

## 1 A CoMeDi Algorithm Details

### 2 A.1 Mixed-play Buffer Collection

3 Mixed-play consists of two phases in each episode: mixed-state generation and self-play. The “input”  
 4 policies are the policy for the convention we are currently training  $\pi_n$  and the partner policy used for  
 5 cross-play optimization  $\pi^*$  (using the same notation from Eq 7). First, we choose a random timestep  
 6 within the episode that represents the length of the first phase (Line 2). Until this timestep occurs,  
 7 we randomly sample the action from self-play or cross-play for both players (Lines 5-9). We do not  
 8 store any of these transitions in the training buffer. Instead, we use the state at the last timestep and  
 9 pretend that this is the initial state of the environment. For the rest of the timesteps, we perform the  
 10 second phase, by taking self-play actions and store that in the buffer (Lines 10-12). When optimizing,  
 11 we treat this new buffer the same as we would treat self-play, but modified with a positive weight  
 12 hyperparameter,  $\beta$ , representing the importance of mixed-play.

---

#### Algorithm 1: Generating Mixed Play Buffer

---

**Input:** policies  $\pi_n, \pi^*$ , MDP  $M$   
**Output:** Replay buffer from running mixed-play.

```

1  $s_0, R \leftarrow s, 0$  // start state, reward
2  $t \sim \text{uniform}(1, T)$  // length of mixed-state generation phase
3 for  $i \leftarrow [0, T)$  do
4    $o^1, o^2 \leftarrow o(s_i)$ 
5   if  $i < t$  then
6      $\pi_1^m \leftarrow \text{Randomly choose } \pi_n \text{ or } \pi^*$ 
7      $\pi_2^m \leftarrow \text{Randomly choose } \pi_n \text{ or } \pi^*$ 
8      $a_i^1, a_i^2 \leftarrow \pi_1^m(o^1), \pi_2^m(o^2)$ 
9      $s_{i+1}, - \leftarrow \text{Step in } M \text{ with } (a_i^1, a_i^2)$ 
10  else
11     $a_i^1, a_i^2 \leftarrow \pi_n(o^1), \pi_n(o^2)$  // self-play
12     $s_{i+1}, r_i \leftarrow \text{Step in } M \text{ with } (a_i^1, a_i^2)$  // keep reward in phase2
Return:  $\text{ReplayBuffer}(s_{t:T}, a_{t:T}^1, a_{t:T}^2, r_{t:T})$ 

```

---

### 13 A.2 Full CoMeDi Algorithm

14 Simplified pseudocode for the CoMeDi algorithm is presented below. Note that conventions are  
 15 generated in a sequential order, with  $\pi_1$  being trained with standard self-play and each  $\pi_i$  being  
 16 trained with the awareness of prior conventions,  $D_{1:i-1}$ . The arg max operation in line 4 is estimated  
 17 empirically by simulating a fixed number of rounds of cross-play in the environment with each  
 18 existing convention and selecting the convention with the highest cross-play as  $\pi^*$ .

---

#### Algorithm 2: Diverse Conventions with CoMeDi

---

**Input:** Number of policies to generate  $n$   
**Output:** Diverse set of conventions  $D$

```

1  $D \leftarrow (\pi_1, \dots, \pi_n)$ , parameterized by  $(\theta_1, \dots, \theta_n)$ 
2 Train  $\pi_1$  with standard self-play
3 for  $i \in \{2, \dots, n\}$  do
4   while policy  $\pi_i$  has not converged do
5      $\pi^* \leftarrow \arg \max_{\pi^* \in D_{1:i-1}} J(\pi_i, \pi^*)$ 
6      $\tau_{SP} \leftarrow \text{GetRollout}(\pi_i, \pi_i)$ 
7      $\tau_{XP} \leftarrow \text{GetRollout}(\pi_i, \pi^*)$ 
8      $\tau_{MP} \leftarrow \text{MixedPlayRollout}(\pi_i, \pi^*)$ 
9     Estimate  $J(\pi_i, \pi_i), J(\pi_i, \pi^*), J_M(\pi_i, \pi^*)$  with  $\tau_{SP}, \tau_{XP}, \tau_{MP}$ 
10     $\theta_i \leftarrow \theta_i + \nabla_{\theta_i} [-J(\pi_i, \pi_i) + \alpha J(\pi_i, \pi^*) - \beta J_M(\pi_i, \pi^*)]$ 
Return:  $D$ 

```

---

### 19 A.3 Implementation Details

20 We base the implementation of our algorithm on the Multi-Agent PPO algorithm (MAPPO) [8].  
21 MAPPO is an actor-critic method which, in standard self-play, trains a single actor network for the  
22 policy and a single critic network for the value function [7]. To adapt MAPPO to train a pool of  $n$   
23 conventions using our proposed mixed-play algorithm, we train  $n$  actor networks,  $n$  self-play critic  
24 networks, and  $n^2 - n$  cross-play critic networks, each representing a cross-play pairing between  
25 the  $n$  conventions. We also use the PantheonRL library [6] to design our environments and training  
26 algorithms since it is designed to handle dynamic training interactions like cross-play and mixed-play.  
27 We have also integrated CoMeDi with a new GPU-accelerated simulation framework, which enables  
28 the collection of large batches of cross-play and mixed-play buffers in parallel with the collection of  
29 self-play buffers (more details to be revealed after the GPU simulation framework is released from  
30 double-blind review).

31 Moreover, instead of training the whole batch of  $n$  diverse agents in parallel, in practice we se-  
32 sequentially grow the set of agents one at a time, keeping the previous agents fixed. We find that  
33 sequential generation leads to more stable training: since the previous agents are fixed, the diversity  
34 regularization term becomes a reward shaping term that is only a function of the policy of the current  
35 agent.

### 36 A.4 Practical Guidelines for Hyperparameter Tuning

37 There are some safe choices for hyperparameters that work well in general, which we used to tune  
38 the hyperparameters for our experiments. First, we observe that directly using the best MAPPO  
39 hyperparameters for the particular environment, like learning rate and the model architecture, transfers  
40 well to CoMeDi. To find the cross-play weight ( $\alpha$ ), fix the mixed-play weight to 0 and find the lowest  
41 value for the cross-play weight such that increasing it further does not significantly increase the  
42 self-play score or decrease the cross-play score. If the cross-play weight is too high, this may cause  
43 training instabilities since the updates to increase the self-play score would be directly counteracted  
44 by the updates to decrease the cross-play score. Finally, choose the value of the mixed-play parameter  
45 such that the average mixed-play score (for the second half) is slightly less than half of the self-play  
46 score, which would indicate that self-play is able to smoothly continue from any mixed-play state.

47 The guidelines for choosing hyperparameters works well in general, but domain knowledge of the  
48 environments also helps. If your specific environment also has some indicators of handshakes, you  
49 can also use those to determine if handshakes are still happening. Furthermore, environments where  
50 partners' actions are not visible, like Blind Bandits, do not require mixed-play at all because hand-  
51 shakes are impossible. In practice, we have seen that CoMeDi is relatively robust to hyperparameters  
52 and it gives reasonable policies with a cross-play weight of 0.5 and a mixed-play weight of 1, even if  
53 they are not perfectly optimal.

54 Choosing a population size is also an art, but due to the sequential nature of CoMeDi, prior conventions  
55 are unaffected by the generation of later conventions. The choice of algorithm for generating a  
56 "convention-aware agent" would likely influence the number of conventions to use for the diverse set.

### 57 A.5 Extending to Larger Teams

58 When using CoMeDi for cooperative games with more than 2 players, we can follow the same  
59 algorithm presented in 2, but we have to be a bit careful when collecting rollouts.

60 To collect the cross-play buffer between  $\pi_i$  and an existing  $\pi^*$  in a  $k$ -player game, we can randomly  
61 assign each player to one of the conventions but keep those assignments consistent throughout the  
62 duration of the episode. The same logic regarding the minimization of cross-play rewards still applies  
63 since semantically similar conventions would result in a high reward even when the team contains a  
64 mix of the conventions.

65 To collect the mixed-play buffer, we would randomly choose between the two conventions for each  
66 player at each timestep during the mixed-state generation phase. We would still treat self-play as  
67 normal by using the convention being trained as the convention for all players.

## 68 B Experiments

### 69 B.1 Choice of Baselines

70 Throughout this work, we compare the performance of CoMeDi against pure MAPPO, a modified  
71 version of ADAP, and a pure cross-play minimization baseline (CoMeDi with  $\beta = 0$ ). We do not  
72 directly use ADAP because we find that its generative capabilities from parameter sharing often  
73 limits the diversity of its agents, resulting in very high cross-play between agents in the population.  
74 By adding its diversity loss to the MAPPO algorithm, we eliminate the confounding variables of  
75 the base PPO implementation and the parameter sharing in the original algorithm. We also do not  
76 directly compare against TrajeDi because a fundamental aspect of its algorithm is the concurrent  
77 generation of a common best response agent, which implicitly optimizes for high cross-play within  
78 the “diverse set” of policies. However, upon analyzing TrajeDi’s diversity loss, we note that it is  
79 very similar in practice to ADAP’s loss. Therefore, we believe that our modified version of ADAP  
80 is the fairest representative of statistical diversity approaches. Concurrent work that implements a  
81 technique similar to our pure cross-play minimization provides some other benchmarks with MAVEN  
82 and TrajeDi which show similar results to what we experienced with our modification of ADAP.

83 We do not compare our work to those in section 2.3, because they require more assumptions regarding  
84 the training domain. In particular, the strength of reward shaping methods (and other domain  
85 engineering techniques) is highly dependent on the manual design of the appropriate environment  
86 parameter space. However, CoMeDi can be interpreted as an *automatic* technique for reward shaping  
87 to form diverse conventions, which can potentially eliminate the need to engineer the environment  
88 parameter space for domain randomization.

### 89 B.2 Hyperparameters

Table 1: Common hyperparameters for agents in Blind Bandits and Balance Beam

hyperparameters	value
fc layer dim	512
num fc	2
activation	ReLU
network	mlp
ppo epochs	15
mini batch	1

Table 2: Hyperparameters in Blind Bandits

hyperparameters	value
buffer length	200
environment timesteps	10000
actor/critic lr	$2 \times 10^{-5}$
linear lr decay	False
entropy coef	0.01
ADAP $\alpha$	0.2
CoMeDi $\alpha$	1.0
CoMeDi $\beta$	0.0

Table 3: Hyperparameters in Balance Beam

hyperparameters	value
buffer length	1250
environment timesteps	50000
actor/critic lr	$2.5 \times 10^{-5}$
linear lr decay	True
entropy coef	0.01
ADAP $\alpha$	0.05
CoMeDi $\alpha$	0.3
CoMeDi $\beta$ (ablation)	0.0, 0.25, 0.5, 1.0

Table 4: Hyperparameters in Overcooked (only training set)

hyperparameters	value
CNN Kernel Size	$3 \times 3$
fc layer dim	64
num fc	2
activation	ReLU
rollout threads	50
buffer length (per thread)	200
environment timesteps	1000000
ppo epochs	10
actor/critic lr	$1 \times 10^{-2}$
linear lr decay	True
entropy coef	0.0
ADAP $\alpha$	0.025
CoMeDi $\alpha$	0.5
CoMeDi $\beta$	0.0, 1.0

Table 5: Hyperparameters in Overcooked (convention-aware agents)

hyperparameters	value
CNN Kernel Size	$3 \times 3$
fc layer dim	64
num fc	2
activation	ReLU
rollout threads	50
buffer length (per thread)	200
environment timesteps (SP phase)	200000
ppo epochs (SP phase)	100
lr	$1 \times 10^{-2}$
entropy coef	$1 \times 10^{-3}$

### 90 B.3 Compute Resources

91 We conducted our experiments with our lab’s internal cluster. We only used the Intel Xeon Silver  
 92 4214R CPU for training in Blind Bandits and Balance Beam. For the Overcooked experiments, we  
 93 used an additional NVIDIA TITAN RTX, which required around 3 hours per configuration. However,  
 94 this time can be reduced significantly by disabling deterministic behavior in CUDA. Current work on  
 95 optimizing performance with the GPU-accelerated simulator also shows that this time can be reduced.

96 **B.4 Blind Bandits Environment Description**

97 The Blind Bandits environment is a collaborative two-player Dec-POMDP where each player takes  $k$   
 98 steps,  $k \geq 2$ , and at each step they have to choose to go left (L) or right (R). Each player is given their  
 99 own past history of actions in the episode, but cannot see the others' actions. There are two ways to  
 100 get a positive score. To get a score of  $S$ , the first player's first step must be L, and the second player's  
 101 last step must be L. To get a score of  $G > S$ , the first player's first step must be R, but all following  
 102 steps by all players must be L except for the second player's last step, which must be R. If the players  
 103 fail to coordinate, they get a score of 0.

104 The policies that converge to  $S$  are all compatible with one another since all of them have agent 1 start  
 105 with L and have agent 2 end with L (shown in blue). If two agents are playing randomly, there is a  
 106 0.25 probability that they will get a reward of  $S$ . Meanwhile, only one trajectory will result in a score  
 107 of  $G$  (shown in orange), so there is only a  $2^{-2k}$  chance that random agents will get a reward of  $G$ .  
 108 Notice how these two conventions ( $S$  and  $G$ ) are entirely incompatible with one another because their  
 109 first and last moves must be different, and are therefore different equilibria. Ideally, we want to find a  
 110 representative policy for both conventions using a technique that finds a diverse set of conventions.

111 **B.5 Blind Bandits Baseline Results**

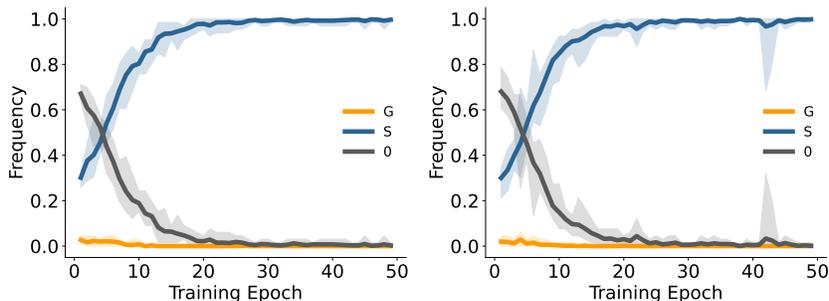


Figure 1: Frequency of self-play scores during the optimization of two conventions using statistical diversity (the ADAP loss) from 10 independent seeds. The blue line indicates an  $S$  convention, the orange line indicates a  $G$  convention, and the black line indicates a score of 0 (all other paths lead to 0). With ADAP, neither convention converges to  $G$ . We choose  $k = 3$  as the number of steps in the environment.

112 In Figure 1, we see that using ADAP to train a set of two agents almost always leads to them both  
 113 converging to  $S$ . For this result, we used the ADAP diversity weight of 0.2, the most common  
 114 value used in the original paper. Increasing this weight sometimes results in discovering a few  $G$   
 115 conventions, but this is inconsistent and results in more unstable training.

116 In fact, training an agent to converge to the  $G$  convention is more difficult than it appears. Off-the-shelf  
 117 MARL algorithms typically converge to the  $S$  convention as well. At the first training epoch, the  
 118 actor chooses random actions, so it will get a score of  $S$  1/4 of the time while it gets a score of  $G$   
 119 with a probability of  $2^{-2k}$ . The critic network (in both the decentralized and centralized setting) will  
 120 learn that the value of choosing to go left in the first state is  $S/2$  while the value of choosing to go to  
 121 the right is  $G/2^k$ . Therefore, the policy updates will favor the  $S$  convention in the first epoch as long  
 122 as  $S > G/2^{k-1}$ , and this separation gets larger until it fully converges to an  $S$  convention.

123 Most existing zero-shot coordination (ZSC) algorithms would also converge to the  $S$  convention. The  
 124 TrajeDi algorithm finds a common best response to a set of conventions that have a wide diversity in  
 125 trajectory, but the  $G$  convention has only one trajectory so it will never be chosen. The Off Belief  
 126 Learning (OBL) algorithm would also only converge to the  $S$  convention because it acts on the belief  
 127 that its partner is a random agent. Higher levels of OBL will be stuck in the  $S$  convention since they  
 128 assume that their partner is from a lower level of OBL. Note that the purpose of ZSC is to have a high  
 129 cross-play between independent runs of the same algorithm, so both of these techniques succeed in  
 130 that manner, but this does not imply that these algorithms would have high cross-play with humans.

## 131 B.6 Balance Beam Environment Description

132 At the start of the game, the location of each player is randomly initialized. At each timestep, both  
133 players take simultaneous actions, and can move 1 or 2 locations to the left or right, but cannot stay  
134 still. If an action leads to them falling off the line, they get a score of -1 multiplied by the number of  
135 remaining timesteps. Otherwise, they get a score of  $-d(s_1, s_2)/5$  where  $d$  is the distance between  
136 the two player’s locations. Finally, if both players are on the same spot at the end of their turn, they  
137 get an extra point. Each episode lasts two timesteps, and the players see the result of the first step’s  
138 actions before making their second move.

139 A perfect convention would always get a score of 2.0, because players can score +1 in each of the  
140 two timesteps. The worst score is -2.0, which occurs when one player moves off of the line at the first  
141 time step.

142 We also hand-code some conventions to see if CoMeDi discovers conventions that are similar to those  
143 that humans follow. There are two simple conventions: a left-biased convention and a right-biased  
144 convention. These dictate how ties are broken when multiple actions are equally optimal. For  
145 example, if the distance between two players is 1 step, like in Figure 6 of the main text, the player  
146 from a left-biased convention would want to move to the open space to the left since staying in one  
147 spot is not an option.

## 148 B.7 Balance Beam Baseline Results

149 We designed the Balance Beam environment to distill the issue of handshakes when minimizing  
150 cross-play, which is why the main text only emphasizes the impact of the  $\beta$  hyperparameter in  
151 CoMeDi. However, we also trained ADAP in this environment to see how it compares to the other  
152 approaches.

153 The first trained convention gets a score of 2.0, and gets an expected score of 0.768 and 1.112 when  
154 paired with the left and right-biased agents respectively.

155 The second trained convention gets a score of 1.808, and gets an expected score of 1.048 and 0.176  
156 when paired with the left and right-biased agents respectively. Its cross play score with the first  
157 trained convention is 0.392.

158 When training the ADAP agents, we observed a very unstable training process resulting in very low  
159 self-play scores with typical values of the diversity weight parameter. For this reason, we had to  
160 choose a relatively low value for the loss parameter (0.05) in order to make a fair comparison with  
161 our technique.

## 162 B.8 Overcooked Agent Generation

163 For our experiments, we generated 2 baseline agents and 2 convention-aware agents using CoMeDi.  
164 The first baseline agent, which we refer to as “SP”, was trained with pure self-play, and we tuned the  
165 hyperparameters to maximize its self-play score. For all other generated agents, we maintained the  
166 same hyperparameters that are inherent to MAPPO, but we would tune the diversity weights specific  
167 to each algorithm. Our second baseline, which we refer to as “ADAP”, is a convention-aware agent  
168 to a population of 8 agents trained with ADAP with no parameter sharing and a diversity weight of  
169 0.025. The XP agent was also a convention-aware agent to a population of 8 agents trained with  
170  $\alpha = 0.5$  and  $\beta = 0.0$ . The CoMeDi agent was the same as XP, but its population was trained with  
171  $\alpha = 0.5$  and  $\beta = 1.0$ .

172 For each layout of Overcooked, we can determine the expected number of dishes to be served by each  
173 agent in self-play. In the Cramped Room, SP averages 4.36 dishes, ADAP averages 2.75 dishes, XP  
174 averages 4.68 dishes, and CoMeDi averages 5.52 dishes. Meanwhile, in the Coordination Ring, SP  
175 averages 3.47 dishes, ADAP averages 1.90 dishes, XP averages 3.06 dishes, and CoMeDi averages  
176 5.36 dishes.

177 We can also calculate the expected reward for the best and worst performing agents in the training sets  
178 of ADAP, XP, and MP. In Cramped Room, ADAP’s average rewards spanned 0 to 5.98, XP’s rewards  
179 spanned 4.12 to 5.96, while MP’s rewards spanned 5.0 to 5.88. In Coordination Ring, ADAP’s  
180 average rewards spanned 0 to 4.965, XP’s rewards spanned 2.75 to 4.56, while MP’s rewards spanned  
181 4.92 to 5.99.

182 When training agents with ADAP, we would frequently see a few policies with very high expected  
183 returns with the remaining policies having low scores. We attempted to tune ADAP’s diversity weight  
184 to enable more balanced generation, but this issue continued to persist even with a final diversity  
185 weight significantly lower than the values presented in the original paper.

## 186 **B.9 Overcooked User Study Setup**

187 The human-AI interaction portion of this research was approved by our IRB.

188 Our total population consisted of 25 paid participants with varying prior experiences in Overcooked,  
189 recruited through Prolific. We did not impose any conditions on participation through Prolific except  
190 for requiring “proficiency in English language.” We paid \$10.33 (US dollars) per participant for  
191 approximately 40 minutes of time, equivalent to \$15.50 per hour. We titled the study “Playing with  
192 AI Agents in the Overcooked Video Game” with the following description:

193 “In this study, participants will play with 4 different AI agents in 2 settings of the Overcooked  
194 game. Our goal is to understand how well trained agents can work with humans in tasks that require  
195 coordination. Your task is to try to work with the AI to cook many “dishes of onion soup” within a  
196 40-second time limit. You will also fill out short surveys to judge the quality of the AI agents. You  
197 will play 20 games in total.

198 To play the game, you need to use a keyboard with arrow keys and a space bar. Desktops, laptops,  
199 and tablets with keyboards may be used. The AI agents runs on your browser and are designed to be  
200 lightweight so they should be compatible with most hardware.

201 The games will be fast-paced, and we may reject submissions that are consistently unable to score  
202 points in the game. As long as we can see that you are trying to score points, your submission will be  
203 accepted.

204 NOTE: Please only accept this study if you have at least 8 GB RAM. The game will take up to 2 GB  
205 of RAM since we are loading models onto your computer so you have a smoother experience. If the  
206 game freezes, please let us know and refresh the page. Please use Firefox, Chrome, or Safari.”

207 To ensure that participants across the world are able to complete the study without excessive latency,  
208 we run all models on the users’ devices directly through tensorflow-js.

209 Each user played 2 games in a layout independently before playing each AI agent for two 40-second  
210 rounds in a random order. The users first played with all agents in the Cramped Room environment  
211 before playing with all agents in the Coordination Ring environment.

212 We also asked the users to fill out a short survey with qualitative questions about each partner after  
213 playing both rounds in any configuration using the 7-point Likert scale.

214 We presented the following 8 statements (in order):

- 215 1. The AI followed my lead when making decisions.
- 216 2. The AI agent frequently blocked my progress.
- 217 3. The AI was consistent in its actions.
- 218 4. The AI always made reasonable actions throughout the game.
- 219 5. I would like to collaborate with this AI in future Overcooked tasks.
- 220 6. The AI’s actions were human-like.
- 221 7. I trusted the AI agent in making good decisions.
- 222 8. The AI agent was better than me at this game.

223 Note that a higher Likert score is better for all of the prompts except for the second question.

224 We also presented an optional free-response section with “Other comments or observations?” that  
225 users could fill out at the end of each survey.

226 To determine statistical significance, we use the paired Student t-test between the scores across  
227 different AI agents for each layout.

228 We also asked a smaller set of users (8 users) to play with an expert human player in-person after  
229 playing with all AI agents to determine what human-level performance would entail without allowing

230 explicit communication. This expert human player was trained to adapt to different play styles, and  
 231 users reported that this human player was “very good” at the game, noting fast reflexes. This set of  
 232 users was mutually exclusive from the Prolific set, since the games needed to be played in person.  
 233 We determine if there is a statistically significant difference between the AI agents and the human  
 234 expert through the unpaired t-test.

### 235 B.10 Cramped Room Overcooked Results

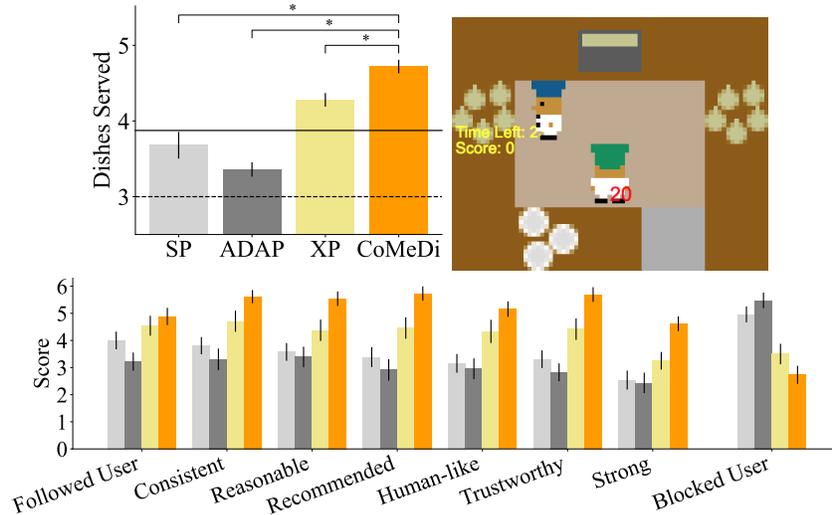


Figure 2: Results for the Cramped Room: average scores (top left), environment visualization (top right), and user survey feedback using the 7-point Likert scale (bottom). Higher average scores are better. The dashed horizontal line indicates the average score when a player is working alone while the solid horizontal line indicates the average score when paired with an expert human. Higher is better for all but the last feedback score. Error bars represent the standard error.

236 In the Cramped Room Environment, players have a very small area to move around in, and they need  
 237 to coordinate around the management of ingredients and the usage of plates. Given the simplicity of  
 238 the setting (relative to the Coordination ring), humans might be able to adapt quickly to AI players  
 239 since the patterns of movement can be very clear after interacting for a short time. In particular, we  
 240 note that this layout has lower convention dependence than the Coordination Ring.

241 The results of the user study in the Cramped Room environment are in Figure 2. In terms of scores,  
 242 we see that CoMeDi (score of 4.72) performs the best, followed by XP (4.28), SP (3.68), and ADAP  
 243 (3.36). In terms of statistical significance, CoMeDi outperforms all AI agents ( $p < 10^{-3}$ ) and the  
 244 expert human ( $p < 10^{-4}$ ). XP also outperforms ADAP, SP, and the expert human ( $p < 0.01$  for  
 245 each).

### 246 B.11 User Free Response Feedback

247 User feedback for SP in the Cramped Room environment:

- 248 • This AI trusted me more than the others. He waited for me to get the plate and deliver the  
 249 soup. The other bots rushed to do it themselves.
- 250 • I could not understand what the AI was doing
- 251 • The AI was slow in decision making
- 252 • Ai was a bit hesitant sometimes

253 User feedback for ADAP in the Cramped Room environment:

- 254 • The AI sometimes blocked me. He held the onion in his hands, and it was already cooked,  
255 and I had the plate in my hands, but I couldn't get to pick up the soup because he was  
256 blocking it. I think he was trying to place the onion even tho there was no space to place it  
257 in.
- 258 • AI was blocking me whole time
- 259 • i didnt like this one...
- 260 • It was somehow worse than S, it blocked the plate for like 7 seconds
- 261 • the ai seemed very confused
- 262 • He tried to do things right, but sometimes he seemed confused.

263 User feedback for XP in the Cramped Room environment:

- 264 • This AI was more lazy. He only held the plate in his hand and waited for me to cook the  
265 soup. The other AI wasn't that way (AI D). He cooked and adapted in getting the plate and  
266 everything. This bot did the exact opposite.
- 267 • I did not like the fact that it could not path find a different way to deliver the full plate.
- 268 • The AI seems to be faster
- 269 • ai was very good at this game
- 270 • Once he picked up a plate after I got one plate, then he never stopped that action until I put  
271 the next 3 onions in the pot.

272 User feedback for CoMeDi in the Cramped Room environment:

- 273 • This bot was by far the best. He was very consistent in everything he did, his moves were  
274 all correct and I felt very good, working with him.
- 275 • this one was not so bad
- 276 • This AI could work by himself, literally, and get the same 100 score, I believe.

277 User feedback for SP in the Coordination Ring environment:

- 278 • The AI did the correct thing, but when it came to pressing the space bar, he struggled. Even  
279 tho he just needed to press the space bar, and he stood right in front of it, and it was just one  
280 space bar away, he decided to wait in front of it for some seconds and then press it, which  
281 concluded us to make less soup overall.
- 282 • the first soup was ok and then AI like "freeze"...it was weird.
- 283 • ok so this one completly stood in my way while i was trying to get the plate he just stood  
284 there with an onion on his hand. Aside from the ai i really likes this study and experiment  
285 and hope one day ill be able to participate in one of them again!
- 286 • At first it looked like the cooperation was working until when it just stopped doing anytihng  
287 and occasionally just blocked the way to the stoves. Towards the end it just standed in the  
288 corner doing nothing.
- 289 • Yes it was consistenly bad
- 290 • He didn't know what to do when picking a plate and there were onions missing in a pot.

291 User feedback for ADAP in the Coordination Ring environment:

- 292 • This bot blocked me a lot of times. He also just stood around the onions and did nothing,  
293 which stopped us both. He sometimes got a plate in his hand and just went up and down all  
294 the time until he decided to play normally again. Would not play with this bot.
- 295 • i also didnt like this ai.
- 296 • This one was very bad at making decisions and pathfinding
- 297 • the AI slowed the process

- 298 • AI was better than the other ones so far, but still got a lot in the way which is frustrating to  
299 play with.
- 300 • He was in the way a lot and almost never knew what to do next, or where to go.

301 User feedback for XP in the Coordination Ring environment:

- 302 • He did very good in the first round, we always went clockwise. It was perfect. In the second  
303 round, he decided to be the lazy chef. He was too lazy to pick up an onion or the plate.
- 304 • It's moves felt very scattered and didn't make sense to me.
- 305 • The AI kept getting in the way and got nothing done.
- 306 • The AI was helpful
- 307 • This one was doing great decisions, but sometimes he seemed confused. In general I liked it.

308 User feedback for CoMeDi in the Coordination Ring environment:

- 309 • This bot is GODLIKE! He did everything correct, he adapted to my moves like he can see  
310 into the future. Just great playing with him. The most efficient by far.
- 311 • it seems sometimes you get into a flow, which the AI breaks after a while
- 312 • He knew exactly what to do next.
- 313 • The experiment was very cool because every AI was unique. The best AI was AI M by far,  
314 my most favorite. If I was a chef, I would definitely hire him!

## 315 C Related Work

316 Concurrent to our work, methods for training diverse agents with respect to reward have been  
317 proposed.

318 In LIPO [3], they also consider cross-play as a diversity metric, but do not address the critical issue of  
319 handshakes that arise when minimizing cross-play. Their results reinforce our observations when the  
320 mixed-play weight,  $\beta$ , is set to 0. However, as we have noticed in our Balance Beam simulation results  
321 and the Overcooked user study, mixed-play has a large impact on creating good-faith conventions.

322 The ADVERSITY [4] work considers a zero-shot coordination framework in the game of Hanabi.  
323 They propose a belief reinterpretation model to address a similar “sabotaging” behavior that we  
324 experienced. This model is designed to tackle the issue of handshakes by finding a plausible  
325 distribution of self-play states that would result in the same observation received from cross-play and  
326 use this while training so agents cannot discriminate between self-play and cross-play observations.  
327 However, it is unclear whether belief reinterpretation would help in the games we examine in this  
328 paper since cross-play can potentially encounter observations that are completely impossible under  
329 self-play. Specifically, in Overcooked, all agents have access to the state from the prior frame, but  
330 conventions exist as a way to manage the workload of tasks between partners. In Hanabi, a core part  
331 of the game is predicting the underlying state of the game given the observation. As such, conventions  
332 are based around communicating information about the state implicitly through actions.

333 ADVERSITY uses the fact that multiple trajectories can generate the same action-observation history,  
334 which enables it to gain very strong results in Hanabi. However, this technique would fail in our  
335 settings because there are observations that are only possible under cross-play and not self-play,  
336 so belief reinterpretation would not be helpful. Therefore, although ADVERSITY and CoMeDi  
337 both attempt to address the problem of handshakes, their core assumptions are entirely different.  
338 Since CoMeDi does not perform explicit belief reinterpretation, it will not be competitive with  
339 ADVERSITY on Hanabi, but it would still be able to train a sequence of agents.

340 In games where implicit communication to predict the underlying state is the core task, ADVERSITY  
341 is a strong choice for training a diverse set of agents. However, CoMeDi would be more effective  
342 in tasks where a team needs to divide a workload or commit to a particular strategy for effective  
343 coordination. We believe that these scenarios are more similar to the tasks that one would encounter  
344 in a robotics domain or typical video game setting.

## 345 **D Limitations**

346 Although our technique of creating a convention-aware agent using CoMeDi was able to surpass  
347 human-level performance in Overcooked, this technique has some drawbacks. Since the policy has no  
348 memory, training a convention-aware agent on a diverse set may lead to the AI breaking conventions  
349 established with humans earlier in the game, which is why one user reported that the CoMeDi  
350 agent sometimes breaks the established flow in Coordination Ring. Also, the BC-based algorithm  
351 for generating the convention-aware agent is often unable to account for human suboptimalities or  
352 transitions between conventions. For instance, some agents in Cramped Room would pick up an  
353 onion and expect the human to pick up a plate. If the human does not comply, it simply stands around  
354 instead of dropping off the onion and picking up a plate itself, because this type of action is never  
355 experienced under self-play for any convention.

356 On a theoretical level, CoMeDi does not provide any guarantees regarding the quality of agents. This  
357 is not unique to our algorithm, as statistical diversity techniques and reward shaping often changes  
358 the environment to the extent that “equilibrium conventions” are no longer guaranteed. Also, our  
359 solution to handshakes, mixed-play, implicitly assumes that re-establishing handshakes will come at  
360 some expense to the self-play scores. If this assumption is violated, as is the case with cheap-talk  
361 signals that aren’t necessary for coordination, handshakes can be re-established at every timestep,  
362 effectively bypassing the mixed-play optimization. The issue of cheap-talk is a very tricky case in  
363 general when attempting to define diversity or robustness, because signals have no implicit meaning.  
364 In these cases, environment designers can remove extraneous cheap-talk signals or add nonuniform  
365 cost to communication, which has been effective in the realm of zero-shot communication [1].

## 366 **E Broader Impact**

367 We believe that CoMeDi can have a positive impact on game design and human-AI interaction in  
368 general. Being able to generate diverse conventions can allow game designers to understand the  
369 different strategies that players might try to use before extensive play-testing. Effective zero-shot  
370 coordination techniques would also help reduce the risk of misaligned conventions. We observe that,  
371 with proper tuning of the mixed-play weight, the convention-aware agent trained with CoMeDi learns  
372 to follow the lead of the human player, as indicated by the “followed user” section of the user survey.  
373 This is important for safety-critical applications like human-robot interaction tasks, because an overly  
374 assertive robot could unintentionally harm a human.

375 As a tool, it can directly be used for harmful ends, such as making it easier to cheat in multi-player  
376 games or generally conduct harm on others. Another potential effect of developing effective artificial  
377 zero-shot collaborators is that it could lead to more social withdrawal. In particular, if people who  
378 play video games start to strongly prefer playing with super-human AI collaborators over other  
379 humans, we may see people play less games with other people, which could counteract the prosocial  
380 benefits of cooperative gaming [5]. We therefore urge potential game designers and publishers who  
381 want to use CoMeDi to generate AI partners to evaluate the impact that artificial agents would have  
382 on their community.

## 383 **F Creative Assets**

384 Custom assets for this paper were digitally created by the authors without the assistance of AI image  
385 generation models. Stylistic inspiration was drawn from the Overcooked figures in [2] under the MIT  
386 license.

## 387 **References**

- 388 [1] K. Bullard, F. Meier, D. Kiela, J. Pineau, and J. N. Foerster. Exploring zero-shot emergent  
389 communication in embodied multi-agent populations. *CoRR*, abs/2010.15896, 2020. URL  
390 <https://arxiv.org/abs/2010.15896>.
- 391 [2] M. Carroll, R. Shah, M. K. Ho, T. Griffiths, S. Seshia, P. Abbeel, and A. Dragan. On the utility of  
392 learning about humans for human-ai coordination. In *Advances in Neural Information Processing*  
393 *Systems*, pages 5175–5186, 2019.

- 394 [3] R. Charakorn, P. Manoonpong, and N. Dilokthanakul. Generating diverse cooperative agents  
395 by learning incompatible policies. In *The Eleventh International Conference on Learning*  
396 *Representations*, 2023. URL [https://openreview.net/forum?id=UkU05G0H7\\_6](https://openreview.net/forum?id=UkU05G0H7_6).
- 397 [4] B. Cui, A. Lupu, S. Sokota, H. Hu, D. J. Wu, and J. N. Foerster. Adversarial diversity in  
398 hanabi. In *The Eleventh International Conference on Learning Representations*, 2023. URL  
399 [https://openreview.net/forum?id=uLE3WF3-H\\_5](https://openreview.net/forum?id=uLE3WF3-H_5).
- 400 [5] T. Greitemeyer and D. O. Mügge. Video games do affect social outcomes: A meta-analytic  
401 review of the effects of violent and prosocial video game play. *Personality and Social Psychology*  
402 *Bulletin*, 40(5):578–589, 2014. doi: 10.1177/0146167213520459. URL [https://doi.org/10.](https://doi.org/10.1177/0146167213520459)  
403 [1177/0146167213520459](https://doi.org/10.1177/0146167213520459). PMID: 24458215.
- 404 [6] B. Sarkar, A. Talati, A. Shih, and S. Dorsa. Pantheonrl: A marl library for dynamic training  
405 interactions. In *Proceedings of the 36th AAAI Conference on Artificial Intelligence (Demo Track)*,  
406 2022.
- 407 [7] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov. Proximal policy optimization  
408 algorithms. *arXiv preprint arXiv:1707.06347*, 2017.
- 409 [8] C. Yu, A. Velu, E. Vinitzky, Y. Wang, A. Bayen, and Y. Wu. The surprising effectiveness of ppo  
410 in cooperative, multi-agent games. *arXiv preprint arXiv:2103.01955*, 2021.