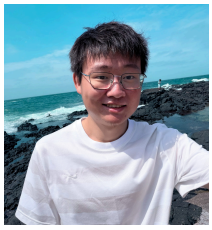

Shiyan Liu (刘师言)

shyliu.china@gmail.com · shiyanliu.com · github.com/shiyan-liu

scholar.google.com/citations?user=ITv29t8AAAAJ

Research Interests

Neural Combinatorial Optimization	Geometric equivariance: MViewRouter Multimodal fusion: VAGPO
Dense Retrieval	Positive-unlabeled: PURE Test-time adaptation: DART
Trustworthy LLM Systems	Probabilistic evaluation: DICE Explainable optimization: VISTA

Education

B.S. in Data Science & Big Data Tech., **Huazhong University of Science and Technology (HUST)** Wuhan, China | Sep, 2022 – Jun, 2026

- Weighted Average: 88/100 (Major Courses: 95/100), Overall Rank: 6/60 (Top 10%)
- Selected Honors: Outstanding Graduate (2026), National Scholarship (2025), Merit Student Scholarship (2025), Tencent Scholarship (2025), Academic Excellence Scholarship (2024)
- Selected Courses: Java (100), Big Data Processing (99), Machine Learning (98), Big Data Analytics (97), Graph Neural Networks (96), Computer Vision (95), Natural Language Processing (95)

Visiting Student, **University of California, Berkeley** Berkeley, US | Jan, 2026 – May, 2026

- Selected Courses: CS188 Artificial Intelligence, CS189 Machine Learning

M.S. in Computing (Artificial Intelligence and Machine Learning), **Imperial College London** London, UK | Sep, 2026 – Dec, 2027

Work Experience

JD.com Beijing, China | Mar, 2025 – Aug, 2025

Java Backend / Agent / Data Development

- **Distributed Systems Development**: Led microservice development for JD Finance App's community module, handling 8M+ DAU high-concurrency traffic; built a Flink real-time computing pipeline using sliding windows for large-scale stream processing and interactive data mining, ensuring high reliability and low latency for hot data.
- **LLM RAG Optimization**: Built a 300,000+ document multi-source RAG knowledge base, completing production-grade deployment of vector storage and embedding models using Elasticsearch and NVIDIA Triton; improved retrieval accuracy through optimized chunking and embedding strategies, introduced the RAGAS framework for quantitative business evaluation; extended the work to develop the DICE evaluation benchmark addressing poor interpretability in existing evaluation systems, published as first author at NeurIPS '25 Workshop.
- **End-to-end Hot Topic Mining**: Built an end-to-end data pipeline from message queues and vectorization to unsupervised learning, applying UMAP dimensionality reduction and HDBSCAN clustering to mine hot topics from 20,000+ daily posts, achieving 90% precision with full-pipeline latency under 5 minutes, significantly reducing operational costs; built an interactive data visualization platform with Plotly supporting multi-dimensional quality assessment and business insights, improving algorithm interpretability and decision-making efficiency.

HUST Reinforcement Learning and Intelligent Computing Lab Wuhan, China | Jun, 2024 – Present

Research Assistant (Advisor: Associate Professor Yan Jin)

- As first author, led multiple frontier research projects including MViewRouter (combinatorial routing) and VAGPO (combinatorial routing), covering neural combinatorial optimization, deep reinforcement learning, and representation learning; results submitted to top international conferences including NeurIPS '26 and UAI '26.
- Served as leader of a Hubei Province College Innovation and Entrepreneurship Training Program, leading the "AI-based Digital Human Design and Implementation" project to excellent completion.

Research Experience

DART, **Test-Time Training for Zero-Resource Dense Retrieval Reranking** [1] Berkeley, US | Mar, 2026 – May, 2026

First Author, **Accepted** at KnowFM @ ACL 2026

- Identified a fundamental dilemma in zero-resource reranking: cross-encoders deliver high accuracy but require costly training and high latency, while unsupervised BM25 reranking consistently degrades performance on most BEIR benchmarks.
- Proposed DART (Dense Adaptive Reranking at Test-time), which adapts a bilinear scoring matrix W per query at inference time using pseudo-labels derived from the dense retrieval ranking itself; introduced confidence-weighted margin loss with AdaMargin and cross-query momentum (MetaInit + EMA) to improve robustness and convergence across the query stream.
- Achieved +2.1% mean per-dataset relative NDCG@10 gain over the dense retrieval baseline on six BEIR benchmarks with under 10ms additional latency per query, outperforming LLM-based training-free rerankers (ASRank, ICR, InstUPR) that require ∼200ms per query.

VISTA, **Exposing and Escaping the Black Box in Reflective Prompt Optimization** [2] Berkeley, US | Jan, 2026 – Mar, 2026

First Author, **Accepted** at ACL 2026 (SRW)

- Systematically identified and analyzed four core limitations of reflective automatic prompt optimization, including uninterpretable black-box trajectories and susceptibility to local optima; experiments revealed that baseline methods (e.g., GEPA) suffer severe performance degradation from 23.81% to 13.50% accuracy with defective initial seeds (normal performance should be 80%+).

- Proposed the VISTA framework, decoupling hypothesis generation and prompt rewriting into a multi-agent system; introduced prior heuristic knowledge based on common error patterns, transforming optimization trajectories from opaque black-box processes into semantically observable evolutionary paths; introduced parallel mini-batch validation to ensure feedback accuracy and optimization stability.
- Innovatively designed a dual-layer exploration-exploitation mechanism combining random restarts and ε-greedy sampling to effectively escape local optima; on mathematical reasoning benchmarks GSM8K and AIME2025, VISTA restored and improved defective seed performance from 13.50% to 87.57%, significantly outperforming baseline models.

MViewRouter, **Geometric-Equivariant Routing via Multi-view Alternating Attention** [3] Wuhan, China | Nov, 2025 – Jan, 2026
First Author, NeurIPS '26 Under Review

- Systematically analyzed performance bottlenecks of traditional deep RL methods in combinatorial routing due to neglect of geometric symmetry; by internalizing geometric equivariance as an inductive bias, the model captures rotation- and flip-invariant high-order topological representations via representation learning, unifying data augmentation in training and inference to improve solution quality while reducing variance across multiple views.
- Innovatively designed a Multi-view Alternating Attention (MAA) mechanism, constructing parallel encoding pipelines in the D_4 symmetry group space; through alternating intra-view relationship modeling and inter-view feature aggregation, using fiberization embedding to align local representations to a unified geometric space, achieving deep integration of geometric equivariance.
- Proposed a Collective Policy Gradient Aggregation (CPGA) algorithm using multi-view consensus gradients to stabilize training and accelerate convergence; experiments on TSP and CVRP benchmarks and real-world TSPLIB instances demonstrate competitive solution quality and strong zero-shot generalization.

DICE, **Discrete Interpretable Comparative Evaluation with Probabilistic Scoring for RAG** [4] Wuhan, China | Aug, 2025 – Nov, 2025
First Author, **Accepted** at ResponsibleFM @ NeurIPS 2025

- Addressing poor interpretability and difficulty in quantifying confidence for RAG evaluation metrics, proposed the DICE framework: decoupling evidence alignment and probabilistic scoring into two stages, representing evaluation decisions as discrete random variables, and quantifying uncertainty using the probability distribution over token set $\{A, B, \text{Tie}\}$.
- Introduced a Swiss-tournament-based mechanism, reducing computational complexity of multi-system comparison from $O(n^2)$ to $O(n \log n)$; validation on a Chinese financial QA dataset shows DICE achieves 85.7% consistency with expert gold-standard evaluation, significantly outperforming RAGAS (45.7% consistency), providing a scalable automated evaluation benchmark for high-reliability RAG systems.

VAGPO, **Vision-augmented Asymmetric Group Preference Optimization for Graph Routing** [5] Wuhan, China | Jan, 2025 – May, 2025
First Author, UAI '26 Under Review

- Addressing limitations of traditional methods in capturing topological features in combinatorial routing, proposed the VAGPO framework: integrating ResNet visual encoder and Transformer sequence model to extract high-order graph structure representations from spatial geometry and temporal dependency dimensions with cross-modal fusion, significantly enhancing perception of complex node distributions.
- Innovatively introduced an asymmetric group preference optimization strategy extending POMO multi-start strategies, guiding policy gradients by contrasting the relative quality of multiple solution groups, overcoming slow convergence and susceptibility to local optima in traditional RL; experiments demonstrate VAGPO outperforms baselines in convergence speed ($\times 6$) and solution quality (TSP-100, 0.03% Gap vs. POMO 0.17% Gap), with stronger zero-shot generalization on complex-distribution and large-scale instances (TSP-1000) without fine-tuning.

Project Experience

PURE, **Positive-Unlabeled Contrastive Retrieval for Graduate Program Recommendation** [6] Wuhan, China | Apr, 2026 – Jun, 2026
Project Leader, EMNLP '26 Under Review, Online Demo

- Addressing the fundamental challenge of missing negative labels in graduate admissions (rejections go unreported), formulated program recommendation as a positive-unlabeled retrieval problem; crawled and built a real-world dataset covering 531 programs and 15,291 admission records, with over half of students receiving multiple offers, exhibiting a many-to-many matching structure.
- Proposed the dual-encoder framework PURE: programs are represented offline by fusing GPA-weighted prototype embeddings with program description text embeddings; applicants are encoded via gated structural residual fusion of structured features (GPA, school tier) and free-text experience descriptions; trained with InfoNCE loss with in-batch positive masking and hard negative mining, with a multi-positive MRR evaluation protocol to correct false-negative bias.
- Experiments show PURE achieves MRR of 26.96%, outperforming BM25 keyword retrieval (+4.9%) and GPT-4o-mini zero-shot ranking (+8.3%), while matching Biencoder + GPT-4o-mini reranking at negligible inference cost, validating the dual-encoder approach in this high-stakes information retrieval setting.

Knowledge Graph-based Heterogeneous RAG Sleep Diagnosis Platform Wuhan, China | Oct, 2024 – Dec, 2024
Project Leader, Global Campus AI Algorithm Competition – National Second Prize

- Addressing the disconnect between unstructured literature and structured knowledge in sleep disorder diagnosis, built a vectorized document and knowledge graph heterogeneous knowledge base; processed 500,000+ characters of Chinese medical literature into Markdown-structured vectorized form, and used LLMs to extract and build a Neo4j knowledge graph with 4,000+ entities and 500+ relationships.
- Developed a multi-stage retrieval pipeline based on LangChain: mapping user natural language symptoms to professional vocabulary via semantic matching, performing multi-hop reasoning on the knowledge graph to obtain related subgraphs (recall rate 80%+), using graph context to guide document retrieval, and aggregating results from all stages as LLM context; effectively solved long-tail medical knowledge gaps and LLM hallucination issues, ensuring diagnostic accuracy in high-risk medical scenarios.
- Developed a 3D knowledge graph real-time visualization interface using Three.js, and completed cross-platform mobile deployment using uni-app to serve both doctors and patients; the project won the National Second Prize at the Global Campus AI Algorithm Competition for its potential in improving healthcare accessibility and diagnostic efficiency.

Competition Awards

Huawei Developer Competition National 1st Prize | 2024
Global Campus AI Competition (KG) National 2nd Prize | 2024
CCF Open Source Competition National 3rd Prize | 2024
Huawei Developer Competition National 3rd Prize | 2024
National Math Modeling Competition Hubei 2nd Prize | 2024
National College English Competition National 1st Prize | 2024
Global Campus AI Competition (AIGC) National 2nd Prize | 2024
National Info. Storage Competition National 3rd Prize | 2025
National CS Systems Competition Excellence Award | 2024
Ruikang Robot Developer Competition Hubei 3rd Prize | 2024

Skills

- **Programming Languages**: Proficient in Python, Java; familiar with C++, C, C#, JavaScript.
- **Machine Learning**: Familiar with common ML algorithm theories and applications, including unsupervised learning (clustering, dimensionality reduction), supervised learning (linear regression, logistic regression, neural networks), probabilistic graphical models (Bayesian networks, HMM), regularization techniques (L1, L2), gradient descent (SGD, Adam), and model evaluation (cross-validation, confusion matrix, ROC curve).
- **Deep Learning**: Familiar with deep learning concepts and mainstream architectures including MLP, CNN, RNN, and Transformer; proficient in PyTorch for model development and tuning.
- **Reinforcement Learning**: In-depth understanding of MDP modeling and RL theory, including Q-learning and Policy Gradient; familiar with mainstream policy optimization algorithms such as PPO, DPO, and GRPO.
- **Agentic AI**: Familiar with Agentic AI concepts and applications; proficient in related techniques including CoT, RAG, MCP, Skill, and Harness implementation and tuning.

Misc

- Languages: Chinese – native speaker; English – proficient (IELTS 7.5, GRE 327)
- Hobbies: Football (Captain of HUST CS School Football Team), sports, travel, etc.

Publications & Preprints

[1] **S. Liu** and Y. Li, "Test-Time Training for Zero-Resource Dense Retrieval Reranking," in *KnowFM @ ACL 2026*, 2026.

[2] **S. Liu**, Q. Xia, Q. Xia, Y. Liu, X. Yu, and R. Qu, "Reflection in the Dark: Exposing and Escaping the Black Box in Reflective Prompt Optimization," in *ACL 2026 (Student Research Workshop)*, 2026.

[3] **S. Liu**, B. Tan, Y. Wu, and Y. Jin, "MViewRouter: Internalizing Geometric Equivariance via Multi-view Alternating Attention for Combinatorial Routing," *under review*, 2026.

[4] **S. Liu**, J. Ma, and R. Qu, "DICE: Discrete Interpretable Comparative Evaluation with Probabilistic Scoring for Retrieval-Augmented Generation," in *ResponsibleFM @ NeurIPS 2025*,

[5] **S. Liu**, B. Tan, Z. Cao, and Y. Jin, "VAGPO: Vision-augmented Asymmetric Group Preference Optimization for Graph Routing Problems," *arXiv preprint arXiv:2508.01774*, 2025.

[6] **S. Liu** and R. Qu, "Positive-Unlabeled Contrastive Retrieval for Graduate Program Recommendation," *under review*, 2026.