

Jugal Gajjar

812jugalgajjar@gmail.com · github.com/JugalGajjar · linkedin.com/in/jugal-gajjar

SUMMARY

AI/ML Engineer with an M.S. in Computer Science (GWU, GPA 3.88, May 2026) and production experience building LLM agent systems, RAG pipelines, and ML training pipelines. Trained SFT + GRPO alignment across LLMs on 75K examples on A100/H100 GPU (RSAT, ACL 2026); built a DQN-based RL policy evaluated on 199,847 data points (Adaptive RAG); operates a production multimodal agent platform at projectmmap.com. Proficient in Python, PyTorch, HuggingFace, LangChain, LangGraph, and cloud ML. 6 peer-reviewed publications. STEM OPT starting June 22, 2026; work-authorized through mid-2029 without sponsorship.

SKILLS

AI & LLMs: LLM Agents, Multi-Agent Orchestration, RAG, Prompt Engineering, LangChain, LangGraph, n8n, AutoGen, Tool Use, Structured Outputs, Evaluation Harnesses, Claude Code, GitHub Copilot

ML & DL: PyTorch, TensorFlow, HuggingFace Transformers, TRL (GRPO/PPO/DPO), PEFT/LoRA, Deep Q-Learning, GNNs (GraphSAGE, GCN, GAT), YOLOv5/v8, CLIP, scikit-learn

ML Ops: MLflow, W&B, SFT + GRPO Training Pipelines, A100/H100 GPU Training, Ablation Studies, Benchmark Design, Model Evaluation, LLM-as-Judge Pipelines, statsmodels

Data & Search: PostgreSQL, Qdrant, Neo4j, OpenSearch, Redis, MongoDB, Apache Parquet, Pandas, NumPy, HF Datasets, Kafka

Languages & Infra: Python, TypeScript, Java, C++, SQL, Bash, FastAPI, REST APIs, SQLAlchemy, Docker, Kubernetes, AWS, GCP, GitHub Actions CI/CD

EXPERIENCE

Graduate Assistant – Big Data & Analytics, The George Washington University *Sep 2025 – Dec 2025*

- Designed big data analytics assignments and built Python automation to evaluate 20+ student results against ground-truth datasets.
- Guided students through cloud ML infrastructure, distributed data pipelines, and data engineering concepts during office hours.

Teaching Assistant – Software Testing, Navrachana University *Jan 2024 – May 2024*

- Developed TDD curriculum and automated testing tooling for 80+ students in Python and Java; conducted weekly code reviews with structured feedback on correctness, API testing, and coverage; maintained course infrastructure with clear documentation.

Freelance Software Developer, Shubhlaxmi Mfg. Co. *Jun 2023 – Jul 2023*

- Gathered requirements from the client, designed a Python/MySQL application, and shipped it end-to-end with full ownership, delivering a 50% operational efficiency improvement through clean, tested, well-documented code.

NVIDIA Jetson AI Project Coordinator, Navrachana University *Mar 2022 – Jun 2022*

- Coordinated NVIDIA Deep Learning Institute AI workshops for 20+ students on Jetson Nano hardware, mentored 5+ student projects from initial setup through production deployment, and built documentation for all sessions.

PROJECTS

Multimodal AI Intelligence Platform: Production LLM Agent System (Live Website) projectmmap.com

- Built a production LLM agent platform at projectmmap.com: LangGraph multi-agent orchestration (routing, retrieval, reasoning, and verification agents with tool use), RAG over PDFs, images, and audio using Qdrant HNSW vector search and a Neo4j knowledge graph, multi-model routing across LLM providers, Whisper Tubro for audio transcription, and NVIDIA Nemotron Nano 2 VL for vision understanding.
- Python/FastAPI async backend with arq workers, Next.js 16/React 19 TypeScript frontend with real-time SSE streaming, enterprise security (JWT, RBAC, bcrypt, rate limiting, CSP/HSTS, prompt injection protection), GitHub Actions CI/CD with nightly Playwright e2e tests, and Prometheus/Grafana/OpenTelemetry observability.

RSAT: SFT + GRPO Post-Training Alignment for LLMs github.com/JugalGajjar/RSAT

- End-to-end LLM training pipeline: SFT on 1,000 curated examples then GRPO reward optimization with a composite NLI-based reward (faithfulness, citation validity, parsimony) across 6 transformer families on 75K examples with A100/H100 GPU, tracked across 40+ ablation runs in MLflow and W&B.
- Resolved reward hacking through ablation and reward redesign, achieving a 3.7x faithfulness improvement; all code and model weights are open-sourced; published as an oral presentation at the SURGeLLM workshop in ACL 2026.

Adaptive RAG: Reinforcement Learning for Dynamic Retrieval Optimization github.com/JugalGajjar/Adaptive-RAG

- Formulated adaptive retrieval as a Markov Decision Process and trained a DQN policy (2-layer MLP, 4-head attention, 150K training examples) with a 3-action policy (retrieve, reformulate, answer) across LLM providers from 3.8B to 120B parameters.
- Evaluated on 199,847 questions across 9 QA datasets, achieving +3.2–6.5% accuracy over fixed-strategy baselines while using 14–37% fewer retrievals. All experiments tracked with MLflow and W&B; full codebase open-sourced.

EDUCATION

M.S. in Computer Science · The George Washington University · GPA: 3.88/4.0 *May 2026*

B.Tech. in Computer Science and Engineering · Navrachana University · CGPA: 9.39/10 *May 2024*

AWARDS & PUBLICATIONS

6 peer-reviewed publications on LLM agents, RAG, trustworthy AI, and software security · ACL 2026 · CVPR 2026 · IEEE BigData 2025 · GW SEAS Dean's Award Scholarship · Sustainability Award, Complex Networks 2025 · NVIDIA-Certified Associate: Generative AI LLMs · Architecting with Google Compute Engine (Coursera) · AWS Academy Graduate: Cloud Foundations