

Yuhan Ma

☎ +1 9296382817 ✉ ym3447@nyu.edu 🌐 github.com/wmyh-0416

Education

New York University, New York

M.S. in Computer Science, Concentration in Artificial Intelligence

Aug. 2025 – May 2027

Anhui University, Hefei, China

B.E. in Computer Science and Technology

GPA: 3.9/4.0

Sep. 2021 – Jun. 2025

GPA: 3.73/4.0

- Outstanding Undergraduate Graduate of Anhui University, Class of 2025 & Exchange Program: Deakin University (Australia).

Skills

Programming Languages: Python, SQL, C++, Bash

Recommender Systems: Multi-Recall, ItemCF, Two-Tower, Graph Embedding, LightGBM, DIN, BPR-MF, GRU4Rec, SASRec, BERT4Rec, Ranking & Re-ranking

Deep Learning / LLMs: PyTorch, Transformer, GPT, RoPE, RMSNorm, SwiGLU, SFT, RLHF, REINFORCE, GRPO

Projects

Four-Stage Recommendation System for Short-Video Feed Based on KuaiRand-Pure

Jan. 2026 – Mar. 2026

- Built a four-stage recommendation system for a short-video feed scenario from scratch on KuaiRand-Pure, covering the full pipeline of multi-channel retrieval, coarse ranking, fine ranking, and re-ranking; adopted time-based data splitting to prevent information leakage and constructed a reproducible recommendation pipeline by materializing intermediate candidates.
- Implemented four retrieval channels: Popular, ItemCF, Two-Tower, and Graph Embedding; constructed an item-item graph from positive feedback and learned representations via random walks, achieving Recall@100 of 0.2476 / 0.2343 (val/test) with candidate coverage close to 1.0.
- Trained a LightGBM coarse-ranking model on retrieved candidates with lightweight features including user/item features, statistical features, source scores, and freshness; reduced per-user candidates from 500 to 100, achieved val AUC=0.8149, and improved top-100 absolute Recall to 0.3523 / 0.3317.
- Developed a DIN-style multi-task fine-ranking model based on user behavioral sequences to jointly predict long watch, finish, and like; achieved finish AUC=0.6663, like AUC=0.6219, and NDCG@20=0.0359 on the test set. Designed a rule-based re-ranking module with author frequency control, tag diversity, and freshness adjustment, improving avg unique tags per user by +0.302 and reducing adjacent same tag rate by 0.0171, while maintaining slight gains in both NDCG@20 and Recall@20.

Retrieval-Ranking Benchmark for Recommender Systems (MovieLens)

Sep. 2025 – Nov. 2025

- Built a two-stage recommender framework (retrieval + ranking) with a unified data processing and evaluation pipeline, including negative sampling, Recall, and NDCG.
- Implemented and benchmarked multiple models, including BPR-MF, GRU4Rec, SASRec, BERT4Rec, and Dual-Tower.
- Conducted large-scale experiments on MovieLens-32M, where BERT4Rec achieved Recall@10=0.97 and NDCG@10=0.80.

GPT-Style Transformer Language Model from Scratch (TinyStories)

Jan. 2026 – Feb. 2026

- Independently implemented a byte-level BPE tokenizer and core GPT-style Transformer modules, including causal self-attention, RoPE, RMSNorm, and SwiGLU.
- Designed a streaming tokenization + uint16 memmap data pipeline, preprocessing TinyStories into 541M training tokens and 5.46M validation tokens, enabling low-memory random-sampling training and checkpoint-based resume.
- Built a reproducible experimentation framework supporting cosine LR warmup, learning-rate sweep, batch-size sweep, and architectural ablations; reduced sampled-dev loss to 1.444 and full-dev loss to 1.475, and analyzed the learning-rate stability boundary as well as the impact of RMSNorm, RoPE, and pre-norm on convergence.

Reinforcement Learning for LLM Reasoning (SFT + GRPO)

Feb. 2026 – Mar. 2026

- Built a training and evaluation pipeline for LLM reasoning based on the Qwen2.5-Math-1.5B series, completing zero-shot, SFT, and GRPO experiments on both MATH and Countdown tasks.
- Implemented REINFORCE, group-relative baseline, clipped GRPO, masked log-prob scoring, as well as micro-batch accumulation and rollout updates.
- Designed a reward function and automatic grader for Countdown, supporting answer-tag parsing, expression evaluation, arithmetic validation, and numeric-constraint checking.
- Completed comparative experiments across training settings and improved Countdown validation accuracy to 32.4%; additionally conducted SFT data-scale experiments and result analysis.

Experience

National Quantum (Hefei, China)

Jan. 2024 – Mar. 2024

Machine Learning Engineer Intern

- Contributed to a magnetic-particle image segmentation project; designed data processing and annotation workflows for a low-data setting, using SAM to generate initial labels followed by manual correction, which improved data preparation efficiency.
- Built and optimized a UNet training pipeline in PyTorch; introduced Albumentations-based augmentation strategies to mitigate overfitting caused by limited data and improve generalization.
- Addressed sparse foreground and class imbalance in the segmentation task by combining Optuna-based hyperparameter search with joint optimization using Focal Loss and Dice Loss, improving validation-set segmentation performance.
- Completed model export and inference deployment; implemented production-side inference with ONNX Runtime, gaining end-to-end engineering experience from CV model training to deployment.