

## PROFESSIONAL SUMMARY

AI Engineer with 2+ years shipping production LLM and autonomous systems — published RAG pipeline validated on a 50MB+ document corpus, deployed on AWS Bedrock, and selected for internal productization targeting TB-scale deployment. Active PhD research in autonomous mission systems and command & control architectures for SAR drone operations. Strong across the full stack: vector databases, agentic workflows, cloud infrastructure, and Python. Looking to bring that depth to an early-stage team building real AI products.

## EDUCATION

**Worcester Polytechnic Institute (WPI)**

*PhD Artificial Intelligence (In Progress, MS completed 2025)*

Worcester, MA

2021 – Present

- **Research focus:** Command and control architectures for SAR drone operations; mission effectiveness modelling for AI-enabled autonomous systems.
- **MS ECE (2025):** Thesis on Context-Aware AI for Requirements Analysis using Production RAG — published *Procedia Computer Science* (2025).
- **BS ECE (2024):** Graduated with Computer Science Minor.

## PROFESSIONAL EXPERIENCE

**Draper**

*AI/ML Engineer (Draper Scholar)*

Cambridge, MA

Jun 2024 – Dec 2025

- Architected and shipped a production RAG system for large-scale technical document analysis — AWS Bedrock (Claude 3 Sonnet), FAISS, Titan Embeddings — validated on a 50MB+ document corpus with sub-15s latency; methodology published in *Procedia Computer Science* (2025) with a second full-system paper in submission.
- System selected for internal productization at Draper — currently advising on architecture as it is prepared for vendor deployment at TB scale.
- Built an agentic document analyzer using chain-of-thought reasoning, structured function calling, and Pydantic schema validation to classify 10+ categories of quality issues with deterministic, auditable outputs.
- Engineered a document ingestion pipeline across PDF, DOCX, TXT, and Markdown — intelligent chunking, metadata extraction, batch embedding generation, and optimized FAISS index management.
- Designed modular software architecture for autonomous drone systems in Python — mission logic, real-time health monitoring, and operator interfaces with clear separation of concerns and graceful fault handling.
- Built Python libraries and CLI tools for RAG orchestration, semantic search, and LLM evaluation, cutting onboarding and iteration time for future AI work on the team.

**Arcfield**

*Software Engineer Intern — Cloud Infrastructure & MLOps*

Chantilly, VA

Jul 2023 – Aug 2023

- Implemented containerized monitoring services using Docker and deployed across Azure GCC-High environments, improving system observability and reducing mean time to detection (MTTD) by 60%.
- Optimized CI/CD pipelines with automated service configuration, deployment scripts, and infrastructure-as-code, reducing release cycles from 2 hours to 20 minutes and enabling faster ML model deployment iterations.
- Designed telemetry pipelines for model performance tracking, drift detection, and automated anomaly alerting using Prometheus, Grafana, and custom Python monitoring agents.

**AbbVie**

*Cloud Engineer Intern — Data Engineering & Analytics*

Chicago, IL

Jun 2022 – Mar 2023

- Built and deployed a 5-node EMR/Spark cluster for clinical analytics workloads, architecting PySpark ETL pipelines that reduced data processing time by 3.3x through query optimization and preprocessing.
- Developed production-grade data pipelines with AWS ETL tools, automated data quality checks, LDAP access permissions, and auto-alerts for performance metric anomalies.
- Designed Grafana dashboards and Prometheus alerting for monitoring pipeline health, query performance, SLA compliance, and data freshness metrics.

## PROJECTS &amp; PUBLICATIONS

**SEMP RAG Analyzer (MS Thesis)** · [github.com/krishpatel1077/semp-rq-debt-analyzer](https://github.com/krishpatel1077/semp-rq-debt-analyzer) — Production RAG system for automated document quality analysis. Claude 3 Sonnet + FAISS vector search + real-time streaming web UI. Published in *Procedia Computer Science*, Vol 268, 2025. Demo: [video link](#).

**AME POC — SAR Mission Simulator** · [github.com/krishpatel1077/AME\\_POC\\_2026](https://github.com/krishpatel1077/AME_POC_2026) — Simulation framework for SAR drone operations with real-time Autonomous Mission Effectiveness (AME) scoring across five sub-dimensions (coverage, timeline, resource, comms, operator). Grounded in two preprints on mission effectiveness for AI-based systems of systems and SE4xAI verification frameworks.

**Publication (in submission):** Full-system paper covering end-to-end production RAG architecture, evaluation methodology, chunking strategies, and agentic reasoning pipeline.

**Preprint:** “A Review of Mission Effectiveness for AI-Based System-of-Systems” — surveys 60+ years of ME literature across probabilistic, qualitative, MBSE, and AI traditions; defines the AME framework.

**Preprint:** “A Systems Engineering Approach for V&V of AI-Enabled Autonomous COTS Drone Systems” (with GMU/WPI collaborators) — SE4xAI framework for V&V of autonomous drone systems using xAI diagnostics. Submitted to CSER 2026.

## TECHNICAL SKILLS &amp; CERTIFICATIONS

**AI/ML Stack:** LLMs (Claude, GPT-4, Bedrock API), RAG Systems, Vector Databases (FAISS, Pinecone), Prompt Engineering, Chain-of-Thought Reasoning, Pydantic, Function Calling, Embeddings, PyTorch, TensorFlow

**Software:** Python (Expert), C++, FastAPI, Flask, Django, Microservices, Async/Await, REST APIs, GraphQL

**Cloud & Infrastructure:** AWS Bedrock/Lambda/S3/EMR, Azure, Docker, Kubernetes, Terraform, Serverless, Infrastructure-as-Code

**Data Engineering:** PySpark, Apache Airflow, ETL Pipelines, SQL, NoSQL, Data Modeling, Stream Processing, Telemetry

**Certifications:** AWS Solutions Architect Professional (SAP-C02), AWS Associate (SAA-C03), Kubernetes (CKAD), Azure (AZ-305, AZ-104), Red Hat (RHCSA)