

# Deepa Tilwani

My research focuses on Neurosymbolic AI to develop trustworthy large language models (LLMs) for interdisciplinary and high-stakes applications. I have received accolades for interdisciplinary research in healthcare by developing plausible methods for interpretable and explainable AI. I have accepted papers in top-tier journals and conferences, including AAAI (core A\*), IEEE Intelligent Systems (IF: 6.7), and IEEE Journal of Biomedical and Health Informatics (IF: 7.7), where I explore the integration of symbolic reasoning with LLMs. I also developed REASONS, an open-source benchmark dataset for developing attribution and reasoning capabilities in LLMs.

## Education

- 2022–Present **Ph.D. in Computer Science and Engineering**, *University of South Carolina*, Columbia, SC, USA
- 2019–2022 **M.Tech in Computer Science and Engineering**, *The LNM Institute of Information Technology*, Jaipur, Rajasthan, India  
Thesis: *Predicting Familial Likelihood of Autism Spectrum Disorder in Infancy Using ECG*

## Skills

- Programming Languages Python, PyTorch, Keras, TensorFlow, Scikit-learn, NumPy, Pandas, CUDA, GIT, C
- Tools Seaborn, Matplotlib, Jupyter, Git, Docker, Matlab, GPU
- Methodologies Machine Learning, Deep Learning, Knowledge Graphs, NeuroSymbolic AI, Signal Processing (EEG, fMRI/MRI, ECG), Large Language Models (LLMs), Reterival Augmentation Generation (RAG)
- Soft Skills Team Leadership, Project Management, Communication Skills, Cross-Functional

## Work Experience

- 2022–Present **Graduate Research Assistant**, *Artificial Intelligence Institute, University of South Carolina*, Columbia, SC, USA
- Leading research projects focused on LLMs for search and attribution.
  - Led a project on *ECG Recordings as Predictors of Very Early Autism Likelihood: A Machine Learning Approach* that resulted in *Trainee Best Research Presentation Winner in SCAND Symposium*.
  - Developed an open source dataset for attribution evaluation *REASONS: A benchmark for REtrieval and Automated citationS Of scieNtific Sentences using Public and Proprietary LLMs*.
  - Collaborated with a multidisciplinary team *to develop machine learning benchmarks for neuroimaging data*.
- June **AI Research And Development Intern**, *Neural Nest, LLC*, Remote, USA
- 2025–August 2025
- Contributed to the development of ARTHEMIS, an AI-driven Virtual Courtroom platform for U.S. arbitration and mediation.
  - Developed multi-agent simulation models to enhance tribunal dynamics and decision-making processes.
  - Designed workflows that reduced case resolution times from over 60 days to approximately 15 days, improving efficiency by 75
  - Supported the creation of secure digital courtroom spaces, increasing accessibility for users nationwide.

- 2021–2022 **Visiting Research Scholar**, *Artificial Intelligence Institute, University of South Carolina, Columbia, SC, USA*
- Facilitated research within a multidisciplinary neuroscience team by contributing AI expertise, helping to bridge the gap between computational methods and neurocognitive studies.
  - Provided valuable insights to support the development of AI tools for analyzing neuroimaging datasets.
  - Gained substantial experience in neuroimaging data (EEG, MRI) processing and analysis, collaborating with neuroscience experts to refine research goals and methodologies.
  - Studied emerging trends in AI, machine learning, and deep learning, including their applications in neuroscience, leading to enhanced understanding and expertise in both fields.
- 2020–2021 **Remote Research Intern**, *Artificial Intelligence Institute, University of South Carolina, Columbia, SC, USA*
- Collaborated on experiments exploring the interaction between ECG and machine learning, assisting in developing new research methodologies.
  - Assisted in designing experiments that tested AI models' effectiveness in predicting Autism likelihood from ECG.

## Publications

- Journal Articles
- **Tilwani, D.**, O'Reilly, C., Riccardi, N., Shalin, V.L., den Ouden, D.B., Fridriksson, J., Shinkareva, S.V., Sheth, A.P. and Desai, R.H., (2025). Benchmarking machine learning models in lesion-symptom mapping for predicting language outcomes in stroke survivors. **Frontiers in Neuroimaging**, 4, p.1573816.
  - Dalal, S., **Tilwani, D.**, Gaur, M., Jain, S., Shalin, V., & Sheth, A. (2024). A Cross Attention Approach to Diagnostic Explainability Using Clinical Practice Guidelines for Depression 2024. **IEEE Journal of Biomedical and Health Informatics (IF: 7.7)** [Paper].
  - **Tilwani, D.**, Venkataramanan, R., & Sheth, A. P. (2024). Neurosymbolic AI Approach to Attribution in Large Language Models. **IEEE Intelligent Systems 2024 (IF 5.6)** [Paper].
  - **Tilwani, D.**, Bradshaw, J., Sheth, A., & O'Reilly, C. (2023). ECG Recordings as Predictors of Very Early Autism Likelihood: A Machine Learning Approach. *Bioengineering*. [Paper]
  - O'Reilly, C., Oruganti, S. D. R., **Tilwani, D.**, & Bradshaw, J. (2023). Model-Driven Analysis of ECG Using Reinforcement Learning. *Bioengineering*. [Paper]
- Conference Proceedings
- **Tilwani, D.**, O'Reilly, C. (2025). Deep Jansen-Rit Parameter Inference for Model-driven Analysis of Brain Activity. **Proceedings of Advances in Signal Processing and Artificial Intelligence**, 92.
  - Mohseni, S., Mohammadi, S, **Tilwani, D.**, Ndawula, G. K., Vema, S., Saxena, Y., Raff, E., Gaur, M. Can LLMs Obfuscate Code? A Systematic Analysis of Large Language Models into Assembly Code Obfuscation. **Proceedings of AAAI 2025**. [Pre-Print]
  - Porwal, S., Patel, K. C., **Tilwani, D.**, & Bansal, S. K. (2021). A Comparative Study and Tool to Early Predict Diabetes Using Machine and Deep Learning Techniques. *Emerging Trends in Data-Driven Computing and Communications*. [Paper]
- Book Chapters
- **Tilwani, D.**, Saxena, Y., Padia, A., Parthasarathy S., Gaur, M. (2026). Neurosymbolic AI for Legal AI-TRISM: Trustworthy, Reliable, Interpretable, Safe Models, In Neurosymbolic AI, Link
- Tutorials
- Gaur, M., **Tilwani, D.**, Raff E., Azimi, I., Chadha, A., Neurosymbolic AI for EGI: Explainable, Grounded, and Instructable Generations, In AAAI 2025

- Posters
- Yang, X., **Tilwani, D.**, O'Reilly, C., & Shinkareva, S. V. (2026). Revisiting the positivity effect: nonlinear age effects on valence perception across adulthood. Paper presented at the 2026 SAS Annual Conference.
  - **Tilwani, D.**, O'Reilly, C. Exploring Neural Dynamics: A Long Short-Term Memory for Brain Effective Connectivity Analysis in EEG. Discover USC, 2024. [Poster]
  - **Tilwani, D.**, Goswami, R., O'Reilly, C., Riccardi, N., Yang, X., Shalin, V., Shinkareva, S., Sheth, A., & Desai, H. R. (2023). Predicting Language Outcomes from MRI Post-Stroke: A Machine Learning Approach. *Organization for Human Brain Mapping*, Montreal, Canada. [Poster]
  - **Tilwani, D.**, O'Reilly, C., Bradshaw, J., & Sheth, A. (2023). Interpretable Machine Learning for Predicting the Likelihood of Autism from Infant ECG Recordings. *SCAND Research Symposium*, Columbia, SC. [Poster, Trainee Best Research Presentation Winner]
- Under Review
- **Tilwani, D.**, Saxena, Y., Mohammadi, A., Raff, E., Sheth, A., Parthasarathy, S., & Gaur, M. (2024). REASONS: A benchmark for REtrieval and Automated citationS Of scieNtific Sentences using Public and Proprietary LLMs. (**KDD 2026**)
  - Yang, X., **Tilwani, D.**, O'Reilly, C., Shinkareva, S. Evaluating Large Language Models as Substitutes for Human Affective Ratings in Naturalistic Paradigms. Under Review in IEEE Transactions on Affective Computing.

## Open Source Contributions

Dataset REASONS: REtrieval and Automated citationS Of scieNtific Sentences (<https://zenodo.org/records/18264033>).

## Awards and Achievements

- 2025 Recipient of the 2025 University Travel Award for attendance at the International Conference on Advances in Signal Processing and Artificial Intelligence (ASPAAI' 2025).
- 2025 AAI-25 Student Scholarship Award
- 2024 EMNLP Diversity and Inclusion Award
- 2023 Trainee Best Research Presentation Winner, SCAND Symposium.
- 2023 Research Symposium Third Place Poster Award, University of South Carolina.
- 2021 Nirmala and Jashwantlal Clerk Memorial Scholarship, AIISC.
- 2020 2nd Prize, LINZ Ars Festival - BR41N.IO Hackathon.
- 2020 2nd Prize, BR41N.IO: Brain-Computer Interface Designers Hackathon.
- 2016 1st Place, Poster Presentation on AR and VR Technology, GWECA.
- 2015 3rd Place, Coding Challenge: Toast to Code - C Language, GWECA.
- 2012 Silver Prize, National Science Olympiad (NSO).

## Advising and Mentoring

- Yash Saxena, Galgotias University, Sept 2023- Sept 2024. Project: "REASON: Reference and Assertions for Consistent Evaluation of Factual/Non-Factual Sentences".
- Nethra Gunti, IIIT SriCity, 2022. Project: "Phase Shift Analysis in Autism Spectrum Disorder: A Video-Based Study of Parent and Object Interactions".
- Sai Durga Rithvik Oruganti, University of South Carolina, 2022. Project: "Phase Shift Analysis in Autism Spectrum Disorder: A Video-Based Study of Parent and Object Interactions".

## Selected Media Coverage

- **USC 2024 Newsletter: Pioneering AI to transform autism diagnosis.** Link

## Teaching Experience

- Teaching Assistant, SCINBRE Machine Learning in Python Workshop 2024, University of South Carolina.
- Instructor, Introduction to Machine Learning, AIISC High School Summer Camp, 2024.
- Instructor, Introduction to Python, AIISC High School Summer Camp, 2023.
- Teaching Assistant (2019-2021), The LNM Institute of Information Technology: Computer Networks, Data Structures, DBMS, and Advanced Programming Labs.

## Community Service

- Conference ○ ICWSM 2026
- Reviewer ○ AAI 2025, 2026
  - EMNLP 2025
  - KDD 2025
  - THE WEB CONFERENCE 2025 (WWW).
  - CIKM, KG-STAR Workshop, 2024.
  - KDD, KIL Workshop, 2024.
- Journal ○ IEEE Internet Computing
- Reviewer ○ Neurosymbolic Artificial Intelligence
  - ACM Transactions on Computing for Healthcare
  - ACM Computing Surveys
  - Scientific Reports
  - Data Mining and Knowledge Discovery
  - Frontiers in Neuroimaging
  - MDPI, Advanced NLP and Machine Translation
- Program ○ KG-STAR, CIKM 2024
- Committee ○ KIL, KDD 2024
  - KIL, KDD 2023
- Voluntary Work ○ Web and Publicity Chair, KG-STAR Workshop CIKM 2024: Organized events, managed communications, and enhanced visibility of the workshop.
  - Coordinator, AIISC Retreat, 2023: Organized the institute's retreat, ensuring participation and facilitating collaborations.
  - Session Moderator and Publicity Chair, ACM KDD Workshop on Knowledge-infused Learning, 2023: Moderated discussions and Q&A sessions.
  - Coordinator, AIISC High School Summer Camp, 2023: Led the planning and execution of the camp, including scheduling and recruitment of instructors.
  - Student Member, AAI (2022-Present).