Utkarsh Tiwari

Email | GitHub | Website

Research Interests

Computational Expressivity of Neural Networks, Machine Learning Theory, Systems, Interpretability

Education

- **Birla Institute of Technology and Science, Pilani** — Rajasthan, India
  Bachelor of Engineering – Computer Science; GPA: 9.08 — *Oct 2022 – July 2026 (Expected)*

Summer School

- **Approximation Algorithms**, IISc Bangalore – Advanced Undergraduate/Graduate Summer School, May 2025

Research Experience

- **Microsoft Research, India** — Bangalore, India
  Research Intern — *Aug 2025 – Jan 2026*
  - **Transformer Expressivity Bounds**: Worked with Dr. Amit Deshpande on theoretical limits of transformers for graph algorithms and on the expressivity of neural networks with oracle access formalizing function calls in LLMs
  - **Systems for ML and GPU Optimization**: Worked in Prof. Manik Varma's group on distributed training and kernel-level optimizations, increasing memory efficiency by over 12x allowing for efficient billion parameter training runs
- **STARS Group, INRIA Université Côte d'Azur** — Nice, France
  Research Intern — *Jul 2024 – Aug 2024*
  - **Weakly Supervised Traffic Anomaly Detection**: Curated weakly supervised datasets, benchmarked SoTA models, and designed transformations modules enhancing SoTA. Published as a workshop paper **(ECCV 2024 ROAM Oral)**

Publications

- **Utkarsh Tiwari**, Sai Soumya Nalli, Amit Deshpande. Modeling Tool Use in Transformers via Computation Oracles, 2026. ***Under Review***
- **Utkarsh Tiwari**, Aviral Gupta, and Michael Hahn. Emergent stack representations in modeling counter languages using transformers, 2025 ***World Models Workshop, ICLR 2025*** [arXiv]
- **Utkarsh Tiwari**, Aryan Seth, Adi Mukherjee, Kaavya Mer, Kavish, Dhruv Kumar. DebateBench: A Challenging Long-Context Reasoning Benchmark for Large Language Models, 2025. ***Student Research Workshop, NAACL 2025*** [arXiv]
- **Utkarsh Tiwari**, Snehashis Majhi, Michal Balazia, and François Brémond. What matters in autonomous driving anomaly detection: A weakly supervised horizon. *Springer*, 2024 ***Oral, ROAM Workshop at ECCV 2024*** [arXiv]

Teaching Assistantships

- **Discrete Mathematics** – Created material and assisted students in algorithms, formal logic and proof techniques.
- **Logic in Computer Science** – Created material for FO, temporal logic, Prolog, and program verification
- **Computer Systems and Performance** – Created material and assisted over 2000 students as the lead TA

Industry Experience

- **Pavo AI** — Machine Learning Intern, March 2025 – May 2025
  Agentic Dataset Creation: Designed agentic systems for automated retrieval and summarisation of research papers.
- **Summer Intern** — Machine Learning Intern, May 2024 – June 2024
  Software Development and Optimisation: Worked with Bihar govt. to optimize community healthcare technology.

Reviewing and Academic Talks

- **Reviewer -** World Models: Understanding, Modelling and Scaling (**ICLR 2025**)
- Talk on **Formal Theory for Interpretability** at GenAI Research Group, BITS Pilani - Slides, Recording
- **ICLR 2025 Workshop Oral presentation** on weakly supervised anomaly detection - Slides

Personal Projects

- **Solving Leetcode in Lambda Calculus** : Blog Post, Code
- **Triton Kernels for Faster Graphics** : Blog Post
- **Ray Tracer in C++** : Blog Post, Code

Awards and Distinctions

- 100% Merit Scholarship on College Tuition – Semester 1 (2022–23)
- BITSAT 2022 Score: 360/390 (Top 80 All-India Rank)

Extracurricular

- **Amateur Radio Hosting**: Programmed and hosted the college radio with regular academic guest interviews
- **Debating**: Active member of BITS Pilani Debating Society; multiple national open breaks.
- **Technical Blogging**: Maintain a regular blog covering ML research, systems, and theoretical CS.