MAANAV SIKARIA

(425) 429-8114 • maanavsikaria@gmail.com • linkedin.com/in/maanav-sikaria • github.com/Maanav-s

EDUCATION

The University of Texas at Austin – Austin, TX
Bachelor of Science in Electrical and Computer Engineering **May 2028**

- GPA: 3.71
- Relevant Coursework: Linear Algebra, Multivariable Calculus, Differential Equations, Linear Systems and Signals, Data Structures, Embedded Systems, Probability and Random Processes, Discrete Math, Operating Systems

EXPERIENCE

CLeAR Lab, The University of Texas at Austin – Austin, TX **05/2026 – Present**
Undergraduate Research Assistant

- Implemented the REINFORCE, GRPO, and PPO reinforcement learning algorithms from scratch in PyTorch, validating PPO by training a policy to a 92% success rate on OpenAI Gymnasium's Lunar Lander
- Designed reward functions and training curriculums for humanoid robots in IsaacLab simulation
- Implemented and trained Mamba (state-space) architecture to encode a history-aware world model latent state encoder for humanoid control

Texas Undergraduate Computational Finance (TUCF) – Austin, TX **01/2026 – Present**
Quantitative Analyst

- Built a live tennis win-probability model (MLP + recursive match-outcome solver) on scraped match data, scoring 0.1777 Brier against the bookmaker's 0.1694 on the same matches
- Built an automated mention-market trading system with real-time speech recognition and API-based order execution, at 322 ms end-to-end from spoken mention to submitted order

SELECTED PROJECTS

Dietary-Menu LLM Agent — fine-tuned a 4B Gemma 4 model (LoRA context-distillation SFT)

- Built a tool-based agent that finds a restaurant's menu online, filters it for dietary restrictions, and returns structured JSON, using a search API, headless Chromium scraper, and SQLite cache
- Distilled from a 235B Qwen3 teacher via LoRA context-distillation SFT on a scraped corpus of thousands of menus, then tuned task-specific GRPO reinforcement-learning rewards atop the SFT checkpoint
- Lifted menu-find accuracy from 74.1% to 81.3% against the teacher's 82.0% at 59x fewer parameters with 66% fewer hallucinated menus, 2.4x throughput, and 28% F1 lift over base
- Designed SQLite schema to manage training data collection and cleaning, built 500 item eval harness from data

Variational Autoencoder — from scratch (Python, NumPy)

- Implemented the full model from scratch — hand-derived backpropagation and gradient updates in pure NumPy, with no PyTorch, or auto-grad library
- Sampled the latent space to generate images (Fashion MNIST) and analyzed latent structure for disentanglement and KL divergence. Studied dead neurons and hyperparameter convergence of trained model

PRE-COLLEGE LEADERSHIP

- FRC Robotics, Strategy & Finance Lead (2021–2025) — lead on a ~50-person team on an ~$80K budget, raised $50,000+ heading a 5-person grants team, and mentored 40+ teammates to a World Championship appearance
- Redmond STEM Center, Executive Director (2021–2025) — raised $10,000+ over 1,000+ student-hours of instruction, oversaw operations for 30+ volunteer, partnered with the City of Redmond to broaden STEM access

SKILLS

- Languages: Python, C, C++, Java, JavaScript, SQL, ARM Assembly
- Data: PyTorch, TensorFlow, Hugging Face Transformers, NumPy, Pandas, Matplotlib, OpenCV, XGBoost, HMMs
- ML/RL: LLM fine-tuning (LoRA, SFT, context distillation), GRPO, PPO, reward design, quantization, Mamba/SSMs
- Methods: Deep learning, computer vision, NLP, model evaluation & calibration, web scraping
- Tools: AWS, RunPod, Linux, Docker, Git, REST APIs