
Ad Hoc Teamwork in the Presence of Adversaries

Ted Fujimoto¹ Samrat Chatterjee¹ Auroop Ganguly^{2,1}

Abstract

Advances in ad hoc teamwork have the potential to create agents that collaborate robustly in real-world applications. Agents deployed in the real world, however, are vulnerable to adversaries with the intent to subvert them. There has been little research in ad hoc teamwork that assumes the presence of adversaries. We explain the importance of extending ad hoc teamwork to include the presence of adversaries and clarify why this problem is difficult. We then propose some directions for new research opportunities in ad hoc teamwork that leads to more robust multi-agent cyber-physical infrastructure systems.

1. Introduction

Ad hoc teamwork (AHT) is a step towards general multi-agent intelligence where decision-making agents (human and computational) are trained to collaborate with other previously unseen agents. That is, AHT assumes that each agent will learn to collaborate without prior coordination or direct control of the other agents (Mirsky et al., 2022). Although the challenge of AHT was first formalized by Stone et al. (2010), there has been a recent resurgence of AHT research that incorporates advances in deep reinforcement learning (RL).

Established works have focused on multi-agent consensus among a network of agents typically with centralized training and decentralized execution, where a fixed set of agents are known to all (Lowe et al., 2017; Rashid et al., 2018). For example, this may be useful in real-world infrastructure protection settings where a team of collaborative defense agents might be tasked with securing an asset of interest. This might not be practical, however, in dynamic, adversarial, and contested environments. In these settings, new agents may need to enter the environment and start competing or

collaborating without access to a centralized policy. Even if we assume such a policy, it would likely lead to brittle agents that are tethered to their training partners for optimal performance (Li et al., 2019). This leads the group to be unable to perform optimally in the presence of adversaries. One appeal of AHT is the potential for multiple people or organizations to train agents that collaborate with each other without the need to retrain with previously unseen agents. Outside of Shafipour & Fallah (2021), AHT in the presence of adversaries has not been explored due to the assumption of full cooperation between agents. Adversarial machine learning (ML) researchers have an opportunity to chart the next path of truly robust multi-agent learning. In this paper, we contribute (1) a summary of past research in AHT, (2) why it could be a necessary next step in high-stakes, real-world multi-agent RL, (3) a list of some challenges in AHT in the presence of adversaries, and (4) a proposal for some new directions for researchers in both adversarial ML and AHT.

2. Ad Hoc Teamwork

2.1. Background

AHT is a subfield of cooperative AI, which is defined as AI methods that help individuals, humans and machines, to find ways that improve joint welfare (Dafoe et al., 2020). As mentioned previously, AHT methods train agents to collaborate without prior coordination or direct control of other agents. In this paper, AHT will mostly be synonymous with RL-based methods in AHT, called zero-shot coordination (ZSC). The main reason is because ZSC has enjoyed some recent success. These successes include a general and differentiable information theoretic objective for training a diverse population of optimal policies (Lupu et al., 2021), training methods that prevent agents from exchanging information through arbitrary conventions (Hu et al., 2021), and graph neural networks that learn agent models and joint-action value models under varying team compositions (Rahman et al., 2021). Some approaches even outperform self-play and behavioral cloning when collaborating with humans in video games (Strouse et al., 2021). ZSC is an evolving field and there is currently some disagreement over how to precisely formulate it (Treutlein et al., 2021). We recommend Mirsky et al. (2022) as an up-to-date survey of AHT.

¹Pacific Northwest National Laboratory, WA, USA

²Northeastern University, MA, USA. Correspondence to: Ted Fujimoto <ted.fujimoto@pnnl.gov>.

2.2. Potential Applications

While AHT has shown some success in robot soccer (Genter et al., 2017), there are still many potential applications yet to be explored. Ideally, AHT can perform more robustly than centralized multi-agent RL methods, so any multi-agent environment can potentially train an AHT application. In the context of secure cyber-physical infrastructure operations (such as transportation and energy distribution), multi-agent learning and control is essential for achieving system resilience goals (Januário et al., 2019; Hou et al., 2021; Phan et al., 2021). Some future AHT applications worth mentioning include important automated infrastructure, like energy and power distribution (Biagioni et al., 2021; Pigott et al., 2021; Wang et al., 2021). AHT would be particularly useful here because of the possibility that different buildings or power plants might need to coordinate with previously unseen agents from other allied power plants. Another application would be human-automated vehicle interaction. It is still difficult for automated cars to learn and appropriately react to humans while driving (Dommès et al., 2021).

3. Potential Challenges with Adversaries in Ad Hoc Teamwork

We now briefly describe some challenges that agents trained using AHT methods could face unless it is framed as an adversarial ML problem.

3.1. Past Research in Adversarial RL

It is possible that some attacks described in past research on adversarial RL (Huang et al., 2017; Gleave et al., 2019; Lin et al., 2020; Sun et al., 2020; Figura et al., 2021; Fujimoto et al., 2021) will have similar negative effects in AHT. Given that there are currently no universally effective defenses against adversarial attacks in supervised learning (Athalye et al., 2018; Short et al., 2019), the same is likely true for RL.

3.2. Balancing the Presence of Allies and Adversaries

Even if we assume adversaries might exist in the environment, it is not clear how to train the agents to collaborate. It might be that the adversary is secretly controlling the policy of one of the agents. We would need new methods that assume some of the training agent’s partners might not be aligned with the objectives of the group.

3.3. Methods of Evaluation

One aspect of AHT that might be overlooked is the evaluation of the agents at test time. An AHT agent might have to act differently in the presence of an adversary instead of a low-skill partner. In general, we would need AHT meth-

ods that are more risk-aware and define what success looks like when there is a high probability that some of the other agents will try to subvert the group’s goal.

4. Future Research Directions

AHT was established with robust multi-agent behavior as the main goal. Some past works in AHT have made steps in that direction. Type-based methods, like in Rahman et al. (2021), attempt to infer the agent’s behavior from its actions. Agents can also use communication to help each other quickly adapt to new scenarios (Macke et al., 2021), which has been successful in defending against attacks in multi-agent deep learning (Tu et al., 2021). Extending past work in AHT can show us which methods have the most potential in being robust against adversaries. Past works in adversarial training in RL (Pinto et al., 2017; Pan et al., 2019) show that robustness can be achieved if the agent is trained with adversarial examples.

Extending past methods alone, however, will likely not be enough. Introducing new agents to other previously trained agents exacerbate the brittleness of current RL methods (Gleave et al., 2019). The success of adversarial AHT, and AHT in general, will depend on a foundation of risk-awareness for multi-agent RL in the presence of adversaries. This foundation needs to account for the “openness” of the environment, which models how new agents enter the group. The presence of adversaries compels us to carefully understand how openness impacts the training of an AHT agent. Openness that includes adversarial agents will require researchers to assume additional formal descriptions of the environment, which include (1) the distribution of new adversarial agents to introduce to the group, (2) the diversity of the adversaries, and (3) the reward function that accounts for the number and types of adversaries. As multi-agent RL is starting to be applied to more real-world tasks (Wang et al., 2021; Bae & Koumoutsakos, 2022), there is now more motivation to train agents that are robust and risk-aware. This creates an incentive to train AHT agents in high-stakes environments. An adversarial ML perspective can make these endeavors more successful.

5. Conclusion

As RL continues to be pursued to accomplish complex, real-world tasks, we will inevitably need to train RL agents that robustly cooperate with others. Since AHT seems to be a promising path toward that direction, we propose accounting for the possible presence of adversaries as a necessary component of that field. We put forward that adversarial ML researchers will be critical to ensuring AHT will advance safe collaboration between both humans and RL agents.

References

- Athalye, A., Carlini, N., and Wagner, D. Obfuscated gradients give a false sense of security: Circumventing defenses to adversarial examples. In *International conference on machine learning*, pp. 274–283. PMLR, 2018.
- Bae, H. J. and Koumoutsakos, P. Scientific multi-agent reinforcement learning for wall-models of turbulent flows. *Nature Communications*, 13(1):1–9, 2022.
- Biagioni, D., Zhang, X., Wald, D., Vaidhyanathan, D., Chintala, R., King, J., and Zamzam, A. S. Powergridworld: A framework for multi-agent reinforcement learning in power systems. *arXiv preprint arXiv:2111.05969*, 2021.
- Dafoe, A., Hughes, E., Bachrach, Y., Collins, T., McKee, K. R., Leibo, J. Z., Larson, K., and Graepel, T. Open problems in cooperative ai. *arXiv preprint arXiv:2012.08630*, 2020.
- Dommès, A., Merlhiot, G., Lobjois, R., Dang, N.-T., Vienne, F., Boulo, J., Oliver, A.-H., Cretual, A., and Cavallo, V. Young and older adult pedestrians’ behavior when crossing a street in front of conventional and self-driving cars. *Accident Analysis & Prevention*, 159:106256, 2021.
- Figura, M., Kosaraju, K. C., and Gupta, V. Adversarial attacks in consensus-based multi-agent reinforcement learning. In *2021 American Control Conference (ACC)*, pp. 3050–3055. IEEE, 2021.
- Fujimoto, T., Doster, T., Attarian, A., Brandenberger, J., and Hodas, N. Reward-free attacks in multi-agent reinforcement learning. *arXiv preprint arXiv:2112.00940*, 2021.
- Genter, K., Laue, T., and Stone, P. Three years of the robocup standard platform league drop-in player competition. *Autonomous Agents and Multi-Agent Systems*, 31(4):790–820, 2017.
- Gleave, A., Dennis, M., Wild, C., Kant, N., Levine, S., and Russell, S. Adversarial policies: Attacking deep reinforcement learning. In *International Conference on Learning Representations*, 2019.
- Hou, J., Wang, F., Wang, L., and Chen, Z. Reinforcement learning based multi-agent resilient control: From deep neural networks to an adaptive law. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pp. 7737–7745, 2021.
- Hu, H., Lerer, A., Cui, B., Pineda, L., Brown, N., and Foerster, J. Off-belief learning. In *International Conference on Machine Learning*, pp. 4369–4379. PMLR, 2021.
- Huang, S., Papernot, N., Goodfellow, I., Duan, Y., and Abbeel, P. Adversarial attacks on neural network policies. *arXiv preprint arXiv:1702.02284*, 2017.
- Januário, F., Cardoso, A., and Gil, P. A distributed multi-agent framework for resilience enhancement in cyber-physical systems. *IEEE Access*, 7:31342–31357, 2019.
- Li, S., Wu, Y., Cui, X., Dong, H., Fang, F., and Russell, S. Robust multi-agent reinforcement learning via minimax deep deterministic policy gradient. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pp. 4213–4220, 2019.
- Lin, J., Dzevaroska, K., Zhang, S. Q., Leon-Garcia, A., and Papernot, N. On the robustness of cooperative multi-agent reinforcement learning. In *2020 IEEE Security and Privacy Workshops (SPW)*, pp. 62–68. IEEE, 2020.
- Lowe, R., Wu, Y. I., Tamar, A., Harb, J., Pieter Abbeel, O., and Mordatch, I. Multi-agent actor-critic for mixed cooperative-competitive environments. *Advances in neural information processing systems*, 30, 2017.
- Lupu, A., Cui, B., Hu, H., and Foerster, J. Trajectory diversity for zero-shot coordination. In *International Conference on Machine Learning*, pp. 7204–7213. PMLR, 2021.
- Macke, W., Mirsky, R., and Stone, P. Expected value of communication for planning in ad hoc teamwork. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pp. 11290–11298, 2021.
- Mirsky, R., Carlucho, I., Rahman, A., Fosong, E., Macke, W., Sridharan, M., Stone, P., and Albrecht, S. V. A survey of ad hoc teamwork: Definitions, methods, and open problems. *arXiv preprint arXiv:2202.10450*, 2022.
- Pan, X., Seita, D., Gao, Y., and Canny, J. Risk averse robust adversarial reinforcement learning. In *2019 International Conference on Robotics and Automation (ICRA)*, pp. 8522–8528. IEEE, 2019.
- Phan, T., Belzner, L., Gabor, T., Sedlmeier, A., Ritz, F., and Linnhoff-Popien, C. Resilient multi-agent reinforcement learning with adversarial value decomposition. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pp. 11308–11316, 2021.
- Pigott, A., Crozier, C., Baker, K., and Nagy, Z. Gridlearn: Multiagent reinforcement learning for grid-aware building energy management. *arXiv preprint arXiv:2110.06396*, 2021.
- Pinto, L., Davidson, J., Sukthankar, R., and Gupta, A. Robust adversarial reinforcement learning. In *International Conference on Machine Learning*, pp. 2817–2826. PMLR, 2017.

- Rahman, M. A., Hopner, N., Christianos, F., and Albrecht, S. V. Towards open ad hoc teamwork using graph-based policy learning. In *International Conference on Machine Learning*, pp. 8776–8786. PMLR, 2021.
- Rashid, T., Samvelyan, M., Schroeder, C., Farquhar, G., Foerster, J., and Whiteson, S. Qmix: Monotonic value function factorisation for deep multi-agent reinforcement learning. In *International Conference on Machine Learning*, pp. 4295–4304. PMLR, 2018.
- Shafipour, E. and Fallah, S. Task-based ad-hoc teamwork with adversary. In *Annual Conference Towards Autonomous Robotic Systems*, pp. 76–87. Springer, 2021.
- Short, A., La Pay, T., and Gandhi, A. Defending against adversarial examples. Technical report, Sandia National Lab.(SNL-NM), Albuquerque, NM (United States), 2019.
- Stone, P., Kaminka, G. A., Kraus, S., and Rosenschein, J. S. Ad hoc autonomous agent teams: Collaboration without pre-coordination. In *Twenty-Fourth AAAI Conference on Artificial Intelligence*, 2010.
- Strouse, D., McKee, K., Botvinick, M., Hughes, E., and Everett, R. Collaborating with humans without human data. *Advances in Neural Information Processing Systems*, 34, 2021.
- Sun, J., Zhang, T., Xie, X., Ma, L., Zheng, Y., Chen, K., and Liu, Y. Stealthy and efficient adversarial attacks against deep reinforcement learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pp. 5883–5891, 2020.
- Treutlein, J., Dennis, M., Oesterheld, C., and Foerster, J. A new formalism, method and open issues for zero-shot coordination. In *International Conference on Machine Learning*, pp. 10413–10423. PMLR, 2021.
- Tu, J., Wang, T., Wang, J., Manivasagam, S., Ren, M., and Urtasun, R. Adversarial attacks on multi-agent communication. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 7768–7777, 2021.
- Wang, J., Xu, W., Gu, Y., Song, W., and Green, T. C. Multi-agent reinforcement learning for active voltage control on power distribution networks. *Advances in Neural Information Processing Systems*, 34:3271–3284, 2021.