

Alan Ritter

CONTACT

E-mail: alan.ritter@cc.gatech.edu
Webpage: <http://aritter.github.io/>
Google Scholar: <https://scholar.google.com/citations?user=unXtH3IAAAAJ>

Natural Language Processing, Large Language Models, Machine Learning

EDUCATION

University of Washington, Seattle, WA USA
Ph.D. in Computer Science & Engineering, 2013
Western Washington University, Bellingham, WA USA
B.S./M.S., Computer Science, 2006

PROFESSIONAL APPOINTMENTS

Georgia Institute of Technology, Atlanta GA USA
Associate Professor - School of Interactive Computing, College of Computing **2020 - Present**

The Ohio State University, Columbus OH USA
Associate Professor - Computer Science and Engineering **2020**
Assistant Professor - Computer Science and Engineering **2014 - 2020**

Carnegie Mellon University, Pittsburgh PA USA
Visiting Faculty - Language Technologies Institute **2019**
Postdoctoral Fellow - Machine Learning Department **2013 - 2014**

Microsoft Research, Redmond WA USA
Consultant - Natural Language Processing Group **2014 - 2016**
Intern - Natural Language Processing Group **2009**
Intern - Knowledge Tools Group **2008**

Naval Research Laboratory, Marine Meteorology Division, Monterey, CA USA
Intern - Data Assimilation Group **2006**

GRANTS AS PI

Ford Motor Company, (\$231,938), Reliable Repair Chatbot, 2024-2026
Sony Group, (\$100k), Adapting LLM-based Multimodal Dialog Across Languages and Cultures, 2024-2025
NSF, (\$500k), CAREER: Large-Scale Learning for Information Extraction, 2019-2024
IARPA, (\$491,467), Stable, Interpretable, and General Neural Attribution for Trust of Users and Robust Effacement (SIGNATURE), 2022-2024
DARPA, (\$800k), Logically Adaptive Molecular Design Automation (LAMDA), 2019-2023
IARPA, (\$850k), Task and User-Aware Representation Learning for Fine-Grained Cross-Lingual Information Retrieval, 2019-2023
NSF, (\$151k), CRII: Learning to Extract Events from Knowledge Base Revisions, 2015-2018
Amazon Research Award, (\$50k) Integrating Data-Driven Conversation in Open-Domain Question Answering, 2018
DARPA (\$600k) Modeling the Spread of Information through Social and Knowledge Graphs, 2017-2021
IARPA (\$507k) Extracting a Realtime Cybersecurity Knowledge Graph from Text, 2016-2019
CMU Software Engineering Institute (\$40k) Events, relationships, and script learning for situational awareness, 1/2017-9/2017

GRANTS AS CO-PI OR SENIOR PERSONNEL	NSF (PI:Sonia Chernova) (\$19,995,808) AI Institute for Collaborative Assistance and Responsive Interaction for Networked Groups (AI-CARING)	
	NSF (PI:Betsy DiSalvo) (\$399,634) The Cost of AI: A Comparative Study of Machine Learning Training Methods 2024-Present	
	NIH (PI:Po-Yin Yen) (\$100k) Supporting systematic review production with article similarity network visualization, 2016-2018	
	Criteo Faculty Research Award (PI:Wei Xu) (\$20k) 2018	
	Battelle (PI:Marie de Marneffe) (\$23k) Trustworthiness in Social Media, 2017-2019	
PRESS COVERAGE	LLMs exhibit significant Western cultural bias, study finds , VentureBeat (March 2024) https://venturebeat.com/ai/large-language-models-exhibit-significant-western-cultural-bias-study-finds/	
	<ul style="list-style-type: none"> This article covers the work described in our ACL 2024 paper, entitled “<i>Having Beer after Prayer? Measuring Cultural Bias in Large Language Models</i>” 	
	GPT-3 mimics human love for ‘offensive’ Reddit comments, study finds , TNW (August 2021) https://thenextweb.com/news/gpt-3-and-humans-twice-as-likely-agree-with-offensive-reddit-comments-chatbots	
	<ul style="list-style-type: none"> This article covers the work described in our EMNLP 2021 paper, entitled “<i>Just Say No: Analyzing the Stance of Neural Dialogue Generation in Offensive Contexts</i>” 	
	Machine Learning Can Use Tweets to Spot Critical Security Flaws , WIRED (March 2019) https://www.wired.com/story/machine-learning-tweets-critical-security-flaws/	
	<ul style="list-style-type: none"> This article covers the work described in our NAACL 2019 paper, entitled “<i>Analyzing the Perceived Severity of Cybersecurity Threats Reported on Social Media</i>” 	
AWARDS	Sony Faculty Innovation Award (2024)	
	Social Impact Award for <i>Having Beer after Prayer? Measuring Cultural Bias in Large Language Models</i> at ACL (2024)	
	Reproduction Award for <i>Do CoNLL-2003 Named Entity Taggers Still Work Well in 2023?</i> at ACL (2023)	
	NSF CAREER Award , (2019)	
	Amazon Research Award , (2018)	
	Lumley Research Award , College of Engineering, Ohio State University (2019)	
	Outstanding Student Organization Advisor , College of Engineering, Ohio State University (2018)	
	Best Student Paper Award for <i>Learning to Generalize for Complex Selection Tasks</i> at the ACM Conference on Intelligent User Interfaces (2009)	
EXTERNAL THESIS COMMITTEES	NDSEG Fellowship (2008)	
	<i>Hassan Shavarani</i>	PhD, December 2024
	Structured Prediction for Compute Efficient and High Accuracy NLP (Advisor: Anoop Sarkar) Simon Frasier University	
	<i>Evangelia Spiliopoulou</i>	PhD, Jan 2023
	Modeling Event Implications via Multi-faceted Entity Representations (Advisor: Eduard Hovy) Carnegie Mellon University	
	<i>Alex Memory</i>	PhD, Oct 2019
	Collective Relational Data Integration with Diverse and Noisy Evidence (Advisor: Lise Getoor) University of Maryland	

Young-Bum Kim

PhD, Feb 2015

Natural Language Technologies for Low-Resource Languages (Advisor: Benjamin Snyder)

University of Wisconsin

INVITED TALKS

University of Tokyo, October 2024

Shonan Seminar on Empowering Local Open Source LLMs, October 2024

University of California San Diego, May 2024

University of California Irvine, March 2024

BrainLink Workshop, December 2023

University of Chicago / TTIC, November 2023

Allen Institute for Artificial Intelligence, October 2023

University of Washington, October 2023

Matching Workshop at ACL, July 2023

University of Pennsylvania, February 2023

Princeton University, October 2022

Columbia University, October 2022

New York University, October 2022

Google New York, October 2022

Johns Hopkins University, May 2022

Megagon Labs, January 2022

Twitter, September 2021

University of Notre Dame, September 2021

University of Maryland, April 2020

UMass Amherst, February 2020

NeurIPS Workshop on Conversational AI, December 2019

Georgia Institute of Technology, December 2019

Twitter, June 2019

University of Michigan, April 2019

University of Maryland, April 2019

Rice University, March 2019

University of Wisconsin, January 2019

Stanford NLP Seminar, May 2018

NeurIPS Workshop on Learning With Limited Labeled Data, December 2017

Microsoft Research Faculty Summit, July 2017

Midwest Machine Learning Symposium, June 2017

University of Colorado, March 2017

University of Edinburgh, July 2016

University of North Carolina, March 2016

Tsinghua University, December 2016

Machine Learning Seminar, Purdue University, December 2015

Statistics Department, Ohio State University, October 2015

School of Computer Science, Carnegie Mellon University, April 2015

CLIP, University of Maryland, March 2015

Clunch, University of Pennsylvania, September 2014

University of Pittsburgh, January 2014

Machine Learning Lunch, Carnegie Mellon University, October 2013

IEEE Intelligence and Security Informatics, June 2013

Toyota Technological Institute at Chicago, April 2013

Microsoft Research, April 2013

University of Texas at Dallas, March 2013

The Ohio State University, March 2013

Machine Learning and Friends Lunch, University of Massachusetts Amherst, March 2013

Human Language Technology Center of Excellence, March 2013

Microsoft Research, September 2011

Pacific Northwest National Laboratory, January 2013
 New York University, September 2012
 Wavii, April 2012
 Facebook, March 2012
 Google Research, September 2011
 Yahoo! Labs, September 2011
 Twitter, May 2011
 Information Sciences Institute, University of Southern California, February 2011
 ICT workshop on Reasoning with Text, February 2011
 UW CSE Industrial Affiliates Meeting, October 2009
 UW/Microsoft Quarterly Symposium on Computational Linguistics, October 2009
 Boeing Research, February 2008

PH.D. STUDENTS Geyang Guo
 Ruohao Guo
 Junmo Kang
 Duong Minh Le
 Nghia Max Le
 Ethan Mendes
 Jungsoo Park
 Jonathan Zheng

ALUMNI Yang Chen. Extracting Knowledge with Multimodal and Multilingual Intelligent Systems. (Ph.D. Georgia Institute of Technology, 2024)
 Fan Bai (Ph.D. Georgia Institute of Technology 2023, M.S. Ohio State University, 2020)
 Sandesh Swamy. Forecasting Event Outcomes from User Predictions on Twitter. (M.S., Ohio State University, 2017)
 Jeniya Tabassum. Information Extraction From User Generated Noisy Texts. co-advised with Wei Xu (Ph.D., Ohio State University, 2020)
 Shi Zong. Extracting Events from Social Media and the Web. (Ph.D., Ohio State University, 2020)

TEACHING GT CS 4650: Natural Language Processing (Fall 2024)
 GT CS 8803: Large Language Models (Spring 2024)
 GT CS 4650: Natural Language Processing (Fall 2023)
 GT CS 7650: Natural Language Processing (Spring 2023)
 GT CS 4650: Natural Language Processing (Fall 2022)
 GT CS 7650: Natural Language Processing (Spring 2022)
 GT CS 4641: Machine Learning (Fall 2021)
 GT CS 7650: Natural Language Processing (Spring 2021)
 OSU CSE 5539: Social Media and Text Analytics (Spring 2020)
 OSU CSE 5525: Speech and Language Processing (Fall 2019)
 OSU CSE 5525: Speech and Language Processing (Spring 2019)
 OSU CSE 5523: Machine Learning and Statistical Pattern Recognition (Fall 2018)
 OSU CSE 5539: Social Media and Text Analytics (Spring 2018)
 OSU CSE 5522: Survey of Artificial Intelligence II: Advanced Techniques (Fall 2017)
 OSU CSE 5523: Machine Learning and Statistical Pattern Recognition (Spring 2017)
 OSU CSE 5525: Speech and Language Processing (Spring 2016)
 OSU CSE 5539: Web Information Extraction (Fall 2015)
 OSU CSE 5522: Survey of Artificial Intelligence II: Advanced Techniques (Spring 2015)
 OSU CSE 5539: NLP and IE for the Social Web (Fall 2014)

SERVICE

Program Chair NAACL 2025

Best Paper Awards Committee EACL 2024

Demo Chair ACL 2023

Organizer ACL/EMNLP Workshop on Noisy User-Generated Text (WNUT 2015, 2016, 2017, 2018, 2019, 2020, 2021), Sentiment Analysis in Twitter Shared Task (SemEval 2013,2014,2015,2016), Named Entity Recognition in Twitter Shared Task (at WNUT 2015, 2016)

Senior Area Chair EMNLP 2020 (Information Extraction), NAACL 2021 (Machine Learning), EMNLP 2021 (Computational Social Science), AAAI 2022, ACL 2022 (Information Extraction), EMNLP 2022) (Information Extraction

Area Chair ACL 2014 (NLP for the web and social media), NAACL 2015 (Information Extraction and Question Answering), EMNLP 2016 (Information Extraction, Information Retrieval and Question Answering), NAACL 2018 (Computational Social Science), EMNLP 2018 (Text Mining and Information Retrieval), CoNLL 2018, ACL 2019 (Information Extraction and Text Mining), EMNLP 2019 (Social Media and Computational Social Science), ACL 2020 (Information Extraction), ACL 2021 (Information Extraction) EMNLP 2023 (Information Extraction)

Steering Committee SemEval 2016,2018,2019,2020

Proposal Reviewing Panelist for U.S. National Science Foundation, Division of Information and Intelligent Systems (CISE/IIS), (Twice in 2017, Twice in 2018, Once in 2019, Once in 2023, Once in 2024), Reviewer for U.S. National Science Foundation, Division of Information and Intelligent Systems (CISE/IIS), (Once in 2020), Independent Reviewer for Hong Kong Research Grants Council (2016), Independent Reviewer for Natural Sciences and Engineering Research Council of Canada (2016), Independent Reviewer U.S. Department of Energy (2015)

Departmental/College Service GT College of Computing Undergraduate Curriculum Committee, 2020-2021, GT Computational Media Undergraduate Curriculum Committee, 2020-2021, GT School of Interactive Computing Faculty Hiring Committee, 2020-2021, OSU CSE Graduate Studies Committee (2016-2019), OSU CSE Curriculum Committee (2019-2020), OSU CSE Graduate Admissions Committee - Artificial Intelligence Group Representative (2014-2020), OSU CSE Prospective Ph.D. Student Visit Day Organizer, (2014-2016)

PATENTS

US Patent US20110010669

“Items Selection Via Automatic Generalization”

- This covers the work from my 2008 internship at Microsoft Research, which also resulted in a best student paper award at the ACM Conference on Intelligent User Interfaces.

PUBLICATIONS

Granular Privacy Control for Geolocation with Vision Language Models
Ethan Mendes, Yang Chen, James Hays, Sauvik Das, Wei Xu, **Alan Ritter**
EMNLP 2024

Schema-Driven Information Extraction from Heterogeneous Tables
Fan Bai, Junmo Kang, Gabriel Stanovsky, Dayne Freitag, Mark Dredze, **Alan Ritter**
Findings of EMNLP 2024

ChatHF: Collecting Rich Human Feedback from Real-time Conversations
Andrew Li, Zhenduo Wang, Ethan Mendes, Duong Minh Le, Wei Xu, **Alan Ritter**
EMNLP 2024 (Demo Track)

Are Language Models Robust Coreference Resolvers?
Nghia T. Le, **Alan Ritter**
COLM 2024

UniIR: Training and Benchmarking Universal Multimodal Information Retrievers
Cong Wei, Yang Chen, Haonan Chen, Hexiang Hu, Ge Zhang, Jie Fu, **Alan Ritter**, Wenhui Chen
ECCV 2024

Constrained Decoding for Cross-lingual Label Projection
Duong Minh Le, Yang Chen, **Alan Ritter**, Wei Xu
ICLR 2024

Tensor Trust: Interpretable Prompt Injection Attacks from an Online Game
Sam Toyer, Olivia Watkins, Ethan Adrian Mendes, Justin Svegliato, Luke Bailey, Tiffany Wang, Isaac Ong, Karim Elmaaroufi, Pieter Abbeel, Trevor Darrell, **Alan Ritter**, Stuart Russell
ICLR 2024 (Spotlight)

Having Beer after Prayer? Measuring Cultural Bias in Large Language Models
Tarek Naous, Michael J. Ryan, **Alan Ritter**, Wei Xu
Proceedings of ACL 2024
Social Impact Award

Meta-Tuning LLMs to Leverage Lexical Knowledge for Generalizable Language Style Understanding
Ruohao Guo, Wei Xu, **Alan Ritter**
Proceedings of ACL 2024

NEO-BENCH: Evaluating Robustness of Large Language Models with Neologisms
Jonathan Zheng, **Alan Ritter**, Wei Xu
Proceedings of ACL 2024

Reducing Privacy Risks in Online Self-Disclosures with Language Models
Yao Dou, Isadora Krsek, Tarek Naous, Anubha Kabra, Sauvik Das, **Alan Ritter**, Wei Xu
Proceedings of ACL 2024

Self-Specialization: Uncovering Latent Expertise within Large Language Models
Junmo Kang, Hongyin Luo, Yada Zhu, Jacob Hansen, James Glass, David Cox, **Alan Ritter**, Rogerio Feris, Leonid Karlinsky
Findings of ACL 2024

Can Pre-trained Vision and Language Models Answer Visual Information-Seeking Questions?
Yang Chen, Hexiang Hu, Yi Luan, Haitian Sun, Soravit Changpinyo, **Alan Ritter**, Ming-Wei Chang
EMNLP 2023

Distill or Annotate? Cost-Efficient Fine-Tuning of Compact Models
Junmo Kang, Wei Xu and **Alan Ritter**
ACL 2023

Duong Minh Le, Ruohao Guo, Wei Xu and **Alan Ritter**
Improved Instruction Ordering in Recipe-Grounded Conversation
ACL 2023

Ethan Mendes, Yang Chen, Wei Xu, **Alan Ritter**
“Human-in-the-loop Evaluation for Early Misinformation Detection: A Case Study of COVID-19 Treatments”

ACL 2023

Shuheng Liu, **Alan Ritter**

“Do CoNLL-2003 Named Entity Taggers Still Work Well in 2023?”

ACL 2023

Reproduction Award

Yang Chen, Chao Jiang, **Alan Ritter**, Wei Xu

“Frustratingly Easy Label Projection for Cross-lingual Transfer”

Findings of ACL 2023

Nghia Le, Fan Bai, **Alan Ritter**

“Few-Shot Anaphora Resolution in Scientific Protocols via Mixtures of In-Context Experts”

EMNLP Findings 2022

Jonathan Qiaoyi Zheng, Ashutosh Baheti, Tarek Naous, Wei Xu, **Alan Ritter**

“Stanceosaurus: Classifying Stance Towards Multicultural Misinformation”

To Appear in Proceedings of EMNLP 2022

Fan Bai, **Alan Ritter**, Peter Madrid, Dayne Freitag and John Niekraz

“SynKB: Semantic Search for Synthetic Procedures”

To Appear in Proceedings of EMNLP 2022 (Demo Track)

Shi Zong, Ashutosh Baheti, Wei Xu, **Alan Ritter**

“Extracting a Knowledge Base of COVID-19 Events from Social Media”

To Appear in Proceedings of Coling 2022

Fan Bai, **Alan Ritter** and Wei Xu

“Pre-train or Annotate? Domain Adaptation with a Constrained Budget”

Proceedings of EMNLP 2021

Ashutosh Baheti, Maarten Sap, **Alan Ritter**, Mark Riedl

“Just Say No: Analyzing the Stance of Neural Dialogue Generation in Offensive Contexts”

Proceedings of EMNLP 2021 [\[Covered by TNW\]](#)

Yang Chen and **Alan Ritter**

“Model Selection for Cross-Lingual Transfer”

Proceedings of EMNLP 2021

Ronen Tamari, Fan Bai, **Alan Ritter** and Gabriel Stanovsky

“Process-Level Representation of Scientific Protocols with Interactive Annotation”

Proceedings of EACL 2021

Wuwei Lan, Yang Chen, Wei Xu, **Alan Ritter**

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

Proceedings of EMNLP 2020 (short paper)

Shi Zong, **Alan Ritter** and Eduard Hovy

“Measuring Forecasting Skill from Text”

Proceedings of ACL 2020

Ashutosh Baheti, **Alan Ritter** and Kevin Small

“Fluent Response Generation for Conversational Question Answering”

Proceedings of ACL 2020

Jeniya Tabassum, Mounica Maddela, Wei Xu and **Alan Ritter**
“Code and Named Entity Recognition in StackOverflow”
Proceedings of ACL 2020

Fan Bai, **Alan Ritter**
“Structured Minimally Supervised Learning for Neural Relation Extraction”
Proceedings of NAACL 2019

Shi Zong, **Alan Ritter**, Graham Mueller and Evan Wright
“Analyzing the Perceived Severity of Cybersecurity Threats Reported on Social Media”
Proceedings of NAACL 2019 [\[Covered by WIRED\]](#)

Ashutosh Baheti, **Alan Ritter**, Jiwei Li, Bill Dolan
“Generating More Interesting Responses in Neural Conversation Models with Distributional Constraints”
Proceedings of EMNLP 2018

Chaitanya Kulkarni, Wei Xu, **Alan Ritter**, Raghu Machiraju
“An Annotated Corpus for Machine Reading of Instructions in Wet Lab Protocols”
Proceedings of NAACL 2018 (short paper)

Jiwei Li, Will Monroe, Tianlin Shi, Sébastien Jean, **Alan Ritter**, Dan Jurafsky
“Adversarial Learning for Neural Dialogue Generation”
Proceedings of EMNLP 2017

Sandesh Swamy, **Alan Ritter** and Marie-Catherine de Marneffe
“i have a feeling trump will win.....: Forecasting Winners and Losers from User Predictions on Twitter”
Proceedings of EMNLP 2017

Alexander Konovalov, Benjamin Strauss, **Alan Ritter** and Brendan O’Connor
“Learning to Extract Events from Knowledge Base Revisions”
Proceedings of WWW 2017

Jiwei Li, Will Monroe, **Alan Ritter**, Michel Galley, Jianfeng Gao, Dan Jurafsky
“Deep Reinforcement Learning for Dialogue Generation”
Proceedings of EMNLP 2016

Jeniya Tabassum, **Alan Ritter**, Wei Xu
“A Minimally Supervised Method for Recognizing and Normalizing Time Expressions in Twitter”
Proceedings of EMNLP 2016

Xinlei Chen, **Alan Ritter**, Abhinav Gupta and Tom Mitchell
“Sense Discovery via Co-Clustering on Images and Text”
Proceedings of CVPR 2015

Alan Ritter, Evan Wright William Casey and Tom Mitchell
“Weakly Supervised Extraction of Computer Security Events from Twitter”
Proceedings of WWW 2015

Tom Mitchell, William Cohen, Estevam Hruschka, Partha Talukdar, Justin Betteridge, Andrew Carlson, Bhavana Dalvi, Matt Gardner, Bryan Kisiel, Jayant Krishnmurthy, Ni Lao, Kathryn Mazaitis, Tahir Mohammad, Ndapa Nakashole, Emmanouil Antonios Platanios, **Alan Ritter**, Mehdi Samadi, Burr Settles, Richard Wang, Derry Wijaya, Abhinav Gupta, Xinlei Chen, Abulhair Saparov, Mal-

colm Greaves and Joel Welling
“Never Ending Learning”
Proceedings of AAAI 2015

Jiwei Li, **Alan Ritter**, Claire Cardie and Eduard Hovy
“Major Life Event Extraction from Twitter based on Congratulations/Condolences Speech Acts”
Proceedings of EMNLP 2014

Jiwei Li, **Alan Ritter** and Eduard Hovy
“Weakly Supervised User Profile Extraction from Twitter”
Proceedings of ACL 2014

Wei Xu, **Alan Ritter**, Chris Callison-Burch, William B. Dolan and Yangfeng Ji
“Extracting Lexically Divergent Paraphrases from Twitter”
Proceedings of TACL 2014

Alan Ritter, Luke Zettlemoyer, Mausam and Oren Etzioni
“Modeling Missing Data in Distant Supervision for Information Extraction”
Proceedings of TACL 2013

Alan Ritter, Mausam, Oren Etzioni, Sam Clark
“Open Domain Event Extraction from Twitter”
Proceedings of KDD 2012

Wei Xu, **Alan Ritter**, Bill Dolan, Ralph Grishman, Colin Cherry
“Paraphrasing for Style”
Proceedings of Coling 2012

Alan Ritter, Sam Clark, Mausam, Oren Etzioni
“Named Entity Recognition in Tweets: An Experimental Study”
Proceedings of EMNLP 2011

Alan Ritter, Colin Cherry, Bill Dolan
“Data-Driven Response Generation in Social Media”
Proceedings of EMNLP 2011

Alan Ritter, Mausam and Oren Etzioni
“A Latent Dirichlet Allocation Method for Selectional Preferences”
Proceedings of ACL 2010

Alan Ritter, Colin Cherry and Bill Dolan
“Unsupervised Modeling of Twitter Conversations”
Proceedings of HLT-NAACL 2010

Alan Ritter and Sumit Basu
“Learning to Generalize for Complex Selection Tasks”
Best Student Paper Award
Proceedings of IUI 2009

Alan Ritter, Doug Downey, Stephen Soderland, and Oren Etzioni
“It’s a Contradiction – No, It’s Not: A Case Study using Functional Relations”
Proceedings of EMNLP 2008