
Equivariant Networks for Zero-Shot Coordination

Darius Muglich
University of Oxford
dariusm1997@yahoo.com

Christian Schroeder de Witt
FLAIR, University of Oxford
cs@robots.ox.ac.uk

Elise van der Pol
Microsoft Research AI4Science
evanderpol@microsoft.com

Shimon Whiteson
University of Oxford
shimon.whiteson@cs.ox.ac.uk

Jakob Foerster
FLAIR, University of Oxford
jakob.foerster@eng.ox.ac.uk

Abstract

Successful coordination in Dec-POMDPs requires agents to adopt robust strategies and interpretable styles of play for their partner. A common failure mode is *symmetry breaking*, when agents arbitrarily converge on one out of many equivalent but mutually incompatible policies. Commonly these examples include partial observability, e.g. waving your right hand vs. left hand to convey a covert message. In this paper, we present a novel equivariant network architecture for use in Dec-POMDPs that prevents the agent from learning policies which break symmetries, doing so more effectively than prior methods. Our method also acts as a “coordination-improvement operator” for generic, pre-trained policies, and thus may be applied at test-time in conjunction with any self-play algorithm. We provide theoretical guarantees of our work and test on the AI benchmark task of Hanabi, where we demonstrate our methods outperforming other symmetry-aware baselines in zero-shot coordination, as well as able to improve the coordination ability of a variety of pre-trained policies. In particular, we show our method can be used to improve on the state of the art for zero-shot coordination on the Hanabi benchmark.

1 Introduction

A popular method for learning strategies in partially-observable cooperative Markov games is via *self-play* (SP), where a joint policy controls all players during training and at test time [21, 37, 45, 53]. While SP can yield highly effective strategies, these strategies often fail in *zero-shot coordination* (ZSC), where independently trained strategies are paired together for one step at test time. A common cause for this failure is mutually incompatible symmetry breaking (e.g. signalling 0 to refer to a “cat” vs. a “dog”). Specifically, this symmetry breaking results in policies that act *differently* in situations which are equivalent with respect to the symmetries of the environment [22, 38].

To address this, we are interested in devising *equivariant* strategies, which we illustrate in Figure 1. Equivariant policies are such that symmetric changes to their observation cause a corresponding change to their output. In doing so, we fundamentally prevent the agent from breaking symmetries over the course of training.

In earlier work, the *Other-Play* (OP) learning algorithm [22] addressed symmetry in this setting by training agents to maximize expected return when randomly matched with symmetry-equivalent policies of their training time partners. This learning rule is implemented by independently applying a random symmetry permutation to the action-observation history of each agent during each episode. OP, however, has some pitfalls, namely: 1) it does not guarantee equivariance; 2) it poses the burden of learning environmental symmetry onto the agent; and 3) it is only applicable during *training time* and cannot be used as a policy-improvement operator.

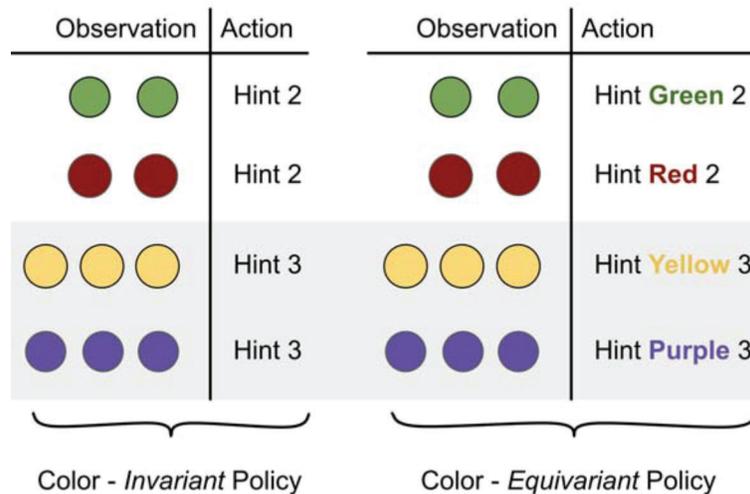


Figure 1: Illustrating different kinds of symmetry-robust policies, where the symmetries in this example are colors. Policies make actions based on the number of circles and the color. Invariant policies act irrespective of the change in color (i.e. only the number of circles matters), while equivariant policies act in correspondence with the change in color (i.e. changing the color will cause a corresponding change to the action).

To address all of this, our approach is to encode symmetry not in the data, but in the *network itself*.

There are many approaches to building equivariant networks in the literature. Most work focuses on analytical derivations of equivariant layers [11, 49, 51, 52], which can prove time-consuming to engineer for each problem instance and inhibits rapid prototyping. Algorithmic approaches to constructing equivariant networks are restricted to layer-wise constraints [10, 17, 18, 35], which require computational overhead to enforce. Furthermore, for complex problems such as ZSC in challenging benchmarks such as Hanabi [5], these approaches are insufficient because 1) they are not scalable; and 2) they usually do not allow for recurrent layers, which are the standard method for encoding action-observation histories in partially observable settings.

In this paper, we propose the *equivariant coordinator* (EQC), which uses our novel, scalable approach towards equivariant modelling in Dec-POMDPs. Specifically, our main contributions are:

1. Our method, EQC, that mathematically guarantees symmetry-equivariance of multi-agent policies, and can be applied solely at test time as a coordination-improvement operator; and
2. Showing that EQC outperforms prior symmetry-robust baselines on the AI benchmark Hanabi.

2 Background

This section formalises the problem setting (2.1), provides a brief primer to some group-theoretic notions (2.2), introduces equivariance (2.3), provides background to prior multi-agent work on symmetry (2.4), and offers a small, concrete example to help tie these notions together and facilitate intuition (2.5).

2.1 Dec-POMDPs

We formalise the cooperative multi-agent setting as a decentralized partially-observable Markov decision process (Dec-POMDP) [32] which is a 9-tuple $(\mathcal{S}, \mathcal{N}, \{\mathcal{A}^i\}_{i=1}^n, \{\mathcal{O}^i\}_{i=1}^n, \mathcal{T}, \mathcal{R}, \{U^i\}_{i=1}^n, T, \gamma)$, for finite sets $\mathcal{S}, \mathcal{N}, \{\mathcal{A}^i\}_{i=1}^n, \{\mathcal{O}^i\}_{i=1}^n$, denoting the set of states, agents, actions and observations, respectively, where a superscript i denotes the set pertaining to agent $i \in \mathcal{N} = \{1, \dots, n\}$ (i.e., \mathcal{A}^i and \mathcal{O}^i are the action and observation sets for agent i , and $a^i \in \mathcal{A}^i$ and $o^i \in \mathcal{O}^i$ are a specific action and observation of agent i). We also write $\mathcal{A} = \times_i \mathcal{A}^i$ and $\mathcal{O} = \times_i \mathcal{O}^i$, the sets of joint actions

and observations, respectively. $s_t \in \mathcal{S}$ is the state at time t and $s_t = \{s_t^k\}_k$, where s_t^k is state feature k of s_t . $a_t \in \mathcal{A}$ is the joint action of all agents taken at time t , which changes the state according to the transition distribution $s_{t+1} \sim \mathcal{T}(s_{t+1} | s_t, a_t)$. The subsequent joint observation of the agents is $o_{t+1} \in \mathcal{O}$, distributed according to $o_{t+1} \sim \mathcal{U}(o_{t+1} | s_{t+1}, a_t)$, where $\mathcal{U} = \times_i \mathcal{U}^i$; observation features of $o_{t+1}^i \in \mathcal{O}^i$ are notated analogously to state features; that is, $o_{t+1}^i = \{o_{t+1}^{i,k}\}_k$. The reward $r_{t+1} \in \mathbb{R}$ is distributed according to $r_{t+1} \sim \mathcal{R}(r_{t+1} | s_{t+1}, a_t)$. T is the horizon and $\gamma \in [0, 1]$ is the discount factor.

Notating $\tau_t^i = (a_0^i, o_1^i, \dots, a_{t-1}^i, o_t^i)$ for the action-observation history of agent i , agent i acts according to a policy $a_t^i \sim \pi^i(a_t^i | \tau_t^i)$. The agents seek to maximize the return, i.e., the expected discounted sum of rewards:

$$J_T := \mathbb{E}_{p(\tau_T)} \left[\sum_{t' \leq T} \gamma^{t'-1} r_{t'} \right], \quad (1)$$

where $\tau_t = (s_0, a_0, o_1, r_1, \dots, a_{t-1}, o_t, r_t, s_t)$ is the trajectory until time t .

2.2 Groups

A group G is a set with an associative binary operation on G , such that, relative to the binary operation, the inverse of each element in the set is also in the set, and the set contains an identity element. For our purposes, each element $g \in G$ can be thought of as an automorphism over some space X (i.e. an isomorphism from X to X), and the binary operation is function composition. The cardinality of the group is referred to as its order. If a subset K of G is also a group under the binary operation of G (a.k.a. a subgroup), we denote it $K < G$.

A group action¹ is a function $G \times X \rightarrow X$ satisfying $ex = x$ (where e is the identity element) and $(g \cdot h)x = g \cdot (hx)$, for all $g, h \in G$. In particular, for a fixed g , we get a function $\alpha_g : X \rightarrow X$ given by $x \mapsto gx$. g can be thought of as “relabeling” the elements of X , which aligns with our definition for Dec-POMDP symmetry introduced Section 2.4.

Groups have specific mathematical properties, such as associativity, containing the inverse of each element, and closure (closure is implicit of the binary operation being over G itself) which are necessary for proving Propositions 1 and 2 and ensuring equivariance.

2.3 Equivariance

Let G be a group. Consider the set of all neural networks Ψ of a given architecture whose first and ultimate layers are linear; that is, $\Psi := \{\psi | \psi = h(\hat{\psi}(f))\}$, where f, h are linear functions, and $\hat{\psi}$ is a non-linear function. We say a network $\psi \in \Psi$ is equivariant (with respect to a group G , or G -equivariant) if

$$\mathbf{K}_g \psi(\mathbf{x}) = \psi(\mathbf{L}_g \mathbf{x}), \text{ for all } g \in G, \mathbf{x} \in \mathbb{R}^d, \quad (2)$$

where \mathbf{L}_g is the group action of $g \in G$ on inputs, \mathbf{K}_g is similarly the group action on outputs, and d is the dimension of our data. If ψ satisfies Equation 2, then we say $\psi \in \Psi_{\text{equiv}}$. In the context of reinforcement learning, one may consider ψ as the neural parameterization of a policy π ², so a policy is G -equivariant if

$$\pi(\mathbf{K}_g(a) | \tau) = \pi(a | \mathbf{L}_g(\tau)), \text{ for all } g \in G \text{ and for all } \tau. \quad (3)$$

If for any choice of $g \in G$ we have that $\mathbf{K}_g = \mathbf{I}$, the identity function/matrix, then we say π (or ψ) is invariant to g .

2.4 Other-Play and Dec-POMDP Symmetry

The OP learning rule [22] of a joint policy, $\pi = (\pi^1, \pi^2)$, is formally the maximization of the following objective (which we state for the two-player case for ease of notation):

$$\pi^* = \arg \max_{\pi} \mathbb{E}_{\phi \sim \Phi} J(\pi^1, \phi(\pi^2)), \quad (4)$$

¹To disambiguate between actions of mathematical groups and actions of reinforcement learning agents, we always refer to the former as “group actions”.

²We use ψ and π for this reason more or less interchangeably.

where Φ is the class of equivalence mappings (symmetries) for the given Dec-POMDP, so that each $\phi \in \Phi$ is an automorphism of $\mathcal{S}, \mathcal{A}, \mathcal{O}$ whose application leaves the Dec-POMDP unchanged up to relabeling (ϕ is shorthand for $\phi = \{\phi_{\mathcal{S}}, \phi_{\mathcal{A}}, \phi_{\mathcal{O}}\}$). The expectation in Equation 4 is taken with respect to the uniform distribution on Φ .

Each $\phi \in \Phi$ acts on the action-observation history as

$$\phi(\tau_t^i) = (\phi(a_0^i), \phi(o_1^i), \dots, \phi(a_{t-1}^i), \phi(o_t^i)), \quad (5)$$

and acts on policies as

$$\hat{\pi} = \phi(\pi) \iff \hat{\pi}(\phi^{-1}(a) \mid \phi(\tau)) = \pi(a \mid \tau). \quad (6)$$

Policies $\pi, \hat{\pi}$ in Equation 6 are said to be symmetry-equivalent to one another with respect to ϕ . To correspond with Equation 3, we have $\mathbf{L}_\phi = \{\phi_{\mathcal{A}}, \phi_{\mathcal{O}}\}$ and $\mathbf{K}_\phi = \phi_{\mathcal{A}}$.

2.5 Idealized Example

Consider a game where you coordinate with an unknown stranger. Each of you must pick a lever from a set of 10 levers. Nine of the levers have 1 written on them, and one has 0.9 written. If you and the stranger pick the same lever then you are paid out the amount that is written on the lever, otherwise you are paid out nothing.

We have $G := S_9 \leq \Phi$, where S_9 is the symmetric group over 9 elements (all possible ways to permute 9 elements), since the nine levers with 1 written on them are symmetric. If you (using policy π^1) pick the 1-point lever l_1 and the stranger (using policy π^2) picks the 1-point lever l_2 , then π^1 and π^2 would be symmetry-equivalent with respect to any $\phi \in G$ that permutes the lever choices (e.g. $\phi(l_1) = l_2$, or equivalently $\mathbf{L}_\phi(l_1) = \mathbf{K}_\phi(l_1) = l_2$, as the levers are both observations and actions). If, as is likely, $l_1 \neq l_2$, you and the stranger leave empty-handed, illustrating symmetry breaking. If during training you and the stranger constrained yourselves to the space of equivariant policies whilst trying to maximize payoff, you would both converge to the policy that picks the 0.9-point lever (which is not only equivariant, but invariant to all $\phi \in G$), and so the choice of using equivariant policies allows symmetry to be maintained and payout to be guaranteed.

3 Method

We now present our methodology and architectural choices, as well as several theoretical results.

3.1 Architecture

We first consider the *symmetrizer* introduced in [35], which is a functional that transforms linear maps to equivariant linear maps. Inspired by this we present a generalisation of the symmetrizer that maps from Ψ to the equivariant subspace, Ψ_{equiv} , namely $S : \Psi \rightarrow \Psi_{\text{equiv}}$, which we define as

$$S(\psi) = \frac{1}{|G|} \sum_{g \in G} \mathbf{K}_g^{-1} \psi(\mathbf{L}_g), \quad (7)$$

so that the first layer of ψ , denoted f , being linear, is composed with each \mathbf{L}_g , and the ultimate layer, denoted h , is similarly composed with \mathbf{K}_g^{-1} (i.e. $\mathbf{K}_g^{-1} \psi(\mathbf{L}_g) = \mathbf{K}_g^{-1} \circ h(\hat{\psi}(f \circ \mathbf{L}_g))$), where $\psi = h(\hat{\psi}(f))$.

Below we establish properties of S that were analogously shown for the linear symmetrizer case [35]:

Proposition 1. (*Symmetric Property*) $S(\psi) \in \Psi_{\text{equiv}}$, for all $\psi \in \Psi$; that is, S maps neural networks to equivariant neural networks.

Proposition 2. (*Fixing Property*) $\Psi_{\text{equiv}} = \text{Ran}(S)$; that is, the range of S covers the entire equivariant subspace.

The proofs of these propositions can be found in the Appendix. The proof of Proposition 1 relies on properties of the group structure (such as closure of the group); it is thus critical for the symmetrizer to sum over permutations that together form a group structure, and not any random collection of permutations.

Propositions 1 and 2 show in combination that the architecture of $S(\psi)$ is indeed equivariant, and that optimizing over the class $\text{Ran}(S)$ (that is, keeping the permutation matrices \mathbf{L}_g and \mathbf{K}_g^{-1} for all g frozen while updating the other weights) is equivalent to optimizing over Ψ_{equiv} . The sum over each $g \in G$ of $S(\psi)$ is highly parallelizable [19], which allows for maintaining tractability. In this way we algorithmically enforce equivariance, mitigating the need to design equivariant layers by hand, as is typical [11, 12, 48, 49, 51, 52].

We note that for a network containing LSTM [20] layers, S will still map the network to an equivariant one no matter the values used for the hidden and cell states; indeed, Propositions 1, 2 and 4 all hold regardless of how the hidden and cell states are treated. However, how the hidden and cell states are treated matter for the overall performance of the network (equivariance is not a sufficient condition itself for success in ZSC, consider the uniformly random policy). We ergo propose the following: for each hidden state $h_{t,g}$ and cell state $c_{t,g}$ produced by the LSTM layer over input $(\mathbf{L}_g(\tau), h_{t-1}, c_{t-1})$, we compute $h_t = \frac{1}{|G|} \sum_{g \in G} h_{t,g}$ and $c_t = \frac{1}{|G|} \sum_{g \in G} c_{t,g}$, and proceed inductively. This is a non-obvious design choice, since each $h_{t,g}$ and $c_{t,g}$ is the output of a non-linear function; a more obvious choice would be to choose $h_t = h_{t,e}$ and $c_t = c_{t,e}$ where $e \in G$ is the identity element, but to ensure symmetry-aware play we experimented with this averaging scheme to “symmetrically adjust” the short and long-term memories. Since all the permutations over the input are computed in parallel with matrix multiplication, this design choice retains tractability. We test this choice empirically in Section 4. See the Appendix for further experimentation surrounding this design choice.

3.2 A Theory for Symmetry And Coordination

The following propositions provide guarantees for EQC’s ability to prevent all symmetry breaking:

Proposition 3. *If π^1 and π^2 are symmetry-equivalent policies with respect to $\phi \in G$ and π^1 is G -equivariant, then $\pi^1 = \pi^2$.*

Proof. We have $\pi^1 = \phi(\pi^2)$ for some $\phi \in G$, meaning $\pi^1(\phi^{-1}(a) | \phi(\tau)) = \pi^2(a | \tau)$. Identifying $\mathbf{L}_\phi = \{\phi_A, \phi_O\}$ and $\mathbf{K}_\phi = \phi_A$, by equivariance (Equation 3) we observe that $\pi^1(a | \tau) = \pi^1(\mathbf{K}_\phi \circ \mathbf{K}_\phi^{-1}(a) | \tau) = \pi^1(\mathbf{K}_\phi^{-1}(a) | \mathbf{L}_\phi(\tau)) = \pi^1(\phi^{-1}(a) | \phi(\tau))$. Thus $\pi^1(a | \tau) = \pi^1(\phi^{-1}(a) | \phi(\tau)) = \pi^2(a | \tau)$. \square

Proposition 4. *If π^1 and π^2 are symmetry-equivalent policies with respect to $\phi \in G$, then $S(\pi^1) = S(\pi^2)$.*

Proof. This is a corollary of the proof of Proposition 1, such that the symmetrizer architecture is fixed under any $\phi \in G$. \square

G -equivariance thus reduces cross-play to self-play between policies that are symmetry-equivalent with respect to any $\phi \in G$. This is critical, as this means equivariance merges policies of large classes together, such that symmetry-equivalent policies now indeed become equivalent in coordination. Put another way, the set of policies that are symmetry-equivalent with respect to any $\phi \in G$ form an equivalence class, and G -equivariance reduces the cardinality of each equivalence class to 1, thus allowing fluent coordination of policies within an equivalence class (as it merely becomes self-play).

3.3 Algorithmic Approach

Based on Propositions 1 and 2, a “naive approach” to obtaining a G -equivariant agent, for a given group G , might then be as follows: 1) initialize a random seed $\psi \in \Psi$; 2) map ψ to $S(\psi)$; 3) train $S(\psi)$ using any self-play algorithm, updating the weights while keeping the permutation matrices \mathbf{L}_g and \mathbf{K}_g^{-1} for all g frozen; 4) $S(\psi)$ is ready for inference. This naive approach is sound, but time and memory complexity inhibit scaling to large choices of G . For this reason we present the following two options (that together summarize EQC):

³Any group can be viewed as a group of permutations, or as a group of “permutation matrices”; this is known as the permutation representation of a group [9]. Also, the terminology “permutation matrix” alludes to the implementation of group actions as matrix products.

1. As a first option, we introduce “ G -OP”, where during training we sample random subsets of G for each mini-batch (i.e. we select a random $g \in G$ for each \mathbf{x} in the mini-batch, so that g group-theoretically acts on \mathbf{x} , and update the weights of ψ) and at test time, we use all the permutations of G by deploying $S(\psi)$.
2. As a second option, we train under any self-play algorithm and at test time deploy $S(\psi)$.

With Option 1 we train over the permutations of G at training time. With Option 2 we train under standard SP and use $S(\cdot)$ to obtain $S(\psi)$ at test time, which can be viewed as a coordination-improvement operator. We empirically validate both options in Section 4. From here on, we refer to $S(\psi)$ as “symmetrizing ψ ”.

We expect compatibility between the G -OP trained agents and the symmetrizer at test time, since at training time, these agents were trained over these permutations. For Option 2, where we have a network ϕ not trained with permutations, it is worthwhile to inspect $S(\phi)$; while we know $S(\phi)$ is equivariant, we know little else regarding the conventions and style of play it would use. While symmetry is a critical component of ZSC, other factors (such as choosing robust actions that are likely to succeed under minimal assumptions [2, 13]) are also significant. The following points justify why S maps any network to one that is not only equivariant, but overall stronger at ZSC: 1) the network averages over how ψ would act under each $g \in G$, thus ensuring play that accounts for multiple symmetries; 2) the group closure property fixes the architecture under any permutation in the group (see the proof of Proposition 1), forcing consistent play across different symmetries. We verify the empirical efficacy of Options 1 and 2 in Section 4.

What remains is the choice of the group G for our G -equivariant agents. One might consider the choice of $G = \Phi$, but this in general scales factorially with the size of the Dec-POMDP (e.g. the automorphism group over an environment with m symmetric features grows asymptotically with the number of permutations over m elements), and so quickly becomes infeasible. Hence, we treat G as a hyperparameter, and in Section 4 we analyze how different choices of G affect ZSC. In any case, picking G as a subgroup of Φ may be motivated in and of itself, when enumerating all possible symmetries of an environment is intractable, or if an environment is obfuscated such that some symmetries cannot be known (as is the case in many large, realistic domains). In this way, our method can relax the assumption of requiring access to all environment symmetries by flexibly choosing equivariance with respect to a subgroup.

4 Experiments

The primary test bed for our methodology is the AI benchmark task Hanabi [5]. Hanabi is a unique and challenging card game that requires agents to formulate informative implicit conventions in order to be successful. Hanabi has served as the primary test bed for many algorithms designed for ZSC [14, 22, 23, 28] due to its representation as a formidable Dec-POMDP task (see the Appendix for an introduction to Hanabi).

To make sense of the ensuing results, recall that a perfect score in Hanabi is 25, and that “bombing out” refers to expending all three fuse tokens and ending the game with a score of 0. Agents thus must be cautious not to expend all fuse tokens.

In Hanabi, the class of environment symmetries are the permutations of the card colors, because assuming there is no supplemental information, a permutation of the colors of the cards leaves the game unchanged. We therefore have $\Phi = S_5$, the symmetric group over 5 elements (all possible permutations over 5 elements). 5 here represents the number of different colors. $|S_5| = 5! = 120$.

In this section we analyze differences in play and the relationship to environmental symmetry between our equivariant agents and OP agents (4.1) and evaluate various pre-trained agents before and after symmetrization (4.2). These subsections evaluate Options 1 and 2, respectively (see Section 3).

For further details of the training setup and hyperparameters used in the following experiments, please refer to the Appendix.

Table 1: Comparing OP agents with our equivariant agents (each G -equivariant agent is trained using G -OP then symmetrized at test time) for their respective cross-play (XP) abilities. Each agent is trained with a different seed. Each pair of agents was evaluated over 5000 games, with the total averages compiled here. The error bars are the standard error of the mean.

XP Stats	OP	C_5 -Equivariant	D_{10} -Equivariant
Average Scores	15.32 ± 0.65	15.45 ± 0.47	16.08 ± 0.42
Average Bombout Rate	30.8%	24.3%	19.7%

4.1 Symmetry-Robust Agents

Here we compare the 12 OP agents from [22] with our equivariant agents. Specifically, we conduct the following ablation study: we fix two subgroups of S_5 : C_5 and D_{10} , which are the cyclic group of order 5 and the dihedral group of order 10, respectively; the geometric and/or algebraic interpretations of these groups is not important, for us groups purely serve as structured collections of permutations [9]. Next, we train 10 agents with the C_5 -OP learning rule (see Section 3), then map these 10 agents to equivariant ones using the C_5 -symmetrizer. Similarly, we train 10 agents with the D_{10} -OP learning rule, then map them with the D_{10} -symmetrizer. Finally, we compare the ZSC performance between these three pools of 10 agents. This is an ablation, because we control for the order of the group by considering both C_5 and D_{10} (in fact, C_5 is a subgroup of D_{10}) so to observe the effect of group order. Each agent uses Recurrent Replay Distributed Deep Q-Networks (R2D2) [25], and is trained using value decomposition networks (VDN) [43] and the simplified action decoding (SAD) algorithm [21]. The average self-play scores achieved by the OP agents is 23.93 ± 0.02 , of the C_5 -equivariant agents is 23.96 ± 0.03 , and the D_{10} -equivariant agents is 23.95 ± 0.03 .

Table 1 shows the findings of our ablation analysis. We see that both equivariant pools outperform the OP agents, in spite of the OP agents being trained over far more permutations than any of our equivariant agents: $|C_5| = 5$, $|D_{10}| = 10$, $|S_5| = 120$ (the OP agents are trained on permutations from S_5). This is consistent with earlier work where equivariant networks outperform data augmentation approaches [36, 47, 49, 50, 52, 55]. We conduct a Monte Carlo permutation test [15, 16] to compare the D_{10} -equivariant agent with the OP agents, for which we bound the derived p -value in a 99% binomial confidence interval.

We also observe that the D_{10} -equivariant agents outperform the C_5 -equivariant agents, which we largely attribute to $C_5 < D_{10}$ and Propositions 3 and 4: if we have a group $G < \Phi$, and choose $\phi \in \Phi$ uniformly at random, then $P(\phi \in G) = \frac{|G|}{|\Phi|}$, and so for two groups G, H with $|G| \leq |H|$, we have $P(\phi \in G) \leq P(\phi \in H)$.

4.2 Equivariance as a Policy-Improvement Operator

Here we consider the extent to which test-time symmetrizing improves ZSC performance across diverse policy types: we train 10 IQL agents [44], and use the 12 OP agents and 12 SAD agents from [22] and the 5 agents from [23] (level 5), and examine the effect of transforming each of these self-play trained agents with our equivariant architecture at test time. This selection of agent types is diverse in terms of both overall performance level as well as the style of play. Furthermore, we consider agents that are both trained for cross-play (OP and OBL) and trained for the self-play setting (SAD and IQL), which is an important dimension to account for in consideration of EQC's applicability. All agents use R2D2 [25], and the SAD and OP agents also use VDN [43]. The average self-play scores achieved by the SAD agents is 23.97 ± 0.04 , the IQL agents is 23.15 ± 0.02 , the OP agents is 23.93 ± 0.02 , and the OBL agents is 24.20 ± 0.01 .

Table 2 summarizes our findings over the policy types, where we consider symmetrization with both the cyclic group of order 5, C_5 , and the dihedral group of order 10, D_{10} . We find that mapping any of these networks through EQC with either choice of group improves its cross-play ability with novel agents at test time. We conduct Monte Carlo permutation tests [15, 16] to compare the D_{10} -symmetrized agents with their unsymmetrized counterparts for each policy type, for which we bound each derived p -value in a 99% binomial confidence interval. The p -values for the respective one-tailed tests for each of the SAD, IQL, OP and OBL agents are found to be within $(0, 0.00028]$, $(0, 0.00028]$, $(0, 0.00028]$, and $[0.0825, 0.0973]$ respectively, making the differences significant at

Table 2: Cross-play (XP) scores and bombout rates of various diverse policy types symmetrized at test time. Each agent is trained with a different seed. Each pair of agents was evaluated over 5000 games, with the total averages compiled here. The error bars are the standard error of the mean.

XP Stats	SAD	IQL	OP	OBL-L5
W/o symmetrizer	2.52 ± 0.34	10.53 ± 0.78	15.32 ± 0.65	23.77 ± 0.06
C_5 -symmetrized	3.61 ± 0.38	13.57 ± 0.66	16.07 ± 0.59	23.77 ± 0.05
D_{10} -symmetrized	3.61 ± 0.39	13.62 ± 0.65	16.48 ± 0.53	23.89 ± 0.04

Bombout Rate Stats	SAD	IQL	OP	OBL-L5
W/o symmetrizer	87.2%	39.4%	30.8%	0.56%
C_5 -symmetrized	81.9%	33.3%	26.1%	0.33%
D_{10} -symmetrized	82.0%	33.3%	24.6%	0.22%

the level $\alpha = 0.01$ for the SAD, IQL, and OP agents, and at the level $\alpha = 0.1$ for the OBL agents. Therefore, the improvement in coordination found by symmetrizing agents is highly statistically significant.

In particular, we see that OP agents improve in ZSC when symmetrized. In the proof for Proposition 2 we showed that equivariant networks are fixed under the symmetrizer. Since these OP agents were trained under the same permutations they are being symmetrized with, yet still improve upon symmetrization, this suggests that the data augmentation of OP does not lead to full equivariance, and so encoding environmental symmetry at the model level is necessary to realise crucial properties such as Propositions 3 and 4 as well as to improve overall performance.

It is especially significant that EQC maps networks not trained with permutations (SAD, IQL, OBL) to ones that are overall better at ZSC; we justify the mechanism by which it accomplishes this in Section 3. In particular, we demonstrate that EQC can improve on OBL, the state of the art for zero-shot coordination on the Hanabi benchmark at the time of writing.

Figure 2 shows the conditional action matrices (i.e. $P(a_t^i | a_{t-1}^j)$) of a symmetrized and unsymmetrized IQL agent (see the Appendix for the plots of other agents). This plot informs us on how differently the agent responds to possible actions of its partner, where large differences indicate specialized play, and therefore likely discrepancy between the conventions used by this agent and potential partners, which would be detrimental to cross-play. We see that the IQL agent has learned specialized conventions (for example, hinting the fourth color means for the partner to discard their first card), but that upon symmetrization its actions and responses become more consistent.

Another observation is that the D_{10} -symmetrizer does not always outperform the C_5 -symmetrizer: namely for the case of the SAD agents in terms of bombout rate. The combination of $C_5 < D_{10}$ and Propositions 3 and 4 implies that the D_{10} -symmetrizer corrects for symmetry breaking better than the C_5 does, so the worse bombout rate may only be attributed to the kind of play the symmetrized policy encourages, with a likely explanation as follows: if a policy is such that it acts very differently in each symmetry, then symmetrizing this policy over a large number of symmetry permutations will result in a policy with higher entropy (i.e. averaging over a large number of very different distributions results in one that is approximately uniform (see the Appendix)). But behaving “randomly” is not in general conducive to effective coordination, where one should try to use unambiguous conventions to facilitate play. *Thus, when symmetrizing an arbitrary network, we tacitly assume some degree of regularity of the network across symmetries* (we control this regularity when we opt for the proposed G -OP learning rule). Nonetheless, both choices of C_5 and D_{10} result in better ZSC than without any symmetrization for all the tested policy types (even with SAD, which learns notoriously specialized play [22, 23]), suggesting that policies in practice do generally satisfy these regularity requirements. This phenomenon, however, should be kept in mind and motivates future research into deriving the optimal choice of G for G -equivariant agents.

5 Related Work

Zero-Shot Coordination. Zero-shot coordination is the problem of constructing independent agents that can effectively engage with one another in cross-play at test time. Many directions to this

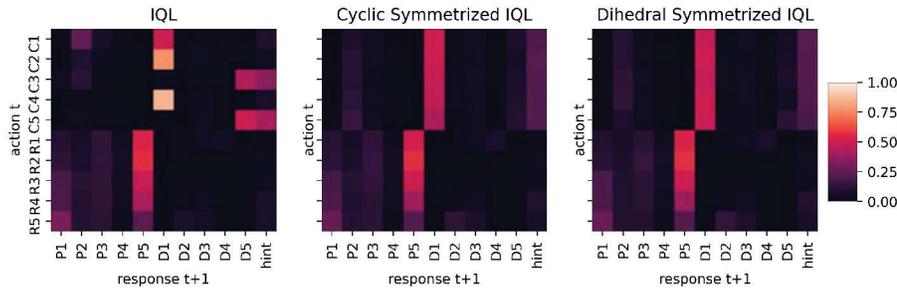


Figure 2: Conditional action matrices of IQL, i.e. $P(a_t^i | a_{t-1}^j)$, unsymmetrized (left) and symmetrized at test time (middle is C_5 -symmetrized and right is D_{10} -symmetrized). The y-axis represents the action taken at timestep t and the x-axis shows the proportion of each action as response at timestep $t + 1$. The matrices show the interactions between color/rank hinting and play/discarding. C1-5 and R1-5 mean hinting the 5 different colors and ranks respectively, and P1-5 and D1-5 mean playing and discarding the 1st-5th cards in the hand. We selected a random agent, and each plot is thereby computed by running 1000 episodes of self-play with the agent to compute the statistics.

problem have been considered: [14, 23] explore controlling the cognitive-reasoning depth so to avoid formation of arbitrary conventions that can complicate cross-play. However, [14] is computationally exacting, requiring days of GPU running time, and [23] requires simulator (environment) access. [38] offers a modular approach to separately learn rule-dependent and convention-dependent behaviors to facilitate coordination, but their method comes at the expense of scalability, with experiments limited to bandit problems and 1-color Hanabi. [34] uses repeated eigenvalues of the Hessian during training to find policies suited for ZSC, but is again inhibited by scalability, with experiments focused on lever games. [31] investigate using life-long learning for ZSC, but their approach requires access to a pool of pre-trained policies. [22], which we have referred to at length throughout this paper, takes to data augmentation to train symmetry-robust agents. EQC can be used along with any of these approaches (even with [22, 23], as we demonstrated in Section 4), so to correct for symmetry breaking and to encourage robust play.

Ad-Hoc Teamwork. Similar to zero-shot coordination is the ad-hoc teamwork setting [6, 30, 41], where an agent is pitted with an existing group of agents, and is tasked with learning to coordinate with this group during interaction. While ZSC and ad-hoc teamwork are closely related problems, they differ in that ad-hoc team play seeks policies that do well with *arbitrary agents* (often an undefined problem setting), while ZSC agents only need to perform well with agents that are optimized for ZSC using a common high-level approach. In principle EQC can also be used to improve performance in ad-hoc teamwork, which we leave for future work.

Equivariance. Equivariant networks are neural networks with symmetry constraints built in, such that for a transformation of the input, the output is appropriately transformed as well [3, 11, 46, 49, 52]. Equivariant policies have shown improved data efficiency in single agent reinforcement learning [1, 29, 33, 35, 40, 47, 54]. In the multi-agent case, prior works use permutation symmetries to ensure agent anonymity [8, 24, 27, 42, 43] or rotational equivariance over multiple agents [36]. While these approaches consider equivariance in multi-agent settings, they do not tackle the problem of ZSC. In terms of architecture, the symmetry ensemble of [39] is related to ours, but we propose adaptations that tailor recurrent layers to multi-agent coordination, thus suiting our design not just for Markov Decision Processes (MDPs) [7] but for more general POMDPs [4] and Dec-POMDPs.

6 Conclusion

In this work we proposed equivariant modelling to solve symmetry breaking. To this aim we presented EQC, which we validated via mathematical proof and experimentation over the complex Dec-POMDP task of Hanabi. We showed that EQC greatly improves on [22] by guaranteeing symmetry-equivariance of policies, and can be used as a policy-improvement operator to prevent symmetry breaking and promote robust play for a diversity of agents. In this way, we also showed the extent to which symmetrizing improves overall performance in coordination. Furthermore, we showed that EQC is “symmetry efficient”, such that it does not require access to all environment

symmetries in order to perform well. In addition, we empirically validated that agents perform well under our proposed architectural choices.

There are many directions for future work. One is to consider how to derive optimal choices of the group G for G -equivariant agents effective at ZSC, relative to a given policy type and task. Another direction is exploring how to efficiently uncover environment symmetries from a domain when they are not given nor assumed to be known. While our work focused on symmetry breaking, future work could explore other fundamental aspects of coordination.

Acknowledgements

The experiments were made possible by a generous equipment grant from NVIDIA. Christian Schroeder de Witt is generously funded by the Cooperative AI Foundation. We thank Jacob Beck for his constructive criticism of the manuscript.

References

- [1] Farzad Abdolhosseini et al. “On Learning Symmetric Locomotion”. In: *ACM SIGGRAPH Motion, Interaction, and Games*. 2019.
- [2] Marina Agronov et al. “Beliefs and endogenous cognitive levels: An experimental study”. In: *Games and Economic Behavior* 75.2 (2012), pp. 449–463.
- [3] Brandon Anderson, Truong Son Hy, and Risi Kondor. “Cormorant: Covariant molecular neural networks”. In: *Advances in neural information processing systems* 32 (2019).
- [4] Karl J Astrom. “Optimal control of Markov decision processes with incomplete state estimation”. In: *J. Math. Anal. Applic.* 10 (1965), pp. 174–205.
- [5] Nolan Bard et al. “The hanabi challenge: A new frontier for ai research”. In: *Artificial Intelligence* 280 (2020), p. 103216.
- [6] Samuel Barrett, Peter Stone, and Sarit Kraus. “Empirical evaluation of ad hoc teamwork in the pursuit domain”. In: *The 10th International Conference on Autonomous Agents and Multiagent Systems-Volume 2*. 2011, pp. 567–574.
- [7] Richard Bellman. “A Markovian decision process”. In: *Journal of mathematics and mechanics* (1957), pp. 679–684.
- [8] Wendelin Böhmer, Vitaly Kurin, and Shimon Whiteson. “Deep coordination graphs”. In: *International Conference on Machine Learning*. PMLR. 2020.
- [9] Arthur Cayley. “VII. On the theory of groups, as depending on the symbolic equation $\theta^n = 1$ ”. In: *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science* 7.42 (1854), pp. 40–47.
- [10] Gabriele Cesa, Leon Lang, and Maurice Weiler. “A Program to Build E (N)-Equivariant Steerable CNNs”. In: *International Conference on Learning Representations*. 2021.
- [11] Taco Cohen and Max Welling. “Group equivariant convolutional networks”. In: *International conference on machine learning*. PMLR. 2016, pp. 2990–2999.
- [12] Taco S Cohen and Max Welling. “Steerable cnns”. In: *arXiv preprint arXiv:1612.08498* (2016).
- [13] Miguel A Costa-Gomes and Vincent P Crawford. “Cognition and behavior in two-person guessing games: An experimental study”. In: *American economic review* 96.5 (2006), pp. 1737–1768.
- [14] Brandon Cui et al. “K-level Reasoning for Zero-Shot Coordination in Hanabi”. In: *Advances in Neural Information Processing Systems* 34 (2021).
- [15] Meyer Dwass. “Modified randomization tests for nonparametric hypotheses”. In: *The Annals of Mathematical Statistics* (1957), pp. 181–187.
- [16] T Eden and F Yates. “On the validity of Fisher’s z test when applied to an actual example of non-normal data.(With five text-figures.)” In: *The Journal of Agricultural Science* 23.1 (1933), pp. 6–17.
- [17] Marc Finzi, Max Welling, and Andrew Gordon Wilson. “A practical method for constructing equivariant multilayer perceptrons for arbitrary matrix groups”. In: *International Conference on Machine Learning*. PMLR. 2021, pp. 3318–3328.
- [18] Jonathan Gordon et al. “Permutation equivariant models for compositional generalization in language”. In: *International Conference on Learning Representations*. 2019.
- [19] Maurice Herlihy and Nir Shavit. *The Art of Multiprocessor Programming, revised first edition*. Morgan Kaufmann, 2012.
- [20] Sepp Hochreiter and Jürgen Schmidhuber. “Long short-term memory”. In: *Neural computation* 9.8 (1997), pp. 1735–1780.
- [21] Hengyuan Hu and Jakob N Foerster. “Simplified action decoder for deep multi-agent reinforcement learning”. In: *arXiv preprint arXiv:1912.02288* (2019).
- [22] Hengyuan Hu et al. ““Other-Play” for Zero-Shot Coordination”. In: *International Conference on Machine Learning*. PMLR. 2020, pp. 4399–4410.
- [23] Hengyuan Hu et al. “Off-belief learning”. In: *International Conference on Machine Learning*. PMLR. 2021, pp. 4369–4379.
- [24] Jiechuan Jiang et al. “Graph convolutional reinforcement learning”. In: *International Conference on Learning Representations*. 2020.
- [25] Steven Kapturowski et al. “Recurrent experience replay in distributed reinforcement learning”. In: *International conference on learning representations*. 2018.

- [26] Diederik P Kingma and Jimmy Ba. “Adam: A method for stochastic optimization”. In: *arXiv preprint arXiv:1412.6980* (2014).
- [27] Iou-Jen Liu, Raymond A. Yeh, and Alexander G. Schwing. “PIC: Permutation invariant critic for multi-agent deep reinforcement learning”. In: *Conference on Robot Learning*. 2019.
- [28] Andrei Lupu et al. “Trajectory diversity for zero-shot coordination”. In: *International Conference on Machine Learning*. PMLR. 2021, pp. 7204–7213.
- [29] Arnab Kumar Mondal, Pratheeksha Nair, and Kaleem Siddiqi. “Group equivariant deep reinforcement learning”. In: *arXiv preprint arXiv:2007.03437* (2020).
- [30] Darius Muglich et al. “Generalized Beliefs for Cooperative AI”. In: *International Conference on Machine Learning*. PMLR. 2022, pp. 16062–16082.
- [31] Hadi Nekoei et al. “Continuous coordination as a realistic scenario for lifelong learning”. In: *International Conference on Machine Learning*. PMLR. 2021, pp. 8016–8024.
- [32] Frans A Oliehoek. “Decentralized pomdps”. In: *Reinforcement Learning*. Springer, 2012, pp. 471–503.
- [33] Jung Yeon Park et al. “Learning Symmetric Representations for Equivariant World Models”. In: (2021).
- [34] Jack Parker-Holder et al. “Ridge rider: Finding diverse solutions by following eigenvectors of the hessian”. In: *Advances in Neural Information Processing Systems* 33 (2020), pp. 753–765.
- [35] Elise van der Pol et al. “MDP homomorphic networks: Group symmetries in reinforcement learning”. In: *Advances in Neural Information Processing Systems* 33 (2020), pp. 4199–4210.
- [36] Elise van der Pol et al. “Multi-Agent MDP Homomorphic Networks”. In: *International Conference on Learning Representations* (2022).
- [37] Arthur L Samuel. “Some studies in machine learning using the game of checkers”. In: *IBM Journal of research and development* 3.3 (1959), pp. 210–229.
- [38] Andy Shih et al. “On the critical role of conventions in adaptive human-AI collaboration”. In: *arXiv preprint arXiv:2104.02871* (2021).
- [39] David Silver et al. “Mastering the game of Go with deep neural networks and tree search”. In: *nature* 529.7587 (2016), pp. 484–489.
- [40] Gregor NC Simm et al. “Symmetry-aware actor-critic for 3d molecular design”. In: *International Conference on Learning Representations* (2021).
- [41] Peter Stone et al. “Ad hoc autonomous agent teams: Collaboration without pre-coordination”. In: *Twenty-Fourth AAAI Conference on Artificial Intelligence*. 2010.
- [42] Sainbayar Sukhbaatar, Arthur Szlam, and Rob Fergus. “Learning multiagent communication with backpropagation”. In: *Advances in Neural Information Processing Systems*. 2016.
- [43] Peter Sunehag et al. “Value-decomposition networks for cooperative multi-agent learning”. In: *arXiv preprint arXiv:1706.05296* (2017).
- [44] Ming Tan. “Multi-agent reinforcement learning: Independent vs. cooperative agents”. In: *Proceedings of the tenth international conference on machine learning*. 1993, pp. 330–337.
- [45] Gerald Tesauro. “TD-Gammon, a self-teaching backgammon program, achieves master-level play”. In: *Neural computation* 6.2 (1994), pp. 215–219.
- [46] Nathaniel Thomas et al. “Tensor field networks: Rotation-and translation-equivariant neural networks for 3d point clouds”. In: *arXiv preprint arXiv:1802.08219* (2018).
- [47] Dian Wang, Robin Walters, and Robert Platt. “SO (2)-equivariant reinforcement learning”. In: *The Tenth International Conference on Learning Representations*. 2022.
- [48] Maurice Weiler and Gabriele Cesa. “General e (2)-equivariant steerable cnns”. In: *Advances in Neural Information Processing Systems* 32 (2019).
- [49] Maurice Weiler et al. “3d steerable cnns: Learning rotationally equivariant features in volumetric data”. In: *Advances in Neural Information Processing Systems* 31 (2018).
- [50] Marysia Winkels and Taco S Cohen. “3D G-CNNs for pulmonary nodule detection”. In: *arXiv preprint arXiv:1804.04656* (2018).
- [51] Daniel Worrall and Max Welling. “Deep scale-spaces: Equivariance over scale”. In: *Advances in Neural Information Processing Systems* 32 (2019).
- [52] Daniel E Worrall et al. “Harmonic networks: Deep translation and rotation equivariance”. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 2017, pp. 5028–5037.

- [53] Chao Yu et al. “The Surprising Effectiveness of PPO in Cooperative, Multi-Agent Games”. In: *arXiv preprint arXiv:2103.01955* (2021).
- [54] Xupeng Zhu et al. “Equivariant Grasp learning In Real Time”. In: (2021).
- [55] Xupeng Zhu et al. “Sample Efficient Grasp Learning Using Equivariant Models”. In: *arXiv preprint arXiv:2202.09468* (2022).

1. For all authors...
 - (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes] See Sections 3 and 4
 - (b) Did you describe the limitations of your work? [Yes] See Sections 3, 4 and 6
 - (c) Did you discuss any potential negative societal impacts of your work? [Yes] See the Appendix
 - (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
 - (a) Did you state the full set of assumptions of all theoretical results? [Yes]
 - (b) Did you include complete proofs of all theoretical results? [Yes] In Section 3 and the Appendix
3. If you ran experiments...
 - (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [Yes] Section 4, the Appendix, and the supplementary material
 - (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes] See Section 4 and the Appendix
 - (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [Yes] See Section 4
 - (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [Yes] See the Appendix
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
 - (a) If your work uses existing assets, did you cite the creators? [Yes]
 - (b) Did you mention the license of the assets? [N/A]
 - (c) Did you include any new assets either in the supplemental material or as a URL? [Yes] In the supplement material
 - (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [N/A]
 - (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A] We use rollouts of Hanabi games played by AI agents.
5. If you used crowdsourcing or conducted research with human subjects...
 - (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]
 - (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
 - (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]