

Bogdan Ouatu

Python Performance & ML Systems Consultant

ouatu.ro | bogdan@ouatu.ro | github.com/ouatu-ro

SUMMARY

I help teams make Python and ML code run faster and more efficiently. Typical results include shorter processing times, higher throughput, and smoother ML deployments. I co-founded and scaled a video translation platform to 60 k+ users on GCP/Kubernetes, focusing on runtime performance, async parallelism, and model-serving optimization. My expertise lies in profiling, GPU/CPU efficiency, and designing pipelines that stay responsive under load.

SELECTED RESULTS

- **Scaled Platform to 60k+ Users:** Architected and led the engineering for Verbalate's low-latency cloud infrastructure on GCP/Kubernetes.
- **~\$120M/yr Revenue Line Impact:** Deployed Adore Me's first ML recommender, now driving 15–20% of recommendations.
- **+\$9M/yr Marketing Spend Enabled:** Co-built an influencer marketing platform that scaled marketing efforts with zero increase in customer acquisition cost (CAC).

EXPERIENCE

Co-founder & CTO — Verbalate

Jul 2023 – Present

- Scaled the platform to 60k+ users by leading product and engineering, focusing on a highly available and cost-effective GCP/Kubernetes architecture.
- Delivered production-grade ML systems (diarization, lip-sync, TTS) used in thousands of daily media processing jobs.
- Shipped a collaborative dubbing editor (React) and a retry-safe billing and metering pipeline (Stripe).

ML & Systems Engineer — Independent Consultant

Jan 2019 – Jul 2023

- **Turnaround Project:** Delivered a real-time production speech system on NVIDIA Jetson in under three weeks for a project that had been stalled for six months.
- **Award-Winning Innovation:** Won a special prize from Google for a real-time money transfer PoC at the Central Europe Open Banking competition.
- **Finance Systems:** Built a backtesting and options pricing platform that enabled the client to validate and commit to their core investment strategy.

Machine Learning Engineer & Full-stack Developer — Adore Me

2020 – Aug 2022

- Optimized and deployed the company's first ML recommender, successfully serving a ~\$120M/year business line.
- Built an aspect-based sentiment analysis pipeline and BI app that directly reshaped the product roadmaps of three teams.

- Automated the sourcing of 100,000 marketing prospects in one year by building an integration with Instagram before official API support.

CONSULTING SERVICES

Python Performance & Cost Optimization

System Profiling (py-spy, Scalene); Latency Reduction (p95/p99); Cloud Cost Analysis (GCP/AWS); Database Query Tuning

Systems Architecture & Reliability

Cloud-Native Design (GCP, AWS); Observability (OpenTelemetry, MLops); Incident Forensics

ML Systems Engineering

Production Model Serving (FastAPI, PyTorch); RAG & Recommender Systems; GPU Pipeline Optimization

Python & Rust Integration

Accelerating critical code with Rust (PyO3); 10-100x speed improvements

EDUCATION

- **University of Bucharest** — MSc, Artificial Intelligence 2022
- **Alexandru Ioan Cuza University** — BSc, Computer Science 2019

AWARDS & RECOGNITION

- **1st place at Quant Challenge** — London Stock Exchange Group 2021
- **Winner at Open Banking Hackathon** — Google Cloud 2020