

Lekhapriya Dheeraj Kashyap

lekha.dk@tamu.edu · (+1) 908-294-0183 · [Website](#) · [Github](#) · [Google scholar](#) · [Linkedin](#)

Summary

AI/ML researcher in Human-Centered AI with a strong quantitative foundation in stochastic optimization, machine learning, operations research, computational and foundational models. Research focuses on interpretable, learning-based decision models (MDPs, POMDPs, semi-Markov processes) to infer latent states, preferences, and policies from behavioral data. Experience applying reinforcement learning and multi-agent learning to human-centered systems, intelligent vehicle systems, finance, manufacturing, power grid operations, and data center operations, with industry experience at Wells Fargo and Hewlett Packard Enterprise.

Education

Texas A&M University PhD in Industrial & Systems Engineering	College Station, TX 2021 – 2026
New Jersey Institute of Technology M.S. in Industrial & Engineering Management	Newark, NJ 2014 – 2016
Visvesvaraya Technological University B.Eng. in Electronics and Communication	Bangalore, India 2010 – 2014

Research Experience

Texas A&M University <i>PhD Student, Industrial & Systems Engineering</i>	College Station, TX 2021 - present
---	---------------------------------------

- **State Representation and Preference Learning from Demonstrations** (*In preparation*, [Github](#))
Developed deep learning methods to learn compact latent representations from high-dimensional, multi-modal human behavioral data. Jointly inferred system dynamics, agent preferences, and utility functions from expert demonstrations to support interpretable human-AI decision modeling.
- **Modeling Multitasking Behavior and Driver Distraction with Structural Inverse Learning under Partial Observability** ([Published](#), [Human Factors](#), [Github](#))
Formulated driver attention and multitasking as a Partially Observable Semi-Markov Decision Process and applied structural inverse reinforcement learning to infer latent distraction states, belief dynamics, and infer reward functions from behavioral data. Enabled personalized modeling and evaluation of in-vehicle human-AI interaction and intervention strategies.
- **Latent-State Modeling of Attention and Decision-Making from Behavioral Data** ([Published](#), [Human Factors](#), [Github](#))
Developed probabilistic latent-state models to infer unobserved context, attentional effort, and reward trade-offs from large-scale, naturalistic human decision data. Applied Bayesian inference in partially observable settings to support interpretable evaluation of human performance, bias, and AI-assisted decision policies.
- **Multi-Agent Reinforcement Learning for Decentralized Power Systems** ([Github](#))
Applied multi-agent reinforcement learning to decentralized power systems, addressing coordination, scalability, and stochastic optimization under uncertainty while preserving local operational autonomy.
- **Reinforcement Learning with Digital Twins for Precision Manufacturing** ([Under review](#))
Developed an actor-critic (PPO) framework for constrained decision-making in multi-stage manufac-

turing. Integrated stochastic digital twins to enable adaptive control and joint optimization of process parameters to achieve precise surface roughness targets while minimizing energy and operational costs.

Indian Institute of Science

Research Associate, Department of Management Studies

Bangalore, India

2019 - 2021

- **Scalable Learning of Hawkes Processes for Event Modeling** (*Published, Journal of Computational Science, Github*)

Derived unbiased stochastic gradient estimators for maximum likelihood optimization of multi-dimensional Hawkes processes with neural network parameterization of excitation kernels. Validated the approach on real-world finance data, demonstrating strong performance in high-frequency event modeling.

Work Experience

Hewlett Packard Enterprise

Research Associate Intern, AI-RL Lab

Milpitas, CA

2023

- **SustainDC: Benchmarking for sustainable data center control** (*Published, Neurips*)
Developed configurable benchmarking environments for multi-agent reinforcement learning in data center operations, modeling coordinated workload scheduling, cooling, and energy storage under shared constraints. Evaluated MARL algorithms to analyze system behavior, energy efficiency, and carbon-aware control strategies for real-world infrastructure.
- **PyDCM: Custom Data Center Models with Reinforcement Learning for Sustainability** (*Published, BuildSys*) Built a high-performance, configurable data center simulator with a Gymnasium-compatible interface, enabling rapid prototyping and evaluation of reinforcement learning-based cooling strategies for sustainable system design.

WellsFargo

Quantitative Intern, Market & Counterparty Risk Analytics

Charlotte, NC

2022

- Leveraged Bloomberg BVAL sector- and rating-level time series to model bond portfolio risk. Implemented Python-based risk engines to compute and assess Value at Risk (VaR) and SVaR. Built SQL pipelines for daily P&L computation and performed P&L attribution and backtesting against BVAL-derived risk factors. Analyzed and reported key drivers of material deviations relative to production benchmarks to senior management.

Teaching Experience

- Quality Engineering (TA) – *Fall 2021, Spring 2022, Fall 2022, Spring 2023, Fall 2025*
- Uncertainty Modeling for Industrial Engineering (TA) – *Fall 2023*
- Decision and Risk Analysis (TA) – *Spring 2025*

Skills

Programming: Python, MATLAB, R, C++, SQL

Libraries/Frameworks: PyTorch, TensorFlow, Gymnasium, pandas, scikit-learn, Hugging Face, RL-Lib

Tools: Git, Docker, Linux