

Wen Yang Ho

Engineering Leader · AI-native teams and agent systems

North York, ON (Toronto) · Canadian · howenyang@gmail.com · linkedin.com/in/howenyang · github.com/peekolo · wenyangho.com

SUMMARY

VP of Engineering with 7 years at Mortgage Automator, a lending SaaS for private lenders in four countries. I led engineering through technical due diligence for the 2024 majority recapitalization by BVP Forge (Bessemer Venture Partners) and through SOC 2, while staying hands-on in the codebase. Since late 2025 I have built production AI systems with fleets of coding agents: multi-agent orchestration, MCP servers, LLM-judge evals and agent-safe deployment. I'm looking to lead AI-native engineering teams built on agent workflows, evals and guardrails.

EXPERIENCE

Mortgage Automator

Toronto, ON

Loan origination & servicing SaaS for private lenders in Canada, US, Australia and New Zealand

VP of Engineering

Feb 2024 – Present

Senior Software Engineer / Tech Lead

Sep 2019 – Feb 2024

Engineering team of up to ~20 (7 today, after restructuring that followed the recapitalization and a CEO change); platform serving 375+ lenders with \$20B+ funded as of 2024.

- Led engineering through technical due diligence for the May 2024 majority recapitalization of the company by BVP Forge, the private-equity arm of Bessemer Venture Partners.
- Led engineering through SOC 2 certification: controls, evidence collection and changes to the development process.
- Rebuilt the core engine of the premium Fund module, which makes it at least 5x faster.
- Designed a third-party integration framework with built-in logging and usage billing. Four more high-traffic integrations were later added on it without changing the framework.
- Brought Docker into local development, cutting new-developer setup from 1–2 days to under 30 minutes. Built a preview platform that runs several feature branches side by side for QA and demos.
- Built internal platform tools that gave leadership visibility into more than 90% of the engineering team's resource use and delivery speed.
- Redesigned engineering hiring with gamified technical screens that cut process bloat by 80%. Formalized onboarding and the training curriculum with the CTO.
- Introduced a double code-review system and review guidelines that spread knowledge and kept quality steady while the team grew to ~20 developers. Mentored engineers from junior to senior.

SELECTED AI BUILDS (INDEPENDENT, LATE 2025–PRESENT)

- **Motherbrain** (open-source release planned): Control plane for every coding-agent session across ~45 repositories. Uses Claude Code lifecycle hooks (fail-open, spooled when the server is down, no LLM calls on the hot path), an 8-tool MCP server, and a full-text knowledge map with secret redaction. Retrieval quality is gated by an eval (recall@5 0.94, MRR 0.78); a full re-index costs about \$0.11.
- **vincevangel** (launching · vincevangel.com): Image, video and music generation platform for AI agents, available as a library, CLI, MCP server and HTTP SaaS over 8+ model families. Includes spend caps with reservations, a content-addressed cache, embedded provenance, per-model licence checks and an LLM image judge. More than 1,100 tests. Paddle billing and WorkOS auth.
- **LLM-judge calibration** (research): Blind pairwise LLM-judge panels across Claude, Gemini and Qwen, run in both orders with swap-consistency checks. On five real Shorts with known view outcomes the pooled panel reached 73% pair accuracy (Kendall τ +0.80). Documented why that is only a smoke test: five items, and topic familiarity confounds the outcome.
- **Multi-agent orchestration** (Waverunner public on GitHub): Built Waverunner (public), which runs Claude and Codex task waves in parallel in isolated git worktrees, and CompanyMaestro, a team of role agents with DAG-based engineering dispatch and gates that return structured verdicts. Used both to build and ship 5+ products.

EARLIER EXPERIENCE

Freelance Software Engineer · Singapore & Toronto · clients incl. Genentech (Roche) Jul 2007 – 2021

Worked directly with Genentech engineers and scientists for four years, turning requirements into internal apps on Google Kubernetes Engine, App Engine and AWS: equipment scheduling, headcount scenario modeling and email-mined delivery tracking. A scheduling app was adopted at Roche headquarters in Germany.

Co-founder & CEO, Lead Developer · PieceOfCake.sg (Singapore) 2007 – 2010

Ran a 3-developer web studio building e-commerce, rental subscription and inventory systems for local businesses.

Earlier career: Pianist · National Ballet School of Canada; School of the Arts Singapore 2007 – 2019

Principal piano accompanist who led a team of 10 musicians. Built interactive music systems in C++, SuperCollider and Python, including Kinect motion tracking and NLP-driven sound.

SKILLS

AI & agents: Claude Code (hooks, skills, subagents), MCP servers, Multi-agent orchestration, LLM evals & LLM-as-judge, Human-in-the-loop pipelines, Anthropic, OpenAI-compatible & DeepInfra APIs, LLM cost controls

Languages: Python, TypeScript / Node.js, PHP, C# / .NET, SQL, Bash, Java, Swift

Frameworks & data: FastAPI, Laravel, Next.js / React, Playwright, PostgreSQL, MySQL, SQLite FTS5, Redis

Infra & security: Docker / Compose, CI/CD (GitHub Actions), Cloudflare, Caddy / nginx, AWS, GCP / GKE, SOC 2

Leadership: Technical due diligence, Hiring & onboarding, Code-review culture, Mentorship, Incident response

EDUCATION

B.Eng. (Hons.) Computer Engineering, National University of Singapore · Software engineering & embedded systems

Master of Music, Royal Conservatoire of Scotland · Three scholarships

Languages: English (native; IELTS 8.5), Mandarin (fluent)