Jonathan Oliveira

Brasilia, Brazil • hello@jonathanoliveira.dev
linkedin.com/in/jonathan-jorge-oliveira

Experience

Senior Machine Learning Engineer — Brasilia, Brazil

TTY2000 | TRF1 - Federal Regional Court of the 1st Region (Remote) — *January 2024 – April 2026*

Senior ML-focused systems analyst defining MLOps foundations and owning the lifecycle of NLP models in a large-scale justice environment.

- Reduced processing time by 25% by refactoring and productionizing legacy ML services used in judicial workflows.
- Enabled model adoption across more than 500 internal users by building Django APIs that embedded ML outputs into analyst workflows.
- Cut model refresh cycles from weeks to days by automating ETL, retraining, and release workflows with Airflow and CI/CD.
- Reduced manual classification effort by 85% by delivering an NLP assistant tailored to analyst review workflows.
- De-risked RAG adoption for legal retrieval by validating embeddings, vector search, and AWS integration paths for LLM pipelines.

Stack: Python, Django, Airflow, MLflow, DVC, BentoML, AWS

Machine Learning Engineer — Brasilia, Brazil

AI.Lab (Hybrid) — *March 2021 – May 2024*

Machine learning engineer delivering legal-tech R and D and production ML systems for Brazil's higher courts, connecting legal stakeholders, data workflows, APIs, MLOps, and technical delivery.

- Improved rollout readiness across judiciary programs by leading legal-tech discovery, KPI definition, and delivery planning with institutional stakeholders.
- Accelerated deployment of legal AI systems by shipping supervised, unsupervised, and semi-supervised models behind production APIs with standardized MLflow and DVC governance.
- Surfaced more than 30 precedent and semantic categories by designing topic-modeling and semantic-similarity pipelines for long-form judicial text.
- Reduced drift and annotation waste by implementing active-learning loops driven by user feedback and evolving legal language.
- Improved auditability and incident response by delivering end-to-end data, observability, and lineage workflows for production legal AI systems.

Stack: Python, FastAPI, scikit-learn, PyTorch, MLflow, DVC, Transformers, Ollama, Docker

Software Engineer Intern — Brasilia, Brazil

Arvvo Tecnologia (On-site, part-time) — *January 2019 – March 2020*

First industry experience as a full-stack developer in a consulting environment building and maintaining an internal Node.js and Vue.js intranet backed by MongoDB.

- Accelerated internal feature delivery by contributing across frontend and backend in a Git-based consulting workflow.
- Improved usability across devices by building responsive internal interfaces with Vue.js and Angular.
- Enabled authentication, CRUD flows, and business rules by implementing REST services in Node.js and Express.
- Stabilized application data flows and common-path performance by modeling MongoDB collections, indexes, and aggregation queries.

Stack: Node.js, Express, Vue.js, Angular, MongoDB

Education

University of Brasilia — Brasilia, Brazil

Bachelor of Software Engineering — *2018 – 2023*

Grade: 4.2768 / 5.0

CEMI - Gama Integrated High School — Brasilia, Brazil

Computer Technician — *2015 – 2017*

Technical secondary education focused on computing fundamentals.

Languages

- Portuguese — Native
- English — Working proficiency
- Spanish — Basic

Skills

- **Programming and Data:** Python, Node.js, SQL, Pandas, scikit-learn, TensorFlow, PyTorch, SQLAlchemy
- **Backend:** FastAPI, Flask, Django, REST API design, Nginx, Microservices
- **MLOps and Infrastructure:** MLflow, DVC, BentoML, Airflow, Docker, GitLab CI/CD, Kubernetes, Terraform
- **LLMs and Retrieval:** RAG, Multi-agent systems, Prompt engineering, Fine-tuning, Knowledge graphs, Context caching, LangChain, LangGraph, LlamaIndex, Ollama, Qdrant, Hugging Face Transformers
- **Datastores and Cloud:** PostgreSQL, MySQL, Oracle, SQL Server, SQLite, MongoDB, AWS, AWS, GCP