

# Valerii Verezhynskyi

## AI Engineer | 2x Founding Engineer

✉ work@valerii.cc 📍 Relocating - London Sep 2026 🌐 [linkedin.com/in/valerianverezhynskyi](https://www.linkedin.com/in/valerianverezhynskyi)

### COMMERCIAL ENGAGEMENTS

---

#### Refact.ai [↗](#)

*Founding Team Member, Data & Backend Engineer*

Nov 2021 – Feb 2024

Joined as the first engineering hire (Employee #1) to build the initial product alongside ex-OpenAI engineer, Oleg Klimov.

Remote

- Built data ingestion scripts to scrape **80 Million** Code Repositories & co-authored in a creation of a Dataset used for training Code Completion model Refact-1.6B-fim (SOTA 2022, 32% HumanEval pass@1)
- Engineered enterprise-grade LLM Inference Backend supporting OSS and on-prem
- Developed a Telemetry & Statistics for Code Completion Model, User Activity

*Founding Team Member, AI Developer*

Feb 2024 – Jan 2025

- Engineered LLM Client in Rust
- Developed RAG for Code Completion & Chat using AST and Vector Search code retrieval
- Implemented Agent Capabilities for LLM: repository exploration, planning, autonomous changes application (SWE bench-compatible)

Remote

#### Coxit, AI Developer, Consulting Engagements [↗](#)

Feb 2025 – Jun 2025

Lead AI Developer across multiple US/EU startup engagements.

Remote

- Developed Named Entity Recognition Model for Personal Information Identification, WASM-compilable
- Engineered MCP Server with RAG for feeding LLM with Technical Documentation in Construction Domain AND Evaluation using LLM-as-a-Judge; Medium [↗](#)

#### BeeSensible, Interim Head of AI (0-to-1 Build) [↗](#)

Jul 2025 – Dec 2025

Architected and shipped complete production-ready AI/ML Stack in <6 months

Remote

- Deployed the complete air-gapped ML Stack: training, storage, deployment, synthetic data generation
- Developed a streaming-first latency-optimized BERT/LLM Inference backend
- Created a reusable pipeline for Synthetic Dataset Generation for fine-tuning LoRAs/BERTs
- Delivered production-ready AI/ML Backend, ML Ops, Synthetic Data Pipeline, Finetuned BERTs & LoRAs (70 labels, production-grade F1)

#### Nautiloid Protocol, Founder, Independent Consulting [↗](#)

Jan 2026 – Present

- Local Voice-to-Voice Cascade Stack [↗](#) | STT+VAD+LLM+TTS | gRPC, websockets, HTTP Real-time, Chat Completions, Audio, Transcriptions APIs implementation with Parakeet V3, gpt-oss-20b, kokoro TTS that produce **sub 400ms TTFA**.
- Zero-dependency MCP Server (Python) [↗](#) | Streamable HTTP, JSON-RPC Custom MCP Server implementation with explicit global state and without opaque decorators unlike reference implementations.

### EDUCATION

---

**B.Sc. Computer-Integrated Optical and Navigation Systems, Kyiv Polytechnic Institute**

2018 – 2022

*Foundation in real-time systems, signal processing, and hardware-software integration*

### STACK

---

**Languages:** Python, Rust, JS

**Tools & Technologies:** uv, fastapi, pydantic, protobuf, websockets, vllm, llama.cpp, onnx, aim, mlflow, otel+clickhouse+grafana

**Environment:** macOS, Proxmox, bare-metal/HPC deployment

### LANGUAGES

---

English C1, Portuguese B2, Russian N, Ukrainian N, French B1, German B1