

Marlon Ma

New York, NY • +1 347-645-9944 • lm5472@nyu.edu • github.com/marlon-damn • linkedin.com/in/your-handle

EDUCATION

New York University — M.S. Computer Engineering

Sep 2024 – Jun 2026

Northeastern University — B.S. Computer Science · Top 1% of cohort

Sep 2020 – Jun 2024

EXPERIENCE

Founding Engineer · Hawala

Jul 2025 – Present

Stablecoin neobank for emerging markets · New York / Remote · First engineering hire. Idea stage to \$3M pre-seed and a live app on iOS + Android in 10 months.

- Shipped the consumer mobile app from blank repo to App Store and Play Store; own architecture, design system, release engineering, and on-call. 2,008 commits over 10 months — ~38% of total repo.
- Shipped every consumer-facing flow — international onboarding & identity verification, multi-rail money movement (cross-chain send/receive, spending cards, on/off-ramps, investing, swap), referrals & rewards — driving performance and funnel completion end-to-end. Authored ~300 PostgreSQL migrations and 178 edge functions; integrated 4 financial partners.
- Designed and built the internal AI operating system on a single-agent multi-role shared-context runtime (memory, tool use, approvals, audit, evals), grounded in retrieval against a Git-backed llms.txt substrate of ~180 source docs. Used daily by 6 internal teams across eng, ops, PM, growth, marketing, and support; +28% retrieval accuracy vs. Notion baseline.
- Two flagship agents running on top: a Tier-1 support agent grounded in live partner-rail transaction context, serving 6 languages — replaced the English-only human queue, median resolution 9–38s. An engineering debug agent kicked off from Slack pulls evidence from Sentry, Supabase, and Git in parallel and drafts root-cause PRs — time-to-RCA hours → <10 min, ~60% merged without rewrite.
- Set the engineering bar as the team grew — coding standards, design system, folder-structure ESLint rules, Zod runtime contracts mobile ↔ backend, migration-first DB writes, 80%+ edge-function test coverage. Reviewed hundreds of PRs.

Undergraduate Researcher · AI Lab, Northeastern University

Sep 2021 – Jun 2024

Computer vision research · industrial QC and medical imaging

Industrial defect detection · glass surfaces

- Sub-2mm defects were getting lost in YOLOv5's deeper feature maps. Designed YOLOv5-G (high-resolution P2 detection head, RepGFPN multi-scale feature fusion, Kmeans++ anchor clustering); small-object mAP +8.3% over the baseline.
- Industrial defect data is scarce by nature — 741 raw images, expanded to 2,124 via a Mosaic + geometric-transform pipeline (rotation, scaling, perspective shift) that preserved defect-class balance.
- Shipped the inspection system end-to-end (Flask inference backend + Vue UI for QC operators) and deployed it at a local glass manufacturer's QC line: detection precision 60% → 82%, manual inspection labor down 60%, fewer defects slipping through to downstream packaging.

Medical image segmentation · hospital CT data

- Reproduced U-Net for medical segmentation on hospital-provided CT scans; matched the paper's reported Dice on a held-out test set.
- Wrote the DICOM preprocessing pipeline (intensity normalization, patch sampling, patient-level data splits to prevent leakage) and ran augmentation ablations — random affine, intensity jitter, elastic deformation — logging per-class Dice and failure modes for the lab's follow-up experiments.

Software Engineer Intern · NetEase

Dec 2022 – Feb 2023

Applied AI for image and design workflows · Beijing

- Built and shipped an AI image-style-transfer service to 800+ DAU. PyTorch + Flask with a Redis + Celery async queue brought inference latency 3–5s → under 1s; engagement up 60%.
- Integrated Stable Diffusion into the design team's workflow; feedback cycles 3 days → 8 hours.

ENTREPRENEURSHIP

Photography Studio SaaS · Founder

Mar 2022 – Jun 2024

Ran a 4-person team on a full-stack booking + portfolio platform for photographers. 60+ paying clients, \$15k+ revenue. Payments, S3 CDN, a CRM that cut studio overhead 50%.

SKILLS

Mobile / Frontend: React Native, Expo, React, Next.js, TypeScript, mobile release engineering (EAS)

Backend / Data: Deno edge functions, Node, PostgreSQL, Supabase, Redis, Kafka, Zod

AI / Agentic Systems: LLM orchestration, tool use, RAG, eval harnesses, agent safety + HITL design, PyTorch

Infra: AWS, Docker, GitHub Actions, Sentry, PostHog