Raj Sanjay Shah

rajsanjayshah@gmail.com | 6282809784 | TSRB 228, 85 5th St NW, Atlanta, GA 30332
www.linkedin.com/in/raj-sanjay-shah | Raj-Sanjay-Shah | https://raj-sanjay-shah.github.io/

About me

PhD student in Interactive Computing at Georgia Tech; working on natural language processing, evaluation, and robustness. My work spans benchmarking and domain adaptation (e.g., financial language modeling, mental health, etc.), robust NLP methods (unlearning, watermarking, continual learning), and human-centered applications supporting well-being. With publications in ACL, NAACL, EMNLP, COLM, and CSCW, and strong experience in large-scale model evaluation, applied NLP algorithms, and interdisciplinary research.

Education

Aug 2023 – **Georgia Institute of Technology**
*Ph.D in **Computer Science*** (GPA: 4.0/ 4.0)
Advisor: Professor Sashank Varma

Sep 2024 – June 2025 **Stanford University**
Visiting Student Researcher
Advisor: Professor Diyi Yang

Aug 2021 – May 2023 **Georgia Institute of Technology**
*MS in **Computer Science** (GPA: 4.0/ 4.0, Specialization: Machine Learning)*

Aug 2017 – Jul 2021 **Birla Institute of Technology and Science, Pilani, India**
Bachelor of Engineering in Computer Science (CGPA: 8.54 / 10)
Minor in Data Science (Minor GPA: 9.8 / 10)

Awards

Georgia Tech President's Fellowship
This award is given to students who have excelled in scholarship and are pursuing doctoral degrees. Selection for this competitive award is based on multiple areas, including but not limited to academic history, standardized test scores, and letters of recommendation. This recognition comes with a stipend of $5,500 per academic year, for a total of 3 years.

Georgia Tech Students Honors Program - The Marshall D. Williamson Fellowship Award 2023
Awarded to a well-rounded, second-year Master's student who best embodies Marshall's values of academic excellence and leadership.

Georgia Tech Students Honors Program - Donald V. Jackson Fellowship Award 2022
Awarded to a well-rounded, first-year Master's student who best embodies the values of academic excellence and leadership.

Birla Institute of Technology and Science best paper award - information retrieval course
Awarded to a computer science undergraduate for rigor in research and academics.

Lalita and Ravindra Nath (LRN) Foundation for student travel grant
Awarded for travel to EMNLP 2022. The award is given to selected student candidates with an accepted paper as a first author at a top-tier International conference outside India and a paper or poster in the COMSNETS conference or one of the associated workshops (in any of the previous COMSNETS).

Employment

May 2025 – Aug 2025 | **Amazon Science** *(Applied Scientist Intern II)*
Developed a clarification-driven dialogue summarization framework that identifies information gaps and generates targeted questions to improve summary completeness and accuracy across clinical, legal, and customer support conversations. Part of Amazon Rufus.

May 2024 – Aug 2024 | **Amazon Science - Healthscribe Team** *(Applied Scientist Intern II)*
Assessing Large Language Models for high-quality behavioral therapy progress note generation. Deployed in Amazon AWS.

May 2023 – Aug 2023 | **IBM Research - Human-Centered AI Team** *(Sr. Research Software Engineer - AI)* Explored different intervention strategies for high-quality human-AI collaborative data creation. Built a full-stack application for user studies. User studies investigate the idea of appropriate user reliance on AI agents using subtle real-time nudges.

May 2022 – Aug 2022 | **Epic Systems -** *(Software Developer - ML Intern)*
Trained Large Language Models to predict workflow actions of Health Care Providers with the goal of bringing relevant patient information to the Provider Dashboard on the fly.

Jan 2021 – Jun 2021 | **JPMorgan Chase -** *(Software Development Intern)* Worked in the Merchant Services team on developing automated testing scripts for the onboarding platform using the Postman tool.

May 2020 – Jun 2020 | **Samsung R&D Institute, India -** *(Research Intern)*
Developed a bidirectional Long-short-term memory-based variable autoencoder to generate melodies and chords for piano music as a part of the User-Experience: Ringtone project.

Leadership and Teaching Experience

Aug 2025 - Dec 2025 | **CS 7641: Human and Machine Learning, Georgia Tech**
Graduate Teaching Assistant - Professor Sashank Varma
- Created homework and assignments for the course CS 7641.
- Assisted students in their course capstone projects.

Jan 2024 - May 2024 | **CS 3790: Intro to Cognitive Science, Georgia Tech**
Graduate Teaching Assistant (TA) - Professor Sashank Varma
- Created homework and assignments for the course CS 3790.

Aug 2023 - Dec 2023 | **CS 7641: Human and Machine Learning, Georgia Tech**
Graduate Teaching Assistant - Professor Sashank Varma
- Assisted students in their course capstone projects.

Jan 2023 - May 2023 | **CS 3790: Intro to Cognitive Science, Georgia Tech**
Graduate Teaching Assistant - Professor Sashank Varma
- Created homework and assignments for the course CS 3790.

Aug 2022 - Dec 2022 | **CS 4650: Natural Language Processing, Georgia Tech**
Graduate Teaching Assistant - Professor Alan Ritter
- Making coding exercises for students to have hands-on NLP experience.

Aug 2020 - Dec 2020 | **CS F469: Information Retrieval, Birla Institute of Technology and Science, Pilani**
Undergraduate Teaching Assistant - Professor Vinti Agarwal
- Created a term paper-based assignment to promote student exposure to research.

Aug 2019 - Dec 2019 | **Academic Counselling Cell, BITS Pilani**
Undergraduate Teaching Assistant - Professor Sangeeta Sharma
- Planned mentorship programs, development seminars, and community bonding activities.

Publications

2025

1. Brachman, M. *et al. [Patent] Guiding a user to interact with an intelligent computing system using best practices* June 2025.
2. Charpentier, L. *et al.* BabyLM Turns 3: Call for papers for the 2025 BabyLM workshop. *arXiv preprint arXiv:2502.10645* (2025).
3. Hsu, S.-L. *et al. Helping the Helper: Supporting Peer Counselors via AI-Empowered Practice and Feedback* in (2025).
4. Jun, J., Marupudi, V., **Shah, R. S.** & Varma, S. *A Neural Network Model of Complementary Learning Systems: Pattern Separation and Completion for Continual Learning* in *Proceedings of the Annual Meeting of the Cognitive Science Society* **47** (2025).
5. Lalai, H. N., Ramakrishnan, A. A., **Shah, R. S.** & Lee, D. *From Intentions to Techniques: A Comprehensive Taxonomy and Challenges in Text Watermarking for Large Language Models* in *Findings of the Association for Computational Linguistics: NAACL 2025* (2025), 6147–6160.
6. Louie, R. *et al.* Can llm-simulated practice and feedback upskill human counselors? a randomized study with 90+ novice counselors. *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems* (2025).
7. **Shah, R. S.**, Huang, J., Murugesan, K., Baracaldo, N. & Yang, D. *The Unlearning Mirage: A Dynamic Framework for Evaluating LLM Unlearning* in *Second Conference on Language Modeling* (2025).
8. **Shah, R. S.** & Varma, S. The potential–and the pitfalls–of using pre-trained language models as cognitive science theories. *arXiv preprint arXiv:2501.12651* (2025).
9. **Shah, R. S.** *et al.* The World According to LLMs: How Geographic Origin Influences LLMs' Entity Deduction Capabilities (2025).
10. **Shah, R. S.** *et al.* TN-Eval: Rubric and Evaluation Protocols for Measuring the Quality of Behavioral Therapy Notes. *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 6: Industry Track)* (2025).

2024

11. Cai, T. *et al. Do Large Language Models Understand Garden-Path Sentences?* in *Proceedings of the Annual Meeting of the Cognitive Science Society* (2024).
12. Chaszczewicz, A. *et al. Multi-Level Feedback Generation with Large Language Models for Empowering Novice Peer Counselors* in *ACL 2024* (2024).
13. Du, J. *et al. LLMs Assist NLP Researchers: Critique Paper (Meta-) Reviewing* in *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing* (2024), 5081–5099.
14. Guo, G., Kang, J. J., **Shah, R. S.**, Pfister, H. & Varma, S. Understanding Graphical Perception in Data Visualization through Zero-shot Prompting of Vision-Language Models. *arXiv preprint arXiv:2411.00257* (2024).
15. **Shah, R. S.**, Bhardwaj, K. & Varma, S. *Development of Cognitive Intelligence in Pre-trained Language Models* in *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing* (2024), 9632–9657.
16. **Shah, R. S.**, Gandhi, A., Marupudi, V. & Varma, S. *Mitigating Catastrophic Interference using Power Law Training Environments* in *27TH EUROPEAN CONFERENCE ON ARTIFICIAL INTELLIGENCE* (2024), (Equal contribution between the first two authors).
17. Sharma, M., Ding, R., **Shah, R. S.** & Varma, S. *Monolingual and Bilingual Language Acquisition in Language Models* in *Submitted to ACL Rolling Review* (2024).
18. Vemuri, S., **Shah, R. S.** & Varma, S. *Evaluating Typicality in Combined Language and Vision Model Concept Representations* in *Proceedings of the Annual Meeting of the Cognitive Science Society* (2024).
19. Yang, W. *et al. What Makes Digital Support Effective? How Therapeutic Skills Affect Clinical Well-Being* in *Proceedings of the ACM on Human-Computer Interaction 6* (2024).

2023

20. Bhardwaj, K., **Shah, R. S.** & Varma, S. *Pre-training LLMs using human-like development data corpus* in *Proceedings of the BabyLM Challenge at the 27th Conference on Computational Natural Language Learning* (eds Warstadt, A. *et al.*) (Association for Computational Linguistics, Singapore, Dec. 2023), 339–345.
21. **Shah, R. S.**, Marupudi, V., Koenen, R., Bhardwaj, K. & Varma, S. *Numeric Magnitude Comparison Effects in Large Language Models* in *Findings of the Association for Computational Linguistics: ACL 2023* (Association for Computational Linguistics, Toronto, Canada, July 2023), 6147–6161.
22. Wang, T. *et al. Metrics for Peer Counseling: Triangulating Success Outcomes for Online Therapy Platforms* in *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems* (Association for Computing Machinery, Hamburg, Germany, 2023).

2022

23. Pathak, A., **Shah, R. S.**, Kumar, V. & Jakhotiya, Y. JARVix at SemEval-2022 Task 2: It Takes One to Know One? Idiomaticity Detection using Zero and One Shot Learning. *CoRR* **abs/2202.02394.** arXiv: 2202.02394 (2022).
24. **Shah, R. S.** *et al. Modeling Motivational Interviewing Strategies On An Online Peer-to-Peer Counseling Platform* in *Proceedings of the ACM on Human-Computer Interaction 6* (2022).
25. **Shah, R. S.** *et al. When FLUE Meets FLANG: Benchmarks and Large Pretrained Language Model for Financial Domain* in *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing (EMNLP)* (Association for Computational Linguistics, 2022).

2021

26. **Shah, R. S.**, Bhatia, A., Gandhi, A. & Mathur, S. *Bitcoin Data Analytics: Scalable techniques for transaction clustering and embedding generation* in *2021 International Conference on COMmunication Systems NETworkS (COMSNETS)* (2021), 1–6.

2020

27. Kristiansen, L.-M., Agarwal, V., Franke, K. & **Shah, R. S.** *CTI-Twitter: Gathering Cyber Threat Intelligence from Twitter using Integrated Supervised and Unsupervised Learning* in *2020 IEEE International Conference on Big Data (Big Data)* (2020), 2299–2308.
28. **Shah, R. S.** & Bhatia, A. *Bitcoin Data Analytics: Exploring Research Avenues and Implementing a Hadoop-Based Analytics Framework* in *Web, Artificial Intelligence and Network Applications* (eds Barolli, L., Amato, F., Moscato, F., Enokido, T. & Takizawa, M.) (Springer International Publishing, Cham, 2020), 178–189.

Ongoing Research Projects

Aug 2025 | **Professor Sashank Varma -** *Part of the Cognitive Architecture Lab at Georgia Institute of technology*
- Human-like learning of LLMs
- The potential – and the pitfalls – of using pre-trained language models as cognitive science theories

May 2025 | **Professor Diyi Yang -** *Part of the SALT LAB at Stanford/ Georgia Institute of technology*
- CARE: Empowering Peer Counselors via Automatic Suggestion Generation
- Evaluating the brittleness of LLM unlearning methods

Service

Ombudsperson - Interactive Computing, Georgia Institute of Technology
The student Ombudsperson / Neutral Third Party is selected by the Associate Chair and Academic Program Manager in the IC Department and serves for a term of 2 years. The primary role of the Student Ombudsperson/Neutral Third Party is to serve as a resource for PhD students who have concerns with a faculty or staff member, or something within the PhD program.

Reviewer for the following conferences
Association of Computational Linguistics Rolling Review 2026, 2025, 2024; International Conference on Machine Learning 2026, Computer-Supported Cooperative Work And Social Computing 2024, 2023; ACM CHI Conference on Human Factors in Computing Systems 2025, 2023; Empirical Methods in Natural Language Processing 2022 (Computational Social Science and Cultural Analytics Track); IEEE BigData 2021.

Student Volunteer for the following conferences
Computer-Supported Cooperative Work And Social Computing 2022, 2021; Annual Meeting of the Association for Computational Linguistics 2021.

Mentorship

- Mentoring experiences with student collaborators:
  - Ifdita Hasan Orney - Senior @ Stanford (August 2024 -)
  - Jenna Kang - PhD @ NYU (August 2024 -)
  - Harsh Lalai - Undergraduate @ BITS Pilani (July 2023 -)
  - James Jun - MS CS @ GT (January 2024 -)
  - Keane Zhang - MS CS @ GT (August 2025 -)
  - Khushi Bhardwaj - Senior @ GT (August 2022 - May 2024, Research Scientist at NVIDIA)
  - Austin Peng - MS CS @ GT (January 2024 - May 2024)
  - Siddhant Narang - MS CS @ GT (January 2024 - May 2024, Senior SWE at BlackRock)
  - Xianle Feng - MS CS @ GT (January 2024 - May 2024)
  - Andrew Li - MS CS @ GT (January 2024 - May 2024)
  - Om Bhatt - MS CS @ GT (January 2024 - May 2024)
  - Dhyan Gandhi - MS CS @ GT (January 2024 - May 2024, SWE at Amazon)
  - Mihir Sharma - BS/MS CS @ GT (January 2024 - May 2024, Research Engineer at NVIDIA)
  - Ziyuan Cao - MS CS @ GT (January 2024 - May 2024, PhD student at OSU)
  - Arpan Parikh - MS CS @ GT (January 2024 - May 2024)
  - Atith Gandhi - MS CS @ GT (January 2023 - May 2024, SWE at Oracle)
  - Ananjan Nandi - MS CS @ Stanford (August 2023 - May 2024, PhD Student at Stanford)
  - Cheng Chang - MS CS @ Stanford (August 2023 - May 2024, now SWE at Google)
  - Siddhartha Vemuri - MS CS @ GT (January 2023 - August 2023)
  - Jason Liu - Undergraduate @ UPenn (August 2021 - August 2022)
  - Damian Rene - Undergraduate @ Swarthmore (August 2021 - August 2022)
- Student mentor and judge at Undergraduate Research Opportunities Program(UROP) @GT

Spoken Languages

English (Native Proficiency), Hindi (Native Proficiency), Gujarati (Native Proficiency), Marathi (Expert Proficiency), and Sanskrit (Intermediate Proficiency).