

Chenhao Tan

Assistant Professor
University of Chicago
Department of Computer Science
Harris School of Public Policy
5730 S. Ellis Avenue
Chicago, IL 60637

Email: chenhao@chenhaot.com
Homepage: <https://chenhaot.com>

Research Interests

Human-centered AI, natural language processing, computational social science.

Professional experience

- Assistant Professor**, 2021-,
University of Chicago, Chicago, IL.
- Assistant Professor**, 2017-2020,
University of Colorado, Boulder, CO.
- Research Associate**, 2016-2017,
University of Washington, Seattle, WA
Supervisor: Noah Smith.

Education

- Ph.D. in Computer Science**, 2010-2016,
Cornell University, Ithaca, NY,
Advisor: Lillian Lee.
- B.E. in Computer Science**, 2006-2010,
Tsinghua University, Beijing, China,
Advisor: Jie Tang.
- B.E. in Economics (second major)**, 2007-2010,
Tsinghua University, Beijing, China.

Selected Publications

- Chacha Chen, Shi Feng, Amit Sharma, and *Chenhao Tan*. *Machine Explanations and Human Understanding*. Transactions on Machine Learning Research (TMLR); FAccT 2023; **Best paper award** at the ICML 2022 Workshop on Human-Machine Collaboration and Teaming.
- Han Liu, Yizhou Tian, Chacha Chen, Shi Feng, Yuxin Chen, and *Chenhao Tan*. *Learning Human-Compatible Representations for Case-Based Decision Support*. ICLR 2023.
- Mourad Heddaya, Solomon Dworkin, *Chenhao Tan*, Rob Voigt, and Alexander Zentefis. *Language of Bargaining*. ACL 2023.

Yangqiaoyu Zhou, Yiming Zhang, and Chenhao Tan. *FLamE: Few-shot Learning from Natural Language Explanations*. ACL 2023.

Vivian Lai, Chacha Chen, Alison Smith-Renner, Q. Vera Liao, and Chenhao Tan. *Towards a Science of Human-AI Decision Making: A Survey of Empirical Studies*. FAccT 2023.

Yiming Zhang, Shi Feng, and Chenhao Tan. *Active Example Selection for In-Context Learning*. EMNLP 2022.

Chenhao Tan. *On the Diversity and Limits of Human Explanations*. NAACL 2022.

Abhinav Kumar, Chenhao Tan, and Amit Sharma. *Probing Classifiers are Unreliable for Concept Removal and Detection*. NeurIPS 2022.

Vivian Lai, Samuel Carton, Rajat Bhatnagar, Q. Vera Liao, Yunfeng Zhang, and Chenhao Tan. *Human-AI Collaboration via Conditional Delegation: A Case Study of Content Moderation*. CHI 2022.

Chao-Chun Hsu and Chenhao Tan. *Decision-focused Summarization*. EMNLP 2021.

Han Liu, Vivian Lai, and Chenhao Tan. *Understanding the Effect of Out-of-distribution Examples and Interactive Explanations on Human-AI Decision Making*. CSCW 2021.

Samuel Carton, Anirudh Rathore, and Chenhao Tan. *Evaluating and Characterizing Human Rationales*. EMNLP 2020.

Ramaravind Kommiya Mothilal, Amit Sharma, and Chenhao Tan. *Explaining Machine Learning Classifiers through Diverse Counterfactual Explanations*. FAccT 2020.

Vivian Lai, Han Liu, and Chenhao Tan. *“Why is ‘Chicago’ deceptive?”: Towards Building Model-Driven Tutorials for Humans*. CHI 2020.

Brian Lubars and Chenhao Tan. *Ask not what AI can do for you, but what AI should do: Towards a framework of task delegability*. NeurIPS 2019 (spotlight presentation).

Vivian Lai and Chenhao Tan. *On Human Predictions with Explanations and Predictions of Machine Learning Models: A Case Study on Deception Detection*. FAccT 2019.

Jason Shuo Zhang, Chenhao Tan, and Qin Lv *“This is why we play”: Characterizing Online Fan Communities of the NBA Teams*. CSCW 2018.

Dallas Card, Chenhao Tan, and Noah A. Smith. *Neural Models for Documents with Metadata*. ACL 2018.

Chenhao Tan. *Tracing Community Genealogy: How New Communities Emerge from the Old*. ICWSM 2018.

Elizabeth Clark, Anne Ross, Chenhao Tan, Yangfeng Ji, Noah A. Smith. *Creative Writing with a Machine in the Loop: Case Studies on Slogans and Stories*. IUI 2018.

Yang Yang, Chenhao Tan, Zongtao Liu, Fei Wu, Yueting Zhuang. *Urban Dreams of Migrants: A Case Study of Migrant Integration in Shanghai..* AAAI 2018.

Chenhao Tan, Vlad Niculae, Cristian Danescu-Niculescu-Mizil, Lillian Lee. *Winning arguments: Interaction dynamics and persuasion strategies in good-faith online discussions*. WWW 2016.

Chenhao Tan, Lillian Lee. *All Who Wander: On the Prevalence and Characteristics of Multi-community Engagement..* WWW 2015.

Chenhao Tan, Lillian Lee, Bo Pang. *The effect of wording on message propagation: Topic- and author-controlled natural experiments on Twitter*. ACL 2014.

Selected Awards and Grants

Outstanding Senior Area Chair at ACL (2023)

Sloan research fellowship (2023)

Best Paper Award at the ICML Workshop on Human-Machine Collaboration and Teaming (2022)

Google, JP Morgan research award (2022)

NSF CAREER award (2020)

IBM, Amazon, Salesforce research award (2020)

Best Paper Award at the ACL RepL4NLP workshop (2018)

Facebook Fellowship (2015)

Selected Talks

Towards Human-Centered AI: Understanding and Improving Human Decision Making through Explanations. 2023.

Language, Information, and Decision Making. Machine Learning in Economics Summer Institute. 2022.

Tutorial on Human-Centered Evaluation of Explanations. NAACL 2022.

Towards Human-Centered AI: Understanding Language in Social Contexts. 2020.

Online social interactions: a lens on humans and a world for humans. 2016.

Selected Service

Steering Committee at FAccT (2023-2026)

General Chair at ACM Conference on Fairness, Accountability, and Transparency (FAccT 2023)

Program chair at the International Conference on Computational Social Science (IC2S2 2022)

Senior area chair at ACL 2023 and NAACL 2021

Co-editor of the Human-Centered AI Medium Publication

Full List of Awards and Grants

Sloan research fellowship (2023)

Outstanding Senior Area Chair at ACL (2023)

NSF-CSIRO: HCC: Small: From Legislations to Action: Responsible AI for Climate Change (2023)

Best Paper Award at the ICML Workshop on Human-Machine Collaboration and Teaming (2022)

Google research scholar award (2022)

JP Morgan faculty research award (2022)

Data and Democracy initial grant: Partisan Polarization and Political Dysfunction (co-PI, 2022)

Data and Democracy initial grant: Understanding and Auditing Content Moderation Policies (co-PI, 2022)

CDAC Discovery Challenge: AI-Driven Tutorials for Radiology Student Training (lead PI, 2021)

FAI: Towards Adaptive and Interactive Post hoc Explanations, National Science Foundation (lead PI, 2021)

NSF CAREER award (2020)

IBM faculty award (2020)

Amazon research award (2020)

Salesforce research award (2020)

AI-Society EAGER: Explanations through Diverse, Feasible, and Interactive Counterfactuals (sole PI), National Science Foundation (2019)

CRII: Harnessing Machine Learning to Improve Human Decision Making: A Case Study on Deceptive Detection (sole PI), National Science Foundation (2019)

CHS Small: A Genealogical Framework to Understand the Emergence of Online Groups (co-PI), National Science Foundation (2019)

Best Paper Award at the ACL RepL4NLP workshop (2018)

BIGDATA Large: Automating Analysis and Feedback to Improve Mathematics Teachers' Classroom Discourse (co-PI), National Science Foundation (2018)

Facebook Global Literacy Challenges, final list (2018)

Facebook Research gift (2018)

NVIDIA GPU grant (2017)

Postdoc research awards, University of Washington (2017)

Selected to participate in Heidelberg Laureate Forum (2016)

OARU travel award for Heidelberg Laureate Forum (2016)

Facebook Fellowship (2015)

Telluride Residential Scholarship (2012-2016)

Yahoo! Key Scientific Challenges (2012)

KDD Student Travel Award (2011, 2010)

Best Undergraduate Thesis of Tsinghua University (Top 1 among ~160) (2010)

Excellent Undergraduate of Tsinghua University (Top 58 among 3144) (2010)

Excellent Undergraduate of Beijing (2010)

National Scholarship (Top 3 among ~160) (2009, 2008)

Full List of Publications

[Google scholar profile]

Gale H. Prinster, C. Estelle Smith, *Chenhao Tan*, and Brian C. Keegan. *Community Archetypes: An Empirical Framework for Guiding Research Methodologies to Reflect User Experiences of Sense of Virtual Community*. CSCW 2024.

Dang Nguyen, Chacha Chen, He He, and *Chenhao Tan*. *Pragmatic Radiology Report Generation*. ML4H 2023.

Xiaoming Shi, Siqiao Xue, Kangrui Wang, Fan Zhou, James Y. Zhang, Jun Zhou, *Chenhao Tan*, and Hongyuan Mei. *Language Models Can Improve Event Prediction by Few-Shot Abductive Reasoning*. NeurIPS 2023.

Karen Zhou and *Chenhao Tan*. Entity-Based Evaluation of Political Bias in Automatic Summarization. Findings of EMNLP 2023 (short paper). Presented at TrustNLP at ACL 2023.

Nan-Jiang Jiang, *Chenhao Tan*, and Marie-Catherine de Marneffe. *Ecologically Valid Explanations for Label Variation in NLI*. Findings of EMNLP 2023 (short paper).

Vivian Lai, Yiming Zhang, Chacha Chen, Q. Vera Liao, and *Chenhao Tan*. *Selective Explanations: Leveraging Human Input to Align Explainable AI*. CSCW 2023.

Raul Castro Fernandez, Aaron J. Elmore, Michael J. Franklin, Sanjay Krishnan, and *Chenhao Tan*. *How large language models will disrupt data management*. VLDB 2023.

Chacha Chen, Shi Feng, Amit Sharma, and *Chenhao Tan*. *Machine Explanations and Human Understanding*. Transactions on Machine Learning Research (TMLR); FAccT 2023; **Best paper award** at the ICML 2022 Workshop on Human-Machine Collaboration and Teaming.

Vivian Lai, Chacha Chen, Alison Smith-Renner, Q. Vera Liao, and *Chenhao Tan*. *Towards a Science of Human-AI Decision Making: A Survey of Empirical Studies*. FAccT 2023.

Mourad Heddaya, Solomon Dworkin, *Chenhao Tan*, Rob Voigt, and Alexander Zentefis. *Language of Bargaining*. ACL 2023.

Yangqiaoyu Zhou, Yiming Zhang, and *Chenhao Tan*. *FLamE: Few-shot Learning from Natural Language Explanations*. ACL 2023.

Han Liu, Yizhou Tian, Chacha Chen, Shi Feng, Yuxin Chen, and *Chenhao Tan*. *Learning Human-Compatible Representations for Case-Based Decision Support*. In Proceedings of ICLR 2023. Presented in the ICML 2022 Workshop on Human-Machine Collaboration and Teaming.

Yiming Zhang, Yangqiaoyu Zhou, Samuel Carton, and *Chenhao Tan*. *Learning to Ignore Adversarial Attacks*. In Proceedings of EACL 2023.

Yuan Yuan, Tracy Xiao Liu, *Chenhao Tan*, Qian Chen, Alex 'Sandy' Pentland, and Jie Tang. *Gift Contagion in Online Groups: Evidence From Virtual Red Packets*. Management Science. 2023.

Xuan Yang, Yang Yang, *Chenhao Tan*, Yinghe Lin, Zhengzhe Fu, Fei Wu, and Yueting Zhuang. *Unfolding and Modeling the Recovery Process after COVID Lockdowns*. Scientific Reports. 2023.

Abhinav Kumar, *Chenhao Tan*, and Amit Sharma. *Probing Classifiers are Unreliable for Concept Removal and Detection*. NeurIPS 2022.

- Yiming Zhang, Shi Feng, and Chenhao Tan. *Active Example Selection for In-Context Learning*. EMNLP 2022.
- Vivian Lai, Samuel Carton, Rajat Bhatnagar, Q. Vera Liao, Yunfeng Zhang, and Chenhao Tan. *Human-AI Collaboration via Conditional Delegation: A Case Study of Content Moderation*. In Proceedings of the ACM CHI Conference on Human Factors in Computing Systems 2022 (CHI 2022).
- Samuel Carton, Surya Kanoria, and Chenhao Tan. *What to Learn, and How: Toward Effective Learning from Rationales*. In Findings of ACL 2022.
- Chenhao Tan. *On the Diversity and Limits of Human Explanations*. NAACL 2022.
- Jamar L. Sullivan, Will Brackenbury, Andrew McNutt, Kevin Bryson, Kwam Byll, Yuxin Chen, Michael Littman, Chenhao Tan, Blase Ur. *Explaining Why: How Instructions and User Interfaces Impact Annotator Rationales When Labeling Text Data*. In Proceedings of the 2022 Annual Conference of the North American Chapter of the Association for Computational Linguistics (NAACL 2022).
- C. Estelle Smith, Irfanul Alam, Chenhao Tan, Brian C. Keegan, Anita Blanchard. *The Impact of Governance Bots on Sense of Virtual Community: Development and Validation of the GOV-BOTs Scale*. In Proceedings of The 25th ACM Conference On Computer-Supported Cooperative Work And Social Computing (CSCW 2022).
- Chao-Chun Hsu and Chenhao Tan. *Decision-focused Summarization*. In Proceedings of The 2021 Conference on Empirical Methods in Natural Language Processing (EMNLP 2021).
- Han Liu, Vivian Lai, and Chenhao Tan. *Understanding the Effect of Out-of-distribution Examples and Interactive Explanations on Human-AI Decision Making*. In Proceedings of 2021 Conference on Computer Supported Cooperative Work and Social Computing (CSCW 2021).
- Madhusudhan Aithal Mahabhaleshwara and Chenhao Tan. *On Positivity Bias in Negative Reviews*. In Proceedings of the Joint Conference of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (ACL-IJCNLP 2021) (short paper).
- Ramaravind Kommiya Mothilal, Divyat Mahajan, Chenhao Tan and Amit Sharma. *Towards Unifying Feature Attribution and Counterfactual Explanations: Different Means to the Same End*. In Proceedings of the Fourth AAAI/ACM conference on Artificial Intelligence, Ethics, and Society (AIES'2021).
- Jason Shuo Zhang, Qin Lv, Brian Keegan, and Chenhao Tan. *Understanding the Diverging User Trajectories in Highly-related Online Communities during the COVID-19 Pandemic*. In Proceedings of the 15th International Conference On Web and Social Media (ICWSM'2021).
- Isil Erel, Lea Stern, Chenhao Tan, Michael Weisbach. *Selecting Directors Using Machine Learning*. Review of Financial Studies. 2021.
- Abhijit Suresh, Jennifer Jacobs, Vivian Lai, Chenhao Tan, Wayne Ward, James H. Martin, Tamara Sumner. *Using Transformers to Provide Teachers with Personalized Feedback on their Classroom Discourse: The TalkMoves Application*. AAAI Spring Symposium Series 2021.
- Chao-Chun Hsu, Shantanu Karnwal, Sendhil Mullainathan, Ziad Obermeyer, and Chenhao Tan. *Characterizing the Value of Information in Medical Notes*. In Findings of EMNLP 2020.
- Samuel Carton, Anirudh Rathore, and Chenhao Tan. *Evaluating and Characterizing Human Rationales*. In Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP'2020).

Vivian Lai, Han Liu, and Chenhao Tan. “Why is ‘Chicago’ deceptive?”: Towards Building Model-Driven Tutorials for Humans. In Proceedings of the ACM CHI Conference on Human Factors in Computing Systems 2020 (CHI’2020).

Ramaravind Kommiya Mothilal, Amit Sharma, and Chenhao Tan. Explaining Machine Learning Classifiers through Diverse Counterfactual Explanations. In Proceedings of ACM FAccT Conference 2020.

Vivian Lai, Samuel Carton, and Chenhao Tan. Harnessing Explanations to Bridge AI and Humans. In Fair & Responsible AI Workshop, CHI 2020.

Brian Lubars and Chenhao Tan. Ask not what AI can do for you, but what AI should do: Towards a framework of task delegability. In Proceedings of Thirty-third Conference on Neural Information Processing Systems (NeurIPS’2019) (spotlight presentation).

David Atkinson, Kumar Bhargav Srinivasan, and Chenhao Tan. What Gets Echoed? Understanding the ‘Pointers’ in Explanations of Persuasive Arguments. In Proceedings of 2019 Conference on Empirical Methods in Natural Language Processing and 9th International Joint Conference on Natural Language Processing (EMNLP’2019).

Vivian Lai, Jon Z. Cai, and Chenhao Tan. Many Faces of Feature Importance: Comparing Built-in and Post-hoc Feature Importance in Text Classification. In Proceedings of 2019 Conference on Empirical Methods in Natural Language Processing and 9th International Joint Conference on Natural Language Processing (EMNLP’2019).

Vivian Lai and Chenhao Tan. On Human Predictions with Explanations and Predictions of Machine Learning Models: A Case Study on Deception Detection. In Proceedings of ACM FAccT Conference 2019.

Kelvin Luu, Chenhao Tan, and Noah A. Smith. Measuring Online Debaters’ Persuasive Skill from Text over Time. Transactions of the Association for Computational Linguistics (TACL’2019).

Xiaochuang Han, Eunsol Choi, and Chenhao Tan. No Permanent Friends or Enemies: Tracking Relationships between Nations from News. In Proceedings of the 2019 Annual Conference of the North American Chapter of the Association for Computational Linguistics (NAACL’2019).

Jason Shuo Zhang, Chenhao Tan, and Qin Lv. Intergroup Contact in the Wild: Characterizing Language Differences between Intergroup and Single-group Members in NBA-related Discussion Forums. In Proceedings of the Conference on Computer-Supported Cooperative Work and Social Computing (CSCW’2019).

Kumar Bhargav Srinivasan, Cristian Danescu-Niculescu-Mizil, Lillian Lee, and Chenhao Tan. Content Removal as a Moderation Strategy: Compliance and Other Outcomes in the ChangeMyView Community. In Proceedings of the Conference on Computer-Supported Cooperative Work and Social Computing (CSCW’2019).

Tiago Cunha, David Jurgens, Chenhao Tan, and Daniel Romero. Are All Successful Communities Alike? Characterizing and Predicting the Success of Online Communities. In Proceedings of The Web Conference 2019 (WWW’2019).

Ziqiang Cheng, Yang Yang, Chenhao Tan, Denny Cheng, Alex Cheng, and Yueting Zhuang. What Makes a Good Team? A Large-scale Study on the Effect of Team Composition in Online Gaming. In Proceedings of The Web Conference 2019 (short papers, WWW’2019).

Jason Shuo Zhang, Chenhao Tan, and Qin Lv. “This is why we play”: Characterizing Online Fan Communities of the NBA Teams. In Proceedings of the Conference on Computer-Supported Cooperative Work and Social Computing (CSCW’2018).

Tal August, Nigini Oliveira, *Chenhao Tan*, Noah A. Smith, and Katharina Reinecke. *Framing Effect: The Choice of Slogans Used to Advertise Online Experiments Can Boost Recruitment and Lead to Sample Biases*. In Proceedings of the Conference on Computer-Supported Cooperative Work and Social Computing (CSCW'2018).

Dallas Card, *Chenhao Tan*, and Noah A. Smith. *Neural Models for Documents with Metadata*. In Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (ACL'2018).

Nelson F. Liu, Omer Levy, Roy Schwartz, *Chenhao Tan*, and Noah A. Smith. *LSTMs Exploit Linguistic Attributes of Data*. In Proceedings of ACL 2018 RepL4NLP Workshop (**Best paper award**).

Chenhao Tan. *Tracing Community Genealogy: How New Communities Emerge from the Old*. In Proceedings of the 12th International AAAI Conference on Web and Social Media (ICWSM'18).

Chenhao Tan, Hao Peng and Noah A. Smith. *"You are no Jack Kennedy": On Media Selection of Highlights from Presidential Debates*. In Proceedings of The Web Conference (WWW'2018).

Yang Yang, Zongtao Liu, *Chenhao Tan*, Fei Wu, Yueting Zhuang and Yafeng Li. *To Stay or to Leave: Churn Prediction for Urban Migrants in the Initial Period*. In Proceedings of The Web Conference (WWW'2018).

Isil Erel, Léa H. Stern, *Chenhao Tan*, and Michael S. Weisbach. *Could Machine Learning Help Companies Select Better Board Directors?* Harvard Business Review. April, 2018. (*The associated manuscript is under submission*).

Elizabeth Clark, Anne Ross, *Chenhao Tan*, Yangfeng Ji, Noah A. Smith. *Creative Writing with a Machine in the Loop: Case Studies on Slogans and Stories*. In Proceedings of the 23rd ACM Conference on Intelligent User Interfaces (IUI'2018).

Yang Yang, *Chenhao Tan*, Zongtao Liu, Fei Wu, Yueting Zhuang. *Urban Dreams of Migrants: A Case Study of Migrant Integration in Shanghai*. In Proceedings of the 32nd AAAI Conference on Artificial Intelligence (AAAI'2018).

Yangfeng Ji, *Chenhao Tan*, Sebastian Martschat, Yejin Choi, Noah A. Smith. *Dynamic Entity Representations in Neural Language Models*. In Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing (EMNLP'2017).

Chenhao Tan, Dallas Card, Noah A. Smith. *Friendships, Rivalries, and Trysts: Characterizing Relations between Ideas in Texts*. In Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (ACL'2017).

Chenhao Tan, Vlad Niculae, Cristian Danescu-Niculescu-Mizil, Lillian Lee. *Winning arguments: Interaction dynamics and persuasion strategies in good-faith online discussions*. In Proceedings of the 25th International World Wide Web Conference (WWW'2016).

Isabel Kloumann, *Chenhao Tan*, Jon Kleinberg, Lillian Lee. *Internet Collaboration on Extremely Difficult Problems: Research versus Olympiad Questions on the Polymath Site*. In Proceedings of the 25th International World Wide Web Conference (WWW'2016).

Chenhao Tan, Adrien Friggeri, Lada Adamic. *Lost in propagation? Unfolding news cycles from the source*. In Proceedings of 10th International AAAI Conference on Weblogs and Social Media (ICWSM'2016).

Jack Hessel, *Chenhao Tan*, Lillian Lee. *Science, AskScience, and BadScience: On the Coexistence of Highly Related Communities*. In Proceedings of 10th International AAAI Conference on Weblogs and Social Media (ICWSM'2016).

Chenhao Tan, Lillian Lee. *All Who Wander: On the Prevalence and Characteristics of Multi-community Engagement*. In Proceedings of the 24th International World Wide Web Conference (WWW'2015).

Chenhao Tan, Lillian Lee, Bo Pang. *The effect of wording on message propagation: Topic- and author-controlled natural experiments on Twitter*. In Proceedings of the Association for Computational Linguistics (ACL'2014).

Chenhao Tan, Lillian Lee. *A Corpus of Sentence-level Revisions in Academic Writing: A Step towards Understanding Statement Strength in Communication*. In Proceedings of the Association for Computational Linguistics (ACL'2014) (short paper).

Jon Kleinberg, Sendhil Mullainathan, Chenhao Tan, Thomas Zimmermann. *Inductive Testing: Theory and an Application to the Disposition Effect*. NBER Behavioral Finance Working Group, 2014.

Chenhao Tan, Ed H. Chi, David Huffaker, Gueorgi Kossinets and Alexander J. Smola. *Instant Foodie: Predicting Expert Ratings From Grassroots*. In Proceedings of 22nd ACM International Conference on Information and Knowledge Management (CIKM'2013).

Daniel M. Romero, Chenhao Tan and Johan Ugander. *On the Interplay between Social and Topical Structure*. In Proceedings of 7th International AAAI Conference on Weblogs and Social Media (ICWSM'2013).

Eunsol Choi, Chenhao Tan, Lillian Lee, Cristian Danescu-Niculescu-Mizil and Jennifer Spindel. *Hedge detection as a lens on framing in the GMO debates: A position paper*. In Proceedings of the Workshop on Extra-Propositional Aspects of Meaning in Computational Linguistics (2012).

Chenhao Tan, Evgeniy Gabrilovich, and Bo Pang. *To Each His Own: Personalized Content Selection based on Text Comprehensibility*. In Proceedings of the Fifth ACM International Conference on Web Search and Data Mining (WSDM'2012).

Wenbin Tang, Jie Tang, Tao Lei, Chenhao Tan, Bo Gao, and Tian Li. *On optimization of expertise matching with various constraints*. In Neurocomputing (July, 2012).

Chenhao Tan, Lillian Lee, Jie Tang, Long Jiang, Ming Zhou, and Ping Li. *User-level sentiment analysis incorporating social networks*. In Proceedings of the Seventeenth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (SIGKDD'2011).

Chenhao Tan, Jie Tang, Jimeng Sun, Quan Lin, and Fengjiao Wang. *Social Action Tracking via Noise Tolerant Time-varying Factor Graphs*. In Proceedings of the Sixteenth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (SIGKDD'2010).

Bin Lu, Chenhao Tan, Claire Cardie and Benjamin K. Tsou. *Joint Bilingual Sentiment Classification with Unlabeled Parallel Corpora*. In Proceedings of the Association for Computational Linguistics (ACL'2011).

Shaomei Wu, Chenhao Tan, Jon Kleinberg, and Michael Macy. *Does Bad News Go Away Faster?* In Proceedings of Fifth International AAAI Conference on Weblogs and Social Media (ICWSM'2011) (short paper).

Fei Wang, Chenhao Tan, Arnd Christian Konig, and Ping Li. *Efficient Document Clustering via Online Nonnegative Matrix Factorizations*. In Proceedings of the SIAM International Conference on Data Mining (SDM'2011).

Wenbin Tang, Jie Tang, and Chenhao Tan. *Expertise Matching via Constrained-based Optimization*. In Proceedings of 2010 IEEE/WIC/ACM International Conferences (WIC'2010).

Invited talks

Towards Human-Centered AI: Understanding and Improving Human Decision Making through Explanations.

Columbia University, 2023.

Georgia Tech, 2023.

MIT, 2023.

University of Washington, 2023.

Machine Learning in Economics Summer Institute, 2023.

Three Lessons Towards Human-Centered Explanations. Keynote at the Artificial Intelligence & Human Computer Interaction workshop at ICML 2023.

The Future of AI is Human-Centered. Keynote at the Research Day at the Governors State University, 2023.

Towards Human-Centered Explanations of AI Predictions.

Stanford University, 2023.

Cambridge University, 2023.

University of Illinois at Chicago, 2023.

Tutorial on Human-Centered Evaluation of Explanations. NAACL 2022.

Towards Human-Centered Explanations of AI Predictions.

University of Pennsylvania, 2022.

University of Toronto, 2022.

New York University, 2022.

University of Texas at Austin, 2022.

TTIC, 2022.

Language, Information, and Decision Making.

Machine Learning in Economics Summer Institute, 2022.

Brown Bag at Booth Healthcare Initiative, 2022.

Why are Explanations Falling Short of their Promises? AAAI workshop on Explainable Agency in Artificial Intelligence. 2022.

Towards Human-Centered AI: Complementing Humans with AI. The Academic Fringe Festival on Designing at Scale with Human-AI Collaboration. 2022.

Characterizing the "Value" of Information, Toyota Technological Institute, 2021

Towards Human-Centered AI: Harnessing Explanations to Improve Human Decision Making in Discovery Tasks, the second ELLIS HCML workshop at ICLR, 2021

Towards Human-Centered AI: Harnessing Explanations to Improve Human Decision Making

University of Wisconsin at Madison, 2021.

Lunch for CAPP at the University of Chicago, 2021.

A Multitude of Communities: Genealogy and Intergroup Contact, University of Notre Dame, 2021

Towards Human-Centered AI: Understanding Human Production and Consumption of Explanations

University of Maryland, 2021

University of Michigan, 2021

Towards Human-Centered AI: Understanding Language in Social Contexts

University of Chicago, 2020

CMU, 2020

Google, 2020

Microsoft Research, 2020

Yale University, 2020

UIUC, 2020

University of Washington, 2020

Towards Human-Centered AI: Harness AI to Improve Human Decision Making

Harris School of Public Policy, University of Chicago, 2020

Tsinghua University, 2020

USC AI Rising Stars Symposium, 2019

Towards Human-Centered AI: Opportunities and Challenges for Natural Language Processing in Policy, Deploying Machine Learning Tools for Public Policy Impact, APPAM pre-conference workshop, 2019.

The Promise of Human-centered AI for Addressing Challenges in Literacy, Facebook, 2019

Towards Human-Centered AI: Understanding Language in Social Contexts, ASAPP, 2019

The Future of Brand Engagement: Community, Trust and Digital Discourse, Reddit, 2018

Towards Understanding and Generating Language for Effective Social Interaction, IST Seminar, Caltech, 2018

Characterizing the Ecosystem of Ideas in Texts

Summer Institute in Computational Social Science, University of Colorado Boulder, 2018

The Fall 2018 Workshop on Quantitative Methods, University of Colorado Boulder, 2018

Department of Computer Science, Tsinghua University, 2018

Computational Semantics Group at University of Colorado Boulder, 2017

UW-MS NLP Symposium, 2017

Winning Arguments: Interaction Dynamics and Persuasion Strategies in Good-Faith online discussions

The 24th Annual Northwest Dispute Resolution Conference, 2018

Colorado Data Science Team at University of Colorado Boulder, 2017

Phrasing Matters: An Interactive Session on Message Virality Based on Twitter, the 24th Annual Northwest Dispute Resolution Conference, 2018

Wandering Individuals and Emerging Communities in a World of Communities

Department of Sociology, Tsinghua University, 2018

INFO/HCC Seminar at University of Colorado Boulder, 2017

It Takes Two to Tango: Understanding the Effects of Wording through Online Interactions

Fudan University, 2018

Toutiao AI lab, 2018

Zhejiang University, 2017

A Short Talk on Research, Zhejiang University, 2017

Online social interactions: a lens on humans and a world for humans

CSSS Seminar at University of Washington, 2016

MIT EECS, 2016

CMU Computer Science Department, 2016

CMU Heinz College, 2016

Max Planck Institute for Software Systems, 2016

University of Michigan, 2016

University of Colorado Boulder, 2016

Rice University, 2016

Northeastern University, 2016

It takes two to tango: Understanding the effects of language via “natural experiments”, AI seminar at Cornell University, 2014

The effect of wording on message propagation: Topic- and author-controlled natural experiments on Twitter, Tsinghua University, 2014

More than Accuracy: Interpretability, Machine learning reading group at Cornell University, 2013

User-level sentiment analysis incorporating social networks

Tsinghua University, 2011

Microsoft Research Asia NLP group seminar, 2011

Professional Service

Steering committee of FAccT (2023-2026)

Faculty advisory board for Cognitive Science, University of Chicago (2021.09 — present)

Executive committee member at the Institute of Cognitive Science, CU Boulder (2019.10 — 2020.12)

General chair:

2023: FAccT

Program Chair:

2022: IC2S2

Senior Area Chair:

2023: ACL

2022: AACL

2021: NAACL

Area Chair (Senior PC):

2021: ACL, CSCW, AKBC

2020: EMNLP

2019: NAACL, EMNLP, AAI

Tutorial Chair for ICWSM 2020**PC Member:**

2022: NAACL, FAccT, EMNLP

2021: CSCW

2020: WWW

2019: WWW, FAT*, ACL, NeurIPS (one of 400 highest-scoring reviewers)

2018: WWW, NAACL, ACL, KDD, NeurIPS, ICWSM, BigNet Workshop

2017: WWW, WSDM, KDD, ACL, EMNLP, CONLL

2016: WWW, KDD, WSDM, IJCAI, NAACL, ICWSM, SDM, CIKM

2015: WWW, KDD, WSDM

2014: WWW, EMNLP, ICWSM, DYAD

2013: IJCAI, IJCNLP, NAACL Student Research Workshop

2012: ADMA

2011: ADMA

Panelists: NSF (2021, 2020, 2019, 2019)**Session chair:** NAACL (2019), WWW (2015)**Reviewer:** JMLR (2023), JMLR(2021), PNAS (2020), ICWSM (2019), HICSS (2019), TACL (2020, 2019, 2018, 2017, 2016), Transactions on Knowledge Discovery from Data (2015, 2012), TOIS (2018), Plos ONE (2018), CHI (2017), CSCW (2016), WWW (2011, 2010), ICDM (2011), KDD (2010), CIKM (2010)

Selected Media Coverage

New York Times (2023): How Easy Is It to Fool A.I.-Detection Tools?

PolitiFact (2023): How improperly using AI could deter you from voting, or hurt your health

CU Boulder Today (2020): As the coronavirus spread, 2 social media communities drifted apart.

The New York Times (2019): How to Argue on the Internet Without Losing Your Mind.

9NEWS (2019): Study of NBA Reddit forum hopes to learn more about how people interact with each other online.

CU Boulder Today (2019): Flagrant fouls: What Reddit's basketball fans can tell us about online discourse.

MIT Technology Review (2019): WeChat is running a natural experiment in human generosity.

Harvard Business Review (2018): Want Less-Biased Decisions? Use Algorithms.

NPR (2017): Change My View On Reddit Helps People Challenge Their Own Opinions.

New Scientist (2017): Phone metadata reveals where city migrants go and who they call.

You are not so smart podcasts (2016): Change My View.

Slate (2016): You Can't Win: How algorithms are changing the ways we argue.

The Washington Post (2016): How to change someone's mind, according to science.

New York Magazine (2016): A Subreddit Sparked a Scientific Inquiry Into How to Change Someone's Mind.

Daily Mail (2016): How to win ANY argument using science: Experts reveal which words to use and how one simple trick can help you get your own way.

The New York Times (2014): A 25-Question Twitter Quiz to Predict Retweets.

The Washington Post (2014): This is the perfect tweet, according to science.

The Physics arXiv Blog (2014): Data Mining Reveals How Wording Influences Tweet Propagation.

engadget (2014): Crafting the perfect tweet is one-third talent, two-thirds science.

The Daily Dot (2014): A data miner's guide to crafting the perfect tweet.

Postdoctorates

Shi Feng (2023) → Research Associate at the New York University → Assistant professor at George Washington University

Sam Carton (2022) → Assistant professor at the University of New Hampshire

Estelle Smith (2022) → Assistant professor at Colorado School of Mines

PhD Students

Chacha Chen

Mourad Heddaya

Chao-Chun Hsu

Han Liu

Karen Zhou

Yangqiaoyu (Rosa) Zhou

Vivian Lai (2022) → Visa Research

Jason Shuo Zhang (2020, co-advised with Qin Lv) → postdoc at the Max Planck Institute → ASOS → Tripadvisor

Master Students (thesis/independent study)

Tejes Srivastava (class of 2025, University of Chicago)
Yixuan Wang (class of 2025, University of Chicago)
Jessica Ji (class of 2024, University of Chicago)
Jiamin Yang (class of 2024, University of Chicago)
Zander Meitus (class of 2023, University of Chicago)
Harry Tian (class of 2023, University of Chicago)
Yiming Zhang (class of 2023, University of Chicago)
Madhusudhan AithalMahabhaleshwara (class of 2021, University of Colorado Boulder).
Rajat Bhatnagar (class of 2021, University of Colorado Boulder).
Anirudh Rathore (class of 2021, University of Colorado Boulder).
Surya Kanoria (class of 2020, University of Colorado Boulder)
Brian Lubars (class of 2019, University of Colorado Boulder).
Shantanu Karnwal (class of 2019, University of Colorado Boulder).
Hunter Wapman (class of 2019, University of Colorado Boulder).
Kumar Srinivasan (class of 2019, University of Colorado Boulder).
Vivian Lai (class of 2018, University of Colorado Boulder).

Undergraduate/High School Advising

Miles Wang (class of 2026, Harvard University)
Tianyu Cao (class of 2024, Zhejiang University)
Lily Wang (class of 2024, University of Chicago)
Dang Nguyen (class of 2024, University of Chicago)
Grace Wang (class of 2024, University of Chicago)
Aditya Krishna (class of 2024, University of Chicago)
Milo Chase (class of 2023, University of Chicago)
Kenan Tang (class of 2023, University of Chicago)
Brent DeVries (class of 2022, University of Chicago)
Ava Bartolome (class of 2022, Clark University → PhD at the University of Minnesota)
Qitian Hu (class of 2022, University of Chicago)
Martin Chou (class of 2021, Tsinghua University)

Anne Mykland (class of 2023, University of Chicago Laboratory School)

Yangqiaoyu Zhou (class of 2020, Carleton College)

Jie Wang (class of 2020, University of Colorado Boulder)

Zhenqi Li (class of 2020, University of Colorado Boulder)

Xiaochuang Han (class of 2019, Georgia Tech).

David Atkinson (class of 2019, University of Colorado Boulder).

Cora Schneck (class of 2018, University of Colorado Boulder).

Kelvin Luu (class of 2015, Cornell).

Eunsol Choi (class of 2012, Cornell, honorable mention for the CRA Outstanding Undergraduate Research Award in 2011).

Committee Members

I served as a committee member for the following students: Jiaming Qu (PhD, University of North Carolina), Afia Khan (MD/PhD, University of Chicago), Nan-Jiang Jiang (PhD, Ohio State University), Joon Sik Kim (PhD, Carnegie Mellon University), Valdemar Danry (master, MIT), Gil Rosenthal (master, University of Chicago), Jingwen Jiang (PhD, University of Chicago), Divyansh Kaushik (PhD, Carnegie Mellon University), Rehan Ahmed (PhD, University of Colorado Boulder), Anupama (master, University of Colorado Boulder), Jon Cai (PhD, University of Colorado Boulder), Nora Connor (PhD, University of Colorado Boulder), Xiaohan Ding (PhD, University of Chicago), Xu Han (PhD, University of Colorado Boulder), Shudong Hao (PhD, University of Colorado Boulder), Daniel Leopold (Psychology PhD, University of Colorado Boulder), Forough Poursabzi (PhD, University of Colorado Boulder), Tyler Scott (undergraduate, University of Colorado Boulder), Saumya Sinha (master, University of Colorado Boulder), Hayeong Song (master, University of Colorado Boulder), Angela Stewart (PhD, University of Colorado Boulder), Shubha Swamy (undergraduate, University of Colorado Boulder), Yichen Wang (PhD, University of Colorado Boulder), Hunter Wapman (PhD, University of Colorado Boulder), Max Wenzel (undergraduate, University of Colorado Boulder), Mark Wilmes (undergraduate, University of Colorado Boulder), and Linzi Xing (master, University of Colorado Boulder).

Industrial Experience

Data Scientist Intern, 2015.6-2015.8,

Facebook, Menlo Park, CA,

Understanding news propagation.

Research Intern, 2012.5-2012.8,

Google, Mountain View, CA

Harmonizing place scores & instant foodie.

Research Intern, 2011.6-2011.8,

Yahoo! Research, Santa Clara, CA

Personalization based on text comprehensibility.

Research Intern, 2010.4-2010.7,

Microsoft Research Asia, Beijing, China,

Sentiment analysis incorporating social networks.

Teaching Experience

Instructor, 03/2022-05/2022, 03/2023-05/2023

University of Chicago, Chicago IL

Machine Learning for Public Policy.

Instructor, 01/2021-03/2021, 01/2022-03/2022, 01/2023-03/2023

University of Chicago, Chicago IL

Human-Centered Machine Learning.

Instructor, 01/2021-03/2021, 01/2022-03/2021

University of Chicago, Chicago IL

Natural language processing.

Instructor, 08/2018-12/2018, 08/2020-12/2020

University of Colorado, Boulder, CO,

Undergraduate Machine Learning, CSCI 4622.

Instructor, 01/2018-05/2018, 01/2019-05/2019

University of Colorado, Boulder, CO,

Human-centered Machine Learning, CSCI 7000.

Instructor, 08/2017-12/2017, 08/2019/2019

University of Colorado, Boulder, CO,

Machine Learning, CSCI 5622.

Teaching Assistant, 08/2013-12/2013,

Cornell University, Ithaca, NY,

Information Retrieval, CS4300/INFO4300.

Teaching Assistant, 08/2011-12/2011,

Cornell University, Ithaca, NY,

Machine Learning, CS4780/5780.