APPROXIMATING PARETO FRONTIER THROUGH BAYESIAN-OPTIMIZATION-DIRECTED ROBUST MULTI-OBJECTIVE REINFORCEMENT LEARNING

Anonymous authors
Paper under double-blind review

ABSTRACT

Many real-word decision or control problems involve multiple conflicting objectives and uncertainties, which requires learned policies are not only Pareto optimal but also robust. In this paper, we proposed a novel algorithm to approximate a representation for robust Pareto frontier through Bayesian-optimization-directed robust multi-objective reinforcement learning (BRMORL). Firstly, environmental uncertainty is modeled as an adversarial agent over the entire space of preferences by incorporating zero-sum game into multi-objective reinforcement learning (MORL). Secondly, a comprehensive metric based on hypervolume and information entropy is presented to evaluate convergence, diversity and evenness of the distribution for Pareto solutions. Thirdly, the agent's learning process is regarded as a black-box, and the comprehensive metric we proposed is computed after each episode of training, then a Bayesian optimization (BO) algorithm is adopted to guide the agent to evolve towards improving the quality of the approximated Pareto frontier. Finally, we demonstrate the effectiveness of proposed approach on challenging high-dimensional multi-objective tasks across several environments, and show our scheme can produce robust policies under environmental uncertainty.

1 INTRODUCTION

Reinforcement learning (RL) algorithms have demonstrated its worth in a series of challenging sequential decision making and control tasks, which train policies to optimize a single scalar reward function (Mnih et al., 2015; Silver et al., 2016; Haarnoja et al., 2018; Hwangbo et al., 2019). However, many real-world tasks are characterized by multiple competing objectives whose relative importance (preferences) is ambiguous in most cases. Moreover, uncertainty or perturbation caused by environment dynamic change, is inevitable in real-world scenarios, which may result in lowered agent performance (Pinto et al., 2017; Ji et al., 2018). For instance, autonomous electric vehicle requires trading off transport efficiency and electricity consumption while considering environmental uncertainty (e.g., vehicle mass, tire pressure and road conditions might vary over time). Consider a decision-making problem for traffic mode. A practitioner or a rule is responsible for picking the appropriate preference among time and cost, and the agent need to determine different policies depending on the chosen trade-off between these two metrics. If time is crucial, the agent tend to choose plan-A that takes less time, but it costs more. On the other hand, if cost is more important matters, the agent will be inclined to select plan-B that requires less cost, but it takes more time. Whereas, the environment contain uncertainty factors related to actions of other agents or to dynamic changes of Nature, which may lead to more randomness in these two metrics, and makes multi-objective decision-making or control more challenging.

In traditional multi-objective reinforcement learning (MORL), one popular way is scalarization, which is to convert the multi-objective reward vector into a single scalar reward through various techniques (e.g., by taking a convex combination), and then adopt standard RL algorithms to optimize this scalar reward (Vamplew et al., 2011). Unfortunately, it is very tricky to determine an appropriate scalarization, because often common approach only learn an 'average' policy over the space of preferences (Yang et al., 2019), or though the obtained policies can be relatively quickly adapted to different preferences between performance objectives but are not necessarily optimal.

Furthermore, these methods almost did not take into account the robustness of the policies under different preferences, which means the agent cannot learn robust Pareto optimal policies.

In this work, we propose a novel approach to approximate well-distributed robust Pareto frontier through BRMORL. This allows our trained single network model to produce the robust Pareto optimal policy for any specified preference. Our algorithm is based on three key ideas, which are also the main contributions of this paper: (1) present a generalized robust MORL framework through modelling uncertainty as an adversarial agent; (2) inspired by Shannon-Wiener diversity index, a novel metric is presented to evaluate diversity and evenness of distribution for Pareto solutions. In addition, combined with hypervolume indicator, a comprehensive metric is designed, which can evaluate the convergence, diversity and evenness for the solutions on the approximated Pareto frontier; (3) regard agent's learning process in each episode as a black-box, and BO algorithm is used to guide agent to evolve towards improving the quality of the Pareto set. Finally, we demonstrate our proposed algorithm outperform competitive baselines on multi-objective tasks across several MuJoCo (Todorov et al., 2012) environments and SUMO (Simulation of Urban Mobility), and show our approach can produce robust policies under environmental uncertainty.

2 Related work

2.1 Multi-objective reinforcement learning

MORL algorithms can be roughly classified into two main categories: single-policy approaches and multiple-policy approaches (Roijers et al., 2013; Liu et al., 2014). Single-policy methods seek to find the optimal policy for a given preference among multiple competing objectives. These approaches convert the multi-objective problem into a single-objective problem through different forms of scalarization, including linear and non-linear ones (Mannor & Shimkin, 2002; Tesauro et al., 2008). The main advantage of scalarization is its simplicity, which can be integrated into single-policy scheme with very little modification. However, the main drawback of these approaches is that the preference among the objectives must be set in advance.

Multi-policy methods aim to learn a set of policies that approximate Pareto frontier under different preference conditions. The most common approaches repeatedly call a single-policy scheme with different preferences (Natarajan & Tadepalli, 2005; Van Moffaert et al., 2013; Zuluaga et al., 2016). Other methods learn a set of policies simultaneously via using a multi-objective extended version of value-based RL (Barrett & Narayanan, 2008; Castelletti et al., 2012; Van Moffaert & Nowé, 2014; Mossalam et al., 2016; Nottingham et al., 2019) or via modifying policy-based RL as a MORL variant (Pirotta et al., 2015; Parisi et al., 2017; Abdolmaleki et al., 2020; Xu et al., 2020). Nevertheless, most of these methods are offen constrained to convex regions of the Pareto front and explicitly maintain sets of policies, which may prevent these schemes from finding the sets of well-distributed Pareto solutions which can represent different preferences. There are also meta-policy methods, which can be relatively quickly adapted to different preferences (Chen et al., 2018; Abels et al., 2019; Yang et al., 2019). Although the above works were successful to some extent, these approaches share the same shortcomings that no attention is paid to the robustness of Pareto-optimal policy over the entire space of preferences. In addition, most approaches still only work in domains with low-dimensional and discrete action spaces. In contrast, our scheme can guarantee the learned policies is robust Pareto-optimal on high-dimensional continuous control tasks.

2.2 Robust reinforcement learning

Robust reinforcement learning (RRL) algorithms can be broadly grouped into three distinct methods (Derman et al., 2020). The first approach focuses on solving robust Markov decision process (MDP) with rectangular uncertainty sets. Some researches proposed RRL algorithms for learning optimal policies using coupled uncertainty sets (Mannor et al., 2012). Other works modeled an ambiguous linear function of a factor matrix as a selection setting from an uncertainty set (Goyal & Grand-Clement, 2018). The second RRL approach considered a distribution over the uncertainty set to mitigate the conservativeness. Yu & Xu (2015) presented the distributional RRL method by supposing the uncertain parameters are random variables following an unknown distribution. Tirinzoni et al. (2018) proposed a RRL scheme using conditioned probability distribution that defines uncertainty sets. A third RRL method mostly concerns adversarial setting in RL. Pinto et al.

(2017) developed a robust adversarial reinforcement learning (RARL) scheme through modeling uncertainties via adversarial agent which applies disturbances to the system. Tessler et al. (2019) proposed an adversarial RRL framework through structuring probabilistic action robust MDP and noisy action robust MDP. Nonetheless, these researches do not take into account the connection between Pareto-optimal policy and robust policy, which leaves room for improving the performance of them in practical applications. In contrast, our scheme can learn robust Pareto-optimal policies through modeling uncertainty as an adversary over the entire space of preferences.

3 BACKGROUND

3.1 MULTI-OBJECTIVE MARKOV DECISION PROCESS

In this work, we consider a MORL problem defined by a multi-objective Markov decision process (MOMDP), which is represented by the tuple $\langle \mathcal{S}, \mathcal{A}, \mathcal{P}, \boldsymbol{R}, \gamma, \Omega, U_{\boldsymbol{\Omega}} \rangle$ with state space S, action space A, state transition probability $\mathcal{P}(s'|s,a)$, vector reward function $\boldsymbol{R}(s,a) = [r_1, ..., r_k]^{\mathrm{T}}$, the space of preferences Ω, and preference functions, e.g., $U_{\boldsymbol{\omega}}(\boldsymbol{R})$ which produces an utility function using preference $\boldsymbol{\omega} \in \Omega$, and a discount factor $\gamma \in [1, 0)$. In MOMDP, a policy π is associated with a vector of expected returns $\boldsymbol{Q}^{\pi}(s,a) = [Q_1^{\pi}, ..., Q_k^{\pi}]^{\mathrm{T}}$, where the action-value function for objective k can be represented as $Q_k^{\pi}(s,a) = \mathbb{E}_{\pi}[\sum_t \gamma^t r_k(s_t, a_t)|s_0 = s, a_0 = a]$.
For MOMDP, a set of non-dominated policies is called as the Pareto frontier.

Definition 1. *A policy π_1 Pareto dominates another policy π_2, i.e., $\pi_1 \succ \pi_2$ when*

$$\exists i : Q_i^{\pi_1}(s,a) > Q_i^{\pi_2}(s,a) \wedge \forall j \neq i : Q_j^{\pi_1}(s,a) \geqslant Q_j^{\pi_2}(s,a).$$

Definition 2. *A policy π is Pareto optimal if and only if it is non-dominated by any other policies.*

3.2 TWO-PERSON ZERO-SUM GAMES

In standard two-person zero-sum games, players have opposite goals—the payoff of a player equals the loss of the opponent (Mazalov, 2014), i.e., $V + \bar{V} = 0$, where V and $\bar{V}$ are payoff of a player and the opponent, respectively.

Definition 3. *A zero-sum game is a normal-form game $\Gamma = \langle I, II, \Pi, \bar{\Pi}, V \rangle$, where Π, $\bar{\Pi}$ indicate the policy sets of player I (protagonist) and player II (adversary), $V(a, \bar{a}) : \Pi \times \bar{\Pi} \to \mathbb{R}$, $a \in \Pi$ and $\bar{a} \in \bar{\Pi}$.*

For two player discounted zero-sum Markov game, assuming protagonist is playing policy π and adversary is playing the policy $\bar{\pi}$, transition kernel $\mathcal{P}(s'|s, a, \bar{a})$ depend on both players. The value function of the game can be represented as $v^{\pi,\bar{\pi}}(\mathbf{s}) \equiv \mathbb{E}^{\pi,\bar{\pi}}[\sum_{t=0}^{\infty} \gamma^t r(s_t, a_t, \bar{a}_t) \mid s_0 = s], \forall s \in \mathcal{S}$. Each player chooses his policy regardless of the opponent. Protagonist attempts to maximize the value function (i.e., total expected discounted reward), and adversary seeks to minimize this function. Nash equilibrium is a key role in game theory, which is one kind of game solution concept. A Nash equilibrium $(\pi^*, \bar{\pi}^*)$ in zero-sum Markov game exists when the following relation holds (Shapley, 1953; Başar & Olsder, 1998): $v^*(\mathbf{s}) = \max_{\pi} \min_{\bar{\pi}} \mathbb{E}^{\pi,\bar{\pi}}[\sum_{t=0}^{\infty} \gamma^t r(s_t, a_t, \bar{a}_t) \mid s_0 = s]$
$= \min_{\bar{\pi}} \max_{\pi} \mathbb{E}^{\pi,\bar{\pi}}[\sum_{t=0}^{\infty} \gamma^t r(s_t, a_t, \bar{a}_t) \mid s_0 = s]$, where π^* and $\bar{\pi}^*$ are the optimal policies of protagonist and adversary respectively, v^* is optimal equilibrium value of the game. In such a situation, neither player can improve their respective returns, and there is an important relation., i.e., $\forall \pi, \bar{\pi}$, $v^{\pi,\bar{\pi}^*} \leq v^* \leq v^{\pi^*,\bar{\pi}}$.

4 BAYESIAN-OPTIMIZATION-DIRECTED ROBUST MORL

4.1 OVERVIEW

We propose a generalized robust MORL framework to learn a single parametric representation for robust Pareto optimal policy over the space of preferences. In Sections 4.2 and 4.3, through incorporating zero-sum game into MORL, environmental uncertainty is modeled as an adversarial agent. As shown in Figure 1, the policy of the adversary evolves in the opposite direction to the policy of the protagonist in each preference. This means that the protagonist needs to learn Pareto

optimal policy under attack from the adversary. In Section 4.4, inspired by Shannon-Wiener diversity index, a novel metric for Pareto quality is presented to evaluate the distribution of Pareto solutions from diversity and evenness. Moreover, combined with hypervolume index, a comprehensive metric is designed, which can evaluate the convergence, diversity and evenness for solutions in Pareto set. In Section 4.5, regard agent's learning process as a black-box, and the comprehensive metric for the approximated Pareto frontier is computed after each episode of training, then BO algorithm is adopted to guide the protagonist to evolve towards improving the Pareto quality (i.e., maximizing the comprehensive metric).

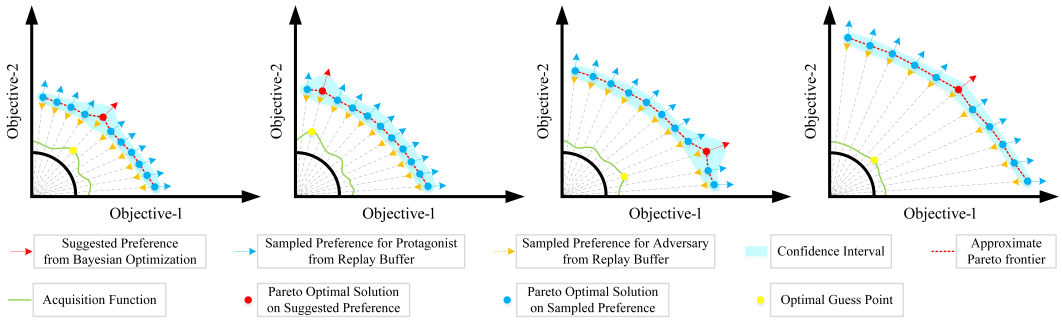

Figure 1: Illustration for process to approximate uniformly distributed robust Pareto frontier through the proposed algorithm.

4.2 Robust multi-objective MDP

In this section, we propose a robust multi-objective MDP (RMO-MDP), which considers both the Pareto optimality and robustness for the learned policies. Probabilistic action robust MDP (PR-MDP) (Tessler et al., 2019) is adopted to improve the robustness of the policies, which can be regarded as a special zero-sum game between a protagonist and an adversary. We refer to the optimal policies of the protagonist as robust Pareto-optimal policies in RMO-MDP, which the difference from the MOMDP is that the action space here includes not only the actions of the protagonist, but also the actions of the adversary with a certain probability.

Definition 4. *A RMO-MDP can be defined by the tuple* $\langle \mathcal{S}, \mathcal{A}^{\text{mix}}, \mathcal{P}, \boldsymbol{R}, \gamma, \Omega, U_{\boldsymbol{\Omega}} \rangle$. $\mathcal{A}^{\text{mix}}$ *is the action space involves the actions of the protagonist with probability* $1-\alpha$ *and the adversary with probability* α*, and* $\alpha \in [0,1]$. *The mixed policy* $\pi_\alpha^{\text{mix}}(\pi, \bar{\pi})$ *is defined as* $\pi_\alpha^{\text{mix}}(a^{\text{mix}} \mid s, \boldsymbol{\omega}) \equiv (1-\alpha)\pi(a \mid s, \boldsymbol{\omega}) + \alpha\bar{\pi}(\bar{a} \mid s, \boldsymbol{\omega}), \forall s \in \mathcal{S}$. π *and* $\bar{\pi}$ *are policies the players can take, and* $a^{\text{mix}} \sim \pi_\alpha^{\text{mix}}(\pi(s), \bar{\pi}(s))$.

In this work, in order to improve the quality of the approximated Pareto frontier, the scalar utility function $U_{\boldsymbol{\Omega}}$ is designed as non-linear combinations of objectives:

$$U_{\boldsymbol{\Omega}}(s, a^{\text{mix}}, \boldsymbol{\omega}) = \boldsymbol{\omega}^\intercal \boldsymbol{Q}^{\pi_\alpha^{\text{mix}}(\pi,\bar{\pi})}(s, a^{\text{mix}}, \boldsymbol{\omega}) + M(s, a^{\text{mix}}, \boldsymbol{\omega}), \tag{1}$$

$$M(s, a^{\text{mix}}, \boldsymbol{\omega}) = k \left\| \frac{\boldsymbol{Q}(s, a^{\text{mix}}, \boldsymbol{\omega})}{\|\boldsymbol{Q}(s, a^{\text{mix}}, \boldsymbol{\omega})\|_2} - \frac{\boldsymbol{\omega}}{\|\boldsymbol{\omega}\|_2} \right\|_2^2, \tag{2}$$

$$\boldsymbol{Q}^{\pi_\alpha^{\text{mix}}(\pi,\bar{\pi})}(s, a^{\text{mix}}, \boldsymbol{\omega}) = (1-\alpha)\boldsymbol{Q}(s, a, \boldsymbol{\omega}) + \alpha \boldsymbol{Q}(s, \bar{a}, \boldsymbol{\omega}), \tag{3}$$

where $M(s, a^{\text{mix}}, \boldsymbol{\omega})$ is a metric, which can evaluate the mismatch between the Pareto optimal solution and the corresponding preference. Figure 2 illustrates metric function $M(s, a^{\text{mix}}, \boldsymbol{\omega})$ in more detail. The distribution of solutions on the Pareto front can be more well-distributed through optimizing function $M(s, a^{\text{mix}}, \boldsymbol{\omega})$. k is a coefficient can adjust the role of $M(s, a^{\text{mix}}, \boldsymbol{\omega})$ in the utility function. For a protagonist, k is a negative, and k is positive for an adversary. This means that the policy with higher preference is more likely to be violently attacked by an adversary, which can makes the policy with higher preference stronger robust.

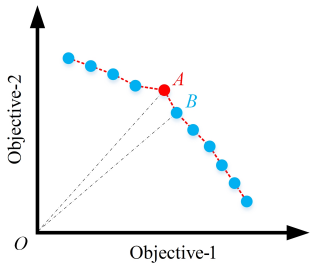

Figure 2: Illustration of the mismatch between the Pareto optimal solution and the corresponding preference. Suppose the point A represents a Pareto optimal solution, which and the origin form the vector $\vec{OA}$. The corresponding preference vector can be represented by $\vec{OB}$. In most cases, the vector $\vec{OA}$ can not parallel to $\vec{OB}$.

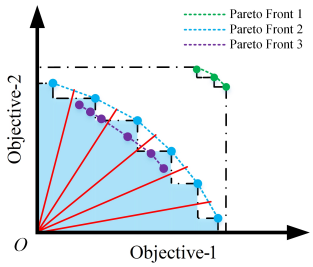

Figure 3: Quality analysis of Pareto frontiers. The Pareto frontiers 1, 2 and 3 are obtained by different approaches. The green, blue and purple points represent the solutions on Pareto frontiers 1, 2 and 3 respectively. The hypervolume formed by the solutions on Pareto front 2 and the reference point O is the blue shaded region.

Under the condition of adversary attack, the utility value of protagonist's policy can be defined as $v_\alpha^\pi \equiv \min_{\bar{\pi}} \mathbb{E}^{\pi_\alpha^{\text{mix}}(\pi,\bar{\pi})}[U_{\boldsymbol{\Omega}}(s, a^{\text{mix}}, \boldsymbol{\omega})]$. Therefore, the robust Pareto optimal policy is optimal policy in RMO-MDP, which can be represent as:

$$\pi_\alpha^* \in \arg\max_{\pi} \min_{\bar{\pi}} \mathbb{E}^{\pi_\alpha^{\text{mix}}(\pi,\bar{\pi})}[U_{\boldsymbol{\Omega}}(s, a^{\text{mix}}, \boldsymbol{\omega})]. \tag{4}$$

The complexity of greedy solution to finding the Nash equilibria policies is exponential in the cardinality of the action spaces, which makes it unworkable in most cases (Schulman et al., 2015). In addition, most two player discounted zero-sum Markov game methods require solving for the equilibrium policy of a minimax action-value function at each iteration. This is a typically intractable optimization problem (Pinto et al., 2017). Instead, we focus on approximating equilibrium solution to avoid this tricky optimization.

4.3 Policy Iteration for RMO-MDP

In this section, we present a policy iteration (PI) approach for solving RMO-MDP called robust multi-objective PI (RMO-PI). RMO-PI algorithm can decompose the RMO-MDP problem into two sub-problems (policy evaluation and policy improvement) and iterate until convergence.

4.3.1 Robust Multi-Objective Policy Evaluation

In this stage, the vectorized $\boldsymbol{Q}$-function is learned to evaluate the policy π of the protagonist. We define the target vectorized $\boldsymbol{Q}$-function as:

$$\boldsymbol{y} \equiv \mathbb{E}^{\pi_\alpha^{\text{mix}}}[\boldsymbol{R} + \gamma \boldsymbol{Q}^{\pi_\alpha^{\text{mix}}(\pi,\bar{\pi})}(s, a^{\text{mix}}, \boldsymbol{\omega}; \phi^-)]. \tag{5}$$

Then, we minimize the following loss function at each step:

$$L_1(\phi) = \mathbb{E}^{\pi_\alpha^{\text{mix}}}\left[\|\boldsymbol{y} - \boldsymbol{Q}(s, a, \boldsymbol{\omega}; \phi)\|_2^2\right], \tag{6}$$

where $\boldsymbol{y} = \mathbb{E}^{\pi_\alpha^{\text{mix}}}[\boldsymbol{R} + \gamma \boldsymbol{Q}^{\pi_\alpha^{\text{mix}}(\pi,\bar{\pi})}(s', a^{\text{mix}}, \boldsymbol{\omega}'; \phi^-)]$, ϕ and ϕ^- are the parameters of the $\boldsymbol{Q}$-function network and the target $\boldsymbol{Q}$-function network respectively. In order to improve the smoothness of the landscape of loss function, the auxiliary loss function is used (Yang et al., 2019):

$$L_2(\phi) = \mathbb{E}^{\pi_\alpha^{\text{mix}}}\left[\|\boldsymbol{\omega}^\intercal \boldsymbol{y} - \boldsymbol{\omega}^\intercal \boldsymbol{Q}(s, a, \boldsymbol{\omega}; \phi)\|_2^2\right]. \tag{7}$$

The final loss function can be written as: $L(\phi) = (1-\beta)L_1(\phi) + \beta L_2(\phi)$, where β is a weighting coefficient to trade off between losses $L_1(\phi)$ and $L_2(\phi)$.

4.3.2 ROBUST MULTI-OBJECTIVE POLICY IMPROVEMENT

In RMO-PI, policy improvement refers to optimizing and updating the policies of a protagonist and an adversary for the given utility function. RMO-PI optimizes both of the agents through the following alternating process. In the first stage, the policy of protagonist is learned while holding the adversary's policy fixed. In the second stage, the policy of protagonist is held constant and the adversary's policy is learned. This learning sequence is repeated until convergence.

The protagonist seeks to maximize the utility function $U_{\boldsymbol{\Omega}}$, and then the policy gradient can be represented as:

$$\begin{aligned}\nabla_\theta L_\pi &\approx \mathbb{E}^{\pi_\alpha^{\text{mix}}(\pi,\bar{\pi})}[\nabla_\theta U_{\boldsymbol{\Omega}}(s, a^{\text{mix}}, \boldsymbol{\omega})] \\ &= \mathbb{E}^{\pi_\alpha^{\text{mix}}(\pi,\bar{\pi})}[(1-\alpha)\nabla_a \boldsymbol{\omega}^\intercal \boldsymbol{Q}(s, a, \boldsymbol{\omega}; \phi)\nabla_\theta \pi(s, \boldsymbol{\omega}; \theta) + \nabla_a M(s, a, \boldsymbol{\omega})\nabla_\theta \pi(s, \boldsymbol{\omega}; \theta)],\end{aligned} \tag{8}$$

where θ is the model parameters of the protagonist. Next, the adversary tries to minimize the utility function $U_{\boldsymbol{\Omega}}$, and the policy gradient can be written as:

$$\begin{aligned}\nabla_{\bar{\theta}} L_{\bar{\pi}} &\approx \mathbb{E}^{\pi_\alpha^{\text{mix}}(\pi,\bar{\pi})}[\nabla_{\bar{\theta}} U_{\boldsymbol{\Omega}}(s, a^{\text{mix}}, \boldsymbol{\omega})] \\ &= \mathbb{E}^{\pi_\alpha^{\text{mix}}(\pi,\bar{\pi})}[\alpha\nabla_{\bar{a}} \boldsymbol{\omega}^\intercal \boldsymbol{Q}(s, \bar{a}, \boldsymbol{\omega}; \phi)\nabla_{\bar{\theta}} \bar{\pi}(s, \boldsymbol{\omega}; \bar{\theta}) + \nabla_{\bar{a}} M(s, \bar{a}, \boldsymbol{\omega})\nabla_{\bar{\theta}} \bar{\pi}(s, \boldsymbol{\omega}; \bar{\theta})],\end{aligned} \tag{9}$$

where $\bar{\theta}$ is the model parameters of the adversary.

4.4 METRICS FOR PARETO REPRESENTATION

Since the true Pareto set is intractable to obtain in complex problems, the goal of MORL is to find the set of policies that best approximates the optimal Pareto front. Hypervolume indicator is widely adopted to evaluate the quality of an approximated Pareto frontier, which can measure the convergence and uniformity for the distribution of Pareto solutions (Zitzler & Thiele, 1999; Xu et al., 2020). From our perspective, this metric may be difficult to accurately measure the uniformity of the Pareto solution distribution. As shown in Figure 3, compared with the Pareto frontiers 2 and 3, although the hypervolume metric formed by the solutions on Pareto frontier 1 and the reference point O is optimal, the distribution of solutions on the frontier 1 is not well-distributed, which makes the valid preferences of the practitioner or the agent to choose is very limited. Moreover, suppose that the solutions on Pareto frontier 1 are very close to each other or even overlap into one solution. At this time, if we adopt the metric proposed in the paper of (Xu et al., 2020) to measure the quality of Pareto frontier 1, the result to have high hypervolume and low sparsity is very ideal. However, such Pareto frontier 1 might not satisfy the needs of the practitioner or the agent. In a word, the high quality of the approximated Pareto frontier is expected to have high hypervolume, and the distribution of solutions is well-distributed. Therefore, in this section, we proposed a novel metric for quality of the approximated Pareto frontier through combining hypervolume metric and evenness metric.

Inspired by Shannon-Wiener diversity index, the diversity metric for the solutions of the Pareto frontier can be expressed as $D(P) = -\sum[p_i \ln(p_i)]$, where P represents the solutions of the Pareto frontier, and p_i is the proportion of the number of non-dominated solutions in the corresponding solution interval to the total number of the solutions on Pareto frontier. The expected diversity of Pareto set D_{max} can be defined as $\ln(S_n)$, and S_n is the number of solution intervals. Then, our evenness metric $E(P)$ can be represented as $D(P)/D_{max}$. For example, in Figure 3, $S_n = 6$, and the evenness metrics for the distribution of the solutions on the Pareto frontiers 1, 2 and 3 are approximately equal to 0.37, 1 and 0.56, respectively. Hence, we can get the following two inferences.

Proposition 1. *As $E(P)$ and S_n increases, the distribution of solutions in Pareto set becomes denser and more uniform, and the Pareto frontier becomes more continuous.*

Proposition 2. *The Pareto frontier is continuous as* $E(P) = 1$ *and* $S_n \to \infty$.

Combined with the hypervolume indicator $H(P)$, we propose a comprehensive metric $I(P)$ that can measure the convergence, diversity and evenness of the solutions: $I(P) = H(P)(1 + \lambda E(P))$, where λ is a weight coefficient.

4.5 Pareto representation improvement based on Bayesian-optimization-directed

In this Section, in order to further improve the representation of the approximated Pareto frontier, the agent's learning process is regarded as a black-box, and the comprehensive metric $I(P)$ is computed after each episode of training, then a BO algorithm is adopted to guide the protagonist to evolve towards maximizing the proposed metric $I(P)$. As shown in Figure 4, the Pareto representation improvement scheme based on BO-directed is illustrated. The value of the objective function $f(\Omega)$ equals the value of the comprehensive metric $I(P)$, which is obtaind after each episode of training. In addition, suggested preferences from BO and sampled preferences from replay buffer are simultaneously used to guide the learning process, which is to avoid the algorithm into a local optimum. The scheme to guide the learning process with Bayesian optimization has high universality for Pareto quality improvement, and does not require much expert experience in the selection of prediction models.

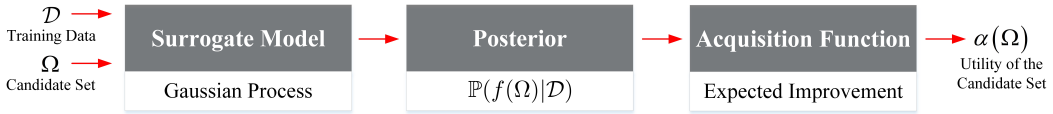

Figure 4: Illustration for Pareto representation improvement scheme based on BO algorithm. The surrogate model for the objective function $f(\Omega)$ is typically a Gaussian Process. Posteriors represent the confidence a model has about the function values at a point or set of points. Acquisition function is employed to evaluate the usefulness of optimal guess point corresponding to posterior distribution over $f(\Omega)$. The expected improvement method chosen to design the acquisition function in our scheme.

5 Experiments

In order to benchmark our proposed scheme, we develop two MORL environments with continuous action space based on SUMO and Mujoco. In addition, we also adopted two MORL environments provided by Xu et al. (2020). The goal of all tasks is to try to optimize the speed of the agent while minimizing energy consumption.

Our algorithm is implemented based on Deep Deterministic Policy Gradient (DDPG) (Lillicrap et al., 2015) framework. In principle, our scheme can be combined with any RL method, regardless of whether it is off-policy or on-policy. Moreover, we implement five baseline methods for comparison and ablation analysis: SMORL represents a MO-DDPG method based on scalarization function, which is a linear combination of rewards in the form of a preference; SRMORL is a RMO-DDPG approach using the scalarization function; RMORL represents a RMO-DDPG approach with the utility function $U_{\mathbf{\Omega}}$. BRMORL is a RMO-DDPG scheme combined with the utility function $U_{\mathbf{\Omega}}$ and BO algorithm.

Figure 5 shows the learning curves and Pareto frontiers comparison results on SUMO. In addition, the results in Table 1 and Table 2 demonstrate that our proposed BRMORL scheme outperforms all the baseline methods on SUMO and Swimmer-v2 environments in hypervolume and evenness metrics.

Table 1: Training results on SUMO.

	Hypervolume	Evenness
SMORL	4547.43 ± 1919.14	0.45 ± 0.30
SRMORL	4904.25 ± 2480.33	0.42 ± 0.37
RMORL	5900.67 ± 2443.58	0.67 ± 0.36
BRMORL	6219.57 ± 2164.15	0.81 ± 0.24

Table 2: Training results on Swimmer-v2.

	Hypervolume	Evenness
SMORL	2309.82 ± 2079.79	0.39 ± 0.36
SRMORL	3276.50 ± 1791.15	0.44 ± 0.21
RMORL	4146.06 ± 2733.65	0.67 ± 0.20
BRMORL	5363.62 ± 1745.60	0.72 ± 0.17

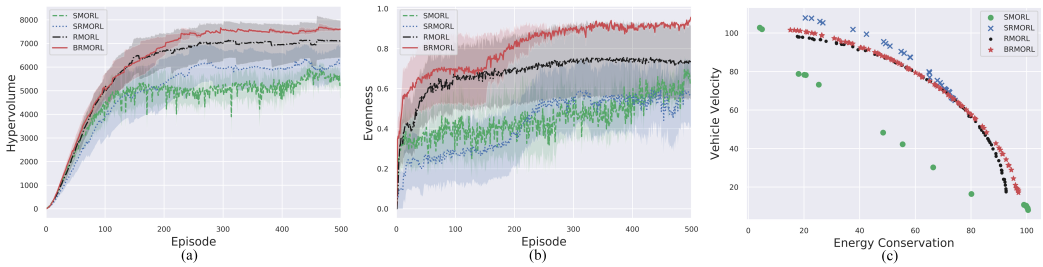

Figure 5: The learning curves and the Pareto frontiers obtained by different algorithms on SUMO.

Figure 6 illustrates that the robustness of different models under different preferences to environmental uncertainty, on Swimmer-v2 domain. We test with jointly varying both mass and disturbance probability. Obviously, BRMORL scheme can lead significantly higher reward values compared to the SMORL method in testing.

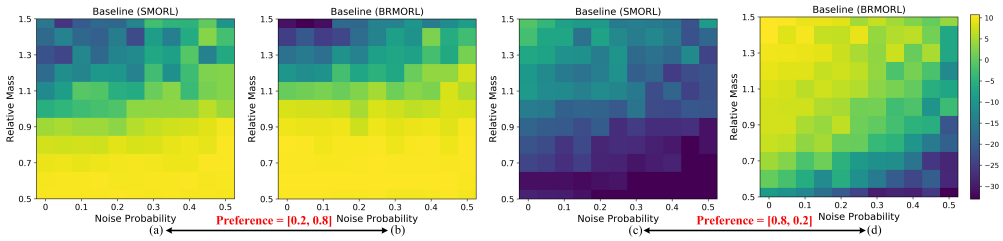

Figure 6: Robustness to environmental uncertainty. Noise probability represents the probability of a random disturbance being played instead of the selected action. Relative mass denotes the ratio of the current agent's mass to its original mass.

In Table 3, we compare our BRMORL scheme with classic and state-of-the-art baselines provided by Xu et al. (2020), and the results show that our approach outperforms the baselines in hypervolume metric, which is among the most widespread indicators in multi-objective task.

Table 3: Comparison of different MORL methods on Walker2d-v2 and HalfCheetah-v2 domains.

	RA	PFA	MOEA/D	RANDOM	META	PGMORL	BRMORL
Hypervolume $\times 10^6$ (Walker2d-v2)	4.15	4.16	4.44	4.11	2.10	4.82	8.84
Hypervolume $\times 10^6$ (HalfCheetah-v2)	5.66	5.75	5.61	5.69	5.18	5.77	12.76

6 Conclusion and Discussion

In this paper, we proposed a novel algorithm to approximate a representation for robust Pareto frontier, which allows our trained single network model to produce the robust Pareto optimal policy for any specified preference. First, environmental uncertainty is modeled as an adversarial agent over the entire space of preferences. Second, a comprehensive metric is constructed through combining with hypervolume and evenness index we proposed. Third, the training process in each episode is regarded as a black-box, then a BO algorithm is adopted to guide the agent's learning process.

Our experiments across four different domains demonstrate that our schemes is effective and advanced. Most importantly, we note that training with appropriate adversarial setting can not only result in robust policies, but also improve the performance even. Although our approach cannot guarantee the learned policies are optimal, it is approximately robust Pareto optimal.

REFERENCES

Abbas Abdolmaleki, Sandy H Huang, Leonard Hasenclever, Michael Neunert, H Francis Song, Martina Zambelli, Murilo F Martins, Nicolas Heess, Raia Hadsell, and Martin Riedmiller. A distributional view on multi-objective policy optimization. *arXiv preprint arXiv:2005.07513*, 2020.

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*, pp. 11–20. PMLR, 2019.

Leon Barrett and Srini Narayanan. Learning all optimal policies with multiple criteria. In *Proceedings of the 25th international conference on Machine learning*, pp. 41–47, 2008.

Tamer Başar and Geert Jan Olsder. *Dynamic noncooperative game theory*. SIAM, 1998.

Andrea Castelletti, Francesca Pianosi, and Marcello Restelli. Tree-based fitted q-iteration for multi-objective markov decision problems. In *The 2012 International Joint Conference on Neural Networks (IJCNN)*, pp. 1–8. IEEE, 2012.

Xi Chen, Ali Ghadirzadeh, Mårten Björkman, and Patric Jensfelt. Meta-learning for multi-objective reinforcement learning. *arXiv preprint arXiv:1811.03376*, 2018.

Esther Derman, Daniel Mankowitz, Timothy Mann, and Shie Mannor. A bayesian approach to robust reinforcement learning. In *Uncertainty in Artificial Intelligence*, pp. 648–658. PMLR, 2020.

Vineet Goyal and Julien Grand-Clement. Robust markov decision process: Beyond rectangularity. *arXiv preprint arXiv:1811.00215*, 2018.

Tuomas Haarnoja, Aurick Zhou, Kristian Hartikainen, George Tucker, Sehoon Ha, Jie Tan, Vikash Kumar, Henry Zhu, Abhishek Gupta, Pieter Abbeel, et al. Soft actor-critic algorithms and applications. *arXiv preprint arXiv:1812.05905*, 2018.

Jemin Hwangbo, Joonho Lee, Alexey Dosovitskiy, Dario Bellicoso, Vassilios Tsounis, Vladlen Koltun, and Marco Hutter. Learning agile and dynamic motor skills for legged robots. *Science Robotics*, 4(26), 2019.

Xuewu Ji, Xiangkun He, Chen Lv, Yahui Liu, and Jian Wu. Adaptive-neural-network-based robust lateral motion control for autonomous vehicle at driving limits. *Control Engineering Practice*, 76: 41–53, 2018.

Timothy P Lillicrap, Jonathan J Hunt, Alexander Pritzel, Nicolas Heess, Tom Erez, Yuval Tassa, David Silver, and Daan Wierstra. Continuous control with deep reinforcement learning. *arXiv preprint arXiv:1509.02971*, 2015.

Chunming Liu, Xin Xu, and Dewen Hu. Multiobjective reinforcement learning: A comprehensive overview. *IEEE Transactions on Systems, Man, and Cybernetics: Systems*, 45(3):385–398, 2014.

Shie Mannor and Nahum Shimkin. The steering approach for multi-criteria reinforcement learning. In *Advances in Neural Information Processing Systems*, pp. 1563–1570, 2002.

Shie Mannor, Ofir Mebel, and Huan Xu. Lightning does not strike twice: Robust mdps with coupled uncertainty. *arXiv preprint arXiv:1206.4643*, 2012.

Vladimir Mazalov. *Mathematical game theory and applications*. John Wiley & Sons, 2014.

Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A Rusu, Joel Veness, Marc G Bellemare, Alex Graves, Martin Riedmiller, Andreas K Fidjeland, Georg Ostrovski, et al. Human-level control through deep reinforcement learning. *nature*, 518(7540):529–533, 2015.

Hossam Mossalam, Yannis M Assael, Diederik M Roijers, and Shimon Whiteson. Multi-objective deep reinforcement learning. *arXiv preprint arXiv:1610.02707*, 2016.

Sriraam Natarajan and Prasad Tadepalli. Dynamic preferences in multi-criteria reinforcement learning. In *Proceedings of the 22nd international conference on Machine learning*, pp. 601–608, 2005.

Kolby Nottingham, Anand Balakrishnan, Jyotirmoy Deshmukh, Connor Christopherson, and David Wingate. Using logical specifications of objectives in multi-objective reinforcement learning. *arXiv preprint arXiv:1910.01723*, 2019.

Simone Parisi, Matteo Pirotta, and Jan Peters. Manifold-based multi-objective policy search with sample reuse. *Neurocomputing*, 263:3–14, 2017.

Lerrel Pinto, James Davidson, Rahul Sukthankar, and Abhinav Gupta. Robust adversarial reinforcement learning. *arXiv preprint arXiv:1703.02702*, 2017.

Matteo Pirotta, Simone Parisi, and Marcello Restelli. Multi-objective reinforcement learning with continuous pareto frontier approximation. In *29th AAAI Conference on Artificial Intelligence, AAAI 2015 and the 27th Innovative Applications of Artificial Intelligence Conference, IAAI 2015*, pp. 2928–2934. AAAI Press, 2015.

Diederik M Roijers, Peter Vamplew, Shimon Whiteson, and Richard Dazeley. A survey of multi-objective sequential decision-making. *Journal of Artificial Intelligence Research*, 48:67–113, 2013.

John Schulman, Sergey Levine, Pieter Abbeel, Michael Jordan, and Philipp Moritz. Trust region policy optimization. In *International conference on machine learning*, pp. 1889–1897, 2015.

Lloyd S Shapley. Stochastic games. *Proceedings of the national academy of sciences*, 39(10): 1095–1100, 1953.

David Silver, Aja Huang, Chris J Maddison, Arthur Guez, Laurent Sifre, George Van Den Driessche, Julian Schrittwieser, Ioannis Antonoglou, Veda Panneershelvam, Marc Lanctot, et al. Mastering the game of go with deep neural networks and tree search. *nature*, 529(7587):484–489, 2016.

Gerald Tesauro, Rajarshi Das, Hoi Chan, Jeffrey Kephart, David Levine, Freeman Rawson, and Charles Lefurgy. Managing power consumption and performance of computing systems using reinforcement learning. In *Advances in Neural Information Processing Systems*, pp. 1497–1504, 2008.

Chen Tessler, Yonathan Efroni, and Shie Mannor. Action robust reinforcement learning and applications in continuous control. *arXiv preprint arXiv:1901.09184*, 2019.

Andrea Tirinzoni, Marek Petrik, Xiangli Chen, and Brian Ziebart. Policy-conditioned uncertainty sets for robust markov decision processes. In *Advances in Neural Information Processing Systems*, pp. 8939–8949, 2018.

Emanuel Todorov, Tom Erez, and Yuval Tassa. Mujoco: A physics engine for model-based control. In *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 5026–5033. IEEE, 2012.

Peter Vamplew, Richard Dazeley, Adam Berry, Rustam Issabekov, and Evan Dekker. Empirical evaluation methods for multiobjective reinforcement learning algorithms. *Machine learning*, 84 (1-2):51–80, 2011.

Kristof Van Moffaert and Ann Nowé. Multi-objective reinforcement learning using sets of pareto dominating policies. *The Journal of Machine Learning Research*, 15(1):3483–3512, 2014.

Kristof Van Moffaert, Madalina M Drugan, and Ann Nowé. Scalarized multi-objective reinforcement learning: Novel design techniques. In *2013 IEEE Symposium on Adaptive Dynamic Programming and Reinforcement Learning (ADPRL)*, pp. 191–199. IEEE, 2013.

Jie Xu, Yunsheng Tian, Pingchuan Ma, Daniela Rus, Shinjiro Sueda, and Wojciech Matusik. Prediction-guided multi-objective reinforcement learning for continuous robot control. In *Proceedings of the 37th International Conference on Machine Learning*, 2020.

Runzhe Yang, Xingyuan Sun, and Karthik Narasimhan. A generalized algorithm for multi-objective reinforcement learning and policy adaptation. In *Advances in Neural Information Processing Systems*, pp. 14636–14647, 2019.

Pengqian Yu and Huan Xu. Distributionally robust counterpart in markov decision processes. *IEEE Transactions on Automatic Control*, 61(9):2538–2543, 2015.

Eckart Zitzler and Lothar Thiele. Multiobjective evolutionary algorithms: a comparative case study and the strength pareto approach. *IEEE transactions on Evolutionary Computation*, 3(4):257–271, 1999.

Marcela Zuluaga, Andreas Krause, and Markus Püschel. ε-pal: an active learning approach to the multi-objective optimization problem. *The Journal of Machine Learning Research*, 17(1): 3619–3650, 2016.