

# 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.

## PUBLICATIONS

\* denotes equal contribution

## Manuscripts and Pre-prints

- 2021 P.1 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**. In *arXiv:2112.02721*.  
arXiv 2021

## Peer-reviewed Conference and Journal Publications

- 2023 P.2 **Tongshuang Wu**, Hua Shen, Jeffrey Heer, Daniel S. Weld, Marco Tulio Ribeiro. **ScatterShot: Interactive In-context Example Curation for Text Transformation**. In *ACM Conference on Intelligent User Interfaces*.  
IUI 2023
- 2022 P.3 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**. In *the 60th Annual Meeting of the Association for Computational Linguistics*.  
ACL 2022
- P.4 **Tongshuang Wu\***, Alexis Ross\*, Hao Peng, Matthew E. Peters, Matt Gardner. **Tailor: Generating and Perturbing Text with Semantic Controls**. In *the 60th Annual Meeting of the Association for Computational Linguistics*.  
ACL 2022
- P.5 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**. In *the 60th Annual Meeting of the Association for Computational Linguistics*.  
ACL 2022
- P.6 Hua Shen, **Tongshuang Wu**, Wenbo Guo, Ting-Hao 'Kenneth' Huang. **Are Shortest Rationales the Best Explanations for Human Understanding?**. In *the 60th Annual Meeting of the Association for Computational Linguistics*.  
ACL 2022
- P.7 Jiao Sun, **Tongshuang Wu**, Yue Jiang, Ronil Awalegaonkar, Xi Victoria Lin, Diyi Yang. **Pretty Princess vs. Successful Leader: Gender Roles in Greeting Card Messages**. In *the 2022 Conference on Human Factors in Computing Systems*.  
CHI 2022 🏆 Honorable Mention
- P.8 **Tongshuang Wu**, Michael Terry, Carrie J. Cai. **AI Chains: Transparent and Controllable Human-AI Interaction by Chaining Large Language Model Prompts**. In *the 2022 Conference on Human Factors in Computing Systems*.  
CHI 2022
- P.9 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**. In *the 2022 Conference on Human Factors in Computing Systems*.  
CHI 2022
- P.10 Yun Wang, Zhitao Hou, Leixian Shen, **Tongshuang Wu**, Jiaqi Wang, He Huang, Haidong Zhang, Dongmei Zhang. **Towards Natural Language-Based Visualization Authoring**. In *IEEE Transactions on Visualization and Computer Graphics (Proc. InfoVis 2022)*.  
TVCG 2022
- 2021 P.11 **Tongshuang Wu**, Marco Tulio Ribeiro, Jeffrey Heer, Daniel S. Weld. **Polyjuice: Generating Counterfactuals for Explaining, Evaluating, and Improving Models**. In *the 59th Annual Meeting of the Association for Computational Linguistics*.  
ACL 2021
- P.12 **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**. In *the 2021 Conference on Human Factors in Computing Systems*.  
CHI 2021
- P.13 Xingbo Wang, Yao Ming, **Tongshuang Wu**, Haipeng Zeng, Yong Wang, Huamin Qu. **DeHumor: Visual Analytics for Decomposing Humor**. In *IEEE Transactions on Visualization and Computer Graphics*.  
TVCG 2021

- 2020 P.14 Marco Tulio Ribeiro, **Tongshuang Wu**, Carlos Guestrin, Sameer Singh. **Beyond Accuracy: Behavioral Testing of NLP Models with CheckList**. In the 58th Annual Meeting of the Association for Computational Linguistics. ACL 2020 🏆 Best Paper
- P.15 **Tongshuang Wu**, Kanit (Ham) Wongsuphasawat, Donghao Ren, Kayur Patel, Chris DuBois. **Tempura: Query Analysis with Structural Templates**. In the 2020 Conference on Human Factors in Computing Systems. CHI 2020
- P.16 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**. In the 2020 Conference on Human Factors in Computing Systems. CHI 2020
- P.17 **Tongshuang Wu\***, Zhihang Dong\*, Sicheng Song, Mingrui Zhang. **Interactive Attention Model Explorer for Natural Language Processing Tasks with Unbalanced Data Sizes**. In IEEE Pacific Visualization Symposium 2020. PacificVis 2020
- 2019 P.18 **Tongshuang Wu**, Marco Tulio Ribeiro, Jeffrey Heer, Daniel S. Weld. **Errudite: Scalable, Reproducible, and Testable Error Analysis**. In the 57th Annual Meeting of the Association for Computational Linguistics. ACL 2019
- P.19 Yang Shi, Maoran Xu, Rongwen Zhao, Hao Fu, **Tongshuang Wu**, Nan Cao. **Interactive Context-Aware Anomaly Detection Guided by User Feedback**. In IEEE Transactions on Human-Machine Systems. THMS 2019
- P.20 **Tongshuang Wu**, Daniel S. Weld, Jeffrey Heer. **Local Decision Pitfalls in Interactive Machine Learning: An Investigation into Feature Selection in Sentiment Analysis**. In ACM Transactions on Computer-Human Interaction 26.4. TOCHI 2019
- 2016 P.21 Yun Wang, **Tongshuang Wu**, Zhutian Chen, Huamin Qu, Qiong Luo. **STAC: Enhancing Stacked Graphs for Time Series Analysis**. In IEEE Pacific Visualization Symposium 2016. PacificVis 2016
- P.22 **Tongshuang Wu**, Yuan Yao, Yuqing Duan, Xinzhi Fan, Huamin Qu. **NetworkSeer: Visual Analysis for Social Network in MOOCs**. In IEEE Pacific Visualization Symposium 2016. PacificVis 2016 🏆 Best Paper
- P.23 **Tongshuang Wu**, Yingcai Wu, Conglei Shi, Huamin Qu, Weiwei Cui. **PieceStack: Toward Better Understanding of Stacked Graphs**. In IEEE Transactions on Visualization and Computer Graphics (Proc. PacificVis 2016). TVCG 2016 🏆 Honorable Mention
- P.24 Qiaomu Shen, **Tongshuang Wu**, Haiyan Yang, Yanhong Wu, Huamin Qu, Weiwei Cui. **NameClarifier: A Visual Analytics System for Author Name Disambiguation**. In IEEE Transactions on Visualization and Computer Graphics (Proc. VAST 2016). TVCG 2016

## Posters, Extended Abstracts, Workshop Papers and Technical Reports

- 2022 W.1 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**. In AAAI Artificial Intelligence for Education Workshop. AAAI 2022
- W.2 Chenyang Yang, Rachel Brower-Sinning, Grace A. Lewis, Christian Kästner, **Tongshuang Wu**. **Capabilities for Better ML Engineering**. In The AAAI's Workshop on Artificial Intelligence Safety. AAAI SafeAI 2023 2022
- W.3 **Tongshuang Wu\***, Ellen Jiang\*, Aaron Donsbach, Jeff Gray, Alejandra Molina, Michael Terry, Carrie J. Cai. **PromptChainer: Chaining Large Language Model Prompts through Visual Programming**. In the 2022 Conference on Human Factors in Computing Systems (Late Breaking Work). CHI LBW 2022
- 2021 W.4 **Tongshuang Wu**. **Principles and Interactive Tools for Evaluating and Improving the Behavior of Natural Language Processing models**. In CHI 2022 Doctoral Consortium. CHI DC 2021

- 2018 W.5 Halden Lin, Tongshuang Wu, Kanit (Ham) Wongsuphasawat, Yejin Choi, Jeffrey Heer. **Visualizing Attention in Sequence-to-Sequence Summarization Models**. In *IEEE Conference on Visual Analytics Science and Technology. VAST 2018*

## Patent

- 2022 PT.1 Ajit Narayanan, Subhashini Venugopalan, Tongshuang Wu, Shanqing Cai, Michael Terry, Meredith Morris, Carrie Cai. **Providing Suggestions of Expanded Text from Abbreviated Text Input**.  
US Patent (Defensive Publication) 2022
- PT.2 Carrie Cai, Tongshuang Wu, Michael Terry. **Transparent and Controllable Human-AI Interaction via Chaining of Machine-Learned Language Models**.  
US Patent US63/251,264 (submitted), 2021 2022

## SELECTED HONORS AND AWARDS

- 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

## TEACHING EXPERIENCE

### Instructor

- 2023 05-499/899 **Human-Centered NLP** (CMU)  
2023 Spring. Designed this special topic course on HCI+NLP research.

### Guest Lecture

- 2022 **Visualization and Machine Learning** (Carnegie Mellon University)  
*In 05-898 Data Visualization*
- Interacting with Large Language Models** (Carnegie Mellon University)  
*In 05-318/618 Human-AI Interaction*
- Designing Human-Centered, AI-Infused Software** (Carnegie Mellon University)  
*In 05-391/891: Designing Human-Centered Software*
- Visualizing Text Summarization Models** (Carnegie Mellon University)  
*In 05-120: Introduction to Human-Computer Interaction*
- 2021 **HCI+AI Interaction** (University of Notre Dame)  
*In CSE 60427: Human-Centered Computing Research*
- 2019 **Model Interpretability** (University of Washington)  
*In CSE 512: Data Visualization*

### Teaching Assistant

- 2019 CSE 512 **Data Visualization** (University of Washington)  
*Graduate level course, 112 students (Instructor: Jeffrey Heer)*
- 2018 CSE 442 **Data Visualization** (University of Washington)  
*Undergraduate level course, 109 students (Instructor: Jeffrey Heer)*

## MENTORING EXPERIENCE

### Advisees

- PhD*      **Christina Ma** (CMU HCII, co-advisor: Ken Koedinger)  
 Human-AI Interaction for Education  
**Vijay Viswanathan** (CMU LTI, co-advisor: Graham Neubig)  
 Human-in-the-loop NLP Model Fixing
- Master*    **Yiyang (Diana) Wang** (CMU HCII)  
 End-User Prompt Disambiguation  
**Yuanchen (Sophie) Bai** (CMU Heinz)  
 NLP dataset characterization  
**Raoyi (Cathy) Huang** (CMU Heinz)  
 NLP dataset characterization

## Thesis Committee

- PhD*      **Kundan Krishna** (CMU)  
 Improving the reliability of summarization models  
**Hua Shen** (Penn State)  
 Towards Useful AI Interpretability via Interactive AI Explanations
- Master*    **Shreya Bali** (CMU)  
 Tools to facilitate working on Machine Learning in the Industry

## 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.

## INVITED TALKS

- 2022      **Advanced NLP: Prompting**  
 Peking University, Wangxuan Institute of Computer Technology (2022.3)
- Model-in-the-loop Data Collection: What Roles does the Model Play?**  
 DADC: Dynamic Adversarial Data Collection, NAACL 2022 (2022.7)
- Interactive AI Model Debugging and Correction**  
 Carnegie Mellon University, Ameet Group Meeting (2022.9)  
 Carnegie Mellon University, LTI Colloquium (2022.9)  
 University of Washington, DUB Seminar (2022.7)  
 University of Washington, Paul G. Allen School of CSE (2022.6)  
 MIT Computer Science & Artificial Intelligence Laboratory (2022.3)  
 Stanford University, Computer Science Department (2022.3)  
 Princeton University, Computer Science Department (2022.3)  
 Cornell University, Computer Science Department (2022.3)  
 Carnegie Mellon University, Human-Computer Interaction Institute (2022.3)  
 Peking University, Wangxuan Institute of Computer Technology (2022.3)  
 UT Austin, Computer Science Department (2022.2)  
 University of Chicago, Data Science Institute (2022.2)  
 Hong kong University of Science and Technology, Computer Science Department (2022.2)  
 PennState, Young Achievers Symposium (2022.1)

- 2021 **Principles and Tools for Evaluating and Improving NLP Models**  
 Google Research (2021.7)  
 ACM CHI Doctoral Consortium (2021.5)  
 NLP with Friends (2021.4)  
 Interactive Data Exploration System (IDEAS) lab, Shandong University (2021.4)
- Transparent and Controllable Collaboration with Large Language Models**  
 Google Film Sprint: Fluid Language Integrating Muse (2021.1)  
 Google PAIR: People+AI Research Initiative (2021.7)
- Human+AI: the Relationships, the Goals, the Challenges**  
 University of Notre Dame, Human-Centered Computing Research (2021.1)  
 Hong Kong University of Science and Technology, VisLab (2021.1)
- Machines in the Loop: Explainability, Transparency, and Rich Interaction**  
 ACL InterNLP 2021, invited panelist (2021.8)
- Generating and Perturbing Text with Semantic Controls**  
 Allen Institute for Artificial Intelligence, All-AI2 Meeting (2021.8)
- 2020 **Behavioral Testing of NLP Models**  
 AI Technology Review (2020.8)  
 AI Time PhD (2020.7)  
 UCLA, Center for Vision, Cognition, Learning and Autonomy (VCLA) (2020.7)
- 2019 **Scalable, Reproducible, and Testable Error Analysis**  
 UW CSE 512, as guest lecturer (2019.5)  
 Allen Institute for Artificial Intelligence, All-AI2 Meeting (2019.5)  
 Apple Inc., Knowledge Graph Team Seminar (2019.3)  
 Robust AI Event: Research & Reality (2019.5)

## PROFESSIONAL SERVICE

### Organizing Committees

- 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 AAI 2022, AAI HCOMP 2022, ACM FAccT 2022, NeurIPS XAI 2021  
 NLP ACL 2023  
 HCI CHI 2023, 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, EAACL 2021, NAACL 2021  
 AI NeurIPS 2022, AAI 2022, AKBC 2021, ACM Computing Surveys 2021  
 Viz. IEEE VIS 2017-21, TVCG 2021, EuroVis 2021, PacificVis 2018/20, ChinaVis 2017-19

### Community Service

- 2021 **Course design mentor**, UW AVELA (A Vision for Electronic Literacy & Access)  
 Mentored undergraduate students to develop curriculum for high-school web development courses
- Co-organizer**, UW Allen School Pre-Application Mentoring Service (PAMS)  
 A program supporting 107 potential CS PhD applicants, with 80% from underrepresented communities.

- Co-organizer**, *UW Allen School Women's Research Day*  
*An outreach event to women and nonbinary people in research.*
- Coordinator**, *UW Allen School Diverse Genders in Research Events*
- 2020 **Student contributor**, *UW Allen School Strategic Plan for Diversity, Equity & Inclusion*  
**Subcommittee Student Assistant**, *ACM SIGCHI 2021*  
**Student volunteer**, *IEEE VIS 2020*
- 2019 **Reviewer**, *UW Allen School Graduate Admission Committee*
- 2018 **Co-leader**, *UW Interactive Systems Seminar*
- 2013 **Community tutor**, *HKUST Connect*

## **MEDIA COVERAGE**

- 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*