

Muhammad Ahmed Mohsin

muahmed@stanford.edu
git/ahmd-mohsin
building conoid.ai

RESEARCH INTERESTS

Preference Optimization and Alignment for LLMs. Designing sample-efficient preference-learning and reinforcement-learning methods for aligning language-model reasoning: continuous-utility and general-preference formulations (CU-DPO, General Preference RL), active alignment under Bayesian general preference models with calibrated uncertainty, and sycophnacy reduction.

Test-Time Scaling and Test-Time Discovery. Developing adaptive test-time compute methods for reasoning and scientific discovery under uncertainty: dynamic control of inference depth, tool invocation, and verification under strict budgets; stratified scaling search for diffusion language models; and epistemic-uncertainty-driven test-time training for discovery.

Interactive and Multi-Agent Systems. Building self-evolving multi-agent ecosystems in which agents accumulate scoped memory, earn reputation through Bayesian posteriors, and coordinate over an evolving social graph. Related work on learning from code-agent trajectories via causal redundancy analysis, and on privacy, security, and shared context in collaborative agentic reasoning.

EDUCATION

Stanford University, Bay Area, California 2024–Present
Ph.D. in Electrical Engineering (LLM Post Training, Test Time Scaling, Reinforcement Learning)

- **CGPA:** 4.0/4.0
- **Advisors:** Dr. Emily Fox and Dr. John Cioffi — Research collaboration with Google DeepMind, Meta, Amazon AGI and Microsoft Core AI.

National University of Sciences & Technology (NUST), Islamabad, Pakistan Aug. 2020 – June 2024

Bachelor of Engineering in Electrical Engineering

- **CGPA:** 3.87/4.0 | **Specialization GPA:** 4.0/4.0
- **Senior Design Thesis:** Reinforcement learning for resource allocation optimization in dynamic environments
- **Rector's Gold Medal:** Awarded Rector's Gold Medal for best final year project cohort (2024)

RESEARCH AND PROFESSIONAL EXPERIENCE

Stanford Artificial Intelligence Laboratory (SAIL) September 2025 – Present

Advisor: Prof. Emily Fox

Project: Internet of Evolving Agents

- Co-developed a modular multi-agent ecosystem where autonomous agents evolve their capabilities, reputation, and social connections over time through Bayesian reputation updates, dynamic team formation, and social graph evolution. The framework enables emergent specialization and self-organizing collaboration for complex task execution (NeurIPS 2026, in progress).

Project: Test-Time Compute and Reasoning in Large Language Models

- Submitted work on adaptive test-time compute allocation for LLM reasoning, focusing on dynamic control of inference depth, tool invocation, and verification under strict compute budgets, and analyzing principled trade-offs among accuracy, latency, and reliability.

Project: Bayesian Preference Alignment for Mathematical Reasoning

- Developed active learning frameworks for Bayesian General Preference Models and CU-DPO to align small language models for mathematical reasoning, enabling sample-efficient preference learning with calibrated uncertainty (ICML 2026, CoLM 2026).

Intel Corporation, Researcher

September 2024 – December 2024

Advisor: Dr. John M. Cioffi

Project: Neural Gaussian Radio Fields for Environment Perception

- Developed a 3D computer vision-based channel estimation framework for next-generation wireless networks, implementing a CUDA-based differentiable real-time pipeline achieving 1 ms inference latency (KDD 2026 submission).

Samsung Semiconductors, Researcher

June 2024 – September 2025

Advisor: Dr. John M. Cioffi

Project: Graph Neural Networks for Accelerating Low-Rank SDP Solvers (JMLR 2026)

- Developed a constraint-graph representation of SDPs with a Graph Attention Network encoder to predict rank trajectories, integrating learned rank schedules into low-rank solvers to eliminate manual heuristics and achieve up to $3\times$ speedups on large-scale benchmarks.

SELECTED PUBLICATIONS

[C0] **Distilling Disagreement at Test Time** EMNLP 2026
Muhammad Ahmed Mohsin, Muhammad Umer, Ahsan Bilal, John M. Cioffi, Emily Fox

[C1] **General Preference Reinforcement Learning** NeurIPS 2026
Muhammad Umer, M. A. Mohsin*, A. Bilal*, E. Fox*

[C2] **Epistemic Uncertainty for Test-Time Discovery** NeurIPS 2026
Muhammad Ahmed Mohsin, Kainat Riaz*, Muhammad Umer, Ahsan Bilal, John M. Cioffi, Emily Fox*

[C3] **Internet of Evolving Agents** EMNLP 2026
Z. Ali, M. A. Mohsin*, M. Umer, A. Bilal, E. Fox*

[C4] **Sycophancy Disentanglement in LLMs via Reward Decomposition** CoLM 2026
M. A. Mohsin, A. Bilal*, M. Umer, E. Fox*

[C5] S^3 : **Stratified Scaling Search for Test-Time in Diffusion Language Models** CoLM 2026
A. Bilal, M. A. Mohsin*, M. Umer, D. F. Hougen*

[C6] **Active Alignment with Bayesian General Preference Models** CoLM 2026
M. Umer, M. A. Mohsin*, A. Bilal, Ellen Vitercik, J. M. Cioffi*

[C7] **Continuous-Utility Direct Preference Optimization** ICML 2026
M. A. Mohsin, M. Umer, A. Bilal, Ellen Vitercik, J. M. Cioffi

[C8] **What If We Allocate Test-Time Compute Adaptively?** ICML 2026
A. Bilal, M. A. Mohsin*, M. Umer, D. F. Hougen, J. M. Cioffi*

[C9] **Graph Neural Networks for Accelerating Low-Rank SDP Solvers** JMLR 2026
M. A. Mohsin, M. Umer, A. Bilal, J. M. Cioffi, Ellen Vitercik

¹*Equal contributing authors.

- [C10] **Neural Gaussian Radio Fields for Channel Estimation** KDD 2026
M. A. Mohsin, M. Umer*, A. Bilal, J. M. Cioffi*
- [C11] **Continual Learning for Wireless Channel Prediction** ICML 2025
M. A. Mohsin, M. Umer, A. Bilal, J. M. Cioffi
- [C12] **Automated Prompt Optimization Improves Robustness of Medical Benchmarks for LLMs** ML4H 2025
Asad Aali, Muhammad Ahmed Mohsin, V. Bikia, A. Singhvi, S. Bedi, H. Cui, M. A. F. Hernandez, A. Unell, Y. Mai, M. A. Pfeffer, S. Koyejo, R. Daneshjou, E. Alsentzer, C. Potts, A. S. Chaudhari
- [C13] **On the Fundamental Limits of Large Models at Scale** TMLR 2025
Muhammad Ahmed Mohsin, Muhammad Umer, Ahsan Bilal, Z. Memon, M. I. Qadir, S. Bhattacharya, H. Rizwan, A. R. Gorle, M. Z. Kazmi, A. Mohsin, A. Subhan, M. U. Rafique, Z. He, P. Mehta, J. Han, M. A. Jamshed, D. F. Hougén, J. M. Cioffi
- [C14] **Task-Aware Distributed Source Coding for Correlated Audio Signals Using Perceptual Loss** AAAI 2024
M. A. Mohsin, A. Bilal, S. Bhattacharya, J. M. Cioffi
- [C15] **PyramidTabNet: Transformer-Based Table Recognition in Image-Based Documents** ICDAR 2023
M. Umer, M. A. Mohsin, A. Ul-Hasan, F. Shafait
- [C16] **Conditional Prior-Based Non-Stationary Channel Estimation Using Accelerated Diffusion Model-Based Denoising** ICASSP 2025
M. A. Mohsin, A. Bilal, M. Umer, A. Aali, M. A. Jamshed, J. M. Cioffi
- [C17] **Successive Interference Cancellation-Aided Diffusion Models for Joint Channel Estimation and Data Detection** ICASSP 2024
S. Bhattacharya, M. A. Mohsin, K. Rajabalifardi, J. M. Cioffi

REVIEWER AND GRANTS

Conference Reviewer: AAAI, ICDAR, NeurIPS, ICML, ICLR, KDD, CoLM

Journal Reviewer: IEEE Transactions on AI, Transactions on Machine Learning Research, Applied Artificial Intelligence.

Travel Grants: IEEE GLOBECOM 2024 (Cape Town, South Africa), IEEE ICC 2025 (Canada), Stanford Conference Travel Grant 2025.

Best Paper/Poster Awards: 6G Summit Abu Dhabi (Best Poster Nomination), IEEE ICC Canada (Best Workshop Paper Award), IEEE Communications Society Competition (honorary mention), IEEE FIT 2026 Best Main Conference Paper Award.

Technical Program Committee: NeurIPS 2025, ICASSP 2026.

Workshop Co-Chair: IEEE ICC 2026.

Area Chair: ICASSP 2026, NeurIPS 2025.

HONORS AND ACHIEVEMENTS

Fellowship: Awarded Stanford Graduate Fellowship.

Knight-Hennessy Fellowship: Finalist.

President's Medal: Third position nationwide in pre-engineering group.

Scholarships: HSSC Federal Board Scholarship (PKR 200,000); STEP-ECAT Scholarship for top 10 engineering candidates nationwide (PKR 200,000).