

Umer Siddique

Department of Electrical and Computer Engineering
University of Texas at San Antonio, San Antonio, TX, USA

umersiddique297@gmail.com | umersiddique94.github.io | Google Scholar | GitHub

Research Focus

My research focuses on **AI alignment through reinforcement learning (RL) for fair, safe, and socially aligned decision making**. I develop RL methods that enable autonomous decision-making systems to account for competing objectives, stakeholder preferences, and safety requirements. My work spans multi-objective and preference-based RL, multi-agent systems, and alignment of learned decision-making systems, with applications to autonomous and real-world systems.

Education

PhD, Electrical Engineering *2022–Dec 2026 (expected)*

University of Texas at San Antonio, USA

Advisor: Prof. Yongcan Cao

Dissertation: *AI Alignment through Reinforcement Learning: Fairness, Safety, and Social Welfare Optimization*

MS, Information and Communication Engineering *2018–2021*

University of Michigan–Shanghai Jiao Tong University Joint Institute, Shanghai, China

Advisor: Prof. Paul Weng

Research and Professional Experience

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

Sponsored projects:

- **Resource-Aware RL for Autonomous Systems**, Office of Naval Research *2022–Present*
Lead Graduate Researcher. Lead research on resource-aware RL for autonomous systems, including event-triggered and self-triggered learning methods that reduce policy-update computation by approximately 40% relative to fixed-interval baselines at matched return. Prepare quarterly and annual technical reports and present research findings at Navy CAST program reviews (2023–2026).
- **Human-Aligned RL for Decision Support**, Army Research Laboratory *2022–Present*
Graduate Researcher. Co-lead research on fairness and bias mitigation in RL and multi-agent decision-support systems.

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

Built and deployed RL agents for financial portfolio optimization and strategic game-playing.

Research Assistant, UM–SJTU Joint Institute, Shanghai, China *2018–2021*

Fair multi-objective and decentralized multi-agent RL with Prof. Paul Weng (ICML 2020, ICML 2021).

Research Awards and Honors

- **Jamshidi Fellowship**, University of Texas at San Antonio, 2026. Awarded annually to one outstanding researcher; \$3,500 award.
- **Gold Reviewer Award**, International Conference on Machine Learning (ICML), 2026.
- **Best Paper Award**, Adaptive Learning Agents (ALA) Workshop at AAMAS, 2023. Sponsored by *Neural Computing and Applications*.
- **Student Travel and Volunteer Award**, AAAI, 2024.
- **Student Award**, Reinforcement Learning Conference (RLC), 2024. Sponsored by the Cooperative AI Foundation.

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)*, 2025.
- [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*, 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.

Refereed Conference Proceedings

- [C1] **Umer Siddique**, P. Li, C. Wallace, Y. Cao. “Inference-Time Policy Alignment for Fair Reinforcement Learning.” *Reinforcement Learning Conference (RLC)*, 2026.
- [C2] **Umer Siddique**, A. Sinha, Y. Cao. “Adaptive Event-triggered Policy Gradient for Multi-Agent Reinforcement Learning.” *American Control Conference (ACC)*, 2026.
- [C3] **Umer Siddique**, P. Kumar, A. Sinha, Y. Cao. “Deep Reinforcement Learning for Target Enclosing Under Limited Sensing and Autopilot Lag.” *AIAA SCITECH*, 2026.
- [C4] M. A. Manzoor, D. Azizov, D. Orel, **Umer Siddique**, Z. M. Mujahid, Y. Hou, P. Nakov. “A Multi-View Media Profiling Suite: Resources, Evaluation, and Analysis.” *Annual Meeting of the Association for Computational Linguistics (ACL)*, 2026.
- [C5] **Umer Siddique**, P. Li, Y. Cao. “Learning Fair Pareto-Optimal Policies in Multi-Objective Reinforcement Learning.” *Reinforcement Learning Conference (RLC)*, 2025.
- [C6] P. Li, **Umer Siddique**, Y. Cao. “From Explainability to Interpretability: Interpretable Reinforcement Learning Via Model Explanations.” *Reinforcement Learning Conference (RLC)*, 2025.
- [C7] C. Wallace, **Umer Siddique**, Y. Cao. “TransAM: Transformer-Based Agent Modeling for Multi-Agent Systems via Local Trajectory Encoding.” *Reinforcement Learning Conference (RLC)*, 2025.
- [C8] **Umer Siddique**, A. Sinha, Y. Cao. “Adaptive Event-triggered Reinforcement Learning Control for Complex Nonlinear Systems.” *American Control Conference (ACC)*, 2025.
- [C9] G. Yu, **Umer Siddique**, P. Weng. “Fair Deep Reinforcement Learning with Preferential Treatment.” *European Conference on Artificial Intelligence (ECAI)*, 2023.
- [C10] **Umer Siddique**, A. Sinha, Y. Cao. “On Deep Reinforcement Learning for Target Capture Autonomous Guidance.” *AIAA SCITECH*, 2023.
- [C11] M. Zimmer, C. Glanois, **Umer Siddique**, P. Weng. “Learning Fair Policies in Decentralized Cooperative Multi-Agent Reinforcement Learning.” *ICML*, 2021.
- [C12] **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].

Preprints

- [U1] **Umer Siddique**, P. Li, Y. Cao. “Inference-Time Safety Alignment for Reinforcement Learning.”
- [U2] P. Li, **Umer Siddique**, Y. Cao. “SPID: Fidelity-Preserving Symbolic Policy Distillation for Interpretable Reinforcement Learning.”
- [U3] P. Li, **Umer Siddique**, Y. Cao. “Automated Interpretability for Diagnosing Reinforcement Learning Failures.”

- [U4] Conor Wallace, **Umer Siddique**, Y. Cao, P. Li. “BayesToM: Retrieval-Augmented Behavior Models for Calibrated Partner Identification.”

Refereed Workshop Papers and Extended Abstracts

- [W1] **Umer Siddique**, P. Li, Y. Cao. “Right Goal, Wrong Body: Re-grounding Frozen VLAs under Actuation Shift.” *Workshop on The Science of Physical AI Safety, CoRL*, 2026.
- [W2] **Umer Siddique**, P. Li, Y. Cao. “Beyond the Weighted Sum: Rubric Aggregation for Multi-Criterion LLM Alignment.” *Workshop on Agent Behavior, COLM*, 2026.
- [W3] P. Li, **Umer Siddique**, Y. Cao. “Symbolic Policy Distillation for Interpretable Reinforcement Learning.” *Mechanistic Interpretability Workshop at NeurIPS*, 2025.
- [W4] C. Wallace, **Umer Siddique**, Y. Cao. “RECOLLAB: Retrieval-Augmented LLMs for Cooperative Ad-hoc Teammate Modeling.” *Bridging Language, Agent, and World Models Workshop at NeurIPS*, 2025.
- [W5] J. Cofield, **Umer Siddique**, Y. Cao. “MODIFY: A Scalable End-to-end Multi-Agent Simulation for Unmanned Aerial Vehicles.” *26th International Workshop on Multi-Agent-Based Simulation (MABS)*, 2025.
- [W6] **Umer Siddique**, P. Li, Y. Cao. “Towards Fair and Efficient Policy Learning in Cooperative Multi-Agent Reinforcement Learning.” *AAMAS [Extended Abstract]*, 2025.
- [W7] **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 Workshop at AAMAS*, 2025.
- [W8] C. Wallace, **Umer Siddique**, Y. Cao. “Opponent Transformer: Modeling Opponent Policies as a Sequence Problem.” *Coordination and Cooperation for Multi-Agent RL at RLC*, 2024.
- [W9] **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 Workshop at RLC*, 2024.
- [W10] **Umer Siddique**, A. Sinha, Y. Cao. “Fairness in Preference-based Reinforcement Learning.” *The Many Facets of Preference-Based Learning Workshop at ICML*, 2023.
- [W11] G. Yu, **Umer Siddique**, P. Weng. “Fair Deep Reinforcement Learning with Generalized Gini Welfare Functions.” *Adaptive Learning Agents (ALA) Workshop at AAMAS*, 2023. **(Best Paper Award)**

Teaching Experience

Lecturer, Computer Science, Government College Satiana Road, Faisalabad, Pakistan 2017–2018
Instructor of record for Introduction to Computer Science, Databases, and Programming.

Lecturer (Part-time), Lyallpur International College, Faisalabad, Pakistan 2016
Instructor for Computing, Calculus I, Discrete Mathematics, and Python Programming.

Seminar and Tutorial Organizer, University of Texas at San Antonio 2022–Present
Organize and deliver research seminars and hands-on tutorials on RL, RLHF, and AI alignment.

Research Mentoring and Advising

- **Jeremy Cofield** (Undergraduate, UTSA). Mentored through his first research project on multi-agent UAV simulation, from initial design to publication. Resulted in a first-authored MABS 2025 paper (MODIFY) and a pending U.S. patent. Now at SpaceX.
- **Sergio DeLeon Jr.** (Undergraduate, UTSA). Mentored on the design of water-based multi-agent simulation environments for RL. Now at Raytheon.
- **Peilang Li** (PhD student, UTSA). Mentored on interpretable RL, guiding problem formulation, experimental design, and paper development. Resulted in co-authored papers at RLC 2025 and the NeurIPS 2025 Mechanistic Interpretability Workshop, on which he is first author.

Media and Public Engagement

- Interviewed for “The New Wild West: Texans on the AI Frontier,” *Texas Standard* (Texas public radio), 2026. Discussed research at the UTSA Unmanned Systems Lab on training robots from failed and synthetic data so they generalize at test time to tasks they were not trained on.

Conference Presentations

- [T1] “Adaptive Event-triggered Policy Gradient for Multi-Agent Reinforcement Learning.” American Control Conference (ACC), 2026.
- [T2] “Deep Reinforcement Learning for Target Enclosing Under Limited Sensing.” AIAA SCITECH, 2026.
- [T3] “Adaptive Event-triggered Reinforcement Learning Control for Complex Nonlinear Systems.” American Control Conference (ACC), 2025.
- [T4] “Fairness in Multi-Objective Reinforcement Learning.” Adaptive Learning Agents Workshop at AAMAS, 2025.
- [T5] “Multi-Agent Reinforcement Learning for Traffic Light Control.” Multi-Agent RL for Transportation Autonomy Workshop at AAMAS, 2025.
- [T6] “Reinforcement Learning for Target Capture Autonomous Guidance.” AIAA SCITECH, 2024.
- [T7] Navy CAST Program Review presentations, Office of Naval Research, 2023, 2024, 2025, 2026.

Professional Service

Conference Reviewing

NeurIPS (2024–2026) · ICML (2023–2026; Gold Reviewer, 2026) · ICLR (2024–2026) · AAAI (2024–2026) · AISTATS (2025) · IJCAI (2024–2025) · ECAI (2024–2025) · IROS (2024)

Journal Reviewing

Automatica · *IEEE Transactions on Cybernetics* · *Neural Computing and Applications*

Scholarships and Early Honors

- **Fully Funded PhD Scholarship**, University of Texas at San Antonio, 2022–Present. Supported by ONR, ARL, and ARO.
- **Fully Funded Master’s Scholarship**, UM–SJTU Joint Institute, 2018–2021. Chinese Government Scholarship program.
- **Excellence International Student Award**, UM–SJTU 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**, Pakistan, 2014–2017.

References

Prof. Yongcan Cao

Professor, Department of Electrical and Computer Engineering, UT San Antonio
yongcan.cao@utsa.edu

Prof. Paul Weng

Associate Professor, Duke Kunshan University
paul.weng@dukekunshan.edu.cn

Dr. Abhinav Sinha

Assistant Professor, Department of Aerospace Engineering, University of Cincinnati
sinhaab@ucmail.uc.edu

Prof. Chunjiang Qian

Professor, Department of Electrical and Computer Engineering, UT San Antonio
chunjiang.qian@utsa.edu