

Zhuoran Lu

Updated January 18, 2025

Email: lu800@purdue.edu

Website: <https://zhuoranlu.github.io/>

Phone: +1 (217)-419-2369

RESEARCH INTERESTS Human-AI Interaction, Human-centered AI, Human-AI Collaboration, Responsible AI, Applied Machine Learning, Computational Social Science

EDUCATION **Purdue University**, West Lafayette, IN Aug 2019 - May 2025 (Expected)
Ph.D. in Computer Science
Advisor: Dr. Ming Yin.
Committee: Dr. Ninghui Li, Dr. Tianyi Zhang, Dr. Sooyeon Jeong.

University of Illinois at Urbana-Champaign, Champaign, IL Jan 2017 - May 2019
B.S. in Computer Science and Statistics; Minor in Mathematics
Graduated with Highest Distinction
Advisor: Dr. Roxana Girju and Dr. Richard Sowers

HONORS & SCHOLARSHIPS Bilsland Dissertation Fellowship, Purdue University 2024
Special Recognition for Outstanding Reviews, ACM CHI 2023
Best Paper Award, ACM CSCW 2022
Frederick N. Andrews Fellowship, Purdue University 2019
Best Project of the Year, Illinois Geometry Lab, UIUC 2018

PUBLICATIONS * indicates equal contribution

Preprints

[P.3] Sheshera Mysore, **Zhuoran Lu**, Mengting Wan, Longqi Yang, Steve Menezes, Tina Baghaee, Emmanuel Barajas Gonzalez, Jennifer Neville, Tara Safavi. Personalizing Large Language Model Writing Assistants for Workplace Social Media via Generation-Calibrated Retrievers. *Preprint*

[P.2] **Zhuoran Lu**, Sheshera Mysore, Tara Safavi, Jennifer Neville, Longqi Yang, Mengting Wan. Augmenting Workplace Social Media Composition via Personalized LLM-powered Writing Assistance. *Preprint*

[P.1] **Zhuoran Lu***, Gionnieve Lim*, Simon Tangi Perrault, Ming Yin. Understanding the Effects of Large Language Model-driven Adversarial Social Influence in Information Spread. *Preprint*

Peer-reviewed Conference Papers

- [C.16] **Zhuoran Lu**, Qian Zhou, Yi Wang. WhatELSE: Shaping Narrative Space at Configurable Level of Abstraction for AI-bridged Interactive Storytelling. In *Proc. of the 43rd ACM Conference on Human Factors in Computing Systems (CHI)*, Yokohama, May 2025.
- [C.15] **Zhuoran Lu**, Patrick Li, Weilong Wang, Ming Yin. The Effects of AI-based Credibility Indicators on the Detection and Spread of Misinformation under Multiple Information Sources. In *Proc. of the 43rd ACM Conference on Human Factors in Computing Systems (CHI)*, Yokohama, May 2025.
- [C.14] Zhuoyan Li, Hangxiao Zhu, **Zhuoran Lu**, Ziang Xiao, Ming Yin. From Text to Trust: Empowering AI-assisted Decision Making with Adaptive LLM-powered Analysis. In *Proc. of the 43rd ACM Conference on Human Factors in Computing Systems (CHI)*, Yokohama, May 2025.
- [C.13] **Zhuoran Lu**, Dakuo Wang, Ming Yin. Does More Advice Help? The Effects of Second Opinions in AI-Assisted Decision Making. In *Proc. of the 27th ACM Conference on Computer-Supported Cooperative Work and Social Computing (CSCW)*, San José, November 2024.
- [C.12] **Zhuoran Lu**, Syed Hasan Amin Mahmood, Zhuoyan Li, Ming Yin. Mix and Match: Characterizing Heterogeneous Human Behavior in AI-assisted Decision Making. In *Proc. of the 12th AAAI Conference on Human Computation and Crowdsourcing (HCOMP)*, Pittsburg, October 2024
- [C.11] **Zhuoran Lu***, Syed Hasan Amin Mahmood*, Ming Yin. Designing Behavior-Aware AI to Optimize Human-AI Team Decision Making. In *Proc. of the 32nd International Joint Conference on Artificial Intelligence (IJCAI)*, Jeju, August 2024
- [C.10] Chun-wei Chiang, **Zhuoran Lu**, Zhuoyan Li, Ming Yin. Enhancing AI-Assisted Group Decision Making through LLM-Powered Devil’s Advocate. In *Proc. of the 29th ACM Conference on Intelligent User Interface (IUI)*, Greenville, March 2024.
- [C.9] Zhuoyan Li, **Zhuoran Lu**, Ming Yin. Decoding AI’s Nudge: A Unified Framework to Predict Human Behavior in AI-assisted Decision Making. In *Proc. of the 38th Conference on Artificial Intelligence (AAAI)*, Vancouver, February 2024.
- [C.8] Zhuoyan Li, Hangxiao Zhu, **Zhuoran Lu**, Ming Yin. Synthetic Data Generation with Large Language Models for Text Classification: Potential and Limitations. In *Proc. of the 2023 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, Singapore, December 2023.
- [C.7] **Zhuoran Lu***, Zhuoyan Li*, Chun-Wei Chiang, Ming Yin. Strategic Adversarial Attacks in AI- assisted Decision Making to Reduce Human Trust and Reliance. In *Proc. of the 32nd International Joint Conference on Artificial Intelligence (IJCAI)*, Macao, August 2023.

[C.6] Chun-Wei Chiang, **Zhuoran Lu**, Zhuoyan Li, Ming Yin. Are Two Heads Better Than One in AI-Assisted Decision Making? Comparing the Behavior and Performance of Groups and Individuals in Human-AI Collaborative Recidivism Risk Assessment. In *Proc. of the 41st ACM Conference on Human Factors in Computing Systems (CHI)*, Hamburg, May 2023.

[C.5] **Zhuoran Lu***, Zhuoyan Li*, Ming Yin. Modeling Human Trust and Reliance in AI-assisted Decision Making: A Markovian Approach. In *Proc. of the 37th Conference on Artificial Intelligence (AAAI)*, D.C., July 2023. **[Oral Presentation]**

[C.4] **Zhuoran Lu**, Patrick Li, Weilong Wang, Ming Yin. The Effects of AI-based Credibility Indicators on the Detection and Spread of Misinformation under Social Influence. In *Proc. of the 25th ACM Conference on Computer-Supported Cooperative Work and Social Computing (CSCW)*, Online, November 2022. **[Best Paper Award]**

[C.3] **Zhuoran Lu***, Zhuoyan Li*, Ming Yin. Towards Better Detection of Biased Language with Scarce, Noisy, and Biased Annotations. In *Proc. of the 5th AAAI-ACM Conference on AI, Ethics, and Society (AIES)*, Oxford, August 2022.

[C.2] **Zhuoran Lu***, Xinru Wang*, Ming Yin. Will You Accept the AI Recommendation? Predicting Human Behavior in AI-Assisted Decision Making. In *Proc. of the ACM Web Conference (WWW)*, Online, April 2022.

[C.1] **Zhuoran Lu**, Ming Yin. Human Reliance on Machine Learning Models When Performance Feedback is Limited: Heuristics and Risks. In *Proc. of the 39th ACM Conference on Human Factors in Computing Systems (CHI)*, Online, May 2021.

Peer-reviewed Journal Articles

[J.1] Shan Gao, **Zhuoran Lu**, Hao Luan, Ming Yin, Lei Wang. AI Pilot in the Cockpit: An Investigation of Public Acceptance. In *Proc. of the International Journal of Human-Computer Interaction (IJHCI)*.

Lightly Reviewed Workshop Papers

[W.3] Sheshera Mysore, **Zhuoran Lu**, Mengting Wan, Longqi Yang, Steve Menezes, Tina Baghaee, Emmanuel Barajas Gonzalez, Jennifer Neville, Tara Safavi. PEARL: Personalizing Large Language Model Writing Assistants with Generation-Calibrated Retrievers. In *the EMNLP Customizable NLP Workshop*, Miami, November 2024.

[W.2] **Zhuoran Lu***, Syed Hasan Amin Mahmood*, Ming Yin. Give Weight to Human Reactions: Optimizing Complementary AI in Practical Human-AI Teams. In *the ICML AI and HCI Workshop*, Hawaii, July 2023.

[W.1] **Zhuoran Lu**, Dakuo Wang, Ming Yin. Does More Advice Help? The Effects of Second Opinions from Peers in AI-Assisted Decision Making. In *the CHI Workshop on Trust and Reliance in AI-Assisted Tasks (TRAIT)*, Hamburg, May 2023.

PATENT

[P.T.1] Yi Wang, Qian Zhou, **Zhuoran Lu**. Authoring Narrative Spaces for AI-bridged Interactive Storytelling. US Patent Application

TEACHING
EXPERIENCE

Graduate Teaching Assistant, Purdue University

CS 373: Data Mining & Machine Learning Fall 2021, Spring 2022

- Designed homework, exams, and slides, and held tutoring sessions.

CS 180: Object-oriented Programming Fall 2019

- Graded homework, and taught in lab sessions.

Guest Lecturer, Purdue University

CS 373: Data Mining & Machine Learning Spring 2022

- Topic: Trustworthy Machine Learning - A Human-centric View

CS 590: Human-AI Interaction Fall 2022, Spring 2024

- Topic: Human Reliance on AI - Internal and External Signals

INDUSTRY
EXPERIENCE

Autodesk Research May 2024 - Aug 2024

Research Intern, From Intent to Experience Group, *Toronto, ON*

Advised by: Dr. Yi Wang & Dr. Qian Zhou

- Large language model (LLM)-driven interactive narrative generation.

Microsoft Research May 2023 - Sep 2023

Research Intern, Office of Applied Research, *Redmond, WA*

Advised by: Dr. Mengting Wan & Dr. Longqi Yang

- Personalizing large language model (LLM)-assisted workplace social media composition.

Honda Research Institute Jan 2023 - Apr 2023

Research Intern, Human-centered Intelligence Group, *San Jose, CA*

Advised by: Dr. Shashank Mehrotra

- Large-scale survey analysis on human trust in intelligent driving assistance.

National Center of Supercomputing Applications Jan 2018 - May 2018

Software Engineer Intern, Advanced Reactors and Fuel Cycles, *Urbana, IL*

Advised by: Dr. Kathryn D. Huff

- Machine learning approaches to accelerate particle transportation simulation.

INVITED
TALKS

Understand and Intervene Human-AI Decision-making

Systems Software & Security Lab, Georgia Tech

Nov 2024

Human Reliance on Machine Learning - Internal and External Signals

Information Management Seminar, Nanjing University Dec 2023

Information Management Young Scholar Forum, Nanjing University Dec 2022

Personalized LLM-empowered Writing Assistant for Workplace Social Media

Microsoft Research Sep 2023

Experiments and User Studies in Information Credibility Research

Information Science Methodology Seminar, Wuhan University Aug 2023

SERVICE

Associate Chair

ACM CHI Late-Breaking Work: 2024, 2025

Program Committee Member

ACM Conference on Fairness, Accountability, and Transparency (FAccT): 2023, 2024, 2025

ACM International Conference on Conversational User Interfaces (CUI): 2024

ACM International Joint Conference on Pervasive and Ubiquitous Computing (Ubi-comp): 2023

Conference Reviewer

ACM Conference on Human Factors in Computing System (CHI): 2023, 2024, 2025

AAAI Conference On Web and Social Media (ICWSM): 2023, 2024

ACM Conference on Computer-Supported Cooperative Work and Social Computing (CSCW): 2023, 2024, 2025

ACM Conference on Designing Interactive Systems (DIS): 2022, 2023, 2024

ACM Conference on Intelligent User Interface (IUI): 2023, 2024

ACM Interaction Design and Children (IDC) Conference: 2023, 2024

Meeting of the Cognitive Science Society (CogSci): 2023, 2024

Computer-Human Interaction of Australia (OzCHI): 2024

ACM Symposium on User Interface Software and Technology (UIST): 2023

Symposium on Computer-Human Interaction in Play (CHI PLAY): 2023

ACM International Conference on Conversational User Interfaces (CUI): 2023

ACM International Conference on Mobile Human-computer Interaction (MobileHCI): 2022

ACM Symposium on Spatial User Interaction (SUI): 2022

International Conference on Human-Agent Interaction (HAI): 2022

Nordic Human-Computer Interaction Conference (NordiCHI): 2022

ACM International Conference on Multimodal Interaction (ICMI): 2022

ACM Interactive Surfaces and Spaces Conference (ISS): 2023

ACM International Joint Conference on Pervasive and Ubiquitous Computing (Ubi-comp): 2022

Journal Reviewer

ACM Transactions on Knowledge Discovery from Data (TKDD)

ACM Transactions on Interactive Intelligent Systems (TiiS)

International Journal of Human-Computer Interaction (IJHCI)
Information Processing and Management (IPM)
Human Factors and Ergonomics in Manufacturing & Service Industries (HFE)
Computational Brain & Behavior

MENTORING
EXPERIENCE

Hangxiao Zhu, M.S. student at Washinton University at St. Louis, worked on [C.8, P.4]
Vivian Yan, B.S. student at Purdue University
Tony Xu, B.S. student at Purdue University
Taylor Gong, B.S. student at University of Illinois at Urbana-Champaign

SKILLS

Programming: Python, R, JavaScript, C, C++, Java
Machine Learning: Pytorch, TensorFlow, Keras, Sklearn, Langchain, Semantic Kernel
Web Development: Meteor.js, Node.js, Flask, React, d3.js, AWS, MongoDB, SQL
HCI: Quantitative and Qualitative methods