

WASSA 2023

**The 13th Workshop on Computational Approaches to
Subjectivity, Sentiment, & Social Media Analysis**

Proceedings of the Workshop

July 14, 2023

The WASSA organizers gratefully acknowledge the support from the following sponsors.

Gold



©2023 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 978-1-959429-87-6

Introduction

While the first editions of the Workshop on Computational Approaches to Subjectivity Analysis (WASSA) focused mainly on sentiment and subjectivity analysis, the scope of this workshop has become both wider and deeper over the years, expanding the number of languages and tasks that are considered, while simultaneously looking at more specific subjective tasks.

The topics of this edition of WASSA range from exploring how harmful language proliferates in social media and news, to looking at the effects of using large pretrained language models on well-known sentiment/emotion analysis tasks, or creating resources for new tasks that focus on subjective language. Regarding languages, in this edition English is unsurprisingly the most represented (in 29 papers), with 2 papers each including Mandarin Chinese, Dutch, German, Hindi, and Spanish. This is followed by a long tail of languages that appear in only one paper (Algerian Arabic, Basque, Czech, Dari, French, Gujarati, Italian, Japanese, Kannada, Korean, Malayalam, Maltese, Modern Standard Arabic, Norwegian, and Tamil). Two papers also address multilingual emotion detection.

This year, WASSA offered two shared tasks: Shared-Task 1: Empathy Detection and Emotion Classification, which builds upon the previous shared task with a new track on empathy, emotion and self-disclosure detection in conversation at the speech-turn level, using a new, unpublished dataset. Shared-Task 2: Emotion Classification on code-mixed text messages (Roman Urdu + English), on the other hand, is a novel shared task focusing on multi-class and multi-label emotion classification on code-mixed (Roman Urdu + English) text messages. Two papers also submitted work on multilingual emotion analysis.

For the main workshop, we received 71 papers, one of the largest number of submissions for WASSA in its history. While we accepted a comparable large number, namely 43 (60.6%), we would like to point out that the competition has been quite tough – most submitted papers were of very high quality. For Shared Task 1 we received 12 system description paper submissions, while for Shared Task 2 we received 7 papers. In total, 61 papers will be presented at the workshop.

We also decided to award the best paper. Starting with the reviewer scores, the organizers of the workshop anonymously proposed candidates and then voted on them. The winner of the Best Paper Awards of this year’s WASSA is *Multilingual Language Models are not Multicultural: A Case Study in Emotion by Shreya Havaldar, Sunny Rai, Bhumika Singhal, Langchen Liu Sharath Chandra Guntuku, and Lyle Ungar*.

On top of the main workshop and shared task presentations, we are happy to have two invited speakers who will hold keynotes at WASSA 2023: Professor Emily Öhman from Wakeda University and Professor David Jurgens from the University of Michigan.

We would like to thank the ACL Organizers and Workshop chairs for their help and support during the preparation. We thank Google for their gold sponsorship. We also thank the OpenReview support team for their technical support. Finally, we especially thank the program committee for the time and effort they spent on reviewing, especially to those who were willing to perform emergency reviews. The quality of the program is a reflection of this hard work.

Jeremy Barnes, Orphée De Clercq, Roman Klinger

WASSA 2023 Chairs

Organizing Committee

Program Chairs

Jeremy Barnes, University of the Basque Country UPV/EHU, Spain
Orphée De Clercq, Ghent University, Belgium
Roman Klinger, University of Stuttgart, Germany

Shared Task 1 Chairs

Valentin Barriere, Centro Nacional de Inteligencia Artificial, Chile
Salvatore Giorgi, University of Pennsylvania, USA
Joao Sedoc, New York University, USA
Shabnam Tafreshi, University of Maryland, USA

Shared Task 2 Chairs

Iqra Ameer, Yale University, USA
Necva Bölücü, Hacettepe University, Republic of Türkiye
Hua Xu, Yale University, USA
Ali Al Bataineh, Norwich University, USA

Program Committee

Program Committee

Rodrigo Agerri, University of the Basque Country
Wazir Ali, College of Business Management (CBM)
Murali Balusu, Georgia Institute of Technology
Sivaji Bandyopadhyay, Jadavpur University
Tilman Beck, TU Darmstadt
Emerson Cabrera Paraiso, Pontificia Universidade Católica do Paraná
Luna De Bruyne, Universiteit Gent
Loic De Langhe, Universiteit Gent
Koji Eguchi, Hiroshima University
Annerose Eichel, University of Stuttgart, Universität Stuttgart
Thomas Nikolaus Haider, Universität Passau
Mohamad Hardyman Barawi, Universiti Malaysia Sarawak
Yoshihiko Hayashi, Waseda University
Kokil Jaidka, National University of Singapore
Salud María Jiménez-Zafra, Universidad de Jaén
Abhinav Joshi, Indian Institute of Technology, Kanpur
Akbar Karimi, Phillips-Universität Marburg
Elma Kerz, Rheinisch Westfälische Technische Hochschule Aachen
Mayank Kulkarni, Amazon
Sofie Labat, Universiteit Gent
Els Lefever, Ghent University
Natalia V Loukachevitch, Lomonosov Moscow State University
Petter Mæhlum, University of Oslo
Aaron Maladry, Universiteit Gent
María-Teresa Martín-Valdivia, Universidad de Jaén
Eugenio Martinez-Camara, Universidad de Jaén
Amita Misra, Amazon
Ashutosh Modi, IIT Kanpur
Constantin Orasan, University of Surrey
Jessica Ouyang, University of Texas at Dallas
Flor Miriam Plaza del Arco, Bocconi University
Joan Plepi, Math and CS, Marburg University
Anil Ramakrishna, Amazon
Irene Russo, Consiglio Nazionale delle Ricerche
Iñaki San Vicente, Orai NLP Technologies
Salim Sazed, Old Dominion University
Pranaydeep Singh, Universiteit Gent
Efsthios Stamatatos, University of the Aegean
Zeerak Talat, Simon Fraser University
Sai Kiran Tanikella, IIT Kanpur, Indian Institute of Technology, Kanpur
Samia Touileb, University of Bergen, Norway
Amine Trabelsi, Université de Sherbrooke
Enrica Troiano, Vrije Universiteit Amsterdam
L. Alfonso Ureña, Universidad de Jaén
Aswathy Velutharambath, University of Stuttgart, Universität Stuttgart
Jacky Visser, University of Dundee

Jin Wang, Yunnan University
Charles Welch, University of Marburg
Michael Wiegand, Alpen-Adria Universität Klagenfurt
Tak-Lam Wong, Douglas College
Arkaitz Zubiaga, Queen Mary University of London

Invited Speakers

David Jurgens, University of Michigan
Emily Öhman, Wakeda University

Keynote Talk: Invited 1

David Jurgens

University of Michigan

2023-07-14 12:30:00 – Room: **Harbour C**

Abstract: NLP studies of communication often focus on the individual: What we say, when we say it, and how we say it. Yet, the larger social context beyond the individual also plays an important role in our communication — just think of things you can say to your friends but not your parents. How does the social context influence our communication style and content? In this talk, I will describe recent work from my group studying the influence of this context by examining how we choose who to communicate with, how we interpret messages, and how we phrase messages. Across these studies, I will motivate a causal approach for NLP when studying communication behavior to move beyond descriptive analyses to more precise estimates of the effects of social context.

Bio: David Jurgens is an assistant professor at the University of Michigan School of Information where he leads the Blablablab. He holds a PhD in Computer Science from the University of California, Los Angeles. His research focuses on the intersection between NLP and computational social science venues and has won the Cozzarelli Prize, Cialdini Prize, best paper at ICWSM and W-NUT, and best paper nomination at ACL and Web Science.

Keynote Talk: Invited 2

Emily Öhman

Waseda University

2023-07-14 16:00:00 – Room: Harbour C

Abstract: Affective Datafication of Narratives: measuring affect, emotion, and mood in literary texts
Our understanding of affect, emotion, and mood - despite the distinct nuances each term holds - often becomes blurred, leading to a usage that is almost interchangeable, particularly within sentiment analysis and NLP. In contrast, traditional fields such as literary studies hold on to more rigid definitions of these terms and how they are understood both in theory and practice. This can easily foster a disconnect between emerging fields such as computational literary studies and the more established qualitative counterparts. This disconnect unfortunately hinders the free exchange of innovative research ideas and methodologies. This talk aims to bridge this gap, highlighting the unique roles of affect, emotion, and mood in narratives and how we can attempt to robustly measure them. We will delve into the interplay of these terms, exploring how they shape and are shaped by authors, readers, and researchers focusing on the operationalization and translation involved in the analysis of emotion-laden phenomena. This exploration will underscore the need for a more comprehensive and nuanced understanding, encouraging synergy between tradition and innovation in emotion detection in general and literary research in particular.

Bio: Emily Öhman is currently a tenure-track Assistant professor of Digital Humanities at Waseda University. She received her PhD in Language Technology from the University of Helsinki, where her work centered on building multilingual emotion detection resources for downstream tasks.

Her research interests lie within digital humanities and NLP, more specifically sentiment analysis and emotion detection, often doing collaborations with various disciplines such as history, literature, and political science. Her recent projects have focused on negative emotions in literature using affect as a proxy for the literary concept of mood and most recently contrasting the semantic spaces of shame and guilt in Japanese and English social media posts.

Table of Contents

<i>PESTO: A Post-User Fusion Network for Rumour Detection on Social Media</i> Erxue Min and Sophia Ananiadou	1
<i>Sentimental Matters - Predicting Literary Quality by Sentiment Analysis and Stylometric Features</i> Yuri Bizzoni, Pascale Feldkamp Moreira, Mads Rosendahl Thomsen and Kristoffer Nielbo ...	11
<i>Instruction Tuning for Few-Shot Aspect-Based Sentiment Analysis</i> Siddharth Varia, Shuai Wang, Kishaloy Halder, Robert Vacareanu, Miguel Ballesteros, Yassine Benajiba, Neha Anna John, Rishita Anubhai, Smaranda Muresan and Dan Roth	19
<i>You Are What You Read: Inferring Personality From Consumed Textual Content</i> Adam Sutton, Almog Simchon, Matthew Edwards and Stephan Lewandowsky	28
<i>UNIDECOR: A Unified Deception Corpus for Cross-Corpus Deception Detection</i> Aswathy Velutharambath and Roman Klinger	39
<i>Discourse Mode Categorization of Bengali Social Media Health Text</i> Salim Sazzed	52
<i>Emotion and Sentiment Guided Paraphrasing</i> Justin J Xie and Ameeta Agrawal	58
<i>Emotions in Spoken Language - Do we need acoustics?</i> Nadine Probol and Margot Mieskes	71
<i>Understanding Emotion Valence is a Joint Deep Learning Task</i> Gabriel Roccabruna, Seyed Mahed Mousavi and Giuseppe Riccardi	85
<i>Czech-ing the News: Article Trustworthiness Dataset for Czech</i> Matyas Bohacek, Michal Bravansky, Filip Trhlík and Vaclav Moravec	96
<i>Towards Detecting Harmful Agendas in News Articles</i> Melanie Subbiah, Amrita Bhattacharjee, Yilun Hua, Tharindu Sandaruwan Kumarage, Huan Liu and Kathleen McKeown	110
<i>GSAC: A Gujarati Sentiment Analysis Corpus from Twitter</i> Monil Gokani and Radhika Mamidi	129
<i>A Dataset for Explainable Sentiment Analysis in the German Automotive Industry</i> Andrea Zielinski, Calvin Spolwind, Henning Kroll and Anna Grimm	138
<i>Examining Bias in Opinion Summarisation through the Perspective of Opinion Diversity</i> Nannan Huang, Lin Tian, Haytham M. Fayek and Xiuzhen Zhang	149
<i>Fluency Matters! Controllable Style Transfer with Syntax Guidance</i> Ji-Eun Han and Kyung-Ah Sohn	162
<i>ChatGPT for Suicide Risk Assessment on Social Media: Quantitative Evaluation of Model Performance, Potentials and Limitations</i> Hamideh Ghanadian, Isar Nejadgholi and Hussein Al Osman	172
<i>Unsupervised Domain Adaptation using Lexical Transformations and Label Injection for Twitter Data</i> Akshat Gupta, Xiaomo Liu and Sameena Shah	184

<i>Transformer-based cynical expression detection in a corpus of Spanish YouTube reviews</i> Samuel Gonzalez-Lopez and Steven Bethard	194
<i>Multilingual Language Models are not Multicultural: A Case Study in Emotion</i> Shreya Havaladar, Bhumika Singhal, Sunny Rai, Langchen Liu, Sharath Chandra Guntuku and Lyle Ungar	202
<i>Painsight: An Extendable Opinion Mining Framework for Detecting Pain Points Based on Online Customer Reviews</i> Yukyung Lee, Jaehee Kim, Doyoon Kim, Yookyung Kho, Younsun Kim and Pilsung Kang ..	215
<i>Context-Dependent Embedding Utterance Representations for Emotion Recognition in Conversations</i> Patrícia Pereira, Helena Silva Moniz, Isabel Dias and Joao Paulo Carvalho	228
<i>Combining Active Learning and Task Adaptation with BERT for Cost-Effective Annotation of Social Media Datasets</i> Jens Lemmens and Walter Daelemans	237
<i>Improving Dutch Vaccine Hesitancy Monitoring via Multi-Label Data Augmentation with GPT-3.5</i> Jens Van Nooten and Walter Daelemans	251
<i>Emotion Analysis of Tweets Banning Education in Afghanistan</i> Mohammad Ali Hussiny and Lilja Øvrelid	271
<i>Identifying Slurs and Lexical Hate Speech via Light-Weight Dimension Projection in Embedding Space</i> Sanne Hoeken, Sina Zarriß and Ozge Alacam	278
<i>Sentiment and Emotion Classification in Low-resource Settings</i> Jeremy Barnes	290
<i>Analyzing Subjectivity Using a Transformer-Based Regressor Trained on Naïve Speakers' Judgements</i> Elena Savinova and Fermin Moscoso Del Prado	305
<i>A Fine Line Between Irony and Sincerity: Identifying Bias in Transformer Models for Irony Detection</i> Aaron Maladry, Els Lefever, Cynthia Van Hee and Veronique Hoste	315
<i>ChatGPT is fun, but it is not funny! Humor is still challenging Large Language Models</i> Sophie Jentzsch and Kristian Kersting	325
<i>How to Control Sentiment in Text Generation: A Survey of the State-of-the-Art in Sentiment-Control Techniques</i> Michela Lorandi and Anya Belz	341
<i>Transformer-based Prediction of Emotional Reactions to Online Social Network Posts</i> Irene Benedetto, Moreno La Quatra, Luca Cagliero, Luca Vassio and Martino Trevisan	354
<i>Transfer Learning for Code-Mixed Data: Do Pretraining Languages Matter?</i> Kushal Tatariya, Heather Lent and Miryam De Lhoneux	365
<i>Can ChatGPT Understand Causal Language in Science Claims?</i> Yuheun Kim, Lu Guo, Bei Yu and Yingya Li	379
<i>Systematic Evaluation of GPT-3 for Zero-Shot Personality Estimation</i> Adithya V Ganesan, Yash Kumar Lal, August Håkan Nilsson and H. Schwartz	390
<i>Utterance Emotion Dynamics in Children's Poems: Emotional Changes Across Age</i> Daniela Teodorescu, Alona Fyshe and Saif M. Mohammad	401

<i>Annotating and Training for Population Subjective Views</i>	
Maria Alexeeva, Caroline Hyland, Keith Alcock, Allegra A. Beal Cohen, Hubert Kanyamahanga, Isaac Kobby Anni and Mihai Surdeanu	416
<i>Exploration of Contrastive Learning Strategies toward more Robust Stance Detection</i>	
Udhaya Kumar Rajendran and Amine Trabelsi	431
<i>Adapting Emotion Detection to Analyze Influence Campaigns on Social Media</i>	
Ankita Bhaumik, Andy Bernhardt, Gregorios A Katsios, Ning Sa and Tomek Strzalkowski . .	441
<i>Not Just Iconic: Emoji Interpretation is Shaped by Use</i>	
Brianna O’Boyle and Gabriel Doyle	452
<i>The Paradox of Multilingual Emotion Detection</i>	
Luna De Bruyne	458
<i>Sadness and Anxiety Language in Reddit Messages Before and After Quitting a Job</i>	
Molly E. Ireland, Micah Iserman and Kiki Rose Adams	467
<i>Communicating Climate Change: A Comparison Between Tweets and Speeches by German Members of Parliament</i>	
Robin Schaefer, Christoph Maximilian Abels, Stephan Lewandowsky and Manfred Stede . . .	479
<i>Modelling Political Aggression on Social Media Platforms</i>	
Akash Rawat, Nazia Nafis, Dnyaneshwar Bhadane, Diptesh Kanojia and Rudra Murthy	497
<i>Findings of WASSA 2023 Shared Task on Empathy, Emotion and Personality Detection in Conversation and Reactions to News Articles</i>	
Valentin Barriere, João Sedoc, Shabnam Tafreshi and Salvatore Giorgi	511
<i>YNU-HPCC at WASSA-2023 Shared Task 1: Large-scale Language Model with LoRA Fine-Tuning for Empathy Detection and Emotion Classification</i>	
Yukun Wang, Jin Wang and Xuejie Zhang	526
<i>AdityaPatkar at WASSA 2023 Empathy, Emotion, and Personality Shared Task: RoBERTa-Based Emotion Classification of Essays, Improving Performance on Imbalanced Data</i>	
Aditya Patkar, Suraj Chandrashekhar and Ram Mohan Rao Kadiyala	531
<i>Curtin OCAI at WASSA 2023 Empathy, Emotion and Personality Shared Task: Demographic-Aware Prediction Using Multiple Transformers</i>	
Md Rakibul Hasan, Md Zakir Hossain, Tom Gedeon, Susannah Soon and Shafin Rahman . . .	536
<i>Team_Hawk at WASSA 2023 Empathy, Emotion, and Personality Shared Task: Multi-tasking Multi-encoder based transformers for Empathy and Emotion Prediction in Conversations</i>	
Addepalli Sai Srinivas, Nabarun Barua and Santanu Pal	542
<i>NCUEE-NLP at WASSA 2023 Shared Task 1: Empathy and Emotion Prediction Using Sentiment-Enhanced RoBERTa Transformers</i>	
Tzu-Mi Lin, Jung-Ying Chang and Lung-Hao Lee	548
<i>Domain Transfer for Empathy, Distress, and Personality Prediction</i>	
Fabio Gruschka, Allison Lahnala, Charles Welch and Lucie Flek	553
<i>Converge at WASSA 2023 Empathy, Emotion and Personality Shared Task: A Transformer-based Approach for Multi-Label Emotion Classification</i>	
Aditya Paranjape, Gaurav Kolhatkar, Yash Patwardhan, Omkar Gokhale and Shweta Dharmadhikari	558

<i>PICT-CLRL at WASSA 2023 Empathy, Emotion and Personality Shared Task: Empathy and Distress Detection using Ensembles of Transformer Models</i>	
Tanmay Chavan, Kshitij Deshpande and Sheetal Sonawane	564
<i>Team Bias Busters at WASSA 2023 Empathy, Emotion and Personality Shared Task: Emotion Detection with Generative Pretrained Transformers</i>	
Andrew Nedilko and Yi Chu	569
<i>HIT-SCIR at WASSA 2023: Empathy and Emotion Analysis at the Utterance-Level and the Essay-Level</i>	
Xin Lu, Zhuojun Li, Yanpeng Tong, Yanyan Zhao and Bing Qin	574
<i>VISU at WASSA 2023 Shared Task: Detecting Emotions in Reaction to News Stories Using Transformers and Stacked Embeddings</i>	
Vivek Kumar, Prayag Tiwari and Sushmita Singh	581
<i>Findings of WASSA 2023 Shared Task: Multi-Label and Multi-Class Emotion Classification on Code-Mixed Text Messages</i>	
Iqra Ameer, Necva Bölücü, Hua Xu and Ali Al Bataineh	587
<i>Emotion classification on code-mixed text messages via soft prompt tuning</i>	
Jinghui Zhang, Dongming Yang, Siyu Bao, Lina Cao and Shunguo Fan	596
<i>PrecogIITH@WASSA2023: Emotion Detection for Urdu-English Code-mixed Text</i>	
Bhaskara Hanuma Vedula, Prashant Kodali, Manish Shrivastava and Ponnurangam Kumaraguru	
601	
<i>BpHigh at WASSA 2023: Using Contrastive Learning to build Sentence Transformer models for Multi-Class Emotion Classification in Code-mixed Urdu</i>	
Bhavish Pahwa	606
<i>YNU-HPCC at WASSA 2023: Using Text-Mixed Data Augmentation for Emotion Classification on Code-Mixed Text Message</i>	
Xuqiao Ran, You Zhang, Jin Wang, Dan Xu and Xuejie Zhang	611
<i>Generative Pretrained Transformers for Emotion Detection in a Code-Switching Setting</i>	
Andrew Nedilko	616

Program

Friday, July 14, 2023

09:10 - 09:00 *Opening Remarks*

09:10 - 10:25 *Session 1: Sentiment, Emotion, and Irony Recognition*

A Fine Line Between Irony and Sincerity: Identifying Bias in Transformer Models for Irony Detection

Aaron Maladry, Els Lefever, Cynthia Van Hee and Veronique Hoste

Instruction Tuning for Few-Shot Aspect-Based Sentiment Analysis

Siddharth Varia, Shuai Wang, Kishaloy Halder, Robert Vacareanu, Miguel Ballesteros, Yassine Benajiba, Neha Anna John, Rishita Anubhai, Smaranda Muresan and Dan Roth

Unsupervised Domain Adaptation using Lexical Transformations and Label Injection for Twitter Data

Akshat Gupta, Xiaomo Liu and Sameena Shah

Context-Dependent Embedding Utterance Representations for Emotion Recognition in Conversations

Patrícia Pereira, Helena Silva Moniz, Isabel Dias and Joao Paulo Carvalho

Multilingual Language Models are not Multicultural: A Case Study in Emotion

Shreya Havaladar, Bhumika Singhal, Sunny Rai, Langchen Liu, Sharath Chandra Guntuku and Lyle Ungar

10:30 - 11:00 *Coffee Break*

12:30 - 11:00 *Hybrid Poster Session*

12:30 - 13:30 *Invited Talk 1: David Jurgens*

13:30 - 14:30 *Lunch Break*

14:30 - 15:30 *Shared task session*

Findings of WASSA 2023 Shared Task on Empathy, Emotion and Personality Detection in Conversation and Reactions to News Articles

Valentin Barriere, João Sedoc, Shabnam Tafreshi and Salvatore Giorgi

Friday, July 14, 2023 (continued)

HIT-SCIR at WASSA 2023: Empathy and Emotion Analysis at the Utterance-Level and the Essay-Level

Xin Lu, Zhuojun Li, Yanpeng Tong, Yanyan Zhao and Bing Qin

Findings of WASSA 2023 Shared Task: Multi-Label and Multi-Class Emotion Classification on Code-Mixed Text Messages

Iqra Ameer, Necva Bölücü, Hua Xu and Ali Al Bataineh

YNU-HPCC at WASSA 2023: Using Text-Mixed Data Augmentation for Emotion Classification on Code-Mixed Text Message

Xuqiao Ran, You Zhang, Jin Wang, Dan Xu and Xuejie Zhang

15:30 - 16:00 *Coffee Break*

16:00 - 17:00 *Invited Talk 2: Emily Öhman*

17:00 - 18:00 *Session 2: Harmful Speech and Dissatisfaction Recognition*

Identifying Slurs and Lexical Hate Speech via Light-Weight Dimension Projection in Embedding Space

Sanne Hoeken, Sina Zarriß and Ozge Alacam

Modelling Political Aggression on Social Media Platforms

Akash Rawat, Nazia Nafis, Dnyaneshwar Bhadane, Diptesh Kanojia and Rudra Murthy

Towards Detecting Harmful Agendas in News Articles

Melanie Subbiah, Amrita Bhattacharjee, Yilun Hua, Tharindu Sandaruwan Kumarage, Huan Liu and Kathleen McKeown

Painsight: An Extendable Opinion Mining Framework for Detecting Pain Points Based on Online Customer Reviews

Yukyung Lee, Jaehee Kim, Doyoon Kim, Yookyung Kho, Younsun Kim and Pilsung Kang

18:00 - 18:15 *Closing Remarks and Best Paper Awards*