

# Naveen Prashanna Gurumurthy

[gnavveen1509@gmail.com](mailto:gnavveen1509@gmail.com) | +1 (945) 527-5193 | [portfolio-naveen-sde.vercel.app](https://portfolio-naveen-sde.vercel.app) | [linkedin.com/in/naveen015](https://linkedin.com/in/naveen015)

## Summary

- Software engineer with 2+ years of experience designing, building, and deploying scalable Web and ML systems.
- Expertise in developing trading systems, end-to-end ML pipelines, and CI/CD workflows for Gen-AI applications.
- Proven ability to bridge the gap between data science & software engineering to deliver production-grade solutions.

## Education

- Master of Science, University of Texas at Dallas (UTD)** | Dallas, TX **Aug 2023 - May 2025**
- Major: Computer Science | Specialization: Intelligent System
- Bachelor + Master of Technology, Indian Institute of Technology (IIT) Madras** | India **Aug 2017 - May 2022**
- Major: Mechanical Engineering | Minor: Artificial Intelligence & Machine Learning

## Technical Skills

- **Languages:** C++, Python, Javascript, TypeScript, SQL, R, MATLAB, Bash
- **Frameworks:** PyTorch, Tensorflow, Scikit-learn, LangChain, Vue.js, React.js, FastAPI, Hadoop, Spark
- **Software & Tools:** GCP, DataBricks, Docker, Kubernetes, AWS, Snowflake, MongoDB, PostgreSQL, Jenkins

## Professional Experience

- AI Engineer, Kahana Group Inc** | United States **Feb 2025 - Present**
- Spearheaded the development of user-centric UI/UX features using React.js, to create a intuitive and responsive interface that improved user engagement by 12% and overall usability.
  - Integrated robust data analytics tools into our core product, providing real-time insights that directly influenced a product strategy shift which resulted in a 10% increase in subscription revenue.
  - Engineered a LangChain-based AI browser assistant to let users perform actions with natural language commands, using LangGraph to model state transitions and LangSmith for real-time activity tracking and system refinement.
- Software Engineer, Quantitative Brokers** | Chennai, India **Jul 2022 - Jan 2023**
- Engineered a proprietary internal tool to automate the creation and transmission of FIX order messages, which eliminated manual processes and streamlined high-frequency trading operations.
  - Enhanced the core FIX messaging platform by integrating robust support for Multi-Leg trade orders, which boosted operational efficiency by 15% and enabled the firm to expand its product offerings.
  - Designed and optimized decentralized database solutions using PostgreSQL, enhancing database efficiency and reducing query latency by 10% to support real-time data needs.
- Software Engineer Intern, Quantitative Brokers** | Chennai, India **Jun 2021 - Jan 2022**
- Developed a highly interactive Vue.js web application to visualize operational costs across global trading exchanges and network infrastructure, providing real-time insights that led to a 6% reduction in annual network expenditures.
  - Led the strategic integration of SonarQube and BlackDuck into the Jenkins CI/CD pipeline, which automated code quality analysis and vulnerability scanning, enhancing application security and reducing critical risk factors by 20%.
- Machine Learning Engineer Intern, Big Data Science Research** | Bangalore, India **Apr 2020 - Jun 2020**
- Developed and deployed a machine learning pipeline using Python and scikit-learn to predict urban traffic patterns, enhancing model accuracy by 15% through the integration of geospatial data features.
  - Implemented a novel map-matching algorithm to refine GPS data, significantly improving the precision of traffic flow models and reducing prediction error by 20%.

## Projects

- Generative AI Engineer: Personal Assistant Chatbot** | Personal Project **Apr 2025 - May 2025**
- Fine-tuned the Mistral LLM using a custom QLoRa approach on a specialized dataset, which reduced the model's memory footprint by 40% while maintaining 98% of its original performance.
  - Integrated a Retrieval-Augmented Generation (RAG) pipeline to retrieve embeddings from personal documents, which decreased response latency by 160ms and reduced factual inaccuracies (hallucinations) by 75%.
- ML Engineer: Anthropic's Computer Control** | Personal Project **Feb 2025 - Mar 2025**
- Developed a LangChain-based system for Mac control, successfully replicating core functionalities of Anthropic's Computer Control tool to enable natural language-based desktop automation.
  - Implemented custom tools for text simulation, mouse automation, and application management, which reduced average task completion time by 30% by eliminating manual steps.
- RL Engineer: Trajectory-Aware Human Feedback for Hierarchical RL** | UTD **Sep 2024 - Dec 2024**
- Proposed a novel Hierarchical Reinforcement Learning framework to improve subgoal generation in complex tasks, demonstrating a new approach for solving long-horizon problems.
  - Deployed the Deep-RL framework in the FetchReach environment, resulting in a 10% increase in task success rates by effectively breaking down complex goals into manageable sub-goals.