

# JINHAO PAN

(979)209-9726 ◊ jpan23@gmu.edu ◊ Fairfax, Virginia, 22030

## EDUCATION

---

**George Mason University, Fairfax**  
Ph.D. in Computer Science

*August 2024 - May 2028 (Expected)*  
Advisor: Dr. Ziwei Zhu

**Texas A&M University, College Station**  
Master of Science (thesis Master degree) in Computer Science

*December 2023 | GPA: 4.0/4.0*  
Advisor: Dr. James Caverlee

**Texas A&M University, College Station**  
Bachelor of Science in Computer Engineering

*May 2021 | GPA: 3.7/4.0*

## RESEARCH INTERESTS

---

I am broadly interested in **Large Language Models (LLMs)**, **data mining**, **machine learning**, and **information retrieval**, specifically in enhancing AI-powered systems' responsibility, fairness, and trustworthiness to serve both users and society.

## PUBLICATIONS

---

**WSDM25** **Jinhao Pan**, James Caverlee, Ziwei Zhu, "Combating Heterogeneous Model Biases in Recommendations via Boosting", The 18th ACM International Conference on Web Search and Data Mining, 2025.

**ECIR24** **Jinhao Pan**, Ziwei Zhu, Jianling Wang, Allen Lin, James Caverlee, "Countering Mainstream Bias via End-to-End Adaptive Local Learning", The 46th European Conference on Information Retrieval, 2024.

## EXPERIENCE

---

**INFO Lab, George Mason University, Fairfax**  
*Research*

*August 2024 - Present*

- Explore social/stereotype biases in Large Language Models (LLMs);
- Construct a more novel and robust social bias benchmark for LLMs;
- Seek various approaches and mechanisms to analyze, explain and interpret biases in LLMs.

**INFO Lab, Texas A&M University, College Station**  
*Student Researcher*

*August 2021 - December 2023*

- Analyzed the various types of bias in the recommender system to achieve the fairness goal;
- Understood and demonstrated the popularity bias and mainstream bias from both the user side and the item side in the recommender system. And alleviated popularity bias and mainstream bias issues by implementing and analyzing novel machine learning and deep learning models;
- Developed an end-to-end fashion model to alleviate mainstream bias and a unified boosting method to mitigate multiple biases from both user and item sides in the recommender system.

## SKILLS & COURSES

---

**Skills:** Python, C/C++, Java, PyTorch, Tensorflow, Scikit-Learn, Numpy, Pandas, AWS, SQL, Golang.

**Courses:** Data Mining, Machine Learning, Deep Learning, Artificial Intelligent, Information Retrieval.

## ACADEMIC ACHIEVEMENTS & SERVICES

---

- **Journal Reviewer:** ACM Transactions on Recommender Systems (TORS), ACM Transactions on Intelligent Systems and Technology (TIST), Transactions on Knowledge and Data Engineering (TKDE)
- **Conference Reviewer:** KDD 2024, ACML 2024