SUCHIT GUPTE

gupte.31@osu.edu
website: suchitgupte.github.io

Education

2023-Present **PhD CSE at The Ohio State University**.
PhD in Computer Science
- *Coursework*: Computer Vision, Data Visualization, Advanced NLP, Advanced Algorithms, Data Mining, Machine Learning, CyberSecurity.

2021-2023 **Visiting Student at the University of Pennsylvania**.
- *Coursework*: Machine Perception, Applied Machine Learning, Database Management Systems, Engineering Probability, Data Structures & Algorithms

2021-2023 **BS CS at Rutgers University–Camden**, *GPA: 3.98/4*.
Bachelors in Computer Science
- *Coursework*: Operating Systems, Software Engineering, C & Unix Systems Programming.

2019-2021 **BSc CS at Mumbai University**, *GPA: 10/10*.
Bachelors of Science in Computer Science - Transferred to Rutgers University–Camden in Fall 2021
- *Coursework*: Operating Systems, Web Programming, Computer Networks, Software Engineering, Computer Graphics & Animation, Mobile Programming.

Publications and Presentations

arXiv **On the transferability of Sparse Autoencoders for interpreting compressed models**,
Suchit Gupte, Vishnu Kabir Chhabra, Mohammad Mahdi Khalili.
arXiv preprint arXiv:2507.15977

SIGMOD 2025 **Understanding the Black Box: A Deep Empirical Dive into Shapley Value Approximations for Tabular Data**,
Suchit Gupte, John Paparrizos.
ACM SIGMOD/PODS International Conference on Management of Data, June 22-27, 2025 Berlin, Germany

SIGMOD 2025 **ShapX Engine: A Demonstration of Shapley Value Approximations**,
Suchit Gupte, John Paparrizos.
ACM SIGMOD/PODS International Conference on Management of Data, June 22-27, 2025 Berlin, Germany

EnCORE 2022 **Detecting Invisible Keypoints**, *University of California, San Diego*,
Suchit Gupte, Nimish Jayakar, Zanming Huang, Eshed Ohn-Bar, Ruizhao Zhu.
Presented our work at the Institute for Emerging CORE Methods in Data Science.

Select Research Experience

2025-Present **Interpreting Low-Rank Compressed GPT-2-like models using SAEs**.
The Ohio State University | Graduate Research Associate
- *Mentor*: Prof. Mohammad Mahdi Khalili
- *Domain*: Mechanistic Interpretability
- *Work*: Exploring the effects of low-rank compression on the internal representations of GPT -2-like models using Sparse Autoencoders. We focus on identifying changes in feature activations across layers and understanding the interpretability trade-offs introduced by compression.

2025 **Off-Policy Reinforcement Learning**.
Amazon | Applied Science Intern at the AGI LLM Foundations team
- *Mentor*: Hadrien Glaude
- *Domain*: Agentic Reinforcement Learning
- *Work*: Worked on reinforcement learning for large language models, focusing on building and improving multi-turn RL training pipelines to evaluate new algorithms across long-horizon tasks.

2023-2024 **Shapley Value Explanations**.
The Ohio State University | Graduate Research Associate
- *Mentor*: Prof. John Paparrizos
- *Domain*: Data-centric AI, Interpretability
- *Work*: A benchmarking survey of Shapley value approximation methods for feature attribution in tabular data. The study evaluates their accuracy, computational efficiency, and trade-offs across different datasets and tasks. This work provides a comprehensive comparison to guide the selection of efficient and reliable Shapley-based interpretability methods.

2022-2023 **Detecting Invisible Keypoints**.

Boston University | Undergraduate Research Intern

- *Mentor*: Prof. Eshed Ohn-Bar
- *Domain*: Computer Vision, 3D modeling
- *Work*: Developed a large-scale synthetic benchmark in CARLA by incorporating 62 vehicle keypoints across diverse perspectives, towns, and weather conditions. Trained and analyzed ViTPose for fine-grained keypoint detection under severe occlusion and truncation. Proposed a novel geometric loss over a gradual training process, improving the average precision by 16%.

Select Teaching Experience

Graduate Teaching Assistant

2025 **CSE 5234: Data Mining**.
2025 Spring: instructed by Prof. Thomas Bihari

2024 **CSE 5523: Machine Learning**.
2024 Fall: instructed by Prof. Mohammad Mahdi Khalili

2024 **CSE 3241: Introduction to Database Systems**.
2024 Spring: instructed by Prof. Jeffery Eden

2023 **CSE 1223: Introduction to Java Programming**.
2023 Fall: instructed by Prof. Peter Gerstmann

Undergraduate Teaching Assistant

2025 **CS 213: Data Structures & Algorithms**.
2023 Spring: instructed by Prof. Rajiv Gandhi

2024 **CS 171: Mathematical foundations of Computer Science**.
2022 Fall: instructed by Prof. Rajiv Gandhi

Outreach

2022 **Program in Algorithmic and Combinatorial Thinking**.
2022 Summer: instructed by Prof. Rajiv Gandhi

- Mentored a small group of students in an algorithmic and combinatorial thinking program, fostering problem-solving skills and mathematical reasoning. Guided them through complex concepts, algorithm design, and combinatorial problem-solving techniques. Encouraged analytical thinking and structured approaches to tackle challenging problems.

Technical Skills

Languages	Python, C++, C, Java, SQL, MATLAB, Bash
DL Tools	PyTorch, TensorFlow, Keras, OpenCV, Hugging Face, Scikit-learn, CUDA
Web-Dev	HTML, CSS, JavaScript, TypeScript, PHP, React.js, Node.js, StreamLit
Databases	Oracle DB, MongoDB, MySQL
Deployment	AWS (EC2, S3, Lambda, SageMaker), Docker, Kubernetes
Visualization	NumPy, Pandas, Matplotlib, Seaborn, Tableau
Others	Linux, Git, GitHub, Visual Studio, VS Code, Jupyter, LaTeX

Achievements

2023 Recipient of the Computer Science Academic Achievement Award, Rutgers–Camden
2023 Dean's list Spring 2023, Rutgers–Camden.
2022 Dean's list Fall 2022, Rutgers–Camden.
2022 Dean's list Spring 2022, Rutgers–Camden.
2021 Dean's list Fall 2021, Rutgers–Camden.

Select Projects

2024 **Examining the Reversal Curse on Logical Equivalence**.

- *Technologies*: LLMs
- *Course*: Advanced NLP, CSE 5539
- *Description*: The Reversal Curse paper (link) highlights a simple task that these models fail at. If the model has seen A is B, it is not guaranteed that the model can generalize B is A - this is coined as Reversal Curse in the paper. Expanding on this, my work investigates whether this limitation extends to logical equivalences—specifically, whether models trained on "A implies B" can infer "not B implies not A." This research evaluates the model's ability to reason beyond memorization.

2023 **BirdViz**, *Data Visualization Tool for Bird Sightings in US Ecological Centers*.

- *Technologies*: HTML, CSS, JavaScript, PyDeck
- *Course*: Data Visualization, CSE 5544
- *Description*: BirdViz is an interactive data visualization tool developed to explore and analyze bird sightings data. It utilizes datasets from the Macaulay Library and NEON Breeding Birds Dataset to provide interactive maps, species identification analysis, and graphical visualizations. Designed for researchers, conservationists, and bird enthusiasts, BirdViz transforms complex ecological data into meaningful insights for a better understanding of avian biodiversity.

2021 **Automatic Equation Images to LaTeX Math Conversion**.

- *Domain*: Computer Vision and Natural Language Processing
- *Course*: Applied Machine Learning, CIS 5190
- *Description*: Supervised by Prof. Dinesh Jayaraman (UPenn), trained an optical character recognition model pix2tex to convert images of mathematical expressions into accurate LaTeX code using a supervised learning approach.