

MICHAEL J.Q. ZHANG

michaelzhang@nyu.edu | <https://mikejqzhang.github.io> | 781-710-4239

EDUCATION

New York University , New York, NY Ph.D. in Computer Science <i>Advisor: Eunsol Choi</i>	<i>2024 - Spring 2026</i>
University of Texas at Austin , Austin, TX Ph.D. in Computer Science (<i>Transferred to NYU</i>) <i>Advisor: Eunsol Choi</i>	<i>2020 - 2024</i>
University of Washington , Seattle, WA B.S. in Computer Science and Engineering <i>Advisor: Noah A. Smith</i>	<i>2016 - 2020</i>

INDUSTRY EXPERIENCE

Allen Institute for AI (AI2) , Seattle, WA <i>Research Intern</i> , with Valentina Pyatkin	<i>June 2024 - September 2024</i>
Google Research , San Francisco, CA <i>Student Researcher</i> , with Jeremy R. Cole and Bhuwan Dhingra	<i>June 2022 - December 2022</i>
Microsoft Dynamics 365 AI Research , Seattle, WA <i>Research Intern</i> , with Yu Cheng, Zhe Gan, and Jingjing Liu	<i>June 2020 - August 2020</i>
Unbabel , Lisbon, Portugal <i>Research Intern</i> , with André F.T. Martins	<i>June 2019 - September 2019</i>
Legalist , San Francisco, CA <i>Data Scientist</i> , Y Combinator S16	<i>June 2016 - August 2017</i>

PUBLICATIONS

- [1] User Feedback in Human-LLM Dialogues: A Lens to Understand Users But Noisy as a Learning Signal
Yuhan Liu, **Michael J.Q. Zhang**, Eunsol Choi
EMNLP 2025
- [2] Improving LLM-as-a-Judge Inference with the Judgment Distribution
Victor Wang, **Michael J.Q. Zhang**, Eunsol Choi
EMNLP 2025 Findings
- [3] Diverging Preferences: When do Annotators Disagree and do Models Know?
Michael J.Q. Zhang, Zhilin Wang, Jena D. Hwang, Yi Dong, Olivier Delalleau, Yejin Choi, Eunsol Choi, Xiang Ren, and Valentina Pyatkin.
ICML 2025
- [4] Modeling Future Conversation Turns to Teach LLMs to Ask Clarifying Questions
Michael J.Q. Zhang, W. Bradley Knox, and Eunsol Choi.
ICLR 2025
- [5] Clarify When Necessary: Resolving Ambiguity Through Interaction with LMs
Michael J.Q. Zhang and Eunsol Choi
NAACL Findings 2025
- [6] Selectively Answering Ambiguous Questions
Jeremy R. Cole, **Michael J.Q. Zhang**, Daniel Gillick, Julian Martin Eisenschlos, Bhuwan Dhingra, and Jacob Eisenstein
EMNLP 2023
- [7] Mitigating Temporal Misalignment by Discarding Outdated Facts
Michael J.Q. Zhang and Eunsol Choi
EMNLP 2023

- [8] Propagating Knowledge Updates to LMs Through Distillation
Shankar Padmanabhan, Yasumasa Onoe, **Michael J.Q. Zhang**, Greg Durrett, and Eunsol Choi
NeurIPS 2023
- [9] Can LMs Learn New Entities from Descriptions? Challenges in Propagating Injected Knowledge
Yasumasa Onoe, **Michael J.Q. Zhang**, Shankar Padmanabhan, Greg Durrett, and Eunsol Choi
ACL 2023
- [10] DIFFQG: Generating Questions to Summarize Factual Changes
Jeremy R. Cole, Palak Jain, Julian Martin Eisenschlos, **Michael J.Q. Zhang**, Eunsol Choi, and Bhuwan Dhingra
EACL 2023
- [11] Rich Knowledge Sources Bring Complex Knowledge Conflicts in QA: Recalibrating Models to Reflect Conflicting Evidences
Hungting Chen, **Michael J.Q. Zhang**, and Eunsol Choi
EMNLP 2022
- [12] Entity Cloze by Date: What LMs Know About Unseen Entities
Yasumasa Onoe, **Michael J.Q. Zhang**, Eunsol Choi, and Greg Durrett
NAACL Findings 2022
- [13] CREAK: A Dataset for Commonsense Reasoning over Entity Knowledge
Yasumasa Onoe, **Michael J.Q. Zhang**, Eunsol Choi, and Greg Durrett
NeurIPS Datasets and Benchmarks 2021
- [14] SituatedQA: Incorporating Extra-Linguistic Contexts into QA
Michael J.Q. Zhang and Eunsol Choi
EMNLP 2021, Outstanding Paper Award
- [15] Document-level Neural MT: A Systematic Comparison
António V. Lopes, M. Amin Farajian, Rachel Bawden, **Michael Zhang**, and André F. T. Martins
EAMT 2020
- [16] Deep Weighted Averaging Classifiers
Dallas Card, **Michael Zhang**, and Noah A. Smith.
ACM FAccT 2019

TEACHING ASSISTANTSHIPS

CS 3033: Emerging Topics in Natural Language Processing , New York University Instructor: Professor Eunsol Choi	<i>Spring 2025</i>
CS 395T: Topics in Natural Language Processing , University of Texas at Austin Instructor: Professor Eunsol Choi	<i>Autumn 2020</i>
CSE 447: Natural Language Processing , University of Washington Instructor: Professor Yejin Choi	<i>Winter 2020</i>
CSE 446: Machine Learning , University of Washington Instructor: Professor Sewoong Oh	<i>Autumn 2019</i>
CSE 446: Machine Learning , University of Washington Instructors: Professor Kevin Jamieson and Professor Anna Karlin	<i>Spring 2019</i>
CSE 351: The Hardware / Software Interface , University of Washington Instructor: Professor Justin Hsia	<i>Autumn 2016</i>
CSE 351: The Hardware / Software Interface , University of Washington Instructor: Professor Justin Hsia	<i>Summer 2016</i>

SERVICE

Reviewer ARR (EMNLP, ACL, NAACL, etc.), NeurIPS, ICLR, ICML
Co-Organizer Workshop on Towards Knowledgeable Language Models at EMNLP 2024
Instructor High school diversity and outreach at UT Computer Science (Academy for All)

AWARDS

Outstanding Paper Award at EMNLP, 2021
 NAACL Scholarship for JSALT Workshop, 2018