

# **$Q$ -PENSIEVE: BOOSTING SAMPLE EFFICIENCY OF MULTI-OBJECTIVE RL THROUGH MEMORY SHARING OF $Q$ -SNAPSHOTS**

**Anonymous authors**

Paper under double-blind review

## **ABSTRACT**

Many real-world continuous control problems are in the dilemma of weighing the pros and cons, multi-objective reinforcement learning (MORL) serves as a generic framework of learning control policies for different preferences over objectives. However, the existing MORL methods either rely on multiple passes of explicit search for finding the Pareto front and therefore are not sample-efficient, or utilizes a shared policy network for coarse knowledge sharing among policies. To boost the sample efficiency of MORL, we propose  $Q$ -Pensieve, a policy improvement scheme that stores a collection of  $Q$ -snapshots to jointly determine the policy update direction and thereby enables data sharing at the policy level. We show that  $Q$ -Pensieve can be naturally integrated with soft policy iteration with convergence guarantee. To substantiate this concept, we propose the technique of  $Q$  replay buffer, which stores the learned  $Q$ -networks from the past iterations, and arrive at a practical actor-critic implementation. Through extensive experiments and an ablation study, we demonstrate that with much fewer samples, the proposed algorithm can outperform the benchmark MORL methods on a variety of MORL benchmark tasks.

## **1 INTRODUCTION**

Many real-world sequential decision-making problems involve the joint optimization of multiple objectives, while some of them may be in conflict. For example, in robot control, it is expected that the robot can run fast while consuming as little energy as possible; nevertheless, we inevitably need to use more energy to make the robot run fast, regardless of how energy-efficient the robot motion is. Moreover, various other real-world continuous control problems are also multi-objective tasks by nature, such as congestion control in communication networks (Ma et al., 2022) and diversified portfolios (Abdolmaleki et al., 2020). Moreover, the relative importance of these objectives could vary over time (Roijers and Whiteson, 2017). For example, the preference over energy and speed in robot locomotion could change with the energy budget; network service providers need to continuously switch service among various networking applications (e.g., on-demand video streaming versus real-time conferencing), each of which could have preferences over latency and throughput.

To address the above practical challenges, multi-objective reinforcement learning (MORL) serves as one classic and popular formulation for learning optimal control strategies from vector-valued reward signal and achieve favorable trade-off among the objectives. In the MORL framework, the goal is to learn a collection of policies, under which the attained return vectors recover as much of the Pareto front as possible. One popular approach to addressing MORL is to explicitly search for the Pareto front with an aim to maximize the hypervolume associated with the reward vectors, such as evolutionary search (Xu et al., 2020) and search by first-order stationarity (Kyriakis et al., 2022). While being effective, explicit search algorithms are known to be rather sample-inefficient as the data sharing among different passes of explicit search is rather limited. As a result, it is typically difficult to maintain a sufficiently diverse set of optimal policies for different preferences within a reasonable number of training samples. Another way to address MORL is to implicitly search for non-dominated policies through linear scalarization, i.e., convert the vector-valued reward signal to a single scalar with the help of a linear preference and thereafter apply a conventional single-objective RL algorithm for iteratively improving the policies (e.g., (Abels et al., 2019; Yang et al., 2019)).

To enable implicit search for diverse preferences simultaneously, a single network is typically used to express a whole collection of policies. As a result, some level of data sharing among policies of different preferences is done implicitly through the shared network parameters. However, such sharing is clearly not guaranteed to achieve policy improvement for all preferences. Therefore, there remains one critical open research question to be answered: *How to boost the sample efficiency of MORL through better policy-level knowledge sharing?*

To answer this question, we revisit MORL from the perspective of *memory sharing* among the policies learned across different training iterations and propose *Q-Pensieve*, where a “*Pensieve*”, as illustrated in the novel *Harry Potter*, is a magical device used to store pieces of personal memories, which can later be shared with someone else. By drawing an analogy between the memory sharing among humans and the knowledge sharing among policies, we propose to construct a *Q-Pensieve*, which stores snapshots of the *Q*-functions of the policies learned in the past iterations. Upon improving the policy for a specific preference, we expect that these *Q*-snapshots could help jointly determine the policy update direction. In this way, we explicitly enforce knowledge sharing on the policy level and thereby enhance the sample use in learning optimal policies for various preferences. To substantiate this idea, we start by considering *Q-Pensieve* memory sharing in the tabular planning setting and integrate *Q-Pensieve* with the soft policy iteration for entropy-regularized MDPs. Inspired by (Yang et al., 2019), we leverage the envelope operation and propose the *Q-Pensieve* policy iteration for MORL, which we show would preserve the similar convergence guarantee as the standard single-objective soft policy iteration. Based on this result, we propose a practical implementation that consists of two major components: (i) We introduce the technique of *Q* replay buffer. Similar to the standard replay buffer of state transitions, a *Q* replay buffer is meant to achieve sample reuse and improve sample efficiency, but notably at the policy level. Through the use of *Q* replay buffer, we can directly obtain a large collection of *Q* functions, each of which corresponds to a policy in a prior training iteration, without any additional efforts or computation in forming the *Q-Pensieve*. (ii) We convert the *Q-Pensieve* policy iteration into an actor-critic off-policy MORL algorithm by adapting the soft actor critic to the multi-objective setting and using it as the base of our implementation.

The main contributions of this paper can be summarized as:

- We identify the critical sample inefficiency issue in MORL and address this issue by proposing *Q-Pensieve*, which is a policy improvement scheme for enhancing knowledge sharing on the policy level. We then present *Q-Pensieve* policy iteration and establish its convergence property.
- We substantiate the concept of *Q-Pensieve* policy iteration by proposing the technique of *Q* replay buffer and arrive at a practical actor-critic type practical implementation.
- We evaluate the proposed algorithm in various benchmark MORL environments, including Deep Sea Treasure and MuJoCo. Through extensive experiments and an ablation study, we demonstrate the the proposed *Q-Pensieve* can indeed achieve significantly better empirical sample efficiency than the popular benchmark MORL algorithms, in terms of multiple common MORL performance metrics, including hypervolume and utility.

## 2 PRELIMINARIES

**Multi-Objective Markov Decision Processes (MOMDPs).** We consider the formulation of MOMDP defined by the tuple  $(\mathcal{S}, \mathcal{A}, \mathcal{P}, \mathbf{r}, \gamma, \mathcal{D}, \mathfrak{S}_{\lambda}, \Lambda)$ , where  $\mathcal{S}$  denotes the state space,  $\mathcal{A}$  is the action space,  $\mathcal{P} : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \rightarrow [0, 1]$  is the transition kernel of the environment,  $\mathbf{r} : \mathcal{S} \times \mathcal{A} \rightarrow [-r_{\max}, r_{\max}]^d$  is the vector-valued reward function with  $d$  as the number of objectives,  $\gamma \in (0, 1)$  is the discount factor,  $\mathcal{D}$  is the initial state distribution,  $\mathfrak{S}_{\lambda} : \mathbb{R}^d \rightarrow \mathbb{R}$  is the scalarization function (under some preference vector  $\lambda \in \mathbb{R}^d$ ), and  $\Lambda$  denotes the set of all preference vectors. In this paper, we focus on the *linear reward scalarization* setting, i.e.,  $\mathfrak{S}_{\lambda}(\mathbf{r}) = \lambda^\top \mathbf{r}(s, a)$ , as commonly adopted in the MORL literature (Abels et al., 2019; Yang et al., 2019; Kyriakis et al., 2022). Without loss of generality, we let  $\Lambda$  be the unit simplex. If  $d = 1$ , an MOMDP would degenerate to a standard MDP, and we simply use  $r(s, a)$  to denote the scalar reward. At each time step  $t \in \mathbb{N} \cup \{0\}$ , the learner receives the observation  $s_t$ , takes an action  $a_t$ , and receives a reward vector  $\mathbf{r}_t$ . We use  $\pi : \mathcal{S} \rightarrow \Delta(\mathcal{A})$  to denote a stationary randomized policy, where  $\Delta(\mathcal{A})$  denotes the set of all probability distributions over the action space. Let  $\Pi$  be the set of all such policies.

**Single-Objective Entropy-Regularized RL.** In the standard framework of single-objective entropy-regularized RL (Haarnoja et al., 2017; 2018; Geist et al., 2019), the goal is to learn an optimal policy for an entropy-regularized MDP, where an entropy regularization term is augmented to the original reward function. For a policy  $\pi \in \Pi$ , the regularized value functions  $V^\pi : \mathcal{S} \rightarrow \mathbb{R}$  and  $Q^\pi : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$  can be characterized through the regularized Bellman equations as

$$Q^\pi(s, a) = r(s, a) + \gamma \mathbb{E}_{s' \sim \mathcal{P}(\cdot|s, a)} [V^\pi(s')], \quad (1)$$

$$V^\pi(s) = \mathbb{E}_{a \sim \pi(\cdot|s)} [Q^\pi(s, a) - \alpha \log \pi(a|s)], \quad (2)$$

where  $\alpha$  is a temperature parameter that specifies the relative importance of the entropy regularization term. In this setting, the goal is to learn an optimal policy  $\pi^*$  such that  $Q^{\pi^*}(s, a) \geq Q^\pi(s, a)$ , for all  $(s, a)$ , for all  $\pi \in \Pi$ . An optimal policy can be obtained through soft policy iteration, which alternates between soft policy evaluation and soft policy improvement: (i) Soft policy evaluation: For a policy  $\pi$ , the soft  $Q$ -function of  $\pi$  can be obtained by iteratively applying the corresponding soft Bellman backup operator  $\mathcal{T}^\pi$  defined as

$$\mathcal{T}^\pi Q(s, a) = r(s, a) + \gamma \mathbb{E}_{s' \sim \mathcal{P}(\cdot|s, a)} [V(s')], \quad (3)$$

where  $V(s') = \mathbb{E}_{a' \sim \pi(\cdot|s')} [Q(s', a') - \alpha \log \pi(a' | s')]$ . (ii) Soft policy improvement: In each iteration  $k$ , the policy is updated towards an energy-based policy induced by the soft  $Q$ -function, i.e.,

$$\pi_{k+1} = \arg \min_{\pi' \in \tilde{\Pi}} \text{D}_{\text{KL}} \left( \pi'(\cdot | s) \left\| \frac{\exp \left( \frac{1}{\alpha} Q^{\pi_k}(s, \cdot) \right)}{Z^{\pi_k}(s)} \right) \right), \quad (4)$$

where  $\tilde{\Pi}$  is the set of parameterized policies of interest and  $Z^{\pi_k}$  is the normalization term.

**Multi-Objective Entropy-Regularized RL.** We extend the standard single-objective RL with entropy regularization to the multi-objective setting. For each policy  $\pi \in \Pi$ , we define the multi-objective regularized value functions via the following multi-objective version of entropy-regularized Bellman equations as follows:

$$Q^\pi(s, a) = r(s, a) + \gamma \mathbb{E}_{s' \sim \mathcal{P}(\cdot|s, a)} [V^\pi(s')], \quad (5)$$

$$V^\pi(s) = \mathbb{E}_{a \sim \pi(\cdot|s)} [Q^\pi(s, a) - \alpha \log \pi(a|s) \mathbf{1}_d], \quad (6)$$

where  $\mathbf{1}_d$  denotes a  $d$ -dimensional vector of all ones.

In this paper, our goal is to learn a preference-dependent policy  $\pi(\cdot | \cdot; \lambda)$  such that for any preference  $\lambda \in \Lambda$ ,  $\lambda^\top Q^{\pi(\cdot | \cdot; \lambda)}(s, a; \lambda) \geq \lambda^\top Q^{\pi'}(s, a)$ , for all  $(s, a)$ , for all  $\pi' \in \Pi$ . For ease of notation, we let  $V^{\pi(\cdot | \cdot; \lambda)}(s; \lambda) \equiv V^\pi(s; \lambda)$  and  $Q^{\pi(\cdot | \cdot; \lambda)}(s, a; \lambda) \equiv Q^\pi(s, a; \lambda)$  in the sequel.

### 3 ALGORITHMS

In this section, we propose our **Q-Pensieve** learning algorithm for boosting the sample efficiency of multi-objective RL. We first describe the idea of **Q-Pensieve** in the tabular planning setting by introducing **Q-Pensieve** soft policy iteration. We then extend the idea to develop a practical deep reinforcement learning algorithm.

#### 3.1 NAIVE MULTI-OBJECTIVE SOFT POLICY ITERATION

To solve MORL in the entropy-regularized setting, one straightforward approach is to leverage the single-objective soft policy improvement with the help of linear scalarization. That is, in each iteration  $k$ , the policy can be updated by

$$\pi_{k+1}(\cdot, \cdot; \lambda) = \arg \min_{\pi' \in \tilde{\Pi}} \text{D}_{\text{KL}} \left( \pi'(\cdot | s) \left\| \frac{\exp \left( \frac{1}{\alpha} \lambda^\top Q^{\pi_k}(s, \cdot; \lambda) \right)}{Z_\lambda^{\pi_k}(s)} \right) \right). \quad (7)$$

While (7) serves as a reasonable approach, designing a learning algorithm based on the update scheme in (7) could suffer from *sample inefficiency* due to the *lack of policy-level knowledge sharing*: In (7), the policy for each preference  $\lambda$  is updated completely separately based solely on the  $Q$ -function under  $\lambda$ . Moreover, as the update (7) relies on an accurate estimate of the  $Q$ -function, the critic

learning for the policy of each individual preference would typically require at least a moderate number of samples. These issues could be particularly critical for a large preference set in practice. While the use of a conditioned policy network (e.g., (Abels et al., 2019)), a commonly-used network architecture in the MORL literature, could somewhat mitigate this issue, it remains unclear whether the knowledge sharing induced by the conditioned network can indeed achieve policy improvement across various preferences. As a result, a systematic approach is needed for boosting the sample efficiency in MORL.

### 3.2 $Q$ -PENSIEVE SOFT POLICY ITERATION

To boost the sample efficiency of MORL, we propose to enhance the policy-level knowledge sharing by constructing a  $Q$ -Pensieve for *memory sharing across iterations*. Specifically, a  $Q$ -Pensieve is a collection of  $Q$ -snapshots obtained from the past iterations, and it is formed to boost the policy improvement update with respect to the  $Q$ -function in the current iteration as these  $Q$ -snapshots could offer potentially better policy improvement directions under linear scalarization. Moreover, one major computational benefit of  $Q$ -Pensieve is that these  $Q$ -snapshots are obtained without the need for any updates or additional samples from the environment (and hence are for free) as they already exist during training. We substantiate this idea by first introducing the  $Q$ -Pensieve soft policy iteration in the tabular setting (i.e.,  $|S|$  and  $|A|$  are finite) as follows:

**$Q$ -Pensieve Policy Improvement.** In the policy improvement step of the  $k$ -th iteration, for each specific  $\lambda$ , we update the policy as

$$\pi_{k+1}(\cdot|\cdot; \lambda) = \arg \min_{\pi' \in \Pi} D_{\text{KL}} \left( \pi'(\cdot|s; \lambda) \left\| \frac{\exp \left( \sup_{\lambda' \in W_k(\lambda), Q' \in \mathcal{Q}_k} \frac{1}{\alpha} \lambda^\top Q'(s, \cdot; \lambda') \right)}{Z_{\mathcal{Q}_k}(s)} \right. \right), \quad (8)$$

where  $Z_{\mathcal{Q}_k}$  is again the normalization term,  $W_k(\lambda) \subset \Lambda$  is a set of preference vectors, and  $\mathcal{Q}_k$  is a set of  $Q$ -snapshots. The two sets  $W_k(\lambda)$  and  $\mathcal{Q}_k$  are to be selected as follows:

- For  $W_k(\lambda)$ , the only requirement is that  $\lambda \in W_k(\lambda)$ , for all  $k$ . The preference sets can be different in different iterations.
- Similarly, for  $\mathcal{Q}_k$ , the only requirement is that  $Q^{\pi_k} \in \mathcal{Q}_k$ , for all  $k$ . The set of  $Q$ -snapshots can also be different in different iterations. Hence, the choice of  $\mathcal{Q}_k$  is rather flexible.

When choosing  $W_k(\lambda) = \{\lambda\}$  and  $\mathcal{Q}_k = \{Q^{\pi_k}\}$ , one would recover the update in (7).

**Policy Evaluation.** In the policy evaluation step, we evaluate the policy that corresponds to each preference  $\lambda$  by iteratively applying the multi-objective softmax Bellman backup operator  $\mathcal{T}_{\text{MO}}^\pi$  as

$$(\mathcal{T}_{\text{MO}}^\pi Q)(s, a; \lambda) = r(s, a) + \gamma \mathbb{E}_{s' \sim \mathcal{P}(\cdot|s, a), a' \sim \pi(\cdot|s'; \lambda)} [Q(s', a'; \lambda) - \alpha \log \pi(a'|s'; \lambda) \mathbf{1}_d]. \quad (9)$$

**Remark 1** The  $Q$ -Pensieve update in (8) is inspired by the envelope Q-learning (EQL) technique (Yang et al., 2019), where in each iteration  $k$ , the Q-learning update takes into account the envelope formed by the  $Q$ -functions of the current policy  $\pi_k$  for different preferences. The fundamental difference between  $Q$ -Pensieve and EQL is that  $Q$ -Pensieve further achieves *memory sharing across training iterations* through the use of  $Q$ -snapshots from the past iterations, and EQL focuses mainly on the use of the  $Q$ -function of the current iteration.

**Convergence of  $Q$ -Pensieve Soft Policy Iteration.** Another nice feature of the  $Q$ -Pensieve policy improvement step is that it preserves the similar convergence result as the standard single-objective soft policy iteration, as stated below. The proof of Theorem 3.1 is provided in Appendix A.

**Theorem 3.1** *Under the  $Q$ -Pensieve soft policy iteration given by (8) and (9), the sequence of preference-dependent policies  $\{\pi_k\}$  converges to a policy  $\pi^*$  such that  $\lambda^\top Q^{\pi^*}(s, a; \lambda) \geq \lambda^\top Q^\pi(s, a)$  for all  $\pi \in \Pi$ , for all  $(s, a) \in \mathcal{S} \times \mathcal{A}$  and for all  $\lambda \in \Lambda$ .*

### 3.3 PRACTICAL IMPLEMENTATION OF $Q$ -PENSIEVE

In this section, we present the implementation of proposed  $Q$ -Pensieve algorithm for learning policies with function approximation for the general state and action spaces.

**Q Replay Buffer.** Based on (8), we know that the policy update of *Q*-Pensieve would involve both the current *Q*-function and the *Q*-snapshots from the past iterations. To implement this, we introduce *Q* replay buffer, which could store multiple *Q*-networks in a predetermined manner (e.g., first-in first-out). Notably, unlike the conventional experience replay buffer (Mnih et al., 2013) of state transitions, *Q* replay buffer stores the learned *Q*-networks in past iterations as candidates for forming the *Q*-Pensieve. On the other hand, while each *Q*-network would require a moderate amount of memory usage, we found that in practice a rather small *Q* replay buffer is already effective enough for boosting the sample efficiency. We further illustrate this observation through the experimental results in Section 4.

Next, we convert the *Q*-Pensieve soft policy iteration into an actor-critic off-policy MORL algorithm. Specifically, we adapt the idea of soft actor critic to *Q*-Pensieve by minimizing the residual of the multi-objective soft *Q*-function: Let  $\theta$  and  $\phi$  be the parameters of the policy network and the critic network, respectively. Then, the critic network is updated by minimizing the following loss

$$\mathcal{L}_Q(\phi; \lambda) = \mathbb{E}_{(s,a) \sim \mu} \left[ \lambda^\top \left( Q_\phi(s, a; \lambda) - (r(s, a) + \gamma \mathbb{E}_{s' \sim \mathcal{P}(\cdot|s,a)} [V_{\bar{\phi}}(s')]) \right)^2 \right], \quad (10)$$

where  $\bar{\phi}$  is the parameter of the target network and  $\mu$  is the sampling distribution of the state-action pairs (e.g., a distribution induced by a replay buffer of state transitions). On the other hand, based on (8), the policy network is updated by minimizing the following objective

$$\mathcal{L}_\pi(\theta; \lambda) = \mathbb{E}_{s \sim \mu} \left[ \mathbb{E}_{a \sim \pi_\theta} \left[ \sup_{\lambda' \in W(\lambda), Q' \in \mathcal{Q}} \{ \alpha \log(\pi_\theta(a | s; \lambda)) - \lambda^\top Q'(s, a; \lambda') \} \right] \right]. \quad (11)$$

The overall architecture of *Q*-Pensieve is provided in Figure 1. The pseudo code of the *Q*-Pensieve algorithm is described in Algorithm 1 in Appendix.




Figure 1: The architecture of *Q*-Pensieve.

## 4 EXPERIMENTS

In this section, we demonstrate the effectiveness of *Q*-Pensieve on various benchmark RL tasks and discuss how *Q*-Pensieve boosts the sample efficiency through an extensive ablation study.

### 4.1 EXPERIMENTAL CONFIGURATION

**Popular Benchmark Methods.** We compare the proposed algorithm against various popular benchmark methods, including the Conditioned Network with Diverse Experience Replay (CN-DER) in (Abels et al., 2019), the Prediction-Guided Multi-Objective RL (PGMORL) in (Xu et al., 2020), the Pareto Following Algorithm (PFA) in (Parisi et al., 2014), and SAC (Haarnoja et al., 2018). For

CN-DER, as the original CN-DER is built on deep Q-networks (DQN) for discrete actions, we modify the source code of Abels et al. (2019) for continuous control by implementing CN-DER on top of DDPG. Moreover, we follow the same DER technique, which uses a diverse replay buffer and gives priority according to how much the samples increase the overall diversity of the buffer. For PGMORL and PFA, we use the open-source implementation of (Xu et al., 2020) for the experiments. As these explicit search methods typically require more samples before reaching a comparable performance level, we evaluate the performance PGMORL and PFA under both 1 times and  $\beta$  times ( $\beta > 1$ ) of the number of samples used by  $Q$ -Pensieve to demonstrate the sample efficiency of  $Q$ -Pensieve. For SAC, as the MORL problem reduces a single-objective one under a fixed preference, we train multiple models using single-objective SAC (one model for each fixed preference) as a performance reference for other MORL methods.

**Performance Metrics.** In the evaluation, we consider the following three commonly-used performance metrics for MORL:

- **HyperVolume (HV):** Let  $\mathcal{R}$  be a set of return vectors attained and  $\mathbf{r}_0 \in \mathbb{R}^d$  be a reference point. Then, we define the HyperVolume as  $HV := \int_{H(\mathcal{R})} \mathbb{I}\{z \in H(\mathcal{R})\} dz$ , where  $H(\mathcal{R}) := \{z \in \mathbb{R}^d : \exists \mathbf{r} \in \mathcal{R}, \mathbf{r}_0 \prec z \prec \mathbf{r}\}$  and  $\mathbb{I}$  is the indicator function.
- **Utility (UT):** To further evaluate the performance under linear scalarization, we define the utility metric as  $UT := \mathbb{E}_{\lambda} \left[ \sum_{t=0}^T \lambda^\top \mathbf{r}_t \right]$ , where the preference  $\lambda$  is sampled uniformly from  $\Lambda$ .
- **Episodic Dominance (ED):** To compare the performance of a pair of algorithms, we define Episodic Dominance as  $ED_{1,2} := \mathbb{E}_{\lambda} [\mathbb{I}\{\sum_{t=0}^{T_1} \lambda^\top \mathbf{r}_t^1 > \sum_{t=0}^{T_2} \lambda^\top \mathbf{r}_t^2\}]$ , where  $\mathbf{r}_t^1, \mathbf{r}_t^2$  are the return vectors, and  $T_1, T_2$  are the episode lengths of algorithm 1 and 2, respectively. ED serves as a useful metric for pairwise comparison in those problems where the return vectors under different preferences can differ by a lot in magnitude (in this case, HV and UT could be dominated by the return vectors of a few preferences).

**Evaluation Domains.** We evaluate the algorithms in the following domains: (i) Continuous Deep Sea Treasure (DST): a two-objective continuous control task modified from the original DST environment. (ii) Multi-Objective Continuous LunarLander: a four-objective task modified from the classic control task in the OpenAI gym. (iii) Multi-Objective MuJoCo: modified benchmark locomotion tasks with either two or three objectives.

**Configuration of  $Q$ -Pensieve.** For  $Q$ -Pensieve, at each policy update, we set the size of the preference set  $W_k(\lambda)$  to be 5 (including  $\lambda$  and another four preferences drawn randomly) and set the size of the  $Q$  replay buffer to be 4, unless stated otherwise.

## 4.2 EXPERIMENTAL RESULTS

**Does  $Q$ -Pensieve achieve better sample efficiency than the MORL benchmark methods?** Table 1 shows the performance of  $Q$ -Pensieve and the benchmark methods in terms of the three metrics. For each algorithm, we report the mean and the standard deviation over five random seeds. We can observe that  $Q$ -Pensieve consistently enjoys higher HV, UT, and ED in almost all the domains. More importantly,  $Q$ -Pensieve indeed exhibits superior sample efficiency as it still outperforms the explicit search methods (i.e., PFA and PGMORL) even if these methods are given 10 times of the number of samples used by  $Q$ -Pensieve. Moreover, we can observe that the explicit search methods (i.e., PFA and PGMORL) often have larger HV than the implicit search method (such as CN-DER), while implicit search methods tend to have larger UT. This manifests the design principles and the characteristics of the two families of approaches, where explicit search is designed mainly for achieving large HV and implicit search typically aims for larger scalarized return.

**How much improvement in sample efficiency can  $Q$ -Pensieve achieve compared to training multiple single-objective SAC models separately?** To answer this question, we conduct experiments on 2-objective MuJoCo tasks and consider a whole range of 19 preference vectors  $([0.05, 0.95], [0.1, 0.9], [0.15, 0.85], \dots, [0.95, 0.05])$ . We train 19 models by using single-objective SAC, one model for each individual preference. Each model is trained for 1.5M steps (and hence the total number of steps under SAC is 28.5M steps). By contrast,  $Q$ -Pensieve only uses 1.5M steps in total in learning policies for all the preferences. Figure 2 shows the return vectors attained by  $Q$ -Pensieve and the collection of 19 SAC models.  $Q$ -Pensieve can achieve comparable or better

Environments	Metrics	PFA (1.5M steps)	PFA (1.5 × $\beta$ M steps)	PGMORL (1.5M steps)	PGMORL (1.5 × $\beta$ M steps)	CN-DER (1.5M steps)	$Q$ -Pensieve (1.5M steps)
DST	HV( $\times 10^2$ )	7.43 $\pm$ 3.68	8.67 $\pm$ 1.49	8.10 $\pm$ 1.57	8.13 $\pm$ 1.61	5.36 $\pm$ 4.71	<b>10.21<math>\pm</math>1.40</b>
	UT	-9.27 $\pm$ 6.03	-6.86 $\pm$ 6.06	4.90 $\pm$ 0.44	5.02 $\pm$ 0.35	-5.10 $\pm$ 15.73	<b>7.31<math>\pm</math>0.91</b>
	ED	0.13 $\pm$ 0.11	0.10 $\pm$ 0.08	0.25 $\pm$ 0.18	0.28 $\pm$ 0.18	0.21 $\pm$ 0.17	<b>0.54<math>\pm</math>0.11</b>
LunarLander	HV( $\times 10^8$ )	-	-	0.32 $\pm$ 0.11	0.38 $\pm$ 0.11	1.50 $\pm$ 0.60	<b>2.10<math>\pm</math>0.10</b>
	UT( $\times 10$ )	-	-	-0.26 $\pm$ 0.27	1.10 $\pm$ 0.50	3.60 $\pm$ 2.90	<b>5.10<math>\pm</math>0.30</b>
	ED	-	-	0.02 $\pm$ 0.01	0.04 $\pm$ 0.04	0.21 $\pm$ 0.12	<b>0.49<math>\pm</math>0.05</b>
HalfCheetah	HV( $\times 10^7$ )	0.73 $\pm$ 0.19	1.31 $\pm$ 0.26	0.53 $\pm$ 0.17	0.28 $\pm$ 0.29	2.08 $\pm$ 0.54	<b>3.82<math>\pm</math>0.27</b>
	UT( $\times 10^3$ )	0.31 $\pm$ 0.20	1.02 $\pm$ 0.40	-0.28 $\pm$ 0.94	0.09 $\pm$ 0.17	5.09 $\pm$ 3.57	<b>5.61<math>\pm</math>0.31</b>
	ED	0.08 $\pm$ 0.10	0.10 $\pm$ 0.06	0.01 $\pm$ 0.00	0.11 $\pm$ 0.05	0.02 $\pm$ 0.01	<b>0.54<math>\pm</math>0.08</b>
Hopper	HV( $\times 10^6$ )	0.49 $\pm$ 0.46	1.01 $\pm$ 0.62	0.63 $\pm$ 0.48	1.31 $\pm$ 0.48	0.56 $\pm$ 0.16	<b>1.33<math>\pm</math>0.20</b>
	UT( $\times 10^2$ )	2.89 $\pm$ 1.93	3.50 $\pm$ 1.85	1.94 $\pm$ 2.46	3.70 $\pm$ 1.78	1.42 $\pm$ 1.00	<b>4.08<math>\pm</math>1.10</b>
	ED	0.31 $\pm$ 0.17	0.41 $\pm$ 0.10	0.31 $\pm$ 0.25	0.31 $\pm$ 0.11	0.04 $\pm$ 0.03	<b>0.43<math>\pm</math>0.09</b>
Hopper3d	HV( $\times 10^9$ )	-	-	0.29 $\pm$ 0.37	0.91 $\pm$ 1.39	3.70 $\pm$ 0.81	<b>9.56<math>\pm</math>0.60</b>
	UT( $\times 10^3$ )	-	-	0.19 $\pm$ 0.16	0.31 $\pm$ 0.26	0.72 $\pm$ 0.16	<b>1.39<math>\pm</math>0.15</b>
	ED	-	-	0.02 $\pm$ 0.03	0.03 $\pm$ 0.03	0.07 $\pm$ 0.03	<b>0.55<math>\pm</math>0.08</b>
Ant	HV( $\times 10^6$ )	0.17 $\pm$ 0.05	0.77 $\pm$ 0.53	0.14 $\pm$ 0.03	0.13 $\pm$ 0.04	5.03 $\pm$ 3.60	<b>10.01<math>\pm</math>1.86</b>
	UT( $\times 10^2$ )	-0.06 $\pm$ 0.01	0.14 $\pm$ 0.14	-0.21 $\pm$ 0.15	-0.18 $\pm$ 0.38	3.68 $\pm$ 2.34	<b>14.04<math>\pm</math>3.03</b>
	ED	0.22 $\pm$ 0.03	0.22 $\pm$ 0.02	0.21 $\pm$ 0.02	0.21 $\pm$ 0.03	0.21 $\pm$ 0.08	<b>0.60<math>\pm</math>0.07</b>
Ant3d	HV( $\times 10^8$ )	-	-	0.41 $\pm$ 0.48	0.68 $\pm$ 0.62	13.00 $\pm$ 4.11	<b>21.87<math>\pm</math>1.07</b>
	UT( $\times 10^3$ )	-	-	0.18 $\pm$ 0.05	0.25 $\pm$ 0.05	0.49 $\pm$ 0.23	<b>1.14<math>\pm</math>0.22</b>
	ED	-	-	0.02 $\pm$ 0.02	0.03 $\pm$ 0.03	0.28 $\pm$ 0.14	<b>0.56<math>\pm</math>0.07</b>
Walker2d	HV( $\times 10^6$ )	0.52 $\pm$ 0.20	1.05 $\pm$ 0.44	0.83 $\pm$ 0.42	<b>1.28<math>\pm</math>0.66</b>	0.42 $\pm$ 0.09	1.12 $\pm$ 0.36
	UT( $\times 10^2$ )	0.23 $\pm$ 0.13	0.95 $\pm$ 0.55	0.38 $\pm$ 0.24	1.20 $\pm$ 0.67	3.17 $\pm$ 0.53	<b>6.37<math>\pm</math>1.42</b>
	ED	0.32 $\pm$ 0.06	0.37 $\pm$ 0.09	0.30 $\pm$ 0.10	0.34 $\pm$ 0.12	0.21 $\pm$ 0.11	<b>0.48<math>\pm</math>0.10</b>

Table 1: Comparison of  $Q$ -Pensieve and other benchmark algorithms in terms of the three metrics across eight domains. We report the mean and standard deviation over five random seeds. The ED is calculated through comparing each algorithm to a multi-objective version of SAC (equivalent to  $Q$ -Pensieve with the size of the preference set equal to 1 and without  $Q$  replay buffer). We set  $\beta = 10$  for HalfCheetah, Ant, Ant3d, and Hopper3d, set  $\beta = 5$  for LunarLander, and set  $\beta = 3$  for DST, Hopper, and Walker2d.

returns than the collection of SAC models with only 1/19 of the samples. This further demonstrate the sample efficiency of  $Q$ -Pensieve.

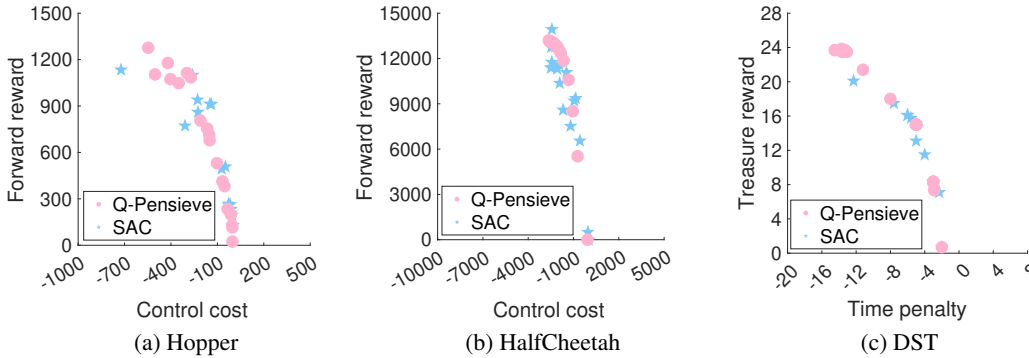


Figure 2: Return vectors attained by  $Q$ -Pensieve and the collection of single-objective SAC models under 19 preferences.

**Why can  $Q$ -Pensieve outperform single-objective SAC in some cases?** From Figures 2(a) and (c), we see that  $Q$ -Pensieve can attain some return vectors that are strictly better than those of the single-objective SAC models. The reasons behind this phenomenon are mainly two-fold: (i) Under single-objective SAC, despite that we train one model for each individual preference, it could still occur that single-objective SAC gets stuck at a sub-optimal policy under some preferences. (ii) By contrast,  $Q$ -Pensieve has a better chance of escaping from these sub-optimal policies with the help of the  $Q$ -snapshots in the  $Q$  replay buffer.

To verify the above argument, we design a hybrid SAC algorithm as follows: (a) For the first  $10^5$  time steps, this algorithm simply follows the single-objective SAC. (b) At time step  $10^5$ , it switches

to the update rule of  $Q$ -Pensieve based on the  $Q$ -snapshots stored in the  $Q$  replay buffer of another model trained under  $Q$ -Pensieve algorithm in parallel. Figure 3 shows the performance of this hybrid algorithm in DST and HalfCheetah. Clearly, the  $Q$ -Pensieve update could help the SAC model escape from the sub-optimal policies, under various preferences.




Figure 3: Comparison of standard single-objective SAC and the hybrid SAC assisted by another  $Q$ -Pensieve model trained in parallel.

**An ablation study on  $Q$  replay buffer.** To verify the effectiveness of the technique of  $Q$  replay buffer, we compare the performance of  $Q$ -Pensieve with buffer size equal to 4 and that without using  $Q$  replay buffer (termed “Vanilla” in Figures 4 and 5). Figure 4 and 5 show the attained return vectors and HV of both methods. We can see that  $Q$  replay buffer indeed leads to a better policy improvement behavior, in terms of both HV and the scalarized returns. However, these figures may sometimes oscillate a lot in the end period. It is because our algorithm finds solutions from another  $Q$ -vector, and their inner product of  $Q$  and preference may be quite close. We can check the points are in the same contour.

## 5 RELATED WORK

The multi-objective RL problems have been extensively studied from two major perspectives:

**Explicit search.** A plethora of prior works on MORL updates a policy or a set of policies by explicitly searching for the Pareto front of the reward space. To learn policies under time-varying preferences, (Natarajan and Tadepalli, 2005) presented to store a set of policies, which are to be used in searching for a proper policy for a new preference without learning from scratch. (Lizotte et al., 2012) leveraged linear value function approximation to search for optimal policies. (Van Moffaert and Nowé, 2014) proposed Pareto Q-learning, which stores the immediate rewards and the non-dominated future return vectors separately and leverage the Pareto dominance for selecting the actions in Q-learning. (Parisi et al., 2014) presented a policy gradient approach to search for non-dominated policies. (Mossalam et al., 2016) solves MORL via scalarized Q-learning along with the concept of prioritizing the




Figure 4: Return vectors attained under preference  $\lambda = [0.5, 0.5]$  at different training stages. (We also plot return vectors under others preference in Figure 7 and Figure 8 in Appendix.) A number  $x$  on the red or blue marker indicates that the model is obtained at  $100 \cdot x$  thousand steps.



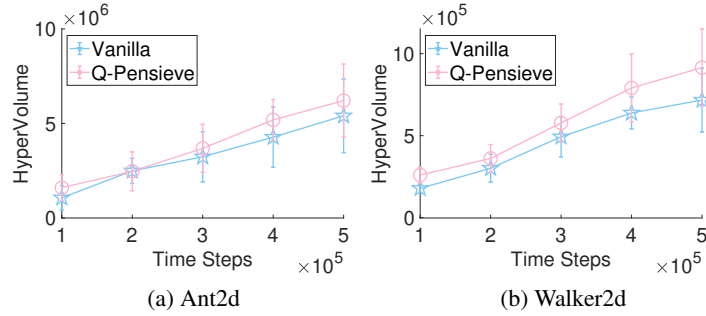


Figure 5: A comparison in HV between  $Q$ -Pensieve with buffer size equal to 4 and that without using  $Q$  replay buffer at different training stages.

corner weights for selecting the preference of the scalarized problem. (Xu et al., 2020) proposed an evolutionary approach to search for the Pareto set of policies, with the help of a prediction model for determining the search direction. (Kyriakis et al., 2022) presented a policy gradient method by approximating the Pareto front via a first-order necessary condition. However, the above explicit search algorithms are known to be rather sample-inefficient as the knowledge sharing among different passes of search is limited.

**Implicit search.** Another class of algorithms are designed to improve policies for multiple preferences through implicit search. For example, (Abels et al., 2019) presents Conditioned Network, which extends the standard single-objective DQN to learning preference-dependent multi-objective  $Q$ -functions. To achieve scale-invariant MORL, (Abdolmaleki et al., 2020) proposed to first learn the  $Q$ -functions for different objectives and encode the preference through constraints. While there is implicit Recently, (Yang et al., 2019) proposes envelope  $Q$ -learning to encourage knowledge sharing among the  $Q$  functions of different the current multi-objective  $Q$ -values that any policy can benefit from other preferences’ experiences, that make training more efficiently, and (Zhou et al., 2020) proposed model-based envelope value iteration base on envelope  $Q$ -learning, it provides an efficient way to get optimal multi-objective  $Q$  function. Despite that our method is inspired by (Yang et al., 2019), the main difference between our work and theirs is that we boost the sample efficiency of MORL via explicit memory sharing among policies learned during training.

## 6 CONCLUSION

This paper proposes  $Q$ -Pensieve, which significantly enhances the policy-level data sharing through in order to boost the sample efficiency of MORL problems. We substantiate the idea by presenting  $Q$ -Pensieve soft policy iteration in the tabular setting and show that it preserves the global convergence property. Then, to implement the  $Q$ -Pensieve policy improvement step, we introduce the  $Q$  replay buffer technique, which offers a simple yet effective way to maintain  $Q$ -snapshot. Our experiments demonstrate that  $Q$ -Pensieve is a promising approach in that it can outperform the state-of-the-art MORL methods with much fewer samples in a variety of MORL benchmark tasks.

## REFERENCES

- Yiqing Ma, Han Tian, Xudong Liao, Junxue Zhang, Weiyan Wang, Kai Chen, and Xin Jin. Multi-objective congestion control. In *Proceedings of the Seventeenth European Conference on Computer Systems*, pages 218–235, 2022.
- Abbas Abdolmaleki, Sandy Huang, Leonard Hasenclever, Michael Neunert, Francis Song, Martina Zambelli, Murilo Martins, Nicolas Heess, Raia Hadsell, and Martin Riedmiller. A distributional view on multi-objective policy optimization. In Hal Daumé III and Aarti Singh, editors, *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pages 11–22. PMLR, 13–18 Jul 2020. URL <https://proceedings.mlr.press/v119/abdolmaleki20a.html>.

- Diederik M Roijers and Shimon Whiteson. Multi-objective decision making. *Synthesis Lectures on Artificial Intelligence and Machine Learning*, 11(1):1–129, 2017.
- Jie Xu, Yunsheng Tian, Pingchuan Ma, Daniela Rus, Shinjiro Sueda, and Wojciech Matusik. Prediction-guided multi-objective reinforcement learning for continuous robot control. In Hal Daumé III and Aarti Singh, editors, *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pages 10607–10616. PMLR, 13–18 Jul 2020. URL <https://proceedings.mlr.press/v119/xu20h.html>.
- Panagiotis Kyriakis, Jyotirmoy Deshmukh, and Paul Bogdan. Pareto policy adaptation. In *International Conference on Learning Representations*, 2022.
- Axel Abels, Diederik Roijers, Tom Lenaerts, Ann Nowé, and Denis Steckelmacher. Dynamic weights in multi-objective deep reinforcement learning. In *International Conference on Machine Learning*, pages 11–20. PMLR, 2019.
- Runzhe Yang, Xingyuan Sun, and Karthik Narasimhan. A generalized algorithm for multi-objective reinforcement learning and policy adaptation. *Advances in Neural Information Processing Systems*, 32, 2019.
- Tuomas Haarnoja, Haoran Tang, Pieter Abbeel, and Sergey Levine. Reinforcement learning with deep energy-based policies. In *International Conference on Machine Learning*, pages 1352–1361, 2017.
- Tuomas Haarnoja, Aurick Zhou, Pieter Abbeel, and Sergey Levine. Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. In *International conference on machine learning*, pages 1861–1870. PMLR, 2018.
- Matthieu Geist, Bruno Scherrer, and Olivier Pietquin. A theory of regularized markov decision processes. In *International Conference on Machine Learning*, pages 2160–2169, 2019.
- Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Alex Graves, Ioannis Antonoglou, Daan Wierstra, and Martin Riedmiller. Playing atari with deep reinforcement learning. *arXiv:1312.5602*, 2013.
- Simone Parisi, Matteo Pirotta, Nicola Smacchia, Luca Bascetta, and Marcello Restelli. Policy gradient approaches for multi-objective sequential decision making. In *International Joint Conference on Neural Networks (IJCNN)*, pages 2323–2330, 2014.
- Sriraam Natarajan and Prasad Tadepalli. Dynamic preferences in multi-criteria reinforcement learning. In *International Conference on Machine Learning*, pages 601–608, 2005.
- Daniel J Lizotte, Michael Bowling, and Susan A Murphy. Linear fitted-q iteration with multiple reward functions. *Journal of Machine Learning Research*, 13(1):3253–3295, 2012.
- Kristof Van Moffaert and Ann Nowé. Multi-objective reinforcement learning using sets of pareto dominating policies. *Journal of Machine Learning Research*, 15(1):3483–3512, 2014.
- Hossam Mossalam, Yannis M. Assael, Diederik M. Roijers, and Shimon Whiteson. Multi-objective deep reinforcement learning, 2016. URL <https://arxiv.org/abs/1610.02707>.
- Dongruo Zhou, Jiahao Chen, and Quanquan Gu. Provable multi-objective reinforcement learning with generative models. *arXiv preprint arXiv:2011.10134*, 2020.
- Richard S Sutton and Andrew G Barto. *Reinforcement learning: An introduction*. MIT press, 2018.

## APPENDIX

## A PROOF OF THEOREM 3.1

Before proving Theorem 3.1, we first present two supporting lemmas as follows. To begin with, we establish the policy improvement property of the  $Q$ -Pensieve update. Recall that the  $Q$ -Pensieve policy update is that for each preference  $\lambda \in \Lambda$ ,

$$\pi_{k+1}(\cdot|s; \lambda) = \arg \min_{\pi' \in \Pi} \underbrace{D_{\text{KL}}\left(\pi'(\cdot|s; \lambda) \parallel \frac{\exp\left(\sup_{\lambda' \in W_k(\lambda), Q' \in \mathcal{Q}_k} \frac{1}{\alpha} \lambda^\top Q'(s, \cdot; \lambda')\right)}{Z_{\mathcal{Q}}(s)}\right)}_{=: L(\pi'; \lambda)}. \quad (12)$$

**Lemma 1 ( $Q$ -Pensieve Policy Improvement)** *Under the  $Q$ -Pensieve policy improvement update, we have  $\lambda^\top Q^{\pi_k}(s, a; \lambda) \leq \lambda^\top Q^{\pi_{k+1}}(s, a; \lambda)$ , for all state-action pairs  $(s, a) \in \mathcal{S} \times \mathcal{A}$ , for all preference vectors  $\lambda \in \Lambda$ , and for all iteration  $k \in \mathbb{N} \cup \{0\}$ .*

**Proof (Lemma 1)** By the update rule in (12), we know that  $\pi_{k+1}$  is a minimizer of  $L(\pi'; \lambda)$  and hence  $L(\pi_{k+1}; \lambda) \leq L(\pi_k; \lambda)$ . This implies that for each state  $s \in \mathcal{S}$ , we have

$$\begin{aligned} & \mathbb{E}_{a \sim \pi_{k+1}(\cdot|s)} \left[ \lambda^\top \mathbf{1}_d \cdot \log \pi_{k+1}(a|s; \lambda) - \frac{1}{\alpha} \sup_{\lambda' \in W_k(\lambda), Q' \in \mathcal{Q}_k} \lambda^\top Q'(s, a; \lambda') + \log Z_{\mathcal{Q}_k}(s) \right] \\ & \leq \mathbb{E}_{a \sim \pi_k(\cdot|s)} \left[ \lambda^\top \mathbf{1}_d \cdot \log \pi_k(a|s; \lambda) - \frac{1}{\alpha} \sup_{\lambda' \in W_k(\lambda), Q' \in \mathcal{Q}_k} \lambda^\top Q'(s, a; \lambda') + \log Z_{\mathcal{Q}_k}(s) \right]. \end{aligned} \quad (13)$$

Since  $Z_{\mathcal{Q}_k}$  only depends on the state, the inequality (13) reduces to

$$\begin{aligned} & \mathbb{E}_{a \sim \pi_{k+1}(\cdot|s)} \left[ \lambda^\top \mathbf{1}_d \cdot \log \pi_{k+1}(a|s; \lambda) - \frac{1}{\alpha} \sup_{\lambda' \in W_k(\lambda), Q' \in \mathcal{Q}_k} \lambda^\top Q'(s, a; \lambda') \right] \\ & \leq \mathbb{E}_{a \sim \pi_k(\cdot|s)} \left[ \lambda^\top \mathbf{1}_d \cdot \log \pi_k(a|s; \lambda) - \frac{1}{\alpha} \sup_{\lambda' \in W_k(\lambda), Q' \in \mathcal{Q}_k} \lambda^\top Q'(s, a; \lambda') \right]. \end{aligned} \quad (14)$$

Next, we proceed to consider the multi-objective soft Bellman equation as follows:

$$\lambda^\top Q^{\pi_k}(s_0, a_0; \lambda) - \lambda^\top \mathbf{r}(s_0, a_0) \quad (15)$$

$$= \gamma \lambda^\top \mathbb{E}_{s_1 \sim P(\cdot|s_0, a_0), a_1 \sim \pi_k(\cdot|s_1; \lambda)} [Q^{\pi_k}(s_1, a_1; \lambda) - \alpha \cdot \lambda^\top \mathbf{1}_d \log(\pi_k(a_1|s_1; \lambda))] \quad (16)$$

$$\leq \gamma \mathbb{E}_{s_1 \sim P(\cdot|s_0, a_0), a_1 \sim \pi_k(\cdot|s_1; \lambda)} \left[ \sup_{\lambda' \in W_k(\lambda), Q' \in \mathcal{Q}_k} \lambda^\top Q'(s_1, a_1; \lambda') - \alpha \cdot \lambda^\top \mathbf{1}_d \log(\pi_k(a_1|s_1; \lambda)) \right] \quad (17)$$

$$\leq \gamma \mathbb{E}_{s_1 \sim P(\cdot|s_0, a_0), a_1 \sim \pi_{k+1}(\cdot|s_1; \lambda)} \left[ \sup_{\lambda' \in W_k(\lambda), Q' \in \mathcal{Q}_k} \lambda^\top Q'(s_1, a_1; \lambda') - \alpha \cdot \lambda^\top \mathbf{1}_d \log(\pi_{k+1}(a_1|s_1; \lambda)) \right] \quad (18)$$

$$\begin{aligned} & \leq -\gamma \mathbb{E}_{s_1 \sim P(\cdot|s_0, a_0), a_1 \sim \pi_{k+1}(\cdot|s_1; \lambda)} \left[ \alpha \cdot \lambda^\top \mathbf{1}_d \log(\pi_{k+1}(a_1|s_1; \lambda)) \right] \\ & \quad + \gamma \mathbb{E}_{s_1 \sim P(\cdot|s_0, a_0), a_1 \sim \pi_{k+1}(\cdot|s_1; \lambda)} \left[ \lambda^\top \mathbf{r}(s_1, a_1) - \gamma \mathbb{E}_{s_2 \sim P(\cdot|s_1, a_1), a_2 \sim \pi_k(\cdot|s_2)} [\alpha \lambda^\top \mathbf{1}_d \log(\pi_k(a_2|s_2; \lambda))] \right. \\ & \quad \left. + \gamma \sup_{\lambda' \in W_k(\lambda), Q' \in \mathcal{Q}_k} \mathbb{E}_{s_2 \sim P(\cdot|s_1, a_1), a_2 \sim \pi_k(\cdot|s_2)} [\lambda^\top Q'(s_2, a_2; \lambda')] \right] \end{aligned} \quad (19)$$

$$\begin{aligned} & \leq -\gamma \mathbb{E}_{s_1 \sim P(\cdot|s_0, a_0), a_1 \sim \pi_{k+1}(\cdot|s_1; \lambda)} \left[ \alpha \cdot \lambda^\top \mathbf{1}_d \log(\pi_{k+1}(a_1|s_1; \lambda)) \right] \\ & \quad + \gamma \mathbb{E}_{s_1 \sim P(\cdot|s_0, a_0), a_1 \sim \pi_{k+1}(\cdot|s_1; \lambda)} \left[ \lambda^\top \mathbf{r}(s_1, a_1) - \gamma \mathbb{E}_{s_2 \sim P(\cdot|s_1, a_1), a_2 \sim \pi_{k+1}(\cdot|s_2)} [\alpha \lambda^\top \mathbf{1}_d \log(\pi_{k+1}(a_2|s_2; \lambda))] \right. \\ & \quad \left. + \gamma \sup_{\lambda' \in W_k(\lambda), Q' \in \mathcal{Q}_k} \mathbb{E}_{s_2 \sim P(\cdot|s_1, a_1), a_2 \sim \pi_{k+1}(\cdot|s_2)} [\lambda^\top Q'(s_2, a_2; \lambda')] \right] \end{aligned} \quad (20)$$

(21)

$$\begin{aligned} &\leq -\gamma \mathbb{E}_{s_1 \sim P(\cdot|s_0, a_0), a_1 \sim \pi_{k+1}(\cdot|s_1; \lambda)} \left[ \alpha \cdot \lambda^\top \mathbf{1}_d \log(\pi_{k+1}(a_1|s_1; \lambda)) \right] \\ &\quad + \mathbb{E}_{P, \pi_{k+1}} \left[ \sum_{t \geq 1} \gamma^t \mathbb{E} \left[ \lambda^\top \mathbf{r}(s_t, a_t) \right. \right. \end{aligned} \quad (22)$$

$$\left. \left. - \gamma \mathbb{E}_{s_{t+1} \sim P(\cdot|s_t, a_t), a_{t+1} \sim \pi_{k+1}(\cdot|s_{t+1})} \left[ \alpha \lambda^\top \mathbf{1}_d \log(\pi_{k+1}(a_{t+1}|s_{t+1}; \lambda)) \right] | s_t, a_t \right] \right] \quad (23)$$

$$= \lambda^\top \mathbf{Q}^{\pi_{k+1}}(s_0, a_0; \lambda) - \lambda^\top \mathbf{r}(s_0, a_0), \quad (24)$$

where (16) follows from the multi-objective soft Bellman equation, (17) holds by the sup operation and the fact that  $\mathbf{Q}^{\pi_k} \in \mathcal{Q}_k$ , (18) follows from (14), (19) holds by applying the multi-objective soft Bellman equation to  $\mathbf{Q}'(s_1, a_1; \lambda)$ , (20) again follows from the inequality in (14), (23) is obtained by unrolling the whole trajectory, and (24) holds by the definition of  $\mathbf{Q}^\pi$ .  $\square$

**Lemma 2 (Multi-Objective Soft Policy Evaluation)** *Under the multi-objective soft Bellman backup operator  $\mathcal{T}_{MO}^\pi$  with respect to a policy  $\pi$  and some  $\mathbf{Q}^{(0)} : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}^d$ , the sequence of intermediate  $\mathbf{Q}$ -functions  $\{\mathbf{Q}^{(i)}\}$  during policy evaluation is given by  $\mathbf{Q}^{(i+1)} = \mathcal{T}_{MO}^\pi \mathbf{Q}^{(i)}$ , for all  $i \in \mathbb{N} \cup \{0\}$ . Then,  $\mathbf{Q}^{(i)}$  converges to the soft  $\mathbf{Q}$ -function of  $\pi$ , as  $i \rightarrow \infty$ .*

**Proof (Lemma 2)** This can be directly obtained from the standard convergence property of iterative policy evaluation (Sutton and Barto, 2018) in two steps: (i) Define the entropy-augmented reward as  $\mathbf{r}(s, a; \pi) := \mathbf{r}(s, a) + \gamma \mathbb{E}_{s' \sim \mathcal{P}(\cdot|s, a), a' \sim \pi(\cdot|s)} [\alpha \log \pi(a|s) \mathbf{1}_d]$ , which is a bounded function. (ii) Then, rewrite the policy evaluation update as

$$\mathbf{Q}^\pi(s, a) \leftarrow \mathbf{r}(s, a; \pi) + \gamma \mathbb{E}_{s' \sim \mathcal{P}(\cdot|s, a), a' \sim \pi(\cdot|s)} [\mathbf{Q}^\pi(s', a'; \lambda)]. \quad (25)$$

This completes the proof.  $\square$

Now we are ready to prove Theorem 3.1.

**Proof (Theorem 3.1)** Note that by Lemma 1, the sequence  $\lambda^\top \mathbf{Q}^{\pi_k}$  is monotonically increasing. As each element in  $\mathbf{Q}^\pi$  is bounded above for all  $\pi \in \Pi$  given the boundedness of both the reward and the entropy term, the sequence of policies shall converge to some policy  $\pi^*$ . The remaining thing is to show that  $\pi^*$  is optimal: (i) Define  $L_{\pi^*}(\pi) := D_{KL} \left( \pi(\cdot | s) \left\| \frac{\exp \left( \sup_{\lambda \in W(\lambda), \mathbf{Q}' \in \mathcal{Q}} \mathbf{Q}'^\pi(s, \cdot; \lambda) \right)}{Z_{\mathcal{Q}}} \right) \right)$ . (ii) Upon convergence, we shall have  $L_{\pi^*}(\pi^*(\cdot|s)) \leq L_{\pi^*}(\pi(\cdot|s))$  for all  $\pi \in \Pi$ . Using the same iterative argument as in the proof of Lemma 1, we get  $\lambda^\top \mathbf{Q}^{\pi^*}(s, a; \lambda) \geq \lambda^\top \mathbf{Q}^\pi(s, a; \lambda)$  for all  $(s, a) \in \mathcal{S} \times \mathcal{A}$  and all  $\lambda \in \Lambda$ .  $\square$

## B DETAILED CONFIGURATION OF EXPERIMENTS

### B.1 DETAILS ON THE EVALUATION DOMAINS

- **Continuous Deep Sea Treasure (DST):** DST is a classical multi-objective reinforcement learning environment. We control the agent to find the treasure, while the further the treasure is, the higher its value. In other words, the agent needs to spend more resources (-1 penalty for each action) to get the more precious treasure. To extend DST to continuous space, we modify the simple four direction movement to the movement in a circle, we set the  $\beta$  of DST to 3.
- **Multi-Objective Continuous LunarLander:** We modify LunarLander to the multi-objective version by dismantling the reward to main engine cost, side engine cost, shaping reward, and result reward. Since the past MORL methods were conducted in environments with 2 or 3 objectives, we created an environment with four objectives to show our method can be used in high dimension objectives environments, we set the  $\beta$  of Lunar Lander to 5.
- **MuJoCo:** We divide the scalar reward in MuJoCo environments into vector rewards. What's more, we amplify the weight of the control cost to make the magnitude of each reward element similar.
  - **HalfCheetah:** 2 objectives as forward speed, control cost ( $\mathcal{S} \subseteq \mathbb{R}^{17}, \mathcal{A} \subseteq \mathbb{R}^6$ ), 1000 times for control cost, and  $\beta = 10$ .

- **Hopper**: 2 objectives: forward speed, control cost ( $\mathcal{S} \subseteq \mathbb{R}^{11}, \mathcal{A} \subseteq \mathbb{R}^3$ ), 1500 times for control cost, and  $\beta = 3$ .
- **Hopper3d**: 3 objectives: forward speed, jump reward, control cost ( $\mathcal{S} \subseteq \mathbb{R}^{11}, \mathcal{A} \subseteq \mathbb{R}^3$ ), 1500 times for control cost. The jump reward is 15 times of the difference between current height and initial height, and  $\beta = 10$ .
- **Ant**: 2 objectives: forward speed, control cost ( $\mathcal{S} \subseteq \mathbb{R}^{11}, \mathcal{A} \subseteq \mathbb{R}^8$ ), 1 times for control cost, and  $\beta = 10$ .
- **Ant3d**: 3 objectives: forward speed, control cost, healthy reward ( $\mathcal{S} \subseteq \mathbb{R}^{11}, \mathcal{A} \subseteq \mathbb{R}^8$ ), 1 times for control cost, 1 times for healthy reward, and  $\beta = 10$ .
- **Walker2d**: 2 objectives: forward speed, control cost ( $\mathcal{S} \subseteq \mathbb{R}^{17}, \mathcal{A} \subseteq \mathbb{R}^6$ ), 1000 times for control cost, and  $\beta = 3$ .

## B.2 HYPERPARAMETERS

### B.2.1 HYPERPARAMETERS OF Q-PENSIEVE

We conduct all experiment on baselines with following hyperparameters.

Parameter	Value
Optimizer	Adam
Learning Rate	0.0003
Discount Factor	0.99
Replay Buffer Size	1000000
Depth of Hidden Layers	2
Number of Hidden Units per Layer	256
Number of Samples per Minibatch	256
Nonlinearity	ReLU
Target Smoothing Coefficient	0.005
Target Update Interval	1
Gradient Steps	1

Table 2: Hyperparameters of  $Q$ -Pensieve

### B.2.2 HYPERPARAMETERS OF PGMORL AND PFA

For PGMORL and PFA, we use the hyperparameters as provided in Table 3:

- $n$ : the number of reinforcement learning tasks.
- $\text{total\_steps}$ : the total number of environment training steps.
- $m_w$ : the number of iterations in warm-up stages.
- $m_t$ : the number of iterations in evolutionary stages.
- $P_{\text{num}}$ : the number of performance buffers.
- $P_{\text{size}}$ : the size of each performance buffer.
- $n_{\text{weight}}$ : the number of sampled weights for each policy.
- $\text{sparsity}$ : the weight of sparsity metric.

## C PSEUDO CODE OF Q-PENSIEVE

We provide the pseudo code in Algorithm 1 as follows.

Environments	$n$	total_steps	$m_w$	$m_t$	$P_{\text{num}}$	$P_{\text{size}}$	$n_{\text{weight}}$	sparsity
DST	5	$1.5 \times 10^6$	80	20	100	2	7	-1
LunarLander	35	$7.5 \times 10^6$	40	10	400	2	7	$-10^6$
HalfCheetah	5	$1.5 \times 10^7$	80	20	100	2	7	-1
Hopper	5	$4.5 \times 10^6$	200	40	100	2	7	-1
Hopper3d	15	$1.5 \times 10^7$	200	40	200	2	7	$-10^6$
Ant	5	$1.5 \times 10^7$	200	40	100	2	7	-1
Ant3d	15	$1.5 \times 10^7$	200	40	200	2	7	$-10^6$
Walker2d	5	$4.5 \times 10^6$	80	20	100	2	7	-1

Table 3: Hyperparameters of PGMORL and PFA

**Algorithm 1:**  $Q$ -Pensieve


**Input :**  $\phi_1, \phi_2, \theta$ , preference sampling distribution  $\mathcal{P}_\lambda$ , number of preference vectors  $N_\lambda$ , the soft update coefficient  $\tau$ , actor learning rates  $\eta_\pi$ , critic learning rates  $\eta_Q$

**Output :**  $\phi_1, \phi_2, \theta$

```

1  $\bar{\phi}_1 \leftarrow \phi_1, \bar{\phi}_2 \leftarrow \phi_2$ ;
2  $\mathcal{M} \leftarrow \emptyset$  ▷ Initialize replay buffer;
3  $\mathcal{B} \leftarrow \emptyset$ ;
4 for each iteration do
5   sample weight  $\lambda$  from  $\Lambda$  according  $\mathcal{P}_\lambda$  ;
6   for each environment step do
7      $a_t \sim \pi_\theta(\cdot | s_t; \lambda)$ ;
8      $s_{t+1} \sim P(\cdot | s_t, a_t)$ ;
9      $\mathcal{M} \leftarrow \mathcal{M} \cup \{(s_t, a_t, r(s_t, a_t), s_{t+1})\}$ ;
10  for each gradient step do
11    sample  $N_\lambda - 1$  preferences and add them to set  $W$ ;
12     $W \leftarrow W \cup \{\lambda\}$ ;
13     $\phi_i \leftarrow \phi_i - \eta_Q \hat{\nabla}_{\phi_i} \mathcal{L}_Q(\phi_i; \lambda), \mathcal{B} \leftarrow \mathcal{B} \cup \{\phi_i; \lambda\}$  for  $i \in \{1, 2\}$ ;
14    compute  $\hat{\nabla}_\theta$  with eq. 11 with  $W$ ;
15     $\theta \leftarrow \theta - \eta_\pi \hat{\nabla}_\theta \mathcal{L}_\pi(\theta; \lambda)$ ;
16     $\bar{\phi}_i \leftarrow \tau \phi_i + (1 - \tau) \bar{\phi}_i$  for  $i \in \{1, 2\}$ ;

```

Figure 6: Return vectors attained under different  $Q$  replay buffer sizes of  $Q$ -Pensieve




Figure 7: Return vectors attained under preference  $\lambda = [0.9, 0.1]$  at different training stages. A number  $x$  on the red or blue marker indicates that the model is obtained at  $10 \cdot x$  million steps.




Figure 8: Return vectors attained under preference  $\lambda = [0.1, 0.9]$  at different training stages. A number  $x$  on the red or blue marker indicates that the model is obtained at  $10 \cdot x$  million steps.