

Yizhong Wang

☎ (+1) 949-241-0294 ✉ yizhongw@cs.uw.edu 🏠 [yizhongw.github.io](https://github.com/yizhongw) 🌐 [/yizhongw](https://www.yizhongw.com/)

EDUCATION

University of Washington, Seattle, USA *Sep 2019 - Present*

Ph.D. student in Computer Science

Advisors: Hannaneh Hajishirzi, Noah Smith

Peking University, Beijing, China *Sep 2016 - July 2019*

M.S. in Computer Science

Advisor: Sujian Li

Shanghai Jiao Tong University, Shanghai, China *Sep 2012 - July 2016*

B.S. in Computer Science and Technology (IEEE Pilot Class)

Thesis advisor: Kenny Q. Zhu

EMPLOYMENT

University of Washington *Sep 2019 - Present*

Research Assistant, Supervisors: Hannaneh Hajishirzi, Noah Smith

Allen Institute for Artificial Intelligence *Jan 2023 - Present*

Part-time Research Intern, Collaborators: Daniel Khashabi, Pradeep Dasigi

Facebook AI Research *Sep 2021 - Dec 2021*

Research Intern, Mentor: Sebastian Riedel

Allen Institute for Artificial Intelligence *Dec 2020 - May 2021*

Part-time Research Intern, Collaborators: Tushar Khot, Ashish Sabharwal

Institute of Computational Linguistics, Peking University *Sep 2016 - July 2019*

Research Assistant, Supervisor: Sujian Li

Allen Institute for Artificial Intelligence *Oct 2018 - Feb 2019*

Research Intern, Mentors: Matt Gardner, Sameer Singh

Microsoft Research Asia *May 2018 - Sep 2018*

Research Intern, Mentor: Furu Wei

Baidu NLP *June 2017 - Feb 2018*

Research Intern, Mentors: Kai Liu, Yajuan Lyu

HONORS AND AWARDS

Geekwire Innovation of the Year 2024 (for OLMo [28])

Best Paper Award of CCL 2020 (for LiveQA [11])

Outstanding Paper Award of ACL 2017 (for staged discourse parser [3])

Founder Scholarship, Peking University, 2017-2018

Chun-Tsung Scholarship (by Nobel laureate Tsung-Dao Lee), Shanghai Jiao Tong Univ., 2014-2016

PUBLICATIONS

Google Scholar: <https://scholar.google.com/citations?user=y5zpqdAAAAAJ>

Semantic Scholar: <https://www.semanticscholar.org/author/Yizhong-Wang/1705260>

* denotes equal contributions

- [33] Unpacking DPO and PPO: Disentangling Best Practices for Learning from Preference Feedback
Hamish Ivison, Yizhong Wang, Jiacheng Liu, Zeqiu Wu, Valentina Pyatkin, Nathan Lambert, Noah A Smith, Yejin Choi, Hannaneh Hajishirzi
arXiv, 2024
- [32] Retrieval Head Mechanistically Explains Long-Context Factuality
Wenhao Wu, Yizhong Wang, Guangxuan Xiao, Hao Peng, Yao Fu
arXiv, 2024
- [31] Tur[k]ingBench: A Challenge Benchmark for Web Agents
Kevin Xu, Yeganeh Kordi, Kate Sanders, Yizhong Wang, Adam Byerly, Jack Zhang, Benjamin Van Durme, Daniel Khashabi
arXiv, 2024
- [30] Can Language Models Act as Knowledge Bases at Scale?
Qiyuan He, Yizhong Wang, Wenya Wang
arXiv, 2024
- [29] Set the Clock: Temporal Alignment of Pretrained Language Models
Bowen Zhao*, Zander Brumbaugh*, Yizhong Wang*, Hannaneh Hajishirzi, Noah A. Smith
ACL Findings, 2024
- [28] OLMo: Accelerating the Science of Language Models
Dirk Groeneveld, Iz Beltagy, Pete Walsh, Akshita Bhagia, Rodney Kinney, Oyvind Tafjord, Ananya Harsh Jha, Hamish Ivison, Ian Magnusson, Yizhong Wang, and many others
ACL, 2024
- [27] Tuning Language Models by Proxy
Alisa Liu, Xiaochuang Han, Yizhong Wang, Yulia Tsvetkov, Yejin Choi, Noah A Smith
CoLM, 2024
- [26] Fine-grained Hallucination Detection and Editing for Language Models
Abhika Mishra, Akari Asai, Vidhisha Balachandran, Yizhong Wang, Graham Neubig, Yulia Tsvetkov, Hannaneh Hajishirzi
CoLM, 2024
- [25] Self-RAG: Learning to Retrieve, Generate, and Critique through Self-Reflection
Akari Asai, Zeqiu Wu, Yizhong Wang, Avirup Sil, Hannaneh Hajishirzi
ICLR, 2024
- [24] BTR: Binary Token Representations for Efficient Retrieval Augmented Language Models
Qingqing Cao, Sewon Min, Yizhong Wang, Hannaneh Hajishirzi
ICLR, 2024
- [23] Camels in a Changing Climate: Enhancing LM Adaptation with Tulu 2
Hamish Ivison*, Yizhong Wang*, Valentina Pyatkin, Nathan Lambert, Matthew Peters, Pradeep Dasigi, Joel Jang, David Wadden, Noah A Smith, Iz Beltagy, Hannaneh Hajishirzi
arXiv, 2023

- [22] Personalized Soups: Personalized Large Language Model Alignment via Post-hoc Parameter Merging
Joel Jang, Seungone Kim, Bill Yuchen Lin, Yizhong Wang, Jack Hessel, Luke Zettlemoyer, Hannaneh Hajishirzi, Yejin Choi, Prithviraj Ammanabrolu
arXiv, 2023
- [21] How Far Can Camels Go? Exploring the State of Instruction Tuning on Open Resources
Yizhong Wang*, Hamish Ivison*, Pradeep Dasigi, Jack Hessel, Tushar Khot, Khyathi Raghavi Chandu, David Wadden, Kelsey MacMillan, Noah A. Smith, Iz Beltagy, Hannaneh Hajishirzi
NeurIPS (Datasets and Benchmarks Track), 2023 (**Spotlight**)
- [20] TIFA: Accurate and interpretable text-to-image faithfulness evaluation with question answering
Yushi Hu, Benlin Liu, Jungo Kasai, Yizhong Wang, Mari Ostendorf, Ranjay Krishna, Noah A. Smith
ICCV, 2023
- [19] Self-Instruct: Aligning Language Models with Self-Generated Instructions
Yizhong Wang, Yeganeh Kordi, Swaroop Mishra, Alisa Liu, Noah A. Smith, Daniel Khashabi, Hannaneh Hajishirzi
ACL, 2023 (**Most influential paper top 1 recognized by Paper Digest**)
- [18] HINT: Hypernetwork Instruction Tuning for Efficient Zero-Shot Generalisation
Hamish Ivison, Akshita Bhagia, Yizhong Wang, Hannaneh Hajishirzi, Matthew Peters
ACL, 2023
- [17] One Embedder, Any Task: Instruction-Finetuned Text Embeddings
Hongjin Su, Weijia Shi, Jungo Kasai, Yizhong Wang, Yushi Hu, Mari Ostendorf, Wen-tau Yih, Noah A. Smith, Luke Zettlemoyer, Tao Yu
ACL, 2023
- [16] Super-NaturalInstructions: Generalization via Declarative Instructions on 1600+ NLP Tasks
Yizhong Wang*, Swaroop Mishra*, Pegah Alipoormolabashi, Yeganeh Kordi, . . . , Chitta Baral, Yejin Choi, Hannaneh Hajishirzi, Noah A. Smith, Daniel Khashabi
EMNLP, 2022 (**Most influential paper top 2 recognized by Paper Digest**)
- [15] The Web Is Your Oyster - Knowledge-Intensive NLP against a Very Large Web Corpus
Aleksandra Piktus, Fabio Petroni, Yizhong Wang, Vladimir Karpukhin, Dmytro Okhonko, Samuel Broscheit, Gautier Izacard, Patrick Lewis, Barlas Oğuz, Edouard Grave, Wen-tau Yih, Sebastian Riedel
EMNLP, 2022
- [14] Probing Across Time: What Does RoBERTa Know and When?
Leo Z. Liu*, Yizhong Wang*, Jungo Kasai, Hannaneh Hajishirzi, Noah A. Smith
EMNLP Findings, 2021
- [13] MultiModalQA: Complex Question Answering over Text, Tables and Images
Alon Talmor, Ori Yoran, Amnon Catav, Dan Lahav, Yizhong Wang, Akari Asai, Gabriel Ilharco, Hannaneh Hajishirzi, Jonathan Berant
ICLR, 2021
- [12] Automated Lay Language Summarization of Biomedical Scientific Reviews
Yue Guo, Wei Qiu, Yizhong Wang, Trevor Cohen
AAAI, 2021
- [11] LiveQA: A Question Answering Dataset over Sports Live
Qianying Liu, Sicong Jiang, Yizhong Wang, Sujian Li
CCL, 2020 (**Best Paper Award**)

- [10] Dataset Cartography: Mapping and Diagnosing Datasets with Training Dynamics
Swabha Swayamdipta, Roy Schwartz, Nicholas Lourie, Yizhong Wang, Hannaneh Hajishirzi,
Noah A. Smith and Yejin Choi
EMNLP, 2020
- [9] Do Neural NLP Models Know Numbers? Probing Numeracy in Embeddings
Eric Wallace*, Yizhong Wang*, Sujian Li, Sameer Singh and Matt Gardner
EMNLP-IJCNLP, 2019
- [8] DROP: A Reading Comprehension Benchmark Requiring Discrete Reasoning Over Paragraphs
Dheeru Dua, Yizhong Wang, Pradeep Dasigi, Gabriel Stanovsky, Sameer Singh, Matt Gardner
NAACL, 2019
- [7] Toward Fast and Accurate Neural Discourse Segmentation
Yizhong Wang, Sujian Li
EMNLP, 2018
- [6] Multi-Passage Machine Reading Comprehension with Cross-Passage Answer Verification
Yizhong Wang, Kai Liu, Jing Liu, Wei He, Yajuan Lyu, Hua Wu, Sujian Li, Haifeng Wang
ACL, 2018
- [5] Bag-of-Words as Target for Neural Machine Translation
Shuming Ma, Xu Sun, Yizhong Wang, Junyang Lin
ACL, 2018
- [4] DuReader: a Chinese Machine Reading Comprehension Dataset from Real-world Applications
Wei He, Kai Liu, Jing Liu, Yajuan Lyu, Shiqi Zhao, Xinyan Xiao, Yuan Liu, Yizhong Wang, Hua Wu,
Qiaoqiao She, Xuan Liu, Tian Wu, Haifeng Wang
ACL Workshop on Machine Reading for Question Answering, 2018
- [3] A Two-stage Parsing Method for Text-level Discourse Analysis
Yizhong Wang, Sujian Li, Houfeng Wang
ACL, 2017 (**Outstanding Paper Award**)
- [2] Tag-Enhanced Tree-Structured Neural Networks for Implicit Discourse Relation Classification
Yizhong Wang, Sujian Li, Jingfeng Yang, Xu Sun and Houfeng Wang
IJCNLP, 2017
- [1] Towards Non-projective High-Order Dependency Parser
Wenjing Fang, Kenny Q. Zhu, Yizhong Wang, Jia Tan.
COLING 2016, System Demonstration

PROFESSIONAL SERVICE

Workshop organization:

- Organizer of NeurIPS 2023 Workshop on Instruction Tuning and Instruction Following
- Co-chair of ACL 2020 Student Research Workshop

Area Chair:

- NLP: EMNLP 2024

Reviewer / Program Committee:

- NLP: ACL 2019-24, EMNLP 2019-24, NAACL 2020-24, COLING 2022, AACL 2020

- ML/AI: ICLR 2022-24, NeurIPS 2021-24, ICML 2023-24, CoLM 2024, AAAI 2020-21, AKBC 2020-22

Departmental service:

- Co-organizer of UWNLP Seminar Series 2023-24
- UW CSE PhD admission committee member 2021-23
- Co-organizer of UWNLP Retreat 2022

TEACHING EXPERIENCE

Introduction to Artificial Intelligence, University of Washington Autumn 2023
Head Teaching Assistant, Instructor: Hannaneh Hajishirzi

Natural Language Processing Capstone, University of Washington Spring 2022
Teaching Assistant, Instructor: Noah Smith

Introduction to Artificial Intelligence, Peking University Summer 2018
Teaching Assistant, Instructor: Vincent Ng

Discrete Math, Peking University Fall 2016 / Fall 2017
Teaching Assistant, Instructor: Sujian Li

STUDENT MENTORING

- Mickel Liu, Visiting student@UW→PhD student@UW, 01/2024–06/2024
- Bowen Zhao, MS student@UW, 07/2023–Present (Publication: [29])
- Zander Brumbaugh, BS/MS student@UW, 05/2023–03/2024 (Publication: [29])
- Abhika Mishra, BS/MS student@UW, 01/2023–06/2024 (Publication: [26])
- Hamish Ivison, PYI@AI2→PhD student@UW, 04/2022–09/2023 (Publications: [18, 21])
- Yeganeh Kordi, BS@Tehran Polytechnic→PhD student@Brown, 06/2022–09/2023 (Publications: [16, 19])
- Zeyu Liu, BS/MS@UW→AI resident@Meta→PhD student@UT Austin, 12/2019–04/2021 (Publication: [14])