Utkarsh Tiwari

Email | GitHub | Website

Research Interests

Machine Learning Theory, Theoretical CS, Computation and Complexity Theory

Education

- **Toyota Technological Institute, Chicago** — Chicago, Illinois
  Ph.D., Machine Learning and CS Theory — *Incoming, 2026*
- **Birla Institute of Technology and Science, Pilani** — Rajasthan, India
  Bachelor of Engineering – Computer Science; GPA: 8.95 — *Oct 2022 – July 2026*

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

- **On the Hardness of "Infinite Energy" Puzzles**, 2026. **Utkarsh Tiwari**, Utkarsh Thanvi, Armaan Sethi, Aalok Thakkar, Neeldhara Misra.
  ***Under Review**, early results presented at JCDCG^3 2026.*
- **Modeling Tool Use in Transformers via Computation Oracles**, 2026. **Utkarsh Tiwari**, Sai Soumya Nalli, Amit Deshpande.
  ***Under Review**, early results presented at Workshop on Latent & Implicit Thinking, ICLR 2026, and FLaNN workshop at Yale University*
- **Improving Reachability on Reasoning Puzzles**, 2026. Sukruta Prakash Midigeshi, Sai Soumya Nalli, **Utkarsh Tiwari**, Amit Deshpande, Nagarajan Natarajan, Vineeth N. Balasubramanian, Amit Sharma, Gaurav Sinha.
  Workshop on LLM Reasoning, ICLR 2026. [Openreview]
- **Emergent Stack Representations in Modeling Counter Languages using Transformers**, 2025. **Utkarsh Tiwari**, Aviral Gupta, Michael Hahn.
  World Models Workshop, ICLR 2025. [arXiv]
- **DebateBench: A Challenging Long-Context Reasoning Benchmark for Large Language Models**, 2025. **Utkarsh Tiwari**, Aryan Seth, Adi Mukherjee, Kaavya Mer, Kavish, Dhruv Kumar.
  Student Research Workshop, NAACL 2025 [arXiv]
- **What Matter in Autonomous Drivinf Anomaly Detection: a Weakly Supervised Horizon**, 2024. **Utkarsh Tiwari**, Snehashis Majhi, Michal Balazia, Francois Bremond.
  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.

Research Talks

- 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