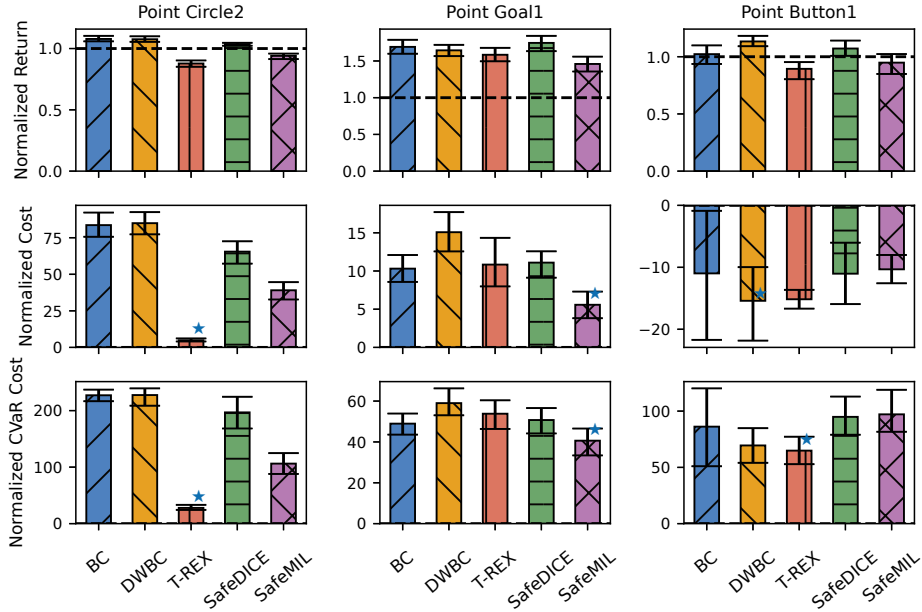


Figure 2: Performance Comparison. We report the final bootstrapped mean performance of the algorithm on Point-Circle2, Point-Goal1, Point-Button1 tasks after 1 million training steps. Mean and 95% CIs over 5 seeds. We observe that the Point-Goal1 environment, our algorithm performs better, while maintaining competitive performance in the remaining environments. Fig. 9 (Burnwal, Bhatt, and Ravindran 2025) reports the learning curves for all the algorithms.

ized Conditional Value at Risk performance (CVaR) 20% Cost, similar to *Normalized Cost* it scales the policy’s mean episodic cost of the worst 20% runs by subtracting the Constrained-RL policy’s episodic cost.

Our primary focus of the comparison is to evaluate safety (*Normalized Cost*, *Normalized CVaR@20% Cost*). We consider a policy to have successfully recovered the preferred behavior if its performance closely matches the Constrained-RL policy. This is evidenced by the *Normalized Cost* less than or close to 0, and the *Normalized Return* being greater than or close to 1. All plots are generated by averaging the performance of 50 trajectories generated from the learned policy. To assess statistical significance, we generate 1000 bootstrap samples from the data, using results from 5 different random seeds, and we plot the resulting 95% confidence intervals.

Through our experiments, we focus on answering the following questions:

1. *Performance Comparison*: How does our algorithm compare to other baselines in terms of learning safer policies that satisfy cost constraints while maintaining reward performance?
2. Sensitivity to the choice of bag size: How does our algorithm’s performance vary with different bag sizes K ?
3. Sensitivity to the choice of trajectory length: What is the impact of trajectory length on our algorithm’s performance?
4. Effect of different weighting scheme in policy learning: How does our algorithm’s performance change if we de-

fine weights on individual state-action pairs instead of the entire trajectory when learning the policy?

Performance Comparison

In both velocity-constrained tasks (Fig. 1) and navigation tasks (Fig. 2), we verify that when the unlabeled dataset contains non-preferred (constraint-violating) trajectories, the standard BC policy has a higher cost and therefore can exhibit non-preferred behavior.

SafeMIL can outperform all baseline algorithms in velocity-constrained tasks. It is also able to recover preferred behavior from the unlabeled dataset, with safety performance closely approximating that of Constrained-RL. SafeMIL learns a safer policy for navigation tasks in Point-Goal1 and is competitive compared to other baselines in Point-Circle2 and Point-Button1. In Point-Circle2, T-REX learns a policy with a better constraint-satisfying policy. SafeMIL achieves a median performance of $3.7\times$ better than the best baseline algorithm across all environments. Appendix F. (Burnwal, Bhatt, and Ravindran 2025) contains additional supporting results and the learning curves for all the algorithms.

Sensitivity Analysis

This section investigates how the bag’s size and the trajectory length impact our algorithm’s performance. We demonstrate that increasing the bag size and using partial trajectory length does not alter the algorithm’s safety or the overall reward achieved.




Figure 3: Sensitivity to Bag Size. We report the final mean performance of the algorithm on the Swimmer-Velocity environment for different bag sizes $K = \{1, 8, 16, 64, 128\}$, after training for 1 million steps. Mean and 95% CIs over 5 seeds. We observe that increasing the bag size (K) lead to a higher probability of finding preferred trajectories within the unlabeled bag. Thereby, improving the algorithm’s safety performance while maintaining reasonable episode return.




Figure 4: Sensitivity to Trajectory Length. We report the final mean performance of the algorithm on the Swimmer-Velocity environment for different partial trajectory lengths $H = \{1, 5, 10\}$, after training for 1 million steps. Mean and 95% CIs over 5 seeds. For a sufficient bag size of 128, we observe that safety performance is stable across different trajectory lengths.

Sensitivity to Bag Size: We train the SafeMIL algorithm with these different bag sizes, $K = \{1, 8, 16, 64, 128\}$, for 1M training steps and report the final performance in Fig. 3. Based on Theorem 1, we expect that increasing the bag size (K) would lead to a higher probability of finding preferred trajectories within the unlabeled bag. Therefore, increasing bag size would improve the algorithm’s safety performance while maintaining reasonable episode return. As shown in Fig. 3, as K increases, the safety performance also improves (lower cost) and becomes stable after some bag size.

Sensitivity to Trajectory Length: To evaluate the effect of trajectory length on our SafeMIL algorithm, we performed experiments on the Swimmer-Velocity environments for sufficiently large bag size ($K = 128$) with varying trajectory length, $H = \{1, 5, 10\}$. The SafeMIL algorithm is trained with these trajectory lengths for 1M training steps, and the final performance results are summarized in Fig. 4. We observe that the safety performance is stable irrespective of the trajectory length. This finding supports using partial trajectories for training, as they offer a computationally efficient alternative to utilizing full-length trajectories without compromising overall performance.

Effect of Different Weighing Scheme

This section analyzes the effect of different weighing scheme in policy learning. Specifically, we want to verify the algorithm’s performance when we define weights on in-

dividual state-action pair transitions instead of the entire trajectory when learning the policy. Since \hat{c}_θ estimates the likelihood of a state-action pair being risky, then we can learn the safe policy by defining the weight in the BC loss function at each state-action (s, a) pair as: $(1 - \hat{c}_\theta(s, a))$. The loss function is then defined as:

$$\min_{\pi} \sum_{(s,a) \in \mathcal{D}^U} (1 - \hat{c}_\theta(s, a)) \mathcal{L}_\pi(s, a) \quad (14)$$

We refer to this loss function as transition weighted BC loss function and call equation 12 as trajectory weighted BC loss function. Fig. 12, 13 (Burnwal, Bhatt, and Ravindran 2025) reports the result of using transition and trajectory weighted BC loss function to train safe policy. We observe that both weighing schemes for policy learning have similar safety and return performance for Swimmer-Velocity and Point-Goal1 tasks.

Conclusion

This work introduces SafeMIL, an offline safe imitation learning algorithm. We frame the problem of learning a cost function as a MIL task. Based on the intuition that non-preferred trajectories incur higher costs than unlabeled ones on average, we learn the cost function that explains this intuition. Subsequently, we employ this learned cost function to identify preferred behavior and use it to learn safe policy. The experiment demonstrates that SafeMIL can learn safer policies in the constrained RL benchmarks.

Acknowledgments

Returaj Burnwal acknowledges financial support from the Prime Minister’s Research Fellowship (PMRF), Ministry of Education, India.

References

- Altman, E. 1999. Constrained Markov Decision Processes.
- Andrews, S.; Tsochantaridis, I.; and Hofmann, T. 2002. Support vector machines for multiple-instance learning. *Advances in neural information processing systems*, 15.
- Baker, B.; Kanitscheider, I.; Markov, T.; Wu, Y.; Powell, G.; McGrew, B.; and Mordatch, I. 2020. Emergent Tool Use From Multi-Agent Autocurricula. In *International Conference on Learning Representations*.
- Bradley, R. A.; and Terry, M. E. 1952. Rank Analysis of Incomplete Block Designs: I. The Method of Paired Comparisons. *Biometrika*, 39: 324.
- Brown, D.; Goo, W.; Nagarajan, P.; and Niekum, S. 2019. Extrapolating beyond suboptimal demonstrations via inverse reinforcement learning from observations. In *International conference on machine learning*, 783–792. PMLR.
- Brown, D. S.; Goo, W.; and Niekum, S. 2020. Better-than-demonstrator imitation learning via automatically-ranked demonstrations. In *Conference on robot learning*, 330–359. PMLR.
- Burnwal, R.; Bhatt, N. P.; and Ravindran, B. 2025. SafeMIL: Learning Offline Safe Imitation Policy from Non-Preferred Trajectories. *arXiv preprint arXiv:2511.08136*.
- Carbonneau, M.-A.; Cheplygina, V.; Granger, E.; and Gagnon, G. 2018. Multiple instance learning: A survey of problem characteristics and applications. *Pattern Recognition*, 77: 329–353.
- Chen, L.; Paleja, R.; and Gombolay, M. 2020. Learning from suboptimal demonstration via self-supervised reward regression. In *Conference on robot learning*, 1262–1277. PMLR.
- Dewey, D. 2014. Reinforcement learning and the reward engineering principle. In *2014 AAAI Spring Symposium Series*.
- Elkan, C.; and Noto, K. 2008. Learning classifiers from only positive and unlabeled data. In *Proceedings of the 14th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 213–220. New York, NY, USA: Association for Computing Machinery. ISBN 9781605581934.
- Haarnoja, T.; Zhou, A.; Abbeel, P.; and Levine, S. 2018. Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. In *International conference on machine learning*, 1861–1870. PMLR.
- Ilse, M.; Tomczak, J.; and Welling, M. 2018. Attention-based deep multiple instance learning. In *International conference on machine learning*, 2127–2136. PMLR.
- Jang, Y.; Kim, G.-H.; Lee, J.; Sohn, S.; Kim, B.; Lee, H.; and Lee, M. 2023. SafeDICE: Offline Safe Imitation Learning with Non-Preferred Demonstrations. In *Thirty-seventh Conference on Neural Information Processing Systems*.
- Ji, J.; Zhang, B.; Zhou, J.; Pan, X.; Huang, W.; Sun, R.; Geng, Y.; Zhong, Y.; Dai, J.; and Yang, Y. 2023. Safety Gymnasium: A Unified Safe Reinforcement Learning Benchmark. In *Thirty-seventh Conference on Neural Information Processing Systems Datasets and Benchmarks Track*.
- Kim, G.-H.; Seo, S.; Lee, J.; Jeon, W.; Hwang, H.; Yang, H.; and Kim, K.-E. 2022. DemoDICE: Offline Imitation Learning with Supplementary Imperfect Demonstrations. In *International Conference on Learning Representations*.
- Kiran, B. R.; Sobh, I.; Talpaert, V.; Mannion, P.; Al Salhab, 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.
- Kober, J.; Bagnell, J. A.; and Peters, J. 2013. Reinforcement learning in robotics: A survey. *The International Journal of Robotics Research*, 32(11): 1238–1274.
- Kumar, A.; Zhou, A.; Tucker, G.; and Levine, S. 2020. Conservative q-learning for offline reinforcement learning. *Advances in Neural Information Processing Systems*, 33: 1179–1191.
- 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 *International Conference on Learning Representations*.
- Lee, K.; Smith, L.; Dragan, A.; and Abbeel, P. 2021. B-Pref: Benchmarking Preference-Based Reinforcement Learning. In *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 1)*.
- Lee, K.; Smith, L. M.; and Abbeel, P. 2021. PEBBLE: Feedback-Efficient Interactive Reinforcement Learning via Relabeling Experience and Unsupervised Pre-training. In *International Conference on Machine Learning*.
- Liu, Z.; Guo, Z.; Lin, H.; Yao, Y.; Zhu, J.; Cen, Z.; Hu, H.; Yu, W.; Zhang, T.; Tan, J.; and Zhao, D. 2024. Datasets and Benchmarks for Offline Safe Reinforcement Learning. *Journal of Data-centric Machine Learning Research*.
- Liu, Z.; Guo, Z.; Yao, Y.; Cen, Z.; Yu, W.; Zhang, T.; and Zhao, D. 2023. Constrained decision transformer for offline safe reinforcement learning. In *International Conference on Machine Learning*, 21611–21630. PMLR.
- 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.
- Polosky, N.; Da Silva, B. C.; Fiterau, M.; and Jagannath, J. 2022. Constrained offline policy optimization. In *International Conference on Machine Learning*, 17801–17810. PMLR.
- Pomerleau, D. A. 1991. Efficient training of artificial neural networks for autonomous navigation. *Neural computation*, 3(1): 88–97.
- Qin, R.-J.; Zhang, X.; Gao, S.; Chen, X.-H.; Li, Z.; Zhang, W.; and Yu, Y. 2022. NeoRL: A near real-world benchmark for offline reinforcement learning. *Advances in Neural Information Processing Systems*, 35: 24753–24765.

- Ray, A.; Achiam, J.; and Amodei, D. 2019. Benchmarking Safe Exploration in Deep Reinforcement Learning.
- Schaal, S. 1996. Learning from demonstration. *Advances in neural information processing systems*, 9.
- Schulman, J.; Wolski, F.; Dhariwal, P.; Radford, A.; and Klimov, O. 2017. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*.
- Shin, D.; Dragan, A.; and Brown, D. S. 2023. Benchmarks and Algorithms for Offline Preference-Based Reward Learning. *Transactions on Machine Learning Research*.
- Silver, D.; Huang, A.; Maddison, C. J.; Guez, A.; Sifre, L.; Van Den Driessche, G.; Schrittwieser, J.; Antonoglou, I.; Panneershelvam, V.; Lanctot, M.; et al. 2016. Mastering the game of Go with deep neural networks and tree search. *Nature*, 529(7587): 484–489.
- Xu, H.; Zhan, X.; Yin, H.; and Qin, H. 2022. Discriminator-weighted offline imitation learning from suboptimal demonstrations. In *International Conference on Machine Learning*, 24725–24742. PMLR.
- Zaheer, M.; Kottur, S.; Ravanbakhsh, S.; Póczos, B.; Salakhutdinov, R. R.; and Smola, A. J. 2017. Deep sets. *Advances in neural information processing systems*, 30.
- Zhang, Q.; and Goldman, S. 2001. EM-DD: An improved multiple-instance learning technique. *Advances in neural information processing systems*, 14.
- Zhou, Z.-H.; and Zhang, M.-L. 2007. Solving multi-instance problems with classifier ensemble based on constructive clustering. *Knowledge and Information Systems*, 11: 155–170.
- Żołna, K.; Saharia, C.; Boussioux, L.; Hui, D. Y.-T.; Chevalier-Boisvert, M.; Bahdanau, D.; and Bengio, Y. 2021. Combating false negatives in adversarial imitation learning. In *2021 International Joint Conference on Neural Networks (IJCNN)*, 1–9. IEEE.