

Jugal Gajjar

812jugalgajjar@gmail.com | github.com/JugalGajjar | linkedin.com/in/jugal-gajjar | jugalgajjar.github.io | Google Scholar

RESEARCH SUMMARY

My research focuses on trustworthy and reasoning-centric systems, particularly at the intersection of large language models (LLMs), retrieval-augmented generation (RAG), agentic workflows, and human-AI collaboration. I am interested in developing structured and verifiable AI systems that combine neural representations with symbolic reasoning, graph-based knowledge, and grounded validation for reliable decision-making, business and scientific automation, and code intelligence.

EDUCATION

M.S. in Computer Science (Machine Intelligence & Cognition)

The George Washington University, *Washington, DC*

May 2026

GPA: 3.88/4.0

B.Tech. in Computer Science and Engineering

Navrachana University, *Vadodara, India*

May 2024

CGPA: 9.39/10

RESEARCH PUBLICATIONS

Verify Before You Fix: Agentic Execution Grounding for Trustworthy Cross-Language Code Analysis

Solo Author; Under Review at NeurIPS 2026

Dec. 2026

RSAT: Structured Attribution Makes Small Language Models Faithful Table Reasoners

First Author; SURGeLLM Workshop (Proceedings Track) at ACL 2026

Jul. 2026

Feature Attribution Stability Suite: How Stable Are Post-Hoc Attributions?

Co-Author; XAI4CV Workshop (Proceedings Track) at CVPR 2026

Jun. 2026

Bridging Semantics & Structure for Software Vulnerability Detection using Hybrid Network Models

First Author; 14th International Conference on Complex Networks and their Applications

Dec. 2025

HyperCompleX: Adaptive Multi-Space Knowledge Graph Embeddings

First Author; 2025 IEEE International Conference on Big Data (IEEE BigData 2025)

Dec. 2025

SecureFixAgent: A Hybrid LLM Agent for Automated Python Static Vulnerability Repair

First Author; 2025 International Conference on Machine Learning and Applications (ICMLA)

Dec. 2025

MLCPD: A Multi-Language Code Parser Dataset with Universal AST Representation

First Author; arXiv preprint (cs.SE)

Oct. 2025

MalCodeAI: Autonomous Vulnerability Detection and Remediation via Language Agnostic Code Reasoning

First Author; IEEE 26th International Conference on Information Reuse and Integration (IRI 2025)

Aug. 2025

Multimodal Sentiment Analysis on CMU-MOSEI Dataset using Transformer-based Models

First Author; arXiv preprint (cs.CL)

May 2025

PROFESSIONAL SERVICE

Program Committee Member:

- 38th IEEE International Conference on Tools with Artificial Intelligence (ICTAI)

Nov. 2026

Reviewer:

- ICMLA 2026	Oct. 2026
- ACL 2026 (GEM Workshop, SURGeLLM Workshop, CDL Workshop)	Jul. 2026
- IEEE USBEREIT 2026	May. 2026
- ICMLA 2025 (ROSE-LLM Special Session)	Dec. 2025

SELECTED RESEARCH PROJECTS

RSAT: Structured Attribution for Faithful Table Reasoning in Small LLMs Dec. 2025 – Mar. 2026

- Designed a two-phase SFT + GRPO training pipeline enabling small language models (1–8B) to generate step-by-step reasoning with explicit cell-level citations.
- Formulated a multi-objective reward function combining faithfulness, citation validity, and reasoning efficiency for GRPO-based policy optimization.
- Applied GRPO-based policy optimization to improve reasoning faithfulness and attribution quality under structured reward signals.
- Achieved $3.7\times$ improvement in attribution faithfulness over supervised fine-tuning baselines.

HyperComplEx: Adaptive Multi-Space Knowledge Graph Embeddings Apr. 2025 – Oct. 2025

- Proposed a representation learning framework for knowledge graphs combining hyperbolic, complex, and Euclidean spaces through relation-adaptive attention.
- Developed a scalable multi-space model with consistency regularization for heterogeneous relational reasoning.
- Achieved 0.612 MRR on a 45M-triple knowledge graph with near-linear scalability across heterogeneous relations.

FASS: Feature Attribution Stability Suite for Robust XAI Evaluation Nov. 2025 – Feb. 2026

- Developed a benchmark for evaluating attribution stability under prediction-invariant perturbations (geometric, photometric, compression).
- Introduced a three-axis stability decomposition: structural similarity, rank correlation, and top- k overlap.
- Enabled systematic comparison of post-hoc explanation methods for robustness and reliability analysis.

Cross-Language Vulnerability Detection, Validation, and Repair (M.S. Thesis) Jun. 2025 – Mar. 2026

- Built a uAST-based agentic framework for vulnerability detection and repair, achieving a 69.74% end-to-end resolution rate on a 120K-sample dataset.
- Designed a GNN + transformer fusion model, reaching 74.4–78.2% zero-shot cross-language F1.
- Developed an end-to-end pipeline with execution-based validation, reducing false positives by up to 62% and successfully repairing 81–87% of cases.

VulnGraph: Graph + LLM Embeddings for Vulnerability Detection Jul. 2025 – Sep. 2025

- Engineered a multimodal fusion model integrating AST/CFG graph embeddings with LLM-based semantic representations.
- Integrated structural program representations (AST/CFG) as symbolic constraints to guide neural-LLM feature fusion via a two-way gating mechanism for effective alignment.
- Achieved 93.57% accuracy, outperforming GNN-only (+8.36%) and LLM-only (+17.81%) baselines with interpretable outputs.

TECHNICAL SKILLS

Programming: Python, C++, Java, SQL, Bash**Machine Learning & AI:** PyTorch, TensorFlow, Hugging Face Transformers and TRL, Reinforcement Learning (RL), Reinforcement Learning from Human Feedback (RLHF), Preference Optimization (GRPO/DPO), Representation Learning, Large Language Models (LLMs), Retrieval-Augmented Generation (RAG), GNNs, AI Agents**Systems & Infrastructure:** Linux, Git, Docker, Apache Spark, Hadoop, AWS, Google Cloud**Research Interests:** Reinforcement Learning, Interactive Learning, AI Alignment, Preference Learning, Decision Making Under Uncertainty, Large Language Models, Graph Representation Learning, Explainable AI (XAI), Agentic AI Systems, Knowledge Graph Reasoning

RELEVANT EXPERIENCE

Graduate Assistant

The George Washington University, *Washington, DC*

Sep. 2025 – Dec. 2025

- Supported an upper undergrad- & graduate-level Big Data and Analytics course (CSCI 4907/6444).
- Guided and graded students on assignments and projects, focusing on tools like Spark, Hadoop, and S3 Bucket, and programming in languages like Python, Java, and Linux-Bash.
- Assisted with analytical methods, including machine learning, and provided support for cloud-based and distributed systems projects.

Teaching Assistant

Navrachana University, *Vadodara, India*

Jan. 2024 – May 2024

- Mentored undergraduate students through code review, grading, and structured feedback on 750+ lab reports.
- Assisted in designing rubrics, evaluating technical writing, and improving student understanding of test-driven development.
- Supported curriculum delivery along with faculty, focusing on code quality, documentation, and debugging.

NVIDIA Jetson AI Project Coordinator

Navrachana University, *Vadodara, India*

Mar. 2022 – Jun. 2022

- Conducted workshops on CV and NLP using the Jetson Nano platform, introducing students to embedded AI and real-time inference.
- Mentored 5+ student research projects selected by the NVIDIA Deep Learning Institute.
- Promoted student research initiatives and hands-on learning in edge AI systems.

CERTIFICATIONS & HONORS

Sustainability Award (Complex Networks 2025 Conference)	2025
NVIDIA-Certified Associate: Generative AI LLMs	2025
GW SEAS Dean's Award Scholarship	2024
Runner-Up, Tinkerthon 2.0 Hackathon	2024
SSIP Innovation Grant (Project EzyCart)	2023
NVIDIA Jetson AI Specialist	2022