

ANTHONY QIN

antqin27@gmail.com | linkedin.com/in/anthony-qin | github.com/antqin | San Francisco, CA

EDUCATION

Stanford University, Stanford, CA

June 2025

M.S. & B.S., Computer Science (Artificial Intelligence, Computer Systems); Minor in Music

- GPA 4.05/4.0 · Graduated with Distinction · Phi Beta Kappa · Tau Beta Pi · Head TA, CS 194 (Senior Capstone)
- Coursework: Machine Learning, Natural Language Processing, Computer Vision, Compilers, Operating Systems, Matrix Theory

PUBLICATIONS & PATENTS

- C. Ly, E. Peng, K. Liu, **A. Qin**, G. Howe, A. Y. Cheng, A. Cuadra. *Museum in the Classroom: Engaging Students with Augmented Reality Museum Artifacts and Generative AI*. CHI EA '25 (ACM CHI Conference on Human Factors in Computing Systems, Extended Abstracts), 2025. [doi:10.1145/3706599.3719787](https://doi.org/10.1145/3706599.3719787)
- Co-inventor, U.S. provisional patent: *System and Method for Zonal and Combinatorial Sports Prediction* (The Beautiful Game Entertainment).

EXPERIENCE

Co-founder — The Beautiful Game Entertainment, Stanford, CA

Mar 2025 – Sept 2026

- Co-founded and launched a patent-pending, soccer-native sports prediction platform built on a proprietary zonal-odds model; scaled to 16K+ registered users, 59K+ entries, 2.4K prediction-market trades (206K+ contracts), and \$116K+ in customer deposits and purchases in its first year
- Built the product end-to-end (Next.js, React Native/Expo, Python/FastAPI, Aurora PostgreSQL, Redis, Docker on AWS ECS/Fargate): prediction-market and contest engine, dual-currency wallet, payments, and 7-locale localization; live in the US, UK, Canada, Mexico, Brazil, and Denmark for the 2026 World Cup
- Raised angel capital at an \$8M pre-money valuation, led partnership talks with Sportradar, FanDuel, and Underdog, and ran the seed and M&A process

AI Engineering Intern — LinkedIn (Feed AI), Sunnyvale, CA

June 2024 – Sept 2024

- Built causal inference models (S-learner and T-learner uplift neural nets) predicting each member's response to Feed ranking weights, replacing rule-based cohorts with per-member personalized scoring; trained on ~190M members' behavioral features from a live A/B experiment
- Best model lifted long-dwell +1.29% and professional interactions +0.91% over the global baseline in off-policy replay evaluation (p < 1e-70); designed the Qini-curve and OPE evaluation framework for multi-treatment uplift

Research Assistant & Software Engineer — Stanford Institute for Human-Centered AI (HAI)

Apr 2023 – Jan 2024

- Built "Museum in the Classroom," an applied-AI iOS app combining generative AI and AR for classroom learning; shipped to 100+ TestFlight users; ran a middle-school user study (advised by Prof. James Landay), published at CHI EA '25

Software Engineer Intern — Ginger Labs (Notability), San Francisco, CA

June 2023 – Sept 2023

- Shipped redesign features to production every two weeks for 40M+ users; owned the empty-note experience and single-page view, and pitched and built a new two-page view from scratch

Software Engineer Intern — Meta (Messenger), Menlo Park, CA

June 2022 – Sept 2022

- Built Messenger's first fully interactive inbox feature end-to-end (frontend + backend) as a working prototype that became the reference model for subsequent inbox projects

PROJECTS

- **Sage: Voice AI Therapist** (React Native, Swift, TypeScript, GPT-4o Realtime, Llama 3 70B, Whisper, Deepgram) — built an iOS AI therapist with text and voice modes and memory across sessions; wrote a custom speech-to-text → LLM → text-to-speech pipeline before OpenAI's Realtime API existed, then migrated to Realtime within 4 weeks of its launch; ran a 40+ participant user study and built an LLM-as-a-judge eval based on the HAQ-II therapeutic alliance scale
- **Generalized Conditioning of Diffusion Models for Indoor Scenes** (PyTorch, Stable Diffusion) — fine-tuned Stable Diffusion with LoRA/DreamBooth and ControlNet for fine-grained control of interior scene generation; 10% improvement over baseline
- **BERT Multitask Fine-Tuning with Mean Teacher** (PyTorch) — fine-tuned BERT jointly on sentiment (SST), paraphrase (Quora), and STS with Mean Teacher consistency regularization and balanced round-robin sampling; dev average 0.409 → 0.664 over single-task baseline
- **GyML: Smart Fitness Trainer** (HMR 2.0, PyTorch, scikit-learn) — 99%-accurate exercise-form classifier on 3D human mesh recovery, with natural-language coaching feedback
- **WorldDex** (Swift, PyTorch, GroundingDINO, AWS EC2, OpenAI, ElevenLabs) — "real-life PokéDex" iOS app for collecting memories; Best Blockchain Project & Best Web3 Hack, Cal Hacks 10.0 ([Devpost](#))

SKILLS

Languages: Python, TypeScript/JavaScript, C++, C, Swift, SQL, Java

ML / AI: PyTorch, TensorFlow, scikit-learn, Spark; LLM apps & evals, OpenAI Realtime API, diffusion models, causal/uplift modeling

Systems / Product: React/Next.js, React Native/Expo, FastAPI, Node/Express, PostgreSQL, Redis, Docker, AWS, iOS (Swift, ARKit)