Luiz Gustavo

Brasília, Brazil | luiz.tapioca@icloud.com |

Summary

Backend/Data engineer with experience building event-driven microservices (**FastAPI**, **Kafka**) and data pipelines (**PySpark**), with emphasis on reliability, testing, and scalable processing.

Education

Universidade do Distrito Federal - Jorge Amaury Maia Nunes — Brasília, Brazil
B.Sc. in Information Systems — *Jun 2023 — Present*

Experience

Software Development Intern — Jun 2025 — Present
Banco do Brasil — *Brasília, Brazil*

- End-to-end implementation of core features for an internal AI-driven KPI design platform used by executives, including requirement alignment with stakeholders.
- Development of the "Adherence Matrix" (internal): a vector-based scoring model over indicator proposals (capabilities, aspirations, strategic drivers) using DataFrame computations (thousands of calculations per run) and dot-product visualization to support decision-making.
- Delivery of the "Hive" view (internal) to model capability hierarchies with time-effective (period-based) factors and approval gates; development of "Pandora" (internal), a **FastAPI** API hub for shared microservices (email + Microsoft Teams notifications).

Backend Engineer (Scholarship) — Sep 2025 — Present
eLattes Platform — *Remote*

- Development of an asynchronous microservice pipeline: **FastAPI** for analysis requests, a coordinator publishing stage messages to **Kafka**, and workers processing researcher ZIP/XML inputs into **Parquet** intermediates.
- Implementation of **PySpark** aggregation stages and a finalizer producing a single JSON artifact per analysis (researchers, publications, advising, patents, inferred connections).
- Establishment of an agent-assisted TDD workflow with deterministic integration/regression tests (versioned fixtures + repeatable assertions) and a reproducible Docker Compose runtime (**Postgres**, **Kafka**, **Spark**).

Backend & MLOps Researcher (Scholarship) — Mar 2025 — Nov 2025
IBICT — *Remote*

- Development of an asynchronous **FastAPI** API backed by **Redis** queues, decoupling ingestion from heavy processing to keep the service responsive under load.
- Implementation of Portuguese text normalization (**Enelvo**) and BERT-based hate-speech classification (**Hugging Face Transformers** / **PyTorch**); inference optimized via batching, reaching 200–400 messages/min on NVIDIA GPU (CUDA).
- Implementation of error queues, retry strategy, and graceful shutdown to prevent message loss; stability validated with **Docker** and **K6** load tests.

Projects

Software Engineering / AI, Neural Network from Scratch in Go

- Development of a neural network from scratch in **Go** (matrix ops, forward pass, backpropagation) and training to classify handwritten digits using the MNIST dataset.

Technical Skills

- **Languages:** Python, Go, Java, JavaScript/TypeScript, Rust, C.
- **Backend/Data:** FastAPI, Kafka, PySpark (Spark), SQLAlchemy, Pandas, PostgreSQL, Redis, Docker.
- **ML:** PyTorch, Transformers (BERT), Enelvo.
- **Systems/Testing:** Linux, Nix Flakes, Windows Server, K6.