Learning Flexible Heterogeneous Coordination With Capability-Aware Shared Hypernetworks†

Extended Abstract

Kevin Fu
Georgia Institute of Technology
Atlanta, United States
kfu38@gatech.edu

Pierce Howell*
Georgia Institute of Technology
Atlanta, United States
pierce.howell@gatech.edu

Shalin Jain*
Georgia Institute of Technology
Atlanta, United States
sjain441@gatech.edu

Harish Ravichandar
Georgia Institute of Technology
Atlanta, United States
harish.ravichandar@cc.gatech.edu

ABSTRACT

Cooperative heterogeneous multi-agent tasks require agents to effectively coordinate their behaviors while accounting for their relative capabilities. Learning-based solutions span between two extremes with opposing tradeoffs: i) shared-parameter solutions, which encode diverse behaviors within a single architecture by assigning an ID to each agent, are sample-efficient but result in limited behavioral diversity; ii) independent solutions, which learn a separate policy for each agent, show greater behavioral diversity but lack sample-efficiency. Prior work has also explored *selective* parameter-sharing, allowing for a compromise between diversity and efficiency. None of these approaches, however, effectively generalize to unseen agents or teams. We present Capability-Aware Shared Hypernetworks (CASH), a novel architecture for heterogeneous multi-agent coordination that generates behavioral diversity while staying sample-efficient via *soft parameter-sharing* hypernetworks. CASH allows the team to learn common strategies using a shared encoder, which are then *adapted* according to the team's capabilities with a hypernetwork, allowing for *zero-shot generalization* to unseen teams and agents. We conduct experiments across two heterogeneous coordination tasks and three learning paradigms (imitation learning, on- and off-policy reinforcement learning). Results show that CASH outperforms baseline architectures in success rate and sample efficiency when evaluated on unseen teams and agents despite using 60-80% fewer learnable parameters.

KEYWORDS

Heterogeneity; Multi-Agent Learning; Coordination

ACM Reference Format:
Kevin Fu, Pierce Howell*, Shalin Jain*, and Harish Ravichandar. 2025. Learning Flexible Heterogeneous Coordination With Capability-Aware Shared Hypernetworks†: Extended Abstract. In *Proc. of the 24th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2025), Detroit, Michigan, USA, May 19 – 23, 2025*, IFAAMAS, 3 pages.

Proc. of the 24th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2025), Y. Vorobeychik, S. Das, A. Nowé (eds.), May 19 – 23, 2025, Detroit, Michigan, USA. © 2025 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org).

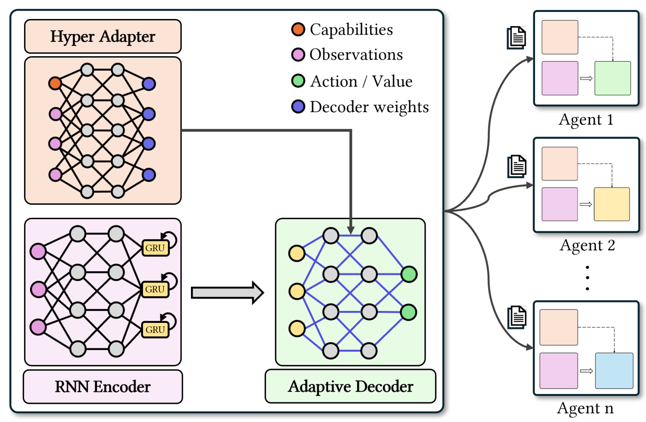

Figure 1: We present Capability-Aware Shared Hypernetworks (CASH), an architecture for flexible and decentralized heterogeneous multi-agent coordination. CASH uses hypernetworks to achieve *soft parameter sharing*, and conditions agent behavior on individual and collective capabilities to generate sufficient behavioral diversity. Our implementation can be found at https://github.com/GT-STAR-Lab/CASH.

1 INTRODUCTION

Consider a firefighting scenario where a team of heterogeneous robots is tasked with putting out several wildfires in a large forest. Robots in this team have different capabilities (e.g. speed, firefighting capacity) and thus must reason about their collective capabilities in order to effectively cooperate. Suppose we do not know which specific agents are available until runtime, or want to account for individual capabilities changing during runtime. This motivates the need for a single policy to readily adapt to changing team compositions and capabilities, without requiring retraining. Moreover, such scenarios require a decentralized solution since centralized coordination might not be possible.

*Equal Contribution
†This work was supported in part by the Army Research Lab under Grant W911NF-17-2-0181 (DCIST CRA).

Learning-based approaches to this decentralized heterogeneous coordination problem span between two extremes. The dominant approach has been to append a unique ID to each individual agent [10, 14, 21, 23], preserving the sample-efficiency and scalability benefits of shared-parameter multi-agent policy learning [8, 20] while enabling agents to exhibit diverse behavior. The opposite extreme is to learn an *independent* policy for each agent, improving diversity in agent behavior but sacrificing efficiency and scalability [4, 5, 12]. *Selective* parameter sharing, in which subsets of agents are given unique IDs, has proven to be a useful compromise between the two extremes [7, 19, 22]. However, none of these approaches reason over the impact of agent capabilities on behavior, and thus cannot effectively generalize to unseen agents or teams.

In this work, we present Capability-Aware Shared Hypernetworks (CASH), a novel middle-ground approach to decentralized heterogeneous coordination: *soft weight sharing* using hypernetworks [9]. Motivated by robotics, we are interested in efficiently learning decentralized coordination strategies that can *zero-shot generalize* to changes in team composition and capabilities.

2 CAPABILITY-AWARE SHARED HYPERNETWORKS

Inspired by prior work in trait-based task allocation, we start with the assumption that agent capabilities can be represented with a single vector of scalars [10, 15]. To best adapt agent behavior to capabilities, we employ hypernetworks, or networks which generate the parameters of another “target” network [9]. Hypernetworks can be seen as *soft weight sharing* [6, 9], as one hypernetwork can approximate an ensemble of individually trained target networks, and indeed there are multiple recent examples of hypernetworks flexibly adapting network behavior for continual learning [11], transfer learning [16], and meta-RL [2, 3].

CASH has three components: the RNN Encoder, the Hyper Adapter, and the Adaptive Decoder (Fig. 1). First, the RNN Encoder generates a latent embedding from observations. Then, the Hyper Adapter, conditioned on both local observations and individual and team capabilities, produces the parameters of the Adaptive Decoder. Empirically, we found that adding LayerNorm [1] before each activation in the Hyper Adapter was critical for stabilizing training. Finally, the Adaptive Decoder, given the latent embedding of the RNN Encoder, produces an action (or value, depending on the learning paradigm). Since the RNN Encoder and Hyper Adapter are both fully shared-parameter, CASH retains many of the efficiency benefits of shared-parameter approaches. However, the Adaptive Decoder’s parameters are dynamically modified for each agent at each timestep, enabling diverse and adaptive behaviors.

3 RESULTS

We evaluate CASH on two novel heterogeneous coordination tasks implemented with JaxMARL’s MPE [13, 18] (`Firefighting` and `Transport`) and three learning paradigms: DAgger [17], representing imitation learning, MAPPO [23], for on-policy RL, and QMIX [14], for off-policy RL. We compare CASH against two shared-parameter baseline approaches that do not leverage hypernetworks: RNN-IMP, which does not have access to capabilities and must infer them from observations alone, and RNN-EXP, where capability

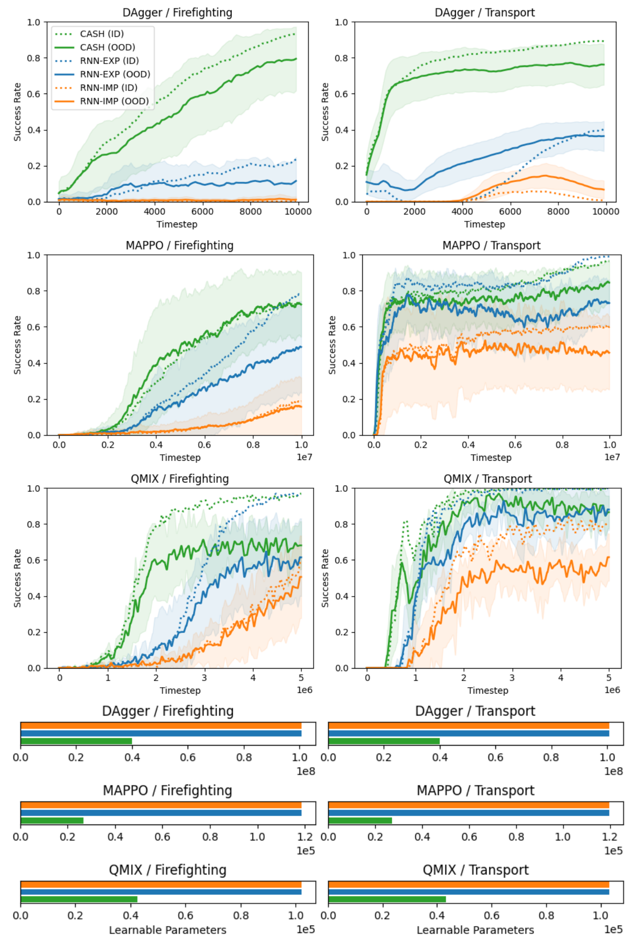

Figure 2: Success rates (↑, top) and learnable parameters (↓, bottom) across two tasks and three learning paradigms. Solid curves/shaded area are mean/stdev when evaluated on unseen teams & agents (OOD), while dotted curves are when evaluated on training teams (ID). CASH outperforms both RNN-IMP and RNN-EXP on *unseen* teams & agents with 60-80% fewer learnable parameters, even when RNN-EXP matches CASH’s performance on *training* teams. Results are with 10 seeds for QMIX/MAPPO, 3 for DAgger.

vectors are explicitly concatenated to inputs. We do not compare against independent learning approaches as it is not possible to apply those methods to unseen agents.

Despite having 60%-80% fewer learnable parameters to work with, CASH outperforms the baselines in both final success rate and sample efficiency when generalizing to unseen team compositions and out-of-distribution capabilities, even when CASH and RNN-EXP both perform well on training teams (Fig. 2). Our results show that Capability-Aware Shared Hypernetworks are a powerful architecture for sample-efficient heterogeneous multi-agent learning. This is an extended abstract; for more detailed methodology and results, please see our full paper at http://arxiv.org/abs/2501.06058.

REFERENCES

[1] Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E. Hinton. 2016. Layer Normalization. arXiv:1607.06450 [stat.ML] https://arxiv.org/abs/1607.06450

[2] Jacob Beck, Matthew Thomas Jackson, Risto Vuorio, and Shimon Whiteson. 2023. Hypernetworks in Meta-Reinforcement Learning. 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, 1478–1487. https://proceedings.mlr.press/v205/beck23a.html

[3] Jacob Beck, Risto Vuorio, Zheng Xiong, and Shimon Whiteson. 2023. Recurrent Hypernetworks are Surprisingly Strong in Meta-RL. In *Advances in Neural Information Processing Systems*, A. Oh, T. Naumann, A. Globerson, K. Saenko, M. Hardt, and S. Levine (Eds.), Vol. 36. Curran Associates, Inc., 62121–62138. https://proceedings.neurips.cc/paper_files/paper/2023/file/c3fa3a7d50b34732c6d08f6f66380d75-Paper-Conference.pdf

[4] Matteo Bettini, Ajay Shankar, and Amanda Prorok. 2023. Heterogeneous Multi-Robot Reinforcement Learning. In *Proceedings of the 2023 International Conference on Autonomous Agents and Multiagent Systems (AAMAS '23)*.

[5] Matteo Bettini, Ajay Shankar, and Amanda Prorok. 2023. System Neural Diversity: Measuring Behavioral Heterogeneity in Multi-Agent Learning. arXiv:2305.02128 (May 2023). https://doi.org/10.48550/arXiv.2305.02128

[6] Vinod Kumar Chauhan, Jiandong Zhou, Ping Lu, Soheila Molaei, and David A. Clifton. 2023. A Brief Review of Hypernetworks in Deep Learning. arXiv:2306.06955 (Aug. 2023). https://doi.org/10.48550/arXiv.2306.06955

[7] Filippos Christianos, Georgios Papoudakis, Muhammad A Rahman, and Stefano V Albrecht. 2021. Scaling Multi-Agent Reinforcement Learning with Selective Parameter Sharing. In *Proceedings of the 38th International Conference on Machine Learning (Proceedings of Machine Learning Research, Vol. 139)*. PMLR.

[8] Jayesh K Gupta, Maxim Egorov, and Mykel Kochenderfer. 2017. Cooperative multi-agent control using deep reinforcement learning. In *Autonomous Agents and Multiagent Systems: AAMAS 2017 Workshops, Best Papers, São Paulo, Brazil, May 8-12, 2017, Revised Selected Papers 16*. Springer, 66–83.

[9] David Ha, Andrew Dai, and Quoc V. Le. 2016. HyperNetworks. arXiv:1609.09106 (Dec. 2016). https://doi.org/10.48550/arXiv.1609.09106

[10] Pierce Howell, Max Rudolph, Reza Joseph Torbati, Kevin Fu, and Harish Ravichandar. 2023. Generalization of Heterogeneous Multi-Robot Policies via Awareness and Communication of Capabilities. In *Proceedings of The 7th Conference on Robot Learning (Proceedings of Machine Learning Research, Vol. 229)*, Jie Tan, Marc Toussaint, and Kourosh Darvish (Eds.). PMLR, 2772–2790. https://proceedings.mlr.press/v229/howell23a.html

[11] Yizhou Huang, Kevin Xie, Homanga Bharadhwaj, and Florian Shkurti. 2021. Continual Model-Based Reinforcement Learning with Hypernetworks. In *2021 IEEE International Conference on Robotics and Automation (ICRA)*. 799–805. https://doi.org/10.1109/ICRA48506.2021.9560793

[12] 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*, M. Ranzato, A. Beygelzimer, Y. Dauphin, P.S. Liang, and J. Wortman Vaughan (Eds.), Vol. 34. Curran Associates, Inc., 3991–4002. https://proceedings.neurips.cc/paper_files/paper/2021/file/20aee3a5f4643755a79ee5f6a73050ac-Paper.pdf

[13] Ryan Lowe, YI WU, Aviv Tamar, Jean Harb, OpenAI Pieter Abbeel, and Igor Mordatch. 2017. Multi-Agent Actor-Critic for Mixed Cooperative-Competitive Environments. In *Advances in Neural Information Processing Systems*, I. Guyon, U. Von Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett (Eds.), Vol. 30. Curran Associates, Inc. https://proceedings.neurips.cc/paper_files/paper/2017/file/68a9750337a418a86fe06c1991a1d64c-Paper.pdf

[14] 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.

[15] Harish Ravichandar, Kenneth Shaw, and Sonia Chernova. 2020. STRATA: unified framework for task assignments in large teams of heterogeneous agents. *Autonomous Agents and Multi-Agent Systems* 34 (2020), 1–25.

[16] Sahand Rezaei-Shoshtari, Charlotte Morissette, Francois R. Hogan, Gregory Dudek, and David Meger. 2023. Hypernetworks for zero-shot transfer in reinforcement learning. In *Proceedings of the Thirty-Seventh AAAI Conference on Artificial Intelligence and Thirty-Fifth Conference on Innovative Applications of Artificial Intelligence and Thirteenth Symposium on Educational Advances in Artificial Intelligence (AAAI'23/IAAI'23/EAAI'23)*. AAAI Press, Article 1077, 9 pages. https://doi.org/10.1609/aaai.v37i8.26146

[17] Stéphane Ross and Drew Bagnell. 2010. Efficient reductions for imitation learning. In *Proceedings of the thirteenth international conference on artificial intelligence and statistics*. JMLR Workshop and Conference Proceedings, 661–668.

[18] Alexander Rutherford, Benjamin Ellis, Matteo Gallici, Jonathan Cook, Andrei Lupu, Gardar Ingvarsson, Timon Willi, Akbir Khan, Christian Schroeder de Witt, Alexandra Souly, Saptarashmi Bandyopadhyay, Mikayel Samvelyan, Minqi Jiang, Robert Tjarko Lange, Shimon Whiteson, Bruno Lacerda, Nick Hawes, Tim Rocktaschel, Chris Lu, and Jakob Nicolaus Foerster. 2023. JaxMARL: Multi-Agent RL Environments in JAX. arXiv:2311.10090 (Dec. 2023). https://doi.org/10.48550/arXiv.2311.10090 arXiv:2311.10090 [cs].

[19] Esmaeil Seraj, Zheyuan Wang, Rohan Paleja, Daniel Martin, Matthew Sklar, Anirudh Patel, and Matthew Gombolay. 2022. Learning Efficient Diverse Communication for Cooperative Heterogeneous Teaming. 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, 1173–1182.

[20] Ming Tan. 1993. Multi-agent reinforcement learning: Independent vs. cooperative agents. In *Proceedings of the tenth international conference on machine learning*.

[21] J. K. Terry, Nathaniel Grammel, Sanghyun Son, Benjamin Black, and Aakriti Agrawal. 2023. Revisiting Parameter Sharing in Multi-Agent Deep Reinforcement Learning. arXiv:2005.13625 [cs.LG] https://arxiv.org/abs/2005.13625

[22] Ceyer Wakilpoor, Patrick J. Martin, Carrie Rebhuhn, and Amanda Vu. 2020. Heterogeneous Multi-Agent Reinforcement Learning for Unknown Environment Mapping. arXiv:2010.02663 [cs.MA] https://arxiv.org/abs/2010.02663

[23] Chao Yu, Akash Velu, Eugene Vinitsky, Jiaxuan Gao, Yu Wang, Alexandre Bayen, and YI WU. 2022. The Surprising Effectiveness of PPO in Cooperative Multi-Agent Games. In *Advances in Neural Information Processing Systems*, S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh (Eds.), Vol. 35. Curran Associates, Inc., 24611–24624. https://proceedings.neurips.cc/paper_files/paper/2022/file/9c1535a02f0ce079433344e14d910597-Paper-Datasets_and_Benchmarks.pdf