

Leveraging Partial Symmetry for Multi-Agent Reinforcement Learning

Xin Yu¹, Rongye Shi^{2*}, Pu Feng¹, Yongkai Tian¹, Simin Li¹, Shuhao Liao¹, Wenjun Wu²

¹School of Computer Science and Engineering, Beihang University, Beijing, China

²Institute of Artificial Intelligence, Beihang University, Beijing, China

{nlsdeyuxin, shirongye, fengpu, tianyk, lisiminsimon, lsh_iai, wwj}@buaa.edu.cn

Abstract

Incorporating symmetry as an inductive bias into multi-agent reinforcement learning (MARL) has led to improvements in generalization, data efficiency, and physical consistency. While prior research has succeeded in using perfect symmetry prior, the realm of partial symmetry in the multi-agent domain remains unexplored. To fill in this gap, we introduce the partially symmetric Markov game, a new subclass of the Markov game. We then theoretically show that the performance error introduced by utilizing symmetry in MARL is bounded, implying that the symmetry prior can still be useful in MARL even in partial symmetry situations. Motivated by this insight, we propose the Partial Symmetry Exploitation (PSE) framework that is able to adaptively incorporate symmetry prior in MARL under different symmetry-breaking conditions. Specifically, by adaptively adjusting the exploitation of symmetry, our framework is able to achieve superior sample efficiency and overall performance of MARL algorithms. Extensive experiments are conducted to demonstrate the superior performance of the proposed framework over baselines. Finally, we implement the proposed framework in real-world multi-robot testbed to show its superiority.

1 Introduction

Multi-Agent Reinforcement Learning (MARL) is increasingly gaining attention due to its capabilities in handling complex tasks (Li et al. 2023). Such tasks often necessitate strategic interaction and rivalry among various entities (Yu et al. 2021b; Feng et al. 2023). However, a notorious limitation of most MARL approaches is their substantial data dependency, necessitating vast amounts of data to build an efficient model. This limitation severely narrows the scope of MARL’s practical application in real-life settings (Shi, Steenkiste, and Veloso 2021). How to develop strategies to improve MARL’s sample efficiency has become an important and long-standing research topic.

Strategies to augment sample efficiency often involve integrating external knowledge to accelerate MARL’s training. Various methods incorporating extra knowledge have been proposed in recent literature (Shi, Mo, and Di 2021; Shi et al. 2022). In the realm of MARL, prior research has highlighted

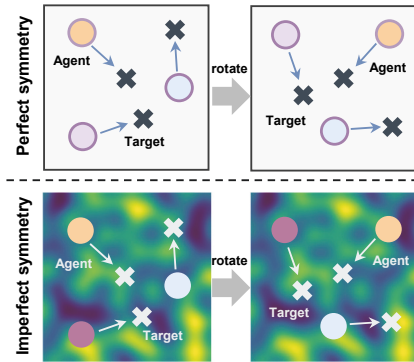


Figure 1: Illustration of symmetry disruption in a non-uniform field: despite the spatial symmetry of the multi-agent system, the introduction of a non-uniform field, such as uneven terrain or a wind field, disrupts this symmetry and the symmetry assumption does not strictly hold everywhere. Colors denote varying intensities of the field.

the advantage of employing permutation invariance. Permutation invariance asserts that the systemic behavior remains unaffected by any changes in the order of agent consideration (Jianye et al. 2023). The application of permutation invariance encourages extensive parameter sharing among agents, thereby augmenting data efficiency. Additionally, the most common symmetry in multi-agent systems is rotation symmetry, as illustrated in Figure 1. In this context, rotating the global state results in a rotation of the optimal joint policy. Some studies have ameliorated data efficiency by designing inherent network structures that satisfy this property (van der Pol et al. 2021).

Current techniques often assume the existence of perfect permutation invariance or perfect spatial symmetry. However, such ideal conditions are rare in real-world scenarios. For instance, again in Figure 1, multiple agents attempt to approach a target point where each agent can sense the environment, including information about other agents, obstacles, and the target point. Such problems, conditioned on the perfect symmetry transition function and symmetry reward function, are defined as symmetric Markov game in (van der Pol et al. 2021; Yu et al. 2023). Unfortunately, in the real world, there might exist imperfections in the environment,

*Corresponding author: Rongye Shi.

e.g., uneven ground, wind, and other non-uniform fields acting on the agents. The non-uniform fields can deviate the system’s transition dynamics or reward functions from perfect spatial symmetry to a certain extent. Specifically, when we rotate the state-action pairs of our agents, we cannot rotate the non-uniform fields, accordingly. As a result, despite the multi-agent system having a spatially symmetrical structure, its response to the non-uniform fields can no longer employ perfect symmetry. Furthermore, slight variations in elements such as power supply or physical structures could make the agents slightly heterogeneous, thus violating the principle of perfect permutation invariance. This violation poses a significant challenge to real-world implementations of symmetry-prior-based reinforcement learning methods.

Regrettably, existing studies, either single-agent or multi-agent ones, have seldomly explored such scenarios of partial symmetry, neither from a theoretical nor from a practical point of view. To emphasize the necessity of such a study, we evaluated the performance of the perfect symmetry network proposed in (van der Pol et al. 2021) under various symmetry-breaking conditions. As depicted in Figure 2, the performance of their EQ-MPN and MPN methods is evaluated under three distinct noise levels, which signifies the extent of symmetry-breaking introduced into the system. We found that as symmetry breaks, the performance of the network with embedded symmetry, EQ-MPN, deteriorates. Motivated by these challenges, we delve into the partial symmetry scenarios, targeting at a new methodology that relaxes the requirement of strict symmetry with a theoretical performance bound guaranteed.

In this paper, we first define the partially symmetric Markov game. It gives rise to a new class of symmetry Markov game with slack symmetry constraints while partially maintaining favorable inductive biases for learning. We theoretically show that the performance errors introduced by leveraging symmetry under partially symmetric Markov game are bounded. Our theoretical analysis can be seamlessly applied to a variety of symmetries, including permutation invariance, rotational equivariance, etc. Upon this setting, we introduce a Partial Symmetry Exploitation framework (PSE). PSE first quantifies the extent/level of symmetry in the environment using a dedicated symmetry quantification module and then selects an appropriate training pipeline according to that symmetry level. The PSE involves several technical components to adaptively incorporate symmetry into the training process. Our main contributions are listed as follows:

- Formally define the concept of partial equivariance and generalize symmetry Markov game to partially symmetric Markov game;
- For partially symmetric Markov game, theoretically show that the performance error introduced by utilizing symmetry in MARL is bounded;
- Motivated by the error bound, propose a novel PSE framework to adaptively incorporate and leverage symmetry prior in MARL;
- Demonstrate our framework’s superiority over baselines in both simulated tasks and real-world robot experiments.




Figure 2: Performance of the EQ-MPN and MPN under varying noise levels. As the noise intensity (the degree of symmetry-breaking) increases, the perfect symmetry network, EQ-MPN, exhibits a declining trend.

2 Related Work

2.1 Symmetries in Single-Agent RL

The methods for exploiting symmetry in RL can be broadly classified into two major categories: data augmentation and network structure design. Data augmentation in single-agent RL is to generate additional data through image transformations during the training phase of the model (Laskin et al. 2020; Yarats, Kostrikov, and Fergus 2020; Lin et al. 2020; Amadio, Colomé, and Torras 2019). Alternatively, symmetry can be introduced through a contrastive learning framework by enforcing consistencies between an image and its augmented version (Laskin, Srinivas, and Abbeel 2020). The network design method is to design specialized architectures that implicitly embed prior knowledge relevant to the task (Ravindran and Barto 2001). For instance, symmetries in the joint state-action space can be expressed through the implementation of policy networks (van der Pol et al. 2020; Wang, Walters, and Platt 2022). Our paper explores the realm of partial symmetry, extending beyond the scope of the approaches commonly employed.

2.2 Symmetries in Multi-Agent RL

In the realm of multi-agent systems, fewer studies have explored the use of data augmentation techniques. To our knowledge, the most related work to our work is the data augmentation method proposed in (Ye et al. 2021). This method generates additional data by implementing permutation transformations for homogeneous agents, interpreting data augmentation from the perspective of permutation invariance. In a similar vein, the need for more extensive integration of prior knowledge into MARL is apparent. Multi-Agent MDP Homomorphic Networks have been developed to embed symmetries, thus enhancing data efficiency (van der Pol et al. 2021). However, these methods impose strict constraints on symmetry, which hinders their applicability in real-world scenarios characterized by partial symmetry. In contrast, we treat symmetry as an additional objective and incorporate it through soft constraints such as data augmentation and regularization. Our approach is able to adjust to different symmetry levels, thereby improving algorithmic performance in scenarios with partial symmetry.

3 Preliminaries

3.1 Cooperative Markov Game

An n -agent cooperative Markov game (Boutilier 1996) can be defined as 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 agent $i = 1, \dots, n$. Let $A = A_1 \times A_2 \times \dots \times A_n$ be the joint action space, and $T : S \times A \times S \rightarrow [0, 1]$ be 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) and take independent action (a_1, \dots, a_N) relying 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 $\mathbb{E}_\pi [\sum_{t=0}^{\infty} \gamma^t R(s_t, a_t)]$, where $0 < \gamma < 1$ 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 $\eta : S \rightarrow [0, 1]$.

3.2 Groups and Transformations

This section offers an overview of the concepts of groups and transformations (Bronstein et al. 2021). A group G is a set equipped with a binary operator that has four mathematical properties: identity, inverse, closure, and associativity. Our discussion primarily revolves around the group $SO(2)$ and its cyclic subgroup C_n . Specifically, $SO(2)$ represents the group of continuous rotations $\{R_\theta : 0 \leq \theta < 2\pi\}$. Meanwhile, C_n stands for the discrete subgroup, defined as $C_n = \{R_\theta : \theta \in \{\frac{2\pi i}{n} \mid 0 \leq i < n\}\}$. A rotation matrix illustrates the act of rotating within Euclidean space (Fillmore 1984). For a specific rotation set $\{0^\circ, 90^\circ, 180^\circ, 270^\circ\}$, the rotation matrix is formulated as:

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

The four group axioms are satisfied in the case of a rotation transformation.

3.3 Equivariance and Invariance

In multi-agent systems, the symmetries are commonly referred to as equivariance and invariance (Yu et al. 2023). 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:

$$K_g[f(x)] = f(L_g[x]),$$

where $g \in G$ and G is a mathematical group, then, function f is equivariant to the transformation L_g . The operators L_g and K_g can be used to describe the same transformation, but in different spaces. 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 (upper half) shows the equivariance of the optimal policy, rotating the state globally results in a transformation of the optimal policy. Given two states s and $L_g[s]$, the optimal policy π^* is equivariant to its transformation which is denoted by $K_g[\pi^*(s)] = \pi^*(L_g[s])$. Without special notice, the transformations L_g, K_g are assumed to be *bijective* in this paper.

4 Defining and Characterizing Partially Symmetric Markov Game

4.1 Partial Equivariance and Invariance

Real-world dynamics may not satisfy the strict equivariance though such a strict assumption has been commonly used for simplicity in literature (Wang, Walters, and Platt 2022; van der Pol et al. 2020). In this paper, we introduce a definition of partial equivariance and invariance to fill the gap.

Definition 1 (Partial 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:

$$\|K_g[f(x)] - f(L_g[x])\| \leq \epsilon,$$

where $g \in G$ and G is a mathematical group, we say f is ϵ -partially equivariant to the transformation L_g and K_g . A related notion to ϵ -partial equivariance is the ϵ -partial invariance: If for any choice of $g \in G$ we have $K_g = I$, then function f is ϵ -partially invariant to transformation L_g . Note that strict equivariance or invariance are special cases of partial ones with $\epsilon = 0$.

4.2 Partially Symmetric Markov Game

In this subsection, we formally define the partially symmetric Markov game, a subclass of the cooperative Markov game characterized by partial symmetry.

Definition 2 (Partially Symmetric Markov game). The partially symmetric Markov game $\mathcal{M}_g = (N, S, \{A_i\}_{i=1}^n, R, T, \Psi, g, \epsilon, \delta)$ is a cooperative Markov game that satisfies the conditions of partial reward invariance and partial transition invariance. The state and action transformation 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 transformed state-action pairs as (gs, ga) . Here, $gs = L_g(s)$ and $ga = K_g(a)$ for short. The partial reward invariance is characterized with:

$$|R(s, a) - R(gs, ga)| \leq \epsilon. \quad (1)$$

The partial transition invariance is characterized using the Maximum Mean Discrepancy (MMD) between the distributions $T(s'|s, a)$ and $T(gs'|gs, ga)$:

$$\text{MMD}_{\mathcal{F}}(T(s'|s, a), T(gs'|gs, ga)) \leq \delta. \quad (2)$$

Please note that $T(gs'|gs, ga)$ is also a distribution of s' with the sampling process $s' \sim T(gs'|gs, ga)$ being defined as 1) $gs' \sim T(gs'|gs, ga)$ and then 2) $s' = L_g^{-1}(gs')$.

The MMD is defined as a measure of distance used to quantify the discrepancy between two distributions under a general class of bounded mapping functions $f \in \mathcal{F}$ (Gretton et al. 2006). Eq.(2) can be expressed as:

$$\begin{aligned} \text{MMD}_{\mathcal{F}}(T(s'|s, a), T(gs'|gs, ga)) \\ = \sup_{f \in \mathcal{F}} |\mathbb{E}_{s' \sim T(s'|s, a)}[f(s')] - \mathbb{E}_{s' \sim T(gs'|gs, ga)}[f(s')]|. \end{aligned}$$

We believe that the symmetry property is a common occurrence in many real-world multi-agent tasks. Our formal definition of this issue holds theoretical and practical significance.

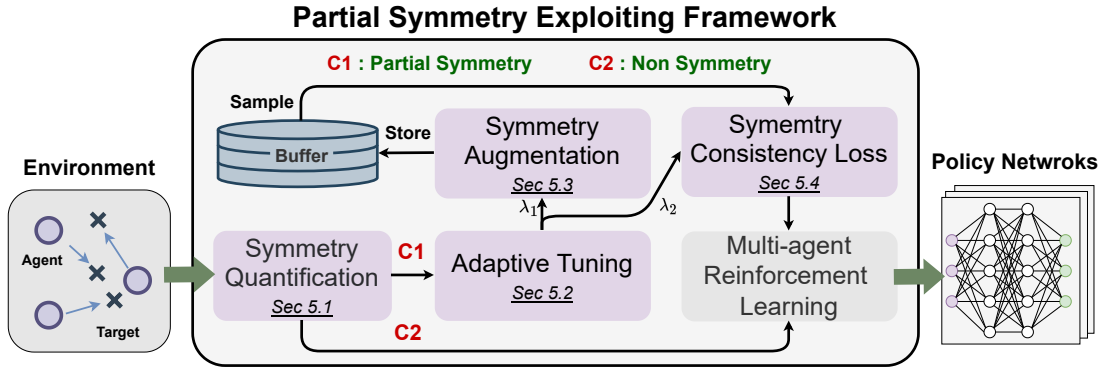


Figure 3: The overall framework of the proposed PSE. The framework is composed of four key modules: 1) Symmetry Quantification, which measures the level of symmetry in the environment, 2) Adaptive Tuning, which serves as the annealing coefficient modulating the continuous degree of symmetry utilization. 3) Symmetry Augmentation, which manipulates the data based on the quantified symmetry, and 4) Symmetry Loss, a specially crafted function that optimizes the policy network with respect to the symmetry.

4.3 Performance Error Analysis

We show that if a problem can be formulated as a partially symmetric Markov game, the performance error introduced by using symmetry-augmented data in training is bounded. 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). Based on the definition in (Ravindran and Barto 2001), the m -step optimal discounted action value function recursively for all $(s, a) \in \Psi$ and for all non-negative integers m is defined as follows:

$$Q_m(s, a) = R(s, a) + \gamma \sum_{s' \in S} [T(s'|s, a) \max_{a' \in A} Q_{m-1}(s', a')].$$

The optimal action-value function $Q^*(s, a)$ is the limit of $Q_m(s, a)$ as m approaches infinity. We now define the performance error for \mathcal{M}_g , which measures the error of a Q-function trained on symmetry-augmented data.

Definition 3 (Performance Error for \mathcal{M}_g when using symmetry-augmented data). Let $Q^*(s, a)$ be the optimal action-value function, g be the transformation associated with \mathcal{M}_g . Then, the performance error of \mathcal{M}_g is defined as:

$$Error_{\mathcal{M}_g} = |Q^*(s, a) - Q^*(gs, ga)|,$$

where $Q^*(gs, ga)$ is the action-value function trained by the symmetry-augmented data.

Proposition 1 (Performance Error Bound). *If a partially symmetric Markov game \mathcal{M}_g satisfies the conditions in equations (1) and (2) for all $(s, a) \in \Psi$, then the performance error $Error_{\mathcal{M}_g} = |Q^*(s, a) - Q^*(gs, ga)|$ is bounded by $\frac{\epsilon}{1-\gamma} + \frac{\gamma\delta}{1-\gamma}$.*

As stated in Prop 1, for the Partially Symmetric Markov game \mathcal{M}_g , the error introduced by incorporating symmetry samples is bounded. The proof of Prop 1 can be found in Section 1 of the Appendix¹. Prop 1 implies that the symmetry-augmented data are useful in MARL to a bounded extent even in partial symmetry situations like \mathcal{M}_g .

¹Video demonstrations and Supplementary materials are available at the project website <https://xinyu-site.github.io/PSE/>.

5 Framework of the Partial Symmetry Exploitation

This paper focuses on solving the following problem: *In the context of partially symmetric Markov game \mathcal{M}_g , how can we appropriately leverage the symmetry prior to improving sample efficiency and performance of MARL?* On top of Prop 1, we propose a general framework, called Partial Symmetry Exploitation (PSE), for exploiting the symmetry prior, properly. The PSE framework is designed to adaptively utilize symmetry and is composed of four key modules: Symmetry Quantification, Adaptive Tuning, Symmetry Augmentation, and Symmetry Consistency Loss.

5.1 Symmetry Quantification

We propose a quantification method to measure the symmetry in \mathcal{M}_g . This method is applicable to various symmetries inherent in multi-agent systems, including permutation invariance, rotational equivariance, and translational invariance. We employ a transformed environment to assess the degree of symmetry by comparing the system’s responses in both the original and transformed environments. For a partially symmetric Markov game \mathcal{M}_g , consider (s, a, s') with an associated transformation g . In the transformed environment, action ga is applied to the state gs , leading to a new state s' . We define a function $D(gs', \bar{s}')$ to measure the degree of symmetry in MARL:

$$D(gs', \bar{s}') = 1 - \frac{1}{2} \frac{\|gs' - \bar{s}'\|_2^2}{\|gs'\|_2^2 + \|\bar{s}'\|_2^2}, \quad (3)$$

where the numerator represents the Euclidean norm of the difference between the vectors, and the denominator is the sum of individual Euclidean norms. Given a vector \mathbf{v} , its Euclidean norm is defined as $\|\mathbf{v}\|_2 = \sqrt{\sum_i v_i^2}$. D is a scalar and its values lie within the interval $[0, 1]$. A value of 1 indicates perfect symmetry and the lower bound of D is attained when $gs' = -\bar{s}'$. While we utilize the Euclidean norm in this context, other norms can also be employed.

Upon obtaining the measure of symmetry, we can introduce threshold values to determine the level of symmetry, categorizing the environment into Partial Symmetry (C1) and Non-Symmetry (C2). It's worth noting that perfect symmetry is a special case of partial symmetry. The threshold can be tuned according to the specific requirements of the problem and the performance trade-offs acceptable.

5.2 Adaptive Tuning

In the early stages of training in \mathcal{M}_g , symmetry can assist the model to converge more swiftly and reduce the loss faster. However, as the model progressively adapts to the training environment and starts to capture more nuanced features of the data, an over-reliance on symmetry might have a negative effect on the model's training. To tackle this issue, as training advances, our PSE gradually reduces the dependence on symmetry. To this end, we present the following function:

$$\lambda(D, k) = De^{-\beta k}, \quad (4)$$

which serves as the annealing coefficient at the k^{th} iteration, with D signifying the degree of symmetry in Eq. (3) and β denoting the decay rate. In the follow-up stage, as will be discussed later, the coefficient $\lambda(D, k)$ serves as both 1) a probability to decide on whether the symmetry-augmented data is used and 2) a coefficient in the objective function to weigh the component of symmetry constraints. This auto-tuning approach strikes an adaptive balance on the extent to which the symmetry is leveraged in different training phases.

5.3 Symmetry Augmentation

One straightforward way to leverage symmetry is through data augmentation. Motivated by Prop 1, we present a data augmentation strategy designed to adaptively leverage symmetry. Based on Eq. (4), we obtain a coefficient $\lambda_1 = De^{-\beta_1 k}$ that starts with a value equal to the degree of symmetry and decreases over training iterations. Specifically, the coefficient λ_1 acts as a probabilistic threshold. If a random number r drawn from a uniform distribution between $[0,1]$ is less than λ_1 , data augmentation is applied in that iteration of training. This strategy ensures an expedited training process in the early phases. As training progresses, reliance on symmetry-augmented samples is reduced, thereby mitigating potential performance errors they might introduce.

5.4 Symmetry Consistency Loss

In multi-agent settings, using data augmentation to improve sample efficiency can be challenging. The reason is that when multiple agents are considered, more sources of variance are introduced, making the training unstable. The proposed symmetry consistency loss provides mitigation to this challenge (see Section 2 in the Appendix for more details). For a clean presentation, the MAPPO is used as an example to introduce symmetry consistency loss.

Symmetry Consistency Loss. The policy consistency loss term $S_\pi(\theta)$ is defined as

$$S_\pi = KL[\pi_\theta(ga | gs) | \pi_\theta(a | s)], \quad (5)$$

aiming to constrain distribution $\pi_\theta(ga | gs)$ to be close to $\pi_\theta(a | s)$. This helps guide the training process according




Figure 4: The simulated tasks considered in the experiments.

to the symmetry prior. Assume that $V_\psi(s)$ represents an approximate value for state s , the symmetry consistency loss for value function is designed as

$$S_V = \mathbb{E}_s [(V_\psi(s) - V_\psi(gs))^2], \quad (6)$$

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 Eqs. (5) and (6) as the symmetry consistency loss.

MARL with Symmetry Consistency Loss. Rather than having a fixed coefficient in the loss function, we utilize $\lambda_2 = De^{-\beta_2 k}$ calculated by Eq. (4) to dynamically adjust the coefficient of symmetric loss. For the MAPPO, our PSE method optimizes the following loss objective:

$$J_{PSE} = J_{MAPPO} - \lambda_2(S_\pi + S_V). \quad (7)$$

The J_{MAPPO} is the objective function for MAPPO and can be found in Section 2 of the Appendix.

6 Experiments

This section demonstrates our PSE's superiority via experiments in both simulated tasks and real-world robot systems.

6.1 Environmental Settings

We conducted experiments in several tasks, including Predator-Prey (PP), Cooperative Navigation (CN), Wildlife Monitoring, and Formation Change (FC). CN and PP is a classic scenario implemented in multi-agent particle environment (Mordatch and Abbeel 2017). The wildlife monitoring is a grid-world-based environment, where a set of drones has to coordinate to accomplish the task (van der Pol et al. 2020). The goal is to trap poachers by having one drone hover above them while the other assists from the side. The details of FC are provided in Section 3 of the Appendix. Figure 4 shows parts of the tasks.

Symmetry Breaking. The original classic setups of these four tasks were designed to have perfect spatial symmetry. We deliberately modified these tasks to incorporate partial symmetry by integrating noise into the transition dynamics. The multi-robot FC task was conducted in the high-precision robot simulation environment Webots. As shown in Figure 4c, robots were required to learn to avoid each other as well as the obstacles and coordinate to reach their destinations. We set an uneven terrain for this task to simulate partial symmetry. The severity of these uneven conditions can be

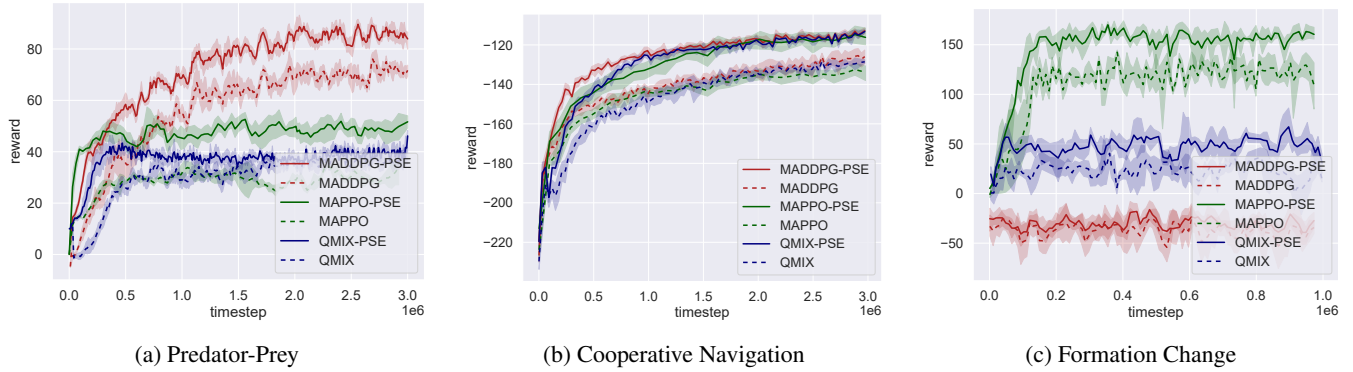


Figure 5: Learning curves of the baseline and their versions with the PSE framework on the three multi-agent tasks.




Figure 6: The learning curve of MPN, EQ-MPN and MPN-PSE under different symmetry-breaking in the Wildlife Monitoring.

modified, allowing us to explore the influences of environmental uncertainties on MARL in depth. For more details on symmetry breaking, please see Section 4 of the Appendix.

Baselines. The proposed PSE framework was applied to several baselines, including Multi-Agent Deep Deterministic Policy Gradient (MADDPG), Monotonic Value Function Factorisation for Deep MARL (QMIX), and Multi-Agent Proximal Policy Optimization (MAPPO), which are mainstream MARL approaches (Rashid et al. 2018; Lowe et al. 2017; Yu et al. 2021a).

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 10 different random seeds, and the final experimental results under partial symmetry are shown in Figure 5. The results show that the MARL algorithms adopting the PSE 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 5a, the proposed PSE framework outperformed the baseline methods significantly. The results indicated that the proposed framework could improve the data efficiency, convergence speed, and performance in terms of evaluation rewards.

Cooperative Navigation. The cooperative navigation was a fully cooperative environment, where 3 agents (cir-

cles) cooperated to reach 3 landmarks (crosses) under a minimum number of collisions. Similarly, as shown in Figure 5b, the results show that the proposed framework can improve data efficiency and performance 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, 8 robots started in a square formation and had their destinations set on the opposite side. The experimental results show that the MADDPG and QMIX could not learn a useful policy in this task, whereas the agents trained by the MAPPO-PSE and MAPPO could reach the destination while avoiding collisions with each other and obstacles. As presented in Figure 5c, the algorithms enhanced by the proposed framework obtained higher rewards than the original versions. This indicates that the proposed PSE can further improve the performance of MARL algorithms in challenging environments.

Wildlife Monitoring. We conducted a comparison of three graph network-based methods: the Message Passing Network (MPN, a classic graph convolutional network), the EQ-MPN (an advanced baseline network embedding perfect symmetry as proposed in (van der Pol et al. 2021)) and MPN with the PSE framework (MPN-PSE). The Wildlife Monitoring is specially analyzed here because it is the chosen environment in the literature well suited to the EQ-MPN, featuring pixel-based and grid-world states. Advanced re-




Figure 7: Convergence reward under varying noise intensities of various models across PP, CN, and FC scenarios.

sults are achieved by EQ-MPN when the perfect symmetry in the environment holds. Here, we assessed the three models’ performance across varying degrees of symmetry-breaking. Under conditions of perfect symmetry (see Figure 6a), the EQ-MPN with the perfect symmetry prior embedded in its network structure demonstrates superior performance compared to the classic MPN. However, our PSE framework surpasses both EQ-MPN and MPN in terms of convergence speed and the quality of the final convergence. As shown in Figure 6b, under partial symmetry, the performance of the EQ-MPN deteriorates, even to a level worse than the classic MPN. In contrast, our PSE-based method continues to enhance the performance of MPN. In Figure 6c, where the environment is completely devoid of any symmetry, all three methods face challenges in learning an acceptable policy. Yet, our PSE framework still maintains a discernible advantage.

6.3 Impact of Different Degrees of Symmetry and Ablation Analysis

We analyzed four distinct algorithm variations in our study: 1) MAPPO, which stands for the most primitive version of the algorithm. 2) MAPPO-SE, which retains only our symmetry augmentation and loss function components within the original MAPPO, and where the coefficient of the loss function is fixed to 0.5. 3) the MAPPO-PSE, our comprehensive framework representing the entirety of our proposed enhancements. Figure 7 denotes each algorithm type with a distinct color, and different algorithm variations are highlighted by varying line types. It is observed that the PSE framework consistently excels across different degrees of symmetry-breaking. Interestingly, MAPPO-SE, which sticks to leveraging the perfect symmetry, experiences a substantial performance decline as noise intensity increases, even deteriorating to a level worse than the classic MAPPO.

The results provide two insights: 1) a strong dependency on embedding perfect symmetry may seriously hamper the training and the final performance when the symmetry keeps breaking, and 2) our PSE framework can adapt to various symmetry-breaking conditions and consistently enhance the performance of mainstream multi-agent algorithms. The PSE enjoys this advantage due to the framework’s symmetry quantification and adaptive tuning components. The same




Figure 8: Real-world formation change on a swarm of robots. The robots successfully switched their positions to the antipodal points by achieving collision avoidance.

experiments are also conducted based on QMIX and MADDPG, which are included in Figure 7, and similar observations and conclusions can be obtained. The exception is that the MADDPG does not perform well in the FC task.

6.4 Real World Experiments

As shown in Figure 8, the real-world version of formation change presented in Section 6.1 was considered in this experiment. The trained policies were deployed on the Epuck, which is a small, lightweight robot platform. We followed a direct sim2real paradigm to deploy the policy network (De Souza et al. 2021). By incorporating our PSE 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 MAPPO-PSE is 2.1%, while the rate for MAPPO is 5.6%. Details are provided in Section 5 of the Appendix.

7 Conclusion

In this paper, we newly introduce the partially symmetric Markov game. We then theoretically show that the corresponding performance error is bounded. Based on the bounded property, we propose a novel PSE framework to adaptively leverage symmetry prior in MARL. Experimental results support the superiority of PSE over baselines. In the future, we plan to extend the PSE to systems with heterogeneous agents, whose sensitivity to the symmetry-breaking conditions is different.

Acknowledgments

This work was supported by the National Key Research and Development Program of China (Grant No. 2022ZD0117801), and the National Natural Science Foundation of China (Grant No. 62306023).

References

- Amadio, F.; Colomé, A.; and Torras, C. 2019. Exploiting symmetries in reinforcement learning of bimanual robotic tasks. *IEEE Robotics and Automation Letters*, 4(2): 1838–1845.
- Boutillier, C. 1996. Planning, learning and coordination in multiagent decision processes. In *TARK*, volume 96, 195–210. Citeseer.
- Bronstein, M. M.; Bruna, J.; Cohen, T.; and Velicković, P. 2021. Geometric deep learning: Grids, groups, graphs, geodesics, and gauges. *arXiv preprint arXiv:2104.13478*.
- De Souza, C.; Newbury, R.; Cosgun, A.; Castillo, P.; Vidolov, B.; and Kulić, D. 2021. Decentralized multi-agent pursuit using deep reinforcement learning. *IEEE Robotics and Automation Letters*, 6(3): 4552–4559.
- Feng, P.; Yu, X.; Liang, J.; Wu, W.; and Tian, Y. 2023. MACT: Multi-agent Collision Avoidance with Continuous Transition Reinforcement Learning via Mixup. In *International Conference on Swarm Intelligence*, 74–85. Springer.
- Fillmore, J. P. 1984. A Note on Rotation Matrices. *IEEE Computer Graphics and Applications*, 4(2): 30–33.
- Gretton, A.; Borgwardt, K.; Rasch, M.; Schölkopf, B.; and Smola, A. 2006. A kernel method for the two-sample problem. *Advances in neural information processing systems*, 19.
- Jianye, H.; Hao, X.; Mao, H.; Wang, W.; Yang, Y.; Li, D.; Zheng, Y.; and Wang, Z. 2023. Boosting Multiagent Reinforcement Learning via Permutation Invariant and Permutation Equivariant Networks. In *The Eleventh International Conference on Learning Representations*.
- Laskin, M.; Lee, K.; Stooke, A.; Pinto, L.; Abbeel, P.; and Srinivas, A. 2020. Reinforcement Learning with Augmented Data. *Advances in Neural Information Processing Systems*, 33.
- Laskin, M.; Srinivas, A.; and Abbeel, P. 2020. Curl: Contrastive unsupervised representations for reinforcement learning. In *International Conference on Machine Learning*, 5639–5650. PMLR.
- Li, S.; Guo, J.; Xiu, J.; Yu, X.; Wang, J.; Liu, A.; Yang, Y.; and Liu, X. 2023. Byzantine Robust Cooperative Multi-Agent Reinforcement Learning as a Bayesian Game. *arXiv preprint arXiv:2305.12872*.
- Lin, Y.; Huang, J.; Zimmer, M.; Guan, Y.; Rojas, J.; and Weng, P. 2020. Invariant transform experience replay: Data augmentation for deep reinforcement learning. *IEEE Robotics and Automation Letters*, 5(4): 6615–6622.
- Lowe, R.; Wu, Y.; Tamar, A.; Harb, J.; Abbeel, P.; and Mordatch, I. 2017. Multi-agent actor-critic for mixed cooperative-competitive environments. *arXiv preprint arXiv:1706.02275*.
- Mordatch, I.; and Abbeel, P. 2017. Emergence of Grounded Compositional Language in Multi-Agent Populations. *arXiv preprint arXiv:1703.04908*.
- Rashid, T.; Samvelyan, M.; Schroeder, C.; Farquhar, G.; Foerster, J.; and Whiteson, S. 2018. Qmix: Monotonic value function factorisation for deep multi-agent reinforcement learning. In *International Conference on Machine Learning*, 4295–4304. PMLR. ISBN 2640-3498.
- Ravindran, B.; and Barto, A. G. 2001. Symmetries and Model Minimization in Markov Decision Processes. Technical report, University of Massachusetts, Amherst, MA, United States.
- Shi, R.; Mo, Z.; and Di, X. 2021. 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, 540–547.
- Shi, R.; Mo, Z.; Huang, K.; Di, X.; and Du, Q. 2022. A physics-informed deep learning paradigm for traffic state and fundamental diagram estimation. *IEEE Transactions on Intelligent Transportation Systems*, 23: 11688–11698.
- Shi, R.; Steenkiste, P.; and Veloso, M. M. 2021. Improving the on-vehicle experience of passengers through SC-M*: A scalable multi-passenger multi-criteria mobility planner. *IEEE Transactions on Intelligent Transportation Systems*, 22(2): 1026–1040.
- van der Pol, E.; van Hoof, H.; Oliehoek, F. A.; and Welling, M. 2021. Multi-Agent MDP Homomorphic Networks. *arXiv preprint arXiv:2110.04495*.
- van der Pol, E.; Worrall, D.; van Hoof, H.; Oliehoek, F.; and Welling, M. 2020. MDP homomorphic networks: Group symmetries in reinforcement learning. *Advances in Neural Information Processing Systems*, 33.
- Wang, D.; Walters, R.; and Platt, R. 2022. SO(2)-Equivariant Reinforcement Learning. *arXiv preprint arXiv:2203.04439*.
- Yarats, D.; Kostrikov, I.; and Fergus, R. 2020. Image augmentation is all you need: Regularizing deep reinforcement learning from pixels. In *International Conference on Learning Representations*.
- Ye, Z.; Chen, Y.; Jiang, X.; Song, G.; Yang, B.; and Fan, S. 2021. Improving sample efficiency in Multi-Agent Actor-Critic methods. *Applied Intelligence*, 1–14.
- Yu, C.; Velu, A.; Vinitsky, E.; Wang, Y.; Bayen, A.; and Wu, Y. 2021a. The Surprising Effectiveness of PPO in Cooperative, Multi-Agent Games. *arXiv preprint arXiv:2103.01955*.
- Yu, X.; Shi, R.; Feng, P.; Tian, Y.; Luo, J.; and Wu, W. 2023. ESP: Exploiting Symmetry Prior for Multi-Agent Reinforcement Learning. In *ECAI 2023*, 2946–2953. IOS Press.
- Yu, X.; Wu, W.; Feng, P.; and Tian, Y. 2021b. Swarm inverse reinforcement learning for biological systems. In *2021 IEEE International Conference on Bioinformatics and Biomedicine (BIBM)*, 274–279. IEEE.