

WASSA 2026

**15th Workshop on Computational Approaches to
Subjectivity, Sentiment and Social Media Analysis**

Proceedings of the Workshop

March 29, 2026

©2026 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 979-8-89176-378-4

Introduction

This edition marks the 15th iteration of the Workshop on Computational Approaches to Subjectivity, Sentiment and Social Media Analysis (WASSA), a crystal jubilee!

We can summarize the workshop as focusing on topics at the intersection of emotion, subjectivity, persuasion, and trust. Crystal is fragile, and many topics at this year's WASSA can be situated within a social or political context, reflecting the brittle state our world is in. Regarding methodologies, Large Language Models (LLMs) prevail: as tools for simplification or for simulating social media users, as synthetic data generators, and as an object of study. This year, WASSA also featured a Special Track on Multilinguality and Social Bridges between High- and Lesser-Resourced Languages and Communities. In the proceedings, one paper takes a truly multilingual approach to quantify social sentiment in hostels. Only three papers do not mention English data at all, but instead focus on subjective tasks in Farsi, emotion preservation when translating from Spanish to Basque, and assessing the capability of LLMs to predict the emotional ratings of Russian words. Other languages that are being studied alongside English are German, Danish, and Luxembourgish.

For the main workshop, we received 32 papers and 9 ARR commitments. From the direct submissions, we accepted 17 (53%), and from ARR, another 6 (66%), bringing the total to 23 accepted papers. The high quality of much of the work made the decision process difficult, as this year, WASSA is held as a half-day workshop, reducing the room we had for accepted papers.

For our jubilee, we decided to award both a Best Paper and a Best Theme Paper Award. Starting with the reviewer scores, the organizers of the workshop anonymously proposed candidates and then voted on them. The winner of the Best Paper Award of this year's WASSA is *The Impact of Highlighting Subjective Language on Perceived News Trustworthiness* by Mohammad Shokri, Vivek Sharma, Emily Klapper, Shweta Jain, Elena Filatova, and Sarah Ita Levitan. The winner of the Best Theme Paper Award is *Crowd-Based Evaluation of Emotion Intensity Preservation in Spanish–Basque Tweet Machine Translation* by Nora Aranberri. Congratulations!

On top of the main workshop, we are happy to have José Camacho-Collados from the Cardiff NLP group as our invited speaker. In his keynote he will discuss social media analysis in the language model era, and this from an interdisciplinary perspective.

We would like to thank the EACL Organizers and Workshop chairs for their help and support during the preparation. We also thank the OpenReview support team for their technical support. Finally, we 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 your hard work.

Finally, we would like to dedicate WASSA 2026 to Shabnam Tafreshi. Shab was one of the driving forces behind many previous iterations of WASSA and the shared tasks on emotion and empathy prediction.

Jeremy Barnes, Valentin Barriere, Orphée De Clercq, Roman Klinger, Debora Nozza, Célia Nouri & Pranaydeep Singh

WASSA 2026 Chairs

Acknowledgements

This work is partially supported by the Basque Government (IXA excellence research group IT1570-22) and the European Union under Horizon Europe (Project LUMINOUS, grant number 101135724). Debora Nozza's research is supported by the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation program (grant agreement No. 101116095, PERSONAE). Valentin's work is partially funded by ANID grant Fondecyt de Iniciacion 11251024. Pranaydeep's work is supported by a BOF post-doctoral research fellowship (BOF.PDO.2025.0011.01) from Ghent University.

Organizing Committee

Program Chairs

Jeremy Barnes, University of the Basque Country UPV/EHU, Spain

Valentin Barriere, Universidad de Chile, Chile

Orphée De Clercq, Ghent University, Belgium

Roman Klinger, Otto-Friedrich Universität Bamberg, Germany

Debora Nozza, Bocconi University, Italy

Célia Nouri, INRIA and Sciences Po, France

Pranaydeep Singh, Ghent University, Belgium

Program Committee

Reviewers

Gavin Abercrombie, Wazir Ali, Liesbeth Allein

Murali Balusu, Sivaji Bandyopadhyay, Delphine Battistelli

Jose Camacho-Collados, Alessandra Teresa Cignarella

Luna De Bruyne, Loic De Langhe, Quanqi Du

Natalia Evgrafova

Neele Falk, Pascale Feldkamp, Elena Filatova

Gaël Guibon

Thomas Haider, Junko Hayashi, Jet Hoek, Andreas Hotho

Imezadelajara, Sekh Mainul Islam, Kazuhiro Ito, Gabriel Iturra-Bocaz

Aditya Joshi

Akbar Karimi, Yannis Karmim, Elma Kerz, Khalid Al Khatib, Marek Kubis, Kemal Kurniawan

Sofie Labat, Oier Lopez De Lacalle, Els Lefever, Natalia V Loukachevitch

Cinthia Sanchez Macias, Aaron Maladry, Eugenio Martínez-Cámara, Petter Mæhlum

Emerson Cabrera Paraiso

Tamara Quiroga

Ananya Raval, David Rossouw, Irene Russo, Egil Rønningstad

Hamidreza Saffari, Iñaki San Vicente, Elena Savinova, Mohammadamin Shafiei, Vivek Sharma, Mohammad Shokri, Phillip Smith, Efstathios Stamatatos, Benjamin Steel

Samia Touileb, Enrica Troiano

L. Alfonso Ureña

Cynthia Van Hee, Eva Maria Vecchi, Aswathy Velutharambath, David Vilares

Shoko Wakamiya, Jin Wang, Jingjing Wang, Michael Wiegand, Tak-Lam Wong

Zehong Yan

Lei Zhang

Keynote Talk

Social Media Analysis in the Language Model Era: An Interdisciplinary Perspective

Jose Camacho-Collados

Cardiff University, UK

29 March, 2026 – Time: 14:30 – 15:30 – Room: SALLE Le Chellah

Abstract: Understanding what is going on in social media is not a simple task, even after the significant progress in NLP in recent years. Given the specific nature of social media and the large amount of content generated, efficient and specialised solutions are often necessary.

At the same time, language models, and in particular large generative language models (LLMs) are currently ubiquitous across the NLP research landscape and beyond. However, it isn't clear how to effectively leverage them to process large amounts of unstructured data, such as the type usually found on social media, especially for non-expert users. There are also additional challenges when processing large amounts of text, such as ensuring transparency in the process and dealing with long context.

In this talk, I will explain my journey when it comes to social media analysis, from proposing efficient solutions for tasks such as sentiment analysis, hate speech detection or emotion detection, to sharing tips and lessons learned from interdisciplinary collaborations. I will try to revisit the importance of high-quality data and reliable evaluation benchmarks in the current LLM landscape, and some recent results suggesting promising areas for future research.

Bio: Jose Camacho-Collados is a UKRI Future Leaders Fellow and Professor at the School of Computer Science of Cardiff University, where he co-founded the Cardiff Natural Language Processing group (Cardiff NLP). Before joining Cardiff University, he completed his PhD in Sapienza University of Rome and was a Google AI PhD Fellow.

Jose has worked in multiple NLP areas with a particular focus on semantics, multilinguality and computational social science with an interdisciplinary perspective. In this area, he has been developing specialised and efficient NLP models for social media applications, such as TweetNLP and related efforts. His work has received several recognitions, including awards at top NLP conferences, or the 2023 AIJ Prominent Paper Award. He is also the co-author of the “Embeddings in Natural Language Processing” book.

Table of Contents

<i>Council of LLMs: Evaluating Capability of Large Language Models to Annotate Propaganda</i> Vivek Sharma, Shweta Jain, Mohammad Shokri, Sarah Ita Levitan and Elena Filatova	1
<i>Emoji Reactions on Telegram: Unreliable Indicators of Emotional Resonance</i> Serena Tardelli, Lorenzo Alvisi, Lorenzo Cima, Stefano Cresci and Maurizio Tesconi	13
<i>Quantifying Social Sentiment in Hostels Using A Domain-Specific Transformer Pipeline</i> Ian W. McMurry	24
<i>Predicting Convincingness in Political Speech: How Emotional Tone Shapes Persuasive Strength</i> Bhuvanesh Verma, Mounika Marreddy and Alexander Mehler	37
<i>Measuring LLMs’ Sensitivity to Paraphrased Opinion Prompts</i> Bushra Alhetelah and Irfan Ahmad	52
<i>The Impact of Highlighting Subjective Language on Perceived News Trustworthiness</i> Mohammad Shokri, Vivek Sharma, Emily Klapper, Shweta Jain, Elena Filatova and Sarah Ita Levitan	60
<i>Appraisal Trajectories in Narratives Reveal Distinct Patterns of Emotion Evocation</i> Johannes Schäfer, Janne Wagner and Roman Klinger	73
<i>Exploring Subjective Tasks in Farsi: A Survey Analysis and Evaluation of Language Model</i> Donya Rooein, Flor Miriam Plaza-del-Arco, Debora Nozza and Dirk Hovy	83
<i>Emotional Lexicons: How Large Language Models Predict Emotional Ratings of Russian Words</i> Polina V. Iaroshenko and Natalia V Loukachevitch	96
<i>Emotion-aware text simplification of user generated content using LLMs</i> Anastasiia Bezobrazova, Daria Sokova and Constantin Orasan	107
<i>Crowd-Based Evaluation of Emotion Intensity Preservation in Spanish–Basque Tweet Machine Translation</i> Nora Aranberri	123
<i>A Position Paper on Toxic Reasoning: Grounding Categories of Toxic Language in Implications and Attitudes</i> Stefan F. Schouten, Ilia Markov and Piek Vossen	134
<i>Is Sentiment Banana-Shaped? Exploring the Geometry and Portability of Sentiment Concept Vectors</i> Laurits Lyngbaek, Pascale Feldkamp, Yuri Bizzoni, Kristoffer Nielbo and Kenneth Enevoldsen	146
<i>Disentangling Emotion Understanding and Generation in Large Language Models</i> Sadegh Jafari, Els Lefever and Veronique Hoste	161
<i>News Credibility Assessment by LLMs and Humans: Implications for Political Bias</i> Pia Wenzel Neves, Charlott Jakob and Vera Schmitt	172
<i>Towards Simulating Social Media Users with LLMs: Evaluating the Operational Validity of Conditioned Comment Prediction</i> Nils Schwager, Simon Münker, Alistair Plum and Achim Rettinger	208

<i>Label-Consistent Data Generation for Aspect-Based Sentiment Analysis Using LLM Agents</i> Mohammad Hossein Akbari Monfared, Lucie Flek and Akbar Karimi	222
<i>Antisocial Behavior Prediction: A Survey and Practical Guide</i> Anaïs Ollagnier	235
<i>Real-Time Mitigation of Negative Emotion in Customer Care Calls</i> Surupendu Gangopadhyay and Mahnoosh Mehrabani	252
<i>Says Who? Argument Convincingness and Reader Stance Are Correlated with Perceived Author Personality</i> Sabine Weber, Lynn Greschner and Roman Klinger	265
<i>A Transformer and Prototype-based Interpretable Model for Contextual Sarcasm Detection</i> Ximing Wen and Rezvaneh Rezapour	278
<i>Multimodal Claim Extraction for Fact-Checking</i> Joycelyn Teo, Rui Cao, Zhenyun Deng, Zifeng Ding, Michael Sejr Schlichtkrull and Andreas Vlachos	289
<i>A Multi-Aspect Evaluation Framework for Synthetic Data: Case Study on Irony and Sarcasm</i> Laura Majer, Ana Barić, Florijan Sandalj, Ivan Unković, Bojan Puvača and Jan Šnajder	305