Sherry Tongshuang Wu

HCII, School of Computer Science
Carnegie Mellon University
Newell-Simon Hall 3525
Pittsburgh, PA 15213

+1 206-910-2784
sherryw@cs.cmu.edu
http://cs.cmu.edu/~sherryw
@tongshuangwu

ACADEMIC EXPERIENCE

2022– Carnegie Mellon University, Assistant Professor
Human-Computer Interaction Institute (CMU HCII)
Language Technology Institute (CMU LTI)

2016–22 University of Washington, Research Assistant
with Jeffrey Heer, Dan Weld
Pitfalls in *status quo* human-AI interactions.
Principles and tools for enhanced NLP model analysis.
Controllable generators for model analysis and improvement.

EDUCATION

2016–22 Ph.D. in Computer Science and Engineering
University of Washington, Seattle, WA
Thesis: Interactive AI Model Debugging and Correction
Advisor: Jeffrey Heer, Dan Weld
Committee: Marco Tulio Ribeiro, Noah Smith, Mari Ostendorf

2016–18 M.S. in Computer Science and Engineering
University of Washington, Seattle, WA

2012–16 B.Eng. in Computer Science and Engineering
Hong Kong University of Science and Technology, Hong Kong, Hong Kong
Advisor: Huamin Qu

2014 Exchange student in Computer Science and Engineering
University of Michigan, Ann Arbor, MI

INDUSTRY EXPERIENCE

2021 Google Brain/PAIR, Research Intern & Part-time Student Researcher
with Carrie Cai, Michael Terry
Transparent & controllable human-AI collaborations via multi-step problem-solving.

2019 Microsoft Research, Research Intern
with Marco Tulio Ribeiro
Behavioral testing for NLP models covering broad model capabilities.

2018–19 Apple Inc., Full-time Intern & Part-time Intern
with Chris DuBois, Kayur Patel, Kanit Wongsuphasawat, Donghao Ren, Charlie Maalouf
Structural analysis for unstructured text datasets.

2017 Microsoft Research, Research Intern
with Bongshin Lee, Ece Kamar, Saleema Amersh
Uncertainty-aware data labeling and visual refinement.

2015 Microsoft Research Asia, Research Intern
with Weiwei Cui
De-cluttering statistical graphs.

SELECTED HONORS AND AWARDS

2024 Amazon Research Awards

2024 AIED 2024 Best Paper Award

2024 AIED 2024 Honorable Mention Award

2024 AIED 2024 Best Interactive Event Award,

2023 CSCW 2023 Best Demo Award

2023 IUI 2023 Honorable Mention Award

2022 CHI 2022 Honorable Mention Award

2020 Rising Stars in EECS Workshop (UC Berkeley)
A highly selective workshop based on academic excellence and commitment to advancing equity and inclusion.

2020 ACL 2020 Best Paper Award

2016–17 Faithful Steward Endowed Fellowship in Computer Science & Engineering

2012–16 Scholarship Scheme for Continuing Undergraduate Students

2016 IEEE PacificVis 2016 Honorable Mention Award

2016 IEEE PacificVis 2016 Best Notes Paper

PUBLICATIONS

** denotes equal contribution*

Manuscripts and Pre-prints

2024 P.1 Tongshuang Wu, Haiyi Zhu, Maya Albayrak, Alexis Axon, Amanda Bertsch, Wenxing Deng, Ziqi Ding, Bill Guo, Sireesh Gururaja, Tzu-Sheng Kuo, Jenny T Liang, Ryan Liu, Ihita Mandal, Jeremiah Milbauer, Xiaolin Ni, Namrata Padmanabhan, Subhashini Ramkumar, Alexis Sudjianto, Jordan Taylor, Ying-Jui Tseng, Patricia Vaidos, Zhijin Wu, Wei Wu, Chenyang Yang. LLMs as Workers in Human-Computational Algorithms? Replicating Crowdsourcing Pipelines with LLMs. ArXiv 2024

P.2 Qianou Ma, Weirui Peng, Hua Shen, Kenneth Koedinger, Tongshuang Wu. What You Say = What You Want? Teaching Humans to Articulate Requirements for LLMs. ArXiv 2024

2023 P.3 Yujia Qin, Shengding Hu, Yankai Lin, Weize Chen, Ning Ding, Ganqu Cui, Zheni Zeng, Yufei Huang, Chaojun Xiao, Chi Han, Yi Ren Fung, Yusheng Su, Huadong Wang, Cheng Qian, Runchu Tian, Kunlun Zhu, Shihao Liang, Xingyu Shen, Bokai Xu, Zhen Zhang, Yining Ye, Bowen Li, Ziwei Tang5, Jing Yi, Yuzhang Zhu, Zhenning Dai, Lan Yan, Xin Cong, Yaxi Lu, Weilin Zhao, Yuxiang Huang, Junxi Yan, Xu Han, Xian Sun, Dahai Li, Jason Phang, Cheng Yang, Tongshuang Wu, Heng Ji, Zhiyuan Liu, Maosong Sun. Tool Learning with Foundation Models. Computing Surveys (Submitted) 2023

Peer-reviewed Journal Publications

2024 P.4 Shayne Longpre, Robert Mahari, Anthony Chen, Naana Obeng-Marnu, Damien Sileo, William Brannon, Niklas Muennighoff, Nathan Khazam, Jad Kabbara, Kartik Perisetla, Xinyi (Alexis) Wu, Enrico Shippole, Kurt Bollacker, Tongshuang Wu, Luis Villa, Sandy Pentland, Deb Roy, Sara Hooker. A Large Scale Audit of Dataset Licensing and Attribution in AI. Nature Machine Intelligence 2024

P.5 Lindia Tjuatja, Valerie Chen, Tongshuang Wu, Ameet Talwalkar, Graham Neubig. Do LLMs Exhibit Human-Like Response Biases? A Case Study in Survey Design. TACL 2024

2023 P.6 Vijay Viswanathan, Kiril Gashteovski, Carolin Lawrence, Tongshuang Wu, Graham Neubig. Large Language Models Enable Few-Shot Clustering. TACL 2023

P.7 Patrick Fernandes, Aman Madaan, Emmy Liu, António Farinhas, Pedro Henrique Martins, Amanda Bertsch, José G. C. de Souza, Shuyan Zhou, Tongshuang Wu, Graham Neubig, André F. T. Martins. Bridging the Gap: A Survey on Integrating (Human) Feedback for Natural Language Generation. TACL 2023

2022 P.8 Yun Wang, Zhitao Hou, Leixian Shen, Tongshuang Wu, Jiaqi Wang, He Huang, Haidong Zhang, Dongmei Zhang. Towards Natural Language-Based Visualization Authoring. TVCG 2022

2019 P.9 Yang Shi, Maoran Xu, Rongwen Zhao, Hao Fu, Tongshuang Wu, Nan Cao. Interactive Context-Aware Anomaly Detection Guided by User Feedback. THMS 2019

P.10 Tongshuang Wu, Daniel S. Weld, Jeffrey Heer. Local Decision Pitfalls in Interactive Machine Learning: An Investigation into Feature Selection in Sentiment Analysis. TOCHI 2019

2016 P.11 Qiaomu Shen, Tongshuang Wu, Haiyan Yang, Yanhong Wu, Huamin Qu, Weiwei Cui. NameClarifier: A Visual Analytics System for Author Name Disambiguation. TVCG 2016

P.12 Tongshuang Wu, Yingcai Wu, Conglei Shi, Huamin Qu, Weiwei Cui. PieceStack: Toward Better Understanding of Stacked Graphs. TVCG 2016 Honorable Mention

Peer-reviewed Conference Publications

2024 P.13 Xinran Zhao, Hongming Zhang, Xiaoman Pan, Wenlin Yao, Dong Yu, Tongshuang Wu, Jianshu Chen. Fact-and-Reflection (FaR) Improves Confidence Calibration of Large Language Models. ACL Findings 2024

P.14 Saumya Gandhi, Ritu Gala, Vijay Viswanathan, **Tongshuang Wu,** Graham Neubig. Better Synthetic Data by Retrieving and Transforming Existing Datasets. ACL Findings 2024

P.15 Atharva Naik, Jessica Ruhan Yin, Anusha Kamath, Qianou Ma, **Sherry Tongshuang Wu,** Charles Murray, Majd Sakr, Carolyn P. Rose. Generating Situated Reflection Triggers About Alternative Solution Paths: A Case Study in Generative AI for Computer-Supported Collaborative Learning. AIED 2024

P.16 Qiaomu Ma, Hua Shen, Kenneth Koedinger, **Tongshuang Wu**. How to Teach Programming in the AI Era? Using LLMs as a Teachable Agent for Debugging. AIED 2024 Best Paper

P.17 Chenyang Yang, Yining Hong, Grace A. Lewis, **Tongshuang Wu,** Christian Kästner. What Is Wrong with My Model? Identifying Systematic Problems with Semantic Data Slicing. ASE 2024

P.18 Michael Xieyang Liu, **Tongshuang Wu,** Tianying Chen, Franklin Mingzhe Li, Aniket Kittur, Brad A. Myers. Selenite: Scaffolding Online Sensemaking with Comprehensive Overviews Elicited from Large Language Models. CHI 2024

P.19 Tzu-Sheng Kuo, Aaron Halfaker, Zirui Cheng, Jiwoo Kim, Meng-Hsin Wu, **Tongshuang Wu,** Ken Holstein, Haiyi Zhu. Wikibench: Community-Driven Data Curation for AI Evaluation on Wikipedia. CHI 2024

P.20 Xinran Zhao, Tong Chen, Sihao Chen, Hongming Zhang, **Tongshuang Wu**. Beyond Relevance: Evaluate and Improve Retrievers on Perspective Awareness. CoLM 2024

P.21 Chenyang Zhao, Xueying Jia, Vijay Viswanathan, Graham Neubig, **Tongshuang Wu**. Self-Guide: Better Task-Specific Instruction Following via Self-Synthetic Finetuning. CoLM 2024

P.22 Cheng Qian, Xinran Zhao, **Tongshuang Wu**. "Merge Conflicts!" Exploring the Impacts of External Distractors to Parametric Knowledge Graphs. CoLM 2024

P.23 Chenglei Si, Navita Goyal, **Tongshuang Wu,** Chen Zhao, Shi Feng, Hal Daumé III, Jordan Boyd-Graber. Large Language Models Help Humans Verify Truthfulness – Except When They are Convincingly Wrong. NAACL 2024

2023 P.24 Vijay Viswanathan, Luyu Gao, **Tongshuang Wu,** Pengfei Liu, Graham Neubig. DataFinder: Scientific Dataset Recommendation from Natural Language Descriptions. ACL 2023

P.25 Logan Stapleton, Jordan Taylor, Sarah Fox, **Tongshuang Wu,** Haiyi Zhu. Seeing Seeds Beyond Weeds: Green Teaming Generative AI for Beneficial Uses. ArXiv 2023

P.26 Yiming Zhang, Sravani Nanduri, Liwei Jiang, **Tongshuang Wu,** Maarten Sap. BiasX: "Thinking Slow" in Toxic Content Moderation with Explanations of Implied Social Biases. EMNLP 2023

P.27 Vijay Viswanathan, Chenyang Zhao, Amanda Bertsch, **Tongshuang Wu,** Graham Neubig. Promp2Model: Generating Deployable Models from Natural Language Instructions. EMNLP Demo Track 2023

P.28 Jeremiah Milbauer, Ziqi Ding, Zhijin Wu, **Tongshuang Wu**. From Nuisance to News Sense: Augmenting the News with Cross-document Evidence and Context. EMNLP Demo Track 2023

P.29 Chenyang Yang, Rishabh Rustogi, Rachel Brower-Sinning, Grace Lewis, Christian Kaestner, **Tongshuang Wu**. Beyond Testers' Biases: Guiding Model Testing with Knowledge Bases using LLMs. EMNLP Findings 2023

P.30 **Tongshuang Wu,** Hua Shen, Jeffrey Heer, Daniel S. Weld, Marco Tulio Ribeiro. ScatterShot: Interactive In-context Example Curation for Text Transformation. IUI 2023 Honorable Mention

P.31 Kaustubh D Dhole, Varun Gangal, Sebastian Gehrmann, Aadesh Gupta, Zhenhao Li, Saad Mahamood, Abinaya Mahendiran, Simon Mille, Ashish Srivastava, Samson Tan, **Tongshuang Wu,** Jascha Sohl-Dickstein, Jinho D Choi, Eduard Hovy, Ondrej Dusek, Sebastian Ruder, et al.. NL-Augmenter: A Framework for Task-Sensitive Natural Language Augmentation. NEJLT 2023

P.32 Hyeonsu Kang, **Tongshuang Wu,** Joseph Chee Chang, Aniket Kittur. Synergi: A Mixed-Initiative System for Scholarly Synthesis and Sensemaking. UIST 2023

2022 P.33 Ying Xu, Dakuo Wang, Mo Yu, Daniel Ritchie, Bingsheng Yao, **Tongshuang Wu,** Zheng Zheng, Toby Jia-Jun Li, Nora Bradford, Branda Sun, Tran Bao Hoang, Yisi Sang, Yufang Hou, Xiaojuan Ma, Diyi Yang, Nanyun Peng, Zhou Yu, Mark Warschauer. Fantastic Questions and Where to Find Them: FairytaleQA -- An Authentic Dataset for Narrative Comprehension. ACL 2022

P.34 **Tongshuang Wu*,** Alexis Ross*, Hao Peng, Matthew E. Peters, Matt Gardner. Tailor: Generating and Perturbing Text with Semantic Controls. ACL 2022

P.35 Bingsheng Yao, Dakuo Wang, **Tongshuang Wu,** Toby Jia-Jun Li, Mo Yu, Ying Xu. It is AI's Turn to Ask Humans a Question: Question and Answer Pair Generation for Children Storybooks with FairytaleQA Dataset. ACL 2022

P.36 Hua Shen, **Tongshuang Wu,** Wenbo Guo, Ting-Hao 'Kenneth' Huang. Are Shortest Rationales the Best Explanations for Human Understanding?. ACL 2022

P.37 **Tongshuang Wu,** Michael Terry, Carrie J. Cai. AI Chains: Transparent and Controllable Human-AI Interaction by Chaining Large Language Model Prompts. CHI 2022

P.38 Zheng Zhang, Ying Xu, Bingsheng Yao, Daniel Ritchie, **Tongshuang Wu,** Mo Yu, Dakuo Wang, Toby Jia-Jun Li. StoryBuddy: A Human-AI Collaborative Agent for Parent-Child Interactive Storytelling with Flexible Parent

Involvement. CHI 2022

P.39 Jiao Sun, Tongshuang Wu, Yue Jiang, Ronil Awalegaonkar, Xi Victoria Lin, Diyi Yang. Pretty Princess vs. Successful Leader: Gender Roles in Greeting Card Messages. CHI 2022 Honorable Mention

P.40 Tongshuang Wu*, Ellen Jiang*, Aaron Donsbach, Jeff Gray, Alejandra Molina, Michael Terry, Carrie J. Cai. PromptChainer: Chaining Large Language Model Promptsthrough Visual Programming. CHI LBW 2022

2021 P.41 Tongshuang Wu, Marco Tulio Ribeiro, Jeffrey Heer, Daniel S. Weld. Polyjuice: Generating Counterfactuals for Explaining, Evaluating, and Improving Models. ACL 2021

P.42 Tongshuang Wu*, Gagan Bansal*, Joyce Zhou+, Raymond Fok+, Besmira Nushi, Ece Kamar, Marco Tulio Ribeiro, Daniel S. Weld. Does the Whole Exceed its Parts? The Effect of AI Explanations on Complementary Team Performance . CHI 2021

P.43 Xingbo Wang, Yao Ming, Tongshuang Wu, Haipeng Zeng, Yong Wang, Huamin Qu. DeHumor: Visual Analytics forDecomposing Humor. TVCG 2021

2020 P.44 Marco Tulio Ribeiro, Tongshuang Wu, Carlos Guestrin, Sameer Singh. Beyond Accuracy: Behavioral Testing of NLP Models with CheckList. ACL 2020 Best Paper

P.45 Tongshuang Wu, Kanit (Ham) Wongsuphasawat, Donghao Ren, Kayur Patel, Chris DuBois. Tempura: Query Analysis with Structural Templates. CHI 2020

P.46 Alison Smith-Renner, Ron Fan, Melissa Birchfield, Tongshuang Wu, Jordan Boyd-Graber, Daniel S. Weld, Leah Findlater. No Explainability without Accountability: An Empirical Study of Explanations and Feedback in Interactive ML. CHI 2020

P.47 Tongshuang Wu*, Zhihang Dong*, Sicheng Song, Mingrui Zhang. Interactive Attention Model Explorer for Natural Language Processing Tasks with Unbalanced Data Sizes. PacificVis 2020

2019 P.48 Tongshuang Wu, Marco Tulio Ribeiro, Jeffrey Heer, Daniel S. Weld. Errudite: Scalable, Reproducible, and Testable Error Analysis. ACL 2019

2016 P.49 Yun Wang, Tongshuang Wu, Zhutian Chen, Huamin Qu, Qiong Luo. STAC: Enhancing Stacked Graphs for Time Series Analysis. PacificVis 2016

P.50 Tongshuang Wu, Yuan Yao, Yuqing Duan, Xinzhi Fan, Huamin Qu. NetworkSeer: Visual Analysis for Social Network in MOOCs. PacificVis 2016 Best Paper

Posters, Extended Abstracts, Workshop Papers and Technical Reports

2023 W.1 Chenyang Yang, Rachel Brower-Sinning, Grace A. Lewis, Christian Kästner, Tongshuang Wu. Capabilities for Better ML Engineering. AAAI SafeAI 2023

W.2 Yuanchen Bai, Raoyi Huang, Vijay Viswanathan, Tzu-Sheng Kuo, Tongshuang Wu. Measuring Adversarial Datasets. AACL ART of Safety 2023

W.3 Qianou Christina Ma, Tongshuang Wu, Kenneth Koedinger. Is AI the Better Programming Partner? Human-Human Pair Programming vs. Human-AI pAIr Programming. AIED2023 Empowering Education with LLMs 2023

W.4 Hua Shen, Tongshuang Wu. Parachute: Evaluating Interactive Human-LM Co-writing Systems. CHI In2Writing 2023

W.5 Hua Shen, Chieh-Yang Huang, Tongshuang Wu, Ting-Hao (Kenneth) Huang. ConvXAI: Delivering Heterogeneous AI Explanations via Conversations to Support Human-AI Scientific Writing. CSCW Demo Track 2023

2022 W.6 Zheng Zheng, Ying Xu, Yanhao Wang, Tongshuang Wu, Bingsheng Yao, Daniel Ritchie, Mo Yu, Dakuo Wang, Toby Jia-Jun Li. Building a Storytelling conversational Agent through Parent-AI Collaboration. AAAI AI4ED 2022

2021 W.7 Tongshuang Wu. Principles and Interactive Tools for Evaluating and Improvingthe Behavior of Natural Language Processing models. CHI DC 2021

2018 W.8 Halden Lin, Tongshuang Wu, Kanit (Ham) Wongsuphasawat, Yejin Choi, Jeffrey Heer. Visualizing Attention in Sequence-to-Sequence Summarization Models. VAST 2018

Patent

2023 PT.1 Carrie Cai, Tongshuang Wu, Michael Terry. Transparent and Controllable Human-AI Interaction via Chaining of Machine-Learned Language Models. US Patent US 2023/0112921 A1 2023

2022 PT.2 Ajit Narayanan, Subhashini Venugopalan, Tongshuang Wu, Shanqing Cai, Michael Terry, Meredith Morris, Carrie Cai. Providing Suggestions of Expanded Text from Abbreviated Text Input. (Defensive Publication) 2022

TEACHING EXPERIENCE

Instructor

2024 17-445/645/745 Machine Learning in Production (CMU)

2023-24 05-391/891 Designing Human Centered Software (CMU)

2023 05-499/899 Human-Centered NLP (CMU)

Guest Lecture

2024 Interacting with Large Language Models (Carnegie Mellon University)
Human Interactions with Code Gen Models (Carnegie Mellon University)

2023 Human-Centerered AI (University of South California)

2022 Visualization and Machine Learning (Carnegie Mellon University)
Interacting with Large Language Models (Carnegie Mellon University)
Visualizing Text Summarization Models (Carnegie Mellon University)
Designing Human-Centered, AI-Infused Software (Carnegie Mellon University)

2021 HCI+AI Interaction (University of Notre Dame)

2019 Model Interpretability (University of Washington)

Conference Tutorial

2024 EMNLP 2024: Human-AI Interaction in the Age of LLMs Models

2023 EMNLP 2023: Designing, Evaluating, and Learning from Humans Interacting with NLP Models

Teaching Assistant

2019 CSE 512 Data Visualization (University of Washingon)

2018 CSE 442 Data Visualization (University of Washingon)

MENTORING EXPERIENCE

Advisees

PhD Vijay Viswanathan (CMU LTI, *co-advisor: Graham Neubig*). *Democratization of NLP Development and Evaluation*
Christina Ma (CMU HCII, *co-advisor: Ken Koedinger*). *Preparing Students for Effective Human-LLM Partnerships*
Chenyang Yang (CMU S3D, *co-advisor: Christian Kästner*). *Human-Centered ML Engineering*
Xinran Zhao (CMU LTI). *Information Seeking and Retrieval for Complex Tasks*
Jessie Mindel (CMU HCII). *Simulated Agents and Collective Sensemaking*

Master Yiyang (Diana) Wang (CMU HCII). *End-User Prompt Disambiguation. Now PhD student at Georgia Tech.*
Yuanchen (Sophie) Bai (CMU Heinz). *NLP dataset characterization*
Raoyi (Cathy) Huang (CMU Heinz). *NLP dataset characterization. Now PhD student at Cornell.*
Atharva Naik (CMU LTI). *LLM in CS education. Now PhD student at CMU.*
Jushaan Kalra (CMU MIIS). *Multi-domain Retrieval*
Yilin Zhang (CMU MIIS). *Code Retrieval with AST*

Visit Cheng Qian (Tsinghua University). *LLM hullucination. Now PhD student at UIUC.*

Undergrad Alex Cheung (CMU IS). *LLM sensemaking copilot*
Samriddhi Bhardwaj (CMU CS). *LLM sensemaking copilot*
Alina Chen *LLM sensemaking copilot*
Yashika Batra (CMU CS). *LLM sensemaking copilot*
Shaan Lehal (CMU CS). *LLM sensemaking copilot*
Cassandra Shi (CMU CS). *Requirement-driven LLMs*

Thesis Committee

PhD Steven Moore (CMU). *Creating and Evaluating Pedagogically Valid Assessments at Scale*
Yoonjoo Lee (KAIST). *Aligning AI Agents with How Humans Understand Knowledge*
Hyeonsu Kang (CMU). *Accelerating Innovation through AI-Powered Conceptual Abstraction and Interaction Design*
Kundan Krishna (CMU). *Improving the reliability of summarization models*
Hua Shen (Penn State). *Towards Useful AI Interpretability via Interactive AI Explanations*
Jason Wu (CMU). *Computational Understanding of User Interfaces*

Master Shreya Bali (CMU). *Tools to facilitate working on Machine Learning in the Industry*
Ihita Mandal (CMU). *Accessible Descriptions for Surprising Clusters in Scatterplots*

Prior to CMU

PhD Yi Guo (Tongji University). *Co-supervised with Nan Cao. Natural-language-based visualization generation.*
Sebastin Santy (UW). *The design and creation of an HCI+NLP research playbook.*
Jiao Sun (USC). *Co-supervised with Diyi Yang. Gender bias in NLP datasets.*

Master Joyce Zhou (UW; Now at Cornell). *Co-supervised with Dan Weld & Gagan Bansal. Human-AI teaming.*
Halden Lin (UW; Now at Apple Inc.). *Attention visualization for NLP models.*
Akshat Shrivastava (UW; Now at Meta). *Active learning for sequence labeling.*

PROFESSIONAL SERVICE

Organizing Committees

2024 TREW: Workshop on Trust and Reliance in Evolving Human-AI Workflows (CHI 2024)

2024 Tutorial: Human-AI Interaction in the Age of LLMs (NAACL 2024)

2023 Tutorial: Designing, Evaluating, and Learning from Humans Interacting with NLP Models (EMNLP 2023)

2022 SSLL: Sharing Stories and Lessons Learned Workshop (EMNLP 2022)

2022 TRAIT: Workshop on Trust and Reliance in AI-Human Teams (CHI 2022-23)

2022 NL-Augmenter (part of GEM: Workshop for Generation, Evaluation, Metrics, ACL 2021)

Program Committees

AI AAAI 2022, AAAI HCOMP 2022-23, ACM FAccT 2022, NeurIPS XAI 2021

NLP ACL 2023, EMNLP 2023, COLM 2024

HCI CHI 2023-24, ACM IUI 2022-23, IUI TExSS 2022, CHI HCXAI 2021

Paper Reviewing

HCI ACM CHI 2019-22, TOCHI 2021, UIST 2018/20/22, IUI 2020, CSCW 2020, TiiS 2022
Special recognition for outstanding reviews ACM CHI, IUI

NLP ACL 2020, EACL 2021, NAACL 2021

AI NeurIPS 2022, AAAI 2022, AKBC 2021, ACM Computing Surveys 2021

Viz. IEEE VIS 2017-23, TVCG 2021, EuroVis 2021, PacificVis 2018/20, ChinaVis 2017-19

Community Service

2024-25 Committee member, CMU Faculty Senate
Represented the HCII department in the CMU Faculty Senate.

2024 Reviewer, Department of Energy Office Proposal Panel

2023-24 Leader, Postdoc Mentoring Group
Led a bi-weekly mentorship group for Postdocs within PhD HCII.

Committee member, CMU HCII PhD Admission Committee

Committee member, CMU HCII Undergraduate Admission Committee

2023 Committee member, CMU SCS Doctoral Dissertation Award
Selected CMU SCS level internal dissertation awardees, as well as candidates for AAAI and ACM SIGAI thesis award.

Reviewer, NSF Proposal Panel

Committee member, CMU K&L Gates Award Selection Committee
Selected awardees who have inspired their fellow students to love learning through a combination of intellect, high scholarly achievement, engagement with others, and character.

2022-23 Committee member, ACM/SCS Thesis Nomination & Award Committee
Read various candidate PhD dissertations and selected awardees for both CMU and ACM level thesis award.

2021 Co-organizer, UW Allen School Women's Research Day
An outreach event to women and nonbinary people in research.

Co-organizer, UW Allen School Pre-Application Mentoring Service (PAMS)
A program supporting 107 potential CS PhD applicants, with 80% from underrepresented communities.

Coordinator, UW Allen School Diverse Genders in Research Events

Course design mentor, UW AVELA (A Vision for Electronic Literacy & Access)
Mentored undergraduate students to develop curriculum for high-school web development courses.

2020 Student volunteer, IEEE VIS 2020

Student contributor, UW Allen School Strategic Plan for Diversity, Equity & Inclusion

Subcomittee Student Assistant, ACM SIGCHI 2021

2019 Reviewer, UW Allen School Graduate Admission Committee

2018 Co-leader, UW Interactive Systems Seminar

2013 Community tutor, HKUST Connect

MEDIA COVERAGE

2023 MIT, Cohere for AI, others launch platform to track and filter audited AI datasets
VentureBeat, 2023.1

AI Researchers Uncover Ethical, Legal Risks to Using Popular Data Sets
The Washington Post, 2023.1

15% of datasets for fine-tuning language models use Wikipedia
Wikimedia, 2023.11

CMU & Tsinghua U's Prompt2Model Generates Deployable Models Following Natural Language Instructions
Synced Review, 2023.8

Debugging Imperfect AI
CMU User Experience Association (UXA) Newsletter, 2023.2

Bigger isn't always better when it comes to large language models
Axios, 2023.12

Researchers from CMU and Tsinghua University Propose Prompt2Model: A General Purpose Method that Generates Deployable AI Models from Natural Language Instructions
MarkTechPost, 2023.8

2020 AI researchers create testing tool to find bugs in NLP from Amazon, Google, and Microsoft
VentureBeat, 2020.7

Allen School researchers earn Best Paper Award at ACL 2020
Allen School News, 2020.8

How Should We Do Error Analysis? A Lesson for NLP developers [Chinese]
AI Technology Review, 2020.8

2019 Experimental results and error reporting, Ethics and NLP, Distillation
Distillation vol. 2,, SemEval 2020