Clara Zeng

213-248-9142 | qiang_zeng@berkeley.edu | Berkeley, CA | **LinkedIn** | **GitHub**

EXPERIENCE

Cmind AI | ***Associate Data Scientist*** **May 2024 – Dec 2025**

- Architected a modular microservices architecture using Python and FastAPI to decouple feature engineering from model inference, increasing system maintainability and reducing deployment downtime by 40%.
- Engineered automated CI/CD pipelines for model retraining and deployment; leveraged MLflow to manage version control and production monitoring, ensuring high availability and 99.9% reproducibility across environments.
- Led development of a scalable LLM-as-a-Service interface using FinBERT (financial domain) and OpenAI APIs for tailored Q&A with explanations; designed Oracle Cloud ingestion for cost-efficient legacy-system compatibility and AWS S3 storage to support a smooth migration to modern client-facing analytics workflows.
- Enhanced the Earnings Per Share Surprise predictive analytics system by implementing automated retraining pipelines, training XGBoost, SVM, & RNN models on financial data to achieve 90% quarterly forecast accuracy.

Bluebono | ***Data Analyst*** **June – Aug 2025**

- Designed and implemented an end-to-end backend pipeline for residential property evaluation, managing the full software lifecycle from initial project scoping to production development.
- Built a real-time data ingestion engine by integrating the Trestle API satisfying the ultra low latency business requirement
- Optimized ETL pipeline to handle and manipulate 11.3M property records with 1,000+ market features.
- Implemented unsupervised clustering algorithm to reduce data dimensionality by 65%, significantly improving system interpretability and processing efficiency.

Guangfa Futures | ***Quantitative Reasercher*** **June - Aug 2023**

- Developed a hybrid trading system in Python combining K-Means regime classification, regression-based forecasting, and deep reinforcement learning to optimize allocation across futures: spanning energy, metals, agriculture, and financials.
- Implemented an end-to-end ML pipeline in PyTorch and Stable-Baselines3 featuring FinBERT-based sentiment analysis, transformer attention modules, walk-forward backtesting, and multi-seed reproducibility for rigorous model validation.

EDUCATION

University of California, Berkeley **May 2026 (expected)**

Masters in Analytics GPA: 4.0/4.0

University of Southern California **May 2025**

B.S. in Computer Science & B.A. in Applied and Computational Mathematics GPA: 3.94/4.0

PROJECTS & COMPETITIONS

Google x UC Berkeley Analytics Hackathon — Top 3 Finalist **Dec 2025**

- Engineered a diagnostic tool using Python, pandas, and NumPy to analyze a 3x customer churn spike, processing CRM, support tickets and usage logs to identify trends revealing small companies drove 75% of churn volume.
- Performed statistical analysis to validate usage drop patterns. Identified 58% of churned tickets as onboarding gaps vs. technical bugs, delivering data-driven recommendation to executive panel addressing customer success friction.

Tesla Global Sourcing Strategy Case Competition — Semi-Finalist **Sep 2025**

- Built cost-optimization model analyzing $200M+ sourcing decision across US, Mexico, and China under 25% tariff scenarios, incorporating capacity constraints, logistics costs, and currency fluctuations
- Applied Monte Carlo simulation to model tariff escalation probabilities and supply chain disruption risks; quantified $18.7M in operational flexibility value through strategic multi-plant switching options

LLM-Powered Career Advisor (RAG Pipeline) **Aug - Dec 2024**

- Built a RAG pipeline over 1.3M+ O*NET and LinkedIn career entries using vector embeddings and similarity search; integrated GPT and Llama3 APIs to generate personalized career recommendations and skill gap analysis

LLM Prompt Recovery Competition **Feb - Apr 2024**

- Fine-tuned Mixtral-8x7B using QLoRA with 4-bit quantization; implemented a custom Sharpened Cosine Similarity loss function using sentence-T5 embeddings to optimize prompt prediction from text transformation pairs.

TECHNICAL SKILLS

Programming & Tools: Python (pandas, NumPy, SciPy, TensorFlow, PyTorch), SQL, R, C++, Java, Git/GitHub, MLflow, Tableau, Streamlit, n8n, RESTful APIs (OpenAI, Trestle), Web Scraping, ETL/ELT pipelines

Machine Learning: RAG, SVM, Tree Ensembles (DT/RF/XGBoost), kNN, Naive Bayes, A/B Testing, Unsupervised (k-Means, GMM, PCA), Deep Learning (LSTM, Transformers, BERT/FinBERT, GPT, seq2seq), GAT, Q-Learning

Cloud & Databases: AWS (S3, Glue), GCP (Cloud SQL, BigQuery), Snowflake, Azure, Oracle Cloud, MySQL, Redshift, MongoDB