Zhiyuan Wang

wangzhiyuanstan@gmail.com | (862) 363-0998 | hanyuuu.com |linkedin.com/in/zhiyuanwangstan

EDUCATION

Columbia University — **New York, NY**
M.S. Electrical Engineering — Sep. 2026 – Dec. 2027

Rensselaer Polytechnic Institute — **Troy, NY**
B.S. Computer Science and Mathematics — Sep. 2022 – Dec. 2025

TECHNICAL SKILLS

Programming Languages: Python, C++, Java, Go, SQL, Kotlin
Databases and Cloud: MySQL, PostgreSQL, AWS, GCP
Web/Mobile Development: React, Node.js, Spring, HTML/CSS, FastAPI, Android Studio
Machine Learning: PyTorch, Transformers, Reinforcement Learning, GNN, Curriculum Learning, Combinatorial Optimization
AI Systems & Tools: RAG, FAISS, Recommendation Systems, Computer Vision, IBM Qiskit

PROFESSIONAL EXPERIENCE

Learning-to-Anneal for Combinatorial Optimization, Tensor & Deep Learning Lab, Columbia University — New York, NY
Research Assistant — *Jan. 2026 – Present*

- Implemented L2A, a **Transformer-based** annealing framework for MaxCut, from the method specification into a full training and inference pipeline, including blockwise autoregressive sampling, curriculum-based local-search filtering, and policy-gradient training.
- Batched **RL training** across graphs and trajectories, achieving 12× faster training per epoch.
- Built an incremental **CUDA-graph–accelerated** tabu local search, reducing single-graph inference time by up to 10× and GPU memory usage by 3.5×.
- Designed a multi-scale MaxCut benchmark over Barabási–Albert(BA) and Erdős–Rényi(ER) graphs with 100–1,500 nodes and 30 held-out instances per size; used disjoint train/test seeds and cross-scale evaluation, comparing all methods on identical instances by cut value, optimality gap, and wall-clock time against Gurobi (1-h limit), ECO-DQN, S2V-DQN and X2GNN.
- Matched or exceeded Gurobi on all 960 BA (+0.6%) and ER (+0.1%) instances with 100–1,200 nodes and outperformed Gurobi (1-h limit), ECO-DQN, S2V-DQN and X2GNN on every instance with 100–1,200 nodes.

Reinforcement Learning for Ising Models: Dataset and Benchmark , RPI — Troy, NY
Research Assistant — *May. 2025 – Sep. 2025*

- Established the first RL benchmark for Ising-model optimization in **ML4Physics**, standardizing evaluation across 3 RL algorithms and industrial-grade MIP solvers.
- Curated and released around 190K Ising instances spanning 5 dimensional settings and 7 problem families, with problem sizes ranging from 3 to 8,000 spins.
- Developed a **Transformer-based RL policy** for large-scale Ising **optimization**, achieving 1%–9.11% solution-quality gaps relative to industrial-grade solvers while scaling to large problem instances.
- Built a reproducible benchmarking pipeline for consistent comparison of solution quality, scalability, and computational performance across RL and classical optimization methods.

Quantum Circuit Design via Reinforcement Learning, RPI — Troy, NY
Research Assistant — *Sep. 2024 – May. 2025*

- Led a 5-member research team to formulate quantum circuit synthesis as 3 distinct mathematical/MDP representations, and evaluated them using **Q-Learning and Deep Q-Networks** across 10 circuit-design tasks.
- Demonstrated that RL could recover target quantum circuits across all 10 task categories, with the best formulations achieving up to 100% success over 100 training runs on Bell-state synthesis and 91% on the multi-qubit Toffoli gate.
- Identified Reverse Matrix and **Tensor Network (TN)** representations as the most effective formulations for different circuit classes, demonstrating the impact of mathematical state representation on RL convergence and solution quality.

PROJECTS

SitePulse: Web Status Monitoring Platform

- Built an asynchronous monitoring platform with **FastAPI, Celery, and Redis**, processing up to 200 URLs/job with 10 concurrent checks, live progress, and 8-way HTTP/network status classification.
- Engineered fault-tolerant execution using late acknowledgements, **PostgreSQL** advisory locks, retries, and idempotency keys, preventing lost, double-counted, and duplicate jobs under worker failures and repeated submissions.
- Added DNS-level SSRF protection, Redis rate limiting, distributed request tracing, **Dockerized** deployment, and 31 automated CI tests.

Spoiler-Free RAG E-Book Reader

- Built a spoiler-safe **RAG** e-book reader with paragraph-level retrieval boundaries, preventing unread content from entering LLM context while returning source-linked answers.
- Designed a hierarchical retrieval pipeline with **Qwen3** embeddings, **pgvector**, hybrid search, and reranking, using full context up to 60K tokens or reducing 40 retrieved candidates to 8 evidence chunks for larger books.
- Built asynchronous EPUB ingestion, long-term chapter/character memory, streaming **DeepSeek** responses, and multi-turn context over the last 6 turns, backed by 63 unit tests.
- Implemented versioned index configurations to preserve retrieval consistency across changes to embedding models and chunking strategies.

StreamLens: Twitch Content Discovery and Recommendation Platform

- Built a Twitch discovery platform with 18 **Spring Boot** REST endpoints, supporting search and personalized recommendations across 3 content types: live streams, videos, and clips.
- Integrated the Twitch Helix API via **OpenFeign** and **OAuth2**, with automatic token refresh and persistence through a 4-table **MySQL** schema on **AWS RDS** for users, favorites, and content metadata.
- Developed a content-based recommender that derives preferences from a user's top 3 favorite games and generates 20-item feeds per content type, filtering saved items and falling back to trending games for cold-start users.
- Added **Caffeine caching** across 4 high-traffic paths to bypass repeated Helix API calls within a 1-minute TTL, then secured and deployed the system with Spring Security, Docker, AWS App Runner, and Vercel.

EchoDocs AI: Real-Time Document RAG and Voice Assistant

- Built a full-stack document assistant using Next.js and its **Backend-for-Frontend (BFF)** layer to support conversational queries and AI response generation.
- Integrated the **Vercel AI SDK** with **Google Gemini 3.8 Flash** to orchestrate document-based question answering and streaming responses.
- Implemented bidirectional voice interaction using the browser-native **Web Speech API**, including speech recognition and text-to-speech output.
- Developed a responsive conversational interface using shadcn/ui and **Tailwind CSS**.

PUBLICATIONS

- Lin, L. , **Wang, Z.** , et al. Reinforcement Learning for Ising Models: Datasets and Benchmark, NeurIPS ML4PS Workshop, 2025.
- **Wang, Z.** , et al. Reinforcement Learning for Quantum Circuit Design: Using Matrix Representations, AAAI-25 Workshop on Structured Low-Rank Representations in AI (CoLoRAI), 2025