

Guanchu (Gary) Wang

Ph.D. Candidate in Computer Science
Rice University
6100 Main Street Duncan Hall 3122 Houston, TX 77005-1892

(+1)832-875-9593
gw22@rice.edu
<https://guanchuwang.github.io/home>

Research Interests & Highlights [Google Scholar]

My research centers on artificial intelligence (AI), machine learning (ML), including the following topics:

- **Trustworthy Artificial Intelligence:** Interpretable AI, Explainable AI, ML Safety.
- **Large Language Models (LLM):** LLM interpretability/Reasoning, Multi-agent LLMs.
- **Health Informatics:** AI for Healthcare, Retrieval-augmented Generation.

The highlights of my research are as follows:

- **CIKM'22 Best Demo Award, CIKM'23 Best Demo Honorable Mention, ICML'22 Spotlight.**
- 23 peer-reviewed research papers (12 first-author) published in top venues.
- 800 citations; h-index: 14; Ken Kennedy Institute Fellowship.

Education

- Rice University** *Aug 2020 - May 2025 (expected)*
Ph.D. candidate in Computer Science, Advisor: Dr. Xia (Ben) Hu and Dr. Vladimir Braverman
- University of Science and Technology of China** *Sep 2016 - May 2019*
M.S. in Electrical and Computer Engineering, Advisor: Dr. Chen Gong and Dr. Zhengyuan Xu
- Dalian University of Technology** *Sep 2012 - May 2016*
B.S. in Electrical and Computer Engineering

Publications [Google Scholar]

Conference and Journal Publications: (* denotes equal contributions)

- [ICML' 24] **TVE: Learning Meta-attribution for Transferable Vision Explainer**
Guanchu Wang, Yu-Neng Chuang, Fan Yang, Mengnan Du, Chia-Yuan Chang, Shaochen Zhong, Zirui Liu, Zhaozhuo Xu, Kaixiong Zhou, Xuanting Cai, Xia Hu
- [ICML' 23] **DIVISION: Memory Efficient Training via Dual Activation Precision**
Guanchu Wang, Zirui Liu, Zhimeng Jiang, Ninghao Liu, Na Zou, Xia Hu
- [ECML' 23] **Mitigating Algorithmic Bias with Limited Annotations**
Guanchu Wang, Mengnan Du, Ninghao Liu, Na Zou, Xia Hu
- [ICML' 22] **Accelerating Shapley Explanation via Contributive Cooperator Selection**
Guanchu Wang, Yu-Neng Chuang, Mengnan Du, Fan Yang, Quan Zhou, Pushkar Tripathi, Xuanting Cai, Xia Hu
Spotlight Paper
- [EMNLP' 24] **Secured Weight Release for Large Language Models via Taylor Expansion**
Guanchu Wang*, Yu-Neng Chuang*, Ruixiang Tang, Shaochen Zhong, Jiayi Yuan, Hongye Jin, Zirui Liu, Vipin Chaudhary, Shuai Xu, James Caverlee, Xia Hu
- [ICLR' 23] **CoRTX: Contrastive Learning for Real-time Explanations**
Guanchu Wang*, Yu-Neng Chuang*, Fan Yang, Quan Zhou, Pushkar Tripathi, Xuanting Cai, Xia Hu

- [NeurIPS' 23] **Winner-Take-All Column Row Sampling for Memory Efficient Adaptation of Language Model**
Guanchu Wang*, Zirui Liu*, Shaochen Zhong, Zhaozhuo Xu, Daochen Zha, Ruixiang Tang, Zhimeng Jiang, Kaixiong Zhou, Vipin Chaudhary, Shuai Xu, Xia Hu
- [CIKM' 22 Demo] **BED: A Real-Time Object Detection System for Edge Devices**
Guanchu Wang*, Zaid Pervaiz Bhat*, Zhimeng Jiang*, Yi-Wei Chen*, Daochen Zha*, Alfredo Costilla Reyes*, Afshin Niktash, Gorkem Ulkar, Erman Okman, Xuanting Cai, Xia Hu
Best Demo Award
- [TKDD' 23] **Efficient GNN Explanation via Learning Removal-based Attribution**
Yao Rong, Guanchu Wang, Qizhang Feng, Ninghao Liu, Zirui Liu, Enkelejda Kasneci, Xia Hu
ACM Transactions on Knowledge Discovery from Data
- [CIKM' 23 Demo] **DiscoverPath: A Knowledge Refinement and Retrieval System for Interdisciplinarity on Biomedical Research**
Yu-neng Chuang, Guanchu Wang, Chia-Yuan Zhang, Kwei-Herng Lai, Ruixiang Tang, Fan Yang, Alfredo Costilla-Reyes, Kaixiong Zhou, Xiaoqian Jiang and Xia Hu
Best Demo Final-list
- [AISTATS' 24] **Personalizing Low-Rank Bayesian Neural Networks Via Federated Learning**
Boning Zhang, Dongzhu Liu, Osvaldo Simeone, Guanchu Wang, Dimitrios Pezaros, Guangxu Zhu
- [EMNLP' 24 Findings] **KV Cache Compression, But What Must We Give in Return? A Comprehensive Benchmark of Long Context Capable Approaches**
Jiayi Yuan, Hongyi Liu, Shaochen Zhong, Yu-Neng Chuang, Songchen Li, Guanchu Wang, Duy Le, Hongye Jin, Vipin Chaudhary, Zhaozhuo Xu, Zirui Liu, Xia Hu
- [NeurIPS' 23] **Chasing Fairness under Distribution Shift: a Model Weight Perturbation Approach**
Zhimeng Jiang, Xiaotian Han, Hongye Jin, Guanchu Wang, Rui Chen, Na Zou, Xia Hu
- [NeurIPS' 21] **Fairness via Representation Neutralization**
Mengnan Du, Subhabrata Mukherjee, Guanchu Wang, Ruixiang Tang, Ahmed Hassan Awadallah, Xia Hu
- [NeurIPS' 21] **Revisiting Time Series Outlier Detection: Definitions and Benchmarks**
Kwei-Herng Lai, Daochen Zha, Junjie Xu, Yue Zhao, Guanchu Wang, Xia Hu
- [IJCNN' 21] **Learning Transitional Skills with Intrinsic Motivation**
Qiangxing Tian, Jinxin Liu, Guanchu Wang, Donglin Wang
- [AAAI' 21 Demo] **TODS: An Automated Time Series Outlier Detection System**
Kwei-Herng Lai, Daochen Zha, Guanchu Wang, Junjie Xu, Yue Zhao, Junjie Xu, Yue Zhao, Devesh Kumar, Yile Chen, Purav Zumkhawaka, Minyang Wan, Diego Martinez, Xia Hu
- [IJCAI' 20] **Independent Skill Transfer for Deep Reinforcement Learning**
Qiangxing Tian, Guanchu Wang, Jinxin Liu, Donglin Wang
- [JOCN' 23] **Graph-based Conflict-free MAC Protocol and Conflict Analysis for Two-layer Ultraviolet Communication Network**
Yuchen Pan, Guanchu Wang, Yubo Zhang, Jingyin Tang, Chen Gong, Zhengyuan Xu
IEEE Journal of Optical Communications and Networking
- [TCOM' 21] **On the Achievable Rate and Capacity for a Sample-based Practical Photon-counting Receiver**

Zhimeng Jiang, Chen Gong, **Guanchu Wang**, Zhengyuan Xu
IEEE Transaction on Communication

[PJ' 19] **Multi-layer Superimposed Transmission for Optical Wireless Scattering Communication**
Guanchu Wang, Chen Gong, Zhimeng Jiang, Zhengyuan Xu
IEEE Photonics Journal

[TCOM' 18] **Signal Characterization for Multiple Access Non-line of Sight Scattering Communication**
Guanchu Wang, Chen Gong, Zhengyuan Xu
IEEE Transaction on Communication

[PJ' 18] **A 1Mbps Real-time NLOS UV Scattering Communication System with Receiver Diversity over 1km**
Guanchu Wang, Kun Wang, Chen Gong, Difan Zou, Zhimeng Jiang, Zhengyuan Xu
IEEE Photonics Journal

[Globcom' 17] **Signal Detection and Achievable Rates for Multiple Access Optical Wireless Scattering Communication**
Guanchu Wang, Chen Gong, Zhengyuan Xu
IEEE Global Communication Conference

Preprints:

[Arxiv] **Stop Overthinking: A Survey on Efficient Reasoning for Large Language Models**
Yang Sui, Yu-Neng Chuang, **Guanchu Wang**, Jiamu Zhang, Tianyi Zhang, Jiayi Yuan, Hongyi Liu, Andrew Wen, Shaochen Zhong, Hanjie Chen, Xia Hu

[Under review, ICML'25] **Confident or Seek Stronger: Exploring Uncertainty-Based Small LM Routing From Benchmarking to Generalization**
Yu-Neng Chuang, Leisheng Yu, **Guanchu Wang**, Lizhe Zhang, Zirui Liu, Xuanting Cai, Yang Sui, Vladimir Braverman, Xia Hu

[Under review, ARR] **MAIN-RAG: Multi-Agent Filtering Retrieval-Augmented Generation**
Chia-Yuan Chang, Zhimeng Jiang, Vineeth Rakesh, Menghai Pan, Chin-Chia Michael Yeh, **Guanchu Wang**, Mingzhi Hu, Zhichao Xu, Yan Zheng, Mahashweta Das, Na Zou

[Under review, ARR] **A Decoupled Multi-Agent Framework for Complex Text Style Transfer**
LingXi Zhang, Yu-Neng Chuang, **Guanchu Wang**, Ruixiang Tang, Xuanting Cai, Rajesh Shenoy, Xia Hu

[Under review, bioinformatics advances] **Survey and Improvement Strategies for Gene Prioritization with Large Language Models**
Guanchu Wang*, Matthew B. Neeley*, Guantong Qi*, Ruixiang Tang, Dongxue Mao, Chaozhong Liu, Sasidhar Pasupuleti, Bo Yuan, Fan Xia, Hugo Bellen, Pengfei Liu, Zhandong Liu, Xia Hu

[Under review, JBI] **Assessing and Enhancing Large Language Models in Rare Disease Question-answering**
Guanchu Wang, Junhao Ran, Ruixiang Tang, Chia-Yuan Chang, Yu-Neng Chuang, Zirui Liu, Vladimir Braverman, Zhandong Liu, Xia Hu
Under review of Journal of Biomedical Informatics

[Under review, ARR] **FaithLM: Towards Faithful Explanations for Large Language Models**
Yu-Neng Chuang, **Guanchu Wang**, Chia-Yuan Chang, Ruixiang Tang, Shaochen Zhong, Fan Yang, Mengnan Du, Xuanting Cai, Xia Hu

[Under review, TMLR] **Understanding Different Design Choices in Training Large Time Series Models**
Yu-Neng Chuang*, Songchen Li*, Jiayi Yuan*, **Guanchu Wang***, Kwei-Heng Lai, Leisheng Yu, Sirui Ding, Chia-Yuan Chang, Qiaoyu Tan, Daochen Zha, Xia Hu

- [Under review, AAAI] **LEMO: Learning Shapley Manifold from Feature Ablation**
Guanchu Wang, Yu-Neng Chuang, Huiyuan Chen, Yuzhong Chen, Zhimeng Jiang, Zirui Liu, Jiayi Yuan, Xia Hu
- [Under review, KDD] **DISPEL: Domain Generalization via Domain-Specific Liberating**
Chia-Yuan Chang, Yu-Neng Chuang, **Guanchu Wang**, Mengnan Du, Na Zou
- [Under review, TKDD] **Efficient XAI Techniques: A Taxonomic Survey**
Yu-Neng Chuang, **Guanchu Wang**, Fan Yang, Zirui Liu, Xuanting Cai, Mengnan Du, Xia Hu

Research Proposal Experience

ReDDDoT Phase 2: Responsible Multi-Modal AI Systems for Multi-Hazard Resilience and Situational Awareness 2024

PI: Dr. Jamie Padgett; Co-PI: Dr. Xia Hu

Status of Support: Current

Source of Support: National Science Foundation

Contribution: Proposal writing in research thrust: Foundational time-series models for hazard analysis

Collaborative Research: III: Medium: Towards Effective Detection and Mitigation for Shortcut Learning: A Data Modeling Framework 2023

PI: Dr. Xia Hu

Status of Support: Current

Source of Support: National Science Foundation

Contribution: Proposal writing in research objective: Shortcut detection via DNN interpretation

Awards

- Invited Attendee to Future Leader Summit Feb 2024
- Best Demo Honorable Mention in CIKM' 23 Oct 2023
- Ken Kennedy Institute Fellowship Dec 2022
- Best Demo Award in CIKM' 22 Oct 2022
- Spotlight in ICML' 22 Sep 2022
- Doctoral Forum Travel Award in SDM' 24 Mar 2024
- Graduate Fellowship of Rice University Aug 2021

Academic & Industrial Experience

Research Scientist Intern May 2024 - Aug 2024

Mentor: Dr. Yuzhong Chen and Dr. Huiyuan Chen VISA Research, Foster City, CA, USA

Industrial Collaboration with Visa Aug 2023 - present

Mentor: Dr. Huiyuan Chen

Industrial Collaboration with Meta Platforms Aug 2021 - present

Mentor: Dr. Xuanting Cai

Graduate Research Assistant Aug 2021 - present

Advisor: Dr. Xia (Ben) Hu Rice University, Houston, TX, USA

Research Assistant Aug 2019 - Jan 2020

Advisor: Dr. Donglin Wang Westlake University, Hangzhou, Zhejiang, P.R.C.

Teaching & Talks

Introduction to Information Retrieval Teaching Assistant	Spring 2023 Rice University, Houston, TX, USA
Machine Learning with Graph Teaching Assistant	Spring 2022 Rice University, Houston, TX, USA
Introduction of Large Language Models Guest Lecture in Graduate Research Seminar	Sep 2024 Rice University, Houston, TX, USA
Efficient Algorithms of Shapley Values Guest Lecture in Graduate Seminar of CS Department	Mar 2023 Rice University, Houston, TX, USA
Matrix Analysis Teaching Assistant	Fall 2018 University of Science and Technology of China, P.R.C.

Mentorship

James Zhong Ph.D. Student, Rice University	Aug 2024 - now Uncertainty Quantification for Large Language Models
Haochen Zhang Ph.D. Student, Rice University	Aug 2024 - Dec 2024 Efficiency of Large Language Models
Junhao Ran M.S. Student, Rice University	May 2024 - Aug 2024 Large Language Models for Rare Disease Diagnosis
Lingxi Zhang Ph.D. Student, Rice University	May 2024 - Present Large Language Models for Text-style Transfer
Songchen Li M.S. Student, Rice University	Oct 2023 - July 2024 Foundational Models for Time-series Forecasting
Khushbu Pahwa Ph.D. Student, Rice University	Oct 2023 - Jan 2024 Large Language Models for Medical Applications
Yao Rong Ph.D. Student, Technical University of Munich	Aug 2022 - May 2023 Explainable Graph Neural Networks
Chia-yuan Chang Ph.D. Student, Texas A&M University	Aug 2022 - May 2023 Domain Generalization of Deep Neural Networks
Ryan Beckwith Undergraduate Student, Tufts University	June 2022 - Aug 2022 Medical Anomaly Detection
Yu-Neng Chuang Ph.D. Student, Rice University	Mar 2022 - Sep 2022 Real-time Explanations of Deep Neural Networks
Zaid Pervaiz Bhat M.S. Student, Texas A&M University	Sep 2021 - Jan 2022 Real-Time Object Detection on Edge Devices

Academic Services

Program Committee Member: NeurIPS, ICML, KDD, AAAI, SIGIR, CIKM.

Journal Reviewer: ACM TIST, TAI, TCDS; IEEE IS, CL.