

Artem Moshnin

Full-Stack Software & Machine Learning Engineer with 3+ years building production-grade enterprise systems across distributed architectures, AI pipelines and cloud infrastructure. Delivered scalable SaaS platforms (Fluxo), high-concurrency transactional engines (LedgerX) & advanced RAG systems (DocuMind). Experience spans architecting ML systems at AMD & ST Engineering and leading full-stack development at SQA and Yapp Agency. Technical expertise: end-to-end ML pipelines, event-driven architectures, distributed systems, Docker/Kubernetes & AWS. Graduating NUS - Electrical Engineering & Data Analytics, GPA 4.25/5, May 2026. **SG Work Eligibility:** Authorized for Employment Pass upon graduation (via MOE Tuition Grant; supported by [NUS and MOM letter](#))

artem.moshnin@u.nus.edu

+65 8922 6877

Singapore

artemmoshnin.com

github.com/amoshnin

linkedin.com/in/amoshnin

Work Experience

AMD (Advanced Micro Devices) > Semiconductor Technology Corporation

Singapore, Singapore

Machine Learning Engineer Intern, Semiconductor Reliability Lab

Aug 2024 - Jan 2025 (Internship)

- Spearheaded all design, development & deployment of ML systems at AMD's Reliability Engineering, enhancing testing reliability & efficiency.
- Developed real-time data pipelines & predictive alerting system, achieving >85% precision using threshold-based, statistical & ML methods.
- Built automated ML infrastructure with Random Forest & XGBoost models, with online training, feature selection and hyperparameter tuning.
- Automated data analysis & visualisation through production-grade dashboards, enabling analyses, pattern recognition and trend visualisation.
- Led technical presentations, meetings and collaboration with IT, vendors & engineers. Authored technical documentations & training sessions.

ST Engineering > Multinational Technology & Engineering Group

Singapore, Singapore

Machine Learning Engineer Intern, Vision-Language Models

May 2024 - Aug 2024 (Internship)

- Developed and enhanced ST Engineering's VisionX platform, a vision-language framework for real-world challenges.
- Developed a PDF analysis tool to extract insights from embedded images, enabling text & image-based query answering with 93% accuracy.
- Implemented a traffic incident analysis system, integrating: image analysis, video analysis and scene monitoring, achieving 86% accuracy.
- Designed & developed an Image Retrieval System by Text/Image Queries. Improved search performance by 95% by transitioning to VectorDB.
- Recognised for outstanding contributions and high performance with the Intern@ST Engineering Award (2024).

SQA Consulting > IT Consultancy Company

Murcia, Spain

Full-Stack Software Engineer

May 2022 - Sep 2022 (Full-time)

- Developed secure backend applications in Java using Spring Boot, applying TDD principles, resulted in 20% reduction in post-release defects.
- Utilised AWS services, including Fargate, API Gateway, Cognito, to containerise app, expose REST APIs and implement secure authentication.
- Collaborated with teams of software engineers & managers with Jira, fostering effective communication & streamlined project management.
- Led migration of monolithic backend to microservices architecture with Spring Boot, enhancing scalability & reducing response times by 32%.

Yapp Agency > Software Development Agency

Madrid, Spain

Co-Founder & Lead Software Engineer

May 2020 - May 2022 (Full-time)

- Directed the project lifecycle of client requirement collection, project estimation, technical specification formulation and development.
- Developed and maintained over 10+ cross-platform mobile applications (iOS and Android) using React Native & SwiftUI frameworks.
- Created and managed 5+ front-end and backend applications with React, Next.JS and Node.JS, ensuring responsive & scalable solutions.

Technical Projects

• **Fluxo (Workflow Automation & Orchestration Platform)** > TypeScript, Next.js, tRPC, Prisma, PostgreSQL, Inngest, Sentry [Live Demo](#) | [Docs](#)

Enterprise workflow orchestration SaaS with 30+ AI/integration nodes. Architected event-driven execution engine (Inngest), visual graph editor (React Flow), and multi-tenant RBAC. Implemented Human-in-the-Loop governance: centralized review inbox, white-labeled client portals, policy-based routing and AI prompt learning from edits. Built complete SaaS stack: tRPC API, Prisma/PostgreSQL, encrypted credentials, Stripe billing, signed webhooks, real-time streaming. Production system handles concurrent executions with audit trails and organization-scoped access control.

• **LedgerX (High-Concurrency Transactional Ledger Engine)** > Java, Spring Boot, PostgreSQL, Docker, Testcontainers [Live Demo](#) | [Docs](#) | [GitHub](#)

ACID-compliant payment engine handling 500+ TPS with zero integrity violations. Implemented pessimistic locking with deadlock prevention, idempotent API design, and double-entry accounting. Built async audit pipeline with transaction-safe logging. Spring Boot/PostgreSQL backend, Next.js dashboard, validated with Testcontainers and k6 load tests achieving 100% success rate under high concurrency.

• **DocuMind (Hybrid RAG Document Intelligence Platform)** > Python, FastAPI, LangChain, ChromaDB [Live Demo](#) | [Docs](#) | [GitHub](#)

RAG platform combining Chroma vector embeddings and BM25 for hallucination-resistant PDF analysis. Architected hybrid retrieval with query reformulation, streaming generation with structured citations, and vendor-agnostic LLM routing (Groq vs OpenAI). Implemented BYOK security and LLM-as-Judge evaluation. FastAPI backend with sub-second TTFT, Next.js SSE streaming frontend.

Education

• **National University of Singapore (NUS)** > 8th Best University Worldwide (QS Ranking)

Aug 2022 - May 2026 (Present) (GPA: 4.25/5)

Bachelor of Engineering (B.Eng.) in Electrical Engineering (Primary Major)

(Singapore, Singapore) [Grades Transcript](#)

Relevant coursework: Electronic Circuits, Electrical Energy Systems, Electromagnetics, Signals and Systems, Microcontroller Programming and Interfacing, Digital Circuit Design, Embedded Hardware System Design, Electrical Engineering Practices & Principles, Calculus, Differential Equations & Linear Algebra.

Bachelor of Science (B.Sc.) in Data Analytics (Second Major)

Relevant coursework: Foundations of Data Science, Machine Learning, Data Structures and Algorithms, Probability, Mathematical and Digital Statistics, Data Visualisation Tools, Computer Vision, Image Processing and Analysis, Regression Analysis, Convex Optimisation and Systems Thinking & Dynamics.

• **Computer Science Bootcamp (Educative.io, Online)**

May 2025 - Feb 2026

Self-Directed Studies via Educative.io (72 courses, 33 skill paths)

[Certifications Dashboard](#)

Systems & Architecture: Distributed systems, microservices, advanced concurrency, scalability & database internals. AI/ML Engineering: LLM engineering, agentic AI & MCP, RAG, MLOps, model evaluation, deep learning and NLP. Infrastructure and DevOps: AWS (Solutions Architect, Developer and SysOps), Kubernetes, Kafka, Docker, Terraform, CI/CD and serverless. Software Engineering: Java, Python, C++, Go, Rust, OOP, low-level design and testing patterns.

Languages

- Fluent (C1) - English, Spanish, Russian
- Intermediate (B1) - Chinese, German
- Elementary (A2) - French