

UMANG GARG

+91 9717182251 ◊ umanggarg28@gmail.com ◊ linkedin.com/in/umanggarg28 ◊ github.com/umanggarg28

PROFESSIONAL SUMMARY

Software Engineer with 8+ years of experience building scalable, production-grade applications across full-stack environments. Track record of delivering complex systems end-to-end — spanning backend APIs, frontend applications, and AI-integrated products — across healthcare AI, biotech, and enterprise domains. Combines strong software engineering fundamentals with applied AI engineering experience, including production ML systems, LLM-powered tooling, and agentic workflow design.

TECHNICAL SKILLS

AI/LLMs: RAG, Hybrid Search, HyDE, Cross-encoder Reranking, LangChain, LangGraph, Multi-agent Systems, MCP, Claude API / Anthropic SDK, Prompt Engineering, LLMOps, Vector Databases (Qdrant, ChromaDB)

ML / Vision: PyTorch, scikit-learn, Computer Vision, OpenCV, OpenSlide, TorchServe, tree-sitter

Languages: Python, JavaScript (ES6), TypeScript, SQL, HTML/CSS, Java

Frameworks: Django, React.js, Node.js, FastAPI, Redux, Express, GraphQL, OpenCV

Infrastructure: AWS (S3, ECS, EKS, Lambda, SQS, CloudWatch), Docker, Apache Kafka, Git, Vercel, HuggingFace Spaces

EXPERIENCE

Full Stack Developer 2

Jun 2025 – Feb 2026

PictorLabs.ai — AI-powered digital pathology platform (UCLA spinoff, venture-backed)

Remote, India

- Developed LLM-powered internal tooling using Claude API and LangGraph - RAG over clinical documentation, multi-step agent workflows for automated reasoning and task execution.
- Fine-tuned Google's PathFoundation (histopathology vision foundation model) in PyTorch; built end-to-end training pipeline: OpenSlide patch extraction, tissue masking, Macenko stain normalization, augmentation; integrated into validation pipeline for automated stain type detection.
- Led delivery of ClearStain - AI virtual HE staining pipeline for unstained brightfield WSIs; owned Django API routing logic (organ/species/diagnosis-based predictor selection), Kafka job orchestration, and TorchServe model serving integration.
- Built adaptive chunked multipart upload (React Web Workers + S3) with concurrent worker pool, exponential backoff retry, and resumable support for multi-tenant pathology SaaS.
- Worked within event-driven microservices: Kafka/MSK, PostgreSQL, S3, ECS/EKS; Docker multi-stage builds, CI/CD pipelines.

Software Engineer 2

May 2024 – Nov 2024

Medhaz Corp

Charlotte, NC

- Developed a candidate data management portal using Python/Django REST Framework and React.js, processing structured workforce data across organizations.
- Led migration from React 15 to React 18, optimizing performance via Concurrent Rendering and React Suspense; refactored legacy code and ensured backward compatibility for a seamless transition.

Software Engineer 2

Jun 2022 – Mar 2024

Asurion

Charlotte, NC

- Lead developer of the **Entitlement Engine**, an API engine for insurance claim eligibility determination serving 40,000+ users weekly; achieved 99% uptime through optimized Node.js architecture.
- Contributed to migration from monolithic to micro-frontend architecture, enabling faster and independently deployable release cycles with improved fault isolation.
- Designed and maintained serverless event-driven Python-based data processing pipelines using AWS Lambda, API Gateway, S3, SQS, and CloudWatch, automating downstream data workflows at scale.
- Built React functional components with React Hooks; enhanced UI with responsive HTML5/CSS3 layouts.

React Developer

Oct 2021 – Jun 2022

HYR Global Source, Client — Kroger Co.

Charlotte, NC

- Developed cross-functional frontend architecture and reusable React components deployed across multiple Kroger entities in an enterprise agile environment.
- Accelerated development using component-based architecture with React (ES6) and Redux for state management; enforced code quality via ESLint/Webpack tooling.
- Built comprehensive test suites with Jest and React Testing Library, increasing coverage by 50%.

Software Developer

Mar 2019 – Jul 2021

GTR Event Technology

Charlotte, NC

- Built and maintained client event registration and CEU platforms with fully automated CI/CD pipelines using React, Node.js, Express, MongoDB, and RESTful APIs.
- Developed a computer vision-based automated image classification system using Python and OpenCV; applied feature extraction and spatial analysis techniques to detect and categorize booth layouts from event floor plan images, eliminating 50+ hours of manual work per event.
- Integrated GraphQL mutations/queries for Swapcard API connectivity with the event registration platform.

Software Engineering Co-Op

Jun 2018 – Dec 2018

Rochester Regional Health

Rochester, NY

- Developed backend REST APIs for a web-based Test Catalog application using Java, Hibernate OGM, and MongoDB; implemented SSO and full CRUD/file services.
- Integrated OpenAPI (Swagger 2.0) for auto-generated REST API documentation, improving API discoverability and documentation quality.

EDUCATION

Master of Science in Software Engineering

2016 – 2018

Rochester Institute of Technology

GPA: 3.94

B.Tech in Mechanical Engineering

2010 – 2014

ITM University

PROJECTS

Cartographer - Codebase Reasoning Graph

github.com/umanggarg28/Cartographer

- Indexes any GitHub repo and answers natural-language questions about code with source citations. Built without LangChain or LlamaIndex - every retrieval layer implemented manually.
- Retrieval: tree-sitter AST chunking, contextual LLM descriptions, Qdrant hybrid search (dense + BM25 + RRF), Cohere cross-encoder reranking, HyDE + query expansion.
- ReAct agent via MCP (FastMCP) - 12 tools, working memory, parallel async execution, streaming thought traces to UI.
- Stack: FastAPI, React/Vite, Qdrant Cloud, multi-provider LLM cascade. Deployed: HuggingFace Spaces (Docker) + Vercel via GitHub Actions.

CERTIFICATIONS

Anthropic (2026): MCP Advanced Topics, Introduction to MCP, Claude Code in Action, Claude with Amazon Bedrock, Introduction to Agent Skills, Introduction to Subagents

DeepLearning.AI: Retrieval Augmented Generation (RAG), Agentic AI, PyTorch and Deep Learning