

Wei Xu

✉ wei.xu@cc.gatech.edu 🌐 <https://cocoxu.github.io/> 📍 CODA Tech Square, Atlanta, GA, 30308 U.S.A.
✉ @cocoweixu 📄 <https://scholar.google.com/citations?user=Bf0dG-oAAAAJ>

RESEARCH	Natural Language Processing, Machine Learning, Social Media	
CITIZENSHIP	United States	
APPOINTMENT	Associate Professor, Georgia Institute of Technology, Atlanta, GA	Aug 2023 – Present
	<i>College of Computing, School of Interactive Computing, Machine Learning Center</i>	
	Assistant Professor, Georgia Institute of Technology, Atlanta, GA	Aug 2020 – July 2023
	<i>College of Computing, School of Interactive Computing, Machine Learning Center</i>	
EDUCATION	Ph.D. in Computer Science, New York University, New York, NY	2014
	Advisor: Ralph Grishman; Committee: Satoshi Sekine, Ernest Davis, Bill Dolan (Microsoft Research), Luke Zettlemoyer (University of Washington / Meta)	
	B.S./M.S. in Computer Science, Tsinghua University, Beijing, CHINA	2004/2007
SELECTED AWARDS	Sony Faculty Innovation Award, 2026	
	Google Academic Research Award, 2024	
	Best Social Impact Paper Award, ACL, 2024	
	Sony Faculty Innovation Award, co-PI, 2024	
	Best Paper Award Honorable Mention, ACL, 2023	
	NSF CAREER Award, 2022	
	NSF CRII Award, 2018	
	Best Paper Award, COLING, 2018	
	Criteo Faculty Research Award, 2018	
	CrowdFlower AI for Everyone Award, 2018	
PROFESSIONAL SERVICES	NYU MacCracken PhD Fellowship, 2007 – 2012	
	Executive Board Member: NAACL (2023-2024); Best Paper Award Committee: ACL (2026), EMNLP (2024, 2018); Senior Area Chair: ACL (2026, 2025), EMNLP (2026, 2024, 2022), NAACL (2025, 2024, 2022, 2021), ACL (2020); Area Chair: COLM (2024), ACL (2023, 2019), EMNLP (2021, 2020, 2018, 2016), AAAI (2020), NAACL (2019), COLING (2018); Action Editor: ACL Rolling Review (2021-2022); Program Committee: ACL (2021, 2018, 2017, 2015, 2014, 2013), EMNLP (2017, 2015, 2014), NAACL (2015), WWW (2017, 2016, 2015), AAAI (2016, 2015, 2012), KDD (2015), COLING (2014); Publicity Chair: EMNLP (2019), NAACL (2018, 2016); Workshop Chair: ACL (2017); Journal Reviewer: Transactions of the Association for Computational (TACL), Journal of Artificial Intelligence Research (JAIR); Organizer: Workshop on Human-centered Evaluation and Auditing of Language Models at CHI (2026, 2025, 2024), Workshop on Text Simplification, Accessibility, and Readability at EMNLP (2022); Workshop on Natural Language Generation, Evaluation, and Metrics at ACL (2021); Workshop on Noisy User-generated Text at NAACL (2025), EACL (2023), EMNLP (2026, 2021, 2020, 2019, 2018, 2017), COLING (2022, 2016), ACL (2015); Mid-Atlantic Student Colloquium on Speech, Language and Learning (2016).	
RESEARCH GRANTS	NSF Grant <i>Robust and Responsible Reasoning for Privacy-Preserving Image Geolocation</i>	2026 – 2029
	co-PI, total \$997,565.00	
	Sony <i>Multi-Agent Reinforcement Learning for Culture, Knowledge, and Context Adaptive Translation</i>	2026
	PI, total \$100,000	
	Google <i>Improving Accessibility of Legal Documents with AI to Support Civil Rights Litigation</i>	2025
	PI, total \$75,000	
	Sony <i>Adapting LLM-based Multimodal Dialog Across Languages and Cultures</i>	2025
	co-PI, total \$100,000	
	NIH R01 <i>Building the Safety Guards for the Automatic Simplification of Medical Documents</i>	2024 – 2028
	MPI, total \$1,374,201	
	NSF Grant <i>The Cost of AI: A Comparative Study of Machine Learning Training Methods</i>	2024 – 2028
	co-PI, total \$399,634	
	NSF CAREER <i>An Integrated Framework for Controllable Text Generation</i>	2022 – 2027
	sole PI, total \$537,527	

NSF AI Institute *Collaborative Assistance and Responsive Interaction for Networked Groups* 2021 – 2026
Senior Personnel, total \$19,995,808

IARPA Grant *Human Interpretable Attribution of Text Using Underlying Structure* 2022 – 2024
co-PI, total \$491,467

IARPA Grant *Better Extraction from Text Towards Enhanced Retrieval* 2019 – 2023
co-PI, total \$850,000

NSF Grant *Collaborative Research: Automatic Text-Simplification and Reading-Assistance to Support Self-Directed Learning by Deaf and Hard-of-Hearing Computing Workers* 2018 – 2022
PI, total \$375,732

NSF CRII RI: *Learning a Timely Semantic Resource from Social Media Data* 2018 – 2021
sole PI, total \$183,000

DARPA Grant *Modeling the Spread of Information through Social and Knowledge Graphs* 2017 – 2021
co-PI, total \$600,000

PUBLICATIONS (Underline is used to indicate student advisees.)

Learning to Route Languages for Multilingual Preference Optimization

Geyang Guo, Hiromi Wakaki, Yuki Mitsufuji, Alan Ritter, Wei Xu

ICML 2026

Flipping the Dialogue: Training and Evaluating User Language Models

Tarek Naous, Philippe Laban, Wei Xu, Jennifer Neville

ICLR 2026

Do Vision-Language Models Respect Contextual Integrity in Location Disclosure?

Ruixin Yang, Ethan Mendes, Arthur Wang, James Hays, Sauvik Das, Wei Xu, Alan Ritter

ICLR 2026

Supporting Informed Self-Disclosure: Design Principles for Presenting AI-Estimates of Privacy Risks to Users

Isadora Krsek, Meryl Ye, Wei Xu, Alan Ritter, Laura Dabbish, and Sauvik Das

CHI 2026

GeoRC: A Benchmark for Geolocation Reasoning Chains

Mohit Talreja, Joshua Diao, Jim Thannikary James, Radu Casapu, Tejas Santanam, Ethan Mendes, Alan Ritter, Wei Xu, James Hays

ACL 2026

Faithfulness vs. Safety: Evaluating LLM Behavior Under Counterfactual Medical Evidence

Kaijie Mo, Siddhartha Venkatayogi, Chantal Shaib, Ramez Kouzy, Wei Xu, Byron C. Wallace, Junyi Li

ACL Findings 2026

Probabilistic Reasoning with LLMs for Privacy Risk Estimation

Jonathan Zheng, Sauvik Das, Alan Ritter, Wei Xu

NeurIPS 2025, main track (acceptance rate 24.52%)

CARE: Assessing the Impact of Multilingual Human Preference Learning on Cultural Awareness

Geyang Guo, Tarek Naous, Hiromi Wakaki, Yukiko Nishimura, Yuki Mitsufuji, Alan Ritter, Wei Xu

EMNLP 2025 (acceptance rate 22.16%)

USimBench: Are User Simulators Reliable Proxies for Multi-Turn Evaluation of AI Assistants?

Yao Dou, Michel Galley, Baolin Peng, Chris Kedzie, Weixin Cai, Alan Ritter, Chris Quirk, Wei Xu, Jianfeng Gao

EMNLP 2025 (acceptance rate 22.16%)

What are Foundation Models Cooking in the Post-Soviet World?

Anton Lavrouk, Tarek Naous, Alan Ritter, Wei Xu

EMNLP 2025 (acceptance rate 22.16%)

How to Protect Yourself from 5G Radiation? Investigating LLM Responses to Implicit Misinformation

Ruohao Guo, Wei Xu, Alan Ritter

EMNLP 2025 (acceptance rate 22.16%)

Evaluating LLMs on Chinese Idiom Translation

Cai Yang, Yao Dou, David Heineman, Xiaofeng Wu, Wei Xu

COLM 2025

Beyond the Reported Cutoff: Where Large Language Models Fall Short on Financial Knowledge

Agam Shah, Liqin Ye, Sebastian Jaskowski, Wei Xu, Sudheer Chava

COLM 2025

On The Origin of Cultural Biases in Language Models: From Pre-training Data to Linguistic Phenomena
Tarek Naous, Karl Stratos, Wei Xu
 NAACL 2025, long paper (acceptance rate 22%)

The Impact of Visual Information in Chinese Characters
Xiaofeng Wu, Karl Stratos, Wei Xu
 NAACL 2025, long paper (acceptance rate 22%)

CROSSNEWS: A Cross-Genre Authorship Verification and Attribution Benchmarks
Marcus Ma, Duong Minh Le, Junmo Kang, Yao Dou, John Cadigan, Dayne Freitag, Alan Ritter, Wei Xu
 AAAI 2025 (acceptance rate 23.4%)

Measuring, Modeling, and Helping People Account for Privacy Risks in Online Self-Disclosures with AI
 Isadora Krsek, Anubha Kabra, Yao Dou, Tarek Naous, Laura A. Dabbish, Alan Ritter, Wei Xu, Sauvik Das
 CSCW 2025

Computer-aided Design Code Generation and Automatic Refinement with Large Language Models and Geometric Solvers
 Kamel Alrashedy, Pradyumna Tambwekar, Zulfiqar Haider Zaidi, Megan Langwasser, Wei Xu, Matthew Gombolay
 ICLR 2025 (acceptance rate 32.08%)

Improving Minimum Bayes Risk Decoding with Multi-Prompt
 David Heineman, Yao Dou, Wei Xu
 EMNLP 2024, long paper (acceptance rate 20.8%)

ReadMe++: Benchmarking Multilingual Language Models for Multi-Domain Readability Assessment
 Tarek Naous, Michael J. Ryan, Anton Lavrouk, Mohit Chandra, Wei Xu
 EMNLP 2024, long paper (acceptance rate 20.8%)

MedReadMe: A Systematic Study for Fine-grained Sentence Readability in Medical Domain
 Chao Jiang, Wei Xu
 EMNLP 2024, long paper (acceptance rate 20.8%)

Granular Privacy Control for Geolocation with Vision Language Models
 Ethan Mendes, Yang Chen, James Hays, Sauvik Das, Wei Xu, Alan Ritter
 EMNLP 2024, long paper (acceptance rate 20.8%)

ChatHF: Collecting Rich Human Feedback from Real-time Conversations
 Andrew Li, Zhenduo Wang, Ethan Mendes, Duong Minh Le, Wei Xu, Alan Ritter
 EMNLP 2024, short paper (acceptance rate 33.77%)

GPT-4 Jailbreaks Itself with Near-Perfect Success Using Self-Explanation
Govind Ramesh, Yao Dou, Wei Xu
 EMNLP 2024, system demonstration (acceptance rate 20.8%)

Having Beer after Prayer? Measuring Cultural Bias in Large Language Models
Tarek Naous, Michael J. Ryan, Alan Ritter, Wei Xu
 ACL 2024, long paper (acceptance rate 21.3%)
Best Social Impact Award, selection rate $3/4407 = 0.07\%$; **Press Coverage by VentureBeat**

Reducing Privacy Risks in Online Self-Disclosures with Language Models
Yao Dou, Isadora Krsek, Tarek Naous, Anubha Kabra, Sauvik Das, Alan Ritter, Wei Xu
 ACL 2024, long paper (acceptance rate 21.3%)

NEO-BENCH: Evaluating Robustness of Large Language Models with Neologisms
Jonathan Zheng, Alan Ritter, Wei Xu
 ACL 2024, long paper (acceptance rate 21.3%)

Meta-Tuning LLMs to Leverage Lexical Knowledge for Generalizable Language Style Understanding
 Ruohao Guo, Wei Xu, Alan Ritter
 ACL 2024, long paper (acceptance rate 21.3%)

FactPICO: Factuality Evaluation for Plain Language Summarization of Medical Evidence
 Sebastian Antony Joseph, Lily Chen, Jan Trienes, Hannah Louisa Göke, Monika Coers, Wei Xu, Byron C Wallace, Junyi Jessy Li
 ACL 2024, long paper (acceptance rate 21.3%)

InfoLossQA: Characterizing and Recovering Information Loss in Text Simplification.
 Jan Trienes, Sebastian Joseph, Jörg Schlötterer, Christin Seifert, Kyle Lo, Wei Xu, Byron Wallace, Jessy Li
 ACL 2024, long paper (acceptance rate 21.3%)

Constrained Decoding for Cross-lingual Label Projection

Duong Minh Le, Yang Chen, Alan Ritter, Wei Xu
ICLR 2024 (acceptance rate 30.8%)

Design and Evaluation of an Automatic Text Simplification Prototype with Deaf and Hard-of-hearing Readers
Oliver Alonzo, Sooyeon Lee, Akhter Al Amin, Mounica Maddela, Wei Xu, Matt Huenerfauth
ASSETS 2024 (acceptance rate 30.0%)

Dancing Between Success and Failure: Edit-level Simplification Evaluation using SALSA
David Heineman, Yao Dou, Mounica Maddela, Wei Xu
EMNLP 2023, long paper (acceptance rate 23.3%)

Multilingual Simplification of Medical Texts
Sebastian Joseph, Kathryn Kazanas, Keziah Reina, Vishnesh Ramanathan, Wei Xu, Byron Wallace, Junyi Jessy Li
EMNLP 2023, long paper (acceptance rate 23.3%)

Thresh: A Unified, Customizable and Deployable Platform for Fine-Grained Text Evaluation
David Heineman, Yao Dou, Wei Xu
EMNLP 2023, system demonstration (acceptance rate 25%)

A Computational Interface to Translate Strategic Intent from Unstructured Language in a Low-Data Setting
Pradyumna Tambwekar, Lakshita Dodeja, Nathan Vaska, Wei Xu, Matthew Gombolay
Findings of EMNLP 2023, long paper (acceptance rate 46.2%)

LENS: A Learnable Evaluation Metric for Text Simplification
Mounica Maddela, Yao Dou, David Heineman, Wei Xu
ACL 2023, long paper (acceptance rate 23.5%)

Distill or Annotate? Cost-Efficient Fine-Tuning of Compact Models
Junmo Kang, Wei Xu, Alan Ritter
ACL 2023, long paper (acceptance rate 23.5%)

Improved Instruction Ordering in Recipe-Grounded Conversation
Duong Minh Le, Ruohao Guo, Wei Xu, Alan Ritter
ACL 2023, long paper (acceptance rate 23.5%)

Revisiting non-English Text Simplification: A Unified Multilingual Benchmark
Michael Ryan, Tarek Naous, Wei Xu
ACL 2023, long paper (acceptance rate 23.5%)

Best Paper Award Honorable Mention

Frustratingly Easy Label Projection for Cross-lingual Transfer
Yang Chen, Chao Jiang, Alan Ritter, Wei Xu
Findings of ACL 2023, long paper

Improving Large-scale Paraphrase Acquisition and Generation
Yao Dou, Chao Jiang, Wei Xu
EMNLP 2022, long paper (acceptance rate 22.1%)

arXivEdits: Understanding the Human Revision Process in Scientific Writing
Chao Jiang, Wei Xu, Sam Stevens
EMNLP 2022, long paper (acceptance rate 22.1%)

Extracting a Knowledge Base of COVID-19 Events from Social Media
Shi Zong, Ashutosh Baheti, Wei Xu, Alan Ritter
COLING 2022, long paper (acceptance rate 31.2%)

BiSECT: Learning to Split and Rephrase Sentences with Bitexts
Joongwon Kim*, Mounica Maddela*, Reno Kriz, Wei Xu, Chris Callison-Burch (*equal contribution)
EMNLP 2021, long paper (acceptance rate 23.3%)

Pre-train or Annotate? Domain Adaptation with a Constrained Budget
Fan Bai, Alan Ritter, Wei Xu
EMNLP 2021 (acceptance rate 23.3%)

Neural semi-Markov CRF for Monolingual Word Alignment
Wuwei Lan*, Chao Jiang*, Wei Xu (*equal contribution)
ACL 2021, long paper (acceptance rate 21.2%)

Controllable Text Simplification with Explicit Paraphrasing
Mounica Maddela, Fernando Alva-Manchego, Wei Xu
NAACL 2021, long paper (acceptance rate 28%)

The GEM Benchmark: Natural Language Generation, its Evaluation and Metrics

Sebastian Gehrmann, Tosin Adewumi, Karmanya Aggarwal, Pawan Sasanka Ammanamanchi, Anuoluwapo Aremu, Antoine Bosselut, Khyathi Raghavi Chandu, Miruna-Adriana Clinciu, Dipanjan Das, Kaustubh Dhole, Wanyu Du, Esin Durmus, Ondřej Dušek, Chris Chinenye Emezue, Varun Gangal, Cristina Garbacea, Tatsunori Hashimoto, Yufang Hou, Yacine Jernite, Harsh Jhamtani, Yangfeng Ji, Shailza Jolly, Mihir Kale, Dhruv Kumar, Faisal Ladhak, Aman Madaan, Mounica Maddela, Khyati Mahajan, Saad Mahamood, Bodhisattwa Prasad Majumder, Pedro Henrique Martins, Angelina McMillan-Major, Simon Mille, Emiel van Miltenburg, Moin Nadeem, Shashi Narayan, Vitaly Nikolaev, Andre Niyongabo Rubungo, Salomey Osei, Ankur Parikh, Laura Perez-Beltrachini, Niranjan Ramesh Rao, Vikas Raunak, Juan Diego Rodriguez, Sashank Santhanam, João Sedoc, Thibault Sellam, Samira Shaikh, Anastasia Shimorina, Marco Antonio Sobrevilla Cabezudo, Hendrik Strobelt, Nishant Subramani, Wei Xu, Diyi Yang, Akhila Yerukola, Jiawei Zhou (alphabetically ordered except for the first author)

ACL 2021 Workshop on Natural Language Generation, Evaluation, and Metrics

An Empirical Study of Pre-trained Transformers for Arabic Information Extraction

Wuwei Lan, Yang Chen, Wei Xu, Alan Ritter

EMNLP 2020, short paper (acceptance rate 16.7%)

WNUT-2020 Task 1 Overview: Extracting Entities and Relations from Wet Lab Protocols

Jeniya Tabassum, Sydney Lee, Wei Xu, Alan Ritter

EMNLP 2020 Workshop on Noisy User-generated Text (shared-task overview)

Neural CRF Model for Sentence Alignment in Text Simplification

Chao Jiang, Mounica Maddela, Wuwei Lan, Yang Zhong, Wei Xu

ACL 2020, long paper (acceptance rate 25.2%)

An Empirical Study of Named Entity Recognition in StackOverflow

Jeniya Tabassum, Mounica Maddela, Wei Xu, Alan Ritter

ACL 2020, long paper (acceptance rate 25.2%)

Generalizing Natural Language Analysis through Span-relation Representations

Zhengbao Jiang, Wei Xu, Jun Araki, Graham Neubig

ACL 2020, long paper (acceptance rate 25.2%)

Learning Relation Entailment with Structured and Textual Information

Zhengbao Jiang, Jun Araki, Donghan Yu, Ruohong Zhang, Wei Xu, Yiming Yang, Graham Neubig

AKBC 2020, long paper

Discourse Level Factors for Sentence Deletion in Text Simplification

Yang Zhong, Chao Jiang, Wei Xu, Junyi Jessy Li

AAAI 2020, long paper (acceptance rate 20.6%; oral presentation)

Multi-task Pairwise Neural Ranking for Hashtag Segmentation

Mounica Maddela, Wei Xu, Daniel Preotiuc-Pietro

ACL 2019, long paper (acceptance rate 25.7%)

A Word-Complexity Lexicon and A Neural Readability Ranking Model for Lexical Simplification

Mounica Maddela, Wei Xu

EMNLP 2018, long paper (acceptance rate 25.8%; oral presentation)

Neural Network Models for Paraphrase Identification, Semantic Textual Similarity, Natural Language Inference, and Question Answering

Wuwei Lan, Wei Xu

COLING 2018, long paper (**Best Paper Award**; selection rate $8/888 = 0.90\%$)

An Annotated Corpus for Machine Reading of Instructions in Wet Lab Protocols

Chaitanya Kulkarni, Wei Xu, Alan Ritter, Raghu Machiraju

NAACL 2018, short paper (acceptance rate 29%)

Character-based Neural Networks for Sentence Pair Modeling

Wuwei Lan, Wei Xu

NAACL 2018, short paper (acceptance rate 29%)

A Continuously Growing Dataset of Sentential Paraphrases

Wuwei Lan, Siyu Qiu, Hua He, Wei Xu

EMNLP 2017, long paper (acceptance rate 25.8%)

From Shakespeare to Twitter: What are Language Styles all about?

Wei Xu

EMNLP 2017 Workshop on Stylistic Variation

A Minimally Supervised Method for Recognizing and Normalizing Time Expressions in Twitter

Jeniya Tabassum, Alan Ritter, Wei Xu

EMNLP 2016, long paper (acceptance rate 26%; oral presentation)
Optimizing Statistical Machine Translation for Simplification
 Wei Xu, Courtney Napoles, Ellie Pavlick, Quanze Chen, Chris Callison-Burch
 TACL 2016, long paper (oral presentation at ACL 2016)
Discovering User Attribute Stylistic Differences via Paraphrasing
 Daniel Preotiuc-Pietro, Wei Xu, Lyle Ungar
 AAAI 2016, long paper (acceptance rate 26%; oral presentation)
Results of the WNUT16 Named Entity Recognition Shared Task
 Benjamin Strauss, Bethany Toma, Alan Ritter, Marie-Catherine de Marneffe, Wei Xu
 COLING 2016 Workshop on Noisy User-generated Text (shared-task overview)
Problems in Current Text Simplification Research: New Data Can Help
 Wei Xu, Chris Callison-Burch, Courtney Napoles
 TACL 2015, long paper (oral presentation at EMNLP 2015)
Cost Optimization for Crowdsourcing Translation
 Mingkun Gao, Wei Xu, Chris Callison-Burch
 NAACL 2015, long paper (acceptance rate 29%)
SemEval-2015 Task 1: Paraphrase and Semantic Similarity in Twitter
 Wei Xu, Chris Callison-Burch, William B. Dolan
 SemEval 2015, long paper (shared-task overview)
Shared Tasks of the 2015 Workshop on Noisy User-generated Text: Twitter Lexical Normalization and Named Entity Recognition
 Timothy Baldwin, Marie Catherine de Marneffe, Bo Han, Young-Bum Kim, Alan Ritter, Wei Xu
 ACL 2015 Workshop on Noisy User-generated Text (shared-task overview; author ordered alphabetically)
Data-driven Approaches for Paraphrasing Across Language Variations
 Wei Xu
 Ph.D. Thesis 2014
Extracting Lexically Divergent Paraphrases from Twitter
 Wei Xu, Alan Ritter, Chris Callison-Burch, William B. Dolan, Yangfeng Ji
 TACL 2014, long paper (oral presentation at NAACL 2015)
Infusion of Labeled Data into Distant Supervision for Relation Extraction
 Maria Pershina, Bonan Min, Wei Xu, Ralph Grishman
 ACL 2014, short paper (acceptance rate 25.2%; oral presentation)
Filling Knowledge Base Gaps for Distant Supervision of Relation Extraction
 Wei Xu, Raphael Hoffmann, Le Zhao, Ralph Grishman
 ACL 2013, short paper (acceptance rate 24%)
Gathering and Generating Paraphrases from Twitter with Application to Normalization
 Wei Xu, Alan Ritter, Ralph Grishman
 ACL 2013 Workshop on Building and Using Comparable Corpora
A Preliminary Study of Tweet Summarization using Information Extraction
 Wei Xu, Ralph Grishman, Adam Meyers, Alan Ritter
 NAACL 2013 Workshop on Language Analysis in Social Media
Paraphrasing for Style
 Wei Xu, Alan Ritter, Bill Dolan, Ralph Grishman, Colin Cherry
 COLING 2012, long paper (acceptance rate 25%)
Exploiting Syntactic and Distributional Information for Spelling Correction with Web-Scale N-grams Models
 Wei Xu, Joel Tetreault, Martin Chodorow, Ralph Grishman, Le Zhao
 EMNLP 2011, long paper (acceptance rate 23.7%)
New York University 2011 System for KBP (Knowledge Base Population) Slot Filing
 Ang Sun, Ralph Grishman, Wei Xu, Bonan Min
 TAC 2011 (best performance system in NIST KBP-2011 evaluation)
Passage Retrieval for Information Extraction using Distant Supervision
 Wei Xu, Ralph Grishman, Le Zhao
 IJCNLP 2011, long paper (acceptance rate 36%)
Who, What, When, Where, Why? Comparing Multiple Approaches to the Cross-Lingual 5W Task
 Kristen Parton, Kathleen McKeown, Bob Coyne, Mona Diab, Ralph Grishman, Dilek Hakkani-Tür, Mary Harper, Heng Ji, Weiyun Ma, Adam Meyers, Sara Stolbach, Ang Sun, Gokhan Tur, Wei Xu, Sibel Yaman

	ACL 2009, long paper (acceptance rate 21%; oral presentation)	
	<i>A Parse-and-Trim Approach with Information Significance for Chinese Sentence Compression</i>	
	Wei Xu, Ralph Grishman	
	ACL Workshop on Language Generation and Summarisation 2009	
	<i>Transducing Logical Relations from Automatic and Manual Annotation</i>	
	Adam Meyers, Michiko Kosaka, Heng Ji, Nianwen Xue, Mary Harper, Ang Sun, Wei Xu, Shasha Liao	
	ACL Workshop on Linguistic Annotation 2009	
	<i>Automatic Recognition of Logical Relations for English, Chinese and Japanese in the GLARF Framework</i>	
	Adam Meyers, Michiko Kosaka, Nianwen Xue, Heng Ji, Ang Sun, Shasha Liao, Wei Xu	
	SemEval 2009, long paper	
	<i>Extractive Summarization using Inter- and Intra- Event Relevance</i>	
	Wenjie Li, Wei Xu, Mingli Wu, Chunfa Yuan, Qin Lu	
	ACL 2006, long paper (acceptance rate 23%; oral presentation)	
TUTORIALS	Automatic and Human-AI Interactive Text Generation	
	ACL 2024 (https://acl2024-text-generation-tutorial.github.io/)	August 2024
INVITED TALKS	Beyond Logic and Math: Probabilistic Reasoning for Real-World Decision Making	
	Oracle Distinguished Lecture Series (virtual)	May 2026
	NeurIPS 2025 Workshop on Efficient Reasoning	December 2025
	Probabilistic Reasoning and Multicultural Alignment in LLMs	
	Apple ML Research (virtual)	August 2025
	Reasoning, Privacy, and Cultural Sensitivity: Advancing LLMs for Real-World Use	
	JPMorgan AI Research (virtual)	July 2025
	Empowering Everyday Users to Protect Their Privacy in the Age of AI	
	Keynote at NAACL 2025 PrivateNLP Workshop	May 2025
	Multilingual and Multicultural Capabilities of Large Language Models	
	Sungkyunkwan University, South Korea	May 2025
	Google, Mountain View, CA	Feb 2025
	Cultural Bias and Privacy Protection in Large Language Models	
	University of Pennsylvania (CLUNCH Seminar), Philadelphia, PA	Mar 2025
	Improving Multilingual Capabilities and Error Detection in Large Language Models	
	Bloomberg (CTO Data Science Speaker series), New York, NY	Oct 2024
	Cultural Biases, World Languages, and User Privacy in Large Language Models	
	Stony Brook University, New York, NY	Oct 2024
	MIT (Embodied Intelligence Seminar), Boston, MA	Sep 2024
	Human-AI Collaboration in Evaluating Large Language Models	
	University of Massachusetts Lowell (CS Colloquium), Lowell, MA	Feb 2025
	University of California Berkeley (NLP Seminar), Berkeley, CA	Feb 2025
	Northeastern University, Boston, MA	Sep 2024
	HEAL Workshop (Keynote) at CHI 2024	May 2024
	Cultural Biases and Multilingual Capabilities of Large Language Models	
	Megagon Labs (virtual)	Aug 2024
	Amazing Multilingual Capabilities and Concerning Cultural Biases in Large Language Models	
	University of California, Los Angeles	Mar 2024
	University of South California	Mar 2024
	Google Research, India (virtual)	Mar 2024
	Amplifying LLM's Cross-lingual Ability with Label Projection	
	Brainlink, South Korea	Dec 2023
	GPT-3 vs Humans: Rethinking Evaluation of Natural Language Generation	
	Johns Hopkins University (CLIP Seminar), Baltimore, MD	Feb 2023
	Capturing Human Language Diversity and (Mis-)Information Spreading Online	
	VinAI, Vietnam (virtual)	Apr 2023
	University of Chicago (virtual)	Feb 2023
	Columbia University, New York, NY	Oct 2022
	Importance of Data and Controllability in Neural Language Generation	
	Cornell Tech (LMSS Seminar), New York, NY	Sep 2022

Dataminr, New York, NY	Nov 2021
Nanjing University, Nanjing, China	Oct 2021
SimpleText workshop at CLEF 2021	Sep 2021
Stanford University (NLP Seminar), Stanford, CA	Aug 2021
University of California, Los Angeles (Big Data and ML Seminar)	Jun 2021
Importance of Data and Linguistics in Neural Language Generation	
New York University, New York, NY (NLP and Text-as-Data Speaker Series)	May 2021
Carnegie Mellon University, Pittsburgh, PA (LTI Colloquium)	Nov 2020
Google Research, New York (virtual)	Oct 2020
Natural Language Understanding for Noisy Text	
University of Sheffield, Sheffield, United Kingdom (NLP Seminar)	Oct 2020
USC Information Sciences Institute, Los Angeles, CA (NLP Seminar)	Oct 2020
Automatic Text Simplification	
University of Pittsburgh, Pittsburgh, PA (NLP Seminar)	Oct 2020
Understanding and Generating Human Language	
Emory University, Atlanta, GA (CS Department Seminar)	Sep 2020
University of Maryland, College Park, MD (CS Colloquium)	Feb 2020
University of Massachusetts, Amherst, MA	Jan 2020
Georgia Institute of Technology, Atlanta, GA	Dec 2019
Learning for Unlimited Human Language	
Peking University, Beijing, China	Dec 2018
Learning Large-scale Paraphrases for Natural Language Understanding and Generation	
Midwest Machine Learning Symposium, Chicago, IL	Jun 2018
Facebook, Menlo Park, CA	May 2018
Stanford Research Institute, Menlo Park, CA	May 2018
Twitter, San Francisco, CA	May 2018
IBM Thomas J. Watson Research Center, New York, NY	Nov 2017
How does AI Understand Language?	
Women in Analytics Conference, Columbus, OH (Main Stage Panel)	Mar 2018
Can Paraphrase be a Ultimate Solution for NLU and NLG?	
Google Research, New York, NY	Jul 2017
Paraphrase $\approx$ Monolingual Translation	
Amazon, Berlin, Germany	Aug 2016
Multiple Instance Learning from Unlimited Text	
Microsoft Research Asia, Beijing, China	Dec 2016
University of Delaware, Newark, DE	Sep 2016
University of Edinburgh, Edinburgh, United Kingdom	May 2016
Ohio State University, Columbus, OH	Apr 2016
University of North Carolina, Chapel Hill, NC	Apr 2016
Arizona State University, Tempe, AZ	Mar 2016
Vanderbilt University, Nashville, TN	Mar 2016
Imperial College London, London, United Kingdom	Mar 2016
University of Waterloo, Waterloo, ON, Canada (CS Seminar)	Mar 2016
Indiana University, Bloomington, IN (Computer Science Colloquium Series)	Feb 2016
Washington University, St Louis, MI (Computer Science & Engineering Colloquia Series)	Feb 2016
Simon Fraser University, Vancouver, BC, Canada	Feb 2016
University of Alberta, Edmonton, AB, Canada	Feb 2016
Yale University, New Haven, CT (CS Talk)	Feb 2016
University of Maryland, College Park, MD (CLIP Colloquium)	Oct 2015
Ohio State University, Columbus, OH (Clippers Seminar)	Oct 2015
Large-scale Paraphrase Acquisition from Twitter	
DARPA's DEFT Project Meeting, Boulder, CO	May 2015
Learning and Generating Paraphrases from Twitter and Beyond	
Carnegie Mellon University, Pittsburgh, PA	Apr 2015
Columbia University, New York, NY (NLP Talk)	Apr 2015
Johns Hopkins University, Baltimore, MD (CLSP Colloquium)	Feb 2015
Paraphrases in Twitter	

	Twitter, San Francisco, CA	Feb 2015
	Modeling Lexically Divergent Paraphrases in Twitter (and Shakespeare!)	
	The City University of New York, New York, NY (NLP Seminar)	Mar 2015
	IBM Research - Almaden, San Jose, CA	Feb 2015
	University of California, Berkeley, CA	Feb 2015
	The University of Texas, Austin, TX (Forum for Artificial Intelligence)	Feb 2015
	Yahoo!, New York, NY	Dec 2014
	Carnegie Mellon University, Pittsburgh, PA (CL+NLP Lunch Seminar)	Nov 2014
	Microsoft Research, Seattle, WA (Visiting Speaker Series)	Aug 2014
	Incremental Information Extraction	
	Stanford Research Institute, Palo Alto, CA	Apr 2012
	IARPA's KDD Project Meeting, San Diego, CA	May 2011
	Event-based Summarization	
	Thomson Reuters, Eagan, Minnesota, MN	Nov 2009
TEACHING	<i>CS 8803-LLM Large Language Models (graduate level)</i>	
	Teaching eval: 4.9/5.0 Fall 2024	
	<i>CS 7650 Natural Language Processing (graduate level)</i>	
	Teaching eval: 4.3/5.0 Fall 2025, 4.7/5.0 Spring 2024, 4.3/5.0 Fall 2022, 4.3/5.0 Fall 2021	
	<i>CS 8803 Advanced Natural Language Processing (graduate level)</i>	
	Teaching eval: 4.9/5.0 Fall 2023	
	<i>CS 4650 Natural Language Processing (undergraduate level)</i>	
	Teaching eval: 4.4/5.0 Spring 2025, 4.7/5.0 Spring 2023, 4.6/5.0 Spring 2022, 3.7/5.0 Spring 2021	
	<i>CSE 5539 Social Media and Text Analytics</i>	
	Teaching eval: 4.13/5.00 Fall 2019, 4.40/5.00 Fall 2017, 4.60/5.00 Fall 2016; 5.72/6.00 at NASSLLI 2015	
	<i>CSE 5522 Artificial Intelligence II: Advanced Techniques (mixed undergraduate and graduate level)</i>	
	Teaching eval: 4.85/5.00 Fall 2018, 4.50/5.00 Spring 2018	
	<i>CSE 5525 Speech and Language Processing (mixed undergraduate and graduate level)</i>	
	Teaching eval: 4.42/5.00 Spring 2020, 3.80/5.00 Spring 2017	