

# Constraint-Adaptive Policy Switching for Offline Safe Reinforcement Learning

Yassine Chemingui<sup>1</sup>, Aryan Deshwal<sup>2</sup>, Honghao Wei<sup>1</sup>, Alan Fern<sup>3</sup>, Jana Doppa<sup>1</sup>

<sup>1</sup>Washington State University

<sup>2</sup>University of Minnesota

<sup>3</sup>Oregon State University

{yassine.chemingui,honghao.wei,jana.doppa}@wsu.edu, adeshwal@umn.edu, alan.fern@oregonstate.edu

## Abstract

Offline safe reinforcement learning (OSRL) involves learning a decision-making policy to maximize rewards from a fixed batch of training data to satisfy *pre-defined* safety constraints. However, adapting to varying safety constraints during deployment without retraining remains an under-explored challenge. To address this challenge, we introduce constraint-adaptive policy switching (CAPS), a wrapper framework around existing offline RL algorithms. During training, CAPS uses offline data to learn multiple policies with a shared representation that optimize different reward and cost trade-offs. During testing, CAPS switches between those policies by selecting at each state the policy that maximizes future rewards among those that satisfy the current cost constraint. Our experiments on 38 tasks from the DSRL benchmark demonstrate that CAPS consistently outperforms existing methods, establishing a strong wrapper-based baseline for OSRL.

## 1 Introduction

Online reinforcement learning (RL) has shown great successes in applications where the agent can continuously interact with the environment to collect new feedback data to improve its decision-making policy (Mnih et al. 2015; Silver et al. 2018; Li 2017). However, in many real-world domains such as agriculture, smart grid and healthcare, executing exploratory decisions is costly and/or dangerous, but we have access to pre-collected datasets. Offline RL (Levine et al. 2020) is an emerging paradigm to learn decision policies *solely* from such offline datasets, eliminating the need for additional interaction with the environment. The key challenge in offline RL is addressing the distributional shift between the state-action distribution of the offline dataset and the learned policy, as this shift results in extrapolation error. Existing offline RL methods typically employ the principle of pessimism in the face of uncertainty (i.e., some form of penalization to deal with out-of-distribution states/actions) to address this challenge (Levine et al. 2020).

In many safety-critical domains (e.g., agriculture and smart grid), the agent behavior needs to also satisfy some cost constraints in addition to maximizing reward. Most of the prior work on offline RL focuses on the unconstrained setting where the policy is optimized to maximize rewards. There is

relatively less work on offline constrained/safe RL and the majority of these methods assume that the cost constraint is known during the training phase (as elaborated in the related work section). However, in many applications, the cost constraints can change at deployment depending on the use-case scenario. For example, in agricultural drones, the reward is the area covered during spraying, and the cost is the battery/power consumption. The goal is to maximize efficiency for a given agricultural field by optimizing coverage based on the available battery power.

This paper studies the problem of offline safe RL (OSRL) where the safety/cost constraints can vary after deployment. We propose *Constraint-Adaptive Policy Switching (CAPS)*, an algorithm that can be wrapped around existing offline RL methods and is easy to implement. CAPS handles the challenge of unknown cost constraint threshold as follows. At training time, we learn multiple policies with a shared representation with the goal of achieving different reward and cost trade-offs: one centered on reward maximization, one on cost minimization, and other policies optimized for varied linear scalarizations of costs and rewards. At testing time, CAPS switches between the decisions from the trained policies at each state, selecting the one that maximizes future rewards from the ones that satisfy the cost constraint. We also theoretically analyze the conditions under which CAPS decision-making process has safety guarantees.

We perform experiments on 38 tasks from the DSRL benchmark (Liu et al. 2024) to compare CAPS with state-of-the-art OSRL baselines and our main findings are as follows. First, CAPS wrapped around two qualitatively different offline RL algorithms exhibits safe behavior (first-order objective) on a larger fraction of tasks compared to baselines across different cost threshold constraints. Prior methods struggle even with known cost constraints, potentially because of estimation errors in value functions and/or instability in the Lagrange multiplier. Second, the implicit Q-learning instantiation of CAPS demonstrates safe performance in **34/38 tasks (89%)** and achieves the highest rewards in **18** of those tasks compared to the baselines.. Third, CAPS using shared representation performs significantly better than CAPS with independently trained policies. Finally, increasing the number of policies improves the overall performance of CAPS, but CAPS with two policies (one for reward maximization and one for cost minimization) is competitive, has the advantage

of removing extra hyperparameters, and stands as a strong minimalist approach for effective offline safe RL.

**Contributions.** The key contribution is the development and evaluation of CAPS framework for offline safe RL with varying cost constraints. Specific contributions include:

- Development of CAPS to handle any test-time cost threshold constraint. CAPS can be wrapped around existing offline RL algorithms and easy to implement.
- Theoretical analysis to understand the conditions under which CAPS decision policy has safety guarantees.
- Experiments and ablations on 38 DSRL tasks, demonstrate the effectiveness of CAPS over prior methods. The code/appendices are available at <https://github.com/yassineCh/CAPS>.

## 2 Problem Setup

RL problems with safety constraints are naturally formulated within the Constrained Markov Decision Process (CMDP) framework. A CMDP is defined by a tuple  $M = (S, A, P, r, c, \mu_0)$ , where  $S$  is the state space,  $A$  is the action space,  $P : S \times A \times S \rightarrow [0, 1]$  is the state transition probability function,  $r : S \times A \rightarrow \mathbb{R}$  defines the reward function, and  $c : S \rightarrow [0, C_{\max}]$  quantifies the costs associated with states, where  $C_{\max}$  is the maximum possible cost, and  $\mu_0 : S \rightarrow [0, 1]$  is the initial state distribution. For simplicity of notation, we define cost function on states noting that it is easy to define an equivalent CMDP with costs associated with state-action pairs.

In offline safe RL problems, we are given a fixed pre-collected dataset  $\mathcal{D} = \{(s, a, r, c, s')\}_{i=1}^n$  from one or more (unknown) behavior policies, where each training example  $i$  contains the action  $a$  taken at state  $s$ , reward received  $r$ , cost incurred  $c$ , and the next state  $s'$ . The goal is to learn a policy  $\pi : S \rightarrow A$  from the offline dataset  $\mathcal{D}$  to maximize the expected reward while satisfying a specified cost/safety constraint. This problem is mathematically formulated as:

$$\max_{\pi} \mathbb{E}_{\tau \sim \pi}[R(\tau)] \quad \text{subject to} \quad \mathbb{E}_{\tau \sim \pi}[C(\tau)] \leq \kappa,$$

Here,  $\kappa \in [0, +\infty)$  is the cost threshold for safety constraint,  $\tau = \{s_1, a_1, r_1, c_1, \dots, s_T, a_T, r_T, c_T\}$  denotes a trajectory sampled by executing the policy  $\pi$ ,  $T$  is the length of episode,  $R(\tau) = \sum_{t=1}^T r_t$  is the total accumulated reward, and  $C(\tau) = \sum_{t=1}^T c_t$  is the total incurred cost.

The majority of prior work on OSRL assumes that the cost threshold  $\kappa$  for safety constraint is known at the training time and remains same at deployment. However, this assumption is violated in many real-world applications where the safety requirements can change during deployment based on the operating conditions or environmental factors. For example, in agricultural crop management, the goal may be to maximize yield under constraints on water and/or nitrogen application depending on time-varying regulations. Additionally, in a decision-support scenario, it can be useful for decision makers to quickly explore the impact of different cost thresholds without requiring retraining.

Other application domains where this problem setting arises include healthcare, power systems, and robotics. Consequently, OSRL methods that learn a single policy tailored to a specific cost constraint are ill-equipped to handle varying test-time cost constraints.

The goal of this paper is to fill this knowledge gap in the current literature by studying a variant of the OSRL problem where the cost constraint  $\kappa$  can vary during deployment.

## 3 Related Work

**Online Safe RL.** RL with safety constraints in the online setting, where the agent can interact with the environment to collect feedback data, is extensively studied (Garcia and Fernández 2015; Gu et al. 2024). However, most of these methods assume a fixed constraint threshold during both training and deployment. Constraint-conditioned policy optimization (Yao et al. 2023), addresses this limitation and adapts to changing constraints using two strategies: approximating value functions under unseen constraints and encoding arbitrary constraint thresholds during policy optimization.

**Offline RL.** The goal in offline RL is to learn effective policies from fixed static datasets without any additional interaction with the environment. Prior work has tackled the key challenge of distribution shift in offline RL (Levine et al. 2020; Figueiredo Prudencio, Maximo, and Colombini 2024). Some representative techniques include incorporating regularization in value function estimation (Fujimoto and Gu 2021; Kostrikov, Nair, and Levine 2022; Kumar et al. 2019; Lyu et al. 2022; Yang et al. 2022), sequential modeling formulations (Janner, Li, and Levine 2021; Wang et al. 2022), uncertainty-aware methods (An et al. 2021; Bai et al. 2022), divergence based policy constraints (Wu, Tucker, and Nachum 2020; Jaques et al. 2020; Wu et al. 2022), and extensions of model based RL (Kidambi et al. 2020; Yu et al. 2020; Rigter, Lacerda, and Hawes 2022) and imitation learning (Xu et al. 2022). SfBC (Chen et al. 2023) and IDQL (Hansen-Estruch et al. 2023) filter candidate actions using Q-function evaluations and reweight them for selection. In contrast, CAPS avoids reweighting by directly selecting the best safe action using two Q-functions. PEX (Zhang, Xu, and Yu 2023) is a offline-to-online RL approach that adds a new expansion policy to the offline policy, unlike CAPS, which dynamically switches between pre-trained policies for safety.

**Offline Safe RL.** Recent work has studied the offline safe RL problem, aiming to learn a safe policy from a given offline dataset (Liu et al. 2024). These methods include finding cost-conservative policies for better actual constraint satisfaction via some form of constrained policy optimization formulation (Polosky et al. 2022; Lee et al. 2022), Lagrangian based approaches to handle cost constraints (Xu, Zhan, and Zhu 2022), and the constrained decision transformer (CDT) (Liu et al. 2023), which exploits the advances in sequential modeling by learning from a dataset of trajectories. TREBI (Lin et al. 2023) and FISOR (Zheng et al. 2024) leverage diffusion models for creating safe policies: TREBI generates safe trajectories, while FISOR uses a diffusion actor to select actions within feasible regions. However, these methods, with

the exception of CDT, rely on fixed cost thresholds during training, limiting their flexibility for deployment scenarios with varying cost threshold constraints.

## 4 Constraint-Adaptive Policy Switching

In this section, we first describe the reasoning procedure behind the *Constraint-Adaptive Policy Switching (CAPS)* approach of handling unknown constraints at testing time. Next, we describe the training process and the design choices that enable effective decision-making in CAPS.

### 4.1 Decision-Making in CAPS

One way to handle dynamically changing cost constraints would be to: 1) train a large set of  $K$  policies that span a wide spectrum of cost-reward trade-offs, and 2) at each test-time state, select the best of those policies that satisfies the current cost constraint. While this makes sense in concept, there are at least two key practical issues. First, there is a high computational cost involved in optimizing  $K$  policies via safe offline RL. Second, and more importantly, step (2) requires accurately assessing the safety of a given policy and also selecting the best among the safe ones. This requires accurate cost and value functions for each policy, which would need to be estimated using off-policy evaluation (OPE) techniques with the available dataset. Unfortunately, as our experimental results show, current OPE techniques are not reliable enough to ensure test-time safety (Figueiredo Prudencio, Maximo, and Colombini 2024). The key idea of CAPS is to address these two issues by creating a set of  $K$  diverse policies at a much smaller computational cost, while avoiding a reliance on OPE to provide test-time safety.

CAPS relies on two key components for its decision-making to handle unknown cost constraints. First, a set of  $K \geq 2$  policies  $\mathcal{P} = \{\pi^r, \pi^1, \dots, \pi^{K-2}, \pi^c\}$  that are trained using the offline dataset  $\mathcal{D}$  to cover actions with different reward and cost trade-offs. At one extreme,  $\pi^r$  is trained to maximize reward via standard Offline-RL, while ignoring the cost objective. At the other extreme,  $\pi^c$  is trained to minimize cost via standard Offline-RL, while ignoring reward. In this sense,  $\pi^c$  is maximally safe, while  $\pi^r$  is cost agnostic. The remaining  $K - 2$  policies are trained with the intent of covering actions that achieve a variety of reward and cost trade-offs, while avoiding the full computational cost of Offline-RL for each one. Second, CAPS retains from Offline-RL the learned Q-functions  $Q_t^r$  and  $Q_t^c$  which we assume are non-stationary given the finite time horizon:

$$Q_t^r(s, a) = \mathbb{E}\left[\sum_t^T r_t | s_t = s, a_t = a\right] \quad (1)$$

$$Q_t^c(s, a) = \mathbb{E}\left[\sum_t^T c_t | s_t = s, a_t = a\right], \quad (2)$$

where the expectation is taken over the corresponding policy with respect to randomness from reward/cost functions and state transitions. Intuitively, CAPS will use  $Q_t^c$  to help filter out unsafe decisions from the  $K$  policies and  $Q_t^r$  is employed to select the best decision from the unfiltered ones. Note

that these Q-functions are an artifact of policy optimization, rather than OPE of an arbitrary policy. This is a subtle, but important, aspect of CAPS, since it has been observed that value estimates of arbitrary policies via OPE tends to be less reliable (Figueiredo Prudencio, Maximo, and Colombini 2024).

Given  $K$  policies  $\mathcal{P} = \{\pi^r, \pi^1, \dots, \pi^{K-2}, \pi^c\}$ , two Q-functions  $Q^r$  and  $Q^c$ , cost threshold constraint  $\kappa$ , and accumulated cost  $c_{\leq t}$  up to current state  $s_t$ , CAPS selects an action as follows:

1. **Filter:** Identify the subset of feasible actions  $\mathcal{F}_t \subseteq \{\pi^r(s_t), \pi^1(s_t), \dots, \pi^{K-2}(s_t), \pi^c(s_t)\}$  selected by the  $K$  policies and are safe in  $s_t$  according to  $Q^c$  and cost threshold  $\kappa$ :

$$\mathcal{F}_t = \{\pi(s_t) \mid Q_t^c(s_t, \pi(s_t)) + c_{\leq t} \leq \kappa, \pi \in \mathcal{P}\} \quad (3)$$

2. **Select:** Among the feasible actions  $\mathcal{F}_t$ , pick the one that maximizes the estimated reward according to  $Q^r$ :

$$a_t^* = \arg \max_{a \in \mathcal{F}_t} Q_t^r(s_t, a) \quad (4)$$

If the filter step results in an empty set of feasible actions ( $\mathcal{F}_t = \emptyset$ ), CAPS selects the action associated with the policy that minimizes the estimated cost:

$$a_t^* = \pi_t^*(s_t), \quad \text{where } \pi_t^* = \arg \min_{\pi \in \mathcal{P}} Q_t^c(s_t, \pi(s_t))$$

Intuitively, the first filtering step keeps any action selected by a policy such that it is safe to take that action and then follow  $\pi^c$  thereafter. That is, after the action is taken there is an assurance in expectation that  $\pi^c$  is a safe fallback policy. The selection step then heuristically selects the action that looks best from a reward maximization point of view.

### 4.2 Training Algorithm for CAPS

Given an offline dataset  $\mathcal{D} = \{(s_i, a_i, r_i, c_i, s'_i)\}_{i=1}^n$ , the CAPS approach aims to efficiently learn a set of policies  $\mathcal{P} = \{\pi^r, \pi^1, \dots, \pi^{K-2}, \pi^c\}$  and the two Q-value functions  $Q^r$  and  $Q^c$  that yield strong test time performance. A key feature of CAPS is that training is done via a reduction to standard Offline-RL. In particular, CAPS can be combined with any Offline-RL algorithm that produces a Q-function and a mechanism to extract a policy for that Q-function. In what follows, we first describe the high-level training schema and then provide two concrete instantiations that are used in our experiments. Finally, we discuss the network architecture design and shared structure among the  $K$  policies.

**Reduction to Offline RL.** The training approach for CAPS can be wrapped around existing offline RL algorithms with minimal changes, making it straightforward to implement. The general recipe consists of the following two steps:

1. Train the reward-only value function  $Q^r$  and cost-only value function  $Q^c$ . For  $Q^c$ , we train using cost data from the offline dataset by defining the rewards as costs.

2. Extract the reward maximizing policy  $\pi^r$  from  $Q^r$ , the cost minimizing policy  $\pi^c$  from  $Q^c$ , and the other  $K - 2$ ,  $\pi^k$ , policies with different reward and cost trade-offs from the mixture Q-value function  $Q^r - \lambda_k Q^c$  following the same policy extraction procedure in the offline RL method, where  $\lambda_k$  is the scalarization parameter.

Thus, the approach requires just two full runs of an Offline RL algorithm and  $K$  runs of the much cheaper policy extraction step. It is important to note that the Q-functions for costs and rewards are not learned for each  $\pi^k$ , which means that the precise cost-reward trade-off of each policy is unknown. Thus, this approach to producing each  $\pi^k$  is a fast heuristics for identifying a set of policies that are likely to span a range of cost-reward trade-offs. Another advantage of this reduction approach is that we can leverage better offline RL algorithms to further improve the performance of CAPS and our experiments demonstrate this advantage. While it is possible to apply OPE methods to learn these Q-functions and use them for decision making, our experiments show that while safety is maintained, the rewards are significantly lower compared to our CAPS decision procedure.

**Two CAPS Instantiations.** Our two instantiations of CAPS training are based on two state-of-the-art offline RL methods: IQL and SAC+BC.

*IQL Instantiation:* We employ the Implicit Q-Learning (IQL) algorithm (Kostrikov, Nair, and Levine 2022) as follows. Given the offline dataset  $\mathcal{D} = \{(s_i, a_i, r_i, c_i, s'_i)\}_{i=1}^n$ , we additionally train a value function  $V^c$  and a Q-value function  $Q^c$  for costs. These functions are trained similarly to the original IQL reward value  $V^r$  and Q-value  $Q^r$  functions by minimizing the expectile loss, thereby avoiding explicit value queries for unseen actions. Next, we extract each policy  $\pi \in \mathcal{P}$  by maximizing the following objective.

$$J_{\pi^k}(\phi) = \mathbb{E}_{s, a \sim \mathcal{D}} [\exp(\beta(A^r - \lambda_k A^c)) \log \pi_\phi(a|s)] \quad (5)$$

where the advantages are defined as:

$$A^r = Q_\theta^r(s, a) - V_\psi^r(s), \quad A^c = Q_\theta^c(s, a) - V_\psi^c(s) \quad (9)$$

where  $\theta$  and  $\psi$  correspond to the parameters of the Q-value and value function network representation. Note that for  $\pi^r$  and  $\pi^c$ , we only use  $A^r$  and  $-A^c$  respectively.

*SAC+BC Instantiation:* (Fujimoto and Gu 2021) recently described a general framework for developing offline RL algorithms by combining an off-policy RL with a behavior cloning (BC) regularization term. We follow the same principle but employ soft-actor-critic (SAC) (Haarnoja et al. 2018) algorithm as the off-policy RL component instead of the TD3 algorithm used originally. Following the offline reduction recipe, we update the actor network to incorporate  $K$  heads and define an additional Q-value function for costs. We then learn each policy  $\pi^k$  by maximizing the following objective:

$$J_{\pi^k}(\phi) = \mathbb{E}_{s, a \sim \mathcal{D}, \tilde{a} \sim \pi_\phi(\cdot|s)} \left[ Q^k(s, \tilde{a}) - \alpha \log \pi_\phi(\tilde{a}|s) - (\tilde{a} - a)^2 \right] \quad (6)$$

where  $Q^k(s, \tilde{a}) = Q_\theta^r(s, \tilde{a}) - \lambda_k Q_\theta^c(s, \tilde{a})$ . For  $\pi^r$  and  $\pi^c$ , we use  $Q^r$  and  $-Q^c$  to update their corresponding heads.

**Improving CAPS via Shared Actor Representation.** To enhance knowledge transfer and/or training efficiency, we utilize a shared neural network architecture for the  $K$  different policies. Instead of training separate networks, we employ a common body with parameters  $\phi_s$  that learns a unified state representation  $f_{\phi_s}(s)$ . This body is paired with distinct output heads (one for each of the  $K$  policies), each with its own parameters  $\phi_k$ , which specialize in predicting the action distributions for their respective policies:  $\pi^k(a|s) = g_{\phi_k}(f_{\phi_s}(s)), \forall k \in [K]$ .

There are two synergistic benefits of a shared policy representation. First, this approach captures general features relevant to both reward and cost objectives, leading to knowledge transfer across policies. Second, CAPS requires only one round of training to adapt to different cost constraint thresholds. This saves significant training time compared to methods that need re-training for each cost constraint.

## 5 Safety Guarantee

This section describes conditions under which the CAPS decision-making process has safety guarantees. For simplicity, we consider the case where the optimal-cost Q-function is perfectly estimated, noting that it is straightforward to modify our result to account for bounded estimation error.

We focus on the finite-horizon setting with horizon  $T$  and let  $Q_t^c$  and  $V_t^c$  denote the cost optimal Q and value functions for time step  $t \in \{0, \dots, T\}$  respectively. Here we treat these functions as positive cost accumulation functions that are minimized during optimization. We consider policies of the form  $\pi_t(s, c_{<t})$  where  $s$  is the current state and  $c_{<t}$  is the accumulated cost before arriving at time step  $t$ . For a given cost safety bound  $\kappa$ , our guarantee applies to any policy that is  $\kappa$ -admissible.

**Assumption 1** *A policy  $\pi$  is  $\kappa$ -admissible if for any state  $s$ , any time step  $t$ , and any accumulated cost  $c_{<t}$ , we have  $Q_t^c(s, \pi_t(s, c_{<t})) \leq \max\{V_t^c(s), \kappa - c_{<t}\}$ .*

Note that by assuming a perfect  $Q^c$ , the switching policy defined for CAPS is  $\kappa$ -admissible since it either picks an action  $a$  such that  $Q_t^c(s, a) \leq \kappa - c_{<t}$  if possible, or resorts to choosing the least cost action, which has a value of  $V_t^c(s)$  since  $Q^c$  and  $V^c$  are cost optimal. This is true of CAPS for any set of  $K$  policies  $\mathcal{P}$  that contains the cost optimal policy  $\pi_c$  and any reward value function  $Q_t^r$ .

Our safety guarantees are specified in terms of a parameter  $\epsilon$  that characterizes the underlying Markov decision process (MDP). In particular, safety in terms of expected cost is challenging when stochastic state transitions can result in a set of possible next states that have wildly different values of  $V^c$ . Indeed, counter examples to safety can be constructed that involve transitions with very small probabilities of reaching enormous values of  $V^c$  while the remaining probability mass results in very small costs. For this purpose, we define the notion of *optimal-cost variation*.

**Assumption 2** *Given an MDP  $M$ , let  $N(s, a) = \{s' \mid P(s, a, s') > 0\}$  be the set of possible next states after taking action  $a$  in state  $s$ .  $M$  has an optimal-cost variation of  $\epsilon$  if each state-action-timestep tuple  $(s, a, t)$  satisfies  $\max_{s' \in N(s, a)} V_t^c(s') - \min_{s' \in N(s, a)} V_t^c(s') \leq \epsilon$ .*

Note that any deterministic MDP has  $\epsilon = 0$ . Also note that this assumption places no constraints between the optimal costs of states resulting from different  $(s, a, t)$  tuples. We can now state our main result which bounds the total cost of a policy  $\pi$ , denoted  $V_t^{\pi,c}$ .

**Theorem 1** *For any MDP  $M$  with optimal-cost variation  $\epsilon$  and any  $\kappa$ -admissible policy  $\pi$ , we have that for any state  $s$ , time step  $t$ , and accumulated cost  $c_{<t}$ ,*

$$V_t^{\pi,c}(s) \leq \max\{V_t^c(s), \kappa - c_{<t}\} + (T - t)\epsilon$$

See the supplementary material for the full proof.

This result says that up to a constant factor either the expected total cost of policy  $\pi$  will be bounded by  $\kappa$  over a horizon of  $T$  or if that is not possible, it will be bounded by the lowest possible expected cost  $V_t^c(s)$ . Based on this result, if one has an estimate of  $\epsilon$  for the system, an appropriate value of  $\kappa$  can be selected to provide a desired guarantee. The constant factor  $(T - t)\epsilon$  is simply a constant of  $\epsilon$  accumulated over the remaining horizon  $T - t$ . Note that for time step 0, we get a bound of  $\kappa$  when it is possible to achieve that safety factor. Also, note that in the deterministic case, the bound is tight since  $\epsilon = 0$ .

## 6 Experiments and Results

This section presents experimental evaluation of the CAPS approach and comparison with state-of-the-art methods.

### 6.1 Experimental Setup

**Benchmarks.** We employ 38 sequential decision-making benchmarks of varying difficulty from Safety-Gymnasium (Ray, Achiam, and Amodei 2019; Ji et al. 2024), Bullet-Safety-Gym (Gronauer 2022), and MetaDrive (Li et al. 2022) within the DSRL framework (Liu et al. 2024). Further details are provided in Appendix C.1.

**Configuration of CAPS.** We consider two instantiations of CAPS training with different offline RL methods: CAPS (IQL) employs IQL (Kostrikov, Nair, and Levine 2022) and CAPS (SAC+BC) employs soft-actor-critic based off-policy RL combined with a behavior cloning regularization term (Haarnoja et al. 2018; Fujimoto and Gu 2021) as described in the technical section. We consider CAPS with different number of policies  $K \in \{2, 4, 8\}$  to determine the best configuration in terms of computational cost and quality of decisions. We provide the details of the neural network structure used for value and Q-functions, policy heads, and hyper-parameters in the Appendix C.3.

**Baseline methods.** We compare CAPS with several state-of-the-art baseline methods. **1)** BC (Behavior Cloning) serves as a straightforward baseline where policies are learned by mimicking the behavior observed in the training data. **2)** BEAR-Lag integrates BEAR (Wu, Tucker, and Nachum 2020), an offline RL method, with a Lagrangian approach to address safety constraints. **3)** Constraints-penalized Q-learning (CPQ) treats out-of-distribution actions as unsafe and updates the Q-value function using only safe state-action pairs (Xu, Zhan, and Zhu 2022). **4)** COptiDICE builds upon the distribution correction estimation (DICE) method to

incorporate cost constraints in offline RL (Lee et al. 2022). **5)** CDT employs a decision transformer framework to consider safety constraints for safe decision-making (Liu et al. 2023). This is the only baseline that allows varying cost constraints at test-time. We run each over three seeds.

**Evaluation methodology.** We evaluate CAPS and the above-mentioned baseline methods using normalized return and normalized cost, where a normalized cost below 1 indicates that safety constraints are satisfied. In line with the DSRL guidelines, safety is our primary evaluation metric and our goal is to maximize rewards while adhering to cost constraints. We evaluate each algorithm using three different target cost threshold configurations, three random seeds, and twenty episodes. The configurations are:  $\{10, 20, 40\}$ ,  $\{20, 40, 80\}$ , where the second cost limit set applies to the more challenging Safety-Gymnasium environments. We compute the average normalized reward and cost to assess the performance. Note that, for all baseline methods except CDT, a separate agent is trained for each cost threshold.

### 6.2 Results and Discussion

**Ablation for number of policies.** Tables 3 and 4 in the Appendix B.1 show the results of CAPS- $K$  with  $K \in \{2, 4, 8\}$  heads/policies. We make the following observations. **1)** CAPS-4 with four policies achieves better results than CAPS-2 with two policies. CAPS-8 with eight policies despite its potential for increased capacity, does not lead to substantial improvements over CAPS-4. The diminishing returns demonstrate that the incremental benefit of adding more policies is decreasing, with CAPS-4 providing an optimized balance of flexibility and complexity. Eight policies introduce some redundancy and potentially exacerbates stochasticity, rather than providing additional value. **2)** CAPS-2 with two policies also demonstrates impressive results and is preferred as our main approach due to its simplicity and computational efficiency. Switching between reward-maximizing and cost-minimizing policies is very effective for managing varying test-time cost constraints. In contrast, CAPS variant with more than two heads involve optimizing additional policies for combined reward and cost objectives with extra scalarization hyper-parameters  $\lambda_k$ . Therefore, we use CAPS with two policies (one for reward maximization and one for cost minimization) for the rest of our experimental analysis.

**Ablation for shared policy representation.** We investigate the advantages of using a shared policy representation architecture compared to training separate agents for reward and cost objectives respectively. A shared representation allows different heads/policies to utilize common state features, to potentially improve the overall performance. The intuition behind this idea is that the shared network enables the cost head to achieve higher rewards and the reward head to incur lower costs compared to independently learned policies.

For instance, Figure 1 compares the performance of separate agents and shared agents in the CarGoal1 environment using the two instantiations: SAC+BC (left) and IQL (right). The focus is on understanding how sharing impacts the behavior of cost agents and reward agents when run uniformly.

Task	BC		BEAR-Lag		CPQ		CDT		COptiDICE		CAPS (SAC+BC)		CAPS (IQL)	
	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 $\downarrow$
PointButton1	0.11	1.06	0.24	1.95	0.64	3.31	0.53	2.99	0.12	1.43	0.10	1.23	<b>0.03</b>	<b>0.50</b>
PointButton2	0.25	2.30	0.41	2.63	0.57	4.47	0.45	2.85	0.17	1.93	0.21	1.54	<b>0.14</b>	<b>0.75</b>
PointCircle1	0.81	4.97	0.69	3.04	0.45	1.14	<b>0.58</b>	<b>0.82</b>	0.85	5.39	<b>0.53</b>	<b>0.36</b>	<b>0.50</b>	<b>0.78</b>
PointCircle2	0.71	5.49	0.70	4.64	0.09	4.97	0.63	1.33	0.87	8.70	0.61	1.38	<b>0.51</b>	<b>0.80</b>
PointGoal1	<b>0.56</b>	<b>0.91</b>	0.74	1.20	<b>0.57</b>	<b>0.66</b>	0.71	1.03	0.49	1.54	<b>0.17</b>	<b>0.11</b>	<b>0.47</b>	<b>0.53</b>
PointGoal2	0.54	2.72	0.70	3.12	0.52	1.86	0.55	1.98	0.41	1.83	<b>0.20</b>	<b>0.55</b>	<b>0.33</b>	<b>0.80</b>
PointPush1	<b>0.18</b>	<b>0.77</b>	<b>0.21</b>	<b>0.92</b>	<b>0.26</b>	<b>0.75</b>	<b>0.26</b>	<b>0.94</b>	0.13	1.08	<b>0.19</b>	<b>0.38</b>	<b>0.19</b>	<b>0.29</b>
PointPush2	0.15	1.39	0.17	1.16	0.09	1.57	0.21	1.13	0.04	1.17	<b>0.15</b>	<b>0.70</b>	<b>0.13</b>	<b>0.64</b>
CarButton1	0.00	1.54	0.22	2.76	0.42	10.20	0.21	2.71	-0.09	1.36	0.01	1.17	<b>-0.01</b>	<b>0.30</b>
CarButton2	-0.03	1.38	0.00	2.00	0.36	12.65	0.19	3.53	-0.06	1.27	-0.14	1.17	<b>-0.08</b>	<b>0.32</b>
CarCircle1	0.71	5.67	0.76	5.39	0.09	1.10	0.59	2.90	0.70	5.86	0.64	3.24	0.54	1.51
CarCircle2	0.73	7.60	0.74	6.68	0.47	2.34	0.65	4.22	0.75	8.34	0.66	4.98	0.50	1.55
CarGoal1	<b>0.39</b>	<b>0.42</b>	0.60	1.10	0.80	1.51	0.67	1.18	<b>0.38</b>	<b>0.53</b>	<b>0.28</b>	<b>0.47</b>	<b>0.33</b>	<b>0.38</b>
CarGoal2	<b>0.23</b>	<b>0.78</b>	0.29	1.27	0.56	3.64	0.48	2.17	<b>0.22</b>	<b>0.82</b>	<b>0.21</b>	<b>0.75</b>	<b>0.16</b>	<b>0.62</b>
CarPush1	<b>0.19</b>	<b>0.39</b>	<b>0.21</b>	<b>0.45</b>	0.11	1.11	<b>0.31</b>	<b>0.81</b>	<b>0.16</b>	<b>0.46</b>	<b>0.22</b>	<b>0.34</b>	<b>0.20</b>	<b>0.31</b>
CarPush2	0.13	1.27	0.11	1.06	0.19	4.59	0.19	1.88	0.09	1.48	<b>0.14</b>	<b>0.82</b>	<b>0.07</b>	<b>0.75</b>
SwimmerVelo	0.51	2.62	0.27	1.83	0.18	2.04	<b>0.63</b>	<b>0.88</b>	0.61	6.31	<b>0.41</b>	<b>0.51</b>	0.43	1.58
HopperVelo	0.54	2.80	0.36	5.47	0.12	1.73	<b>0.71</b>	<b>0.75</b>	0.10	1.45	<b>0.62</b>	<b>0.71</b>	<b>0.41</b>	<b>0.70</b>
HalfCheetahVelo	0.96	5.26	0.98	6.55	<b>0.34</b>	<b>0.66</b>	<b>0.98</b>	<b>0.16</b>	<b>0.62</b>	<b>0.00</b>	<b>0.95</b>	<b>0.21</b>	<b>0.94</b>	<b>0.77</b>
Walker2dVelo	0.74	1.85	0.89	4.16	<b>0.03</b>	<b>0.23</b>	<b>0.79</b>	<b>0.10</b>	<b>0.12</b>	<b>0.86</b>	<b>0.80</b>	<b>0.04</b>	<b>0.80</b>	<b>0.62</b>
AntVelo	0.98	2.58	<b>-1.01</b>	<b>0.00</b>	<b>-1.01</b>	<b>0.00</b>	<b>0.99</b>	<b>0.50</b>	1.00	2.99	<b>0.89</b>	<b>0.67</b>	<b>0.95</b>	<b>0.64</b>
<b>SafetyGym Avg</b>	0.45	2.56	0.39	2.73	0.28	2.88	0.54	1.66	0.37	2.61	0.37	1.02	<b>0.36</b>	<b>0.72</b>
BallRun	0.67	3.16	0.02	3.11	0.30	2.36	0.38	1.02	0.57	3.38	<b>0.24</b>	<b>0.60</b>	<b>0.19</b>	<b>0.94</b>
CarRun	<b>0.97</b>	<b>0.11</b>	0.40	7.65	<b>0.93</b>	<b>0.94</b>	<b>0.98</b>	<b>0.63</b>	<b>0.93</b>	<b>0.00</b>	<b>0.97</b>	<b>0.00</b>	<b>0.97</b>	<b>0.25</b>
DroneRun	0.32	1.44	0.42	3.85	0.36	2.62	<b>0.62</b>	<b>0.74</b>	0.67	4.27	<b>0.54</b>	<b>0.39</b>	0.47	2.19
AntRun	0.72	2.76	<b>0.04</b>	<b>0.08</b>	<b>0.02</b>	<b>0.01</b>	<b>0.72</b>	<b>0.82</b>	<b>0.62</b>	<b>1.00</b>	<b>0.53</b>	<b>0.25</b>	<b>0.61</b>	<b>0.90</b>
BallCircle	0.76	2.70	0.80	2.22	<b>0.62</b>	<b>0.57</b>	0.78	1.20	0.70	2.63	<b>0.59</b>	<b>0.19</b>	<b>0.69</b>	<b>0.59</b>
CarCircle	0.62	3.15	0.76	2.78	<b>0.71</b>	<b>0.27</b>	<b>0.75</b>	<b>0.91</b>	0.48	2.93	<b>0.61</b>	<b>0.29</b>	<b>0.69</b>	<b>0.65</b>
DroneCircle	0.71	2.70	0.80	3.93	<b>-0.23</b>	<b>0.36</b>	0.62	1.02	0.25	1.02	<b>0.42</b>	<b>0.14</b>	<b>0.55</b>	<b>0.67</b>
AntCircle	0.65	4.67	0.64	5.23	<b>0.00</b>	<b>0.00</b>	0.48	2.34	0.18	4.33	<b>0.47</b>	<b>0.42</b>	<b>0.41</b>	<b>0.15</b>
<b>BulletGym Avg</b>	0.68	2.59	0.49	3.61	<b>0.34</b>	<b>0.89</b>	0.67	1.09	0.55	2.45	<b>0.55</b>	<b>0.29</b>	<b>0.57</b>	<b>0.79</b>
easysparse	<b>0.18</b>	<b>0.82</b>	<b>0.03</b>	<b>0.71</b>	<b>-0.06</b>	<b>0.08</b>	<b>0.25</b>	<b>0.26</b>	0.92	5.33	0.54	2.14	<b>0.11</b>	<b>0.34</b>
easymeans	<b>0.05</b>	<b>0.18</b>	<b>0.02</b>	<b>0.50</b>	<b>-0.06</b>	<b>0.06</b>	0.49	1.18	0.73	4.35	<b>0.16</b>	<b>0.41</b>	<b>0.01</b>	<b>0.20</b>
easydense	<b>0.20</b>	<b>0.31</b>	<b>-0.01</b>	<b>0.39</b>	<b>-0.06</b>	<b>0.08</b>	<b>0.41</b>	<b>0.51</b>	0.53	3.02	<b>0.11</b>	<b>0.15</b>	<b>0.10</b>	<b>0.19</b>
mediumsparse	0.65	1.56	<b>-0.01</b>	<b>0.31</b>	<b>-0.09</b>	<b>0.06</b>	0.55	1.03	0.79	2.65	<b>0.65</b>	<b>0.99</b>	<b>0.60</b>	<b>0.74</b>
mediummean	0.74	1.55	<b>-0.04</b>	<b>0.10</b>	<b>-0.08</b>	<b>0.07</b>	0.48	1.31	0.78	2.52	<b>0.13</b>	<b>0.23</b>	<b>0.66</b>	<b>0.94</b>
mediumdense	0.63	1.08	<b>0.08</b>	<b>0.63</b>	<b>-0.07</b>	<b>0.05</b>	<b>0.21</b>	<b>0.29</b>	0.63	2.17	<b>0.69</b>	<b>0.80</b>	<b>0.69</b>	<b>0.56</b>
hardsparse	<b>0.19</b>	<b>0.54</b>	<b>0.02</b>	<b>0.24</b>	<b>-0.04</b>	<b>0.08</b>	<b>0.26</b>	<b>0.63</b>	0.38	2.31	<b>0.10</b>	<b>0.18</b>	<b>0.45</b>	<b>0.72</b>
hardmean	0.47	2.32	<b>0.01</b>	<b>0.26</b>	<b>-0.05</b>	<b>0.07</b>	<b>0.19</b>	<b>0.35</b>	0.36	2.37	<b>0.19</b>	<b>0.30</b>	<b>0.28</b>	<b>0.25</b>
harddense	0.36	1.62	0.07	3.86	<b>-0.04</b>	<b>0.05</b>	<b>0.28</b>	<b>0.69</b>	0.23	1.46	<b>0.22</b>	<b>0.45</b>	<b>0.37</b>	<b>0.67</b>
<b>MetaDrive Avg</b>	0.39	1.11	<b>0.02</b>	<b>0.78</b>	<b>-0.06</b>	<b>0.07</b>	<b>0.35</b>	<b>0.69</b>	0.59	2.91	<b>0.31</b>	<b>0.63</b>	<b>0.36</b>	<b>0.51</b>

Table 1: Results for normalized rewards and costs. The cost threshold is 1. The  $\uparrow$  symbol denotes that higher the rewards, the better. The  $\downarrow$  symbol denotes that the lower the costs (up to threshold 1), the better. Each value is averaged over 3 distinct cost thresholds, 20 evaluation episodes, and 3 random seeds. **Bold**: Safe agents whose normalized cost  $\leq 1$ . Gray: Unsafe agents. **Blue**: Safe agent with the highest reward. We compare CAPS instantiations with other baselines, not with each other.

In the case of separately trained agents, joined by the red line in both plots, the reward agent maximizes rewards without much consideration for the cost objective, resulting in high costs for achieving high rewards. For instance, in the SAC+BC algorithm, the reward agent’s point shows a high cost of around 80 and a moderate reward of about 0.1, indicating a reckless pursuit of rewards. Similarly, in the IQL algorithm, the reward agent incurs a cost of about 25 to achieve a reward of approximately 0.34. Conversely, the cost agent in the separate agents’ scenario aggressively minimizes costs, leading to low rewards as it avoids actions that could increase costs. This conservative behavior is evident the SAC+BC case, where the cost agent achieves low rewards of around -0.1 and costs near 0, and the IQL, where the cost agent’s rewards are around 0.3 with costs around 10.

In contrast, the shared agents, joined by the blue lines, ex-

hibit a more balanced approach. The shared policy representation helps cost agent become less conservative, encouraging them to seek higher rewards even if it means incurring some costs. For example, in the SAC+BC algorithm, the shared cost agent achieves more rewards of about 0.33 with increased costs of around 20, showing a willingness to balance both objectives. Similarly, in the IQL algorithm, the shared cost agent achieves higher rewards of approximately 0.36 with moderate costs of around 15. A similar balanced approach is also evident for the reward agent under sharing.

Tables 6, 7, 8 and 9 in the Appendix B.2 shows the results for shared vs. independent policies ablations on all 38 benchmarks for IQL and SAC+BC. Additionally, the performance of CAPS with and without sharing is detailed in Table 5. These results demonstrate the benefits of the shared policy representation. Therefore, we use CAPS with shared

Shared vs Separate Agents: CarGoal1 Environment

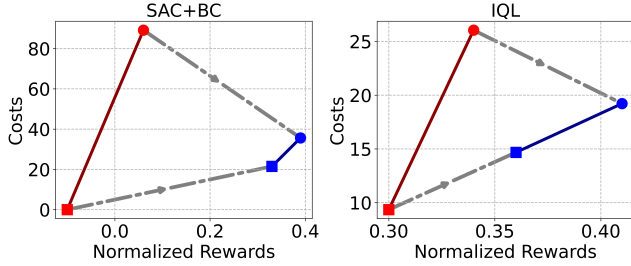


Figure 1: Ablation results for **shared** architecture vs. **independently** trained reward ● and cost ■ optimized policies.

Cost Limits	CAPS(IQL)	CAPS(SAC+BC)	CDT
{ 5, 10 }	18 / 38	15 / 38	11 / 38
{ 10, 20 }	28 / 38	24 / 38	16 / 38
{ 15, 30 }	31 / 38	29 / 38	19 / 38
{ 20, 40 }	33 / 38	32 / 38	20 / 38
{ 30, 60 }	33 / 38	34 / 38	22 / 38
{ 40, 80 }	33 / 38	35 / 38	25 / 38

Table 2: Performance of CAPS(IQL) and CAPS(SAC+BC) for different cost limit configurations.

representation for two policies for baseline comparison.

**Comparison with baselines.** Table 1 compares the two CAPS variants, namely, CAPS(IQL) and CAPS(SAC+BC) against all baseline methods. Each method’s performance is displayed across two objectives: reward maximization and cost minimization. The table includes results for various tasks and aggregated performance for the three environment categories (SafetyGym Avg, BulletGym Avg, MetaDrive Avg). We make the following observations. **1)** Both CAPS instantiations significantly outperform the baselines, meeting cost constraints in 34 out of 38 tasks for CAPS(IQL) and 30 out of 38 tasks for CAPS(SAC+BC). This is remarkable given that safety is the primary metric. In contrast, the second-best method, CDT, only manages to satisfy the constraints in 19 out of 38 tasks. **2)** CAPS(IQL) achieves the highest reward in 18 out of 38 tasks, demonstrating superior performance in reward maximization while meeting the safety/cost constraints. This result shows that CAPS can leverage better offline RL algorithms to further improve the performance in OSRL. **3)** Results in Appendix C.2 show that the training time of CAPS with increasing number of heads/policies is significantly lower than CDT. Other baseline methods will be worse as they need to be trained for every cost constraint. **4)** CAPS with two policies is a minimalist approach and a strong wrapper based baseline for OSRL.

**Ablation on cost limits.** To evaluate the performance of CAPS under various cost limits, we compare both CAPS variants with CDT. The results in Table 2 demonstrate that CAPS handles different cost constraints better than CDT.

For a cost limit configuration of {5, 10}, where the cost limit is 10 for SafetyGym environments and 5 for the other two categories, CAPS(IQL) maintains safety in 18 out of 38

tasks, while CAPS(SAC+BC) achieves safety in 15 out of 38 tasks (compared to 11 for CDT). Although these numbers might appear modest, they are reasonable given that the benchmark was designed with higher cost thresholds for evaluation. When the cost limit is increased to {10, 20}, both CAPS variants show significant improvement: CAPS(IQL) ensures safety in 28 out of 38 tasks, and CAPS(SAC+BC) in 24 out of 38 tasks. With a configuration of {15, 30}, CAPS(IQL) achieves safety in 31 out of 38 tasks, while CAPS(SAC+BC) secures safety in 29 out of 38 tasks. Performance peaks at the {20, 40}, {30, 60}, and {40, 80} configurations. CAPS(IQL) maintains safety in 33 out of 38 tasks at all three configurations, while CAPS(SAC+BC) shows consistent improvement, achieving safety in 32, 34, and 35 out of 38 tasks, respectively. The results demonstrate that both CAPS instantiations can effectively adapt to various cost limits without requiring retraining. In contrast to other methods, which often necessitate specific training for each cost limit, both CAPS variants demonstrate flexibility and robustness, consistently outperforming CDT across all cost limit configurations. Full results comparing CDT and CAPS(IQL) can be found in Appendix B.3.

**Off-policy evaluation.** We implemented a variant of CAPS with off-policy evaluation for the two policies/heads configuration, comparing this method to our original approach. We trained two Q-functions for each head/policy (i.e.,  $\pi_r$  and  $\pi_c$ ) using fitted Q-evaluation (FQE) (Le, Voloshin, and Yue 2019):  $\hat{Q}^r$  for rewards and  $\hat{Q}^c$  for costs, applying the same logic for policy switching. In Appendix B.4, we present results comparing the original CAPS method to two variants: 1) CAPS with FQE based  $\hat{Q}^r$  and  $\hat{Q}^c$  for both  $\pi_r$  and  $\pi_c$  (reward-cost FQE), and 2) CAPS with FQE based  $\hat{Q}^r$  for  $\pi_r$  and  $\pi_c$  (reward FQE). The reward-cost FQE method yields overly conservative results with low rewards. The learned  $\hat{Q}^c$  corresponding to respective policies  $\pi_r$  and  $\pi_c$  frequently estimates that actions exceed the cost limit, defaulting to selecting the action with the lowest cost Q-value. Our results also show that reward FQE method which use  $\hat{Q}^r$  functions from FQE and  $Q^c$  from offline RL method yields better performance than reward-cost FQE and is comparable to the original CAPS (uses  $Q^c$  and  $Q^r$  from offline RL). We hypothesize that this behavior is due to estimation errors in off-policy evaluation as noted by prior literature (Figueiredo Prudencio, Maximo, and Colombini 2024). The FQE based  $\hat{Q}^r$  functions effectively differentiate between actions without needing a precise estimation (unlike cost where precision matters to check safety of decisions at a state). These findings support the robust choice of applying  $Q^c$  in CAPS’s filtering stage.

## 7 Summary and Future Work

This paper introduces CAPS, a novel offline safe RL framework that dynamically adapts to varying safety constraints by training multiple diverse policies. CAPS switches between actions from policies during deployment, ensuring safety while maximizing rewards. Empirical results show that CAPS outperforms existing methods. Future work could integrate CAPS into online learning for real-time gains.

## Acknowledgements

The authors gratefully acknowledge the in part support by USDA-NIFA funded AgAID Institute award 2021-67021-35344. The views expressed are those of the authors and do not reflect the official policy or position of the USDA-NIFA.

## References

- An, G.; Moon, S.; Kim, J.-H.; and Song, H. O. 2021. Uncertainty-based offline reinforcement learning with diversified Q-ensemble. *Advances in neural information processing systems*, 34: 7436–7447.
- Bai, C.; Wang, L.; Yang, Z.; Deng, Z.-H.; Garg, A.; Liu, P.; and Wang, Z. 2022. Pessimistic Bootstrapping for Uncertainty-Driven Offline Reinforcement Learning. In *International Conference on Learning Representations*.
- Chen, H.; Lu, C.; Ying, C.; Su, H.; and Zhu, J. 2023. Offline Reinforcement Learning via High-Fidelity Generative Behavior Modeling. In *The Eleventh International Conference on Learning Representations*.
- Figueiredo Prudencio, R.; Maximo, M. R. O. A.; and Colombini, E. L. 2024. A Survey on Offline Reinforcement Learning: Taxonomy, Review, and Open Problems. *IEEE Transactions on Neural Networks and Learning Systems*, 35(8): 10237–10257.
- Fujimoto, S.; and Gu, S. S. 2021. A minimalist approach to offline reinforcement learning. *Advances in neural information processing systems*, 34: 20132–20145.
- Garcia, J.; and Fernández, F. 2015. A Comprehensive Survey on Safe Reinforcement Learning. *Journal of Machine Learning Research*, 16(1): 1437–1480.
- Gronauer, S. 2022. Bullet-Safety-Gym: A Framework for Constrained Reinforcement Learning. Technical report, mediaTUM.
- Gu, S.; Yang, L.; Du, Y.; Chen, G.; Walter, F.; Wang, J.; and Knoll, A. 2024. A Review of Safe Reinforcement Learning: Methods, Theories and Applications. *IEEE Transactions on Pattern Analysis and Machine Intelligence*.
- Haarnoja, T.; Zhou, A.; Hartikainen, K.; Tucker, G.; Ha, S.; Tan, J.; Kumar, V.; Zhu, H.; Gupta, A.; Abbeel, P.; et al. 2018. Soft actor-critic algorithms and applications. *arXiv preprint arXiv:1812.05905*.
- Hansen-Estruch, P.; Kostrikov, I.; Janner, M.; Kuba, J. G.; and Levine, S. 2023. Idql: Implicit Q-learning as an actor-critic method with diffusion policies. *arXiv preprint arXiv:2304.10573*.
- Janner, M.; Li, Q.; and Levine, S. 2021. Offline reinforcement learning as one big sequence modeling problem. *Advances in neural information processing systems*, 34: 1273–1286.
- Jaques, N.; Ghandeharioun, A.; Shen, J. H.; Ferguson, C.; Lapedriza, A.; Jones, N.; Gu, S.; and Picard, R. 2020. Way Off-Policy Batch Deep Reinforcement Learning of Human Preferences in Dialog.
- Ji, J.; Zhou, J.; Zhang, B.; Dai, J.; Pan, X.; Sun, R.; Huang, W.; Geng, Y.; Liu, M.; and Yang, Y. 2024. OmniSafe: An Infrastructure for Accelerating Safe Reinforcement Learning Research. *Journal of Machine Learning Research*, 25(285): 1–6.
- Kidambi, R.; Rajeswaran, A.; Netrapalli, P.; and Joachims, T. 2020. MOREL: Model-Based Offline Reinforcement Learning. In Larochelle, H.; Ranzato, M.; Hadsell, R.; Balcan, M.; and Lin, H., eds., *Advances in Neural Information Processing Systems*, volume 33, 21810–21823. Curran Associates, Inc.
- Kostrikov, I.; Nair, A.; and Levine, S. 2022. Offline Reinforcement Learning with Implicit Q-Learning. In *International Conference on Learning Representations*.
- Kumar, A.; Fu, J.; Soh, M.; Tucker, G.; and Levine, S. 2019. Stabilizing Off-Policy Q-Learning via Bootstrapping Error Reduction. *Advances in neural information processing systems*, 32.
- Le, H.; Voloshin, C.; and Yue, Y. 2019. Batch policy learning under constraints. In *International Conference on Machine Learning*, 3703–3712. PMLR.
- Lee, J.; Paduraru, C.; Mankowitz, D. J.; Heess, N.; Precup, D.; Kim, K.-E.; and Guez, A. 2022. COptiDICE: Offline Constrained Reinforcement Learning via Stationary Distribution Correction Estimation. In *International Conference on Learning Representations*.
- Levine, S.; Kumar, A.; Tucker, G.; and Fu, J. 2020. Offline reinforcement learning: Tutorial, review, and perspectives on open problems. *arXiv preprint arXiv:2005.01643*.
- Li, Q.; Peng, Z.; Feng, L.; Zhang, Q.; Xue, Z.; and Zhou, B. 2022. Metadrive: Composing diverse driving scenarios for generalizable reinforcement learning. *IEEE transactions on pattern analysis and machine intelligence*, 45(3): 3461–3475.
- Li, Y. 2017. Deep reinforcement learning: An overview. *arXiv preprint arXiv:1701.07274*.
- Lin, Q.; Tang, B.; Wu, Z.; Yu, C.; Mao, S.; Xie, Q.; Wang, X.; and Wang, D. 2023. Safe offline reinforcement learning with real-time budget constraints. In *International Conference on Machine Learning*, 21127–21152. PMLR.
- Liu, Z.; Guo, Z.; Lin, H.; Yao, Y.; Zhu, J.; Cen, Z.; Hu, H.; Yu, W.; Zhang, T.; Tan, J.; and Zhao, D. 2024. Datasets and Benchmarks for Offline Safe Reinforcement Learning. *Journal of Data-centric Machine Learning Research*.
- Liu, Z.; Guo, Z.; Yao, Y.; Cen, Z.; Yu, W.; Zhang, T.; and Zhao, D. 2023. Constrained decision transformer for offline safe reinforcement learning. In *International Conference on Machine Learning*, 21611–21630. PMLR.
- Lyu, J.; Ma, X.; Li, X.; and Lu, Z. 2022. Mildly conservative Q-learning for offline reinforcement learning. *Advances in Neural Information Processing Systems*, 35: 1711–1724.
- Mnih, V.; Kavukcuoglu, K.; Silver, D.; Rusu, A. A.; Veness, J.; Bellemare, M. G.; Graves, A.; Riedmiller, M.; Fidjeland, A. K.; Ostrovski, G.; et al. 2015. Human-level control through deep reinforcement learning. *Nature*, 518(7540): 529–533.
- Polosky, N.; Da Silva, B. C.; Fiterau, M.; and Jagannath, J. 2022. Constrained offline policy optimization. In *International Conference on Machine Learning*, 17801–17810. PMLR.

Ray, A.; Achiam, J.; and Amodei, D. 2019. Benchmarking safe exploration in deep reinforcement learning. *arXiv preprint arXiv:1910.01708*.

Rigter, M.; Lacerda, B.; and Hawes, N. 2022. Rambo-rl: Robust adversarial model-based offline reinforcement learning. *Advances in neural information processing systems*, 35: 16082–16097.

Silver, D.; Hubert, T.; Schrittwieser, J.; Antonoglou, I.; Lai, M.; Guez, A.; Lanctot, M.; Sifre, L.; Kumaran, D.; Graepel, T.; et al. 2018. A general reinforcement learning algorithm that masters chess, shogi, and Go through self-play. *Science*, 362(6419): 1140–1144.

Wang, K.; Zhao, H.; Luo, X.; Ren, K.; Zhang, W.; and Li, D. 2022. Bootstrapped transformer for offline reinforcement learning. *Advances in Neural Information Processing Systems*, 35: 34748–34761.

Wu, J.; Wu, H.; Qiu, Z.; Wang, J.; and Long, M. 2022. Supported policy optimization for offline reinforcement learning. *Advances in Neural Information Processing Systems*, 35: 31278–31291.

Wu, Y.; Tucker, G.; and Nachum, O. 2020. Behavior Regularized Offline Reinforcement Learning.

Xu, H.; Jiang, L.; Jianxiong, L.; and Zhan, X. 2022. A policy-guided imitation approach for offline reinforcement learning. *Advances in Neural Information Processing Systems*, 35: 4085–4098.

Xu, H.; Zhan, X.; and Zhu, X. 2022. Constraints penalized Q-learning for safe offline reinforcement learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, 8753–8760.

Yang, R.; Bai, C.; Ma, X.; Wang, Z.; Zhang, C.; and Han, L. 2022. RORL: Robust Offline Reinforcement Learning via Conservative Smoothing. In Oh, A. H.; Agarwal, A.; Belgrave, D.; and Cho, K., eds., *Advances in Neural Information Processing Systems*.

Yao, Y.; Liu, Z.; Cen, Z.; Zhu, J.; Yu, W.; Zhang, T.; and Zhao, D. 2023. Constraint-Conditioned Policy Optimization for Versatile Safe Reinforcement Learning. In *Thirty-seventh Conference on Neural Information Processing Systems*.

Yu, T.; Thomas, G.; Yu, L.; Ermon, S.; Zou, J. Y.; Levine, S.; Finn, C.; and Ma, T. 2020. MOPO: Model-based Offline Policy Optimization. In Larochelle, H.; Ranzato, M.; Hadsell, R.; Balcan, M.; and Lin, H., eds., *Advances in Neural Information Processing Systems*, volume 33, 14129–14142. Curran Associates, Inc.

Zhang, H.; Xu, W.; and Yu, H. 2023. Policy Expansion for Bridging Offline-to-Online Reinforcement Learning. In *The Eleventh International Conference on Learning Representations*.

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. In *The Twelfth International Conference on Learning Representations*.