Mingzhe Hu

mingzhe.hu@columbia.edu | (+1) 646-925-0794 | https://superbtum.github.io

EDUCATION

Columbia University — **New York, US**
M.S. Electrical Engineering, GPA: 3.9 / 4.0 — *Feb 2023*
Relevant Courses: C++ Design, Big Data Analysis, Machine Learning, High Performance ML, DL, NLP, RL

Southeast University — **Nanjing, CN**
B.Eng. Information Engineering, GPA: 3.6 / 4.0 | Exchange @ Computer Science, TUM — *Jun 2020*
Relevant Courses: Computer Vision, Intro to Database, Computer Networks, Computer Graphics Seminar

EXPERIENCE

SoFi Inc. — **Salt Lake City, UT**
Software Engineer — *Jun 2023 – Present*

- Designed a portal for data science driven automated loan processing pipeline with event-based reporting using Airflow SNS-SQS-Lambda, **Kotlin/Spring** (backend), **JS/React** (frontend), and AWS S3 (cloud storage) for 10+ investors and 2 billion revenues
- Developed a streamlit app for loan details inspection with query template and Snowflake Cortex Analyst
- Designed critical loan data management microservices with state machine logic by rules to ensure accurate ownership transfers and prevent premature processing and support million-loan transfer
- Reduced streaming latency to near real-time on million-row scale via JPA batch processing, **debezium** and JDBC sink connectors for ETL (Postgres/Mongo to Snowflake) and JDBC source connector for reverse-ETL (Snowflake to Kafka)
- Integrated an **LLM-based** AI code review bot into Gitlab CI/CD pipelines in 100+ codebases across company, ensuring consistent code standards across teams and improving code quality with context-aware feedback

Columbia Engineering — **New York, NY**
Associate in Teaching Assistance — *Jan 2023 – May 2023*

ModelShare.ai — **New York, NY**
Software Intern — *Sep 2022 – Dec 2022*

- Implemented data cleaning and preprocessing pipeline and ONNX transformation in scikit-learn, PyTorch, Keras
- Generated online machine learning models with Amazon Web Service (AWS) **Lambda** microservices through serverless Restful API as an Amazon S3 proxy in API gateway
- Deployed and dockerize tunable AutoML (Dabl & FLAML) plugins with model ensemble and Ray parallelism

NVIDIA Corporation — **New York, NY**
Software Intern — *May 2022 – Sep 2022*

- Established multi-camera object tracking through camera calibration and homographic mapping with **OpenCV**
- Conducted spatial-temporal-appearance association, and self-adaptive thresholds for intra/inter-sensor DBSCAN clustering, with 80%+ IDF1 and 5~10% improvement of single camera tracking on warehouse and retail stores
- Developed a micro-batch data pipeline for streaming videos with **Apache Kafka** and unit-tested with coverage 85%+

SKILLS & HONORS

Programming Languages: Java/Kotlin, C/C++/Cython, Python, CUDA/OpenCL, HTML/CSS/JavaScript/Jinja2
Platforms & Tools: Google Cloud Platform, Amazon Web Service, CMake, Airflow, Snowflake, DataDog, Apache Spark/Kafka, BigQuery, React/Flask/Django, Spring/SpringBoot, Terraform, Git, Docker
Packages: DBT, PyTorch, TensorFlow, OpenMP, MongoDB/Postgres, TensorRT, Wandb

PROJECTS

A Web Application Demo of Earthquake Forecast — **New York, NY**
Project Leader, Columbia Hackathon w/. Best Beginner Hack — *Feb 2023*

- Designed layout with **HTML/CSS/JS/Jinja2** and visualized with Basemap and **Google Maps API**
- Built the website with **Flask** form requests and URL redirect in Python
- Trained the earthquake prediction model with data augmentation and autoML

Exam Paper and Question Bank Management System — **New York, NY**
Project Leader, Attached Course: C++ Design — *Sep 2022 – Dec 2022*

- Designed a **Qt6** based multi-user single-server system on GCP, with multithreaded JSON-based communication
- Built a POSIX-based non-blocking TCP socket communication with epoll callback and secured it with **OpenSSL**
- Managed data with multi-threading **SQLite** & **MongoDB**, improved speed by 30% with lock and clustered index
- Integrated **C++20** new features with 8X faster response (80 ms) and unit-tested and compiled with **CMake/Ninja**