

SABARISH RAJA RAMESH RAJA

Chicago, IL | 872-899-0544 | sabarishds03@gmail.com | LinkedIn | Portfolio | Github

PROFESSIONAL SUMMARY

AI/ML Engineer with hands-on experience building agentic AI systems, multi-agent LLM pipelines, and production analytics. MS in Computer Science from Illinois Institute of Technology, Chicago (GPA 3.8). Track record shipping end-to-end AI products across multi-agent orchestration, prompt engineering, and generative AI.

TECHNICAL SKILLS

Languages: Python, JavaScript, TypeScript, SQL, R, C++

AI Agents & LLMs: AI Agents, Natural Language Understanding, Dialogue Management, LLM Orchestration, Multi-Agent Systems, Agentic AI, Prompt Engineering, RAG, Vector Search, Schema-Constrained Generation, NLP, Transformers

ML & Frameworks: Model Development, Evaluation & Deployment, PyTorch, TensorFlow, Scikit-learn, Hugging Face, LangChain, LangGraph, CrewAI, Pydantic, Anthropic Claude, Google Gemini

Cloud & Backend: Google Cloud Platform, BigQuery, Supabase, PostgreSQL, MongoDB, Docker, REST APIs, Node.js

Data & Analytics: Statistical Analysis, Data Pipelines, ETL, GA4, Power BI, Tableau, Pandas, NumPy

PROFESSIONAL EXPERIENCE

SEO & Data Intern, SubjectToClimate

Oct 2025 – May 2026

Remote, Chicago, IL | EdTech (K-12 climate education platform)

- Designed an automated KPI dashboard using Supermetrics and Google Analytics 4 (GA4), tracking 20+ metrics across 8 state hub dashboards, reducing manual reporting time by 75% and eliminating 10+ hours per week of data extraction.
- Developed an n8n automation workflow that fetches daily climate news across 4 grade bands, uses an LLM to score relevance and auto-assign content categories, then publishes approved articles to the CMS and routes low-confidence results to the content team via Slack.
- Engineered an AI analytics platform unifying GA4, Google Search Console, FullStory, and Meilisearch into a Supabase data warehouse, enabling natural-language querying via the Claude API with automated SQL generation across multi-million-row tables.

PROJECTS

AI Marketing-Video Generation Pipeline

June 2026 – Present

- Orchestrated a multi-agent pipeline (CrewAI, Anthropic Claude, Google Gemini on Vertex AI) that turns a product URL into a finished short video with synced AI voiceover, generated visuals, and motion using a Remotion (React/TypeScript) renderer.
- Built a fidelity layer with Gemini Vision QA guard, automatic retry on garbled text, guaranteed product screenshot placement, and per-project background-style controls; validated by a 227-test pytest suite.

StyleSense – AI Wardrobe & Outfit Planner

Mar 2026 – Present

- Solo-engineered a production-grade React Native app with a Supabase backend, integrating a multi-model Claude API (Sonnet 4.5 vision and Haiku 4.5 chat) for clothing classification and conversational outfit generation.
- Architected a multi-stage outfit pipeline: natural-language input, PostgreSQL filtering by formality and season, Claude Sonnet 4.5 visual reasoning over wardrobe images, and 3 ranked outfit combinations with gap analysis and affiliate shopping links.

Claude-Hindsight – Open-Source Developer Tooling for Claude Code

July 2026

- Built an open-source CLI (TypeScript/Node.js) that indexes Claude Code session transcripts into SQLite and surfaces insights via a terminal briefing, three-view web dashboard, and a dual-row statusline wired into Claude Code's settings.json hook system; covered by a Vitest suite.
- Engineered an Instruction Audit that cross-examines CLAUDE.md rules against session transcript corpora to classify each as followed / violated / dead / unchecked, and a living Architecture doc feature that runs an agentic codebase exploration and injects a plain-English project map at SessionStart.

EDUCATION

Illinois Institute of Technology, Chicago, IL

May 2026

Master of Science, Computer Science | GPA 3.8

Relevant Coursework: Machine Learning, Advanced Database Organization, Medical Informatics & AI.

RESEARCH & PUBLICATIONS

Scalable Ensemble Learning for Credit Card Fraud Detection

Jan 2025

- Built an ensemble model (Isolation Forest, XGBoost, SVM) for highly imbalanced credit-card fraud detection, achieving 97.0% Precision-Recall and 96.3% ROC-AUC with ADASYN and NearMiss resampling; published in AIP Conference Proceedings (ICRAIC 2024).