

Andreas Andronikakis

Full-Stack, AI & Data Engineer | Founding Engineer

Manchester, UK | aandronikakis@gmail.com | linkedin.com/in/andreas-andronikakis | github.com/aandronikakis | aandronikakis.github.io

SUMMARY

Founding Backend & AI Engineer with 9+ years building software, data and AI systems across research, high-performance and fast-growing startup environments. Moves fast, learns quickly and thrives on technical intensity, shipping practical production AI – multi-agent platforms, agent orchestration, large-scale data and enrichment pipelines, and cloud-native backends – using Anthropic, OpenAI, Azure and TypeScript, with hands-on machine-learning evaluation, fine-tuning and rapid experimentation through modern AI-assisted engineering workflows. Research background in graph databases, large-scale experimentation on noisy, experimental datasets, runtime optimisation and high-performance computing, combining deep technical ownership with end-to-end delivery from data infrastructure to production. Most recently, shipped an AI-assisted document-automation product solo, architecture through to signed release infrastructure and production support.

EXPERIENCE

Founder & Sole Engineer – Full-Stack, Applied AI & Product Jul 2026 – Sep 2026
Ktimatologio Pipeline / *Agentic automation platform (own IP, paying customer)* / *Remote*

- Shipped an **agentic land-registry automation platform** as **sole engineer** – **~29k lines** of Node.js/Python whose agents reconcile **1,797 cadastral parcels** across **240 client portfolios** in a **~4-hour unattended run**.
- Built the AI layer: a **vision-LLM agent** (Anthropic/OpenAI) that takes over when PDF parsing fails, **geometrically self-verified** against each document's declared area so no unverified output reaches a legal filing.
- Engineered **Playwright browser agents** that autonomously collect **4,100+ official documents** into **2,065 property-right records**, reconciled against the source's declared totals to surface **184 silently missing rights**.
- Delivered the full product alone: **agents propose, humans approve** via a multi-tenant **39-route** review platform with **AES-256-GCM** vaults, in a self-contained **Windows distribution** with **Ed25519-signed** auto-updates.

Founding Software Engineer – Backend & AI Sep 2025 – Jul 2026
PG:AI / *Remote (Manchester, UK)*

- Led the backend and AI rebuild of PG:AI from a CRM-centric product into an AI-native account-intelligence platform as the company's **sole Backend and AI Engineer**, owning architecture, delivery and production operations end to end.
- Designed and shipped a **multi-agent architecture** using Anthropic, OpenAI, Azure Functions, MongoDB, Xano and MCP tooling, enabling enterprise users to generate account strategies, executive briefings, stakeholder analyses and personalised outreach from continuously updated company intelligence.
- Built the core agentic execution framework in TypeScript on the Anthropic Messages API, powering Agent, Ask and Canvas experiences through intelligent tool routing, context management and multi-agent coordination.
- Developed AI-powered enrichment and research pipelines using Azure Functions, MongoDB, Apollo and LinkedIn/financial data sources, enabling automated analysis across **hundreds of thousands of companies and contacts**.
- Designed and shipped a public API platform of **700+ REST endpoints, 300+ Azure Functions and 20+ webhooks**, exposing account intelligence, monitoring, contact discovery and workflow automation to customers and enterprise integrations.
- Architected workflow orchestration for Studio using reusable workflow definitions and agent pipelines, automating account planning, SWOT analysis, executive-briefing generation and meeting preparation.
- Migrated the production agent platform from the Claude Agent SDK to the Anthropic Messages API, improving control over tool execution, streaming responses, observability and operational reliability.
- Integrated Salesforce, HubSpot, AgentMail, StackOne, Apollo and Slack into a unified intelligence platform, enabling automated CRM synchronisation, enrichment, outreach and workflow execution.
- Instrumented end-to-end observability and evaluation with Braintrust, Dash0 and PostHog, improving traceability, debugging and operational visibility across AI services and production workflows.

Research Engineer – AI Systems & Runtime Optimisation Jan 2025 – Jun 2025
University of Manchester / *EU Horizon Project INCODE (Grant No. 101093069)*

- Designed an AI-assisted optimisation framework for TornadoVM using Java, Python and large language models,

automating recommendation of compiler and runtime optimisation strategies for heterogeneous (GPU/FPGA/CPU) environments.

- Built automated data-generation pipelines that transformed benchmark executions, execution plans, hardware configurations and performance measurements into structured JSONL training datasets for model training and evaluation.
- Evaluated GPT-4o, DeepSeek-R1, Llama and Mistral (including local Ollama-hosted models) with custom benchmarking workflows, comparing optimisation recommendations across runtime-engineering tasks.
- Investigated LoRA fine-tuning and local inference optimisation using Unsloth, vLLM, Ollama and GGUF model formats to build domain-specific optimisation assistants.
- Implemented closed-loop execution of LLM-generated optimisation recommendations within TornadoVM, enabling automated evaluation of AI-generated runtime strategies against measured performance.
- Contributed enhancements to the open-source TornadoVM platform, extending its benchmarking infrastructure and AI-assisted optimisation capabilities.

Research Data Engineer

Jan 2023 – Dec 2024

University of Manchester / EU Horizon Project INCODE (Grant No. 101093069)

- Engineered an end-to-end experimentation and profiling pipeline in Java, Python and Bash (NUMAProfiler, PerfUtil) that benchmarked the Neo4j graph database on a ~100 GB LDBC Social Network dataset, generating and visualising large structured datasets of application- and hardware-level performance metrics across NUMA configurations.
- Ran large-scale benchmarking sweeps over 14 complex Cypher queries (656 parameter combinations, thread scaling T1–T16, 5-run averaging across ~15-day cycles) on dual-node 384 GB NUMA servers to quantify memory locality, remote-access overhead and scalability bottlenecks.
- Built a decision-tree classification model that predicts NUMA speedup vs. slowdown from profiling features (shared-access ratio, LLC miss ratio, remote-memory access %), producing a query-agnostic framework to guide memory-placement and thread-scheduling optimisation.
- Authored the profiling methodologies, statistical analyses and microarchitectural datasets that contributed to peer-reviewed publications at ACM SIGPLAN ISMM 2023 and MPLR 2023.

Research Software Engineer / Graph Database Research

2017 – 2022

University of Manchester / Horizon 2020 ACTiCLOUD (No. 732366), ENCRYPT (No. 101070670) & ELEGANT (No. 101004274)

- Extended the MaxineVM JVM profiling stack (NUMAProfiler, PerfUtil) in Java/C with an object-driven, control-flow profiling policy, dominant-thread isolation, and flamegraph visualisation of Neo4j code paths, scoping data collection precisely to the graph database's query-execution phase and cutting profiling overhead and noise versus all-object profiling.
- Engineered fine-grained hardware-counter instrumentation over the Linux perf API – adding multi-core PerfUtil operations and new last-level-cache and node prefetch event IDs – to capture per-thread, microarchitectural performance datasets (cache, prefetch, locality, CPI) for the Neo4j graph database.
- Designed and ran large-scale benchmarking and deep data analysis of Neo4j on MaxineVM across three EU Horizon research programmes, engineering structured application- and microarchitectural-level datasets (validated to within 1% of reference profilers) that underpinned data-driven, model-based performance characterisation.
- Ported the Maxine Virtual Machine to the AArch64 instruction set, debugged its C1X and T1X JIT compilers, and optimised JVM internals and memory layouts for graph-database execution, with Python and Java/JUnit test-automation validating the ARM port.
- Co-authored peer-reviewed publications (ACM Programming 2018; Springer 2019) and contributed to open-source research virtual machines, establishing rigorous, reproducible performance-engineering methodologies for managed graph-database research.

Software Engineer

Jan – Jun 2016

Focus Photography / Chania, Greece

- Designed and built a Java reporting and performance-analytics platform on top of a centralised SQL database that consolidated the company's operational data, transforming raw records into productivity and performance metrics and giving management data-driven visibility into photographer output and team performance.

TECHNICAL SKILLS

Programming: Java, Python, TypeScript, JavaScript, C/C++, C#.NET, Go/Golang, R, Bash, SQL, Cypher, XanoScript, Prolog, PHP, XML/XSL, Visual Basic.NET, ASP.NET, MIPS Assembly, VHDL, Pascal, Arduino Wiring, Flex/Bison.

AI & Machine Learning: LLMs, generative AI, multi-agent systems, agent orchestration, MCP tooling, prompt engineering, AI-assisted code optimisation, LoRA fine-tuning, dataset curation, model evaluation (Braintrust), model deployment & monitoring, automated ML workflows & pipelines, decision-tree / classification modelling; Anthropic, OpenAI, Google and open-source / local LLMs (Hugging Face, Ollama, Unsloth, vLLM); PyTorch; vision-language models, document AI / intelligent document processing, OCR & scanned-document extraction, verified LLM output (geometric validation), LLM fallback pipelines.

Data Science & Analytics: data analysis, statistical modelling, A/B testing & experiment design, customer lifetime value analysis, data visualisation & storytelling, Power BI.

Data Engineering: Neo4j, MongoDB, MySQL, SQLite, Cypher, graph data modelling, data & enrichment pipelines, JSONL dataset generation, large-scale / experimental datasets, data collection & benchmarking (LDBC), performance profiling (NUMAProfiler, PerfUtil), storage solutions, dbt, Airflow, MLflow, Apache Kafka, Spark, data warehousing, Terraform, PDF data extraction (poppler, pdf-parse), computational geometry (shoelace integration, polygon overlap, EGSA'87), CAD / DXF generation (ezdxf), data-integrity reconciliation.

Backend & Cloud: REST API design (700+ endpoints), Azure Functions, webhooks, Xano platform, Kong, Docker, Azure, AWS, Google Cloud Platform, Git/GitHub, Linux shell, Node.js, Playwright browser automation / RPA, desktop application packaging (Inno Setup), Ed25519 code signing, AES-256-GCM / scrypt encryption, multi-tenant architecture, release engineering & auto-update infrastructure.

Parallel & Heterogeneous Computing: TornadoVM, MaxineVM, OpenCL, CUDA, NUMA architectures, parallel programming, heterogeneous computing (GPU/FPGA/CPU), compiler & runtime optimisations, JVM optimisations.

Integrations & Tools: Salesforce, HubSpot, Apollo, StackOne, AgentMail, Slack, Dash0, PostHog, Linear, Atlassian, Notion, Google Drive, Figma, Loom, tl;dv, TheirStack, Cursor, Devin, JUnit, Intel PIN, DRAMSim2, Vivado, Xilinx IDE, MAGIC VLSI, Spice.

Web Development: HTML, CSS, JavaScript, WeWeb, Xano platform, XanoScript, Apache Web Server, Tomcat.

Other: academic research, research software engineering, technical writing, open-source development, EU research projects, strategic thinking, business acumen, cross-functional collaboration, mentoring & knowledge sharing, $\text{\LaTeX}$ typesetting, product ownership, end-to-end product delivery, technical documentation in Greek, GDPR-sensitive data handling.

Languages: English (C1, State Certificate), German (B2, Goethe-Zertifikat), Greek (native).

EDUCATION

M.Phil, Computer Science 2025
University of Manchester – Performance analysis for memory optimisations of managed graph databases on NUMA architectures.

Ph.D. Research (part-time), Computer Science 2017 – 2025
University of Manchester – Neo4j on MaxineVM, NUMA performance analysis, compiler and GC optimisation, and graph-database co-design. Concluded with the award of an MPhil after a server failure caused the irrecoverable loss of three years of experimental data.

Diploma (5-year, MEng-equivalent), Electronic & Computer Engineering 2017
Technical University of Crete – Thesis: memory-system evaluation for disaggregated cloud servers.

SELECTED PUBLICATIONS

- *Performance Analysis for Memory Optimisations of Managed Graph Databases on NUMA Architectures* – 2025.
- *Scaling Up Performance of Managed Applications on NUMA Systems* – ACM SIGPLAN ISMM, 2023.
- *A Multifaceted Memory Analysis of Java Benchmarks* – ACM SIGPLAN MPLR, 2023.
- *dReDBox: A Disaggregated Architectural Perspective for Data Centers* – Springer, 2019.
- *On the Future of Research VMs: A Hardware/Software Perspective* – ACM, 2018.
- *Enabling Virtual Machine Research on ARM-based IoT Devices* – 2018.

AWARDS & SCHOLARSHIPS

Google Compiler Summit 2019 (full sponsorship, Munich) | HiPEAC Advanced Computer Architecture summer-school scholarship 2018 | University of Manchester PhD scholarship (full tuition + stipend) 2017.

REFERENCES

Available upon request.