

HCI+NLP 2024

**Third Workshop on Bridging Human–Computer Interaction
and Natural Language Processing**

Proceedings of the Workshop

June 21, 2024

The HCI+NLP organizers gratefully acknowledge the support from the following sponsors.

Sponsored by



©2024 Association for Computational Linguistics

Order copies of this and other ACL proceedings from:

Association for Computational Linguistics (ACL)
317 Sidney Baker St. S
Suite 400 - 134
Kerrville, TX 78028
USA
Tel: +1-855-225-1962
acl@aclweb.org

ISBN 979-8-89176-111-7

Introduction

Welcome to the Third Workshop on Bridging Human–Computer Interaction and Natural Language Processing!

The rapid advancement of natural language processing (NLP) research such as recent large language models has led to a variety of language technologies spanning a wide range of domains, such as conversational search and writing assistance. Those models are trained on vast amounts of data generated by people and rely on human feedback for continual improvement. While this widespread adoption ignites excitement, it raises pressing concerns and challenges in NLP research, such as real-world evaluation, bias and fairness, and model interpretability and explainability. Meanwhile, the field of human–computer interaction (HCI) develops rigorous methods for 1) studying and understanding human behavior to design technologies and 2) understanding how people use those technologies. Such a human-centered approach manifested in substantial efforts to understand the socio-cultural dynamics of data curation, to develop frameworks and tools to audit biases and ethical issues in intelligent systems, and to study people’s interaction with language technologies and its impact on people’s behavior.

This workshop aims to bridge the NLP and HCI communities to allow members of the NLP community to learn why, whether, and how methods and theories from HCI might be useful in advancing core NLP work, as well as allowing members of the HCI community to learn how advances in NLP might shape HCI research and practice centered around language technologies.

We are delighted to continue the effort of two previous editions of this HCI+NLP workshop at EACL 2021 and NAACL 2022 and bring the third edition to NAACL 2024. In this workshop we present eleven papers, of which nine are archival papers, and two are non-archival papers to be presented at the workshop but not included in the proceedings.

We would like to thank everyone who submitted their work to this workshop, as well as the program committee for their insightful review and feedback. We would also like to thank our invited speakers: Lydia Chilton and Sherry Tongshuang Wu.

We hope you find this workshop enjoyable! — Su Lin Blodgett, Amanda Cercas Curry, Sunipa Dev, Michael Madaio, Ani Nenkova, Diyi Yang, and Ziang Xiao

Organizing Committee

Organizing Committee

Su Lin Blodgett, Microsoft Research, Canada
Amanda Cercas Curry, Bocconi University, Italy
Sunipa Dev, Google Research, USA
Michael Madaio, Google Research, USA
Ani Nenkova, Adobe Research, USA
Diyi Yang, Stanford University, USA
Ziang Xiao, Johns Hopkins University, USA

Program Committee

Program Committee

Özge Alacam, Bielefeld University
Maurice Jakesch, Cornell University
Mahmood Jasim, University of Massachusetts Amherst
Geza Kovacs, Google
Nina Markl, University of Essex
Marianna J. Martindale, University of Maryland, College Park and U.S. DoD
David Mimno, Cornell University
Tatiana Passali, CERTH/ITI and Aristotle University of Thessaloniki
Indira Sen, Rheinisch Westfälische Technische Hochschule Aachen
Qinlan Shen, Oracle
Zijie J. Wang, Georgia Institute of Technology
Austin P Wright, Georgia Institute of Technology

Invited Speakers

Lydia B. Chilton, Columbia University, USA
Sherry Tongshuang Wu, Carnegie Mellon University, USA

Table of Contents

<i>Examining Prosody in Spoken Navigation Instructions for People with Disabilities</i> Cathy Jiao, Aaron Steinfeld and Maxine Eskenazi	1
<i>Properties and Challenges of LLM-Generated Explanations</i> Jenny Kunz and Marco Kuhlmann	13
<i>This Reference Does Not Exist: An Exploration of LLM Citation Accuracy and Relevance</i> Courtnei Byun, Piper Vasicek and Kevin Seppi	28
<i>Combining Multiple Metrics for Evaluating Retrieval-Augmented Conversations</i> Jason Ingyu Choi, Marcus D. Collins, Eugene Agichtein, Oleg Rokhlenko and Shervin Malmasi 40	
<i>How Much Annotation is Needed to Compare Summarization Models?</i> Chantal Shaib, Joe Barrow, Alexa Siu, Byron C Wallace and Ani Nenkova	51
<i>An Interactive Co-Pilot for Accelerated Research Ideation</i> Harshit Nigam, Manasi Patwardhan, Lovekesh Vig and Gautam Shroff	60
<i>Sensemaking of Socially-Mediated Crisis Information</i> Vrushali Koli, Jun Yuan and Aritra Dasgupta	74
<i>Blind Spots and Biases: Exploring the Role of Annotator Cognitive Biases in NLP</i> Sanjana Gautam and Mukund Srinath	82
<i>LLMCheckup: Conversational Examination of Large Language Models via Interpretability Tools and Self-Explanations</i> Qianli Wang, Tatiana Anikina, Nils Feldhus, Josef Van Genabith, Leonhard Hennig and Sebastian Möller	89