BASAB GHOSH

LinkedIn — GitHub — Website — Email — Bengaluru, India — 6295799403

Research Interests

Reinforcement Learning, LLMs and Agents, Representation Learning, Interpretability, Risk-Sensitive Decision Making, NLP

Education

M.Tech (Research), Computer Science & Automation *July 2025 – Present*
Indian Institute of Science (IISc), Bangalore
- Working on Non-stationary RL and preference changes in LLMs through the lens of RL

B.Tech, Computer Science & Engineering *Dec. 2021 – May 2025*
Indian Institute of Information Technology, Design & Manufacturing (IIITDM), Kancheepuram **GPA: 8.69/10**

Research & Industry Experience

AI Research, cmul8 *2026 – Present*
- Building multi-agent systems and societies of AI agents to transform industries at scale, leading research on agent coordination, reliability, and evaluation.

ML Intern, FastCode AI *Apr. 2026 – July 2026*
Remote
- Worked on multi-agent architectures and social-environment simulation for a national road agency.

Research Intern, IIIT Hyderabad *Feb. 2025 – Aug. 2025*
- Studied risk-sensitive reinforcement learning beyond expected return, combining risk objectives with optimal stopping.
- **Manuscript under review at AAAI 2027.**

Research Intern, IIT Guwahati *May 2024 – July 2024*
- Implemented PPO and TRPO for drone navigation and aerial base-station placement in continuous control environments.

Research Projects

Eval-Awareness in LLMs
- Investigated whether evaluation awareness is encoded in LLM activations using linear probes and controlled interventions.
- Showed that strong benchmark performance largely reflected dataset-specific artifacts rather than evaluation awareness.

Spotlight: Multi-Agent AI Security Engineer
- Built a multi-agent security system with sandboxed autonomous reproduction, verifiable action provenance, and human approval.

When Five Agents Behave Like One: Multi-Agent LLM Systems
- Studied correlated errors in LLM committees, showing that five-agent committees provide only 1.2–1.8 effective independent voters and motivating a cheaper calibration-based committee strategy.

A TiH-IoT funded project on "Design and Development of Microservice-Based Fog-Enabled Infrastructure for Smart Agriculture"
- Worked in the intersection of drone programming and computer vision for agricultural purposes, specifically drone navigation and detection and automating task.

interpretability in POMDPs
- Studying internal representations and hidden-state dynamics of DRQN agents in partially observable environments.
- GitHub

Publication

DPO Kernels: A Semantically-Aware, Kernel-Enhanced, and Divergence-Rich Paradigm for Direct Preference Optimization *ACL 2025*

Technical Skills

Areas: Reinforcement Learning, Deep Learning, LLMs, NLP, Representation Learning
Languages: Python, C++, C **Frameworks:** PyTorch, TensorFlow, Scikit-learn, NumPy, Pandas