

John Holik

Chicago, IL · 239-989-2566 · johnholikiii@outlook.com · linkedin.com/in/john-holik-iii · github.com/John-Holik

AI engineer focused on LLM application development and rapid prototyping. I build LLM systems end to end (prompt design, structured output, data/API integration), validate their feasibility with rigorous evaluation against human ground truth, and document them for clean handoff to production teams. I work in AI coding tools daily, using Claude Code as my primary environment and extending it with custom skills, agents, and multi-agent workflows to prototype quickly with minimal infrastructure.

TECHNICAL SKILLS

Languages: Python, C++, SQL, Bash, PowerShell

AI and LLM: OpenAI, Anthropic (Claude), and Gemini APIs; prompt engineering; structured JSON-schema output; LLM evaluation vs human ground truth; AI coding tools (Claude Code) with custom skill, agent, and multi-agent workflow development; RAG and embeddings

ML and Data: scikit-learn, pandas, NumPy, SciPy, model evaluation (ROC-AUC, precision/recall/F1, calibration, Cohen's kappa), uncertainty quantification (Bootstrap, EnKF, GLUE), clustering (K-means, UMAP, HDBSCAN)

Web and Backend: React, FastAPI, Node.js/Express, PostgreSQL

Tools: Git/GitHub, Jupyter, matplotlib/plotly, LaTeX, Linux

EXPERIENCE

Undergraduate Research Assistant, then Independent Researcher (ElectriAI)

Aug 2025 – Present

U.A. Whitaker College of Engineering, Florida Gulf Coast University

- Built and iterated an end-to-end LLM pipeline over an electrical-construction media corpus: data collection, GPT-based classification, human expert annotation, and rigorous evaluation.
- Tested feasibility and performance of the model's behavior against human ground truth: 0.8965 Cohen's kappa human-vs-model agreement with per-class precision, recall, and F1 on a held-out gold set.
- Redesigned the approach from unsupervised LDA topics to a supervised, human-validated classification schema after early prototypes proved unreliable.
- Built a reproducible evaluation harness (fixed seeds, versioned data) and documented method decisions so the work hands off cleanly; authoring a journal manuscript (in progress).

PROJECTS

CapSight AI, LLM Pitch-Deck Grading Platform

React, FastAPI, PostgreSQL, Node/TypeScript, OpenAI

- Prototyped and shipped a three-service application that scores uploaded pitch decks across six categories with an LLM, persisted to PostgreSQL with a history view.
- Engineered prompts for reliability: enforced structured output via OpenAI `json_schema` and made grading deterministic (seed 42, temperature 0) for repeatable results.
- Hardened the grading core for handoff with adaptive retry, score clamping, an unreadable-deck sentinel, a fallback result, and non-blocking database writes so grading still returns if the database is down.

Retrieval-Augmented (RAG) Chatbot, ElectriAI "Ask ElectriAI" (deployed)

Python, Gemini embeddings, cosine retrieval, Cloudflare Workers/Pages

- Built and deployed a retrieval-augmented chatbot grounded in a knowledge-base wiki: pre-computed `gemini-embedding-001` vectors over the corpus, cosine-similarity retrieval, and grounded generation via `gemini-2.5-flash`.
- Shipped to production on Cloudflare Pages with a Worker proxy that isolates the API key from the client and a KV namespace for rate limiting; serves a static build with committed embedding artifacts (no build step).

Credit-Risk ML, Loan-Default Prediction

Python, scikit-learn

- Benchmarked 14 models on 32,416 loans with an explicit train-vs-test overfit check; selected Gradient Boosting at 0.95 test ROC-AUC with a small generalization gap.
- Went beyond accuracy with probability calibration (Brier 0.05), ROC/PR curves, partial dependence, and permutation importance; fully reproducible (seed 42, pinned library versions).

EDUCATION

Florida Gulf Coast University, Fort Myers, FL

May 2026

B.S. in Computer Software Engineering · Dean's List (Junior and Senior Years)

Relevant coursework: Data Structures and Algorithms, Machine Learning, Data Mining, Data Engineering, Operating Systems, Computer Security, Software Engineering, AI Entrepreneurship