



BRANDON MINER

San Francisco, CA

brandonm333@outlook.com +1 408-335-5436  

Machine Learning Engineer with 2 years of experience in machine learning, cloud computing, and data systems

Education

University of San Francisco

M.S. in Data Science and Artificial Intelligence

San Francisco, CA

Jun 2026

San Diego State University

B.S. in Statistics with Emphasis in Data Science, Minor in Mathematics

San Diego, CA

Dec 2025

Experience

Machine Learning Engineer

ACLU

San Francisco, CA

Oct 2025 - Present

- Designed and built an end-to-end **ML inference platform** from scratch (**Python, PyTorch, FastAPI, PostgreSQL, Docker**) to automate analysis of unstructured police bodycam footage, projected to eliminate **300+ hours** of manual review per monthly cycle.
- Engineered a **GPU-accelerated ETL and inference pipeline** using **FFmpeg** and **Whisper ASR** for batch transcription of **600+ videos per month**; developed an **LLM-driven auto-labeling** system to generate training data for fine-tuning multi-stage text classification models.
- Architected a **multi-stage NLP classification system** with LLM-assisted labeling, iterative fine-tuning, and structured output parsing to extract actionable audit signals from raw transcripts at scale.
- Deployed the full ML system on a **private-server backend** with a PySide6 desktop GUI, enabling reliable model serving and inference for non-technical end-users without cloud dependency.

Data Scientist, Lead Intern

San Diego County Taxpayers Association

San Diego, CA

Mar 2024 - Aug 2025

- Led a team of **4 analysts** to build data pipelines delivering statistical analysis of civil workforce compensation and homelessness intervention programs.
- Developed an **data standardization pipeline** combining iterative web scraping and text normalization to unify messy public wage records; applied regression modeling to quantify a **19.7% police wage increase** per **1%** staffing reduction.
- Applied **mixed-effects regression models** with multicollinearity reduction (VIF analysis) to regional homelessness expenditure data, surfacing high-impact interventions such as rapid rehousing to guide resource allocation decisions.

Technical Projects

California Grape ETL Pipeline & Analytics Dashboard

- Engineered a **multi-source ETL pipeline** in Python to extract, transform, and consolidate disparate agricultural datasets into a unified schema ready for downstream ML and analytics workloads.
- **Containerized the full-stack application** using **Docker Compose** and deployed to **Google Cloud Run** as a serverless instance, enabling scalable, reproducible model and app serving.
- Built an interactive **Streamlit front-end** dashboard for dynamic visualization and filtering of regional production metrics; demonstrates end-to-end MLOps deployment from pipeline to user-facing interface.
- View Live Dashboard

Technical Skills

Computer Science: Python (PyTorch, tensorflow, Scikit-Learn, Pandas, NumPy, FastAPI, Transformers, Hugging Face), SQL (PostgreSQL), R (tidyverse), Bash, Spark SQL

AI/ML: Deep Learning (NNs, CNNs, Transformers, fine-tuning), LLM Integration & Prompting, natural language processing (NLP)(ASR, Text Classification, Named Entity Recognition), Recommender systems, Regression, Ensemble Methods (Random Forests, XGBoost, AdaBoost), Bayesian Methods, PCA, DBSCAN, Explainable AI (XAI)

MLOps & Infrastructure: Docker, GCP, Cloud Computing, Git, GPU-accelerated inference, GitHub, FFmpeg, OpenAI Whisper, REST APIs (FastAPI), Distributed Computing (Spark), MongoDB, Web Scraping