

ESP: Exploiting Symmetry Prior for Multi-Agent Reinforcement Learning

Xin Yu^a, Rongye Shi^{a,*}, Pu Feng^a, Yongkai Tian^a, Jie Luo^a and Wenjun Wu^a

^aBeihang University, Beijing, China

Abstract. Multi-agent reinforcement learning (MARL) has achieved promising results in recent years. However, most existing reinforcement learning methods require a large amount of data for model training. In addition, data-efficient reinforcement learning requires the construction of strong inductive biases, which are ignored in the current MARL approaches. Inspired by the symmetry phenomenon in multi-agent systems, this paper proposes a framework for exploiting prior knowledge by integrating data augmentation and a well-designed consistency loss into the existing MARL methods. In addition, the proposed framework is model-agnostic and can be applied to most of the current MARL algorithms. Experimental tests on multiple challenging tasks demonstrate the effectiveness of the proposed framework. Moreover, the proposed framework is applied to a physical multi-robot testbed to show its superiority. Video demonstrations and Supplementary material are available at the project website <https://sites.google.com/view/esp-marl>.

1 Introduction

Artificial intelligence (AI) has been extensively applied across diverse scientific disciplines [29, 30]. A variety of AI challenges can be formulated as multi-agent reinforcement learning (MARL) problems [9]. With recent advances in MARL, many achievements in constructing intelligent agents have been realized, thus overcoming complex challenges, such as biology [40], multiplayer games [2] and multi-robot tasks [8]. However, the majority of the reinforcement learning (RL) methods strongly rely on massive amounts of data for model development. From a practical perspective, although a parallel and accelerated simulation environment enables agents to solve complex tasks within a reasonable amount of time, learning in real-world applications suffers from physical conditions-related constraints. Similarly, in simulations, due to limitations on the rendering speed, data efficiency is critical for realizing rapid experimental iterations [28, 23]. Therefore, improving the sample efficiency of the existing MARL methods is essential for both theoretical and practical research.

Improving sample efficiency entails reducing interactions with the environment in order to learn a specific policy more effectively. From the perspective of representation learning, achieving data-efficient reinforcement learning requires constructing stronger inductive biases, which poses a challenging task [14]. Recent literature has proposed various methods to address the data inefficiency issue in deep RL, broadly categorized as the data augmentation approach or network structure design approach. Specifically, data augmentation has



Figure 1: Illustration of multi-agent symmetry. Rotating the state globally in the cooperative navigation task results in a permutation of the optimal joint policy.

been widely used in recent studies in single-agent RL, most of the proposed solutions focus on image-based data augmentation [22]. A technique called RL with augmented data (RAD) incorporates multiple data augmentations into visual observations, such as image crop, enabling single-agent RL pipelines to outperform prior state-of-the-art methods [17]. Despite the importance of data augmentation, to the best of our knowledge, theoretically founded data augmentation methods in MARL have been scarcely investigated. Second, network structure design approaches usually aim to design specialized network architectures that implicitly embed prior knowledge of a given task. Network structure design methods use specific sets of basis functions chosen according to the type of transformation group [5, 6]. For example, circular harmonic functions and Hermite polynomials have been used to achieve rotation and scale equivariance, respectively [35, 31]. However, previous studies have shown that this type of constraint can overly restrict the model and result in sub-optimal performance on learning tasks [15]. In addition, the sets of basis functions vary for different transformation groups, making it challenging to design composite networks that are equivariant to multiple transformation groups such as rotations, reflections, and changes in scale. Compared to network structure design methods, our method enables the simultaneous utilization of multiple forms of symmetry, providing a more straightforward implementation.

* Rongye Shi is the corresponding author. Email: shirongye@buaa.edu.cn.

The most general symmetry in the multi-agent systems is the global symmetry shown in Figure 1, where rotating the global state results in a permutation of the optimal joint policy. Motivated by the symmetric phenomenon in multi-agent domains, we propose a general framework for MARL to increase sample efficiency. The proposed framework consists of two parts, a symmetry augmentation module, and a well-designed symmetry consistency loss. The symmetry augmentation module performs transformations on the trajectories to generate extra data before storing them in a replay buffer. Moreover, a symmetry consistency loss, which serves as an auxiliary module, is introduced to bring symmetry prior into an RL agent.

We conduct extensive experiments on several challenging MARL benchmarks (e.g., Predator-Prey, Cooperative Navigation, and Formation Change) to demonstrate the superior performance of the proposed framework regarding data efficiency and higher evaluation rewards. Experimental results demonstrate that our proposed method improves the performance of MARL on both simple and complex tasks. Compared to network structure design methods, our approach offers superior performance and is more straightforward to implement. In addition, we tested the performance of our methods using physical robots on the Formation Change task.

2 Related work

2.1 Data augmentation in RL

Introducing data augmentation into RL has primarily aimed at enhancing data efficiency [12]. A natural approach to exploit data augmentation in single-agent RL is to obtain more data via image transformation during model training [17, 37, 18, 1]. Another type of approach introduces data augmentation through an innovative contrastive learning framework called the CURL. The CURL learns representations that can improve data efficiency by enforcing consistencies between an image and its augmented version through instance contrastive loss [16]. Although improving data efficiency is a widely studied problem in the RL field, fewer studies focused on exploiting the data augmentation methods in multi-agent domains. To the best of the authors’ knowledge, the most relevant algorithm from the available literature to this work is the data augmentation method for MARL proposed in [38]. This method generates extra data by performing permutation transformations for homogeneous agents. The approach of performing data augmentation using permutation invariance is essentially another interpretation perspective of the parameter-sharing learning paradigm. However, this study goes beyond permutation symmetries to a broader group of symmetries in multi-agent RL.

2.2 Integrating inductive bias into network

Currently, the major approach for exploiting inductive bias in neural networks is to design specialized architectures that implicitly embed prior knowledge associated with a given task [3, 27, 41]. In the context of multi-agent reinforcement learning (MARL), graph neural networks (GNNs) are commonly used to model interactions between agents and allow for a graph-based representation [13, 20]. However, more prior knowledge needs to be integrated into MARL beyond the graph structure. The content relevant to this work focuses on embedding symmetry into network. For single-agent RL, symmetries in the joint state-action space can be expressed in MDP homomorphic networks [33]. SO(2)-Equivariant RL enforce symmetry within the structure of their convolutional layers [34]. As for MARL, Multi-Agent MDP Homomorphic Networks integrate the symmetries into a

neural network, thus improving data efficiency [32]. These methods essentially rely on the construction of filters from linear combinations of certain sets of basis functions. Previous studies have shown that constraining the weight parameterization space to specific basis sets overly constrains the model, and although this strategy makes the model equivariant to the desired transformation group, the performance on the actual learning task may be sub-optimal [15]. Furthermore, the method of network structure embedding requires re-designing standard networks by incorporating sets of complex basis functions that tend to vary for different transformation groups. This makes it difficult to design composite networks that are equivariant to multiple transformation groups, such as rotations, reflections, and changes in scale [36, 31]. In contrast, our method enables the simultaneous utilization of multiple forms of symmetry and avoids the limitations of network structure embedding in terms of complexity and design.

3 Background

3.1 Group and transformations

We provide a brief introduction to the concepts of groups and transformations in this section [4]. A group G is a set with a binary operator that have specific mathematical properties: identity, inverse, closure, and associativity. We will refer extensively to the group SO(2) and its cyclic subgroup C_n . SO(2) is the group of continuous rotations $\{R_\theta : 0 \leq \theta < 2\pi\}$. C_n is the discrete subgroup $C_n = \{R_\theta : \theta \in \{\frac{2\pi i}{n} \mid 0 \leq i < n\}\}$ of rotations.

A rotation matrix is a transformation matrix that describes performing a rotation in Euclidean space [10]. For a set of rotations $\{0^\circ, 90^\circ, 180^\circ, 270^\circ\}$, the rotation matrix is defined as follows:

$$R(\theta) = \begin{bmatrix} \cos \theta & -\sin \theta \\ \sin \theta & \cos \theta \end{bmatrix}. \quad (1)$$

The four group axioms are satisfied in the case of a rotation transformation. This paper primarily utilizes the C_4 group, while our method can be applied to other groups as well.

3.2 Equivariance and invariance

The symmetry appearing in multi-agent systems can be denoted as equivariance and invariance. Given a transformation operator $L_g : \mathcal{X} \rightarrow \mathcal{X}$ and a mapping function $f : \mathcal{X} \rightarrow \mathcal{Y}$, if there exists a second transformation operator $K_g : \mathcal{Y} \rightarrow \mathcal{Y}$ in the output space of f such that:

$$K_g[f(x)] = f(L_g[x]), \quad (2)$$

where $g \in G$ and G is a mathematical group, then, function f is equivariant to the transformation L_g . A related notion to equivariance is invariance. If for any choice of $g \in G$ we have that $K_g = I$, the identity function, then we say function f is invariant to transformation L_g .

Figure 1 shows the equivariance of the optimal policy, rotating the state globally results in a transformation of the optimal joint policy. Given two states s and $L_g[s]$, the optimal policy π^* is equivariant to the transformation L_g which is denoted by $K_g[\pi^*(s)] = \pi^*(L_g[s])$.

3.3 Multi-agent proximal policy optimization

Multi-agent Proximal Policy Optimization (MAPPO) is a variant of PPO which is specialized for multi-agent settings. In the default formulation of MAPPO under the Decentralized Partially Observable

Markov Decision Process (Dec-POMDP) framework, the agents employ the trick of policy sharing. As such, the MAPPO learns policy π_θ by optimizing the following objective:

$$J_\pi(\theta) = \sum_{i=1}^n E_{\pi_{\theta_{\text{old}}}} \left[\frac{\pi_\theta(a^i | s)}{\pi_{\theta_{\text{old}}}(a^i | s)} A^\pi(s, a^i) \right], \quad (3)$$

where $\pi_{\theta_{\text{old}}}$ is the behavior policy used to collect trajectories, π_θ is the policy we want to optimize, and $A^\pi(s, a^i)$ denotes the advantage function. And the critic network is trained to optimize the following objective:

$$J_{V^\pi}(\psi) = - \sum_{i=1}^n E_\pi [V_{\text{target}}^\pi(s) - V_\psi^\pi(s)]^2. \quad (4)$$

In the above equations, $V_{\text{target}}^\pi(s)$ denotes the target value computed by Generalized Advantage Estimation [39]. For ease of clarity, we simplified the description of the objective. We use s to denote the input of both policy and value functions, while the actual input may not be the same.

4 Problem statement

4.1 Cooperative markov game

We formulate the cooperative MARL problem as cooperative Markov game [19]. An n -agent cooperative Markov game can be represented using a tuple $(N, S, \{A_i\}_{i=1}^n, R, T, \Psi)$, where N denotes the set of agents, S is the state space, and A_i is the action space of an agent $i = 1, \dots, n$. Let $A = A_1 \times A_2 \times \dots \times A_n$ be the joint action space. $T : S \times A \times S \rightarrow [0, 1]$ is the transition function. Ψ is the set of admissible state-action pairs. At time step t , the agents are at state s_t (which may not be fully observable); the agents take independent action (a_1, \dots, a_n) rely on their policy. Then, the environment emits the bounded joint reward R and moves to the next state s_{t+1} . The agents aim to maximize the expected joint return, defined as $E_\pi [\sum_{t=0}^{\infty} \gamma^t R(s_t, a_t)]$, where γ is the discount factor, by selecting actions according to the policy $\pi_i : S \times A_i \rightarrow [0, 1]$. The initial states are determined by a distribution $\eta : S \rightarrow [0, 1]$.

4.2 Symmetry markov game

We propose Symmetry Markov game which is a subclass of the cooperative Markov game endowed with symmetry characterization. The Symmetry Markov game $(N, S, \{A_i\}_{i=1}^n, R, T, \Psi, G)$ is a cooperative Markov Game that satisfies the conditions of reward invariance and transition invariance. The state transformation and action transformation in the Symmetry Markov game are defined as $L_g : S \rightarrow S$ and $K_g : A \rightarrow A$, respectively. For state-action pairs $(s, a) \in \Psi$, we denote the transformation of group g on s and a as $(L_g[s], K_g[a]) \in \Psi$. Further, for all $g \in G, s \in S, a \in A$, the conditions of reward invariance and the transition invariance are given by:

$$R(s, a) = R(L_g[s], K_g[a]), \quad (5)$$

$$T(s, a, s') = T(L_g[s], K_g[a], L_g[s']). \quad (6)$$

We assert that the symmetry property is a common occurrence in many real-world multi-agent tasks, albeit often requiring effort to identify. One approach to identifying symmetry, for instance, is to observe whether an optimal action remains optimal for a transformed state. If so, the symmetry property may be present. It should be emphasized that our method is not restricted to the Symmetry Markov

game and can also be applied to Dec-POMDPs with symmetry. Additional details on this topic are available in the Sec.4 of the Supplementary Material.

5 Methods

In this section, we introduce our proposed framework, the Symmetry Prior Exploitation (ESP), for improving the performance of existing MARL algorithms. It is important to note that although fully utilizing symmetry is a desirable property in certain deep learning applications, our primary objective is to maximize the performance of agents in reinforcement learning tasks. To achieve this goal, we avoid using a hard constraint approach that embeds symmetry into the network structure. Instead, we treat symmetry as an additional objective and incorporate it through soft constraints such as data augmentation and regularization. This allows us to improve the performance of the reinforcement learning algorithm while benefiting from the advantages of symmetry. ESP comprises of two main components: a symmetry augmentation module and a symmetry consistency loss module.

5.1 Symmetry augmentation

The most straightforward way to utilize symmetry is through data augmentation. This approach has been widely used in deep learning and has been shown to be effective in various tasks. In the context of multi-agent reinforcement learning, data augmentation can be used to generate additional samples that reflect the symmetry property of the environment, which can improve the training efficiency of the agents. Although data augmentation has been widely used in supervised learning, it still lacks a theoretical framework to prove the correctness of exploiting symmetry to augment state-action pairs in MARL. This study proposes a multi-agent optimal value function equivalence proposition for the symmetric Markov game.

In a Symmetry Markov game, the transition and reward functions are invariant to group elements $g \in G$ acting on the state and action space. In the following, variables without a subscript i denote the concatenation of all variables for all agents. (e.g. a denotes the joint actions of all agents).

Proposition 1 (Optimal value equivalence). *Consider a Symmetry Markov game $(N, S, \{A_i\}_{i=1}^n, \{R_i\}_{i=1}^n, T, \Psi, G)$ with a invariant transformation $h = (L_g, K_g)$ where $g \in G$. For any $(s, a) \in \Psi$, it holds that $Q^*(s, a) = Q^*(L_g[s], K_g[a])$.*

A detailed proof of proposition 1 is provided in Sec.1 of the Supplementary Material. Proposition 1 indicates that using transformation data to train the policy in a Symmetry Markov game is reasonable. The process of utilizing data augmentation in MARL is demonstrated in Figure 2, where the policy interacts with the environment to generate samples that are subsequently augmented. These trajectories can then be stored in the replay buffer and used for training purposes. This will be explained in the example of a multi-agent navigation task, which is shown in Figure 1. In this task, agents need to cover three target locations in a two-dimensional space simultaneously while avoiding collisions with each other. Agents observe the relative positions of other agents and landmarks. When performing a rotation operation on a state, each element can be multiplied by the rotation matrix given in Equation (1). After multiplying the action vector by the permutation matrix, a set of trajectories consisting of true trajectories and their symmetric trajectories is obtained. It is worth mentioning that the proposed symmetry augmentation method is not limited by the type of algorithm and can be directly applied to most MARL algorithms.

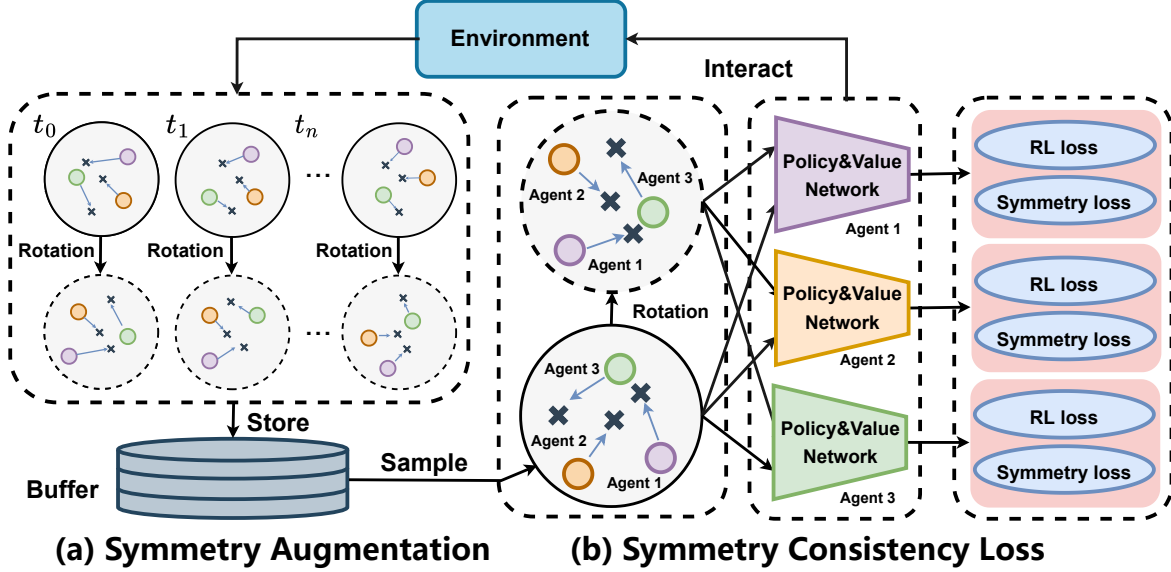


Figure 2: The overall framework of the proposed ESP. The ESP consists of the symmetry augmentation module and symmetry consistency loss module. symmetry augmentation module: the policy interacts with the environment to generate samples, which are then augmented and deposited into a buffer. symmetry consistency loss: transitions are sampled from the buffer, and the data are transformed to obtain symmetric transition pairs, which are subsequently used to compute symmetry consistency loss and further update the agent network.

5.2 Symmetry consistency loss

While data augmentation has been previously utilized in RL, we argue that there are issues with its application in practice. For clarity of presentation, the MAPPO is used as an example to introduce the symmetry consistency loss. If transformation (L_g, K_g) is applied to the algorithm, the MAPPO objective changes, and equation (3) is replaced by

$$J_\pi(\theta) = \sum_{i=1}^n E_{\pi_{\theta_{\text{old}}}} \left[\frac{\pi_\theta(K_g[a^i] | L_g[s])}{\pi_{\theta_{\text{old}}}(a^i | s)} A_\psi^\pi(s, a^i) \right]. \quad (7)$$

However, the right hand is not a sound estimate of the left hand because $\pi_\theta(K_g[a] | L_g[s]) \neq \pi_\theta(a | s)$, particularly in the early stage of training. In fact a certain transformation (L_g, K_g) for data augmentation can result in an arbitrarily large ratio $\pi_\theta(K_g[a] | L_g[s]) / \pi_{\theta_{\text{old}}}(a | s)$. Please see the Sec.2 of the Supplementary Material for more details. In multi-agent settings, using data augmentation to improve sample efficiency can be challenging. The reason for this is that when multiple agents are considered, more sources of variance are introduced, making the training severely unstable as shown in equation (7). The proposed symmetry consistency loss provides a solution to this challenge and helps to make the use of data augmentation successful in multi-agent settings.

Symmetry consistency loss. In order to properly utilize data augmentation in multi-agent reinforcement learning, we have designed a symmetry consistency loss for the policy and value function. The policy consistency loss term $S_\pi(\theta)$ is defined as

$$S_\pi = KL[\pi_\theta(K_g[a] | L_g[s]) | \pi_\theta(a | s)], \quad (8)$$

aims to constrain distribution $\pi_\theta(K_g[a] | L_g[s])$ to be close to $\pi_\theta(a | s)$. This helps guide the training process according to the symmetry prior.

Assume that $V_\psi(s)$ represents an approximate value for state s ,

the symmetry consistency loss for value function is shown as

$$S_V = E_{s,a} [(V_\psi(s) - V_\psi(L_g[s]))^2], \quad (9)$$

designed to minimize the discrepancy between the outputs of the value function when provided with the original input and the symmetry-transformed input. Therefore, we regard Equation (8) and (9) as the symmetry consistency loss.

MARL with symmetry consistency loss. It is convenient to combine the symmetry consistency loss with the existing MARL algorithms. For the policy-based algorithm MAPPO, our ESP framework maximizes the following objective:

$$J_{ESP} = J_{MAPPO} - c(S_\pi + S_V), \quad (10)$$

where c is the hyperparameter that controls the symmetry consistency loss coefficient. J_{MAPPO} are defined in Equation (3) and (4). As shown in Figure 2, data sampled from the replay buffer is transformed to $(L_g[s], K_g[a])$, and the paired data are used to calculate the symmetry consistency loss. We can adjust hyperparameters c according to the strength of symmetry to attempt to embed symmetry into the learning process.

5.3 Algorithm description

The training procedure of our ESP framework is described in Figure 2. We start by initializing the group transformation $h = (L_g, K_g)$ alongside several training parameters, including transformation group g , agent numbers, maximum steps, and batch size. Then the transformation h is selected according to group g . After that, the real trajectories generated by agents are augmented by symmetric transformation. Both real and augmented samples are stored in an experience replay buffer. Finally, data are sampled from the replay buffer to update the agent network using the symmetry consistency loss. Our method is based on data augmentation and is not limited by fully observable or partial observable. When facing partially observable problems, we can rotate the state globally during the

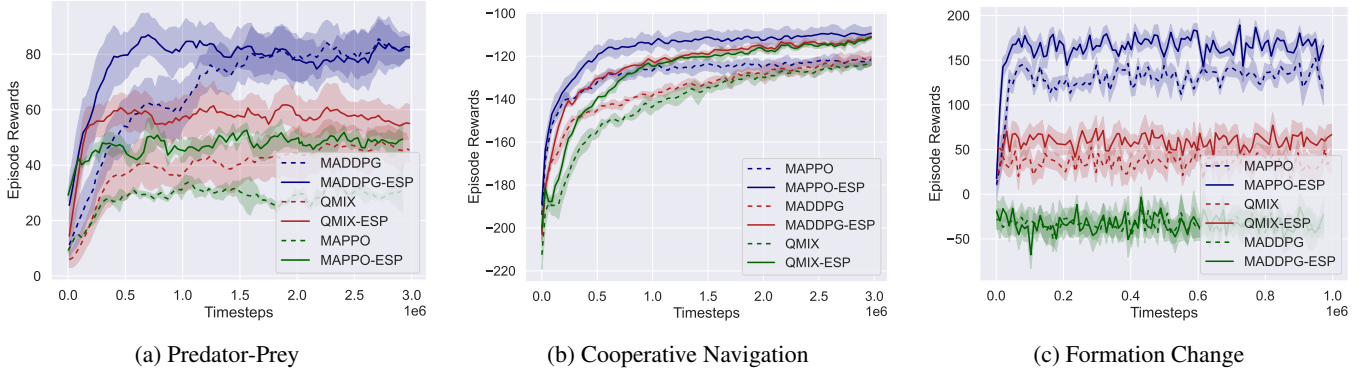


Figure 3: Learning curves of the original MADDPG, QMIX, and MAPPO and their versions with the ESP framework on the three multi-agent tasks. Each experiment was executed ten times with different random seeds. The solid lines denote the algorithm employing ESP framework.

centralized training phase, and then augment the agent’s trajectories which can be used to train the algorithm. In the execution phase, the agents can perform the policy in a distributed manner.

6 Experiments

6.1 Environmental setting

The effectiveness of our framework was verified by experiments on four different tasks, cooperation navigation (CN), predator-prey (PP), formation change (FC), and wildlife monitoring. About the detailed settings of the task, such as reward function, state space, and action space, please refer to the Sec.5 of the Supplementary Material.

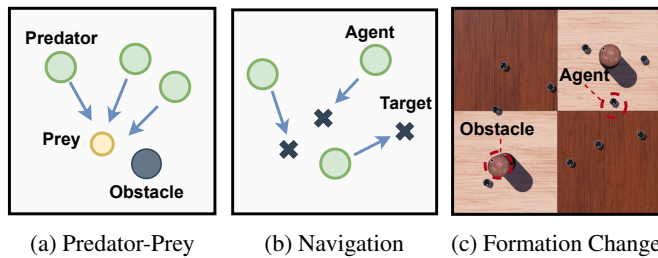


Figure 4: The schematics of the experimental tasks.

Predator-Prey. Predator-Prey is a classic scenario implemented in a multi-agent particle environment [25]. In this task, collaborating agents need to catch prey. Predators receive a reward at every time step when at least one predator is on top of the prey. The Predator-Prey task is presented in Figure 4a.

Cooperative Navigation. Cooperative navigation is a classic scenario implemented in a multi-agent particle environment. In this task, agents need to cover target landmarks whilst avoiding collisions with each other. Both agents and targets are spawned at random locations at the start of each episode. A classic cooperative navigation scenario with three agents is presented in Figure 4b.

Formation Change. A multi-robot formation change task was designed inspired by [8]. This task is shown in Figure 4c, in this task, all robots started on a square formation and had their goals set to the opposite side. The Epuck was selected as a model in the open-source mobile robot simulator Webots [24]. The Epuck was controlled by assigning the wheel speed. Robots were required to learn to avoid obstacles and each other and coordinate to cover the destinations. Robots took actions according to the current state in each step and

obtained the corresponding rewards, and then made the next state. When the maximum step number was reached, the current episode ended, and the next one started.

Wildlife monitoring. The wildlife monitoring is a grid-world based environment, where a set of drones has to coordinate to accomplish the task [33]. The goal is to trap poachers by having one drone hover above them while the other assists from the side. Two drones cannot be in the same location at the same time. This task was used to compare our ESP with the network design method.

Baselines. The proposed framework was applied to several baselines, including Multi-Agent Deep Deterministic Policy Gradient (MADDPG), Monotonic Value Function Factorisation for Deep Multi-Agent Reinforcement Learning (QMIX), and Multi-Agent Proximal Policy Optimization (MAPPO), which are mainstream MARL approaches [26, 21, 39].

6.2 Main results

This section presents the experimental results obtained using the setup described in Section 6.1. The performance of each algorithm was evaluated with ten different random seeds, and the final experimental results are shown in Figure 3 and Table 1. The results show that the algorithms adopting the ESP framework achieved different degrees of advantage over their original versions.


Predator-Prey. In this scenario, there were three predators and one prey. As shown in Figure 3a, the proposed ESP framework outperformed the baseline methods significantly. The results indicated that the proposed framework could increase both the data efficiency and the convergence reward.

Cooperative Navigation. The cooperative navigation was a fully cooperative environment, where 3 agents (circles) cooperated to reach 3 landmarks (crosses) under a minimum number of collisions. Similarly, as shown in Figure 3b, the results show that the proposed framework can improve data efficiency and convergence rewards in this task.

Formation Change. To evaluate the proposed method in complex tasks, experiments were conducted on the multi-robot formation change task in the Webots simulator, as shown in Figure 4c. In this scenario, eight robots started in a square formation and had their goals set on the opposite side. The experimental results showed that the MADDPG and QMIX could not learn a useful policy in this task, whereas the agents trained by the MAPPO-ESP and MAPPO could reach the goal while avoiding collisions with each other and obstacles. As presented in Figure 3c, the algorithms enhanced by

Table 1: Episode rewards of the ESP and baselines on the three tasks. '500k' and '3000k' represent the number of training steps of the algorithms. The error bars are the standard error of the mean.

Task	Steps	MADDPG	MADDPG-ESP	QMIX	QMIX-ESP	MAPPO	MAPPO-ESP
Predator-Prey	500k	56.81±11.26	80.71±9.12	38.93±8.96	59.95±4.86	28.35±6.76	45.22±5.45
	3000k	85.43±8.14	87.06±4.87	48.78±7.75	57.75±6.81	32.54±8.21	59.54±4.32
Cooperative Navigation	500k	-145.23±3.71	-132.22±1.27	-154.24±3.16	-140.12±2.25	-132.65±2.89	-122.25±1.62
	3000k	-122.63±6.04	-112.21±2.75	-124.67±3.05	-113.84±1.43	-126.97±2.12	-110.22±2.12
Formation Change	200k	-45.63±13.96	-41.13±11.17	48.53±20.12	51.23±14.71	122.56±15.12	163.18±14.14
	1000k	-42.13±10.71	-44.13±12.31	49.44±18.71	59.47±14.35	140.64±18.45	170.85±10.2



(a) Learning curve of MADDPG using augmentation solely with different numbers.

(b) Learning curve of MADDPG using different augmentation types.

(c) Parameter sensitivity of symmetry loss coefficient at different training steps.

Figure 5: The ablation study and parameter sensitivity of our ESP framework.

the proposed framework obtained higher rewards than their original versions. This indicates that the proposed method can effectively improve the performance of multi-agent reinforcement learning algorithms in challenging environments.

In Table 1, the episode rewards of different algorithms at different training steps in the three task scenarios are presented. The superiority of the proposed framework is reflected in two main aspects. First, the proposed algorithm outperforms baseline algorithms by giving higher converged rewards. Additionally, it can be observed that the proposed algorithm exhibits faster convergence during training, possibly due to the utilization of prior knowledge within the framework. In contrast, the baseline trains the network through pure trial and error, which has a high computational cost. Consequently, the proposed method has strong versatility and can perform well in different tasks.

6.3 Ablation study

To verify the effectiveness of each part in the proposed framework, ablation experiments were conducted in cooperative navigation task. Our experiments demonstrate that: 1) Using a single data augmentation significantly improves performance, whereas adding more of the same type has only a minor effect. 2) Using diverse transformation types is more effective than using multiple transformations of the same type. 3) The data augmentation and symmetry loss modules are both effective and improve performance when used in combination.

Impact of the number and type of transformations. In this section, we investigate the impact of using different numbers and types of symmetric transformations on performance without using symmetry loss. Firstly, we examine the impact of using different quantities of the same group of augmentations. As shown in Figure 5a, we used one to three rotation augmentations. We observed that the convergence speed of MADDPG gradually increased with an increasing number of rotations of the same type, and convergence reward

slightly improved. Secondly, we analyze the impact of using different types of symmetry transformations. We denote the flip around the x-axis as Flip. As illustrated in Figure 5b, the simultaneous use of rotation and flip has a more pronounced effect on performance than using only rotation transformations.

Effect of Symmetry Consistency Loss Coefficient. In this subsection, we analyze the impact of the symmetry consistency loss coefficient. We use the transformation of rotation by 90 degrees for the remaining experiments. The ESP was tested with different values of the symmetry consistency loss coefficient. In Figure 5c, the horizontal axis denotes the value of the symmetry consistency loss coefficient, and the vertical axis shows the testing episode reward averaged over ten seeds. The color of a line denotes a different step of the training steps. When the symmetry consistency loss coefficient equaled zero, the algorithm degraded to the baseline algorithm MADDPG. The results indicated that as the coefficient value increased, the episode rewards of the early training stage (i.e., 500k steps) first increased correspondingly and then decreased when the coefficient was larger than 0.5. However, the coefficient had little effect on the convergence episode rewards, at the training step of 3,000k step.

Effect of different modules. The effect of data augmentation and symmetry consistency loss on the MADDPG performance was investigated in this subsection. We choose $c=0.5$ to balance the traditional RL loss and the symmetry consistency loss. Four variations were analyzed: 1) MADDPG which is the vanilla algorithm; 2) MADDPG + loss, which used only the symmetry consistency loss in the MADDPG; 3) MADDPG + Data aug, which used only the symmetry augmentation in the training process; 4) MADDPG-ESP, which used the symmetry consistency loss and the symmetry augmentation in the MADDPG. As shown in Figure 6a, when used separately, both of the proposed modules improved the convergence speed and the final rewards. When used together, their combined performance was




Figure 6: Additional experimental results for validating the effectiveness of the ESP framework.

even better. The empirical results suggested the importance of the two modules in improving multi-agent task performance.

6.4 Comparison with network design method

In [32], the authors integrated symmetry into the MARL process by designing a specialized network structure, namely Multi-Agent MDP Homomorphic Networks, for wildlife monitoring tasks. However, Multi-Agent MDP Homomorphic Networks are limited to grid-world environments and rely on pixel-based states, making them challenging to extend to more complex scenarios. The three task scenarios shown in Figure 4 involve complex, continuous state spaces that cannot be handled by the network structure proposed in [32]. Therefore, we compared our proposed ESP method with their approach in the wildlife monitoring task.

Authors of [32] designed Multi-Agent MDP Homomorphic Networks (Equivariant MPN), which is built upon the standard MPN. Similarly, we employed the ESP framework on the PPO algorithm used in Equivariant MPN and denoted as MPN-ESP. As depicted in Figure 6b, the proposed ESP exhibits comparable performance to that of using a Multi-Agent MDP Homomorphic Network. We conclude that our approach demonstrates similar results as network structure design methods, but it is simpler to implement and more easily adaptable to different tasks.

7 Demonstration on robots

As shown in Figure 7, the trained policies were deployed on the Epuck, which is a small, lightweight, open-source robot platform. The formation change task presented in Section 6.1 was used in this experiment. We followed a direct sim2real paradigm to deploy the policy network [7]. Trajectories from a successful demonstration are displayed in Figure 6c. Different colors denote the trajectories of different robots, and color transparency indicates the temporal state. By incorporating our ESP approach into the MAPPO algorithm, the agents are able to complete tasks with fewer risky states. Risky states are defined as those in which the distance between agents is less than 5 centimeters, and the rate of risky states is the proportion of risky states to all states. The rate of risky states for ESP-MAPPO is 2.2%, while the rate for MAPPO is 5.8%. Additional details and videos are provided in Sec.6 of the Supplementary Material.




Figure 7: Real-world formation change on a swarm of robots. The robots successfully switched their positions to the antipodal points by achieving collision avoidance.

8 Conclusions

In this paper, we proposed a general framework named ESP that utilizes prior knowledge to increase sample efficiency in multi-agent reinforcement learning. Our framework includes two key components: symmetry augmentation and symmetry consistency loss. By performing symmetric transformations on state-action pairs, our method generates additional data, which is stored in the replay buffer for training. The symmetry consistency loss serves as an auxiliary module to incorporate symmetry representation into the RL agent. Extensive empirical experiments on various multi-agent tasks demonstrate that our ESP framework can effectively enhance the performance of existing MARL algorithms. Moreover, our approach outperforms network structure design methods while being more straightforward to implement.

However, our approach is currently limited to scenarios where the presence of symmetry is known. Future work can extend our approach to more complex real-world tasks with unknown prior knowledge by incorporating automatic symmetry discovery functionality. Additionally, considering problems that only have local symmetry or break global symmetry to a certain degree could be a promising direction. Moreover, the integration of symmetry principles in the enhancement of robustness within multi-agent systems emerges as a noteworthy direction of research[11]. Overall, our work sheds light on the application of reinforcement learning in complex physical environments and provides a useful framework for incorporating prior knowledge in multi-agent settings.

Acknowledgements

This work is supported by the National Key Research and Development Project of China (No. 2022ZD0117801).

References

- [1] Fabio Amadio, Adrià Colomé, and Carme Torras, 'Exploiting symmetries in reinforcement learning of bimanual robotic tasks', *IEEE Robotics and Automation Letters*, **4**(2), 1838–1845, (2019).
- [2] Christopher Berner, Greg Brockman, Brooke Chan, Vicki Cheung, Przemysław D. Biak, Christy Dennison, David Farhi, Quirin Fischer, Shariq Hashme, Chris Hesse, et al., 'Dota 2 with large scale deep reinforcement learning', *arXiv preprint arXiv:1912.06680*, (2019).
- [3] Michael M Bronstein, Joan Bruna, Taco Cohen, and Petar Veličković, 'Geometric deep learning: Grids, groups, graphs, geodesics, and gauges', *arXiv preprint arXiv:2104.13478*, (2021).
- [4] Michael M Bronstein, Joan Bruna, Taco Cohen, and Petar Veličković, 'Geometric deep learning: Grids, groups, graphs, geodesics, and gauges', *arXiv preprint arXiv:2104.13478*, (2021).
- [5] Taco Cohen and Max Welling, 'Group equivariant convolutional networks', in *International conference on machine learning*, pp. 2990–2999. PMLR, (2016).
- [6] Taco S Cohen and Max Welling, 'Steerable cnns', *arXiv preprint arXiv:1612.08498*, (2016).
- [7] Cristino De Souza, Rhys Newbury, Akansel Cosgun, Pedro Castillo, Boris Vidolov, and Dana Kulić, 'Decentralized multi-agent pursuit using deep reinforcement learning', *IEEE Robotics and Automation Letters*, **6**(3), 4552–4559, (2021).
- [8] Tingxiang Fan, Pinxin Long, Wenxi Liu, and Jia Pan, 'Distributed multi-robot collision avoidance via deep reinforcement learning for navigation in complex scenarios', *The International Journal of Robotics Research*, **39**(7), 856–892, (2020).
- [9] Pu Feng, Xin Yu, Junkang Liang, Wenjun Wu, and Yongkai Tian, 'Mact: Multi-agent collision avoidance with continuous transition reinforcement learning via mixup', in *International Conference on Swarm Intelligence*, pp. 74–85. Springer, (2023).
- [10] Jay P. Fillmore, 'A note on rotation matrices', *IEEE Computer Graphics and Applications*, **4**(2), 30–33, (1984).
- [11] Jun Guo, Yonghong Chen, Yihang Hao, Zixin Yin, Yin Yu, and Simin Li, 'Towards comprehensive testing on the robustness of cooperative multi-agent reinforcement learning', in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 115–122, (2022).
- [12] Hengyuan Hu, Adam Lerer, Alex Peysakhovich, and Jakob Foerster, "'other-play" for zero-shot coordination', in *International Conference on Machine Learning*, pp. 4399–4410. PMLR, (2020).
- [13] Jiechuan Jiang, Chen Dun, Tiejun Huang, and Zongqing Lu, 'Graph convolutional reinforcement learning', *arXiv preprint arXiv:1810.09202*, (2018).
- [14] George Em Karniadakis, Ioannis G Kevrekidis, Lu Lu, Paris Perdikaris, Sifan Wang, and Liu Yang, 'Physics-informed machine learning', *Nature Reviews Physics*, **3**(6), 422–440, (2021).
- [15] Naman Khetan, Tushar Arora, Samee Ur Rehman, and Deepak K Gupta, 'Implicit equivariance in convolutional networks', *arXiv preprint arXiv:2111.14157*, (2021).
- [16] Michael Laskin, Aravind Srinivas, and Pieter Abbeel, 'Curl: Contrastive unsupervised representations for reinforcement learning', in *International Conference on Machine Learning*, pp. 5639–5650. PMLR, (2020).
- [17] Misha Laskin, Kimin Lee, Adam Stooke, Lerrel Pinto, Pieter Abbeel, and Aravind Srinivas, 'Reinforcement learning with augmented data', *Advances in Neural Information Processing Systems*, **33**, (2020).
- [18] Yijiong Lin, Jiancong Huang, Matthieu Zimmer, Yisheng Guan, Juan Rojas, and Paul Weng, 'Invariant transform experience replay: Data augmentation for deep reinforcement learning', *IEEE Robotics and Automation Letters*, **5**(4), 6615–6622, (2020).
- [19] Michael L Littman, 'Markov games as a framework for multi-agent reinforcement learning', in *Machine learning proceedings 1994*, 157–163, Elsevier, (1994).
- [20] Yong Liu, Weixun Wang, Yujing Hu, Jianye Hao, Xingguo Chen, and Yang Gao, 'Multi-agent game abstraction via graph attention neural network', in *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pp. 7211–7218, (2020).
- [21] Ryan Lowe, Yi Wu, Aviv Tamar, Jean Harb, Pieter Abbeel, and Igor Mordatch, 'Multi-agent actor-critic for mixed cooperative-competitive environments', *arXiv preprint arXiv:1706.02275*, (2017).
- [22] Guozheng Ma, Zhen Wang, Zhecheng Yuan, Xueqian Wang, Bo Yuan, and Dacheng Tao, 'A comprehensive survey of data augmentation in visual reinforcement learning', *arXiv preprint arXiv:2210.04561*, (2022).
- [23] Viktor Makoviychuk, Lukasz Wawrzyniak, Yunrong Guo, Michelle Lu, Kier Storey, Miles Macklin, David Hoeller, Nikita Rudin, Arthur Allshire, Ankur Handa, et al., 'Isaac gym: High performance gpu-based physics simulation for robot learning', *arXiv preprint arXiv:2108.10470*, (2021).
- [24] Francesco Mondada, Michael Bonani, Xavier Raemy, James Pugh, Christopher Cianci, Adam Klaptocz, Stephane Magnenat, Jean-Christophe Zufferey, Dario Floreano, and Alcherio Martinoli, 'The e-puck, a robot designed for education in engineering', in *Proceedings of the 9th conference on autonomous robot systems and competitions*, volume 1, pp. 59–65. IPCB: Instituto Politécnico de Castelo Branco, (2009).
- [25] Igor Mordatch and Pieter Abbeel, 'Emergence of grounded compositional language in multi-agent populations', *arXiv preprint arXiv:1703.04908*, (2017).
- [26] Tabish Rashid, Mikayel Samvelyan, Christian Schroeder, Gregory Farquhar, Jakob Foerster, and Shimon Whiteson, 'Qmix: Monotonic value function factorisation for deep multi-agent reinforcement learning', in *International Conference on Machine Learning*, pp. 4295–4304. PMLR, (2018).
- [27] B. Ravindran and A. G. Barto, 'Symmetries and model minimization in markov decision processes', Technical report, USA, (2001).
- [28] Shital Shah, Debadepta Dey, Chris Lovett, and Ashish Kapoor, 'Airsim: High-fidelity visual and physical simulation for autonomous vehicles', in *Field and service robotics*, pp. 621–635. Springer, (2018).
- [29] Rongye Shi, Zhaobin Mo, and Xuan Di, 'Physics-informed deep learning for traffic state estimation: A hybrid paradigm informed by second-order traffic models', in *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pp. 540–547, (2021).
- [30] Rongye Shi, Zhaobin Mo, Kuang Huang, Xuan Di, and Qiang Du, 'A physics-informed deep learning paradigm for traffic state and fundamental diagram estimation', *IEEE Transactions on Intelligent Transportation Systems*, **23**(8), 11688–11698, (2022).
- [31] Ivan Sosnovik, Michał Szmaja, and Arnold Smeulders, 'Scale-equivariant steerable networks', *arXiv preprint arXiv:1910.11093*, (2019).
- [32] Elise van der Pol, Herke van Hoof, Frans A Oliehoek, and Max Welling, 'Multi-agent mdp homomorphic networks', *arXiv preprint arXiv:2110.04495*, (2021).
- [33] Elise van der Pol, Daniel Worrall, Herke van Hoof, Frans Oliehoek, and Max Welling, 'Mdp homomorphic networks: Group symmetries in reinforcement learning', *Advances in Neural Information Processing Systems*, **33**, (2020).
- [34] Dian Wang, Robin Walters, and Robert Platt, 'SO(2)-equivariant reinforcement learning', *arXiv preprint arXiv:2203.04439*, (2022).
- [35] Maurice Weiler, Fred A Hamprecht, and Martin Storath, 'Learning steerable filters for rotation equivariant cnns', in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 849–858, (2018).
- [36] Maurice Weiler, Fred A Hamprecht, and Martin Storath, 'Learning steerable filters for rotation equivariant cnns', in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 849–858, (2018).
- [37] Denis Yarats, Ilya Kostrikov, and Rob Fergus, 'Image augmentation is all you need: Regularizing deep reinforcement learning from pixels', in *International Conference on Learning Representations*, (2020).
- [38] Zhenhui Ye, Yining Chen, Xiaohong Jiang, Guanghua Song, Bowei Yang, and Sheng Fan, 'Improving sample efficiency in multi-agent actor-critic methods', *Applied Intelligence*, 1–14, (2021).
- [39] Chao Yu, Akash Velu, Eugene Vinitzky, Yu Wang, Alexandre Bayen, and Yi Wu, 'The surprising effectiveness of ppo in cooperative, multi-agent games', *arXiv preprint arXiv:2103.01955*, (2021).
- [40] Xin Yu, Wenjun Wu, Pu Feng, and Yongkai Tian, 'Swarm inverse reinforcement learning for biological systems', in *2021 IEEE International Conference on Bioinformatics and Biomedicine (BIBM)*, pp. 274–279. IEEE, (2021).
- [41] Martin Zinkevich and Tucker Balch, 'Symmetry in markov decision processes and its implications for single agent and multi agent learning', in *In Proceedings of the 18th International Conference on Machine Learning*. Citeseer, (2001).