

CHRISTINA JIA

New York, NY | (646) 812-1231 | yj2923@columbia.edu | [LinkedIn](#) | [Github](#)

EDUCATION

Columbia University M.S. in Computer Engineering GPA: 4.0 Relevant Coursework: Performance Evaluation & Metrics, DL on the Edge, DL for Robotic Manipulation, Computational Images	New York, NY Dec 2026
University of Michigan - Shanghai Jiao Tong University Joint Institute B.E. in Electrical and Computer Engineering Core Courses: Artificial Intelligence, Computer Architecture, Data Science and Analytics (Python), Data Structures and Algorithms, High-Performance Machine Learning (HPML) Systems, Computer Security, Software Engineering	Shanghai, China Aug 2025
New York University (Exchange Program) Core Courses: Operating Systems, Web Systems, Computer Networks, AIoT	New York, NY May 2024

INTERNSHIP

Johnson & Johnson Medical – CDO MT <i>Innovative Research Intern</i>	Shanghai, China Jul 2025 - Aug 2025
- Designed and built a Python + SQL data processing pipeline to ingest and normalize heterogeneous AI product data from 20+ global smart surgery companies across 5+ domains. - Developed rule-based and ML-driven decision workflows enabling automated benchmarking, reducing manual evaluation effort by ~60% and enabling repeatable, ML-driven analysis workflows. - Identified performance bottlenecks in batch data processing and SQL query performance and optimized to support iterative evaluations and downstream ML-driven analysis workflows used by digital scientists, reducing end-to-end latency by ~2.5×.	
iQIYI Shanghai Media Co., Ltd, Intelligent Platform Department <i>Applied Machine Learning Engineer Intern</i>	Shanghai, China Feb 2025 - Jul 2025
- Designed and built agent-driven image generation pipelines in PyTorch, refactoring preprocessing, inference, and post-processing stages to reduce latency by ~30% and improve stability for batch production, resolving 15+ production issues. - Fine-tuned and managed multiple LoRA adapters for style-controlled image generation, and built a reusable, batch-oriented inference pipeline enabling rapid style and version switching and cross-project reuse. - Productionized large-model image synthesis modules and integrated them into internal content workflows, supporting automated asset generation adopted by 3+ commercial projects with one generating ¥10K+ incremental revenue over a weekend. - Integrated LLM-based translation and video restoration modules into a unified agent workflow, comparing video enhancement methods and LLM-based translation workflows, achieving ~15–20% perceptual quality improvement (PSNR/SSIM) on legacy film content and improving subtitle translation consistency and fluency by ~25% in internal evaluations.	

TikTok

<i>Data-TNS-Engineering-Algorithm Infra Intern</i>	June 2026 - Present
- Migrated model-serving from personal devboxes to a team-shared deployment platform, removing single-owner dependency and standardizing releases so the team can jointly operate and reuse 25+ production content-safety models. - Automated multi-datacenter deployment with operator auto-registration and PSM switching over 120+ daily deploys. - Built a wiki CLI in Cursor syncing personal and team wikis, cutting engineers' knowledge-lookup time by ~40%. - Improved Alice's agent-evaluation accuracy ~15% via region-aware scoring and Top-K case surfacing before launch.	

TECHNICAL SKILLS

LLM & Agent Systems: Large Language Models (LLMs), Small Language Models (SLMs), Agentic AI, Retrieval-Augmented Generation (RAG), Prompt Engineering, Tool-augmented Reasoning, Multi-step Task Execution
ML & AI: Deep Learning, Computer Vision, NLP, Multimodal Learning, Generative Models, Transformers
ML Systems: Distributed ML Systems, ML Pipelines, Vector Search (FAISS), Kafka, Spark, Redis
Backend & Infrastructure: Python, FastAPI, Django, REST APIs, Docker, Kubernetes, AWS, Linux

PROJECT EXPERIENCES

Conversational AI Agent Evaluation Platform (Walmart Global Tech)	
• Built turn-level attribute mining and concept clustering pipeline for production conversational AI agents, extracting recurring failure patterns from thousands of real customer dialogues via NER, intent classification, and knowledge-grounding signals. • Designed a custom ranking algorithm with <code>selection_score</code>, <code>priority_score</code>, <code>failure_lift</code>, and <code>success_lift</code> terms to surface the highest-impact agent errors, validated against human annotations on 100+ sampled conversations across 5 rating tiers. • Built a user simulator that synthesizes dialogues whose attribute distribution matches real production traffic, enabling closed-loop agent improvement testing without dependence on fresh live customer conversations.	