Skill-Aligned Fairness in Multi-Agent Learning for Collaboration in Healthcare

Fairness in multi-agent reinforcement learning (MARL) is often framed as a workload balance problem, where fairness is defined as distributing an approximately equal number of subtasks or equal effort across agents, regardless of their skill levels. However, this overlooks the expertise of agents and the structured coordination required in real-world domains. In healthcare, equitable task allocation requires workload balance or expertise alignment to prevent burnout and overuse of highly skilled agents. To better account for these issues, we propose FairSkillMARL, a framework that defines fairness as the dual objective of workload balance and skill-task alignment. Prior work has quantified fairness as reward equality in social dilemmas [1], worst-case performance guarantees in traffic scheduling [2], balanced throughput in network control [3,4], equitable effort distribution in cooperative navigation tasks [5], cumulative reward parity in hierarchical learning [6] and demographic parity in resource allocation [7]. However, these definitions do not jointly account for the skill-task alignment and overutilization of agents who are repeatedly tasked with the most demanding actions, which is of relevance to healthcare applications. More broadly, the intersection of MARL, healthcare, and fairness is an understudied area.

Our approach includes designing two reward functions: 1) Workload balance-based reward (R1) and 2) FairSkillMARL framework, where the reward function incorporates both workload balance and skill alignment (R2), where **α** in [0,1] is a tunable parameter specifying the trade-off between workload balance and expertise alignment. The **workload imbalance** is measured using the Gini Index, which captures how unevenly subtasks are distributed among agents, where 0 indicates perfect equality and 1 indicates complete inequality. The **skill-task misalignment** is measured as the proportion of subtasks completed relative to that agent's skill levels in each subgoal. In other words, it measures the deviation from optimal task-to-skill assignments.

We conducted experiments comparing the contributions of three heterogeneous agents with varying skill levels in a medical-inspired resuscitation task, utilising the QMIX MARL algorithm with R1, and with FairSkillMARL reward with R2. For evaluation, we compare this to FEN, a framework for fairness in MARL, which measures agent resource utilization and penalizes agents when they deviate from the average utilization of all agents. Our findings reveal that neither pure workload balancing nor pure skill alignment produces the best coordination. Instead, intermediate balance points consistently outperform both extremes, indicating that fairness in MARL is inherently a multi-objective problem rather than one that can be reduced to a single dimension. At the same time, we identified the need for FairSkillMARL to be tuned as very high penalties can impede the agents' ability to learn the task and reach the goal. Additionally, Skill-task alignment and workload balance can conflict due to the framing of the linear combination of both fairness dimensions might be overly simplistic. Hence, we propose an adaptive fairness scheduler that is low initially, allowing the agents to learn the task and then increased progressively along the episode. We advocate for fairness in MARL to be expanded to consider the context of the application domain and heterogeneity of the agents. Further work will include more MARL baselines and include optimal parameter tuning of the fairness penalty to balance task efficiency and fairness.

References:

1. Hughes, E.; Leibo, J. Z.; Phillips, M. G.; Tuyls, K.; Dueñez-Guzmán, E. A.; Castañeda, A. G.; Dunning, I.; Zhu, T.; McKee, K. R.; Koster, R.; Roff, H.; and Graepel, T. 2018. Inequity aversion improves cooperation in intertemporal social dilemmas.
2. Yuan, M.; Cao, Q.; Pun, M.-O.; and Chen, Y. 2021. Multiagent reinforcement learning-based fairness-aware scheduling for bursty traffic. In 2021 IEEE Global Communications Conference (GLOBECOM), 1–6. IEEE.
3. Fang, W.; Zhao, X.; and Zhang, C. 2024. Fairness-aware multi-agent reinforcement learning and visual perception for adaptive traffic signal control. Optoelectronics Letters, 20(12): 764–768.
4. Huang, Z., Wu, T., and Boulet, B. (2022). Fairness-Aware Model-Based Multi-Agent Reinforcement Learning for Traffic Signal Control.
5. Aloor, J. J.; Nayak, S.; Dolan, S.; and Balakrishnan, H. 2024. Cooperation and Fairness in Multi-Agent Reinforcement Learning.
6. Jiang, J.; and Lu, Z. 2019b. Learning Fairness in MultiAgent Systems. In Advances in Neural Information Processing Systems, volume 32. Curran Associates, Inc.
7. Malfa, G.; Zhang, J.; Luck, M.; and Black, E. 2025. Fairness Aware Reinforcement Learning via Proximal Policy Optimization.