




Allen School of Computer Science & Engineering  
University of Washington  
3800 E Stevens Way NE, CSE2 233  
Seattle, WA, USA 98195

(760) 689-2466   
[@cquanze](https://twitter.com/cquanze)   
[cqz@cs.washington.edu](mailto:cqz@cs.washington.edu)   
[homes.cs.washington.edu/~cqz](https://homes.cs.washington.edu/~cqz) 

# QUAN ZE CHEN

## APPOINTMENTS

2023–2025 **University of Washington**, Seattle, WA  
Postdoctoral Scholar, Computer Science & Engineering  
Advisors: Amy X. Zhang, Franziska Roesner

## EDUCATION

2018–2023 **University of Washington**, Seattle, WA  
Ph.D. Computer Science & Engineering  
Advisor: Amy X. Zhang

2016–2018 **University of Washington**, Seattle, WA  
M.S. Computer Science & Engineering

2012–2016 **University of Pennsylvania**, Philadelphia, PA  
B.S.E. Networked and Social Systems Engineering, Summa Cum Laude

## PUBLICATIONS

### Conference and Journal Papers

- [c.17] **Quan Ze Chen**, Amy X. Zhang. Case Law Grounding: Using Precedents to Align Decision-Making for Humans and AI. *ACM Conference On Collective Intelligence (CI '25)*. La Jolla, CA. August 2025.
- [c.16] **Quan Ze Chen**, K.J. Kevin Feng, Chan Young Park, Amy X. Zhang. SPICA: Retrieving Scenarios for Pluralistic In-Context Alignment. *ACL Findings 2025*. Vienna, Austria. July 2025.
- [c.15] Aviv Ovadya, Kyle Redman, Oliver Smith, Luke Thorburn, **Quan Ze Chen**, Flynn Devine, Andrew Konya, Smitha Milli, Manon Revel, K. J. Kevin Feng, Amy X. Zhang, Bilva Chandra, Michiel A. Bakker, Atoosa Kasirzadeh. Position: Democratic AI is Possible. The Democracy Levels Framework Shows How It Could Work. *The 42nd International Conference on Machine Learning*. Vancouver, Canada. July 2025.
- [c.14] Tzu-Sheng Kuo, **Quan Ze Chen**, Amy X. Zhang, Jane Hsieh, Haiyi Zhu, Kenneth Holstein. PolicyCraft: Supporting Collaborative and Participatory Policy Design through Case-Grounded Deliberation. *ACM Conference on Human Factors in Computing Systems (CHI '25)*. Yokohama, Japan. April 2025.
- [c.13] Leijie Wang, Kathryn Yurechko, Pranati Dani, **Quan Ze Chen**, Amy X. Zhang. End User Authoring of Personalized Content Classifiers: Comparing Example Labeling, Rule Writing, and LLM Prompting. *ACM Conference on Human Factors in Computing Systems (CHI '25)*. Yokohama, Japan. April 2025.
- [c.12] Inyoung Cheong, King Xia, K.J. Kevin Feng, **Quan Ze Chen**, Amy X. Zhang. (A)I Am Not a Lawyer, But...: Engaging Legal Experts towards Responsible LLM Policies for Legal Advice. *ACM Conference on Fairness, Accountability, and Transparency (FAccT '24)*. Rio de Janeiro, Brazil. June 2024.

- [c.11] Kianna Bolante, Kevin Chen, **Quan Ze Chen**, Amy X. Zhang. Bringing Social Computing to Secondary School Classrooms. *ACM Special Interest Group on Computer Science Education Technical Symposium (SIGCSE TS 2024)*. Portland, OR. March 2024.
- [c.10] Andre Ye, **Quan Ze Chen**, Amy X. Zhang. Confidence Contours: Uncertainty-Aware Annotation for Medical Semantic Segmentation. *AAAI Conference on Human Computation and Crowdsourcing (HCOMP '23)*. Delft, Netherlands. November 2023. **Honorable Mention**
- [c.9] **Quan Ze Chen**, Amy X. Zhang. Judgment Sieve: Reducing Uncertainty in Group Judgments through Interventions Targeting Ambiguity versus Disagreement. *ACM Conference On Computer-Supported Cooperative Work And Social Computing (CSCW '23)*. Minneapolis, MN. October 2023.
- [c.8] Shagun Jhaver, Alice Zhang, **Quan Ze Chen**, Detlef Knauss, Amy X. Zhang. Understanding User Preferences for Personal Content Moderation Controls. *ACM Conference On Computer-Supported Cooperative Work And Social Computing (CSCW '23)*. Minneapolis, MN. October 2023.
- [c.7] Teanna Barrett, **Quan Ze Chen**, Amy X. Zhang. Skin Deep: Investigating Subjectivity in Skin Tone Annotations for Computer Vision Benchmark Datasets. *ACM Conference on Fairness Accountability and Transparency (FAccT '23)*. Chicago, IL. June 2023.
- [c.6] Shagun Jhaver, **Quan Ze Chen**, Detlef Knauss, Amy X. Zhang. Designing Word Filter Tools for Creator-led Comment Moderation. *ACM Conference on Human Factors in Computing Systems (CHI '22)*. New Orleans, LA. April 2022.
- [c.5] **Quan Ze Chen**, Tobias Schnabel, Besmira Nushi, Saleema Amershi. HINT: Integration Testing for AI-based features with Humans in the Loop. *ACM Conference on Intelligent User Interfaces (IUI '22)*. Virtual. March 2022.
- [c.4] **Quan Ze Chen**, Daniel S. Weld, Amy X. Zhang. Goldilocks: Consistent Crowdsourced Scalar Annotations with Relative Uncertainty. *Proceedings of the ACM on Human-Computer Interaction (CSCW '21)*. Virtual. October 2021.
- [c.3] **Quan Ze Chen**, Jonathan Bragg, Lydia B. Chilton, Daniel S. Weld. Cicero: Multi-Turn, Contextual Argumentation for Accurate Crowdsourcing. *ACM Conference on Human Factors in Computing Systems (CHI '19)*. Glasgow, United Kingdom. May 2019.
- [c.1] Wei Xu, Courtney Napoles, Ellie Pavlick, **Quan Ze Chen**, Chris Callison-Burch. Optimizing Statistical Machine Translation for Simplification. *Transactions of the Association for Computational Linguistics (ACL '16)*. Berlin, Germany. August 2016.

### Posters and Work-in-Progress Papers

- [p.6] Aviv Ovadya, Kyle Redman, Oliver Smith, Luke Thorburn, **Quan Ze Chen**, Flynn Devine, Andrew Konya, Smitha Milli, Manon Revel, K. J. Kevin Feng, Amy X. Zhang, Bilva Chandra, Michiel A. Bakker, Atoosa Kasirzadeh. The Democracy Levels Framework: Towards Integrating Democratic Infrastructure into the AI Ecosystem. *First Workshop on Sociotechnical AI Governance: Opportunities and Challenges for HCI (STAIG Workshop at CHI '25)*. Yokohama, Japan. Apr 2025.
- [p.5] K. J. Kevin Feng, **Quan Ze Chen**, Inyoung Cheong, Amy X. Zhang. Policy Prototyping for LLMs: Pluralistic Alignment via Interactive and Collaborative Policymaking. *Workshop on Bidirectional Human-AI Alignment (Bi-Align Workshop at ICLR '25)*. Singapore. Apr 2025.
- [p.4] Andrew Konya, Aviv Ovadya, Kevin Feng, **Quan Ze Chen**, Lisa Schirch, Colin Irwin, Amy X Zhang. Chain of Alignment: Integrating Public Will with Expert Intelligence for Language Model Alignment. *Pluralistic Alignment, Conference on Neural Information Processing Systems (Plural-*

*istic Alignment Workshop at NeurIPS '23*). Vancouver, Canada. Dec 2024.

- [p.3] K. J. Kevin Feng, **Quan Ze Chen**, Inyoung Cheong, King Xia, Amy X. Zhang. Case Repositories: Towards Case-Based Reasoning for AI Alignment. *First Workshop on AI meets Moral Philosophy and Moral Psychology, Conference on Neural Information Processing Systems (MP2 Workshop at NeurIPS '23)*. New Orleans, LA. Dec 2023.
- [p.2] John Akers, Gagan Bansal, Gabriel Cadamuro, Christine Chen, **Quan Ze Chen**, Lucy Lin, Phoebe Mulcaire, Rajalakshmi Nandakumar, Matthew Rockett, Lucy Simko, John Toman, Tongshuang Wu, Eric Zeng, Bill Zorn, Franziska Roesner. Technology-enabled disinformation: Summary, lessons, and recommendations. *arXiv 2018*. Digital Preprint.
- [p.1] **Quan Ze Chen**, Chenyang Lei, Wei Xu, Ellie Pavlick, Chris Callison-Burch. Poetry of the Crowd: A Human Computation Algorithm to Convert Prose into Rhyming Verse *AAAI Conference on Human Computation and Crowdsourcing (HCOMP '14)*. Pittsburgh, PA. Sept 2014.

### Doctoral Symposium

- [ds.1] **Quan Ze Chen**. Disentangling Inherent Ambiguity and Disagreement in Crowdsourced Annotation. *AAAI Conference on Human Computation and Crowdsourcing (HCOMP '21)*. Virtual. Nov 2021.

### INDUSTRY EXPERIENCE AND INTERNSHIPS

- Summer 2021 Applied Scientist Intern, **SageMaker Ground Truth, Amazon AI**, Palo Alto, CA  
*Mentors: Mike Schaekermann, Matt Lease, Erran Li*
- Spring 2020 Research Intern, **Human-AI eXperiences (HAX) Team, Microsoft Research**, Redmond, WA  
*Mentors: Besmira Nushi, Tobias Schnabel, Saleema Amershi*
- Spring 2019 Research Intern, **Semantic Scholar, AI2**, Seattle, WA  
*Mentor: Dan Weld*
- Summer 2015 Software Engineer Intern, **Google Inc.**, Los Angeles, CA
- 2010-2012 High School Research Intern, **Institute of Microbiology, Chinese Academy of Sciences**, Beijing, China  
*Mentor: Baoli Zhu*

### GRANTS

- 2023 *Democratic Inputs to AI, OpenAI Inc.*, Case Law for AI Policy: Democratically Aligning AI with Human Preferences through Case Repositories

### TEACHING Courses

- Autumn 2022 Teaching Assistant, **CSE 440: Introduction To HCI**, UW CSE  
Graduate TA for undergraduate special topics course.
- Spring 2021 Teaching Assistant, **CSEP 590B: Enterprise Chatbots**, UW CSE  
Graduate TA for MS level special topics course.
- Winter 2021 Teaching Assistant, **CSE 415: Introduction to Artificial Intelligence**, UW CSE  
Graduate TA for non-majors AI intro course.

Winter 2019	Teaching Assistant, <b>CSE 473: Introduction to Artificial Intelligence</b> , UW CSE Graduate TA for majors AI intro course.
Spring 2019	Teaching Assistant, <b>CSEP 573: Introduction to Artificial Intelligence</b> , UW CSE Graduate TA for MS AI intro course.
2013 - 2016	Teaching Assistant, <b>CIS 121: Data Structures and Algorithms</b> , University of Pennsylvania Undergraduate TA, Head TA for undergraduate intro CS course.

## INVITED TALKS

Nov 2023	<b>Colloquium</b> , University of Texas, Austin, TX Working with Uncertainty to Align Humans and AI on Socially-Constructed Concepts
Sept 2023	<b>Democratic Inputs to AI</b> , OpenAI Inc., San Francisco, CA Case Law for AI Policy
July 2023	<b>ASE Events Invited Talk</b> , Apple Inc., Seattle, WA Tools and Workflows for Aligning Machine Learning Systems with Human Uncertainty
Aug 2022	<b>Stanford HCI Lunch</b> , Stanford University, Virtual Understanding and Addressing Uncertainty of the Crowd
June 2022	<b>Expertise@Scale Salon</b> , Carnegie Mellon University, Virtual Understanding and Addressing Uncertainty of the Crowd

## SERVICE

### Reviewer

Conferences: CHI '20-'24, CSCW '21-'25, WWW '20, CI '25

Journals: JAIR

Workshops: CHI '22-'23 TRAIT, CHI '24 TREW, NeurIPS '24 Pluralistic Alignment, CHI '25 STAIG

## STUDENT ADVISING

### Undergraduate Students

2023	Kathryn Yurechko
2022	Andre Ye (Now PhD Student at MIT) Teanna Barrett (Now PhD Student at UW) Kianna Bolante