

TANG MINH VU

AI Engineer · Agentic AI Systems

📞 0974 944 902 ✉️ tangminhvu2212@gmail.com 📍 Ninh Binh, Việt Nam

🐙 github.com/tang-vu 🔗 linkedin.com/in/tang-vu



SUMMARY

AI Engineer with **3+ years shipping AI products from 0 to production**. Author of **ContribAI (244 ★)**, an autonomous AI agent in Rust/Tokio with multi-LLM routing (Gemini / OpenAI / Claude), and **reasoning-receipt**, a 5-agent reasoning system that received a **Standouts award at the Agora Agents Hackathon 2026**. Focus on building, deploying, and cost-optimizing scalable **agentic AI systems** end-to-end, from backend and API to production integration on cloud platforms, and operating them reliably as sole owner. Comfortable working with both technical and non-technical stakeholders; independent, fast learner.

Highlights: 244 ★ ContribAI · 300 ★ on GitHub · 🏆 Standouts, Agora 2026 · 10 OSS PRs merged into 60k+ ★ projects · 2 npm packages · 5+ AI-agent hackathons

CORE SKILLS

Agentic AI & LLM Autonomous agent architecture, multi-agent ensembles (Supervisor / Critic), multi-LLM routing (Gemini / OpenAI / Claude), RAG, LLM fine-tuning, tool / function calling, prompt & cost optimization, verifiable reasoning, Vertex AI

Languages Python, TypeScript / JavaScript, Rust, Java, Kotlin, Dart, C/C++

Backend & API FastAPI, Node.js / Express, Tokio, REST, WebSocket, Cloudflare Workers

Frontend React, Next.js, Astro, Svelte; Tauri / Electron (desktop)

Cloud & DevOps Docker / Compose, Linux (Ubuntu / Debian), Nginx reverse proxy, DNS / SSL, Cloudflare (Workers / D1 / R2 / Pages), Vercel, CI/CD, monitoring & alerting, backup & recovery

Data & MLOps scikit-learn, time-series ML, model deployment, performance / scalability / reliability tuning

Databases PostgreSQL, MongoDB, Redis, SQLite / Cloudflare D1

EXPERIENCE

Independent AI Engineer & Founder

2023 - Present

- Design, build, and deploy autonomous AI agents and full-stack AI products; own the full lifecycle from problem framing to production, monitoring, and cost optimization.
- Built a reusable multi-LLM integration layer (OpenAI / Claude / Gemini) powering content, automation, and agent products across the portfolio.
- Sole system administrator for **10+ production systems 24/7** (Docker / Compose, Linux, Nginx, DNS, SSL/TLS, DB backup, monitoring & alerting) at **~100% uptime**.

Founder & Full-Stack Developer - VibePressNet

2025 - Present

- Multilingual SaaS blog & SEO platform: system architecture, Node.js + Python backend, PostgreSQL, REST API, React / TypeScript frontend.
- Built the SEO engine, auto-crawler, and an LLM-based content pipeline; deployed via Docker on Linux.

Founder & Full-Stack Developer - FastBuff (SaaS Marketing)

2025

- SaaS digital-marketing platform: automated order processing with multi-platform + payment API integration (Node.js + React), serving business customers with stable revenue.

Backend Developer - XGame Studio

2024

- Owned the backend of a game project: REST API design and server-side performance optimization, collaborating cross-team with frontend, game client, and QA.

SELECTED AI PROJECTS

ContribAI

244 ★ · Rust, Tokio, GitHub API · v6.8.0

Autonomous AI agent that discovers repositories, analyzes code, and submits pull requests through a 5-stage pipeline: Discovery → Analysis (13-language AST, 27 security / quality / performance skills) → LLM generation + self-review + scoring → automated PR + CI monitoring → auto-fix of review feedback. Rewrote from Python to Rust/Tokio; multi-LLM routing; cross-platform CLI. **602 tests · 44+ PRs submitted, 10 merged · 21+ repositories · 13 languages**.

reasoning-receipt 🏆 Standouts - Agora Agents Hackathon 2026

Python / FastAPI, Next.js, Vertex AI, multi-agent

A verifiable multi-agent reasoning system. **Ensemble of 5 LLM agents** (parallel analysts in isolated context → Bayesian Supervisor with Brier-weighted priors → auditing Critic gate) producing Merkle-rooted, byte-verifiable reasoning traces. Cost-optimized at **~\$0.0007 per verified result; 4,558 results generated and audited**.

Open-Self

7 ★ · Node.js, JavaScript, multi-LLM · MIT

Open-source, self-hosted AI “personality clone” that learns from a user’s chat history to mirror their voice. BYOK across Claude / GPT / DeepSeek / Ollama, runs 100% locally (privacy-first), and integrates with Discord / Telegram / WhatsApp.

AI Novel Studio

Astro 5, React 19, Cloudflare D1 / R2, Drizzle, Better Auth

Edge-first multilingual (VN / EN) AI reading platform: SSR + React islands, LLM-assisted content, authentication, bookmarks, four themes, full-text search, and SEO (sitemap / JSON-LD). Deployed on Cloudflare edge (Workers / D1 / R2).

Conzit - AI Content Generator

TypeScript, Node.js, React, OpenAI API

End-to-end content engine: crawl → LLM rewrite → auto-publish, with category, SEO, and scheduling management.

TTS SaaS Platform

Next.js, TypeScript, OpenAI TTS, Vercel

Text-to-speech AI service with credit / token management and PayPal payments; deployed on Vercel.

OPEN SOURCE CONTRIBUTIONS

Upstream repository	Stars	Merged
Worldmonitor (global intelligence dashboard)	61k ★	3
Maigret (OSINT username search, 3000+ sites)	35k ★	3
AI-Research-SKILLs (Orchestra Research)	10k ★	1
s-tui (terminal system monitor)	5k ★	1
HolmesGPT (AIOps investigation)	2k ★	1

Published packages: @tangvu/airules (sync AI coding rules across 15+ tools) · social-posting-skills (6 ★, AI-agent skills for 12 social platforms).

OTHER SELECTED PROJECTS

- **OCR-PDF** (React / TS, Tesseract.js) - Converts scanned PDFs to searchable text 100% in the browser; no upload.
- **Memory Card Toolkit** (11 ★, Electron) - Desktop disk manager: format / partition / extend / shrink for USB / SD / HDD.
- **DevBoot** (Rust, Tauri, React) - Native desktop startup manager that auto-runs projects on boot.
- **Bot Automation Ecosystem** (Python, JS) - Production Telegram / social bots: forwarding, group management, trend tracking, and auto-posting.
- **X Algorithm Playbook** (30 ★) - Reverse-engineering of the X / Twitter recommendation algorithm with optimization guidance.

RESEARCH

- **LLM4SE Bibliometric Mapping Study (2020-2026)** - Replication data and analysis scripts for a bibliometric mapping study of Large Language Models for Software Engineering.
- **AI-Agent Security Benchmark (2026)** - Benchmark evaluating the security of autonomous AI agents, with accompanying strategy analysis.

AWARDS & RECOGNITION

- 🏆 **Standouts - Agora Agents Hackathon 2026** (Canteen × Circle × Arc; \$50k pool; top 10-12 teams) for reasoning-receipt.
- **300 ★** total on GitHub across 86 public repositories · **5+ AI-agent hackathons** in 2026.

EDUCATION & LANGUAGES

B.Sc. Computer Science, Hanoi University of Science and Technology (HUST) School of ICT

Graduated

Coursework projects: web reading platform, a Training-Management System, and 10+ Android / Kotlin apps.

Languages: Vietnamese (native) · English - professional working proficiency (international hackathons, technical documentation, client-facing work).