

Cross-environment Cooperation Enables Zero-shot Multi-agent Coordination

Kunal Jha

University of Washington, Seattle, Washington, United States

wilka carvalho

Harvard University, Cambridge, Massachusetts, United States

Yancheng Liang

University of Washington, Seattle, Washington, United States

Simon Du

University of Washington, Seattle, Washington, United States

Natasha Jaques

University of Washington, Seattle, Washington, United States

Max Kleiman-Weiner

University of Washington, Seattle, Washington, United States

Abstract

Zero-shot coordination (ZSC)—the ability to adapt to new partners in a cooperative task—is critical for human-compatible AI. While prior work has focused on training agents to cooperate on a single task, these specialized models fail to generalize to new tasks, even if similar. We study how reinforcement learning on a distribution of environments with a single partner induces general cooperative skills that support ZSC with many new partners on many new problems. We introduce two Jax-based procedural generators that create billions of solvable coordination challenges. We develop a new paradigm called Cross-Environment Cooperation (CEC), and show that it outperforms baselines quantitatively and qualitatively when collaborating with real people. Our findings suggest that learning to collaborate across diverse scenarios encourages agents to develop general norms effective for collaboration. Together, our results suggest a new route toward designing generalist cooperative agents that interact with humans without requiring human data.