

Noé Jager

ML Engineer | Luxembourg | noe.p.jager@gmail.com | +33 6 95 90 25 02 | elnukakujo.github.io | French / Luxembourgish

SUMMARY

ML Engineer with **strong software engineering foundations** — from **RDBMS-backed REST APIs** and **full-stack TypeScript/React** apps to cloud infrastructure on **AWS** and **containerised CI/CD pipelines**. Comfortable owning the **full delivery lifecycle: designing, shipping, and maintaining production systems**. Equally at **home fine-tuning LLMs, building on-premise AI agents**, and **wiring up DevOps automation**. Winner of **HackOS4** (1st place, McGill University). Eager to bring sovereign, **open-source AI solutions to European SMEs**.

SKILLS

- **Languages:** Python, TypeScript, SQL
 - **Cloud & DevOps:** AWS (VPC, EC2, Aurora RDS, S3, CloudFront, Amplify), Docker, Kubernetes, Git, MLflow, Optuna
 - **Backend & Data Engineering:** Flask, SQLAlchemy, PostgreSQL, SQLite, REST API design, ETL pipelines
 - **Frontend:** React, Next.js, TypeScript, d3.js, Plotly
 - **AI / ML:** PyTorch, HuggingFace, LangChain, RAG, LLM Agents, Transformers, Contrastive Learning, open-source LLMs (LLaMA, Mistral)
 - **Databases:** PostgreSQL, SQLite, Neo4j — relational design, migrations, Aurora RDS
-

EXPERIENCE

Software & Cloud Engineer (CDD) — TPQuants

May 2026 – Aug 2026

- Led the **full rewrite** of a Replit prototype into a **production-grade React / TypeScript + PostgreSQL client-server application** for **automated Transfer Pricing Audit** report generation.
- **Designed and deployed** an **AWS infrastructure from scratch: VPC, EC2 instances, Aurora RDS cluster, and S3 buckets** — including network segmentation, IAM policies, and secure connectivity.
- Drove ongoing **migration, maintenance, and monitoring** of cloud resources, establishing **repeatable deployment patterns** for future engineering team.

Tools & methods: React, TypeScript, PostgreSQL, AWS (VPC, EC2, Aurora RDS, S3), Node.js

ML/Data Intern — AppOX

May 2025 - February 2026

- Built a **benchmarking framework** comparing **distance-based** (Gower/Anonymeter) vs. **learned-representation attacks** to assess **sensitive attribute inference risk** in synthetic datasets.
- Trained **contrastive learning** and **VAE models** to produce structured tabular **embeddings**, improving attack accuracy on high-dimensional feature spaces where distance metrics fail.
- **Structured** all experiments with **MLflow tracking** and **Optuna hyperparameter search** for fully reproducible, production-ready evaluation pipelines.

Tools & methods: Python, PyTorch, MLflow, Optuna, Contrastive Learning, VAE, GAN, Synthetic Data, Data Privacy

PROJECTS

Fluence — Full-Stack AI Application

May 2025 – present

- Engineered a modular Flask backend using an application factory pattern with blueprints, SQLAlchemy models across 8+ entity types, and Pydantic schemas — exposing a fully documented REST API via Swagger/Flasgger.
- Implemented APScheduler background jobs for automated backups, on-device TTS generation (Qwen TTS), and AI-assisted sentence generation (HuggingFace Transformers) — no external API calls.
- Delivered a Next.js 15 / React 19 / TypeScript frontend with 7 exercise modes (translate, fill-in-the-blank, matching, organize, essay, true/false, answer) and Markdown rendering for grammar notes.

Tools & methods: Python, Flask, SQLAlchemy, Pydantic, PostgreSQL, HuggingFace, Next.js, React, TypeScript, Tailwind CSS

Home Lab — Self-Hosted Infrastructure on Recycled Hardware

May 2026 – present

- Built a **two-node home server cluster** from **repurposed gaming laptops**, running **Proxmox VE** for bare-metal hypervisor management and hardware resource isolation.
- Deployed a **Kubernetes cluster** across both nodes, orchestrating **Dockerised workloads** including **local LLM inference, media and file serving, and self-hosted productivity services**.
- Hands-on experience with on-premise LLM inference on constrained hardware — directly relevant to sovereign AI deployment patterns.

Tools & methods: Proxmox VE, Kubernetes (K8s), Docker, local LLM inference, self-hosted services

EDUCATION

Master in Computer Science at University of Montréal, with AI specialization

September 2022 - March 2026

Bachelor of Computer Science at the University of Luxembourg, with honors

September 2019 - July 2022

Middle School and High School at Saint Pierre Chanel, with honors

September 2012 - July 2019