EACL 2021

The Workshop on Computational Approaches to Subjectivity, Sentiment and Social Media Analysis

Proceedings of the Eleventh Workshop

April 19, 2021
Held online due to the COVID situation

©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 978-1-954085-18-3

Introduction

After the tenth edition of WASSA in 2019, which came only seven months after WASSA 2018, it was decided to take a one-year break from organising the workshop to give the community some time to breathe... Little did we know what COVID-19 had in mind. 2020 has been a year full of sentiment and emotion, to say the least. The pandemic has dominated the news headlines all around the world and evoked a variety of emotions amongst the general public. Understanding these emotions not only provides insights into the way the public responds to the COVID-19 pandemic in itself and to the media coverage of the disease, but might help to encourage health promotion measures.

Research in automatic subjectivity and sentiment analysis remains a popular research task in the field of computational linguistics with a great application potential. Over the years the problem of dealing with affect in text has evolved, making it a very challenging research area with many research questions that still need to be answered and often requiring interdisciplinary approaches.

The aim of the Eleventh Workshop on Computational Approaches to Subjectivity, Sentiment and Social Media Analysis (WASSA 2021) was to continue the line of the previous editions and bring together researchers working on Subjectivity, Sentiment Analysis, Emotion Detection and Classification and their applications to other NLP or real-world tasks (e.g. public health messaging, fake news, media impact analysis) and researchers working on interdisciplinary aspects of affect computation from text. We also welcomed submissions that specifically tackled sentiment or emotion detection and classification in the context of the COVID-19 pandemic.

Though the timing was rather tight, we decided to also organise a shared task on Predicting Empathy and Emotion in Reaction to News Stories (https://competitions.codalab.org/competitions/28713). This task aimed at developing models which can predict empathy (Track I) and emotion (Track II) based on essays written in reaction to news articles which reported on harm caused to a person, a group, or other situations. Five teams participated in the shared task, with three teams submitting predictions for both tracks. For track I, empathy prediction, four teams submitted a system and the best result obtained was an average Pearson correlation of 0.545. For track II, emotion label prediction, four teams submitted a system and the best result was a macro F-1 of 55.3%.

For the main workshop, we accepted 15 papers as long and another 9 as short papers, leading to a total of 24/32 accepted papers (acceptance rate of 75%). For the shared task we received 6 system description paper submissions, out of which we accepted 5. Thus, in total 29 papers will be presented at the workshop, together with the additional contribution from our invited speaker Lyle Ungar, professor of Computer and Information Science at the University of Pennsylvania.

Accepted papers deal with topics including implicit and explicit sentiment analysis, emotion detection or classification and the detection of hate speech, stance or sarcasm. A large number of papers deal with languages other than English, including multilingual approaches but also work conducted on Italian, Dutch, code-mixed Hindi-English and even less-resourced languages such as Sindhi, Marathi and Arabizi. The dominance of COVID-19 in the headlines did not translate to a high number of COVID-related papers, but we gladly included one paper scrutinising resistance to COVID-19 directives.

This year we also asked the reviewers for recommendations for a best paper award and are thrilled to announce that the paper "Lightweight Models for Multimodal Sequential Data" by Soumya Sourav and Jessica Ouyang wins this year's award.

We would like to thank the EACL 2021 Organizers and Workshop Chairs for their help and support at the different stages of the workshop organisation process. We are also especially grateful to the Program Committee members for the time and effort spent to thoroughly review and assess the papers. Finally, we would like to extend our thanks to our invited speaker – Prof. Lyle Ungar - for accepting the invitation to deliver the keynote talk.

Orphée De Clercq, Alexandra Balahur, João Sedoc, Valentin Barriere, Shabnam Tafreshi, Sven Buechel and Veronique Hoste

WASSA 2021 Chairs

Keynote talk: Using language to study emotional contagion

Lyle Ungar
Computer and Information Science, University of Pennsylvania

Abstract

The words people use not only reveal their happiness, anger, depression, and empathy toward others; they also influence the people they communicate with, changing their moods and language. Language thus drives emotional contagion and allows us to measure it. We present case studies in which people experience different amounts of emotional contagion based on two factors: 1) Their empathy style: The words people use on Facebook, when correlated with their scores on empathy-measuring questionnaires, reveal empathy-driven emotional contagion. 2) Their level of depression: SMS messages from cell phones show that although depressed people use more sad, negative, and angry language, the texts they receive only show more anger than texts to non-depressed people, suggesting that anger may be more contagious than sadness.

Organizers

Organizers

Orphée de Clercq – Ghent University, Belgium
Alexandra Balahur – European Commission Joint Research Centre
João Sedoc – New York University, U.S.A
Valentin Barriere – European Commission Joint Research Centre
Shabnam Trafreshi – Georgetown University & IBM, U.S.A.
Sven Buechel – Friedrich Schiller University Jena, Germany
Veronique Hoste – Ghent University, Belgium

Program Committee

Liesbeth Allein – European Commission Joint Research Centre
Jeremy Barnes – University Pompeu Fabra, Spain
Sabine Bergler – Concordia University, Canada
Cristina Bosco – University of Torino, Italy
Nicoletta Calzolari – CNR Pisa, Italy
Erik Cambria – Nanyang Technological University, Singapore
Sergio Consoli – European Commission Joint Research Centre
Montse Cuadros – Vicomtech, Spain
Luna De Bruyne – Ghent University, Belgium
Lingjia Deng – Bloomberg, U.S.A.
Antske Fokkens – VU Amsterdam, The Netherlands
Michael Gamon – Microsoft, U.S.A.
Lorenzo Gatti – University of Twente, The Netherlands
Matthias Hartung – Semalytix GmbH, Germany
Dirk Hovy – Bocconi University, Italy
Stefano Maria Iacus – European Commission Joint Research Centre
Carlos A. Iglesias – Universidad Politécnica de Madrid, Spain
Ruben Izquierdo Bevia – Nuance, Spain
Gilles Jacobs – Ghent University, Belgium
Aditya Joshi – Notiv, Australia
Evgeny Kim – University of Stuttgart, Germany
Manfred Klenner – University of Zuerich, Switzerland
Roman Klinger – University of Stuttgart, Germany
Emiel Krahmer – Tilburg University, The Netherlands
Mayank Kulkarni – Bloomberg, U.S.A.
Els Lefever – Ghent University, Belgium
Edison Marrese-Taylor – University of Tokyo, Japan
Saif M. Mohammad – National Research Council Canada, Canada
Karo Moilanen – University of Oxford, U.K.
Guenter Neumann – DFKI, Germany
Malvina Nissim – University of Groningen, The Netherlands
Laura Ana Maria Oberländer – University of Stuttgart, Germany
Constantin Orasan – University of Surrey, U.K.
Sean Papay – University of Stuttgart, Germany
Viviana Patti – University of Torino, Italy

Viktor Pekar – University of Wolverhampton, U.K.
Jose Manuel Perea-Ortega – University of Extremadura, Spain
Barbara Plank – IT University of Copenhagen, Denmark
Daniel Preotiuc-Pietro – Bloomberg, U.S.A.
Paolo Rosso – Technical University of Valencia, Spain
Pranaydeep Singh – Ghent University, Belgium
Josef Steinberger – West Bohemia University Prague, The Czech Republic
Carlo Strapparava – Fondazione Bruno Kessler, Italy
Mike Thelwall – University of Wolverhampton, U.K
Dan Tufis – RACAI, Romania
Cynthia Van Hee – Ghent University, Belgium
Tony Veale – University College Dublin, Ireland
Charles Welch – University of Michigan, U.S.A
Michael Wiegand – Saarland University, Germany
Taras Zagibalov – Brantwatch, U.K.

Table of Contents

ToxCCIn: Toxic Content Classification with Interpretability
Tong Xiang, Sean MacAvaney, Eugene Yang and Nazli Goharian 1

Language that Captivates the Audience: Predicting Affective Ratings of TED Talks in a Multi-Label Classification Task
Elma Kerz, Yu Qiao and Daniel Wiechmann 13

Partisanship and Fear are Associated with Resistance to COVID-19 Directives
Mike Lindow, David DeFranza, Arul Mishra and Himanshu Mishra 25

Explainable Detection of Sarcasm in Social Media
Ramya Akula and Ivan Garibay 34

Emotion Ratings: How Intensity, Annotation Confidence and Agreements are Entangled
Enrica Troiano, Sebastian Padó and Roman Klinger 40

Disentangling Document Topic and Author Gender in Multiple Languages: Lessons for Adversarial Debiasing
Erenay Dayanik and Sebastian Padó 50

Universal Joy A Data Set and Results for Classifying Emotions Across Languages
Sotiris Lamprinidis, Federico Bianchi, Daniel Hardt and Dirk Hovy 62

FEEL-IT: Emotion and Sentiment Classification for the Italian Language
Federico Bianchi, Debora Nozza and Dirk Hovy 76

An End-to-End Network for Emotion-Cause Pair Extraction
Aaditya Singh, Shreeshail Hingane, Saim Wani and Ashutosh Modi 84

WASSA 2021 Shared Task: Predicting Empathy and Emotion in Reaction to News Stories
Shabnam Tafreshi, Orphee De Clercq, Valentin Barriere, Sven Buechel, João Sedoc and Alexandra Balahur 92

PVG at WASSA 2021: A Multi-Input, Multi-Task, Transformer-Based Architecture for Empathy and Distress Prediction
Atharva Kulkarni, Sunanda Somwase, Shivam Rajput and Manisha Marathe 105

WASSA@IITK at WASSA 2021: Multi-task Learning and Transformer Finetuning for Emotion Classification and Empathy Prediction
Jay Mundra, Rohan Gupta and Sagnik Mukherjee 112

Analyzing Curriculum Learning for Sentiment Analysis along Task Difficulty, Pacing and Visualization Axes
Anvesh Rao Vijjini, Kaveri Anuranjana and Radhika Mamidi 117

Lightweight Models for Multimodal Sequential Data
Soumya Sourav and Jessica Ouyang 129

Exploring Implicit Sentiment Evoked by Fine-grained News Events
Cynthia Van Hee, Orphee De Clercq and Veronique Hoste 138

Exploring Stylometric and Emotion-Based Features for Multilingual Cross-Domain Hate Speech Detection
Ilia Markov, Nikola Ljubešić, Darja Fišer and Walter Daelemans 149

Emotion-Aware, Emotion-Agnostic, or Automatic: Corpus Creation Strategies to Obtain Cognitive Event Appraisal Annotations
Jan Hofmann, Enrica Troiano and Roman Klinger 160

Hate Towards the Political Opponent: A Twitter Corpus Study of the 2020 US Elections on the Basis of Offensive Speech and Stance Detection
Lara Grimminger and Roman Klinger 171

Synthetic Examples Improve Cross-Target Generalization: A Study on Stance Detection on a Twitter corpus.
Costanza Conforti, Jakob Berndt, Mohammad Taher Pilehvar, Chryssi Giannitsarou, Flavio Toxvaerd and Nigel Collier 181

Creating and Evaluating Resources for Sentiment Analysis in the Low-resource Language: Sindhi
Wazir Ali, Naveed Ali, Yong Dai, Jay Kumar, Saifullah Tumrani and Zenglin Xu 188

Towards Emotion Recognition in Hindi-English Code-Mixed Data: A Transformer Based Approach
Anshul Wadhawan and Akshita Aggarwal 195

Nearest neighbour approaches for Emotion Detection in Tweets
Olha Kaminska, Chris Cornelis and Veronique Hoste 203

L3CubeMahaSent: A Marathi Tweet-based Sentiment Analysis Dataset
Atharva Kulkarni, Meet Mandhane, Manali Likhitkar, Gayatri Kshirsagar and Raviraj Joshi . . . 213

Multi-Emotion Classification for Song Lyrics
Darren Edmonds and João Sedoc 221

ONE: Toward ONE model, ONE algorithm, ONE corpus dedicated to sentiment analysis of Arabic/Arabizi and its dialects
Imane Guellil, Faical Azouaou, Fodil Benali and Hachani Ala-Eddine 236

Me, myself, and ire: Effects of automatic transcription quality on emotion, sarcasm, and personality detection
John Culnan, Seongjin Park, Meghavarshini Krishnaswamy and Rebecca Sharp 250

Emotional RobBERT and Insensitive BERTje: Combining Transformers and Affect Lexica for Dutch Emotion Detection
Luna De Bruyne, Orphee De Clercq and Veronique Hoste 257

EmpNa at WASSA 2021: A Lightweight Model for the Prediction of Empathy, Distress and Emotions from Reactions to News Stories
Giuseppe Vettigli and Antonio Sorgente 264

MilaNLP @ WASSA: Does BERT Feel Sad When You Cry?
Tommaso Fornaciari, Federico Bianchi, Debora Nozza and Dirk Hovy 269

Team Phoenix at WASSA 2021: Emotion Analysis on News Stories with Pre-Trained Language Models
Yash Butala, Kanishk Singh, Adarsh Kumar and Shrey Shrivastava 274

Conference Program

Monday, April 19, 2021

9:00–9:15 *Opening of the WASSA workshop*
Orphée De Clercq

9:15–11:00 ORAL SESSION 1

9:15–9:40 *ToxCCIn: Toxic Content Classification with Interpretability*
Tong Xiang, Sean MacAvaney, Eugene Yang and Nazli Goharian

9:40–10:05 *Language that Captivates the Audience: Predicting Affective Ratings of TED Talks in a Multi-Label Classification Task*
Elma Kerz, Yu Qiao and Daniel Wiechmann

10:05–10:30 *Partisanship and Fear are Associated with Resistance to COVID-19 Directives*
Mike Lindow, David DeFranza, Arul Mishra and Himanshu Mishra

10:30–10:45 *Explainable Detection of Sarcasm in Social Media*
Ramya Akula and Ivan Garibay

10:45–11:00 *Emotion Ratings: How Intensity, Annotation Confidence and Agreements are Entangled*
Enrica Troiano, Sebastian Padó and Roman Klinger

11:00–11:30 ***Coffee break***

11:30–12:50 ORAL SESSION 2

11:30–11:55 *Disentangling Document Topic and Author Gender in Multiple Languages: Lessons for Adversarial Debiasing*
Erenay Dayanik and Sebastian Padó

11:55–12:20 *Universal Joy A Data Set and Results for Classifying Emotions Across Languages*
Sotiris Lamprinidis, Federico Bianchi, Daniel Hardt and Dirk Hovy

12:20–12:35 *FEEL-IT: Emotion and Sentiment Classification for the Italian Language*
Federico Bianchi, Debora Nozza and Dirk Hovy

Monday, April 19, 2021 (continued)

12:35–12:50 *An End-to-End Network for Emotion-Cause Pair Extraction*
Aaditya Singh, Shreeshail Hingane, Saim Wani and Ashutosh Modi

12:50–13:45 ***Lunch break***

13:45–14:30 SHARED TASK SESSION

13:45–14:00 *WASSA 2021 Shared Task: Predicting Empathy and Emotion in Reaction to News Stories*
Shabnam Tafreshi, Orphee De Clercq, Valentin Barriere, Sven Buechel, João Sedoc and Alexandra Balahur

14:00–14:15 *PVG at WASSA 2021: A Multi-Input, Multi-Task, Transformer-Based Architecture for Empathy and Distress Prediction*
Atharva Kulkarni, Sunanda Somwase, Shivam Rajput and Manisha Marathe

14:15–14:30 *WASSA@IITK at WASSA 2021: Multi-task Learning and Transformer Finetuning for Emotion Classification and Empathy Prediction*
Jay Mundra, Rohan Gupta and Sagnik Mukherjee

14:30–15:30 ***INVITED TALK: Using language to study emotional contagion***
Lyle Ungar

15:30–15:45 ***Coffee break***

15:45–16:35 ORAL SESSION 3

15:45–16:10 *Analyzing Curriculum Learning for Sentiment Analysis along Task Difficulty, Pacing and Visualization Axes*
Anvesh Rao Vijjini, Kaveri Anuranjana and Radhika Mamidi

16:10–16:35 *Lightweight Models for Multimodal Sequential Data*
Soumya Sourav and Jessica Ouyang

16:35–17:45 POSTER SESSION

Exploring Implicit Sentiment Evoked by Fine-grained News Events
Cynthia Van Hee, Orphee De Clercq and Veronique Hoste

Exploring Stylometric and Emotion-Based Features for Multilingual Cross-Domain Hate Speech Detection
Ilia Markov, Nikola Ljubešić, Darja Fišer and Walter Daelemans

Emotion-Aware, Emotion-Agnostic, or Automatic: Corpus Creation Strategies to Obtain Cognitive Event Appraisal Annotations
Jan Hofmann, Enrica Troiano and Roman Klinger

Hate Towards the Political Opponent: A Twitter Corpus Study of the 2020 US Elections on the Basis of Offensive Speech and Stance Detection
Lara Grimminger and Roman Klinger

Synthetic Examples Improve Cross-Target Generalization: A Study on Stance Detection on a Twitter corpus.
Costanza Conforti, Jakob Berndt, Mohammad Taher Pilehvar, Chryssi Giannitsarou, Flavio Toxvaerd and Nigel Collier

Creating and Evaluating Resources for Sentiment Analysis in the Low-resource Language: Sindhi
Wazir Ali, Naveed Ali, Yong Dai, Jay Kumar, Saifullah Tumrani and Zenglin Xu

Towards Emotion Recognition in Hindi-English Code-Mixed Data: A Transformer Based Approach
Anshul Wadhawan and Akshita Aggarwal

Nearest neighbour approaches for Emotion Detection in Tweets
Olha Kaminska, Chris Cornelis and Veronique Hoste

L3CubeMahaSent: A Marathi Tweet-based Sentiment Analysis Dataset
Atharva Kulkarni, Meet Mandhane, Manali Likhitkar, Gayatri Kshirsagar and Raviraj Joshi

Multi-Emotion Classification for Song Lyrics
Darren Edmonds and João Sedoc

ONE: Toward ONE model, ONE algorithm, ONE corpus dedicated to sentiment analysis of Arabic/Arabizi and its dialects
Imane Guellil, Faical Azouaou, Fodil Benali and Hachani Ala-Eddine

Monday, April 19, 2021 (continued)

Me, myself, and ire: Effects of automatic transcription quality on emotion, sarcasm, and personality detection
John Culnan, Seongjin Park, Meghavarshini Krishnaswamy and Rebecca Sharp

Emotional RobBERT and Insensitive BERTje: Combining Transformers and Affect Lexica for Dutch Emotion Detection
Luna De Bruyne, Orphee De Clercq and Veronique Hoste

EmpNa at WASSA 2021: A Lightweight Model for the Prediction of Empathy, Distress and Emotions from Reactions to News Stories
Giuseppe Vettigli and Antonio Sorgente

MilaNLP @ WASSA: Does BERT Feel Sad When You Cry?
Tommaso Fornaciari, Federico Bianchi, Debora Nozza and Dirk Hovy

Team Phoenix at WASSA 2021: Emotion Analysis on News Stories with Pre-Trained Language Models
Yash Butala, Kanishk Singh, Adarsh Kumar and Shrey Shrivastava

17:45–18:00 ***Best paper award and closing***