Research Statement

Qi Deng, Email: dengqi22@std.uestc.edu.cn

I am a third-year master's student at the University of Electronic Science and Technology of China supervised by Prof. Lijun Wu, majoring in computer technology. My research focuses on reinforcement learning, multimodal machine learning, game theory, and LLM-based agents, which are the tangible expressions of my passion for realizing artificial general intelligence. During my postgraduate studies, I have systematically studied relevant courses in the field of artificial intelligence, actively engaged in research, and published related articles. I possess a deep understanding and practical experience with various classical multimodal machine learning and reinforcement learning algorithms. All my research is implemented in Python with PyTorch on Linux, and I have solid skills in both quantitative and qualitative analysis. I am capable of using statistical methods and machine learning techniques to process and interpret complex data, and accurately model problems.

Past and Current Research

In my research on multimodality, published at the esteemed international conference IJCNN 2024, I introduced a novel hierarchical fusion training framework for multimodal dialogue response generation. Two analogous tasks have emerged in multimodal dialogue research: multimodal dialogue response generation and multimodal task-oriented dialogue. Both tasks share the goal of response multi-round, interactive content based on the multimodal dialogue history, but the latter focuses on accomplishing specific objectives which can be viewed as the former fine-tuned. The fine-tuning strategy may cause catastrophic forgetting and overfitting on few well-annotated data. Despite considerable progress in both areas, many existing works rely on retrieval-based approaches and additional auxiliary knowledge bases. To address these issues, I propose a Hierarchical Fusion Framework (HFF) for multimodal dialogue response generation. HFF blends these two tasks to learn a generation model from a data-driven perspective by introducing multi-dataset learning scheme, achieving a balance between generalization and expertise. In this work, multi-dataset learning is cast as a multi-objective optimization problem due to potential conflicts between datasets, necessitating a trade-off based on data distribution during training. Hierarchical fusion is performed sequentially between modalities and datasets, which could efficiently establish clear cross-modal relationships and integrate knowledge from multi-dataset. Specifically, HFF aligns extracted unimodal features (image and text) before fusing them through cross-modal attention and integrates them into multimodal encoder-decoder for generating responses. By optimizing the fusion between corpora from multi-dataset as conflicting objectives to satisfy Pareto optimality, my approach effectively facilitates both multimodal task-oriented and task-unoriented dialogues.

Another research endeavor primarily explores the optimization of adaptive traffic signal control

using reinforcement learning, which is also the topic of my master's thesis. Recently, Adaptive Traffic Signal Control (ATSC) has become a main research direction in intelligent transportation systems to improve transportation efficiency and mitigate traffic congestion at signalized intersections. With multi-agent deep reinforcement learning being extensively employed for multi-intersection traffic signal control, ATSC is commonly considered as a fully cooperative problem between agents, presupposing that all agents are committed to achieving a collective optimal solution. However, such altruistic cooperation is often challenging to implement in practice. Additionally, as the number of agents escalates, issues like the curse of dimensionality and environmental non-stationarity emerge, complicating the learning process. I present a novel perspective, framing ATSC as a competitive-cooperative game trade-off scenario, and design a multi-agent Neighborhood Coordinated and Holistic Optimized Actor-Critic framework (NcHo-AC) to address this challenge. NcHo-AC firstly introduces a novel traffic state representation that augments the representable information for single intersection. Secondly, graph attention networks and multi-head attention mechanism are utilized to effectively focus on critical local regions, enabling neighborhood cooperation and communication. Thirdly, the multi-agent proximal policy optimization algorithm is employed to search solution space. Crucially, I utilize centralized critic network leveraging mean field theory to estimate the separate reward of each agent, rather than the global reward of the entire group. This design not only facilitates the learning of the desired Nash equilibrium but also helps to mitigate the noise induced by the exploratory behaviors of agents, thereby alleviating the non-stationarity issue.

Future Research

During my previous study and research experience, I was immersed in the work at hand and excited by the challenges I solved. These endeavors ignited an unquenchable passion for academia and strengthened my determination to explore the field of cutting-edge artificial intelligence agents. Looking ahead, my research vision is to integrate reinforcement learning, multimodal machine learning, and game theory, aiming to develop artificial intelligence agents that perceive their environment, learn from multimodal data, improve from human feedback or preferences, coordinate actions for common goals, and adapt dynamically.

One key aspect of my research agenda is considering reputation systems as tools for altruistic cooperation in complex multi-agent reinforcement learning environments. Cooperation is important in natural and artificial systems. It allows for agents with individual goals to reach beneficial group outcomes, even when group and individual incentives are not perfectly aligned. If cooperation is costly but the benefits of cooperation can be enjoyed by all agents, the temptation to pay no cost is a dominant strategy and cooperation is hard to establish and maintain, unless a specific mechanism is in place to foster cooperation. While reputation systems help other agents evaluate and predict future interactions by recording and disseminating the historical behavior of agents. Reputation not only affects trust and expectations between agents, but also motivates individuals to cooperate.