

Continual Multi-Objective Reinforcement Learning via Reward Model Rehearsal

Lihe Li^{1,2}, Ruotong Chen^{1,2}, Ziqian Zhang^{1,2}, Zhichao Wu^{1,2}, Yi-Chen Li^{1,2,3},
Cong Guan^{1,2}, Yang Yu^{1,2,3*}, Lei Yuan^{1,2,3}

¹National Key Laboratory for Novel Software Technology, Nanjing University

²School of Artificial Intelligence, Nanjing University

³Polixir Technologies

{lilh, zhangzq, liyc, guanc, yuanl}@lamda.nju.edu.cn, {chenrt, zcwu}@smail.nju.edu.cn,
yuy@nju.edu.cn

Abstract

Multi-objective reinforcement learning (MORL) approaches address real-world problems with multiple objectives by learning policies maximizing returns weighted by different user preferences. Typical methods assume the objectives remain unchanged throughout the agent’s lifetime. However, in some real-world situations, the agent may encounter dynamically changing learning objectives, i.e., different vector-valued reward functions at different learning stages. This issue of evolving objectives has not been considered in problem formulation or algorithm design. To address this issue, we formalize the setting as a continual MORL (CMORL) problem for the first time, accounting for the evolution of objectives throughout the learning process. Subsequently, we propose Continual Multi-Objective Reinforcement Learning via Reward Model Rehearsal (CORE3), incorporating a dynamic agent network for rapid adaptation to new objectives. Moreover, we develop a reward model rehearsal technique to recover the reward signals for previous objectives, thus alleviating catastrophic forgetting. Experiments on four CMORL benchmarks showcase that CORE3 effectively learns policies satisfying different preferences on all encountered objectives, and outperforms the best baseline by 171%, highlighting the capability of CORE3 to handle situations with evolving objectives.

1 Introduction

Reinforcement learning (RL) has garnered prominent attention in recent years [Wang *et al.*, 2022], and made exciting progress in various real-world sequential decision-making problems, like robotic control [Singh *et al.*, 2022], autonomous driving [Kiran *et al.*, 2022], and aligning large language models with human values [Kaufmann *et al.*, 2023], etc. These problems could be formalized as completing a specific objective, i.e., maximizing a scalar return. Nevertheless, numerous real-world problems encompass multiple,

possibly conflicting objectives. An illustrative example is an autonomous driving agent tasked with multiple objectives, such as maximizing speed and minimizing fuel consumption. To tackle such complex scenarios, multi-objective reinforcement learning (MORL) approaches [Rojers *et al.*, 2013; Liu *et al.*, 2014] have been introduced to concurrently learn these multiple objectives with a vector-valued reward function, obtain a set of policies maximizing returns weighted by different preferences, and employ the appropriate policy given a user preference during testing, like driving fast for a passenger with urgent matters, or driving slowly and smoothly for a leisurely traveler.

Given the extensive and significant applications of MORL, many approaches have been developed [Hayes *et al.*, 2022]. Typical methods primarily concentrate on developing various techniques to enhance the learning efficiency over the multiple objectives, including applying meta-learning techniques [Chen *et al.*, 2019], utilizing prediction-guided evolutionary learning [Xu *et al.*, 2020], developing preference-conditioned network architecture [Abels *et al.*, 2019], proposing novel Bellman update strategies [Basaklar *et al.*, 2023], etc. Beyond that, these methods assume that throughout the agent’s entire lifetime, the learning objectives remain unchanged. However, the agent may encounter dynamically changing learning objectives, i.e., altering vector-valued reward functions. On one hand, new objectives and their reward signals may appear at different learning stages, as it is natural and inevitable for human users to propose new requirements. On the other hand, some objectives may be removed from the agent’s perspective, as their reward signals become inaccessible after a limited training period, yet the agent must retain the ability to complete them [Chang *et al.*, 2021]. For instance, the reward for the objective “stay in the middle lane” will be unavailable after the corresponding sensors run out of power or malfunction [Shaheen *et al.*, 2022; Wang *et al.*, 2023], while the autonomous driving agent is still obliged to manage it for certain passengers. Without considering these scenarios, the underlying assumption of fixed objectives made by existing methods hinders the further development of MORL and its practical application in scenarios with evolving objectives.

To bridge the gap between existing methods and the mentioned scenarios, and endow MORL agents with continual learning ability, we take the issues into consideration

*Corresponding Author

and formalize the setting as a continual MORL (CMORL) problem for the first time, where the agent encounters a sequence of MORL tasks with objectives altering continually, and should learn to complete all encountered objectives (Figure 1). Furthermore, we propose Continual Multi-Objective Reinforcement Learning via Reward Model Rehearsal (CORE3), incorporating a dynamic agent network with a multi-head architecture [Kessler *et al.*, 2022; Zhang *et al.*, 2023a], which enables flexible expansion for new objectives and effective knowledge transfer [Zhu *et al.*, 2023]. To avoid catastrophic forgetting about objectives lacking reward signals, inspired by the concept of rehearsal in continual learning [Parisi *et al.*, 2019], we develop a reward model rehearsal approach to recover these signals and use them to update the policy. This approach not only provides a theoretical guarantee but also effectively alleviates the unique challenge of catastrophic forgetting in CMORL, ensuring the agent’s high performance on all encountered objectives.

We conduct experiments within four CMORL benchmarks featuring evolving objective sets. The results showcase that CORE3 effectively learns to accomplish all encountered objectives in accordance with different preferences, outperforming the best baseline by 171% on Hypervolume, a widely-used metric in MORL [Hayes *et al.*, 2022]. More results further provide insight into how CORE3 successfully tackles the CMORL problem with the proposed technologies. These findings highlight the capability of CORE3 to effectively handle real-world situations involving evolving objectives.

2 Related Work

Multi-objective reinforcement learning (MORL) extends the conventional RL framework with a single objective to multi-objective settings [Rojers *et al.*, 2013; Liu *et al.*, 2014], and can be applied to real-world problems like hyperparameter tuning [Chen *et al.*, 2021], canal control [Ren *et al.*, 2021], etc. Among the approaches, the single-policy series [Pan *et al.*, 2020; Siddique *et al.*, 2020; Hwang *et al.*, 2023] predefines objective preferences, converting the problem into a single-objective one which is solvable by traditional RL methods. When the preference cannot be known in advance, multi-policy approaches aim to learn a set of policies that approximates the Pareto front of solutions, and select the optimal one for testing given a preference. One classical work is PG-MORL [Xu *et al.*, 2020], which updates a policy population using an evolutionary algorithm to approximate the Pareto front. Another line of works like Envelope [Yang *et al.*, 2019] and PD-MORL [Basaklar *et al.*, 2023], train a single preference-conditioned network over multiple preferences with different Bellman update strategies. To release the burden of learning multiple policies simultaneously, meta-policy approaches [Chen *et al.*, 2019; Zhang *et al.*, 2023b] first train a meta-policy and finetune it with a small amount of update steps to derive the solution for a given preference. To test the learning efficiency of various MORL methods, researchers also develop different benchmarks [Xu *et al.*, 2020; Felten *et al.*, 2023]. Despite these progress, existing MORL approaches and benchmarks focus on learning a fixed set of objectives, while our work is the first to explore the MORL

problem with objectives evolving continually.

Continual reinforcement learning (CRL) [Ring, 1995; Khetarpal *et al.*, 2022] focuses on enabling agents to learn a sequence of different RL tasks and balance the stability-plasticity dilemma, i.e., alleviating catastrophic forgetting of old tasks while adapting to new tasks. One classic method is EWC [Kirkpatrick *et al.*, 2017], which adds an l_2 -distance-based regularizer to constrain the update of the agent network. Another task-agnostic approach CLEAR [Rolnick *et al.*, 2019] stores the transitions of every encountered task in the buffer, and uses them to rehearse the agent. Other works [Huang *et al.*, 2021; Kessler *et al.*, 2023] learn a world model to assist the agent’s continual learning. OWL [Kessler *et al.*, 2022] and DaCoRL [Zhang *et al.*, 2023a] utilize a multi-head architecture to enhance the agent’s continual learning ability. CSP [Gaya *et al.*, 2023] builds a subspace of policies continually. CPPO [Anonymous, 2024] continually learns from human preferences based on PPO [Schulman *et al.*, 2017]. Other researchers also develop benchmarks [Wolczyk *et al.*, 2021; Powers *et al.*, 2022] to test the continual learning ability of different CRL methods. While these methods focus on the standard RL tasks with only one objective, i.e., scalar reward signals, our work presents the first exploration into MORL in the continual learning setting.




Figure 1: The overall workflow of continual MORL (CMORL).

3 Problem Formulation

A multi-objective sequential decision making problem can be formulated as a Multi-Objective Markov Decision Process (MOMDP), with tuple $\mathcal{M} = \langle \mathcal{S}, \mathcal{A}, P, \mathcal{K}, \mathbf{R}, \Omega, f_\omega, \gamma \rangle$. Here, \mathcal{S} and \mathcal{A} denote the state space and the action space. $P : \mathcal{S} \times \mathcal{A} \rightarrow \text{Pr}(\mathcal{S})$ is the transition function. $\mathcal{K} = \{k_1, \dots, k_m\}$ is the set of m different objectives, a subset of all possible objectives \mathbb{K} . $\mathbf{R} : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}^m$ is the vector-valued reward function. Ω is the preference space. $f_\omega : \mathbb{R}^m \rightarrow \mathbb{R}$ is the scalarization function under preference $\omega \in \Omega$, and γ is the discount factor. We focus on the linear reward scalarization setting, i.e., $f_\omega(\mathbf{R}(s, a)) = \omega^\top \mathbf{R}(s, a)$, for $s \in \mathcal{S}, a \in \mathcal{A}$, and $\Omega = \Delta^m$ is the unit simplex, aligning with the established practices in the MORL literature [Hayes *et al.*, 2022]. At each time step t , the agent observes state s_t , takes action a_t , and receives a reward vector $\mathbf{r}_t = \mathbf{R}(s_t, a_t) = (R_{k_1}(s_t, a_t), \dots, R_{k_m}(s_t, a_t))$, where $R_{k_i}(s_t, a_t)$ is the reward signal of objective k_i ($i = 1, \dots, m$). The agent knows the objective that each scalar reward corresponds to. Here, we learn a preference-conditioned agent policy $\pi :$

$S \times \Delta^m \rightarrow \Pr(\mathcal{A})$. The goal is to find the optimal policy π^* , such that $\forall \omega \in \Delta^m, \pi^* = \arg \max_{\pi} \omega^\top J(\pi(\cdot|\cdot; \omega))$, where $J(\pi) = \mathbb{E}_{\tau \sim \pi} [\sum_{t=0}^{\infty} \gamma^t \mathbf{R}(s_t, a_t)]$ is the value vector, and $\tau = (s_0, a_0, s_1, a_1, \dots)$ is the trajectory.

In this work, we aim to solve the CMORL problem where the agent encounters a sequence of multi-objective tasks: $(\mathcal{M}_1, \dots, \mathcal{M}_n, \dots)$, as shown in Figure 1. Each task \mathcal{M}_n is a MOMDP with specific objectives $\mathcal{K}_n = \{k_1^n, \dots, k_{m_n}^n\} \subset \mathbb{K}$, reward function $\mathbf{R}_n = (R_1^n, \dots, R_{m_n}^n)$, and preference space $\Omega_n = \Delta^{m_n}$. Here, R_i^n denotes $R_{k_i^n}$ ($i = 1, \dots, m_n$) for simplicity. The agent knows if the task has changed. However, it can only receive the reward signals of the current objective set \mathcal{K}_n , and is evaluated on all encountered ones, i.e., $\mathcal{K}_n^{\text{all}} = \cup_{i=1}^n \mathcal{K}_i = \{k_1^1, \dots, k_{m_1}^1\} \cup \dots \cup \{k_1^n, \dots, k_{m_n}^n\} = \{k_1, \dots, k_{M_n}\}$, where M_n represents the total number of objectives introduced up to and including task \mathcal{M}_n . Note that $M_n \leq \sum_{i=1}^n m_i$ because some objectives may appear repeatedly in different tasks (e.g., $k_1^1 = k_1^2 = k_1$). These settings of CMORL necessitate the agent’s rapid adaptation to new objectives while maintaining the ability to accomplish the previously encountered ones.

4 Method

In this section, we present the detailed design of our proposed method CORE3. First, we introduce a novel dynamic agent network architecture that adaptively expands as new objectives arrive, facilitating rapid adaptation and circumventing the need to learn the entire network from scratch. Next, we utilize a multi-objective reward model to recover reward signals of previously encountered objectives, thereby alleviating the phenomenon of catastrophic forgetting. Through the incorporation of these techniques, CORE3 effectively learns the sequence of multi-objective tasks in a continual manner.

4.1 Dynamic Network for Evolving Objectives

Under the CMORL setting, the learning objectives alter continually along with the task sequence, leading to an increase in the number of all encountered objectives. As a result, the conventional network architecture with fixed input and output shapes will no longer work. To address this problem, we design a dynamic and expandable network architecture.

Concretely, for discrete action spaces, we learn a vectorized Q network $\mathbf{Q}(s, a, \omega; \theta)$ to approximate the action-value function and use it as the policy. We design the Q network architecture as a combination of a feature extractor $\mathcal{E}(s, a, \omega; \xi)$ and distinct heads $\{h(e; \psi_i)\}_{i=1}^{M_n}$ for all encountered objectives. The feature extractor first transforms the varying-length preference ω into a fixed-length embedding z_ω using a GRU [Chung *et al.*, 2014]. Then, a multi-layer perceptron (MLP) takes the state, action, and embedding z_ω as input and outputs feature $e = \mathcal{E}(s, a, \omega; \xi)$. Then, each MLP head $h(e; \psi_i)$ takes e as input, and outputs the Q value of objective k_i , forming the overall output Q vector $(h(e; \psi_1), \dots, h(e; \psi_{M_n})) = \mathbf{Q}(s, a, \omega; \theta)$ for all objectives. With the Q vector, the CMORL agent is able to select the optimal action $a = \arg \max_{\tilde{a} \in \mathcal{A}} \omega^\top \mathbf{Q}(s, \tilde{a}, \omega; \theta)$ w.r.t. the given preference ω . During training task \mathcal{M}_n , the Q network is optimized to minimize the following loss:

$$\mathcal{L}(\theta) = \mathbb{E}_{(s, a, \mathbf{r}, s', \omega) \sim \mathcal{D}} [\|\mathbf{y} - \mathbf{Q}(s, a, \omega; \theta)\|_2], \quad (1)$$

where $\mathbf{Q}(s, a, \omega; \theta) = (h(e; \psi_1), \dots, h(e; \psi_{M_n}))$,

and \mathcal{D} is the replay buffer, $\theta = (\xi, \{\psi_i\}_{i=1}^{M_n})$ are the updated parameters of the networks, $\mathbf{y} = \mathbf{r} + \gamma \mathbf{Q}(s', a', \omega; \theta^-)$ denotes the target value vector which is obtained using target Q network’s parameters θ^- and the action $a' = \arg \max_{\tilde{a} \in \mathcal{A}} \omega^\top \mathbf{Q}(s', \tilde{a}, \omega; \theta^-)$. When new objectives arrive, we dynamically create new objective heads and expand the network. For continuous action spaces, we apply $\mathbf{Q}(s, a, \omega; \theta)$ as the critic and learn an actor $\pi(a|s, \omega; \phi)$. When computing the target value vector \mathbf{y} to optimize \mathbf{Q} , the action a' is obtained by the actor: $a' \sim \pi(\cdot|s', \omega; \phi)$. The actor $\pi(\cdot|s, \omega; \phi)$ is also equipped with a GRU to process the varying-length preference, and is optimized to maximize $\omega^\top \mathbf{Q}$. Such network architecture not only handles the increasing objective number, but also facilitates efficient knowledge transfer between tasks with different objectives. More details about the network architecture and the optimization process are provided in Appendix C.1.

4.2 Multi-Objective Reward Model Rehearsal

With the designed network architecture and loss function, the agent can learn to complete all encountered objectives effectively. Nevertheless, the computation of the loss function $\mathcal{L}(\theta)$ defined in Equation 1 requires access to the reward signals of all previous objectives, i.e., $\mathcal{K}_n^{\text{all}} \setminus \mathcal{K}_n$. Unfortunately, in the CMORL setting, these reward signals are unavailable from the environment. Simply neglecting these objectives would lead to the undesirable consequence of the agent losing its capability to accomplish them, i.e., catastrophic forgetting phenomenon. These issues pose a great challenge to the agent’s learning process.

To solve these issues, we propose learning a multi-objective reward model $\hat{\mathbf{R}}$ to recover the reward signals for all previous objectives $\mathcal{K}_n^{\text{all}} \setminus \mathcal{K}_n$. Concretely, when learning task \mathcal{M}_n , the environment provides reward signals $\mathbf{r}_t = (R_1^n(s_t, a_t), \dots, R_{m_n}^n(s_t, a_t))$ corresponding to objectives $\mathcal{K}_n = \{k_1^n, \dots, k_{m_n}^n\}$. Next, we predict reward vector $\hat{\mathbf{r}}_t = \hat{\mathbf{R}}(s_t, a_t) = (\hat{R}_1(s_t, a_t), \dots, \hat{R}_{M_n}(s_t, a_t))$. Note that if objective $k_i \in \mathcal{K}_n$ ($i = 1, \dots, M_n$), we directly set $\hat{R}_i(s_t, a_t)$ as the ground-truth reward in \mathbf{r}_t . Then, $\hat{\mathbf{r}}_t$ is employed to compute the target value vector in Equation 1, to rehearse the agent about previous objectives. From the agent’s perspective, the reward signals for all encountered objectives persist throughout the entire learning process. Consequently, the agent’s performance significantly depends on the quality of the multi-objective reward model. Here, we theoretically analyze their relationship.

Theorem 1 (Bounded Performance Gap). *Consider two MOMDPs, \mathcal{M} and $\hat{\mathcal{M}}$, differing only in their reward functions: $\mathbf{R} = (R_1, \dots, R_m)$ and $\hat{\mathbf{R}} = (\hat{R}_1, \dots, \hat{R}_m)$. The value vectors of a policy π in \mathcal{M} and $\hat{\mathcal{M}}$ are $J_{\mathcal{M}}(\pi) = \mathbb{E}_{\tau \sim \pi} [\sum_{t=0}^{\infty} \gamma^t \mathbf{R}(s_t, a_t)]$ and $J_{\hat{\mathcal{M}}}(\pi) = \mathbb{E}_{\tau \sim \pi} [\sum_{t=0}^{\infty} \gamma^t \hat{\mathbf{R}}(s_t, a_t)]$, respectively. The optimal policies for \mathcal{M} and $\hat{\mathcal{M}}$ are denoted as π^* and $\hat{\pi}^*$, satisfy-*

ing $\forall \omega \in \Delta^m, \pi^* = \arg \max_{\pi} \omega^\top J_{\mathcal{M}}(\pi)$, and $\hat{\pi}^* = \arg \max_{\pi} \omega^\top J_{\hat{\mathcal{M}}}(\pi)$. Define the difference between these reward functions as $\delta = \sup_{s \in \mathcal{S}, a \in \mathcal{A}} \|\mathbf{R}(s, a) - \hat{\mathbf{R}}(s, a)\|_\infty$. Based on these definitions, the following inequality holds,

$$\forall \omega \in \Delta^m, \omega^\top J_{\mathcal{M}}(\pi^*) - \omega^\top J_{\mathcal{M}}(\hat{\pi}^*) \leq \frac{2\delta}{1-\gamma}. \quad (2)$$

The proof is provided in Appendix A. As indicated by Theorem 1, the performance gap between the policy learned from the recovered rewards $\hat{\mathbf{r}}_t$ and the policy learned from the ground-truth rewards \mathbf{r}_t is bounded by $\frac{2\delta}{1-\gamma}$, with δ representing the prediction error of the multi-objective reward model in our context. So, we can minimize the performance gap by improving the reward model’s prediction accuracy. To achieve this goal and circumvent the necessity of training a separate reward model for each objective, we also design a multi-head architecture for the reward model network. A dynamic encoder \mathcal{F} takes the current state and action as input, and outputs the dynamic feature $z = \mathcal{F}(s, a; \vartheta)$. For each objective $k_i (i = 1, \dots, M_n)$, a reward prediction head $g(\cdot; \varphi_i)$ takes z as input, and predicts reward $\hat{R}_i(s, a) = g(z; \varphi_i)$. New heads will be created when novel objectives appear. We also learn a state prediction head to predict the next state $\hat{s}' = g(z; \varphi_s)$, thereby stabilizing training and enhancing the dynamic encoder’s ability to capture the fundamental information of the environment. Let $\Phi = (\vartheta, \{\varphi_i^n\}_{i=1}^{m_n}, \varphi_s)$ denote the network parameters. The model is optimized in an end-to-end manner to minimize the following prediction loss term:

$$\begin{aligned} \mathcal{L}_{\text{pred}}(\Phi) = & \mathbb{E}_{(s, a, \mathbf{r}_{\text{env}}, s') \sim \mathcal{D}} \left[\|s' - g(z; \varphi_s)\|_2 \right. \\ & \left. + \|\mathbf{r}_{\text{env}} - (g(z; \varphi_1^n), \dots, g(z; \varphi_{m_n}^n))\|_2 \right], \end{aligned} \quad (3)$$

where \mathcal{D} is the replay buffer, $\mathbf{r}_{\text{env}} = \mathbf{R}_n(s, a)$ is the ground-truth reward vector received from the environment. Specifically, only the reward heads $\{\varphi_i^n\}_{i=1}^{m_n}$ corresponding to the objectives $\{k_1^n, \dots, k_{m_n}^n\}$ of task \mathcal{M}_n are updated when $n > 1$. Empirical results showcase that it does not weaken the expressiveness and accuracy of the model. We also apply the ensemble technique to improve the robustness of the model. More details about model learning are provided in Appendix C.2. With the accurate learned multi-objective reward model, we can recover the lost reward signals to rehearse the agent about previously encountered objectives, alleviating catastrophic forgetting and ensuring the performance on all objectives encountered so far.

4.3 Overall Learning Procedure

Through the incorporation of the technologies mentioned in Section 4.1 and 4.2, we formulate our overall method, Continual Multi-Objective Reinforcement Learning via Reward Model Rehearsal (CORE3), which can effectively learn the sequence of multi-objective tasks in a continual manner. In this section, we outline the overall learning procedure of CORE3 in Algorithm 1.

For each learning task \mathcal{M}_n , we use a replay buffer to store and only store its transition data (Line 3), obviating the explicit storage of transition data of each encountered task, and

Algorithm 1 CORE3

Input: Multi-objective task $(\mathcal{M}_1, \mathcal{M}_2, \dots)$, policy π , multi-objective reward model $\hat{\mathbf{R}}$, and replay buffer \mathcal{D} .

Parameter: Q function $\theta = (\xi)$, actor ϕ (for continuous action spaces), and multi-objective reward model $\Phi = (\vartheta, \varphi_s)$.

```

1:  $\mathcal{K}_0^{\text{all}} \leftarrow \emptyset, M_n \leftarrow 0$ .
2: for  $n = 1, 2, \dots$  do
3:   Clear replay buffer  $\mathcal{D}$ .
4:    $\mathcal{K}_n^{\text{all}} \leftarrow \mathcal{K}_{n-1}^{\text{all}} \cup \mathcal{K}_n, m_n \leftarrow |\mathcal{K}_n|, M_n \leftarrow |\mathcal{K}_n^{\text{all}}|$ .
5:   if  $\mathcal{K}_n^{\text{all}} \setminus \mathcal{K}_{n-1}^{\text{all}} \neq \emptyset$  then
6:     // Dynamic Network Expansion
7:     Create objective heads  $\{\psi_i\}_{i=M_{n-1}+1}^{M_n}$ , add to  $\theta$ .
8:     Create reward heads  $\{\varphi_i\}_{i=M_{n-1}+1}^{M_n}$ , add to  $\Phi$ .
9:   end if
10:  while task  $\mathcal{M}_n$  has not ended do
11:    Uniformly sample preference  $\omega \sim \Delta^{M_n}$ .
12:    Execute  $\pi(\cdot|\cdot, \omega)$  in task  $\mathcal{M}_n$  to collect transitions  $\{(s, a, \mathbf{r}_{\text{env}}, s', \omega)\}$  and append to  $\mathcal{D}$ .
13:    Update  $\Phi = (\vartheta, \{\varphi_i^n\}_{i=1}^{m_n}, \varphi_s)$  of the multi-objective reward model  $\hat{\mathbf{R}}$  to minimize  $\mathcal{L}_{\text{pred}}$  in Equation 2. // Reward Model Learning
14:    Sample transitions  $\{(s, a, \mathbf{r}_{\text{env}}, s', \omega)\} \sim \mathcal{D}$ , replace  $\mathbf{r}_{\text{env}}$  with prediction  $\hat{\mathbf{r}}$ , as mentioned in Section 4.2. // Reward Model Rehearsal
15:    Update  $\theta = (\xi, \{\psi_i\}_{i=1}^{M_n})$  of the Q function to minimize  $\mathcal{L}$  in Equation 1. // Policy Learning
16:    (For continuous action spaces) Update actor  $\phi$ .
17:  end while
18:  Evaluate  $\pi$  on all encountered objectives  $\mathcal{K}_n^{\text{all}}$ .
19: end for

```

ensuring the scalability of our approach. If novel objectives appear in the current task \mathcal{K}_n , we dynamically expand the policy network and the multi-objective reward model network (Lines 5 ~ 9). Then, the agent explores \mathcal{M}_n , and collect transitions into the replay buffer \mathcal{D} . With the collected data, we finetune the reward model so that it can predict the reward of novel objectives (Line 13). Leveraging this reward model, we recover the reward signals of previously encountered objectives, and utilize them to train and rehearse the agent (Line 14). Finally, the policy is optimized based on the reward signals for all encountered objectives, alleviating the catastrophic forgetting phenomenon (Lines 15 ~ 16). To assess the efficacy of CORE3 in addressing the CMORL problem, we evaluate the learned policy on all encountered objectives (Line 18), as elaborated in Section 5.

5 Experiments

In this section, we conduct experiments to answer the following questions: (1) Can CORE3 better accomplish all encountered objectives compared with the baselines, and alleviate forgetting about previous objectives lacking reward signals (Section 5.2) ? (2) How does CORE3 adapt to new objectives, and learn previous objectives from the perspective of Pareto fronts (Section 5.3) ? (3) How does the hyperparameters influence the performance of CORE3 (Section 5.4) ?

Envs	Metrics	Finetune	CLEAR	EWC	CORE3 (Oracle)	CORE3 (Ours)
FTN	Hv (\uparrow)	1.00 ± 0.07	0.85 ± 0.07	1.04 ± 0.07	1.10 ± 0.04	1.06 ± 0.00
	SP (\downarrow)	1.00 ± 0.14	0.73 ± 0.02	0.58 ± 0.06	0.72 ± 0.01	0.82 ± 0.02
	BWT (\uparrow)	-0.12 ± 0.02	0.00 ± 0.01	-0.14 ± 0.01	0.03 ± 0.00	0.02 ± 0.01
Grid	Hv (\uparrow)	1.00 ± 0.82	4.21 ± 0.66	2.26 ± 1.34	4.60 ± 0.26	4.54 ± 0.60
	SP (\downarrow)	1.00 ± 0.58	0.42 ± 0.11	0.47 ± 0.10	0.37 ± 0.02	0.44 ± 0.10
	BWT (\uparrow)	-0.44 ± 0.07	-0.14 ± 0.17	-0.58 ± 0.05	-0.07 ± 0.12	-0.17 ± 0.15
Ant	Hv (\uparrow)	1.00 ± 0.42	1.78 ± 0.37	1.03 ± 0.60	2.30 ± 0.18	1.89 ± 0.35
	SP (\downarrow)	1.00 ± 0.56	0.33 ± 0.06	1.43 ± 1.78	1.06 ± 0.57	0.73 ± 0.72
	BWT (\uparrow)	-0.27 ± 0.02	0.01 ± 0.18	-0.35 ± 0.1	0.34 ± 0.18	0.11 ± 0.16
Hopper	Hv (\uparrow)	1.00 ± 1.22	2.03 ± 3.49	1.01 ± 1.43	7.84 ± 0.29	8.23 ± 1.42
	SP (\downarrow)	1.00 ± 0.76	7.59 ± 2.98	4.76 ± 3.54	0.35 ± 0.11	0.21 ± 0.10
	BWT (\uparrow)	-0.40 ± 0.07	-0.38 ± 0.22	-0.37 ± 0.07	0.11 ± 0.09	0.06 ± 0.05
API (%) (\uparrow)	Hv	\	121.72	33.61	296.03	293.18
	SP	\	-126.81	-80.85	37.61	45.37
	BWT	\	69.09	-17.87	145.36	109.09

Table 1: Metric values are presented as mean \pm std for different methods across four CMORL benchmarks. For simplicity, we re-scale Hv and SP values by taking the Finetune’s results as an anchor and present the average performance improvement (API, %) relative to it. The best results except the Oracle baseline in each row is highlighted in **bold**. \uparrow indicates better performance with higher values, and \downarrow the opposite.

5.1 Benchmarks, Baselines, and Metrics

To provide a testbed for CORE3, we first design four CMORL benchmarks with both discrete and continuous action spaces, by extending widely-used MORL benchmarks into continual learning settings. The first benchmark is **Fruit Tree Navigation (FTN)** [Yang *et al.*, 2019], a task with a full binary tree, and an agent navigating from the root node to one of the leaf node to receive a 6-dimension reward, representing the amount of six different nutrients. Next, inspired by the four-room [Alegre *et al.*, 2022] environment, we design a more complicated **Grid** world benchmark with five goal positions. The agent is rewarded when navigating closer to the goals, and should navigate to the optimal position given a preference on the goals. Both FTN and Grid are environments with discrete action spaces, and we design two continuous control CMORL benchmarks based on MuJoCo physics engine [Todorov *et al.*, 2012; Xu *et al.*, 2020]. In the **Ant** benchmark, the agent controls the ant robot to complete five objectives, including saving energy and moving in $\{\pm x, \pm y\}$ -axes. Another MuJoCo benchmark is **Hopper**, where the five objectives of the hopper robot include saving energy, moving forward, moving backward, jumping high, and staying low. In the CMORL setting, the agent learns a sequence of five MOMDPs. Each MOMDP is trained for 100k steps in FTN and 500k steps otherwise, with two objectives drawn from the objectives mentioned above. More details about the benchmarks are in Appendix B.1.

Next, we compare CORE3 with multiple strong baselines. To evaluate the impact of reward model rehearsal on the agent’s learning process, we compare CORE3 with **Finetune**, which can be seen as an ablation of CORE3 without reward model rehearsal, and directly tunes the policy on the current task, ignoring previous objectives lacking reward signals. Additionally, we extend representative works in single-objective CRL into multi-objective settings as baselines. **CLEAR** [Rolnick *et al.*, 2019] is a task-agnostic method which stores transitions of all previous tasks and use them to rehearse the agent. **EWC** [Kirkpatrick *et al.*, 2017]

adds an l_2 -distance-based regularizer to constrain the update of the agent network. We further investigate **CORE3 (Oracle)**, representing a performance upper bound of CORE3, as it has access to the rewards of previous objectives, i.e., CORE3 with ground-truth multi-objective reward model. We run each method for five distinct seeds. More details about the baselines are in Appendix B.2.

To evaluate the methods’ performance on completing all encountered objectives, we randomly sample multiple preferences to test the agent, deriving a set V of all non-dominated return vectors, i.e., the approximated Pareto front. To measure the quality of V , we adopt two metrics widely used in the MORL literature [Hayes *et al.*, 2022]. The first one is **Hypervolume (Hv)** $:= \int_{H(V)} \mathbb{I}\{z \in H(V)\} dz$, where $H(V) := \{z \in \mathbb{R}^m : \exists \mathbf{v} \in V, \mathbf{v} \succ z \succ \mathbf{v}_0\}$, m is the number of objectives, \mathbf{v}_0 is a predefined reference point, \succ is the Pareto dominance relation, and \mathbb{I} is the indicator function. A larger Hv indicates a better approximation of the optimal Pareto front. The second one **Sparsity (SP)** $:= \frac{1}{|V|-1} \sum_{j=1}^m \sum_{i=1}^{|V|-1} (\tilde{V}_j(i) - \tilde{V}_j(i+1))^2$ measures the density of V , where $\tilde{V}_j(i)$ is the i -th value in the sorted list for the j -th objective values in V . Given two Pareto fronts with close Hv, the one with smaller SP is denser and considered better. As we study a continual learning problem, it is necessary to measure the forgetting phenomenon. We therefore calculate **Backward Transfer (BWT)** [Wang *et al.*, 2023] $:= \frac{1}{N-1} (\sum_{n=2}^N \frac{1}{n-1} \sum_{j=1}^{n-1} [(\alpha_n^j - \alpha_j^j) / \alpha_j^j])$, where $N = 5$ is the number of tasks, α_i^j is the Hv tested on objectives \mathcal{K}_j after training on task \mathcal{M}_i with objectives \mathcal{K}_i . BWT evaluates the influence of learning new objectives on completing the previous ones. A larger BWT indicates a better anti-forgetting.

5.2 Competitive Results

In this section, we compare CORE3 with the baselines on the four CMORL benchmarks. The overall results on all three metrics are shown in Table 1, and the learning curves on the primary metric, Hv of all objectives encountered so

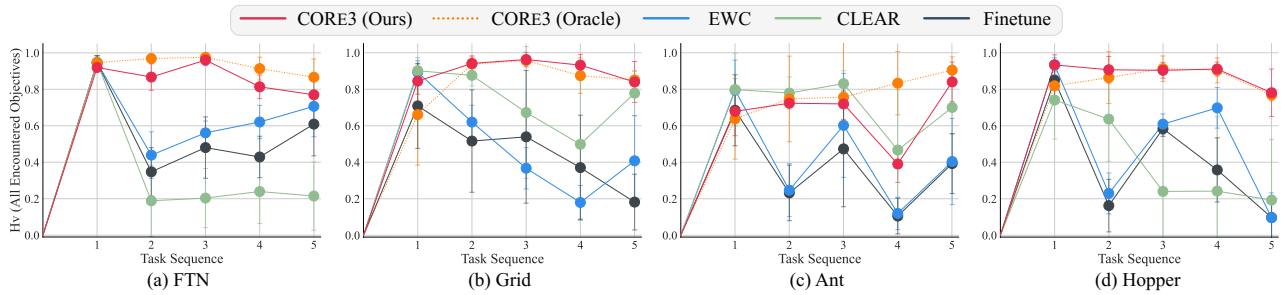


Figure 2: Performance comparison during the learning process on metric HV in four benchmarks. Each benchmark has a sequence of five MORL tasks, and each plot denotes the HV value on all objectives encountered so far. As the numerical range of HV varies greatly with the number of objectives, we re-scale the HV values in each task to the range of $[0, 1]$ for a comprehensive visualization.

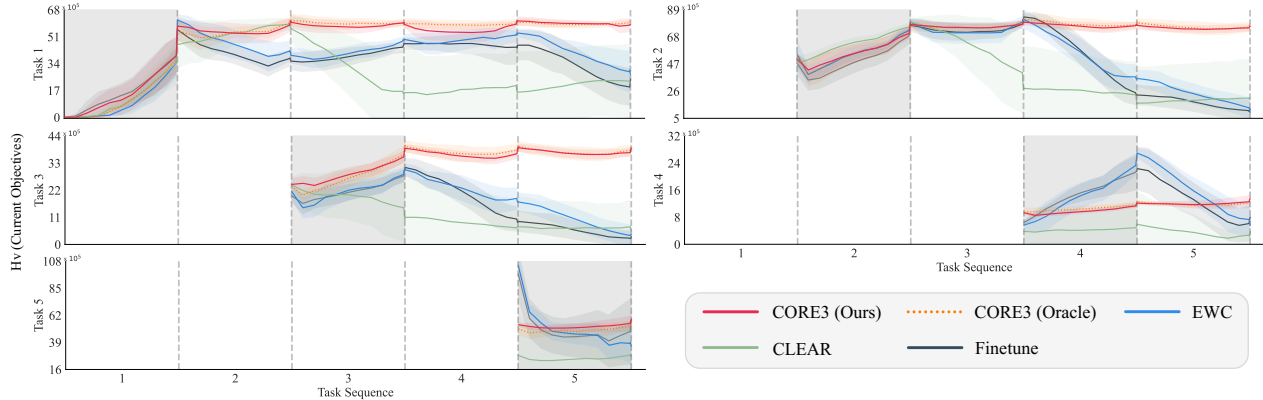


Figure 3: The learning curves of HV on *each* encountered task in the Hopper benchmark. Each of the five plots corresponds to a specific task, and a gray background within each plot signifies the period during which the agent is training on that particular task. There is blank space in the plots for task $n > 1$ because the task has not yet appeared. The results on other benchmarks are provided in Appendix D.1.

far, are displayed in Figure 2. We find that Finetune has the most inferior overall performance on HV and suffers from forgetting according to a low BWT, showing that ignoring the lost reward signals cannot perform well in CMORL settings, even with a well-designed multi-head policy network. To alleviate forgetting, CLEAR uses transitions of previous tasks to rehearse the agent, but still suffers from performance degradation, because the transitions trading off objectives in different tasks are still lacking. Another widely-used approach for single-objective CRL, EWC, also cannot perform well, demonstrating the necessity of specific considerations for multi-objective scenarios. Having access to ground-truth reward signals for all encountered objectives, CORE3 (Oracle) can be seen as an upper bound of performance, and achieves the best overall results, demonstrating the reward rehearsal strategy can solve the CMORL problem while conventional CRL approaches fail. Our approach CORE3, achieves comparable performance to Oracle on all four benchmarks and all three metrics, and outperforms the best baseline CLEAR on the average HV improvement by $293.18\% - 121.72\% = 171.46\%$, indicating the effectiveness of reward model rehearsal.

We further display the learning curves of HV on *each* encountered task in the Hopper benchmark in Figure 3. We observe the catastrophic forgetting phenomenon, where baselines suffer from performance degradation on tasks after training on them. Take task 1 as an example, all methods dis-

play similar learning curves when training on it, but baselines Finetune, EWC, and CLEAR suffer from performance degradation after starting to train on task 2 or task 3. On the contrary, CORE3 shows non-decreasing learning curves, which are comparable to CORE3 (Oracle) on every task, indicating that our approach successfully alleviates the catastrophic forgetting phenomenon, and maintains the agent’s ability to complete previous objectives. The results on other benchmarks are provided in Appendix D.1.

5.3 Pareto Front Analysis

To provide a more comprehensive illustration of how CORE3 adapts to new tasks and alleviates catastrophic forgetting, we further analyze the learning process by visualizing the Pareto fronts in Figure 4. As shown in Figure 4(a), after learning previous tasks, the CMORL agent is now confronted with a new task with two objectives. One objective (moving in $+y$ axis) has been learned before, and the other one (moving in $-x$ axis) is new. To provide reference and comparison, we introduce Learning-From-Scratch (LFS), which trains a randomly initialized policy network to complete these two objectives from scratch, and expands the Pareto front starting from a position near $(0, 0)$. The markers with different colors in the figure ($\bullet \rightarrow \bullet$) demonstrate the progression of the learning process. We then display the learning process of CORE3 in the same manner ($\blacklozenge \rightarrow \blacklozenge$). At the early beginning of the task, CORE3 already achieves a Pareto front that surpasses

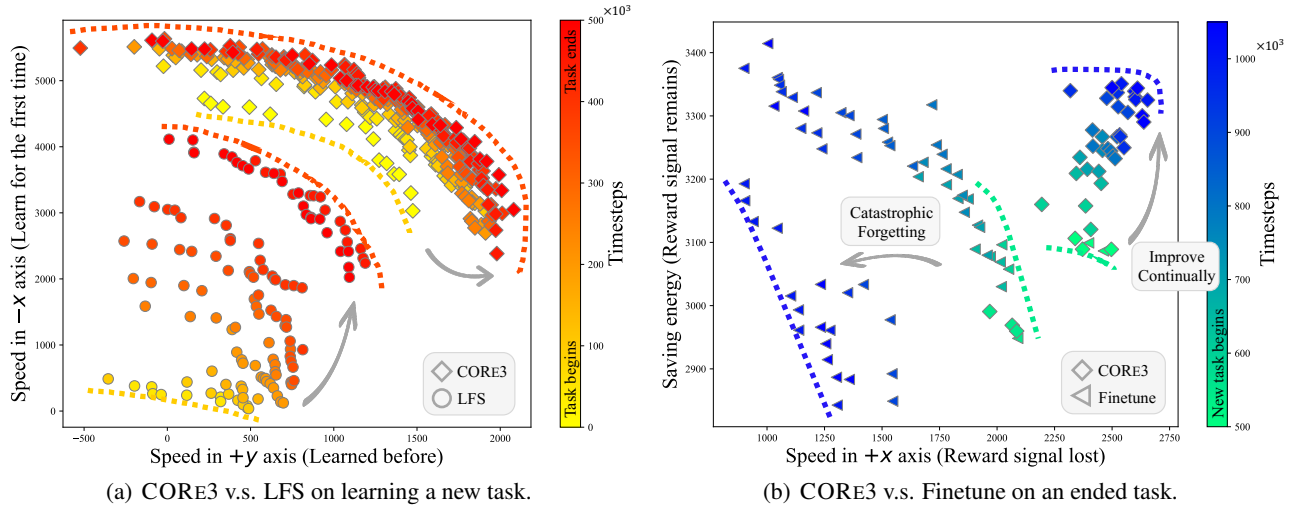


Figure 4: Comparison of Pareto fronts on the Ant benchmark. The x, y -axes represent the returns of two different objectives. Each marker is a found solution, and markers in the same color form an approximation of the optimal Pareto front. Different colors represent different training periods. (a) CORE3 v.s. Learning-From-Scratch (LFS) on learning a new task. LFS trains a random initialized policy network from scratch. (b) CORE3 v.s. Finetune on a task whose training period has ended. The displayed training period begins at 500k steps (lower right corner) instead of 0, indicating the solutions are induced by policies learning subsequent tasks with objectives different from the ones in x, y -axes.

the one obtained by LFS at the task’s conclusion. It indicates that, by learning from previous objectives, CORE3 has acquired fundamental and transferable task knowledge, like controlling the joints of the Ant robot in a reasonable way. Armed with this knowledge and the continual learning of the new task, CORE3 further effectively expands the Pareto front to a broader and denser curve, better approximating the optimal Pareto front of the task.

Next, we display the Pareto fronts of a task after training on it. As shown in Figure 4(b), the reward signal of one objective (moving in $+x$ axis) is lost in the subsequent task while the other remains. Without specific consideration, Finetune suffers from a significant performance degradation about 50% of the objective lacking reward signal ($\blacktriangleleft \rightarrow \blacktriangleright$), i.e., catastrophic forgetting. On the contrary, with the reward model rehearsal technique, CORE3 retains and continually improves the ability of completing this objective ($\blacklozenge \rightarrow \blackdiamond$), explaining how CORE3 successfully alleviates catastrophic forgetting.

5.4 Sensitivity Studies

The performance of CORE3 significantly depends on the quality of the multi-objective reward model, which is learned through established world model learning techniques [Luo *et al.*, 2022]. As the learning process includes multiple hyperparameters, we here conduct experiments on the Hopper benchmark to investigate how each one influences the performance of CORE3. One important hyperparameter is the ensemble size, as we learn multiple ensemble models to improve robustness. On one hand, a small ensemble size is insufficient for robustness. On the other hand, learning an excessively large number of ensemble models will reduce the overall efficiency. As shown in Figure 5, we find that CORE3 achieves the highest HV and lowest sparsity when the ensemble size is equal to 5. CORE3 with 10 ensemble models performs similar to the former, meaning that an ensemble size of 5 is

enough to learn a robust multi-objective reward model. When the ensemble size is less than 5, the reward model is not stable enough, resulting in worse performance. More detailed analysis of other important hyperparameters like the batch size of model learning is provided in Appendix D.2.

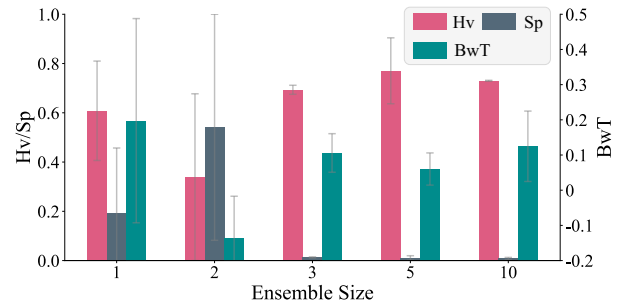


Figure 5: Sensitivity studies on model ensemble size in Hopper.

6 Final Remarks

Recognizing the significance of continual learning, this work takes a further step towards continual multi-objective RL (CMORL). We formulate this problem for the first time, where learning objectives may evolve throughout the agent’s lifetime. Then, we propose Continual Multi-Objective Reinforcement Learning via Reward Model Rehearsal (CORE3), which utilizes a dynamic agent network for adaptation to new objectives, and a multi-objective reward model to rehearse the agents about objectives lacking reward signals. Experiments demonstrate the effectiveness of our approach. Furthermore, an intriguing avenue for future exploration is extending multi-objective RL to the domain of multi-agent reinforcement learning (MARL) [Yuan *et al.*, 2023], thereby fostering progress in both these vital fields.

Acknowledgments

This work is supported by National Science Foundation of China (61921006). We would like to thank all reviewers for their valuable suggestions.

References

- [Abels *et al.*, 2019] Axel Abels, Diederik M. Roijers, Tom Lenaerts, Ann Nowé, and Denis Steckelmacher. Dynamic weights in multi-objective deep reinforcement learning. In *ICML*, pages 11–20, 2019.
- [Alegre *et al.*, 2022] Lucas N. Alegre, Ana L. C. Bazzan, and Bruno C. da Silva. Optimistic linear support and successor features as a basis for optimal policy transfer. In *ICML*, pages 394–413, 2022.
- [Anonymous, 2024] Anonymous. CPPO: Continual learning for reinforcement learning with human feedback. In *ICLR*, 2024.
- [Basaklar *et al.*, 2023] Toygun Basaklar, Suat Gumussoy, and Ümit Y. Ogras. PD-MORL: preference-driven multi-objective reinforcement learning algorithm. In *ICLR*, 2023.
- [Chang *et al.*, 2021] Yifan Chang, Wenbo Li, Jian Peng, Bo Tang, Yu Kang, Yinjie Lei, Yuanmiao Gui, Qing Zhu, Yu Liu, and Haifeng Li. Reviewing continual learning from the perspective of human-level intelligence. *arXiv preprint arXiv:2111.11964*, 2021.
- [Chen *et al.*, 2019] Xi Chen, Ali Ghadirzadeh, Mrarten Björkman, and Patric Jensfelt. Meta-learning for multi-objective reinforcement learning. In *IROS*, pages 977–983, 2019.
- [Chen *et al.*, 2021] SenPeng Chen, Jia Wu, and XiYuan Liu. Emorl: Effective multi-objective reinforcement learning method for hyperparameter optimization. *Engineering Applications of Artificial Intelligence*, 104:104315, 2021.
- [Chung *et al.*, 2014] Junyoung Chung, Caglar Gulcehre, KyungHyun Cho, and Yoshua Bengio. Empirical evaluation of gated recurrent neural networks on sequence modeling. *arXiv preprint arXiv:1412.3555*, 2014.
- [Felten *et al.*, 2023] Florian Felten, Lucas Nunes Alegre, Ann Nowe, Ana L. C. Bazzan, El Ghazali Talbi, Grégoire Danoy, and Bruno Castro da Silva. A toolkit for reliable benchmarking and research in multi-objective reinforcement learning. In *NeurIPS*, 2023.
- [Gaya *et al.*, 2023] Jean-Baptiste Gaya, Thang Doan, Lucas Caccia, Laure Soulier, Ludovic Denoyer, and Roberta Raileanu. Building a subspace of policies for scalable continual learning. In *ICLR*, 2023.
- [Hayes *et al.*, 2022] Conor F Hayes, Roxana Rădulescu, Eugenio Bargiacchi, Johan Källström, Matthew Macfarlane, Mathieu Reymond, Timothy Verstraeten, Luisa M Zintgraf, Richard Dazeley, Fredrik Heintz, et al. A practical guide to multi-objective reinforcement learning and planning. *Autonomous Agents and Multi-Agent Systems*, 36(1):26, 2022.
- [Huang *et al.*, 2021] Yizhou Huang, Kevin Xie, Homanga Bharadhwaj, and Florian Shkurti. Continual model-based reinforcement learning with hypernetworks. In *ICRA*, pages 799–805, 2021.
- [Hwang *et al.*, 2023] Minyoung Hwang, Luca Weihs, Chanwoo Park, Kimin Lee, Aniruddha Kembhavi, and Kiana Ehsani. Promptable behaviors: Personalizing multi-objective rewards from human preferences. *arXiv preprint arXiv:2312.09337*, 2023.
- [Kaufmann *et al.*, 2023] Timo Kaufmann, Paul Weng, Viktor Bengs, and Eyke Hüllermeier. A survey of reinforcement learning from human feedback. *arXiv preprint arXiv:2312.14925*, 2023.
- [Kessler *et al.*, 2022] Samuel Kessler, Jack Parker-Holder, Philip J. Ball, Stefan Zohren, and Stephen J. Roberts. Same state, different task: Continual reinforcement learning without interference. In *AAAI*, pages 7143–7151, 2022.
- [Kessler *et al.*, 2023] Samuel Kessler, Mateusz Ostaszewski, Michał Bortkiewicz, Mateusz Żarski, Maciej Wołczyk, Jack Parker-Holder, Stephen J. Roberts, and Piotr Miłoś. The effectiveness of world models for continual reinforcement learning. In *CoLLAs*, pages 184–204, 2023.
- [Khetarpal *et al.*, 2022] Khimya Khetarpal, Matthew Riemer, Irina Rish, and Doina Precup. Towards continual reinforcement learning: A review and perspectives. *Journal of Artificial Intelligence Research*, 75:1401–1476, 2022.
- [Kiran *et al.*, 2022] B. Ravi Kiran, Ibrahim Sobh, Victor Talpaert, Patrick Mannion, Ahmad A. Al Sallab, Senthil Kumar Yogamani, and Patrick Pérez. Deep reinforcement learning for autonomous driving: A survey. *IEEE Transactions on Intelligent Transportation Systems*, 23(6):4909–4926, 2022.
- [Kirkpatrick *et al.*, 2017] James Kirkpatrick, Razvan Pascanu, Neil Rabinowitz, Joel Veness, Guillaume Desjardins, Andrei A Rusu, Kieran Milan, John Quan, Tiago Ramalho, Agnieszka Grabska-Barwinska, et al. Overcoming catastrophic forgetting in neural networks. *Proceedings of the national academy of sciences*, 114(13):3521–3526, 2017.
- [Liu *et al.*, 2014] 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.
- [Luo *et al.*, 2022] Fan-Ming Luo, Tian Xu, Hang Lai, Xiong-Hui Chen, Weinan Zhang, and Yang Yu. A survey on model-based reinforcement learning. *SCIENCE CHINA Information Sciences*, 2022.
- [Pan *et al.*, 2020] Anqi Pan, Wenjun Xu, Lei Wang, and Hongliang Ren. Additional planning with multiple objectives for reinforcement learning. *Knowledge-Based Systems*, 193:105392, 2020.
- [Parisi *et al.*, 2019] German I. Parisi, Ronald Kemker, Jose L. Part, Christopher Kanan, and Stefan Wermter.

- Continual lifelong learning with neural networks: A review. *Neural Networks*, 113:54–71, 2019.
- [Powers *et al.*, 2022] Sam Powers, Eliot Xing, Eric Kolve, Roozbeh Mottaghi, and Abhinav Gupta. Cora: Benchmarks, baselines, and metrics as a platform for continual reinforcement learning agents. In *CoLLAs*, pages 705–743, 2022.
- [Ren *et al.*, 2021] Tao Ren, Jianwei Niu, Jiahe Cui, Zhenchao Ouyang, and Xuefeng Liu. An application of multi-objective reinforcement learning for efficient model-free control of canals deployed with iot networks. *Journal of Network and Computer Applications*, 182:103049, 2021.
- [Ring, 1995] Mark B. Ring. *Continual learning in reinforcement environments*. PhD thesis, University of Texas at Austin, TX, USA, 1995.
- [Roijers *et al.*, 2013] 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.
- [Rolnick *et al.*, 2019] David Rolnick, Arun Ahuja, Jonathan Schwarz, Timothy Lillicrap, and Gregory Wayne. Experience replay for continual learning. In *NeurIPS*, pages 350–360, 2019.
- [Schulman *et al.*, 2017] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.
- [Shaheen *et al.*, 2022] Khadija Shaheen, Muhammad Abdullah Hanif, Osman Hasan, and Muhammad Shafique. Continual learning for real-world autonomous systems: Algorithms, challenges and frameworks. *Journal of Intelligent & Robotic Systems*, 105(1):9, 2022.
- [Siddique *et al.*, 2020] Umer Siddique, Paul Weng, and Matthieu Zimmer. Learning fair policies in multiobjective (deep) reinforcement learning with average and discounted rewards. In *ICML*, pages 8905–8915, 2020.
- [Singh *et al.*, 2022] Bharat Singh, Rajesh Kumar, and Vinay Pratap Singh. Reinforcement learning in robotic applications: a comprehensive survey. *Artificial Intelligence Review*, 55(2):945–990, 2022.
- [Todorov *et al.*, 2012] Emanuel Todorov, Tom Erez, and Yuval Tassa. Mujoco: A physics engine for model-based control. In *IROS*, pages 5026–5033, 2012.
- [Wang *et al.*, 2022] Xu Wang, Sen Wang, Xingxing Liang, Dawei Zhao, Jincai Huang, Xin Xu, Bin Dai, and Qiguang Miao. Deep reinforcement learning: a survey. *IEEE Transactions on Neural Networks and Learning Systems*, 2022.
- [Wang *et al.*, 2023] Liyuan Wang, Xingxing Zhang, Hang Su, and Jun Zhu. A comprehensive survey of continual learning: Theory, method and application. *arXiv preprint arXiv:2302.00487*, 2023.
- [Wolczyk *et al.*, 2021] Maciej Wolczyk, Michal Zajac, Razvan Pascanu, Lukasz Kucinski, and Piotr Milos. Continual world: A robotic benchmark for continual reinforcement learning. In *NeurIPS*, pages 28496–28510, 2021.
- [Xu *et al.*, 2020] Jie Xu, Yunsheng Tian, Pingchuan Ma, Daniela Rus, Shinjiro Sueda, and Wojciech Matusik. Prediction-guided multi-objective reinforcement learning for continuous robot control. In *ICML*, pages 10607–10616, 2020.
- [Yang *et al.*, 2019] Runzhe Yang, Xingyuan Sun, and Karthik Narasimhan. A generalized algorithm for multi-objective reinforcement learning and policy adaptation. In *NeurIPS*, page 14636–14647, 2019.
- [Yuan *et al.*, 2023] Lei Yuan, Ziqian Zhang, Lihe Li, Cong Guan, and Yang Yu. A survey of progress on cooperative multi-agent reinforcement learning in open environment. *arXiv preprint arXiv:2312.01058*, 2023.
- [Zhang *et al.*, 2023a] Tiantian Zhang, Zichuan Lin, Yuxing Wang, Deheng Ye, Qiang Fu, Wei Yang, Xueqian Wang, Bin Liang, Bo Yuan, and Xiu Li. Dynamics-adaptive continual reinforcement learning via progressive contextualization. *IEEE Transactions on Neural Networks and Learning Systems*, 2023.
- [Zhang *et al.*, 2023b] Zizhen Zhang, Zhiyuan Wu, Hang Zhang, and Jiahai Wang. Meta-learning-based deep reinforcement learning for multiobjective optimization problems. *IEEE Transactions on Neural Networks and Learning Systems*, 34(10):7978–7991, 2023.
- [Zhu *et al.*, 2023] Zhuangdi Zhu, Kaixiang Lin, Anil K. Jain, and Jiayu Zhou. Transfer learning in deep reinforcement learning: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(11):13344–13362, 2023.