

Youdahe Asfaw

| youdaheasfaw@gmail.com | [linkedin.com/in/youdahe](https://www.linkedin.com/in/youdahe) | github.com/youdahe |

EDUCATION

Gustavus Adolphus College

Bachelor of Science in Computer Science & Statistics

June 2028

3.5/4.0

- **Relevant Coursework:** Introduction to Data, Object-Oriented Programming, Data Structures and Algorithms

TECHNICAL SKILLS

Languages: Python, TypeScript, Java, JavaScript, HTML/CSS, Go, Rust

Frameworks: FastAPI, LangChain, Tesseract.js, React.js, Express, GraphQL, Spring, Three.js

Developer Tools: Postman, Git, Selenium, Postman, Node.js, Jenkins, MongoDB, Docker, AWS, SQL (MySQL), Kubernetes, Cassandra, Ansible

EXPERIENCE

Incoming Software Engineer Intern

Liberty Mutual Insurance

May 2026 – Aug 2026

Portsmouth, NH

Software Engineer Intern

Medica

Aug 2025 – Present

Minnetonka, MN

- Built and deployed internal employee-enrollment systems using **Java, Spring Boot, REST APIs, and Maven**, improving workflow automation and reducing manual processing dependencies.
- **Digitized paper-based onboarding** by integrating automated data validation and persistence layers, **boosting employee enrollment speed by 65%**, improving **system reliability**, and enhancing overall **operational efficiency** across HR systems.

Machine Learning Research Assistant

Gustavus Adolphus College

Feb 2025 – Jun 2025

Saint Peter, MN

- Built an ML-powered, agentic voice-guided form-completion system using **Python, Azure Form Recognizer, LayoutLMv3**, reducing completion time from 60 to 15 minutes.
- Improved form completion for visually impaired users **from 45–60 minutes down to 15 minutes**, enabling faster, more accurate submissions via automated email or print outputs.

Machine Learning Researcher

Hugging Face

Sep 2024 – Jan 2025

San Francisco, CA

- Implemented LLM evaluation metrics in Python and ran **benchmarks across 10+ model configurations**, measuring reasoning accuracy, long-context performance, and memory retention over multi-turn interactions.
- Improved reproducible model evaluation workflows, reducing manual analysis time and enabling **consistent comparison of LLM performance** through documented open-source contributions

Software Engineer Intern

Medtronic

Aug 2023 – Aug 2024

Mounds View, MN

- Engineered automated backend pipelines by **replacing manual Excel processes with GoLang and SQL-based services**; **designed modular APIs, optimized database queries with indexing and batching**, and implemented concurrency to reduce data processing time by **50%** while improving system reliability for **8+ R&D teams**
- Spearheaded end-to-end pipeline automation, **implementing data validation and error-handling routines**, coordinating 12+ user acceptance tests, and increasing processing efficiency by **30%** for teams tracking **over \$2.5M** in R&D projects

PROJECTS

Docere | *React, TypeScript, Node.js, Express, SQL, Auth0, Redis, Docker, Kubernetes*

June 2025 – Present

- Architected a **memory-augmented AI tutoring platform** serving **100+ students across 2 schools**, featuring a three-source RAG layer (Qdrant vector search, PostgreSQL student profiles, interaction history) with token-budgeted retrieval, a UCB1 multi-armed bandit that autonomously selects and evolves teaching strategies, and Bayesian Knowledge Tracing for per-concept mastery estimation—improving AI response relevance by 60% and student task completion by 40%.
- Built a full-stack async pipeline (**FastAPI, React 19, Redis/ARQ task queues, LangGraph agent orchestration**) with dual LMS adapters (Canvas + Moodle via LTI 1.3), Docker containerization, and GitHub Actions CI/CD with automated staging/production deployments **securing \$10,000 in funding** and launching pilots that ingested real gradebooks, course materials, and Google Calendar scheduling

Taylor.ai (HackHarvard 2025) | *Python, FastAPI, JavaScript, Google Gemini, Selenium*

Oct 2025 – Oct 2025

- Built an AI-powered fashion crawler with **Python/FastAPI and Selenium**, using **Google Gemini to turn user prompts into real-time searches across 30B+ items** and surface cached results instantly via Elasticsearch-style querying.
- Engineered personalized outfit generation with Gemini Nano, mapping items onto user photos with **90%+ alignment accuracy** and delivering pricing, product links, and dynamic recommendations for users.