Umer Siddique

Personal Website | umersiddique297@gmail.com | Google Scholar | LinkedIn

Research Interests

Reinforcement Learning (RL), Multi-Agent RL (MARL), AI Safety and Alignment (Fairness, Bias, Transparency), Sample-Efficient RL, RL from Human-Feedback (RLHF), Large Language Models (LLMs), and applications in Robotics, Autonomous Systems with a particular focus in defense, aerospace, and critical infrastructure domains.

Education

PhD in Electrical Engineering, University of Texas at San Antonio, USA *2022–Present*
Advisor: Prof. Yongcan Cao

MS in Information and Communication Engineering, UM-SJTU Joint Institute, China *2018–2021*
Advisor: Prof. Paul Weng

BS in Information Technology, University of Agriculture Faisalabad, Pakistan *2013–2017*
Advisor: Prof. Tasleem Mustafa

Dissertation

Title: *AI Alignment through Reinforcement Learning: Fairness, Safety, and Social Welfare Optimization*
Advisor: Prof. Yongcan Cao **Expected Completion:** Spring 2026

Experience

Research Experience

Research Assistant, University of Texas at San Antonio *2022–Present*

- Led an Office of Naval Research (ONR)-funded project to develop resource-aware RL algorithms, reducing computational requirements by 40% using event-triggered control for supporting U.S. Navy initiatives.
- Co-led an Army Research Lab (ARL)-funded project to develop human-aligned RL and MARL systems that address fairness and mitigate bias in decision support systems.
- Contributed to multiple federal research projects (ONR, ARL) focused on sample-efficient RL/MARL methods for real-world robotics, control, and defense systems.
- Developed and deployed novel RL-based controllers for complex applications, including aerospace systems, nonlinear control problems, and a custom end-to-end multi-agent UAV drone simulation.
- Designed multi-objective alignment strategies to ensure safety, fairness, and transparency in RL/MARL systems, including new methods to reduce bias in RLHF and LLMs.

Research Assistant, University of Michigan-Shanghai Jiao Tong University, China *2018–2021*

- Introduced a new fairness paradigm in RL by defining novel welfare functions to measure and enforce fairness in agent policies.
- Developed state-of-the-art fair RL algorithms that demonstrably improved equity among users in multi-objective scenarios.
- Proposed and validated novel methods for ensuring fairness in decentralized MARL, achieving superior performance over baseline models.
- Established a theoretical and practical link between fair solutions and compromise solutions in RL.

Industry Experience

AI Engineer, Neurog, Islamabad *Dec 2021–May 2022*

- Deployed production RL agents for financial portfolio optimization and strategic game-playing systems.
- Achieved 20% improvement in win rates over baseline methods in internal football simulation.
- Optimized training pipelines and managed GPU resource allocation for efficient model deployment in resource-constrained environments.

Teaching Experience

Computer Science Lecturer, Government College Satiana Road, Faisalabad, Pakistan *2017–2018*

- **Courses:** Introduction to Computer Science, Databases, Programming.

Lecturer (Part Time), Lyallpur International College, Faisalabad, Pakistan *2016*

- **Courses:** Computing, Calculus I, Discrete Mathematics, Python Programming.

Publications

Refereed Journal Articles

[J1] **Umer Siddique**, A. Sinha, Y. Cao. "Learning Fair Policies in Multi-Objective Preference-based Reinforcement Learning." *Machine Learning Journal (MLJ)*, 2026.

[J2] **Umer Siddique**, P. Li, Y. Cao. "Fair Multi-Agent Reinforcement Learning for Traffic Control." *Journal on Autonomous Transportation Systems (JATS)*, 2025.

[J3] J. Qian, **Umer Siddique**, G. Yu, P. Weng. "From Fair Solutions to Compromise Solutions in Multi-Objective Deep Reinforcement Learning." *Neural Computing and Applications (NCAA)*, 2025.

[J4] S. Abahussein, D. Ye, C. Zhu, Z. Cheng, **Umer Siddique**, S. Shen. "Multi-Agent Reinforcement Learning for Online Food Delivery with Location Privacy Preservation." *Information*, 2023.

Peer-Reviewed Conference and Workshop Articles

[C1] **Umer Siddique**, A. Sinha, Y. Cao. "Adaptive Event-triggered Policy Gradient for Multi-Agent Reinforcement Learning." *American Control Conference (ACC)*, 2026.

[C2] **Umer Siddique**, A. Sinha, Y. Cao. "Adaptive Event-triggered Reinforcement Learning Control for Complex Nonlinear Systems." *American Control Conference (ACC)*, 2025.

[C3] **Umer Siddique**, P. Li, and Y. Cao. "Learning Fair Pareto-Optimal Policies in Multi-Objective Reinforcement Learning." *Reinforcement Learning Conference (RLC)*, 2025.

[C4] **Umer Siddique**, P. Li, Y. Cao. "Towards Fair and Efficient Policy Learning in Cooperative Multi-Agent Reinforcement Learning." *International Conference on Autonomous Agents and Multiagent Systems (AAMAS)* [Extended Abstract], 2025.

[C5] **Umer Siddique**, P. Kumar, A. Sinha, Y. Cao. "Deep Reinforcement Learning for Target Enclosing Under Limited Sensing and Autopilot Lag." *American Institute of Aeronautics and Astronautics (AIAA) SCITECH*, 2026.

[C6] P. Li, **Umer Siddique**, Y. Cao. "Symbolic Policy Distillation for Interpretable Reinforcement Learning." *Mechanistic Interpretability Workshop at NeurIPS*, 2025.

[C7] C. Wallace, **Umer Siddique**, Y. Cao. "RECOLLAB: Retrieval-Augmented LLMs for Cooperative Ad-hoc Teammate Modeling." *Bridging Language, Agent, and World Models for Reasoning and Planning Workshop at NeurIPS*, 2025.

[C8] P. Li, **Umer Siddique**, Y. Cao. "From Explainability to Interpretability: Interpretable Reinforcement Learning Via Model Explanations." *Reinforcement Learning Conference (RLC)*, 2025.

[C9] C. Wallace, **Umer Siddique**, Y. Cao. "TransAM: Transformer-Based Agent Modeling for Multi-Agent Systems via Local Trajectory Encoding." *Reinforcement Learning Conference (RLC)*, 2025.

[C10] J. Cofield, **Umer Siddique**, Y. Cao. "MODIFLY: A Scalable End-to-end Multi-Agent Simulation for Unmanned Aerial Vehicles." *The 26th International Workshop on Multi-Agent-Based Simulation (MABS)*, 2025.

[C11] **Umer Siddique**, P. Li, Y. Cao. "Fairness in Traffic Control: Decentralized Multi-agent Reinforcement Learning with Generalized Gini Welfare Functions." *Multi-Agent RL for Transportation Autonomy at AAMAS*, 2025.

[C12] C. Wallace, **Umer Siddique**, Y. Cao. "Opponent Transformer: Modeling Opponent Policies as a Sequence Problem." *Coordination and Cooperation for Multi-Agent RL Methods at RLC*, 2024.

[C13] **Umer Siddique**, P. Li, Y. Cao. "Towards Fair and Equitable Policy Learning in Cooperative Multi-Agent Reinforcement Learning." *Coordination and Cooperation for Multi-Agent RL Methods at RLC*, 2024.

[C14] G. Yu, **Umer Siddique** and P. Weng. "Fair Deep Reinforcement Learning with Preferential Treatment." *European Conference on Artificial Intelligence (ECAI)*, 2023.

[C15] **Umer Siddique**, A. Sinha, and Y. Cao. "Fairness in Preference-based Reinforcement Learning." *The Many Facets of Preference-Based Learning at ICML*, 2023.

[C16] **Umer Siddique**, A. Sinha, and Y. Cao. "On Deep Reinforcement Learning for Target Capture Autonomous Guidance." *American Institute of Aeronautics and Astronautics (AIAA) SCITECH*, 2023.

[C17] G. Yu, **Umer Siddique** and P. Weng. "Fair Deep Reinforcement Learning with Generalized Gini Welfare Functions." *Adaptive Learning Agents (ALA)*, 2023 **(Best Paper Award)**.

[C18] M. Zimmer, C. Glanois, **Umer Siddique** and P. Weng. "Learning Fair Policies in Decentralized Multi-Agent Reinforcement Learning." *International Conference on Machine Learning (ICML)*, 2021.

[C19] **Umer Siddique**, P. Weng, M. Zimmer. "Learning Fair Policies in Multi-Objective (Deep) Reinforcement Learning with Average and Discounted Rewards." *International Conference on Machine Learning (ICML)*, 2020.

Patents

[P1] J. Cofield, **Umer Siddique**, Y. Cao. *Scalable End-to-End Multi-Agent Simulation for Unmanned Vehicles.* U.S. Patent [Pending].

Talks

[T1] "Event-triggered for Multi-Agent Reinforcement Learning." American Control Conference (ACC), 2026.

[T2] "Deep Reinforcement Learning for Target Enclosing Under Limited Sensing." American Institute of Aeronautics and Astronautics (AIAA) SCITECH, 2026.

[T3] "Event-triggered Reinforcement Learning for Complex Nonlinear Control Systems." American Control Conference (ACC), 2025.

[T4] "Fairness in Multi-Objective Reinforcement Learning." Adaptive Learning Agents at International Conference on Autonomous Agents and Multiagent Systems (AAMAS), 2025.

[T5] "Multi-Agent Reinforcement Learning for Traffic Light Control." Multi-Agent RL for Transportation Autonomy at AAAI Conference on Artificial Intelligence (AAAI), 2025.

[T6] "Reinforcement Learning for Target Capture Autonomous Guidance." American Institute of Aeronautics and Astronautics (AIAA) SCITECH, 2024.

Awards and Honors

- **AAAI Student Travel and Volunteer Award**, 2024
  Given by the Association for the Advancement of Artificial Intelligence (AAAI).
- **Student Registration Award**, Reinforcement Learning Conference (RLC), 2024
  Sponsored by the Cooperative AI Foundation.
- **Best Paper Award**, Adaptive Learning Agents (ALA) @ AAMAS, 2023
  Awarded by the ALA Workshop organizers, sponsored by *Neural Computing and Applications* journal.
- **Fully Funded PhD Scholarship**, University of Texas at San Antonio, 2022–Present
  Supported by projects funded by the Office of Naval Research, Army Research Lab, and Army Research Office.
- **Fully Funded Master's Scholarship**, University of Michigan–Shanghai Jiao Tong Joint Institute, 2018–2021
  Awarded under the Chinese Government Scholarship program.
- **Excellence International Student Award**, University of Michigan–Shanghai Jiao Tong Joint Institute, 2018.
- **Gold Medalist**, BS in Information Technology, University of Agriculture Faisalabad, 2017.
- **Young Motivational Teacher Award**, Government College, Satiana Road Faisalabad, 2018.
- **University Merit-Based Scholarship**, University of Agriculture Faisalabad, 2013–2017.
- **Higher Education Commission (HEC) Merit and Need-Based Scholarship**, 2014–2017.

Professional Services

Journal Reviewer

- Neural Computing and Applications
- Automatica
- IEEE Transactions on Cybernetics
- IEEE Robotics and Automation Letters

Conference Reviewer

- Conference on Neural Information Processing Systems (NeurIPS)
- International Conference on Machine Learning (ICML)
- International Conference on Learning Representations (ICLR)
- AAAI Conference on Artificial Intelligence (AAAI)
- International Conference on Artificial Intelligence and Statistics (AISTATS)
- IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)
- International Joint Conference on Artificial Intelligence (IJCAI)
- European Conference on Artificial Intelligence (ECAI)

Other Academic Service

- Workshop Co-chair: *Good Data for Generative AI*, AAAI 2025.

Technical Skills

Programming: Python (expert), MATLAB, C++, PHP
ML/AI: PyTorch, JAX, TensorFlow, scikit-learn, Keras, NumPy, Pandas, Hugging Face, Transformers
RLHF/LLM: TRL, vLLM, VeRL, PEFT, LoRA, DPO
RL Tools: Stable-Baselines3, Gymnasium, MuJoCo, ROS, SUMO
DevOps: Docker, AWS, CI/CD, Slurm