Mykyta Bulatnikov

AI Engineer & Founder

Prague, Czech Republic · m.bulatnikov@icloud.com · linkedin.com/in/mykyta-bulatnikov · github.com/razornne · skriptly.io

SUMMARY

I build Skriptly, an AI transcription and dictation product now in private beta, and I am the only person who owns data and AI at a proptech platform. I decide what to build and how it should behave, use AI tooling to ship it fast, and measure the result: evaluation sets, cost per hour of audio, and checks on every number a language model states.

Looking for an early-stage startup where I can own data and AI end to end. Always glad to talk with people building applied AI and with investors who follow the space.

EXPERIENCE

Founder & Engineer — Skriptly (skriptly.io) May 2026 – present

AI transcription and dictation. Web app, Windows app, subscription plans. Solo: product, engineering, pricing, billing, support. In private beta with a limited group of testers.

- Built it from idea to private beta: speech-to-text in 60+ languages with speaker recognition, summaries and action items, a hold-to-talk dictation app for Windows, and Free, Pro and Team plans with Stripe billing.
- Rebuilt the engine twice. A hosted API first, then an open-source pipeline (faster-whisper, pyannote, an LLM correction pass) on serverless GPUs, then Soniox after a five-language benchmark: 2.4% against 3.7% word error rate, no dropped words against 1.3%, and about four times cheaper per audio hour.
- Works from unit economics: an hour of call costs about $0.30 to process end to end. Switched off two features after measuring that they cost more than they earned (GPU-generated titles, a live transcript that doubled the cost of an hour).
- Dictation built for latency: a Rust app that streams audio straight to the speech provider with a short-lived key, pastes LLM-cleaned text in about a second, and learns each user's vocabulary from their own edits.
- Handles 3–4 hour recordings with silence-aware chunking, parallel GPU workers and speaker re-identification across chunks. Quality is measured, not felt: an evaluation set tracks word error rate and deleted words.

Data & AI Engineer — proptech SaaS company (B2B, property management) Jun 2025 – present

Prague. Sole owner of data and AI, reporting to the CEO. Part-time intern Jun–Sep 2025, full-time since Oct 2025.

- Built the analytics platform from zero: PostgreSQL into BigQuery, a dbt warehouse of 100+ models, and a metric registry of 150+ metrics defined once in git, so a dashboard card without a definition fails the build.
- Replaced Looker Studio with four custom Next.js dashboards served from static exports of the dbt marts. Cut reporting latency from about a day to 15 minutes.
- Built ML prototypes for building risk, churn and payment behaviour, validated out of time with point-in-time features.
- Built an MCP server over the metric layer and an LLM analytics agent that answers through it, with every number it states checked against the query it came from.
- Built a proactive analyst agent that works as a member of the team: it watches all metrics around the clock and flags anomalies without being asked.
- Audited production data instead of trusting dashboards: found a collection-rate metric that was off by several times, exclusion logic duplicated across views, and a device ID being counted as a person.
- Built monitoring and alerting on Cloud Functions and Slack, and cut its alert volume by about two thirds after auditing a week of its own output.

OPEN-SOURCE PROJECTS

lease-validator — an agentic RAG engine that checks lease agreements against YAML rules and returns no verdict it cannot quote the document for. Violation F1 0.976 against a 0.692 keyword baseline, 100% violation recall, on 180 labelled pairs. github.com/razornne/lease-validator

metrics-mcp — a semantic layer an LLM agent has to go through: metrics defined once in YAML, SQL built from dbt, and every number the agent states verified against its source query. The guard is tested against deliberately lying stubs. github.com/razornne/metrics-mcp

SKILLS

Applied AI: Anthropic, Gemini and speech-to-text APIs, MCP servers, RAG and agent loops, evaluation sets, faster-whisper, pyannote, embeddings and pgvector, self-hosted open-weight models

Engineering: Python, TypeScript, Next.js, Rust (Tauri), Modal, GCP (Cloud Functions, Cloud Scheduler), GitHub Actions, Stripe, Git

Data: SQL, BigQuery, PostgreSQL, dbt, dimensional modelling, metric governance, product analytics, anomaly detection

EDUCATION

Vysoká škola finanční a správní, Prague — BSc Economics and Management 2025 – 2028 (expected)

LANGUAGES AND WORK AUTHORIZATION

Ukrainian and Russian (native), English (C1), Czech (A2). Based in Prague. Available as an employee, or as a contractor (Czech sole-trader registration takes about a week). Open to remote, hybrid and Prague-based roles.