

# Theory of Mind Guided Strategy Adaptation for Zero-Shot Coordination

Andrew Ni  
Carnegie Mellon University  
Pittsburgh, United States  
acni@andrew.cmu.edu

Simon Stepputtis  
Virginia Tech  
Blacksburg, United States  
stepputtis@vt.edu

Stefanos Nikolaidis  
University of Southern California  
Los Angeles, United States  
nikolaid@usc.edu

Michael Lewis  
University of Pittsburgh  
Pittsburgh, United States  
mlewis@sis.pitt.edu

Katia P. Sycara  
Carnegie Mellon University  
Pittsburgh, United States  
sycara@andrew.cmu.edu

Woojun Kim  
Carnegie Mellon University  
Pittsburgh, United States  
woojunk@andrew.cmu.edu

## ABSTRACT

A central challenge in multi-agent reinforcement learning is enabling agents to adapt to previously unseen teammates in a zero-shot fashion. Prior work in zero-shot coordination often follows a two-stage process, first generating a diverse training pool of partner agents, and then training a best-response agent to collaborate effectively with the entire training pool. While many previous works have achieved strong performance by devising better ways to diversify the partner agent pool, there has been less emphasis on how to leverage this pool to build an adaptive agent. One limitation is that the best-response agent may converge to a *static, generalist* policy that performs reasonably well across diverse teammates, rather than learning a more *adaptive, specialist* policy that can better adapt to teammates and achieve higher synergy. To address this, we propose an adaptive ensemble agent that uses Theory-of-Mind-based best-response selection to first infer its teammate’s intentions and then select the most suitable policy from a policy ensemble. We conduct experiments in the Overcooked environment to evaluate zero-shot coordination performance under both fully and partially observable settings. The empirical results demonstrate the superiority of our method over a single best-response baseline.

## KEYWORDS

Zero-shot Coordination; Theory of Mind; Multi-Agent RL

### ACM Reference Format:

Andrew Ni, Simon Stepputtis, Stefanos Nikolaidis, Michael Lewis, Katia P. Sycara, and Woojun Kim. 2026. Theory of Mind Guided Strategy Adaptation for Zero-Shot Coordination. In *Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026)*, Paphos, Cyprus, May 25 – 29, 2026, IFAAMAS, 10 pages. <https://doi.org/10.65109/UMCV8852>


## 1 INTRODUCTION

Effective coordination with previously unseen partners is vital for multi-agent systems deployed in real-world settings. One promising approach is zero-shot coordination (ZSC) [13, 31, 40], in which agents must align with novel collaborators without any additional

learning or explicit communication at test time. Achieving ZSC, however, is challenging because agents trained in self-play, where they are jointly trained, tend to overfit to shared conventions that emerge during learning. Consequently, they struggle to infer a new partner’s intent and often lack a principled strategy for adapting their behavior. This challenge is further exacerbated in deep learning-based approaches, as deep neural networks often behave unpredictably when exposed to unseen inputs [4, 5, 7, 10], ultimately resulting in coordination failures.

A common approach for ZSC is the population-based approach [3, 28, 37, 46], which first generates a diverse partner population, e.g., by using different seeds or by maximizing entropy across populations, and then trains a best-response (BR) agent<sup>1</sup> to maximize the expected return across the partner pool. During training, the BR agent interacts with partners sampled from the population, observes their behaviors, and learns to select actions that maximize the expected team reward averaged across all possible partners. In other words, the BR agent implicitly infers each partner’s policy through experience and adapts its own actions to achieve the highest overall return when paired with diverse partners. Despite reasonable performance in many tasks, this approach still has a limitation: the BR agent is trained to fit all partners with a single policy and may converge to a more *static, generalist* policy that does not fully adapt to the new partners. Intuitively, the static generalist policy represents an easy-to-learn local optimum that hinders the BR agent from learning a truly adaptive policy. Effective collaboration, however, inherently requires adapting to the behaviors of one’s partners, rather than relying on fixed conventions or averaged responses. The ability to adjust one’s policy based on the partner’s behavior is essential for achieving robust coordination with diverse or unseen partners. Our goal is to develop a method that enables an agent to accurately infer a partner’s intent and adapt its actions in a zero-shot manner.

In this paper, we propose the Theory of Mind-based Best Response Selection (TBS) framework to address the adaptivity gap in population-based BR methods. TBS leverages the concept of Theory of Mind (ToM)—the ability to infer others’ beliefs and intentions—to enable agents to reason about their partners’ behavior and adapt their own actions accordingly. This capability enhances adaptivity and leads to more effective coordination with novel or dynamically

 This work is licensed under a Creative Commons Attribution International 4.0 License.

*Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026)*, C. Amato, L. Dennis, V. Mascardi, J. Thangarajah (eds.), May 25 – 29, 2026, Paphos, Cyprus. © 2026 International Foundation for Autonomous Agents and Multiagent Systems ([www.ifaamas.org](http://www.ifaamas.org)). <https://doi.org/10.65109/UMCV8852>

<sup>1</sup>In game theory, the term best-response refers to the strategy which produces the most favorable outcome, when other players’ strategies are fixed [6, 9]

changing partners. Building on this insight, TBS constructs a diverse partner pool, clusters partners into behaviorally distinct groups through self-tuning spectral clustering, and trains a specialized BR policy for each cluster. At execution time, a ToM-based adaptive selection module infers a partner’s intent and identifies the most similar cluster, selecting the corresponding BR policy for real-time coordination. By integrating ToM reasoning with population-based training, TBS bridges the adaptivity gap and achieves robust zero-shot coordination with previously unseen partners.

We evaluate TBS in terms of ZSC performance in the Overcooked environment [3], a fully cooperative two-player cooking game. In addition, we conduct ablation studies assessing the performance of our method with different training population sizes and with different hyperparameter choices such as different number of clusters. Empirical results show that the proposed method outperforms the population-based Best-Response method [37, 46], with a performance gap that increases with the size of the training pool. Overall, our results demonstrate the utility of explicit ToM based adaptation for zero-shot coordination. Our contributions include:

- Automatic identification of strategy clusters within a diverse partner population using self-tuning spectral clustering of cross-play performance.
- A Theory-of-Mind-based adaptive selection module that infers a partner’s intent and dynamically selects the most suitable best-response policy from an ensemble.
- Demonstration of TBS’s adaptivity and generalization, showing that combining strategy-specific BR agents with ToM-guided policy selection yields robust ZSC in the Overcooked environment.

## 2 BACKGROUND AND RELATED WORKS

### 2.1 Multi-Agent Reinforcement Learning

A decentralized partially observable Markov decision process (Dec-POMDP) [2, 33] formalizes fully cooperative multi-agent decision-making under partial observability. At each timestep  $t$ , each agent  $i$  receives a local observation  $o_t^i$  of the global state  $s_t$  and selects an action according to its policy  $a_t^i \sim \pi^i(\cdot | o_t^i)$ . The environment then transitions to a new state  $s_{t+1}$  and all agents receive a shared reward  $r_t$  that depends on the joint action. The objective is to maximize the expected discounted return under the joint policy  $(\pi^1, \dots, \pi^N)$ :

$$J(\pi^1, \dots, \pi^N) = \mathbb{E}_{\tau \sim (\pi^1, \dots, \pi^N)} \left[ \sum_{t=0}^{\infty} \gamma^t r_t \right], \quad (1)$$

where  $\gamma$  and  $\tau$  denote the discount factor and a trajectory sampled from the trajectory distribution induced by the joint policy  $(\pi^1, \dots, \pi^N)$  interacting with the environment dynamics, respectively. In standard cooperative MARL, these policies are typically trained jointly in *self-play*, where all agents learn simultaneously within the same environment instance.

A common framework for Dec-POMDPs is centralized training with decentralized execution (CTDE), where agents are trained with access to global information but execute policies using only local observations [16, 22]. Under CTDE, a key research direction in cooperative MARL focuses on enhancing coordination among agents, such as maximizing mutual information [20, 27] learning to communicate [17, 19], and value decomposition methods [36, 38]. VDN [38] is a value decomposition method for multi-agent credit

assignment that approximates the joint Q-function as a sum of individual value functions, and trains it by minimizing the temporal-difference error.

### 2.2 Zero-Shot Coordination

Despite the success of various approaches in cooperative multi-agent environments, the learned policies may fail to cooperate even with other agents trained using the same algorithm but different random seeds. This limitation motivates the problem of zero-shot coordination [3, 14], where agents should be able to cooperate with novel partners without additional training. This stems from the fact that, as described in Eq. 1, standard cooperative MARL optimizes the expected return  $J(\pi^1, \dots, \pi^N)$  through self-play, where all agents are trained jointly and their policies co-adapt to one another. Consequently, the trajectory distribution  $\tau \sim (\pi^1, \dots, \pi^N)$  used during training is determined by these co-adapted policies. When an agent is later paired with an independently trained partner  $\hat{\pi}^j$ , the induced trajectory distribution  $\tau' \sim (\pi^1, \dots, \hat{\pi}^j, \dots, \pi^N)$  differs from the one optimized during training. As a result, the optimization objective no longer matches the evaluation setting, leading to a shift in the trajectory distribution and potential coordination failure.

For evaluation, ZSC performance is commonly measured using the cross-play (XP) and inter-algorithm cross-play (inter-XP) objectives. In the two-agent case, the XP objective assesses how well agents trained independently under the same learning algorithm  $\mathcal{A}$  can coordinate:

$$J_{\text{XP}}(\mathcal{A}) = \mathbb{E}_{(\pi_A^1, \pi_A^2), (\pi_B^1, \pi_B^2) \stackrel{\text{iid}}{\sim} \mathcal{A}} \left[ \frac{J(\pi_A^1, \pi_B^2) + J(\pi_B^1, \pi_A^2)}{2} \right]. \quad (2)$$

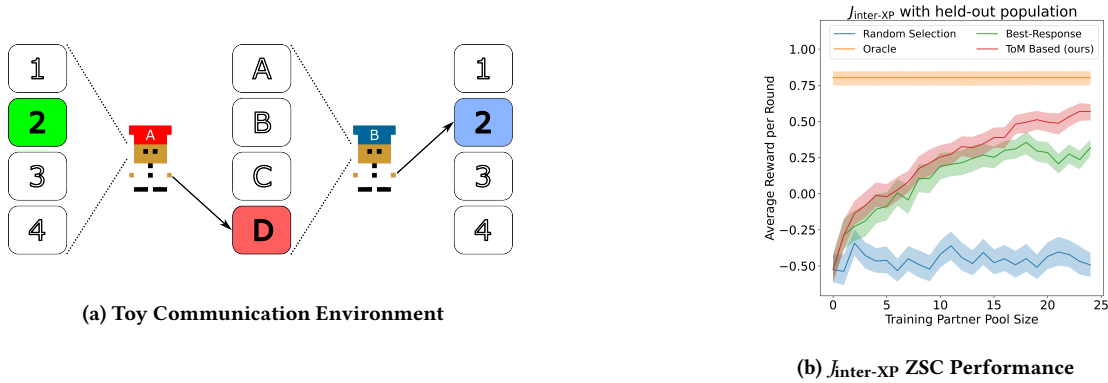
Here,  $(\pi_A^1, \pi_A^2)$  and  $(\pi_B^1, \pi_B^2)$  denote two independent policy pairs obtained from separate training runs of the same algorithm  $\mathcal{A}$  with different random seeds. The expectation thus measures how well independently trained agents from  $\mathcal{A}$  can coordinate when paired across runs.

Prior works such as Other-Play [14] and Off-Belief Learning [12] optimize their methods to improve performance under this metric. More recently, Any-Play [30] extends the evaluation to partners trained with different algorithms, defining the inter-XP objective as  $J_{\text{inter-XP}}(\mathcal{A}) =$

$$\mathbb{E}_{(\pi_A^1, \pi_A^2) \sim \mathcal{A}, (\pi_B^1, \pi_B^2) \sim \mathcal{B}} \left[ \frac{J(\pi_A^1, \pi_B^2) + J(\pi_B^1, \pi_A^2)}{2} \right] \quad (3)$$

where  $\mathcal{A}$  and  $\mathcal{B}$  denote two distinct multi-agent learning algorithms, each providing independently trained policy pairs. To improve ZSC performance, a variety of approaches have been proposed, including convention-avoidance methods, which aim to reduce agents’ reliance on arbitrary conventions formed during self-play, and population-based methods [3, 28, 37, 46], which train cooperative agents against a diverse set of partners to promote generalization and adaptability. We next discuss population-based methods in more detail.

**Population-based ZSC** has been widely investigated as a framework for enabling adaptive agents to achieve ZSC with novel partners at test time [15, 24]. The PBT approach typically involves (1) creating a *diverse population* of simulated agents, and (2) training an adaptive cooperator agent as a *best response* to this population



**Figure 1: Left: Toy communication environment. Alice (in red) is shown a random number and can take actions A through D or bail out. Bob (in blue) sees Alice’s action and can guess the random number or bail out. After the round is over, the random number is revealed. Right:  $J_{\text{inter-XP}}$  of different PBT ZSC algorithms as a function of partner agent pool size in a toy environment. Shaded regions represent bootstrapped 95% confidence intervals capturing variability over evaluation seeds.**

to mitigate overfitting to specific strategies and promote generalization across diverse partners. Methods for constructing a diverse partner population [25, 32, 37, 46] include Fictitious Co-Play [37], which trains multiple agent pairs with different seeds and takes checkpoints from the start, middle, and end of training to represent partners with diverse skill levels, Maximum Entropy Population-based (MEP) training [46], which maximizes the entropy of the average policy of the population, random reward shaping [26, 42, 43], which models implicit biases by adding a random reward shaping to each agent during self-play training, and other diversity-promoting techniques such as training a generative model to simulate infinitely many partner agents [28]. In this work, we use random reward shaping to generate a diverse population of agents.

Once the population is constructed, most population-based approaches train a single cooperator, i.e., BR agent, to learn a common best response to the entire population of simulated partners [3, 28, 37, 46]. However, this design may limit the BR agent’s ability to adapt to novel partners, as it tends to learn a static generalist policy rather than an adaptive one [29, 45]. To improve adaptability, Lou et al. [29] and He et al. [11] proposed conditioning the BR policy on a representation of the partner derived from prior interactions with that partner. Specifically, PECAN [29] learns a few-shot, level-based best response conditioned on the approximate skill level of a given partner, which is estimated by a context encoder using trajectory data from previous episodes. In contrast, PALM [11] introduces a framework that adapts to unknown partners by conditioning the BR policy on a latent embedding of a partner’s policy, which is fine-tuned at test time using data from prior interactions. Similarly, TSF [26] and MESH [45] construct a library of diverse strategies and adapt to unknown partners within a single episode by evaluating how well each strategy explains the partner’s recent actions. TSF builds its strategy library through a mix of heuristic agents and RL with reward shaping, selecting an appropriate strategy to execute depending on partner strategy, whereas MESH uses apprenticeship learning from clustered human-human trajectories and performs a weighted combination of strategies during execution.

### 2.3 Theory of Mind

Theory of Mind [34] refers to the ability of an agent to reason about the mental states of itself and others, including their intentions, beliefs, and preferences. Recent work has shown that such capabilities can be approximated by deep neural networks. For example, Theory of Mind Networks (ToMNETs) [35] model other agents solely from observed behavior, enabling a learned prior over plausible behavioral patterns. These networks can improve their predictive accuracy after observing a few timesteps of another agent’s behavior and infer false beliefs in partially observable scenarios.

Forming such internal models of other agents, considered as agent modeling, is crucial for adaptive multi-agent interaction [1, 21]. In this work, we focus on a form of ToM that infers an agent’s intentions, represented as a set of interpretable concepts  $c[i]_{i=1}^M$ , where each  $c[i]$  corresponds to one of  $M$  predefined, environment-specific intentions [8, 18, 23]. For instance, in the Overcooked environment, a concept may represent the “intention to pick up an onion from the onion pile.” The ToM network outputs the probability of each concept being active given an observation, and is trained using binary cross-entropy loss against the corresponding ground-truth concept labels obtained from the environment. The loss function for the ToM parameters  $\theta_{\text{ToM}}$  is given by

$$\mathcal{L}(\theta_{\text{ToM}}) = \mathbb{E}_{(o_t, c_t)} \left[ c_t \log \hat{c}_t + (1 - c_t) \log(1 - \hat{c}_t) \right] \quad (4)$$

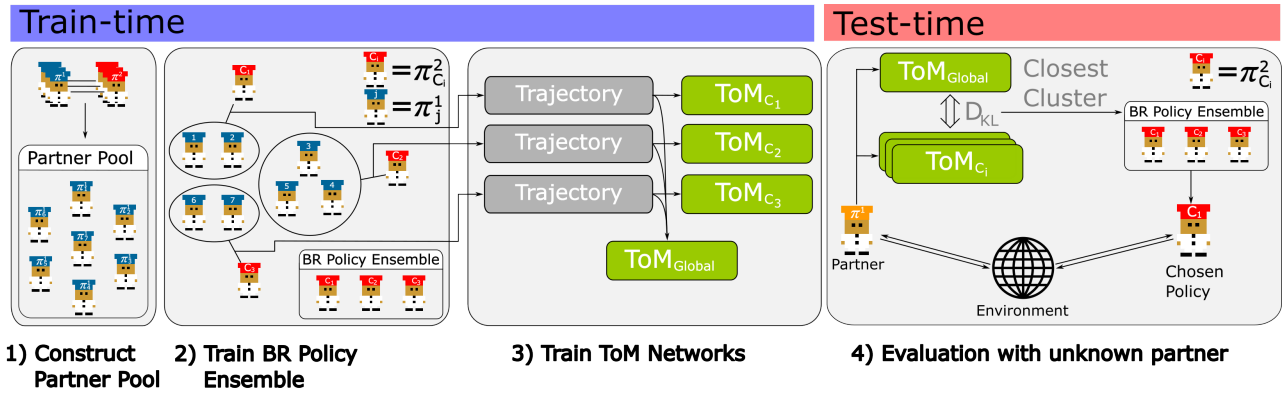
where  $\hat{c}_t = \text{ToM}(o_t; \theta_{\text{ToM}})$ .

This objective encourages the ToMNet to accurately predict the partner’s high-level intentions from observations

## 3 METHODOLOGY

### 3.1 Motivation

One limitation of training a BR against an entire population of partners is that it often converges to a static, generalist policy rather than an adaptive one [29, 45]. Because adaptive behavior requires learning multiple distinct strategies and maintaining long-term memory for partner modeling, such generalist policies can



**Figure 2: Overview of TBS: During training, (1) a diverse partner pool is constructed, (2) partners are clustered by behavioral similarity and specialist BR policies are trained per cluster, and (3) global and cluster-specific ToM networks are learned. During testing, the cooperators infer an unseen partner’s cluster and selects the corresponding BR policy for adaptive coordination.**

easily become local optima during BR training. Moreover, as the diversity and size of the partner pool increase, learning an adaptive BR becomes increasingly difficult, diminishing the potential ZSC benefits of population diversity.

To illustrate these limitations, we consider a simple two-agent communication environment (Fig. 1 (a)) with 16 rounds. In each round, the red agent (Alice) is shown a random number from 1 to 4 and can either take one of four actions (A-D) or bail out for a reward of 0. The blue agent (Bob) observes Alice’s action and must either guess Alice’s number, receiving a reward of 1 if correct and -1 if incorrect, or bail out for a reward of 0; after each round, the true number is revealed to Bob. This task can be viewed as a signaling game: Alice’s action (A-D) serves as a discrete message intended to convey the observed number (1-4). Each self-played pair develops its own private convention mapping numbers to actions (e.g., Alice 1→A, 2→B, 3→C, 4→D with Bob decoding A→1, B→2, C→3, D→4), while a different pair may learn a different convention (e.g., 1→A, 2→C, 3→B, 4→D). Consequently, ZSC across pairs is non-trivial: Bob must infer the partner’s convention online after observing Alice’s signal.

In the toy environment, we consider the task of adapting as Bob to an unseen Alice agent. We use MEP [46] to generate a training pool of Alice agents and a separate held-out pool for evaluation. We compare four methods: (i) *Random Selection*, which selects a random strategy from the pool; (ii) *BR agent*, which trains a single agent to cooperate with any agent in the Alice pool; (iii) *ToM-based Best-Response Selection (TBS, ours)*, which infers Alice’s intentions via ToM and selects the specialized Bob agent that best matches them; and (iv) *Oracle*, which uses the exact partner co-trained with the teammate, providing an upper bound on ZSC performance. Fig. 1 (b) shows the inter-XP performance  $J_{\text{inter-XP}}$  of the four methods as a function of the number of training agents. When the pool is small, our method performs comparably to the BR baseline. However, as the pool grows, the BR agent struggles to adapt to the increasing diversity of strategies and instead converges to a static generalist policy. In contrast, our approach, through ToM-based reasoning, retains strong adaptation capability and outperforms both the BR and random baselines.

### 3.2 ToM-based Best Response Selection

To overcome the limitations of static generalist BR policy, we propose an adaptive **Theory of Mind-based Best Response Selection (TBS)** framework. TBS trains an ensemble of specialist BR policies and adaptively selects one to act based on the inferred intent of a new teammate through ToM reasoning. To this end, TBS includes a cooperators agent that collaborates with partners at test time. It comprises three main components: (1) a diverse partner pool, (2) an ensemble of specialist BR policies, and (3) an adaptive selection module that dynamically selects the most suitable policy from the ensemble according to the observed behavior of the teammate. Fig. 2 provides an overview of the proposed TBS framework.

**(1) Partner Pool Construction:** We begin by constructing a diverse pool of simulated partner agents using population-based training methods such as Fictitious Co-Play (FCP) [37], Maximum Entropy Population-based training [46], and random reward shaping [39, 42, 43], following prior BR approaches. Rather than training a single generalist BR policy against this population, we train an ensemble of BR policies, each specialized to cooperate with a distinct strategy cluster within the partner pool.

**(2) BR Policy Ensemble:** Having constructed a diverse partner pool, we now train an ensemble of BR policies, each specialized to collaborate with behaviorally similar groups of partners. Rather than training a separate BR policy for every individual partner, we group partners into clusters to improve computational efficiency and promote generalization across similar strategies. For this, we first cluster the partner pool into distinct strategy clusters using *self-tuning spectral clustering* [44], which automatically determines the number of clusters based on the policy similarity matrix. Specifically, we construct the similarity matrix using cross-play performance, where each entry is defined as

$$s(A, B) = \frac{J(\pi_A^1, \pi_B^2) + J(\pi_B^1, \pi_A^2)}{J(\pi_A^1, \pi_A^2) + J(\pi_B^1, \pi_B^2)}. \tag{5}$$

Intuitively, the better two agent pairs  $A$  and  $B$  perform in cross-play, the more similar they are, and thus the closer they are in the similarity space. Applying self-tuning spectral clustering (STSC) [44] to the cross-play similarity matrix yields  $k$  clusters, denoted by  $C_1, \dots, C_k$ ,




Figure 3: Onion soup environment layouts. From left to right: *Large Room*, *Counter Circuit*, *Bothway Coordination*, *Asymmetric Advantages*, *Forced Coordination*, *Coordination Ring*, *Cramped Room*.

where each cluster contains a subset of agent pairs with similar behavior. Specifically, given an  $n \times n$  similarity matrix  $S$ , where  $s(A, B)$  denotes each entry, STSC [44] performs eigen-decomposition of the graph Laplacian of  $S$  and collects the top  $k$  eigenvectors  $x_1, \dots, x_k$  to form the matrix  $X = [x_1, \dots, x_k] \in \mathbb{R}^{n \times k}$ . In the ideal case where the Laplacian is strictly block-diagonal, each row of  $X$  has only one nonzero entry (up to some rotation  $XR$ , with  $R \in \mathbb{R}^{k \times k}$ ). STSC defines the alignment cost for a given number of clusters  $k$  and rotation matrix  $R$  as

$$J(R, k) = \sum_{i=1}^n \sum_{j=1}^k \frac{(XR)_{ij}^2}{[\max_j (XR)_{ij}]^2}, \quad (6)$$

and determines the optimal number of clusters by first minimizing this cost over all rotation matrices  $R$  for each  $k$ , and then selecting the  $k$  that yields the lowest alignment cost. Here,  $k$  corresponds to the number of strategy clusters used in our BR ensemble.

We then train a BR policy for each cluster. Specifically, for the BR policy corresponding to cluster  $i$ , with parameters  $\theta_{\text{BR}}^i$ , we optimize its policy  $\pi_{\theta_{\text{BR}}^i}$  to maximize the expected return when paired with partners from cluster  $i$ . The training objective is defined as

$$J_{\text{BR}}(\theta_{\text{BR}}^i) = \mathbb{E}_{A \sim U(C_i), \tau \sim [\pi_A^1, \pi_{\theta_{\text{BR}}^i}^2]} \left[ \sum_{t=0}^{\infty} \gamma^t r_t \right], \quad (7)$$

where  $U(C_i)$  denotes the uniform distribution over the agent pairs in cluster  $C_i$ . Note that the BR policies are complementary to their corresponding clusters.

**(3) Adaptive Selection:** For adaptation, the cooperators agent should be able to infer the partner’s intent and act accordingly. In the adaptive selection module, we use ToM models to infer the partner’s intent and determine the most behaviorally similar cluster, allowing the agent to select and execute the corresponding complementary BR policy.

To do this, we build a total of  $k+1$  ToM models:  $k$  cluster-specific ToM models, denoted as  $\text{ToM}_{C_i}$  for each cluster  $i$ , which infer the partner’s intent within their respective clusters, and one global ToM model, denoted as  $\text{ToM}_{\text{global}}$ , which infers partners across all clusters. Here, all ToM models are implemented as recurrent networks. In addition, each  $\text{ToM}_{C_i}$  is trained to output a concept representation (i.e., a high-level action) corresponding to the intention of a partner agent within cluster  $i$ , whereas  $\text{ToM}_{\text{global}}$  is trained to predict a concept representation of the partner’s most likely intent, inferred from the history and generalized across clusters. To train the ToM models, we generate trajectories from each cluster and use them to train  $\text{ToM}_{C_i}$  with its respective trajectories, and  $\text{ToM}_{\text{global}}$  with trajectories aggregated across all clusters. All ToM models are optimized according to Eq. 4, as described in Sec. 2.3.

The outputs of the  $k+1$  ToM models are leveraged at test time in the following online adaptation phase to infer which cluster the new partner corresponds to.

### 3.3 Online Adaptation to New Partners

The rationale behind TBS’s adaptation to new partners is that effective coordination requires understanding the partner’s intent and aligning one’s behavior accordingly. By leveraging ToM reasoning to infer the partner’s underlying intention and identify its corresponding cluster, TBS enables the agent to select an appropriate BR policy that matches the partner’s behavior. Specifically, the corresponding cluster is identified by comparing the output of the global ToM model with those of the cluster-specific ToM models; the cluster whose ToM output is closest to that of the global model is regarded as the one corresponding to the new partner. Formally, we compute the divergence between the Bernoulli distributions of the concept predictions from the cluster-specific and global ToM models as  $D(\text{ToM}_{C_i}, \text{ToM}_{\text{global}}) =$

$$\sum_{t'=1}^t \sum_{j=1}^M D_{\text{KL}}(\text{Bernoulli}(c_{t'}^i[j]) \parallel \text{Bernoulli}(c_{t'}^{\text{global}}[j])), \quad (8)$$

where  $c_{t'}^i$  and  $c_{t'}^{\text{global}}$  denote the concept distributions predicted by  $\text{ToM}_{C_i}$  and  $\text{ToM}_{\text{global}}$ , respectively, and  $D_{\text{KL}}(\cdot \parallel \cdot)$  denotes the Kullback–Leibler divergence.


We then select the index  $\hat{i}$  corresponding to the minimum distance,  $\hat{i} = \arg \min_i D(\text{ToM}_{C_i}, \text{ToM}_{\text{global}})$ , and assign the  $\hat{i}$ -th BR policy to collaborate with the new partner. Since the BR policies in our ensemble are specialized for distinct strategy clusters identified by the clustering algorithm, this process can be intuitively viewed as identifying the teammate’s strategy and choosing the most appropriate response. At the beginning of each episode, the cooperators agent randomly selects one policy from the ensemble, as some interaction history is required before the partner’s intent can be inferred.

## 4 EXPERIMENTS

In this section, we evaluate the proposed TBS against several baselines in the Overcooked environment, focusing on ZSC.

### 4.1 Experiment Setup

**Environment.** We evaluate our method in the Overcooked-AI environment [3], a widely used benchmark for ZSC. Specifically, we use the *Onion Soup* scenario, in which two agents must cooperate to prepare and deliver onion soup dishes, mimicking the gameplay dynamics of the commercial game *Overcooked*. Preparing a dish



**Figure 4: ZSC Performance of different methods in 7 Onion Soup layouts in fully observable (left) and partially observable (right) settings. The final “Overall” bar in each plot denotes the average performance across all layouts. Error bars indicate bootstrapped 95% confidence intervals capturing variability over different evaluation seeds. Our TBS method equals or improves on the baseline BR method in almost all layouts.**

requires placing three onions into a pot, waiting 20 frames for the soup to cook, retrieving the cooked soup with a plate, and serving it at the delivery counter. Each successful delivery yields a reward of 20, and the objective is to maximize the number of deliveries within a 400-frame episode. Each agent receives a spatial observation encoding the positions and states of key items and agents, and selects actions from movement directions and a context-dependent *interact* action for picking up, placing, and delivering items.

We assess performance across seven distinct layouts, each designed to elicit different coordination challenges and cooperative behaviors. For example, the *Forced Coordination* layout requires the red agent to pass onions across the counter to the blue agent, whereas the *Cramped Room* layout emphasizes fine-grained movement coordination in tight spaces. All layouts are evaluated under both fully observable and partially observable conditions. The seven layouts are illustrated in Fig. 3 and described in Appendix E.

**Concept Definition.** As discussed in Sec. 2.3, ToM models predict concepts that capture the high-level intentions of the partner. In the Overcooked environment, each concept corresponds to a semantically meaningful high-level action that reflects the agent’s current subgoal (e.g., “onion\_pickup\_from\_pile,” “onion\_drop\_in\_pot,” “dish\_delivery”). These concepts represent distinct uses of the *interact* button (e.g., picking up, placing, or delivering items) and provide interpretable supervision signals for learning partner intent. A detailed description of all concept categories used in our experiments is provided in Appendix C.

**Baselines.** (i) *Random Selection*: randomly selects an agent from our training pool to play out the entire episode, used to assess the effectiveness of adaptive selection. (ii) *Best-Response* (BR): trains a single best-response agent to coordinate with the entire training pool. (iii) *Oracle*: uses the exact policy that was co-trained with the unknown partner, essentially measuring  $J_{SP}$  of the base algorithm, and thus serves as an upper bound. (iv) *CEM*: Uses the cross-entropy metric described in [26] to choose an agent from the training pool to cooperate with the partner. (v) *CBPR*: Uses the intra-episode bayesian belief update proposed in [41] to choose a specialist best-response agent to cooperate with the teammate.

Both the baselines and the proposed TBS require constructing an agent pool, which we create by training 10 agents. *Random Selection* and *CEM* do not require additional RL training, whereas *Best-Response*, *CBPR* and *TBS* involve further training. Specifically, *Best-Response* trains a single cooperator agent by pairing it with one of the 10 agents in the pool, randomly selected at the beginning of each training episode, and *CBPR* and *TBS* pair each cooperator agent with agents randomly selected from its corresponding strategy cluster. We provide the training details in Appendix D.

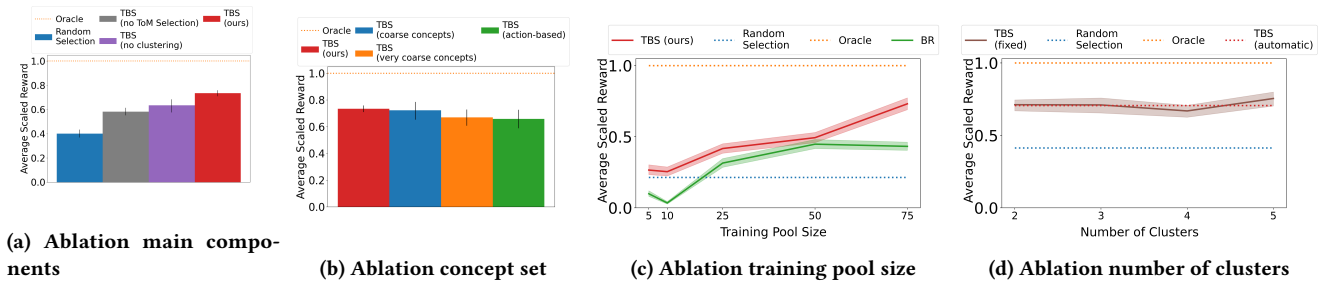
**Evaluation.** We evaluate the any-play performance,  $J_{inter-XP}$ , of the proposed TBS and the baselines. Specifically, this measures performance when playing with held-out agents that are independently trained using VDN with randomly shaped rewards. Note that the held-out agents follow different strategies, as they are trained with distinct reward functions. For this purpose, we generate 10 different held-out agents and use them for evaluation.

Code and experiment scripts are available at: <https://github.com/andrewni2002/ToMZSC>. Appendix is available at: <https://arxiv.org/pdf/2602.12458>.


## 4.2 Experimental Results

We provide results under both fully and partially observable settings, as shown in Fig. 4(a) and (b), respectively. The figures also report the average performance across the considered layouts. Across both settings, TBS achieves higher average performance than the baseline methods, demonstrating superior adaptability in ZSC.

In the fully observable setting, the proposed TBS method matches or exceeds the performance of the BR baseline in six out of seven layouts. While the BR agent tends to learn a static generalist policy that fails to adapt to unseen teammates, TBS leverages explicit ToM-based adaptation through its ensemble of specialized BR policies, enabling rapid identification of the teammate’s strategy and effective adaptation to it, which leads to higher ZSC performance overall. In the partially observable setting, the overall ZSC scores decrease due to the increased difficulty of coordination under limited observability. Nevertheless, TBS matches or outperforms the BR baseline across all layouts, demonstrating its robustness and adaptability even in partially observable environments.



**Figure 5: Experiments ablating the main components of TBS (left), the number of agents in the training pool (middle) and the number of strategy clusters when using spectral clustering with a fixed number of clusters (right). Error bars and shaded regions indicate bootstrapped 95% confidence intervals capturing variability over different evaluation seeds Left: Average ZSC performance of different methods over all 7 onion soup layouts. Both clustering and ToM-based strategy selection improve ZSC performance. Middle: ZSC performance on the difficult *Counter Circuit* layout as a function of the number of agents in the training pool. Larger training pools result in better ZSC performance for both TBS and BR, but TBS improves more. Right: Average performance of TBS with different numbers of clusters over all 7 onion soup layouts. Average scaled rewards are all around 0.7 to 0.8, indicating that our method is robust to the number of clusters.**



**Figure 6: Frequency of different methods selecting the same action as that of the agent co-trained with the unknown teammate, averaged over 7 simple soup layouts**

In both settings, the Random Selection baseline performs significantly worse, highlighting the necessity of adaptive selection and the effectiveness of the ToM-based adaptive module in TBS.

### 4.3 Analysis

In this section, we provide several in-depth analyses for better understanding of the proposed method. Our analysis consists of two main parts. The first focuses on the adaptive selection module, analyzing how effectively it identifies and responds to the partner’s strategy. The second focuses on the partner pool and the corresponding BR policies, examining how the proposed spectral clustering operates and investigating the effects of the number of clusters and the size of the partner pool. We also include a sensitivity analysis on the concept sets used for ToM training to assess their impact on performance.

**4.3.1 Adaptive Selection.** We first evaluate how often each method selects the same action as the oracle, defined as the co-trained agent. Since the oracle was trained jointly with the unknown partner, its policy can be viewed as a reference for the correct cooperative action at each timestep. The results are shown in Fig. 6. TBS achieves higher action accuracy than all baselines except CBPR.


We further perform an ablation on the adaptive selection module. Fig. 5(a) shows the results of TBS without the adaptation module and a random cross-play variant that removes both cross-play clustering and ToM-based selection. Both perform worse than TBS, indicating that the proposed ToM-based adaptive selection module is essential for effective ZSC. Additional ablation studies on window size and steps per selection are provided in Appendix A.

**4.3.2 Clustering Operation.** We first describe how our proposed self-tuning spectral clustering operates. Fig. 7 illustrates a walk-through of our automatic strategy clustering procedure in the *Large Room* layout. As discussed in Sec. 3.2, we first construct the cross-play matrix over the training pool and then compute the similarity score matrix  $s(A, B)$  as defined in Eq. 5. Finally, we apply self-tuning spectral clustering [44] on the similarity matrix, minimizing the cost function from Eq. 6 to determine the number of clusters.

As shown in the left and middle panels of Fig. 7, three clusters (boxed in blue) are identified, indicating that the policies within each blue box exhibit similar behaviors. The number of clusters is determined to be three, as shown in the rightmost panel of Fig. 7, where the STSC cost function reaches its minimum when  $k = 3$ . It is observed that agents with relatively higher cross-play performance tend to be grouped together, which aligns with the intended objective of clustering behaviorally similar strategies. Additional clustering results for other layouts are provided in Appendix B.

In addition, we conduct an ablation on the clustering component. Fig. 5(a) shows the result of TBS without clustering. Removing the clustering step leads to a clear drop in performance, highlighting its importance in the overall framework.

**4.3.3 Size of Partner Pool.** Having a diverse set of partners is one of the most important factors for achieving effective ZSC—exposing agents to a wide range of partners during training improves their ability to collaborate with unseen partners. Therefore, it is crucial to leverage a sufficiently large partner population within a given computational budget, such as memory and compute capacity. As observed in the toy example in Sec. 3.1, performance generally improves as the partner pool size increases. However, while the BR



**Figure 7: Self-tuning spectral clustering in the *Large Room*. Left: XP matrix  $J(\pi_A^1, \pi_B^2)$ . Middle: similarity matrix  $s(A, B)$ . Right: STSC cost minimized at  $k = 3$  (red circle).**

baseline struggles to effectively leverage larger populations, TBS can continue to benefit from increased partner diversity. To verify this observation in a more realistic setting, we conduct experiments in the *Overcooked* environment, varying the size of the partner pool between 5 and 75 in the most challenging layout, *Counter Circuit*, and report the average ZSC performance of both TBS and the BR baseline on 25 held-out evaluation agents.

The corresponding results are shown in Fig. 5(c). It is observed that (i) TBS consistently outperforms BR in ZSC performance, in line with our previous findings; (ii) the ZSC performance of both methods increases with the training population size, which aligns with intuition and prior studies [28, 29, 37, 46]; and (iii) TBS benefits more from larger partner populations, as it can effectively exploit agent diversity to construct an adaptive ensemble, whereas BR learns a static, generalist policy that limits its performance gains from increased population size. This aligns with the result observed in the simple toy example, described in Sec 3.1.

**4.3.4 Number of Clusters.** While we have used self-tuning spectral clustering [44], which automatically determines the number of clusters, we can also treat the number of strategy clusters as a tunable hyperparameter. In this analysis, instead of using self-tuning spectral clustering, we manually set the number of clusters  $k$  and then perform vanilla spectral clustering to construct the strategy clusters. We train a population of 10 training agents and 10 held-out evaluation agents using VDN, vary  $k$  between 2 and 5, and report the average ZSC performance of our method over the seven fully observable *Simple Soup* layouts described above.

The results are shown in Fig. 5(d). TBS (fixed) indicates TBS with a preset number of clusters, whereas TBS (automatic) refers to the original method using self-tuning spectral clustering. The average performance of TBS (fixed) stays around 0.7–0.8, indicating that our method is robust to the choice of  $k$ . In addition, using self-tuning spectral clustering eliminates the need for manual hyperparameter tuning while maintaining comparable performance.

**4.3.5 Concept Set.** The ToM models in TBS predicts a partner’s intention as a distribution over a predefined set of concepts. To understand how the choice of this concept set affects performance, we investigate the sensitivity of TBS to different levels of concept granularity. While we have used a granular concept set containing a total of 44 concepts for TBS, in this section we explore using TBS with less granular and different kinds of concept sets. Specifically,

the TBS (coarse concepts) method uses a coarser concept set of size 11 that does not distinguish between key items at different locations. The TBS (very coarse concepts) uses an even lower resolution concept set containing only 6 concepts. Finally the TBS (action-based) method directly uses the 6 low-level actions as its concept set.

The average ZSC performance of each method over 7 fully observable *Onion Soup* layouts is displayed in Fig. 5b. The average ZSC performances of the four methods are similar, indicating the robustness of TBS to concept set specification. However, we also see that high-level intention-based concept sets outperform low-level action based concept sets, and that more granular concept sets outperform less granular ones, as concept sets that are less granular and less aligned with the problem of high-level coordination are more likely to conflate distinct strategies and choose the wrong specialist best-response agent.

## 5 CONCLUSION

In this work, we presented TBS, a Theory-of-Mind-based Best Response Selection framework designed to enhance zero-shot coordination with novel partners. Unlike conventional population-based approaches that rely on a single static best-response policy, TBS leverages behavioral clustering and ToM-guided policy selection to achieve adaptive collaboration. During evaluation, the agent infers its partner’s behavioral intent and selects the most compatible best-response policy in real time, enabling robust adaptation to unseen strategies. Empirical evaluations in the *Overcooked* environment demonstrate that TBS significantly improves coordination performance across diverse conditions, underscoring the potential of explicit Theory-of-Mind reasoning as a general mechanism for zero-shot coordination.

**Limitation** One limitation of our work is the lack of theoretical analysis on coordination and generalization. We leave this for future work. In addition, our method depends on the specific concept sets used, and underspecification of concept sets could lead to different strategies being conflated by our ToM-based ensemble coordinator.

## ACKNOWLEDGMENTS

This work was supported in part by the DARPA EMHAT Program under Agreement No. HR00112490409, the ONR Award No. N00014-23-1-2840, the Okawa Foundation Research Grant, and the NSF CAREER Award No. 2145077.

## REFERENCES

- [1] Stefano V. Albrecht and Peter Stone. 2018. Autonomous agents modelling other agents: A comprehensive survey and open problems. *Artificial Intelligence* 258 (2018), 66–95. <https://doi.org/10.1016/j.artint.2018.01.002>
- [2] Daniel S. Bernstein, Robert Givan, Neil Immerman, and Shlomo Zilberstein. 2002. The Complexity of Decentralized Control of Markov Decision Processes. *Mathematics of Operations Research* 27, 4 (2002), 819–840. <https://doi.org/10.1287/moor.27.4.819.297> arXiv:<https://doi.org/10.1287/moor.27.4.819.297>
- [3] Micah Carroll, Rohin Shah, Mark K. Ho, Thomas L. Griffiths, Sanjit A. Seshia, Pieter Abbeel, and Anca Dragan. 2019. On the utility of learning about humans for human-AI coordination. In *Proceedings of the 33rd International Conference on Neural Information Processing Systems*. Curran Associates Inc., Red Hook, NY, USA, Article 465, 12 pages.
- [4] Logan Engstrom, Brandon Tran, Dimitris Tsipras, Ludwig Schmidt, and Aleksander Madry. 2019. Exploring the Landscape of Spatial Robustness. In *Proceedings of the 36th International Conference on Machine Learning (Proceedings of Machine Learning Research, Vol. 97)*, Kamalika Chaudhuri and Ruslan Salakhutdinov (Eds.). PMLR, 1802–1811. <https://proceedings.mlr.press/v97/engstrom19a.html>
- [5] Kevin Eykholt, Ivan Evtimov, Earlene Fernandes, Bo Li, Amir Rahmati, Chaowei Xiao, Atul Prakash, Tadayoshi Kohno, and Dawn Song. 2018. Robust Physical-World Attacks on Deep Learning Visual Classification. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- [6] D. Fudenberg and J. Tirole. 1991. *Game Theory*. MIT Press. <https://books.google.com/books?id=3KnuDwAAQBAJ>
- [7] Tobias Gessler, Tin Dizdarevic, Ani Calinescu, Benjamin Ellis, Andrei Lupu, and Jakob Nicolaus Foerster. 2025. OvercookedV2: Rethinking Overcooked for Zero-Shot Coordination. In *The Thirteenth International Conference on Learning Representations*. <https://openreview.net/forum?id=hlvLM3GX8R>
- [8] Amirata Ghorbani, James Wexler, James Y Zou, and Been Kim. 2019. Towards Automatic Concept-based Explanations. In *Advances in Neural Information Processing Systems*, H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett (Eds.), Vol. 32. Curran Associates, Inc. [https://proceedings.neurips.cc/paper\\_files/paper/2019/file/77d2fbc31f6493e350fca61764efb9a-Paper.pdf](https://proceedings.neurips.cc/paper_files/paper/2019/file/77d2fbc31f6493e350fca61764efb9a-Paper.pdf)
- [9] Robert Gibbons et al. 1992. *A primer in game theory*. Harvester Wheatsheaf New York.
- [10] Ian J. Goodfellow, Jonathon Shlens, and Christian Szegedy. 2015. Explaining and Harnessing Adversarial Examples. In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7–9, 2015, Conference Track Proceedings*, Yoshua Bengio and Yann LeCun (Eds.). <http://arxiv.org/abs/1412.6572>
- [11] Jerry Zhi-Yang He, Zackory Erickson, Daniel S. Brown, Aditi Raghunathan, and Anca Dragan. 2023. Learning Representations that Enable Generalization in Assistive Tasks. In *Proceedings of The 6th Conference on Robot Learning (Proceedings of Machine Learning Research, Vol. 205)*, Karen Liu, Dana Kulic, and Jeff Ichnowski (Eds.). PMLR, 2105–2114. <https://proceedings.mlr.press/v205/he23a.html>
- [12] Hengyuan Hu, Adam Lerer, Brandon Cui, Luis Pineda, Noam Brown, and Jakob Foerster. 2021. Off-Belief Learning. In *Proceedings of the 38th International Conference on Machine Learning (Proceedings of Machine Learning Research, Vol. 139)*, Marina Meila and Tong Zhang (Eds.). PMLR, 4369–4379. <https://proceedings.mlr.press/v139/hu21c.html>
- [13] Hengyuan Hu, Adam Lerer, Alex Peysakhovich, and Jakob Foerster. 2020. “otherplay” for zero-shot coordination. In *International Conference on Machine Learning*. PMLR, 4399–4410.
- [14] Hengyuan Hu, Adam Lerer, Alex Peysakhovich, and Jakob Foerster. 2020. “OtherPlay” for Zero-Shot Coordination. In *Proceedings of the 37th International Conference on Machine Learning (Proceedings of Machine Learning Research, Vol. 119)*, Hal Daumé III and Aarti Singh (Eds.). PMLR, 4399–4410. <https://proceedings.mlr.press/v119/hu20a.html>
- [15] Max Jaderberg, Valentin Dalibard, Simon Osindero, Wojciech M. Czarnecki, Jeff Donahue, Ali Razavi, Oriol Vinyals, Tim Green, Iain Dunning, Karen Simonyan, Chrisantha Fernando, and Koray Kavukcuoglu. 2017. Population Based Training of Neural Networks. arXiv:1711.09846 [cs.LG] <https://arxiv.org/abs/1711.09846>
- [16] Jeewon Jeon, Woojun Kim, Whiyoung Jung, and Youngchul Sung. 2022. Maser: Multi-agent reinforcement learning with subgoals generated from experience replay buffer. In *International conference on machine learning*. PMLR, 10041–10052.
- [17] Jiechuan Jiang and Zongqing Lu. 2018. Learning attentional communication for multi-agent cooperation. *Advances in neural information processing systems* 31 (2018).
- [18] Been Kim, Martin Wattenberg, Justin Gilmer, Carrie Cai, James Wexler, Fernanda Viegas, and Rory sayres. 2018. Interpretability Beyond Feature Attribution: Quantitative Testing with Concept Activation Vectors (TCAV). In *Proceedings of the 35th International Conference on Machine Learning (Proceedings of Machine Learning Research, Vol. 80)*, Jennifer Dy and Andreas Krause (Eds.). PMLR, 2668–2677. <https://proceedings.mlr.press/v80/kim18d.html>
- [19] Woojun Kim, Myungsik Cho, and Youngchul Sung. 2019. Message-dropout: An efficient training method for multi-agent deep reinforcement learning. In *Proceedings of the AAAI conference on artificial intelligence*, Vol. 33. 6079–6086.
- [20] Woojun Kim, Whiyoung Jung, Myungsik Cho, and Youngchul Sung. 2023. A Variational Approach to Mutual Information-Based Coordination for Multi-Agent Reinforcement Learning. In *Proceedings of the 2023 International Conference on Autonomous Agents and Multiagent Systems*. 40–48.
- [21] Woojun Kim, Jongeui Park, and Youngchul Sung. 2021. Communication in Multi-Agent Reinforcement Learning: Intention Sharing. In *International Conference on Learning Representations*.
- [22] Woojun Kim and Youngchul Sung. 2023. An adaptive entropy-regularization framework for multi-agent reinforcement learning. In *International Conference on Machine Learning*. PMLR, 16829–16852.
- [23] Pang Wei Koh, Thao Nguyen, Yew Siang Tang, Stephen Mussmann, Emma Pierson, Been Kim, and Percy Liang. 2020. Concept Bottleneck Models. In *Proceedings of the 37th International Conference on Machine Learning (Proceedings of Machine Learning Research, Vol. 119)*, Hal Daumé III and Aarti Singh (Eds.). PMLR, 5338–5348. <https://proceedings.mlr.press/v119/koh20a.html>
- [24] Benjamin Li, Shuyang Shi, Lucia Romero, Huao Li, Yaqi Xie, Woojun Kim, Stefanos Nikolaidis, Michael Lewis, Katia Sycara, and Simon Stepputtis. 2025. Modeling Latent Partner Strategies for Adaptive Zero-Shot Human-Agent Collaboration. *arXiv preprint arXiv:2507.05244* (2025).
- [25] Chenghao Li, Tonghan Wang, Chengjie Wu, Qianchuan Zhao, Jun Yang, and Chongjie Zhang. 2021. Celebrating Diversity in Shared Multi-Agent Reinforcement Learning. In *Advances in Neural Information Processing Systems*, A. Beygelzimer, Y. Dauphin, P. Liang, and J. Wortman Vaughan (Eds.). <https://openreview.net/forum?id=CO87OIEOGU8>
- [26] Huao Li, Tianwei Ni, Siddharth Agrawal, Fan Jia, Suhas Raja, Yikang Gui, Dana Hughes, Michael Lewis, and Katia Sycara. 2021. Individualized Mutual Adaptation in Human-Agent Teams. *IEEE Transactions on Human-Machine Systems* PP (09 2021). <https://doi.org/10.1109/THMS.2021.3107675>
- [27] Pengyi Li, Hongyao Tang, Tianpei Yang, Xiaotian Hao, Tong Sang, Yan Zheng, Jianye Hao, Matthew E Taylor, Wenyuan Tao, and Zhen Wang. 2022. PMIC: Improving Multi-Agent Reinforcement Learning with Progressive Mutual Information Collaboration. In *International Conference on Machine Learning*. PMLR, 12979–12997.
- [28] Yancheng Liang, Daphne Chen, Abhishek Gupta, Simon Shaolei Du, and Natasha Jaques. 2024. Learning to Cooperate with Humans using Generative Agents. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*. <https://openreview.net/forum?id=v4dXL3LsGX>
- [29] Xingzhou Lou, Jiaxian Guo, Junge Zhang, Jun Wang, Kaiqi Huang, and Yali Du. 2023. PECAN: Leveraging Policy Ensemble for Context-Aware Zero-Shot Human-AI Coordination. In *Proceedings of the 2023 International Conference on Autonomous Agents and Multiagent Systems (London, United Kingdom) (AAMAS '23)*. International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, 679–688.
- [30] Keane Lucas and Ross E. Allen. 2022. Any-Play: An Intrinsic Augmentation for Zero-Shot Coordination. In *Proceedings of the 21st International Conference on Autonomous Agents and Multiagent Systems (Virtual Event, New Zealand) (AAMAS '22)*. International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, 853–861.
- [31] Andrei Lupu, Brandon Cui, Hengyuan Hu, and Jakob Foerster. 2021. Trajectory diversity for zero-shot coordination. In *International conference on machine learning*. PMLR, 7204–7213.
- [32] Andrei Lupu, Brandon Cui, Hengyuan Hu, and Jakob Foerster. 2021. Trajectory Diversity for Zero-Shot Coordination. In *Proceedings of the 38th International Conference on Machine Learning (Proceedings of Machine Learning Research, Vol. 139)*, Marina Meila and Tong Zhang (Eds.). PMLR, 7204–7213. <https://proceedings.mlr.press/v139/lupu21a.html>
- [33] Frans A Oliehoek, Christopher Amato, et al. 2016. *A concise introduction to decentralized POMDPs*. Vol. 1. Springer.
- [34] David Premack and Guy Woodruff. 1978. Does the chimpanzee have a theory of mind? *Behavioral and Brain Sciences* 1, 4 (1978), 515–526. <https://doi.org/10.1017/S0140525X00076512>
- [35] Neil Rabinowitz, Frank Perbet, Francis Song, Chiyuan Zhang, S. M. Ali Eslami, and Matthew Botvinick. 2018. Machine Theory of Mind. In *Proceedings of the 35th International Conference on Machine Learning (Proceedings of Machine Learning Research, Vol. 80)*, Jennifer Dy and Andreas Krause (Eds.). PMLR, 4218–4227. <https://proceedings.mlr.press/v80/rabinowitz18a.html>
- [36] Tabish Rashid, Mikayel Samvelyan, Christian Schroeder de Witt, Gregory Farquhar, Jakob Foerster, and Shimon Whiteson. 2020. Monotonic Value Function Factorisation for Deep Multi-Agent Reinforcement Learning. *Journal of Machine Learning Research* 21, 178 (2020), 1–51. <http://jmlr.org/papers/v21/20-081.html>
- [37] DJ Strouse, Kevin R. McKee, Matthew Botvinick, Edward Hughes, and Richard Everett. 2021. Collaborating with Humans without Human Data. In *Advances in Neural Information Processing Systems*, A. Beygelzimer, Y. Dauphin, P. Liang, and J. Wortman Vaughan (Eds.). <https://openreview.net/forum?id=79zWncwO2p>
- [38] Peter Sunehag, Guy Lever, Audrunas Gruslys, Wojciech Marian Czarnecki, Vinicius Zambaldi, Max Jaderberg, Marc Lanctot, Nicolas Sonnerat, Joel Z. Leibo, Karl Tuyls, and Thore Graepel. 2018. Value-Decomposition Networks For Cooperative

- Multi-Agent Learning Based On Team Reward. In *Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems* (Stockholm, Sweden) (AAMAS '18). International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, 2085–2087.
- [39] Zhenggang Tang, Chao Yu, Boyuan Chen, Huazhe Xu, Xiaolong Wang, Fei Fang, Simon Shaolei Du, Yu Wang, and Yi Wu. 2021. Discovering Diverse Multi-Agent Strategic Behavior via Reward Randomization. In *International Conference on Learning Representations*. [https://openreview.net/forum?id=lvRTC669EY\\_](https://openreview.net/forum?id=lvRTC669EY_)
- [40] Johannes Treutlein, Michael Dennis, Caspar Oesterheld, and Jakob Foerster. 2021. A new formalism, method and open issues for zero-shot coordination. In *International Conference on Machine Learning*. PMLR, 10413–10423.
- [41] Haoming Wang, Zhaoming Tian, Yunpeng Song, Xiangliang Zhang, and Zhongmin Cai. 2024. Beyond Single Stationary Policies: Meta-Task Players as Naturally Superior Collaborators. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*. <https://openreview.net/forum?id=HpN4xeDJQF>
- [42] Xihuai Wang, Shao Zhang, Wenhao Zhang, Wentao Dong, Jingxiao Chen, Ying Wen, and Weinan Zhang. 2024. ZSC-Eval: An Evaluation Toolkit and Benchmark for Multi-agent Zero-shot Coordination. In *Advances in Neural Information Processing Systems*, A. Globerson, L. Mackey, D. Belgrave, A. Fan, U. Paquet, J. Tomczak, and C. Zhang (Eds.), Vol. 37. Curran Associates, Inc., 47344–47377. [https://proceedings.neurips.cc/paper\\_files/paper/2024/file/54a7139c548c88e288aa0fcd2bcbeceb-Paper-Datasets\\_and\\_Benchmarks\\_Track.pdf](https://proceedings.neurips.cc/paper_files/paper/2024/file/54a7139c548c88e288aa0fcd2bcbeceb-Paper-Datasets_and_Benchmarks_Track.pdf)
- [43] Chao Yu, Jiaxuan Gao, Weilin Liu, Botian Xu, Hao Tang, Jiaqi Yang, Yu Wang, and Yi Wu. 2023. Learning Zero-Shot Cooperation with Humans, Assuming Humans Are Biased. In *The Eleventh International Conference on Learning Representations*. <https://openreview.net/forum?id=TrwE8l9ajzs>
- [44] Lih Zelnik-manor and Pietro Perona. 2004. Self-Tuning Spectral Clustering. In *Advances in Neural Information Processing Systems*, L. Saul, Y. Weiss, and L. Bottou (Eds.), Vol. 17. MIT Press. [https://proceedings.neurips.cc/paper\\_files/paper/2004/file/40173ea48d9567f1f393b20c855bb40b-Paper.pdf](https://proceedings.neurips.cc/paper_files/paper/2004/file/40173ea48d9567f1f393b20c855bb40b-Paper.pdf)
- [45] Michelle Zhao, Reid Simmons, and Henny Admoni. 2022. Coordination With Humans Via Strategy Matching. In *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. 9116–9123. <https://doi.org/10.1109/IROS47612.2022.9982277>
- [46] Rui Zhao, Jinming Song, Yufeng Yuan, Haifeng Hu, Yang Gao, Yi Wu, Zhongqian Sun, and Wei Yang. 2023. Maximum Entropy Population-Based Training for Zero-Shot Human-AI Coordination. *Proceedings of the AAAI Conference on Artificial Intelligence* 37, 5 (Jun. 2023), 6145–6153. <https://doi.org/10.1609/aaai.v37i5.25758>