Yuyu Liu

Stony Brook, New York, USA

yuyu.liu@stonybrook.edu | Personal Page | Google Scholar | (631) 406-0344

EDUCATION BACKGROUND

State University of New York at Stony Brook, Ph.D. — **New York, USA**
Major: Computer Science — *Aug. 2025- present*

Southern University of Science and Technology (SUSTech), B.Sc. — **Guangdong, China**
Major: Mathematics — *Sep. 2020- Jul. 2024*

SELECTED PUBLICATIONS

- **HyperGuide: Hyperbolic Guidance for Efficient Multi-Step Reasoning in Large Language Models**. Yuyu Liu, Haotian Xu, Yanan He, Sarang Rajendra Patil, Mengjia Xu, Tengfei Ma. *Preprint*.
- **HypEHR: Hyperbolic Modeling of Electronic Health Records for Efficient Question Answering**. Yuyu Liu, Sarang Rajendra Patil, Mengjia Xu, Tengfei Ma. *Findings of ACL 2026*.
- **Efficient Imputation for Patch-based Missing Single-cell Data via Cluster-regularized Optimal Transport**. Yuyu Liu, Jiannan Yang, Ziyang Yu, Weishen Pan, Fei Wang, Tengfei Ma. *ACM BCB 2026*.

WORK EXPERIENCE

Machine Learning Engineer Intern — **Meituan, Inc.**
Search Department, Retrieval Group — *Nov. 2024 - Apr. 2025*

- Built an end-to-end generative item retrieval channel for Meituan app search using Sentence-T5, adapting the dense retrieval framework of LIGER (Meta, 2025); on the NC dataset, improved Hit@100 / Hit@500 by xxx% / xxx% over traditional retrieval baselines.
- Developed a user query category prediction model with a Qwen2.5-0.5B backbone to guide candidate generation, improving retrieval relevance and coverage for long-tail and ambiguous queries.

Machine Learning Engineer Intern — **Baidu, Inc.**
Digital Human Group — *Jan. 2024 - May. 2024*

- **Real-Time Voice Cloning for Live Streaming**
  Built a voice-cloning model that faithfully replicates a target speaker's voice with natural, diverse speech while running efficiently under limited compute.
  - *Architecture*: Replaced the Bert-VITS2 backbone with Chinese-pretrained open-source models (e.g., chinese-roberta-wwm-ext-large); deployed with PyTorch on 8×A100 GPUs to fit target timbres.
  - *Method*: Introduced the few-shot GPT-SoVITS pipeline with LoRA fine-tuning, enabling timbre cloning from under 10 minutes of speech.
- **Digital-Human Lip-Sync Post-Processing**
  Improved the lip motion of 3D digital humans rendered in Unreal Engine 5.
  - *Algorithm*: Implemented a C++ routine that computes the distance between upper- and lower-lip blendshape (BS) vectors and snaps the teeth closed below a threshold, fixing a teeth-not-closing modeling artifact.

PROJECTS

Fair-Aurora: Comparing Fairness Strategies for Reinforcement Learning-Based Congestion Control in Multi-Flow Environments

Yuyu Liu, Thomas Mbrice. — *Feb. 2026 - May 2026*

Studied the fairness of Aurora, a deep RL congestion controller, when multiple flows share a bottleneck, and evaluated three post-hoc strategies (reward shaping, observation augmentation, and loss-sensitivity tuning) using Jain's fairness index on a custom shared-bottleneck simulator. Showed that modest reward shaping yields the best fairness while preserving aggregate throughput, achieving fairness through bandwidth redistribution rather than reduction.

AWARDS & ACHIEVEMENTS

- College Special Scholarship (Second Class, 2022)
- College Special Scholarship (First Class, 2021)
- Jia Zhaoye · Special Scholarship of Principal(36 out of 1092, 2020)