

Pareto Set Learning for Multi-Objective Reinforcement Learning

Erlong Liu, Yu-Chang Wu, Xiaobin Huang, Chengrui Gao,
Ren-Jian Wang, Ke Xue, Chao Qian

National Key Laboratory for Novel Software Technology, Nanjing University, Nanjing 210023, China
School of Artificial Intelligence, Nanjing University, Nanjing 210023, China
{liuel, wuyc, huangxb, gaocr, wangrj, xuek, qianc}@lamda.nju.edu.cn

Abstract

Multi-objective decision-making problems have emerged in numerous real-world scenarios, such as video games, navigation and robotics. Considering the clear advantages of Reinforcement Learning (RL) in optimizing decision-making processes, researchers have delved into the development of Multi-Objective RL (MORL) methods for solving multi-objective decision problems. However, previous methods either cannot obtain the entire Pareto front, or employ only a single policy network for all the preferences over multiple objectives, which may not produce personalized solutions for each preference. To address these limitations, we propose a novel decomposition-based framework for MORL, Pareto Set Learning for MORL (PSL-MORL), that harnesses the generation capability of hypernetwork to produce the parameters of the policy network for each decomposition weight, generating relatively distinct policies for various scalarized subproblems with high efficiency. PSL-MORL is a general framework, which is compatible for any RL algorithm. The theoretical result guarantees the superiority of the model capacity of PSL-MORL and the optimality of the obtained policy network. Through extensive experiments on diverse benchmarks, we demonstrate the effectiveness of PSL-MORL in achieving dense coverage of the Pareto front, significantly outperforming state-of-the-art MORL methods in the hypervolume and sparsity indicators.

Introduction

Reinforcement Learning (RL) [Sutton and Barto 2018] has become a prevalent technique for addressing decision-making and control problems in a wide range of scenarios, including video games [Mnih et al. 2015], navigation [Yang, Sun, and Narasimhan 2019], robotics [Schulman et al. 2017], dynamic algorithm configuration [Xue et al. 2022], and financial cloud services [Yang et al. 2024]. In standard RL, the main objective is obtaining an optimal policy that achieves the best expected returns. However, in many real-world scenarios, spanning from robotics, where the balance between speed and energy efficiency is paramount [Todorov, Erez, and Tassa 2012], to smart systems design, where the trade-offs between competing metrics must be dynamically managed [Azzouz et al. 2018], the

problems may involve multiple conflicting objectives. Contrary to standard RL, Multi-Objective RL (MORL) aims to generate an optimal policy tailored to each preference over multiple objectives, constructing a diverse set of optimal policies that collectively form a Pareto set.

Traditional MORL methods transform the multi-dimensional objective space into a scalar representation [Mannor and Shimkin 2001] through the static assignment of weights to each objective, achieving a trade-off between multiple objectives. Such approaches, however, struggle with objectives of varying magnitudes and require extensive domain knowledge to set the appropriate weights/preferences [Amodei et al. 2016]. Moreover, they inherently limit the solution to a singular policy for a predetermined weight, disregarding the demand for multiple trade-off solutions in real-world scenarios. When extending such approaches to find the dense Pareto set, the inefficiency of repetitively retraining the policy is unacceptable [Parisi et al. 2014]. Recent studies have proposed several multi-policy methods, which simultaneously learn a set of policies over multiple preferences using a single network [Parisi, Pirotta, and Peters 2017, Chen et al. 2019, Xu et al. 2020, Basaklar, Gumussoy, and Ogras 2023]. However, these methods either cannot construct a continuous Pareto front, or will suffer from the curse of dimensionality and poor performance when confronting with the conflicting objectives due to the reason that they only obtain a single network for all the preferences.

To address these limitations, in this work, we integrate Pareto Set Learning (PSL) [Lin et al. 2022, Lin, Yang, and Zhang 2022] within the context of MORL, utilizing the specialty of PSL in covering the whole preference space. Considering that the solutions to MORL problems are generated by the policy networks, we propose to employ hypernetwork [Ha and Le 2017] as the PSL model to generate a continuum of Pareto-optimal policy parameters using a decomposition-based method, and thus generate the relatively personalized policy network for each preference. Since the trained hypernetwork can generate policy parameters for every possible weight without any additional cost (e.g., evaluation and fine-tuning), our proposed framework can dynamically align with user-specified preferences [Abels et al. 2019]. It is worth noting that our method is a general MORL framework, which can be equipped with

any standard RL algorithm. We also perform theoretical analysis, proving the superiority of the model capacity of PSL-MORL by using Rademacher complexity [Golowich, Rakhlin, and Shamir 2018], and the optimality of the obtained policy network by using Banach Fixed Point Theorem [Daskalakis, Tzamos, and Zampetakis 2018].

We conduct experiments to demonstrate the effectiveness of PSL-MORL across several benchmarks, containing the Fruit Tree Navigation (FTN) task with discrete state-action spaces [Yang, Sun, and Narasimhan 2019] and multiple continuous state-action control tasks MO-MuJoCo [Todorov, Erez, and Tassa 2012, Xu et al. 2020]. Since PSL-MORL is a general framework, we employ Double Deep Q-Network (DDQN) [Van Hasselt, Guez, and Silver 2016] for discrete environments and Twin Delayed Deep Deterministic Policy Gradient (TD3) [Fujimoto, Hoof, and Meger 2018] for continuous environments as the lower-level RL algorithms. The experimental results demonstrate that our method is superior to state-of-the-art MORL methods, and the ablation study also verifies the effectiveness of the parameter fusion technique employed by PSL-MORL.

Our contributions are three-fold:

1. We propose a novel MORL method called PSL-MORL, which, to the best of our knowledge, is the first MORL method that covers all preferences over multiple objectives and outputs a personalized policy network for each preference. PSL-MORL is a general framework that can be integrated with any single-objective RL algorithm.
2. The theoretical results guarantee the superiority of the model capacity of PSL-MORL over previous methods, as well as the optimality of the generated policy network.
3. The experiment results show that PSL-MORL outperforms all the other MORL baselines both in the hyper-volume and sparsity metrics.

Background

MORL Problem

A multi-objective sequential decision problem can be characterized as a Multi-Objective Markov Decision Process (MOMDP) [Hayes et al. 2022], represented by a tuple $\langle S, A, T, \gamma, \mu, \mathbf{R} \rangle$, where S is the state space, A is the action space, $T: S \times A \times S \rightarrow [0, 1]$ is a probabilistic transition function, $\gamma = [\gamma_1, \gamma_2, \dots, \gamma_m] \in [0, 1]^m$ is a discount factor vector, $\mu: S \rightarrow [0, 1]$ is a probability distribution over initial states, and $\mathbf{R} = [r_1, r_2, \dots, r_m]^T: S \times A \times S \rightarrow \mathbb{R}^m$ is a vector-valued reward function, specifying the immediate reward for each of the considered m ($m \geq 2$) objectives. When $m = 1$, it is specialized to a single-objective MDP.

In MOMDPs, an agent behaves according to a policy $\pi_\theta \in \Pi$, where Π is the set of all possible policies. A policy is a mapping $\pi_\theta: S \rightarrow A$, which selects an action according to a certain probability distribution for any given state. For brevity, we use π to denote π_θ . The performance of the policy is measured by the associated vector of expected returns, i.e., $\mathbb{F}(\pi) = J^\pi = [J_1^\pi, J_2^\pi, \dots, J_m^\pi]^T$ with

$$J_i^\pi = \mathbb{E} \left[\sum_{k=0}^{\infty} \gamma_i^k r_i(s_k, a_k, s_{k+1}) \mid \pi, \mu \right]. \quad (1)$$

The MORL problem is then formulated as

$$\max_{\pi} \mathbb{F}(\pi) = \max_{\pi} [J_1^\pi, J_2^\pi, \dots, J_m^\pi]^T. \quad (2)$$

In single-objective settings, the expected return offers a complete ordering over the policy space, i.e., for any two policies π and π' , J^π will either be greater than, equal to, or lower than $J^{\pi'}$. Thus, it is sufficient to find an optimal policy π^* that maximizes the expected cumulative discounted reward. However, in MORL, there exists no single optimal policy that can maximize all the objectives simultaneously, and we need the following Pareto concepts [Qian, Yu, and Zhou 2013, Qian et al. 2019, Zhou, Yu, and Qian 2019].

Definition 1 (Pareto Dominance). *For $\pi, \pi' \in \Pi$, π is said to weakly dominate π' ($\pi \succeq \pi'$) if and only if $\forall i \in \{1, \dots, m\}, J_i^\pi \geq J_i^{\pi'}$; π is said to dominate π' ($\pi \succ \pi'$) if and only if $\pi \succeq \pi'$ and $\exists j \in \{1, \dots, m\}, J_j^\pi > J_j^{\pi'}$.*

Definition 2 (Pareto Optimality). *A solution $\pi^* \in \Pi$ is Pareto optimal if $\nexists \hat{\pi} \in \Pi$ such that $\hat{\pi} \succ \pi^*$.*

Definition 3 (Pareto Set/Front). *The set of all Pareto optimal solutions is called the Pareto set, and the image of the Pareto set in the objective space is called the Pareto front.*

Given a fixed utility function u mapping the multi-objective expected returns of a policy to a scalar value, i.e., $u(J^\pi, \omega) = \omega^T J^\pi$, where $\omega = [\omega_1, \dots, \omega_m]^T, \omega_i \geq 0$ and $\sum_i \omega_i = 1$, the problem will be reduced to a single-objective decision-making problem, which can be solved using standard RL methods to generate a Pareto optimal solution corresponding to the preference ω . However, in many complex scenarios, the size of the Pareto front is extremely large. Thus, the goal of MORL is often to efficiently obtain a good approximation of the Pareto front.

MORL Algorithms

Previous MORL approaches can be mainly divided into single-policy, multi-policy, and meta-policy approaches. Single-policy approaches transform a multi-objective problem into a single-objective problem by combining rewards into a single scalar reward using a utility function, and then use standard RL methods to maximize the scalarized return [Roijsers et al. 2013]. However, these approaches require domain-specific knowledge and predefined preferences [Van Moffaert, Drugan, and Nowé 2013, Abdolmaleki et al. 2020]. Multi-policy approaches aim to obtain a good approximated Pareto set, and the most widely used approach is running a single-policy algorithm repeatedly over various preferences [Roijsers et al. 2014, Mossalam et al. 2016, Zuluaga, Krause, and Püschel 2016]. A more recent study proposes an efficient evolutionary learning algorithm to update a population of policies simultaneously in each run to improve the approximation to the Pareto front [Xu et al. 2020]. However, it suffers from the low efficiency of acquiring dense Pareto optimal solutions and handling a large number of objectives [Pirota, Parisi, and Restelli 2015, Parisi, Pirota, and Peters 2017, Chen et al. 2019].

Recent studies have proposed multi-objective Q-learning approaches that simultaneously learn a set of policies over

multiple preferences, using a single network that takes preferences as inputs and uses vectorized value function updates [Abels et al. 2019, Yang, Sun, and Narasimhan 2019]. Another orthogonal approach frames MORL as a meta-learning problem using a task distribution given by a distribution over the preference space [Chen et al. 2019]. It first trains a meta-policy to approximate the Pareto front implicitly, and then obtains the Pareto optimal solution of a given preference by fine-tuning the meta-policy with a few gradient updates. But it often fails because the meta-policy may not be optimal, leading to the ultimate policy being suboptimal. PD-MORL [Basaklar, Gumussoy, and Ogras 2023] is a recently proposed method that utilizes a universal neural network to obtain the entire Pareto front. However, these approaches obtain a single policy network for all the trade-off weights, making it challenging to simultaneously address all scalarized subproblems. Compared to previous methods, our proposed PSL-MORL method generates a policy network for each preference, thereby alleviating conflicts and enhancing the overall quality of the solution set.

Multi-Objective Evolutionary Algorithms (MOEAs), inspired by natural evolution, are another sort of algorithms that can solve MORL problems. They are competitive approaches in solving complicated multi-objective optimization problems [Qian, Yu, and Zhou 2015a,b,c, Qian et al. 2017, Hong et al. 2019] and can be applied to some continuous control problems [Qian and Yu 2021]. They obtain a high quality solution set by iteratively using crossover, mutation, and selection operators. For instance, MOEA/D (MOEA based on Decomposition) [Zhang and Li 2007] is one of the most common MOEA algorithms, which partitions a multi-objective optimization problem into a series of single-objective optimization or less complex multi-objective optimization components. It then employs a heuristic search approach to concurrently and collaboratively refine these easier optimization sub-tasks. However, MOEAs face the challenges in scaling up with the dimensionality of parameters and are not ideally suited for the training of large-scale neural networks in RL.

Pareto Set Learning and Hypernetworks

Previous multi-objective optimization algorithms can only attain a finite Pareto set/front, while PSL [Lin et al. 2020, Navon et al. 2021] encompasses the utilization of a neural network that is tailored to user preferences, enabling the acquisition of the complete Pareto set/front. A standard PSL problem is formulated as $\min_{\theta \in \Theta} g_{WS}(x = h_{\theta}(\lambda) \mid \lambda), \forall \lambda \in \Lambda$, where the function g_{WS} is a weighted sum of the objective functions, i.e., $g_{WS} = \sum_{i=1}^m \lambda_i f_i(x)$, Λ is the preference distribution, and h_{θ} is the Pareto set model and is often implemented as a hypernetwork. If we take the expectation over all preferences, it turns out to be

$$\min_{\theta \in \Theta} \mathbb{E}_{\lambda \sim \Lambda} g_{WS}(x = h_{\theta}(\lambda) \mid \lambda),$$

which can be computed using stochastic optimization, i.e., $\min_{\theta \in \Theta} \sum_i g_{WS}(x = h_{\theta}(\lambda_i) \mid \lambda_i), \{\lambda_i\} \sim \Lambda$. PSL has been successfully applied to multi-objective Bayesian optimization [Lin et al. 2022, Lu, Li, and Zhou 2024], multi-objective neural combinatorial optimization [Lin, Yang,

and Zhang 2022], and drug design [Jain et al. 2023]. Recently, [Zhang et al. 2023] provided a novel geometric perspective for PSL and recognized its equivalence to hypervolume maximization.

The concept of hypernetwork was first introduced in [Ha and Le 2017]. Instead of directly training a primary network to derive the optimized parameters, the hypernetwork takes a different approach: The weights within the primary network remain static, while the adaptable parameters of the hypernetwork are refined through backpropagation and a preferred update mechanism. A key advantage of hypernetworks lies in their dynamic model generation capabilities, which allow for creating customized models for a given input. This is achieved through a single learnable network that provides a versatile and efficient approach to generate a diverse array of models tailored to specific needs.

Considering that PSL is a strong tool in covering the whole preference space and meanwhile hypernetwork has a strong ability in enhancing the model representation ability, we adopt PSL for MORL and implement it as a hypernetwork to generate the parameters for the main policy network.

Pareto Set Learning for MORL

In this section, we introduce the Pareto Set Learning method for MORL (PSL-MORL), an efficient framework for approximating the whole Pareto front of MORL. The overall workflow of PSL-MORL is intuitively illustrated in Figure 1. PSL-MORL consists of two primary modules: a hypernetwork ϕ and a main policy network. The hypernetwork, when given a preference vector ω as input, generates the parameters for the main policy network. This policy network is then used to control an agent in response to various states. During each training episode, a random preference vector ω is drawn from a uniform distribution. After inputting the preference into the hypernetwork, the agent is run in the environment using the generated policy network parameters $\phi(\omega)$. The hypernetwork parameters, along with the target network, are subsequently updated through gradient backpropagation of the RL loss. In practice, to enhance training stability, we co-train another policy network with the hypernetwork using the parameter fusion technique [Ortiz, Guttag, and Dalca 2024].

In the following, we will first introduce the details of the general PSL-MORL framework. Then, we give the instantiations of PSL-MORL with DDQN, and prove the optimality of the generated policy network. Finally, we theoretically analyze the model capacity of PSL-MORL, and prove its superiority over the state-of-the-art method PD-MORL [Basaklar, Gumussoy, and Ogras 2023], which thus provides theoretical justification for its greater ability of generating more personalized policies for various preference vectors.

Details of PSL-MORL

As shown in Algorithm 1, PSL-MORL takes the preference distribution Λ , fundamental elements of RL (i.e., environment \mathcal{E} , replay buffer \mathcal{D} , number E of episodes, and batch size N), and number K of weights per episode as input, and outputs the optimized hypernetwork parameters ϕ and the

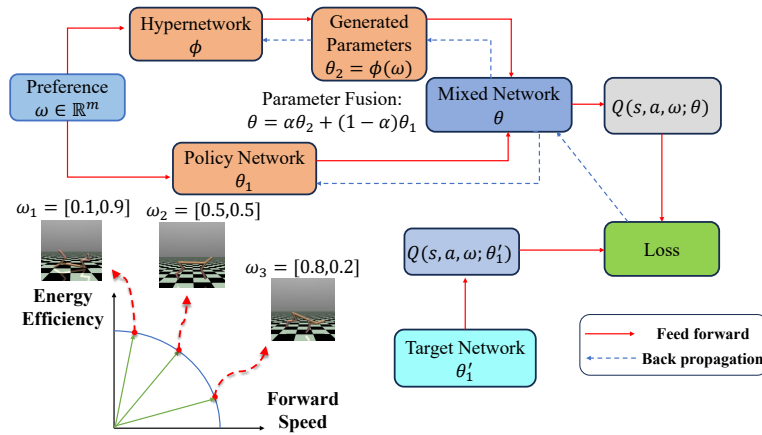


Figure 1: Illustration of the proposed PSL-MORL method. The ultimate parameters of the policy network is composed of two portions, one is the parameters $\theta_2 = \phi(\omega)$ generated by the hypernetwork, and the other is the initial parameters θ_1 of the policy network. The left part is the input of the whole PSL-MORL, i.e., the preference randomly sampled from the uniform distribution. The middle part is the two portions, and the right part is the parameter fusion. Through mixing the parameters, we can get the final policy network and derive the loss to update our hypernetwork and the policy network. The output is the optimal parameters for our hypernetwork and the policy network parameters θ_1 .

Algorithm 1: PSL-MORL

Input: Preference distribution Λ , environment \mathcal{E} , number E of episodes, number K of weights per episode, replay buffer \mathcal{D} , batch size N

Output: Hypernetwork parameters ϕ and primary policy network parameters θ_1

- 1: Initialize the hypernetwork parameters ϕ and policy network parameters θ_1 ;
 - 2: Initialize replay buffer \mathcal{D} ;
 - 3: **for** episode $e = 1$ to E **do**
 - 4: **for** each step in the episode **do**
 - 5: // Parallel child processes for $k \in \{1, 2, \dots, K\}$
 - 6: $\omega_k \sim \text{SamplePreference}(\Lambda)$;
 - 7: Generate parameters of the policy network by hypernetwork, i.e., $\theta_2 = \phi(\omega_k)$;
 - 8: Obtain the mixed θ using the **Parameter Fusion** method on θ_1 and θ_2 ;
 - 9: Observe state s ; select action a by the policy π_θ ;
 - 10: Execute action a in environment \mathcal{E} , and observe reward r and next state s' ;
 - 11: Store transition $(s, a, r, s', d, \omega_k)$ in \mathcal{D} ;
 - 12: // Main process
 - 13: Sample N transitions from \mathcal{D} ;
 - 14: Update ϕ and θ_1 by conducting a single-objective RL algorithm to maximize the scalarized return by weight $\omega_i, i \in \{1, 2, \dots, N\}$
 - 15: **end for**
 - 16: **end for**
-

primary policy network parameters θ_1 . In this work, we use a simple MLP, $\phi(\omega) = \text{MLP}(\omega \mid \phi)$, as our hypernetwork to generate the policy network parameters conditioned on the preference. The detailed structure of the MLP model can be found in Appendix A in the full version.

The objective of training hypernetwork for Pareto set approximation is to maximize the expected scalarized return of MORL, that is, $\max_\phi \mathbb{E}_{\omega \sim \Lambda} [u(J^{\pi_{\phi(\omega)}}), \omega]$, where $\pi_{\phi(\omega)}$ is a policy parameterized by $\phi(\omega)$, and u is the linearized utility function. The optimization can be based on a single-objective RL algorithm. We divide the optimization process into two phases: the parallel child process and the main process. First, we sample multiple weights from the distribution Λ and use the policies generated by the hypernetwork ϕ to collect trajectories. These trajectories, along with the corresponding weights, are then stored in the buffer. In the second phase, a single-objective RL algorithm is applied to optimize the utility function $u(J^{\pi_{\phi(\omega)}}, \omega)$ for each sampled weight. Note that PSL-MORL is compatible with any single-objective RL algorithm.

To be more specific, PSL-MORL first initializes the hypernetwork parameters ϕ , policy network parameters θ_1 , and experience replay buffer \mathcal{D} in lines 1–2 of Algorithm 1. For each step in each episode, we randomly sample K preferences $\{\omega_1, \dots, \omega_K\} \sim \Lambda$ in line 6 for the parallel child process. In line 7, we utilize the hypernetwork to obtain the generated policy parameters θ_2 , and get the mixed policy parameters $\theta = (1 - \alpha)\theta_1 + \alpha\theta_2$ using the parameter fusion technique (which will be presented in detail in the next paragraph) in line 8. After that, we use the policy π_θ to update the environment and store the transitions $(s, a, r, s', d, \omega_k)$ to \mathcal{D} in lines 9–11. In the main process (lines 12–14), we randomly sample N transitions from the buffer \mathcal{D} , and use an RL algorithm (e.g., DDQN and TD3) to update the hypernetwork ϕ and the original policy network θ_1 .

Parameter Fusion for Stable Training As observed in [Ortiz, Gutttag, and Dalca 2024], the performance of hypernetwork prediction has large fluctuations and is highly sensitive to the norm of the input vector, thus leading to unstable training. To address this issue and stabilize the hyper-

network training, we design a parameter fusion technique to blend the parameters predicted by the hypernetwork and the main policy network parameters, i.e.,

$$\theta = (1 - \alpha) \cdot \theta_1 + \alpha \cdot \phi(\omega),$$

where θ_1 denotes the main policy network parameters that are simultaneously optimized with the hypernetwork ϕ , and α controls the proportionality of parameter magnitudes. In this work, we set α to a proper value for different benchmarks, considering the mixing of the generated parameters by the hypernetwork and the parameters from the primary policy network. Different settings of α will be empirically compared in Appendix A.6.

Instantiation of PSL-MORL with DDQN

In this subsection, we introduce an instantiation of PSL-MORL by taking DDQN [Van Hasselt, Guez, and Silver 2016] as the single-objective RL algorithm. The pseudo code is shown in Appendix A.1 due to space limitation. The chosen replay buffer is the Hindsight Experience Replay (HER) [Andrychowicz et al. 2017], which is used to overcome the bias towards over-represented preferences. Inspired by [Basaklar, Gumussoy, and Ogras 2023], we adopt the idea of dividing the preference space into several subspaces to improve the sample efficiency of the algorithm and using the multi-dimensional preference interpolator $I(\omega)$ to align the generated solution with the preference vector.

PSL-MORL with DDQN first initializes the hypernetwork and HER. The parallel child processes are the same as that in Algorithm 1. After that, it operates on a batch of preferences. For each preference, the multi-dimensional interpolator $I(\omega)$ is used to obtain the interpolated preference w_p , and a modified Q-learning update scheme (with the cosine similarity term S_c proposed in [Basaklar, Gumussoy, and Ogras 2023]) is used to calculate the target value. The empirical expected loss is then obtained by calculating the DDQN loss for each preference. Finally, an optimizer (e.g., SGD) is employed to update the parameters of both the hypernetwork and target network.

The following theorem [Yang, Sun, and Narasimhan 2019, Basaklar, Gumussoy, and Ogras 2023] shows that, the Q-network space of the policy is a complete metric space, and the existence of a contraction mapping \mathcal{C} ensures that \mathbf{Q}^* under the optimal policy of any preference is the unique fixed point of this contraction mapping. In other words, for PSL-MORL using DDQN as the basis RL algorithm, the generated policy network can converge to the optimal policy under any preference ω .

Theorem 1. [Yang, Sun, and Narasimhan 2019, Basaklar, Gumussoy, and Ogras 2023] (\mathcal{Q}, d) is a complete metric space, where

$$\forall \mathbf{Q}, \mathbf{Q}' \in \mathcal{Q}, d(\mathbf{Q}, \mathbf{Q}') := \sup_{s \in \mathcal{S}, a \in \mathcal{A}, \omega \in \Lambda} |\omega^T(\mathbf{Q}(s, a, \omega) - \mathbf{Q}'(s, a, \omega))|.$$

Let $\mathbf{Q}^* \in \mathcal{Q}$ be the optimal multi-objective value function, such that it takes multi-objective Q-value corresponding to the supremum of expected discounted rewards under a policy

π . If $\mathcal{C} : \mathcal{Q} \rightarrow \mathcal{Q}$ is a contraction on \mathcal{Q} with modulus γ , then \mathbf{Q}^* is a unique fixed point such that $\mathbf{Q}^* = \mathcal{C}(\mathbf{Q}^*)$.

In the experiments, we also adopt PSL-MORL with TD3 [Dankwa and Zheng 2019], whose details are shown in Appendix A.2. Here, we use the hypernetwork to generate only the parameters of actor network and also conduct the parameter fusion technique. Similar to [Basaklar, Gumussoy, and Ogras 2023], we include the directional angle term to both the actor and critic loss function, which depicts the alignment relationship between the preference vectors and the Q-values.

Theoretical Analysis for PSL-MORL

In this subsection, we theoretically analyze the model representation ability of PSL-MORL, which is characterized by the Rademacher complexity in Definition 4. A model with higher Rademacher complexity will have a larger model capacity [Vapnik 2000, Zhu, Rogers, and Gibson 2009, Shalev-Shwartz and Ben-David 2014, Neyshabur et al. 2019].

Definition 4. (Rademacher Complexity [Golowich, Rakhlin, and Shamir 2018]) Given a real-valued function class \mathcal{H} and a set of data points $\mathbf{x}_1, \dots, \mathbf{x}_n \in \mathcal{X}$, the (empirical) Rademacher complexity $\hat{\mathcal{R}}_n(\mathcal{H})$ is defined as

$$\hat{\mathcal{R}}_n(\mathcal{H}) = \mathbb{E}_\epsilon \left[\sup_{h \in \mathcal{H}} \frac{1}{n} \sum_{i=1}^n \epsilon_i h(\mathbf{x}_i) \right], \quad (3)$$

where $\epsilon = [\epsilon_1, \dots, \epsilon_n]$ is a vector uniformly distributed in $\{-1, +1\}^n$.

We prove in Theorem 2 that PSL-MORL has a higher Rademacher complexity (thus a better model capacity) than PD-MORL [Basaklar, Gumussoy, and Ogras 2023], the most competitive baseline among all the MORL algorithms. This implies a greater ability of PSL-MORL to provide more personalized policies for various preference vectors, as will also be shown in the experiments. The theoretical analysis relies on the common assumption that the activation function σ of a neural network satisfies the 1-Lipschitz property and is positive homogeneous.

Theorem 2. PSL-MORL has a better model capacity compared to PD-MORL, i.e., $\hat{\mathcal{R}}_2 > \hat{\mathcal{R}}_1$, where $\hat{\mathcal{R}}_2$ and $\hat{\mathcal{R}}_1$ denote the Rademacher complexity of PSL-MORL and PD-MORL, respectively.

The proof details of Theorem 2 are provided in Appendix B.3, and we only provide a proof sketch here. By inequality scaling, we first establish a lower bound R_2 for the Rademacher complexity $\hat{\mathcal{R}}_2$ of the hypernetwork MLP used in PSL-MORL, and an upper bound R_1 for the Rademacher complexity $\hat{\mathcal{R}}_1$ of the MLP used in PD-MORL. Subsequently, we prove that the ratio R_2/R_1 has a lower bound greater than 1. Then we have $\hat{\mathcal{R}}_2/\hat{\mathcal{R}}_1 \geq R_2/R_1 > 1$, implying $\hat{\mathcal{R}}_2 > \hat{\mathcal{R}}_1$.

Experiments

In this section, we will give the experimental settings and results. The experiments aim to answer the following two research questions: (1) How is the approximation performance

of PSL-MORL compared to other state-of-the-art MORL algorithms? (2) How does the parameter fusion technique help PSL-MORL find better policies? Due to space limitation, some details are shown in Appendix A.

Experimental Settings

To examine the performance of PSL-MORL, we conduct experiments on two popular MORL benchmarks: a continuous benchmark MO-MuJoCo [Xu et al. 2020] and a discrete benchmark Fruit Tree Navigation (FTN) [Yang, Sun, and Narasimhan 2019]. MO-MuJoCo is a popular MORL benchmark based on the MoJocCo physics simulation environment, consisting of several continuous control tasks. We conduct experiments under five different environments, including MO-HalfCheetah-v2, MO-Hopper-v2, MO-Ant-v2, MO-Swimmer-v2, and MO-Walker-v2. The number of objectives is two for all environments in MO-MuJoCo. FTN is a discrete MORL benchmark with six objectives, whose goal is to navigate the tree to harvest fruit to optimize six nutritional values on specific preferences. Full details of the benchmarks are provided in Appendix A.4.

To evaluate the performance of PSL-MORL, we compare the following seven state-of-the-art methods: PG-MORL [Xu et al. 2020], PD-MORL [Basaklar, Gumussoy, and Ogras 2023], Radial Algorithm (RA) [Parisi et al. 2014], Pareto Front Adaption (PFA) [Parisi et al. 2014], MOEA/D [Zhang and Li 2007], RANDOM [Xu et al. 2020], and META [Chen et al. 2019]. Please refer to Appendix A.5 for details.

An MORL algorithm aims to find a set of policies to approximate the underlying optimal Pareto set and Pareto front. Following previous works [Chen et al. 2019, Xu et al. 2020, Liu and Qian 2021, Basaklar, Gumussoy, and Ogras 2023], we use two metrics: (i) Hypervolume and (ii) Sparsity [Hayes et al. 2022], to compare the quality of the set of policies obtained by different methods.

Hypervolume (HV) Let Q be an approximated Pareto front in an m -dimensional objective space and contain M solutions, more precisely, the objective vectors of M solutions. Let $\mathbf{r}_1 \in \mathbb{R}^m$ be a reference point. The hypervolume indicator is defined as:

$$\text{HV}(Q) := \alpha(H(Q, \mathbf{r}_1)),$$

where $H(Q, \mathbf{r}_1) = \{\mathbf{w} \in \mathbb{R}^m \mid \exists 1 \leq j \leq M, \mathbf{r}_1 \preceq \mathbf{w} \preceq Q_j\}$ with Q_j being the j -th solution in Q , and \preceq is the weak Pareto dominance operator as defined in Definition 1. Note that $\alpha(\cdot)$ is the Lebesgue measure with $\alpha(H(Q, \mathbf{r}_1)) = \int_{\mathbb{R}^m} 1_{H(Q, \mathbf{r}_1)}(\mathbf{w}) d\mathbf{w}$, where $1_{H(Q, \mathbf{r}_1)}$ is the indicator function of $H(Q, \mathbf{r}_1)$.

Sparsity (SP) For an approximated Pareto front Q , sparsity is defined as:

$$\text{Sparsity}(Q) := \frac{1}{M-1} \sum_{k=1}^m \sum_{j=1}^{M-1} (Q_j^k - Q_{j+1}^k)^2,$$

where for each objective, the solutions are sorted in descending order according their values on this objective, and Q_j^k denotes the j -th largest value on the k -th objective. Note that

if there is only one solution in Q , we set the sparsity metric to be N/A.

The hypervolume metric can measure the convergence performance of the approximated Pareto front. However, it cannot fully describe the density of the approximated Pareto front. A dense policy set is always favored for better Pareto approximation [Chen et al. 2019, Xu et al. 2020, Liu and Qian 2021, Basaklar, Gumussoy, and Ogras 2023]. Therefore, it is necessary to consider the sparsity metric, which measures the density of the approximated Pareto front in a multi-objective control problem. As a larger hypervolume implies a better approximated Pareto set and a lower sparsity indicates a denser approximated Pareto set, the goal is to obtain an approximated Pareto set with a high hypervolume metric and meanwhile a low sparsity metric. Note that hypervolume and sparsity are the most frequently used metrics to measure the performance of MORL algorithms [Chen et al. 2019, Xu et al. 2020, Liu and Qian 2021, Basaklar, Gumussoy, and Ogras 2023].

Results on MORL Benchmarks with Continuous State-Action Spaces

We first evaluate the methods on the popular multi-objective continuous control benchmark MO-MuJoCo. We run each method with six different random seeds, consistent with previous works [Xu et al. 2020, Basaklar, Gumussoy, and Ogras 2023]. We do not compare with Envelope Q-Learning [Yang, Sun, and Narasimhan 2019] on MO-MuJoCo since it cannot be applied to continuous action spaces. For a fair comparison, we use the same seeds and hyperparameters for other baseline algorithms. As shown in Table 1, PSL-MORL achieves the best average rank for both metrics, demonstrating the high quality of the obtained Pareto set. PD-MORL is the runner-up, performing worse than PSL-MORL which may be due to the limited capacity and diversity of its single policy model. PG-MORL involves a large optimization space for training multiple policies simultaneously, achieving the third average rank. RANDOM performs worse than PG-MORL, because it does not consider the influence of the choice of weights. Due to the fixed weights and the low efficiency of training separately, RA doesn't have a good performance on the hypervolume metric. PFA also performs poorly since it trains the policy separately and the transfer may not be optimal. The low sample efficiency of the gradient-free optimization in MOEA/D leads to its poor performance. META has the worst hypervolume, because it can only obtain a sub-optimal policy.

Results on MORL Benchmarks with Discrete State-Action Spaces

We evaluate the proposed PSL-MORL on the commonly used benchmark Fruit Tree Navigation (FTN) with different depths, i.e., $d = 5, 6, 7$. As shown in Table 3, PSL-MORL achieves the best performance in both metrics in all the environments. When the depths are 5 and 6, Envelope [Yang, Sun, and Narasimhan 2019] and PD-MORL [Basaklar, Gumussoy, and Ogras 2023] fail to find the whole Pareto set but find only one policy, resulting in an N/A sparsity value.

	Metrics	RA	PFA	MOEA/D	RANDOM	PG-MORL	META	PD-MORL	PSL-MORL
MO-Walker2d-v2	HV ($\times 10^6$)	4.82	4.16	4.44	4.11	4.82	2.10	5.41	5.36
	SP ($\times 10^4$)	0.04	0.37	1.28	0.07	0.04	2.10	0.03	0.01
MO-HalfCheetah-v2	HV ($\times 10^6$)	5.66	5.75	5.61	5.69	5.77	5.18	5.89	5.92
	SP ($\times 10^3$)	15.87	3.81	16.96	1.09	0.44	2.13	0.49	0.16
MO-Ant-v2	HV ($\times 10^6$)	5.98	6.23	6.28	5.54	6.35	2.40	7.48	8.63
	SP ($\times 10^4$)	5.50	1.56	1.97	1.13	0.37	1.56	0.78	0.68
MO-Swimmer-v2	HV ($\times 10^4$)	2.33	2.35	2.42	2.38	2.57	1.23	3.21	3.22
	SP ($\times 10^1$)	4.43	2.49	5.64	1.94	0.99	2.44	0.57	0.42
MO-Hopper-v2	HV ($\times 10^7$)	1.96	1.90	2.03	1.88	2.02	1.25	1.88	1.95
	SP ($\times 10^4$)	5.99	3.96	2.73	1.20	0.50	4.84	0.30	0.68
Average Rank	HV	5.1	5.2	4.2	6.1	2.9	8.0	2.7	1.8
	SP	6.7	5.9	7.0	4.2	2.3	6.1	2.2	1.6

Table 1: Performance comparison between PSL-MORL and other state-of-the-art algorithms on the popular multi-objective continuous control benchmark MO-MuJoCo in terms of hypervolume and sparsity. We run each algorithm with six random seeds and report the mean value of the two metrics. Bold numbers are the best in each row. The reference points for hypervolume calculation are set to $(0, 0)$.

	MO-Walker2d-v2		MO-HalfCheetah-v2		MO-Ant-v2		MO-Swimmer-v2		MO-Hopper-v2	
	HV ($\times 10^6$)	SP ($\times 10^4$)	HV ($\times 10^6$)	SP ($\times 10^3$)	HV ($\times 10^6$)	SP ($\times 10^4$)	HV ($\times 10^4$)	SP ($\times 10$)	HV ($\times 10^7$)	SP ($\times 10^4$)
PSL-MORL(Ours)	5.36	0.01	5.92	0.16	8.63	0.68	3.22	0.42	1.95	0.68
PSL-MORL-gen	1.47	0.55	4.65	0.75	1.39	1.95	2.11	3.64	0.83	6.88
PSL-MORL-add	2.63	1.26	5.62	1.56	4.27	2.13	2.67	3.68	1.28	8.19

Table 2: Ablation study on the parameter fusion technique on the multi-objective continuous control benchmarks. We run all algorithms on each problem with 6 different random seeds and report the mean value of the hypervolume and sparsity metrics. The bold number is the best in each column. The reference points for hypervolume calculation are set to $(0, 0)$.

	FTN (d=5)		FTN(d=6)		FTN (d=7)	
	HV	SP	HV	SP	HV	SP
Envelope	6920.58	N/A	8427.51	N/A	6395.27	N/A
PD-MORL	6920.58	N/A	9299.15	N/A	11419.58	N/A
PSL-MORL	6920.58	0.01	9302.38	0.01	11786.10	N/A

Table 3: Comparison on the multi-objective discrete benchmark FTN in terms of hypervolume and sparsity. The reference points for hypervolume calculation are set to $(0, 0)$.

However, PSL-MORL can find a better approximated Pareto set and get a 0.01 sparsity. The results demonstrate that PSL-MORL has a better exploration capability than Envelope and PD-MORL. Furthermore, as the depth increases from 5 to 7, the hypervolume gap between PSL-MORL and the other two methods also increases, highlighting the greater advantage in more complex environments.

Ablation Study

The above experiments have shown the advantage of PSL-MORL, which employs the parameter fusion technique, i.e., $(1 - \alpha)\theta_1 + \alpha\theta_2$, to merge the parameters from the policy θ_2 generated by the hypernetwork ϕ and the initial original policy θ_1 . For the choice of α , we pick the value that owns the best performance in grid search experiments, as shown in Appendix A.6. This brings two natural questions: (i) Whether the parameter fusion technique (i.e., setting $\alpha < 1$) really improves the performance? (ii) How about adding the two policy parameters (i.e., $\theta_1 + \theta_2$) directly?

To answer the questions, we implement two algorithms on Mo-MuJoCo: (i) PSL-MORL-gen, which uses only the parameters of the policy network generated by the hypernetwork, i.e., without using the parameter fusion technique. (ii) PSL-MORL-add, where the parameters in the policy network are obtained by directly adding the parameters generated by the hypernetwork and the initial policy network. The results are shown in Table 2. We can observe that when obtaining the parameters entirely from the generation of hypernetwork or directly adding the two portions, the performance of the algorithm will drop drastically, demonstrating the effectiveness of our parameter fusion technique.

Conclusion

In this paper, we introduce the PSL-MORL method, which employs the idea of Pareto Set Learning in MORL and can be viewed as adopting the decomposition idea to deal with all the trade-off preferences. PSL-MORL can allow a decision-maker to make real-time decisions without additional procedures. Experimental results show that PSL-MORL outperforms previous methods, and more specifically, PSL-MORL achieves the best hypervolume and sparsity in most benchmarks. Furthermore, the theoretical results guarantee the superiority of the model capacity and the optimality of the generated policy network. One limitation of this work is that we only consider the linearized reward setting. It would be interesting to research on the non-linear scalarization functions and other advanced models, which may bring further performance improvement.

Acknowledgements

The authors want to thank the anonymous reviewers for their helpful comments and suggestions. This work was supported by the National Science and Technology Major Project (2022ZD0116600), the National Science Foundation of China (62276124), and the Fundamental Research Funds for the Central Universities (14380020). Chao Qian is the corresponding author.

References

- Abdolmaleki, A.; Huang, S.; Hasenclever, L.; Neunert, M.; Song, F.; Zambelli, M.; Martins, M.; Heess, N.; Hadsell, R.; and Riedmiller, M. 2020. A distributional view on multi-objective policy optimization. In *Proceedings of the 37th International Conference on Machine Learning*, 11–22. Vienna, Austria.
- Abels, A.; Roijers, D. M.; Lenaerts, T.; Nowé, A.; and Steckelmacher, D. 2019. Dynamic weights in multi-objective deep reinforcement learning. In *Proceedings of the 36th International Conference on Machine Learning*, 11–20. Long Beach, CA.
- Amodei, D.; Olah, C.; Steinhardt, J.; Christiano, P.; Schulman, J.; and Mané, D. 2016. Concrete problems in AI safety. *arXiv:1606.06565*.
- Andrychowicz, M.; Wolski, F.; Ray, A.; Schneider, J.; Fong, R.; Welinder, P.; McGrew, B.; Tobin, J.; Pieter Abbeel, O.; and Zaremba, W. 2017. Hindsight experience replay. In *Advances in Neural Information Processing Systems 30*, 5055–5065. Long Beach, CA.
- Azzouz, R.; Bechikh, S.; Said, L. B.; and Trabelsi, W. 2018. Handling time-varying constraints and objectives in dynamic evolutionary multi-objective optimization. *Swarm and Evolutionary Computation*, 39: 222–248.
- Basaklar, T.; Gumussoy, S.; and Ogras, U. 2023. PD-MORL: Preference-driven multi-objective reinforcement learning algorithm. In *Proceedings of the 11th International Conference on Learning Representations*. Kigali, Rwanda.
- Chen, X.; Ghadirzadeh, A.; Björkman, M.; and Jensfelt, P. 2019. Meta-learning for multi-objective reinforcement learning. In *Proceedings of the 2019 International Conference on Intelligent Robots and Systems*, 977–983. Macau, China.
- Dankwa, S.; and Zheng, W. 2019. Twin-delayed DDPG: A deep reinforcement learning technique to model a continuous movement of an intelligent robot agent. In *Proceedings of the 3rd International Conference on Vision, Image and Signal Processing*, 1–5. Vancouver, Canada.
- Daskalakis, C.; Tzamos, C.; and Zampetakis, M. 2018. A converse to Banach’s fixed point theorem and its CLS-completeness. In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing, Los Angeles, CA*, 44–50.
- Fujimoto, S.; Hoof, H.; and Meger, D. 2018. Addressing function approximation error in actor-critic methods. In *Proceedings of the 35th International Conference on Machine Learning*, 1587–1596. Stockholm, Sweden.
- Golowich, N.; Rakhlin, A.; and Shamir, O. 2018. Size-independent sample complexity of neural networks. In *Proceedings of the 31st Conference on Learning Theory*, 297–299. Stockholm, Sweden.
- Ha, A. D., David; and Le, Q. V. 2017. HyperNetworks. In *Proceedings of the 5th International Conference on Learning Representations*. Toulon, France.
- Hayes, C. F.; Rădulescu, R.; Bargiacchi, E.; Källström, J.; Macfarlane, M.; Reymond, M.; Verstraeten, T.; Zintgraf, L. M.; Dazeley, R.; Heintz, F.; et al. 2022. A practical guide to multi-objective reinforcement learning and planning. *Autonomous Agents and Multi-Agent Systems*, 36(1): 1–59.
- Hong, W.; Tang, K.; Zhou, A.; Ishibuchi, H.; and Yao, X. 2019. A scalable indicator-based evolutionary algorithm for large-scale multiobjective optimization. *IEEE Transactions on Evolutionary Computation*, 23(3): 525–537.
- Jain, M.; Raparthy, S. C.; Hernández-Garcia, A.; Rector-Brooks, J.; Bengio, Y.; Miret, S.; and Bengio, E. 2023. Multi-objective gflownets. In *Proceedings of the 40th International Conference on Machine Learning*, 14631–14653. Honolulu, HI.
- Lin, X.; Yang, Z.; and Zhang, Q. 2022. Pareto set learning for neural multi-objective combinatorial optimization. In *Proceedings of the 10th International Conference on Learning Representations*. Virtual.
- Lin, X.; Yang, Z.; Zhang, Q.; and Kwong, S. 2020. Controllable Pareto multi-task learning. *arXiv:2010.06313*.
- Lin, X.; Yang, Z.; Zhang, X.; and Zhang, Q. 2022. Pareto set learning for expensive multi-objective optimization. In *Advances in Neural Information Processing Systems 35*. New Orleans, LA.
- Liu, F.; and Qian, C. 2021. Prediction guided meta-learning for multi-objective reinforcement learning. In *Proceedings of the 2021 IEEE Congress on Evolutionary Computation*, 2171–2178. Kraków, Poland.
- Lu, Y.; Li, B.; and Zhou, A. 2024. Are you concerned about limited function evaluations: Data-augmented Pareto set learning for expensive multi-objective optimization. In *Proceedings of the 38th AAAI Conference on Artificial Intelligence*, 14202–14210. Vancouver, Canada.
- Mannor, S.; and Shimkin, N. 2001. The steering approach for multi-criteria reinforcement learning. In *Advances in Neural Information Processing Systems 14*, 1563 – 1570. Vancouver, Canada.
- 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.
- Mossalam, H.; Assael, Y. M.; Roijers, D. M.; and Whiteson, S. 2016. Multi-objective deep reinforcement learning. *arXiv:1610.02707*.
- Navon, A.; Shamsian, A.; Fetaya, E.; and Chechik, G. 2021. Learning the Pareto front with hypernetworks. In *Proceedings of the 9th International Conference on Learning Representations*. Vienna, Austria.

- Neyshabur, B.; Li, Z.; Bhojanapalli, S.; LeCun, Y.; and Srebro, N. 2019. The role of over-parametrization in generalization of neural networks. In *Proceedings of the 7th International Conference on Learning Representations*. New Orleans, LA.
- Ortiz, J. J. G.; Gutttag, J.; and Dalca, A. V. 2024. Magnitude invariant parametrizations improve hypernetwork learning. In *Proceedings of the 12th International Conference on Learning Representations*. Vienna, Austria.
- Parisi, S.; Pirotta, M.; and Peters, J. 2017. Manifold-based multi-objective policy search with sample reuse. *Neurocomputing*, 263: 3–14.
- Parisi, S.; Pirotta, M.; Smacchia, N.; Bascetta, L.; and Restelli, M. 2014. Policy gradient approaches for multi-objective sequential decision making. In *Proceedings of the 2014 International Joint Conference on Neural Networks*, 2323–2330. Beijing, China.
- Pirotta, M.; Parisi, S.; and Restelli, M. 2015. Multi-objective reinforcement learning with continuous Pareto frontier approximation. In *Proceedings of the 29th AAAI Conference on Artificial Intelligence*, 2928–2934. Austin, TX.
- Qian, C.; Shi, J.-C.; Yu, Y.; Tang, K.; and Zhou, Z.-H. 2017. Subset selection under noise. In *Advances in Neural Information Processing Systems 30*, 3563–3573. Long Beach, CA.
- Qian, C.; Yu, Y.; Tang, K.; Yao, X.; and Zhou, Z.-H. 2019. Maximizing submodular or monotone approximately submodular functions by multi-objective evolutionary algorithms. *Artificial Intelligence*, 275: 279–294.
- Qian, C.; Yu, Y.; and Zhou, Z.-H. 2013. An analysis on recombination in multi-objective evolutionary optimization. *Artificial Intelligence*, 204: 99–119.
- Qian, C.; Yu, Y.; and Zhou, Z.-H. 2015a. On constrained Boolean Pareto optimization. In *Proceedings of the 24th International Joint Conference on Artificial Intelligence*, 389–395. Buenos Aires, Argentina.
- Qian, C.; Yu, Y.; and Zhou, Z.-H. 2015b. Pareto ensemble pruning. In *Proceedings of the 29th AAAI Conference on Artificial Intelligence*, 2935–2941. Austin, TX.
- Qian, C.; Yu, Y.; and Zhou, Z.-H. 2015c. Subset selection by Pareto optimization. In *Advances in Neural Information Processing Systems 28*, 1765–1773. Montreal, Canada.
- Qian, H.; and Yu, Y. 2021. Derivative-free reinforcement learning: A review. *Frontiers of Computer Science*, 15(6): 156336.
- Roijers, D. M.; Vamplew, P.; Whiteson, S.; and Dazeley, R. 2013. A survey of multi-objective sequential decision-making. *Journal of Artificial Intelligence Research*, 48: 67–113.
- Roijers, D. M.; Whiteson, S.; Oliehoek, F. A.; et al. 2014. Linear support for multi-objective coordination graphs. In *Proceedings of the 2014 International Conference on Autonomous Agents & Multiagent Systems*, 1297–1304. Paris, France.
- Schulman, J.; Wolski, F.; Dhariwal, P.; Radford, A.; and Klimov, O. 2017. Proximal policy optimization algorithms. *arXiv:1707.06347*.
- Shalev-Shwartz, S.; and Ben-David, S. 2014. *Understanding Machine Learning - From Theory to Algorithms*. Cambridge University Press.
- Sutton, R. S.; and Barto, A. G. 2018. *Reinforcement Learning: An Introduction*. MIT press.
- Todorov, E.; Erez, T.; and Tassa, Y. 2012. Mujoco: A physics engine for model-based control. In *Proceedings of the 2012 International Conference on Intelligent Robots and Systems*, 5026–5033. Algarve, Portugal.
- Van Hasselt, H.; Guez, A.; and Silver, D. 2016. Deep reinforcement learning with double Q-learning. In *Proceedings of the 30th AAAI Conference on Artificial Intelligence*, 2094–2100. Phoenix, AZ.
- Van Moffaert, K.; Drugan, M. M.; and Nowé, A. 2013. Scalarized multi-objective reinforcement learning: Novel design techniques. In *Proceedings of the 2013 IEEE Symposium on Adaptive Dynamic Programming and Reinforcement Learning*, 191–199. Singapore.
- Vapnik, V. 2000. *The Nature of Statistical Learning Theory*. Springer.
- Xu, J.; Tian, Y.; Ma, P.; Rus, D.; Sueda, S.; and Matusik, W. 2020. Prediction-guided multi-objective reinforcement learning for continuous robot control. In *Proceedings of the 37th International Conference on Machine Learning*, 10607–10616. Vienna, Austria.
- Xue, K.; Xu, J.; Yuan, L.; Li, M.; Qian, C.; Zhang, Z.; and Yu, Y. 2022. Multi-agent dynamic algorithm configuration. In *Advances in Neural Information Processing Systems 35*. New Orleans, LA.
- Yang, P.; Zhang, L.; Liu, H.; and Li, G. 2024. Reducing idleness in financial cloud services via multi-objective evolutionary reinforcement learning based load balancer. *Science China Information Sciences*, 67(2): 120102.
- Yang, R.; Sun, X.; and Narasimhan, K. 2019. A generalized algorithm for multi-objective reinforcement learning and policy adaptation. In *Advances in Neural Information Processing Systems 32*, 14636–14647. Vancouver, Canada.
- Zhang, Q.; and Li, H. 2007. MOEA/D: A multiobjective evolutionary algorithm based on decomposition. *IEEE Transactions on Evolutionary Computation*, 11(6): 712–731.
- Zhang, X.; Lin, X.; Xue, B.; Chen, Y.; and Zhang, Q. 2023. Hypervolume maximization: A geometric view of pareto set learning. In *Advances in Neural Information Processing Systems 36*. New Orleans, LA.
- Zhou, Z.-H.; Yu, Y.; and Qian, C. 2019. *Evolutionary Learning: Advances in Theories and Algorithms*. Singapore: Springer.
- Zhu, X.; Rogers, T. T.; and Gibson, B. R. 2009. Human rademacher complexity. In *Advances in Neural Information Processing Systems 23*, 2322–2330. Vancouver, British Columbia.
- Zuluaga, M.; Krause, A.; and Püschel, M. 2016. ϵ -pal: An active learning approach to the multi-objective optimization problem. *Journal of Machine Learning Research*, 17(1): 3619–3650.