

# Felipe Mazzuco Kauer

GenAI Engineer | Data Scientist

📍 Porto Alegre, Brazil (Open to Remote)

✉ fmkauer@gmail.com

🌐 Felipe Mazzuco Kauer

🔗 fmkauer

## PROFESSIONAL SUMMARY

AI Engineer specializing in the design and deployment of production-grade Generative AI solutions for industry leaders like Airbnb and Charter Communications. Expert in leveraging Large Language Models (LLMs), RAG pipelines, and vector databases to build novel AI-powered features that drive user engagement and business value. Full-stack experience building on robust cloud platforms (GCP/AWS, Kubernetes), taking AI concepts from prototype to scalable, low-latency production systems.

## TECHNICAL SKILLS

**Generative AI & LLMs:** LangChain, LlamaIndex, Vector Databases, OpenAI API, MCP, RAG (Retrieval-Augmented Generation), Fine-Tuning

**Cloud & MLOps:** AWS, GCP, Kubernetes, Docker, Terraform, MLflow

**Programming & Data:** Python (Pandas, Scikit-learn, FastAPI), Go, SQL, TypeScript, C, Spark

**Languages:** Portuguese (Native), English (C2 - Professional Fluency), Spanish (B1)

## PROFESSIONAL EXPERIENCE

### Platform AI Engineer

TELUS Digital (via Poatek/WillowTree)

Remote

11/2024 – Present

As a platform engineer within a leading tech consulting firm, I am integrated with globally recognized companies to lead the design and implementation of their mission-critical AI platforms and features.

### Client Engagement: Charter Communications

- Co-designed a proof-of-concept for a Text-to-SQL Generative AI tool to empower non-technical users with rapid, self-service data discovery and analysis.
- Engineered the core RAG pipeline that automatically connects to live databases, then extracts, enriches, and embeds complex schemas into a context-aware knowledge base using **pgvector**.
- Implemented the retrieval and generation logic using LLM, enabling the system to translate natural language questions into accurate, executable SQL queries and deliver answers directly to business users.

### Client Engagement: Government of British Columbia

- Directed the successful conversion of the official B.C. Building Code from a static PDF into a fully structured, machine-readable JSON format, unlocking the data for digital search and compliance applications.
- Deployed a novel hybrid data extraction pipeline, pioneering the use of a Visual Language Model (SmolDocling) to interpret and transcribe complex layouts and diagrams previously intractable to standard parsers.
- Combined VLM analysis with programmatic PDF text and table extraction based on positional data, ensuring a comprehensive and high-fidelity conversion of the entire legal document.

## Junior Developer

*Villela Brasil Bank*

Porto Alegre, Brazil

09/2023 – 10/2024

- Solely responsible for the conception, development, and launch of Vick, an internal Generative AI chatbot for legal analysis.
- Developed the backend application (FastAPI) and NodeJS. Integrated payment and notification systems, contributing to the full product launch.
- Applied machine learning for lead generation and created internal data analysis tools that improved management decision-making and reduced operational costs.

## Data Science and PO Intern

## DataGlass

*Porto Alegre, Brazil*

06/2023 – 09/2023

- Developed and optimized Tableau dashboards presenting complex healthcare data from DataSUS and ANS, improving clarity and usability for stakeholders.

## Education

## Bachelor of Science in Production Engineering

Universidade Federal do Rio Grande do Sul (UFRGS)

Porto Alegre, Brazil

*Expected 2027*

## Associate Degree in Data Science

Universidade Estácio de Sá

Brazil

Completed 2024

## CERTIFICATIONS

- CS50: Introduction to Computer Science - Harvard University (2020)