

João Victor Carvalho

📍 Kansas City, MO — 📞 +1 (816) 217-9408 — ✉ joaocarvoli@hotmail.com
🌐 linkedin.com/in/joaocarvoli — 🐙 github.com/joaocarvoli — 🌐 joaocarvoli.com

Summary — Machine Learning Engineer and technical lead focused on speech AI for low-resource and unwritten languages, with 5 years across AI research and software engineering. At OBT Lab I own the technical side of the project, lead the engineering team and build the core stack myself: self-supervised acoustic units, unit vocoders, LLM-driven voice apps and the shared backend behind five field products.

Skills

Speech & ML Self-supervised speech units (XEUS), unit vocoders (HiFi-GAN), ASR/TTS, PyTorch, LLM agents & evaluation (Claude API), RAG	Backend & Cloud FastAPI, SQLAlchemy, PostgreSQL, Alembic, Docker, GCP (Cloud Run, Secret Manager), AWS, CI/CD
Languages Python, Dart (Flutter), TypeScript (React), Kotlin, Java (Spring Boot), SQL	Practice Software architecture, microservices, test-first development, multi-agent dev workflows, tech leadership

Experience

Youth With A Mission (OBT Lab)

December 2025 – Present

Machine Learning Engineer · Technical Lead

Kansas City, MO

- Lead the project's technical work: plan and review the engineers' work and turn the linguistics lead's method into specifications, tickets and acceptance criteria.
- Built the Tripod Method speech pipeline for unwritten languages, text-free end to end: XEUS-based acoustic tokenizer (100 discrete units), BPE motif tokenizer, pitch-conditioned HiFi-GAN v2 unit vocoder (16 kHz), and Tripod Lab, the experiment-tracking and training dashboard with staged quality gates.
- Designed and own the shared FastAPI/SQLAlchemy/PostgreSQL backend on Cloud Run behind five products (Internalization Room, Facilitator Desk, Sound Necklace, Oral Collector, Annotation Studio): auth, app-scoped access, Alembic migrations, GCP Secret Manager.
- Shipped voice-first Flutter apps: Oral Collector (monolingual audio capture on iOS, Android, web) and the Internalization Room, a screen-text-free room where LLM Guide and Validator roles on Claude, TTS and server-sent events walk a team through a passage.
- Built the LLM evaluation harness (golden-session runner with an LLM judge, per-role cost accounting, budget gates) and React/TypeScript consoles: Tripod Console, Meaning Map UI, and Annotation Studio for XEUS layer-wise studies.
- Run data collection and pilots with speakers of Indigenous languages of Brazil (Terena, Sateré-Mawé); orchestrate multi-agent Claude Code delivery workflows (isolated worktrees, dedicated reviewers, Linear).

Maria Educação

November 2024 – November 2025

Machine Learning Engineer · Tech Lead (Jan–Mar 2025) · Software Engineer

São Paulo, Brazil

- Engineered multi-agent architectures and monitoring pipelines over enterprise content indexing, with human-in-the-loop validation to raise LLM answer accuracy and operational efficiency.
- As Tech Lead, partnered with the CTO to scale ML-driven products, bridging product and engineering squads; ran code reviews and automation to keep systems scalable.
- Optimized personalized content generation for 50+ B2B clients with microservices and non-blocking APIs; automated CI/CD for testing and linting.

Projects & Awards

Insight Data Science Lab (Samsung/CRIA)

August 2021 – November 2023

- Researched Explainable AI (XAI) and built health-insight features for Samsung smartwatches from sensor data and microservices.

Awards & Publications

- **1st Place** – AWS DeepRacer for undergraduates, São Paulo, Brazil; **5th in Brazil** – AWS DeepRacer qualifiers, CSBC national competition.
- **Published Paper** – “Comparação de Desempenho entre Soluções de Interoperabilidade” (CSBC).

Education

USP - Universidade de São Paulo

Visiting Student: NLP and Deep Learning

February 2025 – July 2026

São Carlos, SP

Federal University of Ceará

Bachelor's degree, Software Engineering · GPA 3.6/4.0 (9.0/10)

February 2020 – July 2025

Ceará, Brazil

Languages: Portuguese (native), English (advanced)