YANG HU

yang-i-hu.github.io ◇ +1-805-637-9661 ◇ yang.h@nyu.edu ◇ Brooklyn, NY

EDUCATION

New York University Sep. 2024 - May 2026
Master of Science - *Computer Science*, GPA: 3.89/4.0
University of California, Santa Barbara Sep. 2020 - Mar. 2024
Bachelor of Science, Double Major in *Mathematics* and *Statistics & Data Science*

SKILLS

Languages	Fluent with C++, Python, SQL, R; Experienced with HTML/CSS
Frameworks	PyTorch, CUDA, OpenCV, Scikit-Learn, Langchain, Faiss, Qlib, Spark, PostgreSQL, Django
Tools	Google Cloud Platform, AWS, Travis CI, Redis, Git, SVN, Shell, LaTeX, Agile

EXPERIENCE

Quantitative Research Intern Jul. 2025 – Aug. 2025
Abama Private Fund Investment Management Co., Ltd. *Shanghai, China*

- Designed a parallelized data pipeline to transform raw event-level financial data into 300+ engineered features to feed high-frequency trading ML models, cutting data preparation from 8 h to 15 min.
- Delivered a live-updating factor store database powering large-scale model training and backtests; increased experiment throughput by $\sim 10\times$ and enabled apples-to-apples model comparison during strategy reviews.
- Fine-tuned Transformers for trading signal prediction, improving annualized return by 3% in backtests.
- Built an automated feature-discovery framework that iteratively surfaces noise-robust, signal-strengthening factors, accelerating model iteration and reducing manual screening.

R&D Software Engineering Intern May 2024 - Aug. 2024
Unity *Shanghai, China*

- Scaled the RAG knowledge base for Unity's AI assistant via community and forum ingestion and LLM-based quality filtering; expanded answer coverage for Editor/API issues by $\sim 10\times$, improving first-response resolution.
- Built *MuseBench*, an evaluation platform that pairs an LLM-as-a-judge with a curated benchmark of Q&A pairs, enabling consistent scoring of RAG results across different model versions and cutting evaluation time by 95%.
- Deployed a local-LLM pipeline to parse 100k+-line Unity Cloud Build logs and extract root-cause signals; reduced triage from hours to minutes and accelerated time-to-fix for recurring build failures.

Machine Learning Undergraduate Researcher Jun. 2023 - Mar. 2024
The WAVES Lab, University of California, Santa Barbara *Santa Barbara, CA*

- First-authored the *NeurIPS '23* Workshop paper "Segment-then-Classify: Few-shot instance segmentation for environmental remote sensing", a computationally efficient model with a novel structure.
- Proposed Segment-then-Classify: use SAM for zero-shot mask proposals, then a lightweight ViT classifier—achieved peak performance with $\sim 3\times$ less training data than the state-of-the-art model on large-scale satellite imagery.
- Recognized with the Schmidt Family Foundation Research Mentorship Award.

Computer Vision Lab Research Assistant Jun. 2022 - May. 2023
School of Data Science & Engineering, East China Normal University *Remote*

- Co-authored the *ACM MM '24* paper "Compacter: A Lightweight Transformer for Image Restoration", achieving state-of-the-art PSNR performance across different Image Restoration tasks with $\sim 50\% - 65\%$ fewer parameters.
- Co-designed Compact Adaptive Self-Attention, enabling omnidirectional spatial–channel information flow through cross-modulation of global context to strengthen long-range dependencies while preserving local detail.
- Proposed a Dual Selective Gated Module that dynamically injects global context into each pixel for context-adaptive aggregation, amplifying informative features and suppressing noise.
- Built the PyTorch training/benchmarking pipeline and ran ablation experiments, enabling reproducible results and efficient comparison to baselines.