EACL 2021

Bridging Human–Computer Interaction and Natural Language Processing

Proceedings of the First Workshop

April 20, 2021

©2021 The Association for Computational Linguistics

Order copies of this and other ACL proceedings from:

Association for Computational Linguistics (ACL)
209 N. Eighth Street
Stroudsburg, PA 18360
USA
Tel: +1-570-476-8006
Fax: +1-570-476-0860
acl@aclweb.org

ISBN 978-1-954085-17-6

Introduction

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

The rapid advancement of natural language processing (NLP) research has led to a variety of applications spanning a wide range of domains. As NLP applications are being used by more people in their everyday lives, and are increasingly powered by data generated by people, it is more important than ever that NLP researchers and practitioners adopt and develop methods to incorporate people into their work in meaningful ways. Perspectives from human–computer interaction (HCI) can enable NLP researchers and practitioners to advance the field of NLP in ways that are aligned with people's needs, raising novel questions and research directions for both NLP and HCI.

The workshop brings together researchers and practitioners from both disciplines to discuss shared research interests, highlight work at the intersection of these fields, and identify challenges and opportunities for productive interdisciplinary collaborations.

We are delighted to present eighteen papers spanning reports of empirical work, research proposals, provocations, and surveys, along with seven 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 feedback. We would also like to thank our invited speakers: Marti Hearst, Niloufar Salehi, and Burr Settles.

We hope you enjoy the workshop!

—Su Lin Blodgett, Michael Madaio, Brendan O'Connor, Hanna Wallach, and Qian Yang

Organizers:

Su Lin Blodgett, Microsoft Research

Michael Madaio, Microsoft Research

Brendan O'Connor, University of Massachusetts Amherst

Hanna Wallach, Microsoft Research

Qian Yang, Cornell University

Program Committee:

Heloisa Candello

Stevie Chancellor

Sindhu Ernala

Mitchell Gordon

Huda Khayrallah

Geza Kovacs

Narges Mahyar

David Mimno

Swati Mishra

Carolyn Rosé

Koustuv Saha

Joseph Seering

Qinlan Shen

Alison Smith-Renner

Adam Trischler

Austin Wright

Michael Yoder

Zhou Yu

Justine Zhang

Invited Speakers:

Marti Hearst, University of California Berkeley

Niloufar Salehi, University of California Berkeley

Burr Settles, Duolingo

Table of Contents

Towards Hybrid Human-Machine Workflow for Natural Language Generation
Neslihan Iskender, Tim Polzehl and Sebastian Möller 1

Spellchecking for Children in Web Search: a Natural Language Interface Case-study
Casey Kennington, Jerry Alan Fails, Katherine Landau Wright and Maria Soledad Pera 8

Capturing Covertly Toxic Speech via Crowdsourcing
Alyssa Lees, Daniel Borkan, Ian Kivlichan, Jorge Nario and Tesh Goyal 14

Towards Human-Centered Summarization: A Case Study on Financial News
Tatiana Passali, Alexios Gidiotis, Efstathios Chatzikyriakidis and Grigorios Tsoumakas 21

Methods for the Design and Evaluation of HCI+NLP Systems
Hendrik Heuer and Daniel Buschek 28

Context-sensitive evaluation of automatic speech recognition: considering user experience & language variation
Nina Markl and Catherine Lai 34

Embodied Multimodal Agents to Bridge the Understanding Gap
Nikhil Krishnaswamy and Nada Alalyani 41

Putting Humans in the Natural Language Processing Loop: A Survey
Zijie J. Wang, Dongjin Choi, Shenyu Xu and Diyi Yang 47

Can You Distinguish Truthful from Fake Reviews? User Analysis and Assistance Tool for Fake Review Detection
Jeonghwan Kim, Junmo Kang, Suwon Shin and Sung-Hyon Myaeng 53

Challenges in Designing Games with a Purpose for Abusive Language Annotation
Federico Bonetti and Sara Tonelli 60

Challenges in Designing Natural Language Interfaces for Complex Visual Models
Henrik Voigt, Monique Meuschke, Kai Lawonn and Sina Zarrieß 66

AfriKI: Machine-in-the-Loop Afrikaans Poetry Generation
Imke van Heerden and Anil Bas 74

Situation-Specific Multimodal Feature Adaptation
Özge Alacam 81

Machine Translation Believability
Marianna Martindale, Kevin Duh and Marine Carpuat 88

Provocation: Contestability in Large-Scale Interactive NLP Systems
Tanushree Mitra 96

An IDR Framework of Opportunities and Barriers between HCI and NLP
Nanna Inie and Leon Derczynski 101

RE-AIMing Predictive Text
Matthew Higgs, Claire McCallum, Selina Sutton and Mark Warner 109

How do people interact with biased text prediction models while writing?
Advait Bhat, Saaket Agashe and Anirudha Joshi 116

Non-Archival Papers

Three Directions for the Design of Human-Centered Machine Translation
Samantha Robertson, Wesley Hanwen Deng, Timnit Gebru, Margaret Mitchell, Daniel Liebling, Michal Lahav, Katherine Heller, Mark Díaz, Samy Bengio and Niloufar Salehi

Humans predictably fail to detect fake self-presentations generated by AI
Benjamin Carson, Maurice Jakesch and Mor Naaman

What Human Rationales Tell Us about Machine Explanations
Indira Sen, Mattia Samory, Fabian Floeck and Claudia Wagner

An HCI+NLP Collaboration on Supporting Distributed Teamwork with Machine Translation
Yongle Zhang, Dennis Owusu, Emily Gong, Shaan Chopra, Marine Carpuat and Ge Gao

Influenced by Stochastic Parrots: Does AI-Assisted Writing Change Users' Views?
Yujie Shao, Maurice Jakesch and Mor Naaman

Spaced Repetition Meets Representation Learning
Matthew Shu, Shi Feng and Jordan Boyd-Graber

Wordcraft: a Human-AI Collaborative Editor for Story Writing
Andy Coenen, Luke Davis, Daphne Ippolito, Emily Reif and Ann Yuan