

Jordan Boyd-Graber

Legal Name: Jordan Boyd-Graber Ying

jordan.ying@ntu.edu.sg • (+65) 9148 1138 • <https://boydgraber.org>

Office S3-01c-103, CCDS, Nanyang Technological University, 50 Nanyang Avenue, Singapore 639798

Academic Appointments

Nanyang Technological University	SINGAPORE
Full Professor: Computing and Data Science	2026–present
University of Maryland	COLLEGE PARK, MD
Full Professor: Computer Science (tenure home), UMIACS, LSC	2024–2026
Associate Professor: Computer Science (tenure home), UMIACS, LSC, iSchool ¹	2017–2024
Assistant Professor, College of Information Studies	2010–2014
Postdoc (Advisor: Philip Resnik)	2009–2010
University of Colorado Boulder	BOULDER, CO
Associate Professor, ² Computer Science	2017
Assistant Professor, Computer Science	2014–2017

Other Employment

Google Switzerland, GmbH	ZÜRICH, CH
Visiting Research Scientist	2019–2020
Princeton University	PRINCETON, NJ
Writing Fellow, Princeton Writing Center	2007–2008
Google	NEW YORK, NY
Intern, Google Books	2007
University of California Los Angeles	LOS ANGELES, CA
Digital Humanities Programmer	2004
California Institute of Technology	PASADENA, CA
Newsprint Researcher / Programmer, Einstein Papers Project	2003–2004
Peer Tutor, Hixon Writing Center	2001–2004
Lab Technician, Caltech Digital Media Center	2001–2003
Berlin-Brandenburg Akademie der Wissenschaften	BERLIN, GERMANY
Praktikant, Digitales Wörterbuch der Deutschen Sprache	2002

Educational Background

Princeton University	PRINCETON, NJ
Ph.D., Computer Science	2010
Thesis: Linguistic Extensions of Topic Models (Advisor: David Blei)	
M.A., Computer Science	2007
California Institute of Technology	PASADENA, CA
B.S., Computer Science	2004
B.S., History	2004

Immigration status: U.S. citizen (born in Eagle County, Colorado)

Fellowships, Prizes, and Awards

- Outstanding Area Chair: ACL Rolling Review, January 2026
- Outstanding Area Chair: ACL Rolling Review, July 2025
- Outstanding Paper Award: NAACL, 2025
- Best Theme Paper Award: EMNLP, 2023
- Best Student Paper Honorable Mention (Alison Smith): IUI, 2018
- NSF CAREER Award, 2017
- Quora Top Writer: 2016, 2017, 2018

¹Information Studies appointment decreased to courtesy in 2022.

²Tenure announced 2017 but never awarded (I resigned from CU before it was effective.)

- Best Paper Award: NAACL, 2016
- Best Demonstration Award: NeurIPS, 2015
- Karen Spärk Jones Award, 2015
- Best Paper Award: CoNLL, 2015
- Honorable Mention, Best Student Paper: NeurIPS, 2009
- Computing Innovation Postdoctoral Fellowship, 2009 (declined)
- Richter Undergraduate Research Fellowship: 2001 and 2002
- Caltech Jorgensen Scholarship: 2001–2004
- AAAI Research Award: International Science and Engineering Fair, 2000

Publications

Students directly advised or co-advised in underline.

Books

1. **Jordan Boyd-Graber**, Yuening Hu, and David Mimno. **Applications of Topic Models**. 2017, 153 pages.

Chapters in Books

1. **Jordan Boyd-Graber**, Shi Feng, and Pedro Rodriguez. **Human-Computer Question Answering: The Case for Quizbowl**. *The NIPS '17 Competition: Building Intelligent Systems*, 2018, 10 pages.
2. Evgeny Klochikhin and **Jordan Boyd-Graber**. **Text Analysis**. *Big Data and Social Science Research: Theory and Practical Approaches*, 2016, 27 pages.
3. **Jordan Boyd-Graber**, David Mimno, and David Newman. **Care and Feeding of Topic Models: Problems, Diagnostics, and Improvements**. *Handbook of Mixed Membership Models and Their Applications*, 2014, 39 pages.
4. Sonya S. Nikolova, **Jordan Boyd-Graber**, and Christiane Fellbaum. **Collecting Semantic Similarity Ratings to Connect Concepts in Assistive Communication Tools**. *Modeling, Learning and Processing of Text Technological Data Structures*, 2011, 11 pages.

Refereed Journal Articles

1. Ishani Mondal, Yiwen Song, Mihir Parmar, Palash Goyal, **Jordan Boyd-Graber**, Tomas Pfister, and Yale Song. **CANVAS: Continuity-Aware Narrative Generation via Visual Agentic Storyboarding**. *Findings of Empirical Methods in Natural Language Processing*, 2026 (29.7% Acceptance Rate).
2. Senyi Zhang, Tasnim Kabir, and **Jordan Boyd-Graber**. **Um, Actually: Timely and Calibrated Fact-Checking with DIVER**. *Findings of Empirical Methods in Natural Language Processing*, 2026 (29.7% Acceptance Rate).
3. Maharshi Gor, Yoo Yeon Sung, Yu Hou, Eve Fleisig, Zhu Irene Ying, Tianyi Zhou, and **Jordan Boyd-Graber**. **AI, Take the Wheel: What Drives Delegation and Trust in Human-Computer Cooperative Question Answering?**. *Findings of the Association for Computational Linguistics*, 2026.
4. Tasnim Kabir, Dmytro Kurdydyk, Aadi Palnitkar, Liam Dorn, Ahmed Haj Ahmed, and **Jordan Boyd-Graber**. **AUDITA: A New Dataset to Audit Whether Humans or AI are Better at Audio QA**. *Findings of the Association for Computational Linguistics*, 2026.
5. Feng Gu, Zongxia Li, Carlos R. Colon, Benjamin Evans, Ishani Mondal, and **Jordan Boyd-Graber**. **Large Language Models Are Effective Human Annotation Assistants, But Not Good Independent Annotators**. *Findings of the Association for Computational Linguistics*, 2026.
6. Wichayaporn Wongkamjan, Yanze Wang, Feng Gu, Denis Peskoff, Jonathan K. Kummerfeld, Jonathan May, and **Jordan Boyd-Graber**. **Should I Trust You? Detecting Deception in Negotiations using Counterfactual RL**. *Findings of the Association for Computational Linguistics*, 2025.
7. Xiyang Wu, Tianrui Guan, Dianqi Li, Shuaiyi Huang, Xiaoyu Liu, Xijun Wang, Ruiqi Xian, Abhinav Shrivastava, Furong Huang, **Jordan Boyd-Graber**, Tianyi Zhou, and Dinesh Manocha. **AUTOHALLUSION: Automatic Generation of Hallucination Benchmarks for Vision-Language Models**. *Findings of the Empirical Methods in Natural Language Processing*, 2024.
8. Zongxia Li, Ishani Mondal, Huy Nghiem, Yijun Liang, and **Jordan Boyd-Graber**. **PEDANTS (Precise Evaluations of Diverse Answer Nominee Text for Skinflints): Use Evaluation Metrics Wisely—Efficient Evaluation Analysis and Benchmarking for Open-Domain Question Answering**. *Findings of the Empirical Methods in Natural Language Processing*, 2024.
9. Ishani Mondal, Zongxia Li, Yufang Hou, Anandhavelu Natarajan, Aparna Garimella, Sambaran Bandyopadhyay, and **Jordan Boyd-Graber**. **SciDoc2Diagrammer-MAF: Towards Generation of Scientific Diagrams from Documents guided by Multi-Aspect Feedback Refinement**. *Findings of the Empirical Methods in Natural Language Processing*, 2024 (21% Acceptance Rate).

10. Quynh C. Nguyen, Elizabeth M. Aparicio, Michelle Jasczynski, Amara Channell Doig, Xiaohe Yue, Heran Mane, [Neha Pundlik Srikanth](#), Francia Ximena Marin Gutierrez, Nataly Delcid, Xin He, and **Jordan Boyd-Graber**. **Randomized Pilot of Rosie, a Health Education Question-and-Answer Chatbot for New Mothers**. *Journal of Medical Internet Research: Journal of Formative Research*, 2024.
11. [Chenglei Si](#), Weijia Shi, Chen Zhao, Luke Zettlemoyer, and **Jordan Boyd-Graber**. **Getting MoRE out of Mixture of Language Model Reasoning Experts**. *Findings of Empirical Methods in Natural Language Processing*, 2023 (45% Acceptance Rate).
12. Mane, Heran Y., Channell Doig, Amara, Marin Gutierrez, Francia Ximena, Jasczynski, Michelle, Yue, Xiaohe, [Neha Pundlik Srikanth](#), Mane, Sourabh, Sun, Abby, Moats, Rachel Ann, Patel, Pragat, He, Xin, **Jordan Boyd-Graber**, Aparicio, Elizabeth M., and Nguyen, Quynh C.. **Practical Guidance for the Development of Rosie, a Health Education Question-and-Answer Chatbot for New Mothers**. *Journal of Public Health Management and Practice*, 2023.
13. Wanrong He, [Andrew Mao](#), and **Jordan Boyd-Graber**. **Cheater’s Bowl: Human vs. Computer Search Strategies for Open-Domain QA**. *Findings of Empirical Methods in Natural Language Processing*, 2022, 9 pages.
14. [Chenglei Si](#), Chen Zhao, Sewon Min, and **Jordan Boyd-Graber**. **Re-Examining Calibration: The Case of Question Answering**. *Findings of Empirical Methods in Natural Language Processing*, 2022, 9 pages.
15. [Fenfei Guo](#), Chen Zhang, Zhirui Zhang, Qixin He, Kejun Zhang, Jun Xie, and **Jordan Boyd-Graber**. **Automatic Song Translation for Tonal Languages**. *Findings of the Association for Computational Linguistics*, 2022, 9 pages (31% Acceptance Rate).
16. [Denis Peskov](#), Viktor Hangya, **Jordan Boyd-Graber**, and Alexander Fraser. **Adapting Entities across Languages and Cultures**. *Findings of Empirical Methods in Natural Language Processing*, 2021, 5 pages (37% Acceptance Rate).
17. Wenyan Li, Alvin Grissom II, and **Jordan Boyd-Graber**. **An Attentive Recurrent Model for Incremental Prediction of Sentence-final Verbs**. *Findings of EMNLP*, 2020, 9 pages.
18. Tianze Shi, [Chen Zhao](#), **Jordan Boyd-Graber**, Hal Daumé III, and Lillian Lee. **On the Potential of Lexico-logical Alignments for Semantic Parsing to SQL Queries**. *Findings of EMNLP*, 2020, 9 pages.
19. Dasha Pruss, [Yoshinari Fujinuma](#), Ashlynn Daughton, Michael Paul, Brad Arnot, Danielle Szafir, and **Jordan Boyd-Graber**. **Zika discourse in the Americas: A multilingual topic analysis of Twitter**. *PlosOne*, 2019, 23 pages.
20. [Eric Wallace](#), Pedro Rodriguez, [Shi Feng](#), Ikuya Yamada, and **Jordan Boyd-Graber**. **Trick Me If You Can: Human-in-the-loop Generation of Adversarial Question Answering Examples**. *Transactions of the Association for Computational Linguistics*, 2019, 14 pages.
21. Aaron Gerow, Yuening Hu, **Jordan Boyd-Graber**, David M. Blei, and James A. Evans. **Measuring Discursive Influence Across Scholarship**. *Proceedings of the National Academies of Science*, 2018.
22. Tak Yeon Lee, [Alison Smith](#), Kevin Seppi, Niklas Elmqvist, **Jordan Boyd-Graber**, and Leah Findlater. **The Human Touch: How Non-expert Users Perceive, Interpret, and Fix Topic Models**. *International Journal of Human-Computer Studies*, 2017, 27 pages.
23. **Jordan Boyd-Graber**. **Humans and Computers Working Together to Measure Machine Learning Interpretability**. *The Bridge*, 2017, 5 pages.
24. [Alison Smith](#), Tak Yeon Lee, [Forough Poursabzi-Sangdeh](#), **Jordan Boyd-Graber**, Kevin Seppi, Niklas Elmqvist, and Leah Findlater. **Evaluating Visual Representations for Topic Understanding and Their Effects on Manually Generated Labels**. *Transactions of the Association for Computational Linguistics*, 2017, 15 pages.
25. [Yuening Hu](#), **Jordan Boyd-Graber**, Brianna Satinoff, and [Alison Smith](#). **Interactive Topic Modeling**. *Machine Learning*, 2014, 56 pages.
26. [Viet-An Nguyen](#), **Jordan Boyd-Graber**, Philip Resnik, Deborah Cai, Jennifer Midberry, and Yuanxin Wang. **Modeling Topic Control to Detect Influence in Conversations using Nonparametric Topic Models**. *Machine Learning*, 2014, 48 pages.
27. [Ke Zhai](#), **Jordan Boyd-Graber**, and Shay B. Cohen. **Online Adaptor Grammars with Hybrid Inference**. *Transactions of the Association for Computational Linguistics*, 2014, 12 pages.
28. [Viet-An Nguyen](#), **Jordan Boyd-Graber**, and Stephen Altschul. **Dirichlet Mixtures, the Dirichlet Process, and the Structure of Protein Space**. *Journal of Computational Biology*, 2013, 48 pages.
29. Alexander Geyken and **Jordan Boyd-Graber**. **Automatic classification of multi-word expressions in print dictionaries**. *Linguisticae Investigationes*, 2003, 16 pages.

Preprints / Working Papers

1. Benjamin Börschinger, **Jordan Boyd-Graber**, Christian Buck, Jannis Bulian, Massimiliano Ciaramita, Michelle Chen Huebscher, Wojciech Gajewski, Yannic Kilcher, Rodrigo Nogueira, and Lierni Sestorain Saralegu. **Meta Answering for Machine Reading**. *ArXiv*, 2020, 10 pages.

2. [Pedro Rodriguez, Shi Feng, Mohit Iyyer, He He, and Jordan Boyd-Graber. Quizbowl: The Case for Incremental Question Answering. *ArXiv*, 2020, 55 pages.](#)

Refereed Conference Proceedings

1. [Nishant Balepur, Bhavya Rajasekaran, Hyunjin Jane Oh, Michael Xie, Atrey Desai, Vipul Gupta, Steven James Moore, Eunsol Choi, Rachel Rudinger, and Jordan Boyd-Graber. BenchMarker: An Education-Inspired Toolkit for Highlighting Flaws in Multiple-Choice Benchmarks. *Association for Computational Linguistics*, 2026.](#)
2. [Nishant Balepur, Malachi Hamada, Varsha Kishore, Sergey Feldman, Amanpreet Singh, Pao Siangliulue, Joseph Chee Chang, Eunsol Choi, Jordan Boyd-Graber, and Aakanksha Naik. Language Models Don't Know What You Want: Evaluating Personalization in Deep Research Needs Real Users. *Association for Computational Linguistics*, 2026.](#)
3. [HyoJung Han, Nishant Balepur, Jordan Boyd-Graber, and Marine Carpuat. Measuring User's Mental Models of Speech Translation in Human-MT Collaboration. *Association for Computational Linguistics*, 2026.](#)
4. [Nishant Balepur, Paiheng Xu, Wei Ai, Eunsol Choi, Rachel Rudinger, and Jordan Boyd-Graber. Check The Scoreboard: An Analysis of Scoring Schemes on Multiple-Choice Evaluation. *Empirical Methods in Natural Language Processing*, 2026 \(15.4% Acceptance Rate\).](#)
5. [Neha Pundlik Srikanth, Jordan Boyd-Graber, and Rachel Rudinger. DiscoTrace: Representing and Comparing Answering Strategies of Humans and LLMs in Information-Seeking Question Answering. *Empirical Methods in Natural Language Processing*, 2026 \(15.4% Acceptance Rate\).](#)
6. [Ishani Mondal, Javad Baghirov, and Jordan Boyd-Graber. Helping Figures Tell their Story! Paper-Grounded Video Generation Explaining Complex Scientific Figures. *Empirical Methods in Natural Language Processing*, 2026 \(15.4% Acceptance Rate\).](#)
7. [Ishani Mondal, Meera Bharadwaj, Ayush Roy, Aparna Garimella, and Jordan Boyd-Graber. SMART-Editor: A Multi-Agent Framework for Human-Like Design Editing with Structural Integrity. *European Association for Computational Linguistics*, 2026.](#)
8. [Zongxia Li, Wenhao Yu, Chengsong Huang, Rui Liu, Zhenwen Liang, Fuxiao Liu, Jingxi Che, Dian Yu, Jordan Boyd-Graber, Haitao Mi, and Dong Yu. Self-Rewarding Vision-Language Model via Reasoning Decomposition. *International Conference on Learning Representations*, 2026.](#)
9. [Nishant Balepur, Alexa Siu, Nedim Lipka, Franck Dernoncourt, Tong Sun, Jordan Boyd-Graber, and Puneet Mathur. MoDS: Moderating a Mixture of Document Speakers to Summarize Debatable Queries in Document Collections. *Annual Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics*, 2025.](#)
10. [Feng Gu, Wichayaporn Wongkamjan, Jonathan K. Kummerfeld, Denis Peskoff, Jonathan May, and Jordan Boyd-Graber. Personalized Help for Optimizing Low-Skilled Users' Strategy. *Annual Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics*, 2025.](#)
11. [Nishant Balepur, Feng Gu, Abhilasha Ravichander, Shi Feng, Jordan Boyd-Graber, and Rachel Rudinger. Reverse Question Answering: Can an LLM Write a Question so Hard \(or Bad\) that it Can't Answer?. *Annual Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics*, 2025.](#)
12. [Yoo Yeon Sung, Eve Fleisig, Yu Hope, Ishan Upadhyay, and Jordan Boyd-Graber. GRACE: A Granular Benchmark for Evaluating Model Calibration against Human Calibration. *Association for Computational Linguistics*, 2025.](#)
13. [Zongxia Li, Lorena Calvo-Bartolomé, Alexander Miserlis Hoyle, Paiheng Xu, Daniel Kofi Stephens, Juan Francisco Fung, Alden Dima, and Jordan Boyd-Graber. LLMs Struggle to Describe the Haystack without Human Help: A Social Science-Inspired Evaluation of Topic Models. *Association for Computational Linguistics*, 2025.](#)
14. [Neha Pundlik Srikanth, Rachel Rudinger, and Jordan Boyd-Graber. No Questions are Stupid and but some are Poorly Posed: Understanding Poorly-Posed Information-Seeking Questions. *Association for Computational Linguistics*, 2025.](#)
15. [Alexander Hoyle, Lorena Calvo-Bartolomé, Jordan Boyd-Graber, and Philip Resnik. ProxAnn: Use-Oriented Evaluations of Topic Models and Document Clustering. *Association for Computational Linguistics*, 2025.](#)
16. [Nishant Balepur, Rachel Rudinger, and Jordan Boyd-Graber. Which of These Best Describes Multiple Choice Evaluation with LLMs? A\) Forced B\) Flawed C\) Fixable D\) All of the Above. *Association for Computational Linguistics*, 2025.](#)
17. [Nishant Balepur, Vishakh Padmakumar, Fumeng Yang, Shi Feng, Rachel Rudinger, and Jordan Boyd-Graber. Whose Boat Does it Float? Improving Personalization in Preference Tuning via Inferred User Personas. *Association for Computational Linguistics*, 2025.](#)
18. [Nishant Balepur, Matthew Shu, Yoo Yeon Sung, Seraphina Goldfarb-Tarrant, Shi Feng, Fumeng Yang, Rachel Rudinger, and Jordan Boyd-Graber. A Good Plan is Hard to Find: Aligning Models with Preferences is Misaligned with What Helps Users. *Empirical Methods in Natural Language Processing*, 2025.](#)

19. Lorena Calvo-Bartolomé, Valérie Aldana, Karla Cantarero, Alonso Madroñal de Mesa, Jerónimo Arenas-García, and **Jordan Boyd-Graber**. **Discrepancy Detection at the Data Level: Toward Consistent Multilingual Question Answering**. *Empirical Methods in Natural Language Processing*, 2025.
20. **Ishani Mondal**, Jack W. Stokes, Sujay Kumar Jauhar, Longqi Yang, Mengting Wan, Xiaofeng Xu, Xia Song, **Jordan Boyd-Graber**, and Jennifer Neville. **Group Preference Alignment: Customizing LLM Responses from In-Situ Conversations Only When Needed**. *Empirical Methods in Natural Language Processing (Industry)*, 2025.
21. Zongxia Li, Xiyang Wu, Guangyao Shi, Yubin Qin, Hongyang Du, Tianyi Zhou, Dinesh Manocha, and **Jordan Boyd-Graber**. **VideoHallu: Evaluating and Mitigating Multi-modal Hallucinations on Synthetic Video Understanding**. *Neural Information Processing Systems*, 2025.
22. **Yoo Yeon Sung**, **Maharshi Gor**, Eve Fleisig, **Ishani Mondal**, and **Jordan Boyd-Graber**. **ADVSCORE: A Metric for the Evaluation and Creation of Adversarial Benchmarks**. *North American Association for Computational Linguistics*, 2025.
23. **Wichayaporn Wongkamjan**, Feng Gu, Yanze Wang, Ulf Hermjakob, Jonathan May, Brandon M. Stewart, Jonathan K. Kummerfeld, Denis Peskoff, and **Jordan Boyd-Graber**. **More Victories, Less Cooperation: Assessing Cicero's Diplomacy Play**. *Association for Computational Linguistics*, 2024.
24. **Nishant Balepur**, Matthew Shu, Alexander Hoyle, Alison Robey, Shi Feng, Seraphina Goldfarb-Tarrant, and **Jordan Boyd-Graber**. **A SMART Mnemonic Sounds like "Glue Tonic": Mixing LLMs with Student Feedback to Make Mnemonic Learning Stick**. *Empirical Methods in Natural Language Processing*, 2024.
25. **Maharshi Gor**, Hal Daumé III, Tianyi Zhou, and **Jordan Boyd-Graber**. **Do great minds think alike? Investigating Human-AI Complementarity in Question Answering with CAIMIRA**. *Empirical Methods in Natural Language Processing*, 2024.
26. Matthew Shu, **Nishant Balepur**, Shi Feng, and **Jordan Boyd-Graber**. **KARL: Knowledge-Aware Retrieval and Representations aid Retention and Learning in Students**. *Empirical Methods in Natural Language Processing*, 2024.
27. **Tasnim Kabir**, **Yoo Yeon Sung**, Saptarashmi Bandyopadhyay, Hao Zou, Abhranil Chandra, and **Jordan Boyd-Graber**. **You Make me Feel like a Natural Question: Training QA Systems on Transformed Trivia Questions**. *Empirical Methods in Natural Language Processing*, 2024.
28. **Ishani Mondal**, Shwetha S, Anandhavelu Natarajan, Aparna Garimella, Sambaran Bandyopadhyay, and **Jordan Boyd-Graber**. **Presentations by the People, for the People: Harnessing LLMs for Generating Persona-Aware Slides from Documents**. *European Association for Computational Linguistics*, 2024 (21% Acceptance Rate).
29. **Zongxia Li**, Andrew Mao, Daniel Kofi Stephens, Pranav Goel, Emily Walpole, Juan Francisco Fung, Alden Dima, and **Jordan Boyd-Graber**. **TENOR: Topic Enabled Neural Organization and Recommendation: Evaluating Topic Models in Task Based Settings**. *European Association for Computational Linguistics*, 2024 (21% Acceptance Rate).
30. Alvin Grissom II, Jo Shoemaker, Benjamin Goldman, Ruikang Shi, Craig Stewart, C. Anton Rytting, Leah Findlater, **Jordan Boyd-Graber**, Wenyan Li, Alvin Grissom II, and **Jordan Boyd-Graber**. **Rapidly Piloting Real-time Linguistic Assistance for Simultaneous Interpreters with Untrained Bilingual Surrogates**. *Linguistic Resources and Evaluation Conference*, 2024, 8 pages.
31. Chenglei Si, Navita Goyal, Tongshuang Wu, Chen Zhao, Shi Feng, Hal Daumé III, and **Jordan Boyd-Graber**. **Large Language Models Help Humans Verify Truthfulness—Except When They Are Convincingly Wrong**. *North American Association for Computational Linguistics*, 2024.
32. **Neha Pundlik Srikanth**, Rupak Sarkar, Mane, Heran Y., Aparicio, Elizabeth M., Nguyen, Quynh C., Rachel Rudinger, and **Jordan Boyd-Graber**. **Pregnant Questions: The Importance of Pragmatic Awareness in Maternal Health Question Answering**. *North American Association for Computational Linguistics*, 2024.
33. Anna Rogers, Marzena Karpinska, **Jordan Boyd-Graber**, and Naoaki Okazaki. **Program Chairs' Report on Peer Review at ACL 2023**. *Association for Computational Linguistics*, 2023, 33 pages.
34. **HyoJung Han**, Marine Carpuat, and **Jordan Boyd-Graber**. **Automatic Explicitation to Bridge the Background Knowledge Gap in Translation and its Evaluation with Multilingual QA**. *Empirical Methods in Natural Language Processing*, 2023, 9 pages (23% Acceptance Rate).
35. Sander V Schulhoff, Jeremy Pinto, Anaum Khan, Louis-François Bouchard, **Chenglei Si**, **Jordan Boyd-Graber**, Svetlana Anati, Valen Tagliabue, Anson Liu Kost, and Christopher R Carnahan. **Ignore This Title and Hack-APrompt: Exposing Systemic Vulnerabilities of LLMs Through a Global Prompt Hacking Competition**. *Empirical Methods in Natural Language Processing*, 2023, 9 pages (23% Acceptance Rate).
36. **Yoo Yeon Sung**, Naeemul Hassan, and **Jordan Boyd-Graber**. **Not all Fake News is Written: A Dataset and Analysis of Misleading Video Headlines**. *Empirical Methods in Natural Language Processing*, 2023, 9 pages (23% Acceptance Rate).
37. **Chenglei Si**, Zhe Gan, Zhengyuan Yang, Shuohang Wang, Jianfeng Wang, **Jordan Boyd-Graber**, and Lijuan Wang. **Prompting GPT-3 To Be Reliable**. *International Conference on Learning Representations*, 2023.

38. [Michelle Yuan](#), Patrick Xia, Chandler May, Benjamin Van Durme, and **Jordan Boyd-Graber**. **Adapting Coreference Resolution Models through Active Learning**. *Association for Computational Linguistics*, 2022, 9 pages (21% Acceptance Rate).
39. Yoshinari Fujinuma, **Jordan Boyd-Graber**, and Katharina Kann. **How Does Multilingual Pretraining Affect Cross-Lingual Transferability?**. *Association for Computational Linguistics*, 2022, 9 pages (21% Acceptance Rate).
40. Shi Feng and **Jordan Boyd-Graber**. **Learning to Explain Selectively: A Case Study on Question Answering**. *Empirical Methods in Natural Language Processing*, 2022, 9 pages.
41. [HyoJung Han](#), Marine Carpuat, and **Jordan Boyd-Graber**. **SimQA: Detecting Simultaneous MT Errors through Word-by-Word Question Answering**. *Empirical Methods in Natural Language Processing*, 2022, 9 pages.
42. [Pedro Rodriguez](#), Joe Barrow, Alexander Hoyle, John P. Lalor, Robin Jia, and **Jordan Boyd-Graber**. **Evaluation Examples Are Not Equally Informative: How Should That Change NLP Leaderboards?**. *Association for Computational Linguistics*, 2021, 9 pages (21% Acceptance Rate).
43. [Chen Zhao](#), Chenyan Xiong, **Jordan Boyd-Graber**, and Hal Daumé III. **Distantly-Supervised Dense Retrieval Enables Open-Domain Question Answering without Evidence Annotation**. *Empirical Methods in Natural Language Processing*, 2021, 11 pages (26% Acceptance Rate).
44. [Pedro Rodriguez](#) and **Jordan Boyd-Graber**. **Evaluation Paradigms in Question Answering**. *Empirical Methods in Natural Language Processing*, 2021, 5 pages (18% Acceptance Rate).
45. Maharshi Gor, Kellie Webster, and **Jordan Boyd-Graber**. **Toward Deconfounding the Influence of Subject’s Demographic Characteristics in Question Answering**. *Empirical Methods in Natural Language Processing*, 2021, 9 pages (26% Acceptance Rate).
46. [Chenglei Si](#), [Chen Zhao](#), and **Jordan Boyd-Graber**. **What’s in a Name? Answer Equivalence For Open-Domain Question Answering**. *Empirical Methods in Natural Language Processing*, 2021, 6 pages (18% Acceptance Rate).
47. Alexander Hoyle, Pranav Goel, [Denis Peskov](#), Andrew Hian-Cheong, **Jordan Boyd-Graber**, and Philip Resnik. **Is Automated Topic Model Evaluation Broken?: The Incoherence of Coherence**. *Neural Information Processing Systems*, 2021, 10 pages (26% Acceptance Rate).
48. Julian Martin Eisenschlos, Bhuwan Dhingra, Jannis Bulian, Benjamin Börschinger, and **Jordan Boyd-Graber**. **Fool Me Twice: Entailment from Wikipedia Gamification**. *North American Association for Computational Linguistics*, 2021, 10 pages (28% Acceptance Rate).
49. [Chen Zhao](#), Chenyan Xiong, Hal Daumé III, and **Jordan Boyd-Graber**. **Multi-Step Reasoning Over Unstructured Text with Beam Dense Retrieval**. *North American Association for Computational Linguistics*, 2021, 6 pages (23% Acceptance Rate).
50. [Chen Zhao](#), Chenyan Xiong, Xin Qian, and **Jordan Boyd-Graber**. **Complex Factoid Question Answering with a Free-Text Knowledge Graph**. *ACM International Conference on World Wide Web*, 2020, 11 pages (19.2% Acceptance Rate).
51. [Denis Peskov](#), Benny Cheng, [Ahmed Elgohary Ghoneim](#), Joe Barrow, Cristian Danescu-Niculescu-Mizil, and **Jordan Boyd-Graber**. **It Takes Two to Lie: One to Lie and One to Listen**. *Association for Computational Linguistics*, 2020 (25.4% Acceptance Rate).
52. **Jordan Boyd-Graber** and Benjamin Börschinger. **What Question Answering can Learn from Trivia Nerds**. *Association for Computational Linguistics*, 2020 (25.4% Acceptance Rate).
53. [Mozhi Zhang](#), Yoshinari Fujinuma, Michael J. Paul, and **Jordan Boyd-Graber**. **Why Overfitting Isn’t Always Bad: Retrofitting Cross-Lingual Word Embeddings to Dictionaries**. *Association for Computational Linguistics*, 2020 (17.6% Acceptance Rate).
54. [Mozhi Zhang](#), [Yoshinari Fujinuma](#), and **Jordan Boyd-Graber**. **Exploiting Cross-Lingual Subword Similarities in Low-Resource Document Classification**. *Association for the Advancement of Artificial Intelligence*, 2020, 9 pages (20.6% Acceptance Rate).
55. [Alison Smith](#), **Jordan Boyd-Graber**, Ron Fan, Melissa Birchfield, Tongshuang Wu, Dan Weld, and Leah Findlater. **No Explainability without Accountability: An Empirical Study of Explanations and Feedback in Interactive ML**. *Computer-Human Interaction*, 2020 (24% Acceptance Rate).
56. [Michelle Yuan](#), Hsuan-Tien Lin, and **Jordan Boyd-Graber**. **Cold-start Active Learning through Self-Supervised Language Modeling**. *Empirical Methods in Natural Language Processing*, 2020, 9 pages (25% Acceptance Rate).
57. [Michelle Yuan](#), [Mozhi Zhang](#), Benjamin Van Durme, Leah Findlater, and **Jordan Boyd-Graber**. **Interactive Refinement of Cross-Lingual Word Embeddings**. *Empirical Methods in Natural Language Processing*, 2020, 9 pages (25% Acceptance Rate).
58. [Alison Smith](#), Varun Kumar, **Jordan Boyd-Graber**, Kevin Seppi, and Leah Findlater. **Digging into User Control: Perceptions of Adherence and Instability in Transparent Models**. *Intelligent User Interfaces*, 2020, 12 pages (23% Acceptance Rate).
59. [Fenfei Guo](#), **Jordan Boyd-Graber**, Mohit Iyyer, and Leah Findlater. **Which Evaluations Uncover Sense Representations that Actually Make Sense?**. *Linguistic Resources and Evaluation Conference*, 2020, 10 pages.

60. [Yoshinari Fujinuma](#), Michael Paul, and **Jordan Boyd-Graber**. **A Resource-Free Evaluation Metric for Cross-Lingual Word Embeddings Based on Graph Modularity**. *Association for Computational Linguistics*, 2019, 10 pages (26% Acceptance Rate).
61. [Mozhi Zhang](#), Keyulu Xu, Ken-ichi Kawarabayashi, Stefanie Jegelka, and **Jordan Boyd-Graber**. **Are Girls Neko or Shōjo? Cross-Lingual Alignment of Non-Isomorphic Embeddings with Iterative Normalization**. *Association for Computational Linguistics*, 2019 (18.3% Acceptance Rate).
62. Jeffrey Lund, Piper Armstrong, Wilson Fearn, Stephen Cowley, Courtnei Byun, **Jordan Boyd-Graber**, and Kevin Seppi. **Automatic and Human Evaluation of Local Topic Quality**. *Association for Computational Linguistics*, 2019, 10 pages (26% Acceptance Rate).
63. [Eric Wallace](#), [Shi Feng](#), and **Jordan Boyd-Graber**. **Misleading Failures of Partial-input Baselines**. *Association for Computational Linguistics*, 2019, 6 pages (18% Acceptance Rate).
64. Varun Kumar, [Alison Smith](#), Leah Findlater, Kevin Seppi, and **Jordan Boyd-Graber**. **Why Didn't You Listen to Me? Comparing User Control of Human-in-the-Loop Topic Models**. *Association for Computational Linguistics*, 2019, 6 pages (18% Acceptance Rate).
65. [Denis Peskov](#), Joe Barrow, [Pedro Rodriguez](#), Graham Neubig, and **Jordan Boyd-Graber**. **Mitigating Noisy Inputs for Question Answering**. *Conference of the International Speech Communication Association*, 2019, 5 pages.
66. [Weiwei Yang](#), **Jordan Boyd-Graber**, and Philip Resnik. **A Multilingual Topic Model for Learning Weighted Topic Links Across Incomparable Corpora**. *Empirical Methods in Natural Language Processing*, 2019, 6 pages.
67. [Ahmed Elgohary Ghoneim](#), [Denis Peskov](#), and **Jordan Boyd-Graber**. **Can You Unpack That? Learning to Rewrite Questions-in-Context**. *Empirical Methods in Natural Language Processing*, 2019, 6 pages.
68. [Shi Feng](#) and **Jordan Boyd-Graber**. **What AI can do for me: Evaluating Machine Learning Interpretations in Cooperative Play**. *Intelligent User Interfaces*, 2019 (25% Acceptance Rate).
69. Craig Stewart, Nikolai Vogler, Junjie Hu, **Jordan Boyd-Graber**, and Graham Neubig. **Automatic Estimation of Simultaneous Interpreter Performance**. *Association for Computational Linguistics*, 2018, 5 pages (24% Acceptance Rate).
70. [Ahmed Elgohary Ghoneim](#), [Chen Zhao](#), and **Jordan Boyd-Graber**. **Dataset and Baselines for Sequential Open-Domain Question Answering**. *Empirical Methods in Natural Language Processing*, 2018, 6 pages (23% Acceptance Rate).
71. [Shi Feng](#), [Eric Wallace](#), Alvin Grissom II, [Pedro Rodriguez](#), Mohit Iyyer, and **Jordan Boyd-Graber**. **Pathologies of Neural Models Make Interpretation Difficult**. *Empirical Methods in Natural Language Processing*, 2018, 9 pages (26% Acceptance Rate).
72. [Alison Smith](#), Varun Kumar, **Jordan Boyd-Graber**, Kevin Seppi, and Leah Findlater. **User-Centered Design and Evaluation of a Human-in-the-Loop Topic Modeling System**. *Intelligent User Interfaces*, 2018, 12 pages (23% Acceptance Rate).
73. Paul Felt, Eric Ringger, Kevin Seppi, and **Jordan Boyd-Graber**. **Learning from Measurements in Crowdsourcing Models: Inferring Ground Truth from Diverse Annotation Types**. *International Conference on Computational Linguistics*, 2018, 10 pages (37% Acceptance Rate).
74. [Michelle Yuan](#), Benjamin Van Durme, and **Jordan Boyd-Graber**. **Multilingual Anchoring: Interactive Topic Modeling and Alignment Across Languages**. *Neural Information Processing Systems*, 2018, 10 pages (21% Acceptance Rate).
75. Mohit Iyyer, Varun Manjunatha, **Jordan Boyd-Graber**, and Larry Davis. **Learning to Color from Language**. *North American Association for Computational Linguistics*, 2018, 6 pages (29% Acceptance Rate).
76. Shudong Hao, Michael J. Paul, and **Jordan Boyd-Graber**. **Lessons from the Bible on Modern Topics: Multilingual Topic Model Evaluation on Low-Resource Languages**. *North American Association for Computational Linguistics*, 2018, 9 pages (35% Acceptance Rate).
77. Jeff Lund, Connor Cook, Kevin Seppi, and **Jordan Boyd-Graber**. **Tandem Anchoring: A Multiword Anchor Approach for Interactive Topic Modeling**. *Association for Computational Linguistics*, 2017, 10 pages (22% Acceptance Rate).
78. [Mohit Iyyer](#), Varun Manjunatha, Anupam Guha, Yogarshi Vyas, **Jordan Boyd-Graber**, Hal Daumé III, and Larry Davis. **The Amazing Mysteries of the Gutter: Drawing Inferences Between Panels in Comic Book Narratives**. *Computer Vision and Pattern Recognition*, 2017, 10 pages (30% Acceptance Rate).
79. [Weiwei Yang](#), **Jordan Boyd-Graber**, and Philip Resnik. **Adapting Topic Models using Lexical Associations with Tree Priors**. *Empirical Methods in Natural Language Processing*, 2017, 6 pages (18% Acceptance Rate).
80. Khanh Nguyen, **Jordan Boyd-Graber**, and Hal Daumé III. **Reinforcement Learning for Bandit Neural Machine Translation with Simulated Human Feedback**. *Empirical Methods in Natural Language Processing*, 2017, 11 pages (26% Acceptance Rate).

81. You Lu, Jeff Lund, and **Jordan Boyd-Graber**. **Why ADAGRAD Fails for Online Topic Modeling**. *Empirical Methods in Natural Language Processing*, 2017, 6 pages (18% Acceptance Rate).
82. **Weiwei Yang**, **Jordan Boyd-Graber**, and Philip Resnik. **A Discriminative Topic Model using Document Network Structure**. *Association for Computational Linguistics*, 2016, 10 pages (28% Acceptance Rate).
83. **Forough Poursabzi-Sangdeh**, **Jordan Boyd-Graber**, Leah Findlater, and Kevin Seppi. **ALTO: Active Learning with Topic Overviews for Speeding Label Induction and Document Labeling**. *Association for Computational Linguistics*, 2016 (28% Acceptance Rate).
84. Hadi Amiri, Philip Resnik, **Jordan Boyd-Graber**, and Hal Daumé III. **Learning Text Pair Similarity with Context-sensitive Autoencoders**. *Association for Computational Linguistics*, 2016 (28% Acceptance Rate).
85. **Alvin Grissom II**, Naho Orita, and **Jordan Boyd-Graber**. **Incremental Prediction of Sentence-final Verbs**. *Conference on Computational Natural Language Learning*, 2016, 10 pages (20% Acceptance Rate).
86. **He He**, **Jordan Boyd-Graber**, Kevin Kwok, and Hal Daumé III. **Opponent Modeling in Deep Reinforcement Learning**. *International Conference on Machine Learning*, 2016, 10 pages (24% Acceptance Rate).
87. Md Arafat Sultan, **Jordan Boyd-Graber**, and Tamara Sumner. **Bayesian Supervised Domain Adaptation for Short Text Similarity**. *North American Association for Computational Linguistics*, 2016, 11 pages (24% Acceptance Rate).
88. **Mohit Iyyer**, Anupam Guha, Snigdha Chaturvedi, **Jordan Boyd-Graber**, and Hal Daumé III. **Feuding Families and Former Friends: Unsupervised Learning for Dynamic Fictional Relationships**. *North American Association for Computational Linguistics*, 2016, 11 pages (24% Acceptance Rate).
89. **He He**, **Jordan Boyd-Graber**, and Hal Daumé III. **Interpretese vs. Translationese: The Uniqueness of Human Strategies in Simultaneous Interpretation**. *North American Association for Computational Linguistics*, 2016, 6 pages (29% Acceptance Rate).
90. Daniel Peterson, **Jordan Boyd-Graber**, Martha Palmer, and Daisuke Kawahara. **Leveraging VerbNet to build Corpus-Specific Verb Clusters**. *Proceedings of the Fifth Joint Conference on Lexical and Computational Semantics*, 2016, 5 pages.
91. **Mohit Iyyer**, Varun Manjunatha, **Jordan Boyd-Graber**, and Hal Daumé III. **Deep Unordered Composition Rivals Syntactic Methods for Text Classification**. *Association for Computational Linguistics*, 2015, 11 pages (25% Acceptance Rate).
92. Vlad Niculae, Srikan Kumar, **Jordan Boyd-Graber**, and Cristian Danescu-Niculescu-Mizil. **Linguistic Harbingers of Betrayal: A Case Study on an Online Strategy Game**. *Association for Computational Linguistics*, 2015, 10 pages (25% Acceptance Rate).
93. **Viet-An Nguyen**, **Jordan Boyd-Graber**, Philip Resnik, and Kristina Miler. **Tea Party in the House: A Hierarchical Ideal Point Topic Model and Its Application to Republican Legislators in the 112th Congress**. *Association for Computational Linguistics*, 2015, 11 pages (25% Acceptance Rate).
94. Paul Felt, Eric Ringger, **Jordan Boyd-Graber**, and Kevin Seppi. **Making the Most of Crowdsourced Document Annotations: Confused Supervised LDA**. *Conference on Computational Natural Language Learning*, 2015, 10 pages (30% Acceptance Rate).
95. **Weiwei Yang**, **Jordan Boyd-Graber**, and Philip Resnik. **Birds of a Feather Linked Together: A Discriminative Topic Model using Link-based Priors**. *Empirical Methods in Natural Language Processing*, 2015, 5 pages (28% Acceptance Rate).
96. Yi Yang, Doug Downey, and **Jordan Boyd-Graber**. **Efficient Methods for Incorporating Knowledge into Topic Models**. *Empirical Methods in Natural Language Processing*, 2015, 9 pages (24% Acceptance Rate).
97. **He He**, **Alvin Grissom II**, **Jordan Boyd-Graber**, and Hal Daumé III. **Syntax-based Rewriting for Simultaneous Machine Translation**. *Empirical Methods in Natural Language Processing*, 2015 (24% Acceptance Rate).
98. Stephen H. Bach, Bert Huang, **Jordan Boyd-Graber**, and Lise Getoor. **Paired-Dual Learning for Fast Training of Latent Variable Hinge-Loss MRFs**. *International Conference on Machine Learning*, 2015, 10 pages (20% Acceptance Rate).
99. **Jordan Boyd-Graber**, **Mohit Iyyer**, **He He**, and Hal Daumé III. **Interactive Incremental Question Answering**. *Neural Information Processing Systems*, 2015.
100. **Thang Nguyen**, **Jordan Boyd-Graber**, Jeff Lund, Kevin Seppi, and Eric Ringger. **Is your anchor going up or down? Fast and accurate supervised topic models**. *North American Association for Computational Linguistics*, 2015, 10 pages (26% Acceptance Rate).
101. Anupam Guha, **Mohit Iyyer**, Danny Bouman, and **Jordan Boyd-Graber**. **Removing the Training Wheels: A Coreference Dataset that Entertains Humans and Challenges Computers**. *North American Association for Computational Linguistics*, 2015, 11 pages (26% Acceptance Rate).
102. **Thang Nguyen**, **Yuening Hu**, and **Jordan Boyd-Graber**. **Anchors Regularized: Adding Robustness and Extensibility to Scalable Topic-Modeling Algorithms**. *Association for Computational Linguistics*, 2014, 10 pages (26% Acceptance Rate).

103. [Mohit Iyyer](#), Peter Enns, [Jordan Boyd-Graber](#), and Philip Resnik. **Political Ideology Detection Using Recursive Neural Networks**. *Association for Computational Linguistics*, 2014, 10 pages (26% Acceptance Rate).
104. [Yuening Hu](#), [Ke Zhai](#), Vlad Eidelman, and [Jordan Boyd-Graber](#). **Polylingual Tree-Based Topic Models for Translation Domain Adaptation**. *Association for Computational Linguistics*, 2014, 11 pages (26% Acceptance Rate).
105. [Mohit Iyyer](#), [Jordan Boyd-Graber](#), Leonardo Claudino, Richard Socher, and Hal Daumé III. **A Neural Network for Factoid Question Answering over Paragraphs**. *Empirical Methods in Natural Language Processing*, 2014, 12 pages (26% Acceptance Rate).
106. [Alvin Grissom II](#), [He He](#), [Jordan Boyd-Graber](#), John Morgan, and Hal Daumé III. **Don't Until the Final Verb Wait: Reinforcement Learning for Simultaneous Machine Translation**. *Empirical Methods in Natural Language Processing*, 2014, 11 pages (30% Acceptance Rate).
107. [Viet-An Nguyen](#), [Jordan Boyd-Graber](#), and Philip Resnik. **Sometimes Average is Best: The Importance of Averaging for Prediction using MCMC Inference in Topic Modeling**. *Empirical Methods in Natural Language Processing*, 2014, 6 pages (30% Acceptance Rate).
108. [Viet-An Nguyen](#), [Jordan Boyd-Graber](#), Philip Resnik, and Jonathan Chang. **Learning a Concept Hierarchy from Multi-labeled Documents**. *Neural Information Processing Systems*, 2014, 9 pages (25% Acceptance Rate).
109. Kimberly Glasgow, Clay Fink, and [Jordan Boyd-Graber](#). **Our grief is unspeakable: Measuring the community impact of a tragedy**. *The International AAAI Conference on Weblogs and Social Media*, 2014, 9 pages (20% Acceptance Rate).
110. [Jordan Boyd-Graber](#), Kimberly Glasgow, and Jackie Sauter Zajac. **Spoiler Alert: Machine Learning Approaches to Detect Social Media Posts with Revelatory Information**. *ASIST 2013: The 76th Annual Meeting of the American Society for Information Science and Technology*, 2013, 9 pages.
111. [Ke Zhai](#) and [Jordan Boyd-Graber](#). **Online Topic Models with Infinite Vocabulary**. *International Conference on Machine Learning*, 2013, 9 pages (20% Acceptance Rate).
112. [Yuening Hu](#), [Jordan Boyd-Graber](#), Hal Daumé III, and Z. Irene Ying. **Binary to Bushy: Bayesian Hierarchical Clustering with the Beta Coalescent**. *Neural Information Processing Systems*, 2013, 9 pages (25% Acceptance Rate).
113. [Viet-An Nguyen](#), [Jordan Boyd-Graber](#), and Philip Resnik. **Lexical and Hierarchical Topic Regression**. *Neural Information Processing Systems*, 2013, 10 pages (25% Acceptance Rate).
114. [Viet-An Nguyen](#), [Yuening Hu](#), [Jordan Boyd-Graber](#), and Philip Resnik. **Argviz: Interactive Visualization of Topic Dynamics in Multi-party Conversations**. *North American Association for Computational Linguistics*, 2013, 4 pages (50% Acceptance Rate).
115. Naho Orita, Rebecca McKeown, Naomi H. Feldman, Jeffrey Lidz, and [Jordan Boyd-Graber](#). **Discovering Pronoun Categories using Discourse Information**. *Proceedings of the Cognitive Science Society*, 2013, 6 pages.
116. [Ke Zhai](#), [Jordan Boyd-Graber](#), Nima Asadi, and Mohamad (Jude) Alkhouja. **Mr. LDA: A Flexible Large Scale Topic Modeling Package using Variational Inference in MapReduce**. *ACM International Conference on World Wide Web*, 2012, 10 pages (12% Acceptance Rate).
117. [Yuening Hu](#) and [Jordan Boyd-Graber](#). **Efficient Tree-Based Topic Modeling**. *Association for Computational Linguistics*, 2012, 5 pages (21% Acceptance Rate).
118. Vladimir Eidelman, [Jordan Boyd-Graber](#), and Philip Resnik. **Topic Models for Dynamic Translation Model Adaptation**. *Association for Computational Linguistics*, 2012, 5 pages (21% Acceptance Rate).
119. [Viet-An Nguyen](#), [Jordan Boyd-Graber](#), and Philip Resnik. **SITS: A Hierarchical Nonparametric Model using Speaker Identity for Topic Segmentation in Multiparty Conversations**. *Association for Computational Linguistics*, 2012, 10 pages (19% Acceptance Rate).
120. [Jordan Boyd-Graber](#), [Brianna Satinoff](#), [He He](#), and Hal Daumé III. **Besting the Quiz Master: Crowdsourcing Incremental Classification Games**. *Empirical Methods in Natural Language Processing*, 2012, 12 pages (25% Acceptance Rate).
121. [Yuening Hu](#), [Ke Zhai](#), Sinead Williamson, and [Jordan Boyd-Graber](#). **Modeling Images using Transformed Indian Buffet Processes**. *International Conference on Machine Learning*, 2012, 8 pages (27% Acceptance Rate).
122. Asad B. Sayeed, [Jordan Boyd-Graber](#), Bryan Rusk, and Amy Weinberg. **Grammatical structures for word-level sentiment detection**. *North American Association for Computational Linguistics*, 2012, 10 pages (31% Acceptance Rate).
123. [Yuening Hu](#), [Jordan Boyd-Graber](#), and [Brianna Satinoff](#). **Interactive Topic Modeling**. *Association for Computational Linguistics*, 2011, 10 pages (25% Acceptance Rate).
124. Clay Templeton, Kenneth R. Fleischmann, and [Jordan Boyd-Graber](#). **Simulating Audiences: Automating Analysis of Values, Attitudes, and Sentiment**. *IEEE International Conference on Social Computing*, 2011, 4 pages (10% Acceptance Rate).
125. Clay Templeton, Kenneth R. Fleischmann, and [Jordan Boyd-Graber](#). **Comparing Values and Sentiment Using Mechanical Turk**. *iConference*, 2011, 2 pages.

126. Kenneth R. Fleischmann, Clay Templeton, and **Jordan Boyd-Graber**. **Modeling Diverse Standpoints in Text Classification: Learning to Be Human by Modeling Human Values**. *iConference*, 2011, 2 pages.
127. **Jordan Boyd-Graber** and Philip Resnik. **Holistic Sentiment Analysis Across Languages: Multilingual Supervised Latent Dirichlet Allocation**. *Empirical Methods in Natural Language Processing*, 2010, 11 pages (25% Acceptance Rate).
128. Eric Hardisty, **Jordan Boyd-Graber**, and Philip Resnik. **Modeling Perspective using Adaptor Grammars**. *Empirical Methods in Natural Language Processing*, 2010, 10 pages (25% Acceptance Rate).
129. Sonya S. Nikolova, **Jordan Boyd-Graber**, Christiane Fellbaum, and Perry Cook. **Better Vocabularies for Assistive Communication Aids: Connecting Terms using Semantic Networks and Untrained Annotators**. *ACM Conference on Computers and Accessibility*, 2009, 8 pages (31% Acceptance Rate).
130. Xiaojuan Ma, **Jordan Boyd-Graber**, Sonya S. Nikolova, and Perry Cook. **Speaking Through Pictures: Images vs. Icons**. *ACM Conference on Computers and Accessibility*, 2009, 8 pages (31% Acceptance Rate).
131. Jonathan Chang, **Jordan Boyd-Graber**, and David M. Blei. **Connections between the Lines: Augmenting Social Networks with Text**. *Knowledge Discovery and Data Mining*, 2009, 9 pages (9% Acceptance Rate).
132. Jonathan Chang, **Jordan Boyd-Graber**, Chong Wang, Sean Gerrish, and David M. Blei. **Reading Tea Leaves: How Humans Interpret Topic Models**. *Neural Information Processing Systems*, 2009, 9 pages (24% Acceptance Rate).
133. **Jordan Boyd-Graber** and David M. Blei. **Multilingual Topic Models for Unaligned Text**. *Uncertainty in Artificial Intelligence*, 2009, 8 pages (31% Acceptance Rate).
134. **Jordan Boyd-Graber** and David M. Blei. **Syntactic Topic Models**. *Neural Information Processing Systems*, 2008, 8 pages (25% Acceptance Rate).
135. **Jordan Boyd-Graber**, David M. Blei, and Xiaojin Zhu. **A Topic Model for Word Sense Disambiguation**. *Empirical Methods in Natural Language Processing*, 2007, 10 pages (27% Acceptance Rate).
136. **Jordan Boyd-Graber**, Sonya S. Nikolova, Karyn A. Moffatt, Kenrick C. Kin, Joshua Y. Lee, Lester W. Mackey, Marilyn M. Tremaine, and Maria M. Klawe. **Participatory design with proxies: Developing a desktop-PDA system to support people with aphasia**. *Computer-Human Interaction*, 2006, 10 pages (23% Acceptance Rate).
137. **Jordan Boyd-Graber**, Christiane Fellbaum, Daniel Osherson, and Robert Schapire. **Adding Dense, Weighted, Connections to WordNet**. *Proceedings of the Global WordNet Conference*, 2006, 10 pages.

Refereed Workshops

1. Peter Jansen and **Jordan Boyd-Graber**. **Picard understanding Darmok: A Dataset and Model for Metaphor-Rich Translation in a Constructed Language**. *Figurative Language Workshop 2022 @EMNLP*, 2022.
2. Diggelmann, Thomas, **Boyd-Graber, Jordan**, Bulian, Jannis, Ciaramita, Massimiliano, and Leippold, Markus. **CLIMATE-FEVER: A Dataset for Verification of Real-World Climate Claims**. *NIPS Workshop on Tackling Climate Change with Machine Learning*, 2020.
3. Francesco Saverio Varini, **Jordan Boyd-Graber**, Massimiliano Ciaramita, and Markus Leippold. **ClimaText: A Dataset for Climate Change Topic Detection**. *NeurIPS Workshop on Tackling Climate Change with Machine Learning*, 2020.
4. [Eric Wallace](#) and **Jordan Boyd-Graber**. **Trick Me If You Can: Adversarial Writing of Trivia Challenge Questions**. *ACL Student Research Workshop*, 2018, 6 pages.
5. [Mozhi Zhang](#), [Yoshinari Fujinuma](#), and **Jordan Boyd-Graber**. **Exploiting Cross-Lingual Subword Similarities in Low-Resource Document Classification**. *ACL Workshop on Deep Learning Approaches for Low-Resource Natural Language Processing*, 2018, 6 pages.
6. [Shi Feng](#), [Eric Wallace](#), and **Jordan Boyd-Graber**. **Interpreting Neural Networks with Nearest Neighbors**. *EMNLP Workshop on BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, 2018, 7 pages.
7. [Alison Smith](#), Varun Kumar, **Jordan Boyd-Graber**, Kevin Seppi, and Leah Findlater. **Accounting for Input Uncertainty in Human-in-the-Loop Systems**. *CHI 2017 Designing for Uncertainty Workshop*, 2017.
8. [Alison Smith](#), Tak Yeon Lee, [Forough Poursabzi-Sangdeh](#), **Jordan Boyd-Graber**, Kevin Seppi, Niklas Elmqvist, and Leah Findlater. **Human-Centered and Interactive: Expanding the Impact of Topic Models**. *CHI Human Centred Machine Learning Workshop*, 2016.
9. [Weiwei Yang](#), **Jordan Boyd-Graber**, and Philip Resnik. **Birds of a Feather in the Same Nest: A Discriminative Topic Model using Block-based Priors**. *Mid-Atlantic Student Colloquium on Speech, Language, and Learning*, 2016.
10. Anupam Guha, [Mohit Iyyer](#), and **Jordan Boyd-Graber**. **A Distorted Skull Lies in the Bottom Center: Identifying Paintings from Text Descriptions**. *NAACL Human-Computer Question Answering Workshop*, 2016.
11. [Forough Poursabzi-Sangdeh](#) and **Jordan Boyd-Graber**. **Speeding Document Annotation with Topic Models**. *NAACL Student Research Workshop*, 2015.
12. Philip Resnik, William Armstrong, Leonardo Claudino, [Thang Nguyen](#), [Viet-An Nguyen](#), and **Jordan Boyd-Graber**. **Beyond LDA: Exploring Supervised Topic Modeling for Depression-Related Language in Twitter**. *NAACL Workshop on Cognitive Modeling and Computational Linguistics*, 2015.

13. Naho Orita, Naomi Feldman, and **Jordan Boyd-Graber**. **Quantifying the role of discourse topicality in speakers' choices of referring expressions**. *ACL Workshop on Cognitive Modeling and Computational Linguistics*, 2014.
14. [Alison Smith](#), Jason Chuang, [Yuening Hu](#), **Jordan Boyd-Graber**, and Leah Findlater. **Concurrent Visualization of Relationships between Words and Topics in Topic Models**. *ACL Workshop on Workshop on Interactive Language Learning, Visualization, and Interfaces*, 2014.
15. [Ke Zhai](#), **Jordan Boyd-Graber**, and Shay B. Cohen. **Hybrid Online Inference with Adaptor Grammars**. *NIPS Workshop on Advances in Variational Inference*, 2014.
16. Jason Chuang, John D. Wilkerson, Rebecca Weiss, Dustin Tingley, Brandon M. Stewart, Margaret E. Roberts, [Forough Poursabzi-Sangdeh](#), Justin Grimmer, Leah Findlater, **Jordan Boyd-Graber**, and Jeffrey Heer. **Computer-Assisted Content Analysis: Topic Models for Exploring Multiple Subjective Interpretations**. *NIPS Workshop on Human-Propelled Machine Learning*, 2014.
17. [Mohit Iyyer](#), **Jordan Boyd-Graber**, and Hal Daumé III. **Generating Sentences from Semantic Vector Space Representations**. *NIPS Workshop on Learning Semantics*, 2014.
18. Thang Nguyen, [Yuening Hu](#), and **Jordan Boyd-Graber**. **Evaluating Regularized Anchor Words**. *NIPS Workshop on Topic Models: Computation, Application, and Evaluation*, 2013.
19. [Yuening Hu](#), [Ke Zhai](#), Vlad Edelman, and **Jordan Boyd-Graber**. **Topic Models for Translation Domain Adaptation**. *NIPS Workshop on Topic Models: Computation, Application, and Evaluation*, 2013.
20. [Viet-An Nguyen](#), **Jordan Boyd-Graber**, Jonathan Chang, and Philip Resnik. **Tree-Based Label Dependency Topic Models**. *NIPS Workshop on Topic Models: Computation, Application, and Evaluation*, 2013.
21. [Yuening Hu](#) and **Jordan Boyd-Graber**. **Suggesting Constraints for Interactive Topic Modeling**. *ICML Workshop on Machine Learning in Human Computation and Crowdsourcing*, 2012.
22. [Yuening Hu](#) and **Jordan Boyd-Graber**. **Bayesian Hierarchical Clustering with Beta Coalescents**. *Mid-Atlantic Student Colloquium on Speech, Language, and Learning*, 2012.
23. [Ke Zhai](#) and **Jordan Boyd-Graber**. **Online Topic Model with Infinite Vocabulary**. *Mid-Atlantic Student Colloquium on Speech, Language, and Learning*, 2012.
24. [Viet-An Nguyen](#), **Jordan Boyd-Graber**, and Philip Resnik. **"I Want to Talk About, Again, My Record On Energy ...": Modeling Topic Control in Conversations using Speaker-centric Nonparametric Topic Models**. *Mid-Atlantic Student Colloquium on Speech, Language, and Learning*, 2012.
25. Clay Templeton, Travis Brown, Sayan Battacharyya, and **Jordan Boyd-Graber**. **Mining the Dispatch under Supervision: Using Casualty Counts to Guide Topics from the Richmond Daily Dispatch Corpus**. *Chicago Colloquium on Digital Humanities and Computer Science*, 2011, 7 pages.
26. **Jordan Boyd-Graber**. **Linguistic Resource Creation in a Web 2.0 World**. *NSF Workshop on Collaborative Annotation*, 2011, 7 pages.
27. Pranav Anand, Joseph King, **Jordan Boyd-Graber**, Earl Wagner, Craig Martell, Douglas W. Oard, and Philip Resnik. **Believe Me: We Can Do This!**. *The AAAI 2011 workshop on Computational Models of Natural Argument*, 2011, 5 pages.
28. [Brianna Satinoff](#) and **Jordan Boyd-Graber**. **Trivial Classification: What features do humans use for classification?**. *Workshop on Crowdsourcing Technologies for Language and Cognition Studies*, 2011.
29. Nitin Madnani, **Jordan Boyd-Graber**, and Philip Resnik. **Measuring Transitivity Using Untrained Annotators**. *Creating Speech and Language Data With Amazon's Mechanical Turk*, 2010, 6 pages.
30. Sonya S. Nikolova, **Jordan Boyd-Graber**, and Perry Cook. **The Design of ViVA: A Mixed-initiative Visual Vocabulary for Aphasia**. *Proceedings of the 27th international conference extended abstracts on Human factors in computing systems*, 2009, 6 pages.
31. Jonathan Chang, **Jordan Boyd-Graber**, and David M. Blei. **Discovering social networks from free text**. *3rd Annual Machine Learning Symposium*, 2008.
32. **Jordan Boyd-Graber** and David M. Blei. **Multilingual Topic Models**. *NIPS Workshop on Unsupervised Latent Variable Models*, 2008.
33. **Jordan Boyd-Graber** and David M. Blei. **PUTOP: Turning Predominant Senses into a Topic Model for WSD**. *4th International Workshop on Semantic Evaluations*, 2007, 5 pages.

Sponsored Research and Programs

Active Grants

Interactive Information Seeking

Investigators: Jordan Boyd-Graber et al.

Award: \$70,000 (Share: \$70,000)

2021–2025 (Adobe Corporation)

Rosie the Chatbot: Leveraging Automated and Personalized Health Information Communication to Reduce Disparities in Maternal and Child Health 09/2021–06/2026 (NIH)

Investigators: Quynh Nguyen (PI), Liz Aparicio (co-PI), Jordan Lee Boyd-Graber (co-PI), Xin He (co-PI)

Award: \$3,921,898³ (Share: \$980,475)

Collaborative Research: RI: Medium: Hard Data to the Model: Personalized, Diverse Preferences for Language Models 2024–2028 (NSF)

Investigators: Jordan Boyd-Graber (PI)

Award: \$1,200,000 (Share: \$385,000)

Collaboration: usc, Notre Dame, Haverford

DARPA-PA-23-04-02-FACT-FP-003, Complementing Human Intelligence to Recognize Opponent Narratives (CHIRON) 2024–2025 (DARPA)

Investigators: Jordan Boyd-Graber (PI)

Award: \$293,000 (Share: \$293,000)

Collaboration: usc

Completed Funding

Leaderboard and Competition for Human–Computer Adversarial Question Answering 2022–2024 (Meta)

Investigators: Jordan Boyd-Graber PI

Award: \$51,601 (Share: \$51,601)

ATAC—Foundational Data: Human-Computer Pipelines for Information Extraction for Terrorism-related News Articles 2023–2025 (DoD)

Investigators: Devin Hayes (PI), Michael Jensen (co-PI), Amy Pate (co-PI), Jordan Boyd-Graber (co-PI)

Award: \$133,298 (Share: \$66,650)

ARO BAA W911NF-23-S-0007, Strengthening analytic products through theories of Discourse, Sensemaking, and Nonmonotonic Interference 2024–2027 (IARPA)

Investigators: Rachel Rudinger (PI), Jordan Boyd-Graber (co-PI)

Award: \$Cancelled (Share: \$Negotiated, but not Awarded)

Collaboration: Johns Hopkins

ALLAN: Agents Learning Lying And Negotiation

01/2021–07/2023 (IARPA)

Investigators: Jordan Boyd-Graber PI

Award: \$984,859 (Share: \$250,000)

Collaboration: DARPA

BETTER: Multilingual Fine-grained Compositional Analysis

08/2019–02/2023 (IARPA)

Investigators: Jordan Boyd-Graber co-PI

Award: \$414,884 (Share: \$414,884)

Collaboration: Johns Hopkins

Collaborative Research: Interactive, Multilingual Representation Learning

9/2017–8/2023 (HLT COE)

Investigators: Jordan Boyd-Graber co-PI

Award: \$600,000 (Share: \$225,000)

CAREER: Human-Computer Cooperation for Word-by-Word Question Answering

2/2017–1/2023 (NSF)

Investigators: Jordan Boyd-Graber PI

Award: \$500,000 (Share: \$500,000)

Strategies for Investigating and Eliciting INformation from Nuanced Attackers (SIENNA) 10/2018–03/2020 (DARPA ASED via BBN)

Investigators: Jordan Boyd-Graber PI

Award: \$164,660 (Share: \$164,660)

RI: EAGER: Collaborative Research: Adaptive Heads-up Displays for Simultaneous Interpretation 10/2017–2/2019 (NSF)

Investigators: Jordan Boyd-Graber co-PI

Award: \$150,000 (Share: \$75,000)

Collaboration: Carnegie Mellon, University of Washington (only UMD portion shown)

CHS: Medium: Hyperlocal and Hypertemporal Information in Mass Emergencies Events: Next Generation Crisis Informatics Data Collection & Analytics 8/2016–8/2020 (NSF)

Investigators: Ken Anderson PI, Leysia Palen co-PI, and Jordan Boyd-Graber co-PI

Award: \$1,200,000 (Share: \$300,000)

Multilingual Interactive Topic Modeling

8/2015–7/2019 (DARPA LORELEI)

Investigators: Jordan Boyd-Graber (PI) and Mans Hulden (co-PI)

Award: \$426,654 (Share: \$325,000)

Collaboration: Raytheon BBN (prime) with University of Maryland, Johns Hopkins University, and University of Washington (only Colorado portion shown)

³\$221,898 supplement awarded in 2023 for increasing the robustness of the chatbot.

Closing the User-Model Loop for Understanding Topics in Large Document Collections⁴

8/2014–7/2018 (NSF)

Investigators: Jordan Boyd-Graber (PI) and Leah Findlater (CO-PI)

Award: \$650,000 (Share: \$325,000)

Collaboration: Brigham Young University and University of Maryland (only Maryland portion shown)

eTASC: Empirical Evidence for a Theoretical Approach to Semantic Components 12/2015–11/2018 (DTRA)

Investigators: Martha Palmer (PI), Laura Michaelis (CO-PI) and Jordan Boyd-Graber (CO-PI)

Award: \$1,250,000 (Share: \$200,000)

Collaboration: Brandeis, Princeton, Stanford

Temporal Relation Discovery for Clinical Text

9/2015–9/2018 (NIH)

Investigators: Martha Palmer (PI) and Jordan Boyd-Graber (CO-PI)

Award: \$531,328 (Share: \$76,896)

Collaboration: Harvard, University of Alabama (Only Colorado portion shown)

Scaling Insight into Science: Assessing the value and effectiveness of machine assisted classification within a statistical system

8/2014–7/2017 (NSF)

Investigators: Jordan Boyd-Graber (PI)

Award: \$195,000 (Share: \$195,000)

Collaboration: University of Chicago and American Institutes for Research (only Colorado portion shown)

Bayesian Thinking on Your Feet—Embedding Generative Models in Reinforcement Learning for Sequentially Revealed Data

8/2013–7/2016 (NSF)

Investigators: Jordan Boyd-Graber (PI) and Hal Daumé III (CO-PI)

Award: \$500,000 (Share: \$250,000)

Collaboration: Grant located at University of Maryland

Sentiment Analysis in Social Media: Political Spin and Cultural Biases

8/2013–8/2014 (CASL)

Investigators: Philip Resnik PI and Jordan Boyd-Graber CO-PI

Award: \$100,000 (Share: \$50,000)

Cross-Language Bayesian Models for Web-Scale Text Analysis

9/2009–8/2014 (NSF)

Investigators: Jimmy Lin (PI), Philip Resnik (CO-PI), Jordan Boyd-Graber⁵ (CO-PI)

Award: \$350,000 (Share: \$175,000)

Language Evidence for Social Goals

8/2009–10/2012 (IARPA)

Investigators: Philip Resnik (PI), Pranav Anand (CO-PI), Jordan Boyd-Graber (CO-PI), Deborah Cai (CO-PI), Craig Martell (CO-PI), Doug Oard (CO-PI), Marilyn Walker (CO-PI)

Award: \$1,454,439 (Share: \$100,000)

Center for Language and Cultural Analysis

9/2009–8/2012 (ARL)

Investigators: Amy Weinberg (PI), Jordan Boyd-Graber (CO-PI), Michele Gelfand (CO-PI), Philip Resnik (CO-PI, later PI)

Award: \$ 735,050 (Share: \$100,000)

Advanced Open Source Exploitation Models

4/2011–12/2011 (Lockheed Martin)

Investigators: Philip Resnik (PI), Jordan Boyd-Graber (CO-PI)

Award: \$60,000 (Share: \$30,000)

Social Media Scanning

5/2011–12/2011 (Optimal Solutions Group)

Investigators: Philip Resnik (PI), Jordan Boyd-Graber (CO-PI)

Award: \$29,849 (Share: \$14,925)

Teaching, Extension, Mentoring, and Advising

Courses Taught

CMSC 470: Natural Language Processing

UMD, Fall 2025

96 students

CMSC 723: Natural Language Processing

UMD, Fall 2024

90 students

CMSC 470: Natural Language Processing

UMD, Spring 2024

96 students

CMSC 848Q: How and Why Artificial Intelligence Answers Questions

UMD, Fall 2023

30 students

⁴After I moved to Colorado, Leah Findlater assumed the role of PI to enable a new subcontract to Colorado. Then, amusingly enough, Leah left for University of Washington and I returned to UMD; I again became PI. The original grant as awarded is provided here to minimize confusion.

⁵I wrote this grant while a postdoc working with Philip Resnik; the vast majority of the text and the entirety of the scientific ideas were my own. However, I could not serve as PI while still a postdoc. I was added to the grant as PI after it was awarded and served as sole research advisor to the students funded by the grant while the other PIs were both on sabbatical.

CMSC 470: Natural Language Processing 70 students	<i>UMD, Spring 2023</i>
CMSC 723: Graduate Natural Language Processing 80 students	<i>UMD, Fall 2022</i>
CMSC 848Q: How and Why Artificial Intelligence Answers Questions 60 students	<i>UMD, Spring 2022</i>
CMSC 470: Natural Language Processing 40 students	<i>UMD, Fall 2021</i>
CMSC 723: Computational Linguistics I 65 students	<i>UMD, Spring 2021</i>
INST 808: Quantitative Methods (Deep Learning) 8 students	<i>UMD, Fall 2020</i>
CMSC 470: Natural Language Processing 40 students, First Offering of Course with Permanent Number	<i>UMD, Spring 2019</i>
CMSC 723: Computational Linguistics I 60 students	<i>UMD, Fall 2018</i>
CMSC 389A: Practical Deep Learning 30 students, Student-led course initiative: classroom instruction by Sujith Vishwajith	<i>UMD, Spring 2018</i>
INST 414: Data Science Methods 50 students	<i>UMD, Spring 2018</i>
CMSC 726: Machine Learning 60 students	<i>UMD, Fall 2017</i>
CSCI 7000: Advanced Machine Learning for Natural Language Processing 24 students	<i>Colorado, Spring 2017</i>
CSCI 3022: Introduction to Data Science Algorithms 100 students	<i>Colorado, Fall 2016</i>
CSCI 5622: Machine Learning 104 students	<i>Colorado, Fall 2015</i>
CSCI 5622: Machine Learning 58 students	<i>Colorado, Spring 2015</i>
CSCI/LING 5832: Natural Language Processing 32 students	<i>Colorado, Fall 2014</i>
INST 737: Digging into Data 29 students	<i>UMD, Spring 2014</i>
CMSC/LING 723 / INST 735: Computational Linguistics I 45 students	<i>UMD, Fall 2013</i>
LING 848B / CMSC 828B: Bayesian Nonparametrics 15 students	<i>UMD, Spring 2013</i>
INST 737: Digging into Data 30 students	<i>UMD, Spring 2013</i>
LBSC 690: Introduction to Information Technology 30 students	<i>UMD, Fall 2012</i>
INST728C / CMSC 773 / LING 773: Computational Linguistics II 11 Students	<i>UMD, Spring 2012</i>
LBSC 690: Introduction to Information Technology 30 students	<i>UMD, Fall 2011</i>
INFM 718G: Web Scale Information Processing Applications 12 students	<i>UMD, Spring 2011</i>
LBSC 690: Introduction to Information Technology 30 students	<i>UMD, Fall 2010</i>
COS/LIN 280: Computational Linguistics 40 students, Taught by Christiane Fellbaum (I developed homeworks)	<i>Princeton, Fall 2008</i>

Course or Curriculum Development

- Chair of committee developing **CSAI: A New Bachelor of Science in Artificial Intelligence**. I proposed the development of a structured AI curriculum that integrates technical foundations with interdisciplinary perspectives. The proposed approach introduces students to the computational and mathematical tools necessary to build AI systems while simultaneously exposing them to the ethical, legal, linguistic, and societal contexts in which those systems operate.
- Developed new course, *CMSC 848Q: How and Why Artificial Intelligence Answers Questions* with new lectures and homeworks
- First online offering of *CMSC 723: Computational Linguistics*
- First offering of *CMSC 470: Natural Language Processing* with permanent course number
- Developed new undergraduate course, *INST 414: Data Science Methods*
- Developed new undergraduate course, *CSCI 3022: Data Science Algorithms*, a 100-person first offering without TA support
- New offering of *CSCI 5622: Machine Learning* (Spring / Fall 2015) as a flipped classroom
- Significant revisions to *LBSC 690: Information Technology* (Fall 2012)
- Chair of committee developing new undergraduate Information Science program at Universities at Shady Grove for University of Maryland (2011-2013)
- Developed new course *INST 737: Digging into Data* (Spring 2013), and recorded lectures for “flipped” classroom in 2014.
- Redesigned both elements of Computational Linguistics I-II sequence (2012 and 2013), and created a “flipped” classroom in 2013 for Computational Linguistics I.

Advising: Research Direction (Undergrad)

Current Students

1. Rohan Chintakindi (UMD CS, 2025–): Baseline Agents for Training RL Negotiation
2. Nyle Masood (UMD CS, 2025–): Distilling Human Thought Process for GeoGuesser
3. Jane Oh (UMD CS, 2025–): MCQA Scoring Schemes
4. Ahmed Haj Ahmed (Haverford CS, 2024): Manchester Paradigm Audio Questions
5. Liam Dorn (Columbia CS, 2024): Manchester Paradigm Audio Questions
6. Dmytro Kurdydyk (Davidson CS, 2024–): Manchester Paradigm Audio Questions
7. Aadi Palnitkar (UMD CS, 2024–): Item Response Theory for Incremental Questions
8. Bijan Naimi (UMD CS, 2023–2024): Interface for Audio QA
9. Matthew Shu (Yale CS, 2019–2024): Representation Learning for Spaced Repetition

Former Students

1. Runze Li (Nanjing CS, 2024): Identifying Locations from Pictures [PhD Student]
2. Zheyuan Zhang (Tsinghua CS, 2024): Identifying Locations from Pictures [PhD Student, CMU]
3. Konstantine Kahadze (UMD, 2022–2023): Creating Lies in Diplomacy
4. Sander Schulhoff (UMD, 2020–2023): Adversarial Prompt Engineering [Startup]
5. Feng Guo (UMD, 2022–2023): Playing Centaur Diplomacy [MS Student, UMD]
6. Chenglei Si (UMD, 2020–2023): Answer Equivalence [PhD Student, Stanford]
7. Andrew Mao (UMD, 2020–2022): Interactive Information Retrieval for Online Academic Competitions [Engineer, Meta]
8. Naveen Raman (UMD, 2020–2022): Active Learning for Coreference [PhD Student, CMU: Winner, Goldwater Scholarship]
9. Arjun Akkiraju (UMD, 2020–2021): Generating Adversarial QA Examples [Apple]
10. Eric Wallace (UMD, 2018–2019): Generating challenging questions for question answering [PhD Student, Berkeley]
11. Sujith Vishwajith (UMD, 2017–2018): Answer equivalence for question answering [AirBnB]
12. Davis Yoshida (CU, 2015–2017): Active feature solicitation for question answering [PhD Student, TTI]
13. Henrik Larsen (CU, 2015–2016): Verb prediction [University of Oslo]
14. Stephanie Hwa (UMD, 2013–2014): Vector word representations for named entities in question answering [LinkedIn]
15. Danny Bouwman (UMD, 2013–2014): Crowdsourced coreference annotation [Web Developer, The Canton Group]
16. Kenrick Kin, Joshua Lee, Lester Mackey (Princeton, 2005–2006): Assistive vocabulary devices for people with aphasia [Kenrick: Pixar; Joshua: Parachute Health; Lester: MSR]

Advising: Research Direction (Masters)

Chair or Co-Chair (former)

1. Feng Guo (UMD, 2023–2025)
2. Davis Yoshida (CU APMA, 2016–2017): Domain Adaptation for Question Answering [First position: PhD student, TTI-Chicago]
3. You Lu (CU CSCI, 2016–2017): Task-based Evaluation of Topic Models [First position: PhD student, Virginia Tech]

4. Alison Smith (UMD CMSC, 2011–2013): Evaluating Interfaces for Interactive Topic Modeling [First position: Dataminr]
5. Brianna Satinoff (UMD CMSC, 2009–2011): Incremental Models for Text Classification [First position: Well-point]

On Committee

1. Thomas Diggelmann (CS, 2020 ETH)
2. Varini Francesco Saverio (CS, 2020 ETH)
3. Jordan Hoskins (German, 2015 Colorado)
4. Bradley Skaggs (CMSC, 2011 UMD)

Guest lectures

- 2018 CMSC 828P: Topic Models
- 2017, PSYC 6200: Topic Models
- 2017, INFO 2301: Clustering
- 2016, HIST 6546: Topic Models
- 2016, CSCI 7000 (Data Science Team): Topic Models as Features
- 2016, CSCI 6000: Machine Learning Research
- 2015, CSCI 5832: Topic Models
- 2012, CMSC 421: Topic Models
- 2012, CMSC 726: Topic Models
- 2011, LING 773: Topic Models
- 2010, CMSC 726: Topic Models

Advising: Research Direction (Doctoral)

Chair or Co-chair (current)

1. Feng Gu (UMD CMSC, 2025–): Language Games
2. John Kanu (UMD CMSC, 2024–): IRT for Personalization
3. Nishant Balepur (UMD CMSC, 2023–): NLP for Education
4. Maharshi Gor (UMD CMSC, 2023–): Human–Computer Teaming
5. Tasnim Kabir (UMD CMSC, 2023–): Manchester Paradigm Questions
6. Ishani Mondal (UMD CMSC, 2022–): Interactive Information Extraction
7. Neha Srikanth (UMD CMSC, 2022–): Pragmatic Assumptions of Question Answering
8. Zongxia Li (UMD CMSC, 2023–): Combining LLM and Human Expertise
9. Hyojung Han (UMD CMSC, 2022–): Multilingual Question Answering
10. Wichayaporn Wongkamjan (UMD CMSC, 2022–): Reinforcement Learning for Negotiation
11. Mozhi Zhang (UMD CMSC, 2017–): Multilingual Classification for Low Resource Languages
12. Fenfei Guo (UMD CMSC, 2017–):⁶ Interactive Embedding Learning

Chair or Co-chair (former)

1. Kyle Seelman (UMD CMSC, 2022–2025): Interactive Neural Topic Models [Roblox]
2. YooYeon Sung (iSchool, 2021–2025): Adversarial Examples for Question Answering and Disinformation [Hippocratic AI]
3. Michelle Yuan (UMD CMSC, 2018–2022): Multilingual Interactive Topic Modeling [Amazon]
4. Chen Zhao (UMD CMSC, 2018–2022): Algorithms with Structure for Answering Complex Questions [New York University Shanghai, Assistant Professor]
5. Yoshinari Fujinuma (Colorado CMSC, 2016–2021): Multilingual Dense Representations [Amazon]
6. Denis Peskoff⁷ (UMD CMSC, 2017–2021): Human-in-the-loop Data Curation [Princeton University, Postdoc → Berkeley University, Postdoc]
7. Pedro Rodriguez (UMD CMSC, 2017–2021): Evaluating Machine Intelligence with Question Answering [Meta FAIR]
8. Shi Feng (UMD CMSC, 2017–2021): User Modeling of Knowledge and Ability for Question Answering [George Washington University, Assistant Professor⁸]
9. Ahmed Elgohary (UMD CMSC, 2018–2021): Human Feedback for Question Answering [Microsoft Research]
10. Shudong Hao⁹ (CU CSCI, 2015–2017): Interactive Multilingual Topic Modeling [Stevens Institute of Technology, Assistant Professor]
11. Alison Renner-Smith¹⁰ (UMD CMSC, 2017–2019): Interactive Topic Modeling [Dataminr]
12. Weiwei Yang (UMD CSCI, 2014–2019; co-advised by Philip Resnik): Linked Topic Models [Meta]
13. Forough Poursabzi-Sangdeh (CU CSCI, 2014–2018): Active Labeling with Topic Models [Microsoft]
14. Kim Glasgow¹¹ (UMD iSchool 2011–2014): Social Action in Social Media [JHU APL]

⁶Was formerly a PhD student at Colorado, enrolled 2015

⁷Formerly Denis Peskov, but changed his name to distinguish himself from the son of Kremlin spokesperson

⁸Starting Fall 2024

⁹Later advised by Michael Paul

¹⁰Originally advised by Leah Findlater

¹¹Later advised by Yla Tausczik

15. Thang Nguyen (UMD CMSC, 2012–2018): Anchor-Based Topic Inference [CGI Federal]
16. Alvin Grissom II (CU CSCI, 2012–2017): Reinforcement Learning for Feature-wise Language Tasks [Haverford College, Assistant Professor]
17. Mohit Iyyer (UMD CMSC, 2012–2017; co-advised by Hal Daumé): Deep Learning for Question Answering [University of Massachusetts Amherst, Associate Professor]
18. He He (UMD CMSC: 2011–2016; co-advised by Hal Daumé): Algorithms that Trade-Off Speed and Accuracy [New York University, Assistant Professor]
19. Viet-An Nguyen (UMD CMSC, 2010–2015; co-advised by Philip Resnik): Detecting Influence in Text [Meta]
20. Ke Zhai (UMD CMSC, 2009–2014): Large Scale Bayesian Inference [Microsoft]
21. Yuening Hu (UMD CMSC: 2009–2014): Interactive Topic Modeling [Google]

On committee

1. Lingjun Hao (CMSC)
2. Paiheng Xu (CMSC)
3. Rachel Sexton (MENG, 2025)
4. Angelos Mavrogiannis (CMSC, 2025)
5. Md Main Uddin Rony (INFO, 2025)
6. Yuhang Zhou (INFO, 2025)
7. Md Naimul Hoque (INFO, 2024)
8. Chujun Song (CMSC, 2024)
9. Alexander Hoyle (CMSC, 2024)
10. Pranav Goel (CMSC, 2024)
11. Xin Qian (INST)
12. Shraey Bhatia (Melbourne, Computer Science 2024)
13. Martin Fajčík (Brno, Computer Science 2024)
14. Suraj Nair (CMSC, 2023)
15. Noel Warford (CMSC, 2023)
16. Sweta Argawal (CMSC, 2023)
17. Jason Fan (CMSC, 2023)
18. Joeseeph Barrow (CMSC, 2022)
19. Varun Suryan (CMSC, 2022)
20. Sunandita Patra (CMSC, 2020)
21. Xing Niu (CMSC, 2019)
22. Yogarshi Vyas (CMSC, 2019)
23. Varun Manjunatha (CMSC, 2018)
24. Sudha Rao (CMSC, 2018)
25. Jinfeng Rao (CMSC, 2018)
26. Anupam Guha (CMSC, 2017)
27. Nicholas Dronen (CU CSCI 2017)
28. Bill Foland (CU CSCI 2017)
29. Sam Way (CU CSCI 2017)
30. Shweta Bhandare (CU CSCI 2017)
31. Nicole Beckage (CU CSCI 2017)
32. Karl Ridgeway (CU CSCI)
33. Amir Ghasemianlangroodi (CU CSCI 2017)
34. Brett Roads (CU CSCI 2017)
35. Abbie Jacobs (CU CSCI 2017)
36. Md Arafat Sultan (CU CSCI, 2016)
37. Ben London (UMD CMSC, 2014)
38. Irene Eleta (UMD INFO, 2014)
39. Kevin Dayaratna (UMD STAT, 2014)
40. Jiarong Jiang (UMD CMSC, 2014)
41. Jagadeesh Jagarlamudi (UMD CMSC, 2013)
42. Amit Goyal (UMD CMSC, 2013)
43. Piyush Rai (Utah CMSC, 2012)
44. Arvind Agarwal (UMD CMSC, 2012)
45. Elena Zheleva (UMD CMSC, 2011)
46. Asad Sayeed (UMD CMSC, 2011)

Service and Outreach

Reviewing and Editing for Journals and Presses: Articles

- Reviewer for *Science*: 2022
- Reviewer for *Journal of Machine Learning Research*: 2011, 2012, 2019, 2021, 2022

- Reviewer for *Computational & Mathematical Organization Theory*: 2021
- Reviewer for *Machine Learning*: 2019
- Action Editor for *Transactions of the Association for Computational Linguistics*: 2017–2019
- Reviewer for *Journal of AI Research*: 2016
- Reviewer for *International Journal on Digital Libraries*: 2016
- Reviewer for *Machine Learning Journal*: 2014
- Reviewer for *Transactions of the Association of Computational Linguistics*: 2011,2012,2012,2013,2013,2013,2014
- Reviewer for *IEEE Transactions on Pattern Analysis and Machine Intelligence*: 2013
- Reviewer for *Computational Linguistics*: 2013
- Reviewer for *Scientometrics*: 2012
- Reviewer for *Information Visualization*: 2012
- Reviewer for *Transactions on Knowledge Discovery from Data*: 2011
- Reviewer for *Annals of Applied Statistics*: 2011
- Reviewer for *Elsivier Computer Speech and Language*: 2007

Reviewing for Journals and Presses: Books

- Reviewer for *R Programming and Data Science*, Chapman and Hall

Reviewing activities for Agencies and Foundations

- NSF IIS Review Panel (2024)
- NSF IIS Review Panel (2023)
- NSF IIS Review Panel (2022)
- Hasler Stiftung: Expertengutachten (2021)
- NSF External Review (2021)
- NSF III Review Panel (2021)
- NSF IIS Review Panel (2021)
- NSF External Review (2020)
- NSF III Review Panel (2018)
- NSF IIS Review Panel (2017)
- NSERC External Review (2017)
- NSF IIS External Review (2015)
- NSF IIS Review Panel (2015)
- NSF External Review (2014)
- NSF IIS Review Panel (2012)
- NSF BIGDATA Review Panel (2012)

Reviewing Activities for Conferences

- Program Committee *LREC*: 2023
- Program Committee *International Conference on Learning Representations*: 2021
- Program Committee *Association for Computational Linguistics*: 2020, 2019, 2014, 2012, 2011, 2010
- Program Committee *Empirical Methods in Natural Language Processing*: 2019, 2014, 2013, 2012, 2011, 2008
- Best Paper Committee, *ACL* 2018
- Program Committee *2014 Workshop on Language Technologies and Computational Social Science*
- Program Committee *Neural Information Processing Systems*: 2014, 2013, 2012, 2011, 2010, 2009
- Program Committee *International Conference of Machine Learning*: 2014, 2013, 2012, 2011, 2010, 2009
- Program Committee *International Conference on the Web and Social Media*: 2014
- Program Committee *World Wide Web Conference*: 2014
- Program Committee *AISTATS*: 2012, 2011
- Program Committee *North American Association for Computational Linguistics*: 2012
- Program Committee *NIPS 2010 Workshop on Computational Social Science and the Wisdom of Crowds*
- Program Committee *NAACL 2010 Workshop on Creating Speech and Text Language Data With Amazon's Mechanical Turk*
- Reviewer for *COLING* 2010
- Program Committee *Global WordNet Association Conference*: 2010, 2008, 2006
- Assistant Reviewer for *UAI* 2007
- Reviewer, Works in Progress 2006 *SIGCHI*

Professional and Campus Service

Departmental Service

Computer Science, Maryland

- Chair, Instructional Faculty Search Committee (2024)
- Faculty Teaching Review Committee (2020-2023)
- Faculty Mentor, Technica Research Bitcamp (2021)
- Faculty Mentor, Exploration Trail Research Bitcamp for Underrepresented Minorities (2021)

- Teaching Review Committee (2020–2021)
- Member: Search Committee (2018–2019)

Departmental Service

Institute for Advanced Computer Studies, UMD

- Merit Review Committee (2022)
- Computer Science Liaison (Graduate Admissions, 2020–2022)
- APT Member (2017–2018, 2020–2021)
- Director: Computational Linguistics and Information Processing Lab (2017–2018, 2021–2022)
- Hardware Czar: Computational Linguistics and Information Processing Lab (2017–2018, 2021)
- Member: Appointments and Promotion (2012–2013, 2017–2018)
- Data Czar: Computational Linguistics and Information Processing Lab (2011–2014)
- Coordinator: Computational Linguistics and Information Processing Lab Colloquium (2010–2012)

College Service

College of Computer, Mathematical, and Natural Sciences UMD

- Member: Rankings and Reputation Committee (2022–2024)

College Service

College of Information Studies, UMD

- Member: Search Committee (joint with Journalism) for Computational Journalism (2017–2018, 2018–2019)
- Chair: College of Information Studies Undergraduate Education Committee (2011–2013)
- Member: College of Information Studies Undergraduate Education Committee (2011–2014)
- Secretary: College of Information Studies Assembly (2011–2012)
- Member: College of Information Studies Programs, Courses, Curriculum Committee (2011–2013)
- Member: College of Information Studies Research Committee (2010–2013)

University Service

UMD

- Member: Language Science Undergraduate Committee (2021–2023)
- Member: Language Science Executive Council (2017–2023)
- Faculty Advisor: Maryland Academic Quiz Team (2010–2014, 2017–)
 - National Champion, 2019 Division II Intercollegiate Championship Tournament
 - 5th place, 2019 Division I Intercollegiate Championship Tournament
 - 8th place, 2018 Division I Intercollegiate Championship Tournament
 - Champion, 2017 Academic Competition Federation
 - 4th place, 2017 Division I Intercollegiate Championship Tournament
 - 4th place, 2014 Division I Intercollegiate Championship Tournament
 - 4th place, 2012 Division I Intercollegiate Championship Tournament

Departmental Service

Computer Science, Colorado

- Member: Graduate Committee (2014–2017)
- Member: Search Committee (Machine Learning, 2014–2015)

College Service

Engineering and Applied Science, Colorado

- Yellowshirt Interviewer, 2015

University Service

Colorado

Faculty Advisor: cu Trivia Buffs (20th place, ICT D2 2016)

- Tournament Director: Rocky Mountain Region Academic Quiz Tournament 2017 (College)
- Tournament Director: Rocky Mountain Region Academic Quiz Tournament 2016 (College)
- Tournament Director: Colorado State Academic Quiz Tournament 2016 (High School)
- Tournament Director: Colorado State Academic Quiz Tournament 2015 (High School)

Offices and Committee Memberships—Professional

- Board of the Association for Computational Linguistics (At Large Elected Member), 2025–2028

Leadership Roles in Meetings and Conferences

- **Member** Best Paper Committee, *ICML 2026*
- **Area Chair**, *ARR/EMNLP 2026*
- **Area Chair**, *ARR/EMNLP 2025*
- **Area Chair** for Position Papers, *ICML 2025*
- **Senior Area Chair**, *ACL 2025*
- **Area Chair**, *NeurIPS 2024*
- **Area Chair** for Question Answering *EMNLP 2024*
- **Area Chair** for Large Language Models and the Future of NLP *EMNLP 2023*
- **Program Chair** *ACL 2023*

- **Poster Chair** EMNLP 2022
- **Area Chair** for Question Answering EMNLP 2022
- **Senior Area Chair** for Question Answering NAACL 2022
- **Area Chair** for Question Answering EMNLP 2021
- **Area Chair** for Question Answering ACL 2021
- **Area Chair** for Question Answering NAACL 2021
- **Organizer** Efficient QA Competition, *NeurIPS* 2020
- **Senior Area Chair** for Machine Learning, EMNLP 2018
- **Area Chair**, ICML 2018
- **Area Chair**, *NeurIPS* 2018
- **Oranizer**, *NIPS 2017 Human-Computer Question Answering Competition*
- **Area Chair**, ICML 2017
- **Area Chair** for Machine Learning, EMNLP 2017
- **Tutorial Co-Chair**, ACL 2017
- **Area Chair** for Machine Learning, EMNLP 2015
- **Area Chair**, ICML 2015
- **Area Chair** for Document Classification and Topic Clustering, NAACL 2015
- **Co-organizer** for ACL 2014 *Student Research Workshop*
- **Co-organizer** for NIPS 2013 *Workshop on Topic Models*
- **Computational Committee** *North American Computational Linguistics Olympiad* 2012-2014
- **Area Chair** for Document Classification and Topic Clustering, NAACL 2012
- **Co-organizer** for NIPS 2009 *Workshop on Topic Model Applications: Text and Beyond*

Unpaid services to local, state, and federal agencies

- Panelist for *Under the Hood: A Deeper Dive Into How AI Technology Works*, organized by Computer and Communications Industry Association for Congressional Staffers (2023)
- Consultant for *Interactive Topic Modeling*, National Institutes for Health (2010)
- Collaboration on *Dirichlet Process Protein Clustering*, National Institutes for Health (2012)
- Collaboration on *Nonparametric Beta Coalescent Clustering*, US Department of Agriculture (2013)
- Consultant for *Interactive Topic Modeling*, National Institute of Food and Agriculture (2013)

Paid Consulting

- Judge (with honararium): Rich Context Competition, 2019
- Consultant: Norwegian Research Council, 2014
- Consultant: Barquin International, 2013-2014
- Consultant: New Brand Analytics, 2012-2014

Non-Research Presentations

Outreach Presentations

1. QANTA Trivia All-Star Competition, 2018 (College Park, MD)
2. OUSIA vs. Trivia All-Stars, 2018 (Reston, VA)
3. OUSIA vs. Trivia All-Stars, 2017 (Long Beach, CA)
4. QANTA vs. All-Star Quiz Bowl Team, 2017 (Atlanta, GA)
5. OUSIA vs. CA NASAT Team, 2016 (San Diego, CA)
6. QANTA vs. All-Star Quiz Bowl Team, 2016 (Dallas, TX)
7. QANTA vs. Ken Jennings, 2015 (Seattle, WA)
8. QANTA vs. Jeopardy Champions, 2014 (Chicago, IL)

External Visibility

Keynotes

1. AI and Tech Workshop 2025: From Muppet Models to an an Interdisciplinary AI-first Education: Sustaining and Celebrating Maryland's AI Contributions
2. Open Data Science Conference East 2023: If We Want AI to be Interpretable, We Need to Measure Interpretability
3. KnowledgeNLP workshop at AAAI 2023: Raw Knowledge vs. Understanding: What Adversarial QA Reveals about the Limits of AI
4. School Science and Mathematics Association 2020: Question Answering isn't a Game Show (but maybe it should be)
5. Frontiers of Engineering (National Academies of Engineering) 2017: Humans and Computers Working Together to Measure Machine Learning Interpretability - Jordan Boyd-Graber
6. European Conference on Information Retrieval 2016: Machine Learning Shouldn't be a Black Box

7. **Yandex Data Science 2013: Big Data Analysis with Topic Models: Human Interaction, Streaming Computation, and Social Science Applications**

Invited Talks

1. **The Questions Computers (still) can't Answer:** Planet Word, 2025
2. **Helpful AI Models: You can't always get what you want, but you might get what you need:** Universität des Saarlandes, 2025; Universität Zürich; Aleph Alpha; Nanyang Technical University, 2025
3. **Language Games:** Gothenburg University 2025; ITU Copenhagen, 2025
4. **AI's Certain Uncertainty Problem:** Tsinghua University, 2024; Getting Aligned On AI Alignment Workshop at UMD, 2024
5. **The questions that computers still cannot answer: how to find them, why they're fun, and what they mean for AI:** Max Planck-Cornell-Maryland Summer School, 2023
6. **The State of Generative AI in 2023:** Computer & Communications Industry Association AI Informational Event for Congressional Staffers
7. **The Benefits of Unnatural Data: Adversarial Examples for Question Answers, Fact Checking, and Detecting Deception:** Northwestern University, 2023
8. **If We Want AI to be Interpretable, We Need to Measure Interpretability:** University of Toronto, 2022; Ruhr Universität Bochum, 2022; University of Chicago, 2022; University of Illinois at Urbana-Champaign, 2022; Northwestern University, 2023
9. **Manchester vs. Cranfield: Why do we do question answering and how can we do it better?:** ByteDance Research, 2021; University of Pennsylvania, 2022; UMD Search Mastery Series, 2022; NAACL Dynamic Adversarial Dataset Collection Workshop (invited talk and panel)
10. **Active Learning and Refining Representations for NLP:** HLT CoE Tech Exchange, 2021
11. **Cooperative and Competitive Machine Learning through Question Answering:** MBZUAI Distinguished Colloquium, 2021; University of Würzburg, 2021; NLP Zurich, 2020; TTI Chicago, 2018; DC Data Science (Washington, DC), 2018; Georgetown Computer Science (Washington, DC), 2017; New York Text as Data (New York City, NY), 2017
12. **Question Answering isn't a Game Show (but maybe it should be):** UMass CS Colloquium, 2020
13. **What Question Answering can Learn from Trivia Nerds:** Machine Reading Question Answering Workshop (EMNLP Hong Kong), 2019
14. **Engaging Hobbyist Communities to Deceive Machines and Each Other:** Text as Data (Zürich), 2019; Aggregating and analysing crowdsourced annotations for NLP (EMNLP Hong Kong), 2019
15. **Opening the black box of machine learning: Interactive, interpretable interfaces for exploring linguistic tasks:** Laboratory of Telecommunication Science (College Park, MD), 2018; University of Maryland Language Science Center Winter Storm (College Park, MD), 2018
16. **Defining Artificial Intelligence, Data Science, and Machine Learning:** Office of the Inspector General of the US Postal Service (Arlington, VA), 2018
17. **Computational Modeling of Relationships, Spin, and Betrayal:** Center for Bioinformatics and Computational Biology Seminar, 2019 (College Park, MD); Stanford University, 2017 (Palo Alto, CA)
18. **Machine Learning Shouldn't be a Black Box:** Ludwig Maximilian Universität, 2016 (München, Germany); ECIR Keynote, 2016 (Padua, Italy); Cambridge Statistics Colloquium, 2016 (Cambridge, England); Royal Institute of Technology (Stockholm, Sweden), 2016; Carnegie Mellon University Language Technology Institute Colloquium (Pittsburgh, PA), 2016; University of Maryland, 2016 (College Park, MD)
19. **Interactive Topic Modeling and The US Tea Party:** New Frontiers of Automated Content Analysis in the Social Sciences, 2015 (Zürich, CH)
20. **Thinking on your Feet: Reinforcement Learning for Incremental Language Tasks:** Colorado School of Mines, 2014 (Golden, CO); Harvey Mudd College, 2014 (Claremont, CA); California Institute of Technology, 2014 (Pasadena, CA); Front Range NLP (Boulder, CO); EECS Colloquium, Colorado School of Mines, 2014 (Golden, CO); Brigham Young University, 2014 (Provo, UT); Peking University, 2014 (Beijing, PRC); Darmstadt University, 2014 (Darmstadt, Germany); Hong Kong University of Science and Technology, 2014 (Hong Kong); Cornell University, 2015 (Ithaca, NY); 2015 Yandex School of Data Science (Berlin, Germany); Facebook 2017 (Menlo Park, CA)
21. **Big Data Analysis with Topic Models: Human Interaction, Streaming Computation, and Social Science Applications:** University of Colorado Boulder Computer Science Colloquium, 2013 (Boulder, CO); Yandex Machine Learning Conference, 2013 (**Invited Keynote**, Moscow, Russia); DC NLP Meetup, 2014 (Washington, DC); Yahoo! Labs, 2014 (New York, NY); Northwestern University (Evanston, IL); Renmin University, 2014 (Beijing, PRC); Tsinghua University, 2014 (Beijing, PRC); University of Maryland, 2017 (College Park, MD)
22. **Incorporating Human Knowledge and Insights into Probabilistic Models of Text:** Brigham Young University Department of Computer Science Colloquium, 2012 (Provo, UT)
23. **Besting the Quiz Master: Crowdsourcing Incremental Classification Games:** Rutgers University, 2012 (New Brunswick, NJ); Brigham Young University, 2012 (Provo, UT)
24. **Making Topic Models More Human(e):** Colorado University, 2012 (Boulder, CO); University of Maryland Institute for Technology and Humanities, 2012 (College Park, MD)
25. **When Topic Models Go Bad: Diagnosing and Improving Models for Exploring Large Corpora:** Johns Hopkins

- University, 2011 (Baltimore, MD); Rutgers University, 2011 (New Brunswick, NJ)
26. **Inference and Validation of Probabilistic Models of Language in the Cloud:** UMD Winter Storm, 2011 (College Park, MD)
 27. **Interactive Topic Models:** Harvard University's New Directions in Text Analysis Symposium, 2011 (Cambridge, MA); Princeton University, 2011 (Princeton, NJ); Maryland Institute for Technology and the Humanities: Topic Modeling and the Humanities Workshop, 2012 (College Park, MD)
 28. **Putting Words Together: Crowdsourcing Data Collection for Lexical Similarity and Topical Coherence:** University of Massachusetts, 2010 (Amherst, Massachusetts)
 29. **Topic Models, Mechanical Turk, and WordNet:** Harvard University, 2010 (Cambridge, MA)
 30. **Topic Models and Hierarchical Models:** Johns Hopkins Summer Workshop for SMT, 2010 (Baltimore, MD)
 31. **Linguistic Extensions to Topic Models:** University of Massachusetts, 2009 (Amherst, Massachusetts); Center for Communications Research, 2009 (Princeton, NJ); John Hopkins Human Language Technologies Center of Excellence, 2009 (Baltimore, MD); Columbia University, 2009 (New York, NY)

Press Coverage

1. Chris Carroll. **Can AI and People Play Nice?: A Trivia Nerd's Quizbowl Quest Sheds Light on When to Trust Computers, or Ourselves.** *EnTERPrise*, 2026.
2. Matteo Wong. **Could ChatGPT Secretly Tell You How to Vote?.** *Atlantic*, 2025.
3. Jean Marbella. **Can AI Surpass Human Intelligence?.** *Baltimore Sun*, 2025.
4. Chris Carroll. **Can AI and People Play Nice?.** *Baltimore Sun*, 2025.
5. Maryland Today Staff. **At New AI Institute's Celebration, a Question of 'Who's at the Table'.** *Maryland Today*, 2024.
6. Maryland Today Staff. **Pocket-Sized Support for New Parents.** *Maryland Today*, 2024.
7. Courtney Ryan. **NSF grant will help address AI's diversity problem.** *Mendoza College of Business*, 2024.
8. Stephanie Lee. **Why does AI beat humans at the strategy game Diplomacy?.** *Tech Xplore*, 2024.
9. Maria Herd. **Can a Strategy Game Help AI Learn to Spot Scammers?.** *Maryland Today*, 2021.
10. Tom Ventsias. **Humans vs. Machines: Examining the Effectiveness of Automated Topic Modeling Evaluations.** *UMIACS*, 2021.
11. Maria Herd. **Two CLIP Graduate Students Named 2021 Computing Innovation Fellows.** *UMIACS*, 2021.
12. Kaner and Amby. **The Lying Interview.** *DiplomacyGames*, 2020.
13. Matt Gardner and Waleed Ammar. **Pathologies of Neural Models Make Interpretation Difficult.** *AI2 NLP Highlights*, 2019.
14. Kimbra Cutlip. **Seeing How Computers "Think" Helps Humans Stump Machines and Reveals Artificial Intelligence Weaknesses.** *Science Daily*, 2019.
15. Gino Dino. **Terminology Assistance Coming to a Simultaneous Interpreter Near You.** *Slator*, 2019.
16. Sam Charrington. **Pathologies of Neural Models and Interpretability with Alvin Grissom II.** *TWiML Talk*, 2019.
17. Graham Neubig. **Does Not Compute?.** *The Linguist*, 2019.
18. Melissa Brachfeld. **Boyd-Graber, Feng Present Paper on Machine Learning Interpretability at IUI 2019.** *UMIACS*, 2019.
19. Matt Gardner and Waleed Ammar. **A Discussion of Question Answering.** *AI2 NLP Highlights*, 2018.
20. Sydney Worth. **CU Boulder team to compete at national quiz tourney for second year in a row.** *Daily Camera*, 2018.
21. Yomihiro Katabuti. **The Impact of Fast-Moving AI: Dominating Trivia Champs.** *IT Media*, 2018.
22. Matt Early Wright. **Inside AI's "Black Box".** *Maryland Today*, 2018.
23. Sala Levin. **What Is ... a Research-Inspired Route to "Jeopardy!"?.** *Maryland Today*, 2018.
24. Esther Bond. **Simultaneous Interpreters May Soon Get Real-Time Help Just When They Need It.** *Slator*, 2018.
25. Samir Singh. **Summary of Pathologies of Neural Models Make Interpretations Difficult.** *UCI NLP*, 2018.
26. Brandi Adams. **Associate Professor Jordan Boyd-Graber to appear on Jeopardy on September 26th 2018.** *UMD Computer Science*, 2018.
27. Melissa Brachfeld. **Boyd-Graber Publishes Paper in PNAS that Assesses Scholarly Influence.** *UMIACS*, 2018.
28. Rob Mitchum. **New model reveals forgotten influencers and 'sleeping beauties' of science.** *University of Chicago*, 2018.
29. Melissa Brachfeld. **UMD Computerized System Beats Human Quiz Bowl Team at Atlanta Exhibition.** *UMIACS*, 2017.
30. Inderjeet Mani. **When robots read books.** *Aeon*, 2016.

31. Luke Dormehl. **Like parents from the 1950s, AI still can't understand comics. Here's why.** *Digital Trends*, 2016.
32. Katherine Gorman. **Automatic Translation and t-SNE.** *Talking Machines*, 2016.
33. Christopher Chabris. **Game On: The Psychology of Betrayal.** *Wall Street Journal*, 2016.
34. Melissa Dahl. **Should you worry about people who are too polite?.** *CNN*, 2015.
35. Emily Adams. **Professor's Quiz Bowl Robot Goes Head-To-Head With Humans.** *Colorado Computer Science Newsletter*, 2015.
36. Morgan Eichensehr. **STUDY: Changes in language, tone could reveal impending betrayal.** *Diamondback*, 2015.
37. Lexie Schapitl. **UMD researchers' computer beats 'Jeopardy!' star Ken Jennings at trivia.** *Diamondback*, 2015.
38. Leigh Weingus. **What You Should Know About Really Polite People.** *Huffington Post*, 2015.
39. Melissa Dahl. **Here's a Good Reason to Be Wary of Overly Polite People.** *New York Magazine*, 2015.
40. Nathan Collins. **Detecting a Coming Betrayal.** *Pacific Standard*, 2015.
41. Bill Steele. **Language analysis predicts a coming betrayal.** *Phys.org*, 2015.
42. Rachel Ehrenberg. **A few key signs betray betrayal.** *Science News*, 2015.
43. Liam Farrell. **Question-Answering System Built by UMD, UC Boulder Bests Ken Jennings.** *Terp*, 2015.
44. **Question: Which language scientists designed first competitive quiz bowl playing NLP system?.** *UMD Language Science Newsletter*, 2015.
45. Melissa Brachfeld. **Computerized Question-Answering System Built by UMD, UC Boulder Bests "Jeopardy!" Champion.** *UMIACS*, 2015.
46. Marcus Smith. **Prof. Ringger and Natural Language Processing.** *Thinking Aloud*, 2014.