Hao Liang

148 1st St, Jersey City, NJ
+1 608-770-5554 haoliang428@outlook.com LinkedIn Google Scholar

AI/ML Engineer specializing in production LLM systems, cloud-native data platforms, and MLOps for large-scale environments. Experienced in end-to-end system design spanning modeling, optimization, governance, and deployment.

EDUCATION

Northwestern University | Evanston, IL **Sep 2022 – Dec 2023**
M.S. in Machine Learning and Data Science (formerly Analytics) **GPA**: 3.95/4.00

University of Wisconsin – Madison | Madison, WI **Sep 2018 – Dec 2021**
B.S. in Statistics & B.S. in Economics (Scholarship: Lenell Goodman Ammerman Scholarship in Economics) **GPA**: 4.00/4.00

WORK EXPERIENCE

Roche **Mar 2024 – Present**
Data Engineer Branchburg, NJ
Design and deliver end-to-end AI/ML, data engineering, and optimization systems supporting Roche's R&D and product pipelines

AI & LLM Engineering and Optimization

- Architect production-grade LLM knowledge systems by building RAG pipelines and benchmarking NLP models (LDA, BERT, BART) to automate R&D data synthesis
- Integrate evaluation frameworks (Ragas) to quantify retrieval accuracy and generation quality, ensuring high-fidelity outputs for sensitive scientific workflows
- Build OCR and embedding-based parsing pipelines to convert compliance documents into structured data, automating form-filling and reducing manual effort
- Apply graph-based optimization algorithms (e.g., Dijkstra's) to automate combinatorial oligonucleotide design, translating scientific constraints into scalable decision logic
- Optimize manufacturing logistics and reduce line failures using Optimization to balance inventory and production targets

Data Architecture and Governance

- Engineer Data Vault 2.0 on Snowflake, integrating SAP, MES, and SQL Server sources via Qlik, Talend, and dbt to establish a governed, analytics-ready data foundation
- Build a data governance stack integrating Monte Carlo, Collibra, and Immuta to enable observability, metadata management, and secure access across AI workloads

Cloud Infrastructure and MLOps

- Modernize production infrastructure by migrating 10+ projects to AWS ECS/Fargate using GitLab CI/CD and Infrastructure as Code (IaC)

Analytics Platforms & BI Automation

- Pioneer the first Posit Connect Server deployment to host Streamlit applications, replacing manual workflows with containerized, production-ready analytics environments
- Lead multi-site Tableau rollouts by automating ingestion from Monday.com and Smartsheet into Snowflake, improving operational visibility and reporting consistency

Cisco **Sep – Dec 2023**
Data Science Consultant San Jose, CA
Empowered supply chain teams with an end-to-end analytics platform to detect market shifts and potential disruption risks

- Built ELT pipelines to aggregate semiconductor news and commodity price data via web scraping (Beautiful Soup) and external APIs, creating a unified signal layer for downstream analysis
- Leveraged fine-tuned prompt engineering with LLM APIs to automate news summarization and topic modeling
- Distilled risk signals using NLP (NER, sentiment, lemmatization) in NLTK and spaCy, improving intent and risk models
- Developed Tableau dashboards to track emerging risks and support data-driven decisions for supply chain stakeholders

H-E-B **Jun – Aug 2023**
Data Scientist Intern Austin, TX
Applied ML, AI, and optimization to solve large-scale supply chain challenges, improving efficiency and driving innovation

- Engineered 100M+ record datasets with PySpark SQL, building scalable ETL pipelines and data-quality checks for forecasting and optimization

- Developed demand forecasting models using time-series methods (e.g., SARIMAX) and deep learning architectures (e.g., GRU, LSTM), incorporating feature-driven signals to capture demand dynamics
- Designed an inventory simulation framework to evaluate policy performance under demand uncertainty and operational constraints; applied Hyperopt to tune forecasting and inventory parameters
- Optimized inventory strategies with Google OR-Tools and heuristics, producing actionable supply-chain recommendations

Sabre **Sep 2022 – Jun 2023**
Data Science Consultant Dallas, TX
Developed and deployed statistical and ML solutions to support dynamic pricing and market strategy in the airline industry

- Built airline price elasticity models using Poisson Regression, XGBoost, Random Forest, and LSTM, improving demand sensitivity estimation across markets and fare granularities to inform revenue and market entry strategies
- Performed market segmentation using K-means, Agglomerative Clustering, and DBSCAN, enabling personalized prediction and differentiated pricing strategies across customer and route segments
- Architected automated ML pipelines using YAML, Docker, and Python-based UIs for model training and real-time inference, accelerating experimentation cycles and enabling production-grade predictions

Refinitiv (London Stock Exchange Group) **Apr – Jul 2021**
Data Scientist Intern Beijing, China
Delivered data-driven insights on user behavior and product engagement to support analytics and platform strategy for Eikon

- Designed and ran A/B experiments on Eikon webpage layouts, driving a 16% increase in Daily Active Users (DAU)
- Built dashboards for bond credit rating monitoring using Python (Bokeh), reaching 5,000+ views in the first month
- Analyzed subscriber behavior and retention with R to inform product and marketing strategy
- Authored and delivered Python tutorials on Eikon API usage; trained 1,000+ users in programmatic data access and analysis

Minsheng Securities **Aug – Oct 2020**
Data Analyst Intern Beijing, China
Built data pipelines, analytical models, and forecasting tools to inform trading and client-profitability analysis

- Built and maintained ETL pipelines for Private Placement transaction data using SQL, extracting key trading and client metrics; performed feature engineering and exploratory data analysis in Python to prepare modeling-ready datasets
- Developed time-series forecasting models (e.g., ARIMA) in R (forecast) to predict Private Placement trading volume, achieving a MAPE of approximately 15% and supporting short-term market activity planning
- Analyzed client profitability and portfolio relationships using factor-exposure analysis to identify concentration risk
- Translated analytical findings into structured insights for stakeholders, informing trading and investment decisions

Dongxing Securities **Jun – Aug 2020**
Risk Data Analyst Intern Beijing, China
Applied quantitative modeling and analytics to assess portfolio risk, support equity research, and inform client investment insights

- Quantified portfolio risk by computing and comparing VaR and ES, enabling robust cross-method risk assessment
- Developed risk analysis workflows to evaluate portfolio sensitivity under different market scenarios, synthesizing results into structured risk assessment reports for internal teams and external clients
- Performed equity valuation and factor analysis, implementing Discounted Cash Flow (DCF) models and estimating equity beta via linear regression in R to support fundamental research and risk-adjusted return analysis
- Built visualizations to analyze stock-price dynamics and key indicators (RSI, MACD) for equity-research insight

PROJECTS

REFOP (Renewable Energy Forecasting & Optimization Platform) **Mar 2025 – Present**

- Spearhead the technical development of a SaaS platform helping renewable energy companies navigate regulatory frameworks (RPS, Carbon Pricing) and optimize operations
- Orchestrate end-to-end MVP delivery, from business planning to technical architecture across seven feature modules
- Architect a containerized microservices platform on AWS ECS with real-time ingestion pipelines by Kafka and Spark
- Engineer backend with Python and Node.js while building interactive React.js and D3.js dashboards for scenario modeling
- Direct development of the core AI engine, implementing models for time-series forecasting, policy analysis with BERT, and optimization of carbon and REC trading

Deep Learning Framework for Medical Image Enhancement and Automated COVID-19 Diagnosis **Jan – Jun 2023**

- Preprocessed X-ray images by resizing dimensions, balancing class distributions, and normalizing pixel intensities
- Enhanced image quality using denoising autoencoders and OpenCV to reduce noise and emphasize subtle features

- Benchmarked a custom CNN against transfer learning models (ResNet, VGG, EfficientNet, DenseNet), achieving 96% accuracy and improving diagnostic speed and precision

Machine Learning for Automated Credit Risk and Customer Segmentation Modeling **Jan – Jun 2023**

- Built client segmentation models using K-means, hierarchical clustering, and K-prototypes to enable tailored loan strategies
- Benchmarked default risk prediction models (Logistic Regression, Random Forest, XGBoost) to mitigate credit risk
- Deployed models as containerized services on AWS (ECR, ECS) with a Streamlit app for real-time risk assessment

RESEARCH EXPERIENCE

UW-Madison Business School Madison, WI

Research Assistant under Prof. Michael Hernke **Sep 2019 – Dec 2021**

- Constructed a break-even analytical framework to investigate the Jevons' Paradox threshold, quantifying the energy-efficiency gains required to achieve net environmental savings in alternative energy vehicle markets
- Designed and executed a multi-source data integration architecture, harmonizing high-dimensional datasets from state and federal agencies to establish normalized energy features for longitudinal econometric analysis
- Utilized input-output analysis to model ecological footprints, establishing a quantitative framework for assessing cross-sectoral environmental impacts of emerging transportation technologies
- Authored peer-reviewed research on the rebound effect in EV adoption, presented at international conferences

Research Assistant under Prof. Margie Rosenberg **Sep 2019 – May 2020**

- Implemented a Difference-in-Differences (DiD) model to quantify the causal effect of economic downturns on population-level health indicators across demographic cohorts
- Engineered a harmonized longitudinal database from 20+ public health surveys (including BRFSS and NHANES) using survey-weighted adjustments and cell-based weighting
- Authored the research proposal defining the technical methodology and led the initial data-engineering architecture

JOURNAL PUBLICATIONS

- Liu M. , Zhang W., & **Liang H.** Value effect of AI innovation zones: Green premium and cost reduction pathways in environmental disclosure, *Research in International Business and Finance*, 2026
- Li, Y., Huynh, L., Xu, Y., & **Liang, H.** The Forecast Ability of a Belief-Based Momentum Indicator in Full-Day, Daytime, and Nighttime Volatilities of Chinese Oil Futures, *Energy Economics*, 2023

CONFERENCE PROCEEDINGS

- **Liang, H.** & Hernke M. Quantifying the Environmental Impact of Electric Vehicle Adoption: An Analysis of California's Alternative Energy Vehicle Market, *International Conference on Frontiers of Energy and Environment Engineering*, Cairns, Australia, Dec 2-4, 2024

REVIEWER FOR JOURNALS

- Journal of Innovation & Knowledge
- Finance Research Letters
- International Review of Financial Analysis
- Research in International Business and Finance
- International Review of Economics and Finance
- Evaluation Review
- International Finance
- Review of Accounting and Finance
- Humanities and Social Sciences Communications

SKILLS

- **AI, Machine Learning & LLM Systems:** PyTorch, TensorFlow, Keras, Hugging Face, LangChain, LlamaIndex, Ragas, Pinecone, DeepSpeed, RLlib, Ollama, Dify, scikit-learn, statsmodels
- **Programming Languages:** Python, Java, JavaScript, TypeScript, SQL, R, Stata, HTML/CSS
- **Data Engineering & Platforms:** Snowflake, Databricks, Spark, Hive, Kafka, MySQL, MongoDB, Redis, Elasticsearch
- **Cloud & DevOps:** AWS (ECS/Fargate, EC2, S3, Lambda), Docker, CI/CD, Linux
- **Web & Application Frameworks:** Node.js, React.js, Streamlit, R Shiny
- **Data Visualization & BI:** Tableau, Power BI, D3.js, Matplotlib, Seaborn, Plotly, Bokeh, Vega-Lite, ggplot2

CERTIFICATIONS

- **Financial Risk Manager (FRM)** — Global Association of Risk Professionals (GARP)
- **Certificate in Sustainable Investing** — CFA Institute (formerly ESG Investing)