

John Doan

Backend / API + Data Engineering (Python)

Brief resume for quick review — full resume available upon request.

Remote / Open to relocate | john@johnvdoan.com | linkedin.com/in/johnvdoan | jdoan5.github.io/home/

Summary

Associate-level engineer focused on building reliable backend APIs and data pipelines in Python. Comfortable with FastAPI, SQL, ETL workflows, and project documentation.

Experience

Independent Developer — Self-Directed Projects (Jan 2023 – Present)

- Built portfolio projects (Fraud Detection, Customer Churn) using Python/SQL: data prep, feature engineering, baselines, and threshold trade-offs.
- Packaged FastAPI scoring endpoints with tests, structured logging, Docker, and CI; published short runbooks.
- Produced lightweight model notes and stakeholder views (confusion matrix at threshold, basic lift/decile table).

Selected Projects

- Customer Metrics Pipeline & API — Built an ETL pipeline and exposed scoring via a FastAPI endpoint with Swagger/OpenAPI docs.
- RAG Mini Chat (Stages 1–3) — Implemented staged retrieval workflows and basic run logging for debuggable, grounded answers.
- Support Ticket Analytics — Built a small DuckDB-based data mart and analyzed ticket text to surface recurring operational themes.

Skills

Backend / API: Python, FastAPI, Pydantic, REST, OpenAPI/Swagger

Data: SQL, ETL, DuckDB, SQLite, Pandas

ML (project work): scikit-learn, classification baselines, evaluation/logging, retrieval (RAG basics)

Professional: clear documentation, root-cause analysis, stakeholder communication, cross-team coordination, ownership

Certifications

- IBM Generative AI Engineering Professional Certificate (Coursera/IBM)
- AI Agents Using RAG and LangChain (Coursera/IBM)
- PCEP™ — Certified Entry-Level Python Programmer

Education

B.S. Computer Science