

UCB-driven Utility Function Search for Multi-objective Reinforcement Learning

Yucheng Shi¹, David Lynch², and Alexandros Agapitos² (✉)

¹ Trinity College Dublin, Ireland shiy2@tcd.ie

² Huawei Ireland Research Centre, Ireland {david.lynch1, alexandros.agapitos}@huawei.com

Abstract. In Multi-objective Reinforcement Learning (MORL) agents are tasked with optimising decision-making behaviours that trade-off between multiple, possibly conflicting, objectives. MORL based on decomposition is a family of solution methods that employ a number of utility functions to decompose the multi-objective problem into individual single-objective problems solved simultaneously in order to approximate a Pareto front of policies. We focus on the case of linear utility functions parametrised by weight vectors \mathbf{w} . We introduce a method based on Upper Confidence Bound to efficiently search for the most promising weight vectors during different stages of the learning process, with the aim of maximising the hypervolume of the resulting Pareto front. The proposed method demonstrates consistency and strong performance across various MORL baselines on Mujoco benchmark problems. The code is released in: <https://github.com/SYCAMORE-1/ucb-MOPPO>

Keywords: Multi-objective Reinforcement Learning · Upper Confidence Bound · Mujoco benchmark problems.

1 Introduction

In many real-world control and planning problems, multiple and often conflicting objectives arise. These objectives are interrelated, requiring trade-offs that significantly affect the overall quality of decision-making [30, 24]. Such learning objectives are usually represented as weighted reward signals, where conflicting rewards can lead to divergent optimisation directions [35]. Consequently, the classic reinforcement learning (RL) methods are inadequate, as training individual policies to align with each preference weight vector across multiple rewards results in an impractical computational burden [13, 10]. Therefore, Multi-objective reinforcement learning (MORL) has become an increasingly recognised methodology for handling tasks with conflicting objectives, enabling the simultaneous optimisation of multiple criteria [31, 3, 36]. In cases where the objectives conflict with each other, the best trade-offs among the objectives are defined in terms of Pareto optimal. A policy π is said to dominate π' if $(\forall i : \mathbf{V}_i^\pi \geq \mathbf{V}_i^{\pi'}) \wedge (\exists i : \mathbf{V}_i^\pi > \mathbf{V}_i^{\pi'})$, and it is Pareto optimal if there is no other policy that dominates it. The set of Pareto optimal policies is called the Pareto front (PF).

In many real-world applications, an approximation to the PF is required by a decision maker in order to select a preferred policy [13]. A real-world use case is the multi-objective optimisation of wireless networks [11]. In this case, four primary objective classes, i.e., throughput, coverage, energy efficiency and utilisation are considered. This is a very challenging problem that motivates further research to advance the state-of-the-art (SOTA) in MORL.

One prevalent category of MORL approaches is Multi-Objective Reinforcement Learning with Decomposition (MORL/D) [10]. A Pareto-optimal solution to a multi-objective sequential decision-making problem can be interpreted as an optimal solution to a single-objective problem, defined with respect to a utility function $u : \mathbb{R}^m \rightarrow \mathbb{R}$. This utility function aggregates m objectives into a scalar reward, mapping the multi-objective value of the policy into a scalar value, expressed as $V_u^\pi = u(\mathbf{V}^\pi)$ [8]. Therefore, the approximation of the PF can be reformulated as a set of scalar reward sub-problems, each defined by a utility function. This approach unifies MORL methods [13] with decomposition-based multi-objective optimisation techniques [34], collectively termed MORL/D.

MORL/D offers several advantages, including scalability to many objectives, flexibility in employing various scalarisation techniques, and the ability to parallelise sub-problem training, thereby reducing computational overhead [9, 10]. SOTA MORL/D methods can be broadly categorised into single-policy and multi-policy approaches [10]. Single-policy methods [1, 2, 18, 32] aim to learn a single policy conditioned on the parameters of the utility function, i.e., scalarisation vector \mathbf{w} , where $\forall i : w_i \geq 0$, and $\sum_i w_i = 1.0$. However, these methods face two major challenges: (1) single-policy methods require large neural networks and extensive training time to accurately reconstruct the optimal PF; and (2) they may struggle to generalise effectively to unseen \mathbf{w} [33]. In contrast, multi-policy approaches maintain a separate policy for each \mathbf{w} [17, 19]. Parameter sharing improves the sample efficiency of single-policy methods and enhances generalisation to new objective preferences [1, 32]. However, multi-policy methods often require maintaining various suboptimal policies to adequately cover the preference space, which can be highly memory-intensive [17]. Existing MORL/D approaches share common limitations. As the number of objectives increases or the granularity of the weight space becomes finer, the total preference space grows exponentially. This exponential growth renders prefixed scalarisation weights inefficient and degrades the quality of the approximated PF.

To address the limitations of MORL/D, particularly in the multi-policy paradigm, this work introduces an Upper Confidence Bound (UCB) acquisition function to identify promising scalarisation vectors $\mathbf{w} \in \mathbf{W}$ at different learning stages, guided by the Pareto front quality metric of *hypervolume* (HV) [9]. Extending [1], we use a weight-conditioned Actor-Critic network $\pi_\theta(s, w)$ trained with Proximal Policy Optimisation (PPO) [26] for C iterations, instead of a Q-network $Q_\theta(s, a, \mathbf{w})$.

We adopt a multi-policy approach, each $\pi_\theta(s, w)$ specializes in a sub-space of \mathbf{W} , enabling time-efficient and parallel optimisation with a compact parameter size. Inspired by [17], we frame HV maximisation as a surrogate-assisted opti-

misation problem, where a data-driven surrogate predicts changes in objective values to select weight vectors that improve HV. We extend [17] by incorporating prediction uncertainty into the UCB acquisition function, balancing exploration and exploitation. Further, we replace large Pareto archives with scalarisation-vector-conditioned policies $\pi_\theta(s, w)$, reducing memory overhead while ensuring efficient PF coverage.

In summary, two contributions distinguish our work from previous systems of MORL/D:

1. A two-layer decomposition enables policies to specialize in different subspaces of the scalarisation vector space, with conditioning refining sub-problems within each subspace. This facilitates generalisation across local scalarisation neighbourhoods, enhancing HV metrics [9].
2. The use of a UCB acquisition function for selecting from a finite set of evenly distributed scalarisation vectors to balance exploration and exploitation.

The rest of the paper is structured as follows. Background knowledge is introduced in Section 2. Prior work on gradient-based MORL is reviewed in Section 3. The proposed method is described in Section 4. Experiment configurations on six multi-objective benchmark problems are outlined in Section 5, and SOTA results on these problems are discussed in Section 6. Finally, we conclude in Section 7 with directions for future research.

2 Preliminaries

2.1 Policy-gradient Reinforcement Learning

Policy gradient methods are a class of RL algorithms that optimise policies directly by computing gradients of an objective function with respect to the policy parameters. The objective function $J(\theta)$ for a parameterised policy $\pi_\theta(a|s)$ can be expressed as the expected return, the policy gradient is then computed as:

$$\nabla_\theta J(\theta) = \mathbb{E}_{\tau \sim \pi_\theta} \left[\nabla_\theta \log \pi_\theta(a_t|s_t) \hat{A}_t \right],$$

where \hat{A}_t is the advantage function [29], which measures the relative value of taking action a_t in state s_t .

Proximal Policy Optimisation (PPO) [26] improves upon traditional policy gradient methods by introducing a surrogate objective that prevents overly large policy updates, ensuring stable learning. The PPO objective is defined as:

$$L^{\text{PPO}}(\theta) = \mathbb{E}_t \left[\min \left(r_t(\theta) \hat{A}_t, \text{clip}(r_t(\theta), 1 - \epsilon, 1 + \epsilon) \hat{A}_t \right) \right],$$

where $r_t(\theta) = \frac{\pi_\theta(a_t|s_t)}{\pi_{\theta_{\text{old}}}(a_t|s_t)}$ is the probability ratio, and ϵ is a hyperparameter controlling the clipping range. By clipping the ratio, PPO limits deviations from the current policy, leading to improved stability and sample efficiency. These advancements have made PPO one of the most widely used algorithms for continuous and high-dimensional control tasks.

2.2 Definition of MORL

A multi-objective sequential decision making problem can be formulated as a multi-objective Markov Decision Process (MOMDP) defined by the tuple $\langle S, A, P, \gamma, \rho_0, \mathbf{r} \rangle$ with state space S , action space A , state transition probability $P : S \times A \times S \rightarrow [0, 1]$, discount factor γ , initial state distribution ρ_0 , and vector-valued reward function $\mathbf{r} : S \times A \rightarrow \mathbb{R}^m$ specifying one-step reward for each of the m objectives. A decision-making policy $\pi : S \rightarrow A$ maps states into actions, for which a vector-valued value function \mathbf{V}^π is defined as:

$$\mathbf{V}^\pi = \mathbb{E} \left[\sum_{t=0}^H \gamma^t \mathbf{r}(s_t, a_t) \mid s_0 \sim \rho_0, a_t \sim \pi \right] \quad (1)$$

where H is the length of the horizon.

The space of utility functions for a MORL problem is typically populated by linear and non-linear functions of \mathbf{V}^π . In this work we focus on the highly prevalent case where the utility functions are linear, taking the form of $u(\mathbf{V}^\pi) = \mathbf{w}^\top \mathbf{V}^\pi$, where scalarisation vector \mathbf{w} provides the parametrisation of the utility function. Each element of $\mathbf{w} \in \mathbb{R}^m$ specifies the relative importance (preference) of each objective. The space of linear scalarisation vectors \mathbf{W} is the m -dimensional simplex: $\sum_i w_i = 1, w_i \geq 0, i = 1, \dots, m$. For any given \mathbf{w} , the original MOMDP is reduced to a single-objective MDP.

2.3 Convex Coverage Set (CCS)

Linear utility functions enable MORL to generate a Convex Coverage Set (CCS) of policies, which is a subset of all possible policies such as there exists a policy π in the set that is optimal with respect to any linear scalarisation vector \mathbf{w} :

$$CCS \equiv \left\{ \mathbf{V}^\pi \in \Pi \mid \exists \mathbf{w} \text{ s.t. } \forall \mathbf{V}^{\pi'} \in \Pi, \mathbf{V}^\pi \mathbf{w}^\top \geq \mathbf{V}^{\pi'} \mathbf{w}^\top \right\} \quad (2)$$

CCS is a subset of PF defined for monotonically increasing utility functions [13]. At the same time, the underlying linear utility function space (space of scalarisation vectors \mathbf{W}) is easier to search than that of arbitrarily-composed utility functions, where one has to search for the overall function composition out of primitive function elements. The sampling of scalarisation vectors $\mathbf{w} \in \mathbf{W}$ impacts the quality of the CCS in approximating the PF. Vectors can be sampled uniformly at random [1, 6, 18, 20, 32], selected from a predefined set [4, 17], or adapted via search methods [2, 19]. A small step size is required for high-quality CCS when using predefined vectors [25].

3 Related Work

A thorough review of MORL can be found in [13]. In this section we highlight previous work that employs gradient-based with a linear utility function, which can be classified into three main categories.

In the first category, a single policy [21, 27] is trained using a linear scalarisation function when the user’s preferences are known in advance. For example, [4] proposes a novel algorithm that trains a single universal network to cover the entire preference space. The approach utilises preferences to guide network parameter updates and employs a novel parallelisation strategy to enhance sample efficiency.

In the second category, when user preferences are unknown or difficult to define, a CCS of policies is computed by training multiple independent policies with different scalarisation vectors to capture various trade-offs between objectives [16, 18]. The work of [19] proposes Optimistic Linear Support, a method that adaptively selects the weights of the linear utility function via the concepts of corner weights and estimated improvement to prioritise those corner weights. [2] introduces a sample-efficient MORL algorithm that leverages Generalised Policy Improvement (GPI) to prioritize training on specific preference weights. By focusing on corner weights with higher GPI priority, the method iteratively learns a set of policies whose value vectors approximate the CCS. The work of [17] maintains a Pareto archive of policies by focusing on those scalarisation vectors that are expected to improve the hypervolume and sparsity metrics of the resulting PF the most.

The third category of methods maintains a single policy conditioned on the scalarisation vector \mathbf{w} . Several works aim to train such policies for few-shot adaptation to varying objective preferences [1, 6, 32]. [1] uses a scalarisation-vector-conditioned Q-network trained on randomly sampled \mathbf{w} to solve single-objective RL sub-problems, generalizing across changing objective preferences. Similarly, [32] applies Envelope Q-learning to train a Q-network for few-shot adaptation to new scalarisation vectors. [6] proposes a PPO-based meta-policy trained collaboratively with data from policies specialized to sampled scalarisation vectors, enabling PF construction via few-shot fine-tuning. Other works directly approximate the CCS using a single scalarisation-conditioned policy. [18] enhances Soft Actor Critic with an entropy term to train scalarisation-conditioned policies and Q-networks, while [4] extends [17] by incorporating scalarisation-conditioning, replacing a Pareto archive with a single policy.

Existing methods demonstrate notable strengths but face challenges in scalability, computational cost, generalisation, and performance consistency. By comparing these methods to UCB-MOPPO, we highlight its key advantages: 1) Preference space decomposition enables scalable, efficient parallel training with compact policies. 2) The UCB-driven surrogate model selects scalarisation vectors to maximise hypervolume. 3) Integration of PPO enhances stability and sample efficiency. 4) The weight-conditioned network reduces memory overhead.

4 Methods

The key technical characteristics of the proposed MORL/D algorithm, named UCB-MOPPO, are as follows:

MORL/D as scalar RL sub-problems. The overall scalarisation weight \mathbf{W} is divided into K sub-spaces as shown in Figure 2a. A separate policy π_k is trained for each sub-problem, conditioned on scalarisation vectors sampled from the corresponding sub-space $\mathbf{W}_k \subset \mathbf{W}$, where $k = 1, \dots, K$.

Scalarisation-vector-conditioned Actor-Critic. A policy network π_θ , a value network v_ϕ^π , and a scalarisation vector \mathbf{w} ($\sum_i w_i = 1, w_i \geq 0, i = 1, \dots, m$) are used to maximise the weighted-sum of expected rewards, denoted as $J(\theta, \phi, \mathbf{w}) = \sum_{i=1}^m w_i V_i^\pi$. Following [1, 18, 32], both $\pi_\theta(s, \mathbf{w})$ and $v_\phi^\pi(s, \mathbf{w})$ are conditioned on \mathbf{w} . This enables a single policy to express different trade-off between objectives by generalising across a neighbourhood of scalarisation vectors.

Surrogate-assisted maximisation of CCS hypervolume. An acquisition function based on UCB [28] is used to select scalarisation vectors for training from each sub-space \mathbf{W}_k . During training, the chosen scalarisation vectors are those expected to maximise the hypervolume of the resulting CCS most effectively.

4.1 Two-layer decomposition of MORL problem into scalar RL sub-problems

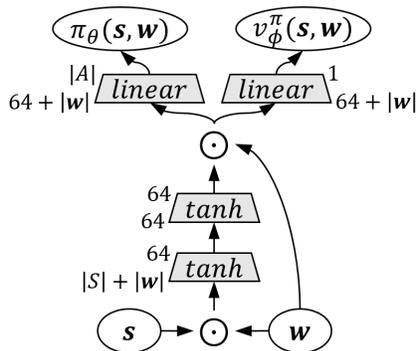
At the first layer of problem decomposition, a set of K evenly distributed scalarisation vectors, named as *pivots*, are defined within the scalarisation vector space \mathbf{W} . These pivot vectors effectively divide \mathbf{W} into K different sub-spaces as shown in Figure 2a. The system allows to independently train multiple pivot policies with different random seeds in order to improve the density of the resulting CCS. At the second layer of problem decomposition, for each sub-space \mathbf{W}_k , $k = 1, \dots, K$, a number of M evenly distributed scalarisation vectors are defined in turn. Therefore the overall MORL problem is decomposed into a total of $K * M$ scalar RL sub-problems, defined as follows for a fixed \mathbf{w} :

$$\pi(\cdot, \mathbf{w}) = \operatorname{argmax}_{\pi'(\cdot, \mathbf{w})} \mathbb{E} \left[\sum_{t=1}^H \gamma^t \mathbf{w}^\top \mathbf{r}(s_t, \alpha_t) \right] \quad (3)$$

A separate policy π_k (i.e., *pivot* policy) is independently trained for each sub-space $\mathbf{W}_k \subset \mathbf{W}$ by conditioning on the corresponding scalarisation vectors. As an example, for a bi-objective problem with $\mathbf{w} = [w_1, w_2]$ in Figure 2a, for two sub-spaces $\mathbf{W}_i, \mathbf{W}_j$ with $i > j$, $(w_1^{i,n} > w_1^{j,n}) \wedge (w_2^{i,n} < w_2^{j,n})$, $n = 1, \dots, M$. The solutions of the $K * M$ sub-problems compose a CCS.

4.2 Scalarisation-vector-conditioned Actor-Critic

The neural network(NN) architecture of the Actor-Critic is illustrated in Figure 1. The state vector \mathbf{s} concatenated with the scalarisation vector \mathbf{w} form the input layer. \mathbf{w} is also concatenated with the last shared layer through a residual connection [14]. Residual connections aim at improving the sensitivity of NN output to changes in \mathbf{w} in a similar vein with reward-conditioned policies in [15].

Fig. 1: Scalarisation vector \mathbf{w} conditioned actor-critic network.

4.3 Surrogate-assisted maximisation of CCS hypervolume

Overview The learning algorithm proceeds in stages of C iterations each. At every stage, K pivot policies π_k , $k = 1, \dots, K$ are trained in parallel using PPO, each conditioned on a subset of N scalarisation vectors out of M in total that are defined for each corresponding sub-space \mathbf{W}_k , $k = 1, \dots, K$, with $N \ll M$. The selection of N scalarisation vectors to train on at each stage is performed via surrogate-assisted maximisation of the hypervolume of the CCS. First, a data-driven uncertainty-aware surrogate model is built in the scalarisation vector space to predict the expected change in each objective after training π_k conditioned on said scalarisation vectors for C iterations. Second, an acquisition function is defined as the UCB of the CCS hypervolume that is expected by including a scalarisation vector to the policy’s conditioning set at each training stage. Maximising the acquisition function selects the scalarisation vectors that are expected to improve CCS hypervolume the most. Pseudo code can be provided in supplementary material, section 1, algorithm 3.

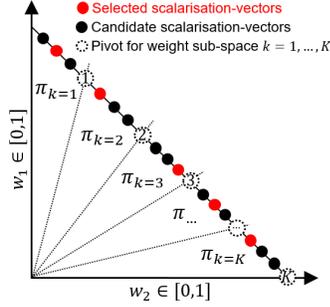
Warm-up phase The algorithm starts with a warm-up stage performed using Algorithm 1. For an m -objective problem, a set of K evenly distributed pivot vectors $\{\mathbf{w}_i\}_{i=1}^K$ are generated, where $\sum_i w_i = 1, w_i \geq 0, i = 1, \dots, m$. Accordingly, a set of K pivot policies, each conditioned on a separate pivot vector, is trained using Algorithm 1 for a number of epochs.

Surrogate model Let $\pi_{k,z+1}$ be the policy that results from the k^{th} policy $\pi_{k,z}$ during the z^{th} training stage of C iterations. Let $\Delta V_{j,\mathbf{w}}^{\pi_k,(z \rightarrow z+1)} = V_{j,\mathbf{w}}^{\pi_{k,z+1}} - V_{j,\mathbf{w}}^{\pi_{k,z}}$ be the change in the value of the j^{th} objective for the k^{th} policy conditioned on scalarisation vector \mathbf{w} trained with PPO for C iterations. For each pivot policy π_k , $k = 1, \dots, K$, and for each objective $j = 1, \dots, m$, a separate dataset

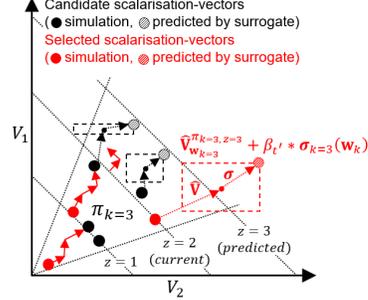
$D_{surrogate}^{k,j}$ is created using policy's π_k evaluation data of objective j that are collected from the simulation environment during a number of consecutive training stages Z , as follows:

$$D_{surrogate}^{k,j} = \left\{ \left(\mathbf{w}, (V_{j,\mathbf{w}}^{\pi_{k,z+1}} - V_{j,\mathbf{w}}^{\pi_{k,z}}) \right) \right\}_{z=1}^Z \quad (4)$$

A surrogate model $f_{bagging}^{k,j} : \mathbb{R}^m \rightarrow \mathbb{R}$ is trained on $D_{surrogate}^{k,j}$ to predict $\Delta V_{j,\mathbf{w}}^{\pi_{k,(z \rightarrow z+1)}}$ as a function of the scalarisation vector \mathbf{w} , for m number of objectives. The training is incremental while additional tuples are appended to $D_{surrogate}$ from consecutive training stages. The surrogate model takes the form of Bagging [5] of linear models $f_{\psi}(\mathbf{w}) = \sum_{i=1}^m \psi_i w_i + \psi_0$ with elastic net regularisation [37]. Bagging trains independently B linear models $\{f_{\psi}^b\}_{b=1}^B$ on B bootstrap samples of the original training data and predicts using their average, that is $f_{bagging}^{k,j}(\mathbf{w}) = \frac{1}{B} \sum_{i=1}^B f_{\psi}^b(\mathbf{w})$. An estimate of the epistemic uncertainty of the prediction can be computed using the variance of the component model predictions [12], that is $\sigma_{k,j}^2(\mathbf{w}) = \frac{1}{B} \sum_{i=1}^B (f_{\psi}^b(\mathbf{w}) - f_{bagging}^{k,j}(\mathbf{w}))^2$.



(a) Decomposition of MORL problem into scalar sub-problems.



(b) Surrogate-assisted maximisation of CCS hypervolume.

Fig. 2: Overview of the proposed approach for UCB-driven utility function search for MORL/D.

UCB acquisition function maximisation For each policy k and each objective j a surrogate model can predict the expected objective values by conditioning a policy on scalarisation vector \mathbf{w} during a training stage z as follows:

$$\hat{V}_{j,\mathbf{w}}^{\pi_{k,z+1}} = V_{j,\mathbf{w}}^{\pi_{k,z}} + f_{bagging}^{k,j}(\mathbf{w}) \quad (5)$$

with a corresponding vectorised prediction, denoted as $\hat{\mathbf{V}}_{\mathbf{w}}^{\pi_{k,z+1}} = \mathbf{V}_{\mathbf{w}}^{\pi_{k,z}} + \mathbf{f}_{bagging}^k(\mathbf{w})$ over j objectives and the accompanying vectorised uncertainty estimate $\sigma_k^2(\mathbf{w})$.

At the beginning of each training stage z the algorithm needs to select those N scalarisation vectors out M evenly distributed vectors in each sub-space \mathbf{W}_k , $k = 1, \dots, K$ that are predicted via the surrogate model to improve the hypervolume of the resulting $\text{HV}(\text{CCS})$ the most. $\text{HV}(\cdot)$ is a function that computes hypervolume as in [9]. The scalarisation vectors are selected one at a time without replacement in a sequence of N invocations of a process that maximises a UCB acquisition function defined on $\text{HV}(\text{CCS})$ as:

$$\begin{aligned} \mathbf{w}_k^* &= \underset{\mathbf{w}_{k,j} \in \{\mathbf{w}_{k,j}\}_{j=1}^M}{\text{argmax}} \text{HV}(\text{Pareto}(\mathbf{V}_{k,j})), \\ \mathbf{V}_{k,j} &= \hat{\mathbf{V}}_{\mathbf{w}_{k,j}}^{\pi_{k,z+1}} + \beta_{t'} \cdot \boldsymbol{\sigma}_k(\mathbf{w}_{k,j}), \quad k = 1, \dots, K. \end{aligned} \quad (6)$$

where $\text{Pareto}(\cdot)$ is a function computing the CCS from a set of objective vectors. The dynamic parameter of current training step, i.e., $\beta_{t'} = \sqrt{\frac{\log(2t')}{t'}}$ ensures that scalarisation vector with higher uncertainty (standard deviation) are explored early, but the focus shifts towards exploiting the ones with high mean rewards as confidence improves, assuming maximisation of objectives. Once selected, a scalarisation vector \mathbf{w}_k^* is removed from the candidate set $\{\mathbf{w}_{k,j}\}_{j=1}^M$ for the current training stage. The process of selecting scalarisation vectors that are expected to maximise the hypervolume of the resulting CCS is illustrated in Figure 2b.

4.4 Baseline Methods

Algorithm 1 Fixed-MOPPO

- 1: **Input:** State s_t , weights \mathbf{w}
 - 2: **Initialize:** K Actor-Critic networks π_k, v^{π_k} ; scalarisation spaces \mathbf{W}_k ; buffer \mathcal{E} of size D .
 - 3: **for** $k = 1$ **to** K **do**
 - 4: **for** $t = 1$ **to** D **do**
 - 5: $\mathbf{w}_{\text{pivot}} \leftarrow \text{get_pivot_weight}(\mathbf{W}_k)$
 - 6: $a_t \sim \pi_k(s_t, \mathbf{w}_{\text{pivot}})$, $s_{t+1}, \mathbf{r}_t \leftarrow \text{simulator}(s_t, a_t)$
 - 7: $\mathcal{E} \leftarrow \mathcal{E} \cup \langle s_t, a_t, \mathbf{w}_{\text{pivot}}, \mathbf{r}_t, s_{t+1} \rangle$, $s_t \leftarrow s_{t+1}$
 - 8: **end for**
 - 9: Sample $\langle s_t, a_t, \mathbf{w}_{\text{pivot}}, \mathbf{r}_t, s_{t+1} \rangle \sim \mathcal{E}$
 - 10: $\theta \leftarrow \theta + \eta \nabla_{\theta} \log \pi_{\theta}(s_t, a_t; \mathbf{w}_{\text{pivot}}) A^{\pi}(s_t, a_t; \mathbf{w}_{\text{pivot}})$
 - 11: $\phi \leftarrow \phi + \eta \|V^{\pi_k}(s_t; \mathbf{w}_{\text{pivot}}) - (\mathbf{r}_t + \gamma V^{\pi_k}(s_{t+1}; \mathbf{w}_{\text{pivot}}))\|^2$
 - 12: $\mathcal{E} \leftarrow \emptyset$
 - 13: **end for**
-

Seven baseline methods are introduced for comparison, including three proposed **MOPPO** baseline methods for ablation studies and four SOTA MORL

methods: 1) **PG-MORL** [17], 2) **PD-MORL** [4], 3) **CAPQL** [18], and 4) **GPI-LS** [2]. These baselines are used to evaluate the benefits introduced by PPO and the UCB-driven search of the scalarisation vector space. A brief description of the proposed **MOPPO baselines** is as follows:

Fixed-MOPPO in Algorithm 1, which trains K policies, each conditioned on a fixed scalarisation vector $\mathbf{w}_{pivot} \in \mathbf{W}_k$ corresponding to subspace k (see Figure 2a).

Random-MOPPO (see supplementary material, section 1, algorithm 2) which is similar to Fixed-MOPPO, K policies are trained, with each π_k conditioned on scalarisation vectors uniformly sampled from \mathbf{W}_k . These vectors are periodically re-sampled to ensure diversity during training.

Mean-MOPPO which is identical to UCB-MOPPO, except that the acquisition function ignores uncertainty by setting $\beta_{t'} := 0, \forall t'$ in Equation 6.

5 Experiment Setup

We evaluate the performance of all seven baselines on six continuous control multi-objective RL problems: **Swimmer-V2**, **Halfcheetah-V2**, **Walker2d-V2**, **Ant-V2**, **Hopper-V2**, and **Hopper-V3** (the only problem with three objectives). Detailed descriptions of the objectives, as well as the state and action spaces, are provided in the supplementary material, section 2. Each method is evaluated on three different random seeds for each problem, consistent with [17]. For a fair comparison, we use raw reward outputs and assess the results using the *hypervolume* (HV) and *Expected Utility* (EU) metrics [2].

Table 1: Decomposition setup for two-objective and three-objective problems.

| Parameter | Objective | UCB | Mean | Random | Fixed |
|---------------------------|-------------|------|------|--------|-------|
| K (pivot vectors) | 2-objective | 10 | 10 | 10 | 10 |
| | 3-objective | 36 | 36 | 36 | 36 |
| N (sub-space selection) | 2-objective | 10 | 10 | 10 | 1 |
| | 3-objective | 10 | 10 | 10 | 1 |
| M (sub-space vectors) | 2-objective | 100 | 100 | 10 | 10 |
| | 3-objective | 117 | 117 | 36 | 36 |
| Step-size (layer 1) | 2-objective | 0.1 | 0.1 | 0.1 | 0.1 |
| | 3-objective | 0.1 | 0.1 | 0.1 | 0.1 |
| Step-size (layer 2) | 2-objective | 0.01 | 0.01 | 0.1 | 0.1 |
| | 3-objective | 0.05 | 0.05 | 0.1 | 0.1 |

For UCB-MOPPO and the proposed MOPPO baselines, hyperparameters for policy neural network initialisation and PPO are aligned with the values recommended in Stable-Baselines3 [22]. Detailed hyperparameters are provided

in the supplementary material, section 3. Proposed problem decomposition is based on scalarisation vectors generated via discretisation of space \mathbf{W} with a fixed step size. The overall parameters of problem decomposition for 2- and 3-objective problems are summarised in Table 1. SOTA baselines are reproduced using the MORL-Baseline library [9].

6 Results Analysis

6.1 Quality of Pareto Front

Pareto front(PF) are evaluated in **HV** and **EU**, which are reported in Table 2. **PF coverage** and **HV convergence speed** are visualized in Figure 3.

Table 2: Evaluation of HV and EU metrics for continuous MORL tasks over three independent runs. Due to table size limitations, environment and method names are **abbreviated**, and only **mean** values from three runs are reported. The full detailed table is provided in supplementary material, section 4.

| | Metric | UCB | Mean | Rand- | Fix- | PG- | PD- | CAP- | GPI- |
|-------------|------------------|-------------|-------------|-------------|------|------|------|-------------|------|
| Swimmer | HV (10^4) | 5.60 | 4.72 | 4.56 | 4.45 | 1.67 | 1.77 | 2.40 | 4.75 |
| | EU (10^2) | 2.18 | 2.16 | 2.08 | 2.11 | 1.23 | 1.24 | 1.79 | 2.14 |
| Halfcheetah | HV (10^7) | 1.78 | 1.60 | 1.21 | 1.12 | 0.58 | 0.62 | 2.20 | 2.16 |
| | EU (10^3) | 3.88 | 3.58 | 3.55 | 3.44 | 2.30 | 2.42 | 4.47 | 4.30 |
| Walker2d | HV (10^7) | 1.40 | 1.29 | 1.15 | 1.13 | 0.44 | 0.56 | 0.18 | 1.22 |
| | EU (10^3) | 3.54 | 3.42 | 3.34 | 3.30 | 1.98 | 2.24 | 1.50 | 3.40 |
| Ant | HV (10^7) | 1.07 | 0.92 | 0.65 | 0.60 | 0.61 | 0.66 | 0.45 | 0.81 |
| | EU (10^3) | 2.96 | 2.61 | 2.35 | 2.21 | 2.28 | 2.41 | 2.03 | 2.84 |
| Hopper2d | HV (10^7) | 0.84 | 0.81 | 0.79 | 0.76 | 0.23 | 0.25 | 0.25 | 0.80 |
| | EU (10^3) | 2.77 | 2.84 | 2.77 | 2.71 | 1.48 | 1.51 | 1.46 | 2.75 |
| Hopper3d | HV (10^{10}) | 3.62 | 2.83 | 2.90 | 2.64 | 0.63 | 0.11 | 0.31 | 0.65 |
| | EU (10^3) | 3.20 | 3.19 | 3.28 | 2.90 | 1.70 | 1.74 | 1.52 | 2.14 |

UCB-MOPPO delivers consistent performance and outperforms all baselines except in Halfcheetah-V2. Notably, **PGMORL** produces a dense PF by accumulating a large number of policies in the archive, yet it achieves significantly lower HV and EU compared to the proposed MOPPO baselines. **PDMORL** achieves slightly higher HV and EU than PGMORL while using only a single network; however, it struggles with three-objective problems. **CAPQL** achieves the best performance in Halfcheetah-V2 but struggles significantly on Walker2d-V2 and Hopper-V3 due to learning only a limited dynamics of those problems. **GPI-LS** delivers the most competitive results and requires the longest GPU time for training. it performs better than the proposed MOPPO baselines in Halfcheetah-V2; however, it fails to explore a spread CCS in the Ant-V2 and Hopper-2d/-3d environments.

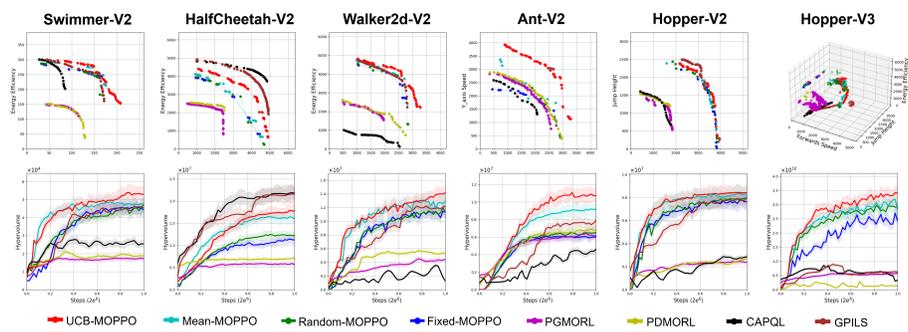


Fig. 3: Comparison of PF and HV evaluation for six MORL tasks. The top row shows PF plots, and the bottom row shows HV comparisons. Shaded areas in HV plots represent the standard deviation over three seeds.

The consistent performance of the proposed method can be partially attributed to the integration of PPO, which provides enhanced stability and sample efficiency by employing trust region enforcement compared to vanilla policy gradient methods. As shown in Figure 3, UCB-MOPPO demonstrates superior convergence across most tasks (except for HalfCheetah-V2). The proposed method achieves rapid early-stage convergence by efficiently identifying scalarisation vectors that improve hypervolume. Additionally, the MOPPO baselines exhibit stable convergence without overfitting and dropping, ensuring robust performance across diverse multi-objective tasks.

6.2 Analysis of Ablation Experiments

Fixed-MOPPO is the simplest baseline, conditioning policies on pivot scalarisation vectors. It provides quick insights into a new MORL problem but results in a sparse PF.

Random-MOPPO achieves higher hypervolume and a denser PF than Fixed-MOPPO by randomly selecting scalarisation vectors under sub-space during training. However, it is inefficient, failing to explore promising regions effectively.

Mean-MOPPO improves search efficiency but yields a less dense PF than UCB-driven methods. The less coverage of PF can be attributed to the greedy selection of scalarisation vectors based solely on the mean predicted hypervolume improvement from the surrogates, without awareness of the uncertainty.

6.3 Comparison of the Policy Archive

UCB-MOPPO constructs a satisfactory PF with a consistently low number of policies compared to the baselines across all tested environments. The total number of policies after convergence is shown in Table 3. Some challenging problems

Table 3: Number of archived policies after the hypervolume converges.

| Environment | PGMORL | CAPQL | GIPLS | UCB-MOPPO |
|----------------|--------|-------|-------|-----------|
| Swimmer-v2 | 168 | 33 | 41 | 30 |
| Halfcheetah-v2 | 285 | 35 | 83 | 30 |
| Walker2d-v2 | 412 | 51 | 52 | 30 |
| Ant-v2 | 64 | 26 | 27 | 30 |
| Hopper-v2 | 206 | 30 | 45 | 30 |
| Hopper-v3 | 4023 | 82 | 89 | 108 |

result in a large number of policies in the archive (e.g., Walker2d-V2 and Hopper-V3).

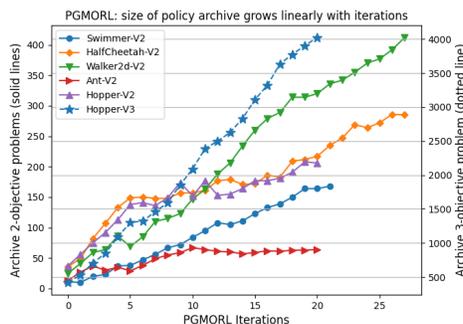


Fig. 4: Growth in policy archive size in PGMORL.

A key advantage of all the MOPPO methods is that they maintain a small, fixed set of policies, unlike PGMORL, where the policy archive size grows linearly, as depicted in Figure 4. Therefore, an efficient search in UCB-MOPPO is complemented by a memory-efficient implementation, which can be an important consideration in production environments with tight resource constraints.

6.4 Interpolating in scalarisation vectors spaces

An analysis was carried out to assess the ability of policies trained via UCB-MOPPO to interpolate in a more fine-grained discretisation of vector space \mathbf{W} , discretised using smaller step-sizes than those considered during training.

Each of the K policies, corresponding to the K sub-spaces, is evaluated on $N = \{10, 20, \dots, 50\} \cup \{100, 150, \dots, 500\}$ scalarisation vectors. The results are shown in Figure 5, where we present the **hypervolume** and **sparsity** curves. The plots reveal a general trend: as N increases, the hypervolume improves, indicating a better approximation of the PF. Simultaneously, sparsity decreases,

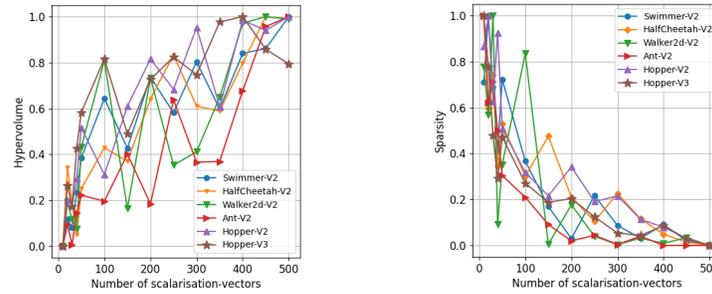


Fig. 5: Hypervolume and sparsity achieved by UCB-MOPPO improve with more scalarisation vectors conditioning the policies.

signifying a more uniform PF coverage. This suggests that increasing the granularity of scalarisation vectors leads to a more comprehensive and higher-quality representation of the PF, all without requiring additional training. This advantage is not shared by PGMORL or other MORL/D approaches, which often need further training to achieve similar results. Thus, the proposed method offers a distinct benefit in improving PF quality with minimal computational overhead.

7 Conclusion

This paper presents a method for efficiently searching scalarisation vectors that maximize the quality of the CCS. The key findings are: (1) the proposed method outperforms competitive baselines in terms of CCS hypervolume and Expected Utility in most cases, (2) it requires maintaining a minimal number of policies to produce high-quality CCS, making it well-suited for resource-constrained environments, and (3) the CCS hypervolume and sparsity metrics improve as the scalarisation vector step-size decreases, demonstrating effective generalisation across scalarisation vector neighbourhoods. For future work, we plan to explore acquisition functions that leverage additional Pareto front quality indicators [7] and search algorithms designed for non-linear utility function spaces [23].

References

1. Abels, A., Roijers, D., Lenaerts, T., Nowé, A., Steckelmacher, D.: Dynamic weights in multi-objective deep reinforcement learning. In: International conference on machine learning. pp. 11–20. PMLR (2019)
2. Alegre, L.N., Bazzan, A.L., Roijers, D.M., Nowé, A., da Silva, B.C.: Sample-efficient multi-objective learning via generalized policy improvement prioritization. arXiv preprint arXiv:2301.07784 (2023)
3. Basaklar, F., Sen, T.: Multi-objective reinforcement learning: Methods and applications. *Expert Systems with Applications* **202**, 117016 (2022)
4. Basaklar, T., Gumussoy, S., Ogras, U.Y.: Pd-morl: Preference-driven multi-objective reinforcement learning algorithm. arXiv preprint arXiv:2208.07914 (2022)

5. Breiman, L.: Bagging predictors. *Machine Learning* **24**, 123–140 (1996), <https://api.semanticscholar.org/CorpusID:47328136>
6. Chen, X., Ghadirzadeh, A., Björkman, M., Jensfelt, P.: Meta-learning for multi-objective reinforcement learning. In: 2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS). pp. 977–983 (2019). <https://doi.org/10.1109/IROS40897.2019.8968092>
7. Falcón-Cardona, J.G., Coello, C.A.C.: Indicator-based multi-objective evolutionary algorithms: A comprehensive survey. *ACM Computing Surveys (CSUR)* **53**(2), 1–35 (2020)
8. Feinberg, E.A., Shwartz, A.: Constrained markov decision models with weighted discounted rewards. *Math. Oper. Res.* **20**, 302–320 (1995), <https://api.semanticscholar.org/CorpusID:16510079>
9. Felten, F., Alegre, L.N., Nowé, A., Bazzan, A.L.C., Talbi, E.G., Danoy, G., da Silva, B.C.: A toolkit for reliable benchmarking and research in multi-objective reinforcement learning. In: *Neural Information Processing Systems* (2023), <https://api.semanticscholar.org/CorpusID:268042105>
10. Felten, F., Talbi, E.G., Danoy, G.: Multi-objective reinforcement learning based on decomposition: A taxonomy and framework. *Journal of Artificial Intelligence Research* **79**, 679–723 (2024)
11. Gao, X., Yi, W., Agapitos, A., Wang, H., Liu, Y.: Coverage and capacity optimization in star-riss assisted networks: A machine learning approach. In: 2023 IEEE Wireless Communications and Networking Conference (WCNC). pp. 1–6. IEEE (2023)
12. Gawlikowski, J., Tassi, C.R.N., Ali, M., Lee, J., Humt, M., Feng, J., Kruspe, A., Triebel, R., Jung, P., Roscher, R., et al.: A survey of uncertainty in deep neural networks. *Artificial Intelligence Review* **56**(Suppl 1), 1513–1589 (2023)
13. Hayes, C.F., Rădulescu, R., Bargiacchi, E., Källström, J., Macfarlane, M., Raymond, M., Verstraeten, T., Zintgraf, L.M., Dazeley, R., Heintz, F., et al.: A practical guide to multi-objective reinforcement learning and planning. *Autonomous Agents and Multi-Agent Systems* **36**(1), 26 (2022)
14. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 770–778 (2016)
15. Kumar, A., Peng, X.B., Levine, S.: Reward-conditioned policies. *arXiv preprint arXiv:1912.13465* (2019)
16. Li, K., Zhang, T., Wang, R.: Deep reinforcement learning for multiobjective optimization. *IEEE Transactions on Cybernetics* **51**, 3103–3114 (2019), <https://api.semanticscholar.org/CorpusID:174802898>
17. Liu, F., Qian, C.: Prediction guided meta-learning for multi-objective reinforcement learning. 2021 IEEE Congress on Evolutionary Computation (CEC) pp. 2171–2178 (2021), <https://api.semanticscholar.org/CorpusID:236981140>
18. Lu, H., Herman, D., Yu, Y.: Multi-objective reinforcement learning: Convexity, stationarity and pareto optimality. In: *International Conference on Learning Representations* (2023), <https://api.semanticscholar.org/CorpusID:260108193>
19. Mossalam, H., Assael, Y.M., Roijers, D.M., Whiteson, S.: Multi-objective deep reinforcement learning (2016)
20. Natarajan, S., Tadepalli, P.: Dynamic preferences in multi-criteria reinforcement learning. *Proceedings of the 22nd international conference on Machine learning* (2005), <https://api.semanticscholar.org/CorpusID:15370837>
21. Pan, A., Xu, W., Wang, L., Ren, H.: Additional planning with multiple objectives for reinforcement learning. *Knowledge-Based Systems* **193**, 105392 (2020)

22. Raffin, A., Hill, A., Gleave, A., Kanervisto, A., Ernestus, M., Dormann, N.: Stable-baselines3: Reliable reinforcement learning implementations. *Journal of Machine Learning Research* **22**(268), 1–8 (2021)
23. Reymond, M., Hayes, C.F., Steckelmacher, D., Roijers, D.M., Nowé, A.: Actor-critic multi-objective reinforcement learning for non-linear utility functions. *Autonomous Agents and Multi-Agent Systems* **37**(2), 23 (2023)
24. Roijers, D.M., Vamplew, P., Whiteson, S., Dazeley, R.: A survey of multi-objective sequential decision-making. *Journal of Artificial Intelligence Research* **48**, 67–113 (2013)
25. Roijers, D.M., Whiteson, S., Oliehoek, F.A.: Point-based planning for multi-objective pomdps. In: *International Joint Conference on Artificial Intelligence* (2015), <https://api.semanticscholar.org/CorpusID:9116349>
26. Schulman, J., Wolski, F., Dhariwal, P., Radford, A., Klimov, O.: Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347* (2017)
27. Siddique, U., Weng, P., Zimmer, M.: Learning fair policies in multi-objective (deep) reinforcement learning with average and discounted rewards. In: *International Conference on Machine Learning*. pp. 8905–8915. PMLR (2020)
28. Srinivas, N., Krause, A., Kakade, S.M., Seeger, M.: Gaussian process optimization in the bandit setting: No regret and experimental design. *arXiv preprint arXiv:0912.3995* (2009)
29. Sutton, R.S., McAllester, D., Singh, S., Mansour, Y.: Policy gradient methods for reinforcement learning with function approximation. In: *Proceedings of the 12th International Conference on Neural Information Processing Systems*. pp. 1057–1063 (2000)
30. Vamplew, P., Dazeley, R., Foale, C., Firmin, S., Mummery, S.: Empirical evaluation methods for multiobjective reinforcement learning algorithms. *Machine learning* **84**(1), 51–80 (2011)
31. Xu, Y., Zhou, B., Sun, X.: Reinforcement learning for multi-objective optimization: A survey. *Artificial Intelligence Review* **53**(6), 4123–4156 (2020)
32. Yang, R., Sun, X., Narasimhan, K.: A generalized algorithm for multi-objective reinforcement learning and policy adaptation. In: *Neural Information Processing Systems* (2019), <https://api.semanticscholar.org/CorpusID:201319046>
33. Yang, Y., Roijers, D.M., Foster, C.: Generalized multiple objectives for reinforcement learning. In: *Advances in Neural Information Processing Systems (NeurIPS)*. vol. 32, pp. 5194–5204 (2019)
34. Zhang, Q., Li, H.: Moea/d: A multiobjective evolutionary algorithm based on decomposition. *IEEE Transactions on Evolutionary Computation* **11**(6), 712–731 (2007)
35. Zhang, Z., Yang, Y., Wang, J., Roijers, D.M., Nowé, A., Foster, C.: Convex hull value iteration for multi-objective reinforcement learning. *Advances in Neural Information Processing Systems* **33**, 20837–20848 (2020)
36. Zhu, W., Li, Q., Wang, Z.: A survey on multi-objective reinforcement learning: Advances and challenges. *ACM Computing Surveys* **55**(1), 1–32 (2023)
37. Zou, H., Hastie, T.: Regularization and variable selection via the elastic net. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **67**(2), 301–320 (2003)