

HARRY (JIAJUN) HU

New York, Open for Relocation | 646-684-5950 | jiajunhuu@gmail.com | [LinkedIn](#) | [Personal Website](#)

EDUCATION

Washington University in St. Louis, St. Louis, MO

Aug 2024 - Dec 2025

M.Sc. in Data Analytics and Stats at McKelvey School of Engineering

New York University, New York, NY

Jan 2021 - May 2024

B.Sc. in Data Science and Math at College of Arts & Science

Relevant Courses: Machine Learning, Deep Learning, Advanced Computer Vision, Data Structure and Algorithms, Advanced AI Algorithms, Causal Inference, Database, Mathematical Statistics, Cloud Computing, Data Mining

SKILLS

Programming: Python (PyTorch, Scikit-learn, Pandas, NumPy, HuggingFace Transformers), SQL, Java, Git

Algorithms & Models: Deep Learning, Neural Networks, Computer Vision (Object Detection, Instance Segmentation, OCR), Transformer (Self-Attention, Cross-Attention, Masked-Attention), Transfer Learning, Ensemble Learning, Classification, Regression

Techniques: Feature Engineering, Hyperparameter Tuning, Cross-Validation, Model Deployment, A/B Testing, Causal Inference, Statistical Analysis, RAG, LLM Integration, Prompt Engineering, SHAP/LIME Interpretability, Data Preprocessing, Data Visualization

Tools & Platforms: AWS (S3, SageMaker), Docker, Kubernetes, HuggingFace Inference Endpoints, Power BI, CI/CD, Data Pipelines

PROFESSIONAL EXPERIENCE

Machine Learning Engineer, FPM Department, Washington University in St. Louis, MO

May 2025 - Dec 2025

- Designed and deployed a multi-model outlier-detection pipeline (rule baseline + Isolation Forest + DBSCAN + XGBoost); optimized Precision@K through model fusion and feature engineering; delivered SHAP-based interpretability analysis, identified 14% under-utilized space and projected ~\$45M revenue uplift.
- Built automated data-drift detection and scheduled retraining mechanism (PSI/KS tests + cron jobs), ensuring sustained model accuracy in production.
- Optimized form UX via A/B testing—cutting auditor input time by 25% and halving error rates; built live Power BI dashboards that raised daily audit coverage by 20%.

Machine Learning Engineer, Glodon, Beijing, China

June 2024 - Sept 2024

- Fine-tuned a pre-trained CNN model (PyTorch, transfer learning + rigorous K-Fold cross-validation) on semi-structured blueprint OCR data, slashing tally time from 60 min to 30 sec (99% improvement) with 97% component recognition accuracy.
- Diagnosed class imbalance and noise via EDA pipelines; engineered multi-strategy data augmentation (rotation, flip, Mixup) that boosted minority-class recall by 20%; monitored bias metrics alongside model assumptions documentation.
- Created a cloud-based NoSQL repository to centralize datasets and automate model label generation, ensuring scalability and audit compliance.

Machine Learning Engineer, Sigtica, Hong Kong

Jan 2023 - May 2023

- Refined a classification pipeline for historical documents using OCR + LLM-assisted outlier detection, achieving 87% precision and 93% recall across 3,000 labeled PDF images.
- Established data quality controls and a lightweight labeling QA loop to ensure model reliability in limited-data conditions.
- Packaged a repeatable inference workflow, unlocking data mining across 50,000 historical PDFs and translating results into actionable business recommendations.

Data Scientist, The 11th Impact LLC, New York, NY

May 2022 - Sept 2022

- Uncovered ad campaign opportunities through funnel analysis (views, clicks, revenue), boosting ROAS by 18%; analyzed user engagement patterns and recommended strategy adjustments to reduce drop-off at key stages.
- Built regression-based models to quantify the impact of ad features (creative type, placement, timing) on downstream revenue, guiding budget reallocation decisions.
- Designed and executed multiple end-to-end A/B tests (sample sizing, significance testing, controlled design) ensuring statistical rigor for CTR evaluation; built time-series forecasting models to optimize promotional strategies; developed Power BI + ETL pipelines for self-serve analytics.

PROJECTS

Cloud-Native Deep Vision Pipeline: Real-Time Detection & Segmentation, (Capstone at WashU) Spring - Fall 2025

- Developed an end-to-end CV pipeline: YOLO object detection + Mask2Former instance segmentation; explored Transformer-based encoder-decoder architecture, masked cross-attention mechanism, and multi-scale feature fusion strategies.
- Deployed on AWS S3 + SageMaker for distributed fine-tuning in Dockerized HuggingFace containers, improving mAP from 19% to 43% (+126%) through learning-rate scheduling, data augmentation, and systematic hyperparameter search.
- Served via HuggingFace Inference Endpoints (GPU autoscaling), delivering real-time segmentation and stable cloud inference.

Bank Loan Decision Prediction, (Capstone in Responsibility in DS at NYU)

Spring 2024

- Built LR / Random Forest / XGBoost ensemble models; applied SMOTE + threshold tuning + bagging variance reduction to handle class imbalance, achieving 92% recall on defaults and PR-AUC of 0.84.
- Applied SHAP, LIME, and Permutation Importance to identify key default drivers; delivered regulator-ready interpretability reports; audited model fairness across borrower demographics to inform equitable lending policies.