

Maria Angelika Agutaya.

AI Engineer · agents, evaluation, and the interface

mariaangelika.com · github.com/mariaangelikabuilds
linkedin.com/in/mariaangelikaagutaya
angelika.f.agutaya@gmail.com
Metro Manila · remote · UTC+8, wide-offset friendly

I build agents on my own time, put them into production, and design the screens they run from. **28 builds, 26 running live**, each documented with what it runs unattended, where it stops for a person, and what it costs per run. I do not assert my systems are safe, I measure them: the same agent runs ungoverned as a control arm, and the numbers make the case.

Builds

warrant

MCP authority gateway · public · measured against a control arm

An MCP server between an agent and the systems it would otherwise reach: the agent proposes one action, warrant executes or refuses it, and seals every decision into a hash chain a third party can verify. Same model, same prompt, same tools, nineteen scenarios: ungoverned, an unauthorized action got through in **19 of 19**, including a 40,000-row delete and a production deploy. Governed, **0 of 19**, and the legitimate work still finished in all nineteen.

19 of 19 to 0 of 19 · precision 96.7% / recall 100% · +18% tokens, the measured cost of the gate

Meridian Ops

agent fleet · live 24/7 on fly.io since August

A fleet (triage, runbook, comms, and a skeptic that can only raise severity) works incidents around the clock across a lab book of invented clients and a production book of my real deployed sites. Nothing sends or executes until a human approves from the web board, monday.com, or Discord, and every approval leaves a real email and a ledger row with the run's exact token cost.

\$0.056 per closed incident · 5 sites on real uptime monitors · 3 approval surfaces, one gate

labor-ph

RAG · public · 1,199 chunks of Philippine labour law

Hybrid retrieval, pgvector plus BM25 with reciprocal rank fusion, over statute text. A system that always answers is not grounded, it is confident, so it refuses below a threshold that is not hand-picked: a committed sweep moved it from 0.4 to a calibrated **0.58**, with the working in the repo.

hybrid pgvector + BM25 · calibrated refusal at 0.58 · sweep committed alongside the code

Cywareness Simulation Pipeline

content engineering system · 2025 to 2026

AI design engineer at a security-training company: built the system behind roughly **300 phishing-simulation assets** across four difficulty tiers and 8+ markets, each bound to a written ruleset on which red flags belong in which tier. Claude Code skills hold the rules, an orchestrator runs eight steps in order, twelve sub-agents do the writing that needs judgment, and Python packages every output identically.

~300 assets · 4 difficulty tiers · 8+ markets · manual body of work turned production line

The other 24, each with a demo and its receipts: mariaangelika.com

Experience

AI Design Engineer · Cywareness.io

Nov 2025 to May 2026

UI/UX Lead · Codebility

Sep to Dec 2025

UX/UI Designer · freelance

Sep 2023 to Nov 2025

Founder · Add+ Coffee

Dec 2021 to present

Technical

Agentic AI	agents on the Anthropic Messages API and Azure OpenAI, hand-written tool loops, MCP as client and server, LangGraph with real interrupts, structured outputs, prompt engineering, Claude Code as daily driver
Retrieval	retrieval-augmented generation: hybrid pgvector + BM25 with reciprocal rank fusion, Azure OpenAI embeddings, calibrated refusal thresholds, per-claim citations
Evaluation	control-arm testing, golden sets, mutation testing of quality gates, cost and latency per model call, traced in Langfuse
Build	Python, TypeScript, JavaScript, React, Next.js, FastAPI, n8n, Postgres, Supabase
Cloud & engineering	Microsoft Azure (Functions, AI Foundry, Azure OpenAI), REST APIs, Git, Docker, CI/CD on GitHub Actions, fly.io, Vercel, Cloudflare Pages, webhooks
Design	Figma, design systems, UI/UX; four shipped case studies (ARI, Ambani, shoshanamd, LOIS)

Education & Certifications

BS Electronics & Communications Engineering

University of Santo Tomas
2014 to 2020

RAG & Agentic AI Professional Certificate

IBM · Apr 2026 · credential I2NFHOROH13C
Google AI Professional Certificate · Feb 2026
UX/UI: UpVersion 2025 · Frost Design 2024 · CIIT 2023 · Backbone Creatives 2023