

William Yang Wang

Mellichamp Professor of Artificial Intelligence
Director, Responsible Machine Learning Center
Director, UCSB NLP Group
University of California, Santa Barbara
Henley Hall 2005
Santa Barbara, California, 93106-5110

Phone: (412) 370-8296
Email: william@cs.ucsb.edu
Web URL: <http://www.cs.ucsb.edu/~william/>

Education

- **Ph.D., Computer Science - Language Technologies**
Carnegie Mellon University, 2011 - 2016
Advisor: William W. Cohen (Machine Learning Department).
Additional thesis committee members: Tom Mitchell, Christos Faloutsos, Eric Horvitz.
- **M.S., Computer Science - Natural Language Processing**
Columbia University, 2009 - 2011
Advisors: Kathleen McKeown and Julia Hirschberg.
- **B.Eng., Computer Science**
School of Information Engineering, Shenzhen University, 2005 - 2009.

Research Interests

- Natural Language Processing (Large Language Models, Language & Vision)
- Machine Learning (Science of LLMs and Generative AI)
- Artificial Intelligence (Neuro-Symbolic AI & Reasoning, Responsible AI)

Experience

- **University of California, Santa Barbara, CA** 06/2016 – now.
Assistant Professor (2016-2021), Associate Professor (2021-2023), and Professor (2023-now). Director of UCSB Responsible Machine Learning Center, Director of UCSB NLP Group and Mellichamp Professor of Artificial Intelligence and Designs
 1. I have organized a research agenda centered around both theories and practices of natural language processing and machine learning, and we made some good progress in self-supervised data-to-knowledge generation, open-domain question answering, vision-and-language generation, knowledge graph reasoning, science of large language models, neuro-symbolic reasoning, and responsible AI
 2. I am founding director of UCSB's Center for Responsible Machine Learning, which is the interdisciplinary center for 60+ faculty from 20+ departments. I also co-direct UCSB's Mind and Machine Intelligence Initiative, which centers three departments including Economics, Psychology & Brain Sciences, and Computer Science.

3. I have received several prestigious awards such as the 2018 DARPA Young Faculty Award, 2020 IEEE AI's 10 to Watch, 2021 NSF CAREER Award, 2022 Karen Spärck Jones Award, 2023 CRA-E Undergraduate Research Faculty Mentoring Award, 2024 Simon-Pierre Laplace Award from IEEE SPS, and faculty research awards from Google, Facebook, Adobe, Intel, Amazon, JP Morgan, CISCO, Apple, and IBM.
 4. I built a very productive NLP/ML/AI research group in two years, and according to CSRankings.org, we are top-8 in NLP in the year of 2018-2022 in the America.
 5. I have raised USD \$20M+ in research funding in the past few years.
- **Carnegie Mellon University**, Pittsburgh, PA 01/2013 – 05/2016.
Research Fellow, Advisor: William W. Cohen
 - **Yahoo! Labs**, New York, NY 05/2015 – 08/2015.
Research Scientist Intern, Mentors: Amanda Stent, Drago Radev, and Yashar Mehdad.
 - **Microsoft Research**, Redmond, WA 05/2012 – 08/2012 and 10/2012 – 01/2013.
Research Intern, Mentors: Dan Bohus, Ece Kamar, and Eric Horvitz.
 - **Carnegie Mellon University**, Pittsburgh, PA 09/2011 – 05/2012.
R. K. Mellon Presidential Fellow, Advisors: Justine Cassell and Alan W Black.
 - **Columbia University**, New York, NY 09/2009 – 08/2011.
 - *Research Assistant*, Natural Language Processing Group, Advisor: Kathleen McKeown. 09/2010 – 05/2011
 - *Research Assistant*, Spoken Language Processing Group, Advisor: Julia Hirschberg. 09/2009 – 05/2011
 - *Visiting Research Assistant*, Biomedical Informatics, Host: Noemie Elhadad. 05/2011 – 07/2011 notes.
 - *Visiting Research Assistant*, School of International and Public Affairs, Host: Suresh Naidu. 07/2011 – 08/2011

Major Awards and Honors

- IEEE SPS Pierre-Simon Laplace Early Career Technical Achievement Award, 2024
- CRA-E Undergraduate Research Faculty Mentoring Award, CRA 2023
- Senior Member, AAAI 2023
- Amazon Alexa Prize Runner-Up, Simbot, 2023
- Karen Spärck Jones Award - British Computer Society, 2022
- Google Ads Faculty Research Award, 2022
- ACM Senior Member, 2022

- IEEE Senior Member, 2022
- IJCAI Early Career Spotlight (declined), 2022
- AMiner AI 2000 Most Influential Scholar Honorable Mention, 2022
- Amazon Alexa Prize Finalist, Taskbot, 2021
- NSF CAREER Award, 2021
- JP Morgan Chase Faculty Research Award, 2021
- AI's 10 to Watch: the Future of AI Award, by IEEE Intelligent Systems, 2020
- Google Focused Research Award, 2020
- Duncan and Suzanne Mellichamp Chair in Artificial Intelligence and Designs, 2020
- JP Morgan Chase Faculty Research Award, 2020
- Amazon AWS Machine Learning Research Award, 2020
- CVPR Best Student Paper Award, 2019 (awarded to 1 out of 5160 submissions)
- Facebook Research Award, 2019
- Google Faculty Research Award, 2019-2020
- IBM Faculty Award, 2019
- Member, ACM Future of Computing Academy, 2019
- Google Faculty Research Award, 2018
- DARPA Young Faculty Award, 2018
- IBM Faculty Award, 2018
- Intel Faculty Grants, 2018
- Facebook Research Award, 2018
- Adobe Research Award, 2018
- IBM Faculty Award, 2017
- Notable Data Set Award, EMNLP 2015 (awarded to 2 out of 1300 submissions)
- Best Paper Honorable Mention Award, CIKM 2013 (awarded to 3 out of 848 submissions)
- Best Student Paper Award, ASRU 2013 (awarded to 1 out of 172 submissions)

Research Grants, Sponsorships, and Awards

Total amount: \$20M+.

- 2023 CISCO Faculty Research Award.
- 2023 Google Deepmind Continuing Award.
- 2023 Google GenAI Faculty Award.
- 2023 Army UARC Grant.
- 2023 Noyce Foundation Grant.
- 2023 JP Morgan Chase Faculty Award.
- 2022 IEE IGSB Grant.
- 2022 Amazon AWS Faculty Award.
- 2022 Google Ads Faculty Research Award.
- 2022 Google AI Continuing Grant.
- 2021 IARPA COVID-19 Seedling.
- 2021 Facebook AI Sponsored Research Project.
- 2021 Noyce Trust Mind and Machine Intelligence.
- 2021 NSF CAREER Award.
- 2021 Amazon Alexa Prize - Taskbot.
- 2021 Amazon Alexa Prize - Simbot.
- 2021 Amazon Alexa AI - Knowledge Award.
- 2021 JP Morgan Chase Faculty Research Award.
- 2020 Google Focused Research Award.
- 2020 Google Ads Faculty Research Grant.
- 2020 UCSB Academic Senate Research Grant.
- 2020 JP Morgan Chase Faculty Research Award.
- 2020 Amazon AWS Machine Learning Research Award.
- 2019-2020 Google Faculty Research Award.
- 2019 Facebook Research Award.
- 2019 IEE Seed Grant.
- 2019 IBM Faculty Award.
- 2019 UCSB Academic Senate Research Grants.
- 2019 Google Gift Award.

- 2019 Army NSRDEC Grant.
- 2019 Google Faculty Research Award, Google.
- 2018 DARPA Young Faculty Award, DARPA.
- 2018 “Building NLP Tools for Healthy Online Conversations” and “Knowledge Guided Natural Language Generation”, Intel Faculty Grants.
- 2018 “Cross-Lingual and Open-World Task-Oriented Dialogue Schema Induction and Generation”, a Gift Fund from a leading Internet company.
- 2018 “Knowledge-Driven Neural Semantic Matching”, a leading software company.
- 2018 “Life-Long Open-World Knowledge Graph Reasoning”, IBM Faculty Award.
- 2018 “Low-Resource Machine Translation”, Facebook Research Award.
- 2018 “Language, Vision and Knowledge”, a Gift Fund from a leading Internet company.
- 2018 “Knowledge Graph Construction and Meeting Summarization”, LogMeIn Gift Fund.
- 2018 “Deep Inverse Reinforcement Learning for Visual Storytelling”, Adobe Research Award.
- 2018 NVIDIA Sponsorship: SoCal NLP Symposium.
- 2018 Disney Research Sponsorship: KBCOM workshop.
- 2018 Google AI Sponsorship: SoCal NLP Symposium.
- 2018 “Trust Score Inference from Networks”, UCSB Academic Senate Faculty Research Grant.
- 2018 Amazon AWS Cloud Credits for Research (formerly AWS Research Grants).
- 2017 “Reinforced Semi-Supervised Learning”, a gift fund from a leading Internet company.
- 2017 “Modeling Long-Term Semantic Dependencies for Dialogue”, a gift fund from a leading Internet company.
- 2017 “Multimodal Information Extraction”, IBM Faculty Award.
- 2017 “Automated Fake News Detection”, UCSB Academic Senate Faculty Research Grant.
- 2017 Google Cloud Platform Education Grant.
- 2016 Amazon AWS Cloud Credits for Research (formerly AWS Research Grants).
- 2016, 2017 NVIDIA GPU Grants, hardware gift.

Other Honors

- Best Reviewer Award, NAACL 2015 (awarded to a small fraction of reviewers)
- Best Presentation Runner-Up, LTI Student Research Symposium, Carnegie Mellon University, 2011

Peer-Reviewed Conference, Journal Papers, and Book Chapters

1. Jiachen Li, Qian Long, Jian Zheng, Xiaofeng Gao, Robinson Piramuthu, Wenhui Chen, William Yang Wang, "T2V-Turbo-v2: Enhancing Video Model Post-Training through Data, Reward, and Conditional Guidance Design", to appear in Proceedings of The Thirteenth International Conference on Learning Representations (**ICLR 2025**), Singapore.
2. Xinyi Wang, Antonis Antoniadis, Yanai Elazar, Alfonso Amayuelas, Alon Albalak, Kexun Zhang, William Yang Wang, "Generalization v.s. Memorization: Tracing Language Models' Capabilities Back to Pretraining Data", to appear in Proceedings of The Thirteenth International Conference on Learning Representations (**ICLR 2025**), Singapore.
3. Antonis Antoniadis, Albert Örwall, Kexun Zhang, Yuxi Xie, Anirudh Goyal, William Yang Wang, "Enhancing Software Agents with Monte Carlo Tree Search and Hindsight Feedback", to appear in Proceedings of The Thirteenth International Conference on Learning Representations (**ICLR 2025**), Singapore.
4. Xuehai He, Weixi Feng, Kaizhi Zheng, Yujie Lu, Wanrong Zhu, Jiachen Li, Yue Fan, Jianfeng Wang, Linjie Li, Zhengyuan Yang, Kevin Lin, William Yang Wang, Lijuan Wang, Xin Eric Wang, "MMWorld: Towards Multi-discipline Multi-faceted World Model Evaluation in Videos", to appear in Proceedings of The Thirteenth International Conference on Learning Representations (**ICLR 2025**), Singapore.
5. Wenda Xu, Rujun Han, Zifeng Wang, Long Le, Dhruv Madeka, Lei Li, William Yang Wang, Rishabh Agarwal, Chen-Yu Lee, Tomas Pfister, "Speculative Knowledge Distillation: Bridging the Teacher-Student Gap Through Interleaved Sampling", to appear in Proceedings of The Thirteenth International Conference on Learning Representations (**ICLR 2025**), Singapore.
6. Mian Zhang, Xianjun Yang, Xinlu Zhang, Travis Labrum, Jamie C. Chiu, Shaun M. Eack, Fei Fang, William Yang Wang, Zhiyu Chen, "CBT-Bench: Evaluating Large Language Models on Assisting Cognitive Behavior Therapy", in Proceedings of the 2025 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (**NAACL 2025**), Albuquerque, New Mexico, United States.
7. Kyle Wong, Alfonso Amayuelas, Liangming Pan, William Yang Wang, "Investigating the Transferability of Code Repair for Low-Resource Programming Languages", in Findings of the 2025 Annual Conference of the Nations of the Americas Chapter of the ACL (Findings of **NAACL 2025**), Albuquerque, New Mexico, United States.
8. Kexun Zhang, Shang Zhou, Danqing Wang, William Yang Wang, Lei Li, "Scaling LLM Inference Efficiently with Optimized Sample Compute Allocation", in Findings of the 2025 Annual Conference of the Nations of the Americas Chapter of the ACL (Findings of **NAACL 2025**), Albuquerque, New Mexico, United States.
9. Xinlu Zhang, Zhiyu Chen, Xi Ye, Xianjun Yang, Lichang Chen, William Yang Wang, Linda Ruth Petzold, "Unveiling the Impact of Coding Data Instruction Fine-Tuning on Large Language Models Reasoning", to appear in Proceedings of The 39th Annual AAAI Conference on Artificial Intelligence (**AAAI 2025**), Philadelphia, Pennsylvania, USA.
10. Shengqiong Wu, Hao Fei, Liangming Pan, William Yang Wang, Shuicheng YAN, Tat-Seng Chua, "Combating Multimodal LLM Hallucination via Bottom-up Holistic Reasoning", to appear in Proceedings of The 39th Annual AAAI Conference on Artificial Intelligence (**AAAI 2025**), Philadelphia, Pennsylvania, USA.

11. Michael Saxon, Fatima Jahara, Mahsa Khoshnoodi, Yujie Lu, Aditya Sharma, William Yang Wang, "Who Evaluates the Evaluations? Objectively Scoring Text-to-Image Prompt Coherence Metrics with T2IScoreScore (TS2)", to appear in Proceedings of The Thirty-eighth Annual Conference on Neural Information Processing Systems (**NeurIPS 2024**), Vancouver, Canada. Spotlight
12. Jiachen Li, Weixi Feng, Tsu-Jui Fu, Xinyi Wang, S Basu, Wenhui Chen, William Yang Wang, "T2V-Turbo: Breaking the Quality Bottleneck of Video Consistency Model with Mixed Reward Feedback", to appear in Proceedings of The Thirty-eighth Annual Conference on Neural Information Processing Systems (**NeurIPS 2024**), Vancouver, Canada.
13. Xiaobao Wu, Thong Thanh Nguyen, Delvin Ce Zhang, William Yang Wang, Anh Tuan Luu, "FASTopic: Pretrained Transformer is a Fast, Adaptive, Stable, and Transferable Topic Model", to appear in Proceedings of The Thirty-eighth Annual Conference on Neural Information Processing Systems (**NeurIPS 2024**), Vancouver, Canada.
14. Yujie Lu, Dongfu Jiang, Wenhui Chen, William Yang Wang, Yejin Choi, Bill Yuchen Lin, "WildVision: Evaluating Vision-Language Models in the Wild with Human Preferences", to appear in Proceedings of The Thirty-eight Conference on Neural Information Processing Systems Datasets and Benchmarks Track (**NeurIPS 2024 Datasets and Benchmarks Track**), Vancouver, Canada.
15. Yujie Lu, Pan Lu, Zhiyu Chen, Wanrong Zhu, Xin Eric Wang, William Yang Wang, "Multimodal Procedural Planning via Dual Text-Image Prompting", to appear in Findings of The 2024 Conference on Empirical Methods in Natural Language Processing (**Findings of EMNLP 2024**), Miami, Florida, ACL.
16. Xianjun Yang, Liangming Pan, Xuandong Zhao, Haifeng Chen, Linda Ruth Petzold, William Yang Wang, Wei Cheng, "A Survey on Detection of LLMs-Generated Content", to appear in Findings of The 2024 Conference on Empirical Methods in Natural Language Processing (**Findings of EMNLP 2024**), Miami, Florida, ACL.
17. Xiaobao Wu, Liangming Pan, William Yang Wang, Anh Tuan Luu, "AKEW: Assessing Knowledge Editing in the Wild", to appear in Proceedings of The 2024 Conference on Empirical Methods in Natural Language Processing (**EMNLP 2024**), Miami, Florida, ACL.
18. Alfonso Amayuelas, Xianjun Yang, Antonis Antoniadis, Wenyue Hua, Liangming Pan, William Yang Wang, "MultiAgent Collaboration Attack: Investigating Adversarial Attacks in Large Language Model Collaborations via Debate", to appear in Findings of The 2024 Conference on Empirical Methods in Natural Language Processing (**Findings of EMNLP 2024**), Miami, Florida, ACL.
19. Wenda Xu, Jiachen Li, William Yang Wang, Lei Li, "BPO: Staying Close to the Behavior LLM Creates Better Online LLM Alignment", to appear in Proceedings of The 2024 Conference on Empirical Methods in Natural Language Processing (**EMNLP 2024**), Miami, Florida, ACL.
20. Aditya Sharma, Michael Saxon, William Yang Wang, "Losing Visual Needles in Image Haystacks: Vision Language Models are Easily Distracted in Short and Long Contexts", to appear in Findings of The 2024 Conference on Empirical Methods in Natural Language Processing (**Findings of EMNLP 2024**), Miami, Florida, ACL.
21. Rujun Han, Yuhao Zhang, Peng Qi, Yumo Xu, Jenyuan Wang, Lan Liu, William Yang Wang, Bonan Min, Vittorio Castelli, "RAG-QA Arena: Evaluating Domain Robustness for Long-form Retrieval Augmented Question Answering", to appear in Proceedings of The 2024 Conference on Empirical Methods in Natural Language Processing (**EMNLP 2024**), Miami, Florida, ACL.
22. Jiachen Li, Weixi Feng, Wenhui Chen, William Yang Wang, "Reward Guided Latent Consistency Distillation", to appear in Transactions on Machine Learning Research (**TMLR**), 2024.

23. Zhiyu Chen, Jing Ma, Xinlu Zhang, Nan Hao, An Yan, Armineh Nourbakhsh, Xianjun Yang, Julian McAuley, Linda Ruth Petzold, William Yang Wang, "A Survey on Large Language Models for Critical Societal Domains: Finance, Healthcare, and Law", to appear in *Transactions on Machine Learning Research (TMLR)*, 2024. Survey Certification.
24. Alon Albalak, Yanai Elazar, Sang Michael Xie, Shayne Longpre, Nathan Lambert, Xinyi Wang, Niklas Muennighoff, Bairu Hou, Liangming Pan, Haewon Jeong, Colin Raffel, Shiyu Chang, Tatsunori Hashimoto, William Yang Wang, "A Survey on Data Selection for Language Models", to appear in *Transactions on Machine Learning Research (TMLR)*, 2024.
25. Xuehai He, Weixi Feng, Tsu-Jui Fu, Varun Jampani, Arjun Reddy Akula, Pradyumna Narayana, S Basu, William Yang Wang, Xin Eric Wang, "Discffusion: Discriminative Diffusion Models as Few-shot Vision and Language Learners", to appear in *Transactions on Machine Learning Research (TMLR)*, 2024.
26. Michael Saxon, Ari Holtzman, Peter West, William Yang Wang, Naomi Saphra, "Benchmarks as Microscopes: A Call for Model Metrology", to appear in *Proceedings of the First Conference on Language Modeling (COLM 2024)*, Philadelphia, PA, 2024.
27. Xinyi Wang, Lucas Caccia, Oleksiy Ostapenko, Xingdi Yuan, William Yang Wang, Alessandro Sordani, "Guiding Language Model Reasoning with Planning Tokens", to appear in *Proceedings of the First Conference on Language Modeling (COLM 2024)*, Philadelphia, PA, 2024.
28. Xingyu Fu, Muyu He, Yujie Lu, William Yang Wang, Dan Roth, "Commonsense-T2I Challenge: Can Text-to-Image Generation Models Understand Commonsense?", to appear in *Proceedings of the First Conference on Language Modeling (COLM 2024)*, Philadelphia, PA, 2024.
29. Haotian Zhang, Haoxuan You, Philipp Dufter, Bowen Zhang, Chen Chen, Hong-You Chen, Tsu-Jui Fu, William Yang Wang, Shih-Fu Chang, Zhe Gan, Yinfei Yang, "Ferret-v2: An Improved Baseline for Referring and Grounding with Large Language Models", to appear in *Proceedings of the First Conference on Language Modeling (COLM 2024)*, Philadelphia, PA, 2024.
30. Wenda Xu, Guanglei Zhu, Xuandong Zhao, Liangming Pan, Lei Li, William Yang Wang, "Perils of Self-Feedback: Self-Bias Amplifies in Large Language Models", to appear in *Proceedings of The 62nd Annual Meeting of the Association for Computational Linguistics (ACL 2024)*, Bangkok, Thailand, 2024.
31. Shuo Zhang, Liangming Pan, Junzhou Zhao, William Yang Wang, "The Knowledge Alignment Problem: Bridging Human and External Knowledge for Large Language Models", to appear in *Findings of The 62nd Annual Meeting of the Association for Computational Linguistics (Findings of ACL 2024)*, Bangkok, Thailand, 2024.
32. Kexun Zhang, Yee Man Choi, Zhenqiao Song, Taiqi He, William Yang Wang, Lei Li, "Hire a Linguist!: Learning Endangered Languages with In-Context Linguistic Descriptions", to appear in *Findings of The 62nd Annual Meeting of the Association for Computational Linguistics (Findings of ACL 2024)*, Bangkok, Thailand, 2024.
33. Alfonso Amayuelas, Kyle Wong, Liangming Pan, Wenhu Chen, William Yang Wang, "Knowledge of Knowledge: Exploring Known-Unknowns Uncertainty with Large Language Models", to appear in *Findings of The 62nd Annual Meeting of the Association for Computational Linguistics (Findings of ACL 2024)*, Bangkok, Thailand, 2024.
34. Danqing Wang, Antonis Antoniadis, Kha-Dinh Luong, Edwin Zhang, Mert Kosan, Jiachen Li, Ambuj Singh, William Yang Wang, Lei Li, "Global Human-guided Counterfactual Explanations for Molecular Properties via Reinforcement Learning", to appear in *Proceedings of 30th SIGKDD Conference on Knowledge Discovery and Data Mining (KDD 2024)*, Barcelona, Spain, 2024.

35. Xinyi Wang, Alfonso Amayuelas, Kexun Zhang, Liangming Pan, Wenhua Chen, William Yang Wang, "Understanding Reasoning Ability of Language Models From the Perspective of Reasoning Paths Aggregation", to appear in Proceedings of 41st International Conference on Machine Learning (**ICML 2024**), Vienna, Austria, July, 2024.
36. Jiachen Li, Qiaozhi Gao, Michael Johnston, Xiaofeng Gao, Xuehai He, Suhaila Shakiah, Hangjie Shi, Reza Ghanadan, William Yang Wang, "Mastering Robot Manipulation with Multimodal Prompts through Pretraining and Multi-task Fine-tuning", to appear in Proceedings of 41st International Conference on Machine Learning (**ICML 2024**), Vienna, Austria, July, 2024.
37. Iain Xie Weissburg, Mehir Arora, Xinyi Wang, Liangming Pan, William Yang Wang, "Position Paper: Understanding the Role of Social Media Influencers in AI Research Visibility", to appear in Proceedings of 41st International Conference on Machine Learning (**ICML 2024**), Vienna, Austria, July, 2024.
38. Huang et al., "Position Paper: TrustLLM: Trustworthiness in Large Language Models", to appear in Proceedings of 41st International Conference on Machine Learning (**ICML 2024**), Vienna, Austria, July, 2024.
39. Michael Saxon, Yiran Lawrence Luo, Sharon Levy, Chitta Baral, Yezhou Yang, William Yang Wang, "Lost in Translation? Translation Errors and Challenges for Fair Assessment of Text-to-Image Models on Multilingual Concepts", to appear in Proceedings of 2024 Annual Conference of the North American Chapter of the Association for Computational Linguistics (**NAACL 2024**), Mexico City, Mexico, ACL, 2024.
40. Wenda Xu, Daniel Deutsch, Mara Finkelstein, Juraj Juraska, Biao Zhang, Zhongtao Liu, William Yang Wang, Lei Li, Markus Freitag, "LLMMend: Pinpointing and Refining Large Language Models via Fine-Grained Actionable Feedback", to appear in Findings of 2024 Annual Conference of the North American Chapter of the Association for Computational Linguistics (**Findings of NAACL 2024**), Mexico City, Mexico, ACL, 2024.
41. Liangming Pan, Michael Saxon, Wenda Xu, Deepak Nathani, Xinyi Wang, William Yang Wang, "Automatically Correcting Large Language Models: Surveying the landscape of diverse self-correction strategies", to appear in Transactions of the Association for Computational Linguistics (**TACL 2024**), journal, 2024.
42. Antonis Antoniadis, Yiyi Yu, Joe S Canzano, William Yang Wang, Spencer Smith, "Neuroformer: Multimodal and Multitask Generative Pretraining for Brain Data", to appear in Proceedings of the Twelfth International Conference on Learning Representations (**ICLR 2024**), Vienna, Austria, 2024.
43. Xianjun Yang, Wei Cheng, Yue Wu, Linda Ruth Petzold, William Yang Wang, Haifeng Chen, "DNA-GPT: Divergent N-Gram Analysis for Training-Free Detection of GPT-Generated Text", to appear in Proceedings of the Twelfth International Conference on Learning Representations (**ICLR 2024**), Vienna, Austria, 2024.
44. Edwin Zhang, Yujie Lu, Shinda Huang, William Yang Wang, Amy Zhang, "Language Control Diffusion: Efficiently Scaling through Space, Time, and Tasks", to appear in Proceedings of the Twelfth International Conference on Learning Representations (**ICLR 2024**), Vienna, Austria, 2024.
45. Tsu-Jui Fu, Wenze Hu, Xianzhi Du, William Yang Wang, Yinfei Yang, Zhe Gan, "Guiding Instruction-based Image Editing via Multimodal Large Language Models", to appear in Proceedings of the Twelfth International Conference on Learning Representations (**ICLR 2024**), Spotlight paper, Vienna, Austria, 2024.

46. Raphael Schumann, Wanrong Zhu, Weixi Feng, Tsu-Jui Fu, Stefan Riezler, William Yang Wang, "VELMA: Verbalization Embodiment of LLM Agents for Vision and Language Navigation in Street View", to appear in Proceedings of the 38th Annual AAAI Conference on Artificial Intelligence (AAAI 2024), Vancouver, Canada, 2024.
47. Sudhanshu Srivastava, William Yang Wang, Miguel Eckstein, "Emergent human-like covert attention in feedforward convolutional neural networks", to appear in Current Biology, Volume 34, Issue 3, 5 February 2024, Pages 579-593.e12.
48. Wanrong Zhu, Xinyi Wang, Yujie Lu, Tsu-Jui Fu, Xin Eric Wang, Miguel Eckstein, William Yang Wang, "Collaborative Generative AI: Integrating GPT-k for Efficient Editing in Text-to-Image Generation", to appear in Proceedings of The 2023 Conference on Empirical Methods in Natural Language Processing (EMNLP 2023), Singapore, 2023.
49. Siqi Liu, Weixi Feng, Tsu-Jui Fu, Wenhui Chen, William Yang Wang, "EDIS: Entity-Driven Image Search over Multimodal Web Content", to appear in Proceedings of The 2023 Conference on Empirical Methods in Natural Language Processing (EMNLP 2023), Singapore, 2023.
50. Tsu-Jui Fu, Wenhan Xiong, Yixin Nie, Jingyu Liu, Barlas Oguz, William Yang Wang, "Text-guided 3D Human Generation from 2D Collections", to appear in Findings of The 2023 Conference on Empirical Methods in Natural Language Processing (EMNLP 2023 Findings), Singapore, 2023.
51. Liangming Pan, Alon Albalak, Xinyi Wang, William Yang Wang, "Logic-LM: Empowering Large Language Models with Symbolic Solvers for Faithful Logical Reasoning", to appear in Findings of The 2023 Conference on Empirical Methods in Natural Language Processing (EMNLP 2023 Findings), Singapore, 2023.
52. Hui Liu, Qingyu Yin, Zhengyang Wang, Chenwei Zhang, Haoming Jiang, Yifan Gao, Zheng Li, Xian Li, Chao Zhang, Bing Yin, William Yang Wang, Xiaodan Zhu, "Knowledge-Selective Pretraining for Attribute Value Extraction", to appear in Findings of The 2023 Conference on Empirical Methods in Natural Language Processing (EMNLP 2023 Findings), Singapore, 2023.
53. Alex Mei, Sharon Levy, William Yang Wang, "ASSERT: Automated Safety Scenario Red Teaming for Evaluating the Robustness of Large Language Models", to appear in Findings of The 2023 Conference on Empirical Methods in Natural Language Processing (EMNLP 2023 Findings), Singapore, 2023.
54. Yikang Pan, Liangming Pan, Wenhui Chen, Preslav Nakov, Min-Yen Kan, William Yang Wang, "On the Risk of Misinformation Pollution with Large Language Models", to appear in Findings of The 2023 Conference on Empirical Methods in Natural Language Processing (EMNLP 2023 Findings), Singapore, 2023.
55. Deepak Nathani, David Wang, Liangming Pan, William Yang Wang, "MAF: Multi-Aspect Feedback for Improving Reasoning in Large Language Models", to appear in Proceedings of The 2023 Conference on Empirical Methods in Natural Language Processing (EMNLP 2023), Singapore, 2023.
56. Zhiyu Chen, Yujie Lu, William Yang Wang, "Empowering Psychotherapy with Large Language Model: Cognitive Distortion Detection through Diagnosis of Thought Prompting", to appear in Findings of The 2023 Conference on Empirical Methods in Natural Language Processing (EMNLP 2023 Findings), Singapore, 2023.
57. Wenda Xu, Danqing Wang, Liangming Pan, Zhenqiao Song, Markus Freitag, William Yang Wang, Lei Li, "INSTRUCTSCORE: Towards Explainable Text Generation Evaluation with Automatic Feedback", to appear in Proceedings of The 2023 Conference on Empirical Methods in Natural Language Processing (EMNLP 2023), Singapore, 2023.

58. Kaiser Sun, Peng Qi, Yuhao Zhang, Lan Liu, William Yang Wang, ziheng huang, "Tokenization Consistency Matters for Generative Models on Extractive NLP Tasks", to appear in Findings of The 2023 Conference on Empirical Methods in Natural Language Processing (**EMNLP 2023 Findings**), Singapore, 2023.
59. Vaishnavi Himakunthala, Andy Ouyang, Daniel Philip Rose, Ryan He, Alex Mei, Yujie Lu, Chinmay Sonar, Michael Saxon, William Yang Wang, "Let's Think Frame by Frame with VIP: A Video Infilling and Prediction Dataset for Evaluating Video Chain-of-Thought", to appear in Proceedings of The 2023 Conference on Empirical Methods in Natural Language Processing (**EMNLP 2023**), Singapore, 2023.
60. Weixi Feng, Wanrong Zhu, Tsu-Jui Fu, Varun Jampani, Arjun Reddy Akula, Xuehai He, S Basu, Xin Eric Wang, William Yang Wang, "LayoutGPT: Compositional Visual Planning and Generation with Large Language Models", to appear in Proceedings of Thirty-seventh Conference on Neural Information Processing Systems (**NeurIPS 2023**), New Orleans, LA, 2023.
61. Alon Albalak, Colin Raffel, William Yang Wang, "Improving Few-Shot Generalization by Exploring and Exploiting Auxiliary Data", to appear in Proceedings of Thirty-seventh Conference on Neural Information Processing Systems (**NeurIPS 2023**), New Orleans, LA, 2023.
62. Xinyi Wang, Wanrong Zhu, Michael Saxon, Mark Steyvers, William Yang Wang, "Large Language Models Are Implicitly Topic Models: Explaining and Finding Good Demonstrations for In-Context Learning", to appear in Proceedings of Thirty-seventh Conference on Neural Information Processing Systems (**NeurIPS 2023**), New Orleans, LA, 2023.
63. Yujie Lu, Xianjun Yang, Xiujun Li, Xin Eric Wang, William Yang Wang, "LLMScore: Unveiling the Power of Large Language Models in Text-to-Image Synthesis Evaluation", to appear in Proceedings of Thirty-seventh Conference on Neural Information Processing Systems (**NeurIPS 2023**), New Orleans, LA, 2023.
64. Kexun Zhang, Danqing Wang, Jingtao Xia, William Yang Wang, Lei Li, "ALGO: Synthesizing Algorithmic Programs with Generated Oracle Verifiers", to appear in Proceedings of Thirty-seventh Conference on Neural Information Processing Systems (**NeurIPS 2023**), New Orleans, LA, 2023.
65. Zih-Yun Chiu, Yi-Lin Tuan, William Yang Wang, Michael C. Yip, "Flexible Attention-Based Multi-Policy Fusion for Efficient Deep Reinforcement Learning", to appear in Proceedings of Thirty-seventh Conference on Neural Information Processing Systems (**NeurIPS 2023**), New Orleans, LA, 2023.
66. Wanrong Zhu, Jack Hessel, Anas Awadalla, Samir Yitzhak Gadre, Jesse Dodge, Alex Fang, Youngjae Yu, Ludwig Schmidt, William Yang Wang, Yejin Choi, "Multimodal C4: An Open, Billion-scale Corpus of Images Interleaved with Text", to appear in Proceedings of Thirty-seventh Conference on Neural Information Processing Systems Datasets and Benchmarks Track (**NeurIPS 2023**), New Orleans, LA, 2023.
67. Liangming Pan, Wenhua Chen, Min-Yen Kan and William Yang Wang, "Attacking Open-domain Question Answering by Injecting Misinformation", to appear in Proceedings of the 13th International Joint Conference on Natural Language Processing and the 3rd Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics (**IJCNLP-AACL 2023**), Bali, Indonesia, 2023, **ACL Area Chair's Award**
68. Wenda Xu, Xian Qian, Mingxuan Wang, Lei Li, and William Yang Wang, "Retrieval Augmented Pretraining for Text Generation Evaluation", to appear in Proceedings of The 61st Annual Meeting of the Association for Computational Linguistics (**ACL 2023**), Toronto, Canada, 2023, **ACL**.

69. Michael Saxon and William Yang Wang, "Multilingual Conceptual Coverage in Text-to-Image Models", to appear in Proceedings of The 61st Annual Meeting of the Association for Computational Linguistics (**ACL 2023**), Toronto, Canada, 2023, ACL.
70. Liangming Pan, Xiaobao Wu, Xinyuan Lu, Anh Tuan Luu, William Yang Wang, Min-Yen Kan, and Preslav Nakov, "Fact-Checking Complex Claims with Program-Guided Reasoning", to appear in Proceedings of The 61st Annual Meeting of the Association for Computational Linguistics (**ACL 2023**), Toronto, Canada, 2023, ACL.
71. Yi-Lin Tuan, Alon Albalak, Wenda Xu, Michael Saxon, Connor F. Pryor, Lise Getoor, and William Yang Wang, "CausalDialogue: Modeling Utterance-level Causality in Conversations", to appear in Findings of The 61st Annual Meeting of the Association for Computational Linguistics (**ACL 2023 Findings**), Toronto, Canada, 2023, ACL.
72. Alex Mei, Sharon Levy, and William Yang Wang, "Foveate, Attribute, and Rationalize: Towards Safe and Trustworthy AI", to appear in Findings of The 61st Annual Meeting of the Association for Computational Linguistics (**ACL 2023 Findings**), Toronto, Canada, 2023, ACL.
73. Rujun Han, Peng Qi, Yuhao Zhang, Lan Liu, Juliette Burger, William Yang Wang, zhiheng huang, Bing Xiang and Dan Roth, "RobustQA: Benchmarking the Robustness of Domain Adaptation for Open-Domain Question Answering", to appear in Findings of The 61st Annual Meeting of the Association for Computational Linguistics (**ACL 2023 Findings**), Toronto, Canada, 2023, ACL.
74. Manoj Ghuhana Arivazhagan, Lan Liu, Peng Qi, Xinchu Chen, William Yang Wang and zhiheng huang, "Hybrid Hierarchical Retrieval for Open-Domain Question Answering", to appear in Findings of The 61st Annual Meeting of the Association for Computational Linguistics (**ACL 2023 Findings**), Toronto, Canada, 2023, ACL.
75. Xiyang Hu, Xinchu Chen, Peng Qi, Deguang Kong, Kunlun Liu, William Yang Wang and zhiheng huang, "Language Agnostic Multilingual Information Retrieval with Contrastive Learning", to appear in Findings of The 61st Annual Meeting of the Association for Computational Linguistics (**ACL 2023 Findings**), Toronto, Canada, 2023, ACL.
76. Xingyu Fu, Sheng Zhang, Gukyeong Kwon, Pramuditha Perera, Henghui Zhu, Yuhao Zhang, Alexander Hanbo Li, William Yang Wang, Zhiguo Wang, Vittorio Castelli, Patrick Ng, Dan Roth and Bing Xiang, "Generate then Select: Open-ended Visual Question Answering Guided by World Knowledge", to appear in Findings of The 61st Annual Meeting of the Association for Computational Linguistics (**ACL 2023 Findings**), Toronto, Canada, 2023, ACL.
77. Jifan Chen, Yuhao Zhang, Lan Liu, Rui Dong, Xinchu Chen, Patrick Ng, William Yang Wang and zhiheng huang, "Improving Cross-task Generalization of Unified Table-to-text Models with Compositional Task Configurations", to appear in Findings of The 61st Annual Meeting of the Association for Computational Linguistics (**ACL 2023 Findings**), Toronto, Canada, 2023, ACL.
78. Sijia Wang, Alexander Hanbo Li, Henghui Zhu, Sheng Zhang, Pramuditha Perera, Chung-Wei Hang, Jie Ma, William Yang Wang, Zhiguo Wang, Vittorio Castelli, Bing Xiang and Patrick Ng, "Benchmarking Diverse-Modal Entity Linking with Generative Models", to appear in Proceedings of The 61st Annual Meeting of the Association for Computational Linguistics (**ACL 2023**), Toronto, Canada, 2023, ACL.
79. Alexander Hanbo Li, Mingyue Shang, Evangelia Spiliopoulou, Jie Ma, Patrick Ng, Zhiguo Wang, Bonan Min, William Yang Wang, Kathleen McKeown, Vittorio Castelli, Dan Roth and Bing Xiang, "Few-Shot Data-to-Text Generation via Unified Representation and Multi-Source Learning", to appear in Proceedings of The 61st Annual Meeting of the Association for Computational Linguistics (**ACL 2023**), Toronto, Canada, 2023, ACL.

80. Kexun Zhang, Xianjun Yang, William Yang Wang, Lei Li, "ReDi: Efficient Learning-Free Diffusion Inference via Trajectory Retrieval", to appear in Proceedings of Fortieth International Conference on Machine Learning (ICML 2023), Honolulu, Hawaii, July 23 - July 29, ICML.
81. Jiachen Li, Edwin Zhang, Ming Yin, Qinxun Bai, Yu-Xiang Wang, William Yang Wang, "Offline Reinforcement Learning with Closed-Form Policy Improvement Operators", to appear in Proceedings of Fortieth International Conference on Machine Learning (ICML 2023), Honolulu, Hawaii, July 23 - July 29, ICML.
82. An Yan, Yu Wang, Yiwu Zhong, Chengyu Dong, Zexue He, Yujie Lu, William Yang Wang, Jingbo Shang, Julian McAuley, "Learning Concise and Descriptive Attributes for Visual Recognition", to appear in International Conference on Computer Vision (ICCV 2023), Paris, France, Oct 2-6, 2023, CVF.
83. Connor Pryor, Charles Dickens, Eriq Augustine, Alon Albalak, William Yang Wang, Lise Getoor, "NeuPSL: Neural Probabilistic Soft Logic", to appear in Proceedings of The 32nd International Joint Conference on Artificial Intelligence (IJCAI 2023), Macau, China, Aug 19-25, IJCAI.
84. Matthew Ho, Aditya Sharma, Justin Chang, Michael Saxon, Sharon Levy, Yujie Lu, William Yang Wang, "WikiWhy: Answering and Explaining Cause-and-Effect Questions", to appear in Proceedings of Eleventh International Conference on Learning Representations (ICLR 2023), Oral Paper: **Top 5% out of all 4019 submissions**, Kigali, Rwanda, May 1st to 5th.
85. Yujie Lu, Weixi Feng, Wanrong Zhu, Wenda Xu, Xin Eric Wang, Miguel Eckstein, William Yang Wang, "Neuro-Symbolic Procedural Planning with Commonsense Prompting", to appear in Proceedings of Eleventh International Conference on Learning Representations (ICLR 2023), Spotlight Paper, Kigali, Rwanda, May 1st to 5th.
86. Weixi Feng, Xuehai He, Tsu-Jui Fu, Varun Jampani, Arjun Reddy Akula, Pradyumna Narayana, Sugato Basu, Xin Eric Wang, William Yang Wang, "Training-Free Structured Diffusion Guidance for Compositional Text-to-Image Synthesis", to appear in Proceedings of Eleventh International Conference on Learning Representations (ICLR 2023), Kigali, Rwanda, May 1st to 5th.
87. Xinyi Wang, Michael Saxon, Jiachen Li, Hongyang Zhang, Kun Zhang, William Yang Wang, "Causal Balancing for Domain Generalization", to appear in Proceedings of Eleventh International Conference on Learning Representations (ICLR 2023), Kigali, Rwanda, May 1st to 5th.
88. Shuaichen Chang, Jun Wang, Mingwen Dong, Lin Pan, Henghui Zhu, Alexander Hanbo Li, Wuwei Lan, Sheng Zhang, Jiarong Jiang, Joseph Lilien, Steve Ash, William Yang Wang, Zhiguo Wang, Vittorio Castelli, Patrick Ng, Bing Xiang, "Dr.Spider: A Diagnostic Evaluation Benchmark towards Text-to-SQL Robustness", to appear in Proceedings of Eleventh International Conference on Learning Representations (ICLR 2023), Oral Paper: **Top 5% out of all 4019 submissions**, Kigali, Rwanda, May 1st to 5th.
89. Danilo Neves Ribeiro, Shen Wang, Xiaofei Ma, Henghui Zhu, Rui Dong, Deguang Kong, Juliette Burger, Anjelica Ramos, zhiheng huang, William Yang Wang, George Karypis, Bing Xiang, Dan Roth, "Street: a Multi-Task Structured Reasoning and Explanation benchmark", to appear in Proceedings of Eleventh International Conference on Learning Representations (ICLR 2023), Spotlight Paper, Kigali, Rwanda, May 1st to 5th.
90. Donghan Yu, Sheng Zhang, Patrick Ng, Henghui Zhu, Alexander Hanbo Li, Jun Wang, Yiqun Hu, William Yang Wang, Zhiguo Wang, Bing Xiang, "DecAF: Joint Decoding of Answers and Logical Forms for Question Answering over Knowledge Bases", to appear in Proceedings of Eleventh International Conference on Learning Representations (ICLR 2023), Kigali, Rwanda, May 1st to 5th.

91. Wanrong Zhu, An Yan, Yujie Lu, Wenda Xu, Xin Eric Wang, Miguel Eckstein and William Yang Wang, "Visualize Before You Write: Imagination-Guided Open-Ended Text Generation", in Proceedings of The 17th Conference of the European Chapter of the Association for Computational Linguistics (**EACL 2023**), long paper, Dubrovnik, Croatia, May 2-6 2023, ACL.
92. Wanrong Zhu, Xin Eric Wang, An Yan, Miguel Eckstein and William Yang Wang, "ImaginE: An Imagination-Based Automatic Evaluation Metric for Natural Language Generation", in Proceedings of The 17th Conference of the European Chapter of the Association for Computational Linguistics (**EACL 2023**), long paper, Dubrovnik, Croatia, May 2-6 2023, ACL.
93. Michael S. Saxon, Xinyi Wang, Wenda Xu and William Yang Wang, "PECO: Examining Single Sentence Label Leakage in Natural Language Inference Datasets through Progressive Evaluation of Cluster Outliers", in Proceedings of The 17th Conference of the European Chapter of the Association for Computational Linguistics (**EACL 2023**), long paper, Dubrovnik, Croatia, May 2-6 2023, ACL.
94. Kung-Hsiang Huang, Siffi Singh, Xiaofei Ma, Wei Xiao, Feng Nan, Nicholas Dingwall, William Yang Wang and Kathleen McKeown, "SWING: Mitigating Missing Information for Dialogue Summarization", in Proceedings of The 17th Conference of the European Chapter of the Association for Computational Linguistics (**EACL 2023**), long paper, Dubrovnik, Croatia, May 2-6 2023, ACL.
95. Tsu-Jui Fu, Licheng Yu, Ning Zhang, Cheng-Yang Fu, Jong-Chyi Su, William Yang Wang, Sean Bell, "Tell Me What Happened: Unifying Text-guided Video Completion via Multimodal Masked Video Generation", in Proceedings of The Thirty-Fourth IEEE/CVF Conference on Computer Vision and Pattern Recognition (**CVPR 2023**), long paper, Vancouver, Canada, June 18-22, 2023, IEEE/CVF.
96. Tsu-Jui Fu, Linjie Li, Zhe Gan, Kevin Lin, William Yang Wang, Lijuan Wang, Zicheng Liu, "An Empirical Study of End-to-End Video-Language Transformers with Masked Visual Modeling", in Proceedings of The Thirty-Fourth IEEE/CVF Conference on Computer Vision and Pattern Recognition (**CVPR 2023**), long paper, Vancouver, Canada, June 19-22, 2023, IEEE/CVF.
97. Alon Albalak, Sharon Levy and William Yang Wang, "Addressing Issues of Cross-Linguality in Open-Retrieval Question Answering Systems For Emergent Domains", in Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics (**EACL 2023**), demo track, Dubrovnik, Croatia, May 2-6 2023, ACL.
98. Daniel Xiaodan Zhou, Lan Liu, William Yang Wang, Zhiheng Huang, Anmol Anubhai, Maansi Shandilya and Steph Sigalas, "Beyond Accurate Answers: Evaluating Open-Domain Question Answering in Enterprise Search", to appear in Proceedings of the 2023 ACM SIGIR Conference on Human Information Interaction and Retrieval (**CHIIR 2023**), Short paper, Austin, Texas, USA March 19-23, ACM.
99. Alon Albalak, Yi-Lin Tuan, Pegah Jandaghi, Connor Pryor, Luke Yoffe, Deepak Ramachandran, Lise Getoor, Jay Pujara and William Yang Wang, "FETA: A Benchmark for Few-Sample Task Transfer in Open-Domain Dialogue", to appear in Proceedings of The 2022 Conference on Empirical Methods in Natural Language Processing (**EMNLP 2022**), full paper, Abu Dhabi, UAE, 2022.
100. Wenda Xu, Yi-Lin Tuan, Yujie Lu, Michael S. Saxon, Lei Li and William Yang Wang, "Not All Errors are Equal: Learning Text Generation Metrics using Stratified Error Synthesis", to appear in Findings of The 2022 Conference on Empirical Methods in Natural Language Processing (**Findings of EMNLP 2022**), full paper, Abu Dhabi, UAE, 2022.
101. Gyuwan Kim, Jinhyuk Lee, Barlas Oguz, Wenhan Xiong, Yizhe Zhang, Yashar Mehdad and William Yang Wang, "Bridging the Training-Inference Gap for Dense Phrase Retrieval", to appear in Findings of The 2022 Conference on Empirical Methods in Natural Language Processing (**Findings of EMNLP 2022**), full paper, Abu Dhabi, UAE, 2022.

102. Sharon Levy, Emily Allaway, Melanie Subbiah, Lydia Chilton, Desmond Patton, Kathleen McKeown and William Yang Wang, "SafeText: A Benchmark for Exploring Physical Safety in Language Models", to appear in Proceedings of The 2022 Conference on Empirical Methods in Natural Language Processing (**EMNLP 2022**), full paper, Abu Dhabi, UAE, 2022.
103. Weixi Feng, Tsu-Jui Fu, Yujie Lu and William Yang Wang, "ULN: Towards Underspecified Vision-and-Language Navigation", to appear in Proceedings of The 2022 Conference on Empirical Methods in Natural Language Processing (**EMNLP 2022**), full paper, Abu Dhabi, UAE, 2022.
104. Zhiyu Chen, Shiyang Li, Charese Smiley, Zhiqiang Ma, Sameena Shah and William Yang Wang, "ConvFinQA: Exploring the Chain of Numerical Reasoning in Conversational Finance Question Answering", to appear in Proceedings of The 2022 Conference on Empirical Methods in Natural Language Processing (**EMNLP 2022**), full paper, Abu Dhabi, UAE, 2022.
105. Alex Mei, Anisha Kabir, Sharon Levy, Melanie Subbiah, Emily Allaway, John N. Judge, Desmond Patton, Bruce Bimber, Kathleen McKeown and William Yang Wang, "Mitigating Covertly Unsafe Text within Natural Language Systems", to appear in Findings of The 2022 Conference on Empirical Methods in Natural Language Processing (**Findings of EMNLP 2022**), full paper, Abu Dhabi, UAE, 2022.
106. Xuehai He, Diji Yang, Weixi Feng, Tsu-Jui Fu, Arjun Akula, Varun Jampani, Pradyumna Narayana, Sugato Basu, William Yang Wang and Xin Eric Wang, "CPL: Counterfactual Prompt Learning for Vision and Language Models", to appear in Proceedings of The 2022 Conference on Empirical Methods in Natural Language Processing (**EMNLP 2022**), full paper, Abu Dhabi, UAE, 2022.
107. Tsu-Jui Fu, Xin Eric Wang, William Yang Wang, "Language-Driven Artistic Style Transfer", to appear in Proceedings of European Conference on Computer Vision 2022 (**ECCV 2022**), full paper, Tel Aviv, Israel, Oct 23-27, 2022.
108. Yujie Lu, Wanrong Zhu, Xin Eric Wang, Miguel Eckstein, William Yang Wang, "Imagination-Augmented Natural Language Understanding", to appear in Proceedings of 2022 Annual Conference of the North American Chapter of the Association for Computational Linguistics (**NAACL 2022**), long paper, Seattle, WA, July 10-15, 2022, USA.
109. Wanrong Zhu, Yuankai Qi, Pradyumna Narayana, Kazoo Sone, Sugato Basu, Xin Eric Wang, Qi Wu, Miguel P. Eckstein, William Yang Wang, "Diagnosing Vision-and-Language Navigation: What Really Matters", to appear in Proceedings of 2022 Annual Conference of the North American Chapter of the Association for Computational Linguistics (**NAACL 2022**), long paper, Seattle, WA, July 10-15, 2022, USA.
110. Zhiyu Chen, Bing Liu, Seungwhan Moon, Chinnadhurai Sankar, Paul A. Crook, William Yang Wang, "KETOD: Knowledge-Enriched Task-Oriented Dialogue", to appear in Findings of 2022 Annual Conference of the North American Chapter of the Association for Computational Linguistics (**Findings of NAACL 2022**), long paper, Seattle, WA, July 10-15, 2022, USA.
111. Tsu-Jui Fu, Xin Eric Wang, Scott Grafton, Miguel Eckstein, William Yang Wang, "M3L: Language-based Video Editing via Multi-Modal Multi-Level Transformers", to appear in Proceedings of Thirty-Third IEEE/CVF Conference on Computer Vision and Pattern Recognition (**CVPR 2022**), long paper, New Orleans, LA, USA.
112. Kai Nakamura, Sharon Levy, Yi-Lin Tuan, Wenhui Chen, William Yang Wang, "HybriDialogue: An Information-Seeking Dialogue Dataset Grounded on Tabular and Textual Data", to appear in Findings of 60th Annual Meeting of the Association for Computational Linguistics (**ACL 2022**), long paper, Dublin, Ireland.

113. Yi-Lin Tuan, Sajjad Beygi, Maryam Fazel-Zarandi, Qiaozi Gao, Alessandra Cervone, William Yang Wang, "Towards Large-Scale Interpretable Knowledge Graph Reasoning for Dialogue Systems", to appear in Findings of 60th Annual Meeting of the Association for Computational Linguistics (**ACL 2022**), long paper, Dublin, Ireland.
114. Wenda Xu, Michael Saxon, Misha Sra, William Yang Wang, "Self-Supervised Knowledge Assimilation for Expert-Layman Text Style Transfer", to appear in Proceedings of the Thirty-Sixth AAAI Conference on Artificial Intelligence (**AAAI 2022**), long paper, Vancouver, BC, Canada.
115. Tsu-Jui Fu, William Yang Wang, Daniel McDuff, Yale Song, "DOC2PPT: Automatic Presentation Slides Generation from Scientific Documents", to appear in Proceedings of the Thirty-Sixth AAAI Conference on Artificial Intelligence (**AAAI 2022**), long paper, Vancouver, BC, Canada.
116. Samhita Honnavalli, Aesha Parekh, Lily Ou, Sophie Groenwold, Sharon Levy, Vicente Ordonez and William Yang Wang, "Towards Understanding Gender-Seniority Compound Bias in Natural Language Generation", to appear in Proceedings of The 2022 International Conference on Language Resources and Evaluation (**LREC 2022**), Oral paper, Marseille, France, June 20-25,2022.
117. Alex Mei, Anisha Kabir, Rukmini Bapat, John N. Judge, Tony Sun and William Yang Wang, "Learning to Prioritize: Precision-Driven Sentence Filtering for Long Text Summarization", to appear in Proceedings of The 2022 International Conference on Language Resources and Evaluation (**LREC 2022**), Oral paper, Marseille, France, June 20-25,2022.
118. Xinyi Wang, Wenhua Chen, Michael Saxon, William Yang Wang, "Counterfactual Maximum Likelihood Estimation for Training Deep Networks", to appear in Proceedings of Thirty-fifth Conference on Neural Information Processing Systems (**NeurIPS 2021**), long paper, online.
119. Yi-Lin Tuan, Connor Pryor, Wenhua Chen, Lise Getoor, William Yang Wang, "Local Explanation of Dialogue Response Generation", to appear in Proceedings of Thirty-fifth Conference on Neural Information Processing Systems (**NeurIPS 2021**), long paper, online.
120. Wenhua Chen, Xinyi Wang, William Yang Wang, "A Dataset for Answering Time-Sensitive Questions", to appear in Proceedings of Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (**NeurIPS 2021**), long paper, online.
121. Linjie Li, Jie Lei, Zhe Gan, Licheng Yu, Yen-Chun Chen, Rohit Pillai, Yu Cheng, Luowei Zhou, Xin Eric Wang, William Yang Wang, Tamara Lee Berg, Mohit Bansal, Jingjing Liu, Lijuan Wang, Zicheng Liu, "VALUE: A Multi-Task Benchmark for Video-and-Language Understanding Evaluation", to appear in Proceedings of Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (**NeurIPS 2021**), long paper, online.
122. Zhiyu Chen, Wenhua Chen, Charese Smiley, Sameena Shah, Iana Borova, Dylan Langdon, Reema Moussa, Matt Beane, Ting-Hao Huang, Bryan R. Routledge, William Yang Wang, "FinQA: A Dataset of Numerical Reasoning over Financial Data", to appear in Proceedings of The 2021 Conference on Empirical Methods in Natural Language Processing (**EMNLP 2021**), long paper, online, ACL.
123. Michael Saxon, Sharon Levy, Xinyi Wang, Alon Albalak and William Yang Wang, "Modeling Disclosive Transparency in NLP Application Descriptions", to appear in Proceedings of The 2021 Conference on Empirical Methods in Natural Language Processing (**EMNLP 2021**), long paper, online, ACL.
124. Alexander Jones, William Yang Wang and Kyle Mahowald, "A Massively Multilingual Analysis of Cross-linguality in Shared Embedding Space", to appear in Proceedings of The 2021 Conference on Empirical Methods in Natural Language Processing (**EMNLP 2021**), long paper, online, ACL.

125. Sharon Levy, Kevin Mo, Wenhan Xiong and William Yang Wang, "Open-Domain Question-Answering for COVID-19 and Other Emergent Domains", to appear Proceedings of The 2021 Conference on Empirical Methods in Natural Language Processing (**EMNLP 2021**), demos track, online, ACL.
126. Liangming Pan, Wenhui Chen, Wenhan Xiong, Min-Yen Kan and William Yang Wang, "Zero-Shot Fact Verification by Claim Generation", to appear 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, online, ACL.
127. Qingfu Zhu, Wei-Nan Zhang, Ting Liu and William Yang Wang, "Neural Stylistic Response Generation with Disentangled Latent Variables", to appear 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**), long paper, online, ACL.
128. Sharon Levy, Michael Saxon and William Yang Wang, "Investigating Memorization of Conspiracy Theories in Text Generation", to appear in Findings 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 (**Findings of ACL-IJCNLP 2021**), long paper, online, ACL.
129. Liangming Pan, Wenhui Chen, Wenhan Xiong, Min-Yen Kan and William Yang Wang, "Unsupervised Multi-hop Question Answering by Question Generation", to appear in Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics - Human Language Technologies (**NAACL-HLT 2021**), long paper, ACL.
130. Tsu-Jui Fu and William Yang Wang, "Semi-Supervised Policy Initialization for Playing Games with Language Hints", to appear in Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics - Human Language Technologies (**NAACL-HLT 2021**), short paper, ACL.
131. Wenhui Chen, Ming-Wei Chang, Eva Schlinger, William Yang Wang, William Cohen, "Open Question Answering over Tables and Text", to appear in Proceedings of the International Conference on Learning Representations (**ICLR 2021**), full paper, ICLR.
132. Wenhan Xiong, Xiang Lorraine Li, Srini Iyer, Jingfei Du, Patrick Lewis, William Yang Wang, Yashar Mehdad, Wen-tau Yih, Sebastian Riedel, Douwe Kiela, Barlas Oguz, "Answering Complex Open-Domain Questions with Multi-Hop Dense Retrieval", to appear in Proceedings of the International Conference on Learning Representations (**ICLR 2021**), full paper, ICLR.
133. Wenhan Xiong, Hong Wang and William Yang Wang, "Progressively Pretrained Dense Corpus Index for Open-Domain Question Answering", to appear in Proceedings of the 2021 Conference of the European Chapter of the Association for Computational Linguistics (**EACL 2021**), long paper, ACL.
134. Yijun Xiao and William Yang Wang, "On Hallucination and Predictive Uncertainty in Conditional Language Generation", to appear in Proceedings of the 2021 Conference of the European Chapter of the Association for Computational Linguistics (**EACL 2021**), long paper, ACL.
135. An Yan, Xin Wang, Tsu-Jui Fu and William Yang Wang, "L2C: Describing Visual Differences Needs Semantic Understanding of Individuals", to appear in Proceedings of the 2021 Conference of the European Chapter of the Association for Computational Linguistics (**EACL 2021**), short paper, ACL.
136. Wanrong Zhu, Xin Wang, Tsu-Jui Fu, An Yan, Pradyumna Narayana, Kazuo Sone, Sugato Basu and William Yang Wang, "Multimodal Text Style Transfer for Outdoor Vision-and-Language Navigation", to appear in Proceedings of the 2021 Conference of the European Chapter of the Association for Computational Linguistics (**EACL 2021**), long paper, ACL.

137. Xiyu Zhou, Zhiyu Chen, Xiaoyong Jin and William Yang Wang, "HULK: An Energy Efficiency Benchmark Platform for Responsible Natural Language Processing", to appear in Proceedings of the 2021 Conference of the European Chapter of the Association for Computational Linguistics (EACL 2021), demo paper, ACL.
138. Wenhui Chen, Zhe Gan, Linjie Li, Yu Cheng, William Yang Wang, Jingjing Liu, "Meta Module Network for Compositional Visual Reasoning", to appear in Proceedings of Winter Conference on Applications of Computer Vision (WACV 2021), oral paper, IEEE/CVF. **Best Student Paper Award Honorable Mention**
139. Sophie Groenwold, Lily Ou, Aesha Parekh, Samhita Honnavalli, Sharon Levy, Diba Mirza and William Yang Wang, "Investigating African-American Vernacular English in Transformer-Based Text Generation", to appear in Proceedings of Conference on Empirical Methods in Natural Language Processing (EMNLP 2020), Short Paper, ACL.
140. Wenhui Chen, Yu Su, Xifeng Yan and William Yang Wang, "KGPT: Knowledge-Grounded Pre-Training for Data-to-Text Generation", to appear in Proceedings of Conference on Empirical Methods in Natural Language Processing (EMNLP 2020), Long Paper, ACL.
141. Wanrong Zhu, Xin Wang, Pradyumna Narayana, Kazuo Sone, Sugato Basu, William Yang Wang, "Towards Understanding Sample Variance in Visually Grounded Language Generation: Evaluations and Observations", to appear in Proceedings of Conference on Empirical Methods in Natural Language Processing (EMNLP 2020), Short Paper, ACL.
142. Tsu-Jui Fu, Xin Wang, Scott Grafton, Miguel Eckstein, William Yang Wang, "Iterative Language-Based Image Editing via Self-Supervised Counterfactual Reasoning", to appear in Proceedings of Conference on Empirical Methods in Natural Language Processing (EMNLP 2020), Long Paper, ACL.
143. Qingfu Zhu, Wei-Nan Zhang, Ting Liu and William Yang Wang, "Counterfactual Off-Policy Training for Neural Dialogue Generation", to appear in Proceedings of Conference on Empirical Methods in Natural Language Processing (EMNLP 2020), Long Paper, ACL.
144. Zhiyu Chen, Wenhui Chen, Hanwen Zha, Xiyu Zhou, Yunkai Zhang, Sairam Sundaresan, William Yang Wang, "Logic2Text: High-Fidelity Natural Language Generation from Logical Forms", to appear in Proceedings of Conference on Empirical Methods in Natural Language Processing (EMNLP 2020), Findings of EMNLP, ACL.
145. Wenhui Chen, Hanwen Zha, Zhiyu Chen, Wenhan Xiong, Hong Wang, William Wang, "HybridQA: A Dataset of Multi-Hop Question Answering over Tabular and Textual Data", to appear in Proceedings of Conference on Empirical Methods in Natural Language Processing (EMNLP 2020), Findings of EMNLP, ACL.
146. Jiannan Xiang, Xin Wang and William Yang Wang, "Learning to Stop: A Simple yet Effective Approach to Urban Vision-Language Navigation", to appear in Proceedings of Conference on Empirical Methods in Natural Language Processing (EMNLP 2020), Findings of EMNLP, ACL.
147. Tsu-Jui Fu, Xin Wang, Matthew F Peterson, Scott T. Grafton, Miguel Eckstein, William Yang Wang, "Counterfactual Vision-and-Language Navigation via Adversarial Path Sampler", to appear in Proceedings of the 16th European Conference on Computer Vision (ECCV 2020), **Spotlight Paper: Top 5% out of 5025 submissions**, Springer.
148. Xin Wang, Vihan Jain, Eugene Ie, William Yang Wang, Zornitsa Kozareva, Sujith Ravi, "Environment-agnostic Multitask Learning for Natural Language Grounded Navigation", to appear in Proceedings of the 16th European Conference on Computer Vision (ECCV 2020), full paper, Springer.

149. Sharon Levy, Wenhan Xiong, Elizabeth Belding, and William Yang Wang, "SafeRoute: Learning to Navigate Streets Safely in an Urban Environment", to appear in *ACM Transactions on Intelligent Systems and Technology (ACM TIST)*, journal paper, ACM, 2020.
150. Wenhua Chen, Jianshu Chen, Yu Su, Zhiyu Chen and William Yang Wang, "Logical Natural Language Generation from Open-Domain Tables", to appear in *Proceedings of The 58th Annual Meeting of the Association for Computational Linguistics (ACL 2020)*, full paper, ACL.
151. Chen Wu, Prince Zizhuang Wang and William Yang Wang, "On the Encoder-Decoder Incompatibility in Variational Text Modeling and Beyond", to appear in *Proceedings of The 58th Annual Meeting of the Association for Computational Linguistics (ACL 2020)*, full paper, ACL.
152. Andrew Gaut, Tony Sun, Shirlyn Tang, Yuxin Huang, Jing Qian, Mai ElSherief, Jieyu Zhao, Diba Mirza, Elizabeth Belding, Kai-Wei Chang and William Yang Wang, "Towards Understanding Gender Bias in Relation Extraction", to appear in *Proceedings of The 58th Annual Meeting of the Association for Computational Linguistics (ACL 2020)*, full paper, ACL.
153. Zhiyu Chen, Harini Eavani, Wenhua Chen, Yinyin Liu and William Yang Wang, "Few-Shot NLG with Pre-Trained Language Model", to appear in *Proceedings of The 58th Annual Meeting of the Association for Computational Linguistics (ACL 2020)*, short paper, ACL.
154. Xin Wang, Qiuyuan Huang, Asli Celikyilmaz, Jianfeng Gao, Dinghan Shen, Yuan-Fang Wang, William Yang Wang, Lei Zhang, "Vision-Language Navigation Policy Learning and Adaptation", accepted to *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, IEEE Computer Science Society. Journal Version of the CVPR 2019 Best Student Paper.
155. Wenhan Xiong, Jingfei Du, William Yang Wang, Veselin Stoyanov, "Pretrained Encyclopedia: Weakly Supervised Knowledge-Pretrained Language Model", to appear in *Proceedings of the International Conference on Learning Representations (ICLR 2020)*, full paper, Addis Ababa, Ethiopia, Apr. 30, ICLR.
156. Wenhua Chen, Hongmin Wang, Jianshu Chen, Yunkai Zhang, Hong Wang, Shiyang Li, Xiyu Zhou, William Yang Wang, "TabFact: A Large-scale Dataset for Table-based Fact Verification", to appear in *Proceedings of the International Conference on Learning Representations (ICLR 2020)*, full paper, Addis Ababa, Ethiopia, Apr. 30, ICLR.
157. Yuankai Qi, Qi Wu, Peter Anderson, Xin Wang, William Yang Wang, Chunhua Shen, Anton van den Hengel, "REVERIE: Remote Embodied Visual Referring Expression in Real Indoor Environments", to appear in *Proceedings of the 2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR 2020)*, full paper, Seattle, WA, June 14-19, IEEE/CVF.
158. Juncheng Li, Xin Wang, Siliang Tang, Haizhou Shi, Fei Wu, Yueting Zhuang, William Yang Wang, "Unsupervised Reinforcement Learning of Transferable Meta-Skills for Embodied Navigation", to appear in *Proceedings of the 2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR 2020)*, full paper, Seattle, WA, June 14-19, IEEE/CVF.
159. Pengda Qin, Xin Wang, Wenhua Chen, Chunyun Zhang, Weiran Xu, and William Yang Wang, "Generative Adversarial Zero-Shot Relational Learning for Knowledge Graphs", to appear in *Proceedings of the Thirty-Fourth AAAI Conference on Artificial Intelligence (AAAI 2020)*, full oral paper, New York, New York, AAAI.
160. Shaolei Wang, Wanxiang Che, Qi Liu, Pengda Qin, Ting Liu, and William Yang Wang, "Multi-Task Self-Supervised Learning for Disfluency Detection", to appear in *Proceedings of the Thirty-Fourth AAAI Conference on Artificial Intelligence (AAAI 2020)*, full oral paper, New York, New York, AAAI.

161. Kai Nakamura, Sharon Levy and William Yang Wang, "Fakeddit: A New Multimodal Benchmark Dataset for Fine-grained Fake News Detection", to appear in Proceedings of 12th International Conference on Language Resources and Evaluation (**LREC 2020**), full paper, Marseille, France, May 11-16, 2020, ELRA.

162. Ray Oshikawa, Jing Qian and William Yang Wang, "A Survey on Natural Language Processing for Fake News Detection", to appear in Proceedings of 12th International Conference on Language Resources and Evaluation (**LREC 2020**), full paper, Marseille, France, May 11-16, 2020, ELRA.

163. Prince Zizhuang Wang and William Yang Wang, "Neural Gaussian Copula for Variational Autoencoder", to appear in Proceedings of 2019 Conference on Empirical Methods in Natural Language Processing and 9th International Joint Conference on Natural Language Processing (**EMNLP-IJCNLP 2019**), full paper, Hong Kong, Nov. 3-7, ACL.
164. Jiawei Wu, Wenhan Xiong and William Yang Wang, "Learning to Learn and Predict: A Meta-Learning Approach for Multi-Label Classification", to appear in Proceedings of 2019 Conference on Empirical Methods in Natural Language Processing and 9th International Joint Conference on Natural Language Processing (**EMNLP-IJCNLP 2019**), full paper, Hong Kong, Nov. 3-7, ACL.
165. Jing Qian, Anna Bethke, Yinyin Liu, Elizabeth Belding and William Yang Wang, "A Benchmark Dataset for Learning to Intervene in Online Hate Speech", to appear in Proceedings of 2019 Conference on Empirical Methods in Natural Language Processing and 9th International Joint Conference on Natural Language Processing (**EMNLP-IJCNLP 2019**), full paper, Hong Kong, Nov. 3-7, ACL.
166. Siyao Li, Deren Lei, Pengda Qin and William Yang Wang, "Deep Reinforcement Learning with Distributional Semantic Rewards for Abstractive Summarization", to appear in Proceedings of 2019 Conference on Empirical Methods in Natural Language Processing and 9th International Joint Conference on Natural Language Processing (**EMNLP-IJCNLP 2019**), short paper, Hong Kong, Nov. 3-7, ACL.
167. Xin Wang, Jiawei Wu, Junkun Chen, Lei Li, Yuan-Fang Wang, and William Yang Wang, "VATEX: A Large-Scale, High-Quality Multilingual Dataset for Video-and-Language Research", to appear in Proceedings of the 17th CVF/IEEE International Conference on Computer Vision (**ICCV 2019**), full paper, Seoul, Korea, Oct. 27 - Nov. 2, CVF/IEEE.
168. Wenhu Chen, Jianshu Chen, Pengda Qin, Xifeng Yan and William Yang Wang, "Semantically Conditioned Dialog Response Generation via Hierarchical Disentangled Self-Attention", to appear in Proceedings of The 57th Annual Meeting of the Association for Computational Linguistics (**ACL 2019**), long paper, Florence, Italy, July 28th - August 2nd, ACL.
169. Jiawei Wu, Xin Wang and William Yang Wang, "Self-Supervised Dialogue Learning", to appear in Proceedings of The 57th Annual Meeting of the Association for Computational Linguistics (**ACL 2019**), long paper, Florence, Italy, July 28th - August 2nd, ACL.
170. Pushkar shukla, Carlos Elmadjian, Richika Sharan, Vivek Kulkarni, Matthew Turk and William Yang Wang, "What Should I Ask? Using Conversationally Informative Rewards for Goal-Oriented Visual Dialog", to appear in Proceedings of The 57th Annual Meeting of the Association for Computational Linguistics (**ACL 2019**), long paper, Florence, Italy, July 28th - August 2nd, ACL.
171. Wenhan Xiong, Mo Yu, Shiyu Chang, Xiaoxiao Guo and William Yang Wang, "Improving Question Answering over Incomplete KBs with Knowledge-Aware Reader", to appear in Proceedings of The 57th Annual Meeting of the Association for Computational Linguistics (**ACL 2019**), long paper, Florence, Italy, July 28th - August 2nd, ACL.

172. Wenhan Xiong, Jiawei Wu, Hong Wang, Vivek Kulkarni, Mo Yu, Xiaoxiao Guo, Shiyu Chang and William Yang Wang, "TWEETQA: A Social Media Focused Question Answering Dataset", to appear in Proceedings of The 57th Annual Meeting of the Association for Computational Linguistics (**ACL 2019**), long paper, Florence, Italy, July 28th - August 2nd, ACL.
173. Hong Wang, Xin Wang, Wenhan Xiong, Mo Yu, Xiaoxiao Guo, Shiyu Chang and William Yang Wang, "Self-Supervised Learning for Contextualized Extractive Summarization", to appear in Proceedings of The 57th Annual Meeting of the Association for Computational Linguistics (**ACL 2019**), short paper, Florence, Italy, July 28th - August 2nd, ACL.
174. Hui Liu, Qingyu Yin and William Yang Wang, "Towards Explainable NLP: A Generative Explanation Framework for Text Classification", to appear in Proceedings of The 57th Annual Meeting of the Association for Computational Linguistics (**ACL 2019**), full paper, Florence, Italy, July 28th - August 2nd, ACL.
175. Tony Sun, Andrew Gaut, Shirlyn Tang, Yuxin Huang, Mai ElSherief, Jieyu Zhao, Diba Mirza, Kai-Wei Chang and William Yang Wang, "Debiasing Gender in Natural Language Processing: Literature Review", to appear in Proceedings of The 57th Annual Meeting of the Association for Computational Linguistics (**ACL 2019**), full paper, Florence, Italy, July 28th - August 2nd, ACL.
176. Xin Wang, Qiuyuan Huang, Asli Celikyilmaz, Jianfeng Gao, Dinghan Shen, Yuan-Fang Wang, William Yang Wang, Lei Zhang, "Reinforced Cross-Modal Matching and Self-Supervised Imitation Learning for Vision-Language Navigation", to appear in Proceedings of the 31st IEEE/CVF Conference on Computer Vision and Pattern Recognition (**CVPR 2019**), full oral paper, Long Beach, CA, USA, June 16-20, CVF/IEEE. (**Best Student Paper Award, 1/5160=0.02%**)
177. Hong Wang, Wenhan Xiong, Mo Yu, Xiaoxiao Guo, Shiyu Chang, and William Yang Wang, "Sentence Embedding Alignment for Lifelong Relation Extraction", to appear in Proceedings of the 2019 Annual Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (**NAACL-HLT 2019**), long paper, Minneapolis, MN, USA, June 3-5, ACL.
178. Prince Zizhuang Wang and William Yang Wang, "Riemannian Normalizing Flow on Variational Wasserstein Autoencoder for Text Modeling", to appear in Proceedings of the 2019 Annual Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (**NAACL-HLT 2019**), long oral paper, Minneapolis, MN, USA, June 3-5, ACL.
179. Wenhan Xiong, Jiawei Wu, Deren Lei, Mo Yu, Shiyu Chang, Xiaoxiao Guo, and William Yang Wang, "Imposing Label-Relational Inductive Bias for Extremely Fine-Grained Entity Typing", to appear in Proceedings of the 2019 Annual Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (**NAACL-HLT 2019**), long paper, Minneapolis, MN, USA, June 3-5, ACL.
180. Jiawei Wu, Xin Wang and William Yang Wang, "Extract and Edit: An Alternative to Back-Translation for Unsupervised Neural Machine Translation", to appear in Proceedings of the 2019 Annual Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (**NAACL-HLT 2019**), long oral paper, Minneapolis, MN, USA, June 3-5, ACL.
181. Jing Qian, Mai ElSherief, Elizabeth Belding and William Yang Wang, "Learning to Decipher Hate Symbols", to appear in Proceedings of the 2019 Annual Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (**NAACL-HLT 2019**), long oral paper, Minneapolis, MN, USA, June 3-5, ACL.

182. Wenhua Chen, Yu Su, Yilin Shen, Zhiyu Chen, Xifeng Yan, and William Yang Wang, "How Large A Vocabulary Does Text Classification Need? A Variational Approach on Vocabulary Selection", to appear in Proceedings of the 2019 Annual Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (**NAACL-HLT 2019**), long oral paper, Minneapolis, MN, USA, June 3-5, ACL.
183. Xin Wang, Jiawei Wu, Da Zhang, Yu Su, William Yang Wang, "Learning to Compose Topic-Aware Mixture of Experts for Zero-Shot Video Captioning", to appear in Proceedings of the Thirty-Third AAAI Conference on Artificial Intelligence (**AAAI 2019**), long oral paper, Honolulu, HI, USA, Jan 27 - Feb 1, 2019, AAAI. (Oral paper acceptance rate: 8%)
184. Yijun Xiao and William Yang Wang, "Quantifying Uncertainties in Natural Language Processing Tasks", to appear in Proceedings of the Thirty-Third AAAI Conference on Artificial Intelligence (**AAAI 2019**), long paper, Honolulu, HI, USA, Jan 27 - Feb 1, 2019, AAAI.
185. Wenhan Xiong, Mo Yu, Shiyu Chang, Xiaoxiao Guo and William Yang Wang, "One-Shot Relational Learning for Knowledge Graphs", to appear in Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing (**EMNLP 2018**), long paper, Brussels, Belgium, Oct 31-Nov 4, ACL.
186. Wenhua Chen, Jianshu Chen, Yu Su, Xin Wang, Dong Yu, Xifeng Yan and William Yang Wang, "XL-NBT: A Cross-lingual Neural Belief Tracking Framework", to appear in Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing (**EMNLP 2018**), long paper, Brussels, Belgium, Oct 31-Nov 4, ACL.
187. Jing Qian, Mai ElSherief, Elizabeth Belding and William Yang Wang, "Hierarchical CVAE for Fine-Grained Hate Speech Classification", to appear in Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing (**EMNLP 2018**), long paper, Brussels, Belgium, Oct 31-Nov 4, ACL.
188. Vivek Kulkarni, Junting Ye, Steve Skiena and William Yang Wang, "Multi-view Models for Political Ideology Detection of News Articles", to appear in Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing (**EMNLP 2018**), long paper, Brussels, Belgium, Oct 31-Nov 4, ACL.
189. Qingyu Yin, Yu Zhang, Wei-Nan Zhang, Ting Liu, and William Yang Wang, "Zero Pronoun Resolution with Attention-based Neural Network", in Proceedings of the 27th International Conference on Computational Linguistics (**COLING 2018**), full paper, Santa Fe, New Mexico, USA, August 20-26, ACL.
190. Xin Wang, Wenhan Xiong, Hongmin Wang, William Yang Wang, "Look Before You Leap: Bridging Model-Free and Model-Based Reinforcement Learning for Planned-Ahead Vision-and-Language Navigation", in Proceedings of the European Conference on Computer Vision (**ECCV 2018**), full paper, Munich, Germany, Sept. 8-14, Springer.
191. Pengda Qin, Weiran Xu, and William Yang Wang, "DSGAN: Generative Adversarial Training for Robust Distant Supervision Relation Extraction", in Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (**ACL 2018**), full paper, Melbourne, Australia, July 15-20, 2018, ACL.
192. Qingyu Yin, Yu Zhang, Wei-Nan Zhang, Ting Liu, and William Yang Wang, "Deep Reinforcement Learning for Chinese Zero Pronoun Resolution", in Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (**ACL 2018**), full paper, Melbourne, Australia, July 15-20, 2018, ACL.

193. Xin Wang, Wenhua Chen, Yuan-Fang Wang, and William Yang Wang, "No Metrics Are Perfect: Adversarial Reward Learning for Visual Storytelling", in Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (**ACL 2018**), full paper, Melbourne, Australia, July 15-20, 2018, ACL.
194. Pengda Qin, Weiran Xu, and William Yang Wang, "Robust Distant Supervision via Deep Reinforcement Learning", in Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (**ACL 2018**), full paper, Melbourne, Australia, July 15-20, 2018, ACL.
195. Xianda Zhou and William Yang Wang, "MOJITALK: Generating Emotional Responses at Scale", in Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (**ACL 2018**), full paper, Melbourne, Australia, July 15-20, 2018, ACL.
196. Wenhan Xiong, Xiaoxiao Guo, Mo Yu, Shiyu Chang, Bowen Zhou, and William Yang Wang, "Scheduled Policy Optimization for Natural Language Communication with Intelligent Agents", in Proceedings of the 27th International Joint Conference on Artificial Intelligence and the 23rd European Conference on Artificial Intelligence (**IJCAI-ECAI 2018**), full oral paper, Stockholm, Sweden, July 13-19, 2018, IJCAI.
197. Mai Elshierief, Vivek Kulkarni, Dana Nguyen, William Yang Wang, and Elizabeth Belding, "Hate Lingo: A Target-based Linguistic Analysis of Hate Speech in Social Media", in Proceedings of the International AAAI Conference on Web and Social Media (**ICWSM 2018**), full paper, Stanford, CA, USA, June 25-28, 2018, AAAI.
198. Xin Wang, Wenhua Chen, Jiawei Wu, Yuan-Fang Wang, and **William Yang Wang**, "Video Captioning via Hierarchical Reinforcement Learning", in Proceedings of the Thirtieth IEEE/CVF Conference on Computer Vision and Pattern Recognition (**CVPR 2018**), full paper, Salt Lake City, Utah, USA, June 18 - June 22, 2018, IEEE/CVF.
199. Jing Qian, Mai ElSherief, Elizabeth Belding, and William Yang Wang, "Leveraging Intra-User and Inter-User Representation Learning for Automated Hate Speech Detection", in Proceedings of The 16th Annual Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (**NAACL HLT 2018**), short paper, New Orleans, LA, USA, June 1 - June 6, 2018, ACL.
200. Xin Wang, Yuan-Fang Wang, and **William Yang Wang**, "Watch, Listen, and Describe: Globally and Locally Aligned Cross-Modal Attentions for Video Captioning", in Proceedings of The 16th Annual Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (**NAACL HLT 2018**), short paper, New Orleans, LA, USA, June 1 - June 6, 2018, ACL.
201. Wenhua Chen, Wenhan Xiong, Xifeng Yan and **William Yang Wang**, "Variational Knowledge Graph Reasoning", in Proceedings of The 16th Annual Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (**NAACL HLT 2018**), long oral paper, New Orleans, LA, USA, June 1 - June 6, 2018, ACL.
202. Jiawei Wu, Lei Li, and **William Yang Wang**, "Reinforced Co-Training", in Proceedings of The 16th Annual Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (**NAACL HLT 2018**), long paper, New Orleans, LA, USA, June 1 - June 6, 2018, ACL.
203. Vivek Kulkarni and **William Yang Wang**, "Simple Models for Word Formation in Slang", in Proceedings of The 16th Annual Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (**NAACL HLT 2018**), long paper, New Orleans, LA, USA, June 1 - June 6, 2018, ACL.

204. Liwei Cai and **William Yang Wang**, "KBGAN: Adversarial Learning for Knowledge Graph Embeddings", in Proceedings of The 16th Annual Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (**NAACL HLT 2018**), long oral paper, New Orleans, LA, USA, June 1 - June 6, 2018, ACL.
205. Ke Ni, and **William Yang Wang**, "Learning to Explain Non-Standard English Words and Phrases", in Proceedings of the 8th International Joint Conference on Natural Language Processing (**IJCNLP 2017**), short paper, Taipei, Taiwan, Nov.27-Dec.1, AFNLP.
206. Wenhan Xiong, Thien Hoang, **William Yang Wang**, "DeepPath: A Reinforcement Learning Method for Knowledge Graph Reasoning", in Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing (**EMNLP 2017**), long oral paper, Copenhagen, Denmark, Sept. 7-11, ACL.
207. YiYao Huang, and **William Yang Wang**, "Deep Residual Learning for Weakly-Supervised Relation Extraction", in Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing (**EMNLP 2017**), short paper, Copenhagen, Denmark, Sept. 7-11, ACL.
208. **William Yang Wang**, "Liar, Liar Pants on Fire": A New Benchmark Dataset for Fake News Detection, in Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (**ACL 2017**), short paper, Vancouver, BC, Canada, July 30-August 4, ACL.
209. **William Yang Wang**, and William W. Cohen, "Learning First-Order Logic Embeddings via Matrix Factorization", in Proceedings of the 25th International Joint Conference on Artificial Intelligence (**IJCAI 2016**), full paper for oral presentation, New York, NY, U.S.A., July 9-15, IJCAI.
210. **William Yang Wang**, Yashar Mehdad, Dragomir Radev, and Amanda Stent, "A Low-Rank Approximation Approach to Learning Joint Embeddings of News Stories and Images for Timeline Summarization", in Proceedings of the 15th Annual Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (**NAACL HLT 2016**), long paper, San Diego, CA, U.S.A., June 13-15, ACL.
211. **William Yang Wang** and Diyi Yang, "That's So Annoying!!!: A Lexical and Frame-Semantic Embedding Based Data Augmentation Approach to Automatic Categorization of Annoying Behaviors using #petpeeve Tweets", in Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing (**EMNLP 2015**), short paper, Lisbon, Portugal, Sept 19-21, ACL. (**EMNLP 2015 Inaugural Notable Data Set Award**, $2/1300=0.2\%$)
212. **William Yang Wang**, Kathryn Mazaitis, and William W. Cohen, "A Soft Version of Predicate Invention Based on Structured Sparsity", in Proceedings of the 24th International Joint Conference on Artificial Intelligence (**IJCAI 2015**), full paper for oral presentation, Buenos Aires, Argentina, July 25-31, IJCAI.
213. **William Yang Wang** and William W. Cohen, "Joint Information Extraction and Reasoning: A Scalable Statistical Relational Learning Approach", in Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and The 7th International Joint Conference of the Asian Federation of Natural Language Processing (**ACL-IJCNLP 2015**), long paper for oral presentation, Beijing, China, July 26-31, ACL.
214. Yun-Nung Chen, **William Yang Wang**, Anatole Gershman and Alexander Rudnicky, "Matrix Factorization with Knowledge Graph Propagation for Unsupervised Spoken Language Understanding", in Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and The 7th International Joint Conference of the Asian Federation of Natural Language Processing (**ACL-IJCNLP 2015**), long paper for oral presentation, Beijing, China, July 26-31, ACL.

215. **William Yang Wang**, Kathryn Mazaitis, Ni Lao, William W. Cohen, "Efficient Inference and Learning in a Large Knowledge Base: Reasoning with Extracted Information using a Locally Groundable First-Order Probabilistic Logic", in *Machine Learning Journal (MLJ)*, 2015, Springer.
216. **William Yang Wang**, Miaomiao Wen, "I Can Has Cheezburger? A Nonparanormal Approach to Combining Textual and Visual Information for Predicting and Generating Popular Meme Descriptions", in *Proceedings of the 2015 Conference of the North American Chapter of the Association for Computational Linguistics - Human Language Technologies (NAACL HLT 2015)*, long paper, Denver, CO, May 31 - June 5, ACL. Covered by U.S. and Spanish media.
217. Yun-Nung Chen, **William Yang Wang**, and Alex Rudnicky, "Jointly Modeling Inter-Slot Relations by Random Walk on Knowledge Graphs for Unsupervised Spoken Language Understanding", in *Proceedings of the 2015 Conference of the North American Chapter of the Association for Computational Linguistics - Human Language Technologies (NAACL HLT 2015)*, long paper, Denver, CO, May 31 - June 5, ACL.
218. Yun-Nung Chen, **William Yang Wang**, and Alex Rudnicky, "Learning Semantic Hierarchy with Distributed Representations for Unsupervised Spoken Language Understanding", in *Proceedings of the 16th Annual Conference of the International Speech Communication Association (INTERSPEECH 2015)*, Dresden, Germany, Sept. 6-10, ISCA.
219. **William Yang Wang**, Kathryn Mazaitis, William W. Cohen, "Structure Learning via Parameter Learning", in *Proceedings of the 23rd ACM International Conference on Information and Knowledge Management (CIKM 2014)*, full oral paper, Shanghai, China, Nov. 3-7th, ACM.
220. **William Yang Wang**, Lingpeng Kong, Kathryn Mazaitis, William W. Cohen, "Dependency Parsing for Weibo: An Efficient Probabilistic Logic Programming Approach", in *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP 2014)*, short paper, Doha, Qatar, Oct. 25-29, ACL.
221. **William Yang Wang**, and Zhenhao Hua, "A Semiparametric Gaussian Copula Regression Model for Predicting Financial Risks from Earnings Calls", in *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (ACL 2014)*, long paper, Baltimore, MD, June 22-27, ACL.
222. **William Yang Wang**, Kathryn Mazaitis, William W. Cohen, "ProPPR: Efficient First-Order Probabilistic Logic Programming for Structure Discovery, Parameter Learning, and Scalable Inference", in *Proceedings of the AAAI 2014 International Workshop on Statistical Relational AI (StarAI 2014)*, position paper, Québec City, Québec, Canada, July 27-31, AAAI.
223. Yun-Nung Chen, **William Yang Wang**, and Alex Rudnicky, "Leveraging Frame Semantics and Distributional Semantics for Unsupervised Semantic Slot Induction for Spoken Dialogue Systems", in *Proceedings of IEEE SLT 2014*, Lake Tahoe, California and Nevada, Dec 7-10, IEEE.
224. **William Yang Wang**, Kathryn Mazaitis, William W. Cohen, "Programming with Personalized PageRank: A Locally Groundable First-Order Probabilistic Logic", in *Proceedings of the 22nd ACM International Conference on Information and Knowledge Management (CIKM 2013)*, full oral paper, Burlingame, CA, USA, Oct. 27th - Nov. 1st, ACM. (**Best Paper Honorable Mention Award**)
225. **William Yang Wang**, Edward Lin, John Kominek, "This Text has the Scent of Starbucks: A Laplacian Structured Sparsity Model for Computational Branding Analytics", in *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing (EMNLP 2013)*, full paper, Seattle, WA, USA, Oct. 18-21, ACL.
226. Di Wang, Chenyan Xiong, **William Yang Wang**, "Automatic Domain Partitioning for Multi-Domain Learning", in *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing (EMNLP 2013)*, short paper, Seattle, WA, USA, Oct. 18-21, ACL.

227. Yun-Nung Chen, **William Yang Wang**, and Alex Rudnicky, "Unsupervised Induction and Filling of Semantic Slots for Spoken Dialogue Systems using Frame-Semantic Parsing", in Proceedings of the IEEE ASRU 2013, Olomouc, Czech Republic, Dec. 8-12, 2013, IEEE. (Best Student Paper Award)
228. Yun-Nung Chen, **William Yang Wang**, and Alex Rudnicky, "An Empirical Investigation of Sparse Log-Linear Models for Improved Dialogue Act Classification", in Proceedings of 38th International Conference on Acoustics, Speech, and Signal Processing (ICASSP 2013), Vancouver, May 26-31, IEEE.
229. **William Yang Wang**, Dan Bohus, Ece Kamar, and Eric Horvitz, "Crowdsourcing the Acquisition of Natural Language Corpora: Methods and Observations", in Proceedings of IEEE SLT 2012, Miami, Florida, Dec 2-5, IEEE.
230. **William Yang Wang**, Elijah Mayfield, Suresh Naidu, and Jeremiah Dittmar, "Historical Analysis of Legal Opinions with a Sparse Mixed-Effects Latent Variable Model", in Proceedings of the 50th Annual Meeting of the Association for Computational Linguistics (ACL 2012), full oral paper, Jeju Island, Korea, July 8-14, ACL. (Oral Presentation Long Paper Acceptance Rate: 14%)
231. **William Yang Wang**, Samantha Finkelstein, Amy Ogan, Alan W Black and Justine Cassell, "'Love ya, jerkface': using Sparse Log-Linear Models to Build Positive (and Impolite) Relationships with Teens", in Proceedings of the 13th annual SIGdial Meeting on Discourse and Dialogue (SIGDIAL 2012), full oral paper, Seoul, Korea, July 5-6, ACL.
232. **William Yang Wang**, Fadi Biadisy, Andrew Rosenberg, Julia Hirschberg, "Automatic Detection of Speaker State: Lexical, Prosodic, and Phonetic Approaches to Level-of-Interest and Intoxication Classification", in the **Computer Speech and Language** journal, 2012.
233. **William Yang Wang** and Kallirroi Georgila, "Automatic Detection of Unnatural Word-Level Segments in Unit-Selection Speech Synthesis", in Proceedings of IEEE ASRU 2011, Hawaii, USA, Dec. 11-15, IEEE.
234. **William Yang Wang**, Kapil Thadani and Kathleen McKeown, "Identifying Event Descriptions using Co-training with Online News Summaries", in Proceedings of the 5th International Joint Conference on Natural Language Processing (IJCNLP 2011), full oral paper, Chiang Mai, Thailand, Nov. 8-13, ACL-AFNLP.
235. **William Yang Wang**, Ron Artstein, Anton Leuski, and David Traum, "Improving Spoken Dialogue Understanding Using Phonetic Mixture Models", in Cross-Disciplinary Advances in Applied Natural Language Processing: Issues and Approaches, IGI Global. (This is an invited book chapter based on the FLAIRS 2011 paper.)
236. Fadi Biadisy, **William Yang Wang**, Andrew Rosenberg, Julia Hirschberg, "Intoxication Detection using Phonetic, Phonotactic and Prosodic Cues", in Proceedings of the 12th Annual Conference of the International Speech Communication Association (INTERSPEECH 2011), oral paper, Florence, Italy 28-31 Aug., 2011.
237. **William Yang Wang** and Julia Hirschberg, "Detecting Levels of Interest from Spoken Dialog with Multistream Prediction Feedback and Similarity Based Fusion Learning", in Proceedings of the 12th Annual SIGdial Meeting on Discourse and Dialogue (SIGDIAL 2011), oral paper, Portland, Oregon, USA 17-18 June 2011.
238. **William Yang Wang**, Ron Artstein, Anton Leuski, and David Traum, "Improving spoken dialogue understanding using phonetic mixture models", in Proceedings of the 24th International Florida Artificial Intelligence Research Society Conference (FLAIRS-24), oral paper, Palm Beach, Florida, USA 18-20 May 2011 (Best Paper Honorable Mention Award)

239. **William Yang Wang** and Kathleen McKeown, “Got You!: Automatic Vandalism Detection in Wikipedia with Web-based Syntactic-Semantic Modeling”, in Proceedings of The 23rd International Conference on Computational Linguistics (**COLING 2010**) (Oral Presentation Paper Acceptance Rate: 19%)

Professional Activities

- Elected Member, IEEE Speech and Language Technical Committee IEEE SLTC, 2021-2023.
- Diversity, Inclusion, and Accessibility Chair, the Thirty-eighth Annual Conference on Neural Information Processing Systems (NeurIPS), 2024.
- Area Chair, The International Conference on Learning Representations (ICLR), 2023-2024.
- Senior Area Chair, 2022 Conference on Empirical Methods in Natural Language Processing (EMNLP) 2022 and 2023
- Senior Area Chair, Language Grounding to Vision, Robotics, and Beyond, the 60th Annual Meeting of the Association for Computational Linguistics (ACL) 2022 and 2024
- Senior Area Chair, The 2nd Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics and the 12th International Joint Conference on Natural Language Processing (AAACL-IJCNLP) 2022
- Senior Area Chair, Information Extraction, 2021 Annual Conference of the North American Chapter of the Association for Computational Linguistics (NAACL HLT) 2021
- Senior Area Chair, Interpretability and Analysis of Models for NLP, *Conference on Empirical Methods in Natural Language Processing* (EMNLP) 2020
- Best Paper Award Committee, the Thirty-Sixth AAAI Conference on Artificial Intelligence (AAAI) 2022
- Area Chair, the Thirty-Sixth AAAI Conference on Artificial Intelligence (AAAI) 2022
- Area Chair, the 31st International Joint Conference on Artificial Intelligence (IJCAI-ECAI) 2022
- Area Chair, the 2022 IEEE Spoken Language Technology Workshop (SLT) 2022
- Reviewer, Nature 2023
- Reviewer, IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) 2022
- Meta Reviewer, the 46th international Conference on Acoustics, Speech, & Signal Processing (ICASSP) 2021
- Area Chair, *the 58th Annual Meeting of the Association for Computational Linguistics* (ACL) 2020
- Area Chair, Machine Learning, *Conference on Empirical Methods in Natural Language Processing* (EMNLP) 2019
- Area Chair, *the Thirty-Fourth AAAI Conference on Artificial Intelligence* (AAAI) 2020
- Area Chair, *the 57th Annual Meeting of the Association for Computational Linguistics* (ACL) 2019
- Area Chair, *the 33rd AAAI Conference on Artificial Intelligence* (AAAI) 2019
- Area Chair, *the 17th Annual Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies* (NAACL-HLT) 2019

- Best Paper Committee, *Conference on Empirical Methods in Natural Language Processing (EMNLP)* 2018
- Area Chair, Information Extraction and Question Answering, *Conference on Empirical Methods in Natural Language Processing (EMNLP)* 2018
- Area Chair, Computational Social Science, *the 16th Annual Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT)* 2018
- Senior Program committee member, *the 27th International Joint Conference on Artificial Intelligence and the 23rd European Conference on Artificial Intelligence (IJCAI-ECAI)* 2018
- Area Chair, Machine Learning, *the 55th Annual Meeting of the Association for Computational Linguistics (ACL)* 2017
- Publicity Co-Chair, the 28th ACM International Conference on Information and Knowledge Management (CIKM), 2019.
- Publicity Co-Chair, the IEEE International Conference on Data Mining (ICDM), 2018.
- Proposal reviewer, Schmidt Futures 2023
- Selection Committee, Alexander von Humboldt Endowed Chair Professorship, 2021
- Proposal reviewer, Czech Science Foundation, 2021
- Proposal reviewer, Singapore National Research Foundation (NRF), 2020.
- Proposal reviewer, Research Grants Council (RGC) of Hong Kong, 2020
- Panelist, National Science Foundation, 2021.
- Proposal reviewer, Natural Sciences and Engineering Research Council of Canada, 2018.
- Proposal reviewer, Chilean National Science and Technology Commission (CONICYT - Chile), 2018.
- Panelist, National Science Foundation, Directorate for Computer & Information Science & Engineering, Core Programs (NSF CISE) 2017
- Proposal reviewer, Research Foundation Flanders (formerly National Fund for Scientific Research), Belgium 2017
- Area Chair, Information Extraction, *China National Conference on Computational Linguistics (CCL)* 2017
- Publicity Co-chair, *China National Conference on Computational Linguistics* 2018
- Publicity Co-chair, *Natural Language Processing and Chinese Computing (NLPCC)* 2017
- Area Chair, *Natural Language Processing and Chinese Computing (NLPCC)* 2018
- Session chair for two text mining sessions, *the 23rd ACM International Conference on Information and Knowledge Management (CIKM)* 2014
- Session chair for machine learning and data mining, *the 22nd ACM International Conference on Information and Knowledge Management (CIKM)* 2013
- PC co-chair, *the 11th annual Language Technologies Institute Student Research Symposium (SRS)* 2013
- Program committee member, *International Conference on Learning Representations* 2019
- Program committee member, *the 24th International Joint Conference on Artificial Intelligence (IJCAI)* 2015

- Program committee member, *the Conference on Empirical Methods in Natural Language Processing (EMNLP)* 2015-2017
- Program committee member, *the SIGNLL Conference on Computational Natural Language Learning (CoNLL)* 2018
- Technical program committee member, *the 16th Annual Conference of the International Speech Communication Association (INTERSPEECH)* 2015-2018
- Program committee member, *ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD)* 2018
- Program committee member, *the North American Chapter of the Association for Computational Linguistics - Human Language Technologies (NAACL-HLT)* 2015, 2016
- Program committee member, *the International Conference on Computational Linguistics (COLING)* 2016
- Program committee member, *the AAAI Conference on Artificial Intelligence (AAAI)* 2018
- Program committee member, *the Annual Meeting of the Association for Computational Linguistics (ACL)* 2018
- Program committee member, *the International Conference on Machine Learning (ICML)* 2018, 2019
- Reviewer, *Neural Information Processing Systems (NIPS)* 2018
- Scientific committee member, *the International Conference on Language Resources and Evaluation (LREC)* 2018
- Program committee member, *the 1st IEEE Workshop on Fake Multimedia (FakeMM)* 2018
- Program committee member, *1st Workshop on Trolling, Aggression and Cyberbullying (TRAC-1) at 27th International Conference on Computational Linguistics (COLING-2018) (TRAC-1)* 2018
- Program committee member, *the International Joint Conference on Natural Language Processing (IJCNLP)* 2017
- Program committee member, *the International Conference on Affective Computing and Intelligent Interaction (ACII)* 2017
- Program committee member, *the Annual Meeting of the Association for Computational Linguistics Student Research Workshop (ACL SRW)* 2016, 2017
- Program committee member, *the Annual Meeting of the Association for Computational Linguistics Demo Session (ACL Demo)* 2017, 2018
- Program committee member, *the Conference on Empirical Methods in Natural Language Processing (EMNLP Demo)* 2018
- Program committee member, *European Chapter of the Association for Computational Linguistics (EACL SRW)* 2017
- Program committee member, *the Women and Underrepresented Minorities in NLP Workshop (WiNLP)* 2017, 2018
- Program committee member, *International Workshop on Natural Language Processing for Social Media (SocialNLP) at IJCAI/EMNLP/EACL* 2016, 2017

- Program committee member, *Workshop on Declarative Learning Based Programming (DeLBP)* at IJCAI 2017 and AAAI 2018
- Program committee member, *ACM SIGSPATIAL Workshop on AI and Deep Learning for Geographic Knowledge Discovery* at ACM SIGSPATIAL 2017-2018
- Program committee member, *the 5th CCF Conference on Natural Language Processing & Chinese Computing (NLPCC)* 2016
- Program committee member, *the ACM International Conference on Information and Knowledge Management (CIKM)* 2014,2015,2016
- Program committee member, *the 2nd Mid-Atlantic Student Colloquium on Speech, Language, and Learning (MASC-SLL)* 2012
- Journal reviewer, *IEEE Transactions on Affective Computing (IEEE TAC)* 2010
- Journal reviewer, *IEEE Signal Processing Letters (IEEE SPL)* 2011
- Journal reviewer, *IEEE Communications Letters (IEEE CL)* 2010, 2011
- Journal reviewer, *Journal of the American Society for Information Science and Technology (JASIST)* 2012, 2014 x 2
- Journal reviewer, *Educational Research Review (ERR)* 2012
- Journal reviewer, *IEEE Transactions on Audio, Speech, and Language Processing (IEEE TASL)* 2013 x 3, 2014
- Journal reviewer, *Bioinformatics* 2013
- Journal reviewer, *Journal of Medical Internet Research (JMIR)* 2013, 2014,2016
- Journal reviewer, *Artificial Intelligence (AIJ)* 2014
- Journal reviewer, *Journal of Artificial Intelligence Research (JAIR)* 2016
- Journal reviewer, *Speech Communication* 2016
- Journal reviewer, *Natural Language Engineering (NLE)* 2015,2016,2017,2018
- Journal reviewer, *Knowledge-Based Systems* 2017
- Journal reviewer, *Machine Learning Journal* 2017
- Journal reviewer, *IEEE Transactions on Knowledge and Data Engineering* 2017
- Journal reviewer, *Data Mining and Knowledge Discovery* 2018
- Reviewer, *the IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)* 2015-2019
- Reviewer, *the 18th ACM Conference on Computer Supported Cooperative Work and Social Computing (CSCW)* 2015,2016
- Reviewer, *the ACM CHI Conference on Human Factors in Computing Systems (CHI)* 2015
- Reviewer, *the 20th ACM Conference on Intelligent User Interfaces (IUI)* 2015

- Reviewer, *the 17th ACM Conference on Computer Supported Cooperative Work and Social Computing (CSCW)* 2014
- Organizing committee, *the 7th Young Researchers' Roundtable on Spoken Dialog Systems (YRRSDS)* 2011
- Organizing committee, *the 10th and 11th annual LTI Student Research Symposium (SRS)* 2012, 2013
- Volunteer, *the 2013 Conference on Empirical Methods in Natural Language Processing (EMNLP)* 2013
- Volunteer, *the 12th and 13th annual SIGdial Meeting on Discourse and Dialogue (SIGDIAL)* 2011, 2012
- Student member, *Association for Computational Linguistics (ACL)*, 2011-
- Student member, *Association for Computing Machinery (ACM)*, 2013-
- Graduate student member, *IEEE* 2011-
- Student member, *IEEE Signal Processing Society (IEEE SPS)*, 2011-
- Affiliate member, *Speech and Language Processing Technical Committee, IEEE Signal Processing Society*, 2011-
- Member, *SIGdial/SIGSEM/SIGNLL/SIGDAT*, 2011-

Recent Presentations and Invited Talks

1. Principles of Reasoning: Compositional and Collaborative Generative AI, Invited Talks at various universities, 2024.
2. Self-Supervised Language-and-Vision Reasoning, Invited Talks at Amazon Alexa AI, 2020.
3. Deep Adversarial Learning for NLP?, NAACL 2019 Tutorial with Sameer Singh.
4. Automatic Fact Verification with Semi-Structured Knowledge, Singapore Management University, EMNLP 2019 Fever Workshop, 2019.
5. Self-Supervised Natural Language Processing, CMU, Yale, Maryland, GaTech, USC, UChicago, UCLA, Singapore NLP Day, 2019.
6. "Deep Reinforcement Learning for NLP: Classification, Generation, and Reasoning", Peking University, iFlyTek-HIT Lab, ByteDance Inc, Beijing, China, Summer 2018.
7. "Frontiers in Task-Oriented Dialogue Systems: Scalability, Empathy, and Multilinguality", Lenovo AI Research, Beijing, China, Summer 2018.
8. "Deep Knowledge Graph Reasoning", Tsinghua, CASIA, Microsoft Research Asia, CCKS 2018, CIPS Summer School, Beijing and Tianjin, China, Summer 2018.
9. "Deep Reinforcement Learning for NLP", tutorial at ACL 2018 with Jiwei Li and Xiaodong He, Melbourne, Victoria, Australia, July 15, 2018.
10. "Artificial Intelligence: Past, Presence, and Future", Invited Talk at the UC Legislative Roundtable in the CA State Capitol, Sacramento, CA, June 20th, 2018.
11. NAACL 2018 Tutorial on Scalable Knowledge Base Construction and Reasoning, with Xiang Ren and Nanyun Peng (USC), New Orleans, LA, June 1st, 2018.

12. "KBGAN: Adversarial Learning for Negative Sampling", NAACL 2018, New Orleans, LA, June 4th, 2018.
13. "Reinforced Semi-Supervised Learning", Keynote at **Google I/O Extended at Santa Barbara**, Santa Barbara, CA, May 9th, 2018.
14. "Deep Reinforcement Learning for Knowledge Graph Reasoning", at **Facebook**, Menlo Park, CA, Oct 20th, 2017.
15. "Natural Language Processing and Automated Fake News Detection", at **Sonoma State University**, Rohnert Park, CA, Oct 19th, 2017.
16. "Deep Reinforcement Learning for Knowledge Graph Reasoning", at **UCSB PStat Seminar**, Santa Barbara, CA, Oct 11th, 2017.
17. "Deep Reinforcement Learning for Knowledge Graph Reasoning", at **IBM T. J. Watson Research Center**, Yorktown Heights, NY, July 26th, 2017.
18. "Scalable Learning and Reasoning for Large Knowledge Graphs", at **University of Southern California**, Los Angeles, CA, Sept. 25, 2016.
19. "Scalable Learning and Reasoning for Large Knowledge Graphs", at **University of Virginia**, Charlottesville, VA, Apr. 25, 2016.
20. "Scalable Learning and Reasoning for Large Knowledge Graphs", at **Toyota Technological Institute at Chicago**, Chicago, IL, Apr. 13, 2016.
21. "Scalable Learning and Reasoning for Large Knowledge Graphs", at **Northeastern University**, Boston, MA, Apr. 8, 2016.
22. "Scalable Learning and Reasoning for Large Knowledge Graphs", at **University of Rochester**, Rochester, NY, Mar. 31, 2016.
23. "Scalable Learning and Reasoning for Large Knowledge Graphs", at **Rutgers University**, New Brunswick, NJ, Mar. 29, 2016.
24. "Scalable Learning and Reasoning for Large Knowledge Graphs", at **University of Texas, Dallas**, Richardson, TX, Mar. 21, 2016.
25. "Scalable Learning and Reasoning for Large Knowledge Graphs", at **University of California, Santa Barbara**, Santa Barbara, CA, Mar. 16, 2016.
26. "Scalable Learning and Reasoning for Large Knowledge Graphs", at **Pennsylvania State University**, State College, PA, Mar. 2, 2016.
27. "Scalable Learning and Reasoning for Large Knowledge Graphs", at **New York University**, New York, NY, Nov. 11, 2015.
28. "Scalable Learning and Reasoning for Large Knowledge Graphs", at **Columbia University**, New York, NY, Nov. 9, 2015.
29. "Interdisciplinary Approaches to Computational Social Science", at **Academia Sinica**, Taipei, Taiwan (R.O.C.), Oct 29th, 2015.
30. "Scalable Learning and Reasoning for Large Knowledge Graphs", at **Academia Sinica**, Taipei, Taiwan (R.O.C.), Oct 28th., 2015.

31. "Joint Information Extraction and Reasoning: A Scalable Statistical Relational Learning Approach", at **Microsoft Research Asia**, Beijing, China, July 31st, 2015.
32. "Joint Information Extraction and Reasoning: A Scalable Statistical Relational Learning Approach", at **ACL-IJCNLP 2015**, Beijing, China, July 27th, 2015.
33. "A Low-Rank Approximation Approach to Learning Joint Embeddings of News Stories and Images for Timeline Summarization", at **Yahoo! Summer Science Week**, Sunnyvale, CA, July 21th, 2015.
34. "Joint Information Extraction and Reasoning: A Scalable Statistical Relational Learning Approach", at **Yahoo! NLP group**, New York, NY, July 16th, 2015.
35. "I Can Has Cheezburger? A Nonparanormal Approach to Combining Textual and Visual Information for Predicting and Generating Popular Meme Descriptions", at **Yahoo! Labs NYC** lunch talk, New York, NY, June 25th, 2015.
36. "Structure Learning via Parameter Learning", at the 23rd ACM International Conference on Information and Knowledge Management (**CIKM 2014**), Shanghai, China, Nov. 6th, 2014.
37. "Programming with Personalized PageRank: a Locally Groundable First-Order Probabilistic Logic", at the 22nd ACM International Conference on Information and Knowledge Management (**CIKM 2013**), Burlingame, CA, USA, Oct. 30th, 2013.
38. "Programming with Personalized PageRank: a Locally Groundable First-Order Probabilistic Logic", at the 2013 ICML workshop on Inferring: Interactions between Inference and Learning. (**ICML 2013 workshop**), Atlanta, GA, Jun. 20th, 2013.
39. "Crowdsourcing Natural Language Corpora: Methods and Observations", at the 2013 Microsoft TechFest (**TechFest 2013**), Redmond, WA, Mar. 6-7th, 2013.
40. "Crowdsourcing Natural Language Corpora: Methods and Observations", at the 2013 CMU LTI Student Open House poster session, Pittsburgh, PA, Mar. 1st, 2013.
41. "Learning to Recognize Things Unknown: Studies of Out-of-Domain Detection", **Microsoft Research**, Redmond, Jan. 11th, 2013.
42. "Crowdsourcing Natural Language Corpora: Methods and Observations", **Microsoft Online Services Division**, Bellevue, WA, Dec. 11th, 2012.
43. "Structured Natural Language Elicitation via Crowdsourcing: Methods and Observations", at the second Mid-Atlantic Student Colloquium on Speech, Language and Learning (**MASC 2012**), University of Maryland, College Park, Oct. 19th.
44. "Historical Analysis of Legal Opinions with a Sparse Mixed-Effects Latent Variable Model", at the 50th Annual Meeting of the Association for Computational Linguistics (**ACL 2012**), Jeju Island, Korea, July 10th, ACL.
45. "Love ya, jerkface: using Sparse Log-Linear Models to Build Positive (and Impolite) Relationships with Teens", at the 13th annual SIGdial Meeting on Discourse and Dialogue (**SIGDIAL 2012**), Seoul National University, Seoul, Korea, July 5, 2012.
46. "Crowd-sourcing Natural Language Corpora for Time-Efficient Construction of Natural Language Processing Systems", **Microsoft Research**, Redmond, WA. June 29, 2012.
47. "Automatic Intoxication Detection using Accent Identification Approaches", at the 2011 Mid-Atlantic Student Colloquium on Speech, Language and Learning (**MASC**), Baltimore, MD, Sept. 23th, 2011.

48. “Automatic Intoxication Detection using Accent Identification Approaches”, at the 2011 CMU LTI Student Research Symposium (**SRS 2011**), Pittsburgh, PA, Sept. 9th, 2011.
49. “Detecting Levels of Interest from Spoken Dialog with Multistream Prediction Feedback and Similarity Based Hierarchical Fusion Learning”, at **SIGDIAL 2011**, Portland, OR., June 18th, 2011.
50. “Improving Spoken Language Understanding with Phonetic Mixture Models”, Spoken Language Processing Group of **Columbia University**, NY, USA, Feb. 14th, 2011.
51. “Automatic Vandalism Detection in Wikipedia”, invited talk at the Spoken Language Processing Group of **Columbia University**, NY, USA, Oct. 26th, 2010.
52. “Got You!: Automatic Vandalism Detection in Wikipedia with Web-based Syntactic-Semantic Modeling”, presented at the 23rd International Conference on Computational Linguistics (**COLING 2010**), Beijing, China, Aug. 27th, 2010.
53. “Automatic Vandalism Detection in Wikipedia”, invited talk at Information Science Institute (**ISI**), CA, USA, July 30th, 2010.
54. “Improving Spoken Language Understanding with Phonetic Mixture Models”, invited talk at Institute for Creative Technology, University of Southern California (**USC ICT**), CA, USA, June 29th, 2010.

Postdocs and PhD Students

- Vivek Kulkarni (Postdoc, 2017-2018, now at Twitter)
- Liangming Pan (Postdoc, 2023-2024, now Assistant Professor, Arizona)
- Mai Elshierief (PhD, 2019; Assistant Professor, Northeastern)
- Xin Wang (PhD, 2020; Assistant Professor, UC Santa Cruz)
- Wenhui Chen (PhD, 2021; Assistant Professor, Univ. Waterloo)
- Wenhan Xiong (PhD, 2021; Research Scientist, xAI)
- Zhiyu Chen (PhD, 2022; Assistant Professor, U Texas - Dallas)
- Yijun Xiao (PhD, 2022; Research Scientist, Facebook/Meta AI)
- Sharon Levy (PhD, 2023; Assistant Professor, Rutgers)
- Yi-Lin Tuan (PhD, 2024; now at a startup)
- Tsu-Jui Fu (PhD, 2024; now at Apple)
- Wanrong Zhu (PhD, 2024; now at Adobe)
- Michael Saxon (PhD student, NSF Fellow, B.S. M.S. ASU)
- Alon Albalak (PhD, 2024; now at SynthLabs)
- Kexun Zhang (PhD student)
- Weixi Feng (PhD student)
- Deepak Nathani (PhD student)

- Wenda Xu (PhD student)
- Gyuwan Kim (PhD student)
- Yujie Lu (PhD, 2025; now at Meta GenAI)
- Xinyi Wang (PhD student)
- Jiachen Li (PhD, 2025; now at xAI)
- Qingfu Zhu (Visiting PhD student 2019-2020, HIT)
- Qingyu Yin (Visiting PhD student 2017-2018, HIT)
- Shaolei Wang (Visiting PhD student 2018-2019, HIT)
- Pengda Qin (Visiting PhD student, BUPT)

PhD and MS Committees

- Xin Wang (Co-Chair, MAE, PhD Thesis Proposal, and Defense Committee)
- Pankaj Gupta (PhD Thesis Committee, University of Munich, Germany)
- Vivek Gupta (PhD Thesis Committee, University of Utah)
- Richard Jiang (MAE Committee)
- Yi Ding (MAE Committee)
- Liu Liu (MAE Committee)
- Keqian Li (MAE, PhD Thesis Proposal, and Defense Committee)
- Xinyi Zhang (MAE, PhD Thesis Proposal, and Defense Committee)
- Shiliang Tang (MAE, PhD Thesis Proposal, and Defense Committee)
- Izzeddin Gur (MAE, PhD Thesis Proposal, and Defense Committee)
- Semih Yavuz (MAE, PhD Thesis Proposal, and Defense Committee)
- Bay-Yuan Hsu (MAE, PhD Thesis Proposal, and Defense Committee)
- Anurag Pallaprolu (ECE PhD Committee)
- Jieli Qiu (CMU CSD PhD Committee)
- Wei-Tsung Lin (MAE Committee)
- Brandon Huynh (MAE Committee)
- Maohua Zhu (ECE PhD Thesis Committee)
- Guillaume Bellegarda (ECE PhD Thesis Committee)
- Steve Bako (ECE Thesis Proposal Committee)
- Shuangchen Li (ECE Qualify Exam and Defense)

- Christian Bueno (Math Doctoral Candidacy Exam)
- Nidhi Hiremath (Chair, MS Committee)
- Vivek Pradhan (Chair, MS Committee)
- Shayan Sadigh (MS Committee)
- Zenghui Yan (MS Committee)
- Sharath Rao (MS Committee)
- Dong-Anh Nguyen (MS Committee)

Undergraduate & Masters Student Mentoring

- Mahnaz Koupaee (UCSB CS MS, 2017-2018)
- Sharon Levy (UCSB BS/MS, 2017-2018)
- Ashwini Patil (UCSB MS, 2018)
- Chani Jindal (UCSB MS, 2017)
- Sai Nikhil Maram (UCSB MS, 2018)
- Vivek Pradhan (UCSB MS, 2018-2019)
- Conner Vercellino (UCSB CCS, 2017)
- Brian Lim (UCSB CS, 2018)
- Lawrence Lim (UCSB CS, 2018, co-advised with Divy Agrawal)
- Yuhao Zhang (UCSB CS, 2017-2018)
- Deren Lei (UCSB CS, 2017-2019)
- Alexander Sniff (UCSB Linguistics, 2017-2019)
- Samantha Lee (UCSB Statistics, 2017)
- Yuan Yao (UCSB Computer Science, 2017)
- Thien Hoang (UCSB College of Creative Studies, 2016-2017)
- Daniel Spokoyny (UCSB CCS, 2016-now, NSF Graduate Fellow, Chancellor's Award in Undergraduate Research, going to CMU for PhD)
- Darren Huang (National Taiwan University, 2016-17)
- Tyler Vuong (UCSB Electrical Engineering, 2016-17, going to CMU for PhD)
- Ke Ni (UCSB Computer Science, 2016-now, CRA Undergraduate Outstanding Undergraduate Researcher Award Honorable Mention, 2018)
- Shao-Yi Zhang (UCSB CS & Stats, 2016-17)
- Tristan Hull (Cornell, 2014)
- Chao-Yuan Wu (MLD, CMU, 2014, now at UT-Austin)
- Sneha Chaudhari (LTI, CMU, 2014-2015)
- Adele Chase (Columbia University, 2010)

Undergraduate & Highschool Interns

- Sahil Jain (2018 Summer)
- Akul Gupta (Evergreen Valley High School, San Jose, 2017 Summer, now UCSC)
- Xianda Zhou (Tsinghua, 2017 Summer, now UT-Austin)
- Yuteng Ma (Tsinghua, 2017 Summer)
- Xuanyue Yang (BUPT, 2017 Summer, now CMU)
- Fengyu Deng (SJTU, 2017 Summer, now UT-Austin)
- Zikun Lin (Tsinghua, 2017 Summer, now Columbia)
- Shiyi Han (Beihang, 2017 Summer, now Brown)
- Liwei Cai (Tsinghua, 2017 Summer, now CMU)
- Zheyu Wang (Tsinghua, 2017 Summer, now USC)
- Fangjun Zhang (HKU, 2017 Summer, now NYU)

Teaching Experience

- NAACL 2019 Tutorial on Deep Adversarial Learning for NLP.
- ACL 2018 Tutorial on Deep Reinforcement Learning for NLP.
- NAACL 2018 Tutorial on Knowledge Graph Reasoning.
- UC Santa Barbara, Spring 2017
Instructor, CS 292F, *Deep Learning for Natural Language Processing*.
- UC Santa Barbara, Winter 2017
Instructor, CS 190I, *Introduction to Natural Language Processing*.
- The 15th Annual Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, San Diego, USA Summer 2016
Tutorial Speaker, NAACL HLT 2016, *Scalable Statistical Relational Learning for NLP*.
- International Joint Conference on Artificial Intelligence, New York, USA Summer 2016
Tutorial Speaker, IJCAI 2016, *Scalable Probabilistic Logics*.
- Peking University, China, Summer 2015
Instructor, CIPS Summer School, *Information Extraction*, 250 students.
- Carnegie Mellon University, USA, Spring 2014
Head Teaching Assistant, *Machine Learning for Large Datasets*, 150 students.
- Carnegie Mellon University, USA, Fall 2013
Head Teaching Assistant, *Machine Learning*, 60 students.
- Shenzhen University, China, Spring 2008
Teaching Assistant, *Programming Language in Visual Basics*, 150 students.

Media Coverage and Invited Stories

- A Few Social Media Influencers Are Shaping AI Tweets lead to attention, which leads to higher citation numbers, Julianne Pepitone, *IEEE Spectrum*, 03/02/2024.
- Apple releases 'MGIE', a revolutionary AI model for instruction-based image editing, Michael Nuñez, *VentureBeat*, 2/6/2024.
- How does ChatGPT decide what to say next? Here's a quick explainer., Sindhu Sundar, *Business Insider*, 08/31/2023.
- Scammers Used ChatGPT to Unleash a Crypto Botnet on X, Will Knight, *Wired*, 08/21/2023.
- AI2 researchers release new multimodal approach to boost AI capabilities using images and audio, RICHARD YONCK, *GeekWire*, 04/29/2023.
- TechScape: Google and Microsoft are in an AI arms race - who wins could change how we use the internet, Chris Stokel-Walker, *The Guardian*, 02/21/2023.
- AIs become smarter if you tell them to think step by step, Chris Stokel-Walker, *New Scientist*, 10/26/2022.
- Lecture focuses on responsible AI, Katherine Zehnder, *Santa Barbara News-Press*, 1/31/2022.
- Google Fixes Two Annoying Quirks in Its Voice Assistant, Julian Chokkattu, *Wired*, 04/28/2021.
- Artificial Intelligence Learns to Talk Back to Bigots, Chris Intagliata, *Scientific American*, 10/10/2019.
- How marketers are increasingly using A.I. to persuade you to buy, David Z. Morris, **Fortune**, Jan. 31, 2020.
- Reddit and Gab's most toxic communities inadvertently train AI to combat hate speech, **The Next Web**, Sept 25, 2019, (Other media coverage Synced and NPR MarketPlace Tech).
- Modeling Misinformation, Sonia Fernandez, **The UC Santa Barbara Current**, July 29, 2019.
- CVPR 2019 Attracts 9K Attendees; Best Papers Announced; ImageNet Honoured 10 Years Later, **Synced**, June 18, 2019 (Other media coverage VentureBeat, IEEE Computer Society, Noozhawk, UCSB Current).
- An Algorithm That Can Spot When People Lie to the Police, Emiliano Rodriguez Mega, **Scientific American**, Feb 1, 2019.
- Artificial Intelligence: What's Next?, Tom Lydon, **NASDAQ.com**, Sept 25, 2018.
- What's Next: How to Spot an ICO Scam Without Really Trying, **Law.com**, Mar. 20, 2018.
- Researchers tap AI in the fight against ICO scams, Tristan Greene, **The Next Web**, Mar. 13, 2018.
- How to combat fake news and disinformation, Darrell M. West, **The Brookings Institution**, Dec. 18, 2017.
- 'Siri, why have you fallen behind other digital assistants?', Tom Simonite, **Wired**, Oct. 5, 2017.
- The iPhone, Siri and Conversation, Sonia Fernandez, **The UC Santa Barbara Current**, July 7, 2017.
- Teaching Machines to Detect Fake News Is Really Hard, Daniel Oberhaus, **VICE Motherboard**, May 7th, 2017.

- Google Can Finally Understand Your Drunk Texts, John Pavlus, **Fast Company**, Apr. 21st, 2017.
- Machine Learning: The Present and Future, Yehua Zhu, **Science & Technology Review**, 34(7), Apr. 15th, 2016, China Association for Science and Technology.
- AlphaGo Beats Lee Sedol, Scientists: Not Surprising, Shiqi Zhu, **IFENG.com**, Mar. 9th, 2016.
- ‘Software’ con sentido del humor, o cómo crear memes automáticos (y graciosos), José Manuel Blanco, **El Confidencial**, Oct. 20, 2015.
- Can a Computer Make a Meme?, Gretchen McCulloch, **Mental Floss**, Sept. 22, 2015.
- The Theory of “Technological Singularity”: A Scientific Theory of Artificial Intelligence? Or Junk Science?, William Wang, **Communications of the China Computer Federation**, June Issue, 2015.
- Proof: the Science of Booze, Adams Roger, **Mariner Books**, June, 2014. Best Science Book of 2014 by Amazon, Wired, the Guardian, and NBC, Winner of the 2014 Gourmand Award for Best Spirits Book in the United States.
- Big Data and Machine Learning: a Retrospective of ICML 2013, **Programmer Magazine**, Sept. issue, 2013.
- Acquiring Corpora using Crowdsourcing, Matt Marge, **IEEE SLTC Newsletter**, May, 2013.
- Detecting Intoxication in Speech, Matt Marge, **IEEE SLTC Newsletter**, Oct., 2011.

Public Policy and Outreach

- UCOP AI Congress at UCLA, 2024.
- UC AI Congressional Briefing at the Capitol, 2023.
- Senate Commerce Committee AI Briefing, 2023.
- White House OSTP AI Briefing, 2023.
- UC AI Briefing at CA State Capitol, 2018.
- Host Sahil Jain (The Harker School in California) for UCSB Summer Research Mentorship Program, Summer 2018.
- Host a visit for McNair Scholars Program, 04/20/2018.
- Host a visit from Laguna Blanca School, Santa Barbara, CA 03/08/2018.
- Hosted Akul Gupta (Evergreen Valley High School, San Jose, 2017 Summer, Now B.S. UCSC) for a research internship, 07/2017.
- Invited talk at the Santa Barbara Humanist Society, 2017.
- AI microblogs on Twitter and Weibo: 170,000+ followers, and 4,000,000+ monthly views.