

RESEARCH INTERESTS

Controllable Generation, Adversarial Robustness, Representation Learning, Interpretability

EDUCATION

Georgia Institute of Technology

Atlanta, GA

Ph.D. in Computer Science

Aug 2021 – Present

- **Advisor:** Prof. Srijan Kumar, Assistant Professor, CSE
- **Committee:** Prof. Srijan Kumar, Prof. Chao Zhang, Prof. Bo Dai, Dr. Rakshit Trivedi
- **Thesis Topic:** *Intention in tension:* Towards Robust and Controllable Machine Learning

Indian Institute of Technology, Delhi

New Delhi, India

Bachelor of Technology in Computer Science and Engineering. GPA: 9.0/10.0

Jul 2017 – May 2021

- **Advisor:** Prof. Sayan Ranu, Associate Professor, CSE
- **UG Thesis topic:** Adversarial attacks on Graph Neural Networks

WORK EXPERIENCE

Amazon Science

Seattle, WA

Applied Scientist Intern

May 2025 – Aug 2025

- Designed SPECS to advance zero-shot LLM classification by evolving prompts with label-supporting justifications
- Accepted at Amazon's internal machine learning conference (AMLC 2025) prompt optimization workshop for oral

Microsoft Research

Redmond, WA

Research Intern

May 2024 – Aug 2024

- Developed OG-RAG, an efficient retrieval method on documents grounded in a pre-defined ontology
- OG-RAG gives higher accuracy of responses and faster fact verification by users

Visa Research

Palo Alto, CA

Research Intern

May 2023 – Aug 2023

- Devised a novel method to classify nodes using a node-specific scope of the local neighborhood
- Experiments show improved accuracy, robustness, and depth for graph neural networks

Tensor Dynamics Pvt. Ltd.

New Delhi, India

Founding Member

Sep 2018 – Jul 2021

- Developed initial code-base for a green-tech startup to provide reliable solar energy forecasts
- Forecasting models deployed by Load Dispatch Centre to monitor over 20 solar stations across India

PUBLICATIONS

- Inner Speech as Behavior Guides: Steerable Imitation of Diverse Behaviors for Human-AI coordination Rakshit Trivedi*, **Kartik Sharma***, David C. Parkes. *Advances in Neural Information Processing Systems (NeurIPS) Spotlight*, 2025
- OG-RAG: Ontology-Grounded Retrieval-Augmented Generation For Large Language Models **Kartik Sharma**, Peeyush Kumar, Yunqing Li. *Empirical Methods in Natural Language Processing (EMNLP) Main*, 2025
- *Who Speaks Matters:* Analysing the Influence of the Speaker's Ethnicity on Hate Classification Ananya Malik, **Kartik Sharma**, Shaily Bhatt, Lynnette Hui Xian Ng. *Empirical Methods in Natural Language Processing (EMNLP) Short Findings*, 2025
- Personalized Layer Selection for Graph Neural Networks **Kartik Sharma**, Vineeth Rakesh, Yingtong Dou, Srijan Kumar, Mahashweta Das. *Transactions on Machine Learning Research (TMLR)*, 2025
- Diffuse, Sample, Project: Plug-And-Play Controllable Graph Generation **Kartik Sharma**, Srijan Kumar, Rakshit Trivedi. *International Conference on Machine Learning (ICML)*, 2024
- A Survey of Graph Neural Networks for Social Recommender Systems. **Kartik Sharma***, Yeon-Chang Lee*, Sivagami Nambi, Aditya Salian, Shlok Shah, Sang-Wook Kim, Srijan Kumar. *ACM Computing Surveys (October 2024)*
- Mysterious Projections: Multimodal LLMs Gain Domain-Specific Visual Capabilities Without Richer Cross-Modal Projections. Gaurav Verma, Minje Choi, **Kartik Sharma**, Jamelle Watson-Daniels, Sejoon Oh, Srijan Kumar. *Annual Meeting of the Association for Computational Linguistics (ACL) Short*, 2024

- Representation Learning in Continuous-Time Dynamic Signed Networks. **Kartik Sharma***, Mohit Raghavendra*, Yeon Chang Lee, Anand Kumar M, Srijan Kumar. *ACM International Conference on Information and Knowledge Management (CIKM)*, 2023
- A survey on explainability of graph neural networks. Jaykumar Kakkad, Jaspal Jannu, **Kartik Sharma**, Charu Aggarwal, Sourav Medya. *IEEE Data Engineering Bulletin*, 2023
- Temporal Dynamics-Aware Adversarial Attacks on Discrete-Time Dynamic Graph Models. **Kartik Sharma**, Rakshit Trivedi, Rohit Sridhar, Srijan Kumar. *ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD), Research Track*, 2023 (also *NeurIPS Workshop on Temporal Graph Learning (TGL)*, 2022)
- Predicting Information Pathways Across Online Communities. Yiqiao Jin, Yeon-Chang Lee, **Kartik Sharma**, Meng Ye, Karan Sikka, Ajay Divakaran, Srijan Kumar. *ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD), Research Track*, 2023
- Task and Model Agnostic Adversarial Attack on Graph Neural Networks. **Kartik Sharma**, Samidha Verma, Sourav Medya, Sayan Ranu, and Arnab Bhattacharya. *AAAI Conference on Artificial Intelligence*, 2023
- Balance maximization in Signed Networks via Edge Deletions. **Kartik Sharma**, Iqra Altaf Gillani, Sourav Medya, Sayan Ranu, Amitabha Bagchi. *ACM International Conference on Web Search and Data Mining (WSDM)*, 2021
- Clause Final Verb Prediction in Hindi: Evidence for Noisy Channel Model of Communication. **Kartik Sharma**, Niyati Bafna, Samar Husain. *Cognitive Modeling & Computational Linguistics (CMCL) Workshop*, 2021
- What determines the order of verbal dependents in Hindi ? Effects of efficiency in production & comprehension. **Kartik Sharma**, Richard Futrell, Samar Husain. *Cognitive Modeling & Computational Linguistics (CMCL)*, 2020
- Can Greenbergian universals be induced from language networks ? **Kartik Sharma**, Kaivalya Swami, Aditya Shete, Samar Husain. *Treebanks and Linguistic Theories (TLT) Workshop*, 2019

PREPRINTS

- COLD-Steer: Steering Large Language Models via in-Context One-Step Learning Dynamics **Kartik Sharma**, Rakshit Trivedi. *Under Review*
- Efficient Knowledge Probing of Large Language Models by Adapting Pre-trained Embeddings **Kartik Sharma**, Yiqiao Jin, Rakshit Trivedi, Srijan Kumar. *arXiv:2508.06030*
- Sysformer: Safeguarding Frozen Large Language Models with Adaptive System Prompts **Kartik Sharma**, Yiqiao Jin, Vineeth Rakesh, Yingdong Dou, Menghai Pan, Mahashweta Das, Srijan Kumar. *arXiv:2506.15751*
- *A Thousand Words or An Image*: Studying the Influence of Persona Modality in Multimodal LLMs Julius Broomfield*, **Kartik Sharma***, Srijan Kumar. *arXiv:2502.20504*
- SARA: Selective and Adaptive Retrieval-augmented Generation with Context Compression Yiqiao Jin, **Kartik Sharma**, Vineeth Rakesh, Yingdong Dou, Menghai Pan, Mahashweta Das, Srijan Kumar. *arXiv:2507.05633*

HONORS & AWARDS

- **Outstanding Reviewer**: EMNLP 2024, KDD 2025 (Aug cycle)
- **Suresh Chandra Memorial Trust Award**: Best undergraduate thesis in CSE, IIT Delhi
- **Summer Undergraduate Research Award**: For credible research in summers 2019, IIT Delhi
- **Semester Merit Award**: Awarded in 2017 by IIT Delhi
- **All India Rank 197**: Joint Entrance Examination Advanced 2017 among 100k qualified candidates
- **All India Rank 250**: Joint Entrance Examination Mains 2017 among 1M candidates

TEACHING & VOLUNTEERING

- **PC Member/Reviewer**: ICML 2024-25; ICLR 2024-26; NeurIPS 2024-25; KDD 2022-2025; WWW 2025; ACL ARR 2024-25; SDM 2024; AAAI 2024,26; LOG 2023-24; TNNLS 2023; SPIGM@ICML 2023; TGL@NeurIPS 2022-23
- **Student volunteer**: AAAI 2023
- **Teaching Assistant, CSE 6240, Georgia Tech**: Course included in Course-Instructor Opinion Survey Honor Roll
- **Teaching Assistant, I-STEM**: For visually-impaired students in an online Data Structures course

SKILLS

Languages: Python, C++, R, OCaml, Java, MATLAB, SPARQL

Libraries: Langchain, LlamaIndex, Huggingface, Diffusers, TorchGeometric, TorchGeometricTemporal, Pytorch, Tensorflow, Sklearn, Scipy, Numpy, Pandas

REFERENCES

- **Prof. Srijan Kumar**: Assistant Professor, Georgia Institute of Technology
- **Dr. Rakshit Trivedi**: Postdoctoral Associate, Massachusetts Institute of Technology
- **Prof. Sayan Ranu**: Associate Professor, Indian Institute of Technology, Delhi (IIT Delhi)
- **Prof. Sourav Medya**: Assistant Professor, University of Illinois, Chicago (UIC)