A Cordial Sync: Going Beyond Marginal Policies for Multi-Agent Embodied Tasks

ECCV 2020 (Spotlight)

Unnat Jain^{1*}, Luca Weihs^{2*}, Eric Kolve², Ali Farhadi³, Svetlana Lazebnik¹, Aniruddha Kembhavi^{2,3}, Alexander Schwing¹ * Equal contribution by UJ and LW



Code, data, and pretrained models at: https://unnat.github.io/cordial-sync/

Continuous coordination task

1. Furniture Moving for embodied agents



MARL beyond marginal policies

2. Cordial SYNC policies



Preview of contributions

1. Furniture Moving task

2. Decentralized MARL beyond marginal policies







Jain* and Weihs* et al. "Two Body Problem: Collaborative Visual Task Completion" in CVPR 2019



Centralized MARL



Centralized MARL

Expressive but introduces issues:

Joint policy and model complexity scale exponentially Require high-bandwidth communication channel



Decentralized MARL



Decentralized MARL



Decentralized MARL

Previous methods: Single marginal policy per agent



 $\varPi = \pi^1 \otimes \pi^2 \quad \text{Rank-1}$

One policy per agent (rank-1)



0.29	0	0	0.43
0	0	0	0.06
0.03	0	0	0.05
0	0	0	0.14

Many policies per agent (high-rank)



Marginal agents



Mixture head



Generate m policies per agent



Use communication symbols



Generate mixture weights



Synchronized sampling



Select the same policy j across agents



High-Rank

 $\Pi = \sum_{j=1}^{m} \alpha_j \cdot \pi_j^1 \otimes \pi_j^2$













 $156/169 \approx 92.3\%$ of action pairs will *always* fail.

Qualitative runs

Field of view: Triangles denote field of view & orientation

Trajectories:

- Agent 1 trajectory in red
- Agent 2 trajectory in green
- TV trajectory in blue

Top-down view



Marginal Agents

Top-down view (Not available to agents)



Agent 1's view



Agent 2's view



Cordial SYNC Agents

Top-down view (Not available to agents)



Agent 1's view Agent 2's view



Cordial SYNC Agents



Quantitative results



1. Rank-1 restriction of marginal agents

	Marginal Agents					
$\Pi = \pi^1 \bigotimes \pi^2 =$ Effective Joint Policy	0.23	0	0	0.49		
	0.02	0	0	0.04		
	0.03	0	0	0.05		
	0.04	0	0	0.1		
			γ			

Rank 1



- 1. Rank-1 restriction of marginal agents
- 2. Mixture-of-marginals

Mixture-of-Marginals							
$\sum_{i=1}^{2} \alpha_i \cdot (\pi_i^1 \otimes \pi_i^2) = \begin{array}{c} \alpha_1 \cdot (\pi_1^1 \otimes \pi_1^2) \\ + \alpha_2 \cdot (\pi_2^1 \otimes \pi_2^2) \end{array} =$	0.29	0	0	0.43			
	$\begin{array}{rl} \alpha_1 \cdot (\pi_1^1 \otimes \pi_1^2) \\ + \alpha_2 \cdot (\pi_2^1 \otimes \pi_2^2) \end{array} = \end{array}$	0	0	0	0.06		
		0.03	0	0	0.05		
		0	0	0	0.14		

- 1. Rank-1 restriction of marginal agents
- 2. Mixture-of-marginals
- 3. SYNC policies



- 1. Rank-1 restriction of marginal agents
- 2. Mixture-of-marginals
- 3. SYNC policies
- 4. FurnMove task



- 1. Rank-1 restriction of marginal agents
- 2. Mixture-of-marginals
- 3. SYNC policies
- 4. FurnMove task
- 5. Qualitative results



A Cordial Sync:

Going Beyond Marginal Policies for Multi-Agent Embodied Tasks <u>https://unnat.github.io/cordial-sync/</u>

Mirrored Gridworld Agents

Interpreting Communication

