

Chinguun Ganbaatar

San Francisco, CA · cgchinguun@gmail.com · +1 415 670 0598
[linkedin.com/in/chinguun-ganbaatar/](https://www.linkedin.com/in/chinguun-ganbaatar/) · github.com/chinguun101

EDUCATION

Minerva University
B.S. Computer Science

Sep 2021 – May 2025
San Francisco, CA

EXPERIENCE

Starch (B2B office supplies marketplace)
Forward Deployed Engineer

Sep 2025 – Present
New York, NY

- Built an LLM-powered extraction pipeline parsing 5,000+ product pages into a normalized, unified schema
- Shipped “any URL to structured item” system using OCR + LLM parsing with automated QA validation
- Deployed monitoring and drift-detection layers to maintain data accuracy across continuously changing sources
- Scoped and shipped data integrations for 34 enterprise customers, iterating on extraction pipelines weekly

Mandarin (Bill Sharing App)
Co-Founder

Nov 2024 – Present
San Francisco, CA

- Architected real-time data sync backend serving 1,000+ users with <150ms latency across all operations
- Built logging, tracing, and guardrails pipeline to detect and surface data inconsistencies in production
- Engineered deterministic settlement + ledger reconciliation ensuring correctness across multiple regions
- Owned full system end-to-end: backend services, deployment pipelines, monitoring, and incident resolution

Stora (E-commerce)
Product Manager

May 2023 – Jan 2024
Mongolia

- Drove a 43% revenue lift by optimizing cross-border bundles, shipping logic, and the onboarding conversion flow
- Instrumented key funnels and retention metrics with event tracking, enabling 15–20% faster iteration cycles

Independent Marketing Agency
Founder

Jun 2018 – Sep 2022
Mongolia

- Ran a growth studio serving startups across beverages, film, and real estate, owning multi-client delivery
- Shipped end-to-end campaigns driving measurable business results, including \$3.7M in real estate sales

PROJECTS

Self-Evolving Agent Hack — *Hairfit*, Finalist (83 teams)

San Francisco, CA

- Built a multimodal LLM pipeline (Gemini Flash) generating 8 hairstyle variants from a single selfie input
- Designed self-evolving prompt optimization with ELO-style ranking driven by real-time user feedback loops
- Implemented an evaluation framework to systematically measure and improve model output quality over time

AWS: MCP Agents Hackathon — *Ease*, Finalist (63 teams)

San Francisco, CA

- Built a LangGraph agent orchestrating Linkup, Apify, and ArXiv APIs for RAG-style document retrieval
- Implemented multi-source data aggregation with GPT-4 for natural-language query understanding and synthesis
- Shipped end-to-end agent pipeline: intent parsing, multi-source retrieval, relevance ranking, and audio output

AWS: MCP and A2A Hackathon — *Napify*, Best Use of Arcade Prize (65 teams)

San Francisco, CA

- Built a multi-tool agent orchestrating web search, Spotify, translation, and TTS APIs into one pipeline
- Designed a feedback-driven ranking system that continuously adapts content based on per-user signals
- Shipped a serverless agent framework handling dynamic tool selection, execution, and output synthesis

SKILLS

Languages: Python, TypeScript, JavaScript, SQL

ML & AI: PyTorch, LLMs (GPT-4, Gemini), LangChain, LangGraph, multimodal models, fine-tuning, RAG

Frameworks: React, Next.js, Node.js, FastAPI

Infrastructure: AWS, GCP, Docker, serverless, data pipelines