Hamed Khosravi

Ph.D. Student, Machine Learning
H. Milton Stewart School of Industrial & Systems Engineering, Georgia Institute of Technology
Atlanta, GA, USA · (304) 216-0412 · hkhosravi7@gatech.edu
Personal Website · Google Scholar · GitHub · LinkedIn

Education

Georgia Institute of Technology, Atlanta, GA Aug 2025 – present

Ph.D. in Machine Learning, H. Milton Stewart School of Industrial & Systems Engineering. GPA 4.0/4.0.
Advisor: Prof. Xiaoming Huo.

West Virginia University, Morgantown, WV

M.S. in Industrial Engineering (Data Science), Department of Industrial and Management Systems Engineering. GPA 4.0/4.0.
Advisor: Dr. Imtiaz Ahmed · Co-advisor: Dr. Srinjoy Das.
Thesis: *Attentive Optimism in the Face of Uncertainty: Balancing Exploration–Exploitation in Multi-Armed Bandits.*

Iran University of Science and Technology, Tehran, Iran

B.S. in Industrial Engineering. Graduated as a Distinguished Student.

Research Interests

Statistical machine learning; reinforcement learning; LLM and AI-agent evaluation; anytime-valid inference and sequential testing; reinforcement learning with verifiable rewards (RLVR) for LLMs; evidence-grounded reasoning and benchmarking for agents; human–AI teaming; leaderboard and ranking validity; multi-armed and contextual bandits; non-stationary and online learning; sequential decision making; conformal prediction and uncertainty quantification; causal inference and experimental design.

Research Projects

Ph.D. research at the Georgia Institute of Technology, advised by Prof. Xiaoming Huo. Bracketed numbers refer to the Publications list.

LLM Agent Reasoning: Reliability Guardrails and Benchmarking Oct 2026 – present

PROBLEM LLM agents that check records against an organization's requirements report conclusions their own evidence does not justify, or settle for a weaker conclusion when another route would still have succeeded—and there is no benchmark that scores an agent on what its evidence actually establishes.

APPROACH Fix the evidence routes for each conclusion before the agent looks at a case, then have a separate program record what each check establishes and authorize only the conclusions those results support, pointing the agent to the checks that would justify a stronger one. Evaluated on SUFFICE, a benchmark of evidence-based review tasks built from real public records with a deterministic scorer.

PAPERS [5]

LLM and AI Model Evaluation: Reliable Rankings and Leaderboards Sep 2026 – present

PROBLEM Public AI leaderboards are read continuously and updated as votes arrive, so a ranking that looks settled at one glance can be wrong at the next; because every standing is public, an adversary can watch each claim and rig the votes behind it.

APPROACH Rank confidence sequences that keep every pairwise claim valid at all times, however often the leaderboard is read, and a certified corruption budget published with each claim: a tolerance such that the claim is correct unless more records than that were corrupted.

PAPERS [6], [7]

Statistical Evaluation & Deployment of LLMs and AI Agents Mar 2026 – present

PROBLEM Organizations must decide which LLM or AI agent handles which task—and whether a human–AI team beats either alone—with limited evaluation data and budget.

APPROACH Decision-targeted evaluation methods that direct testing toward the comparisons that could actually change the deployment choice.

PAPERS [3], [2], [4]

Reinforcement Learning × LLMs for Reliable and Adaptive AI Systems Nov 2025 – present

PROBLEM Language models that keep learning from online feedback must stay reliable under uncertainty and distribution shift.

APPROACH Anytime-valid risk control for RLVR-trained LLMs; an LLM-observed multi-agent reinforcement learning framework that reads operational evidence to decide when deployed models should be maintained, retrained, or updated; and a decision layer for constrained reinforcement learning that certifies when a single reading of a natural-language safety rule is sufficient to train under, and otherwise enforces multiple readings.

PAPERS [1], [8], [10]

Adaptive Personalization Under Changing User Preferences Oct 2025 – present

PROBLEM Recommendation and ranking systems learn from sequential feedback while the preferences they are learning drift.

APPROACH Contextual bandit methods that balance exploring new options against exploiting known ones, with low-rank structure that tracks an evolving latent preference space under non-stationarity.

PAPERS [9]

Publications

---

* equal contribution (co-first author) · † corresponding author

Peer-Reviewed Conference & Workshop Papers

1. **Hamed Khosravi**, Xiaoming Huo. Conformal Selective Acting: Anytime-Valid Risk Control for RLVR-Trained LLMs. **NeurIPS 2026** (main track). [paper] [code]
2. **Hamed Khosravi**, Xiaoming Huo. Which LLM for Which Work? Budgeted Model Allocation under Uncertain Evaluation. *Can We Trust AI Evaluation? Workshop (TAE)*, **NeurIPS 2026**. [paper] [code]
3. **Hamed Khosravi**, Xiaoming Huo. Beyond "AI Helps Humans": Decision-Targeted Evaluation Design for Human–Agent Teams in the Agentic Era. *Human–AI Coevolution Workshop (HAIC)*, **NeurIPS 2026**. [paper] [code]
4. **Hamed Khosravi**, Xiaoming Huo. TEAM-Audit: A Two-Endpoint Audit of Margins. *Workshop on Human–AI Teaming (HAT 2026), AAAI Conference on Human Computation and Crowdsourcing / Collective Intelligence (HCOMP/CI)*, 2026. [paper] [code]

Preprints and Papers Under Review

5. **Hamed Khosravi**, Xiaoming Huo. PRIME: Towards Verifiable Evidence-Grounded Reasoning in LLM Agents. *arXiv:2610.22426*, 2026. Under review at *NAACL 2027*. [paper] [code] [benchmark]
6. **Hamed Khosravi**, Xiaoming Huo. Certified Corruption Budgets: Anytime-Valid Leaderboard Claims under Adaptive Rigging. *arXiv:2610.10597*, 2026. Under review at *AISTATS 2027*. [paper] [code]
7. **Hamed Khosravi**, Xiaoming Huo. Rank Confidence Sequences: Anytime-Valid Leaderboards. *arXiv:2609.32211*, 2026. Under review at *ICLR 2027*. [paper] [code]
8. **Hamed Khosravi**, Xiaoming Huo. LLM-Observed Multi-Agent Maintenance with Amortized Reinforcement Learning Policies. *SSRN*, 2026. Under review at *IISE Transactions*. [paper] [code]
9. **Hamed Khosravi**, Xiaoming Huo. Catching a Moving Subspace: Low-Rank Bandits Beyond Stationarity. *arXiv:2605.20269*, 2026. Under review at *ICLR 2027*. [paper] [code]
10. **Hamed Khosravi**, Xiaoming Huo. When Is One Reading Enough? Decision-Relevant Ambiguity in Constrained Reinforcement Learning. Under review at *ICLR 2027*. [paper] [code]
11. **Hamed Khosravi**, Milad Khanchi, Mobina Noori, Srinjoy Das, Abdullah Al-Mamun, Imtiaz Ahmed. Binary

Gaussian Copula Synthesis: An LLM-Powered Data Augmentation Framework for Early Dialysis Prediction in Chronic Kidney Disease. *arXiv:2403.00965*, 2024. Under review at *PLOS ONE*. [paper]

Peer-Reviewed Journal Papers

12. **Hamed Khosravi**[*], Mohammad Reza Shafie[*], Morteza Hajiabadi[*], Mobina Noori, Imtiaz Ahmed. KANGURA: Kolmogorov–Arnold Network-Based Geometry-Aware Learning with Unified Representation Attention for 3D Modeling of Complex Structures. *Manufacturing Letters*, 2026. [paper]
13. **Hamed Khosravi**, Ahmed Shoyeb Raihan, Farzana Islam, Ashish Nimbarte, Imtiaz Ahmed. A Comprehensive Approach to CO_2 Emissions Analysis in High-Human-Development-Index Countries Using Statistical and Time Series Approaches. *Sustainability*, 17(2), 2025. [paper]
14. **Hamed Khosravi**, Taofeeq Olajire, Ahmed Shoyeb Raihan, Imtiaz Ahmed. A Data-Driven Sequential Learning Framework to Accelerate and Optimize Multi-Objective Manufacturing Decisions. *Journal of Intelligent Manufacturing*, 35(8), 2024. [paper]
15. **Hamed Khosravi**, Sarah Farhadpour, Manikanta Grandhi, Ahmed Shoyeb Raihan, Srinjoy Das, Imtiaz Ahmed. Strategic Data Augmentation with CTGAN for Smart Manufacturing: Enhancing ML Predictions of Paper Breaks in Pulp-and-Paper Production. *Manufacturing Letters*, 41, 2024. [paper]
16. **Hamed Khosravi**, Imtiaz Ahmed, Avishek Choudhury. Predicting Suicidal Ideation, Planning, and Attempts among the Adolescent Population of the United States. *Healthcare*, 12(13), 2024. [paper]
17. **Hamed Khosravi**, Mohammad Reza Shafie, Morteza Hajiabadi, Ahmed Shoyeb Raihan, Imtiaz Ahmed. Chatbots and ChatGPT: A Bibliometric Analysis and Systematic Review of Publications in Web of Science and Scopus Databases. *International Journal of Data Mining, Modelling and Management*, 16(2), 2024. [paper]
18. **Hamed Khosravi**, Babak Amiri, Navid Yazdanjue, Vahide Babaiyan. An Improved Group Teaching Optimization Algorithm Based on Local Search and Chaotic Map for Feature Selection in High-Dimensional Data. *Expert Systems with Applications*, 204, 2022. [paper]
19. Mohammad Reza Shafie, **Hamed Khosravi**[†], Sarah Farhadpour, Srinjoy Das, Imtiaz Ahmed. A Cluster-Based Human Resources Analytics for Predicting Employee Turnover Using Optimized Artificial Neural Networks and Data Augmentation. *Decision Analytics Journal*, 11, 2024. [paper]
20. Avijit Saha Asru, **Hamed Khosravi**, Imtiaz Ahmed, Abdullahil Azeem. From Automation to Autonomy in Smart Manufacturing: A Bayesian Optimization Framework for Modeling Multi-Objective Experimentation and Sequential Decision Making. *The International Journal of Advanced Manufacturing Technology*, 137(9–10), 2025. [paper]
21. Mohammad Mahdi Barati Jozan, **Hamed Khosravi**, Aynaz Lotfata, Krzysztof J. Cios, Hamed Tabesh. Predicting Firefighter Injury and Entrapment in Urban Firefighting Operations: An Investigation Into the Effectiveness of Modified Fire Time Stages and Machine Learning. *Health Science Reports*, 8(3), 2025. [paper]
22. Ahmed Shoyeb Raihan, **Hamed Khosravi**, Srinjoy Das, Imtiaz Ahmed. Accelerating Material Discovery with a Threshold-Driven Hybrid Acquisition Policy-Based Bayesian Optimization. *Manufacturing Letters*, 41, 2024. [paper]
23. Ahmed Shoyeb Raihan, **Hamed Khosravi**, Tanveer Hossain Bhuiyan, Imtiaz Ahmed. An Augmented Surprise-Guided Sequential Learning Framework for Predicting the Melt Pool Geometry. *Journal of Manufacturing Systems*, 75, 2024. [paper]

Presentations

Talks

- Statistical Foundations for Credible Human–AI Teaming Claims — TEAM-Audit: A Two-Endpoint Audit of Margins. Workshop on Human–AI Teaming (HAT), HCOMP/CI 2026, Alexandria, VA, Sep 2026.
- Bridging Misspecification Gaps in Linear Contextual Bandits. Ph.D. Research Seminar, Georgia Institute of Technology, Feb 2026.
- Attention Meets Geometry: Advancing 3D Object Learning with Balanced Local and Global Enhancements. INFORMS QSR Data Competition (Finalist, **3rd place**), INFORMS Annual Meeting, Seattle, Oct 2024.
- LNUCB-TA: Linear–Nonlinear Hybrid Bandit Learning with Temporal Attention. INFORMS Annual Meeting, Seattle, Oct 2024.
- Binary Gaussian Copula Synthesis: Pioneering Data Augmentation for Early Dialysis Intervention in CKD Patients. INFORMS Annual Meeting, Seattle, Oct 2024.
- Strategic Data Augmentation with CTGAN for Smart Manufacturing: Enhancing Machine Learning Predictions of

Paper Breaks in Pulp-and-Paper Production. North American Manufacturing Research Conference (NAMRC 52), Knoxville, Jun 2024 (NSF Travel Award).
- AI-Driven Innovations: Transforming Healthcare and Manufacturing. DSCI 101 (Introduction to Data Science), West Virginia University, May 2024.
- Artificial Intelligence in Chronic Kidney Disease: Uncovering Patient Phenotypes. Healthcare Systems Process Improvement Conference (HSPI), Atlanta, Feb 2024.
- A Data-Driven Sequential Learning Framework to Accelerate and Optimize Multi-Objective Manufacturing Decisions. INFORMS Annual Meeting, Phoenix, Oct 2023.

Posters
- Risk-Controlled Decision-Making for Large Language Models. ML@GT Conference, Georgia Institute of Technology, Oct 2026.
- Attentive Optimism in the Face of Uncertainty: Balancing Exploration–Exploitation in Multi-Armed Bandits. INFORMS QSR Poster Competition, INFORMS Annual Meeting, Seattle, Oct 2024.
- Binary Gaussian Copula Synthesis: A Novel Data Augmentation Technique to Advance Machine Learning Prediction of Dialysis among CKD Patients. INFORMS Poster Competition (**3rd place**, 100+ posters), INFORMS Annual Meeting, Seattle, Oct 2024.
- Early Prediction of Dialysis among Chronic Kidney Disease Patients Using Advanced Machine Learning Models with Data Augmentation. Research Week Student Research Poster Symposium, WVU Statler College, Apr 2024.
- Efficient Frontier: Introducing an Active Learning Bayesian Framework for Cost-Effective Multi-Objective Manufacturing Optimization. INFORMS Annual Meeting, Phoenix, Oct 2023.
- Inclusive AI in Chronic Kidney Disease: Generative Adversarial Networks on Patient Care. All Voices as One! conference, WVU Statler College, Sep 2023.
- From Automation to Autonomy in Smart Manufacturing Using Sequential Learning. Graduate Student Research Symposium, WVU Statler College, Mar 2023.

Honors & Awards

National & Professional-Society Awards
- **Second Place**, K-12 STEM Spark Poster Competition, **INFORMS**, Oct 2026 — *Your Feed Changed. Can AI Keep Up?*
- **Finalist**, QSR Best Student Poster Competition, **INFORMS**, Oct 2026 — *Conformal Selective Acting: Anytime-Valid Risk Control for RLVR-Trained LLMs.*
- **Most Innovative Approach Award**, **NSF** Future Manufacturing Data Challenge, Jul 2026.
- **Finalist**, PG&E Energy Analytics Challenge Competition, **IISE**, May 2025.
- **NSF Travel Award** to attend the IISE Annual Conference, Apr 2025.
- **3rd Place**, Manufacturing AI Competition, **INFORMS** QSR Section, Oct 2024.
- **3rd Place**, Poster Competition (100+ posters), **INFORMS** Annual Meeting, Oct 2024.
- **NSF Travel Award** to attend NAMRC 52, May 2024.

University Fellowships & Honors
- Student Government Association (SGA) Conference Fund Travel Award, **Georgia Institute of Technology**, Aug 2026.
- College of Engineering (COE) Travel Award, **Georgia Institute of Technology**, Aug 2026.
- **John Morris Fellowship**, **Georgia Institute of Technology**, Aug 2025 and Jan 2026.
- **Excellence in Research Award**, **West Virginia University**, Nov 2024.
- **Benjamin M. Statler College of Engineering and Mineral Resources Fellowship**, **West Virginia University**, Aug 2023.
- **Distinguished Student**, B.S. degree, **Iran University of Science and Technology**, Feb 2022.

Industry Experience

Tesla — Machine Learning Engineer Intern Jan – Aug 2025

Cell Engineering, Analytics & Software · California, USA

- Developed a modular Python library for automated machine learning workflows, integrating data preprocessing, model selection, interpretability, and reporting; built to scale within the analytical ecosystem for fully automated classification, regression, and time series modeling.
- Designed and deployed ML pipelines on Kubeflow, using its Kubernetes-based infrastructure to run and manage containerized steps for scalable, reliable model training and evaluation.
- Integrated conformal prediction and contrastive learning into analytics workflows to improve anomaly scoring reliability and decision support.
- Applied machine learning to real-world manufacturing data, collaborating with software and quality teams to balance model performance, interpretability, and practical constraints.

Technical Skills

Languages	Python, R, SQL, Bash
ML & deep learning	PyTorch, Hugging Face Transformers and PEFT (LoRA), scikit-learn, XGBoost, LightGBM, statsmodels
Methods	Contextual bandits and reinforcement learning (Gymnasium); conformal prediction and anytime-valid inference; Bayesian optimization and Gaussian processes; causal and off-policy evaluation; generative data augmentation (CTGAN, Gaussian copulas); Kolmogorov–Arnold networks; LLM post-training (RLVR), agent scaffolding, and evaluation
Scientific computing	NumPy, pandas, SciPy, CVXPY, Matplotlib
MLOps & tooling	Kubeflow, Kubernetes, Docker, Git, pytest, Streamlit

Selected Coursework

Georgia Institute of Technology	Machine Learning; Deep Reinforcement Learning; Mathematical Foundations of Machine Learning; Theory of Statistics; Advanced Nonlinear Programming; Convexity
West Virginia University	Pattern Recognition (Computer Science); Math for Machine Learning (Math and Data Sciences); Design of Experiments; Applied Linear Programming (Industrial Engineering); Statistical Methods (Statistics); Business Statistical Methods (Business Data Analytics)
Iran Univ. of Science & Technology	Engineering Data Analytics; Computer Programming (Industrial Engineering)

Service & Leadership

- **Co-Session Chair, INFORMS Annual Meeting** — Led and moderated technical sessions, facilitating discussions on cutting-edge research in decision sciences and machine learning.
- **Conference Service, INFORMS Annual Meeting** — Contributed to peer engagement and scholarly exchange across machine learning and decision analytics tracks.
- **Academic Reviewer** — Reviewer for more than 15 journals and conferences, including NeurIPS 2026, *Scientific Reports*, *Expert Systems with Applications*, *Artificial Intelligence Review*, *Discover Artificial Intelligence*, *Signal, Image and Video Processing*, *Neural Processing Letters*, *Machine Learning Research*, and *Machine Learning with Applications*; 40+ peer reviews completed.
- **Graduate Mentorship & Supervision** — Guided three master's students through research design and execution to successful project completion.
- **Research Team Representative, INFORMS QSR Data Competition, 2024** — Represented and presented finalist research at a national-level data competition.
- **Invited Academic Speaker, DSCI 101 (Introduction to Data Science), 2024** — Delivered a talk on AI-driven innovations in healthcare and manufacturing.
- **Curriculum Development** — Enhanced the *Advanced Analysis of Engineering Data* course project (R programming), motivating undergraduate interest in applied research.