

Learning Diverse Risk Preferences in Population-Based Self-Play

**Yuhua Jiang^{*}, Qihan Liu^{*}, Xiaoteng Ma, Chenghao Li, Yiqin Yang,
Jun Yang[†], Bin Liang, Qianchuan Zhao[†]**

Department of Automation, Tsinghua University
 {jiangyh22, lqh20, ma-xt17, lch18, yangyiqi19}@mails.tsinghua.edu.cn
 {yangjun603, bliang, zhaoqc}@tsinghua.edu.cn

Abstract

Among the remarkable successes of Reinforcement Learning (RL), self-play algorithms have played a crucial role in solving competitive games. However, current self-play RL methods commonly optimize the agent to maximize the expected win-rates against its current or historical copies, resulting in a limited strategy style and a tendency to get stuck in local optima. To address this limitation, it is important to improve the diversity of policies, allowing the agent to break stalemates and enhance its robustness when facing with different opponents. In this paper, we present a novel perspective to promote diversity by considering that agents could have diverse risk preferences in the face of uncertainty. To achieve this, we introduce a novel reinforcement learning algorithm called Risk-sensitive Proximal Policy Optimization (RPPO), which smoothly interpolates between worst-case and best-case policy learning, enabling policy learning with desired risk preferences. Furthermore, by seamlessly integrating RPPO with population-based self-play, agents in the population optimize dynamic risk-sensitive objectives using experiences gained from playing against diverse opponents. Our empirical results demonstrate that our method achieves comparable or superior performance in competitive games and, importantly, leads to the emergence of diverse behavioral modes. Code is available at <https://github.com/Jackory/RPBT>.

Introduction

Reinforcement Learning (RL) has witnessed significant advancements in solving challenging decision problems, particularly in competitive games such as Go (Silver et al. 2016), Dota (Berner et al. 2019), and StarCraft (Vinyals et al. 2019). One of the key factors contributing to these successes is self-play, where an agent improves its policy by playing against itself or its previous policies. However, building expert-level AI solely through model-free self-play remains challenging, as it requires no access to a dynamic model or the prior knowledge of a human expert. One of the difficulties is that agents trained in self-play only compete against themselves, leading to policies that often become stuck in local optima and struggle to generalize to different opponents (Bansal et al. 2018).

^{*}These authors contributed equally.

[†]Corresponding authors.

Copyright © 2024, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

In this study, we focus on the problem of training agents that are both robust and high-performing against any type of opponent. We argue that the key to addressing this problem lies in producing a diverse set of training opponents, or in other words, learning diverse strategies in self-play. A promising paradigm to address this problem is population-based approaches (Vinyals et al. 2019; Jaderberg et al. 2019), where a collection of agents is gathered and trained against each other. However, the diversity achieved in such methods mainly stems from various hyperparameter setups and falls short. Recent research has increasingly focused on enhancing population diversity (Lupu et al. 2021; Zhao et al. 2022; Parker-Holder et al. 2020; Wu et al. 2023; Liu et al. 2021). These methods introduce diversity as a learning objective and incorporate additional auxiliary losses on a population-wide scale. Nonetheless, this population-wide objective increases implementation complexity and necessitates careful hyperparameter tuning to strike a balance between diversity and performance.

Here, we propose a novel perspective on introducing diversity within the population: agents should possess diverse risk preferences. In this context, risk refers to the uncertainty arising from stochastic transitions within the environment, while risk preference encompasses the agent's sensitivity to such uncertainty, typically including risk-seeking, risk-neutral, and risk-averse tendencies. Taking inspiration from the fact that humans are risk-sensitive (Tversky and Kahneman 1992), we argue that the lack of diversity in self-play arises from agents solely optimizing their expected winning rate against opponents without considering other optimization goals that go beyond the expectation, such as higher-order moments of the winning rate distribution. Intuitively, the population of agents should be diverse, including conservative (risk-averse) agents that aim to enhance their worst performance against others, as well as aggressive (risk-seeking) agents that focus on improving their best winning record.

To achieve risk-sensitive learning with minimal modifications, we introduce the *expectile Bellman operator*, which smoothly interpolates between the worst-case and best-case Bellman operators. Building upon this operator, we propose *Risk-sensitive Proximal Policy Optimization* (RPPO), which utilizes a risk level hyperparameter to control the risk preference during policy learning. To strike a balance between

bias and variance in policy learning, we extend the expectile Bellman operator to its multi-step form and adapt the Generalized Advantage Estimation (GAE) (Schulman et al. 2018). It is important to highlight that this extension is non-trivial due to the nonlinearity of the expectile Bellman operator.

RPPO offers a simple and practical approach that can be easily integrated into existing population-based self-play frameworks. We provide an implementation called RPBT, where a population of agents is trained using RPPO with diverse risk preference settings, without introducing additional population-level optimization objectives. The risk preferences of the agents are automatically tuned using exploitation and exploration techniques drawn from Population-based Training (PBT) (Jaderberg et al. 2017). By gathering the experience of agents competing against each other in the diverse population, RPBT effectively addresses the issue of overfitting to specific types of opponents.

Related Work

Self-Play RL Training agents in multi-agent games requires instantiations of opponents the in environment. A solution could be self-play RL, where an agent is trained by playing against its own policy or its past versions. Self-play variants have proven effective in some multi-agent games (Tesauro et al. 1995; Silver et al. 2016; Berner et al. 2019; Vinyals et al. 2019; Jaderberg et al. 2019). A key advantage of self-play is that the competitive multi-agent environment provides the agents with an appropriate curriculum (Bansal et al. 2018; Liu et al. 2019; Baker et al. 2020), which facilitates the emergence of complex and interesting behaviors. However, real-world games often involve non-transitivity and strategic cycles (Czarnecki et al. 2020), and thus self-play cannot produce agents that generalize well to different types of opponents. Population-based self-play (Jaderberg et al. 2019; Garnelo et al. 2021; Yu et al. 2023; Strouse et al. 2021) improves self-play by training a population of agents, all of whom compete against each other. While population-based self-play is able to gather a substantial amount of match experience to alleviate the problem of overfitting to specific opponents, it still requires techniques to introduce diversity among agents to stabilize training and facilitate robustness (Jaderberg et al. 2019).

Population-Based Diversity Population-based methods demonstrate a strong connection with evolutionary algorithms (Mouret and Clune 2015). These population-based approaches have effectively addressed black-box optimization challenges (Loshchilov and Hutter 2016). Their primary advantages encompass the ability to obtain high-performing hyperparameter schedules in a single training run, which leads to great performance across various environments (Liu et al. 2019; Li et al. 2019; Espeholt et al. 2018). Recently, diversity across the population has drawn great interests (Shen et al. 2020; Khadka et al. 2019; Liu et al. 2021, 2022; Wu et al. 2023). Reward randomizations (Tang et al. 2021; Yu et al. 2023) were employed to discover diverse policies in multi-agent games. Other representative works formulated individual rewards and diversity among agents into a multi-objective optimization problem. More specifically,

DvD (Parker-Holder et al. 2020) utilizes the determinant of the kernel matrix of action embedding as a population diversity metric. TrajeDi (Lupu et al. 2021) introduces trajectory diversity by approximating Jensen-Shannon divergence with action discounting kernel. MEP (Zhao et al. 2022) maximizes a lower bound of the average KL divergence within the population. Instead of introducing additional diversity-driven objectives across the entire population, our RPBT approach trains a population of agents with different risk preferences. Each individual agent maximizes its own returns based on its specific risk level, making our method both easy to implement and adaptable for integration into a population-based self-play framework.

Game Theory Fictitious play (Brown 1951) and double oracle (McMahan, Gordon, and Blum 2003) have been studied to achieve approximate Nash equilibrium in normal-form games. Fictitious self-play (FSP) (Heinrich, Lanctot, and Silver 2015) extends fictitious play to extensive-form games. Policy-space response oracle (PSRO) (Lanctot et al. 2017) is a natural generalization of double oracle, where the choices become the policies in meta-games rather than the actions in games. PSRO is a general framework for solving games, maintaining a policy pool and continuously adds the best responses. Our methods fall under PSRO framework at this level. However, in practice, PSRO necessitates the computation of the meta-payoff matrix between policies in the policy pool, which is computationally intensive (McAfee et al. 2020) in real-world games since the policy pool is pretty large. Various improvements (Balduzzi et al. 2019; Perez-Nieves et al. 2021; Liu et al. 2021) have been made upon PSRO by using different metrics based on the meta-payoffs to promote diversity. However, most of these works have confined their experiments to normal-form games or meta-games.

Risk-Sensitive RL Our risk-sensitive methods draw from risk-sensitive and distributional RL, with comprehensive surveys (Bellemare, Dabney, and Rowland 2023) available. Key distributional RL studies (Bellemare, Dabney, and Munos 2017) highlight the value of learning return distributions over expected returns, enabling approximation of value functions under various risk measures like Wang (Müller 1997) and CVaR (Rockafellar, Uryasev et al. 2000; Chow and Ghavamzadeh 2014; Qiu et al. 2021) for generating risk-averse or risk-seeking policies. However, these methods' reliance on discrete samples for risk and gradient estimation increases computational complexity. Sampling-free methods (Tang, Zhang, and Salakhutdinov 2019; Yang et al. 2021; Ying et al. 2022) have been explored for CVaR computation, but CVaR focuses solely on risk aversion, neglecting best-case scenarios, which may not align with competition objectives. A risk-sensitive RL algorithm (Delétang et al. 2021) balances risk aversion and seeking, but assumes gaussian data generation. In contrast, our RPPO algorithm requires no data assumptions and minimal code modifications on PPO to learn diverse risk-sensitive policies.

Preliminary

Problem Definition We consider fully competitive games, which can be regarded as Markov games (Littman 1994). A Markov game for N agents is a partially observable Markov decision process (POMDP), which can be defined by: a set of states \mathcal{S} , a set of observations $\mathcal{O}^1, \dots, \mathcal{O}^N$ of each agent, a set of actions of each agent $\mathcal{A}^1, \dots, \mathcal{A}^N$, a transition function $p(s'|s, a_1, \dots, a_n): \mathcal{S} \times \mathcal{A}^1 \times \dots \times \mathcal{A}^N \rightarrow \Delta(\mathcal{S})$ determining distribution over next states, and a reward function for each agent $r^i: \mathcal{S} \times \mathcal{A}^i \times \mathcal{A}^{-i} \times \mathcal{S} \rightarrow \mathbb{R}$. Each agent chooses its actions based on a stochastic policy $\pi_{\theta_i}: \mathcal{O}^i \rightarrow \Delta(\mathcal{A})$, where θ_i is the policy parameter of agent i .

In the self-play setting, we have control over a single agent known as the main agent, while the remaining agents act as opponents. These opponents are selected from the main agent's checkpoints. The main agent's primary objective is to maximize its expected returns, also referred to as discounted cumulative rewards, denoted as $\mathbb{E}[\sum_{t=0}^T \gamma^t r_t^i]$, where γ represents the discount factor, and T signifies the time horizon. By considering the opponents as part of the environmental dynamics, we can view the multi-agent environment as a single-agent stochastic environment from the main agent's perspective. Consequently, we can employ single-agent RL methods such as PPO to approximate the best response against all opponents. However, it should be noted that the opponents are sampled from the main agent, which itself lacks diversity as it is generated by single-agent RL methods. In this study, we aim to construct a strong and robust policy, and generating diverse policies serves this goal.

Risk-sensitive PPO

In this section, we present our novel RL algorithm, Risk-sensitive Proximal Policy Optimization (RPPO), which involves minimal modifications to PPO. The algorithm is shown in Algorithm 1. RPPO utilizes an expectile Bellman operator that interpolates between a worst-case Bellman operator and a best-case Bellman operator. This allows learning to occur with a specific risk preference, ranging from risk-averse to risk-seeking. Additionally, we extend the expectile Bellman operator into a multi-step form to balance bias and variance in RPPO. The theoretical analysis of the expectile Bellman operator and its multi-step form, as well as a toy example demonstrating the potential of RPPO to learn risk preferences, are presented in the following subsections.

Expectile Bellman Operator

Given a policy π and a risk level hyperparameter $\tau \in (0, 1)$, we consider expectile Bellman operator defined as follows:

$$\mathcal{T}_\tau^\pi V(s) := V(s) + 2\alpha \mathbb{E}_{a,s'} [\tau[\delta]_+ + (1-\tau)[\delta]_-], \quad (1)$$

where α is the step size which we set to $\frac{1}{2\max\{\tau, 1-\tau\}}$, δ refers to $\delta(s, a, s') = r(s, a, s') + \gamma V(s') - V(s)$ which is one-step TD error, $[\cdot]_+ = \max(\cdot, 0)$ and $[\cdot]_- = \min(\cdot, 0)$. This operator draws inspiration from expectile statistics (Newey and Powell 1987; Rowland et al. 2019; Ma

et al. 2022), and thus the name. We can see that the standard Bellman operator is a special case of expectile Bellman operator when $\tau = 1/2$. We have the following theoretical properties to guide the application of expectile Bellman operator in practice. Please refer to the Appendix for the proof.

Proposition 1. For any $\tau \in (0, 1)$, \mathcal{T}_τ^π is a γ_τ -contraction, where $\gamma_\tau = 1 - 2\alpha(1-\gamma) \min\{\tau, 1-\tau\}$.

Proposition 1 guarantees the convergence of the value function in the policy evaluation phase.

Proposition 2. Let V_τ^* denotes the fixed point of \mathcal{T}_τ^π . For any $\tau, \tau' \in (0, 1)$, if $\tau' \geq \tau$, we have $V_{\tau'}^*(s) \geq V_\tau^*(s), \forall s \in \mathcal{S}$.

Proposition 2 guarantees the fixed point of expectile Bellman operator is monotonic with respect to τ .

Proposition 3. Let V_τ^* , V_{best}^* , and V_{worst}^* respectively denote the fixed point of expectile Bellman operator, best-case Bellman operator and worst-case Bellman operator. We have

$$V_\tau^* = \begin{cases} V_{worst}^* & \text{if } \tau \rightarrow 0 \\ V_{best}^* & \text{if } \tau \rightarrow 1. \end{cases} \quad (2)$$

The worst-case Bellman operator and best-case Bellman operator are defined by:

$$\begin{aligned} \mathcal{T}_{best} V(s) &:= \max_{a,s'} [R(s, a, s') + \gamma V(s')], \\ \mathcal{T}_{worst} V(s) &= \min_{a,s'} [R(s, a, s') + \gamma V(s')]. \end{aligned} \quad (3)$$

It is important to highlight the difference between best-case Bellman operator and Bellman optimal operator ($\mathcal{T}V(s) := \max_a [\sum_{s'} (R(s, a, s') + \gamma V(s'))]$). Best-case Bellman operator takes stochastic transitions into account, which is the main source of risk in competitive games when confronting unknown actions of different opponents.

Proposition 3 guarantees expectile Bellman operator approaches best-case Bellman operator as risk level τ approaches 1, and approaches worst-case Bellman operator as τ approaches 0.

Combining Proposition 2 and Proposition 3, we observe that expectile Bellman operator can be used to design an optimization algorithm whose objective is the interpolation between the best-case and the worst-case objective. When $\tau = 1/2$, the objective is equivalent to expected returns, which is the scenario of risk-neutral. And risk level $\tau < 1/2$ and $\tau > 1/2$ represent the risk-averse and risk-seeking cases, respectively. As τ varies, the learning objective varies, and hence diverse risk preferences arise.

We can see the potential of using the expectile Bellman operator for designing risk-sensitive RL algorithms. PPO, widely utilized in self-play, is favored for its ease of use, stable performance, and parallel scalability. By extending PPO with the expectile Bellman operator, we introduce the RPPO algorithm. Moreover, generalizing this operator to other RL algorithms is not difficult. In practice, we define the advantage function as

$$A_\tau^\pi(s, a) := 2\alpha \mathbb{E}_{s'} [\tau[\delta]_+ + (1-\tau)[\delta]_-], \quad (4)$$

Algorithm 1 Risk-sensitive PPO (RPPO)

Input: initial network parameter θ , ϕ , horizon T , update epochs K , risk level τ .

- 1: **for** $i = 1, 2, \dots$ **do**
- 2: Collect trajectories by running policy $\pi_{\theta_{old}}$ in the environment for T time steps.
- 3: Compute λ -variant returns $\hat{V}_{\tau, \lambda}$ according to Eq.10
- 4: Compute advantages $\hat{A}_{\tau, \lambda}$ according to Eq.11
- 5: **for** $k = 1, 2, \dots, K$ **do**
- 6: Update the θ by maximizing the surrogate function with Eq.5 via some gradient algorithms.
- 7: Update the ϕ by minimizing mean-squared error with Eq.6 via some gradient algorithms.
- 8: **end for**
- 9: $\theta_{old} = \theta$
- 10: **end for**

With a batch of data $\{(s_t, a_t, r_t, s_{t+1})\}_{t=0}^{T-1}$ collected by π , we train the policy with the clip variant PPO loss,

$$\mathcal{L}_{\pi}^{\text{CLIP}}(\theta) = \frac{1}{T} \sum_{t=0}^{T-1} \left[\min \left(\omega_t(\theta) \hat{A}_{\tau}(s_t, a_t), \right. \right. \\ \left. \left. \text{clip}(\omega_t(\theta), 1 - \varepsilon, 1 + \varepsilon) \hat{A}_{\tau}(s_t, a_t) \right) \right], \quad (5)$$

where $\omega_t(\theta) = \frac{\pi_{\theta}(a_t | s_t)}{\pi(a_t | s_t)}$ is the importance sampling ratio. Meanwhile, we train the value network with mean squared loss,

$$\mathcal{L}_V(\phi) = \frac{1}{2T} \sum_{t=0}^{T-1} (V_{\phi}(s_t) - \hat{V}_{\tau, t})^2, \quad (6)$$

where ϕ is the parameter of the value network, and $\hat{V}_{\tau, t} := V_{\phi}(s_t) + \hat{A}_{\tau}(s_t, a_t)$ is the target value.

Multi-Step Expectile Bellman Operator

While RPPO is sufficient for providing decent risk preferences, we still require some techniques to balance the bias and variance when estimating the advantage function. The original PPO uses the Generalized Advantage Estimation (GAE) technique (Schulman et al. 2018) to reduce variance by using an exponentially-weighted estimator of the advantage function, at the cost of introducing some bias. In this section, we extend GAE to RPPO along this line of work. However, it is non-trivial for direct incorporation of GAE into RPPO due to the *non-linearity* of the expectile Bellman operator. To address this issue, we define the multi-step expectile Bellman operator as:

$$\mathcal{T}_{\tau, \lambda}^{\pi} V(s) := (1 - \lambda) \sum_{n=1}^{\infty} \lambda^{n-1} (\mathcal{T}_{\tau}^{\pi})^n V(s). \quad (7)$$

We can derive that multi-step expectile Bellman operator still remains the contraction and risk-sensitivity properties. Please refer to Appendix for the proof.

Proposition 4. For any $\tau \in (0, 1)$, $\mathcal{T}_{\tau, \lambda}^{\pi}$ is a $\gamma_{\tau, \lambda}$ -contraction, where $\gamma_{\tau, \lambda} = \frac{(1-\lambda)\gamma_{\tau}}{1-\lambda\gamma_{\tau}}$.

Proposition 5. Let $V_{\tau, \lambda}^*$ denote the fixed point of $\mathcal{T}_{\tau, \lambda}^{\pi}$. For any $\tau, \tau' \in (0, 1)$, if $\tau' \geq \tau$, we have $V_{\tau', \lambda}^*(s) \geq V_{\tau, \lambda}^*(s), \forall s \in S$.

Proposition 6. Let $V_{\tau, \lambda}^*$ denote the fixed point of $\mathcal{T}_{\tau, \lambda}^{\pi}$, we have

$$\lim_{\tau} V_{\tau, \lambda}^* = \begin{cases} V_{\text{worst}}^* & \text{if } \tau \rightarrow 0 \\ V_{\text{best}}^* & \text{if } \tau \rightarrow 1 \end{cases}. \quad (8)$$

Despite the fact that $\mathcal{T}_{\tau, \lambda}^{\pi}$ has the property of contraction, it is hard to be estimated with trajectories. Hence, we introduce another sample form operator

$$\hat{\mathcal{T}}_{\tau}^{\pi, \hat{V}} V(s) := V(s) + 2\alpha \left[\tau[\hat{\delta}]_+ + (1 - \tau)[\hat{\delta}]_- \right], \quad (9)$$

where $\hat{\delta} := r(s, a, s') + \gamma \hat{V}(s') - V(s)$. Here \hat{V} is an estimate of target value. When we choose $\hat{V} = V$ and take the expectation, we recover the $\mathcal{T}_{\tau}^{\pi}(s) = \mathbb{E}_{a \sim \pi(\cdot | s), s' \sim p(\cdot | s, a)} \hat{\mathcal{T}}_{\tau}^{\pi, V} V(s)$. Furthermore, The multi-step sample form operator is

$$\hat{V}_{\tau, \lambda}^{\pi}(s) = (1 - \lambda) \sum_{n=1}^{\infty} \lambda^{n-1} \hat{\mathcal{T}}_{\tau}^{\pi, \hat{V}_n} V(s). \quad (10)$$

where $\hat{V}_n(s) = \hat{\mathcal{T}}_{\tau}^{\pi, \hat{V}_{n-1}} V(s)$ and $\hat{V}_0(s) = V(s)$. Finally, we estimate the advantage as

$$\hat{A}_{\tau, \lambda}^{\pi}(s, a) = \hat{V}_{\tau, \lambda}^{\pi}(s) - V(s). \quad (11)$$

When $\tau = 1/2$, we recover the original form of GAE. However, if $\tau \neq 1/2$, we introduce extra bias¹. The details of computing multi-step expectile Bellman operator in practice are provided in Appendix.

With sample form multi-step expectile Bellman operator in place, we propose the whole RPPO algorithm, as shown in Algorithm 1. The implementation of RPPO requires only a few lines of code modifications based on PPO.

Toy Example

We use a grid world (Delétang et al. 2021) shown in Fig. 1 to illustrate the ability of RPPO to learn diverse risk preferences with different risk levels τ . The grid world consists of an agent, a flag, water, and strong wind. The agent's goal is to reach the flag position within 25 steps while avoiding stepping into the water. When reaching the flag, the agent receives a +1 bonus and is done, but each step spent in the water results in a -1 penalty. In addition, the strong wind randomly pushes the agent in one direction with 50% probability. In this world, risk (uncertainty) comes from the strong wind.

We train 9 RPPO agents with $\tau \in \{0.1, 0.2, \dots, 0.9\}$ to investigate how risk level τ affects the behaviors of policies. We illustrate the state-visitation frequencies computed from the 1,000 rollouts, as shown in Fig. 1. Three modes

¹The fixed points of $\mathcal{T}_{\tau, \lambda}^{\pi}$ and \mathcal{T}_{τ}^{π} are the same one. However, the sample form operator introduces extra bias for value estimation, which means the fixed points of $\hat{\mathcal{T}}_{\tau, \lambda}^{\pi}$ and $\mathcal{T}_{\tau, \lambda}^{\pi}$ are not the same one. It is caused by the non-linearity of the operator with respect to the noise.

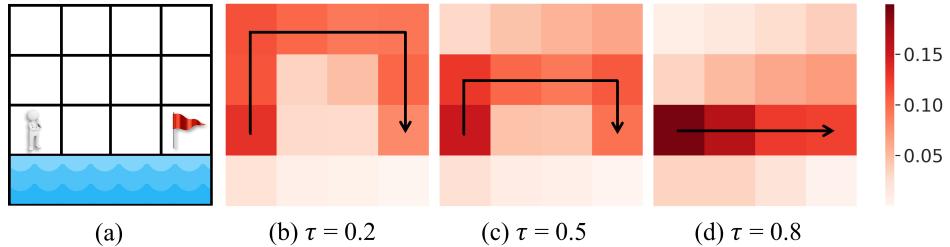


Figure 1: Experiments of the toy example. (a) Illustration of the toy example. The task is to pick up the bonus located at the flag while avoiding the penalty of stepping into water in a 4×4 grid world. A strong wind pushes the agent into a random direction 50% of the time. (b), (c) and (d) States visitation frequencies for the three policies that have different risk preferences. The black arrow indicates the deterministic policy when there is no wind.

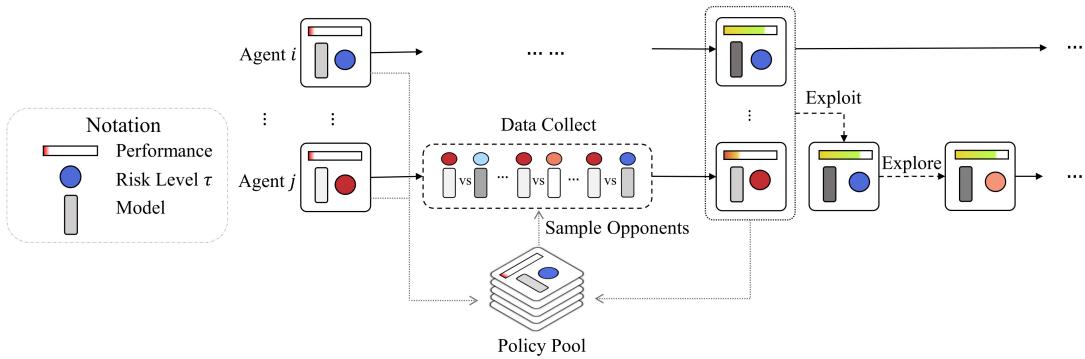


Figure 2: The framework of RPBT. We train a population of agents with different initialization of risk levels. During a training round, each agent spawns a number of subprocesses to collect data against randomly selected opponents from the policy pool. We use RPPO to update models under the specific risk level. The policies updated are added to the policy pool. If an agent in the population is under-performing, it will *exploit* (copy) the model parameters and the risk level τ of a better-performing agent, and it will *explore* a new risk level by adding a perturbation to the better-performing agent's risk level for the following training.

of policies emerge here: risk-averse ($\tau \in \{0.2, 0.3, 0.4\}$), taking the longest path away from the water; risk-neutral ($\tau \in \{0.5, 0.6, 0.7\}$), taking the middle path, and risk-seeking ($\tau \in \{0.8, 0.9\}$), taking the shortest route along the water. Interestingly, agents with $\tau = 0.1$ are too conservative and do not always reach the flag, which is consistent with risk-averse behavior styles.

RPBT

In this section, we present our population-based self-play method based on RPPO, which we refer to as RPBT, as illustrated in Fig. 2. Our proposed RPBT stabilizes the learning process by concurrently training a diverse population of agents who learn by playing with each other. Moreover, different agents in the population have varying risk preferences in order to promote the learned policies to be diverse and not homogeneous. In contrast to previous work (Parker-Holder et al. 2020; Lupu et al. 2021; Zhao et al. 2022), RPBT does not include population-level diversity objectives. Instead, each agent acts and learns to maximize its individual rewards under the specific risk preference settings. This makes RPBT more easier to implement and scale.

In our RPBT, the value of the risk level $\tau \in (0, 1)$ controls

the agent's risk preference. Since it is impossible to cover all possible risk levels within the population, we dynamically adjust the risk levels during training, particularly for those with poor performance. We adopt the exploitation and exploration steps in PBT (Jaderberg et al. 2017) to achieve this auto-tuning goal. In the exploitation step, a poorly performing agent can directly copy the model parameters and risk level τ from a well-performing agent to achieve equal performance. Then in the exploration step, the newly copied risk level is randomly perturbed by noise to produce a new risk level. In our practice, we introduce a simple technique:

- *Exploitation.* We rank all agents in the population according to ELO score, indicating performance. If an agent's ELO score falls below a certain threshold, it is considered an underperforming agent, and its model parameters and risk level will be replaced by those from a superior agent.
- *Exploration.* The risk level of the underperforming agent is further updated by adding a noise term varying between -0.2 and 0.2.

We find that this technique allows RPBT to explore almost all possible values of τ during training. At the later stage of training, the values of τ will converge to an interval with



Figure 3: Illustration of diverse risk preferences in Slimevolley. Risk-seeking agent (left) stands farther from the fence and hit the ball at a lower angle while risk-averse agent (right) does the opposite.



Figure 4: Illustration of diverse risk preferences in Sumoants. The risk-seeking agent (red) constantly attacks, while the risk-averse agent (green) assumes a defensive stance.

better performance while maintaining diversity and avoiding harmful values.

Additionally, we treat all agents of the population and their historical copies as a policy pool from which opponents are uniformly sampled for self-play. Actually, this approach of sampling opponents is commonly referred to as FSP (Vinyals et al. 2019), which ensures that an agent must be able to defeat random old versions of any other agent in the population in order to continue learning adaptively (Bansal et al. 2018). Since any approach of sampling opponents can be adapted to our method, we do not dwell on more complicated opponent sampling techniques (Vinyals et al. 2019; Berner et al. 2019). Moreover, we utilize FSP for all our experiments to ensure fair comparisons.

Fig. 2 illustrated one training round of RPBT. Each training round consists of uniformly sampling opponents from the policy pool, collecting game experience, updating models using RPPO, and performing exploitation and exploration on risk levels. After training is ended, we select the agent with the highest ELO score from the population to serve as the evaluation agent.

Experiments

In our experiments, we aim to answer the following questions: **Q1**, can RPPO generate policies with diverse risk preferences in competitive games? **Q2**, how does RPBT perform compared to other methods in competitive games? Some results of the ablation experiments on RPBT are presented in the Appendix. All the experiments are conducted with one 64-core CPU and one GeForce RTX 3090 GPU.

Environment Setup

We consider two competitive multi-agent benchmarks: **Slimevolley** (Ha 2020) and **Sumoants** (Al-Shedivat et al. 2018). Slimevolley is a two-agent volleyball game where the action space is discrete. The goal of each agent is to land the ball on the opponent’s field, causing the opponent to lose lives. Sumoants is a two-agent game based on MuJoCo where the action space is continuous. Two ants compete in

a square area, aiming to knock the other agent to the ground or push it out of the ring. More details about the two benchmarks are given in Appendix.

Diverse Risk Preferences Illustration

We trained RPPO agents via self-play, using various risk levels. As for Slimevolley, we pit the risk-seeking agent with $\tau = 0.9$ against the risk-averse agent with $\tau = 0.1$, as shown in Fig. 3. We observe that the risk-seeking agent prefers to adjust its position to spike the ball so that the height of the ball is lower, while the risk-averse agent takes a more conservative strategy and just focuses on catching the ball. To further quantify this phenomenon, we pit the three agents $\tau \in \{0.1, 0.5, 0.9\}$ versus the $\tau = 0.5$ agent and compute two metrics average over 200 rollouts: the distance to the fences (denoted as d) and hitting angle (denoted as β). As τ increases, β shows a monotonically decreasing trend, and d shows a monotonically increasing trend. Because the physics engine of the volleyball game is a bit “dodgy”, the agent can hit the ball with a low height only when standing far away from the fence.

As for Sumoants, a typical play between a risk-seeking agent with $\tau = 0.7$ and a risk-averse agent with $\tau = 0.3$ is illustrated in Fig. 4. We observed that the risk-averse agent tends to maintain a defensive stance: four legs spread out, lowering its center of gravity and holding the floor tightly to keep it as still as possible. In contrast, the risk-seeking agent frequently attempts to attack. We also provide multiple videos for further observation on our code website.

(a) Slimevolley			(b) Sumoants			
	RPBT	PP	SP	RPBT	PP	SP
RPBT	-	56%	64%	-	63%	67%
PP	44%	-	56%	37%	-	53%
SP	36%	44%	-	33%	47%	-

Table 1: RPBT performing against basic baselines.

(a) Slimevolley						(b) Sumoants								
	RPBT	MEP	TrajeDi	DvD	RR	PSRO		RPBT	MEP	TrajeDi	DvD	RR	PSRO	
RPBT(ours)	-	64%	59%	48%	60%	53%			58%	63%	54%	56%	53%	
MEP	36%	-	46%	32%	39%	32%		MEP	42%	-	54%	53%	51%	50%
TrajeDi	41%	53%	-	38%	38%	45%		TrajeDi	37%	46%	-	44%	48%	37%
DvD	52%	63%	61%	-	58%	63%		DvD	46%	47%	56%	-	52%	43%
RR	37%	52%	63%	42%	-	54%		RR	44%	49%	52%	48%	-	43%
PSRO	48%	68%	55%	37%	46%	-		PSRO	47%	50%	63%	57%	57%	-

Table 2: RPBT performing against different methods.

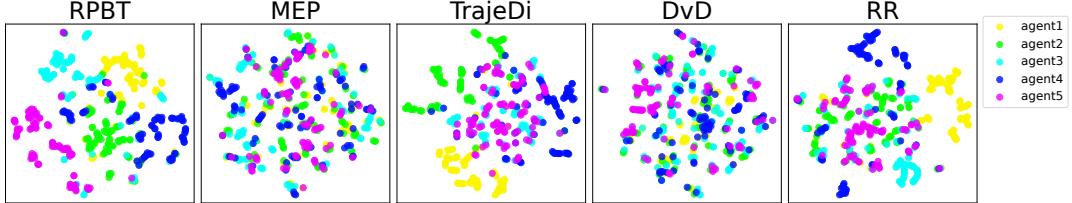


Figure 5: Comparing diversity of different methods with a population size of 5.

Results Compared With Other Methods

We trained RPBT with population size 5 and set initial risk levels to $\{0.1, 0.4, 0.5, 0.6, 0.9\}$ for all the experiments. To ensure fairness in our comparison, all the baselines were integrated with PPO and set to the same population size if necessary. See Appendix for more implementation details. The baseline methods includes:

- Basic baselines: self-play (**SP**), where agents learn solely through competing with themselves, and population-based self-play (**PP**), where a population of agents are co-trained through randomly competing against one another.
- Population-based methods: **MEP** (Zhao et al. 2022), **TrajeDi** (Lupu et al. 2021), **DvD** (Parker-Holder et al. 2020), and **RR** (Reward Randomization, similar ideas with (Tang et al. 2021)).
- Game-theoretic method: **PSRO** (Lanctot et al. 2017).

For each method, we trained 3 runs using different random seeds and selected the one with the highest ELO score for evaluation. Then, we calculated the win rate matrix between the methods, recording the average win rate corresponding to 200 plays. Tab.1 shows the results among RPBT and basic methods. We observe that RPBT exhibits superior performance compared to the baselines. Therefore, it is crucial to integrate diversity-oriented methods like RPBT in self-play. Tab.2 shows the results among RPBT, population-based methods, and game-theoretic methods. We observed that RPBT outperforms MEP, TrajeDi, and RR. Moreover, RPBT achieves performance comparable to that of DvD in Slimevolley but better in Sumoants. Furthermore, RPBT performs slightly better than PSRO. However, PSRO is more computationally intensive (While PSRO requires 34 hours for a single run, other methods only require almost 12 hours).

In order to further compare the diversity in population-based methods, we let each agent in the population play

against itself. We then extract the first 100 states from the game trajectories and use t-SNE (Van der Maaten and Hinton 2008) to reduce the states to 2 dimensions for visualizing. Fig.5 shows the results. We observed that the trajectories of agents in the RPBT population are distinct from each other. In comparison, TrajeDi and RR exhibit a certain level of diversity, whereas MEP and DvD show no indications of diversity. Notably, these population-based methods require the introduction of an additional diversity objective at the population level. This causes higher complexity to implementation as well as hyperparameter tuning due to the trade-off between diversity and performance. Overall, RPBT is a favorable approach for both learning diverse behaviors and enhancing performance.

Conclusion

Observing that learning to be diverse is essential for addressing the overfitting problem in self-play, we propose that diversity could be introduced in terms of risk preferences from different agents. Specifically, we propose a novel Risk-sensitive PPO (RPPO) approach to learn policies that align with desired risk preferences. Furthermore, we incorporate RPPO into a population-based self-play framework, which is easy to implement, thereby introducing the RPBT approach. In this approach, a population of agents, with diverse and dynamically adjusted risk preferences, compete with one another. We showed that our method can generate diverse modes of policies and achieve comparable or superior performance over existing methods. To the best of our knowledge, this is the first work that uses a risk-oriented approach to train a diverse population of agents. Future directions include developing a meta-policy that can adaptively select the optimal risk level according to the opponent’s strategy and exploring the applicability of RPPO to safe reinforcement learning scenarios.

Acknowledgments

This work was supported by National Natural Science Foundation of China under Grant No. 62192751, in part by the National Science and Technology Innovation 2030 - Major Project (Grant No. 2022ZD0208804), in part by Key R&D Project of China under Grant No. 2017YFC0704100, the 111 International Collaboration Program of China under Grant No. BP2018006, and in part by the BNRist Program under Grant No. BNR2019TD01009, the National Innovation Center of High Speed Train R&D project (CX/KJ-2020-0006), in part by the InnoHK Initiative, The Government of HKSAR; and in part by the Laboratory for AI-Powered Financial Technologies.

References

Al-Shedivat, M.; Bansal, T.; Burda, Y.; Sutskever, I.; Mordatch, I.; and Abbeel, P. 2018. Continuous Adaptation via Meta-Learning in Nonstationary and Competitive Environments. In *International Conference on Learning Representations*.

Baker, B.; Kanitscheider, I.; Markov, T.; Wu, Y.; Powell, G.; McGrew, B.; and Mordatch, I. 2020. Emergent Tool Use From Multi-Agent Autocurricula. In *International Conference on Learning Representations*.

Balduzzi, D.; Garnelo, M.; Bachrach, Y.; Czarnecki, W.; Perolat, J.; Jaderberg, M.; and Graepel, T. 2019. Open-ended learning in symmetric zero-sum games. In *International Conference on Machine Learning*, 434–443. PMLR.

Bansal, T.; Pachocki, J.; Sidor, S.; Sutskever, I.; and Mordatch, I. 2018. Emergent Complexity via Multi-Agent Competition. In *International Conference on Learning Representations*.

Bellemare, M. G.; Dabney, W.; and Munos, R. 2017. A distributional perspective on reinforcement learning. In *International conference on machine learning*, 449–458. PMLR.

Bellemare, M. G.; Dabney, W.; and Rowland, M. 2023. *Distributional Reinforcement Learning*. MIT Press.

Berner, C.; Brockman, G.; Chan, B.; Cheung, V.; Debiak, P.; Dennison, C.; Farhi, D.; Fischer, Q.; Hashme, S.; Hesse, C.; Józefowicz, R.; Gray, S.; Olsson, C.; Pachocki, J. W.; Petrov, M.; de Oliveira Pinto, H. P.; Raiman, J.; Salimans, T.; Schlatte, J.; Schneider, J.; Sidor, S.; Sutskever, I.; Tang, J.; Wolski, F.; and Zhang, S. 2019. Dota 2 with Large Scale Deep Reinforcement Learning. *ArXiv preprint*, abs/1912.06680.

Brown, G. W. 1951. Iterative Solution of Games by Fictitious Play. In Koopmans, T. C., ed., *Activity Analysis of Production and Allocation*. New York: Wiley.

Chow, Y.; and Ghavamzadeh, M. 2014. Algorithms for CVaR optimization in MDPs. *Advances in neural information processing systems*, 27.

Czarnecki, W. M.; Gidel, G.; Tracey, B.; Tuyls, K.; Omidshafiei, S.; Balduzzi, D.; and Jaderberg, M. 2020. Real world games look like spinning tops. *Advances in Neural Information Processing Systems*, 33: 17443–17454.

Delétang, G.; Grau-Moya, J.; Kunesch, M.; Genewein, T.; Brekelmans, R.; Legg, S.; and Ortega, P. A. 2021. Model-free risk-sensitive reinforcement learning. *arXiv preprint arXiv:2111.02907*.

Espeholt, L.; Soyer, H.; Munos, R.; Simonyan, K.; Mnih, V.; Ward, T.; Doron, Y.; Firoiu, V.; Harley, T.; Dunning, I.; et al. 2018. Impala: Scalable distributed deep-rl with importance weighted actor-learner architectures. In *International conference on machine learning*, 1407–1416. PMLR.

Garnelo, M.; Czarnecki, W. M.; Liu, S.; Tirumala, D.; Oh, J.; Gidel, G.; van Hasselt, H.; and Balduzzi, D. 2021. Pick your battles: Interaction graphs as population-level objectives for strategic diversity. *arXiv preprint arXiv:2110.04041*.

Ha, D. 2020. Slime Volleyball Gym Environment.

Heinrich, J.; Lanctot, M.; and Silver, D. 2015. Fictitious Self-Play in Extensive-Form Games. In Bach, F.; and Blei, D., eds., *Proceedings of the 32nd International Conference on Machine Learning*, volume 37 of *Proceedings of Machine Learning Research*, 805–813. Lille, France: PMLR.

Jaderberg, M.; Czarnecki, W. M.; Dunning, I.; Marrs, L.; Lever, G.; Castañeda, A. G.; Beattie, C.; Rabinowitz, N. C.; Morcos, A. S.; Ruderman, A.; Sonnerat, N.; Green, T.; Deason, L.; Leibo, J. Z.; Silver, D.; Hassabis, D.; Kavukcuoglu, K.; and Graepel, T. 2019. Human-level performance in 3D multiplayer games with population-based reinforcement learning. *Science*, 364(6443): 859–865.

Jaderberg, M.; Dalibard, V.; Osindero, S.; Czarnecki, W. M.; Donahue, J.; Razavi, A.; Vinyals, O.; Green, T.; Dunning, I.; Simonyan, K.; et al. 2017. Population based training of neural networks. *arXiv preprint arXiv:1711.09846*.

Khadka, S.; Majumdar, S.; Nassar, T.; Dwiel, Z.; Tumer, E.; Miret, S.; Liu, Y.; and Tumer, K. 2019. Collaborative evolutionary reinforcement learning. In *International conference on machine learning*, 3341–3350. PMLR.

Lanctot, M.; Zambaldi, V.; Gruslys, A.; Lazaridou, A.; Tuyls, K.; Pérolat, J.; Silver, D.; and Graepel, T. 2017. A unified game-theoretic approach to multiagent reinforcement learning. *Advances in neural information processing systems*, 30.

Li, A.; Spyra, O.; Perel, S.; Dalibard, V.; Jaderberg, M.; Gu, C.; Budden, D.; Harley, T.; and Gupta, P. 2019. A generalized framework for population based training. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, 1791–1799.

Littman, M. L. 1994. Markov Games as a Framework for Multi-Agent Reinforcement Learning. In *Machine Learning Proceedings 1994*, 157–163. Elsevier. ISBN 978-1-55860-335-6.

Liu, S.; Lever, G.; Heess, N.; Merel, J.; Tunyasuvunakool, S.; and Graepel, T. 2019. Emergent Coordination Through Competition. In *International Conference on Learning Representations*.

Liu, X.; Jia, H.; Wen, Y.; Hu, Y.; Chen, Y.; Fan, C.; Hu, Z.; and Yang, Y. 2021. Towards unifying behavioral and response diversity for open-ended learning in zero-sum games. *Advances in Neural Information Processing Systems*, 34: 941–952.

Liu, Z.; Yu, C.; Yang, Y.; Wu, Z.; Li, Y.; et al. 2022. A Unified Diversity Measure for Multiagent Reinforcement Learning. *Advances in Neural Information Processing Systems*, 35: 10339–10352.

Loshchilov, I.; and Hutter, F. 2016. CMA-ES for hyper-parameter optimization of deep neural networks. *arXiv preprint arXiv:1604.07269*.

Lupu, A.; Cui, B.; Hu, H.; and Foerster, J. 2021. Trajectory Diversity for Zero-Shot Coordination. In *Proceedings of the 38th International Conference on Machine Learning*, 7204–7213. PMLR.

Ma, X.; Yang, Y.; Hu, H.; Yang, J.; Zhang, C.; Zhao, Q.; Liang, B.; and Liu, Q. 2022. Offline Reinforcement Learning with Value-based Episodic Memory. In *International Conference on Learning Representations*.

McAleer, S.; Lanier, J. B.; Fox, R.; and Baldi, P. 2020. Pipeline psro: A scalable approach for finding approximate nash equilibria in large games. *Advances in neural information processing systems*, 33: 20238–20248.

McMahan, H. B.; Gordon, G. J.; and Blum, A. 2003. Planning in the Presence of Cost Functions Controlled by an Adversary. In *Proceedings of the Twentieth International Conference on International Conference on Machine Learning*, ICML'03, 536–543. AAAI Press. ISBN 1577351894.

Mouret, J.-B.; and Clune, J. 2015. Illuminating search spaces by mapping elites. *arXiv preprint arXiv:1504.04909*.

Müller, A. 1997. Integral Probability Metrics and Their Generating Classes of Functions. *Advances in Applied Probability*, 29(2): 429–443.

Newey, W. K.; and Powell, J. L. 1987. Asymmetric least squares estimation and testing. *Econometrica: Journal of the Econometric Society*, 819–847.

Parker-Holder, J.; Pacchiano, A.; Choromanski, K. M.; and Roberts, S. J. 2020. Effective diversity in population based reinforcement learning. *Advances in Neural Information Processing Systems*, 33: 18050–18062.

Perez-Nieves, N.; Yang, Y.; Slumbers, O.; Mguni, D. H.; Wen, Y.; and Wang, J. 2021. Modelling behavioural diversity for learning in open-ended games. In *International Conference on Machine Learning*, 8514–8524. PMLR.

Qiu, W.; Wang, X.; Yu, R.; Wang, R.; He, X.; An, B.; Obraztsova, S.; and Rabinovich, Z. 2021. RMIX: Learning risk-sensitive policies for cooperative reinforcement learning agents. *Advances in Neural Information Processing Systems*, 34: 23049–23062.

Rockafellar, R. T.; Uryasev, S.; et al. 2000. Optimization of conditional value-at-risk. *Journal of risk*, 2: 21–42.

Rowland, M.; Dadashi, R.; Kumar, S.; Munos, R.; Bellemare, M. G.; and Dabney, W. 2019. Statistics and Samples in Distributional Reinforcement Learning. In *International Conference on Machine Learning*, 5528–5536. PMLR.

Schulman, J.; Moritz, P.; Levine, S.; Jordan, M.; and Abbeel, P. 2018. High-Dimensional Continuous Control Using Generalized Advantage Estimation. *arXiv:1506.02438 [cs]*.

Shen, R.; Zheng, Y.; Hao, J.; Meng, Z.; Chen, Y.; Fan, C.; and Liu, Y. 2020. Generating Behavior-Diverse Game AIs with Evolutionary Multi-Objective Deep Reinforcement Learning. In *Twenty-Ninth International Joint Conference on Artificial Intelligence*, volume 4, 3371–3377.

Silver, D.; Huang, A.; Maddison, C. J.; Guez, A.; Sifre, L.; Van Den Driessche, G.; Schrittwieser, J.; Antonoglou, I.; Panneershelvam, V.; and Lanctot, M. 2016. Mastering the Game of Go with Deep Neural Networks and Tree Search. *nature*, 529(7587): 484–489.

Strouse, D.; McKee, K.; Botvinick, M.; Hughes, E.; and Everett, R. 2021. Collaborating with humans without human data. *Advances in Neural Information Processing Systems*, 34: 14502–14515.

Tang, Y. C.; Zhang, J.; and Salakhutdinov, R. 2019. Worst cases policy gradients. *arXiv preprint arXiv:1911.03618*.

Tang, Z.; Yu, C.; Chen, B.; Xu, H.; Wang, X.; Fang, F.; Du, S. S.; Wang, Y.; and Wu, Y. 2021. Discovering Diverse Multi-Agent Strategic Behavior via Reward Randomization. In *International Conference on Learning Representations*.

Tesrauro, G.; et al. 1995. Temporal difference learning and TD-Gammon. *Communications of the ACM*, 38(3): 58–68.

Tversky, A.; and Kahneman, D. 1992. Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and uncertainty*, 5: 297–323.

Van der Maaten, L.; and Hinton, G. 2008. Visualizing data using t-SNE. *Journal of machine learning research*, 9(11).

Vinyals, O.; Babuschkin, I.; Czarnecki, W. M.; Mathieu, M.; Dudzik, A.; Chung, J.; Choi, D. H.; Powell, R.; Ewalds, T.; Georgiev, P.; et al. 2019. Grandmaster level in StarCraft II using multi-agent reinforcement learning. *Nature*, 575(7782): 350–354.

Wu, S.; Yao, J.; Fu, H.; Tian, Y.; Qian, C.; Yang, Y.; FU, Q.; and Wei, Y. 2023. Quality-Similar Diversity via Population Based Reinforcement Learning. In *The Eleventh International Conference on Learning Representations*.

Yang, Q.; Simão, T. D.; Tindemans, S. H.; and Spaan, M. T. 2021. WCSAC: Worst-case soft actor critic for safety-constrained reinforcement learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, 10639–10646.

Ying, C.; Zhou, X.; Su, H.; Yan, D.; Chen, N.; and Zhu, J. 2022. Towards safe reinforcement learning via constraining conditional value-at-risk. *arXiv preprint arXiv:2206.04436*.

Yu, C.; Gao, J.; Liu, W.; Xu, B.; Tang, H.; Yang, J.; Wang, Y.; and Wu, Y. 2023. Learning Zero-Shot Cooperation with Humans, Assuming Humans Are Biased. In *The Eleventh International Conference on Learning Representations*.

Zhao, R.; Song, J.; Yuan, Y.; Haifeng, H.; Gao, Y.; Wu, Y.; Sun, Z.; and Wei, Y. 2022. Maximum Entropy Population-Based Training for Zero-Shot Human-AI Coordination. *arxiv:2112.11701*.