

Qamil Mirza

(+1) 341-388-6936 | qamilmirza@berkeley.edu | <https://www.linkedin.com/in/qamil> | [portfolio](#)

EDUCATION

University Of California, Berkeley

Bachelor of Arts in Data Science & Statistics; GPA: 3.97

Berkeley, CA

August 2023 – May 2027

Relevant Coursework: Foundations of Data Science, Data Structures, Discrete Math & Probability, Multivariable Calculus, Abstract Lin. Algebra, Real Analysis, Computational Chemistry, Data Science for Biology

EXPERIENCE

AI Search Optimization Engineer

Maxis Berhad

June 2025 – Present

Kuala Lumpur, Malaysia

- Developed a fully customizable spin-the-wheel React application, now deployed in 15,200+ Maxis retail stores nationwide to support customer engagement campaigns
- Designed and deployed a serverless AWS-based document storage platform, migrating 200K documents across multiple departments, eliminating third-party dependency and reducing annual infrastructure costs by RM100,000+
- Hardened backend APIs through comprehensive security audits, remediating high-risk vulnerabilities including SQL injection points and outdated cryptographic algorithms (MD5, SHA1)
- Revamped the ETL pipeline for Maxis's AI search agents, boosting document retrieval precision by **8%** and enhancing support performance for over 3,000 employees

Machine Learning Research Assistant

Merck Group

Jan 2025 – Present

Berkeley, CA

- Built a modular, end-to-end model development pipeline by organizing the project on GitHub and scripting automated workflows for data extraction, feature engineering, model training, and evaluation
- Engineered Directed Message Passing Neural Networks (D-MPNNs) to predict drug potency (IC50/EC50), leveraging protein-ligand interaction fingerprints as model features
- Streamlined molecular docking workflow for 18,000+ ligands, optimizing docking parameters to reduce protocol noise by **25%** and significantly improving signal quality in downstream modeling

Fullstack Software Engineer

Supa

May 2024 – Aug 2024

Selangor, Malaysia

- Developed and deployed a PERN stack onboarding application on AWS EC2, streamlining the registration of 100+ data annotators onto Supa's data labeling platform
- Implemented a dynamic skill-based tagging system to assist project managers in efficiently assigning 500+ annotators to projects based on relevant expertise
- Engineered a browser-based image embedding visualization tool supporting over 1 million datapoints, enabling the team to meet a key client requirement and secure a RM10K contract

PROJECTS

ML-Guided Directed Evolution of Activation Domain | Python, PyTorch, Transformers

April 2025

- Trained regression models (Random Forest, Gradient Boosting) on amino acid features and ESM-2 embeddings to predict gene activation strength of intrinsically disordered 40-amino-acid protein sequences
- Engineered feature representations using net charge, hydrophobicity, motif counts, and per-residue disorder to combine domain-specific priors with learned embeddings
- Implemented in silico directed evolution to iteratively mutate low-activity sequences, filtering candidates by predicted activity and mean disorder score ≥ 0.5

Optimizing Equity Derivative Hedging Strategies | Python, Gurobi, Pandas, NumPy, Jupyter

Feb 2025

- Designed a Mixed-Integer Linear Programming (MILP) model to hedge equity derivatives across SPY, QQQ, and IWM, ensuring a minimum +\$10M exposure under all market scenarios
- Minimized daily and cumulative option premium costs while maintaining non-negative exposure throughout a simulated trading horizon
- Awarded **2nd place out of 20+ teams** in the Berkeley IEOR + Wells Fargo BADSS Case Competition for technical rigor and practical feasibility of the solution
- Set up modular codebase with Conda environments, shell scripts, and reproducible result pipelines for collaborative model development and analysis

TECHNICAL SKILLS

Languages: Python, JavaScript/TypeScript, SQL, Java, Go

Frameworks & Tools: React, Node.js, Flask, FastAPI, Docker, Jest, Git

ML & Data: PyTorch, TensorFlow, Scikit-Learn, Gurobi, pandas, NumPy, ESM-2, MetaPredict

Cloud & DevOps: AWS (EC2, S3, Lambda, Step Functions, DynamoDB, RDS, EFS, Amplify, Load Balancer), Supabase, Docker