ALEX MITKOV

Software Engineer · Backend Systems, OOP & Real-Time Architectures
Las Vegas, NV · (702) 758-3026 · alex.mitkov.dev@gmail.com · amitkov.dev · linkedin.com/in/amitkov · github.com/amitkov-sudo

SUMMARY

Software engineer with a strong OOP foundation and production experience across C++, Java, and Python. Designs and ships client-server systems, REST APIs, and event-driven backends in performance-sensitive environments. Practiced in TDD, CI/CD workflows, profiling, and debugging. Applies clean architecture and modular design to build maintainable, team-ready codebases.

SKILLS

Languages: Java, C++, Python, JavaScript, Fortran, TypeScript, SQL, Bash
OOP & Architecture: Object-oriented design, client-server architecture, REST APIs, Kafka event processing, modular library design
Backend Frameworks: Spring Boot, FastAPI, Flask, Node.js, Next.js, React
Debugging & Performance: Profiling, log-based debugging, query optimization, B-tree indexing, latency benchmarking
Testing & CI/CD: TDD, JUnit, pytest, unit & integration testing, GitHub Actions, CI/CD pipelines, Docker
Dev Tools: Visual Studio, IntelliJ, Git, Jira, Linux, Docker
Database: PostgreSQL, H2, SQL, MySQL, schema design, query optimization, vector search
AI & ML: LlamaIndex, FAISS, Chroma, RAG, PyTorch, scikit-learn, XGBoost
Computer Vision & OCR: OpenCV, PIL, Tesseract, PaddleOCR, EasyOCR, PyMuPDF, pdfplumber

EXPERIENCE

AI/ML Engineering Extern · Pfizer (via Extern HQ) *Mar 2026 – Present*

- Architected a multi-engine OCR document processing pipeline in Python (OpenCV, PIL, Tesseract, PaddleOCR, EasyOCR): preprocessing scanned documents and emitting structured JSON with field-level coordinates across multiple document types
- Benchmarked three OCR engines against real pharmaceutical scans; identified accuracy divergence by document type and delivered a data-backed engine recommendation
- Designed document classification and routing logic in Python to categorize incoming files and dispatch them to the correct extraction pipeline: reducing processing ambiguity and enabling modular extension
- Built a RAG retrieval system using LlamaIndex, FAISS/Chroma, and metadata filtering over Gemini and open-source LLMs, surfaced through a Gradio/Streamlit interface

Graduate Researcher & Instructor · University of California, Merced *Aug 2024 – May 2026*

- Implemented quantum computing and molecular simulations in Python, C++, and Fortran; supported instruction for 100+ students across Applied Mathematics and Chemistry and communicated research findings to both technical and non-technical audiences

PROJECTS

RailCast AI: Real-Time Transit Decision System · *Java, Spring Boot, FastAPI, PostgreSQL, Python, Docker, Airflow, React*

- Designed a Spring Boot REST API serving a React client: resolving live GTFS-RT stop times, falling back to an XGBoost regressor (val MAE 1.22 min) when real-time data was unavailable, and defaulting to static schedule as a last resort across 40,968 ingested rows
- Reduced FastAPI inference to sub-13ms mean over 500 warm calls by profiling the prediction path end-to-end and adding B-tree indexes on high-cardinality PostgreSQL columns
- Automated the full data and model lifecycle with a daily Airflow DAG: enforced releases with a CI script spanning Python, Maven, and frontend builds

JPMorgan Chase Advanced Software Engineering Program · *Java, Spring Boot, Kafka, H2, JPA, REST APIs*

- Built Midas Core in Spring Boot: implemented a Kafka listener deserializing Transaction messages into typed Java objects, enforced validation logic against user existence and account balances before commit, and persisted state via Spring Data JPA with H2
- Integrated with an external Incentive API via HTTP POST, applied conditional incentive logic to recipient balances, and exposed a REST endpoint returning JSON-serialized account state

RetentionIQ: Containerized Inference Service · *Python, FastAPI, Docker, AWS EC2, scikit-learn*

- Improved churn model ROC-AUC by 10–15% over baseline across 7K+ records through feature engineering, model selection, and hyperparameter tuning with automated drift detection
- Deployed a containerized FastAPI service on AWS EC2: Pydantic schema validation, sub-100ms response under single-instance load, full CI/CD pipeline

Additional projects available at github.com/amitkov-sudo

EDUCATION

M.S. Computer Science: Artificial Intelligence · Monash University · Part-Time *Exp. May 2028*
M.S. Computational Chemistry: Quantum Computing · UC Merced *May 2026*
B.S. Natural Science · American Public University · GPA: 3.92 · Summa Cum Laude *Aug 2024*