

Ketaki Dabade

(651) 384-8787 | kvd2112@columbia.edu | linkedin.com/in/ketakidabade | github.com/ketaki3 | ketaki.dev

EDUCATION

Columbia University

New York, NY

Master of Science in Computer Science (Machine Learning Track)

Aug. 2025 – Dec. 2026

- **Coursework:** AI for Operations Research & Financial Engineering, Natural Language Processing, Neural Networks & Deep Learning, Applied Data Science, Startup Studio, Databases

Dr. Vishwanath Karad MIT World Peace University

Pune, IN

B.Tech. in Computer Science and Engineering, CGPA: 3.74/4.0

Jul. 2021 – Jul. 2025

- Published **3** peer-reviewed papers (1 IEEE, 2 Springer)

EXPERIENCE

Columbia University, QMSS

New York, NY

Teaching Assistant, Natural Language Processing (GR5067)

Sep. 2026 – Dec. 2026

- Teaching assistant for the graduate NLP course: hold office hours, grade assignments, and advise students on final NLP projects

Carlson Private Capital Partners (Private Equity)

Minnetonka, MN

AI Engineering Intern

Jun. 2026 – Sep. 2026

- Ran 1:1s with the full deal team to map where time goes in sourcing, screening, and document review, turned the pain points into a prioritized build plan, and wrote the docs and prompting patterns for every tool shipped
- Shipped an NDA redlining agent through **7 releases** of associate and VP testing on live deals, plus a meeting-note pipeline in daily use syncing **45 interactions per week** into the firm's CRM through its REST API, behind a human-confirmation gate
- Built the question-answering agent for RelGraph, the firm's relationship graph (LangGraph, LLM tool use): plans retrieval across **8** graph, SQL, and vector tools, finds warm-intro paths up to **4 hops** out, and cites every claim to a database record
- Wrote the eval harness gating every prompt and retrieval change in CI on retrieval recall and citation integrity
- Shipped the deal intake workflow on the firm's internal AI platform (React, Express, PostgreSQL, Docker, Azure): a subagent reads the shared inbox, extracts deal fields from emails and PDFs, and writes them into the CRM (~**6,300** companies) after an associate confirms

Complex Resilient Intelligent Systems Lab, Columbia University

New York, NY

Research Assistant under Prof. Venkat Venkatasubramanian

Sep. 2025 – May 2026

- Built the document-to-knowledge pipeline behind the lab's hierarchical sparse autoencoder (SAE) interpretability work: **3,000+** pages of technical PDFs (MinerU), **17,000+** embeddings (Qwen3-Embedding), and **493** topics surfaced with BERTopic/HDBSCAN

AI4M Technology Private Limited

Pune, IN

Deep Learning Engineer Intern

Jul. 2024 – Dec. 2024

- Trained and deployed computer vision models for manufacturing defect detection on NVIDIA Jetson with TensorRT quantization (**3x** inference speedup, **25%** lower latency), served through Flask REST APIs across **3 production lines** on a Dockerized AWS backend

PROJECTS

PaperTrail: SEC Filing Contradiction Detector | Python, LangChain, FinBERT, Kafka, Neo4j, pgvector

2026

- Streams EDGAR ingestion of S&P 500 filings through Kafka-backed microservices at **under 5-minute** latency; fine-tuned FinBERT on **5,000** labeled financial sentences
- Orchestrated an LLM agent with **5 custom tools** over hybrid retrieval: pgvector search plus a Neo4j temporal knowledge graph

ExtendingMemoryLLM: Eviction Strategies for Continual Learning | Python, PyTorch, MemoryLLM

2026

- Studied what a latent memory pool should forget: a surprise-gated eviction policy held **50%** exact match at update step 4 against **30%** for the default policy, retaining injected facts across more steps

Patrona: AI Voice Safety Companion | 1st Place, Columbia AI for Good Hackathon

2026

- Voice-first AI agent (ElevenLabs, WebRTC, Supabase, live GPS) built in **36 hours**, with silence escalation and an NLP safe-word trigger dispatching geolocated emergency SMS

Quant Portfolio Returns Dashboard | Python, SciPy, SQL, Streamlit, Docker

2025

- Computes **15+** risk metrics (Sharpe, Sortino, VaR, CVaR, Beta, Max Drawdown), a SciPy mean-variance optimizer tracing the efficient frontier, and a **1,000+** path Monte Carlo engine

TECHNICAL SKILLS

Languages: Python, SQL, TypeScript, Bash

AI: LLM APIs and tool use, agentic workflows, retrieval (RAG, pgvector), prompt engineering, LoRA fine-tuning, PyTorch

Data & Delivery: PostgreSQL, Neo4j, Kafka, Pandas, Docker, Azure, AWS, Git, CI/CD, FastAPI, React, REST APIs