

Krish Bhatt

408-460-0935 | kbhatt20@terpmail.umd.edu | [linkedin.com/in/kbhatt20](https://www.linkedin.com/in/kbhatt20) | github.com/KrishB-2005 | krish-bhatt-eight.vercel.app

EDUCATION

University of Maryland

Aug. 2023 – May 2027

Bachelor of Science in Computer Science, Minor in Business Analytics

College Park, MD

- GPA: 3.67
- **Teaching Assistant:** CMSC131 Object-Oriented Programming I
- **Research:** Abellon Clean Energy, Waste-to-Energy: which feedstock properties predict efficient conversion, over **100,000+** rows of sensor data. Deployed to operations, **+23%** productivity.

EXPERIENCE

EY

May 2026 – Present

Software Engineer

Remote

- Built a full-stack GenAI platform (FastAPI, React, LangChain) that extracts and enriches enterprise KPIs from unstructured client documentation at **89%** extraction accuracy, cutting per-client turnaround from roughly a month of manual analysis to a matter of days.
- Shipped to production and adopted by consultants across the engagement team, processing **100+** clients in its first two months against a manual baseline of roughly one client per month.
- Designed a **3-stage** validation engine and chunked Map-Reduce pipelines so documents exceeding a single model context window process in parallel and recombine without losing cross-section references.
- Automated the downstream design step end to end: requirement documents in, structured Snowflake DDLs and SQL transformation logic out, removing a manual schema-authoring stage from delivery.

Radyant Inc.

May 2025 – Aug. 2025

Software Engineer Intern

Worthington, OH

- Built a Python document intelligence pipeline using LLM APIs and multi-stage prompt chaining, scaling VA claims processing from **10 to 50+** documents/day and cutting average claim review time roughly **60%**.
- Deployed the pipeline into live VA adjudication workflows behind a Flask REST API, iterating on prompt templates and output schemas to hold extraction accuracy across **15+** distinct document formats.

PROJECTS

Keepsake Diary | Co-Founder | 3D Printing, Local LLM Fine-Tuning, Guardrails | [site](#)

March 2026 – Present

- Co-founded a hardware startup building a **3D-printed interactive diary** for young Harry Potter fans, powered by a **locally fine-tuned open-source LLM** with custom safety guardrails.
- Owned **product and technology**, driving **30M+** views and **1.2M+** likes on launch content, building a waitlist of **1,000+** customers and fulfilling the first **20** orders.

ResuMatch | Python, Pydantic, OpenAI/Anthropic APIs, BM25 | [live demo](#) | [github](#)

Aug. 2025 – Present

- Built a resume matcher that recognises when a posting and a resume use different words for the same skill, across a **166-entry** canonical vocabulary, and flags the gap an automated keyword screen would reject.
- Validated it against two hand-labelled benchmarks that exposed six defects **140** passing unit tests had missed, debugged each to root cause, and held the Python and TypeScript implementations to byte-identical output.

TaxOrchestra | Python, AWS Bedrock, Elasticsearch, pypdf | [live demo](#) | [github](#)

June 2025 – Present

- Built a multi-agent system that reads W-2s and 1099s and fills a genuine IRS Form 1040 end to end, grounding each value against a RAG pipeline over an Elasticsearch index of IRS rules: **22/22** fields correct, verified by reading the written PDF back.
- Inferred meaning for the form's **229** undocumented fields from page layout alone, generalising to **11** further IRS forms and **736** fields at **80%** label recovery.

G.E.O.P.A.L | Python, OpenCV, Google Earth Engine, Verbwire API | [devpost](#)

Sept. 2023

- **1st place of 30+ teams**, HopHacks 2023. Full-stack platform surfacing Google Earth Engine satellite insights for NGO decision-making, with OpenCV facial-recognition auth and on-chain identity verification.

TECHNICAL SKILLS

Languages: Python, Java, C/C++, TypeScript, JavaScript, SQL

AI/ML: RAG pipelines, vector search & embeddings, multi-agent orchestration, LLM fine-tuning, LangChain, LangGraph, Scikit-learn

Backend/Web: FastAPI, Flask, Spring Boot, Node.js, React, HTML/CSS, REST APIs, GraphQL, Pydantic

Data: PostgreSQL, MongoDB, Redis, DynamoDB, Elasticsearch, Snowflake, Pandas, NumPy

Infrastructure: AWS (Bedrock, S3), Docker, Kubernetes, Kafka, Terraform, CI/CD (GitHub Actions), Git, Linux, Vercel