HANNANEH HAJISHIRZI

hannaneh.ai
last updated: Jun 2025

EDUCATIONAL HISTORY

University of Illinois at Urbana-Champaign, Urbana-Champaign, IL
Ph.D., Computer Science
September 2011
Thesis: Action-centered Reasoning for Probabilistic Dynamic Systems
Sharif University of Technology, Tehran, Iran
B.Sc., Computer Engineering (Software)

EMPLOYMENT HISTORY

Senior Director, Research Oct 2021 – present
Allen Institute for AI, Seattle, WA
Professor Sep 2025 – present
Paul G. Allen School for Computer Science & Engineering, University of Washington, Seattle, WA
Associate Professor Sep 2022 – Sep 2025
Paul G. Allen School for Computer Science & Engineering, University of Washington, Seattle, WA
Assistant Professor Sep 2018 – Sept 2022
Paul G. Allen School for Computer Science & Engineering, University of Washington, Seattle, WA
AI Fellow Sep 2018 – Oct 2021
Allen Institute for AI, Seattle, WA
Research Assistant Professor Jan 2015 – Sep 2018
Electrical Engineering, University of Washington, Seattle, WA
Research Scientist Sep 2012 – Jan 2015
Electrical Engineering, University of Washington, Seattle, WA
Postdoctoral Associate Sep 2011 – Sep 2012
Disney Research, Pittsburgh, PA

TEAM LEADING

Senior Director of AI at Ai2

1. Manage the open echosystem team
2. Lead the OLMo project, pre-training LLMs
3. Lead the Tulu project, post-training LLMs

AWARDS AND HONORS

1. Best demo paper award, ACL 2025.
2. Best paper honorable mention, CVPR 2025.
3. Uncommon thinker award by GeekWire, 2024.
4. VentureBeat Women of AI award finalist, 2024, 2025.
5. Best theme paper award, ACL 2024.
6. Best Resource paper award, ACL 2024.
7. NAIRR Inaugural Awardee, 2024

8. Frontier of Engineering Committee at National Academy of Engineering, 2024.
9. Innovation of the year award by GeekWire, 2024
10. UIUC CS Early Career Academic Achievement Alumni Award, 2023
11. Torode Family Career Development Professorship, 2022
12. NSF CAREER award, 2021
13. Intel Rising Star faculty award, 2020
14. Sloan Fellowship, 2020
15. Facebook Research Award, 2020
16. Best paper honorable mention, AKBC, 2020
17. Allen Distinguished Investigator Award, 2019
18. Best paper award finalist,
19. Amazon Research Faculty award, 2018
20. Bloomberg Data Science award, 2017
21. Google Research Faculty award, 2015, 2018
22. Allen Distinguished Investigator Award, 2014
23. SigDial Best Paper Award, 2012
24. CS/AI Award, Beckman Institute, UIUC, 2010

PUBLICATIONS

Citation indices (per Google Scholar, as of June 2024): citations (41,019), h-index (88), i10-index (177) Google scholar link has the most recent preprints: `https://scholar.google.com/citations?user=LOV6_WIAAAAJ&hl=en&oi=ao`

Conferences & Journals

[2025]

1. Nathan Lambert, Jacob Morrison, Valentina Pyatkin, Shengyi Huang, Hamish Ivison, Faeze Brahman, Lester James Validad Miranda, Alisa Liu, Nouha Dziri, Xinxi Lyu, Yuling Gu, Saumya Malik, Victoria Graf, Jena D. Hwang, Jiangjiang Yang, Ronan Le Bras, Oyvind Tafjord, Christopher Wilhelm, Luca Soldaini, Noah A. Smith, Yizhong Wang, Pradeep Dasigi, Hannaneh Hajishirzi
   "Tulu 3: Pushing Frontiers in Open Language Model Post-Training."
   Conference on Language Modeling (COLM), 9 pages, 2025.
2. Zhiyuan Zeng, Yizhong Wang, Hannaneh Hajishirzi, Pang Wei Koh.
   " EvalTree: Profiling Language Model Weaknesses via Hierarchical Capability Trees."
   Conference on Language Modeling (COLM), 9 pages, 2025.
3. Akshita Bhagia, Jiacheng Liu, Alexander Wettig, David Heineman, Oyvind Tafjord, Ananya Harsh, Luca Soldaini, Noah A. Smith, Dirk Groeneveld, Pang Wei Koh, Jesse Dodge, Hannaneh Hajishirzi.
   " Establishing Task Scaling Laws via Compute-Efficient Model Ladders."
   Conference on Language Modeling (COLM), 9 pages, 2025.
4. Tong Chen, Faeze Brahman, Jiacheng Liu, Niloofar Mireshghallah, Weijia Shi, Pang Wei Koh, Luke Zettlemoyer, Hannaneh Hajishirzi.
   "ParaPO: Aligning Language Models to Reduce Verbatim Reproduction of Pre-training Data."
   Conference on Language Modeling (COLM), 9 pages, 2025.
5. OLMo Team, Evan Pete Walsh, Luca Soldaini, Dirk Groeneveld, Kyle Lo, Shane Arora, Akshita Bhagia, Yuling Gu, Shengyi Huang, Matt Jordan, Nathan Lambert, Dustin Schwenk, Oyvind Tafjord, Taira Anderson, David Atkinson, Faeze Brahman, Christopher Clark, Pradeep Dasigi, Nouha Dziri, Allyson Ettinger, Michal Guerquin, David Heineman, Hamish Ivison, Pang Wei Koh, Jiacheng Liu, Saumya Malik, William Merrill, Lester James Validad Miranda, Jacob Morrison, Tyler Murray, Crystal Nam, Jake Poznanski, Valentina Pyatkin, Aman Rangapur, Michael Schmitz, Sam Skjonsberg, David Wadden, Christopher Wilhelm, Michael Wilson, Luke Zettlemoyer, Ali Farhadi, Noah A. Smith, Hannaneh Hajishirzi.
   "2 OLMo 2 Furious"
   Conference on Language Modeling (COLM), 9 pages, 2025.

6. Matt Deitke, Christopher Clark, Sangho Lee, Rohun Tripathi, Yue Yang, Jae Sung Park, Mohammadreza Salehi, Niklas Muennighoff, Kyle Lo, Luca Soldaini, Jiasen Lu, Taira Anderson, Erin Bransom, Kiana Ehsani, Huong Ngo, YenSung Chen, Ajay Patel, Mark Yatskar, Chris Callison-Burch, Andrew Head, Rose Hendrix, Favyen Bastani, Eli VanderBilt, Nathan Lambert, Yvonne Chou, Arnavi Chheda, Jenna Sparks, Sam Skjonsberg, Michael Schmitz, Aaron Sarnat, Byron Bischoff, Pete Walsh, Chris Newell, Piper Wolters, Tanmay Gupta, Kuo-Hao Zeng, Jon Borchardt, Dirk Groeneveld, Crystal Nam, Sophie Lebrecht, Caitlin Wittlif, Carissa Schoenick, Oscar Michel, Ranjay Krishna, Luca Weihs, Noah A. Smith, Hannaneh Hajishirzi, Ross Girshick, Ali Farhadi, Aniruddha Kembhavi.
   "Molmo and PixMo: Open Weights and Open Data for State-of-the-Art Multimodal Models."
   CVPR, 9 pages, 2025
   Best paper honorable mention
7. Valentin Hofmann, David Heineman, Ian Magnusson, Kyle Lo, Jesse Dodge, Maarten Sap, Pang Wei Koh, Chun Wang, Hannaneh Hajishirzi, Noah A. Smith.
   "Fluid Language Model Benchmarking."
   Conference on Language Modeling (COLM), 9 pages, 2025.
8. Yuling Gu, Oyvind Tafjord, Bailey Kuehl, Dany Haddad, Jesse Dodge, Hannaneh Hajishirzi.
   "OLMES: A Standard for Language Model Evaluations."
   Findings of NAACL, 9 pages, 2025.
9. David Wadden, Kejian Shi, Jacob Morrison, Aakanksha Naik, Shruti Singh, Nitzan Barzilay, Kyle Lo, Tom Hope, Luca Soldaini, Zejiang Shen, Doug Downey, Hannaneh Hajishirzi, Arman Cohan.
   SciRIFF: A Resource to Enhance Language Model Instruction-Following over Scientific Literature
   ACL, 9 pages, 2025.
10. Lester James V. Miranda*, Yizhong Wang*, Yanai Elazar, Sachin Kumar, Valentina Pyatkin, Faeze Brahman, Noah A. Smith, Hannaneh Hajishirzi, and Pradeep Dasigi.
    " Hybrid Preferences: Learning to Route Instances for Human vs. AI Feedback."
    ACL, 9 pages, 2025.
11. Jiacheng Liu, Taylor Blanton, Yanai Elazar, Sewon Min, YenSung Chen, Arnavi Chheda-Kothary, Huy Tran, Byron Bischoff, Eric Marsh, Michael Schmitz, Cassidy Trier, Aaron Sarnat, Jenna James, Jon Borchardt, Bailey Kuehl, Evie Cheng, Karen Farley, Sruthi Sreeram, Taira Anderson, David Albright, Carissa Schoenick, Luca Soldaini, Dirk Groeneveld, Rock Yuren Pang, Pang Wei Koh, Noah A. Smith, Sophie Lebrecht, Yejin Choi, Hannaneh Hajishirzi, Ali Farhadi, Jesse Dodge.
    "OLMoTrace: Tracing Language Model Outputs Back to Trillions of Training Tokens. "
    ACL Demo Track, 9 pages, 2025.
12. Yuling Gu, Oyvind Tafjord, Bailey Kuehl, Dany Haddad, Jesse Dodge, Hannaneh Hajishirzi.
    "OLMES: A Standard for Language Model Evaluations."
    Findings of NAACL, 9 pages, 2025.
13. Nathan Lambert, Valentina Pyatkin, Jacob Morrison, Lester James Validad Miranda, Bill Yuchen Lin, Khyathi Chandu, Nouha Dziri, Sachin Kumar, Tom Zick, Yejin Choi, Noah A. Smith, Hannaneh Hajishirzi.
    "RewardBench: Evaluating Reward Models for Language Modeling."
    Findings of NAACL, 9 pages, 2025.
14. Sachin Kumar, Chan Young Park, Yulia Tsvetkov, Noah A. Smith, Hannaneh Hajishirzi.
    "ComPO: Community Preferences for Language Model Personalization."
    NAACL, 9 pages, 2025.
15. Joongwon Kim, Anirudh Goyal, Aston Zhang, Bo Xiong, Rui Hou, Melanie Kambadur, Dhruv Mahajan, Hannaneh Hajishirzi, Liang Tan.
    "A Systematic Examination of Preference Learning through the Lens of Instruction-Following."
    NAACL, 9 pages, 2025.
16. Niklas Muennighoff, Luca Soldaini, Dirk Groeneveld, Kyle Lo, Jacob Morrison, Sewon Min, Weijia Shi, Evan Pete Walsh, Oyvind Tafjord, Nathan Lambert, Yuling Gu, Shane Arora, Akshita Bhagia, Dustin Schwenk, David Wadden, Alexander Wettig, Binyuan Hui, Tim Dettmers, Douwe Kiela, Ali Farhadi, Noah A. Smith, Pang Wei Koh, Amanpreet Singh, Hannaneh Hajishirzi.
    "OLMoE: Open Mixture-of-Experts Language Models."
    International Conference of Learning Represenations (ICLR), Oral, 2025.
17. Sarah Wiegreffe, Oyvind Tafjord, Yonatan Belinkov, Hannaneh Hajishirzi, Ashish Sabharwal.
    "Answer, Assemble, Ace: Understanding How LMs Answer Multiple Choice Questions."

International Conference of Learning Represenations (ICLR), Spotlight, 2025.

[2024]

18. Hamish Ivison, Yizhong Wang, Jiacheng Liu, Zeqiu Wu, Valentina Pyatkin, Nathan Lambert, Noah A. Smith, Yejin Choi, Hannaneh Hajishirzi
"Unpacking DPO and PPO: Disentangling Best Practices for Learning from Preference Feedback."
NeurIPS, 9 pages, 2024.
19. Fnu Devvrit, Sneha Kudugunta, Aditya Kusupati, Tim Dettmers, Kaifeng Chen, Inderjit S Dhillon, Yulia Tsvetkov, Hannaneh Hajishirzi, Sham M. Kakade, Ali Farhadi, Prateek Jain
"MatFormer: Nested Transformer for Elastic Inference."
NeurIPS, 9 pages, 2024.
20. Ruizhe Shi, Yifang Chen, Yushi Hu, Alisa Liu, Hannaneh Hajishirzi, Noah A. Smith, Simon Shaolei Du
"Decoding-Time Language Model Alignment with Multiple Objectives."
NeurIPS, 9 pages, 2024.
21. Faeze Brahman, Sachin Kumar, Vidhisha Balachandran, Pradeep Dasigi, Valentina Pyatkin, Abhilasha Ravichander, Sarah Wiegreffe, Nouha Dziri, Khyathi Chandu, Jack Hessel, Yulia Tsvetkov, Noah A. Smith, Yejin Choi, Hannaneh Hajishirzi
"The Art of Saying No: Contextual Noncompliance in Language Models."
NeurIPS Datasets and Benchmarks Track, 9 pages, 2024.
22. Ian Magnusson, Akshita Bhagia, Valentin Hofmann, Luca Soldaini, Ananya Harsh Jha, Oyvind Tafjord, Dustin Schwenk, Evan Pete Walsh, Yanai Elazar, Kyle Lo, Dirk Groeneveld, Iz Beltagy, Hannaneh Hajishirzi, Noah A. Smith, Kyle Richardson, Jesse Dodge
"Paloma: A Benchmark for Evaluating Language Model Fit."
NeurIPS Datasets and Benchmarks Track, 9 pages, 2024.
23. Mohammadreza Salehi, Jae Sung Park, Aditya Kusupati, Ranjay Krishna, Yejin Choi, Hannaneh Hajishirzi, Ali Farhadi
"ActionAtlas: A VideoQA Benchmark for Domain-specialized Action Recognition ."
NeurIPS Datasets and Benchmarks Track, 9 pages, 2024.
24. Gary (Jiacheng) Liu, Sewon Min, Luke Zettlemoyer, Yejin Choi, Hannaneh Hajishirzi
"Infini-gram: Scaling Unbounded n-gram Language Models to a Trillion Tokens."
Conference on Language Modeling (COLM), 9 pages, 2024.
25. Abhika Mishra, Akari Asai, Vidhisha Balachandran, Yizhong Wang, Graham Neubig, Yulia Tsvetkov, and Hannaneh Hajishirzi
"Fine-grained Hallucination Detection and Editing for Language Models."
Conference on Language Modeling (COLM), 9 pages, 2024.
26. Michael Duan, Anshuman Suri, Niloofar Mireshghallah, Sewon Min, Weijia Shi, Luke Zettlemoyer, Yulia Tsvetkov, Yejin Choi, David Evans, and Hannaneh Hajishirzi
"Membership Inference Attacks Work on Large Language Models?."
Conference on Language Modeling (COLM), 9 pages, 2024.
27. Gary (Jiacheng) Liu, Andrew Cohen, Ramakanth Pasunuru, Yejin Choi, Hannaneh Hajishirzi, and Asli Celikyilmaz
" Don't throw away your value model! Generating more preferable text with Value-Guided Monte-Carlo Tree Search decoding."
Conference on Language Modeling (COLM), 9 pages, 2024.
28. Dirk Groeneveld, Iz Beltagy, Pete Walsh, Akshita Bhagia, Rodney Kinney, Oyvind Tafjord, Ananya Harsh Jha, Hamish Ivison, Ian Magnusson, Yizhong Wang, Shane Arora, David Atkinson, Russell Authur, Khyathi Raghavi Chandu, Arman Cohan, Jennifer Dumas, Yanai Elazar, Yuling Gu, Jack Hessel, Tushar Khot, William Merrill, Jacob Morrison, Niklas Muennighoff, Aakanksha Naik, Crystal Nam, Matthew E. Peters, Valentina Pyatkin, Abhilasha Ravichander, Dustin Schwenk, Saurabh Shah, Will Smith, Emma Strubell, Nishant Subramani, Mitchell Wortsman, Pradeep Dasigi, Nathan Lambert, Kyle Richardson and Luke Zettlemoyer, Jesse Dodge, Kyle Lo, Luca Soldaini, Noah A. Smith, and Hannaneh Hajishirzi,
"OLMo: Accelerating the Science of Language Models"
Associations of Computational Linguistics (ACL), 9 pages, 2024.
Best paper award
Innovator of the year award by Geekwire
29. Luca Soldaini, Rodney Kinney, Akshita Bhagia, Dustin Schwenk, David Atkinson, Russell Authur, Ben Bogin,

Khyathi Chandu, Jennifer Dumas, Yanai Elazar, Valentin Hofmann, Ananya Harsh Jha, Sachin Kumar, Li Lucy, Xinxi Lyu, Nathan Lambert, Ian Magnusson, Jacob Morrison, Niklas Muennighoff, Aakanksha Naik, Crystal Nam, Matthew E Peters, Abhilasha Ravichander, Kyle Richardson, Zejiang Shen, Emma Strubell, Nishant Subramani, Oyvind Tafjord, Pete Walsh, Luke Zettlemoyer, Noah A Smith, Hannaneh Hajishirzi, Iz Beltagy, Dirk Groeneveld, Jesse Dodge, Kyle Lo
"Dolma: An Open Corpus of Three Trillion Tokens for Language Model Pretraining Research."
Associations of Computational Linguistics (ACL), 9 pages, 2024.
Best Resource Paper Award

30. Bowen Zhao, Zander Brumbaugh, Yizhong Wang, Hannaneh Hajishirzi, Noah A Smith
"Set the Clock: Temporal Alignment of Pretrained Language Models."
Associations of Computational Linguistics (ACL), 9 pages, 2024.
31. Bowen Zhao, Hannaneh Hajishirzi, Qingqing Cao
"APT: Adaptive pruning and tuning pretrained language models for efficient training and inference."
International Conference on Representation Learning (ICLR), 9 pages, 2024.
32. Y Fu, R Panda, X Niu, X Yue, H Hajishirzi, Y Kim, H Peng
"Data Engineering for Scaling Language Models to 128K Context."
International Conference on Machine Learning (ICML), 9 pages, 2024.
33. Akari Asai, Zeqiu Wu, Yizhong Wang, Avirup Sil, Hannaneh Hajishirzi
"Self-RAG: Learning to Retrieve, Generate, and Critique through Self-Reflection ."
International Conference on Representation Learning (ICLR), 9 pages, 2024.
34. Pan Lu, Hritik Bansal, Tony Xia, Jiacheng Liu, Chunyuan Li, Hannaneh Hajishirzi, Hao Cheng, Kai-Wei Chang, Michel Galley, Jianfeng Gao
"MathVista: Evaluating Mathematical Reasoning of Foundation Models in Visual Context ."
International Conference on Representation Learning (ICLR), 9 pages, 2024.
35. Yanai Elazar, Akshita Bhagia, Ian Helgi Magnusson, Abhilasha Ravichander, Dustin Schwenk, Alane Suhr, Evan Pete Walsh, Dirk Groeneveld, Luca Soldaini, Sameer Singh, Hannaneh Hajishirzi, Noah A. Smith, Jesse Dodge
"What's In My Big Data?"
International Conference on Representation Learning (ICLR), 9 pages, 2024.
36. Sewon Min, Suchin Gururangan, Eric Wallace, Weijia Shi, Hannaneh Hajishirzi, Noah A. Smith, Luke Zettlemoyer
"SILO Language Models: Isolating Legal Risk In a Nonparametric Datastore."
International Conference on Representation Learning (ICLR), 9 pages, 2024.
37. Qingqing Cao, Sewon Min, Yizhong Wang, Hannaneh Hajishirzi
"BTR: Binary Token Representations for Efficient Retrieval Augmented Language Models."
International Conference on Representation Learning (ICLR), 9 pages, 2024.

[2023]

38. Sewon Min, Kalpesh Krishna, Xinxi Lyu, Mike Lewis, Wen-tau Yih, Pang Koh, Mohit Iyyer, Luke Zettlemoyer, Hannaneh Hajishirzi
"FActScore: Fine-grained Atomic Evaluation of Factual Precision in Long Form Text Generation."
Empirical Methods in Natural Language Processing (EMNLP), 9 pages, 2023.
39. Joongwon (Daniel) Kim, Akari Asai, Gabriel Ilharco, Hannaneh Hajishirzi
"TaskWeb: Selecting Better Source Tasks for Multi-task NLP."
Empirical Methods in Natural Language Processing (EMNLP), 9 pages, 2023.
40. Jiacheng Liu, Ramakanth Pasunuru, Hannaneh Hajishirzi, Yejin Choi, Asli Celikyilmaz
"Crystal: Introspective Reasoners Reinforced with Self-Feedback."
Empirical Methods in Natural Language Processing (EMNLP), 9 pages, 2023
41. Jiacheng Liu, Wenya Wang, Dianzhuo Wang, Noah A. Smith, Yejin Choi, Hannaneh Hajishirzi
"Vera: A General-Purpose Plausibility Estimation Model for Commonsense Statements."
Empirical Methods in Natural Language Processing (EMNLP), 9 pages, 2023
42. Dung Thai, Dhruv Agarwal, Mudit Chaudhary, Wenlong Zhao, Rajarshi Das, Jay-Yoon Lee, Hannaneh Hajishirzi, Manzil Zaheer, Andrew McCallum
"Machine Reading Comprehension using Case-based Reasoning."
Findings of Empirical Methods in Natural Language Processing (EMNLP), 9 pages, 2023
43. Mohammadreza Salehi, Sachin Mehta, Aditya Kusupati, Ali Farhadi, Hannaneh Hajishirzi

"SHARCS: Efficient Transformers Through Routing with Dynamic Width Sub-networks."
Findings of Empirical Methods in Natural Language Processing (EMNLP), 5 pages, 2023

44. Zeqiu Wu*, Yushi Hu*, Weijia Shi, Nouha Dziri, Alane Suhr, Prithviraj Ammanabrolu, Noah A. Smith, Mari Ostendorf, Hannaneh Hajishirzi
"Fine-Grained Human Feedback Gives Better Rewards for Language Model Training."
Neural Information Processing (NeurIPS), 2023
45. Dhruba Ghosh, Hannaneh Hajishirzi, Ludwig Schmidt
"GenEval: An Object-Focused Framework for Evaluating Text-to-Image Alignment" Neural Information Processing (NeurIPS) Datasets and Benchmarks Track, 2023
46. Samir Yitzhak Gadre, Gabriel Ilharco, Alex Fang, Jonathan Hayase, Georgios Smyrnis, Thao Nguyen, Ryan Marten, Mitchell Wortsman, Dhruba Ghosh, Jieyu Zhang, Eyal Orgad, Rahim Entezari, Giannis Daras, Sarah Pratt, Vivek Ramanujan, Yonatan Bitton, Kalyani Marathe, Stephen Mussmann, Richard Vencu, Mehdi Cherti, Ranjay Krishna, Pang Wei Koh, Olga Saukh, Alexander Ratner, Shuran Song, Hannaneh Hajishirzi, Ali Farhadi, Romain Beaumont, Sewoong Oh, Alex Dimakis, Jenia Jitsev, Yair Carmon, Vaishaal Shankar, Ludwig Schmidt
"DataComp: In search of the next generation of multimodal datasets."
Neural Information Processing (NeurIPS) Datasets and Benchmarks Track, 2023
47. Yizhong Wang*, Hamish Ivison*, Pradeep Dasigi, Jack Hessel, Tushar Khot, Khyathi Raghavi Chandu, David Wadden, Kelsey MacMillan, Noah A. Smith, Iz Beltagy, Hannaneh Hajishirzi
"How Far Can Camels Go? Exploring the State of Instruction Tuning on Open Resources."
Neural Information Processing (NeurIPS) Datasets and Benchmarks Track, 2023
48. Yizhong Wang, Yeganeh Kordi, Swaroop Mishra, Alisa Liu, Noah A Smith, Daniel Khashabi, Hannaneh Hajishirzi
"Self-Instruct: Aligning Language Model with Self-Generated Instructions."
Associations of Computational Linguistics (ACL), 9 pages, 2023.
49. Alex Mallen, Akari Asai, Victor Zhong, Rajarshi Das, Daniel Khashabi, Hannaneh Hajishirzi
" When Not to Trust Language Models: Investigating Effectiveness of Parametric and Non-Parametric Memories."
Associations of Computational Linguistics (ACL), 9 pages, 2023.
50. Akari Asai, Timo Schick, Patrick Lewis, Xilun Chen, Gautier Izacard, Sebastian Riedel, Hannaneh Hajishirzi, Wen-tau Yih,
"Task-Aware Retrieval with Instructions."
Findings of Associations of Computational Linguistics (ACL), 9 pages, 2023.
51. Qingqing Cao, Bhargavi Paranjape, Hannaneh Hajishirzi
"PuMer: Pruning and Merging Tokens for Efficient Vision Language Models."
Associations of Computational Linguistics (ACL), 9 pages, 2023.
52. Xinyan Velocity Yu, Sewon Min, Luke Zettlemoyer, Hannaneh Hajishirzi
"CREPE: Open-Domain Question Answering with False Presuppositions."
Associations of Computational Linguistics (ACL), 9 pages, 2023.
53. Xinxi Lyu, Sewon Min, Iz Beltagy, Luke Zettlemoyer, Hannaneh Hajishirzi
"Z-ICL: Zero-Shot In-Context Learning with Pseudo-Demonstrations."
Associations of Computational Linguistics (ACL), 9 pages, 2023.
54. Sewon Min, Weijia Shi, Mike Lewis, Xilun Chen, Wen-tau Yih, Hannaneh Hajishirzi, Luke Zettlemoyer
"Nonparametric Masked Language Modeling."
Findings of Associations of Computational Linguistics (ACL), 9 pages, 2023.
55. Hamish Ivison, Akshita Bhagia, Yizhong Wang, Hannaneh Hajishirzi, Matthew Peters
" HINT: Hypernetwork Instruction Tuning for Efficient Zero-Shot Generalisation."
Associations of Computational Linguistics (ACL), 9 pages, 2023.
56. Wenya Wang, Vivek Srikumar, Hanna Hajishirzi, Noah A. Smith
"Elaboration-Generating Commonsense Question Answering at Scale."
Associations of Computational Linguistics (ACL), 9 pages, 2023.
57. Hamish Ivison, Noah A. Smith, Hannaneh Hajishirzi, and Pradeep Dasigi
"Data-Efficient Finetuning Using Cross-Task Nearest Neighbors."
Findings of Associations of Computational Linguistics (ACL), 9 pages, 2023.
58. Qinyuan Ye, Iz Beltagy, Matthew E. Peters, Xiang Ren, Hannaneh Hajishirzi
"FiD-ICL: A Fusion-in-Decoder Approach for Efficient In-Context Learning."

Associations of Computational Linguistics (ACL), 9 pages, 2023.

59. Rajkumar Ramamurthy*, Prithviraj Ammanabrolu*, Kianté Brantley, Jack Hessel, Rafet Sifa, Christian Bauckhage, Hannaneh Hajishirzi, Yejin Choi
"Is Reinforcement Learning (Not) for Natural Language Processing: Benchmarks, Baselines, and Building Blocks for Natural Language Policy Optimization."
International Conference on Representation Learning (ICLR), 9 pages, 2023.
60. Gabriel Ilharco, Marco Tulio Ribeiro, Mitchell Wortsman, Suchin Gururangan, Ludwig Schmidt, Hannaneh Hajishirzi, Ali Farhadi
"Editing Models with Task Arithmetic."
International Conference on Representation Learning (ICLR), 9 pages, 2023.
61. Zeqiu Wu, Ryu Parish, Hao Cheng, Sewon Min, Prithviraj Ammanabrolu, Mari Ostendorf, Hannaneh Hajishirzi
"INSCIT: Information-Seeking Conversations with Mixed-Initiative Interactions."
Transaction of Association for Computational Linguistics (TACL), 2023

[2022]

62. Jiacheng Liu, Skyler Hallinan, Ximing Lu, Pengfei He, Sean Welleck, Hannaneh Hajishirzi, Yejin Choi
"Rainier: Reinforced Knowledge Introspector for Commonsense Question Answering ."
Empirical Methods in Natural Language Processing (EMNLP), 9 pages, 2022.
63. Akari Asai, Mohammadreza Salehi, Mathew Peters, Hannaneh Hajishirzi
"ATTEMPT: Parameter-Efficient Multi-task Tuning via Attentional Mixtures of Soft Prompts ."
Empirical Methods in Natural Language Processing (EMNLP), 9 pages, 2022.
64. Yizhong Wang, Swaroop Mishra, Pegah Alipoormolabashi, Yeganeh Kordi, Amirreza Mirzaei, Anjana Arunkumar, Arjun Ashok, Arut Selvan Dhanasekaran, Atharva Naik, David Stap, Eshaan Pathak, Giannis Karamanolakis, Haizhi Gary Lai, Ishan Purohit, Ishani Mondal, Jacob Anderson, Kirby Kuznia, Krima Doshi, Maitreya Patel, Kuntal Kumar Pal, Mehrad Moradshahi, Mihir Parmar, Mirali Purohit, Neeraj Varshney, Phani Rohitha Kaza, Pulkit Verma, Ravsehaj Singh Puri, Rushang Karia, Shailaja Keyur Sampat, Savan Doshi, Siddhartha Mishra, Sujan Reddy, Sumanta Patro, Tanay Dixit, Xudong Shen, Chitta Baral, Yejin Choi, Noah A. Smith, Hannaneh Hajishirzi, Daniel Khashabi
"Super-NaturalInstructions: Generalization via Declarative Instructions on 1600+ NLP Tasks ."
Empirical Methods in Natural Language Processing (EMNLP), 9 pages, 2022.
65. Sewon Min, Xinxi Lyu, Ari Holtzman, Mikel Artetxe, Mike Lewis, Hannaneh Hajishirzi, Luke Zettlemoyer
"Rethinking the Role of Demonstrations: What makes In-context Learning Work? ."
Empirical Methods in Natural Language Processing (EMNLP), 9 pages, 2022.
66. Zeqiu Wu, Yi Luan, Hannah Rashkin, David Reitter, Hannaneh Hajishirzi, Mari Ostendorf, Gaurav Singh Tomar
"CONQRR: Conversational Query Rewriting for Retrieval with Reinforcement Learning ."
Empirical Methods in Natural Language Processing (EMNLP), 9 pages, 2022.
67. David Wadden, Kyle Lo, Bailey Kuehl, Arman Cohan, Iz Beltagy, Lucy Lu Wang, Hannaneh Hajishirzi
"SciFact-Open: Towards open-domain scientific claim verification ."
Findings of Empirical Methods in Natural Language Processing (EMNLP), 9 pages, 2022.
68. Anas Awadalla, Mitchell Wortsman, Gabriel Ilharco, Sewon Min, Hannaneh Hajishirzi, Ludwig Schmidt
"Exploring The Landscape of Distributional Robustness for Question Answering Models ."
Findings of Empirical Methods in Natural Language Processing (EMNLP), 9 pages, 2022.
69. Sean Welleck, Jiacheng Liu, Ximing Lu, Hannaneh Hajishirzi, Yejin Choi
"NaturalProver: Grounded Mathematical Proof Generation with Language Models."
Neural Information Processing (NeurIPS), 2022.
70. Gabriel Ilharco, Mitchell Wortsman, Samir Yitzhak Gadre, Shuran Song, Hannaneh Hajishirzi, Simon Kornblith, Ali Farhadi, Ludwig Schmidt
"Patching open-vocabulary models by interpolating weights."
Neural Information Processing (NeurIPS), 2022.
71. Sewon Min, Mike Lewis, Luke Zettlemoyer, Hannaneh Hajishirzi.
"MetaICL: Learning to Learn In Context."
North American Chapter of the Association of Computational Linguistics (NAACL), 9 pages, 2022.
72. Daniel Khashabi, Xinxi Lyu, Sewon Min, Lianhui Qin, Kyle Richardson, Sameer Singh, Sean Welleck, Hannaneh Hajishirzi, Tushar Khot, Ashish Sabharwal, Yejin Choi.
"Prompt Waywardness: The Curious Case of Discretized Interpretation of Continuous Prompts."
North American Chapter of the Association of Computational Linguistics (NAACL), 9 pages, 2022.

73. Akari Asai, Matt Gardner, and Hannaneh Hajishirzi.
"Evidentiality-guided Generation for Knowledge-Intensive NLP Tasks."
North American Chapter of the Association of Computational Linguistics (NAACL), 9 pages, 2022.
74. Prithviraj Ammanabrolu, Liwei Jiang, Maarten Sap, Hannaneh Hajizhirzi, and Yejin Choi.
"Aligning to Social Norms and Values in Interactive Narratives."
North American Chapter of the Association of Computational Linguistics (NAACL), 9 pages, 2022.
75. David Wadden, Kyle Lo, Lucy Lu Wang, Arman Cohan, Iz Beltagy, Hannaneh Hajishirzi.
"MultiVerS: Improving scientific claim verification with weak supervision and full-document context."
Findings of North American Chapter of the Association of Computational Linguistics (NAACL), 9 pages, 2022.
76. Mitchell Wortsman, Gabriel Ilharco, Jong Wook Kim, Mike Li, Simon Kornblith, Rebecca Roelofs, Raphael Gontijo Lopes, Hannaneh Hajishirzi, Ali Farhadi, Hongseok Namkoong, Ludwig Schmidt.
"Robust fine-tuning of zero-shot models."
Computer Vision and Pattern Recognition (CVPR) 9 pages, 2022.
Best Paper Nomination
77. Swaroop Mishra, Daniel Khashabi, Chitta Baral, Hannaneh Hajishirzi
"Cross-Task Generalization via Natural Language Crowdsourcing Instructions."
Associations of Computational Linguistics (ACL), 2022.
78. Sewon Min, Mike Lewis, Hannaneh Hajishirzi, Luke Zettlemoyer
"Noisy Channel Language Model Prompting for Few-Shot Text Classification."
Associations of Computational Linguistics (ACL), 2022.
79. Jungsoo Park, Sewon Min, Jaewoo Jang, Luke Zettlemoyer, Hannaneh Hajishirzi
"FaVIQ: FAct Verification from Information-seeking Questions."
Associations of Computational Linguistics (ACL), 2022.
80. Swaroop Mishra, Daniel Khashabi, Chitta Baral, Yejin Choi, Hannaneh Hajishirzi
"Reframing Instructional Prompts to GPTk's Language."
Findings of Associations of Computational Linguistics (ACL), 2022.
81. Jiacheng Liu, Alisa Liu, Ximing Lu, Sean Welleck, Peter West, Ronan Le Bras, Yejin Choi, Hannaneh Hajishirzi
"Generated Knowledge Prompting for Commonsense Reasoning."
Associations of Computational Linguistics (ACL), 2022.

[2021]

82. Akari Asai, Xinyan Yu, Jungo Kasai, and Hannaneh Hajishirzi
"One Question Answering Model for Many Languages with Cross-lingual Dense Passage Retrieval."
Neural Information Processing (NeurIPS), 2021.
83. Sewon Min*, Kenton Lee, Ming-Wei Chang, Kristina Toutanova and Hannaneh Hajishirzi
"Joint Passage Ranking for Diverse Multi-Answer Retrieval."
Empirical Methods in Natural Language Processing (EMNLP), 9 pages, 2021.
84. Zeqiu Wu*, Bo-Ru Lu, Hannaneh Hajishirzi and Mari Ostendorf
"DIALKI: Knowledge Identification in Conversational Systems through Dialogue-Document Contextualization."
Empirical Methods in Natural Language Processing (EMNLP), 9 pages, 2021.
85. Christopher Clark, Jordi Salvador, Dustin Schwenk, Derrick Bonafilia, Mark Yatskar, Eric Kolve, Alvaro Herrasti, Jonghyun Choi, Sachin Mehta*, Sam Skjonsberg, Carissa Schoenick, Aaron Sarnat, Hannaneh Hajishirzi, Aniruddha Kembhavi, Oren Etzioni and Ali Farhadi
"Iconary: A Pictionary-based Game for Testing Multimodal Communication with Drawings and Text."
Empirical Methods in Natural Language Processing (EMNLP), 9 pages, 2021.
86. Daniel Khashabi, Amos Ng, Tushar Khot, Ashish Sabharwal, Hannaneh Hajishirzi and Chris Callison-Burch
"GooAQ: Open Question Answering with Diverse Answer Types."
Findings of Empirical Methods in Natural Language Processing (EMNLP), 9 pages, 2021.
87. Zeyu Liu, Yizhong Wang*, Jungo Kasai, Hannaneh Hajishirzi and Noah A. Smith
"Probing Across Time: What Does RoBERTa Know and When?."
Findings of Empirical Methods in Natural Language Processing (EMNLP), 9 pages, 2021.
88. Rahul Nadkarni, David Wadden*, Iz Beltagy, Noah Smith, Hannaneh Hajishirzi, Tom Hope
"Scientific Language Models for Biomedical Knowledge Base Completion: An Empirical Study."
Automated Knowledge Base Construction (AKBC), 2021.

89. Ikuya Yamada, Akari Asai*, Hannaneh Hajishirzi
"Efficient Passage Retrieval with Hashing for Open-domain Question Answering."
Associations of Computational Linguistics (ACL), short, p 979-986, 5 pages, 2021.

90. Bhargavi Paranjape*, Julian Michael, Marjan Ghazvininejad, Luke Zettlemoyer, Hannaneh Hajishirzi.
"Prompting Contrastive Explanations for Commonsense Reasoning Tasks."
Findings of Associations of Computational Linguistics (ACL), p4179-4192, 9 pages, 2021.

91. Akari Asai*, Jungo Kasai, Jonathan H. Clark, Kenton Lee, Eunsol Choi, Hannaneh Hajishirzi.
"XOR QA: Cross-lingual Open-Retrieval Question Answering."
North American Chapter of the Association of Computational Linguistics (NAACL), p547-564, 9 pages, 2021.

92. Gabriel Ilharco*, Rowan Zellers, Ali Farhadi, Hannaneh Hajishirzi.
"Probing Text Models for Common Ground with Visual Representations."
North American Chapter of the Association of Computational Linguistics (NAACL), p5367-5377, 9 pages, 2021.

93. Aida Amini*, Tom Hope, David Wadden, Madeleine van Zuylen, Eric Horvitz, Daniel S. Weld, Roy Schwartz, Hannaneh Hajishirzi.
"Extracting A Knowledge Base Of Mechanisms From COVID-19 Papers."
North American Chapter of the Association of Computational Linguistics (NAACL), p4489-4503, 9 pages, 2021.

94. Sachin Mehta*, Marjan Ghazvininejad, Srinivasan Iyer, Luke Zettlemoyer, Hannaneh Hajishirzi.
"DeLighT: Deep and Light-weight Transformer."
International Conference on Representation Learning (ICLR), 9 pages, 2021.

95. Alon Talmor, Ori Yoran, Amnon Catav, Dan Lahav, Yizhong Wang*, Akari Asai*, Gabriel Ilharco*, Hannaneh Hajishirzi, Jonathan Berant.
"MultiModalQA: Complex Question Answering over Text, Tables and Images."
International Conference on Representation Learning (ICLR), 9 pages, 2021.

96. Zeqiu Wu*, Michel Galley, Chris Brockett, Yizhe Zhang, Xiang Gao, Chris Quirk, Rik Koncel-Kedziorski, Jianfeng Gao, Hannaneh Hajishirzi, Mari Ostendorf, Bill Dolan. "A Controllable Model of Grounded Response Generation."
Advancement in Artificial Intelligence (AAAI), p14085-14093, 7 pages, 2021.

[2020]

97. Sachin Mehta*, Hannaneh Hajishirzi, Mohammad Rastegari."DiCENet: Dimension-wise convolutions for efficient networks."
IEEE Transactions on Pattern Analysis and Machine Intelligence, 9 pages, 2020.

98. Daniel Khashabi, Sewon Min*, Tushar Khot, Ashish Sabharwal, Oyvind Tafjord, Peter Clark, Hannaneh Hajishirzi.
"UnifiedQA: Crossing Format Boundaries With a Single QA System."
Findings of Empirical Methods in Natural Language Processing (EMNLP), p1896-1907, 9 pages, 2020.

99. Bhargavi Paranjape*, Mandar Joshi, John Thickstun, Hannaneh Hajishirzi, Luke Zettlemoyer.
"An Information Bottleneck Approach for Controlling Conciseness in Rationale Extraction."
Empirical Methods in Natural Language Processing (EMNLP), 9 pages, p1938-1952, 2020.

100. David Wadden*, Shanchuan Lin, Kyle Lo, Lucy Lu Wang, Madeleine van Zuylen, Arman Cohan, Hannaneh Hajishirzi.
"Fact or Fiction: Verifying Scientific Claims."
Empirical Methods in Natural Language Processing (EMNLP), p7534-7550, 9 pages, 2020.
Media Coverage: MIT Technology Review, etc.

101. Sewon Min*, Julian Michael, Hannaneh Hajishirzi, Luke Zettlemoyer.
"AmbigQA: Answering Ambiguous Open-domain Questions."
Empirical Methods in Natural Language Processing (EMNLP), p5783-5797, 9 pages, 2020.

102. Swabha Swayamdipta, Roy Schwartz, Nicholas Lourie, Yizhong Wang*, Hannaneh Hajishirzi, Noah A. Smith and Yejin Choi.
"Dataset Cartography: Mapping and Diagnosing Datasets with Training Dynamics."
Empirical Methods in Natural Language Processing (EMNLP), p9275-9293, 9 pages, 2020.

103. Jaemin Cho, Jiasen Li, Hannaneh Hajishirzi, Ani Kembhavi.
"X-LXMERT: Paint, Caption and Answer Questions with Multi-Modal Transformers."
Empirical Methods in Natural Language Processing (EMNLP), p8785-8805, 9 pages, 2020.
Media coverage: MIT Technology Review, etc.

104. James Ferguson*, Matt Gardner, Hannaneh Hajishirzi, Tushar Khot, Pradeep Dasigi. "IIRC: A Dataset of Incomplete Information Reading Comprehension Questions."
Empirical Methods in Natural Language Processing (EMNLP), 9 pages, 2020.

105. Matt Gardner, Yoav Artzi, Victoria Basmova, Jonathan Berant, Ben Bogin, Sihao Chen, Pradeep Dasigi, Dheeru Dua, Yanai Elazar, Ananth Gottumukkala, Nitish Gupta, Hanna Hajishirzi, Gabriel Ilharco*, Daniel Khashabi, Kevin Lin, Jiangming Liu, Nelson F. Liu, Phoebe Mulcaire, Qiang Ning, Sameer Singh, Noah A. Smith, Sanjay Subramanian, Reut Tsarfaty, Eric Wallace, Ally Zhang, Ben Zhou.
"Evaluating NLP models via contrast sets."
in Findings of Empirical Methods in Natural Language Processing (EMNLP), p1307-1323, 9 pages, 2020.

106. Sanjay Subramanian, Lucy Lu Wang, Ben Bogin, Sachin Mehta*, Madeleine van Zuylen, Sravanthi Parasa, Sameer Singh, Matt Gardner, Hannaneh Hajishirzi. "MedICaT: A Dataset of Medical Images, Captions, and Textual References."
Findings of Empirical Methods in Natural Language Processing (EMNLP), p2112-2120, short (5 pages), 2020.

107. Aida Amini*, Antoine Bosselut, Bhavana Dalvi Mishra, Yejin Choi, Hannaneh Hajishirzi.
"Procedural Reading Comprehension with Attribute-Aware Context Flow. "
Automated Knowledge Base Construction, 10 pages, 2020.
Best paper honorable mention

108. Colin Lockard*, Prashant Shiralkar, Xin Luna Dong, Hannaneh Hajishirzi.
"ZeroShotCeres: Zero-Shot Relation Extraction from Semi-Structured Webpages."
Association for Computational Linguistics (ACL), p8105-8117, 9 pages, 2020.

109. Akari Asai*, Hannaneh Hajishirzi.
"Logic-Guided Data Augmentation and Regularization for Consistent Question Answering."
Association for Computational Linguistics (ACL), p5642-5650, short (5 pages), 2020.

110. Jinhyuk Lee, Minjoon Seo*, Hannaneh Hajishirzi, Jaewoo Kang.
"Contextualized Sparse Representation with Rectified N-Gram Attention for Open-Domain Question Answering."
Association for Computational Linguistics (ACL), p912-919, short (5 pages), 2020.

111. Sarthak Jain^{+}, Madeleine van Zuylen, Hannaneh Hajishirzi, Iz Beltagy.
'SciREX: A Challenge Dataset for Document-Level Information Extraction."
Association for Computational Linguistics (ACL), p7506-7516, 9 pages, 2020.

112. Sachin Mehta*, Rik Koncel-Kedziorski*, Mohammad Rastegari, Hannaneh Hajishirzi.
"DeFINE: Deep Factorized Input Word Embeddings for Neural Sequence Modeling."
International Conference on Representation Learning (ICLR), 9 pages, 2020.

113. Akari Asai*, Kazuma Hashimoto, Hannaneh Hajishirzi, Richard Socher, and Caiming Xiong.
"Learning to Retrieve Reasoning Paths over Wikipedia Graph for Question Answering."
International Conference on Representation Learning (ICLR), 9 pages, 2020.

[2019]

114. Jaemin Cho, Minjoon Seo* and Hannaneh Hajishirzi.
"Mixture Content Selection for Diverse Sequence Generation."
Empirical Methods in Natural Language Processing (EMNLP), p3119-3129, 9 pages, 2019.

115. David Wadden*, Ulme Wennberg, Yi Luan* and Hannaneh Hajishirzi.
"Entity, Relation, and Event Extraction with Contextualized Span Representations."
Empirical Methods in Natural Language Processing (EMNLP), p5783-5788, short (5 pages) 2019.

116. Sewon Min*, Danqi Chen, Hannaneh Hajishirzi and Luke Zettlemoyer.
"A Discrete Hard EM Approach for Weakly Supervised Question Answering."
Empirical Methods in Natural Language Processing (EMNLP), p2851-2864, 9 pages, 2019.

117. Minjoon Seo*, Jinhyuk Lee, Tom Kwiatkowski, Ankur P. Parikh, Ali Farhadi, Hannaneh Hajishirzi.
"Real-Time Open-Domain Question Answering with Dense-Sparse Phrase Index."
Association for Computational Linguistics (ACL), p4430-4441, 9 pages, 2019.

118. Sewon Min*, Victor Zhong, Luke Zettlemoyer, Hannaneh Hajishirzi.
"Multi-hop Reading Comprehension through Question Decomposition and Rescoring."
Association for Computational Linguistics (ACL), p6097-6109, 9 pages, 2019.

119. Sewon Min*, Eric Wallace, Sameer Singh, Matt Gardner, Hannaneh Hajishirzi, Luke Zettlemoyer.
"Compositional Questions Do Not Necessitate Multi-hop Reasoning."

Association for Computational Linguistics (ACL), p4249-4257, short (5 pages), 2019.

120. Yi Luan*, Dave Wadden*, Luheng He, Mari Ostendorf, Hannaneh Hajishirzi.
"A General Framework for Information Extraction Using Dynamic Span Graphs."
North American Chapter of the Association for Computational Linguistics (NAACL), p3036-3046, 9 pages, 2019.
121. Rik Koncel-Kedziorski*, Dhanush Bekal*, Yi Luan*, Mirella Lapata, Hannaneh Hajishirzi.
"Text Generation from Knowledge Graphs with Graph Transformers."
North American Chapter of the Association for Computational Linguistics (NAACL), p2284-2293, 9 pages, 2019.
122. Aida Amini*, Saadia Gabriel, Yejin Choi, Hannaneh Hajishirzi.
"MathQA:Towards Interpretable Math Word Problem Solving with Operation-Based Formalisms."
North American Chapter of the Association for Computational Linguistics (NAACL), p2357-2367, 9 pages, 2019.
123. Sachin Mehta*, Mohammad Rastegari, Linda Shapiro, Hannaneh Hajishirzi.
"ESPNetv2: A Light-weight, Power Efficient, and General Purpose Convolutional Neural Network."
Computer Vision and Pattern Recognition (CVPR), p9190-9200, 8 pages, 2019.

[2018]

124. Yi Luan*, Luheng He, Mari Ostendorf and Hannaneh Hajishirzi.
"Multi-Task Identification of Entities, Relations, and Coreference for Scientific Knowledge Graph Construction."
Empirical Methods in Natural Language Processing (EMNLP), p3219-3232, 9 pages, 2018.
125. Sachin Mehta*, Rik Koncel-Kedziorski*, Mohammad Rastegari and Hannaneh Hajishirzi.
"Pyramidal Recurrent Unit for Language Modeling."
Empirical Methods in Natural Language Processing (EMNLP), p4620-4630, 9 pages, 2018.
126. Minjoon Seo*, Tom Kwiatkowski, Ankur Parikh, Ali Farhadi and Hannaneh Hajishirzi.
"Phrase-Indexed Question Answering:A New Challenge for Scalable Document Comprehension."
Empirical Methods in Natural Language Processing (EMNLP), p559-564, short (5 pages), 2018.
127. Sachin Mehta*, Mohammad Rastegari, Anat Caspi, Linda Shapiro, Hannaneh Hajishirzi.
"ESPNet: Efficient Spatial Pyramid of Dilated Convolutions for Semantic Segmentation."
European Conference in Computer Vision (ECCV), p561-580, 16 pages, 2018.
128. James Ferguson*, Colin Lockard*, Dan Weld, Hannaneh Hajishirzi.
"Semi-Supervised Event Extraction with Paraphrase Clusters."
North American Chapter of the Association for Computational Linguistics (NAACL), p359-364, short (5 pages), 2018.
129. Sewon Min*, Minjoon Seo*, Ali Farhadi, Hannaneh Hajishirzi.
"Neural Speed Reading via Skim-RNN."
International Conference on Representation Learning (ICLR), 9 pages, 2018.

[2017]

130. Yi Luan*, Mari Ostendorf, Hannaneh Hajishirzi.
"Scientific Information Extraction with Semi-supervised Neural Tagging."
Empirical Methods in Natural Language Processing (EMNLP), p2641-2651, 9 pages, 2017.
131. Sewon Min*, Minjoon Seo*, Hannaneh Hajishirzi.
"Question Answering through Transfer Learning from Large Fine-grained Supervision Data."
Associations for Computational Linguistics (ACL), p510-517, short (5 pages), 2017.
132. Minjoon Seo*, Aniruddha Kembhavi, Ali Farhadi, Hannaneh Hajishirzi.
"Bidirectional Attention Flow for Machine Comprehension."
International Conference on Representation Learning (ICLR), 9 pages, 2017.
133. Minjoon Seo*, Ali Farhadi, Hannaneh Hajishirzi.
"Query-Reduction Networks for Question Answering."
International Conference on Representation Learning (ICLR), 9 pages, 2017.
134. Aniruddha Kembhavi, Minjoon Seo*, Eric Klove, Dustin Schwenk, Hannaneh Hajishirzi, Ali Farhadi.
"Are You Smarter Than A Sixth Grader? Textbook Question Answering for Multimodal Machine Comprehension."
Computer Vision and Pattern Recognition (CVPR), p5376-5384, 8 pages, 2017.

[2016]

135. Rik Koncel-Kedziorski*, Ioannis Konstas, Luke Zettlemoyer, Hannaneh Hajishirzi.

"A Theme-Rewriting Approach for Generating Algebra Word Problems."
Empirical Methods in Natural Language Processing (EMNLP), p1617-1628, 9 pages, 2016.

136. Aniruddha Kembhavi, Mike Salvato, Eric Kolve, Minjoon Seo*, Hannaneh Hajishirzi, and Ali Farhadi. %hfill (86 citations)
"A Diagram is Worth a Dozen Images."
European Conference in Computer Vision (ECCV), p235-251, 15 pages, 2016

137. Antoine Bosselut, Jianfu Chen, David Warren, Hannaneh Hajishirzi and Yejin Choi.
"Learning Prototypical Event Structure from Photo Albums."
Associations of Computational Linguistics (ACL), 9 pages, 2016.

138. Yi Luan*, Yangfeng Ji, Hannaneh Hajishirzi and Boyang Li.
"Multiplicative Representations for Unsupervised Semantic Role Induction."
Association of Computational Linguistics (ACL), short (5 pages), 2016.

139. Rik Koncel-Kedziorski*, Subhro Roy, Aida Amini*, Nate Kushman, Hannaneh Hajishirzi."MAWPS: A Math Word Problem Repository."
North American Chapter of the Association for Computational Linguistics (NAACL), p1152-1157, short (5 pages), 2016.

140. Roozbeh Mottaghi, Hannaneh Hajishirzi, and Ali Farhadi.
"A Task-Oriented Approach for Cost-sensitive Recognition."
Computer Vision and Pattern Recognition (CVPR), p2203-2211, 8 pages, 2016

141. Rik Koncel-Kedziorski*, Hannaneh Hajishirzi, Ashish Sabharwal, Oren Etzioni, and Siena Dumas Ang*.
" Parsing Algebraic Word Problems into Equations."
Transactions of the Association for Computational Linguistics (TACL), p5858-597, 10 pages, 2016.

142. Hessam Bagherinezhad, Hannaneh Hajishirzi, Yejin Choi, and Ali Farhadi.
"Are Elephants Bigger than Butterflies? Reasoning about Sizes of Objects."
the Conference in Artificial Intelligence (AAAI), p3449-3456, 7 pages, 2016.

143. Victoria Zayats, Mari Ostendorf, Hannaneh Hajishirzi.
"Disfluency detection using a bidirectional LSTM."
Interspeech, 2523-2527, 4 pages, 2016.

[2015]

144. Minjoon Seo*, Hannaneh Hajishirzi, Ali Farhadi, and Oren Etzioni, Clint Malcolm*
"Solving Geometry Problems: Combining Text and Diagram Interpretations."
Empirical Methods in Natural Language Processing (EMNLP), p1466-1476, 9 pages, 2015
Media Coverage: New York Times, Washington Post, Geekwire, etc.

145. Aaron Jaech, Vicky Zayats, Hao Fang, Mari Ostendorf, Hannaneh Hajishirzi.
"Talking to the crowd: What do people react to in online discussions?"
Empirical Methods in Natural Language Processing (EMNLP), p2026-2031, short (5 pages), 2015.

146. Hamid Izadinia, Fereshteh Sadeghi, Santosh Divvala, Hannaneh Hajishirzi, Yejin Choi, and Ali Farhadi (17 citations)
and Ali Farhadi, "Segment-Phrase Table for Semantic Segmentation, Visual Entailment, and Paraphrasing."
International Conference in Computer Vision, (ICCV), p10-18, 8 pages, 2015

147. Mohammad Rastegari^{+}, Hannaneh Hajishirzi, and Ali Farhadi.
"Discriminative and Consistent Similarities in Instance-Level Multiple Instance Learning", Computer Vision and Pattern Recognition (CVPR), p740-748, 8 pages, 2015.

148. Ben Hixon, Peter Clark, Hannaneh Hajishirzi. 'Learning Knowledge Graphs for Question Answering through Conversational Dialog."
North American Chapter of the Association for Computational Linguistics (NAACL HLT), p851-861, 9 pages, 2015.

149. William Hwang*, Hannaneh Hajishirzi, Mari Ostendorf, and Wei Wu.
"Aligning Sentences from Standard Wikipedia to Simple Wikipedia."
North American Chapter of the Association for Computational Linguistics (NAACL HLT), p211-217, short (5 pages), 2015.

150. Victoria Zayats, Mari Ostendorf, Hannaneh Hajishirzi, "Unediting: Detecting Disfluencies Without Careful Transcripts."
North American Chapter of the Association for Computational Linguistics (NAACL HLT), short (5 pages), p1410-

1415, 2015.

[2014]

151. Rik Koncel-Kedziorsk*, Hannaneh Hajishirzi, and Ali Farhadi.
"Multi-resolution Language Grounding with Weak Supervision."
Empirical Methods in Natural Language Processing (EMNLP), p386-396, 9 pages, 2014.
152. Mohammad Javad Hosseini*, Hannaneh Hajishirzi, Oren Etzioni, and Nate Kushman.
"Learning to Solve Arithmetic Word Problems with Verb Categorization. "
Empirical Methods in Natural Language Processing (EMNLP), p523-533, 9 pages, 2014.
153. Victoria Zayats, Mari Ostendorf, Hannaneh Hajishirzi.
"Multi-domain Disfluency and Repair Detection."
Interspeech, p 2907-2911, 4 pages, 2014.
154. Minjoon Seo*, Hannaneh Hajishirzi, Ali Farhadi, and Oren Etzioni.
"Diagram Understanding in Geometry Questions."
Conference in Artificial Intelligence (AAAI), p2831-2838, 7 pages, 2014.

[2013]

155. Hannaneh Hajishirzi, Leila Zilles, Dan Weld, and Luke Zettlemoyer.
"Joint Coreference Resolution and Named-Entity Linking with Multi-pass Sieves."
Empirical Methods in Natural Language Processing (EMNLP), p289-299, 9 pages, 2013.
156. Iolanda Leite^{+}, Hannaneh Hajishirzi, Jill F. Lehman, and Sean Andrist.
"Managing Chaos: Models of Turn-Taking Models in Charachter-Multichild Interactions."
International Conference on Multimodal Interactions (ICMI), p43-50, 8 pages, 2013.

[2007-2012]

157. Hannaneh Hajishirzi, Mohammad Rastegari^{+}, and Jessica Hodgins.
"Semantic Understanding of Professional Soccer Commentaries. "
28th Conference on Uncertainty in Artificial Intelligence (UAI), p326-335, 9 pages, 2012.
158. Hannaneh Hajishirzi, Jill F. Lehman, and Jessica Hodgins.
"Using Group History to Identify Character Directed Utterances in Multi-child Interactions. "
13th SIGdial conference on Discourse and Dialogue (SigDial), p207-216, 8 pages, 2012.
Best paper award
159. Hannaneh Hajishirzi. "Action-centered Reasoning for Probabilistic Dynamic Systems. "
PhD Thesis, University of Illinois at Urbana-Champaign, 2011.
160. Hannaneh Hajishirzi, Julia Hockenmaier, Erik T. Mueller, and Eyal Amir.
"Reasoning about RoboCup Soccer Narratives. "
27th Conference on Uncertainty in Artificial Intelligence (UAI), p291-300, 8 pages, 2011.
161. Hannaneh Hajishirzi, Wen-tau Yih, and Aleksander Kolcz.
"Adaptive Near Duplicate Detection via Similarity Learning."
33rd Annual ACM Conference on Information Retrieval (SIGIR), p419-426, 8 pages, 2010.
162. Hannaneh Hajishirzi and Eyal Amir.
"Reasoning about Deterministic Actions with Probabilistic Prior, Application to Stochastic Filtering"
12th International Conference on Principles of Knowledge Representation and Reasoning (KR), 10 paes, 2010.
163. Hannaneh Hajishirzi, Afsaneh Shirazi, Jaesic Choi, and Eyal Amir.
"Greedy Algorithms for Sensing Decision."
21st International Joint Conference on Artificial Intelligence (IJCAI), p1908-1915, 7 pages, 2009.
164. Hannaneh Hajishirzi and Eyal Amir.
"Sampling First Order Logical Particles. "
24th conference on Uncertainty in Artificial Intelligence (UAI), p248-255, 8 pages, 2008.
165. Hannaneh Hajishirzi and Eyal Amir.
"Stochastic Filtering in a Probabilistic Action Model."
22nd National Conference on Artificial Intelligence (AAAI), p999-1006, 7 pages, 2007.

Preprint and Archival

1. Hamish Ivison, Yizhong Wang, Valentina Pyatkin, Nathan Lambert, Matthew Peters, Pradeep Dasigi, Joel Jang, David Wadden, Noah A Smith, Iz Beltagy, Hannaneh Hajishirzi
"Camels in a changing climate: Enhancing lm adaptation with Tülu 2."
arXiv preprint arXiv:2311.10702, 2023.
2. Jesse Dodge, Gabriel Ilharco*, Roy Schwartz, Ali Farhadi, Hannaneh Hajishirzi, Noah Smith.
"Fine-tuning pretrained language models: Weight initializations, data orders, and early stopping".
arXiv preprint arXiv:2002.06305, 8 pages, 2019.
3. Sewon Min, Danqi Chen, Luke Zettlemoyer, Hannaneh Hajishirzi. "Knowledge guided text retrieval and reading for open domain question answering".
arXiv preprint arXiv:1911.03868, 8 pages, 2019.
4. Matt Gardner, Jonathan Berant, Hannaneh Hajishirzi, Alon Talmor, Sewon Min.
"Question answering is a format; When is it useful?"
arXiv preprint arXiv:1909.11291, 8 pages, 2019.

Workshops and Symposiums

1. Matt Gardner, Jonathan Berant, Hannaneh Hajishirzi, Alon Talmor, Sewon Min*.
"On making reading comprehension more comprehensive"
2nd Workshop on Machine Reading for Question Answering, 2019.
2. Mark Hopkins, Ronan Le Bras, Christan Petrescu-Prahova, Gabriel Stanovsky, Hannaneh Hajishirzi, Rik Koncel-Kediorski*.
"Semeval-2019 task 10: math question answering."
13th International Workshop on Semantic Evaluation, 2019.
3. Yi Luan*, Mari Ostendorf, Hannaneh Hajishirzi.
"Neural relation extraction model with selectively incorporated concept embeddings"
12th International Workshop on Semantic Evaluation (SemEval), p 788-792, 2018.
4. Iolanda Leite, Hannaneh Hajishirzi, Jill F. Lehman, and Sean Andrist.
"Take or Wait? Learning Turn-Taking from Multiparty Data."
in Late breaking developments at Conference in Artificial Intelligence (AAAI), 2013.
5. Hannaneh Hajishirzi , Jill F. Lehman, Kenichi Kumatani, Leonid Segal, and Jessica Hodgins.
"Recognizing Character Directed Utterances in Multi-child Interactions."
in late breaking report section of Human Robot Interaction (HRI), 2012.
6. Hannaneh Hajishirzi, Erik T. Mueller, and Eyal Amir.
"Symbolic Probabilistic Reasoning about Narratives."
in AAAI Symposium on Logical Formalization on Commonsense Reasoning (Commonsense), 2011.
7. Hannaneh Hajishirzi, Julia Hockenmaier, Erik T. Mueller, and Eyal Amir.
"Symbolic Reasoning in Narratives."
in Workshop of Planning and Learning in ICAPS, 2011.
8. Hannaneh Hajishirzi and Eyal Amir.
"Reasoning Probabilistic Action Models. "
in 9th Bar-Ilan Symposium on the Foundations of AI, 2007.

MISCELLANEOUS

Patents

1. Wen-tau Yih, Christopher A. Meek, Hannaneh Hajishirzi.
Learning Element Weighting for Similarity Measures.
Patent No.328335.01, filed by Microsoft corporation, 2010.
2. Hannaneh Hajishirzi and Jill Lehman.
Addressee Identification of Speech in Small Groups of Children and Adults.
Patent No.0260339, filed by Disney Research, 2012.

Other Scholarly Activity

Invited Lectures and Seminars

1. Keynote speech at EMNLP, Nov. 2025.
2. Guest speaker at Laude's inaugural Ship Your Research Impact Summit, June 2025.
3. Guest speaker at Apple Workshop on Natural Language and Interactive Systems, April 2025.
4. Invited speaker at Oracle, March 2025.
5. Lecture at the MOOC class on Large Language Model Agents, Feb. 2025.
6. Keynote speaker at AI Agent Reasoning and Decision-Making Workshop 2025
7. Keynote speech at efficient Natural Language and Speech Processing (ENLSP) workshop at NeurIPS, Dec. 2024.
8. Keynote speech at self-supervised Learning Workshop at NeurIPS, Dec. 2024.
9. Distinguished lecture at Columbia University, Oct. 2024.
10. Grace Hopper distinguished Lecture at University of Pennsylvania, Oct. 2024.
11. Keynote Speech at COLM Conference, Oct. 2024.
12. Featured Keynote speech at PyTorch Conference, Sept. 2024.
13. Panelist at Compound AI Systems Workshop at DataBricks Summit, June 2024.
14. Panelist at Visionary AI Research at SnowFlake Summit, June 2024.
15. Keynote speech at AI Expo for National Competitiveness, DC, May 2024.
16. Keynote speech at Opportunities at the AI Research Frontier event at the White House, DC, May 2024.
17. Keynote, Data for Foundation Models at ICLR Conference, May 2024.
18. Keynote and Panelist at Mathematical and Empirical Understanding of Foundation Models, ICLR Conference, May 2024.
19. Featured Keynote speech, Open-Source Summit by Linux Foundation, April 2024.
20. Keynote speech, Open-Source Generative AI workshop (OSGAI), March 2024
21. Keynote speech, Northwest Database Society Annual Meeting, Jan 2024
22. Panelist at UW AI Symposium, Dec. 2023
23. Keynote speaker at CS Affiliates at UW CSE, Nov 2023
24. Colloquium speaker, CIIR talk series, University of Massachusetts at Amherst, Nov 2023
25. Keynote speech, Workshop of DialDoc at ACL, Jul 2023
26. Keynote speech, Workshop of Representation Learning for NLP (Rep4NLP) at ACL, Jul 2023
27. Colloquium speaker, University of Illinois at Urbana-Champaign, April 2023
28. Invited speaker, Open Data Science Conference, Nov 2022
29. Colloquium speaker, University of Pittsburgh, Nov. 2022
30. Colloquium speaker, University of Pennsylvania, Nov. 2022
31. Colloquium speaker, Stanford University, Oct. 2022
32. Keynote speech, Workshop of NeuroAI at UW, Sept. 2022
33. Summer School, Jul 2022
34. Keynote speech, Workshop of Learning with Natural Language Supervision at ACL, May 2022.
35. Keynote speech, Workshop on semi-parametric methods in NLP at ACL, May 2022.
36. Invited speaker, PNNL, March, 2022
37. Colloquium speaker, Johns Hopkins University, Feb 2022
38. Keynote speech, Workshop of structured and unstructured knowledge integration, NAACL, Jul 2022
39. Keynote speech, Workshop of Robustness in Sequential Data in CVPR, Jul 2022
40. Keynote speech, Workshop of Math AI for Education (MathAI4Ed) at NurIPS, Dec. 2021.
41. Keynote speach and Panel speaker, Workshop of Machine Reading and Question Answering at EMNLP, Nov. 2021
42. Colloquium speaker, University of Texas at Austin, Nov. 2021.
43. NLP Colloquium speaker, Carnegie Mellon University, Oct. 2021
44. Colloquium spearker, George Town Univeristy, Oct. 2021
45. Keynote speech, Workshop of Self-Supervised Learning and reasoning workshop at ICML, Jul 2021, *"Robust*

Knowledge-Rich Neural Text Comprehension and Reasoning"

46. Keynote speech, Workshop of Scholarly Document Processing Workshop at NAACL, June 2021, *"Knowledge Extraction from Unstructured Scientific Text"*
47. Intel's Rising Star Tech Talk, April 2021, *Knowledge-Rich Neural Text Comprehension and Reasoning*
48. Panel speaker, IE/IR TAC KBP, March 2021
49. Keynote speech, Workshop of Knowledge Discovery from Financial Documents at AAAI, Feb 2021, *"Knowledge Extraction from Unstructured Scientific Text"*
50. Panel speaker, Mega AI Panel (PNNL and Universities), March 2021
51. NLP seminar, University of Texas Austin, Dec 2020, *Knowledge-Rich Neural Text Comprehension and Reasoning*
52. NLP seminar, University of Columbia, Nov 2020, *Knowledge-Rich Neural Text Comprehension and Reasoning*
53. Keynote speech, Annual West Coast NLP Summit, Oct 2020, *"Efficient transformers for natural language processing"*
54. CS Colloquium, George Town University, Oct 2020, *Knowledge-Rich Neural Text Comprehension and Reasoning*
55. Keynote speaker, Women+Data Science seminar at University of Michigan, Sept 2020, *"Knowledge Extraction from Unstructured Scientific Text"*
56. Guest speaker, AI2, Sept 2020, *Knowledge-Rich Neural Text Comprehension and Reasoning*
57. Panel speaker, DB affiliates at University of Washington, July 2019
58. Keynote speech, Workshop of Adaptive and Multitask Learning: Algorithms & Systems at ICML, June 2019, *"Toward Robust AI Systems for Understanding and Reasoning Over Multimodal Data"*
59. Keynote speech, Workshop of Structured Prediction for NLP at NAACL, Jun 2019, *"Toward Robust AI Systems for Understanding and Reasoning over Multimodal Data"*
60. NLP seminar, New York University Data Science center, May 2019, *"Toward Robust AI Systems for Understanding and Reasoning over Multimodal Data"*
61. Panel speaker, UW Cybersecurity Artificial Intelligence Panel, Feb 2019.
62. Panel speaker, UW MSR summer workshop, Jul 2019
63. Guest speaker, Google, Jul 2018, *"Learning to reason in multiple modalities"*
64. Keynote speech, Workshop of UW DB Affiliates, Jul 2018, *"Learning to reason in multiple modalites"*
65. UW CSE colloquium, Jan 2018, *"Learning to reason in multiple modalities"*
66. Guest speaker, Big Data Workshop, Aug 2016, *"Learning to read, ground, and reason in multimodal text"*
67. UW AI seminars, April 2016, *"Reasoning-based Question Answering"*
68. Jelenik summer workshop (JSALT), June 2015, *"Grounded language learning with weak supervision"*
69. Guest speaker, Allen Institute of AI, Feb 2014, *"Diagram Understanding in Geometry Questions"*
70. UW MSR Summer Institute Invited Talk, July 2013, *"Language Grounding with Weak Supervision"*
71. UW MSR NLP colloquium, Feb 2013, *"Language Grounding with Weak Supervision"*
72. University of Washington, Yahoo! Machine Learning Seminar, Dec 2012
73. Disney Research, May 2011, *"Reasoning in Dynamic Systems: Theory and Application"*
74. AI seminar, March 2011, *"Reasoning in Dynamic Systems: Theory and Application"*
75. IBM Research Almaden, March 2011, *"Reasoning in Dynamic Systems"*
76. University of Zurich, Switzerland , Feb 2011, *"Reasoning in Dynamic Systems"*
77. Microsoft Research, May 2010, *"Logical Particle Filtering and Text Understanding"*

Professional Society Memberships

1. Member of Association for Computational Linguistics
2. Member of AAAI association for Artificial Intelligence

Other

Program Chair:

Automated Knowledge Base Construction (AKBC) (2020)

Board Member:

North American Chapter of ACL (NAACL) (2024-present)

Committee Member:

National Academy of Engineering Frontier of Engineering (2024)

Senior Area Chair:

Associations for Computational Linguistics (ACL) (2021, 2019, 2017)
North American Chapter of the Association for Computational Linguistics (NAACL) (2024, 2021, 2020)
Empirical Methods on Natural Language Processing (EMNLP) (2023)
ARR (2021-present)

Area Chair:

Conference on Language Modeling (COLM), 2025.
Empirical Methods on Natural Language Processing (EMNLP) (2021, 2020, 2019, 2017)
North American Chapter of the Association for Computational Linguistics (NAACL) (2012)
NeurIPS Datasets and Benchmark Track (2023, 2025)
Associations for Computational Linguistics (ACL) (2018)
AAAI Association of Artificial Intelligence (2020, 2016)

Editorial Board:

Journal of Artificial Intelligence Research (JAIR) (2021 — present)

Program Committee or Journal Reviewer:

1. Workshops: WiNLP, RoboNLP 2019, Embodied grounding 2017, Cognitum 2016, 2017.
2. Conferences: NeurIPS, ACL, NAACL, EMNLP, AAAI, ICML, UAI, IJCAI, CoNLL.

External reviewer for over 25+ conferences and journals.

Task and Workshop Organizer

Co-organizer, *"3rd Workshop on Knowledge-Augmented NLP "*, 2024
Co-organizer, *"5th Workshop on Representation Learning for NLP (RepL4NLP)"*, 2020
Task Co-organizer, *Semeval-2019 task 10: math question answering*
Workshop at IJCAI 2019, *Fourth Workshop on Declarative Learning Based Programming*

ADVISING

Alumni Postdocs

1. Sarah Wiegreffe, Summer 2025 ⟶ University of Maryland, Assistant Professor
2. Sachin Kumar, Summer 2024 ⟶ Ohio State University, Assistant Professor
3. Qingqing Cao, Fall 2024 ⟶ Apple Inc., Research Scientist
4. Prithviraj Ammanabrolou, Summer 2023 ⟶ University of California at San Diego (UCSD), Assistant Professor
5. Wenya Wang, Summer 2023, ⟶ Nanyang Technological University, Singapore, Assistant Professor
6. Hao Peng, Summer 2023 ⟶ University of Illinois at Urbana Champaign (UIUC), Assistant Professor
7. Daniel Khashabi, Summer 2022 ⟶ Johns Hopkins University, Assistant Professor
8. Rajarshi Das, Winter 2022⟶ Amazon Inc., Research Scientist

Chaired Doctoral Degrees

1. Akari Asai, ⟶ Assistant Professor at Carnegie Mellon University
2. Yizhong Wang ⟶ Assistant Professor at University of Texas Austin
3. Sewon Min, co-chair ⟶ Assistant Professor at University of California at Berkeley
4. Ellen Wu, co-chair, Spring 2024 ⟶ Member of Technical Staff at Microsoft AI
5. Bhargavi Paranjape, co-chair, Winter 2024 ⟶ Research Scientist at Meta AI
6. Gabriel Illharco, co-chair, Winter 2024⟶ Member of Technical Staff at XAI
7. Dave Wadden, chair, Summer 2022 ⟶ Research Scientist at AI2
8. James Ferguson, chair , Summer 2022 ⟶ Amazon Inc., Applied Research Scientist
9. Aida Amini , chair , Summer 2022 ⟶ Google Research Inc., Research Engineer
10. Colin Lockard, chair *"Information Extraction from Semi-Structured Websites"*, Spring 2021 ⟶ Amazon Inc., Research Scientist

11. Sachin Mehta, co-chair *"Efficient Deep Learning for Visual and Textual Data"*, Winter 2021 ⟶ Apple Inc., Research Scientist
12. Min Joon Seo, co-chair, *"Web-scale Neural Memory towards Universal Knowledge Interface"*, *recipient of 2019 Facebook Fellowship and 2018 AI2 Key Scientific Challenges Award*, Spring 2020 ⟶ Assistant Professor, Korean AI Institute
13. Yi Luan, co-chair, *"Multi-task graph-based information extraction with global context"*, *recipient of Yang Award for Outstanding Doctoral Student*, Spring 2019 ⟶ Google Research, Research Scientist
14. Rik Koncel-Kedziorski, co-chair, *"Understanding and Generating Multi-Sentence Texts"*, *recipient of Excellence in Linguistics Research Graduate Fellowship*, Spring 2019 ⟶ Amazon Inc., Research Scientist

Current PostDocs

1. Valentina Pyatkin
2. Varsha Kishore

Current Doctoral Students

1. Yizhong Wang, co-chair, 5th year, passed general exam
2. Akari Asai, chair, 5th year, passed qualifying exam
3. Jiacheng Liu, co-chair, 3rd year, passed qualifying exam
4. Mohammadreza Salehi, co-chair, 3rd year, passed qualifying exam
5. Tong Chen, co-chair, 2nd year
6. Joongwon Kim, chair, 2nd year
7. Hamish Ivison, 1st year

Chaired Masters Degrees

1. Dhanush Bekal Kannagola, chair, Coursework only, summer 2019 ⟶ Amazon, Software Engineer

Other Significant Student Supervision

1. *Interns at AI2 and Disney Research:* Sarthak Jain, Swaroop Mishra, Jaemin Cho, Iolanda Leite (now Associate Professor at KTH Sweden), Mohammad Rastegari (now at Apple Inc.)
2. *Thesis Committee:* Pengxiang Cheng (UT-Austin), Christopher Clark (CSE, UW), Junyuan Xie (CSE, UW), Jennifer Oritz (CSE, UW), Parmita Mehta (CSE, UW), Shrainik Jain (CSE, UW), Junyuan Xie (CSE, UW), Rafal Kocielnik (HCDE, UW), Hao Fang (ECE, UW), Trang Tan (ECE, UW), Vicky Zayats (ECE, UW), Chandrashekhar Lavania(ECE, UW), Farah Nadeem (ECE, UW)
3. *Past and current undergraduates:* Siena Dumas Ang, William Hwang (*winner of Mary Gates fellowship and UW Dean's medal for academics excellence*), Clint Malcolm, Amy Shah, Peter (Shanchuan) Li, Edan Sneh

TEACHING

Courses Taught

1. CSE 599, Advanced topics in Reasoning in Large Language Models [Spring 2025]
2. CSE 473, Introduction to AI [Fall 2023]
3. CSE 573, Artificial Intelligence [Winter 2022]
4. CSE 473, Introduction to AI [Spring 2021]
5. CSE 573, Artificial Intelligence [Winter 2021]
6. CSE 573, Artificial Intelligence [Winter 2020]
7. CSE 473, Introduction to AI [Fall 2019]
8. CSE 599, Advanced NLP [Spring 2019]
9. CSE 573, Artificial Intelligence [Winter 2019]
10. EE 511, Statistical Machine learning [Winter 2018]

Other Teaching Contributions

Conference Tutorials

1. Co-teach a tutorial at NAACL 2021, *"Beyond Paragraphs: NLP for Long Sequences"*
2. Co-teach a tutorial at ACL 2020, *"Multi-modal Information Extraction from Text, Semi-structured, and Tabular Data on the Web"*
3. Co-teach a tutorial at KDD 2020, *"Multi-modal Information Extraction from Text, Semi-structured, and Tabular Data on the Web"*
4. Co-teach a tutorial at WSDM 2020, *"Web-scale Knowledge Collection"*
5. Co-teach a tutorial at EMNLP 2018, *Standardized tests as benchmarks for Artificial Intelligence*

Professional Society and Other Service

1. Student Research Workshop Faculty advisor, ACL 2019.
2. Student mentoring sessions EMNLP, ACL, NAACL 2020, 2021
3. Serve as program chair, senior area chair, and area chair at multiple conferences
4. Serve as the best paper committee member at NAACL 2021.

International, National, or Governmental Service

1. **Panel reviewer for** NSF 2020, 2017, 2016, Department of Energy, 2016.