
Adversarially Trained Weighted Actor-Critic for Safe Offline Reinforcement Learning

Honghao Wei
Washington State University
honghao.wei@wsu.edu

Xiyue Peng
ShanghaiTech University
pengxy2024@shanghaitech.edu.cn

Arnob Ghosh
New Jersey Institute of Technology
arnob.ghosh@njit.edu

Xin Liu
ShanghaiTech University
liuxin7@shanghaitech.edu.cn

Abstract

We propose WSAC (Weighted Safe Actor-Critic), a novel algorithm for Safe Offline Reinforcement Learning (RL) under functional approximation, which can robustly optimize policies to improve upon an arbitrary reference policy with limited data coverage. WSAC is designed as a two-player Stackelberg game to optimize a refined objective function. The actor optimizes the policy against two adversarially trained value critics with small importance-weighted Bellman errors, which focus on scenarios where the actor’s performance is inferior to the reference policy. In theory, we demonstrate that when the actor employs a no-regret optimization oracle, WSAC achieves a number of guarantees: (i) For the first time in the safe offline RL setting, we establish that WSAC can produce a policy that outperforms **any** reference policy while maintaining the same level of safety, which is critical to designing a safe algorithm for offline RL. (ii) WSAC achieves the optimal statistical convergence rate of $1/\sqrt{N}$ to the reference policy, where N is the size of the offline dataset. (iii) We theoretically show that WSAC guarantees a safe policy improvement across a broad range of hyperparameters that control the degree of pessimism, indicating its practical robustness. Additionally, we offer a practical version of WSAC and compare it with existing state-of-the-art safe offline RL algorithms in several continuous control environments. WSAC outperforms all baselines across a range of tasks, supporting the theoretical results.

1 Introduction

Online safe reinforcement learning (RL) has found successful applications in various domains, such as autonomous driving (Isele et al., 2018), recommender systems (Chow et al., 2017), and robotics (Achiam et al., 2017). It enables the learning of safe policies effectively while satisfying certain safety constraints, including collision avoidance, budget adherence, and reliability. However, collecting diverse interaction data can be extremely costly and infeasible in many real-world applications, and this challenge becomes even more critical in scenarios where risky behavior cannot be tolerated. Given the inherently risk-sensitive nature of these safety-related tasks, data collection becomes feasible only when employing behavior policies satisfies all the safety requirements.

To overcome the limitations imposed by interactive data collection, offline RL algorithms are designed to learn a policy from an available dataset collected from historical experiences by some behavior policy, which may differ from the policy we aim to learn. A desirable property of an effective offline algorithm is the assurance of robust policy improvement (RPI), which guarantees that a learned policy is always at least as good as the baseline behavior policies (Fujimoto et al., 2019; Laroche et al., 2019;

Kumar et al., 2019; Siegel et al., 2020; Chen et al., 2022a; Zhu et al., 2023; Bhardwaj et al., 2024). We extend the property of RPI to offline safe RL called safe robust policy improvement (SRPI), which indicates the improvement should be *safe* as well. This is particularly important in offline safe RL. For example, in autonomous driving, an expert human driver operates the vehicle to collect a diverse dataset under various road and weather conditions, serving as the behavior policy. This policy is considered both effective and safe, as it demonstrates proficient human driving behavior while adhering to all traffic laws and other safety constraints. Achieving a policy that upholds the SRPI characteristic with such a dataset can significantly mitigate the likelihood of potential collisions and other safety concerns.

In offline RL, we represent the state-action occupancy distribution of policy π over the dataset distribution μ using the ratio $w^\pi = d^\pi / \mu$. A commonly required assumption is that the ℓ_∞ concentrability $C_{\ell_\infty}^\pi$ is bounded, which is defined as the infinite norm of w^π for **all** policies (Liu et al., 2019; Chen and Jiang, 2019; Wang et al., 2019; Liao et al., 2022; Zhang et al., 2020). A stronger assumption requires a uniform lower bound on $\mu(a|s)$ (Xie and Jiang, 2021). However, such all-policy concentrability assumptions are difficult to satisfy in practice, particularly for offline safe RL, as they essentially require the offline dataset to have good coverage of **all** unsafe state-action pairs. To address the full coverage requirement, other works (Rashidinejad et al., 2021; Zhan et al., 2022; Chen and Jiang, 2022; Xie et al., 2021; Uehara and Sun, 2021) adapt conservative algorithms by employing the principle of pessimism in the face of uncertainty, reducing the assumption to the best covered policy (or optimal policy) concerning ℓ_∞ concentrability. Zhu et al. (2023) introduce ℓ_2 concentrability to further relax the assumption, indicating that ℓ_∞ concentrability is always an upper bound of ℓ_2 concentrability (see Table 1 for detailed comparisons with previous works). While provable guarantees are obtained using single policy concentrability for unconstrained MDP as Table 1 suggests for the safe RL setting, all the existing studies (Hong et al., 2024; Le et al., 2019) *still* require the coverage on **all** the policies. Further, as Table 1 suggests, the above papers do not guarantee robust safe policy improvement. Our main contributions are summarized below:

1. We prove that our algorithm, which uses average Bellman error, enjoys an optimal statistical rate of $1/\sqrt{N}$ under partial data coverage assumption. *This is the first work that achieves such a result using only single-policy ℓ_2 concentrability.*
2. We propose a novel offline safe RL algorithm, called Weighted Safe Actor-Critic (WSAC), which can robustly learn policies that improve upon any behavior policy with controlled relative pessimism. We prove that under standard function approximation assumptions, when the actor incorporates a no-regret policy optimization oracle, WSAC outputs a policy that never degrades the performance of a reference policy (including the behavior policy) for a range of hyperparameters (defined later). *This is the first work that provably demonstrates the property of SRPI in offline safe RL setting.*
3. We point out that primal-dual-based approaches Hong et al. (2024) may require **all**-policy concentrability assumption. Thus, unlike, the primal-dual-based approach, we propose a novel rectified penalty-based approach to obtain results using **single-policy** concentrability. Thus, we need novel analysis techniques to prove results.
4. Furthermore, we provide a practical implementation of WSAC following a two-timescale actor-critic framework using adversarial frameworks similar to Cheng et al. (2022); Zhu et al. (2023), and test it on several continuous control environments in the offline safe RL benchmark (Liu et al., 2023a). WSAC outperforms all other state-of-the-art baselines, validating the property of a safe policy improvement.




Figure 1: Comparison between WSAC and the behavior policy in the tabular case. The behavior policy is a mixture of the optimal policy and a random policy, with the mixture percentage representing the proportion of the optimal policy. The cost threshold is set to 0.1. We observe that WSAC consistently ensures a safely improved policy across various scenarios, even when the behavior policy is not safe.

Table 1: Comparison of algorithms for offline RL (safe RL) with function approximation. The parameters $C_{\ell_2}^\pi, C_{\ell_\infty}^\pi, C_{Bellman}^\pi$ refer to different types of concentrabilities, it always hold $C_{\ell_2}^\pi \leq C_{\ell_\infty}^\pi$ and under certain condition $C_{\ell_2}^\pi \leq C_{Bellman}^\pi$, detailed definitions and more discussions can be found in Section 3.3.

Algorithm	Safe RL	Coverage assumption	Policy Improvement	Suboptimality
Xie and Jiang (2021)	No	all policy, $C_{\ell_2}^\pi$	Yes	$\mathcal{O}(1/\sqrt{N})$
Xie et al. (2021)	No	single-policy, $C_{Bellman}^\pi$	Yes	$\mathcal{O}(1/\sqrt{N})$
Cheng et al. (2022)	No	single-policy, $C_{Bellman}^\pi$	Yes & Robust	$\mathcal{O}(1/N^{1/3})$
Ozdaglar et al. (2023)	No	single-policy, $C_{\ell_\infty}^\pi$	No	$\mathcal{O}(1/\sqrt{N})$
Zhu et al. (2023)	No	single-policy, $C_{\ell_2}^\pi$	Yes & Robust	$\mathcal{O}(1/\sqrt{N})$
Le et al. (2019)	Yes	all policy, $C_{\ell_\infty}^\pi$	No	$\mathcal{O}(1/\sqrt{N})$
Hong et al. (2024)	Yes	all policy, $C_{\ell_2}^\pi$	No	$\mathcal{O}(1/\sqrt{N})$
Ours	Yes	single-policy, $C_{\ell_2}^\pi$	Yes & Robust	$\mathcal{O}(1/\sqrt{N})$

2 Related Work

Offline safe RL: Deep offline safe RL algorithms (Lee et al., 2022; Liu et al., 2023b; Xu et al., 2022; Chen et al., 2021; Zheng et al., 2024) have shown strong empirical performance but lack theoretical guarantees. To the best of our knowledge, the investigation of policy improvement properties in offline safe RL is relatively rare in the state-of-the-art offline RL literature. Wu et al. (2021) focus on the offline constrained multi-objective Markov Decision Process (CMOMDP) and demonstrate that an optimal policy can be learned when there is sufficient data coverage. However, although they show that CMDP problems can be formulated as CMOMDP problems, they assume a linear kernel CMOMDP in their paper, whereas our consideration extends to a more general function approximation setting. Le et al. (2019) propose a model-based primal-dual-type algorithm with deviation control for offline safe RL in the tabular setting. With prior knowledge of the slackness in Slater’s condition and a constant on the concentrability coefficient, an (ϵ, δ) -PAC error is achievable when the number of data samples N is large enough ($N = \tilde{\mathcal{O}}(1/\epsilon^2)$). These assumptions make the algorithm impractical, and their computational complexity is much higher than ours. Additionally, we consider a more practical, model-free function approximation setting. In another concurrent work (Hong et al., 2024), a primal-dual critic algorithm is proposed for offline-constrained RL settings with general function approximation. However, their algorithm requires ℓ_2 concentrability for **all** policies, which is not practical as discussed. The reason is that the dual variable optimization in their primal-dual design requires an accurate estimation of all the policies used in each episode, which necessitates coverage over all policies. Moreover, they cannot guarantee the property of SRPI. Moreover, their algorithm requires an additional offline policy evaluation (OPE) oracle for policy evaluation, making the algorithm less efficient.

3 Preliminaries

3.1 Constrained Markov Decision Process

We consider a Constrained Markov Decision Process (CMDP) \mathcal{M} , denoted by $(\mathcal{S}, \mathcal{A}, \mathcal{P}, R, C, \gamma, \rho)$. \mathcal{S} is the state space, \mathcal{A} is the action space, $\mathcal{P} : \mathcal{S} \times \mathcal{A} \rightarrow \Delta(\mathcal{S})$ is the transition kernel, where $\Delta(\cdot)$ is a probability simplex, $R : \mathcal{S} \times \mathcal{A} \rightarrow [0, 1]$ is the reward function, $C : \mathcal{S} \times \mathcal{A} \rightarrow [-1, 1]$ is the cost function, $\gamma \in [0, 1)$ is the discount factor and $\rho : \mathcal{S} \rightarrow [0, 1]$ is the initial state distribution. We assume \mathcal{A} is finite while allowing \mathcal{S} to be arbitrarily complex. We use $\pi : \mathcal{S} \rightarrow \Delta(\mathcal{A})$ to denote a stationary policy, which specifies a distribution over actions for each state. At each time, the agent

observes a state $s_t \in \mathcal{S}$, takes an action $a_t \in \mathcal{A}$ according to a policy π , receives a reward r_t and a cost c_t , where $r_t = R(s_t, a_t)$, $c_t = C(s_t, a_t)$. Then the CMDP moves to the next state s_{t+1} based on the transition kernel $\mathcal{P}(\cdot | s_t, a_t)$. Given a policy π , we use $V_r^\pi(s)$ and $V_c^\pi(s)$ to denote the expected discounted return and the expected cumulative discounted cost of π , starting from state s , respectively.

$$V_r^\pi(s) := \mathbb{E} \left[\sum_{t=0}^{\infty} \gamma^t r_t \mid s_0 = s, a_t \sim \pi(\cdot | s_t) \right] \quad (1)$$

$$V_c^\pi(s) := \mathbb{E} \left[\sum_{t=0}^{\infty} \gamma^t c_t \mid s_0 = s, a_t \sim \pi(\cdot | s_t) \right]. \quad (2)$$

Accordingly, we also define the Q -value function of a policy π for the reward and cost as

$$Q_r^\pi(s, a) := \mathbb{E} \left[\sum_{t=0}^{\infty} \gamma^t r_t \mid (s_0, a_0) = (s, a), a_t \sim \pi(\cdot | s_t) \right] \quad (3)$$

$$Q_c^\pi(s, a) := \mathbb{E} \left[\sum_{t=0}^{\infty} \gamma^t c_t \mid (s_0, a_0) = (s, a), a_t \sim \pi(\cdot | s_t) \right], \quad (4)$$

respectively. As rewards and costs are bounded, we have that $0 \leq Q_r^\pi \leq \frac{1}{1-\gamma}$, and $-\frac{1}{1-\gamma} \leq Q_c^\pi \leq \frac{1}{1-\gamma}$. We let $V_{\max} = \frac{1}{1-\gamma}$ to simplify the notation. We further write

$$J_r(\pi) := (1 - \gamma) \mathbb{E}_{s \sim \rho} [V_r^\pi(s)], \quad J_c(\pi) := (1 - \gamma) \mathbb{E}_{s \sim \rho} [V_c^\pi(s)]$$

to represent the normalized average reward/cost value of policy π . In addition, we use $d^\pi(s, a)$ to denote the normalized and discounted state-action occupancy measure of the policy π :

$$d^\pi(s, a) := (1 - \gamma) \mathbb{E} \left[\sum_{t=0}^{\infty} \gamma^t \mathbf{1}(s_t = s, a_t = a) \mid a_t \sim \pi(\cdot | s_t) \right],$$

where $\mathbf{1}(\cdot)$ is the indicator function. We also use $d^\pi(s) = \sum_{a \in \mathcal{A}} d^\pi(s, a)$ to denote the discounted state occupancy and we use \mathbb{E}_π as a shorthand of $\mathbb{E}_{(s,a) \sim d^\pi}[\cdot]$ or $\mathbb{E}_{s \sim d^\pi}[\cdot]$ to denote the expectation with respect to d^π . Thus The objective in safe RL for an agent is to find a policy such that

$$\pi \in \arg \max J_r(\pi) \quad \text{s.t. } J_c(\pi) \leq 0. \quad (5)$$

Remark 3.1. For ease of exposition, this paper exclusively focuses on a single constraint. However, it is readily extendable to accommodate multiple constraints.

3.2 Function Approximation

In complex environments, the state space \mathcal{S} is usually very large or even infinite. We assume access to a policy class $\Pi \subseteq (\mathcal{S} \rightarrow \Delta(\mathcal{A}))$ consisting of all candidate policies from which we can search for good policies. We also assume access to a value function class $\mathcal{F} \subseteq (\mathcal{S} \times \mathcal{A} \rightarrow [0, V_{\max}])$ to model the reward Q -functions, and $\mathcal{G} \subseteq (\mathcal{S} \times \mathcal{A} \rightarrow [-V_{\max}, V_{\max}])$ to model the cost Q -functions of candidate policies. We further assume access to a function class $\mathcal{W} \in \{w : \mathcal{S} \times \mathcal{A} \rightarrow [0, B_w]\}$ that represents marginalized importance weights with respect to the offline data distribution. Without loss of generality, we assume that the all-one function is contained in \mathcal{W} .

For a given policy $\pi \in \Pi$, we denote $f(s', \pi) := \sum_{a'} \pi(a' | s') f(s', a')$ for any $s \in \mathcal{S}$. The Bellman operator $\mathcal{T}_r^\pi : \mathbb{R}^{\mathcal{S} \times \mathcal{A}} \rightarrow \mathbb{R}^{\mathcal{S} \times \mathcal{A}}$ for the reward is defined as

$$(\mathcal{T}_r^\pi f)(s, a) := R(s, a) + \gamma \mathbb{E}_{\mathcal{P}(s'|s,a)} [f(s', \pi)],$$

The Bellman operator $\mathcal{T}_c^\pi : \mathbb{R}^{\mathcal{S} \times \mathcal{A}} \rightarrow \mathbb{R}^{\mathcal{S} \times \mathcal{A}}$ for the cost is

$$(\mathcal{T}_c^\pi f)(s, a) := C(s, a) + \gamma \mathbb{E}_{\mathcal{P}(s'|s,a)} [f(s', \pi)].$$

Let $\|\cdot\|_{2,\mu} := \sqrt{\mathbb{E}_\mu[(\cdot)^2]}$ denote the Euclidean norm weighted by distribution μ . We make the following standard assumptions in offline RL setting (Xie et al., 2021; Cheng et al., 2022; Zhu et al., 2023) on the representation power of the function classes:

Assumption 3.2 (Approximate Realizability). Assume there exists $\epsilon_1 \geq 0$, s.t. for any given policy $\pi \in \Pi$, we have $\min_{f \in \mathcal{F}} \max_{\text{admissible } \nu} \|f - T_r^\pi f\|_{2,\nu}^2 \leq \epsilon_1$, and $\min_{f \in \mathcal{F}} \max_{\text{admissible } \nu} \|f - T_c^\pi f\|_{2,\nu}^2 \leq \epsilon_1$, where ν is the state-action distribution of any admissible policy such that $\nu \in \{d^\pi, \forall \pi \in \Pi\}$.

Assumption 3.2 assumes that for any policy $\pi \in \Pi$, Q_r^π and Q_c^π are approximately realizable in \mathcal{F} and \mathcal{G} . When ϵ_1 is small for all admissible ν , we have $f_r \approx Q_r^\pi$, and $f_c \approx Q_c^\pi$. In particular, when $\epsilon_1 = 0$, we have $Q_r^\pi \in \mathcal{F}$, $Q_c^\pi \in \mathcal{G}$ for any policy $\pi \in \Pi$. Note that we do not need Bellman completeness assumption Cheng et al. (2022).

3.3 Offline RL

In offline RL, we assume that the available offline data $\mathcal{D} = \{(s_i, a_i, r_i, c_i, s'_i)\}_{i=1}^N$ consists of N samples. Samples are i.i.d. (which are common assumptions in unconstrained Cheng et al. (2022), as well as constrained setting Hong et al. (2024)), and the distribution of each tuple (s, a, r, c, s') is specified by a distribution $\mu \in \Delta(\mathcal{S} \times \mathcal{A})$, which is also the discounted visitation probability of a behavior policy (also denoted by μ for simplicity). In particular, $(s, a) \sim \mu$, $r = R(s, a)$, $c = C(s, a)$, $s' \sim \mathcal{P}(\cdot|s, a)$. We use $a \sim \mu(\cdot|s)$, to denote that the action is drawn using the behavior policy and $(s, a, s') \sim \mu$ to denote that $(s, a) \sim \mu$, and $s' \sim \mathcal{P}(\cdot|s, a)$.

For a given policy π , we define the marginalized importance weights $w^\pi(s, a) := \frac{d^\pi(s, a)}{\mu(s, a)}$ which is the ratio between the discounted state-action occupancy of π and the data distribution μ . This ratio can be used to measure the concentrability of the data coverage (Xie and Jiang, 2020; Zhan et al., 2022; Rashidinejad et al., 2022; Ozdaglar et al., 2023; Lee et al., 2021).

In this paper we study offline RL with access to a dataset with limited coverage. The coverage of a policy π is the dataset can be measured by the weighted ℓ_2 single policy concentrability coefficient (Zhu et al., 2023; Yin and Wang, 2021; Uehara et al., 2024; Hong et al., 2024):

Definition 3.3 (ℓ_2 Concentrability). Given a policy π , define $C_{\ell_2}^\pi = \|w^\pi\|_{2,\mu} = \|d^\pi/\mu\|_{2,\mu}$.

Remark 3.4. The definition here is much weaker than the **all policy** concentrability used in offline RL (Chen and Jiang, 2019) and safe offline RL (Le et al., 2019; Hong et al., 2024), which requires the ratio $\frac{d^\pi(s, a)}{\mu(s, a)}$ to be bounded for all $s \in \mathcal{S}$ and $a \in \mathcal{A}$ and **all** policies π . In particular, the all-policy concentrability assumption essentially requires the dataset to have full coverage of all policies ((nearly all the state action pairs). This requirement is often violated in practical scenarios. This requirement is even impossible to meet in safe offline RL because it would require collecting data from **every** dangerous state and actions, which clearly is impractical.

In the following lemma, we compare two variants of single-policy concentrability definition with the ℓ_2 defined in Definition 3.3.

Lemma 1 (Restate Proposition 2.1 in Zhu et al. (2023)). Define the ℓ_∞ single policy concentrability (Rashidinejad et al., 2021) as $C_{\ell_\infty}^\pi = \|d^\pi/\mu\|_\infty$ and the Bellman-consistent single-policy concentrability (Xie et al., 2021) as $C_{Bellman}^\pi = \max_{f \in \mathcal{F}} \frac{\|f - \mathcal{T}^\pi f\|_{2,d^\pi}^2}{\|f - \mathcal{T}^\pi f\|_{2,\mu}^2}$ (\mathcal{T} could be \mathcal{T}_r or \mathcal{T}_c in our setting) Then, it always holds $(C_{\ell_2}^\pi)^2 \leq C_{\ell_\infty}^\pi$, $C_{\ell_2}^\pi \leq C_{\ell_\infty}^\pi$ and there exist offline RL instances where $(C_{\ell_2}^\pi)^2 \leq C_{Bellman}^\pi$, $C_{\ell_2}^\pi \leq C_{Bellman}^\pi$.

Remark 3.5. It is easy to observe that the ℓ_2 variant is bounded by ℓ_∞ and $C_{Bellman}^\pi$ under some cases. There is an example (Example 1) in Zhu et al. (2023) showing that $C_{\ell_2}^\pi$ is bounded by a constant $\sqrt{2}$ while $C_{\ell_\infty}^\pi$ could be arbitrarily large. For the case when the function class \mathcal{F} is highly expressive, $C_{Bellman}^\pi$ could be close to $C_{\ell_\infty}^\pi$ and thus possibly larger than $C_{\ell_2}^\pi$. Intuitively, $C_{\ell_2}^\pi$ implies that only $\mathbb{E}_{d^\pi}[w^\pi(s, a)]$ is bounded, rather, $w^\pi(s, a)$ is bounded for all (s, a) in ℓ_∞ concentrability bound.

Given the definition of the concentrability, we make the following assumption on the weight function class \mathcal{W} and a single-policy realizability:

Assumption 3.6 (Boundedness of \mathcal{W} in ℓ_2 norm). For all $w \in \mathcal{W}$, assume that $\|w\|_{2,\mu} \leq C_{\ell_2}^*$.

Assumption 3.7 (Single-policy realizability of w^π). For some policy π that we would like to compete with, assume that $w^\pi \in \mathcal{W}$.

In this paper, we want to study the robust policy improvement on any reference policy, then we assume that we are provided a reference policy π_{ref} . Note that in many applications (e.g., scheduling,

networking) we indeed have a reference policy. We want that while applying a sophisticated RL policy it should do better and be safe as well. This is one of the main motivations behind this assumption.

Assumption 3.8 (Reference Policy). We assume access to a reference policy $\pi_{\text{ref}} \in \Pi$, which can be queried at any state.

In many applications such as networking, scheduling, and control problems, there are existing good enough reference policies. In these cases, a robust and safe policy improvement over these reference policies has practical value. If π_{ref} is not provided, we can simply run a behavior cloning on the offline data to extract the behavior policy as π_{ref} accurately, as long as the size of the offline data set is large enough. More discussion can be found in Section C in the Appendix.

4 Actor-Critic with Importance Weighted Bellman Error

Our algorithm design builds upon the constrained actor-critic method, in which we iteratively optimize a policy and improve the policy based on the evaluation of reward and cost. Consider the following actor-critic approach for solving the optimization problem (5):

$$\mathbf{Actor:} \hat{\pi}^* \in \arg \max_{\pi \in \Pi} f_r^\pi(s_0, \pi) \quad s.t. \quad f_c^\pi(s_0, \pi) \leq 0$$

$$\mathbf{Critic:} f_r^\pi \in \arg \min_{f \in \mathcal{F}} \mathbb{E}_\mu [((f - \mathcal{T}_r f)(s, a))^2], \quad f_c^\pi \in \arg \min_{f \in \mathcal{G}} \mathbb{E}_\mu [((f - \mathcal{T}_c f)(s, a))^2],$$

where we assume that s_0 is a fixed initial state, and $f_r(s, \pi) = \sum_{a \in \mathcal{A}} \pi(a|s) f_r(s, a)$, $f_c(s, \pi) = \sum_{a \in \mathcal{A}} \pi(a|s) f_c(s, a)$. The policy is optimized by maximizing the reward q function f_r while ensuring that f_c satisfies the constraint, and the two functions are trained by minimizing the Bellman error. However, this formulation has several disadvantages. 1) It cannot handle insufficient data coverage, which may fail to provide an accurate estimation of the policy for unseen states and actions. 2) It cannot guarantee robust policy improvement. 3) The actor training step is computationally intractable especially when the policy space is extremely large.

To address the insufficient data coverage issue, as mentioned in [Xie et al. \(2021\)](#) the critic can include a Bellman-consistent pessimistic evaluation of π , which selects the most pessimistic function that approximately satisfies the Bellman equation, which is called absolute pessimism. Then later as indicated by [Cheng et al. \(2022\)](#), instead of using an absolute pessimism, a relative pessimism approach by considering competing to the behavior policy can obtain a robust improvement over the behavior policy. However, this kind of approach can only achieve a suboptimal statistical rate of $N^{1/3}$, and fails to achieve the optimal statistical rate of $1/\sqrt{N}$, then later a weighted average Bellman error ([Uehara et al., 2020; Xie and Jiang, 2020; Zhu et al., 2023](#)) could be treated as one possible solution for improving the order. We remark here that all the discussions here are for the traditional *unconstrained* offline RL. Regarding safety, *no existing efficient algorithms in safe offline RL have theoretically demonstrated* the property of robust policy improvement with optimal statistical rate.

Can Primal-dual based approaches achieve result using only single policy coverability?: The most commonly used approach for addressing safe RL problems is primal-dual optimization ([Efroni et al., 2020; Altman, 1999](#)). As shown in current offline safe RL literature ([Hong et al., 2024; Le et al., 2019](#)), the policy can be optimized by maximizing a new unconstrained “reward” Q -function $f_r^\pi(s_0, \pi) - \lambda f_c^\pi(s_0, \pi)$ where λ is a dual variable. Then, the dual-variable can be tuned by taking gradient descent step. As we discussed in the introduction, all these require **all** policy concentrability which is not practical especially for safe RL. Important question is whether all policy concentrability assumption can be relaxed. Note that primal-dual algorithm relies on solving the min-max problem $\min_\lambda \max_\pi f_r^\pi(s_0, \pi) - \lambda f_c^\pi(s_0, \pi)$. Recent result ([Cui and Du, 2022](#)) shows that single policy concentrability assumption is *not* enough for offline min-max game. Hence, we *conjecture* that using the primal-dual method we can not relax the all policy concentrability assumption. Intuitively, the primal-dual based method ([Hong et al., 2024](#)) rely on bounding the regret in dual domain $\sum_k (\lambda_k - \lambda^*) (f_c^{\pi_k} - 0)$, hence, all the policies $\{\pi_k\}_{k=1}^K$ encountered throughout the iteration must be supported by the dataset to evaluate the dual value $\lambda^* (f_c^{\pi_k} - 0)$ where λ^* is the optimal dual value.

Our novelty: In contrast, we propose an aggression-limited objective function $f_r(s_0, \pi) - \lambda \cdot [f_c(s_0, \pi)]_+$ to control aggressive policies, where $\{\cdot\}_+ := \max\{\cdot, 0\}$. The high-level intuition behind this aggression-limited objective function is that by appropriately selecting a λ (usually large enough), we penalize all the policies that are not safe. As a result, the policy that maximizes the

objective function is the optimal safe policy. This formulation is fundamentally different from the traditional primal-dual approach as it does not require dual-variable tuning, and thus, does not require all policy concentrability. In particular, we only need to bound the primal domain regret which can be done as long as the reference policy is covered by the dataset similar to the unconstrained setup.

Combining all the previous ideas together provides the design of our main algorithm named WSAC (**W**eighted **S**afe **A**ctor-**C**ritic). In Section 5, we will provide theoretical guarantees of WSAC and discuss its advantages over existing approaches in offline safe RL. WSAC aims to solve the following optimization problem:

$$\begin{aligned} \hat{\pi}^* \in \arg \max_{\pi \in \Pi} \mathcal{L}_\mu(\pi, f_r^\pi) - \lambda \{\mathcal{L}_\mu(\pi, f_c^\pi)\}_+ \\ \text{s.t. } f_r^\pi \in \arg \min_{f_r \in \mathcal{F}} \mathcal{L}_\mu(\pi, f_r) + \beta \mathcal{E}_\mu(\pi, f_r), \quad f_c^\pi \in \arg \min_{f_c \in \mathcal{G}} -\lambda \mathcal{L}_\mu(\pi, f_c) + \beta \hat{\mathcal{E}}_\mu(\pi, f_c), \end{aligned} \quad (6)$$

where $\mathcal{L}_\mu(\pi, f) := \mathbb{E}_\mu[f(s, \pi) - f(s, a)]$, and $\mathcal{E}_\mu(\pi, f) := \max_{w \in \mathcal{W}} |\mathbb{E}_\mu[w(s, a)((f - T_r^\pi f)(s, a))]|$, $\hat{\mathcal{E}}_\mu(\pi, f) := \max_{w \in \mathcal{W}} |\mathbb{E}_\mu[w(s, a)((f - T_c^\pi f)(s, a))]|$. This formulation can also be treated as a Stackelberg game (Von Stackelberg, 2010) or bilevel optimization problem. We penalize the objective function only when the approximate cost Q -function f_c^π of the policy π is more perilous than the behavior policy ($f_c^\pi(s, \pi) \geq f_c^\pi(s, a)$) forcing our policy to be as safe as the behavior policy. Maximization over w in for training the two critics can ensure that the Bellman error is small when averaged over measure $\mu \cdot w$ for any $w \in \mathcal{W}$, which turns out to be sufficient to control the suboptimality of the learned policy.

In the following theorem, we show that the solution of the optimization problem (6) is not worse than the behavior policy μ in both performance and safety for any $\beta \geq 0, \lambda > 0$ than the policy μ under Assumption 3.2 with $\epsilon_1 = 0$.

Theorem 4.1. *Assume that Assumption 3.2 holds with $\epsilon_1 = 0$, and the behavior policy $\mu \in \Pi$, then for any $\beta \geq 0, \lambda > 0$ we have $J_r(\hat{\pi}^*) \geq J_r(\mu)$, and $\{J_c(\hat{\pi}^*)\}_+ \leq \{J_c(\mu)\}_+ + \frac{1}{\lambda}$.*

The result in Theorem 4.1 shows that by selecting λ large enough, for any $\beta \geq 0$, the solution can achieve better performance than the behavior policy while maintaining safety that is arbitrarily close to that of the behavior policy. The Theorem verifies the design of our framework which has the potential to have a robust safe improvement.

In the next section, we will introduce our main algorithm WSAC and provide its theoretical guarantees.

5 Theoretical Analysis of WSAC

5.1 Main Algorithm

In this section, we present the theoretical version of our new model-free offline safe RL algorithm WSAC. Since we only have access to a dataset \mathcal{D} instead of the data distribution. WSAC solves an empirical version of (6):

$$\begin{aligned} \hat{\pi} \in \arg \max_{\pi \in \Pi} \mathcal{L}_\mathcal{D}(\pi, f_r^\pi) - \lambda \{\mathcal{L}_\mathcal{D}(\pi, f_c^\pi)\}_+ \\ \text{s.t. } f_r^\pi \in \arg \min_{f_r \in \mathcal{F}} \mathcal{L}_\mathcal{D}(\pi, f_r) + \beta \mathcal{E}_\mathcal{D}(\pi, f_r), \quad f_c^\pi \in \arg \min_{f_c \in \mathcal{G}} -\lambda \mathcal{L}_\mathcal{D}(\pi, f_c) + \beta \hat{\mathcal{E}}_\mathcal{D}(\pi, f_c), \end{aligned} \quad (7)$$

where

$$\begin{aligned} \mathcal{L}_\mathcal{D}(\pi, f) &:= \mathbb{E}_\mathcal{D}[f(s, \pi) - f(s, a)] \\ \mathcal{E}_\mathcal{D}(\pi, f) &:= \max_{w \in \mathcal{W}} |\mathbb{E}_\mathcal{D}[w(s, a)(f(s, a) - r - \gamma f(s', \pi))]| \\ \hat{\mathcal{E}}_\mathcal{D}(\pi, f) &:= \max_{w \in \mathcal{W}} |\mathbb{E}_\mathcal{D}[w(s, a)(f(s, a) - c - \gamma f(s', \pi))]|. \end{aligned} \quad (8)$$

As shown in Algorithm 1, at each iteration, WSAC selects f_r^k maximally pessimistic and f_c^k maximally optimistic for the current policy π_k with a weighted regularization on the estimated Bellman error for reward and cost, respectively (Line 4 and 6) to address the worse cases within reasonable range. In order to achieve a safe robust policy improvement, the actor then applies a no-regret policy optimization oracle to update the policy π_{k+1} by optimizing the aggression-limited objective function compared with the reference policy (Line 7) $f_r^k(s, a) - \lambda \{f_c^k(s, a) -$

Algorithm 1 Weighted Safe Actor-Critic (WSAC)

- 1: **Input:** Batch data \mathcal{D} , coefficient β, λ . Value function classes \mathcal{F}, \mathcal{G} , importance weight function class \mathcal{W} , Initialize policy π_1 randomly. Any reference policy π_{ref} . No-regret policy optimization oracle **PO** (Definition 5.1).
- 2: **for** $k = 1, 2, \dots, K$ **do**
- 3: Obtain the reward state-action value function estimation of π_k :
- 4: $f_r^k \leftarrow \arg \min_{f_r \in \mathcal{F}} \mathcal{L}_{\mathcal{D}}(\pi_k, f_r) + \beta \mathcal{E}_{\mathcal{D}}(\pi_k, f_r)$
- 5: Obtain the cost state-action value function estimation of π_k :
- 6: $f_c^k \leftarrow \arg \min_{f_c \in \mathcal{G}} -\lambda \mathcal{L}_{\mathcal{D}}(\pi_k, f_c) + \beta \hat{\mathcal{E}}_{\mathcal{D}}(\pi_k, f_c)$
- 7: Update policy: $\pi_{k+1} \leftarrow \mathbf{PO}(\pi_k, f_r^k(s, a) - \lambda \{f_c^k(s, a) - f_c^k(s, \pi_{\text{ref}})\}_+, \mathcal{D})$. // $\mathcal{L}_{\mathcal{D}}, \mathcal{E}_{\mathcal{D}}, \hat{\mathcal{E}}_{\mathcal{D}}$ are defined in (5.1)
- 8: **end for**
- 9: **Output:** $\bar{\pi} = \text{Unif}(\pi_1, \dots, \pi_K)$. // Uniformly mix π_1, \dots, π_K

$f_c^k(s, \pi_{\text{ref}})\}_+$. Our algorithm is very computationally efficient and tractable compared with existing approaches (Hong et al., 2024; Le et al., 2019), since we do not need another inner loop for optimizing the dual variable with an additional online algorithm or offline policy evaluation oracle. The policy improvement process relies on a no-regret policy optimization oracle, a technique commonly employed in offline RL literature (Zhu et al., 2023; Cheng et al., 2022; Hong et al., 2024; Zhu et al., 2023). Extensive literature exists on such methodologies. For instance, approaches like soft policy iteration (Pirotta et al., 2013) and algorithms based on natural policy gradients (Kakade, 2001; Agarwal et al., 2021) can function as effective no-regret policy optimization oracles. We now formally define the oracle:

Definition 5.1 (No-regret policy optimization oracle). An algorithm **PO** is called a no-regret policy optimization oracle if for any sequence of functions f^1, \dots, f^K with $f^k : \mathcal{S} \times \mathcal{A} \rightarrow [0, V_{\max}]$, $\forall k \in [K]$. The policies π_1, \dots, π_K produced by the oracle **PO** satisfy that for any policy $\pi \in \Pi$:

$$\epsilon_{opt}^{\pi} \triangleq \frac{1}{K} \sum_{k=1}^K \mathbb{E}_{\pi} [f^k(s, \pi) - f^k(s, \pi_k)] = o(1) \quad (9)$$

There indeed exist many methods that can serve as the no-regret oracle, for example, the mirror-descent approach (Geist et al., 2019) or the natural policy gradient approach (Kakade, 2001) of the form $\pi_{k+1}(a|s) \propto \pi_k(a|s) \exp(\eta f^k(s, a))$ with $\eta = \sqrt{\frac{\log |\mathcal{A}|}{2V_{\max}^2 K}}$ (Even-Dar et al., 2009; Agarwal et al., 2021). In the following define ϵ_{opt}^{π} as the error generated from the oracle **PO** by considering $f_r^k(s, a) - \lambda \{f_c^k(s, a) - f_c^k(s, \pi)\}_+$ as the sequence of functions in Definition 5.1, then we have the following guarantee.

Lemma 2. Applying a no-regret oracle **PO** for K episodes with $(f_r^k(s, a) - \lambda \{f_c^k(s, a) - f_c^k(s, \pi)\}_+)$ for an arbitrary policy π , can guarantee

$$\frac{1}{K} \sum_{k=1}^K \mathbb{E}_{\pi} [f_r^k(s, \pi) - f_r^k(s, \pi_k)] \leq \epsilon_{opt}^{\pi} \quad (10)$$

$$\frac{1}{K} \sum_{k=1}^K \mathbb{E}_{\pi} [\{f_c^k(s, \pi_k) - f_c^k(s, \pi)\}_+] \leq \epsilon_{opt}^{\pi} + \frac{V_{\max}}{\lambda}. \quad (11)$$

Lemma 2 establishes that the policy outputted by **PO** with considering the aggression-limited “reward” can have a strong guarantee on the performance of both reward and cost when λ is large enough., which is comparable with any competitor policy. This requirement is critical to achieving the performance guarantee of our algorithm and the safe and robust policy improvement. The detailed proof is deferred to Appendix B.2 due to page limit.

5.2 Theoretical Guarantees

We are now ready to provide the theoretical guarantees of WSAC Algorithm 1. The complete proof is deferred to Appendix B.3.

Theorem 5.2 (Main Theorem). *Under Assumptions 3.2 and 3.6, let the reference policy $\pi_{ref} \in \Pi$ in Algorithm 1 be any policy satisfying Assumption 3.7, then with probability at least $1 - \delta$,*

$$J_r(\pi_{ref}) - J_r(\bar{\pi}) \leq \mathcal{O}\left(\epsilon_{stat} + C_{\ell_2}^* \sqrt{\epsilon_1}\right) + \epsilon_{opt}^{\pi_{ref}} \quad (12)$$

$$J_c(\bar{\pi}) - J_c(\pi_{ref}) \leq \mathcal{O}\left(\epsilon_{stat} + C_{\ell_2}^* \sqrt{\epsilon_1}\right) + \epsilon_{opt}^{\pi_{ref}} + \frac{V_{\max}}{\lambda}, \quad (13)$$

where $\epsilon_{stat} := V_{\max} C_{\ell_2}^* \sqrt{\frac{\log(|\mathcal{F}||\Pi||W|/\delta)}{N}} + \frac{V_{\max} B_w \log(|\mathcal{F}||\Pi||W|/\delta)}{N}$, and $\bar{\pi}$ is the policy returned by Algorithm 1 with $\beta = 2$ and π_{ref} as input.

Remark 5.3. When $\epsilon_1 = 0$, i.e., no model misspecification, which states that the true value function belongs to the function class being used to approximate it (the function class is right enough), let π_{ref} be the optimal policy, the results in Theorem 5.2 achieve an optimal dependence statistical rate of $\frac{1}{\sqrt{N}}$ (for large k), which matches the best existing results. Our algorithm is both statistically optimal and computationally efficient with only **single-policy** assumption rather relying much stronger assumptions of **all** policy concentrability [Hong et al. \(2024\)](#); [Le et al. \(2019\)](#). Hence, if the behavior policy or the reference policy is safe, our result indicates that the policy returned by our algorithm will also be safe (nearly). *Such a guarantee was missing in the existing literature.*

Remark 5.4. We also do not need a completeness assumption, which requires that for any $f \in \mathcal{F}$ or \mathcal{G} and $\pi \in \Pi$, it approximately holds that $\mathcal{T}_r f \in \mathcal{F}$, $\mathcal{T}_c f \in \mathcal{F}$ as required in [Xie et al. \(2021\)](#); [Chen et al. \(2022b\)](#). They need this assumption to address over-estimation issues caused by the ℓ_2 square Bellman error, but our algorithm can get rid of the strong assumption by using a weighted Bellman error which is a simple and unbiased estimator.

Remark 5.5. Our algorithm can compete with any reference policy $\pi_{ref} \in \Pi$ as long as $w^{\pi_{ref}} = d^{\pi_{ref}}/\mu$ is contained in \mathcal{W} . The importance ratio of the behavior policy is $w^\mu = d^\mu/\mu = \mu/\mu = 1$ which always satisfies this condition, implying that our algorithm can have a safe robust policy improvement (in Theorem 5.6 discussed below).

5.3 A Safe Robust Policy Improvement

A Robust policy improvement (RPI) ([Cheng et al., 2022](#); [Zhu et al., 2023](#); [Bhardwaj et al., 2024](#)) refers to the property of an offline RL algorithm that the offline algorithm can learn to improve over the behavior policy, using a wide range of hyperparameters. In this paper, we introduce the property of Safe Robust policy improvement (SRPI) such that the offline algorithm can learn to improve over the behavior policy in both return and safety, using a wide range of hyperparameters. In the following Theorem 5.6 we show that as long as the hyperparameter $\beta = o(\sqrt{N})$, our algorithm can, with high probability, produce a policy with vanishing suboptimality compared to the behavior policy.

Theorem 5.6 (SRPI). *Under Assumptions 3.2 and 3.6, then with probability at least $1 - \delta$,*

$$J_r(\mu) - J_r(\bar{\pi}) \leq \mathcal{O}\left(\epsilon_{stat}^\pi + C_{\ell_2}^* \sqrt{\epsilon_1}\right) + \epsilon_{opt}^\mu \quad (14)$$

$$J_c(\bar{\pi}) - J_c(\mu) \leq \mathcal{O}\left(\epsilon_{stat}^\pi + C_{\ell_2}^* \sqrt{\epsilon_1}\right) + \epsilon_{opt}^\mu + \frac{V_{\max}}{\lambda}, \quad (15)$$

where $\epsilon_{stat} := V_{\max} C_{\ell_2}^* \sqrt{\frac{\log(|\mathcal{F}||\Pi||W|/\delta)}{N}} + \frac{V_{\max} B_w \log(|\mathcal{F}||\Pi||W|/\delta)}{N}$, and $\bar{\pi}$ is the policy returned by Algorithm 1 with $\beta \geq 0$ and μ as input.

The detailed proofs are deferred to Appendix B.4.

6 Experiments

6.1 WSAC-Practical Implementation

We introduce a deep RL implementation of WSAC in Algorithm 2 (in Appendix), following the key structure of its theoretical version (Algorithm 1). The reward, cost Q -functions f_r, f_c and the policy network π are all parameterized by neural networks. The critic losses (line 4) $l_{reward}(f_r)$ and

Table 2: The normalized reward and cost of WSAC and other baselines. The Average line shows the average situation in various environments. The cost threshold is 1. Gray: Unsafe agent whose normalized cost is greater than 1. Blue: Safe agent with best performance

Environment	BC		Safe-BC		BCQL		BEARL		CPQ		COptDICE		WSAC	
	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
BallCircle	0.70	0.95	0.61	0.49	0.73	0.82	0.80	1.23	0.62	0.76	0.71	1.13	0.75	0.27
CarCircle	0.57	1.43	0.57	0.65	0.79	1.19	0.84	1.87	0.67	0.28	0.49	1.52	0.68	0.59
PointButton	0.26	1.75	0.12	0.69	0.36	1.76	0.32	1.71	0.43	3.10	0.15	1.92	0.13	0.67
PointPush	0.13	0.67	0.20	1.35	0.16	1.01	0.12	0.90	-0.01	2.39	0.07	1.18	0.07	0.52
Average	0.42	1.20	0.38	0.80	0.51	1.12	0.52	1.43	0.36	1.63	0.36	1.44	0.41	0.51

$l_{cost}(f_c)$ are calculated based on the principles of Algorithm 1, on the minibatch dataset. Optimizing the actor aims to achieve a no-regret optimization oracle, we use a gradient based update on the actor loss (line 5) $l_{actor}(\pi)$. In the implementation we use adaptive gradient descent algorithm ADAM (Kingma and Ba, 2015) for updating two critic networks and the actor network. Algorithm follows standard two-timescale first-order algorithms (Fujimoto et al., 2018; Haarnoja et al., 2018) with a fast learning rate η_{fast} on update critic networks and a slow learning rate η_{slow} for updating the actor.

6.2 Simulations

We present a scalable deep RL version of WSAC in Algorithm 2, following the principles of Algorithm 1. We evaluate WSAC and consider Behavior Cloning (BC), safe Behavior Cloning (Safe-BC), Batch-Constrained deep Q-learning with Lagrangian PID (BCQL) Fujimoto et al. (2019); Stooke et al. (2020), bootstrapping error accumulation reduction with Lagrangian PID (BEARL) Kumar et al. (2019); Stooke et al. (2020), Constraints Penalized Q-learning (CPQ) Xu et al. (2022) and one of the state-of-the-art algorithms, COptDICE (Lee et al., 2022) as baselines.

We study several representative environments and focus on presenting “BallCircle”. In BallCircle, it requires the ball on a circle in a clockwise direction without leaving the safety zone defined by the boundaries as proposed by Achiam et al. (2017). The ball is a spherical-shaped agent which can freely move on the xy-plane. The reward is dense and increases by the car’s velocity and by the proximity towards the boundary of the circle. The cost is incurred if the agent leaves the safety zone defined by the boundaries.

We use the offline dataset from Liu et al. (2019), where the corresponding expert policy are used to interact with the environments and collect the data. To better illustrate the results, we normalize the reward and cost. Our simulation results are reported in Table 2, we observe that WSAC can guarantee that all the final agents are safe, which is most critical in safe RL literature. Even in challenging environments such as PointButton, which most baselines fail to learn safe policies. WSAC has the best results in 3 of the environments. Moreover, WSAC outperforms all the baselines in terms of the average performance, demonstrating its ability to learn a safe policy by leveraging an offline dataset. The simulation results verify our theoretical findings. We also compared WSAC with all the baselines in the case where the cost limits are different, WSAC still outperforms all the other baselines and ensures a safe policy. We further include simulations to investigate the contribution of each component of our algorithm, including the weighted Bellman regularizer, the aggression-limited objective, and the no-regret policy optimization which together guarantee the theoretical results. More details and discussions are deferred to the Appendix D due to page limit.

7 Conclusion

In this paper, we explore the problem of offline Safe-RL with a single policy data coverage assumption. We propose a novel algorithm, WSAC, which, for the first time, is proven to guarantee the property of safe robust policy improvement. WSAC is able to outperform any reference policy, including the behavior policy, while maintaining the same level of safety across a broad range of hyperparameters. Our simulation results demonstrate that WSAC outperforms existing state-of-the-art offline safe-RL algorithms. Interesting future work includes combining WSAC with online exploration with safety guarantees and extending the approach to multi-agent settings to handle coupled constraints.

References

- Achiam, J., Held, D., Tamar, A., and Abbeel, P. (2017). Constrained policy optimization. In *Int. Conf. Machine Learning (ICML)*, volume 70, pages 22–31. JMLR.
- Agarwal, A., Kakade, S. M., Lee, J. D., and Mahajan, G. (2021). On the theory of policy gradient methods: Optimality, approximation, and distribution shift. *Journal of Machine Learning Research*, 22(98):1–76.
- Altman, E. (1999). *Constrained Markov decision processes*, volume 7. CRC Press.
- Bhardwaj, M., Xie, T., Boots, B., Jiang, N., and Cheng, C.-A. (2024). Adversarial model for offline reinforcement learning. *Advances in Neural Information Processing Systems*, 36.
- Chen, F., Zhang, J., and Wen, Z. (2022a). A near-optimal primal-dual method for off-policy learning in cmdp. In *Advances Neural Information Processing Systems (NeurIPS)*, volume 35, pages 10521–10532.
- Chen, J. and Jiang, N. (2019). Information-theoretic considerations in batch reinforcement learning. In *Int. Conf. Machine Learning (ICML)*, pages 1042–1051. PMLR.
- Chen, J. and Jiang, N. (2022). Offline reinforcement learning under value and density-ratio realizability: the power of gaps. In *Uncertainty in Artificial Intelligence*, pages 378–388. PMLR.
- Chen, L., Jain, R., and Luo, H. (2022b). Learning infinite-horizon average-reward markov decision process with constraints. In *Int. Conf. Machine Learning (ICML)*, pages 3246–3270. PMLR.
- Chen, L., Lu, K., Rajeswaran, A., Lee, K., Grover, A., Laskin, M., Abbeel, P., Srinivas, A., and Mordatch, I. (2021). Decision transformer: Reinforcement learning via sequence modeling. *Advances in neural information processing systems*, 34:15084–15097.
- Cheng, C.-A., Kolobov, A., and Agarwal, A. (2020). Policy improvement via imitation of multiple oracles. In *Advances Neural Information Processing Systems (NeurIPS)*, volume 33, pages 5587–5598.
- Cheng, C.-A., Xie, T., Jiang, N., and Agarwal, A. (2022). Adversarially trained actor critic for offline reinforcement learning. In *International Conference on Machine Learning*, pages 3852–3878. PMLR.
- Chow, Y., Ghavamzadeh, M., Janson, L., and Pavone, M. (2017). Risk-constrained reinforcement learning with percentile risk criteria. *The Journal of Machine Learning Research*, 18(1):6070–6120.
- Cui, Q. and Du, S. S. (2022). When are offline two-player zero-sum markov games solvable? *Advances in Neural Information Processing Systems*, 35:25779–25791.
- Efroni, Y., Mannor, S., and Pirootta, M. (2020). Exploration-exploitation in constrained MDPs. *arXiv preprint arXiv:2003.02189*.
- Even-Dar, E., Kakade, S. M., and Mansour, Y. (2009). Online markov decision processes. *Mathematics of Operations Research*, 34(3):726–736.
- Fujimoto, S., Meger, D., and Precup, D. (2019). Off-policy deep reinforcement learning without exploration. In *International conference on machine learning*, pages 2052–2062. PMLR.
- Fujimoto, S., van Hoof, H., and Meger, D. (2018). Addressing function approximation error in actor-critic methods. In *Int. Conf. Machine Learning (ICML)*, pages 1582–1591.
- Geist, M., Scherrer, B., and Pietquin, O. (2019). A theory of regularized markov decision processes. In *International Conference on Machine Learning*, pages 2160–2169. PMLR.
- 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 *Int. Conf. Machine Learning (ICML)*, pages 1861–1870.

- Hong, K., Li, Y., and Tewari, A. (2024). A primal-dual-critic algorithm for offline constrained reinforcement learning. In *Int. Conf. Artificial Intelligence and Statistics (AISTATS)*, pages 280–288. PMLR.
- Isele, D., Nakhaei, A., and Fujimura, K. (2018). Safe reinforcement learning on autonomous vehicles. In *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 1–6. IEEE.
- Kakade, S. and Langford, J. (2002). Approximately optimal approximate reinforcement learning. In *Int. Conf. Machine Learning (ICML)*, pages 267–274.
- Kakade, S. M. (2001). A natural policy gradient. In *Advances Neural Information Processing Systems (NeurIPS)*.
- Kingma, D. P. and Ba, J. (2015). Adam: A method for stochastic optimization. In Bengio, Y. and LeCun, Y., editors, *Int. Conf. on Learning Representations (ICLR)*.
- 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.
- Kumar, A., Hong, J., Singh, A., and Levine, S. (2022). Should i run offline reinforcement learning or behavioral cloning? In *Int. Conf. on Learning Representations (ICLR)*.
- Laroche, R., Trichelair, P., and Des Combes, R. T. (2019). Safe policy improvement with baseline bootstrapping. In *International conference on machine learning*, pages 3652–3661. PMLR.
- Le, H., Voloshin, C., and Yue, Y. (2019). Batch policy learning under constraints. In *International Conference on Machine Learning*, pages 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 Meila, M. and Zhang, T., editors, *Int. Conf. Machine Learning (ICML)*, volume 139 of *Proceedings of Machine Learning Research*, pages 6120–6130. PMLR.
- Lee, J., Paduraru, C., Mankowitz, D. J., Heess, N., Precup, D., Kim, K.-E., and Guez, A. (2022). Cop-tidice: Offline constrained reinforcement learning via stationary distribution correction estimation. *arXiv preprint arXiv:2204.08957*.
- Liao, P., Qi, Z., Wan, R., Klasnja, P., and Murphy, S. A. (2022). Batch policy learning in average reward markov decision processes. *Annals of statistics*, 50(6):3364.
- Liu, B., Cai, Q., Yang, Z., and Wang, Z. (2019). Neural trust region/proximal policy optimization attains globally optimal policy. *Advances in neural information processing systems*, 32.
- 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 *International Conference on Machine Learning*, pages 21611–21630. PMLR.
- Ozdoglar, A. E., Pattathil, S., Zhang, J., and Zhang, K. (2023). Revisiting the linear-programming framework for offline rl with general function approximation. In *International Conference on Machine Learning*, pages 26769–26791. PMLR.
- Pirotta, M., Restelli, M., Pecorino, A., and Calandriello, D. (2013). Safe policy iteration. In *Int. Conf. Machine Learning (ICML)*, pages 307–315. PMLR.
- Rajaraman, N., Yang, L., Jiao, J., and Ramchandran, K. (2020). Toward the fundamental limits of imitation learning. *Advances Neural Information Processing Systems (NeurIPS)*, 33:2914–2924.
- Rashidinejad, P., Zhu, B., Ma, C., Jiao, J., and Russell, S. (2021). Bridging offline reinforcement learning and imitation learning: A tale of pessimism. In *Advances Neural Information Processing Systems (NeurIPS)*, volume 34, pages 11702–11716.

- Rashidinejad, P., Zhu, H., Yang, K., Russell, S., and Jiao, J. (2022). Optimal conservative offline rl with general function approximation via augmented lagrangian. *arXiv preprint arXiv:2211.00716*.
- Siegel, N. Y., Springenberg, J. T., Berkenkamp, F., Abdolmaleki, A., Neunert, M., Lampe, T., Hafner, R., Heess, N., and Riedmiller, M. (2020). Keep doing what worked: Behavioral modelling priors for offline reinforcement learning. *arXiv preprint arXiv:2002.08396*.
- Stooke, A., Achiam, J., and Abbeel, P. (2020). Responsive safety in reinforcement learning by pid lagrangian methods. In *Int. Conf. Machine Learning (ICML)*, pages 9133–9143. PMLR.
- Uehara, M., Huang, J., and Jiang, N. (2020). Minimax weight and q-function learning for off-policy evaluation. In *International Conference on Machine Learning*, pages 9659–9668. PMLR.
- Uehara, M., Kallus, N., Lee, J. D., and Sun, W. (2024). Offline minimax soft-q-learning under realizability and partial coverage. *Advances in Neural Information Processing Systems*, 36.
- Uehara, M. and Sun, W. (2021). Pessimistic model-based offline reinforcement learning under partial coverage. *arXiv preprint arXiv:2107.06226*.
- Von Stackelberg, H. (2010). *Market structure and equilibrium*. Springer Science & Business Media.
- Wang, L., Cai, Q., Yang, Z., and Wang, Z. (2019). Neural policy gradient methods: Global optimality and rates of convergence. *arXiv preprint arXiv:1909.01150*.
- Wu, R., Zhang, Y., Yang, Z., and Wang, Z. (2021). Offline constrained multi-objective reinforcement learning via pessimistic dual value iteration. In *Advances Neural Information Processing Systems (NeurIPS)*, volume 34, pages 25439–25451.
- Xie, T., Cheng, C.-A., Jiang, N., Mineiro, P., and Agarwal, A. (2021). Bellman-consistent pessimism for offline reinforcement learning. *Advances in neural information processing systems*, 34:6683–6694.
- Xie, T. and Jiang, N. (2020). Q^* approximation schemes for batch reinforcement learning: A theoretical comparison. In *Conference on Uncertainty in Artificial Intelligence*, pages 550–559. PMLR.
- Xie, T. and Jiang, N. (2021). Batch value-function approximation with only realizability. In *Int. Conf. Machine Learning (ICML)*, pages 11404–11413. PMLR.
- Xu, H., Zhan, X., and Zhu, X. (2022). Constraints penalized q-learning for safe offline reinforcement learning. In *AAAI Conf. Artificial Intelligence*, volume 36, pages 8753–8760.
- Yin, M. and Wang, Y.-X. (2021). Towards instance-optimal offline reinforcement learning with pessimism. *Advances in neural information processing systems*, 34:4065–4078.
- Zhan, W., Huang, B., Huang, A., Jiang, N., and Lee, J. (2022). Offline reinforcement learning with realizability and single-policy concentrability. In *Proc. Conf. Learning Theory (COLT)*, pages 2730–2775. PMLR.
- Zhang, J., Koppel, A., Bedi, A. S., Szepesvari, C., and Wang, M. (2020). Variational policy gradient method for reinforcement learning with general utilities. *Advances in Neural Information Processing Systems*, 33:4572–4583.
- Zheng, Y., Li, J., Yu, D., Yang, Y., Li, S. E., Zhan, X., and Liu, J. (2024). Safe offline reinforcement learning with feasibility-guided diffusion model. *arXiv preprint arXiv:2401.10700*.
- Zhu, H., Rashidinejad, P., and Jiao, J. (2023). Importance weighted actor-critic for optimal conservative offline reinforcement learning. *arXiv preprint arXiv:2301.12714*.

Supplementary Material

A Auxiliary Lemmas

In the following, we first provide several lemmas which are useful for proving our main results.

Lemma 3. *With probability at least $1 - \delta$, for any $f_r \in \mathcal{F}$, $f_c \in \mathcal{G}$, $\pi \in \Pi$ and $w \in \mathcal{W}$, we have*

$$\begin{aligned} & \left| \mathbb{E}_\mu[(f_r - \mathcal{T}_r^\pi f)w] \right| - \left| \frac{1}{N} \sum_{(s,a,r,s')} w(s,a)(f_r(s,a) - r - \gamma f_r(s',\pi)) \right| \\ & \leq \mathcal{O} \left(V_{\max} \sqrt{\frac{\log(|\mathcal{F}||\Pi||\mathcal{W}|/\delta)}{N}} + \frac{V_{\max} B_w \log(|\mathcal{F}||\Pi||\mathcal{W}|/\delta)}{N} \right) \end{aligned} \quad (16)$$

$$\begin{aligned} & \left| \mathbb{E}_\mu[(f_c - \mathcal{T}_c^\pi f)w] \right| - \left| \frac{1}{N} \sum_{(s,a,r,s')} w(s,a)(f_c(s,a) - c - \gamma f_c(s',\pi)) \right| \\ & \leq \mathcal{O} \left(V_{\max} \sqrt{\frac{\log(|\mathcal{G}||\Pi||\mathcal{W}|/\delta)}{N}} + \frac{V_{\max} B_w \log(|\mathcal{G}||\Pi||\mathcal{W}|/\delta)}{N} \right) \end{aligned} \quad (17)$$

The proofs can be found in Lemma 4 in [Zhu et al. \(2023\)](#).

Lemma 4. *With probability at least $1 - 2\delta$, for any $f_r \in \mathcal{F}$, $f_c \in \mathcal{G}$ and $\pi \in \Pi$, we have*

$$|\mathcal{E}_\mu(\pi, f_r) - \mathcal{E}_\mathcal{D}(\pi, f_r)| \leq \epsilon_{stat} \quad (18)$$

$$|\mathcal{E}_\mu(\pi, f_c) - \mathcal{E}_\mathcal{D}(\pi, f_c)| \leq \epsilon_{stat}, \quad (19)$$

where $\epsilon_{stat} := V_{\max} C_{\ell_2}^* \sqrt{\frac{\log(|\mathcal{F}||\mathcal{G}||\Pi||\mathcal{W}|/\delta)}{N}} + \frac{V_{\max} B_w \log(|\mathcal{F}||\mathcal{G}||\Pi||\mathcal{W}|/\delta)}{N}$.

Proof. Condition on the high probability event in Lemma 3, for any $f_r \in \mathcal{F}$, $f_c \in \mathcal{G}$, $\pi \in \Pi$, define

$$w_{\pi,f}^* = \arg \max_{w \in \mathcal{W}} \mathcal{E}_\mu(\pi, f_r) = \arg \max_{w \in \mathcal{W}} |\mathbb{E}_\mu[w(s,a)(f_r - \mathcal{T}_r^\pi f_r)(s,a)]|$$

and define

$$\hat{w}_{\pi,f_r} = \arg \max_{w \in \mathcal{W}} \mathcal{E}_\mathcal{D}(\pi, f_r) = \arg \max_{w \in \mathcal{W}} \left| \frac{1}{N} \sum_{(s,a,r,s') \in \mathcal{D}} w(s,a)(f_r(s,a) - r - \gamma f_r(s',\pi)) \right|.$$

Then we can have

$$\begin{aligned} & \mathcal{T}_\mu(\pi, f_r) - \mathcal{E}_\mathcal{D}(\pi, f_r) \\ & = |\mathbb{E}_\mu[w_{\pi,f_r}^*(s,a)(f_r - \mathcal{T}_r^\pi f_r)(s,a)]| - \left| \frac{1}{N} \sum_{(s,a,r,s')} \hat{w}_{\pi,f_r}(s,a)(f_r(s,a) - r - \gamma f_r(s',\pi)) \right| \\ & = |\mathbb{E}_\mu[w_{\pi,f_r}^*(s,a)(f_r - \mathcal{T}_r^\pi f_r)(s,a)]| - |\mathbb{E}_\mu[\hat{w}_{\pi,f_r}(s,a)(f_r - \mathcal{T}_r^\pi f_r)(s,a)]| \\ & \quad + |\mathbb{E}_\mu[\hat{w}_{\pi,f_r}(s,a)(f_r - \mathcal{T}_r^\pi f_r)(s,a)]| - \left| \frac{1}{N} \sum_{(s,a,r,s')} \hat{w}_{\pi,f_r}(s,a)(f_r(s,a) - r - \gamma f_r(s',\pi)) \right| \\ & \geq 0 - \epsilon_{stat} = -\epsilon_{stat}, \end{aligned}$$

where the inequality is true by using the definition of w_{π,f_r}^* and Lemma 3. Thus

$$\begin{aligned} & \mathcal{E}_\mu(\pi, f_r) - \mathcal{E}_\mathcal{D}(\pi, f_r) \\ & = |\mathbb{E}_\mu[w_{\pi,f_r}^*(s,a)(f_r - \mathcal{T}_r^\pi f_r)(s,a)]| - \left| \frac{1}{N} \sum_{(s,a,r,s')} w_{\pi,f_r}^*(s,a)(f_r(s,a) - r - \gamma f_r(s',\pi)) \right| \\ & \quad + \left| \frac{1}{N} \sum_{(s,a,r,s')} w_{\pi,f_r}^*(s,a)(f_r(s,a) - r - \gamma f_r(s',\pi)) \right| \\ & \quad - \left| \frac{1}{N} \sum_{(s,a,r,s')} \hat{w}_{\pi,f_r}(s,a)(f_r(s,a) - r - \gamma f_r(s',\pi)) \right| \\ & \leq \epsilon_{stat} \end{aligned}$$

The proof for the case $|\mathcal{E}_\mu(\pi, f_c) - \mathcal{E}_\mathcal{D}(\pi, f_c)| \leq \epsilon_{stat}$ is similar. \square

Lemma 5. (Empirical weighted average Bellman Error) With probability at least $1 - 2\delta$, for any $\pi \in \Pi$, we have

$$\mathcal{E}_{\mathcal{D}}(\pi, f_r^\pi) \leq C_{\ell_2}^* \sqrt{\epsilon_1} + \epsilon_{stat} \quad (20)$$

$$\mathcal{E}_{\mathcal{D}}(\pi, f_c^\pi) \leq C_{\ell_2}^* \sqrt{\epsilon_1} + \epsilon_{stat}, \quad (21)$$

where

$$f_r^\pi := \arg \min_{f_r \in \mathcal{F}} \sup_{\text{admissible } \nu} \|f_r - \mathcal{T}_r^\pi f_r\|_{2,\nu}^2, \forall \pi \in \Pi$$

$$f_c^\pi := \arg \min_{f_c \in \mathcal{G}} \sup_{\text{admissible } \nu} \|f_c - \mathcal{T}_c^\pi f_c\|_{2,\nu}^2, \forall \pi \in \Pi.$$

Proof. Condition on the high probability event in Lemma 4, we have

$$\begin{aligned} \mathcal{E}_\mu(\pi, f_r^\pi) &= \max_{w \in \mathcal{W}} |\mathbb{E}_\mu[w(s, a)((f - \mathcal{T}_r^\pi f_r^\pi)(s, a))]| \\ &\leq \mathcal{E}_\mu(\pi, f_r^\pi) = \max_{w \in \mathcal{W}} \|w\|_{2,\mu} \|f - \mathcal{T}_r^\pi f_r^\pi\|_{2,\mu} \\ &\leq C_{\ell_2}^* \sqrt{\epsilon_1}, \end{aligned}$$

where the first inequality is true because of Cauchy-Schwarz inequality and the second inequality comes from the definition of f_r^π and Assumption 3.2, thus we can obtain

$$\mathcal{E}_{\mathcal{D}}(\pi, f_r^\pi) \leq \mathcal{E}_\mu(\pi, f_r^\pi) + \epsilon_{stat} \leq C_{\ell_2}^* \sqrt{\epsilon_1} + \epsilon_{stat}. \quad (22)$$

Following a similar proof we can have

$$\hat{\mathcal{E}}_{\mathcal{D}}(\pi, f_c^\pi) \leq \mathcal{E}_\mu(\pi, f_c^\pi) + \epsilon_{stat} \leq C_{\ell_2}^* \sqrt{\epsilon_1} + \epsilon_{stat}. \quad (23)$$

□

Lemma 6. (Performance difference decomposition, restate of Lemma 12 in Cheng et al. (2022)) For an arbitrary policy $\pi, \hat{\pi} \in \Pi$, and f be an arbitrary function over $\mathcal{S} \times \mathcal{A}$. Then we have,

$$\begin{aligned} &J_\diamond(\pi) - J_\diamond(\hat{\pi}) \\ &= \mathbb{E}_\mu[(f - \mathcal{T}_\diamond^{\hat{\pi}})(s, a)] + \mathbb{E}_\pi[(\mathcal{T}_\diamond^{\hat{\pi}} f - f)(s, a)] + \mathbb{E}_\pi[f(s, \pi) - f(s, \hat{\pi})] + \mathcal{L}_\mu(\hat{\pi}, f) - \mathcal{L}_\mu(\hat{\pi}, Q_\diamond^{\hat{\pi}}), \end{aligned} \quad (24)$$

where $\diamond := r$ or c .

Proof. We prove the case when $\diamond := r$, the other case is identical. Let $R^{f, \hat{\pi}}(s, a) := f(s, a) - \gamma \mathbb{E}_{s'| (s, a)}[f(s', \hat{\pi})]$ be a virtual reward function for given f and $\hat{\pi}$. According to performance difference lemma (Kakade and Langford, 2002), We first have that

$$\begin{aligned} (J_r(\hat{\pi}) - J_r(\mu)) &= \mathcal{L}_\mu(\hat{\pi}, Q_r^{\hat{\pi}}) \\ &= \Delta(\hat{\pi}) + \mathcal{L}_\mu(\hat{\pi}, f) \quad (\Delta(\hat{\pi}) := \mathcal{L}_\mu(\hat{\pi}, Q_r^{\hat{\pi}}) - \mathcal{L}_\mu(\hat{\pi}, f)) \\ &= \Delta(\hat{\pi}) + \mathbb{E}_\mu[f(s, \hat{\pi}) - f(s, a)] \\ &= \Delta(\hat{\pi}) + (1 - \gamma)(J_{R^{f, \hat{\pi}}}(\hat{\pi}) - J_{R^{f, \hat{\pi}}}(\mu)) \\ &= \Delta(\hat{\pi}) + (1 - \gamma)Q_{R^{f, \hat{\pi}}}^{\hat{\pi}}(s_0, \hat{\pi}) - \mathbb{E}_\mu[R^{\hat{\pi}, f}(s, a)] \\ &= \Delta(\hat{\pi}) + (1 - \gamma)f(s_0, \hat{\pi}) - \mathbb{E}_\mu[R^{\hat{\pi}, f}(s, a)], \end{aligned}$$

where the last equality is true because that

$$Q_{R^{f, \hat{\pi}}}^{\hat{\pi}}(s, a) = (\mathcal{T}_{R^{f, \hat{\pi}}}^\pi f)(s, a) = R^{f, \hat{\pi}} + \gamma \mathbb{E}_{s'| (s, a)}[f(s', \hat{\pi})] = f(s, a).$$

Thus we have

$$\begin{aligned} (J_r(\pi) - J_r(\hat{\pi})) &= (J_r(\pi) - J_r(\mu)) - (J_r(\hat{\pi}) - J_r(\mu)) \\ &= (J_r(\pi) - f(d_0, \hat{\pi})) + \left(\mathbb{E}_\mu[R^{\hat{\pi}, f}(s, a)] - J_r(\mu) \right) - \Delta(\hat{\pi}). \end{aligned} \quad (25)$$

For the first term, we have

$$\begin{aligned}
(J_r(\pi) - f(d_0, \hat{\pi})) &= (J_r(\pi) - f(s_0, \hat{\pi})) && \text{(deterministic initial state)} \\
&= J_r(\pi) - \mathbb{E}_{d^\pi}[R^{\hat{\pi}, f}(s, a)] + \mathbb{E}_{d^\pi}[R^{\hat{\pi}, f}(s, a)] - f(s_0, \hat{\pi}) \\
&= \mathbb{E}_{d^\pi}[R(s, a) - R^{\hat{\pi}, f}(s, a)] + \mathbb{E}_{d^\pi}[f(s, \pi) - f(s, \hat{\pi})] \\
&= \mathbb{E}_{d^\pi}[(\mathcal{T}_r^{\hat{\pi}} f - f)(s, a)] + \mathbb{E}_{d^\pi}[f(s, \pi) - f(s, \hat{\pi})], \tag{26}
\end{aligned}$$

where the second equality is true because

$$\begin{aligned}
&\mathbb{E}_{d^\pi}[R^{\hat{\pi}, f}(s, a)] - f(s_0, \hat{\pi}) \\
&= \mathbb{E}_{d^\pi}[f(s, a) - \gamma \mathbb{E}_{s'|s, a}[f(s', \hat{\pi})]] - f(s_0, \hat{\pi}) \\
&= \mathbb{E}_{d^\pi}[f(s, \pi)] - \sum_s \sum_{t=1}^{\infty} \gamma^t \Pr(s_t = s | s_0 \sim d_0, \pi) f(s, \hat{\pi}(s)) - f(s_0, \hat{\pi}) \\
&= \mathbb{E}_{d^\pi}[f(s, \pi)] - \sum_s \sum_{t=0}^{\infty} \gamma^t \Pr(s_t = s | s_0 \sim d_0, \pi) f(s, \hat{\pi}(s)) \\
&= \mathbb{E}_{d^\pi}[f(s, \pi)] - \sum_{s, a} \sum_{t=0}^{\infty} \gamma^t \Pr(s_t = s, a_t = a | s_0 \sim d_0, \pi) f(s, \hat{\pi}(s)) \\
&= \mathbb{E}_{d^\pi}[f(s, \pi) - f(s, \hat{\pi})].
\end{aligned}$$

For the second term we have

$$\begin{aligned}
&\mathbb{E}_\mu[R^{\hat{\pi}, f}(s, a)] - J_r(\mu) \\
&= \mathbb{E}_\mu[R^{\hat{\pi}, f}(s, a) - R(s, a)] \\
&= \mathbb{E}_\mu[(f - \mathcal{T}_r^{\hat{\pi}} f)(s, a)]. \tag{27}
\end{aligned}$$

Therefore plugging 26 and (27) into Eq. (25), we have

$$\begin{aligned}
&J_r(\pi) - J_r(\hat{\pi}) \\
&= \mathbb{E}_\mu[(f - \mathcal{T}_r^{\hat{\pi}})(s, a)] + \mathbb{E}_\pi[(\mathcal{T}_r^{\hat{\pi}} f - f)(s, a)] + \mathbb{E}_\pi[f(s, \pi) - f(s, \hat{\pi})] + \mathcal{L}_\mu(\hat{\pi}, f) - \mathcal{L}_\mu(\hat{\pi}, Q_r^{\hat{\pi}}).
\end{aligned}$$

The proof is completed. \square

Lemma 7. *With probability at least $1 - 2\delta$, for any $f_r \in \mathcal{F}$, $f_c \in \mathcal{G}$, and $\pi \in \Pi$, we have:*

$$|\mathcal{L}_\mu(\pi, f_r) - \mathcal{L}_\mathcal{D}(\pi, f_r)| \leq \epsilon_{stat} \tag{28}$$

$$|\mathcal{L}_\mu(\pi, f_c) - \mathcal{L}_\mathcal{D}(\pi, f_c)| \leq \epsilon_{stat} \tag{29}$$

where $\epsilon_{stat} := V_{\max} C_{\ell_2}^* \sqrt{\frac{\log(|\mathcal{F}||\mathcal{G}||\Pi|W/\delta)}{N}} + \frac{V_{\max} B_w \log(|\mathcal{F}||\mathcal{G}||\Pi|W/\delta)}{N}$.

Proof. Recall that $\mathbb{E}_\mu[\mathcal{L}_\mathcal{D}(\pi, f_r)] = \mathcal{L}_\mu(\pi, f)$ and $|f_r(s, \pi) - f_r(s, a)| \leq V_{\max}$. For any $f_r \in \mathcal{F}$, policy $\pi \in \Pi$, applying a Hoeffding's inequality and a union bound we can obtain with probability $1 - \delta$,

$$|\mathcal{L}_\mu(\pi, f_r) - \mathcal{L}_\mathcal{D}(\pi, f_r)| \leq \mathcal{O}\left(V_{\max} \sqrt{\frac{\log(|\mathcal{F}||\Pi|/\delta)}{N}}\right) \leq \epsilon_{stat}. \tag{30}$$

The inequality for proving the f_c, π is the same. \square

B Missing Proofs

B.1 Proof of Theorem 4.1

Proof. According to the performance difference lemma (Kakade and Langford, 2002), we have

$$\begin{aligned}
&(J_r(\pi) - J_r(\mu)) - \lambda\{J_c(\pi) - J_c(\mu)\}_+ \\
&= \mathcal{L}_\mu(\pi, Q_r^\pi) - \lambda\{\mathcal{L}_\mu(\pi, Q_c^\pi)\}_+ \\
&= \mathcal{L}_\mu(\pi, Q_r^\pi) + \beta \mathcal{E}_\mu(\pi, Q_r^\pi) - \lambda\{\mathcal{L}_\mu(\pi, Q_c^\pi)\}_+ + \beta \hat{\mathcal{E}}_\mu(\pi, Q_c^\pi) \\
&\geq \mathcal{L}_\mu(\pi, f_r^\pi) + \beta \mathcal{E}_\mu(\pi, f_r^\pi) - \lambda\{\mathcal{L}_\mu(\pi, f_c^\pi)\}_+ + \beta \hat{\mathcal{E}}_\mu(\pi, f_c^\pi) \\
&\geq \mathcal{L}_\mu(\pi, f_r^\pi) - \lambda\{\mathcal{L}_\mu(\pi, f_c^\pi)\}_+, \tag{31}
\end{aligned}$$

where the second equality is true because $\mathcal{E}_\mu(\pi, Q_r^\pi) = \hat{\mathcal{E}}_\mu(\pi, Q_c^\pi) = 0$ by Assumption 3.2, and the first inequality comes from the selection of f_r^π and f_c^π in optimization (6).

Therefore, we can obtain

$$\begin{aligned} J_r(\hat{\pi}^*) - J_r(\mu) &\geq (\mathcal{L}_\mu(\hat{\pi}^*, f_r^{\hat{\pi}^*}) - \lambda\{\mathcal{L}_\mu(\hat{\pi}^*, f_c^{\hat{\pi}^*})\}_+) + \lambda\{J_c(\hat{\pi}^*) - J_c(\mu)\}_+ \\ &\geq (\mathcal{L}_\mu(\mu, f_r^\mu) - \lambda\{\mathcal{L}_\mu(\mu, f_c^\mu)\}_+) + \lambda\{J_c(\hat{\pi}^*) - J_c(\mu)\}_+ \\ &\geq \lambda\{J_c(\hat{\pi}^*) - J_c(\mu)\}_+ \geq 0 \end{aligned} \quad (32)$$

and

$$\{J_c(\hat{\pi}^*)\}_+ - \{J_c(\mu)\}_+ \leq \{J_c(\hat{\pi}^*) - J_c(\mu)\}_+ \leq \frac{1}{\lambda}(J_r(\hat{\pi}^*) - J_r(\mu)) \leq \frac{1}{\lambda}. \quad (33)$$

□

B.2 Proof of Lemma 2

Proof. Denote π_{ref} as π . First according to the definition for the no-regret oracle 5.1, we have

$$\begin{aligned} \frac{1}{K} \sum_{k=1}^K \mathbb{E}_\pi [f_r^k(s, \pi) - f_r^k(s, \pi_k) - \lambda\{f_c^k(s, \pi) - f_c^k(s, \pi)\}_+ \\ + \lambda\{f_c^k(s, \pi_k) - f_c^k(s, \pi)\}_+] \leq \epsilon_{opt}^\pi \end{aligned} \quad (34)$$

Therefore,

$$\begin{aligned} \frac{1}{K} \sum_{k=1}^K \mathbb{E}_\pi [f_r^k(s, \pi) - f_r^k(s, \pi_k)] \\ \leq \epsilon_{opt}^\pi + \frac{1}{K} \sum_{k=1}^K \mathbb{E}_\pi [\lambda\{f_c^k(s, \pi) - f_c^k(s, \pi)\}_+ - \lambda\{f_c^k(s, \pi_k) - f_c^k(s, \pi)\}_+] \leq \epsilon_{opt}^\pi, \end{aligned} \quad (35)$$

and

$$\frac{1}{K} \sum_{k=1}^K \mathbb{E}_\pi [\{f_c^k(s, \pi_k) - f_c^k(s, \pi)\}_+] \leq \epsilon_{opt}^\pi - \frac{1}{\lambda K} \sum_{k=1}^K \mathbb{E}_\pi [f_r^k(s, \pi) - f_r^k(s, \pi_k)] \leq \epsilon_{opt}^\pi + \frac{V_{\max}}{\lambda}. \quad (36)$$

We finish the proof. □

B.3 Proof of Theorem 5.2

Theorem (Restate of Theorem 5.2). *Under Assumptions 3.2 and 3.6, let the reference policy $\pi_{\text{ref}} \in \Pi$ be any policy satisfying Assumption 3.7, then with probability at least $1 - \delta$,*

$$J_r(\pi_{\text{ref}}) - J_r(\bar{\pi}) \leq \mathcal{O}\left(\epsilon_{\text{stat}} + C_{\ell_2}^* \sqrt{\epsilon_1}\right) + \epsilon_{opt}^\pi \quad (37)$$

$$J_c(\bar{\pi}) - J_c(\pi_{\text{ref}}) \leq \mathcal{O}\left(\epsilon_{\text{stat}} + C_{\ell_2}^* \sqrt{\epsilon_1}\right) + \epsilon_{opt}^\pi + \frac{V_{\max}}{\lambda}, \quad (38)$$

where $\epsilon_{\text{stat}} := V_{\max} C_{\ell_2}^* \sqrt{\frac{\log(|\mathcal{F}||\mathcal{G}||\Pi||W|/\delta)}{N}} + \frac{V_{\max} B_w \log(|\mathcal{F}||\mathcal{G}||\Pi||W|/\delta)}{N}$, and $\bar{\pi}$ is the policy returned by Algorithm 1 with $\beta = 2$ and π_{ref} as input.

Proof. Denote π_{ref} as π . According to the definition of $\bar{\pi}$, and Lemma 6 we have

$$\begin{aligned}
J_r(\pi) - J_r(\bar{\pi}) &= \frac{1}{K} \sum_{k=1}^K (J_r(\pi) - J_r(\pi_k)) \\
&= \frac{1}{K} \sum_{k=1}^K \left(\underbrace{\mathbb{E}_\mu [f_r^k - \mathcal{T}_r^{\pi_k} f_r^k]}_{\text{(I)}} + \underbrace{\mathbb{E}_\pi [\mathcal{T}_r^{\pi_k} f_r^k - f_r^k]}_{\text{(II)}} \right. \\
&\quad \left. + \underbrace{\mathbb{E}_\pi [f_r^k(s, \pi) - f_r^k(s, \pi_k)]}_{\text{(III)}} + \underbrace{\mathcal{L}_\mu(\pi_k, f_r^k) - \mathcal{L}_\mu(\pi_k, Q^{\pi_k})}_{\text{(IV)}} \right) \tag{39}
\end{aligned}$$

Condition on the high probability event in , we have

$$\text{(I)} + \text{(II)} \leq 2\mathcal{E}_\mu(\pi_k, f_r^k) \leq 2\mathcal{E}_\mathcal{D}(\pi_k, f_r^k) + 2\epsilon_{\text{stat}} \tag{40}$$

According to a similar argument as that in the Lemma 13 in [Cheng et al. \(2022\)](#), we have that

$$\begin{aligned}
&|\mathcal{L}_\mu(\pi_k, Q_r^{\pi_k}) - \mathcal{L}_\mu(\pi_k, f_r^{\pi_k})| \\
&= |\mathbb{E}_\mu[Q_r^{\pi_k}(s, \pi_k) - Q_r^{\pi_k}(s, a)] - \mathcal{L}_\mu(\pi_k, f_r^{\pi_k})| \\
&= |(J_r(\pi_k) - J_r(\mu)) - \mathcal{L}_\mu(\pi_k, f_r^{\pi_k})| \\
&= |(f_r^{\pi_k}(s_0, \pi_k) - J_r(\mu)) + (J_r(\pi_k) - f_r^{\pi_k}(s_0, \pi_k)) - \mathcal{L}_\mu(\pi_k, f_r^{\pi_k})| \\
&= |\mathbb{E}_\mu[f_r^{\pi_k}(s, \pi_k) - (\mathcal{T}_r^{\pi_k} f_r^{\pi_k})(s, a)] + \mathbb{E}_{d^{\pi_k}}[(\mathcal{T}_r^{\pi_k} f_r^{\pi_k})(s, a) - f_r^{\pi_k}(s, a)] - \mathcal{L}_\mu(\pi_k, f_r^{\pi_k})| \\
&\quad \text{(by the extension of performance difference lemma (Lemma 1 in [Cheng et al. \(2020\)](#)))} \\
&= |\mathcal{L}_\mu(\pi_k, f_r^{\pi_k}) + \mathbb{E}_\mu[f_r^{\pi_k}(s, a) - (\mathcal{T}_r^{\pi_k} f_r^{\pi_k})(s, a)] + \mathbb{E}_{d^{\pi_k}}[(\mathcal{T}_r^{\pi_k} f_r^{\pi_k})(s, a) - f_r^{\pi_k}(s, a)] - \mathcal{L}_\mu(\pi_k, f_r^{\pi_k})| \\
&\leq \|f_r^{\pi_k}(s, a) - (\mathcal{T}_r^{\pi_k} f_r^{\pi_k})(s, a)\|_{2, \mu} + \|(\mathcal{T}_r^{\pi_k} f_r^{\pi_k})(s, a) - f_r^{\pi_k}(s, a)\|_{2, d^{\pi_k}} \\
&\leq \mathcal{O}(\sqrt{\epsilon_1}), \tag{41}
\end{aligned}$$

where $f_r^\pi := \arg \min_{f_r \in \mathcal{F}} \sup_{\text{admissible } \nu} \|f_r - \mathcal{T}_r^\pi f_r\|_{2, \nu}^2, \forall \pi \in \Pi$. By using Lemma 7, we have

$$|\mathcal{L}_\mu(\pi_k, f_r^k) - \mathcal{L}_\mathcal{D}(\pi_k, f_r^k)| + |\mathcal{L}_\mu(\pi_k, f_r^{\pi_k}) - \mathcal{L}_\mathcal{D}(\pi_k, f_r^{\pi_k})| \leq \mathcal{O}(\epsilon_{\text{stat}}). \tag{42}$$

Therefore

$$\text{(I)} + \text{(II)} + \text{(IV)} \leq \mathcal{L}_\mu(\pi_k, f_r^k) + 2\mathcal{E}_\mu(\pi_k, f_r^k) + 2\epsilon_{\text{stat}} - \mathcal{L}_\mu(\pi_k, f_r^{\pi_k}) + \mathcal{O}(\sqrt{\epsilon_1}) \tag{43}$$

$$\leq \mathcal{L}_\mathcal{D}(\pi_k, f_r^k) + 2\mathcal{E}_\mathcal{D}(\pi_k, f_r^k) + \mathcal{O}(\epsilon_{\text{stat}}) - \mathcal{L}_\mathcal{D}(\pi_k, f_r^{\pi_k}) + \mathcal{O}(\sqrt{\epsilon_1}) \tag{44}$$

$$\leq \mathcal{L}_\mathcal{D}(\pi_k, f_r^{\pi_k}) + 2\mathcal{E}_\mathcal{D}(\pi_k, f_r^{\pi_k}) + \mathcal{O}(\epsilon_{\text{stat}}) - \mathcal{L}_\mathcal{D}(\pi_k, f_r^{\pi_k}) + \mathcal{O}(\sqrt{\epsilon_1}) \tag{45}$$

$$\leq \mathcal{O}(\epsilon_{\text{stat}} + C_{\ell_2}^* \sqrt{\epsilon_1}), \tag{46}$$

where the third inequality holds by the selection of f_r^k , and the last inequality holds by Lemma 5. Therefore by using Lemma B.2 we obtain

$$J_r(\pi) - J_r(\bar{\pi}) \leq \mathcal{O}(\epsilon_{\text{stat}} + C_{\ell_2}^* \sqrt{\epsilon_1}) + \epsilon_{\text{opt}}. \tag{47}$$

Following a similar argument, we have that

$$J_c(\bar{\pi}) - J_c(\pi) = \frac{1}{K} \sum_{k=1}^K (J_c(\pi_k) - J_c(\pi)) \leq \mathcal{O}(\epsilon_{\text{stat}} + C_{\ell_2}^* \sqrt{\epsilon_1}) + \epsilon_{\text{opt}}^\pi + \frac{V_{\text{max}}}{\lambda}. \tag{48}$$

□

B.4 Proof of Theorem 5.6

Theorem (Restate of Theorem 5.2). *Under Assumptions 3.2 and 3.6, let the reference policy $\pi_{\text{ref}} \in \Pi$ be any policy satisfying Assumption 3.7, then with probability at least $1 - \delta$,*

$$J_r(\mu) - J_r(\bar{\pi}) \leq \mathcal{O}\left(\epsilon_{\text{stat}}^\pi + C_{\ell_2}^* \sqrt{\epsilon_1}\right) + \epsilon_{\text{opt}}^\mu \tag{49}$$

$$J_c(\bar{\pi}) - J_c(\mu) \leq \mathcal{O}\left(\epsilon_{\text{stat}}^\pi + C_{\ell_2}^* \sqrt{\epsilon_1}\right) + \epsilon_{\text{opt}}^\mu + \frac{V_{\text{max}}}{\lambda}, \tag{50}$$

where $\epsilon_{stat} := V_{\max} C_{\ell_2}^* \sqrt{\frac{\log(|\mathcal{F}||\Pi||W|/\delta)}{N}} + \frac{V_{\max} B_w \log(|\mathcal{F}||\Pi||W|/\delta)}{N}$, and $\bar{\pi}$ is the policy returned by Algorithm 1 with $\beta \geq 0$ and μ as input.

Proof. Following a similar proof in Theorem 5.2. But when the reference policy is the behavior policy, we have (I) + (II) = 0. Therefore we have have

$$\begin{aligned}
(\text{IV}) &= \mathcal{L}_\mu(\pi_k, f_r^k) - \mathcal{L}_\mu(\pi_k, Q^{\pi_k}) \\
&\leq \mathcal{L}_\mu(\pi_k, f_r^k) - \mathcal{L}_\mu(\pi_k, Q^{\pi_k}) + \beta \mathcal{E}_{\mathcal{D}}(\pi_k, f_r^k) \\
&\leq \mathcal{L}_\mu(\pi_k, f_r^k) - \mathcal{L}_\mu(\pi_k, Q^{\pi_k}) + \beta \mathcal{E}_{\mathcal{D}}(\pi_k, f_r^k) - \beta \mathcal{E}_{\mathcal{D}}(\pi, f_{\pi_k}) + \beta C_{\ell_2}^* \sqrt{\epsilon_1} + \beta \epsilon_{stat} \quad (\text{Lemma 5}) \\
&\leq \mathcal{L}_{\mathcal{D}}(\pi_k, f_r^k) + \beta \mathcal{E}_{\mathcal{D}}(\pi_k, f_r^k) - \mathcal{L}_{\mathcal{D}}(\pi_k, f_r^{\pi_k}) - \beta \mathcal{E}_{\mathcal{D}}(\pi, f_{\pi_k}) + (\beta + 1)(\epsilon_{stat} + C_{\ell_2}^* \sqrt{\epsilon_1}) \\
&\leq (\beta + 1)(\epsilon_{stat} + C_{\ell_2}^* \sqrt{\epsilon_1}).
\end{aligned}$$

We finish the proof. \square

C Discussion on obtaining the behavior policy

To extract the behavior policy when it is not provided, we can simply run behavior cloning on the offline data. In particular, we can estimate the learned behavior policy $\hat{\pi}_\mu$ as follows: $\forall s \in \mathcal{D}$, $\hat{\pi}_\mu(a|s) \leftarrow \frac{n(s,a)}{n(s)}$, and $\forall s \notin \mathcal{D}$, $\hat{\pi}_\mu(a|s) \leftarrow \frac{1}{|A|}$, where $n(s, a)$ is the number of times (s, a) appears in the offline dataset \mathcal{D} . Essentially, the estimated BC policy matches the empirical behavior policy on states in the offline dataset and takes uniform random actions outside the support of the dataset. It is easy to show that the gap between the learned policy $\hat{\pi}_\mu$ and the behavior policy π_μ is upper bounded by $\mathcal{O}(\min\{1, |\mathcal{S}|/N\})$ (Kumar et al., 2022; Rajaraman et al., 2020). We can have a very accurate estimate as long as the size of the dataset is large enough.

D Experimental Supplement

D.1 Practical Algorithm

The practical version of our algorithm WSAC is shown in Algorithm 2.

Algorithm 2 WSAC - Practical Version

- 1: **Input:** Batch data \mathcal{D} , policy network π , network for the reward critic f_r , network for the cost critic f_c , $\beta > 0$, $\lambda > 0$.
- 2: **for** $k = 1, 2, \dots, K$ **do**
- 3: Sample minibatch $\mathcal{D}_{\text{mini}}$ from the dataset \mathcal{D} .
- 4: **Update Critic Networks:**

$$\begin{aligned}
l_{\text{reward}}(f_r) &:= \mathcal{L}_{\mathcal{D}_{\text{mini}}}(\pi, f_r) + \beta \mathcal{E}_{\mathcal{D}_{\text{mini}}}(\pi, f_r), \\
f_r &\leftarrow \text{ADAM}(f_r - \eta_{\text{fast}} \nabla l_{\text{reward}}(f_r)), \\
l_{\text{cost}}(f_c) &:= -\lambda \mathcal{L}_{\mathcal{D}_{\text{mini}}}(\pi, f_c) + \beta \mathcal{E}_{\mathcal{D}_{\text{mini}}}(\pi, f_c), \\
f_c &\leftarrow \text{ADAM}(f_c - \eta_{\text{fast}} \nabla l_{\text{cost}}(f_c)).
\end{aligned}$$

- 5: **Update Policy Network:**

$$\begin{aligned}
l_{\text{actor}}(\pi) &:= -\mathcal{L}_{\text{mini}}(\pi, f_r) + \lambda \{\mathcal{L}_{\text{mini}}(\pi, f_c)\}_+, \\
\pi &\leftarrow \text{ADAM}(\pi - \eta_{\text{slow}} \nabla l_{\text{actor}}(\pi)).
\end{aligned}$$

- 6: **end for**
 - 7: **Output:** π
-

D.2 Environments Description

Besides the ‘‘BallCircle’’ environment, we also study several representative environments as follows. All of them are shown in Figure 2 and their offline dataset is from Liu et al. (2023a).



Figure 2: BallCircle and CarCircle (left), PointButton (medium), PointPush(right) .

- **CarCircle:** This environment requires the car to move on a circle in a clockwise direction within the safety zone defined by the boundaries. The car is a four-wheeled agent based on MIT’s race car. The reward is dense and increases by the car’s velocity and by the proximity towards the boundary of the circle and the cost is incurred if the agent leaves the safety zone defined by the two yellow boundaries, which are the same as "CarCircle".
- **PointButton:** This environment requires the point to navigate to the goal button location and touch the right goal button while avoiding more gremlins and hazards. The point has two actuators, one for rotation and the other for forward/backward movement. The reward consists of two parts, indicating the distance between the agent and the goal and if the agent reaches the goal button and touches it. The cost will be incurred if the agent enters the hazardous areas, contacts the gremlins, or presses the wrong button.
- **PointPush:** This environment requires the point to push a box to reach the goal while circumventing hazards and pillars. The objects are in 2D planes and the point is the same as "PointButton". It has a small square in front of it, which makes it easier to determine the orientation visually and also helps point push the box.

D.3 Implementation Details and Experimental settings

We run all the experiments with NVIDIA GeForce RTX 3080 Ti 8–Core Processor.

The normalized reward and cost are summarized as follows:

$$R_{normalized} = \frac{R_{\pi} - r_{min}(\mathcal{M})}{r_{max}(\mathcal{M}) - r_{min}(\mathcal{M})} \quad (51)$$


$$C_{normalized} = \frac{C_{\pi} + \epsilon}{\kappa + \epsilon}, \quad (52)$$

where $r(\mathcal{M})$ is the empirical reward for task \mathcal{M} , κ is the cost threshold, ϵ is a small number to ensure numerical stability. Thus any normalized cost below 1 is considered as safe. We use R_{π} and C_{π} to denote the cumulative rewards and cost for the evaluated policy, respectively. The parameters of $r_{min}(\mathcal{M})$, $r_{max}(\mathcal{M})$ and κ are environment-dependent constants and the specific values can be found in the Appendix D. We remark that the normalized reward and cost only used for demonstrating the performance purpose and are not used in the training process. The detailed value of the reward and costs can be found in Table 3. To mitigate the risk of unsafe scenarios, we introduce a hyperparameter


Table 3: Hyperparameters of WSAC

Parameters	BallCircle	CarCircle	PointButton	PointPush
β_c	30.0	38.0	30.0	30.0
β_r	10.0	12.0	10.0	10.0
UB_{Q_C}	30.0	28.0	32.0	30.0
λ	[1.0, 20.0]			
Batch size	512			
Actor learning rate	0.0001			
Critic learning rate	0.0003			
κ	40			
$r_{min}(\mathcal{M})$	0.3831	3.4844	0.0141	0.0012
$r_{max}(\mathcal{M})$	881.4633	534.3061	42.8986	14.6910


UB_{Q_C} to the cost Q -function as an overestimation when calculating the actor loss. We use two separate β_r, β_c for reward and cost Q functions to make the algorithm more flexible.




(a) BallCircle



(b) CarCircle



(c) PointButton



(d) PointPush

Figure 3: The moving average of evaluation results is recorded every 500 training steps, with each result representing the average over 20 evaluation episodes and three random seeds. A cost threshold 1 is applied, with any normalized cost below 1 considered safe.

We use different β for the reward and cost critic networks and different UB_{QC} for the actor-network to make the adversarial training more stable. We also let the key parameter λ within a certain range balance reward and cost during the training process. Their values are shown in Table 3. In experiments, we take $\mathcal{W} = \{0, C_\infty\}$ for computation effective. Then we can reduce $\mathcal{E}_{\mathcal{D}}(\pi, f)$ to $C_\infty \mathbb{E}_{\mathcal{D}}[(f(s, a) - r - \gamma f(s', \pi))^2]$ and reduce $\hat{\mathcal{E}}_{\mathcal{D}}(\pi, f)$ to $C_\infty \mathbb{E}_{\mathcal{D}}[(f(s, a) - c - \gamma f(s', \pi))^2]$. In this case, C_∞ can be considered as a part of the hyperparameter $\beta_r(\beta_c)$.

D.4 Experimental results details and supplements

The evaluation performances of the agents in each environment after 30000 update steps of training are shown in Table 2, and the performance of average rewards and costs are shown in Figure 3. From the results, we observe that WSAC achieves a best reward performance with significantly lowest costs against all the baselines. It suggests WSAC can establish a safe and efficient policy and achieve a steady improvement by leveraging the offline dataset.

D.5 Simulations under different cost limits

To further evaluate the performance of our algorithm under varying situations. We further compare our algorithm with baselines under varying cost limits, we report the average performance of our method and other baselines in Table 4. Specifically, cost limits of [10, 20, 40] are used for the BallCircle and CarCircle environments, and [20, 40, 80] for the PointButton and PointPush environments, following the standard setup outlined by Liu et al. (2023a). Our results demonstrate that WSAC maintains safety across all environments, and its performance is either comparable to or superior to the best baseline in each case. These suggest that WSAC is well-suited for adapting to tasks of varying difficulty.

Table 4: The normalized reward and cost of WSAC and other baselines for different cost limits. Each value is averaged over 3 distinct cost limits, 20 evaluation episodes, and 3 random seeds. The Average line shows the average situation in various environments. The cost threshold is 1. Gray: Unsafe agent whose normalized cost is greater than 1. Blue: Safe agent with best performance. The performance of all the baselines is reported according to the results in Liu et al. (2023a).

	BC		Safe-BC		CDT		BCQL		BEARL		CPQ		COptiDICE		WSAC	
	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
BallCircle	0.74	4.71	0.52	0.65	0.77	1.07	0.69	2.36	0.86	3.09	0.64	0.76	0.70	2.61	0.74	0.51
CarCircle	0.58	3.74	0.50	0.84	0.75	0.95	0.63	1.89	0.74	2.18	0.71	0.33	0.49	3.14	0.65	0.55
PointButton	0.27	2.02	0.16	1.10	0.46	1.57	0.40	2.66	0.43	2.47	0.58	4.30	0.15	1.51	0.11	0.55
PointPush	0.18	0.91	0.11	0.80	0.21	0.65	0.23	0.99	0.16	0.89	0.11	1.04	0.02	1.18	0.07	0.61
Average	0.44	2.85	0.32	0.85	0.55	1.06	0.49	1.98	0.55	2.16	0.51	1.61	0.34	2.11	0.39	0.56

D.6 Ablation studies

To investigate the contribution of each component of our algorithm, including the weighted Bellman regularizer, the aggression-limited objective, and the no-regret policy optimization (which together guarantee our theoretical results), we performed an ablation study in the tabular setting. The results, presented in Table 5, indicate that the weighted Bellman regularization ensures the safety of the algorithm, while the aggression-limited objective fine-tunes the algorithm to achieve higher rewards without compromising safety.

Table 5: Ablation study under tabular case (cost limit is 0.1) over 10 repeat experiments

Components	cost	reward
ALL	0.014 \pm 0.006	0.788 \pm 0.004
W/O no-regret policy optimization	0.014 \pm 0.006	0.788 \pm 0.004
W/O Aggression-limited objective	0.014 \pm 0.006	0.788 \pm 0.005
W/O Weighted Bellman regularizer	0.323 \pm 0.061	0.684 \pm 0.017

D.7 Sensitivity Analysis of Hyper-Parameters

We provide the rewards and costs under different sets of $\beta_r = \beta_c \in \{1, 0.5, 0.05\}$ and $\lambda \in \{[0, 1], [0, 2], [1, 2]\}$ (since our λ only increases, the closed interval here represents the initial value




Figure 4: Sensitivity Analysis of Hyperparameters in the Tabular Case. The left figure illustrates tests conducted with various β values (For the sake of discussion, we denote $\beta = \beta_r = \beta_c$) with $\lambda = [0, 2]$, while the right figure presents tests across different ranges of λ with $\beta_r = \beta_c = 2.0$ and the upper bound of λ) to demonstrate the robustness of our approach in the tabular setting in Figure 4. We can observe that the performance is almost the same under different sets of parameters and different qualities of behavior policies.

NeurIPS Paper Checklist

The checklist is designed to encourage best practices for responsible machine learning research, addressing issues of reproducibility, transparency, research ethics, and societal impact. Do not remove the checklist: **The papers not including the checklist will be desk rejected.** The checklist should follow the references and follow the (optional) supplemental material. The checklist does NOT count towards the page limit.

Please read the checklist guidelines carefully for information on how to answer these questions. For each question in the checklist:

- You should answer [Yes], [No], or [NA].
- [NA] means either that the question is Not Applicable for that particular paper or the relevant information is Not Available.
- Please provide a short (1–2 sentence) justification right after your answer (even for NA).

The checklist answers are an integral part of your paper submission. They are visible to the reviewers, area chairs, senior area chairs, and ethics reviewers. You will be asked to also include it (after eventual revisions) with the final version of your paper, and its final version will be published with the paper.

The reviewers of your paper will be asked to use the checklist as one of the factors in their evaluation. While "[Yes]" is generally preferable to "[No]", it is perfectly acceptable to answer "[No]" provided a proper justification is given (e.g., "error bars are not reported because it would be too computationally expensive" or "we were unable to find the license for the dataset we used"). In general, answering "[No]" or "[NA]" is not grounds for rejection. While the questions are phrased in a binary way, we acknowledge that the true answer is often more nuanced, so please just use your best judgment and write a justification to elaborate. All supporting evidence can appear either in the main paper or the supplemental material, provided in appendix. If you answer [Yes] to a question, in the justification please point to the section(s) where related material for the question can be found.

IMPORTANT, please:

- **Delete this instruction block, but keep the section heading "NeurIPS paper checklist",**
- **Keep the checklist subsection headings, questions/answers and guidelines below.**
- **Do not modify the questions and only use the provided macros for your answers.**

1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [Yes]

Justification: The claims made in Abstract and Introduction reflect the paper's contributions.

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [Yes]

Justification: The limitations have been properly discussed in Conclusion and in Section 4.

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. Theory Assumptions and Proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [Yes]

Justification: All the assumptions have been properly stated. The complete proofs are in Appendix.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. Experimental Result Reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes]

Justification: Our experimental results are reproducible.

Guidelines:

- The answer NA means that the paper does not include experiments.

- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
 - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
 - (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
 - (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
 - (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [Yes]

Justification: We provide the data and code.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so “No” is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).

- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. Experimental Setting/Details

Question: Does the paper specify all the training and test details (e.g., data splits, hyper-parameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [Yes]

Justification: The Experimental setting and details are in Appendix.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

7. Experiment Statistical Significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [Yes]

Justification: In the experiment, we used multiple random seeds to ensure the statistical significance of the results.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. Experiments Compute Resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [Yes]

Justification: We specified the computational resources we used in the Appendix.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.

- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. Code Of Ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines?>

Answer: [Yes]

Justification: This paper is theoretical in nature, and it has been conducted with the NeurIPS code of Ethics.

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. Broader Impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [Yes]

Justification: This paper has tremendous positive societal impact as it develops RL algorithms with provable safety guarantee using offline data. Such a guarantee is essential for practical implementation of RL algorithms.

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification:

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [NA]

Justification:

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, paperswithcode.com/datasets has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. New Assets

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [NA]

Justification:

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. Crowdsourcing and Research with Human Subjects

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer:[NA]

Justification:

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. Institutional Review Board (IRB) Approvals or Equivalent for Research with Human Subjects

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification:

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.