CONSTRAINT 2022

Second Workshop on Combating Online Hostile Posts in Regional Languages during Emergency Situation

Proceedings of the Workshop

May 27, 2022

©2022 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 978-1-955917-31-5

Preface

The advent of Web 2.0 induced the evolution of what has traditionally been described as a "participatory Web". From pop-culture music to Black Friday becoming a global phenomenon, and movements like BlackLivesMatter turning into a powerful instrument of global resistance, the Internet and social media have played a pivotal role. As much as we relish the connectedness facilitated by social media, the sentient being in all of us cannot remain obscured by the perils of the unabated misuse of the very free speech that these platforms aim to empower. Within the shadows of a transparent yet anonymous social media, lurk those disguising themselves as pseudo-flag-bearers of free speech, and pounce on every opportunity they get to spread vile content, detrimental to society. Such miscreants are desperate to misuse those 280 character sound bites to further their anti-openness agendas in the form of hate speech, disinformation, and ill-intended propaganda. Such menace experiences flare-ups during emergency situations such as the COVID-19 outbreak and geopolitically conflicting global order.

There have been numerous efforts toward addressing some of these problems computationally, but with evolving complexities of online harmful content, more robust solutions are needed. Some of these challenges stem from linguistic diversity, abstract semiotics, multimodality, anonymity of the real instigators, etc. Thus, there is a pressing need to start a discussion around such aspects, which are more inclusive than conventional efforts. With this in mind, and motivated by the success of the first edition of the CONSTRAINT Workshop on Combating Online Hostile Posts in Regional Languages during Emergency Situation, we have launched the second edition in hybrid mode, with a special focus on Multimodal Low-Resource Language Processing to Combat COVID-19 Related Online Hostile Content.

The workshop additionally highlighted three major points:

1. Regional languages: offensive posts may be written in low-resource regional languages, e.g., Tamil, Urdu, Bangali, Polish, Czech, Lithuanian, etc.
2. Emergency situations: The proposed solutions should be able to tackle misinformation during emergency situations where, due to the lack of enough historical data, machine learning models need to adopt additional intelligence to handle emerging and novel posts.
3. Early detection: Since the impact of misinformation during emergency situations can be highly detrimental to society (e.g., health-related misadvice during a pandemic may take human's life), we encourage solutions that can detect such hostile posts as early as possible after they have been posted in social media.

Our workshop also features a shared task titled: Hero, Villain and Victim: Dissecting harmful memes for Semantic role labelling of entities. The objective is to determine the role of the entities referred to within a meme: hero vs. villain vs. victim vs. other. The meme is to be analyzed from the perspective of its author. The datasets released as part of this shared task span memes from two domains: COVID-19 and US Politics. The complex and engaging nature of the shared task led to a total of 6 unique final submissions for evaluation, from amongst 105 total registered participants.

We accepted a total of ten papers: four for the regular track and six for the shared task. The workshop papers cover topics ranging from detecting multimodal/unimodal fake news (Choi et al., 2022; Lucas et al., 2022) to aggressive content (Sharif et al., 2022), with additional fine-grained analysis and sub-tasks like document retrieval towards mitigating misinformation (Sundriyal et al., 2022). On the other hand, the accepted papers for the shared task proposed various multimodal fusion strategies including state-of-the-art encoder models such as variants of ViT, BERT, and CLIP (Nandi et al., 2022; Kun et al., 2022; Montariol et al., 2022), with ensembling playing a key role in the overall performance enhancement. Consequently, diverse strategies for addressing the task along with their limitations are elucidated via the contributions made hereupon.

We are glad to have 3 eminent invited speakers: (i) Smaranda Muresan, Research Scientist at the Data Science Institute (DSI) and the Department of Computer Science at Columbia University, and Amazon, (2) Isabelle Augenstein, Associate Professor at the University of Copenhagen, Department of Computer Science, where she heads the Copenhagen Natural Language Understanding research group as well as the Natural Language Processing section, and (iii) Andreas Vlachos, Associate Professor at the Natural

Language and Information Processing group at the Department of Computer Science and Technology at the University of Cambridge and a member of the European Lab for Learning and Intelligent Systems.
We thank the authors and the task participants for their interest in the workshop. We would also like to thank the program committee for their help with reviewing the papers and with advertising the workshop.
The work was partially supported by a Wipro research grant, Ramanujan Fellowship, the Infosys Centre for AI, IIIT Delhi, India, and ihub-Anubhuti-iiitd Foundation, set up under the NM-ICPS scheme of the Department of Science and Technology, India.
It is also part of the Tanbih mega-project, which is developed at the Qatar Computing Research Institute, HBKU, and aims to limit the impact of fake news, propaganda, and media bias by making users aware of what they are reading, thus promoting media literacy and critical thinking.

The CONSTRAINT 2022 Organizers: Tanmoy Chakraborty, Md. Shad Akhtar, Kai Shu, H. Russell Bernard, Maria Liakata, and Preslav Nakov
Website: http://lcs2.iiitd.edu.in/CONSTRAINT-2022/

Organizing Committee

Program Committee Chairs

Tanmoy Chakraborty, IIIT Delhi, India
Md. Shad Akhtar, IIIT Delhi, India
Kai Shu, Illinois Institute of Technology, USA
H. Russell Bernard, Arizona State University, USA
Maria Liakata, Queen Mary, University of London, UK
Preslav Nakov, Qatar Computing Research Institute, HBKU, Qatar

Web Chair

Aseem Srivastava, IIIT Delhi, India

Invited Speakers

Isabelle Augenstein, University of Copenhagen, Denmark
Smaranda Muresan, Columbia University, USA
Andreas Vlachos, University of Cambridge, UK

Program Committee

Program Committee

Amila Silva, The University of Melbourne
Andreas Vlachos, University of Cambridge
Anoop Kunchukuttan, Microsoft
Arkaitz Zubiaga, Queen Mary University of London
Balaji Vasan Srinivasan, Adobe Research
Firoj Alam, Qatar Computing Research Institute, HBKU
Marc Spaniol, Université de Caen
Matt Lease, University of Texas at Austin
Monojit Choudhury, Microsoft Research
Tracy King, Adobe Sensei and Search
Paolo Papotti, EURECOM
Paolo Rosso, Universitat Politècnica de València
Pushpak Bhattacharya, IIT Bombay
Roy Ka-Wei Lee, Singapore University of Technology and Design
Xinyi Zhou, Syracuse University
Yingtong Dou, University of Illinois at Chicago
Reza Zafarani, Syracuse University
Nitin Agarwal, University of Arkansas at Little Rock
Victoria Rubin, Western University
Francesco Barbieri, Snap Research
Ashique KhudaBukhsh, Carnegie Mellon University
Ugur Kursuncu, Georgia State University
Vagelis Papalexakis, University of California Riverside
Sibel Adali, Rensselaer Polytechnic Institute
Shivam Sharma, IIIT Delhi, Wipro AI Research
Chhavi Sharma, Wipro AI Research
Shivani Kumar, IIIT Delhi
Yash Kumar Atri, IIIT Delhi
Sarah Masud, IIIT Delhi
Sunil Saumya, IIIT Dharwad
Megha Sundriyal, IIIT Delhi
Karan Goyal, IIIT Delhi
Anam Fatima, IIIT Delhi

Table of Contents

Findings of the CONSTRAINT 2022 Shared Task on Detecting the Hero, the Villain, and the Victim in Memes
Shivam Sharma, Tharun Suresh, Atharva Kulkarni, Himanshi Mathur, Preslav Nakov, Md. Shad Akhtar and Tanmoy Chakraborty 1

DD-TIG at Constraint@ACL2022: Multimodal Understanding and Reasoning for Role Labeling of Entities in Hateful Memes
Ziming Zhou, Han Zhao, Jingjing Dong, Jun Gao and Xiaolong Liu 12

Are you a hero or a villain? A semantic role labelling approach for detecting harmful memes.
Shaik Fharook, Syed Sufyan Ahmed, Gurram Rithika, Sumith Sai Budde, Sunil Saumya and Shankar Biradar 19

Logically at the Constraint 2022: Multimodal role labelling
Ludovic Kun, Jayesh Bankoti and David Kiskovski 24

Combining Language Models and Linguistic Information to Label Entities in Memes
Pranaydeep Singh, Aaron Maladry and Els Lefever 35

Detecting the Role of an Entity in Harmful Memes: Techniques and their Limitations
Rabindra Nath Nandi, Firoj Alam and Preslav Nakov 43

Fine-tuning and Sampling Strategies for Multimodal Role Labeling of Entities under Class Imbalance
Syrielle Montariol, Étienne Simon, Arij Riabi and Djamé Seddah 55

Document Retrieval and Claim Verification to Mitigate COVID-19 Misinformation
Megha Sundriyal, Ganeshan Malhotra, Md Shad Akhtar, Shubhashis Sengupta, Andrew Fano and Tanmoy Chakraborty 66

M-BAD: A Multilabel Dataset for Detecting Aggressive Texts and Their Targets
Omar Sharif, Eftekhar Hossain and Mohammed Moshiul Hoque 75

How does fake news use a thumbnail? CLIP-based Multimodal Detection on the Unrepresentative News Image
Hyewon Choi, Yejun Yoon, Seunghyun Yoon and Kunwoo Park 86

Detecting False Claims in Low-Resource Regions: A Case Study of Caribbean Islands
Jason Lucas, Limeng Cui, Thai Le and Dongwon Lee 95

Program

Friday, May 27, 2022

09:00 - 09:10 *Opening Remarks*

09:10 - 10:10 *Keynote 1: Isabelle Augenstein — Automatically Detecting Scientific Misinformation*

10:10 - 10:30 *Regular Paper Session - I*

M-BAD: A Multilabel Dataset for Detecting Aggressive Texts and Their Targets
Omar Sharif, Eftekhar Hossain and Mohammed Moshiul Hoque

10:30 - 11:00 *Coffee break*

11:00 - 12:00 *Keynote 2: Andreas Vlachos — Fact-Checking Using Structured and Unstructured Information*

13:00 - 12:00 *Regular Paper Session - II*

How does fake news use a thumbnail? CLIP-based Multimodal Detection on the Unrepresentative News Image
Hyewon Choi, Yejun Yoon, Seunghyun Yoon and Kunwoo Park

Detecting False Claims in Low-Resource Regions: A Case Study of Caribbean Islands
Jason Lucas, Limeng Cui, Thai Le and Dongwon Lee

Document Retrieval and Claim Verification to Mitigate COVID-19 Misinformation
Megha Sundriyal, Ganeshan Malhotra, Md Shad Akhtar, Shubhashis Sengupta, Andrew Fano and Tanmoy Chakraborty

13:00 - 14:00 *Lunch Break*

14:00 - 15:00 *Keynote 3: Smaranda Muresan — The Role of Text Generation in Fighting Hostile Posts*

15:00 - 15:30 *Coffee Break*

15:30 - 17:15 *Shared Task Session*

Findings of the CONSTRAINT 2022 Shared Task on Detecting the Hero, the Villain, and the Victim in Memes
Shivam Sharma, Tharun Suresh, Atharva Kulkarni, Himanshi Mathur, Preslav Nakov, Md. Shad Akhtar and Tanmoy Chakraborty

DD-TIG at Constraint@ACL2022: Multimodal Understanding and Reasoning for Role Labeling of Entities in Hateful Memes
Ziming Zhou, Han Zhao, Jingjing Dong, Jun Gao and Xiaolong Liu

Are you a hero or a villain? A semantic role labelling approach for detecting harmful memes.
Shaik Fharook, Syed Sufyan Ahmed, Gurram Rithika, Sumith Sai Budde, Sunil Saumya and Shankar Biradar

Logically at the Constraint 2022: Multimodal role labelling
Ludovic Kun, Jayesh Bankoti and David Kiskovski

Combining Language Models and Linguistic Information to Label Entities in Memes
Pranaydeep Singh, Aaron Maladry and Els Lefever

Detecting the Role of an Entity in Harmful Memes: Techniques and their Limitations
Rabindra Nath Nandi, Firoj Alam and Preslav Nakov

Fine-tuning and Sampling Strategies for Multimodal Role Labeling of Entities under Class Imbalance
Syrielle Montariol, Étienne Simon, Arij Riabi and Djamé Seddah

16:50 - 17:15 *Closing*