

KHANG NGUYEN

snguyen@usc.edu | (310) 985-5459 | Los Angeles, CA | github.com/Steve6378

EDUCATION

University of Southern California

Aug 2025 – Dec 2026 (expected)

M.S., *Mathematical Data Science*

GPA 3.91

Selected coursework: Introduction to Online Optimization (CSCI 659; online learning and bandits), Mathematics of Machine Learning (MATH 447), Deep Learning and Its Applications (CSCI 566), Mathematical Statistics for Data Science (MATH 546), Foundations of Mathematical Data Science (MATH 549)

University of California, Los Angeles

Sep 2021 – Jun 2025

B.S., *Data Theory*

GPA 3.50; Dean's List (Fall 2021, Winter 2022, Winter 2023)

PUBLICATIONS AND PREPRINTS

* denotes equal contribution (alphabetical author order)

1. K. Nguyen, V. Iverson, S. Wijetunga, W. Chang, G. Wang. "Tight rates of approximation of mixed Nash equilibria by entropy regularization in continuous games." Proceedings of the 42nd Conference on Uncertainty in Artificial Intelligence (UAI), PMLR 337:4855–4872, 2026.
2. B. Chen, K. Nguyen, J. Liu, X. Li, Y. Luo, W. Chang. "Gap-Independent Regret for Multi-Agent Combinatorial Semi-Bandits." European Workshop on Reinforcement Learning (EWRL), 2026.
3. W. Chang*, C. S. Goodrich*, V. Iverson*, K. Nguyen*. "A pathological property of nonlocal discrete operators." Proceedings of the American Mathematical Society, Series B, 13 (2026), 20–34.
4. M. S. Lin, K. Nguyen, B. Chen, J. Q. Loi, V. Bansal, W. Chang. "Hybrid Offline-Online Follower Manipulation in General-Sum Stackelberg Games." NeurIPS 2026, under review.
5. A. Kulkarni, K. Nguyen, R. Parada, K. Guo, W. Chang, Y. Dai. "Learning Adversarial Continuous MDPs with Bandit Feedback and Unknown Transitions." NeurIPS 2026, under review.
6. K. Nguyen, K. Guo, W. Chang. "Policy Optimization for Corrupted Markov Decision Processes." AAAI 2027, under review.
7. K. Nguyen, L. Xu, L. Jiang, W. Chang. "Best Arm Identification in Lipschitz Bandits: Uniform and Adaptive Strategies." AAAI 2027, under review.
8. J. Q. Loi, B. Chen, K. Nguyen, W. Chang, N. Weinberger. "Pure Exploration for Curriculum Bandits." AAAI 2027, under review.
9. T. Iyadomi, L. Xu, K. Nguyen, R. Parada, W. Chang, S. Yogamani. "ZoomQ: Adaptive Action Discretization for Continuous-Action RL." AAAI 2027, under review.
10. K. Nguyen, L. Jones, W. Chang. "Minimum-Energy Optimal Gaze Control for K-Eye Systems with Three-Axis Rotation." IFAC Conference on Cyber-Physical & Human Systems (CPHS) 2026, under review.

RESEARCH EXPERIENCE

BruinML, UCLA Mathematics

Jun 2024 – Present

Applied Researcher

Hybrid, Los Angeles, CA

- First author on 4 papers and co-author on 6 in William Chang's group (3 published or accepted, 7 under review; see Publications), with first-author work on mixed Nash equilibria of continuous games, corrupted MDPs, Lipschitz best-arm identification, and optimal gaze control.
- Led the UAI 2026 paper establishing tight rates for approximating mixed Nash equilibria of continuous games via entropy regularization.
- Act as the group's internal reviewer on nearly every paper: adversarially check proofs for gaps, rewrite flawed arguments and exposition, and run numerical experiments where the theory admits them.
- Lead author responses across the group's submissions, including two where reviewer scores rose after rebuttal.
- Lead the experiments team across concurrent projects, including a PPO baseline for autonomous driving on the Waymo Open Dataset via GPUDrive.

Institute for Computational Science & AI, Van Lang University

Jun 2023 – Sep 2023

Undergraduate Researcher

Ho Chi Minh City, Vietnam

- Developed LSTM models for VN-Index stock prediction, testing 20+ technical indicators across 10 years of data.
- Applied two decomposition methods (EMD and wavelet transforms) to separate trend and cyclical components before forecasting.

ENGINEERING EXPERIENCE

Coursistant

Dec 2025 – Present

Machine Learning Engineer (contract since Feb 2026; solo prototype Dec 2025 – Jan 2026)

Los Angeles, CA

- Sole AI engineer for a classroom-management LLM agent: 25+ function-calling tools on GPT-4o cover the course lifecycle through chat, with citation-aware RAG and per-course student visibility filtering; now running on the production service ahead of market release.
- Designed the agent's action-gating model: every state-changing action requires an explicit tool call plus an out-of-band human confirmation the model cannot produce; nine write-action rules are codified in the spec, each enforced by a named regression test.

- Built defenses against adversarial student inputs (hidden-text prompt injection in submissions, jailbreak attempts), role-based tool dispatch with defense-in-depth re-checks, and fuzzy name resolution with confirmation prompts to prevent wrong-target actions.
- Maintain a 300+ test backend suite covering tool-call privacy gating, anonymity invariants, and security regressions.

TEACHING AND TRAINING

- Teaching Assistant, DSCI 560 Data Science Professional Practicum (algorithmic trading), University of Southern California, Spring 2026.
- ARENA 2.0 (UCLA AI Safety club, 2023) – worked through the AI safety and mechanistic interpretability curriculum.

PROJECTS

Yorru – AI-powered event-planning platform with RAG

Oct 2025 – Mar 2026

Group project, lead engineer · yorru.net

- Built the full-stack platform with pgvector-backed RAG (OpenAI text-embedding-3-small) and a two-tier query system combining keyword matching with semantic search; the AI passively extracts guest preferences (dietary, budget) from group chat.
- Implemented real-time Socket.io chat across 3 interaction modes: 1-on-1 guest AI assistant, observer-mode group chat, and host dashboard.

Pitch Prediction with LSTM and Conservative Q-Learning

- Built an offline RL pipeline for pitch selection from 8,400+ Statcast pitches (state: count plus last two pitch types; reward: negative change in run expectancy), sweeping CQL conservatism alpha and temporal reweighting with early stopping.
- Found that high-alpha configurations reported apparent policy improvements of up to 0.94 runs per pitch, several times the reward standard deviation of 0.20; cross-checked against training TD error and diagnosed the gain as Q-function miscalibration, since the metric is computed under the learned Q rather than a true value estimate.

Steam Games Insight Engine (SGIE)

- End-to-end NLP pipeline over scraped Steam reviews: KNN-based fake-review filtering, aspect-based sentiment analysis (NMF topic modeling + DistilBERT), and an ensemble of 3 fine-tuned transformers (GPT-2, RoBERTa, DeBERTa) with Jensen–Shannon disagreement analysis.

CUDA Convolutional Neural Network

- CNN implemented from scratch in CUDA C/C++ (convolution, pooling, activation, and fully connected layers with forward and backward passes) with a CPU-vs-GPU benchmarking suite on MNIST.

SKILLS

Programming: Python (PyTorch, TensorFlow, scikit-learn, NumPy, pandas, XGBoost), CUDA C/C++, R, SQL, TypeScript/JavaScript

Systems: FastAPI, React, Next.js, PostgreSQL/pgvector, Node.js, Docker, Railway, Vercel, Cloudflare R2, Git, HPC/SLURM

Research areas: Reinforcement learning, multi-armed bandits and online learning, Markov decision processes, game theory, statistical learning theory, optimal transport