

Efficient Reinforcement Learning for Zero-Shot Coordination in Evolving Games

Bingyu Hui, Lebin Yu, Quanming Yao, Yunpeng Qu, Xudong Zhang, Jian Wang*

Department of Electronic Engineering, Beijing National Research Center for Information Science and Technology, Tsinghua University, Beijing 100084, China

huiy23@mails.tsinghua.edu.cn, paladinee15@gmail.com, qyaoaa@tsinghua.edu.cn, qyp21@mails.tsinghua.edu.cn, zhangxd@tsinghua.edu.cn, jian-wang@tsinghua.edu.cn

Abstract

Zero-shot coordination (ZSC), a key challenge in multi-agent game theory, has become a hot topic in reinforcement learning (RL) research recently, especially in complex evolving games. It focuses on the generalization ability of agents, requiring them to coordinate well with collaborators from a diverse, potentially evolving, pool of partners that are not seen before without any fine-tuning. Population-based training, which approximates such an evolving partner pool, has been proven to provide good zero-shot coordination performance; nevertheless, existing methods are limited by computational resources, mainly focusing on optimizing diversity in small populations while neglecting the potential performance gains from scaling population size. To address this issue, this paper proposes the Scalable Population Training (ScaPT), an efficient RL training framework comprising two key components: a meta-agent that efficiently realizes a population by selectively sharing parameters across agents, and a mutual information regularizer that guarantees population diversity. To empirically validate the effectiveness of ScaPT, this paper evaluates it along with representational frameworks in Hanabi cooperative game and confirms its superiority.

Introduction

As Multi-Agent Reinforcement Learning (MARL) continues to achieve remarkable success in a wide range of complex games and domains, its scalability is increasingly hindered by the greedy growth of computational resource demands, in the scenarios like agent MOBA games (Gao et al. 2023), robotic soccer (Zhu et al. 2024), and intelligent manufacturing systems (Do, Jang, and Kim 2025) owing to the demand for training a diverse and evolving set of agents. These challenges highlight the need for more scalable and resource-efficient MARL frameworks.

We address this problem in the context of Zero-Shot Coordination (ZSC)—a setting where agents must collaborate with unseen partners. This scenario is central to what we term “evolving games”: environments where the pool of potential partners is not fixed, but rather represents a continuously shifting or evolving set of strategies. While self-play (Lowe et al. 2017) is widely used for training cooperation, it




Figure 1: The diagram of different training paradigms in evolving games.

often leads to overfitting, limiting agents’ ability to generalize. To overcome this, Hu et al. (2020) introduced the ZSC problem, where agents are evaluated on their ability to coordinate without prior exposure to their partners. Population-based training (Jaderberg et al. 2017) offers a promising solution by exposing agents to a diverse set of behaviors during training, thus encouraging more generalizable coordination strategies (Charakorn, Manoonpong, and Dilokthanakul 2023). In principle, if the training population covers all potential partner strategies, ZSC would be effectively solved. However, approaching this ideal requires training a large and diverse population of agents, which is itself constrained by limited computational resources, hindering the development of efficient and scalable RL methods.

Recent methods focus on improving Zero-Shot Coordination (ZSC) performance by enhancing population diversity, which serves as a key proxy for strategic variation in these evolving games. For example, Trajectory Diversity for Zero-Shot Coordination (Lupu et al. 2021) maximizes trajectory-level diversity to foster generalizable strategies. Similarly, Maximum Entropy Population-Based Training (Zhao et al. 2023) uses entropy maximization to broaden behavioral coverage and improve human-AI coordination. Heterogeneous Multi-Agent Zero-Shot Coordination by Coevolution (Xue et al. 2024) tackles real-world heterogeneity by co-evolving two distinct populations through iterative pairing and selection, enabling robust cooperation with diverse partners.

However, constrained by limited computational resources, existing studies typically maintain relatively small population sizes and attempt to improve the zero-shot coordination performance of trained agents by optimizing

*Corresponding author

objective functions to encourage policy “diversity” within these small populations. Nevertheless, focusing solely on individual-level behavioral diversity while neglecting the potential impact of population scale may limit the overall effectiveness of population-based methods in truly capturing the breadth of an evolving game’s strategy space. In fact, the constrained population size could be one of the critical bottlenecks restricting current ZSC performance.

To address this challenge, we propose an efficient RL with scalable population-based training framework: ScaPT. Figure 1 illustrates the differences among self-play, common population-based training, and our proposed framework in evolving games. Our goal is to enable ZSC training to scale more efficiently to larger populations, mitigating the constraints imposed by computational resources. To achieve this, our framework is designed to meet two key criteria: (1) alleviate the coupling between population size and computational and parameter overhead, thereby mitigating the linear growth of resource consumption with increasing population size; (2) maintain comparable or even superior ZSC performance relative to conventional population-based methods.

To satisfy these requirements, ScaPT employs a hierarchical meta-agent architecture that simulates arbitrarily large populations within a single network, thereby avoiding redundant training across individuals and effectively suppressing the linear growth of computational costs with population size. To promote behavioral diversity, ScaPT maximizes the conditional mutual information between the actions generated by sub-decision modules under given observations and the modules themselves.

In experiments conducted on two standard ZSC benchmark environments, matrix games and Hanabi, ScaPT achieves significantly better performance than existing algorithms, primarily due to its scalable population design that allows it to leverage larger agent pools under the same memory resources constraints. Furthermore, when comparing different algorithms under the same scalable population framework, ScaPT demonstrates outstanding zero-shot coordination game capabilities.

Our contributions are summarized below:

- We introduce ScaPT, a scalable population training framework for zero-shot coordination that surpasses the population size limits of conventional methods.
- ScaPT designs the scalable population using a hierarchical meta-agent architecture to simulate large populations, achieving a high degree of decoupling between population size and computational resources, while employing conditional mutual information to ensure diversity within the scalable population.
- Under constrained memory budgets, ScaPT outperforms prior methods by leveraging significantly larger agent populations, achieving superior zero-shot coordination performance with the same or lower computational cost.

Related Work

Self-play (Yu et al. 2022) is widely used in MARL, and can efficiently train cooperative agents, but it typically fails under zero-shot coordination. To address this issue, recent

works introduce additional reasoning abilities, such as break symmetries of the task to keep agents from learning specified strategies (Hu et al. 2020; Treutlein et al. 2021; Muglich et al. 2022), conduct multi-level reasoning for higher-level consensus (Cui et al. 2021; Hu et al. 2021) and predict partners’ actions (Lucas and Allen 2022; Yan et al. 2024). However, these approaches often lead to algorithm-level consensus or overfitting to a limited set of partners.

Population-based training enhances generalization by exposing a main agent to diverse partners. Common approaches include reducing the collaboration scores of different partner agents within the population to make them behave differently (Charakorn, Manoonpong, and Dilokthanakul 2023; Rahman et al. 2023), improving trajectory diversity of agents (Lupu et al. 2021) and increasing policy entropy of the population (Zhao et al. 2023). However, most approaches train distinct networks per agent or require differentiable action distributions (Lupu et al. 2021; Zhao et al. 2023), limiting efficiency and applicability; In contrast, ScaPT employs a meta-agent for efficient population training, and designs a generic mutual information term to guarantee population divergence. This differs from MI-based skill-discovery methods such as DIAYN (Eysenbach et al. 2018) and DADS (Sharma et al. 2019), which increase intra-agent skill diversity rather than inter-agent policy diversity.

Notably, the meta-agent in ScaPT is different from the agents in meta-RL (Nagabandi et al. 2018; Gupta et al. 2018). Traditional meta-RL aims to train agents that can quickly adapt to various tasks, the diversity of which is innate and invariable. In contrast, ScaPT’s meta-agent is designed to exhibit various policies, the diversity of which is variable and what we seek to augment.

Scalable Population Training Framework for Zero-Shot Coordination

We present ScaPT, a scalable population-training framework that achieves diverse zero-shot coordination.

Hierarchical Meta-Agent for Scalable Population

A major drawback of population training is its high resource demand. Inspired by meta-RL and multi-task learning (Ma et al. 2018; Sagiurova, Kuratov, and Burtsev 2025), task-related parameters can be shared, while behavior-specific submodules are optimized independently.

Based on the aforementioned idea, we design a hierarchical meta-agent—a simple yet effective architecture—to realize a scalable population. Fig. 2 show the details of the proposed architecture. Here, h and c denote the hidden and cell states of the LSTM, respectively, and o^t is the observation input from the environment. The variable v^t represents the estimated state value. The submodules P_1, P_2, \dots, P_K are used to simulate different individuals in the population. Each of P_i ’s output $a_{u=i}^t$ serves as the policy output for the i -th agent in the population, where K denotes the population size.

As can be seen, the hierarchical architecture of the meta-agent greatly reduces the parameters that need to be opti-

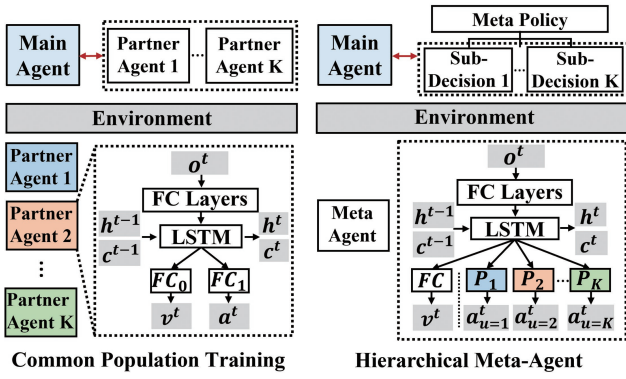


Figure 2: Comparison of common population and hierarchical meta-agent population

mized (K complete sets \rightarrow 1 complete set + K subsets). In this way, one can simulate populations of arbitrary size simply by adjusting the number of subsets, without the need to independently train each agent from scratch. In our implementation, we simply define the P_i as a two-layer fully connected network; however, this choice is not mandatory. More complex architectures could potentially yield further performance improvements, though this comes with higher computational cost.

Conditional Mutual Information Maximization

If all the agents in a population have similar policies, the best-response agent can only learn to cooperate with partners of one kind of policy, and in this way the advantage of population training disappears (See Ablation study section for details). Therefore, increasing the diversity of population (i.e. different agents in the population behave variously) has always been a key concern in this research area. It is noteworthy that a diverse population means different agents in the population act differently, and this can be achieved by making the meta-agent output distinct actions with different sub-decision modules given an observation and history trajectory. This operation can be formulated as maximizing the conditional mutual information:

$$I(A; U|H) = \int \int \int p(a, u, h) \log \frac{p(h)p(a, u, h)}{p(u, h)p(a, h)} du ds da \quad (1)$$

where H represents observation input (containing current observation and historical trajectory used for decision), U represents the index of sub-decision module that outputs action A , and $p(a, u, h)$ is the joint probability density function. Considering that reinforcement learning frameworks commonly estimate integrals using the Monte Carlo method, that is, sampling transitions from replay buffer to calculate, the unbiased estimation of mutual information $\hat{I}(A; U|H)$ can be written as:

$$\begin{aligned} \hat{I}(A; U|H) &= \frac{1}{N} \sum_{j=1}^N \log \frac{p(a_j|u_j, h_j)}{p(a_j|h_j)} \\ &= \frac{1}{N} \sum_{j=1}^N \left[\log p(a_j|u_j, h_j) \right. \\ &\quad \left. - \log \sum_{i=1}^K p(u_i|h_j) \cdot p(a_j|u_i, h_j) \right] \end{aligned} \quad (2)$$

where K is the total number of sub-decision modules, N is the number of transitions, a_j, u_j, h_j are action, sub-decision module index and observation input of the j -th transition, and $p(a_j|u_j, h_j)$ is the conditional probability that an agent takes action a_j given u_j, h_j . For brevity, use I_j to denote the j -th term in $\hat{I}(A; U|H)$. The equation is: $I_j := \log p(a_j|u_j, h_j) - \log \sum_{i=1}^K p(u_i|h_j)p(a_j|u_i, h_j)$.

Notably, directly maximizing $\hat{I}(A; U|H)$ with gradient-based methods is not a preferable choice for two reasons. Firstly, The posterior probability $p(u_i|h_j)$ is hard to calculate. Secondly, the gradient of $p(a_j|u_i, h_j)$ is almost always equal to zero for many RL policies. For example, value-based policies commonly output actions that maximize the action-value function $Q(a, u, h)$ (or use ϵ -greedy for exploration). In this way, the value of $p(a_j|u_i, h_j)$ is determined only by whether $a_j = \arg \max_a Q(a, u_i, h_j)$. Consequently, the derivatives of $p(a_j|u_i, h_j)$ with respect to the neural network parameters are equal to zero.

We propose to optimize an alternative objective $\bar{I}(A; U|H)$, which have the following two properties:

1. $\bar{I}(A; U|H)$ provides gradients that can be used for neural network training, whether the meta-agent makes decisions by outputting action distributions or maximizing Q-functions.
2. Increasing $\bar{I}(A; U|H)$ also increases $\hat{I}(A; U|H)$.

The definition of $\bar{I}(A; U|H)$ is given below:

$$\bar{I}(A; U|H) = -\frac{1}{N} \sum_{j=1}^N \sum_{i=1, i \neq u_j}^K F(u_i, h_j, a_j) \quad (3)$$

where $F(u_i, h_j, a_j)$ represents the favor of the meta-agent for a_j given u_i, h_j and is required to be the direct output of a neural network so that its gradients can be used for gradient-based training. Below are several possible forms of $F(u_i, h_j, a_j)$:

1. If the neural network used for decision directly outputs action distribution (common for policy-gradient-based methods such as PPO (Schulman et al. 2017)), then $F(u_i, h_j, a_j)$ represents the probability of choosing a_j given u_i, h_j , which is $p(a_j|u_i, h_j)$;
2. If the neural network used for decision outputs advantage functions (or Q-values, common for value-based methods such as Dueling-DQN (Wang et al. 2016)), then $F(u_i, h_j, a_j) = A(u_i, h_j, a_j)$ (or $Q(u_i, h_j, a_j)$).

Moreover, we provide certain theoretical guarantee for maximizing $\bar{I}(A; U|H)$.

Theorem 1. Given $F(u_j, h_j, a_j)$, if F is update to F' such that:

$$\begin{aligned} \exists v \text{ s.t. } \quad & \arg \max_a F(u_v, h_j, a) = a_j \wedge \\ & F'(u_v, h_j, a_j) < F(u_v, h_j, a_j) \wedge \\ & \forall i \neq v, F'(u_i, h_j, a_j) = F(u_i, h_j, a_j) \end{aligned} \quad (4)$$

then the corresponding term I_j in $\hat{I}(A; U|H)$ is updated to I'_j and satisfies $I'_j \geq I_j$.

The proof is presented in the Appendix (Hui et al. 2025).

Computing $\bar{I}(A; U|H)$ requires obtaining all agents' actions under the same observation, which is efficiently achieved via the meta-agent with a single forward pass—unlike standard population training that needs N slow LSTM passes. Notably, ScaPT's two components are complementary: the meta-agent enables efficient mutual information computation and training, while the conditional mutual information term guides the learning of a diverse population.

Instantiation

ScaPT only specifies how to build an efficient and diverse population, and is compatible with multiple base RL frameworks that optimize agent policies. Our instantiation is based on an value-based approach because it is confirmed that this kind of method is suitable for our experimental task Hanabi (Hu and Foerster 2019; Bard et al. 2020). The main agent only needs to cooperate well with all the agents in the population (which are realized using the partner meta-agent), and thus it is required to minimize the base TD-error (Van Hasselt, Guez, and Silver 2016):

$$\begin{aligned} L_m = \frac{1}{N} \sum_{j=1}^N r_j + \gamma \max_a Q_{\theta'_m}(h'_j, u_j, a) \\ - Q_{\theta_m}(h_j, u_j, a_j) \end{aligned} \quad (5)$$

where θ_m and θ'_m represent the parameters in the online Q-net and target Q-net of the main agent respectively. Notably, the main agent has the same neural network architecture as a normal agent shown in Fig. 2. In comparison, the meta-agent not only needs to learn coordination, but also needs to become diverse by maximizing $\bar{I}(A; U|H)$. Consequently, $\bar{I}(A; U|H)$ is added to the base TD-loss with a weight α , which controls the balance between cooperation ability and population diversity:

$$\begin{aligned} L_p = \frac{1}{N} \sum_{j=1}^N \left[r_j + \gamma \max_a Q_{\theta'_p}(u_j, h'_j, a) \right. \\ \left. - Q_{\theta_p}(u_j, h_j, a_j) + \alpha \sum_{\substack{i=1 \\ i \neq j}}^K Q_{\theta_p}(u_i, h_j, a_j) \right] \end{aligned} \quad (6)$$

where θ_p represent the parameters in the online Q-net of the partner meta-agent. To accelerate convergence, prioritized replay (Schaul et al. 2015) and dueling-net (Wang et al. 2016) is also utilized.

Another key component of instantiation is the training mode, which specifies the pairs of agents interacting with the environment and the corresponding transitions for training. Table 1 summarizes six feasible modes, with details in the Appendix. Take Mode-III as an example: it has act groups MP, PP , meaning agents generate transitions in two groups: [main agent, partner agent] and [partner agent, partner agent]. The training objectives require the main agent to cooperate well with the partner, while the partner only optimizes self-play scores without adapting. Each mode emphasizes different aspects, and we investigate their performance experimentally.

Experiments

Matrix Game: Task Complexity vs. Population Size

To investigate how population size influences the ZSC performance of the main agent, we conducted comparative evaluations using collaborative single-step matrix games (Lupu et al. 2021). In each game, player 1 selects a row and player 2 selects a column independently. After both players have made their choices, the selected actions are revealed, and the agents receive the reward corresponding to the intersection of the chosen row and column. Specifically, we compared our approach against representative baselines, including **ET3** (Yan et al. 2024), **MEP** (Zhao et al. 2023), and **TrajeDi** (Lupu et al. 2021), **Individual agents**, where agents are trained independently, and **baseline** (Lupu et al. 2021). These methods cover both population-based and non-population-based paradigms, providing both a comparison across different algorithmic architectures and an in-depth view of how population size impacts the ZSC performance of population-based methods.

Fig. 3a presents the experimental results for various baseline algorithms under a matrix dimension of 50 and a fixed population size of 2. It can be observed that while all methods achieved strong performance in the self-play setting, their scores dropped significantly in the cross-play evaluation, indicating a lack of generalization between agents. Notably, ScaPT retained a non-negligible degree of ZSC capability, outperforming other baselines in this regard.

As illustrated in Fig. 3b, when the population size was increased, ScaPT's performance in cross-play improved markedly, demonstrating a strong resurgence in ZSC ability. This suggests that scaling the population size can enhance the diversity and representational capacity of the training process, thereby equipping the main agent with more generalized coordination strategies. The appendix includes examples of matrices with varying dimensions as well as the remaining experimental results.

ScaPT in Hanabi

We conducted experiments on Hanabi (Bard et al. 2020), a cooperative card game in which players must play cards in order by color and rank. Players can only see their teammates' cards, not their own, and must infer intentions and communicate through limited actions. As a result, self-play agents perform well due to shared training, but collaboration with unfamiliar agents is challenging, making Hanabi a

Index	Act Group	Objective for π_m	Objective for π_{pi}
I	<i>MP</i>	$\sum_{i=1}^N J(\pi_m, \pi_{pi})$	$J(\pi_m, \pi_{pi})$
II	<i>MM, MP</i>	$J(\pi_m, \pi_m) + \sum_{i=1}^N J(\pi_m, \pi_{pi})$	$J(\pi_m, \pi_{pi})$
III	<i>MP, PP</i>	$\sum_{i=1}^N J(\pi_m, \pi_{pi})$	$J(\pi_{pi}, \pi_{pi})$
IV	<i>MM, MP, PP</i>	$J(\pi_m, \pi_m) + \sum_{i=1}^N J(\pi_m, \pi_{pi})$	$J(\pi_{pi}, \pi_{pi})$
V	<i>MP, PP</i>	$\sum_{i=1}^N J(\pi_m, \pi_{pi})$	$J(\pi_{pi}, \pi_{pi}) + J(\pi_m, \pi_{pi})$
VI	<i>MM, MP, PP</i>	$J(\pi_m, \pi_m) + \sum_{i=1}^N J(\pi_m, \pi_{pi})$	$J(\pi_{pi}, \pi_{pi}) + J(\pi_m, \pi_{pi})$

Table 1: Feasible population training modes

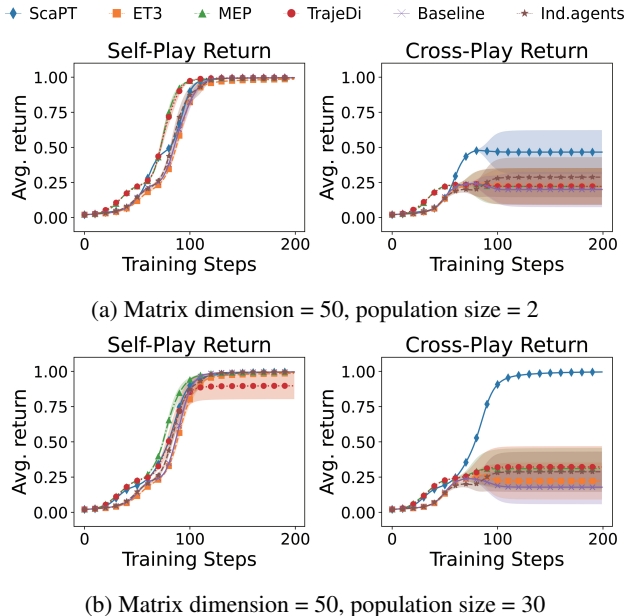


Figure 3: Results for matrix dimension = 50 with different population sizes

standard benchmark for zero-shot coordination.

Compared Methods and Evaluation Criteria

We compare our method with five representative methods, including one self-play methods: **SP** — self-play training with parameter sharing, serving as a baseline; two non-population-based algorithms designed for ZSC: **OP** (Hu et al. 2020), **OBL** (Hu et al. 2021); and two population-based ZSC algorithms, **TrajeDi** (Lupu et al. 2021) and **MEP** (Zhao et al. 2023).

Zero-shot coordination requires agents to cooperate well with collaborators that have not been seen before, that is, “strangers”. The corresponding evaluation criteria are slightly different from those of normal MARL. Below we give a formulaic representation in two-agent scenarios.

Use $J(\pi_1, \pi_2)$ to denote the expected cumulative discounted return obtained by the collaboration of π_1 and π_2 . Use π_i^M to denote the policy obtained with training framework M and random seed i . The earliest metric to evaluate zero-shot cooperation performance is intra-algorithm cross-

play (abbreviated as Intra-XP) score (Hu et al. 2020):

$$S_{intra-XP}(M) = \mathbb{E}[J(\pi_i^M, \pi_j^M) | i \neq j] \quad (7)$$

Cross-play is a commonly used metric that evaluates agent’s ZSC performance with partner agents trained using the same algorithm but different random seeds. However, since all agents share the same architecture, the test agent may retain implicit prior knowledge about its partners. Thus, relying solely on cross-play does not provide a rigorous assessment of zero-shot coordination.

In order to address the deficiencies of the aforementioned metric, Lucas and Allen (2022) propose one-sided zero-shot coordination (abbreviated as 1ZSC-XP) score:

$$S_{1ZSC-XP}(M) = \mathbb{E}[J(\pi^M, \pi^{M_t})] \quad (8)$$

where M_t refers to a set of algorithms that are not specially designed for zero-shot coordination. The shortcoming of this criterion is that π^{M_t} still cannot represent all feasible “strangers”, and the results may be biased.

Neither of the metrics is perfect, hence the following sections display the above two metrics for more comprehensive evaluation.

To enhance the credibility of the evaluation, we train models with several fixed random seeds (time consumption: 0-4 (2-player), 1-4 (5-player)) under each experiments, and test Intra-XP and 1ZSC-XP scores. Specifically, 1ZSC-XP scores are obtained by pairing the tested zero-shot coordination agents with 40 non-ZSC agents obtained with four kinds of self-play frameworks: IQL (Tan 1993), VDN (Sunehag et al. 2018), SAD and SAD+AUX (Hu and Foerster 2019).

Experiments with Different Frameworks under Resource Constraints

Firstly, in Hanabi, we evaluate each algorithm within its native framework under a strict resource constraints (128GB RAM and 48GB GPU memory), reflecting typical small-scale research servers.

We compare our method only with population methods designed for ZSC: TrajeDi and MEP on 5-player game. Table 2 shows the performance upper bounds of different methods under constrained RAM and GPU memory. Specifically, under these limits, TrajeDi and MEP can only support a population size of 1 due to the need to allocate separate replay buffers for each agent. In contrast, ScaPT achieves the best performance for both metrics by flexibly scaling the population size to 5 and 8.

Metrics	TrajeDi	MEP	ScaPT
Intra-XP	4.17±1.34	5.13±2.11	11.07±1.65
1ZSC-XP	4.22±1.47	6.69±1.62	11.39±1.60

Table 2: Best performance under Resource Constraints in 5-player hanabi game




Figure 4: Comparison of common population and meta-agent population: (a) training time consuming with population size increase; (b) resource consumption with population size increase; (c) sum of 1ZSC-XP scores over 40 experiments with different architectures using the MEP method.

To validate our architecture and ensure fair comparison, we evaluate MEP under common and scalable population architectures in computational cost, training time, and ZSC performance on 2-player Hanabi. As shown in Fig. 4, the scalable population scales more efficiently with slower growth in resource use and training time, while the common framework grows approximately linearly. Both approaches achieve comparable ZSC performance, demonstrating the scalability and effectiveness of our method for complex tasks requiring large population agents.

Experiments with Scalable Population Framework

We conduct a 2-player Hanabi experiment with population size 5 for all population-based algorithms. Table 3 reports mean and standard error of evaluation metrics. ScaPT achieves the highest 1ZSC-XP, demonstrating strong coordination with 40 unseen partners and ranks top on Intra-XP except for OBL. Although OBL scores high on Intra-XP, its 1ZSC-XP score is only 3.8, suggesting that OBL lacks genuine zero-shot coordination capability since it performs well mainly when cooperating with agents that are trained using the same algorithm as itself. Considering both metrics together, all population-based methods demonstrate solid performance. Among them, ScaPT delivers the most outstanding results, further confirming its superiority.

Besides, Fig. 5 visualizes the detailed pair-wise 1ZSC-XP cooperation scores of Table 3 using a heat-map. The heat-maps of three compared population-based methods are presented along for reference, and the heat-maps of other frameworks are presented in the Appendix. In each sub-figure, each row represents the coordination scores of a testing main agent pairing with all non-ZSC agents, and deeper




Figure 5: Detailed pair-wise 1ZSC-XP scores of TrajeDi, MEP and CMIMP. Deeper colors represent higher scores and each row represents the coordination scores of testing a main agent pairing with 40 non-ZSC agent, thus forming a 5×40 heat-map.

colors represent higher scores. Different rows in a sub-figure correspond to agents trained with different random number seeds under the same framework. There are two phenomena worth noting: Firstly, the differences between columns are consistent for these two frameworks. The reason is that some non-ZSC agents are relatively easy to cooperate with (e.g. column 7,8), while some others are not (e.g. column 2,3). Secondly, the performance stability of ScaPT is good: 1ZSC-XP scores of ScaPT models with different random seeds vary little, while TrajeDi and MEP demonstrate certain instability.

To validate the relationship between population size and ZSC performance, we conducted experiments on the Hanabi benchmark with all the population-based algorithms. By increasing the number of players from 2 to 5, we raised task difficulty to evaluate algorithm performance, while also varying population size to assess its effect on ZSC outcomes.

As shown in Table 4, increasing the number of players from 2 to 5, which raises task complexity, causes all algorithms to experience performance changes. TrajeDi suffers the largest drop, while MEP and ScaPT maintain relatively better performance. Overall, MEP performs strongest, with ScaPT showing comparable but slightly lower scores and more stable results across random seeds. Increasing the population size from 2 to 5 improves performance for all methods considerably—specifically, the 1ZSC-XP metric improves by 3.83 for TrajeDi, 1.38 for MEP, and 3.19 for ScaPT. As the population size increases from 5 to 8, the overall performance of all methods continues to improve, though with reduced marginal gains. However, a further increase to 15 results in performance degradation, indicating the presence of algorithm-specific optimal population sizes—an aspect that merits deeper exploration. ScaPT achieves the best overall results on both metrics, demonstrating strong generalization and underscoring the critical importance of population size for ZSC.

Ablation Study

As stated in equation 6, the training objective includes a weighted mutual information term that encourages behavioral differences among sub-modules. To examine its impact, Fig. 7 plots four key metrics across training epochs in

	SP	OP	OBL	TrajeDi*	MEP*	ScaPT
Intra-XP	3.96±0.49	14.94±0.67	23.80±0.03	15.95±2.47	20.09±0.69	20.77±0.40
1ZSC-XP	7.68±0.39	13.48±0.19	3.80±0.07	14.13±3.20	15.64±3.02	15.95±3.05

Table 3: Zero-shot coordination performance of different algorithms on the 2-player Hanabi game (population-based methods: TrajeDi, MEP, ScaPT, with population size 5). Algorithms marked * are implemented under the proposed scalable framework.

ps	Metrics	TrajeDi*	MEP*	ScaPT
ps=2	Intra-XP	6.08±1.86	8.13±2.11	7.88±1.72
	1ZSC-XP	3.71±0.60	7.08±1.62	6.18±1.53
ps=5	Intra-XP	9.91±1.48	9.51±1.64	11.07±1.65
	1ZSC-XP	8.38±0.75	8.84±0.93	11.35±0.58
ps=8	Intra-XP	9.04±0.44	10.98±0.33	10.61±0.93
	1ZSC-XP	9.86±1.73	10.48±1.60	11.39±1.60
ps=15	Intra-XP	8.98±0.87	6.78±0.37	8.40±0.83
	1ZSC-XP	9.01±1.65	8.25±1.09	9.10±1.81

Table 4: ZSC performance of population-based algorithms on 5-player Hanabi, under the scalable framework with different population sizes (PS).

	I	II	III	IV	V	VI
Intra-XP	1.16	20.77	8.93	12.43	12.07	19.22
1ZSC-XP	5.63	15.95	10.93	12.59	12.41	15.09

Table 5: Zero-shot coordination performance under different training modes (mean values only).

2-player hanabi environment under different training modes.

MM Score: The self-play score of the main agent;

MP Score: The cooperation score of the main agent and the partner agent, one of the key objectives of population training;

PP Score: The self-play score of the partner agents;

Diff Prob: The probability that different agents select the same action given the same observation—lower values indicate higher population diversity.

As is shown in Fig. 6, when the mutual information term $\bar{I}(A; U|H)$ is ignored (i.e. $\alpha = 0$), **Diff Prob** quickly rises to 1, meaning that the partner agents in the population act similarly. As a result, the generalization performance of the main agent is reduced (low XP scores during testing) despite that the training process goes smoothly (high MM/MP/PP scores during training). When α is set to a proper value ($\alpha = 1$), **Diff Prob** maintains relatively low, indicating a diverse population. A larger α ($\alpha = 10$) does not further increase diversity and reduces partner self-play performance (PP score), suggesting that excessive weight can harm rationality and coordination. To sum up, $\bar{I}(A; U|H)$ can help build a diverse population and thereby improve the zero-shot coordination, while assigning it a too large weight might have negative effect.

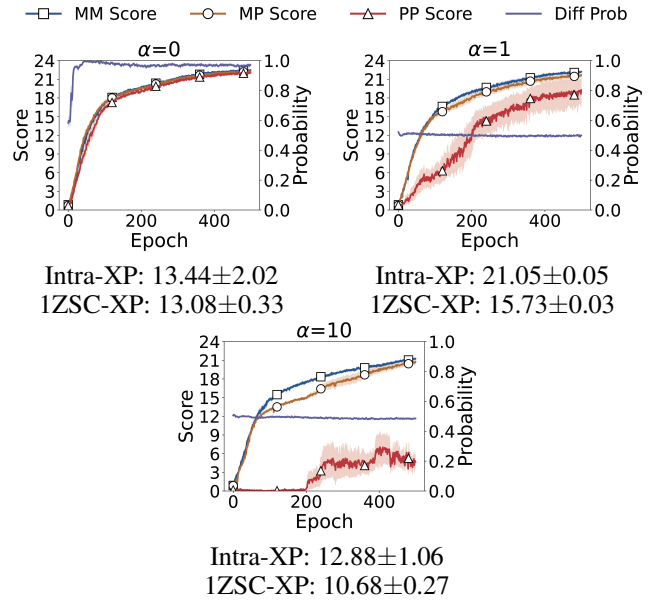


Figure 6: Training curves and testing scores of ScaPT with different α .

Comparison of Different Training Modes

Table 1 introduces six feasible training modes for population-based training. Then which mode is the best? Tab. 5 confirms that zero-shot coordination performance of different training modes varies a lot, and such differences are brought by the settings of secondary objectives. Mode-I is the worst, indicating that only optimizing the primary objective ($J(\pi_m, \pi_p)$) is not enough. Mode-II is the best, confirming the necessity of adding the self-play objective for the main agent. Notably, Mode-IV and Mode-VI additionally require increasing self-play scores for partner agents on the basis of Mode-II, and this operation is of no benefit judging from the results. The training curves for the different modes are provided in the appendix for reference.

Conclusion

In this paper, we analyzed limitations of population-based RL methods for ZSC, which struggle to scale efficiently enough to represent the diversity of evolving games, and proposed a population-size-free RL training framework. Experiments demonstrated its effectiveness and revealed a strong correlation between population size and ZSC performance, further validating the significance of our study.

Acknowledgments

This work is supported by National Key Research and Development Program of China (under Grant No.2023YFB2903904), and BNRist projects (No.BNR20231880004 and No.BNR2024TD03003).

References

- Bard, N.; Foerster, J. N.; Chandar, S.; Burch, N.; Lanctot, M.; Song, H. F.; Parisotto, E.; Dumoulin, V.; Moitra, S.; Hughes, E.; et al. 2020. The hanabi challenge: A new frontier for ai research. *Artificial Intelligence*, 280: 103216.
- Charakorn, R.; Manoonpong, P.; and Dilokthanakul, N. 2023. Generating diverse cooperative agents by learning incompatible policies. In *The Eleventh International Conference on Learning Representations*.
- Cui, B.; Hu, H.; Pineda, L.; and Foerster, J. 2021. K-level Reasoning for Zero-Shot Coordination in Hanabi. *Advances in Neural Information Processing Systems*, 34: 8215–8228.
- Do, H.; Jang, J.; and Kim, J. 2025. Heterogeneous multi-robot system mission planning with cooperative replenishment through data-driven rendezvous point selection. *Intelligent Service Robotics*, 18(1): 61–73.
- Eysenbach, B.; Gupta, A.; Ibarz, J.; and Levine, S. 2018. Diversity is all you need: Learning skills without a reward function. *arXiv preprint arXiv:1802.06070*.
- Gao, Y.; Liu, F.; Wang, L.; Lian, Z.; Wang, W.; Li, S.; Wang, X.; Zeng, X.; Wang, R.; Wang, J.; et al. 2023. Towards effective and interpretable human-agent collaboration in moba games: A communication perspective. *arXiv preprint arXiv:2304.11632*.
- Gupta, A.; Mendonca, R.; Liu, Y.; Abbeel, P.; and Levine, S. 2018. Meta-reinforcement learning of structured exploration strategies. *Advances in neural information processing systems*, 31.
- Hu, H.; and Foerster, J. N. 2019. Simplified Action Decoder for Deep Multi-Agent Reinforcement Learning. In *International Conference on Learning Representations*.
- Hu, H.; Lerer, A.; Cui, B.; Pineda, L.; Brown, N.; and Foerster, J. 2021. Off-belief learning. In *International Conference on Machine Learning*, 4369–4379. PMLR.
- Hu, H.; Lerer, A.; Peysakhovich, A.; and Foerster, J. 2020. “Other-Play” for Zero-Shot Coordination. In *International Conference on Machine Learning*, 4399–4410. PMLR.
- Hui, B.; Yu, L.; Yao, Q.; Qu, Y.; Zhang, X.; and Wang, J. 2025. Efficient Reinforcement Learning for Zero-Shot Coordination in Evolving Games. *arXiv:2511.11083*.
- 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*.
- Lowe, R.; Wu, Y. I.; Tamar, A.; Harb, J.; Pieter Abbeel, O.; and Mordatch, I. 2017. Multi-agent actor-critic for mixed cooperative-competitive environments. *Advances in neural information processing systems*, 30.
- Lucas, K.; and Allen, R. E. 2022. Any-Play: An Intrinsic Augmentation for Zero-Shot Coordination. In *Proceedings of the 21st International Conference on Autonomous Agents and Multiagent Systems*, 853–861.
- Lupu, A.; Cui, B.; Hu, H.; and Foerster, J. 2021. Trajectory diversity for zero-shot coordination. In *International Conference on Machine Learning*, 7204–7213. PMLR.
- Ma, J.; Zhao, Z.; Yi, X.; Chen, J.; Hong, L.; and Chi, E. H. 2018. Modeling task relationships in multi-task learning with multi-gate mixture-of-experts. In *Proceedings of the 24th ACM SIGKDD international conference on knowledge discovery & data mining*, 1930–1939.
- Muglich, D.; Schroeder de Witt, C.; van der Pol, E.; Whiteson, S.; and Foerster, J. 2022. Equivariant networks for zero-shot coordination. *Advances in Neural Information Processing Systems*, 35: 6410–6423.
- Nagabandi, A.; Clavera, I.; Liu, S.; Fearing, R. S.; Abbeel, P.; Levine, S.; and Finn, C. 2018. Learning to Adapt in Dynamic, Real-World Environments through Meta-Reinforcement Learning. In *International Conference on Learning Representations*.
- Rahman, A.; Fosong, E.; Carlucho, I.; and Albrecht, S. V. 2023. Generating Teammates for Training Robust Ad Hoc Teamwork Agents via Best-Response Diversity. *Transactions on Machine Learning Research*.
- Sagirova, A.; Kuratov, Y.; and Burtsev, M. 2025. SRMT: shared memory for multi-agent lifelong pathfinding. *arXiv preprint arXiv:2501.13200*.
- Schaul, T.; Quan, J.; Antonoglou, I.; and Silver, D. 2015. Prioritized experience replay. *arXiv preprint arXiv:1511.05952*.
- Schulman, J.; Wolski, F.; Dhariwal, P.; Radford, A.; and Klimov, O. 2017. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*.
- Sharma, A.; Gu, S.; Levine, S.; Kumar, V.; and Hausman, K. 2019. Dynamics-aware unsupervised discovery of skills. *arXiv preprint arXiv:1907.01657*.
- Sunehag, P.; Lever, G.; Gruslys, A.; Czarnecki, W. M.; Zambaldi, V.; Jaderberg, M.; Lanctot, M.; Sonnerat, N.; Leibo, J. Z.; Tuyls, K.; et al. 2018. Value-Decomposition Networks For Cooperative Multi-Agent Learning Based On Team Reward. In *Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems*, 2085–2087.
- Tan, M. 1993. Multi-agent reinforcement learning: Independent vs. cooperative agents. In *Proceedings of the tenth international conference on machine learning*, 330–337.
- Treutlein, J.; Dennis, M.; Oesterheld, C.; and Foerster, J. 2021. A new formalism, method and open issues for zero-shot coordination. In *International Conference on Machine Learning*, 10413–10423. PMLR.
- Van Hasselt, H.; Guez, A.; and Silver, D. 2016. Deep reinforcement learning with double q-learning. In *Proceedings of the AAAI conference on artificial intelligence*, volume 30.
- Wang, Z.; Schaul, T.; Hessel, M.; Hasselt, H.; Lanctot, M.; and Freitas, N. 2016. Dueling network architectures for deep

reinforcement learning. In *International conference on machine learning*, 1995–2003. PMLR.

Xue, K.; Wang, Y.; Guan, C.; Yuan, L.; Fu, H.; Fu, Q.; Qian, C.; and Yu, Y. 2024. Heterogeneous multi-agent zero-shot coordination by coevolution. *IEEE Transactions on Evolutionary Computation*.

Yan, X.; Guo, J.; Lou, X.; Wang, J.; Zhang, H.; and Du, Y. 2024. An Efficient End-to-End Training Approach for Zero-Shot Human-AI Coordination. *Advances in Neural Information Processing Systems*, 36.

Yu, C.; Velu, A.; Vinitsky, E.; Gao, J.; Wang, Y.; Bayen, A.; and Wu, Y. 2022. The surprising effectiveness of ppo in cooperative multi-agent games. *Advances in Neural Information Processing Systems*, 35: 24611–24624.

Zhao, R.; Song, J.; Yuan, Y.; Hu, H.; Gao, Y.; Wu, Y.; Sun, Z.; and Yang, W. 2023. Maximum entropy population-based training for zero-shot human-ai coordination. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, 6145–6153.

Zhu, H.; Zhong, H.; Yan, B.; and Li, X. 2024. Dynamic Formation Planning and Control for Robot Soccer Game with Multi-Agent Reinforcement Learning and Behavioral Model. In *2024 International Conference on Advanced Robotics and Mechatronics (ICARM)*, 837–842. IEEE.