

# Kunal Mahato

[kunalmahato.tech](http://kunalmahato.tech) | +1-5712659470 | [linkedin.com/in/kunalmahato](https://linkedin.com/in/kunalmahato) | [github.com/mahatokunal](https://github.com/mahatokunal) | [mahatokunal41@gmail.com](mailto:mahatokunal41@gmail.com)

## WORK EXPERIENCE

---

**Virginia Tech, USA** | Research Assistant

July 2025 - April 2026

- Built data-processing pipelines for LLMs in Python integrating OpenAI, Llama, and Gemini APIs through implementation of data-selection strategies including seed-based, embedding-based, alongside domain expert in the loop for ~1k records.
- Designed an agentic RAG pipeline (PyTorch + SentenceTransformers) with tool-use for top-K selection, improving relevance and recall and achieving ~80% F1 in classification tasks.
- Automated reporting with Pandas/NumPy, producing results across 12 model setups and reducing run time by 3x.

**Turvo Inc, India** | Software Engineer

May 2021 - June 2023

- Enhanced global search and filter efficiency by 10% through advanced Elasticsearch indexing and GraphQL API implementation, improving data retrieval speed and accuracy for supply-chain operations.
- Architected a resilient Python-based ETL pipeline to launch Intermodal shipment mode, migrating 216,000+ multi-leg records between heterogeneous data stores (MySQL to MongoDB) with zero data loss during a critical schema redesign.
- Managed reorganization of the Rules Framework into a dynamic Business Rule Management System, enforcing complex role-mode-status compliance constraints across 10+ enterprise clients while decoupling business logic from the core codebase.
- Built Jenkins CI/CD pipelines for isolated dev setups, cutting deployment time by 40% and improving release reliability.
- Implemented a resilient Python library for HashiCorp Vault, ensuring secure secret injection with 99.99% uptime for 10 clients.

**Samsung, South Korea** | R&D Engineer

December 2019

- Integrated backend UI components in Java using Android Framework with system services for foldable-specific notifications, contributing to the global launch of the world's first foldable phone at Samsung HQ.

**Samsung, India** | R&D Engineer

June 2018 - April 2021

- Maintained and optimized full-stack codebase for 5 OS iterations, enhancing performance for 1.48+ billion devices.
- Optimized System UI (Quick Panel & Notifications) in Java for low end devices, cutting memory use 20% and battery drain 20%, reducing crash rate 85% and boosting user satisfaction to 4.4/5 across 1M+ devices.
- Handled 5 Application and Framework modules - Settings, Calendar, QuickPanel, Status-Bar and warning notifications.
- Upgraded the customer satisfaction metrics by 10% through implementation of Search feature for Quick Panel.
- Collaborated cross-functionally with teams to build backend data-porting solutions integrating 20+ network operators.

## EDUCATION

---

**Virginia Tech**

Virginia, USA

Master Of Engineering in Computer Science & Applications - GPA: 3.93/4

August 2023 - May 2025

**Motilal Nehru National Institute Of Technology Allahabad (NIT Allahabad)**

Prayagraj, UP, India

Bachelor Of Technology in Computer Science & Engineering - GPA: 3.53/4

July 2014 - May 2018

## PROJECT EXPERIENCE

---

**U-Pass Manager** | Virginia Tech | [GitHub](#) | [Demo](#) | [VTechworks](#)

January 2025 - May 2025

- Architected Kubernetes-managed Next.js, React, Node.js web app, ensuring 99.9% uptime for 10K+ student records.
- Cut manual entry 95% with AWS Lambda automating S3 to RDS ingestion and boosted pickups using SES notifications.
- Improved API speed 50% and secured role-based access via optimized GraphQL/REST endpoints (Express, Apollo, JWT).
- Halved student wait times with a 100 ms NFC validation pipeline using Socket.IO for real-time distribution.

**LifexAI** | Bitcamp-Hackathon of University of Maryland | [GitHub](#) | [Demo](#)

April 2025

- Built LifexAI - AI detective web app integrating Google Gemini Vision API, Capital One Nessie API, and real-time Chart.js dashboards with Tailwind CSS and Framer Motion, boosting anomaly detection accuracy by 30%.
- Accelerated data loading 50% with Typescript React Query caching & Next.js APIs on Cloudflare Workers.
- Ensured 99.9% uptime and global delivery via Vercel-hosted frontend and automated CI/CD with Docker.

**DoDEx (Winner for People's Choice Hack)** | Bitcamp-Hackathon of University of Maryland | [GitHub](#) | [Devpost](#)

April 2024

- Designed an AI-driven ingestion pipeline using Python and RAG to parse unstructured Defense Contract Notices. This architecture extracted structured compliance data with 30% higher accuracy than regex methods.
- Developed a ground-truth validation framework to benchmark LLM, preventing hallucinations in mission-critical reporting.

## SKILLS

---

**Programming Languages and Databases:** Java, Python, Go, C++, JavaScript, PostgreSQL, MongoDB, MySQL.

**AI & Machine Learning:** LangChain, LLMs, RAG, PyTorch, HuggingFace, TensorFlow, Pandas, NumPy, Streamlit.

**Frameworks and Technologies:** React, Angular, Next.js, TypeScript, Django, GraphQL, Spring, Node.js, Express, FastAPI.

**Cloud & DevOps:** AWS, Docker, Kubernetes, Terraform, CI/CD.

**Tools & Platforms:** Git, Figma, JIRA, IntelliJ, Kibana, UNIX, Perforce/P4V.